

Dynamical Systems via Domains

**Toward a Unified Foundation of
Symbolic and Non-symbolic
Computation**

Levin Hornischer

Dynamical Systems via Domains

Toward a Unified Foundation of
Symbolic and Non-symbolic
Computation

ILLC Dissertation Series DS-2021-10



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Science Park 107
1098 XG Amsterdam
phone: +31-20-525 6051
e-mail: illc@uva.nl
homepage: <http://www.illc.uva.nl/>

Copyright © 2021 by Levin Hornischer

Cover woodcut and photo by Monika Schaber
Printed and bound by GVO drukkers & vormgevers B.V.

ISBN: 978-94-6332-789-3

Dynamical Systems via Domains

Toward a Unified Foundation of Symbolic and Non-symbolic Computation

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K.I.J. Maex

ten overstaan van een door het College voor Promoties ingestelde
commissie, in het openbaar te verdedigen in de Agnietenkapel
op woensdag 13 oktober 2021, te 10.00 uur

door

Levin Adrian Hornischer

geboren te Filderstadt

Promotiecommissie

<i>Promotores:</i>	prof. dr. M. van Lambalgen	Universiteit van Amsterdam
	prof. dr. F. Berto	Universiteit van Amsterdam

<i>Overige leden:</i>	prof. dr. J.F.A.K. van Benthem	Universiteit van Amsterdam
	prof. dr. S.J.L. Smets	Universiteit van Amsterdam
	prof. dr. A. Betti	Universiteit van Amsterdam
	prof. dr. S. Abramsky	University of Oxford
	prof. dr. H. Leitgeb	Ludwig-Maximilians-Universität München
	dr. L. Incurvati	Universiteit van Amsterdam

Faculteit der Geesteswetenschappen



The research for this doctoral thesis received financial assistance from the *Netherlands Organisation for Scientific Research* (NWO) as part of the research programme *PhDs in the Humanities*, project number 322-20-017.

Für Moni, Winni und Jele

Contents

Acknowledgments	xix
1 Introduction	1
1.1 Motivation	2
1.2 Outline	8

Part One: Symbolic computation

2 Trajectory domains 1: Construction	15
2.1 Introduction	15
2.2 Background	18
2.2.1 Labeled transition systems	18
2.2.2 Domain and order theory	19
2.3 Two guiding examples	20
2.3.1 Observing a black box system	20
2.3.2 Concurrent computation	28
2.3.3 Summary and outlook	32
2.4 Pre-behavioral transition systems	33
2.4.1 Definition	33
2.4.2 Comments	34
2.4.3 Example constructions	36
2.5 Information containment of behaviors	38
2.5.1 Three definitions of information containment	38
2.5.2 . . . and how they are united	40
2.6 The characterization theorem	42
2.6.1 Statement	42
2.6.2 Proof	45
2.7 Behavioral transition systems	49

2.7.1	Definition	50
2.7.2	Simplifying assumptions	50
2.7.3	Examples	53
2.8	Trajectory domains	56
2.9	Generalizations of information systems	57
2.9.1	Scott information systems	57
2.9.2	. . . and their generalizations as BTSs	59
2.10	Conclusion	62
3	Trajectory domains 2: Category	67
3.1	Introduction	67
3.2	Background	71
3.2.1	Category of labeled transition systems	71
3.2.2	Domain theory	74
3.2.3	Category theory	76
3.2.4	Recap from the previous chapter	76
3.3	Category of behavioral transition systems	78
3.3.1	Definition	78
3.3.2	Basic properties	80
3.3.3	Embedding labeled transition systems	82
3.3.4	Removing non-approximable behavior	83
3.4	Trajectory domain functor	85
3.5	Adjunction between systems and domains	87
3.5.1	Extensionalizing	87
3.5.2	Unlabeling and reflexing	91
3.5.3	Adjunction to domains	93
3.6	Toward incorporating labels on domains	97
3.6.1	Marked domains	98
3.6.2	An interpretation of relevance logic	100
3.7	Conclusion	103

Part Two: Non-symbolic computation

4	Systems and domains 1: Model	109
4.1	Introduction	109
4.2	Background	119
4.2.1	Domain theory	119
4.2.2	Dynamical and topological systems	123
4.3	Observing dynamical systems	127
4.3.1	Basis or ‘set of possible observations’	127
4.3.2	The index set or ‘set of observation parameters’	128
4.3.3	Observed system	129

4.3.4	Refining observations	131
4.3.5	Observation probabilities	132
4.3.6	Summary	135
4.4	Dynamical domains	140
4.4.1	Dynamical dcpo's	140
4.4.2	Dynamical expanding systems	144
4.4.3	The limit theorem	147
4.4.4	Definition of dynamical domains	160
4.5	The system modeled by a dynamical domain	160
4.6	Dynamical domain models for systems	162
4.6.1	For dynamical systems	162
4.6.2	For topological systems	167
4.7	Conclusion	170
5	Systems and domains 2: Category	173
5.1	Introduction	173
5.2	The categories	180
5.2.1	Background	181
5.2.2	Categories of dynamical systems	182
5.2.3	Categories of measured topological systems	186
5.2.4	Categories of dynamical domains	188
5.2.5	Recap from chapter 4	193
5.2.6	Categories of based measured topological systems	196
5.2.7	Categories of max-reflective dynamical domains	198
5.3	The bottom layer of the main diagram	201
5.3.1	Dynamical systems as category of fractions	201
5.3.2	Compactification of a system: informally	202
5.3.3	Compactification of a system: formally	205
5.4	The system and domain functors	214
5.4.1	The system functor	215
5.4.2	The domain functor	216
5.4.3	Computational and logical compactification coincide	223
5.5	The systems-domains adjunction	227
5.5.1	The counit and unit	228
5.5.2	Triangle identities	232
5.6	Analyzing the systems-domains adjunction	234
5.6.1	Restricting to equivalence	234
5.6.2	Max-reflecting a dynamical domain	237
5.7	Conclusion	239

6	Systems and domains 3: Application	243
6.1	Introduction	243
6.2	Background	243
6.2.1	Recap dynamical systems and dynamical domains	244
6.2.2	Metric entropy	245
6.2.3	Topological entropy	246
6.3	Domain-entropy	247
6.3.1	Definition of domain-entropy	247
6.3.2	Main theorem on domain-entropy	249
6.3.3	Normal form for domain-entropy	252
6.4	Max-entropy	258
6.4.1	Definition of max-entropy	258
6.4.2	Main theorem on max-entropy	259
6.5	Conclusion	262

Part Three: Stability

7	Interlude: symbolic vs. non-symbolic	267
7.1	Non-symbolic computation as limit of symbolic computation	267
7.2	Non-symbolic realization of symbolic computation	269
7.2.1	Symbolic approximation	270
7.2.2	Ergodicity	272
7.2.3	Randomness	277
7.2.4	Stability	281
8	Stability: Fitch’s paradox and AI-safety	283
8.1	Introduction	283
8.2	Examples of stability	287
8.2.1	Verifiability and falsifiability (observation)	287
8.2.2	Safety (epistemology)	289
8.2.3	Safety (artificial intelligence)	290
8.2.4	Stability of belief (probabilistic reasoning)	292
8.2.5	Significance (mathematical modeling)	293
8.2.6	Further examples	295
8.3	Four principles of stability	296
8.3.1	A logic to reason about stability	296
8.3.2	Formalization and motivation of the principles	301
8.3.3	The duality between falsification and verification	305
8.3.4	Constructing sets of questions	306
8.4	Impossibility via a novel interpretation of Fitch’s paradox	307
8.4.1	Reinterpretation of Fitch’s paradox	307
8.4.2	Impossibility	308

8.5	Impossibility via semantics	310
8.5.1	Kripke semantics	310
8.5.2	Topological semantics	312
8.6	Applications	315
8.6.1	An extension of Fitch’s paradox	315
8.6.2	A limitation for AI-safety	319
8.7	Conclusion	322
9	Conclusion	327
A	Systems as a category of fractions	331
A.1	Statement of the theorem	331
A.2	Topological realizations of systems	335
A.3	The key lemma	338
A.4	Calculus of fractions	343
A.5	Equivalence	344
B	Dynamical domain example	349
B.1	A dynamical domain of binary sequences	349
B.2	More facts about the dynamical domain	355
B.3	Words on the components	356
B.4	Computing max-entropy	360
	Bibliography	363
	Index	385
	List of symbols	389
	Samenvatting	393
	Summary	395

List of Figures

1.1	An operational semantics for the bubble sort algorithm	3
1.2	Computation as dynamical systems	6
2.1	The North-South map	22
2.2	The trajectory domain of the North-South map	28
2.3	The independence diamond	29
3.1	Summary of the results	104
5.1	The main diagram	174
5.2	(Non-) examples of max-reflective finite Scott domains	199
5.3	The compactification functor	210
5.4	Notational conventions of this subsection	216
5.5	Proof of naturality	226
5.6	Overview of the results	239
7.1	Non-symbolic computation as limit of symbolic computation	268
7.2	The predator-prey dynamics	271
7.3	A converging system	276
8.1	The four principles and their duals	302
A.1	Calculus of fractions	332
A.2	Equivalence of spans	333
A.3	Composition of spans	333
A.4	Inverses of morphisms in the localization	334
A.5	Visualization of lemma A.3.1	339
B.1	Example of word realization	357

Publications

All chapters of the thesis are single-authored and have not been published previously (at the time of submission). Papers that aren't part of this thesis but have been published within this PhD project are the following:

- L. Hornischer (2019). “Toward a Logic for Neural Networks”. In: *The Logica Yearbook 2018*. Ed. by I. Sedlár and M. Blicha. London: College Publications, pp. 133–148

This is an early exploration of the guiding idea of the thesis: using domain theory to develop a semantics also for non-symbolic computation including neural networks. Some of the ideas eventually lead to chapter 2 (though the chapter itself is quite different) and to the last paper below.

- L. Hornischer (2020). “Logics of Synonymy”. In: *Journal of Philosophical Logic* 49, pp. 767–805

This investigates the notion of synonymy (or content identity or strong equivalence). It axiomatizes various notions of synonymy and characterizes them as good benchmarks. This paper is cited in chapter 2 when asking whether the generalization of Scott information systems presented there can be seen as moving to a stronger underlying notion of equivalence.

- L. Hornischer (2021). “The Logic of Information in State Spaces”. In: *The Review of Symbolic Logic* 14.1, pp. 155–186

This uses domain theory to describe the information contained in state spaces (dynamical systems, possible worlds, etc.). This provides a semantics to the logic HYPE and is applied to study information fusion. This paper is cited in chapter 8 as a potential tool to further study the notion of ‘state space stability’ discussed in that chapter.

Acknowledgments

First of all, I would like to thank my supervisors Michiel van Lambalgen and Franz Berto: Michiel, thank you for sharing some of your seemingly unbounded amount of knowledge (even if it can be intimidating when you start quoting page numbers from books you read 30 years ago); for teaching me how to look at a problem from many perspectives and to check with the (cognitive) data; for never going easy on me with critical questions while still being encouraging; for treating me like a peer in the many courses that we taught together; for trusting me with a lot of freedom in my research and making plenty of time for me; and for sharing a passion for music and mountains. Franz, thank you for your refreshing approach to philosophy (your motto of not taking oneself too seriously stuck with me); for your enthusiasm and the encouragement that comes with it; for magically being available 24/7; for being—and I cannot stress this enough—incredibly supportive; and for your humor that cannot fail to entertain even in impossible worlds. I am indebted to you two for your support, guidance, and encouragement making this PhD project possible!

Next, I would like to thank my committee: Samson Abramsky, Johan van Benthem, Arianna Betti, Luca Incurvati, Hannes Leitgeb, and Sonja Smets. I feel very lucky and honored that you have agreed to this task. Your work is a great inspiration to me, and I am deeply grateful for all the time that you have devoted to this thesis—especially given its long and often abstract nature. Thank you!

I have been very fortunate to have spent my PhD at ILLC with its inspiring, collegial, and interdisciplinary atmosphere. This is made by its people: First, I would like to thank Arnold, Dean, Ivar, Kaibo, Nadine, and Thomas, for making—at various time periods over the past four years—the austere sounding ‘Science Park Room F1.11’ a wonderful and welcoming office. Special thanks to Dean for the many enlightening philosophy comments. Over the years, I had the pleasure to meet many inspiring, intimidatingly smart, but unbelievably kind people at ILLC. Thank you for the lunches, coffees, chats, discussions, and much more: in particular, Ana Lucia, Anna, Anthi, Arnold, Arthur, Aybüke, Bas, Bastiaan, Dean,

Dieuwke, Ece, Frederik, Gianluca, Giorgio, Giovanni, Ilaria, Iris, Ivar, Julian, Kaibo, Karolina, Leïla, Malvin, Mario G., Milica, Nadine, Peter H., Riccardo, Robert P., Ronald, Sam, Samira, Simon, Sirin, Thom, Thomas, Tom, Yvette, and Zoi, as well as you fellow MoLers. A big part of the special atmosphere at ILLC is also due to its extraordinary staff members: Thank you for keeping a—proverbial and literal—open door. In particular, for the inspiring encounters over the years, thank you, Arianna, Bahareh, Benedikt, Benno, Dick, Floris, Jakub, Jelle, Johan, Luca, Maria, Nick, Robert v.R., and Sonja. And for making everything work and for always being there, a big thank you to Caitlin, Debbie, Jenny, Karine, Patty, Peter, and Tanja.

During my PhD, I also had the privilege to do two research visits: At the end of 2018, I was at the Department of Computer Science at the University of Oxford visiting Samson Abramsky and his group. Thank you for your great hospitality, the opportunity to present my work, and the introduction to contextuality. In this thesis, I've tried to live up to Samson's wise words about using category theory to systematize. Special thanks to Sivert and Rui for welcoming me in their office, and to Karen Barnes for a perfect organization. In winter 2019–20, I was at the Munich Center of Mathematical Philosophy at LMU Munich visiting Hannes Leitgeb. Thank you, Hannes, for making this possible, for the insightful meetings, for making me feel at home, and for encouraging me to present my papers. A heartfelt thank you to all the MCMPerS that made this such a wonderful stay: in particular, Alessandra, Conrad, Ivano, Lorenzo, Martin, Matteo, Norbert, Sander, Sena, Timo, and Ursula. On the topic of visits, I would also like to thank Richard Evans for the truly inspiring conversations during his visit in Amsterdam. In the same vein, thank you, Laurenz Hudetz, for the stimulating chats during your time at ILLC. Finally, Franz, you have set up a perfect visit for me to St Andrews, which we then unfortunately had to cancel last minute due to the pandemic. But I sincerely hope we can make up for this soon!

Further, I am thankful to the audiences of the events where I could present some (earlier versions) of the material in this thesis. This includes, among others, the LoC Seminar in Amsterdam, TULIPS in Utrecht, the Algebra|Coalgebra Seminar in Amsterdam, OASIS in Oxford, the Logic Colloquium in Munich, and Logica 2018 in Hejnice. Many thanks to Aybüke Özgün for comments on what is now chapter 8, to Bastiaan van der Weij for comments on chapter 1, and to Benno van den Berg for comments on chapter 5. Thank you, James Grayot, for the philosophy discussions and answering my language questions. Thank you, Bastiaan and Michiel, for the excellent Dutch translation of the summary. And also thank you, Bastiaan, for all the other invaluable thesis advice. That said, of course no one but me is responsible for any mistakes. Finally, a big thank you to my mum, Monika Schaber, for the beautiful cover!

A shout-out to my Amsterdam-related friends: in particular, Ana, Bastiaan, Dean, Felix, Iris, Jakob, James, Janine, Joannes, Jur, Karin, Moriz, Petra, Ronald, and Zoi. (And, Zoi, this also extends to your family: you have shown me Greek

hospitality that no words can describe.) Thank you for filling my life with joy, warmth, and home, with brunches, cakes, and dinners, and with biking, Frisbees, and bouldering. For being fantastic flatmates during these PhD years, thank you Bastian (even pandemic-proven!), Ana and Zoi, and Esteban and Max.

I also would like to thank my friends far away from Amsterdam—especially Conner, Dennis, Leo, Magge, Max, Samu, and Tobi. Thank you for staying close, for the memorable bike rides and hiking trips, for the now-way-too-long-ago jam sessions, for keeping digital company during the pandemic, and for all the fun we had and will have.

Moreover, I am deeply thankful for my loving family: the Hornischers and the Schabers, including the Abbrechts and the Würzs. I would like to mention you all here, but to keep it short: thank you, Irme, for your constant interest in my life and work; thank you, grandma and grandpa Hornischer, for your continuous support and for being proud of me; and thank you, grandma Schaber, for your wise and calming words!

In particular, I would like to thank my parents, Moni and Winni, and my sister, Jele. There are so many things that I could say why I'm fortunate, grateful, proud—and simply happy—to call you family. Knowing that you are always there for me—from far and close—is invaluable. Even though it has been 'far' much too long during the pandemic, our family calls were a great source of energy for me. This thesis is dedicated to you. You mean the world to me!

Amsterdam
August, 2021.

Levin Hornischer

Chapter 1

Introduction

Computation can be distinguished into symbolic and non-symbolic. Symbolic computation is what computers do: A computer program is a more or less human-readable description of how to manipulate, step by step, a ‘symbolic’ input (like the list of numbers 5, 2, 7) to obtain a certain output (say, the ordered list 2, 5, 7 or the sum 14). This is what we typically think of as ‘computation’. But there also is a broader sense: non-symbolic computation cannot (only) be viewed as a rule-based manipulation of symbols.¹

The paradigm examples are neural networks: Given the signals from sensory neurons, your nervous system computes whether that flying something is a dangerous wasp or just a harmless fly. But nowhere in that process do ‘symbols’ occur: There are only neurons, synapses, electrical signals, etc. Symbols like ‘black and yellow’ or ‘poisonous’ at best emerge as high-level descriptions of the process. This is not only true for biological neural networks like our nervous system, but also for artificial neural networks. They are behind the recent boom of artificial intelligence (AI). For instance, if you use your phone to translate a sentence, chances are there is a neural network involved. But it doesn’t compute this translation ‘symbolically’ using grammatical rules and dictionaries. Rather it uses non-symbolic representations of words as vectors (i.e., long lists of real numbers). Similar things can be said for other applications of artificial neural networks like speech recognition or image classification. But there also are more mundane examples: An old-school mechanical thermostat computes how much hot water should flow into the radiator to maintain a desired room temperature—without

¹Cf., e.g., Bauer (2000), Blum et al. (1998), Bournez and Campagnolo (2008), Edalat (1997), Hoyrup and Rojas (2009), Pour-El and Richards (1989), Siegelmann and Fishman (1998), and Weihrauch (2000). This includes analog computation (overviewed by MacLennan 2009). Some might rather speak of ‘dynamical systems’ than of ‘(non-symbolic) computation’ (Van Gelder 1998). The thesis can also be read as simply being about dynamical systems without the computational motivation described in this introduction (we come back to this in the ‘reading guide’ below).

any digital (i.e., ‘symbolic’) representation of the temperature.²

Traditionally, symbolic computation and non-symbolic computation are taken to be diametrically opposed:³ While symbolic computation is human-readable but cannot generalize beyond clearly defined boundaries, non-symbolic computation is not interpretable but can learn well from noisy real-world data. While symbolic computation allows safety verification but needs domain knowledge, non-symbolic computation doesn’t come with complete certainty but only needs big data. (In actuality, things are, of course, more nuanced.) Only fairly recently, the focus shifted to combining the two approaches to obtain ‘the best of both worlds’. Especially due to the recent proliferation of neural networks, it is becoming increasingly important to understand their behavior—ideally by relating it to well-understood symbolic computation. This is what explainable artificial intelligence (XAI) is all about.⁴

In this thesis, we work toward a unified foundation of symbolic and non-symbolic computation. This introduction first explains what that means and sketches an idea to achieve it (section 1.1). Then we outline how the thesis develops this idea, including a reading guide (section 1.2). In addition to the informal main text, there are many footnotes and endnotes with further references and more technical topics.⁵

1.1 Motivation

A good starting point is to ask why we have a good understanding of symbolic computation.

The starting idea Symbolic computation is usually described by computer programs in some programming language (or more abstractly as, e.g., Turing machines). Such a program could, for example, implement a sorting algorithm like bubble sort which, given as input a list of numbers, transforms this input into

²A classical, but more involved example is the Watt governor which computes the appropriate setting of a steam engine to generate the desired speed of a flywheel (Van Gelder 1995). Carmantini (2017, p. 2) discusses a digital thermostat as an example of (symbolic, open-ended, reactive) computation.

³A testament to this is the at times fierce debate in cognitive science and AI between the symbolic camp (‘classicists’ or ‘Good Old-Fashioned AI’) and the non-symbolic camp (‘connectionists’ or ‘subsymbolic AI’). Appeasing voices are, e.g., Marr (1982/2010) or Smolensky (1988) taking the symbolic to be at a higher level of description than the non-symbolic.

⁴In as little as the duration of this PhD project, XAI has risen from a rather niche topic to a main field of AI. To mention but some recent surveys: Adadi and Berrada (2018), Besold et al. (2017), Doran, Schulz, and Besold (2017), Goebel et al. (2018), Murdoch et al. (2019), and Samek et al. (2019). Earlier work is done, e.g., by D’Avila Garcez, Lamb, and Gabbay (2009).

⁵Endnotes are indicated by Roman numerals and are collected at the end of this chapter. Compared to the footnotes, they contain some longer, less immediate comments. In the electronic version, they are clickable (both to get there and to get back).




Figure 1.1: An operational semantics for the bubble sort algorithm.

an ordered list of numbers. However, given such a piece of program code, we don't yet *understand* what exactly it does: this is all too common when looking at some code in an unfamiliar language. We need a *semantics* which assigns meaning to code: i.e., a description of the behavior of the program. There are two standard ways of doing this: operational semantics and denotational semantics.

Operational semantics describes the program by the changes in the states of a machine running the program (Plotkin 1981/2004). For bubble sort, this could look like in figure 1.1: For the list $\langle 5, 2, 7 \rangle$, the machine would start in state 1 with reading this input. Then it moves to state 2 and goes through the two adjacent pairs $(5, 2)$ and $(2, 7)$. Since $5 > 2$, it only swaps the first, ending up with $\langle 2, 5, 7 \rangle$. Since a swap occurred, it goes again to state 2. Now both adjacent pairs are ordered correctly, so no swap occurs. Thus, it goes to state 3 where it terminates with the correct list $\langle 2, 5, 7 \rangle$. Such an operational description can be done at various levels of detail. We could have split up state 2 into several states further describing the subprocess of comparing adjacent pairs. We could even go all the way down to the 'machine level' where a state describes the memory-entries and the processor-state of the computer.

Denotational semantics, on the other hand, describes the program by the function that it computes (Scott 1970). In the sorting example, this is a function $f : D \rightarrow D$ where D is the set (or data type) of finite lists of integers: f maps a finite list to the ordered version of that list. To still provide *some* information on how this function can be computed, denotational semantics also describes how the function can be approximated by *finite* partial functions.

Operational semantics is dynamic and closer to the implementation (or execution) of the program in terms of machine states. The advantage of the denotational semantics is that it is static and fairly independent of the implementation (e.g., which sorting algorithm exactly is used). Thus, denotational semantics is particularly suited for a structural mathematical theory and analysis (Ong 1995). After all, there the 'meaning' of the program is 'directly given' (and doesn't need to be 'dynamically constructed') and it is not restrained by implementational details.¹

Domain theory was developed as a 'mathematical theory of computation' providing a denotational semantics for programming languages (Scott 1970).

Thus, the starting idea of this thesis is: can this method be generalized also to non-symbolic computation like neural networks? In other words, can we extend domain-theoretic semantics beyond computational processes described by a precise programming language to also include more general computational processes?

On understanding Before embarking on this question, let’s discuss how this helps in the aim of understanding neural networks and non-symbolic computation more generally. There are two senses of ‘understanding’: *specific* or *structural*.

In the specific sense, we (aim to) understand some specific neural network very well: Why exactly did it classify this image as a stop sign—was it its shape, color, or location? How did it learn this concept—was it easy or hard? Can we say which of its weights store this information—making them meaningful to us? This is analogous to understanding a specific program: not just its computed function and operational description at various levels of detail, but also its efficiency, its safety, its required resources, etc.

In the structural sense, we (aim to) understand a whole class of neural networks: When should two neural networks be considered equivalent—to transfer knowledge about one to the other? Among equivalent ones, is there a most simple one—offering a ‘best explanation’ of the data? What are principled ways of combining networks—to avoid retraining? What are the limits of these networks—to assess their capabilities? The analogous questions for programs are answered by a semantics: Two programs are equivalent if they get assigned the same ‘meaning’, i.e., compute the same function; assessing the capabilities amounts to determining the class of computed functions, etc.

To be sure, this is not a sharp distinction and there is no better or worse between these two kinds of understanding: they are complementary parts of a holistic theory. In symbolic computation, these two understandings are largely achieved: For a specific understanding of a computer program the tools of, e.g., software verification can be used, while a structural understanding is provided by, e.g., computability theory and domain theory. In non-symbolic computation, XAI arguably is more focused on the specific understanding, while this thesis is concerned with the structural understanding of symbolic and non-symbolic computation.⁶

⁶For other work in this direction (mostly on neural networks), see, e.g., the computational complexity theory for neural networks (overviewed by Šíma and Orponen 2003), category-theoretic approaches (Fong, Spivak, and Tuyeras 2019; Jacobs and Sprunger 2019), the ‘(programming) language of machine learning’ (Cheung et al. 2018; Porter 1994), program induction/synthesis (Evans and Grefenstette 2018; Penkov and Ramamoorthy 2017), algebraic topology and topological data analysis (Naitzat, Zhitnikov, and Lim 2020; Reimann et al. 2017), dynamical systems approaches (Carmantini 2017; Milne 2019; Saxe, McClelland, and Ganguli 2014), statistical mechanics (Bahri et al. 2020), or statistical learning theory (Vapnik 2000). The latter (on p. xii) motivates this endeavor by the principle ‘nothing is more practical than a good theory’ (going back to Kurt Lewin).

Unified foundation So we work toward a ‘unified foundation’ of symbolic and non-symbolic computation via a domain-theoretic semantics. More precisely, we’d at least expect the following:

1. *Framework*: To talk about symbolic and non-symbolic computation in a unified manner, we need to capture them in a single framework. So, any computational process—be it symbolic or non-symbolic—can be described in this framework.

We’ll argue below that dynamical systems provide this framework.

2. *Behavior*: To systematically understand symbolic and non-symbolic computation in a uniform manner, we need to assign to any computational process (of the framework) a description of its behavior (semantics). This description should abstract away as much as possible from the implementational details of the computation and rather focus on the observable behavior or output.

We’ll use domain theory: in part 1, for symbolic computation and, in part 2, for non-symbolic computation.

3. *Relationship*: To understand non-symbolic computation in terms of symbolic computation and vice versa, we need to specify what the relationship is between the (behavior of the) two types of computation.

We’ll discuss this in part 3: non-symbolic computation is, in a sense, the limit of symbolic computation, and it can realize it if it has enough stability.

In the remainder of the present ‘motivation’ section, we describe computation as dynamical systems (the first expectation). In the next ‘outline’ section, we describe how the thesis establishes the other two expectations. Before we start, though, we review dynamical systems. (Those in the know may skip the next paragraph.)

Crash-course dynamical systems There are many formal notions of dynamical systems, but they share the following intuition. A dynamical system consists of two ingredients: a *state space* and a *dynamics*. The state space describes the possible states that the system could be in, and the dynamics describes how the system changes its states over time. A sequence of states following the dynamics is called a *trajectory* (or *orbit*).

Here are some examples: In the case of the thermostat, a state is described by the current room temperature and the amount of hot water that flows into the radiator. The dynamics is such that if the system is, for example, in a state with low room temperature and much inflowing hot water, then, at a later time, the system is in a state with high room temperature (the hot water heated up the room) and little inflowing water (the thermostat lowered the inflow since the room temperature is high). Another example is given by the sorting system of

Computation	\leftrightarrow	Dynamical systems
—symbolic	\leftrightarrow	—time- and space-discrete, possibly non-deterministic
—non-symbolic	\leftrightarrow	—time-discrete, space-continuous, deterministic

Figure 1.2: Computation as dynamical systems.

figure 1.1: it has three states and the dynamics between them (which states are possible after the current state) is described by the arrows. For the input $\langle 5, 2, 7 \rangle$, we've seen the trajectory of states $1 \rightarrow 2 \rightarrow 2 \rightarrow 3$.

There are some useful conceptual distinctions. A dynamical system is *state-discrete* if its state space is discrete (or 'countable') and otherwise *state-continuous*. (In the continuous case, there is usually more structure on the state space like a topology or a probability measure.) For example, the sorting system is state-discrete (there are only finitely many states) and the thermostat is state-continuous (the states are given by pairs of real numbers which form a continuum).

A dynamical system is *time-discrete* if its dynamics takes place in time steps: for any state we can specify the set of its possible immediate successors under the dynamics. The sorting system, for instance, is time-discrete. Thus, a trajectory is a (finite or infinite) sequence x_0, x_1, x_2, \dots of states such that each x_{n+1} is a successor of x_n . So the whole numbers $\mathbb{N} = \{0, 1, 2, \dots\}$ play the role of time. The system is *time-continuous* if, in contrast, the dynamics describes continuous change: a state doesn't have a next state but rather closer and closer states reached after smaller and smaller time intervals. The thermostat system, for instance, is time-continuous. Thus, a (maximal) trajectory is a sequence $(x_t)_{t \in \mathbb{R}}$ of states where, for $t < t'$, state $x_{t'}$ can be reached from x_t in time $t' - t$. So the real numbers \mathbb{R} play the role of time. A common trick to study a time-continuous system is to fix an 'updating time interval' and study the resulting time-discrete system. In the thermostat example, we could take this interval to be 1 second and declare the successor of a state to be the state reached after 1 second. For this reason, we'll restrict our attention to time-discrete systems.

Finally, we say a time-discrete system is *deterministic* if each state has a unique successor state. Otherwise, we call the system *non-deterministic*.

Computation as dynamical systems Now we get to expectation (1): how computation can be described by dynamical systems as in figure 1.2.⁷

Let's start with symbolic computation. We've already seen that bubble sort can be described as the time- and space-discrete dynamical system of figure 1.1. This holds for other examples as well: the general argument is the following. Turing machines are regarded as *the* model of symbolic computation. And they can be described as time- and space-discrete dynamical systems: A state of a

⁷Although this idea has been around for a while (see, e.g., Giunti 1997; Siegelmann and Fishman 1998), it arguably deserves more appreciation.

Turing machine is described by: (i) what is written on its tape (i.e., its memory), (ii) which part of the tape is currently observed (and to be altered), and (iii) the internal state of the machine.⁸ So there are countably many states $x = (a, b, c)$ where a describes the tape, b the observed part, and c the internal state. The program of the machine describes the possible transitions. There is a transition from $x = (a, b, c)$ to $x' = (a', b', c')$ iff, roughly, the program says: when in internal state c reading symbol $a(b)$ at location b of the tape a , change it to $a'(b)$ and leave everything else unchanged (so $a' = a$ outside of b) and go to position b' and into internal state c' . If the Turing machine is non-deterministic, these dynamics are non-deterministic, too. Thus, any symbolic computation can be seen as a time- and space-discrete dynamical system that possibly is non-deterministic. We'll also allow some labeling of the system: we can label a state as the initial state and add labels to transitions between states. So we can describe this class of dynamical systems as the well-known (countable) labeled transition systems.

Now, let's consider non-symbolic computation and how it can be described as time-discrete, space-continuous, and deterministic dynamical systems. Our paradigm examples are such systems: We've already seen the thermostat. And, importantly, also artificial neural networks are such systems. A state of the network describes the activation that each neuron has at that moment. And the dynamics is given by how this activation propagates through the network: The activation of any neuron at the next time step is determined by how much input it receives from its neighboring neurons weighted by the weight on their connection. So this system is deterministic, time-discrete, and space-continuous (activation is usually given by real numbers). This describes the 'run time' dynamics of the network. But also its learning dynamics—i.e., computing the best approximation to some observed data—can be seen as such a system. Then a state describes the weights on the connections between neurons. And the dynamics is given by the learning algorithm (e.g., backpropagation): Given some observed data, the current weight-state is updated according to the algorithm to a new weight-state in which the network better approximates the observed data. Finally, the general answer is that 'real-world systems'—be it physical, chemical, biological, or engineering—usually are described via differential equations as dynamical systems.⁹ This makes them deterministic, space-continuous and, after fixing an update time interval, also time-discrete.¹¹

Thus, dynamical systems do indeed provide a framework to describe (presum-

⁸Turing (1936-7, p. 250) writes: "We know the state of the system if we know the sequence of symbols on the tape, which of these are observed by the computer . . . , and the state of mind of the computer". The internal state or 'state of mind' or ' m -configuration' is not to be confused with the system's state.

⁹Obtaining such a description for the nervous system is a central topic of the field of computational neuroscience. The dynamics of the nervous system can be observed, e.g., through EEG and fMRI. Anderson (2015, pp. 22–25) describes what thus can be observed when the nervous system implements symbolic computation: namely, equation solving.

ably) anything that can reasonably be said to be computation—both symbolic and non-symbolic.¹⁰ They capture the essence of a program: what states the system implementing it can be in (state space) and how they are transformed according to the program (dynamics). The ‘symbolicity’ of the computation is reflected in the discreteness of the state space. We deliberately leave open the converse question: whether any such dynamical system also constitutes computation. In other words, whether dynamical systems provide an *explication* of symbolic and non-symbolic computation. This is a deeply philosophical question.¹¹

Nonetheless, we do want to provide some evidence that dynamical systems are, if not an explication, at least a good (qualitative) *approximation* to non-symbolic computation.¹² The argument sketch is this: If anything counts as non-symbolic computation, this should include neural networks. By the universal approximation theorems (Cybenko 1989; Hornik, Stinchcombe, and White 1989), neural networks can approximate any dynamical system arbitrarily well.^{III} And, plausibly, the class of non-symbolic computational processes is closed under such approximation (if effective).¹³

1.2 Outline

With this framework in place, we outline the thesis: how it develops a domain-theoretic semantics for the dynamical systems describing symbolic computation (part 1) and non-symbolic computation (part 2). This can then be used to relate symbolic and non-symbolic computation (part 3).

Part 1 (symbolic computation) In chapter 2, we develop the trajectory domain construction. It assigns to each time- and space-discrete dynamical system (representing some symbolic computation) its trajectory domain. This is a structure in the sense of domain theory describing the behavior of the system.

¹⁰The reason for saying ‘presumably’ is that this is a thesis, not something that can be proven. A counterexample would need to provide a clearly computational processes that can in no way be seen as a dynamical system. This is hard to imagine. In any case, this determines an extremely general class of systems.

¹¹For some discussion, see the literature on ‘physical computation’ cited in footnote 18 below. For more general literature on this nexus of dynamical systems, computation, cognition, and logic, see, e.g., Van Gelder (1995), Giunti (1997), Van Gelder (1998), beim Graben (2004), Leitgeb (2005), Tabor (2009), and Dewhurst (2016).

¹²One may distinguish between a ‘qualitative’ and ‘quantitative’ study of computability (cf. e.g. Abramsky and Jung 1994, sec. 1.1). In the former, one studies, e.g., the topological, algebraic, or order-theoretic structure of computation—just as any other structure of classical mathematics. In the latter, one adds a ‘computability’ or ‘effectiveness’ structure to determine the ‘computable’ elements of these classical structures. (Cf. Bauer (2000) and Pour-El and Richards (1989).) Here we’re mostly concerned with a qualitative study of computation. Though, adding effective structure in domain theory is well-understood (Edalat 1997).

¹³Cf. computable real numbers are closed under effective convergence (Turing 1936-7, p. 256).

The idea behind the construction is as follows. In the case of programming languages, we usually still have information, for example, about what type of input-output function the program aims to compute. This is not available anymore for general time- and space-discrete systems, but we still have the trajectories. Each trajectory is an instance of the behavior of the system. So the rough idea is that the trajectory domain is the set of the finite and infinite trajectories ordered by extension: the infinite trajectories are the infinite ('ideal') limit-behaviors which are approximated by the finite ('real') behaviors. (This idea is refined by taking into account that we often want to exclude some trajectories and consider others to be instances of the same type of behavior; so, really, the trajectory domain consists of equivalence classes of some set of trajectories.) Thus, the trajectory domain can be said to describe the (types of) behavior of the system. Chapter 2 describes and axiomatizes this construction.

A main reason why a semantics provides structural understanding is that it not only assigns meaning to programs ('syntax') but also preserves relations between them (this is known as *compositionality*). In our case, the crucial relation between systems is that of *simulation*: one system S can be simulated by another S' if there is a function $f : S \rightarrow S'$ assigning states of S to states of S' preserving the dynamics.¹⁴ So we would expect that the trajectory domain semantics respects this. Indeed, in chapter 3, we show that the simulation f can be assigned to an appropriate function from the trajectory domain of S to the trajectory domain of S' .¹⁵ Category-theory provides the language to describe this more precisely: the trajectory domain construction is a functor from systems to domains that, in many cases, even forms a so-called adjunction. This is in line with the general idea that syntax (formal) and semantics (conceptual) should be adjoint.¹⁶

Part 2 (non-symbolic computation) When we move to state-continuous systems, we cannot 'access' the states anymore: they are infinitely precise points in a continuous state space. All we can do is measurements. For example, if the thermostat system is in state $(20.071\dots, 0.183\dots)$, we cannot precisely determine this but only measure that, say, the room temperature is $20^\circ C$ plus or minus $1^\circ C$ error in measurement and the incoming hot water flow is $0.2 \frac{kg}{s}$ plus or minus 0.1 error. So we can only determine the *area* of the state space in which the system is in: namely, $[19, 21] \times [0.1, 0.3]$.

These areas (that correspond to measurements) act much like the states in symbolic computation. While, in the symbolic case, repeated observation yielded a sequence of states (i.e., a trajectory), it now yields a sequence of areas of the state space. Thus, we'll use these sequences to build a domain describing the behavior of the system with respect to the available observations. We may also increase

¹⁴I.e., if $x \rightarrow y$ is a transition in S , then $f(x) \rightarrow f(y)$ is a transition in S' .

¹⁵The 'appropriate functions' of domain theory are known as Scott-continuous functions.

¹⁶This is expressed by Lawvere (1969) and explained by Smith (Unpublished).

the precision of our measurements and the observation time—thus refining the available observations. This then also refines the corresponding domain. The main result of chapter 4 is that, as we keep refining the observations, we obtain in the limit again a domain. We call this the *observation domain* of the system. It is a model of the system since, based on the observations, we can also define a dynamics on it and, when restricted to its ‘ideal’ elements, this dynamics is isomorphic to the original system. This model is ‘computational’ since these ideal elements are approximated by the ‘real’ elements given by finite observations.

Thus, the observation domains provide a semantics (or behavior description) to the dynamical systems describing non-symbolic computation. Chapter 5 again verifies that this semantics preserves simulations between dynamical systems. In fact, it establishes a translation (i.e., categorical equivalence) between dynamical systems and certain purely domain-theoretic structures that we call dynamical domains (of which observation domains are examples). This may well be regarded as the main formal result of the thesis.

To sketch applications of this translation, chapter 6 provides a domain-theoretic perspective on a central concept of dynamical systems theory: entropy.

Part 3 (stability) Given these semantics for symbolic and non-symbolic computation, we can turn to expectation (3): What is the relationship between the two types of computation? Chapter 7 is an interlude where we first discuss this informally.

On the one hand, our semantics suggests the thesis that (the behavior of) non-symbolic computation is the limit of (the behavior of) symbolic computation:¹⁷ roughly, the observation domain can be expressed as a limit of trajectory domains describing observation sequences.^{IV}

On the other hand, this raises the question of when non-symbolic computation can be regarded as realizing symbolic computation. As a guiding intuition, we suggest that the system’s behavior should be fairly stable. For example, from the behavior of our nervous system recognizing a wasp we can extract the symbolic rule “if it flies and is black and yellow, it is dangerous”. The reason seems to be that in most of the ‘continuously many’ inputs where the system recognizes something flying that is black and yellow, it will compute that it is a dangerous wasp. So this assessment is stable under a wide range of input states. Other examples are physical realizations of symbolic computation (like my laptop): as physical systems these are (described as) continuous dynamical systems and a terminating computation usually corresponds to a stable state of the system.¹⁸

The interlude chapter sketches what our results together with deep results

¹⁷In analogy with profinite groups, one might say that non-symbolic computation is pro-symbolic computation.

¹⁸For literature on physical computation, see Chalmers (2011), Fredkin and Toffoli (1982), Gandy (1980), Lloyd (2000), Piccinini (2015), Piccinini (2017), Pitowsky (1990), and Sieg (2002).

from ergodic theory can already say about this kind of stable behavior. And it explores how (algorithmic) randomness may help to ensure this kind of stability.

Chapter 8 then begins investigating these ideas in detail. It starts at the foundations with a philosophical analysis of the involved concept of stability. This is done with an eye toward AI-safety: demanding that, like the nervous system, also artificial neural networks should be stable under small perturbations of the input.

Reading guide This thesis has grown rather long—apologies! In an attempt to make up for it, it allows for modular reading. Despite its monographic structure, each chapter can be read independently. If material from previous chapters is needed, it is summarized. Every chapter starts with a non-technical introduction that motivates and summarizes the results. This should allow, for example, skipping more technical parts of a chapter while still getting the gist of it. The suggested order of reading is, non-surprisingly, in order of appearance. For a less formal track, one might skip the category-theoretic chapters 3 and 5 and/or the entropy-theoretic chapter 6. (The usage of category theory is largely restricted to formulating the results concisely rather than actually using the theory.) For a less philosophical track, one might skip chapter 8 on stability.

For clearer (and more common) terminology, we use ‘labeled transition system’ in the case of symbolic computation and simply say ‘dynamical system’ for (typically time-discrete and space-continuous) deterministic dynamical systems.

Skeptics of the sketched view on computation can also read the thesis without this computational interpretation. This may include classical computationalists who like computation but think that, by definition, it cannot be non-symbolic; or embodied cognitive scientists who like to view cognition as a dynamical system but wouldn’t call it ‘computation’. Labeled transition systems and dynamical systems are important mathematical structures in their own right, so our semantics/representation should be useful regardless.

- I. Semantics for programming languages is somewhat analogous to semantics for formal languages or sufficiently regimented fragments of natural languages—as they are considered, e.g., in the philosophy of language (Speaks 2021). Both assign meaning (or ‘semantic content’) to expressions of the language. The classical semantic theories à la Frege, Russell, or Tarski resemble denotational semantics: assigning to an expression a static (mathematical) object describing its meaning. Inferentialist (or proof-theoretic) semantics loosely resemble operational semantics: describing the meaning of an expression by its inferential interaction with other expressions.
- II. As mentioned, we’ll focus on time-discrete systems: Partly since they approximate the time-continuous systems and partly since artificial neural networks are commonly time-discrete only. Nonetheless, future work should investigate whether our results extend to time-continuous systems. One might try adapting our approximation process described in part 2 below by not only approximating space (through measurements) but also time (increasingly finer discretizations). (Also cf. the generator theorems for flows (Eberlein 1974).) For an overview of models of continuous time computation see Bournez and Campagnolo (2008) and Orponen (1997).
- III. More precisely: Given a dynamical system (X, T) with X an uncountable standard Borel space and $T : X \rightarrow X$ Borel-measurable (which are very minimal assumptions), we can assume, by the Borel isomorphism theorem (see e.g. Kechris 1995, thm. 15.6), that $X = \mathbb{R}$. Let $\epsilon > 0$ be the precision to which we want to approximate (X, T) with a neural network. By the universal approximation theorem, there is a feedforward neural network N with one input neuron, one hidden layer, and one output neuron such that N , regarded as a function $N : X \rightarrow X$, is ϵ -close to T (for an appropriate choice of metric). Let M be the recurrent neural network obtained from N by feeding the output into the input. Now, consider the activation dynamics of M : Since the activation-state of M is determined just by the input neuron, we identify an activation-state of M with the activation $s \in \mathbb{R} = X$ of the input neuron. So, the state space is X . Regarding the dynamics, if, M is in state s , then the activation $N(s)$ of the output layer will be fed into the input layer, so the new state is $s' = N(s)$. Hence the activation dynamics of M is the dynamical system (X, N) . And (X, N) approximates the original (X, T) up to precision $< \epsilon$. (For more on universal approximation, see Kratsios (2020).)
- IV. In this light, one may view the result of Pour-El and Richards (1981) that computable initial conditions of physical systems can lead to non-computable solutions. This is taken to show that these (non-symbolic) systems cannot be simulated by (symbolic) Turing machines (Pitowsky 2002, S169).

Part One

Symbolic computation

Chapter 2

Trajectory domains 1: Construction

Abstract With the aim of providing a denotational semantics (or behavior description) to the widely used labeled transition systems (LTS), we introduce the notion of a behavioral transition system (BTS). These are structures $M = (A, T, \equiv)$ where A is an LTS, T a set of trajectories (or paths) in A , and \equiv an equivalence relation on T satisfying five axioms. While any trajectory is ‘locally possible’, T describes which are ‘globally possible’; and \equiv describes when two trajectories are instances of the same type of behavior—so the equivalence classes represent possible behaviors. The main result is that, for countable systems, there is, roughly, a unique way of defining an information containment order between behaviors and this yields an ω -algebraic domain. We call this the trajectory domain and think of it as the denotation of M . We also show that BTSs (and their trajectory domains) generalize both Scott information systems and various models of concurrent computation (and their respective domain constructions).

2.1 Introduction

We’re concerned with providing denotational semantics (or behavior description) to labeled transition systems. Let’s explain:

A *labeled transition system* (LTS) is a structure (S, i, L, \rightarrow) where S is a set of states, $i \in S$ is the initial state, L is a set of labels (or actions), and $\rightarrow \subseteq S \times L \times S$ is relation, written $s \xrightarrow{\alpha} s'$. LTSs are a general model of *computing systems*. They include ‘sequential’ computing like Turing machines: intuitively, a state consists of the values stored in the memory of the machine at a given time step, and a label is a command that can be executed to processes some stored values leading to a new state. But they also include ‘non-sequential’ computing as in reactive systems: the system (a standard example is a vending machine) interacts with—i.e., reacts to—a non-deterministic environment (users can insert coins and select items) in an open-ended way (the system doesn’t aim to compute a specific outcome).¹ Due to

¹LTSs can also be regarded as time-discrete dynamical systems: they consist of a state space

this generality, LTSs are prominently used in *model checking* which is a standard technique to formally verify that a computing system behaves as intended (Baier and Katoen 2008). This ranges from the standard examples of ensuring safety in money transfers or space flight to examples that recently gained prominence: the verification of neural networks as a way to address the safety concerns raised by their intransparency.²

As computational models, LTSs describe *how* the computation proceeds: an *operational* description that is dynamic and close to ‘machine implementation’ (i.e., the states of the execution of the computation). But, we may ask, *what* is it that they compute: is there a *denotational* description of their behavior that is static and more ‘machine-independent’ (i.e., abstracting away implementational details and facilitating mathematical analysis)? In the case of programming languages, these two complementary advantages are associated with operational and denotational semantics, respectively (Ong 1995). Roughly speaking, LTSs operationally describe programming code by how it transforms states of the computer (Plotkin 1981/2004), and domain theory denotationally describes programming code by the function that it computes and how it is obtained from other functions (Scott 1970). (If the semantics coincide, one speaks of full abstraction (Cardone 2021; Ong 1995).³)

Given these advantages, we’d like to develop a denotational semantics (or behavior description) for *any* LTS. However, in general—as, e.g., with reactive systems—we neither have available a programming language (or typed meta-language) nor an input-output description.⁴ What we still have, though, are the *trajectories*: the (finite or infinite) sequences of the form

$$s_0 \xrightarrow{\alpha_1} s_1 \xrightarrow{\alpha_2} s_2 \xrightarrow{\alpha_3} s_3 \xrightarrow{\alpha_4} \dots$$

This is an instance of a possible behavior of the system. However, depending on the level of abstraction at which we analyze the behavior of the system, we may want to identify some trajectories as instantiating the same type of behavior (e.g., two concurrent computations). Moreover, some trajectories may only be locally possible (from each s_{i-1} one can move to s_i via label α_i) but not globally possible (e.g., some action can only be applied a certain number of times due to, say, memory constraints).

Thus, to describe the behavior of an LTS A , we’re lead to also specify a set T of ‘globally possible’ trajectories in A and an equivalence relation \equiv on T . So the quotient $\mathbb{T} := T / \equiv$ is the set of possible behaviors and can be regarded as the denotation of A —or, rather, of (A, T, \equiv) .

S together with a dynamics \rightarrow describing how the system can transform from one state into another.

²For the former, see the textbook of Baier and Katoen (2008). For the latter, to mention but two references, see Kuper et al. (2018) and Vengertsev and Sherman (2020).

³In a somewhat more abstract setting, also see Abramsky (1991).

⁴Cf., e.g., Winskel and Nielsen (1995, p. 2) or Carmantini (2017, p. 2).

For a satisfying treatment, however, we should expect some more structure on these denotations (as it also is the case in domain-theoretic semantics for programming languages). Indeed, we intuitively also would expect \mathbb{T} to be ordered by information containment: behavior $[t]$ is informationally contained in behavior $[t']$ if each instance t_0 of $[t]$ can be extended to an instance t_1 of $[t']$. This poses the question that we investigate: When can such a notion of information containment be appropriately defined on \mathbb{T} and when does this then form a domain—so \mathbb{T} can satisfyingly be said to be the denotation of the LTS. The answer will be: We provide some axioms for the structures (A, T, \equiv) to define an appropriate notion of information containment. For countable systems, it turns out that this notion is, in a sense, unique and turns \mathbb{T} into a domain—indicating that we’ve found a stable axiomatization. We call structures (A, T, \equiv) satisfying these axioms *behavioral transition system (BTS)* and, in the countable case, we call \mathbb{T} their *trajectory domain*.

The chapter is structured as follows: In section 2.2, we provide the relevant background on labeled transition systems and domain theory.

In section 2.3, we discuss two guiding examples: First, LTSs arising from observing a ‘black box’ system like those in statistical mechanics or neural networks. Second, LTSs arising as models of concurrent computation. These provide more concrete motivation for studying the structures (A, T, \equiv) and their set of behaviors \mathbb{T} —in addition to the abstract motivation above.

In section 2.4, we introduce pre-behavioral transition system (pre-BTS) as structures (A, T, \equiv) satisfying a minimal set of axiom capturing that \equiv describes ‘trajectory equivalence’. In section 2.5, we consider various natural ways of defining information containment on \mathbb{T} and find that they coincide once (A, T, \equiv) is what we’ll call limit-respecting.

In section 2.6, we then show the main result: Roughly, for a countable pre-BTS (A, T, \equiv) and a preorder \preceq on \mathbb{T} , the following are equivalent: (a) \preceq satisfies some rather weak properties that we’d expect from an information containment, (b) the partial order induced by (\mathbb{T}, \preceq) is a domain (the trajectory domain), and (c) the system is limit-respecting and \preceq is one of the coinciding natural notions of information containment.

This then suggests defining BTSs as pre-BTSs that are limit-respecting. In section 2.7, we investigate this notion and discuss several examples. In section 2.8, we describe, for countable systems, their trajectory domain and show that every ω -algebraic domain arises as the trajectory domain of a system.

In section 2.9, we show that countable BTSs and their trajectory domains can be regarded as a generalization of the well-known Scott information systems and their induced Scott domains. Finally, in section 2.10, we conclude with some open questions.

Further related work is discussed in the subsection on concurrent computation (section 2.3.2): In short, for various models of concurrent computation, it has been shown that the computation traces or sequences form the domain of concurrent

computations under a certain partial order (see that section for references). Our BTSs provide a general framework containing these models and generalize this idea considerably (in fact, the main result determines just how much it can be generalized).

2.2 Background

2.2.1 Labeled transition systems

There is a huge amount of literature on (labeled) transition systems. Here we follow the handbook article of Winskel and Nielsen (1995) since it offers a particularly systematic treatment: it not only describes labeled transition systems but also their connections to other computational models in a structural way (using category theory).

A *transition system* is a structure (S, \rightarrow) where S is a set and \rightarrow a binary relation on S . In other words, (S, \rightarrow) is a directed graph. The elements of S are called *states* and $s \rightarrow s'$ a state *transition*. Often, one also singles out an *initial state* $i \in S$ and writes (S, i, \rightarrow) . Labeled transition systems are obtained—as the name suggests—by adding labels:

2.2.1. DEFINITION. A *labeled transition system* (LTS) A is a structure (S, i, L, \rightarrow) where S is a set of *states* with *initial state* i , L is a set of *labels*, and $\rightarrow \subseteq S \times L \times S$ is the *transition relation*. We write $s \xrightarrow{\alpha} s'$ for $(s, \alpha, s') \in \rightarrow$. Given an LTS A , we use S_A , i_A , L_A , and \rightarrow_A to refer to its set of states, initial state, set of labels, and transition relation, respectively. We call A *countable* if both S and L are countable sets.

Sometimes, LTSs are *defined* to be countable. This is indeed the typical case—especially when regarding LTSs as models of *symbolic* computation. But, generally speaking, it is advisable to distinguish considerations of structure from those of cardinality.⁵ So rather than generally making the countability assumption, we develop much of our theory without it and explicitly mention the assumption if we need it.

We use the usual notation for sequences: Formally, a finite or infinite sequence σ over a set A is a partial function $\sigma : \omega \rightarrow A$ whose domain is of the form $\{n \in \omega : 0 \leq n < l\}$ where $0 \leq l \leq \omega$ is the *length* of the sequence, denoted $|\sigma|$.⁶ If $l = 0$, then σ is the *empty* sequence ϵ . If $l = \omega$, then σ is *infinite*; otherwise σ is *finite*. (So by an ‘infinite sequence’ we always mean a sequence of length ω , i.e., we won’t consider sequences whose length is an ordinal number $> \omega$.) We

⁵The countability assumption would, for example, preclude taking limits of LTSs—just based on issues of cardinality, not due to structural constraints.

⁶Here ω denotes the first infinite ordinal (so ω can be thought of as the set of natural numbers $\{0, 1, 2, \dots\}$).

often just write σ as $\sigma(0)\sigma(1)\dots$. For $n \in \omega$, we define the *restriction* $\sigma \upharpoonright n$ as the restriction of the partial function σ to the set $\{m \in \omega : m < n\}$. So if $|\sigma| \geq n$, then $|\sigma \upharpoonright n| = n$, and if $|\sigma| \leq n$, then $\sigma \upharpoonright n = \sigma$. A sequence σ' is an *extension* of another sequence σ (written $\sigma \preceq \sigma'$) if, for all $n \in \omega$, if $\sigma(n)$ is defined, then $\sigma'(n)$ is defined and $\sigma(n) = \sigma'(n)$. We write $\sigma \prec \sigma'$ if $\sigma \preceq \sigma'$ and $\sigma \neq \sigma'$.

2.2.2. DEFINITION. Let $A = (S, i, L, \rightarrow)$ be an LTS. An *A-trajectory* is a sequence

$$t = (s_0, \alpha_0, s'_0), (s_1, \alpha_1, s'_1), \dots, (s_n, \alpha_n, s'_n), \dots$$

of elements of \rightarrow such that $s'_i = s_{i+1}$. We then write $s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots$. If t is nonempty, we call s_0 the *starting state* of t and, if t also is finite, we call the s' of the last entry the *ending* or *last state* of t , which we refer to by ‘last(t)’. We refer to $s(t) := s_0, s'_0, s'_1, s'_2 \dots$ and $l(t) := \alpha_0, \alpha_1, \dots$ as the *state sequence* and *label sequence* of t , respectively.

One can also consider morphisms between LTSs: that one system can simulate the other. Thus, one can form the category of LTSs, but for our present purposes we don’t need to do this.

2.2.2 Domain and order theory

We recall the basic concepts from order and domain theory that we’ll use. A standard reference is Abramsky and Jung (1994).

A *preorder* is a structure (P, \leq) where P is a set and $\leq \subseteq P \times P$ a reflexive ($\forall x : x \leq x$) and transitive ($\forall x, y, z : x \leq y, y \leq z \Rightarrow x \leq z$) relation. A subset $A \subseteq P$ is *directed* if it is nonempty and, for all $x, y \in A$, there is $z \in A$ with $x, y \leq z$. A *partial order* is a preorder (P, \leq) that is antisymmetric ($\forall x, y : x \leq y, y \leq x \Rightarrow x = y$).

If (P, \leq) is a preorder, the *induced* partial order $(\overline{P}, \overline{\leq})$ is the quotient under the equivalence relation $x \sim y$ iff $x \leq y$ and $y \leq x$: To be precise, \overline{P} is the set of \sim -equivalence classes, which we denote $[x]_{\leq}$, and $[x]_{\leq} \overline{\leq} [y]_{\leq}$ iff $x \leq y$ (this is independent of the representatives x and y).

Let (P, \leq) be a partial order. It has a *least element* if there is $x \in P$ such that, for any $y \in P$, $x \leq y$. If existent, such x is unique and usually denoted \perp . A subset $A \subseteq P$ has a *least upper bound* (or *supremum*) if there is $x \in P$ that is an upper bound ($\forall a \in A : a \leq x$) and that is the least one (if y also is an upper bound of A , then $x \leq y$). If existent, such x is unique and denoted $\bigvee A$. A *directed complete partial order* (*dcpo*) is a partial order in which every directed subset has a least upper bound.

Let (D, \leq) be a dcpo. An element $x \in D$ is *compact* if, for all directed subsets A of D , if $x \leq \bigvee A$, then there is $a \in A$ such that $x \leq a$. We write $K(D)$ for the set of compact elements of D . Finally, (D, \leq) is *algebraic* if, for all $x \in D$, the set $\{y \in K(D) : y \leq x\}$ is directed and its least upper bound is x . If $K(D)$ also is

countable, we call (D, \leq) an ω -algebraic domain. (The more general concept is that of a continuous domain, but we don't need that here.)

A function $f : D \rightarrow E$ between dcpos is (*Scott-*) *continuous* if it is monotone ($\forall x, y : x \leq y \Rightarrow f(x) \leq f(y)$) and preserves directed suprema (for all directed $A \subseteq D$, $f(\bigvee A) = \bigvee f(A)$). (Note that $f(A) := \{f(a) : a \in A\}$ is directed by monotonicity.) Two dcpos D and E are *isomorphic* iff they are order isomorphic, i.e., there is a surjective $f : D \rightarrow E$ such that, for all $x, y \in D$, $x \leq y$ iff $f(x) \leq f(y)$ (the latter implies injectivity, so f is bijective).⁷

Let (P, \leq) be a preorder. A subset $I \subseteq P$ is an *ideal* if it is a downset ($\forall x, y : x \leq y, y \in I \Rightarrow x \in I$) and directed. An ideal is *principal* if it is of the form $\downarrow x := \{y \in P : y \leq x\}$. The *ideal completion* $\text{Idl}(P, \leq)$ of (P, \leq) is the set of ideals ordered by inclusion. If (D, \leq) is an algebraic dcpo, then (D, \leq) is isomorphic to $\text{Idl}(K(D), \leq)$ (Abramsky and Jung 1994, prop. 2.2.25).

We'll use the following simple but fundamental fact about countable directed preorders (often without explicitly mentioning it).

2.2.3. LEMMA. *Let (P, \leq) be a countable and directed preorder. Then P has a cofinal chain $C = \{c_0, c_1, \dots\} \subseteq P$, i.e., $c_0 \leq c_1 \leq \dots$ and, for all $x \in P$, there is n with $x \leq c_n$.*

PROOF. Since P is countable, write $P = \{x_0, x_1, \dots\}$. Construct $c_0 \leq c_1 \leq \dots$ by: $c_0 := x_0$ (note P is nonempty), and given c_n , let k be the least index such that $x_k \geq c_n, x_n$ (such upper bounds exists by directedness, whence there also is one with least index), and define $c_{n+1} := x_k$. Then, by construction, C is indeed a chain and cofinal in P .⁸ \square

2.3 Two guiding examples

We describe two examples to motivate the abstract structures that we subsequently investigate. To keep to the point, the discussion will be more intuitive and not strictly formal.

2.3.1 Observing a black box system

In this subsection, we describe our initial motivation for the present work: observing a 'black box' system like a neural network.

⁷Equivalently, D and E are isomorphic in the category consisting of dcpos and Scott-continuous functions (i.e., there are Scott-continuous functions $f : D \rightleftharpoons E : g$ that compose to the respective identity functions).

⁸This proof is given by Schweber (2016). Also see Abramsky and Jung (1994, prop. 2.2.13).

Black box system As a guiding example, we consider the following situation: We're given a 'black box' and 'low-level' deterministic system (X, f) and we'd like to make sense of it at a higher level through observations. So X is a set of (low-level) states and $f : X \rightarrow X$ is a function. We can also write this as a (unlabeled) transition system with state space X and transitions $s \rightarrow s'$ iff $f(s) = s'$. The intuitive terms 'black box' and 'high/low level' are best illustrated by examples.

First, statistical mechanics: A state s is, say, a list of the position and momentum of each gas particle in a box of gas, and $s \rightarrow s'$ iff, whenever the system is in state s , the laws of classical mechanics determine s' as the state in the next time step. This is a deterministic dynamical system whose laws we fully understand, but it is a 'black box' system in the sense that it is not feasible to determine the exact state of the system at a given time. Statistical mechanics is about relating the microscopic or 'low-level' description of system states to macroscopic or 'high-level' states like temperature or pressure that are more meaningful to a human observer.

Second, neural networks: A state s is a list of the value of each weight of a neural network during a training process, and $s \rightarrow s'$ iff, whenever the system is in state s , the learning algorithm (e.g., backpropagation) determines s' as the next state given a data point d . (Here we could take d as the label for the transition $s \rightarrow s'$.) This, too, is a dynamical system whose laws we fully understand (we can even program it) and which is deterministic once the order of data points is fixed. But it is a 'black box' system in the sense that it is very difficult to relate the macroscopic or high-level properties of the system (e.g., whether the network classifies this image as depicting a cat) to the microscopic or low-level properties of the states (e.g., which value a certain weight has).⁹

Given the generality of the structure of (X, f) , *many* more examples are possible. For illustrative purposes, we consider the well-known North-South map (Walters 1982, sec. 5.1, ex. 8). It is much simpler than, say, the neural network example, but it still displays important qualitative similarities: e.g., stable fixed points (convergence) and non-stable fixed points (divergence).

The North-South map is the system (X, f) depicted on the left of figure 2.1: The state space X consists of the points on the circle and the dynamics $f : X \rightarrow X$ is defined as follows: If $x = p$ is the 'North Pole', then $f(x) = x$. Otherwise, draw a line from p through x and go to where it intersects the real line (the horizontal line), then go to the midpoint from the origin q (the South Pole), and draw a line back to the North Pole: the intersection of this line with the circle is the new state $f(x)$. Thus, any state $x \notin \{p, q\}$ will move under the dynamics closer and closer to the South Pole q . Moreover, both the North Pole p and the South Pole q are fixed points. However, the North Pole is unstable in the sense that every close-by

⁹For an overview of the connections between statistical mechanics and deep neural networks, see (Bahri et al. 2020). Also see our chapter 7.




Figure 2.1: The North-South map (left) and the observed system under the partitioning into the four sets N, E, S, W (right).

state $x \neq p$ (to which the system might be perturbed to from p) will move away from p , while the South Pole is stable in the sense that all close-by states $x \neq q$ will converge back to q .

Observed system Next, what does it mean to observe the system? For simplicity, we'll identify possible observations (i.e., observable properties) with subsets of the state space: To make observation P if the system is in state s corresponds to coming to know that s has property P , i.e., that s is in the set P of states that have property P .

For example, the observable property that the network classifies a given picture correctly as depicting a cat corresponds to the set P of weight-states where the network shows this classification behavior. In the North-South map, assume we have a way to observe whether the system is in the *North*, *East*, *South* or *West* arc as indicated in the figure.

When we regard the system through these possible observations, we see the *observed system*: Its states are the possible observations that we can make, and $P \rightarrow Q$ iff there is $x \in P$ with $f(x) \in Q$, i.e., if we make observation P now, then we can make observation Q next. For the North-South map, it is depicted on the right of figure 2.1. In particular, we neither can have $N \rightarrow S$ (any orbit starting in N has to go through W or E to get to S), nor $W \rightarrow W$ (starting in W will take one outside W), nor $W \rightarrow E$ (once orbiting down the left side of the circle, one cannot go to the right side anymore). Also note that unlike the original deterministic system, the observed system need not be deterministic.

Observation topology A more general perspective on the observed system is as follows. (This more technical paragraph can be skipped.) We've said that observable properties are subsets of X , but which subsets are observable properties? We take it to be those subsets P for which we have finite decision procedure

to tell whether the system, in a given state, has the property P or not. In the ‘cat picture’ example, we have such a procedure: given weight state s , input the picture to the neural network, let it run and see whether it provides a positive answer (i.e., it has P) or a negative answer (i.e., it doesn’t have P).

Let’s write \mathcal{B} for the set of these ‘decisively observable properties’. We treat it as a variable, but given this intended interpretation it makes sense to demand:

- \mathcal{B} is a Boolean algebra: if we can decisively observe P and also Q , then we also can decisively observe P^c , $P \cap Q$, and $P \cup Q$,
- \mathcal{B} is closed under f -preimage: if we can decisively observe P , we can decisively observe $f^{-1}(P)$, because to see whether s is in $f^{-1}(P)$, we see whether $f(s)$ has P , i.e., we wait one time step and see whether the system has property P ,
- \mathcal{B} is countable: the decision procedures need to be accessible to us, so we at least need to be able to enumerate them.¹⁰

Note that we’re considering *decidable* observable properties and not *semi-decidable* observable properties which, famously, form a topology rather than a Boolean algebra (Smyth 1983; Vickers 1989). In the North-South map, \mathcal{B} could be the closure under Boolean operations and f -preimages of $\{N, E, S, W\}$.

Given our collection \mathcal{B} of decisively observable properties, we wonder what are the possible ways things can be according to these observations. In other words, what are the possible complete and consistent collections of properties that the system could have at a given point in time? These are known as the ultrafilters of \mathcal{B} . Every state s induces such an ultrafilter (the set of properties $P \in \mathcal{B}$ that s has). If things go well, also every ultrafilter F determines a unique state (the state s which has all the properties $P \in F$), and if not, we may think of F as an ‘imaginary’ state that ‘logically completes’ the state space of the system. The set of ultrafilters of \mathcal{B} is denoted $\mathbf{Spec}(\mathcal{B})$ and we may call it the set of *logical* states of our system (X, f) .

This set $\mathbf{Spec}(\mathcal{B})$ of logical states has a natural topology induced by the basic open sets of the form $\{F : P \in F\}$ for $P \in \mathcal{B}$. This is a Stone space: zero-dimensional, compact, and Hausdorff. (That is the classic Stone duality: the correspondence between Boolean algebras and Stone spaces.) Since \mathcal{B} is countable, $\mathbf{Spec}(\mathcal{B})$ also is second-countable and hence a compact metrizable space.

Moreover, the dynamics $f : X \rightarrow X$ naturally extends to a dynamics $\bar{f} : \mathbf{Spec}(\mathcal{B}) \rightarrow \mathbf{Spec}(\mathcal{B})$ on the logical states: Since \mathcal{B} is closed under preimage, the function $h : \mathcal{B} \rightarrow \mathcal{B}$ given by $P \mapsto f^{-1}(P)$ is well-defined. It is a Boolean algebra homomorphism and hence determines, by Stone duality, the continuous function $\bar{f} : \mathbf{Spec}(\mathcal{B}) \rightarrow \mathbf{Spec}(\mathcal{B})$ given by $F \mapsto h^{-1}(F) = \{P \in \mathcal{B} : f^{-1}(P) \in F\}$.¹¹

¹⁰Cf. there are only countably many Turing machines.

¹¹We’ll see this construction again in section 5.3.2 of chapter 5.

For all we ever can observe, the ‘logical’ system $(\text{Spec}(\mathcal{B}), \bar{f})$ simply *is* the system (X, f) that we’re observing: any difference between them can, in a sense, never be observed by us. Thus, we may assume without loss of generality that (X, f) is a zero-dimensional topological system: X is a zero-dimensional compact metrizable space and $f : X \rightarrow X$ is continuous. And \mathcal{B} essentially consists of clopen subsets of X . (We also could have taken these systems as our starting point, since the study of zero-dimensional topological systems is an important subfield of dynamical systems theory (Downarowicz and Karpel 2016).)

Trajectories The possible sequences of observations are those trajectories in the observed system A that result from observing the orbit of some state x of the underlying system (X, f) :

$$T := \left\{ t : t \text{ is an } A\text{-trajectory followed by some } x \in X \right\},$$

where we say x follows t iff, for $k = 0, \dots, |t|$, $f^k(x) \in s(t)(k)$, i.e., if the system starts in x , then, after k time steps, we can make the observation $s(t)(k)$ (the k -th state of t).

We can expect T to have two crucial properties:

- (a) If $t \preceq t' \in T$ and t is nonempty, then $t \in T$.

In words: T is closed under nonempty prefixes.

- (b) For all infinite A -trajectories t , if, for all n , $t \upharpoonright n \in T$, then $t \in T$.

In words: T is ‘finitary’ or ‘compact’: if t is not in A , we can realize this after some finite amount of time.

Indeed, concerning (a), if x follows t' , then it also follows the initial segment t . Concerning (b), in the observed system of the North-South map, we’re in the fortunate case that all trajectories are followed by some x , so this property holds vacuously. In the general case of the preceding paragraph, X is compact and decisively observable properties are clopen. Define $A_n := \bigcap_{k=0}^n f^{-k}s(t)(k)$, i.e., the set of those $x \in X$ that follow $t \upharpoonright n$. Then $(A_n)_n$ is a decreasing sequence of non-empty closed subsets of X , so, by compactness, there is $x \in \bigcap_n A_n$, whence $t \in T$.

Trajectory equivalence Each trajectory represents a possible behavior of the system, but often we want to move to a higher *level of abstraction* (or *explanation*) where we consider some distinct trajectories to be instances of the same behavior. For instance, in the neural network example we might want to investigate whether a certain initial value range for some weights is predictive of a certain classification behavior at the end of learning (a microscopic explanation of a macroscopic property). So we consider two observation trajectories t and t' equivalent if,

intuitively, initially their values of the weights in question lie in the same range and the network ends up with the same classification behavior on the test data set. Thus, trajectory equivalence represents a level of abstraction where we ignore information that we don't deem relevant for the intended explanation of the macroscopic properties.

There are many trajectory equivalence relations that can be defined, and this chapter is about axiomatizing those that provide a 'good' level of abstraction. To provide some concrete examples, for the North-South map, we can consider two trajectories equivalent if they have the same length and visit the same sets of states (which here implies having the same start and end). Intuitively, equivalence then represents predicability of observations within a certain number of time steps.

In the general case, we may define, for $t, t' \in T$, that $t \equiv t'$ iff $|t| = |t'|$ and if $|t| > 0$, there is $1 \leq i \leq |t|$ such that

- (i) Same start: $s(t)(0) = s(t')(0)$.
- (ii) Consistent observations: $\bigcap_{k=0}^i f^{-k}(s(t)(k)) = \bigcap_{k=0}^i f^{-k}(s(t')(k))$. (So we might make different observations along $t \upharpoonright i$ and $t' \upharpoonright i$, but we cannot deduce a difference in microscopic states.)
- (iii) Same end: $\text{last}(t \upharpoonright i) = \text{last}(t' \upharpoonright i)$ and, for all $n \geq 0$, $t(i+n) = t'(i+n)$ whenever defined.¹²

Both in the North-South map and in the general case, we can expect \equiv to have two crucial properties:

- (c) For all $t, t' \in T$, if $t \equiv t'$, then $|t| = |t'|$ and there is $i \geq 1$ such that, for all $n \geq 0$, $t \upharpoonright i+n \equiv t' \upharpoonright i+n$.¹³

In words: Equivalent trajectories have the same length and, after some finite time, become (and stay) equivalent.

- (d) For all nonempty finite $t, t' \in T$ with $t \equiv t'$, if $tt'' \in T$ is finite, then $t't'' \in T$ and $tt'' \equiv t't''$.¹⁴ (If t , t' , or t'' are empty, this holds trivially.)

¹²This is indeed an equivalence relation: Reflexivity and symmetry are clear. For transitivity, assume $t \equiv t'$ and $t' \equiv t''$. So $|t| = |t'| = |t''|$, and if > 0 , then there are $1 \leq i, j \leq |t|$ such that the conditions (i)–(iii) are satisfied for (t, t') and (t', t'') , respectively. Let $l := \max(i, j)$ and show that it satisfies (i)–(iii) for (t, t'') .

¹³Proof: North-South: If $t \equiv t'$, then, by definition, $|t| = |t'|$. If they are finite, choose $i := |t| + 1$. If they are infinite, they either both have a tail of N -states or of S -states, and we choose i large enough such that $t(i) = t'(i)$. General: If $t \equiv t'$, then, by definition, $|t| = |t'|$. If t is empty, choose $i = 1$, and if t is nonempty, choose i as in the definition of $t \equiv t'$.

¹⁴Proof: North-South: The crucial part is to show $t't'' \in T$. Since $t \equiv t'$, they have the same last state M . Let M' be the first state of t'' (if t'' is empty the claim is trivial). Since $tt'' \in T$, $M \rightarrow M'$ is an A -transition. So the paths t' and t'' can be concatenated, i.e., $t't''$ is an A -trajectory. So $t't'' \in T$, since any A -trajectory is followed by some state. General: Since $|t| = |t'| > 0$, let $1 \leq i \leq |t|$ be as in the definition of $t \equiv t'$. Write $t_0 := t''$ and $t_1 := t't''$.

In words: Extending equivalent trajectories in the same way yields equivalent trajectories as soon as one extension is in T .

Possible behavior So we're looking at a structure (A, T, \equiv) where A is a countable transition system, T is a set of A -trajectories, and \equiv is an equivalence relation on T such that (a)–(d) are satisfied. An equivalence class $[t]$ describes a possible behavior of A *at the level of abstraction* represented by \equiv .

To understand these behaviors, we're lead to study the structure of the set of possible behaviors $\mathbb{T} := T / \equiv$. It is useful to start with the subset $\mathbb{T}_{\text{fin}} := \{t \in T : t \text{ finite}\} / \equiv$. On there we have a natural order of *information containment* between behaviors: For $[t], [t'] \in \mathbb{T}_{\text{fin}}$, define

$$[t] \leq [t'] :\Leftrightarrow \forall t_0 \in [t] \exists t_1 \in [t'] : t_0 \preceq t_1.$$

A natural way to extend this to infinite behaviors $[t], [t'] \in \mathbb{T}$ is: $[t] \sqsubseteq [t']$ iff, for all n there is m such that $[t \upharpoonright n] \leq [t' \upharpoonright m]$. (The main result of this chapter will show that this essentially also is the only natural way.) This definition makes sense:

2.3.1. LEMMA. 1. $(\mathbb{T}_{\text{fin}}, \leq)$ is a preorder.

2. For finite $t, t' \in T$, if $t \preceq t'$, then $[t] \leq [t']$.

3. The definition of \sqsubseteq is independent of the representative.

4. \sqsubseteq and \leq coincide on \mathbb{T}_{fin} .

PROOF. Concerning (1), this is immediate. Concerning (2), let $t_0 \in [t]$ and write $tt'' = t' \in T$. So (d) implies $t_0 \preceq t_0t'' =: t_1 \in T$ and $t_1 \equiv t'$.

Concerning (3), we show: If $t_0 \in [t]$ and $t_1 \in [t']$ and $[t] \sqsubseteq [t']$, then $[t_0] \sqsubseteq [t_1]$. So given n , find m such that $[t_0 \upharpoonright n] \leq [t_1 \upharpoonright m]$. Since $t_0 \equiv t$, use (c) and let $i \geq 1$ be such that, for all $k \geq 0$, $t_0 \upharpoonright i+k \equiv t \upharpoonright i+k$. Choose some $k \geq i, n$. So, by (2), $[t_0 \upharpoonright n] \leq [t_0 \upharpoonright k] = [t \upharpoonright k]$. Since $[t] \sqsubseteq [t']$, there is j such that $[t \upharpoonright k] \leq [t' \upharpoonright j]$. Since $t' \equiv t_1$, use (c) as above and get $m \geq j$ such that $t' \upharpoonright m \equiv t_1 \upharpoonright m$. So, by (2), $[t' \upharpoonright j] \leq [t' \upharpoonright m] = [t_1 \upharpoonright m]$, as needed.

Concerning (4), let $t, t' \in T$ be finite. If $[t] \sqsubseteq [t']$, let $n := |t|$, so there is m with $[t] = [t \upharpoonright n] \leq [t' \upharpoonright m] \leq [t']$, where the last step follows by (2). Conversely, if $[t] \leq [t']$, then, for any n , let $m := |t'|$ and we have, by (2), that $[t \upharpoonright n] \leq [t] \leq [t'] = [t' \upharpoonright m]$. \square

We have $t_1 \in T$ because $\bigcap_{k=0}^{|t_1|} f^{-k}t_1(k) = \bigcap_{k=0}^{|t_0|} f^{-k}t_0(k)$ (since until i , the intersections are identical by assumption, and after i the trajectories are identical), and the latter is nonempty. And i also witnesses $tt'' \equiv t't''$.

Trajectory domain Now, the key insight into the structure of the set of possible behaviors \mathbb{T} is stated in the following theorem. Its terminology was reviewed in section 2.2.2.

2.3.2. THEOREM. *The partial order $(\overline{\mathbb{T}}, \overline{\sqsubseteq})$ induced by $(\mathbb{T}, \sqsubseteq)$ is isomorphic to the ideal completion of $(\mathbb{T}_{\text{fin}}, \leq)$, and hence an ω -algebraic dcpo.*

In section 2.3.2, we discuss in detail the origins of the proof and the surrounding ideas in the different setting of concurrent computation. The short answer will be: although different in setting and detail, the essential idea of the proof is provided by Droste (1990, thm. 2.3) and Stark (1990, thm. 3). Since the theorem is a consequence of our main result (theorem 2.6.3 below), we only provide a proof sketch.

PROOF SKETCH. We show that the mapping $[[t]_{\equiv}]_{\sqsubseteq} \mapsto I(t) := \{[t'] \in \mathbb{T}_{\text{fin}} : \exists m . [t'] \leq [t \upharpoonright m]\}$ is an order-isomorphism. It is readily seen to be well-defined (i.e., independent of the representative and $I(t)$ is an ideal in \mathbb{T}_{fin}) and an order-embedding (i.e., $[t]_{\sqsubseteq} \leq [t']_{\sqsubseteq}$ iff $I(t) \subseteq I(t')$).

So the key is surjectivity. Since the system is assumed to be countable, \mathbb{T}_{fin} is countable, too. So if D is an ideal of \mathbb{T}_{fin} , it is a countable directed set and hence has a cofinal chain $C = [t_0] \leq [t_1] \leq \dots$. By definition of \leq , we can pick the t_i such that each t_i is an extension of the previous t_j 's. Let t be the trajectory having all t_i as initial segments. By (b), it is in T . Then $I(t) = D$: If $[t'] \in I(t)$, then $[t'] \leq [t \upharpoonright m]$ for some m , so, since D is a downset, $[t'] \in D$. If $[t'] \in D$, then, since the chain is cofinal, there is m such that $[t \upharpoonright m] \geq [t']$, whence $[t'] \in I(t)$. \square

As an example, let's consider the trajectory domain of the North-South map as shown in figure 2.2. We focus on trajectories starting with N (plus the empty trajectory) and abbreviate trajectories thus: $N \rightarrow N \rightarrow E \rightarrow S$ becomes $N^2E^1S^1$. The ellipses 'hide' a more complicated order involving equivalence classes of the form written inside the ellipse.

What does the trajectory domain tell us about the system's behavior? Here are three examples: First, the fact that it has very few noncompact elements relates to the system being very 'convergent' or 'non-chaotic'. Second, it highlights consistent and inconsistent behavior: On the one hand, $[[N^1W^1]]$ and $[[N^1E^1]]$, for instance, are inconsistent (i.e., aren't both informationally contained in some behavior) which reflects that there are no transitions between W and E . On the other hand, $[[N^5]]$ and $[[N^4W^1]]$, for instance, are informationally incomparable, but they both are contained in $[[N^5W^1S^1]] = [[N^4W^1S^2]]$. Third, the fact that the North Pole is an unstable fixed point is reflected in the fact that its infinite fixed point behavior $[[N^\omega]]$ is dominated by the infinite non-fixed point behaviors $[[N^nW^1S^\omega]]$ and $[[N^nE^1S^\omega]]$: because any initial segment N^n can also be realized by a sufficiently close state $x \neq p$ which, however, will eventually evolve into an initial segment $N^nW^1S^1$ or $N^nE^1S^1$ of the infinite non-fixed point behavior.




Figure 2.2: A sketch of the trajectory domain of the North-South map.

2.3.2 Concurrent computation

Curiously, in the study of concurrent computation, we can also find the structure (A, T, \equiv) of an LTS A together with a set of A -trajectories T and an equivalence relation \equiv on T satisfying properties (a)–(d).

Concurrency Concurrency is a vast field of computer science, and it is usually sketched along the following lines (Lampert 2015; Winskel and Nielsen 1995). In sequential computation—as performed, e.g., by Turing machines—the computing system performs one task after the other as dictated by its program. In concurrent computation—e.g., electrical circuits, the internet, or (artificial) neural networks—, many computing units form a network mutually influencing each other and usually performing a joint task.

As a result, several execution paths of this system of computing units may be seen as performing the same task (or computation). To illustrate this, consider the well-known situation of figure 2.3 (Winskel and Nielsen 1995). Assume the system is in the ‘global’ state s which describes the state of each computing unit. Then it could perform either the action α of, say, updating unit 1 or the action β of updating unit 2. This respectively yields the two new states s_0 and s_1 . In each of these states the respective other unit can be updated, and this happens to be such that either order of updating yields the same global state s' . Thus, we’d consider the two *distinct* execution paths $t = s \xrightarrow{\alpha} s_0 \xrightarrow{\beta} s'$ and $t' = s \xrightarrow{\beta} s_1 \xrightarrow{\alpha} s'$ to be behaviorally *equivalent*.




Figure 2.3: The independence diamond.

Models of concurrency There is a plethora of formal models to describe and reason about the behavior of concurrent systems (Baier and Katoen 2008; Sassone, Nielsen, and Winskel 1996; Winskel and Nielsen 1995). We'll mention some (roughly in increasing generality) that are particularly suited to describe the above intuition of equivalence of execution paths.

First, Mazurkiewicz trace languages.¹⁵ These are structures of the form (M, L, I) where L is a set of *actions*, $I \subseteq L \times L$ is a symmetric and irreflexive relation, called the *independence* relation, and $M \subseteq L^{<\omega}$ is a nonempty set of strings over L that is prefix closed (for all $t \in L^{<\omega}$ and $\alpha \in L$, if $t\alpha \in M$, then $t \in M$) and I -closed (for all $t, t' \in L^{<\omega}$ and $\alpha, \beta \in L$, if $t\alpha\beta t' \in M$ and $\alpha I \beta$, then $t\beta\alpha t' \in M$). Thus, we think of M as the set of possible finite sequences of actions (from the set L) that the system can perform, and the independence relation I describes which actions can occur concurrently. In the independence diamond, we'd have $\alpha I \beta$, so the sequences of actions $\alpha\beta$ and $\beta\alpha$ would be considered equivalent. More generally, one defines an equivalence relation \simeq on M as the smallest equivalence relation such that $t\alpha\beta t' \simeq t\beta\alpha t'$ whenever $\alpha I \beta$. Its equivalence classes are called *traces*. A natural preorder on M is $t \leq t'$ iff $\exists t'' : tt'' \simeq t'$ which becomes a preorder on traces when quotienting under \simeq . (For extensions, see, e.g., Katz and Peled 1992 loosening the constraint on the independence relation to be fixed for all actions.)

Second, asynchronous transition systems.¹⁶ Their idea is to specify the transition system that gives rise to the possible strings of labels (of a Mazurkiewicz trace language). So they are structures (A, I) where A is an LTS and $I \subseteq L \times L$ is an irreflexive and symmetric relation, called the independence relation, such that the following axioms are satisfied: (i) Every label occurs in a transition, (ii) performing the action described by a label yields a unique state, and (iii) the independence diamond is respected, i.e., the lower half of the diamond can be completed to the upper half, and the left half of the diamond can be completed to the right half. (See the references for a formal statement.)

¹⁵They were introduced by Mazurkiewicz (in 1977); for references and an overview, see Winskel and Nielsen (1995, sec. 7).

¹⁶They were introduced independently by Bednarczyk (in 1988) and Shields (in 1985); for references and an overview see Winskel and Nielsen (1995, sec. 10).

Third, automata with concurrency.¹⁷ Their main idea is to generalize the independence relation to be relative to the state of the transition system and not fixed for all labels (cf. the extension of trace languages above). Concretely, an automaton with concurrency relations is a structure $(A, (I_s)_{s \in S_A})$ where A is a countable LTS and each $I_s \subseteq L \times L$ is a irreflexive and symmetric relation such that the axioms (i) and (ii) are satisfied and if $\alpha I_s \beta$, then the independence diamond of figure 2.3 can be formed (for details see the references).

Fourth, labeled transition systems with independence.¹⁸ They have the same structure (A, I) as asynchronous transition systems, but they are governed by different axioms (see the references for details). The main difference is that they allow defining two transitions to be occurrences of the same event if, roughly, they participate in an appropriate independence diamond. (Other than being an equivalence relation, the exact definition will not be important for us.)

(We haven't mentioned two other important models of concurrency: Petri nets and event structures. Winskel and Nielsen (1995) discuss their close connections to the models mentioned above.)

Generalization We show that these models essentially have the structure (A, T, \equiv) of an LTS A together with a set of A -trajectories T and an equivalence relation \equiv on T :

First, given a Mazurkiewicz trace language (M, L, I) we can think of it as consisting of the one-state LTS $A = (\{i\}, i, L, \rightarrow)$, where $\rightarrow := \{i\} \times L \times \{i\}$ is the trivial relation, together with the set of A -trajectories T with label sequences that are in M :

$$T := \{t \text{ } A\text{-trajectory} : \forall n . l(t \upharpoonright n) \in M\}.$$

The equivalence relation \equiv on T is the natural extension of \simeq : $t \equiv t'$ iff $\exists i \forall n : l(t \upharpoonright i + n) \simeq l(t' \upharpoonright i + n)$.¹⁹

Moreover, this satisfies properties (a)–(d): Properties (a) and (b) are satisfied by construction. Concerning (c), if $t \equiv t'$, then $\exists i \forall n : l(t \upharpoonright i + n) \simeq l(t' \upharpoonright i + n)$, so $|t| = |t'|$ since \simeq implies having the same length and, for all $n \geq 0$, $t \upharpoonright i + n \equiv t' \upharpoonright i + n$. Concerning (d), this is a basic feature of \simeq in trace languages (see e.g. Winskel and Nielsen 1995, prop. 7.1.3).²⁰

¹⁷See Bracho and Droste (1994) and Droste (1990).

¹⁸See Sassone, Nielsen, and Winskel (1996) and Winskel and Nielsen (1995). Related models are: the concurrent transition systems of Stark (1990), the geometric approaches of Fajstrup, Raußen, and Goubault (2006), Goubault and Jensen (1992), and Pratt (1991), or the transition systems with independence and multi-arcs of Hildebrandt and Sassone (1997).

¹⁹This is an equivalence relation: It clearly is reflexive and symmetric since \simeq is. For transitivity, if i and j witness the equivalence of (t, t') and (t', t'') , then $k := \max(i, j)$ witnesses the equivalence of (t, t'') : for $n \geq 0$ we have $l(t \upharpoonright k + n) = l(t \upharpoonright i + (k - i + n)) \simeq l(t' \upharpoonright i + (k - i + n)) = l(t' \upharpoonright j + (k - j + n)) \simeq l(t'' \upharpoonright j + (k - j + n)) = l(t'' \upharpoonright k + n)$.

²⁰Let $t, t' \in T$ be nonempty finite with $t \equiv t'$ and $tt'' \in T$ finite. Hence, $l(t) \simeq l(t')$ and

Second, let's consider asynchronous transition systems and automata with concurrency. Because of their similarity, we'll only discuss the latter. With slightly different terminology, we follow Bracho and Droste (1994) and Droste (1990); for similar constructions see Stark (1989), Stark (1990), and Katz and Peled (1992). Let (A, I_s) be an automaton with concurrency relations. Let T be the set of all A -trajectories starting in the initial state. Trajectory of equivalence of finite $t, t' \in T$ is given by: the reflexive and transitive closure \sim of $t \sim_0 t'$ iff t and t' only differ by an independence diamond, i.e., they are of the form

$$\begin{array}{cccccccc} t & = & t(0) & \dots & t(i-1) & (s, \alpha, s_1) & (s_1, \beta, s') & t(i+2) & \dots & t(n) \\ t' & = & t'(0) & \dots & t'(i-1) & (s, \beta, s_2) & (s_2, \alpha, s') & t'(i+2) & \dots & t'(n) \end{array}$$

for some $\alpha I_s \beta$. We preorder T by $t \leq t'$ iff $\forall n \exists m : \exists t_1 \in T : t \upharpoonright n \leq t_1 \sim t' \upharpoonright m$. Bracho and Droste (1994) now take two trajectories t and t' to be equivalent if $t \leq t'$ and $t' \leq t$. However, we define the finer relation $t \equiv t'$ iff $\exists i \forall n : t \upharpoonright i+n \sim t' \upharpoonright i+n$.

Notice that \sim for finite trajectories essentially is \simeq for trace languages and that \leq is the natural extension of the preorder of traces languages to infinite trajectories. A difference to trace languages is that, apart from the restriction to start with the initial state, the set of possible trajectories is not constrained any further.

This, too, satisfies (a)–(d): Properties (a)–(b) are satisfied by construction. Concerning (c), if $t \equiv t'$, then $\exists i \forall n : t \upharpoonright i+n \sim t' \upharpoonright i+n$ so $|t| = |t'|$ (since \sim implies having the same length) and, for all n , $t \upharpoonright i+n \equiv t' \upharpoonright i+n$. Concerning (d), if $t, t' \in T$ are nonempty finite with $t \equiv t'$ and $tt'' \in T$ finite, then $t't''$ is an A -trajectory (since $t \sim t'$ which implies $\text{last}(t) = \text{last}(t')$) that starts in i (since t' does), so $t't'' \in T$, and $tt'' \sim t't''$ whence $tt'' \equiv t't''$.

Third, for transition systems with independence we may also consider another notion of trajectory equivalence: $t \equiv t'$ iff $|t| = |t'|$ and, for $n < |t|$, $t(n)$ and $t'(n)$ are occurrences of the same event. We'll show that this implies a weaker version of property (d) that we'll introduce as part of our more general axiomatization (see example 2.7.6).

Connections to domain theory There are various connections between models of concurrency and domain theory.

Based on the partial order of traces, one can construct an event structure (Winskel and Nielsen 1995, sec. 8.3), and event structures, in turn, can represent various classes of domains: Winskel and Nielsen (1995, p. 125) provide a brief summary.

More directly, assume (A, I_s) is an automaton with concurrency relations. Then, for the set of trajectories T with the preorder \leq defined as above, the induced partial order $(\overline{T}, \overline{\leq})$ is a domain whose compact elements are given by

$l(tt'') = l(t)l(t'') \in M$. By said proposition, $l(t't'') = l(t')l(t'') \in M$ and $l(tt'') \simeq l(t't'')$. Hence, $t't'' \in T$ (since M is prefixed closed) and $tt'' \equiv t't''$ (choose $i := |tt''|$).

equivalence classes of finite trajectories in T —see Droste (1990), Bracho and Droste (1994), and Stark (1990).

Let’s compare this to the connection to domain theory from the black box system example. There we’ve motivated the move from trajectories to their equivalence classes as moving from concrete instances of the system’s behavior to the behavior at the level of abstraction of interest. Here, in the case of concurrent computation, the motivation is well put by Stark:

“concurrency is reflected in the [domain] through the existence of nontrivial upper bounds. Since our goal is to make concurrency explicit, one might argue that concurrent computations [i.e., equivalence classes], rather than computation sequences [i.e., trajectories], ought to be the main focus of attention” (Stark 1990, p. 54).

Moreover, we’ll distinguish the concepts of *trajectory equivalence* \equiv and *information containment* \sqsubseteq , i.e., we don’t *define* equivalence as mutual information containment. We’ll find this separation conceptually useful in searching for the right axiomatization.

2.3.3 Summary and outlook

To summarize, we’ve discussed two examples: observing black box systems and concurrent computation. Both have the structure (A, T, \equiv) of an LTS A together with a set of A -trajectories T and an equivalence relation \equiv on T such that properties (a)–(d) were satisfied. This allowed two things: (i) there is a natural information containment order on the set of possible behaviors T/\equiv , and (ii) this in fact forms an ω -algebraic domain.²¹

Given these examples, it is natural to ask how general they are: In the black box system example, we ask which abstractions are good ones. Which equivalence relations on the set of possible trajectories provide a ‘well-structured’ representation of the types of behaviors of interest? In the concurrency example, we ask which other notions of concurrency are plausible. Which equivalence relations on the set of possible trajectories provide the ‘hallmark’ structure identified in the literature: that the equivalence classes—i.e., concurrent computations—form a domain? For example, can we circumvent the restriction that concurrent computations *cannot* differ in their computation time?

Thus, we ask: What are the minimal demands on the structure (A, T, \equiv) such that (i) we can define a natural information containment, and what is additionally

²¹In hindsight, it may not be too surprising that, despite their distinct appearance, both examples had this structure: we could think of equivalent trajectories in the observed system as being ‘concurrent’ computations done by the black box system (under our interpretation). In the other direction, it is an interesting question—somewhat similar to hidden-variable theory discussions in physics—in which cases it is possible to think of concurrent execution paths as equivalent observation trajectories of an underlying low-level deterministic black box system.

needed that (ii) this forms a domain? In other words, we don't ask the *analytic* question of building further equivalence relations, but rather the *synthetic* question of what the right axioms are for these structures.

As already sketched in the introduction, the answer will be this: First, we define a 'pre-BTS' as the structure (A, T, \equiv) with some very minimal axioms. Then we investigate what it takes to satisfy the (i)-demand and find one more axiom which, when added, yields a 'BTS'. The characterization theorem then says that satisfying the (i)-demand is essentially equivalent to the much stronger (ii)-demand.

2.4 Pre-behavioral transition systems

We introduce pre-behavioral transition systems (pre-BTS) as structures (A, T, \equiv) where A is an LTS, T a set of A -trajectories, and \equiv an equivalence relation on T satisfying some minimal axioms. They are 'pre' in the sense that we'll later explain why we should add one more axiom, which then yields BTSs. After the formal definition, we discuss the axioms; in particular, how they relate to the guiding examples. And we construct basic examples of pre-BTSs.

2.4.1 Definition

Before we define pre-behavioral transition systems, we need the definition of information containment between finite behaviors whose importance we've already encountered in the examples.

2.4.1. DEFINITION. Let (A, T, \equiv) be a structure where A is an LTS, T a set of A -trajectories, and \equiv an equivalence relation on T . Write $[t]$ for the \equiv -equivalence classes. For finite $t, t' \in T$, define

$$[t] \leq [t'] :\Leftrightarrow \forall t_0 \in [t] \exists t_1 \in [t'] : t_0 \preceq t_1.$$

In words, every realization of behavior $[t]$ can be extended to a realization of behavior $[t']$.

Now we can state the definition of a pre-behavioral transition system.

2.4.2. DEFINITION. A *pre-behavioral transition system* (pre-BTS) is a triple $M = (A, T, \equiv)$ where A is an LTS, T is a set of A -trajectories, and \equiv is an equivalence relation on T such that:

1. For all $t \in T$, if t' is a nonempty finite initial segment of t , then $t' \in T$.
2. For all infinite A -trajectories t , if $0 < n_0 < n_1 < \dots$ with $t \upharpoonright n_i \in T$ and $[t \upharpoonright n_i] \leq [t \upharpoonright n_{i+1}]$ (for all $i \geq 0$), then $t \in T$.

3. For all $t, t' \in T$ with $t \equiv t'$, if t is empty, then t' is empty, and if t is finite, then t' is finite.
4. For all infinite $t, t' \in T$, if $t \equiv t'$, then there is $i, j \geq 1$ such that, for all $n \geq 0$, $t \upharpoonright i + n \equiv t' \upharpoonright j + n$.

If M is a pre-BTS, we write $M = (A_M, T_M, \equiv_M)$ and $A_M = (S_M, i_M, L_M, \rightarrow_M)$ and call \equiv_M the trajectory equivalence of M . We call M *countable* if A_M is countable. It will be useful to have a name for the following stronger version of axiom (2):

- (2)* For all infinite A -trajectories t , if $t \notin T$, then there is $n \geq 1$ such that $t \upharpoonright n \notin T$.

To be more precise, we should probably call such M a pre-behavioral *labeled* transition systems, but the current name already is enough of a mouthful, so we omit the term ‘labeled’.

2.4.2 Comments

Finite information containment First, technically, definition 2.4.1 of information containment \leq between finite behaviors would also work for infinite trajectories. However, then any infinite $[t]$ is maximal: If $[t] \leq [t']$, then $t \in [t]$ can be extended to $t_1 \in [t']$, but, since t is infinite, $t = t_1$, whence $[t] = [t']$. This discards too much structure of T/\equiv . (In example 2.7.7 below, we discuss this in more detail.) In section 2.5, we discuss how to extend this definition appropriately to infinite behaviors.

Second, the formal idea behind the definition of \leq is to ‘lift’ the extension preorder \preceq on T to the preorder \leq on equivalence classes (i.e., certain subsets of) T . Such constructions are well-known: For example, to characterize the Hoare powerdomain of a continuous domain D with basis (B, \ll) , one ‘lifts’ the relation \ll from B to finite subsets of B by defining $X \ll_H Y$ iff $\forall x \in X \exists y \in Y : x \ll y$. (The Hoare powerdomain of D then is isomorphic to the ideal completion of the set of finite subsets of B ordered by \ll_H .) For the Plotkin powerdomain, one also demands the ‘dual’: $\forall y \in Y \exists x \in X : x \ll y$.²² This is then an instance of the Egli–Milner relation lifting.²³ However, in our setting, adding this additional ‘dual’ demand seems to be too strong in general: for $[t]$ to be informationally contained in $[t']$ it need not be the case that every realization of $[t']$ is an extension of a realization of $[t]$.

²²See e.g. Abramsky and Jung (1994) for these results about powerdomains.

²³This is yet a special case of the notion of ‘relation lifting’ in coalgebra (Kurz and Velebil 2016, example 2.8).

Comparison to the guiding examples In the guiding examples, we’ve identified the properties (a)–(d) as restricting the structures (A, T, \equiv) , so let’s discuss how the axioms (1)–(4) of a pre-BTS are generalizations thereof.

First, axiom (1) is just verbatim property (a). Note that we can equivalently demand that *any* nonempty initial segment t' of $t \in T$ is in T (because if t' is infinite, then $t' = t \in T$).

Second, property (b) is the stronger version (2)* of axiom (2). The reason for opting for the weaker version as axiom is that (i) it is enough for the desired results, (ii) we want the axioms to be as weak as possible, and (iii) it allows systematically disregarding ‘non-approximable’ behavior as we’ll see in the next chapter. However, if the pre-BTS has the property of being bisimulative (definition 2.7.2), the two versions are equivalent.

Third, axioms (3)–(4) are a weakening of property (c). While property (c) required equivalent trajectories to have the same length, axiom (3) only demands them to have the same ‘cardinality type’: one is empty (resp., finite, infinite) iff the other is empty (resp., finite, infinite). Axiom (4) then is similar to property (c) but now restricted to infinite trajectories and taking into account the lack of a ‘global time’ by allowing distinct offsets i and j (rather than a single i).

Fourth, property (d) has no analogue axiom, and section 2.5 will be about finding a much weaker version of property (d) that can then serve as an axiom to turn a pre-BTS into a BTS.

General motivation for pre-BTSs Independent of the guiding examples, the axioms of a pre-BTS can be motivated generally as follows.

First, the notion of an LTS specifies *locally* which trajectories are possible. However, not every trajectory that is locally possible is *globally* possible.²⁴ Consider, for example, an action α that requires a certain amount of some bounded resource like storage space. Then there is a bound on how often α can be performed which becomes relevant only at a global scale but not at a local one. Yet, it is the globally possible trajectories that we’re interested in when we want to know what the ‘possible behavior’ of the system is.²⁵ So, we need to specify the globally possible trajectories explicitly as a subset T of the set of all locally possible trajectories.

This motivates axioms (1) and (2): Regarding (1), nonempty initial segments of globally possible trajectories should be globally possible as well. (We discuss the empty trajectory below.)

Regarding (2), its stronger version (2)* demands that if an infinite locally possible trajectory t is not globally possible, then already some finite initial

²⁴This general situation of local possibility/consistency and global impossibility/inconsistency is known as contextuality (Abramsky and Brandenburger 2011; Abramsky et al. 2015).

²⁵For example, such additional global constraints on the possible trajectories play a role when considering liveness property of the system (Van Glabbeek and Höfner 2018; Manna and Pnueli 1991).

segment of t fails to be globally possible. Thus, the property of ‘globally possible’ is refutative: if it is false, we can eventually discover that it is false—this is necessary for it to be a constructive or “finitary” concept. In the example, if an infinite trajectory t exceeds the storage space with its α -applications, then this happens already after some finite amount of time. Again, as an axiom, we only demand the weaker version (2).

Second, as already discussed, globally possible trajectories exhibit possible behavior of the system. However, they may be instances of the same type of behavior—as described by an equivalence relation. Two constraints seem fundamental for this notion of trajectory equivalence: First, axiom (3) requires that an infinite trajectory is essentially different from a finite one, and a nonempty trajectory is essentially different from the empty one. Second, axiom (4) requires that if two infinite trajectories are equivalent, then there is a point from which on they are (and remain) equivalent. This again is necessary for the notion of trajectory equivalence to be finitary: we exclude the possibility of two infinite trajectories that are equivalent without us ever being able to observe that (i.e., we can never find two of their finite initial segments that are and remain equivalent).²⁶

On the empty trajectory As already seen in the guiding examples, axiom (1) doesn’t play a major role in the proofs, but it is very plausible and it makes things neater (if we consider an initial segment of a trajectory we don’t need to additionally check that it is in T). However, one might wonder: Why the restriction to *nonempty* initial segment? Why not count the empty trajectory as ‘vacuously’ globally possible (at least as soon as T is nonempty)? The answer is: This is very much *allowed*, but for greater generality we don’t *require* it per axiom. The reason is that if it is in T , then it will always be the least element in the ‘behavior order’, while, in the current phrasing, we could also consider behavior preorders without (or ‘removed’) least element.

2.4.3 Example constructions

We’ve already seen that the structures (A, T, \equiv) of the guiding examples (section 2.3) are pre-BTSs. So let’s consider some general constructions of a pre-BTS starting from an LTS A .

First, here are some natural examples for T .

2.4.3. EXAMPLE. Let A be an LTS. The set T of all (nonempty) A -trajectories (starting in i_A) satisfies axioms (1) and (2)*, whence also (2).

Second, let’s consider possible choices for \equiv . We can always choose \equiv to be the identity relation on T . More interesting examples are obtained by starting with

²⁶This is reminiscent of the idea of learning in the limit (Gold 1967).

an equivalence on finite trajectories and extending them to infinite trajectories guided by axiom (4) as a definition.

2.4.4. PROPOSITION. *Let A be an LTS and T a set of A -trajectories. Let \equiv_0 be an equivalence relation on $\{t \in T : t \text{ nonempty finite}\}$. For $t, t' \in T$, define $t \equiv t'$ iff*

- (a) *both t and t' are empty, or*
- (b) *both t and t' are nonempty finite and $t \equiv_0 t'$, or*
- (c) *both t and t' are infinite and there are $i, j \geq 1$ such that, for all $n \geq 0$, $t \upharpoonright i + n \equiv_0 t' \upharpoonright j + n$.*

Then, if T satisfies axioms (1) and (2) (which are stated without reference to \equiv), then (A, T, \equiv) is a pre-BTS.*

PROOF. We first show that \equiv is an equivalence relation on T . Reflexivity and symmetry are immediate in each of the cases (a)–(c). For transitivity, assume $t \equiv t' \equiv t''$, and show $t \equiv t''$. If one of the trajectories is finite, the others must be finite, too, and $t \equiv t''$ follows since both ‘being empty’ and \equiv_0 are transitive. So assume that all trajectories are infinite. There are $i, j \geq 1$ such that, for all $n \geq 0$, $t \upharpoonright i + n \equiv_0 t' \upharpoonright j + n$, and there are $k, l \geq 1$ such that, for all $m \geq 0$, $t' \upharpoonright k + m \equiv_0 t'' \upharpoonright l + m$. Without loss of generality, $j \leq k$ (the case $j \geq k$ is analogous). Define $n_0 := k - j$. Set $i' := i + n_0$ and $j' := l$. Then we have for $n' \geq 0$ that

$$\begin{aligned} t \upharpoonright i' + n' &= t \upharpoonright i + (n_0 + n') \equiv_0 t' \upharpoonright j + (n_0 + n') = t' \upharpoonright j + ((k - j) + n') \\ &= t' \upharpoonright k + n' \equiv_0 t'' \upharpoonright l + n' = t'' \upharpoonright j' + n', \end{aligned}$$

whence $t \equiv t''$, as needed. Now, axioms (1) and (2) hold by assumption, and axioms (3) and (4) hold by construction. \square

The following are some concrete trajectory equivalences built in this manner.

2.4.5. DEFINITION. Let A be an LTS and T a set of A -trajectories. Consider the following equivalence relations on $\{t \in T : t \text{ nonempty finite}\}$:

- (a) $t \equiv_1 t'$ iff $\text{last}(t) = \text{last}(t')$
- (b) $t \equiv_2 t'$ iff $|t| = |t'|$ and $\text{last}(t) = \text{last}(t')$
- (c) $t \equiv_3 t'$ iff $l(t) = l(t')$ and $\text{last}(t) = \text{last}(t')$.

The equivalence relation induced on T (as defined in proposition 2.4.4) by (a), (b), and (c) is called the *extensional*, *temporal*, and *intensional* equivalence on T , respectively.

2.5 Information containment of behaviors

As mentioned, to understand the behavior of a pre-BTS $M = (A, T, \equiv)$, we want to understand the structure of the set of possible behaviors T / \equiv . To give it some notation:

2.5.1. DEFINITION. Let $M = (A, T, \equiv)$ be a pre-BTS. Define $\mathbb{T} := T / \equiv = \{[t] : t \in T\}$ and $\mathbb{T}_{\text{fin}} := \{[t] : t \in T \text{ finite}\}$. We call the elements of \mathbb{T}_{fin} *finite behaviors* and those of $\mathbb{T} \setminus \mathbb{T}_{\text{fin}}$ *infinite behaviors*.

We've already seen that \mathbb{T}_{fin} has the natural information containment preorder $[t] \leq [t']$. The crucial question is how to sensibly extend this to \mathbb{T} . This is the topic of this section: We first provide three natural definition of such extensions (section 2.5.1) and then we identify a condition which makes them all equivalent (section 2.5.2).

2.5.1 Three definitions of information containment . . .

We'll discuss three natural candidates for a definition of information containment also on infinite behaviors. For all of them, the following notion of approximation is crucial.

2.5.2. DEFINITION. Let $M = (A, T, \equiv)$ be a pre-BTS. For $[t] \in \mathbb{T}$, an *approximation* to $[t]$ is a pair $\tau = (t^\dagger, (n_i)_{i \geq 0})$ with $t^\dagger \in [t]$ and $0 < n_0 < n_1 < \dots$ an infinite sequence such that $[t^\dagger \upharpoonright n_0] \leq [t^\dagger \upharpoonright n_1] \leq \dots$. We call $[t]$ *approximable* iff there is an approximation to $[t]$.

Comments: First, this is an example where axiom (1) is handy: We have $t^\dagger \in [t] \subseteq T$, and if t^\dagger is empty, then $t^\dagger \upharpoonright n_i = t^\dagger \in T$, and if t^\dagger is nonempty, each $t^\dagger \upharpoonright n_i$ is a nonempty initial segment of t^\dagger and hence in T , so we can indeed consider the equivalence classes $[t^\dagger \upharpoonright n_i]$.

Second, also note that whether $[t]$ is approximable doesn't depend on the representative t , so it makes sense to say that $[t]$ (as opposed to t) is approximable.

Third, in general, not every $[t] \in \mathbb{T}$ is approximable, but any $[t] \in \mathbb{T}_{\text{fin}}$ has an approximation $\tau = (t, (|t| + 1 + i)_{i \geq 0})$ for which $[t] = [t \upharpoonright n_0] = [t \upharpoonright n_1] = \dots$. Non-approximable behaviors are—in a sense—completely 'out of reach', and we'll reflect this in our definitions of information containment below by demanding that a non-approximable behavior cannot be informationally contained in an approximable behavior.

The first two definitions for information containment are the following.

2.5.3. DEFINITION. Let $M = (A, T, \equiv)$ be a pre-BTS. For $[t] \in \mathbb{T}_{\text{fin}}$ and $[t'] \in \mathbb{T}$, we define:

1. $[t] \Vdash_{\forall} [t']$ iff for all approximations $(t^\dagger, (n_i)_{i \geq 0})$ to $[t']$, there is an $i \geq 0$ such that $[t] \leq [t^\dagger \upharpoonright n_i]$.
2. $[t] \Vdash_{\exists} [t']$ iff either $[t']$ is not approximable or there is an approximation $(t^\dagger, (n_i)_{i \geq 0})$ to $[t']$ and $i \geq 0$ with $[t] \leq [t^\dagger \upharpoonright n_i]$.

Then, for $[t], [t'] \in \mathbb{T}$, we define:

3. $[t] \sqsubseteq_{\forall} [t']$ iff (a) for all $[t_0] \in \mathbb{T}_{\text{fin}}$, if $[t_0] \Vdash_{\forall} [t]$, then $[t_0] \Vdash_{\forall} [t']$, and (b) if $[t]$ is not approximable, then $[t']$ is not approximable.
4. $[t] \sqsubseteq_{\exists} [t']$ iff for all $[t_0] \in \mathbb{T}_{\text{fin}}$, if $[t_0] \Vdash_{\exists} [t]$, then $[t_0] \Vdash_{\exists} [t']$, and (b) if $[t]$ is not approximable, then $[t']$ is not approximable.

The following two lemmas collect some facts about \sqsubseteq_{\forall} and \sqsubseteq_{\exists} , respectively, that show that they indeed are plausible generalization of \leq . We move their straightforward but somewhat technical proofs to an appendix.

2.5.4. LEMMA. *Let $M = (A, T, \equiv)$ be a pre-BTS. Then*

1. $(\mathbb{T}, \sqsubseteq_{\forall})$ is a preorder.
2. \sqsubseteq_{\forall} and \leq coincide on \mathbb{T}_{fin} .
3. For $[t] \in \mathbb{T}_{\text{fin}}$ and $[t'] \in \mathbb{T}$, we have $[t] \Vdash_{\forall} [t']$ iff $[t] \sqsubseteq_{\forall} [t']$.
4. For $[t], [t'] \in \mathbb{T}$ with $[t]$ approximable, we have $[t] \sqsubseteq_{\forall} [t']$ iff for all $[t_0] \in \mathbb{T}_{\text{fin}}$, if $[t_0] \sqsubseteq_{\forall} [t]$, then $[t_0] \sqsubseteq_{\forall} [t']$.
5. If $[t] \in \mathbb{T}$ doesn't have an approximation, then $[t]$ is infinite and for all $[t'] \in \mathbb{T}$, $[t] \sqsubseteq_{\forall} [t']$.

2.5.5. LEMMA. *Let $M = (A, T, \equiv)$ be a pre-BTS. The statements of lemma 2.5.4 remain true after replacing each subscript \forall by \exists .*

A third definition for information containment is the following.

2.5.6. DEFINITION. Let $M = (A, T, \equiv)$ be a pre-BTS. For $[t], [t'] \in \mathbb{T}$, define $[t] \sqsubseteq_{\text{dom}} [t']$ iff

- (a) for all approximations $\tau = (t^\dagger, (n_i))$ to $[t]$ and $\tau' = (t'^{\dagger}, (m_j))$ to $[t']$, τ' dominates τ , i.e., $\forall i \geq 0 \exists j \geq 0 : [t^\dagger \upharpoonright n_i] \leq [t'^{\dagger} \upharpoonright m_j]$, and
- (b) if $[t]$ is not approximable, then $[t']$ is not approximable.

In general, however, this is not a preorder. This raises the question of how these three candidates for information containment can be united.

2.5.2 ... and how they are united

We face a rather messy situation: We have two natural preorders \sqsubseteq_{\forall} and \sqsubseteq_{\exists} and a natural attempt \sqsubseteq_{dom} which, however, doesn't always work. The following notion (definition 2.5.7) provides the precise condition to bring order to this mess—as the subsequent proposition 2.5.8 shows.

2.5.7. DEFINITION. Let $M = (A, T, \equiv)$ be a pre-BTS. We say M is *limit-respecting* if for all infinite $t \in T$ and for all infinite sequences $0 < n_0 < n_1 < \dots$ and $0 < m_0 < m_1 < \dots$, if $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$ and $[t \upharpoonright m_0] \leq [t \upharpoonright m_1] \leq \dots$, then the latter dominates the former, i.e., for all $i \geq 0$, there is $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$.

This is the minimal definition, but we could also phrase it symmetrically: A pre-BTS $M = (A, T, \equiv)$ is limit-respecting iff, for all infinite $t \in T$ and $0 < n_0 < n_1 < \dots$ and $0 < m_0 < m_1 < \dots$, if $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$ and $[t \upharpoonright m_0] \leq [t \upharpoonright m_1] \leq \dots$, then they *mutually* dominate each other: i.e., $\forall i \exists j : [t \upharpoonright n_i] \leq [t \upharpoonright m_j]$ and $\forall k \exists l : [t \upharpoonright m_k] \leq [t \upharpoonright n_l]$.²⁷

2.5.8. PROPOSITION. Let $M = (A, T, \equiv)$ be a pre-BTS. The following are equivalent:

1. M is limit-respecting.
2. $\sqsubseteq_{\forall} = \sqsubseteq_{\exists}$.
3. For all $[t] \in \mathbb{T}$ and approximations $(t^\dagger, (n_i))$ to $[t]$, we have, for all $i \geq 0$, that $[t^\dagger \upharpoonright n_i] \Vdash_{\forall} [t]$.
4. $\sqsubseteq_{\forall} = \sqsubseteq_{\text{dom}}$.
5. \sqsubseteq_{dom} is a preorder.
6. \sqsubseteq_{dom} is reflexive.

In particular, M is limit-respecting iff all the relations $\sqsubseteq_{\forall}, \sqsubseteq_{\exists}, \sqsubseteq_{\text{dom}}$ coincide and thus provide a single natural preorder on \mathbb{T} .

PROOF. (1) \Rightarrow (2). Assume that M is limit-respecting. To show $\sqsubseteq_{\forall} = \sqsubseteq_{\exists}$, it suffices, by definition, to show $\Vdash_{\forall} = \Vdash_{\exists}$. So let $[t] \in \mathbb{T}_{\text{fin}}$ and $[t'] \in \mathbb{T}$ and show $[t] \Vdash_{\forall} [t']$ iff $[t] \Vdash_{\exists} [t']$. If $[t']$ doesn't have an approximation, both sides are true, so let τ'' be an approximation to $[t']$. If $[t] \Vdash_{\forall} [t']$, then all approximations to $[t']$ dominate $[t]$, so in particular, τ'' dominates $[t]$, whence $[t] \Vdash_{\exists} [t']$. So assume $[t] \Vdash_{\exists} [t']$ and show $[t] \Vdash_{\forall} [t']$.

²⁷Proof: The right-to-left direction is immediate, and for the left-to-right direction apply the definition of being limit-respecting first to $((n_i), (m_j))$ and then to $((m_k), (n_i))$.

If $[t']$ is finite, then, since both \Vdash_{\exists} and \Vdash_{\forall} coincide, by lemmas 2.5.4 and 2.5.5, with \leq on finite trajectories, we have that $[t] \Vdash_{\exists} [t']$ iff $[t] \Vdash_{\forall} [t']$. So let $[t']$ be infinite.

To show $[t] \Vdash_{\forall} [t']$, let $\tau' = (t^{\ddagger}, (m_j))$ be an approximation to $[t']$, and find $j \geq 0$ such that $[t] \leq [t^{\ddagger} \upharpoonright m_j]$. Since $[t] \Vdash_{\exists} [t']$ and $[t']$ is approximable,

(*) there is an approximation $\tau = (t^{\ddagger}, (n_i))$ to $[t']$ and $i \geq 0$ with $[t] \leq [t^{\ddagger} \upharpoonright n_i]$.

Since $t^{\ddagger} \equiv t' \equiv t^{\ddagger}$ are infinite by axiom (3), there are, by axiom (4), $k, l \geq 1$ such that, for all $n \geq 0$, $[t^{\ddagger} \upharpoonright k + n] = [t^{\ddagger} \upharpoonright l + n]$. Let $i_0 \geq 0$ be such that $n_{i_0} > k$. And let $j_0 \geq 0$ be such that $m_{j_0} > l$. For $j \geq j_0$, define $n(j) := m_j - l \geq 0$. Define $(n'_i)_{i \geq 0} := (n_{i+i_0})_{i \geq 0}$ and $(m'_j)_j := (k + n(j + j_0))_{j \geq 0}$. Then $0 < n'_0 < n'_1 < \dots$ and $0 < m'_0 < m'_1 < \dots$ ²⁸ Moreover, $[t^{\ddagger} \upharpoonright n'_i] = [t^{\ddagger} \upharpoonright n_{i+i_0}] \leq [t^{\ddagger} \upharpoonright n_{i+1+i_0}] = [t^{\ddagger} \upharpoonright n'_{i+1}]$. Note that

$$\begin{aligned} [t^{\ddagger} \upharpoonright m'_j] &= [t^{\ddagger} \upharpoonright k + n(j + j_0)] = [t^{\ddagger} \upharpoonright l + n(j + j_0)] \\ &= [t^{\ddagger} \upharpoonright l + (m_{j+j_0} - l)] = [t^{\ddagger} \upharpoonright m_{j+j_0}]. \end{aligned}$$

Hence, $[t^{\ddagger} \upharpoonright m'_j] = [t^{\ddagger} \upharpoonright m_{j+j_0}] \leq [t^{\ddagger} \upharpoonright m_{j+1+j_0}] = [t^{\ddagger} \upharpoonright m'_{j+1}]$.

Now, we apply the property that M is limit-respecting to t^{\ddagger} and i from (*), and we obtain that there is $j_1 \geq 0$ such that $[t^{\ddagger} \upharpoonright n'_i] \leq [t^{\ddagger} \upharpoonright m'_{j_1}]$. Thus, we have

$$[t] \leq [t^{\ddagger} \upharpoonright n_i] \leq [t^{\ddagger} \upharpoonright n_{i+i_0}] = [t^{\ddagger} \upharpoonright n'_i] \leq [t^{\ddagger} \upharpoonright m'_{j_1}] = [t^{\ddagger} \upharpoonright m_{j_1+j_0}].$$

Hence, for $j := j_1 + j_0$, we have $[t] \leq [t^{\ddagger} \upharpoonright m_j]$, as needed.

(2) \Rightarrow (3). If $\sqsubseteq_{\forall} = \sqsubseteq_{\exists}$, then we have, by lemmas 2.5.4 and 2.5.5, that $\Vdash_{\forall} = \Vdash_{\exists}$. So it suffices to show (3) for \Vdash_{\exists} instead of \Vdash_{\forall} . But this is immediate: If $[t] \in \mathbb{T}$ has approximation $(t^{\ddagger}, (n_i))$, we have, for $i \geq 0$, that $[t^{\ddagger} \upharpoonright n_i] \leq [t^{\ddagger} \upharpoonright n_i]$, so $[t^{\ddagger} \upharpoonright n_i] \Vdash_{\exists} [t]$.

(3) \Rightarrow (4). Let $[t], [t'] \in \mathbb{T}$ and show $[t] \sqsubseteq_{\forall} [t']$ iff $[t] \sqsubseteq_{\text{dom}} [t']$.

(\Rightarrow) To show $[t] \sqsubseteq_{\text{dom}} [t']$ we need to show properties (a) and (b). Concerning (a), consider the approximations $\tau = (t^{\ddagger}, (n_i))$ to $[t]$ and $\tau' = (t^{\ddagger}, (m_j))$ to $[t']$. Let $i \geq 0$ and find $j \geq 0$ such that $[t^{\ddagger} \upharpoonright n_i] \leq [t^{\ddagger} \upharpoonright m_j]$. Indeed, by (3), we have $[t^{\ddagger} \upharpoonright n_i] \Vdash_{\forall} [t]$. Since $[t] \sqsubseteq_{\forall} [t']$, we hence have $[t^{\ddagger} \upharpoonright n_i] \Vdash_{\forall} [t']$. So, for the approximation τ' to $[t']$, there is $j \geq 0$ such that $[t^{\ddagger} \upharpoonright n_i] \leq [t^{\ddagger} \upharpoonright m_j]$, as needed. Concerning (b), this is implied by $[t] \sqsubseteq_{\forall} [t']$.

(\Leftarrow) To show $[t] \sqsubseteq_{\forall} [t']$, clause (b) of \sqsubseteq_{\forall} is given by clause (b) of \sqsubseteq_{dom} , so let $[t_0] \in \mathbb{T}_{\text{fin}}$ with $[t_0] \Vdash_{\forall} [t]$, and show $[t_0] \Vdash_{\forall} [t']$. So let $\tau' = (t^{\ddagger}, (m_j))$ be an approximation to $[t']$ and find $j \geq 0$ such that $[t_0] \leq [t^{\ddagger} \upharpoonright m_j]$. Since $[t']$ is approximable, also $[t]$ is by clause (b) of \sqsubseteq_{dom} . So let $\tau = (t^{\ddagger}, (n_i))$ be an approximation to $[t]$. Since $[t_0] \Vdash_{\forall} [t]$, there is $i \geq 0$ such that $[t_0] \leq [t^{\ddagger} \upharpoonright i]$. Since $[t] \sqsubseteq_{\text{dom}} [t']$, there is $j \geq 0$ such that $[t^{\ddagger} \upharpoonright i] \leq [t^{\ddagger} \upharpoonright m_j]$. Together this yields $[t_0] \leq [t^{\ddagger} \upharpoonright m_j]$, as needed.

²⁸For the latter: $m'_0 = k + n(j_0) = k + (m_{j_0} - l) \geq m_{j_0} - l > 0$ and, since $m_i < m_{i+1}$ we have $m'_j = k + n(j + j_0) = k + m_{j+j_0} - l < k + m_{j+1+j_0} - l = k + n(j + 1 + j_0) = m'_{j+1}$.

(4) \Rightarrow (5). This holds since $\sqsubseteq_{\vee} = \sqsubseteq_{\text{dom}}$ is a preorder.

(5) \Rightarrow (6). This is trivial.

(6) \Rightarrow (1). Let $t \in T$ be infinite, and consider the strictly increasing $n_i, m_j > 0$ with $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$ and $[t \upharpoonright m_0] \leq [t \upharpoonright m_1] \leq \dots$. Let $i \geq 0$ and find $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$. Note that $\tau = (t, (n_i))$ and $\tau' = (t, (m_j))$ are approximations to $[t]$. Since \sqsubseteq_{dom} is reflexive, we have, by clause (a), that τ' dominates τ , so there is $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$, as needed. \square

2.6 The characterization theorem

Now we get to the main result of this chapter: the characterization theorem. For countable systems it roughly says that, once united, the natural information containment preorders \sqsubseteq from the previous section not only are the only sensible ones, but they also form the algebraic domain $(\overline{\mathbb{T}}, \sqsubseteq)$. We first state and discuss the theorem (section 2.6.1) and then prove it (section 2.6.2).

2.6.1 Statement

We first introduce two more bits of notation.

2.6.1. DEFINITION. Let $M = (A, T, \equiv)$ be a pre-BTS. For a preorder \preceq on \mathbb{T} , define the induced partial order $\mathbb{T}(M, \preceq) := (\overline{\mathbb{T}}, \preceq)$. We denote the elements $[[t]_{\equiv}]_{\preceq}$ or, if clear from context, simply $[[t]]$. For $[t] \in \mathbb{T}$, we define $I_{\preceq}([t]) := \{[t_0] \in \mathbb{T}_{\text{fin}} : [t_0] \preceq [t]\}$.

2.6.2. DEFINITION. Let $M = (A, T, \equiv)$ be a pre-BTS. We write $\text{ldl}(\mathbb{T}_{\text{fin}}, \leq)$ for the ideal completion of $(\mathbb{T}_{\text{fin}}, \leq)$. We define $\overline{\text{ldl}}(\mathbb{T}_{\text{fin}}, \leq)$ to be $\text{ldl}(\mathbb{T}_{\text{fin}}, \leq)$ if all $[t] \in \mathbb{T}$ are approximable, and we define it to be $\text{ldl}(\mathbb{T}_{\text{fin}}, \leq)$ with an added top element \top if M has non-approximable trajectories. So $\overline{\text{ldl}}(\mathbb{T}_{\text{fin}}, \leq)$ is an algebraic domain (if existent, \top is a compact element). We still use \subseteq to denote the order relation in $\overline{\text{ldl}}(\mathbb{T}_{\text{fin}}, \leq)$.

Now the characterization theorem reads as follows. After stating it, we discuss how it indeed provides the answers we were looking for.

2.6.3. THEOREM. *Let $M = (A, T, \equiv)$ be a countable pre-BTS. Let $\preceq \subseteq \mathbb{T} \times \mathbb{T}$ be a relation. The following are equivalent.*

1. (a) \preceq is a preorder that coincides with \leq on \mathbb{T}_{fin} .
- (b) For $[t] \in \mathbb{T}$ not approximable and $[t'] \in \mathbb{T}$, (i) $[t'] \preceq [t]$, and (ii) if $[t] \preceq [t']$, then $[t']$ is not approximable, as well.

- (c) For infinite $t \in T$ and $0 < n_0 < n_1 < \dots$ such that $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$, we have (i) $[t]$ is an \sqsubseteq -upper bound, i.e., for all $i \geq 0$, $[t \upharpoonright n_i] \sqsubseteq [t]$, and (ii) for all $[t_0] \in \mathbb{T}_{\text{fin}}$, if $[t_0] \sqsubseteq [t]$, then there is $i \geq 0$ such that $[t_0] \sqsubseteq [t \upharpoonright n_i]$.
- (d) For approximable $[t], [t'] \in \mathbb{T}$, if, for all $[t_0] \in \mathbb{T}_{\text{fin}}$, $[t_0] \sqsubseteq [t]$ implies $[t_0] \sqsubseteq [t']$, then $[t] \sqsubseteq [t']$.
2. (a) For all approximable $[t] \in \mathbb{T}$, $I_{\sqsubseteq}([t])$ is an ideal in $(\mathbb{T}_{\text{fin}}, \leq)$.
- (b) For $[t] \in \mathbb{T}$ not approximable and $[t'] \in \mathbb{T}$, (i) $[t'] \sqsubseteq [t]$, and (ii) if $[t] \sqsubseteq [t']$, then $[t']$ is not approximable, as well.
- (c) For all approximable $[t], [t'] \in \mathbb{T}$, $[t] \sqsubseteq [t']$ iff for all $[t_0] \in \mathbb{T}_{\text{fin}}$, if $[t_0] \sqsubseteq [t]$, then $[t_0] \sqsubseteq [t']$.
- (d) For all trajectories $t \in T$ and $0 < n_0 < n_1 < \dots$ such that $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$, we have for all $[t_0] \in \mathbb{T}_{\text{fin}}$, $[t_0] \sqsubseteq [t]$ iff $[t_0] \leq [t \upharpoonright n_i]$ for some i .
3. \sqsubseteq is a preorder that coincides with \leq on \mathbb{T}_{fin} and the mapping

$$\iota : \mathbb{T}(M, \sqsubseteq) \rightarrow \overline{\text{Idl}}(\mathbb{T}_{\text{fin}}, \leq)$$

$$[[t]] \mapsto \begin{cases} I_{\sqsubseteq}([t]) & \text{if } [t] \text{ is approximable} \\ \top & \text{otherwise} \end{cases}$$

is a well-defined function and:

- (a) ι is an isomorphism. In particular, $\mathbb{T}(M, \sqsubseteq)$ is an ω -algebraic domain.
- (b) $K(\mathbb{T}(M, \sqsubseteq)) = \{[[t]] : [t] \in \mathbb{T}_{\text{fin}}\} \cup \{[[t]] \in \mathbb{T} : [t] \text{ not approximable}\}$.
- (c) For all $[t] \in \mathbb{T}$ and $0 < n_0 < n_1 < \dots$, if $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$, then $[[t]]$ is the least upper bound in $\mathbb{T}(M, \sqsubseteq)$ of the directed subset $\{[[t \upharpoonright n_0]], [[t \upharpoonright n_1]], \dots\}$ of $\mathbb{T}(M, \sqsubseteq)$.
4. One of the following holds:
- (a) $\sqsubseteq \in \{\sqsubseteq_{\forall}, \sqsubseteq_{\exists}, \sqsubseteq_{\text{dom}}\}$ and M is limit-respecting.
- (b) $\sqsubseteq = \sqsubseteq_{\forall} = \sqsubseteq_{\exists}$.
- (c) $\sqsubseteq = \sqsubseteq_{\forall} = \sqsubseteq_{\text{dom}}$.
- (d) $\sqsubseteq = \sqsubseteq_{\text{dom}}$ is reflexive.
5. $\sqsubseteq = \sqsubseteq_{\forall} = \sqsubseteq_{\exists} = \sqsubseteq_{\text{dom}}$ and M is limit-respecting.

Here is how the theorem answers the question of which preorders can sensibly provide an ‘information containment’ ordering on the set of behaviors of a system. The theorem considers any possible preorder \sqsubseteq (or in fact just a relation) that

one might have on \mathbb{T} . The first two items of the theorem are two formulations of rather weak demands on \trianglelefteq that we would like to be satisfied if \trianglelefteq is to provide any sensible ‘information containment’ relation. The third item shows that these minimal demands actually are enough to yield some very strong demands that we would expect in the best case of an informational order of behaviors. The fourth and fifth item show that there in fact is only one way of defining this information containment ordering and that this puts an additional demand on the underlying system: it should be limit-respecting. Let’s discuss this in a bit more detail.

Item 1 lists minimal assumptions on what it means that a preorder \trianglelefteq is a ‘sensible extension’ of the information containment preorder \leq of finite behaviors: (a) is the demand that \trianglelefteq actually is an extension of \leq . (b) captures the idea that non-approximable behaviors are ‘completely out of reach’: they form an \trianglelefteq -cluster which is strictly above any other \trianglelefteq -cluster. (c) says that if we consider an infinite trajectory t such that its behavior can be approximated by the behaviors realized by its initial segments, then (i) each of these initial segments $[t \upharpoonright n_i]$ is informationally contained in $[t]$, and (ii) if a finite behavior $[t_0]$ is informationally contained in $[t]$, then this can in principle be observed (i.e., $[t_0]$ is already contained in some step of the approximation). (d) says that, for two approximable behaviors $[t]$ and $[t']$, if $[t]$ is not informationally contained in $[t']$, then this again can in principle be observed (i.e., there is some finite behavior that is informationally contained in $[t]$ but not in $[t']$).

Item 2 similarly lists an equivalent set of minimal requirements for an information containment preorder.

Item 3 states a very strong demand on an information containment \trianglelefteq : that the partial order that it induces actually is (a) an ω -algebraic domain where (b) the compact (i.e., ‘real’ or finitely accessible) elements precisely are the finite behaviors and the non-compact (i.e., the ‘ideal’) elements are the infinite behaviors, and (c) where the limit of a chain of finite behaviors is precisely given by the infinite behavior realized by a trajectory extending each of the finite behaviors. In short, ‘finite behavior of the system’ precisely corresponds to ‘compact element of the domain’ and ‘limit behavior of the system’ precisely corresponds to ‘limit in the domain’. Thus, questions about the system’s behavior (e.g., concerning finite observability, consistency, limit behavior) correspond precisely to domain-theoretic notions (e.g., compactness, being upper bounded, least upper bound). Of course, we could simply *define* the ‘trajectory domain’ of a system to be the ideal completion of \mathbb{T}_{fin} , but this then would precisely lack this correspondence which is what provides meaning to the ideal completion.

Item 4 says that—focusing on the (a)-condition—that the information containment preorder \trianglelefteq actually has to be one of natural ones \sqsubseteq_{\forall} , \sqsubseteq_{\exists} , \sqsubseteq_{dom} and the system has to have the unifying property of being limit-respecting. Similarly for the other conditions (b)–(d).

Item 5 states that all three natural information containment preorders actually collapse to the information containment preorder \trianglelefteq under consideration and the

system is limit-respecting.

2.6.2 Proof

To show the equivalences, we'll show the following: (3) \Leftrightarrow (2) \Rightarrow (4) \Rightarrow (5) \Rightarrow (1) \Rightarrow (2).

(3) \Rightarrow (2) Ad (2)(a). This follows since ι is well-defined.

Ad (2)(b). If $[t] \in \mathbb{T}$ is not approximable and $[t'] \in \mathbb{T}$, then $\iota([[t']]) \subseteq \top = \iota([[t]])$, so, qua isomorphism, $[[t']] \sqsubseteq [[t]]$, so $[t'] \sqsubseteq [t]$. And if $[t] \sqsubseteq [t']$, then $[[t]] \sqsubseteq [[t']]$, so, qua isomorphism, $\top = \iota([[t]]) \subseteq \iota([[t']])$, so, qua top element, $\iota([[t']]) = \top$, so $[t']$ must be non-approximable (if it were approximable, $\iota([[t']]) = I_{\sqsubseteq}([t'])$ would be in $\text{Idl}(\mathbb{T}_{\text{fin}})$ and hence not the top element of $\overline{\text{Idl}}(\mathbb{T}_{\text{fin}})$).

Ad (2)(c). By the order isomorphism condition we have, for all approximable $[t], [t'] \in \mathbb{T}$ that

$$[t] \sqsubseteq [t'] \Leftrightarrow [[t]] \sqsubseteq [[t']] \Leftrightarrow I_{\sqsubseteq}([t]) \subseteq I_{\sqsubseteq}([t']) \Leftrightarrow \forall [t_0] \in \mathbb{T}_{\text{fin}} : [t_0] \sqsubseteq [t] \Rightarrow [t_0] \sqsubseteq [t'].$$

Ad (2)(d). Under the assumptions, $I := \{[t'] \in \mathbb{T}_{\text{fin}} : \exists i : [t'] \leq [t \upharpoonright n_i]\}$ is an ideal: It is nonempty, since $[t \upharpoonright n_0] \in I$. It is a downset and directed by construction. Moreover, by (3)(c), $[[t]]$ is the least upper bound of $A := \{[[t \upharpoonright n_0]], [[t \upharpoonright n_1]], \dots\}$ in $\mathbb{T}(M, \sqsubseteq)$. Since ι is an isomorphism and since $[t]$ is approximable by assumption,

$$I_{\sqsubseteq}([t]) = \iota([[t]]) = \iota(\bigvee A) = \bigvee_{a \in A} \iota(a) = \bigcup_i I_{\sqsubseteq}([t \upharpoonright n_i]) = I,$$

where for the last identity we use that \sqsubseteq agrees with \leq on \mathbb{T}_{fin} . Hence, for all $[t_0] \in \mathbb{T}_{\text{fin}}$, we have $[t_0] \sqsubseteq [t]$ iff $[t_0] \in I_{\sqsubseteq}([t])$ iff $[t_0] \in I$ iff $[t_0] \leq [t \upharpoonright n_i]$ for some i .

(3) \Leftarrow (2) First note that, by (2)(b)–(c), \sqsubseteq is a preorder (reflexive and transitive).²⁹ And by (2)(d), \sqsubseteq coincides with \leq on \mathbb{T}_{fin} because: Given $[t], [t'] \in \mathbb{T}_{\text{fin}}$, choose $n_i := |t'| + 1 + i$, then $0 < n_0 < n_1 < \dots$ and $[t' \upharpoonright n_i] \leq [t' \upharpoonright n_{i+1}]$ (because they all equal $[t']$), whence, for $[t_0] := [t]$, we have $[t] \sqsubseteq [t']$ iff there is i with $[t] \leq [t' \upharpoonright n_i] = [t']$.

Next, we show that the mapping ι is well-defined. By (2)(a), $I_{\sqsubseteq}([t])$ is indeed an ideal, whence in $\overline{\text{Idl}}(\mathbb{T}_{\text{fin}})$. And the mapping is unique: For $[t], [t'] \in \mathbb{T}$, assume $[t] \sqsubseteq [t']$ and $[t'] \sqsubseteq [t]$ and show either both are non-approximable (and thus both

²⁹Reflexive: Given $[t] \in \mathbb{T}$, if $[t]$ is not approximable, then $[t] \sqsubseteq [t]$ by (2)(b), and if $[t]$ is approximable, then $[t] \sqsubseteq [t]$ by (2)(c). Transitive: Given $[t] \sqsubseteq [t'] \sqsubseteq [t'']$, if $[t'']$ is not approximable, then, by (2)(b), $[t] \sqsubseteq [t'']$, as needed, so let $[t'']$ be approximable. Then, since $[t'] \sqsubseteq [t'']$, (2)(b) implies that $[t']$ is approximable, too, and similarly $[t] \sqsubseteq [t']$ implies that $[t]$ is approximable. So we can apply (2)(c) to $[t] \sqsubseteq [t'] \sqsubseteq [t'']$ and get that, for all $[t_0] \in \mathbb{T}_{\text{fin}}$, if $[t_0] \sqsubseteq [t]$, then $[t_0] \sqsubseteq [t'']$, so, by (2)(c), $[t] \sqsubseteq [t'']$.

get mapped to \top) or both are approximable and $I_{\triangleleft}([t]) = I_{\triangleleft}([t'])$. If $[t]$ is not approximable, then, since $[t] \trianglelefteq [t']$, by (2)(b), also $[t']$ is not approximable. So assume $[t]$ is approximable. Then also $[t']$ is approximable (for otherwise $[t'] \trianglelefteq [t]$ implies, by (2)(b), that $[t]$ is not approximable). Then we have $I_{\triangleleft}([t]) = I_{\triangleleft}([t'])$, because if $[t_0] \in I_{\triangleleft}([t])$, then $[t_0] \in \mathbb{T}_{\text{fin}}$ and $[t_0] \trianglelefteq [t] \trianglelefteq [t']$, so, by transitivity, $[t_0] \in I_{\triangleleft}([t'])$, and similarly for the other direction.

Ad (3)(a) We need to show that ι is surjective and an order-isomorphism (this implies injectivity).

Surjective. If $\overline{\text{Idl}}(\mathbb{T}_{\text{fin}})$ has an added top element \top , then \mathbb{T} has a non-approximable element $[t]$, and \top has a preimage, namely $[[t]]$. So we need to show that, given an ideal $I \subseteq \mathbb{T}_{\text{fin}}$, there is an approximable $[t] \in \mathbb{T}$ such that $I = I_{\triangleleft}([t])$ (since then $\iota([[t]]) = I$).

Since the system M is assumed to be countable, there are only countable many finite A -trajectories, so \mathbb{T}_{fin} is countable, too. So I is a directed subset of the countable preorder $(\mathbb{T}_{\text{fin}}, \leq)$, so there is a cofinal sequence $C = ([t_i])_{i \geq 0}$.

If C stagnates with $[t_n]$ (i.e., for all $i \geq n$, $[t_i] = [t_n]$), then, for all $[t'] \in I$, $[t'] \leq [t_n]$ (by cofinality), so $I \subseteq \{[t'] \in \mathbb{T}_{\text{fin}} : [t'] \leq [t_n]\}$. We also have \supseteq since $[t_n] \in I$ and I is a downset. Thus, since \trianglelefteq and \leq agree on \mathbb{T}_{fin} , we have $I = \{[t'] \in \mathbb{T}_{\text{fin}} : [t'] \leq [t_n]\} = \{[t'] \in \mathbb{T}_{\text{fin}} : [t'] \trianglelefteq [t_n]\} = I_{\triangleleft}([t_n])$. So $[t_n]$, which is approximable qua finite trajectory, is the required element of \mathbb{T} .

So assume C doesn't stagnate (i.e., for all $[t_i]$ there is j with $[t_i] < [t_j]$). Without loss of generality, assume that $[t_0] \neq [\epsilon]$ and $[t_i] < [t_{i+1}]$ for all $i \geq 0$ (otherwise we pick a subsequence of C with this property which then still is cofinal in I). We construct $t \in T$ and $0 < n_0 < n_1 < \dots$ as follows. Set $t'_0 := t_0$ and $n_0 := |t_0| > 0$ (if $n_0 = 0$, then $[t_0] = [\epsilon]$). Given $t'_i \in [t_i]$ and $n_i = |t'_i|$, we can, since $[t_i] < [t_{i+1}]$, extend t'_i to some $t'_{i+1} \in [t_{i+1}]$, and since $[t_i] \neq [t_{i+1}]$ we have $n_i = |t'_i| < |t'_{i+1}| =: n_{i+1}$. We define the sequence $t(k) = t'_i(k)$ for some i such that $|t'_i| > k$. This is well-defined since the t'_i get arbitrarily long and extend each other. Moreover, we claim that $t \in T$. First, t is infinite and an A -trajectory since each $t(k)$ is in \rightarrow_A and the ending state of $t(k)$ is the starting state of $t(k+1)$. Second, for $i \geq 0$, we have $t \upharpoonright n_i = t'_i \in T$ and $[t \upharpoonright n_i] = [t_i] \leq [t_{i+1}] = [t \upharpoonright n_{i+1}]$. So axiom (2) implies $t \in T$.

Thus, $[t] \in \mathbb{T}$ which has the approximation $\tau := (t, (n_i))$. So it remains to show $I = I_{\triangleleft}([t])$. Indeed, by 2(d) we have for all $[t_0] \in \mathbb{T}_{\text{fin}}$, $[t_0] \trianglelefteq [t]$ iff $[t_0] \leq [t \upharpoonright n_i]$ for some i . Thus, if $[t_0] \in I$, then, by cofinality, $[t_0] \leq [t_i] = [t \upharpoonright n_i]$ for some i , so $[t_0] \trianglelefteq [t]$, whence $[t_0] \in I_{\triangleleft}([t])$. And if $[t_0] \in I_{\triangleleft}([t])$, then $[t_0] \trianglelefteq [t]$, so, for some $i \geq 0$, $[t_0] \leq [t \upharpoonright n_i] = [t_i] \in I$, so, since I is an ideal, $[t_0] \in I$.

Order-isomorphism. For $[t], [t'] \in \mathbb{T}$, we need to show that $[[t]] \trianglelefteq [[t']]$ iff $\iota([[t]]) \subseteq \iota([[t']])$.

If $[t']$ is not approximable, then, by 2(b), $[t] \trianglelefteq [t']$, so $[[t]] \trianglelefteq [[t']]$ and $\iota([[t]]) \subseteq \top = \iota([[t']])$. So assume that $[t']$ is approximable.

First, assume $[t]$ is not approximable. Then we cannot have $[t] \trianglelefteq [t']$ by 2(b) and we also cannot have $\iota([[t]]) \subseteq \iota([[t']])$ since the former is \top but the latter is in

$\text{Idl}(\mathbb{T}_{\text{fin}}, \leq)$.

So assume that both $[t]$ and $[t']$ are approximable. Then we have

$$\begin{aligned} [[t]] \sqsubseteq [[t']] &\Leftrightarrow [t] \leq [t'] \stackrel{2(c)}{\Leftrightarrow} \forall [t_0] \in \mathbb{T}_{\text{fin}} : [t_0] \leq [t] \Rightarrow [t_0] \leq [t'] \\ &\Leftrightarrow I_{\leq}([t]) \subseteq I_{\leq}([t']) \Leftrightarrow \iota([[t]]) \subseteq \iota([[t']]), \end{aligned}$$

as needed.

Ad (3)(b). The compact elements of $\overline{\text{Idl}}(\mathbb{T}_{\text{fin}})$ are the principal ideals (those of the form $\downarrow[t_0] = \{[t'] \in \mathbb{T}_{\text{fin}} : [t'] \leq [t_0]\}$ for $[t_0] \in \mathbb{T}_{\text{fin}}$) and, if existent, the top element. Qua isomorphism, $K(\mathbb{T}(M, \leq)) = \iota^{-1}(K(\overline{\text{Idl}}(\mathbb{T}_{\text{fin}})))$. So we show

$$A := \iota^{-1}(K(\overline{\text{Idl}}(\mathbb{T}_{\text{fin}}))) = \{[[t]] : [t] \in \mathbb{T}_{\text{fin}}\} \cup \{[[t]] \in \mathbb{T} : [t] \text{ not approx.}\} =: B.$$

(\subseteq) If $[[t]] \in A$, then $\iota([[t]])$ is compact. If $[t]$ is not approximable, then $[[t]] \in B$, so let $[t]$ be approximable. Then $\iota([[t]]) = I_{\leq}([t])$ is an ideal (and not the top element), so, qua compactness, $\iota([[t]]) = \downarrow[t_0]$ for some $[t_0] \in \mathbb{T}_{\text{fin}}$. Since \leq and \leq agree on \mathbb{T}_{fin} and $[t_0]$ is approximable, we further have

$$\iota([[t]]) = \downarrow[t_0] = \{[t'] \in \mathbb{T}_{\text{fin}} : [t'] \leq [t_0]\} = I_{\leq}([t_0]) = \iota([[t_0]]),$$

so, since ι is injective, $[[t]] = [[t_0]] \in B$.

(\supseteq) Let $[[t]] \in B$, i.e., $[t] \in \mathbb{T}_{\text{fin}}$ or $[t]$ is not approximable, and show $\iota([[t]])$ is compact. If $[t]$ is not approximable, then $\iota([[t]]) = \top$ is compact, so $[[t]] \in A$. So $[t] \in \mathbb{T}_{\text{fin}}$. Then, since \leq and \leq agree on \mathbb{T}_{fin} , we have $\iota([[t]]) = I_{\leq}([t]) = \{[t'] \in \mathbb{T}_{\text{fin}} : [t'] \leq [t]\} = \downarrow[t]$ is compact, so $[[t]] \in A$.

Ad (3)(c). Let $[t] \in \mathbb{T}$ and $0 < n_0 < n_1 < \dots$ with $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$. In particular, $[t]$ is approximable. Write $A := \{[[t \upharpoonright n_0]], [[t \upharpoonright n_1]], \dots\}$. Also, each $[t \upharpoonright n_i]$ is, qua finite trajectory, approximable. Then, since ι is an isomorphism and \leq and \leq agree on \mathbb{T}_{fin} ,

$$\begin{aligned} \iota\left(\bigvee_{a \in A} a\right) &= \bigvee_{a \in A} \iota(a) = \bigcup_i I_{\leq}([t \upharpoonright n_i]) = \{[t_0] \in \mathbb{T}_{\text{fin}} : \exists i . [t_0] \leq [t \upharpoonright n_i]\} \\ &\stackrel{2(d)}{=} \{[t_0] \in \mathbb{T}_{\text{fin}} : [t_0] \leq [t]\} = I_{\leq}([t]) = \iota([[t]]). \end{aligned}$$

Since ι is injective, $[[t]] = \bigvee A$, as needed.

(2) \Rightarrow (4) We show that (4)(a) holds by showing that (i) $\leq = \sqsubseteq_{\exists}$ and (ii) M is limit-respecting. We first show that

(*) For all $[t] \in \mathbb{T}_{\text{fin}}$ and approximable $[t'] \in \mathbb{T}$, we have $[t] \Vdash_{\exists} [t']$ iff $[t] \leq [t']$.

(\Rightarrow) If $[t] \Vdash_{\exists} [t']$, then, since $[t']$ is approximable, there is an approximation $(t^{\ddagger}, (m_j))$ to $[t']$ and $j \geq 0$ such that $[t] \leq [t^{\ddagger} \upharpoonright m_j]$. By (2)(d) applied to t^{\ddagger} and $[t_0] := [t]$, we have $[t] \leq [t^{\ddagger}] = [t']$.

(\Leftarrow) Assume $[t] \sqsubseteq [t']$. Since $[t']$ is approximable, let $(t^\ddagger, (m_j))$ be an approximation to $[t']$. By (2)(d) applied to t^\ddagger and $[t_0] := [t]$, we have that $[t] \sqsubseteq [t'] = [t^\ddagger]$ implies that there is j such that $[t] \leq [t^\ddagger \upharpoonright m_j]$. So $[t] \Vdash_{\exists} [t']$.

Ad (i). Now, let $[t], [t'] \in \mathbb{T}$ and show $[t] \sqsubseteq [t']$ iff $[t] \sqsubseteq_{\exists} [t']$. If $[t']$ is not approximable, then, by (2)(b), $[t] \sqsubseteq [t']$. And by lemma 2.5.5, $[t] \sqsubseteq_{\exists} [t']$. So assume that $[t']$ is approximable. If $[t]$ is not approximable, then, by (2)(b), we cannot have $[t] \sqsubseteq [t']$ (otherwise $[t']$ is not approximable). And, by definition of \sqsubseteq_{\exists} , we also cannot have $[t] \sqsubseteq_{\exists} [t']$ (since clause (b) is violated).

So let both $[t]$ and $[t']$ be approximable. Then we have

$$\begin{aligned} [t] \sqsubseteq [t'] &\stackrel{(2)(c)}{\Leftrightarrow} \forall [t_0] \in \mathbb{T}_{\text{fin}} : [t_0] \sqsubseteq [t] \Rightarrow [t_0] \sqsubseteq [t'] \\ &\stackrel{(*)}{\Leftrightarrow} \forall [t_0] \in \mathbb{T}_{\text{fin}} : [t_0] \Vdash_{\exists} [t] \Rightarrow [t_0] \Vdash_{\exists} [t'] \Leftrightarrow [t] \sqsubseteq_{\exists} [t'], \end{aligned}$$

where the last equivalence holds since the clause (b) in the definition of \sqsubseteq_{\exists} is trivially satisfied if $[t]$ is approximable.

Ad (ii). Let $t \in T$ be infinite and $0 < n_0 < n_1 < \dots$ and $0 < m_0 < m_1 < \dots$ such that $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$ and $[t \upharpoonright m_0] \leq [t \upharpoonright m_1] \leq \dots$. Let $i \geq 0$ and find $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$.

Note that $[t \upharpoonright n_i] \Vdash_{\exists} [t]$ since $\tau = (t, (n_i))$ is an approximation to $[t]$ and $[t \upharpoonright n_i] \leq [t \upharpoonright n_i]$. In particular, $[t]$ is approximable, so, by (*), we have $[t \upharpoonright n_i] \sqsubseteq [t]$. Now, by (2)(d) applied to t and $0 < m_0 < m_1 < \dots$ and $[t_0] = [t \upharpoonright n_i]$, the fact that $[t \upharpoonright n_i] \sqsubseteq [t]$ implies that there is $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$, as needed.

(4) \Rightarrow (5) Each of the conditions (4)(a)–(d) ensures that \sqsubseteq is in $\{\sqsubseteq_{\forall}, \sqsubseteq_{\exists}, \sqsubseteq_{\text{dom}}\}$ and one of the conditions (1)–(6) of proposition 2.5.8 is satisfied. Thus, that proposition implies $\sqsubseteq_{\forall} = \sqsubseteq_{\exists} = \sqsubseteq_{\text{dom}}$ and M is limit-respecting, and the claim follows.

(5) \Rightarrow (1) Ad 1(a). This follows from $\sqsubseteq = \sqsubseteq_{\forall}$ and lemma 2.5.4 (1) and (2).

Ad 1(b). Let $[t] \in \mathbb{T}$ be non-approximable and $[t'] \in \mathbb{T}$. Concerning (i), $[t'] \sqsubseteq [t]$ follows from $\sqsubseteq = \sqsubseteq_{\forall}$ and lemma 2.5.4 (5). Concerning (ii), if $[t] \sqsubseteq [t']$, then, by $\sqsubseteq = \sqsubseteq_{\forall}$ and the definition of \sqsubseteq_{\forall} , also $[t']$ is not approximable.

Ad 1(c). Let $t \in T$ be infinite and $0 < n_0 < n_1 < \dots$ such that $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$.

Concerning (i), let $i \geq 0$ and show $[t \upharpoonright n_i] \sqsubseteq [t]$. Indeed, $(t, (n_i))$ is an approximation to $[t]$ and $[t \upharpoonright n_i] \leq [t \upharpoonright n_i]$. Hence $[t \upharpoonright n_i] \Vdash_{\exists} [t]$, so, by lemma 2.5.5, $[t \upharpoonright n_i] \sqsubseteq_{\exists} [t]$, so the claim follows from $\sqsubseteq = \sqsubseteq_{\exists}$.

Concerning (ii), let $[t_0] \in \mathbb{T}_{\text{fin}}$ with $[t_0] \sqsubseteq [t]$. Show that there is $i \geq 0$ such that $[t_0] \sqsubseteq [t \upharpoonright n_i]$. Since $\sqsubseteq = \sqsubseteq_{\forall}$, we have, by lemma 2.5.4 (3), that $[t_0] \Vdash_{\forall} [t]$. Since $(t, (n_i))$ is an approximation to $[t]$, there is $i \geq 0$ such that $[t_0] \leq [t \upharpoonright n_i]$. Since \leq coincides with $\sqsubseteq_{\forall} = \sqsubseteq$ on \mathbb{T}_{fin} , $[t_0] \sqsubseteq [t \upharpoonright n_i]$, as needed.

Ad 1(d). Let $[t], [t'] \in \mathbb{T}$ be approximable such that, for all $[t_0] \in \mathbb{T}_{\text{fin}}$, $[t_0] \leq [t]$ implies $[t_0] \leq [t']$. Show $[t] \leq [t']$. Indeed, since $\leq = \sqsubseteq_{\vee}$, we have, by lemma 2.5.4 (4), that $[t] \sqsubseteq_{\vee} [t']$, i.e., $[t] \leq [t']$.

(1) \Rightarrow (2) We prove the items in a different order than stated. Ad (2)(b). This is implied by—or, rather, identical to—(1)(b).

Ad (2)(d). Let $t \in T$ and $0 < n_0 < n_1 < \dots$ be with $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$. Let $[t_0] \in \mathbb{T}_{\text{fin}}$ and show $[t_0] \leq [t]$ iff $[t_0] \leq [t \upharpoonright n_i]$ for some i .

If t is finite, then, there is j such that $n_j > |t|$ and, by (1)(a), we have $[t_0] \leq [t]$ iff $[t_0] \leq [t] = [t \upharpoonright n_j]$ iff $[t_0] \leq [t \upharpoonright n_i]$ for some i , where the reverse direction holds since $[t \upharpoonright n_i] \leq [t \upharpoonright n_j]$: if $i \leq j$ this holds by assumption, and if $i \geq j$, then $t \upharpoonright n_i = t \upharpoonright n_j$.

So let t be infinite. If $[t_0] \leq [t]$, then, by (1)(c)(ii), there is i such that $[t_0] \leq [t \upharpoonright n_i]$, whence, by (1)(a), $[t_0] \leq [t \upharpoonright n_i]$. If $[t_0] \leq [t \upharpoonright n_i]$ for some i , then, by (1)(c)(i), $[t_0] \leq [t \upharpoonright n_i] \leq [t]$, so, by (1)(a), $[t_0] \leq [t]$.

Ad (2)(c). One direction already follows from (1)(d). For the other direction, let $[t], [t'] \in \mathbb{T}$ be approximable with $[t] \leq [t']$. Let $[t_0] \in \mathbb{T}_{\text{fin}}$ with $[t_0] \leq [t]$. Show $[t_0] \leq [t']$.

Since $[t]$ is approximable, let $(t^\dagger, (n_i))$ be an approximation to $[t]$. Since $[t_0] \leq [t] = [t^\dagger]$, there is, by (2)(d) applied to t^\dagger , some $i \geq 0$ such that $[t_0] \leq [t^\dagger \upharpoonright n_i]$.

If t^\dagger is finite, there is $j \geq i$ such that $n_j > |t^\dagger|$, so $[t_0] \leq [t^\dagger \upharpoonright n_i] \leq [t^\dagger \upharpoonright n_j] = [t^\dagger] = [t] \leq [t']$. Thus, by (1)(a), $[t_0] \leq [t']$.

If t^\dagger is infinite, then, by applying (1)(c)(i) to t^\dagger and (n_j) , we have $[t_0] \leq [t^\dagger \upharpoonright i] \leq [t^\dagger] = [t] \leq [t']$. Thus, by (1)(a), $[t_0] \leq [t']$.

Ad (2)(a). Let $[t] \in \mathbb{T}$ be approximable and show that $I_{\leq}([t])$ is an ideal in \mathbb{T}_{fin} . Let $(t^\dagger, (n_i))$ be an approximation to $[t]$.

It is nonempty: We have $[t^\dagger \upharpoonright n_0] \leq [t^\dagger \upharpoonright n_0]$, so, by (2)(d) applied to t^\dagger , we have $[t^\dagger \upharpoonright n_0] \leq [t]$, whence $[t^\dagger \upharpoonright n_0] \in I_{\leq}([t])$.

It is a downset: Let $[t''] \leq [t'] \in I_{\leq}([t])$. So $[t''] \leq [t'] \leq [t]$, whence, by (1)(a), $[t''] \leq [t]$, so $[t''] \in I_{\leq}([t])$.

It is directed: Let $[t'], [t''] \in I_{\leq}([t])$ and find $[t_0] \in I_{\leq}([t])$ such that $[t'], [t''] \leq [t_0]$. So $[t'] \leq [t] = [t^\dagger]$ and $[t''] \leq [t] = [t^\dagger]$. Hence, by (2)(d) applied to t^\dagger , there is $i \geq 0$ and $j \geq 0$ such that $[t'] \leq [t^\dagger \upharpoonright n_i]$ and $[t''] \leq [t^\dagger \upharpoonright n_j]$. Let $k := \max(i, j)$ and define $[t_0] := [t^\dagger \upharpoonright k]$. Then $[t'], [t''] \leq [t_0]$ and $[t_0] \leq [t^\dagger \upharpoonright k]$, so, by (2)(d) applied to t^\dagger , $[t_0] \leq [t^\dagger] = [t]$. Hence $[t_0] \in I_{\leq}([t])$.

This completes the proof of theorem 2.6.3.

2.7 Behavioral transition systems

Guided by the previous sections, we define behavioral transition systems (BTSs) as pre-BTSs that are limit-respecting, and we discuss examples and basic properties.

2.7.1 Definition

The preceding two sections strongly suggest one additional axiom for the notion of a pre-BTS $M = (A, T, \equiv)$: namely, requiring it to be limit-respecting.

First, this will make the three natural candidates for information containment coincide (and be a preorder). Thus, being limit-respecting ensures that we *can* define a satisfying notion of information containment. This was the (i)-requirement for the structures (A, T, \equiv) that we've identified in section 2.3.3.

Second, for countable pre-BTSs, being limit-respecting ensures that this natural notion of information containment also is the only one satisfying the rather mild constraints laid out in item 1 of the characterization theorem (theorem 2.6.3). So the information containment of the structure (A, T, \equiv) is unique in a certain sense.

