Taming the Infinity Quantifier: On Well-Behaved Fragments of First-Order Logic with the Quantifier 'There are Infinitely Many'

MSc Thesis *(Afstudeerscriptie)*

written by

Thibault Rushbrooke
(born 16th August, 1998 in London)

under the supervision of **Prof. Dr. Yde Venema**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	**Members of the Thesis Committee:**
28th October, 2022	Dr. Benno van den Berg (Chair)
	Prof. Dr. Yde Venema (Supervisor)
	Dr. Alexandru Baltag
	Dr. Balder ten Cate

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

This thesis investigates a range of fragments of FOL^{∞}, that is, first-order logic extended with a quantifier expressing 'there are infinitely many...' We extend several results of Bellas Acosta [1], concerning the logic ML^{∞}. The logic ML^{∞} extends the basic modal language with an infinite modality, asserting that there are infinitely many successors satisfying a certain property. We introduce logics which extend ML^{∞} with a backwards modality and a backwards infinite modality, with a global modality, and with both backwards modalities and the global modality. We show that each of these logics may be characterised, as fragments of FOL^{∞}, by invariance under a suitable notion of bisimulation. On the way to obtaining these results, we prove an FOL^{∞} analogue of Gaifman's theorem for first-order logic [9]. In a second line of research, we establish that the set of validities of the logic with forwards and backwards infinite modalities is decidable, and admits a finite axiomatisation. We also introduce a guarded fragment of FOL^{∞}, and conjecture that, restricted to a certain vocabulary, this logic is also decidable.

Acknowledgements

The most important person I would like to thank is my supervisor, Yde, for all your help and support in writing this thesis. You were unwaveringly generous with your time, and always returned my draft writings with perceptive and helpful comments. Whenever I thought I'd proved something, you played the important role of making me write the argument out rigorously (more often than not, I found I'd overlooked something); and whenever I was stuck, you played the even more important role of giving me new ideas and suggesting new directions. I came out of every meeting with you feeling that I'd learnt something about logic. For all this, and more, I'm immensely grateful.

I would also like to thank the members of my Defense committee, Benno, Alexandru, and Balder, for taking the time to read my thesis in full, and for your insightful questions and comments. I hope I can pursue some of your suggestions in future work.

Next, I want to thank all the friends I've made in Amsterdam, who've made my time in the Master of Logic so enjoyable through stimulating conversations, bad logic puns and warm affection. The cardinality of this set is too great for me to thank you all individually, much though I'd like to. However, I would like to single out Rodrigo, for sportingly agreeing to be my adversary in a demonstration of the Infinite Bisimilarity game in my Thesis presentation, preparing for which made me realise a key detail; Vasily, for your steadfast companionship in the computer area of Science Park, where most of this thesis was written; and Nathan, for ensuring I ate plenty of herring (it's good for the brain).

Finally, thank you to my family for all your love and support, and for patiently listening to me ramble about infinitely many dogs and cats.

Contents

Notational Conventions

The aim of this section is not to explain all the pieces of notation which are specific to this thesis; these will be explicitly defined whenever they are introduced. They may also be found in the 'Index of Notation', located at the end of the thesis. Rather, this section explains some of the uses of notation which are fairly conventional, but may nevertheless differ across different authors or be to some degree unclear.

I have tried to consistently use calligraphic font for Kripke models, e.g. $\mathcal{M}, \mathcal{A}$, and fraktur font for structures more generally, which may or may not be Kripke models, e.g. $\mathfrak{M}, \mathfrak{A}$. This will not prevent us from interpreting modal formulas on first-order structures, or formulas of FOL^{∞} on Kripke structures, making free use of the standard translation. The point of the different notation is that where a calligraphic letter such as $\mathcal{M}$ is used, it may explicitly be assumed that the model in question has a basic modal signature- it has no relations of arity greater than 2, and exactly one relation of arity 2. So definitions of, say, bisimilarity between structures $\mathcal{A}, \mathcal{B}$ are guaranteed to make sense. (This rule will be broken in Chapter 6, where I use calligraphic letters for polymodal Kripke structures with multiple binary relations). Where fraktur letters are used, the models may very well have relations of arity greater than 2.

When I use a calligraphic or fraktur letter for a structure, I will generally use the corresponding, regular font capital letter for the universe of that structure, e.g. using A to denote the universe of $\mathfrak{A}$. If $\mathcal{M}$ is a Kripke model, I may also use S (the set of 'states') to denote the universe of $\mathcal{M}$.

As an abbreviation, I use overline notation for finite tuples of elements within a model, $\overline{a}, \overline{b}$ etc., but also for tuples of variables, $\overline{x}, \overline{y}$ etc. These tuples may generally be assumed to be ordered. However, I will also use notation for ordered tuples that strictly speaking only makes sense for sets. For example, where $\overline{a} = (a_1, \ldots, a_n)$, I may write $\overline{a} \subseteq A$ to mean that for every i with $1 \leq i \leq n$, it holds that $a_i \in A$.

I generally use round brackets for ordered tuples. However, in the case of 'paths' through Kripke structures (see Definition 4.3), I try to use angled brackets: $\pi = \langle a_1, \ldots, a_n \rangle$, even though these are technically just ordered tuples as well. I generally use lower-case Greek letters to refer to paths. To denote the concatenation of a path or ordered tuple π with a path or ordered tuple ρ, I write: $\pi^\frown\rho$.

Chapter 1

Introduction

1.1 Background and Motivation

Consider the following inference **I**.

- **P1**: There are infinitely many dogs, each of which sees a cat.
- **P2**: There is no cat which is seen by infinitely many dogs.
- **C**: There are infinitely many cats.

This inference is not trivial, but one may see that it is logically valid using the fact that a finite union of finite sets is finite. As a logically valid inference, **I** ought to be representable within a formal logical system. However, it is very difficult to see how **I** could be formalised in classical first-order logic. The most natural way to capture it is to enrich the first-order language with an *infinity quantifier*, a means of expressing that there are infinitely many entities satisfying a certain property.

The infinity quantifier falls under the broader category of *generalised quantifiers*. The study of generalised quantifiers was introduced by Mostowski [14]. His idea was to think of a unary quantifier as a function which associates, with any given set A, a set $Q_A \subseteq \wp(A)$ of subsets of A. The formula $Qx\phi(x)$ is then defined to be true in a structure $\mathfrak{A}$ iff $\{a \in A : \phi^{\mathfrak{A}}(a)\} \in Q_A$. For example, one may think of the classical existential quantifier $\exists$ as mapping any set A to the set of non-empty subsets of A. Similarly, one may define the infinity quantifier as the function which maps any set A to the set

$$\{X \subseteq A : |X| \geq \aleph_0\}$$

and indeed, for any ordinal α, one may define the quantifier

$$Q_\alpha : A \mapsto \{X \subseteq A : |X| \geq \aleph_\alpha\}$$

Thus, the infinity quantifier may be viewed as one of a whole family of cardinality quantifiers Q_α. (Under the above notation, the infinity quantifier would be denoted by Q_0, though we choose to denote it by $\exists^\infty$ for the majority of this thesis.) Further, for each Q_α, one may obtain a new logic, $\mathcal{L}(Q_\alpha)$, by extending classical first-order logic with the quantifier Q_α.[1]

Now, it was soon observed that $\mathcal{L}(Q_0)$ has various model-theoretic properties which were considered "less than satisfactory" ([4], p.9). In particular, as Fuhrken [8] observes, $\mathcal{L}(Q_0)$ is able to completely characterise the order type $(\omega, <)$ of the natural numbers. This is achieved by the $\mathcal{L}(Q_0)$ sentence

$$(\forall x \neg \exists^\infty y(y < x)) \wedge (\text{`} < \text{ is a strict linear order without endpoint'})$$

From this fact, it quickly follows that $\mathcal{L}(Q_0)$ is not compact, and the set of its valid sentences (in any vocabulary of reasonable size) is not recursively enumerable. (A further negative result, due to Mostowski [13], was that $\mathcal{L}(Q_0)$ lacks both the Craig interpolation property and the Beth definability property.) In striking contrast, Vaught [18] showed that the set of validities of $\mathcal{L}(Q_1)$ *is* recursively enumerable, while Fuhrken [8] showed that $\mathcal{L}(Q_1)$ is countably compact, meaning that every countable, finitely satisfiable set of sentences is satisfiable. An explicit axiomatisation for the validities of $\mathcal{L}(Q_1)$ was later given by Keisler [11]. The moral seems to have been drawn that few interesting model-theoretic results concerning $\mathcal{L}(Q_0)$ were to be had, and consequently, research attention shifted away from $\mathcal{L}(Q_0)$, towards other abstract logics such as $\mathcal{L}(Q_1)$ and its extensions.

However, a more recent line of research is motivated by the observation that negative model-theoretic properties of a logic, such as $\mathcal{L}(Q_0)$, are not necessarily inherited by its *fragments*. For example, it was already shown by Mostowski [14] that the monadic fragment of $\mathcal{L}(Q_0)$ is decidable. Carreiro et al. [6] obtain further nice results about the model theory of this monadic fragment, providing syntactic characterisations of various semantically given fragments of this logic. Another fragment of $\mathcal{L}(Q_0)$ shown to enjoy desirable model-theoretic properties is the logic $ML^\bullet$, introduced by ten Cate, van Benthem and Väänänen [7]. These authors show that the logic $ML^\bullet$ is compact, and has the Löwenheim-Skolem property.

In a similar spirit, Bellas Acosta [1] studies a modal fragment, ML^∞, of $\mathcal{L}(Q_0)$, which extends the basic modal language with an 'infinite modality' asserting the existence of infinitely many successors with a certain property. This logic is shown to be semantically characterisable as the bisimulation-invariant fragment of $\mathcal{L}(Q_0)$ (for a suitable notion of bisimulation), and is further shown to admit a complete,

[1] More precisely, this logic would normally be denoted $\mathcal{L}_{\omega\omega}(Q_\alpha)$, to distinguish it from abstract logics which allow formulas of infinite length. However, this thesis will not be at all concerned with extensions of *FOL* with infinitely long formulas, and therefore we adopt the lighter notation since there is no danger of confusion.

finite axiomatisation. Bellas Acosta and Venema [2] also study a 'graded' extension, GML^{∞}, of ML^{∞} which includes counting modalities $\Diamond^{\geq n}$ for every natural number n. This logic is also characterised by invariance under a suitable notion of bisimulation, and is shown to be decidable (implying decidability for the strictly weaker logic ML^{∞} as well).

The present work takes its inspiration from the two papers [1] and [2] just cited. We introduce an assortment of fragments of $\mathcal{L}(Q_0)$ of a modal character, each properly extending ML^{∞}, and we try to obtain analogous results to those of [1] and [2] for each of these fragments. We now present a chapter-by-chapter overview of the contents of this thesis.

1.2 Structure of the Thesis

Chapter 2 consists of basic definitions, and a few straightforward results which are needed at a later stage of the thesis. The syntax and semantics of several logics to be studied, namely T^{∞}, G^{∞} and GT^{∞}, are defined, and we introduce a notion of bisimilarity to correspond to the expressive power of each of these logics. We also recapitulate the syntax and semantics of $\mathcal{L}(Q_0)$, which we choose to refer to in the remainder of this thesis as 'FOL^{∞}'.

In **Chapter 3** we state, and prove, an FOL^{∞} analogue of Gaifman's theorem for first-order logic, demonstrating that, in some sense, FOL^{∞} is only able to express 'local' notions. The theorem is needed to obtain semantic characterisations of T^{∞}, G^{∞}, and GT^{∞}, but it also seems to be of independent interest.

Chapter 4 completes the proof that T^{∞}, G^{∞} and GT^{∞} may all be characterised, as fragments of FOL^{∞}, by invariance under a suitable notion of bisimulation. The proof makes use of an 'upgrading' technique, developed in the setting of finite model theory to cope with the absence of the compactness theorem.

Chapter 5 introduces the logic GF^{∞}, which stands to FOL^{∞} roughly as the Guarded Fragment introduced in [3] stands to *FOL*. An attractive feature of this logic is that it is strong enough to capture the inference **I** from the beginning of this introduction. We define the syntax of GF^{∞}, introduce a bisimilarity relation $\simeq_{GF}$ corresponding to its expressivity features, and show that formulas of GF^{∞} are invariant under $\simeq_{GF}$. Finally, we explain the obstacles that make it difficult to show that GF^{∞} is characterised by invariance under $\simeq_{GF}$.

In **Chapter 6**, we turn away from bisimulation characterisation, and investigate the complexity of some of the logics introduced. In the first section of the chapter, we prove that T^∞ is decidable, and provide a complete finite axiomatisation for it. In the second half, we conjecture that, when restricted to a certain vocabulary, GF^∞ is also decidable, and we indicate a possible proof strategy.

Chapter 2

Preliminaries

In this chapter, we will introduce a series of modal languages. These languages all extend previously studied modal languages by the inclusion of an infinite modality, that is, a modality allowing the assertion that there are infinitely many accessible states satisfying a certain property. We also recapitulate the logic FOL^∞, and show how to view each of the modal languages under consideration as fragments of this logic. Finally, we introduce suitable notions of bisimulation for each modal language, and show that formulas of each language are invariant under the corresponding bisimilarity relation.

For the remainder of this section we fix an arbitrary, countable set of proposition letters, P, as well as an enumeration $\{p_0, p_1, \ldots\}$ of this set.

2.1 ML^∞, T^∞, G^∞, GT^∞

We begin with a recapitulation of the logic ML^∞, which is defined and studied in [1].

Definition 2.1 (Syntax of ML^∞)**.** ML^∞ is the least set satisfying the following conditions:

- $\mathsf{P} \subseteq ML^\infty$
- $\bot \in ML^\infty$
- If $\phi, \psi \in ML^\infty$ then $\phi \wedge \psi \in ML^\infty$ and $\neg\phi \in ML^\infty$
- If $\phi \in ML^\infty$ then $\Diamond\phi \in ML^\infty$
- If $\phi \in ML^\infty$ then $\Diamond^\infty\phi \in ML^\infty$

If $\phi \in ML^\infty$, we say that ϕ is a *well-formed formula* of ML^∞. We make use of the standard abbreviations $\phi \vee \psi := \neg(\neg\phi \wedge \neg\psi)$, $\Box\phi := \neg\Diamond\neg\phi$. We additionally define $\Box^\infty\phi := \neg\Diamond^\infty\neg\phi$.

Definition 2.2 (Semantics of ML^∞)**.** The satisfaction relation for ML^∞ is defined inductively. Let $\mathcal{M} = (S, R, V)$ be a Kripke model.

- $\mathcal{M}, s \Vdash p$ iff $s \in V(p)$
- $\mathcal{M}, s \Vdash \bot$ never
- $\mathcal{M}, s \Vdash \phi \wedge \psi$ iff $\mathcal{M}, s \Vdash \phi$ and $\mathcal{M}, s \Vdash \psi$
- $\mathcal{M}, s \Vdash \neg\phi$ iff it is not the case that $\mathcal{M}, s \Vdash \phi$
- $\mathcal{M}, s \Vdash \Diamond\phi$ iff there is a $t \in S$ such that sRt and $\mathcal{M}, t \Vdash \phi$
- $\mathcal{M}, s \Vdash \Diamond^\infty\phi$ iff there are infinitely many $t \in S$ such that sRt and $\mathcal{M}, t \Vdash \phi$

Taking inspiration from the two-way extension of basic modal logic known as temporal logic (see e.g. [5]), we now define a two-way analogue of ML^∞, which we call T^∞.

Definition 2.3 (Syntax of T^∞)**.** T^∞ is the least set such that $ML^\infty \subseteq T^\infty$, and T^∞ is additionally closed under the following:

- If $\phi \in T^\infty$ then $\Diamond_\leftarrow\phi \in T^\infty$
- If $\phi \in T^\infty$ then $\Diamond^\infty_\leftarrow\phi \in T^\infty$

If $\phi \in T^\infty$, we say that ϕ is a *well-formed formula* of T^∞. We use the same abbreviations as for ML^∞, as well as the obvious analogues $\Box_\leftarrow\phi := \neg\Diamond_\leftarrow\neg\phi$, $\Box^\infty_\leftarrow\phi := \neg\Diamond^\infty_\leftarrow\neg\phi$.

Definition 2.4 (Semantics of T^∞)**.** The satisfaction relation for T^∞, for Kripke models, is also defined inductively. The clauses are identical to those for ML^∞, but with the following two additions:

- $\mathcal{M}, s \Vdash \Diamond_\leftarrow\phi$ iff there is some $t \in S$ such that tRs and $\mathcal{M}, t \Vdash \phi$
- $\mathcal{M}, s \Vdash \Diamond^\infty_\leftarrow\phi$ iff there are infinitely many $t \in S$ such that tRs and $\mathcal{M}, t \Vdash \phi$

Another extension to the basic modal language which has received considerable attention is the so-called *global modality* E (again, see [5]). In the same spirit, we may add this modality to ML^∞. We call the resulting logic G^∞.

Definition 2.5 (Syntax of G^∞)**.** G^∞ is the least set such that $ML^\infty \subseteq G^\infty$, and G^∞ is additionally closed under the following:

- If $\phi \in G^\infty$ then $E\phi \in G^\infty$

If $\phi \in G^\infty$, we say that ϕ is a *well-formed formula* of G^∞.

Definition 2.6 (Semantics of G^∞)**.** The definition is the same as for ML^∞, but with the following inductive clause added:

- $\mathcal{M} \Vdash E\phi$ iff there is a $t \in S$ such that $\mathcal{M}, t \Vdash \phi$

Notice also that we do *not* let G^∞ include an 'infinite global modality', i.e. a modality which asserts that a property holds at infinitely many states in the model. The reason for this will be explained at a later stage (see Chapter 5).

Finally, we may combine the features enjoyed by T^∞ and G^∞, extending ML^∞ with both converse modalities and a global modality. The resulting logic, which we call GT^∞, properly extends both G^∞ and T^∞ (although we omit a proof of this fact).

Definition 2.7 (Syntax of GT^∞)**.** GT^∞ is the least set such that $T^\infty \subseteq GT^\infty$, and GT^∞ is additionally closed under the following:

- If $\phi \in GT^\infty$ then $E\phi \in GT^\infty$

If $\phi \in GT^\infty$, we say that ϕ is a *well-formed formula* of GT^∞.

Definition 2.8 (Semantics of GT^∞)**.** The definition of the satisfaction relation for GT^∞ is given simply by combining the clauses for T^∞ and G^∞.

2.2 FOL^∞

The logic FOL^∞ is obtained from classical first-order logic, *FOL*, by adding to the language a new quantifier which expresses that there are infinitely many objects having a certain property.

Definition 2.9 (Syntax of FOL^∞)**.** Let τ be a signature which may include constants, predicates of any arity, and function symbols. We define $\mathfrak{Form}^\infty(\tau)$ to be the least set satisfying the following conditions:

- If P is an n-ary predicate symbol in τ, or is the identity symbol, $=$, and $t_1, \ldots, t_n$ are terms of τ, then $Pt_1 \ldots t_n$ is a member of $\mathfrak{Form}^\infty(\tau)$ and is an *atomic formula*

- $\perp$ is a member of $\mathfrak{Form}^\infty(\tau)$ and is an atomic formula

- If $\phi, \psi \in \mathfrak{Form}^\infty(\tau)$ then $\phi \wedge \psi \in \mathfrak{Form}^\infty(\tau)$ and $\neg\phi \in \mathfrak{Form}^\infty(\tau)$
- If $\phi \in \mathfrak{Form}^\infty(\tau)$ and v is a variable then $\exists v\phi \in \mathfrak{Form}^\infty(\tau)$ and $\exists^\infty v\phi \in \mathfrak{Form}^\infty(\tau)$

If there is some τ for which $\phi \in \mathfrak{Form}^\infty(\tau)$, we say that ϕ is a *well-formed formula* of FOL^∞, and write $\phi \in FOL^\infty$. If in addition ϕ has no free variables, we say that ϕ is a *sentence* of FOL^∞. We define $\vee$ in terms of the other Booleans, as in the modal case. We also abbreviate $\forall v\phi := \neg\exists v\neg\phi$, and $\forall^\infty v\phi := \neg\exists^\infty v\neg\phi$. The *quantifier depth* $qd(\phi)$ of a formula ϕ is defined inductively in the usual way.

Definition 2.10 (Semantics of FOL^∞)**.** Let τ be any signature, $\phi \in \mathfrak{Form}^\infty(\tau)$, $\mathfrak{M} = \langle D, I\rangle$ a τ-structure, and σ a variable assignment. For any $d \in D$, we let $\sigma[d/v]$ denote the assignment which is identical to σ, except possibly for mapping the variable v to the object d. The satisfaction relation is inductively defined, with the same clauses as for classical FOL, along with the following additional clause:

- $\mathfrak{M}, \sigma \models \exists^\infty v\phi$ iff there are infinitely many $d \in D$ such that $\mathfrak{M}, \sigma[d/v] \models \phi$

Where ψ is a formula containing at most the variables $v_1, \ldots, v_n$ free, we will sometimes use $\mathfrak{M}, a_1, \ldots, a_n \models \psi(x_1, \ldots, x_n)$ as an abbreviation for:
$\mathfrak{M}, \sigma[a_1/v_1, \ldots, a_n/v_n] \models \psi$.

Definition 2.11 (Equivalence of structures)**.** Let $\mathfrak{M}_0, \mathfrak{M}_1$ be structures, and let $\overline{a}_0, \overline{a}_1$ be ordered n-tuples with $\overline{a}_0 \subseteq M_0, \overline{a}_1 \subseteq M_1$. We write $\mathfrak{M}_0, \overline{a}_0 \equiv \mathfrak{M}_1, a_1$ iff for all $\phi \in FOL^\infty$, $\mathfrak{M}_0, \overline{a}_0 \models \phi \iff \mathfrak{M}_1, \overline{a}_1 \models \phi$. We write $\mathfrak{M}_0, \overline{a}_0 \equiv^k \mathfrak{M}_1, \overline{a}_1$ iff for all $\phi \in FOL^\infty$ such that $qd(\phi) \leq k$, $\mathfrak{M}_0, \overline{a}_0 \models \phi$ iff $\mathfrak{M}_1, \overline{a}_1 \models \phi$

Remark. The relations $\equiv$ and $\equiv^k$ are not to be confused with first-order equivalence between structures. It would have been natural to include a subscript or superscript indicating that the notation is for equivalence in *FOL infinity*, which is strictly stronger. However, we have not done so, because this thesis is almost exclusively concerned with FOL^∞, so there is little danger of confusion; and omitting the infinity symbol keeps notation lighter.

Definition 2.12 (Entailment)**.** Let Γ be a set of FOL^∞-formulas and ϕ an FOL^∞-formula. We say that Γ *entails* ϕ, notation: $\Gamma \models \phi$, iff every structure-assignment pair $\mathfrak{M}, \sigma$ which satisfies every member of Γ also satisfies ϕ. (Note that this definition also covers the case of formulas which may have free variables.)

We now state an observation which provides a very handy way of viewing the semantics of FOL^∞. We need a few preliminaries first.

Definition 2.13. Let $(q_0, \ldots, q_n) \subseteq \{1, \omega\}$. We let $\mathbb{F}(q_0, \ldots, q_n)$ denote the forest structure which has q_0-many roots, and has branching degree q_1 at depth 1, branching degree q_2 at depth 2, and so on, up to the nodes of depth $n+1$ which are leaves. For instance, $\mathbb{F}(\omega, 1, \omega)$ is isomorphic to $\{(p) : p \in \mathbb{N}^+\} \cup \{(p, 0) : p \in \mathbb{N}^+\} \cup \{(p, 0, q) : p, q \in \mathbb{N}^+\}$, under the *proper initial segment* relation.

Given any forest structure $\mathbb{F} = (F, R)$ and any $t \in \mathbb{F}$, we let $\uparrow_{\mathbb{F}}(t)$ denote the set of R-descendants of t in $\mathbb{F}$. We usually omit the subscript $\mathbb{F}$ if it is obvious from context. We further define an equivalence relation $\sim_{\mathbb{F}}$ by: $t \sim_{\mathbb{F}} t'$ iff either t has the same immediate predecessor as t', or both t, t' have no predecessor (they are both 'roots').

Let $Q_0, \ldots, Q_n$ be a string of quantifiers, with each $Q_i \in \{\exists, \exists^\infty\}$. The *index sequence* for $Q_0, \ldots, Q_n$ is the sequence $q_0, \ldots, q_n$ such that, for each $0 \leq i \leq n$, $q_i = 1$ iff $Q_i = \exists$, and $q_i = \omega$ iff $Q_i = \exists^\infty$.

We let $(\omega)^n$ be shorthand for the ordered n-tuple all of whose elements are ω. So $\mathbb{F}((\omega)^n)$ denotes the forest structure of depth n with ω many roots, and with every node having branching degree ω.

Observation 2.14. Let ϕ be a formula of FOL^∞, $\mathfrak{M} = \langle D, I \rangle$ a structure, σ an assignment, $n \in \mathbb{N}$. Let $Q_0, \ldots, Q_n \in \{\exists, \exists^\infty\}$, and let $(q_0, \ldots, q_n)$ be the index sequence for $Q_0, \ldots, Q_n$. Then $\mathfrak{M}, \sigma \models Q_0 x_0 \ldots Q_n x_n \phi$ iff there exists a function $f : \mathbb{F}(q_0, \ldots, q_n) \to D$ such that:

1. For any $\sim_{\mathbb{F}}$-equivalence class $[t] \subseteq \mathbb{F}(q_0, \ldots, q_n)$, $f \restriction [t]$ is injective.
2. For any path $\pi = (t_0 R^{\mathbb{F}} \ldots R^{\mathbb{F}} t_n)$ through $\mathbb{F}(q_0, \ldots, q_n)$, we have $\mathfrak{M}, \sigma[f(t_0)/x_0, \ldots, f(t_n)/x_n] \models \phi$.

Proof. We argue by induction on n.

Base case: $n = 0$. There are two subcases to consider.

Case 1: $Q_0 = \exists$. Suppose $\mathfrak{M}, \sigma \models \exists x_0 \phi$. Then there is some $a \in M$ such that $\mathfrak{M}, \sigma[a/x_0] \models \phi$. Now, $\mathbb{F}(1)$ is just the 'forest' with a single element, t. Define the map f by: $f(t) = a$. Then $f : \mathbb{F}(q_0) \to M$ is a map satisfying conditions 1. and 2.

Conversely, suppose there is a map $f : \mathbb{F}(1) \to M$ satisfying the given conditions. Let t be the sole element of $\mathbb{F}(1)$, and let $f(t) = a$. Then $\mathfrak{M}, \sigma[a/x_0] \models \phi$, hence $\mathfrak{M}, \sigma \models \exists x \phi$.

Case 2: $Q_0 = \exists^\infty$. Suppose $\mathfrak{M}, \sigma \models \exists^\infty x_0 \phi$. Then there is an infinite set $A \subseteq M$ such that for each $a \in A, \mathfrak{M}, \sigma[a/x_0] \models \phi$. $\mathbb{F}(\omega)$ is the forest which simply consists of ω-many disconnected elements, so let f be any injection from $\mathbb{F}(\omega)$ to A. Then f satisfies conditions 1. and 2.