Third, this also implies that the partial order induced by the information containment preorder is an ω -algebraic domain. (We'll call it the trajectory domain and study it in the next section.) This was the (ii)-requirement for the structures (A, T, \equiv) that we've identified in section 2.3.3.

In short, being limit-respecting is (a) *necessary* to be able to define a sensible notion of information containment and (b) it also already is *sufficient* for that notion to be in a sense unique and to yield a domain of behaviors. In other words, there is no other possible axiom in between satisfying the weaker (i)-demand and the stronger (ii)-demand.

This stability suggests that a good axiomatization for the structures (A, T, \equiv) is to be a pre-BTS that is limit-respecting.

2.7.1. DEFINITION. A *behavioral transition system* (BTS) is a pre-BTS $M = (A, T, \equiv)$ that is limit-respecting (see definition 2.5.7). We call $\sqsubseteq_M := \sqsubseteq_{\forall} = \sqsubseteq_{\exists} = \sqsubseteq_{\text{dom}}$ the *information containment* preorder of M (see definitions 2.5.3 and 2.5.6). We drop the subscript ' M ' when clear from context. We call M *countable* if A is countable.

Now that we've defined BTSs, it's high time to consider examples: Both the 'black box' and 'concurrency' examples from section 2.3 and the examples of pre-BTSs from definition 2.4.5 (extensional, temporal, and intensional equivalence) are, in fact, BTSs. To see this, and many more examples, it will be useful to first introduce some simplifying properties.

2.7.2 Simplifying assumptions

We introduce the following simplifying assumptions on a pre-BTS which help to show being limit-respecting.

2.7.2. DEFINITION. Let $M = (A, T, \equiv)$ be a pre-BTS. We call M

1. *bisimulative* if, for all nonempty finite $t, t' \in T$, if $t \equiv t'$ and $tt_0 \in T$ extends t by one element, then there is a finite extension $t't_1 \in T$ of t' such that $tt_0 \equiv t't_1$. (Note that $t't_1$ need not be a one-element extension.)
2. *extendable* if, for all nonempty finite $t, t' \in T$, if $t \equiv t'$ and $tt'' \in T$ is finite, then $t't'' \in T$ and $tt'' \equiv t't''$.
3. *restrictable* if, for all nonempty finite $t, t', tt_0, t't_1 \in T$, if $t \equiv t'$ and $tt_0 \equiv t't_1$, then, for any $t \preceq t_2 \preceq tt_0$, there is $t' \preceq t_3 \preceq t't_1$ with $t_2 \equiv t_3$.³⁰
4. *full* (resp., *full $_\epsilon$*) if T is the set of all (nonempty) A -trajectories.
5. *extensional* if \equiv is extensional equivalence.

We've encountered (2) as property (d) in section 2.3. The following proposition states how these properties are simplifying.

2.7.3. PROPOSITION. *Let $M = (A, T, \equiv)$ be a pre-BTS. Then*

1. *M is bisimulative iff for all finite $t, t' \in T$, if $t \preceq t'$, then $[t] \leq [t']$.*
2. *We have the following implications:*

$$\begin{array}{ccccccc}
 \text{full}_\epsilon \mathcal{E} \text{ extens.} & \Rightarrow & & & & & \\
 & & \text{extendable} & \Rightarrow & \text{bisimulative} & \Rightarrow & \\
 \text{full } \mathcal{E} \text{ extens.} & \Rightarrow & & & & & \text{limit-respecting} \\
 & & & & \text{restrictable} & \Rightarrow &
 \end{array}$$

3. *If M is bisimulative, then, for all $t \in T$, $[t]$ is approximable.*
4. *If M is bisimulative, the information containment \sqsubseteq is well-defined (since M is limit-respecting) and, for all $t, t' \in T$,*

$$[t] \sqsubseteq [t'] \Leftrightarrow \forall n \geq 0 \exists m \geq 0 : [t \upharpoonright n] \leq [t' \upharpoonright m].$$

PROOF. Ad (1). (\Rightarrow) If t is empty, then $[t] \leq [t']$, so let t be nonempty. And if $t = t'$, then $[t] \leq [t']$, so let $t \prec t'$. To show $[t] \leq [t']$, let $t_0 \in [t]$ and find $t_1 \in [t']$ with $t_0 \preceq t_1$. Consider the one-step extension $tt'(n+1)$ of t where $n := |t| - 1$. Since $t \equiv t_0$, there is a finite extension $t_0t^1 \in T$ such that $tt'(n+1) \equiv t_0t^1$. If $tt'(n+1) = t'$, we choose $t_1 := t_0t^1$, and if not we continue this way: Consider the one-step extension $tt'(n+1)t'(n+2)$ of $tt'(n+1)$. Since $tt'(n+1) \equiv t_0t^1$, there is a finite extension $t_0t^1t^2 \in T$ such that $tt'(n+1)t'(n+2) \equiv t_0t^1t^2$. Since t' is finite, we will eventually obtain an extension $t_1 \in T$ of t_0 such that $t' \equiv t_1$, as needed.

³⁰Note that t_2 and t_3 are nonempty initial segments of trajectories in T and hence in T .

(\Leftarrow) To show that M is bisimulative, let $t, t' \in T$ be finite with $t \equiv t'$ and $tt_0 \in T$ a one-step extension. So tt_0 is finite, too, and, by the assumption, $[t] \leq [tt_0]$. Hence, $t' \in [t]$ can be extended to some $t't_1 \in T$ with $t't_1 \equiv tt_0$.

Ad (2). (full & extensional \Rightarrow extendable) Let $t, t' \in T$ be nonempty finite with $t \equiv t'$ and $tt'' \in T$ finite. By extensionality, t and t' have the same last state. Thus, also $t't''$ is an A -trajectory. Since M is full, $t't'' \in T$. Since tt'' and $t't''$ have the same last state and are finite, extensionality implies $tt'' \equiv t't''$.

(full $_{\epsilon}$ & extensional \Rightarrow extendable) As above: now $t't''$ is in T since it is a nonempty A -trajectory (since t' is nonempty).

(extendable \Rightarrow bisimulative) Let $t, t' \in T$ be nonempty finite with $t \equiv t'$ and $tt_0 \in T$ a one-step extension. In particular, $tt_0 \in T$ is finite. By being extendable, $t't_0 \in T$ is a finite extension of t' and $tt_0 \equiv t't_0$, as needed.

(bisimulative \Rightarrow limit-respecting) Let $t \in T$ be infinite, and let $(n_i)_i$ and $(m_i)_i$ be strictly increasing sequences of positive integers such that $([t \upharpoonright n_i])_i$ and $([t \upharpoonright m_j])_j$ are \leq -increasing. (In particular, $t \upharpoonright n_i, t \upharpoonright m_j \in T$ for all i and j .) Let $i \geq 0$ and find $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$. Choose $j \geq 0$ such that $m_j \geq n_i$. Then $t \upharpoonright n_i \preceq t \upharpoonright m_j$ are finite trajectories in T . Since M is bisimulative, the equivalent condition from (1) implies $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$, as needed.

(restrictable \Rightarrow limit-respecting) Let $t \in T$ be infinite, and let $(n_i)_i$ and $(m_i)_i$ be strictly increasing sequences of positive integers such that $([t \upharpoonright n_i])_i$ and $([t \upharpoonright m_j])_j$ are \leq -increasing. Let $i \geq 0$ and find $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$. Choose $j \geq 0$ and $k \geq i$ such that $n_i \leq m_j \leq n_k$. To show $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$, let $t_a \in [t \upharpoonright n_i]$ and find $t_b \in T$ with $t_a \preceq t_b \equiv t \upharpoonright m_j$. Since $[t \upharpoonright n_i] \leq [t \upharpoonright n_k]$, there is $t_a t_1 \in T$ with $t_a t_1 \equiv t \upharpoonright n_k$. Write $t \upharpoonright n_k = t \upharpoonright n_i t_0 \in T$ and $t_2 := t \upharpoonright m_j$. So $t \upharpoonright n_i \equiv t_a$ and $t \upharpoonright n_i t_0 \equiv t_a t_1$ and $t \upharpoonright n_i \preceq t_2 \preceq t \upharpoonright n_i t_0$. By being restrictable, there is $t_a \preceq t_b \preceq t_a t_1$ such that $t_2 \equiv t_b$. So t_b has the required properties.

Ad (3). Since, by (1), extension implies \leq , $(t, (i+1)_{i \geq 0})$ is an approximation to $[t]$.

Ad (4). (\Rightarrow) Let $n \geq 0$ and find $m \geq 0$ such that $[t \upharpoonright n] \leq [t' \upharpoonright m]$. If $n = 0$, choose $m := 0$, so let $n > 0$. As just seen, $\tau = (t, (n_i))$ with $n_i = i + 1$ and $\tau' = (t', (m_j))$ with $m_j = j + 1$ are approximations to $[t]$ and $[t']$, respectively. Since $[t] \sqsubseteq [t']$ and $\sqsubseteq = \sqsubseteq_{\text{dom}}$, τ' dominates τ , so, for $i := n - 1 \geq 0$, there is $j \geq 0$ such that $[t \upharpoonright n] = [t \upharpoonright n_i] \leq [t' \upharpoonright m_j]$, and we choose $m := m_j \geq 0$.

(\Leftarrow) We show $[t] \sqsubseteq_{\exists} [t']$. Clause (b) of \sqsubseteq_{\exists} is vacuously satisfied, since any $[t]$ is approximable. For clause (a), let $[t_0] \in \mathbb{T}_{\text{fin}}$ with $[t_0] \Vdash_{\exists} [t]$ and show $[t_0] \Vdash_{\exists} [t']$. Since M is a limit-respecting, the orders \sqsubseteq_{\exists} and \sqsubseteq_{\forall} agree, so, by lemmas 2.5.5 and 2.5.4, $[t_0] \Vdash_{\exists} [t]$ implies $[t_0] \Vdash_{\forall} [t]$. Again, $\tau = (t, (i+1)_i)$ and $\tau' = (t', (j+1)_j)$ are approximations to $[t]$ and $[t']$, respectively. Thus, since $[t_0] \Vdash_{\forall} [t]$, there is $i \geq 0$ such that $[t_0] \leq [t \upharpoonright i + 1]$. By the assumption, there is $m \geq 0$ such that $[t \upharpoonright i + 1] \leq [t' \upharpoonright m]$. Hence, for $j := m$, $[t_0] \leq [t' \upharpoonright m] \leq [t' \upharpoonright j + 1]$, whence $[t_0] \Vdash_{\exists} [t']$. \square

2.7.3 Examples

We discuss several examples (and non-examples) of BTSs. The first three are ‘positive’ in the sense of providing BTSs, while the last three are ‘negative’ in the sense of showing that various assumptions that we’ve discussed are not vacuous.

We start with the two guiding examples of ‘black box systems’ and ‘concurrency’ and the extensional, temporal, and intensional equivalence construction: they are all extendable BTSs.

2.7.4. EXAMPLE. (1). In the guiding examples from section 2.3, we’ve considered structures (A, T, \equiv) where A is an LTS, T a set of A -trajectories, and \equiv an equivalence relation on T such that properties (a)–(d) are satisfied. Such structures are extendable BTSs: As seen before, properties (a)–(c) ensure that they are pre-BTSs and property (d) is that of being extendable.

(2). If A is an LTS, T the set of all (nonempty) A -trajectories (starting in i_A) as in example 2.4.3, and \equiv any of the examples from definition 2.4.5 (extensional, temporal, or intensional equivalence), then $M := (A, T, \equiv)$ is a extendable BTS:

Indeed, we know already that M is a pre-BTS, so we need to show that M is extendable. This is done as in the ‘full & extensional’ case above: If $t, t' \in T$ are nonempty finite with $t \equiv t'$ and $tt'' \in T$ finite, then, for all the above choices of \equiv , we have $\text{last}(t) = \text{last}(t')$. Thus, also $t't''$ is an A -trajectory. And $t't''$ is a (nonempty if t' is nonempty) A -trajectory (that starts in i_A if t' starts in i_A), so $t't''$ is again in T . Finally, $tt'' \equiv t't''$ for any of the above choices for \equiv : since $t \equiv t'$, tt'' and $t't''$ have the same last state, and the same length if t and t' have the same length, and the same label-sequence if t and t' have the same label-sequence.

There is a natural way to generalize extensional equivalence: rather than demanding the last states to be identical, one can demand them to be in a bisimulation relation. (This is a generalization since the identity relation on states is a bisimulation.) This yields bisimulative BTSs and does justice to the term ‘bisimulative’.

2.7.5. EXAMPLE. Let A be an LTS and let $\approx \subseteq S_A \times S_A$ be a bisimulation (see e.g. Sangiorgi 2012, ch. 1): for all $s \approx s'$ and $\alpha \in L_A$,

- Forth: If $s \xrightarrow{\alpha} s_0$, then there is $s_1 \in S_A$ with $s' \xrightarrow{\alpha} s_1$ and $s_0 \approx s_1$.
- Back: If $s' \xrightarrow{\alpha} s_1$, then there is $s_0 \in S_A$ with $s \xrightarrow{\alpha} s_0$ and $s_0 \approx s_1$.

Also assume that \approx is an equivalence relation. This is the case if \approx is identity (as in extensional equivalence). The coarsest choice is bisimilarity: $s \sim s'$ iff there is a bisimulation \approx such that $s \approx s'$.

Let T be the set of all A -trajectories and let \equiv be generated (in the sense of proposition 2.4.4) by: for $t, t' \in T$ nonempty finite, $t \equiv t'$ iff $\text{last}(t) \approx \text{last}(t')$.

So $M = (A, T, \equiv)$ is a pre-BTS, and we see that it is bisimulative: Assume $t, t' \in T$ are nonempty finite with $t \equiv t'$ and $tt_0 \in T$ is a one-step extension. Write $t_0 = (s, \alpha, s_0)$. Then $s = \text{last}(t) \approx \text{last}(t') =: s'$. By the forth condition, there is $s_1 \in S_A$ with $s' \xrightarrow{\alpha} s_1$ and $s_0 \approx s_1$. Let $t't_1 := t'(s', \alpha, s_1)$. Since T is the set of all A -trajectories, $t't_1 \in T$, and, since $\text{last}(tt_0) = s_0 \approx s_1 = \text{last}(t't_1)$, we have $tt_0 \equiv t't_1$.

Next, here are examples of restrictable BTSs:

2.7.6. EXAMPLE. (1). In section 2.3.2, we've introduced transition systems A with independence I and said that the independence relation I induces an equivalence relation on the set T of all A -trajectories: $t \equiv t'$ iff $|t| = |t'|$ and, for $n < |t|$, $t(n)$ and $t'(n)$ are occurrences of the same event. This straightforwardly yields a pre-BTS $M = (A, T, \equiv)$, and M is restrictable: Assume $t, t', tt_0, t't_1 \in T$ are nonempty finite with $t \equiv t'$ and $tt_0 \equiv t't_1$. If $t \preceq t_2 \preceq tt_0$, let n be such that $t_2 = tt_0 \upharpoonright n$, then, since tt_0 and $t't_1$ are the same sequences of events, also $t_2 = tt_0 \upharpoonright n$ and $t_3 := t't_1 \upharpoonright n$ are the same sequences of events, and, since $|t| = |t'|$, $t' \preceq t_3 \preceq t't_1$.

(2). In model checking, one adds to an LTS A an interpretation function I assigning each state $s \in S_A$ a subset $I(s)$ of a set of atomic propositions (see e.g. Baier and Katoen 2008). Intuitively, $I(s)$ is the set of basic properties of s (or observations about s). The trace of a trajectory $t = s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots$ then is the sequence $I(s_0), I(s_1), \dots$. So we can choose T as the set of all A -trajectories and \equiv as having the same trace (i.e., being 'observationally equivalent'). This straightforwardly yields a pre-BTS $M = (A, T, \equiv)$, and M is restrictable similarly as in (1) above.

Now to the 'negative' examples. First, we now can see more precisely why we generally cannot appropriately define information containment for infinite behaviors just like for finite behaviors:

2.7.7. EXAMPLE. Consider the following LTS A :



and let T be the set of all A -trajectories and \equiv extensional equivalence. In particular, $M = (A, T, \equiv)$ is a countable bisimulative pre-BTS.

Assume we'd define information containment for all behaviors like for finite ones: $[t] \sqsubseteq [t']$ iff $\forall t_0 \in [t] \exists t_1 \in [t_1] : t_0 \preceq t_1$. As discussed, saying that this is an appropriate definition of information containment means that \sqsubseteq satisfies one of the equivalent items (1)–(5) of theorem 2.6.3. But then $\sqsubseteq = \sqsubseteq_M$ and, by proposition 2.7.3 (4), $[t] \sqsubseteq [t']$ iff $\forall n \exists m : [t \upharpoonright n] \leq [t' \upharpoonright m]$. So these two characterizations of $[t] \sqsubseteq [t']$ should be equivalent.

However, consider the two infinite A -trajectories $t = i \rightarrow s \rightarrow i \rightarrow s \rightarrow \dots$ and $t' = i \rightarrow i \rightarrow i \rightarrow \dots$. Then $\forall n \exists m : [t \upharpoonright n] \leq [t' \upharpoonright m]$ (since there is always a path


from $\text{last}(t \upharpoonright n)$ to $\text{last}(t' \upharpoonright m)$). But we do *not* have $\forall t_0 \in [t] \exists t_1 \in [t_1] : t_0 \preceq t_1$: Otherwise t can be extended to $t_1 \equiv t'$, whence $t = t_1$ and $t \equiv t'$, so t and t' would have the same tail.

The following example shows that the countability assumption in theorem 2.6.3 is necessary. The assumption was used when showing that ι is surjective by employing the fact that a countable directed set has a cofinal chain. This fact fails for uncountable directed sets, and the usual counterexample inspires the example (see e.g. Abramsky and Jung 1994, exercise 2.3.9 (6)).

2.7.8. EXAMPLE. Consider the LTS $A = (S, i, L, \rightarrow)$ where S is the set of finite subsets of the real numbers (so A is uncountable), $i := \emptyset$, $L := \{\alpha\}$, and $s \xrightarrow{\alpha} s'$ iff $s \subseteq s'$. Let $M := (A, T, \equiv)$ where T is the set of all A -trajectories and \equiv is extensional equivalence. So M is an *uncountable* BTS, and we show that $(\overline{\mathbb{T}}, \underline{\equiv})$ is not an ω -algebraic domain: it even fails to be a dcpo. Indeed, $A := \{[t] : t \in T \text{ nonempty finite}\}$ is a directed subset of $\overline{\mathbb{T}}$, but A cannot have an upper bound $[t] \in \overline{\mathbb{T}}$ for some $t \in T$, since then $\bigcup_n s(t)(n) = \mathbb{R}$ (using proposition 2.7.3) would be countable.

Finally, here is a small example to illustrate how being limit-respecting can fail in a pre-BTS.

2.7.9. EXAMPLE. Consider the following unlabeled transition system A where state s_1 has the color red and the states s_2 and s_3 have the color green:



Let T be the set of all A -trajectories starting in i and \equiv is generated by: for $t, t' \in T$ nonempty finite, $t \equiv t'$ iff $\text{last}(t)$ and $\text{last}(t')$ have the same color.

So $M := (A, T, \equiv)$ is a pre-BTS, but it is not limit-respecting: Consider $t = i \rightarrow s_1 \rightarrow s_2 \rightarrow s_1 \rightarrow \dots \in T$. Let (n_i) be the sequence $2 < 4 < \dots$ of even numbers, and let (m_j) be the sequence $1 < 3 < \dots$ of odd numbers. So $t \upharpoonright n_i$ ends in the green state s_2 and $t \upharpoonright m_j$ ends in the red state s_1 . Hence $t \upharpoonright n_i \equiv t \upharpoonright n_{i+1}$ and, in particular, $[t \upharpoonright n_i] \leq [t \upharpoonright n_{i+1}]$. Similarly, $[t \upharpoonright m_j] \leq [t \upharpoonright m_{j+1}]$. However, for $i := 0$ there is no $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$ because $t_0 := i \rightarrow s_3 \in T$ ends in a green state and hence is in $[t \upharpoonright n_i]$, but the only extensions are of the form $t_0 := i \rightarrow s_3 \rightarrow s_3 \rightarrow \dots$ and hence never end in a red state, so cannot be in $[t \upharpoonright m_j]$.

2.8 Trajectory domains

As mentioned, the additional axiom of being limit-respecting is not only enough to define a sensible information containment ordering, it also is sufficient, in the countable case, for this ordering to yield a domain of behaviors:

2.8.1. DEFINITION. Let $M = (A, T, \equiv)$ be a countable BTS. In the notation of theorem 2.6.3, we call the ω -algebraic dcpo $\mathbb{T}(M) := \mathbb{T}(M, \sqsubseteq) = (\overline{\mathbb{T}}, \overline{\sqsubseteq})$ the *trajectory domain* of M .

This raises the question: which ω -algebraic domains can be obtained (up to isomorphism) as trajectory domains of countable BTS? The answer is: all of them.³¹

2.8.2. THEOREM. *For every ω -algebraic domain D , there is a countable BTS M such that D is isomorphic to $\mathbb{T}(M)$. Moreover, M can be chosen to be full_ϵ and extensional.*

PROOF. If D is the empty domain, choose $A := (\{i\}, i, \emptyset, \emptyset)$ and T as the set of all nonempty A -trajectories (i.e., $T = \emptyset$) and \equiv as extensional equivalence (i.e., $\equiv = \emptyset$). So $M := (A, T, \equiv)$ is a BTS that is countable full_ϵ and extensional and $\mathbb{T}(M) = \emptyset \cong D$.

So let D be nonempty. Define the LTS $A = (S, i, L, \rightarrow)$ by:

- $S := K(D)$ (since D is nonempty, the set of compact elements $K(D)$ is nonempty),
- i is any fixed element of $K(D)$.
- $L = \{\alpha\}$ for some object α .
- $s \xrightarrow{\alpha} s'$ iff $s \leq s'$ (in D).

Let T be the set of all nonempty A -trajectories, and let \equiv be extensional equivalence. So $M := (A, T, \equiv)$ is a full_ϵ and extensional BTS and M is countable since A is countable (because $K(D)$ is countable).

We show that $(\mathbb{T}_{\text{fin}}, \leq) \cong (K(D), \leq)$. Then, by theorem 2.6.3 and the fact that M has no non-approximable elements,

$$\mathbb{T}(M) \cong \overline{\text{Idl}}(\mathbb{T}_{\text{fin}}) = \text{Idl}(\mathbb{T}_{\text{fin}}) \cong \text{Idl}(K(D)) \cong D,$$

³¹In the context of concurrent computation (section 2.3.2), this question has been investigated for the domain constructions used there: see Droste (1990), Bracho and Droste (1994), and Stark (1990). Given that BTSs are a considerable generalization of the transition systems considered there, we should expect a considerable larger class of domains—indeed, the largest possible as the theorem shows.

where the last isomorphism is a basic fact about algebraic domains.

We define $\iota : \mathbb{T}_{\text{fin}} \rightarrow K(D)$ by $\iota([t]) = \text{last}(t)$. This is well-defined: If $[t] \in \mathbb{T}_{\text{fin}}$ for $t \in T$, then t is finite nonempty, so $\text{last}(t)$ is defined. And if $[t] = [t']$, then $\text{last}(t) = \text{last}(t')$ by extensional equivalence. This is injective: If $[t] \neq [t']$, then $t \neq t'$, so $\text{last}(t) \neq \text{last}(t')$. And surjective: If $x \in K(D)$, then $t := x \xrightarrow{\alpha} x$ is a nonempty finite trajectory in A , so $[t] \in \mathbb{T}_{\text{fin}}$ and $\iota([t]) = \text{last}(t) = x$.

It remains to show that ι is an order-isomorphism. Let $[t], [t'] \in \mathbb{T}_{\text{fin}}$, and show $[t] \leq [t']$ iff $\iota([t]) \leq \iota([t'])$.

Assume $[t] \leq [t']$. Hence t can be extended to $t_1 \equiv t'$. So there is a trajectory from $\text{last}(t)$ to $\text{last}(t_1) = \text{last}(t')$. Since $s \xrightarrow{\alpha} s'$ implies $s \leq s'$, we have $\iota([t]) = \text{last}(t) \leq \text{last}(t') = \iota([t'])$.

Assume $\iota([t]) \leq \iota([t'])$. Then $\text{last}(t) = \iota([t]) \leq \iota([t']) = \text{last}(t')$. So $t_1 := t(\text{last}(t), \alpha, \text{last}(t')) \in T$ is an extension of t with $t_1 \equiv t'$. Since M is bisimulative, $[t] \leq [t_1] = [t']$. \square

An important corollary is that for every BTS there is a particularly simple one which has the same behavior:

2.8.3. COROLLARY. *For every countable BTS M there is a countable full_ϵ and extensional BTS N such that M and N have the same behavior in the sense that their trajectory domains are isomorphic.*

2.9 Generalizations of information systems

We argue that we can regard the notion of a BTS and their induced trajectory domains as a generalization of the well-known notion of a Scott information system and their induced Scott domains.

2.9.1 Scott information systems . . .

Scott information systems were introduced by Scott (1982). They are important both as a technical tool for ‘doing domain theory’ (by representing Scott domains through their more manageable bases) and as a conceptual tool for motivating domains and providing connections to other fields (event structures, logic, locale theory, etc.). For references see Winskel (1993, sec. 12) and Abramsky and Jung (1994, sec. 8.1.4), and for a general categorical treatment see Edalat and Smyth (1993). Here we’ll use the definition of Winskel (1993, ch. 12).

2.9.1. DEFINITION. An *information system* is a triple $I = (U, \text{Con}, \vdash)$ where U is a countable set (information *tokens*), Con is a non-empty class of finite subsets of U (*consistent sets*), and $\vdash \subseteq (\text{Con} \setminus \{\emptyset\}) \times U$ (*entailment relation*) such that

1. If $X \subseteq Y \in \text{Con}$, then $X \in \text{Con}$.

2. If $a \in U$, then $\{a\} \in \mathbf{Con}$.
3. If $X \vdash a$, then $X \cup \{a\} \in \mathbf{Con}$.
4. If $a \in X \in \mathbf{Con}$, then $X \vdash a$.
5. If $X, Y \in \mathbf{Con}$ and $X \vdash Y$ (i.e., $X \vdash b$ for all $b \in Y$), then $Y \vdash a$ implies $X \vdash a$.

An *element* of I is a subset $x \subseteq U$ such that

1. $x \neq \emptyset$
2. If $X \subseteq x$ is finite, then $X \in \mathbf{Con}$.
3. If $X \subseteq x$ and $X \vdash a$, then $a \in x$.

The set of elements of I is denoted $|I|$. For $X \in \mathbf{Con}$, define $\overline{X} := \{a \in U : X \vdash a\}$.

The point of information systems is that they induce domains (see e.g. Winskel 1993, prop. 12.8): $\mathbf{D}_S(I) := (|I|, \subseteq)$ is an ω -algebraic dcpo where every nonempty subset with an upper bound has a least upper bound.³² The compact elements are of the form \overline{X} for $\emptyset \neq X \in \mathbf{Con}$.

Some useful basic facts are the following:

2.9.2. LEMMA. *Let $I = (U, \mathbf{Con}, \vdash)$ be an information system. Then*

1. *Monotonicity: For $Y \subseteq X \in \mathbf{Con}$, if $Y \vdash a$, then $X \vdash a$.*
2. *If $X \in \mathbf{Con}$ and $X \vdash \{a_1, \dots, a_n\}$ ($n \geq 1$), then $X \cup \{a_1, \dots, a_n\} \in \mathbf{Con}$.*
3. *If $X, Y \in \mathbf{Con}$ and $X \dashv\vdash Y$ (i.e., $X \vdash Y$ and $Y \vdash X$), then $\overline{X} = \overline{Y}$.*

PROOF. Ad (1). By axiom 1, $Y \in \mathbf{Con}$. By axiom 4, $X \vdash b$ for all $b \in Y \subseteq X$, so $X \vdash Y$. Since $Y \vdash a$, axiom 5 implies $X \vdash a$.

Ad (2). First, we have $X \vdash a_1$, so, by axiom 3, $X \cup \{a_1\} \in \mathbf{Con}$. Now, we proceed inductively for $i = 2, \dots, n$: Assume $X \cup \{a_1, \dots, a_{i-1}\} \in \mathbf{Con}$, and show $X \cup \{a_1, \dots, a_i\} \in \mathbf{Con}$. We have $X \vdash a_i$. By monotonicity, since $X \subseteq X \cup \{a_1, \dots, a_{i-1}\} \in \mathbf{Con}$, also $X \cup \{a_1, \dots, a_{i-1}\} \vdash a_i$. By axiom 3, $X \cup \{a_1, \dots, a_{i-1}\} \cup \{a_i\} \in \mathbf{Con}$, as needed.

Ad (3). Let $a \in U$ and show $X \vdash a$ iff $Y \vdash a$. If $X \vdash a$, then, since $Y \vdash X$, axiom 5 implies $Y \vdash a$. Similarly for the other direction. \square

³²Such domains hence are Scott domains without, possibly, a least element, whence they also are called Scott predomains.

2.9.2 ... and their generalizations as BTSs

We show that we can interpret an information system I as a countable BTS M_I such that the trajectory domain $\mathsf{T}(M_I)$ of M_I is isomorphic to the domain $\mathsf{D}_S(I)$ induced by I . In that sense, we can regard BTSs and the trajectory domain construction as a generalization of information systems and the ‘set of elements’ construction.

The main intuition for the BTS M_I that interprets $I = (U, \mathsf{Con}, \vdash)$ is to think of the consistent sets $X \in \mathsf{Con}$ as trajectories (modulo order) through the space of information tokens U that satisfy the global constraint of being consistent. Formally, we do this as follows.

2.9.3. DEFINITION. Let $I = (U, \mathsf{Con}, \vdash)$ be an information system. Define $M_I := (A, T, \equiv)$ as follows:

- $S_A := U \cup \{i\}$ where i is some object not in U .
- $i_A := i$.
- $L_A := \{\alpha\}$ (i.e., A essentially is ‘unlabeled’ and we omit labels in \rightarrow).
- $a \rightarrow b$ iff $a, b \in S_A$ (so $\rightarrow = S_A \times S_A$ is the trivial relation).
- $T :=$ the set of all A -trajectories t with the following properties:
 - (a) t is nonempty (i.e., $|t| > 0$) and of the form $i \rightarrow a_1 \rightarrow a_2 \rightarrow \dots$ for $a_i \in U$. (Hence, if t is finite, then $\text{last}(t) = a_{|t|}$.)
 - (b) For all $n \geq 1$, if $n \leq |t|$, then $\{a_1, \dots, a_n\} \in \mathsf{Con}$.
- (For $t \in T$ finite, let $S(t)$ be the set of states occurring in t and $S_i(t) := S(t) \setminus \{i\}$. Note $S_i(t) \in \mathsf{Con}$ by (b).)
- \equiv is the equivalence induced by: for $t, t' \in T$ finite nonempty, $t \equiv t'$ iff $\overline{S_i(t)} = \overline{S_i(t')}$.

Thus, the ‘globally possible’ trajectories through the space S of information tokens (together with an additional starting state i) are precisely those with consistent initial segments. And two such finite trajectories are behaviorally equivalent if they contain the same information, i.e., the information that can be deduced from the information tokens that they visit is the same.

Thus, one way that BTSs generalize Scott information systems, is as follows: In information systems, two consistent trajectories t and t' of information tokens are considered to be equivalent if, roughly, they are entailment equivalent: $t \dashv\vdash t'$. This ‘logic’ is monotonic and insensitive to count and order of premises. Thus, one could move to non-monotonic or resource sensitive logics (like linear logic) and their

respective notion of equivalence \equiv , to obtain more general (BTS representations of) information systems.³³

We discuss this further in the next section, but now let's prove the two announced claims:

2.9.4. PROPOSITION. *If I is an information system, then M_I is a countable bisimulative BTS.*

PROOF. We first show that M_I is a pre-BTS. We show that T satisfies axioms (1) and (2)*; then the claim follows by proposition 2.4.4.

Concerning axiom (1), let t' be a nonempty finite initial segment of $t \in T$, and show $t' \in T$. By definition, t is of the form $i \rightarrow a_1 \rightarrow a_2 \rightarrow \dots$ and t' is of the form $i \rightarrow a_1 \rightarrow a_2 \rightarrow \dots a_n$ for some $1 \leq n \leq |t|$. And, since $t \in T$, we in particular have, for all $1 \leq m \leq n = |t'| \leq |t|$, that $\{a_1, \dots, a_m\} \in \mathbf{Con}$. Hence $t' \in T$.

Concerning (2)*, let t be an infinite A -trajectory such that $t \notin T$, and find $n \geq 1$ such that $t \upharpoonright n \notin T$. So t fails to have property (a) or (b). If it fails (a), it is not of the form $i \rightarrow a_1 \rightarrow a_2 \rightarrow \dots$, i.e., t either doesn't start with i or it goes back to i after having started with i . Thus, some nonempty initial segment of t fails to be of the form $i \rightarrow a_1 \rightarrow a_2 \rightarrow \dots$, and we can choose $n \geq 1$ large enough such that $t \upharpoonright n$ includes that initial segment, whence $t \upharpoonright n \notin T$. So assume t has (a) but fails (b), whence t is of the form $i \rightarrow a_1 \rightarrow a_2 \rightarrow \dots$ but there is $n \geq 1$ with $n \leq |t|$ and $\{a_1, \dots, a_n\} \notin \mathbf{Con}$. Then $t \upharpoonright n$ fails to have property (b), so $t \upharpoonright n \notin T$.

Next, note that M_I is countable since A is countable (since S_A is countable because U is countable). So it remains to show that M_I is bisimulative.

So let $t, t' \in T$ be finite nonempty with $t \equiv t'$ and $tt_0 \in T$ a one-step extension. We need to find an extension $t't_1 \in T$ such that $tt_0 \equiv t't_1$. (In fact, we'll show that we can choose $t't_1$ as one-step extension, too.)

Since $t, t' \in T$ are finite, they are of the form $t = i \rightarrow a_1 \rightarrow a_2 \rightarrow \dots \rightarrow a_n$ and $t' = i \rightarrow a'_1 \rightarrow a'_2 \rightarrow \dots \rightarrow a'_m$ for $n, m \geq 1$ and $t_0 = a_n \rightarrow b$.

We claim that $t't_1 := i \rightarrow a'_1 \rightarrow a'_2 \rightarrow \dots \rightarrow a'_m \rightarrow b$ is in T . Since \rightarrow is the trivial relation, $t't_1$ is an A -trajectory and it satisfies (a). So we need to show that it satisfies (b). It suffices to show $\{a'_1, \dots, a'_m, b\} \in \mathbf{Con}$. Since $tt_0 \in T$, we know that $\{a_1, \dots, a_n, b\} \in \mathbf{Con}$. Since $t \equiv t'$, we have $\overline{S_i(t)} = \overline{S_i(t')}$. So, for $a'_j \in S_i(t')$, we have $S_i(t) \vdash a'_j$. Since $S_i(t) = \{a_1, \dots, a_n\} \subseteq \{a_1, \dots, a_n, b\} \in \mathbf{Con}$, monotonicity implies $S_i(t) \cup \{b\} \vdash a'_j$. Hence, by lemma 2.9.2 (2), $S_i(t) \cup \{b\} \cup \{a'_1, \dots, a'_m\} \in \mathbf{Con}$. Then $\{a'_1, \dots, a'_m, b\}$ is a subset of a set in \mathbf{Con} and hence in \mathbf{Con} by axiom 1.

So it remains to show $tt_0 \equiv t't_1$. Note that $S_i(tt_0) = S_i(t) \cup \{b\}$ and $S_i(t't_1) = S_i(t') \cup \{b\}$, so it suffices to show $\overline{S_i(t) \cup \{b\}} = \overline{S_i(t') \cup \{b\}}$. Since $S_i(t') \cup \{b\}$

³³For some benchmark axiomatizations of logical equivalence (or synonymy) that are finer than classical logic, see Hornischer (2020).

and $S_i(t) \cup \{b\}$ are in \mathbf{Con} this follows from lemma 2.9.2 (3) once we can show $S_i(t') \cup \{b\} \dashv\vdash S_i(t) \cup \{b\}$.

Concerning \vdash , let $c \in S_i(t) \cup \{b\}$ and show $S_i(t') \cup \{b\} \vdash c$. If $c = b$, then $c \in S_i(t') \cup \{b\} \in \mathbf{Con}$, and the claim follows by axiom 4. So let $c \in S_i(t) \in \mathbf{Con}$. Then, again by axiom 4, $S_i(t) \vdash c$, so $c \in \overline{S_i(t)} = \overline{S_i(t')}$, so $S_i(t') \vdash c$. Since $S_i(t') \subseteq S_i(t') \cup \{b\} \in \mathbf{Con}$, monotonicity implies $S_i(t') \cup \{b\} \vdash c$. The other direction is shown analogously. \square

2.9.5. PROPOSITION. *If I is an information system, then $\mathbb{T}(M_I) \cong \mathbb{D}_S(I)$.*

PROOF. Write $M_I = (A, T, \equiv)$. We claim that $\iota : K(\mathbb{T}(M_I)) \rightarrow K(\mathbb{D}_S(I))$ defined by $[[t]] \mapsto \overline{S_i(t)}$ is a well-defined isomorphism. Then the claim follows since $\mathbb{T}(M_I)$ and $\mathbb{D}_S(I)$ are algebraic domains (so they are isomorphic to the ideal completion of their compact elements). To do so, we show:

1. If $t \in T$ is finite, then $\emptyset \neq S_i(t) \in \mathbf{Con}$, whence $\overline{S_i(t)}$ is a compact element of $\mathbb{D}_S(I)$.
2. For all finite $t, t' \in T$, $[t] \leq [t']$ iff $\overline{S_i(t)} \subseteq \overline{S_i(t')}$.
3. If x is a compact element of $\mathbb{D}_S(I)$, then there is a finite trajectory $t \in T$ with $\overline{S_i(t)} = x$.

Then (1) shows that, for $[[t]] \in K(\mathbb{T}(M_I))$, we have that $\iota([[t]]) \in K(\mathbb{D}_S(I))$ and (2) shows that the mapping is well-defined. Moreover, ι is injective by (2) and surjective by (3). Finally, it is an order-isomorphism by (2).

Ad (1). So t is of the form $i \rightarrow a_1 \rightarrow \dots \rightarrow a_n$ with $n \geq 1$ and $S_i(t) = \{a_1, \dots, a_n\} \in \mathbf{Con}$ is nonempty.

Ad (2). Let $t, t' \in T$ be finite. So $t = i \rightarrow a_1 \rightarrow \dots \rightarrow a_n$ and $t' = i \rightarrow a'_1 \rightarrow \dots \rightarrow a'_m$ for $n, m \geq 1$.

(\Rightarrow) If $[t] \leq [t']$, then t can be extended to $t_1 \in T$ with $t_1 \equiv t'$. Hence $S_i(t) \subseteq S_i(t_1)$ are in \mathbf{Con} , so, by monotonicity, $\overline{S_i(t)} \subseteq \overline{S_i(t_1)} = \overline{S_i(t')}$.

(\Leftarrow) Assume $\overline{S_i(t)} \subseteq \overline{S_i(t')}$. Consider

$$t_1 := t \ a_n \rightarrow a'_1 \rightarrow \dots \rightarrow a'_m.$$

This is an A -trajectory since t ends in a_n and \rightarrow_A is the trivial relation. We will show that $t_1 \in T$ and $t_1 \equiv t'$. Then, because M_I is bisimulative, $[t] \leq [t_1] = [t']$, as needed.

To do so, we'll first show that

$$(*) \ S_i(t') \vdash S_i(t_1).$$

Indeed, let $a \in S_i(t_1) = S_i(t) \cup S_i(t')$ and show $S_i(t') \vdash a$. If $a \in S_i(t') \in \mathbf{Con}$, then, by axiom 4, $S_i(t') \vdash a$. So let $a \in S_i(t)$. Then, again by axiom 4, $S_i(t) \vdash a$. So $a \in \overline{S_i(t)} \subseteq \overline{S_i(t')}$, whence $S_i(t') \vdash a$.

In particular, $S_i(t_1) \in \mathbf{Con}$: By (*), we have $S_i(t') \vdash \{a_1, \dots, a_n\}$, so, by lemma 2.9.2 (2), $S_i(t_1) = S_i(t') \cup \{a_1, \dots, a_n\} \in \mathbf{Con}$.

Now we show $t_1 \in T$. Indeed, the A -trajectory t_1 is of the right form, whence it satisfies (a), and, since $S_i(t_1) \in \mathbf{Con}$, also subsets thereof are in \mathbf{Con} , so t_1 satisfies (b).

So it remains to show that $t_1 \equiv t'$: We have $S_i(t_1), S_i(t') \in \mathbf{Con}$ and $S_i(t') \vdash S_i(t_1)$ by (*) and $S_i(t_1) \vdash S_i(t')$ because of axiom 4 and $S_i(t') \subseteq S_i(t_1) \in \mathbf{Con}$. So lemma 2.9.2 (3) implies $\overline{S_i(t_1)} = \overline{S_i(t')}$, as needed.

Ad (3). If x is a compact element of $\mathbf{D}_S(I)$, then $x = \overline{X}$ for $\emptyset \neq X \in \mathbf{Con}$. In particular, $X = \{a_1, \dots, a_n\}$ is finite nonempty. Then $t := i \rightarrow a_1 \rightarrow \dots \rightarrow a_n$ is a finite A -trajectory satisfying (a) and (b), so $t \in T$. And $\overline{S_i(t)} = \overline{X} = x$. \square

2.10 Conclusion

We conclude with six open questions for future work.

First, arguably the most pressing question by now is about the *category* of BTSs: We've introduced and studied BTSs as objects, but how do they relate to each other, i.e., what are morphisms between BTSs? Does this capture the common idea of one system being simulated in another? Does the trajectory domain construction respect these relations, i.e., is functorial?³⁴ After all, a lesson from Winskel and Nielsen (1995) is that only with this categorical structure can we consider BTSs as a computational model that we can fruitfully relate to other computation models. We'll study these (and more) questions in the next chapter and provide a positive answer.

Second, we've seen that BTSs generalize Scott information systems. Further work should investigate this, for example, by considering (as indicated) various classes of 'generalized information systems' that correspond to processing information according to various finer substructural logics. We explore one direction in the next chapter by showing that the trajectory domains provide an interpretation to relevance logic. Another direction could be to consider the closely related linear logic:

Third, in the game semantics for linear logic of Abramsky and Jagadeesan (1994), the meaning of a formula is a game and a proof for the formula is a

³⁴We may regard this as one implication of the equivalence between operational and denotations semantics that the full abstraction problem (mentioned in the introduction) asks for: equivalence (i.e., isomorphism) in the operational semantics implies equivalence (i.e., isomorphism) in the denotational semantics. We discuss this further in the next chapter.

winning strategy for this game.³⁵ Using their notation for a game, it is tempting to try to view a game as a BTS (A, T, \equiv) : the state space $S = M \times \{P, O\}$ is the set of moves M labeled by whether it is a move of Player or Opponent and the transition relation is the trivial one, $T = \pi \cup \pi^\infty$ consists of the set π of finite trajectories that are possible in the game (where Player and Opponent are alternating) together with the set π^∞ of infinite trajectories all whose initial segments are in π , and trajectory equivalence is chosen in way to capture strategies (maybe indistinguishability by strategies?). Can this game semantics fruitfully be captured this way? And, to come full circle, how does this relate to the solution to the full abstraction problem (mentioned in the introduction) which this semantics provided (Abramsky and McCusker 1999)?

Fourth, how do these different logical perspectives relate to existing logics for LTSs like linear temporal logic (see e.g. Baier and Katoen 2008), and could they provide a domain theory for trajectory domains ‘in logical form’ (Abramsky 1991)?

Fifth, Bratteli–Vershik diagrams play an important role in the study of zero-dimensional topological systems. (See Downarowicz and Karpel 2016 for a brief introduction and references.) At least superficially, they have some ‘BTS-like’ structure: they are certain graphs, their space of infinite paths represents dynamical systems, and also orders on the space of all finite and infinite paths are considered. Can they be fruitfully captured as BTSs?

Sixth, an algebraic way to analyze a graph (i.e., an unlabeled transition system) is through its Leavitt path algebra (Abrams, Ara, and Siles Molina 2017): These algebras can be seen as algebraic analogues of C^* -algebras and are constructed based on the idea of identifying certain paths of the underlying graph. Is there a connection?

Appendix

PROOF OF LEMMA 2.5.4. Item (1), that $(\mathbb{T}, \sqsubseteq_{\forall})$ is a preorder, is immediate from the definition. Before getting to the other items, we show two claims:

(C1) For all $[t] \in \mathbb{T}_{\text{fin}}$, $[t] \Vdash_{\forall} [t]$.

Proof: To show $[t] \Vdash_{\forall} [t]$, let $\tau = (t^\dagger, (n_i))$ be an approximation to $[t]$, and find $i \geq 0$ such that $[t] \leq [t^\dagger \upharpoonright n_i]$. Indeed, since $[t] = [t^\dagger]$ and t is finite, also t^\dagger is finite. Since (n_i) is indefinitely increasing, there is i such that $n_i \geq |t^\dagger|$. Then we have $[t] = [t^\dagger] = [t^\dagger \upharpoonright n_i]$, which implies $[t] \leq [t^\dagger \upharpoonright n_i]$.

(C2) For $[t], [t'] \in \mathbb{T}_{\text{fin}}$, $[t] \Vdash_{\forall} [t']$ implies $[t] \leq [t']$.

Proof: Assume $[t] \Vdash_{\forall} [t']$. Consider the approximation $\tau = (t', (n_i))$ to $[t']$ with $n_i := |t'| + 1 + i$. Then there is $i \geq 0$ such that $[t] \leq [t' \upharpoonright n_i] = [t']$, as needed.

³⁵For more on the topic of ‘games in logic’ see Van Benthem (2014) or Hodges and Väänänen (2019).

Ad (2). Let $[t], [t'] \in \mathbb{T}_{\text{fin}}$, and show $[t] \sqsubseteq_{\forall} [t']$ iff $[t] \leq [t']$. Assume $[t] \sqsubseteq_{\forall} [t']$. For $[t_0] := [t]$ we have, by (C1), $[t_0] \Vdash_{\forall} [t]$. Hence $[t_0] = [t] \Vdash_{\forall} [t']$. By (C2), $[t] \leq [t']$.

Conversely, assume $[t] \leq [t']$. To show $[t] \sqsubseteq_{\forall} [t']$, first observe that condition (b) is trivially satisfied since $[t]$ is finite and hence approximable. For condition (a), let $[t_0] \in \mathbb{T}_{\text{fin}}$ with $[t_0] \Vdash_{\forall} [t]$, and show $[t_0] \Vdash_{\forall} [t']$. So let $(t^{\ddagger}, (m_j))$ be an approximation to $[t']$ and find $j \geq 0$ such that $[t_0] \leq [t^{\ddagger} \uparrow m_j]$. Since t' is finite, also t^{\ddagger} is finite, so there is $j \geq 0$ such that $m_j \geq |t^{\ddagger}|$. By (C2), $[t_0] \Vdash_{\forall} [t]$ implies $[t_0] \leq [t]$. Hence, $[t_0] \leq [t] \leq [t'] = [t^{\ddagger}] = [t^{\ddagger} \uparrow m_j]$, as needed.

Ad (3). Let $[t] \in \mathbb{T}_{\text{fin}}$ and $[t'] \in \mathbb{T}$ and show $[t] \Vdash_{\forall} [t']$ iff $[t] \sqsubseteq_{\forall} [t']$. If $[t] \sqsubseteq_{\forall} [t']$, then, by (C1), $[t] \Vdash_{\forall} [t]$, so, by condition (a), $[t] \Vdash_{\forall} [t']$.

Conversely, assume $[t] \Vdash_{\forall} [t']$. To show $[t] \sqsubseteq_{\forall} [t']$, condition (b) is trivially satisfied since $[t]$ is finite and hence approximable, and for condition (a) let $[t_0] \in \mathbb{T}_{\text{fin}}$ with $[t_0] \Vdash_{\forall} [t]$, and show $[t_0] \Vdash_{\forall} [t']$. So let $(t^{\ddagger}, (m_j))$ be an approximation to $[t']$ and find $j \geq 0$ such that $[t_0] \leq [t^{\ddagger} \uparrow m_j]$. Since $[t] \Vdash_{\forall} [t']$, there is $j \geq 0$ such that $[t] \leq [t^{\ddagger} \uparrow m_j]$. By (C2), $[t_0] \Vdash_{\forall} [t]$ implies $[t_0] \leq [t]$. So $[t_0] \leq [t] \leq [t^{\ddagger} \uparrow m_j]$, as needed.

Ad (4). This follows from (3) and the definition of \sqsubseteq_{\forall} which, for approximable $[t], [t'] \in \mathbb{T}$, reduces to just condition (a).

Ad (5). If $[t] \in \mathbb{T}$ doesn't have an approximation, then $[t]$ is infinite (since all finite trajectories have an approximation) and, for $[t'] \in \mathbb{T}$, we have $[t'] \sqsubseteq_{\forall} [t]$ because condition (a) holds vacuously since $[t_0] \Vdash_{\forall} [t]$ holds vacuously, and condition (b) holds vacuously since $[t]$ is not approximable. \square

PROOF OF 2.5.5. Item (1), that $(\mathbb{T}, \sqsubseteq_{\forall})$ is a preorder, is immediate from the definition. Before getting to the other items, we show two claims:

(C1) If $[t] \in \mathbb{T}_{\text{fin}}$, then $[t] \Vdash_{\exists} [t]$.

Proof: We have that $(t, (|t| + 1 + i)_{i \geq 0})$ is an approximation to $[t]$ and $[t] \leq [t \uparrow |t| + 1 + 0]$, whence $[t] \Vdash_{\exists} [t]$.

(C2) For $[t], [t'] \in \mathbb{T}_{\text{fin}}$, if $[t] \Vdash_{\exists} [t']$, then $[t] \leq [t']$.

Proof: Since $[t']$ is finite, it is approximable, so $[t] \Vdash_{\exists} [t']$ holds because there is an approximation $(t^{\ddagger}, (m_j))$ to $[t']$ and $j \geq 0$ such that $[t] \leq [t^{\ddagger} \uparrow m_j]$. Choose $k \geq j$ big enough such that $m_k > |t^{\ddagger}|$ (t^{\ddagger} is finite since it is equivalent to the finite t'). Then $[t] \leq [t^{\ddagger} \uparrow m_j] \leq [t^{\ddagger} \uparrow m_k] = [t^{\ddagger}] = [t']$, as needed.

Ad (2). Let $[t], [t'] \in \mathbb{T}_{\text{fin}}$. And show $[t] \sqsubseteq_{\exists} [t']$ iff $[t] \leq [t']$. Assume $[t] \sqsubseteq_{\exists} [t']$. By (C1), $[t] \Vdash_{\exists} [t]$, so, by condition (a), $[t] \Vdash_{\exists} [t']$, so, by (C2), $[t] \leq [t']$.

Conversely, assume $[t] \leq [t']$. To show $[t] \sqsubseteq_{\exists} [t']$, condition (b) is satisfied since $[t]$ is approximable, and for condition (a), let $[t_0] \in \mathbb{T}_{\text{fin}}$ with $[t_0] \Vdash_{\exists} [t]$, and show $[t_0] \Vdash_{\exists} [t']$. Consider the approximation $(t', (|t'| + 1 + j)_j)$ to $[t']$ and $j := 0$. Then $[t] \leq [t'] = [t' \uparrow |t'| + 1 + j]$, as needed.

Ad (3). Let $[t] \in \mathbb{T}_{\text{fin}}$ and $[t'] \in \mathbb{T}$, and show $[t] \Vdash_{\exists} [t']$ iff $[t] \sqsubseteq_{\exists} [t']$. Assume $[t] \sqsubseteq_{\exists} [t']$. For $[t_0] := [t]$ we have, by (C1), $[t_0] \Vdash_{\exists} [t]$, so $[t] = [t_0] \Vdash_{\exists} [t']$.

Conversely, assume $[t] \Vdash_{\exists} [t']$. To show $[t] \sqsubseteq_{\exists} [t']$, condition (b) is satisfied since $[t]$ is approximable, and for condition (a), let $[t_0] \in \mathbb{T}_{\text{fin}}$ with $[t_0] \Vdash_{\exists} [t]$, and show $[t_0] \Vdash_{\exists} [t']$. If $[t']$ is not approximable, then $[t_0] \Vdash_{\exists} [t']$, so let $[t']$ be approximable

and $(t^\ddagger, (m_j))$ an approximation to $[t']$. Since $[t] \Vdash_{\exists} [t']$ (and $[t']$ is approximable), there is $j \geq 0$ such that $[t] \leq [t^\ddagger \upharpoonright m_j]$. By (C2), $[t_0] \Vdash_{\exists} [t]$ implies $[t_0] \leq [t]$. So $[t_0] \leq [t^\ddagger \upharpoonright m_j]$, as needed.

Ad (4). This follows from (3) and the definition of \sqsubseteq_{\exists} which, for approximable elements, reduces to just condition (a).

Ad (5). If $[t] \in \mathbb{T}$ doesn't have an approximation, then $[t]$ is infinite (since all finite trajectories have an approximation) and, for $[t'] \in \mathbb{T}$, we have $[t'] \sqsubseteq_{\exists} [t]$ because condition (a) holds vacuously since $[t_0] \Vdash_{\exists} [t]$ holds vacuously, and condition (b) holds vacuously since $[t]$ is not approximable. \square

Chapter 3

Trajectory domains 2: Category

Abstract In the previous chapter, we provided a denotational semantics to labeled transition systems (LTS): We introduced the notion of a behavioral transition system (BTS) which extends an LTS by some structure to specify its behavior, and, for countable systems, we constructed their trajectory domain which serves as their denotation (or ‘behavior description’). In this chapter, we complete this construction category-theoretically: We introduce the category $(\omega)\text{BTS}$ of (countable) BTSs and show that the trajectory domain construction extends to a functor $\mathbb{T} : \omega\text{BTS} \rightarrow \omega\text{ALG}$ into the category of ω -algebraic domains. The main result is that we build an adjunction between a subcategory of ωBTS and a version of ωALG : thus, the well-known ‘computational model’ of ω -algebraic domains can be embedded into (i.e., can be abstracted from) the computational model of BTSs. We also note that the trajectory domain construction naturally leads to a new interpretation of relevance logic in terms of LTSs.

3.1 Introduction


Labeled transition systems (LTS) are a widely used computational model providing operational semantics to systems (or processes): in the previous chapter, we’ve mentioned as examples computer programs (or Turing machines more generally), reactive systems interacting with a nondeterministic environment, model checking, concurrent computation, or observing dynamical systems. So they can be seen as a general model of symbolic computation. An LTS provides operational meaning in the sense of describing the possible states of the system and their dynamics, i.e., how the system can transform from one state to another.¹ In the previous chapter, we constructed a corresponding denotational semantics: a more system-independent and static description of the possible behavior of the system that facilitates mathematical analysis.

¹The states could be fairly ‘low-level’ (e.g., describing the tape and internal state of a Turing machine) or more ‘high-level’ (e.g., bundling together low-level states with a similar function).

To this end, we introduced the notion of a behavioral transition system (BTS). This extends an LTS A by two more entities to specify its behavior: First, a set T of A -trajectories that not only are ‘locally’ possible (each step being a possible transition in the LTS), but also ‘globally’ possible (e.g., reflecting memory constraints); and second, an equivalence relation on T to say that two trajectories are instances of the same (type of) behavior (e.g., two concurrent versions of the same computation). So an equivalence class describes a possible behavior of A (relative to T and \equiv), and the set of equivalence classes T/\equiv describes the possible behavior—and thus acts like a denotation of the LTS. We defined BTSs as such structures $M = (A, T, \equiv)$ satisfying five axioms capturing the intended interpretation. We’ve shown that, for countable systems, there is essentially a unique way of defining an information containment order on T/\equiv turning it into an ω -algebraic domain (a well-behaved partial order studied in domain theory). We wrote $\mathbb{T}(M)$ for this domain and called it the trajectory domain of M .

This left open the issue of extending this to a category-theoretic treatment—which is the purpose of this chapter. But why is this important? The short answer is: only then do we have a complete description of BTSs as a computational model. This is necessary to show that the trajectory domain semantics is ‘compositional’ (as we’d expect of a semantics) and to understand the relations to other computational models. But it also is needed for a structural understanding of the class of BTSs: for example, to see whether a BTS suggested as a model for the safety verification of a reactive system is equivalent to another (simpler) one or to analyze it into subsystems. Let’s explain.

Computational models as categories A lesson from Winskel and Nielsen (1995) is that computational models (like LTSs) are fruitfully regarded as categories: the objects are instances of the computational model (i.e., any LTS) and the morphisms are simulations between instances of the model (one LTS being simulated in another).² For example, consider the following two LTSs: A on the left and B on the right.



We can simulate A in B by mapping $s_n \mapsto r$ and $s' \mapsto r'$ (and $\alpha \mapsto \alpha, \beta \mapsto \beta$): then any transition in A is simulated by a transition in B . (For the precise definition see section 3.2.1.) Thus, we can form the category LTS in which we can not just talk about LTSs (objects) but also about their relationships given by simulations (morphisms).

²Also see Sassone, Nielsen, and Winskel (1996).

Somewhat more liberally, we may also think of the category of ω ALG of ω -algebraic domains with Scott-continuous functions (formally defined in section 3.2.2) as a computational model.³ A domain D can be regarded as the data type of the possible (interim) outputs of a (type of) computational process.⁴ For example, such a process may be that of computing an increasingly precise binary representation of a real number x in the unit interval $[0, 1]$, say $x := \sqrt{2}/2 = 0.7071\dots$, so the interim outputs are 1, 10, 101, \dots ⁵ Then D may be taken as the set of all finite or finite binary sequences ordered by extension. A morphism, i.e., Scott-continuous function $f : D \rightarrow E$ maps D -outputs to E -outputs in a computational way: to obtain a finite approximation to the output $f(x)$ we only need a finite approximation to the input x .

Two advantages of this view of a computational model as a category are the following (Winskel and Nielsen 1995). First, constructions within the computational model (e.g., forming products of LTSs or domains) can now be characterized category-theoretically: i.e., purely ‘structurally’ without reference to the notational details of the computational model. Second, one can compare computational models even if they are stated in very different terms: If \mathbf{C} and \mathbf{D} are categories representing two computational models, a functor $F : \mathbf{C} \rightarrow \mathbf{D}$ turns an instance A of \mathbf{C} into an instance $F(A)$ of \mathbf{D} , and it turns a simulation $f : A \rightarrow B$ in \mathbf{C} into a simulation $F(f) : F(A) \rightarrow F(B)$ in \mathbf{D} . Most importantly, we can also formally reconstruct the idea that model \mathbf{D} is more *abstract* than (i.e., can be embedded into) model \mathbf{C} : we also have functor $G : \mathbf{D} \rightarrow \mathbf{C}$ in the other direction such that, roughly, if we start with B in \mathbf{D} and build $G(B)$ to go to \mathbf{C} and then build $F(G(B))$ to go back to \mathbf{D} , then we’re back to where we started. Formally, F and G form a (co-) reflective adjunction (as defined in section 3.2.3).

Inspired by this, our categorical treatment of BTSs and their trajectory domains establishes the following four results.

Result 1 We define the category **BTS** of BTSs where the morphisms are also based on the notion of simulation (as for LTSs). We show that BTSs do indeed structurally extend LTSs: We have the forgetful functor $G : \mathbf{BTS} \rightarrow \mathbf{LTS}$ that assigns each BTS to its underlying LTS, and this is part of a coreflective adjunction. So, the computational model **LTS** can indeed be embedded into **BTS**. We also show that we can ‘systematically’ ignore the ‘pathological’ non-approximable behavior in a BTS (i.e., infinite behavior that cannot be represented as limit of finite behavior). In categorical terms, the inclusion from the category \mathbf{BTS}_a of

³This is not explicitly mentioned by Winskel and Nielsen (1995), but see their section 8 on event structures (in particular, the domain of configurations).

⁴See, e.g., Scott (1970), Abramsky and Jung (1994) or Stoltenberg-Hansen, Lindström, and Griffor (1994, esp. the preface).

⁵The output starts with 1 since x is in the upper half of $[0, 1]$, i.e., $x \in [\frac{1}{2}, 1]$. It continues with 0 since x is in the lower half of $[\frac{1}{2}, 1]$, i.e., $x \in [\frac{1}{2}, \frac{3}{4}]$. It then continues with 1 since x is in the upper half of $[\frac{1}{2}, \frac{3}{4}]$, etc.

approximable BTSs (where every behavior is approximable) to the category \mathbf{BTS} has a right adjoint (and hence forms a reflective adjunction).

Result 2 We show that the trajectory domain construction \mathbb{T} is functorial: As mentioned, for a countable BTS M , the trajectory domain $\mathbb{T}(M)$ is an ω -algebraic domain. Here we show that \mathbb{T} also maps simulations between systems to Scott-continuous functions between their trajectory domains. In categorical terms, writing $\omega\mathbf{BTS}$ for the full subcategory of \mathbf{BTS} consisting of countable BTS, $\mathbb{T} : \omega\mathbf{BTS} \rightarrow \omega\mathbf{ALG}$ is a functor. Thus, the denotational semantics provided by the trajectory domains is ‘compositional’: it preserves the fundamental simulation relations between BTSs. A corollary is that equivalence in the operational semantics (i.e., isomorphism between LTSs) implies equivalence in the denotational semantics (i.e., isomorphism of trajectory domains). We further discuss this point in the open questions.

Result 3 Thus, we may ask whether the computational model $\omega\mathbf{ALG}$ is an abstraction of the computational model $\omega\mathbf{BTS}$, obtained through the trajectory domain functor.

We tackle this question for the mildly restricted subcategory $\omega\mathbf{BTS}_a^s$ of $\omega\mathbf{BTS}$: First, as justified before, we restrict us to approximable BTSs (hence the a). Second, instead of the general *partial* simulations where transitions may be simulated by inaction, we restrict us to *synchronous* simulations where transitions are always simulated by transitions (hence the s). We’re also led to a mild restriction on $\omega\mathbf{ALG}$: First, the Scott-continuous functions between trajectory domains that come from simulations are always compactness preserving (they map finite behavior to finite behavior). Second, the distinguishedness of the initial state of the system is, in some cases, reflected by the distinguishedness of an element of the trajectory domain (namely the behavior ending in the initial state). This leads us to consider the category \mathbf{iALG} whose objects are pairs (D, c) of an ω -algebraic domain D with a distinguished compact element c and whose morphisms are Scott-continuous functions preserving compactness and the distinguished element.

With these details out of the way, we show that there is indeed an adjunction

$$\omega\mathbf{BTS}_a^s \begin{array}{c} \xrightarrow{\text{Abs}} \\ \perp \\ \xleftarrow{\text{Emb}} \end{array} \mathbf{iALG}$$

which we obtain as a composition of three reflective adjunctions. Thus, we can indeed think of the computational model \mathbf{iALG} as an abstraction of the computational model $\omega\mathbf{BTS}_a^s$.

Result 4 Fourth, while speculating on how this adjunction may be extended to partial simulations, we make the surprising observation that LTSs and their

trajectory domains provide an interpretation of relevance logic. The importance of this is that relevance logic is often criticized for only having a formal but not a ‘concrete’ semantics.

Related work Much of the related work that we’ve already discussed in chapter 2 is also done at a categorical level: For example, the work on the correspondence of operational and denotational semantics for programming languages (Cardone 2021; Ong 1995) or the connections between concurrent computation and domain theory (Bracho and Droste 1994; Winskel and Nielsen 1995). As mentioned, here we consider denotational semantics for LTSs directly (without recourse to a programming language), and BTSs may be viewed as a generalization of various LTS-based models of concurrency. Also, our focus here is not on providing a categorical equivalence between ‘system-based’ categories and ‘domain-based’ categories (cf. Bracho and Droste 1994). Rather, we focus on reflective adjunctions (which, as discussed above, still have a strong computational interpretation) with the aim of establishing connections to categories of domains that are close to the standard ones of domain theory.

Outline The chapter is structured as follows. In section 3.2, we make sure that this chapter is self-contained: we provide the relevant background on labeled transition systems and domain theory, and we summarize the previous chapter.

In section 3.3, we define the category BTS of behavioral transition systems and show the adjunctions $\text{LTS} \rightleftharpoons \text{BTS}$ and $\text{BTS} \rightleftharpoons \text{BTS}_a$. In section 3.4, we show that the trajectory domain construction is a functor $\text{T} : \omega\text{BTS} \rightarrow \omega\text{ALG}$.

In section 3.5, we develop the adjunction $\omega\text{BTS}_a^s \rightleftharpoons \text{iALG}$. And in section 3.6, we speculate on possible extensions of the adjunction and sketch the interpretation of relevance logic. In section 3.7, we conclude with some open questions. A summary of the categories and their established connections is given in figure 3.1.

3.2 Background

We provide the relevant background on labeled transition systems (section 3.2.1), domain theory (section 3.2.2), and category theory (section 3.2.3). Then we summarize the relevant parts from the previous chapter (section 3.2.4).

3.2.1 Category of labeled transition systems

In the previous chapter, we’ve already recalled the notion of a labeled transition system (LTS) following the handbook article of Winskel and Nielsen (1995). In this chapter, we continue following this article and use the same standard notion for sequences: if σ is a finite or infinite sequence, $|\sigma| \leq \omega$ is its length, $\sigma \upharpoonright n$ is the restriction to its first n elements, and \preceq denotes sequence extension.

3.2.1. DEFINITION. A *labeled transition system (LTS)* A is a structure (S, i, L, \rightarrow) where S is a set of *states* with *initial state* i , L is a set of *labels*, and $\rightarrow \subseteq S \times L \times S$ is the *transition relation*. We write $s \xrightarrow{\alpha} s'$ for $(s, \alpha, s') \in \rightarrow$. Given an LTS A , we use S_A , i_A , L_A , and \rightarrow_A to refer to its set of states, initial state, set of labels, and transition relation, respectively. We call A *countable* if both S and L are countable sets. An *A-trajectory* is a sequence

$$t = (s_0, \alpha_0, s'_0), (s_1, \alpha_1, s'_1), \dots, (s_n, \alpha_n, s'_n), \dots$$

of elements of \rightarrow such that $s'_i = s_{i+1}$. We then write $s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots$. If t is nonempty, we call s_0 the *starting state* of t and, if t also is finite, we call the s' of the last entry the *ending* or *last state* of t , which we refer to by ‘ $\text{last}(t)$ ’.

A natural notion of morphism between LTSs is given by the idea of a *simulation*: A simulation of an LTS A in the LTS B (or an interpretation of A in B) is a way to map the states and labels of A to states and labels of B such that an A -transition is mimicked by a B -transition under this mapping. There at least two ways to understand ‘mimicked’. The most general way is that of a *partial simulation*: an A -transition either is mapped to a B -transition or is ignored and hence interpreted as ‘inaction’. A more specific way is that of a *synchronous simulation*: here we don’t allow the ‘inaction’ interpretation, whence every action in A is interpreted by an action in B . Thus, the original LTS A and the host LTS B (in which A is simulated) run ‘in sync’. Formally, this is spelled out as follows. (For more background, see Winskel and Nielsen (1995).)

3.2.2. DEFINITION. Let $A = (S_A, i_A, L_A, \rightarrow_A)$ and $B = (S_B, i_B, L_B, \rightarrow_B)$ be two LTSs. An *LTS-morphism* $f : A \rightarrow B$ is a pair (σ, λ) where $\sigma : S_A \rightarrow S_B$ is a total function and $\lambda : L_A \rightarrow L_B$ is a partial function such that

1. $\sigma(i_A) = i_B$
2. if $s \xrightarrow{\alpha}_A s'$, then $\sigma(s) \xrightarrow{\lambda(\alpha)}_B \sigma(s')$ if $\lambda(\alpha)$ is defined, and otherwise $\sigma(s) = \sigma(s')$.

If f is an LTS-morphism, we write $f = (\sigma_f, \lambda_f)$. We call f *synchronous* if λ_f is total.

3.2.3. DEFINITION. Labeled transition systems together with their morphisms form the category **LTS**. The identity morphism id_A is $(\text{id}_{S_A}, \text{id}_{L_A})$ (where id_X denotes the identity function on the set X). Morphism composition is pairwise function composition: $g \circ f = (\sigma_g \circ \sigma_f, \lambda_g \circ \lambda_f)$.⁶

Note that an LTS-morphism $f : A \rightarrow B$ sends A -trajectories to B -trajectories: If t is an A -trajectory, it is of the following form

⁶Composition of two partial functions is defined by $\lambda_g \circ \lambda_f(\alpha) := \lambda_g(\lambda_f(\alpha))$ if both $\lambda_f(\alpha)$ and $\lambda_g(\lambda_f(\alpha))$ are defined, and otherwise $\lambda_g \circ \lambda_f(\alpha)$ is undefined.

$$\begin{aligned}
t &= t(0) \quad t(1) \quad t(2) \quad \dots \\
&= s_0 \xrightarrow{\alpha_0} s'_0 \quad s_1 \xrightarrow{\alpha_1} s'_1 \quad s_2 \xrightarrow{\alpha_2} s'_2 \quad \dots
\end{aligned}$$

with $s'_i = s_{i+1}$. For each $t(n)$, we have $f(t(n)) := \sigma(s_n) \xrightarrow{\alpha_n} \sigma(s'_n)$ if $\lambda(\alpha_n)$ is defined and otherwise $f(t(n)) := (\sigma(s_n), \sigma(s'_n))$ is a pair of two identical elements, which we call an *idle pair*.⁷ We write $f^*(t)$ for the sequence $f(t(0))f(t(1))f(t(2))\dots$ and we write $f(t)$ for the B -trajectory obtained from $f^*(t)$ after removing all idle pairs.

Here are some basic facts (which we often use without explicit reference).

3.2.4. LEMMA. *Let $f : A \rightarrow B$ and $g : B \rightarrow C$ be LTS-morphisms, and let t and t' be A -trajectories. Then*

1. *If $t \preceq t'$, then $f(t) \preceq f(t')$.*
2. *$|f(t)| \leq |t|$.*
3. *For $n \geq 0$, we have $f(t \upharpoonright n) \preceq f(t) \upharpoonright n$.*
4. *For all $n \geq 0$, there is $m \geq 0$ such that $f(t) \upharpoonright n = f(t \upharpoonright m)$. In words: an initial segment of $f(t)$ is determined by an initial segment of t .*
5. *If f is synchronous, then, for all $n \geq 0$, $f(t) \upharpoonright n = f(t \upharpoonright n)$.*
6. *$g(f(t)) = g \circ f(t)$. In words: applying g to the B -trajectory $f(t)$ is the same as applying $g \circ f$ to the A -trajectory t .*

PROOF. Ad (1). If $t \preceq t'$, then $f^*(t) \preceq f^*(t')$, so $f(t) \preceq f(t')$.

Ad (2). Since idle pairs are only deleted but never added, we have $|f(t)| \leq |f(t(0))f(t(1))\dots| = |t|$.

Ad (3). We have $t \upharpoonright n \preceq t$, so, by (1), $f(t \upharpoonright n) \preceq f(t)$. Since, by (2), $|f(t \upharpoonright n)| \leq |t \upharpoonright n| \leq n$, we have $f(t \upharpoonright n) \preceq f(t) \upharpoonright n$.

Ad (4). By induction on n : If $n = 0$, we choose $m := 0$. For $n + 1$, if $f(t)(n)$ is not defined, then $f(t) \upharpoonright n + 1 = f(t) \upharpoonright n$ and the claim follows by induction hypothesis. So assume $f(t)(n)$ is defined. So $f(t) \upharpoonright n + 1 = (f(t) \upharpoonright n)f(t)(n)$. By induction hypothesis, let m_n be such that $f(t) \upharpoonright n = f(t \upharpoonright m_n)$. If, for all $m \geq m_n$, $t(m)$ is a transition whose label is not in the domain of λ , then none of these transitions will contribute to $f(t)$, whence $f(t) = f(t \upharpoonright m_n) = f(t) \upharpoonright n$, so $f(t)(n)$

⁷ The name alludes to the concept of an idle transition (Winskel and Nielsen 1995): one fixes a symbol $*$ (which no LTS is allowed to use as a label) which is interpreted as the ‘do nothing action’. So every LTS can be extended by adding all transitions of the form $s \xrightarrow{*} s$ which are called idle transitions. Then partial simulations can be rephrased as mapping (proper) transitions to (proper) transitions if defined or to idle transitions if undefined (and the latter are essentially the idle pairs above). Albeit elegant, we don’t use this to keep notation minimal (but we’ll encounter this idea again in section 3.6).

wouldn't be defined. So let $m \geq m_n$ be minimal such that the label of $t(m)$ is in the domain of λ . Then $f(t \upharpoonright m) = f(t \upharpoonright m_n)f(t(m)) = (f(t) \upharpoonright n)f(t)(n) = f(t) \upharpoonright n+1$, as needed.

Ad (5). If f is synchronous, each $\lambda_f(\alpha)$ is defined, so no $f(t(n))$ is idle, so $f(t \upharpoonright n) = f(t) \upharpoonright n$.

Ad (6). If $f(t(n)) = (s, s)$ is an idle pair, define $g(f(t(n))) := (\sigma_g(s), \sigma_g(s))$, and accordingly write

$$\begin{array}{rcccccc} t & = & t(0) & t(1) & t(2) & t(3) & \dots \\ f^*(t) & = & f(t(0)) & f(t(1)) & f(t(2)) & f(t(3)) & \dots \\ g^*(f^*(t)) & := & g(f(t(0))) & g(f(t(1))) & g(f(t(2))) & g(f(t(3))) & \dots \end{array}$$

For each $n \geq 0$ with $t(n) = s \xrightarrow{\alpha} s'$ defined, we have the following equivalences:

$g(f(t(n)))$ is idle *iff* $f(t(n))$ is idle or it is a transition but g is not defined on it *iff* $\lambda_f(\alpha)$ is not defined or it is defined but $\lambda_g(\lambda_f(\alpha))$ is not defined *iff* $\lambda_{g \circ f}(\alpha)$ is not defined *iff* $g \circ f(t(n))$ is idle.

And if $g(f(t(n)))$ and, equivalently, $g \circ f(t(n))$ are not idle (i.e., are transitions), then

$$\begin{aligned} g(f(t(n))) &= g(\sigma_f(s) \xrightarrow{\lambda_f(\alpha)} \sigma_f(s')) = \sigma_g(\sigma_f(s)) \xrightarrow{\sigma_g(\lambda_f(\alpha))} \sigma_g(\sigma_f(s')) \\ &= \sigma_{g \circ f}(s) \xrightarrow{\sigma_{g \circ f}(\alpha)} \sigma_{g \circ f}(s') = g \circ f(t(n)). \end{aligned}$$

Hence $g^*(f^*(t)) = (g \circ f)^*(t)$, whence $g(f(t)) = g \circ f(t)$. \square

Thus, while trajectory length—i.e., ‘computation time’—may get shorter along a partial simulation, it remains the same along a synchronous simulation (which, again, explains the name).

3.2.2 Domain theory

We recall some basic domain theory. A standard reference is Abramsky and Jung (1994). A partial order (D, \leq) is *directed complete* (in short, a *dcpo*) if any directed subset $A \subseteq D$ has a least upper bound $\bigvee A$ (also called supremum). (A is *directed* if A is nonempty and any two elements of A have an upper bound in A .) An element c of a dcpo D is *compact* if, for all directed subsets $A \subseteq D$, if $\bigvee A \geq c$, there is $a \in A$ with $a \geq c$. The set of compact elements of D is written $K(D)$. A dcpo D is *algebraic* if, for all $x \in D$, the set $\{c \in K(D) : c \leq x\}$ is directed and has supremum x . Finally, an ω -algebraic dcpo is an algebraic dcpo where $K(D)$ is countable.

A function $f : D \rightarrow E$ between dcpos is *Scott-continuous* if it is monotone and preserves all directed suprema, i.e., if $A \subseteq D$ is directed, then $f(\bigvee A) = \bigvee f(A)$.

We write ωALG for the category of ω -algebraic dcpos with Scott-continuous functions.

A useful fact to establish continuity is the following.

3.2.5. LEMMA. *Let $f : D \rightarrow E$ be a monotone function between two ω -algebraic domains. Assume that for every ω -chain $C \subseteq K(D)$ we have $f(\bigvee C) \leq \bigvee f(C)$. Then f is continuous.*

PROOF. Let $A \subseteq D$ be directed and show $f(\bigvee A) = \bigvee f(A)$. Since f is monotone, we have \geq , and for \leq we show that, if $A' \subseteq K(D)$ is directed, then $f(\bigvee A') \leq \bigvee f(A')$: Indeed, A' is directed and countable, so it has a cofinal chain $C = a_0 \leq a_1 \leq \dots$, so $\bigvee A' = \bigvee C$, whence, by assumption, $f(\bigvee A') = f(\bigvee C) \leq \bigvee f(C) \leq \bigvee f(A')$. Now take $A' := \{x \in K(D) : \exists a \in A. x \leq a\}$: by algebraicity, A' is still directed and $\bigvee A' = \bigvee A$, so $f(\bigvee A) = f(\bigvee A') \leq \bigvee f(A') \leq \bigvee f(A)$. \square

We'll also use the following two facts on reconstructing Scott-continuous functions between algebraic domains from monotone functions between their compact elements (i.e., their bases). (For the more general theory on reducing domains to bases see Abramsky and Jung (1994, sec. 2.2.6).)

3.2.6. LEMMA. *Let D and E be algebraic domains and $f : K(D) \rightarrow K(E)$ an order-isomorphism. Then $\hat{f} : D \rightarrow E$ defined by*

$$\hat{f}(x) := \bigvee \{f(c) : x \geq c \in K(D)\}$$

is a well-defined order isomorphism extending f .

PROOF. Well-defined: Since D is algebraic, $\{c \in K(D) : x \geq c\}$ is directed, so, since f is monotone, $\{f(c) : x \geq c \in K(D)\}$ is a directed subset of E and hence has a least upper bound.

Monotone: If $x \leq y$, then $\{f(c) : x \geq c \in K(D)\} \subseteq \{f(c) : y \geq c \in K(D)\}$, so $\hat{f}(x) = \bigvee \{f(c) : x \geq c \in K(D)\} \leq \bigvee \{f(c) : y \geq c \in K(D)\} = \hat{f}(y)$.

Surjective: If $y \in E$, then $B := \{d \in K(E) : y \geq d\}$ is directed with $\bigvee B = y$. Since f is an order-isomorphism, $A := f^{-1}(B)$ is directed in D . Let $x := \bigvee A$. We show $\hat{f}(x) = y$, i.e., $\bigvee \{f(c) : x \geq c \in K(D)\} = \bigvee B$. Concerning \leq , given $z = f(c)$ for some $\bigvee A = x \geq c \in K(D)$, we have, since c is compact, $c \leq a$ for some $a \in A = f^{-1}(B)$, so $z = f(c) \leq f(a) \in B$, so $z \leq \bigvee B$. Concerning \geq , given $d \in B \subseteq K(E)$, note that $c := f^{-1}(d)$ is in A since $f(c) = d \in B$. So $d \in \{f(c) : x \geq c \in K(D)\}$, whence $d \leq \bigvee \{f(c) : x \geq c \in K(D)\}$.

Order-respecting: Let $x, y \in D$ with $\hat{f}(x) \leq \hat{f}(y)$ and show $x \leq y$. It suffices to show, for $c \in K(D)$, that $c \leq x$ implies $c \leq y$ (then $x = \bigvee \{c \in K(D) : c \leq x\} \leq \bigvee \{c \in K(D) : c \leq y\} = y$). If $c \leq x$, then $f(c) \leq \bigvee \{f(c) : x \geq c \in K(D)\} \leq \bigvee \{f(c) : y \geq c \in K(D)\}$. Since $f(c)$ is compact in E , there is $y \geq c' \in K(D)$ with $f(c) \leq f(c')$, so, since f is an order-isomorphism, $c \leq c' \leq y$, as needed.

Extension: If $x \in D$ is compact, we have $\hat{f}(x) = \bigvee f(\{c \in K(D) : c \leq x\}) = f(x)$ since $f(x)$ is defined and, by monotonicity, an upper bound of $f(\{c \in K(D) : c \leq x\})$, and, since $f(x)$ is in this set, it also is a least upper bound. \square

3.2.7. LEMMA. *Let $f, g : D \rightarrow E$ be Scott-continuous functions between algebraic domains. If f and g agree on compact elements, then $f = g$.*

PROOF. For $x \in D$, we have

$$\begin{aligned} f(x) &= f\left(\bigvee\{c \in K(D) : c \leq x\}\right) = \bigvee f(\{c \in K(D) : c \leq x\}) \\ &= \bigvee g(\{c \in K(D) : c \leq x\}) = g\left(\bigvee\{c \in K(D) : c \leq x\}\right) = g(x), \end{aligned}$$

as needed. \square

3.2.3 Category theory

We only use the basic concepts of a category, a functor, and an adjunction; as found in standard references like Leinster (2014) or the classic Mac Lane (1998). We follow the slightly more general terminology of Sassone, Nielsen, and Winskel (1996) and Winskel and Nielsen (1995) and call an adjunction *reflective* (resp., *co-reflective*) if the counit (resp., unit) is a natural isomorphism. This generalizes the usual terminology of a subcategory \mathbf{D} of a category \mathbf{C} being reflective (resp., co-reflective) if the inclusion $\mathbf{I} : \mathbf{D} \rightarrow \mathbf{C}$ has a left adjoint (resp., right adjoint).

3.2.4 Recap from the previous chapter

The previous chapter provides an extensive discussion and motivation of the notion of a behavioral transition system (BTS) and its axiomatization. As already mentioned, the main idea was to extend an LTS A by a set T of ‘globally possible’ trajectories and an equivalence relation \equiv on T indicating when two trajectories are instances of the same (type of) behavior. To support this interpretation, the resulting structures (A, T, \equiv) should satisfy five axioms:

3.2.8. DEFINITION. A *behavioral transition system (BTS)* is a structure $M = (A, T, \equiv)$ where A is an LTS, T is a set of A -trajectories, and \equiv is an equivalence relation on T such that the following holds. (For finite $t, t' \in T$, $[t] \leq_M [t'] :\Leftrightarrow \forall t_0 \in [t] \exists t_1 \in [t'] : t_0 \preceq t_1$; we just write \leq if clear from context.)

1. For all $t \in T$, if t' is a nonempty finite initial segment of t , then $t' \in T$.

2. For all infinite A -trajectories t , if $0 < n_0 < n_1 < \dots$ with $t \upharpoonright n_i \in T$ and $[t \upharpoonright n_i] \leq [t \upharpoonright n_{i+1}]$ (for all $i \geq 0$), then $t \in T$.
3. For all $t, t' \in T$ with $t \equiv t'$, if t is empty, then t' is empty, and if t is finite, then t' is finite.
4. For all infinite $t, t' \in T$, if $t \equiv t'$, there is $i, j \geq 1$ such that, for all $n \geq 0$, $t \upharpoonright i + n \equiv t' \upharpoonright j + n$.
5. For all infinite $t \in T$ and $0 < n_0 < n_1 < \dots$ and $0 < m_0 < m_1 < \dots$, if $[t \upharpoonright n_0] \leq [t \upharpoonright n_1] \leq \dots$ and $[t \upharpoonright m_0] \leq [t \upharpoonright m_1] \leq \dots$, then, for all $i \geq 0$, there is $j \geq 0$ such that $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$.

We call M *countable* if A is countable.

A simple construction of a BTS from an LTS A is as follows: Let T be a set of A -trajectories that is closed under nonempty finite initial segments (axiom 1) and satisfies the following strengthening of axiom 2:

- (2)* For all infinite A -trajectories t , if $t \notin T$, then there is $n \geq 1$ such that $t \upharpoonright n \notin T$.

For example, T could be the set of all A -trajectories. Define \equiv as *extensional equivalence*: for $t, t' \in T$, define $t \equiv t'$ iff

- both t and t' are empty, or
- both t and t' are nonempty finite and $\text{last}(t) = \text{last}(t')$, or
- both t and t' are infinite and there are $i, j \geq 0$ such that, for all $n \geq 0$, $t(i + n) = t'(j + n)$.

Then, as shown in the previous chapter, $M := (A, T, \equiv)$ is a BTS.

In a BTS, we can define an ‘information containment’ order. (In the previous chapter, we’ve discussed various equivalent definitions.)

3.2.9. DEFINITION. Let $M = (A, T, \equiv)$ be a BTS. Let $t \in T$. We write $[t] := \{t' \in T : t' \equiv t\}$. An *approximation* to $[t]$ is a pair $(t^\dagger, (n_i)_{i \geq 0})$ with $t^\dagger \in [t]$ and (n_i) a strictly increasing sequence of positive integers such that the sequence $([t^\dagger \upharpoonright n_i])_i$ is \leq_M -increasing. We call t *approximable* if there is an approximation to $[t]$. For $t, t' \in T$ we define $[t] \sqsubseteq_M [t']$ iff

- (a) For all approximations $\tau = (t^\dagger, (n_i))$ to $[t]$ and $\tau' = (t^\ddagger, (m_j))$ to $[t']$, τ' *dominates* τ , i.e., $\forall i \geq 0 \exists j \geq 0 : [t^\dagger \upharpoonright n_i] \leq_M [t^\ddagger \upharpoonright m_j]$.
- (b) If $[t]$ is not approximable, then $[t']$ is not approximable.

As shown in the previous chapter, \sqsubseteq_M is a preorder on T/\equiv that coincides with \leq on equivalence classes of finite trajectories. We call it the *information containment* preorder of M and just write \sqsubseteq if M is clear from context.

This definition simplifies if $M = (A, T, \equiv)$ is *bisimulative*: i.e., for all nonempty finite $t, t' \in T$, if $t \equiv t'$ and $t_0 \in T$ extends t by one element, then there is a finite extension $t_1 \in T$ of t' such that $t_0 \equiv t_1$. (Equivalently, for all finite $t, t' \in T$, if $t \preceq t'$, then $[t] \leq [t']$.) As shown in the previous chapter, then, for all $t, t' \in T$, we have

$$[t] \sqsubseteq [t'] \Leftrightarrow \forall n \geq 0 \exists m \geq 0 : [t \upharpoonright n] \leq [t' \upharpoonright m].$$

The characterization theorem of the previous chapter shows that, for countable BTSs, the information containment preorder is, in a sense, unique and the partial order induced by the information containment preorder is an ω -algebraic domain.

3.2.10. DEFINITION. Let $M = (A, T, \equiv)$ be a countable BTS. Let $\mathsf{T}(M)$ be the partial order induced by $(T/\equiv, \sqsubseteq_M)$: its elements are equivalence classes $[[t]_{\equiv}]_{\sqsubseteq} := \{[t'] \in T/\equiv : [t'] \sqsubseteq_M [t], [t] \sqsubseteq_M [t']\}$, which often just denote $[[t]]$, and they are ordered by $[[t]] \sqsubseteq_M [[t']]$ iff $[t] \sqsubseteq_M [t']$. We often write $\overline{\sqsubseteq}_M = \sqsubseteq_M$. We call $\mathsf{T}(M)$ the *trajectory domain* of M .

Below we see, as this notation suggests, that the trajectory domain construction $\mathsf{T}(M)$ extends to a functor.

3.3 Category of behavioral transition systems

We define the category **BTS** of behavioral transition systems (section 3.3.1) and we prove some basic facts about its morphisms (section 3.3.2). Then we show that the category **LTS** can be ‘embedded’ into **BTS** (section 3.3.3) and that we can ignore non-approximable behavior (section 3.3.4).

3.3.1 Definition

The notion of morphism for LTSs extends naturally to BTSs by requiring that they additionally preserve the structure we care about: globally possible trajectories should be mapped to globally possible trajectories and information containment should be preserved.

3.3.1. DEFINITION. Let $M = (A_M, T_M, \equiv_M)$ and $N = (A_N, T_N, \equiv_N)$ be BTSs. A *BTS-morphism* $f : M \rightarrow N$ is an LTS-morphism $f : A_M \rightarrow A_N$ such that

1. For all $t \in T_M$, $f(t) \in T_N$.

2. For all $t, t' \in T_M$, if $[t] \sqsubseteq_M [t']$, then $[f(t)] \sqsubseteq_N [f(t')]$.

We call f *synchronous* if it is a synchronous LTS-morphism.

3.3.2. PROPOSITION. *We can form the category **BTS** whose objects are BTSs and whose morphisms are BTS-morphisms. The identity morphism is $\text{id}_M = (\text{id}_{S_M}, \text{id}_{L_M})$ and morphism composition is given by pairwise function composition.*

PROOF. Since LTS-morphisms already form a category, we need to check that (a) the identity LTS-morphism indeed satisfies the additional conditions (1) and (2) on BTS-morphism, and that (b) compositions of BTS-morphisms are again BTS-morphism. Now, (a) is immediate, so let $f : M \rightarrow N$ and $g : N \rightarrow K$ be BTS-morphism and show $g \circ f = (\sigma_g \circ \sigma_f, \lambda_g \circ \lambda_f)$ again satisfies conditions (1) and (2). Indeed, concerning (1), if $t \in T_M$, then $f(t) \in T_N$, so $g(f(t)) \in T_K$, so, using lemma 3.2.4, $g \circ f(t) = g(f(t)) \in T_K$. And concerning (2), if $t, t' \in T_M$ and $[t] \sqsubseteq_M [t']$, then $[f(t)] \sqsubseteq_N [f(t')]$, so $[g(f(t))] \sqsubseteq_K [g(f(t'))]$. So, since $g(f(t)) = g \circ f(t)$ and $g(f(t')) = g \circ f(t')$, we have $[g \circ f(t)] \sqsubseteq_K [g \circ f(t')]$. \square

We define various subcategories of **BTS** that we'll use below.

3.3.3. DEFINITION. Let BTS^s be the (wide) subcategory of **BTS** where morphisms are required to be synchronous. Let ωBTS be the full subcategory of **BTS** consisting of countable BTS. For further properties \mathfrak{p} of BTSs, let $\text{BTS}_{\mathfrak{p}}$ be the full subcategory of **BTS** whose objects have property \mathfrak{p} . Examples of \mathfrak{p} that we'll use are the following: If $M = (A, T, \equiv)$ is a BTS, we say M is

f *full* if T is the set of all A -trajectories.

e *extensional* if \equiv is extensional equivalence.

u *unlabeled* if the label set L_A is a singleton.

r *reflexive* if, for all $s \in S_A$ and $\alpha \in L_A$, $s \xrightarrow{\alpha} s$.

a *approximable* if every $t \in T$ is approximable.

y *antisymmetric* if \leq_M is antisymmetric (i.e., a partial order).⁸

Thus, for example, $\omega\text{BTS}_{\mathfrak{a}}$ is the full category of **BTS** consisting of countable and approximable BTSs. And $\omega\text{BTS}_{\text{fey}}^s$ is the full subcategory of BTS^s consisting of countable, full, extensional, and antisymmetric BTSs. More generally, the naming pattern is this: Categories are denoted by three upper case, sans serif letters. The countability restriction on objects is so prominent to deserve a place at the front

⁸Since the letters **a** (as in *antisymmetric*) and **s** (as in *antisymmetric*) are already taken, the next best mnemonic seems to be the letter **y** which appears rather idiosyncratically in 'antisymmetric'.

(i.e., a prefixed ω). Restrictions on morphisms are noted as suffixed superscripts. And restrictions on objects (other than the countability restriction) are noted as suffixed subscripts. If there are several properties, we don't need any notation to separate them since a single letter stands for a unique property.

In the previous chapter we've established $\text{f\&e} \Rightarrow \text{bisimulative} \Rightarrow \text{a}$.

3.3.2 Basic properties

We show two basic properties about BTS-morphisms: First, that their definition simplifies considerably for various subcategories of BTS. And second, that they preserve approximability.

3.3.4. PROPOSITION. *Let $M = (A_M, T_M, \equiv_M)$ and $N = (A_N, T_N, \equiv_N)$ be BTSs and $f : A_M \rightarrow A_N$ an LTS-morphism. Then:*

1. *Assume M is approximable and N is bisimulative. Then f is a BTS-morphism iff*

(a) *for all $t \in T_M$, $f(t) \in T_N$, and*

(b) *for all nonempty finite $t, t' \in T_M$, if $[t] \leq_M [t']$, then $[f(t)] \leq_N [f(t')]$.*

2. *In (1), clause (b) is implied by*

(c) *for all nonempty finite $t, t' \in T_M$, if $t \equiv_M t'$, then $f(t) \equiv_N f(t')$.*

Moreover, if N additionally is antisymmetric, then (b) also implies (c).

3. *If M and N are full and extensional, then f already is a BTS-morphism.*

PROOF. Ad (1). (\Rightarrow) If f is a BTS-morphism, it has property (a) by definition. Concerning (b), let $t, t' \in T_M$ be nonempty finite with $[t] \leq_M [t']$. Since \sqsubseteq coincides with \leq on finite trajectories, we have $[t] \sqsubseteq_M [t']$, whence, since f is a BTS-morphism, $[f(t)] \sqsubseteq_N [f(t')]$, whence, since $f(t)$ and $f(t')$ are finite, $[f(t)] \leq_N [f(t')]$.

(\Leftarrow) Assume f satisfies properties (a) and (b). By property (a), clause (1) of being a BTS-morphism is satisfied. For clause (2), let $t, t' \in T_M$ with $[t] \sqsubseteq_M [t']$, and show $[f(t)] \sqsubseteq_N [f(t')]$. As noted in section 3.2.4, we have, since N is bisimulative, $[f(t)] \sqsubseteq_N [f(t')]$ iff $\forall n \exists m : [f(t) \upharpoonright n] \leq_N [f(t') \upharpoonright m]$. So let $n \geq 0$ and find $m \geq 0$ such that $[f(t) \upharpoonright n] \leq_N [f(t') \upharpoonright m]$.

If t is empty, then we can choose $m := 0$ since then $[f(t) \upharpoonright n] = [\epsilon] \leq_N [\epsilon] = [f(t') \upharpoonright m]$. So let t be nonempty.

Since M is approximable, let $(t^\dagger, (n_i))$ and $(t^\dagger, (m_j))$ be approximations to $[t]$ and $[t']$, respectively. Also let $k \geq 0$ be such that $f(t) \upharpoonright n = f(t \upharpoonright k)$. Let $i \geq 0$ be big enough such that $n_i > k \geq 0$. Since N is bisimulative and $f(t \upharpoonright k) \preceq f(t \upharpoonright n_i)$,

we have $[f(t) \upharpoonright n] = [f(t \upharpoonright k)] \leq_N [f(t \upharpoonright n_i)]$. Since $[t] \sqsubseteq_M [t']$, there is $j \geq 0$ such that $[t \upharpoonright n_i] \leq_M [t' \upharpoonright m_j]$. We claim that we can choose $m := m_j$.

If $t' \upharpoonright m_j$ is empty, then also $t \upharpoonright n_i$ is empty (otherwise it cannot be extended to a trajectory equivalent to $t' \upharpoonright m_j$), so t is empty (otherwise, since $n_i > 0$, also $t \upharpoonright n_i$ is nonempty). Hence also $f(t)$ is empty, so $[f(t) \upharpoonright n] = [\epsilon] \leq_N [f(t') \upharpoonright m_j]$, as needed.

So assume $t' \upharpoonright m_j$ is nonempty. Since $n_i > 0$ and t is nonempty, also $t \upharpoonright n_i$ is nonempty. And $t \upharpoonright n_i$ and $t' \upharpoonright m_j$ are in T_M qua nonempty initial segments of the trajectories t and t' in T_M , respectively. Since $[t \upharpoonright n_i] \leq_M [t' \upharpoonright m_j]$, clause (b) implies

$$[f(t) \upharpoonright n] = [f(t \upharpoonright k)] \leq_N [f(t \upharpoonright n_i)] \leq_N [f(t' \upharpoonright m_j)] \leq_N [f(t') \upharpoonright m_j],$$

where the last step follows since $f(t' \upharpoonright m_j) \preceq f(t') \upharpoonright m_j$ and N is bisimulative.

Ad (2). First, we show, in the setting of (1), that (c) \Rightarrow (b).

Indeed, let $t, t' \in T_M$ be nonempty finite with $[t] \leq_M [t']$. So t can be extended to $t_1 \in T_M$ with $t_1 \equiv_M t'$. In particular, t_1 also is nonempty finite. So, by (c), $f(t) \preceq f(t_1) \equiv_N f(t')$. Since N is bisimulative, $[f(t)] \leq_N [f(t_1)] = [f(t')]$.

Next, assume that N additionally is antisymmetric and show (b) \Rightarrow (c).

Indeed, let $t, t' \in T_M$ be nonempty finite with $t \equiv_M t'$. By reflexivity of \sqsubseteq , $[t] \sqsubseteq_M [t']$ and $[t'] \sqsubseteq_M [t]$, so, by (b), we have $[f(t)] \sqsubseteq_N [f(t')]$ and $[f(t')] \sqsubseteq_N [f(t)]$. Since \sqsubseteq_N coincides with \leq_N on finite trajectories and \leq_N is antisymmetric, we have $[f(t)] = [f(t')]$, so $f(t) \equiv_N f(t')$.

Ad (3). Let M and N be full and extensional. In particular, M is approximable and N is bisimulative. By (1) and (2), it suffices to show that clauses (a) and (c) are satisfied. Indeed, (a) is satisfied since N is full. For (c), let $t, t' \in T_M$ be nonempty finite with $t \equiv_M t'$. Since \equiv_M is extensional equivalence, $\text{last}(t) = \text{last}(t')$. Hence

$$\text{last}(f(t)) = \sigma_f(\text{last}(t)) = \sigma_f(\text{last}(t')) = \text{last}(f(t')),$$

so, since \equiv_N is extensional equivalence, $f(t) \equiv_N f(t')$. \square

3.3.5. PROPOSITION. *Let $f : M \rightarrow N$ be a BTS-morphism. If $[t]$ is approximable in M , then $[f(t)]$ is approximable in N .*

PROOF. By assumption, there is an approximation $(t^\dagger, (n_i))$ to $[t]$. It suffices to show that $f(t^\dagger)$ is approximable in N : Then, since $t \equiv t^\dagger$, we have, by reflexivity of \sqsubseteq , $[t] \sqsubseteq_M [t^\dagger]$, whence, since f preserves \sqsubseteq , $[f(t)] \sqsubseteq_N [f(t^\dagger)]$. By definition of \sqsubseteq , this implies that, if $[f(t)]$ is non-approximable, also $[f(t^\dagger)]$ is non-approximable. So if $[f(t^\dagger)]$ is approximable, also $[f(t)]$ is.

If $f(t^\dagger)$ is finite, it is approximable, so let it be infinite (so also t^\dagger is infinite). Hence $|f(t^\dagger \upharpoonright n_i)|$ grows unboundedly (otherwise there is m such that all transitions $t^\dagger(m')$ with $m' \geq m$ get mapped by f to undefined transitions, so $f(t^\dagger)$ is finite). Let $(n_{i_j})_{j \geq 0}$ be a subsequence such that $0 < |f(t^\dagger \upharpoonright n_{i_j})| < |f(t^\dagger \upharpoonright n_{i_{j+1}})|$.

Now, define $m_j := |f(t^\dagger \upharpoonright n_{i_j})|$. Note that $f(t^\dagger \upharpoonright n_{i_j}) = f(t^\dagger) \upharpoonright m_j$.⁹ Then $0 < m_0 < m_1 < \dots$ and, for any $j \geq 0$, we have $[t^\dagger \upharpoonright n_{i_j}] \leq_M [t^\dagger \upharpoonright n_{i_{j+1}}]$, so, since f preserves \sqsubseteq which coincides with \leq on finite trajectories, we have

$$[f(t^\dagger) \upharpoonright m_j] = [f(t^\dagger \upharpoonright n_{i_j})] \leq_N [f(t^\dagger \upharpoonright n_{i_{j+1}})] = [f(t^\dagger) \upharpoonright m_{j+1}].$$

Hence, $(f(t^\dagger), (m_j))$ is an approximation to $f(t^\dagger)$ in N . \square

3.3.3 Embedding labeled transition systems

We have the forgetful functor $\mathbf{G} : \mathbf{BTS} \rightarrow \mathbf{LTS}$ that maps a BTS $M = (A, T, \equiv)$ to the underlying LTS A and that maps a BTS-morphism $f : M \rightarrow N$ to $\mathbf{G}(f) := f : A_M \rightarrow A_N$. Conversely, we show that there also is an optimal way of turning an LTS into a BTS, i.e., the forgetful functor \mathbf{G} has a left adjoint \mathbf{F} :

$$\mathbf{LTS} \begin{array}{c} \xrightarrow{\mathbf{F}} \\ \perp \\ \xleftarrow{\mathbf{G}} \end{array} \mathbf{BTS}$$

and the unit of the adjunction is an isomorphism. Thus, the computational model LTS can be abstracted from (i.e., embedded into) the computational model BTS. Spelled out, this means the following.

3.3.6. PROPOSITION. *The forgetful functor $\mathbf{G} : \mathbf{BTS} \rightarrow \mathbf{LTS}$ is a right adjoint: For each B in LTS there is $\mathbf{F}(B)$ in BTS and an isomorphism $\eta_B : B \rightarrow \mathbf{G}(\mathbf{F}(B))$ such that, for every M in BTS and every $g : B \rightarrow \mathbf{G}(M)$, there is a unique morphism $f : \mathbf{F}(B) \rightarrow M$ with $\mathbf{G}(f) \circ \eta_B = g$.*

$$\begin{array}{ccc} B & \xrightarrow{\eta_B} & \mathbf{G}(\mathbf{F}(B)) \\ & \searrow g & \downarrow \mathbf{G}(f) \\ & & \mathbf{G}(M) \end{array}$$

PROOF. *Construction of $\mathbf{F}(B)$.* Define $\mathbf{F}(B) := (B, T_B, \equiv_B)$ with $T_B := \emptyset$ and $\equiv_B := \emptyset$. This is a BTS: B is an LTS, T_B is a set of B -trajectories, and \equiv_B is an equivalence relation on T_B , and it vacuously satisfies the axioms (1)–(5).

Construction of η_B . Define $\eta_B := \text{id}_B = (\text{id}_{S_B}, \text{id}_{L_B}) : B \rightarrow B = \mathbf{G}(\mathbf{F}(B))$. This, in particular, is an isomorphism in LTS.

Universality. Now, let $M = (A, T, \equiv)$ be in BTS and let $g : B \rightarrow \mathbf{G}(M)$ be a morphism, and find a unique $f : \mathbf{F}(B) \rightarrow M$ with $\mathbf{G}(f) \circ \eta_B = g$. Uniqueness is immediate: if f, f' are such morphisms, then $f = f' \circ \text{id}_B = \mathbf{G}(f') \circ \eta_B = g =$

⁹Proof: Since $t^\dagger \upharpoonright n_{i_j} \preceq t^\dagger$, we have $f(t^\dagger \upharpoonright n_{i_j}) \preceq f(t^\dagger)$, so $f(t^\dagger \upharpoonright n_{i_j}) = f(t^\dagger) \upharpoonright |f(t^\dagger \upharpoonright n_{i_j})| = f(t^\dagger) \upharpoonright m_j$.

$\mathbf{G}(f') \circ \eta_B = f' \circ \text{id}_B = f'$. For existence, it suffices to show that $f := g : F(B) \rightarrow M$ is a BTS-morphism (since it automatically has the property $\mathbf{G}(f) \circ \eta_B = g \circ \text{id}_B = g$). Indeed, it is an LTS-morphism $B \rightarrow \mathbf{G}(M)$ of the underlying LTSs and, since $T_B = \emptyset$, it vacuously satisfies the axioms (1)–(2) of BTS-morphisms. \square

3.3.4 Removing non-approximable behavior

We show that we can systematically ignore the ‘pathological’ non-approximable behavior: The operation \mathbf{A} of ‘removing’ non-approximable trajectories from a BTS yields an approximable BTS and is optimal in the sense of being right-adjoint to the inclusion

$$\text{BTS} \begin{array}{c} \xrightarrow{\mathbf{A}} \\ \xleftarrow{\mathbf{I}} \\ \text{BTS}_a \end{array}$$

Spelled out, this means the following.

3.3.7. THEOREM. *The inclusion $\mathbf{I} : \text{BTS}_a \rightarrow \text{BTS}$ is a left adjoint: For each M in BTS there is $\mathbf{A}(M)$ in BTS_a and $\epsilon_M : \mathbf{A}(M) \rightarrow M$ such that, for every N in BTS_a and every $f : N \rightarrow M$, there is a unique morphism $g : N \rightarrow \mathbf{A}(M)$ with $\epsilon_M \circ g = f$.*

$$\begin{array}{ccc} N & & \\ \downarrow g & \searrow f & \\ \mathbf{A}(M) & \xrightarrow{\epsilon_M} & M \end{array}$$

PROOF. *Construction of $\mathbf{A}(M)$.* Write $M = (A, T, \equiv)$. Define

$$T' := \{t \in T : [t] \text{ approximable in } M\}.$$

Define $\mathbf{A}(M) := (A, T', \equiv')$, where \equiv' is the restriction of the equivalence relation \equiv on T to the subset T' . So T' is a set of A -trajectories and \equiv' an equivalence relation on T' , so, to verify that $\mathbf{A}(M)$ is in BTS_a , we need to show that it satisfies axioms (1)–(5) and is approximable.

We signal notions in $\mathbf{A}(M)$ by an apostrophe (e.g., \leq' or $[t]'$). We first show four claims.

(C1). For finite $t, t' \in T'$, we have $[t]' \leq' [t']'$ iff $[t] \leq [t']$.

Proof: For finite $t \in T'$, we have, since finite trajectories are approximable and since finite trajectories can only be equivalent to finite trajectories, that $[t]' = \{t' \in T' : t' \equiv' t\} = \{t' \in T : t' \equiv t\} = [t]$. Hence, $[t]' \leq' [t']'$ iff $\forall t_0 \in [t]' \exists t_1 \in [t']' : t_0 \preceq t_1$ iff $\forall t_0 \in [t] \exists t_1 \in [t'] : t_0 \preceq t_1$ iff $[t] \leq [t']$.

(C2). If $(t^\dagger, (n_i))$ is an approximation to $[t]$ in M , it is also an approximation to $[t]'$ in $\mathbf{A}(M)$, and vice versa.

Proof: So the $n_i > 0$ are strictly increasing and $([t^\dagger \upharpoonright n_i])$ is \leq -increasing and $t^\dagger \in [t]$. In particular, $t^\dagger \in T$ also is approximable, so $t^\dagger \in T'$, whence also all $t^\dagger \upharpoonright n_i$ are in T' . Hence $t^\dagger \in [t]'$ and $([t^\dagger \upharpoonright n_i]')_i$ is, by (C1), \leq' -increasing. So $(t^\dagger, (n_i))$ is an approximation to $[t]'$ in $\mathbf{A}(M)$.

Conversely, if $(t^\dagger, (n_i))$ is an approximation to $[t]'$ in $\mathbf{A}(M)$, then (n_i) is strictly increasing and $([t^\dagger \upharpoonright n_i]')$ is \leq' -increasing and $t^\dagger \in [t]'$. So also $t^\dagger \in [t]$ and, by (C1), $([t^\dagger \upharpoonright n_i])$ is \leq -increasing, so $(t^\dagger, (n_i))$ is an approximation to $[t]$ in M .

(C3). In particular, each $t \in T'$ is approximable in $\mathbf{A}(M)$.

Proof: By definition of T' , $[t]$ has an approximation in M , which is, by (C2), an approximation in $\mathbf{A}(M)$, so t is approximable in $\mathbf{A}(M)$.

(C4). For $t, t' \in T'$, $[t] \sqsubseteq [t']$ iff $[t] \sqsubseteq [t']$.

Proof: (\Rightarrow) Assume $[t] \sqsubseteq [t']$, and show $[t] \sqsubseteq [t']$. Since $t \in T'$, $[t]$ is approximable in M , so the (b)-condition of \sqsubseteq is satisfied, and we need to show the (a)-condition. So let $(t^\dagger, (n_i))$ and $(t^\ddagger, (m_j))$ be approximations in M to $[t]$ and $[t']$, respectively, and let $i \geq 0$. By (C2), these also are approximations in $\mathbf{A}(M)$ to $[t]'$ and $[t']'$, respectively. Since $[t] \sqsubseteq [t']$, there is $j \geq 0$ such that $[t^\dagger \upharpoonright n_i] \leq [t^\ddagger \upharpoonright m_j]$, so, by (C1), $[t^\dagger \upharpoonright n_i] \leq [t^\ddagger \upharpoonright m_j]$, as needed. (\Leftarrow) Assume $[t] \sqsubseteq [t']$, and show $[t] \sqsubseteq [t']$. Since, by (C3), $[t]'$ is approximable in $\mathbf{A}(M)$, the (b)-condition of \sqsubseteq' is satisfied, and we need to show the (a)-condition. So let $(t^\dagger, (n_i))$ and $(t^\ddagger, (m_j))$ be approximations in $\mathbf{A}(M)$ to $[t]'$ and $[t']'$, respectively, and let $i \geq 0$. By (C2), these are also approximations in M to $[t]$ and $[t']$, respectively. Since $[t] \sqsubseteq [t']$, there is $j \geq 0$ such that $[t^\dagger \upharpoonright n_i] \leq [t^\ddagger \upharpoonright m_j]$, so, by (C1), $[t^\dagger \upharpoonright n_i] \leq [t^\ddagger \upharpoonright m_j]$.

Now, concerning axiom (1), assume $t' \preceq t \in T'$ with t' nonempty finite. So t' is approximable, and it is in T qua nonempty initial segment of $t \in T' \subseteq T$, whence $t' \in T'$.

Concerning axiom (2), assume t is an infinite A -trajectory and $0 < n_0 < n_1 < \dots$ with $t \upharpoonright n_i \in T'$ and $[t \upharpoonright n_i]' \leq [t \upharpoonright n_{i+1}]'$. Then $t \upharpoonright n_i \in T$ and, by (C1), $[t \upharpoonright n_i] \leq [t \upharpoonright n_{i+1}]$. This implies that $[t]$ is approximable in M , and, since M satisfies this axiom (2), $t \in T$. Hence $t \in T'$.

Concerning axiom (3), if $t, t' \in T'$ with $t \equiv' t'$, then $t, t' \in T$ with $t \equiv t'$, so, if t is empty, also t' is empty, and if t is finite, also t' is finite, as needed.

Concerning axiom (4), if $t, t' \in T'$ are infinite with $t \equiv' t'$, then $t, t' \in T$ are infinite with $t \equiv t'$, so there is $i, j \geq 1$ such that, for all $n \geq 0$, $t \upharpoonright i+n \equiv t' \upharpoonright j+n$, whence, since these trajectories are in T' (qua nonempty finite initial segments of trajectories in T'), $t \upharpoonright i+n \equiv t' \upharpoonright j+n$, as needed.

Concerning axiom (5), assume $t \in T'$ is infinite and (n_i) and (m_j) are strictly increasing with $([t \upharpoonright n_i]')_i$ and $([t \upharpoonright m_j]')_j$ \leq' -increasing. By (C1), $([t \upharpoonright n_i]')_i$ and $([t \upharpoonright m_j]')_j$ are \leq -increasing. So, for all $i \geq 0$, there is $j \geq 0$ with $[t \upharpoonright n_i] \leq [t \upharpoonright m_j]$, so, by (C1), $[t \upharpoonright n_i]' \leq [t \upharpoonright m_j]'$, as needed.

Finally, $\mathbf{A}(M)$ is approximable by (C3).

Construction of ϵ_M . Write $A = (S, i, L, \rightarrow)$ for the underlying LTS of M . Let $\epsilon_M := (\text{id}_S, \text{id}_L)$ be the identity LTS-morphism. To show that it is a BTS-morphism, we need to verify properties (1) and (2). Concerning (1), if $t \in T'$,

then $\epsilon_M(t) = t \in T$ since $T' \subseteq T$. Concerning (2), if $t, t' \in T'$ with $[t]' \sqsubseteq [t']'$, then, by (C4), $[\epsilon_M(t)] = [t] \sqsubseteq [t'] = [\epsilon_M(t')]$.

Universality. Now, let N be in BTS_a and $f : N \rightarrow M$ a BTS-morphism. We need to find a unique morphism $g : N \rightarrow \mathbf{A}(M)$ with $\epsilon_M \circ g = f$.

Uniqueness is immediate: if g, g' are such morphisms, we have, since ϵ_M is the identity LTS-morphism, $g = \epsilon_M \circ g = f = \epsilon_M \circ g' = g'$. For existence, we need to show that $g := f : N \rightarrow \mathbf{A}(M)$ is a BTS-morphism. It is an LTS-morphism $A_N \rightarrow A$, so we need to show that it satisfies properties (1) and (2) of a BTS-morphism.

Concerning (1), if $t \in T_N$, then, qua BTS-morphism $N \rightarrow M$, $f(t) \in T$. Since $[t]$ is approximable in N (since N is in BTS_a), $[f(t)]$ is, by proposition 3.3.5, approximable in M , so $f(t) \in T'$.

Concerning (2), assume $t, t' \in T_N$ with $[t] \sqsubseteq_N [t']$. Since f is a BTS-morphism $N \rightarrow M$, we have $[f(t)] \sqsubseteq [f(t')]$. Since $f(t), f(t') \in T'$, we have, by (C4), that $[f(t)]' \sqsubseteq [f(t')]'$, as needed. \square

We also note that the adjunction restricts to the countable case.

$$\omega\text{BTS}^s \begin{array}{c} \xrightarrow{A} \\ \Upsilon \\ \xleftarrow{I} \end{array} \omega\text{BTS}_a^s$$

Indeed, if M is in ωBTS^s , then the LTS A underlying M is countable, so, since A is also the LTS underlying $\mathbf{A}(M)$, also $\mathbf{A}(M)$ is countable and hence in ωBTS^s . Moreover, the morphism $\epsilon_M : \mathbf{A}(M) \rightarrow M$ is in ωBTS^s since it is synchronous (λ_{ϵ_M} is the identity).

3.4 Trajectory domain functor

We show that the trajectory domain construction is functorial: it naturally extends to a functor from the category ωBTS to the category ωALG of ω -algebraic domains with Scott-continuous functions. This will follow easily from the following proposition.

3.4.1. PROPOSITION. *Let M and N be in ωBTS and $f : M \rightarrow N$ a BTS-morphism. Then the function $\Upsilon(f) : \mathbb{T}(M) \rightarrow \mathbb{T}(N)$ given by $[[t]] \mapsto [[f(t)]]$ is well-defined and Scott-continuous.*

PROOF. Well-defined: Since $t \in T_M$, $f(t) \in T_N$, so $[[f(t)]] \in \mathbb{T}(N)$, and if $[[t]] = [[t']]$, then $[t] \sqsubseteq [t']$ and $[t'] \sqsubseteq [t]$, so, by clause (2) of BTS-morphisms, $[f(t)] \sqsubseteq [f(t')]$ and $[f(t')] \sqsubseteq [f(t)]$, so $[[f(t)]] = [[f(t')]]$. Similarly, we see that $\Upsilon(f)$ is monotone.

Thus, $\Upsilon(f)$ is a monotone function between the two ω -algebraic domains $\mathbb{T}(M)$ and $\mathbb{T}(N)$. So to show that it is continuous, it suffices, by lemma 3.2.5, to show that, for an ω -chain $C \subseteq K(\mathbb{T}(M))$, we have $\Upsilon(f)(\bigvee C) \sqsubseteq \bigvee \Upsilon(f)(C)$.

Now, C is of the form $[[t_0]] \sqsubseteq [[t_1]] \sqsubseteq \dots$ for $t_i \in T_M$, whence $[t_0] \leq [t_1] \leq \dots$. If C has a greatest element (i.e., ‘stagnates’ with some $[t_k]$), the claim is immediate by monotonicity. So we can assume without loss of generality that the chain is strictly increasing, doesn’t start with $[\epsilon]$, and all t_k are finite (if some t_k were infinite, it must be non-approximable since it is compact, so it would be maximal).

Let $t'_0 := t_0 \in [t_0]$ and $n_0 := |t'_0| > 0$. Then we can extend t'_0 to $t'_1 \in [t_1]$ and have $n_1 := |t'_1| > |t'_0| = n_0$. We continue and extend t'_1 to $t'_2 \in [t_2]$, etc., and define $t \in T$ by: $t(n) := t'_k(n)$ for some k with $|t'_k| > n$. Then $C = [[t \upharpoonright n_0]] \leq [[t \upharpoonright n_1]] \leq \dots$, so, by the characterization theorem (from the previous chapter), $\bigvee C = [[t]]$. Moreover, $\mathbb{T}(f)(C) = \{[[f(t \upharpoonright n_0)]], [[f(t \upharpoonright n_1)]], \dots\}$.

By monotonicity, $[f(t)]$ is a \sqsubseteq -upper bound of $\mathbb{T}(f)(C)$. So if $f(t)$ is finite, there is a big enough n_i such that $f(t) = f(t \upharpoonright n_i)$, so $[[f(t)]] \in \mathbb{T}(f)(C)$ is the least upper bound, i.e., $\mathbb{T}(f)(\bigvee C) = \mathbb{T}(f)([[t]]) = \bigvee \mathbb{T}(f)(C)$.

So assume $f(t)$ is infinite. For each n_i define $m_i := |f(t \upharpoonright n_i)|$. Then $f(t \upharpoonright n_i) = f(t) \upharpoonright m_i$ (since $f(t \upharpoonright n_i)$ is an initial segment of $f(t)$ of length m_i). Note that $m_i \leq m_{i+1}$ and $[f(t) \upharpoonright m_i] = [f(t \upharpoonright n_i)] \leq [f(t \upharpoonright n_{i+1})] = [f(t) \upharpoonright m_{i+1}]$. And the m_i grow unboundedly (if not, $f(t)$ would be finite). Pick a subsequence $(m_{i_j})_j$ that is strictly increasing with $m_{i_0} > 0$. Then we have, by the characterization theorem, that $[[f(t)]] = \bigvee_i [[f(t) \upharpoonright m_i]]$. And since $[f(t) \upharpoonright m_i] \in \mathbb{T}(f)(C)$, we have $\mathbb{T}(f)(\bigvee C) = \mathbb{T}(f)([[t]]) = [[f(t)]] = \bigvee [f(t) \upharpoonright m_i] \sqsubseteq \bigvee \mathbb{T}(f)(C)$, as needed. \square

3.4.2. THEOREM. *We have the trajectory domain functor $\mathbb{T} : \omega\text{BTS} \rightarrow \omega\text{ALG}$ which sends a BTS M to its trajectory domain $\mathbb{T}(M)$ and which sends a BTS-morphism $f : M \rightarrow N$ to the Scott-continuous function $\mathbb{T}(f) : \mathbb{T}(M) \rightarrow \mathbb{T}(N)$ defined by $[[t]] \mapsto [[f(t)]]$.*

PROOF. It remains to check that \mathbb{T} satisfies the compositionality conditions. Indeed, $\mathbb{T}(\text{id}_M)$ maps $[[t]]$ to $[[\text{id}_M(t)]] = [[t]]$ and hence is the identity on $\mathbb{T}(M)$. And if $f : M \rightarrow N$ and $g : N \rightarrow K$ are BTS-morphisms, then we have, for all $t \in T_M$, that $g \circ f(t) = g(f(t))$, hence

$$\begin{aligned} \mathbb{T}(g) \circ \mathbb{T}(f)([[t]]) &= \mathbb{T}(g)(\mathbb{T}(f)([[t]])) = \mathbb{T}(g)([[f(t)]]) \\ &= [[g(f(t))]] = [[g \circ f(t)]] = \mathbb{T}(g \circ f)([[t]]), \end{aligned}$$

so $\mathbb{T}(g \circ f) = \mathbb{T}(g) \circ \mathbb{T}(f)$. \square

Three comments: First, when restricting to approximable BTSs, any simulation between BTSs is turned by \mathbb{T} into a Scott-continuous function that preserves compactness (i.e., maps compact elements to compact elements): If $f : M \rightarrow N$ is in ωBTS_a and $[[t]] \in \mathbb{T}(M)$ is compact, then t is a finite trajectory, so $\mathbb{T}(f)([[t]]) = [[f(t)]] \in \mathbb{T}(N)$ is compact since $f(t)$ is finite.

Second, the fact that $\mathbb{T}(f)$ preserves compactness is, in a sense, the consequence of LTS-morphisms being ‘uniform’ or ‘context insensitive’: Whether a state s

or label α in A is mapped to a state s' or label α' in B has to be determined without reference to the context—i.e., trajectory—in which s and α occur. Thus, one might consider *generalized* BTS-morphisms that can be sensitive to context (but are insensitive to informationally equivalent trajectories) as Scott-continuous functions $\mathbb{T}(M) \rightarrow \mathbb{T}(N)$. Here, though, we stick to the standard definition.

Third, since \mathbb{T} is a functor, it maps isomorphisms to isomorphisms. So equivalence in operational semantics in the sense of isomorphism of countable BTSs implies equivalence in denotational semantics in the sense of isomorphism of the trajectory domains. In the context of the discussion of full abstraction (Cardone 2021; Ong 1995), this is the difficult direction in establishing the coincidence of operational and denotational semantics (since in that setup the denotational semantics usually is too rich). Here it is the other way round since the denotation abstracts away information as will become clear in the next section. We discuss this further in section 3.7.

3.5 Adjunction between systems and domains

As motivated in the introduction, the functor $\mathbb{T} : \omega\text{BTS} \rightarrow \omega\text{ALG}$ invites the question whether the computational model ωALG is an abstraction of the computational model ωBTS . In this section, we tackle this question—as explained in the introduction—by establishing an adjunction $\omega\text{BTS}_a^s \rightleftarrows \text{iALG}$ obtained as a composition of the following three reflective adjunctions:

$$\omega\text{BTS}_a^s \begin{array}{c} \xrightarrow{E} \\ \perp \\ \xleftarrow{I} \end{array} \omega\text{BTS}_{\text{fey}}^s \begin{array}{c} \xrightarrow{U} \\ \perp \\ \xleftarrow{I} \end{array} \omega\text{BTS}_{\text{feyur}}^s \begin{array}{c} \xrightarrow{T_i} \\ \perp \\ \xleftarrow{B} \end{array} \text{iALG}$$

The following three subsections establish these three adjunctions in turn (from left to right) and also formally define (and recall) the involved categories and functors.

Thus, we can indeed think of the computational model iALG as an abstraction of the computational model ωBTS_a^s .

3.5.1 Extensionalizing

Recall that if a BTS M is full and extensional, it in particular is approximable, so we have the inclusion $I : \omega\text{BTS}_{\text{fey}}^s \rightarrow \omega\text{BTS}_a^s$. (To recall, f stands for full, e for extensional, and y for antisymmetric.) In this subsection, we show that this is a right adjoint: i.e., there is an optimal way of rendering an M in ωBTS_a^s full, extensional, and antisymmetric.

$$\omega\text{BTS}_a^s \begin{array}{c} \xrightarrow{E} \\ \perp \\ \xleftarrow{I} \end{array} \omega\text{BTS}_{\text{fey}}^s$$

Spelled out, this means the following.

3.5.1. PROPOSITION. *The inclusion $\mathbb{1} : \omega\text{BTS}_{\text{fey}}^s \rightarrow \omega\text{BTS}_a^s$ is a right adjoint: For each M in ωBTS_a^s there is $\mathbf{E}(M)$ in $\omega\text{BTS}_{\text{fey}}^s$ and $\eta_M : M \rightarrow \mathbf{E}(M)$ such that, for every N in $\omega\text{BTS}_{\text{fey}}^s$ and every $g : M \rightarrow N$, there is a unique morphism $f : \mathbf{E}(M) \rightarrow N$ with $f \circ \eta_M = g$.*

$$\begin{array}{ccc} M & \xrightarrow{\eta_M} & \mathbf{E}(M) \\ & \searrow g & \downarrow \text{---} f \\ & & N \end{array}$$

PROOF. *Construction of $\mathbf{E}(M)$.* Write $M = (A, T, \equiv)$ and $A = (S, i, L, \rightarrow)$.

First, we define a preliminary LTS A_0 : Define the equivalence relation \sim_0 on S by

$$s \sim_0 s' :\Leftrightarrow s = s' \text{ or } \exists t, t' \in T \text{ nonempty finite : } \text{last}(t) = s, \text{last}(t') = s', t \equiv t'.^{10}$$

Write the equivalence classes as $[s]_0$. Now, we define $A_0 := (S_0, i_0, L_0, \rightarrow_0)$ as:

- $S_0 := S / \sim_0$, $i_0 := [i]_0$, $L_0 := L$.
- $[s]_0 \xrightarrow{\alpha} [s']_0$ iff $\exists s_0 \in [s]_0 \exists s_1 \in [s']_0 : s_0 \xrightarrow{\alpha} s_1$.

Next we define the actual LTS A_1 : Define an equivalence relation \sim_1 on S_0 by: $[s]_0 \sim_1 [s']_0$ iff, roughly, there is a (possibly empty) path in A_0 from $[s]_0$ to $[s']_0$ and one from $[s']_0$ back to $[s]_0$. Precisely:

$$[s]_0 \sim_1 [s']_0 \text{ iff } [s]_0 = [s']_0 \text{ or there is a nonempty finite } A_0\text{-trajectory } t \text{ starting in } [s]_0 \text{ and ending in } [s']_0 \text{ and there is a nonempty finite } A_0\text{-trajectory } t' \text{ starting in } [s']_0 \text{ and ending in } [s]_0.^{11}$$

(This is a very familiar concept: if we think of states as topological spaces and trajectories as continuous functions between these spaces with a trivial notion of homotopy between functions, then $[s]_0 \sim_1 [s']_0$ means that $[s]_0$ and $[s']_0$ are homotopy equivalent.) So, essentially, \sim_1 clusters S_0 into its connected components.

Write the equivalence classes as $[[s]_0]_1$. We define $A_1 := (S_1, i_1, L_1, \rightarrow_1)$ as:

¹⁰This is indeed an equivalence relation: It is reflexive by construction. Symmetry is immediate since \equiv is symmetric. And for transitivity: if $s \sim_0 s'$ and $s' \sim_0 s''$, then if one or both of these relations hold due to identity, we immediately get $s \sim_0 s''$, so assume these relations hold since there are nonempty finite $t, t', t'' \in T$ with $\text{last}(t) = s$, $\text{last}(t') = s'$, $\text{last}(t'') = s''$, $t \equiv t'$, and $t' \equiv t''$, then we have, by transitivity of \equiv , that $t \equiv t''$, so $s \sim_0 s''$, as needed.

¹¹This is indeed an equivalence relation: It is reflexive by construction. Symmetry is immediate (swap t and t'). And for transitivity: if $[s]_0 \sim_1 [s']_0$ and $[s']_0 \sim_1 [s'']_0$, then if one or both of these relations hold due to identity, we immediately get $[s]_0 \sim_1 [s'']_0$, and if we have the loop (t_0, t'_0) between $[s]_0$ and $[s']_0$ and the loop (t_1, t'_1) between $[s']_0$ and $[s'']_0$, then $(t_0 t_1, t'_1 t'_0)$ is a loop between $[s]_0$ and $[s'']_0$.

- $S_1 := S_0 / \sim_1$, $i_1 := [i_0]_1$, $L_1 := L_0 = L$.
- $[[s]_0]_1 \xrightarrow{\alpha} [[s']_0]_1$ iff $\exists [s_a]_0 \in [[s]_0]_1 \exists [s_b]_0 \in [[s']_0]_1 : [s_a]_0 \xrightarrow{\alpha} [s_b]_0$.

Then we define $\mathbf{E}(M) := (A_1, T_1, \equiv_1)$ where T_1 is the set of all A_1 -trajectories and \equiv_1 is extensional equivalence.

So $\mathbf{E}(M)$ is a full and extensional BTS (as mentioned in section 3.2.4), and it is countable (since $L_1 = L$ is countable and $|S_1| \leq |S_0| \leq |S| \leq \omega$ is countable). So we need to show that it is antisymmetric:

Assume $t, t' \in T_1$ are finite with $[t] \leq [t']$ and $[t'] \leq [t]$, and show $[t] = [t']$. If t or t' are empty, this implies that both are empty, so $[t] = [t']$. So let both be nonempty. If $\text{last}(t) = \text{last}(t')$, then $t \equiv_1 t'$ and the claim follows. So let $\text{last}(t) \neq \text{last}(t')$. Then there is nonempty path t_a in A_1 from $\text{last}(t)$ to $\text{last}(t')$ and a nonempty path t_b in A_1 from $\text{last}(t')$ to $\text{last}(t)$. Write

$$\begin{aligned} t_a : \text{last}(t) &= [[s^0]_0]_1 \xrightarrow{\alpha_1} [[s^1]_0]_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} [[s^n]_0]_1 = \text{last}(t') \\ t_b : \text{last}(t') &= [[r^0]_0]_1 \xrightarrow{\alpha'_1} [[r^1]_0]_1 \xrightarrow{\alpha'_2} \dots \xrightarrow{\alpha'_m} [[r^m]_0]_1 = \text{last}(t). \end{aligned}$$

We show that $[s^0]_0 \sim_1 [s^n]_0$, so $\text{last}(t) = \text{last}(t')$, whence $t \equiv_1 t'$, as needed.

By definition, for $i = 0, \dots, n-1$, there is $[s_a^i]_0 \in [[s^i]_0]_1$ and $[s_b^{i+1}]_0 \in [[s^{i+1}]_0]_1$ with $[s_a^i]_0 \xrightarrow{\alpha_{i+1}} [s_b^{i+1}]_0$. Similarly, for $j = 0, \dots, m-1$, there is $[r_a^j]_0 \in [[r^j]_0]_1$ and $[r_b^{j+1}]_0 \in [[r^{j+1}]_0]_1$ with $[r_a^j]_0 \xrightarrow{\alpha'_{j+1}} [r_b^{j+1}]_0$. Thus,

$$\begin{aligned} [s^0]_0 &\sim_1 [s_a^0]_0 \xrightarrow{\alpha_1} [s_b^1]_0 \sim_1 [s_a^1]_0 \xrightarrow{\alpha_2} [s_b^2]_0 \dots \xrightarrow{\alpha_n} [s_b^n]_0 \sim_1 [s^n]_0 \sim_1 [r^0]_0 \\ &\sim_1 [r_a^0]_0 \xrightarrow{\alpha'_1} [r_b^1]_0 \sim_1 [r_a^1]_0 \xrightarrow{\alpha'_2} [r_b^2]_0 \dots \xrightarrow{\alpha'_m} [r_b^m]_0 \sim_1 [r^m]_0 \sim_1 [s^0]_0. \end{aligned}$$

Note that $[s]_0 \sim_1 [s']_0$ in particular means that there is a (possibly empty) path in A_0 from $[s]_0$ to $[s']_0$. So we have:

- from $[s^0]_0$ there is a (possibly empty) A_0 -path to $[s_a^0]_0$, from which there is a (one-step) A_0 -path to $[s_b^1]_0$, from which there is (possibly empty) A_0 -path to $[s_a^1]_0$, from which \dots , from which there is a (one-step) A_0 -path to $[s_b^n]_0$, from which there is (possibly empty) A_0 -path to $[s^n]_0$
- from $[s^n]_0$ there is a (possibly empty) A_0 -path to $[r_a^n]_0$, from which there is a (one-step) A_0 -path to $[r_b^{n-1}]_0$, from which there is (possibly empty) A_0 -path to $[r_a^{n-1}]_0$, from which \dots , from which there is a (one-step) A_0 -path to $[r_b^0]_0$, from which there is (possibly empty) A_0 -path to $[s^0]_0$.

In sum, there is a nonempty A_0 -path from $[s^0]_0$ to $[s^n]_0$, and there is a nonempty A_0 -path from $[s^n]_0$ to $[s^0]_0$. Hence $[s^0]_0 \sim_1 [s^n]_0$, as needed.

Construction of η_M . We define $\eta_M := (\sigma, \lambda) : M \rightarrow \mathbf{E}(M)$ by $\sigma : S \rightarrow S_1$, $s \mapsto [[s]_0]_1$ and $\lambda : L \rightarrow L_1$, $\alpha \mapsto \alpha$.

This is an LTS-morphism: First, it maps i to $[[i]_0]_1 = i_1$. Second, assume $s \xrightarrow{\alpha} s'$ in A . Note that $\lambda(\alpha)$ always is defined, so we need to show $\sigma(s) \xrightarrow{\lambda(\alpha)} \sigma(s')$. Indeed, since $\alpha \in L$ and $s \in [s]_0$ and $s' \in [s']_0$, we have, by definition, $[s]_0 \xrightarrow{\alpha} [s']_0$. Similarly, $[[s]_0]_1 \xrightarrow{\alpha} [[s']_0]_1$, i.e., $\sigma(s) \xrightarrow{\lambda(\alpha)} \sigma(s')$.

We see that it is an BTS-morphism $M \rightarrow \mathbf{E}(M)$ as follows. Since M is approximable and $\mathbf{E}(M)$ is bisimulative (since full and extensional), we can apply proposition 3.3.4. So it suffices to show: (a) for all $t \in T$, $\eta_M(t) \in T_1$, and (c) for all nonempty finite $t, t' \in T$, if $t \equiv t'$, then $\eta_M(t) \equiv_1 \eta_M(t')$.

Since T_1 is the set of all A' -trajectories, (a) is trivial. Concerning (c), assume $t, t' \in T$ are nonempty finite with $t \equiv t'$. Note that, since λ is total, $\eta_M(t)$ and $\eta_M(t')$ are nonempty finite, too. Write $s := \text{last}(t)$ and $s' := \text{last}(t')$. So $s \sim_0 s'$, whence $\sigma(s) = [[s]_0]_1 = [[s']_0]_1 = \sigma(s')$, so

$$\text{last}(\eta_M(t)) = \sigma_{\eta_M}(\text{last}(t)) = \sigma(s) = \sigma(s') = \sigma_{\eta_M}(\text{last}(t')) = \text{last}(\eta_M(t')).$$

Since \equiv_1 is extensional equivalence, we have $\eta_M(t) \equiv_1 \eta_M(t')$.

Finally, since λ is total, η_M is synchronous, whence a morphism in ωBTS_a^s .

Universality. Now, let N be in $\omega\text{BTS}_{\text{fev}}^s$ and let $g : M \rightarrow N$ be a morphism, and find a unique $f : \mathbf{E}(M) \rightarrow N$ with $f \circ \eta_M = g$.

Concerning uniqueness, Assume f, f' are such morphisms. On labels, since λ_{η_M} is the identity, we have $\lambda_f = \lambda_f \circ \lambda_{\eta_M} = \lambda_g = \lambda_{f'} \circ \lambda_{\eta_M} = \lambda_{f'}$. On states, we have, for a state $[[s]_0]_1$ in $\mathbf{E}(M)$, that

$$\sigma_f([[s]_0]_1) = \sigma_f \circ \sigma_{\eta_M}(s) = \sigma_g(s) = \sigma_{f'} \circ \sigma_{\eta_M}(s) = \sigma_{f'}([[s]_0]_1).$$

Concerning existence, define $f = (\sigma_f, \lambda_f)$ as follows: $\lambda_f := \lambda_g$, and $\sigma_f : S_{\mathbf{E}(M)} \rightarrow S_N$ is defined by

$$\sigma_f([[s]_0]_1) := \sigma_g(s).$$

To show that this is well-defined, we show (a) if $s \sim_0 s'$, then $\sigma_g(s) = \sigma_g(s')$, and (b) if $[s]_0 \sim_1 [s']_0$, then $\sigma_g(s) = \sigma_g(s')$.

Concerning (a), assume $s \sim_0 s'$ and show $\sigma_g(s) = \sigma_g(s')$. If $s = s'$, this is trivial, so assume $s \neq s'$. Then, by definition of \sim_0 , there are nonempty finite $t, t' \in T$ with $\text{last}(t) = s$, $\text{last}(t') = s'$, and $t \equiv t'$. In particular, $[t] \sqsubseteq_M [t']$ and $[t'] \sqsubseteq_M [t]$. Since BTS-morphisms preserve \sqsubseteq , $[g(t)] \sqsubseteq_N [g(t')]$ and $[g(t')] \sqsubseteq_N [g(t)]$. Since \sqsubseteq agrees with \leq on finite trajectories, $[g(t)] \leq_N [g(t')]$ and $[g(t')] \leq_N [g(t)]$. Since N is antisymmetric, $[g(t)] = [g(t')]$, whence $g(t) \equiv_N g(t')$. Moreover, $g(t)$ and $g(t')$ also are nonempty finite since g is synchronous, so, since N is extensional, they have the same last state. So

$$\sigma_g(s) = \sigma_g(\text{last}(t)) = \text{last}(g(t)) = \text{last}(g(t')) = \sigma_g(\text{last}(t')) = \sigma_g(s').$$

Concerning (b), assume $[s]_0 \sim_1 [s']_0$ and show $\sigma_g(s) = \sigma_g(s')$. If $[s]_0 = [s']_0$, then, by (a), $\sigma_g(s) = \sigma_g(s')$, so assume $[s]_0 \neq [s']_0$. Then, by definition of \sim_1 ,

there is a nonempty finite A_0 -trajectory t (resp., t') starting in $[s]_0$ (resp., $[s']_0$) and ending in $[s']_0$ (resp., $[s]_0$). We show that there is a nonempty finite path t_a in N from $\sigma_g(s)$ to $\sigma_g(s')$ and another one t_b back. Then, since N is extensional and full, $[t_a] \leq [t_a t_b] = [t_b]$ and $[t_b] \leq [t_b t_a] = [t_a]$, so $t_b \equiv t_a$, so, since N is extensional, $\sigma_g(s) = \text{last}(t_b) = \text{last}(t_a) = \sigma_g(s')$, as needed.

Write t as $[s^0]_0 \xrightarrow{\alpha_1} [s^1]_0 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} [s^n]$ with $s^0 = s$ and $s^n = s'$. For each $i = 0, \dots, n-1$, there is, by definition, $s_a^i \in [s^i]_0$ and $s_b^{i+1} \in [s^{i+1}]_0$ with $s_a^i \xrightarrow{\alpha_{i+1}} s_b^{i+1}$. Thus,

$$s = s^0 \sim_0 s_a^0 \xrightarrow{\alpha_1} s_b^1 \sim_0 s^1 \sim_0 s_a^1 \xrightarrow{\alpha_2} s_b^2 \sim_0 s^2 \sim_0 s_a^2 \xrightarrow{\alpha_3} \dots \xrightarrow{\alpha_n} s_b^n \sim_0 s^n = s'$$

By applying the synchronous LTS-morphism $g = (\sigma_g, \lambda_g)$ we get, using (a), that

$$\begin{aligned} \sigma_g(s) &= \sigma_g(s^0) = \sigma_g(s_a^0) \xrightarrow{\lambda_g(\alpha_1)} \sigma_g(s_b^1) = \sigma_g(s^1) = \sigma_g(s_a^1) \xrightarrow{\lambda_g(\alpha_2)} \\ &\quad \sigma_g(s_b^2) = \sigma_g(s^2) = \sigma_g(s_a^2) \xrightarrow{\lambda_g(\alpha_3)} \dots \xrightarrow{\lambda_g(\alpha_n)} \sigma_g(s_b^n) = \sigma_g(s^n) = \sigma_g(s') \end{aligned}$$

This is a nonempty path t_a in N from $\sigma_g(s)$ to $\sigma_g(s')$. Similarly, the nonempty finite A_0 -trajectory t' from $[s']_0$ to $[s]_0$ yields a nonempty path t_b in N from $\sigma_g(s')$ to $\sigma_g(s)$, as needed.

Next, we show that f is an LTS-morphism. First, we have $\sigma_f(i_1) = \sigma_f([[i]_0]_1) = \sigma_g(i) = i_N$. Second, assume $[[s]_0]_1 \xrightarrow{\alpha} [[s']_0]_1$ in $\mathbf{E}(M)$. Since $\lambda_f = \lambda_g$ is always defined (qua synchronous morphism), we need to show $\sigma_f([[s]_0]_1) \xrightarrow{\lambda_f(\alpha)} \sigma_f([[s']_0]_1)$.

Indeed, by definition, there is $[s_a]_0 \in [[s]_0]_1$ and $[s_b]_0 \in [[s']_0]_1$ such that $[s_a]_0 \xrightarrow{\alpha} [s_b]_0$. So, again by definition, there is $s_c \in [s_a]_0$ and $s_d \in [s_b]_0$ such that $s_c \xrightarrow{\alpha} s_d$. Note that $[s_c]_0 = [s_a]_0$, so $[[s_c]_0]_1 = [[s_a]_0]_1 = [[s]_0]_1$, and similarly $[[s_d]_0]_1 = [[s']_0]_1$. Since g is an LTS-morphism, we have

$$\sigma_f([[s]_0]_1) = \sigma_f([[s_c]_0]_1) = \sigma_g(s_c) \xrightarrow{\lambda_g(\alpha) = \lambda_f(\alpha)} \sigma_g(s_d) = \sigma_f([[s_d]_0]_1) = \sigma_f([[s']_0]_1).$$

Also note that, by construction, $f \circ \eta_M = g$ (as LTS-morphisms): On labels, $\lambda_f \circ \lambda_{\eta_M} = \lambda_f = \lambda_g$. On states, for $s \in S_M$, $\sigma_f \circ \sigma_{\eta_M}(s) = \sigma_f([[s]_0]_1) = \sigma_g(s)$.

So it remains to show that f is a BTS-morphism (we already know it to be synchronous). Since $\mathbf{E}(M)$ and N are extensional and full, this follows by proposition 3.3.4. \square

3.5.2 Unlabeling and reflexing

In this subsection, we show that there is an optimal way of rendering an M in $\omega\text{BTS}_{\text{fey}}^5$ unlabeled and reflexive (recall that **u** stands for unlabeled and **r** for reflexive):

$$\omega\text{BTS}_{\text{fey}}^s \begin{array}{c} \xrightarrow{\text{U}} \\ \perp \\ \xleftarrow{\text{I}} \end{array} \omega\text{BTS}_{\text{feyur}}^s$$

Spelled out, this means the following.

3.5.2. PROPOSITION. *The inclusion $\text{I} : \omega\text{BTS}_{\text{feyur}}^s \rightarrow \omega\text{BTS}_{\text{fey}}^s$ is a right adjoint: For each M in $\omega\text{BTS}_{\text{fey}}^s$ there is $\text{U}(M)$ in $\omega\text{BTS}_{\text{feyur}}^s$ and $\eta_M : M \rightarrow \text{U}(M)$ such that, for every N in $\omega\text{BTS}_{\text{feyur}}^s$ and every $g : M \rightarrow N$, there is a unique morphism $f : \text{U}(M) \rightarrow N$ with $f \circ \eta_M = g$.*

$$\begin{array}{ccc} M & \xrightarrow{\eta_M} & \text{U}(M) \\ & \searrow g & \downarrow f \\ & & N \end{array}$$

PROOF. *Construction of $\text{U}(M)$.* Write $M = (A, T, \equiv)$ and $A = (S, i, L, \rightarrow)$. Define $\text{U}(M) = (A', T', \equiv')$ with $A' = (S', i', L', \rightarrow')$ as

- $S' := S, i' := i, L' := \{\cdot\}$ (where \cdot is some object),
- $s \rightarrow' s'$ iff (a) $s = s'$ or (b) $\exists \alpha \in L : s \xrightarrow{\alpha} s'$.
- T' is the set of all A' -trajectories
- \equiv' is extensional equivalence.

Then $\text{U}(M)$ is a countable, full and extensional BTS. It is unlabeled and reflexive by construction. So we need to show that it still is antisymmetric:

Let $t, t' \in T'$ be finite with $[t] \leq [t']$ and $[t'] \leq [t]$ in $\text{U}(M)$. If t or t' is empty, this implies that both are empty, whence $[t] = [t']$, so let both be nonempty. Write $s := \text{last}(t)$ and $s' := \text{last}(t')$ and show $s = s'$ (whence $t \equiv t'$).

By the assumption, there is a (possibly empty) A' -path t_0 from s to s' and a (possibly empty) A' -path t_1 from s' to s . If t_0 or t_1 are empty, then $s = s'$, as needed. So assume both are nonempty. Write $t_0 = t_0(0) \dots t_0(n)$ and $t_1 = t_1(0) \dots t_1(m)$. Let t_0^* be the result of deleting those $t_0(i)$ of the form (s, \cdot, s) . Thus, t_0^* is still an A' -trajectory, and if t_0^* is empty, then $s = s'$, as needed. Similarly, let t_1^* be the result of deleting those $t_1(j)$ of the form (s, \cdot, s) . Thus, t_1^* is still an A' -trajectory, and if t_1^* is empty, then $s = s'$, as needed.

So assume both t_0^* and t_1^* are nonempty. In particular, they still start in s (resp., s') and end in s' (resp., s). Write $n^* := |t_0^*| \geq 1$ and $m^* := |t_1^*| \geq 1$. Then each $t_0^*(i)$ (with $i = 0, \dots, n^* - 1$) must be due to clause (b), i.e., of the form $s_i \rightarrow' s'_i$ such that there is $\alpha_i \in L$ with $s_i \xrightarrow{\alpha_i} s'_i$. Similarly, each $t_1^*(j)$ (with

$j = 0, \dots, m^* - 1$) must be due to clause (b), i.e., of the form $r_j \dot{\rightarrow} r'_j$ such that there is $\alpha'_j \in L$ with $r_j \xrightarrow{\alpha'_j} r'_j$. But then we have the following A -trajectories

$$\begin{aligned} t_a : s &= s_0 \xrightarrow{\alpha_0} s'_0 = s_1 \xrightarrow{\alpha_1} s'_1 = s_2 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_{n^*-1}} s'_{n^*-1} = s' \\ t_b : s' &= r_0 \xrightarrow{\alpha'_0} r'_0 = r_1 \xrightarrow{\alpha'_1} r'_1 = r_2 \xrightarrow{\alpha'_2} \dots \xrightarrow{\alpha'_{m^*-1}} r'_{m^*-1} = s \end{aligned}$$

So, since M is full and extensional, $[t_b] \leq [t_b t_a] = [t_a]$ and $[t_a] \leq [t_a t_b] = [t_b]$, so, since M is antisymmetric, $t_b \equiv t_a$, so $s = \text{last}(t_b) = \text{last}(t_a) = s'$, as needed.

Construction of η_M . We define $\eta_M = (\sigma, \lambda) : M \rightarrow \mathbf{U}(M)$ as follows: $\sigma : S_M \rightarrow S_{\mathbf{U}(M)}$, $s \mapsto s$ and $\lambda : L_M \rightarrow L_{\mathbf{U}(M)}$, $\alpha \mapsto \cdot$.

This is an LTS-morphism: It maps i_M to $i_M = i_{\mathbf{U}(M)}$ and if $s \xrightarrow{\alpha} s'$, then $\lambda(\alpha) = \cdot$ is defined and, by clause (b), $s \xrightarrow{\lambda(\alpha) = \cdot} s'$.

Moreover, it is synchronous and, qua LTS-morphism between full and extensional systems, also a BTS-morphism.

Universality. Now let N be in $\omega\text{BTS}_{\text{feyur}}^s$ and let $g : M \rightarrow N$ be a morphism. Find a unique morphism $f : \mathbf{U}(M) \rightarrow N$ with $f \circ \eta_M = g$.

Uniqueness: Let f, f' be two such morphisms. On labels, they have to map the one label of $\mathbf{U}(M)$ to the one label of N qua synchronous morphisms between unlabeled systems. On states, let $s \in S_{\mathbf{U}(M)}$. Then $\sigma_f(s) = \sigma_f \circ \sigma_{\eta_M}(s) = \sigma_g(s) = \sigma_{f'} \circ \sigma_{\eta_M}(s) = \sigma_{f'}(s)$.

Existence: Define $f = (\sigma_f, \lambda_f) : \mathbf{U}(M) \rightarrow N$ by $\sigma_f : S_{\mathbf{U}(M)} = S_M \rightarrow S_N$, $s \mapsto \sigma_g(s)$ and λ_f as the unique function from the singleton $L_{\mathbf{U}(M)}$ to the singleton L_N .

This is an LTS-morphism: First, it maps $i_{\mathbf{U}(M)} = i_M$ to $\sigma_g(i_M) = i_N$. Second, assume $s \dot{\rightarrow} s'$, and show, since $\lambda_f(\cdot) = \cdot_N$ is defined, that $\sigma_f(s) \xrightarrow{\lambda_f(\cdot)} \sigma_f(s')$. If $s \dot{\rightarrow} s'$ is due to clause (a), then $s = s'$, so $\sigma_f(s) = \sigma_f(s')$, whence, since N is reflexive, $\sigma_f(s) \xrightarrow{\lambda_f(\cdot)} \sigma_f(s')$. If $s \dot{\rightarrow} s'$ is due to clause (b), then there is $\alpha \in L_M$ with $s \xrightarrow{\alpha} s'$. Since g is a synchronous LTS-morphism, $\lambda_g(\alpha) = \cdot_N$ is defined and $\sigma_f(s) = \sigma_g(s) \xrightarrow{\cdot_N = \lambda_f(\cdot)} \sigma_g(s') = \sigma_f(s')$.

Moreover, f is synchronous and, qua LTS-morphism between full and extensional systems, also a BTS-morphism. \square

3.5.3 Adjunction to domains

In this subsection, we establish the remaining reflective adjunction. We show that with a slight extension \mathbb{T}_i of the trajectory domain construction we can go from $\omega\text{BTS}_{\text{feyur}}^s$ to \mathbf{iALG} , and we show that there is an optimal way back: i.e.,

$$\omega\text{BTS}_{\text{feyur}}^s \begin{array}{c} \xrightarrow{\mathbb{T}_i} \\ \perp \\ \xleftarrow{\mathbb{B}} \end{array} \mathbf{iALG}$$

is a reflective adjunction.

First, we need to define \mathbf{iALG} and then \mathbb{T}_i . In the introduction, we've already motivated and intuitively defined the category \mathbf{iALG} . The formal definition is as follows.

3.5.3. DEFINITION. An *initialized domain* (or, in full, an initialized ω -algebraic dcpo) is a pair (D, c) where D is an ω -algebraic dcpo and $c \in K(D)$. We call c the *initial* element of D .

A *morphism* $f : (D, c) \rightarrow (E, d)$ between initialized domains is a Scott-continuous function $f : D \rightarrow E$ that preserves compactness (if $x \in K(D)$, then $f(x) \in K(E)$) and the initial element ($f(c) = d$).

Let \mathbf{iALG} be the category of initialized domains and their morphisms (the identity morphism is the identity function and morphism composition is function composition).

Intuitively, in a domain with a least element, the least element acts like an ‘initial’ element. But in the absence of a least element, many choices of an initial element are possible, and the notion of an initialized domain makes these choices explicit.

In particular, we have a forgetful functor $\mathbf{G}_i : \mathbf{iALG} \rightarrow \omega\mathbf{ALG}$ sending (D, c) to D and $f : (D, c) \rightarrow (E, d)$ to $f : D \rightarrow E$.

The trajectory domain of a system in $\omega\mathbf{BTS}_{\text{feyur}}^s$ naturally yields an initialized domain.

3.5.4. LEMMA. *The following defines a functor $\mathbb{T}_i : \omega\mathbf{BTS}_{\text{feyur}}^s \rightarrow \mathbf{iALG}$:*

- For M in $\omega\mathbf{BTS}_{\text{feyur}}^s$, define $\mathbb{T}_i(M) := (\mathbb{T}(M) \setminus \{[[\epsilon]]\}, [[i_M \xrightarrow{M} i_M]])$.
- For $f : M \rightarrow N$ in $\omega\mathbf{BTS}_{\text{feyur}}^s$, define $\mathbb{T}_i(f) := \mathbb{T}(f) \setminus \{([\epsilon], [\epsilon])\}$.

PROOF. First, we show that $\mathbb{T}_i(M)$ is in \mathbf{iALG} . We already know that $\mathbb{T}(M)$ is an ω -algebraic domain and since M is full, $\epsilon \in T_M$, so $[[\epsilon]]$ is the least element of $\mathbb{T}(M)$ (since M is bisimulative, it suffices to note that, for every $n \geq 0$, there is $m := 0 \geq 0$ such that $[\epsilon \upharpoonright n] = [\epsilon] \leq [t \upharpoonright m]$). So, after removing the least element, $\mathbb{T}(M) \setminus \{[[\epsilon]]\}$ still is an ω -algebraic domain, and its compact elements are those of $\mathbb{T}(M)$ minus the least element.¹² Moreover, since M is full and reflexive, $i \dot{\rightarrow} i \in T$, whence $[[i \dot{\rightarrow} i]] \in K(\mathbb{T}(M) \setminus \{[[\epsilon]]\})$.

¹²To show this in detail: Let D be an ω -algebraic dcpo with least element \perp , and show that $D' := D \setminus \{\perp\}$ (with the inherited order) is again an ω -algebraic dcpo. Indeed, D' is a partial order, and it is directed: If $A \subseteq D'$ is directed, then also $A \subseteq D$ is directed, so $x := \bigvee A$ exists in D and $x \in D'$ (because there is, since A is nonempty, some $a \in A \subseteq D'$, so $\perp < a \leq x$) and x is a least upper bound of A in D' . Also note that if $A \subseteq D$ is directed with $\bigvee A \neq \perp$, then $A' := A \setminus \{\perp\} \subseteq D'$ is directed and $\bigvee A = \bigvee A'$. So, for all $x \in D'$, we have: x is compact in D' iff x is compact in D . Hence $K(D') = K(D) \setminus \{\perp\}$. In particular, $K(D')$ is countable. So it remains to show algebraicity: If $x \in D'$, then $A := \{c \in K(D) : c \leq x\}$ is directed and $\bigvee A = x > \perp$, so $A \setminus \{\perp\} = \{c \in K(D') : c \leq x\}$ is directed and $\bigvee A \setminus \{\perp\} = \bigvee A = x$, as needed.

Second, we show that $\mathsf{T}_i(f) : \mathsf{T}_i(M) \rightarrow \mathsf{T}_i(N)$ is a morphism between initialized domains. First, it is well-defined: $\mathsf{T}(f) : \mathsf{T}(M) \rightarrow \mathsf{T}(N)$ is a Scott-continuous function that maps $[[\epsilon]]$ to $[[\epsilon]]$ and any $[[t]]$ with t nonempty to $[[f(t)]]$ with $f(t)$ nonempty since f is synchronous. So $\mathsf{T}(f) \setminus \{([\epsilon], [[\epsilon]])\} : \mathsf{T}(M) \setminus \{([\epsilon], [[\epsilon]])\} \rightarrow \mathsf{T}(N) \setminus \{([\epsilon], [[\epsilon]])\}$ is a well-defined function and it still is Scott-continuous.¹³ Moreover, if $[[t]]$ is compact, then $t \in T_M$ is finite, so $f(t) \in T_N$ is finite, so $\mathsf{T}(f)([[f(t)]])$ is compact. And we have $\mathsf{T}(f)([[i_M \xrightarrow{M} i_M]]) = [[f(i_M \xrightarrow{M} i_M)]] = [[i_N \xrightarrow{N} i_N]]$.

Third, the functor conditions are satisfied: We have

$$\mathsf{T}_i(\text{id}_M) = \mathsf{T}(\text{id}_M) \setminus \{([\epsilon], [[\epsilon]])\} = \text{id}_{\mathsf{T}(M)} \setminus \{([\epsilon], [[\epsilon]])\} = \text{id}_{\mathsf{T}_i(M)}$$

and

$$\begin{aligned} \mathsf{T}_i(g \circ f) &= \mathsf{T}(g \circ f) \setminus \{([\epsilon], [[\epsilon]])\} = (\mathsf{T}(g) \circ \mathsf{T}(f)) \setminus \{([\epsilon], [[\epsilon]])\} \\ &= \mathsf{T}(g) \setminus \{([\epsilon], [[\epsilon]])\} \circ \mathsf{T}(f) \setminus \{([\epsilon], [[\epsilon]])\} = \mathsf{T}_i(g) \circ \mathsf{T}_i(f), \end{aligned}$$

as needed. \square

In fact, this functor $\mathsf{T}_i : \omega\text{BTS}_{\text{feyur}}^s \rightarrow \text{iALG}$ has a right adjoint.

3.5.5. PROPOSITION. *The functor $\mathsf{T}_i : \omega\text{BTS}_{\text{feyur}}^s \rightarrow \text{iALG}$ is a left adjoint: For each (D, c) in iALG there is $\mathsf{B}(D, c)$ in $\omega\text{BTS}_{\text{feyur}}^s$ and an isomorphism $\epsilon_{(D,c)} : \mathsf{T}_i\mathsf{B}(D, c) \rightarrow (D, c)$ such that, for every M in $\omega\text{BTS}_{\text{feyur}}^s$ and every $f : \mathsf{T}_i(M) \rightarrow (D, c)$, there is a unique morphism $g : M \rightarrow \mathsf{B}(D, c)$ with $\epsilon_{(D,c)} \circ \mathsf{T}_i(g) = f$.*

$$\begin{array}{ccc} M & & \mathsf{T}_i(M) \\ \downarrow g & & \downarrow \mathsf{T}_i(g) \\ \mathsf{B}(D, c) & & \mathsf{T}_i\mathsf{B}(D, c) \end{array} \quad \begin{array}{c} \searrow f \\ \xrightarrow{\epsilon_{(D,c)}} (D, c) \end{array}$$

PROOF. *Construction of $\mathsf{B}(D, c)$.* We define $\mathsf{B}(D, c) := (A, T, \equiv)$ with $A := (S, i, L, \rightarrow)$ as follows:

- $S := K(D)$, $i := c$, $L := \{\cdot\}$,
- $s \dot{\rightarrow} s'$ iff $s \leq s'$.
- T is the set of all A -trajectories,
- \equiv is extensional equivalence.

¹³If $A \subseteq \mathsf{T}(M) \setminus \{([\epsilon], [[\epsilon]])\}$ is directed, also $A \subseteq \mathsf{T}(M)$ is directed, so $\mathsf{T}_i(f)(\bigvee A) = \mathsf{T}(f)(\bigvee A) = \bigvee \mathsf{T}(f)(A) = \bigvee \mathsf{T}_i(f)(A)$.

So $\mathbf{B}(D, c)$ is a full and extensional BTS. It is countable since $K(D)$ is countable (D is ω -algebraic). It is unlabeled and reflexive by construction. And it is antisymmetric: For $t, t' \in T$, if $[t] \leq [t']$ and $[t'] \leq [t]$, then, if t or t' are empty, $[t] = [\epsilon] = [t']$, and if both t and t' are nonempty, there is a (possibly empty) path from $\text{last}(t)$ to $\text{last}(t')$ and one from $\text{last}(t')$ to $\text{last}(t)$, which, by definition of \rightarrow , means $\text{last}(t) \leq \text{last}(t') \leq \text{last}(t)$, so, since (D, \leq) is a partial order, $\text{last}(t) = \text{last}(t')$, so $t \equiv t'$. Hence $\mathbf{B}(D, c)$ is in $\omega\text{BTS}_{\text{feyur}}^s$.

Construction of $\epsilon_{(D,c)}$. We first show that the function $\varphi : K(\mathbf{T}(\mathbf{B}(D, c)) \setminus \{[\epsilon]\}) \rightarrow K(D)$ given by $[[t]] \mapsto \text{last}(t)$ is a well-defined order-isomorphism (then we'll define $\epsilon_{(D,c)}$ as an extension of φ). We've essentially given the proof already in the previous chapter, but we repeat it here for convenience.

Well-defined: Note that $[[t]] \neq [[\epsilon]]$ so t is nonempty finite, whence $\text{last}(t) \in S = K(D)$. Moreover, if $t, t' \in T_{\mathbf{B}(D,c)}$ are finite with $[[t]] = [[t']]$, then, by antisymmetry, $t \equiv t'$, so $\text{last}(t) = \text{last}(t')$.

Surjective: Let $x \in K(D)$, then $t := x \dot{\rightarrow} x$ is in T and $\varphi([[t]]) = \text{last}(t) = x$.

Monotone: Assume $[[t]] \leq [[t']]$ for $t, t' \in T$ finite nonempty, and show $\varphi([[t]]) \leq \varphi([[t']])$. Since t can be extended to a trajectory equivalent to t' , there is a trajectory from $\text{last}(t)$ to $\text{last}(t')$, whence, since $\rightarrow = \leq$, we have $\varphi([[t]]) = \text{last}(t) \leq \text{last}(t') = \varphi([[t']])$.

Order-respecting: Assume $\varphi([[t]]) \leq \varphi([[t']])$ for $t, t' \in T$ finite nonempty, and show $[[t]] \leq [[t']]$. Then $s := \text{last}(t) = \varphi([[t]]) \leq \varphi([[t']]) = \text{last}(t') =: s'$. So we have $[t] \leq [t']$: if $t_0 \in [t]$, then $\text{last}(t_0) = \text{last}(t) = s$, so $t_1 := t_0 s \dot{\rightarrow} s' \in T$ is an extension of t_0 with $t_1 \in [t']$ since $\text{last}(t_1) = s' = \text{last}(t')$.

Now, we define $\epsilon_{(D,c)} : \mathbf{T}_i \mathbf{B}(D, c) \rightarrow (D, c)$ by

$$\begin{aligned} \epsilon_{(D,c)}([[t]]) &:= \bigvee \{ \varphi([[t']]) : [[t']] \in K(\mathbf{T}_i \mathbf{B}(D, c)), [[t']] \sqsubseteq [[t]] \} \\ &= \bigvee \{ \text{last}(t') : t' \in T_{\mathbf{B}(D,c)} \text{ finite}, [[t']] \sqsubseteq [[t]] \} \end{aligned}$$

which is an order-isomorphism $\mathbf{T}(\mathbf{B}(D, c)) \setminus \{[\epsilon]\} \rightarrow D$ by lemma 3.2.6. In particular, $\epsilon_{(D,c)}$ is Scott-continuous and preserves compactness. And $\epsilon_{(D,c)}$ preserves the initial element: Since the initial element $[[i \dot{\rightarrow} i]]$ is compact, $\epsilon_{(D,c)}([[i \dot{\rightarrow} i]]) = \varphi([[i \dot{\rightarrow} i]]) = \text{last}(i \dot{\rightarrow} i) = i = c$. Hence also the inverse $\epsilon_{(D,c)}^{-1}$ is Scott-continuous, preserves compactness, and preserves the initial element. So $\epsilon_{(D,c)}$ is an isomorphism in \mathbf{iALG} .

Universality. Now, let M be in $\omega\text{BTS}_{\text{feyur}}^s$ and $f : \mathbf{T}_i(M) \rightarrow (D, c)$ a morphism, and find a unique morphism $g : M \rightarrow \mathbf{B}(D, c)$ with $\epsilon_{(D,c)} \circ \mathbf{T}_i(g) = f$.

Uniqueness: Assume g, g' are such morphisms. On labels, they both are, qua synchronous BTS-morphisms between unlabeled systems, the unique function from the singleton $L_M = \{\cdot_M\}$ to the singleton $L_{\mathbf{B}(D,c)} = \{\cdot\}$. On states, let $s \in S_M$, and show $\sigma_g(s) = \sigma_{g'}(s)$. Since M is reflexive and full, $t := s \xrightarrow{M} s \in T_M$

is nonempty. So $[[g(t)]]$ is compact in $\mathsf{T}(\mathsf{B}(D, c)) \setminus \{[[\epsilon]]\}$ and

$$\begin{aligned}\sigma_g(s) &= \sigma_g(\text{last}(t)) = \text{last}(g(t)) = \epsilon_{(D,c)}([[g(t)]]) \\ &= \epsilon_{(D,c)} \circ \mathsf{T}_i(g)([[t]]) = f([[t]]) = \epsilon_{(D,c)} \circ \mathsf{T}_i(g')([[t]]) \\ &= \epsilon_{(D,c)}([[g'(t)]]) = \text{last}(g'(t)) = \sigma_{g'}(\text{last}(t)) = \sigma_{g'}(s).\end{aligned}$$

Existence: Define $g = (\sigma, \lambda) : M \rightarrow \mathsf{B}(D, c)$ as follows: λ is the unique function from the singleton L_M to the singleton $L_{\mathsf{B}(D,c)}$ and $\sigma : S_M \rightarrow S_{\mathsf{B}(D,c)}$ is defined by

$$\sigma(s) := f([[s \xrightarrow{M} s]])$$

This is well-defined: Since M is reflexive and full, $t := s \xrightarrow{M} s \in T_M$, so $[[s \xrightarrow{M} s]] \in \mathsf{T}(M) \setminus \{[[\epsilon]]\}$ on which f is defined. Since f preserves compactness, $f([[s \xrightarrow{M} s]]) \in K(D) = S_{\mathsf{B}(D,c)}$.

We show that g is an LTS-morphism: First, it maps i_M to $f([[i_M \xrightarrow{M} i_M]]) = c$ since f preserves the initial element. Second, assume $s \xrightarrow{M} s'$, and show, since λ is total, $\sigma(s) \xrightarrow{\lambda(M)=} \sigma(s')$. We have $[s \xrightarrow{M} s] \leq [s' \xrightarrow{M} s']$ because if $t_0 \in [s \xrightarrow{M} s]$, then, since M is extensional, t_0 ends in s , so $t_1 := t_0(s \xrightarrow{M} s')$ is, since M is full, in T_M , and $\text{last}(t_1) = s'$, so $t_0 \preceq t_1 \in [s' \xrightarrow{M} s']$, as needed. Hence, since f is monotone, $\sigma(s) = f([[s \xrightarrow{M} s]]) \leq f([[s' \xrightarrow{M} s']]) = \sigma(s')$, i.e., $\sigma(s) \dot{\rightarrow} \sigma(s')$.

Finally, g is a BTS-morphism since it is an LTS-morphism between extensional and full BTSs, and it is synchronous by construction.

So it remains to show $\epsilon_{(D,c)} \circ \mathsf{T}_i(g) = f$. Since both sides are Scott-continuous functions $\mathsf{T}(M) \setminus \{[[\epsilon]]\} \rightarrow D$, it is enough to show that they agree on compact elements (lemma 3.2.7). Indeed, given $[[t]]$ with $t \in T_M$ finite nonempty, write $\text{last}(t) = s$, whence, since M is extensional, $[[t]] = [[s \xrightarrow{M} s]]$. Then

$$\begin{aligned}\epsilon_{(D,c)} \circ \mathsf{T}_i(g)([[t]]) &= \epsilon_{(D,c)} \circ \mathsf{T}_i(g)([[s \xrightarrow{M} s]]) = \epsilon_{(D,c)}([[g(s \xrightarrow{M} s)]) \\ &= \text{last}(g(s \xrightarrow{M} s)) = \sigma(\text{last}(s \xrightarrow{M} s)) = \sigma(s) = f([[s \xrightarrow{M} s]]) = f([[t]]),\end{aligned}$$

as needed. □

3.6 Toward incorporating labels on domains

The trajectory domain of a BTS abstracts away labels. (Though, depending on the choice of trajectory equivalence, information about labels may be ‘hidden’ in equivalence classes). So if we think of the trajectory domain as denotations of LTSs, we may wonder whether we can appropriately add explicit information about labels. In this section, we show how this might be done and that this curiously leads to an interpretation of relevance logic.

3.6.1 Marked domains

In fact, there actually are two reasons for considering labels. The first is the one just mentioned: In other words, now that we know that \mathbf{iALG} is a computational model that is more abstract than $\omega\mathbf{BTS}^s$, we may ask whether we can bring them closer together by adding labels. The second reason is that if we want to extend the adjunction $\omega\mathbf{BTS}_a^s \rightleftarrows \mathbf{iALG}$ to the partial simulation case, it seems like we have to keep some information about the labels: namely, on which labels the label function is defined.

Looking at trajectory domains, there is a suggestive idea of how to add labels: Given a countable BTS $M = (A, T, \equiv)$, assume we have finite $t, t' \in T$ such that t' extends t by one transition, i.e., $t' = ts \xrightarrow{\alpha} s'$. If we have $[[t]] \sqsubseteq [[t']]$ in the trajectory domain $\mathbf{T}(M)$, it then is natural to think of the order *interval* $([[t]], [[t']]) := \{x \in \mathbf{T}(M) : [[t]] \sqsubseteq x \sqsubseteq [[t']]\}$ as being *marked* by the label α . More generally, we think of $([[t]], [[t']])$ as being marked by a label α (abbreviated as $([[t]], [[t']]) \mathbf{m} \alpha$), if there are representative t_a and t_b of $[[t]]$ and $[[t']]$, respectively, that are of the form $t_b = t_a \text{last}(t_a) \xrightarrow{\alpha} \text{last}(t_b)$.¹⁴

We can extend this idea by adding the concept of an *idle* transition. (See Winskel and Nielsen (1995) and footnote 7 above.) We fix a symbol $*$ (which no LTS is allowed to use as a label) and interpret it as the ‘do nothing action’. Thus, we can extend each LTS by adding all transitions of the form $s \xrightarrow{*} s$, which we call *idle* transitions. Then we can think of the trivial intervals in the trajectory domain—i.e., those $([[t]], [[t']])$ with $[[t]] = [[t']]$ —as always being marked by the idle label $*$ since we can ‘extend’ a representative of $[[t]]$ by the ‘do nothing action’ and obtain a representative of $[[t']]$. Since the idle label cannot occur in other transitions, we have $([[t]], [[t']]) \mathbf{m} *$ iff $[[t]] = [[t']]$.

This, then, suggests a general idea of adding labels to an (ω -algebraic) domain in a domain-theoretic fashion: We have a domain D (e.g., the trajectory domain) and a countable set of labels L (e.g., from the countable BTS) with an additional label $*$, together with a relation $(x, y) \mathbf{m} a$ between pairs (x, y) of elements in $K(D)$ that are in the \leq -relation and elements a of $L \cup \{*\}$. Now, $L \cup \{*\}$ naturally forms a domain: A common way to represent a (countable) set L in domain theory (e.g., the natural numbers) is as the *flat domain* L_\perp consisting of L with the discrete order ($x \leq y$ iff $x = y$) together with a least element $\perp = *$. If L is countable, this is indeed an ω -algebraic domain. Thus, we get the following purely domain-theoretic definition.

3.6.1. DEFINITION. A *marked domain* is a structure (D, \mathbf{m}, F) where D is an ω -algebraic domain, F is a countable flat domain, and $\mathbf{m} \subseteq (\leq_D \upharpoonright K(D)) \times F$ is a relation such that

1. for all $(x, y) \in \leq_D \upharpoonright K(D)$, we have $(x, y) \mathbf{m} \perp_F$ iff $x = y$.

¹⁴Also cf. the labelled domains of Bracho and Droste (1994) for another way of adding labels to domains (in the context of automata with concurrency).

We read $(x, y) \mathbf{m} a$ as ‘the interval (x, y) is *marked* with a ’.

A morphism $f : (D, \mathbf{m}, F) \rightarrow (E, \mathbf{n}, G)$ between marked domains is a pair (α, β) of Scott-continuous functions $\alpha : D \rightarrow E$ and $\beta : F \rightarrow G$ such that

1. α preserves compactness,
2. $\beta(\perp) = \perp$, and
3. if $(x, y) \mathbf{m} a$, then $(\alpha(x), \alpha(y)) \mathbf{n} \beta(a)$.

We write $f = (\alpha_f, \beta_f)$ and call it a *marked domain morphism*.

Let \mathbf{mALG} be the category of marked domains with their morphisms. The identity morphism is $\mathbf{id}_{(D, \mathbf{m}, F)} = (\mathbf{id}_D, \mathbf{id}_F)$ and composition is component-wise: $g \circ f = (\alpha_g \circ \alpha_f, \beta_g \circ \beta_f)$.

Two comments: First, condition (1) on morphisms ensures that condition (3) ‘type-checks’: if $x \leq y$ are compact, then, since α is monotone and preserves compactness, $\alpha(x) \leq \alpha(y)$ are compact.

Second, the partiality of simulations of BTSs is mirrored on the domain-side as follows: if $f : (D, \mathbf{m}, F) \rightarrow (E, \mathbf{n}, G)$ is a morphism and $(x, y) \mathbf{m} a$ with $\beta_f(a) = \perp$ (i.e., is ‘undefined’), then $(\alpha(x), \alpha(y)) \mathbf{n} \beta(a)$, so $\alpha(x) = \alpha(y)$.

Next, we show that the observation about the structure of trajectory domains that motivated the above definition of marked domains does indeed yield a marked domain—in a functorial way.

3.6.2. PROPOSITION. *The following defines a functor $\mathbb{T}_m : \omega\mathbf{BTS}_a \rightarrow \mathbf{mALG}$:*

- M is sent to $(\mathbb{T}(M), \mathbf{m}, (L_M)_\perp)$, where $([[t]], [[t']]) \mathbf{m} a$ with t, t' finite iff (a) $[[t]] = [[t']]$ and $a = \perp$ or (b) $[[t]] \sqsubseteq [[t']]$ and there is $t_a \in [t_0] \in [[t]]$ and $t_b \in [t_1] \in [[t']]$ with $t_b = t_a \text{last}(t_a) \xrightarrow{\alpha} \text{last}(t_b)$.
- $f : M \rightarrow N$ is sent to $(\mathbb{T}(f), \beta)$ where β maps \perp to \perp and $\alpha \in L_M$ to $\lambda_f(\alpha)$ if defined and otherwise to \perp .

PROOF. First, note that $(\mathbb{T}(M), \mathbf{m}, (L_M)_\perp)$ is indeed a marked domain: $\mathbb{T}(M)$ is an ω -algebraic domain, $(L_M)_\perp$ is a countable flat domain, and, by construction $\mathbf{m} \subseteq (\sqsubseteq \upharpoonright K(\mathbb{T}(M)) \times (L_M)_\perp)$ with $(x, y) \mathbf{m} \perp$ iff $x = y$.

Second, note that $(\mathbb{T}(f), \beta)$ satisfies requirements (1)–(3): Concerning (1), since M is approximable, $\mathbb{T}(f)$ preserves compactness. Concerning (2), by definition, $\beta(\perp) = \perp$. Concerning (3), assume $([[t]], [[t']]) \mathbf{m} a$ and show

$$([[f(t)], [f(t')]]) \mathbf{m} \beta(a).$$

If $a = \perp$, this always follows.¹⁵ So let $a \neq \perp$. Hence $\alpha := a \in L_M$ and $[[t]] \sqsubseteq [[t']]$ and there is $t_a \in [t_0] \in [[t]]$ and $t_b \in [t_1] \in [[t']]$ with $t_b = t_a \text{last}(t_a) \xrightarrow{\alpha} \text{last}(t_b)$. Since $\mathbb{T}(f)$ is monotone and well-defined, $[[f(t_a)]] = [[f(t)]] \sqsubseteq [[f(t')]] = [[f(t_b)]]$.

¹⁵If $(x, y) \mathbf{m} a$ and $a = \perp$, then $x = y$, so $\alpha(x) = \alpha(y)$, so $(\alpha(x), \alpha(y)) \mathbf{n} \perp = \beta(\perp)$.

If $\lambda_f(\alpha)$ is undefined, then $f(t_b) = f(t_a)$, so $[[f(t)]] = [[f(t')]]$ and $\beta(\alpha) = \perp$, so $([[f(t)], [[f(t')]]) \mathbf{m} \beta(a)$, as needed.

So assume $\lambda_f(\alpha) = \beta(\alpha)$ is defined. Then $[[f(t)]] \sqsubseteq [[f(t')]]$ and $f(t_a) \in [f(t_a)] \in [[f(t_a)]] = [[f(t)]]$ and $f(t_b) \in [f(t_b)] \in [[f(t_b)]] = [[f(t')]]$ and

$$f(t_b) = f(t_a) \underbrace{\text{last}(f(t_a))}_{=\sigma_f(\text{last}(t_a))} \xrightarrow{\lambda_f(\alpha)=\beta(\alpha)} \underbrace{\text{last}(f(t_b))}_{=\sigma_f(\text{last}(t_b))},$$

so $([[f(t)], [[f(t')]]) \mathbf{m} \beta(a)$, as needed.

Third, the functor conditions are satisfied: Concerning identity, $\mathbb{T}_m(\text{id}_M) = (\mathbb{T}(\text{id}_M), \beta)$ where $\mathbb{T}(\text{id}_M) = \text{id}_{\mathbb{T}(M)}$ is the identity on $\mathbb{T}(M)$ and β maps the bottom element to the bottom element and $\alpha \in L_M$ to $\text{id}_{L_M}(\alpha) = \alpha$, so it is the identity on $(L_M)_\perp$.

Concerning composition, let $f : M \rightarrow N$ and $g : N \rightarrow K$ be in ωBTS_a . Write $\mathbb{T}_m(f) = (\mathbb{T}(f), \beta)$ and $\mathbb{T}_m(g) = (\mathbb{T}(g), \beta')$ and $\mathbb{T}_m(g \circ f) = (\mathbb{T}(g \circ f), \beta'')$, and show $\mathbb{T}_m(g) \circ \mathbb{T}_m(f) = (\mathbb{T}(g) \circ \mathbb{T}(f), \beta' \circ \beta) = (\mathbb{T}(g \circ f), \beta'') = \mathbb{T}_m(g \circ f)$. Since \mathbb{T} is a functor, we have $\mathbb{T}(g \circ f) = \mathbb{T}(g) \circ \mathbb{T}(f)$. So it remains to show $\beta' \circ \beta = \beta''$, which is readily seen.¹⁶ \square

3.6.2 An interpretation of relevance logic

Relevance logic (or relevant logic) aims at providing a conditional $\varphi \rightarrow \psi$ where the antecedent φ is *relevant* to the consequence ψ . Notoriously, classical logic (or, more precisely, the material conditional or the strict conditional) cannot provide this: a sentence like $\varphi \rightarrow (\psi \rightarrow \psi)$ is logically valid, although the antecedent φ doesn't need to provide a reason for—or be relevant to—the consequence $\psi \rightarrow \psi$.

Various logical systems have been developed that provide such 'relevant conditionals'—both proof-theoretically and semantically. A common semantics is the ternary relation semantics (for an overview see Mares 2020). As in the usual Kripke semantics for modal logics, formulas are interpreted at possible worlds, but instead of a binary relation (interpreting the necessity operator) one uses a ternary relation R on worlds to interpret the conditional: $a \models \varphi \rightarrow \psi$ iff, for all worlds b and c , if $Rabc$ and $b \models \varphi$, then $c \models \psi$.

This provides a powerful formal semantics, though a common criticism is that it doesn't have a clear intended interpretation that provides contentful (as opposed to 'formal') meaning. (See e.g. Beall et al. (2012) and Mares (2020) for discussion.) Several such interpretations have been suggested revolving around

¹⁶By construction, both sides preserve the bottom element, so let $\alpha \in L_M$ and show $\beta' \circ \beta(\alpha) = \beta''(\alpha)$. If $\lambda_f(\alpha)$ is not defined, then also $\lambda_{g \circ f}(\alpha) = \lambda_g \circ \lambda_f(\alpha)$ is not defined, so $\beta' \circ \beta(\alpha) = \beta'(\perp) = \perp = \beta''(\alpha)$. So let $\lambda_f(\alpha) =: \alpha'$ be defined. If $\lambda_g(\alpha')$ is not defined, then also $\lambda_{g \circ f}(\alpha) = \lambda_g \circ \lambda_f(\alpha)$ is not defined, so $\beta' \circ \beta(\alpha) = \beta'(\alpha') = \perp = \beta''(\alpha)$. So let $\lambda_g(\alpha')$ be defined, then also $\lambda_{g \circ f}(\alpha) = \lambda_g \circ \lambda_f(\alpha)$ is defined, so $\beta' \circ \beta(\alpha) = \lambda_g \circ \lambda_f(\alpha) = \lambda_{g \circ f}(\alpha) = \beta''(\alpha)$.

the notion of information (for an overview see Mares 2020). Here we sketch a different concrete interpretation in terms of the behavior of labeled transition systems (or the abstracted version as marked domains).

Interpretation Surprisingly, a marked domain (D, \mathbf{m}, F) essentially has the structure of a frame in the simplified semantics of relevance logic (Priest and Sylvan 1992; Restall 1993; Restall and Tony 2009). We get the frame (g, W, R) —i.e., an interpretation minus the assignment of truth-values—with the base world $g := \perp_F$, the set of worlds $W := K(D) \cup F$, and the ternary relation $R \subseteq W^3$ defined by

$$(*) \quad Rabc \text{ iff (i) } (b, c) \mathbf{m} a \text{ or (ii) } a = g \text{ and } b = c \in F.$$

In particular, clause (1) of marked domains ensures that we have the condition on frames that $Rgbc$ iff $b = c$. (This explains the somewhat technical additional clause (ii) above.) This condition makes the truth conditions for the conditional univocal (otherwise one would need to distinguish in $a \models \varphi \rightarrow \psi$ between $a = g$ and $a \neq g$).

Thus, the conditional $\alpha \models p \rightarrow q$ says: for any order interval (x, y) that is marked with α , if x has property p , then y has property q . In the case of a trajectory domain of a system, this says: whenever behavior $[[t']]$ can be obtained from behavior $[[t]]$ by a single α -action, if behavior $[[t]]$ has property p , then $[[t']]$ has property q .

Under this interpretation, the notorious sentence $p \rightarrow (q \rightarrow q)$ can be falsified—as is typical, though not defining, for a relevance logic. Indeed, consider the following marked domain (D, \mathbf{m}, F) : D is the chain consisting of two elements x and y (the \leq -order is indicated by lines below), F is the flat domain of the singleton $\{\alpha\}$ together with the least element g , and \mathbf{m} relates $(x, y) \mathbf{m} \alpha$ (indicated as dotted lines below) and has otherwise only the ‘trivial’ relations $(x, x) \mathbf{m} \perp$ and $(y, y) \mathbf{m} \perp$. The atomic sentence p is set to be true at α and the atomic sentence q is set to be true at x . Visualized:

$$\begin{array}{ccccc} & & y & \cdots & \alpha & p \\ & & | & & | & \\ q & & x & \cdots & g & \\ & & & & & \end{array} \quad (3.1)$$

Then $g \not\models p \rightarrow (q \rightarrow q)$ since $Rg\alpha\alpha$ and $\alpha \models p$ but $\alpha \not\models q \rightarrow q$ because $R\alpha xy$ and $x \models q$ but $y \not\models q$. A very simple BTS that realizes this marked domain is $M = (A, T, \equiv)$ where A is the following LTS

$$i \xrightarrow{\alpha} s$$

and T is the set of all trajectories and \equiv is extensional equivalence. So $x = [[\epsilon]]$ and $y = [[i \xrightarrow{\alpha} s]]$. Intuitively, q could be the property of being the empty behavior and p could be the property of being a label.

Permuting The interpretation of R as in (*) above is somewhat rigid: for $x, y \in D$ and $\alpha \in F$ we can at most have $R\alpha xy$, but never, say, $Rx\alpha y$. (As a result, for any $x \in D$, every conditional $\varphi \rightarrow \psi$ is trivially true at x .) Though this ‘permutation’ $Rx\alpha y$ also would have a suggestive interpretation: $x \models p \rightarrow q$ meaning whenever x forms an interval with y that is marked by α and α has property p , then y has property q .

This suggests allowing permutations of the α, x, y as long as they form a ‘marking triangle’: To be precise, a set $\{a, b, c\} \subseteq W$ forms a *marking* if there is $x, y \in D$ and $\alpha \in F$ such that $(x, y) \mathfrak{m} \alpha$ and $\{a, b, c\} = \{x, y, \alpha\}$. Then we define the ‘closure’ of R under these permutations:

(**) $\overline{R}abc$ iff (i) $\{a, b, c\}$ forms a marking or (ii) $a = g$ and $b = c \in F$.

(Of course, one might also consider only allowing some but not all permutations.) If we also loosen clause (ii) under some permutation, say, add “ $a \neq g, b = g, c = a \in F$ ”, then we have $Rabc \Rightarrow Rbac$ which makes the assertion axiom $\varphi \rightarrow ((\varphi \rightarrow \psi) \rightarrow \psi)$ true (Restall 1993, thm. 2). (This axiom is part of the famous relevance logic R.) Note that both with and without this extension of (ii) we still have $Rgbc \Leftrightarrow b = c$.

Containment More correspondences between common axioms of various relevance logics and semantic conditions are obtained by extending the simplified semantics with a notion of containment (Restall 1993, p. 498): a relation \leq on W such that propositional atoms are monotone along \leq and, given $a \leq b$, if $a \neq g$, then $Rbcd \Rightarrow Racd$, and if $a = g$, then $Rbcd \Rightarrow c \leq d$.

We also have the domain-theoretic information containment order on W : the ‘merging’ of the order of D and the order of F , i.e., $a \leq b$ iff $a \leq_D b$ or $a \leq_F b$. We assume that D and F are (made) disjoint.

So it stands to reason that domain-theoretic information containment interprets the notion of containment in relevance logic. Indeed, this is the case for any R coming from a marked domain as in (*) together with any interpretation of atoms that is monotone along \leq :

Assume $a \leq b$. First, if $a \neq g$ and $Rbcd$, then the latter implies $b \in F$ (both in case (i) and in case (ii) the first entry of R is in F), whence $a \leq b$ implies $a \in F$, so $a \neq g = \perp_F$ implies that a is maximal in F , whence $a \leq b$ implies $a = b$, so $Racd$. Second, if $a = g$ and $Rbcd$, then, if $c \in D$, $Rbcd$ must be due to clause (i), whence, by definition of the marking relation, $c \leq d$, and if $c \in F$, then $Rbcd$ must be due to clause (ii), so $c = d$, whence $c \leq d$.

This prompts some discussion: First, note that in the second part we didn’t use the assumption $a = g$, whence $Rabc \Rightarrow b \leq c$. This implies, for any \leq -monotone interpretation of atoms, that $\varphi \rightarrow (\psi \rightarrow \psi)$ is valid (Restall 1993, thm. 11). Whether this validity is welcomed should be discussed: On the one hand, although not defining, the non-validity of this sentence is nonetheless typical

for ‘proper’ relevance logics. On the other hand, however, this crucially hinges on the monotonicity assumption: we’ve seen in (3.1) above that with a *non-monotone* assignment of atoms, this formula can be violated (there $x \models q$ but $x \leq y \not\models q$). And for the example interpretation of q as ‘is the empty behavior’ we indeed shouldn’t expect this property to be monotone along \leq . Thus, if we restrict us to the *simpler* logic of properties that are monotone along the domain order, we get a *stronger* logic that validates $\varphi \rightarrow (\psi \rightarrow \psi)$ —and this logic may be interesting in its own right.

Second, with the (**)-interpretation as \overline{R} , the above reasoning doesn’t go through. So one may explore variants or special cases of \overline{R} or \leq that deliver a containment relation.

Morphisms Furthermore, the conditions on a marked domain morphism $f = (\alpha, \beta) : (D, \mathbf{m}, F) \rightarrow (E, \mathbf{n}, G)$ ensures that it translates into a ‘frame-morphism’, i.e., a function between the set of worlds that preserves the base-world and the ternary relation:

$$\begin{aligned} \varphi : K(D) \cup F &\rightarrow K(E) \cup G \\ w &\mapsto \begin{cases} \alpha(w) & \text{if } w \in K(D) \\ \beta(w) & \text{if } w \in F. \end{cases} \end{aligned}$$

Indeed, φ maps $g_D = \perp_F$ to $g_E = \perp_G$. And if $R_D abc$, then $R_E \varphi(a)\varphi(b)\varphi(c)$: If $R_D abc$ is due to clause (i), then $(b, c) \mathbf{m} a$, so $(\alpha(b), \alpha(c)) \mathbf{n} \beta(a)$, so $R_E \varphi(a)\varphi(b)\varphi(c)$. If $R_D abc$ is due to clause (ii), then $a = g_D$ and $b = c \in F$, so $\varphi(a) = g_E$ and $\varphi(b) = \varphi(c) \in G$, so $R_E \varphi(a)\varphi(b)\varphi(c)$.

Similarly for \overline{R} : For clause (ii) we reason identically, and if $\overline{R}_D abc$ is because $\{a, b, c\}$ forms a marking $(x, y) \mathbf{m} \alpha$, then $\{\varphi(x), \varphi(y), \varphi(\alpha)\} = \{\varphi(a), \varphi(b), \varphi(c)\}$ forms a marking, so $\overline{R}_E \varphi(a)\varphi(b)\varphi(c)$.

Open questions This interpretation poses many interesting questions: Which models of relevance logic can be represented this way—both with (*) and with (**)? What is (an axiomatization of) the relevance logic of this interpretation (i.e., the set of sentences valid on it)? How do additional relevance logic axioms correspond to restriction on the domains and the systems giving rise to them? Does this then yield a logic for LTSs and BTS? And how does it compare to the usual one: linear temporal logic?

3.7 Conclusion

To summarize, we’ve established the categorical connections depicted in figure 3.1: Each adjunction either is reflective or co-reflective. The dotted arrows are to indicate the adjunction $\omega\text{BTS}_a^s \rightleftharpoons \text{iALG}$ masking the three adjunctions from which

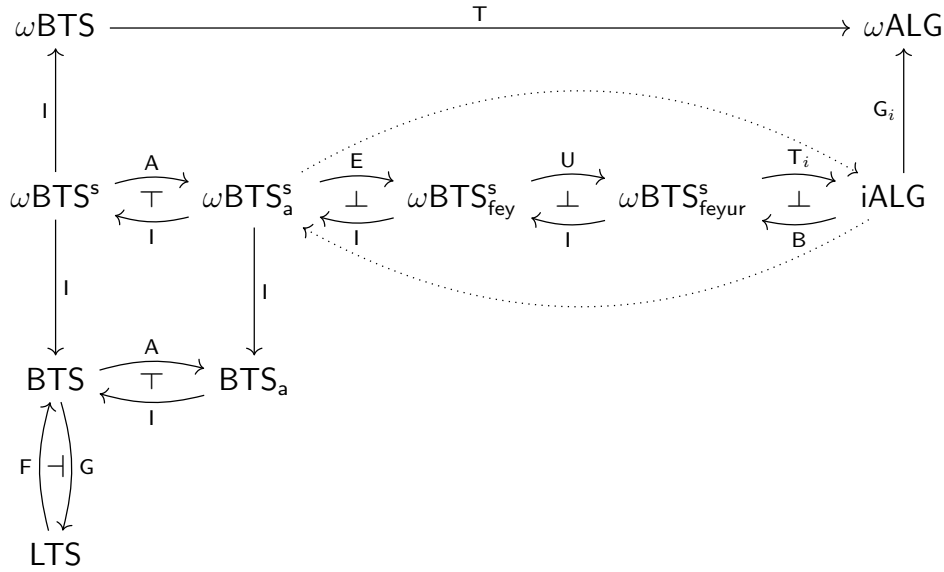


Figure 3.1: Summary of the results.

it is built. Intuitively, moving from left to right in the figure, the categories gradually become less ‘system-like’ and more ‘domain-like’. Note that commutativity is not claimed: While the small square involving the functor A commutes trivially, there is no reason why the big square (or rather rectangle) involving the functor T should commute.

We end with four open questions.

First, can the adjunction $\omega\text{BTS}_a^s \rightleftharpoons i\text{ALG}$ be extended to partial simulations? One approach may be using the above concept of a marked domain, and another approach may be using the representation of domains via Scott information systems which are generalized, as seen in the previous chapter, as BTSs.

Second, the issue of coincidence of the operational and denotational semantics (full abstraction) should be discussed further. We’ve seen that the strong sense of operational equivalence as isomorphism between countable BTSs implies the natural sense of denotational equivalence as trajectory domain isomorphism. What about weaker senses of operational equivalence like bisimulation? (For a categorical treatment of bisimulation, see Joyal, Nielsen, and Winskel (1996).) Approached from the other direction, is there an ‘operational’ equivalent to ‘having the same trajectory domain’, and is this related to bisimulation?

Third, we’ve seen that the denotational semantics provided by trajectory domains has some compositionality (in virtue of being a functor). Which further compositionality properties does it have? For example, is the denotation of a product of BTSs the product of the denotations of the BTSs? Some such preservation properties are already given by the established adjunctions, though what more can be said? In particular, what are the categorical constructions and

properties of **BTS**?

Fourth, we've already listed several open questions on the interpretation on relevance logic. Moreover, in the previous chapter, we've already asked whether the generalization of Scott information systems provided by BTSs can be seen as generalizing the underlying logic to substructural logics like relevance logic or linear logic: This is motivated by the seeming connection to the game semantics of linear logic discussed there and by the above interpretation of relevance logic.

Part Two

Non-symbolic computation

Chapter 4

Systems and domains 1: Model

Abstract With the aim of providing a new general tool for analyzing dynamical systems, we define the category of *dynamical domains*. These are structures in the sense of domain theory and can be seen as computational models for dynamical systems. We show that every dynamical system is isomorphic to the dynamical system modeled by some dynamical domain.

4.1 Introduction

Dynamical systems are tremendously important and ubiquitous in all areas of science. Hence, a lot of effort has been put into developing tools to understand dynamical systems. In this chapter (and the following ones), we wish to contribute to this ongoing effort: For every dynamical system \mathfrak{X} , we construct what we will call a *dynamical domain* \mathfrak{D} . This is a mathematical structure in the sense of domain theory (which is a mathematical theory of computation). Intuitively, it consists of ‘basic’ elements that represent increasingly finer observations of the system \mathfrak{X} together with the ‘limits’ of these basic elements. The additional domain-theoretic structure on \mathfrak{D} is such that it induces a dynamical system on these limit elements that is isomorphic to \mathfrak{X} . Thus, to every dynamical system \mathfrak{X} we associate the dynamical domain \mathfrak{D} which is a *computational model* for \mathfrak{X} . So we can translate questions about dynamical systems into questions about corresponding dynamical domains, to which the rich domain theory can be applied.

In this introduction, we motivate the class of dynamical systems that we’ll consider, and we sketch both how we construct the dynamical domain for a system and how we define the category of dynamical domains independently in a domain-theoretic spirit. We state our main result that every dynamical system can be modeled by a dynamical domain, and we mention two motivations for it. Finally, we discuss related work and outline the structure of the chapter.

The chapter is self-contained and provides all relevant background from both dynamical systems theory and domain theory. The length of this introduction is

due to providing all the conceptual explanation and motivation for the formal results and definitions to follow. The separation in paragraphs should help to skip the parts that may be less relevant to some readers.

Dynamical systems There are many (formal) notions of dynamical systems. What they have in common is that a dynamical system consists of a *state space* together with a *dynamics* which is a collection of ‘transformations’ of the state space indexed by a *time* parameter. These notions differ in what structure the state space has (measurable space, probability space, topological space, manifold), which of it is preserved by the transformations (measurable, measure-preserving, continuous, diffeomorphism), and what the time parameters are (the integers, the reals, a group).

Ergodic theory is about studying the qualitative behavior of dynamical systems in these various senses. In doing so, the abstract setting of a measure-preserving transformation of a probability space has proven to be particularly fruitful. (Formal definitions of the concepts to follow are given in section 4.2.2.) Thus, a dynamical system is a structure $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a probability space and $T : X \rightarrow X$ is a measure-preserving function.¹ Usually, one additionally assumes that (X, \mathcal{A}, μ) is a standard probability space (also called Lebesgue space) and that T is invertible. In that case, we call \mathfrak{X} a *standard* dynamical system.

However, in many applications (we’ll mention an example below), the natural measure on the state space may not be preserved by the dynamics: measure-preservation is a theorem (e.g., the Liouville theorem in classical mechanics) rather than an obvious axiom. Thus, by an *abstract* dynamical system we mean a structure $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a probability space and $T : X \rightarrow X$ is measurable (but not necessarily measure-preserving). To develop our representation result also in this setting, we’ll eventually also add a (somewhat milder) ‘standardness’ assumption on the underlying probability space: namely, to be a standard Borel space. (As in the Lebesgue case, this allows for a unified theory of isomorphism.) Thus, we call an abstract dynamical system (X, \mathcal{A}, μ, T) *general* if (X, \mathcal{A}) is a standard Borel space.

To summarize, abstract dynamical systems include both the standard and general ones and their main difference is that the latter don’t make any assumptions about the dynamics except for being measurable. A topological perspective on dynamical systems will eventually also be useful, in which case we call a structure (X, τ, μ, T) a *measured topological system* if (X, τ) is a Polish topology (i.e., separable and completely metrizable), μ a measure on its Borel σ -algebra, and $T : X \rightarrow X$ a continuous function.

¹I.e., T is measurable (if $A \in \mathcal{A}$, then $T^{-1}(A) \in \mathcal{A}$) and, for all $A \in \mathcal{A}$, we have $\mu(T^{-1}(A)) = \mu(A)$.

Example (learning) Ergodic theory supplies many examples of measure-preserving transformations of probability spaces. So, rather than reciting them here, we refer to, e.g., Petersen (1983) and Walters (1982). However, we mention an example that motivates generalizing our setting to transformations that may *not* be measure-preserving (and also not bijective). The example comes from *learning*, as, for instance, in neural networks or, more generally, in stochastic gradient descent (in optimization).

At each stage of the learning (or training or optimizing) process, the machine (or neural network or model) that we're optimizing is characterized by a set of parameters w (e.g., the *weights* of the neural network). Given a *data* point d , the optimization algorithm (e.g., backpropagation or gradient descent) produces a new set of parameters $w' = L(w, d)$ (' L ' as in *learning*). The whole point of the algorithm is that the machine in state w' is (or aims to be) a better approximation to the phenomenon from which the data points are sampled than it was before in the state w .

Thus, we have a set W of sets of parameters w and a set D of data points d and a function $L : W \times D \rightarrow W$. Usually, the set W is the \mathbb{R}^n , so we may, at the very least, assume that it is a Polish space. And the data set D usually is a finite set (finitely many samples), but, to account for the potential infinity of sampling, we'll only assume that D is countable (in fact, for our purposes here, D could be any Polish space).

Now, to understand this learning process, we're obviously interested in the (statistical) long-term behavior of the learning dynamics (as is the general motivation for ergodic theory). So, for an infinite sequence of data points $\delta = \langle d_0, d_1, \dots \rangle$ and an initial state w_0 , we're interested in the sequence:

$$w_0, L(w_0, d_0) =: w_1, L(w_1, d_1) =: w_2, L(w_2, d_2) =: w_3, \dots \quad (4.1)$$

Does it converge to some w (i.e., learn)? Does it get stuck in a non-optimal area of the state space? Is it 'all over the place' (e.g., dense in W) and hence doesn't work at all? However, due to, for example, noise or necessary imprecision in implementing the algorithm on a computer, we can only ask these questions statistically: Assuming a probability distribution p on W representing the likelihood of our choice of initial state w_0 , and a probability distribution q on D representing the likelihood of a datapoint in d (and hence of δ), we ask what is the probability of a yes-answer to the above questions?

We can write this setting more conveniently: Let $X := W \times D^\omega$, which, qua countable product of Polish spaces, is Polish. And the product measure $\mu := p \times q^\omega$ is a probability distribution on X . Define $T : X \rightarrow X$ by

$$T(w, \delta) := \left(L(w, \delta(0)), S(\delta) \right),$$

where $S(\delta) := \delta(1)\delta(2)\dots$ is the shift function. Thus, the (first entry of the) iterates $T^k(w, \delta)$ correspond to the sequence (4.1). Moreover, if L is measurable

(resp. continuous), then T is measurable (resp. continuous).² Measurability of L is a very weak demand, and continuity of L is a very plausible demand if the optimization algorithm is to be computable (according to the well-known slogan that ‘computability implies continuity’).

So our learning dynamics is the general dynamical system $\mathfrak{X} := (X, \mathcal{B}(X), \mu, T)$. (Here $\mathcal{B}(X)$ denotes the Borel σ -algebra of the space X .) And there is, of course, no reason to expect T to be measure-preserving, i.e., \mathfrak{X} to be standard. However, this poses the question of when (and what) preserved measures exists, and our general framework provides a good framework to investigate this (we come back to this in section 4.7 and chapter 7).³

Moreover, if L is continuous, then (X, τ, μ, T) is a measured topological system (where τ is the topology on X). If we, e.g., restrict $W = \mathbb{R}^n$ to the irrational numbers and work with a countable discrete D , this even is a measured zero-dimensional topological system. These systems, where the topological structure is reduced to a minimum, become important below as well.⁴

The construction We outline the construction of the dynamical domain for a given abstract dynamical system $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ (the details are in section 4.3).

A measurable subset A of X can be regarded as an observation or measurement that we can make about the system: if the system is in a state $x \in A$, making measurement A comes out positive. So if we have a finite set \mathcal{C} of measurable sets that cover the state space X (i.e., every state of X is in some set of \mathcal{C}), it provides a finite and non-deterministic dynamical system that ‘reflects’ the original system: the states are the elements from \mathcal{C} and there is a connection from

²Note that $h : W \times D^\omega \rightarrow W \times D$ defined by $h(w, \delta) := (w, \delta(0))$ is continuous, so $T_0 := L \circ h$ is measurable (resp. continuous). And $S : D^\omega \rightarrow D^\omega$ is continuous (for a subbasic open $p_n^{-1}V$, the preimage $S^{-1}(p_n^{-1}V) = p_{n+1}^{-1}V$ is open), so $T_1 : W \times D^\omega \rightarrow D^\omega$ defined by $T_1(w, \delta) := S(\delta)$ is continuous. So $T = (T_0, T_1)$ is measurable (resp. continuous).

³To give an idea of how this can be developed further: In this learning setting, there are various *loss* (or *cost*) *functions* that assign each weight state w a real number that indicates how ‘well’ the machine approximates the real phenomenon. Thus, we can regard such a cost function as a measurable function $f : X \rightarrow \mathbb{R}$. The move to such ‘observables’ is in striking analogy with the situation in physics that has led to the operator approach to ergodic theory (Eisner et al. 2015). So the tools from there—most notably the Koopman operator—can be used to analyze the learning dynamics. Also see chapter 7. Moreover, this idea of the data D ‘acting on’ the weights W may be compared to the notion of a nonautonomous dynamical system (see, e.g., Berger and Siegmund 2003).

⁴A few more remarks on general dynamical systems like the given example not belonging to the standard structures found in ergodic theory: As just seen, they cannot be regarded as measure-preserving transformations. Since X is not compact, they also don’t belong to the usual structures of topological dynamics. Measurable dynamics is somewhat of a middle ground (Weiss 1984): X is a standard Borel space, but the role of measure zero sets is here taken over by the sigma algebra generated by the wandering sets, yet this may not be the preferred notion of negligible sets given by the measure. In addition to this, there also is the issue that the dynamics is not necessarily bijective.

state A to state B if there is $x \in A$ with $T(x) \in B$. For each ‘observation length’ $n \geq 0$, a state $x \in X$ induces a set $\mathcal{O}_{\mathcal{C}}^n(x)$ of trajectories in this observed system: namely those $t = A_0, A_1, \dots, A_{n-1}$ such that $T^k(x) \in A_k$ (for $k = 0, \dots, n-1$). We can call $\mathcal{O}_{\mathcal{C}}^n(x)$ an observation history, and let $\mathbf{H}_{\mathcal{C}}^n$ be the set of observation histories. The dynamics T naturally induces a multi-valued function $f_{\mathcal{C}}^n$ on $\mathbf{H}_{\mathcal{C}}^n$: it maps $\mathcal{O}_{\mathcal{C}}^n(x)$ to the set $\{\mathcal{O}_{\mathcal{C}}^n(T(y)) : \mathcal{O}_{\mathcal{C}}^n(y) = \mathcal{O}_{\mathcal{C}}^n(x)\}$ of possible ‘next’ observation histories. We can turn this into a ‘usual’ function by moving to the set $D_{\mathcal{C}}^n$ of nonempty subsets of $\mathbf{H}_{\mathcal{C}}^n$ ordered by reverse inclusion and define the monotone function $f_{\mathcal{C}}^n(M) = \{\mathcal{O}_{\mathcal{C}}^n(T(y)) : \mathcal{O}_{\mathcal{C}}^n(y) \in M\}$. In domain-theoretic terminology (which we introduce in section 4.2.1), we’ve built the Smyth powerdomain $D_{\mathcal{C}}^n$ (which here is a Scott domain). It is a tool to study the original multi-valued (i.e., non-deterministic) function by the Scott-continuous function $f_{\mathcal{C}}^n : D_{\mathcal{C}}^n \rightarrow D_{\mathcal{C}}^n$. We also can induce a valuation $v_{\mathcal{C}}^n$ of the Scott-open sets of $D_{\mathcal{C}}^n$ by assigning each $\mathcal{O}_{\mathcal{C}}^n(x)$ the value $\mu\{y \in X : \mathcal{O}_{\mathcal{C}}^n(y) = \mathcal{O}_{\mathcal{C}}^n(x)\}$.

Thus, for the observation parameter $i = (n, \mathcal{C})$, we’ve obtained the structure $\mathfrak{D}_i = (D_i, v_i, f_i)$ of a finite Scott domain with a valuation v_i and a Scott-continuous function $f_i : D_i \rightarrow D_i$. Now, we can also refine our observation parameter to $j = (m, \mathcal{D})$ where $n \leq m$ and \mathcal{D} refines the cover \mathcal{C} (with a slight twist to the usual definition, see definition 4.3.2). And we have a function $p_{ij} : D_j \rightarrow D_i$ induced by mapping the finer observation history $\mathcal{O}_j(x)$ to the coarser $\mathcal{O}_i(x)$. Let’s write \mathcal{B} for the set of measurable subsets of X that we are prepared to count as ‘possible observation’, and let $I(\mathcal{B})$ be the set of observation parameters (n, \mathcal{C}) that we can build using this set. Then $(\mathfrak{D}_i, p_{ij})_{I(\mathcal{B})}$ is a diagram (or inverse system) of the domain-theoretic structures \mathfrak{D}_i . We want to take the (inverse) limit of the diagram, which will then eventually yield the dynamical domain that models the system \mathfrak{X} . However, to do so, we first have to specify what the category is in which we build the diagram (and the limit).

Dynamical domains The category in which we can build the required limit will be the category of dynamical domains. We motivate it in a purely domain-theoretic way (with the above construction in the back of our mind). A common way to define categories of domains (e.g., bifinite domains) is to first specify a collection of finite domains (e.g., finite pointed posets) and define the desired category to consist of those objects that are obtained as appropriate limits of appropriate diagrams built with appropriate finite domains and appropriate morphisms (e.g., Scott-continuous projection).⁵ This idea, of course, is more general: for example, profinite graphs and groups (Ribes 2017) and Stone spaces (Johnstone 1982) are defined similarly; and the reason is that the resulting category usually has very pleasant properties.

⁵The morphisms of the desired category are then taken as the morphisms that may occur in the diagrams, but one drops the requirement of being a projection (i.e., in the case of bifinite domains, these then are the Scott-continuous functions).

We proceed similarly here: Our ‘appropriate finite domains’ are the structures $\mathfrak{D} = (D, v, f)$ where D is a finite Scott domain, v is a valuation on it with $v(D) = 1$ and all value ‘sits’ in the maximal elements, and $f : D \rightarrow D$ is a Scott-continuous function. Let’s call these finite max-normalized dynamical Scott domains. Our ‘appropriate morphisms’ are Scott-continuous projections that additionally satisfy some (not entirely obvious) properties having to do with the appropriate preservation of v and f (see definition 4.4.2). And our ‘appropriate diagrams’ are diagrams of these ‘finite domains’ with some (again not entirely obvious) constraints on their shape ultimately having to do with constraining the function f in the limit to model a dynamical system (see definition 4.4.7).

If we’re given an appropriate diagram $(\mathfrak{D}_i, p_{ij})_I$ of finite max-normalized dynamical Scott domains $\mathfrak{D}_i = (D_i, v_i, f_i)$, how do we build the appropriate limit $\mathfrak{D} = (D, v, f)$? We build $D = \{a \in \prod_I D_i : p_{ij}(a(j)) = a(i)\}$ as the usual limit of an expanding sequence of domain (Abramsky and Jung 1994, sec. 3.3). Fortunately, there also exist results on building the valuation v on D in a unique way from the v_i ’s (Goubault-Larrecq 2018, thm. 4.2). However, the issue is with the function f . The straightforward thing would be to define $f : D \rightarrow D$ elementwise by $f(a) := \langle f_i(a(i)) : i \in I \rangle$. But then f won’t, in general, preserve maximality (i.e., $f(\max D) \subseteq \max D$), since the f_i , in general, don’t preserve maximality. However, if we want that (D, v, f) models a dynamical system, we, in particular, want that f induces a transformation on $\max D$ (i.e., the space modeled by D ; more on this below) and, to do so, f has to be max-preserving. It turns out (theorem 4.4.8) that there is a canonical way of selecting a (maximal, if a is maximal) element above each $f_i(a(i))$ in a way that builds an element in D : There is a largest function $f : D \rightarrow D$ that is Scott-continuous and max-preserving such that, for all $a \in D$ and $i \in I$, $f(a)(i) \geq f_i(a(i))$. This will be our function f . One can then show that \mathfrak{D} is a ‘restricted’ limit in the sense that it has the following universal property: $(\mathfrak{D}, p_i)_I$ is a cone for $(\mathfrak{D}_i, p_{ij})_I$ with a max-preserving function f , and for any cone $(\mathfrak{E}, q_i)_I$ with a max-preserving function, there is exactly one morphism $\alpha : \mathfrak{E} \rightarrow \mathfrak{D}$ such that $q_i = p_i \circ \alpha$.

The system modeled by a dynamical domain We’ve already hinted at the well-known idea that a domain D is a computational model for the space $\max D$ of maximal elements (with the relative Scott topology).⁶ With our dynamical domains $\mathfrak{D} = (D, v, f)$, we can extend this to dynamical systems: the state space modeled by \mathfrak{D} is $\max D$ (which will be a compact zero-dimensional Polish space), the continuous dynamics is $f \upharpoonright \max D$, and the valuation v uniquely determines a probability measure μ_v on $\max D$ (see theorem 4.5.1). So $(\max D, \mathcal{B}(\max D), \mu_v, f \upharpoonright \max D)$ is a general dynamical system which we call the dynamical system modeled by \mathfrak{D} .

We’ll define standard dynamical domains as those dynamical domains that are obtained as restricted limits of diagrams that satisfy some additional (again

⁶See e.g. Edalat and Heckmann (1998), Lawson (1997), and Martin (1998).

not entirely obvious) conditions detailed in definition 4.4.7. These ensure that the dynamical systems modeled by those standard dynamical domains are standard, i.e., $f \upharpoonright \max D$ is bijective and preserves the measure μ_v determined by the valuation.

Roughly speaking, then, the setup of dynamical domains is general enough to handle both deterministic and non-deterministic dynamics, and restricting in a category-theoretic way to the appropriate diagrams and limits corresponds to restricting to deterministic dynamics which are the object of study in ergodic theory.

The main result While the main conceptual (and technical) contribution of this chapter is to define the category of dynamical domains, the main *result* is to show that:

For every (standard) dynamical system \mathfrak{X} , there is a (standard) dynamical domain \mathfrak{D} such that the (standard) dynamical system modeled by \mathfrak{D} is isomorphic to \mathfrak{X} (corollary 4.6.4 below).

We'll now describe two interpretations of this result that motivate it: the representational interpretation and the computational interpretation.

The representational interpretation To understand a given class of dynamical systems, representation results are crucial: Given a class \mathcal{C} , find a (more restricted class) class \mathcal{D} of dynamical systems such that every system in \mathcal{C} is isomorphic to a system from \mathcal{D} . In other words, \mathcal{D} realizes all the isomorphism types of \mathcal{C} .

Here are three examples from the literature. First, the Jewett–Krieger theorem (see, e.g., Petersen 1983, sec. 4.4) states that every ergodic measure-preserving transformation on a Lebesgue space (forming the class \mathcal{C}) is isomorphic to a minimal, uniquely ergodic homeomorphism of a compact metric space (forming the class \mathcal{D}). This is a ‘topological representation’ result: representing (measure-theoretic) dynamical systems as coming from a topological dynamics. Second, the Krieger Generator Theorem says that the class \mathcal{D} of all finite subshifts (i.e., dynamical systems whose state space is the set of infinite sequences over a finite alphabet and the dynamics is the ‘shift’ operator) is complete for the class \mathcal{C} of ergodic measure-preserving transformation with finite entropy over a Lebesgue space. Third, as discussed by Weiss (1989), Rokhlin’s theorem states that every ergodic aperiodic measure-preserving transformation is isomorphic to a shift space on a countable alphabet. (See Weiss (1989) for similar results for the class of measurable systems and the class of topological systems.) The theorems of Krieger and Rokhlin are instances of ‘symbolic representation’: representing dynamical systems as subshifts over a finite or countable alphabet.

In our result, \mathcal{C} is the class of (standard) dynamical systems and \mathcal{D} is the class of dynamical systems modeled by a (standard) dynamical domain. This shares

with the Jewett–Krieger theorem that the representing systems have compact metric state spaces with continuous dynamics. Its conclusion is weaker, of course, but (and since) it also applies to a vastly more general class of dynamical systems—so this is a very general representation (and ‘topological realization’) result. And if we restrict us to standard dynamical systems, the conclusion strengthens to the representing system being a homeomorphism of a compact metric space. This, of course, also is weaker than the Jewett–Krieger theorem, but (and since) we also didn’t require any additional assumptions like ergodicity.

The computational interpretation The guiding idea behind the mentioned idea of a domain D being a model for the space $\max D$ is this: The domain consists of ‘approximate’ elements that approximate, when moving up in the order of the domain, the ‘maximal’ or ‘ideal’ elements of $\max D$. The classic example is the domain D of closed intervals $[\underline{x}, \bar{x}] \subseteq \mathbb{R}$ ordered by reverse inclusion: the ‘approximate reals’—i.e., $[\underline{x}, \bar{x}]$ with $\underline{x} < \bar{x}$ —approximate the ‘maximal reals’—i.e., $[\underline{x}, \bar{x}]$ with $\underline{x} = \bar{x}$ —, whence D is a model for $\max D \cong \mathbb{R}$ (Scott 1970, p. 16). Moreover, a manipulation of the ideal elements (e.g., a function $f : \mathbb{R} \rightarrow \mathbb{R}$) is ‘computational’ if it can be approximated by a manipulation of the ‘approximate’ elements. Formally, this is described by Scott-continuity of the extension $\bar{f} : D \rightarrow D$ of f . (For more on the idea of domain theory as a mathematical theory of computation, see section 4.2.1.)

Thus, a dynamical domain $\mathfrak{D} = (D, v, f)$ —obtained as a limit of some $\mathfrak{D}_i = (D_i, v_i, f_i)$ —is a computational model for the dynamical system $\mathfrak{X} = (\max D, \mathcal{B}(\max D), \mu_v, f \upharpoonright \max D)$ in the following sense: First, the elements of $\max D$ can be approximated by ‘finitary’ or ‘compact’ elements of D which in turn are determined by the elements of the finite D_i . In fact, D is a Scott domain and hence a particularly well-behaved domain (especially as a domain model for a space). Second, the measure μ_v is entirely determined by the valuation v , which in turn is determined by the finite v_i . Third, the dynamics $f \upharpoonright \max D$ is modeled by (i.e., extended by) the Scott-continuous $f : D \rightarrow D$, so the dynamics $f \upharpoonright \max D$ can be approximated in a computable way by its action on the finitary elements. A general theme will be that the concepts that we define for dynamical domains will be finitary in the sense that they can be expressed purely as a condition on the finitary diagrams from which the dynamical domains are constructed.

In that sense, our result says that every (standard) dynamical system has (up to isomorphism) a computational model. The general motivation for domain-theoretic computational models for classical mathematical structures is to provide effective models to make these structures constructive and to provide new proofs and algorithms using domain-theoretic tools (Edalat 1995a).

Finally, dynamical systems themselves may be seen as computing systems (see chapter 1). However, they are ‘non-symbolic’ in the sense that they (usually) act on a continuous state space (where a state is an infinite object) rather than a

discrete one (as, e.g., in a Turing machine). Thus, our computational models for dynamical systems may be regarded as providing a *symbolic* computational model for the *non-symbolic* computation performed by the dynamical system. After all, the domains are described by the discrete (i.e., countably many) compact elements which approximate the dynamical system to arbitrary precision—more on this in chapter 7. Note that the ‘symbolic representation’ of a dynamical system is—contrary to what the name may suggest—*not* a symbolic computational model of the dynamical system in this sense: its states are, qua infinite sequences of symbols, still infinite (and not discrete) objects.

Related work Concerning the construction of $(\mathfrak{D}_i, p_{ij})_{I(\mathcal{B})}$ from a system \mathfrak{X} , some of its aspects are found in the following references. Polish spaces with a distinguished basis play an important role for Danos and Garnier (2015) and Dahlqvist, Danos, and Garnier (2016) in the finitary analysis of natural transformations between functors on the category of Polish spaces.⁷ Building the index set using finite partitions of a space figures prominently in the proof that a topological space is profinite (i.e., the projective limit of finite discrete spaces) iff it is a compact Hausdorff totally disconnected space (i.e., a Stone space): see Borceux and Janelidze (2001, thm. 3.4.7) and, in the (second-) countable case, Danos and Garnier (2015); for more context, see Johnstone (1982, especially sec. VI 2.3). Some smaller differences are: We work, in general, with measurable bases over probability spaces and not just with open bases over Polish spaces, and we consider covers and not just partition. But the two main differences are: The objects of our diagrams are finite Scott domains with additional structure and not just finite spaces (whence we are in a different ambient category), and, most importantly, we additionally consider dynamics on the spaces; this is, as we’ll see in the proof of the theorem 4.4.8, the lion’s share of the work.

Concerning the representational interpretation, there are results in topological dynamics on obtaining certain topological dynamical systems (X, T) as an inverse limit of finite graphs: see Gambaudo and Martens (2006), Kucharski (2020), Shimomura (2014), and Shimomura (2020). There also is the well-known result that a topological system is zero-dimensional iff it is (topologically conjugate to) an inverse limit of subshifts (Downarowicz and Karpel 2016, thm. 2.21). However, as with the ‘symbolic representation’ of measure-theoretic dynamical systems, the elements of these subshifts are not discrete but infinite objects, so the representation as limits of finite graphs has the advantage, as our result, of being symbolic in the sense of approximating the system by finite means. There also is a natural graph structure on the maximal elements of our approximating domains D_i given by aE_ib iff $b \geq f_i(a)$. But, especially on the level of morphisms,

⁷ This has interesting applications to the Dirichlet distribution (as demonstrated in the just cited papers) and to Bayesian learning (Clerc et al. 2017; Gagné and Panangaden 2018). This suggests exploring whether there is a fruitful fusion of these and our ideas in the learning example mentioned above. We leave this as an open question (mentioned in section 4.7).

the constructions seem to be different.⁸ Moreover, our construction also takes measures into account (which are absent in topological dynamics) and hence also works for measure-theoretic dynamical systems, and, by using limits of domains rather than limits of graphs, it connects more directly to a theory of computation (i.e., domain theory).

Concerning the computational interpretation, the first paper that probably comes to mind is Edalat (1995b) introducing domain theory to dynamical systems.⁹ It investigates topological systems (X, f) by looking at the induced hyperspaces (X is not required to be compact). There are several hyperspace constructions that one could choose. (Also see Edalat and Heckmann (1998) for a construction using open balls.) The upper space is particularly suited for computational models: The upper space UX consists of all nonempty compact subsets of X with the topology given by the basic opens $\{C \in UX : C \subseteq A\}$ for $A \subseteq X$ open. If X is locally compact, second countable, and Hausdorff, then UX under reversed inclusion is an ω -continuous dcpo and its topology coincides with the Scott topology, whence the function $Uf : UX \rightarrow UX$ defined by $Uf(A) := f(A)$ is Scott-continuous. (Moreover, if X is zero-dimensional compact, then UX is a Scott domain.) Meaningful dynamic behavior of (X, f) like attractors then are, in the interesting cases, fixed points of the domain Uf (and fixed points are a central concept in domain theory).

Since Uf is max-preserving, the ω -continuous dcpo UX and Scott-continuous function Uf model the topological system $(\max UX, Uf \upharpoonright \max UX)$, and it is readily seen that this system is isomorphic to the original (X, f) via the conjugate homeomorphism $\varphi : X \rightarrow \max UX$ defined by $\varphi(x) = \{x\}$.¹⁰ Thus, this very elegantly provides a computational model for topological systems (X, f) with X locally compact, second-countable, Hausdorff and f continuous.

Our construction is, as outlined above, rather different, so let's compare the two. Our construction is more general in the sense that it works for arbitrary probability spaces (rather than the above topological spaces) and measurable functions (rather than continuous functions), and that, with the valuations, it has built in a domain-theoretic representation of measures. Moreover, the constructed modeling domain always is a Scott domain (rather than, generally, just ω -continuous). It models a compact topological system with a continuous dynamics, but, due to the greater generality, we can only ask the isomorphism to the original system to be a measure-

⁸For example, in our construction, there is no direct analogue of the '+ directionality' condition of Shimomura (2014, p. 184), which would say, for the above relation, if $aE_j b$ and $aE_j b'$, then $p_{ij}(b) = p_{ij}(b')$.

⁹This is part of a broader research programme, an overview of which is provided by Edalat (1997).

¹⁰The maximal elements of UX are precisely the singletons of X , so φ is well-defined and bijective. For an open $A \subseteq X$, we have $\varphi(A) = \{\{x\} : x \in A\} = \{C \in UX : C \subseteq A\} \cap \max UX$, so φ is open. Also, the preimage of the basic open $\{C \in UX : C \subseteq A\} \cap \max UX$ hence is $\varphi^{-1}\varphi(A) = A$, whence open, so φ is continuous. Finally, to see that φ is conjugate, we have, for $x \in X$, that $\varphi(f(x)) = \{f(x)\} = f(\{x\}) = Uf(\{x\}) = Uf(\varphi(x))$.

theoretic one (rather than a topological one). In section 4.6.2, we see that the price for the ‘emergent’ topological structure on the modeled system is that, if we restrict our setting to that of (zero-dimensional) topological systems, we, roughly, only have a dense embedding of the original system into the modeled system (so the modeled system acts as a ‘compactification’). But we do have a topological isomorphism if the original system was zero-dimensional compact (which is the case where the upper space model yields a Scott domain).

Outline of the chapter In section 4.2, we provide the relevant background from domain theory and dynamical systems theory. In section 4.3, we detail the above construction of observing an (abstract) dynamical system. In section 4.4, we build the category of dynamical domains. In section 4.5, we show that every dynamical domain models a dynamical system. In section 4.6, we show the main result that every system is modeled (up to isomorphism) by some dynamical domain. In section 4.7, we conclude with some open questions.

4.2 Background

We provide the relevant background in domain theory (section 4.2.1) and in dynamical systems theory (section 4.2.2).

4.2.1 Domain theory

We provide a brief—but self-contained—introduction to domain theory by recalling the basic notions that we’ll need. We follow the standard reference on domain theory: Abramsky and Jung (1994). Only the last notion (that of a valuation) is not covered there, for which we provide separate references below.

Dcpo. Let (D, \leq) be a partial order (abbreviated to *poset*).¹¹ (A subset $A \subseteq D$ is *directed* if it is nonempty and for any two $a, b \in A$, there is $c \in A$ such that $a, b \leq c$. If every directed subset A of D has a least upper bound $\bigvee A$, then (D, \leq) is *directed complete*. A *dcpo* is a directed complete partial order.

Scott domain. An element c of a dcpo (D, \leq) is *compact* if for all directed subsets A of D , if $\bigvee A \geq c$, then there is $a \in A$ such that $a \geq c$. The set of compact elements is denoted $K(D)$. A dcpo (D, \leq) is *ω -algebraic* if $K(D)$ is countable and, for all $a \in D$, the set $\{b \in K(D) : b \leq a\}$ is directed and has supremum a . Finally, a dcpo (D, \leq) is *bounded complete* if any subset $B \subseteq D$ that has an upper bound also has a least upper bound. A *Scott domain* is a non-empty ω -algebraic and bounded-complete dcpo (D, \leq) . Note that Scott domains have a least element: since the empty set has an upper bound, it hence has a least

¹¹I.e., D is a set and $\leq \subseteq D \times D$ is reflexive (for all $a \in D$, $a \leq a$), transitive (for all $a, b, c \in D$, if $a \leq b$ and $b \leq c$, then $a \leq c$), and antisymmetric (for all $a, b \in D$, if $a \leq b$ and $b \leq a$, then $a = b$). If we don’t demand antisymmetry, (D, \leq) is a *preorder*.

upper bound, which must be the least element. (In different contexts, one works with different classes of dcpos and calls the dcpos under consideration simply ‘domains’.)

Order-theoretic notation. Let (D, \leq) be a poset. A subset $A \subseteq D$ is an *upset* if, for all $a, b \in D$, $a \leq b$ and $a \in A$ implies $b \in A$. We call $a \in D$ *maximal* if, for all $b \in D$, if $b \geq a$, then $b = a$. The set of maximal elements of D is denoted $\max D$. More generally, if $A \subseteq D$ is a subset, $\max A := \{a \in A : \forall b \in A. b \geq a \Rightarrow b = a\}$.¹² Also, for a subset $A \subseteq D$, we define $\uparrow A := \{b \in D : \exists a \in A. b \geq a\}$ and $\downarrow A := \{b \in D : \exists a \in A. b \leq a\}$. If $A = \{a\}$ is a singleton, we write $\uparrow a := \uparrow A$ and $\downarrow a := \downarrow A$. By Zorn’s lemma, any element of a dcpo has a maximal element above it.¹³

Intuition. The guiding intuition is to think of the maximal elements of a dcpo as the ‘ideal’ elements that are approximated by the ‘real’ non-maximal elements. This intuition is made more precise by the way-below relation and continuous domains. We don’t need to define them here, since, in our setting of Scott domains, these more general concept can be described using compact elements only. Thus, the compact elements are the ‘real’, ‘finitary’, or ‘directly accessible’ elements of the Scott domain, and the maximal elements (or, more generally, the non-compact elements) are the ‘ideal’ elements that are obtained as limits of approximating them with the ‘real’ compact elements. As an example, consider the set $D := 2^{\leq \omega}$ of binary sequences of length $\leq \omega$ (the first infinite ordinal) ordered by extension. This forms a Scott domain: the compact elements are the finite sequences $2^{< \omega}$ and the maximal elements are the infinite sequences 2^ω .

The more general intuition, and the reason why domain theory is motivated as a ‘mathematical theory of computation’ (Scott 1970), is that domains can be regarded as providing *denotational semantics* to computational processes: In the above example, D can be regarded as the ‘data type’ of binary sequences which contains the output (or denotations) of computational processes that specify certain binary sequence (e.g., the process “print 0, print 1, repeat”). More importantly, we could also consider a computational process that takes binary sequences a as input and produces another one $f(a)$ as output. For this mapping to indeed be ‘computable’, we would expect that, if we want to approximate the output $f(a)$ to some finite degree $b_0 \leq f(a)$, we only need to know the input a up to some finite degree $a_0 \leq a$ such that $f(a_0) \geq b_0$ (or rather $f(\uparrow a_0) \subseteq \uparrow b_0$). The denotation of that process would then be an element of the domain $[D \rightarrow D]$ of ‘computable’ functions on D . Domain constructions like this function domain are central to domain theory, but will only implicitly play a role here (e.g., limits and powerdomains), but we’ll come back to this in the conclusion. What is important now is that specifying this kind of ‘(qualitative) computability’ is done

¹²If $A \subseteq D$ is an upset, then $\max A := \{a \in A : \forall b \in D. b \geq a \Rightarrow b = a\}$.

¹³I.e., if D is a dcpo, then $\forall x \in D \exists y \in \max D : x \leq y$. Proof: Since D is directed complete, any chain in the poset $P := \uparrow x$ (with the order inherited from D) has a (least) upper bound, so, by Zorn’s lemma, P has a maximal element y . Then $y \in \max D$ with $y \geq x$, as needed.

by continuity in an appropriate topology.

Topology. The most important topology on a dcpo (D, \leq) is the *Scott topology*. The open sets are those $U \subseteq D$ that are upsets and, whenever $A \subseteq D$ is directed and $\bigvee A \in U$, there is $a \in A$ such that $a \in U$. The intuition is that the Scott-open sets are the (finitely) observable properties of elements of D : If the limit $\bigvee A$ of an approximation has property U , then this will already be seen at a finite stage $a \in A$. The Scott topology has an important refinement: the *Lawson topology* which is the join of the Scott topology and the lower topology on D (generated by the sets of the form $D \setminus \uparrow a$ for $a \in D$). The Scott topology is denoted $\Sigma(D)$ and the Lawson topology is denoted $\Lambda(D)$; we drop ‘ D ’ if it is clear from context. The Scott-continuous functions $f : D \rightarrow D$ are those that are continuous with respect to the Scott topology. Equivalently, they are the \leq -monotone functions that preserve the supremum of directed subsets. (For continuous, and hence algebraic, domains, this is equivalent to the ‘finite approximation’ property sketched above.)

Projections. The most concise definition of a *projection* is as a surjective monotone function $p : Q \rightarrow P$ between posets such that preimages of principal upsets are again principal upsets (i.e., for all $a \in P$, $p^{-1}(\uparrow a) = \uparrow b$ for some $b \in Q$). This understanding should suffice for the chapter, however, the conceptually apt definition is that projections are one half of a pair of monotone functions $p : Q \rightleftarrows P : e$ that form an adjunction: so we’ll now explain that here as well.

Abstractly, a partial order (P, \leq) can be considered as a category (the objects of P are the elements of P and there is a single morphism from a to b iff $a \leq b$). A category-theoretic adjunction $l : P \rightleftarrows Q : u$ between two partial orders then is, concretely, a pair of monotone functions $l : P \rightarrow Q$ (called the *left* or *lower adjoint*) and $u : Q \rightarrow P$ (called the *right* or *upper adjoint*) such that, for all $a \in P$ and $b \in Q$, $l(a) \leq b$ iff $a \leq u(b)$. The category-theoretic fact that the right adjoint determines the left adjoint then becomes: For all $a \in P$, $l(a)$ is the least element of $u^{-1}(\uparrow a)$ (dually for the left adjoint determining the right adjoint). And the category theoretic fact that right adjoint functors preserve limits then becomes: u preserves existing infima (dually, l preserves existing suprema). Moreover, u is surjective iff $u \circ l = \text{id}_P$ iff l is injective (where id_X denotes the identity function on the set X). Now it’s not hard to see that a surjective monotone function $p : Q \rightarrow P$ between posets is a projection (in the above sense) iff it is an upper adjoint, i.e., there is a (uniquely determined) monotone function $e : P \rightarrow Q$ such that $e : P \rightleftarrows Q : p$ form an adjunction—then e is injective and is called an *embedding*. (Dually, we can define an embedding $e : P \rightarrow Q$ as an injective monotone function that is a lower adjoint.)

Powerdomains. Powerdomains have been developed to describe non-deterministic processes: While a deterministic computational process over a data type D maps an input $a \in D$ to a unique output $f(a) \in D$, a non-deterministic process maps an input $a \in D$ to a *set* of possible outputs $F(a) \subseteq D$. Intuitively, a powerdomain $\mathbf{P}(D)$ of D is a collection of subsets of D that can sensibly occur as non-deterministic outputs. Moreover, $\mathbf{P}(D)$ is ordered in a way that provides

domain-theoretic (or computable) structure. And the non-deterministic process f extends to a deterministic function on $\mathbf{P}(D)$: sending a set $M \subseteq D$ to $\bigcup_{a \in M} f(a)$.

There are several powerdomain constructions that make this precise (Abramsky and Jung 1994, sec. 6.2). Here we're using one of the common ones: the *Smyth powerdomain* \mathbf{P} . In fact, we'll only build it for *finite* and *discrete* nonempty dcpos (D, \leq) , i.e., where D is finite nonempty, and $a \leq b$ iff $a = b$. (But, conceptually, it is worth knowing that there is a more general construction behind it.) In this case, $\mathbf{P}(D)$ is defined as the finite partial order $(\mathcal{P}(D) \setminus \{\emptyset\}, \supseteq)$, i.e., the powerset of D ordered by reverse inclusion with the top element \emptyset removed.¹⁴ This is a finite Scott domain, and a multi-function $F : D \rightarrow D$ (mapping each $a \in D$ to a nonempty subset $F(a) \subseteq D$) becomes a Scott-continuous function $f : \mathbf{P}(D) \rightarrow \mathbf{P}(D)$ mapping M to $f(M) = \bigcup_{a \in M} F(a)$.

Valuations. In domain theory, the notion of a valuation plays a crucial role in the theory of probabilistic powerdomains (see e.g. Edalat 1995a; Jones and Plotkin 1989; Lawson 1982). But it also is treated in a more general topological setting (see e.g. Alvarez-Manilla, Edalat, and Saheb-Djahromi 2000; Keimel and Lawson 2005).¹⁵

If (D, \leq) is a dcpo, then a function $v : \Sigma(D) \rightarrow [0, \infty]$ is a *valuation* on (D, \leq) if, for all $U, V \in \Sigma(D)$,

1. Strictness: $v(\emptyset) = 0$,
2. Monotonicity: if $U \subseteq V$, then $v(U) \leq v(V)$, and
3. Modularity: $v(U \cup V) + v(U \cap V) = v(U) + v(V)$.

The valuation v is *continuous* if, whenever $(U_j)_{j \in J}$ is a directed family in $\Sigma(D)$, then $v(\bigcup_j U_j) = \sup_j v(U_j)$. And v is *normalized* if $v(D) = 1$.

Intuitively, (normalized) continuous valuations are the domain-theoretic analogues of (probability) measures.¹⁶ In the 'computable' spirit of domain theory, the Scott-open sets of a domain are the observable properties of the data type represented by the domain. A valuation then assigns probabilities to making the observations represented by the open sets. Thus, the open sets are the events to which we can assign probabilities, but, unlike the case of probability theory, we shouldn't expect to be able to do this for all the Borel sets generated by the

¹⁴Here is how this is a special case of the Smyth powerdomain: Since D is finite nonempty, this is a continuous dcpo (the only directed subsets are singletons). The Smyth powerdomain of a continuous domain is (isomorphic to) the collection of non-empty Scott-compact saturated subsets ordered by reverse inclusion (Abramsky and Jung 1994, thm. 6.2.14). Since D is finite and discrete, these are precisely the non-empty subsets of \mathbf{H}_i .

¹⁵As mentioned by Alvarez-Manilla, Edalat, and Saheb-Djahromi (2000, p. 629), the term 'valuation' goes back to the notion of a valuation on a lattice (Birkhoff 1973, ch. X).

¹⁶It is important to keep in mind that the analogy is not perfect (Keimel and Lawson 2005): Borel measures restrict to valuations, but not every valuation can be extended to a measure. However, as we'll see below, for the more well-behaved domains this is true.

Scott-open sets, since the complement of an observable property need not be observable anymore.¹⁷

4.2.2 Dynamical and topological systems

Dynamical systems. Before defining dynamical systems (in the sense of ergodic theory), we first recap their underlying spaces. There are three kinds: (1) probability spaces, which include both (2) standard Borel spaces with a probability measure and (3) Lebesgue spaces. These are defined as follows.

(1) As usual, a *probability space* is a triple (X, \mathcal{A}, μ) where X is a set, \mathcal{A} is a σ -algebra, and $\mu : \mathcal{A} \rightarrow [0, 1]$ is measure with $\mu(X) = 1$. A probability space (X, \mathcal{A}, μ) is *complete* if, for all $A \subseteq B \in \mathcal{A}$, if $\mu(B) = 0$, then $A \in \mathcal{A}$. The *completion* of (X, \mathcal{A}, μ) is denoted $(X, \mathcal{A}_\mu, \mu)$.¹⁸

(2) A *standard Borel space* is a pair (X, \mathcal{A}) such that there is a Polish (i.e., separable and completely metrizable) topology τ on X with $\mathcal{A} = \mathcal{B}(\tau)$, where $\mathcal{B}(\tau)$ denotes the Borel σ -algebra of the topology τ (Kechris 1995, def. 12.5). A probability measure μ on (X, \mathcal{A}) is then often called a Borel probability measure and (X, \mathcal{A}, μ) a Borel probability space.

(3) *Lebesgue spaces* (or standard probability spaces) can be defined in two equivalent ways.¹⁹ First, a *Lebesgue space* is a complete probability space (X, \mathcal{A}, μ) such that there is a second-countable topology τ on X with $\tau \subseteq \mathcal{A}$, $\mathcal{B}(\tau)_\mu = \mathcal{A}$, and μ inner regular, i.e., for $A \in \mathcal{A}$, $\mu(A) = \sup \mu(K)$ where the supremum is taken over all τ -compact subsets K of A (de la Rue 1993, def. 1-1). Second, a Lebesgue space is a complete probability space that is isomorphic mod 0 to the ordinary Lebesgue space of an interval $[0, a] \subseteq \mathbb{R}$ together with countably many point masses (Petersen 1983, def. 4.5, Walters 1982, def. 2.3). (For a proof of the equivalence, see de la Rue (1993, thm. 4-3).) The latter definition probably is more common and its intuition is that the unit interval with the Lebesgue measure (plus countably many point masses) serves as the canonical probability space. The first definition is, in a sense, conceptually more pure and its intuition is—analogue to standard Borel spaces—to consider probability spaces arising from well-behaved topological spaces. As a simple consequence of these definitions, any completion of a standard Borel space with a probability measure is a Lebesgue space.²⁰ And

¹⁷Similar motivation is given by, e.g., Jones and Plotkin (1989) and Keimel and Lawson (2005).

¹⁸Here \mathcal{A}_μ is the σ -algebra of sets of the form $A \cup N$ for $A \in \mathcal{A}$ and $N \subseteq M$ for some $M \in \mathcal{A}$ with $\mu(M) = 0$, and $\mu(A \cup N)$ is taken to mean $\mu(A)$.

¹⁹Regarding references: As noted by Petersen (1983, p. 17) and Eisner et al. (2015, remark 7.22), an early systematic treatment of Lebesgue spaces is provided by Rokhlin and by Halmos and von Neumann (see references therein). For a concise more modern treatment, see de la Rue (1993) or the entry “Standard probability space” of the *Encyclopedia of Mathematics*: http://encyclopediaofmath.org/index.php?title=Standard_probability_space&oldid=24675 (accessed 31 January 2021).

²⁰Proof: If $(X, \mathcal{A}_\mu, \mu)$ is the completion of the standard Borel space (X, \mathcal{A}) with a probability

any Lebesgue space is isomorphic mod 0 to the completion of a standard Borel space with a Borel probability measure.²¹

A major reason for restricting attention to the subclasses (2) and (3) of probability spaces is that then different natural notions of isomorphism of probability spaces—and hence dynamical systems built over them—coincide (Walters 1982, ch. 2).

Now we can define dynamical systems in the sense of ergodic theory. (For general references on ergodic theory, see, e.g., Petersen (1983) and Walters (1982).) In the most abstract sense, these are structures (X, \mathcal{A}, μ, T) where (X, \mathcal{A}, μ) is a probability space and $T : X \rightarrow X$ a measurable function.²² The standard setting of ergodic theory, however is more concrete in that it additionally assumes that the probability space is a Lebesgue space and that the transformation is measure-preserving and bijective (i.e., invertible). As motivated in the introduction, we aim for a treatment of dynamical systems that is general enough to not assume measure-preservation and bijectiveness from the start (and rather leave it as an option to obtain these as theorems). But it should also be reasonably concrete (to, e.g., allow for a unified theory of isomorphisms) and in line with other strands of dynamical systems theory, like measurable dynamics or descriptive dynamics. Thus, our ‘general’ setting is that of a Borel probability space with a Borel-measurable dynamics. As a result, a dynamical system in the standard sense is not verbatim a dynamical system in the general sense but only ‘modulo isomorphism and completion’. (Implicitly, this distinction between abstract, standard, and general is also found in Walters (1982, esp. ch. 2).) The formal definition reads as follows.

4.2.1. DEFINITION. An *abstract dynamical system* is a structure $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a probability space and $T : X \rightarrow X$ is measurable (i.e., for $A \in \mathcal{A}$, $T^{-1}(A) \in \mathcal{A}$). A *(general) dynamical system* is an abstract dynamical system $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}) is a standard Borel space. A *standard dynamical system* is an abstract dynamical system $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a Lebesgue space and $T : X \rightarrow X$ is bijective and measure-preserving (i.e., measurable and for $A \in \mathcal{A}$, $\mu(T^{-1}(A)) = \mu(A)$). We often omit the term ‘general’.

measure μ on it, then there is a Polish topology τ on X such that $\mathcal{A} = \mathcal{B}(\tau)$. So (X, \mathcal{A}, μ) is a complete probability space with a second-countable topology τ on X such that $\tau \subseteq \mathcal{A}_\mu$, $\mathcal{B}(\tau)_\mu = \mathcal{A}_\mu$, and to show inner regularity one uses that Borel measures on Polish spaces are inner regular.

²¹If (X, \mathfrak{A}, μ) is a Lebesgue space it is, as shown in the proof of the equivalence by de la Rue (1993, thm. 4-3), isomorphic mod 0 to the space (Y, \mathfrak{B}, ν) where Y is of the form $[0, a] \cup \{a_1, a_2, \dots\} \subseteq \mathbb{R}$, ν is the Lebesgue measure on $[0, a]$ and the point mass on the a_n , and $\mathfrak{B} = \mathcal{B}(Y)_\nu$. Note that Y is a G_δ subset (countable intersection of open sets) of the Polish space \mathbb{R} and hence Polish. So (X, \mathfrak{A}, μ) is isomorphic mod 0 to the completion of $(Y, \mathcal{B}(Y), \nu)$ which is a standard Borel space with a Borel probability measure.

²²The only further generalization is to move from the iterations of one dynamics T to more general group actions.

The usual notion of isomorphism between dynamical systems is the following (Walters 1982, ch. 2).

4.2.2. DEFINITION. Two abstract dynamical systems $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ and $\mathfrak{Y} = (Y, \mathcal{B}, \nu, S)$ are (*metrically*) *isomorphic* if there is a partial function $\varphi : X \rightarrow Y$ with domain $M \subseteq X$ and codomain $N \subseteq Y$ such that

1. $\varphi : M \rightarrow N$ is a bijective function,
2. M and N are invariant sets of full measure: i.e., $M \in \mathcal{A}$, $\mu(M) = 1$, $T(M) \subseteq M$, and $N \in \mathcal{B}$, $\nu(N) = 1$, $S(N) \subseteq N$,
3. φ is measure-preserving: i.e., for $B \in \mathcal{B}$, we have $\varphi^{-1}(B) \in \mathcal{A}$ and $\mu(\varphi^{-1}(B)) = \nu(B)$, and
4. φ is equivariant: i.e., for $x \in M$, $\varphi(T(x)) = S(\varphi(x))$.²³

For a discussion of when this coincides with the generally weaker conjugacy (equivariant isomorphism of measure algebras), see Walters (1982, ch. 2). For a discussion of the (history of) spatial vs. spectral approaches to the definition of isomorphisms, see Rédei and Werndl (2012).

Topological systems. We can look at a dynamical system (X, \mathcal{A}, μ, T) with a greater level of detail if we also consider topological information about the state space X and not just the measure-theoretic information. The resulting notion then is the following.

4.2.3. DEFINITION. A (*general*) *measured topological system* is a structure $\mathfrak{X} = (X, \tau, \mu, T)$ where (X, τ) is a Polish space, μ is a probability measure on $\mathcal{B}(\tau)$, and $T : X \rightarrow X$ is continuous. It is *standard* if, additionally, T is a homeomorphism and measure-preserving. It is zero-dimensional (resp. compact) if (X, τ) is. We usually omit the term ‘general’.

Comments: First, in topological dynamics, one usually doesn’t consider measures, whence we add the term ‘measured’ to stress the presence of a measure. The standard setting in topological dynamics is that X is a compact metric space (hence Polish) and T is a homeomorphism. Here, however, we’ll also discuss non-compact state spaces and non-bijective dynamics.

Second, a paradigm example of a zero-dimensional compact Polish space is the Cantor space. A paradigm example of a zero-dimensional (non-compact) Polish space is the space of irrational numbers considered as a subspace of the real numbers, which is homeomorphic to the Baire space. (See, e.g., Kechris (1995, ch. 7) for a discussion of these spaces and their ‘paradigmness’.)

Third, there are two perspectives on zero-dimensionality: From a topological perspective, as described by Hjorth and Molberg (2006, p. 1117), this means

²³Note that $\varphi(T(x))$ is defined because $x \in M$ implies $\varphi(x) \in M$.

minimizing the topological influence of the state space on the dynamical system (e.g., trivializing homotopy and homology), whence the complexity of the system comes from the dynamics. From a logical or computational perspective, the clopen (closed and open) sets of the state space act as ‘(finitely) decidable’ properties of states: Under the well-known computational interpretation of topology (Smyth 1983; Vickers 1989), the open sets of a topology are the ‘semi-decidable properties’ of the points of the space. Thus, the sets that not only are open but also have an open complement—i.e., the clopen sets—are the ‘decidable properties’ of the space. So, from this logico-computational perspective, the assumption that the clopen sets form a basis (i.e., zero-dimensionality) means that we can describe the states with these decidable properties. Moreover, from this perspective, zero-dimensionality can be seen as a ‘without loss of generality assumption’: If we start with a countable basis of our space (e.g., the intervals of \mathbb{R} with rational endpoints), then the points that lie exactly on the boundary of the basic opens (i.e., the rational numbers) are very much ‘non-typical’ points, whence they may be ignored, yielding a zero-dimensional space (the irrational numbers). This is a common idea, for example, in algorithmic randomness (Downey and Hirschfeldt 2010).

Fourth, note that, if the state space X is compact, it is enough to check that T is bijective and continuous to conclude that it is a homeomorphism.

Unlike the case of (measure-theoretic) dynamical systems, the notion of (iso-)morphisms between topological systems is straightforward.

4.2.4. DEFINITION. If $\mathfrak{X} = (X, \tau, \mu, T)$ and $\mathfrak{Y} = (Y, \sigma, \nu, S)$ are measured topological systems, a *morphism* $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ of measured topological systems is a function $\varphi : X \rightarrow Y$ that is continuous (i.e., if $V \in \sigma$, then $\varphi^{-1}(V) \in \tau$), measure-preserving (i.e., if $B \in \mathcal{B}(\sigma)$, then $\mu(\varphi^{-1}(B)) = \nu(B)$)²⁴, and equivariant (i.e., $\varphi \circ T = S \circ \varphi$). If φ additionally is a homeomorphism, then φ is an *isomorphism* of measured topological systems.

Every topological system induces a dynamical system: If $\mathfrak{X} = (X, \tau, \mu, T)$ is a measured topological system, then

$$J(\mathfrak{X}) := (X, \mathcal{B}(\tau), \mu, T)$$

is a general dynamical system.²⁵ And if \mathfrak{X} additionally is standard, then

$$\bar{J}(\mathfrak{X}) := (X, \mathcal{B}(\tau)_\mu, \mu, T)$$

is a standard dynamical system.²⁶

²⁴Note that continuity implies $\varphi^{-1}(B) \in \mathcal{B}(\tau)$ so $\mu(\varphi^{-1}(B))$ is defined.

²⁵Proof: Since τ is a Polish topology on X , $(X, \mathcal{B}(\tau))$ is a standard Borel space with probability measure μ and $T : X \rightarrow X$ is measurable since it is continuous.

²⁶Proof: As above, $(X, \mathcal{B}(\tau), \mu)$ is a Borel probability space, so its completion $(X, \mathcal{B}(\tau)_\mu, \mu)$ is a Lebesgue space. Since T is continuous, it is measurable. And, by assumption, T is bijective and measure-preserving.

A final word on notation: As a rule of thumb (i.e., when feasible), elements of sets are denoted by lower-case letters (like x), sets are denoted by upper-case letters (like Y), sets of sets are denoted by calligraphic letters (like \mathcal{A}), and structures are denoted by Fraktur letters (like \mathfrak{X}).

4.3 Observing dynamical systems

In this section, we describe the structure of possible ways of observing an abstract dynamical system $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$.

4.3.1 Basis or ‘set of possible observations’

Intuitively, possible observations (or measurements or experiments) that we can make about the system \mathfrak{X} correspond to subsets U of the state space X : to make observation U is to realize that the system’s current state is in the set U . The following describes some minimal properties that a set \mathcal{B} —which we’ll call a basis—of all possible observations (under consideration) should satisfy.

4.3.1. DEFINITION. If X is a set, a *basis* for X is a collection \mathcal{B} of subsets of X closed under finite intersection.²⁷ (The empty intersection equals X , so $X \in \mathcal{B}$.) If (X, τ) is a topological space, a *topological basis* for (X, τ) is a basis \mathcal{B} for X consisting of open sets (i.e., $\mathcal{B} \subseteq \tau$) such that the topology generated by \mathcal{B} is τ (i.e., every $U \in \tau$ can be written as unions of elements of \mathcal{B}).²⁸ If (X, \mathcal{A}) is a measurable space (i.e., X a set and \mathcal{A} a σ -algebra on X), a *measurable basis* for (X, \mathcal{A}) is a basis \mathcal{B} for X consisting of measurable sets (i.e., $\mathcal{B} \subseteq \mathcal{A}$).

If X is a set, \mathcal{B} a basis for X , and $T : X \rightarrow X$ a function, we call \mathcal{B} :

- *forward T -closed* if $T\mathcal{B} \subseteq \mathcal{B}$ (i.e., if $U \in \mathcal{B}$, then $T(U) \in \mathcal{B}$).
- *backward T -closed* if $T^{-1}\mathcal{B} \subseteq \mathcal{B}$ (i.e., if $U \in \mathcal{B}$, then $T^{-1}(U) \in \mathcal{B}$).
- *countable* if \mathcal{B} is a countable set.
- *separating* if, for all $x \neq y$ in X , there is $U \in \mathcal{B}$ such that $x \in U$ but $y \notin U$.

This terminology naturally extends to systems: For example, if $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ is an abstract dynamical system, then a backward (T -) closed measurable basis \mathcal{B} for \mathfrak{X} is a measurable basis for (X, \mathcal{A}) that is backward T -closed. Or if $\mathfrak{X} = (X, \tau, \mu, T)$ is standard measured topological system, then a countable topological basis \mathcal{B} for \mathfrak{X} is a topological basis for (X, τ) that is countable.

²⁷I.e., if $B_1, \dots, B_n \in \mathcal{B}$ for $n \geq 0$, then $\bigcap_{k=1}^n B_k \in \mathcal{B}$. If $n = 0$, the convention is that $\bigcup_{k=1}^n B_k := X$.

²⁸Note that the usual definition of a basis for a topology (e.g. Munkres 2000, p. 78) is a bit more general and doesn’t require closure under intersection. But, for our purposes, the chosen notion will be more convenient.

Comments: First, eventually we'll either use topological or measurable bases, but for much of this section we can work with the general (set-theoretic) concept of a basis.

Second, it is tempting to demand (some of) the above additional properties of bases from the start, together with further *logical* closure conditions that \mathcal{B} is not only closed under finite intersection (conjunction), but also under finite union (disjunction) and under complement (negation). We often can build bases with (some of) these properties, but, at this general level, we'll keep our assumptions minimal. In particular, on the 'possible observations' interpretation of a basis, closure under finite intersection means that we can form conjunctions of observations, which is rather uncontroversial. However, other logical combinations, like classical negation or disjunction, may be more controversial, which is why we won't require them (but allow them).

Third, it may also be tempting to additionally demand, in analogy with topological bases, that a measurable basis \mathcal{B} for (X, \mathcal{A}) generates the σ -algebra \mathcal{A} (in the sense that \mathcal{A} is the smallest σ -algebra containing \mathcal{B}). However, there are two reasons not to: First, we don't have to, and second, in complete probability spaces, a countable \mathcal{B} may fail to generate the σ -algebra, but only do so *after* completion. Moreover, in standard Borel spaces, countable separating measurable bases automatically generate (Mackey 1957, thm. 3.3), and, in Lebesgue spaces, countable separating measurable bases generate after completion (de la Rue 1993, thm. 3-4).

Fourth, de la Rue (1993, def. 3-3) uses the term 'basis' for a countable subset of the σ -algebra of a Lebesgue space that separates points. And, as mentioned, Polish spaces with a distinguished basis play an important role for Danos and Garnier (2015) and Dahlqvist, Danos, and Garnier (2016).

Fifth, in the extreme cases, the powerset $\mathcal{P}(X)$ of a set X is always a basis for X , and the only basis for $X = \emptyset$ is $\mathcal{B} = \{\emptyset\}$.

4.3.2 The index set or 'set of observation parameters'

Given a basis (or set of possible observations) for system \mathfrak{X} , we define the index set $I(\mathcal{B})$ to consist of tuples (n, \mathcal{C}) describing the observation parameters: n is the observation length (i.e., the number of update steps that we observe) and \mathcal{C} is the granularity of our observation (i.e., the areas of the state space that we can observe the system to be in).

4.3.2. DEFINITION. Let X be a set and \mathcal{B} a basis for X . A (*finite*) \mathcal{B} -cover \mathcal{C} of X is a (finite) set of elements of \mathcal{B} whose union is X . A \mathcal{B} -cover \mathcal{D} *accurately refines* a \mathcal{B} -cover \mathcal{C} , written $\mathcal{C} \preceq \mathcal{D}$, if

1. Refinement: for all $D \in \mathcal{D}$, there is $C \in \mathcal{C}$ such that $D \subseteq C$, and
2. Accuracy: for all $x \in C \in \mathcal{C}$, there is $D \in \mathcal{D}$ such that $x \in D \subseteq C$.

We define the *index set* $I(\mathcal{B}) := \mathbb{N} \times \text{FCov}(\mathcal{B})$ where \mathbb{N} is the set of non-negative integers with the usual order \leq and $\text{FCov}(\mathcal{B})$ is the set of finite \mathcal{B} -covers of X ordered by \preceq . We often just write I if \mathcal{B} is clear from context. We order I by the product order $\leq \times \preceq$ which we'll also denote \leq .

Comments: First, clause 1 is the usual definition of refinements of (open) covers. However, it will turn out that the additional clause 2 will be crucial for our purposes. To stress its presence we'll speak of *accurate* refinement, but since this is the only notion of refinement that we use, we'll usually omit the term 'accurate'.

Second, in words, clause 1 says that every set of \mathcal{D} is contained in a set of \mathcal{C} , while clause 2 says (in a sense conversely) that every set of \mathcal{C} can be written as a union of sets from \mathcal{D} .

Third, as mentioned, considering finite partitions of a space plays an important role in the characterization of profinite spaces. However, since we're also considering dynamics, we not only need to take into account which *areas* of the state space we can observe (\mathcal{C}), but also for how *long* we're observing the dynamics (n).

4.3.3. LEMMA. *In the notation of definition 4.3.2, $(I(\mathcal{B}), \leq)$ is a nonempty directed preorder.*

PROOF. We first show that (I, \leq) is a preorder. For that we need to show that \leq and \preceq is a preorder (since the product is again a preorder). For \leq this is clear, so we need to show that \preceq is reflexive and transitive. That conditions (1) and (2) are reflexive is clear. So assume $\mathcal{C} \preceq \mathcal{D} \preceq \mathcal{E}$ and show $\mathcal{C} \preceq \mathcal{E}$. Ad (1). Let $E \in \mathcal{E}$. Then there is $D \in \mathcal{D}$ such that $E \subseteq D$. So there is $C \in \mathcal{C}$ such that $D \subseteq C$. Whence $E \subseteq C \in \mathcal{C}$, as needed. Ad (2). Let $x \in C \in \mathcal{C}$. Then there is $D \in \mathcal{D}$ such that $x \in D \subseteq C$. So there is $E \in \mathcal{E}$ such that $x \in E \subseteq D$. So $x \in E \subseteq C$ for $E \in \mathcal{E}$, as needed.

Nonempty: since $X \in \mathcal{B}$, $\{X\}$ is a finite \mathcal{B} -cover, whence $(0, \{X\}) \in I$.

Directed: If $(n, \mathcal{C}), (m, \mathcal{D}) \in I$, consider $(\max(n, m), \mathcal{C} \vee \mathcal{D})$ where $\mathcal{C} \vee \mathcal{D} := \{C \cap D : C \in \mathcal{C}, D \in \mathcal{D}\}$. To show that $(\max(n, m), \mathcal{C} \vee \mathcal{D})$ is in I , we need to show that $\mathcal{C} \vee \mathcal{D}$ is a finite \mathcal{B} -cover: Since \mathcal{B} is closed under intersection, the elements of $\mathcal{C} \vee \mathcal{D}$ are in \mathcal{B} . Clearly, it is finite. And the union of $\mathcal{C} \vee \mathcal{D}$ is X since: for $x \in X$, there is, since \mathcal{C} and \mathcal{D} are covers, some $C \in \mathcal{C}$ with $x \in C$ and some $D \in \mathcal{D}$ with $x \in D$, so $x \in C \cap D \in \mathcal{C} \vee \mathcal{D}$. To show $(n, \mathcal{C}), (m, \mathcal{D}) \sqsubseteq (\max(n, m), \mathcal{C} \vee \mathcal{D})$, we need to show, since $n, m \leq \max(n, m)$, that $\mathcal{C}, \mathcal{D} \preceq \mathcal{C} \vee \mathcal{D}$. We show it for \mathcal{C} since for \mathcal{D} is similar. Ad (1). Let $C \cap D$ be in $\mathcal{C} \vee \mathcal{D}$ with $C \in \mathcal{C}$ and $D \in \mathcal{D}$. Then $C \cap D \subseteq C$ for $C \in \mathcal{C}$, as needed. Ad (2). Let $x \in C \in \mathcal{C}$. Since \mathcal{D} is a cover, there is $D \in \mathcal{D}$ with $x \in D$. Then $x \in C \cap D \subseteq C$ for $C \cap D \in \mathcal{C} \vee \mathcal{D}$. \square

4.3.3 Observed system

Given an observation parameter (n, \mathcal{C}) , the following captures the possible sequences of observations that we can make (realized by a single state).

4.3.4. DEFINITION. Let X be a set with basis \mathcal{B} and $T : X \rightarrow X$ a function. Let $i = (n, \mathcal{C}) \in I(\mathcal{B})$. For $x \in X$ and $t = (U_0, U_1, \dots, U_{n-1}) \in \mathcal{C}^n$, we say x follows t if $T^k(x) \in t(k)$ for all $0 \leq k < n$. We define

$$\mathcal{O}_i(x) := \mathcal{O}_{\mathcal{C}}^n(x) := \{t \in \mathcal{C}^n : x \text{ follows } t\}.$$

We call $\mathcal{O}_i(x)$ the *observation history* that x gives rise to (so each $t \in \mathcal{O}_i(x)$ is an instantiation of that history). Define

$$\mathbf{H}_i := \mathbf{H}_{\mathcal{C}}^n := \{\mathcal{O}_i(x) : x \in X\}$$

We call \mathbf{H}_i the *set of observation histories* of (X, \mathcal{B}, T) .

Comments: First, we can think of t as a trajectory in the transition system $(\mathcal{C}, \rightarrow)$ where $U \rightarrow V$ iff there is $x \in U$ such that $T(x) \in V$. This is the system that we observe when observing the underlying dynamical system \mathfrak{X} through ‘the lens of’ parameter (n, \mathcal{C}) .

Second, note that the empty trajectory ϵ is the only trajectory of length 0 and any $x \in X$ follows ϵ (qua vacuous quantification), so for any $x \in X$, $\mathcal{O}_{\mathcal{C}}^0(x) = \{\epsilon\}$. Also note that \mathbf{H}_i is finite: Since \mathcal{C}^n is finite, also $\mathbf{H}_i = \{\mathcal{O}_i(x) : x \in X\} \subseteq \mathcal{P}(\mathcal{C}^n)$ is finite. Also, \mathbf{H}_i is nonempty if X is nonempty.

Third, we could also consider a ‘cumulative’ definition of $\mathbf{H}_{\mathcal{C}}^n$ as $\{\mathcal{O}_{\mathcal{C}}^k(x) : x \in X, 0 \leq k \leq n\}$ and partially order $\mathbf{H}_{\mathcal{C}}^n$ by $\mathcal{O}_{\mathcal{C}}^k(x) \leq \mathcal{O}_{\mathcal{C}}^l(y)$ iff $k \leq l$ and $\mathcal{O}_{\mathcal{C}}^k(x) = \mathcal{O}_{\mathcal{C}}^k(y)$. It would have a least element $\mathcal{O}_{\mathcal{C}}^0(x)$ (for any $x \in X$), and its maximal elements are $\{\mathcal{O}_{\mathcal{C}}^n(x) : x \in X\}$. Much of our representation result could also be developed with this idea, but the chosen one is simpler.

Fourth, note that the transition system dynamics extends from $(\mathcal{C}, \rightarrow)$ to \mathbf{H}_i . Given observation history $\mathcal{O}_i(x)$ at the current time step, the observation histories that are possible in the next time step are precisely the $\mathcal{O}_i(T(y))$ with $\mathcal{O}_i(y) = \mathcal{O}_i(x)$. We think of this ‘observation dynamics’ as a non-deterministic computational process that assigns each ‘input’ state $\mathcal{O}_i(x)$ the set of possible ‘output’ states $\{\mathcal{O}_i(T(y)) : \mathcal{O}_i(y) = \mathcal{O}_i(x)\}$. Thus, we’re in the setting of powerdomain theory, as described in section 4.2.1. As mentioned there, we can use the Smyth powerdomain \mathbf{P} applied to the special case of a finite discrete order. Concretely, this is described as follows.

4.3.5. LEMMA. *Let X be a nonempty set with basis \mathcal{B} and $T : X \rightarrow X$ a function. Then $\mathbf{P}(\mathbf{H}_i)$ is the finite partial order $(\mathbf{P}(\mathbf{H}_i) \setminus \{\emptyset\}, \supseteq)$ which is a Scott domain and*

$$\begin{aligned} f_i : \mathbf{P}(\mathbf{H}_i) &\rightarrow \mathbf{P}(\mathbf{H}_i) \\ M &\mapsto \{\mathcal{O}_i(T(y)) : \mathcal{O}_i(y) \in M\} \end{aligned}$$

is a well-defined Scott-continuous function.

PROOF. We consider H_i as a nonempty finite and discrete dcpo. Then, as shown in section 4.2.1, the Smyth powerdomain $P(H_i)$ is a Scott domain and f_i is induced by the multi-valued function $\mathcal{O}_i(x) \mapsto \{\mathcal{O}_i(T(y)) : \mathcal{O}_i(y) = \mathcal{O}_i(x)\}$ and hence well-defined and continuous. \square

4.3.4 Refining observations

If we increase observation parameters, $(n, \mathcal{C}) \leq (m, \mathcal{D})$, then we can compare the observations $\mathcal{O}_{\mathcal{D}}^m(x)$ from the finer level to those from the coarser level $\mathcal{O}_{\mathcal{C}}^n(x)$. The following lemma states that we can do this in a functional way:

4.3.6. LEMMA. *Let X be a set with basis \mathcal{B} and $T : X \rightarrow X$ a function. Let $i = (n, \mathcal{C}) \leq (m, \mathcal{D}) = j$ in $I(\mathcal{B})$. Then, for all $x, y \in X$, if $\mathcal{O}_j(x) = \mathcal{O}_j(y)$, then $\mathcal{O}_i(x) = \mathcal{O}_i(y)$.*

PROOF. Let $t \in \mathcal{O}_i(x)$ and show $t \in \mathcal{O}_i(y)$ (the other direction is analogous). So $t = (C_0, \dots, C_{n-1}) \in \mathcal{C}^n$ with $T^k(x) \in C_k \in \mathcal{C}$ for $k = 0, \dots, n-1$. Since $\mathcal{C} \preceq \mathcal{D}$ we have, by the ‘accuracy’ clause (2), that, for $k = 0, \dots, n-1$, there are $D_k \in \mathcal{D}$ with $T^k(x) \in D_k \subseteq C_k$. Moreover, since \mathcal{D} is a cover, there are, for $k = n, \dots, m-1$, some $D_k \in \mathcal{D}$ with $T^k(x) \in D_k$. Let $t' := (D_0, \dots, D_{n-1}, D_n, \dots, D_{m-1})$. So x follows t' , whence $t' \in \mathcal{O}_{\mathcal{D}}^m(x) = \mathcal{O}_{\mathcal{D}}^m(y)$, so y also follows t' . So $T^k(y) \in D_k \subseteq C_k$ for $k = 0, \dots, n-1$. So y follows t , i.e., $t \in \mathcal{O}_i(y)$. \square

Note that here we made crucial use of the ‘accuracy’ clause in our notion of cover refinement (definition 4.3.2).

Due to this lemma, we can define the surjective function $h_{ij} : H_j \rightarrow H_i$ by mapping $\mathcal{O}_j(x)$ to $\mathcal{O}_i(x)$.²⁹ Conveniently, the move to powerdomains to capture the non-deterministic dynamics on the ‘observation system’ also ensures that this function h_{ij} lifts to a projection on the powerdomains:

4.3.7. LEMMA. *Let X be a nonempty set with basis \mathcal{B} and $T : X \rightarrow X$ a function. Let $i \leq j$ in $I(\mathcal{B})$. Then*

$$\begin{aligned} p_{ij} : P(H_j) &\rightarrow P(H_i) \\ M &\mapsto h_{ij}(M) := \{\mathcal{O}_i(x) : \mathcal{O}_j(x) \in M\}, \end{aligned}$$

is a Scott-continuous projection.

PROOF. Write $p := p_{ij}$, $h := h_{ij}$ and define $e : P(H_i) \rightarrow P(H_j)$ by $e(M) := h^{-1}(M)$; since h is surjective, this is indeed a nonempty subset of H_j . Qua image and preimage, we have, for $\emptyset \neq M \subseteq N \subseteq H_j$, that $p(M) \subseteq p(N)$, and for $\emptyset \neq M \subseteq$

²⁹Surjective: Given $\mathcal{O}_i(x) \in H_i$ we have $\mathcal{O}_j(x) \in H_j$, and $h(\mathcal{O}_j(x)) = \mathcal{O}_i(x)$.

$N \subseteq H_i$, that $e(M) \subseteq e(N)$. Hence, p and e are monotone and thus, since $P(H_i)$ and $P(H_j)$ are finite, also Scott-continuous.

To show that (e, p) is an embedding-projection pair, we show $p \circ e = \text{id}_{H_i}$ and $e \circ p \leq \text{id}_{H_j}$. (This is an equivalent way of saying that (e, p) is an embedding-projection pair, see Abramsky and Jung (1994, sec. 3.1.3–4).) We have, for $M \in P(H_i)$, that $p \circ e(M) = h(h^{-1}(M)) = M$ since h is surjective.³⁰ And, for $M \in P(H_j)$, we have $e \circ p(M) = h^{-1}(h(M)) \supseteq M$,³¹ so $e \circ p(M) \leq M$ in $P(H_j)$. \square

In fact, there is a more general statement of this lemma: If $h : Q \rightarrow P$ is a surjective monotone function between two posets with least elements, then $p : P(Q) \rightarrow P(P)$ given by $p(M) := \uparrow h(M)$ is a projection with the embedding $e(M) := h^{-1}(M)$. But for our purposes the above is enough.

4.3.5 Observation probabilities

We show how we can assign probabilities to the possible observations in a domain-theoretic manner using valuations.

We start by defining observational equivalence.

4.3.8. DEFINITION. Let X be a set with basis \mathcal{B} and $T : X \rightarrow X$ a function. For $i = (n, \mathcal{C}) \in I(\mathcal{B})$, we define the i -observational equivalence relation on X by

$$\begin{aligned} x \approx_i y &:\text{iff } \mathcal{O}_i(x) = \mathcal{O}_i(y) \\ &\text{iff } \forall U \in \mathcal{C} \forall k \in \{0, \dots, n-1\} : T^k x \in U \Leftrightarrow T^k y \in U. \end{aligned} \quad ^{32}$$

We denote the equivalence classes $[x]_i := \{y \in X : x \approx_i y\}$.

These equivalence classes are well-behaved in the measurable and topological setting, respectively:

4.3.9. LEMMA. 1. Let (X, \mathcal{A}) be a measurable space, \mathcal{B} a measurable basis, and $T : X \rightarrow X$ a measurable function. Then, for $i \in I(\mathcal{B})$, $[x]_i \in \mathcal{A}$.

³⁰If $b \in M \subseteq H_i$, let, by surjectivity, $a \in H_j$ with $h(a) = b$. Then $a \in h^{-1}(M)$, so $b \in h(h^{-1}(M))$. The other direction is trivial: If $b \in h(h^{-1}(M))$, then $b = h(a)$ for $a \in h^{-1}(M)$, so $b = h(a) \in M$.

³¹If $b \in M$, then $h(b) \in h(M)$, so $b \in h^{-1}(h(M))$.

³²Proof: (\Rightarrow) Assume $\mathcal{O}_i(x) = \mathcal{O}_i(y)$. Let $U \in \mathcal{C}$ and $k \in \{0, \dots, n-1\}$. Assume $T^k x \in U$ and show $T^k y \in U$ (the other direction is analogous). Since \mathcal{C} is a cover, there is $t = (U_0, \dots, U_{k-1}, U, U_{k+1}, \dots, U_{n-1}) \in \mathcal{C}^n$ with $T^l(x) \in U_l$ for $l \in \{0, \dots, n-1\} \setminus \{k\}$. So x follows t , whence $t \in \mathcal{O}_i(x) = \mathcal{O}_i(y)$, so y follows t , so $T^k(y) \in U$.

(\Leftarrow) Let $t \in \mathcal{O}_i(x)$ and show $t \in \mathcal{O}_i(y)$ (the other direction is analogous). Write $t = (U_0, \dots, U_{n-1}) \in \mathcal{C}^n$. Since x follows t , we have, for $k = 0, \dots, n-1$, that $T^k(x) \in U_k$, so, by the assumption, $T^k(y) \in U_k$. So also y follows t , whence $t \in \mathcal{O}_i(y)$.

2. Let (X, τ) be a zero-dimensional topological space, \mathcal{B} a topological basis of clopen sets, and $T : X \rightarrow X$ a continuous function. Then, for $i = (n, \mathcal{C}) \in I(\mathcal{B})$, $[x]_i$ can be written as Boolean combination of sets from $\bigcup_{k=0}^{n-1} T^{-k}\mathcal{C}$ and hence is, in particular, a clopen subset of X .

PROOF. Let's first only assume what is common to both claim (1) and (2): that X is a set with basis \mathcal{B} and $T : X \rightarrow X$ a function. We first describe $[x]_i$ in this general setting with $i = (n, \mathcal{C})$: For $x \in X$, define $\llbracket x \rrbracket^+ := \{t \in \mathcal{C}^n : x \in \bigcap_{k=0}^{n-1} T^{-k}(t(k))\}$ and $\llbracket x \rrbracket^- := (\llbracket x \rrbracket^+)^c$. We claim that

$$[x]_i = \bigcap_{t \in \llbracket x \rrbracket^+} \bigcap_{k=0}^{n-1} T^{-k}(t(k)) \cap \bigcap_{t \in \llbracket x \rrbracket^-} \left(\bigcap_{k=0}^{n-1} T^{-k}(t(k)) \right)^c.$$

Indeed, first note that, by definition of \approx_i , we have for $x, y \in X$:

$$x \approx_{(n, \mathcal{C})} y \Leftrightarrow \forall t \in \mathcal{C}^n : x \in \bigcap_{k=0}^{n-1} T^{-k}(t(k)) \text{ iff } y \in \bigcap_{k=0}^{n-1} T^{-k}(t(k)). \quad (4.2)$$

This is readily seen to imply the claimed identity.³³ Now, the two claims follow:

Concerning claim (1), since $t(k) \in \mathcal{C} \subseteq \mathcal{B} \subseteq \mathcal{A}$ and T is measurable (and hence also its compositions), $T^{-k}(t(k)) \subseteq X$ is in \mathcal{A} . Hence, qua finite intersection, $\bigcap_{k=0}^{n-1} T^{-k}(t(k))$ is in \mathcal{A} , and, qua complement, $\left(\bigcap_{k=0}^{n-1} T^{-k}(t(k)) \right)^c$ is in \mathcal{A} . Since \mathcal{C} is finite, also \mathcal{C}^n is finite, so $\llbracket x \rrbracket^+$ and $\llbracket x \rrbracket^-$ are finite. Hence, $[x]_i$ is a finite intersection of sets in \mathcal{A} and hence in \mathcal{A} , as needed.

Concerning claim (2), since $t(k) \in \mathcal{C}$, each $T^{-k}(t(k))$ is in $\bigcup_{k=0}^{n-1} T^{-k}\mathcal{C}$, so $[x]_i$ is a Boolean combination of sets from $\bigcup_{k=0}^{n-1} T^{-k}\mathcal{C}$. Since $\mathcal{C} \subseteq \mathcal{B}$ is a set of clopen sets and T continuous, each set in $\bigcup_{k=0}^{n-1} T^{-k}\mathcal{C}$ is clopen, whence, qua Boolean combination of such sets, $[x]_i$ is clopen. \square

Given the measurability of the equivalence classes, we can define the following valuation.

³³(\subseteq). Assume $y \in [x]_i$. To show $y \in \bigcap_{t \in \llbracket x \rrbracket^+} \bigcap_{k=0}^{n-1} T^{-k}(t(k))$, let $t \in \llbracket x \rrbracket^+$ be given. Then $t \in \mathcal{C}^n$ and $x \in \bigcap_{k=0}^{n-1} T^{-k}(t(k))$. Since $x \approx_i y$, we have, by (4.2), that $y \in \bigcap_{k=0}^{n-1} T^{-k}(t(k))$.

To show $y \in \bigcap_{t \in \llbracket x \rrbracket^-} \left(\bigcap_{k=0}^{n-1} T^{-k}(t(k)) \right)^c$, let $t \in \llbracket x \rrbracket^-$ be given. Then $t \in \mathcal{C}^n$ and $x \notin \bigcap_{k=0}^{n-1} T^{-k}(t(k))$. Since $x \approx_i y$, we have, by (4.2), that $y \notin \bigcap_{k=0}^{n-1} T^{-k}(t(k))$. So $y \in \left(\bigcap_{k=0}^{n-1} T^{-k}(t(k)) \right)^c$.

(\supseteq). Assume $y \in \bigcap_{t \in \llbracket x \rrbracket^+} \bigcap_{k=0}^{n-1} T^{-k}(t(k))$ and $y \in \bigcap_{t \in \llbracket x \rrbracket^-} \left(\bigcap_{k=0}^{n-1} T^{-k}(t(k)) \right)^c$. To show $x \approx_i y$ via (4.2), let $t \in \mathcal{C}^n$ and show $x \in \bigcap_{k=0}^{n-1} T^{-k}(t(k))$ iff $y \in \bigcap_{k=0}^{n-1} T^{-k}(t(k))$. (\Rightarrow) If $x \in \bigcap_{k=0}^{n-1} T^{-k}(t(k))$, then $t \in \llbracket x \rrbracket^+$, so $y \in \bigcap_{k=0}^{n-1} T^{-k}(t(k))$. (\Leftarrow) If $x \notin \bigcap_{k=0}^{n-1} T^{-k}(t(k))$, then $t \in \llbracket x \rrbracket^-$, so $y \in \left(\bigcap_{k=0}^{n-1} T^{-k}(t(k)) \right)^c$, whence $y \notin \bigcap_{k=0}^{n-1} T^{-k}(t(k))$.

4.3.10. LEMMA. *Let $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ be an abstract dynamical system. Let \mathcal{B} be a measurable basis for \mathfrak{X} . Then, for $i \in I(\mathcal{B})$ and $D_i := \mathbf{P}(\mathbf{H}_i)$, the following defines a function $v_i : \Sigma(D_i) \rightarrow [0, 1]$:*

$$v_i(U) := \sum_{k=1}^m \mu([x_k]_i) \quad \text{if } \exists m \geq 0 \exists x_1, \dots, x_m \in X : \\ \max U = \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\} \text{ and } \mathcal{O}_i(x_k) \neq \mathcal{O}_i(x_l) \text{ for } k \neq l$$

This is a normalized continuous valuation with $v_i(\max D_i) = 1$.

PROOF. We first show that v_i is a well-defined function. First, note that $X \neq \emptyset$ (since $\mu(X) = 1$), so D_i is a Scott domain, and, by lemma 4.3.9, $[x_k]_i$ is in \mathcal{A} , so $\mu([x_k]_i)$ is defined. Moreover, since $\mathcal{O}_i(x_k) \neq \mathcal{O}_i(x_l)$ for $k \neq l$, the equivalence classes $[x_k]_i$ and $[x_l]_i$ are disjoint, so $\sum_{k=1}^m \mu([x_k]_i) = \mu(\bigcup_{k=1}^m [x_k]_i) \in [0, 1]$.

Second, note that, if $U \in \Sigma(D_i)$, then we can find such m, x_1, \dots, x_m : If $U = \emptyset$, then choose $m := 0$. So let $U \neq \emptyset$. Then $\max U$ is, qua upset, a finite nonempty subset of $\max D_i = \{\{\mathcal{O}_i(x) : x \in X\}\}$. So $\max U = \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\}$ for some $m \geq 1$ and $x_1, \dots, x_m \in X$ with $\{\mathcal{O}_i(x_k)\} \neq \{\mathcal{O}_i(x_l)\}$ for $k \neq l$.

Third, to show that the function is independent of the choice of m, x_1, \dots, x_m , let $m, m' \geq 0$ and $x_1, \dots, x_m, x'_1, \dots, x'_{m'} \in X$ with $\mathcal{O}_i(x_k) \neq \mathcal{O}_i(x_l)$ for $k \neq l$ in $\{1, \dots, m\}$ and $\mathcal{O}_i(x'_{k'}) \neq \mathcal{O}_i(x'_{l'})$ for $k' \neq l'$ in $\{1, \dots, m'\}$ and

$$\{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\} = \max U = \{\{\mathcal{O}_i(x'_1)\}, \dots, \{\mathcal{O}_i(x'_{m'})\}\}.$$

Then $m = m'$ and there is a bijection $b : \{1, \dots, m\} \rightarrow \{1, \dots, m\}$ with $\{\mathcal{O}_i(x_k)\} = \{\mathcal{O}_i(x'_{b(k)})\}$. Hence $[x_k]_i = [x'_{b(k)}]_i$. So

$$\sum_{k=1}^m \mu([x_k]_i) = \sum_{k=1}^m \mu([x'_{b(k)}]_i) = \sum_{k=1}^{m'} \mu([x'_k]_i),$$

as needed.

Next we show that v_i is a valuation. Concerning (i), we have $v_i(\emptyset) = \sum_{k=1}^0 \mu([x_k]_i) = 0$.

Concerning (ii), let $U \subseteq V$ and show $v(U) \leq v(V)$. Let $m, m' \geq 0$ and $x_1, \dots, x_m, y_1, \dots, y_{m'} \in X$ with $\max U = \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\}$ (pairwise distinct) and $\max V = \{\{\mathcal{O}_i(y_1)\}, \dots, \{\mathcal{O}_i(y_{m'})\}\}$ (pairwise distinct). Since $U \subseteq V$ are upsets, $\max U \subseteq \max V$. So $m = |\max U| \leq |\max V| = m'$ and we can write

$$\max V = \{\{\mathcal{O}_i(x'_1)\}, \dots, \{\mathcal{O}_i(x'_{m'})\}\}$$

with $(x'_1, \dots, x'_m) = (x_1, \dots, x_m)$ and $\{x'_{m+1}, \dots, x'_{m'}\}$ are the y_k with $\mathcal{O}_i(y_k) \in \max V \setminus \max U$. Hence

$$v_i(U) = \sum_{k=1}^m \mu([x_k]_i) \leq \sum_{k=1}^{m'} \mu([x'_k]_i) = v_i(V).$$

Concerning (iii), let $U, V \in \Sigma(D_i)$ and show $v_i(U \cup V) + v_i(U \cap V) = v_i(U) + v_i(V)$. Let $m, m' \geq 0$ and $x_1, \dots, x_m, y_1, \dots, y_{m'} \in X$ with

$$\begin{aligned} \max U &= \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\} && \text{(pairwise distinct)} \\ \max V &= \{\{\mathcal{O}_i(y_1)\}, \dots, \{\mathcal{O}_i(y_{m'})\}\} && \text{(pairwise distinct)}. \end{aligned}$$

Let $K^+ := \{k \in \{1, \dots, m'\} : \{\mathcal{O}_i(y_k)\} \in \max U\}$ and $K^- := \{k \in \{1, \dots, m'\} : \{\mathcal{O}_i(y_k)\} \notin \max U\}$. Then

$$\begin{aligned} \max(U \cup V) &= \max U \cup \max V = \{\{\mathcal{O}_i(x_l)\}, \{\mathcal{O}_i(y_k)\} : l \in \{1, \dots, m\}, k \in K^-\} \\ \max(U \cap V) &= \max U \cap \max V = \{\{\mathcal{O}_i(y_k)\} : k \in K^+\}. \end{aligned}$$

So, since $K^+ \cup K^- = \{1, \dots, m'\}$, we have

$$\begin{aligned} v_i(U \cup V) + v_i(U \cap V) &= \left(\sum_{k=1}^m \mu[x_k]_i + \sum_{k \in K^-} \mu[y_k]_i \right) + \left(\sum_{k \in K^+} \mu[y_k]_i \right) \\ &= \sum_{k=1}^m \mu[x_k]_i + \sum_{k=1}^{m'} \mu[y_k]_i = v_i(U) + v_i(V). \end{aligned}$$

Finally, we observe that v_i automatically is continuous since $\Sigma(D_i)$ is finite. To see $v_i(D_i) = 1 = v_i(\max D_i)$, note that $\max D_i \in \Sigma(D_i)$ (qua upset of D_i) and write $\max D_i = \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\}$ (pairwise distinct). Note that $X = \bigcup_{k=1}^m [x_k]_i$ since each $[x_k]_i$ is a subset of X and if $x \in X$, then $\{\mathcal{O}_i(x)\} \in \max D_i$, so $\{\mathcal{O}_i(x)\} = \{\mathcal{O}_i(x_k)\}$ for some $k \in \{1, \dots, m\}$, so $x \in [x_k]_i$. Since $\max(\max D_i) = \max D_i$, we have $v_i(D_i) = v_i(\max D_i) = \sum_{k=1}^m \mu[x_k]_i = \mu(\bigcup_{k=1}^m [x_k]_i) = \mu(X) = 1$. \square

4.3.6 Summary

We summarize the preceding results in the following theorem. We also add further properties that will play an important role in section 4.4 below.

4.3.11. THEOREM. *Let $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ be an abstract dynamical system. Let \mathcal{B} be a measurable basis for \mathfrak{X} . For $i \in I = I(\mathcal{B})$, we have*

1. $D_i := \mathbf{P}(\mathbf{H}_i)$ is a finite Scott domain.
2. $v_i : \Sigma(D_i) \rightarrow [0, 1]$ defined by

$$v_i(U) := \sum_{k=1}^m \mu([x_k]_i) \quad \text{where } \max U = \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\}$$

is a normalized continuous valuation with $v_i(\max D_i) = 1$.

3. $f_i : D_i \rightarrow D_i$ with $f_i(M) := \{\mathcal{O}_i(T(x)) : \mathcal{O}_i(x) \in M\}$ is Scott-continuous.

For $i \leq j$ in I , define $p_{ij} : D_j \rightarrow D_i$ by $p_{ij}(M) := \{\mathcal{O}_i(x) : \mathcal{O}_j(x) \in M\}$. Then

4. p_{ij} is a Scott-continuous projection.

5. $p_{ii} : D_i \rightarrow D_i$ is the identity function.

6. $p_{ik} = p_{ij} \circ p_{jk}$ if $i \leq j \leq k$.

7. If $a \in \max D_j$, then $p_{ij}(a) \in \max D_i$.

8. If $a \in D_j$ and $p_{ij}(a) \leq e \in \max D_i$, then there is $a \leq d \in \max D$ with $p_{ij}(d) = e$.

9. For all $V \in \Sigma(D_i)$, $v_i(V) = v_j(p_{ij}^{-1}(V))$.

10. For $a \in \max D_j$, $p_{ij}(f_j(a)) \geq f_i(p_{ij}(a))$.

11. For all $i \in I$, if $\exists a_i, b_i \neq b'_i \in \max D_i : b_i, b'_i \geq f_i(a_i)$, then there is $j \geq i$ in I such that $\forall a_j, b_j, b'_j \in \max D_j : \text{if } p_{ij}(a_j) = a_i, p_{ij}(b_j) = b_i, p_{ij}(b'_j) = b'_i$,³⁴ then $b_j \not\geq f_j(a_j)$ or $b'_j \not\geq f_j(a_j)$.

If, additionally, T is bijective and μ -preserving, and \mathcal{B} is forward T -closed, then

12. for all $i \in I$, if $b \in \max D_i$, then there is $a \in \max D_i$ such that $b \geq f_i(a)$.

13. For all $i \in I$, if $\exists a_i \neq a'_i, b_i \in \max D_i : b_i \geq f_i(a_i), f_i(a'_i)$, then there is $j \geq i$ in I such that $\forall a_j, a'_j, b_j \in \max D_j : \text{if } p_{ij}(a_j) = a_i, p_{ij}(a'_j) = a'_i, p_{ij}(b_j) = b_i$, then $b_j \not\geq f_j(a_j)$ or $b_j \not\geq f_j(a'_j)$.

14. (a) For all $i \in I$ and $U_i \in \Sigma(D_i)$, there is $j_0 \geq i$ such that, for all $j \geq j_0$, we have $v_j(f_j^{-1}(p_{ij}^{-1}(\downarrow \max U_i)) \cap \max D_j) = v_j(p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j)$ ($= v_i(U_i)$).³⁵ (b) For all $i \leq j$ in I , if $a_i, b_i \in \max D_i$ with $f_i(a_i) \leq b_i$, then there is $a_j, b_j \in \max D_j$ such that $p_{ij}(a_j) = a_i$ and $p_{ij}(b_j) = b_i$ and $f_j(a_j) \leq b_j$.

PROOF. Items (1)–(4) are summaries of the preceding results (again, $X \neq \emptyset$ since $\mu(X) = 1$).

Ad (5). We have $p_{ii}(M) = \{\mathcal{O}_i(x) : \mathcal{O}_i(x) \in M\} = M$.

Ad (6). If $a \in p_{ik}(M)$, then $a = \mathcal{O}_i(x)$ with $\mathcal{O}_k(x) \in M$. So $\mathcal{O}_j(x) \in p_{jk}(M)$, whence $a = \mathcal{O}_i(x) \in p_{ij}(p_{jk}(M))$, so $a \in p_{ij} \circ p_{jk}(M)$. Conversely, if $a \in p_{ij} \circ p_{jk}(M)$, then $a = \mathcal{O}_i(x)$ for $b := \mathcal{O}_j(x) \in p_{jk}(M)$. So there is $\mathcal{O}_k(y) \in M$ with $b = \mathcal{O}_j(y)$. So $\mathcal{O}_j(x) = \mathcal{O}_j(y)$, whence also $a = \mathcal{O}_i(x) = \mathcal{O}_i(y)$. So $a \in p_{ik}(M)$.

³⁴Note that this implies $b_j \neq b'_j$ since otherwise $b_i = p_{ij}(b_j) = p_{ij}(b'_j) = b'_i$.

³⁵Note that, qua sets of maximal elements of finite domains, both sets are Scott-open.

Ad (7). If $a \in \max D_j$, then $a = \{\mathcal{O}_j(x)\}$ for some $x \in X$, and $p_{ij}(a) = p_{ij}(\{\mathcal{O}_j(x)\}) = \{\mathcal{O}_i(x)\} \in \max D_i$.

Ad (8). Let $a \in D_j$ and $p_{ij}(a) \leq e \in \max D_i$. So $e = \{\mathcal{O}_i(x)\} \subseteq \{\mathcal{O}_i(y) : \mathcal{O}_j(y) \in a\}$. So $\mathcal{O}_i(x) = \mathcal{O}_i(y)$ for some $\mathcal{O}_j(y) \in a$. Set $d := \{\mathcal{O}_j(y)\}$. Then $a \leq d \in \max D_j$ and $p_{ij}(d) = \{\mathcal{O}_i(y)\} = \{\mathcal{O}_i(x)\} = e$.

Ad (9). Let $V \in \Sigma(D_i)$. Let

$$\begin{aligned} \max V &= \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\} && \text{(pairwise distinct)} \\ \max p_{ij}^{-1}(V) &= \{\{\mathcal{O}_j(y_1)\}, \dots, \{\mathcal{O}_j(y_{m'})\}\} && \text{(pairwise distinct)}. \end{aligned}$$

We claim that

$$\bigcup_{k=1}^m [x_k]_i = \bigcup_{l=1}^{m'} [y_l]_j$$

(\subseteq) Let $z \in [x_k]_i$. Then $\{\mathcal{O}_j(z)\} \in \max D_j$ and $p_{ij}(\{\mathcal{O}_j(z)\}) = \{\mathcal{O}_i(z)\} = \{\mathcal{O}_i(x_k)\} \in V$. So $\{\mathcal{O}_j(z)\} \in \max p_{ij}^{-1}(V)$, whence $\{\mathcal{O}_j(z)\} = \{\mathcal{O}_j(y_l)\}$ for some $l \in \{1, \dots, m'\}$, so $z \in [y_l]_j$.

(\supseteq) Let $z \in [y_l]_j$. Then $\{\mathcal{O}_j(z)\} = \{\mathcal{O}_j(y_l)\}$, so

$$\{\mathcal{O}_i(z)\} = p_{ij}(\{\mathcal{O}_j(z)\}) = p_{ij}(\{\mathcal{O}_j(y_l)\}) \in V.$$

Since $\{\mathcal{O}_i(z)\} \in \max D_i$, it is in $\max V$. So $\{\mathcal{O}_i(z)\} = \{\mathcal{O}_i(x_k)\}$ for some $k \in \{1, \dots, m\}$. So $z \in [x_k]_i$.

Then

$$v_i(V) = \sum_{k=1}^m \mu[x_k]_i = \mu \bigcup_{k=1}^m [x_k]_i = \mu \bigcup_{l=1}^{m'} [y_l]_j = \sum_{l=1}^{m'} \mu[y_l]_j = v_j(p_{ij}^{-1}(V)).$$

Ad (10). We actually prove the claim for any $a \in D_j$. We have

$$\begin{aligned} p_{ij}(f_j(a)) &= \{\mathcal{O}_i(y) : \mathcal{O}_j(y) \in f_j(a)\} \\ &= \{\mathcal{O}_i(y) : \mathcal{O}_j(y) \in \{\mathcal{O}_j(Tx) : \mathcal{O}_j(x) \in a\}\} =: A \\ f_i(p_{ij}(a)) &= \{\mathcal{O}_i(Tx) : \mathcal{O}_i(x) \in p_{ij}(a)\} =: B \end{aligned}$$

So we need to show that $A \subseteq B$ (then $A \geq B$). If $u \in A$, then $u = \mathcal{O}_i(y)$ for $\mathcal{O}_j(y) \in \{\mathcal{O}_j(Tx) : \mathcal{O}_j(x) \in a\}$. So there is $\mathcal{O}_j(x) \in a$ such that $\mathcal{O}_j(y) = \mathcal{O}_j(Tx)$. So $\mathcal{O}_i(x) \in p_{ij}(a)$ and $u = \mathcal{O}_i(y) = \mathcal{O}_i(Tx) \in B$.³⁶

³⁶ Here is why the straightforward proof for the other direction ($A \supseteq B$) doesn't go through: If $u \in B$, then $u = \mathcal{O}_i(Tx)$ for some $\mathcal{O}_i(x) \in p_{ij}(a)$. So $\mathcal{O}_i(x) = \mathcal{O}_i(z)$ for some $\mathcal{O}_j(z) \in a$. So $\mathcal{O}_j(Tz) \in \{\mathcal{O}_j(Tx) : \mathcal{O}_j(x) \in a\}$, whence $\mathcal{O}_i(Tz) \in A$. However, even if $\mathcal{O}_i(x) = \mathcal{O}_i(z)$, it is not clear that we have $\mathcal{O}_i(Tx) = \mathcal{O}_i(Tz)$.

Ad (11). Let $i = (n, \mathcal{C}) \in I$ and $a_i = \{\mathcal{O}_i(x_i)\}, b_i = \{\mathcal{O}_i(y_i)\}, b'_i = \{\mathcal{O}_i(y'_i)\}$ (for $x_i, y_i, y'_i \in X$) be in $\max D_i$ with $b_i \neq b'_i \geq f_i(a_i)$. So $\mathcal{O}_i(y_i), \mathcal{O}_i(y'_i) \in \{\mathcal{O}_i(Tz_i) : \mathcal{O}_i(z_i) \in a_i\}$. So there is $z_i, z'_i \in X$ with $\mathcal{O}_i(z_i) = \mathcal{O}_i(x_i) = \mathcal{O}_i(z'_i)$ and $\mathcal{O}_i(y_i) = \mathcal{O}_i(Tz_i)$ and $\mathcal{O}_i(y'_i) = \mathcal{O}_i(Tz'_i)$.

We have $i = (n, \mathcal{C}) \leq (n+1, \mathcal{C}) =: j \in I$. To see that j has the required properties, let $a_j = \{\mathcal{O}_j(x_j)\}, b_j = \{\mathcal{O}_j(y_j)\}, b'_j = \{\mathcal{O}_j(y'_j)\}$ (for $x_j, y_j, y'_j \in X$) be in $\max D_i$ with $p_{ij}(a_j) = a_i, p_{ij}(b_j) = b_i, p_{ij}(b'_j) = b'_i$. Assume for contradiction that $b_j, b'_j \geq f_j(a_j)$.

First, note that $\mathcal{O}_i(y_j) = \mathcal{O}_i(Tz_i)$ and $\mathcal{O}_i(y'_j) = \mathcal{O}_i(Tz'_i)$: Indeed, we have

$$\{\mathcal{O}_i(Tz_i)\} = \{\mathcal{O}_i(y_i)\} = b_i = p_{ij}(b_j) = p_{ij}(\{\mathcal{O}_j(y_j)\}) = \{\mathcal{O}_i(y_j)\},$$

and similarly for y'_j .

Second, since $b_j, b'_j \geq f_j(a_j)$, we have, as above, $z_j, z'_j \in X$ with $\mathcal{O}_j(z_j) = \mathcal{O}_j(x_j) = \mathcal{O}_j(z'_j)$ and $\mathcal{O}_j(y_j) = \mathcal{O}_j(Tz_j)$ and $\mathcal{O}_j(y'_j) = \mathcal{O}_j(Tz'_j)$. The latter two identities imply, together with the first observation, $\mathcal{O}_i(Tz_i) = \mathcal{O}_i(y_j) = \mathcal{O}_i(Tz_j)$ and $\mathcal{O}_i(Tz'_i) = \mathcal{O}_i(y'_j) = \mathcal{O}_i(Tz'_j)$.

Third, we claim $\mathcal{O}_i(Tz_j) = \mathcal{O}_i(Tz'_j)$. Indeed, let $t = (C_0, \dots, C_{n-1}) \in \mathcal{C}$ with Tz_j following t , i.e., $T^k(Tz_j) \in C_k$ for $k = 0, \dots, n-1$. Let $C \in \mathcal{C}$ with $z_j \in C$. Then z_j follows $t' = (C, C_0, \dots, C_{n-1}) \in \mathcal{C}^{n+1}$. Since $\mathcal{O}_j(z_j) = \mathcal{O}_j(z'_j)$, also z'_j follows t' . In particular, for $k = 0, \dots, n-1$, we have $T^k(Tz'_j) \in C_k$. So Tz'_j also follows t . Similarly for the other direction.

Finally, putting everything together, we obtain

$$\mathcal{O}_i(Tz_i) = \mathcal{O}_i(Tz_j) = \mathcal{O}_i(Tz'_j) = \mathcal{O}_i(Tz'_i)$$

which contradicts $\{\mathcal{O}_i(Tz_i)\} = b_i \neq b'_i = \{\mathcal{O}_i(Tz'_i)\}$.

For the last three items, we now assume T to be bijective and measure-preserving, and \mathcal{B} to be forward T -closed.

Ad (12). Let $b = \{\mathcal{O}_i(y)\}$ be in $\max D_i$ (for some $y \in X$). Since T is surjective, let $x \in X$ with $T(x) = y$. Let $a := \{\mathcal{O}_i(x)\} \in \max D_i$. We have $b = \{\mathcal{O}_i(y)\} = \{\mathcal{O}_i(Tx)\} \subseteq \{\mathcal{O}_i(Tx) : \mathcal{O}_i(x) \in a\} = f_i(a)$, so $b \geq f_i(a)$, as needed.

Ad (13). Let $i = (n, \mathcal{C}) \in I$ and $a_i = \{\mathcal{O}_i(x_i)\}, a'_i = \{\mathcal{O}_i(x'_i)\}, b_i = \{\mathcal{O}_i(y_i)\}$ (for $x_i, x'_i, y_i \in X$) be in $\max D_i$ with $a_i \neq a'_i$ and $b_i \geq f_i(a_i), f_i(a'_i)$.

Since $a_i \neq a'_i$, we have $x_i \not\approx_i x'_i$, so there is, without loss of generality (the other case is analogous), $U \in \mathcal{C}$ and $k \in \{0, \dots, n-1\}$ (so $n \geq 1$) such that $T^k(x_i) \in U$ but $T^k(x'_i) \notin U$. Since $U \in \mathcal{C} \subseteq \mathcal{B}$ and \mathcal{B} is T -closed, $T(U) \in \mathcal{B}$, so $j_0 := (n, \{T(U), X\}) \in I$. By directedness, let $j \geq i, j_0$ be in I .

To see that j has the required properties, assume for contradiction that there are $a_j = \{\mathcal{O}_j(x_j)\}, a'_j = \{\mathcal{O}_j(x'_j)\}, b_j = \{\mathcal{O}_j(y_j)\}$ (for $x_j, x'_j, y_j \in X$) in $\max D_j$ with $p_{ij}(a_j) = a_i, p_{ij}(a'_j) = a'_i, p_{ij}(b_j) = b_i$ but $b_j \geq f_j(a_j), f_j(a'_j)$.

The former implies $\mathcal{O}_i(x_j) = \mathcal{O}_i(x_i), \mathcal{O}_i(x'_j) = \mathcal{O}_i(x'_i), \mathcal{O}_i(y_j) = \mathcal{O}_i(y_i)$. The latter implies $\mathcal{O}_j(y_j) = \mathcal{O}_j(Tz_j)$ for some $\mathcal{O}_j(z_j) = \mathcal{O}_j(x_j)$ and $\mathcal{O}_j(y_j) = \mathcal{O}_j(Tz'_j)$

for some $\mathcal{O}_j(z'_j) = \mathcal{O}_j(x'_j)$. Since $i \leq j$ we, in particular, have $\mathcal{O}_i(z_j) = \mathcal{O}_i(x_j) = \mathcal{O}_i(x_i)$ and $\mathcal{O}_i(z'_j) = \mathcal{O}_i(x'_j) = \mathcal{O}_i(x'_i)$.

Since $U \in \mathcal{C}$, $k \in \{0, \dots, n-1\}$, and $x_i \approx_i z_j$, the fact that $T^k x_i \in U$ hence implies $T^k z_j \in U$. Similarly, $T^k x'_i \notin U$ implies $T^k z'_j \notin U$. Thus, $T^k T z_j \in T(U)$. And we cannot have $T^k T z'_j \in T(U)$ since otherwise $T^k T z'_j = T u$ for some $u \in U$, so, by injectivity of T , $T^k z'_j = u \in U$. Hence $T z_j$ and $T z'_j$ can be separated in j_0 , so $\mathcal{O}_{j_0}(T z_j) \neq \mathcal{O}_{j_0}(T z'_j)$, so, since $j_0 \leq j$,

$$\mathcal{O}_j(y_j) = \mathcal{O}_j(T z_j) \neq \mathcal{O}_j(T z'_j) = \mathcal{O}_j(y_j),$$

which is a contradiction.

Ad (14). (a). Let $(n, \mathcal{C}) = i \in I$ and $U_i \in \Sigma(D_i)$. Without loss of generality, $U_i \neq \emptyset$ (otherwise both evaluated sets are empty, and hence both have the value 0). Write $\max U_i = \{\{\mathcal{O}_i(y_1)\}, \dots, \{\mathcal{O}_i(y_r)\}\}$ for $y_1, \dots, y_r \in X$ ($r \geq 1$). Let $j_0 := (n+1, \mathcal{C}) \geq i$. To show that this has the required property, let $j \geq j_0$, and show $v_i(f_j^{-1}(p_{ij}^{-1}(\downarrow \max U_i)) \cap \max D_j) = v_j(p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j) = v_i(U_i)$.

We first show the second equality: Since, by (7), p_{ij} preserves maximality, we have $v_j(p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j) = v_j(p_{ij}^{-1}(\max U_i) \cap \max D_j)$. Since, by (2), $v_j(\max D_j) = 1$ and v_j is a normalized valuation, this further equals $v_j(p_{ij}^{-1}(\max U_i))$.³⁷ By (9), this equals $v_i(\max U_i)$. Again, since v_i is normalized, $v_i(\max U_i) = v_i(U_i \cap \max D_i) = v_i(U_i)$.

Concerning the first equality, write $\max D_j = \{\{\mathcal{O}_j(x_1)\}, \dots, \{\mathcal{O}_j(x_m)\}\}$ (for some $m \geq 1$), and let

$$K := \{k \in \{1, \dots, m\} : \exists s \in \{1, \dots, r\} \exists z \in [x_k]_j. T(z) \in [y_s]_i\}.$$

We claim that $T^{-1}(\bigcup_{s=1}^r [y_s]_i) = \bigcup_{k \in K} [x_k]_j$. Indeed, if $x \in T^{-1}(\bigcup_{s=1}^r [y_s]_i)$, then, since the $[x_k]_j$'s partition X , there is $k \in \{1, \dots, m\}$ with $x \in [x_k]_j$, so it remains to show $k \in K$: we have $z := x \in [x_k]_j$ with $T(z) = T(x) \in [y_s]_i$ for some $s \in \{1, \dots, r\}$. Conversely, if $x \in [x_k]_j$ for some $k \in K$, then there is $s \in \{1, \dots, r\}$ and $z \in [x_k]_j$ with $T(z) \in [y_s]_i$. In particular, $z \approx_j x$. Since $j \geq j_0$ this implies $z \approx_{j_0} x$. Since $j_0 = (n+1, \mathcal{C})$, we have

$$\forall C \in \mathcal{C} \forall k = 0, \dots, n : T^k x \in C \Leftrightarrow T^k z \in C.$$

Since $T(z) \approx_i y_s$ and $i = (n, \mathcal{C})$, we have

$$\forall C \in \mathcal{C} \forall k = 0, \dots, n-1 : T^k(Tx) \in C \Leftrightarrow T^k(Tz) \in C \Leftrightarrow T^k(y_s) \in C.$$

Hence $T(x) \in [y_s]_i$, so $x \in T^{-1} \bigcup_{s=1}^r [y_s]_i$.

³⁷ If $v : \Sigma(D) \rightarrow [0, 1]$ is a normalized valuation and $A, B \in \Sigma(D)$ with $v(B) = 1$, then $v(A) = v(A \cap B)$. Proof: By modularity, $v(A \cup B) + v(A \cap B) = v(A) + v(B)$. Since v is normalized and monotone, $1 \geq v(A \cup B) \geq v(B) = 1$, so $v(A \cup B) = 1$. Since also $v(B) = 1$, the modularity yields $v(A \cap B) = v(A)$.

Next, we write $U_j := f_j^{-1}(p_{ij}^{-1}(\downarrow \max U_i)) \cap \max D_j$. We claim that $U_j = \{\{\mathcal{O}_j(x_k)\} : k \in K\}$. Indeed, for any $\{\mathcal{O}_j(x_k)\} = a \in \max D_j$ we have: $a \in U_j$ iff $\exists s \in \{1, \dots, r\} : p_{ij}(f_j(a)) \leq \{\mathcal{O}_i(y_s)\}$ iff $\exists s \in \{1, \dots, r\} : \{\mathcal{O}_i(Tz) : \mathcal{O}_j(z) = \mathcal{O}_j(x_k)\} \supseteq \{\mathcal{O}_i(y_s)\}$ iff $\exists s \in \{1, \dots, r\} \exists z \in [x_k]_j : T(z) \in [y_s]_i$ iff $k \in K$ iff $a \in \{\{\mathcal{O}_j(x_k)\} : k \in K\}$.

Hence, since T is μ -preserving, we have

$$\begin{aligned} v_j(f_j^{-1}(p_{ij}^{-1}(\downarrow \max U_i)) \cap \max D_j) &= v_j(U_j) = \sum_{k \in K} \mu([x_k]_j) = \mu\left(\bigcup_{k \in K} [x_k]_j\right) \\ &= \mu(T^{-1} \bigcup_{s=1}^r [y]_i) = \mu\left(\bigcup_{s=1}^r [y]_i\right) = \sum_{s=1}^r \mu([y]_i) = v_i(U_i), \end{aligned}$$

which, by the already established second equality, equals $v_j(p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j)$, as needed.

(b). Let $i \leq j$ in I and let $a_i = \{\mathcal{O}_i(x)\}$ and $b_i = \{\mathcal{O}_i(y)\}$ be in $\max D_i$ with $f_i(a_i) \leq b_i$. Then $\{\mathcal{O}_i(y)\} \subseteq \{\mathcal{O}_i(Tz) : \mathcal{O}_i(z) = \mathcal{O}_i(x)\}$. So there is $z \in X$ with $\mathcal{O}_i(y) = \mathcal{O}_i(Tz)$ and $\mathcal{O}_i(z) = \mathcal{O}_i(x)$. Choose $a_j := \{\mathcal{O}_j(z)\}$ and $b_j = \{\mathcal{O}_j(Tz)\}$ in $\max D_j$. Then $p_{ij}(a_j) = \{\mathcal{O}_i(z)\} = \{\mathcal{O}_i(x)\} = a_i$ and $p_{ij}(b_j) = \{\mathcal{O}_i(Tz)\} = \{\mathcal{O}_i(y)\} = b_i$ and $b_j = \{\mathcal{O}_j(Tz)\} \subseteq \{\mathcal{O}_j(Tw) : \mathcal{O}_j(w) = \mathcal{O}_j(z)\} = f_j(a_j)$, so $f_j(a_j) \leq b_j$, as needed. \square

Note that, in the proofs of the last three items, we've only assumed T to be surjective for (12), we've only assumed T to be injective and \mathcal{B} to be forward T -closed for (13), and we've only assumed T to be measure-preserving for (14). So we could be more precise and specify more classes of dynamical systems between general and standard (injective, surjective, measure-preserving) and link them to the respective properties above. However, to avoid introducing even more distinctions, we won't do this explicitly.

4.4 Dynamical domains

This section is written from a purely domain-theoretic perspective: motivating a certain category of domains in a domain-theoretic way. However, given the last section, many definitions should be natural (additionally highlighted by using the same notation). So the domain-theoretic definitions could also be motivated by dynamical systems.

4.4.1 Dynamical dcpos

We first fix some terminology. Given a dcpo D , recall that $\max D$ is the set of maximal elements of D . We call a function $f : D \rightarrow E$ between dcpos *max-preserving* if $f(\max D) \subseteq \max E$ (i.e., if $a \in \max D$, then $f(a) \in \max E$).

Recall that, intuitively, (normalized) continuous valuations are domain-theoretic analogues of (probability) measures. We call a valuation v on D *max-normalized* if v is normalized and $\max D$ can be written as a countable intersection of Scott-open sets with v -value 1.³⁸ This condition captures the idea that the valuation is a domain-theoretic description of a probability measure on $\max D$, i.e., the space that the domain D is a computational model for.

Note that a normalized valuation v on a finite dcpo D is max-normalized iff $v(\max D) = 1$ (since D is finite, $\max D$ is Scott-open).³⁹

4.4.1. DEFINITION. A *dynamical dcpo* is a triple $\mathfrak{D} = (D, v, f)$ where D is a dcpo, $v : \Sigma(D) \rightarrow [0, \infty]$ is a continuous valuation, and $f : D \rightarrow D$ is Scott-continuous. We call \mathfrak{D} :

- *finite* if D is finite
- *max-normalized* if v is max-normalized
- *max-preserving* if f is max-preserving
- a *dynamical Scott domain* if D is a Scott domain
- *max-surjective* if, for all $b \in \max D$, there is $a \in \max D$ such that $b \geq f(a)$.
- *valuation-preserving* if, for all $U \in \Sigma(D)$, $v(f^{-1}(U)) = v(U)$.

Comments: First, this definition of a dynamical dcpo mimics that of a dynamical system: D is like the state space, v like the (probability) measure, and f like the transformation. However, there is an important difference: a dcpo contains both ‘real’ and ‘ideal’ elements while the state space of a dynamical system only contains ‘ideal’ points.

Second, in general we don’t require f to be max-preserving (i.e., ‘idealness’-preserving). So f may map an ideal (or informationally complete) state to only a non-ideal (or informationally non-complete) state. As we’ve seen in the (D_i, v_i, f_i) coming from the finite observations of a dynamical system (section 4.3), this may be interpreted as non-determinism: Even given the knowledge of the maximal state $a = \{\mathcal{O}_i(x)\}$, there still is some uncertainty which of the $\{\mathcal{O}_i(T(y))\}$ for

³⁸A simple example of a dcpo (in fact, Scott domain) with a normalized continuous valuation that is not max-normalized is $D := 2 = \{0, 1\}$ with the natural order and $v : \Sigma(D) \rightarrow [0, 1]$ given by $v(\emptyset) = 0$, $v(\{1\}) = \frac{1}{2}$, and $v(\{0, 1\}) = 1$.

³⁹Proof: If $v(\max D) = 1$, then $\max D$ can trivially be written as a countable intersection of Scott-open sets with v -value 1 (and v is already assumed to be normalized). Conversely, if v is max-normalized, then $\max D$ can be written as a countable intersection of Scott-open sets with v -value 1. Since there are only finitely many Scott-opens (since D is finite), this, in fact, is a finite intersection. As in footnote 37, for a normalized valuation the modularity condition implies that a finite intersection of sets of v -value 1 has v -value 1.

$\mathcal{O}_i(y) = \mathcal{O}_i(x)$ will be the successor state, so f only determines an informationally non-complete successor state $f(a)$.

Third, and dually, f being max-preserving may be interpreted as being deterministic: from the complete information about the current state (a maximal element a of D), the dynamics determines the complete information about the successor state (the maximal element $f(a)$). If \mathfrak{D} is to be a computational model for a dynamical system, we want that $(\max D, f \upharpoonright \max D)$ is a dynamical system, so we need max-preservation of f and the valuation should be like a probability measure on $\max D$ (which is the max-normalized condition).

Fourth, the definition of a dynamical dcpo is an instance of the more general idea of adding additional *domain-theoretic* structure to the order-theoretic (and induced topological) structure of a domain: here a (domain-theoretic) valuation and a (domain-theoretic) function.

Next, we define morphisms between dynamical dcpos. Some parts of the definition are not obvious, but we'll offer some (preliminary) explanation below.

4.4.2. DEFINITION. A *dynamical morphism* $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ between two dynamical dcpos $\mathfrak{D} = (D, v, f)$ and $\mathfrak{E} = (E, w, g)$ is a function $\alpha : D \rightarrow E$ such that

1. Scott-continuous: α is Scott-continuous.
2. Max-preserving: α is max-preserving.
3. Max-bisimulative: For all $a \in D$ and $e \in \max E$, if $\alpha(a) \leq e$, then there is $d \in \max D$ such that $d \geq a$ and $\alpha(d) = e$.⁴⁰
4. Valuation-preserving: For all $V \in \Sigma(E)$, $w(V) = v(\alpha^{-1}(V))$.
5. Max-semi-equivariant: For all $a \in \max D$, $\alpha(f(a)) \geq g(\alpha(a))$.

Note that, if $g : D \rightarrow D$ is max-preserving, then α is *max-equivariant*, i.e., for all $a \in \max D$, $\alpha(f(a)) = g(\alpha(a))$.

Among the obvious conditions are: Scott-continuity as *the* requirement for a morphism between domains, and valuation-preserving as the straightforward counterpart to measure-preservation of morphisms between dynamical systems.

Among the 'semi'-obvious conditions are: max-semi-equivariance as 'half' of the equivariance of morphisms of dynamical systems, and max-preserving as being 'deterministic' (in the above sense of mapping an 'ideal information-complete' state to another such state). Concerning the former, an indication for why we should only expect this inequality rather than full equality is the following: if f is max-preserving but g is not (which may happen), then $f(a)$, and hence also

⁴⁰Note that we can equivalently *only* require that $d \in D$ (instead of $d \in \max D$). We can then choose $d' \in \uparrow d \cap \max D$ and we get that $d' \in \max D$ with $d' \geq a$ and $\alpha(d') = e$ (since e is maximal and α monotone).

$\alpha(f(a))$, is maximal, but $g(\alpha(a))$ may fail to be maximal (even though $\alpha(a)$ is), so they cannot be identical. More concretely, in the dynamical dcpos coming from observing a dynamical system, we saw in theorem 4.3.11 that we could only expect this ‘semi-max-equivariance’ of the projections $p_{ij} : D_j \rightarrow D_i$ (especially footnote 36).

The least obvious may be the max-bisimulation condition. Presumably, its best explanation—although unsatisfactory at this stage—is that it makes things work. We’ll see it being used in several places below, but its *raison d’être* will only become completely apparent in the next chapter in proving that the ‘observation domain’ construction is adjoint to the ‘model of a dynamical domain’ construction.

Now, with dynamical dcpos and their morphisms at hand, we can define their categories.

4.4.3. PROPOSITION. *The following define categories:*

1. **dDCP**: *the objects are dynamical dcpos and the morphisms are the dynamical morphisms. The identity morphism is the identity function and morphism composition is function composition.*
2. **dDCP^P**: *the wide (also called lluf) subcategory of dDCP with the same objects and as morphisms those dynamical morphisms that also are projections.*
3. **dSCO**: *the full subcategory of dDCP with dynamical Scott domains as objects.*
4. **dSCO^P**: *the full subcategory of dDCP^P with dynamical Scott domains as objects.*

We can build further categories indicated by the following suffixed subscripts:

- f** *restricting to finite dynamical dcpos.*
- m** *restricting to max-preserving dynamical dcpos.*
- n** *restricting to max-normalized dynamical dcpos.*

For example, **dSCO_{nf}^P** is the full subcategory of **dSCO^P** consisting of max-normalized finite dynamical Scott domains.

PROOF. Ad (1). For dynamical dcpos $\mathfrak{D} = (D, v, f)$, $\mathfrak{E} = (E, w, g)$, $\mathfrak{F} = (F, u, h)$, and dynamical morphisms $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ (so α is a function from D to E) and $\beta : \mathfrak{E} \rightarrow \mathfrak{F}$ (so β is a function from E to F), let $\beta \circ \alpha$ be the composition of the two functions α and β . This is again a dynamical morphism: Ad (1), a composition of (Scott) continuous functions is again (Scott) continuous. Ad (2), if $a \in \max D$, then $\alpha(a) \in \max E$, so $\beta \circ \alpha(a) = \beta(\alpha(a)) \in \max F$. Ad (3), if $a \in D$ and $f \in \max F$ with $\beta \circ \alpha(a) \leq f$, then (since $\alpha(a) \in E$ and $f \in \max F$ with $\beta(\alpha(a)) \leq f$) there is $e \in \max E$ such that $e \geq \alpha(a)$ and $\beta(e) = f$. Hence there is

$d \in \max D$ such that $d \geq a$ and $\alpha(d) = e$, whence also $\beta \circ \alpha(d) = \beta(\alpha(d)) = f$. Ad (4), for $W \in \Sigma(F)$,

$$u(W) = w(\beta^{-1}(W)) = v\left((\alpha^{-1}(\beta^{-1}(W)))\right) = v((\beta \circ \alpha)^{-1}(W)).$$

Ad (5), we have, for $a \in \max D$, that (since α is max-semi-equivariant): $\alpha(f(a)) \geq g(\alpha(a))$. Since β is monotone, we get

$$(\beta \circ \alpha)(f(a)) = \beta\left(\alpha(f(a))\right) \geq \beta\left(g(\alpha(a))\right).$$

Since α is max-preserving, $\alpha(a) \in \max E$. So, since β is max-semi-equivariant, this continues by

$$\geq h\left(\beta(\alpha(a))\right) = h((\beta \circ \alpha)(a)),$$

as needed.

This composition operation satisfies, qua function composition, the associativity axiom, and also the identity axiom: For a dynamical dcpo $\mathfrak{D} = (D, v, f)$, it is readily seen that the identity function $\text{id}_{\mathfrak{D}} := \text{id}_D : D \rightarrow D$ is a dynamical morphism. And we clearly have $\text{id}_{\mathfrak{D}} \circ \alpha = \alpha$ (resp. $\beta \circ \text{id}_D = \beta$) for dynamical morphism α (resp. β) with target (resp. source) \mathfrak{D} .

Ad (2). $\mathbf{dDCP}^{\mathbf{p}}$ is indeed a wide subcategory of \mathbf{dDCP} since the identity function in particular is a projection and the composition of projections is again a projection.

Ad (3) and (4), this is immediate. □

4.4.2 Dynamical expanding systems

To construct the dynamical domains that we'll be working with, we proceed in a way that is common in domain theory: roughly, they are defined as the limits of certain diagrams of finite domains. So they are, as mentioned in the introduction, certain 'profinite' objects. We first recall this standard construction (remark 4.4.4) and explain why there is an additional twist to it in our case (remark 4.4.5 and definition 4.4.6). Then we present the definition of the appropriate diagrams of finite domains (definition 4.4.7).

4.4.4. REMARK ('Profinite' domains). SFP domains (Plotkin 1976) and, more generally, bifinite domains (see Abramsky and Jung 1994), are defined as follows.⁴¹ We start with a 'background' category \mathbf{C} of domains. Adding a superscript \mathbf{p} means considering the wide subcategory $\mathbf{C}^{\mathbf{p}}$ where the morphisms are additionally

⁴¹The acronym SFP stands for Sequences of Finite inductive Partial orders.

required to be projections. Adding the subscript f means restricting to the full subcategory C_f of finite domains. Both for bifinite and SEP domains, C is the category of dcpos with least element together with Scott-continuous functions. For us, $C = \text{dSCO}_n$.

Next, we define a notion of expanding system of finite domains as a diagram in the category C_f^p with possibly additional properties. In the case of bifinite domains, these are the diagrams in the category C_f^p with a directed index set. For SFP domains, the index set is additionally required to be the naturals \mathbb{N} . For us, this will be the diagrams in the category dSCO_{nf}^p with countable directed index set that have an additional property that we call ‘upward deterministic’ (and some more in the ‘standard case’).