Conversely, let $f : \mathbb{F}(\omega) \to M$ be a map satisfying conditions 1. and 2. By 1., $A := ran(f)$ is infinite. By 2., $\mathfrak{M}, \sigma[a/x_0] \models \phi$ for each $a \in A$. Hence $\mathfrak{M}, \sigma \models \exists^\infty x_0 \phi$ by the semantics for FOL^∞.

Induction step: we inductively assume that for any $Q_0, \ldots, Q_k \in \{\exists, \exists^\infty\}$, for *any* structure $\mathfrak{M}$ and assignment σ, $\mathfrak{M}, \sigma \models Q_0 x_0 \ldots Q_k x_k \phi$ iff there exists a function $f : \mathbb{F}(q_0, \ldots, q_k) \to M$ satisfying conditions 1. and 2. We show that this statement also holds for $k+1$. There are two subcases to consider.

Case 1: $Q_0 = \exists$. Suppose $\mathfrak{M}, \sigma \models \exists x_0 Q_1 x_1 \ldots Q_{k+1} x_{k+1} \phi$. Then there is some $a \in M$ such that $\mathfrak{M}, \sigma[a/x_0] \models Q_1 x_1 \ldots Q_{k+1} x_{k+1} \phi$. By IH, there is a function $f : \mathbb{F}(q_1, \ldots, q_{k+1}) \to M$ (with each q_i the index for Q_i) satisfying conditions 1. and 2.

Now let t_0 denote the root of $\mathbb{F}(q_0 = 1, q_1, \ldots, q_{k+1})$, and identify $\mathbb{F}(q_1, \ldots, q_{k+1})$ with the structure induced by $\mathbb{F}(1, q_1, \ldots, q_{k+1}) \setminus \{t_0\}$. Then define a function $f' : \mathbb{F}(1, q_1, \ldots, q_{k+1}) \to M$ by: $f'(t_0) = a$, $f'(t) = f(t)$ for all $t \neq t_0 \in \mathbb{F}(1, q_1, \ldots, q_{k+1})$. Then f' will satisfy conditions 1. and 2., as required.

Conversely, let f be a map from $\mathbb{F}(1, q_1, \ldots, q_{k+1}) \to M$ satisfying conditions 1. and 2. Let t_0 again denote the root of $\mathbb{F}(1, q_1, \ldots, q_{k+1})$, and let a denote $f(t_0)$.

Consider the map f' resulting from restricting f to the structure induced by $\mathbb{F}(1, q_1, \ldots, q_{k+1}) \setminus \{t_0\}$. For each path $t_1, \ldots, t_{k+1}$ through $\mathbb{F}(1, q_1, \ldots, q_{k+1}) \setminus \{t_0\}$, we have $\mathfrak{M}, \sigma[a/x_0, f(t_1)/x_1, \ldots, f(t_{k+1})/x_{k+1}] \models \phi$. Then by IH, $\mathfrak{M}, \sigma[a/x_0] \models Q_1 x_1 \ldots Q_{k+1} x_{k+1} \phi$. It follows that $\mathfrak{M}, \sigma \models \exists x_0 Q_1 x_1 \ldots Q_{k+1} x_{k+1} \phi$, as required.

Case 2: $Q_0 = \exists^\infty$. Suppose $\mathfrak{M}, \sigma \models \exists^\infty x_0 Q_1 x_1 \ldots Q_{k+1} x_{k+1} \phi$. Then there is an infinite set $A \subseteq M$ such that for each $a \in A, \mathfrak{M}, \sigma[a/x_0] \models Q_1 x_1 \ldots Q_{k+1} x_{k+1} \phi$. WLOG we can assume A is countable (otherwise, take a countable subset of A). Enumerate A as $\{a_0, a_1, \ldots\}$. By IH, for each $a_i \in A$, there is some map, $f_i : \mathbb{F}(q_0, \ldots, q_k) \to M$, satisfying condition 1. and such that for any path $t_0, \ldots, t_k$ through this tree,

$$\mathfrak{M}, \sigma[a_i/x_0, f_i(t_0)/x_1, \ldots, f_i(t_k)/x_{k+1}] \models \phi$$

We now define a map f on $\mathbb{F}(\omega, q_1, \ldots, q_{k+1})$. Let

$$T_0 := \{t \in \mathbb{F}(\omega, q_1, \ldots, q_{k+1}) : t \text{ has no predecessor}\}$$

$|T_0| = \omega$, so enumerate it as $\{t_0^0, t_0^1, \ldots\}$. For each $t_0^i \in T_0$, define $f(t_0^i) := a_i$. Then, for each $t_0^i \in T_0$, identify $\mathbb{F}(\omega, q_1, \ldots, q_{k+1}) \upharpoonright \uparrow(t_0^i)$ with $\mathbb{F}(q_0, \ldots, q_k)$, and define $f \upharpoonright \uparrow(t_0^i) := f_i$. Then f will satisfy conditions 1. and 2., as required.

Conversely, let $f : \mathbb{F}(\omega, q_1, \ldots, q_{k+1}) \to M$ be a map satisfying conditions 1. and

2. Let T_0 denote the set of $\mathbb{F}(\omega, q_1, \dots, q_{k+1})$-elements with no predecessors, and enumerate T_0 as $\{t_0^0, t_0^1, \dots\}$.

For each $i \in \mathbb{N}$, let f_i denote the map $f \restriction \uparrow(t_0^i)$. Now consider an arbitrary $t_0^i \in T_0$. Let a_i denote $f(t_0^i)$. Observe that for each path $t_1, \dots, t_{k+1}$ through $\uparrow(t_0^i)$, we have

$$\mathfrak{M}, \sigma[a_i/x_0, f(t_1)/x_1, \dots, f(t_{k+1})/x_{k+1}] \models \phi$$

By IH, we may conclude that $\mathfrak{M}, \sigma[a_i/x_0] \models Q_1 x_1 \dots Q_{k+1} x_{k+1} \phi$.

Since i was arbitrary, we may conclude that the above holds for every $a_i \in f[t_0]$. Also, $f[t_0]$ is infinite, because f is injective on $\sim_{\mathbb{F}}$ equivalence classes. Hence $\mathfrak{M}, \sigma \models \exists^\infty x_0 Q_1 x_1 \dots Q_{k+1} x_{k+1} \phi$ as required.

This completes the induction. □

We now state another useful fact about the semantics of *FOL*$^\infty$. This fact illustrates a certain similarity between the infinity quantifier and the regular existential quantifier: like $\exists$, $\exists^\infty$ distributes over disjunction.

Proposition 2.15 (Distributivity of $\exists^\infty$)**.** *Let $\phi, \psi \in FOL^\infty$. Then $\exists^\infty x(\phi \vee \psi)$ is logically equivalent to $\exists^\infty x\phi \vee \exists^\infty x\psi$.*

Proof. The implication from $\exists^\infty x\phi \vee \exists^\infty x\psi$ to $\exists^\infty x(\phi \vee \psi)$ is trivial, so we only consider the other direction. Suppose $\mathfrak{M}, \sigma \models \exists^\infty x(\phi \vee \psi)$. Let $A \subseteq M$ be an infinite subset of M, such that for every $a \in A$, $\mathfrak{M}, \sigma[a/x] \models \phi \vee \psi$. Then for each $a \in A$, either $\mathfrak{M}, \sigma[a/x] \models \phi$ or $\mathfrak{M}, \sigma[a/x] \models \psi$.

Suppose the set

$$A' := \{a' \in A : \mathfrak{M}, \sigma[a'/x] \models \phi\}$$

is infinite. Then $\mathfrak{M}, \sigma \models \exists^\infty x\phi$. On the other hand, suppose the set A' is finite. Then $A \setminus A'$ is infinite, and for every $a \in A \setminus A'$, we have $\mathfrak{M}, \sigma[a/x] \models \psi$, hence $\mathfrak{M}, \sigma \models \exists^\infty x\psi$. In either case, $\mathfrak{M}, \sigma \models \exists^\infty x\phi \vee \exists^\infty x\psi$, as required. □

We now show that GT^∞ may be viewed as a fragment of *FOL*$^\infty$, by providing a map from the well-formed formulas of GT^∞ to the well-formed formulas of *FOL*$^\infty$. We assume an enumeration of our countable stock of variables: $\{x_0, x_1, \dots\}$. We also assume an enumeration of our proposition letters P, and introduce a corresponding enumerated set of unary predicate letters, $\{P_0, P_1, \dots\}$

Definition 2.16 (Standard Translation)**.** We extend the usual standard translation of modal logic into *FOL*, by adding clauses to handle the infinity modalities. Let x_i be a variable, $\phi \in GT^\infty$. We define $ST(x_i, \phi)$ by induction on the construction of ϕ.

- $ST(x_i, p_j) = P_j x_i$

- $ST(x_i, \bot) = \bot$
- $ST(x_i, \phi \wedge \psi) = ST(x_i, \phi) \wedge ST(x_i, \psi)$, $ST(x_i, \neg\phi) = \neg ST(x_i, \phi)$
- $ST(x_i, \Diamond\phi) = \exists x_{i+1}(Rx_ix_{i+1} \wedge ST(x_{i+1}, \phi))$
- $ST(x_i, \Diamond^{\infty}\phi) = \exists^{\infty} x_{i+1}(Rx_ix_{i+1} \wedge ST(x_{i+1}, \phi))$
- $ST(x_i, \Diamond_{\leftarrow}\phi) = \exists x_{i+1}(Rx_{i+1}x_i \wedge ST(x_{i+1}, \phi))$
- $ST(x_i, \Diamond^{\infty}_{\leftarrow}\phi) = \exists^{\infty} x_{i+1}(Rx_{i+1}x_i \wedge ST(x_{i+1}, \phi))$
- $ST(x_i, E\phi) = \exists x_{i+1}(ST(x_{i+1}, \phi))$

The key fact about this translation is that the translated formulas are satisfied in the same structures as their originals. More precisely, let $\tau = \{P_0, P_1, \ldots\} \cup \{R\}$. Given a Kripke model $\mathcal{M} = (S, R, V)$, we define a τ-structure $\mathfrak{M} = \langle D, I \rangle$ by: $D = S, I(R) = R \subseteq S \times S$, and for each $i \in \mathbb{N}$, $I(P_i) = V(p_i)$.

Proposition 2.17. *For any $\phi \in GT^{\infty}$, for any Kripke model $\mathcal{M}$ and any $s \in S$, we have: $\mathcal{M}, s \Vdash \phi$ iff $\mathfrak{M}, \sigma[s/x_i] \models ST(x_i, \phi)$.*

The same mapping, $ST(x_i, \cdot)$, shows that ML^{∞}, T^{∞} and G^{∞} may all also be viewed as fragments of FOL^{∞}, since they are all contained within GT^{∞}.

In view of this proposition, we will sometimes be less precise in the remainder of the thesis, and view Kripke models as τ-structures, and vice versa. Thus, we might write things like $\mathfrak{M}, s \Vdash \phi$, with ϕ a GT^{∞}-formula, to mean that $\mathcal{M}, s \Vdash \phi$; or $\mathcal{M}, s \models ST(x_i, \phi)$, to mean that $\mathfrak{M}, s \models ST(x_i, \phi)$.

2.3 Bisimulations

One of the most important concepts for any modal system is that of *bisimulation.* We provide a game-theoretic definition of a bisimilarity relation for each of the modal fragments of FOL^{∞} defined thus far.

Definition 2.18 (Infinity bisimulation game)**.** Following Definition 3.2 of [1], we define a bisimulation game, Bis, played by two players, Spoiler and Duplicater. The game is played over two pointed Kripke models, $\mathcal{M}_0, s_0$ and $\mathcal{M}_1, s_1$, in a series of rounds. Each round has a *configuration* $(s_0, s_1) \in S_0 \times S_1$ (we also refer to these s_0 and s_1 as the *focus elements*). In each round, Spoiler moves first and is permitted to make two types of move:

- *Forward move*: Spoiler chooses a structure, $\mathcal{M}_i$, and selects an R_i-successor s'_i of s_i. Duplicater must respond by selecting an R_{1-i}-successor s'_{1-i} of s_{1-i} in the corresponding structure $\mathcal{M}_{1-i}$. The game moves on to the next round, with the configuration (s'_i, s'_{1-i}).

- *Forward infinity move*: Spoiler chooses a structure $\mathcal{M}_i$, and selects an infinite set X_i of R_i-successors of the focus element s_i. Duplicater must respond by selecting an infinite set X_{1-i} of R_{1-i}-successors of s_{1-i}. Spoiler then selects an element $s'_{1-i} \in X_{1-i}$, and in response, Duplicater must select an element $s'_i \in X_i$. The game moves on to the next round, with the configuration (s'_i, s'_{1-i}).

Spoiler wins the game if Duplicater is ever unable to make a legal move, or if a configuration (s_0, s_1) is reached such that s_0 and s_1 do not satisfy the same proposition letters from P. Otherwise, Duplicater wins.

We write $\mathsf{Bis}(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$ to denote the game played on $\mathcal{M}_0, \mathcal{M}_1$ which begins with a round with configuration (s_0, s_1), and continues for indefinitely many rounds. We write $\mathsf{Bis}^k(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$ to denote the game over the same structures and with the same opening-round configuration, but played for exactly k rounds. We write $\mathcal{M}_0, s_0 \simeq \mathcal{M}_1, s_1$ iff Duplicater has a winning strategy in the game $\mathsf{Bis}(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$, and $\mathcal{M}_0, s_0 \simeq^k \mathcal{M}_1, s_1$ iff Duplicater has a winning strategy in the game $\mathsf{Bis}^k(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$.

Remark. We have not included an infinity symbol in our notation for the $\simeq$ relation, because this thesis is exclusively concerned with logics with an infinity modality, so no confusion is likely to ensue, and it keeps notation lighter. Cf. the remark after Definition 2.11.

Definition 2.19 (Two-way infinity bisimulation game)**.** The two-way infinity bisimulation game, Bis_T, is like Bis, except that Spoiler is allowed to make the following additional types of move each round:

- *Backward move*: Spoiler chooses a structure, $\mathcal{M}_i$, and selects an R_i-predecessor s'_i of s_i. Duplicater must respond by selecting an R_i-predecessor s'_{1-i} of s_{1-i} in the corresponding structure $\mathcal{M}_{1-i}$. The game moves on to the next round, with the configuration (s'_i, s'_{1-i}).

- *Backward infinity move*: Spoiler chooses a structure $\mathcal{M}_i$, and selects an infinite set X_i of R_i-predecessors of the focus element s_i. Duplicater must respond by selecting an infinite set X_{1-i} of R_{1-i}-predecessors of s_{1-i}. Spoiler then selects an element $s'_{1-i} \in X_{1-i}$, and in response, Duplicater must select an element $s'_i \in X_i$. The game moves on to the next round, with the configuration (s'_i, s'_{1-i}).

We let $\mathsf{Bis}_T(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$ denote the Bis_T game played on $\mathcal{M}_0, \mathcal{M}_1$ with starting configuration (s_0, s_1), and we let $\mathsf{Bis}_T^k(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$ denote the k-round version of this game. We write $\mathcal{M}_0, s_0 \simeq_T \mathcal{M}_1, s_1$ iff Duplicater has a winning strategy in $\mathsf{Bis}_T(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$, and $\mathcal{M}_0, s_0 \simeq_T^k \mathcal{M}_1, s_1$ iff Duplicater has a winning strategy in the k-round version of this game.

Definition 2.20 (Global infinity bisimulation game)**.** The global infinity bisimulation game, Bis_G, is played on two Kripke models, $\mathcal{M}_0, \mathcal{M}_1$. The game begins with an *initial round*, with the following rules:

- Global move: Spoiler chooses a model, $\mathcal{M}_i$, and may select any element s_i from S_i. Duplicater must select an element s_{1-i} from S_{1-i}. The new configuration is (s_i, s_{1-i}).

After the initial round, the game continues with the players simply playing the game $\mathsf{Bis}(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$. We call (s_0, s_1) the initial configuration, since it is the first configuration reached in the game.

The game Bis_G^k is defined similarly, except that after the initial round, the players play the game $\mathsf{Bis}^k(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$ instead of $\mathsf{Bis}(\mathcal{M}_0, s_0\mathcal{M}_1, s_1)$. (Note that this means the game Bis_G^k actually lasts for $k+1$ rounds- one of which is the initial round.)

We write $\mathcal{M}_0 \simeq_G \mathcal{M}_1$ iff Duplicater has a winning strategy in $\mathsf{Bis}_G(\mathcal{M}_0, \mathcal{M}_1)$, and $\mathcal{M}_0 \simeq_G^k \mathcal{M}_1$ iff Duplicater has a winning strategy in $\mathsf{Bis}_G^k(\mathcal{M}_0, \mathcal{M}_1)$.

Remark. The Bis_G game would correspond more closely to the syntax of G^∞ if global moves were allowed in *every* round of the game, not just the initial round. However, it is straightforward to verify that this way of defining the game would be equivalent to the definition given above: allowing global moves in later rounds would not actually make the game any easier for Spoiler. Further, the definition given above makes proofs about the $\simeq_G$ relation simpler, so we adopt it here.

A separate question is why we do not allow versions of the game which start from a particular configuration, (s_0, s_1), and where Spoiler may play a non-global move (e.g. a forwards move) in the opening round. There is a technical reason behind this, but it is not of much interest, and therefore we choose not to explain it in detail.

Definition 2.21 (Global two-way infinity bisimulation game)**.** The global two-way infinity bisimulation game, $\mathsf{Bis}_{GT}(\mathcal{M}_0, \mathcal{M}_1)$,, is just like Bis_G, except that after the initial round, the players play $\mathsf{Bis}_T(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$ instead of $\mathsf{Bis}(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$. Analogously, Bis_{GT}^k is just like Bis_G^k except that after the initial round, Bis_T^k is played instead of Bis^k.

We write $\mathcal{M}_0 \simeq_{GT} \mathcal{M}_1$ iff Duplicater has a winning strategy in $\mathsf{Bis}_{GT}(\mathcal{M}_0, \mathcal{M}_1)$, and $\mathcal{M}_0 \simeq^k \mathcal{M}_1$ iff Duplicater has a winning strategy in $\mathsf{Bis}_{GT}^k(\mathcal{M}_0, \mathcal{M}_1)$.

Analogous remarks to those made after Definition 2.20 also apply here.

These bisimilarity relations have each been cooked up to match the different expressivity features of the modal languages ML^∞, T^∞, G^∞, and GT^∞. The following proposition partially illustrates this correspondence.

Proposition 2.22 (Bisimulation invariance). *ML^∞, T^∞, G^∞, GT^∞ are each invariant under their corresponding notion of bisimulation. That is:*

1. *For any $\phi \in ML^\infty$, for any $\mathcal{M}_0, s_0 \simeq \mathcal{M}_1, s_1$, $\mathcal{M}_0, s_0 \Vdash \phi$ iff $\mathcal{M}_1, s_1 \Vdash \phi$.*
2. *For any $\phi \in T^\infty$, for any $\mathcal{M}_0, s_0 \simeq_T \mathcal{M}_1, s_1$, $\mathcal{M}_0, s_0 \Vdash \phi$ iff $\mathcal{M}_1, s_1 \Vdash \phi$.*
3. *For any $\phi \in G^\infty$ such that $ST(x, \phi)$ has no free variables, for any $\mathcal{M}_0 \simeq_G \mathcal{M}_1$, $\mathcal{M}_0 \Vdash \phi$ iff $\mathcal{M}_1 \Vdash \phi$.*
4. *For any $\phi \in GT^\infty$ such that $ST(x, \phi)$ has no free variables, for any $\mathcal{M}_0 \simeq_{GT} \mathcal{M}_1$, $\mathcal{M}_0 \Vdash \phi$ iff $\mathcal{M}_1 \Vdash \phi$.*

Proof. For a proof of 1, we refer the reader to [1]. We prove item 2 by induction on the construction of ϕ.

Base case: $\phi = p$ for some $p \in \mathsf{P}$. If s_0 and s_1 did not agree on all proposition letters, then Duplicater would immediately lose the game $\mathsf{Bis}_T(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$. So s_0, s_1 must agree on proposition letters, and in particular $\mathcal{M}_0, s_0 \Vdash p$ iff $\mathcal{M}_1, s_1 \Vdash p$.

Base case: $\phi = \bot$. Then neither s_0 nor s_1 satisfies $\bot$, so $\mathcal{M}_0, s_0 \Vdash \bot$ iff $\mathcal{M}_1, s_1 \Vdash \bot$.

The induction step for the Boolean connectives is trivial, and is therefore omitted for brevity.

Induction step: $\phi = \Diamond_{\leftarrow}\psi$. The induction hypothesis is that for any $\mathcal{M}_0, t_0 \simeq_T \mathcal{M}_1, t_1$, $\mathcal{M}_0, t_0 \Vdash \psi$ iff $\mathcal{M}_1, t_1 \Vdash \psi$.

Suppose $\mathcal{M}_0, s_0 \Vdash \Diamond_{\leftarrow}\psi$. Then there is $t_0 \in S_0$ such that $t_0 R s_0$ and $\mathcal{M}_0, t_0 \Vdash \psi$. Since $t_0 R s_0$, t_0 would be a legal 'backwards move' for Spoiler in $\mathsf{Bis}_T(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$. Since $\mathcal{M}_0, s_0 \simeq_T \mathcal{M}_1, s_1$, there must be some winning response $t_1 \in R^{-1}[s_1]$ for Duplicater in response to Spoiler selecting t_0. We now have $\mathcal{M}_0, t_0 \Vdash \psi$ and $\mathcal{M}_0, t_0 \simeq_T \mathcal{M}_1, t_1$, so by the inductive hypothesis, $\mathcal{M}_1, t_1 \Vdash \psi$. Since $t_1 R s_1$, we have $\mathcal{M}_1, s_1 \Vdash \Diamond_{\leftarrow}\psi$, as required.

The argument for the converse direction is analogous.

Induction step: $\phi = \Diamond_{\leftarrow}^{\infty}\psi$. The induction hypothesis is that for any $\mathcal{M}_0, t_0 \simeq_T \mathcal{M}_1, t_1$, $\mathcal{M}_0, t_0 \Vdash \psi$ iff $\mathcal{M}_1, t_1 \Vdash \psi$.

Suppose $\mathcal{M}_0, s_0 \Vdash \Diamond^{\infty}_{\leftarrow}\psi$. Then there is an infinite set T_0 of R-predecessors of s_0, such that for each $t_0 \in T_0, \mathcal{M}_0, t_0 \Vdash \psi$. Further, Spoiler would be able to select T_0 as a 'backwards infinity move' in the game $\mathsf{Bis}_T(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$. Let $T_1 \subseteq R^{-1}[s_1]$ be the set prescribed as a response for Duplicater by the winning strategy for Duplicater in this game.

Now let $t_1 \in T_1$. Then t_1 would be a legal follow-up for Spoiler to his backward infinity move. Let t_0 be the element prescribed by the winning strategy for Duplicater in response to this follow-up. Then $\mathcal{M}_0, t_0 \simeq_T \mathcal{M}_1, t_1$. By IH, it follows that $\mathcal{M}_1, t_1 \Vdash \psi$. But t_1 was arbitrary, so we can conclude that every member of T_1 satisfies ψ. Since T_1 was infinite, it follows that $\mathcal{M}_1, s_1 \Vdash \Diamond^{\infty}_{\leftarrow}\psi$.

The argument for the converse direction is analogous.

The induction steps where $\phi = \Diamond\psi$ and $\phi = \Diamond^{\infty}\psi$ may be taken care of by the same reasoning as the $\Diamond_{\leftarrow}$ and $\Diamond^{\infty}_{\leftarrow}$ cases, with 'predecessor' replaced by 'successor'. This concludes the induction, and completes the proof of item 2. □

We now turn to item 3. First of all, it is easy to see that if $\phi \in G^{\infty}$ is such that $ST(x, \phi)$ has no free variables, then ϕ is logically equivalent to a Boolean combination of formulas from the set $\{E\psi : \psi \in ML^{\infty}\}$. We therefore assume that ϕ is written in this form, and argue by induction on the construction of ϕ, as follows.

Base case: $\phi = E\psi$, with $\psi \in ML^{\infty}$. Suppose $\mathcal{M}_0 \Vdash E\psi$. Then there is some $s_0 \in S_0$ such that $\mathcal{M}_0, s_0 \Vdash \psi$. Further, s_0 would be a legal selection for Spoiler in the 'initial round' of the game $\mathsf{Bis}_G(\mathcal{M}_0, \mathcal{M}_1)$. Since $\mathcal{M}_0 \simeq_G \mathcal{M}_1$, there must be some winning response s_1 for Duplicater in reply to Spoiler choosing s_0. We then have $\mathcal{M}_0, s_0 \simeq \mathcal{M}_1, s_1$, so by item 1, $\mathcal{M}_1, s_1 \Vdash \psi$. Hence $\mathcal{M}_1 \Vdash E\psi$, which was to show.

The induction steps are: $\phi = \psi \wedge \chi, \phi = \neg\psi$. Both arguments are routine, so we omit them. This concludes the induction, and completes the proof of item 3.

The proof of item 4 is identical to the proof of item 3, except that we rewrite ϕ as a Boolean combination of formulas from $\{E\psi : \psi \in T^{\infty}\}$ and use item 2 rather than item 1 in the proof of the base case. □

We round off this chapter with some facts about our k-step bisimilarity relations, and the modal-depth-k- fragments of the corresponding logics. These facts will be very important later on.

Definition 2.23 (Modal depth)**.** Let $\phi \in GT^{\infty}$. We let $md(\phi)$ denote the *modal depth* of ϕ, which is defined inductively by:

- $md(p) = md(\bot) = 0$ for all $p \in \mathsf{P}$
- $md(\neg\psi) = md(\psi)$
- $md(\psi \wedge \chi) = \max(md(\psi), md(\chi))$
- $md(\heartsuit\psi) = md(\psi) + 1$ for each $\heartsuit \in \{\Diamond, \Diamond^\infty, \Diamond_\leftarrow, \Diamond_\leftarrow^\infty, E\}$

We let GT_n^∞ denote the fragment $\{\phi \in GT^\infty : md(\phi) \leq n\}$. We adopt analogous notation for T^∞ and G^∞.

Definition 2.24 (FOL_n^∞)**.** For each $n \in \mathbb{N}$, we let FOL_n^∞ denote $\{\phi \in FOL^\infty : qd(\phi) \leq n\}$.

Definition 2.25 (Depth n equivalence)**.** Let $\mathcal{M}_0, s_0$ and $\mathcal{M}_1, s_1$ be pointed Kripke structures. We write: $\mathcal{M}_0, s_0 \equiv_T^n \mathcal{M}_1, s_1$ iff for all $\phi \in T_n^\infty$, $\mathcal{M}_0, s_0 \Vdash \phi$ iff $\mathcal{M}_1, s_1 \Vdash \phi$.