Finally, we show that any such expanding system has a limit in the category C^p , and define the category D of domains (that we’re interested in) as the full subcategory of C whose objects are those domains that are the limits of the expanding systems of finite domains. This way one obtains the bifinite and SEP domains, respectively. However, for us, there will be an additional twist: we will have to take a ‘restricted limit’ as we’ll explain next.

4.4.5. REMARK (Intuition for ‘restricted’ limit). The intuition for the definition of D is that the domains in D are those that are completely described in a finitary way: If D is in D , i.e., a limit of an expanding system of finite domains, then the expanding system is the finitary description of D . And the fact that D is the limit means that D is precisely described by that expanding system. (Being a cone of the diagram means that D contains all the information in the diagram, and being limiting means that D also doesn’t contain more information.)

For our purposes, however, we not only want that the domains that we work with have such a complete finitary description, we also need them to model a deterministic dynamical system. So we want that (at least) they are max-preserving dynamical Scott domains. (We’ll later see that this is enough to imply that they model a deterministic system in a much stronger sense.)

Thus, we define the objects of D as the limits of expanding systems of finite domains *subject to* the condition of being max-preserving. So, they are not cones that are limiting among *all* the cones for the expanding systems of finite domains, but they rather are limiting among all the cones built with max-preserving dynamical Scott domains.

Formally, this idea of a ‘restricted’ limit is spelled out in a category-theoretic way in the following definition.⁴² Technically, it is not necessary to understand the chapter: it can be replaced with the concrete description of the restricted limit provided in theorem 4.4.8 below. However, conceptually it is important as it formalizes an important part of the intuition behind the definition of a dynamical

⁴²Given the preceding remark, this concept occurs quite naturally in our setting, but we’re not aware of it having been defined before.

domain. (As a reference for the basic category-theoretic concepts used in the definition, see, e.g., Leinster (2014).)

4.4.6. DEFINITION. If \mathbf{C} is a category, \mathbf{D} a full subcategory, and $F : \mathbf{I} \rightarrow \mathbf{C}$ a diagram (i.e., a functor), a \mathbf{D} -*limit* of F is a cone (A, f_i) to F in \mathbf{C} ,⁴³ with A an object in \mathbf{D} , and the following universal property:

for any cone (B, g_i) to F in \mathbf{C} , if $B \in \mathbf{D}$, then there is a unique morphism $u : B \rightarrow A$ (in \mathbf{C} and hence also in \mathbf{D}) such that $f_i \circ u = g_i$ for all objects i in \mathbf{I} .

Note that, if it exists, (A, f_i) is unique up to unique isomorphism in \mathbf{D} : If (A', f'_i) is another \mathbf{D} -limit to F , there is a unique isomorphism $u : A \rightarrow A'$ in \mathbf{D} with $f'_i \circ u = f_i$.⁴⁴ Thus, we can also speak of ‘the’ \mathbf{D} -limit. If the categories are clear from context, we also say the *restricted limit*.

Finally, we state the definition of the appropriate notion of an expanding system of domains.

4.4.7. DEFINITION. An *expanding system of dynamical dcpos* is a structure $(\mathfrak{D}_i, p_{ij})_I$ where (I, \leq) is a directed preorder (called the *index set*), the $\mathfrak{D}_i = (D_i, v_i, f_i)$ are dynamical dcpos, and, for $i \leq j$ in I , $p_{ij} : \mathfrak{D}_j \rightarrow \mathfrak{D}_i$ is a dynamical morphism such that,

1. For all $i \leq j$ in I , $p_{ij} : D_j \rightarrow D_i$ is a projection.
2. For all $i \in I$, $p_{ii} = \text{id}_{\mathfrak{D}_i}$.
3. For all $i \leq j \leq k$ in I , $p_{ik} = p_{ij} \circ p_{jk}$.

An expanding system of dynamical dcpos $(\mathfrak{D}_i, p_{ij})_I$ is:

4. *upward deterministic* :iff for all $i \in I$, if $\exists a_i, b_i \neq b'_i \in \max D_i : b_i, b'_i \geq f_i(a_i)$, then there is $j \geq i$ in I such that $\forall a_j, b_j, b'_j \in \max D_j : \text{if } p_{ij}(a_j) = a_i, p_{ij}(b_j) = b_i, p_{ij}(b'_j) = b'_i$, then $b_j \not\geq f_j(a_j)$ or $b'_j \not\geq f_j(a_j)$.
5. *downward deterministic* :iff for all $i \in I$, if $\exists a_i \neq a'_i, b_i \in \max D_i : b_i \geq f_i(a_i), f_i(a'_i)$, then there is $j \geq i$ in I such that $\forall a_j, a'_j, b_j \in \max D_j : \text{if } p_{ij}(a_j) = a_i, p_{ij}(a'_j) = a'_i, p_{ij}(b_j) = b_i$, then $b_j \not\geq f_j(a_j)$ or $b_j \not\geq f_j(a'_j)$.

⁴³I.e., A is an object of \mathbf{C} and, for each object i of \mathbf{I} , $f_i : A \rightarrow F(i)$ is a morphism in \mathbf{C} such that, for every morphism $\iota : i \rightarrow j$ in \mathbf{I} , $F(\iota) \circ f_i = f_j$.

⁴⁴The usual proof: Since (A, f_i) is a cone to F in \mathbf{C} with A in \mathbf{D} and (A', f'_i) is a \mathbf{D} -limit, there is a unique morphism $u : A \rightarrow A'$ such that $f'_i \circ u = f_i$. Similarly, there is a unique morphism $u' : A' \rightarrow A$ such that $f_i \circ u' = f'_i$. Now, (A, f_i) is a cone to F in \mathbf{C} with A in \mathbf{D} and both $u' \circ u : A \rightarrow A$ and $\text{id}_A : A \rightarrow A$ are morphisms in \mathbf{D} with $f_i \circ (u' \circ u) = (f_i \circ u') \circ u = f'_i \circ u = f_i$ and $f_i \circ \text{id}_A = f_i$. Since (A, f_i) is limiting, there is a unique such morphism, so $u' \circ u = \text{id}_A$. Similarly, $u \circ u' = \text{id}_{A'}$. Hence $u : A \rightarrow A'$ is an isomorphism. And, by definition of (A', f'_i) being a \mathbf{D} -limit, it is unique with the property $f'_i \circ u = f_i$.

6. *eventually valuation-preserving* :iff (a) all \mathfrak{D}_i are finite and, for all $i \in I$ and $U_i \in \Sigma(D_i)$, there is $j_0 \geq i$ such that, for all $j \geq j_0$, we have

$$v_j \left(f_j^{-1} \left(p_{ij}^{-1} (\downarrow \max U_i) \right) \cap \max D_j \right) = v_j \left(p_{ij}^{-1} (\downarrow \max U_i) \cap \max D_j \right),^{45}$$

and (b) for all $i \leq j$ in I , if $a_i, b_i \in \max D_i$ with $f_i(a_i) \leq b_i$, then there is $a_j, b_j \in \max D_j$ such that $p_{ij}(a_j) = a_i$ and $p_{ij}(b_j) = b_i$ and $f_j(a_j) \leq b_j$.

A *finitary dynamical expanding system* is an upward deterministic expanding system of dynamical dcpos $(\mathfrak{D}_i, p_{ij})_I$ where I is countable and each \mathfrak{D}_i is a finite max-normalized dynamical Scott domain. It is *standard* if, additionally, the \mathfrak{D}_i are max-surjective and $(\mathfrak{D}_i, p_{ij})_I$ is downward deterministic and eventually valuation-preserving.

Comments: First, the term ‘finitary’ is to stress that we’re assuming the index set to be countable and the domains to be finite.

Second, at a very intuitive level, being upward deterministic means that if f_i fails, on input a_i , to uniquely pick out a maximal element above it, then this will be eventually remedied at some higher level j . Similarly, being downward deterministic means that if b_i informationally completes two images $f_i(a_i)$ and $f_i(a'_i)$ of informationally complete and distinct states, then the underlying inconsistency of these image will eventually be apparent at some higher level j .

Third, as usual, we can regard an expanding system $(\mathfrak{D}_i, p_{ij})_I$ as a diagram as follows: Consider the preorder I as a category with the elements of I as objects and with a single morphism $\iota : i \rightarrow j$ iff $i \leq j$. Then $(\mathfrak{D}_i, p_{ij})_I$ corresponds to the functor $F : I^{\text{op}} \rightarrow \mathbf{dDCP}^{\text{p}}$ that sends $i \in I$ to $F(i) := \mathfrak{D}_i$ and that sends $\iota^{\text{op}} : j \rightarrow i$ (i.e., $i \leq j$) in I to $F(\iota^{\text{op}}) := p_{ij} : \mathfrak{D}_j \rightarrow \mathfrak{D}_i$.⁴⁶

Fourth, just to be sure: the domain-theoretic notion of an ‘expanding system’ has nothing to do with the notion of an ‘expanding (or expansive) dynamical system’ from dynamical systems theory.

4.4.3 The limit theorem

The main technical contribution of this chapter is to show that the desired restricted limits of finitary dynamical expanding system do indeed exist.

4.4.8. THEOREM. *Let $(\mathfrak{D}_i, p_{ij})_I$ be a finitary dynamical expanding system. Write $\mathfrak{D}_i = (D_i, v_i, f_i)$. Then*

⁴⁵Qua sets of maximal elements of finite domains, both sets are Scott-open.

⁴⁶This is a functor: By clause (2), $F(\text{id}_i) = F(i \rightarrow i) = p_{ii} = \text{id}_{\mathfrak{D}_i}$. By clause (3), if $\kappa^{\text{op}} : k \rightarrow j$ and $\iota^{\text{op}} : j \rightarrow i$ are in I^{op} (so $k \geq j \geq i$), then $F(\iota^{\text{op}} \circ \kappa^{\text{op}}) = F(k \rightarrow i) = p_{ik} = p_{ij} \circ p_{jk} = F(j \rightarrow i) \circ F(k \rightarrow j) = F(\iota^{\text{op}}) \circ F(\kappa^{\text{op}})$.

1. $D := \{ \langle a(i) : i \in I \rangle \in \prod_{i \in I} D_i : \forall i \leq j \in I . a(i) = p_{ij}(a(j)) \}$ with the pointwise order⁴⁷ is a Scott domain and $\max D$ is closed in the Lawson topology. For $i \in I$, the function $p_i : D \rightarrow D_i$ defined by $p_i(a) := a(i)$ is a max-preserving Scott-continuous projection.
2. There is a unique continuous valuation $v : \sigma(D) \rightarrow [0, \infty]$ such that, for all $U_i \in \Sigma(D_i)$, we have $v_i(U_i) = v(p_i^{-1}(U_i))$. Moreover, v is max-normalized.
3. There is a largest (in the pointwise ordering) function $f : D \rightarrow D$ that is Scott-continuous and max-preserving such that, for all $a \in D$ and $i \in I$, $f(a)(i) \geq f_i(a(i))$. (The proof provides a concrete description of f .)

Hence $\mathfrak{D} := (D, v, f)$ is a max-normalized and max-preserving dynamical Scott domain. Moreover, (\mathfrak{D}, p_i) is a $\text{dSCO}_{\text{nm}}^{\text{P}}$ -limit of the diagram $(\mathfrak{D}_i, p_{ij})_I$ in $\text{dSCO}_{\text{n}}^{\text{P}}$:

4. (\mathfrak{D}, p_i) is a cone to the diagram with \mathfrak{D} in $\text{dSCO}_{\text{nm}}^{\text{P}}$ (so the p_i are morphisms in $\text{dSCO}_{\text{n}}^{\text{P}}$ such that $p_{ij} \circ p_j = p_i$ for all $i \leq j$ in I).⁴⁸
5. If $(\mathfrak{E}, \beta_i)_I$ is a cone to the diagram with \mathfrak{E} in $\text{dSCO}_{\text{nm}}^{\text{P}}$, then there is a unique morphism $\beta : \mathfrak{E} \rightarrow \mathfrak{D}$ in $\text{dSCO}_{\text{n}}^{\text{P}}$, which is defined by $\beta(e) := \langle \beta_i(e) : i \in I \rangle$, such that $p_i \circ \beta = \beta_i$ for all $i \in I$.

If $(\mathfrak{D}_i, p_{ij})_I$ additionally is standard, then

6. f is bijective on $\max D$.
7. \mathfrak{D} is valuation-preserving (i.e., $v(f^{-1}(U)) = v(U)$ for $U \in \Sigma(D)$).

The proof will take up the rest of this subsection. Since it is rather long, we don't write it in a proof environment, but rather divide it up in paragraphs corresponding to the items of the theorem (which then can have lemmas with their own proof environments). To avoid repeating notation, we'll always work in the setting of the theorem.

Claim (1) This is the standard theory of limits in domain theory (Abramsky and Jung 1994, sec. 3.3.1): (D_i, p_{ij}) is a diagram with countable and directed index set in the category of dcpos with continuous projections as morphisms and (D, p_i) is a limiting cone. Since the category of Scott domains is closed under these (bi)limits, D again is a Scott domain (for ω -algebraicity, note that the index set is countable).

Next, we show that p_i is max-preserving: Otherwise, there are $a \in \max D$ and $i \in I$ with $a(i) \notin \max D_i$. Since I is a countable directed poset, we can find a

⁴⁷I.e., for $a, b \in D$, $a \leq b$ iff $\forall i \in I : a(i) \leq b(i)$.

⁴⁸Note that this agrees with the definition of a cone when we regard the expanding system as a functor $F : I^{\text{op}} \rightarrow \text{dSCO}_{\text{n}}^{\text{P}}$: then this requires that, for $\iota^{\text{op}} : j \rightarrow i$, we have $p_i = F(\iota^{\text{op}}) \circ p_j = p_{ij} \circ p_j$.

cofinal chain $i = j_0 < j_1 < \dots$ in I . We recursively define a sequence $(d_{j_n})_{n \geq 0}$ with $p_{j_{n-1}j_n}(d_{j_n}) = d_{j_{n-1}}$ and $a(j_n) \leq d_{j_n} \in \max D_{j_n}$. This determines an element $d \in D$ with $a \leq d$. Since a is maximal, this implies $a = d$, so $a(i) = d(i) \in \max D_i$, contradiction.

For $n = 0$, define d_{j_0} as some maximal element of D_i above $a(i)$. Given $a(j_n) \leq d_{j_n} \in \max D_{j_n}$, we define $d_{j_{n+1}}$ as follows. Since $a(j_{n+1}) \in D_{j_{n+1}}$ and $p_{j_n j_{n+1}}(a(j_{n+1})) = a(j_n) \leq d_{j_n} \in \max D_{j_n}$, there is, since $p_{j_n j_{n+1}}$ is max-bisimulative, $d_{j_{n+1}} := d \in \max D_{j_{n+1}}$ such that $a(j_{n+1}) \leq d$ and $p_{j_n j_{n+1}}(d) = d_{j_n}$.

Finally, we show that $\max D$ is Scott-closed. Since the p_i are max-preserving, we have $\max D = \bigcap_{i \in I} p_i^{-1}(\max D_i)$. Since the $\max D_i$ are Lawson-closed (the D_i are finite and the Lawson topology on a finite dcpo is the discrete topology), it remains to show that the p_i are Lawson-continuous (so $\max D$ is an intersection of Lawson-closed sets). Indeed, qua projection, the p_i -preimages of upsets are upsets, so p_i is continuous in the lower topology. Since p_i is Scott-continuous and the Lawson topology is the join of the Scott topology and the lower topology, p_i is Lawson-continuous.

Claim (2) We'll use a result of Goubault-Larrecq (2018, thm. 4.2) on projective limits of valued topological spaces: it states that, given an ep-system $(X_i, (p_{ij}, e_{ij}))$ in the category of topological spaces with projective limit (X, p_i) together with valuations v_i on X_i with the obvious compatibility condition, there is a unique valuation v on X with the obvious compatibility condition. Let's precisely state and verify the assumption for our case.

Each D_i with its Scott topology is a topological space whose specialization ordering is precisely the partial order on the set D_i (Abramsky and Jung 1994, prop. 2.3.2). And the $p_{ij} : D_j \rightarrow D_i$ are continuous with respect to the Scott topologies on D_j and D_i (since p_{ij} is Scott-continuous). In fact, since the p_{ij} are projections (and the D_i are T_0 spaces), they have uniquely determined corresponding embeddings $e_{ij} : D_i \rightarrow D_j$, turning $(D_i, (p_{ij}, e_{ij}))_I$ into an ep-system. Moreover, qua bilimit of continuous domains, the Scott topology on D coincides with the relative product topology, which is the projective topology of the projective system (D_i, p_{ij}) (Abramsky and Jung 1994, ex. 3.3.12 (18)). Further, each $v_i : \Sigma(D_i) \rightarrow [0, 1]$ is a continuous valuation on the topological space D_i turning (D_i, v_i) into a valued space. And, for $i \leq j$ and $U_i \in \Sigma(D_i)$, we have, since $p_{ij} : D_j \rightarrow D_i$ is valuation-preserving (qua dynamical morphism), that $v_i(U_i) = v_j(p_{ij}^{-1}(U_i))$.

The result on valued spaces mentioned above now implies that there is a unique continuous valuation $v : \Omega(D) \rightarrow [0, \infty]$ on D (where $\Omega(D) = \Sigma(D)$ is the set of open sets in the projective topology) such that, for each $i \in I$ and $U_i \in \Sigma(U_i)$, $v_i(U_i) = v(p_i^{-1}(U_i))$.

It remains to show that v is max-normalized. Indeed, to show that v is normalized, choose any $i \in I$ (recall that I is nonempty since it is directed), and

we have, since v_i is normalized, that $v(D) = v(p_i^{-1}(D_i)) = v_i(D_i) = 1$. To show that v is max-normalized, note that $\max D = \bigcap_{i \in I} p_i^{-1}(\max D_i)$. Since D_i is finite and I countable, this is a countable intersection of Scott-open sets and, since v_i is max-normalized and D_i finite, $v(p_i^{-1}(\max D_i)) = v_i(\max D_i) = 1$.

Claim (3) We define $f : D \rightarrow D$ as follows: First we define a function $\bar{f} : \max D \rightarrow \max D$. Next, we show that it is continuous. Then we extend it to a function $f : D \rightarrow D$. Finally, we show that this f has the required properties.

To define \bar{f} we observe the following.

4.4.9. LEMMA. 1. For each $a \in \max D$, there is some $b \in \max D$ such that $\forall i \in I : b(i) \geq f_i(a(i))$.

2. For each $a \in \max D$, there is at most one $b \in \max D$ such that $\forall i \in I : b(i) \geq f_i(a(i))$.

PROOF. Ad (1). For $i \in I$, define $F_i := p_i^{-1}(\uparrow f_i(a(i))) \cap \max D$. Since $\max D$ is Lawson-closed and $\uparrow f_i(a(i)) \subseteq D_i$ is Lawson-closed and p_i Lawson-continuous, also F_i is Lawson-closed. Moreover, F_i is nonempty: Let $b_i \in D_i$ be a maximal element above $f_i(a(i))$. Since p_i is surjective and monotone, there is a maximal element $b \in D$ with $p_i(b) = b_i$. So $b \in F_i$. Finally, if $i \leq j$, then $F_i \supseteq F_j$: If $b \in F_j$, then $b \in \max D$ with $p_j(b) \geq f_j(a(j))$. Since p_{ij} is monotone and max-semi-equivariant (and $a(j) \in \max D_j$), we have

$$b(i) = p_{ij}(p_j(b)) \geq p_{ij}(f_j(a(j))) \geq f_i(p_{ij}(a(j))) = f_i(a(i)).$$

Hence $b \in F_i$.

Thus, $\{F_i : i \in I\}$ is a family of Lawson-closed subsets of D with the finite intersection property. Since the Lawson topology on a Scott domain is compact, there is $b \in \bigcap_I F_i$, so $b \in \max D$ with $b(i) \geq f_i(a(i))$ for all $i \in I$, as needed.

Ad (2). Otherwise, there is $a, b, b' \in \max D$ with $b \neq b'$ and $\forall i : b(i), b'(i) \geq f_i(a(i))$. Since $b \neq b'$, there is $i \in I$ with $b(i) \neq b'(i)$, and $a(i), b(i), b'(i) \in \max D_i$ (since the p_i are max-preserving) and $b(i), b'(i) \geq f_i(a(i))$. However, for all $j \geq i$, the elements $a(j), b(j), b'(j) \in \max D_j$ are such that $p_{ij}(a(j)) = a(i)$, $p_{ij}(b(j)) = b(i)$, $p_{ij}(b'(j)) = b'(i)$, but $b(j), b'(j) \geq f_j(a(j))$. This contradicts the fact that the expanding system is upward deterministic. \square

Hence, we can define the function $\bar{f} : \max D \rightarrow \max D$ by mapping $a \in \max D$ to the unique $b \in \max D$ with $\forall i : b(i) \geq f_i(a(i))$. To establish the continuity of \bar{f} , we observe the following.

4.4.10. LEMMA. 1. For all $i \in I$ and $a \in \max D$, there is $j_a = j \geq i$ such that, for all $b \in \max D$, if $b(j) \geq f_j(a(j))$, then $b(i) = \bar{f}(a)(i)$.

2. For all $i \in I$ there is $j \geq i$ such that, for all $a, b \in \max D$, if $a(j) = b(j)$, then $\bar{f}(a)(i) = \bar{f}(b)(i)$.
3. For all $i \in I$ and $U_i \subseteq D_i$, there is $j_0 \geq i$ such that, for all $j \geq j_0$,

$$\bar{f}^{-1}(p_i^{-1}U_i \cap \max D) = \bigcup_{a_j \in R_{ij}^{-1}U_i} p_j^{-1}(a_j) \cap \max D$$

where $R_{ij}^{-1}U_i := \{a_j \in \max D_j : \exists a \in \max D. a(j) = a_j, \bar{f}(a)(i) \in U_i\}$.

Note that (1) is like the ϵ - δ definition of continuity and (2) is like uniform continuity. Indeed, the proof of (2) from (1) is, since $\max D$ is a compact space, essentially (the proof of) the Heine-Cantor theorem.

PROOF. Ad (1). Let $i \in I$ and $a \in \max D$. To show the claim, assume for contradiction that for all $j \geq i$ there is $b \in \max D$ with $b(j) \geq f_j(a(j))$ but $b(i) \neq \bar{f}(a)(i)$. For $j \geq i$ consider the sets

$$F_j := \{b \in \max D : b(j) \geq f_j(a(j)), b(i) \neq \bar{f}(a)(i)\}.$$

Then F_j is a Lawson-closed subset of D : We have $F_j = \max D \cap p_j^{-1}(\uparrow f_j(a(j))) \cap (p_i^{-1}(\bar{f}(a)(i)))^c$, and (i) $\max D$ is Lawson-closed, (ii) p_j is Lawson-continuous and $\uparrow f_j(a(j)) \subseteq D_j$ Lawson-closed, and (iii) p_i Lawson-continuous and $\{\bar{f}(a)(i)\} \subseteq D_i$ is Lawson-open, so $p_i^{-1}(\bar{f}(a)(i))$ is Lawson-open and its complement Lawson-closed. Moreover, by assumption, F_j is nonempty. And if $j \leq k$, then $F_j \supseteq F_k$: If $b \in F_k$, then $b \in \max D$ with $b(i) \neq \bar{f}(a)(i)$ and $b(k) \geq f_k(a(k))$. The latter implies, since p_{jk} is monotone and max-semi-equivariant (and $a(k) \in \max D_k$ since p_k is max-preserving), that

$$b(j) = p_{jk}(b(k)) \geq p_{jk}(f_k(a(k))) \geq f_j(p_{jk}(a(k))) = f_j(a(j)),$$

so $b \in F_j$. Hence, $\{F_j : j \geq i\}$ is a family of Lawson-closed subsets of D with the finite intersection property, so, since the Lawson topology on D is compact, there is $b \in \bigcap_{j \geq i} F_j$. But then, for all $k \in I : b(k) \geq f_k(a(k))$ (given $k \in I$, let, by directedness, $j \geq k, i$, so $b(j) \geq f_j(a(j))$, and $b(k) \geq f_k(a(k))$ follows as above by applying p_{kj} to both sides). So, by definition of \bar{f} , $b = \bar{f}(a)$, which contradicts $b(i) \neq \bar{f}(a)(i)$.

Ad (2). Let $i \in I$. First, we show that, for all $a \in \max D$, there is $j_a \geq i$ such that

$$(*) \text{ For all } b \in \max D, \text{ if } b(j_a) = a(j_a), \text{ then } \bar{f}(a)(i) = \bar{f}(b)(i).$$

Indeed, by (1), there is $j_a \geq i$ such that, for all $c \in \max D$, if $c(j_a) \geq f_{j_a}(a(j_a))$, then $c(i) = \bar{f}(a)(i)$. To show (*), let $b \in \max D$ with $b(j_a) = a(j_a)$. So, for $c := \bar{f}(b) \in \max D$, we have $c(j_a) = \bar{f}(b)(j_a) \geq f_{j_a}(b(j_a)) = f_{j_a}(a(j_a))$. Hence $\bar{f}(b)(i) = c(i) = \bar{f}(a)(i)$, as claimed.

Thus, for $a \in \max D$, we define $U_a := p_{j_a}^{-1}(a(j_a)) \cap \max D$, which is an open set of $\max D$ with $a \in U_a$. Thus, $\{U_a : a \in \max D\}$ is an open cover of the compact space $\max D$, so there is a finite subcover $\mathcal{U} = \{U_{a_1}, \dots, U_{a_m}\}$ (with $m \geq 1$). By directedness of I , let $j \geq j_{a_1}, \dots, j_{a_m} \geq i$.

We claim that $j \geq i$ has the required property: let $a, b \in \max D$ with $a(j) = b(j)$. Since \mathcal{U} is a cover, $a \in U_{a_k}$ for some $k \in \{1, \dots, m\}$. So $a(j_{a_k}) = a_k(j_{a_k})$. Since $j \geq j_{a_k}$ and $a(j) = b(j)$, also $a(j_{a_k}) = b(j_{a_k})$. So $a(j_{a_k}) = a_k(j_{a_k}) = b(j_{a_k})$, whence (*) implies $\bar{f}(a)(i) = \bar{f}(a_k)(i) = \bar{f}(b)(i)$, as needed.

Ad (3). Let $i \in I$ and $U_i \subseteq D_i$. Let $j_0 \geq i$ be as in (2). Let $j \geq j_0$. If $a \in \bar{f}^{-1}(p_i^{-1}U_i \cap \max D)$, then $a_j := a(j) \in R_{i_j}^{-1}U_i$ (since $\bar{f}(a)(i) \in U_i$) and $a \in p_j^{-1}(a_j) \cap \max D$. For the other direction, let $a \in p_j^{-1}(a_j) \cap \max D$ for some $a_j \in R_{i_j}^{-1}U_i$. So, by definition of $R_{i_j}^{-1}U_i$, there is $a' \in \max D$ such that $a'(j) = a_j$ and $\bar{f}(a')(i) \in U_i$. So $a(j) = a_j = a'(j)$, so, since $j \geq j_0$, also $a(j_0) = a'(j_0)$, whence, by (2), $\bar{f}(a)(i) = \bar{f}(a')(i) \in U_i$. So $a \in \bar{f}^{-1}(p_i^{-1}U_i \cap \max D)$. \square

Note that (3) shows that $\bar{f} : \max D \rightarrow \max D$ is continuous (in the relative Scott topology on $\max D$): For the subbasic open sets $p_i^{-1}(U_i) \cap \max D$ of $\max D$ (with $i \in I$ and $U_i \subseteq D_i$ Scott-open), $\bar{f}^{-1}(p_i^{-1}(U_i) \cap \max D)$ is a union of subbasic open sets in $\max D$, and hence open.

Now, we can use a characteristic theorem about Scott domains to extend the function $\bar{f} : \max D \rightarrow \max D \subseteq D$ to the whole of D .

4.4.11. LEMMA. *Define the function $f : D \rightarrow D$ by*

$$f(a) := \bigvee \left\{ \bigwedge \bar{f}(U \cap \max D) : a \in U \in \Sigma(D) \right\}.$$

Then f is the largest Scott-continuous function extending $\bar{f} : \max D \rightarrow \max D$. In particular, f is max-preserving.

PROOF. We'll use the characterization of Scott domains (regarded as topological spaces with the Scott topology) as precisely those T_0 spaces that are densely injective (Gierz et al. 2003, Prop. II-3.11). Since D is a Scott domain, it is densely injective, which means that every continuous map $h : A \rightarrow D$ extends continuously to any space B containing A as a dense subspace. We choose $A := \max D$ and $h : \max D \rightarrow D$ as $h(z) := \bar{f}(z)$. By the previous lemma, h is indeed continuous. We choose $B := D$. Then A is a dense subspace of B : It is a subspace, so we need to show that $\max D$ is dense in D . Indeed, given a point $a \in D$ and a Scott-open set $a \in U \subseteq D$, we show that $U \cap A$ is nonempty: since D is a dcpo, there is a maximal element $a' \geq a$, whence, since U is an upset, $a' \in U \cap \max D$. Hence, h has an extension to a continuous $h^* : D \rightarrow D$ which is given as $h^*(z) := \bigvee \left\{ \bigwedge h(U \cap \max D) : y \in U \in \Sigma(D) \right\}$ (Gierz et al. 2003, Prop. II-3.9). Thus, $h^* = f$ is a Scott-continuous extension of \bar{f} . \square

Finally, we check that $f : D \rightarrow D$ has the required properties.

4.4.12. LEMMA. 1. *The function f from lemma 4.4.11 has the property*

$$\forall a \in D \forall i \in I : f(a)(i) \geq f_i(a(i)) \quad (4.3)$$

2. *The function f is the largest Scott-continuous and max-preserving function with property (4.3).*

PROOF. Ad (1). Let $a \in D$ and $i \in I$. Since p_i is Scott-continuous and a projection, we have

$$\begin{aligned} f(a)(i) &= p_i \bigvee \{ \bigwedge \bar{f}(U \cap \max D) : a \in U \in \Sigma(D) \} \\ &\stackrel{\text{cont.}}{=} \bigvee \{ p_i \bigwedge \bar{f}(U \cap \max D) : a \in U \in \Sigma(D) \} \\ &\stackrel{\text{proj.}}{=} \bigvee \{ \bigwedge p_i \circ \bar{f}(U \cap \max D) : a \in U \in \Sigma(D) \} \end{aligned}$$

Consider $U := p_i^{-1}(\uparrow a(i))$. Then $a \in U \in \Sigma(D)$ (since p_i is Scott-continuous and $\uparrow a(i)$ is an upset in a finite dcpo and hence Scott-open). So $f(a)(i) \geq \bigwedge p_i \circ \bar{f}(U \cap \max D)$. Next, we show:

$$p_i \circ \bar{f}(U \cap \max D) \subseteq \uparrow f_i(a(i)) \cap \max D_i.$$

Indeed, let $b_i \in p_i \circ \bar{f}(U \cap \max D)$. So $b_i = b(i)$ for $b = \bar{f}(c)$ for some $c \in U \cap \max D$. So $b_i \in \max D_i$ (since p_i is max-preserving and $b = \bar{f}(c) \in \max D$) and, by definition of \bar{f} , $b_i = b(i) = \bar{f}(c)(i) \geq f_i(c(i)) \geq f_i(a(i))$, where the last step follows since $c(i) \geq a(i)$ (since $c \in U$) and f_i is monotone.

Note that $\bigwedge \uparrow f_i(a(i)) \cap \max D_i$ exists since any nonempty set has an infimum in a Scott domain. Now, we have

$$f(a)(i) \geq \bigwedge p_i \circ \bar{f}(U \cap \max D) \geq \bigwedge \uparrow f_i(a(i)) \cap \max D_i \geq f_i(a(i))$$

where the last step follows since $f_i(a(i))$ is a lower bound of the set over which the infimum is taken.

Ad (2). By now, we've established that the function f is Scott-continuous and max-preserving with property (4.3), so we need to show that it is the largest. Assume f' is another such function. Since f' is max-preserving, $f' \upharpoonright \max D$ is a function from $\max D$ to $\max D$, and it has, by assumption, the property $\forall a \in \max D \forall i \in I : f' \upharpoonright \max D(a)(i) \geq f_i(a(i))$. Hence $f' \upharpoonright \max D = \bar{f}$. Thus, f' is a Scott-continuous extension of \bar{f} . Since, by lemma 4.4.11, f is the largest Scott-continuous extension of \bar{f} , we have $f' \leq f$, as needed. \square

The second item of the lemma states that $f : D \rightarrow D$ has the required properties, which finishes the proof of claim (3).

Claim (4) We already know that the p_i are projections and we have $p_{ij} \circ p_j = p_i$. So it remains to show that p_i is a dynamical morphism. In (1), we've already seen that p_i is Scott-continuous and max-preserving, so it remains to check conditions (3)–(5) of being a dynamical morphism.

Concerning condition (3), let $a \in D$, $p_i(a) = a(i) \leq e \in \max D_i$. We need to find $d \in \max D$ with $d \geq a$ and $d(i) = e$. Since I is a countable directed poset, we can find a cofinal chain $i = j_0 < j_1 < \dots$ in I . We recursively define a sequence $(d_{j_n})_{n \geq 0}$ with $p_{j_{n-1}j_n}(d_{j_n}) = d_{j_{n-1}}$ and $a(j_n) \leq d_{j_n} \in \max D_{j_n}$ and $d_i = e$. This determines an element $d \in \max D$ with $a \leq d$ and $d(i) = e$, as needed.

For $n = 0$, define $d_{j_0} := e \in \max D_i$, so $a(j_0) \leq d_{j_0}$. Given $a(j_n) \leq d_{j_n} \in \max D_{j_n}$, we define $d_{j_{n+1}}$ as follows. Since $a(j_{n+1}) \in D_{j_{n+1}}$ and $p_{j_n j_{n+1}}(a(j_{n+1})) = a(j_n) \leq d_{j_n} \in \max D_{j_n}$, there is, since $p_{j_n j_{n+1}}$ is max-bisimulative, $d_{j_{n+1}} := d \in \max D_{j_{n+1}}$ such that $a(j_{n+1}) \leq d$ and $p_{j_n j_{n+1}}(d) = d_{j_n}$.

Concerning condition (4), by the defining property of v , we have, for $U_i \in \Sigma(D_i)$, that $v_i(U_i) = v(p_i^{-1}(U_i))$.

Concerning condition (5), for $a \in \max D$, we have, by definition of f , that $p_i(f(a)) = f(a)(i) \geq f_i(a(i)) = f_i(p_i(a))$.

Claim (5) Let $(\mathfrak{E}, \beta_i)_I$ be another cone to the diagram $(\mathfrak{D}_i, p_{ij})_I$ in \mathbf{dSCO}_n^p with \mathfrak{E} in \mathbf{dSCO}_{nm}^p . Write $\mathfrak{E} = (E, w, g)$. Set $\beta : \mathfrak{E} \rightarrow \mathfrak{D}, \alpha(e) := \langle \beta_i(e) : i \in I \rangle$. From the standard theory of limits in domain theory (Abramsky and Jung 1994, sec. 3.3.1), it is known that $\beta : E \rightarrow D$ is the unique continuous projection such that $\beta_i = p_i \circ \beta$ for all $i \in I$. So we need to show that β is a dynamical morphism.

Concerning condition (1), as just mentioned, β is Scott-continuous.

Concerning condition (2), if $e \in \max E$, then $\beta(e) = \langle \beta_i(e) : i \in I \rangle$ is a thread of maximal elements (since all β_i are max-preserving), so $\beta(e)$ is maximal in D .

Concerning condition (3), let $a \in E$ and $\beta(a) \leq e \in \max D$. For $i \in I$, consider $C_i := \uparrow a \cap \beta_i^{-1}(e(i))$.

This is a Lawson-closed subset of E : First, by definition of the Lawson topology, $E \setminus \uparrow a$ is open (so $\uparrow a$ is closed). Second, we have $\uparrow e(i) = \{e(i)\}$ since $e(i) = p_i(e)$ is maximal because p_i is max-preserving. Further, β_i is a projection since it is a morphism in \mathbf{dSCO}_n^p . Hence, the β_i preimage of $\uparrow e(i)$ also is a principal upset, so $\beta_i^{-1}(e(i))$ is Lawson-closed.

Moreover, C_i is nonempty: Since β_i is max-bisimulative and $\beta_i(a) \leq e(i)$ (since $\beta(a) \leq e$), there is $d \in \max E$ with $d \geq a$ and $\beta_i(d) = e(i)$, so $d \in C_i$.

Finally, if $i \leq j$, then $C_i \supseteq C_j$: If $d \in C_j$, then $d \in E$ with $d \geq a$ and $\beta_j(d) = e(j)$, so, since the β_i 's are part of a cone,

$$\beta_i(d) = p_{ij}(\beta_j(d)) = p_{ij}(e(j)) = e(i),$$

whence $d \in C_i$.

Hence, $\{C_i : i \in I\}$ is a family of Lawson-closed subsets of E with the finite intersection property (given C_{i_1}, \dots, C_{i_n} , let $I \ni j \geq i_1, \dots, i_n$, then $\emptyset \neq C_j \subseteq$

$C_{i_1} \cap \dots \cap C_{i_n}$). Since E is a Scott domain, its Lawson topology is compact. Hence $\bigcap_{i \in I} C_i$ is nonempty. Let $d \in \bigcap_I C_i$. Then $d \geq a$ and $\beta(d) = \langle \beta_i(d) : i \in I \rangle = \langle e(i) : i \in I \rangle = e$, as needed.

Concerning condition (4), we first show that for the basic opens $V = p_{i_1}^{-1}(U_{i_1}) \cap \dots \cap p_{i_n}^{-1}(U_{i_n})$ of D (with $i_k \in I$ and $U_{i_k} \in \Sigma(D_{i_k})$ for $k = 1, \dots, n$), we have $v(V) = w(\beta^{-1}(V))$.

Since I is directed, let $j \geq i_1, \dots, i_n$. So

$$V = \bigcap_{k=1}^n p_{i_k}^{-1}(U_{i_k}) \stackrel{p_{i_k j} \circ p_j = p_{i_k}}{=} \bigcap_{k=1}^n p_j^{-1}(p_{i_k j}^{-1}(U_{i_k})) = p_j^{-1}\left(\bigcap_{k=1}^n p_{i_k j}^{-1}(U_{i_k})\right).$$

Hence, since β_j is valuation-preserving,

$$\begin{aligned} v(V) &= v\left(p_j^{-1}\left(\bigcap_{k=1}^n p_{i_k j}^{-1}(U_{i_k})\right)\right) = v_j\left(\bigcap_{k=1}^n p_{i_k j}^{-1}(U_{i_k})\right) \\ &= w\left(\beta_j^{-1}\left(\bigcap_{k=1}^n p_{i_k j}^{-1}(U_{i_k})\right)\right) = w\left(\bigcap_{k=1}^n \beta_j^{-1}(p_{i_k j}^{-1}(U_{i_k}))\right) \\ &\stackrel{p_{i_k j} \circ \beta_j = \beta_{i_k}}{=} w\left(\bigcap_{k=1}^n \beta_{i_k}^{-1}(U_{i_k})\right) \stackrel{\beta_i = p_i \circ \beta}{=} w\left(\bigcap_{k=1}^n \beta^{-1}(p_{i_k}^{-1}(U_{i_k}))\right) \\ &= w\left(\beta^{-1}\left(\bigcap_{k=1}^n p_{i_k}^{-1}(U_{i_k})\right)\right) = w\left(\beta^{-1}(V)\right). \end{aligned}$$

Now we need to show the claim for arbitrary opens V of D . So $V = \bigcup_{k \in K} V_k$ for basic opens V_k and an index set K . Without loss of generality, $\{V_k : k \in K\}$ is directed (otherwise consider the family of finite unions of the V_k 's). Note that then also $\beta^{-1}(V_k)$ is a directed family of open sets in E . Then

$$\begin{aligned} v(V) &= v\left(\bigcup_{k \in K} V_k\right) = \sup_{k \in K} v(V_k) = \sup_{k \in K} w(\beta^{-1}(V_k)) \\ &= w\left(\bigcup_{k \in K} \beta^{-1}(V_k)\right) = w\left(\beta^{-1}\left(\bigcup_{k \in K} V_k\right)\right) = w\left(\beta^{-1}(V)\right). \end{aligned}$$

Concerning condition (5), let $a \in \max E$ and show $\beta(g(a)) \geq f(\beta(a))$. Since \mathfrak{E} is in $\text{dSCO}_{\text{nm}}^{\text{p}}$, g is max-preserving, so $g(a) \in \max E$. Since β is max-preserving, $\beta(g(a)) \in \max D$. Since β_i is max-semi-equivariant, we have, for $i \in I$,

$$\beta(g(a))(i) = \beta_i(g(a)) \geq f_i(\beta_i(a)) = f_i(\beta(a)(i)).$$

Thus, $\beta(a), \beta(g(a)) \in \max D$ are such that, for all $i \in I$, $\beta(g(a))(i) \geq f_i(\beta(a)(i))$. Hence, by definition of f , $f(\beta(a)) = \beta(g(a))$, as needed.

Claim (6) Injective: Let $a \neq a'$ be in $\max D$ and assume, for contradiction, that $f(a) = f(a') =: b$. Then there is $i \in I$ such that $a(i) \neq a'(i)$ in $\max D_i$ with $b(i) = f(a)(i) \geq f_i(a(i))$ and $b(i) = f(a')(i) \geq f_i(a'(i))$. However, for any $j \geq i$ we have $a(j), a'(j), b(j) \in \max D_j$ with $p_{ij}(a(j)) = a(i)$, $p_{ij}(a'(j)) = a'(i)$, $p_{ij}(b(j)) = b(i)$, but both $b(j) = f(a)(j) \geq f_j(a(j))$ and $b(j) = f(a')(j) \geq f_j(a'(j))$. This contradicts $(\mathfrak{D}_i, p_{ij})_I$ being downward deterministic.

Surjective: Let $b \in \max D$ and find $a \in \max D$ such that $f(a) = b$. For $i \in I$, define

$$C_i := \{a \in \max D : b(i) \geq f_i(a(i))\}.$$

Closed: We have $C_i = p_i^{-1}(f_i^{-1}(\downarrow b(i))) \cap \max D$ and $f_i^{-1}(\downarrow b(i)) \subseteq D_i$ is closed in the Lawson topology (since D_i is finite and hence discrete in the Lawson topology), so, since p_i is Lawson-continuous, C_i is closed in the relative Lawson topology, which coincides with the relative Scott topology.

Monotone: If $i \leq j$ in I , then $C_i \supseteq C_j$: if $a \in C_j$, then $b(j) \geq f_j(a(j))$, so

$$b(i) = p_{ij}(b(j)) \geq p_{ij}(f_j(a(j))) \geq f_i(p_{ij}(a(j))) = f_i(a(i)),$$

so $a \in C_i$.

Nonempty: Given $i \in I$, we have $b(i) \in \max D_i$. Since \mathfrak{D}_i is max-surjective, there is $a_i \in \max D_i$ such that $b(i) \geq f_i(a_i)$. Since p_i is surjective and monotone, there is $a \in \max D$ with $p_i(a) = a_i$. Hence $a \in C_i$.

Since $\max D$ is compact, there is $a \in \bigcap_{i \in I} C_i$. So, for all $i \in I$, $b(i) \geq f_i(a(i))$, whence, by definition of f on maximal elements, $b = f(a)$, as needed.

Claim (7) We need to show $v(f^{-1}(U)) = v(U)$ for $U \in \Sigma(D)$. The main task is to show this for the subbasic opens $U = p_i^{-1}(U_i)$ (for $i \in I$ and $U_i \in \Sigma(D_i)$). This will imply the claim by a simple and general argument.

We first analyze the set $\bar{f}^{-1}(p_i^{-1}U_i)$. We'll use lemma 4.4.10 (3) which says that, under appropriate assumptions, $\bar{f}^{-1}(p_i^{-1}U_i) = \bigcup_{a_j \in R_{ij}^{-1}U_i} p_j^{-1}(a_j) \cap \max D$. To analyze (the v -value of) the union, we start by observing that requirement (b) of being eventually valuation-preserving extends to the following.

4.4.13. LEMMA. *For all $i \in I$ and $a_i, b_i \in \max D_i$ with $b_i \geq f_i(a_i)$, there is $a \in \max D$ with $a(i) = a_i$ and $f(a)(i) = b_i$.*

PROOF. Since I is countable and directed, let $i = i_0 \leq i_1 \leq \dots$ be a cofinal chain in I . Set $a_{i_0} := a_i$ and $b_{i_0} := b_i$. Given a_{i_n} and b_{i_n} in $\max D_{i_n}$ with $b_{i_n} \geq f_{i_n}(a_{i_n})$, choose, using requirement (b) of being eventually valuation-preserving, $a_{i_{n+1}}$ and $b_{i_{n+1}}$ in $\max D_{i_{n+1}}$ with $p_{i_n i_{n+1}}(a_{i_{n+1}}) = a_{i_n}$ and $p_{i_n i_{n+1}}(b_{i_{n+1}}) = b_{i_n}$ and $b_{i_{n+1}} \geq f_{i_{n+1}}(a_{i_{n+1}})$. Then $\langle a_{i_n} \rangle$ and $\langle b_{i_n} \rangle$ determine a and b in $\max D$, respectively, with $a(i) = a(i_0) = a_{i_0} = a_i$ and $b(i) = b(i_0) = b_{i_0} = b_i$ and, for all $j \in I$,

$b(j) \geq f_j(a(j))$.⁴⁹ So $\bar{f}(a) = b$, whence $f(a)(i) = b(i) = b_i$. \square

We can use this now to analyze the union $\bigcup_{a_j \in R_{ij}^{-1}U_i} p_j^{-1}(a_j) \cap \max D$.

4.4.14. LEMMA. 1. For $i \leq j$ in I and $U_i \text{ in } \Sigma(D_i)$ we have, for $a_j \in \max D_j$,

$$\begin{aligned} a_j \in R_{ij}^{-1}U_i &\stackrel{\text{def}}{\Leftrightarrow} \exists a \in \max D : a(j) = a_j, f(a)(i) \in U_i \\ &\Leftrightarrow \exists b_i \in \max U_i : p_{ij}(f_j(a_j)) \leq b_i \end{aligned}$$

2. For $i \leq j$ in I and $U_i \text{ in } \Sigma(D_i)$,

$$\bigcup_{a_j \in R_{ij}^{-1}U_i} p_j^{-1}(a_j) \cap \max D = p_j^{-1}(f_j^{-1}p_{ij}^{-1}(\downarrow \max U_i)) \cap \max D$$

PROOF. Ad (1). (\Rightarrow) Let $a \in \max D$ with $a(j) = a_j$ and $b_i := f(a)(i) \in U_i$. Since $f(a)(i)$ is maximal, $b_i \in \max U_i$. We have $f_j(a(j)) \leq f(a)(j)$, so

$$p_{ij}(f_j(a_j)) \leq p_{ij}(f(a)(j)) = f(a)(i) = b_i.$$

(\Leftarrow) Let $b_i \in \max U_i$ with $p_{ij}(f_j(a_j)) \leq b_i$. Since $f_j(a_j) \in D_j$ with $p_{ij}(f_j(a_j)) \leq b_i \in \max D_i$ (since U_i is an upset) and p_{ij} is max-bisimulative, there is $b_j \in \max D_j$ with $f_j(a_j) \leq b_j$ and $p_{ij}(b_j) = b_i$. By lemma 4.4.13, there is $a \in \max D$ with $a(j) = a_j$ and $f(a)(j) = b_j$. So $f(a)(i) = p_{ij}(f(a)(j)) = p_{ij}(b_j) = b_i \in U_i$, as needed.

Ad (2). By (1), we have, for $c \in \max D$,

$$\begin{aligned} &\exists a_j \in R_{ij}^{-1}U_i : c(j) = a_j \\ &\Leftrightarrow \exists a_j \in \max D_j : \exists b_i \in \max U_i : p_{ij}(f_j(a_j)) \leq b_i \text{ and } c(j) = a_j \\ &\stackrel{50}{\Leftrightarrow} \exists b_i \in \max U_i : p_{ij}(f_j(c(j))) \leq b_i \\ &\Leftrightarrow c \in p_j^{-1}(f_j^{-1}p_{ij}^{-1}(\downarrow \max U_i)), \end{aligned}$$

so the claim follows. \square

Using requirement (a) of being eventually valuation-preserving, we can simplify the v -value of the set $p_j^{-1}(f_j^{-1}p_{ij}^{-1}(\downarrow \max U_i))$.

4.4.15. LEMMA. For $i \in I$ and $U_i \in \Sigma(U_i)$, there is $j_0 \geq i$ such that, for all $j \geq j_0$,

$$v\left(p_j^{-1}(f_j^{-1}p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j)\right) = v_i(U_i).$$

⁴⁹For the latter, given $j \in I$, let, by cofinality, $i_n \geq j$. Then, since $b_{i_n} \geq f_{i_n}(a_{i_n})$, we have $b(j) = p_{ji_n}(b(i_n)) \geq p_{ji_n}(f_{i_n}(a_{i_n})) \geq f_j(p_{ji_n}(a_{i_n})) = f_j(a(j))$.

⁵⁰For the reverse direction, take $a_j := c(j)$.

PROOF. By requirement (a) of being eventually valuation-preserving, there is $j_0 \geq i$ such that, for all $j \geq j_0$,

$$v_j(f_j^{-1}p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j) = v_j(p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j).$$

As in the proof of 4.3.11 (14), this further equals $v_i(U_i)$: Since p_{ij} is max-preserving and valuation-preserving, and since v_j is max-normalized, we have $v_j(p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j) = v_j(p_{ij}^{-1}(\max U_i) \cap \max D_j) = v_j(p_{ij}^{-1}(\max U_i)) = v_i(\max U_i) = v_i(U_i \cap \max D_i) = v_i(U_i)$. Hence

$$v(p_j^{-1}(f_j^{-1}p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j)) = v_j(f_j^{-1}p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j) = v_i(U_i),$$

as needed. \square

We need one last, rather general lemma.

4.4.16. LEMMA. *For $U, V \in \Sigma(D)$, if $U \cap \max D = V \cap \max D$, then $v(U) = v(V)$.*

PROOF. This proof anticipates an argument that we'll use below (in theorem 4.5.1). Since v is a normalized continuous valuation on a Scott domain (hence continuous depo), it can be extended to a (unique) measure μ on the Borel σ -algebra $\mathcal{B}(\Sigma(D))$ of D with the Scott topology (Alvarez-Manilla, Edalat, and Saheb-Djahromi 2000, cor. 4.3). (Below, we'll use the similar result, specialized to bifinite domains, which include Scott domains, by Abbes and Keimel (2006, thm. 2); for a general treatment of these kinds of extension results, see Keimel and Lawson (2005).) Since v is normalized, $\mu(D) = v(D) = 1$, so μ is a probability measure on D . Since v is max-normalized, $\max D$ is a countable intersection of Scott-open sets with v -value (and hence μ -value) 1. So $\max D \in \mathcal{B}(\Sigma(D))$ and, since μ is a probability measure, $\mu(\max D) = 1$. Hence

$$v(U) = \mu(U) = \mu(U \cap \max D) = \mu(V \cap \max D) = \mu(V) = v(V),$$

as needed. \square

Now, we can show the main lemma.

4.4.17. LEMMA. *For $i \in I$ and $U_i \in \Sigma(D_i)$, $v(f^{-1}(p_i^{-1}(U_i))) = v(p_i^{-1}(U_i))$.*

PROOF. By lemma 4.4.15, there is $j_0 \geq i$ such that for all $j \geq j_0$,

$$v\left(p_j^{-1}(f_j^{-1}p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j)\right) = v_i(U_i) \quad (4.4)$$

Moreover, by lemma 4.4.10 (3), there is $j_1 \geq i$ such that, for all $j \geq j_1$,

$$\bar{f}^{-1}(p_i^{-1}(U_i) \cap \max D) = \bigcup_{a_j \in R_{ij}^{-1}U_i} p_j^{-1}(a_j) \cap \max D. \quad (4.5)$$

By directedness, let $j \geq j_0, j_1 \geq i$. By lemma 4.4.14 (2) and \bar{f} being the restriction of f to $\max D$, we have, using (4.5),

$$\begin{aligned} f^{-1}(p_i^{-1}(U_i)) \cap \max D &= \bar{f}^{-1}(p_i^{-1}(U_i) \cap \max D) \\ &= \bigcup_{a_j \in R_{i_j}^{-1}U_i} p_j^{-1}(a_j) \cap \max D \\ &= p_j^{-1}(f_j^{-1}p_{i_j}^{-1}(\downarrow \max U_i)) \cap \max D \\ &= p_j^{-1}(f_j^{-1}p_{i_j}^{-1}(\downarrow \max U_i) \cap \max D_j) \cap \max D \end{aligned}$$

So by lemma 4.4.16, after removing ' $\cap \max D$ ' on both sides, the two resulting open sets have the same v -value, so, using (4.4),

$$\begin{aligned} v\left(f^{-1}(p_i^{-1}(U_i))\right) &= v\left(p_j^{-1}(f_j^{-1}p_{i_j}^{-1}(\downarrow \max U_i) \cap \max D_j)\right) \\ &= v_i(U_i) = v(p_i^{-1}(U_i)), \end{aligned}$$

as needed. \square

As mentioned, the general case now follows by a simple and general argument.

4.4.18. LEMMA. *For $U \in \Sigma(D)$, $v(f^{-1}U) = v(U)$.*

PROOF. First, let's assume U is a basic open set: $U = \bigcap_{k=1}^m p_{i_k}^{-1}(U_{i_k})$ with $i_k \in I$ and $U_{i_k} \in \Sigma(D_{i_k})$ for $k = 1, \dots, m$. Since I is directed, let $j \geq i_1, \dots, i_m$. So

$$U = \bigcap_{k=1}^m p_j^{-1}(p_{i_k j}^{-1}(U_{i_k})) = p_j^{-1}\left(\bigcap_{k=1}^m p_{i_k j}^{-1}(U_{i_k})\right).$$

So $U_j := \bigcap_{k=1}^m p_{i_k j}^{-1}(U_{i_k}) \in \Sigma(D_j)$, whence, by lemma 4.4.17, $v(U) = v(p_j^{-1}U_j) = v(f^{-1}p_j^{-1}U_j) = v(f^{-1}U)$.

Now, let U be an arbitrary Scott-open of D . So $U = \bigcup_{l \in L} U_l$ for basic opens U_l and an index set L . Without loss of generality, $\{U_l : l \in L\}$ is directed (otherwise consider the family of finite unions of the U_l 's). Note that then also $f^{-1}(U_l)$ is a directed family of open sets in D . Then

$$\begin{aligned} v(U) &= v\left(\bigcup_{l \in L} U_l\right) = \sup_{l \in L} v(U_l) = \sup_{l \in L} v(f^{-1}(U_l)) \\ &= v\left(\bigcup_{l \in L} f^{-1}(U_l)\right) = v\left(f^{-1}\left(\bigcup_{l \in L} U_l\right)\right) = v(f^{-1}(U)), \end{aligned}$$

as needed. \square

This shows that $\mathfrak{D} = (D, v, f)$ is valuation-preserving. Thus, the proof of theorem 4.4.8 is complete.

Again, note that, to conclude that $f \upharpoonright \max D$ is injective and surjective (claim (6) above), we've only used that the diagram is downward deterministic and that the \mathfrak{D}_i are max-surjective, respectively. And to conclude that f is valuation-preserving (claim (7) above), we've only used that the diagram is eventually valuation-preserving. So we could be more precise and specify further cases between general and standard, but, again, won't do so explicitly.

4.4.4 Definition of dynamical domains

In accordance with the tradition of calling the kinds of dcpos under study simply 'domains', we define:

4.4.19. DEFINITION. A *dynamical domain* is a dynamical dcpo \mathfrak{D} that is the $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of a finitary dynamical expanding system. A *standard dynamical domain* is a dynamical dcpo \mathfrak{D} that is the $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of a standard finitary dynamical expanding system.

The full subcategory of dSCO_{nm} whose objects are dynamical domains is denoted dDOM . The full subcategory of dSCO_{nm} whose objects are standard dynamical domains is denoted dDOM_{s} . (Note that morphisms in dDOM and dDOM_{s} are not required to be projections.)

In particular, if $\mathfrak{D} = (D, v, f)$ is a dynamical domain, f is max-preserving, and if \mathfrak{D} also is standard, then f is bijective on $\max D$ and valuation-preserving.⁵¹

Regarding examples, we'll see next—as already suggested in section 4.3—how dynamical systems induce dynamical domains through observation. But in appendix B (esp. section B.1), we also provide a detailed example of a dynamical domain constructed in a purely domain-theoretic way.

4.5 The system modeled by a dynamical domain

We show that a (standard) dynamical domain is a computational model for a (standard) dynamical system.

4.5.1. THEOREM. *Let $\mathfrak{D} = (D, v, f)$ be a dynamical domain. Then*

⁵¹The former follows since \mathfrak{D} is in $\text{dSCO}_{\text{nm}}^{\text{p}}$. For the latter, since \mathfrak{D} is a $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of a standard finitary dynamical expanding system, \mathfrak{D} is isomorphic in $\text{dSCO}_{\text{nm}}^{\text{p}}$ to the $\mathfrak{D}' = (D', v', f')$ constructed in theorem 4.4.8. By that theorem, f' is bijective on $\max D'$ and preserves the valuation v' . And these two properties are preserved by an isomorphism $\alpha : \mathfrak{D}' \rightarrow \mathfrak{D}$: Since f and f' are max-preserving, max-semi-equivariance implies max-equivariance, so, on $\max D'$, $\alpha \circ f' = f \circ \alpha$, whence $f = \alpha \circ f' \circ \alpha^{-1}$. Hence $f \upharpoonright \max D = \alpha \upharpoonright \max D' \circ f' \upharpoonright \max D' \circ \alpha^{-1} \upharpoonright \max D$ is bijective qua composition of bijective functions. And f is valuation-preserving qua composition of valuation-preserving functions.

1. $\max D$ with the relative Scott topology is a compact zero-dimensional Polish space,
2. f restricts to a continuous function on $\max D$,
3. v determines a unique probability measure μ_v on $\mathcal{B}(D, \Lambda)$ extending v ,
4. $\mathcal{B}(\max D) \subseteq \mathcal{B}(D, \Lambda)$ and $\mu_v \upharpoonright \mathcal{B}(\max D)$ is a probability measure on $\max D$.

Thus, we obtain the compact zero-dimensional measured topological system

$$\mathsf{S}(\mathfrak{D}) := (\max D, \Sigma(D) \upharpoonright \max D, \mu_v \upharpoonright \mathcal{B}(\max D), f \upharpoonright \max D)$$

which induces the general dynamical system $\mathsf{JS}(\mathfrak{D})$. Moreover, if \mathfrak{D} is standard, then both the topological system $\mathsf{S}(\mathfrak{D})$ and the dynamical system $\bar{\mathsf{J}}(\mathfrak{D})$ are standard.

We call $\mathsf{S}(\mathfrak{D})$ the *topological system modeled by \mathfrak{D}* and $\mathsf{JS}(\mathfrak{D})$ (resp., $\bar{\mathsf{J}}(\mathfrak{D})$ in case \mathfrak{D} is standard) the *dynamical system model by \mathfrak{D}* .

PROOF. By definition, \mathfrak{D} is the $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of a finitary dynamical expanding system $(\mathfrak{D}_i, p_{ij})_I$. So \mathfrak{D} is isomorphic in $\text{dSCO}_{\text{nm}}^{\text{p}}$ to the $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit $\mathfrak{D}' = (D', v', f')$ described in theorem 4.3.11. In particular, D and D' are isomorphic as deposes, so, since $\max D'$ is Lawson-closed, also $\max D$ is.

Ad (1). The relative Scott topology on the maximal elements of a Scott domain coincides with the relative Lawson topology and forms a Polish space (Lawson 1997, also see Edalat and Heckmann 1998; Martin 1998) which, due to algebraicity, also is zero-dimensional (Flagg and Kopperman 1997). Finally, since $\max D$ is Lawson-closed and the Lawson topology of a Scott domain is compact, $\max D$ also is compact. Hence, also $\max D$ with the relative Scott topology is compact.

Ad (2). Since f is max-preserving $f \upharpoonright \max D$ is a function from $\max D$ to $\max D$, and it is continuous with respect to the relative Scott topology since $f : D \rightarrow D$ is continuous in the Scott topology on D .

Ad (3). As shown by Abbes and Keimel (2006, thm. 2), for every continuous valuation v on a Scott domain D (or any bifinite domain), there exists a unique Radon measure μ on $\mathcal{B}(D, \Lambda(D))$ extending v on $\Sigma(D)$. Since D is ω -algebraic and coherent, the Lawson topology on D is second-countable, compact, Hausdorff, whence Polish, so any measure on $\mathcal{B}(D, \Lambda)$ is Radon, whence μ_v is the unique measure on $\mathcal{B}(D, \Lambda)$ extending v . And since $\mu_v(D) = v(D) = 1$ (because the valuation is normalized), μ_v is a probability measure.

Ad (4). Since $\max D$ is Lawson-closed, $\mathcal{B}(D, \Lambda)$ is a σ -algebra that, among others, contains all sets of the form $U \cap \max D$ for $U \in \Lambda(D)$. Now, $\mathcal{B}(\max D)$ is the smallest such σ -algebra, whence $\mathcal{B}(\max D) \subseteq \mathcal{B}(D, \Lambda)$. So it remains to show $\mu_v(\max D) = 1$. Indeed, since v is max-normalized, $\max D$ can be written as countable intersection of Scott-open sets with v -value 1, and hence μ_v -measure 1, so $\max D$ is a countable intersection of sets of full measure and hence has full measure, too.

For the ‘moreover’ part, since \mathfrak{D} is standard, $f \upharpoonright \max D$ is bijective and f is valuation-preserving. So, for $U \in \Sigma(D)$,

$$\begin{aligned} \mu_v((f \upharpoonright \max D)^{-1}(U \cap \max D)) &= \mu_v(f^{-1}(U) \cap \max D) = \mu_v(f^{-1}(U)) \\ &= v(f^{-1}(U)) = v(U) = \mu_v(U) = \mu_v(U \cap \max D). \end{aligned}$$

So the Borel probability measures

$$\mu_v \upharpoonright \mathcal{B}(\max D) \quad \text{and} \quad \mu_v \upharpoonright \mathcal{B}(\max D)(f \upharpoonright \max D)^{-1}$$

on $\mathcal{B}(\max D)$ agree on the open sets of $\max D$, and hence on all of $\mathcal{B}(\max D)$ (see, e.g., Bogachev 2007a, lem. 7.1.2, p. 68). So $f \upharpoonright \max D$ is measure-preserving. \square

Note that, to conclude that μ_v is preserved by $f \upharpoonright \max D$, we’ve only used that f is valuation-preserving.

4.6 Dynamical domain models for systems

Now we get to putting all the pieces together and build dynamical domains for both dynamical and topological systems.

4.6.1 For dynamical systems

4.6.1. DEFINITION. Let \mathfrak{X} be an abstract dynamical system and \mathcal{B} a countable measurable basis for \mathfrak{X} . Build $\mathfrak{D}_i = (D_i, v_i, f_i)$ and $p_{ij} : D_j \rightarrow D_i$ as in theorem 4.3.11. Then $(\mathfrak{D}_i, p_{ij})_{I(\mathcal{B})}$ is a finitary dynamical expanding system.⁵² Let $\mathfrak{D} = (D, v, f)$ be the $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of $(\mathfrak{D}_i, p_{ij})_{I(\mathcal{B})}$ as constructed in theorem 4.4.8. So $\text{D}(\mathfrak{X}, \mathcal{B}) := \mathfrak{D}$ is a dynamical domain which we call the *observation domain* of $(\mathfrak{X}, \mathcal{B})$. We call

$$\begin{aligned} \varphi : X &\rightarrow \max D \\ x &\mapsto \langle \{\mathcal{O}_i(x) : i \in I(\mathcal{B})\} \rangle \end{aligned}$$

the *canonical embedding* of \mathfrak{X} into $\text{S}(\text{D}(\mathfrak{X}, \mathcal{B}))$.⁵³

Here are the main facts about the canonical embedding.

⁵²Theorem 4.3.11 states, in the terminology of section 4.4, that $(\mathfrak{D}_i, p_{ij})_I$ is a finitary dynamical expanding system: the \mathfrak{D}_i are max-normalized finite dynamical Scott domains, the p_{ij} are dynamical morphisms and projections commuting appropriately, whence $(\mathfrak{D}_i, p_{ij})_I$ is a dynamical expanding system of finite max-normalized dynamical Scott domains which, moreover, has a countable index set $I(\mathcal{B})$ (since \mathcal{B} is countable) and is upward deterministic.

⁵³This is a well-defined function because, for $i \leq j$ in I , we have $\{\mathcal{O}_i(x)\} \in \max D_i$ and $p_{ij}(\{\mathcal{O}_j(x)\}) = \{\mathcal{O}_i(x)\}$, whence $\varphi(x) \in \max D$.

4.6.2. LEMMA. *Let $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ be an abstract dynamical system and \mathcal{B} a countable measurable basis for \mathfrak{X} . Write $I := I(\mathcal{B})$. Let $\mathfrak{D} = (D, v, f)$ be the observation domain built over diagram $(\mathfrak{D}_i, p_{ij})_I$ with $\mathfrak{D}_i = (D_i, v_i, f_i)$. Let $\varphi : X \rightarrow \max D$ be the canonical embedding. Then*

1. *For a nonempty basic open $U \cap \max D$ of $\max D$ with $U = \bigcap_{k=1}^m p_{i_k}^{-1} U_{i_k}$ (where $i_k \in I$ and $U_{i_k} \in \Sigma(D_{i_k})$ for $k = 1, \dots, m$), we have*

$$\varphi^{-1}(U \cap \max D) = \bigcap_{k=1}^m \bigcup_{l=1}^{n_k} [x_l^{(k)}]_{i_k}$$

for some $n_k \geq 1$ and $x_1^{(k)}, \dots, x_{n_k}^{(k)} \in X$ ($k = 1, \dots, m$).

2. *φ is measurable: for $B \in \mathcal{B}(\max D)$, $\varphi^{-1}(B) \in \mathcal{A}$.*
3. *φ is measure-preserving: for $B \in \mathcal{B}(\max D)$, $\mu(\varphi^{-1}(B)) = \mu_v(B)$.*
4. *φ is equivariant: for $x \in X$, $\varphi(T(x)) = (f \upharpoonright \max D)(\varphi(x))$.*
5. *$\varphi(X) \subseteq \max D$ is dense.*

PROOF. Ad (1). Since $U \cap \max D$ is nonempty, each U_{i_k} is nonempty, so $\max U_{i_k} \subseteq \max D_{i_k}$ is finite nonempty. Write $\max U_{i_k} = \{\{\mathcal{O}_{i_k}(x_1^{(k)})\}, \dots, \{\mathcal{O}_{i_k}(x_{n_k}^{(k)})\}\}$ with $n_k \geq 1$ and $x_1^{(k)}, \dots, x_{n_k}^{(k)} \in X$. Hence

$$\begin{aligned} x \in \varphi^{-1}(U \cap \max D) &\Leftrightarrow \forall k \in \{1, \dots, m\} : \varphi(x)(i_k) \in U_{i_k} \\ &\Leftrightarrow \forall k \in \{1, \dots, m\} : \{\mathcal{O}_{i_k}(x)\} \in \max U_{i_k} \\ &\Leftrightarrow \forall k \in \{1, \dots, m\} \exists l \in \{1, \dots, n_k\} : \mathcal{O}_{i_k}(x) = \mathcal{O}_{i_k}(x_l^{(k)}) \\ &\Leftrightarrow \forall k \in \{1, \dots, m\} \exists l \in \{1, \dots, n_k\} : x \in [x_l^{(k)}]_{i_k}. \end{aligned}$$

So, $\varphi^{-1}(U \cap \max D) = \bigcap_{k=1}^m \bigcup_{l=1}^{n_k} [x_l^{(k)}]_{i_k}$.

Ad (2). We need to show that φ -preimages of open sets are Borel. So let $U \subseteq \max D$ be open. Since $\max D$ is second-countable, U can be written as countable union of basic open sets $U_n \cap \max D$. So $\varphi^{-1}(U) = \bigcup_n \varphi^{-1}(U_n \cap \max D)$. Without loss of generality, all $U_n \cap \max D$ are nonempty (if, after discarding all empty ones, nothing remains, $\varphi^{-1}U = \emptyset \in \mathcal{A}$ since \mathcal{A} is a σ -algebra). Recall from lemma 4.3.9, the equivalence classes $[x]_i$ are in \mathcal{A} . So, by (1), the $\varphi^{-1}(U_n \cap \max D)$ are in \mathcal{A} qua finite intersection of finite unions of elements from \mathcal{A} . Hence $\varphi^{-1}(U)$ is in \mathcal{A} qua countable union of elements from \mathcal{A} .

Ad (3). Since $\varphi : X \rightarrow \max D$ is measurable where $\max D$ is equipped with the relative Scott—i.e., relative Lawson—topology, also $\varphi : X \rightarrow D$ is measurable where D is equipped with the Lawson topology. Since μ is a probability measure on X , $\kappa := \mu(\varphi^{-1}(\cdot))$ is a probability measure on $\mathcal{B}(D, \Lambda)$. Since (D, Σ) is second

countable (by ω -algebraicity), any probability measure κ on $\mathcal{B}(D, \Lambda)$ restricts to a continuous valuation on $\Sigma(D)$.⁵⁴ So $w : \Sigma(D) \rightarrow [0, \infty]$, $w(U) := \kappa(U)$ is a continuous valuation.

Moreover, we claim that, for $i \in I$ and $U_i \in \Sigma(D_i)$, we have $v_i(U_i) = w(p_i^{-1}(U_i))$. Indeed, write $\max(U_i) = \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\}$. Then

$$\bigcup_{k=1}^m [x_k]_i = \varphi^{-1}(p_i^{-1}(U_i))$$

because $y \in [x_k]_i$ for some $1 \leq k \leq m$ iff $\mathcal{O}_i(y) = \mathcal{O}_i(x_k)$ for some $1 \leq k \leq m$ iff $\varphi(y)(i) = \{\mathcal{O}_i(y)\} \in \max U_i$ iff $y \in \varphi^{-1}(p_i^{-1}(U_i))$. Hence

$$v_i(U_i) = \sum_{k=1}^m \mu[x_k]_i = \mu \bigcup_{k=1}^m [x_k]_i = \mu \varphi^{-1}(p_i^{-1}(U_i)) = \kappa(p_i^{-1}(U_i)) = w(p_i^{-1}(U_i)).$$

Now, v is the unique valuation on D with $v_i(U_i) = v(p_i^{-1}(U_i))$, so $v = w$. Moreover, μ_v is the unique measure on $\mathcal{B}(D, \Lambda)$ extending v , and κ also is a measure on $\mathcal{B}(D, \Lambda)$ extending $w = v$, so $\mu_v = \kappa$. In particular, for $B \in \mathcal{B}(\max D) \subseteq \mathcal{B}(D, \Lambda)$, we have $\mu_v(B) = \kappa(B) = \mu(\varphi^{-1}(B))$, as needed.

Ad (4). Let $x \in X$, and show $\varphi \circ T(x) = (f \upharpoonright \max D) \circ \varphi(x)$. Recall that, for $a \in \max D$, $f(a)$ is defined as the unique $b \in \max D$ such that, for all $i \in I$, $b(i) \geq f_i(a(i))$. Since $a := \varphi(x) \in \max D$, it hence suffices to show that, for $b := \varphi(T(x)) \in \max D$ and $i \in I$, we have $b(i) \geq f_i(a(i))$. Indeed, we have

$$b(i) = \{\mathcal{O}_i(Tx)\} \subseteq \{\mathcal{O}_i(Ty) : \mathcal{O}_i(y) \in \{\mathcal{O}_i(x)\}\} = f_i(\varphi(x)(i)) = f_i(a(i)).$$

Ad (5). Let $a \in \max D$ and $U \cap \max D$ a basic open set with $U = \bigcap_{k=1}^m p_{i_k}^{-1} U_{i_k}$ (where $i_k \in I$ and $U_{i_k} \in \Sigma(D_{i_k})$ for $k = 1, \dots, m$). Assume $a \in U \cap \max D$ and show $(U \cap \max D) \cap \varphi(X) \neq \emptyset$. Since I is directed, let $j \geq i_1, \dots, i_m$. Write $\{\mathcal{O}_j(x)\} = a(j) \in \max D_j$ for some $x \in X$. We show $\varphi(x) \in U \cap \max D$. So, for $k \in \{1, \dots, m\}$ we have to show $\varphi(x) \in p_{i_k}^{-1} U_{i_k}$. Indeed,

$$\varphi(x)(i_k) = \{\mathcal{O}_{i_k}(x)\} = p_{i_k j}(\{\mathcal{O}_j(x)\}) = p_{i_k j}(a(j)) = a(i_k) \in U_{i_k},$$

⁵⁴ Since κ is a Borel measure on (D, Λ) and $\Sigma \subseteq \Lambda$, the function $v := \kappa \upharpoonright \Sigma : \Sigma \rightarrow [0, \infty]$ satisfies the strictness, monotonicity, and modularity axiom for valuations (also see Keimel and Lawson 2005, p. 58). So we need to check that v is continuous (also see Edalat 1995b, cor. 5.3). Let $\{U_k\}_{k \in K}$ be a directed family in $\Sigma(D)$, and show $v(\bigcup_k U_k) \leq \sup_k v(U_k)$ (we have \geq by monotonicity). Let's first assume K is countable, so, after possibly reindexing, $K = \omega$. Define $V_0 := U_0$ and $V_k := U_k \setminus (U_0 \cup \dots \cup U_{k-1})$. So $V_k \in \mathcal{B}(D, \Lambda)$ and $\bigcup_k V_k$ is a disjoint union that equals $\bigcup_k U_k$. So $v(\bigcup_k U_k) = \kappa(\bigcup_k V_k) = \sum_k \kappa(V_k) = \lim_{k \rightarrow \infty} \sum_{i=0}^k \kappa(V_i) = \lim_{k \rightarrow \infty} \kappa(V_0 \cup \dots \cup V_k) \leq \sup_k v(U_k)$, where the last step follows since, for $k \geq 0$, there is, by directedness, $U_n \supseteq U_0, \dots, U_k \supseteq V_0, \dots, V_k$, so $\kappa(V_0 \cup \dots \cup V_k) \leq \kappa(U_n) = v(U_n) \leq \sup_k v(U_k)$. Now, let K be arbitrary. Since Σ is second countable, there is $K_0 \subseteq K$ countable with $\bigcup_K U_k = \bigcup_{K_0} U_k$. Let $\{V_l\}_L$ be the family of finite unions of sets from $\{U_k : k \in K_0\}$. Then $\{V_l\}$ is directed in Σ with countable index set. So $v(\bigcup_K U_k) = v(\bigcup_{K_0} U_k) = v(\bigcup_L V_l) \leq \sup_L v(V_l) \leq \sup_K v(U_k)$, where the last step follows since, for $V_l = U_{k_1} \cup \dots \cup U_{k_m}$, there is, by directedness, $U_n \supseteq U_{k_1}, \dots, U_{k_m}$, so $v(V_l) \leq v(U_n) \leq \sup_K v(U_k)$.

as needed. \square

To render the canonical embedding injective, we'll consider separating measurable bases. Recall that \mathcal{B} is separating if, for all $x \neq y$ in X , there is $U \in \mathcal{B}$ such that $x \in U$ but $y \notin U$. In the 'non-pathological' dynamical systems that are built over standard Borel spaces, we can always find such bases: If (X, \mathcal{A}) is a standard Borel space, then it has a countable and separating measurable basis.⁵⁵ In the Lebesgue case, we need to disregard null sets: A measurable basis \mathcal{B} for a standard dynamical system $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ is *separating mod 0* if there is $M \in \mathcal{A}$ with $\mu(M) = 1$ and, for all $x \neq y$ in M , there is $U \in \mathcal{B}$ such that $x \in U$ but $y \notin U$. Then \mathfrak{X} has a countable and separating mod 0 measurable basis.⁵⁶

Now we can state the main theorem on building dynamical domain models for dynamical systems.

4.6.3. THEOREM. 1. *Let \mathfrak{X} be a general dynamical system and \mathcal{B} a countable and separating measurable basis. Let \mathfrak{D} be the observation domain. Then the canonical embedding φ is an isomorphism $\mathfrak{X} \rightarrow \mathbf{JS}(\mathfrak{D})$ of general dynamical systems.*

2. *Let \mathfrak{X} be a standard dynamical system and \mathcal{B} a countable and separating mod 0 measurable basis that is forward closed (under the dynamics of \mathfrak{X}). Let \mathfrak{D} be the observation domain. Then \mathfrak{D} is standard and the canonical embedding φ is an isomorphism $\mathfrak{X} \rightarrow \overline{\mathbf{JS}}(\mathfrak{D})$ of standard dynamical systems.*

PROOF. Ad (1). Write (X, \mathcal{A}, μ, T) and $\mathfrak{D} = (D, v, f)$. We know that $\mathbf{J}(\mathbf{S}(\mathfrak{D}))$ is a general dynamical system. So we need to show that $\varphi : X \rightarrow \max D$ is an isomorphism of (abstract) dynamical system.

From lemma 4.6.2, we already know that φ (considered with domain $M := X$ which is invariant and of full measure) is measure-preserving and equivariant. It also is injective:

Assume $x \neq y$ in X . Since \mathcal{B} is separating, there is $U \in \mathcal{B}$ such that $x \in U$ and $y \notin U$. Define $\mathcal{C} := \{U, X\}$ (which is a finite \mathcal{B} -cover) and $n := 1$, whence $i := (n, \mathcal{C}) \in I(\mathcal{B})$. We have $\mathcal{O}_i(x) \neq \mathcal{O}_i(y)$ since $t := \langle U \rangle \in \mathcal{O}_i(x)$ but $t \notin \mathcal{O}_i(y)$. So $\varphi(x)(i) = \{\mathcal{O}_i(x)\} \neq \{\mathcal{O}_i(y)\} = \varphi(y)(i)$, so $\varphi(x) \neq \varphi(y)$.

⁵⁵By definition (e.g. Kechris 1995, def. 12.5), there is a Polish topology τ on X such that $\mathcal{A} = \mathcal{B}(\tau)$. Since Polish spaces are second-countable, let $\mathcal{B} = \{U_0, U_1, \dots\}$ be a countable basis for (X, τ) , which we can assume to be closed under finite intersection. So $\mathcal{B} \subseteq \mathcal{B}(\tau)$, and every open set can be written as a countable union of elements from \mathcal{B} , whence $\sigma(\mathcal{B}) = \mathcal{B}(\tau)$. So \mathcal{B} is a countable measurable basis for (X, \mathcal{A}) and it is separating since Polish spaces are Hausdorff.

⁵⁶We know that (X, \mathcal{A}, μ) is isomorphic mod 0 to the completion of a Borel probability space $(Y, \mathcal{B}(\sigma), \nu)$. Let $\varphi : X \rightarrow Y$ be that isomorphism with domain $M \subseteq X$. Let \mathcal{C} be a countable and separating measurable basis for (Y, \mathcal{B}) . Set $\mathcal{B} := \varphi^{-1}\mathcal{C} \cup \{X\}$. Since \mathcal{C} is countable and closed under finite intersection, also \mathcal{B} is. Since φ is measurable, \mathcal{B} is measurable. And \mathcal{B} separates points in M : if $x \neq y$ in M , then $\varphi(x) \neq \varphi(y)$ in Y , so there is $C \in \mathcal{C}$ with $\varphi(x) \in C$ and $\varphi(y) \notin C$, so $\varphi^{-1}(C) \in \mathcal{B}$ with $x \in \varphi^{-1}(C)$ and $y \notin \varphi^{-1}(C)$.

Finally, define the codomain $N := \varphi(X)$, so $\varphi : M \rightarrow N$ is bijective. Thus, to show that φ is an isomorphism, it remains to show that N is invariant and of full measure.

Indeed, since $\varphi : X \rightarrow \max D$ is an injective Borel-measurable function between standard Borel spaces, $\varphi(X)$ is Borel (e.g. Kechris 1995, cor. 15.2). Since φ is measure-preserving, we have, since $X \subseteq \varphi^{-1}\varphi(X)$,

$$\mu_v(\varphi(X)) = \mu\varphi^{-1}\varphi(X) \geq \mu(X) = 1,$$

so $\mu_v(\varphi(X)) = 1$. Finally, since φ is equivariant, $N = \varphi(X)$ is invariant (if $\varphi(x) \in N$, then $(f \upharpoonright \max D)(\varphi(x)) = \varphi(T(x)) \in \varphi(X) = N$).

Ad (2). Under the assumptions, T is bijective and measure-preserving, and \mathcal{B} is a measurable basis that is forward T -closed, so, by theorem 4.3.11, the finitary dynamical expanding system $(\mathfrak{D}_i, p_{ij})_{I(\mathcal{B})}$ with which \mathfrak{D} is built is standard, so \mathfrak{D} is a standard dynamical domain. Hence, $\bar{J}(\mathcal{S}(\mathfrak{D}))$ is a standard dynamical system.

Let $M \in \mathcal{A}$ be the set of full measure on which \mathcal{B} is separating. We can assume that M is invariant: Otherwise take $M' := \bigcap_{k \in \mathbb{Z}} T^k(M)$. Then \mathcal{B} is still separating on M' (since $M' \subseteq M$), M' has full measure (qua countable intersection of sets of full measure, using that T is measure-preserving), and $T(M') = M'$ (by construction and since T is bijective).

So it remains to show that φ with domain M and codomain $N := \varphi(M)$ is an isomorphism between (abstract) dynamical systems. Its domain M is invariant and of full measure. By lemma 4.6.2, φ is measure-preserving and equivariant (on M). As above, $\varphi \upharpoonright M$ is injective: If $x \neq y$ in M , then, since \mathcal{B} is separating on M , there is $U \in \mathcal{B}$ such that $x \in U$ and $y \notin U$, whence $i := (1, \{U, X\}) \in I(\mathcal{B})$ yields $\varphi(x)(i) \neq \varphi(y)(i)$. So we need to show that N is invariant and of full measure. Since $M \subseteq X$ is of full measure, $(M, \mathcal{A} \upharpoonright M, \mu)$ is again a Lebesgue space (since we only removed a null set, it still is isomorphic mod 0 to an ordinary Lebesgue interval with countably many point masses). So $\varphi \upharpoonright M$ is a measure-preserving injection between Lebesgue spaces, so its image $N = \varphi(M)$ is measurable (de la Rue 1993, thm. 3.5). As above, since φ is measure-preserving and equivariant, this image has full measure and is invariant. \square

By using the existence of separating bases, we can state the above theorem more concisely and as our main result:

4.6.4. COROLLARY. *For every (standard) dynamical system \mathfrak{X} , there is a (standard) dynamical domain \mathfrak{D} such that the (standard) dynamical system modeled by \mathfrak{D} is metrically isomorphic to \mathfrak{X} .*

PROOF. For the general case, i.e., if \mathfrak{X} is a general dynamical system, there is, as noted above, a countable and separating measurable basis \mathcal{B} for \mathfrak{X} , so, by theorem 4.6.3 (1), there is a dynamical domain \mathfrak{D} such that \mathfrak{X} is metrically isomorphic to the dynamical system $\mathcal{J}\mathcal{S}(\mathfrak{D})$ modeled by \mathfrak{D} .

For the standardized case, i.e., if $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ is a standard dynamical system, we need to show that there is a countable and separating mod 0 measurable basis \mathcal{B} for \mathfrak{X} that is forward T -closed. Then theorem 4.6.3 (2) implies that there is a standard dynamical domain \mathfrak{D} such that \mathfrak{X} is metrically isomorphic to the standard dynamical system $\text{JS}(\mathfrak{D})$ modeled by \mathfrak{D} .

Indeed, since (X, \mathcal{A}, μ) is a Lebesgue space, it has, as noted above, a countable and separating mod 0 measurable basis \mathcal{B}_0 . Since $T : X \rightarrow X$ is an injective measure-preserving functions between two Lebesgue spaces, it preserves measurability (de la Rue 1993, thm. 3-5). So, for all $A \in \mathcal{A}$, we have $T(A) \in \mathcal{A}$. Let \mathcal{B} be the set of finite intersections of sets in $\bigcup_{k \geq 0} T^k \mathcal{B}_0$. So \mathcal{B} is a basis (by construction), countable (qua finite intersections of sets from a countable union of countable sets), separating mod 0 (it contains \mathcal{B}_0 which is separating mod 0), measurable (its elements are in \mathcal{A}), and forward T -closed (if $A \in \mathcal{B}$, then $A = T^{k_1} A_1 \cap \dots \cap T^{k_n} A_n$, so $T(A) = T^{k_1+1} A_1 \cap \dots \cap T^{k_n+1} A_n \in \mathcal{B}$ using that T is injective). \square

4.6.2 For topological systems

We state the result with some more detail that is available in the topological setting.

If $\mathfrak{X} = (X, \tau, \mu, T)$ is a measured topological system and \mathcal{B} a countable topological basis for \mathfrak{X} , then \mathcal{B} is a countable and separating measurable basis for the abstract dynamical system $\text{J}(\mathfrak{X}) = (X, \mathcal{B}(\tau), \mu, T)$. We define the *observation domain* $\text{D}(\mathfrak{X}, \mathcal{B})$ as the observation domain of $\text{J}(\mathfrak{X})$ with basis \mathcal{B} . In particular, we can still define the canonical embedding $\varphi : X \rightarrow \max D$. Here are its topological properties:

4.6.5. LEMMA. *Let $\mathfrak{X} = (X, \tau, \mu, T)$ be a zero-dimensional measured topological system and \mathcal{B} a countable topological basis for \mathfrak{X} consisting of clopen sets. Let $\mathfrak{D} = \text{D}(\mathfrak{X}, \mathcal{B}) = (D, v, f)$ be the observation domain and $\varphi : X \rightarrow \max D$ the canonical embedding. Then*

1. *The φ -preimages of clopen sets of $\text{S}(\mathfrak{D})$ can be written as Boolean combinations of equivalence classes $[x]_i \subseteq X$ (where $x \in X$ and $i \in I(\mathcal{B})$) and hence of sets in $\bigcup_{k \geq 0} T^{-k} \mathcal{B}$.⁵⁷*
2. *φ is relatively open (i.e., if $U \subseteq X$ is open, then there is an open $V \subseteq \max D$ such that $\varphi(U) = V \cap \varphi(X)$).*
3. *φ is an injective and relatively open homomorphism $\mathfrak{X} \rightarrow \text{S}(\mathfrak{D})$ of topological systems whose image is dense in $\max D$.*

⁵⁷The latter part follows from lemma 4.3.9 (2).

4. φ is surjective iff φ is a homeomorphism iff X is compact.

PROOF. Ad (1). First, we show that the clopen sets of $\max D$ are precisely finite unions of basic open sets $\bigcap_{k=1}^m p_{i_k}^{-1}(U_{i_k}) \cap \max D$. Indeed, these basic opens are clopen (whence also finite unions thereof) since the D_i are finite. Conversely, assume $U \subseteq \max D$ is clopen. Since it is open, it is a union of basic open sets $\{U_i\}$, whence the U_i cover U . Since U is a closed subset of a compact space, it is compact, so there is a finite subcover of $\{U_i\}$. Hence, U can be written as finite union of basic open sets.

Next, note that the claim holds for basic opens U : if U is empty, then $\varphi^{-1}(U) = \emptyset$ trivially is a Boolean combination (the triviality or empty disjunction), and if U is nonempty, lemma 4.6.2 (1) yields the claim.

Now, if $U \subseteq \max D$ is clopen, it is a finite union $U_1 \cup \dots \cup U_n$ of basic open sets, so $\varphi^{-1}(U) = \varphi^{-1}(U_1) \cup \dots \cup \varphi^{-1}(U_n)$ is (a finite union of) Boolean combinations of equivalence classes $[x]_i$.

Ad (2). It is enough to check this for the opens of the basis \mathcal{B} . So let $U \in \mathcal{B}$. If $U = \emptyset$, this is trivial, so let $x_0 \in U$. Consider $i := (1, \{U, X\}) \in I(\mathcal{B})$. Let $a_i := \{\mathcal{O}_i(x_0)\} \in \max D_i$ and $U_i := \{a_i\} \in \Sigma(D_i)$. Set $V := p_i^{-1}(U_i) \cap \max D$ which is open in $\max D$. We show $\varphi(U) = V \cap \varphi(X)$.

Indeed, if $a \in \varphi(U)$, then $a = \varphi(x)$ for some $x \in U$. So $\mathcal{O}_i(x) = \{\langle U \rangle, \langle X \rangle\} = \mathcal{O}_i(x_0)$, so

$$a(i) = \varphi(x)(i) = \{\mathcal{O}_i(x)\} = \{\mathcal{O}_i(x_0)\} = a_i \in U_i,$$

so $a \in V \cap \varphi(X)$. Conversely, if $a \in V \cap \varphi(X)$, then $a = \varphi(x) \in V$ for some $x \in X$. So $\{\mathcal{O}_i(x)\} = \varphi(x)(i) = a(i) \in U_i = \{a_i\}$, so $\{\mathcal{O}_i(x)\} = a_i = \{\mathcal{O}_i(x_0)\}$. So $\langle U \rangle \in \mathcal{O}_i(x_0) = \mathcal{O}_i(x)$, whence x follows $\langle U \rangle$, so $x \in U$. Hence $a = \varphi(x) \in \varphi(U)$.

Ad (3). By (1), φ is continuous: the clopens form a basis for $\max D$ (since it is zero-dimensional) and their preimages are Boolean combinations of equivalence classes $[x]_i$, which are clopen by lemma 4.3.9 (2). By (2), φ is relatively open. By lemma 4.6.2, φ is measure-preserving, and equivariant, with dense image. So it remains to show that φ is injective: This follows as in the proof of theorem 4.6.3 above since the topological basis is separating (if $x \neq y$ there is, since X is Hausdorff, a basic open set U such that $x \in U$ and $y \notin U$).

Ad (4). If φ is surjective, then it is open by (2) and by (3) it is a continuous bijection $X \rightarrow \max D$, so φ is a homeomorphism. If φ is a homeomorphism, then, since $\max D$ is compact, also X is compact. So the following remains: Assume X is compact, and show that φ is surjective.

Let $a \in \max D$ and find $x \in X$ such that $\varphi(x) = a$. Since a is maximal, $a = \langle \{\mathcal{O}_i(x_i)\} : i \in I(\mathcal{B}) \rangle$ for some $x_i \in X$.

We claim that $\{[x_i]_i : i \in I\}$ is a family of closed subsets of X with the finite intersection property. Indeed, by lemma 4.3.9 (2), each such equivalence class is closed (in fact clopen). And given $[x_{i_1}]_{i_1}, \dots, [x_{i_n}]_{i_n}$, let, by directedness,

$j \geq i_1, \dots, i_n$. Then $x_j \in \bigcap_{k=1}^n [x_{i_k}]_{i_k}$, since, for $k = 1, \dots, n$, we have $\{\mathcal{O}_{i_k}(x_j)\} = p_{i_k, j}(a(j)) = a(i_k) = \{\mathcal{O}_{i_k}(x_{i_k})\}$, so $x_j \in [x_{i_k}]_{i_k}$.

Now, since X is compact, there is some $x \in \bigcap_{i \in I} [x_i]_i$. Then, for each $i \in I$, $\varphi(x)(i) = \{\mathcal{O}_i(x)\} = \{\mathcal{O}_i(x_i)\} = a(i)$, so $\varphi(x) = a$. \square

Comments: First, the mathematical relevance of (1) is that it implies the continuity of the canonical embedding. The philosophical relevance (and the reason for stating it explicitly) is that the clopen sets of $\mathbf{S}(\mathfrak{D})$ are the ‘computable’ or ‘observable’ properties of $\max D$, so (1) states that the canonical embedding respects this ‘computability structure’: φ -preimages of clopens in $\mathbf{S}(\mathfrak{D})$ are Boolean combinations of T -preimages of elements of \mathcal{B} . In particular, if \mathcal{B} is logically and backward dynamically closed, then φ -preimages of clopens in $\mathbf{S}(\mathfrak{D})$ are possible \mathcal{B} -observations that can be made about the system \mathfrak{X} .

Second, note that item (4) shows that, as soon as the state space X is not compact (e.g., a dynamics on the irrational numbers), the canonical embedding φ fails to be surjective: there is $a = \langle \{\mathcal{O}_i(x_i)\} : i \in I(\mathcal{B}) \rangle$ in $\max D$ such that $a \neq \varphi(x)$ for all x . In words, a is a ‘refining’ sequence of observation histories that cannot be generated by a single state x and rather must come from different states x_i (whose difference we hence, in a sense, cannot observe).

As in the case of dynamical systems, the following now is immediate from lemma 4.6.5.

4.6.6. THEOREM. *1. If \mathfrak{X} is a zero-dimensional measured topological system and \mathcal{B} a countable clopen topological basis, then $\mathfrak{D} := \mathbf{D}(\mathfrak{X}, \mathcal{B})$ is a dynamical domain and the canonical embedding φ is an injective, dense, and relatively open morphism $\mathfrak{X} \rightarrow \mathbf{S}(\mathbf{D}(\mathfrak{X}, \mathcal{B}))$ of measured topological systems.*

2. If \mathfrak{X} is a standard zero-dimensional measured topological system and \mathcal{B} a countable clopen topological basis that is forward closed, then $\mathfrak{D} := \mathbf{D}(\mathfrak{X}, \mathcal{B})$ is a standard dynamical domain and the canonical embedding φ is an injective, dense, and relatively open morphism $\mathfrak{X} \rightarrow \mathbf{S}(\mathbf{D}(\mathfrak{X}, \mathcal{B}))$ of standard measured topological systems.

And again as in the case of dynamical systems, we get a more concise formulation using the existence of appropriate bases.

4.6.7. COROLLARY. *For every (standard) zero-dimensional measured topological system \mathfrak{X} , there is a (standard) dynamical domain \mathfrak{D} such that \mathfrak{X} can be densely and relatively openly embedded into the (standard) zero-dimensional and compact measured topological system modeled by \mathfrak{D} .*

PROOF. We need to show that every (standard) zero-dimensional measured topological system \mathfrak{X} has a countable clopen topological basis \mathcal{B} (that is closed under the dynamics of \mathfrak{X}). Then the corollary follows from theorem 4.6.6. Indeed, in

the general case, since X is zero-dimensional Polish, it has a basis of clopens which, by second-countability, can be chosen to be countable. After closing it under finite intersections, it is a countable clopen topological basis \mathcal{B} for \mathfrak{X} . In the standard case, we consider the family of finite intersections of $\bigcup_{k \geq 0} T^k \mathcal{B}$. This is a basis (closed under intersection), countable, forward closed (since T is injective), generates the topology (since it contains the basis \mathcal{B}), and clopen (since T is a homeomorphism it maps clopen sets to clopen sets). \square

Things come together particularly neatly in the following topological setting.

4.6.8. COROLLARY. *Let $\mathfrak{X} = (X, \tau, \mu, T)$ be a (standard) compact and zero-dimensional measured topological system and $\mathcal{B} = \text{Clp}(X)$ the set of clopen sets of X . Let $\mathfrak{D} = \text{D}(\mathfrak{X}, \mathcal{B})$ be the (standard) observation domain. Then the canonical embedding $\varphi : \mathfrak{X} \rightarrow \text{S}(\mathfrak{D})$ is an isomorphism of (standard) topological systems (measure-preserving and equivariant homeomorphism).*

PROOF. First, note that, since X is compact and second-countable, $\text{Clp}(X)$ is countable.⁵⁸ So \mathcal{B} is indeed a countable clopen topological basis of \mathfrak{X} and the observation domain can be build. If \mathfrak{X} is standard, then $\text{Clp}(X)$ is closed under T since the homeomorphism T maps clopen sets to clopen sets, hence \mathfrak{D} is standard. We already know that the canonical embedding is a morphism $\varphi : \mathfrak{X} \rightarrow \text{S}(\mathfrak{D})$ of (standard) topological systems, and, since X is compact, it is a homeomorphism by lemma 4.6.5 (4). \square

4.7 Conclusion

We've defined the category of dynamical domains in a domain-theoretic way as restricted limits of dynamical expanding systems of finite dynamical Scott domains. And we've shown that every (standard) dynamical system is modeled up to isomorphism by the (standard) dynamical domain built as observation domain from some countable and separating measurable basis of the system.

We end with some questions for future work.

(1) Category theory: Arguably the most pressing question by now is whether the two constructions S (the system modeled by a dynamical domain) and D (the observation domain of a system) actually form functors between the category of dynamical systems and the category of dynamical domains. After phrasing this more carefully, the answer will be yes, and, much better, they will be adjoint

⁵⁸Proof: Let (U_n) be a countable basis for X . If $U \in \text{Clp}(X)$, then, since U is open, it is a union of basic open sets, which can be written as finite union since U is compact (qua closed subset of a compact space). Thus, we can assign to each $U \in \text{Clp}(X)$ a unique finite sequence of elements of (U_n) . Since there are only countably many such sequences, $\text{Clp}(X)$ is countable.

(restricting to an equivalence for a natural subcategory of \mathbf{dDOM}). This will be the main result of the next chapter.

(2) Dynamical systems theory: The next pressing question is how our result can be used to understand dynamical systems. Can we, for example, transfer dynamical system concepts to the domain-theoretic setting and vice versa? Doing this for one of the most central concepts of dynamical system theory—namely, entropy—will be the topic of chapter 6.

(3) Computability theory: Can we give more content to the idea that dynamical domains provide *computational* models for dynamical systems? Computability- and recursion-theoretic concepts can be developed in domain theory by fixing an enumeration of the countable basis of a domain (for a summary, see Abramsky and Jung 1994, sec. 8.1.1). Such a basis is available to us since we work over Scott domains. Can we thus provide a computability theory for dynamical systems, analogously to how Edalat (1995b, sec. 3.2) provides an effective structure for the topological dynamical systems via their upper spaces? How would such a theory compare to existing ones like those of Blum, Shub, and Smale (e.g. 1989) or Pour-El and Richards (1989)?

(4) Domain theory: As indicated in section 4.2.1, to ‘do domain theory’ means, to an approximation, performing domain constructions (Abramsky and Jung 1994, sec. 7.2.1, p. 124). So how do the usual domain constructions—product, function space, limits, fixed points, powerdomains, etc.—lift to dynamical domains? Also, are there other characterizations of dynamical domains without using a limit (e.g., bifinite domains also have an order-theoretic characterization)? Moreover, is there other domain-theoretic structure that can be added to a dynamical \mathbf{dcpo} $\mathfrak{D} = (D, v, f)$? Specifically, inspired by Martin and Panangaden (2011), can the domain-theoretic notion of *measurement* be used to represent *geometric* (as opposed to topological or measure-theoretic) information—e.g., when considering diffeomorphisms on manifolds? Further, can dynamical domains be described—as common in domain theory—as being obtained from a few basic dynamical domains together with a few constructions, or from a ‘universal’ dynamical domain? Interpreted in terms of (non-symbolic) computation, would these constructions be a sign of a ‘programming language of machine learning’ (Cheung et al. 2018)? Would that be a start for a ‘program logic’ (or axiomatic semantics) for neural networks?⁵⁹ And would a universal dynamical domain be a non-symbolic analogue of a universal Turing machine?

(5) Measurement theory: In our construction of the observation domain of a dynamical system (X, \mathcal{A}, μ, T) , we used a *binary* notion of observation: a system state either is or is not in the measurable set representing the observation. In other words, we only used those measurable functions $f : X \rightarrow \mathbb{R}$ that are characteristic functions of some measurable set. What would an observation domain (and corresponding notion of dynamical domain) look like if we’d allow all measurable

⁵⁹Also cf. ‘domain theory in logical form’ (Abramsky 1991).

functions $f : X \rightarrow \mathbb{R}$ (of some L^p space)? Thus, one is in the setting of the operator-theoretic approach to ergodic theory (Eisner et al. 2015). How does this compare to the observational logic of Abramsky and Vickers (1993)?

(6) Learning theory: To come full circle, let's revisit the learning example from the introduction. Following up on footnote 7, a continuous dynamics T on a Polish space X naturally induces a dynamics $G(T) : G(X) \rightarrow G(X)$ on the Polish space $G(X)$ of probability measures on X with the weak topology; and this Giry functor G plays a crucial role the papers mentioned in the footnote. This also connects to the Krylov–Bogolioubov theory (see, e.g., Walters 1982, ch. 6) where, for compact metrizable X , the ‘averaged’ action of T on $G(X)$ is used to construct preserved measures. Can we use this (together with our dynamical domain model) to obtain meaningful preserved measures in the learning example? For uniqueness, can we restrict the state space of \mathfrak{X} to the ‘ μ -random’ elements to obtain that the ergodic theorem holds surely, rather than almost surely (Galatolo, Hoyrup, and Rojas 2010), and then conclude that the measure must be uniquely ergodic? (We come back to ergodicity and randomness in chapter 7.) Following up on footnote 3, can the ideas of question (5) be used to include loss functions of the learning algorithm into the domain representation?

Chapter 5

Systems and domains 2: Category

Abstract In the previous chapter, we've shown that every dynamical system is, up to isomorphism, modeled by some dynamical domain (a structure in the sense of domain theory). This provides a tool to analyze *individual* dynamical systems, but it doesn't capture yet *relationships* between dynamical systems (in the category of dynamical systems). In this chapter, we complete this tool by describing in detail the category-theoretic connections between systems and domains. The main result is that, roughly, the construction of the dynamical domain for a system and the construction of the system modeled by a dynamical domain are adjoint. Moreover, the category of dynamical systems is a localization of the category of compact zero-dimensional measured topological systems which, in turn, is equivalent to an interesting full and reflective subcategory of the category of dynamical domains.

5.1 Introduction

In the previous chapter, we've defined, as a general tool to analyze dynamical systems, the category of *dynamical domains*. These are structures in the sense of domain theory (a mathematical theory of computation). They can be seen as computational models for dynamical systems: for every dynamical domain \mathfrak{D} , we can naturally define the dynamical system $\mathbf{S}(\mathfrak{D})$ modeled by \mathfrak{D} . We've also shown that, for every dynamical system \mathfrak{X} , we can construct a dynamical domain $\mathbf{D}(\mathfrak{X})$, which we called the observation domain, such that \mathfrak{X} is isomorphic to $\mathbf{S}(\mathbf{D}(\mathfrak{X}))$.

This poses the question whether these two constructions, \mathbf{S} and \mathbf{D} , can be extended in a category-theoretic way: Do they form functors between the category of dynamical systems and the category of dynamical domains? If so, do they even form an adjunction, indicating the optimality of these constructions? The importance of the question is that, given a positive answer, we can not only translate questions about particular dynamical systems into domain-theoretic questions. We can also do this for questions about relationships between dynamical systems: for example, can one be decomposed into two factors of a particular kind

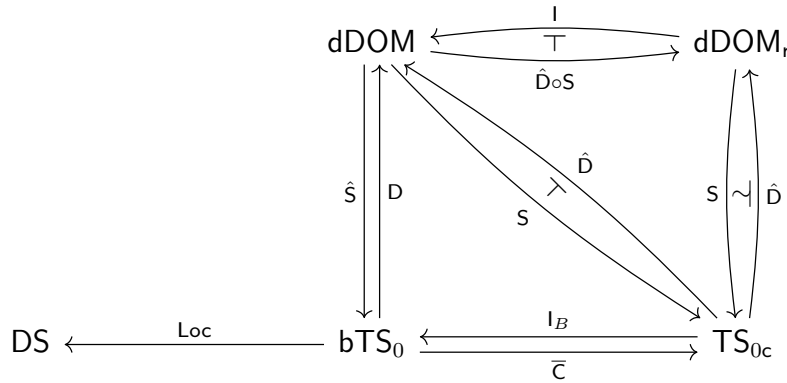


Figure 5.1: The main diagram: The categories and their functorial connections (here \vdash denotes adjunction, \sim adjoint equivalence, and Loc a localization). The diagram commutes up to natural isomorphism. It restricts from the general to the standard setting by adding ‘s’ to every category.

(as it occurs, e.g., in the weak Pinsker conjecture). Thus, as a slogan, all questions about dynamical systems can be translated into domain-theoretic questions, where the rich domain theory can be applied (and also vice versa).

The short summary of this chapter is that the answer is a resounding ‘yes’. The precise answer is given in the category-theoretic diagram of figure 5.1. We’ll refer to it as the *main diagram* and explain it in this introduction.

Even though this chapter builds on the previous, it is self-contained: In section 5.2.1, we provide a brief background of the required domain theory, dynamical systems theory, and category theory, and, in section 5.2.5, we provide a summary of the required results from the previous chapter. Moreover, since the present chapter is rather long, the chapter outline at the end of this introduction provides some suggested selective reading.

Explanation of the main diagram The bottom layer of the diagram consists of categories of dynamical systems (DS, bTS_0, TS_{0c}), while the top layer consists of categories of dynamical domains ($dDOM, dDOM_r$). All formal definitions of the categories are in section 5.2. The bottom layer constitutes the ‘foundation’ on which we can build the connection to dynamical domains given by the top layer. We first explain the diagram in a rough and ‘high-level’ way, and then we explain in more detail the two layers in turn.

The high-level explanation. In the previous chapter, we’ve established a translation between dynamical systems and dynamical domains via the ‘observation domain’ construction D and the ‘modeled system’ construction S . In this chapter, we’re concerned with extending this translation to morphisms between dynamical systems (or factors) and morphisms between dynamical domains. This works

particularly neatly if we additionally assume the dynamical systems to carry a zero-dimensional compact topology that determines its measurable sets and that is respected by the dynamics. (This is the category \mathbf{TS}_{0c} .) Then the constructions \mathbf{D} and \mathbf{S} extend to a translation that also takes morphisms into account: in category-theoretic terms, they form an adjunction. This is the diagonal of the main diagram and the main result of this chapter—and hence will be the minimal suggested reading. (The triangle above it analyzes this translation more precisely.)

However, since we’re concerned with relating dynamical systems and dynamical domains, we should also provide an answer without making the topological assumption. This is what the rest of the main diagram—i.e., the bottom layer—does. On the object level, the topological assumption provides no restriction: in the previous chapter, we’ve seen that, roughly, a dynamical system \mathfrak{X} is isomorphic to $\mathbf{SD}(\mathfrak{X})$, which satisfies the topological assumption. On the morphism level, however, the issue is that morphisms between dynamical systems are only defined up to sets of full measure and hence may ignore sets of states with probability 0. But this leeway is not allowed for the morphisms between systems that naturally come with the topological assumption. This mismatch is reconciled using the category-theoretic concept of a localization together with the topological concept of a compactification. Thus, we get a connection between the (unconstrained) dynamical systems and the dynamical systems with the topological assumption (where the above translation to dynamical domains starts). This connection is the content of the bottom layer.

The bottom layer. The category \mathbf{DS} of (measure-theoretic) dynamical systems consists of objects $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a probability space and $T : X \rightarrow X$ is a measurable function. Morphisms are, after possibly discarding null sets, measure-preserving functions that commute with the dynamics. To allow for a well-behaved isomorphism theory, we’ll make the mild assumption that the probability space is Borel (i.e., there is a Polish topology on X that generates the σ -algebra \mathcal{A}). This is a *very* general definition of a measure-theoretic dynamical system. (In chapter 4, we’ve motivated this choice as including systems, like learning dynamics, that otherwise wouldn’t be covered.) The standard definition (in ergodic theory) is to additionally assume that the probability space is a Lebesgue space (roughly, a completion of a Borel probability space) and that the dynamics T is both measure-preserving and invertible. In that case, we speak of a *standard* dynamical system. (In fact, all our categories have a standard version and the main diagram will restrict to them; but we’ll first explain the general case and then get to the standard case below.)

As mentioned, the previous chapter establishes a topological realization result (in the spirit of the Jewett–Krieger theorem in ergodic theory): every dynamical system \mathfrak{X} is, after a choice of measurable basis, isomorphic to the dynamical system induced by the compact zero-dimensional topological system $\mathbf{SD}(\mathfrak{X})$. Thus, as far as objects are concerned, we can do dynamical systems theory on measured topological systems $\mathfrak{X} = (X, \tau, \mu, T)$ where (X, τ) is a compact zero-dimensional

Polish space, μ a probability measure on $\mathcal{B}(\tau)$ (the induced Borel σ -algebra), and $T : X \rightarrow X$ is a continuous function. They naturally form the category \mathbf{TS}_{0c} where the morphisms are continuous, measure-preserving functions that commute with the dynamics.

The mentioned issue with the morphisms is that any injective (modulo null sets) morphism in \mathbf{TS}_{0c} will be an isomorphism in \mathbf{DS} but need not be one in \mathbf{TS}_{0c} . So we cannot expect the two categories to be equivalent, but we may expect them to be equivalent *after* adding inverses to those injective morphisms in \mathbf{TS}_{0c} . We show that this is indeed true: Formally, adding inverses is precisely the task of localizations in category theory.

There is a subtlety in constructing such a localization, and this explains the third category \mathbf{bTS}_0 in the bottom layer. If we have two topological (realizations of dynamical) systems \mathfrak{X} and \mathfrak{Y} and a measure-theoretic morphism $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$, we also need to ‘topologically realize’ this morphism as a continuous one. We can use the (Borel-) measurability of φ to refine the topology τ of \mathfrak{X} into a zero-dimensional Polish topology τ' such that φ becomes continuous:

$$\begin{array}{ccc} \mathfrak{X} \text{ with } \tau' & & \\ \text{id}_X \text{ cont} \downarrow & \searrow^{\varphi \text{ cont}} & \\ \mathfrak{X} \text{ with } \tau & \xrightarrow{\varphi \text{ meas}} & \mathfrak{Y} \end{array}$$

so we’ve realized the measurable φ as a span of continuous morphisms. (This observation will be at the heart of the formal localization result below.) However, the finer τ' will, in general, not be compact (no strictly finer topology of a compact Hausdorff space is compact). So we need to add the intermediate category of zero-dimensional (not necessarily compact) measured topological systems \mathbf{TS}_0 , and localize that to obtain \mathbf{DS} .

Compact zero-dimensional \mathfrak{X} have, with their sets of clopens $\mathbf{Clp}X$, a canonical countable basis that serves as the set of possible observations with which the observation domain $\mathbf{D}(\mathfrak{X})$ is built. However, this is no longer true when compactness is dropped. So we make the choice of basis explicit and work with pairs $(\mathfrak{X}, \mathcal{B})$. They form the objects of the category \mathbf{bTS}_0 and the morphisms are the usual ones between measured topological systems except that they now also are required to respect bases: the preimage of a basis element is a basis element.

Fortunately, after this addition we can still form the desired localization: First, the class W of morphisms in \mathbf{bTS}_0 that are injective (technically, on a clopen invariant set of full measure) admits what is called a right calculus of fractions. This means that there is a concrete way of adding inverses (in the form of spans) to those morphisms in W and obtain the category $\mathbf{bTS}_0[W^{-1}]$ with the same objects but with those additional added inverses as morphisms. Second, the category of dynamical systems \mathbf{DS} is indeed equivalent to that localization $\mathbf{bTS}_0[W^{-1}]$. Remarkably, the equivalence between spans that is used in the construction of

the localization then coincides with the equivalence modulo null sets between measure-theoretic morphisms of dynamical systems.

This explains the localization Loc in the diagram. The other two functors I_B and \bar{C} show that we basically obtained the original naive aim of localizing the category of compact systems TS_{0c} directly: With \bar{C} we can compactify any system in bTS_0 and obtain one in TS_{0c} .¹ And with I_B we can naturally regard a compact system \mathfrak{X} in TS_{0c} as one in bTS_0 with its canonical basis $\text{Clp}(X)$.

Finally, we may wonder: if the observation domain construction works more generally for any (standard) dynamical system \mathfrak{X} with a countable and separating measurable basis \mathcal{B} , why don't we define the domain functor on DS ? (Or, in analogy to bTS_0 , on the category 'bDS' whose objects are the $(\mathfrak{X}, \mathcal{B})$'s and morphisms additionally respect bases.) The reason is that in bTS_0 it is cleaner to show that the construction is functorial on morphisms: since the morphisms then are total, there is no fuzz in choosing appropriate representatives, etc. And with the localization in place, we can restrict to bTS_0 , in a sense, without loss of generality.

The top layer. By now, we've established a solid foundation on which we can build the top layer of dynamical domains. As constructed in chapter 4, dDOM is the category of dynamical domains: Its objects are structures in the sense of domain theory (domains with additional domain-theoretic structure), and its morphisms are those in the sense of domain theory (Scott-continuous functions that respect the additional domain-theoretic structure).

First, we establish that the constructions S (the system modeled by a dynamical domain) and D (the observation domain of a system \mathfrak{X} relative to a choice of basis \mathcal{B}) extend to functors $S : \text{dDOM} \rightarrow \text{TS}_{0c}$ and $D : \text{bTS}_0 \rightarrow \text{dDOM}$. We use the I_B functor to extend (resp., restrict) those functors so that they are of the corresponding type: $\hat{S} := I_B \circ S : \text{dDOM} \rightarrow \text{bTS}_0$ (resp., $\hat{D} := D \circ I_B : \text{TS}_{0c} \rightarrow \text{dDOM}$).

Next, we show the result anticipated at the very beginning: that the domain functor \hat{D} is right adjoint to the system functor S —in fact, it is a reflective adjunction in the sense that the counit is a natural isomorphism.

This establishes the 'lower' triangle $\text{dDOM} \text{---} \text{bTS}_0 \text{---} \text{TS}_{0c}$. So it remains to explain the 'upper' triangle: Given a (reflective) adjunction, it is natural to analyze what it takes to restrict it to an equivalence (i.e., that also the unit is a natural isomorphism). We'll make the surprising observation that this amounts to having a simple and purely domain-theoretic property that we'll call max-reflective. (The maximal elements of the domain 'reflect', in a sense, the non-maximal elements, see definition 5.2.19.) The full subcategory dDOM_r of dynamical domains whose underlying domain has this property will then turn out to be the right restriction of dDOM such that $S : \text{dDOM}_r \rightleftarrows \text{TS}_{0c} : \hat{D}$ forms an (adjoint) equivalence. We'll

¹We'll further analyze \bar{C} into $C \circ \bar{\cdot}$ where $\bar{\cdot}$ is a right-adjoint functor which closes bases under logical and dynamical operations and C is a left-adjoint functor which compactifies the state space in a way familiar from Stone duality.

also show that the functor $\hat{D} \circ S$ is optimal in turning a dynamical domain \mathfrak{D} into a max-reflective one, i.e., it is left-adjoint to the inclusion functor.

The standard case. So far we've discussed dynamical systems in our general sense, but we've also mentioned that the usual setting of ergodic theory is that of standard dynamical systems. They form the category \mathbf{sDS} (with the same notion of morphism). In a similar spirit, we define a measured topological system to be *standard* if its dynamics is a homeomorphism that is measure-preserving. Thus, \mathbf{TS}_{0cs} is the full subcategory of \mathbf{TS}_{0c} consisting of standard systems. Similarly, \mathbf{bTS}_{0s} is the full subcategory of \mathbf{bTS}_0 consisting of those $(\mathfrak{X}, \mathcal{B})$ where \mathfrak{X} is standard and, additionally, \mathcal{B} is closed under images of the dynamics. The reason for the latter condition is the following. In chapter 4, we've developed a domain-theoretic condition for being a *standard* dynamical domain and showed that it corresponds to the system-theoretic notion. (In the sense that, for a standard dynamical domain \mathfrak{D} , the system $S(\mathfrak{D})$ is standard, and if \mathfrak{X} is a standard system and \mathcal{B} closed under images of the dynamics, then $D(\mathfrak{X}, \mathcal{B})$ is standard.) Thus, we write \mathbf{dDOM}_s (resp., \mathbf{dDOM}_{rs}) for the full subcategory of \mathbf{dDOM} (resp., \mathbf{dDOM}_r) of standard dynamical domains. We'll show that the main diagram remains valid in this standard setting, i.e., after adding an \mathbf{s} to every category (see figure 5.6 for an explicit depiction of the resulting diagram). Much of this is immediate (from the functors restricting appropriately), but for the localization this requires more work.

Interpretation of the main diagram As mentioned at the very beginning, the main motivation for the main diagram is to *complete* the tool of dynamical domains: to have a full translation between systems and domains. We'll add three further remarks on interpretation.

First, the localization (together with the compactification) may also be regarded as determining in categorical language what the common assumption of 'working modulo null sets' in dynamical systems really amounts to. We can think of \mathbf{bTS}_0 (resp., \mathbf{TS}_{0c}) as the 'honest toil' category of dynamical systems: (1) To show that a set is a null set we need to first show that it is a set that we can 'grasp' (i.e., is Borel) and then that its measure is 0; it is not enough to just note that it is an (arbitrarily complicated) subset of such a set. (2) To show that a function is a morphism we need to check its properties everywhere (e.g., equivariance, i.e., commutativity with the dynamics); we cannot just move to a partial function where we discard a troublesome null set.²

Second, in the previous chapter, we've motivated that dynamical domains provide, in a sense, *symbolic* computational models (qua domain-theoretic structures) to the *non-symbolic* computation realized by dynamical systems ('non-symbolic'

²One might be tempted to ask: If \mathbf{TS}_{0c} is the 'honest toil' category of dynamical systems, what, then, is \mathbf{DS} ? Well, that depends on one's views: presumably either mathematical paradise (inspired by the words of Hilbert (1926, p. 170)) or the advantage of theft (inspired by the words of Russell (1919, p. 71)).

in the sense that the state space is continuous rather than discrete). As already discussed in chapter 3, the reflective adjunction $S : \mathbf{dDOM} \rightleftarrows \mathbf{TS}_{0c} : \hat{D}$ thus has a computational interpretation (Sassone, Nielsen, and Winskel 1996; Winskel and Nielsen 1995): It means that the ‘non-symbolic’ model of computation \mathbf{TS}_{0c} is more abstract than (i.e., can be embedded in) the ‘symbolic’ model \mathbf{dDOM} . What is abstracted away is precisely what violates being max-reflective, since the symbolic model \mathbf{dDOM}_r is equivalent to the non-symbolic model \mathbf{TS}_{0c} . We discuss this further in chapter 7.

Third, an important consequence of the fact that the main diagram commutes is that the two functors $S \circ D, \bar{C} : \mathbf{bTS}_0 \rightarrow \mathbf{TS}_{0c}$ are naturally isomorphic. Both compactify a system \mathfrak{X} relative to a basis \mathcal{B} , but they do so in very different ways. The first builds the computational model for \mathfrak{X} with respect to the possible observations given by \mathcal{B} and then takes the system modeled by this model—so we may call this the *computational compactification*. The second builds the space of logical models (the Stone space) that are possible according to the set of properties \mathcal{B} that can be observed—so we may call this the *logical compactification*.³ Two important consequences are: (1) we obtain two very different but (measure-theoretically) isomorphic topological realizations of the original dynamical system, and (2) this alignment of domain theory and logic may hint at a ‘dynamical domain theory in logical form’ à la Abramsky (1991).

Related work Since we’ve discussed work relating to the dynamical domain construction in the previous chapter, we only add here work relating to the categorical constructions of the present chapter.

There is, of course, work on a general category-theoretic approach to dynamical systems. Inspiring earlier references are, for example, Lawvere (1986) and Niefield (1996). More recently, there is, for example, Behrisch et al. (2017) and Schultz, Spivak, and Vasilakopoulou (2020). Here we are in a more concrete setting by explicitly working with the structures of ergodic theory.

Concerning the categories, the categories \mathbf{DS} and \mathbf{sDS} are standard, except for the more general notion of dynamical system in the former. Though, the standard references usually don’t explicitly define them as categories, so we check in some detail that they indeed form a category and that the usual system-theoretic notion of isomorphism coincides with the category-theoretic one. The categories \mathbf{TS}_{0c} and \mathbf{TS}_{0cs} are standard, too: the usual setting of topological dynamics is a compact metrizable space with a continuous dynamics that often is assumed to be invertible, and a measure is added when considering topological realizations. Concerning the category \mathbf{bTS}_0 , as already discussed in chapter 4, Polish spaces with a distinguished basis play an important role for Danos and Garnier (2015) and Dahlqvist, Danos, and Garnier (2016). In the latter paper (in def. 3.1 on page

³Before performing this logical compactification, though, \mathcal{B} is closed under logical and dynamical operations.

87), they also make the choice of basis explicit and form the category of based Polish spaces where the morphisms need to be, as they call it, ‘base-preserving’. The main difference here is that we also deal with dynamics.

Concerning the localization, viewing (measure-theoretic) dynamical systems as a localization of topological dynamical systems is, as we hope to succeed in motivating, a rather natural construction, but so far we haven’t yet found a reference. A reason may be that localizations are usually used in a different field: homological algebra and derived categories (Yekutieli 2020).

Concerning the compactification, we use the standard Wallman compactification theory and extend it to include dynamics. This is done in sections 5.3.2–5.3.3 where we’ll provide references and discuss commonalities with (and differences to) a construction of Danos and Garnier (2015).

Concerning the system and domain functors—i.e., the heart of this chapter—, they are, to the best of our knowledge, new.

Overview of the chapter In section 5.2, we define all the categories that we’ll deal with. In section 5.3, we establish the bottom layer of the main diagram. In section 5.4, we construct the system and domain functors, and show, in section 5.5, that they are adjoint. In section 5.6, we analyze the adjunction into forming an adjoint equivalence with the reflective subcategory of max-reflective dynamical domains. In section 5.7, we conclude that we have established the main diagram (and that it restricts to the standard case).

As suggested selective reading, we recommend sections 5.4 (minus subsection 5.4.3) and 5.5 for the main result: the diagonal of the main diagram. With a bit more time, this can be continued with section 5.6. As preparation one reads 5.2 but skips the categories one is not interested in (e.g., \mathbf{DS} and \mathbf{dDOM}_r). Section 5.3 on the bottom layer can be read largely independently. It is recommended to those who are not satisfied with taking the category \mathbf{TS}_{0c} of compact zero-dimensional measured topological systems as a good enough category of dynamical systems but rather want to see its connection to the ‘real’ category of dynamical systems \mathbf{DS} .

5.2 The categories

After introducing some background (section 5.2.1), we define the categories that occur in the main diagram: dynamical systems (section 5.2.2), measured topological systems (section 5.2.3), dynamical domains (section 5.2.4), based measured topological systems (section 5.2.6), and max-reflective dynamical domains (section 5.2.7). In section 5.2.5, we recap the main results from chapter 4.

5.2.1 Background

Domain theory We refer to section 4.2.1 of the previous chapter for a self-contained introduction to the domain theory that we need. Here we'll just recap the main concepts. Unless noted otherwise, these are found in the standard reference by Abramsky and Jung (1994).

A *Scott domain* (D, \leq) is a directed-complete partial order (dcpo) that is ω -algebraic and bounded-complete. For $A \subseteq D$, we write $\max A := \{a \in A : \forall b \in A. b \geq a \Rightarrow b = a\}$ and $\uparrow A := \{b \in D : \exists a \in A. b \geq a\}$ and $\downarrow A := \{b \in D : \exists a \in A. b \leq a\}$ (and $\uparrow a := \uparrow\{a\}$ and $\downarrow a := \downarrow\{a\}$). The Scott topology on D is denoted $\Sigma = \Sigma(D)$ and the Lawson topology is denoted $\Lambda = \Lambda(D)$.

A function $f : D \rightarrow E$ between dcpos is Scott-continuous iff it is monotone and preserves the supremum of directed subsets. A function $f : Q \rightarrow P$ between posets is a *projection* if it is surjective, monotone, and preimages of principal upsets (i.e., sets of the form $\uparrow x$) are principal upsets. Equivalently, there is a monotone function $e : P \rightarrow Q$ (the *embedding* determined by f) such that $f \circ e = \text{id}_P$ and $e \circ f \leq \text{id}_Q$ (i.e., $e \circ f(y) \leq y$ for all $y \in Q$).

For a dcpo D , a function $v : \Sigma(D) \rightarrow [0, \infty]$ is a *continuous valuation* (see e.g. Edalat 1995a; Jones and Plotkin 1989; Lawson 1982) if, for all $U, V \in \Sigma(D)$, (a) $v(\emptyset) = 0$, (b) if $U \subseteq V$, then $v(U) \leq v(V)$, (c) $v(U \cup V) + v(U \cap V) = v(U) + v(V)$, and (d) if $(U_j)_{j \in J}$ is a directed family in $\Sigma(D)$, then $v(\bigcup_j U_j) = \sup_j v(U_j)$. And v is *normalized* if $v(D) = 1$.

Dynamical systems theory Before defining dynamical systems in the next subsection, we recap here their underlying spaces. (Again we just state the definitions and refer to section 4.2.2 of the previous chapter for some more details.) There are three kinds:

(1) As usual, a *probability space* is a triple (X, \mathcal{A}, μ) where X is a set, \mathcal{A} is a σ -algebra, and $\mu : \mathcal{A} \rightarrow [0, 1]$ is measure with $\mu(X) = 1$. A probability space (X, \mathcal{A}, μ) is *complete* if, for all $A \subseteq B \in \mathcal{A}$, if $\mu(B) = 0$, then $A \in \mathcal{A}$. The *completion* of (X, \mathcal{A}, μ) is denoted $(X, \mathcal{A}_\mu, \mu)$.⁴

(2) A probability space (X, \mathcal{A}, μ) is a *Borel probability space* if there is a Polish (i.e., separable and completely metrizable) topology τ on X with $\mathcal{A} = \mathcal{B}(\tau)$, where $\mathcal{B}(\tau)$ denotes the Borel σ -algebra of the topology τ (see e.g. Kechris 1995, def. 12.5). In that case, (X, \mathcal{A}) is called a standard Borel space.

(3) A probability space (X, \mathcal{A}, μ) is a *Lebesgue space* (or standard probability space) if it is complete and there is a second-countable topology τ on X with $\tau \subseteq \mathcal{A}$, $\mathcal{B}(\tau)_\mu = \mathcal{A}$, and μ inner regular (de la Rue 1993, def. 1-1). Equivalently (de la Rue 1993, thm. 4-3), a Lebesgue space is a complete probability space that is isomorphic mod 0 to the ordinary Lebesgue space of a (possibly empty) interval $[0, a] \subseteq \mathbb{R}$ together with countably many point masses. (See Walters (1982, def. 2.3)

⁴Here \mathcal{A}_μ is the σ -algebra of sets of the form $A \cup N$ for $A \in \mathcal{A}$ and $N \subseteq M$ for some $M \in \mathcal{A}$ with $\mu(M) = 0$, and $\mu(A \cup N)$ is taken to mean $\mu(A)$.

or Petersen (1983, def. 4.5).) Thus, any completion of a standard Borel probability space is a Lebesgue space, and any Lebesgue space is isomorphic mod 0 to the completion of a standard Borel probability space.

When restricting to the subclasses (2) and (3) of probability spaces, different natural notions of isomorphism of probability spaces coincide (Walters 1982, ch. 2).

Category theory We only use basic category theory: the definitions of categories, functors, limits, adjunctions, and equivalences; as found in standard references like Leinster (2014) or the classic Mac Lane (1998). Concerning notation, we use suffixed subscripts to denote restrictions (of the objects) to full subcategories, and suffixed superscripts denote restrictions (of the morphisms) to wide subcategories.

5.2.2 Categories of dynamical systems

On the object level, dynamical systems are defined as follows.

5.2.1. DEFINITION. An *abstract dynamical system* is a structure $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a probability space and $T : X \rightarrow X$ is measurable (i.e., for $A \in \mathcal{A}$, $T^{-1}(A) \in \mathcal{A}$). A *(general) dynamical system* is an abstract dynamical system $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a Borel probability space. A *standard dynamical system* is an abstract dynamical system $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a Lebesgue space and $T : X \rightarrow X$ is bijective and measure-preserving (i.e., measurable and for $A \in \mathcal{A}$, $\mu(T^{-1}(A)) = \mu(A)$). We often omit the term ‘general’.

On the morphism level, we can use the usual notion of homomorphism of systems. (See e.g. Walters (1982, p. 61) or Petersen (1983, p. 11).) Recall that, for two sets X and Y , a partial function $\varphi : X \rightarrow Y$ with domain $M \subseteq X$ and codomain $N \subseteq Y$ is (formally) a relation $\varphi \subseteq X \times Y$ such that for every $x \in M$ there is exactly one $y \in N$ such that $x\varphi y$.⁵

5.2.2. DEFINITION. A *(system) homomorphism* or *factor* $\varphi : (X, \mathcal{A}, \mu, T) \rightarrow (Y, \mathcal{B}, \nu, S)$ between abstract dynamical systems is a partial function $\varphi : X \rightarrow Y$ with domain $M \subseteq X$ and codomain $N \subseteq Y$ such that

1. Domain: $M \in \mathcal{A}$, $T(M) \subseteq M$, and $\mu(M) = 1$.
2. Codomain: $N \in \mathcal{B}$, $S(N) \subseteq N$, and $\nu(N) = 1$.
3. Measurable: For all $B \in \mathcal{B}$, $\varphi^{-1}(B) \in \mathcal{A}$.⁶

⁵So $\varphi \upharpoonright M \times N$ is a function $M \rightarrow N$. Note that the function is not required to be surjective: the image $\varphi(M)$ is only included in the codomain N and need not be identical to it.

⁶So $\varphi : M \rightarrow N$ is measurable with respect to the sub- σ -algebras $\mathcal{A} \upharpoonright M$ and $\mathcal{B} \upharpoonright N$.

4. Measure-preserving: For all $B \in \mathcal{B}$, $\mu(\varphi^{-1}(B)) = \nu(B)$.
5. Equivariant: For all $x \in M$, $\varphi(T(x)) = S(\varphi(x))$.

We identify two partial maps $\varphi, \psi : X \rightarrow Y$ iff they are identical on an invariant set of full measure, i.e., there is $A \in \mathcal{A}$ such that $T(A) \subseteq A$, $\mu(A) = 1$ and, for all $x \in A$, both $\varphi(x)$ and $\psi(x)$ are defined and equal.⁷

One may wonder why bother with the complication of partial functions and not just define morphisms to be total functions that preserve the relevant structure everywhere rather than just almost everywhere. This is because, with the above choice of morphisms, the resulting category-theoretic notion of isomorphism is precisely the one that is the usual notion of isomorphism in dynamical systems theory, as we'll show below.

But first we check that this choice of objects and morphisms does indeed yield a category. Since the standard references don't do this explicitly, we provide a full proof.

5.2.3. PROPOSITION. *The following forms the category \mathbf{aDS} : the objects are abstract dynamical systems and the morphisms are homomorphisms between them. The identity morphism is the identity function (modulo equivalence) and composition of morphisms is function composition (modulo equivalence). We write \mathbf{DS} and \mathbf{sDS} for the full subcategories of general and standard dynamical systems, respectively.*

Note that \mathbf{sDS} is not a (full) subcategory of \mathbf{DS} , but intuitively it is 'modulo' completion. (So we don't write \mathbf{DS}_s .) Since we're interested in dynamical systems over 'well-behaved' spaces, we will be interested in \mathbf{DS} and \mathbf{sDS} . Thus, for us \mathbf{aDS} is only a convenient ambient category containing both categories of study.

PROOF. We first show that composition is well-defined: Let $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$, $\mathfrak{Y} = (Y, \mathcal{B}, \nu, S)$, $\mathfrak{Z} = (Z, \mathcal{C}, \kappa, R)$, and let $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ and $\psi : \mathfrak{Y} \rightarrow \mathfrak{Z}$ be homomorphisms in \mathbf{aDS} where φ has domain $M \subseteq X$ and codomain $N \subseteq Y$ and ψ has domain $K \subseteq Y$ and codomain $L \subseteq Z$.

Let $\chi := \psi \circ \varphi$ be the composition of partial functions: $\chi := \{(x, z) \in X \times Z : \exists y \in Y : (x, y) \in \varphi, (y, z) \in \psi\}$. Then χ is a partial function with domain $M' := M \cap \varphi^{-1}(K) = \{x \in X : x \in M, \varphi(x) \in K\}$ and codomain L . In particular, $\psi \circ \varphi(x)$ is defined (and equals z) iff $\varphi(x)$ is defined and $\psi(\varphi(x))$ is defined (and equals z). We need to show that $\chi : \mathfrak{X} \rightarrow \mathfrak{Z}$ is a homomorphism, i.e., check conditions (1)–(5).

Concerning (1), we have $M \in \mathcal{A}$ and, since φ is measurable and $K \in \mathcal{B}$, $\varphi^{-1}(K) \in \mathcal{A}$, also $M' = M \cap \varphi^{-1}(K) \in \mathcal{A}$. To show $T(M') \subseteq M'$, let $x \in M'$ and show $Tx \in M'$: Since $x \in M$ and $T(M) \subseteq M$, we have $Tx \in M$, so we need to

⁷If T is measure-preserving, we can drop the condition $T(A) \subseteq A$ (and thus recover the usual 'identical mod 0' equivalence relation of functions) since we can take $A' := \bigcap_{n=0}^{\infty} T^{-n}(A)$.

show $Tx \in \varphi^{-1}(K)$. Indeed, we have $\varphi Tx = S\varphi x$ and $\varphi x \in K$. Since $S(K) \subseteq K$, $S\varphi x \in K$. Hence $\varphi Tx \in K$, as needed. Finally, to show $\mu(M') = 1$, note that $\mu(M) = 1$ and $\mu(\varphi^{-1}K) = \nu(K) = 1$ and the intersection of sets of full measure is of full measure.

Concerning (2), we have, by assumption, $L \in \mathcal{C}$, $R(L) \subseteq L$, $\kappa(L) = 1$, as needed.

Concerning (3), if $C \in \mathcal{C}$, then $\chi^{-1}(C) = \varphi^{-1}(\psi^{-1}(C))$.⁸ Since ψ is measurable, $\psi^{-1}(C) \in \mathcal{B}$, so, since φ is measurable, $\chi^{-1}(C) = \varphi^{-1}(\psi^{-1}(C)) \in \mathcal{A}$.

Concerning (4), if $C \in \mathcal{C}$, then,

$$\mu(\chi^{-1}C) = \mu(\varphi^{-1}(\psi^{-1}C)) = \nu(\psi^{-1}C) = \kappa(C).$$

Concerning (5), For $x \in M'$, we have $Tx \in M'$, so $\chi(Tx) = \psi(\varphi(Tx))$ is defined, and $\varphi(x) \in K$, whence $S\varphi(x) \in K$, so

$$\chi(T(x)) = \psi(\varphi(Tx)) = \psi(S(\varphi(x))) = R(\psi(\varphi(x))) = R(\chi(x)).$$

Finally, note that composition of partial functions is associative,⁹ and the (total) identity function $\text{id}_X : X \rightarrow X$ is a homomorphism $\text{id}_{\mathfrak{X}} : \mathfrak{X} \rightarrow \mathfrak{X}$. And this can be lifted to the equivalence classes: If φ is equivalent to φ' (identical on the invariant set $A \subseteq X$ of full measure) and ψ is equivalent to ψ' (identical on the invariant set $B \subseteq Y$ of full measure), then also $\chi = \psi \circ \varphi$ is equivalent to $\chi' = \psi' \circ \varphi'$ (identical on the invariant set $A \cap \varphi^{-1}(B)$ of full measure).¹⁰ \square

Next, as promised, we check that the category-theoretic notion of isomorphism coincides with the usual notion of isomorphism in dynamical systems theory. For this we first prove a lemma which also will be useful later on.

5.2.4. LEMMA. *1. Let $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ be in DS. Assume there is an invariant set A of full measure on which φ is defined and injective. Then φ is an isomorphism in DS (i.e., has an inverse).*

2. The same holds in sDS.

⁸We have, for $x \in X$, that: $x \in \chi^{-1}(C)$ iff $\chi(x)$ is defined and $\chi(x) \in C$ iff $x \in M'$ and $\chi(x) \in C$ iff $\varphi(x)$ is defined and $\psi(\varphi(x))$ is defined and in C iff $x \in \varphi^{-1}(\psi^{-1}(C))$.

⁹We have: $\chi \circ (\psi \circ \varphi)(x)$ is defined iff $\psi \circ \varphi(x)$ is defined and $\chi(\psi \circ \varphi(x))$ is defined iff $\varphi(x)$ is defined and $\psi(\varphi(x))$ is defined and $\chi(\psi(\varphi(x)))$ is defined iff $\varphi(x)$ is defined and $(\chi \circ \psi)(\varphi(x))$ is defined iff $(\chi \circ \psi) \circ \varphi(x)$ is defined. And if defined, the values are identical.

¹⁰Full measure: Since φ is measurable and measure-preserving, $A \cap \varphi^{-1}B$ is an intersection of measurable sets with full measure and hence of full measure as well. Invariant: Assume $x \in A \cap \varphi^{-1}B$. Then $Tx \in A$ since A is T -invariant, so we need to show $\varphi Tx \in B$. Since $x \in A$ and $\varphi x \in B$, and since φ is defined on A and B is S -invariant, $\varphi Tx = S\varphi x \in B$. Identical: If $x, x' \in A \cap \varphi^{-1}B$, then $\varphi(x) = \varphi(x')$ and ψ is defined for them, so $\chi(x) = \psi(\varphi(x)) = \psi(\varphi(x')) = \chi(x')$.

PROOF. We deal with DS in the main text and with sDS in square brackets. Write $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ and $\mathfrak{Y} = (Y, \mathcal{B}, \nu, S)$. By assumption, $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is a homomorphism with domain $M \supseteq A$ and codomain N . Define $B := \varphi(A) \subseteq Y$ and $\psi : Y \rightarrow X$ as the partial function with domain $B \subseteq Y$ and codomain $A \subseteq X$ which assigns every $y \in B$ the unique $x \in A$ such that $\varphi(x) = y$. We claim that ψ is a homomorphism—so we need to check conditions (1)–(5).

First an observation: Since (X, \mathcal{A}, μ) is a Borel probability space [resp., Lebesgue space] and $A \in \mathcal{A}$ has full measure, also $(A, \mathcal{A} \upharpoonright A, \mu)$ is a Borel probability space [resp., Lebesgue space].¹¹ Now, $\varphi : A \rightarrow Y$ is an injective Borel-measurable [resp., measurable and measure-preserving] function between standard Borel spaces [resp., Lebesgue spaces], so, for all $C \in \mathcal{A} \upharpoonright A$, we have $\varphi(C) \in \mathcal{B}$ by Kechris (1995, p. 15.2) [resp., by de la Rue (1993, thm. 3-5)].

Concerning (1), By the observation, $B \in \mathcal{B}$. And B is S -invariant: If $y \in B$, then $y = \varphi(x)$ for some $x \in A$, so $Tx \in A$ and, since $x \in M$, we have $S(y) = S(\varphi(x)) = \varphi(Tx) \in \varphi(A) = B$. Finally, B is of full measure: $\nu(B) = \nu(\varphi(A)) = \mu(\varphi^{-1}\varphi(A)) \geq \mu(A) = 1$, so $\nu(B) = 1$.

Concerning (2), by assumption, $A \in \mathcal{A}$, $T(A) \subseteq A$, $\mu(A) = 1$.

Concerning (3), for $C \in \mathcal{A}$, we have, $\psi^{-1}(C) = \varphi(C \cap A)$.¹² And, by the observation, this is in \mathcal{B} since $C \cap A \in \mathcal{A} \upharpoonright A$.

Concerning (4), for $C \in \mathcal{A}$, we have, because φ is measure-preserving, that $\nu(\psi^{-1}(C)) = \mu(\varphi^{-1}\psi^{-1}(C))$. By the previous step, this equals $\mu(\varphi^{-1}\varphi(C \cap A))$. By definedness and injectivity on A , this further equals $\mu(C \cap A)$.¹³ Since A is of full measure, this equals $\mu(C)$, as needed. ◦ Concerning (5), for $y \in B$, let $x := \psi(y) \in A$. So, by definition, $\varphi(x) = y$. Since φ is equivariant (on $M \supseteq A$), $\varphi(T(x)) = S(\varphi(x)) = S(y)$. So $T(x)$ is the element of A whose φ -image is $S(y)$, whence $\psi(S(y)) = T(x) = T(\psi(y))$, as needed.

To finish the proof, we see that ψ is an inverse to φ : First, $\psi \circ \varphi = \text{id}_{\mathfrak{X}}$ on the T -invariant set $A \in \mathcal{A}$ of full measure. Second, $\varphi \circ \psi = \text{id}_{\mathfrak{Y}}$ on the set B which, by (1) above, is an S -invariant set of full measure. ◻

5.2.5. PROPOSITION. *1. \mathfrak{X} and \mathfrak{Y} are isomorphic in DS iff there is a partial function $\varphi : X \rightarrow Y$ with domain $M \subseteq X$ and codomain $N \subseteq Y$ such that*

¹¹Proof: For the Borel case, this follows since, if (X, \mathcal{A}) is a standard Borel space and $A \in \mathcal{A}$, also $(A, \mathcal{A} \upharpoonright A)$ is standard (Kechris 1995, p. 13.4), and μ still is a probability measure. For the Lebesgue case, this follows since, if (X, \mathcal{A}, μ) is a complete probability space that is isomorphic mod 0 to an interval with Lebesgue measure together with countably many point masses, then, after discarding a measure null set, this is still true for $(A, \mathcal{A} \upharpoonright A, \mu)$.

¹²For $y \in Y$, we have $y \in \psi^{-1}(C)$ iff $y \in B$ and $\psi(y) \in C$ iff $y \in B$ and the unique $x \in A$ with $\varphi(x) = y$ is in C iff $y \in \varphi(C \cap A)$. (For the reverse direction of the last step: if $y \in \varphi(C \cap A)$, then $y = \varphi(x)$ for $x \in C \cap A$, so $y \in \varphi(A) = B$ and $x \in A$ is such that $\varphi(x) = y$, whence x is unique the unique element of A with $\varphi(x) = y$ and it is in C .)

¹³If $x \in C \cap A$, then $\varphi(x)$ is defined and in $\varphi(C \cap A)$, so $x \in \varphi^{-1}\varphi(C \cap A)$. Conversely, if $x \in \varphi^{-1}\varphi(C \cap A)$, then $\varphi(x)$ is defined and in $\varphi(C \cap A)$. So there is $x' \in C \cap A$ with $\varphi(x) = \varphi(x')$. By injectivity on A , $x = x'$. So $x \in C \cap A$.

- (a) $\varphi : M \rightarrow N$ is a bijective function and $M \in \mathcal{A}$, $T(M) \subseteq M$, $\mu(M) = 1$, and $N \in \mathcal{B}$, $S(N) \subseteq N$, $\nu(N) = 1$.
- (b) Measurable: For all $B \in \mathcal{B}$, $\varphi^{-1}(B) \in \mathcal{A}$.
- (c) Measure-preserving: For all $B \in \mathcal{B}$, $\mu(\varphi^{-1}(B)) = \nu(B)$.
- (d) Equivariant: For all $x \in M$, $\varphi(T(x)) = S(\varphi(x))$.

2. The analogous claim holds in sDS.

PROOF. We deal with DS and sDS simultaneously. (\Leftarrow) By the assumption, $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is a homomorphism with domain M and codomain N which is injective on the invariant set M of full measure. By lemma 5.2.4, φ is an isomorphism, so \mathfrak{X} and \mathfrak{Y} are isomorphic. Note that this reasoning applies both to DS and sDS.

(\Rightarrow) Let $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ and $\psi : \mathfrak{Y} \rightarrow \mathfrak{X}$ be homomorphisms such that $\psi \circ \varphi = \text{id}_{\mathfrak{X}}$ on a T -invariant set $A \in \mathcal{A}$ of full measure and $\varphi \circ \psi = \text{id}_{\mathfrak{Y}}$ on a S -invariant set $B \in \mathcal{B}$ of full measure. Note that $\psi \circ \varphi$ is defined on A , so φ is defined on A and ψ is defined on $\varphi(A)$. Similarly for B .

Set $M := A \cap \varphi^{-1}(B)$ and $N := \varphi(M)$. We show that the partial function $\varphi' := \varphi \upharpoonright M : X \rightarrow Y$ with domain M and codomain N has properties (1a)–(1d).

Concerning (1a), by construction, $\varphi' : M \rightarrow N$ is surjective and it also is injective: for $x, x' \in M$, if $\varphi'(x) = \varphi'(x')$, then $\varphi'(x) = \varphi(x) \in B$ and $\varphi'(x') = \varphi(x') \in B$, so $\psi(\varphi(x)) = \psi(\varphi(x'))$ are defined and, since $x, x' \in A$,

$$x = \text{id}_{\mathfrak{X}}(x) = \psi(\varphi(x)) = \psi(\varphi(x')) = \text{id}_{\mathfrak{X}}(x') = x'.$$

Since φ is measurable and measure-preserving, $M = A \cap \varphi^{-1}(B)$ is of full measure, and it is T -invariant: If $x \in M$, then $T(x) \in A$ (since A is T -invariant) and $T(x) \in \varphi^{-1}(B)$ since $\varphi T(x) = S\varphi(x) \in B$ (since $\varphi(x) \in B$ and B is S -invariant).

We claim $\varphi(M) = B \cap \psi^{-1}(A)$. (Then it follows, similarly as above, that $N = B \cap \psi^{-1}(A)$ is of full measure and S -invariant.) If $y \in \varphi(M)$, then there is $x \in A \cap \varphi^{-1}(B)$ with $\varphi(x) = y$. So $y = \varphi(x) \in B$ and $\psi(\varphi(x)) = \text{id}_{\mathfrak{X}}(x) = x \in A$. Hence $y \in B \cap \psi^{-1}(A)$. Conversely, if $y \in B \cap \psi^{-1}(A)$, then $x := \psi(y) \in A$ is defined. And $\varphi(x) = \varphi(\psi(y)) = \text{id}_{\mathfrak{Y}}(y) = y \in B$. So $x \in M$ and $\varphi(x) = y$, whence $y \in \varphi(M)$.

Concerning (1b), if $C \in \mathcal{B}$, then $\varphi'^{-1}(C) = \varphi^{-1}(C) \cap M \in \mathcal{A}$.

Concerning (1c), if $C \in \mathcal{B}$, then, since M is of full measure, $\mu\varphi'^{-1}(C) = \mu(\varphi^{-1}(C) \cap M) = \mu\varphi^{-1}(C) = \nu(C)$.

Concerning (1d), if $x \in M$, $\varphi'(T(x)) = \varphi(T(x)) = S(\varphi(x)) = S(\varphi'(x))$.

Again, this reasoning applies both to DS and sDS. \square

5.2.3 Categories of measured topological systems

Usually, a topological system is defined as (X, T) where X is a compact metric space (hence Polish) and $T : X \rightarrow X$ continuous (and, often, bijective, whence

homeomorphic). Here, we'll be more general and don't assume compactness (and being homeomorphic only in the standard case). We also consider measures and indicate this by the term 'measured'.

5.2.6. DEFINITION. A (*general*) *measured topological system* is a structure $\mathfrak{X} = (X, \tau, \mu, T)$ where (X, τ) is a Polish space, μ a probability measure on $\mathcal{B}(\tau)$ (the Borel σ -algebra generated by τ), and $T : X \rightarrow X$ is continuous. It is *standard* if, additionally, T is a homeomorphism and measure-preserving. It is *zero-dimensional* (resp., *compact*) if X is.

The notion of morphism is straightforward.

5.2.7. DEFINITION. If $\mathfrak{X} = (X, \tau, \mu, T)$ and $\mathfrak{Y} = (Y, \sigma, \nu, S)$ are measured topological systems, a *morphism* $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ of measured topological systems is a function $\varphi : X \rightarrow Y$ that is continuous (if $V \in \sigma$, then $\varphi^{-1}(V) \in \tau$), measure-preserving (if $B \in \mathcal{B}(\sigma)$, then $\mu(\varphi^{-1}(B)) = \nu(B)$)¹⁴, and equivariant ($\varphi \circ T = S \circ \varphi$).

Now, the following is immediate.

5.2.8. PROPOSITION. *We have the category \mathbf{TS} consisting of measured topological systems with their morphisms. We form full subcategories with the following restrictions on objects:*

- 0 *restricting to zero-dimensional measured topological systems*
- c *restricting to compact measured topological systems*
- s *restricting to standard measured topological systems*

For example, \mathbf{TS}_0 (resp., \mathbf{TS}_{0c}) is the full subcategory of \mathbf{TS} consisting of (compact) and zero-dimensional measured topological systems. Moreover, $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is an isomorphism in \mathbf{TS} iff φ is a morphism in \mathbf{TS} that is a homeomorphism on the underlying spaces.

PROOF. Composition in \mathbf{TS} is simply function composition and the identity morphism is simply the identity function. So we need to show the 'moreover' part. Write $\mathfrak{X} = (X, \tau, \mu, T)$ and $\mathfrak{Y} = (Y, \sigma, \nu, S)$

Assume $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is an isomorphism in \mathbf{TS} . Let φ^{-1} be its inverse in \mathbf{TS} . Since φ in particular is a morphism, we need to show that it is a homeomorphism $X \rightarrow Y$. But this follows since it has a continuous inverse function φ^{-1} .

Assume φ is a morphism in \mathbf{TS} that is a homeomorphism on the underlying space. So φ has a continuous inverse function $\varphi^{-1} : Y \rightarrow X$. We need to show that $\psi := \varphi^{-1}$ is measure-preserving and equivariant. Indeed, for $A \in \mathcal{B}(\tau)$, we have, since φ is measure-preserving and $\psi^{-1}(A) \in \mathcal{B}(\sigma)$ (since $\psi : Y \rightarrow X$ is continuous

¹⁴Note that continuity implies $\varphi^{-1}(B) \in \mathcal{B}(\tau)$ so $\mu(\varphi^{-1}(B))$ is defined.

and hence Borel-measurable), that $\nu(\psi^{-1}(A)) = \mu(\varphi^{-1}\psi^{-1}(A)) = \mu(\varphi^{-1}\varphi(A)) = \mu(A)$. And for $y \in Y$, write $x := \varphi^{-1}(y)$, so $\varphi(T(x)) = S(\varphi(x)) = S(y)$, whence $T(x)$ is the φ -preimage of $S(y)$, so $\varphi^{-1}(S(y)) = T(x) = T(\varphi^{-1}(y))$, as needed. \square

Every topological system induces a dynamical system in a functorial way:

5.2.9. PROPOSITION. 1. We can define the functor $J : \mathbf{TS} \rightarrow \mathbf{DS}$ which sends $\mathfrak{X} = (X, \tau, \mu, T)$ to $J(\mathfrak{X}) := (X, \mathcal{B}(\tau), \mu, T)$ and $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ to $J(\varphi) := \varphi$.

2. We can define the functor $\bar{J} : \mathbf{TS}_s \rightarrow \mathbf{sDS}$ which sends $\mathfrak{X} = (X, \tau, \mu, T)$ to $\bar{J}(\mathfrak{X}) := (X, \mathcal{B}(\tau)_\mu, \mu, T)$ and $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ to $\bar{J}(\varphi) := \varphi$.

PROOF. Ad (1). Note that $J(\mathfrak{X}) = (X, \mathcal{B}(\tau), \mu, T)$ is a general dynamical system: Since τ is a Polish topology on X , $(X, \mathcal{B}(\tau))$ is a standard Borel space with probability measure μ and $T : X \rightarrow X$ is measurable since it is continuous. And since $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is continuous on the underlying state spaces, measure-preserving and equivariant, $J(\varphi) = \varphi$, regarded as ‘partial’ function with total domain and codomain, is (Borel-) measurable, measure-preserving, and equivariant, so a morphism in \mathbf{DS} . Since J is the identity on morphisms, it preserves composition and identity, whence is a functor.

Ad (2). Note that $\bar{J}(\mathfrak{X}) = (X, \mathcal{B}(\tau)_\mu, \mu, T)$ is a standard dynamical system: As above, $(X, \mathcal{B}(\tau), \mu)$ is a Borel probability space, so its completion $(X, \mathcal{B}(\tau)_\mu, \mu)$ is a Lebesgue space. Since T is continuous, it is measurable. And, by assumption, T is bijective and measure-preserving. Again, $J(\varphi) = \varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$, regarded as partial function with total domain and codomain, is (Borel-) measurable and measure-preserving on the Borel σ -algebra and hence also on its completion,¹⁵ and it is equivariant, so a morphism in \mathbf{sDS} . Since \bar{J} is the identity on morphisms, it preserves composition and identity, whence is a functor. \square

5.2.4 Categories of dynamical domains

We summarize the definition of the category \mathbf{dDOM} of dynamical domains (and the subcategory \mathbf{dDOM}_s of standard dynamical domains) that was motivated in chapter 4. In short, dynamical domains are certain limits of certain finite dynamical dcpos, where a dynamical dcpo is a domain-theoretic analogue of a dynamical system. We also add some new lemmas about dynamical domains, mostly about their categorical properties.

¹⁵If $B = A \cup N$ is in the completion of $\mathcal{B}(\sigma)$ with $A \in \mathcal{B}(\sigma)$ and $N \subseteq M \in \mathcal{B}(\sigma)$ and $\nu(M) = 0$, then $\varphi^{-1}(B) = \varphi^{-1}(A) \cup \varphi^{-1}(N)$ with $\varphi^{-1}(A) \in \mathcal{B}(\tau)$ and $\varphi^{-1}(N) \subseteq \varphi^{-1}(M) \in \mathcal{B}(\tau)$ with $\mu(\varphi^{-1}(M)) = \nu(M) = 0$. So $\varphi^{-1}(B)$ is in the completion of $\mathcal{B}(\tau)$ and we have $\mu(\varphi^{-1}(B)) = \mu(\varphi^{-1}(A)) = \nu(A) = \nu(B)$, as needed.

Dynamical dcpos Recalling from chapter 4, a function $f : D \rightarrow E$ between dcpos is *max-preserving* if $f(\max D) \subseteq \max E$ (i.e., if $a \in \max D$, then $f(a) \in \max E$). And a valuation v on D is *max-normalized* if v is normalized and $\max D$ can be written as a countable intersection of Scott-open sets with v -value 1.

5.2.10. DEFINITION. A *dynamical dcpo* is a triple $\mathfrak{D} = (D, v, f)$ where D is a dcpo, $v : \Sigma(D) \rightarrow [0, \infty]$ is a continuous valuation, and $f : D \rightarrow D$ is Scott-continuous. We call \mathfrak{D} :

- *finite* if D is finite
- *max-normalized* if v is max-normalized
- *max-preserving* if f is max-preserving
- a *dynamical Scott domain* if D is a Scott domain
- *max-surjective* if, for all $b \in \max D$, there is $a \in \max D$ such that $b \geq f(a)$.
- *valuation-preserving* if, for all $U \in \Sigma(D)$, $v(f^{-1}(U)) = v(U)$.

5.2.11. DEFINITION. A *dynamical morphism* $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ between two dynamical dcpos $\mathfrak{D} = (D, v, f)$ and $\mathfrak{E} = (E, w, g)$ is a function $\alpha : D \rightarrow E$ such that

1. Scott-continuous: α is Scott-continuous.
2. Max-preserving: α is max-preserving.
3. Max-bisimulative: For all $a \in D$ and $e \in \max E$, if $\alpha(a) \leq e$, then there is $d \in \max D$ such that $d \geq a$ and $\alpha(d) = e$.
4. Valuation-preserving: For all $V \in \Sigma(E)$, $w(V) = v(\alpha^{-1}(V))$.
5. Max-semi-equivariant: For all $a \in \max D$, $\alpha(f(a)) \geq g(\alpha(a))$.

Note that, if $g : D \rightarrow D$ is max-preserving, then α is *max-equivariant*: for all $a \in \max D$, $\alpha(f(a)) = g(\alpha(a))$.

5.2.12. DEFINITION. We define **dDCP** as the category of dynamical dcpos with dynamical morphisms. We define **dSCO** as the full subcategory of **dDCP** with dynamical Scott domains as objects. We can build further categories indicated by the following prefixes and suffixes:

- ^p restricting the morphisms to additionally be projections.
- _m restricting to max-preserving dynamical dcpos.
- _n restricting to max-normalized dynamical dcpos.

For example, $\mathbf{dSCO}_{\text{nm}}^{\text{P}}$ is the subcategory of \mathbf{dSCO} consisting of max-normalized and max-preserving dynamical Scott domains where the dynamical morphisms also are projections.

The following characterizes isomorphisms in a ‘one-sided’ way:

5.2.13. PROPOSITION. *Let $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ be a morphism in \mathbf{dDCP} with $\mathfrak{D} = (D, v, f)$ and $\mathfrak{E} = (E, w, g)$. Then α is an isomorphism in \mathbf{dDCP} iff $\alpha : D \rightarrow E$ is an order isomorphism (monotone, order-reflecting, surjective) and α is max-equivariant (for $a \in \max D$, $\alpha(f(a)) = g(\alpha(a))$).¹⁶ This also holds for \mathbf{dDCP}^{P} and all full subcategories of \mathbf{dDCP} and \mathbf{dDCP}^{P} .*

PROOF. (\Rightarrow) If α is an isomorphism, it has an inverse α^{-1} . So α is monotone (since it is Scott-continuous), surjective (since it has an inverse), and order-reflecting (since its inverse is monotone: if $\alpha(a) \leq \alpha(a')$, then $a = \alpha^{-1}\alpha(a) \leq \alpha^{-1}\alpha(a') = a'$). Moreover, α is max-equivariant: For $a \in \max D$, we have, since α is max-semi-equivariant, $\alpha(f(a)) \geq g(\alpha(a))$, so we need to show \leq . Since $b := \alpha(a) \in \max E$ and α^{-1} is max-semi-equivariant, $\alpha^{-1}(g(b)) \geq f(\alpha^{-1}(b)) = f(a)$. By applying α to both sides, we get $g(\alpha(a)) \geq \alpha(f(a))$, as needed.

(\Leftarrow). Since $\alpha : D \rightarrow E$ is an order isomorphism, let $\alpha^{-1} : E \rightarrow D$ be its inverse, which hence again is an order isomorphism. We have to show that it is a dynamical morphism, too.

Continuous: Qua order isomorphism, α^{-1} is Scott-continuous.

Max-preserving: Qua order isomorphism, α^{-1} maps maximal elements to maximal elements.

Max-bisimulative: If $b \in E$ and $\alpha^{-1}(b) \leq d \in \max D$, define $e := \alpha(d) \in \max E$ and we have $e \geq b$ (since $\alpha^{-1}(b) \leq d$ and α is monotone we get $b = \alpha\alpha^{-1}(b) \leq \alpha(d) = e$) and $\alpha^{-1}(e) = \alpha^{-1}\alpha(d) = d$.

Valuation-preserving: Let $U \in \Sigma(D)$ and show $v(U) = w((\alpha^{-1})^{-1}(U))$. We have, since α is valuation-preserving,

$$w((\alpha^{-1})^{-1}(U)) = w(\alpha(U)) = v(\alpha^{-1}(\alpha(U))) = v(U).$$

Max-semi-equivariant: For $b \in \max E$, write $a := \alpha^{-1}(b) \in \max D$, so we have, since α is max-equivariant,

$$\alpha^{-1}(g(b)) = \alpha^{-1}(g(\alpha(a))) = \alpha^{-1}(\alpha(f(a))) = f(a) = f(\alpha^{-1}(b)),$$

so, in particular, $\alpha^{-1}(g(b)) \geq f(\alpha^{-1}(b))$.

In \mathbf{dDCP}^{P} we can use the same reasoning since the order isomorphism α^{-1} in particular is a projection. And the claim holds in full subcategories, since α^{-1} will then again be in that subcategory. \square

¹⁶Again note that the second requirement is satisfied as soon as \mathfrak{E} is max-preserving.