We write $\mathcal{M}_0 \equiv_G^n \mathcal{M}_1$ iff for every $\phi \in G_{n+1}^\infty$ such that $ST(x, \phi)$ is a sentence, $\mathcal{M}_0 \Vdash \phi$ iff $\mathcal{M}_1 \Vdash \phi$. We write $\mathcal{M}_0 \equiv_{GT}^n \mathcal{M}_1$ iff for every $\phi \in GT_{n+1}^\infty$ such that $ST(x, \phi)$ is a sentence, $\mathcal{M}_0 \Vdash \phi$ iff $\mathcal{M}_1 \Vdash \phi$.

Proposition 2.26. *Let τ be an arbitrary finite modal signature (i.e., a finite subset of the proposition letters P). Then for any $n \in \mathbb{N}$, ML^∞, T_n^∞, G_n^∞ and GT_n^∞ are all finite up to logical equivalence relative to τ.*

Proof. By Lemma 4.19 of [1], FOL_n^∞ is finite up to logical equivalence relative to any finite signature. But any $\phi \in GT_n^\infty$ can be expressed in FOL_n^∞, by the Standard Translation. □

Proposition 2.27. *Let τ be a fixed finite modal signature. Then*

1. *For any $n \in \mathbb{N}$, for any pointed τ-structures $\mathcal{M}_0, s_0, \mathcal{M}_1, s_1$, $\mathcal{M}_0, s_0 \simeq_T^n \mathcal{M}_1, s_1$ iff $\mathcal{M}_0, s_0 \equiv_T^n \mathcal{M}_1, s_1$*
2. *For any $n \in \mathbb{N}$, for any τ-structures $\mathcal{M}_0, \mathcal{M}_1$, $\mathcal{M}_0 \simeq_G^n \mathcal{M}_1$ iff $\mathcal{M}_0 \equiv_G^n \mathcal{M}_1$, and $\mathcal{M}_0 \simeq_{GT}^n \mathcal{M}_1$ iff $\mathcal{M}_0 \equiv_{GT}^n \mathcal{M}_1$*

Proof. In each case, we only consider the direction from modal equivalence to bisimulation. The converse direction goes by a routine induction argument analogous to the proof of Proposition 2.22.

We begin with item 1. We argue by induction on n.

Base case: $n = 0$. Since $\mathcal{M}_0, s_0 \equiv_T^0 \mathcal{M}_1, s_1$, we have that s_0 satisfies all the same proposition letters as s_1. But then Duplicater wins the game $\mathsf{Bis}_T^0(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$ automatically, so $\mathcal{M}_0, s_0 \simeq_T^0 \mathcal{M}_1, s_1$.

Induction step: we inductively assume that for any $\mathcal{M}_0, s_0, \mathcal{M}_1, s_1$, if $\mathcal{M}_0, s_0 \equiv^n_T \mathcal{M}_1, s_1$ then $\mathcal{M}_0, s_0 \simeq^n_T \mathcal{M}_1, s_1$. We show that if $\mathcal{M}_0, s_0 \equiv^{n+1}_T \mathcal{M}_1, s_1$ then $\mathcal{M}_0, s_0 \simeq^{n+1}_T \mathcal{M}_1, s_1$.

Let $\mathcal{M}_0, s_0 \equiv^{n+1}_T \mathcal{M}_1, s_1$. We must show that Duplicater has a winning strategy in $\mathsf{Bis}^n_T(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$. We consider the possible moves Spoiler can make from the starting configuration.

- Forward move: Without loss of generality, suppose Spoiler chooses $\mathcal{M}_0$, and selects a successor t_0 of s_0. Since T^∞_n is finite up to logical equivalence, we may write a characteristic formula χ which completely characterises t_0 (with respect to T^∞_n). We then have $\mathcal{M}_0, s_0 \Vdash \Diamond\chi$. Since $\mathcal{M}_0, s_0 \equiv^{n+1}_T \mathcal{M}_1, s_1$, we also have $\mathcal{M}_1, s_1 \Vdash \Diamond\chi$. So there is some $t_1 \in R[s_1]$ such that $\mathcal{M}_1, t_1 \Vdash \chi$.

  Since χ fully characterises t_0 and t_1, we have $\mathcal{M}_0, t_0 \equiv^{n+1}_T \mathcal{M}_1, t_1$. Then by the inductive hypothesis, $\mathcal{M}_0, t_0 \simeq^{n+1}_T \mathcal{M}_1, t_1$. So t_1 is a winning response for Duplicater in response to Spoiler choosing t_0.

- Forward infinity move: Suppose Spoiler selects some infinite set $T_0 \subseteq R[s_0]$. The fact that T^∞_n is finite up to logical equivalence immediately implies that $\equiv^n_T$ has finite index. So there must be some infinite set $T'_0 \subseteq T_0$ such that for all $x, y \in T'_0$, $\mathcal{M}_0, x \equiv^n_T \mathcal{M}_0, y$. Again using the fact that T^∞_n is finite up to logical equivalence, let χ be a characteristic formula which completely determines the T^∞_n-theory of each $x \in T'_0$.

  We now have $\mathcal{M}_0, s_0 \Vdash \Diamond^\infty\chi$, and therefore $\mathcal{M}_1, s_1 \Vdash \Diamond^\infty\chi$. So there is some $T_1 \subseteq R[s_1]$ such that for all $t_1 \in T_1$, $\mathcal{M}_1, t_1 \Vdash \chi$. Duplicater should select the set T_1. Whichever $t_1 \in T_1$ is selected by Spoiler, Duplicater may select any element t_0 of T'_0. Because t_0, t_1 both satisfy χ, we then have $\mathcal{M}_0, t_0 \equiv^n_T \mathcal{M}_1, t_1$. So by the inductive hypothesis, we then have $\mathcal{M}_0, t_0 \simeq^n_T \mathcal{M}_1, t_1$. So we have illustrated a winning response for Duplicater.

- The cases where Spoiler plays a backward move or a backwards infinity move are analogous, using the expressive power of $\Diamond_\leftarrow$ and $\Diamond^\infty_\leftarrow$.

We can conclude that Duplicater has a winning strategy in $\mathsf{Bis}^{n+1}_T(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$, which was to show. This completes the induction. $\square$

We now turn to item 2. We omit the argument for G^∞, since it is analogous to the argument for GT^∞, which we now show.

Suppose $\mathcal{M}_0 \equiv^n_{GT} \mathcal{M}_1$. We give a winning strategy for Duplicater in the game $\mathsf{Bis}^n_{GT}(\mathcal{M}_0, \mathcal{M}_1)$ by showing how Duplicater should respond to Spoiler's initial, global move. Without loss of generality, suppose Spoiler selects the structure $\mathcal{M}_0$, and selects an element s_0. Since T_n^∞ is finite, let $\chi \in T_n^\infty$ be a formula which fully characterises $\mathcal{M}_0, s_0$ relative to T_n^∞. Then $\mathcal{M}_0 \Vdash E\chi$, which implies $\mathcal{M}_1 \Vdash E\chi$ since $\mathcal{M}_0 \equiv_{GT}^n \mathcal{M}_1$. This means there is some $s_1 \in S_1$ such that $\mathcal{M}_1, s_1 \Vdash \chi$. Duplicater should respond to Spoiler's move by selecting s_1. Then since s_0 and s_1 both satisfy χ, we have $\mathcal{M}_0, s_0 \equiv_T^n \mathcal{M}_1, s_1$. By item 1 this implies that Duplicater has a winning strategy in $\mathsf{Bis}_T^n(\mathcal{M}_0, s_0, \mathcal{M}_1, s_1)$, as required.

We may conclude that $\mathcal{M}_0 \simeq_{GT}^n \mathcal{M}_1$. □

Proposition 2.28. *Let τ be a fixed finite modal signature, $n \in \mathbb{N}$. Then*

1. *The relation $\simeq_T^n$ has finite index on the class of pointed τ-structures*
2. *The relations $\simeq_G^n$ and $\simeq_{GT}^n$ both have finite index on the class of τ-structures*

Proof. Each result follows immediately from Propositions 2.26 and 2.27. □

Chapter 3

Gaifman Theorem for FOL^{∞}

In this chapter, we state the first of our two main results in this thesis, namely that the logics T^{∞}, G^{∞} and GT^{∞} may be characterised in terms of invariance under bisimulation. We give an overview of the proof method used to obtain analogous characterisation results in the finite model theory setting, and explain how this method may be adapted to handle our modal fragments of FOL^{∞}. A crucial component of the finite-model-theory strategy is Gaifman's Theorem for first-order logic, and in order to apply the method to fragments of FOL^{∞}, we need an analogous theorem for FOL^{∞}. The main body of this chapter will be concerned with proving this theorem.

3.1 Characterisating T^{∞}, G^{∞} and GT^{∞} through Bisimulation Invariance

Proposition 2.22 from the previous chapter shows that there is a connection between each of our modal systems, and its corresponding notion of bisimulation: each formula of T^{∞} is invariant under the relation $\simeq_T$, for example. A natural question to ask is whether this connection holds in the opposite direction. If a formula of FOL^{∞} is invariant under a particular bisimilarity relation, must it be equivalent to a formula in the modal language corresponding to this notion of bisimulation?

In the context of classical modal logic and classical FOL, the celebrated van Benthem Characterisation Theorem (see e.g. [5]) gives an affirmative answer to this question. Van Benthem's original proof uses the Compactness of first-order logic. However, as noted in Chapter 1, FOL^{∞} is not compact. Therefore, in order to prove analogues of the van Benthem Characterisation Theorem for fragments of FOL^{∞}, one needs a substantially different method from van Benthem's original approach.

Fortunately, alternative approaches have been developed in the setting of *finite model theory*. The finite-model analogue of van Benthem's theorem states that an *FOL*- formula is invariant over bisimulation (across finite models) if and only if it is equivalent (with respect to finite models) to a modal formula. Van Benthem's original proof does not go through for this statement, because it appeals to Compactness, and *FOL* is not compact with respect to the class of finite models. However, Rosen shows in [16] that the theorem still holds with respect to finite models. His approach is constructive: given a formula ϕ, he shows that if two pointed structures $\mathcal{M}_0, s_0, \mathcal{M}_1, s_1$ are n-step bisimilar (in the basic modal sense), then they must have bisimilar companion structures $\mathcal{M}'_0, s'_0, \mathcal{M}'_1, s'_1$ satisfying a stronger equivalence relation which must preserve the truth value of ϕ. If ϕ is also bisimulation invariant then ϕ must in fact be n-step bisimulation invariant, which implies that ϕ is modally definable. (This 'Upgrading' strategy will be explained in more detail in Chapter 4.)

Furthermore, Bellas Acosta in [1] adapts Rosen's approach to show that ML^∞ is characterised by invariance under the bisimilarity relation $\simeq$. He uses the same constructive upgrading idea to show that formulas of FOL^∞ which are invariant under $\simeq$, must be invariant under $\simeq^k$ for some $k \in \mathbb{N}$, which implies that they are definable within ML^∞_k. In the present work, we will see that Bellas Acosta's proof extends easily to the case of T^∞, showing that T^∞ is characterised by invariance under $\simeq_T$.

In the cases of G^∞ and GT^∞, there is a technical issue. Recall that in Proposition 2.22, we showed only that every formula of G^∞ *with no free variables* is invariant under bisimulation; likewise for GT^∞. As a technical convenience, we confine our attention in this thesis to characterising the formulas of G^∞ and GT^∞ whose standard translations have no free variables. So, the result we aim to prove is that for any *sentence* ϕ of FOL^∞, if ϕ is invariant under $\simeq_G$ then ϕ is expressible within G^∞, and if ϕ is invariant under $\simeq_{GT}$ then ϕ is expressible within GT^∞. However, speaking loosely, we may refer to these as 'characterisation results' for G^∞ and GT^∞, respectively.

We will indeed provide proofs of these characterisation results. However, the global modality makes matters trickier, and requires techniques going significantly beyond those used by Bellas Acosta in [1]. In order to prove characterisation results for these logics, we will adapt the methods of Otto in [15], who proved analogous results for modal fragments of *FOL* which include a global modality.

We may isolate two important components of Otto's approach. Firstly, he observes that by Gaifman's theorem for first-order logic, for every first-order formula ϕ, we can find a suitable equivalence relation which is, intuitively, 'finitely bounded', and which preserves the truth value of ϕ. The second component is another 'Upgrading' argument: given two structures which are k-step $\simeq_G$- or

$\simeq_{GT}$-bisimilar, we can find fully bisimilar companion structures which satisfy the stronger equivalence relation yielded by Gaifman's theorem. This means that invariance under bisimulation implies invariance under sufficiently large finite-step bisimulation. But classes of structures which are closed under finite-step bisimulation are modally definable, implying the desired result that formulas which are invariant under bisimulation, are modally definable.

We adopt the same strategy here. But Gaifman's theorem only covers classical first-order logic, whereas we are concerned with fragments of FOL^∞. Hence, our first task is to prove an analogue of Gaifman's theorem for FOL^∞. This task is taken up in the remainder of this chapter.

3.2 Preliminaries

In order to even state Gaifman's theorem, and its FOL^∞ analogue, we require various definitions. Our definitions closely follow those of Gaifman [9].

Definition 3.1 (Gaifman distance)**.** Let $\mathfrak{M}$ be a τ- structure, $a, b \in M$. The *Gaifman distance* between a and b in $\mathfrak{M}$, notation: $d^{\mathfrak{M}}(a, b)$, is defined inductively by the following clauses:

- $d^{\mathfrak{M}}(a, b) = 0$ iff $a = b$
- $d^{\mathfrak{M}}(a, b) \leq 1$ iff $a = b$, or there is some predicate symbol $R \in \tau$, and some tuple $\overline{c} \subseteq M$ with a, b among $\overline{c}$, such that $\mathfrak{M} \models R(\overline{c})$
- $d^{\mathfrak{M}}(a, b) \leq n + 1$ iff there is some $c \in M$ with $d^{\mathfrak{M}}(a, c) \leq n$ and $d^{\mathfrak{M}}(c, b) \leq 1$
- $d^{\mathfrak{M}}(a, b) = n + 1$ iff $d^{\mathfrak{M}}(a, b) \leq n + 1$, but it is not the case that $d^{\mathfrak{M}}(a, b) \leq n$
- $d^{\mathfrak{M}}(a, b) = \infty$ iff there is no $n \in \mathbb{N}$ such that $d^{\mathfrak{M}}(a, b) = n$

Further, for any τ-structure $\mathfrak{M}$, any tuple of objects $\overline{a} \subseteq M$ and any $b \in M$, we define $d^{\mathfrak{M}}(\overline{a}, b) := min(\{d^{\mathfrak{M}}(a, b) : a \in \overline{a}\})$.

Observe that, relative to any fixed finite signature τ and to the class of τ-structures, we may write a first-order formula $d^\tau(\overline{x}, y) \leq n$ expressing that the Gaifman distance from some objects $\overline{a}$ to an object b is less than or equal to n. For example, if we are concerned with the signature $\{P, R\}$, with P a unary and R a binary predicate, then

$$x = y \vee Rxy \vee \exists z((Rxz \wedge Rzy) \vee (Rxz \wedge Ryz) \vee (Rzx \wedge Rzy) \vee (Rzx \wedge Ryz))$$

expresses, relative to this signature, that the Gaifman distance from a to b is at most 2. Further, we may express that the Gaifman distance from $\overline{a}$ to b is strictly less than n, exactly n, at least n, or strictly more than n, using Boolean combinations of formulas of the form $d^\tau(\overline{a}, b) \leq m$, for appropriate m.

In this chapter, we will only be concerned with finite, relational signatures τ. We will therefore use $d^\tau(\overline{x}, y) \leq n$ to denote an FOL-formula which is satisfied by $\overline{a}, b$ in a τ- structure $\mathfrak{M}$ iff $d^{\mathfrak{M}}(\overline{a}, b) \leq n$. We will use $d^\tau(\overline{x}, y) < n$, $d^\tau(\overline{x}, y) = n$, $d^\tau(\overline{x}, y) \geq n$, and $d^\tau(\overline{x}, y) > n$ as abbreviations for appropriate Boolean combinations of formulas of the form $d(\overline{x}, y) \leq m$. In practice, we will often suppress mention of the signature τ.

Definition 3.2 (Neighbourhood)**.** Let $\mathfrak{M}$ be a τ- structure, $a \in M$. The k-*neighbourhood of* a *in* $\mathfrak{M}$ is the set

$$V^k_{\mathfrak{M}}(a) := \{b \in M : d^M(a, b) \leq k\}$$

Generalising this notation, for a finite tuple $\overline{a} \subseteq \mathfrak{M}$, we define

$$V^k_{\mathfrak{M}}(\overline{a}) := \{b \in M : d^{\mathfrak{M}}(\overline{a}, b) \leq k\}$$

We will sometimes drop the subscript $\mathfrak{M}$ if the intended structure is obvious.

We will often want to consider the substructure of $\mathfrak{M}$ induced by the neighbourhood of some object, or finite set of objects, in $\mathfrak{M}$. We therefore introduce a new piece of notation: where $\overline{a}$ is a finite tuple of objects from $\mathfrak{M}$, $k \in \mathbb{N}$, we let $\mathfrak{M}{\restriction}(\overline{a}, k)$ denote $\mathfrak{M}{\restriction}V^k_{\mathfrak{M}}(\overline{a})$, the structure whose universe is $V^k_{\mathfrak{M}}(\overline{a})$ and whose interpretation of n-ary predicates is given by $I(R) = R^{\mathfrak{M}} \cap (V^k_{\mathfrak{M}}(\overline{a}))^n$.

Definition 3.3 (Local formula)**.** Let $\phi \in FOL^\infty$ with $FV(\phi) = (v_1, \ldots, v_n)$. We say that ϕ is k-*local* iff, for any structure $\mathfrak{M}$ and any n-tuple $\overline{a} \subseteq M$:

$$\mathfrak{M}, \overline{a} \models \phi \iff \mathfrak{M}{\restriction}(\overline{a}, k), \overline{a} \models \phi$$

We say that ϕ is *local* iff there is some $k \in \mathbb{N}$ such that ϕ is k-local.

Observe that $\phi(\overline{v})$ is k-local iff it is equivalent to a formula $\phi'(\overline{v})$, such that $FV(\phi') = \overline{v} = FV(\phi)$, and all quantifiers in ϕ' are relativised to the k-neighbourhood of some $v \in \overline{v}$ (using the fact that 'being in the k-neighbourhood of' is first-order definable). We will freely make use of this equivalence in the remainder of this chapter.

Observe also that locality is closed under Boolean operations: a Boolean combination of k-local formulas is still k-local.

Definition 3.4 (Basic local sentences)**.** Let ϕ be a sentence of FOL^∞. We say that ϕ is a *basic local sentence* iff there is some signature τ, some $k, n \in \mathbb{N}$ and some k-local formula ψ with one free variable such that ϕ is of the syntactic form:

$$\exists v_1 \dots \exists v_n (\bigwedge_{1 \leq i < j \leq n} d^\tau(v_i, v_j) \geq k \wedge \bigwedge_{1 \leq i \leq n} \psi(v_i))$$

In this case we say that k is the *locality rank* of ϕ and n is its *scattering rank*.

We say that ϕ is a *basic local infinity sentence* iff there is some signature τ, some $k, n \in \mathbb{N}$ and some k-local formula ψ with one free variable such that ϕ is of the syntactic form:

$$\exists^\infty v_1 \dots \exists^\infty v_n (\bigwedge_{1 \leq i < j \leq n} d^\tau(v_i, v_j) \geq k \wedge \bigwedge_{1 \leq i \leq n} \psi(v_i))$$

Again, k is then the locality rank of ϕ, and s is its scattering rank.

We also use this section to state and prove some lemmas which will be useful later on.

Lemma 3.5. *For any $k \in \mathbb{N}$, the formula $d(\overline{x}, y) \leq 2k$ is k-local, and the formula $d(\overline{x}, y) \leq 2k + 1$ is $k + 1$-local.*

Proof. We first show that $d(\overline{x}, y) \leq 2k$ is k- local.

Suppose $d^{\mathfrak{M}}(\overline{a}, b) \leq 2k$. This means that, for some a_j among $\overline{a}$, there is a 'path' of elements $a_j = c_0, c_1, \dots, c_k, \dots, c_{2k-1}, c_{2k} = b$ such that, for each $i < 2k$, we have c_i, c_{i+1} and possibly some auxiliary elements $\overline{d_i}$ linked by some predicate R. Further, for each $i < k$, we have $c_i \in V^k(\overline{a})$ and $\overline{d_i} \subseteq V^k(\overline{a})$, and for each $i \geq k$, $c_i \in V^k(b)$ and $\overline{d_i} \subseteq V^k(b)$. Hence, all of the elements occurring as some c_i or among some $\overline{d_i}$ are within $V^k(\overline{a}, b)$. This means that the elements c_i and the auxiliary elements $\overline{d_i}$ can also witness that $\mathfrak{M}{\restriction}(\overline{a}, b, k) \models d(\overline{a}, b) \leq 2k$, which was to show.

The converse direction is immediate, by the observation that Gaifman distance does not increase when passing to extensions: if $d^{\mathfrak{M}}(\overline{a}, b) \leq m$ and $\mathfrak{M}$ is a submodel of $\mathfrak{M}'$, then also $d^{\mathfrak{M}'}(\overline{a}, b) \leq m$, because the 'witnesses' showing that $d^{\mathfrak{M}}(\overline{a}, b) \leq m$ will also exist in $\mathfrak{M}'$ and witness $d^{\mathfrak{M}'}(\overline{a}, b) \leq k$.

We now show that $d(\overline{x}, y) \leq 2k + 1$ is $k + 1$- local.

Suppose $d^{\mathfrak{M}}(\overline{a}, b) \leq 2k + 1$. As before, this means there is a 'path' of elements $a_j = c_0, c_1, \dots, c_k, c_{k+1}, \dots, c_{2k}, c_{2k+1} = b$ such that, for each $i < 2k + 1$, c_i, c_{i+1} and possibly some auxiliary elements $\overline{d_i}$ are linked by some predicate R. Further, for each $i \leq k$, $c_i \in V^k(\overline{a})$ and for each $i \geq k + 1$, $c_i \in V^k(b)$. Now, we might *not* have $\overline{d_k} \subseteq V^k(\overline{a}, b)$. The elements $\overline{d_k}$ link c_k to c_{k+1}, so they may be further from $\overline{a}$ than c_k and further from b than c_{k+1}. However, there will certainly be a path $a_j, c_1, \dots, c_k, d_k$ linking each $d_k \in \overline{d_k}$ to $\overline{a}$. So we have $\overline{d_k} \subseteq V^{k+1}(\overline{a}, b)$.

We may conclude that every element c_i and every element occurring within some $\overline{d}_i$ is within $\mathfrak{M}\restriction(\overline{a}, b, k+1)$. So, as before, we have $\mathfrak{M}\restriction(\overline{a}, b, k+1) \models d(\overline{a}, b) \leq 2k+1$, which was to show.

The converse direction is again immediate by the observation that Gaifman distance does not increase when passing to extensions. □

Observe that $d(\overline{x}, y) > m$ may be defined by $\neg d(\overline{x}, y) \leq m$, and recall that (for any m) Boolean operations preserve m-locality. So, we obtain as an easy corollary that $d(\overline{x}, y) > 2k$ is also k-local, and $d(\overline{x}, y) > 2k+1$ is also $k+1$-local.

Lemma 3.6. *Let $\psi(\overline{x}, \overline{y})$ be an l-local formula, and let $m \geq k+l$. Then*

1. *The formula $\phi_1 := \exists y_1 \in V^k(\overline{x}) \ldots \exists y_n \in V^k(\overline{x})\psi$ is m-local*

2. *The formula $\phi_2 := \exists^\infty y_1 \in V^k(\overline{x}) \ldots \exists^\infty y_n \in V^k(\overline{x})\psi$ is m-local*

Proof. We first show item 1. Let ϕ_1, ψ, k, l, m be as in the statement of the lemma. We must show that, for any structure $\mathfrak{M}$ and tuple $\overline{a} \subseteq M$, $\mathfrak{M}, \overline{a} \models \phi_1$ iff $\mathfrak{M}\restriction(\overline{a}, m), \overline{a} \models \phi_1$.

Suppose $\mathfrak{M}, \overline{a} \models \phi_1$. Then there are $b_1, \ldots, b_n$, each within $V^k_{\mathfrak{M}}(\overline{a})$, such that $\mathfrak{M}, \overline{a}, \overline{b} \models \psi$. Since ψ is l-local, we can infer that $\mathfrak{M}\restriction(\overline{a}, \overline{b}, l), \overline{a}, \overline{b} \models \psi$. We now claim that

$$\mathfrak{M}\restriction(\overline{a}, m)]\restriction(\overline{a}, \overline{b}, l) = \mathfrak{M}\restriction(\overline{a}, \overline{b}, l) \tag{3.1}$$

First note that $[\mathfrak{M}\restriction(\overline{a}, m)]\restriction(\overline{a}, \overline{b}, l)$ is well-defined, because $\overline{b} \subseteq V^m_{\mathfrak{M}}(\overline{a})$. It suffices to show that $V^l_{\mathfrak{M}\restriction(\overline{a},m)}(\overline{a}, \overline{b}) = V^l_{\mathfrak{M}}(\overline{a}, \overline{b})$. The left-to-right inclusion is immediate from the fact that Gaifman distance does not increase when passing to extensions. For the right-to-left inclusion, if $c \in V^l_{\mathfrak{M}}(\overline{a}, \overline{b})$ then there is some $d \in \overline{a} \cup \overline{b}$ such that $d^{\mathfrak{M}}(c, d) \leq l$. But by Lemma 3.5, it follows that $\mathfrak{M}\restriction(c, d, l) \models d(c, d) \leq l$. Since $\mathfrak{M}\restriction(\overline{a}, m)$ is an extension of $\mathfrak{M}\restriction(c, d, l)$, we have $\mathfrak{M}\restriction(\overline{a}, m) \models d(c, d) \leq l$ and hence $c \in V^l_{\mathfrak{M}\restriction(\overline{a},m)}(\overline{a}, \overline{b})$, as required. This proves the claim (3.1).

We now have $[\mathfrak{M}\restriction(\overline{a}, m)]\restriction(\overline{a}, \overline{b}, l), \overline{a}, \overline{b} \models \psi$. Since ψ is l-local, it follows that $\mathfrak{M}\restriction(\overline{a}, m), \overline{a}, \overline{b} \models \psi$. Finally, by Lemma 3.5 and because $k \leq m$, we have $\mathfrak{M}\restriction(\overline{a}, m) \models d(\overline{a}, b_i) \leq k$ for each b_i. This means that $\mathfrak{M}\restriction(\overline{a}, m), \overline{a} \models \phi_1$, which was to show.

Conversely, suppose $\mathfrak{M}\restriction(\overline{a}, m), \overline{a} \models \phi_1$. Then there are $b_1, \ldots, b_n \in V^k_{\mathfrak{M}\restriction(\overline{a},m)}(\overline{a})$ such that $\mathfrak{M}\restriction(\overline{a}, m), \overline{a}, \overline{b} \models \psi$. ψ is l-local, so $[\mathfrak{M}\restriction(\overline{a}, m)]\restriction(\overline{a}, \overline{b}, l), \overline{a}, \overline{b} \models \psi$. Now we may re-use (3.1) to infer that $\mathfrak{M}\restriction(\overline{a}, \overline{b}, l), \overline{a}, \overline{b} \models \psi$. But again, ψ is l-local, hence $\mathfrak{M}, \overline{a}, \overline{b} \models \psi$.
Finally, each $b_i \in V^k_{\mathfrak{M}}(\overline{a})$, because distance does not increase when passing to extensions. Hence $\mathfrak{M}, \overline{a} \models \phi_1$, as required.

We now move to item 2. We must show that, for any structure $\mathfrak{M}$ and tuple $\overline{a} \subseteq M$, $\mathfrak{M}, \overline{a} \models \phi_2$ iff $\mathfrak{M}\restriction(\overline{a}, m), \overline{a} \models \phi_2$.

Suppose $\mathfrak{M}, \overline{a} \models \phi_2$. Then by Observation 2.14 there is a function $f : \mathbb{F}((\omega)^n) \to V^k_{\mathfrak{M}}(\overline{a})$ such that f is injective on $\sim_{\mathbb{F}}$-equivalence classes, and for every path $t_1, \ldots, t_n$ through $\mathbb{F}((\omega)^n)$, $\mathfrak{M}, \overline{a}, f(t_1), \ldots, f(t_n) \models \psi$. Now let $(t_1, \ldots, t_n)$ be an arbitrary path through $\mathbb{F}((\omega)^n)$, and let b_1 denote $f(t_1)$, ..., b_n denote $f(t_n)$. We then have $\overline{b} \subseteq V^k(\overline{a})$ and $\mathfrak{M}, \overline{a}, \overline{b} \models \psi$. By the exact same argument as used for item 1 of this lemma, we can conclude that $\mathfrak{M}\restriction(\overline{a}, m), \overline{a}, \overline{b} \models \psi$.

Since $t_1, \ldots, t_n$ was arbitrary, the function f is such that for every path $t_1, \ldots, t_n$ through $\mathbb{F}((\omega)^n)$, we have

$$\mathfrak{M}\restriction(\overline{a}, m), \overline{a}, f(t_1), \ldots, f(t_n) \models \psi$$

By Observation 2.14 we may conclude that $\mathfrak{M}\restriction(\overline{a}, m), \overline{a} \models \phi_2$, which was to show.

The argument for the converse direction is analogous. We reuse the argument in the proof of the converse direction in item 1 to infer from $\mathfrak{M}\restriction(\overline{a}, m), \overline{a}, \overline{b} \models \psi$ to $\mathfrak{M}, \overline{a}, \overline{b} \models \psi$. $\square$

3.3 Statement of the Theorem

We now state Gaifman's original theorem, and its analogue for *FOL*$^\infty$ which will be proved in this chapter.

Theorem 3.7 (Gaifman's theorem for first-order logic). *Every first-order formula ϕ is equivalent to a Boolean combination of first-order local formulae, and basic local sentences.*

Theorem 3.8 (Gaifman theorem for *FOL*$^\infty$). *Every formula ϕ of FOL$^\infty$ is equivalent to a Boolean combination of local formulae, basic local sentences, and basic local infinity sentences.*

For a proof of Theorem 3.7, see Gaifman [9]. We will give a proof of Theorem 3.8 which closely follows Gaifman's original proof in [9]. (As a matter of fact, the proof we will give illustrates all the reasoning required in the proof of Theorem 3.7.)

3.4 Proof of Gaifman theorem for *FOL*$^\infty$

We prove Theorem 3.8 by induction on the construction of formulas. In order to handle one of the induction steps, we will need the following crucial lemma:

Lemma 3.9. *Let $\beta \in FOL^{\infty}$ be a k-local formula for some $k \in \mathbb{N}$, and let $FV(\beta) = (v_1, \ldots, v_m, z)$. Then*

(a) $\exists z\beta$ is logically equivalent to a Boolean combination of local formulae, and basic local sentences

(b) $\exists^{\infty} z\beta$ is logically equivalent to a Boolean combination of local formulae, and basic local infinity sentences.

Almost all the hard work required to prove Theorem 3.8 consists in proving this lemma. We will proceed by proving a long series of claims, eventually reducing the formulas in (a) and (b) to formulas in the desired form.

3.4.1 Proof of Lemma 3.9

We begin with item (a). Assume $\beta, k, \overline{v}, z$ are as in the statement of the lemma. $\exists z\beta(\overline{v}, z)$ is equivalent to

$$\exists z(d(\overline{v}, z) \leq 2k + 1 \wedge \beta) \vee \exists z(d(\overline{v}, z) > 2k + 1 \wedge \beta)$$

Let β_1 denote the first disjunct, β_2 the second. Recall that β is k-local. Therefore, by Lemma 3.6, β_1 is a $3k + 1$-local formula. So we have reduced the problem to showing that β_2 can be written in the desired form.

Let $\mathfrak{M}, \overline{a}, b$ satisfy $d(\overline{a}, b) > 2k+1$. Then the substructures $\mathfrak{M}{\restriction}(\overline{a}, k)$, $\mathfrak{M}{\restriction}(b, k)$ are disjoint. Call these structures $\mathfrak{M}_1$ and $\mathfrak{M}_2$, respectively. Then $\mathfrak{M}{\restriction}(\overline{a}, b, k)$ is the disjoint sum of $\mathfrak{M}_1$ and $\mathfrak{M}_2$.

Claim 1. There exist sets of formulas $\Gamma = \{\gamma_i(\overline{v}) : i \in I\}, \Delta = \{\delta_j(z) : j \in J\}$, and a formula $\rho(\beta)$ such that:

i. All members of Γ, Δ are k-local

ii. $\rho(\beta)$ is a Boolean combination of formulas from $\Gamma \cup \Delta$

iii. $\mathfrak{M}, \overline{a}, b \models \beta$ iff $\mathfrak{M}, \overline{a}, b \models \rho(\beta)$

Proof. Define the map ρ inductively, as follows:

- $\rho(P(\overline{u})) := (\neg u_0 = u_0)$ if the variables $\overline{u}$ include both z and one of the variables from $\overline{v}$, $P(\overline{u})$ otherwise
- $\rho(\phi \wedge \chi) := \rho(\phi) \wedge \rho(\chi)$
- $\rho(\neg\phi) := \neg\rho(\phi)$

- For each $Q \in \{\exists, \exists^{\infty}\}$, $\rho(Qx\phi) = Qx(d(\overline{v}, x) \leq k \wedge \phi) \vee Qx(d(z, x) \leq k \wedge \phi)$

Note that since $\beta(\overline{v}, z)$ is k-local, we may assume that for any subformula ψ of β such that $FV(\psi) \subseteq (\overline{v}, \beta)$, the formula ψ is also k-local. Setting

$$\Gamma := \{\phi \in \mathit{Subf}(\rho(\beta)) : \text{ a variable from } \overline{v} \text{ occurs free in } \phi\}$$

and

$$\Delta := \{\phi \in \mathit{Subf}(\rho(\beta)) : z \text{ occurs free in } \phi\}$$

we see that $\rho(\beta)$ satisfies the conditions i. and ii. It remains to show that

$$\mathfrak{M}, \overline{a}, b \models \beta \iff \mathfrak{M}, \overline{a}, b \models \rho(\beta)$$

By locality of β, $\mathfrak{M}, \overline{a}, b \models \beta$ iff $\mathfrak{M}_1 \oplus \mathfrak{M}_2, \overline{a}, b \models \beta$. A simple induction on β shows that $\mathfrak{M}_1 \oplus \mathfrak{M}_2, \overline{a}, b \models \beta$ iff $\mathfrak{M}_1 \oplus \mathfrak{M}_2, \overline{a}, b \models \rho(\beta)$. The atomic case is ensured by the fact that no formula of the form $P(\overline{a}, b, \overline{c})$ may hold in $\mathfrak{M}_1 \oplus \mathfrak{M}_2$, as $d^{\mathfrak{M}}(\overline{a}, b) > 2k + 1$.

Finally, the construction of $\rho(\beta)$ makes it clearly k-local, so $\mathfrak{M}_1 \oplus \mathfrak{M}_2, \overline{a}, b \models \rho(\beta)$ iff $\mathfrak{M}, \overline{a}, b \models \rho(\beta)$. This completes the proof of iii. □

By Claim 1, we may rewrite β_2 as $\exists z(d(\overline{v}, z) > 2k+1 \wedge \rho(\beta))$. Putting $\rho(\beta)$ into disjunctive normal form, and distributing the disjunctions first over conjunction and then over the quantifier $\exists$, we obtain:

$$\bigvee_{i \in I'} \exists z(d(\overline{v}, z) > 2k + 1 \wedge \gamma_i'(\overline{v}) \wedge \delta_i'(z))$$

(with each γ', δ' a k-local formula).

It suffices to rewrite each disjunct in the desired form, so consider any disjunt $\exists z(d(\overline{v}, z) > 2k + 1 \wedge \gamma_i'(\overline{v}) \wedge \delta_i'(z))$. Since the variable z does not occur free in $\gamma_i'(\overline{v})$, we may rewrite the disjunct as $\gamma_i'(\overline{v}) \wedge \exists z(d(\overline{v}, z) > 2k+1 \wedge \delta_i'(z))$. We know that γ_i' is k-local, so we may focus on

$$\eta(\overline{v}) := \exists z(d(\overline{v}, z) > 2k + 1 \wedge \delta_i'(z))$$

So we have reduced the problem to writing $\eta(\overline{v})$ as a Boolean combination of local formulae and basic local sentences.

For ease of notation we will start writing $\delta(z)$ rather than $\delta_i'(z)$; all that matters is that $\delta(z)$ is a k-local formula with z its sole free variable.

We now define a variety of new formulas and sentences. The idea is that we will be able to rewrite $\eta(\overline{v})$ in terms of these new formulas. For $n \in \mathbb{N}$, define $A_n(z_1, \ldots, z_n)$ to be the formula:

$$A_n := \bigwedge_{1 \le i < j \le n} d(z_i, z_j) > 2(2k+1) \wedge \bigwedge_{1 \le i \le n} \delta(z_i)$$

Then define B_n to be the sentence:

$$B_n := \exists z_1 \ldots \exists z_n A_n(z_1, \ldots, z_n)$$

and $C_n(\overline{v})$ the formula:

$$C_n := \exists z_1 \in V^{2k+1}(\overline{v}) \ldots \exists z_n \in V^{2k+1}(\overline{v}) A_n(z_1, \ldots, z_n),$$

that is, the analogue of B_n with the leading existential quantifiers relativised to $V^{2k+1}(\overline{v})$.

By Lemma 3.5 (and its corollary), each A_n is a $2k+1$-local formula. Each B_n is clearly a basic local sentence. Intuitively, B_n says that there is a set of n different things, all satisfying δ, and all separated from each other by a distance of at least $4k+2$. C_n says that there is such a set within the $2k+1$-neighbourhood of $\overline{v}$. Also, because each A_n is a $2k+1$- local formula, each C_n is a $4k+2$-local formula by Lemma 3.6.

Our strategy to rewrite the formula $\eta(\overline{v})$ will be to write a disjunction, $D = (D_1 \vee E_1) \vee \ldots \vee (D_m \vee E_m) \vee D_{m+1}$, such that each disjunct entails η, but on the other hand, η entails that one of the disjuncts of D must hold. First, we define

$$D_{m+1} := B_{m+1}$$

Claim 2. $D_{m+1} \models \eta$.

Proof. Let $\mathfrak{M}$ be a structure and σ an assignment such that $\mathfrak{M}, \sigma \models B_{m+1}$, with $\{b_1, \ldots, b_{m+1}\} \subseteq M$ a $4k+2$- scattered set witnessing this. Now suppose, for contradiction, that $\mathfrak{M}, \sigma \not\models \eta(v_1, \ldots, v_m)$. This means that every element of $\mathfrak{M}$ satisfying δ is in the $2k+1$-neighbourhood of the elements $\{\sigma(v_1), \ldots, \sigma(v_m)\}$.

By the pigeon-hole principle, there is some v such that the $2k+1$-neighbourhood of $\sigma(v)$ contains two distinct b_i, b_j. But if $d^{\mathfrak{M}}(b_i, \sigma(v)) \le 2k+1$ and $d^{\mathfrak{M}}(\sigma(v), b_j) \le 2k+1$, then $d^{\mathfrak{M}}(b_i, b_j) \le 4k+2$, and the set $\{b_1, \ldots, b_{m+1}\}$ was not $4k+2$- scattered after all. This is a contradiction. $\square$

For $i \le m$, define

$$D_i := B_i \wedge \neg B_{i+1} \wedge \neg C_i$$
$$E_i := B_i \wedge \neg B_{i+1} \wedge C_i \wedge \exists z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$$

Intuitively, D_i says that the greatest $4k+2$-scattered set of objects satisfying δ has exactly i elements, but there is no such scattered set within $V^{2k+1}(\overline{v})$. E_i says that the greatest $4k+2$-scattered set of objects satisfying δ has exactly i elements, and there *is* such a scattered set within $V^{2k+1}(\overline{v})$, *and* there is an element satisfying δ within $6k+3$ 'steps' from $\overline{v}$, but more than $2k+1$ steps from $\overline{v}$.

We must verify that the additional formula used in E_i is indeed a local formula.

Claim 3. The formula

$$\exists z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$$

is $7k+3$-local.

Proof. First, suppose $\mathfrak{M}, \sigma \models \exists z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$. Let $(\sigma(v_1), \ldots, \sigma(v_m)) = \overline{a}$. Then there is a $b \in M$ such that $2k+1 < d^{\mathfrak{M}}(\overline{a}, b) \leq 6k+3$, and $\mathfrak{M}, b \models \delta$.

Then $b \in \mathfrak{M}\restriction(\overline{a}, 7k+3)$. Further, every member of the path witnessing that $d^{\mathfrak{M}}(\overline{a}, b) \leq 6k+3$ must be within $V^{7k+3}(\overline{a})$, therefore $\mathfrak{M}\restriction(\overline{a}, 7k+3) \models d(\overline{a}, b) \leq 6k+3$. On the other hand, $\mathfrak{M}\restriction(\overline{a}, 7k+3) \models d(\overline{a}, b) > 2k+1$ must hold, because distance does not increase when moving to an extension.

Also, δ is k-local, so $\mathfrak{M}\restriction(b, k), b \models \delta$. But $[\mathfrak{M}\restriction(\overline{a}, 7k+3)]\restriction(b, k) = \mathfrak{M}\restriction(b, k)$ because $V^k(b) \subseteq V^{7k+3}(\overline{a})$. Hence $[\mathfrak{M}\restriction(\overline{a}, 7k+3)]\restriction(b, k), b \models \delta$, and by k-locality of δ, it follows that $\mathfrak{M}\restriction(\overline{a}, 7k+3), b \models \delta$. We may conclude that $\mathfrak{M}\restriction(\overline{a}, 7k+3), \overline{a} \models \exists z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$, which was to show.

Conversely, suppose $\mathfrak{M}\restriction(\overline{a}, 7k+3), \overline{a} \models \exists z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$. Then there is $b \in V^{6k+3}_{\mathfrak{M}\restriction(\overline{a}, 7k+3)}(\overline{a})$ such that $\mathfrak{M}\restriction(\overline{a}, 7k+3) \models d(\overline{a}, b) > 2k+1 \wedge \delta(b)$. δ is k-local, hence $[\mathfrak{M}\restriction(\overline{a}, 7k+3)]\restriction(b, k), b \models \delta$.

Again, $[\mathfrak{M}\restriction(\overline{a}, 7k+3)]\restriction(b, k) = \mathfrak{M}\restriction(b, k)$, so we can infer that $\mathfrak{M}\restriction(b, k), b \models \delta$. Also, we will have $b \in V^{6k+3}_{\mathfrak{M}}(\overline{a})$ because distance does not increase when moving to extensions.

It only remains to show that $d^{\mathfrak{M}}(\overline{a}, b) > 2k+1$. Suppose for contradiction that this is not the case. Then there is a path $a_i = c_0, c_1, \ldots, c_{2k+1} = b$ linking $\overline{a}$ to b. Every member of this path must be in $V^{2k+1}(\overline{a})$ and hence in $\mathfrak{M}\restriction(\overline{a}, 7k+3)$, so we would also have $d^{\mathfrak{M}\restriction(\overline{a}, 7k+3)}(\overline{a}, b) \leq 2k+1$. But this is a contradiction.

We may conclude that $\mathfrak{M}, \overline{a} \models \exists z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$, as required. $\square$

We now return to showing that each disjunt of D entails the formula η.

Claim 4. For each $i \leq m$, $D_i \models \eta$ and $E_i \models \eta$.

Proof. Suppose $\mathfrak{M}, \sigma$ satisfies D_i, and let $\{b_1, \ldots, b_i\} \subseteq M$ be a $4k+2$-scattered set with each b_j satisfying δ. Let $\overline{a} = \{\sigma(v_1), \ldots, \sigma(v_m)\}$. Since $\mathfrak{M}, \sigma \not\models C_i$, it cannot be that $\overline{b}$ is contained within $V^{2k+1}(\overline{a})$, so there is some b_j with $\mathfrak{M}, \sigma \models \delta(b_j)$, but $d^{\mathfrak{M}}(\overline{a}, b_j) > 2k+1$. Hence $\mathfrak{M}, \sigma \models \eta$.

Moving to E_i: E_i entails its conjunct, $\exists z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$, which in turn clearly entails η. So: $E_i \models \eta$.

This proves Claim 4.

□

So, we have shown that every disjunct of D entails η. Once we show that η entails that one of the disjuncts must hold, we are done.

Claim 5. $\eta \models D$.

Proof. Suppose $\mathfrak{M}, \sigma \models \eta(\overline{v})$, and let $(\sigma(v_1), \ldots, \sigma(v_m)) = \overline{a}$. Then there is some $b \in M$, at a distance strictly greater than $2k+1$ from $\overline{a}$, which satisfies δ.

Suppose D_{m+1} does not hold in $\mathfrak{M}$. Then the size of the largest $4k+2$-scattered set of elements of M satisfying δ is exactly i, for some natural number $i \leq m$. We claim that in this case, $\mathfrak{M}, \sigma \models (D_i \vee E_i)$.

Suppose $\mathfrak{M}, \sigma \not\models D_i$. We know that $\mathfrak{M}, \sigma \models B_i \wedge \neg B_{i+1}$, so it must be that $\mathfrak{M}, \sigma \models C_i$. Let $\overline{b} = \{b_1, \ldots, b_i\}$ be a $4k+2$-scattered set $\subseteq V^{2k+1}(\overline{a})$, with each b_j satisfying δ. Certainly $b \notin \overline{b}$, as $d(\overline{a}, b) > 2k+1$. However, the set $\overline{b} \cup \{b\}$ is not $4k+2$-scattered, as i is the maximum cardinality of such a set. So there is some b_j such that $d^{\mathfrak{M}}(b_j, b) \leq 4k+2$. Since $d^{\mathfrak{M}}(\overline{a}, b_j) \leq 2k+1$, we have $d^{\mathfrak{M}}(\overline{a}, b) \leq 6k+3$.

But then $\mathfrak{M}, \sigma \models \exists z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$, with b the witness. Hence $\mathfrak{M}, \sigma \models E_i$. □

Claim 5 completes the proof of part (a) of the Lemma. □

We now tackle part (b). Assume $\beta \in FOL^\infty$ is a k-local formula, with $FV(\beta) = (\overline{v}, z)$. Our task is to rewrite $\exists^\infty z \beta$ as a Boolean combination of local formulae and basic local infinity sentences.

The argument used in part (a), up until the problem was reduced to rewriting the formula $\eta(\overline{v})$, still goes through for $\exists^\infty$. We only need to use Proposition 2.15, which says that $\exists^\infty$ distributes over disjunction, at various stages. In the $\exists^\infty$ case, we reach a formula of the form

$$\eta^\infty(\overline{v}) := \exists^\infty z(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$$

with δ a k-local formula.

Now define $A_n^\infty := A_n$. However, define B_n^∞ as:

$$B_n^\infty := \exists^\infty z_1 \ldots \exists^\infty z_n A_n^\infty$$

and C_n^∞ as:

$$C_n^\infty := \exists^\infty z_1 \in V^{2k+1}(\overline{v}) \ldots \exists^\infty z_n \in V^{2k+1}(\overline{v}) A_n^\infty$$

i.e., the analogue of B_n^∞ with the leading string of infinity quantifiers relativised to $V^{2k+1}(\overline{v})$. Clearly each B_n^∞ is a basic local infinity sentence, and the same argument as in the proof of item (a) shows that each C_n^∞ is $4k+2$-local.

We use the same strategy as for item (a), of rewriting η^∞ as a disjunction, $D^\infty := (D_1^\infty \vee E_1^\infty) \vee \ldots \vee (D_m^\infty \vee E_m^\infty) \vee D_{m+1}^\infty$. We first show that each disjunct entails η^∞. Define

$$D_{m+1}^\infty := B_{m+1}^\infty$$

Claim 6. $D_{m+1}^\infty \models \eta^\infty$.

Proof. Suppose $\mathfrak{M}, \sigma \models B_{m+1}^\infty$. Applying Observation 2.14, let $f : \mathbb{F}((\omega)^{m+1}) \to M$ be a map such that for any path $(t_1, \ldots, t_{m+1})$ through $\mathbb{F}((\omega)^{m+1})$, $\mathfrak{M}, \sigma \models A_{m+1}^\infty(f(t_1), \ldots, f(t_{m+1}))$. We claim that in fact, there exists an infinite family $\mathcal{F}$ of ordered $m+1$- tuples $\overline{b} \subseteq M$ such that:

- $\mathcal{F}$ is pairwise disjoint, i.e. for each $\overline{b}, \overline{b}' \in \mathcal{F}$, for all $b \in \overline{b}, b' \in \overline{b}'$, $b \neq b'$
- For each $\overline{b} \in \mathcal{F}$, $\mathfrak{M} \models A_{m+1}(\overline{b})$

To see this, build $\mathcal{F}$ in stages. Let $\mathcal{F}_0 = \emptyset$. At stage $i+1$, choose a path $\pi^{i+1} = \langle p_1^{i+1}, \ldots, p_{m+1}^{i+1} \rangle$ through $\mathbb{F}(q_1, \ldots, q_{m+1})$, taking care that each object in π^{i+1} is mapped by f to a 'fresh' element of M that hasn't been used by any previous $\mathcal{F}_j$. (Because f is injective on $\sim_{\mathbb{F}}$ equivalence classes, we have infinitely many elements of M to choose from at every step, while only finitely many elements of M have been used in some previous $\mathcal{F}_j$.) Let $\mathcal{F}_{i+1} := \mathcal{F}_i \cup \{(f(p_1^{i+1}), \ldots, f(p_{m+1}^{i+1}))\}$. Take $\mathcal{F} := \bigcup_{i \in \omega} \mathcal{F}_i$, then $\mathcal{F}$ satisfies the desired conditions. Enumerate $\mathcal{F}$ as $t_0, t_1, \ldots$ (t for 'tuple').

Let σ be an assignment and let $\overline{a}$ denote $\sigma[\overline{v}]$. Then for each tuple $t_i \in \mathcal{F}$, by the same pigeon-hole argument as used in part (a), there exists some $b_i \in t_i$ such that $b_i \notin V^{2k+1}(\overline{a})$. Because $\mathcal{F}$ is pairwise disjoint, we know that the set $\{b_i : i \in \mathbb{N}\}$ is infinite. Hence, it witnesses $\exists^\infty z(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$ holding in $\mathfrak{M}, \sigma$, and therefore $\mathfrak{M}, \sigma \models \eta^\infty$. So D_{m+1}^∞ entails η^∞. $\square$

For $i \leq m$, define

$$D_i^\infty := B_i^\infty \wedge \neg B_{i+1}^\infty \wedge \neg C_i^\infty$$
$$E_i^\infty := B_i^\infty \wedge \neg B_{i+1}^\infty \wedge C_i^\infty \wedge \exists^\infty z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$$

Claim 7. The formula

$$\exists^\infty z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$$

is $7k+3$-local.

Proof. The proof is similar to that of the analogous Claim 3 within the proof of part (a). We first suppose $\mathfrak{M}, \overline{a} \models \exists^\infty z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$. Then there is an infinite set $\mathbf{b}$ such that for each $b \in \mathbf{b}$, we have $2k+1 < d^{\mathfrak{M}}(\overline{a}, b) \leq 6k+3$ and $\mathfrak{M} \models \delta(b)$. For each $b \in \mathbf{b}$, we may apply exactly the same reasoning as in the earlier proof, to conclude that $b \in \mathfrak{M}{\restriction}(\overline{a}, 7k+3)$ and that $\mathfrak{M}{\restriction}(\overline{a}, 7k+3) \models 2k+1 < d(\overline{a}, b) \leq 6k+3 \wedge \delta(b)$.

For the other direction, we suppose $\mathfrak{M}{\restriction}(\overline{a}, 7k+3), \overline{a} \models \exists^\infty z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$. Then there is an infinite set $\mathbf{b}$ of witnesses. Reasoning as before for each $b \in \mathbf{b}$, we can conclude that $\mathfrak{M} \models 2k+1 < d(\overline{a}, b) \leq 6k+3 \wedge \delta(b)$, for each $b \in \mathbf{b}$. This suffices to prove the claim. □

We now return to showing that each disjunct of D^∞ entails the formula η^∞.

Claim 8. For each $i \leq m$, we have $D_i^\infty \models \eta^\infty$ and $E_i^\infty \models \eta^\infty$.

Proof. Suppose $\mathfrak{M}, \sigma \models D_i^\infty$. Then $\mathfrak{M}, \sigma \models B_i^\infty$, and so

$$\mathfrak{M}, \sigma \models \exists^\infty z_1 \ldots \exists^\infty z_i [A_i \wedge (\bigwedge_{1 \leq j \leq i} (d(\overline{v}, z) \leq 2k+1) \vee \neg \bigwedge_{1 \leq j \leq i} (d(\overline{v}, z) \leq 2k+1))]$$

since this formula is logically equivalent to B_i^∞. Distributing $\wedge$ over $\vee$ yields

$$\begin{aligned} \mathfrak{M}, \sigma \models \exists^\infty z_1 \ldots \exists^\infty z_i [(A_i \wedge \bigwedge_{1 \leq j \leq i} (d(\overline{v}, z) \leq 2k+1)) & \\ \vee (A_i \wedge \neg \bigwedge_{1 \leq j \leq i} (d(\overline{v}, z_j) \leq 2k+1))] & \end{aligned}$$

By Proposition 2.15, this is equivalent to

$$\begin{aligned} \mathfrak{M}, \sigma \models \exists^\infty z_1 \in V^{2k+1}(\overline{v}) \ldots \exists^\infty z_i \in V^{2k+1}(\overline{v})[A_i] & \\ \vee \exists^\infty z_1 \ldots \exists^\infty z_i [A_i \wedge \neg \bigwedge_{1 \leq j \leq i} (d(\overline{v}, z_j) \leq 2k+1)] & \end{aligned}$$

But the first of these disjuncts is precisely the formula C_i^∞, and since $\mathfrak{M}, \sigma \models D_i^\infty$ we have $\mathfrak{M}, \sigma \models \neg C_i^\infty$. Therefore $\mathfrak{M}, \sigma$ must satisfy the second disjunct, which implies

$$\mathfrak{M}, \sigma \models \exists^\infty z_1 \ldots \exists^\infty z_i [\bigwedge_{1 \leq j \leq i} (\delta(z_j)) \wedge \neg \bigwedge_{1 \leq j \leq i} (d(\overline{v}, z_j) \leq 2k+1)]$$

by the definition of A_i and conjunction elimination.