Finitary dynamical expanding system Before considering limits of dynamical dcpos, we need to define the appropriate diagrams over which the limits are taken. In domain theory, these are known as expanding systems, and they generalize to dynamical dcpos as follows.

5.2.14. DEFINITION. An *expanding system of dynamical dcpos* is a structure $(\mathfrak{D}_i, p_{ij})_I$ where (I, \leq) is a directed preorder (called the *index set*), the \mathfrak{D}_i are dynamical dcpos, and, for $i \leq j$ in I , $p_{ij} : \mathfrak{D}_j \rightarrow \mathfrak{D}_i$ is a dynamical morphism such that,

1. For all $i \leq j$ in I , $p_{ij} : D_j \rightarrow D_i$ is a projection.
2. For all $i \in I$, $p_{ii} = \text{id}_{\mathfrak{D}_i}$.
3. For all $i \leq j \leq k$ in I , $p_{ik} = p_{ij} \circ p_{jk}$.

An expanding system of dynamical dcpos $(\mathfrak{D}_i, p_{ij})_I$ is:

4. *upward deterministic* :iff for all $i \in I$, if $\exists a_i, b_i \neq b'_i \in \max D_i : b_i, b'_i \geq f_i(a_i)$, then there is $j \geq i$ in I such that $\forall a_j, b_j, b'_j \in \max D_j : \text{if } p_{ij}(a_j) = a_i, p_{ij}(b_j) = b_i, p_{ij}(b'_j) = b'_i$, then $b_j \not\geq f_j(a_j)$ or $b'_j \not\geq f_j(a_j)$
5. *downward deterministic* :iff for all $i \in I$, if $\exists a_i \neq a'_i, b_i \in \max D_i : b_i \geq f_i(a_i), f_i(a'_i)$, then there is $j \geq i$ in I such that $\forall a_j, a'_j, b_j \in \max D_j : \text{if } p_{ij}(a_j) = a_i, p_{ij}(a'_j) = a'_i, p_{ij}(b_j) = b_i$, then $b_j \not\geq f_j(a_j)$ or $b_j \not\geq f_j(a'_j)$.
6. *eventually valuation-preserving* :iff (a) all \mathfrak{D}_i are finite and, for all $i \in I$ and $U_i \in \Sigma(D_i)$, there is $j_0 \geq i$ such that, for all $j \geq j_0$, we have

$$v_j \left(f_j^{-1} (p_{ij}^{-1} (\downarrow \max U_i)) \cap \max D_j \right) = v_j \left(p_{ij}^{-1} (\downarrow \max U_i) \cap \max D_j \right),^{17}$$

and (b) for all $i \leq j$ in I , if $a_i, b_i \in \max D_i$ with $f_i(a_i) \leq b_i$, then there is $a_j, b_j \in \max D_j$ such that $p_{ij}(a_j) = a_i$ and $p_{ij}(b_j) = b_i$ and $f_j(a_j) \leq b_j$.

A *finitary dynamical expanding system* is an upward deterministic expanding system of dynamical dcpos $(\mathfrak{D}_i, p_{ij})_I$ where I is countable and each \mathfrak{D}_i is a finite max-normalized dynamical Scott domain. It is *standard* if, additionally, the \mathfrak{D}_i are max-surjective and $(\mathfrak{D}_i, p_{ij})_I$ is downward deterministic and eventually valuation-preserving.

To state the main result on limits of these diagrams, we've introduced the notion of a *restricted limit*: If \mathbf{C} is a category, \mathbf{D} a full subcategory, and $\mathbf{F} : \mathbf{l} \rightarrow \mathbf{C}$ a diagram, a *D-limit* of \mathbf{F} is a cone (A, f_i) to \mathbf{F} in \mathbf{C} ,¹⁸ with A an object in \mathbf{D} , such

¹⁷Qua sets of maximal elements of finite domains, both sets are Scott-open.

¹⁸I.e., A is an object of \mathbf{C} and, for each object i of \mathbf{l} , $f_i : A \rightarrow \mathbf{F}(i)$ is a morphism in \mathbf{C} such that, for every morphism $\iota : i \rightarrow j$ in \mathbf{l} , $\mathbf{F}(\iota) \circ f_i = f_j$.

that, for any cone (B, g_i) to F in C , if $B \in D$, then there is a unique morphism $u : B \rightarrow A$ (in C and hence also in D) such that $f_i \circ u = g_i$ for all objects i in I . Note that, as the usual proof shows, if it exists, (A, f_i) is unique up to unique isomorphism in D : If (A', f'_i) is another D -limit of F , there is a unique isomorphism $u : A \rightarrow A'$ in D with $f'_i \circ u = f_i$. Thus, if the categories are clear from context, we call A the *restricted limit*.

5.2.15. THEOREM. *Let $(\mathfrak{D}_i, p_{ij})_I$ be a finitary dynamical expanding system. Write $\mathfrak{D}_i = (D_i, v_i, f_i)$. Then*

1. $D := \{ \langle a(i) : i \in I \rangle \in \prod_{i \in I} D_i : \forall i \leq j \in I . a(i) = p_{ij}(a(j)) \}$ with the pointwise order¹⁹ is a Scott domain and $\max D$ is closed in the Lawson topology. For $i \in I$, the function $p_i : D \rightarrow D_i$ defined by $p_i(a) := a(i)$ is a max-preserving Scott-continuous projection.
2. There is a unique continuous valuation $v : \sigma(D) \rightarrow [0, \infty]$ such that, for all $U_i \in \Sigma(D_i)$, we have $v_i(U_i) = v(p_i^{-1}(U_i))$. Moreover, v is max-normalized.
3. There is a largest (in the pointwise ordering) function $f : D \rightarrow D$ that is Scott-continuous and max-preserving such that, for all $a \in D$ and $i \in I$, $f(a)(i) \geq f_i(a(i))$. If $a \in \max D$, then $f(a)$ is the (unique) element $b \in \max D$ with $b \geq f_i(a(i))$ for all $i \in I$.

Hence $\mathfrak{D} := (D, v, f)$ is a max-normalized and max-preserving dynamical Scott domain. Moreover, (\mathfrak{D}, p_i) is a $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of the diagram $(\mathfrak{D}_i, p_{ij})_I$ in $\text{dSCO}_{\text{n}}^{\text{p}}$:

4. (\mathfrak{D}, p_i) is a cone to the diagram with \mathfrak{D} in $\text{dSCO}_{\text{nm}}^{\text{p}}$.
5. If $(\mathfrak{E}, \beta_i)_I$ is a cone to the diagram with \mathfrak{E} in $\text{dSCO}_{\text{nm}}^{\text{p}}$, then there is a unique morphism $\beta : \mathfrak{E} \rightarrow \mathfrak{D}$ in $\text{dSCO}_{\text{n}}^{\text{p}}$, which is defined by $\beta(e) := \langle \beta_i(e) : i \in I \rangle$, such that $\beta_i = p_i \circ \beta$ for all $i \in I$.

If $(\mathfrak{D}_i, p_{ij})_I$ additionally is standard, then

6. f is bijective on $\max D$.
7. \mathfrak{D} is valuation-preserving.

Dynamical domains In accordance with the tradition of calling the kinds of depos under study simply ‘domains’, we define:

5.2.16. DEFINITION. A *dynamical domain* is a dynamical depo \mathfrak{D} that is the $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of a finitary dynamical expanding system. A *standard dynamical*

¹⁹I.e., for $a, b \in D$, $a \leq b$ iff $\forall i \in I : a(i) \leq b(i)$.

domain is a dynamical dcpo \mathfrak{D} that is the $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of a standard finitary dynamical expanding system.

The full subcategory of dSCO_{nm} whose objects are dynamical domains is denoted dDOM . The full subcategory of dSCO_{nm} whose objects are standard dynamical domains is denoted dDOM_s . (Note that morphisms hence are not required to be projections.)

We also note the following useful lemma for later on.

5.2.17. LEMMA. *Let $(\mathfrak{D}, p_i)_I$ be a $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limit of a finitary dynamical expanding system $(\mathfrak{D}_i, p_{ij})_I$. Write $\mathfrak{D} = (D, v, f)$. Let $A \subseteq \max D$ be closed in the relative Scott topology and $x \in \max D$. If, for all $i \in I$, $p_i(x) \in p_i(A)$, then $x \in A$.*

PROOF. Without loss of generality, we can assume that \mathfrak{D} is the limit constructed in theorem 5.2.15.²⁰ From the assumption, for each $i \in I$, there is $y_i \in A$ such that $p_i(x) = p_i(y_i)$. Hence $(y_i)_I$ is a net in the topological space $\max D$. It suffices to show that (y_i) converges to x , because the limit x of a net whose elements are in a closed set A also is in A .

So let U be a basic open set of $\max D$ with $x \in U$, and show that there is $j \in I$ such that, for all $i \geq j$, $y_i \in U$. The Scott topology on D coincides with the relative product topology. (See Abramsky and Jung (see 1994, ex. 3.3.12 (18)); the \mathfrak{D}_i are, in particular, continuous.) So also the relative Scott topology on $\max D$ coincides with relative product topology, whence $U = \bigcap_{k=1}^n p_{i_k}^{-1}(U_{i_k}) \cap \max D$ for some $i_1, \dots, i_n \in I$ and $U_{i_k} \in \Sigma(D_{i_k})$ (for $k = 1, \dots, n$). Since I is directed, let $I \ni j \geq i_1, \dots, i_n$. Then we have, for $i \geq j$, that, for $k = 1, \dots, n$,

$$p_{i_k}(y_i) = p_{i_k i} p_i(y_i) = p_{i_k i} p_i(x) = p_{i_k}(x) \in U_{i_k},$$

so $y_i \in \bigcap_{k=1}^n p_{i_k}^{-1}(U_{i_k}) \cap \max D = U$, as needed. \square

5.2.5 Recap from chapter 4

Before we introduce the last two categories, we recap the results from the previous chapter.

²⁰To be precise: If $(\mathfrak{D}, p_i)_I$ and $(\mathfrak{D}', p'_i)_I$ are two $\text{dSCO}_{\text{nm}}^{\text{p}}$ -limits of $(\mathfrak{D}_i, p_{ij})_I$, there is a (unique) isomorphism $u : \mathfrak{D}' \rightarrow \mathfrak{D}$ with $p_i \circ u = p'_i$. Assume the claim holds for $(\mathfrak{D}, p_i)_I$ and show that it holds for $(\mathfrak{D}', p'_i)_I$. So let $\emptyset \neq A' \subseteq \max D'$ be closed, $x' \in \max D'$, and $\forall i : p'_i(x') \in p'_i(A')$. Then, since $u : D' \rightarrow D$ is an order isomorphism, $A := u(A')$ is a nonempty closed subset of $\max D$, $x := u(x') \in \max D$, and for $i \in I$, $p_i(x) = p_i \circ u(x') = p'_i(x') \in p'_i(A') = p_i \circ u(A') = p_i(A)$. So $x \in A$, whence $u(x') = x = u(y)$ for some $y \in A'$, so, by injectivity of u , $x' = y \in A'$.

From domains to systems The point of dynamical domains is that they are computational models for dynamical systems: the maximal (‘ideal’) elements of the domain form a dynamical system which is computationally modeled by the other (‘real’) elements of the domain. Formally, recalling from theorem 4.5.1 in chapter 4, if $\mathfrak{D} = (D, v, f)$ is in \mathbf{dDOM} , then

1. $\max D$ with the relative Scott topology $\Sigma(D) \upharpoonright \max D$ is a compact zero-dimensional Polish space,
2. f restricts to a continuous function on $\max D$,
3. v determines a unique probability measure μ_v on $\mathcal{B}(D, \Lambda)$ extending v ,
4. $\mathcal{B}(\max D) \subseteq \mathcal{B}(D, \Lambda)$ and $\mu_v \upharpoonright \mathcal{B}(\max D)$ is a probability measure on $\max D$.

Thus, we obtain the compact zero-dimensional measured topological system

$$\mathbf{S}(\mathfrak{D}) := (\max D, \Sigma(D) \upharpoonright \max D, \mu_v \upharpoonright \mathcal{B}(\max D), f \upharpoonright \max D)$$

which induces the general dynamical system $\mathbf{JS}(\mathfrak{D})$. Moreover, if \mathfrak{D} is standard, then both the topological system $\mathbf{S}(\mathfrak{D})$ and the dynamical system $\mathbf{JS}(\mathfrak{D})$ are standard. We call $\mathbf{S}(\mathfrak{D})$ the *topological system modeled by \mathfrak{D}* and $\mathbf{JS}(\mathfrak{D})$ (resp., $\overline{\mathbf{JS}}(\mathfrak{D})$ in case \mathfrak{D} is standard) the *dynamical system model by \mathfrak{D}* .

From systems to domains Conversely, we want to build such computational models for any given dynamical system: we build finite dynamical dcpos \mathfrak{D}_i from observing the original system \mathfrak{X} in such a way that, as we keep refining the ‘observation granularity’ i , we obtain, in the limit, a dynamical domain \mathfrak{D} which models a system $\mathbf{S}(\mathfrak{D})$ that is isomorphic to the original system \mathfrak{X} . Formally this is done as follows. Here we’ll state it only for topological systems (since this is all we’ll need), but it also works for general and standard dynamical systems.

Let $\mathfrak{X} = (X, \tau, \mu, T)$ be in \mathbf{TS}_0 and let \mathcal{B} be a countable clopen topological basis for \mathfrak{X} in the sense of chapter 4, i.e., \mathcal{B} is a countable basis for (X, τ) consisting of clopen sets that is closed under finite intersection. We called \mathcal{B} *forward* (resp., *backward*) *closed* if, for all $U \in \mathcal{B}$, we have $T(U) \in \mathcal{B}$ (resp., $T^{-1}(U) \in \mathcal{B}$). To stress the dependence on the dynamics, we also say ‘dynamically closed’ or ‘ T -closed’. [If \mathfrak{X} is standard, additionally assume that \mathcal{B} is forward closed.]

Define $I(\mathcal{B})$ as the set of pairs (n, \mathcal{C}) where $n \in \omega$ (the set of non-negative integers) and \mathcal{C} is a finite \mathcal{B} -cover (i.e., $\mathcal{C} \subseteq \mathcal{B}$ is finite and every $x \in X$ is in some element of \mathcal{C}). Order $I(\mathcal{B})$ by $(n, \mathcal{C}) \leq (m, \mathcal{D})$ iff $n \leq m$ and

1. for all $D \in \mathcal{D}$, there is $C \in \mathcal{C}$ such that $D \subseteq C$, and
2. for all $x \in C \in \mathcal{C}$, there is $D \in \mathcal{D}$ such that $x \in D \subseteq C$.

(We write $\mathcal{C} \preceq \mathcal{D}$ if these two conditions are satisfied for \mathcal{C} and \mathcal{D} .) Then $I(\mathcal{B})$ is a directed preorder.

For $i = (n, \mathcal{C}) \in I(\mathcal{B})$ and $x \in X$, define the ‘observation history’ of x :

$$\mathcal{O}_i(x) := \left\{ t \in \mathcal{C}^n : x \in \bigcap_{k=0}^{n-1} T^{-k}(t(k)) \right\}.$$

Define $\mathbf{H}_i := \{\mathcal{O}_i(x) : x \in X\}$, which is a finite set. And define

$$D_i := \mathbf{P}(\mathbf{H}_i)$$

as the set of nonempty subsets of \mathbf{H}_i ordered by reverse inclusion, whence D_i is a finite Scott domain. One reason for moving to this so-called powerdomain, which is a domain-theoretic tool for analyzing non-deterministic functions, is that the naturally induced dynamics on \mathbf{H}_i is not deterministic but it is on D_i : We define the Scott-continuous function $f_i : D_i \rightarrow D_i$ by

$$f_i(M) := \left\{ \mathcal{O}_i(Tx) : \mathcal{O}_i(x) \in M \right\}.$$

To define the valuation on D_i , consider observational equivalence: $x \approx_i x'$ iff $\mathcal{O}_i(x) = \mathcal{O}_i(x')$. This partitions X into finitely many equivalence classes $[x]_i$ that are clopen in X : they can be written as Boolean combinations of sets in $\bigcup_{k \in \omega} T^{-k}\mathcal{B}$. So we define the valuation $v_i : \Sigma(D_i) \rightarrow [0, 1]$ by

$$v_i(U) := \sum_{k=1}^m \mu([x_k]_i) \quad \text{where } \max U = \left\{ \{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\} \right\}.$$

Then $\mathfrak{D}_i := (D_i, v_i, f_i)$ is a max-normalized finite dynamical Scott domain.

For $i \leq j$ in $I(\mathcal{B})$, define $p_{ij} : D_j \rightarrow D_i$ by

$$p_{ij}(M) := \left\{ \mathcal{O}_i(x) : \mathcal{O}_j(x) \in M \right\}.$$

Then $(\mathfrak{D}_i, p_{ij})_{I(\mathcal{B})}$ is a finitary dynamical expanding system, so we can construct the $\text{dSCO}_{\text{nm}}^{\text{P}}$ -limit

$$\mathbf{D}(\mathfrak{X}, \mathcal{B}) := \mathfrak{D} := (D, v, f)$$

as in theorem 5.2.15. Hence $\mathbf{D}(\mathfrak{X}, \mathcal{B})$ is in dDOM . [If \mathfrak{X} is standard, then $(\mathfrak{D}_i, p_{ij})_{I(\mathcal{B})}$ is standard, so $\mathbf{D}(\mathfrak{X}, \mathcal{B})$ is standard.] We call $\mathbf{D}(\mathfrak{X}, \mathcal{B})$ the *observation domain* of $(\mathfrak{X}, \mathcal{B})$.

Finally, the *canonical embedding* $\varphi_X : X \rightarrow \max D$ defined by

$$\varphi_X(x) := \langle \{\mathcal{O}_i(x)\} : i \in I(\mathcal{B}) \rangle$$

is an injective and relatively open morphism $\mathfrak{X} \rightarrow \mathbf{S}(\mathfrak{D})$ in TS_0 whose image is dense in $\mathbf{S}(\mathfrak{D})$. The φ_X -preimages of clopen sets of $\mathbf{S}(\mathfrak{D})$ can be written as Boolean combinations of equivalence classes $[x]_i$. If X is compact and $\mathcal{B} = \text{Clp}(X)$, then $\varphi : \mathfrak{X} \rightarrow \mathbf{S}(\mathfrak{D})$ is an isomorphism in $\text{TS}_{0\text{c}}$ (in this case we also call φ_X the canonical homeomorphism). If context allows, we may drop the subscript of φ_X .

5.2.6 Categories of based measured topological systems

In this subsection, we define the two categories \mathbf{bTS}_0 and \mathbf{bTS}_{0s} (and their subcategories $\overline{\mathbf{bTS}}_0$ and $\overline{\mathbf{bTS}}_{0s}$) of based measured topological systems.

Given that, as seen above, the construction of the observation domain of a measured topological system is relative to a choice of basis, it is suggestive to make that choice explicit by working with pairs $(\mathfrak{X}, \mathcal{B})$ of a topological system with an appropriate basis. Such a pair $(\mathfrak{X}, \mathcal{B})$ of a measured topological system \mathfrak{X} and a topological basis for it may be called a *based* measured topological systems. And we're interested in those for which the above construction of the observation domain goes through:

5.2.18. PROPOSITION. *The following define categories:*

1. \mathbf{bTS}_0 : *Objects are pairs $(\mathfrak{X}, \mathcal{B})$ with \mathfrak{X} in \mathbf{TS}_0 and \mathcal{B} a countable clopen topological basis for \mathfrak{X} .²¹ And a morphism $(\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Y}, \mathcal{B}_Y)$ is a morphism $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ in \mathbf{TS}_0 such that $\varphi^{-1}\mathcal{B}_Y \subseteq \mathcal{B}_X$ (i.e., for all $B \in \mathcal{B}_Y$, $\varphi^{-1}(B) \in \mathcal{B}_X$).*
2. \mathbf{bTS}_{0s} : *The full subcategory of \mathbf{bTS}_0 whose objects $(\mathfrak{X}, \mathcal{B})$ are such that \mathfrak{X} is in \mathbf{TS}_{0s} and \mathcal{B} is forward closed.*
3. $\overline{\mathbf{bTS}}_0$: *The full subcategory of \mathbf{bTS}_0 whose objects $(\mathfrak{X}, \mathcal{B})$ are such that \mathcal{B} is backward closed and closed under Boolean operations (intersection, union, and complement).²²*
4. $\overline{\mathbf{bTS}}_{0s}$: *The full subcategory of \mathbf{bTS}_0 whose objects $(\mathfrak{X}, \mathcal{B})$ are such that \mathfrak{X} is in \mathbf{TS}_{0s} and \mathcal{B} is forward closed, backward closed, and closed under Boolean operations.*

Morphism composition is function composition and identity morphisms are identity functions.

PROOF. We need to show that \mathbf{bTS}_0 is a well-defined category: If $\varphi : (\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Y}, \mathcal{B}_Y)$ and $\psi : (\mathfrak{Y}, \mathcal{B}_Y) \rightarrow (\mathfrak{Z}, \mathcal{B}_Z)$ are morphisms in \mathbf{bTS}_0 , then $\psi \circ \varphi : \mathfrak{X} \rightarrow \mathfrak{Z}$ is a morphism in \mathbf{TS}_0 and $(\psi \circ \varphi)^{-1}(\mathcal{B}_Z) = \varphi^{-1}(\psi^{-1}(\mathcal{B}_Z)) \subseteq \mathcal{B}_X$. So $\psi \circ \varphi : (\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Z}, \mathcal{B}_Z)$ is a morphism in \mathbf{bTS}_0 . Moreover, given $(\mathfrak{X}, \mathcal{B}_X)$ in \mathbf{bTS}_0 , the identity morphism $\text{id}_{\mathfrak{X}} : \mathfrak{X} \rightarrow \mathfrak{X}$ in \mathbf{TS}_0 is a morphism $(\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{X}, \mathcal{B}_X)$ since $\text{id}_{\mathfrak{X}}^{-1}\mathcal{B}_X = \mathcal{B}_X$. This choice of morphism composition and identity morphism satisfies the identity and associativity axioms qua function composition and identity function. \square

²¹Recall that we assume topological bases to be closed under finite intersection.

²²By assumption, \mathcal{B} is already closed under intersection. So being closed under Boolean operations just requires \mathcal{B} to be closed under complement (if $U \in \mathcal{B}$, then $U^c = X \setminus U \in \mathcal{B}$); because then it also is closed under union: if $U, V \in \mathcal{B}$, then $U \cup V = (U^c \cap V^c)^c \in \mathcal{B}$.

Comments: First, to recap from chapter 4, we've motivated bases of dynamical systems as the set of possible observations that we can make about the system. The requirement of bases to be closed under finite intersection then means that we can form the conjunction of observations. And also the other properties, 'countable' and 'clopen', are quite natural: Countability is necessary if these observations are accessible to us in a 'computable' way, to make sure that every observation in the basis has an index $n \in \mathbb{N}$ with which we can refer to it. Clopenness expresses the idea that the observations in the basis are finitely decidable: under the well-known computational interpretation of topology (Smyth 1983; Vickers 1989), the open sets of a topology are 'semi-decidable properties' of the points of the space. Thus, the sets that not only are open but also have an open complement—i.e., the clopen sets—are 'decidable properties' of the space.

Second, other closure properties are possible. Logical ones are, for example, closure under Boolean operations: for negation this says that if we can observe whether the system is in a U -state, we can also observe whether it is in a U^c -state (i.e., conclude 'yes' if measuring U yields a negative answer). Dynamical closure properties are, for example, backward or forward closure: backward closure says that if we can observe whether the system is in a U -state, we can also observe whether it is in a $T^{-1}(U)$ state (i.e., wait one time step and then measure U). These further closure properties may be more controversial (thinking of non-classical logics or bounded observation time). Thus, the full subcategories where these closure properties are assumed are notationally highlighted by the 'closure bar' $\bar{\cdot}$. We show below that they are reflective subcategories, i.e., there is an optimal way of closing a basis under the logical and dynamical operations.

Third, thus we can regard bases as 'computability structure' describing in a computational way the possible observations that we can make. So it is natural to demand that the morphisms should preserve this computability structure: for a morphism $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$, if U is a possible observation in \mathfrak{Y} , also $\varphi^{-1}(U)$ is a possible observation in \mathfrak{X} .

Fourth, if the system \mathfrak{X} in \mathbf{TS}_0 has a compact state space (i.e., is in \mathbf{TS}_{0c}), there is a natural choice of basis: the set of all clopen sets $\mathbf{Clp}(X)$. (It consists of clopen sets, is closed under Boolean operation, forms, by zero-dimensionality, a basis, and, by compactness and second-countability, is countable.) The choice is natural in the sense that if \mathcal{B} is another clopen basis for \mathfrak{X} , then one can show that $I(\mathcal{B})$ is a cofinal subset of $I(\mathbf{Clp}(X))$, whence they give rise to the same limit.²³ This naturality is not available for general \mathfrak{X} in \mathbf{TS}_0 : if X has uncountably many clopen subsets—e.g., if X is the Baire space—, then there can be no countable clopen basis \mathcal{B}_0 such that, for any countable clopen basis \mathcal{B} , we

²³Proof: Since $\mathcal{B} \subseteq \mathbf{Clp}(X)$, we have $I(\mathcal{B}) \subseteq I(\mathbf{Clp}(X))$. For cofinality, if (m, \mathcal{D}) is in $I(\mathbf{Clp}(X))$, then each member U of \mathcal{D} can be written, qua open set and \mathcal{B} being a basis, as a union of elements from \mathcal{B} , whence, since U is a closed subset of a compact space, as finite union. Collecting the sets used in these finite unions for the finitely many $U \in \mathcal{D}$ into the \mathcal{B} -cover \mathcal{C} , we have $(m, \mathcal{D}) \leq (m, \mathcal{C}) \in I(\mathcal{B})$.

have $I(\mathcal{B}) \subseteq I(\mathcal{B}_0)$ (let alone cofinal).²⁴ In fact, if, for two bases \mathcal{B} and \mathcal{B}' of \mathfrak{X} , we have $I(\mathcal{B}) \subseteq I(\mathcal{B}')$ cofinal, then \mathcal{B} and \mathcal{B}' are identical after closing under finite union.²⁵ This indicates that, for general \mathfrak{X} in \mathbf{TS}_0 , we have to make a *real* choice of basis.

Fifth, concerning related work we've already discussed the based Polish spaces of Dahlqvist, Danos, and Garnier (2016) in the introduction.

5.2.7 Categories of max-reflective dynamical domains

As the last category, we introduce the full subcategory \mathbf{dDOM}_r of \mathbf{dDOM} (and \mathbf{dDOM}_{rs} of \mathbf{dDOM}_s) whose objects are those dynamical domains that satisfy a property that we call ‘max-reflective’. Toward the end of this chapter, we make the surprising observation that it is precisely this simple domain-theoretic property that turns into an equivalence the adjunction between systems and domains. But, to have the definitions of all the categories in one place, we already introduce this property here.

5.2.19. DEFINITION. We call a dcpo D *max-reflective* if

1. For all $a \in D$, $a = \bigwedge(\uparrow a \cap \max D)$.
2. If $\emptyset \neq A \subseteq \max D$ is closed in the relative Scott topology, then $\bigwedge A$ exists and $(\uparrow \bigwedge A) \cap \max D = A$.

Note that if D is a Scott domain, the requirement that $\bigwedge A$ exists in clause 2 can be omitted since nonempty infima always exist.

The intuition behind this concept (and its name) is that not only, as in every domain, do the non-maximal (or ‘real’) elements of the domain provide a model for the maximal (or ‘ideal’) elements, but also, conversely, the maximal elements ‘reflect’ the non-maximal ones in the following sense: (1) any non-maximal element can be recovered from the maximal elements that it approximates, and (2) a nonempty set A of maximal elements that is closed (‘with respect to approximation’) can be recovered as those maximal elements that contain the ‘information’ that is in A (i.e., that are approximated by the element $\bigwedge A$ that contains precisely the ‘information’ that is in A). The following provides some concrete intuition in the finite case.

²⁴For a short proof that Baire space has continuum many clopen sets, see Scott (2013). For the other claim: Assume there were such \mathcal{B}_0 . Since it is countable but $\mathbf{Clp}(X)$ uncountable, there is a clopen C with $C \notin \mathcal{B}_0$. Let \mathcal{B} be the closure under finite intersection of $\mathcal{B}_0 \cup \{C\}$, which is a countable clopen basis, so $I(\mathcal{B}) \subseteq I(\mathcal{B}_0)$. This implies $\mathcal{B} \subseteq \mathcal{B}_0$ (if $U \in \mathcal{B}$, then $(1, \{U, X\}) \in I(\mathcal{B}) \subseteq I(\mathcal{B}_0)$, so $\{U, X\}$ is a \mathcal{B}_0 -cover, so $U \in \mathcal{B}_0$). But then $C \in \mathcal{B}_0$, contradiction.

²⁵Proof: Write $\mathbf{U}(\mathcal{B})$ for the closure of \mathcal{B} under finite union, similarly for \mathcal{B}' . As in the previous footnote, $\mathcal{B} \subseteq \mathcal{B}' (\subseteq \mathbf{U}(\mathcal{B}'))$. Conversely, we show $\mathcal{B}' \subseteq \mathbf{U}(\mathcal{B})$: If $U \in \mathcal{B}'$, then, by cofinality, there is $(n, \mathcal{C}) \geq (1, \{U, X\})$ in $I(\mathcal{B}')$. So, by clause (2) of cover-refinements, for all $x \in U$, there is $U_x \in \mathcal{C}$ with $x \in U_x \subseteq U$. Since \mathcal{C} is finite, U can hence be written as finite union of elements from $\mathcal{C} \subseteq \mathcal{B}$.




Figure 5.2: Examples of finite Scott domains and their stance on the two conditions (1) and (2) of being max-reflective.

5.2.20. EXAMPLE. 1. Figure 5.2 provides finite Scott domains for all four combinations of satisfying and violating the two conditions (1) and (2) of being max-reflective. Note how (d) is the ‘fusion’ of (b) and (c).

2. If M is a nonempty finite set, then the Smyth powerdomain $\mathbf{P}(M)$ of M (where M is regarded with the discrete order) is a max-reflective Scott domain. (Recall that then $\mathbf{P}(M)$ is the set of nonempty subsets of M ordered by reverse inclusion.) Figure 5.2a depicts $\mathbf{P}(2)$ where $2 = \{0, 1\}$.

Proof: Qua finite lattice without top element (but with least element), $\mathbf{P}(M)$ is a Scott domain. Regarding (1), if $a \in \mathbf{P}(M)$, then $\bigwedge(\uparrow a \cap \max \mathbf{P}(M)) = \bigcup\{\{m\} : \{m\} \subseteq a\} = a$. Regarding (2), if $\emptyset \neq A \subseteq \max \mathbf{P}(M)$, then $A = \{\{m_1\}, \dots, \{m_n\}\}$ for some $m_1, \dots, m_n \in M$ with $n \geq 1$. Then $\bigwedge A = \bigcup A = \{m_1, \dots, m_n\} =: a$ and $(\uparrow \bigwedge A) \cap \max \mathbf{P}(M) = \{\{m\} : \{m\} \subseteq \bigwedge A\} = \{m\} = A$.

We call a dynamical dcpo $\mathfrak{D} = (D, v, f)$ *max-reflective* if D is max-reflective. We show that this is a finitary concept in the sense that if the limit \mathfrak{D} fails to have it, we’ll realize that at some finite stage of construction:

5.2.21. PROPOSITION. *Let \mathfrak{D} be a $\text{dSCO}_{\text{nm}}^{\text{P}}$ -limit of a finitary dynamical expanding system $(\mathfrak{D}_i, p_{ij})_I$. Then the following are equivalent:*

1. Each \mathfrak{D}_i is max-reflective.
2. \mathfrak{D} is max-reflective.

PROOF. We fix the notation $\mathfrak{D} = (D, v, f)$ and $\mathfrak{D}_i = (D_i, v_i, f_i)$. We first observe that, for $a \in D$, we have

$$p_i(\uparrow a \cap \max D) = \uparrow p_i(a) \cap \max D_i. \quad (5.1)$$

Indeed, if $y = p_i(x)$ for $x \in \uparrow a \cap \max D$, then, since p_i is monotone and max-preserving, $p_i(a) \leq p_i(x) \in \max D_i$. So $y = p_i(x) \in \uparrow p_i(a) \cap \max D_i$. Conversely, if $y \in \uparrow p_i(a) \cap \max D_i$, then $a \in D$ with $p_i(a) \leq y \in \max D_i$, so, since p_i is max-bisimulative, there is $a \leq x \in \max D$ with $p_i(x) = y$, so $y = p_i(x) \in p_i(\uparrow a \cap \max D)$.

(1) \Rightarrow (2). Assume all D_i are max-reflective. Without loss of generality, D is constructed as in theorem 5.2.15 (if D were only isomorphic to the limit constructed

there, it, too, would be max-reflective since that is a purely-domain theoretic property). To show that D is max-reflective, we need to check conditions (1)–(2).

Concerning condition (1), let $a \in D$ and show $a = \bigwedge \uparrow a \cap \max D$. Indeed, for $i \in I$, we have, since p_i commutes, qua projection, with existing infima,

$$p_i(\bigwedge \uparrow a \cap \max D) = \bigwedge p_i(\uparrow a \cap \max D) \stackrel{(5.1)}{=} \bigwedge \uparrow p_i(a) \cap \max D_i,$$

which, since D_i is max-reflective, equals $p_i(a)$. So a and $\bigwedge \uparrow a \cap \max D$ are identical at every component, and hence identical, as needed.

Concerning condition (2), let $\emptyset \neq A \subseteq \max D$ be closed. Since nonempty infima exist in Scott domains (and D is a Scott domain), we need to show $\uparrow \bigwedge A \cap \max D = A$. The \supseteq -direction is easy: If $x \in A$, then $x \in \max D$ and $x \geq \bigwedge A$. For the other direction, let $x \in \uparrow \bigwedge A \cap \max D$ and show $x \in A$. The idea is to use lemma 5.2.17:

Note that, for any $i \in I$, $p_i(A)$ is nonempty (since A is nonempty) and a subset of $\max D_i$ (since $A \subseteq \max D$ and p_i is max-preserving). Moreover, D_i is finite, so the Lawson topology on it is discrete, so, since the relative Lawson topology and the relative Scott topology agree, $\max D_i$ is a discrete space, whence $p_i(A)$ also is closed. Since D_i is max-reflective and p_i commutes with nonempty infima,

$$p_i(A) = \uparrow \bigwedge p_i(A) \cap \max D_i = \uparrow p_i(\bigwedge A) \cap \max D_i \stackrel{(5.1)}{=} p_i(\uparrow \bigwedge A \cap \max D).$$

Since $x \in \uparrow \bigwedge A \cap \max D$, we have $p_i(x) \in p_i(\uparrow \bigwedge A \cap \max D) = p_i(A)$. Now, lemma 5.2.17 applies and yields $x \in A$.

(2) \Rightarrow (1). Assume D is max-reflective. Let $i \in I$ and show that D_i is max-reflective.

Concerning condition (1), let $b \in D_i$ and show $b = \bigwedge(\uparrow b \cap \max D_i)$. Since p_i is surjective (qua projection), there is $a \in D$ such that $p_i(a) = b$. Since \mathfrak{D} is max-reflective, $a = \bigwedge \uparrow a \cap \max D$, so, since projections commute with existing infima,

$$\begin{aligned} b = p_i(a) &= p_i(\bigwedge \uparrow a \cap \max D) = \bigwedge p_i(\uparrow a \cap \max D) \\ &\stackrel{(5.1)}{=} \bigwedge \uparrow p_i(a) \cap \max D_i = \bigwedge \uparrow b \cap \max D_i. \end{aligned}$$

Concerning condition (2), let $\emptyset \neq B \subseteq \max D_i$ be closed in the relative Scott topology. Since D_i is a Scott domain, $\bigwedge B$ exists, so we need to show $\uparrow \bigwedge B \cap \max D_i = B$. Since B is closed in the relative Scott topology, we have $B = C \cap \max D_i$ for some Scott-closed $C \subseteq D_i$. Define $A := p_i^{-1}(C) \cap \max D$.

We claim that $p_i(A) = B$. Indeed, if $y = p_i(x)$ for $x \in A$, then, by definition of A , $y = p_i(x) \in C$ and, since $x \in \max D$ and p_i is max-preserving, $y = p_i(x) \in \max D_i$, so $y \in B$. Conversely, if $y \in B$, then, by surjectivity, there is $a \in D$ with $p_i(a) = y$. Choose a maximal $x \geq a$ in D . Then, by monotonicity, $p_i(x) \geq p_i(a) =$

y , whence, by maximality of y , $p_i(x) = y \in B \subseteq C$, so $x \in p_i^{-1}(C) \cap \max D = A$, and $y = p_i(x) \in p_i(A)$.

In particular, $A \neq \emptyset$ (otherwise $B = p_i(A) = \emptyset$), and, since p_i is Scott-continuous, $A \subseteq \max D$ is closed in the relative Scott topology on $\max D$. Since D is max-reflective, $\uparrow \bigwedge A \cap \max D = A$. So, since projections commute with nonempty infima,

$$\begin{aligned} B = p_i(A) &= p_i(\uparrow \bigwedge A \cap \max D) \stackrel{(5.1)}{=} \uparrow p_i(\bigwedge A) \cap \max D_i \\ &= \uparrow \bigwedge p_i(A) \cap \max D_i = \uparrow \bigwedge B \cap \max D_i, \end{aligned}$$

as needed. □

5.2.22. DEFINITION. Let \mathbf{dDOM}_r be the full subcategory of \mathbf{dDOM} consisting of max-reflective dynamical domains (equivalently, that are $\mathbf{dSCO}_{\text{nm}}^{\text{p}}$ -limits of finitary dynamical expanding systems consisting of max-reflective dynamical dcpos). Analogously, we define the full subcategory \mathbf{dDOM}_{rs} of \mathbf{dDOM}_s .

5.3 The bottom layer of the main diagram

In this section, we establish the bottom layer of the main diagram:

$$\mathbf{DS} \xleftarrow{\text{Loc}} \mathbf{bTS}_0 \xrightleftharpoons[\bar{c}]{I_B} \mathbf{TS}_{0c}$$

and similarly for the standard case. We first discuss the left half (section 5.3.1) and then the right half (section 5.3.2–5.3.3).

5.3.1 Dynamical systems as category of fractions

Concerning the relation between the category of dynamical systems \mathbf{DS} and the category of zero-dimensional measured topological systems \mathbf{TS}_0 , we already observed in proposition 5.2.9 that we have the natural functor $J : \mathbf{TS}_0 \rightarrow \mathbf{DS}$ which sends $\mathfrak{X} = (X, \tau, \mu, T)$ to $(X, \mathcal{B}(\tau), \mu, T)$ and $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ to $\varphi : J(\mathfrak{X}) \rightarrow J(\mathfrak{Y})$.

Looking for a closer relationship, the following two questions come naturally. (We’ve also discussed this, in different words, in the introduction.)

First, on the object level, we may wonder whether J ‘hits’ every dynamical system: is every dynamical system generated by some Polish topology, i.e., isomorphic to a system of the form $J(\mathfrak{X})$? (Formally, this means that J is essentially surjective.) We show that this is indeed true: in part by results from the previous chapter and in part by extending a well-known construction on Polish spaces.

Second, on the morphism level, lemma 5.2.4 says that if $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ in \mathbf{TS}_0 is injective on an invariant set of full measure, then $\varphi : J(\mathfrak{X}) \rightarrow J(\mathfrak{Y})$ already is an

isomorphism in \mathbf{DS} , while in \mathbf{TS}_0 this need not be the case. So, for starters, we cannot expect \mathbf{J} to be an equivalence. However, we may wonder whether those φ are the only ones that aren't yet isomorphisms in \mathbf{TS}_0 but become isomorphisms in \mathbf{DS} . Again, we'll show that this is the case: \mathbf{J} becomes an equivalence after those φ are turned into isomorphisms. Category-theoretically, turning a collection of morphisms in a category into isomorphisms is known as a *localization*.

We move the formal development of these result to appendix A. This is because (a) it requires introducing additional theory which is somewhat different from the main strand of the thesis, and (b) these results are self-standing (in particular, are not needed in the remainder of the thesis) and may be of independent interest. In that appendix, we formulate, in section A.1, the above informal claims precisely. In section A.2, we show the first claim on topological realization. In the remaining sections A.3–A.5, we show the second claim about the localization.

5.3.2 Compactification of a system: informally

Now we consider the right half of the bottom layer of the main diagram: ‘compactifying’ a system in \mathbf{bTS}_0 to obtain one in \mathbf{TS}_{0c} . We first motivate and explain this informally in this subsection, before we do it formally in the next subsection.

Let's start by recalling the concept of a compactification from topology. As the name suggests, a compactification of a topological space X is a way c of turning X into a compact space Y : formally, it is a compact topological space Y together with a homeomorphic embedding $c : X \rightarrow Y$ (Engelking 1989, sec. 3.5). A famous example is the Stone–Čech compactification βX : For every Tychonoff space X , there is a compactification $c : X \rightarrow \beta X$ such that every continuous function from X into a compact space Z is continuously extendable to βX (Engelking 1989, sec. 3.6). So the Stone–Čech compactification βX is in a sense optimal in providing a compactification. This optimality is concisely (and even more generally) captured category-theoretically: the Stone–Čech compactification functor β from the category \mathbf{Top} of topological spaces with continuous maps to the category \mathbf{Top}_c^H of compact Hausdorff spaces with continuous maps is left adjoint to the inclusion $\mathbf{l} : \mathbf{Top}_c^H \rightarrow \mathbf{Top}$ (Leinster 2014, ex. 6.3.14).

Can we do something similar for *systems* where, in addition to the topological space, we also have a dynamics (and a measure)? In this section, we provide a positive answer.

There is a natural candidate for such a compactification: If we start with $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 , we can build the observation domain $\mathbf{D}(\mathfrak{X}, \mathcal{B})$. Then we can build the *compact* zero-dimensional measured topological system $\mathbf{SD}(\mathfrak{X}, \mathcal{B})$ that is modeled by the observation domain $\mathbf{D}(\mathfrak{X}, \mathcal{B})$. So we end up in the category \mathbf{TS}_{0c} and have ‘compactified’ the original system \mathfrak{X} with its choice of basis \mathcal{B} .

There also is a natural way of going back from \mathbf{TS}_{0c} to $\bar{\mathbf{bTS}}_0$. To recall, $\bar{\mathbf{bTS}}_0$ is the full subcategory of \mathbf{bTS}_0 defined by requiring its objects $(\mathfrak{X}, \mathcal{B})$ to be such that the basis \mathcal{B} is backward dynamically closed and closed under Boolean operations.

Conceptually, this way back is like an ‘inclusion’ of \mathbf{TS}_{0c} into $\overline{\mathbf{bTS}}_0$:

5.3.1. PROPOSITION. *There is a functor $\mathbf{l}_B : \mathbf{TS}_{0c} \rightarrow \overline{\mathbf{bTS}}_0$ that sends \mathfrak{X} to $(\mathfrak{X}, \mathbf{Clp}(X))$ and $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ to $\varphi : (\mathfrak{X}, \mathbf{Clp}(X)) \rightarrow (\mathfrak{Y}, \mathbf{Clp}(Y))$. It restricts to $\mathbf{TS}_{0cs} \rightarrow \overline{\mathbf{bTS}}_{0s}$.*

PROOF. Write $\mathfrak{X} = (X, \tau, \mu, T)$. As mentioned in section 5.2.6, $\mathbf{Clp}(X)$ is a countable clopen topological basis that is closed under Boolean operations and, since T is continuous, it is also backward closed. So $\mathbf{l}_B(\mathfrak{X})$ is indeed in $\overline{\mathbf{bTS}}_0$. Moreover, since φ is continuous, the preimage of a clopen set is clopen, so $\varphi^{-1}(\mathbf{Clp}(Y)) \subseteq \mathbf{Clp}(X)$, so $\mathbf{l}_B(\varphi) = \varphi$ is indeed a morphism in $\overline{\mathbf{bTS}}_0$. Since \mathbf{l}_B is the identity on morphisms, it preserves composition and identity. Concerning the restriction, if \mathfrak{X} is in \mathbf{TS}_{0cs} , then \mathfrak{X} is in \mathbf{TS}_{0s} and $\mathbf{Clp}(X)$ additionally is forward closed (since T is a homeomorphism). \square

So we wonder whether this provides a compactification in the ‘full-blown’ sense that the SD construction provides a functor that is left adjoint to \mathbf{l}_B . This is almost true. The full story—which we tell in the remainder of this subsection—is a bit more subtle:

We may call the compactification of \mathfrak{X} relative to \mathcal{B} via the SD construction the *computational compactification* since it is obtained via the computational model $\mathbf{D}(\mathfrak{X}, \mathcal{B})$ for \mathfrak{X} .

A standard way of constructing compactifications is as Wallman compactification relative to a base (Johnstone 1982, sec. IV.2.4): Given a topological space X , a *Wallman base* B is a sublattice of the lattice of all open sets of X that is a topological base for X such that, for $x \in U \in B$, there is $V \in B$ with $U \cup V = X$ and $x \notin V$. The *Wallman compactification* of X relative to B then is the space Y of maximal ideals of B with its usual topology. We may call this the *logical compactification* since, as usual in Stone duality, we can regard the points of Y as *logical models* of the *properties* in B (complete and consistent descriptions of how things can be with respect to the properties collected in B).

This suggests to first build a compactification using the standard ‘logical’ techniques and then compare it to the computational compactification using domains. Indeed, the Wallman construction parallels our setting: Given a system \mathfrak{X} with basis \mathcal{B} , we regard the elements of \mathcal{B} as properties with which we can measure the state space X , so it, too, is natural to build the logical compactification of X as consisting of the possible models of \mathcal{B} . However, this requires some logical closure (the Wallman bases are, qua lattices, closed under conjunction and disjunction). And if we want to lift the original dynamics of \mathfrak{X} to a dynamics on these models, we’ll also need some dynamical closure. Thus, this suggests to work over the category $\overline{\mathbf{bTS}}_0$ for the logical compactification.

Fortunately, this is not a restriction since we can close a basis under the logical and dynamical operations in an optimal way: Formally, $\overline{\mathbf{bTS}}_0$ is a reflective subcategory of \mathbf{bTS}_0 .

In the next subsection, we then establish formally what we can now hope for:

5.3.2. THEOREM. *The inclusion $\mathsf{l} : \bar{\mathsf{b}}\mathsf{TS}_0 \rightarrow \mathsf{b}\mathsf{TS}_0$ has a right adjoint, the (logical and dynamical) closure functor $\bar{\cdot} : \mathsf{b}\mathsf{TS}_0 \rightarrow \bar{\mathsf{b}}\mathsf{TS}_0$. And the functor $\mathsf{l}_B : \mathsf{TS}_{0c} \rightarrow \bar{\mathsf{b}}\mathsf{TS}_0$ has a left adjoint, the logical compactification functor $\mathsf{C} : \bar{\mathsf{b}}\mathsf{TS}_0 \rightarrow \mathsf{TS}_{0c}$:*

$$\mathsf{b}\mathsf{TS}_0 \begin{array}{c} \xleftarrow{\mathsf{l}} \\ \xrightarrow{\perp} \\ \xrightarrow{\bar{\cdot}} \end{array} \bar{\mathsf{b}}\mathsf{TS}_0 \begin{array}{c} \xleftarrow{\mathsf{l}_B} \\ \xrightarrow{\top} \\ \xrightarrow{\mathsf{C}} \end{array} \mathsf{TS}_{0c}.$$

This restricts to the standard case (i.e., adding a suffixed subscript s to all three categories). We define $\bar{\mathsf{C}} := \mathsf{C} \circ \bar{\cdot}$ and also write l_B for $\mathsf{l} \circ \mathsf{l}_B$.

The arguments from section 5.2.6 for working with $\mathsf{b}\mathsf{TS}_0$ rather than TS_0 reappear here: The countability of the basis ensures the second-countability of the compactification, rendering it a Polish space. Thus, the ‘canonical’ choice of all (cl)open sets is not available, since this needn’t be countable.²⁶ So if there is no canonical choice available, we better make the choice explicit to obtain a functorial construction.

This leaves open the question of how this relates to the computational compactification. In section 5.4, we show that S and D are indeed functors and, in subsection 5.4.3, we show that the computational compactification SD is naturally isomorphic to the logical compactification $\bar{\mathsf{C}}$.

The practical reason for this twofold approach to compactification is one of compartmentalization: this way the bottom layer of the main diagram is independent of the domain construction. More importantly, though, the conceptual reason is this: It provides two very different, yet equivalent ways of compactifying a zero-dimensional measured topological system. And, it provides two very different topological realizations of a (standard) dynamical system \mathfrak{X} :

- Computational: Choose a countable and separating (mod 0 and forward closed) basis \mathcal{B} for \mathfrak{X} , then \mathfrak{X} is isomorphic to the (standard) dynamical system induced by the compact $\mathsf{SD}(\mathfrak{X}, \mathcal{B})$. The state space of this topological realization consists of the ‘ideal’ elements of the observation domain $\mathsf{D}(\mathfrak{X}, \mathcal{B})$ of the system \mathfrak{X} with respect to the possible observations \mathcal{B} .
- Logical: Using the localization, we may assume (up to isomorphism) that there is a zero-dimensional Polish topology τ on X with countable clopen (forward closed) basis \mathcal{B} that generates the σ -algebra of \mathfrak{X} . Then \mathfrak{X} is isomorphic to the (standard) dynamical system induced by the compact $\bar{\mathsf{C}}(\mathfrak{X}, \mathcal{B})$. The state space of this topological realization consists of the possible models of the properties in (the logical and dynamical closure of) \mathcal{B} that measure the system \mathfrak{X} .

²⁶Since Polish spaces are normal, the choice of all open sets would yield the Stone–Čech compactification (Johnstone 1982, thm. on p. 138f.).

This alignment of domain theory and logic may be a sign of a ‘*dynamical domain theory in logical form*’ à la Abramsky (1991).

Finally, a comment on literature: As mentioned, the Wallman compactification is a standard construction. Specifically, it also is used by Danos and Garnier (2015) and Dahlqvist, Danos, and Garnier (2016) to compactify a zero-dimensional Polish space. (The former also provides, on page 147, some more references on Wallman compactifications, to which we may add Walker (1974) and Engelking (1989, sec. 3.6).) They obtain this compactification by a limit construction of Polish spaces which they then show to correspond to a Wallman compactification. (We, on the other hand, building the Wallman compactification extended by dynamics and then showing that it corresponds to a limit construction of dynamical domains.) The main difference is that we also have to deal with the dynamics, and, as seen in the proof of the limit theorem from the previous chapter, this is the lion’s share of the work. Thus, because of this non-standard addition of dynamics to the Wallman compactification (and also for completeness), we still provide the details of the construction here.

5.3.3 Compactification of a system: formally

In this subsection, we prove theorem 5.3.2: the formal result on compactification motivated in the preceding subsection. We first establish the left half of the diagram of the theorem (in the paragraph ‘closing bases’ below) and then the right half (starting in paragraph ‘compactification’ below).

Closing bases We first show that there is an optimal way of closing a given basis under the logical and dynamical operations, i.e., we have the adjunction:

$$\mathbf{bTS}_0 \begin{array}{c} \xrightarrow{\top} \\ \xleftarrow{\top} \\ \downarrow \end{array} \bar{\mathbf{b}}\mathbf{TS}_0$$

which restricts to the standard case. We first define the construction in a lemma and then show the adjunction in the subsequent proposition.

5.3.3. LEMMA. *Let $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 with $\mathfrak{X} = (X, \tau, \mu, T)$. Define $\bar{\mathcal{B}}$ as the sub-Boolean algebra of $\mathbf{Clp}(X)$ generated by the subset $\bigcup_{k \geq 0} T^{-k}\mathcal{B}$. Then $(\mathfrak{X}, \bar{\mathcal{B}})$ is in \mathbf{bTS}_0 and $\bar{\mathcal{B}}$ is closed under Boolean operations and T -preimages. If $(\mathfrak{X}, \mathcal{B})$ is in \mathbf{bTS}_{0s} , then $(\mathfrak{X}, \bar{\mathcal{B}})$ is in \mathbf{bTS}_{0s} (i.e., $\bar{\mathcal{B}}$ is also closed under T -image).*

PROOF. Write $\mathcal{B}_1 := \bigcup_{k \geq 0} T^{-k}\mathcal{B}$. Note that \mathcal{B}_1 is a countable set that contains \mathcal{B} . Since T is continuous, \mathcal{B}_1 consists of clopen sets. Thus, the sub-Boolean algebra $\bar{\mathcal{B}}$ of $\mathbf{Clp}(X)$ generated by \mathcal{B}_1 exists, and it is countable. It is closed under Boolean operations by construction. It is a basis since it contains the basis \mathcal{B} (and is

closed under finite intersection). So it remains to show that $\overline{\mathcal{B}}$ is closed under T -preimages.

Indeed, let $U \in \overline{\mathcal{B}}$ and show $T^{-1}(U) \in \overline{\mathcal{B}}$. This claim holds by construction if U is among the generators \mathcal{B}_1 of $\overline{\mathcal{B}}$. So it suffices to show that it is preserved by the Boolean operations: Assume the claim holds for $V, W \in \overline{\mathcal{B}}$ and show that (i) if $U = V \cup W \in \overline{\mathcal{B}}$, then $T^{-1}(U) \in \overline{\mathcal{B}}$, and (ii) if $U = V^c \in \overline{\mathcal{B}}$, then $T^{-1}(U) \in \overline{\mathcal{B}}$. (Note that we don't need to consider intersection since it is definable from union and complement.) Concerning (i), we have $T^{-1}(U) = T^{-1}(V) \cup T^{-1}(W)$ which is in $\overline{\mathcal{B}}$ since $T^{-1}(V)$ and $T^{-1}(W)$ are in $\overline{\mathcal{B}}$ by assumption. Concerning (ii), we have $T^{-1}(U) = (T^{-1}(V))^c$ which is in $\overline{\mathcal{B}}$ since $T^{-1}(V)$ is in $\overline{\mathcal{B}}$ by assumption.

Now, assume that $(\mathfrak{X}, \mathcal{B})$ is in \mathbf{bTS}_{0s} , and show that $\overline{\mathcal{B}}$ is also closed under T -image: i.e., if $U \in \overline{\mathcal{B}}$, then $T(U) \in \overline{\mathcal{B}}$. This claim holds if U is among the generators \mathcal{B}_1 of $\overline{\mathcal{B}}$: If $U = T^{-k}(V)$ for $V \in \mathcal{B}$ and $k \geq 0$, then, if $k = 0$, $T(U) = T(V) \in \mathcal{B} \subseteq \overline{\mathcal{B}}$ since \mathcal{B} is closed under T -image, and if $k \geq 1$, then, since T is bijective, $T(U) = T^{-k+1}(V) \in \mathcal{B}_1 \subseteq \overline{\mathcal{B}}$. And the claim is preserved by Boolean operations: Assume the claim holds for $V, W \in \overline{\mathcal{B}}$ and show that (i) if $U = V \cup W \in \overline{\mathcal{B}}$, then $T(U) \in \overline{\mathcal{B}}$, and (ii) if $U = V^c \in \overline{\mathcal{B}}$, then $T(U) \in \overline{\mathcal{B}}$. Concerning (i), we have $T(U) = T(V) \cup T(W)$ which is in $\overline{\mathcal{B}}$ since $T(V)$ and $T(W)$ are in $\overline{\mathcal{B}}$ by assumption. Concerning (ii), we have, since T is bijective, that $T(U) = (T(V))^c$,²⁷ which is in $\overline{\mathcal{B}}$ since $T(V)$ is in $\overline{\mathcal{B}}$ by assumption. \square

5.3.4. PROPOSITION. *The inclusion $\mathbf{l} : \overline{\mathbf{bTS}}_0 \rightarrow \mathbf{bTS}_0$ is a left adjoint functor: For each $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 , $(\mathfrak{X}, \overline{\mathcal{B}})$ is in $\overline{\mathbf{bTS}}_0$ and $\epsilon_{(\mathfrak{X}, \mathcal{B})} := \text{id}_{\mathfrak{X}} : (\mathfrak{X}, \overline{\mathcal{B}}) \rightarrow (\mathfrak{X}, \mathcal{B})$ is a morphism in \mathbf{bTS}_0 such that, for any $(\mathfrak{Y}, \mathcal{C})$ in $\overline{\mathbf{bTS}}_0$ and $\varphi : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{X}, \mathcal{B})$, there is a unique $\psi : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{X}, \overline{\mathcal{B}})$ with $\epsilon_{(\mathfrak{X}, \mathcal{B})} \circ \psi = \varphi$.*

$$\begin{array}{ccc} (\mathfrak{Y}, \mathcal{C}) & & \\ \exists! \psi \downarrow \text{dotted} & \searrow \varphi & \\ (\mathfrak{X}, \overline{\mathcal{B}}) & \xrightarrow{\epsilon_{(\mathfrak{X}, \mathcal{B})}} & (\mathfrak{X}, \mathcal{B}) \end{array}$$

PROOF. In lemma 5.3.3, we've shown that $(\mathfrak{X}, \overline{\mathcal{B}})$ is in $\overline{\mathbf{bTS}}_0$. And $\epsilon_{(\mathfrak{X}, \mathcal{B})} := \text{id}_{\mathfrak{X}} : (\mathfrak{X}, \overline{\mathcal{B}}) \rightarrow (\mathfrak{X}, \mathcal{B})$ is a morphism in \mathbf{bTS}_0 since $\text{id}_{\mathfrak{X}} : \mathfrak{X} \rightarrow \mathfrak{X}$ is a morphism in \mathbf{TS}_0 and it is base-preserving since, for $U \in \mathcal{B}$, $\text{id}_{\mathfrak{X}}^{-1}(U) = U \in \mathcal{B} \subseteq \overline{\mathcal{B}}$.

So let $(\mathfrak{Y}, \mathcal{C})$ be in $\overline{\mathbf{bTS}}_0$ and let $\varphi : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{X}, \mathcal{B})$ be a morphism. We show that $\psi := \varphi : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{X}, \overline{\mathcal{B}})$ is a morphism in $\overline{\mathbf{bTS}}_0$. Then $\epsilon_{(\mathfrak{X}, \mathcal{B})} \circ \psi = \psi = \varphi$, and ψ is unique with this property (if ψ' is another such morphism, then $\psi = \varphi = \epsilon_{(\mathfrak{X}, \mathcal{B})} \circ \psi' = \psi'$).

²⁷For $x \in X$, we have: If $x \in T(U)$, then $x = T(x')$ for $x' \in U = V^c$. If x were in $T(V)$, then $x = T(y)$ for $y \in V$. By injectivity, $x' = y \in V$, contradiction. Conversely, if $x \in (T(V))^c$, then, by surjectivity, there is $x' \in X$ with $T(x') = x$. If x' were not in $U = V^c$, then $x' \in V$, so $x = T(x') \in T(V)$, contradiction. So $x' \in U$ and $x = T(x') \in T(U)$.

Since $\psi = \varphi : \mathfrak{Y} \rightarrow \mathfrak{X}$ is a morphism in \mathbf{TS}_0 , we need to show: if $U \in \overline{\mathcal{B}}$, then $\varphi^{-1}(U) \in \mathcal{C}$. We first show the claim for the generators of $\overline{\mathcal{B}}$ and then that it is preserved under Boolean operations.

So let U be among the generators of $\overline{\mathcal{B}}$. Then $U = T^{-k}(V)$ for $V \in \mathcal{B}$ and $k \geq 0$. So, by equivariance,

$$\varphi^{-1}(U) = \varphi^{-1}(T^{-k}(V)) = (T^k \circ \varphi)^{-1}(V) = (\varphi \circ S^k)^{-1}(V) = S^{-k}(\varphi^{-1}(V))$$

which is in \mathcal{C} since $\varphi^{-1}(V)$ is in \mathcal{C} (since $V \in \mathcal{B}$ and φ is a morphism in \mathbf{bTS}_0) and \mathcal{C} is closed under preimage (since $(\mathfrak{Y}, \mathcal{C})$ is in $\overline{\mathbf{bTS}}_0$).

Now, assume the claim holds for $V, W \in \mathcal{B}$ and show it for $V \cup W$ and V^c . Indeed, we have $\varphi^{-1}(V \cup W) = \varphi^{-1}(V) \cup \varphi^{-1}(W) \in \mathcal{C}$ and $\varphi^{-1}(V^c) = (\varphi(V))^c \in \mathcal{C}$ since $\varphi^{-1}(V), \varphi^{-1}(W) \in \mathcal{C}$ by assumption and \mathcal{C} is closed under Boolean operations. \square

As usual, these data determine the ‘full’ adjunction between \mathbf{bTS}_0 and $\overline{\mathbf{bTS}}_0$, i.e., the functor $\bar{\cdot}$, the counit $\epsilon : \mathbf{1} \rightarrow \mathbf{1}_{\mathbf{bTS}_0}$, etc. (Mac Lane 1998, thm. 2, p. 83). In particular, $\bar{\cdot}$ acts on objects as described in the proposition, and it is the identity on morphisms: If $\varphi : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ is a morphism in \mathbf{bTS}_0 , then the diagram

$$\begin{array}{ccc} (\mathfrak{X}, \overline{\mathcal{B}}) & \xrightarrow{\epsilon_{(\mathfrak{X}, \mathcal{B})}} & (\mathfrak{X}, \mathcal{B}) \\ \bar{\varphi} \downarrow & & \downarrow \varphi \\ (\mathfrak{Y}, \overline{\mathcal{C}}) & \xrightarrow{\epsilon_{(\mathfrak{Y}, \mathcal{C})}} & (\mathfrak{Y}, \mathcal{C}) \end{array}$$

has to commute (by the naturality of the counit), so $\epsilon_{(\mathfrak{Y}, \mathcal{C})} \circ \bar{\varphi} = \varphi \circ \epsilon_{(\mathfrak{X}, \mathcal{B})}$. Since $\epsilon_{(\mathfrak{X}, \mathcal{B})} = \text{id}_{\mathfrak{X}}$, this reduces to $\bar{\varphi} = \varphi$.

Also note that this adjunction restricts to an adjunction

$$\mathbf{bTS}_{0s} \begin{array}{c} \xrightarrow{\bar{\cdot}} \\ \top \\ \xleftarrow{\mathbf{1}} \end{array} \overline{\mathbf{bTS}}_{0s}$$

since, if $(\mathfrak{X}, \mathcal{B})$ is in \mathbf{bTS}_{0s} , then $(\mathfrak{X}, \overline{\mathcal{B}})$ is in $\overline{\mathbf{bTS}}_{0s}$ as shown in lemma 5.3.3.

Compactification After having closed bases under the logical and dynamical operations, we now show that there is an optimal way of rendering the state space compact, i.e., we establish the adjunction

$$\overline{\mathbf{bTS}}_0 \begin{array}{c} \xrightarrow{\mathbf{c}} \\ \perp \\ \xleftarrow{\mathbf{1}_B} \end{array} \mathbf{TS}_{0c}$$

which restricts to the standard case. The first proposition defines the construction and the second proposition establishes the adjunction. The third proposition verifies a further desired property of a compactification: that compactifying a compact space ‘doesn’t do anything’.

First proposition We first recall some basics from Stone duality theory. A standard reference is Johnstone (1982). Given a Boolean algebra B (like the $\bar{\mathcal{B}}$ above), a *filter* on B is a subset $P \subseteq B$ such that: (a) the top element \top of the Boolean algebra is in P , (b) if $y \geq x \in P$, then $y \in P$, and (c) if $x, y \in P$, then $x \wedge y \in P$ (where \wedge is the conjunction of the Boolean algebra). It is an *ultrafilter* if, additionally, (d) if $x \in B$, then exactly one of x and $\neg x$ is in P (where \neg is the negation of the Boolean algebra). The set of all ultrafilters on B is denoted $\text{Spec}(B)$. It carries the topology generated by $D(U) := \{P \in \text{Spec}(B) : U \in P\}$ for $U \in B$ (which forms a basis).²⁸ It is called the *Stone topology* and turns $\text{Spec}(B)$ into a Stone space (i.e., zero-dimensional, compact, Hausdorff).

5.3.5. PROPOSITION. *Let $(\mathfrak{X}, \mathcal{B})$ be in $\bar{\mathbf{b}}\text{TS}_0$ (resp., $\bar{\mathbf{b}}\text{TS}_{0s}$) with $\mathfrak{X} = (X, \tau, \mu, T)$. Let $\eta : X \rightarrow \text{Spec}(\mathcal{B})$ be defined by $\eta(x) := \{U \in \mathcal{B} : x \in U\}$. Define the compactification of \mathfrak{X} as $\mathbf{C}(\mathfrak{X}, \mathcal{B}) := \mathfrak{Y} := (Y, \sigma, \nu, S)$ where*

- $Y := \text{Spec}(\mathcal{B})$ is the set of ultrafilters of the Boolean algebra \mathcal{B}
- σ is the Stone topology on Y
- $\nu : \mathcal{B}(\sigma) \rightarrow [0, 1]$, $\nu(C) := \mu(\eta^{-1}(C))$
- $S : Y \rightarrow Y$, $S(P) := \{U \in \mathcal{B} : T^{-1}(U) \in P\}$.

Then $\mathbf{C}(\mathfrak{X}, \mathcal{B}) = \mathfrak{Y}$ is indeed in TS_{0c} (resp., TS_{0cs}) and $\eta : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \text{Clp}(Y))$ is an injective morphism in $\bar{\mathbf{b}}\text{TS}_0$ (resp., hence also in $\bar{\mathbf{b}}\text{TS}_{0s}$) which is relatively open and has a dense image.

Note that, by lemma 5.2.4, $\eta : (X, \mathcal{B}(\tau), \mu, T) \rightarrow (Y, \mathcal{B}(\sigma), \nu, S)$ is an isomorphism in DS . So, roughly, dynamical systems are invariant under the operation of compactification.

PROOF. We start by checking that $\mathbf{C}(\mathfrak{X}, \mathcal{B})$ is well-defined and in TS_{0c} . Note that, by standard Stone duality, $\eta : X \rightarrow Y$ is a well-defined injective continuous function.²⁹

First, (Y, σ) is a compact zero-dimensional Polish space: Since \mathcal{B} is countable, Y is second-countable. Hence, since it also is compact and Hausdorff (by standard Stone duality), Y is Polish.³⁰

²⁸The collection $C := \{D(U) : U \in B\}$ is closed under finite intersection: It contains $\text{Spec}(B) = D(\top) \in C$, and given $D(U_1), \dots, D(U_n) \in C$ for $U_1, \dots, U_n \in B$, consider $U := U_1 \wedge \dots \wedge U_n \in B$. Then $D(U_1) \cap \dots \cap D(U_n) = D(U) \in C$ because $P \in D(U_1) \cap \dots \cap D(U_n)$ iff $U_1, \dots, U_n \in P$ iff (since P is a filter) $U_1 \wedge \dots \wedge U_n \in P$ iff $P \in D(U)$.

²⁹Well-defined: $\eta(x)$ is readily seen to be an ultrafilter. Injective: If $x \neq y$, there is an open set U of X such that $x \in U$ and $y \notin U$ (X is Hausdorff). Since \mathcal{B} is a basis, we can assume that $U \in \mathcal{B}$. Hence $U \in \eta(x)$ but $U \notin \eta(y)$. So $\eta(x) \neq \eta(y)$. Continuous: For a basic open $D(U)$ of Y (where $U \in \mathcal{B}$), we have $\eta^{-1}(D(U)) = U$ (since $x \in \eta^{-1}(D(U))$ iff $\eta(x) \in D(U)$ iff $U \in \eta(x)$ iff $x \in U$). Since $U \in \mathcal{B}$, U is open, whence $\eta^{-1}(D(U))$ is open in X .

³⁰As a basic fact about Polish spaces, any Hausdorff, second-countable, and locally compact (which is implied by being compact given Hausdorffness) is Polish (see, e.g., Kechris 1995, thm. 5.3, p. 29).

Second, ν is well-defined: Since η in particular is Borel-measurable, ν is a well-defined measure on $\mathcal{B}(\sigma)$.

Third, $S : Y \rightarrow Y$ is well-defined and continuous: Since \mathcal{B} is closed under T -preimage, $F : \mathcal{B} \rightarrow \mathcal{B}$ defined by $F(U) := T^{-1}(U)$ is a well-defined function, and it is a Boolean algebra homomorphism (since preimages commute with complement, union, and intersection). By standard Stone duality, $S : Y \rightarrow Y$ defined by $S(P) := F^{-1}(P) = \{U \in \mathcal{B} : F(U) \in P\} = \{U \in \mathcal{B} : T^{-1}(U) \in P\}$ is well-defined and continuous.

In the standard case, we also need to check that S is measure-preserving and bijective; the latter implies being homeomorphic since Y is compact and Hausdorff and S is continuous.

Injective: If $S(P) = S(Q)$, then, for all $U \in \mathcal{B}$, we have $T^{-1}(U) \in P$ iff $T^{-1}(U) \in Q$. Now, for any $V \in \mathcal{B}$, we have, since $U := T(V) \in \mathcal{B}$, that $V = T^{-1}T(V) = T^{-1}(U) \in P$ iff $V = T^{-1}T(V) = T^{-1}(U) \in Q$. So $P = Q$.

Surjective: Given $P \in Y$, let $Q := \{T^{-1}(U) : U \in P\}$. It is readily seen that this is an ultrafilter on \mathcal{B} and hence in Y .³¹ And we have $S(Q) = P$, since, for $U \in \mathcal{B}$: $U \in S(Q)$ iff $T^{-1}(U) \in Q$ iff $T^{-1}(U) = T^{-1}(V)$ for some $V \in P$ iff $U \in P$, where the non-trivial \Rightarrow -direction of the last equivalence follows since T is bijective: then $U = TT^{-1}(U) = TT^{-1}(V) = V \in P$.

Measure-preserving: To show that the Borel probability measures ν and νS^{-1} are identical on $\mathcal{B}(\sigma)$, it suffices to show that they are identical on the class \mathcal{E} of basic opens $D(U)$ (with $U \in \mathcal{B}$): then the probability measures agree on \mathcal{E} which is closed under finite intersections, so they agree on the σ -algebra generated by \mathcal{E} (Bogachev 2007b, lem. 1.9.4, p. 35), which is $\mathcal{B}(\sigma)$ since σ is second-countable. So let $D(U)$ be a basic open set of Y (for $U \in \mathcal{B}$). Then $P \in S^{-1}(D(U))$ iff $S(P) \in D(U)$ iff $U \in S(P)$ iff $T^{-1}(U) \in P$ iff $P \in D(T^{-1}(U))$. So, since T is measure-preserving and $\eta^{-1}(D(U)) = U$ for $U \in \mathcal{B}$, we have

$$\begin{aligned} \nu\left(S^{-1}(D(U))\right) &= \nu\left(D(T^{-1}(U))\right) = \mu\eta^{-1}\left(D(T^{-1}(U))\right) = \mu(T^{-1}(U)) \\ &= \mu(U) = \mu\eta^{-1}(D(U)) = \nu(D(U)), \end{aligned}$$

as needed.

Next, we check that $\eta : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \text{Clp}(Y))$ is a morphism in $\bar{\mathbf{b}}\mathbf{TS}_0$ (we already know that it is injective), relatively open, and has a dense image.

Base-preserving (and hence continuous): If $V \subseteq Y$ is clopen, then $V = D(U)$ for some $U \in \mathcal{B}$.³² So $\eta^{-1}(V) = \eta^{-1}(D(U)) = U \in \mathcal{B}$.

³¹(a) The top element X of \mathcal{B} is in P , so $X = T^{-1}(X) \in Q$. (b) Given $T^{-1}(U)$ for $U \in P$ and $T^{-1}(U) \subseteq V$ for $V \in \mathcal{B}$, we have, since T is bijective, $U = TT^{-1}(U) \subseteq T(V) \in \mathcal{B}$, so $T(V) \in P$, whence $V = T^{-1}(T(V)) \in Q$. (c) Given $T^{-1}(U)$ and $T^{-1}(V)$ for $U, V \in P$, we have $U \cap V \in P$, whence, by injectivity of T , $T^{-1}(U) \cap T^{-1}(V) = T^{-1}(U \cap V) \in Q$. (d) Given $U \in \mathcal{B}$, we have $V := T(U) \in \mathcal{B}$, so exactly one of V and V^c is in P . If $V \in P$, then $U = T^{-1}(V) \in Q$. If $V^c \in P$, then $U^c = T^{-1}(V^c) = T^{-1}(V^c) \in Q$. And we cannot have U and U^c in Q , since otherwise $\emptyset \in Q$, so $\emptyset = T^{-1}(U)$ for $U \in P$, whence, by surjectivity of T , $U = \emptyset$, but $\emptyset \notin P$.

³²Proof: Since the $D(U)$ form a basis, $V = \bigcup_{i \in I} D(U_i)$. Since V is a closed subset of a

$$\begin{array}{ccc}
(\mathfrak{X}, \mathcal{B}) & \xleftarrow{\eta} & (\mathfrak{Y}, \text{Clp}(Y)) = \mathbb{I}_B(\mathbb{C}(\mathfrak{X}, \mathcal{B})) \\
& \searrow \psi & \downarrow \exists! \varphi (= \mathbb{I}_B(\varphi)) \\
& & (\mathfrak{Z}, \text{Clp}(Z)) = \mathbb{I}_B(\mathfrak{Z})
\end{array}$$

Figure 5.3: The compactification functor as left adjoint to \mathbb{I}_B .

Measure-preserving: By construction, since ν is the pushforward measure of μ .
Equivariant: Let $x \in X$ and show $\eta T(x) = S\eta(x)$. Indeed, for $U \in \mathcal{B}$, we have

$$\begin{aligned}
U \in \eta T(x) &\Leftrightarrow T(x) \in U \Leftrightarrow x \in T^{-1}(U) \Leftrightarrow x \in F(U) \\
&\Leftrightarrow F(U) \in \eta(x) \Leftrightarrow U \in F^{-1}(\eta(x)) = S\eta(x).
\end{aligned}$$

as needed.

Relatively open: For $U \subseteq X$ is open, we have to find an open $V \subseteq Y$ such that $\eta(U) = V \cap \eta(X)$. Indeed, since \mathcal{B} is a basis for X , we have $U = \bigcup_I U_i$ for $U_i \in \mathcal{B}$. We first show that $\eta(U_i) = D(U_i) \cap \eta(X)$: If $P \in D(U_i) \cap \eta(X)$, then $P = \eta(x)$ for some $x \in X$ and $U_i \in P = \eta(x)$. So $x \in U_i$. Hence $P = \eta(x) \in \eta(U_i)$. Conversely, if $P \in \eta(U_i)$, then $P = \eta(x)$ for some $x \in U_i$, so $P \in \eta(X)$ and, since $x \in U_i$, we have $U_i \in \eta(x) = P$, so $P \in D(U_i)$. Now the claim follows: Since the image of a union is the union of the images, $\eta(U) = \bigcup_I \eta(U_i) = \bigcup_I D(U_i) \cap \eta(X)$ and $V := \bigcup_I D(U_i) \subseteq Y$ is open.

Dense image: Since the $D(U)$ (with $U \in \mathcal{B}$) are a basis for Y , it suffices to show that for every nonempty $D(U)$, we have $D(U) \cap \eta(X) \neq \emptyset$. Indeed, let $P \in D(U)$, so $U \in P$. Since P is an ultrafilter, $U \neq \emptyset$, so let $x \in U$. Then $Q := \eta(x) = \{U \in \mathcal{B} : x \in U\}$ is in Y with $U \in Q$, so $Q \in D(U)$ and $Q \in \eta(X)$, so $D(U) \cap \eta(X) \neq \emptyset$. \square

The proof would almost go through for any $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 using $\overline{\mathcal{B}}$ instead of \mathcal{B} . But base-preservation poses a problem: then we have $\eta^{-1}(V) = \eta^{-1}(D(U)) = U \in \overline{\mathcal{B}}$ but we need it to be an element of \mathcal{B} .

Second proposition To show that this compactification construction indeed forms a left adjoint to $\mathbb{I}_B : \mathbf{TS}_{0c} \rightarrow \mathbf{bTS}_0$, we need to show the following—as depicted in figure 5.3.

compact space, it is compact, so, for some n , $V = D(U_{i_1}) \cup \dots \cup D(U_{i_n})$. So $V = D(U)$ for $U := U_{i_1} \cup \dots \cup U_{i_n} \in \mathcal{B}$. (If $P \in V$, then $P \in D(U_{i_k})$ for some $k = 1, \dots, n$, whence $U \supseteq U_{i_k} \in P$, so, since P is a filter, $U \in P$, whence $P \in D(U)$. Conversely, if $P \in D(U)$ but $P \notin V$, then no U_{i_k} is in P , so, since P is an ultrafilter, all $U_{i_k}^c$ are in P , whence, since P is a filter, $U^c = \bigcap_{k=1}^n U_{i_k}^c \in P$, contradicting $U \in P$.)

5.3.6. PROPOSITION. *The functor $\mathsf{l}_B : \mathsf{TS}_{0c} \rightarrow \overline{\mathsf{bTS}}_0$ is a right adjoint functor: Let $(\mathfrak{X}, \mathcal{B})$ be in $\overline{\mathsf{bTS}}_0$. Let $\mathfrak{Y} = \mathsf{C}(\mathfrak{X}, \mathcal{B})$ be the object in TS_{0c} and $\eta = \eta_{(\mathfrak{X}, \mathcal{B})} : (\mathfrak{X}, \mathcal{B}) \rightarrow \mathsf{l}_B(\mathfrak{Y})$ the injective morphism in $\overline{\mathsf{bTS}}_0$ constructed in proposition 5.3.5. Then, for every \mathfrak{Z} in TS_{0c} and $\psi : (\mathfrak{X}, \mathcal{B}) \rightarrow \mathsf{l}_B(\mathfrak{Z})$ in $\overline{\mathsf{bTS}}_0$, there is a unique morphism $\varphi : \mathfrak{Y} \rightarrow \mathfrak{Z}$ such that $\mathsf{l}_B(\varphi) \circ \eta = \psi$.*

PROOF. We fix the notation $\mathfrak{X} = (X, \tau, \mu, T)$, $\mathfrak{Y} = (Y, \sigma, \nu, S)$, and $\mathfrak{Z} = (Z, \rho, \lambda, R)$. We first show existence of φ . This is the main part of the proof. Afterward, uniqueness follows straightforwardly.

Note that, since $\eta : X \rightarrow Y$ is injective, the following is a well-defined function:

$$\begin{aligned} \hat{\varphi} : \eta(X) &\rightarrow Z \\ \eta(x) &\mapsto \psi(x). \end{aligned}$$

It also is continuous (where $\eta(X)$ gets the relative topology from Y): If $U \subseteq Z$ is open, then

$$\hat{\varphi}^{-1}(U) = \{\eta(x) : x \in X \text{ and } \psi(x) \in U\} = \eta(\psi^{-1}(U)).$$

So, since ψ is continuous and η is relatively open, $\eta(\psi^{-1}(U))$ is open in $\eta(X)$ with the relative topology.

Now, we want to define $\varphi : Y \rightarrow Z$ as a continuous extension of $\hat{\varphi} : \eta(X) \rightarrow Z$. Note that such an extension is unique: If φ' were another, then, given $y \in Y$, we can write $y = \lim y_n$ for $y_n \in \eta(X)$ since $\eta(X)$ is dense in Y , so, by continuity, $\varphi(y) = \lim \varphi(y_n) = \lim \hat{\varphi}(y_n) = \lim \varphi'(y_n) = \varphi'(y)$.

To show that such an extension exists, we apply the characterization of extendability of mappings into compact spaces (Engelking 1989, thm. 3.2.1):³³ the continuous function $\hat{\varphi} : \eta(X) \rightarrow Z$ from the dense subspace $\eta(X)$ of the space Y to the compact space Z has a continuous extension $\varphi : Y \rightarrow Z$ iff for any two disjoint closed subsets C and C' of Z , the sets $\hat{\varphi}^{-1}(C)$ and $\hat{\varphi}^{-1}(C')$ have disjoint closures in Y .

So let C, C' be closed subsets of Z , and show that $\hat{\varphi}^{-1}(C)$ and $\hat{\varphi}^{-1}(C')$ have disjoint closures in Y . Since C is closed and $\mathsf{Clp}(Z)$ a basis, we can write $C = \bigcap_{i \in I} U_i$ for $U_i \in \mathsf{Clp}(Z)$. Without loss of generality, $I \neq \emptyset$.³⁴ Since $\psi : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Z}, \mathsf{Clp}(Z))$ is base-preserving, $\psi^{-1}(U_i) \in \mathcal{B}$. So

$$\hat{\varphi}^{-1}(C) = \eta(\psi^{-1}C) = \eta(\psi^{-1} \bigcap_I U_i) = \eta(\bigcap_I \psi^{-1}U_i).$$

³³As noted by Engelking (1989, p. 145), this theorem goes back to Taïmanov and to Eilenberg and Steenrod. It also is the key to continuously extend a continuous function $f : X \rightarrow Z$, with X a T_1 space and Z compact, to the Wallman extension of X (Engelking 1989, thm. 3.6.21).

³⁴If $I = \emptyset$, then $C = Z$, so we can write $C = \bigcap_{i \in I'} U_i$ with $I' = \{i\}$ a singleton set and $U_i := Z \in \mathsf{Clp}(Z)$.

Similarly, we can write $C' = \bigcap_{j \in J} U'_j$ for $U'_j \in \mathbf{Clp}(Z)$ and $J \neq \emptyset$ with $\psi^{-1}(U'_j) \in \mathcal{B}$ and $\hat{\varphi}^{-1}(C') = \eta(\bigcap_J \psi^{-1}U'_j)$. Hence $A := \bigcap_I D(\psi^{-1}U_i)$ and $A' := \bigcap_J D(\psi^{-1}U'_j)$ are closed subsets of Y . And we have $\hat{\varphi}^{-1}(C) \subseteq A$: If $P \in \hat{\varphi}^{-1}(C)$, then $P = \eta(x)$ for $x \in \bigcap_I \psi^{-1}U_i$, so, for any $i \in I$, we have $x \in \psi^{-1}U_i$, so $\psi^{-1}U_i \in \eta(x)$, so $\eta(x) \in D(\psi^{-1}U_i)$, whence $P = \eta(x) \in A$. Similarly, $\hat{\varphi}^{-1}(C') \subseteq A'$.

Thus, it suffices to show that $A \cap A' = \emptyset$ (then the closures of $\hat{\varphi}^{-1}(C)$ and $\hat{\varphi}^{-1}(C')$ are contained in the closed sets A and A' , respectively, and hence are disjoint). Assume for contradiction that there is $P \in A \cap A'$. We argue that $F := \{U_i \cap U'_j : i \in I, j \in J\}$ has the finite intersection property. Since this is a family of closed subsets of the compact space Z , this implies that it has a nonempty intersection, whence we'll get the contradiction $\emptyset \neq \bigcap_{i \in I, j \in J} U_i \cap U'_j = \bigcap_I U_i \cap \bigcap_J U'_j = C \cap C' = \emptyset$.³⁵ So let $U_{i_1} \cap U'_{j_1}, \dots, U_{i_n} \cap U'_{j_n}$ be given and show that their intersection is nonempty. Since $P \in A$, we have in particular that $P \in D(\psi^{-1}U_{i_1}) \cap \dots \cap D(\psi^{-1}U_{i_n})$, so $\psi^{-1}U_{i_1}, \dots, \psi^{-1}U_{i_n} \in P$. Similarly, since $P \in A'$, $\psi^{-1}U'_{j_1}, \dots, \psi^{-1}U'_{j_n} \in P$. Since P is a proper filter,

$$\emptyset \neq \psi^{-1}U_{i_1} \cap \psi^{-1}U'_{j_1} \cap \dots \cap \psi^{-1}U_{i_n} \cap \psi^{-1}U'_{j_n} \in P.$$

So there is $x \in X$ with $\psi(x) \in U_{i_1} \cap U'_{j_1} \cap \dots \cap U_{i_n} \cap U'_{j_n}$, as needed.

Now, we have the continuous function $\varphi : Y \rightarrow Z$. So it remains to show that it is a morphism $\mathfrak{Y} \rightarrow \mathfrak{Z}$ and $\varphi \circ \eta = \psi$ (note that $\mathbf{l}_B(\varphi) = \varphi$). Regarding the latter, we have, for $x \in X$, that, by construction, $\varphi(\eta(x)) = \hat{\varphi}(\eta(x)) = \psi(x)$. Since φ is continuous, it remains to show that φ is measure-preserving and equivariant.

Measure-preserving: For $E \in \mathcal{B}(Z)$ we have, since η and ψ are measure-preserving, that

$$\nu(\varphi^{-1}(E)) = \mu(\eta^{-1}(\varphi^{-1}(E))) = \mu((\varphi \circ \eta)^{-1}(E)) = \mu((\psi)^{-1}(E)) = \lambda(E).$$

Equivariance: Let $y \in Y$ and show $\varphi S(y) = R\varphi(y)$. Let's first assume $y \in \eta(X)$, so $y = \eta(x)$ for some $x \in X$. Then we have, since η and ψ are equivariant,

$$\varphi S(y) = \varphi S\eta(x) = \varphi\eta T(x) = \psi T(x) = R\psi(x) = R\varphi\eta(x) = R\varphi(y).$$

If $y \notin \eta(X)$, then, since $\eta(X) \subseteq Y$ is dense, $y = \lim_n y_n$ for $y_n \in \eta(X)$. Hence, by the above and the continuity of the functions,

$$\varphi S(y) = \varphi S(\lim_n y_n) = \lim_n \varphi S(y_n) = \lim_n R\varphi(y_n) = R\varphi(\lim_n y_n) = R\varphi(y).$$

This concludes the existence of φ . It remains to check uniqueness. Assume $\varphi' : \mathfrak{Y} \rightarrow \mathfrak{Z}$ is another morphism with $\varphi' \circ \eta = \psi$ (again, note that $\mathbf{l}_B(\varphi') = \varphi'$). Then $\varphi' : Y \rightarrow Z$ is a continuous function that extends $\hat{\varphi}$ (because for $\eta(x)$ in

³⁵For the first equality, we use that $I, J \neq \emptyset$.

$\eta(X)$ we have $\varphi'\eta(x) = \psi(x) = \varphi\eta(x) = \hat{\varphi}\eta(x)$). Since φ is, as noted above, the unique such function, $\varphi' = \varphi$. \square

Again, this determines the functor $C : \bar{\mathbf{bTS}}_0 \rightarrow \mathbf{TS}_{0c}$ (Mac Lane 1998, thm. 2, p. 83). On objects, it is described in the proposition, and if $\varphi : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ is a morphism in $\bar{\mathbf{bTS}}_0$, then $C(\varphi) : C(\mathfrak{X}, \mathcal{B}) \rightarrow C(\mathfrak{Y}, \mathcal{C})$ is the unique morphism obtained from factoring $\eta_{(\mathfrak{Y}, \mathcal{C})} \circ \varphi$ through $\eta_{(\mathfrak{X}, \mathcal{B})}$:

$$\begin{array}{ccc} (\mathfrak{X}, \mathcal{B}) & \xrightarrow{\eta_{(\mathfrak{X}, \mathcal{B})}} & l_B C(\mathfrak{X}, \mathcal{B}) & & C(\mathfrak{X}, \mathcal{B}) \\ \varphi \downarrow & & \downarrow C(\varphi) = l_B C(\varphi) & & \downarrow C(\varphi) \\ (\mathfrak{Y}, \mathcal{C}) & \xrightarrow{\eta_{(\mathfrak{Y}, \mathcal{C})}} & l_B C(\mathfrak{Y}, \mathcal{C}) & & C(\mathfrak{Y}, \mathcal{C}) \end{array}$$

which makes the square on the left commute.

In particular, we have the functor $\bar{C} := C \circ \bar{\cdot} : \bar{\mathbf{bTS}}_0 \rightarrow \mathbf{TS}_{0c}$ mapping $(\mathfrak{X}, \mathcal{B})$ to $\bar{C}(\mathfrak{X}, \mathcal{B}) = C(\mathfrak{X}, \bar{\mathcal{B}})$ and mapping $\varphi : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ to $\bar{C}(\varphi) = C(\bar{\varphi}) = C(\varphi)$.

Moreover, this adjunction restricts to an adjunction

$$\bar{\mathbf{bTS}}_{0s} \begin{array}{c} \xrightarrow{C} \\ \perp \\ \xleftarrow{l_B} \end{array} \mathbf{TS}_{0cs}$$

since l_B maps, as noted in proposition 5.3.1, objects from \mathbf{TS}_{0cs} to objects in $\bar{\mathbf{bTS}}_{0s}$, and, as noted in proposition 5.3.5, if $(\mathfrak{X}, \mathcal{B})$ is in $\bar{\mathbf{bTS}}_{0s}$, then $C(\mathfrak{X}, \mathcal{B})$ is in \mathbf{TS}_{0cs} .

Third proposition We also note that, as we would expect of a compactification, if we start with a compact \mathfrak{X} , the compactification ‘doesn’t do anything’:

5.3.7. PROPOSITION. 1. If \mathfrak{X} is in \mathbf{TS}_{0c} , then $\eta_{\mathfrak{X}} := \eta_{(\mathfrak{X}, \text{Clp}(X))}$ regarded as morphism $\mathfrak{X} \rightarrow C(\mathfrak{X}, \text{Clp}(X))$, is an isomorphism in \mathbf{TS}_{0c} .

2. The family $\{\eta_{\mathfrak{X}} : \mathfrak{X} \in \mathbf{TS}_{0c}\}$ is a natural isomorphism $1_{\mathbf{TS}_{0c}} \rightarrow \text{Cl}_B$.

3. The family $\{\eta_{\mathfrak{X}} : \mathfrak{X} \in \mathbf{TS}_{0c}\}$ also is a natural isomorphism $1_{\mathbf{TS}_{0c}} \rightarrow \bar{\text{Cl}}_B$.

PROOF. Ad (1). By construction, $\eta_{(\mathfrak{X}, \text{Clp}(X))}$ is a morphism $(\mathfrak{X}, \text{Clp}(X)) \rightarrow l_B C(\mathfrak{X}, \text{Clp}(X))$ in $\bar{\mathbf{bTS}}_0$, so it is a morphism $\mathfrak{X} \rightarrow C(\mathfrak{X}, \text{Clp}(X))$ in \mathbf{TS}_0 that also preserves bases. In particular, since both objects are compact, it is a morphism in \mathbf{TS}_{0c} .

It suffices to show that $\eta_{\mathfrak{X}}$ is surjective: Then it is a bijective continuous map between compact Hausdorff spaces and hence a homeomorphism, so, since it is a morphism in \mathbf{TS}_{0c} , it is an isomorphism by proposition 5.2.8.

Indeed, the usual argument from Stone duality works: Let P be an element of the state space $\text{Spec}(\text{Clp}(X))$ of $C(\mathfrak{X}, \text{Clp}(X))$. Then P is, qua filter of $\text{Clp}(X)$

that doesn't contain the empty set, a family of closed subset of the compact space X with the finite intersection property, whence there is $x \in \bigcap_{U \in P} U$. Note that, for $U \in \text{Clp}(X)$, we have $x \in U$ iff $U \in P$: If $U \in P$, then $x \in \bigcap_{U \in P} U \subseteq U$; and conversely, if $x \in U$, but $U \notin P$, then, qua ultrafilter, $U^c \in P$, so $x \in \bigcap_{U \in P} U \subseteq U^c$, contradiction. Hence $\eta_{\mathfrak{X}}(x) = P$ since, for $U \in \text{Clp}(X)$, we have $U \in \eta_{\mathfrak{X}}(x)$ iff $x \in U$ iff $U \in P$.

Ad (2). As just seen, $\eta_{\mathfrak{X}} : \mathfrak{X} \rightarrow \text{Cl}_B(\mathfrak{X})$ is an isomorphism in TS_{0c} , so we need to show that it is natural in \mathfrak{X} . So let $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ be in TS_{0c} and show that the diagram

$$\begin{array}{ccc} \mathfrak{X} & \xrightarrow{\eta_{\mathfrak{X}}} & \text{C}(\mathfrak{X}, \text{Clp}(X)) = \text{Cl}_B(\mathfrak{X}) \\ \varphi \downarrow & & \downarrow \text{C}(\varphi) = \text{Cl}_B(\varphi) \\ \mathfrak{Y} & \xrightarrow{\eta_{\mathfrak{Y}}} & \text{C}(\mathfrak{Y}, \text{Clp}(Y)) = \text{Cl}_B(\mathfrak{Y}) \end{array}$$

commutes. Indeed, we've noted above that $\text{C}(\varphi) \circ \eta_{(\mathfrak{X}, \text{Clp}(\mathfrak{X}))} = \eta_{(\mathfrak{Y}, \text{Clp}(\mathfrak{Y}))} \circ \varphi$, as needed.

Ad (3). Note that, for $\mathfrak{X} = (X, \tau, \mu, T)$ in TS_{0c} , $(\mathfrak{X}, \mathcal{B})$ with $\mathcal{B} := \text{Clp}(X)$ is in bTS_0 and $\overline{\text{Clp}(X)} = \text{Clp}(X)$ (since $\text{Clp}(X)$ is the sub-Boolean algebra of $\text{Clp}(X)$ generated by $\bigcup_{k \geq 0} T^{-k}\mathcal{B}$). Hence, by (1),

$$\eta_{\mathfrak{X}} : \mathfrak{X} \rightarrow \text{C}(\mathfrak{X}, \text{Clp}(X)) = \text{C}(\mathfrak{X}, \overline{\text{Clp}(X)}) = \overline{\text{C}}(\mathfrak{X}, \text{Clp}(X))$$

is an isomorphism in TS_{0c} . It is natural in \mathfrak{X} by via the same diagram as above since $\overline{\text{Cl}_B(\varphi)} = \overline{\text{C}}(\varphi) = \text{C}(\varphi)$. \square

5.4 The system and domain functors

In this section, we establish the following part of the main diagram:

$$\begin{array}{ccc} \text{dDOM} & & \\ \uparrow \hat{S} & \swarrow \hat{D} & \\ \text{bTS}_0 & \xrightarrow{\text{S}} & \text{TS}_{0c} \\ \downarrow \text{D} & \xleftarrow{\text{I}_B} & \downarrow \overline{\text{C}} \end{array}$$

In subsections 5.4.1 and 5.4.2, we construct the functors S and D , respectively. (We then define $\hat{\text{S}} := \text{I}_B \circ \text{S}$ and $\hat{\text{D}} := \text{D} \circ \text{I}_B$.) In subsection 5.4.3, we establish a main ingredient to showing that the diagram commutes: that the two compactification functors $\text{S} \circ \text{D}$ and $\overline{\text{C}}$ are naturally isomorphic. (The complete proof of commutativity and the standard case are in section 5.7.)

Those who follow the minimal reading and skipped the previous section only need to know that $\mathsf{l}_B : \mathsf{TS}_{0c} \rightarrow \mathsf{bTS}_0$ is the functor sending \mathfrak{X} to $(\mathfrak{X}, \mathsf{Clp}(X))$ and that is the identity on morphisms (established in proposition 5.3.1). They can ignore the functor $\overline{\mathsf{C}} : \mathsf{bTS}_0 \rightarrow \mathsf{TS}_{0c}$ and skip subsection 5.4.3 below.

5.4.1 The system functor

We define the *system functor* $\mathsf{S} : \mathsf{dDOM} \rightarrow \mathsf{TS}_{0c}$ and its ‘extension’ $\hat{\mathsf{S}} := \mathsf{l}_B \circ \mathsf{S}$.

5.4.1. PROPOSITION. *The following defines a functor $\mathsf{S} : \mathsf{dDOM} \rightarrow \mathsf{TS}_{0c}$:*

- For \mathfrak{D} in dDOM , let $\mathsf{S}(\mathfrak{D})$ be the (compact zero-dimensional measured) topological system modeled by \mathfrak{D} (recalled in section 5.2.5).
- For $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ in dDOM , let $\mathsf{S}(\alpha) := \alpha \upharpoonright \max D : \mathsf{S}(\mathfrak{D}) \rightarrow \mathsf{S}(\mathfrak{E})$.

We also define $\hat{\mathsf{S}} := \mathsf{l}_B \circ \mathsf{S} : \mathsf{dDOM} \rightarrow \mathsf{bTS}_0$.

PROOF. We first verify that this is well-defined: As noted in section 5.2.5, $\mathsf{S}(\mathfrak{D})$ is an object in TS_{0c} . So we need to show that, for $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ in dDOM , the function $\mathsf{S}(\alpha) := \alpha \upharpoonright \max D : \mathsf{S}(\mathfrak{D}) \rightarrow \mathsf{S}(\mathfrak{E})$ is indeed a morphism in TS_{0c} , where we write $\mathfrak{D} = (D, v, f)$ and $\mathfrak{E} = (E, w, g)$.

Since α is, qua dynamical morphism, max-preserving and Scott-continuous, $\alpha \upharpoonright \max D : \max D \rightarrow \max E$ is a well-defined function which is continuous on the relative Scott topology (since $\alpha \upharpoonright \max D$ is the restriction of the continuous $\alpha : D \rightarrow E$).³⁶ Since the morphisms in dDOM are max-equivariant,³⁷ we have, for $a \in \max D$, that $(\alpha \upharpoonright \max D)(f \upharpoonright \max D)(a) = \alpha f(a) = g\alpha(a) = (g \upharpoonright \max E)(\alpha \upharpoonright \max D)(a)$. So it remains to show that $\alpha \upharpoonright \max D$ is measure-preserving: For the open sets $V \cap \max E$ of $\max E$ (where $V \subseteq E$ is Scott-open) we have, since α is valuation-preserving,

$$\begin{aligned} \mu_v \upharpoonright \mathcal{B}(\max D)((\alpha \upharpoonright \max D)^{-1}(V \cap \max E)) &= \mu_v(\alpha^{-1}(V) \cap \max D) \\ &= \mu_v(\alpha^{-1}(V)) = v(\alpha^{-1}(V)) = w(V) = \mu_w(V) \\ &= \mu_w(V \cap \max E) = \mu_w \upharpoonright \mathcal{B}(\max E)(V \cap \max E). \end{aligned}$$

Now, the Borel probability measures $\mu_w \upharpoonright \mathcal{B}(\max E)$ and $\mu_v \upharpoonright \mathcal{B}(\max D)(\alpha \upharpoonright \max D)^{-1}$ on $\mathcal{B}(\max E)$ agree on the open sets of $\max E$, and hence on all of $\mathcal{B}(\max E)$ (see e.g. Bogachev 2007a, lem. 7.1.2, p. 68).

Now, we verify that the functor conditions are satisfied: Regarding identity, $\mathsf{S}(\mathsf{id}_{\mathfrak{D}}) = \mathsf{id}_D \upharpoonright \max D = \mathsf{id}_{\max D} = \mathsf{id}_{\mathsf{S}(\mathfrak{D})}$. Regarding composition, if $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$

³⁶If $V \cap \max E$ is open in $\max E$, then $\alpha \upharpoonright \max D^{-1}(V \cap \max E) = \alpha^{-1}(V) \cap \max D$, which is open in $\max D$.

³⁷The fact that they are max-semi-equivariant already implies that they are max-equivariant since the domain dynamics are max-preserving.

Object in \mathbf{bTS}_0	Limit	of diagram	with
$(\mathfrak{X}, \mathcal{B}_X), \mathfrak{X} = (X, \tau, \mu, T)$	$D(\mathfrak{X}, \mathcal{B}_X) = (D, v, f)$	$(\mathfrak{D}_i, p_{ij}^D)_{I(\mathcal{B}_X)}$	$\mathfrak{D}_i = (D_i, v_i, f_i)$
$(\mathfrak{Y}, \mathcal{B}_Y), \mathfrak{Y} = (Y, \sigma, \nu, S)$	$D(\mathfrak{Y}, \mathcal{B}_Y) = (E, w, g)$	$(\mathfrak{E}_i, p_{ij}^E)_{I(\mathcal{B}_Y)}$	$\mathfrak{E}_i = (E_i, w_i, g_i)$
$(\mathfrak{Z}, \mathcal{B}_Z), \mathfrak{Z} = (Z, \rho, \lambda, R)$	$D(\mathfrak{Z}, \mathcal{B}_Z) = (F, u, h)$	$(\mathfrak{F}_i, p_{ij}^F)_{I(\mathcal{B}_Z)}$	$\mathfrak{F}_i = (F_i, u_i, h_i)$

If clear from context, we drop the subscript from \mathcal{B}_X and the superscripts from $p_{ij}^D, p_{ij}^E, p_{ij}^F$.

Figure 5.4: Notational conventions of this subsection.

and $\beta : \mathfrak{E} \rightarrow \mathfrak{F}$ are morphisms in \mathbf{dDOM} , then $S(\beta \circ \alpha) = (\beta \circ \alpha) \upharpoonright \max D = (\beta \upharpoonright \max E) \circ (\alpha \upharpoonright \max D) = S(\beta) \circ S(\alpha)$. \square

Note that S restricts to $S : \mathbf{dDOM}_s \rightarrow \mathbf{TS}_{0cs}$ since, as noted in section 5.2.5, if \mathfrak{D} is standard, then $S(\mathfrak{D})$ is standard.

5.4.2 The domain functor

We define the *domain functor* $D : \mathbf{bTS}_0 \rightarrow \mathbf{dDOM}$ and its ‘restriction’ $\hat{D} := D \circ I_B : \mathbf{TS}_{0c} \rightarrow \mathbf{dDOM}$.

On objects, we use the observation domain construction recalled in section 5.2.5: for $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 , $D(\mathfrak{X}, \mathcal{B})$ is the observation domain of \mathfrak{X} with respect to \mathcal{B} . Thus, the main task is to extend this construction to morphisms. This is done in the following proposition. To avoid redefining notion over and over again, we fix some notational conventions summarized in figure 5.4.

5.4.2. PROPOSITION. *Let $\varphi : (\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Y}, \mathcal{B}_Y)$ be a morphism in \mathbf{bTS}_0 . Then we have a morphism $\alpha : D(\mathfrak{X}, \mathcal{B}_X) \rightarrow D(\mathfrak{Y}, \mathcal{B}_Y)$ in \mathbf{dDOM} defined by*

$$\alpha(a) := \left\langle \pi_{(m, \mathcal{D})}(a(m, \varphi^{-1}\mathcal{D})) : (m, \mathcal{D}) \in I(\mathcal{B}_Y) \right\rangle,$$

where $\pi_{(m, \mathcal{D})} : \mathbf{P}(\mathbf{H}_{(m, \varphi^{-1}\mathcal{D})}) \rightarrow \mathbf{P}(\mathbf{H}_{(m, \mathcal{D})})$ is given by

$$M \mapsto \left\{ \mathcal{O}_{\mathcal{D}}^m(\varphi(x)) \in \mathbf{H}_{(m, \mathcal{D})} : \mathcal{O}_{\varphi^{-1}\mathcal{D}}^m(x) \in M \right\}.$$

The proof is rather long due to the number details to be checked. But this shouldn’t obscure that the main idea of the proof is rather simple: Show that

$$\left(D(\mathfrak{X}, \mathcal{B}_X), \pi_{(m, \mathcal{D})} \circ p_{(m, \varphi^{-1}\mathcal{D})} \right)_{I(\mathcal{B}_Y)}$$

is a cone to the diagram over which $D(\mathfrak{Y}, \mathcal{B}_Y)$ is obtained as restricted limit, and then obtain α as the mediating morphism $D(\mathfrak{X}, \mathcal{B}_X) \rightarrow D(\mathfrak{Y}, \mathcal{B}_Y)$. (Although

there is a complication to this idea which we discuss and circumvent in the proof.) The idea behind $\pi_{(m, \mathcal{D})}$ is that φ plays two roles: In the backward direction, it maps a \mathcal{B}_Y -cover \mathcal{D} to a \mathcal{B}_X -cover $\varphi^{-1}\mathcal{D}$, and hence an element $(m, \mathcal{D}) \in I(\mathcal{B}_Y)$ to $(m, \varphi^{-1}\mathcal{D}) \in I(\mathcal{B}_X)$. In the forward direction, it then maps an observation history $\mathcal{O}_{\varphi^{-1}\mathcal{D}}^m(x)$ over X to an observation history $\mathcal{O}_{\mathcal{D}}^m(\varphi(x))$ over Y , and hence induces a mapping $\mathbf{P}(\mathbf{H}_{(m, \varphi^{-1}\mathcal{D})}) \rightarrow \mathbf{P}(\mathbf{H}_{(m, \mathcal{D})})$.

PROOF. *Step 1.* The map

$$\begin{aligned} \xi : I(\mathcal{B}_Y) &\rightarrow I(\mathcal{B}_X) \\ (m, \mathcal{D}) &\mapsto (m, \varphi^{-1}\mathcal{D}), \end{aligned}$$

is well-defined and monotone.

Well-defined: To show that $(m, \varphi^{-1}\mathcal{D}) \in I(\mathcal{B}_X)$, we need to show that $\varphi^{-1}\mathcal{D} = \{\varphi^{-1}(V) : V \in \mathcal{D}\}$ is indeed a finite \mathcal{B}_X -cover of X : It is finite since \mathcal{D} is finite, its elements $\varphi^{-1}(V)$ are in \mathcal{B}_X since $\varphi^{-1}(\mathcal{B}_Y) \subseteq \mathcal{B}_X$ and, for any $x \in X$, we have $\varphi(x) \in V$ for some $V \in \mathcal{D}$, whence $x \in \varphi^{-1}(V) \in \varphi^{-1}\mathcal{D}$.

Monotone: Assume $(m, \mathcal{D}) \leq (m', \mathcal{D}')$ and show $(m, \varphi^{-1}\mathcal{D}) \leq (m', \varphi^{-1}\mathcal{D}')$. So we need to show that $\mathcal{D} \preceq \mathcal{D}'$ implies $\varphi^{-1}\mathcal{D} \preceq \varphi^{-1}\mathcal{D}'$. We need to verify conditions (1) and (2) of the definition of \preceq (recalled in section 5.2.5). Concerning (1), given $\varphi^{-1}(V')$ with $V' \in \mathcal{D}'$, there is $V \in \mathcal{D}$ with $V' \subseteq V$, so $\varphi^{-1}(V') \subseteq \varphi^{-1}(V) \in \varphi^{-1}\mathcal{D}$. Concerning (2), given $x \in \varphi^{-1}(V)$ with $V \in \mathcal{D}$, we have $\varphi(x) \in V \in \mathcal{D}$, so there is $V' \in \mathcal{D}'$ such that $\varphi(x) \in V' \subseteq V$, so $\varphi^{-1}(V') \in \varphi^{-1}\mathcal{D}'$ and $x \in \varphi^{-1}(V') \subseteq \varphi^{-1}(V)$.

Step 2. For $j \in I(\mathcal{B}_Y)$, the map

$$\begin{aligned} \theta_j : \mathbf{H}_{\xi(j)} &\rightarrow \mathbf{H}_j \\ \mathcal{O}_{\xi(j)}(x) &\mapsto \mathcal{O}_j(\varphi(x)) \end{aligned}$$

is well-defined: if $\mathcal{O}_{\xi(j)}(x) = \mathcal{O}_{\xi(j)}(x')$, then $\mathcal{O}_j(\varphi(x)) = \mathcal{O}_j(\varphi(x'))$.

Indeed, write $j = (m, \mathcal{D})$, so $\xi(j) = (m, \varphi^{-1}\mathcal{D})$. Let $t = (D_0, \dots, D_{m-1}) \in \mathcal{O}_j(\varphi(x))$. Then, by equivariance and definition, $\varphi T^k(x) = S^k \varphi(x) \in D_k$ for $k = 0, \dots, m-1$. So $T^k(x) \in \varphi^{-1}D_k$ for $k = 0, \dots, m-1$. So x follows $t' := (\varphi^{-1}D_0, \dots, \varphi^{-1}D_{m-1})$. By assumption, then also x' follows t' . So $T^k(x') \in \varphi^{-1}D_k$ for $k = 0, \dots, m-1$. Whence $S^k \varphi(x') = \varphi T^k(x') \in D_k$ for $k = 0, \dots, m-1$. So $\varphi(x')$ follows t , so $t \in \mathcal{O}_j(\varphi(x'))$. The other direction is analogous.

Step 3. For each $j \in I(\mathcal{B}_Y)$, the map

$$\begin{aligned} \pi_j : D_{\xi(j)} = \mathbf{P}(\mathbf{H}_{\xi(j)}) &\rightarrow \mathbf{P}(\mathbf{H}_j) = E_j \\ M &\mapsto \theta_j(M), \end{aligned}$$

is a well-defined dynamical morphism.

Well-defined: Since M is a nonempty subset of $\mathbf{H}_{\xi(j)}$, also $\theta_j(M)$ is a nonempty subset of \mathbf{H}_j .

Scott-continuous: Since the domains are finite, it suffices to show monotonicity. If $M \leq M'$, then $M \supseteq M'$, then $\pi_j(M) = \theta_j(M) \supseteq \theta_j(M') = \pi_j(M')$, so $\pi_j(M) \leq \pi_j(M')$.

Max-preserving: If $M \in D_{\xi(j)}$ is maximal, then $M = \{\mathcal{O}_{\xi(j)}(x)\}$ for some $x \in X$. Hence $\pi_j(M) = \theta_j(M) = \{\mathcal{O}_j(\varphi(x))\}$ is maximal in E_j .

Max-bisimulative: If $M \in D_{\xi(j)}$ and $e \in \max E_j$ with $\pi_j(M) \leq e$, then e is a singleton subset of H_j that is a subset of $\pi_j(M)$, i.e., $e = \{\mathcal{O}_j(\varphi(x))\}$ for some $\mathcal{O}_{\xi(j)}(x) \in M$. Then $d := \{\mathcal{O}_{\xi(j)}(x)\} \in \max D_{\xi(j)}$ is such that $d \geq M$ and $\pi_j(d) = e$.

Valuation-preserving: Let $V \in \Sigma(E_j)$, and show $w_j(V) = v_{\xi(j)}(\pi_j^{-1}(V))$. Write $\max V = \{\{\mathcal{O}_j(y_1)\}, \dots, \{\mathcal{O}_j(y_m)\}\}$. Consider $\varphi^{-1} \bigcup_{k=1}^m [y_k]_j \subseteq X$. This set is partitioned by $\approx_{\xi(j)}$ into finitely many equivalence classes $[x_1]_{\xi(j)}, \dots, [x_n]_{\xi(j)}$ (with $n \geq 0$).³⁸ In particular, $\varphi^{-1} \bigcup_{k=1}^m [y_k]_j = \bigcup_{l=1}^n [x_l]_{\xi(j)}$.³⁹ We have $\max \pi_j^{-1}(V) = \{\{\mathcal{O}_{\xi(j)}(x_1)\}, \dots, \{\mathcal{O}_{\xi(j)}(x_n)\}\}$ since: $\{\mathcal{O}_{\xi(j)}(x)\} \in \max \pi_j^{-1}(V)$ iff $\{\mathcal{O}_j(\varphi(x))\} = \pi_j(\{\mathcal{O}_{\xi(j)}(x)\}) \in V$ iff $\mathcal{O}_j(\varphi(x)) = \mathcal{O}_j(y_k)$ for some $k \in \{1, \dots, m\}$ iff $\varphi(x) \in [y_k]_j$ for some $k \in \{1, \dots, m\}$ iff $x \in [x_l]_{\xi(j)}$ for some $l \in \{1, \dots, n\}$ iff $\{\mathcal{O}_{\xi(j)}(x)\} = \{\mathcal{O}_{\xi(j)}(x_l)\}$ for some $l \in \{1, \dots, n\}$. Thus,

$$\begin{aligned} w_j(V) &= \sum_{k=1}^m \nu([y_k]_j) = \nu \bigcup_{k=1}^m [y_k]_j = \mu \varphi^{-1} \bigcup_{k=1}^m [y_k]_j \\ &= \mu \bigcup_{l=1}^n [x_l]_{\xi(j)} = \sum_{l=1}^n \mu([x_l]_{\xi(j)}) = v_{\xi(j)}(\pi_j^{-1}(V)). \end{aligned}$$

Max-semi-equivariant: Let $a \in \max D_{\xi(j)}$, and show $\pi_j(f_{\xi(j)}(a)) \geq g_j(\pi_j(a))$. Write $a = \{\mathcal{O}_{\xi(j)}(x)\}$. On the left side, we have

$$\begin{aligned} \pi_j(f_{\xi(j)}(a)) &= \pi_j\{\mathcal{O}_{\xi(j)}(T(x')) : \mathcal{O}_{\xi(j)}(x') = \mathcal{O}_{\xi(j)}(x)\} \\ &= \{\mathcal{O}_j(\varphi T(x')) : \mathcal{O}_{\xi(j)}(x') = \mathcal{O}_{\xi(j)}(x)\}. \end{aligned}$$

On the right side, we have

$$g_j(\pi_j(a)) = g_j\{\mathcal{O}_j(\varphi(x))\} = \{\mathcal{O}_j(S(y)) : \mathcal{O}_j(y) = \mathcal{O}_j(\varphi(x))\}.$$

Since \leq is inverse inclusion, we need to show that the set on the left is a subset of the set on the right. Thus, given $\mathcal{O}_j(\varphi T(x'))$ with $\mathcal{O}_{\xi(j)}(x') = \mathcal{O}_{\xi(j)}(x)$, we have to

³⁸More precisely, choose a representative for each of the finitely many $\approx_{\xi(j)}$ -equivalence classes and let $\{x_1, \dots, x_n\}$ be the set of those representatives whose φ -image is in $\bigcup_{k=1}^m [y_k]_j$.

³⁹Proof: (\supseteq) If $x \in [x_l]_{\xi(j)}$ for some $l \in \{1, \dots, n\}$, then $\mathcal{O}_{\xi(j)}(x) = \mathcal{O}_{\xi(j)}(x_l)$, so $\mathcal{O}_j(\varphi(x)) = \mathcal{O}_j(\varphi(x_l))$. By definition, $\varphi(x_l) \in [y_k]_j$ for some $k \in \{1, \dots, m\}$. So also $\varphi(x) \in [y_k]_j$.

(\subseteq) Assume $x \in X$ with $\varphi(x) \in [y_k]_j$ for some $k \in \{1, \dots, m\}$. Let x' be the chosen representative of the equivalence class that x is in. So $\mathcal{O}_{\xi(j)}(x) = \mathcal{O}_{\xi(j)}(x')$, whence $\mathcal{O}_j(\varphi(x)) = \mathcal{O}_j(\varphi(x'))$. Since $\varphi(x) \in [y_k]_j$, also $\varphi(x') \in [y_k]_j$. So x' is among those representatives whose φ -image is in $\bigcup_{k=1}^m [y_k]_j$, i.e., $x' = x_l$ for some $l \in \{1, \dots, n\}$. So $x \in [x_l]_{\xi(j)}$.

show that it is in the set on the right side. Since π_j is a function, $\{\mathcal{O}_j(\varphi(x'))\} = \pi_j\{\mathcal{O}_{\xi(j)}(x')\} = \pi_j\{\mathcal{O}_{\xi(j)}(x)\} = \{\mathcal{O}_j(\varphi(x))\}$. Writing $y := \varphi(x') \in Y$, we have

$$\mathcal{O}_j(\varphi T(x')) = \mathcal{O}_j(S\varphi(x')) = \mathcal{O}_j(Sy) \in \{\mathcal{O}_j(S(y)) : \mathcal{O}_j(y) = \mathcal{O}_j(\varphi(x))\}.$$

Step 4. Recall that the projections $p_{ij}^D : D_j \rightarrow D_i$ in $(\mathfrak{D}_i, p_{ij}^D)$ are given by $M \mapsto \{\mathcal{O}_i(x) : \mathcal{O}_j(x) \in M\}$. Similarly for the projections $p_{ij}^E : E_j \rightarrow E_i$ in $(\mathfrak{E}_i, p_{ij}^E)$. We claim that, for all $i \leq j$ in $I(\mathcal{B}_Y)$, the following square commutes:

$$\begin{array}{ccc} D_{\xi(i)} & \xleftarrow{p_{\xi(i)\xi(j)}^D} & D_{\xi(j)} \\ \pi_i \downarrow & & \downarrow \pi_j \\ E_i & \xleftarrow{p_{ij}^E} & E_j \end{array}$$

Indeed, let $M \in D_{\xi(j)}$ and show $\pi_i \circ p_{\xi(i)\xi(j)}^D(M) = p_{ij}^E \circ \pi_j(M)$. On the left side, we have

$$\pi_i \circ p_{\xi(i)\xi(j)}^D(M) = \pi_i\{\mathcal{O}_{\xi(i)}(x) : \mathcal{O}_{\xi(j)}(x) \in M\} = \{\mathcal{O}_i(\varphi(x)) : \mathcal{O}_{\xi(j)}(x) \in M\}.$$

On the right side, we have

$$p_{ij}^E \circ \pi_j(M) = p_{ij}^E\{\mathcal{O}_j(\varphi(x)) : \mathcal{O}_{\xi(j)}(x) \in M\} = \{\mathcal{O}_i(\varphi(x)) : \mathcal{O}_{\xi(j)}(x) \in M\}.$$

So the two sides are identical.

Step 5. We show that $(D, \alpha_i)_{i \in I(\mathcal{B}_Y)}$ with $\alpha_i := \pi_i \circ p_{\xi(i)}^D : D \rightarrow E_i$ is a cone (in the category of dpos with Scott-continuous maps) to the expanding system of dpos $(E_i, p_{ij}^E)_{i \in I(\mathcal{B}_Y)}$. Note that, qua composition of dynamical morphisms, the α_i are dynamical morphisms.

Indeed, qua dynamical morphism, α_i in particular is Scott-continuous, and, for $i \leq j$ in $I(\mathcal{B}_Y)$, we have

$$p_{ij}^E \circ \alpha_j = (p_{ij}^E \circ \pi_j) \circ p_{\xi(j)}^D = (\pi_i \circ p_{\xi(i)\xi(j)}^D) \circ p_{\xi(j)}^D = \pi_i \circ p_{\xi(i)}^D = \alpha_i.$$

At this stage, we would be done if only the π_i were projections (which is the case if, for example, φ is surjective): Then $(\mathfrak{D}, \alpha_i)_{i \in I(\mathcal{B}_Y)}$ is a cone in \mathbf{dSCO}_n^p to the diagram $(\mathfrak{E}_i, p_{ij}^E)_{i \in I(\mathcal{B}_Y)}$ with \mathfrak{D} in \mathbf{dSCO}_{nm}^p , so $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ would be the mediating morphism according to theorem 5.2.15 (recalled in section 5.2.5), and hence, in particular, a morphism in \mathbf{dDOM} . Nonetheless, we can closely follow the proof of that theorem in our more general case here, too. (The only differences are that, first, we cannot, and need not, conclude anymore that α is a projection, and second, that we need a different argument that the set C_i in step 7 below is closed.) This then also suggests proofs for more general ‘limit existence’ theorems in \mathbf{dDOM} , which, to keep to the point, we leave to future work.

Step 6. It now follows (see Abramsky and Jung 1994, thm. 3.3.1 and thm. 3.3.7) that the function $\alpha : D \rightarrow E$ defined by

$$\alpha(a) := \langle \alpha_i(a) : i \in I(\mathcal{B}_Y) \rangle = \langle \pi_i(a(\xi(i))) : i \in I(\mathcal{B}_Y) \rangle$$

is well-defined and Scott-continuous.⁴⁰ It also is max-preserving: If $a \in \max D$, then each $\alpha_i(a)$ is maximal, so each entry of $\alpha(a)$ is maximal, so $\alpha(a) \in \max E$. Thus, to show that α is a dynamical morphism it remains to show that it is max-bisimulative, valuation-preserving, and max-semi-equivariant.

Step 7. We show that α is max-bisimulative. Let $a \in D$ and $\alpha(a) \leq e \in \max E$. We need to find $d \in D$ with $d \geq a$ and $\alpha(d) = e$. For $i \in I(\mathcal{B}_Y)$, consider $C_i := \uparrow a \cap \alpha_i^{-1}(e(i)) \subseteq D$.

This is a Lawson-closed subset of D : First, by definition of the Lawson topology, $\uparrow a$ is closed. Second, $p_{\xi(i)}^D : D \rightarrow D_{\xi(i)}$ is Lawson-continuous qua Scott-continuous projection,⁴¹ and $\pi_i : D_{\xi(i)} \rightarrow E_i$ is Lawson-continuous since it is a function between finite domains on which the Lawson topology is the discrete topology. Hence $\alpha_i = \pi_i \circ p_{\xi(i)}^D$ is Lawson-continuous. Third, $\{e(i)\} \subseteq E_i$ is Lawson-closed since the Lawson topology on the finite E_i is discrete. Hence $\alpha_i^{-1}(e(i))$ is Lawson-closed.

Moreover, C_i is nonempty: Since α_i is max-bisimulative and $\alpha_i(a) \leq e(i)$ (since $\alpha(a) \leq e$), there is $d \in \max D$ with $d \geq a$ and $\alpha_i(d) = e(i)$, so $d \in C_i$.

Finally, if $i \leq j$, then $C_i \supseteq C_j$: If $d \in C_j$, then $d \in D$ with $d \geq a$ and $\alpha_j(d) = e(j)$, so

$$\alpha_i(d) = p_{ij}(\alpha_j(d)) = p_{ij}(e(j)) = e(i),$$

whence $d \in C_i$.

Hence, $\{C_i : i \in I\}$ is a family of Lawson-closed subsets of D with the finite intersection property. Since D is a Scott domain, its Lawson topology is compact, so $\bigcap_{i \in I(\mathcal{B}_Y)} C_i$ is nonempty. Let $d \in \bigcap_{i \in I(\mathcal{B}_Y)} C_i$. Then $d \geq a$ and $\alpha(d) = \langle \alpha_i(d) : i \in I(\mathcal{B}_Y) \rangle = \langle e(i) : i \in I(\mathcal{B}_Y) \rangle = e$, as needed.

Step 8. We show that α is valuation-preserving. We first show this for the basic Scott-opens of E . Since the Scott topology on a bilimit of continuous domains is the restriction of the product topology (Abramsky and Jung 1994, exercise 3.3.12 (18)), the basic opens of E are of the form $V = p_{i_1}^{-1}(U_{i_1}) \cap \dots \cap p_{i_n}^{-1}(U_{i_n})$ with $i_k \in I(\mathcal{B}_Y)$ and $U_{i_k} \in \Sigma(E_{i_k})$ for $k = 1, \dots, n$. We show $w(V) = v(\alpha^{-1}(V))$.

⁴⁰ This is readily seen directly: It is well-defined, since the α_i commute with the $p_{i_j}^E$. It is monotone, since, if $a \leq a'$, then $\alpha_i(a) \leq \alpha_i(a')$ for all i , so $\alpha(a) \leq \alpha(a')$. It preserves directed suprema, since, if $A \subseteq D$ is directed, then $\alpha(\bigvee A) = \langle \alpha_i(\bigvee A) : i \in I(\mathcal{B}_Y) \rangle = \langle \bigvee \alpha_i(A) : i \in I(\mathcal{B}_Y) \rangle = \bigvee \{ \langle \alpha_i(a) : i \in I(\mathcal{B}_Y) \rangle : a \in A \} = \bigvee \alpha(A)$.

⁴¹ Proof: The Lawson topology on a dcpo D is generated by the Scott-open sets together with complements of principal upsets ($D \setminus \uparrow x$). And the preimage of a Scott-open set (resp. principal upset) under a Scott-continuous projection is again Scott-open (resp. a principal upset).

Since $I(\mathcal{B}_Y)$ is directed, let $j \geq i_1, \dots, i_n$. So

$$V = \bigcap_{k=1}^n p_{i_k}^{-1}(U_{i_k}) = \bigcap_{k=1}^n p_j^{-1}(p_{i_k j}^{-1}(U_{i_k})) = p_j^{-1}\left(\bigcap_{k=1}^n p_{i_k j}^{-1}(U_{i_k})\right).$$

Hence

$$\begin{aligned} w(V) &= w\left(p_j^{-1}\left(\bigcap_{k=1}^n p_{i_k j}^{-1}(U_{i_k})\right)\right) = w_j\left(\bigcap_{k=1}^n p_{i_k j}^{-1}(U_{i_k})\right) \\ &= v\left(\alpha_j^{-1}\left(\bigcap_{k=1}^n p_{i_k j}^{-1}(U_{i_k})\right)\right) = v\left(\bigcap_{k=1}^n \alpha_j^{-1}(p_{i_k j}^{-1}(U_{i_k}))\right) \\ &= v\left(\bigcap_{k=1}^n \alpha_{i_k}^{-1}(U_{i_k})\right) \stackrel{\alpha_i = p_i \circ \alpha}{=} v\left(\bigcap_{k=1}^n \alpha^{-1}(p_{i_k}^{-1}(U_{i_k}))\right) \\ &= v\left(\alpha^{-1}\left(\bigcap_{k=1}^n p_{i_k}^{-1}(U_{i_k})\right)\right) = v\left(\alpha^{-1}(V)\right). \end{aligned}$$

Now we need to show the claim for arbitrary opens V of E . So $V = \bigcup_{k \in K} V_k$ for basic opens V_k and an index set K . Without loss of generality, $\{V_k : k \in K\}$ is directed (otherwise consider the family of finite unions of the V_k 's). Note that then also $\alpha^{-1}(V_k)$ is a directed family of open sets in D . Then

$$\begin{aligned} w(V) &= w\left(\bigcup_{k \in K} V_k\right) = \sup_{k \in K} w(V_k) = \sup_{k \in K} v(\alpha^{-1}(V_k)) \\ &= v\left(\bigcup_{k \in K} \alpha^{-1}(V_k)\right) = v\left(\alpha^{-1}\left(\bigcup_{k \in K} V_k\right)\right) = v\left(\alpha^{-1}(V)\right). \end{aligned}$$

Step 9. We show that α is max-semi-equivariant. Let $a \in \max D$ and show $\alpha(f(a)) \geq g(\alpha(a))$. Since f is max-preserving, $f(a) \in \max D$. Since α is max-preserving, $\alpha(f(a)) \in \max E$. Since α_i is max-semi-equivariant, we have, for $i \in I(\mathcal{B}_Y)$,

$$\alpha(f(a))(i) = \alpha_i(f(a)) \geq g_i(\alpha_i(a)) = g_i(\alpha(a)(i)).$$

Thus, $\alpha(a), \alpha(f(a)) \in \max E$ are such that, for all $i \in I(\mathcal{B}_Y)$, we have $\alpha(f(a))(i) \geq g_i(\alpha(a)(i))$. Hence, by the properties of g (see theorem 5.2.15), $g(\alpha(a)) = \alpha(f(a))$, as needed. \square

Now, we can define the functor D .

5.4.3. PROPOSITION. *The following defines a functor $D : \mathbf{bTS}_0 \rightarrow \mathbf{dDOM}$:*

- For $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 , $D(\mathfrak{X}, \mathcal{B})$ is the observation domain of $(\mathfrak{X}, \mathcal{B})$ which is in \mathbf{dDOM} (as recalled in section 5.2.5).

- If $\varphi : (\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Y}, \mathcal{B}_Y)$ is a morphism in \mathbf{bTS}_0 , then $D(\varphi) := \alpha$ is as in proposition 5.4.2 and hence in \mathbf{dDOM} .

We also define $\hat{D} := D \circ I_B : \mathbf{TS}_{0c} \rightarrow \mathbf{dDOM}$.

PROOF. We need to check that D preserves identity and composition.

Identity. For $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 , we have to show $D(\text{id}_{(\mathfrak{X}, \mathcal{B})}) = \text{id}_{D(\mathfrak{X}, \mathcal{B})}$. So let $a \in D(\mathfrak{X}, \mathcal{B})$ and show $D(\text{id}_{(\mathfrak{X}, \mathcal{B})})(a) = a$. Indeed, we have

$$\begin{aligned} D(\text{id}_{(\mathfrak{X}, \mathcal{B})})(a) &= \langle \pi_{(m, \mathcal{D})}(a(m, \text{id}_{(\mathfrak{X}, \mathcal{B})}^{-1} \mathcal{D})) : (m, \mathcal{D}) \in I(\mathcal{B}) \rangle \\ &= \langle \pi_{(m, \mathcal{D})}(a(m, \mathcal{D})) : (m, \mathcal{D}) \in I(\mathcal{B}) \rangle, \end{aligned}$$

so it suffices to show that $\pi_{(m, \mathcal{D})}(a(m, \mathcal{D})) = a(m, \mathcal{D})$. Indeed, we have

$$\begin{aligned} \pi_{(m, \mathcal{D})}(a(m, \mathcal{D})) &= \{ \mathcal{O}_{\mathcal{D}}^m(\text{id}_{(\mathfrak{X}, \mathcal{B})}(x)) : \mathcal{O}_{\text{id}_{(\mathfrak{X}, \mathcal{B})}^{-1} \mathcal{D}}^m(x) \in a(m, \mathcal{D}) \} \\ &= \{ \mathcal{O}_{\mathcal{D}}^m(x) : \mathcal{O}_{\mathcal{D}}^m(x) \in a(m, \mathcal{D}) \} = a(m, \mathcal{D}). \end{aligned}$$

Composition. Let $\varphi : (\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Y}, \mathcal{B}_Y)$ and $\psi : (\mathfrak{Y}, \mathcal{B}_Y) \rightarrow (\mathfrak{Z}, \mathcal{B}_Z)$ be morphisms in \mathbf{bTS}_0 . We have to show $D(\psi \circ \varphi) = D(\psi) \circ D(\varphi)$. So let $a \in D(\mathfrak{X}, \mathcal{B}_X)$ and $k = (l, \mathcal{E}) \in I(\mathcal{B}_Z)$, and show $D(\psi \circ \varphi)(a)(k) = D(\psi) \circ D(\varphi)(a)(k)$.

We write $\psi\varphi$ for $\psi \circ \varphi$, and we write $\pi^{\psi\varphi}, \pi^\psi, \pi^\varphi$ for the functions occurring in the definition of $D(\psi \circ \varphi), D(\psi), D(\varphi)$, respectively. Thus, on the left side, we have

$$D(\psi\varphi)(a)(k) = \pi_{(l, \mathcal{E})}^{\psi\varphi}(a(l, \psi\varphi^{-1}\mathcal{E}))$$

On the right side, we have

$$\begin{aligned} D(\psi) \circ D(\varphi)(a)(k) &= \pi_{(l, \mathcal{E})}^\psi \left(D(\varphi)(a)(l, \psi^{-1}\mathcal{E}) \right) \\ &= \pi_{(l, \mathcal{E})}^\psi \left(\pi_{(l, \psi^{-1}\mathcal{E})}^\varphi (a(l, \varphi^{-1}\psi^{-1}\mathcal{E})) \right). \end{aligned}$$

Since $\psi\varphi^{-1}\mathcal{E} = \varphi^{-1}\psi^{-1}\mathcal{E}$, it hence suffices to show that $\pi_{(l, \mathcal{E})}^{\psi\varphi} = \pi_{(l, \mathcal{E})}^\psi \circ \pi_{(l, \psi^{-1}\mathcal{E})}^\varphi$. Note that these two functions type-check:

$$\begin{aligned} \pi_{(l, \mathcal{E})}^{\psi\varphi} &: \mathbf{P}(\mathbf{H}_{(l, \varphi^{-1}\psi^{-1}\mathcal{E})}^X) \rightarrow \mathbf{P}(\mathbf{H}_{(l, \mathcal{E})}^Z) \\ \pi_{(l, \mathcal{E})}^\psi \circ \pi_{(l, \psi^{-1}\mathcal{E})}^\varphi &: \mathbf{P}(\mathbf{H}_{(l, \varphi^{-1}\psi^{-1}\mathcal{E})}^X) \rightarrow \mathbf{P}(\mathbf{H}_{(l, \psi^{-1}\mathcal{E})}^Y) \rightarrow \mathbf{P}(\mathbf{H}_{(l, \mathcal{E})}^Z). \end{aligned}$$

Now, for $M \in \mathbf{P}(\mathbf{H}_{(l, \varphi^{-1}\psi^{-1}\mathcal{E})}^X)$ we have

$$\begin{aligned} \pi_{(l, \mathcal{E})}^\psi \circ \pi_{(l, \psi^{-1}\mathcal{E})}^\varphi(M) &= \pi_{(l, \mathcal{E})}^\psi \left\{ \mathcal{O}_{\psi^{-1}\mathcal{E}}^l(\varphi(x)) : \mathcal{O}_{\varphi^{-1}\psi^{-1}\mathcal{E}}^l(x) \in M \right\} \\ &= \left\{ \mathcal{O}_{\mathcal{E}}^l(\psi\varphi(x)) : \mathcal{O}_{\varphi^{-1}\psi^{-1}\mathcal{E}}^l(x) \in M \right\} = \pi_{(l, \mathcal{E})}^{\psi\varphi}(M), \end{aligned}$$

as needed. \square

Note that D restricts to $D : \mathbf{bTS}_{0s} \rightarrow \mathbf{dDOM}_s$ since, as noted in section 5.2.5, if $(\mathfrak{X}, \mathcal{B})$ is in \mathbf{bTS}_{0s} (i.e., in \mathbf{bTS}_0 and, additionally, \mathfrak{X} is standard and \mathcal{B} is forward closed), then $D(\mathfrak{X}, \mathcal{B})$ is standard.

We also make the following observation—which will be useful later on—about the behavior of $D(\varphi)$ on the maximal elements of $D(\mathfrak{X}, \mathcal{B})$.

5.4.4. LEMMA. 1. For $\varphi : (\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Y}, \mathcal{B}_Y)$ in \mathbf{bTS}_0 , we have, in \mathbf{TS}_0 ,

$$SD(\varphi) \circ \varphi_X = \varphi_Y \circ \varphi.$$

2. For $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ in \mathbf{TS}_{0c} , we have, in \mathbf{TS}_{0c} ,

$$S\hat{D}(\varphi) = \varphi_Y \circ \varphi \circ \varphi_X^{-1}.$$

PROOF. Ad (1). For $x \in X$, we have

$$\begin{aligned} SD(\varphi)(\varphi_X(x)) &= D(\varphi) \upharpoonright \max D(\varphi_X(x)) = D(\varphi)(\varphi_X(x)) \\ &= \left\langle \pi_{(m, \mathcal{D})}(\varphi_X(x)(m, \varphi^{-1}\mathcal{D})) : (m, \mathcal{D}) \in I(\mathcal{B}_Y) \right\rangle \\ &= \left\langle \pi_{(m, \mathcal{D})}(\{\mathcal{O}_{\varphi^{-1}\mathcal{D}}^m(x)\}) : (m, \mathcal{D}) \in I(\mathcal{B}_Y) \right\rangle \\ &= \left\langle \{\mathcal{O}_{\mathcal{D}}^m(\varphi(x))\} : (m, \mathcal{D}) \in I(\mathcal{B}_Y) \right\rangle \\ &= \varphi_Y(\varphi(x)). \end{aligned}$$

Ad (2). Since \mathfrak{X} is compact, $\varphi_X : \mathfrak{X} \rightarrow S\hat{D}(\mathfrak{X}) = SD(X, \text{Clp}(X))$ is an isomorphism in \mathbf{TS}_{0c} (recalled in section 5.2.5). So it has an inverse φ_X^{-1} . By (1), we have for $\varphi = \mathbb{I}_B(\varphi) : \mathbb{I}_B(\mathfrak{X}) \rightarrow \mathbb{I}_B(\mathfrak{Y})$ that

$$S\hat{D}(\varphi) \circ \varphi_X = SD(\mathbb{I}_B(\varphi)) \circ \varphi_X = \varphi_Y \circ \mathbb{I}_B(\varphi) = \varphi_Y \circ \varphi.$$

Thus, the claim follows by ‘multiplying’ φ_X^{-1} to the right. \square

5.4.3 Computational and logical compactification coincide

We show that the diagram

$$\begin{array}{ccc} & \mathbf{dDOM} & \\ \mathbf{D} \nearrow & & \searrow \mathbf{S} \\ \mathbf{bTS}_0 & \xrightarrow{\quad \bar{\mathbf{C}} \quad} & \mathbf{TS}_{0c} \end{array}$$

commutes up to natural isomorphism. Thus, as promised in section 5.3.2, the computational compactification $S \circ D : \mathbf{bTS}_0 \rightarrow \mathbf{TS}_{0c}$ coincides with the logical compactification $\bar{\mathbf{C}} : \mathbf{bTS}_0 \rightarrow \mathbf{TS}_{0c}$.

5.4.5. PROPOSITION. *If $(\mathfrak{X}, \mathcal{B})$ is in \mathbf{bTS}_0 , let $\psi_{(\mathfrak{X}, \mathcal{B})}$ be the ψ in the diagram*

$$\begin{array}{ccc} (\mathfrak{X}, \overline{\mathcal{B}}) & \xleftarrow{\eta} & \mathbb{I}_B(\mathbb{C}(\mathfrak{X}, \overline{\mathcal{B}})) = \mathbb{I}_B(\overline{\mathbb{C}}(\mathfrak{X}, \mathcal{B})) \\ & \searrow \varphi_X & \downarrow \mathbb{I}_B(\psi) = \psi \\ & & \mathbb{I}_B(\mathbb{SD}(\mathfrak{X}, \mathcal{B})) \end{array}$$

Then $\psi_{(\mathfrak{X}, \mathcal{B})} : \overline{\mathbb{C}}(\mathfrak{X}, \mathcal{B}) \rightarrow \mathbb{SD}(\mathfrak{X}, \mathcal{B})$ is an isomorphism in \mathbf{TS}_{0c} .

PROOF. Let's first check that the diagram makes sense. As recalled in section 5.2.5, the canonical embedding $\varphi_X : \mathfrak{X} \rightarrow \mathbb{SD}(\mathfrak{X}, \mathcal{B})$ is in \mathbf{TS}_0 . The equivalence classes $[x]_i$ are Boolean combinations of sets in $\bigcup_{k \geq 0} T^{-k}\mathcal{B}$, so they are in $\overline{\mathcal{B}}$. Since φ -preimages of clopen subsets of $\mathbb{SD}(\mathfrak{X}, \mathcal{B})$ can be written as Boolean combinations of such equivalence classes, they, too, are in $\overline{\mathcal{B}}$. Hence $\varphi_X : (\mathfrak{X}, \overline{\mathcal{B}}) \rightarrow \mathbb{I}_B(\mathbb{SD}(\mathfrak{X}, \mathcal{B}))$ is in \mathbf{bTS}_0 . Since \mathbb{C} is left adjoint to \mathbb{I}_B (proposition 5.3.6), we indeed obtain the desired diagram.

To show that ψ is our desired isomorphism, we need to show that it is bijective: it then is a homeomorphism (qua bijective continuous function between compact Hausdorff spaces) and hence, by proposition 5.2.8, an isomorphism. Let's write X, Y, Z for the state spaces of $\mathfrak{X}, \overline{\mathbb{C}}(\mathfrak{X}, \mathcal{B}), \mathbb{SD}(\mathfrak{X}, \mathcal{B})$, respectively, equipped with their respective topologies.

Surjective: Note that $\varphi_X(X) \subseteq \psi(Y)$: If $z \in \varphi_X(X)$, then $z = \varphi_X(x)$ for some $x \in X$, so $y := \eta(x) \in Y$ and $\psi(y) = \psi \circ \eta(x) = \varphi_X(x) = z$, whence $z \in \psi(Y)$. Moreover, since $\psi : Y \rightarrow Z$ is a continuous function from a compact space into a Hausdorff space, it is closed, so $\psi(Y) \subseteq Z$ is closed. Since $\varphi_X(X) \subseteq Z$ is dense (section 5.2.5), it follows that $\psi(X) = Z$ (since $\psi(X)$ is a closed subset of Z that contains a dense set), so ψ is surjective.

Injective: Let $P \neq P'$ be in Y (so P and P' are ultrafilters on $\overline{\mathcal{B}}$), and show $\psi(P) \neq \psi(P')$. Since $\overline{\mathcal{B}}$ is generated by $\bigcup_{k \geq 0} T^{-k}\mathcal{B}$, there is $V = T^{-k}U$ with $k \geq 0$ and $U \in \mathcal{B}$ such that $V \in P$ but $V \notin P'$, or $V \notin P$ but $V \in P'$.⁴² Without loss of generality, assume the former.

We first show that $\psi(P) \in \text{Cl}(\varphi_X(V))$. Since $D(V)$ is an open neighborhood of P and $\eta(X)$ is dense in Y , there are $P_n \in D(V) \cap \eta(X)$ such that $\lim_n P_n = P$. So $P_n = \eta(x_n)$ for some $x_n \in X$ and, in fact, $x_n \in V$ (since $\eta(x_n) = P_n \in D(V)$, we have $V \in \eta(x_n)$, so $x_n \in V$). So, by continuity of ψ ,

$$\psi(P) = \lim_n \psi(P_n) = \lim_n \psi \circ \eta(x_n) = \lim_n \varphi_X(x_n) \in \text{Cl}(\varphi_X(V)).$$

Similarly, $\psi(P') \in \text{Cl}(\varphi_X(V^c))$.

⁴²Proof: An ultrafilter P on a Boolean algebra B essentially is a Boolean algebra homomorphism $B \rightarrow 2$. And if two such homomorphisms agree on the generators B_0 of B , they are identical. So if $P \neq P'$, there is a generator $a \in B_0$ such that $P(a) \neq P'(a)$.

Next we show that $\text{Cl}(\varphi_X(V))$ and $\text{Cl}(\varphi_X(V^c))$ are disjoint, which then implies $\psi(P) \neq \psi(P')$, as needed. For this, it is enough to find a clopen set $A \subseteq Z$ such that $\varphi_X(V) \subseteq A$ and $\varphi_X(V^c) \subseteq A^c$.

To this end, let $i := (k+1, \{U, X\}) \in I(\mathcal{B})$ (recall that $U \in \mathcal{B}$ with $V = T^{-k}U$). So \approx_i partitions V into finitely many equivalence classes: $V = ([x_1]_i \cup \dots \cup [x_m]_i) \cap V$ for $x_1, \dots, x_m \in V$ and $m \geq 1$ (since $V \neq \emptyset$ because $V \in P$). Define

$$A := \bigcup_{l=1}^m p_i^{-1}(\{\mathcal{O}_i(x_l)\}) \cap \max D.$$

So A is clopen qua finite union of clopen sets.⁴³ And we have $\varphi_X(V) \subseteq A$: if $z \in \varphi_X(V)$, then $z = \varphi_X(x)$ for some $x \in V$, so $x \in [x_l]_i$ for some $l \in \{1, \dots, m\}$. So $z(i) = \varphi_X(x)(i) = \{\mathcal{O}_i(x)\} = \{\mathcal{O}_i(x_l)\}$. Hence $z \in A$.

So it remains to show $\varphi_X(V^c) \subseteq A^c$. Assume for contradiction that there is $z \in \varphi_X(V^c)$ but $z \in A$. So $z = \varphi_X(x)$ for some $x \in V^c$. So $x \notin V$, i.e., $T^k(x) \notin U$. Since $z \in A$, there is $l \in \{1, \dots, m\}$ such that $z \in p_i^{-1}(\{\mathcal{O}_i(x_l)\}) \cap \max D$. Hence

$$\{\mathcal{O}_i(x)\} = \varphi_X(x)(i) = z(i) = \{\mathcal{O}_i(x_l)\}.$$

However, we now argue that $\mathcal{O}_i(x) \neq \mathcal{O}_i(x_l)$, which yields the required contradiction. Indeed, consider the trajectory $t = (X, \dots, X, U) \in \{X, U\}^{k+1}$ (where the X 's are repeated k -times). Since $x_l \in V = T^{-k}U$, we have $T^k(x_l) \in U$, so x_l follows t , i.e., $t \in \mathcal{O}_i(x_l)$. However, we cannot have $t \in \mathcal{O}_i(x)$, since this would imply $T^k(x) \in U$. \square

5.4.6. PROPOSITION. *The family of isomorphisms $\psi := (\psi_{(\mathfrak{X}, \mathcal{B})})$ from proposition 5.4.5 is natural in $(\mathfrak{X}, \mathcal{B})$. So ψ is a natural isomorphism from $\overline{\mathcal{C}}$ to SD .*

PROOF. Let $\varphi : (\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Y}, \mathcal{B}_Y)$ be a morphism in \mathbf{bTS}_0 , and show

$$\psi_{(\mathfrak{Y}, \mathcal{B}_Y)} \circ \overline{\mathcal{C}}(\varphi) = \text{SD}(\varphi) \circ \psi_{(\mathfrak{X}, \mathcal{B}_X)}.$$

The proof is, as we'll discuss now, essentially in figure 5.5. First note that, since we want to show that $\psi_{(\mathfrak{Y}, \mathcal{B}_Y)} \circ \overline{\mathcal{C}}(\varphi)$ and $\text{SD}(\varphi) \circ \psi_{(\mathfrak{X}, \mathcal{B}_X)}$ of the (outer) square are identical as functions, we can also view them as being in the supercategory \mathbf{TS}_0 (instead of the full subcategory \mathbf{TS}_{0c}), which allows us to place the other (inner) morphisms $\eta_X, \varphi_X, \varphi, \eta_Y, \varphi_Y$ into the same diagram.

We will first show that the subdiagrams (a), (b), (c), (d) commute. Concerning (a), as seen in the proof of proposition 5.4.5, we can regard φ_X as a morphism

$$(\mathfrak{X}, \overline{\mathcal{B}_X}) \xrightarrow{\varphi_X} \mathbf{I}_B(\text{SD}(\mathfrak{X}, \mathcal{B}_X))$$

⁴³The sets $p_i^{-1}(\{\mathcal{O}_i(x)\}) \cap \max D$ are open in $\max D$ since $\{\mathcal{O}_i(x)\}$ is maximal in D_i and p_i Scott-continuous, and they are closed since the complement is given by the sets of the same form constructed from the finitely many other maximal elements of D_i .

$$\begin{array}{ccc}
\overline{\mathcal{C}}(\mathfrak{X}, \mathcal{B}_X) & \xrightarrow{\psi_{(\mathfrak{X}, \mathcal{B}_X)}} & \text{SD}(\mathfrak{X}, \mathcal{B}_X) \\
\downarrow \overline{\mathcal{C}}(\varphi) & \begin{array}{c} \swarrow \eta_X \quad (a) \quad \searrow \varphi_X \\ \mathfrak{X} \\ \downarrow \varphi \\ \mathfrak{Y} \\ \swarrow \eta_Y \quad (c) \quad \searrow \varphi_Y \end{array} & \downarrow \text{SD}(\varphi) \\
\overline{\mathcal{C}}(\mathfrak{Y}, \mathcal{B}_Y) & \xrightarrow{\psi_{(\mathfrak{Y}, \mathcal{B}_Y)}} & \text{SD}(\mathfrak{Y}, \mathcal{B}_Y)
\end{array}$$

Figure 5.5: Proof that $\psi = (\psi_{(\mathfrak{X}, \mathcal{B})})$ is natural in $(\mathfrak{X}, \mathcal{B})$.

that commutes with

$$(\mathfrak{X}, \overline{\mathcal{B}}_X) \xrightarrow{\eta_X} \mathbb{I}_B(\mathcal{C}(\mathfrak{X}, \overline{\mathcal{B}}_X)) = \mathbb{I}_B(\overline{\mathcal{C}}(\mathfrak{X}, \mathcal{B}_X)) \xrightarrow{\psi_{(\mathfrak{X}, \mathcal{B}_X)}} \mathbb{I}_B(\text{SD}(\mathfrak{X}, \mathcal{B}_X)),$$

so $\varphi_X = \psi_{(\mathfrak{X}, \mathcal{B}_X)} \circ \eta_X$, as needed. Concerning (c), we reason analogously. Concerning (d), this is lemma 5.4.4 (1). Concerning (b), recall from section 5.3.3 that we have

$$\begin{array}{ccccc}
(\mathfrak{X}, \mathcal{B}_X) & \xleftarrow{\epsilon_X = \text{id}_{\mathfrak{X}}} & (\mathfrak{X}, \overline{\mathcal{B}}_X) & \xrightarrow{\eta_X} & \mathbb{I}_B(\mathcal{C}(\mathfrak{X}, \overline{\mathcal{B}}_X)) \\
\downarrow \varphi & & \downarrow \overline{\varphi} = \varphi & & \downarrow \overline{\mathcal{C}}(\varphi) = \mathcal{C}(\overline{\varphi}) \\
(\mathfrak{Y}, \mathcal{B}_Y) & \xleftarrow{\epsilon_Y = \text{id}_{\mathfrak{Y}}} & (\mathfrak{Y}, \overline{\mathcal{B}}_Y) & \xrightarrow{\eta_Y} & \mathbb{I}_B(\mathcal{C}(\mathfrak{Y}, \overline{\mathcal{B}}_Y))
\end{array}$$

and, in particular, the right square commutes, as needed.


Next, we argue that the outer morphisms agree on the subset $\eta_X(X)$ of the state space of $\overline{\mathcal{C}}(\mathfrak{X}, \mathcal{B}_X)$: Given $\eta_X(x)$ for some $x \in \mathfrak{X}$, we have, by commutativity of the subdiagrams,

$$\psi_{(\mathfrak{Y}, \mathcal{B}_Y)} \circ \overline{\mathcal{C}}(\varphi)(\eta_X(x)) = \varphi_Y \circ \varphi(x) = \text{SD}(\varphi) \circ \psi_{(\mathfrak{X}, \mathcal{B}_X)}(\eta_X(x)).$$

Finally, now the continuous functions $\psi_{(\mathfrak{Y}, \mathcal{B}_Y)} \circ \overline{\mathcal{C}}(\varphi)$ and $\text{SD}(\varphi) \circ \psi_{(\mathfrak{X}, \mathcal{B}_X)}$ agree on the dense subset $\eta_X(X)$ of their domain $\overline{\mathcal{C}}(\mathfrak{X}, \mathcal{B}_X)$, so they are identical, as needed. \square

5.5 The systems-domains adjunction

In this section, we establish the following part of the main diagram:



In other words, we show—as motivated in the introduction—that the domain functor \hat{D} is right adjoint to the systems functor S . (The standard case is treated collectively in section 5.7.)

We proceed by explicitly providing the counit $\epsilon : \hat{S}\hat{D} \rightarrow 1_{\mathbb{T}S_{0c}}$ (which will turn out to be a natural isomorphism) and the unit $\eta : 1_{\mathbb{d}DOM} \rightarrow \hat{D}S$ of the adjunction and then show that they satisfy the triangle identities.

Before we start, we note two more concrete corollaries of lemma 5.2.17 stating a sufficient condition for being in a closed subset $A \subseteq \max D$ of \mathfrak{D} in $\mathbb{d}DOM$.

5.5.1. LEMMA. *Let \mathfrak{X} be in $\mathbb{T}S_{0c}$. Let $x \in X$ and $A \subseteq X$ closed. Assume that, for each $i \in I(\text{Clp}X)$, there is $x_i \in A$ such that $\mathcal{O}_i(x) = \mathcal{O}_i(x_i)$. Then $x \in A$.*

PROOF. Let $\mathfrak{D} := \hat{D}(\mathfrak{X})$ and, for $i \in I(\text{Clp}X)$, let $p_i : \mathfrak{D} \rightarrow \mathfrak{D}_i = P(H_i)$ be the limiting morphisms. Let $\varphi_X : X \rightarrow \max D$ the canonical homeomorphism. So $A' := \varphi_X(A) \subseteq \max D$ is closed. Let $x' := \varphi_X(x) \in \max D$. Now, from the assumption, we have, for $i \in I(\text{Clp}X)$, that $x_i \in A$, whence $\varphi_X(x_i) \in A'$, and

$$p_i(x') = p_i(\varphi_X(x)) = \{\mathcal{O}_i(x)\} = \{\mathcal{O}_i(x_i)\} = p_i(\varphi_X(x_i)) \in p_i(A').$$

Hence, lemma 5.2.17 yields $x' \in A'$, so $x = \varphi_X^{-1}(x') \in \varphi_X^{-1}(A') = A$, as needed. \square

As a corollary we get the following.

5.5.2. LEMMA. *Let $\mathfrak{D} = (D, v, f)$ be in $\mathbb{d}DOM$. Let $a \in D$ and $x \in \max D$. For each $j \in I(\text{Clp}(\max D))$, let $x_j \in \uparrow a \cap \max D$ such that $\mathcal{O}_j(x) = \mathcal{O}_j(x_j)$. Then $x \geq a$.*

PROOF. Note that $\mathfrak{X} := S(\mathfrak{D})$ is in $\mathbb{T}S_{0c}$ with state space $X = \max D$. So $x \in X$. And $A := \uparrow a \cap \max D$ is a closed subset of X : Since $\uparrow a \subseteq D$ is Lawson-closed, A is closed in the relative Lawson topology on $\max D$, which coincides, since D is a Scott domain, with the relative Scott topology on $\max D = X$. Now, by assumption, for $j \in I(\text{Clp}X)$, we have $x_j \in A$ with $\mathcal{O}_j(x) = \mathcal{O}_j(x_j)$. So lemma 5.5.1 implies $x \in A$, whence $x \geq a$. \square

5.5.1 The counit and unit

The counit is immediate.

5.5.3. PROPOSITION. *For \mathfrak{X} in \mathbf{TS}_{0c} , let $\epsilon_{\mathfrak{X}} := \varphi_X^{-1}$ be the inverse of the isomorphism $\varphi_X : \mathfrak{X} \rightarrow \hat{\mathbf{S}}(\mathfrak{X})$ in \mathbf{TS}_{0c} . Then $\epsilon : \hat{\mathbf{S}} \rightarrow 1_{\mathbf{TS}_{0c}}$ is a natural isomorphism.*

PROOF. For each \mathfrak{X} in \mathbf{TS}_{0c} , $\epsilon_{\mathfrak{X}} : \hat{\mathbf{S}}(\mathfrak{X}) \rightarrow \mathfrak{X}$ is an isomorphism in \mathbf{TS}_{0c} . So we need to show that the morphisms $\epsilon_{\mathfrak{X}}$ are natural in \mathfrak{X} . Indeed, for a morphism $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ in \mathbf{TS}_{0c} , we have, by lemma 5.4.4 (2),

$$\epsilon_{\mathfrak{Y}} \circ \hat{\mathbf{S}}(\varphi) = \varphi_Y^{-1} \circ \varphi_Y \circ \varphi \circ \varphi_X^{-1} = \varphi \circ \varphi_X^{-1} = 1_{\mathbf{TS}_{0c}}(\varphi) \circ \epsilon_{\mathfrak{X}},$$

as needed. \square

But the unit requires some work. In the first proposition below, we show how to construct, for \mathfrak{D} in \mathbf{dDOM} , a morphism $\eta_{\mathfrak{D}} : \mathfrak{D} \rightarrow \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D})$, and in the second proposition we then show that these are natural in \mathfrak{D} .

5.5.4. PROPOSITION. *Let \mathfrak{D} be in \mathbf{dDOM} . Then the following defines a morphism in \mathbf{dDOM} :*

$$\begin{aligned} \eta_{\mathfrak{D}} := \alpha : \mathfrak{D} &\rightarrow \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) \\ a &\mapsto \left\langle \alpha_j(a) : j \in I(\mathbf{Clp}(\max D)) \right\rangle \end{aligned}$$

where

$$\alpha_j(a) := \{ \mathcal{O}_j(x) : x \in \uparrow a \cap \max D \}.$$

Moreover, for $x \in \max D$, $\alpha(x) = \varphi_{\max D}(x)$. (Here $\varphi_{\max D}$ is the canonical homeomorphism between the state space $\max D$ of $\mathbf{S}(\mathfrak{D})$ and the maximal elements of the domain underlying $\hat{\mathbf{D}}\mathbf{S}(\mathfrak{D})$.)

PROOF. We first fix notation. Write $\mathfrak{D} = (D, v, f)$, and write $\mathfrak{X} = (X, \tau, \mu, T)$ for $\mathbf{S}(\mathfrak{D}) = (\max D, \Sigma(D) \upharpoonright \max D, \mu_v \upharpoonright \mathcal{B}(\max D), f \upharpoonright \max D)$. Write $\hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) = \hat{\mathbf{D}}(\mathfrak{X}) =: \mathfrak{E} = (E, w, g)$, the limit of $(\mathfrak{E}_j, q_{ij})_{j \in I(\mathbf{Clp}X)}$. So $E_j = \mathbf{P}(\mathbf{H}_j^X)$.

Step 1. For $j \in I(\mathbf{Clp}X)$ define

$$\begin{aligned} \alpha_j : D &\rightarrow E_j \\ a &\mapsto \{ \mathcal{O}_j(x) : x \in \uparrow a \cap \max D \}. \end{aligned}$$

This is well-defined: Since $x \in \max D = X$, the set of trajectories $\mathcal{O}_j(x) \in \mathbf{H}_j^X$ is defined. Moreover, since D is a dcpo, $\uparrow a \cap \max D$ is nonempty, so $\alpha_j(a)$ is a nonempty subset of \mathbf{H}_j^X , and hence in E_j .

Step 2. We show that α_j is Scott-continuous. First note that α_j is monotone: If $a \leq a'$ in D , then $\uparrow a \cap \max D \supseteq \uparrow a' \cap \max D$, so $\alpha_j(a) \supseteq \alpha_j(a')$, i.e., since E_j is ordered by \supseteq , $\alpha_j(a) \leq \alpha_j(a')$.

Now, let $A \subseteq D$ be directed, and show $\alpha_j(\bigvee A) \leq \bigvee \alpha_j(A)$ (the other direction, \geq , follows from monotonicity). Since E_j is ordered by \supseteq , we have

$$\begin{aligned} \bigvee \alpha_j(A) &= \bigvee \left\{ \{\mathcal{O}_j(x) : x \in \uparrow a \cap \max D\} : a \in A \right\} \\ &= \bigcap_{a \in A} \{\mathcal{O}_j(x) : x \in \uparrow a \cap \max D\} \\ \alpha_j(\bigvee A) &= \{\mathcal{O}_j(x) : x \in \uparrow \bigvee A \cap \max D\}. \end{aligned}$$

For $a \in A$, write $M_a := \{\mathcal{O}_j(x) : x \in \uparrow a \cap \max D\}$. Also write $M := \bigcap_{a \in A} M_a$ and $N := \{\mathcal{O}_j(x) : x \in \uparrow \bigvee A \cap \max D\}$. Note that $M, N \subseteq \mathbf{H}_j^X$. We need to show $M \subseteq N$. So let $\mathcal{O}_j(y) \in \mathbf{H}_j^X$ (for some $y \in \max D$) with $\mathcal{O}_j(y) \in M$ and show $\mathcal{O}_j(y) \in N$. Consider

$$\mathcal{F} := \{[y]_j \cap (\uparrow a \cap \max D) : a \in A\}.$$

We show that \mathcal{F} is a family of closed subsets of $\max D$ with the finite intersection property.

Closed: We know, from the construction of $\hat{D}(\mathfrak{X})$, that $[y]_j \subseteq X = \max D$ is clopen (see section 5.2.5). Moreover, $\uparrow a \subseteq D$ is closed in the Lawson topology on D , so $\uparrow a \cap \max D$ is closed in the relative Lawson topology on $\max D$, which, since D is a Scott domain, coincides with the relative Scott topology. So $\uparrow a \cap \max D$ is a closed subset of $\max D$.