Using Observation 2.14 again, there exists a map $f : \mathbb{F}((\omega)^i) \to M$ such that for any path $(t_1, \ldots, t_i)$ through $\mathbb{F}((\omega)^i)$,

$$\mathfrak{M}, \sigma \models \neg \bigwedge_{j \leq i} (d(\overline{v}, f(t_j)) \leq 2k+1) \wedge \delta(f(t_j)))$$

As in the proof of Claim 7, we may build an infinite, pairwise disjoint family $\mathcal{F}$ of i- tuples $\overline{b} \subseteq M$, such that for each $\overline{b} \in \mathcal{F}$

$$\mathfrak{M}, \sigma \models \neg \bigwedge_{1 \leq j \leq i} (d(\overline{v}, b_j) \leq 2k+1 \wedge \delta(b_j))$$

Enumerate $\mathcal{F}$ as $t_0, t_1, \ldots$ Then for each $t_i \in \mathcal{F}$, there must be some $b_i \in t_i$ such that $\mathfrak{M}, \sigma \models d(\overline{v}, b_i) > 2k+1$. We may then take the infinite set $\{b_i : i \in \mathbb{N}\}$ as a witness for $\exists^\infty z(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$ holding in $\mathfrak{M}, \sigma$. So D_i^∞ entails η^∞.

The implication from E_i^∞ to η^∞ goes by the same trivial reasoning as in the proof of (a). □

It only remains to show:

Claim 9. η^∞ entails D^∞.

Proof. Suppose $\mathfrak{M}, \sigma \models \eta^\infty(\overline{v})$. Let $(\sigma(v_1), \ldots, \sigma(v_{m+1})) = \overline{a}$. Then there is an infinite set, $\mathbf{b} \subseteq M$, such that for each $b \in \mathbf{b}$, $\mathfrak{M} \models d(\overline{a}, b) > 2k+1 \wedge \delta(b)$.

The sentence B_1^∞ simply asserts the existence of infinitely many b's satisfying δ, and so must certainly hold in $\mathfrak{M}, \sigma$. Now suppose $D_{m+1}^\infty = B_{m+1}^\infty$ does not hold in $\mathfrak{M}, \sigma$. Then there must be some greatest i, with $i \leq m$, such that $\mathfrak{M}, \sigma \models B_i^\infty$.

We have $\mathfrak{M}, \sigma \models B_i^\infty$ and $\mathfrak{M}, \sigma \not\models B_{i+1}^\infty$. Suppose $\mathfrak{M}, \sigma \not\models D_i^\infty$. We will show that $\mathfrak{M}, \sigma \models E_i^\infty$. If $\mathfrak{M}, \sigma \not\models D_i$, this can only be because $\mathfrak{M}, \sigma \models C_i^\infty$, i.e. $\mathfrak{M}, \sigma \models \exists^\infty z_1 \in V^{2k+1}(\overline{v}) \ldots \exists^\infty z_i \in V^{2k+1}(\overline{v}) A_i^\infty$. Using Observation 2.14, let $f : \mathbb{F}((\omega)^i) \to M$ be a map witnessing this. We claim that

$$\mathfrak{M}, \sigma \models \exists^\infty z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$$

To see this, suppose for contradiction that the set $\mathbf{b}' := \{b \in \mathbf{b} : d^{\mathfrak{M}}(\overline{a}, b) \leq 6k+3\}$ is finite. Then $\mathbf{b} \setminus \mathbf{b}'$ is infinite. We now define a map $g : \mathbb{F}((\omega)^{i+1}) \to M$. Identifying $\mathbb{F}((\omega)^i)$ with the restriction of $\mathbb{F}((\omega)^{i+1})$ to its nodes of depth $\leq i$, we may define $g(t) = f(t)$ for any t that is not a leaf of $\mathbb{F}((\omega)^{i+1})$.

To define g on the leaves of $\mathbb{F}((\omega)^{i+1})$, consider any leaf t of $\mathbb{F}((\omega)^i)$. Define $g \restriction R^{\mathbb{F}}[t]$ to be some arbitrary injection from $R^{\mathbb{F}}[t]$ into $\mathbf{b} \setminus \mathbf{b}'$. Doing this for every leaf of $\mathbb{F}((\omega)^{i+1})$ ensures that g is fully defined on $\mathbb{F}((\omega)^{i+1})$. Further, if $s \in \mathbb{F}(q_1, \ldots, q_i)$ and t is a leaf of $\mathbb{F}(q_1, \ldots, q_{i+1})$, then $d^{\mathfrak{M}}(g(s), g(t)) > 4k+2$ (otherwise $g(t)$ would be in $V^{6k+3}(\overline{a})$ and hence in $\mathbf{b}'$). Therefore g satisfies:

- For any $\sim_{\mathbb{F}}$-equivalence class $[t] \subseteq \mathbb{F}((\omega)^{i+1})$, $g\restriction[t]$ is injective
- For any path $t_1, \ldots, t_{i+1}$ through $\mathbb{F}((\omega)^{i+1})$, we have $\mathfrak{M}, \sigma \models A^{\infty}_{i+1}(g(t_1), \ldots, g(t_{i+1}))$

Then by Observation 2.14, $\mathfrak{M}, \sigma \models B^{\infty}_{i+1}$. But this is a contradiction.

We can conclude that $\mathbf{b}'$ is infinite. Then $\mathbf{b}'$ is a witness for

$$\mathfrak{M}, \sigma \models \exists^{\infty} z \in V^{6k+3}(\overline{v})(d(\overline{v}, z) > 2k+1 \wedge \delta(z))$$

It follows that $\mathfrak{M}, \sigma \models E^{\infty}_i$.

We have shown that if $\mathfrak{M}, \sigma \models \eta^{\infty}$, then if $\mathfrak{M}, \sigma \not\models D^{\infty}_{m+1}$ and $\mathfrak{M}, \sigma \not\models D^{\infty}_i$ (where i is maximal such that $\mathfrak{M}, \sigma \models B^{\infty}_i$), then $\mathfrak{M}, \sigma \models E^{\infty}_i$. We may conclude that η^{∞} entails D^{∞}, as desired. $\square$

This completes the proof of item (b), and thus of Lemma 3.9. $\square$

3.4.2 Proof of Main Theorem

Recall the statement to be shown: Every $\phi \in FOL^{\infty}$ is equivalent to a Boolean combination of local formulae, basic local sentences and basic local infinity sentences. We argue by induction on the construction of ϕ.

Base case: ϕ is atomic. The statement clearly holds for atomic formulas, which are 0-local.

The induction steps for the Boolean connectives are also trivial.

Induction step: $\phi = \exists z \psi$. Inductively assume ψ is a Boolean combination of local formulae, basic local sentences and basic local infinity sentences. We may rewrite ψ in disjunctive normal form: $\psi = \bigvee_i \chi_i$. Then ϕ is equivalent to $\bigvee_i \exists z \chi_i$, and we may focus on $\exists z \chi_i$.

Each χ_i is of the form $\bigwedge_{j \in J} \theta_j$, since we put ψ in DNF. Further, by IH, each θ_j is either a local formula, a basic local sentence or a basic local infinity sentence. Let $J' := \{j \in J : \theta_j \text{ is a basic local sentence or basic local infinity sentence}\}$. Then the variable z does not occur free in $\bigwedge_{j \in J'} \theta_j$, so we may rewrite $\exists z \bigwedge_{j \in J} \theta_j$ as $\bigwedge_{j \in J'} \theta_j \wedge \exists z \bigwedge_{j \in J \setminus J'} \theta_j$, and focus on $\exists z \bigwedge_{j \in J \setminus J'} \theta_j$.

We know that for each $j \in J \setminus J'$, θ_j is a local formula. Local formulas are closed under Boolean operations, so our target formula is of the form $\exists z \beta$, where β is a local formula. We can rewrite this formula in the desired form by part (a) of the lemma.

Induction step: $\phi = \exists^\infty z\psi$. The argument is the same as in the $\exists$ case, using the fact that $\exists^\infty$ distributes over disjunctions, and where $z \notin FV(\theta)$, $\exists^\infty z(\theta \wedge \chi)$ is equivalent to $\theta \wedge \exists^\infty z\chi$. We obtain a formula of the form $\exists^\infty z\beta$, where β is local. We may then apply part (b) of the lemma.

This completes the induction. So the theorem is proved. $\square$

3.5 Reflections

We have proved the Gaifman theorem for *FOL*$^\infty$ in this thesis for the specific purpose of proving characterisation theorems for bisimulation invariant fragments of *FOL*$^\infty$. For that reason, we will not make a serious attempt to draw out further consequences of the theorem for *FOL*$^\infty$. However, the theorem does seem to be of independent interest, and so we will take a moment here to briefly discuss it and to make some speculative remarks.

We begin with a comparison of the first-order version of the theorem with its *FOL*$^\infty$ analogue. In [9], Gaifman provided two distinct proofs of Theorem 3.7: the proof which we followed, and a much shorter and easier proof. The shorter proof makes use of the fact that, within the first-order setting, every structure $\mathfrak{M}$ has an ω-saturated elementary extension. However, this no longer holds in the *FOL*$^\infty$ setting, because there are finitely satisfiable *FOL*$^\infty$ types which are unsatisfiable, such as $p(x) := \{\exists^{\geq 1}y(Rxy), \exists^{\geq 2}y(Rxy), \ldots, \neg\exists^\infty y(Rxy)\}$. Let $\mathfrak{M}$ be any structure which finitely realises p; then an ω-saturated *FOL*$^\infty$-elementary extension of $\mathfrak{M}$ would have to realise p, which is impossible. The fact that *FOL*$^\infty$ lacks the property that every structure has an ω-saturated elementary extension seems to make it all the more striking and interesting that *FOL*$^\infty$ nevertheless satisfies Theorem 3.8.

A further point is that it was not obvious, *a priori*, what exactly the analogue of Gaifman's theorem should be for *FOL*$^\infty$. More precisely, it is clear that the role played by 'basic local sentences' in Theorem 3.7 must be played by a wider class of sentences in an *FOL*$^\infty$ analogue of this theorem; but what should this wider class of sentences be? It might seem like we would need to allow sentences such as

$$\exists x_1 \exists^\infty x_2 \exists x_3 \exists^\infty x_4 \big(\bigwedge_{1 \leq i < j \leq 4} (d(x_i, x_j) > k \wedge \psi(x_i))\big)$$

where ψ is a k-local formula; i.e., we might need to allow basic local sentences introduced by arbitrary strings of quantifiers, perhaps with alternating existential and infinity quantifiers. But we have proven that this is not the case. Theorem 3.8 states that we can express every formula of *FOL*$^\infty$ using local formulas, basic local sentences (which are introduced by a string of existential quantifiers) and basic local infinity sentences (which are introduced by a string of infinity quantifiers). Of course, we have *not* shown that we can avoid alternating strings of existential quantifiers and infinity quantifiers in FOL^{∞} altogether. Rather, the point is that every such alternating string can be pushed within one of the 'local' components of the FOL^{∞} formula. But this still feels like a surprisingly strong result.

We now turn to the implications Theorem 3.8 might have for the expressive power of FOL^{∞}. Gaifman originally proved Theorem 3.7 in order to establish some limitations on the expressive power of first-order logic over finite structures. Of course, the Compactness theorem fails in this setting, and so negative results on the expressivity of first-order logic which are classically shown by compactness require a different proof. For example, it is easy to show by compactness that first-order logic cannot express the undirected-graph-property of being connected; Gaifman used his theorem to show that this property is also not first-order definable within the class of finite structures.

As we have already mentioned, FOL^{∞} is another setting in which the Compactness theorem fails. So, it is natural to try using Theorem 3.8 to show some limitations on the expressive power of FOL^{∞}. It seems plausible that Theorem 3.8 easily yields a proof that FOL^{∞} is also unable to define connexity, analogous to Gaifman's proof in [9]. This, in turn, would yield a simple proof that FOL^{∞} is unable to define the transitive closure of a relation R (otherwise we could define connexity by: for any two points x and y, x stands in the transitive closure of R to y). The fact that FOL^{∞} is so expressively powerful in other ways, as discussed in Chapter 1, would make negative results of this kind all the more significant.

Finally, we raise the question of whether analogous results could be obtained for abstract logics other than FOL^{∞}. First, we note that, for any given logic, an issue will arise of what class of sentences should count as the analogue of 'basic local (infinity) sentences'. The case where we replace the infinity quantifier with another generalised unary quantifier Q is easiest; we can consider sentences of the form

$$Qx_1 \dots Qx_n(\bigwedge_{1 \leq i < j \leq n} (d(x_i, x_j) > k \wedge \psi(x_i)))$$

where ψ is a k-local formula, and ask whether every formula can be expressed by a Boolean combination of local formulas, basic local sentences, and sentences of the above form. If Q is another cardinality quantifier, i.e. $Q = Q_\alpha$ for some $\alpha > 0$, there is no obvious reason why the proof given in this chapter would not go through in just the same way. If we allow binary or other non-unary generalised quantifiers, it becomes less obvious which class of sentences we should consider. If we allow infinitely long strings of quantifiers, such as in the logic $\mathcal{L}_{\omega\omega_1}$, then we may need to include sentences introduced by infinitely long strings of quantifiers within our class of analogues of basic local sentences. Despite these complications, a general question suggests itself: under what conditions does an abstract logic have the 'Gaifman property'?

To the best of my knowledge, Theorem 3.8 is new. Keisler and Lotfallah [12] establish a local normal form result for the logic $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^{\omega}$, the logic with arbitrary conjunctions/disjunctions and arbitrary unary quantifiers, but restricted to formulas with finite quantifier rank. However, their results are restricted to finite structures, and so have no direct bearing on the result shown here. (The authors note that their arguments would also go through for countable, locally finite models, but this would still not be sufficiently general to imply the theorem shown in this chapter.)

Chapter 4

Characterisation Results via Upgrading

This chapter completes the proof that each of the logics T^∞, G^∞, GT^∞ is characterised by invariance under the appropriate notion of bisimulation. The chapter will begin by explaining, in very general terms, the idea behind the 'upgrading' strategy used to complete the proof. We will then provide a series of upgrading results, which suffice to prove the desired characterisation theorems.

4.1 The Upgrading Strategy

We begin with a high-level explanation of the proof strategy we will employ in this chapter. Say we are given an arbitrary formula, ϕ, of FOL^∞, that is invariant under some bisimilarity relation $\rightleftharpoons$, and we want to prove that ϕ can be expressed within a syntactically given fragment of FOL^∞ that corresponds to $\rightleftharpoons$. The relation $\rightleftharpoons$ will have finite approximations $\rightleftharpoons^k$, for each $k \in \mathbb{N}$, which have finite index and coincide with logical equivalence in the corresponding modal fragment, restricted to modal depth k. It therefore suffices to show that there is some k such that ϕ is invariant under $\rightleftharpoons^k$.

Now, the Gaifman theorem proved in Chapter 3 shows that our formula ϕ is logically equivalent to some ϕ', the 'Gaifman form' of ϕ, that is a Boolean combination of local formulas and basic local sentences. Since only finitely many local formulas may occur as subformulas of ϕ', there must be some l- the 'maximum' locality rank- such that *every* one of these local formulas is l-local. Also, there must be some maximum quantifier depth q of any local formula occurring in ϕ', and a maximum quantifier depth n of any basic local sentence occurring in ϕ'. (In the case where ϕ is invariant under $\simeq_T$, we can actually say a bit more, but we leave this aside for now.) The upshot is that the information carried by ϕ must be 'bounded', in an intuitive sense, by the finitary parameters l, q, n, and we can therefore define a suitable finitary equivalence relation $\approx_n^{l,q}$, parametric in l, q, n, such that ϕ is invariant under $\approx_n^{l,q}$. (We call $\approx_n^{l,q}$ the 'target' equivalence relation.)

The upgrading idea is to show that ϕ is invariant under $\rightleftharpoons^k$ by showing that for any l, q, n, there is some k such that $\rightleftharpoons^k$ can be 'upgraded' to $\approx_n^{l,q}$ modulo $\rightleftharpoons$. This means that for any pair of Kripke structures $(\mathcal{M}_0, s_0)$, $(\mathcal{M}_1, s_1)$, such that $\mathcal{M}_0, s_0 \rightleftharpoons^k \mathcal{M}_1, s_1$, there are companion structures $\mathcal{M}'_0, s'_0 \rightleftharpoons \mathcal{M}_0, s_0$ and $\mathcal{M}'_1, s'_1 \rightleftharpoons \mathcal{M}_1, s_1$ which satisfy the stronger, target equivalence relation: $\mathcal{M}'_0, s'_0 \approx_n^{l,q} \mathcal{M}'_1, s'_1$. Then the truth value of ϕ must be preserved from $\mathcal{M}_0, s_0$ through $\mathcal{M}'_0, s'_0$ and $\mathcal{M}'_1, s'_1$ to $\mathcal{M}_1, s_1$. In a diagram:

$$
\begin{tikzcd}
\mathcal{M}_0, s_0 \arrow[r, leftrightarrow, "\rightleftharpoons^k"] \arrow[d, "\rightleftharpoons"'] & \mathcal{M}_1, s_1 \arrow[d, "\rightleftharpoons"] \\
\mathcal{M}'_0, s'_0 \arrow[r, leftrightarrow, "\approx_n^{l,q}"'] & \mathcal{M}'_1, s'_1
\end{tikzcd}
$$

Figure 4.1: Upgrading schema

Since this holds for any $\mathcal{M}_0, s_0 \rightleftharpoons^k \mathcal{M}_1, s_1$, we may conclude that ϕ is $\rightleftharpoons^k$-invariant, which in turn implies that ϕ is expressible in the desired modal fragment of FOL^∞.

See also the discussion in Otto [15], p.186. This chapter is closely based on the proofs given by Otto of his Theorem 23, Corollary 25, and Proposition 41 in the paper just cited.

The main work involved in this chapter will be in proving the required upgrading results. We will employ a variety of model construction methods to turn two original structures, $\mathcal{M}_0$ and $\mathcal{M}_1$, into bisimilar companion structures which, in each case, satisfy the desired 'target' equivalence.

4.2 Upgrading from $\simeq_T^k$ to $\simeq_T^{(k)}$

We treat the $\simeq_T$ case separately, because our target equivalence relation for this upgrading argument will be different from in the $\simeq_G$ and $\simeq_{GT}$ cases. In fact, our target equivalence will be given by the following definition.

Definition 4.1 ($\simeq_T^{(k)}$)**.** Let $k \in \mathbb{N}$, and let $\mathcal{M}_0, s_0, \mathcal{M}_1, s_1$ be pointed Kripke structures. We write $\mathcal{M}_0, s_0 \simeq_T^{(k)} \mathcal{M}_1, s_1$ iff $\mathcal{M}_0{\restriction}(s_0, k), s_0 \simeq_T \mathcal{M}_1{\restriction}(s_1, k), s_1$ (recall this notation from Definition 3.2).

So, two pointed structures are $\simeq_T^{(k)}$-bisimilar iff their restrictions to the k-neighbourhoods of the focus points are fully two-way infinity bisimilar. This is not to be confused with the finite-step bisimilarity relation $\simeq_T^k$, which is strictly weaker than $\simeq_T^{(k)}$. The single reflexive point is two-way k-bisimilar to $\mathbb{Z}, 0$ under the successor relation, but the k-neighbourhood of the single reflexive point (i.e., the single reflexive point) has an infinite path, whereas the restriction of $\mathbb{Z}$ to the k-neighbourhood of 0 is guaranteed not to have an infinite path, so they are not fully two-way bisimilar.

Proposition 4.2 (Two-way upgrading)**.** *Modulo $\simeq_T$, $\simeq_T^k$ can be upgraded to $\simeq_T^{(k)}$. In other words: let $(\mathcal{M}_0, s_0), (\mathcal{M}_1, s_1)$ be pointed Kripke structures, with $\mathcal{M}_0, s_0 \simeq_T^k \mathcal{M}_1, s_1$. Then there are companion structures $\mathcal{M}_0', s_0' \simeq_T \mathcal{M}_0, s_0$, $\mathcal{M}_1', s_1' \simeq_T \mathcal{M}_1, s_1$ such that $\mathcal{M}_0', s_0' \simeq_T^{(k)} \mathcal{M}_1', s_1'$.*

To build our companion structures we will make use of the classic model construction method of *unravelling*, with some extra refinement.

Definition 4.3 (Two-way path)**.** Let $\mathcal{M}$ be a structure. A *path* through $\mathcal{M}$ is an ordered tuple $\langle s_0, \ldots, s_n \rangle$, with each $s_i \in S$ and each $s_i R^{\mathcal{M}} s_{i+1}$.

A *two-way path* through $\mathcal{M}$ is an ordered tuple $\langle s_0, D_0, s_1, \ldots, D_{n-1}, s_n \rangle$, with each $s_i \in S$ and each D_i a *direction indicator*, either 'R' or '$\check{R}$', such that if $D_i =$ 'R' then $\pi_i R^{\mathcal{M}} \pi_{i+1}$, and if $D_i =$ '$\check{R}$' then $\pi_{i+1} R^{\mathcal{M}} \pi_i$. From here on we will be sloppy and drop quotation marks for the direction indicators, writing simply R or $\check{R}$.

The *length* of a path π, regular or two-way, is the number of occurrences of objects in π, minus 1. So for example, the length of the two-way path

$$\langle a_0, R, a_1, R, a_1, \check{R}, a_0 \rangle$$

is 3.

If π is a path (regular or two-way) through a structure $\mathcal{M}$, we write $last(\pi)$ to denote the final object to occur in π.

Definition 4.4 (Two-way unravelling)**.** Let $\mathcal{M}, s$ be a pointed structure. The *two-way unravelling* of $\mathcal{M}$ from s, notation: $\mathcal{U}_T(\mathcal{M}, s)$, is the structure whose universe $U_T(\mathcal{M}, s)$ is the set of all two-way paths through $\mathcal{M}$ beginning from s, and which interprets predicate symbols by:

- $\pi \in P^{\mathcal{U}_T(\mathcal{M},s)}$ iff $last(\pi) \in P^{\mathcal{M}}$, for unary P
- $\pi R^{\mathcal{U}_T(\mathcal{M},s)} \pi'$ iff $\pi' = \pi ^\frown \langle R, last(\pi') \rangle$ or $\pi = \pi' ^\frown \langle \check{R}, last(\pi) \rangle$

The *two-way unravelling of depth* n of $\mathcal{M}, s$, written $\mathcal{U}_T^n(\mathcal{M}, s)$, is defined as above, except that $U_T^n(\mathcal{M}, s)$ contains only those paths starting from s of length $\leq n$.

For any structure $\mathcal{M}, s$, if $\mathcal{U}$ is the (full or depth-n) two-way unravelling of $\mathcal{M}, s$, then for any $\pi \in U$ and any $s' \in S$, we say that π is a *representative* of s' iff $s' = last(\pi)$.

Definition 4.5 (Two-way tree)**.** A *two-way tree* is a binary relational structure $\mathbb{T} = (T, R)$ such that

- The Gaifman graph of $\mathbb{T}$ is connected and acyclic
- $R \cap \check{R} = \emptyset$

Note that there is a unique non-repetitive path between *any* two points $t, t' \in T$ in the Gaifman graph of $\mathbb{T}$. Consequently, two-way trees do not have a uniquely determined root. However, we may still use common tree notions *relativised* to thinking of a particular element as the root.

In particular, we say that $t \in T$ is a *leaf* in $\mathbb{T}$ *relative to root* r, iff there is no $t' \in T$ such that the unique path from t' to r passes through t; we may also say 't is a leaf in $(\mathbb{T}, r)$', or simply 't is a leaf in $\mathbb{T}$' if the choice of r is clear from context. Similarly, we say that the *height* of t in $\mathbb{T}$ *relative to root* r, or '$h(t)$ in $(\mathbb{T}, r)$', is simply $d^{\mathbb{T}}(r, t)$. We say that the height of $\mathbb{T}$ relative to r is the supremum of $\{h(t) : t \in T\}$ relative to $\mathbb{T}, r$. We say that t' is a *descendant* of t in $(\mathbb{T}, r)$ iff the unique path from r to t' passes through t, and t' is an *ancestor* of t in $(\mathbb{T}, r)$ iff t is a descendant of t' in $(\mathbb{T}, r)$. t' is a *k-step descendant of* t in $(\mathbb{T}, r)$ iff t' is a descendant of t and $d^{\mathbb{T}}(t, t') = k$; similarly for '$k$-step ancestor'. If $\mathbb{T}$ and r are clear from context, then we will write $\downarrow(t)$ to denote the set of ancestors of t in $(\mathbb{T}, r)$, and $\uparrow(t)$ for the set of descendants of t in $(\mathbb{T}, r)$.

Note that the two-way unravelling or n-depth two-way unravelling of any pointed structure will always be a two-way tree.

Proof of Proposition 4.2

We are given two pointed structures, $\mathcal{M}_0, s_0$ and $\mathcal{M}_1, s_1$, with $\mathcal{M}_0, s_0 \simeq_T^k \mathcal{M}_1, s_1$. We have to define $\mathcal{M}_0'$ and $\mathcal{M}_1'$.

Define $\mathcal{M}_0'$ by first taking $\mathcal{U}_T^{k+1}(\mathcal{M}_0, s_0)$. For each $\pi \in U_T^{k+1}(\mathcal{M}_0, s_0)$ that is a leaf in $\mathcal{U}_T^{k+1}(\mathcal{M}_0, s_0)$ relative to $\langle s_0 \rangle$, let $\mathcal{M}_0^\pi$ be an isomorphic copy of $\mathcal{M}_0$, disjoint from $\mathcal{M}_0$ and also from each $\mathcal{M}_0^\rho$ for all $\rho \neq \pi$. Construct $\mathcal{M}_0'$ by 'gluing' each $\mathcal{M}_0^\pi$ onto the structure $\mathcal{U}_T^{k+1}(\mathcal{M}_0, s_0)$ at the point π, identifying each π with the object $last(\pi)$ in $\mathcal{M}_0^\pi$. Finally, define $s_0' := \langle s_0 \rangle$.

Define $\mathcal{M}_1'$ by exactly the same process, first taking the $k+1$-depth two-way unravelling from s_1 and then gluing isomorphic copies of $\mathcal{M}_1$ to the leaves. Define $s_1' := \langle s_1 \rangle$.

Claim 1. $\mathcal{M}_0', s_0' \simeq_T \mathcal{M}_0, s_0$, and $\mathcal{M}_1', s_1' \simeq_T \mathcal{M}_1, s_1$.

Proof. Without loss of generality, we focus on the $\mathcal{M}_0, \mathcal{M}_0'$ case. It suffices to observe that Duplicater can ensure that for every configuration (t_0, t_0') reached in the game $\mathsf{Bis}_T(\mathcal{M}_0, s_0, \mathcal{M}_0', s_0')$, either

- t_0' is a path with last element t_0, i.e. t_0' is a representative of t_0; or
- t_0' is a copy of t_0 in one of the isomorphic copies $\mathcal{M}_0^\pi$ of $\mathcal{M}_0$.

Verifying this is straightforward, so we leave it to the reader. Clearly a strategy which preserves this configuration-property is winning for Duplicater. □

Claim 2. $\mathcal{M}_0', s_0' \simeq_T^{(k)} \mathcal{M}_1', s_1'$.

Proof. We must show that Duplicater has a winning strategy in

$$\mathsf{Bis}_T((\mathcal{M}_0' {\restriction} (s_0', k), s_0'), (\mathcal{M}_1' {\restriction} (s_1', k), s_1'))$$

Observe that $\mathcal{M}_0' {\restriction} (s_0', k)$ and $\mathcal{M}_1' {\restriction} (s_1', k)$ are both two-way trees of height at most k relative to s_0', s_1' respectively. Further, observe that Duplicater may ensure that for every configuration (t_0', t_1') reached in the game, there are $t_0 \in S_0, t_1 \in S_1$ such that t_0' is a representative of t_0, t_1' is a representative of t_1, and $\mathcal{M}_0, t_0 \simeq_T^m \mathcal{M}_1, t_1$ (where $m = k - h(t_0') = k - h(t_1')$). A strategy which ensures this is clearly winning for Duplicater. □

This completes the proof of Proposition 4.2. □

In order to obtain the desired characterisation, we combine the upgrading result with the following important proposition:

Proposition 4.6. *Let $\phi \in FOL^\infty$ be $\simeq_T$-invariant. Then there is some natural number k such that ϕ is k-local.*

Proof. Bellas Acosta and Venema show in [2] (Proposition 3.7) that if $\phi \in FOL^\infty$ is invariant under taking disjoint unions, then ϕ is k-local for some natural number k. But it is easy to see that for any structures $\mathcal{M}, \mathcal{N}$ and $s \in \mathcal{M}$:

$$\mathcal{M}, s \simeq_T \mathcal{M} \oplus \mathcal{N}, s$$

So any ϕ which is $\simeq_T$- invariant must be invariant under taking disjoint unions. □

We now have all the ingredients we need to prove the following theorem.

Theorem 4.7. *Let $\phi \in FOL^\infty$. Then ϕ is invariant over $\simeq_T$ iff ϕ is logically equivalent to a formula in T^∞.*

Proof. We only need to prove the left-to-right direction. Suppoe ϕ is $\simeq_T$-invariant. Then by Proposition 4.6, there is some k such that ϕ is k-local. We claim that ϕ is invariant over $\simeq_T^{(k)}$. For suppose $\mathcal{M}_0, s_0 \simeq_T^{(k)} \mathcal{M}_1, s_1$. Then

$$\begin{aligned}
\mathcal{M}_0, s_0 \models \phi &\iff \mathcal{M}_0{\restriction}(s_0, k), s_0 \models \phi && (\phi \text{ is } k\text{-local})\\
&\iff \mathcal{M}_1{\restriction}(s_1, k), s_1 \models \phi && (\phi \text{ is } \simeq_T \text{-invariant})\\
&\iff \mathcal{M}_1, s_1 \models \phi && (\phi \text{ is } k\text{-local})
\end{aligned}$$

(this is why we chose $\simeq_T^{(k)}$ as our target equivalence).

But then Proposition 4.2 implies that ϕ is invariant over $\simeq_T^k$. For if $\mathcal{M}_0, s_0 \simeq_T^k \mathcal{M}_1, s_1$, then let $\mathcal{M}_0', s_0' \simeq_T^{(k)} \mathcal{M}_1', s'$ be bisimilar companions as given by Proposition 4.2. Then we have

$$\begin{aligned}
\mathcal{M}_0, s_0 \models \phi &\iff \mathcal{M}_0', s_0' \models \phi && (\text{invariance under } \simeq_T)\\
&\iff \mathcal{M}_1', s_1' \models \phi && (\text{invariance under } \simeq_T^{(k)})\\
&\iff \mathcal{M}_1, s_1 \models \phi && (\text{invariance under } \simeq_T)
\end{aligned}$$

So we can conclude that ϕ is $\simeq_T^k$-invariant.

Now, let τ_ϕ be the signature consisting of the non-logical symbols occurring in ϕ; then τ_ϕ is finite. By Proposition 2.28, $\simeq_T^k$ has finite index relative to τ_ϕ. Let $E_1, \ldots, E_m$ be an enumeration of the cells of this equivalence relation.

Let $i, j \leq m$, with $i \neq j$. By Proposition 2.27, there is a formula $\psi_{i,j} \in T^\infty$ such that ψ is satisfied on all pointed structures within C_i and on no pointed structures within C_j. For each $i \leq m$, define $\chi_i := \bigwedge_{j \neq i} \psi_{i,j}$. Let

$$\mathcal{T}(\phi) := \{i \leq m : \text{ there is some } \mathcal{M}, s \in E_i \text{ such that } \mathcal{M}, s \models \phi\}$$

Since ϕ is $\simeq_T^k$ invariant, ϕ is satisfied by all pointed structures in E_i, for any $i \in \mathcal{T}(\phi)$. Hence ϕ is logically equivalent to $\bigvee_{i \in \mathcal{T}(\phi)} \chi_i$, which is a formula of T^∞. □

For a pictorial representation of the structure of this proof, we repeat Figure 4.1 overleaf, with the specific relations $\simeq_T, \simeq_T^k, \simeq_T^{(k)}$ inserted. Given that ϕ is $\simeq_T$-invariant, the truth value of ϕ must be preserved from $\mathcal{M}_0, s_0$ to $\mathcal{M}_1, s_1$ because it is preserved over both vertical arrows, and over the lower horizontal arrow.

$$\begin{tikzcd}
\mathcal{M}_0, s_0 \arrow[r, leftrightarrow, "\simeq_T^k"] \arrow[d, "\simeq_T"'] & \mathcal{M}_1, s_1 \arrow[d, "\simeq_T"] \\
\mathcal{M}_0', s_0' \arrow[r, leftrightarrow, "\simeq_T^{(k)}"'] & \mathcal{M}_1', s_1'
\end{tikzcd}$$

Figure 4.2: Upgrading from $\simeq_T^k$ to $\simeq_T^{(k)}$

4.3 Upgrading from $\simeq_G^k$ and $\simeq_{GT}^k$

For the proof that T^∞ is characterised by $\simeq_T$-invariance, we chose $\simeq_T^{(k)}$ as our target equivalence relation, using the fact that formulas invariant over $\simeq_T$ must be invariant over taking disjoint unions, and hence k-local for some k. But this does not hold in the case of $\simeq_{GT}$, or even $\simeq_G$. The sentence $\exists x P x$ is invariant over $\simeq_G$, since it is in G^∞ (as the standard translation of Ep). But it is not invariant over disjoint unions: let $\mathcal{M}_0$ be a model consisting of a single point, s_0, which does not satisfy P, and let $\mathcal{M}_1$ be a model consisting of a single point, s_1, which does satisfy P. Then $\mathcal{M}_0 \not\models \exists x P x$, but $\mathcal{M}_0 \oplus \mathcal{M}_1 \models \exists x P x$.

So, FOL^∞ sentences invariant over $\simeq_G$ are not necessarily invariant over $\simeq_G^{(k)}$. We need to use a different target equivalence, and this is where we will use the Gaifman Theorem for FOL^∞ shown in Chapter 3.

Definition 4.8 (Gaifmanisation). Let k, m be natural numbers, and $\psi(x)$ a formula of FOL^∞ with only the variable x free. We define the k, m-*Gaifmanisation* of ψ, notation: $Gaif_{k,m}(\psi)$, to be the FOL^∞-sentence:

$$\exists x_1 \ldots \exists x_m \big(\bigwedge_{1 \leq i < j \leq m} d(x_i, x_j) > k \wedge \bigwedge_{1 \leq i \leq m} \psi(x_i) \big)$$

We define the k, m- *infinity Gaifmanisation* of ψ, notation: $Gaif_{k,m}^\infty(\psi)$, to be the FOL$^\infty$-sentence:

$$\exists^\infty x_1 \ldots \exists^\infty x_m \big(\bigwedge_{1 \leq i < j \leq m} d(x_i, x_j) > k \wedge \bigwedge_{1 \leq i \leq m} \psi(x_i) \big)$$

Definition 4.9 (FOL$^\infty$ equivalences). Let $\mathfrak{M}_0, \mathfrak{M}_1$ be structures. We write:

$$\mathfrak{M}_0 \equiv_{n,\exists}^{k,q} \mathfrak{M}_1$$

iff for every k-local formula $\psi(x)$ with $qd(\psi) \leq q$, for every $m \leq n$, it holds that $\mathfrak{M}_0 \models \mathit{Gaif}_{k,m}(\psi)$ iff $\mathfrak{M}_1 \models \mathit{Gaif}_{k,m}(\psi)$.

We write:

$$\mathfrak{M}_0 \equiv^{k,q}_{n,\exists\infty} \mathfrak{M}_1$$

iff for every k-local formula $\psi(x)$ with $qd(\psi) \leq q$, for every $m \leq n$, it holds that $\mathfrak{M}_0 \models \mathit{Gaif}^{\infty}_{k,m}(\psi)$ iff $\mathfrak{M}_1 \models \mathit{Gaif}^{\infty}_{k,m}(\psi)$.

We write:

$$\mathfrak{M}_0 \equiv^{k,q}_{n} \mathfrak{M}_1$$

iff both $\mathfrak{M}_0 \equiv^{k,q}_{n,\exists} \mathfrak{M}_1$ and $\mathfrak{M}_0 \equiv^{k,q}_{n,\exists\infty} \mathfrak{M}_1$.

Now, consider any sentence ϕ of FOL^{∞}. By Theorem 3.8, ϕ is equivalent to some ϕ', which is a Boolean combination of basic local sentences and basic local infinity sentences. Let k be the maximum of the locality ranks of the Boolean constituents of ϕ', and let n be the maximum of the scattering ranks of the basic local sentences and basic local infinity sentences involved in ϕ'. Let q be the greatest quantifier depth of a local formula occurring as a subformula of ϕ'. Then ϕ' (and hence ϕ) must be invariant over $\equiv^{k,q}_{n}$.

We can conclude that for *every* formula ϕ of FOL^{∞}, there are some $k, q, n \in \mathbb{N}$ such that ϕ is invariant over $\equiv^{k,q}_{n}$. This makes $\equiv^{k,q}_{n}$ a suitable target equivalence for an upgrading argument, as discussed in Section 4.1. Therefore, this section will be devoted to proving the following two main upgrading results.

Proposition 4.10 (Global upgrading)**.** *Modulo $\simeq_G$, $\simeq^{2k}_G$ can be upgraded to $\equiv^{k,q}_{n}$, for any $q, n \in \mathbb{N}$.*

Proposition 4.11 (Global two-way upgrading)**.** *Modulo $\simeq_{GT}$, $\simeq^{k}_{GT}$ can be upgraded to $\equiv^{k,q}_{n}$, for any $q, n \in \mathbb{N}$.*

Our strategy will be to prove a series of upgrading results, starting from the relation $\simeq^{2k}_G$, then upgrading modulo $\simeq_G$ to $\simeq^{k}_{GT}$, and then showing that this relation may be upgraded to progressively stronger equivalence relations modulo $\simeq_{GT}$, eventually reaching the target equivalence, $\equiv^{k,q}_{n}$.

4.3.1 Upgrading from $\simeq^{2k}_G$

In this subsection, we will show the following result:

Lemma 4.12. *Modulo $\simeq_G$, $\simeq^{2k}_G$ can be upgraded to $\simeq^{k}_{GT}$.*

The reader might wonder why we need to start from $\simeq^{2k}_G$, instead of simply $\simeq^{k}_G$. To see why, consider the structures $\mathcal{M}_0$ and $\mathcal{M}_1$ shown overleaf.

(a) $\mathcal{M}_0$ (b) $\mathcal{M}_1$

The reader may easily verify that $\mathcal{M}_0 \simeq_G^1 \mathcal{M}_1$. But any fully $\simeq_G$-bisimilar companion structure $\mathcal{M}'_0$ to $\mathcal{M}_0$ must have a path of length 2, while any fully $\simeq_G$-bisimilar companion structure $\mathcal{M}'_1$ to $\mathcal{M}_1$ cannot have a path of length 2. This means that Spoiler may win the game $\mathsf{Bis}^1_{GT}(\mathcal{M}'_0, \mathcal{M}'_1)$ by opening with a global move, and selecting the middle element of the length-2 path in $\mathcal{M}'_0$. This element has both a successor and a predecessor, but in response, Duplicater will have to select an element with either no successors or no predecessors. Spoiler can win on the next turn by selecting a successor or predecessor, as appropriate, of the element they chose in round 1.

We will find the required companion structures for Proposition 4.10 using another unravelling construction.

Definition 4.13 (Global unravelling)**.** Let $\mathcal{M}$ be a structure. The *global unravelling* of $\mathcal{M}$, notation: $\mathcal{U}_G(\mathcal{M})$, is defined as follows:

- The universe $U_G(\mathcal{M})$ of $\mathcal{U}_G(\mathcal{M})$ is the set of *all* paths through $\mathcal{M}$ (from any starting point), where a path is as defined in Definition 4.3
- Unary predicate symbols P are interpreted in $\mathcal{U}_G(\mathcal{M})$ by: $\pi \in P^{\mathcal{U}_G(\mathcal{M})}$ iff $last(\pi) \in P^{\mathcal{M}}$
- The interpretation of the binary predicate symbol R is: $\pi, \pi' \in R^{\mathcal{U}_G(\mathcal{M})}$ iff $\pi' = \pi^\frown\langle last(\pi')\rangle$

Lemma 4.14 (Bisimilarity across global unravelling)**.** *Let $\mathcal{M}$ be a structure. Then $\mathcal{U}_G(\mathcal{M}) \simeq_G \mathcal{M}$.*

Proof. We give a strategy for Duplicater in the game $\mathsf{Bis}_G(\mathcal{U}_G(\mathcal{M}), \mathcal{M})$ which preserves the property that, for every configuration (π, t) reached in the game, π is a representative of t.

In the initial round, if Spoiler selects an element $\pi \in U_G(\mathcal{M})$, then Duplicater should select $last(\pi) \in M$, whereas if Spoiler selects some $a \in M$, then Duplicater should select the length 0 path $\langle a\rangle \in U_G(\mathcal{M})$.

We now show that Duplicater can maintain the desired configuration-property in non-initial rounds. For brevity, we only discuss a couple of illustrative cases. Assume the game has reached a configuration (π, t) such that $t = last(\pi)$.

- Forward move in $\mathcal{M}$: Spoiler selects some successor t' of t. Then Duplicater should select $\pi^\frown\langle t'\rangle \in U_G(\mathcal{M})$.

- Forward infinity move in $\mathcal{U}_G(\mathcal{M})$: Spoiler selects some infinite set $\Theta \subseteq R[\pi]$. Then every $\theta \in \Theta$ can be written as $\pi^\frown\langle t'\rangle$ for some $t' \in M$. In other words, the paths in Θ will *only* differ on their last element. This means that the set $\{u \in M :$ some $\theta \in \Theta$ is a representative of $u\}$ is also infinite (indeed, has the same cardinality as Θ). Duplicater should select this set in response to Spoiler's move. Then whichever u from Duplicater's set is selected by Spoiler, Duplicater can select a representative θ of u from Θ.

A strategy which preserves the given configuration-property is clearly winning for Duplicater. $\square$

We have already characterised infinity bisimulation between structures in a game-theoretic manner, but it is useful to also characterise it in terms of a relation between the objects of the structures.

Definition 4.15 (Bisimulation relation)**.** Let $\mathcal{M}_0, \mathcal{M}_1$ be structures, $Z \subseteq M_0 \times M_1$. We say that Z is a bisimulation iff it satisfies the following clauses. For any s_0, s_1 such that $s_0 Z s_1$:

- (atomic) For all unary predicate letters P, s_0 satisfies P iff s_1 satisfies P

- (back) For any $s_1' \in R[s_1]$, there is an $s_0' \in R[s_0]$ such that $s_0' Z s_1'$

- (forth) For any $s_0' \in R[s_0]$, there is an $s_1' \in R[s_1]$ such that $s_0' Z s_1'$

- (∞-back) For any infinite set $Y \subseteq R[s_0]$, there is an infinite set $X \subseteq R[s_1]$ such that for any $s_1' \in X$, there is an $s_0' \in Y$ such that $s_0' Z s_1'$

- (∞-forth) For any infinite set $Y \subseteq R[s_1]$, there is an infinite set $X \subseteq R[s_0]$ such that for any $s_0' \in X$, there is an $s_0' \in Y$ such that $s_0' Z s_1'$.

A bisimulation Z between two structures $\mathcal{M}_0, \mathcal{M}_1$ will be called *global* iff for every $s_0 \in M_0$ there is an $s_1 \in M_1$ with $s_0 Z s_1$, and for every $s_1 \in M_1$ there is an $s_0 \in M_0$ with $s_0 Z s_1$. In fact, we will use the term *global* for any relation Z, bisimulation or not, which satisfies this condition.

It is easy to verify that $\mathcal{M}_0, s_0 \simeq \mathcal{M}_1, s_1$, in the sense of Definition 2.18, iff there is a bisimulation relation Z between $\mathcal{M}_0$ and $\mathcal{M}_1$, such that $s_0 Z s_1$. Further, $\mathcal{M}_0 \simeq_G \mathcal{M}_1$ iff there exists a global bisimulation relation Z between $\mathcal{M}_0$ and $\mathcal{M}_1$.

We can also give an equivalent formulation of k-step (global) bisimulation, in terms of sequences of relations.

Definition 4.16 (k-bisimulations)**.** Let $\mathcal{M}_0, \mathcal{M}_1$ be structures. A *k-bisimulation* is an indexed sequence $(Z_i)_{0\leq i\leq k}$ of relations, with each Z_i satisfying the 'atomic' clause from Definition 4.15, and additionally satisfying one-step analogues of the simple and infinitary 'back' and 'forth' clauses. To demonstrate the idea, we give just the infinitary 'back' clause:

- (∞-back) Let $\langle s_0, s_1\rangle \in Z_{i+1}$. For any infinite $Y \subseteq R[s_1]$, there is an infinite $X \subseteq R[s_0]$ such that for any $x \in X$, there is a $y \in Y$ with $\langle x, y\rangle \in Z_i$

We call a k-bisimulation $(Z_i)_{0\leq i\leq k}$ *global* iff Z_k is global.

Again, it is easy to verify that $\mathcal{M}_0, s_0 \simeq^k \mathcal{M}_1, s_1$ iff there exists a k-bisimulation with $\langle s_0, s_1\rangle \in Z_k$, and $\mathcal{M}_0 \simeq_G^k \mathcal{M}_1$ iff there exists a global k-bisimulation between $\mathcal{M}_0$ and $\mathcal{M}_1$.

We now borrow another notion from Otto [15].

Definition 4.17 (Respecting zero in-degree)**.** Let $\mathcal{A}$ be a Kripke structure, $a \in A$. We say that *a has zero in-degree* iff there is no $a' \in A$ such that $a' R^{\mathcal{A}} a$.

Let $\mathcal{A}, \mathcal{B}$ be Kripke structures and $Z \subseteq A \times B$ a relation. We say that Z *respects zero in-degree* iff:

- For every $a \in A$ such that a has zero in-degree, there is a $b \in B$ also with zero in-degree such that aZb

- For every $b \in B$ such that b has zero in-degree, there is an $a \in A$ also with zero in-degree such that aZb

Lemma 4.18. *Let $\mathcal{A}, \mathcal{B}$ be forest structures. If there exists a global $2k$-bisimulation $(Z_i)_{i\leq 2k}$ between $\mathcal{A}$ and $\mathcal{B}$ such that Z_{2k} respects zero in-degree, then $\mathcal{A} \simeq_{GT}^k \mathcal{B}$.*

Proof. We show that there is a winning strategy in $\mathsf{Bis}_{GT}^k(\mathcal{A}, \mathcal{B})$. We begin by explaining how Duplicater should respond to Spoiler's global move in the initial round.

Without loss of generality, suppose Spoiler chooses an element a from $\mathcal{A}$. Consider the unique path leading backwards from a in $\mathcal{A}$. Within a finite number of steps, this backwards path must reach an element with zero in-degree, as $\mathcal{A}$ is a forest structure. So we may write this path as $a = a_0, a_{-1}, \ldots, a_{-l}$ for some natural number l, with a_{-l} having zero in-degree. We distinguish two cases.

Case A: $l < k$. In this case, consider the element a_{-l}. Because Z_{2k}, the first relation in the bisimulation-sequence, is global and respects zero in-degree, there must be some b_{-l} in $\mathcal{B}$ which also has zero in-degree, such that $a_{-l} Z_{2k} b_{-l}$. Further, starting from a_{-l}, b_{-l} and repeatedly exercising the back-and-forth properties of the Z_i's, there must be a sequence $b_{-l} R b_{1-l} R \ldots R b_0$ such that for each $i \leq l$, $(a_{-i}, b_{-i}) \in Z_{2k-(l-i)}$. In particular, then, $a = a_0 \; Z_{2k-l} \; b_0$. Duplicater should choose $b = b_0$ in response to Spoiler's choice of a. Observe that since $l < k$, $a \simeq^k b$.

Case B: $l \geq k$. In this case, consider the element a_{-k}. Again because Z_{2k} is global, there is some b_{-k} in B such that $a_{-k} Z_{2k} b_{-k}$. In fact, there is a sequence $b_{-k} R b_{1-k} R \ldots R b_0$, such that $a_{-i} Z_{k+i} b_{-i}$ for each $i \leq k$, and in particular, $a = a_0 \; Z_k \; b_0$. Duplicater should choose $b = b_0$ in response to Spoiler's choice of a.

We additionally define a rank function on the sets A, B. Define $a^* := a_{-l}$ iff $l < k$, $a^* = a_{-k}$ otherwise. Similarly, define $b^* := b_{-l}$ iff $l < k$, $b^* = b_{-k}$ otherwise. Next, define $\mathit{rank}(a^*) := \mathit{rank}(b^*) := 0$, and for any $x \in A \cup B$, $\mathit{rank}(x) := n+1$ iff there exists an x' with $x' R x$ and $\mathit{rank}(x') = n$. Set $\mathit{rank}(x) := \infty$ iff there is no $n \in \mathbb{N}$ such that $\mathit{rank}(x) = n$. Then every element which could be involved in a play in the game $\mathsf{Bis}_T^k(\mathcal{A}, a, \mathcal{B}, b)$ has a unique natural number as its rank (unique because of the forest structure of $\mathcal{A}$ and $\mathcal{B}$). Also, $\mathit{rank}(a) = \mathit{rank}(b)$, and for any configuration (a', b') that could be reached in the game $\mathsf{Bis}_T(\mathcal{A}, a, \mathcal{B}, b)$, we have $\mathit{rank}(a') = \mathit{rank}(b')$.

We now give a strategy for Duplicater in the remaining rounds of the game. These remaining rounds simply consist of the game $\mathsf{Bis}_T^k(\mathcal{A}, a, \mathcal{B}, b)$. Our strategy will preserve the property that, if (a_n, b_n) is the configuration reached after n non-initial rounds, then for *all* $m \leq n$, $a_m \simeq^{k-m} b_m$ (where (a_m, b_m) denotes the configuration reached after m non-initial rounds). Clearly this condition is satisfied by the initial configuration (a, b).

Suppose the game reaches a configuration (a_n, b_n) such that for all $m \leq n$, $a_m \simeq^{k-m} b_m$. There are various cases to consider, according to what type of move Spoiler makes.

Case 1: Spoiler plays a forward move or forward infinity move, selecting a single successor or infinite set of successors of a_n or b_n. Because $a_n \simeq^{k-n} b_n$, there is a winning strategy for Duplicater in the $k-n$-round version of $\mathsf{Bis}(\mathcal{A}, a_n, \mathcal{B}, b_n)$. Further, Spoiler's move would be a legal move in this game, so the winning strategy for Duplicater must prescribe a response to it. Duplicater should follow this response. It is obvious that this preserves the desired configuration-property.

Case 2: Spoiler plays a backward move, and selects a single predecessor of a_n or b_n. Without loss of generality, suppose it's a predecessor of a_n. Firstly, we argue that b_n also has a predecessor. As observed above, it must hold that $\mathit{rank}(a_n) = \mathit{rank}(b_n)$.

But a_n has a predecessor, and it cannot be that $a_n = a_{-k}$ since then round k would already have been reached, and the game would be over. Therefore, $rank(a_n) > 0$. It follows that $rank(b_n) > 0$, and so b_n has a predecessor. This predecessor must be unique, because $\mathcal{B}$ is a forest structure. Duplicater should select this unique predecessor.

Let a_{n+1} denote the element selected by Spoiler, and b_{n+1} the element selected by Duplicater. To see that $a_{n+1} \simeq^{k-(n+1)} b_{n+1}$, we can make a further case distinction.

Case (i): The configuration (a_{n+1}, b_{n+1}) has not been reached before in the game. This implies that a_{n+1} and b_{n+1} are ancestors of a and b respectively. Because $rank(a_{n+1}) = rank(b_{n+1})$, it follows that $a_{n+1} = a_{-j}$, $b_{n+1} = b_{-j}$ for some $j \leq k$. But then $a_{n+1} Z_{2k-(l-j)} b_{n+1}$ for some $l \leq k$, which implies $a_{n+1} \simeq^{k-(n+1)} b_{n+1}$.
Case (ii): The configuration (a_{n+1}, b_{n+1}) has been reached before. Then $a_{n+1} = a_m, b_{n+1} = b_m$ for some $m < n+1$. But we assume that $a_m \simeq^{k-m} b_m$ for all $m \leq n$, therefore $a_{n+1} \simeq^{k-m} b_{n+1}$, which trivially implies $a_{n+1} \simeq^{k-(n+1)} b_{n+1}$.

In both cases we have $a_{n+1} \simeq^{k-(n+1)} b_{n+1}$, so this wraps up the proof for Case 2.

Case 3: Spoiler selects an infinite set of predecessors of a_n or b_n. But this case is impossible, because $\mathcal{A}$ and $\mathcal{B}$ are forests and so every element has at most one predecessor.