Finite intersection property: Since A is directed, it suffices to show (a) if $a \leq a'$ in A , then $[y]_j \cap (\uparrow a \cap \max D) \supseteq [y]_j \cap (\uparrow a' \cap \max D)$, and (b) for $a \in A$, we have $[y]_j \cap (\uparrow a \cap \max D) \neq \emptyset$. Concerning (a), if $a \leq a'$, then $\uparrow a \supseteq \uparrow a'$, so the claim follows. Concerning (b), since, by assumption, $\mathcal{O}_j(y) \in M$, we have $\mathcal{O}_j(y) \in M_a$, so there is $x \in \uparrow a \cap \max D$ with $\mathcal{O}_j(y) = \mathcal{O}_j(x)$. So $x \in [y]_j \cap (\uparrow a \cap \max D)$.

Since $\max D$ is compact, there is $x_0 \in \bigcap \mathcal{F}$. In particular (since $A \neq \emptyset$ qua directed set), we have $x_0 \in [y]_j$, so $\mathcal{O}_j(y) = \mathcal{O}_j(x_0)$. And, for $a \in A$, we have $x_0 \in \uparrow a \cap \max D$, so $x_0 \geq a$. Hence, $x_0 \geq \bigvee A$. Thus, $\mathcal{O}_j(y) = \mathcal{O}_j(x_0) \in N$, as needed.

Step 3. We show that $(D, \alpha_j)_{j \in I(\text{Clp}X)}$ is a cone (in the category of dcpos with Scott-continuous maps) to the expanding system of dcpos $(E_j, q_{ij})_{j \in I(\text{Clp}X)}$. Indeed, for $i \leq j$ in $I(\text{Clp}X)$, we have, for any $a \in D$,

$$q_{ij} \circ \alpha_j(a) = q_{ij} \{\mathcal{O}_j(x) : x \in \uparrow a \cap \max D\} = \{\mathcal{O}_i(x) : x \in \uparrow a \cap \max D\} = \alpha_i(a).$$

Thus, we can define the Scott-continuous function⁴⁴

$$\begin{aligned} \alpha : D &\rightarrow E \\ a &\mapsto \langle \alpha_j(a) : j \in I(\text{Clp}X) \rangle. \end{aligned}$$

⁴⁴We've used this argument already in the proof of proposition 5.4.2, see footnote 40.

Step 4. We show that, for $x \in \max D = X$, we have $\alpha(x) = \varphi_X(x)$ where $\varphi_X : X \rightarrow \max E$ is the canonical homeomorphism. This, in particular, implies that α is max-preserving.

Indeed, if $x \in \max D$, then $\uparrow x \cap \max D = \{x\}$, so, for any $j \in I(\mathbf{Clp}X)$,

$$\alpha(x)(j) = \alpha_j(x) = \{\mathcal{O}_j(y) : y \in \uparrow x \cap \max D\} = \{\mathcal{O}_j(x)\} = \varphi_X(x)(j).$$

Step 5. We show that α is max-bisimulative: Let $a \in D$ and $\alpha(a) \leq e \in \max E$, then $d := \varphi_X^{-1}(e) \in X = \max D$ and, by step 4, $\alpha(d) = \varphi_X(d) = e$, so it remains to show $d \geq a$. Since $\alpha(a) \leq e$, we have, for any $j \in I(\mathbf{Clp}X)$,

$$\{\mathcal{O}_j(x) : x \in \uparrow a \cap \max D\} = \alpha_j(a) = \alpha(a)(j) \leq e(j) = \varphi_X(d)(j) = \{\mathcal{O}_j(d)\}.$$

Since \leq is reverse inclusion, there hence is, for every $j \in I(\mathbf{Clp}X) = I(\mathbf{Clp} \max D)$, some $x_j \in \uparrow a \cap \max D$ with $\mathcal{O}_j(d) = \mathcal{O}_j(x_j)$. Now, lemma 5.5.2 implies (with $x := d$) that $d \geq a$, as needed.

Step 6. We show that, for all $V \in \Sigma(E)$, $w(V) = v(\alpha^{-1}(V))$. Recall that $\mu = \mu_v \upharpoonright \mathcal{B}(\max D)$ is the measure on $\mathfrak{X} = \mathbf{S}(\mathfrak{D})$ where μ_v is the unique probability measure on $\mathcal{B}(D, \Lambda)$ determined by v . Also, the valuation w of $\mathfrak{E} = \hat{\mathbf{D}}(\mathfrak{X})$ determines a unique probability measure μ_w on $\mathcal{B}(E, \Lambda)$ and $\kappa := \mu_w \upharpoonright \mathcal{B}(\max E)$ is the measure on $\mathbf{SD}(\mathfrak{X})$.

Since $\varphi_X : \mathfrak{X} \rightarrow \mathbf{SD}(\mathfrak{X})$ is measure-preserving, we have

$$w(V) = \mu_w(V) = \mu_w(V \cap \max E) = \kappa(V \cap \max E) = \mu \varphi_X^{-1}(V \cap \max E)$$

Moreover, by step 4, $\varphi_X^{-1}(V \cap \max E) = \alpha^{-1}(V) \cap \max D$. Hence

$$\mu \varphi_X^{-1}(V \cap \max E) = \mu(\alpha^{-1}(V) \cap \max D).$$

Furthermore, by the properties of μ_v , we have, for $U \in \Sigma(D)$, that $v(U) = \mu_v(U) = \mu(U \cap \max D)$. Since $V \in \Sigma(E)$ and α is Scott-continuous, $\alpha^{-1}(V) \in \Sigma(D)$, so

$$\mu(\alpha^{-1}(V) \cap \max D) = v(\alpha^{-1}(V)).$$

Putting everything together, we get $w(V) = v(\alpha^{-1}(V))$, as needed.

Step 7. We show that for all $x \in \max D$, $\alpha(f(x)) = g(\alpha(x))$. Since $x \in \max D$ and f is max-preserving, we have by step 4, $\alpha(f(x)) = \varphi_X(f(x))$. Recall that $T = f \upharpoonright \max D : X \rightarrow X$, and the dynamics on $\mathbf{SD}(\mathfrak{X}) = \mathbf{S}(\mathfrak{E})$ is given by $g \upharpoonright \max E$. Since φ_X is equivariant, we have

$$\alpha(f(x)) = \varphi_X(f(x)) = \varphi_X(T(x)) = (g \upharpoonright \max E)(\varphi_X(x)) = g(\alpha(x)),$$

as needed.

In sum, $\alpha : \mathfrak{D} \rightarrow \mathfrak{E} = \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D})$ is a Scott-continuous, max-preserving, max-bisimulative, valuation-preserving, and max-equivariant function, whence a morphism in \mathbf{dDOM} , which also coincides on $\max D$ with $\varphi_{\max D}$. \square

5.5.5. PROPOSITION. *The family of morphisms $(\eta_{\mathfrak{D}})$ from proposition 5.5.4 constitutes a natural transformation $\eta : 1_{\mathbf{dDOM}} \rightarrow \hat{\mathbf{DS}}$.*

PROOF. We need to show that the morphisms $\eta_{\mathfrak{D}}$ are natural in \mathfrak{D} . So let $\beta : \mathfrak{D} \rightarrow \mathfrak{E}$ be a morphism in \mathbf{dDOM} . Write $\mathfrak{D} = (D, v, f)$ and $\mathfrak{E} = (E, w, g)$, and show, for $a \in D$,

$$\eta_{\mathfrak{E}} \circ 1_{\mathbf{dDOM}}(\beta)(a) = \hat{\mathbf{DS}}(\beta) \circ \eta_{\mathfrak{D}}(a).$$

Both sides are elements of $\hat{\mathbf{DS}}(\mathfrak{E})$ which is the limit of $(\mathbf{P}(\mathbf{H}_i^{\max E}), p_{ij})_{I(\mathbf{Clp}(\max E))}$ (omitting the valuation and dynamics on $\mathbf{P}(\mathbf{H}_i^{\max E})$). So let $j \in I(\mathbf{Clp}(\max E))$, and show

$$\eta_{\mathfrak{E}} \circ \beta(a)(j) = \hat{\mathbf{DS}}(\beta) \circ \eta_{\mathfrak{D}}(a)(j).$$

We'll write $\eta^{\mathfrak{D}} := \eta_{\mathfrak{D}}$ and $\eta^{\mathfrak{E}} := \eta_{\mathfrak{E}}$, so we can use subscripts like $\eta_j^{\mathfrak{D}}$ for the components of $\eta_{\mathfrak{D}}$. Then, on the left side, we have

$$\eta_{\mathfrak{E}} \circ \beta(a)(j) = \eta_j^{\mathfrak{E}}(\beta(a)) = \{\mathcal{O}_j(y) : y \in \uparrow\beta(a) \cap \max E\}.$$

On the right side, write $a' := \eta^{\mathfrak{D}}(a) = \langle \eta_i^{\mathfrak{D}}(a) : i \in I(\mathbf{Clp}(\max D)) \rangle$, which is in $\hat{\mathbf{DS}}(\mathfrak{D})$. Also write $\varphi := \mathbf{S}(\beta) = \beta \upharpoonright \max D : \max D \rightarrow \max E$ and $j = (m, \mathcal{D})$. So

$$\begin{aligned} \hat{\mathbf{DS}}(\beta) \circ \eta_{\mathfrak{D}}(a)(j) &= \hat{\mathbf{D}}(\varphi)(a')(j) \\ &= \pi_{(m, \mathcal{D})}(a'(m, \varphi^{-1}\mathcal{D})) \\ &= \pi_{(m, \mathcal{D})}(\eta_{(m, \varphi^{-1}\mathcal{D})}^{\mathfrak{D}}(a)) \\ &= \pi_{(m, \mathcal{D})}\{\mathcal{O}_{(m, \varphi^{-1}\mathcal{D})}(x) : x \in \uparrow a \cap \max D\} \\ &= \{\mathcal{O}_{(m, \mathcal{D})}(\varphi(x)) : x \in \uparrow a \cap \max D\} \\ &= \{\mathcal{O}_j(\beta(x)) : x \in \uparrow a \cap \max D\}. \end{aligned}$$

So we need to show

$$\{\mathcal{O}_j(y) : y \in \uparrow\beta(a) \cap \max E\} = \{\mathcal{O}_j(\beta(x)) : x \in \uparrow a \cap \max D\}.$$

(\supseteq) Given $\mathcal{O}_j(\beta(x))$ from the right set with $x \in \uparrow a \cap \max D$, we have that $y := \beta(x)$ is maximal in E (since β is max-preserving) and $\geq \beta(a)$ (since $x \geq a$ and β is monotone). Hence $\mathcal{O}_j(\beta(x)) = \mathcal{O}_j(y)$ which is in the left set.

(\subseteq) Given $\mathcal{O}_j(y)$ with $y \in \uparrow\beta(a) \cap \max E$, we have, since β is max-bisimulative that there is $d \in \max D$ such that $d \geq a$ and $\beta(d) = y$. Hence $x := d \in \uparrow a \cap \max D$ with $\beta(x) = y$, so $\mathcal{O}_j(y) = \mathcal{O}_j(\beta(x))$, which is in the right set. \square

Note that in the \subseteq -direction we have crucially used that morphisms in \mathbf{dDOM} are max-bisimulative.

5.5.2 Triangle identities

We need one more observation. Then the triangle identities are immediate.

5.5.6. LEMMA. *Let \mathfrak{X} be in TS_{0c} , write $\hat{D}(\mathfrak{X}) = (D, v, f)$, and let $\varphi_X : X \rightarrow \max D$ be the canonical homeomorphism. Then, for $a \in D$ and $i \in I(\text{Clp}X)$, we have*

$$\{\mathcal{O}_i(\varphi_X^{-1}(z)) : z \in \uparrow a \cap \max D\} = a(i).$$

PROOF. Write $\varphi := \varphi_X$, and write p_i for the Scott-continuous projections of $\hat{D}(\mathfrak{X})$. Denote the left set A . Both A and $a(i)$ are subsets of \mathbf{H}_i^X . So let $\mathcal{O}_i(x) \in \mathbf{H}_i^X$ and show $\mathcal{O}_i(x) \in A$ iff $\mathcal{O}_i(x) \in a(i)$.

(\Rightarrow). Assume $\mathcal{O}_i(x) \in A$. Then $\mathcal{O}_i(x) = \mathcal{O}_i(\varphi^{-1}(z))$ for some $z \in \uparrow a \cap \max D$. By definition of φ ,

$$\{\mathcal{O}_i(\varphi^{-1}(z))\} = \varphi(\varphi^{-1}(z))(i) = z(i) \geq a(i).$$

Since the order on D is reverse inclusion, we have

$$\mathcal{O}_i(x) \in \{\mathcal{O}_i(\varphi^{-1}(z))\} \subseteq a(i),$$

as needed.

(\Leftarrow). Assume $\mathcal{O}_i(x) \in a(i)$ (note that $x \in X$ and $i \in I(\text{Clp}X)$). For each $j \in I(\text{Clp}X)$, define the following subset of X :

$$F_j := [x]_i \cap \{x' \in X : \varphi(x')(j) \geq a(j)\}.$$

We claim that $\{F_j : j \in I(\text{Clp}X)\}$ forms a family of closed subsets of X with the finite intersection property. This then finishes the proof: Since X is compact, let $x_0 \in \bigcap_{j \in I(\text{Clp}X)} F_j$. Define $z := \varphi(x_0) \in \max D$. Then $z \geq a$ since, for $j \in I(B_X)$, $z(j) = \varphi(x_0)(j) \geq a(j)$ since $x_0 \in F_j$. So $z \in \uparrow a \cap \max D$. Further, $\mathcal{O}_i(x_0) = \mathcal{O}_i(x)$, because $x_0 \in [x]_i$, since $x_0 \in F_i$. Hence $\mathcal{O}_i(x) = \mathcal{O}_i(x_0) = \mathcal{O}_i(\varphi^{-1}(z)) \in A$.

Closed: The equivalence classes $[x]_i$ are clopen subsets of X and

$$\begin{aligned} \{x' \in X : \varphi(x')(j) \geq a(j)\} &= \varphi^{-1}\{z \in \max D : z(j) \geq a(j)\} \\ &= \varphi^{-1}(\max D \cap p_j^{-1}(\uparrow a(j))) \end{aligned}$$

is a closed subset of X since $\varphi : X \rightarrow \max D$ is continuous and $\max D \cap p_j^{-1}(\uparrow a(j))$ is, as we'll argue now, a closed subset of $\max D$. Indeed, $\uparrow a(j)$ is a subset of the finite $\mathbf{P}(\mathbf{H}_j^X)$, so it is Lawson-closed (since the Lawson topology on a finite domain is discrete). Hence $p_j^{-1}(\uparrow a(j))$ is a Lawson-closed subset of D (since p_j is a Scott-continuous projection and hence Lawson-continuous).⁴⁵ Since the

⁴⁵We've made this argument already in footnote 41.

relative Lawson topology on $\max D$ coincides with the relative Scott topology, $p_j^{-1}(\uparrow a(j)) \cap \max D$ is closed.

Monotone: Let $j \leq k$ in $I(\mathbf{Clp}X)$ and show $F_j \supseteq F_k$. It suffices to show $\{x' \in X : \varphi(x')(j) \geq a(j)\} \supseteq \{x' \in X : \varphi(x')(k) \geq a(k)\}$. So let $x' \in X$ with $\varphi(x')(k) \geq a(k)$. By monotonicity of p_{jk} , we have

$$p_{jk}(\varphi(x')(k)) \geq p_{jk}(a(k)) = a(j).$$

Moreover,

$$p_{jk}(\varphi(x')(k)) = p_{jk}(\{\mathcal{O}_k(x')\}) = \{\mathcal{O}_j(x')\} = \varphi(x')(j).$$

Hence $x' \in \{x' \in X : \varphi(x')(j) \geq a(j)\}$, as needed.

Nonempty: We show that F_j is nonempty. Since $i, j \in I(\mathbf{Clp}X)$ and $I(\mathbf{Clp}X)$ is directed, there is $k \in I(\mathbf{Clp}X)$ such that $i, j \leq k$. We have

$$\mathcal{O}_i(x) \in a(i) = p_{ik}(a(k)) = \{\mathcal{O}_i(x') : \mathcal{O}_k(x') \in a(k)\}$$

So there is $x' \in X$ with $\mathcal{O}_k(x') \in a(k)$ and $\mathcal{O}_i(x) = \mathcal{O}_i(x')$. So $x' \in [x]_i$. And

$$\mathcal{O}_j(x') \in \{\mathcal{O}_j(x') : \mathcal{O}_k(x') \in a(k)\} = p_{jk}(a(k)) = a(j),$$

whence $\{\mathcal{O}_j(x')\} \subseteq a(j)$. Since \leq is reversed inclusion, $\varphi(x')(j) = \{\mathcal{O}_j(x')\} \geq a(j)$. Hence $x' \in F_j$. \square

As promised, the triangle identities now are immediate:

5.5.7. PROPOSITION. *For \mathfrak{X} in \mathbf{TS}_{0c} and \mathfrak{D} in \mathbf{dDOM} , we have*

$$1. \ 1_{\mathbf{S}(\mathfrak{D})} = \epsilon_{\mathbf{S}(\mathfrak{D})} \circ \mathbf{S}(\eta_{\mathfrak{D}}).$$

$$2. \ 1_{\hat{\mathbf{D}}(\mathfrak{X})} = \hat{\mathbf{D}}(\epsilon_{\mathfrak{X}}) \circ \eta_{\hat{\mathbf{D}}(\mathfrak{X})}.$$

PROOF. Write $\mathfrak{D} = (D, v, f)$ and write $\hat{\mathbf{D}}(\mathfrak{X}) = (E, w, g)$, which is obtained as a limit over the index set $I(\mathbf{Clp}X)$.

Ad (1). We have $\mathbf{S}(\eta_{\mathfrak{D}}) = \eta_{\mathfrak{D}} \upharpoonright \max D$ which, by proposition 5.5.4, equals $\varphi_{\max D}$. Moreover, $\epsilon_{\mathbf{S}(\mathfrak{D})} = \varphi_{\max D}^{-1}$. So the claim follows since $\varphi_{\max D}$ is a homeomorphism.

Ad (2). Write $\epsilon := \epsilon_{\mathfrak{X}} = \varphi_X^{-1} : \max E \rightarrow X$. For $a \in E$ and $i \in I(\mathbf{Clp}X)$, we need to show

$$\hat{\mathbf{D}}(\epsilon) \circ \eta_{\hat{\mathbf{D}}(\mathfrak{X})}(a)(i) = a(i).$$

Indeed, we have the following. (Recall $\xi^\epsilon : I(\mathbf{Clp}X) \rightarrow I(\mathbf{Clp} \max E)$ maps (n, \mathcal{C}) to $(n, \epsilon^{-1}(\mathcal{C}))$.)

$$\begin{aligned}
\hat{D}(\epsilon) \circ \eta_{\hat{D}(\mathfrak{x})}(a)(i) &= \hat{D}(\epsilon) \left(\eta_{\hat{D}(\mathfrak{x})}(a) \right)(i) \\
&= \pi_i^\epsilon \left(\eta_{\hat{D}(\mathfrak{x})}(a)(\xi^\epsilon(i)) \right) \\
&= \pi_i^\epsilon \left(\eta_{\xi^\epsilon(i)}^{\hat{D}(\mathfrak{x})}(a) \right) \\
&= \pi_i^\epsilon \left\{ \mathcal{O}_{\xi^\epsilon(i)}(x) : x \in \uparrow a \cap \max E \right\} \\
&= \left\{ \mathcal{O}_i(\epsilon(x)) : x \in \uparrow a \cap \max E \right\}
\end{aligned}$$

By lemma 5.5.6, this further equals $a(i)$, as needed. \square

5.6 Analyzing the systems-domains adjunction

In this section, we establish the following part of the main diagram:

$$\begin{array}{ccc}
\mathbf{dDOM} & \begin{array}{c} \xleftarrow{\mathbb{I}} \\ \xrightarrow{\mathbb{I}} \\ \hat{D} \circ \mathbb{S} \end{array} & \mathbf{dDOM}_r \\
& \searrow \hat{D} & \uparrow \hat{D} \\
& \searrow \mathbb{S} & \mathbf{TS}_{0c}
\end{array}$$

(That it commutes and restricts to the standard case is shown in section 5.7.)

Thus, we analyze the adjunction $\mathbb{S} : \mathbf{dDOM} \rightleftarrows \mathbf{TS}_{0c} : \hat{D}$ into (a) an equivalence between \mathbf{TS}_{0c} and \mathbf{dDOM}_r (section 5.6.1), and (b) a reflective subcategory \mathbf{dDOM}_r of \mathbf{dDOM} (section 5.6.2).

5.6.1 Restricting to equivalence

Now, we can see that, surprisingly, the simple and purely domain-theoretic property of being max-reflective (definition 5.2.19) turns out to be precisely what is needed to turn the adjunction between systems and domains into an equivalence.

5.6.1. PROPOSITION. *Let \mathfrak{D} be in \mathbf{dDOM} . The following are equivalent.*

1. \mathfrak{D} is isomorphic in \mathbf{dDOM} to $\hat{D}(\mathfrak{x})$ for some system \mathfrak{x} in \mathbf{TS}_{0c} .
2. \mathfrak{D} is max-reflective, i.e., \mathfrak{D} is in \mathbf{dDOM}_r .

3. $\eta_{\mathfrak{D}} : \mathfrak{D} \rightarrow \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D})$ is an isomorphism in \mathbf{dDOM} .

PROOF. Concerning the implication (1) \Rightarrow (2), since \mathfrak{D} is isomorphic in \mathbf{dDOM} to $\hat{\mathbf{D}}(\mathfrak{X})$, their underlying domains are isomorphic (proposition 5.2.13). Since being max-reflective is a purely domain-theoretic property, it hence suffices to show that $\hat{\mathbf{D}}(\mathfrak{X})$ is max-reflective. Indeed, it is, by definition, a $\mathbf{dSCO}_{\text{nm}}^{\text{P}}$ -limit of the finitary dynamical expanding system $(\mathfrak{D}_i, p_{ij})_{I(\text{Clp } X)}$ with underlying domains $D_i = \mathbf{P}(\mathbf{H}_i^X)$. By example 5.2.20 (2), each \mathfrak{D}_i hence is max-reflective (since \mathbf{H}_i^X is a finite set), so, by proposition 5.2.21, also the limit $\mathfrak{D} = \hat{\mathbf{D}}(\mathfrak{X})$ is max-reflective.

The implication (3) \Rightarrow (1) is immediate when taking $\mathfrak{X} := \mathbf{S}(\mathfrak{D})$.

So it remains implication (2) \Rightarrow (3). Write $\alpha := \eta_{\mathfrak{D}} : \mathfrak{D} \rightarrow \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D})$ and $\mathfrak{D} = (D, v, f)$. To show that α is an isomorphism in \mathbf{dDOM} , it suffices to show, by proposition 5.2.13, that α is order-reflecting and surjective (since α is Scott-continuous, it already is montone, and since α is semi-max-equivariant and its codomain $\hat{\mathbf{D}}\mathbf{S}(\mathfrak{D})$ max-preserving, it already is max-equivariant).

Order-reflecting: Let $a, a' \in D$ with $\alpha(a) \leq \alpha(a')$ and show $a \leq a'$. We'll show

$$\uparrow a \cap \max D \supseteq \uparrow a' \cap \max D. \quad (5.2)$$

By condition 1 of being max-reflective (definition 5.2.19), we'll then have

$$a = \bigwedge \uparrow a \cap \max D \leq \bigwedge \uparrow a' \cap \max D = a',$$

as needed.

To show (5.2), let $z \in \uparrow a' \cap \max D$ and show $z \geq a$ (so $z \in \uparrow a \cap \max D$). Since $\alpha(a) \leq \alpha(a')$, we have, for each $j \in I(\text{Clp } \max D)$, that $\alpha_j(a) \leq \alpha_j(a')$, so, since \leq is reverse inclusion,

$$\begin{aligned} \{\mathcal{O}_j(x) : x \in \uparrow a \cap \max D\} &= \alpha_j(a) \supseteq \alpha_j(a') \\ &= \{\mathcal{O}_j(x) : x \in \uparrow a' \cap \max D\} \ni \mathcal{O}_j(z). \end{aligned}$$

Hence, for each $j \in I(\text{Clp } \max D)$, there is $x_j \in \uparrow a \cap \max D$ with $\mathcal{O}_j(z) = \mathcal{O}_j(x_j)$. So lemma 5.5.2 implies $z \geq a$, as needed.

Surjective: Write $\hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) =: \mathfrak{E} = (E, w, g)$ as the limit of $(\mathfrak{E}_j, q_{ij})_J$ with $J := I(\text{Clp } \max D)$, so $E_j = \mathbf{P}(\mathbf{H}_j^{\max D})$. Recall that $\alpha : D \rightarrow E$ is defined, on each component $j \in J$, by $\alpha_j(a) = \{\mathcal{O}_j(x) : x \in \uparrow a \cap \max D\}$. Also, as seen in the proof of proposition 5.5.4, $(D, \alpha_j)_{j \in J}$ is a cone (in the category of dcpos with Scott-continuous maps) to the expanding system of dcpos $(E_j, q_{ij})_J$. We show that each α_j is a projection. It then follows that the mediating map α also is a projection and hence surjective (see Abramsky and Jung 1994, cor. 3.3.10).⁴⁶ (This is not a detour: the hard part is not to show that α_j is an upper adjoint, but to show that it is surjective.)

⁴⁶This can also be shown directly: Write $e_j : E_j \rightarrow D$ for the embedding corresponding to α_j . Define $e : E \rightarrow D$ by $e(b) := \bigvee \{e_j \circ q_j(b) : j \in J\}$. This is directed since $J \neq \emptyset$ and, for $j \leq k$,

So let $j \in J$. Consider the function $e_j : E_j \rightarrow D$ which maps $M \in E_j = \mathbf{P}(\mathbf{H}_j^{\max D})$ to

$$e_j(M) := \bigwedge \{x \in \max D : \mathcal{O}_j(x) \in M\}.$$

This is well-defined since M is nonempty, so it contains some $\mathcal{O}_j(x)$, so the infimum is taken over a nonempty subset of the Scott domain D and hence exists. It also is monotone: if $M \leq M'$, then $M \supseteq M'$, so

$$\{x \in \max D : \mathcal{O}_j(x) \in M\} \supseteq \{x \in \max D : \mathcal{O}_j(x) \in M'\},$$

so the infimum of the left set will be \leq the infimum of the right set, i.e., $e_j(M) \leq e_j(M')$.

To show that e_j is an embedding to the (continuous) projection α_j , we need to show (i) $\alpha_j \circ e_j = \text{id}_{E_j}$ and (ii) $e_j \circ \alpha_j \leq \text{id}_D$.

Concerning (i), let $M \in E_j$ and show $\alpha_j \circ e_j(M) = M$. Define $A := \{x \in \max D : \mathcal{O}_j(x) \in M\}$.

We claim that A is a nonempty closed subset of $\max D$. Indeed, by construction, A is a subset of $\max D$, and it is nonempty since M is nonempty. To see that it is closed, note that, since $M \subseteq \mathbf{H}_j^{\max D}$ is nonempty, we have $M = \{\mathcal{O}_j(x_1), \dots, \mathcal{O}_j(x_n)\}$ for some $x_1, \dots, x_n \in \max D$ (with $n \geq 1$). Now,

$$A = \{x \in \max D : \exists k \in \{1, \dots, n\}. \mathcal{O}_j(x) = \mathcal{O}_j(x_k)\} = \bigcup_{k=1}^n [x_k]_j$$

and the latter is a finite union of clopen subsets of $\max D$ and hence closed.

Now we can apply condition 2 of being max-reflective (definition 5.2.19) and obtain $A = (\uparrow \bigwedge A) \cap \max D$, whence

$$\begin{aligned} \alpha_j \circ e_j(M) &= \alpha_j(\bigwedge A) = \{\mathcal{O}_j(x) : x \in \uparrow \bigwedge A \cap \max D\} \\ &= \{\mathcal{O}_j(x) : x \in A\} = M. \end{aligned}$$

Concerning (ii), let $a \in D$ and show $e_j \circ \alpha_j(a) \leq a$. We have

$$\begin{aligned} e_j \circ \alpha_j(a) &= \bigwedge \{x \in \max D : \mathcal{O}_j(x) \in \alpha_j(a)\} \\ &= \bigwedge \{x \in \max D : \mathcal{O}_j(x) = \mathcal{O}_j(x') \text{ for } x' \in \uparrow a \cap \max D\} \end{aligned}$$

we have $e_j \circ q_j(b) \leq e_k \circ q_k(b)$: indeed, we have

$$e_j \circ q_j = e_j \circ q_{jk} \circ q_k = e_j \circ q_{jk} \circ \alpha_k \circ e_k \circ q_k = e_j \circ \alpha_j \circ e_k \circ q_k \leq e_k \circ q_k.$$

We have $e \circ \alpha \leq \text{id}_D$: For $a \in D$, $e \circ \alpha(a) = \bigvee \{e_j \circ q_j(\alpha(a)) : j \in J\} \leq a$ since $e_j \circ q_j(\alpha(a)) = e_j \circ \alpha_j(a) \leq a$. And we have $\alpha \circ e = \text{id}_E$: Since α preserves suprema, $\alpha \circ e(b) = \bigvee \{\alpha \circ e_j \circ q_j(b) : j \in J\}$ and we show that this equals b . Indeed, b is an upper bound: Given $j \in J$, we have, for any $k \geq j$: $\alpha \circ e_j \circ q_j(b)(k) = \alpha_k(e_j \circ q_j(b)) \leq \alpha_k(e_k \circ q_k(b)) = q_k(b) = b(k)$. (So this holds for any $k' \in J$: pick $k \geq j, k'$, then $\alpha \circ e_j \circ q_j(b)(k') = q_{k'k}(\alpha_k \circ e_j \circ q_j(b)(k)) \leq q_{k'k}(b(k)) = b(k')$.) And b is a least upper bound: If b' is another upper bound, then, for any $j \in J$, we have $b(j) = \alpha_j \circ e_j \circ q_j(b) = \alpha \circ e_j \circ q_j(b)(j) \leq b'(j)$, whence $b \leq b'$.

Note that the set in the last line is a superset of $\uparrow a \cap \max D$. So the last line can be continued with

$$\leq \bigwedge \uparrow a \cap \max D,$$

which, in turn, equals a by condition 1 of being max-reflective (definition 5.2.19). So $e_j \circ \alpha_j(a) \leq a$, as needed. \square

Now, it is a purely categorical matter to obtain the promised result:

5.6.2. THEOREM. *We have the following diagram*

$$\begin{array}{ccc}
 & \text{dDOM}_r & \\
 \hat{D}S \nearrow & & \searrow S \\
 \text{dDOM} & & \text{TS}_{0c} \\
 \downarrow \perp & & \downarrow \hat{D} \\
 \text{dDOM} & \xrightarrow{S} & \text{TS}_{0c} \\
 \downarrow \hat{D} & & \downarrow \hat{D}
 \end{array}$$

where \sim denotes adjoint equivalence.

PROOF. Write $\mathbf{C} := \text{TS}_{0c}$ and $\mathbf{D} := \text{dDOM}$. So the adjunction $S \dashv \hat{D}$ has unit $\eta : 1_{\mathbf{D}} \rightarrow \hat{D}S$ and counit $\epsilon : S\hat{D} \rightarrow 1_{\mathbf{C}}$. By proposition 5.6.1, the full subcategory $\mathbf{D}' := \text{dDOM}_r$ of \mathbf{D} consists of precisely those objects \mathfrak{D} of \mathbf{D} such that $\eta_{\mathfrak{D}}$ is an isomorphism. Since the counit ϵ is a natural isomorphism, every object \mathfrak{X} of \mathbf{C} is such that $\epsilon_{\mathfrak{X}}$ is an isomorphism. Now, it is a well-known fact that an adjunction $(S, \hat{D}, \epsilon, \eta)$ restricts to an adjoint equivalence $(S \upharpoonright \mathbf{D}', \hat{D} \upharpoonright \mathbf{D}', \epsilon \upharpoonright \mathbf{D}', \eta \upharpoonright \mathbf{D}')$ between the fixed points of the adjunction (Leinster 2014, ex. 2.2.11).⁴⁷ Hence $\hat{D} \sim S$. \square

5.6.2 Max-reflecting a dynamical domain

Finally, we show that the functor $\hat{D}S : \text{dDOM} \rightarrow \text{dDOM}_r$ is the *optimal* way of making a dynamical \mathfrak{D} max-reflective: we show that $\hat{D}S$ is left-adjoint to the inclusion $I : \text{dDOM}_r \rightarrow \text{dDOM}$. Note that this question is somewhat analogous to the question of how to compactify a system in \mathbf{bTS}_0 .

We first make the important observation that morphisms in dDOM_r are entirely determined by their behavior on maximal elements. (In fact, that was the guiding intuition in isolating dDOM_r .)

⁴⁷This is also discussed in the *nLab* entry <https://ncatlab.org/nlab/show/fixed+point+of+an+adjunction> (last checked 3 July 2021) and in the *stack exchange* thread(s) <https://math.stackexchange.com/questions/3359118/adjunctions-restrict-to-an-equivalence> (last checked 3 July 2021).

5.6.3. LEMMA. *Let $\beta, \beta' : \mathfrak{D} \rightarrow \mathfrak{E}$ be two morphisms in \mathbf{dDOM}_r . If β and β' agree on the maximal elements of the domain underlying \mathfrak{D} , then $\beta = \beta'$.*

PROOF. Write $\mathfrak{D} = (D, v, f)$ and $\mathfrak{E} = (E, w, g)$. Let $a \in D$ and show $\beta(a) = \beta'(a)$. It suffices to show that

$$\uparrow\beta(a) \cap \max E = \uparrow\beta'(a) \cap \max E,$$

since this implies, because E is max-reflective, that $\beta(a) = \bigwedge \uparrow\beta(a) \cap \max E = \bigwedge \uparrow\beta'(a) \cap \max E = \beta'(a)$.

Indeed, let $y \in \uparrow\beta(a) \cap \max E$ and show $y \geq \beta'(a)$ (the other direction is analogous). Since β is max-bisimulative, there is $x \in \max D$ with $x \geq a$ and $\beta(x) = y$. Since β and β' agree on $\max D$, we have $y = \beta(x) = \beta'(x)$. Since $a \leq x$ and β' is monotone, $\beta'(a) \leq \beta'(x) = y$, as needed. \square

5.6.4. PROPOSITION. *The inclusion $\mathbf{l} : \mathbf{dDOM}_r \rightarrow \mathbf{dDOM}$ is a right adjoint: For every \mathfrak{D} in \mathbf{dDOM} , the morphism $\eta_{\mathfrak{D}} : \mathfrak{D} \rightarrow \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D})$ is such that, for any \mathfrak{E} in \mathbf{dDOM}_r and $\beta : \mathfrak{D} \rightarrow \mathfrak{E}$ in \mathbf{dDOM} , there is a unique $\alpha : \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) \rightarrow \mathfrak{E}$ with $\alpha \circ \eta_{\mathfrak{D}} = \beta$.*

$$\begin{array}{ccc} \mathfrak{D} & \xrightarrow{\eta_{\mathfrak{D}}} & \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) \\ & \searrow \beta & \downarrow \exists! \alpha \\ & & \mathfrak{E} \end{array}$$

PROOF. Existence: Since \mathfrak{E} is in \mathbf{dDOM}_r , $\eta_{\mathfrak{E}} : \mathfrak{E} \rightarrow \hat{\mathbf{D}}\mathbf{S}(\mathfrak{E})$ is an isomorphism. So we can define the morphism $\alpha := \eta_{\mathfrak{E}}^{-1} \circ \hat{\mathbf{D}}\mathbf{S}(\beta) : \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) \rightarrow \mathfrak{E}$ in \mathbf{dDOM}_r . By the naturality of $\eta : \mathbf{1}_{\mathbf{dDOM}} \rightarrow \hat{\mathbf{D}}\mathbf{S}$, we have $\eta_{\mathfrak{E}} \circ \beta = \hat{\mathbf{D}}\mathbf{S}(\beta) \circ \eta_{\mathfrak{D}}$. So

$$\alpha \circ \eta_{\mathfrak{D}} = \eta_{\mathfrak{E}}^{-1} \circ \hat{\mathbf{D}}\mathbf{S}(\beta) \circ \eta_{\mathfrak{D}} = \eta_{\mathfrak{E}}^{-1} \circ \eta_{\mathfrak{E}} \circ \beta = \beta.$$

Uniqueness: Assume $\alpha, \alpha' : \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) \rightarrow \mathfrak{E}$ are such that $\alpha \circ \eta_{\mathfrak{D}} = \beta = \alpha' \circ \eta_{\mathfrak{D}}$. Write $\mathfrak{D} = (D, v, f)$ and $\hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) = (D', v', f')$. So α and α' agree on $\eta_{\mathfrak{D}}(D)$. We claim $\max D' \subseteq \eta_{\mathfrak{D}}(D)$. Indeed, by proposition 5.5.4,

$$\eta_{\mathfrak{D}} \upharpoonright \max D = \varphi_{\max D} : \max D \rightarrow \max D',$$

is the canonical homeomorphism $\mathbf{S}(\mathfrak{D}) \rightarrow \hat{\mathbf{S}}\mathbf{D}\mathbf{S}(\mathfrak{D})$. Thus, if $x' \in \max D'$, there is, by surjectivity, $x \in \max D \subseteq D$ with $x' = \varphi_{\max D}(x) = \eta_{\mathfrak{D}}(x)$, so $x' \in \eta_{\mathfrak{D}}(D)$.

Now, the two morphisms $\alpha, \alpha' : \hat{\mathbf{D}}\mathbf{S}(\mathfrak{D}) \rightarrow \mathfrak{E}$ in \mathbf{dDOM}_r agree on the maximal elements of the domain D' underlying $\hat{\mathbf{D}}\mathbf{S}(\mathfrak{D})$, so lemma 5.6.3 implies $\alpha = \alpha'$. \square





Figure 5.6: Overview of the results in the general case (top) and the standard case (bottom). Both diagrams commute up to natural isomorphism.

5.7 Conclusion

We summarize the results to see that we indeed have established the main diagram in the general case, and we check that it restricts to the standard case. For convenience, the two cases are individually depicted in figure 5.6. For simplicity, we omit from the diagram that we have analyzed the functors \bar{C} and I_B (which provide the compactification operation) into the following subdiagram:



which restricts to the standard case (i.e., adding a suffixed subscript s to all three categories).

The general case The bottom layer $DS—bTS_0—TS_{0c}$ was established in section 5.3. In sections 5.4 and 5.5, we’ve established the functors of the lower triangle $dDOM—bTS_0—TS_{0c}$ and the fact that $\bar{C} \cong S \circ D$ and $\hat{D} \vdash S$. In section 5.6, we established the upper triangle $dDOM—dDOM_r—TS_{0c}$ including the acclaimed adjunctions.

So it remains to check that the lower triangle and the upper triangle commute up to natural isomorphism. For each pair (\mathbf{C}, \mathbf{D}) of categories in the lower triangle (resp., upper triangle), we need to show that the ‘one-step’ functor from \mathbf{C} to \mathbf{D} is naturally isomorphic to the ‘two-step’ functor. For the lower triangle:

- From \mathbf{dDOM} to \mathbf{bTS}_0 , we have $\mathbf{l}_B \circ \mathbf{S} = \hat{\mathbf{S}}$ by definition. From \mathbf{dDOM} to \mathbf{TS}_{0c} , we have $\bar{\mathbf{C}} \circ \hat{\mathbf{S}} \cong (\mathbf{S} \circ \mathbf{D}) \circ (\mathbf{l}_B \circ \mathbf{S}) = \mathbf{S} \circ \hat{\mathbf{D}} \circ \mathbf{S} \cong \mathbf{S}$ where the last step follows since $\mathbf{S}\hat{\mathbf{D}} \cong \mathbf{1}_{\mathbf{TS}_{0c}}$ via the counit of the adjunction $\hat{\mathbf{D}} \vdash \mathbf{S}$.⁴⁸
- From \mathbf{bTS}_0 to \mathbf{TS}_{0c} , we have $\mathbf{S} \circ \mathbf{D} \cong \bar{\mathbf{C}}$, by the already established fact. From \mathbf{bTS}_0 to \mathbf{dDOM} , we have $\hat{\mathbf{D}} \circ \bar{\mathbf{C}} \cong \hat{\mathbf{D}}\mathbf{S} \circ \mathbf{D} \cong \mathbf{D}$ where the last step follows since $\hat{\mathbf{D}}\mathbf{S} \cong \mathbf{1}_{\mathbf{dDOM}_r}$ via the unit of the adjoint equivalence $\hat{\mathbf{D}} \vdash \mathbf{S}$ and since $\mathbf{D} : \mathbf{bTS}_0 \rightarrow \mathbf{dDOM}_r$ (since $\mathbf{D}(\mathcal{X}, \mathcal{B})$ is a restricted limit of a finitary dynamical expanding system of max-reflective domains).
- From \mathbf{TS}_{0c} to \mathbf{dDOM} , we have $\mathbf{D} \circ \mathbf{l}_B = \hat{\mathbf{D}}$ by definition. From \mathbf{TS}_{0c} to \mathbf{bTS}_0 , we have $\hat{\mathbf{S}} \circ \hat{\mathbf{D}} = (\mathbf{l}_B \circ \mathbf{S}) \circ (\mathbf{D} \circ \mathbf{l}_B) \cong \mathbf{l}_B \circ \bar{\mathbf{C}} \circ \mathbf{l}_B \cong \mathbf{l}_B$ since $\bar{\mathbf{C}}\mathbf{l}_B \cong \mathbf{1}_{\mathbf{TS}_{0c}}$ by proposition 5.3.7.

For the upper triangle:

- From \mathbf{dDOM} to \mathbf{dDOM}_r , we trivially have $\hat{\mathbf{D}} \circ \mathbf{S} = (\hat{\mathbf{D}} \circ \mathbf{S})$. From \mathbf{dDOM} to \mathbf{TS}_{0c} , $\mathbf{S} \circ (\hat{\mathbf{D}} \circ \mathbf{S}) = \mathbf{S}\hat{\mathbf{D}} \circ \mathbf{S} \cong \mathbf{S}$ since $\mathbf{S}\hat{\mathbf{D}} \cong \mathbf{1}_{\mathbf{TS}_{0c}}$ by the counit of the adjunction $\hat{\mathbf{D}} \vdash \mathbf{S}$.
- From \mathbf{dDOM}_r to \mathbf{TS}_{0c} , we have $\mathbf{S} \circ \mathbf{l} = \mathbf{S}$ regarded as functors $\mathbf{dDOM}_r \rightarrow \mathbf{TS}_{0c}$. From \mathbf{dDOM}_r to \mathbf{dDOM} , we have $\hat{\mathbf{D}} \circ \mathbf{S} = \hat{\mathbf{D}}\mathbf{S} \circ \mathbf{l} \cong \mathbf{1}_{\mathbf{dDOM}_r} \circ \mathbf{l} = \mathbf{l}$ since $\hat{\mathbf{D}}\mathbf{S} \cong \mathbf{1}_{\mathbf{dDOM}_r}$ by the unit of the adjoint equivalence $\hat{\mathbf{D}} \vdash \mathbf{S}$.
- From \mathbf{TS}_{0c} to \mathbf{dDOM} , trivially $\mathbf{l} \circ \hat{\mathbf{D}} = \hat{\mathbf{D}}$. From \mathbf{TS}_{0c} to \mathbf{dDOM}_r , we have, by the counit of the adjunction $\hat{\mathbf{D}} \vdash \mathbf{S}$, that $(\hat{\mathbf{D}} \circ \mathbf{S}) \circ \hat{\mathbf{D}} = \hat{\mathbf{D}} \circ \mathbf{S}\hat{\mathbf{D}} \cong \hat{\mathbf{D}} \circ \mathbf{1}_{\mathbf{TS}_{0c}} = \hat{\mathbf{D}}$.

The standard case The bottom layer \mathbf{sDS} — \mathbf{bTS}_{0s} — \mathbf{TS}_{0cs} also was established in section 5.3. Moreover, the functors \mathbf{D} and \mathbf{S} restrict to functors $\mathbf{D} : \mathbf{bTS}_{0s} \rightarrow \mathbf{dDOM}_s$ and $\mathbf{S} : \mathbf{dDOM}_s \rightarrow \mathbf{TS}_{0cs}$ (as seen in section 5.4). So all functors are indeed well-defined restrictions to the full subcategories of the original diagram (\mathbf{l}_B and \mathbf{l} restrict appropriately as well). Thus, the original adjunctions also are adjunctions

⁴⁸Here, and several times below, we use the following fact: If $\mathbf{F} \cong \mathbf{G}$ are naturally isomorphic functors $\mathbf{C} \rightarrow \mathbf{D}$ (via the natural isomorphism $\mathbf{F} \xrightarrow{\alpha} \mathbf{G}$) and $\mathbf{F}' \cong \mathbf{G}'$ are naturally isomorphic functors $\mathbf{D} \rightarrow \mathbf{E}$ (via the natural isomorphism $\mathbf{F}' \xrightarrow{\alpha'} \mathbf{G}'$), then $\mathbf{F}'\mathbf{F} \cong \mathbf{G}'\mathbf{G}$ (via the horizontal composition $\alpha' * \alpha$). Indeed, the horizontal composition is a natural transformation $\mathbf{F}'\mathbf{F} \Rightarrow \mathbf{G}'\mathbf{G}$ where $(\alpha' * \alpha)_X$ is defined as $\alpha'_{\mathbf{G}(X)} \circ \mathbf{F}'(\alpha_X) = \mathbf{G}'(\alpha_X) \circ \alpha'_{\mathbf{F}(X)}$ (Leinster 2014, p. 37). Since this is a composition of isomorphisms in \mathbf{E} , $\alpha' * \alpha$ is a natural isomorphism. (This is, e.g., discussed on this *stack exchange* thread: <https://math.stackexchange.com/a/2516416> (last checked 3 July 2021).)

in the restricted diagram. And natural isomorphisms of original functors also are natural isomorphisms in the restricted diagram, so the restricted diagram still commutes up to natural isomorphism.

Further questions In addition to the many open questions already posed in the previous chapter, one may ask for a characterization of the domain-theoretic counterpart of metric isomorphism (i.e., when two dynamical domains model metrically isomorphic dynamical systems). Then one can explore whether this ‘weak equivalence of dynamical domains’ and the established localization of topological systems are the start of a ‘homotopy structure’ (cf. category of fibrant objects).

Chapter 6

Systems and domains 3: Application

Abstract In the previous two chapters, we've developed dynamical domains as a tool to analyze dynamical systems. In this chapter, we extend this by developing domain-theoretic counterparts to the important system-theoretic concepts of metric and topological entropy.

6.1 Introduction

In the previous two chapters, we have introduced dynamical domains and have shown that they are closely connected to dynamical systems. To further develop the idea of dynamical domains as tools to analyze dynamical systems, we develop domain-theoretic counterparts to the system-theoretic concept of entropy. Entropy is a tremendously important concept to study dynamical systems. It provides a quantitative measure of how 'chaotic' the system is. It also is an isomorphism invariant and hence helps in determining whether two systems are, despite their different appearances, isomorphic after all.

The chapter is structured as follows. In section 6.2, we provide the relevant background on dynamical domains, dynamical systems, and entropy. In section 6.3, we develop the concept of *domain-entropy* as a domain-theoretic counterpart to metric entropy. In section 6.4, we develop the concept of *max-entropy* as a domain-theoretic counterpart to topological entropy. In section 6.5, we conclude with some open questions. In appendix B of the thesis, we discuss a detailed example of a dynamical domain and calculate its max-entropy.

6.2 Background

We provide the relevant background on dynamical systems and dynamical domains from the previous two chapters, and we recall the definitions of metric entropy and topological entropy.

6.2.1 Recap dynamical systems and dynamical domains

In the previous two chapters, we've motivated and defined (general and standard) dynamical systems as structures $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a probability space and $T : X \rightarrow X$ is a measurable function (that satisfy some further conditions in the general and standard case, respectively).

We've seen that such a system \mathfrak{X} has a topological realization $\mathfrak{Y} = (Y, \sigma, \nu, S)$: i.e., (Y, σ) is a compact zero-dimensional Polish space with a probability measure ν on the Borel σ -algebra $\mathcal{B}(\sigma)$ and a continuous dynamics $S : Y \rightarrow Y$, and the system \mathfrak{X} is (modulo completion) isomorphic to the system $(Y, \mathcal{B}(Y), \nu, S)$ induced by \mathfrak{Y} . Thus, for our purposes here, we can restrict us to working with the category \mathbf{TS}_{0c} of these zero-dimensional compact measured topological systems. The benefit is that we then can define in one place not only measure-theoretic (i.e., metric) entropy but also topological entropy.

In the previous two chapters, we've introduced the category \mathbf{dDOM} of dynamical domains and the adjunction

$$\mathbf{dDOM} \begin{array}{c} \xrightarrow{S} \\ \perp \\ \xleftarrow{\hat{D}} \end{array} \mathbf{TS}_{0c}.$$

This means the following: A dynamical domain \mathfrak{D} is a structure (D, v, f) where D is a Scott domain, $v : \Sigma(D) \rightarrow [0, 1]$ a valuation, and $f : D \rightarrow D$ a Scott-continuous function, and \mathfrak{D} is, moreover, obtained as a certain limit of finite structures of this form. Each dynamical domain \mathfrak{D} induces (or 'models') a compact zero-dimensional measured topological system $\mathfrak{S}(\mathfrak{D})$: Its state space consists of the maximal elements $\max D$ of D and the dynamics is given by the restriction $f \upharpoonright \max D$; the measure μ_v is determined by the valuation v . In the other direction, each compact zero-dimensional measure topological system \mathfrak{X} induces a dynamical domain $\hat{D}(\mathfrak{X})$. It is called the observation domain since it is constructed from the clopen subsets of X which are viewed as possible observations that we can make about \mathfrak{X} . Here we won't repeat the details, but refer, for undefined notation and terminology, to chapters 4–5.

Since metric entropy is only defined for measure-preserving transformations, we need to restrict to the subcategory \mathbf{TS}_{0cm} of those $\mathfrak{X} = (X, \tau, \mu, T)$ where $T : X \rightarrow X$ is measure-preserving (i.e., for $A \in \mathcal{B}(\tau)$, we have $\mu(T^{-1}(A)) = \mu(A)$).¹ This restriction is not necessary to discuss topological entropy, since it doesn't require a measure.

This restriction is mirrored on the domain-side by restricting to the category \mathbf{dDOM}_v of those dynamical domains that are obtained as restricted limits to diagrams that are, what we've called, eventually valuation-preserving. If \mathfrak{X} is in \mathbf{TS}_{0cm} , then the diagram with which the observation domain $\hat{D}(\mathfrak{X})$ is built

¹We don't assume T to be invertible, so \mathbf{TS}_{0cm} is a supercategory of the category \mathbf{TS}_{0cs} of standard compact zero-dimensional measured topological systems.

has this property (see the comments after theorem 4.3.11 in chapter 4). And if $\mathfrak{D} = (D, v, f)$ is in dDOM_v , then f is valuation-preserving, i.e., $v(f^{-1}(U)) = v(U)$ for all $U \in \Sigma(D)$ (see the comments after theorem 4.4.8 in chapter 4), so $\mathfrak{S}(\mathfrak{D})$ is measure-preserving (see the comments after theorem 4.5.1 in chapter 4). Thus, the adjunction restricts to

$$\text{dDOM}_v \begin{array}{c} \xrightarrow{\mathfrak{S}} \\ \perp \\ \xleftarrow{\mathfrak{D}} \end{array} \text{TS}_{0\text{cm}}.$$

Also recall that when we restrict to what we've called max-reflective dynamical domains, these adjunctions restrict to equivalences.

6.2.2 Metric entropy

Metric entropy (or Kolmogorov-Sinai entropy) is a central concept in dynamical systems theory. For a detailed discussion, we refer to the many references on the topic—for example, Downarowicz (2011). Here we just recall its formal definition.

We'll follow Walters (1982, sec. 4) who presents entropy theory in the general setting of a probability space with a measure-preserving transformation. (So this doesn't presuppose a Lebesgue space or invertible transformation like some other references do for independent reasons.) However, we'll adjust the notation slightly to highlight similarities with the domain-theoretic setting (in particular, with the observation domain).

By 'log' we refer to the logarithm with basis 2, but other bases (like the natural logarithm) are possible and will change the definition of entropy only by a fixed multiplicative constant.

Although the definition can be stated more generally, we work, as motivated in section 6.2.1, with the objects of $\text{TS}_{0\text{cm}}$.

6.2.1. DEFINITION. Let $\mathfrak{X} = (X, \tau, \mu, T)$ be in $\text{TS}_{0\text{cm}}$. A *finite measurable partition* of \mathfrak{X} is a finite subset \mathcal{C} of $\mathcal{B}(\tau)$ that partitions X . For such \mathcal{C} and $n \geq 1$ define,

$$H(n, \mathcal{C}) := - \sum_{(C_0, \dots, C_{n-1}) \in \mathcal{C}^n} \mu \left(\bigcap_{k=0}^{n-1} T^{-k} C_k \right) \log \mu \left(\bigcap_{k=0}^{n-1} T^{-k} C_k \right),$$

with the convention that $0 \log 0 = 0$. The *metric entropy* of \mathfrak{X} is

$$h(\mathfrak{X}) := \sup \left\{ \lim_{n \rightarrow \infty} \frac{1}{n} H(n, \mathcal{C}) : \mathcal{C} \text{ finite measurable partition of } X \right\}.$$

We collect some standard facts about entropy that we'll need for the proofs below. If (X, \mathcal{A}, μ) is a probability space, we say, for two finite measurable partitions \mathcal{C} and \mathcal{D} , that \mathcal{D} *refines* \mathcal{C} (written $\mathcal{C} \preceq \mathcal{D}$) iff every \mathcal{C} -element is a union of \mathcal{D} -elements. A sequence (\mathcal{C}_k) of finite measurable partitions *generates* \mathcal{A} iff $\mathcal{A} = \sigma(\bigcup_k \mathcal{C}_k)$ (where $\sigma(\cdot)$ is the smallest σ -algebra containing \cdot).

6.2.2. PROPOSITION. *Let $\mathfrak{X} = (X, \tau, \mu, T)$ be in $\mathbf{TS}_{0\text{cm}}$ and let \mathcal{C} and \mathcal{D} be finite measurable partitions. Then*

1. *If $\mathcal{C} \preceq \mathcal{D}$ and $n \geq 1$, then $H(n, \mathcal{C}) \leq H(n, \mathcal{D})$.*
2. *$(\frac{1}{n}H(n, \mathcal{C}))_n$ is a decreasing sequence of positive real numbers; in particular, its limit exists.*
3. *If $\mathcal{C} \preceq \mathcal{D}$, then $\lim_n \frac{1}{n}H(n, \mathcal{C}) \leq \lim_n \frac{1}{n}H(n, \mathcal{D})$.*
4. *If (\mathcal{C}_k) is a refining sequence of finite measurable partitions that generates $\mathcal{B}(\tau)$, then $h(\mathfrak{X}) = \lim_k (\lim_n \frac{1}{n}H(n, \mathcal{C}_k))$.*
5. *If $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is a morphism in $\mathbf{TS}_{0\text{cm}}$, then $h(\mathfrak{X}) \geq h(\mathfrak{Y})$.*

The proof is sketched in an appendix at the end of the chapter.

6.2.3 Topological entropy

Topological entropy is analogous to metric entropy but, as the name suggests, it only uses topological information and no measure-theoretic information. It is (usually) defined for topological systems (X, T) where X is a compact metric space and $T : X \rightarrow X$ is continuous. There are several equivalent definitions available (see, e.g., Downarowicz 2011, ch. 6). Here we use the ‘open cover’ definition introduced by Adler, Konheim, and McAndrew (1965), because it best highlights similarities to the observation domain construction.

We again follow Walters (1982, ch. 7) with slightly adjusted notation for emphasis of similarity to the domain-theoretic setting. And, again, although the definition can be stated more generally, we state it for the objects of $\mathbf{TS}_{0\text{c}}$ (ignoring the measure), which we’ve motivated as a convenient setting for our discussion.

6.2.3. DEFINITION. Let $\mathfrak{X} = (X, \tau, \mu, T)$ be in $\mathbf{TS}_{0\text{c}}$. (We ignore the measure μ .) A *finite open cover* \mathcal{C} of \mathfrak{X} is a finite subset of τ whose union is X . For such \mathcal{C} and $n \geq 1$, define

$$H_{\text{top}}(n, \mathcal{C}) := \log N(\mathcal{C}_0^{n-1}),$$

where $\mathcal{C}_0^{n-1} := \{\bigcap_{k=0}^{n-1} T^{-k}U_k : U_0, \dots, U_{n-1} \in \mathcal{C}\}$ and $N(\mathcal{C}_0^{n-1})$ is the minimum cardinality of the subcovers of \mathcal{C}_0^{n-1} . The *topological entropy* of \mathfrak{X} is

$$h_{\text{top}}(\mathfrak{X}) := \sup \left\{ \lim_{n \rightarrow \infty} \frac{1}{n} H_{\text{top}}(n, \mathcal{C}) : \mathcal{C} \text{ finite open cover of } X \right\}.$$

We collect some facts about topological entropy. If \mathcal{C} and \mathcal{D} are two open covers of a topological space, we say \mathcal{D} *refines* \mathcal{C} (written $\mathcal{C} \preceq \mathcal{D}$) if every \mathcal{D} -element is a subset of a \mathcal{C} -element (note the difference to the case of partitions).

6.2.4. PROPOSITION. *Let $\mathfrak{X} = (X, \tau, \mu, T)$ be in \mathbf{TS}_{0c} and let \mathcal{C} and \mathcal{D} be open covers. Then*

1. *If $\mathcal{C} \preceq \mathcal{D}$ and $n \geq 1$, then $H_{\text{top}}(n, \mathcal{C}) \leq H_{\text{top}}(n, \mathcal{D})$.*
2. *$\lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C})$ exists.*
3. *If $\mathcal{C} \preceq \mathcal{D}$, then $\lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C}) \leq \lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{D})$.*
4. *If (\mathcal{C}_k) is a refining sequence of open covers that eventually refines every open cover, then $h_{\text{top}}(\mathfrak{X}) = \lim_k \left(\lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C}_k) \right)$.*
5. *If $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is a surjective morphism in \mathbf{TS}_{0c} , then $h_{\text{top}}(\mathfrak{X}) \geq h_{\text{top}}(\mathfrak{Y})$.*

The proof is sketched in an appendix at the end of the chapter.

6.3 Domain-entropy

In this section, we define the domain-theoretic counterpart to metric entropy—which we’ll call *domain-entropy*. We first state its definition (section 6.3.1) and then show that it indeed is a proper counterpart (section 6.3.2). Finally, we prove that the domain-entropy of a dynamical domain can be expressed in a certain ‘normal form’ (section 6.3.3).

6.3.1 Definition of domain-entropy

Given the adjunction $\mathbf{dDOM}_v \rightleftarrows \mathbf{TS}_{0cm}$, there are two suggestive ways to try to construct domain-theoretic counterparts to metric entropy: First, we start with a dynamical domain \mathfrak{D} and see how the metric entropy of the induced system $\mathfrak{S}(\mathfrak{D})$ may carry over to \mathfrak{D} . Second, we start with a system \mathfrak{X} and see how its metric entropy may carry over to the observation domain $\hat{\mathfrak{D}}(\mathfrak{X})$.

The second way seems appealing given the use of pairs (n, \mathcal{C}) in the definition of entropy (highlighted in our formulation) which is reminiscent of the (index set of) the observation domain. Indeed, the term $H(n, \mathcal{C})$ in the definition of metric entropy of a system \mathfrak{X} is, writing $i = (n, \mathcal{C})$, very close to

$$- \sum_{a \in \max D_i} v_i(a) \log v_i(a)$$

in the setting of the observation domain of the system \mathfrak{X} .

However, for a general dynamical domain $\mathfrak{D} = (D, v, f)$, this doesn’t take into account the *dynamics* f on the domain. And for an abstract index i , it is not clear what the ‘time’ parameter n should be. Moreover, this would only work for dynamical domains that are isomorphic to some observation domain, i.e., it only works for max-reflective dynamical domains.

For these reasons, we first consider the first way. Nonetheless, we'll use the second way in section 6.3.3 to prove a normal form theorem.

If \mathfrak{D} is in \mathbf{dDOM}_v , then by a *diagram F giving rise to \mathfrak{D} via projections p_i* we mean a finitary dynamical expanding system $F = (\mathfrak{D}_i, p_{ij})_I$ that is eventually valuation-preserving together with dynamical morphisms $p_i : \mathfrak{D} \rightarrow \mathfrak{D}_i$ which also are projections such that (\mathfrak{D}, p_i) is the restricted limit to F (see chapter 4 for definitions of these terms). By definition of \mathbf{dDOM}_v , such (F, p_i) always exist. If it is clear from context, we write $\mathfrak{D} = (D, v, f)$ and $\mathfrak{D}_i = (D_i, v_i, f_i)$.

6.3.1. DEFINITION. Let \mathfrak{D} be in \mathbf{dDOM}_v and let $F = (\mathfrak{D}_i, p_{ij})_I$ be a diagram giving rise to \mathfrak{D} via projections p_i . For $i \in I$ and $n \geq 1$, define

$$E(n, i) := - \sum_{\bar{a} \in (\max D_i)^n} v(U_{p_i}(\bar{a})) \log v(U_{p_i}(\bar{a})),$$

where $U_{p_i}(\bar{a}) := \bigcap_{k=0}^{n-1} f^{-k}(p_i^{-1}(\bar{a}(k)))$ is the set of $a \in D$ whose i -th projection follows \bar{a} under the domain dynamics f . We define the *domain-entropy* of \mathfrak{D} as

$$e(\mathfrak{D}) := \sup \left\{ \lim_{n \rightarrow \infty} \frac{1}{n} E(n, i) : i \in I \right\}. \quad (6.1)$$

Below we show that these limits exist. In general, for \mathfrak{D} in \mathbf{dDOM}_v , we define $e(\mathfrak{D})$ with respect to *some* (F, p_i) giving rise to \mathfrak{D} (below we show that this is well-defined, i.e., independent of the choice of (F, p_i)).

Comments: First, for an intuitive interpretation of this definition, let's first recall from information theory that, for a random variable X with finitely many outcomes x_1, \dots, x_n with corresponding probabilities p_1, \dots, p_n , its Shannon entropy $H(X) = -\sum_{k=1}^n p_k \log p_k$ is interpreted as the average uncertainty in the outcomes of variable X . Thus, we can interpret $E(n, i)$ as the average uncertainty in the n -long behavior of the domain dynamics as seen at (i.e., projected to) the i -th component D_i . Thus, $\lim_{n \rightarrow \infty} \frac{1}{n} E(n, i)$ is, in a sense, the time-average of that uncertainty. So the least upper bound $e(\mathfrak{D})$ describes the uncertainty in the domain dynamics that we have to reckon with.

Second, hence the intuition is much like that of metric entropy (see, e.g., Petersen 1983 or Downarowicz 2011). The role of the finite partitions \mathcal{C} in metric entropy is now taken over by the indices i .

Third, one might also try to formulate the definition only relying on the diagram and not on the limit dynamics f (to have 'finitary' definition). However, looking at the characterization of $v(f^{-1}U_i)$ in the limit theorem, this will, at least, complicate notation, so we stick to the above definition. However, in the normal form result below, we'll consider a purely diagram based definition.

6.3.2 Main theorem on domain-entropy

The following is the main theorem about domain-entropy showing, in particular, that it is well-defined (items (1) and (2) below), an isomorphism invariant as any good notion of entropy should be (items (3) and (4) below), and a counterpart to metric entropy (items (5) and (6) below).

- 6.3.2. THEOREM.**
1. *The limits in equation (6.1) exist.*
 2. *The domain-entropy $e(\mathfrak{D})$ is independent of the choice of (F, p_i) giving rise to \mathfrak{D} .*
 3. *If $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ is a morphism in \mathbf{dDOM}_v , then $e(\mathfrak{D}) \geq e(\mathfrak{E})$.*
 4. *In particular, domain-entropy is an isomorphism invariant: two dynamical domains that are isomorphic in \mathbf{dDOM}_v have the same domain-entropy.*
 5. *If \mathfrak{D} is in \mathbf{dDOM}_v , then $e(\mathfrak{D}) = h(\mathbf{S}(\mathfrak{D}))$.*
 6. *If \mathfrak{X} is in $\mathbf{TS}_{0\text{cm}}$, then $h(\mathfrak{X}) = e(\hat{\mathbf{D}}(\mathfrak{X}))$.*

We prove the theorem in the remainder of this subsection. The following lemma contains most of the calculations.

6.3.3. LEMMA. *Let \mathfrak{D} be in \mathbf{dDOM}_v and let $F = (\mathfrak{D}_i, p_{ij})_I$ be a diagram giving rise to \mathfrak{D} via projections p_i . Then*

1. *For each $i \in I$, $\mathcal{C}_i := \{p_i^{-1}(a) \cap \max D : a \in \max D_i\}$ is a finite clopen partition of $\max D$.*
2. *If $i_0 \leq i_1 \leq \dots$ is cofinal in I , then $(\mathcal{C}_{i_k})_k$ is a refining sequence of finite clopen partitions that generates $\mathcal{B}(\max D)$.*
3. *For each $i \in I$ and $n \geq 1$, we have $H(n, \mathcal{C}_i) = E(n, i)$ (where $H(n, \mathcal{C}_i)$ is calculated in the system $\mathbf{S}(\mathfrak{D})$).*
4. *For $i \leq j$ in I and $n \geq 1$, $E(n, i) \leq E(n, j)$.*

PROOF. Ad (1). Since D_i is finite, \mathcal{C}_i is finite. It is a partition, since, for each $a \in \max D$, we have, because p_i is max-preserving, that $p_i(a) \in \max D_i$, so a is in one and only one cell of \mathcal{C}_i . And its cells are clopen because sets of the form $p_i^{-1}(U_i) \cap \max D$ for $U_i \in \Sigma(D_i)$ are clopen in $\max D$.

Ad (2). Refining: For $i := i_k \leq i_l := j$ (with $k \leq l$), we need to show that \mathcal{C}_j refines \mathcal{C}_i (denoted $\mathcal{C}_i \preceq \mathcal{C}_j$), i.e., each \mathcal{C}_i -element is a union of \mathcal{C}_j -elements. Indeed, for $a_i \in \max D_i$, we have

$$p_i^{-1}(a_i) \cap \max D = \bigcup_{a_j \in p_{ij}^{-1}(a_i)} p_j^{-1}(a_j) \cap \max D,$$

since, for $a \in \max D$: if $p_i(a) = a_i$, set $a_j := p_j(a) \in p_{ij}^{-1}(a_i)$; and if $p_j(a) = a_j$ for some $a_j \in p_{ij}^{-1}(a_i)$, then $p_i(a) = p_{ij}(p_j(a)) = p_{ij}(a_j) = a_i$.

Generating: We need to show that $\mathcal{A} := \sigma(\bigcup_{k \geq 0} \mathcal{C}_{i_k}) = \mathcal{B}(\max D)$. Since each cell is (cl)open, we have \subseteq . For the other direction, we need to show that \mathcal{A} contains all the opens of $\max D$. Indeed, qua second-countable space, any open of $\max D$ is a countable union of basic open sets, so it suffices to show that the latter are in \mathcal{A} (since \mathcal{A} is closed under countable union). So let $U = \bigcap_{l=0}^m p_{j_l}^{-1}(U_{j_l}) \cap \max D$ be a basic open set with $j_0, \dots, j_m \in I$ and $U_{j_l} \in \Sigma(D_{j_l})$ (for $l = 0, \dots, m$). Since \mathcal{A} is closed under finite intersection, it suffices to show that each $p_{j_l}^{-1}(U_{j_l}) \cap \max D$ is in \mathcal{A} . Indeed, given j_l , let, since $i_0 \leq i_1 \leq \dots$ is cofinal, $i_k \geq j_l$. Then, similarly as above,

$$p_{j_l}^{-1}(U_{j_l}) \cap \max D = \bigcup_{a_{i_k} \in p_{j_l}^{-1}(U_{j_l})} p_{i_k}^{-1}(a_{i_k}) \cap \max D,$$

which is a finite union (since D_{i_k} is finite) of elements in \mathcal{A} , whence in \mathcal{A} .

Ad (3). Write $\mathcal{C} := \mathcal{C}_i$. Recall that $H(n, \mathcal{C})$ is defined over the system $\mathfrak{S}(\mathfrak{D})$ as

$$H(n, \mathcal{C}) = - \sum_{(C_0, \dots, C_{n-1}) \in \mathcal{C}^n} \mu_v \left[\prod_{k=0}^{n-1} \bar{f}^{-k} C_k \right] \log \mu_v \left[\prod_{k=0}^{n-1} \bar{f}^{-k} C_k \right],$$

where $\bar{f} := f \upharpoonright \max D$ and μ_v is the probability measure on $\mathcal{B}(\max D)$ determined by v .

Note that $\max D_i$ is in bijective correspondence with \mathcal{C} by $a \mapsto b(a) := p_i^{-1}(a) \cap \max D$. So we can write the sum as

$$H(n, \mathcal{C}) = - \sum_{(a_0, \dots, a_{n-1}) \in (\max D_i)^n} \mu_v \left[\prod_{k=0}^{n-1} \bar{f}^{-k} b(a_k) \right] \log \mu_v \left[\prod_{k=0}^{n-1} \bar{f}^{-k} b(a_k) \right].$$

Moreover, we have, since f is max-preserving and \bar{f} the restriction,

$$\begin{aligned} \bar{f}^{-k} b(a_k) &= \bar{f}^{-k} (p_i^{-1}(a_k) \cap \max D) \\ &= f^{-k} (p_i^{-1}(a_k) \cap \max D) \cap \max D \\ &= f^{-k} (p_i^{-1}(a_k)) \cap \max D. \end{aligned}$$

Hence

$$\begin{aligned} \mu_v \left[\prod_{k=0}^{n-1} \bar{f}^{-k} b(a_k) \right] &= \mu_v \left[\prod_{k=0}^{n-1} f^{-k} (p_i^{-1}(a_k)) \cap \max D \right] \\ &= \mu_v \left[\prod_{k=0}^{n-1} f^{-k} (p_i^{-1}(a_k)) \right] \\ &= v \left[\prod_{k=0}^{n-1} f^{-k} (p_i^{-1}(a_k)) \right]. \end{aligned}$$

So we can write the sum as

$$\begin{aligned} H(n, \mathcal{C}) &= - \sum_{(a_0, \dots, a_{n-1}) \in (\max D_i)^n} v \left[\prod_{k=0}^{n-1} f^{-k}(p_i^{-1}(a_k)) \right] \log v \left[\prod_{k=0}^{n-1} f^{-k}(p_i^{-1}(a_k)) \right] \\ &= - \sum_{\bar{a} \in (\max D_i)^n} v[U_{p_i}(\bar{a})] \log v[U_{p_i}(\bar{a})] = E(n, i), \end{aligned}$$

as needed.

Ad (4). By (2), \mathcal{C}_j refines \mathcal{C}_i , so, by fact (1) about entropy (proposition 6.2.2) and by (3),

$$E(n, i) = H(n, \mathcal{C}_i) \leq H(n, \mathcal{C}_j) = E(n, j),$$

as needed. \square

PROOF OF THEOREM 6.3.2. We'll prove the items in a different order than listed, because to show that domain-entropy is independent of the diagram, we need diagram-dependent versions of (3)–(5) which we mark with an asterisk and state precisely below.

Ad (1). In the setting of definition 6.3.1, we need to show, for $i \in I$, that $\lim_n \frac{1}{n} E(n, i)$ exists. By lemma 6.3.3 (3), we have $\lim_n \frac{1}{n} E(n, i) = \lim_n \frac{1}{n} H(n, \mathcal{C}_i)$ and the latter exists by fact (2) about entropy (proposition 6.2.2).

Now, for (F, p_i) giving rise to \mathfrak{D} in \mathbf{dDOM}_v , the extended real $e(\mathfrak{D}, p_i, F)$ in equation (6.1) is well-defined.

Claim (5)*: If (F, p_i) gives rise to \mathfrak{D} , then $e(\mathfrak{D}, p_i, F) = h(\mathbf{S}(\mathfrak{D}))$.

Indeed, let $i_0 \leq i_1 \leq \dots$ be cofinal in the index set I of F (which exists since I is countable and directed). Write $\mathfrak{X} := \mathbf{S}(\mathfrak{D})$. By lemma 6.3.3 (2), $(\mathcal{C}_{i_k})_{k \geq 0}$ is a refining sequence of finite measurable partitions of \mathfrak{X} that generates $\mathcal{B}(\mathfrak{X})$. By fact (4) about entropy, $h(\mathfrak{X}) = \lim_k (\lim_n \frac{1}{n} H(n, \mathcal{C}_{i_k}))$. Since the terms are increasing by fact (3), we can replace the 'lim_k' by 'sup_k'. By lemma 6.3.3 (3), this then further equals $\sup_k (\lim_n \frac{1}{n} E(n, i_k))$. So it remains to show

$$r := \sup_k \left(\lim_n \frac{1}{n} E(n, i_k) \right) = \sup \left\{ \lim_n \frac{1}{n} E(n, i) : i \in I \right\} \quad (= e(\mathfrak{D}, p_i, F)).$$

By the subset relation, we have \leq . For the other direction, let $i \in I$ and show $\lim_n \frac{1}{n} E(n, i) \leq r$. By cofinality, there is $i_k \geq i$. By lemma 6.3.3 (4), we have, for all $n \geq 0$, that $E(n, i) \leq E(n, i_k)$. So $\lim_n \frac{1}{n} E(n, i) \leq \lim_n \frac{1}{n} E(n, i_k) \leq r$.

Claim (3)*: If $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ is in \mathbf{dDOM}_v and (F, p_i) and (G, q_j) giving rise to \mathfrak{D} and \mathfrak{E} , respectively, then $e(\mathfrak{D}, p_i, F) \geq e(\mathfrak{E}, q_j, G)$.

Indeed, then $\mathbf{S}(\alpha) : \mathbf{S}(\mathfrak{D}) \rightarrow \mathbf{S}(\mathfrak{E})$ is a morphism in $\mathbf{TS}_{0\text{cm}}$, so by fact (5) and by (5)*,

$$e(\mathfrak{D}, p_i, F) = h(\mathbf{S}(\mathfrak{D})) \geq h(\mathbf{S}(\mathfrak{E})) = e(\mathfrak{E}, q_j, G).$$

Claim (4)*: If \mathfrak{D} and \mathfrak{E} are isomorphic in \mathbf{dDOM}_v with (F, p_i) and (G, q_j) giving rise to them, respectively, then $e(\mathfrak{D}, p_i, F) = e(\mathfrak{E}, q_j, G)$.

Indeed, then there are, in particular, morphisms $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ and $\beta : \mathfrak{E} \rightarrow \mathfrak{D}$ in \mathbf{dDOM}_v , so, by (3)*, $e(\mathfrak{D}, p_i, F) \geq e(\mathfrak{E}, q_j, G)$ and $e(\mathfrak{E}, q_j, G) \geq e(\mathfrak{D}, p_i, F)$, as needed.

Ad (2). Let (F, p_i) and (G, q_j) give rise to \mathfrak{D} in \mathbf{dDOM}_v . Since \mathfrak{D} is isomorphic to itself, (4)* implies $e(\mathfrak{D}, p_i, F) = e(\mathfrak{D}, q_j, G)$.

Hence, defining, for \mathfrak{D} in \mathbf{dDOM}_v , $e(\mathfrak{D}) := e(\mathfrak{D}, p_i, F)$ for some (F, p_i) giving rise to \mathfrak{D} is well-defined.

Now, (3)–(5) are implied by their asterisked versions. So it remains to show the last item.

Ad (6). Let \mathfrak{X} be in $\mathbf{TS}_{0\text{cm}}$. Since the counit of the adjunction $\hat{\mathbf{D}} \vdash \mathbf{S}$ is a natural isomorphism, \mathfrak{X} is isomorphic in $\mathbf{TS}_{0\text{cm}}$ to $\mathbf{SD}(\mathfrak{X})$. Since h is an isomorphism invariant, $h(\mathfrak{X}) = h(\mathbf{SD}(\mathfrak{X}))$. Further, since $\hat{\mathbf{D}}(\mathfrak{X})$ is in \mathbf{dDOM}_v , we have, by (5), $h(\mathbf{SD}(\mathfrak{X})) = e(\hat{\mathbf{D}}(\mathfrak{X}))$, as needed. \square

6.3.3 Normal form for domain-entropy

Before we can state the normal form theorem, we fix some terminology. Recall that a net $(x_i)_I$ is a function from a directed set I to the reals \mathbb{R} , and

$$\liminf_{i \in I} x_i := \sup_{i \in I} \inf_{j \geq i} x_j,$$

which takes values in the extended reals $[-\infty, +\infty]$ (see, e.g., Beer 1993, p. 2). If I is a directed set, we call, for reasons explained below, a projection $t : I \rightarrow \mathbb{N}$ a *time function* on I . Also, \mathbf{dDOM}_{rv} is the full subcategory of \mathbf{dDOM}_v whose dynamical domains are max-reflective.

6.3.4. THEOREM. *Let \mathfrak{D} be in \mathbf{dDOM}_{rv} . Then there is a diagram $(\mathfrak{D}_i, p_{ij})_I$ giving rise to \mathfrak{D} (via projections p_i) and a time function $t : I \rightarrow \mathbb{N}$ such that*

$$e(\mathfrak{D}) = \liminf_{i \in I_t} \frac{-1}{t(i)} \sum_{a \in \max D_i} v_i(a) \log v_i(a), \quad (6.2)$$

where $I_t := \{i \in I : t(i) \neq 0\}$.²

Comments: First, the term on the right in (6.2) only depends on the diagram $(\mathfrak{D}_i, p_{ij})_I$ and the time function t , but *not* on the projections p_i from the limit. So

²Thus, $\frac{1}{t(i)}$ is well-defined. Note that I_t is still directed (with the inherited order): Since t is a projection, it is surjective, so, in particular, there is $i \in I$ with $t(i) \neq 0$, whence $I_t \neq \emptyset$. And if $i, j \in I_t$, let, since I is directed, $k \in I$ with $k \geq i, j$. Since t is monotone, $t(k) \geq t(i) > 0$, so $k \in I_t$.

it is ‘finitary’ in that it only depends on the finitary representation (the diagram) of the dynamical domain together with an ‘abstract’ time assigned to each element of the index set.

Second, the term on the right in (6.2) is *static* and not *dynamic*: it only depends on the finite domains D_i and their valuations v_i , but *not* on the dynamics f_i on them. Thus, the dynamic domain-entropy $e(\mathfrak{D})$ can be expressed in a static way by an appropriate choice of representation of \mathfrak{D} . Of course, the dynamic information cannot be lost: the proof will show that it is, roughly, encoded into the index set.

Third, here is an interpretation of

$$- \sum_{a \in \max D_i} v_i(a) \log v_i(a) \quad (6.3)$$

(for a given $i \in I$). Recall that we can think of a domain D as the (data type of) possible outputs of a computational process (or a collection of processes); the higher we go in the order of D , the more informationally complete (or closer to termination) the outputs become. Then, the Scott-open sets are the (finitely) observable properties of these processes, and the valuation v describes the probabilities of observing these properties.³ Thus, for our finite domain D_i here, the $v_i(a)$ describe the probability of the computational process to terminate with the (maximal, i.e., informationally complete) output a . Combining this with the usual interpretation of Shannon entropy, we can interpret (6.3) as the average uncertainty in the terminating output of the computational process.

Fourth, here is an interpretation of the time function t . As in the usual interpretation of an expanding system of domains (Abramsky and Jung 1994, sec. 3.1.4 and 3.3.2), we think of the finite domains D_i as increasingly better approximations to the domain D . Let’s assume for a moment that the index set I is \mathbb{N} and $t : I \rightarrow \mathbb{N}$ the identity function. Then D_i is the approximation to D that we’ve reached after $i = t(i)$ many (time-) steps. The limit expression in (6.2) then describes, roughly, the time-average of the average uncertainty in the terminating output—much like $\lim_n \frac{1}{n} H(n, \mathcal{C})$ describes the time-average of the average uncertainty of the n -long behavior of the system seen through partition \mathcal{C} . However, in general, there may be many approximation chains (D_i) to D that are brought together by indexing them over a directed set I rather than just \mathbb{N} . In that case, the time function t assigns each index i the number of time-steps that were needed to reach approximation D_i to D .

Fifth, the ‘lim inf’ merges the two limits (the ‘sup’ and the ‘lim’) found in the definition of both metric entropy and domain-entropy. (Cf. lemma 6.3.8 below.) So domain-entropy is not formulated anymore as one limit over partitions and one limit over time. Rather it is formulated as a ‘unified’ limit over the index set:

³We’ve mentioned these ideas (and references) in the brief introduction to domain theory in chapter 4.

as envisaged in the ‘second way’ of approaching domain-theoretic entropy at the beginning of section 6.3.1.

We prove the theorem in the remainder of this subsection. We start with several preparatory lemmas.

6.3.5. LEMMA. *In the setting of lemma 6.3.3 (2), $J := \{(n, \mathcal{C}_{i_k}) : n, k \in \mathbb{N}\}$ is cofinal in $I(\text{Clp max } D)$.*

PROOF. Since each $\mathcal{C}_k := \mathcal{C}_{i_k}$ is a finite clopen partition of $\text{max } D$, it in particular is a finite $\text{Clp max } D$ -cover of $\text{max } D$, so J is indeed a subset of $I(\text{Clp max } D)$. So we need to show that, for a finite $\text{Clp max } D$ -cover \mathcal{D} of $\text{max } D$, there is some k such that $\mathcal{D} \leq \mathcal{C}_k$.

Write $\mathcal{D} = \{D_0, \dots, D_r\}$. Each D_k is, qua open set, a union of basic (cl)opens and, qua compact set (since it is a closed subset of the compact space) this can be assumed to be a finite union. Let \mathcal{D}' be the cover consisting of the finitely many basic clopens occurring in the unions of the D_k . Then \mathcal{D}' is a finite $\text{Clp max } D$ -cover of $\text{max } D$ consisting of basic clopens with $\mathcal{D} \leq \mathcal{D}'$ (every \mathcal{D}' -element is a subset of a \mathcal{D} -element, and every \mathcal{D} -element can be written as union of \mathcal{D}' -elements).

Write $\mathcal{D}' = \{D'_0, \dots, D'_s\}$. Each D'_l is a finite intersection of sets of the form $p_i^{-1}(U_i) \cap \text{max } D$. Since (i_k) is cofinal in the directed I , we can choose i_k to be bigger than all the finitely many indices i used in the intersections of the D'_l 's. We claim that $\mathcal{C}_k \geq \mathcal{D}'$ (which finishes the proof).

(1) Given a \mathcal{C}_k -element $p_{i_k}^{-1}(a) \cap \text{max } D$ with $a \in \text{max } D_{i_k}$, we need to show that it is a subset of some \mathcal{D}' -element. Let $x \in p_{i_k}^{-1}(a) \cap \text{max } D$ (since p_{i_k} is a projection), so $x \in D'_l$ for some $l \in \{0, \dots, s\}$. Qua basic open, write $D'_l = \bigcap_{t=0}^m p_{j_t}^{-1}(U_{j_t}) \cap \text{max } D$ for $j_0, \dots, j_m \in I$ and $U_{j_t} \in \Sigma(D_{j_t})$. To show $p_{i_k}^{-1}(a) \cap \text{max } D \subseteq D'_l$, let $t \in \{0, \dots, m\}$ and show $p_{i_k}^{-1}(a) \cap \text{max } D \subseteq p_{j_t}^{-1}(U_{j_t}) \cap \text{max } D$. Indeed, if $y \in p_{i_k}^{-1}(a) \cap \text{max } D$, then $y(i_k) = a = x(i_k)$, so, since $j_t \leq i_k$ and $x \in D'_l$, also $y(j_t) = x(j_t) \in U_{j_t}$.

(2) Given $x \in D'_l$ (for $l \in \{0, \dots, s\}$), we need to find a \mathcal{C}_k -element C with $x \in C \subseteq D'_l$. Again, write $D'_l = \bigcap_{t=0}^m p_{j_t}^{-1}(U_{j_t}) \cap \text{max } D$, and let $a := x(i_k)$. So $x \in C := p_{i_k}^{-1}(a) \cap \text{max } D \in \mathcal{C}_k$ and, as seen in (1), $C \subseteq D'_l$ (if $y \in C$, then $y(i_k) = a = x(i_k)$, so, since $j_t \leq i_k$ and $x \in D'_l$, also $y(j_t) = x(j_t) \in U_{j_t}$). \square

6.3.6. LEMMA. *If $(\mathfrak{D}_i, p_{ij})_I$ is a finitary dynamical expanding system that is eventually valuation-preserving and $J \subseteq I$ is cofinal, then $(\mathfrak{D}_i, p_{ij})_J$ is again such a diagram.*

PROOF. Then J is a countable directed index set, each \mathfrak{D}_i is a finite max-normalized dynamical Scott domain, and the p_{ij} are dynamical morphisms with the appropriate commutativity conditions. So we need to show that $(\mathfrak{D}_i, p_{ij})_J$ is (i) upward deterministic and (ii) eventually valuation-preserving.

Ad (i). Let $i \in J$ and $a_i, b_i \neq b'_i \in \max D_i$ with $b_i, b'_i \geq f_i(a_i)$. Since the original diagram is upward deterministic, there is $i \leq j \in I$ with $\forall a_j, b_j, b'_j \in \max D_j$:

$$p_{ij}(a_j) = a_i, p_{ij}(b_j) = b_i, p_{ij}(b'_j) = b'_i \Rightarrow b_j \not\geq f_j(a_j) \text{ or } b'_j \not\geq f_j(a_j). \quad (6.4)$$

Since $J \subseteq I$ is cofinal, let $j \leq k \in J$. Let $a_k, b_k, b'_k \in \max D_k$ with $p_{ik}(a_k) = a_i$, $p_{ik}(b_k) = b_i$, and $p_{ik}(b'_k) = b'_i$. Show $b_k \not\geq f_k(a_k)$ or $b'_k \not\geq f_k(a_k)$. Indeed, consider $a_j := p_{jk}(a_k)$, $b_j := p_{jk}(b_k)$, and $b'_j := p_{jk}(b'_k)$. Since p_{jk} is max-preserving, they are in $\max D_j$. They satisfy the condition of (6.4): $p_{ij}(a_j) = p_{ij}(p_{jk}(a_k)) = p_{ik}(a_k) = a_i$ and similarly for b_j and b'_j . So $b_j \not\geq f_j(a_j)$ or $b'_j \not\geq f_j(a_j)$. This implies the claim: If we had $b_k \geq f_k(a_k)$, then, by monotonicity and max-semi-equivariance of p_{jk} ,

$$b_j = p_{jk}(b_k) \geq p_{jk}(f_k(a_k)) \geq f_j(p_{jk}(a_k)) = f_j(a_j),$$

and similarly for b'_k .

Ad (ii). All \mathfrak{D}_i are finite and, for $i \in J$ and $U_i \in \Sigma(D_i)$, there is, since the original diagram is eventually valuation-preserving, some $i \leq j_0 \in I$ such that, for all $j_0 \leq j \in I$ a certain equation $e(j)$ holds.⁴ By cofinality, let $j_0 \leq j_1 \leq J$. So, in particular, for $j_1 \leq j \in J$, equation $e(j)$ holds, as needed. Moreover, for $i \leq j$ in J , if $a_i, b_i \in \max D_i$ with $f_i(a_i) \leq b_i$, then there are, since the original diagram is eventually valuation-preserving, some $a_j, b_j \in \max D_j$ such that $p_{ij}(a_j) = a_i$ and $p_{ij}(b_j) = b_i$ and $f_j(a_j) \leq b_j$, as needed. \square

6.3.7. LEMMA. *Let \mathfrak{X} be in $\text{TS}_{0\text{cm}}$. Let $(\mathfrak{D}_i, p_{ij})_{I(\text{Clp}X)}$ be the diagram with which the observation domain $\hat{D}(\mathfrak{X})$ is constructed. Then, for any finite clopen partition \mathcal{C} of X and $n \geq 0$, we have $i := (n, \mathcal{C}) \in I(\text{Clp}X)$ and*

$$H(n, \mathcal{C}) = - \sum_{a \in \max D_i} v_i(a) \log v_i(a).$$

PROOF. Since \mathcal{C} is a finite clopen partition of X it, in particular, is a finite $\text{Clp}X$ -cover of X , so $i = (n, \mathcal{C}) \in I(\text{Clp}X)$. Note that, since \mathcal{C} is a partition, each $\mathcal{O}_i(y)$ (for $y \in X$) is of the form $\{(C_0^y, \dots, C_{n-1}^y)\}$ where C_k^y is the *unique* \mathcal{C} -cell that contains $T^k(y)$.⁵ Write $\max D_i = \{\{\mathcal{O}_i(x_1)\}, \dots, \{\mathcal{O}_i(x_m)\}\}$ and $\mathcal{O}_i(x_l) = \{(C_0^l, \dots, C_{n-1}^l)\}$. Then $\bigcap_{k=0}^{n-1} T^{-k}C_k^l = [x_l]_i$.⁶

⁴Namely, $v_j(f_j^{-1}(p_{ij}^{-1}(\downarrow \max U_i)) \cap \max D_j) = v_j(p_{ij}^{-1}(\downarrow \max U_i) \cap \max D_j)$.

⁵Proof: If $t \in \mathcal{O}_i(y)$, then, by definition, $t \in \mathcal{C}^n$ and $y \in \bigcap_{k=0}^{n-1} T^{-k}(t(k))$, so, since C_k^y is the unique cell containing $T^k y$, $t(k) = C_k^y$, so $t = (C_0^y, \dots, C_{n-1}^y)$. Conversely, since $t = (C_0^y, \dots, C_{n-1}^y) \in \mathcal{C}^n$ is such that $y \in \bigcap_{k=0}^{n-1} T^{-k}(C_k^y)$, we have $t \in \mathcal{O}_i(y)$.

⁶If $y \in \bigcap_{k=0}^{n-1} T^{-k}C_k^l$, then each C_k^l is the unique \mathcal{C} -cell that contains $T^k(y)$, so, $\mathcal{O}_i(y) = \{(C_0^l, \dots, C_{n-1}^l)\} = \mathcal{O}_i(x_l)$, so $y \in [x_l]_i$. Conversely, if $y \in [x_l]_i$, then $\mathcal{O}_i(y) = \mathcal{O}_i(x_l)$, so y also follows $(C_0^l, \dots, C_{n-1}^l)$, i.e., $y \in \bigcap_{k=0}^{n-1} T^{-k}C_k^l$.

Define \mathcal{C}_\bullet^n as the set of those $(C_0, \dots, C_{n-1}) \in \mathcal{C}^n$ such that $\bigcap_{k=0}^{n-1} T^{-k}(C_k) \neq \emptyset$. Define a function $b : \max D_i \rightarrow \mathcal{C}_\bullet^n$ by

$$b(\{\mathcal{O}_i(x_l)\}) := (C_0^l, \dots, C_{n-1}^l)$$

We claim that this is a bijection. Surjective: For (C_0, \dots, C_{n-1}) , let $y \in \bigcap_{k=0}^{n-1} T^{-k}(C_k)$ and let x_l be such that $y \in [x_l]_i$ (since these equivalence classes partition X). So $\{(C_0, \dots, C_{n-1})\} = \mathcal{O}_i(y) = \mathcal{O}_i(x_l) = \{(C_0^l, \dots, C_{n-1}^l)\}$, whence $b(\{\mathcal{O}_i(x_l)\}) = (C_0, \dots, C_{n-1})$.

Injective: If $\{\mathcal{O}_i(x_l)\} \neq \{\mathcal{O}_i(x_{l'})\}$, then $\bigcap_{k=0}^{n-1} T^{-k}(C_k^l) = [x_l]_i \neq [x_{l'}]_i = \bigcap_{k=0}^{n-1} T^{-k}(C_k^{l'})$, so $(C_0^l, \dots, C_{n-1}^l) \neq (C_0^{l'}, \dots, C_{n-1}^{l'})$.

Recall that, by definition of v_i , we have $v_i(\{\mathcal{O}_i(x_l)\}) = \mu([x_l]_i)$. Also recall the convention $0 \log 0 = 0$. Hence

$$\begin{aligned} - \sum_{a \in \max D_i} v_i(a) \log v_i(a) &= - \sum_{\{\mathcal{O}_i(x_l)\} \in \max D_i} v_i(\{\mathcal{O}_i(x_l)\}) \log v_i(\{\mathcal{O}_i(x_l)\}) \\ &= - \sum_{\{\mathcal{O}_i(x_l)\} \in \max D_i} \mu([x_l]_i) \log \mu([x_l]_i) \\ &= - \sum_{(C_0, \dots, C_{n-1}) \in \mathcal{C}_\bullet^n} \mu\left(\bigcap_{k=0}^{n-1} T^{-k} C_k\right) \log \mu\left(\bigcap_{k=0}^{n-1} T^{-k} C_k\right) \\ &= - \sum_{(C_0, \dots, C_{n-1}) \in \mathcal{C}^n} \mu\left(\bigcap_{k=0}^{n-1} T^{-k} C_k\right) \log \mu\left(\bigcap_{k=0}^{n-1} T^{-k} C_k\right) \\ &= H(n, \mathcal{C}), \end{aligned}$$

as needed. \square

6.3.8. LEMMA. *Let $\mathfrak{X} = (X, \tau, \mu, T)$ be in TS_{ocm} . Let (\mathcal{C}_k) be a refining sequence of measurable partitions that generates $\mathcal{B}(\tau)$. Then*

$$h(\mathfrak{X}) = \liminf_{(n,k) \in \mathbb{N}_+ \times \mathbb{N}} \frac{1}{n} H(n, \mathcal{C}_k).$$

where $\mathbb{N}_+ \times \mathbb{N} = \{(n, k) \in \mathbb{N}^2 : n > 0, k \geq 0\}$ is equipped with the product order (whence directed).⁷

PROOF. Define $x_{n,k} := \frac{1}{n} H(n, \mathcal{C}_k)$, and show $h(\mathfrak{X}) = \sup_{(n,k)} \inf_{(m,l) \geq (n,k)} x_{m,l}$. Recall from the facts about entropy (proposition 6.2.2) that: (a) for each k , $(x_{n,k})_n$ decreases to $\lim_n x_{n,k}$, (b) for $k \leq l$, $\lim_n x_{n,k} \leq \lim_n x_{n,l}$, and (c) $h(\mathfrak{X}) =$

⁷I.e., $(n, k) \leq (m, l)$ iff $n \leq m$ and $k \leq l$.

$\lim_k (\lim_n x_{n,k})$. By (a), we can replace \lim_n by \inf_n , and, by (b), we can replace \lim_k by \sup_k . Further, for any (n, k) ,

$$\begin{aligned} \inf_{(m,l) \geq (n,k)} x_{m,l} &= \inf \{x_{m,l} : m \geq n, l \geq k\} \\ &= \inf \{x_{m,k} : m \geq n\} && (x_{m,l} \text{ increases in } l) \\ &= \inf \{x_{m,k} : m \geq 1\} && (x_{m,l} \text{ decreases in } m) \\ &= \inf_n x_{n,k}. \end{aligned}$$

Hence

$$\sup_{(n,k)} \inf_{(m,l) \geq (n,k)} x_{m,l} = \sup_{(n,k)} \inf_n x_{n,k} = \sup_k \inf_n x_{n,k} = h(\mathfrak{X}),$$

as needed. \square

PROOF OF THEOREM 6.3.4. Let $(\mathfrak{D}_i, p_{ij})_{I(\text{Clp max } D)}$ be the diagram with which the observation domain $\hat{\mathcal{D}}(\mathcal{S}(\mathfrak{D}))$ is constructed using projections p_i . Since \mathfrak{D} is in dDOM_{rv} , we have, by the category equivalence to $\text{TS}_{0\text{c}}$, that there is an isomorphism $u : \mathfrak{D} \rightarrow \hat{\mathcal{D}}\mathcal{S}(\mathfrak{D})$. Further, since \mathfrak{D} is in dDOM_{rv} , it has a diagram giving rise to it with a cofinal chain in the index set, so we can apply lemma 6.3.3 (2) and obtain a refining sequence (\mathcal{C}_k) of clopen partitions of $\text{max } D$ that generates $\mathcal{B}(\text{max } D)$. We can assume that $\mathcal{C}_0 = \{\text{max } D\}$ is the trivial partition. By lemma 6.3.5 we further have that $I := \{(n, \mathcal{C}_k) : n, k \in \mathbb{N}\}$ is cofinal in $I(\text{Clp max } D)$. By lemma 6.3.6, $F := (\mathfrak{D}_i, p_{ij})_I$ is an eventually valuation-preserving finitary dynamical expanding system. And it gives rise to \mathfrak{D} via projections $q_i := p_i \circ u : \mathfrak{D} \rightarrow \mathfrak{D}_i$.⁸

Define the function $t : I \rightarrow \mathbb{N}$ by $t(n, \mathcal{C}_k) = n$. This is a projection: It clearly is monotone, and we find a corresponding embedding as follows. Define $e : \mathbb{N} \rightarrow I$ by $e(n) := (n, \{\text{max } D\})$. By construction, e is monotone and $t \circ e = \text{id}_{\mathbb{N}}$. And $e \circ t \leq \text{id}_I$ since $e \circ t(n, \mathcal{C}_k) = (n, \{\text{max } D\}) \leq (n, \mathcal{C}_k)$ since $\{\text{max } D\} \leq \mathcal{C}_k$ (every element of \mathcal{C}_k is a subset of $\text{max } D$, and for $x \in \text{max } D$, there is $C \in \mathcal{C}_k$ with $x \in C \subseteq \text{max } D$). Note that $I_t = \{(n, \mathcal{C}_k) \in I : n > 0, k \geq 0\}$.

To show the desired equation, we can apply lemma 6.3.8, since $\mathfrak{X} := \mathcal{S}(\mathfrak{D})$ is in $\text{TS}_{0\text{cm}}$ and (\mathcal{C}_k) a refining sequence of finite measurable partitions that generates

⁸The usual argument: $(\mathfrak{D}, q_i)_I$ is a cone to F with \mathfrak{D} in $\text{dSCO}_{\text{nm}}^{\text{p}}$ (we have $p_{ij} \circ q_j = p_{ij} \circ p_j \circ u = p_i \circ u = q_i$). Assume $(\mathfrak{E}, \beta_i)_I$ is another such cone. Then $(\mathfrak{E}, \beta'_i)_{I(\text{Clp max } D)}$ with $\beta'_i := p_{ik_i} \circ \beta_{k_i}$ (for some $i \leq k_i \in I$ with $k_i = i$ if $i \in I$) is a cone to the original diagram: given $i \leq j$ in $I(\text{Clp max } D)$, let $k_i, k_j \leq l \in I$, so $p_{ij} \circ \beta'_j = p_{ij} \circ p_{jk_j} \circ \beta_{k_j} = p_{ij} \circ p_{jk_j} \circ p_{k_j l} \circ \beta_l = p_{il} \circ \beta_l = p_{ik_i} \circ p_{k_i l} \circ \beta_l = p_{ik_i} \circ \beta_{k_i} = \beta'_i$. So there is a unique morphism $\beta' : \mathfrak{E} \rightarrow \hat{\mathcal{D}}\mathcal{S}(\mathfrak{D})$ with $p_i \circ \beta' = \beta'_i$ (for $i \in I(\text{Clp max } D)$). Define $\beta := u^{-1} \circ \beta' : \mathfrak{E} \rightarrow \mathfrak{D}$. Then, for $i \in I$, $q_i \circ \beta = (p_i \circ u) \circ (u^{-1} \circ \beta') = p_i \circ \beta' = \beta'_i = \beta_i$. And if $\gamma : \mathfrak{E} \rightarrow \mathfrak{D}$ is another morphism with $q_i \circ \gamma = \beta_i$, then $u \circ \gamma : \mathfrak{E} \rightarrow \hat{\mathcal{D}}\mathcal{S}(\mathfrak{D})$ is such that, for $i \in I(\text{Clp max } D)$, $p_i \circ (u \circ \gamma) = p_{ik_i} \circ p_{k_i} \circ u \circ \gamma = p_{ik_i} \circ q_{k_i} \circ \gamma = p_{ik_i} \circ \beta_{k_i} = \beta'_i$. So $u \circ \gamma = \beta'$, so, by ‘multiplying’ u^{-1} to the left, $\gamma = u^{-1} \circ \beta' = \beta$.

$\mathcal{B}(X)$, so

$$e(\mathfrak{D}) \stackrel{\text{thm 6.3.2}}{=} h(\mathcal{S}(\mathfrak{D})) = \liminf_{(n,k) \in \mathbb{N}_+ \times \mathbb{N}} \frac{1}{n} H(n, \mathcal{C}_k).$$

We can re-index this by

$$\liminf_{(n,k) \in \mathbb{N}_+ \times \mathbb{N}} \frac{1}{n} H(n, \mathcal{C}_k) = \liminf_{(n, \mathcal{C}_k) \in I_t} \frac{1}{t(n, \mathcal{C}_k)} H(n, \mathcal{C}_k) = \liminf_{i \in I_t} \frac{1}{t(i)} H(i).$$

Since $(\mathfrak{D}_i, p_{ij})_{I(\text{Clp max } D)}$ is the diagram with which the observation domain $\hat{D}(\mathfrak{X})$ is constructed and each \mathcal{C}_k is a clopen partition of $X = \max D$, we can apply lemma 6.3.7, so

$$\liminf_{i \in I_t} \frac{1}{t(i)} H(i) = \liminf_{i \in I_t} \frac{-1}{t(i)} \sum_{a \in \max D_i} v_i(a) \log v_i(a),$$

as needed. \square

6.4 Max-entropy

In this section, we define the domain-theoretic counterpart to topological entropy—which we'll call max-entropy. We first state its definition (section 6.4.1) and then show that it indeed is a proper counterpart (section 6.4.2). We keep the presentation as parallel as possible to that of domain-entropy, but, for reasons of space, we omit a discussion of normal form.

6.4.1 Definition of max-entropy

Analogously to topological (as opposed to metric) entropy, we can drop the assumption of being valuation-preserving when considering max-entropy and hence work in **dDOM**. Thus, we no longer require the diagrams giving rise to a dynamical domains to be eventually valuation-preserving.

6.4.1. DEFINITION. Let $\mathfrak{D} = (D, v, f)$ be in **dDOM** and let $F = (\mathfrak{D}_i, p_{ij})_I$ be a diagram giving rise to \mathfrak{D} via projections p_i . (We ignore the valuation v .) For $i \in I$ and $n \geq 1$, define

$$W(n, i) := \left\{ \bar{a} \in (\max D_i)^n : \exists x \in \max D \forall k = 0, \dots, n-1. p_i(f^k(x)) = \bar{a}(k) \right\}$$

and $M(n, i) := \log |W(n, i)|$. We define the *max-entropy* of \mathfrak{D} as

$$m(\mathfrak{D}) := \sup \left\{ \lim_{n \rightarrow \infty} \frac{1}{n} M(n, i) : i \in I \right\}. \quad (6.5)$$

Below we show that these limits exist. In general, for \mathfrak{D} in \mathbf{dDOM} , we define $m(\mathfrak{D})$ with respect to *some* (F, p_i) giving rise to \mathfrak{D} (below we show that this is independent of the choice of (F, p_i)).

Comments: First, intuitively, $W(n, i)$ is the set of ‘words’ of length n over the finite alphabet consisting of the maximal elements of D_i that can be realized by the projection of the trajectory of some maximal element of D under the domain dynamics. If the dynamics is very chaotic, their number grows exponentially in n , and if the dynamics is very restricted, their number grows slowly. Thus, its ‘average growth rate’ $\lim_{n \rightarrow \infty} \frac{1}{n} \log |W(n, i)|$ provides a good indication of the complexity of the domain dynamics as seen at (i.e., projected to) the i -th component D_i . So the supremum $m(\mathfrak{D})$ of these growth rates describes the complexity of the domain dynamics that we have to reckon with.

Second, hence the intuition is much like that of the topological entropy of a subshift (see, e.g., Lind and Marcus 1995). Except that, as in the general case of topological entropy, we also take a supremum, but now over indices which take over the role of open covers.

6.4.2 Main theorem on max-entropy

The following is the main theorem about max-entropy showing, in particular, that it is well-defined (items (1) and (2) below), an isomorphism invariant as any good notion of entropy should be (items (3) and (4) below), and a counterpart to topological entropy (items (5) and (6) below).

- 6.4.2. THEOREM.**
1. *The limits in equation (6.5) exist.*
 2. *The max-entropy $m(\mathfrak{D})$ is independent of the choice of diagram giving rise to \mathfrak{D} .*
 3. *If $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ is a surjective morphism in \mathbf{dDOM} , then $m(\mathfrak{D}) \geq m(\mathfrak{E})$.*
 4. *In particular, max-entropy is an isomorphism invariant: two dynamical domains that are isomorphic in \mathbf{dDOM} have the same max-entropy.*
 5. *If \mathfrak{D} is in \mathbf{dDOM} , then $m(\mathfrak{D}) = h_{\text{top}}(\mathbf{S}(\mathfrak{D}))$.*
 6. *If \mathfrak{X} is in \mathbf{TS}_{0c} , then $h_{\text{top}}(\mathfrak{X}) = m(\hat{\mathbf{D}}(\mathfrak{X}))$.*

We prove the theorem in the remainder of this subsection. We start with a lemma. Recall that an open cover \mathcal{D} refines another \mathcal{C} (written $\mathcal{C} \preceq \mathcal{D}$) iff every \mathcal{D} -element is a subset of a \mathcal{C} -element.

6.4.3. LEMMA. *Let \mathfrak{D} be in \mathbf{dDOM} and let $(\mathfrak{D}_i, p_{ij})_I$ be a diagram giving rise to \mathfrak{D} via projections p_i . Then*

1. For each $i \in I$, $\mathcal{C}_i := \{p_i^{-1}(a) \cap \max D : a \in \max D_i\}$ is a finite clopen partition of $\max D$.
2. If $i_0 \leq i_1 \leq \dots$ is cofinal in I , then $(\mathcal{C}_{i_k})_k$ is a refining sequence of finite clopen partitions such that, for any open cover \mathfrak{D} of $\max D$, there is k such that \mathcal{C}_{i_k} refines \mathfrak{D} .
3. For each $i \in I$ and $n \geq 0$, we have $H_{\text{top}}(n, \mathcal{C}_i) = M(n, i)$ (where $H_{\text{top}}(n, \mathcal{C}_i)$ is calculated in the system $\mathfrak{S}(\mathfrak{D})$).
4. For $i \leq j$ in I and $n \geq 0$, $M(n, i) \leq M(n, j)$.

PROOF. Ad (1). This is lemma 6.3.3 (1) (which didn't use that the diagram is eventually valuation-preserving).

Ad (2). By lemma 6.3.3 (2), we know that $(\mathcal{C}_{i_k})_k$ is a refining sequence of finite clopen partitions (again, this didn't use that the diagram is eventually valuation-preserving). So they also are refining as open covers.⁹ So let \mathfrak{D} be an open cover of $\max D$ and find k such that $\mathfrak{D} \preceq \mathcal{C}_{i_k}$.

Since each \mathfrak{D} -element is open, it is a union of basic clopen sets. Let \mathfrak{D}' be the collection of these basic clopens. Then $\mathfrak{D} \preceq \mathfrak{D}'$ since each \mathfrak{D}' -element is a subset of a \mathfrak{D} -element. Moreover, \mathfrak{D}' is an open cover of the compact $\max D$, whence there is a finite subcover $\mathfrak{D}'' = \{D_0, \dots, D_m\}$. We have $\mathfrak{D}' \preceq \mathfrak{D}''$ since, trivially, every \mathfrak{D}'' -element is (a subset of) a \mathfrak{D}' -element.

Since \mathfrak{D}'' is a finite $\text{Clp } \max D$ -cover, $(0, \mathfrak{D}'') \in I(\text{Clp } \max D)$. In lemma 6.3.5, we've show (without using that the diagram is eventually valuation-preserving) that $\{(n, \mathcal{C}_{i_k}) : n, k \in \mathbb{N}\}$ is cofinal in $I(\text{Clp } \max D)$. In particular, there is k such that $\mathcal{C}_{i_k} \succeq \mathfrak{D}''$. This implies $\mathcal{C}_{i_k} \succeq \mathfrak{D}'' \succeq \mathfrak{D}$, as needed.

Ad (3). Write $\mathcal{C} := \mathcal{C}_i$. Recall that $H_{\text{top}}(n, \mathcal{C}) = \log N(\mathcal{C}_0^{n-1})$. We have, writing $\bar{f} := f \upharpoonright \max D$,

$$\begin{aligned} \mathcal{C}_0^{n-1} &= \left\{ \bigcap_{k=0}^{n-1} \bar{f}^{-k} U_k : U_0, \dots, U_{n-1} \in \mathcal{C} \right\} \\ &= \left\{ \bigcap_{k=0}^{n-1} \bar{f}^{-k} (p_i^{-1}(\bar{a}(k)) \cap \max D) : \bar{a} \in (\max D_i)^n \right\}. \end{aligned}$$

Since \mathcal{C} is a (clopen) partition, this is a partition. So, after possibly discarding the empty set, \mathcal{C}_0^{n-1} has no nontrivial subcovers, i.e., $N(\mathcal{C}_0^{n-1}) = |\mathcal{C}_0^{n-1} \setminus \{\emptyset\}|$.

So it suffices to show that there is a bijection $b : W(n, i) \rightarrow \mathcal{C}_0^{n-1} \setminus \{\emptyset\}$. Then $|W(n, i)| = |\mathcal{C}_0^{n-1} \setminus \{\emptyset\}| = N(\mathcal{C}_0^{n-1})$, so $M(n, i) = H_{\text{top}}(n, \mathcal{C})$.

⁹Proof: If \mathfrak{D} is a finite clopen partition that partition-refines another finite clopen partition \mathcal{C} , it also cover-refines it: Let $D \in \mathfrak{D}$. Pick some $x \in D$. Then $x \in C$ for some $C \in \mathcal{C}$. By partition-refinement, write $C = D_1 \cup \dots \cup D_n$. Then $x \in D_i$ for some $i \in \{1, \dots, n\}$. But then $D = D_i$ (if $D_i \neq D$, then x would be in two partition cells). So $D \subseteq C$, as needed.

Indeed, define $b : W(n, i) \rightarrow \mathcal{C}_0^{n-1} \setminus \{\emptyset\}$ by

$$b(\bar{a}) := \bigcap_{k=0}^{n-1} \bar{f}^{-k}(p_i^{-1}(\bar{a}(k)) \cap \max D).$$

Well-defined: If $\bar{a} \in W(n, i)$ there is, by definition, $x \in \max D$ such that, for all $k = 0, \dots, n-1$, $p_i(\bar{f}^k(x)) = p_i(f^k(x)) = \bar{a}(k)$. So $\bigcap_{k=0}^{n-1} \bar{f}^{-k}(p_i^{-1}(\bar{a}(k)) \cap \max D)$ is nonempty and hence in $\mathcal{C}_0^{n-1} \setminus \{\emptyset\}$.

Surjective: If $\bigcap_{k=0}^{n-1} \bar{f}^{-k}(p_i^{-1}(\bar{a}(k)) \cap \max D) \in \mathcal{C}_0^{n-1} \setminus \{\emptyset\}$ for $\bar{a} \in (\max D_i)^n$, then, since this intersection is nonempty, there is $x \in \max D$ such that, for all $k = 0, \dots, n-1$, we have $p_i(f^k(x)) = \bar{a}(k)$, whence $\bar{a} \in W(n, i)$ and $b(\bar{a}) = \bigcap_{k=0}^{n-1} \bar{f}^{-k}(p_i^{-1}(\bar{a}(k)) \cap \max D)$.

Injective: If $\bar{a} \neq \bar{a}'$ in $W(n, i)$, there is $k \in \{0, \dots, n-1\}$ such that $\bar{a}(k) \neq \bar{a}'(k)$. If $b(\bar{a}) = b(\bar{a}')$, then there is $x \in b(\bar{a}) \cap b(\bar{a}')$, whence $\bar{a}(k) = p_i(\bar{f}^k(x)) = \bar{a}'(k)$, contradiction.

Ad (4). By (2), \mathcal{C}_j refines \mathcal{C}_i as open cover, so, by fact (1) about entropy (proposition 6.2.4) and by (3),

$$M(n, i) = H_{\text{top}}(n, \mathcal{C}_i) \leq H_{\text{top}}(n, \mathcal{C}_j) = M(n, j),$$

as needed. □

PROOF OF THEOREM 6.4.2. We'll prove the items in a different order than listed, because to show that max-entropy is independent of the diagram, we need diagram-dependent versions of (3)–(5) which we mark with an asterisk and state precisely below.

Ad (1). In the setting of definition 6.4.1, we need to show, for $i \in I$, that $\lim_n \frac{1}{n} M(n, i)$ exists. By lemma 6.4.3 (3), we have

$$\lim_n \frac{1}{n} M(n, i) = \lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C}_i)$$

and the latter exists by fact (2) about entropy (proposition 6.2.4).

Now, for (F, p_i) giving rise to \mathfrak{D} , the extended real $m(\mathfrak{D}, p_i, F)$ in equation (6.5) is well-defined.

Claim (5)*: If (F, p_i) gives rise to \mathfrak{D} , then $m(\mathfrak{D}, p_i, F) = h(\mathbf{S}(\mathfrak{D}))$.

Indeed, let $i_0 \leq i_1 \leq \dots$ be cofinal in the index set I of F (I is countable and directed). Write $\mathfrak{X} := \mathbf{S}(\mathfrak{D})$. By lemma 6.4.3 (2), (\mathcal{C}_{i_k}) is a refining sequence of finite clopen partitions that eventually refine every cover. By fact (4) about entropy, $h_{\text{top}}(\mathfrak{X}) = \lim_k (\lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C}_{i_k}))$. Since the terms are increasing by fact (3), we can replace the 'lim_k' by 'sup_k'. By lemma 6.4.3 (3), this then further equals $\sup_k (\lim_n \frac{1}{n} M(n, i_k))$. So it remains to show

$$r := \sup_k \left(\lim_n \frac{1}{n} M(n, i_k) \right) = \sup \left\{ \lim_n \frac{1}{n} M(n, i) : i \in I \right\} \quad (= m(\mathfrak{D}, p_i, F)).$$

By the subset relation, we have \leq . For the other direction, let $i \in I$ and show $\lim_n \frac{1}{n} M(n, i) \leq r$. By cofinality, there is $i_k \geq i$. By lemma 6.4.3 (4), we have, for all $n \geq 0$, that $M(n, i) \leq M(n, i_k)$. So $\lim_n \frac{1}{n} M(n, i) \leq \lim_n \frac{1}{n} M(n, i_k) \leq r$.

Claim (3)*: If $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ is in \mathbf{dDOM} with $\alpha(\max D) = \max E$, and (F, p_i) and (G, q_j) give rise to \mathfrak{D} and \mathfrak{E} , respectively, then $m(\mathfrak{D}, p_i, F) \geq m(\mathfrak{E}, q_j, G)$.

Indeed, then $S(\alpha) : S(\mathfrak{D}) \rightarrow S(\mathfrak{E})$ is a surjective morphism in \mathbf{TS}_{0c} , so by fact (5) and by (5)*,

$$m(\mathfrak{D}, p_i, F) = h(S(\mathfrak{D})) \geq h(S(\mathfrak{E})) = m(\mathfrak{E}, q_j, G).$$

Claim (4)*: If \mathfrak{D} and \mathfrak{E} are isomorphic in \mathbf{dDOM} with (F, p_i) and (G, q_j) giving rise to them, respectively, then $m(\mathfrak{D}, p_i, F) = m(\mathfrak{E}, q_j, G)$.

Indeed, then there are, in particular, morphisms $\alpha : \mathfrak{D} \rightarrow \mathfrak{E}$ and $\beta : \mathfrak{E} \rightarrow \mathfrak{D}$ in \mathbf{dDOM}_v with $\alpha(\max D) = \max E$ and $\beta(\max E) = \max D$, so, by (3)*, $m(\mathfrak{D}, p_i, F) \geq m(\mathfrak{E}, q_j, G)$ and $m(\mathfrak{E}, q_j, G) \geq m(\mathfrak{D}, p_i, F)$, as needed.

Ad (2). Let (F, p_i) and (G, q_j) give rise to \mathfrak{D} . Since \mathfrak{D} is isomorphic to itself, (4)* implies $m(\mathfrak{D}, p_i, F) = m(\mathfrak{D}, q_j, G)$.

Hence, defining, for \mathfrak{D} in \mathbf{dDOM} , $m(\mathfrak{D}) := m(\mathfrak{D}, p_i, F)$ for some (F, p_i) giving rise to \mathfrak{D} is well-defined.

Now, (3)–(5) are implied by their asterisked versions. So it remains to show the last item.

Ad (6). Let \mathfrak{X} be in \mathbf{TS}_{0c} . Since the counit of the adjunction $\hat{D} \vdash S$ is a natural isomorphism, \mathfrak{X} is isomorphic in \mathbf{TS}_{0c} to $\hat{S}(\mathfrak{X})$. Since h_{top} is an isomorphism invariant, $h_{\text{top}}(\mathfrak{X}) = h_{\text{top}}(\hat{S}(\mathfrak{X}))$. Further, since $\hat{D}(\mathfrak{X})$ is in \mathbf{dDOM} , we have, by (5), $h_{\text{top}}(\hat{S}(\mathfrak{X})) = m(\hat{D}(\mathfrak{X}))$, as needed. \square

6.5 Conclusion

Our aim was to indicate the potential of the adjunction between dynamical domains and dynamical systems. We did so by determining domain-theoretic counterparts to the important system-theoretic concept of entropy: i.e., independently defined invariants of dynamical domains that turn out to correspond to system-theoretic entropy along the adjunction. (A worked example is given in appendix B of the thesis.) This clearly is but a first step. Here are some directions for future work.

First, how independent is the static representation of entropy given by the normal formal theorem from the diagram and the time function? Second, develop a normal form for max-entropy. Third, how do these domain-theoretic notions of entropy relate to domain-theoretic constructions (product, sum, fixed-point, powerdomain, etc.), and can it be described categorically as a fixed-point operator? Fourth, famously, variational principles relate metric notions of entropy to topological notions of entropy. (The classic theorem is that topological entropy is the supremum of the metric entropies of invariant measures.) What are

analogous principles for the domain-theoretic notions of entropy? Fifth, Ornstein theory describes when entropy is a complete isomorphism invariant. What is the domain-theoretic counterpart? When does identical domain-entropy imply modeling metrically isomorphic systems? Sixth, entropy has also been discussed in the context of computation (Fredkin and Toffoli 1982). How does this relate to the computational interpretation of dynamical domains as computational models for dynamical systems describing, for example, physical computation?

Appendix

PROOF SKETCH OF PROPOSITION 6.2.2. We only mention the main ideas (or facts), but provide detailed references to Walters (1982, sec. 4).

To adjust notation, note that there is a one-to-one correspondence between finite partitions of a probability space and finite sub- σ -algebras (see Walters 1982, sec. 4.1): If \mathcal{C} is a finite measurable partition, let $\mathcal{A}(\mathcal{C})$ be the finite σ -algebra of unions of sets from \mathcal{C} . (Conversely, for each finite sub- σ -algebra \mathcal{A} there is a natural partition $\xi(\mathcal{A})$ and the two operations \mathcal{A} and ξ are inverse to each other.) It is readily seen that $\mathcal{C} \preceq \mathcal{D}$ iff $\mathcal{A}(\mathcal{C}) \subseteq \mathcal{A}(\mathcal{D})$ (Walters 1982, def. 4.2). So entropy could be developed with either concept. Walters (1982, sec. 4) uses the σ -algebra notation more (though not exclusively). But the partition notation will be more convenient for us.

Here is how to adjust the notation: $H(1, \mathcal{C})$ in our notation is $H(\mathcal{A}(\mathcal{C})) := H(\xi \mathcal{A}(\mathcal{C})) = H(\mathcal{C}) = -\sum_{C \in \mathcal{C}} \mu(C) \log \mu(C)$ in the notation of Walters (1982, def. 4.6). And $H(n, \mathcal{C})$ in our notation is $H(\bigvee_{k=0}^{n-1} T^{-k}(\mathcal{A}(\mathcal{C})))$ in the notation of Walters (1982, def. 4.9) where $\bigvee_{k=0}^{n-1} \cdot = \sigma(\bigcup_{k=0}^{n-1} \cdot)$. This is because, as noted by Walters (1982, above def. 4.9), $\xi \bigvee_{k=0}^{n-1} T^{-k}(\mathcal{A}(\mathcal{C}))$ is the partition provided by the sets of the form $\bigcap_{k=0}^{n-1} T^{-k} C_k$ for $C_k \in \mathcal{C}$. As Walters (1982, def. 4.10) notes, entropy can be defined equivalently as supremum over all finite sub- σ -algebras or over all finite partitions, so our $h(\mathfrak{X})$ is identical to $h(T)$ there.

Ad (1). If $\mathcal{C} \preceq \mathcal{D}$, then $\mathcal{A}(\mathcal{C}) \subseteq \mathcal{A}(\mathcal{D})$, so, since $H(\cdot)$ is readily seen to be \preceq -monotone, $H(n, \mathcal{C}) = H(\bigvee_{k=0}^{n-1} T^{-k}(\mathcal{A}(\mathcal{C}))) \leq H(\bigvee_{k=0}^{n-1} T^{-k}(\mathcal{A}(\mathcal{D}))) = H(n, \mathcal{D})$ (Walters 1982, proof of thm. 4.12 (iii)).

Ad (2). Walters (1982, thm. 4.10) shows that $\frac{1}{n} H(\bigvee_{k=0}^{n-1} T^{-k} \mathcal{A}(\mathcal{C})) = \frac{1}{n} H(n, \mathcal{C})$ decreases to $\lim_n \frac{1}{n} H(\bigvee_{k=0}^{n-1} T^{-k} \mathcal{A}(\mathcal{C}))$. The quick way to show that the limit exists (without the ‘decreasing’) is by a subadditivity argument: Show that the sequence $(H(n, \mathcal{C}))_n$ is subadditive (i.e., $H(n+m, \mathcal{C}) \leq H(n, \mathcal{C}) + H(m, \mathcal{C})$) and conclude with Fekete’s Subadditive Lemma that $\lim_n \frac{1}{n} H(n, \mathcal{C})$ exists.

Ad (3). Immediate from (1).

Ad (4). If (\mathcal{C}_i) is a refining sequence of finite partitions generating $\mathcal{B}(\tau)$, then $(\mathcal{A}(\mathcal{C}_k))_k$ is an increasing sequence of finite σ -algebras with $\bigvee_{k=0}^{\infty} \mathcal{A}(\mathcal{C}_k) := \sigma(\bigcup_{k=0}^{\infty} \mathcal{A}(\mathcal{C}_k)) = \sigma(\bigcup_{k=0}^{\infty} \mathcal{C}_k) = \mathcal{B}(\tau)$. A standard result about entropy then is that the entropy of T with respect to the $\mathcal{A}(\mathcal{C}_k)$ converges to the entropy of

T (Walters 1982, thm. 4.22):

$$h(\mathfrak{X}) = h(T) = \lim_k \left(\lim_n \frac{1}{n} H \left(\bigvee_{k=0}^{n-1} T^{-k} \mathcal{A}(\mathcal{C}_k) \right) \right) = \lim_k \left(\lim_n \frac{1}{n} H(n, \mathcal{C}_k) \right).$$

Ad (5). If $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is in $\text{TS}_{0\text{cm}}$ with $\mathfrak{X} = (X, \tau, \mu, T)$ and $\mathfrak{Y} = (Y, \sigma, \nu, S)$, then φ is a factor $(X, \mathcal{B}(\tau), \mu, T) \rightarrow (Y, \mathcal{B}(\sigma), \nu, S)$. It again is a basic property of entropy that it at most decreases along factors (Walters 1982, comment after thm. 4.11). So $h(\mathfrak{X}) = h(T) \geq h(S) = h(\mathfrak{Y})$. \square

PROOF SKETCH OF PROPOSITION 6.2.4. We again only mention the main ideas with detailed references to (Walters 1982, ch. 7). The only notational difference is that we write \mathcal{C}_0^{n-1} instead of $\bigvee_{k=0}^{n-1} T^{-k} \mathcal{C}$.

Ad (1). If $\mathcal{C} \preceq \mathcal{D}$ (as open covers), then also $\mathcal{C}_0^{n-1} \preceq \mathcal{D}_0^{n-1}$, whence $N(\mathcal{C}_0^{n-1}) \leq N(\mathcal{D}_0^{n-1})$ (if $\{V_1, \dots, V_n\}$ is a subcover of \mathcal{D}_0^{n-1} with minimal cardinality, pick, for $i = 1, \dots, n$, some $U_i \in \mathcal{C}_0^{n-1}$ with $U_i \supseteq V_i$, so $\{U_1, \dots, U_n\}$ is a subcover of \mathcal{C}_0^{n-1}), so $H_{\text{top}}(n, \mathcal{C}) \leq H_{\text{top}}(n, \mathcal{D})$. (See Walters (1982), proofs of remarks (3) and (7) on pp. 165–166.)

Ad (2). The existence of the limit $\lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C})$ is again shown by a subadditivity argument. See Walters (1982, thm. 7.1).

Ad (3). Immediate from (1).

Ad (4). Since (\mathcal{C}_k) is refining, the sequence $(\lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C}_k))$ is increasing by (3), so its limit is its supremum. So we need to show $\sup_{\mathcal{C}} \lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C}) = \sup_k \lim_n \frac{1}{n} H_{\text{top}}(n, \mathcal{C}_k)$. Here \geq follows qua subset, and \leq follows by (3) together with (\mathcal{C}_k) eventually refining any open cover. (Also see, e.g., Downarowicz 2011, rmk. 6.1.7.)

Ad (5). If $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is a surjective morphism in $\text{TS}_{0\text{c}}$, then φ is a surjective, continuous, and equivariant map, so (similar to factors in the measure-theoretic case), $h(\mathfrak{X}) \geq h(\mathfrak{Y})$. See Walters (1982, thm. 7.2). \square

Part Three

Stability

Chapter 7

Interlude: symbolic vs. non-symbolic

Abstract In this informal interlude, we wrap up the previous two parts, and we see how this leads us to the questions of the present third and final part. Wrapping up, the work so far suggests the thesis that non-symbolic computation is the limit of symbolic computation (section 7.1 below). This then poses the converse question: when can non-symbolic computation be regarded as realizing symbolic computation (section 7.2)? As a guiding intuition, we suggest that the system’s behavior should be fairly stable. We explore this idea using the concepts of ergodicity and randomness. In the next chapter, we start studying the idea in detail with an analysis of the concept of stability that it involves.

7.1 Non-symbolic computation as limit of symbolic computation

We describe how our results support the intuitive thesis that non-symbolic computation is the limit of symbolic computation. This is summarized in figure 7.1, which we now explain.

Let \mathfrak{X} be a dynamical system. For convenience, say \mathfrak{X} is compact and zero-dimensional, i.e., in \mathbf{TS}_{0c} . (Then we don’t need to specify a choice of basis, but of course our results also hold in the general case when \mathfrak{X} is in \mathbf{DS} .) Then we can apply the ‘semantic denotation’ functor \hat{D} to obtain the dynamical domain $\hat{D}(\mathfrak{X})$, which provides a computational model for \mathfrak{X} .

This already provides a sense in which the non-symbolic computation provided by \mathfrak{X} is a limit of symbolic computation: \mathfrak{X} is, up to isomorphism, the system $S\hat{D}(\mathfrak{X})$ modeled by the dynamical domain $\hat{D}(\mathfrak{X})$. And the states of this system are the maximal elements of the domain D underlying $\hat{D}(\mathfrak{X})$. These maximal elements, in turn, correspond to maximal ideals of compact elements of D . Recall that the compact elements are the ‘finite’ and ‘directly accessible’ elements of the domains and thus are much like symbols. Intuitively, they represent the finite outputs of

	<i>Systems</i>		<i>Domains</i>
<i>Limit (non-symbolic)</i>	\mathfrak{X}	$\stackrel{\hat{D}}{\cong}$	$\hat{D}(\mathfrak{X})$
<i>Approximation (symbolic)</i>	M_i	$\stackrel{T}{\cong}$	$T(M_i) \cong D_i$

Figure 7.1: Non-symbolic computation as limit of symbolic computation.

symbolic computational processes: here, the process of finite observations. So the states of $\mathfrak{X} \cong \hat{S}\hat{D}(\mathfrak{X})$ can be regarded as limits—or infinite sequences—of these symbolic outputs.

We may call this the domain-theoretic sense in which \mathfrak{X} is a limit of symbolic computation. It regards states of \mathfrak{X} as limits of (outputs of) symbolic computation. There also is an additional sense: the category-theoretic one.¹ It regards the whole system \mathfrak{X} as a limit of structures describing symbolic computation. Qua dynamical domain, $\hat{D}(\mathfrak{X})$ is the restricted limit of the finite dynamical dcpos \mathfrak{D}_i obtained from finite observations of the system \mathfrak{X} . In particular, the domain D underlying $\hat{D}(\mathfrak{X})$ is the limit of the domains D_i underlying \mathfrak{D}_i . (In domain-theoretic terminology, D is the bilimit of the D_i because it is, in category-theoretic terminology, both the limit and the colimit.) Since D_i is a finite Scott domain, it is an initialized domain (choosing the least element as initial element), so we can build the BTS $M_i := \mathbf{B}(D_i)$. Its states are the elements of D_i and $x \rightarrow y$ iff $x \leq y$. And the trajectory domain $T(M_i)$ is, after discarding the empty trajectory, isomorphic to D_i . Thus, qua finite BTS, M_i can be regarded as describing symbolic computation. And the M_i can be regarded as approximating in the limit the system \mathfrak{X} in the sense that the limit of their trajectory domains is the domain modeling \mathfrak{X} .²

So far, this shows that non-symbolic computation can be regarded as a limit of symbolic computation: every dynamical system is isomorphic to the system modeled by a restricted limit of a finitary dynamical expanding system of max-reflective dynamical dcpos. In other words, limits of symbolic computation provide an upper bound to what can be done by non-symbolic computation. However, our results also suggest that this upper bound is achieved: every limit of symbolic computation also constitutes non-symbolic computation. Indeed, every restricted limit of a finitary dynamical expanding system (of max-reflective dynamical dcpos) is a (max-reflective) dynamical domain and hence models a dynamical system. (In fact, the category of max-reflective dynamical domains \mathbf{dDOM}_r is isomorphic to the category \mathbf{TS}_{0c} of topological systems.)

This is somewhat similar to ‘pro-categories’ (in the sense of a category of

¹A powerful tool in the theory of solving recursive domain equations is that, in certain important cases (Scott information systems), these two senses are closely related: the category of domains itself can be essentially regarded again as a domain (Winskel 1993, ch. 12).

²Also cf. the results mentioned in the discussion of related work in chapter 4 on obtaining certain topological systems as inverse limits of finite graphs.

pro-objects). Examples are profinite groups, profinite graphs, or profinite spaces. They, too, are the limits of finite groups, graphs, and spaces, respectively. The only difference is that here we take the restricted limit (rather than the limit) and we don't take it of any diagram but only of those diagrams that form a finitary dynamical expanding system. Modulo this caveat, we could sharpen the thesis to:

Non-symbolic computation is profinite symbolic computation.

Or, more precisely: non-symbolic computation is pro-symbolic computation.^{3,4}

Future work should further explore this way of obtaining non-symbolic computation as the limit of symbolic computation. For example, the additional domain-theoretic structure on D_i (i.e., the valuation $v_i : \Sigma(D_i) \rightarrow [0, 1]$ and the domain dynamics $f_i : D_i \rightarrow D_i$) should be represented in the BTS M_i as well. This could be achieved, for example, by (1) adding the transitions $x \xrightarrow{v_i(\{x\})} x$ for those states x of M_i that are maximal elements in D_i , and by (2) adding the transitions $x \xrightarrow{f_i} y$ if $f_i(x) = y$. Moreover, the established adjunctions $\text{TS}_{0c} \rightleftharpoons \text{dDOM}$ and $\omega\text{BTS}_a^s \rightleftharpoons \text{iALG}$ and their limit-preserving properties should also be useful for this.⁵

7.2 Non-symbolic realization of symbolic computation

Concerning the relation between symbolic and non-symbolic computation, the previous section suggests that we can think of non-symbolic computation as the limit of symbolic computation. This naturally poses the question whether there also is a relation in the other direction: When can we think of symbolic computation as being realized by non-symbolic computation?⁶ In other words, when can non-symbolic computation be approximated by (or interpreted as) symbolic computation?

Typically, this is the case when the dynamical system has a fairly stable (i.e., non-chaotic) behavior.⁷ Here are two examples.

³Also cf. profinite automata in automata theory (Pin 2009; Rowland and Yassawi 2017).

⁴In an adventurous mood, this makes one ponder how far this idea reaches: Continuous ('classical') dynamical systems being the limit of finite/discrete ('quantum') order structures is reminiscent of the causal set theory approach to quantum gravity or of quantum cellular automata. Cf. Gogioso, Stasinou, and Coecke (2020).

⁵Also note that the limit Scott domain D can be represented as the BTS describing the Scott information system giving rise to D .

⁶In the introduction (chapter 1, footnote 18), we've already mentioned literature on physical computation, i.e., realizing an abstract symbolic computation in a physical system (whence in a non-symbolic system).

⁷The terms 'stable' and 'non-chaotic' are understood here in an intuitive and not strictly formal sense. (Though there are, of course, various formalizations of these concepts.)

A paradigm case is provided by cognitive systems like us: Crudely put, the neural network involved in our cognition of objects is a physical dynamical system. But we can also describe its behavior symbolically, for example, by “if an object has the shape of a cylinder with a handle, it is a cup”. The reason is that this encapsulates a stable behavior of the system: if the system is in a state with a retinal input containing a cylinder-shaped object with a handle, then—for a wide range of size, color, location, perspective, etc.—the system will usually move to a state where the representation of the concept ‘cup’ is active. More generally, the stability in the behavior of the system allows a symbolic description of the form “on this kind of input state, the system usually moves to that kind of output state”. As indicated by the ‘usual’, this should be taken as a defeasible conditional, i.e., a rule that may have exceptions—see Leitgeb (2005).⁸

Another important and related example is learning. As we see more and more instances (data points) of a concept—say, the concept ‘cup’—we form a more and more refined representation of the concept. And usually this dynamics is fairly stable: the convergence to the concept ‘cup’ isn’t highly dependent on the initial state (e.g., which concepts we already know) and it isn’t highly dependent on the exact instances we see (we all learn the concept of a cup from different examples). This is true both for human cognition but also for learning in artificial neural networks. Although, for artificial neural networks, learning is much of a ‘black box’ process which we’d like to better understand by, e.g., finding a more symbolic description.⁹

In this section, we informally explore this idea of stability allowing a symbolic approximation. In the next chapter, we start doing this in detail by analyzing the underlying concept of stability.

7.2.1 Symbolic approximation

By our results, a dynamical system \mathfrak{X} has a computational model \mathfrak{D} which is obtained as a limit of finite models. So we may think that \mathfrak{X} realizes a symbolic computation if that limit is already reached after finitely many steps: i.e., the dynamical system is modeled by a finite (and hence symbolic) model. Although this cannot literally be true (otherwise the system is finite to start with), this seems to be on the right track. So let’s explore—by means of an example—the question of when a finite model is already a *good enough* approximation to the

⁸For example, as demonstrated by Warrington and Taylor (1973), some subjects may correctly recognize an object—say, a bucket—when seen from a conventional perspective but err on an unconventional perspective—say, a bucket seen from the top.

⁹ One might also explore whether the ‘better understandability’ of stable systems (that allow for a symbolic approximation) has something to do with their reduced complexity. At least in the verification of hybrid systems, checking satisfiability of ‘robust’ formulas/properties (those that are stable under small perturbations) becomes tractable (Franek, Ratschan, and Zgliczynski 2016; Ratschan 2010).

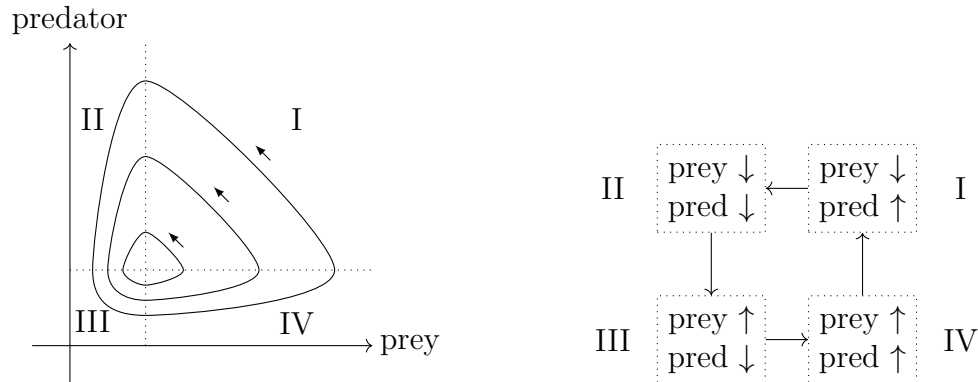


Figure 7.2: The state space of the predator-prey dynamics (left) and its symbolic representation (right).

limit.

As a representative example, let's not choose something as complicated as a cognitive (learning) system, but let's also avoid something too trivially stable. A classic example that strikes this balance is the predator-prey model, as depicted in figure 7.2. (For an overview, see, e.g., Hoppensteadt (2006).) When abstracting away its quantitative description (provided as a solution to a differential equation), this dynamics is qualitatively described as follows. In area I, as the predators feed on the prey, the size of the prey population decreases while the predators multiply. Until, in II, prey becomes scarce and the predator population decreases. With fewer predation, in III, the prey population can recover. So, in IV, the growing prey population can sustain a growing predator population—and the cycle starts anew.

So the dynamical system 'computes' an oscillation of two variables (the prey population size and the predator population size) that depend on each other (there is a phase shift between the oscillations: a prey increase is followed by a later predator increase). More interesting and also more 'computational' examples would be models of a neuron like the Hodgkin–Huxley model or the FitzHugh–Nagumo model that compute various firing patterns (oscillations) of the neuron.¹⁰ However, for our purposes the simpler predator-prey model is enough.

What does the domain-theoretic representation of this dynamical system look like, and what does a finite symbolic approximation to it look like? Since the predator-prey model is a continuous time system, we pick a discretization of time (and a probability measure) to obtain a dynamical system $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ in our sense. Assume for simplicity that the time-discretization is such that the successor

¹⁰The latter also is two-dimensional where one variable describes the membrane voltage of the neuron while the other variable describes a recovery variable. See, e.g., Izhikevich and FitzHugh (2006).

of a state in area I (resp., II, III, IV) is in II (resp., III, IV, I). Assume we observe the system with observation time $n = 1$ and measurements $\mathcal{C} = \{I, II, III, IV\}$. Then the ‘observation history’ $\mathcal{O}_{\mathcal{C}}^n(x)$ of a state x is essentially just the area in which the state is in. And by the assumption about time-discretization, the mapping $\mathcal{O}_{\mathcal{C}}^n(x) \mapsto \mathcal{O}_{\mathcal{C}}^n(T(x))$ is well-defined. So, writing $i = (n, \mathcal{C})$, the observation domain D_i is (isomorphic to) $\mathcal{P}(\{I, II, III, IV\})$, i.e., the power set of $\{I, II, III, IV\}$ minus the empty set ordered by reverse inclusion. And its dynamics is induced by

$$I \mapsto II \qquad II \mapsto III \qquad III \mapsto IV \qquad IV \mapsto I$$

So the maximal elements of D_i correspond to the states of the symbolic representation in figure 7.2. Its dynamics is max-preserving and bijective on the maximal elements. This corresponds to the fact that the transition relation in the symbolic representation is a bijective function.

Thus, it is suggestive to take D_i with this dynamics to be a good enough approximation to the limit D : it provides a good representation of the qualitative behavior of the system and its dynamics is max-preserving (and bijective on maximal elements).

This is a first suggestion for finding good enough symbolic approximation. But before we ask whether this idea generalizes, we should discuss its feasibility.

7.2.2 Ergodicity

Dynamical systems are rarely as well-understood and analytically analyzed as the above predator-prey model. So, given a less-understood system, how can we know whether it has the kind of (stable enough) behavior that allows a symbolic description as, e.g., in the predator-prey model? How can we know that some stage of the limit construction already provides a good symbolic approximation?

To answer this, we need tools to understand the long-term behavior of the system. Ergodic theory precisely provides such tools. So we introduce it briefly to see what it—together with our results—has to say about the ‘fairly stable’ systems that we’re considering here.

We follow the concise story of (Eisner et al. 2015, ch. 1). The starting point of ergodic theory was physical systems like a box of gas. It consists of a state space X where a state $x \in X$ is a vector in \mathbb{R}^{6n} describing the position and momentum in three-dimensional space of each of the n -many gas particles (hence $2 \times 3 \times n = 6n$). The dynamics $T : X \rightarrow X$ is given by the laws of classical mechanics dictating how the particles move over time: if they are in state x at time $t = 0$, they will be in state $T(x)$ at time $t = 1$, and in state $T^2(x)$ at time $t = 2$, etc. While theoretically deep, this is practically not yet very useful due to the following three problems. Their solution then started ergodic theory.

Problem 1 It is practically and physically not possible to uniquely determine the current state x of the system, from which we then could compute the long-term behavior $x, T(x), T^2(x), \dots$ (the *orbit* of x).

Solution All we can do is *measurements*. For example, we could measure the temperature of the gas and see how this evolves over time. More generally, a measurement is a function $f : X \rightarrow \mathbb{R}$: if the system is in state x and we perform measurement f , we measure the value $f(x)$. At the next time step, we then measure $f(T(x))$, and after that, we measure $f(T^2(x))$, etc. So the evolution of the measurements is given by the function $f \circ T^n$.

Problem 2 The time steps in which the system is updated are much shorter than those in which we can perform our measurements. So we cannot measure frequently enough to observe the evolution $f(x), f(T(x)), f(T^2(x)), \dots$ but only some subsequence thereof.

Solution Based on these sampled subsequences of measurements, we only can get a good indication of the average values over time:

$$\frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x))$$

and their limit as $n \rightarrow \infty$. This limit is the *time average* of the measurement f starting with state x .

Problem 3 The time average still depends on the initial state x which we cannot practically determine.

Solution As a response, Boltzmann formulated the *ergodic hypothesis*. The system X has a natural probability measure μ which is preserved by the transformation T (coming from Liouville's theorem in classical mechanics). So Boltzmann hypothesized that the time average of any measurement f starting in any initial state x should be the *space average* $\int f d\mu$, i.e.,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) = \int_X f d\mu. \quad (7.1)$$

If this is true, then we can get our desired knowledge about the long-term behavior of the system: we can get to know the long-term value of an evolution of a measurement starting in an initial state by estimating the expected value of a one-time performance of this measurement (with respect to the probability measure). And this now is independent of the initial state.

Thus, this allows a description of the long-term behavior of the system without reference to specific points in space or time.

The intuition behind the ergodic hypothesis is that the system is 'chaotic' enough so that the orbit of a state x eventually visits every area of the state space—in fact, with a relative frequency given by the measure of that area. This has the

following effect: In the time average, we compute the averages $\frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x))$, and take their limit as n goes to infinity. And in the space average, we compute the value of f in the various areas of the state space and sum them weighted by their probability, and we roughly take the limit as the areas get smaller and smaller. Thus, if the orbit of x visits each area with the relative frequency provided by the measure, we'd expect these calculated values to become more and more similar until they are identical in the limit.

Ergodic theory develops these intuitions into a mathematical theory. Dynamical systems are assumed to be structures $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a (Lebesgue) probability space and $T : X \rightarrow X$ is a measure-preserving bijective function. The system is called *ergodic* if any measurable invariant set has probability 0 or 1 (i.e., if $A \in \mathcal{A}$ with $T^{-1}(A) = A$, then $\mu(A) \in \{0, 1\}$). Thus, being ergodic is one way of spelling out that the system moves through all areas of the state space. If \mathfrak{X} is ergodic, the pointwise ergodic theorem implies that, for any μ -integrable function $f : X \rightarrow \mathbb{R}$, equation (7.1) indeed holds almost everywhere. This formulation generalizes far beyond physical systems: it even finds applications in number theory.

With this knowledge of ergodic theory, let's return to the fairly stable systems that we're considering. As already mentioned in chapters 2 and 4, there is a striking analogy between physical systems motivating ergodic theory and the learning dynamics (one of the paradigm examples): It is not feasible to understand the (time evolution of) the position and momentum states of each and every gas molecules in a box. This is not due to a lack of theoretical understanding, but it is due to the sheer complexity of the number of particles involved. Analogously, the very heart of the interpretability (or black box) problem of neural networks is that it is not feasible to understand the (time evolution of) the weight of each and every connection in a neural network (during learning). This, too, is not due to lack of theoretical understanding (it is described, e.g., by the backpropagation algorithm), but it is due to the sheer number of uninterpreted weights involved. This issue is raised by problems 1 and 2 above.

However, there also is a disanalogy which concerns problem 3. The solution to problem 3 via the ergodic hypothesis is motivated by the system being fairly chaotic. However, the learning dynamics usually is the opposite: namely, fairly stable. This poses the question of finding an analogue of the ergodic hypothesis for learning dynamics and other fairly stable systems.

In the next subsection, we turn to the concept of randomness as a candidate for a solution. But first, for the remainder of this subsection, we see what ergodic theory—together with our results—has to say about fairly stable systems.

With our notion of a general dynamical system, we've covered a much broader class of dynamical systems: namely, structures $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ where (X, \mathcal{A}, μ) is a Borel probability space and $T : X \rightarrow X$ is only assumed to be measurable. So this includes learning dynamics and also the predator-prey model. Due to our representation theorem, \mathfrak{X} is isomorphic to the dynamical system generated

by $S(\hat{D}(\mathfrak{X}, \mathcal{B}))$ for some measurable basis \mathcal{B} . So we can assume that \mathfrak{X} is of the form $(X, \mathcal{B}(\tau), \mu, T)$ where τ is a compact and zero-dimensional Polish topology on X making T continuous. So (X, τ, T) is a topological system, whence Krylov–Bogolioubov theory applies: the sequence of measures $\mu_n := \frac{1}{n} \sum_{k=0}^{n-1} \mu T^{-k}$ on $\mathcal{B}(\tau)$ has a limit point $\bar{\mu}$ and this measure is preserved by T (see e.g. Walters 1982, thm. 6.9). The limit is with respect to the space $M(X)$ of probability measures on $\mathcal{B}(\tau)$ with the weak* topology. So from the original measure μ we can construct a preserved measure $\bar{\mu}$ making $(X, \mathcal{B}(\tau), \bar{\mu}, T)$ a measure-preserving transformation.

Moreover, by the Choquet representation theorem we can express $\bar{\mu}$ as a generalized convex combination of ergodic measures on (X, T) (Walters 1982, rem. 2, p. 153). This is expressed formally as follows: Write $M(X, T)$ (resp., $E(X, T)$) for all the preserved (resp., also ergodic) measures on (X, T) . Then $M(X, T)$ is a compact convex subset of $M(X)$. So there is a unique measure τ on $\mathcal{B}(M(X, T))$ such that $\tau(E(X, T)) = 1$ and, for all continuous $f : X \rightarrow \mathbb{R}$,

$$\int_X f(x) d\bar{\mu}(x) = \int_{E(X, T)} \left(\int_X f(x) d\epsilon(x) \right) d\tau(\epsilon), \quad (7.2)$$

where, for ϵ -almost all $x \in X$, $\int_X f(x) d\epsilon(x) = \lim_n \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x))$.¹¹ This indeed expresses $\bar{\mu}$ as a generalized convex combination of ergodic measures: The integral over a probability measure (here τ) generalizes the idea of summing values weighted by their probability—like a convex combination. In particular, if $E(X, T) = \{\epsilon_1, \dots, \epsilon_n\}$ is finite, then

$$\int_X f(x) d\bar{\mu}(x) = \tau(\epsilon_1) \left(\int_X f(x) d\epsilon_1(x) \right) + \dots + \tau(\epsilon_n) \left(\int_X f(x) d\epsilon_n(x) \right).$$

The fact that we can express a preserved measure $\bar{\mu}$ as a generalized convex combination of ergodic measures as in equation (7.2) is known as *ergodic decomposition*. (See Viana and Oliveira (2016, ch. 5) for discussion.) Its purpose is this: Assume we consider a dynamical system that has a preserved measure which, however, is not ergodic. Then we cannot allude to the ergodic hypothesis to come to know the desired time average of a measurement. However, by estimating the expected value of the measurement, we can still come to know the convex combination of the time averages—and thus still gain valuable information about the long-term behavior.¹²

So far, we've seen what ergodic theory has to say about *any* general dynamical system (in our sense). However, what about those that have a fairly stable

¹¹This is because each ϵ is an ergodic measure, so the ergodic theorem applies. Note that f is ϵ -integrable since it is bounded qua continuous function with compact domain (so its image is compact in \mathbb{R} and hence bounded) and ϵ is a probability measure.

¹²Incidentally, the proof of the (more general) ergodic decomposition theorem presented by Viana and Oliveira (2016, ch. 5) uses the Rokhlin disintegration theorem and a sequence of increasingly refining finite partitions. Thus, it seems worth investigating connections to our observation domain construction.




Figure 7.3: A converging system.

behavior and thus might have a finite symbolic approximation?¹³ Here is an idea: The fact that a system (X, T) is very stable—namely convergent—seems to be related to the assumption that the set $E(X, T)$ of preserved ergodic measures is finite. This is true for the paradigm example of such a system: the North-South map mentioned in chapter 2. (For a short discussion, see Walters (1982, sec. 6.6, ex. 6).) But this works much more generally: Consider the system (X, T) depicted in figure 7.3. It has two attracting fixed points x_0 and x_1 . It could, for example, describe a learning dynamics: If the neural network is initialized with weights coming from the left (resp., right) half of the state space, the learning will converge to the weights represented by x_0 (resp., x_1). And the measurement $f : X \rightarrow \mathbb{R}$ is provided by the loss function: $f(x)$ quantifies the error that the network still makes in weight setting x with respect to some test set.

Now, we can reason informally as follows. Consider the set $\Omega(T)$ of ‘non-wandering’ states: a state $x \in X$ is *non-wandering* if, for every neighborhood U of x , there is $n \geq 1$ such that $T^{-n}U \cap U \neq \emptyset$. So the only non-wandering states are x_0 and x_1 : For any other $x \in X$, choose a neighborhood U of x that is small enough such that T takes any point y of U outside U (i.e., $T(y) \notin U$). Since the dynamics is convergent any further iteration $T^n(y)$ will only get closer to x_0 or x_1 and thus further away from U , so the orbit of y will never return to U , whence $T^{-n}U \cap U = \emptyset$. Now, it is a theorem that any preserved probability measure assigns $\Omega(T)$ measure 1 (see e.g. Walters 1982, thm. 6.15). So the preserved probability measures are precisely those that assign probability p to x_0 and $1 - p$ to x_1 (for $p \in [0, 1]$). Since the ergodic measures are precisely the extreme points of $M(X, T)$, the only two preserved ergodic measures are ϵ_0 assigning probability 1 to x_0 and ϵ_1 assigning probability 1 to x_1 .

To summarize, combining our results with deep results from ergodic theory,

¹³Here we don’t mean that the system has a symbolic representation in the sense of having a finite generator (as implied by the Krieger Generator Theorem). We’ve already discussed this in chapter 4 (in the introduction at the end of the paragraph ‘the computational interpretation’). The issue is that such symbolic representations still have an uncountable state space. But here we’re looking for a ‘good enough’ symbolic approximation which has a finite state space.

we can already say quite a lot also about the behavior of fairly stable systems. However, what can we say for measures that aren't preserved? What is an analogue of the ergodic hypothesis there? To explore this question, we turn to randomness: since, as we'll see, it allows equating statistical behavior with actual behavior.

7.2.3 Randomness

As seen in the previous subsection, since we cannot uniquely determine the states of a system (up to infinite precision), we can study the long-term behavior generally only statistically through repeated measurement of the system. In systems satisfying the ergodic hypothesis—like a box of gas—this still (with probability 1) provides us with the time average of a measurement and thus informs us about the long-term behavior. But in fairly stable systems without a preserved measure—like learning dynamics—we're missing an analogue of the ergodic hypothesis.

In both cases, we face the following questions. Assume we know that *statistically* the system shows a certain long-term behavior: e.g., deduced from the ergodic hypothesis (in the box of gas) or being convergent (in the learning dynamics). How do we know for *certain* that it will do this in the given state? In other words, even if we know that with probability 1 (i.e., *almost surely*) the time average of a measurement will converge to a certain value, how do we know *surely* that in the current state the system will exhibit this behavior? In yet other words, even if we know statistically *that* the system exhibits a certain behavior, is there an explanation *why* we should expect with certainty this behavior in the current state?

As a solution to this problem, we suggest *randomness*. We first state the rough idea and then discuss it with the two examples (the box of gas and the learning dynamics). The idea is that usually there always is some 'random noise' in a system. This ensures that if we find the system in an initial state, this state is 'random' in the sense that it doesn't have any rare statistical properties. So we can expect the system to exhibit the statistically expected behavior when starting in this state. In box of gas example (and in ergodic systems more generally), this provides a way of knowing when the ergodic hypothesis holds for a state (rather than just knowing that it holds with probability 1). In the learning dynamics, this suggests, as we'll discuss below, an analogue of the ergodic hypothesis.

As an example, consider again the box of gas: a physical system with a huge amount of interacting particles. (Another classic example is the weather.) Then we can consider 'typical' states of the system as random: Even if the state where all particles are in the left half of the box and have identical momenta is possible, we wouldn't expect to ever observe it. Just like we would never expect to observe an infinite sequence of tails when repeatedly tossing a coin. It is notoriously hard to formalize this intuitive concept. But it is well researched:

Various formalizations are studied in the field of algorithmic randomness,

but arguably the most common one is Martin-Löf randomness. (As a standard reference for what follows we use Downey and Hirschfeldt (2010); for a history of the development of these notions, see Van Lambalgen (1987).) Such a formalization describes what it means for an infinite object to be random. Typically, these objects are infinite binary sequences. So they can be regarded as representing repeated coin tosses. But they can also encode many other infinite objects. More generally, the theory is also developed for more general spaces of infinite objects (Hoyrup and Rojas 2009). This includes the \mathbb{R}^n : the state space of the box of gas. The idea behind Martin-Löf randomness is that—as already indicated—an infinite object is random if it has no rare effective properties. Rare properties are those subsets U of the state space X that can be described as an intersection of an ‘effective’ sequence of open sets U_n with $\mu(U_n) \leq 2^{-n}$. So U is ‘rare’ because it is a null set; and ‘effective’ means, roughly, that there is an algorithm that takes as input n and enumerates more and more basic open sets whose union is U_n . Such a sequence (U_n) then is also called a Martin-Löf test and an infinite object x is Martin-Löf random if it fails every Martin-Löf test (U_n) , i.e., $x \notin \bigcap_n U_n$. For example, the sequence $x = 000\dots$ of only tossing tails is indeed not Martin-Löf random since it passes the Martin-Löf test (U_n) where U_n is the set of binary sequences starting with n -many 0’s. The notion of Martin-Löf randomness is relative to the measure on the space X and the computability structure (which makes both the basic open sets and the measure ‘computably accessible’). For our informal discussion, we can omit these details (they are found in the above references). It also is worth noting that there are other illuminating characterizations of Martin-Löf randomness in terms of informational incompressibility and unpredicatability.

Equipped with this concept of randomness, we return to the systems of ergodic theory, i.e., measure-preserving transformations of a state space—which, in particular includes the box of gas. In this setting, fairly recent research efforts developed a connection between algorithmic randomness and ergodic theory. See, e.g., Franklin and Towsner (2014), Gács, Hoyrup, and Rojas (2011), Galatolo, Hoyrup, and Rojas (2010), and V’yugin (1998). This includes various theorems relating a state being algorithmically random (e.g., Martin-Löf random) to the state being dynamically typical, i.e., satisfying the ergodic hypothesis (7.1).

Thus, in the case of measure-preserving dynamical systems, randomness indeed provides a solution to our problem (under the reasonable assumptions of the mentioned theorems). If the system starts in a random state, it will exhibit the statistically expected behavior. However, this leaves open the case when there is no natural preserved measure available.

So let’s consider learning dynamics as the paradigm example of a system with fairly stable behavior and hence likely no natural, non-trivializing preserved measure. So, recalling from the introduction of chapter 4, X is the space of parameters of the learning machine (e.g., a neural network) paired with the space of possible sequences of data. And the dynamics $T : X \rightarrow X$ describes the learning: how the current parameters are updated after seeing a data point according to the

learning algorithm of the machine. The measure μ on X describes the probability distribution of the choice of initial parameters of the system (initialization) and the likelihood of the (sequences of) data points. Finally, $f : X \rightarrow \mathbb{R}$ describes some loss function with which we can measure the performance of the learning machine. For our informal discussion, we again omit specifying the computability structure on X : but it should be such that μ , T and f are computable functions—after all, they describe a learning *algorithm*.

The field of statistical learning theory provides conditions ensuring the machine to converge almost surely. See, e.g., Bottou (1998), Nguyen et al. (2019), Saad and Solla (1996), and Vapnik (2000). Thus, this provides an assurance that the learning will work almost always, but it doesn't explain why it works for a given specific initial state. The suggested solution via randomness then claims:

Randomness ensures learning: Assuming the learning dynamics converges statistically, if the learning machine is initialized randomly and samples data points randomly, then it converges certainly.¹⁴

We may regard this as a first step toward the missing analogue of the ergodic hypothesis: From purely statistical knowledge about the system, we can deduce something about the long-term behavior starting in a given initial state. Thus, randomness provides stability: namely ensuring stable convergence with certainty when we only know it statistically. Moreover, if true, this thesis also corroborates theoretically the well-known practical rule of thumb that neural networks show better convergence results when initialized randomly (as opposed to, say, setting all parameters to 0).¹⁵ In the remainder of this subsection, we'll argue for this thesis.

First, we need to formally phrase the convergence assumption. We do this as follows. Assume we know that *statistically* the system minimizes the loss function reasonably fast: The expected average loss converges to zero fast enough such that the accumulated expected average loss is finite. Formally,

$$\sum_{n=1}^{\infty} \int \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) \, d\mu < \infty. \quad (7.3)$$

This is the case if, for example, the expected average loss converges *effectively*, i.e., $\int \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) \, d\mu \leq 2^{-n}$. Surely this is quite a strong convergence

¹⁴There is an interesting subtlety. In our continuous setting, the thesis expresses the intuition that randomness *ensures* learning because randomness excludes statistical outliers where convergence may fail. In a more symbolic setting, Osherson and Weinstein (2008) and Zaffora Blando (2019) explore connections between randomness and learnability. Here the intuition is the opposite: randomness *defies* learning because for a random binary sequence there precisely is no pattern or concept generating it which we could learn through observing the sequence. These intuitions don't clash: rather, we're considering two different learning settings.

¹⁵For more on randomness or noise in neural networks, see, e.g., An (1996), Neelakantan et al. (2015), and Srivastava et al. (2014) or the blog post by Eric Jang at <https://blog.evjang.com/2016/07/randomness-deep-learning.html> (last checked 3 July 2021).

assumption. It is more like the consequence of a theorem and not a condition. But the point here is not how we precisely get the statistical convergence. It rather is how we can use it to guarantee convergence for random starting states. So the assumption only serve as a proof of concept. For actual applications, it would be replaced by more detailed assumptions, for example, provided by statistical learning theory. What is important, though, is that this is a purely statistical assumption: it involves no ‘local’ knowledge about initial states but only about the ‘global’ statistical behavior of the system.

Now, to argue for the thesis, we show that, from the purely statistical convergence assumption (7.3), we can deduce that, if the learning dynamics starts in a random state, it converges: i.e., if $x \in X$ is Martin-Löf random, then $\lim_n \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) = 0$. The idea is to show that, given the statistical convergence assumption, diverging is a rare effective property.

Indeed, fix some $e \in \{0, 1, \dots\}$. For $n \geq 0$, define

$$U_n^e := \left\{ x \in X : \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) > 2^{-e} \right\}.$$

Again without going into details, this is a ‘computable’ set: given x , we compute $\frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x))$ and then check whether it is $> 2^{-e}$ —and we can do this uniformly in n . More precisely, given n , we can computably find the basic open sets constituting the preimage of $(2^{-e}, \infty)$ under the function $\frac{1}{n} \sum_{k=0}^{n-1} f \circ T^k$. (Note we don’t make use of the fact that we can choose U_n to be only semi-computable uniformly in n .)

Moreover, we have $\sum_n \mu(U_n^e) < \infty$. Indeed, write $g_n(x) := \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x))$. By Markov’s inequality,

$$\mu(U_n^e) \leq \mu\left(\left\{x \in X : g_n(x) \geq 2^{-e}\right\}\right) \leq 2^e \int g_n \, d\mu.$$

So $\sum_n \mu(U_n^e) \leq 2^e \sum_n \int g_n \, d\mu$ which is finite since $\sum_n \int g_n \, d\mu < \infty$ by assumption (7.3).

Now, assume $x \in X$ is Martin-Löf random and show $\lim_n \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) = 0$. Indeed, we have shown that, for any $e \geq 0$, the sequence $(U_n^e)_n$ is a Solovay test (which is a version of a Martin-Löf test), so x is in at most finitely many U_n^e , whence there is an N_e such that, for $n \geq N_e$, $x \notin U_n^e$. To show $\lim_n \frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) = 0$, let $\epsilon > 0$. Choose $e \geq 0$ big enough such that $2^{-e} < \epsilon$. Then we have, for any $n \geq N_e$, that $x \notin U_n^e$, so $\frac{1}{n} \sum_{k=0}^{n-1} f(T^k(x)) \leq 2^{-e} < \epsilon$, as needed.^{16,17}

¹⁶Also note the generality of this argument: Apart from the converge assumption, we only made computability assumptions but no measure-preservation assumptions on the system (X, μ, T) . Also, nothing was particular to the learning dynamics. And we could also assume the loss function to converge to some limit function \hat{f} other than the zero function 0 by considering $f' := f - \hat{f}$.

¹⁷Could these ideas on the role of (algorithmic) randomness in non-symbolic computation help

7.2.4 Stability

This section dealt with the question of when a dynamical system can be regarded as realizing—or being approximated by—symbolic computation. The guiding intuition is that this is possible if the system exhibits a fairly stable behavior. The previous subsections explored this idea. However, to begin doing this in detail, we start at the very foundations: namely the underlying concept of stability. According to this idea, if we consider a property of states—like showing converging or oscillating behavior—, then it holds *stably* at a state if it not only holds at that state but also at all sufficiently close ones. Analyzing this notion of stability is the topic of the next chapter.

to incorporate ‘noise’ into analog/continuous computation? Bournez and Campagnolo (2008) and Orponen (1997) describe this as an open problem.

Chapter 8

Stability: Fitch’s paradox and AI-safety

Abstract We investigate the notion of *stability*: a property holds stably at a ‘state’ if it holds at that state and at sufficiently similar ones. Here a ‘state’ is understood in the most general sense: e.g., a state of a dynamical system, a possible world, a model of a theory, an input to an AI, a parameter setting, etc.

We discuss a wide range of examples involving stability (many revolving around AI-safety) and see that they are all instances of this general notion of stability. We describe a formal logic to reason about stability across all these examples.

Then we formulate four desirable principles of stability and prove them inconsistent: both proof-theoretically (via a novel interpretation of Fitch’s lemma) and semantically (via Kripke semantics and topological semantics). Thus, we obtain some consequences for the examples ‘just by logic’: The inconsistency constraints the verifiable sentences (those that we can come to know) and the falsifiable sentences (scientific hypotheses), and it poses fundamental limitations on specifying AI-safety.

8.1 Introduction

When is a property *stable*? The idea is that an object has a property stably if not only the object but also all sufficiently similar ones have the property.

For example, *being heavy* is stable, because if an object—say, a chair—is heavy, then, even if it were slightly heavier or lighter, it would still be heavy. On the contrary, *weighing exactly 21.3 kg* is not stable, because even the tiniest variation in weight means losing that property. More formally, we may say: we consider the set of possible states that the given object might be in (e.g., the state of weighing 5 kg, the state of weighing 0.1 kg more than in the actual world, etc.). And we say that a state has property p if the object in that state has property p . Then the property p of being heavy is a stable property of these states, because if a state has it, also all sufficiently similar ones have it. But the property q of weighing exactly 21.3 kg is not stable.

The advantage of this somewhat cumbersome terminology is that it is much more generally applicable: Instead of just considering

- (1) the possible states of an object

we may also consider, for example,

- (2) the possible states of a dynamical system like the weather
- (3) the possible activation states of a neural network
- (4) the possible inputs to an artificial intelligence
- (5) the possible dynamics that the economy may exhibit
- (6) the possible models of a (scientific or mathematical) theory like the theory of general relativity.

In studying these phenomena, various properties of the states are interesting. For example, atomic properties like: ‘being heavy’ in (1), ‘it is going to rain’ in (2), ‘recognizing a dog’ in (3), ‘identifying a stop sign’ in (4), ‘being ergodic’ in (5), or ‘having a global time’ in (6). This may also include Boolean combinations of atomic properties, like being heavy and not red ($p \wedge \neg q$) in (1), and similarly for the other examples. And, most crucially, we obtain further complex properties by applying the stability claim to existing properties: if p is the property “it is raining in s ” that possible states s of the weather may or may not have, we can also consider the property ‘stably, p ’ (also written $\Box p$): i.e., ‘it is raining in s and, in all states sufficiently similar to s , it also is raining’. This again is a property that some states may or may not have: some may have p but not $\Box p$, and others may have $\Box p$ and hence also p , yet others may fail to have p . Let’s consider two examples.

In (1), a scientist would like to refute the hypothesis p about the objects, for example, that the object weighs ≥ 21.3 kg. Thus, the scientist is really interested in the property $\Box \neg p$: she wants to find an experiment which shows that, even when taking the measurement errors into account, the object doesn’t have property p , i.e., the object stably fails to have property p . The scientist might also fear that the property p cannot be decided: that it is neither false (otherwise it could be falsified qua good scientific hypothesis) nor stably true (otherwise it could be verified by some potential experiment). So the scientist wonders whether $\neg \neg p \wedge \neg \Box p$ is true.

In (4), consider an AI researcher who has built and trained a neural network to identify stop signs in the camera input of a car. To test the safety of this piece of technology, she wants to check that no ‘adversarial attacks’ are possible: We shouldn’t have a camera input s where the AI works correctly, but after a minuscule alteration—like adding a small sticker to a traffic sign—the AI suddenly recognizes a ‘Speed Limit 45’ sign (Eykholt et al. 2017). In other words, if we

write p for the property of identifying a ‘Speed Limit 45’ sign, the AI researcher would like that it is not possible that p fails but holds in very similar inputs. More formally, she is interested in the property $\neg(\neg p \wedge \diamond p)$ (or, equivalently, $p \vee \Box\neg p$) being true at all states.

To summarize (and generalize): In a given phenomenon involving the notion of stability, we have a state space; and complex properties of states can be formed from the given atomic ones using the operations *and* (\wedge), *not* (\neg), and *stably* (\Box). Moreover, some subset Q of the set P of all properties represents the properties that are relevant or sensible (those that the researchers are interested in). Mnemonically, ‘ P ’ as in ‘properties’ and ‘ Q ’ as in ‘questions’. Note that the properties in Q are *not* regarded as possible ‘laws’ governing the state space: they are *not* hypotheses that are meant to hold at every state. Rather, Q simply is a collection of properties that are useful to talk about the state space: just like ‘raining’ is a useful property to distinguish different states of the weather, but it is not intended as a law that it is raining at every state.

We may now consider various general principles about stability and the set of relevant questions. For now, we’ll only introduce these principles briefly and informally. In section 8.3, we’ll formulate them precisely and discuss them in detail.

The first principle simply ensures that stability is a non-trivial notion.

- (S1) *Non-triviality*: Stable truth doesn’t just collapse to simple truth, i.e., there is at least some property φ in Q and some state s such that s has φ but not $\Box\varphi$.

The next two principles state two conditions on the properties on Q that ensure that they are, in a loose sense, ‘accessible’ to us. The first is this:

- (S2) *Falsifiability*: The properties in Q pose good scientific hypotheses, i.e., if they are false at a state, they are stably false, so we can in principle perform some experiment or observation which will show that the property indeed doesn’t hold at the state.

Sometimes, instead of falsifiability as in scientific hypotheses, we may want our properties in Q to be verifiable. As we’ll see, this is true in the example of neural networks. Then, instead of (S2), we consider its dual version:

- ($\bar{S}2$) *Verifiability*: The properties in Q are verifiable, i.e., if they are true at a state, they are stably true, so we can in principle perform some experiment which will show that the property indeed holds at the state.

(In general, we’ll use the ‘bar’ notation to denote dual versions.)

The other principle says that satisfiable properties in Q are stably true at *some* state, so we can consider that state as a paradigm (or standard) case that we can imagine when we imagine a state where the property holds.

(S3) *Standard model*: If φ is in Q and satisfiable (i.e., true at some state), then there is some state where φ is stably true.

For example, the property of weighing ≤ 21.3 kg is stably true at the state where the object weighs 10 kg, even if it can happen that it is not stably true (if the object happens to weigh exactly 21.3 kg).

Again, if we consider—as in the case of neural networks—verifiability instead of falsifiability, then we also consider the dual of (S3), namely:

(\bar{S} 3) *Standard countermodel*: If φ is in Q and not valid (i.e., false at some state), then there is some state where φ is stably false.

The last principle is concerned with how we can form new questions from old ones. Initially, it may seem like the least plausible, but we'll find that, as in the above examples of the scientist and the neural network, it is satisfied in many applications.

(S4) *Moore closure*: If φ is a sensible question, then also $\varphi \wedge \neg \Box \varphi$ (i.e., 'ϕ but not stably so') is a sensible question.¹

(\bar{S} 4) *Dual Moore closure*: If φ is a sensible question, then also $\varphi \vee \Box \neg \varphi$ (i.e., 'ϕ is true or stably false') is a sensible question.

Again, note that these principles are not concerned with the *truth* of the properties involved: For example, (S4) does *not* say that if φ is true at a state, also $\varphi \wedge \neg \Box \varphi$ is. It merely says that if φ is a sensible property to investigate, also $\varphi \wedge \neg \Box \varphi$ is. Moreover, this is in no direct tension with falsifiability: if φ is falsifiable, then—under very mild assumptions—also $\varphi \wedge \neg \Box \varphi$ is falsifiable (we'll formally show this in lemma 8.3.10).

However, we'll show that if taken together, these principles are in (indirect) tension with each other: we'll prove the following impossibility result. Here, the dual version of (S1) is simply (S1) itself, i.e., (\bar{S} 1):=(S1).

8.1.1. THEOREM (Impossibility result). *The principles (S1)–(S4) together are inconsistent, i.e., for any state space and any choice of sensible properties Q , at least one of (S1)–(S4) fails. Moreover, also the dual versions (\bar{S} 1)–(\bar{S} 4) together are inconsistent.*

This impossibility is not intended as a paradox: not all principles are 'obviously true'. However, they are compelling in some applications—like the scientific hypothesis and neural network example—where the inconsistency will then yield interesting consequences that hence can be obtained 'just by logic'.

¹ The name is due to the fact that the sentence $\varphi \wedge \neg \Box \varphi$ famously is known as the Moore sentence in (formal) epistemology.

The remainder of the chapter is structured as follows. In section 8.2, we collect a wide range of examples involving stability. In section 8.3, we describe a formal logic to reason about stability across all these examples. In section 8.4, we show proof-theoretically—via a novel interpretation of Fitch’s lemma—that they jointly are inconsistent. In section 8.5, we illustrate the inconsistency semantically using Kripke semantics and topological semantics. In section 8.6, we apply the impossibility result to obtain an extension of Fitch’s paradox and a fundamental limitation for AI-safety. In section 8.7, we conclude. An appendix at the end of the chapter collects the more technical proofs. In short, we perform one cycle: from examples, we extract a formalization, for which we obtain formal results, which we then apply back to the examples.

A last word before we start: This chapter attempts the difficult—but, arguably, important—task of bridging research in philosophy and artificial intelligence (here via logic). Recent literature on this broad topic is surprisingly scarce, but see, e.g., Buckner (2019) for a ‘call to action’ to remedy this and Evans (2020) for a concrete combination of philosophical thought (Kant) and AI-implementation. It is our hope that the results here are useful in both directions: for AI researchers to get some clarity regarding the *foundations* of central notions like AI-safety; and for philosophers to see novel *applications* of known concepts like Fitch’s paradox or safety conditions of knowledge. Such a cross-disciplinary endeavor is at risk of being too abstract (in its philosophical generality) to be useful for AI and too formal (in its attempt at a precise and AI-amenable formulation) to be philosophically refined enough (for a similar point, see Evans 2020, pp. 18–19). We hope that the fact that these issues neatly come together under the concept of stability—generalized from many concrete examples—goes some way toward alleviating these worries. The aim is to provide a fruitful starting point for discussion rather than decisive answers claiming to end the debate.

8.2 Examples of stability

We collect a wide range of examples involving stability. This should serve as a collection of brief summaries of reasons for demanding stability. It is neither intended to be an exhaustive list nor an in-depth philosophical discussion. Moreover, there are many connections between these ideas, and we attempt to spell out some along the way.

8.2.1 Verifiability and falsifiability (observation)

Maybe the most common motivation for demanding stability is epistemic: stability is needed to make sure that, if a property holds, we also can come to know (or verify) that it holds.

To illustrate this, consider again the example of a chair from the introduction.

The reason why we can come to know (or verify) a property p like 'weighing $> 3.2\text{ kg}$ ' is this: If true—say, because the actual weight-state of the chair is 3.4 kg —, then there is an appropriate measurement that we can perform: e.g., weighing with a scale with a margin of error $\pm 0.1\text{ kg}$. This measurement is appropriate in that it shows that the actual weight-state of the chair must be in an area of the state space—say, the interval $(3.3, 3.5)$ —that is completely contained in the area where p is true. Similarly, we cannot come to know the property q of 'weighing exactly 3.2 kg ' because this is true only in an area consisting of a single state while an area determined by a measurement will consist of an interval of states.

This example is much more general: consider the state space of a dynamical system like the weather. To make predictions about it, we must know what its current state is (before we can apply our knowledge about its dynamics). Again, we can only do so through measurements: observing, e.g., temperature, humidity, and air pressure at different locations. This will delimit the area of the state space in which the system currently is. Only then can we calculate (or simulate) to which areas of the state space the system will evolve from there. As a result, we can, in general, only predict those properties that are stable: if a state has a property, we need to verify it by finding, through observation (and calculation or simulation), an area around that state where the property holds.

Something dual can be said for falsifiability instead of verifiability through measurements: To falsify (or refute) a property means to verify its negation. For example, if the property q of 'weighing exactly 3.2 kg ' is false—say, because the chair weighs 3.4 kg —, we can falsify q by using the scales to determine that the actual weight-state of the chair must be in the area $(3.3, 3.5)$ of the state space which is completely contained in the $\neg q$ area: this verifies $\neg q$ and hence falsifies q . Good scientific hypothesis are usually required to be falsifiable: if they are false, there is an experiment or measurement that shows that they are false—so the negation is stable.²

To summarize, properties that we can come to know whenever they obtain (verifiability) or that are good scientific hypotheses (falsifiability) naturally are properties that we're interested in. However, they need to have a certain stability: to verify or falsify those properties, we need measurements, and these always have some margin of error which ultimately requires the properties to have stability. In

² These ideas figure prominently in the 'topology via logic' approach (e.g. Abramsky 1991; Smyth 1983; Vickers 1989). (Vickers (1989) uses 'affirmative' and 'refutative' for the above 'verifiable' and 'falsifiable', respectively.) Here the axioms of topology are (re-) interpreted as that for a logic of finite observations (or semidecidable properties). As a result, properties that can be confirmed by possible measurements act like open sets: once a state has such a property, also all sufficiently similar states have that property. Thus, in the above terminology, these properties are precisely the stable ones. Dually, properties that can be refuted by measurements act like closed sets: once a state doesn't have such a property, also all sufficiently similar states don't have it.

a slogan: verifiability and falsifiability require stability.

This idea is also closely related to discussions of safety in epistemology and artificial intelligence—as we’ll discuss now.

8.2.2 Safety (epistemology)

In epistemology, there is an extended discussion of so-called safety conditions for knowledge. (See, e.g., Sosa (1999) or Williamson (2000) and, for an overview, Ichikawa and Steup (2018) or Rabinowitz (2020).³) The main idea is that knowledge requires safety in the sense that, if we know that p is the case, we still wouldn’t be wrong in similar cases. A simple example to illustrate the idea is Bertrand Russell’s famous stopped clock (Russell 1948, pp. 170–171): We look at a clock with the intention of coming to know what time it is. Incidentally, at this very moment the clock shows the right time, but, unbeknownst to us, it actually has stopped exactly 24 hours ago. Intuitively, we wouldn’t consider the true belief about the current time that we obtained in this way to be knowledge. And, indeed, in the very similar case where we had looked at the clock just one minute earlier, we would have been wrong, so the safety condition is violated.⁴

There is much discussion on how to formulate such conditions precisely (Rabinowitz 2020; Sosa 1999; Williamson 2000). But it usually involves something like the following:

(SC) If A knows p based on some method M in situation s , then, in all situations s' similar to s , if A believes p in s' based on M , then p is true at s' .

For reasons of space, we won’t enter the discussion of this principle and its formulation. What we’re concerned with here is to indicate how it may provide another argument for the stability of those properties that we can come to know whenever they obtain.

The argument sketch proceeds in two steps. First, we observe that (SC) seems to imply

(*) If A knows p in situation s , then p is true in all situations s' similar to s .

and in fact even the stronger

(**) If A knows p in situation s , then, in all situations s' similar to s , A correctly believes p in s' .

Because: If A knows p in s , then A has obtained the belief that p based on some method M and this belief constitutes knowledge. Hence, in a sufficiently similar

³For a discussion of other stability related notions in epistemology, see Rott (2004).

⁴Closely related to the safety condition is the idea that knowledge shouldn’t involve epistemic luck: in the example, we were just lucky that our belief turned out to be correct (Pritchard 2005).

situation s' , A arguably would still tend to apply method M and would thus come to believe p in s' (otherwise s' wouldn't be sufficiently similar to s). But then, by (SC), p is true in s' , and A correctly believes p in s' .

Second, once we have (*), we seem to get the desired conclusion: Let p be a property that we can come to know whenever it obtains. We want to show that p is stable. Indeed, if p holds at a given state s , then there is something we can do so that we know that p at s . Once we do, (*) implies that p needs to be true also in all similar enough states, whence p is a stable property.⁵

8.2.3 Safety (artificial intelligence)

Especially recently, safety became a prominent concern in the field of artificial intelligence: If an artificial intelligence (AI) judges that a given input s (say, a camera picture of a traffic scene) to be of category c (say, 'a stop sign is present'), then we want that this is a *safe* judgment—which is necessary if we want to trust the AI. In particular, it shouldn't be possible to systematically trick the AI (e.g., by simply placing a sticker on the stop sign) to judge the scene completely differently (say, that a 'Speed Limit 45' sign is present). Such tricks to the AI are called *adversarial attacks*. They recently became infamous since they pose serious problems to neural networks.⁶ In other words, we want that, on all sufficiently similar inputs, the AI still yields the same judgment. Thus, this kind of AI-safety can be formulated as a requirement for 'state space' stability:

(AS) The AI is safe only if, when considering the set of all possible inputs to the AI as state space, the property 'the AI judges input s to be of category c ' is stable (for all categories c of the AI).

This, or some form of it, arguably is a necessary feature of AI-safety, but—as we'll argue below—there also is more to it, so we don't take it as a sufficient feature (hence merely 'only if' and not 'if and only if').⁷

How are the notions of this and the preceding section 8.2.2 related? There seems to be at least a rough analogy between the notion of safety in the sense of

⁵So we've only used the weaker principle (*). The reason why we've also stated (**) is that, first, we'll use it below and, second, the straightforward argument for (*) seems to go via (**).

⁶The literature on adversarial attacks is growing rapidly: some examples are Brown et al. (2017), Cisse et al. (2017), Goodfellow, Shlens, and Szegedy (2014), Goyal et al. (2020), Huang et al. (2017), and Szegedy et al. (2013). The example with the stop sign is from Eykholt et al. (2017). For a philosophical discussion, see Buckner (2019) and Buckner (2020). For a more general overview of AI-safety, see Amodei et al. (2016).

⁷To provide some more background: A crucial part of the engineering of neural networks is to achieve the kind of stability described in (AS): The network shouldn't overfit training data, i.e., fit to the data so well that it represents also their idiosyncrasies at the cost of the general pattern behind it. Various methods (e.g., dropout) guide the training of the network toward learning robust general patterns rendering the network more stable and less overfitted. For an overview, see Buckner (2019, p. 8).

‘safety conditions for knowledge’ and in the sense of ‘AI-safety’. If an AI judges, given some camera input, that there is a stop sign, then we may—in analogy to the epistemological terminology—say that the AI *believes* that there is a stop sign (in the real world from which the camera input originates). (This surely is quite an anthropomorphism, though epistemological principles and terminology usually are intended to be applicable to all epistemic agents: human and artificial.) For the AI to be safe, we would like that, if the AI makes such a judgment, it doesn’t do so carelessly but rather has a high certainty in the judgment, has some good reasons, or has some other form of justification providing sufficient reliability. Thus, in analogy to the epistemological terminology, we might say that the AI *knows* that there is a stop sign.

Thus, the above stability requirement for AI-safety corresponds, via this analogy, to safety conditions for knowledge: Assume that on input s the AI judges c . Then, if the AI is safe, it safely judges c , i.e., according to the analogy, knows c . Hence, by version (**) of the safety condition above, the AI still correctly believes c in similar situations. Thus, according to the analogy, the AI judges c in all similar situations—just as the stability requirement of AI-safety demands. This provides a potentially fruitful application of the extensive philosophical discussions of safety conditions for knowledge to AI-safety. (Or, vice versa, this allows to evaluate safety conditions by AI-safety examples.)

Another application of the analogy may be to better understand the elusive notion of AI-safety (or safe judgment of an AI). Crudely put, the analogy says that safe judgment is to judgment what knowledge is to belief. Thus, we should expect that just as it is notoriously hard (if at all possible) to specify what else is required to turn (justified) true belief into knowledge, it probably is equally hard to specify what else is required to turn a judgment of an AI into a safe judgment. In other words, the analogy casts doubt on the prospect of a fully satisfying conceptual analysis of AI-safety into crisp and operationalizable sufficient and necessary conditions—hence the hesitation to phrase (AS) as ‘if and only if’.⁸ (Besides, there are many more aspects to AI-safety like the avoidance of side effects (Amodei et al. 2016).)

Surely this analogy needs to be developed more carefully to convincingly translate insights between AI-safety and epistemology. To that end, we discuss

⁸The infamous Gettier examples obstruct an analysis of knowledge (see, e.g., Ichikawa and Steup 2018) and adversarial attacks might be seen as the analogical counterparts: Assume an AI built with a deep neural network correctly classifies a camera input as containing a stop sign but could be adversarially attacked. Then the AI has, in the terminology of the analogy, the true belief that there is a stop sign, and it also has some justification for it: after a long learning process, it has built up its current weights which yielded the given classification. So the AI has true justified belief but no knowledge (i.e., safe judgment)—just like a Gettier case. (This kind of justification could be regarded as ‘internalistic’, and an attempt to explain away Gettier examples is to demand that justification shouldn’t be internalistic but externalistic (Ichikawa and Steup 2018); so it would be interesting to consider what the analogical counterpart for AI-safety is.)

two ways of making precise the stability requirement (AS) for AI-safety—the analogical counterpart to the safety condition for knowledge. More specifically, we discuss two ranges under which the behavior of the AI should be stable: (i) under similarity of the input, and (ii) under updating with likely propositions. We come back to (i) when discussing the consequences of our inconsistency result in section 8.6.2. Now, we turn to the stability theory of belief which investigates (ii).

8.2.4 Stability of belief (probabilistic reasoning)

We'll recap the stability theory of belief developed by Leitgeb (2017) and see how it provides another general kind of stability in our sense. We also briefly sketch an application to AI-safety.

The stability theory of belief concerns the question of how qualitative belief ('all-or-nothing' belief) and quantitative belief (degrees of belief) relate. For example, when should we take a probabilistic statement expressing degrees of belief (like 'The agent is 70% certain that this is poisonous') to warrant a definite statement expressing all-or-nothing beliefs (like 'The agent believes that this is poisonous')?

Various answers have been given. The most straightforward one is the *Lockean thesis*: a perfectly rational agent (all-or-nothing) believes a proposition iff her degree of belief in the proposition is higher than a certain threshold (Foley 1993, p. 140). This, however, faces serious problems: most notably, the lottery paradox (for a short overview see Huber 2016, sec. 2.6). These problems (are taken to) show, for example, that the resulting notion of all-or-nothing belief isn't necessarily closed under conjunction and thus violates a logical law that we intuitively would expect of rational beliefs.

The stability theory of belief (Leitgeb 2017) avoids these problems by opting for the *Humean thesis*: a perfectly rational agent believes a proposition p iff she has a stably high degree of belief in it (Leitgeb 2015).⁹ Roughly, this means that for all propositions q of some fixed set \mathcal{Y} of propositions that the agent may suppose or come to learn, if the agent were to update her current degrees of beliefs with q , her degree of belief in p would still be above some fixed threshold ρ (e.g., $\rho = 0.9$).

These ideas may be formulated in the terminology of stability in state spaces as follows. The state space is the set of all possible degrees of beliefs that the agent may have. (Or, to be more precise, the set of all probability measures on an appropriate underlying measurable space.) For $q \in \mathcal{Y}$, we say that the degree of beliefs μ is *q-similar* to the degree of beliefs ν if ν is the result of updating μ with q . (More precisely, ν is the conditional probability measure obtained by conditioning μ on q .) We say μ is *similar* to ν if μ is *q-similar* to ν for some

⁹See Leitgeb (2017) and Leitgeb (2018) for references to related work.

$q \in \mathcal{Y}$.¹⁰ Then it is easy to see that the following are equivalent for all propositions p and degrees of belief μ :

- (a) The agent (all-or-nothing) believes p at μ in the sense of the Humean thesis: i.e., for all $q \in \mathcal{Y}$, if the agent updates μ with q , her degree of belief in p is $> \rho$.
- (b) The property ‘the agent’s degree of belief in p according to μ is $> \rho$ ’ is stably true at μ : i.e., it holds at all states similar to μ .

Thus, the property ‘the agent’s degree of belief in p according to μ is $> \rho$ ’ is stable iff (all-or-nothing) believing p according to the Lockean thesis implies (all-or-nothing) believing p according to the Humean thesis.¹¹

Different choices for \mathcal{Y} are possible (and thus provide different versions of the Humean thesis). A salient one is to take \mathcal{Y} as the set of propositions that the agent believes to a degree $\geq \epsilon$, for some threshold $\epsilon > 0$ which needn’t depend on ρ (e.g., $\epsilon = 0.2$).¹² Choosing \mathcal{Y} to be the set of propositions that the agent deems possible (i.e., doesn’t all-or-nothing believe the negation) turns out to be, in a precise sense, the unifying choice (Leitgeb 2017, thm. 1, p. 85).

We can apply this to AI-safety. In section 8.2.3, we saw the stability requirement for AI-safety: the behavior of the AI is stable under sufficiently small variations of the input. The stability theory of belief—which is encapsulated in the Humean thesis—suggests another, conceptually equally important stability requirement: For the judgment of an AI to be safe, we not only want that the AI has high certainty in its current judgment, but also keeps having high certainty after updating with propositions from the chosen set of propositions.

8.2.5 Significance (mathematical modeling)

Especially in the context of general relativity, it has been argued that mathematical models of a real phenomenon need some stability—which we’ll describe now.

The idea of mathematical modeling is to provide mathematical models for a given real phenomenon with the aim of transferring result from the mathematical investigation of the models into insights about the phenomenon. However, for this to work, one needs to know which mathematical properties of the models are *significant*—i.e., correspond to properties of the phenomenon—and which mathematical properties are merely idiosyncrasies of the mathematical models?

¹⁰A more general similarity structure on S may be obtained by demanding the Kullback–Leibler divergence between ν and $\mu(\cdot|q)$ to be small (i.e., below a fixed threshold).

¹¹Note that, as arguably should be expected, the notion of stability figuring in the Humean thesis need not be transitive: If p is (all-or-nothing) believed at μ according to the Humean thesis, then p is (all-or-nothing) believed at all similar ν according to the Lockean thesis—but in general not again according to the Humean thesis. (But it also can be transitive: see example 8.5.3.)

¹²See Leitgeb (2017, p. 82) for this and more salient choices for \mathcal{Y} .

The idea is that the significant properties are stable (with respect to an appropriate notion of similarity). Here are three examples, which we'll revisit later on.

First, in general relativity, the mathematical models are spacetimes (i.e., four-dimensional, smooth, connected, Lorentzian manifolds). Hawking (1971, p. 395) argued that if a property of spacetimes is physically significant, it has to be stable in some appropriate topology: whenever a spacetime has that property, all sufficiently similar spacetimes have it, too (where 'sufficiently similar' is formally specified by the topology). The argument is that our observations of the physical phenomenon will always be imprecise due to measurement errors and the uncertainty principle, so they will never determine a unique mathematical model. Hence, if the obtaining of a mathematical property is to imply a property of the physical phenomenon, it has to be stable across the uncertainty of the mapping between models and reality. (See Fletcher (2016) for an overview and discussion of this idea of physical significance.)

This argument also seems to apply in many other cases. To illustrate, here are two more examples.

Second, consider (finite-state and time-homogeneous) Markov chains as mathematical models for any of the real phenomena that have successfully been modeled by them (e.g., in physics, biology, chemistry, or economics).¹³ Such a Markov chain over n states is given by a stochastic $n \times n$ matrix A where $A(i, j)$ specifies the probability of the system moving from state i to state j . Again, our observations of the real phenomenon will always be imprecise due to measurement errors, so they will never determine a unique Markov chain (i.e., stochastic matrix). So if we want to infer something about the real phenomenon from a mathematical property of the Markov chains, the property has to be stable across the uncertainty of the mapping between models and reality. So, for a mathematical property p to be significant, it has to be stable: if a Markov chain has it, also all sufficiently similar Markov chains have it. (The natural topology to spell this out formally is the topology on the stochastic matrices given by some matrix norm.)

Third, consider the computer implementation of a mathematical feed-forward neural network. Such a neural network is mathematically specified by a finite list (W_1, \dots, W_n) of real-valued matrices (the network has $n + 1$ layers and matrix W_i describes the weight on the connections between neurons from layer $i - 1$ and i). Any computer implementation of such a network will represent each weight only to a finite approximation (e.g., by floating-point reals). Again, many mathematical models (the network considered as a list of real-valued matrices) will correspond to the 'real' phenomenon (an implementation of the network as a computer program). So, for a mathematical property of the mathematical networks to be significant, it has to be stable: if a mathematical network has it, also all sufficiently similar mathematical networks have it. (Here sufficiently similar means that the difference

¹³For a reference on (general) Markov chains with an eye toward stability, see Meyn and Tweedie (2009).

is beyond what can be expressed in the floating-point reals.)

8.2.6 Further examples

We'll briefly mention three more examples.

First, a central idea of the theory of conceptual spaces (e.g. Gärdenfors 2004) is that (possible) objects can be regarded as points in a space whose dimensions correspond to qualities like weight, color, temperature, size, etc. Concepts then are regions in that space: for example, the concept red is the set of all red objects. Similarly, we can think of the possible states of a given object in this way: each state is given by a vector representing, for example, the weight, color, temperature, or size of the object in that state. Then properties are concepts, i.e., regions in the state space. Famously, convexity seems like a natural condition for concepts: if two objects instantiate a concept, say being red, then any object in between the two also does so (Gärdenfors 2004). We may also think that openness (and hence stability) is a natural condition: if a (possible state of the) object is heavy, red, warm, or large, then also every similarly enough (state of the) object has this property.¹⁴ The motivation for this idea comes from the aforementioned knowability and safety considerations: For an everyday concept or property to be useful, we should (a) be able to come to know that it obtains if it obtains and (b) it should be safe to use, i.e., we cannot be systematically fooled about it. This may be regarded as a transcendental argument for the stability of concepts: our cognition is such that it only learns (and keeps) useful—and hence stable—properties.¹⁵

Second, there is a wide-spread line of thought that objects are characterized by the set of transformations under which they are invariant, i.e., *stable*. Here are some examples from different disciplines. In cognitive science, invariance is theorized to play a central role in cognition. For example, it has been argued that we recognize objects as the things that remain invariant under certain transformation like a change of perspective (Gibson 1979/2015).¹⁶ And it has been argued that we pick up those patterns or concepts that have few variants (Garner 1974/2014).¹⁷ In physics, invariances are known as symmetries and play an important role in conservation laws (Noether's theorem). The idea also is ubiquitous in mathematics:

¹⁴To the best of our knowledge, this idea doesn't seem to be prominently discussed in the theory of conceptual spaces. However, the formalization used by Gärdenfors (2004) is based on the *region connection calculus* (Cohn et al. 1997) where not the points of a space are taken as primitive objects but the regions describing the geometry of the space. (Cf. pointless topology, as it is also used in the 'topology via logic' approach mentioned in footnote 2.) And a standard interpretation of regions in \mathbb{R}^n is as non-empty regular open sets (see e.g. Cohn et al. 1997, 102 f.).

¹⁵ One might also explore whether the usefulness of stable properties has to do with their reduced complexity: see footnote 9 in chapter 7.

¹⁶As far as AI counterparts are concerned, this is somewhat reminiscent of the Generative Query Networks of Eslami et al. (2018).

¹⁷Also see Vigo (2009).

an early example is the Erlanger program of characterizing different geometries by the group of transformations preserving their basic notions (length, angle, parallel, etc.).¹⁸ In logic, this also inspired Tarski to identify the logical constants as those concepts that remain unchanged under all permutations of the domain of discourse (MacFarlane 2017; Tarski 1986).¹⁹ For the role of invariance in philosophy, see, e.g., Nozick (2001).

Third, counterfactuals are sentences of the form ‘If p were the case, then q would be the case’. They are important in philosophy: both as an object of study (aiming to analyze counterfactuals) and as a tool (using counterfactuals to phrase philosophical theories).²⁰ Famously, Stalnaker (1968), Lewis (1973), and others provided a semantic analysis of counterfactuals in terms of similarity: the idea is that a counterfactual ‘if it were that p , then q ’ is true at a possible world s iff in the most similar possible worlds where p is true, also q is true.²¹ Within this analysis, the *limit assumption* (Lewis 1973, pp. 19–20) says that: if φ is false at a given world s , then there is no infinite sequence of φ -worlds that become more and more similar to s . This assumption is usually made since it simplifies the semantics (Lewis 1973; Starr 2019).²² In terms of state space stability, we can consider the possible worlds as states and their similarity relation as providing a notion of similarity between states. Then, intriguingly, the limit assumption says that whenever φ is false at a state s , then it is stably false: all sufficiently similar worlds s' also make false φ (i.e., there aren't arbitrarily close φ -worlds).

8.3 Four principles of stability

We describe a formal logic to reason about stability across all the preceding examples (section 8.3.1). Then we formulate four general principles that we'd like the general notion of stability to have (section 8.3.2). We also make formal the duality between falsification and verification (section 8.3.3). And we describe natural ways of building sets of ‘sensible questions’ (section 8.3.4).

8.3.1 A logic to reason about stability

To describe a logic to reason about stability, we need to: (i) specify a language in which to formulate claims involving stability, (ii) specify a derivability relation which describes which sentences follow logically, and (iii) specify a semantics that assigns meaning to the sentences of the language. Finally, we also need to

¹⁸The founding paper of category theory, regards it as a continuation of the Erlanger program (Eilenberg and MacLane 1945).

¹⁹In set theory, also see the permutation models (Jech 1973, ch. 4).

²⁰For an overview, see Starr (2019).

²¹See Starr (2019, esp. sec. 2.3) for the history of this idea and a summary.

²²Here we won't go into a discussion of the plausibility of the limit assumption. See Lewis (1973, p. 20), Stalnaker (1984, pp. 141–142), and Starr (2019, supplement., sec. B.4).

introduce a formalization of the ‘sensible questions’. We now do these things in turn.

Language What is common to the examples in section 8.2 is that we could formulate the properties of the states using atomic properties (like ‘in state s , the chair weighs between 5.2 kg and 5.6 kg ’) and combining these with Boolean connectives ($\neg, \wedge, \vee, \rightarrow$) and the stability operator (\Box). Intuitively, the stability operator says that the property in its scope is not only true in the current state but also in all sufficiently similar states. The importance of adding the stability operator is that we not only want to reason with the atomic (or Boolean) sentences, but also with claims about their stability (and nested claims of stability). For example, the (validity of the) sentence $p \rightarrow \Box p$ intuitively says that property p is stable: whenever it is true at a state, it is also true at all sufficiently similar states.

Moreover, since \Box expresses ‘local truth’ (truth in a small neighborhood around the current state) it is natural to add a *global truth* (truth at all states) operator \Box . Intuitively, the global truth operator says that the property in its scope is true at every state. So we can phrase the sentence expressing that property p is stable also as $\Box(p \rightarrow \Box p)$.

Thus, our language is built from a set of atomic sentences \mathcal{L}_0 using the operators $\neg, \wedge, \vee, \rightarrow, \Box, \Box$. Formally:

8.3.1. DEFINITION. We fix a nonempty set \mathcal{L}_0 and call it the set of *atomic sentences*, and we use variables p, q, \dots for its elements. The *language* \mathcal{L} is obtained by freely closing \mathcal{L}_0 under the operators $\neg, \wedge, \vee, \Box, \Box$.²³ We call the elements of \mathcal{L} *sentences*. As common, we define (for some $p \in \mathcal{L}_0$)

$$\begin{array}{ll} \varphi \vee \psi := \neg(\neg\varphi \wedge \neg\psi) & \perp := p \wedge \neg p \\ \varphi \rightarrow \psi := \neg\varphi \vee \psi & \Diamond\varphi := \neg\Box\neg\varphi \\ \varphi \leftrightarrow \psi := (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi) & \Box\Diamond\varphi := \neg\Box\neg\varphi \end{array}$$

We neither need further syntactic structure on the atomic properties (as, e.g., first-order quantification) nor cardinality assumptions on the atomic properties, so we only assume that \mathcal{L}_0 is a nonempty set. The language \mathcal{L} is a well-known bimodal language (Shehtman 1999).

Derivability If we want to reason about stability, we need a logic governing the \mathcal{L} -sentences: we need a derivability relation \vdash on \mathcal{L} .

Since we’re interested in formulating the principles about stability at the most general level, we won’t explicitly define a single derivability relation. Rather, we’ll

²³I.e., \mathcal{L} is the smallest set such that, whenever $\varphi, \psi \in \mathcal{L}$, also $\neg\varphi, \varphi \wedge \psi, \Box\varphi, \Box\varphi \in \mathcal{L}$.

axiomatically list the minimal assumptions we make about it. Nonetheless, we'll define concrete (and semantic) derivability relations below as examples. Moreover, we'll only consider whether a sentence is derivable (without premises), and not whether it is derivable from a set of premises. Hence, we'll define the derivability relation \vdash simply as a subset of \mathcal{L} , rather than as a binary relation between subsets of \mathcal{L} and elements from \mathcal{L} .

The formal axiomatic definition of a derivability relation is as follows. Below, we'll briefly motivate the axioms and provide examples. For background on modal logic, see Blackburn, Rijke, and Venema (2001).

8.3.2. DEFINITION. A *derivability relation* on \mathcal{L} is a subset \vdash of \mathcal{L} with the following properties (we write $\vdash \varphi$ for $\varphi \in \vdash$):

- (D1) (a) If φ is a classical tautology,²⁴ then $\vdash \varphi$. (*classical logic*)
 (b) If $\vdash \varphi$ and $\vdash \varphi \rightarrow \psi$, then $\vdash \psi$. (*modus ponens*)
 (c) If $\vdash \varphi$ and ψ is the result of uniformly replacing atomic sentences in φ by arbitrary sentences, then $\vdash \psi$. (*substitution*)
- (D2) (a) $\vdash \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$. (*K-axiom for \Box*)
 (b) If $\vdash \varphi$, then $\vdash \Box\varphi$. (*necessitation for \Box*)
- (D3) (a) $\vdash \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$. (*K-axiom for \Box*)
 (b) If $\vdash \varphi$, then $\vdash \Box\varphi$. (*necessitation for \Box*)
- (D4) $\vdash \Box\varphi \rightarrow \varphi$. (*reflexivity*)
- (D5) $\not\vdash \perp$. (*consistency*)

We call \vdash *transitive* (resp. *symmetric*) if, additionally, $\vdash \Diamond\Diamond\varphi \rightarrow \Diamond\varphi$ (resp. $\vdash \varphi \rightarrow \Box\Diamond\varphi$).

In words, axiom (D1) says that we follow the scientific standard of basing our reasoning on classical logic. Axioms (D2) and (D3) say that \Box and \Box satisfy the basic axioms for a (normal) modal operator. Axiom (D4) says that stable truth implies (simple) truth. Axiom (D5) says that the logic (or derivability relation) is consistent. For much of our purposes, we could replace axiom (D3) by the weaker axiom: if $\vdash \neg\varphi$, then $\vdash \neg\Diamond\varphi$ (which is equivalent to necessitation for \Box). But adding the *K*-axiom for \Box has some 'structural' benefits while coming at no cost for us: With it, the usual tools of modal logic are applicable. And \vdash is closed under equivalence: if $\vdash \varphi \leftrightarrow \psi$, then $\vdash \chi \leftrightarrow \chi[\varphi/\psi]$ where $\chi[\varphi/\psi]$ is the result of

²⁴I.e., φ contains only atomic sentences and Boolean connectives (\neg, \wedge) and is true under any classical valuation.

uniformly replacing the subformula φ of χ by ψ .²⁵ Moreover, if \boxplus has the intended interpretation of a global truth operator, it will satisfy much stronger axioms than just the K -axiom.

Examples for \vdash can be obtained from consistent and reflexive *normal modal logics* Λ (in the bimodal language) by setting $\vdash \varphi$ iff $\varphi \in \Lambda$. Examples of such logics are: $S4 \star S5$, $S4U$, $S4UC$ (Shehtman 1999). However, more intuition can arguably be gained from semantic examples for \vdash , which we'll now cover.

Semantics There are two well-known semantics for modal languages: the Kripke (or relational) semantics and the topological semantics. (For a handbook chapter covering both, see Van Benthem and Bezhanishvili (2007).) For convenience, we recap these here.

8.3.3. DEFINITION. A *Kripke model* is a triple $M = (S, R, V)$ where S is a non-empty set (whose elements are called *states*), $R \subseteq S \times S$ is a binary relation (called *accessibility relation*), and $V : S \rightarrow \mathcal{P}(\mathcal{L}_0)$ is a function (called *valuation*) assigning each state a set of atomic sentences (that, intuitively, are taken to be true at that state). We call M *reflexive* (resp., *transitive*, *symmetric*) if R is a reflexive (resp., transitive, symmetric) relation.²⁶

We recursively define when an \mathcal{L} -sentence φ is *true* at a state s of M , which is denoted $M, s \models \varphi$:

- $M, s \models p$ iff $p \in V(s)$, for $p \in \mathcal{L}_0$.
- $M, s \models \neg\varphi$ iff $M, s \not\models \varphi$.²⁷
- $M, s \models \varphi \wedge \psi$ iff $M, s \models \varphi$ and $M, s \models \psi$.
- $M, s \models \Box\varphi$ iff, for all $s' \in S$, if sRs' , then $M, s' \models \varphi$.
- $M, s \models \boxplus\varphi$ iff, for all $s' \in S$, $M, s' \models \varphi$.

If M is clear from context, we often just write $s \models \varphi$.

8.3.4. DEFINITION. A *topological model* is a triple $M = (S, \tau, V)$ where (S, τ) is a non-empty topological space (whose points are also called *states*),²⁸ and

²⁵Proof sketch: By induction on χ . If χ is atomic, this is immediate. If χ is of the form $\neg\chi_1$ or $\chi_1 \wedge \chi_2$, this follows from the induction hypothesis and axiom (D1). If χ is of the form $\Box\chi_1$ or $\boxplus\chi_1$, we have, by induction hypothesis, $\vdash \chi_1 \rightarrow \chi_1[\varphi/\psi]$, so, by necessitation, $\vdash \Box(\chi_1 \rightarrow \chi_1[\varphi/\psi])$, so, by the K -axiom and axiom (D1), $\vdash \Box\chi_1 \rightarrow \Box\chi_1[\varphi/\psi]$, i.e., $\vdash \chi \rightarrow \chi[\varphi/\psi]$. Similarly for \leftarrow . For \boxplus we reason analogously.

²⁶The relation R is reflexive iff, for all $s \in S$, sRs . The relation R is transitive iff, for all $s, s', s'' \in S$, if sRs' and $s'R s''$, then sRs'' . The relation R is symmetric iff, for all $s, s' \in S$, if sRs' , then $s'R s$.

²⁷Here $M, s \not\models \varphi$ is an abbreviation for it not being the case that $M, s \models \varphi$.

²⁸A *topological space* is a pair (S, τ) where S is a set (whose elements are called *points*) and τ is a set of subsets of S (whose elements are called *open sets*) such that (a) $\emptyset, S \in \tau$ and (b) τ is closed under arbitrary union and under finite intersection.

$V : S \rightarrow \mathcal{P}(\mathcal{L}_0)$ is a function (called *valuation*) assigning each state a set of atomic sentences.

We recursively define when an \mathcal{L} -sentence φ is true at a state s of M , which is denoted $M, s \models \varphi$:

- $M, s \models p$ iff $p \in V(s)$, for $p \in \mathcal{L}_0$.
- $M, s \models \neg\varphi$ iff $M, s \not\models \varphi$.
- $M, s \models \varphi \wedge \psi$ iff $M, s \models \varphi$ and $M, s \models \psi$.
- $M, s \models \Box\varphi$ iff there is $U \in \tau$ such that $s \in U$ and, for all $s' \in U$, $M, s' \models \varphi$.
- $M, s \models \Box\varphi$ iff, for all $s' \in S$, $M, s' \models \varphi$.

If M is clear from context, we often just write $s \models \varphi$. We write \models_K and \models_T , respectively, if we want to distinguish the modeling relation of Kripke semantics from that of topological semantics.

Thus, the main difference between the semantics is their semantic clause for \Box . From our stability perspective, this may be regarded as follows. Both clauses spell out that $\Box\varphi$ is true at a state s if φ is true at all sufficiently similar states. In the Kripke semantics, ‘sufficiently similar’ is spelled out via the accessibility relation: a state s' is sufficiently similar to s if sRs' . In the topological semantics, ‘sufficiently similar’ is spelled out via the topology: there is an open neighborhood U of s whose members are considered sufficiently similar to s ; so we can think of U as a *degree* of similarity. Thus, Kripke semantics can only provide an *absolute* notion of similarity: one that is the same for every sentence to which \Box is applied. While topological semantics provides a *relative* notion of similarity: the degree of similarity may differ among the sentences to which \Box is applied. In this light, Tarski's theorem is all the more astonishing: it shows that the two semantics have the same underlying modal logic.²⁹

These semantics provide examples of derivability relations as follows.

8.3.5. PROPOSITION. *Let C be a nonempty class of reflexive Kripke models or a class of topological models. Then the following defines a derivability relation:*

$$\vdash \varphi \text{ :} \Leftrightarrow \forall M \in C \forall s \in M : M, s \models \varphi.$$

We refer to \vdash as C -validity. In particular, if $C = \{M\}$ is a singleton, $\vdash \varphi$ means that φ is true at every state in M . In that case, we refer to \vdash as M -validity.

²⁹More precisely, reflexive and transitive Kripke models are sound and complete with respect to the very same modal logic (namely $S4$) as topological semantics. See Van Benthem and Bezhanishvili (2007) for a presentation of this result.

PROOF (SKETCH). Axiom (D1) holds since on non-modal formulas the states act like classical valuations. Axioms (D2) and (D3) are the basic axioms for a modal operator and the two semantics are sound for those. Axiom (D4) holds in Kripke semantics by the assumption of reflexivity and in the topological semantics by definition. Axiom (D5) holds since no state can make \perp true. \square

Questions Finally, we get to the last ingredient for a formalization of the discussion of stability: the set of questions Q . As indicated in the introduction and the examples, not every question (or proposition or hypothesis) is ‘sensible’ in a given situation: for example, it is a moot point whether a chair weighs exactly 5.3 *kg* because we can never establish the truth or falsity of this claim since every measuring device has some margin of error. Generally speaking, then, discussions of phenomena involving stability also come with a set Q of sentences that can, in that situation, sensibly be asked, conjectured, claimed, investigated, etc. For example, Q could be the set of good scientific hypotheses, i.e., the set of sentences that, if false, can also be falsified.

However, it is highly debatable what should count as a ‘sensible’ sentence in a given situation, let alone how this may be defined in general. Thus, we’ll treat Q as a parameter and only consider principles about Q . For example, one such principle may be that if a sensible proposition is false, it should be falsifiable.

8.3.6. DEFINITION. By a *set of questions* we simply mean a subset Q of \mathcal{L} -sentences. We refer to the elements of Q also as (sensible) sentences, questions, properties, hypotheses, or propositions.

In section 8.3.4, we present important examples of a set of questions: for example, Q could be the set of good scientific hypotheses $\{\varphi : \vdash \neg\varphi \rightarrow \Box\neg\varphi\}$.

8.3.2 Formalization and motivation of the principles

We now formalize and further motivate the intuitive principles (S1)–(S4) and their duals ($\bar{S}1$)–($\bar{S}4$) from the introduction. We’ll write $(\Sigma 1)$ for the formalization of (S1), and $(\bar{\Sigma} 1)$ for the formalization of ($\bar{S}1$), etc. A summary of the formal principles and their duals is given in figure 8.1.

Non-triviality principle The first principle says that ‘stable truth’ doesn’t trivialize to ‘simple truth’, or—even worse—to ‘possible truth’.

$(\Sigma 1) = (\bar{\Sigma} 1)$ There is $\varphi \in Q$ such that $\not\vdash (\Diamond\varphi \leftrightarrow \varphi) \wedge (\varphi \leftrightarrow \Box\varphi)$.

In other words, this principle demands that there is no *modal collapse* for sensible sentences. This way of putting it, provides an interpretation of the

($\Sigma 1$) $\exists \varphi \in Q : \vdash \varphi \leftrightarrow \diamond \varphi \leftrightarrow \square \varphi.$	($\bar{\Sigma} 1$) $\exists \varphi \in Q : \vdash \varphi \leftrightarrow \diamond \varphi \leftrightarrow \square \varphi.$
($\Sigma 2$) $\forall \varphi \in Q : \vdash \neg \varphi \rightarrow \square \neg \varphi.$	($\bar{\Sigma} 2$) $\forall \varphi \in Q : \vdash \varphi \rightarrow \square \varphi.$
($\Sigma 3$) $\forall \varphi \in Q : \vdash \varphi \rightarrow \diamond \square \varphi.$	($\bar{\Sigma} 3$) $\forall \varphi \in Q : \vdash \neg \varphi \rightarrow \diamond \square \neg \varphi.$
($\Sigma 4$) $\forall \varphi \in Q \exists \psi \in Q : \vdash \psi \leftrightarrow \mathbf{M}\varphi$ where $\mathbf{M}\varphi := \varphi \wedge \neg \square \varphi.$	($\bar{\Sigma} 4$) $\forall \varphi \in Q \exists \psi \in Q : \vdash \psi \leftrightarrow \mathbf{W}\varphi$ where $\mathbf{W}\varphi := \varphi \vee \square \neg \varphi.$

Figure 8.1: The four principles and their duals (formalized).

impossibility result as a modal collapse argument: If we take the other principles below as compelling premises, then the impossibility will imply the modal collapse for sensible sentences.³⁰

Also note that the non-triviality principle in particular implies that the set of sensible questions Q is nonempty.

Falsifiability principle and its dual The next principle demands that the sensible sentences have the kind of stability that they need: The falsifiability version demands that the falsity of sensible sentences is stable, and the dual verifiability version demands that the truth of sensible sentences is stable.

($\Sigma 2$) For all $\varphi \in Q$, $\vdash \neg \varphi \rightarrow \square \neg \varphi.$

($\bar{\Sigma} 2$) For all $\varphi \in Q$, $\vdash \varphi \rightarrow \square \varphi.$

This general motivation becomes more concrete in the various interpretations of stability discussed in section 8.2. Here are some examples.

Consider the epistemic understanding of stability (section 8.2.1). Here \mathcal{L}_0 are the sentences of a basic language to formulate observations and \mathcal{L} is the language obtained by also allowing statements about the stability of the observations. So \mathcal{L} is a ‘language of observation’. Then ($\Sigma 2$) says that sensible observational sentences need to satisfy the Popperian idea of falsifiability of scientific hypotheses: If φ is a sensible hypothesis that in fact is false in the actual state of the world, then it is possible to falsify it by making a fine enough measurement showing that the actual state lies in a $\neg \varphi$ -area. The dual principle ($\bar{\Sigma} 2$) expresses a verificationist standpoint: If φ is a sensible observational claim that is true in the actual state of the world, then it is possible to verify it by making a fine enough measurement showing that the actual state lies in a φ -area.

³⁰Cf. the modal collapse argument of Quine (1960, 181f.): substitution salva veritate in first-order modal logic collapses all modal truths to truth simpliciter. (See Føllesdal (2004) for discussion.)

Consider the significance understanding of stability (section 8.2.5). Then $(\Sigma 2)$ expresses the idea that the falsity of a sensible proposition is significant: If φ is false at a model s , this is not due to an idiosyncrasy of the mathematical model, but due to a structural reason, since sufficiently similar models also falsify φ . So from the falsity of sensible propositions in the mathematical models we can conclude that the corresponding property is also false in the modeled phenomenon (according to the thesis that significance can be spelled out as stability). The dual principle $(\bar{\Sigma} 2)$ says that the truth of sensible propositions is significant.

Similarly, this can be spelled out for other interpretations of stability.

Standard model principle and its dual The next principle demands that the sensible sentences have a standard (counter) model: If they can be made true (false) at all, there is a state where they are stably true (false).

$(\Sigma 3)$ For all $\varphi \in Q$, $\vdash \varphi \rightarrow \diamond \Box \varphi$.

$(\bar{\Sigma} 3)$ For all $\varphi \in Q$, $\vdash \neg \varphi \rightarrow \diamond \Box \neg \varphi$.

To illustrate this, consider again observational verification (section 8.2.1). We consider the possible weight-states x of a given chair. So the state space is S is the set of non-negative real numbers (with their usual topology). Consider the property $\varphi =$ ‘the weight x in kg of the chair is such that $4 < x < 6$ ’. So φ is verifiable (i.e., $\vdash \varphi \rightarrow \Box \varphi$).³¹ Now, it *can* happen that φ is false but cannot be falsified: e.g., in the unfortunate case that the actual weight-state x happens to be $x = 4$. Then x doesn’t have φ but any measurement will be inconclusive since it will delimit an area containing both φ -states and $\neg \varphi$ -states. But the point is that there is *some* state in which φ is stably false: for example, if $x = 3$, then x doesn’t have φ and we can falsify φ with a measurement with precision $\pm 0.5 kg$. Thus, φ satisfies $\vdash \neg \varphi \rightarrow \diamond \Box \neg \varphi$.

Similarly for other interpretations of stability: For example, on the significance interpretation, $(\bar{\Sigma} 3)$ says that for any non-trivial significant property—i.e., one that is a false at some model—, there should be a model where it is stably false. Only then, arguably, can we have an idea of what it means that the property φ ‘meaningfully’ fails—which seems like a plausible requirement for φ to be significant.

Moore closure principle and its dual To state the final principle, we first define the (dual) Moore operator. As noted in footnote 1, the name is due to the Moore sentence $\varphi \wedge \neg \Box \varphi$ in (formal) epistemology.

³¹With the usual argument (e.g. Vickers 1989, ch. 2): If the actual weight-state x of the chair has property φ —i.e., $x \in (4, 6)$ —, then there is an $\epsilon > 0$ such that all weight-states y that are ϵ -close to x have property φ —i.e., $(x - \epsilon, x + \epsilon) \subseteq (4, 6)$ —, so a measurement with precision ϵ will verify that the current state x has property φ .

8.3.7. DEFINITION. The *Moore operator* \mathbf{M} and the *dual Moore operator* \mathbf{W} are defined as

$$\begin{aligned}\mathbf{M}\varphi &:= \varphi \wedge \neg\Box\varphi \\ \mathbf{W}\varphi &:= \varphi \vee \Box\neg\varphi.\end{aligned}$$

They are dual in the sense that $\vdash \mathbf{M}\varphi \leftrightarrow \neg\mathbf{W}\neg\varphi$ for any derivability relation.³²

Now, the (dual) Moore closure principle demands that sensible sentences should be closed—up to equivalence—under the (dual) Moore operator.

($\Sigma 4$) If $\varphi \in Q$, then there is $\psi \in Q$ such that $\vdash \psi \leftrightarrow \mathbf{M}\varphi$.

($\overline{\Sigma} 4$) If $\varphi \in Q$, then there is $\psi \in Q$ such that $\vdash \psi \leftrightarrow \mathbf{W}\varphi$.

In words, if φ is a sensible claim, then also the claim ‘ φ but not stably so’ (respectively, the dual ‘ φ is false but not stably so’) is sensible. At first sight, this may be the least plausible principle, so let’s motivate it with some of our examples.

Consider again observational verifiability. We’ve already seen that the property $\varphi =$ ‘the weight x of the chair is $4 < x < 6$ ’, which intuitively seems sensible, is indeed verifiable ($\vdash \varphi \rightarrow \Box\varphi$) and also satisfies the standard countermodel principle ($\vdash \neg\varphi \rightarrow \Diamond\Box\neg\varphi$). The question that we’re concerned with now is: why should $\mathbf{W}\varphi = \varphi \vee \Box\neg\varphi$ be *sensible* again? It makes the hypothesis that (in the actual weight-state) we can *decide* the property φ by measurement: either φ is true, whence we can verify it by a measurement (since φ is verifiable), or φ is stably false, whence there is a measurement that conclusively shows that φ is false. This seems like a sensible hypothesis: after all, if it is true, it can be verified (since, as just outlined, in either of the two cases φ and $\Box\neg\varphi$ there is a confirming measurement).³³ Note again that the claim is only about the sensibility of $\mathbf{W}\varphi$ and *not* that $\mathbf{W}\varphi$ is true or valid. Indeed, just like the sensible φ *can* be false, also $\mathbf{W}\varphi$ can be false, namely at the boundary states $x = 4$ and $x = 6$.

Let’s also consider the falsifiability version. Similarly to the reasoning in the examples above, the property $\psi =$ ‘the weight x of the chair is ≤ 4 ’ is indeed falsifiable (if false, it is stably false) and also satisfies the standard model principle (there is a state, e.g. $x = 3$ where it stably true). The question that we’re concerned with now is: given ψ is sensible, why should $\mathbf{M}\psi = \psi \wedge \neg\Box\psi$ be sensible again? It makes the hypothesis that (in the actual weight-state) we *cannot decide* the property ψ by measurement: ψ is not false (otherwise we could falsify it) but

³²In \vdash we have the following equivalences: $\neg\mathbf{W}\neg\varphi \leftrightarrow \neg(\neg\varphi \vee \Box\neg\neg\varphi) \leftrightarrow \varphi \wedge \neg\Box\varphi \leftrightarrow \mathbf{M}\varphi$.

³³From an intuitionistic point of view, one may regard $\varphi \vee \Box\neg\varphi$ (which implies $\Box\varphi \vee \Box\neg\varphi$ if φ is stable) as a (classically formulated) form of excluded middle: it claims that we can ‘prove’ (or know or demonstrate) either φ or its negation.

it also is not stably true (so we could detect that with some measurement as well). Again, this seems like a sensible hypothesis: after all, it is falsifiable.³⁴

Consider the AI-safety interpretation of stability (section 8.2.3). Here the state space is given by the possible inputs to an AI and $s \models \varphi$ is interpreted as the AI (correctly) judging input s to have property φ . Then $\mathbb{W}\varphi$ says that, on the current input, the AI cannot be adversarially attacked to judge φ : it cannot be the case that, on the current input, the AI doesn't judge φ (e.g., because it judges the input to depict a stop sign), but on arbitrarily small alternations of the input it does judge φ (i.e., that the input depicts a 'Speed Limit 45' sign). This surely is a sensible (and relevant) hypothesis.

8.3.3 The duality between falsification and verification

There is a well-known duality between falsifiable sentences (if false, they can be shown to be false) and verifiable sentences (if true, they can be shown to be true): φ is falsifiable iff $\neg\varphi$ is verifiable (Vickers 1989, ch. 2). In our setting, this duality can be made formal as follows.

8.3.8. LEMMA. *For any derivability relation \vdash and choice of Q , we have, for $k = 1, \dots, 4$ (with $\neg Q := \{\neg\varphi : \varphi \in Q\}$):*

- (a) *If principle (Σk) holds for Q , then principle $(\bar{\Sigma}k)$ holds for $\neg Q$, and*
- (\bar{a}) *If principle $(\bar{\Sigma}k)$ holds for Q , then principle (Σk) holds for $\neg Q$.*

The straightforward proof is moved to the appendix.

8.3.9. PROPOSITION. *The following are equivalent:*

- (a) *For any derivability relation \vdash and choice of Q , at least one of $(\Sigma 1)$ – $(\Sigma 4)$ fails.*
- (\bar{a}) *For any derivability relation \vdash and choice of Q , at least one of $(\bar{\Sigma} 1)$ – $(\bar{\Sigma} 4)$ fails.*

PROOF. Assume (a). To show (\bar{a}), assume for contradiction that, for \vdash and Q , principles $(\bar{\Sigma} 1)$ – $(\bar{\Sigma} 4)$ are satisfied. Then, by lemma 8.3.8 (\bar{a}), principles $(\Sigma 1)$ – $(\Sigma 4)$ are satisfied for \vdash and $\neg Q$, contradicting (a). The other direction is analogous. \square

Thus, to prove our main result that both (a) and (\bar{a}) hold, it is enough to prove (a) (which will be theorem 8.4.2 below).

³⁴If $\mathbb{M}\psi$ is false because ψ is false, we can falsify it, and if $\mathbb{M}\psi$ is false because ψ is stably true, there is a measurement showing $\Box\psi$, whence falsifying $\mathbb{M}\psi$.

8.3.4 Constructing sets of questions

We provide two important constructions of sets of questions that satisfy principles $(\Sigma 2)$ and $(\Sigma 4)$ (resp. their duals): Following proposition 8.3.11 below, one starts with a ‘basic’ set of sentences that are falsifiable (verifiable) and closes it under a single operation of the (dual) Moore operator. Following proposition 8.3.12 below, one takes all falsifiable (verifiable) sentences. (For this, additional assumptions on the derivability relation apply.)

The key to those results is the following lemma which we prove in the appendix.

8.3.10. LEMMA. *Assume \vdash is a transitive derivability relation. Then*

(a) *If $\vdash \neg\varphi \rightarrow \Box\neg\varphi$, then $\vdash \neg M\varphi \rightarrow \Box\neg M\varphi$.*

(b) $\vdash MM\varphi \leftrightarrow M\varphi$.

(\bar{a}) *If $\vdash \varphi \rightarrow \Box\varphi$, then $\vdash W\varphi \rightarrow \Box W\varphi$.*

(\bar{b}) $\vdash WW\varphi \leftrightarrow W\varphi$.

In (b) and (\bar{b}) transitivity is not needed.

8.3.11. PROPOSITION. *Let \vdash be a transitive derivability relation. Let $Q_0 \subseteq \mathcal{L}$.*

(a) *If $\forall\varphi \in Q_0 : \vdash \neg\varphi \rightarrow \Box\neg\varphi$, then $Q := \{\varphi, M\varphi : \varphi \in Q_0\}$ satisfies principles $(\Sigma 2)$ and $(\Sigma 4)$.*

(\bar{a}) *If $\forall\varphi \in Q_0 : \vdash \varphi \rightarrow \Box\varphi$, then $Q := \{\varphi, W\varphi : \varphi \in Q_0\}$ satisfies principles $(\bar{\Sigma} 2)$ and $(\bar{\Sigma} 4)$.*

PROOF. Ad (a). Concerning $(\Sigma 2)$, let $\chi \in Q$ and show $\vdash \neg\chi \rightarrow \Box\neg\chi$. If $\chi = \varphi$ for some $\varphi \in Q_0$, then the claim follows by the assumption. If $\chi = M\varphi$ for some $\varphi \in Q_0$, then, by the assumption, $\vdash \neg\varphi \rightarrow \Box\neg\varphi$, so, by lemma 8.3.10 (a), $\vdash \neg M\varphi \rightarrow \Box\neg M\varphi$, as needed.

Concerning $(\Sigma 4)$, let $\chi \in Q$ and find $\psi \in Q$ such that $\vdash \psi \leftrightarrow M\chi$. If $\chi = \varphi$ for some $\varphi \in Q_0$, then $\psi := M\chi \in Q$ works. If $\chi = M\varphi$ for some $\varphi \in Q_0$, then $\psi := \chi \in Q$ works since, by lemma 8.3.10 (b), $\vdash \psi = M\varphi \leftrightarrow MM\varphi = M\chi$.

For (\bar{a}) we reason analogously using lemma 8.3.10 (\bar{a})–(\bar{b}). \square

8.3.12. PROPOSITION. *Let \vdash be a transitive derivability relation. Then*

(a) $Q := \{\varphi \in \mathcal{L} : \vdash \neg\varphi \rightarrow \Box\neg\varphi\}$ *satisfies principles $(\Sigma 2)$ and $(\Sigma 4)$.*

(\bar{a}) $Q := \{\varphi \in \mathcal{L} : \vdash \varphi \rightarrow \Box\varphi\}$ *satisfies principles $(\bar{\Sigma} 2)$ and $(\bar{\Sigma} 4)$.*

PROOF. Ad (a). Principle ($\Sigma 2$) holds by definition of Q . Concerning ($\Sigma 4$), if $\varphi \in Q$, choose $\psi := M\varphi$. Then, trivially, $\vdash \psi \leftrightarrow M\varphi$, and $\psi \in Q$ because, by lemma 8.3.10 (a), $\vdash \neg M\varphi \rightarrow \Box \neg M\varphi$.

Ad (\bar{a}). We reason analogously using lemma 8.3.10 (\bar{a}). \square

Note that $Q := \{\varphi \in \mathcal{L} : \vdash \neg\varphi \rightarrow \Box \neg\varphi\}$ violates principle ($\Sigma 1$) iff Q is closed under negation; and similarly for the dual version.³⁵

8.4 Impossibility via a novel interpretation of Fitch's paradox

We sketch Fitch's paradox and our reinterpretation of it in terms of stability (section 8.4.1). Inspired by this, we prove the main formal result: the inconsistency of the principles about stability (section 8.4.2).

8.4.1 Reinterpretation of Fitch's paradox

Fitch's paradox (or the Church-Fitch paradox or the paradox of knowability) says, as it is usually stated nowadays, the following: If all truths can be known ($\forall p : p \rightarrow \Diamond Kp$), then all truths are already known ($\forall p : p \rightarrow Kp$). The contrapositive implication was first published by Fitch (1963) acknowledging an anonymous referee who, much later, was discovered to be Alonzo Church. (See Brogaard and Salerno (2019) for a brief history.) Even though Fitch worked in the context of value concepts, the implication is mostly interpreted in the above form as an objection to verificationism, i.e., the view that all truths are knowable (Brogaard and Salerno 2019).

In more precise words, the *formal* implication

$$\forall p : p \rightarrow \Diamond Kp \Rightarrow \forall p : p \rightarrow Kp,$$

which is also called Fitch's lemma, can be established with fairly minimal proof-theoretic assumptions on the modal operators \Diamond and K ; and it gains philosophical meaning by *interpreting* the quantifier as ranging over all declarative statements and the modal operators as describing metaphysical possibility and knowledge, respectively.

Here we wish to provide another interpretation in terms of stability: First, we make explicit the set of sentences Q that we quantify over. We can recover the original Fitch paradox by choosing Q to be the set of all (declarative) sentences,

³⁵Proof: If Q violates ($\Sigma 1$) and $\varphi \in Q$, then $\vdash \varphi \rightarrow \Box \varphi$, i.e., $\vdash \neg\neg\varphi \rightarrow \Box \neg\neg\varphi$, whence $\neg\varphi \in Q$. Conversely, if Q is closed under negation and $\varphi \in Q$, we need to show $\vdash \Diamond\varphi \leftrightarrow \varphi \leftrightarrow \Box\varphi$. By reflexivity, $\vdash \Box\varphi \rightarrow \varphi$ and $\vdash \varphi \rightarrow \Diamond\varphi$. Since $\varphi \in Q$, $\vdash \neg\varphi \rightarrow \Box\neg\varphi$, i.e., $\vdash \Diamond\varphi \rightarrow \varphi$. Since $\neg\varphi \in Q$, $\vdash \neg\neg\varphi \rightarrow \Box\neg\neg\varphi$, i.e., $\vdash \varphi \rightarrow \Box\varphi$.

but we allow for a more refined analysis of the paradox by restricting the set of 'allowed' sentences. For example, this way we could account for the worry that not all sentences are subject to verificationism but only those of a certain logical form. Second, and most importantly, we suggest a different interpretation of the modal operators. We (re-) interpret the metaphysical possibility operator \diamond as the 'true in some state' operator (i.e., global possibility) \lozenge and we (re-) interpret the knowledge operator K as the 'true in all sufficiently similar states' operator \square (i.e., stability).

Inspired by this interpretation, we'll now prove the impossibility result. But the interpretation could potentially also find other applications: The research on Fitch's paradox of knowability (e.g., ways to avoid it) may yield insights about stability, and vice versa the interpretation in terms of stability may suggest new avenues in the case of knowledge—but we leave that to future research.

8.4.2 Impossibility

Reinterpreted in terms of stability, the Fitch's lemma states: If, for all sentences of Q , truth implies stable truth somewhere, then, for all sentences of Q , truth already implies stable truth. Formally:

8.4.1. LEMMA (Fitch's lemma). *Let \vdash be a derivability relation and Q a set of sentences satisfying $(\Sigma 4)$. Then (i) implies (ii) where*

$$(i) \quad \forall \varphi \in Q : \quad \vdash \varphi \rightarrow \lozenge \square \varphi$$

$$(ii) \quad \forall \varphi \in Q : \quad \vdash \varphi \rightarrow \square \varphi.$$

PROOF. The proof is a slight adaption of the standard proof of Fitch's lemma. (See, e.g., Van Benthem (2004) or Brogaard and Salerno (2019).) Assume (i) and let $\varphi \in Q$. To show $\vdash \varphi \rightarrow \square \varphi$, we show 'by contradiction' that $\vdash \neg(\varphi \rightarrow \square \varphi) \rightarrow \perp$ (using that $(\neg p \rightarrow \perp) \rightarrow p$ is a classical tautology). By $(\Sigma 4)$, let $\psi \in Q$ with $\vdash \psi \leftrightarrow \mathbf{M}\varphi$. Then we can show, in \vdash , the following chain of conditionals:

$$\begin{array}{ll}
 \neg(\varphi \rightarrow \square \varphi) & \rightarrow \quad \varphi \wedge \neg \square \varphi & \text{[classical logic]} \\
 & \rightarrow \quad \psi & [\vdash \psi \leftrightarrow \mathbf{M}\varphi] \\
 & \rightarrow \quad \lozenge \square \psi & [(i)] \\
 & \rightarrow \quad \lozenge \square (\varphi \wedge \neg \square \varphi) & [\vdash \psi \leftrightarrow \mathbf{M}\varphi] \\
 & \rightarrow \quad \lozenge (\square \varphi \wedge \square \neg \square \varphi) & [\vdash \square (p \wedge q) \rightarrow \square p \wedge \square q] \\
 & \rightarrow \quad \lozenge (\square \varphi \wedge \neg \square \varphi) & [\vdash \square p \rightarrow p] \\
 & \rightarrow \quad \perp & [\vdash \lozenge \perp \rightarrow \perp^{36}]
 \end{array}$$

³⁶Since $\vdash \neg \perp$, necessitation implies $\vdash \square \neg \perp$, so $\vdash \neg \lozenge \perp$, so $\vdash \lozenge \perp \rightarrow \perp$.

where we use the closure under equivalence and the fact that, if $\vdash \psi \rightarrow \chi$, then $\vdash \diamond\psi \rightarrow \diamond\chi$.³⁷ \square

This lemma is the key ingredient to obtain the announced impossibility result:

8.4.2. THEOREM. *For any choice of derivability relation \vdash and set of questions Q , the principles $(\Sigma 1)$ – $(\Sigma 4)$ are inconsistent. By proposition 8.3.9, also principles $(\bar{\Sigma} 1)$ – $(\bar{\Sigma} 4)$ are inconsistent.*

PROOF. We show that $(\Sigma 2)$ – $(\Sigma 4)$ imply the negation of $(\Sigma 1)$. By $(\Sigma 4)$, lemma 8.4.1 applies, so $(\Sigma 3)$ implies

$$\forall \varphi \in Q : \vdash \varphi \rightarrow \square\varphi.$$

By $(\Sigma 2)$, for all $\varphi \in Q : \vdash \neg\varphi \rightarrow \square\neg\varphi$, i.e., $\vdash \diamond\varphi \rightarrow \varphi$. By reflexivity, \vdash proves $\square\varphi \rightarrow \varphi \rightarrow \diamond\varphi$. So

$$\forall \varphi \in Q : \vdash \diamond\varphi \leftrightarrow \varphi \leftrightarrow \square\varphi,$$

which is the negation of $(\Sigma 1)$. \square

Note that we have made very few assumptions about the modality \diamond . We didn't even use that it is reflexive, let alone anything that would suggest that it has been intuitively introduced as a global truth operator (for which something like the *S5* axioms would be appropriate). By adding the very mild reflexivity assumption, we can strengthen the inconsistency as follows.

8.4.3. COROLLARY. *Let \vdash be a derivability relation and Q a set of questions. Assume $\vdash \varphi \rightarrow \diamond\varphi$. Then, if principles $(\Sigma 2)$ and $(\Sigma 4)$ are satisfied, exactly one of $(\Sigma 1)$ and $(\Sigma 3)$ is satisfied. Similarly, for the dual principles.*

PROOF. Assume $(\Sigma 2)$ and $(\Sigma 4)$ are satisfied. By the inconsistency, it cannot be that both $(\Sigma 1)$ and $(\Sigma 3)$ are satisfied, so it suffices to show that if $(\Sigma 1)$ is not satisfied, then $(\Sigma 3)$ is. Indeed, if $(\Sigma 1)$ is false, then $\forall \varphi \in Q : \vdash \varphi \rightarrow \square\varphi$. By reflexivity, $\vdash \square\varphi \rightarrow \diamond\square\varphi$. So $(\Sigma 3)$ holds. Similarly for the dual principles (using that falsity of $(\bar{\Sigma} 1)$ implies $\forall \varphi \in Q : \vdash \neg\varphi \rightarrow \square\neg\varphi$). \square

Before coming to applications of this inconsistency, we first investigate it semantically.

³⁷If $\vdash \psi \rightarrow \chi$, then $\vdash \neg\chi \rightarrow \neg\psi$, so $\vdash \boxdot(\neg\chi \rightarrow \neg\psi)$, so, by the *K*-axiom, $\vdash \boxdot\neg\chi \rightarrow \boxdot\neg\psi$, so $\vdash \neg\boxdot\chi \rightarrow \neg\boxdot\psi$, so $\vdash \boxdot\psi \rightarrow \boxdot\chi$.

8.5 Impossibility via semantics

We can also prove the inconsistency for the ‘semantic’ derivability relations from proposition 8.3.5 in a semantic way (rather than proof-theoretic as above). This works both for Kripke semantics and topological semantics.

8.5.1. THEOREM. *Let C be a class of reflexive Kripke models or a class of topological models, and let \vdash be C -validity. Then*

(i) *If $\nVdash \varphi \rightarrow \Box\varphi$, then $\nVdash \mathbf{M}\varphi \rightarrow \Diamond\Box\mathbf{M}\varphi$.*

(ii) *If Q is a set of sentences, then Q and \vdash cannot satisfy all of $(\Sigma 1)$ – $(\Sigma 4)$.*

Similarly for the dual version.

PROOF. Ad (i). By the assumption, there is a model M and a state s such that $M, s \models \varphi \wedge \neg\Box\varphi = \mathbf{M}\varphi$. If we had $\vdash \mathbf{M}\varphi \rightarrow \Diamond\Box\mathbf{M}\varphi$, then $M, s \models \Diamond\Box\mathbf{M}\varphi$. So there is a state s' such that $M, s' \models \Box\mathbf{M}\varphi = \Box(\varphi \wedge \neg\Box\varphi)$. However, that is a contradiction: It implies $s' \models \Box\varphi \wedge \Box\neg\Box\varphi$, which implies, by reflexivity, $s' \models \Box\varphi \wedge \neg\Box\varphi$, but s' cannot make true a contradiction.

Ad (ii). Assume for contradiction that Q and \vdash satisfy $(\Sigma 1)$ – $(\Sigma 4)$. By $(\Sigma 1)$, there is $\varphi \in Q$ such that $\nVdash \varphi \leftrightarrow \Diamond\varphi$ or $\nVdash \varphi \leftrightarrow \Box\varphi$. By (D4), $\vdash \varphi \rightarrow \Diamond\varphi$ and $\vdash \Box\varphi \rightarrow \varphi$. By the contraposition of $(\Sigma 2)$, $\vdash \Diamond\varphi \rightarrow \varphi$. Hence $\nVdash \varphi \rightarrow \Box\varphi$. So, by (i), $\nVdash \mathbf{M}\varphi \rightarrow \Diamond\Box\mathbf{M}\varphi$. Finally, by $(\Sigma 4)$, $\mathbf{M}\varphi$ is equivalent to some $\psi \in Q$, so also $\nVdash \psi \rightarrow \Diamond\Box\psi$, which is a contradiction to $(\Sigma 3)$. \square

In this section, we'll collect some further corollaries specific to the two types of semantics and illustrate them by examples.

8.5.1 Kripke semantics

The following is an example of a more concrete form of the impossibility in the case of Kripke semantics.

8.5.2. COROLLARY. *Let C be a nonempty class of reflexive and transitive Kripke models, let \vdash be C -validity, and let $Q := \{\varphi : \vdash \varphi \rightarrow \Box\varphi\}$. Then the following are equivalent:*

(i) *Q is closed under negation.*

(ii) *For all $\varphi \in Q$ and models $M \in C$, if φ is falsifiable in M , then there is a state s of M such that $s \models \Box\neg\varphi$.*

PROOF. By proposition 8.3.5, \vdash is a derivability relation. By proposition 8.3.12, Q satisfies principles $(\overline{\Sigma} 2)$ and $(\overline{\Sigma} 4)$. By the impossibility result (corollary 8.4.3), exactly one of $(\overline{\Sigma} 1)$ and $(\overline{\Sigma} 3)$ hold. Now the claim follows since (i) is equivalent to the negation of $(\overline{\Sigma} 1)$ and (ii) is equivalent to $(\overline{\Sigma} 3)$. \square

8.5.3. EXAMPLE. Consider again the stability theory of belief (section 8.2.4). Let S be the set of probability measures over a nonempty finite set Ω . So the elements of S are the agent's possible degrees of belief. Let the set of atomic sentences \mathcal{L}_0 be given by the subsets of Ω (the 'propositions'). Let's assume for this example that the set \mathcal{Y} of propositions that the agent may suppose is closed under conjunction and contains the trivial proposition Ω and no propositions with zero probability. Define $\mu R \nu$ iff, for some $q \in \mathcal{Y}$, ν is obtained by updating μ with q , i.e., $\nu = \mu(\cdot|q)$. Finally, the valuation V is defined by: $\mu \models p$ iff p is Locke-believed at μ , i.e., $\mu(p) > \rho$ (where $\frac{1}{2} \leq \rho < 1$ is some fixed threshold).

As discussed in section 8.2.4, Hume-believing p at μ means $\mu \models \Box p$. Naturally, we're also interested in nested claims of Locke- and Hume-belief (i.e., sentences in \mathcal{L} and not just \mathcal{L}_0) and when the two notions (dis)agree: i.e., we want to understand the set $Q := \{\varphi \in \mathcal{L} : \vdash \varphi \rightarrow \Box \varphi\}$.

The corollary helps: The Kripke model (S, R, V) is reflexive (since $\mu = \mu(\cdot|\Omega)$) and transitive (if $\mu_2 = \mu_1(\cdot|q)$ and $\mu_3 = \mu_2(\cdot|q')$, then $\mu_3 = \mu_1(\cdot|q \cap q')$). Thus, (i) and (ii) are equivalent. Here (i) says: for all $\varphi \in Q$, if the agent doesn't have high degree of belief in φ , then no \mathcal{Y} -update will change her mind into a high degree of belief in φ . In contrast, (ii) only says that this *can* happen (but need not happen everywhere): for all $\varphi \in Q$, given low degree of belief in φ is possible at all, the agent *can* have degrees of belief in which she has low belief in φ and this doesn't change after \mathcal{Y} -updates.

8.5.4. EXAMPLE. Let's consider a particular instance of stability as invariance under transformations (section 8.2.6). Let S be a set of models in some given signature (e.g., the set of all countable rings up to isomorphism) and define $A R B$ iff there is a surjective homomorphism $f : A \rightarrow B$ (so B is a factor of A). Let \mathcal{L}_0 be the set of first-order sentences and define $p \in V(A)$ iff $A \models p$. We're interested in the sentences (or properties) whose truth is invariant under the transformation of taking factors. For first-order sentences (or formulas), this is a classical question in model theory whose answer, in our case, is Lyndon's Positivity Theorem (see e.g. Rossman 2008). It says that, among the first-order sentences, it is precisely the positive ones (those without negations) whose truth is invariant under factors. However, the question naturally extends to nested invariance claims. So we ask: what is the set $Q = \{\varphi \in \mathcal{L} : \vdash \varphi \rightarrow \Box \varphi\}$?

Again the corollary provides a partial answer with the equivalence of (i) and (ii). Condition (i) fails rather easily: in the case of rings, we have $\mathbb{Z} R \mathbb{Z}/3$ via the natural homomorphism and, for $p := \exists x(x(1+1) = 1) \in Q$, we have $\mathbb{Z}/3 \models p$ but $\mathbb{Z} \not\models p$, whence $\neg p \notin Q$. So (ii) fails, too: There is a sentence φ which not only is preserved to factors but locally also its negation is preserved to factors (i.e., there is a model A such that, for all factors B , $B \not\models \varphi$).³⁸

³⁸More generally, for combinations of modal logic and set/model theory, see, e.g., Hamkins and Löwe (2008) and Hamkins and Wołoszyn (2020).

In the formulation of the corollary, one could wonder: if the intended interpretation of the accessibility relation of the Kripke models is similarity (so \Box means true in all similar states), then one should rather demand R to be reflexive and symmetric (instead of transitive). However, we now show that this faces an arguably even stronger no-go result.

We call a Kripke model (S, R, V) is *connected* if any two states can be connected by a path, i.e., for all $s, s' \in S$, there are $s_0, s_1, \dots, s_n \in S$ such that $s = s_0$, $s' = s_n$ and $s_i R s_{i+1}$ for all $i = 0, \dots, n - 1$. Intuitively, being connected means that we're considering a 'single' state space and not, really, two or more separate ones. If a (reflexive and symmetric) Kripke model is not connected, one would consider its connected components.

The next proposition then provides the no-go result: In connected Kripke models, the only stable sentences are the trivial ones.

8.5.5. PROPOSITION. *If M is a connected Kripke model, then*

$$\{\varphi \in \mathcal{L} : \vdash \varphi \rightarrow \Box\varphi\} = \{\varphi \in \mathcal{L} : \vdash \varphi \text{ or } \vdash \neg\varphi\},$$

where \vdash is M -validity (i.e., truth at all states of M).

PROOF. The inclusion \supseteq is immediate. For the other inclusion, let $\varphi \in \mathcal{L}$ be such that $\vdash \varphi \rightarrow \Box\varphi$. If φ weren't in the set on the right, there is a state $s \not\models \varphi$ and a state $s' \not\models \neg\varphi$, i.e., $s' \models \varphi$. Since M is connected, there is a path $s' = s_0 R s_1 R \dots R s_n = s$. Since $s' = s_0 \models \varphi$ and φ is stable, also $s_1 \models \varphi$. We continue inductively and obtain $s_n \models \varphi$, which contradicts $s_n = s \not\models \varphi$. \square

This will become relevant in the discussion of AI-safety below (section 8.6.2).

8.5.2 Topological semantics

The beauty of topological semantics is that it relates concepts of modal logic to topological ones (Van Benthem and Bezhanishvili 2007). The following proposition collects some basic such examples: it relates the logical concepts of stability, falsifiability, stable truth, Moore operator, etc., to respective topological concepts: openness, closedness, interior operator, boundary operator, etc.

8.5.6. PROPOSITION (Folklore³⁹). *Let (S, τ, V) be a topological model. Write $\llbracket \varphi \rrbracket := \{s \in S : s \models \varphi\}$. Then*

³⁹It is difficult to exactly determine the origins of these results. They include some of the main conceptual insights of topological semantics (especially (iii) below) going back to Tarski (and McKinsey), and most are found in standard treatments of the topic (Van Benthem and Bezhanishvili 2007). Moreover, the idea expressed in (i) and (ii) that open (resp. closed) sets correspond to verifiable (resp. falsifiable) properties is due to the 'topology via logic' approach (albeit not strictly speaking formulated in the context of topological semantics). If anything, the specific formulation chosen here with an eye toward stability is new.

- (i) $\llbracket \varphi \rrbracket$ is open iff $\varphi \rightarrow \Box \varphi$ is valid on S (i.e., true at every state).
- (ii) $\llbracket \varphi \rrbracket$ is closed iff $\neg \varphi \rightarrow \Box \neg \varphi$ is valid on S .
- (iii) $\llbracket \Box \varphi \rrbracket = \text{Int} \llbracket \varphi \rrbracket$ and $\llbracket \Diamond \varphi \rrbracket = \text{Cl} \llbracket \varphi \rrbracket$.⁴⁰
- (iv) If $\llbracket \varphi \rrbracket$ is closed, then $\llbracket \mathbf{M}\varphi \rrbracket = \delta \llbracket \varphi \rrbracket$.⁴¹
- (v) If (S, τ) is connected and φ non-trivial (neither φ nor $\neg \varphi$ are valid in S), then $(\Diamond \varphi \leftrightarrow \varphi) \wedge (\varphi \leftrightarrow \Box \varphi)$ is not valid in S .

PROOF. Ad (i). $\llbracket \varphi \rrbracket$ is open iff for all $s \in S$, if $s \in \llbracket \varphi \rrbracket$, then there is an open set U such that $s \in U \subseteq \llbracket \varphi \rrbracket$ iff $\vdash \varphi \rightarrow \Box \varphi$ is valid.

Ad (ii). $\llbracket \varphi \rrbracket$ is closed iff $\llbracket \neg \varphi \rrbracket$ is open iff $\vdash \neg \varphi \rightarrow \Box \neg \varphi$ is valid.

Ad (iii). $\llbracket \Box \varphi \rrbracket = \{s \in S : \exists U \in \tau . s \in U \subseteq \llbracket \varphi \rrbracket\} = \bigcup \{U \in \tau : U \subseteq \llbracket \varphi \rrbracket\} = \text{Int} \llbracket \varphi \rrbracket$. And $\llbracket \Diamond \varphi \rrbracket = \llbracket \Box \neg \varphi \rrbracket^c = (\text{Int} \llbracket \varphi \rrbracket^c)^c = \text{Cl} \llbracket \varphi \rrbracket$.

Ad (iv). Since $\llbracket \varphi \rrbracket$ is closed, $\llbracket \mathbf{M}\varphi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \neg \Box \varphi \rrbracket = \text{Cl} \llbracket \varphi \rrbracket \cap (\text{Int} \llbracket \varphi \rrbracket)^c = \delta \llbracket \varphi \rrbracket$.

Ad (v). If the formula were valid, $\llbracket \varphi \rrbracket$ is a closed and open set which is neither \emptyset nor S , which contradicts connectedness. \square

Here are but some deep results about the connection between modal logic and topology relevant to our setting: The global modality increases expressivity (Shehtman 1999): connectedness cannot be defined using only \Box , but it can be defined additionally using \boxplus by the C -axiom $\boxplus(\Box \varphi \vee \Box \neg \varphi) \rightarrow \boxplus \varphi \vee \boxplus \neg \varphi$. The natural ‘basic’ logic for the bimodal logic where \Box is interpreted topologically and \boxplus as global truth is $S4U$ (Bennett 1996; Shehtman 1999): $S4$ for \Box , plus $S5$ for \boxplus , plus the bridge axiom $\boxplus \varphi \rightarrow \Box \varphi$. $S4U + C$ is the logic of any connected dense-in-itself separable metric space (Shehtman 1999). It is also the logic of the algebra of countable unions of convex subsets of the real line (Bezhanishvili and Gehrke 2005). $S4U$ is the logic of the measure algebra of subsets of the Euclidean space with positive Lebesgue measure (Fernández-Duque 2010), which is a semantics originally suggested by Dana Scott (also see Lando 2012; Lando 2015). For a detailed study of expressivity and definability in modal logics for topology, see Cate, Gabelaia, and Sustretov (2009). For adding dynamics, see Kremer and Mints (2007) discussing modal logics for topological dynamical systems.

Using the topological semantics, we can provide yet another, conceptually different proof of the impossibility.

8.5.7. THEOREM. *Let $M = (S, \tau, V)$ be a topological model and $\vdash M$ -validity.*

- (i) *If $\llbracket \varphi \rrbracket$ is closed in S but not open, then $\not\vdash \mathbf{M}\varphi \rightarrow \Diamond \Box \mathbf{M}\varphi$.*
- (ii) *If Q is a set of sentences, then Q and \vdash cannot satisfy all of $(\Sigma 1)$ – $(\Sigma 4)$.*

⁴⁰Here Int is the topological interior and Cl is the topological closure.

⁴¹Here δ is the topological boundary.

Similarly for the dual version.

PROOF. Ad (i). Since $\llbracket \varphi \rrbracket$ is closed and not open, $\emptyset \neq \delta \llbracket \varphi \rrbracket = \llbracket \mathbf{M}\varphi \rrbracket$. If we had $\vdash \mathbf{M}\varphi \rightarrow \diamond \square \mathbf{M}\varphi$, there would be a state $s \models \diamond \square \mathbf{M}\varphi$, so $\emptyset \neq \llbracket \square \mathbf{M}\varphi \rrbracket = \text{Int} \llbracket \mathbf{M}\varphi \rrbracket = \text{Int} \delta \llbracket \varphi \rrbracket = \emptyset$, contradiction.

Ad (ii). Assume for contradiction that Q and \vdash satisfy $(\Sigma 1)$ – $(\Sigma 4)$. By $(\Sigma 1)$, there is $\varphi \in Q$ such that $\not\vdash \square \varphi \leftrightarrow \varphi \leftrightarrow \diamond \varphi$. So $\llbracket \varphi \rrbracket$ is either not closed or not open. Since, by $(\Sigma 2)$, $\llbracket \varphi \rrbracket$ is closed, it is not open. By $(\Sigma 4)$, $\mathbf{M}\varphi$ is equivalent to some $\psi \in Q$. So, by (i), $S \neq \llbracket \mathbf{M}\varphi \rightarrow \diamond \square \mathbf{M}\varphi \rrbracket = \llbracket \psi \rightarrow \diamond \square \psi \rrbracket$, contradicting $(\Sigma 3)$. \square

A more concrete corollary that we'll use below is the following.

8.5.8. COROLLARY. *Let S be a set of states and let $p \subseteq S$ be a non-trivial subset (i.e., $p \neq \emptyset$ and $p \neq S$). We identify the set p with the atomic sentence p (i.e., we choose a valuation V on S with $p \in V(s)$ iff $s \in p$). Let τ be a connected topology on S . If p is open, then it is not entirely but generically true that limits of p -states are p -states, i.e., the set*

$$\llbracket \mathbf{W}p \rrbracket = \{s \in S : \text{if } s \text{ is the limit of states in } p, \text{ then } s \in p\} \quad (8.1)$$

is not S but generic (i.e., open and dense). Similarly for the dual version.

PROOF. Since (S, τ) is connected and $\neg p$ is closed and non-trivial, $\neg p$ cannot be open. By theorem 8.5.7 (i), $\not\vdash \mathbf{M}\neg p \rightarrow \diamond \square \mathbf{M}\neg p$, so there is a state s with $s \models \mathbf{M}\neg p$ and $s \not\models \neg \diamond \square \mathbf{M}\neg p$. The former implies $s \notin \llbracket \neg \mathbf{M}\neg p \rrbracket = \llbracket \mathbf{W}p \rrbracket$ and the latter implies $S = \llbracket \diamond \neg \mathbf{M}\neg p \rrbracket = \text{Cl} \llbracket \mathbf{W}p \rrbracket$. Hence $\llbracket \mathbf{W}p \rrbracket$ is dense. It also is open: $\llbracket \mathbf{W}p \rrbracket = \llbracket p \vee \square \neg p \rrbracket = p \cup \text{Int}(p^c)$. So it remains to show equation (8.1). The set on the right can be written as $\{s \in S : s \notin \text{Cl}p \text{ or } s \in p\}$ which further equals $(\text{Cl} \llbracket p \rrbracket)^c \cup \llbracket p \rrbracket = \llbracket \neg \diamond p \vee p \rrbracket = \llbracket \mathbf{W}p \rrbracket$ as needed. \square

We'll use the corollary in the discussion of AI-safety below (section 8.6.2). But, for now, we end this section by applying it to the example of Markov chains mentioned in the context of significance in mathematical models (section 8.2.5).

8.5.9. EXAMPLE. Recall that (finite-state and time-homogeneous) Markov chains correspond to stochastic matrices (the transition matrices of Markov chains). So, as state space, we consider the set $S := \mathcal{M}_n^s$ of stochastic $n \times n$ matrices (i.e., $n \times n$ matrices with real non-negative entries such that each row sums up to 1). The natural topology on \mathcal{M}_n^s is the subspace topology inherited from the space \mathcal{M}_n of all real-valued $n \times n$ matrices with the topology induced by any matrix norm.

As property p , we'll consider ergodicity: the property that, roughly, we can investigate the system (modeled by the Markov chain) through sampling: we

can come to know the average value of a measurement performed on the system repeatedly over time (‘time average’) though the stochastically expected value of single performance of the measurement (‘space average’). In the case of Markov chains, this property fortunately is equivalent to a simpler one: a Markov chain is ergodic iff its transition matrix A is irreducible, i.e., for all $1 \leq i, j \leq n$ there is $k \geq 1$ such that $A^k(i, j) > 0$ (Petersen 1983, p. 53). So we consider the set $p := \mathcal{M}_n^{\text{irr}}$ of irreducible stochastic matrices.

In the appendix (lemma 8.7.1), we show that (a) $S = \mathcal{M}_n^s$ is a connected topological space, (b) $p = \mathcal{M}_n^{\text{irr}}$ is an open subset of \mathcal{M}_n^s , and (c) p is a non-trivial subset of S (assuming $n \geq 2$). Thus, the property of being ergodic is not only significant ($p \rightarrow \Box p$ is valid) but it also is a generic truth that every Markov chain either is ergodic or ‘significantly non-ergodic’ ($p \vee \Box \neg p$ is valid). So, generically, if during a modeling process we refine our ergodic Markov chain model of a given real phenomenon ever more finely, the limit of the modeling process will again be an ergodic Markov chain.

Another example in section 8.2.5 was provided by general relativity. Here the discussion of Fletcher (2016) is a good illustration of the advantage of using topology to describe stability (as opposed to a binary relation). Its more sophisticated structure allows for a more careful analysis: for example, why to favor one topology (as a description of stability) over another.

8.6 Applications

Among the many examples of stability from section 8.2, we’ll restrict us to applying the impossibility result to two of them: In section 8.6.1, we consider verifiability and falsifiability (or knowability and scientific hypotheses) and obtain an extension of Fitch’s paradox. In section 8.6.2, we consider AI-safety and find some fundamental limitations.

8.6.1 An extension of Fitch’s paradox

Fitch’s paradox is much discussed (Brogaard and Salerno 2019), but one commonly drawn lesson is that we cannot come to know every true sentence. As a response to this, we may ask: What, then, is the set of sentences Q such that, whenever they are true in a situation, we can come to know that they are true? In this subsection, we use the impossibility to find limitations on this set Q of knowable sentences—and thus, in a sense, obtain an extension of Fitch’s paradox. After that, we’ll also discuss the dual case of scientific hypotheses: What is the set Q of good scientific hypotheses? Again, the impossibility result implies some limitations on Q .

Knowability In a sense, Fitch's paradox of knowability is most striking when applied to the most basic instances of knowability: obtaining knowledge about objects through observation. In other words, if we talk about very complicated and hard to grasp topics, it may be less surprising that there are sentences which we cannot come to know. But if we only consider sentences of a basic observational language, this is much more worrying. For this reason—and for concreteness—, we'll take our atomic sentences (i.e., the set \mathcal{L}_0) to be observational statements (like 'object a is red' or 'object b weighs 5 kg'). Thus, the language \mathcal{L} is obtained from the basic observational statements using logical connectives, the stability operator, and the global truth operator.

In the face of Fitch's paradox (and also for independent reasons), we may then wonder which sentences are knowable: what are the observational sentences that have a logical form that ensures that, if true, we can come to know their truth through observation or measurement. As we'll argue now, the impossibility result shows that we have to give up at least one of the following:

- ($\bar{1}$) The safety condition for knowledge: if we know φ , then φ is stably true.
- ($\bar{2}$) Stability ensures knowability: if φ is stably true, there is some measurement that will yield knowledge that φ .
- ($\bar{3}$) Non-triviality: for some knowable sentence φ , we can continuously transform a φ -state into a $\neg\varphi$ state. For example, if an object is red in a state, we can 'imagine' to continuously change the color of the object until it is blue, i.e., we're in a non-red state. We phrase this more precisely below.
- ($\bar{4}$) Standard countermodel: if a knowable sentence is false, there is some state in which we can come to know its falsity.

Thus, for example, given ($\bar{1}$)–($\bar{3}$), we can conclude that there are knowable sentences that can be false in some states, yet we will never come to know their falsity. We may regard this as an extension of (the common conclusion of) Fitch's paradox: Not only are there true sentences that we cannot come to know, but also, even among the sentences that we can come to know whenever they are true, there are sentences that can be false while we cannot come to know their falsity.

We'll now discuss ($\bar{1}$)–($\bar{4}$) and their joint inconsistency in more detail. We write Q for the set of knowable sentences (i.e., if true, we can come to know it) in our 'observation language' \mathcal{L} .

Regarding ($\bar{1}$), we've already discussed the safety condition for knowledge in section 8.2.2. It says that if we know φ in situation s , then φ is not only true in s but also in all sufficiently similar situations s' . The point of ($\bar{1}$) is that it implies that all sentences φ in Q are stable: If φ is true at a state s , we can come to know it (by definition of being in Q), so, since knowledge implies stability, φ is also true in all sufficiently similar states, so $\Box\varphi$ is true at s .

The idea behind $(\bar{2})$ is that the similarity structure on the state space (which determines the notion of stability) ‘reflects’ observability. For example, in the ‘topology via logic’ approach mentioned in footnote 2, the similarity structure is given by a topology (as in topological semantics) and the (basic) open sets in the topology precisely correspond to the possible observations. Thus, if φ is stably true at a state s , then s is in the open set $\llbracket \Box\varphi \rrbracket$, so there is a possible measurement confirming this (i.e., a basic open set U such that $s \in U \subseteq \llbracket \varphi \rrbracket$). The point of $(\bar{2})$ is that it implies that all stable sentences φ are in Q : If φ is true at a state s , it is, by stability, also stably true, so there is some measurement that will yield knowledge that φ . So $\varphi \in Q$: if true, we can come to know it.

To summarize, if $(\bar{1})$ and $(\bar{2})$ are correct, then the set Q of knowable sentences is precisely the set of sentences that are stable: $Q = \{\varphi : \vdash \varphi \rightarrow \Box\varphi\}$ where \vdash is the ‘true at all states’ relation. In other words, the stability operator \Box coincides with the ‘can come to know’ operator. (This is different to the original Fitch paradox where ‘can come to know’ is expressed as $\Diamond K$.)

Next, we argue that $\vdash \Box\varphi \rightarrow \Box\Box\varphi$, i.e., if φ is stably true, also the claim “ φ is stably true” is stably true. We’ve seen (the) two ways to describe a similarity structure on the state space which determines the notion of stability: either absolutely via a binary relation (Kripke semantics) or in degrees via a topology (topological semantics). In the former case, if the binary relation describes a similarity relation, it is natural to assume that it is reflexive (any state is similar to itself) and symmetric (if s is similar to s' , also s' is similar to s). But then, given the plausible assumption of connectedness (cf. $(\bar{3})$ above), proposition 8.5.5 implies that the only stable—i.e., knowable—sentences are the trivial ones.⁴² To avoid this triviality conclusion about the knowable sentences, we better use the topological semantics. But on this semantics, $\vdash \Box\varphi \rightarrow \Box\Box\varphi$ is true no matter which topology we choose to spell out similarity.⁴³

Now, by proposition 8.3.12, Q and \vdash satisfy principles $(\bar{\Sigma}2)$ and $(\bar{\Sigma}4)$. By corollary 8.4.3, exactly one of $(\bar{\Sigma}1)$ and $(\bar{\Sigma}3)$ holds, i.e., exactly one of the following two options holds:

- (a) The state space is ‘disconnected’ with respect to every possible observation $\varphi \in Q$: Whenever a state s has arbitrarily similar states s' making true φ , then s already makes true φ . Intuitively, it hence is never possible to start in a state s making true φ and ‘continuously’ transform it into a state making false φ : the first state along this transformation making false φ has arbitrarily close φ -states.

⁴²Without the connectedness assumption, knowable sentences are trivial on each connected component. Arguably, this already is bad enough.

⁴³And the reason is intuitive enough: If $\Box\varphi$ is true at a state s , then φ is true at all sufficiently similar states s' . To show that s makes true $\Box\Box\varphi$, we now ‘double’ the demand for being sufficiently similar; call this sufficiently* similar. Then all states s' that are sufficiently* similar to s , make true $\Box\varphi$, because any state s'' that is sufficiently* similar to s' is in total sufficiently similar to s and hence makes true φ .

- (b) There are sentences φ in Q which can be false at some state s but in no state can we come to know that they are false (because they nowhere are stably false).

Since (a) is the negation of $(\bar{3})$ above and (b) is the negation of $(\bar{4})$ above, we have argued, using the impossibility result, that $(\bar{1})$ – $(\bar{4})$ are jointly inconsistent. This inconsistency can be used to obtain various more concrete arguments.

We've already mentioned the 'extension of Fitch's paradox': that $(\bar{1})$ – $(\bar{3})$ imply the failure of $(\bar{4})$. Another example is the following. Assume we take $(\bar{2})$ – $(\bar{4})$ for granted: Regarding $(\bar{2})$, we take it, by definition, that similarity and stability on the state space is given by observability. Regarding $(\bar{3})$ and $(\bar{4})$, we think that to understand an observable property, we not only need to be able to 'imagine' what it takes to fail, there also should be a state where we can observe that it fails. But then we must conclude that $(\bar{1})$ fails: the safety condition for knowledge fails when stability is understood as coming from observability.

Scientific hypotheses The dual argument can be made for good scientific hypotheses. (Since it is dual, we'll be more brief.) Assume again that the similarity structure on the state space is given by a topology that reflects observability (as, e.g., in the 'topology via logic' approach). Let Q be the set of those sentences of our observation language \mathcal{L} that are good scientific hypotheses. Then we have to give up at least one of the following:

- (1) Scientific hypotheses are falsifiable: if false, they can be falsified by some measurement or observation, and hence are stably false.
- (2) Falsifiability also is sufficient to be a scientific hypothesis: if φ is a sentence such that, whenever it is false, it is stably false (so the falsity can be observed), then φ is a scientific hypothesis.
- (3) Non-triviality: for some scientific hypothesis φ , we can continuously transform a $\neg\varphi$ -state into a φ -state.
- (4) Standard model: if a scientific hypothesis is true, there is some state in which we can observe that it is true.

Again, (1)–(2) implies $Q = \{\varphi : \vdash \neg\varphi \rightarrow \Box\neg\varphi\}$, whence the impossibility result implies that (3) or (4) fails.

If we take (1)–(3) for granted, we must conclude that (4) fails: there are good scientific hypotheses that are true in some states, but we can never observe this. Similarly, if we take (1), (3), and (4) for granted, we must conclude that (2) fails: there must be more to being a good scientific hypothesis than merely being observationally falsifiable. Arguments for (1) are given by the Popperian idea that scientific hypotheses are falsifiable; and (2) takes this as their defining feature. Arguments for (3) and (4) may proceed as above: to understand a scientific

hypothesis, we must be able to imagine a situation in which (we can observe that) the hypothesis is true. Another argument for (3) is that the topology should be connected (e.g., because it is similar to the topology of the observed physical phenomenon), so any non-trivial scientific hypothesis will satisfy (3).

8.6.2 A limitation for AI-safety

We discuss AI-safety: the idea that for an AI to be safe, its behavior should be stable under small variations of the input. We consider two ways of concretely spelling out this stability. The first results in triviality and the second comes with limitations.

Let's consider the paradigm case where the AI should perform a desired mapping from a continuous input space to a discrete output space: The AI gets 'continuous' inputs, like images or sound recordings, and should produce a discrete output, like a classification label ('this is a cat') or symbolic representation (the transcribed speech). As discussed in section 8.2.3, for such an AI to be safe—e.g., to avoid adversarial attacks—the judgments of the AI should be stable: if it (correctly) judges an input s to have property φ (e.g., depicting a stop sign), then it still does so on sufficiently similar inputs s' (and, e.g., doesn't suddenly judge there to be a 'Speed Limit 45' sign). Naturally, we want to understand which φ have this property. For simplicity, let's call those that have it the 'safe' judgments.

To spell this out more concretely, let S be the input space (i.e., the inputs are the states). The atomic sentences $p \in \mathcal{L}_0$ correspond to the classification labels of the AI, so the valuation V is given by: $s \models p$ iff the AI judges input s to be of category p . To interpret the stability operator \Box , we need to describe the intended similarity relation on the input space. We'll consider two ways of doing that: first, via a similarity relation (Kripke semantics) and second via an appropriate topology (topological semantics). The first way arguably is the most straightforward one but results in triviality. The second way provides a general structure to avoid the triviality, but each instance of this structure has its own limitations.

Stability via a similarity relation The input space is usually (mathematically modeled as) something like the \mathbb{R}^n : for example, the space of images with n pixels each of which can have a real value describing its color. We'll capture this more generally by assuming that the input space $S = (X, d)$ is a path-connected metric space: The metric allows speaking of closeness between inputs (as, e.g., the usual Euclidean distance between two images). The path-connectedness intuitively means that we can continuously transform each input into another one—e.g., continuously change the pixel values of one image into another.

Maybe the most straightforward way to spell out 'sufficiently similar' on this input space $S = (X, d)$ is to pick a 'safety threshold' $\epsilon > 0$ and call two inputs s and s' sufficiently similar if they are ϵ -close (i.e., $d(s, s') < \epsilon$). For example, ϵ could

be the minimal distance between two pixel images that is detectable to the human eye. Then, for an AI to be safe with respect to that notion of similarity means that if the AI judges input s to have property φ , it also judges all ϵ -close inputs s' to have property φ . In that sense, ϵ is a 'safety threshold' since within ϵ -range the AI cannot be adversarially attacked (i.e., it doesn't retract its judgment φ).

However, this approach trivializes the safe judgments to the trivial ones: We have the Kripke model (S, R, V) where sRs' iff $d(s, s') < \epsilon$. Since (X, d) is path-connected, it follows, by a standard topological argument (which we prove in the appendix: lemma 8.7.2), that (S, R, V) is a connected. Hence proposition 8.5.5 implies that the only stable sentences are the trivial ones.

One might object the use of the underlying (Euclidean) metric d on the input space to provide a notion of similarity: We can have two images that are, to the human eye, rather different even though, in terms of pixel similarity, are rather close (e.g., a tiny, but important detail missing)—and also vice versa (e.g., added noise). In short: low-level (pixel) similarity need not coincide with high-level (human) similarity. Note, however, that the triviality argument is not restricted to the Euclidean metric but works for *any* path-connected metric on the input space—which could also attempt to describe high-level closeness of inputs.

Nonetheless, if spelling out AI-safety via a safety threshold trivializes the safe judgments, we need to rethink the assumptions.

Stability via an appropriate topology In addition to objecting the use of a metric, another objection is to not formalize 'sufficiently similar' binarily via a threshold. Rather—as the qualifier 'sufficiently' suggests—we describe 'sufficiently similar' as a matter of degree: that there is some degree of similarity.

What we're left with, then, is that 'sufficiently similar' is made concrete by *some* topology on the input space. There may be many such topologies, and they don't need to be the natural 'low-level' topology of the input space. For example, the appropriate topology that captures similarity between images according to the high-level human cognition may be—as just indicated—different to the 'low-level' Euclidean topology on \mathbb{R}^n describing the pixel values of an image.

One assumption, though, remains plausible: whatever the appropriate topology on the input space, it should be connected. Intuitively, this means that we're dealing with a single input space which cannot be split up into two or more separate spaces (which we could study separately). Here are two considerations that imply connectedness. First, similarly to (3) and (3) in section 8.6.1, we may think that we should be able to continuously transform each input into another one (by continuously changing the parameters determining the input). This path-connectedness and implies connectedness. Second, we may think that high-level similarity 'supervenies' on low-level similarity: if a set of inputs describes a high-level similarity range (i.e., is an open set of the high-level topology), then it can be described as a union of low-level similarity ranges (i.e., as a union of

opens of the low-level topology). So the high-level topology is coarser than the low-level topology, whence, since the latter is connected, also the former is.

By now, one might object that we gave up too many assumptions to still have something that can reasonably be said to formalize similarity and stability. For one, only demanding that there is *some* degree of similarity within which the judgment of the AI is stable may not be practically useful: it may be that this stability range only ensures stability within a distance of 2^{-1000} in the input space \mathbb{R}^n . To remedy this—i.e., to get a practically meaningful notion of stability—we may resort to a finite topology, or, at least, an Alexandroff topology (i.e., every point has a least open neighborhood).⁴⁴

These considerations indicate that there is a real issue in concretely spelling out the kind of stability required for AI-safety: On the one hand, it needs to avoid triviality, which seems to exclude relational approaches and favors topological ones. On the other hand, it should be meaningful and useful, which puts quite some demands on a topology to be ‘appropriate’—i.e., to reasonably capture the notion of similarity and stability intended in AI-safety. However, besides connectedness, we won’t discuss further assumptions on the topology to be ‘appropriate’. Our point here is that the impossibility result yields some general limitations for *any* connected topology—regardless of whether it really is appropriate or eventually turns out not to be appropriate.

In this general case, then, we’re given a connected topology τ on the input space S (which attempts to provide a notion of similarity capturing AI-safety). Thus, the atomic sentences $p \in \mathcal{L}_0$ are still interpreted by the valuation V and the stability operator \Box is interpreted topologically. Now, what can we say about the set of safe judgments $Q = \{\varphi \in \mathcal{L} : \vdash \varphi \rightarrow \Box\varphi\}$?

As already discussed, for $\varphi \in Q$, we can interpret $\mathbb{W}\varphi = \varphi \vee \Box\neg\varphi$ roughly as ‘the AI cannot be φ -adversarially attacked’: it is not the case that φ is false (e.g., the AI doesn’t judge there to be a speed limit sign) but φ is true on arbitrarily close states (so, on some small alternations of the input, the AI judges there to be a speed limit sign). From the impossibility result (corollary 8.5.8), we know:

For every non-trivial safe sentence φ , the AI can be φ -adversarially attacked ($\mathbb{W}\varphi$ is false on some inputs) but for generic inputs it cannot ($\mathbb{W}\varphi$ is generically true).

This is good and bad news. The good news is that we can avoid the no-go result of the ‘safety threshold approach’: Any non-trivial open set of the topology τ provides, when interpreted as a property φ , a non-trivial safe judgment. Moreover, for any safe judgment, the claim $\mathbb{W}\varphi$ that the AI cannot be φ -adversarially attacked is again a safe judgment and it is generically true: the set of inputs where $\mathbb{W}\varphi$ fails

⁴⁴Here it is interesting to note that, by taking the specialization ordering, Alexandroff spaces are in correspondence with $S4$ -frames, i.e., reflexive and transitive Kripke models without the valuation (Van Benthem and Bezhanishvili 2007, sec. 2.4.1). So this would allow a treatment of stability in AI-safety also with Kripke semantics.

is topologically negligible (closed and nowhere dense).⁴⁵ The bad news is that, if φ is non-trivial, there always are some inputs where the AI can be φ -adversarially attacked. Moreover, the AI can also never (i.e., on no input) safely make the 'meta-judgment' that it currently *can* be φ -adversarially attacked: Otherwise, $\mathbb{W}\varphi$ had to be stably false, so the set where $\mathbb{W}\varphi$ fails wouldn't be topologically negligible. Thus, in a sense, the possibility of an adversarial attack cannot safely be detected.

To summarize, even for safe judgments, adversarial attacks can only generically but not absolutely be excluded, and the AI cannot safely make the meta-statement that it cannot be adversarially attacked.

8.7 Conclusion

We end with a brief summary and some open questions. We saw a wide range of examples involving stability and formally described them using modal logic. We formulated four desirable principles about stability and proved them inconsistent: both proof-theoretically (via a novel interpretation of Fitch's paradox) and semantically (via Kripke semantics and topological semantics). We explored two consequences. First, we extended Fitch's paradox: both the set of verifiable sentences (if true, we can come to know them) and the set of falsifiable sentences (scientific hypotheses) are restricted by impossibilities—which could be interpreted, e.g., as there being false verifiable sentences whose falsity we cannot come to know. Second, modeling the stability aspect of AI-safety faces some fundamental difficulties: the straightforward 'safety threshold' approach is likely to trivialize and the topological approach (no matter how it is spelled out in detail) restricts the safe judgments that we can expect.

Regarding open questions, future work should further analyze these difficulties in modeling AI-safety and explore yet other modeling approaches. One might consider neighborhood semantics (Pacuit 2017) or a domain-theoretic semantics as in (Hornischer 2021) and in this thesis. What consequences does the impossibility result have for the other examples of stability from section 8.2? (We've already sketched several in section 8.5.) What implications does this understanding of stability have for the question of how stability allows a symbolic understanding of non-symbolic computation (chapter 7)? More generally, the nexus of philosophy, AI-safety, and logic seems worth exploring: connections between Fitch's paradox,

⁴⁵Note, though, that for concrete topologies τ , one not only needs to check that it provides the intended notion of the stability figuring in AI-safety. One also needs to check that its notion of topological negligibility matches the intended notion of 'negligible'. In particular, the topological sense of negligible and the probabilistic (or measure-theoretic) sense of negligible can come apart (Oxtoby 1980). This becomes particularly relevant if we want to understand 'generic' probabilistically: as almost surely with respect to some probability measure—which might represent random noise present in capturing and processing (image) data.

safety conditions for knowledge (and epistemology more broadly), AI-safety, modal logic, and topology—coming together under the concept of stability.

Appendix

Two proof-theoretic lemmas First, we prove lemma 8.3.8: For any derivability relation \vdash and choice of Q , we have, for $k = 1, \dots, 4$:

- (a) If principle (Σk) holds for Q , then principle $(\bar{\Sigma}k)$ holds for $\neg Q$, and
- (\bar{a}) If principle $(\bar{\Sigma}k)$ holds for Q , then principle (Σk) holds for $\neg Q$.

PROOF OF LEMMA 8.3.8. Ad (a). For $k = 1$, if $\varphi \in Q$ with $\nabla \varphi \leftrightarrow \Diamond \varphi \leftrightarrow \Box \varphi$, then $\neg \varphi \in \neg Q$ and $\nabla \neg \varphi \leftrightarrow \Diamond \neg \varphi \leftrightarrow \Box \neg \varphi$ (otherwise, by applying \neg to all formulas, $\vdash \varphi \leftrightarrow \Box \varphi \leftrightarrow \Diamond \varphi$).

For $k = 2$, let $\varphi' \in \neg Q$ and show $\vdash \varphi' \rightarrow \Box \varphi'$. Since $\varphi' \in \neg Q$, $\varphi' = \neg \varphi$ for some $\varphi \in Q$. So, by assumption, $\vdash \neg \varphi \rightarrow \Box \neg \varphi$, i.e., $\vdash \varphi' \rightarrow \Box \varphi'$, as needed.

For $k = 3$, let $\varphi' \in \neg Q$ and show $\vdash \neg \varphi' \rightarrow \Diamond \Box \neg \varphi'$. Since $\varphi' \in \neg Q$, $\varphi' = \neg \varphi$ for some $\varphi \in Q$. So, by assumption, $\vdash \varphi \rightarrow \Diamond \Box \varphi$. Since $\vdash \varphi \leftrightarrow \neg \neg \varphi$, we have, by closure under equivalence, $\vdash \neg \neg \varphi \rightarrow \Diamond \Box \neg \neg \varphi$, i.e., $\vdash \neg \varphi' \rightarrow \Diamond \Box \neg \varphi'$, as needed.

For $k = 4$, let $\varphi' \in \neg Q$ and find $\psi' \in \neg Q$ such that $\vdash \psi' \leftrightarrow \mathbb{W} \varphi'$. Since $\varphi' \in \neg Q$, $\varphi' = \neg \varphi$ for some $\varphi \in Q$. So, by assumption, there is $\psi \in Q$ such that $\vdash \psi \leftrightarrow \mathbb{M} \varphi$. Now consider $\psi' := \neg \psi \in \neg Q$. By adding a negation to both sides, \vdash proves $\psi' = \neg \psi \leftrightarrow \neg \mathbb{M} \varphi \leftrightarrow \mathbb{W} \neg \varphi = \mathbb{W} \varphi'$, as needed.

Ad (\bar{a}), we reason analogously. □

Second, we prove lemma 8.3.10: If \vdash is a transitive derivability relation, then, without using transitivity in (b) and (\bar{b}):

- (a) If $\vdash \neg \varphi \rightarrow \Box \neg \varphi$, then $\vdash \neg \mathbb{M} \varphi \rightarrow \Box \neg \mathbb{M} \varphi$.
- (b) $\vdash \mathbb{M} \mathbb{M} \varphi \leftrightarrow \mathbb{M} \varphi$.
- (\bar{a}) If $\vdash \varphi \rightarrow \Box \varphi$, then $\vdash \mathbb{W} \varphi \rightarrow \Box \mathbb{W} \varphi$.
- (\bar{b}) $\vdash \mathbb{W} \mathbb{W} \varphi \leftrightarrow \mathbb{W} \varphi$.

PROOF OF LEMMA 8.3.10. Ad (a). Since $\vdash \neg \mathbb{M} \varphi \leftrightarrow (\neg \varphi \vee \Box \varphi)$, we need to show

$$\vdash (\neg \varphi \vee \Box \varphi) \rightarrow \Box (\neg \varphi \vee \Box \varphi)$$

By assumption, we have $\vdash \neg \varphi \rightarrow \Box \neg \varphi$. And by transitivity, we have $\vdash \Box \varphi \rightarrow \Box \Box \varphi$. Since $\vdash \Box p \rightarrow \Box (p \vee q)$, we have

$$\begin{aligned} &\vdash (\neg \varphi \rightarrow \Box \neg \varphi) \wedge (\Box \varphi \rightarrow \Box \Box \varphi) \\ &\vdash (\Box \varphi \rightarrow \Box \Box \varphi) \wedge (\Box \Box \varphi \rightarrow \Box (\neg \varphi \vee \Box \varphi)). \end{aligned}$$

This implies the claim: Informally, both $\neg\varphi$ and $\Box\varphi$ imply $\Box(\neg\varphi \vee \Box\varphi)$.

Ad (b). Note that in \vdash we can show the following chain of implications

$$\begin{aligned} \mathbf{M}\varphi = \varphi \wedge \neg\Box\varphi &\rightarrow \neg\Box\varphi \rightarrow \Diamond\neg\varphi \rightarrow \Diamond(\neg\varphi \vee \Box\varphi) \\ &\rightarrow \Diamond\neg(\varphi \wedge \neg\Box\varphi) \rightarrow \neg\Box(\varphi \wedge \neg\Box\varphi) = \neg\Box\mathbf{M}\varphi. \end{aligned}$$

Hence we can show in \vdash that $\mathbf{M}\mathbf{M}\varphi = (\mathbf{M}\varphi \wedge \neg\Box\mathbf{M}\varphi) \leftrightarrow \mathbf{M}\varphi$.

Ad (\bar{a}). If $\vdash \varphi \rightarrow \Box\varphi$, then, for $\psi := \neg\varphi$, we have $\vdash \neg\psi \rightarrow \Box\neg\psi$. By (a), $\vdash \neg\mathbf{M}\psi \rightarrow \Box\neg\mathbf{M}\psi$. Since $\vdash \mathbf{W}\varphi \leftrightarrow \neg\mathbf{M}\neg\varphi$, we have $\vdash \mathbf{W}\varphi \rightarrow \Box\mathbf{W}\varphi$.

Ad (\bar{b}). Since $\vdash \mathbf{W}\varphi \leftrightarrow \neg\mathbf{M}\neg\varphi$, we can, by (b), show in \vdash that

$$\mathbf{W}\mathbf{W}\varphi \leftrightarrow \neg\mathbf{M}\mathbf{M}\neg\varphi \leftrightarrow \neg\mathbf{M}\neg\varphi \leftrightarrow \mathbf{W}\varphi,$$

as needed. □

Two topological lemmas For the first lemma, recall that we write \mathcal{M}_n^s for the set of stochastic $n \times n$ matrices with the topology inherited from the space \mathcal{M}_n of all real-valued $n \times n$ matrices with the topology induced by any matrix norm. Let $\mathcal{M}_n^{\text{irr}}$ be the irreducible stochastic matrices.

8.7.1. LEMMA. (a) \mathcal{M}_n^s is a connected topological space.

(b) $\mathcal{M}_n^{\text{irr}}$ is an open subset of \mathcal{M}_n^s .

(c) If $n \geq 2$, $\mathcal{M}_n^{\text{irr}}$ a non-trivial subset of \mathcal{M}_n^s (neither empty nor the whole set).

PROOF. Ad (a). We show that \mathcal{M}_n^s is a convex subset of the real vector space \mathcal{M}_n (which implies that it is connected). For $t \in [0, 1]$, we need to show that $C := tA + (1-t)B$ is a stochastic matrix. It again has non-negative entries and, for $1 \leq i \leq n$, the sum of the i -th row is:

$$\sum_{j=1}^n C(i, j) = \sum_{j=1}^n tA(i, j) + (1-t)B(i, j) = t1 + (1-t)1 = 1.^{46}$$

Ad (b). We'll use the max-norm on \mathcal{M}_n , i.e., $\|A\| := \max_{i,j} |A(i, j)|$. Note

$$\mathcal{M}_n^{\text{irr}} = \bigcap_{1 \leq i, j \leq n} \bigcup_{k \geq 1} N_{i,j}^k \quad \text{with} \quad N_{i,j}^k := \{A \in \mathcal{M}_n^s : A^k(i, j) > 0\}.$$

So it suffices to show that $N_{i,j}^k \subseteq \mathcal{M}_n^s$ is open. We do this by induction on k . For $k = 1$, we show that, for any $A \in N_{i,j}^1$, there is $\epsilon > 0$ such that, for any $B \in \mathcal{M}_n^s$,

⁴⁶This is a common argument. For much more context, see Brualdi (2006), in particular p. 19.

if $\|A - B\| < \epsilon$, then $B \in N_{i,j}^1$. Indeed, choose $\epsilon := \frac{1}{2}A(i, j) > 0$. Then for such B we must have $B(i, j) > 0$, since otherwise $B(i, j) = 0$ and we get the contradiction

$$\epsilon > \max_{i,j} |A(i, j) - B(i, j)| \geq |A(i, j) - B(i, j)| = A(i, j).$$

Now, assume that $N_{i,j}^k$ is open and show that $N_{i,j}^{k+1}$ is open. Indeed, we have

$$\begin{aligned} A \in N_{i,j}^{k+1} &\Leftrightarrow 0 < A^{k+1}(i, j) = A^k A(i, j) = \sum_{l=1}^n A^k(i, l)A(l, j) \\ &\Leftrightarrow \exists l \in \{1, \dots, n\} : A^k(i, l) > 0 \text{ and } A(l, j) > 0 \\ &\Leftrightarrow \exists l \in \{1, \dots, n\} : A \in N_{i,l}^k \cap N_{l,j}^1. \end{aligned}$$

Hence $N_{i,j}^{k+1} = \bigcup_{l=1}^n (N_{i,l}^k \cap N_{l,j}^1)$, whence, by the induction hypothesis and the induction base, $N_{i,j}^{k+1}$ is open.

Ad (c). The $n \times n$ matrix where every entry is $\frac{1}{n}$ is in $\mathcal{M}_n^{\text{irr}}$. And, if $n \geq 2$, the identity matrix I is in \mathcal{M}_n^{s} but not in $\mathcal{M}_n^{\text{irr}}$ (for $i := 1$ and $j := 2$, we have, for any $k \geq 1$, that $I^k(i, j) = I(i, j) = 0$). \square

The second lemma is that, in a path-connected metric space, any two points can also be ‘connected’ by a finite sequence of states where the distance between neighboring elements can be chosen arbitrarily small.

8.7.2. LEMMA. *Let (X, d) be a path-connected metric space and $\epsilon > 0$. For all $x, y \in X$, there is $x = x_0, x_1, \dots, x_n = y$ in X such that $d(x_i, x_{i+1}) < \epsilon$ for all $i = 0, \dots, n - 1$.*

PROOF. Since X is path connected, there is a continuous $f : [0, 1] \rightarrow X$ with $f(0) = x$ and $f(1) = y$. Since f is a continuous function from a compact metric space into a metric space, it is uniformly continuous (Heine–Cantor theorem). So there is $\delta > 0$ such that, for all $t, t' \in [0, 1]$, if $|t - t'| < \delta$, then $d(f(t), f(t')) < \epsilon$. Let $n \geq 1$ be big enough such that $\frac{1}{n} < \delta$. For $i = 0, \dots, n$, set $t_i := \frac{i}{n}$ and $x_i := f(t_i)$. Then $x_0 = f(t_0) = f(0) = x$ and $x_n = f(t_n) = f(1) = y$ and, for $i \in \{0, \dots, n - 1\}$, $d(x_i, x_{i+1}) = d(f(t_i), f(t_{i+1})) < \epsilon$ since $|t_i - t_{i+1}| = |\frac{i+1}{n} - \frac{i}{n}| = \frac{1}{n} < \delta$. \square

We end with a brief summary of the main formal and conceptual results, and we mention some of the central open questions. The numbers indicate the corresponding chapters of the listed items.

Results

1. Just like symbolic computation, also non-symbolic computation needs a (denotational) semantics to achieve ‘(structural) understanding’. Domain theory provided a denotational semantics to symbolic computation. This suggests extending it to non-symbolic computation.
1. Albeit known, it deserves more appreciation that symbolic and non-symbolic computation can be described in a unified framework as state-discrete and state-continuous dynamical systems, respectively. This turns the task into providing a domain-theoretic semantics to dynamical systems.
2. The symbolic or state-discrete case is given by labeled transition systems (LTS). A more detailed description of their behavior is provided by the notion of a behavioral transition system (BTS) which we motivated and axiomatized.
2. The trajectory domain construction provides a denotational description of the behavior of a BTS.
3. In fact, it constitutes a functor. This is a form of compositionality of the semantics.
3. After appropriate formulation, the trajectory domain construction can be extended to an adjunction showing that the computational model of ω -algebraic domains can be abstracted from the computational model of BTSs.

3. The trajectory domain construction also leads to a new interpretation of relevance logic in terms of LTSs.
4. Turning to non-continuous computation, we constructed, for every (typically state-continuous) dynamical system, the observation domain. It provides a denotational description of the system's behavior in terms of possible observations of the system.
4. The observation domain is a computational model of the dynamical system in the sense that its compact elements (representing finite observations) approximate the maximal elements which, in turn, are isomorphic to the original system.
4. The observation domain is in the category of dynamical domains which we defined in a purely domain-theoretic way via certain limits of certain finite domains.
5. The observation domain construction constitutes a functor. Again, this is a form of compositionality of the semantics.
5. It is adjoint to the functor which maps a dynamical domain to the dynamical system that it models. This adjunction even is an equivalence of categories after adding the simple domain-theoretic property of being max-reflective.
5. The category of (measure-theoretic) dynamical systems is a localization of the category of not necessarily compact topological systems. These systems, in turn, can be compactified in two equivalent ways: logically (via Stone duality) or computationally (via the observation domain).
5. This establishes a precise correspondence between dynamical systems and dynamical domains.
6. The important system-theoretic concepts of metric and topological entropy have domain-theoretic counterparts: domain-entropy and max-entropy, respectively.
7. We suggested the thesis that non-symbolic computation is profinite symbolic computation.
7. Conversely, having a fairly stable behavior seems crucial for a dynamical system (or non-symbolic computation) to have a symbolic approximation. Ergodicity and (algorithmic) randomness help to use and achieve this stability.
8. However, the general concept of stability is tricky: a reinterpretation of Fitch's paradox shows that it cannot jointly have four desirable properties.

8. In particular, this poses fundamental difficulties to modeling and specifying AI-safety.
8. Established philosophical tools (mostly from epistemology) promise fruitful applications to foundational issues of modern AI.

Questions

- 2,3. Explore the nexus of BTSs, the generalization of Scott information systems, substructural logics, relevance logic, game semantics, and full abstraction.
2. Investigate the links to Leavitt path algebras and Bratteli–Vershik diagrams.
3. Further develop the categorical properties of BTSs and trajectory domains: extend the adjunction and establish (the preservation of) categorical constructions.
4. Develop a unified computability theory for symbolic and non-symbolic computation via effective enumerations of the bases of the semantic domains (trajectory domains and observation domains).
4. Extend observations to real-valued observations and use the tools of the operator-theoretic approach to ergodic theory—in particular, in the case of learning dynamics.
- 4,5. Further develop the domain theory: domain constructions, order-theoretic characterizations of dynamical domains, universal dynamical domains, and how they provide a ‘type structure’ for non-symbolic computation.
5. Characterize the domain-theoretic counterpart of metric isomorphism and explore whether this (and the localization) is the start of a ‘homotopy structure’.
6. Use the domain-theoretic description of metric and topological entropy to find new formulations of known results (e.g., variational principle, isomorphism invariant).
7. Explore in detail the role of (algorithmic) randomness and the (analogue of the) ergodic hypothesis for non-symbolic computation and its symbolic realization.
8. What does the no-go result about stability mean for symbolically describing stable enough dynamical systems?
8. Further investigate specifying the stability aspect of AI-safety in light of the foundational limitations.

Appendix A

Systems as a category of fractions

As motivated in chapter 5 (sections 5.1 and 5.3.1), we show that the category of (measure-theoretic) dynamical systems \mathbf{DS} can be regarded as a localization—or, more precisely, as a category of fractions—of the category of topological systems \mathbf{bTS}_0 :

$$\mathbf{DS} \xleftarrow{\text{Loc}} \mathbf{bTS}_0$$

and similarly for the standard case. The idea was this: Localizing a category means keeping the same objects but adding inverses to a collection of morphisms. On the object level, the two categories \mathbf{DS} and \mathbf{bTS}_0 are ‘essentially the same’ since dynamical systems can be realized as topological ones. On the morphism level, however, morphisms in \mathbf{bTS}_0 that are injective on an invariant set of full measure become isomorphisms in \mathbf{DS} . So it stands to reason that \mathbf{DS} is obtained by turning precisely those morphisms into isomorphisms.

A.1 Statement of the theorem

To state the result formally, let’s start by giving a name to those morphisms φ that have an invariant domain of injectivity with full measure. In the spirit of zero-dimensional topology, we’ll additionally demand that their domain of injectivity is a clopen set.

A.1.1. DEFINITION. A morphism $\varphi : (X, \tau, \mu, T) \rightarrow (Y, \sigma, \nu, S)$ between zero-dimensional measured topological systems is called *injective clop 0* if there is a clopen set $A \subseteq X$ such that $\mu(A) = 1$ and $T(A) \subseteq A$, and, for all $x, x' \in A$, if $x \neq x'$, then $\varphi(x) \neq \varphi(x')$.¹ A morphism $\varphi : (\mathfrak{X}, \mathcal{B}_X) \rightarrow (\mathfrak{Y}, \mathcal{B}_Y)$ in \mathbf{bTS}_0 is injective clop 0 if $\varphi : \mathfrak{X} \rightarrow \mathfrak{Y}$ is injective clop 0.

¹The name is mnemonic for ‘injective mod 0’—i.e., injective on a (T -invariant) domain of full measure—with the additional demand that this domain also is clopen.

$$\begin{array}{ccc}
 V & \xrightarrow{g} & Z \\
 \downarrow w & & \downarrow v \\
 X & \xrightarrow{f} & Y
 \end{array}
 \qquad
 \begin{array}{ccccc}
 V & \xrightarrow{w} & X & \xrightarrow[f]{g} & Y & \xrightarrow{v} & Z
 \end{array}$$

(a) Right Ore condition (b) Right cancellation condition

Figure A.1: Calculus of fractions.

Next, we need to define the localization of a category or, more specifically, the concept of a calculus of fractions. The classic reference is Gabriel and Zisman (1967), and modern treatments—which we use for the presentation here—are found in Borceux (1994, ch. 5), Kashiwara and Schapira (2006, ch. 7), Fritz (2011), Yekutieli (2020, ch. 6), and in the *nlab* entries ‘localization’ and ‘calculus of fractions’.²

As mentioned, the idea of a localization is that, given a category \mathbf{C} and a collection of morphisms W , we find a new category $\mathbf{C}[W^{-1}]$ such that the morphisms in W are isomorphisms in $\mathbf{C}[W^{-1}]$, i.e., have inverses. (The name comes from the localization of a ring where one, roughly, also adds inverses to a set of elements of the original ring.) More precisely, there is a functor $\mathbf{Q} : \mathbf{C} \rightarrow \mathbf{C}[W^{-1}]$ that sends morphisms from W to isomorphisms in $\mathbf{C}[W^{-1}]$ and the pair $(\mathbf{C}[W^{-1}], \mathbf{Q})$ is ‘universal’ with this property. One can define this notion of localization formally and show that such localizations exist, but, at that generality, these often are too abstract to be useful. A more concrete description of a localization in this general sense is available if W admits a *calculus of fractions*. Since this is all we need here, we only define this special case and not the general one.

A.1.2. DEFINITION. Let \mathbf{C} be a category and W a collection of morphisms in \mathbf{C} . We say W admits a *calculus of fractions* if

1. for all X in \mathbf{C} , $\text{id}_X \in W$.
2. if $v : X \rightarrow Y$ and $w : Y \rightarrow Z$ are in W , then $w \circ v$ is in W .
3. As in figure A.1a, if $f : X \rightarrow Y$ is in \mathbf{C} and $v : Z \rightarrow Y$ in W , then there is $g : V \rightarrow Z$ in \mathbf{C} and $w : V \rightarrow X$ in W such that $f \circ w = v \circ g$.
4. As in figure A.1b, if $f, g : X \rightarrow Y$ are in \mathbf{C} and $v : Y \rightarrow Z$ is in W with $v \circ f = v \circ g$, then there is $w : V \rightarrow X$ in W such that $f \circ w = g \circ w$.

Then one defines the *category of fractions* $\mathbf{C}[W^{-1}]$ by having the same objects as \mathbf{C} and a morphism $X \rightarrow Y$ is an equivalence class $[v, Z, f]$ of a span $X \xleftarrow{v} Z \xrightarrow{f} Y$ with v in W and f in \mathbf{C} , where the equivalence of two such spans is defined as:

²See <https://ncatlab.org/nlab/show/localization> and <https://ncatlab.org/nlab/show/calculus+of+fractions>, respectively (accessed 9 Feb 2021).




Figure A.2: Equivalence of spans.




Figure A.3: Composition of spans.


$X \xleftarrow{v} Z \xrightarrow{f} Y$ and $X \xleftarrow{v'} Z' \xrightarrow{f'} Y$ are equivalent iff, as in figure A.2, there is an object V and morphisms $g : V \rightarrow Z$ and $g' : V \rightarrow Z'$ in \mathbf{C} such that $f \circ g = f' \circ g'$ and $v \circ g = v' \circ g' \in W$.

Composition of these equivalence classes is, as shown in figure A.3, done using (3), and the identity equivalence class is $[\text{id}_X, X, \text{id}_X]$. The localization functor $\mathbf{Q} : \mathbf{C} \rightarrow \mathbf{C}[W^{-1}]$ is the identity on objects and maps $f : X \rightarrow Y$ to $[\text{id}_X, X, f]$.

Comments: First, to be more precise, we should say that W admits a calculus of *right* fractions, because we could dually define that W admits a calculus of *left* fractions by reversing all arrows. (Though, as the *nlab* entry notes, there is no agreement in the literature which of these two definitions is ‘left’ and which is ‘right’.) In any case, we’ll only need the above.

Second, note that (1) and (2) are essentially equivalent to saying that W is a wide subcategory of \mathbf{C} (i.e., a subcategory with the same objects).

Third, the intuition behind using spans is this: Assume we have a span $X \xleftarrow{v} Z \xrightarrow{f} Y$ in the category \mathbf{C} with $v \in W$. If we aim to build a category $\mathbf{D} = \mathbf{C}[W^{-1}]$ with the same objects as \mathbf{C} but with an inverse for every element of W , then \mathbf{D} has to have the ‘new’ morphism $X \xrightarrow{v^{-1}} Z \xrightarrow{f} Y$ induced by the span. Moreover, we can also regard any ‘old’ morphism $f : X \rightarrow Y$ in \mathbf{C} as the trivial span $X \xleftarrow{\text{id}_X} X \xrightarrow{f} Y$. This suggests to simply take the new morphisms to be spans of the old morphisms with one leg in W . This goes a long way, but we still rely on a specific choice of C in the composition of spans (figure A.3). A clever observation is that, by moving to the equivalence classes, this can be avoided by



(a) $[v, X, id_X] \circ [id_X, X, v] = [id_X, X, id_X] = id_X$ (b) $[id_X, X, v] \circ [v, X, id_X] = [v, X, v] = [id_Y, Y, id_Y] = id_Y$

Figure A.4: For $v : X \rightarrow Y$ in W , the morphism $Q(v) = [id_X, X, v] : X \rightarrow Y$ has the inverse $Q(v)^{-1} = [v, X, id_X] : Y \rightarrow X$.

making any such choice equivalent. Formally: the defined equivalence of spans is indeed an equivalence relation and the definition of composition is independent of the representative. (For a proof see, e.g., Yekutieli (2020, ch. 6).)

Fourth, and further, one can show that $C[W^{-1}]$ does indeed form a category that, together with the localization functor Q , is a localization in the general sense: Yekutieli (2020, ch. 6) does this by showing that (i) as just mentioned, the above construction of $C[W^{-1}]$ is a well-defined category, (ii) Q is a well-defined functor, and (iii) the pair $(C[W^{-1}], Q)$ is a right Ore localization of C with respect to W , which is a special case of the general sense of a localization.

Fifth, another advantage of moving to equivalence classes, is that, if $v : X \rightarrow Y$ is in W , then $Q(v) = [id_X, X, v] : X \rightarrow Y$ does indeed have an inverse in $C[W^{-1}]$, namely $Q(v)^{-1} = [v, X, id_X] : Y \rightarrow X$, as shown in figure A.4.

Sixth, the dual concept of spans is used—in enriched form, as either structured cospans or decorated cospans—in the modeling of open systems (see, e.g., Baez, Courser, and Vasilakopoulou (2021) and the references therein).

Now, we can state the informally promised results formally. We write $F : \mathbf{bTS}_0 \rightarrow \mathbf{TS}_0$ for the forgetful functor sending $(\mathfrak{X}, \mathcal{B})$ to \mathfrak{X} and which is the identity on morphisms.

A.1.3. THEOREM. *Let W be the class of morphisms in \mathbf{bTS}_0 (resp., in \mathbf{bTS}_{0s}) that are injective clop 0. Then W admits a right calculus of fractions and the category of fractions $\mathbf{bTS}_0[W^{-1}]$ (resp., $\mathbf{bTS}_{0s}[W^{-1}]$) is equivalent to \mathbf{DS} (resp., \mathbf{sDS}) via a functor that is like $J \circ F$ (resp. $\bar{J} \circ F$) on objects.*

Comments: First, here is how this formalizes the informal statement: On the object level, this shows that every dynamical system \mathfrak{X} in \mathbf{DS} is isomorphic to

an object $J \circ F(\mathfrak{X}', \mathcal{B}') = J(\mathfrak{X}')$ for an object $(\mathfrak{X}', \mathcal{B}')$ of the localization and hence of \mathbf{bTS}_0 . Thus, the topological system \mathfrak{X}' is the topological realization of the dynamical system \mathfrak{X} . On the morphism level, this shows that once we turn into isomorphisms the morphisms in \mathbf{bTS}_0 that are injective on a (clopen) invariant set of full measure—i.e., once we form the localization $\mathbf{bTS}_0[W^{-1}]$ —, we essentially obtain the morphisms of \mathbf{DS} .

Second, the reason that, on the side of topological systems, we are working in \mathbf{bTS}_0 rather than the simpler \mathbf{TS}_0 is that, for the domain construction, we need to work with \mathbf{bTS}_0 . However, conceptually this is not too much of a difference.

Third, taking equivalence classes mod 0 of structure-preserving partial functions between dynamical systems (when defining system morphisms) now corresponds precisely to taking equivalence classes of spans (when constructing the localization).

A.2 Topological realizations of systems

A measured topological dynamical system \mathfrak{X} is a *topological realization* (or *topological model*) of a (measure-theoretic) dynamical system \mathfrak{Y} if \mathfrak{Y} is isomorphic to $J(\mathfrak{X})$. Similarly, if both \mathfrak{X} and \mathfrak{Y} are standard, we require that \mathfrak{Y} is isomorphic to $J(\mathfrak{X})$. (For this concept, in the more specific setting of ergodic dynamics and with different notation, see Petersen (1983, sec. 4.4) or Glasner and Weiss (2006, sec. 8).)

As discussed in chapter 4, the main result on topological realization in ergodic dynamics is the Jewett–Krieger theorem every ergodic measure-preserving transformation on a Lebesgue space has a topological realization as a minimal, uniquely ergodic homeomorphism on a compact zero-dimensional metric space (see, e.g., Petersen 1983, sec. 4.4).

Our results from chapter 4 provide a topological realization result in our more general setting: Every (standard) dynamical system has a topological realization as a (standard) compact zero-dimensional measured topological system.

In the standard case, this is precisely what we’ll need in the remainder, but, in the general case, we also need the option to make (cl)open any countable collection of Borel sets in the topological realization (and we don’t need compactness). So we state and prove this as a result on its own.

A.2.1. PROPOSITION. *Let $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ be in \mathbf{DS} and $\mathcal{B} \subseteq \mathcal{A}$ a countable subset. Then there is a zero-dimensional Polish topology τ on X such that (a) $\mathcal{A} = \mathcal{B}(\tau)$, (b) $\mathcal{B} \subseteq \text{Clp}(\tau)$, and (c) $T : X \rightarrow X$ is τ -continuous. In particular, $\mathfrak{X}' = (X, \tau, \mu, T)$ is in \mathbf{TS}_0 . Moreover, if T is injective, we can choose τ such that $T : X \rightarrow X$ is τ -open.*

Comments: First, if we only consider the state space, i.e., the standard Borel space (X, \mathcal{A}) , such a topological realization simply means that (by definition) there is a Polish topology τ on X such that $\mathcal{A} = \mathcal{B}(\tau)$.

Second, if we also take the dynamics T into account, the more general context of results of this kind is the intersection of descriptive set theory and dynamical systems (Becker and KeCHRIS 1996; Foreman et al. 2000). There, one finds results on topologically realizing (universal) Borel actions of a Polish group on a standard Borel space. However, in our setting, we only have a semi-group since T is not assumed to be invertible, and we also have measures. (On the other hand, we only look at the significantly simpler case of a countable semi-group.) To avoid both these complications and additional ‘group action’ terminology, we provide a rather elementary proof.

Third, as mentioned, the Jewett–Krieger theorem is not applicable as we, in general, are far from standard—let alone ergodic. (And we also don’t need that strong of a conclusion.)

Fourth, this doesn’t just provide a topological realization ‘up to isomorphism’ but up to identity.

PROOF. Since (X, \mathcal{A}) is a standard Borel space, there is, by (an equivalent) definition (Kechris 1995, def. 12.5), a Polish topology on τ on X with $\mathcal{A} = \mathcal{B}(\tau)$. We construct a finer zero-dimensional Polish topology τ_B on X with the required properties as a ‘limit’ of a chain of Polish topologies (τ_n) . (So τ_B will be the required τ from the proposition.)

The key tool is the following fundamental fact about Polish spaces (Kechris 1995, ex. 13.5):

- (*) If (X, τ) is a Polish space and $\mathcal{B} \subseteq \mathcal{B}(\tau)$ countable, there is a zero-dimensional Polish topology τ' on X such that $\tau \subseteq \tau'$, $\mathcal{B}(\tau) = \mathcal{B}(\tau')$, and $\mathcal{B} \subseteq \text{Clp}(\tau')$.

Since our (X, τ) is Polish with $\mathcal{B} \subseteq \mathcal{A} = \mathcal{B}(\tau)$ countable, we use (*) to find a zero-dimensional Polish topology τ_0 on X with $\tau \subseteq \tau_0$, $\mathcal{B}(\tau) = \mathcal{B}(\tau_0)$, and $\mathcal{B} \subseteq \text{Clp}(\tau_0)$. Note that, since $\mathcal{B}(\tau_0) = \mathcal{A}$, the function $T : (X, \tau_0) \rightarrow (X, \tau_0)$ is Borel-measurable.

Given a Polish topology τ_n on X such that $T : (X, \tau_n) \rightarrow (X, \tau_n)$ is Borel-measurable, we use (*) to construct τ_{n+1} like this (cf. Kechris 1995, thm. 13.11): Let $(B_k)_{k \geq 0}$ be a countable basis for (X, τ_n) and set $\mathcal{C} := \{T^{-1}(B_k) : k \geq 0\}$ which is a countable subset of $\mathcal{B}(\tau_n)$. [If T also is injective, it is a Borel-measurable injection $(X, \mathcal{B}(\tau_n)) \rightarrow (X, \mathcal{B}(\tau_n))$ between standard Borel spaces, so preserves Borel-measurability, whence we can add the sets $T(B_k)$ to \mathcal{C} .] By (*), there is a zero-dimensional Polish topology τ_{n+1} such that (i) $\tau_n \subseteq \tau_{n+1}$, (ii) $\mathcal{B}(\tau_n) = \mathcal{B}(\tau_{n+1})$, (iii) $T : (X, \tau_{n+1}) \rightarrow (X, \tau_n)$ is continuous (the T -preimage of the basic open B_k is $T^{-1}(B_k) \in \tau_{n+1}$), and (iv) $T : (X, \tau_{n+1}) \rightarrow (X, \tau_{n+1})$ is measurable (since $\mathcal{B}(\tau_{n+1}) = \mathcal{B}(\tau_n)$). [(v) if T also is injective, $T : (X, \tau_n) \rightarrow (X, \tau_{n+1})$ is open (the T -image of the basic open B_k is $T(B_k) \in \tau_{n+1}$).]

Now consider the sequence $(\tau_n)_{n \geq 0}$ of refining zero-dimensional Polish spaces. We have, for any $x \neq y$ in X , that there are, since (X, τ) is Hausdorff, two disjoint open sets $U, V \in \tau \subseteq \tau_0 \subseteq \bigcap_{n \geq 0} \tau_n$ such that $x \in U$ and $y \in V$. It

follows (Srivastava 1998, Obs. 2, p. 93) that the topology τ_B generated by $\bigcup_{n \geq 0} \tau_n$ is Polish.

So it remains to show that (a) τ_B is zero-dimensional, (b) $\mathcal{A} = \mathcal{B}(\tau_B)$, (c) $\mathcal{B} \subseteq \text{Clp}(\tau_B)$, and (d) $T : X \rightarrow X$ is τ_B -continuous [and, if T additionally is injective, then (e) T is τ_B -open.]

Concerning (a), We claim that $S := \bigcup_{n \geq 0} \text{Clp}(\tau_n)$ is a subbase of clopens for τ_B . Since clopens are closed under finite intersection, this implies that τ_B has a basis of clopens and hence is zero-dimensional. If U is clopen in τ_n , then $U, U^c \in \tau_n \subseteq \tau_B$, so U is clopen in τ_B , whence $S \subseteq \text{Clp}(\tau_B)$, and it remains to show that S is a subbase, i.e., writing τ' for the topology generated by S , we need to show $\tau' = \tau_B$. Since $S \subseteq \text{Clp}(\tau_B) \subseteq \tau_B$, we have $\tau' \subseteq \tau_B$. For the other direction, it suffices to show that $\tau' \supseteq \bigcup_{n \geq 0} \tau_n$ (since τ_B is generated by the latter). Indeed, if $U \in \tau_n$ (for some $n \geq 0$), then, since τ_n is zero-dimensional, $\text{Clp}(\tau_n)$ is a basis for τ_n , so U can be written as union of elements from $\text{Clp}(\tau_n) \subseteq S$, whence $U \in \tau'$.

Concerning (b), since $\mathcal{A} = \mathcal{B}(\tau) = \mathcal{B}(\tau_0)$, it is enough to show $\mathcal{B}(\tau_0) = \mathcal{B}(\tau_B)$. Since $\tau_0 \subseteq \tau_B$, we have \subseteq . For the other direction, note that $S \subseteq \mathcal{B}(\tau_0)$: if $U \in \text{Clp}(\tau_n)$ for some $n \geq 0$, then $U \in \mathcal{B}(\tau_n) = \mathcal{B}(\tau_0)$. So $\sigma(S) \subseteq \mathcal{B}(\tau_0)$ (where $\sigma(S)$ is the smallest σ -algebra containing S). As a general fact, the σ -algebra generated by a subbase of a second-countable space is the Borel σ -algebra of the space.³ So $\mathcal{B}(\tau_B) = \sigma(S) \subseteq \mathcal{B}(\tau_0)$.

Concerning (c), we have $\mathcal{B} \subseteq \text{Clp}(\tau_0) \subseteq S \subseteq \text{Clp}(\tau_B)$.

Concerning (d), since $\mathcal{D} := \bigcup_{n \geq 0} \tau_n$ is a subbase for τ_B , it suffices to show that, for $U \in \mathcal{D}$, we have $T^{-1}(U) \in \tau_B$. Indeed, given such U , we have $U \in \tau_n$ for some n , so, since $T : (X, \tau_{n+1}) \rightarrow (X, \tau_n)$ is continuous, $T^{-1}(U) \in \tau_{n+1} \subseteq \tau_B$, as needed.

[Concerning (e), the set of finite intersections of sets from \mathcal{D} forms a basis for τ_B . So it suffices that for $U = U_1 \cap \dots \cap U_m$ with $U_i \in \mathcal{D}$ (for $i = 1, \dots, m$) we have $T(U) \in \tau_B$. Since T is injective, we have $T(U) = T(U_1) \cap \dots \cap T(U_m)$.⁴ And, for each U_i , we have $U_i \in \tau_n$ for some n , so, since $T : (X, \tau_n) \rightarrow (X, \tau_{n+1})$ is open, $T(U_i) \in \tau_{n+1} \subseteq \tau_B$. Hence $T(U) = T(U_1) \cap \dots \cap T(U_m) \in \tau_B$.] \square

We can also state a similar proposition in the setting of TS_0 :

A.2.2. COROLLARY. *Let $\mathfrak{X} = (X, \tau, \mu, T)$ be in TS_0 and $\mathcal{E} \subseteq \mathcal{B}(\tau)$ countable. Then there is a zero-dimensional Polish topology $\tau' \supseteq \tau$ on X such that (X, τ', μ, T) is in TS_0 and $\mathcal{E} \subseteq \text{Clp}(\tau')$; in particular, T is τ' -continuous and $\mathcal{B}(\tau') = \mathcal{B}(\tau)$.⁵*

³Proof: Let S be a subbasis for the second-countable space (X, τ) . Since $S \subseteq \tau$, we have $\sigma(S) \subseteq \mathcal{B}(\tau)$. For the other direction, it suffices to show $\tau \subseteq \sigma(S)$. Let $U \in \tau$. So $U = \bigcup_{i \in I} U_i$ where U_i are finite intersections of elements from S , whence $U_i \in \sigma(S)$. Since τ is second-countable, we can assume I to be countable. So $U \in \sigma(S)$ qua countable union of elements from $\sigma(S)$.

⁴If T is an injective function and U, V sets, then $T(U \cap V) = T(U) \cap T(V)$: The \subseteq -inclusion is trivial, and for the other assume $y \in T(U) \cap T(V)$. Then $y = T(x)$ for $x \in U$ and $y = T(x')$ for $x' \in V$. So $T(x) = T(x')$, whence, by injectivity, $x = x'$. So $x \in U \cap V$ and $y = T(x) \in T(U \cap V)$.

⁵The latter is because μ is a measure on $\mathcal{B}(\tau)$ and must be a measure on $\mathcal{B}(\tau')$.

PROOF. Let \mathcal{B}_X be a countable basis for (X, τ) . So $\mathcal{B}' := \mathcal{B}_X \cup \mathcal{E} \subseteq \mathcal{B}(\tau)$ is countable. Since $(X, \mathcal{B}(\tau), \mu, T)$ is in \mathbf{DS} , we can apply proposition A.2.1 to obtain a zero-dimensional Polish topology τ' on X such that (a) $\mathcal{B}(\tau) = \mathcal{B}(\tau')$, (b) $\mathcal{B}' \subseteq \text{Clp}(\tau')$, and (c) $T : X \rightarrow X$ is τ' -continuous. By (b), $\mathcal{B}_X \subseteq \tau'$, so τ' is a finer topology than τ . By (a), $\mathcal{B}(\tau) = \mathcal{B}(\tau')$, so μ is a probability measure on $\mathcal{B}(\tau')$. By (c), T is τ' -continuous. Hence (X, τ', μ, T) is in \mathbf{TS}_0 . And, by (b), $\mathcal{E} \subseteq \mathcal{B}' \subseteq \text{Clp}(\tau')$. \square

As mentioned, in the standard case, the topological realization that we get from chapter 4 will be enough. We also only aim for realization up to isomorphism and not up to identity anymore: the underlying probability space of \mathfrak{X} in \mathbf{sDS} is isomorphic only mod 0 (but, in general, not strictly) to the completion of a Polish space with a probability measure.

A.2.3. PROPOSITION. *Let \mathfrak{X} be in \mathbf{sDS} . Then \mathfrak{X} is isomorphic to $\bar{\mathbf{J}}(\mathfrak{Y})$ for some \mathfrak{Y} in \mathbf{TS}_0 .*

PROOF. Corollary from chapter 4. \square

A.3 The key lemma

The key lemma is the following one on topologically realizing system homomorphisms. It is visualized in figure A.5 and should be familiar from the right Ore condition (figure A.1a). Again, we first deal with the general case and then with the standard case.

A.3.1. LEMMA (general). *Let $(\mathfrak{X}, \mathcal{B}), (\mathfrak{Y}, \mathcal{C}), (\mathfrak{Z}, \mathcal{D})$ be in \mathbf{bTS}_0 . Let A, B, C be invariant and of full measure in $\mathfrak{X}, \mathfrak{Y}, \mathfrak{Z}$, respectively. Let $v : (\mathfrak{Z}, \mathcal{D}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ be in W . Let $\varphi : \mathbf{J}(\mathfrak{X}) \rightarrow \mathbf{J}(\mathfrak{Y})$ be in \mathbf{DS} . Then there is $\psi : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 and $w : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ in W such that $\varphi \circ w = v \circ \psi$ and, writing V for the state space of \mathfrak{Y} , $\psi(V) \subseteq C$, $w(V) \subseteq A$, and $\varphi \circ w(V) = v \circ \psi(V) \subseteq B$. Moreover, if $\varphi : X \rightarrow Y$ is injective on an invariant set of full measure, then ψ can be chosen in W .*

PROOF. Write $\mathfrak{X} = (X, \tau, \mu, T)$, $\mathfrak{Y} = (Y, \sigma, \nu, S)$, and $\mathfrak{Z} = (Z, \rho, \lambda, R)$. Since $v \in W$, let C' be clopen in \mathfrak{Z} with $\lambda(C') = 1$, $R(C') \subseteq C'$ and v injective on C' . So C and C' are invariant sets of full measure, whence $\bar{C} := C \cap C'$ is, too. Since $\varphi \in \mathbf{DS}$, it is a partial function with domain $M \subseteq X$ such that $M \in \mathcal{B}(\tau)$, $T(M) \subseteq M$, and $\mu(M) = 1$. [If φ is injective on an invariant set of full measure, choose M as that set.]

We first construct \mathfrak{Y} . Note that $v : Z \rightarrow Y$ is a measurable function between two standard Borel spaces, $\bar{C} \subseteq Z$ is Borel, and $v \upharpoonright \bar{C}$ is injective. So $v(\bar{C} \cap D) \subseteq Y$

$$\begin{array}{ccc}
 (\mathfrak{Y}, \mathcal{E}) & \xrightarrow{\psi} & (\mathfrak{Z}, \mathcal{D}) & C \\
 \downarrow w & & \downarrow v & \\
 A & (\mathfrak{X}, \mathcal{B}) \xrightarrow{\varphi} & (\mathfrak{Y}, \mathcal{C}) & B
 \end{array}$$

Figure A.5: Visualization of lemma A.3.1. The quotation marks around φ indicate that, in the lemma, φ is only assumed to be a morphism $J(\mathfrak{X}) \rightarrow J(\mathfrak{Y})$.

is Borel for every $D \in \mathcal{D}$ (Kechris 1995, cor. 15.2). Since $\varphi : X \rightarrow Y$ is measurable (qua morphism in \mathbf{DS}), $\varphi^{-1}(v(\overline{C} \cap D)) \in \mathcal{B}(\tau)$ for every $D \in \mathcal{D}$. In particular, $A' := \varphi^{-1}(v(\overline{C})) \in \mathcal{B}(\tau)$.