We have given a strategy for Duplicater which preserves the desired property, and it is obvious that a strategy which preserves this property must be winning. Thus $\mathcal{A} \simeq_{GT}^k \mathcal{B}$. □

Lemma 4.19. *Let $\mathcal{A}, \mathcal{B}$ be structures such that $\mathcal{A} \simeq_G^{2k} \mathcal{B}$. Then $\mathcal{U}_G(\mathcal{A}) \simeq_{GT}^k \mathcal{U}_G(\mathcal{B})$.*

Proof. Clearly, $\mathcal{U}_G(\mathcal{A})$ and $\mathcal{U}_G(\mathcal{B})$ are forest structures. So it suffices to show that there is a $2k$-bisimulation $(Z_i)_{0 \leq i \leq 2k}$ between $\mathcal{U}_G(\mathcal{A})$ and $\mathcal{U}_G(\mathcal{B})$, such that Z_{2k} is global and respects zero in-degree.

We know that $\mathcal{A} \sim_G^{2k} \mathcal{B}$, so let (Z'_i) be a global $2k$-bisimulation between $\mathcal{A}$ and $\mathcal{B}$. Define $Z_i \subseteq U_G(\mathcal{A}) \times U_G(\mathcal{B})$ by: $\pi Z_i \rho$ iff there are $a \in A, b \in B$ such that π is a representative of a, ρ is a representative of b and $a Z'_i b$. Then the sequence $(Z_i)_{0 \leq i \leq 2k}$ is a $2k$-bisimulation between $\mathcal{U}_G(\mathcal{A})$ and $\mathcal{U}_G(\mathcal{B})$.

Further, any $\pi \in U_G(\mathcal{A})$ is a representative of some a, and because (Z'_i) is global, a must be linked by Z_{2k} to some b. This b has a representative with zero in-degree, namely the 0-step path $\langle b \rangle$. The same holds, *mutatis mutandis*, for any $\rho \in U_G(\mathcal{B})$. This shows both that Z_{2k} is global, and that it respects zero in-degree. □

Proof of Lemma 4.12. Lemmas 4.14 and 4.19 immediately imply the desired result. If $\mathcal{M}_0 \simeq_G^{2k} \mathcal{M}_1$, then we will have $\mathcal{U}_G(\mathcal{M}_0) \simeq_G \mathcal{M}_0$, $\mathcal{U}_G(\mathcal{M}_1) \simeq_G \mathcal{M}_1$, and $\mathcal{U}_G(\mathcal{M}_0) \simeq_{GT}^k \mathcal{U}_G(\mathcal{M}_1)$. So $\mathcal{U}_G(\mathcal{M}_0)$ and $\mathcal{U}_G(\mathcal{M}_1)$ can serve as the desired companion structures. □

Remark. The above argument is very similar to the argument given for Lemma 38 in Otto [15] (compare especially our Lemma 4.18 with Otto's Lemma 40). However, the argument is in fact slightly easier in our case, because we are not restricting ourselves to finite structures. This means that we can take the full global unravelling of our structures $\mathcal{A}$ and $\mathcal{B}$ to be 'upgraded', giving us forest structures as bisimilar companions. This cannot be done in the finite model theory setting, because the unravelling of even a finite structure may well be infinite.

4.3.2 Upgrading from $\simeq_{GT}^k$

Having shown that $\simeq_G^{2k}$ can be upgraded to $\simeq_{GT}^k$ modulo $\simeq_G$, we now turn to showing that $\simeq_{GT}^k$ can be upgraded to $\equiv_n^{k,q}$ modulo $\simeq_{GT}$, as this will imply that $\simeq_G^{2k}$ can be upgraded to this stronger equivalence as well. We will use some further definitions.

Definition 4.20 (Global two-way unravelling)**.** Let $\mathcal{A}$ be a structure. The *global two-way unravelling* of $\mathcal{A}$, notation: $\mathcal{U}_{GT}(\mathcal{A})$, is the structure whose universe is the set of all two-way paths through $\mathcal{A}$ (starting from any point), and which interprets predicate symbols as in the non-global two-way unravelling (see Definition 4.4).

Lemma 4.21 (Bisimilarity across global two-way unravelling)**.** *Let $\mathcal{A}$ be a structure. Then $\mathcal{U}_{GT}(\mathcal{A}) \simeq_{GT} \mathcal{A}$.*

Proof. The argument is almost identical to the proof of Lemma 4.14, with the obvious changes made to handle backward moves and backward steps in paths. Again, the winning strategy for Duplicater is to ensure that for every configuration (π, a) reached in the game, π is a representative of a. For the sake of illustration we discuss one possible case. Assume the game has reached a configuration π, a such that $a = \mathit{last}(\pi)$.

- Forward infinity move in $\mathcal{U}_{GT}(\mathcal{A})$: Spoiler selects an infinite set $\Theta \subseteq R[\pi]$. Then there may be one $\theta \in \Theta$ such that $\pi = \theta^\frown \mathit{last}(\pi)$. But for every $\theta' \in \Theta' := \Theta \setminus \{\theta\}$, we will have $\theta' = \pi^\frown \mathit{last}(\theta')$. Further, Θ' must also be infinite, and for any $\rho, \sigma \in \Theta'$, $\mathit{last}(\rho) \neq \mathit{last}(\sigma)$. So the set $\{\mathit{last}(\theta') : \theta' \in \Theta'\}$ is infinite. Duplicater should select this set. Then whichever object a' Spoiler selects from this set, Duplicater should select $\pi^\frown\langle a'\rangle$.

□

Definition 4.22 (Disjoint copies)**.** Let $\mathfrak{A}$ be a structure. For each cardinal $\kappa \in \mathbb{N} \cup \{\aleph_0\}$, we let $\mathfrak{A} \cdot \kappa$ be the structure with universe $A \times \kappa$, and relations interpreted by: $\mathfrak{A} \cdot \kappa \models R(a_1, i_1) \ldots (a_n, i_n)$ iff $i_1 = \ldots = i_n$ and $\mathfrak{A} \models Ra_1 \ldots a_n$. So $\mathfrak{A} \cdot \kappa$ consists of κ-many disjoint isomorphic copies of $\mathfrak{A}$.

Lemma 4.23 (Bisimilarity across disjoint copies)**.** *Let $\mathcal{A}$ be a Kripke structure, and $\kappa \in \mathbb{N} \cup \{\aleph_0\}$. Then $\mathcal{A} \simeq_{GT} \mathcal{A} \cdot \kappa$.*

Proof. It suffices to note that Duplicater may preserve the property that, for every configuration $((a', j), a)$ reached in $\mathsf{Bis}_{GT}(\mathcal{A} \cdot \kappa, \mathcal{A})$, $a' = a$. The only case that needs care is if Spoiler selects an infinite set X from $\mathcal{A} \cdot \kappa$.

Suppose the game has reached configuration $((a, j), a)$, and Duplicater selects an infinite set $X \subseteq A \times \kappa$ of successors of (a, j) (the predecessor case is analogous). Then by definition of $\mathcal{A} \cdot \kappa$, every $x \in X$ is of the form (a', j) for some $a' \in A$. Hence the set $\{a' \in A : (a', j) \in X\}$ is also infinite, and is thus a legal response for Duplicater. Whichever a' is selected by Spoiler, Duplicater may select (a', j). $\square$

Definition 4.24 (Blowing up multiplicities)**.** Let $\mathfrak{A}$ be a structure, $q \in \mathbb{N}$. We define $\mathfrak{A} \otimes q$ to be the structure whose universe is $A \times \{0, 1, \ldots, q-1\}$, and whose interpretation of predicates is given by: $R^{\mathfrak{A} \otimes q}(\langle a_1, i_1 \rangle, \ldots, \langle a_n, i_n \rangle)$ iff $R^{\mathfrak{A}}(a_1, \ldots, a_n)$.

Note the difference between the operations $\cdot$ and $\otimes$; whereas $\cdot$ results in a certain number of disjoint copies of the original structure, $\otimes$ allows for relation symbols to hold 'across different indices' and so creates copies which are connected to each other.

Lemma 4.25 (Bisimilarity across $\otimes$)**.** *Let $\mathcal{A}$ be a Kripke structure, $q \in \mathbb{N}$. Then $\mathcal{A} \otimes q \simeq_{GT} \mathcal{A}$.*

Proof. It suffices to note that in the game $\mathsf{Bis}_{GT}(\mathcal{A} \otimes q, \mathcal{A})$, Duplicater can preserve the property that, for each configuration $((a', i), a)$ reached in the game, $a' = a$. There is one case which requires some argument. Suppose the game has reached a configuration $((a, i), a)$.

- Infinity move in $\mathcal{A} \otimes q$: Spoiler selects an infinite set X of R-successors or R-predecessors of (a, i). Observe that the image of X under the left-hand projection map must be an infinite set of objects in A, by the finiteness of q. Duplicater should select this infinite set. Then whichever element a is selected by Spoiler, there will be some k such that $(a, k) \in X$. Duplicater should select the element (a, k).

$\square$

Definition 4.26 (q-graded two-way bisimulation)**.** Let $(\mathcal{A}, a), (\mathcal{B}, b)$ be pointed Kripke structures. The *q-graded two-way infinity bisimulation game* on $\mathcal{A}, a, \mathcal{B}, b$, notation: $\mathsf{Bis}_{Tq}(\mathcal{A}, a, \mathcal{B}, b)$, is like $\mathsf{Bis}_{T}(\mathcal{A}, a, \mathcal{B}, b)$, except that Spoiler is allowed to make the following additional types of move:

- Forward graded move: Assume the game has reached a configuration (a', b'). Spoiler may select a set X of successors of a' (or b', *mutatis mutandis*), with $|X| \leq q$. Duplicater must select a set Y of successors of b' (a'), with $|Y| = |X|$. Spoiler chooses an element y from Y, and Duplicater must choose an element x from X. The new configuration is (x, y).
- Backward graded move: As above, but with 'successor' everywhere replaced by 'predecessor'.

We write $\mathcal{A}, a \simeq_{Tq} \mathcal{B}, b$ iff Duplicater has a winning strategy in $\mathsf{Bis}_{Tq}(\mathcal{A}, a, \mathcal{B}, b)$, and $\mathcal{A}, a \simeq_{Tq}^{k} \mathcal{B}, b$ iff Duplicater has a winning strategy in the k-round version of this game.

Definition 4.27 (Ehrenfeucht-Fraissé game of $\exists^{\infty}$)**.** Let $\mathfrak{A}, \mathfrak{B}$ be structures and let $\overline{a}, \overline{b}$ be ordered n-tuples with $\overline{a} \subseteq A, \overline{b} \subseteq B$. The *Ehrenfeucht-Fraissé Game of* $\exists^{\infty}$ of k rounds on $\mathfrak{A}, \overline{a}, \mathfrak{B}, \overline{b}$, notation: $EF_{k}^{\infty}(\mathfrak{A}, \overline{a}, \mathfrak{B}, \overline{b})$, is just like the regular EF-game of k rounds on $\mathfrak{A}, \mathfrak{B}$ beginning from the configuration $(a_1, \ldots, a_n), (b_1, \ldots, b_n)$, except that Spoiler is allowed to make an additional type of move:

- Infinity move: assume a position has been reached with configuration $(a_1, \ldots, a_m), (b_1, \ldots, b_m)$. Spoiler may select an infinite set X from the structure $\mathfrak{A}$ (or $\mathfrak{B}$, *mutatis mutandis*). Duplicater must select an infinite set Y from $\mathfrak{B}$ $(\mathfrak{A})$. Spoiler then selects an element $y \in Y$, and Duplicater must select an element $x \in X$. The new configuration is $(a_1, \ldots, a_m, x), (b_1, \ldots, b_m, y)$.

For any structures $\mathfrak{A}, \mathfrak{B}$ and ordered n-tuples $\overline{a} \subseteq A, \overline{b} \subseteq B$, We write $\mathfrak{A}, \overline{a} \simeq_{EF}^{k} \mathfrak{B}, \overline{b}$ iff Duplicater has a winning strategy in $EF_{k}^{\infty}(\mathfrak{A}, \overline{a}, \mathfrak{B}, \overline{b})$.

This is a particular instance of the *Ehrenfeucht-Fraissé Game of Q* defined in [17] (Definition 10.26), for the monotone generalised quantifier $Q = \exists^{\infty}$. Therefore, we can appeal to Theorem 10.46 of [17] (which covers the general case, for arbitrary monotone Q) to infer:

Lemma 4.28. *Let $\mathfrak{A}, \mathfrak{B}$ be structures and $\overline{a}, \overline{b}$ ordered n-tuples within $\mathfrak{A}, \mathfrak{B}$ respectively. Then the following are equivalent:*

- $\mathfrak{A}, \overline{a} \simeq_{EF}^{k} \mathfrak{B}, \overline{b}$
- $\mathfrak{A}, \overline{a} \equiv^{k} \mathfrak{B}, \overline{b}$

We now turn to the most important and difficult of the upgrading propositions to be proved in this section.

Proposition 4.29 (Upgrading from $\simeq_{GT}^k$ to $\equiv_{1,\exists}^{k,q}$)**.** *Let $k, q \in \mathbb{N}$. Modulo $\simeq_{GT}$, $\simeq_{GT}^k$ can be upgraded to $\equiv_{1,\exists}^{k,q}$.*

To provide the required companion structures, we use the constructions defined above. Let $\mathcal{M}_0, \mathcal{M}_1$ be Kripke structures such that $\mathcal{M}_0 \simeq_{GT}^k \mathcal{M}_1$. We claim:

1. $\mathcal{U}_{GT}(\mathcal{M}_0 \otimes q) \simeq_{GT} \mathcal{M}_0$
2. $\mathcal{U}_{GT}(\mathcal{M}_1 \otimes q) \simeq_{GT} \mathcal{M}_1$
3. $\mathcal{U}_{GT}(\mathcal{M}_0 \otimes q) \equiv_{1,\exists}^{k,q} \mathcal{U}_{GT}(\mathcal{M}_1 \otimes q)$

and hence we may take $\mathcal{U}_{GT}(\mathcal{M}_0 \otimes q)$ and $\mathcal{U}_{GT}(\mathcal{M}_1 \otimes q)$ as our desired companion structures. Further, 1. and 2. are immediate by Lemmas 4.21 and 4.25, so we only need to prove item 3.

To prove item 3 we must show that for every k-local formula ψ with $qd(\psi) \leq q$, $\mathcal{U}_{GT}(\mathcal{M}_0 \otimes q) \models \exists x \psi(x)$ iff $\mathcal{U}_{GT}(\mathcal{M}_1 \otimes q) \models \exists x \psi(x)$. Since we only need to take care of k-local formulas ψ, it suffices to verify that for any $\pi_0 \in \mathcal{U}_{GT}(\mathcal{M}_0 \otimes q)$, there is a $\pi_1 \in \mathcal{U}_{GT}(\mathcal{M}_1 \otimes q)$ such that

$$\mathcal{U}_{GT}(\mathcal{M}_0 \otimes q) \restriction (\pi_0, k), \pi_0 \equiv^q \mathcal{U}_{GT}(\mathcal{M}_1 \otimes q) \restriction (\pi_1, k), \pi_1$$

(By symmetry, we then obtain that the same holds *vice versa*: for any π_1 there is a π_0...)

Let us introduce the notation p_l for the left-hand projection map on ordered pairs, and p_r for the right-hand projection map. Then observe that for every $\pi_0 \in \mathcal{U}_{GT}(\mathcal{M}_0 \otimes q)$, there is a $\pi_1 \in \mathcal{U}_{GT}(\mathcal{M}_1 \otimes q)$ such that $\mathcal{M}_0, p_l(last(\pi_0)) \simeq_T^k \mathcal{M}_1, p_l(last(\pi_1))$. (We can just take $\pi_1 = \langle (a_1', 0) \rangle$, for some $a_1' \simeq_T^k p_l(last(\pi_0))$; such an a_1' must exist because $\mathcal{M}_0$ is *globally* two-way k-bisimilar to $\mathcal{M}_1$.) Finally, observe that for any $\pi_0 \in U_{GT}(\mathcal{M}_0 \otimes q)$,

$$(\mathcal{U}_{GT}(\mathcal{M}_0 \otimes q) \restriction (\pi_0, k), \pi_0)$$

is a two-way tree of height k. Likewise for any $\pi_1 \in U_{GT}(\mathcal{M}_1 \otimes q)$.

In view of these observations, the following two lemmas suffice to prove Proposition 4.29:

Lemma 4.30. *Let $\pi_0 \in U_{GT}(\mathcal{M}_0 \otimes q), \pi_1 \in U_{GT}(\mathcal{M}_1 \otimes q)$ be such that $\mathcal{M}_0, p_l(last(\pi_0)) \simeq_T^k \mathcal{M}_1, p_l(last(\pi_1))$. Then $\mathcal{U}_{GT}(\mathcal{M}_0 \otimes q), \pi_0 \simeq_{Tq}^k \mathcal{U}_{GT}(\mathcal{M}_1 \otimes q), \pi_1$.*

Lemma 4.31. *Let $\mathbb{T}, \mathbb{T}'$ be two-way trees of height at most k relative to roots r, r'. Then $\mathbb{T}, r \simeq_{Tq}^k \mathbb{T}', r'$ implies $\mathbb{T}, r \equiv^q \mathbb{T}', r'$.*

Proof of Lemma 4.30

We give a strategy for Duplicater in the game $\mathsf{Bis}^k_{Tq}(\mathcal{U}_{GT}(\mathcal{M}_0 \otimes q), \pi_0, \mathcal{U}_{GT}(\mathcal{M}_1 \otimes q), \pi_1)$ which preserves the property that, for each configuration ρ_0, ρ_1 reached after m rounds of the game, $\mathcal{M}_0, p_l(last(\rho_0)) \simeq^{k-m}_T \mathcal{M}_1, p_l(last(\rho_1))$. The initial configuration has this property by assumption. We only consider moves from Spoiler in the forwards direction, as the backward cases are analogous. We also neglect the forward simple move case, as it is a special case of the forward graded move. Assume the game has reached configuration ρ_0, ρ_1 after m rounds.

- Forward graded move: Spoiler selects a set X with $|X| \leq q$, and either $X \subseteq R[\rho_0]$ or $X \subseteq R[\rho_1]$. WLOG assume $X \subseteq R[\rho_0]$. Take an arbitrary $x \in X$. Then $p_l(last(x))$ is a successor of $p_l(last(\rho_0))$, so would be a legal move for Spoiler in $\mathsf{Bis}^{k-m}_T(\mathcal{M}_0, p_l(last(\rho_0)), \mathcal{M}_1, p_l(last(\rho_1)))$. Let $a'_1 \in M_1$ be the response prescribed by the winning strategy for Duplicater in this Bis^{k-m}_T game. Define $Y := \{\rho_1^\frown(a'_1, i) : i < |X|\}$. Duplicater should select Y as her response to Spoiler, and whichever $y \in Y$ is chosen by Spoiler, Duplicater should respond by choosing x.

- Forward infinity move: WLOG suppose Spoiler selects an infinite set $X \subseteq R[\rho_0]$. For every distinct $x, x' \in X$, x, x' only differ in their last element. So since q is finite, by the pigeonhole principle, the set $\{p_l(last(x)) : x \in X\}$ is infinite. Therefore, it would be a legitimate forward infinity move for Spoiler in $\mathsf{Bis}^{k-m}_T(\mathcal{M}_0, p_l(last(\rho_0)), \mathcal{M}_1, p_l(last(\rho_1)))$. Let $S \subseteq M_1$ be the response prescribed by the winning strategy for Duplicater in this game. Define $Y := \{\rho_1^\frown(s, 0) : s \in S\}$. Duplicater should choose Y in response to Spoiler's move X. Whichever $\rho_1^\frown(s, 0)$ is chosen by Spoiler, Duplicater should choose an element $x \in X$ such that $p_l(last(x))$ is the response prescribed for Duplicater in $\mathsf{Bis}^{k-m}_T(\mathcal{M}_0, p_l(last(\rho_0)), \mathcal{M}_1, p_l(last(\rho_1)))$ if Spoiler chooses s.

It is easy to see that a strategy which preserves the property stated above is winning for Duplicater. □

Proof of Lemma 4.31

We define a notion of *q-companionship*, which will allow us to link points in different tree models. This notion is inspired by the q-companionship relation defined in [2] (Definition 3.10), although note that the definition used here is subtly different.

Definition 4.32 (q-companions)**.** Let $(\mathbb{T}, r), (\mathbb{T}', r')$ be two-way tree structures. Let $h(\mathbb{T}, r) = h(\mathbb{T}', r') = k$. Let $t \in T, t' \in T'$. Then t, t' are *q-companions*, notation: $t \sim_q t'$, iff the shortest path $\langle r = t_0, t_1, \ldots, t_m = t\rangle$ from r to t and the shortest path $\langle r' = t'_0, t'_1, \ldots, t'_n = t'\rangle$ from r' to t' are such that:

- $m = n$

- For all i with $0 \leq i < m$, either $t_i R t_{i+1}$ and $t'_i R t'_{i+1}$ or $t_{i+1} R t_i$ and $t'_{i+1} R' t'_i$
- For all i with $0 \leq i \leq m$, $t_i \simeq_{Tq}^{k-h(t_i)} t'_i$

Lemma 4.33. *Let $(\mathbb{T}, r), (\mathbb{T}', r')$ be two-way trees of height k. Then $\sim_q$ is an equivalence relation of finite index over $T \cup T'$.*

Proof. It is immediate that $\sim_q$ is an equivalence relation. To see that it has finite index, first observe that every t in $(\mathbb{T}, r)$ has height at most k, and likewise for every t' in $\mathbb{T}', r'$. So it suffices to prove that, for any fixed m, $\sim_q$ has finite index on the set of all members of $T \cup T'$ with height exactly m. We argue by induction on m.

Base case: $m = 0$. Trivial, as only r, r' have height 0.

Induction step: inductively assume that $\sim_q$ has finite index on the set of all elements of $T \cup T'$ with height exactly $m - 1$. We prove the statement for m. To see that the statement holds, note that the q-companionship type of any element $t \in T \cup T'$ is fully determined by:

1. The q-companionship type of the unique 1-step ancestor of t
2. Whether the unique 1-step ancestor of t is an R-successor or R-predecessor of t
3. The $\simeq_{Tq}^{k-h(t)}$-type of t

Now suppose t has height exactly m. Then by the inductive hypothesis there are only finitely many possibilities for item 1; there are exactly 2 possibilities for item 2; and there are only finitely many possibilities for item 3, because $\simeq_{Tq}^{k-m}$ has finite index (by a straightforward adaptation of the argument for Proposition 2.28). Therefore $\sim_q$ has finite index on the set of all elements of height exactly m, which was to show for the induction step.

This concludes the induction, so the proof is complete. □

We now prove a key lemma about the q-companionship relation.

Lemma 4.34. *Let $(\mathbb{T}, r), (\mathbb{T}', r')$ be tree structures of height k. Let $t \in T, t' \in T'$ be such that $t \sim_q t'$. Then the following hold:*

1. *For every set U of descendants of t, with $|U| \leq q$, such that for all $u_0, u_1 \in U$, $u_0 \sim_q u_1$, there is a set U' of descendants of t', with $|U'| = |U|$, such that for all $u' \in U'$, for all $u \in U$, $u' \sim_q u$.*

2. *For every infinite set X of descendants of t such that for all $x_0, x_1 \in X$, $x_0 \sim_q x_1$, there exists an infinite set Y of descendants of t' such that for all $y \in Y$, for all $x \in X$, $y \sim_q x$.*

Proof. Let $\mathbb{T}, \mathbb{T}', t, t'$ be as in the statement of the lemma.

1. Let U be as in the statement of the lemma. Then all $u \in U$ have the same height, and are the same distance from t, so it makes sense to talk about the distance $d(t, U)$ of U from t. We argue by induction on $d(t, U)$.

Base case: $d(t, U) = 1$. Then either $U \subseteq R[t]$ or $U \subseteq \check{R}[t]$.

Case 1: $U \subseteq R[t]$. We have $t \simeq_{Tq}^{k-h(t)} t'$, which means there is a winning strategy for Duplicater in the $k - h(t)$-round version of $\mathsf{Bis}_{Tq}(\mathbb{T}, t, \mathbb{T}', t')$. This strategy must prescribe a move for Duplicater in response to Spoiler selecting the set U. Let U' be the move prescribed. Then for every $u' \in U'$, there is some $u \in U$ such that Duplicater has a winning strategy in the $k-(h(t)+1)$-round version of $\mathsf{Bis}_{Tq}(u, u')$. Since $h(u) = h(t) + 1$, the latter statement means that $u \simeq_{Tq}^{k-h(u)} u'$, which implies $u \sim_q u'$. But $\sim_q$ is an equivalence relation, hence $u' \sim_q u$ for *every* $u \in U$. Since u' was arbitrary, this proves the claim.

Case 2: $U \subseteq \check{R}[t]$. Identical argument.

Induction step: $d(t, U) = k + 1$. The inductive hypothesis is that for any $v \in T, v' \in T'$ with $v \sim_q v'$, for any set V of descendants of v with $|V| \leq q$ such that $v_0 \sim_q v_1$ for all $v_0, v_1 \in V$, there is a set V' of descendants of v' with $|V'| = |V|$, such that for all $v' \in V'$, for all $v \in V$, $v' \sim_q v$.

Because all members of U are q-companions, the unique shortest path leading from t to each $u \in U$ must involve the same type of step- forwards, or backwards- at each stage. In particular, the paths must all include either a forward step from t, or a backward step from t. WLOG assume it is a forward step- the backward case is analogous. Define S_U to be the set

$$S_U := \{v \in R[t] : v \text{ has a descendant in } U\}$$

Then $|S_U| \leq |U| \leq q$, because all members of S_U have the same height. Further, all members of S_U are q-companions, because they are ancestors of q-companions. We may now apply the claim of the Base Case to the set S_U, to find a set $S' \subseteq T'$ such that $|S_U| = |S'|$ and for all $s' \in S'$, for all $s \in S_U$, $s' \sim_q s$.

Let $g : s \mapsto s'$ be some bijective mapping from S_U to S'. For each $s \in S_U$, let $U_s \subseteq U$ be the set of descendants of s which are in U. For each pair s, s' along with the set U_s, we may now apply the inductive hypothesis to obtain a set $U_{s'}$

with $|U_{s'}| = |U_s|$, such that for all $u' \in U_{s'}$, for all $u \in U_s$, $u' \sim_q u$. Further, note that the sets $\{U_{s'} : s' \in S'\}$ are pairwise disjoint. Set

$$U' := \bigcup_{s' \in S'} U_{s'}$$

Then U' is a set of descendants of t', such that $|U'| = |U|$ and for all $u' \in U'$, for all $u \in U$, $u' \sim_q u$. This proves the induction step, so the induction is complete.

2. Let X be an infinite set of descendants of t such that $u \sim_q v$ for all $u, v \in X$. Then all members of X are of the same height, and are the same distance from t. Let $d(t, X)$ denote this common distance. We argue by induction on $d(t, X)$.

Base case: $d(t, X) = 1$. Then either tRx for all $x \in X$ or xRt for all $x \in X$. WLOG assume $X \subseteq R[t]$; the other case is analogous.