Now, define $\mathcal{E}_0 := \mathcal{B} \cup \{\varphi^{-1}(v(\overline{C} \cap D)) : D \in \mathcal{D}\} \cup \{\varphi^{-1}(B), A, A'\}$. Hence $\mathcal{E}_0 \subseteq \mathcal{B}(\tau)$ is countable. Since (X, τ, μ, T) is in \mathbf{TS}_0 , there is, by corollary A.2.2, a zero-dimensional Polish topology $\tau' \supseteq \tau$ on X such that (a) $\mathcal{B}(\tau) = \mathcal{B}(\tau')$, (b) $\mathcal{E}_0 \subseteq \mathbf{Clp}(\tau')$, (c) $T : X \rightarrow X$ is τ' -continuous. Let \mathcal{E} be a countable clopen basis for τ' that includes \mathcal{E}_0 and is closed under finite intersection.⁶

Now, define $V := A' \cap \varphi^{-1}(B) \cap A$. Then we have

- $V \subseteq X$ is τ' -clopen: It is a finite intersection of elements in $\mathcal{E}_0 \subseteq \mathbf{Clp}(\tau')$ and hence in $\mathbf{Clp}(\tau')$.

So $(V, \tau' \upharpoonright V)$ is a zero-dimensional Polish space. And $\mathcal{E} \upharpoonright V$ still is a countable clopen basis for τ' closed under finite intersection.

- And $\mu(V) = 1$ because: Since $\mu(\varphi^{-1}(B)) = \nu(B) = 1$ and $\mu(A) = 1$, we have $\mu(V) = \mu(A') = \nu(v(\overline{C})) = \lambda(v^{-1}v(\overline{C})) \geq \lambda(\overline{C}) = 1$, so $\mu(V) = 1$.

So μ is a probability measure on $\mathcal{B}(\tau' \upharpoonright V)$.

- V is T -invariant: Let $x \in V$, and show $Tx \in V$. We have $x \in \varphi^{-1}(B) \subseteq M$ (the domain of φ), so $T(x) \in M$ and $\varphi(Tx) = S(\varphi(x))$ with $\varphi(x) \in B$. Since $x \in A'$, we have $\varphi(x) \in v(\overline{C})$. Moreover, $v(\overline{C})$ is S -invariant (whence $S\varphi(x) \in v(\overline{C})$): if $y \in v(\overline{C})$, then $y = v(z)$ for $z \in \overline{C}$, so $S(y) = S(v(z)) = v(R(z)) \in v(\overline{C})$ since $R(z) \in \overline{C}$ because \overline{C} is R -invariant. To summarize, $Tx \in A$ (since A is T -invariant), $Tx \in \varphi^{-1}(B)$ (since $\varphi(Tx) = S(\varphi(x)) \in B$ because $\varphi(x) \in B$ and B is S -invariant), and $Tx \in A'$ (since $\varphi(Tx) = S(\varphi(x)) \in v(\overline{C})$). Hence $Tx \in V$.

⁶This exists: Since τ' is second-countable and zero-dimensional, it is not hard to show that τ' has a countable clopen basis \mathcal{E}_1 . (Sketch: τ' has a countable basis B and a zero-dimensional basis B_0 . So each $U \in B$ can be written as a union of B_0 -opens, hence also as a countable union because second-countable spaces are hereditarily Lindelöf. So the collection B_1 of finite intersections of the B_0 -opens occurring in these unions form a countable clopen basis of τ' .) Let \mathcal{E} be the smallest sub-Boolean algebra of $\mathbf{Clp}(\tau')$ generated by $\mathcal{E}_0 \cup \mathcal{E}_1 \subseteq \mathbf{Clp}(\tau')$. This has the required properties.

So $Q := T \upharpoonright V : V \rightarrow V$ is well-defined and continuous (qua restriction of a continuous function to a subspace: If $U \cap V$ is open in V , then $Q^{-1}(U \cap V) = T^{-1}(U) \cap V$ which is open in V).

So we define $\mathfrak{Y} := (V, \tau' \upharpoonright V, \mu, Q)$, whence $(\mathfrak{Y}, \mathcal{E} \upharpoonright V)$ is in \mathbf{bTS}_0 .

Next, we define $w : \mathfrak{Y} \rightarrow \mathfrak{X}$ as the inclusion $V \rightarrow X$. Let's see that w is in W . Base-preserving (and hence continuous): If $U \in \mathcal{B}$, then $w^{-1}(U) = U \cap V \in \mathcal{E} \upharpoonright V$. Measure-preserving: For $C \in \mathcal{B}(\tau)$, we have, since $\mu(V) = 1$, that $\mu(w^{-1}(C)) = \mu(C \cap V) = \mu(C)$, as needed. Equivariant: If $x \in V$, then $w(Q(x)) = w(T(x)) = T(x) = T(w(x))$. And w is injective clop 0 since it is injective on the whole domain.

Next, we define $\psi : \mathfrak{Y} \rightarrow \mathfrak{Z}$. Write v^{-1} for the function $v(\overline{C}) \rightarrow \overline{C}$ that assigns each $y \in v(\overline{C})$ the unique $z \in \overline{C}$ such that $v(z) = y$ (existence follows since $y \in v(\overline{C})$ and uniqueness follows since v is injective on $C' \supseteq \overline{C}$). Note that, if $x \in V$, then $w(x) = x \in V \subseteq A'$, so $\varphi(w(x)) \in v(\overline{C})$. So $v^{-1}(\varphi(w(x)))$ is defined. Thus, we define $\psi := v^{-1} \circ \varphi \circ w : V \rightarrow Z$. We show that ψ is in \mathbf{bTS}_0 .

First note that, for any $U \subseteq Z$, we have $\psi^{-1}(U) = w^{-1}(\varphi^{-1}(v(U \cap \overline{C})))$. Indeed, we have $x \in \psi^{-1}(U)$ iff $v^{-1}(\varphi(w(x))) \in U$ iff the unique v -preimage of $\varphi(w(x))$ in \overline{C} is in U iff $\varphi(w(x)) \in v(\overline{C} \cap U)$ iff $x \in w^{-1}(\varphi^{-1}(v(U \cap \overline{C})))$.

Base-preserving (and hence continuous): For $D \in \mathcal{D}$, we have

$$\psi^{-1}(D) = w^{-1}\varphi^{-1}(v(D \cap \overline{C})) = \varphi^{-1}(v(D \cap \overline{C})) \cap V$$

which is in $\mathcal{E}_0 \upharpoonright V \subseteq \mathcal{E} \upharpoonright V$.

Measure-preserving: For $D \in \mathcal{B}(\rho)$, we have

$$\begin{aligned} \mu(\psi^{-1}(D)) &= \mu(w^{-1}\varphi^{-1}(v(D \cap \overline{C}))) = \mu(\varphi^{-1}(v(D \cap \overline{C}))) = \nu(v(D \cap \overline{C})) \\ &= \lambda(v^{-1}v(D \cap \overline{C})) = \lambda(v^{-1}v(D \cap \overline{C}) \cap \overline{C}) = \lambda(D \cap \overline{C}) = \lambda(D), \end{aligned}$$

where $v^{-1}v(D \cap \overline{C}) \cap \overline{C} = D \cap \overline{C}$ holds since v is injective on \overline{C} .

Equivariant: For $x \in V$, we have $\psi(Q(x)) = v^{-1} \circ \varphi \circ w(Q(x))$. So we need to show that $R(\psi(x))$ is the v -preimage of $\varphi \circ w(Q(x))$ in \overline{C} . Indeed, $R(\psi(x))$ is in \overline{C} (since $\psi(x)$ is and \overline{C} is R -invariant) and

$$v(R\psi(x)) = S(v(\psi(x))) = S(\varphi(w(x))) = \varphi(T(w(x))) = \varphi(w(Q(x))).$$

It remains to check that w and ψ have the required properties. By construction, $v \circ \psi = \varphi \circ w$. Further, ψ maps into \overline{C} , so $\psi(V) \subseteq \overline{C} \subseteq C$. And w maps into V , so $w(V) \subseteq V \subseteq A$. Finally, since w maps into $V \subseteq \varphi^{-1}(B)$, $\varphi \circ w(V) \subseteq B$. Since $\varphi \circ w = v \circ \psi$, also $v \circ \psi(V) \subseteq B$.

Concerning the 'moreover' claim, assume φ is injective on an invariant set of full measure. As indicated in square brackets above, we do the same construction where M now is the domain of injectivity of φ . We need to show that ψ then is in W . We already know that it is in \mathbf{bTS}_0 , so we need to show that it is injective

clop 0, for which it suffices to show that ψ is injective on the whole domain V : Let $x \neq x'$ in V and show $\psi(x) \neq \psi(x')$. We have $w(x) = x \neq x' = w(x')$, which are in $V \subseteq M$ (since $V \subseteq \varphi^{-1}(B) \subseteq M$ because M is the domain of φ). Since φ is injective on M , $\varphi(w(x)) \neq \varphi(w(x'))$. Hence, the v -preimages of these two elements have to be distinct, since otherwise v wouldn't be a function. Hence $\psi(x) = v^{-1}(\varphi(w(x))) \neq v^{-1}(\varphi(w(x'))) = \psi(x')$, as needed. \square

For the standard case, we first make a small observation: Let $\mathfrak{X} = (X, \tau, \mu, T)$ be in TS_{0s} . So $\bar{J}(\mathfrak{X}) = (X, \mathcal{B}(\tau)_\mu, \mu, T)$. Now, if $M \in \mathcal{B}(\tau)_\mu$ has full measure, then there also is $M' \in \mathcal{B}(\tau)$ with $M' \subseteq M$, $\mu(M') = 1$, and $T(M') = M'$.⁷

A.3.2. LEMMA (standard). *Let $(\mathfrak{X}, \mathcal{B}), (\mathfrak{Y}, \mathcal{C}), (\mathfrak{Z}, \mathcal{D})$ be in bTS_{0s} . Let A, B, C be invariant and of full measure in $\mathfrak{X}, \mathfrak{Y}, \mathfrak{Z}$, respectively. Let $v : (\mathfrak{Z}, \mathcal{D}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ be in W . Let $\varphi : \bar{J}(\mathfrak{X}) \rightarrow \bar{J}(\mathfrak{Y})$ be in sDS . Then there is $\psi : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ in bTS_{0s} and $w : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ in W such that $\varphi \circ w = v \circ \psi$ and, writing V for the state space of \mathfrak{Y} , $\psi(V) \subseteq C$, $w(V) \subseteq A$, and $\varphi \circ w(V) = v \circ \psi(V) \subseteq B$. Moreover, if $\varphi : X \rightarrow Y$ is injective on an invariant set of full measure, then ψ can be chosen in W .*

PROOF. Write $\mathfrak{X} = (X, \tau, \mu, T)$, $\mathfrak{Y} = (Y, \sigma, \nu, S)$, and $\mathfrak{Z} = (Z, \rho, \lambda, R)$. Since $v \in W$, let C' be clopen in \mathfrak{Z} with $\lambda(C') = 1$, $R(C') \subseteq C'$ and v injective on C' . So C and C' are invariant sets of full measure, whence $\bar{C} := C \cap C'$ is, too. Since $\varphi \in \text{sDS}$, it is a partial function with invariant domain $M \subseteq X$ of full measure. [If φ is injective on an invariant set of full measure, choose M as this set.]

As before, $v(\bar{C} \cap D) \subseteq Y$ is Borel for every $D \in \mathcal{D}$. For $\mathcal{D} = \{D_0, D_1, D_2, \dots\}$ with $D_0 = Y$, we have $\varphi^{-1}(v(\bar{C} \cap D_n)) = A_n \cup N_n$ with $A_n \in \mathcal{B}(\tau)$ and N_n a μ -nullset. Since $M \setminus \bigcup_n N_n \in \mathcal{B}(\tau)_\mu$ has full measure, there also is, by the observation, $M' \in \mathcal{B}(\tau)$ with $M' \subseteq M \setminus \bigcup_n N_n$, $\mu(M') = 1$, and $T(M') = M'$. Moreover, $\varphi^{-1}(B) = A' \cup N'$ for $A' \in \mathcal{B}(\tau)$ and N' a μ -nullset.

We construct \mathfrak{Y} . Define $\mathcal{E}_0 := \mathcal{B} \cup \{A_n : n \in \omega\} \cup \{A', A, M'\}$. Hence $\mathcal{E}_0 \subseteq \mathcal{B}(\tau)$ is countable. Since (X, τ, μ, T) in particular is in TS_0 , there is, by corollary A.2.2, a zero-dimensional Polish topology $\tau' \supseteq \tau$ on X such that (a) $\mathcal{B}(\tau) = \mathcal{B}(\tau')$, (b) $\mathcal{E}_0 \subseteq \text{Clp}(\tau')$, (c) $T : X \rightarrow X$ is τ' -continuous. Let \mathcal{E} be a countable clopen basis for τ' that includes \mathcal{E}_0 and is closed under finite intersection.

Now, define $V_0 := A' \cap A_0 \cap A \cap M'$. As before, $V_0 \subseteq X$ is τ' -clopen. And V_0 is of full measure: Since $\mu(A') = \mu(\varphi^{-1}(B)) = \nu(B) = 1$, $\mu(A) = 1$, and $\mu(M') = 1$,

⁷Proof: We have $M = M_0 \cup N$ for $M_0 \in \mathcal{B}(\tau)$ and a μ -nullset N . Set $M' := \bigcap_{k \in \mathbb{Z}} T^{-k}(M_0) \subseteq M_0 \subseteq M$. Borel: Since $T : X \rightarrow X$ is continuous, T is Borel-measurable. And since $T : (X, \mathcal{B}(\tau)) \rightarrow (X, \mathcal{B}(\tau))$ is a Borel-measurable injective function between standard Borel spaces, also T -images of Borel sets are Borel. So, since M_0 is Borel, M' is Borel qua countable intersection of Borel sets. Full measure: Since T is measure-preserving and bijective, $\mu(T^{-k}(M_0)) = \mu(M_0) = 1$ (for all $k \in \mathbb{Z}$), so M' is a countable intersection of sets of full measure, whence has full measure. 'Two-sided' invariant: We have $x \in M'$ iff $\forall k \in \mathbb{Z} : T^k x \in M_0$ iff $\forall k \in \mathbb{Z} : T^{k+1} x = T^k T x \in M_0$ iff $T x \in M'$.

we have $\mu(V) = \mu(A_0) = \mu(\varphi^{-1}(v(\overline{C}))) = \nu(v(\overline{C})) = \lambda(v^{-1}v(\overline{C})) \geq \lambda(\overline{C}) = 1$, so $\mu(V) = 1$. As in the observation, we define $V := \bigcap_{k \in \mathbb{Z}} T^{-k}(V_0)$, which is Borel, has full measure, and $T(V) = V$. In fact, since T is a homeomorphism, each $T^{-k}(V_0)$ is clopen, so V is closed. Then we have

- $(V, \tau' \upharpoonright V)$ is a zero-dimensional Polish space qua closed subset of a Polish space. So $\mathcal{E} \upharpoonright V$ still is a countable clopen basis for τ' .
- As noted, $\mu(V) = 1$. So μ is a probability measure on $\mathcal{B}(\tau' \upharpoonright V)$.
- $Q := T \upharpoonright V : V \rightarrow V$ is well-defined, injective (since T is), surjective (since $T(V) = V$), and continuous (qua restriction of a continuous function) and open (if $U \cap V$ is open in V , then, by injectivity, $Q(U \cap V) = T(U \cap V) = T(U) \cap T(V) = T(U) \cap V$ which is open since T is open). So Q is a homeomorphism. Also, Q is measure-preserving: If $U \cap V$ is open in V , then $\mu(Q^{-1}(U \cap V)) = \mu(T^{-1}(U) \cap V) = \mu(T^{-1}(U)) = \mu(U) = \mu(U \cap V)$, so the probability measure μ and μQ^{-1} agree on the open sets of $\tau' \upharpoonright V$, whence they are identical on $\mathcal{B}(\tau' \upharpoonright V)$.

So we define $\mathfrak{Y} := (V, \tau' \upharpoonright V, \mu, Q)$, whence $(\mathfrak{Y}, \mathcal{E} \upharpoonright V)$ is in \mathbf{bTS}_{0s} .

As before, we define $w : \mathfrak{Y} \rightarrow \mathfrak{X}$ as the inclusion $V \rightarrow X$, and see, verbatim as before, that it is in W .

We define $\psi : \mathfrak{Y} \rightarrow \mathfrak{Z}$ as before, too: Write v^{-1} for the function $v(\overline{C}) \rightarrow \overline{C}$ that assigns each $y \in v(\overline{C})$ the unique $z \in \overline{C}$ such that $v(z) = y$. Note that, if $x \in V$, then $w(x) = x \in V \subseteq V_0 \subseteq A_0 = \varphi^{-1}(v(\overline{C}))$, so $\varphi(w(x)) \in v(\overline{C})$. So $v^{-1}(\varphi(w(x)))$ is defined. Thus, we define $\psi := v^{-1} \circ \varphi \circ w : V \rightarrow Z$. We see that ψ is in \mathbf{bTS}_{0s} .

Again, for any $U \subseteq Z$, we have $\psi^{-1}(U) = w^{-1}(\varphi^{-1}(v(U \cap \overline{C})))$. So ψ is base-preserving (and hence continuous): For D_n in \mathcal{D} , we have

$$\psi^{-1}(D_n) = w^{-1}\varphi^{-1}(v(D_n \cap \overline{C})) = \varphi^{-1}(v(D_n \cap \overline{C})) \cap V = A_n \cap V,^8$$

which is in $\mathcal{E}_0 \upharpoonright V \subseteq \mathcal{E} \upharpoonright V$. Measure-preservation and equivariance are seen verbatim as before.

It remains to check that w and ψ have the required properties. By construction, $v \circ \psi = \varphi \circ w$. Further, ψ maps into \overline{C} , so $\psi(V) \subseteq \overline{C} \subseteq C$. And w maps into V , so $w(V) \subseteq V \subseteq V_0 \subseteq A$. Finally, since w maps into $V \subseteq V_0 \subseteq A' \subseteq \varphi^{-1}(B)$, we have $\varphi \circ w(V) \subseteq B$. Since $\varphi \circ w = v \circ \psi$, also $v \circ \psi(V) \subseteq B$.

The ‘moreover’ claim follows verbatim as before. \square

⁸For the last equation, we have \supseteq since $\varphi^{-1}(v(D_n \cap \overline{C})) = A_n \cup N_n \supseteq A_n$. And to show \subseteq , let $x \in \varphi^{-1}(v(D_n \cap \overline{C})) \cap V = (A_n \cup N_n) \cap V$. So $x \in M' \subseteq M \setminus \bigcup_n N_n$. So $x \notin N_n$, whence $x \in A_n$, so $x \in A_n \cap V$.

A.4 Calculus of fractions

A.4.1. PROPOSITION (general). *The class W of \mathbf{bTS}_0 -morphisms that are injective clop 0 admits a calculus of fractions.*

PROOF. We fix the notation $\mathfrak{X} = (X, \tau, \mu, T)$, $\mathfrak{Y} = (Y, \sigma, \nu, S)$, and $\mathfrak{Z} = (Z, \rho, \lambda, R)$. We need to check properties (1)–(4) from definition A.1.2.

Ad (1). For $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_0 , the morphism $\text{id}_{(\mathfrak{X}, \mathcal{B})}$ is the identity function $\text{id}_X : X \rightarrow X$, which is injective on the whole domain, so in W .

Ad (2). Assume $v : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ and $w : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{Z}, \mathcal{D})$ are in W , and show $w \circ v \in W$. Since v is injective clop 0, there is a clopen $A \subseteq X$ such that $\mu(A) = 1$, $T(A) \subseteq A$, and, for $x \neq x'$ in A , $v(x) \neq v(x')$. Since w is injective clop 0, there is a clopen $B \subseteq Y$ such that $\nu(B) = 1$, $S(B) \subseteq B$, and, for $y \neq y'$ in B , $w(y) \neq w(y')$. Set $A' := A \cap v^{-1}(B)$. We show that A' witnesses that $w \circ v$ is injective clop 0, and hence in W . Since v is continuous, also $v^{-1}(B)$ is clopen, so A' is clopen qua finite union of clopen sets. Since v is measure-preserving, $\mu(v^{-1}(B)) = \nu(B) = 1$, so A' is an intersection of sets of full measure, hence $\mu(A') = 1$. Since v is equivariant, if $v(x) \in B$, then, since B is S -invariant, $v(T(x)) = S(v(x)) \in B$, hence, since A is T -invariant, also A' is T -invariant. Finally, if $x, x' \in A'$ with $x \neq x'$, then $x, x' \in A$ and $v(x), v(x') \in B$, so $v(x) \neq v(x')$, whence $w \circ v(x) \neq w \circ v(x')$.

Ad (3). Assume $\varphi : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ is in \mathbf{bTS}_0 and $v : (\mathfrak{Z}, \mathcal{D}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ is in W . In particular, φ , regarded as partial function with total domain X and total codomain Y , is a system homomorphism from $(X, \mathcal{B}(\tau), \mu, T)$ to $(Y, \mathcal{B}(\sigma), \nu, S)$. So lemma A.3.1 applies (choosing X, Y, Z as A, B, C , respectively), whence there is $\psi : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{Z}, \mathcal{D})$ in \mathbf{bTS}_0 and $w : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ in W such that $\varphi \circ w = v \circ \psi$.

Ad (4). Assume $\varphi, \psi : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ are in \mathbf{bTS}_0 and $v : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{Z}, \mathcal{D})$ is in W with $v \circ \varphi = v \circ \psi$. Find $w : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ in W such that $\varphi \circ w = \psi \circ w$.

Since $v \in W$, let B be clopen in \mathfrak{Y} with $\nu(B) = 1$, $S(B) \subseteq B$, and v injective on B . We construct the required $\mathfrak{Y} = (V, \pi, \kappa, Q)$ as follows. Let $V := \varphi^{-1}(B) \cap \psi^{-1}(B) \subseteq X$. Since φ and ψ are continuous, $V \subseteq X$ is closed. Let π be the subspace topology on V . Since $V \subseteq X$ is closed, this is again Polish. Since (X, τ) is zero-dimensional, also the subspace (V, π) is zero-dimensional. Since $\mathcal{B}(\pi) \subseteq \mathcal{B}(\tau)$, we can define $\kappa(A) := \mu(A)$. This is a probability measure since $\mu(\varphi^{-1}B) = \nu(B) = 1$ and $\mu(\psi^{-1}B) = \nu(B) = 1$ so $\kappa(V) = \mu(V) = \mu(\varphi^{-1}(B) \cap \psi^{-1}(B)) = 1$ qua intersection of sets of full measure. Finally, $Q : V \rightarrow V$ is defined as $T \upharpoonright V$. Note that this is well-defined since, if $x \in V$, then $Tx \in V$: Since $\varphi(x) \in B$ and $\psi(x) \in B$, also $\varphi(Tx) = S(\varphi(x)) \in B$ and $\psi(Tx) = S(\psi(x)) \in B$. Further, Q is continuous since T is continuous and π is the relative topology. Hence \mathfrak{Y} is in \mathbf{TS}_0 . Next, let $\mathcal{E} := \{U \cap V : U \in \mathcal{B}\}$. Since \mathcal{B} is a countable clopen basis for (X, τ) and (V, π) is the subspace, \mathcal{E} is a countable clopen basis for (V, π) . So $(\mathfrak{Y}, \mathcal{E})$ is in \mathbf{bTS}_0 .

We define $w : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ as the inclusion. This is base-preserving, and

hence continuous (if $U \in \mathcal{B}$, then $w^{-1}(U) = U \cap V \in \mathcal{E}$), measure-preserving (if $A \in \mathcal{B}(\tau)$, then $\kappa(w^{-1}(A)) = \kappa(A \cap V) = \mu(A \cap V) = \mu(A)$), equivariant (for $x \in V$, $w(Q(x)) = T(x) = T(w(x))$), and injective, whence injective clop 0. So $w \in W$.

Finally, we show $\varphi \circ w = \psi \circ w$: If $x \in V$, then $\varphi(x) \in B$ and $\psi(x) \in B$. Since $v \circ \varphi = v \circ \psi$ by assumption, $v(\varphi(x)) = v(\psi(x))$. Since v is injective on B , $\varphi \circ w(x) = \varphi(x) = \psi(x) = \psi \circ w(x)$, as needed. \square

A.4.2. PROPOSITION (standard). *The class W of \mathbf{bTS}_{0s} -morphisms that are injective clop 0 admits a calculus of fractions.*

PROOF. We fix the notation $\mathfrak{X} = (X, \tau, \mu, T)$, $\mathfrak{Y} = (Y, \sigma, \nu, S)$, and $\mathfrak{Z} = (Z, \rho, \lambda, R)$. We need to check properties (1)–(4) from definition A.1.2.

Ad (1). As before, for $(\mathfrak{X}, \mathcal{B})$ in \mathbf{bTS}_{0s} , the morphism $\text{id}_{(\mathfrak{X}, \mathcal{B})}$ is in W .

Ad (2). Again as before, if $v : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ and $w : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{Z}, \mathcal{D})$ are in W , then $w \circ v \in W$.

Ad (3). This again is done as before, now using lemma A.3.2.

Ad (4). Assume $\varphi, \psi : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ are in \mathbf{bTS}_{0s} and $v : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{Z}, \mathcal{D})$ is in W with $v \circ \varphi = v \circ \psi$. Find $w : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ in W such that $\varphi \circ w = \psi \circ w$.

Since $v \in W$, let B be clopen in \mathfrak{Y} with $\nu(B) = 1$, $S(B) \subseteq B$, and v injective on B . We construct the required $\mathfrak{Y} = (V, \pi, \kappa, Q)$ as follows. Let $V_0 := \varphi^{-1}(B) \cap \psi^{-1}(B) \subseteq X$. Since φ and ψ are continuous and measure-preserving, $V_0 \subseteq X$ is closed and of full measure. Set $V := \bigcap_{k \in \mathbb{Z}} T^{-k}(V_0)$, whence $V \subseteq V_0$ and $T(V) = V$. Since T is a homeomorphism and V_0 closed, V is closed and of full measure (qua countable intersection of closed sets of full measure). So V with the subspace topology π is again zero-dimensional Polish. And $\kappa := \mu \upharpoonright \mathcal{B}(\pi)$ is a probability measure on $\mathcal{B}(\pi) \subseteq \mathcal{B}(\tau)$ since $\mu(V) = 1$. Finally, since $T(V) = V$, we have that $Q := T \upharpoonright V : V \rightarrow V$ is a well-defined homeomorphism that is measure-preserving (as in the third bullet point of the proof of lemma A.3.2). Next, let $\mathcal{E} := \{U \cap V : U \in \mathcal{B}\}$. So $(\mathfrak{Y}, \mathcal{E})$ is in \mathbf{bTS}_{0s} .

We define $w : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{X}, \mathcal{B})$ as the inclusion. As before, this is in W .

Finally, we show $\varphi \circ w = \psi \circ w$: If $x \in V \subseteq V_0$, then $\varphi(x) \in B$ and $\psi(x) \in B$. Since $v \circ \varphi = v \circ \psi$ by assumption, $v(\varphi(x)) = v(\psi(x))$. Since v is injective on B , $\varphi \circ w(x) = \varphi(x) = \psi(x) = \psi \circ w(x)$, as needed. \square

A.5 Equivalence

Recall that $F : \mathbf{bTS}_0 \rightarrow \mathbf{TS}_0$ is the forgetful functor that sends $(\mathfrak{X}, \mathcal{B})$ to \mathfrak{X} and is the identity on morphisms. Also recall that the functor $J : \mathbf{TS}_0 \rightarrow \mathbf{DS}$ sends a topological system (X, τ, μ, T) to the naturally induced dynamical system $(X, \mathcal{B}(\tau), \mu, T)$ and is the identity on morphisms.

A.5.1. PROPOSITION (general). *We define a functor $G : \mathbf{bTS}_0[W^{-1}] \rightarrow \mathbf{DS}$ by*

- *For an object $(\mathfrak{X}, \mathcal{B}) \in \mathbf{bTS}_0[W^{-1}]$, define $G(\mathfrak{X}, \mathcal{B}) := J \circ F(\mathfrak{X}, \mathcal{B})$.*
- *For a morphism $[v, \mathfrak{Z}, \varphi]$ in $\mathbf{bTS}_0[W^{-1}]$, define $G([v, \mathfrak{Z}, \varphi]) := \varphi \circ v^{-1}$ regarded as partial function with domain $M := v(C) \subseteq X$ (where C is a clopen domain of injectivity of v) and codomain $\varphi(C) \subseteq Y$.*

It is essentially surjective, full, and faithful, whence $\mathbf{bTS}_0[W^{-1}]$ and \mathbf{DS} are equivalent categories.

PROOF. Functor. We have the functor $J \circ F : \mathbf{bTS}_0 \rightarrow \mathbf{DS}$, and, for $v \in W$, $J \circ F(v) = v$ is an isomorphism in \mathbf{DS} by lemma 5.2.4 from chapter 5 (it is injective on an invariant set of full measure). Now, by the universal property of the localization $\mathbf{bTS}_0[W^{-1}]$ (see, e.g., Yekutieli 2020, def. 6.1.2), there is a unique factoring functor G such that

$$\begin{array}{ccc} \mathbf{bTS}_0 & \xrightarrow{J \circ F} & \mathbf{DS} \\ \downarrow Q & \nearrow G & \\ \mathbf{bTS}_0[W^{-1}] & & \end{array}$$

commutes. This G has to be given by

- for $(\mathfrak{X}, \mathcal{B})$ in $\mathbf{bTS}_0[W^{-1}]$, $G(\mathfrak{X}, \mathcal{B}) = G \circ Q(\mathfrak{X}, \mathcal{B}) = J \circ F(\mathfrak{X}, \mathcal{B})$, and
- for $[v, \mathfrak{Z}, \varphi]$ in $\mathbf{bTS}_0[W^{-1}]$, $G([v, \mathfrak{Z}, \varphi]) = (J \circ F)(\varphi) \circ (J \circ F)(v)^{-1} = \varphi \circ v^{-1}$.⁹

This is precisely the functor described in the proposition.

Essentially surjective. Let $\mathfrak{X} = (X, \mathcal{A}, \mu, T)$ be in \mathbf{DS} . By proposition A.2.1 (choosing \emptyset as the subset of \mathcal{A}), there is a zero-dimensional Polish topology τ on X with $\mathcal{A} = \mathcal{B}(\tau)$ and $T : X \rightarrow X$ is τ -continuous. Let \mathcal{B} be a countable clopen basis for (X, τ) that is closed under intersection. Then $(\mathfrak{X}', \mathcal{B})$ with $\mathfrak{X}' := (X, \tau, \mu, T)$ is in \mathbf{bTS}_0 and $G(\mathfrak{X}', \mathcal{B}) = (X, \mathcal{B}(\tau), \mu, T) = (X, \mathcal{A}, \mu, T) = \mathfrak{X}$ (which implies isomorphism) in \mathbf{DS} .

Full. Let $(\mathfrak{X}, \mathcal{B})$ and $(\mathfrak{Y}, \mathcal{C})$ be in $\mathbf{bTS}_0[W^{-1}]$. Let $\varphi : J(\mathfrak{X}) = G(\mathfrak{X}, \mathcal{B}) \rightarrow G(\mathfrak{Y}, \mathcal{C}) = J(\mathfrak{Y})$ be a morphism in \mathbf{DS} . We need to find a morphism $[v, (\mathfrak{Z}, \mathcal{D}), \psi] : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ in $\mathbf{bTS}_0[W^{-1}]$ such that $G([v, (\mathfrak{Z}, \mathcal{D}), \psi]) = \varphi$.

Let $\text{id}_{(\mathfrak{Y}, \mathcal{C})} = \text{id}_Y : (\mathfrak{Y}, \mathcal{C}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ be the identity function, which is in W . By lemma A.3.1,

$$\begin{array}{ccc} (\mathfrak{Z}, \mathcal{D}) & \xrightarrow{\psi} & (\mathfrak{Y}, \mathcal{C}) \\ \downarrow v & & \downarrow \text{id}_Y \\ (\mathfrak{X}, \mathcal{B}) & \xrightarrow{\varphi} & (\mathfrak{Y}, \mathcal{C}) \end{array}$$

⁹Since $[v, \mathfrak{Z}, \varphi] = [\text{id}_Z, \mathfrak{Z}, \varphi] \circ [v, \mathfrak{Z}, \text{id}_Z] = Q(\varphi) \circ Q(v)^{-1}$, we have $G([v, \mathfrak{Z}, \varphi]) = GQ(\varphi) \circ GQ(v)^{-1} = JF(\varphi) \circ JF(v)^{-1}$.

i.e., there is $(\mathfrak{Z}, \mathcal{D})$ in \mathbf{bTS}_0 and $v : (\mathfrak{Z}, \mathcal{D}) \rightarrow (\mathfrak{X}, \mathcal{B})$ in W and $\psi : (\mathfrak{Z}, \mathcal{D}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ in \mathbf{bTS}_0 with $\varphi \circ v = \text{id}_V \circ \psi = \psi$. (The quotation marks again indicate that φ is only assumed to be a morphism $J(\mathfrak{X}) \rightarrow J(\mathfrak{Y})$.) Hence

$$(\mathfrak{X}, \mathcal{B}) \xleftarrow{v} (\mathfrak{Z}, \mathcal{D}) \xrightarrow{\psi} (\mathfrak{Y}, \mathcal{C})$$

is a span in \mathbf{bTS}_0 , so $[v, (\mathfrak{Z}, \mathcal{D}), \psi] : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ is in $\mathbf{bTS}_0[W^{-1}]$. And $G([v, (\mathfrak{Z}, \mathcal{D}), \psi]) = \psi \circ v^{-1} = \varphi$, where the last equation follows from $\varphi \circ v = \psi$.

Faithful. Let $(\mathfrak{X}, \mathcal{B})$ and $(\mathfrak{Y}, \mathcal{C})$ be objects in $\mathbf{bTS}_0[W^{-1}]$. Let

$$[v, (\mathfrak{Z}, \mathcal{D}), \varphi], [v', (\mathfrak{Z}', \mathcal{D}'), \varphi'] : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$$

be two morphisms in $\mathbf{bTS}_0[W^{-1}]$. Assume $G([v, (\mathfrak{Z}, \mathcal{D}), \varphi]) = G([v', (\mathfrak{Z}', \mathcal{D}'), \varphi'])$ in \mathbf{DS} and show $[v, (\mathfrak{Z}, \mathcal{D}), \varphi] = [v', (\mathfrak{Z}', \mathcal{D}'), \varphi']$.

Write $\mathfrak{X} = (X, \tau, \mu, T)$, $\mathfrak{Y} = (Y, \sigma, \nu, S)$, $\mathfrak{Z} = (Z, \rho, \lambda, R)$, $\mathfrak{Z}' = (Z', \rho', \lambda', R')$. By the assumption, there is a T -invariant set $A \subseteq X$ of full μ -measure such that the partial functions $G([v, (\mathfrak{Z}, \mathcal{D}), \varphi]) = \varphi \circ v^{-1}$ and $G([v', (\mathfrak{Z}', \mathcal{D}'), \varphi']) = \varphi' \circ v'^{-1}$ are defined and identical on A . Since v^{-1} is the inverse of v in \mathbf{DS} , there is an R -invariant set $C \subseteq Z$ of full λ -measure such that $v^{-1} \circ v = \text{id}_Z$ on C . Similarly, there is an R' -invariant set $C' \subseteq Z'$ of full λ' -measure such that $v'^{-1} \circ v' = \text{id}_{Z'}$ on C' . By lemma A.3.1,

$$\begin{array}{ccc} (\mathfrak{Y}, \mathcal{E}) & \xrightarrow{\psi=w'} & (\mathfrak{Z}', \mathcal{D}') & C' \\ \vdots \downarrow w & & \downarrow v' & \\ C & (\mathfrak{Z}, \mathcal{D}) \xrightarrow{v} & (\mathfrak{X}, \mathcal{B}) & A \end{array}$$

i.e., there is $w : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{Z}, \mathcal{D})$ and $w' := \psi : (\mathfrak{Y}, \mathcal{E}) \rightarrow (\mathfrak{Z}', \mathcal{D}')$ in W such that $v \circ w = v' \circ w'$, $w(V) \subseteq C$, $w'(V) \subseteq C'$, and $v \circ w(V) = v' \circ w'(V) \subseteq A$.

It suffices to show that $\varphi \circ w = \varphi' \circ w'$. Because, since we already have $v \circ w = v' \circ w' \in W$, the diagram

$$\begin{array}{ccccc} & & (\mathfrak{Z}, \mathcal{D}) & & \\ & \swarrow v & \uparrow w & \searrow \varphi & \\ (\mathfrak{X}, \mathcal{B}) & & (\mathfrak{Y}, \mathcal{E}) & & (\mathfrak{Y}, \mathcal{C}) \\ & \swarrow v' & \downarrow w' & \searrow \varphi' & \\ & & (\mathfrak{Z}', \mathcal{D}') & & \end{array}$$

then shows that the spans $(v, (\mathfrak{Z}, \mathcal{D}), \varphi)$ and $(v', (\mathfrak{Z}', \mathcal{D}'), \varphi')$ are equivalent.

Indeed, let $y \in V$, and show $\varphi \circ w(y) = \varphi' \circ w'(y)$. We have $v \circ w(y) = v' \circ w'(y) \in A$. Since the partial functions $\varphi \circ v^{-1}$ and $\varphi' \circ v'^{-1}$ are defined and agree on A , we hence have

$$\varphi \circ v^{-1}(v \circ w(y)) = \varphi' \circ v'^{-1}(v' \circ w'(y)).$$

Since $w(y) \in C$, we have $v^{-1}(v \circ w(y)) = v^{-1} \circ v(w(y)) = w(y)$. Similarly, $v'^{-1}(v' \circ w'(y)) = w'(y)$. Hence the above equation simplifies to $\varphi \circ w(y) = \varphi' \circ w'(y)$, as needed. \square

The standard case is very similar due to our preparation. Recall that the functor $\bar{J} : \text{TS}_{0s} \rightarrow \text{sDS}$ sends a topological system (X, τ, μ, T) to the naturally induced standard dynamical system $(X, \mathcal{B}(\tau)_\mu, \mu, T)$ and is the identity on morphisms.

A.5.2. PROPOSITION (standard). *We define a functor $G : \text{bTS}_{0s}[W^{-1}] \rightarrow \text{sDS}$ by*

- *For an object $(\mathfrak{X}, \mathcal{B}) \in \text{bTS}_0[W^{-1}]$, define $G(\mathfrak{X}, \mathcal{B}) := \bar{J} \circ F(\mathfrak{X}, \mathcal{B})$.*
- *For a morphism $[v, \mathfrak{Z}, \varphi]$ in $\text{bTS}_{0s}[W^{-1}]$, define $G([v, \mathfrak{Z}, \varphi]) := \varphi \circ v^{-1}$ regarded as partial function with domain $M := v(C) \subseteq X$ (where C is a clopen domain of injectivity of v) and codomain $\varphi(C) \subseteq Y$.*

It is essentially surjective, full, and faithful, whence $\text{bTS}_{0s}[W^{-1}]$ and sDS are equivalent categories.

PROOF. Functor. The functor $\bar{J} \circ F : \text{bTS}_0 \rightarrow \text{DS}$ sends elements from W to isomorphisms in sDS by lemma 5.2.4 from chapter 5 (each v is injective on an invariant set of full measure). By the universal property of the localization $\text{bTS}_{0s}[W^{-1}]$, there is G such that

$$\begin{array}{ccc}
 \text{bTS}_{0s} & \xrightarrow{\bar{J} \circ F} & \text{sDS} \\
 \downarrow Q & \nearrow G & \uparrow \tau \\
 \text{bTS}_{0s}[W^{-1}] & &
 \end{array}$$

commutes, and, as before, this G has to be the one stated in the proposition.

Essentially surjective. Let \mathfrak{X} be in sDS . By proposition A.2.3, there is $\mathfrak{Y} = (Y, \sigma, \nu, S)$ in TS_{0s} such that $\bar{J}(\mathfrak{Y})$ is isomorphic to \mathfrak{X} . Let \mathcal{B} be a countable clopen basis for (Y, σ) that is closed under intersection and S -image (their existence was shown in corollary 4.6.7 of chapter 4). Then $(\mathfrak{Y}, \mathcal{B})$ is in bTS_{0s} and $G(\mathfrak{Y}, \mathcal{B}) = \bar{J} \circ F(\mathfrak{Y}, \mathcal{B}) = \bar{J}(\mathfrak{Y})$ is isomorphic to \mathfrak{X} in sDS .

Full. Let $(\mathfrak{X}, \mathcal{B})$ and $(\mathfrak{Y}, \mathcal{C})$ be in $\text{bTS}_{0s}[W^{-1}]$. Let $\varphi : G(\mathfrak{X}, \mathcal{B}) \rightarrow G(\mathfrak{Y}, \mathcal{C})$ be a morphism in sDS . We need to find a morphism $[v, (\mathfrak{Z}, \mathcal{D}), \psi] : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ in $\text{bTS}_{0s}[W^{-1}]$ such that $G([v, (\mathfrak{Z}, \mathcal{D}), \psi]) = \varphi$. Indeed, by lemma A.3.2,

$$\begin{array}{ccc}
 (\mathfrak{Z}, \mathcal{D}) & \xrightarrow{\psi} & (\mathfrak{Y}, \mathcal{C}) \\
 \downarrow v & & \downarrow \text{id}_Y \\
 (\mathfrak{X}, \mathcal{B}) & \xrightarrow{\varphi} & (\mathfrak{Y}, \mathcal{C})
 \end{array}$$

so $[v, (\mathfrak{Z}, \mathcal{D}), \psi] : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$ is in $\mathbf{bTS}_{0s}[W^{-1}]$. And $\mathbf{G}([v, (\mathfrak{Z}, \mathcal{D}), \psi]) = \psi \circ v^{-1} = \varphi$, where the last equation follows from $\varphi \circ v = \psi$.

Faithful. Let $(\mathfrak{X}, \mathcal{B})$ and $(\mathfrak{Y}, \mathcal{C})$ be objects in $\mathbf{bTS}_{0s}[W^{-1}]$. Let

$$[v, (\mathfrak{Z}, \mathcal{D}), \varphi], [v', (\mathfrak{Z}', \mathcal{D}'), \varphi'] : (\mathfrak{X}, \mathcal{B}) \rightarrow (\mathfrak{Y}, \mathcal{C})$$

be two morphisms in $\mathbf{bTS}_{0s}[W^{-1}]$ with identical \mathbf{G} -image, and show that they are identical. We can reason verbatim as before except for now referring to lemma A.3.2 and writing $\mathbf{bTS}_{0s}[W^{-1}]$ and \mathbf{sDS} . \square

Thus, we've proven theorem A.1.3: In the general case, let W be the class of morphisms in \mathbf{bTS}_0 that are injective clop 0. Then W admits a right calculus of fractions (proposition A.4.1) and the category of fractions $\mathbf{bTS}_0[W^{-1}]$ is equivalent to \mathbf{DS} via a functor that is like $\mathbf{J} \circ \mathbf{F}$ on objects (proposition A.5.1). In the standard case, let W be the class of morphisms in \mathbf{bTS}_{0s} that are injective clop 0. Then W admits a right calculus of fractions (proposition A.4.2) and the category of fractions $\mathbf{bTS}_{0s}[W^{-1}]$ is equivalent to \mathbf{sDS} via a functor that is like $\bar{\mathbf{J}} \circ \mathbf{F}$ on objects (proposition A.5.2).

Appendix B

Dynamical domain example

The purpose of this appendix is to discuss in detail a non-trivial but elementary example of a dynamical domain. This domain-theoretic structure was introduced and discussed in part 2. We've seen that every dynamical system induces a dynamical domain (namely its observation domain). This provides a wealth of examples of dynamical domains. In fact, when restricting to the max-reflective ones, all of them are obtained in this way. To also illustrate the other direction, we construct here a non-max-reflective dynamical domain in a 'purely domain-theoretic' way, i.e., not via dynamical systems.

In section B.1, we construct the example dynamical domain \mathfrak{D} . (This essentially only requires the definition of a dynamical domain from chapter 4.) In section B.2, we note that the topological dynamical system $S(\mathfrak{D})$ that \mathfrak{D} models is a Cantor dynamics—i.e., a topological system with state space 2^ω —which is not expansive, so it is not isomorphic to a subshift. (The idea that a dynamical domain models a system also is found in chapter 4.) In sections B.3 and B.4, we determine the max-entropy of \mathfrak{D} as $m(\mathfrak{D}) = \infty$, which also is the topological entropy of $S(\mathfrak{D})$. (This requires chapter 6.) We end with some further questions relating to algorithmic randomness.

B.1 A dynamical domain of binary sequences

Our example dynamical domain will be the following constructed around the well-known Scott domain of finite and infinite binary sequences.

B.1.1. PROPOSITION. *Let $\mathfrak{D} = (D, v, f)$ be defined as follows:*

- $D := 2^{\leq \omega}$ is the set of finite and infinite (i.e., ω -long) binary sequences ordered by extension (denoted \leq).¹

¹A binary sequence is a sequence of 0's and 1's, e.g., $x = 0100111010\dots$. Formally, they are partial functions $\omega \rightarrow 2$ whose domain is of the form $\{n : 0 \leq n \leq l\}$ for $l \in \omega \cup \{\omega\}$ the length of the sequence.

- $v : \Sigma(D) \rightarrow [0, 1]$ is the ‘Lebesgue’ valuation: $v(U) := \lambda(\max U)$ where λ is the Lebesgue measure on Cantor space 2^ω .²
- $f : D \rightarrow D$ is defined by

$$f(x)(n) := \begin{cases} x(n) + x(2n) & \text{if } x(2n) \text{ is defined} \\ \text{undefined} & \text{otherwise,} \end{cases}$$

where ‘+’ is addition modulo 2.³

Then \mathfrak{D} is a dynamical domain (i.e., in \mathbf{dDOM}).

Before we prove this, let’s gather some intuition for the components D , v , and f of \mathfrak{D} .

Concerning the domain D : This is a standard example of a Scott domain. Its compact elements are precisely the finite binary sequences.⁴ So the non-compact elements hence are precisely the maximal ones: $\max D = 2^\omega$. The Scott topology on D has the basis $\{\uparrow a : a \in D \text{ finite}\}$, so the relative Scott topology on $\max D = 2^\omega$ is the usual topology of Cantor space given by the cylinder sets $\llbracket a \rrbracket = \{x \in 2^\omega : a \leq x\} = \uparrow a \cap \max D$.

Concerning the valuation v : We see that this is a valuation as follows. It is strict: $v(\emptyset) = \lambda(\emptyset) = 0$. It is monotone: if $U \subseteq V$ are Scott-open, then $\max U \subseteq \max V$, so $v(U) = \lambda(\max U) \leq \lambda(\max V) = v(V)$. It is modular: if U and V are Scott-open, then

$$\begin{aligned} v(U \cup V) + v(U \cap V) &= \lambda(\max(U \cup V)) + \lambda(\max(U \cap V)) \\ &= \lambda(\max U \cup \max V) + \lambda(\max U \cap \max V) \\ &= \lambda(\max U) + \lambda(\max V) = v(U) + v(V). \end{aligned}$$

It also is continuous: if $(U_j)_{j \in J}$ is a directed family in $\Sigma(D)$, then $(\max U_j)_{j \in J}$ is a directed family in the lattice $\Omega(2^\omega)$ of opens of the Cantor space 2^ω , so, since 2^ω

²Recall that the Lebesgue measure is the ‘uniform’ measure on 2^ω : it is determined by assigning a ‘cylinder set’ $\llbracket a \rrbracket := \{x \in 2^\omega : a \leq x\}$ (for a finite binary string a) the measure $2^{-|a|}$ (see, e.g., Downey and Hirschfeldt 2010, sec. 1.2).

³This is well-defined, i.e., $f(x)$ is in D . Indeed, $f(x) : \omega \rightarrow 2$ is a partial function and if $n \in \omega$ is such that $f(x)(n)$ is not defined, then, for all $m \geq n$, also $f(x)(m)$ is not defined: Since $f(x)(n)$ is not defined, $x(2n)$ is not defined, so, since $2n \leq 2m$, also $x(2m)$ is not defined, so $f(x)(m)$ is not defined.

⁴An infinite sequence $x \in 2^\omega$ cannot be compact: $x = \bigvee_n x \upharpoonright n$ but, for no n , do we have $x \upharpoonright n \geq x$. And conversely, if $c \in 2^{<\omega}$ and $A \subseteq D$ is directed with $\bigvee A \geq c$, then A is a chain (if $a, a' \in A$, there is, by directedness, $a, a' \leq b \in A$, so a and a' are initial segments of b , whence $a \leq a'$ or $a' \leq a$), and $\bigvee A$ is the least binary sequence x extending all $a \in A$, so $x \geq c$ implies, since c is finite, that already some $a \in A$ extends c (otherwise, all $a \in A$ are, qua initial segments of $x \geq c$, also initial segments of c , so c is an upper bound of A , so $x \leq c$, whence $x = c$ is finite, so $x = \bigvee A$ simply is the maximal element of A , whence $x \in A$ with $x \geq c$, contradiction).

is a second-countable space,

$$v\left(\bigcup_J U_j\right) = \lambda\left(\max\bigcup_J U_j\right) = \lambda\left(\bigcup_J \max U_j\right) = \sup_J \left(\lambda(\max U_j)\right) = \sup_J \left(v(U_j)\right).^5$$

It is normalized: $v(D) = \lambda(\max D) = \lambda(2^\omega) = 1$. And it is max-normalized: we can write $\max D$ as the intersection $\bigcap_n U_n$ of opens $U_n = \bigcup_{a \in 2^n} \uparrow a$ with $v(U_n) = \lambda(\max U_n) = \lambda(2^\omega) = 1$.

Concerning the dynamics f : Here are two examples of how to compute $f(x)$

$$\begin{array}{ll} x = 01100 & y = 100011 \\ f(x) = (0+0)(1+1)(1+0) = 001 & f(y) = (1+1)(0+0)(0+1) = 001. \end{array}$$

In particular, f is not injective, and it also is not surjective: For any $x \in D$, $f(x)(0) = x(0) + x(0) = 0$, so the string $y = \langle 1 \rangle$ doesn't have an f -preimage. Moreover, $x = 000\dots$ is a fixpoint of f and $f(111\dots) = 000\dots$. In fact, there are infinitely many fixpoints: If $x \in 2^\omega$ is such that, for all even $n \in \omega$, $x(n) = 0$, then $f(x) = x$. (Proof: For $n \in \omega$, $2n$ is even, so $x(2n) = 0$, whence $f(x)(n) = x(n) + x(2n) = x(n)$.) Here are some more structural facts about f , that we'll also use in the proof of the proposition.

B.1.2. LEMMA. 1. If $x \leq y$ in D , then $f(x) \leq f(y)$.

2. $f : D \rightarrow D$ is Scott-continuous.

3. $f : D \rightarrow D$ is max-preserving.

4. For finite $x \in D$, $|f(x)| = \lfloor \frac{1}{2}|x| \rfloor$ where $\lfloor \cdot \rfloor$ is the floor function.⁶

5. For $i \leq \omega$ and $x \in D$, we have $f(x) \upharpoonright i \geq f(x \upharpoonright i)$.

PROOF. Ad (1). Let $x \leq y$. For $n \in \omega$, we need to show that if $f(x)(n)$ is defined, also $f(y)(n)$ is defined and has the same value. If $f(x)(n)$ is defined, then $x(2n)$ is defined, so $y(2n)$ is defined, so $f(y)(n)$ is defined. Moreover, $x(n) = y(n)$ and $x(2n) = y(2n)$, so $f(x)(n) = x(n) + x(2n) = y(n) + y(2n) = f(y)(n)$.

Ad (2). Let $A \subseteq D$ be directed, and show $f(\bigvee A) = \bigvee f(A)$. As noted above, A is a chain. And if A has a greatest element, the claim is immediate by monotonicity. So assume A doesn't have a greatest element. So A is infinite and

⁵Claim: If $(U_j)_J$ is a directed family of opens of the second-countable space X and μ a measure on $\mathcal{B}(X)$, then $\mu(\bigcup_J U_j) = \sup_J \mu(U_j)$. Proof: Since X is second-countable, there is a countable $J_0 = \{j_0, j_1, \dots\} \subseteq J$ with $\bigcup_J U_j = \bigcup_{j_0} U_j$. Define $V_n := \bigcup_{k=0}^n U_{j_k}$. So (V_n) is an increasing sequence of opens with $\bigcup_n V_n = \bigcup_J U_j$. So, by continuity from below (of measures), $\mu(\bigcup_J U_j) = \mu(\bigcup_n V_n) = \sup_n \mu(V_n)$. And we have $\sup_n \mu(V_n) = \sup_J \mu(U_j)$ because: (\leq) Given $V_n = U_{j_0} \cup \dots \cup U_{j_n}$, consider, by directedness, $U_j \supseteq U_{j_0}, \dots, U_{j_n}$, so $\mu(V_n) \leq \mu(U_j) \leq \sup_J \mu(U_j)$. (\geq) Given U_j , we have $\mu(U_j) \leq \mu(\bigcup_J U_j) = \sup_n \mu(V_n)$.

⁶I.e., for a real number r , $\lfloor r \rfloor$ is the biggest integer $\leq r$.

doesn't include a maximal element. So $A = a_0 < a_1 < \dots$ for $a_k \in 2^{<\omega}$. Thus, for $a := \bigvee A$ we have $a(n) = a_k(n)$ for k big enough such that $a_k(n)$ is defined. Moreover, for $b := \bigvee f(A)$ and $a_k \in A$ we have $f(a_k)(n) = b(n)$ whenever defined. Thus, for $n \in \omega$,

$$f(\bigvee A)(n) = a(n) + a(2n) = a_k(n) + a_k(2n)$$

for some big enough k . Further,

$$a_k(n) + a_k(2n) = f(a_k)(n) = b(n) = \bigvee f(A)(n).$$

Hence $f(\bigvee A) = \bigvee f(A)$, as needed.

Ad (3). If $x \in \max D = 2^\omega$, then also $f(x)$ is infinite (since $f(x)(n)$ is defined for every $n \in \omega$ because every $x(2n)$ is defined), so $f(x) \in 2^\omega = \max D$.

Ad (4). The length of $f(x)$ is n where n is maximal such that $f(x)(n-1)$ is defined. And $f(x)(n-1)$ is defined iff $x(2(n-1))$ is defined iff $2(n-1) < |x|$. So

$$|f(x)| = \max\{n : 2(n-1) < |x|\} = \max\{n : n \leq \frac{1}{2}|x|\} = \lfloor \frac{1}{2}|x| \rfloor.$$

Ad (5). We need to show for all $n \in \omega$: If $f(x \upharpoonright i)(n)$ is defined, then $f(x) \upharpoonright i(n)$ is defined and they have the same value. So assume that $f(x \upharpoonright i)(n)$ is defined. So $x \upharpoonright i(2n)$ is defined, whence $2n < i$. Since $x \upharpoonright i \leq x$, also $x(2n)$ is defined, so $f(x)(n)$ is defined. Since $n \leq 2n < i$, $f(x)(n) = f(x) \upharpoonright i(n)$ is defined. Its value is $f(x) \upharpoonright i(n) = f(x)(n) = x(n) + x(2n) = x \upharpoonright i(n) + x \upharpoonright i(2n) = f(x \upharpoonright i)(n)$. \square

PROOF OF PROPOSITION B.1.1. We need to write \mathfrak{D} as a restricted limit of a finitary dynamical expanding system. We do so in a sequence of steps.

Step 1. As index set we choose the even natural numbers $I := \{2n : n \in \omega\}$, which hence is countable. Working with even numbers (as opposed to all naturals) will make things more convenient later on.

Step 2. The finite domains are $\mathfrak{D}_i := (2^{\leq i}, v_i, f \upharpoonright 2^{\leq i})$ where $2^{\leq i}$ is the set of binary strings of length $\leq i$ (so $2^{\leq 0}$ is the singleton of the empty string) and $v_i : \Sigma(D_i) \rightarrow [0, 1]$ is defined by $v(U_i) := |U_i \cap 2^i|2^{-i}$.

These are indeed finite dynamical Scott domains: First, $2^{\leq i}$ is a nonempty finite partial order that is bounded complete and hence a finite Scott domain. Second, v_i is a valuation (strictness and monotonicity are immediate and for modularity: write $n := |U_i \cap 2^i|$, $m := |V_i \cap 2^i|$, $k := |(U_i \cup V_i) \cap 2^i|$, and $l := |(U_i \cap V_i) \cap 2^i|$, then $n+m = k+l$, so $v(U_i \cup V_i) + v(U_i \cap V_i) = k2^{-i} + l2^{-i} = n2^{-i} + m2^{-i} = v(U_i) + v(V_i)$). And it is continuous (since $\Sigma(D_i)$ is finite), normalized ($v(D_i) = |D_i \cap 2^i|2^{-i} = 2^i 2^{-i} = 1$), and max-normalized (since D_i is finite). Third, $f \upharpoonright 2^{\leq i}$ is well-defined (if $x \in 2^{\leq i}$, then $|f(x)| = \lfloor \frac{1}{2}|x| \rfloor \leq |x| \leq i$, so $f(x) \in 2^{\leq i}$) and monotone (qua restriction of a monotone function) and hence Scott-continuous.⁷

⁷Note that $2^{\leq i}$ is not max-reflective for $2 \leq i < \omega$: It satisfies the first axiom: If $a \in$

Step 3. The connecting and limiting morphisms will be defined as sequence restriction. We'll first collect some facts about it: For $0 \leq i \leq j \leq \omega$, define $r_{ij} : 2^{\leq j} \rightarrow 2^{\leq i}$ by $x \mapsto x \upharpoonright i$. We show that (a) r_{ij} is a Scott-continuous projection, (b) r_{ij} is max-preserving, and (c) r_{ij} is max-bisimulative.

Concerning (a), r_{ij} is monotone: if $x \leq y$, then $x \upharpoonright i \leq y \upharpoonright i$. As corresponding embedding $e : 2^{\leq i} \rightarrow 2^{\leq j}$ we choose $x \mapsto x$. This is monotone and, for $x \in 2^{\leq i}$, we have $r_{ij} \circ e(x) = x \upharpoonright i = x$, and, for $x \in 2^{\leq j}$, we have $e \circ r_{ij}(x) = x \upharpoonright i \leq x$. Finally, r_{ij} is Scott-continuous: If $A \subseteq 2^{\leq j}$ is directed, we need to show $r_{ij}(\bigvee A) \leq \bigvee r_{ij}(A)$ (the \geq -direction follows from monotonicity). If $i = \omega$, then r_{ij} is the identity function which is Scott-continuous. So let $i < \omega$. Then we have $r_{ij}(\bigvee A) = \bigvee A \upharpoonright i = a \upharpoonright i$ for some $a \in A$ (since $\bigvee A$ is above the compact $\bigvee A \upharpoonright i$, so there is $a \in A$ with $a \geq \bigvee A \upharpoonright i$). And $a \upharpoonright i \leq \bigvee_{a \in A} a \upharpoonright i = \bigvee r_{ij}(A)$.

Concerning (b), if $x \in \max 2^{\leq j}$, then $x \upharpoonright i$ has length i , so is in $\max 2^{\leq i}$.

Concerning (c), let $x \in 2^{\leq j}$ and $x \upharpoonright i \leq e \in \max 2^{\leq i}$, and find $d \in \max 2^{\leq j}$ with $x \leq d$ and $d \upharpoonright i = e$. If $|x| < i$, then choose $d := e0 \dots 0$ (with $j - i$ many 0's at the end), so $x \leq e \leq d \in 2^j$ and $d \upharpoonright i = e$. If $|x| \geq i$, then $x \upharpoonright i \in 2^i$ is maximal, so identical to e , so choose $d := x00 \dots$ (with as many 0's at the end to have length j). Then $x \leq d \in 2^j$ and $d \upharpoonright i = x \upharpoonright i = e$.

Step 4. The connecting morphisms are $p_{ij} := r_{ij} : 2^{\leq j} \rightarrow 2^{\leq i}$. As noted in step 3, these are Scott-continuous projections that are max-preserving and max-bisimulative. They are valuation-preserving: Let $U_i \in \Sigma(D_i)$ and show $v_j(p_{ij}^{-1}(U_i)) = v_i(U_i)$. If $U_i = \emptyset$, both sides equal 0, so let U_i be nonempty and write $U_i \cap 2^i = \{a_1, \dots, a_n\}$. Then $p_{ij}^{-1}(U_i) \cap 2^j$ is the disjoint union $\bigcup_{k=1}^n p_{ij}^{-1}(a_k) \cap 2^j$. Note that $|p_{ij}^{-1}(a_k) \cap 2^j| = |\{b \in 2^j : b \upharpoonright i = a_k\}| = 2^{j-i}$. So

$$\begin{aligned} v_j(p_{ij}^{-1}(U_i)) &= |p_{ij}^{-1}(U_i) \cap 2^j| 2^{-j} = \left(\sum_{k=1}^n |p_{ij}^{-1}(a_k) \cap 2^j| \right) 2^{-j} \\ &= (n 2^{j-i}) 2^{-j} = n 2^{-i} = v_i(U_i). \end{aligned}$$

And they are max-semi-equivariant: If $x \in \max 2^{\leq j}$, then $f(x) \upharpoonright i \geq f(x \upharpoonright i)$ by lemma B.1.2 (5). Moreover, they clearly satisfy the compatibility conditions $p_{ii} = \text{id}_{2^{\leq i}}$ and, for $i \leq j \leq k$, $p_{ik} = p_{ij} \circ p_{jk}$.

Step 5. The diagram is upward deterministic. Let $i \in I$ and $a_i, b_i \neq b'_i \in \max 2^{\leq i} = 2^i$ with $b_i, b'_i \geq f(a_i)$. Set $j := 2i \geq i$. Then, for $a_j, b_j \neq b'_j \in \max 2^{\leq j} = 2^j$, if $a_j \upharpoonright i = a_i$, $b_j \upharpoonright i = b_i$, and $b'_j \upharpoonright i = b'_i$, we cannot have $b_j, b'_j \geq f(a_j)$ since otherwise b_j and b'_j agree on the initial segment of length $|f(a_j)| = \frac{1}{2}|a_j| = i$, whence $b_i = b_j \upharpoonright i = b'_j \upharpoonright i = b'_i$, contradicting $b_i \neq b'_i$.

By now we've established that $(\mathfrak{D}_i, p_{ij})_I$ is a finitary dynamical expanding system. Next we construct the limit morphisms $p_i : \mathfrak{D} \rightarrow \mathfrak{D}_i$.

$2^{\leq i}$, then $a = \bigwedge \{a\sigma \in 2^i : \sigma \in 2^{i-|a|}\} = \bigwedge (\uparrow a \cap \max 2^{\leq i})$. But it doesn't satisfy the second: Let $A = \{0\sigma, 1\sigma\}$ for, say, $\sigma = 0 \dots 0 \in 2^{i-1}$. Then $\bigwedge A = \epsilon$ (the empty string), so $(\uparrow \bigwedge A) \cap \max 2^{\leq i} = 2^i \neq A$.

Step 6. The limiting morphisms are $p_i := r_{\omega i} : 2^{\leq \omega} \rightarrow 2^{\leq i}$. As noted in step 3, these are Scott-continuous projections that are max-preserving and max-bisimulative. They are valuation-preserving: For $U_i \in \Sigma(D_i)$, show $v(p_i^{-1}(U_i)) = v_i(U_i)$. If U_i is empty, this is immediate, so write $U_i \cap 2^i = \{a_1, \dots, a_n\}$. Then we have

$$\begin{aligned} v(p_i^{-1}(U_i)) &= \lambda(\max p_i^{-1}(U_i)) = \lambda(\{x \in 2^\omega : x \upharpoonright i \in U_i\}) = \lambda\left(\bigcup_{k=1}^n \llbracket a_k \rrbracket\right) \\ &= \sum_{k=1}^n \lambda(\llbracket a_k \rrbracket) = n2^{-i} = |U_i \cap 2^i|2^{-i} = v_i(U_i). \end{aligned}$$

And they are max-semi-equivariant: If $x \in \max 2^{\leq \omega}$, then $f(x) \upharpoonright i \geq f(x \upharpoonright i)$ by lemma B.1.2 (5). Moreover, they clearly satisfy the compatibility condition that, for $i \leq j$, $p_i = p_{ij} \circ p_j$.

Step 7. \mathfrak{D} is a max-preserving max-normalized dynamical Scott domain: As noted above, D is a Scott domain, v is a max-normalized continuous valuation, and f is a max-preserving and Scott-continuous function $f : D \rightarrow D$ by lemma B.1.2.

Step 8. Thus, (\mathfrak{D}, p_{ij}) is a cone to $(\mathfrak{D}_i, p_{ij})_I$ in \mathbf{dSCO}_n^p with \mathfrak{D} in \mathbf{dSCO}_{nm}^p . Write (\mathfrak{E}, q_i) for the restricted limit constructed in the limit theorem. By its universal property, there is a unique morphism $\beta : \mathfrak{D} \rightarrow \mathfrak{E}$ in \mathbf{dSCO}_n^p with $q_i \circ \beta = p_i$ (for all i), and β is given by $\beta(x) = \langle p_i(x) : i \in I \rangle$. If we can show that $\beta : D \rightarrow E$ is an order isomorphism, then we can conclude that β is an isomorphism (by proposition 5.2.13 in chapter 5), whence \mathfrak{D} is, like \mathfrak{E} , in \mathbf{dDOM} .

Step 9. β indeed is an order isomorphism. Since β is monotone, we only need to show that it is surjective and order-reflecting.

Order-reflecting: Let $x, y \in D$ with $\beta(x) \leq \beta(y)$, and show $x \leq y$. From the assumption, we have, for all $i \in I$, $x \upharpoonright i \leq y \upharpoonright i$. Thus, for any $n \in \omega$, if $x(n)$ is defined, let $i > n$ in I , so $x(n) = x \upharpoonright i(n) = y \upharpoonright i(n) = y(n)$, whence $y(n)$ is defined and has the same value as $x(n)$, so $x \leq y$.

Surjective: If $z \in E$, then $z = \langle z_i : i \in I \rangle$ is a sequence of extending finite binary strings.

Case 1: $z \in \max E$. Then each $q_i(x) = z_i$ is maximal (since the q_i are max-preserving). Let x be the (infinite) sequence extending all z_i . Then, for $i \in I$, $x \upharpoonright i = z_i$ since x extends z_i and z_i is of length i . So $\beta(x) = \langle x \upharpoonright i : i \in I \rangle = \langle z_i : i \in I \rangle = z$.

Case 2: $z \notin \max E$. Then there is a minimal $i \in I$ such that $x := z_i$ is not maximal, and we claim that then $z_i = z_j$ for all $j \geq i$ in I : Otherwise there is $j \geq i$ with $|z_j| > |z_i|$. So, since $z_i = p_{ij}(z_j) = z_j \upharpoonright i$, we get the contradiction

$$|z_i| = |z_j \upharpoonright i| = \min\left(\underbrace{|z_j|}_{>|z_i|}, \underbrace{i}_{>|z_i|}\right) > |z_i|.$$

Now, we have

$$z = \langle z_0, z_2, \dots, z_i, z_i, \dots \rangle = \langle x \upharpoonright 0, x \upharpoonright 2, \dots, x \upharpoonright i, x \upharpoonright i + 1, \dots \rangle = \beta(x),$$

as needed. □

B.2 More facts about the dynamical domain

A first indication that the dynamical domain \mathfrak{D} is more complex than its rather simple definition may suggest is that the topological system that it models is not a subshift (which form a particularly well-behaved class of topological systems).

B.2.1. PROPOSITION. *The Cantor dynamics $(2^\omega, f \upharpoonright 2^\omega)$ modeled by \mathfrak{D} is not expansive and hence not isomorphic to a subshift.*

PROOF. We need to show that $\bar{f} := f \upharpoonright 2^\omega$ is not expansive; then Hedlund's theorem (Hedlund 1969) implies that it cannot be isomorphic to a subshift. Let $\epsilon > 0$ and show that there are $x \neq y$ in 2^ω such that, for all $n \in \omega$, $d(\bar{f}^n(x), \bar{f}^n(y)) < \epsilon$. Indeed, let $N \in \omega$ be big enough such that N is odd and $2^{-N} < \epsilon$. Let $x := 000\dots$ and let y be like x except that it has a '1' at position N , so $x \neq y$. Since both x and y have value '0' for even n , they are, as observed above, fixpoints under f . So $d(\bar{f}^n(x), \bar{f}^n(y)) = d(x, y) = 2^{-N} < \epsilon$. □

We collect some more facts about the dynamics f that we'll use below.

B.2.2. LEMMA. *1. For $n \geq 0$ and $i \in I$, if $x \in D$ has length $2^n i$, then $|f^n(x)| = i$.*

2. For $i \geq 1$ and $x \in D$, $f(x \upharpoonright 2i) = f(x \upharpoonright 2i + 1)$.

3. For $i \in \omega$ and $x \in D$, we have $f(x \upharpoonright 2i) = f(x) \upharpoonright i$.

4. For $k \in \omega$, $i \in \omega$ and $x \in D$, we have $f^k(x \upharpoonright 2^k i) = f^k(x) \upharpoonright i$.

5. For $i \in \omega$ and $x \in D$, $|f(x \upharpoonright 2i)| \leq i$.

PROOF. Ad (1). By induction on n . If $n = 0$, then for $x \in D$ with length $2^n i = i$, we have $|f^n(x)| = |x| = i$. For $n + 1$, given $i \in I$ and $x \in D$ with length $2^{n+1}i$, define $y := f(x)$ which, by lemma B.1.2 (4), has length $\lfloor \frac{1}{2}|x| \rfloor = 2^n i$. So, by induction hypothesis, $|f^{n+1}(x)| = |f^n(y)| = i$.

Ad (2). If $|x| \leq 2i$, then $x \upharpoonright 2i = x \upharpoonright 2i + 1$ and the claim is immediate. So assume $|x| \geq 2i + 1$. Then, by lemma B.1.2 (4),

$$\begin{aligned} |f(x \upharpoonright 2i)| &= \lfloor \frac{1}{2}|x \upharpoonright 2i| \rfloor = \lfloor \frac{1}{2}2i \rfloor = i \\ &= \lfloor i + \frac{1}{2} \rfloor = \lfloor \frac{1}{2}(2i + 1) \rfloor = \lfloor \frac{1}{2}|x \upharpoonright 2i + 1| \rfloor = |f(x \upharpoonright 2i + 1)|. \end{aligned}$$

By monotonicity, $f(x \upharpoonright 2i) \leq f(x \upharpoonright 2i + 1)$, so the claim follows.

Ad (3). Case 1: $|x| < 2i$. Then $x = x \upharpoonright 2i$ and $|f(x)| = \lfloor \frac{1}{2}|x| \rfloor \leq \lfloor \frac{1}{2}2i \rfloor = i$, so $f(x \upharpoonright 2i) = f(x) = f(x) \upharpoonright i$.

Case 2: $|x| \geq 2i$. We have $x \upharpoonright 2i \leq x$, so, by monotonicity, $f(x \upharpoonright 2i) \leq f(x)$, whence $f(x \upharpoonright 2i) = f(x) \upharpoonright l$ where $l = |f(x \upharpoonright 2i)|$. Then the claim follows from $l = |f(x \upharpoonright 2i)| = \lfloor \frac{1}{2}|x \upharpoonright 2i| \rfloor = \lfloor \frac{1}{2}2i \rfloor = i$.

Ad (4). By induction on k . If $k = 0$, then $f^k(x \upharpoonright 2^k i) = x \upharpoonright i = f^k(x) \upharpoonright i$. For $k + 1$, we have

$$f^{k+1}(x \upharpoonright 2^{k+1}i) = f^k(f(x \upharpoonright 2^{k+1}i)).$$

Define $j := 2^k i$, then, by (3),

$$f(x \upharpoonright 2^{k+1}i) = f(x \upharpoonright 2j) = f(x) \upharpoonright j = f(x) \upharpoonright 2^k i$$

Hence

$$f^k(f(x \upharpoonright 2^{k+1}i)) = f^k(f(x) \upharpoonright 2^k i)$$

Set $y := f(x)$. Then by the induction hypothesis,

$$f^k(f(x) \upharpoonright 2^k i) = f^k(y \upharpoonright 2^k i) = f^k(y) \upharpoonright i = f^{k+1}(x) \upharpoonright i,$$

as needed.

Ad (5). By (3), $|f(x \upharpoonright 2i)| = |f(x) \upharpoonright i| \leq i$. □

B.3 Words on the components

The main task in calculating the max-entropy of \mathfrak{D} is to determine the cardinality of the set of ‘words’

$$W(n, i) = \left\{ \sigma \in (2^i)^n : \exists x \in 2^\omega . \forall k = 0, \dots, n-1 . f^k(x) \upharpoonright i = \sigma(k) \right\}$$

(recall p_i is the restriction $x \mapsto x \upharpoonright i$). We think of i as *block length* and n as *time* or number of iterations: an example is given in figure B.1.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	...
x	1	1	0	1	0	0	1	0	1	1	1	0	1	1	1	0	...
$f(x)$	0	1	0	0	1	1	0	1	...								
$f^2(x)$	0	1	1	0	...												

Figure B.1: Example of $x \in 2^\omega$ realizing the word $\sigma = (1101, 0100, 0110) \in (2^i)^n$ for block length $i = 4$ and iteration number $n = 3$.

We'll show that $|W(n, i)| = 2^{i + \frac{(n-1)i}{2}}$. But let's first build some intuition for why to suspect this. We think about it inductively for $n \geq 1$. If $n = 1$, then

$$W(1, i) = \left\{ \sigma \in (2^i)^1 : \exists x \in 2^\omega . \underbrace{f^0(x) \upharpoonright i}_{=x \upharpoonright i} = \sigma(0) \right\} = 2^i,$$

because every $\sigma = (a)$ with $a \in 2^i$ can be realized by, e.g., $x = a00\dots \in 2^\omega$ —hence $N_1 := |W(1, i)| = |2^i| = 2^i$. So given we know the number $N_n = |W(n, i)|$, what can we say about $N_{n+1} = |W(n+1, i)|$? Every element σ of $W(n+1, i)$ is an extension of some element, namely $\sigma \upharpoonright n$, of $W(n, i)$, so we ask: How many extensions does a $\sigma \in W(n, i)$ have to an element of $W(n+1, i)$? As an example, consider $\sigma = (1101, 0100, 0110) \in (2^i)^n$ from figure B.1: For which $a \in 2^i$ can σa be realized by some $x \in 2^\omega$? The first half of a —i.e., $a_l := a(0)a(1)$ —is fixed by the last element b of σ : To satisfy $f^3(x) \upharpoonright i = a$, we must have $a(0) = b(0) + b(0)$ and $a(1) = b(1) + b(2)$. But the second half of a —i.e., $a_r := a(2)a(3)$ —is unconstrained: we fill in the dots in line ' $f^2(x)$ ' with a bit k at position 4 such that $b(2) + k = a(2)$, and at position 6 with a bit l such that $b(3) + l = a(3)$, and make an arbitrary choice for positions 5 and 7. Similarly, we fill in the dots in line ' $f(x)$ ' now with 8 bits to realize the 4 bits we've just added. Then we further propagate this to line ' x ' filling in the dots there now with 16 bits appropriately. This then yields—after adding, say, infinitely many 0's—a new $x' \in 2^\omega$ that realizes σa . Thus, for each element of $W(n, i)$ we add $2^{\frac{i}{2}}$ many new ones, so $N_{n+1} = N_n 2^{\frac{i}{2}}$. In closed form: $N_n = 2^{i + \frac{(n-1)i}{2}}$.

Now for the formal proof. Before making this induction argument, we need some preparatory lemmas. We start by noting that to compute the output of f up to a certain length, we only need a certain length of the input (i.e., a precise version of continuity).

B.3.1. LEMMA. *For $i \in I$, $n \geq 1$, if $x \in 2^\omega$, then, for $x' := x \upharpoonright 2^{n-1}i$, we have, for $k = 0, \dots, n-1$, that $f^k(x') \upharpoonright i = f^k(x) \upharpoonright i$.*

PROOF. By lemma B.2.2 (4), $f^k(x) \upharpoonright i = f^k(x \upharpoonright 2^k i)$. Note that, since x is infinite, $f^k(x) \upharpoonright i$ has length i . Since $k \leq n-1$, we have $2^k i \leq 2^{n-1}i$, so $x \upharpoonright 2^k i \leq x \upharpoonright 2^{n-1}i$. Since f is monotone, also f^k is, so

$$f^k(x \upharpoonright 2^k i) \leq f^k(x \upharpoonright 2^{n-1}i).$$

Since $f^k(x) \upharpoonright i = f^k(x \upharpoonright 2^k i)$ has length i , we have

$$f^k(x) \upharpoonright i = f^k(x \upharpoonright 2^k i) = f^k(x \upharpoonright 2^{n-1} i) \upharpoonright i = f^k(x') \upharpoonright i,$$

as needed. \square

Next, we show that if we have a block $b \in 2^i$ which determines the first half of a block at the next stage, $f(b) = d_l \in 2^{\frac{i}{2}}$, then, for any desired second half d_r to this next block $d_l d_r$, we can find a second block $c \in 2^i$ following b that realizes this second half, $f(bc) = d_l d_r$.

B.3.2. LEMMA. *For $i \in I$, if $b \in 2^i$ and $d_l, d_r \in 2^{\frac{i}{2}}$ with $f(b) = d_l$, then there is $c \in 2^i$ such that $f(bc) = d_l d_r$.*

PROOF. Define $c \in 2^i$ by being 0 on odd indices and for indices $0 \leq 2n' < i$, let $c(2n')$ be a bit b such that $b(\frac{i}{2} + n') + b = d_r(n')$ (which always exists). By lemma B.2.2 (4), $f(bc)$ and $d_l d_r$ have the same length i . So we need to show, for $n = 0, \dots, i-1$ that $f(bc)(n) = d_l d_r(n)$.

Case 1: $n \leq \frac{i}{2} - 1$ (i.e., in the first half of the block length i). Then $2n < i$, so

$$f(bc)(n) = bc(n) + bc(2n) = b(n) + b(2n) = f(b)(n) = d_l(n) = d_l d_r(n).$$

Case 2: $n \geq \frac{i}{2}$ (i.e., in the second half of the block length i). So $n = \frac{i}{2} + n'$ for some $0 \leq n' < \frac{i}{2}$. Hence

$$\begin{aligned} f(bc)(n) &= bc(n) + bc(2n) = b(n) + c(2n - i) = b(\frac{i}{2} + n') + c(2n') \\ &= d_r(n') = d_l d_r(\frac{i}{2} + n') = d_l d_r(n), \end{aligned}$$

as needed. \square

Next, we boost this observation inductively:

B.3.3. LEMMA. *For all $n \geq 1$ and $i \in I$, if $b \in 2^{2^{n-1}i}$ and $d_l, d_r \in 2^{\frac{i}{2}}$ with $f^n(b) = d_l$, then there is $c \in 2^{2^{n-1}i}$ such that $f^n(bc) = d_l d_r$.*

PROOF. By induction on n . If $n = 1$, the claim is given by lemma B.3.2. So assume the claim holds for n and show it for $n + 1$. Let $i \in I$, $b \in 2^{2^n i}$, and $d_l, d_r \in 2^{\frac{i}{2}}$ with $f^{n+1}(b) = d_l$.

We have, by lemma B.2.2 (1), $b' := f^n(b) \in 2^i$, and we have $f(b') = f^{n+1}(b) = d_l \in 2^{\frac{i}{2}}$. So, by lemma B.3.2, there is $c' \in 2^i$ such that $f(b'c') = d_l d_r$.

We apply the induction hypothesis to $b'c'$: Set $j := 2i$. Then $b \in 2^{2^n i} = 2^{2^{n-1}j}$ and $b', c' \in 2^i = 2^{\frac{j}{2}}$ with $f^n(b) = b'$. So the induction hypothesis yields that there is $c \in 2^{2^{n-1}j} = 2^{2^n i}$ such that $f^n(bc) = b'c'$, so

$$f^{n+1}(bc) = f f^n(bc) = f(b'c') = d_l d_r,$$

as needed. □

Now we can characterize the elements of $W(n + 1, i)$ in terms of $W(n, i)$. We write $\text{last}(\sigma)$ to denote the last element of a nonempty finite sequence σ .

B.3.4. LEMMA. *Let $i \in I$, $n \geq 1$ and $\sigma \in (2^i)^{n+1}$. Then*

1. *If $\sigma \in W(n + 1, i)$, then $\sigma \upharpoonright n \in W(n, i)$ and $\sigma(n) = b_l b_r$ with $f(\sigma(n - 1)) = b_l \in 2^{\frac{i}{2}}$ and $b_r \in 2^{\frac{i}{2}}$.*
2. *If $\sigma' \in W(n, i)$ and $b_l, b_r \in 2^{\frac{i}{2}}$ such that $f(\text{last}(\sigma')) = b_l$, then $\sigma := \sigma' b_l b_r \in W(n + 1, i)$.*

PROOF. Ad (1). Since $\sigma \in W(n + 1, i)$, there is $x \in 2^\omega$ such that, for $k = 0, \dots, n$, $f^k(x) \upharpoonright i = \sigma(k)$. So, in particular, this is true for $k = 0, \dots, n - 1$. Hence $\sigma \upharpoonright n \in W(n, i)$. Since $\sigma(n) \in 2^i$ and i is even, we can write $\sigma(n) = b_l b_r$ for $b_l, b_r \in 2^{\frac{i}{2}}$. So it remains to show that $f(\sigma(n - 1)) = b_l$. Indeed, we have, $\sigma(n - 1) = f^{n-1}(x) \upharpoonright i$ and $\sigma(n) = f^n(x) \upharpoonright i$, so, by lemma B.2.2 (3),

$$f(\sigma(n - 1)) = f(f^{n-1}(x) \upharpoonright i) = f(f^{n-1}(x)) \upharpoonright \frac{i}{2} = f^n(x) \upharpoonright \frac{i}{2} = \sigma(n) \upharpoonright \frac{i}{2} = b_l.$$

Ad (2). To show that $\sigma \in W(n + 1, i)$, we need to find $y \in 2^\omega$ such that, for $k = 0, \dots, n$, we have $f^k(y) \upharpoonright i = \sigma(k)$.

Since $\sigma' \in W(n, i)$, there is $x \in 2^\omega$ such that, for $k = 0, \dots, n - 1$, we have $f^k(x) \upharpoonright i = \sigma'(k)$. By lemma B.2.2 (1) and lemma B.3.1, we have, for $x' := x \upharpoonright 2^{n-1}i$, that $|f^{n-1}(x')| = i$ and

$$f^{n-1}(x') = f^{n-1}(x') \upharpoonright i = f^{n-1}(x) \upharpoonright i = \sigma'(n - 1) =: b.$$

So

$$f^n(x') = f f^{n-1}(x') = f(b) = f(\text{last}(\sigma')) = b_l.$$

By lemma B.3.3 (choosing $b := x' \in 2^{2^{n-1}i}$, $d_l := b_l \in 2^{\frac{i}{2}}$, $d_r := b_r \in 2^{\frac{i}{2}}$), there is $y' \in 2^{2^{n-1}i}$ such that $f^n(x'y') = b_l b_r$.

Let y be any element of 2^ω extending $x'y'$. Then we have, for $k = 0, \dots, n - 1$, that $2^k i \leq 2^{n-1}i$, so $y \upharpoonright 2^k i = x' \upharpoonright 2^k i = x \upharpoonright 2^k i$, so, using lemma B.2.2 (4),

$$f^k(y) \upharpoonright i = f^k(y \upharpoonright 2^k i) = f^k(x \upharpoonright 2^k i) = f^k(x) \upharpoonright i = \sigma'(k) = \sigma(k).$$

For $k = n$, we have, using lemma B.2.2 (4),

$$f^n(y) \upharpoonright i = f^n(y \upharpoonright 2^n i) = f^n(x'y') = b_l b_r = \sigma(n),$$

as needed. □

Now, we can prove the announced result.

B.3.5. PROPOSITION. For $i \in I$ and $n \geq 1$, we have $|W(n, i)| = 2^{i + \frac{(n-1)i}{2}}$.

PROOF. By induction on n . If $n = 1$, then, as noted at the beginning of this subsection, $W(1, i) = |2^i| = 2^{i + \frac{(1-1)i}{2}}$. So assume the claim for n and show it for $n + 1$. Using lemma B.3.4, we show that the following function is a well-defined bijection:

$$b : W(n + 1, i) \rightarrow W(n, i) \times 2^{\frac{i}{2}}$$

$$\sigma \mapsto \left(\sigma \upharpoonright n, \sigma(n) \left[\frac{i}{2}, \dots, i - 1 \right] \right),$$

where, for a finite sequence a , we write $a[n, m] := a(n)a(n+1) \dots a(m)$.

Well-defined: Given such σ , we know, by lemma B.3.4 (1), that $\sigma \upharpoonright n \in W(n, i)$ and $\sigma(n)[i, \dots, 2i - 1] \in 2^{\frac{i}{2}}$.

Injective: If $\sigma \neq \sigma'$, there is a least $0 \leq k \leq n$ such that $\sigma(k) \neq \sigma'(k)$. If $k < n$, then $\sigma \upharpoonright n(k) \neq \sigma' \upharpoonright n(k)$, so $b(\sigma) \neq b(\sigma')$. If $k = n$, then $\sigma(n-1) = \sigma'(n-1)$ and $\sigma(n) \neq \sigma'(n)$. Note that

$$f(\sigma(n-1)) = f(\sigma(n-1) \upharpoonright i) = f(\sigma(n-1)) \upharpoonright \frac{i}{2} = \sigma(n) \upharpoonright \frac{i}{2},$$

and similarly for σ' . Hence

$$\sigma(n) \upharpoonright \frac{i}{2} = f(\sigma(n-1)) = f(\sigma'(n-1)) = \sigma'(n) \upharpoonright \frac{i}{2},$$

so, since $\sigma(n) \neq \sigma'(n)$, $\sigma(n)[\frac{i}{2}, \dots, i-1] \neq \sigma'(n)[\frac{i}{2}, \dots, i-1]$, so $b(\sigma) \neq b(\sigma')$.

Surjective: Given $(\sigma', b_r) \in W(n, i) \times 2^{\frac{i}{2}}$, define $b_l := f(\text{last}(\sigma'))$ and $\sigma := \sigma' b_l b_r$. Then, by lemma B.3.4 (2), $\sigma \in W(n+1, i)$ and by construction, $b(\sigma) = (\sigma', b_r)$.

Hence, using the induction hypothesis, we have

$$\begin{aligned} |W(n+1, i)| &= |W(n, i) \times 2^{\frac{i}{2}}| = |W(n, i)| |2^{\frac{i}{2}}| \\ &= 2^{i + \frac{(n-1)i}{2}} 2^{\frac{i}{2}} = 2^{i + \frac{(n-1)i}{2} + \frac{i}{2}} = 2^{i + \frac{(n+1-1)i}{2}}, \end{aligned}$$

as needed. □

B.4 Computing max-entropy

Now we can easily compute the max-entropy $m(\mathfrak{D})$ of the dynamical domain \mathfrak{D} .

B.4.1. PROPOSITION. $m(\mathfrak{D}) = \infty$.

PROOF. By definition, $m(\mathfrak{D}) = \sup_I \left(\lim_n \frac{1}{n} \log |W(n, i)| \right)$. We have, for $i \in I$,

$$\begin{aligned} \lim_n \frac{1}{n} \log |W(n, i)| &= \lim_{n \geq 1} \frac{1}{n} \log 2^{i + \frac{(n-1)i}{2}} \\ &= \lim_{n \geq 1} \frac{1}{n} \left(\left(i + \frac{(n-1)i}{2} \right) \log(2) \right) \\ &= \lim_{n \geq 1} \frac{1}{n} i \log(2) + \lim_{n \geq 1} \frac{1}{n} \frac{(n-1)i}{2} \log(2) \\ &= i \frac{\log(2)}{2} \lim_{n \geq 1} \frac{n-1}{n} \\ &= i \frac{\log(2)}{2}, \end{aligned}$$

where the last step follows since $\lim_{n \geq 1} \frac{n-1}{n} = \lim_{n \geq 1} \frac{n(1-\frac{1}{n})}{n} = \lim_{n \geq 1} (1 - \frac{1}{n}) = 1$. Hence $m(\mathfrak{D}) = \sup \left(i \frac{\log(2)}{2} : i \in I \right) = \infty$. \square

Note that, by theorem 6.4.2 (5), this also means that the modeled Cantor dynamics $(2^\omega, f \upharpoonright 2^\omega)$ has infinite topological entropy.

Further questions The next step would be to construct and study a preserved valuation. Going further, one may explore the connections to algorithmic randomness theory (Downey and Hirschfeldt 2010). This is usually formulated in Cantor space, and an important source of intuition is the idea of randomly sampling a subsequence (Van Lambalgen 1987). For example, given a sequence $x \in 2^\omega$ keep tossing a coin (thus constructing a sequence $y \in 2^\omega$) and build the subsequence $z \in 2^\omega$ by adding the n -th element of x iff the n -th coin toss landed heads (i.e., if $y(n) = 1$). The guiding intuition behind the function $f : 2^\omega \rightarrow 2^\omega$ is similar, although the original sequence x and the sampling sequence y are now ‘jumbled together’: given a sequence $x \in 2^\omega$ build the sequence $z = f(x)$ by choosing $x(n)$ if $y(n) = x(2n) = 0$ and choosing the dual otherwise.

Bibliography

- Abbes, S. and K. Keimel (2006). “Projective topology on bifinite domains and applications”. In: *Theoretical Computer Science* 365.3, pp. 171–183.
- Abrams, G., P. Ara, and M. Siles Molina (2017). *Leavitt Path Algebras*. Vol. 2191. Lecture Notes in Mathematics. Springer.
- Abramsky, S. (1991). “Domain theory in logical form”. In: *Annals of pure and applied logic* 51.1-2, pp. 1–77.
- Abramsky, S. and A. Brandenburger (11/2011). “The sheaf-theoretic structure of non-locality and contextuality”. In: *New Journal of Physics* 13.11, p. 113036.
- Abramsky, S. and R. Jagadeesan (1994). “Games and Full Completeness for Multiplicative Linear Logic”. In: *The Journal of Symbolic Logic* 59.2, pp. 543–574.
- Abramsky, S. and A. Jung (1994). “Domain Theory”. In: *Handbook of Logic in Computer Science*. Ed. by S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum. Corrected and expanded version available at <http://www.cs.bham.ac.uk/~axj/pub/papers/handy1.pdf> (accessed 24 January 2018). Oxford: Oxford University Press.
- Abramsky, S. and G. McCusker (1999). “Game Semantics”. In: *Computational Logic*. Ed. by U. Berger and H. Schwichtenberg. Vol. 165. NATO ASI Series. Proceedings of the NATO Advanced Study Institute on Computational Logic, held in Marktobendorf, Germany, July 29-August 10, 1997. Berlin, Heidelberg: Springer, pp. 1–55.
- Abramsky, S., R. Soares Barbosa, K. Kishida, R. Lal, and S. Mansfield (2015). “Contextuality, Cohomology and Paradox”. In: *24th EACSL Annual Conference on Computer Science Logic (CSL 2015)*. Ed. by S. Kreutzer. Vol. 41.

- Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, pp. 211–228.
- Abramsky, S. and S. Vickers (1993). “Quantales, observational logic and process semantics”. In: *Mathematical Structures in Computer Science* 3.2, pp. 161–227.
- Adadi, A. and M. Berrada (2018). “Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)”. In: *IEEE Access* 6, pp. 52138–52160.
- Adler, R. L., A. G. Konheim, and M. H. McAndrew (1965). “Topological Entropy”. In: *Transactions of the American Mathematical Society* 114.2, pp. 309–319.
- Alvarez-Manilla, M., A. Edalat, and N. Saheb-Djahromi (2000). “An extension result for continuous valuations”. In: *Journal of the London Mathematical Society* 61.2, pp. 629–640.
- Amodei, D., C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané (2016). *Concrete Problems in AI Safety*. arXiv: 1606.06565.
- An, G. (1996). “The Effects of Adding Noise During Backpropagation Training on a Generalization Performance”. In: *Neural Computation* 8.3, pp. 643–674.
- Anderson, J. R. (2015). *Cognitive Psychology and Its Implications*. 8th ed. New York, NY: Worth Publishers.
- Baez, J. C., K. Courser, and C. Vasilakopoulou (2021). *Structured versus Decorated Cospans*. arXiv: 2101.09363.
- Bahri, Y., J. Kadmon, J. Pennington, S. S. Schoenholz, J. Sohl-Dickstein, and S. Ganguli (2020). “Statistical Mechanics of Deep Learning”. In: *Annual review of condensed matter physics* 11.1, pp. 501–528.
- Baier, C. and J.-P. Katoen (2008). *Principles of Model Checking*. Cambridge, Massachusetts: The MIT Press.
- Bauer, A. (09/2000). “The Realizability Approach to Computable Analysis and Topology”. PhD thesis. School of Computer Science, Carnegie Mellon University.
- Beall, J., R. Brady, J. M. Dunn, A. P. Hazen, E. Mares, R. K. Meyer, G. Priest, G. Restall, R. Ripley, J. Slaney, and R. Sylvan (2012). “On the Ternary Relation and Conditionality”. In: *Journal of Philosophical Logic* 41, pp. 595–612.

- Becker, H. and A. S. Kechris (1996). *The Descriptive Set Theory of Polish Group Actions*. London Mathematical Society Lecture Note Series 232. Cambridge: Cambridge University Press.
- Beer, G. (1993). *Topologies on Closed and Closed Convex Sets*. Vol. 268. Mathematics and Its Applications. Dordrecht: Kluwer Academic Publishers.
- Behrisch, M., S. Kerkhoff, R. Pöschel, F. M. Schneider, and S. Siegmund (2017). “Dynamical Systems in Categories”. In: *Applied Categorical Structures* 25, pp. 29–57.
- Bennett, B. (1996). “Modal logics for qualitative spatial reasoning”. In: *Logic Journal of the IGPL* 4.1, pp. 23–45.
- Benthem, J. van (2004). “What One May Come to Know”. In: *Analysis* 64.2, pp. 95–105.
- (2014). *Logic in Games*. Cambridge, Massachusetts: The MIT Press.
- Benthem, J. van and G. Bezhanishvili (2007). “Modal Logics of Space”. In: *Handbook of Spatial Logics*. Ed. by M. Aiello, I. Pratt-Hartmann, and J. van Benthem. Dordrecht: Springer, pp. 217–298.
- Berger, A. and S. Siegmund (2003). “On the Gap Between Random Dynamical Systems and Continuous Skew Products”. In: *Journal of Dynamics and Differential Equations* 15, pp. 237–279.
- Besold, T. R., A. d’Avila Garcez, S. Bader, H. Bowman, P. Domingos, P. Hitzler, K.-U. Kuehnberger, L. C. Lamb, D. Lowd, P. Machado Vieira Lima, L. de Penning, G. Pinkas, H. Poon, and G. Zaverucha (2017). *Neural-Symbolic Learning and Reasoning: A Survey and Interpretation*. arXiv: 1711.03902.
- Bezhanishvili, G. and M. Gehrke (2005). “Completeness of S4 with respect to the real line: revisited”. In: *Annals of Pure and Applied Logic* 131, pp. 287–301.
- Birkhoff, G. (1973). *Lattice Theory*. Third Edition, Second Printing. Vol. XXV. Colloquium Publications. First Edition, 1940. Providence, Rhode Island: American Mathematical Society.
- Blackburn, P., M. Rijke, and Y. Venema (2001). *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge: Cambridge University Press.
- Blum, L., F. Cucker, M. Shub, and S. Smale (1998). *Complexity and Real Computation*. New York: Springer.

- Blum, L., M. Shub, and S. Smale (1989). “On a Theory of Computation and Complexity over the Real Numbers: NP-completeness, Recursive Functions and Universal Machines”. In: *Bulletin of the American Mathematical Society* 21.1, pp. 1–46.
- Bogachev, V. I. (2007b). *Measure Theory*. Vol. 1. Berlin, Heidelberg: Springer.
- (2007a). *Measure Theory*. Vol. 2. Berlin, Heidelberg: Springer.
- Borceux, F. (1994). *Handbook of Categorical Algebra 1: Basic Category Theory*. Encyclopedia of Mathematics and its Applications. Cambridge: Cambridge University Press.
- Borceux, F. and G. Janelidze (2001). *Galois Theories*. Vol. 72. Cambridge studies in advanced mathematics. Cambridge: Cambridge University Press.
- Bottou, L. (1998). “Online Algorithms and Stochastic Approximations”. In: *Online Learning and Neural Networks*. Ed. by D. Saad. Cambridge, UK: Cambridge University Press.
- Bournez, O. and M. L. Campagnolo (2008). “A Survey on Continuous Time Computations”. In: *New Computational Paradigms: Changing Conceptions of What is Computable*. Ed. by S. B. Cooper, B. Löwe, and A. Sorbi. New York, NY: Springer, pp. 383–423.
- Bracho, F. and M. Droste (1994). “Labelled domains and automata with concurrency”. In: *Theoretical Computer Science* 135.2, pp. 289–318.
- Brogaard, B. and J. Salerno (2019). “Fitch’s Paradox of Knowability”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Fall 2019. Metaphysics Research Lab, Stanford University.
- Brown, T. B., D. Mané, A. Roy, M. Abadi, and J. Gilmer (2017). *Adversarial Patch*. arXiv: 1712.09665.
- Brualdi, R. A. (2006). *Combinatorial matrix classes*. Encyclopedia of Mathematics and Its Applications 108. Cambridge: Cambridge University Press.
- Buckner, C. (2019). “Deep learning: A philosophical introduction”. In: *Philosophy Compass* 14.10, e12625.
- (2020). “Understanding adversarial examples requires a theory of artefacts for deep learning”. In: *Nature Machine Intelligence* 2, pp. 731–736.

- Cardone, F. (2021). “Games, Full Abstraction and Full Completeness”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2021. Metaphysics Research Lab, Stanford University.
- Carmantini, G. S. (03/2017). “Dynamical Systems Theory for Transparent Symbolic Computation in Neuronal Networks”. PhD thesis. University of Plymouth, Centre for Robotics and Neural Systems.
- Cate, B. t., D. Gabelaia, and D. Sustretov (2009). “Modal languages for topology: Expressivity and definability”. In: *Annals of Pure and Applied Logic* 159.1, pp. 146–170.
- Chalmers, D. J. (2011). “A Computational Foundation for the Study of Cognition”. In: *Journal of Cognitive Science* 12.4, pp. 323–357.
- Cheung, S., V. Darvariu, D. R. Ghica, K. Muroya, and R. N. S. Rowe (2018). “A Functional Perspective on Machine Learning via Programmable Induction and Abduction”. In: *Functional and Logic Programming. FLOPS 2018*. Ed. by J. Gallagher and M. Sulzmann. Vol. 10818. Lecture Notes in Computer Science. Cham: Springer, pp. 84–98.
- Cisse, M., P. Bojanowski, E. Grave, Y. Dauphin, and N. Usunier (2017). “Parseval Networks: Improving Robustness to Adversarial Examples”. In: *Proceedings of the 34th International Conference on Machine Learning*, pp. 854–863.
- Clerc, F., V. Danos, F. Dahlqvist, and I. Garnier (2017). “Pointless learning”. In: *Foundations of Software Science and Computation Structures (FoSSaCS 2017)*. Ed. by J. Esparza and A. Murawski. Vol. 10203. Lecture Notes in Computer Science. Long version available at <https://hal.archives-ouvertes.fr/hal-01429663v2/document> (accessed 5 Feb 2021). Berlin, Heidelberg: Springer, pp. 355–369.
- Cohn, A. G., B. Bennett, J. Gooday, and N. M. Gotts (1997). “Representing and reasoning with qualitative spatial relations about regions”. In: *Temporal and Spatial Reasoning*. Ed. by O. Stock. Dordrecht: Kluwer, pp. 97–134.
- Cybenko, G. (1989). “Approximations by superpositions of sigmoidal functions”. In: *Mathematics of Control, Signals, and Systems* 2.4, pp. 303–314.
- Dahlqvist, F., V. Danos, and I. Garnier (2016). “Giry and the Machine”. In: *Electronic Notes in Theoretical Computer Science* 325. The Thirty-second Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXII), pp. 85–110.

- Danos, V. and I. Garnier (2015). “Dirichlet is Natural”. In: *Electronic Notes in Theoretical Computer Science* 319. The 31st Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXI), pp. 137–164.
- D’Avila Garcez, A. S., L. C. Lamb, and D. M. Gabbay (2009). *Neural-Symbolic Cognitive Reasoning*. Cognitive Technologies. Berlin, Heidelberg: Springer.
- Dewhurst, J. (2016). “Computing Mechanisms and Autopoietic Systems”. In: *Computing and Philosophy: Selected Papers from IACAP 2014*. Ed. by V. C. Müller. Vol. 375. Studies in Epistemology, Logic, Methodology, and Philosophy of Science. Cham: Springer, pp. 17–26.
- Doran, D., S. Schulz, and T. R. Besold (2017). *What Does Explainable AI Really Mean? A New Conceptualization of Perspectives*. arXiv: 1710.00794.
- Downarowicz, T. (2011). *Entropy in Dynamical Systems*. Cambridge: Cambridge University Press.
- Downarowicz, T. and O. Karpel (2016). *Dynamics in dimension zero: A survey*. arXiv: 1610.02727.
- Downey, R. G. and D. R. Hirschfeldt (2010). *Algorithmic Randomness and Complexity*. New York: Springer.
- Droste, M. (1990). “Concurrency, automata and domains”. In: *Automata, Languages and Programming*. Ed. by M. S. Paterson. Berlin, Heidelberg: Springer, pp. 195–208.
- Eberlein, E. (1974). “A generator theorem for flows”. In: *Boletim da Sociedade Brasileira de Matemática* 5, pp. 45–50.
- Edalat, A. (1995a). “Domain theory and intergration”. In: *Theoretical Computer Science* 151, pp. 163–193.
- (1995b). “Dynamical Systems, Measures, and Fractals via Domain Theory”. In: *Information and Computation* 120.1, pp. 32–48.
- (1997). “Domains for Computation in Mathematics, Physics and Exact Real Arithmetic”. In: *The Bulletin of Symbolic Logic* 3.4, pp. 401–452.
- Edalat, A. and R. Heckmann (1998). “A computational model for metric spaces”. In: *Theoretical Computer Science* 193, pp. 55–73.
- Edalat, A. and M. B. Smyth (1993). “I-Categories as a framework for solving domain equations”. In: *Theoretical Computer Science* 115, pp. 77–106.

- Eilenberg, S. and S. MacLane (1945). “A general theory of natural equivalences”. In: *Trans. Amer. Math. Soc.* 58, pp. 231–294.
- Eisner, T., B. Farkas, M. Haase, and R. Nagel (2015). *Operator Theoretic Aspects of Ergodic Theory*. Graduate Texts in Mathematics 272. Cham: Springer.
- Engelking, R. (1989). *General Topology*. Revised and completed edition. Vol. 6. Sigma Series in Pure Mathematics. Berlin: Heldermann Verlag.
- Eslami, S. M. A., D. Jimenez Rezende, F. Besse, F. Viola, A. S. Morcos, M. Garnelo, A. Ruderman, A. A. Rusu, I. Danihelka, K. Gregor, D. P. Reichert, L. Buesing, T. Weber, O. Vinyals, D. Rosenbaum, N. Rabinowitz, H. King, C. Hillier, M. Botvinick, D. Wierstra, K. Kavukcuoglu, and D. Hassabis (2018). “Neural scene representation and rendering”. In: *Science* 360.6394, pp. 1204–1210.
- Evans, R. (03/2020). “Kant’s Cognitive Architecture”. PhD thesis. Imperial College London, Department of Computing.
- Evans, R. and E. Grefenstette (2018). “Learning Explanatory Rules from Noisy Data”. In: *Journal of Artificial Intelligence Research* 61, pp. 1–64.
- Eykholt, K., I. Evtimov, E. Fernandes, B. Li, A. Rahmati, C. Xiao, A. Prakash, T. Kohno, and D. Song (2017). *Robust Physical-World Attacks on Deep Learning Models*. arXiv: 1707.08945.
- Fajstrup, L., M. Raußen, and E. Goubault (2006). “Algebraic topology and concurrency”. In: *Theoretical Computer Science* 357.1-3, pp. 241–278.
- Fernández-Duque, D. (2010). “Absolute completeness of $S4_u$ for its measure-theoretic semantics”. In: *Advances in modal logic*. Ed. by L. Beklemishev, V. Goranko, and V. Shehtman. Vol. 8. London: College Publications, pp. 100–119.
- Fitch, F. B. (1963). “A Logical Analysis of Some Value Concepts”. In: *The Journal of Symbolic Logic* 28.2, pp. 135–142.
- Flagg, B. and R. Kopperman (1997). “Computational Models for Ultrametric Spaces”. In: *Electronic Notes in Theoretical Computer Science* 6, pp. 151–159.
- Fletcher, S. C. (2016). “Similarity, Topology, and Physical Significance in Relativity Theory”. In: *The British Journal for the Philosophy of Science* 67.2, pp. 365–389.
- Foley, R. (1993). *Working without a Net*. Oxford: Oxford University Press.

- Føllesdal, D. (2004). *Referential opacity and modal logic*. Reprint of the 1961 PhD thesis. New York: Routledge.
- Fong, B., D. Spivak, and R. Tuyeras (2019). “Backprop as Functor: A compositional perspective on supervised learning”. In: *34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*.
- Foreman, M., A. S. Kechris, A. Louveau, and B. Weiss, eds. (2000). *Descriptive Set Theory and Dynamical System*. London Mathematical Society Lecture Note Series 277. Cambridge: Cambridge University Press.
- Franek, P., S. Ratschan, and P. Zgliczynski (2016). “Quasi-decidability of a Fragment of the First-Order Theory of Real Numbers”. In: *Journal of Automated Reasoning* 57.2, pp. 157–185.
- Franklin, J. N. and H. Towsner (2014). “Randomness and Non-Ergodic Systems”. In: *Moscow Mathematical Journal* 14.4. Also see <https://arxiv.org/abs/1206.2682> (last checked 2 July 2021), pp. 711–744.
- Fredkin, E. and T. Toffoli (1982). “Conservative logic”. In: *International Journal of Theoretical Physics* 21, pp. 219–253.
- Fritz, T. (2011). “Categories of fractions revisited”. In: *Morfismos*, 15.2. Also see <https://arxiv.org/abs/0803.2587v2> (accessed 9 Feb 2021), pp. 19–38.
- Gabriel, P. and M. Zisman (1967). *Calculus of Fractions and Homotopy Theory*. Vol. 35. *Ergebnisse der Mathematik und ihrer Grenzgebiete*. New York: Springer-Verlag.
- Gács, P., M. Hoyrup, and C. Rojas (2011). “Randomness on Computable Probability Spaces—A Dynamical Point of View”. In: *Theory of Computing Systems* 48, pp. 465–485.
- Gagné, N. and P. Panangaden (2018). “A Categorical Characterization of Relative Entropy on Standard Borel Spaces”. In: *Electronic Notes in Theoretical Computer Science* 336. The Thirty-third Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXIII), pp. 135–153.
- Galatolo, S., M. Hoyrup, and C. Rojas (2010). “Effective symbolic dynamics, random points, statistical behavior, complexity and entropy”. In: *Information and Computation* 208.1, pp. 23–41.
- Gambaudo, J.-M. and M. Martens (2006). “Algebraic Topology for Minimal Cantor Sets”. In: *Annales Henri Poincaré* 7, pp. 423–446.

- Gandy, R. (1980). “Church’s Thesis and Principles for Mechanisms”. In: *The Kleene Symposium*. Ed. by J. Barwise, H. J. Keisler, and K. Kunen. Amsterdam: North-Holland Publishing Company, pp. 123–148.
- Gärdenfors, P. (2004). *Conceptual Spaces: The Geometry of Thought*. A Bradford book. MIT Press.
- Garner, W. R. (1974/2014). *The Processing of Information and Structure*. New York, NY: Psychology Press.
- Gelder, T. van (1998). “The dynamical hypothesis in cognitive science”. In: *Behavioral and Brain Sciences* 21.5, pp. 615–628.
- Gibson, J. J. (1979/2015). *The Ecological Approach to Visual Perception*. Classic Edition. New York, NY: Psychology Press.
- Gierz, G., K. H. Hoffman, K. Keimel, J. D. Lawson, M. Mislove, and D. S. Scott (2003). *Continuous Lattices and Domains*. Cambridge: Cambridge University Press.
- Giunti, M. (1997). *Computation, Dynamics, and Cognition*. New York, Oxford: Oxford University Press.
- Glabbeek, R. van and P. Höfner (2018). “Progress, Justness and Fairness”. In: arXiv: 1810.07414v1.
- Glasner, E. and B. Weiss (2006). “On the Interplay between Measurable and Topological Dynamics”. In: *Handbook of Dynamical Systems*. Ed. by B. Hasselblatt and A. Katok. Vol. 1. B. Elsevier. Chap. 10, pp. 597–648.
- Goebel, R., A. Chander, K. Holzinger, F. Lecue, Z. Akata, S. Stumpf, P. Kieseberg, and A. Holzinger (2018). “Explainable AI: The New 42?” In: *Machine Learning and Knowledge Extraction. CD-MAKE 2018*. Ed. by A. Holzinger, P. Kieseberg, A. Tjoa, and E. Weippl. Vol. 11015. Lecture Notes in Computer Science. Cham: Springer, pp. 295–303.
- Gogioso, S., M. E. Stasinou, and B. Coecke (2020). *Functorial evolution of quantum fields*. arXiv: 2003.13271.
- Gold, E. M. (1967). “Language identification in the limit”. In: *Information and Control* 10.5, pp. 447–474.
- Goodfellow, I. J., J. Shlens, and C. Szegedy (2014). *Explaining and Harnessing Adversarial Examples*. arXiv: 1412.6572.

- Goubault, E. and T. P. Jensen (1992). “Homology of higher dimensional automata”. In: *CONCUR '92*. Ed. by W. Cleaveland. Berlin, Heidelberg: Springer, pp. 254–268.
- Goubault-Larrecq, J. (2018). *Products and Projective Limits of Continuous Valuations on T_0 Spaces*. arXiv: 1803.05259.
- Gowal, S., C. Qin, J. Uesato, T. Mann, and P. Kohli (2020). *Uncovering the Limits of Adversarial Training against Norm-Bounded Adversarial Examples*. arXiv: 2010.03593.
- Graben, P. beim (2004). “Incompatible Implementations of Physical Symbol Systems”. In: *Mind and Matter* 2.2, pp. 29–51.
- Hamkins, J. D. and B. Löwe (2008). “The modal logic of forcing”. In: *Transactions of the American Mathematical Society* 360.4, pp. 1793–1817.
- Hamkins, J. D. and W. A. Wołoszyn (2020). *Modal model theory*. arXiv: 2009.09394.
- Hawking, S. W. (1971). “Stable and Generic Properties in General Relativity”. In: *General Relativity and Gravitation* 1.4, pp. 393–400.
- Hedlund, G. A. (1969). “Endomorphisms and automorphisms of the shift dynamical system.” In: *Math. Systems Theory* 3, pp. 320–375.
- Hilbert, D. (1926). “Über das Unendliche”. In: *Mathematische Annalen* 95, pp. 161–190.
- Hildebrandt, T. and V. Sassone (1997). “Transition Systems with Independence and Multi-Arcs”. In: *BRICS Report Series* 4.10.
- Hjorth, G. and M. Molberg (2006). “Free continuous actions on zero-dimensional spaces”. In: *Topology and its Applications* 153, pp. 1116–1131.
- Hodges, W. and J. Väänänen (2019). “Logic and Games”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Fall 2019. Metaphysics Research Lab, Stanford University.
- Hoppensteadt, F. (2006). “Predator-prey model”. In: *Scholarpedia* 1.10, p. 1563.
- Hornik, K., M. Stinchcombe, and H. White (1989). “Multilayer feedforward networks are universal approximators”. In: *Neural Networks* 2.5, pp. 359–366.

- Hornischer, L. (2019). “Toward a Logic for Neural Networks”. In: *The Logica Yearbook 2018*. Ed. by I. Sedlár and M. Blicha. London: College Publications, pp. 133–148.
- (2020). “Logics of Synonymy”. In: *Journal of Philosophical Logic* 49, pp. 767–805.
- (2021). “The Logic of Information in State Spaces”. In: *The Review of Symbolic Logic* 14.1, pp. 155–186.
- Hoyrup, M. and C. Rojas (2009). “Computability of probability measures and Martin-Löf randomness over metric spaces”. In: *Information and Computation* 207.7, pp. 830–847.
- Huang, X., M. Kwiatkowska, S. Wang, and M. Wu (2017). “Safety Verification of Deep Neural Networks”. In: *Computer Aided Verification (CAV 2017)*. Ed. by R. Majumdar and V. Kunčák. Vol. 10426. Lecture Notes in Computer Science. Cham: Springer.
- Huber, F. (2016). “Formal Representations of Belief”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2016. Metaphysics Research Lab, Stanford University.
- Ichikawa, J. J. and M. Steup (2018). “The Analysis of Knowledge”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Summer 2018. Metaphysics Research Lab, Stanford University.
- Izhikevich, E. M. and R. FitzHugh (2006). “FitzHugh-Nagumo model”. In: *Scholarpedia* 1.9, p. 1349.
- Jacobs, B. and D. Sprunger (2019). “Neural Nets via Forward State Transformation and Backward Loss Transformation”. In: *Electronic Notes in Theoretical Computer Science* 347. Proceedings of the Thirty-Fifth Conference on the Mathematical Foundations of Programming Semantics, pp. 161–177.
- Jech, T. J. (1973). *The Axiom of Choice*. Amsterdam: North-Holland Publishing Company.
- Johnstone, P. T. (1982). *Stone Spaces*. Cambridge studies in advanced mathematics 3. Cambridge: Cambridge University Press.
- Jones, C. and G. D. Plotkin (1989). “A probabilistic powerdomain of evaluations”. In: *Proceedings. Fourth Annual Symposium on Logic in Computer Science*. IEEE Computer Society, pp. 186–195.

- Joyal, A., M. Nielsen, and G. Winskel (1996). “Bisimulation from Open Maps”. In: *Information and Computation* 127.2, pp. 164–185.
- Kashiwara, M. and P. Schapira (2006). *Categories and Sheaves*. Vol. 332. Grundlehren der mathematischen Wissenschaften. Berlin, Heidelberg: Springer-Verlag.
- Katz, S. and D. Peled (1992). “Defining conditional independence using collapses”. In: *Theoretical Computer Science* 101.2, pp. 337–359.
- Kechris, A. S. (1995). *Classical Descriptive Set Theory*. New York: Springer.
- Keimel, K. and J. D. Lawson (2005). “Measure extension theorems for T₀-spaces”. In: *Topology and its Applications* 149.1, pp. 57–83.
- Kratsios, A. (2020). *Characterizing the Universal Approximation Property*. arXiv: 1910.03344.
- Kremer, P. and G. Mints (2007). “Dynamic Topological Logic”. In: *Handbook of Spatial Logics*. Ed. by M. Aiello, I. Pratt-Hartmann, and J. van Benthem. Dordrecht: Springer, pp. 565–606.
- Kucharski, P. (2020). *Graph covers of higher dimensional dynamical systems*. arXiv: 2011.07299.
- Kuper, L., G. Katz, J. Gottschlich, K. Julian, C. Barrett, and M. Kochenderfer (2018). *Toward Scalable Verification for Safety-Critical Deep Networks*. arXiv: 1801.05950.
- Kurz, A. and J. Velebil (2016). “Relation lifting, a survey”. In: *Journal of Logical and Algebraic Methods in Programming* 85.4, pp. 475–499.
- Lambalgen, M. van (1987). “Random Sequences”. PhD thesis. Amsterdam: University of Amsterdam.
- Lamport, L. (2015). “The Computer Science of Concurrency: The Early Years”. In: *Communications of the ACM* 58.6, pp. 71–76.
- Lando, T. (2012). “Completeness of S₄ for the Lebesgue Measure Algebra”. In: *Journal of Philosophical Logic* 41, pp. 287–316.
- (2015). “First order S₄ and its measure-theoretic semantics”. In: *Annals of Pure and Applied Logic* 166.2, pp. 187–218.

- Lawson, J. D. (1982). “Valuations on continuous lattices”. In: *Continuous Lattices and Related Topics*. Ed. by R.-E. Hoffmann. Vol. 27. Mathematik-Arbeitspapiere. Universität Bremen, pp. 204–225.
- Lawson, J. (1997). “Spaces of maximal points”. In: *Math. Struct. in Comp. Science* 7, pp. 543–555.
- Lawvere, F. W. (1969). “Adjointness in foundations”. In: *Dialectica* 23. Republished in: *Reprints in Theory and Applications of Categories*, No. 16 (2006) pp. 1–16, pp. 281–296.
- (1986). “Taking categories seriously”. In: *Revista Colombiana de Matematicas* XX. Republished in: *Reprints in Theory and Applications of Categories*, No. 8 (2005) pp. 1–24, pp. 147–178.
- Leinster, T. (2014). *Basic Category Theory*. Vol. 143. Cambridge Studies in Advanced Mathematics. Cambridge: Cambridge University Press.
- Leitgeb, H. (2005). “Interpreted Dynamical Systems and Qualitative Laws: from Neural Networks to Evolutionary Systems”. In: *Synthese* 146.1, pp. 189–202.
- (2015). “The Humean Thesis on Belief”. In: *Proceedings of the Aristotelian Society* 89.1, pp. 143–185.
- (2017). *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford: Oxford University Press.
- (2018). “Stability and the Lottery Paradox”. Retrived from https://www.researchgate.net/publication/331533002_Stability_and_the_Lottery_Paradox (last checked 2 July 2021).
- Lewis, D. K. (1973). *Counterfactuals*. Cambridge, Mass.: Harvard University Press.
- Lind, D. and B. Marcus (1995). *An Introduction to Symbolic Dynamics and Coding*. Cambridge: Cambridge University Press.
- Lloyd, S. (2000). “Ultimate physical limits to computation”. In: *Nature* 406, pp. 1047–1054.
- Mac Lane, S. (1998). *Categories for the Working Mathematician*. 2nd ed. First edition from 1971. New York: Springer-Verlag.

- MacFarlane, J. (2017). “Logical Constants”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Winter 2017. Metaphysics Research Lab, Stanford University.
- Mackey, G. W. (1957). “Borel structure in groups and their duals”. In: *Transactions of the American Mathematical Society* 85, pp. 134–165.
- MacLennan, B. J. (2009). “Analog Computation”. In: *Encyclopedia of Complexity and Systems Science*. Ed. by R. A. Meyers. New York, NY: Springer, pp. 271–294.
- Manna, Z. and A. Pnueli (1991). *The Temporal Logic of Reactive and Concurrent Systems: Specification*. New York: Springer-Verlag.
- Mares, E. (2020). “Relevance Logic”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Winter 2020. Metaphysics Research Lab, Stanford University.
- Marr, D. (1982/2010). *Vision*. First published in 1982. Cambridge, Massachusetts: The MIT Press.
- Martin, K. and P. Panangaden (2011). “Domain Theory and General Relativity”. In: *New Structures for Physics*. Ed. by B. Coecke. Vol. 813. Lecture Notes in Physics. Berlin, Heidelberg: Springer, pp. 687–703.
- Martin, K. (1998). “Domain theoretic models of topological spaces”. In: *Electronic Notes in Theoretical Computer Science* 13. Comprox III, Third Workshop on Computation and Approximation, pp. 173–181.
- Meyn, S. and R. L. Tweedie (2009). *Markov Chains and Stochastic Stability*. 2nd ed. Cambridge: Cambridge University Press.
- Milne, T. (2019). *Piecewise Strong Convexity of Neural Networks*. arXiv: 1810.12805.
- Munkres, J. R. (2000). *Topology*. 2nd ed. Upper Saddle River, NJ: Prentice Hall.
- Murdoch, W. J., C. Singh, K. Kumbier, R. Abbasi-Asl, and B. Yu (2019). “Definitions, methods, and applications in interpretable machine learning”. In: *Proceedings of the National Academy of Sciences of the United States of America* 116.44, pp. 22071–22080.
- Naitzat, G., A. Zhitnikov, and L.-H. Lim (2020). “Topology of Deep Neural Networks”. In: *Journal of Machine Learning Research* 21.184, pp. 1–40.

- Neelakantan, A., L. Vilnis, Q. V. Le, I. Sutskever, L. Kaiser, K. Kurach, and J. Martens (2015). *Adding Gradient Noise Improves Learning for Very Deep Networks*. arXiv: 1511.06807.
- Nguyen, L. M., P. H. Nguyen, P. Richtárik, K. Scheinberg, M. Takáč, and M. van Dijk (2019). “New Convergence Aspects of Stochastic Gradient Algorithms”. In: *Journal of Machine Learning Research* 20.176, pp. 1–49.
- Niefield, S. B. (1996). “An Algebraic Approach to Chaos”. In: *Applied Categorical Structures* 4, pp. 423–441.
- Nozick, R. (2001). *Invariances: The Structure of the Objective World*. Cambridge, Massachusetts: Harvard University Press.
- Ong, C.-H. L. (1995). “Correspondence between operational and denotational semantics: the full abstraction problem for PCF”. In: *Handbook of Logic in Computer Science*. Ed. by S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum. Vol. 4. Oxford: Clarendon Press, pp. 269–356.
- Orponen, P. (1997). “A Survey of Continuous-Time Computation Theory”. In: *Advances in Algorithms, Languages, and Complexity*. Ed. by D.-Z. Du and K.-I. Ko. Boston, MA.: Springer.
- Osherson, D. and S. Weinstein (2008). “Recognizing Strong Random Reals”. In: *The Review of Symbolic Logic* 1.1, pp. 56–63.
- Oxtoby, J. C. (1980). *Measure and Category: A Survey of the Analogies between Topological and Measure Spaces*. 2nd ed. New-York: Springer-Verlag.
- Pacuit, E. (2017). *Neighborhood Semantics for Modal Logic*. Cham: Springer.
- Penkov, S. and S. Ramamoorthy (2017). *Using Program Induction to Interpret Transition System Dynamics*. arXiv: 1708.00376.
- Petersen, K. (1983). *Ergodic theory*. Cambridge Studies in Advanced Mathematics 2. Cambridge University Press.
- Piccinini, G. (2015). *Physical Computation: A Mechanistic Account*. Oxford: Oxford University Press.
- (2017). “Computation in Physical Systems”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Summer 2017. Metaphysics Research Lab, Stanford University.

- Pin, J.-E. (2009). “Profinite Methods in Automata Theory”. In: *26th International Symposium on Theoretical Aspects of Computer Science, STACS 2009, February 26-28, 2009, Freiburg, Germany, Proceedings*. Ed. by S. Albers and J. Marion. Vol. 3. LIPIcs. Freiburg, Germany: Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, pp. 31–50.
- Pitowsky, I. (1990). “The Physical Church Thesis and Physical Computational Complexity”. In: *The Jerusalem Philosophical Quarterly* 39, pp. 81–99.
- (2002). “Quantum speed-up of computations”. In: *Philosophy of Science* 69.S3, S168–S177.
- Plotkin, G. D. (1976). “A Powerdomain Construction”. In: *SIAM Journal on Computing* 5, pp. 452–487.
- (1981/2004). “A structural approach to operational semantics”. In: *The Journal of Logic and Algebraic Programming* 60–61. First published in 1981 and known as the ‘Aarhus Notes’, pp. 17–139.
- Porter, T. (1994). “Categorical shape theory as a formal language for pattern recognition?” In: *Annals of Mathematics and Artificial Intelligence* 10, pp. 25–54.
- Pour-El, M. B. and J. I. Richards (1989). *Computability in Analysis and Physics*. Perspectives in Mathematical Logic. Berlin, Heidelberg: Springer.
- Pour-El, M. B. and I. Richards (1981). “The wave equation with computable initial data such that its unique solution is not computable”. In: *Advances in Mathematics* 39.3, pp. 215–239.
- Pratt, V. (1991). “Modeling Concurrency with Geometry”. In: *Proceedings of the 18th ACM Symposium on Principles of Programming Languages*. POPL ’91. New York, NY: Association for Computing Machinery, pp. 311–322.
- Priest, G. and R. Sylvan (1992). “Simplified Semantics for Basic Relevant Logics”. In: *Journal of Philosophical Logic* 21, pp. 217–232.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Quine, W. V. O. (1960). *Word and Object*. Cambridge, Mass.: MIT Press.
- Rabinowitz, D. (2020). “The Safety Condition for Knowledge”. In: *The Internet Encyclopedia of Philosophy*. Available at <https://www.iep.utm.edu/safety-c/> (accessed 27 January 2020).

- Ratschan, S. (2010). “Safety Verification of Non-linear Hybrid Systems Is Quasi-Semidecidable”. In: *Theory and Applications of Models of Computation. TAMC 2010*. Ed. by J. Kratochvíl, A. Li, J. Fiala, and P. Kolman. Vol. 6108. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, pp. 397–408.
- Rédei, M. and C. Werndl (2012). “On the history of the isomorphism problem of dynamical systems with special regard to von Neumann’s contribution”. In: *Archive for history of exact sciences* 66.1, pp. 71–93.
- Reimann, M. W., M. Nolte, M. Scolamiero, K. Turner, R. Perin, G. Chindemi, P. Dłotko, R. Levi, K. Hess, and H. Markram (2017). “Cliques of Neurons Bound into Cavities Provide a Missing Link between Structure and Function”. In: *Frontiers in Computational Neuroscience* 11, p. 48.
- Restall, G. (1993). “Simplified semantics for relevant logics (and some of their rivals)”. In: *Journal of Philosophical Logic* 22, pp. 481–511.
- Restall, G. and R. Tony (2009). “On Permutation in Simplified Semantics”. In: *Journal of Philosophical Logic* 38.3, pp. 333–341.
- Ribes, L. (2017). *Profinite Graphs and Groups*. Vol. 66. Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge / A Series of Modern Surveys in Mathematics. Cham: Springer.
- Rossman, B. (2008). “Homomorphism Preservation Theorems”. In: *Journal of the ACM* 55.3.
- Rott, H. (2004). “Stability, Strength and Sensitivity: Converting Belief into Knowledge”. In: *Erkenntnis* 61, pp. 469–493.
- Rowland, E. and R. Yassawi (2017). “Profinite automata”. In: *Advances in Applied Mathematics* 85, pp. 60–83.
- Rue, T. de la (1993). “Espaces de Lebesgue”. In: *Séminaire de Probabilités XXVII*. Ed. by J. Azéma, P. A. Meyer, and M. Yor. Vol. 1557. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer-Verlag, pp. 15–21.
- Russell, B. (1919). *Introduction to Mathematical Philosophy*. London: Routledge.
- (1948). *Human Knowledge: Its Scope and its Limits*. London: Allen & Unwin.
- Saad, D. and S. A. Solla (1996). “Dynamics of on-line gradient descent learning for multilayer neural networks”. In: *Advances in neural information processing systems 8 (NIPS 1995)*, pp. 302–308.

- Samek, W., G. Montavon, A. Vedaldi, L. K. Hansen, and K.-R. Müller, eds. (2019). *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Cham.
- Sangiorgi, D. (2012). *Introduction to Bisimulation and Coinduction*. Cambridge: Cambridge University Press.
- Sassone, V., M. Nielsen, and G. Winskel (1996). “Models for concurrency: towards a classification”. In: *Theoretical Computer Science* 170.1-2, pp. 297–348.
- Saxe, A. M., J. L. McClelland, and S. Ganguli (2014). *Exact solutions to the nonlinear dynamics of learning in deep linear neural networks*. arXiv: 1312.6120.
- Schultz, P., D. I. Spivak, and C. Vasilakopoulou (2020). “Dynamical Systems and Sheaves”. In: *Applied Categorical Structures* 28, pp. 1–57.
- Schweber, N. (2016). *Existence of a cofinal sequence in a countable directed set*. Mathematics Stack Exchange. URL: <https://math.stackexchange.com/q/1915655> (version: 2016-09-05).
- Scott, B. M. (2013). *Cardinality of the set of clopen subsets of a topological space*. Mathematics Stack Exchange. <https://math.stackexchange.com/q/290890> (version: 2013-01-30).
- Scott, D. S. (1970). *Outline of a Mathematical Theory of Computation*. Tech. rep. PRG02. Available at <https://www.cs.ox.ac.uk/files/3222/PRG02.pdf> (last checked 24 January 2018). Oxford University Computing Laboratory.
- (1982). “Domains for denotational semantics”. In: *Automata, Languages and Programming. ICALP 1982*. Ed. by M. Nielsen and E. M. Schmidt. Vol. 140. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer.
- Shehtman, V. (1999). ““Everywhere” and “here””. In: *Journal of Applied Non-Classical Logics* 9.2-3, pp. 369–379.
- Shimomura, T. (2014). “Special homeomorphisms and approximation for Cantor systems”. In: *Topology and its Applications* 161, pp. 178–195.
- (2020). “Bratteli–Vershik models and graph covering models”. In: *Advances in Mathematics* 367, p. 107127.
- Sieg, W. (2002). “Calculations by man and machine: conceptual analysis”. In: *Reflections on the foundations of mathematics: Essays in honor of Solomon*

- Feferman*. Ed. by W. Sieg, R. Sommer, and C. Talcott. Vol. 15. Lecture Notes in Logic. Natick, Massachusetts: Association for Symbolic Logic, pp. 390–409.
- Siegelmann, H. T. and S. Fishman (1998). “Analog computation with dynamical systems”. In: *Physica D: Nonlinear Phenomena* 120.1-2, pp. 214–235.
- Šíma, J. and P. Orponen (2003). “General-Purpose Computation with Neural Networks: A Survey of Complexity Theoretic Results”. In: *Neural Computation* 15, pp. 2727–2778.
- Smith, P. (Unpublished). *The Galois Connection between Syntax and Semantics*. Available at <https://www.logicmatters.net/resources/pdfs/Galois.pdf> (last checked 1 May 2021).
- Smolensky, P. (1988). “On the proper treatment of connectionism”. In: *Behavioral and brain sciences* 11.1, pp. 1–74.
- Smyth, M. B. (1983). “Powerdomains and predicate transformers: A topological view”. In: *LNCS* 154, pp. 662–675.
- Sosa, E. (1999). “How to Defeat Opposition to Moore”. In: *Philosophical Perspectives* 13, p. 141.
- Speaks, J. (2021). “Theories of Meaning”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2021. Metaphysics Research Lab, Stanford University.
- Srivastava, N., G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov (2014). “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15.56, pp. 1929–1958.
- Srivastava, S. M. (1998). *A Course on Borel Sets*. New York: Springer-Verlag.
- Stalnaker, R. (1968). “A Theory of Conditionals”. In: *Studies in Logical Theory*. Ed. by N. Rescher. Oxford: Blackwell, pp. 98–112.
- (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Stark, E. W. (1989). “Concurrent transition systems”. In: *Theoretical Computer Science* 64.3, pp. 221–269.
- (1990). “Connections between a concrete and an abstract model of concurrent systems”. In: *Mathematical Foundations of Programming Semantics*. Ed. by M. Main, A. Melton, M. Mislove, and D. Schmidt. New York, NY: Springer, pp. 53–79.

- Starr, W. (2019). “Counterfactuals”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. Spring 2019. Metaphysics Research Lab, Stanford University.
- Stoltenberg-Hansen, V., I. Lindström, and E. R. Griffor (1994). *Mathematical Theory of Domains*. Cambridge: Cambridge University Press.
- Szegedy, C., W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus (2013). *Intriguing properties of neural networks*. arXiv: 1312.6199.
- Tabor, W. (2009). “A dynamical systems perspective on the relationship between symbolic and non-symbolic computation”. In: *Cogn Neurodyn* 3, pp. 415–427.
- Tarski, A. (1986). “What are logical notions?” In: *History and Philosophy of Logic* 7.2. Edited by John Corcoran based on a lecture of Tarski from 16 May 1966., pp. 143–154.
- Turing, A. M. (1936-7). “On Computable Numbers, with an Application to the Entscheidungsproblem”. In: *Proceedings of the London Mathematical Society* 42.1, pp. 230–265.
- Van Gelder, T. (1995). “What Might Cognition Be, If Not Computation?” In: *The Journal of Philosophy* 92.7, pp. 345–381.
- Vapnik, V. N. (2000). *The Nature of Statistical Learning Theory*. 2nd ed. New York: Springer.
- Vengertsev, D. and E. Sherman (2020). “Recurrent Neural Network Properties and Their Verification with Monte Carlo Techniques”. In: *CEUR Workshop Proceedings*. Vol. 2560, pp. 178–185.
- Viana, M. and K. Oliveira (2016). *Foundations of Ergodic Theory*. Vol. 151. Cambridge studies in advanced mathematics. Cambridge: Cambridge University Press.
- Vickers, S. (1989). *Topology via Logic*. Cambridge: Cambridge University Press.
- Vigo, R. (2009). “Categorical invariance and structural complexity in human concept learning”. In: *Journal of Mathematical Psychology* 53.4, pp. 203–221.
- V’yugin, V. V. (1998). “Effective convergence in probability and an ergodic theorem for individual random sequences”. In: *Theory of Probability & Its Applications* 42.1, pp. 39–50.

- Walker, R. C. (1974). *The Stone-Čech Compactification*. Vol. 83. Ergebnisse der Mathematik und ihrer Grenzgebiete. Berlin, Heidelberg, New York: Springer-Verlag.
- Walters, P. (1982). *An Introduction to Ergodic Theory*. New York: Springer.
- Warrington, E. K. and A. M. Taylor (1973). “The Contribution of the Right Parietal Lobe to Object Recognition”. In: *Cortex* 9.2, pp. 152–164.
- Weihrauch, K. (2000). *Computable Analysis: An Introduction*. Berlin, Heidelberg: Springer.
- Weiss, B. (1984). “Measurable dynamics”. In: *Conference in Modern Analysis and Probability*. Vol. 26. Contemporary Mathematics, pp. 395–421.
- (1989). “Countable Generators in Dynamics: Universal Minimal Models”. In: *Measure and Measurable Dynamics: Proceedings of a Conference in Honor of Dorothy Maharam Stone held September 17–19, 1987*. Ed. by R. D. Mauldin, R. M. Shortt, and C. E. Silva. Vol. 94. Contemporary Mathematics. Providence, Rhode Island: American Mathematical Society, pp. 321–326.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Winskel, G. (1993). *The Formal Semantics of Programming Languages: An Introduction*. Cambridge, Massachusetts: The MIT Press.
- Winskel, G. and M. Nielsen (1995). “Models for concurrency”. In: *Handbook of Logic in Computer Science*. Ed. by S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum. Vol. 4. Oxford: Clarendon Press.
- Yekutieli, A. (2020). *Derived Categories*. Cambridge: Cambridge University Press.
- Zaffora Blando, F. (2019). “A learning-theoretic characterization of Martin-Löf randomness and Schnorr randomness”. In: *The Review of Symbolic Logic*, pp. 1–19.

- adjunction, 76
 - co-reflective, 76
 - reflective, 76
- adversarial attack, 290
- approximable, 38
- approximation, 38
- asynchronous transition systems, 29
- automaton with concurrency, 30

- basis, 127
 - backward closed, 127
 - countable, 127
 - forward closed, 127
 - measurable, 127
 - separating, 127
 - topological, 127
- behavioral transition system (BTS),
 - 50, 76
 - antisymmetric, 79
 - approximable, 79
 - countable, 50
 - extensional, 79
 - full, 79
 - reflexive, 79
 - unlabeled, 79
- bisimulation, 53
- Bratteli–Vershik diagram, 63
- BTS-morphism, 78

- calculus of fractions, 332
- canonical embedding, 162
- category, 76
 - category of fractions, 332
 - compact element, 19
 - compactification, 202
 - computational, 203
 - functor, 208
 - logical, 203
 - Stone–Čech, 202
 - Wallman, 203
 - computation
 - non-symbolic, 1
 - symbolic, 1
 - computational model, 68
 - concurrency, 28
 - cover, 128
 - finite, 128
 - finite open, 246
 - dcpo, 19
 - algebraic, 19
 - bounded complete, 119
 - max-reflective, 198
 - derivability relation, 298
 - directed, 19
 - domain
 - ω -algebraic, 20
 - bifinite, 144
 - flat, 98
 - initialized, 94
 - marked, 98
 - observation, 162
 - power, 121
 - Scott, 119

- SFP, 144
- domain functor, 216
- domain theory, 3, 119
- domain-entropy, 248
- dynamical dcpo, 141
 - dynamical Scott domain, 141
 - finite, 141
 - max-normalized, 141
 - max-preserving, 141
 - surjective, 141
 - valuation-preserving, 141
- dynamical domain, 160
 - max-reflective, 199
 - standard, 160
- dynamical morphism, 142
- dynamical Scott domain, 141
- dynamical system, 5, 124, 183
 - abstract, 124
 - deterministic, 6
 - ergodic, 274
 - general, 124
 - homomorphism, 182
 - modeled by a dynamical domain, 161
 - non-deterministic, 6
 - standard, 124
 - state-continuous, 6
 - state-discrete, 6
 - time-continuous, 6
 - time-discrete, 6
- embedding, 121
- entropy
 - domain-, 248
 - max-, 258
 - metric, 245
 - topological, 246
- equivalence
 - extensional, 37
 - intensional, 37
 - temporal, 37
- equivariant, 125, 126
- ergodic decomposition, 275
- ergodic hypothesis, 273
- expanding system of dynamical dcpos, 146
 - downward deterministic, 146
 - eventually valuation-preserving, 147
 - upward deterministic, 146
- factor, 182
- falsifiable, 288
- finitary dynamical expanding system, 147
 - standard, 147
- Fitch's paradox, 307
- follows, 130
- full abstraction, 16, 63
- functor, 76
- game semantics, 62
- Giry functor, 172
- Humean thesis, 292
- ideal, 20
 - completion, 20
- index set
 - induced by a basis, 129
- information containment, 50
- initial element, 94
- injective clop 0, 331
- Jewett–Krieger theorem, 115, 335
- Krylov–Bogolioubov theory, 172, 275
- labeled transition system (LTS), 18, 72
 - countable, 18
 - with independence, 30
- language (modal logic), 297
- Lawson topology, 121
- least element, 19
- least upper bound, 19
- Leavitt path algebra, 63
- Lebesgue space, 123
- limit

- restricted, 146
- limit assumption, 296
- linear temporal logic, 63
- localization, 202, 333
- Lockean thesis, 292
- LTS-morphism, 72

- marked domain morphism, 99
- max-entropy, 258
- max-preserving function, 140
- maximal element, 120
 - set of, 140
- Mazurkiewicz trace language, 29
- measure-preserving, 125, 126
- measured topological system, 125, 187
 - based, 196
 - compact, 125
 - general, 125
 - modeled by a dynamical domain, 161
 - morphism, 126
 - standard, 125
 - zero-dimensional, 125
- metric isomorphism, 125

- observation domain, 162
- observation history, 130
- observational equivalence, 132

- partial order, 19
- partition, finite measurable, 245
 - generate, 245
- poset (= partial order), 119
- powerdomain, 121
- pre-behavioral transition system (pre-BTS), 33
 - bisimulative, 51
 - countable, 34
 - extendable, 51
 - extensional, 51
 - full, 51
 - full_c, 51
 - limit-respecting, 40
 - restrictable, 51
- preorder, 19
- principle
 - dual Moore closure, 304
 - falsifiability, 302
 - Moore closure, 304
 - non-triviality, 301
 - standard countermodel, 303
 - standard model, 303
 - verifiability, 302
- probability space, 123
 - Borel, 123
 - completion, 123
 - Lebesgue, 123
- projection, 121

- questions, 301

- randomness, 277
 - algorithmic, 277, 361
- refines, 128
 - accurately, 128
 - open covers, 246
 - partitions, 245
- relevance logic, 100

- safety conditions for knowledge, 289
- Scott information system, 57
- Scott topology, 121
- Scott-continuous function, 20
- semantics, 3
 - denotational, 3
 - Kripke, 299
 - operational, 3
 - topological, 299
- significance, 293
- Smyth powerdomain, 122
- stability theory of belief, 292
- standard Borel space, 123
- Stone duality, 23
- supremum, 19
- system functor, 215

- time function, 252
- topological realization, 335

topology via logic, 288

trajectory, 5, 19

trajectory domain, 56

 functor, 86

upset, 120

valuation, 122

 continuous, 122

 max-normalized, 141

 normalized, 122

verifiable, 288

List of symbols

LTS	Labeled transition system	18
ϵ	Empty trajectory	18
ω	First infinite ordinal	18
last	Last state of a finite nonempty trajectory	19, 72
s	State sequence	19
l	Label sequence	19
$(\overline{P}, \overline{\leq})$	Partial order induced by the preorder (P, \leq)	19
ldl	Ideal completion	20
Spec	Spectrum (set of ultrafilters)	23, 208
pre-BTS	Pre-behavioral transition system	33
\mathbb{T}	Quotient of trajectories	38
\mathbb{T}_{fin}	Quotient of finite trajectories	38
\Vdash_{\forall}	Universal approximation containment	39
\Vdash_{\exists}	Existential approximation containment	39
\sqsubseteq_{\forall}	Universal information containment	39
\sqsubseteq_{\exists}	Existential information containment	39
\sqsubseteq_{dom}	Domination information containment	39
$I_{\triangleleft}([t])$	The \triangleleft -ideal induced by $[t]$	42
$\overline{\text{ldl}}$	Ideal completion with top	42
BTS	Behavioral transition system	50
\mathbb{T}	Trajectory domain functor	56, 86
Con	Consistent sets of information system	57
\vdash	[In information systems] Entailment relation	57
$D_S(I)$	Domain induced by information system I	58
M_I	BTS induced by information system I	59
LTS	Category of LTSs	72
id	Identity function	72, 121
ωALG	Category of ω -algebraic dcpos	75
\mathbb{I}	Inclusion functor	76

BTS	Category of BTSs	79
BTS ^s	Category of BTSs with only synchronous morphisms	79
ω BTS	Category of countable BTSs	79
BTS _a	Category of approximable BTS	80
ω BTS _a	Category of countable approximable BTS	80
ω BTS ^s	Full subcategory of countable objects in BTS ^s	80
ω BTS _a ^s	Full subcategory of approximable objects in ω BTS ^s	80
ω BTS _{fey} ^s	A subcategory of ω BTS ^s	80
ω BTS _{feyur} ^s	A subcategory of ω BTS ^s	80
G	Forgetting behavioral structure	82
F	Freely adding behavioral structure	82
A	Removing non-approximable behavior functor	83
E	Extensionalizing functor	88
U	Unlabeling functor	92
iALG	Category of initialized domains	94
G _i	Forgetting initialized structure	94
T _i	Initialized trajectory domain functor	94
B	Right adjoint to T _i	95
mALG	Category of marked domains	99
$\mathcal{B}(X)$	The Borel σ -algebra of the topological space X	112
P	Smyth powerdomain	122
J	Dynamical system induced by topological system	126, 188
\bar{J}	Completed dynamical system induced by topological system	126, 188
$I(\mathcal{B})$	Index set induced by basis \mathcal{B}	129
H	Observation history	130
max D	Set of maximal element of the dcpo D	140
dDCP	Category of dynamical dcpos	143
dDCP ^P	Category of dynamical dcpos where morphisms also are projections	143
dSCO	Category of dynamical Scott domains where morphisms also are projections	143
dSCO ^P	Category of dynamical Scott domains where morphisms also are projections	143
dSCO _{nf} ^P	Subcategory of dSCO ^P	143
dSCO _{nm}	Subcategory of dSCO	143
dSCO _{nm} ^P	Subcategory of dSCO ^P	143
dSCO _n	Subcategory of dSCO	143
dSCO _n ^P	Subcategory of dSCO ^P	143
dDOM	Category of dynamical domains	160

dDOM_s	Category of standard dynamical domain	160
S	System functor	161, 215
\vdash	[In category theory] Adjunction	174
\sim	Adjoint equivalence	174
Clop	Set of clopens	176
aDS	Category of abstract dynamical systems	183
DS	Category of general dynamical systems	183
sDS	Category of standard dynamical systems	183
$\text{TS}_{(s)}$	Category of (standard) topological systems	187
$\text{TS}_{0(s)}$	Category of zero-dimensional (standard) topological systems	187
$\text{TS}_{0c(s)}$	Category of zero-dimensional compact (standard) topological systems	187
$\text{bTS}_{0(s)}$	Category of based zero-dimensional (standard) topological systems	196
$\bar{\text{bTS}}_{0(s)}$	Category of closed-based zero-dimensional (standard) topological systems	196
dDOM_r	Category of max-reflective dynamical domains	201
dDOM_{rs}	Category of max-reflective standard dynamical domains	201
Loc	Localization functor	201, 331
l_B	The ‘inclusion’ $\text{TS}_{0c} \rightarrow \text{bTS}_0$	203, 204
C	Logical compactification functor	204
\bar{C}	Closing and compactification functor	204
$\bar{\mathcal{B}}$	Closure of the Boolean algebra \mathcal{B}	205
\hat{S}	Extended systems functor	215
D	Domain functor	221
\hat{D}	Restricted domain functor	222
TS_{0cm}	Category of measure-preserving compact zero-dimensional topological systems	244
dDOM_v	Category of dynamical domains with eventually valuation-preserving diagram	244
H	Metric entropy of a partition	245
h	Metric entropy	245
$\sigma(\cdot)$	Smallest σ -algebra containing \cdot	245
H_{top}	Topological entropy of a cover	246
h_{top}	Topological entropy	246
E	Domain-entropy of an index	248
e	Domain-entropy	248
dDOM_{rv}	Full subcategory of dDOM_v consisting of max-reflective dynamical domains	252
M	Max-entropy of an index	258

m	Max-entropy	258
\mathcal{L}	The (bi)modal language	297
\square	Stability operator	297
\boxplus	Global necessity operator	297
\diamond	Dual stability operator	297
\boxminus	Global possibility operator	297
\vdash	[In modal logic] Derivability relation	298
M	Moore operator	304
W	Dual Moore operator	304
Int	Topological interior	313
Cl	Topological closure	313
δ	Topological boundary	313
F	Forgetting the base functor	334

Samenvatting

Dynamische systemen via domeinen: naar een universele grondslag voor symbolische en niet-symbolische informatieverwerking

Niet-symbolische informatieverwerking (zoals voorkomend in b.v. biologische en kunstmatige neurale netwerken) is verbazend krachtig in het leren en verwerken van data waarin ruis aanwezig is. Echter, niet-symbolische informatieverwerking ontbeert een theoretische grondslag zoals die bestaat voor symbolische informatieverwerking (b.v. zoals vastgelegd door programmeertalen). Dit heeft tot gevolg, dat ook indien een systeem voor niet-symbolische informatieverwerking succesvol is, er nog geen garantie is dat we begrijpen waarom en onder welke voorwaarden het succesvol is. Om een dergelijk structureel begrip te bereiken, heeft niet-symbolische informatieverwerking, net als symbolische informatieverwerking, een semantiek—of gedragsbeschrijving—nodig. Domeintheorie levert een algemene semantiek voor symbolische informatieverwerking en dit proefschrift gaat over het uitbreiden daarvan naar niet-symbolische informatieverwerking.

Symbolische en niet-symbolische informatieverwerking kunnen opgevat worden als uitgevoerd door dynamische systemen, waarbij in het eerste geval de verzameling toestanden discreet genomen kan worden, terwijl in het tweede geval ook continue verzamelingen toestanden een rol spelen. Om de gevraagde semantiek te construeren, volstaat het dus een semantiek voor dynamische systemen te geven; wij zullen dat doen middels een constructie die aan een dynamisch systeem een ‘domein’ toekent dat het gedrag van het systeem beschrijft. Een domein is een verzameling van elementen, geordend volgens een abstracte relatie ‘ $_$ bevat minstens zoveel informatie als $_$ ’. In ons geval zijn de elementen observeerbare gedragingen x, y, \dots . We zeggen dat ‘ x bevat minstens zoveel informatie als y ’ indien alles wat over het systeem geleerd kan worden uit x ook geleerd kan worden uit y . Eindig-observeerbare gedragingen zijn dan ‘compacte’ of ‘direct bereikbare’ elementen die de oneindige limiet-gedragingen benaderen.

In deel 1 van dit proefschrift beschrijven we zo een domein-theoretische se-

mantiek voor de ‘symbolische’ dynamische systemen met discrete toestanden (de z.g. gelabelde transitie-systemen). In deel 2 doen we dit voor de ‘niet-symbolische’ dynamische systemen met continue toestanden; deze zijn bekend uit ergodentheorie. Dit is echt een semantiek in de zin dat de gedefinieerde constructies functoren zijn in de zin van categorietheorie en zelfs adjuncties vormen. Sterker nog, we krijgen een categoriale equivalentie in het continue geval: een volledige vertaalbaarheid tussen systemen en domeinen.

In deel 3 verkennen we hoe deze semantiek de twee vormen van informatieverwerking aan elkaar relateert. De semantiek suggereert dat in het algemeen niet-symbolische informatieverwerking de limiet van symbolische informatieverwerking is. Het begrip ‘limiet’ wordt hier gebruikt in ‘pro-eindige’ zin, wat betekent dat een stabiel resultaat pas na een proces van ‘trial and error’ bereikt wordt. In het speciale geval dat het gedrag van het system relatief stabiel is, kan het beschreven worden in termen van berekenbaarheid. Verrassend genoeg zijn het juist de begrippen van ergodiciteit en (algoritmische) willekeurigheid die cruciaal zijn voor het gebruikmaken en bereiken van deze stabiliteit. In het laatste hoofdstuk bestuderen we tenslotte het algemene begrip van stabiliteit: een nieuwe interpretatie van Fitch’s paradox laat zien dat stabiliteit niet tegelijk vier wenselijke eigenschappen kan hebben. Dit heeft gevolgen voor de veiligheid (‘safety’) van kunstmatige intelligentie (KI), het uitsluiten van onbedoelde schadelijke consequenties van in KI gebruikte technologie, zoals neurale netwerken. Immers, wil een neuraal netwerk veilig zijn, dan moeten we aannemen dat het netwerk stabiel is in de zin dat voldoende gelijksoortige invoer leidt tot identieke uitvoer. Het centrale thema in dit hoofdstuk is het verkennen van nieuwe toepassingen van bestaand filosofische gedachtegoed (voornamelijk uit de kentheorie) in de niet-symbolische informatieverwerking binnen kunstmatige intelligentie.

Summary

Dynamical Systems via Domains: Toward a Unified Foundation of Symbolic and Non-symbolic Computation

Non-symbolic computation (as, e.g., in biological and artificial neural networks) is astonishingly good at learning and processing noisy real-world data. However, it lacks the kind of understanding we have of symbolic computation (as, e.g., specified by programming languages). Just like symbolic computation, also non-symbolic computation needs a semantics—or behavior description—to achieve structural understanding. Domain theory has provided this for symbolic computation, and this thesis is about extending it to non-symbolic computation.

Symbolic and non-symbolic computation can be described in a unified framework as state-discrete and state-continuous dynamical systems, respectively. So we need a semantics for dynamical systems: assigning to a dynamical system a ‘domain’ which describes the system’s behavior. A domain is a set of elements ordered by information containment. In our case, the elements are observable behaviors of the system, and one behavior x is informationally contained in another y if what can be learned about the system from x can also be learned from y . Finitely observable behaviors then are ‘compact’ or ‘directly accessible’ elements that can approximate the infinite limit behaviors of the system.

In part 1 of the thesis, we provide this domain-theoretic semantics for the ‘symbolic’ state-discrete systems (i.e., labeled transition systems). And in part 2, we do this for the ‘non-symbolic’ state-continuous systems (known from ergodic theory). This is a proper semantics in that the constructions form functors (in the sense of category theory) and, once appropriately formulated, even adjunctions. Stronger yet, we obtain a categorical equivalence in the continuous case: a complete intertranslatability between systems and domains.

In part 3, we explore how this semantics relates the two types of computation. It suggests that non-symbolic computation is the limit of symbolic computation (in the ‘profinite’ sense). Conversely, if the system’s behavior is fairly stable, it may

be described as realizing symbolic computation. The concepts of ergodicity and (algorithmic) randomness help to use and achieve this stability. In the last chapter, we then study the general concept of stability: A novel interpretation of Fitch's paradox reveals that stability cannot jointly have four desirable properties. This has implications for AI-safety: After all, for a neural network to be safe, we expect it to be stable in the sense of computing the same output on sufficiently similar inputs. The theme here is to explore new applications of established philosophical tools (mostly from epistemology) in the non-symbolic computation of modern AI.

Titles in the ILLC Dissertation Series:

- ILLC DS-2016-01: **Ivano A. Ciardelli**
Questions in Logic
- ILLC DS-2016-02: **Zoé Christoff**
Dynamic Logics of Networks: Information Flow and the Spread of Opinion
- ILLC DS-2016-03: **Fleur Leonie Bouwer**
What do we need to hear a beat? The influence of attention, musical abilities, and accents on the perception of metrical rhythm
- ILLC DS-2016-04: **Johannes Marti**
Interpreting Linguistic Behavior with Possible World Models
- ILLC DS-2016-05: **Phong Lê**
Learning Vector Representations for Sentences - The Recursive Deep Learning Approach
- ILLC DS-2016-06: **Gideon Maillette de Buy Wenniger**
Aligning the Foundations of Hierarchical Statistical Machine Translation
- ILLC DS-2016-07: **Andreas van Cranenburgh**
Rich Statistical Parsing and Literary Language
- ILLC DS-2016-08: **Florian Speelman**
Position-based Quantum Cryptography and Catalytic Computation
- ILLC DS-2016-09: **Teresa Piovesan**
Quantum entanglement: insights via graph parameters and conic optimization
- ILLC DS-2016-10: **Paula Henk**
Nonstandard Provability for Peano Arithmetic. A Modal Perspective
- ILLC DS-2017-01: **Paolo Galeazzi**
Play Without Regret
- ILLC DS-2017-02: **Riccardo Pinosio**
The Logic of Kant's Temporal Continuum
- ILLC DS-2017-03: **Matthijs Westera**
Exhaustivity and intonation: a unified theory
- ILLC DS-2017-04: **Giovanni Cinà**
Categories for the working modal logician
- ILLC DS-2017-05: **Shane Noah Steinert-Threlkeld**
Communication and Computation: New Questions About Compositionality

- ILLC DS-2017-06: **Peter Hawke**
The Problem of Epistemic Relevance
- ILLC DS-2017-07: **Aybüke Özgün**
Evidence in Epistemic Logic: A Topological Perspective
- ILLC DS-2017-08: **Raquel Garrido Alhama**
Computational Modelling of Artificial Language Learning: Retention, Recognition & Recurrence
- ILLC DS-2017-09: **Miloš Stanojević**
Permutation Forests for Modeling Word Order in Machine Translation
- ILLC DS-2018-01: **Berit Janssen**
Retained or Lost in Transmission? Analyzing and Predicting Stability in Dutch Folk Songs
- ILLC DS-2018-02: **Hugo Huurdeman**
Supporting the Complex Dynamics of the Information Seeking Process
- ILLC DS-2018-03: **Corina Koolen**
Reading beyond the female: The relationship between perception of author gender and literary quality
- ILLC DS-2018-04: **Jelle Bruineberg**
Anticipating Affordances: Intentionality in self-organizing brain-body-environment systems
- ILLC DS-2018-05: **Joachim Daiber**
Typologically Robust Statistical Machine Translation: Understanding and Exploiting Differences and Similarities Between Languages in Machine Translation
- ILLC DS-2018-06: **Thomas Brochhagen**
Signaling under Uncertainty
- ILLC DS-2018-07: **Julian Schlöder**
Assertion and Rejection
- ILLC DS-2018-08: **Srinivasan Arunachalam**
Quantum Algorithms and Learning Theory
- ILLC DS-2018-09: **Hugo de Holanda Cunha Nobrega**
Games for functions: Baire classes, Weihrauch degrees, transfinite computations, and ranks

- ILLC DS-2018-10: **Chenwei Shi**
Reason to Believe
- ILLC DS-2018-11: **Malvin Gattinger**
New Directions in Model Checking Dynamic Epistemic Logic
- ILLC DS-2018-12: **Julia Ilin**
Filtration Revisited: Lattices of Stable Non-Classical Logics
- ILLC DS-2018-13: **Jeroen Zuiddam**
Algebraic complexity, asymptotic spectra and entanglement polytopes
- ILLC DS-2019-01: **Carlos Vaquero**
What Makes A Performer Unique? Idiosyncrasies and commonalities in expressive music performance
- ILLC DS-2019-02: **Jort Bergfeld**
Quantum logics for expressing and proving the correctness of quantum programs
- ILLC DS-2019-03: **Andras Gilyen**
Quantum Singular Value Transformation & Its Algorithmic Applications
- ILLC DS-2019-04: **Lorenzo Galeotti**
The theory of the generalised real numbers and other topics in logic
- ILLC DS-2019-05: **Nadine Theiler**
Taking a unified perspective: Resolutions and highlighting in the semantics of attitudes and particles
- ILLC DS-2019-06: **Peter T.S. van der Gulik**
Considerations in Evolutionary Biochemistry
- ILLC DS-2019-07: **Frederik Mollerstrom Lauridsen**
Cuts and Completions: Algebraic aspects of structural proof theory
- ILLC DS-2020-01: **Mostafa Dehghani**
Learning with Imperfect Supervision for Language Understanding
- ILLC DS-2020-02: **Koen Groenland**
Quantum protocols for few-qubit devices
- ILLC DS-2020-03: **Jouke Witteveen**
Parameterized Analysis of Complexity
- ILLC DS-2020-04: **Joran van Apeldoorn**
A Quantum View on Convex Optimization

- ILLC DS-2020-05: **Tom Bannink**
Quantum and stochastic processes
- ILLC DS-2020-06: **Dieuwke Hupkes**
Hierarchy and interpretability in neural models of language processing
- ILLC DS-2020-07: **Ana Lucia Vargas Sandoval**
On the Path to the Truth: Logical & Computational Aspects of Learning
- ILLC DS-2020-08: **Philip Schulz**
Latent Variable Models for Machine Translation and How to Learn Them
- ILLC DS-2020-09: **Jasmijn Bastings**
A Tale of Two Sequences: Interpretable and Linguistically-Informed Deep Learning for Natural Language Processing
- ILLC DS-2020-10: **Arnold Kochari**
Perceiving and communicating magnitudes: Behavioral and electrophysiological studies
- ILLC DS-2020-11: **Marco Del Tredici**
Linguistic Variation in Online Communities: A Computational Perspective
- ILLC DS-2020-12: **Bastiaan van der Weij**
Experienced listeners: Modeling the influence of long-term musical exposure on rhythm perception
- ILLC DS-2020-13: **Thom van Gessel**
Questions in Context
- ILLC DS-2020-14: **Gianluca Grilletti**
Questions & Quantification: A study of first order inquisitive logic
- ILLC DS-2020-15: **Tom Schoonen**
Tales of Similarity and Imagination. A modest epistemology of possibility
- ILLC DS-2020-16: **Ilaria Canavotto**
Where Responsibility Takes You: Logics of Agency, Counterfactuals and Norms