We have $t \simeq_{Tq}^{k-h(t)} t'$, meaning Duplicater has a winning strategy in $\mathsf{Bis}_{Tq}^{k-h(t)}(\mathbb{T}, t, \mathbb{T}', t')$. Since $X \subseteq R[t]$, the set X is a legal move for Spoiler in this game. Let Y be a winning response for Duplicater. Then for every $y \in Y$ there is $x \in X$ with $y \simeq_{Tq}^{k-(h(t)+1)} x$, which implies $y \sim_q x$. But $\sim_q$ is an equivalence relation, hence $y \sim_q x$ for *every* $x \in X$. Since $y \in Y$ was arbitrary, this proves the claim in the base case.

Induction step: $d(t, X) = k + 1$. The inductive hypothesis states that item 2 of the lemma holds for any $t \in \mathbb{T}, t' \in \mathbb{T}'$ when $d(t, X) = k$. We now make a case distinction.

Case 1: There is a 1-step descendant u of t such that there is an infinite set Y of descendants of u such that $Y \subseteq X$. WLOG assume tRu; the case where uRt is analogous. By item 1 of the lemma, there is a descendant u' of t' such that $u \sim_q u'$. (Indeed, this u' must be an R-successor of t'.) Further, we have $d(u, X) = k$. So by the inductive hypothesis, there exists an infinite set Y' of descendants of u' such that for every $y' \in Y'$, for every $y \in Y$, $y' \sim_q y$. Since $\sim_q$ is an equivalence relation this implies that for every $y' \in Y'$, for every $x \in X$, $y' \sim_q x$. All descendants of u' are descendants of t', so this shows the claim.

Case 2: There are infinitely many 1-step descendants of t, all of which have some descendant $x \in X$. Let U be the set of these 1-step descendants. Either $U \subseteq R[t]$ or $U \subseteq \check{R}[t]$. WLOG assume $U \subseteq R[t]$- the other case is analogous. Also WLOG, we may assume that for all $u, v \in U, u \sim_q v$ (if not, we could find some infinite $U_0 \subseteq U$ whose members are all in the same q-companionship class, because $\sim_q$ has finite index).

We have $t \simeq_{Tq}^{k-h(t)} t'$, so there is a winning strategy for Duplicater in $\mathsf{Bis}_{Tq}^{k-h(t)}(\mathbb{T}, t, \mathbb{T}', t')$. Since $U \subseteq R[t]$, the set U is a legal move for Spoiler in this game. Let V be a winning response to this move for Duplicater. Then for each $v \in V$ there is some $u \in U$ such that $v \simeq_{Tq}^{k-(h(t)+1)} u$. We then have $v \sim_q u$, and since $\sim_q$ is an equivalence relation, $v \sim_q u$ for *every* $u \in U$. But v was arbitrary, so for all $v \in V$, for all $u \in U$, $v \sim_q u$.

Choose some particular $u \in U$, and let x_u be some descendant of u that is in X. For each $v \in V$, we may now invoke item 1 of the lemma for the pair u, v, giving a descendant y_v of v such that $y_v \sim_q x_u$ (implying $y_v \sim_q x$ for every $x \in X$). For any distinct $v, v' \in V$, v is neither an ancestor nor a descendant of v', hence v, v' have no common descendants, and $y_v \neq y_{v'}$. It follows that the set $\{y_v : v \in V\}$ is infinite. This shows the claim.

Finally, a simple pigeonhole argument shows that either case 1 or case 2 must obtain. This concludes the induction, and the proof of item 2. □

Definition 4.35. Let $\mathbb{T}, r$ be a two-way tree, and let $u_1, \ldots, u_m$ be an m-tuple within T. We let $\mathbb{T}_{\overline{u}}$ denote the restriction of $\mathbb{T}$ to the set

$$\bigcup_{1 \leq i \leq m} \downarrow(u_i)$$

Definition 4.36. Let $\mathbb{T}, r$ and $\mathbb{T}', r'$ be two-way trees, and let $\overline{u}, \overline{u}'$ be m-tuples within the respective structures. We say that $\overline{u}, \overline{u}'$ are *q-similar*, notation: $\mathbb{T}, \overline{u} \approx_q^{\downarrow} \mathbb{T}', \overline{u}'$, iff there is an isomorphism $f : \mathbb{T}_{\overline{u}} \to \mathbb{T}'_{\overline{u}'}$ such that

- $f(u_i) = u_i'$ for each $1 \leq i \leq m$
- $x \sim_q f(x)$ for all x in $\mathbb{T}_{\overline{u}}$

The following lemma gives the central step in the proof of Lemma 4.31.

Lemma 4.37. *Let $\mathbb{T}, r$ and $\mathbb{T}', r'$ be two-way trees of height k. Let m be a natural number with $m < q$. Let $\overline{u}, \overline{u}'$ be m-tuples within T, T' respectively, such that $\overline{u} \approx_q^{\downarrow} \overline{u}'$. Then the following hold.*

1. *For every $u \in \mathbb{T}$ there is a $u' \in \mathbb{T}'$ such that $(\overline{u}, u) \approx_q^{\downarrow} (\overline{u}', u')$ (and vice versa for $\mathbb{T}', \mathbb{T}$)*
2. *For every infinite $X \subseteq \mathbb{T}$ there exists an infinite $Y \subseteq \mathbb{T}'$ such that for every $y \in Y$, there exists an $x \in X$ with $\overline{u}, x \approx_q^{\downarrow} \overline{u}', y$ (and vice versa).*

Proof. We begin with item 1. We only show the direction from $\mathbb{T}$ to $\mathbb{T}'$; the other direction is analogous. Let $u \in \mathbb{T}$. If $u \in \mathbb{T}_{\overline{u}}$ then $f(u)$ gives the desired element in $\mathbb{T}'$, so assume $u \notin \mathbb{T}_{\overline{u}}$. Define

$$U := \{v \in \overline{u} : v \sim_q u\}$$

Because $\mathbb{T}$ is a tree, there is a unique last common ancestor t of the set $\{u\} \cup U$. Clearly $t \in \mathbb{T}_{\overline{u}}$; let t' denote $f(t)$. Further, the set $\{u\} \cup U$ has cardinality at most $m+1 \leq q$. So by item 1 of Lemma 4.34, there is a set U' of descendants of t', with $|U'| = |\{u\} \cup U|$, whose members are all q-companions of u. However, there can only be m many q-companions of u in the range of f, since there are at most m many q-companions of u in the domain of f. Hence there is some $u' \in U'$ which is not in the range of f.

Thus, the map f can be extended to an isomorphism $f^+ : \mathbb{T}_{\overline{u},u} \to \mathbb{T}'_{\overline{u}',u'}$ linking u with u', and with $x \sim_q f^+(x)$ for all x in $\mathbb{T}_{\overline{u},u}$. (Here we use the fact that if $t \sim_q t'$, then the paths from r, r' to t, t', respectively, must have the same type of step- forwards/backwards- at each stage). So $(\overline{u}, u)$ and $(\overline{u}', u')$ are q-similar, as required.

We now turn to item 2. Again, we only prove the direction from $\mathbb{T}$ to $\mathbb{T}'$; the direction from $\mathbb{T}'$ to $\mathbb{T}$ is analogous. Let $X \subseteq \mathbb{T}$ be infinite. WLOG assume $X \cap \overline{u} = \emptyset$ (otherwise, Spoiler could have chosen a smaller infinite set). Because $\sim_q$ has finite index, there must be some infinite subset Y of X such that for all $y_0, y_1 \in Y, y_0 \sim_q y_1$. So WLOG we may assume that for all $x_0, x_1 \in X$, $x_0 \sim_q x_1$. Define

$$U := \{v \in \overline{u} : v \sim_q x \text{ for all } x \in X\}$$

Since $\mathbb{T}$ is a tree, there must be some last common ancestor of $X \cup U$- call it t. Clearly $t \in \mathbb{T}_{\overline{u}}$; let $t' := f(t)$. By item 2 of Lemma 4.34, there is an infinite set X' of descendants of t' whose members are all in the same q-companionship class as those of X. Since $\overline{u}'$ is finite, there must even be such an infinite set X' which is *disjoint* from $\overline{u}'$.

For any $x' \in X'$, we may in fact take any $x \in X$. We will then have $x \sim_q x'$. It is easy to see that f may be extended to an isomorphism $f^+ : \mathbb{T}_{\overline{u},x} \to \mathbb{T}'_{\overline{u}',x'}$ such that $f(x) = x'$, and $y \sim_q f^+(y)$ for all $y \in \mathbb{T}_{\overline{u},x}$. (Again, we use the restriction on the paths leading from the respective roots to q-companions.) So $(\overline{u}, x)$ and $(\overline{u}', x')$ are q-similar, as required. $\square$

Proof of Lemma 4.31

Let $(\mathbb{T}, r), (\mathbb{T}', r')$ be two-way trees of height k, and suppose $\mathbb{T}, r \simeq_{Tq}^k \mathbb{T}', r'$. It follows that $r \sim_q r'$. It suffices to show that Duplicater has a winning strategy in $EF_q^\infty(\mathbb{T}, r, \mathbb{T}', r')$. We show that there is a strategy which preserves the property that, for each configuration $\overline{u}, \overline{u}'$, $\overline{u}$ is q-similar to $\overline{u}'$. This condition is certainly satisfied in the initial configuration of r, r', because $r \sim_q r'$.

After m rounds (for any $m < q$), at configuration $\overline{u} \approx_q^{\downarrow} \overline{u}'$, suppose Spoiler launches a first-order challenge and selects a single element from T or from T'. Then item 1 of Lemma 4.37 shows that Duplicater can select an element from the other structure, which preserves the desired configuration-property. Alternatively, suppose Spoiler launches a second-order challenge, and selects an infinite set X from T or T'. Then item 2 of Lemma 4.37 shows that Duplicater can respond to this challenge in a way that also preserves the desired configuration-property.

It is easy to see that a strategy which preserves this configuration-property must be winning for Duplicater. □

As discussed, Lemma 4.30 and Lemma 4.31 suffice to prove Proposition 4.29. So we may conclude that Proposition 4.29 holds. □

Proposition 4.38. *Let n be any natural number. Modulo $\simeq_{GT}$, $\equiv_{1,\exists}^{k,q}$ can be upgraded to $\equiv_n^{k,q}$.*

Proof. Let $\mathfrak{A}, \mathfrak{B}$ be structures such that $\mathfrak{A} \equiv_{1,\exists}^{k,q} \mathfrak{B}$. We claim that $\mathfrak{A} \cdot \omega, \mathfrak{B} \cdot \omega$ are the desired companion structures. We already have $\mathfrak{A} \cdot \omega \simeq_{GT} \mathfrak{A}$ and $\mathfrak{B} \cdot \omega \simeq_{GT} \mathfrak{B}$ by Proposition 4.23. So we only need to show that $\mathfrak{A} \cdot \omega \equiv_n^{k,q} \mathfrak{B} \cdot \omega$. This requires us to verify:

1. For all k-local formulas ψ such that $qd(\psi) \leq q$, for all $m \leq n$, $\mathfrak{A} \cdot \omega \models Gaif_{k,m}(\psi)$ iff $\mathfrak{B} \cdot \omega \models Gaif_{k,m}(\psi)$
2. For all k-local formulas ψ such that $qd(\psi) \leq q$, for all $m \leq n$, $\mathfrak{A} \cdot \omega \models Gaif_{k,m}^{\infty}(\psi)$ iff $\mathfrak{B} \cdot \omega \models Gaif_{k,m}^{\infty}(\psi)$

To see that item 1 holds, let ψ be k-local with $qd(\psi) \leq q$, let $m \leq n$, and suppose $\mathfrak{A} \cdot \omega \models \exists x_1 \ldots \exists x_m (\bigwedge_{1 \leq i < j \leq m} d(x_i, x_j) > k \wedge \psi(x_i))$. Then we certainly have $\mathfrak{A} \cdot \omega \models \exists x \psi$. Let (c, i) be a witness; then by locality of ψ, we also have $\mathfrak{A}, c \models \psi$, and hence $\mathfrak{A} \models \exists x \psi$. Since $\mathfrak{A} \equiv_{1,\exists}^{k,q} \mathfrak{B}$, we have $\mathfrak{B} \models \exists x \psi$. Let $b \in \mathfrak{B}$ satisfy the formula ψ. Consider the m elements $(b, 0), (b, 1), \ldots, (b, m-1)$ of $\mathfrak{B} \cdot \omega$. These elements are suitably scattered, and they all satisfy the formula ψ by locality of ψ. So $\mathfrak{B} \cdot \omega$ satisfies the desired basic local sentence.

The direction from $\mathfrak{B} \cdot \omega$ to $\mathfrak{A} \cdot \omega$ proceeds analogously.

To see that item 2 holds, let ψ be k-local with $qd(\psi) \leq q$, let $m \leq n$, and suppose $\mathfrak{A} \cdot \omega \models \exists^{\infty} x_1 \ldots \exists^{\infty} x_m (\bigwedge_{1 \leq i < j \leq m} d(x_i, x_j) > k \wedge \psi(x_i))$. Then clearly $\mathfrak{A} \cdot \omega \models$

$\exists x\psi(x)$. Let (a,i) be a witness for this. Then by locality of ψ, we also have $\mathfrak{A}, a \models \psi$. Since $\mathfrak{A} \equiv_{1\exists}^{k,q} \mathfrak{B}$, we then have $\mathfrak{B} \models \exists x\psi(x)$. Let $b \in \mathfrak{B}$ be an element satisfying ψ. For $1 \leq i \leq m$, we now let S_i be the following infinite set:

$$S_i := \{(b, pm + i) : p \in \mathbb{N}\}$$

Then for each $1 \leq i \leq m$ we have $S_i \subseteq B \times \omega$, and further, if $i \neq j$ then $i' \neq j'$ for any $(b,i') \in S_i, (b,j') \in S_j$. In other words, all members of S_i are suitably scattered from all members of S_j. Because ψ is local, every member of S_i will satisfy ψ (for any i). We may therefore take the sets S_i as witness sets for the quantifications $\exists^\infty x_i$, respectively, showing that $\mathfrak{B} \cdot \omega \models \exists^\infty x_1 \dots \exists^\infty x_m(\bigwedge_{1 \leq i < j \leq m} d(x_i, x_j) > k \wedge \psi(x_i))$.

The direction from $\mathfrak{B} \cdot \omega$ to $\mathfrak{A} \cdot \omega$ proceeds analogously. □

Proof of Proposition 4.10 and Proposition 4.11. We have shown that for any q, $\simeq_{GT}^k$ can be upgraded modulo $\simeq_{GT}$ to $\equiv_{1\exists}^{k,q}$, and for any n, $\equiv_{1\exists}^{k,q}$ can be upgraded modulo $\simeq_{GT}$ to $\equiv_n^{k,q}$. Putting these two results together we obtain that for any q, n, $\simeq_{GT}^k$ can be upgraded modulo $\simeq_{GT}$ to $\equiv_n^{k,q}$, i.e. Proposition 4.11. We also showed that $\simeq_G^{2k}$ can be upgraded modulo $\simeq_G$ to $\simeq_{GT}^k$ (Proposition 4.12). Since $\simeq_{GT}$ is stronger than $\simeq_G$, being able to upgrade modulo $\simeq_{GT}$ implies being able to upgrade modulo $\simeq_G$. So Proposition 4.12 and Proposition 4.11 together yield Proposition 4.10. (See also Figure 4.4 on the following page.)

4.3.3 Characterisation Theorems

Using these Upgrading results, we can provide a semantic characterisation of G^∞ and GT^∞.

Theorem 4.39. *Let ϕ be a sentence of FOL^∞. Then ϕ is invariant under $\simeq_G$ iff ϕ is logically equivalent to a formula in G^∞.*

Proof. We only need to prove the left-to-right direction. Let $\phi \in FOL^\infty$ be a $\simeq_G$-invariant sentence. By the Gaifman theorem for FOL^∞, there are some $k, q, n \in \mathbb{N}$ such that ϕ is invariant over $\equiv_n^{k,q}$. Now let $\mathcal{A}, \mathcal{B}$ be arbitrary structures such that $\mathcal{A} \simeq_G^{2k} \mathcal{B}$. Then by Proposition 4.10, there are companion structures $\mathcal{A}' \simeq_G \mathcal{A}$, $\mathcal{B}' \simeq_G \mathcal{B}$ such that $\mathcal{A}' \equiv_n^{k,q} \mathcal{B}'$. Now the chain of equivalences

$$\begin{aligned} \mathcal{A} \models \phi \iff & \mathcal{A}' \models \phi && \text{(invariance under } \simeq_G \text{)} \\ \iff & \mathcal{B}' \models \phi && \text{(invariance under } \equiv_n^{k,q} \text{)} \\ \iff & \mathcal{B} \models \phi && \text{(invariance under } \simeq_G \text{)} \end{aligned}$$

shows that in fact, ϕ is invariant under $\simeq_G^{2k}$. But then a straightforward argument, analogous to the proof of Theorem 4.7, shows that ϕ is definable in G^∞. $\square$

Theorem 4.40. *Let $\phi \in FOL^\infty$. Then ϕ is invariant under $\simeq_{GT}$ iff ϕ is logically equivalent to a formula in GT^∞.*

Proof. We only need to prove the left-to-right direction. Let $\phi \in FOL^\infty$ be a $\simeq_{GT}$-invariant sentence. By the Gaifman theorem for FOL^∞, there are some $k, q, n \in \mathbb{N}$ such that ϕ is invariant over $\equiv_n^{k,q}$. Now let $\mathcal{A}, \mathcal{B}$ be arbitrary structures such that $\mathcal{A} \simeq_{GT}^k \mathcal{B}$. Then by Proposition 4.11, there are companion structures $\mathcal{A}' \simeq_{GT} \mathcal{A}$, $\mathcal{B}' \simeq_{GT} \mathcal{B}$ such that $\mathcal{A}' \equiv_n^{k,q} \mathcal{B}'$. Now the chain of equivalences

$$\begin{aligned} \mathcal{A} \models \phi &\iff \mathcal{A}' \models \phi && \text{(invariance under } \simeq_{GT}\text{)} \\ &\iff \mathcal{B}' \models \phi && \text{(invariance under } \equiv_n^{k,q}\text{)} \\ &\iff \mathcal{B} \models \phi && \text{(invariance under } \simeq_{GT}\text{)} \end{aligned}$$

shows that in fact, ϕ is invariant under $\simeq_{GT}^k$. But then a straightforward argument, analogous to the proof of Theorem 4.7, shows that ϕ is definable in GT^∞. $\square$

The figure below shows a diagrammatic representation of the ‘upgrading’ component of this proof:

$$\begin{tikzcd} \mathcal{M}_0 \arrow[r, leftrightarrow, "\simeq_G^{2k}"] \arrow[d, "\simeq_G"'] & \mathcal{M}_1 \arrow[d, "\simeq_G"] \\ \mathcal{M}_0' \arrow[r, leftrightarrow, "\simeq_{GT}^k"] \arrow[d, "\simeq_{GT}"'] & \mathcal{M}_1' \arrow[d, "\simeq_{GT}"] \\ \mathcal{M}_0'' \arrow[r, leftrightarrow, "\equiv_{1,\exists}^{k,q}"'] \arrow[d, "\simeq_{GT}"'] & \mathcal{M}_1'' \arrow[d, "\simeq_{GT}"] \\ \mathcal{M}_0''' \arrow[r, leftrightarrow, "\equiv_n^{k,q}"'] & \mathcal{M}_1''' \end{tikzcd}$$

Figure 4.4: Upgrading from $\simeq_G^{2k}$ to $\equiv_n^{k,q}$ via $\simeq_{GT}^k$

The key argument is that a formula ϕ which is invariant across the vertical arrows, and also across the lowest horizontal arrow, must be invariant over all the horizontal arrows.

Chapter 5

A Guarded Fragment of FOL^{∞}

In this chapter, we introduce a fragment of FOL^{∞} inspired by the guarded fragment of first-order logic. The guarded fragment of first-order logic (which I'll refer to as GF) was introduced by Andréka, van Benthem and Németi in [3]. It is defined by means of a syntactic restriction on the use of quantifiers: only quantified formulas that fit the schema $\exists \overline{y}(P(\overline{x}, \overline{y}) \wedge \psi(\overline{x}, \overline{y}))$, where P is a predicate symbol, are allowed. In their paper, Andréka, van Benthem and Németi demonstrate that the guarded fragment shares several important properties with basic modal logic. Firstly, it is characterised as a fragment of FOL by invariance under a suitable notion of bisimulation. Secondly, and even more significantly, it is decidable. In the remainder of this thesis, we ask if we can establish these nice properties for a similarly strong fragment of FOL^{∞}.

The first section of this chapter will define a fragment of FOL^{∞}, which employs a syntactic restriction similar to that of the guarded fragment of FOL. We will then define a suitable notion of bisimulation for this fragment, show that the fragment is invariant under this bisimilarity relation, and consider whether it is characterised by this invariance. We consider the decidability of the newly defined fragment in the following chapter.

5.1 Syntax of GF^{∞}

Definition 5.1. Let τ be a signature which contains no constant or function symbols, but may contain arbitrarily many predicate symbols of any arity. We define $\mathfrak{Form}^{\infty}_{GF}(\tau)$ to be the least set satisfying the following recursive clauses:

- If P is an n-ary predicate symbol of τ, or the identity symbol $=$, and $v_1, \ldots, v_n$ are variables, then $Pv_1 \ldots v_n \in \mathfrak{Form}^{\infty}_{GF}(\tau)$ and is an atomic formula

- $\bot \in \mathfrak{Form}^\infty_{GF}(\tau)$ and is an atomic formula
- If $\phi, \psi \in \mathfrak{Form}^\infty_{GF}(\tau)$ then $\phi \wedge \psi \in \mathfrak{Form}^\infty_{GF}$ and $\neg\phi \in \mathfrak{Form}^\infty_{GF}$
- Let each of $Q_1, \dots, Q_n \in \{\exists, \exists^\infty\}$. Let ϕ be a formula of $\mathfrak{Form}^\infty_{GF}(\tau)$, and α an atomic formula of $\mathfrak{Form}^\infty_{GF}(\tau)$, such that $FV(\phi) \subseteq FV(\alpha)$. Then $Q_1 v_1 \dots Q_n v_n(\alpha \wedge \phi) \in \mathfrak{Form}^\infty_{GF}(\tau)$. We call α the 'guard' of such a formula

We write $\phi \in GF^\infty$ iff there is some (purely relational) signature τ such that $\phi \in \mathfrak{Form}^\infty_{GF}(\tau)$.

The guarded fragment of *FOL* allows arbitrary strings of existential quantifiers before a guard, and we follow this idea in allowing arbitrary strings of existential *or infinity* quantifiers before a guard in GF^∞. But note that this adds to the expressive power of GF^∞ in a way that is disanalogous to GF. For instance, a GF-formula $\exists x_1 \dots \exists x_n(\alpha \wedge \phi)$ would often be abbreviated to $\exists \overline{x}(\alpha \wedge \phi)$, since the order of the quantifiers does not matter; the original formula is equivalent to, e.g., $\exists x_n \dots \exists x_1(\alpha \wedge \phi)$. This makes it tempting to think of the string of existential quantifiers as a single existential quantifier, binding a finite tuple of variables.

We cannot do this for GF^∞, because if a string of quantifiers includes infinity quantifiers, the order of the quantifiers emphatically *does* matter. For example, $\exists x_1 \exists^\infty x_2 \exists^\infty x_3(Px_1x_2x_3 \wedge \top)$ is not equivalent to $\exists x_1 \exists^\infty x_3 \exists^\infty x_2(Px_1x_2x_3 \wedge \top)$. To see this, consider the structure below.

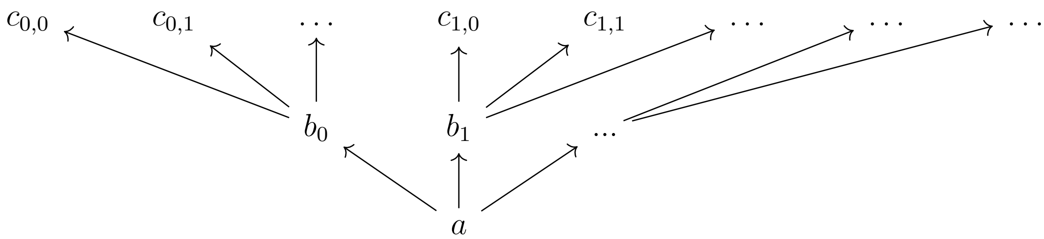

Figure 5.1: $\mathfrak{M}$

Let P be interpreted in $\mathfrak{M}$ by: $(x, y, z) \in P^{\mathfrak{M}}$ iff there is an arrow from x to y, and an arrow from y to z. Then $\mathfrak{M} \models \exists x_1 \exists^\infty x_2 \exists^\infty x_3(Px_1x_2x_3 \wedge \top)$, because a is such that there are infinitely many b's, for each of which there are infinitely many c's, such that $P^{\mathfrak{M}}(a, b, c)$. But $\exists x_1 \exists^\infty x_3 \exists^\infty x_2(Px_1x_2x_3 \wedge \top)$ does not hold in $\mathfrak{M}$, intuitively because, for each c, there is only one b such that $P^{\mathfrak{M}}abc$.

So, it is important to note that allowing strings of quantifiers is significantly less innocent in the GF^∞ case than in the GF case.

Note that nothing needs to be said about the semantics of GF^∞, because it is a fragment of FOL^∞. So, the satisfaction conditions for any formula of GF^∞ are exactly as described by the definition of satisfaction for FOL^∞.

5.2 Bisimulation

Definition 5.2 (Guarded set)**.** Let $\mathfrak{M}$ be a τ-structure, for some relational signature τ. We say that a subset $X \subseteq M$ is *guarded* (in $\mathfrak{M}$) iff there is some predicate symbol $P \in \tau$, and some ordered tuple $(d_1, \ldots, d_n) \in P^{\mathfrak{M}}$, such that every member of X occurs somewhere in $(d_1, \ldots, d_n)$.

Notation. Where π is an ordered n-tuple, and $i \leq n$, I'll use $(\pi)_i$ to denote the ith element of π.

We now define a bisimilarity relation that matches the expressive power of GF^∞. Again, we take a game-theoretic approach.

Definition 5.3 (Guarded infinity bisimulation game)**.** Bis_{GF} is a game played by two players, Spoiler and Duplicater, on two structures, $\mathfrak{M}_0$ and $\mathfrak{M}_1$. The game is played in a series of rounds. Each round has a *configuration*, which consists of finite sets $C_0 \subseteq M_0, C_1 \subseteq M_1$, together with a map $g : C_0 \to C_1$.

In each round, Spoiler moves first and selects one of the structures, $\mathfrak{M}_i$. He then selects a subset $B_i \subseteq C_i$, selects an $n \in \mathbb{N}$, selects an ordered $n+1$-tuple $(a_0, \ldots, a_n)$ with each $a_i \in \{1, \infty\}$, and specifies a map $f_i : \mathbb{F}(a_0, \ldots, a_n) \to M_i$ such that

1. For any $\sim_{\mathbb{F}}$ equivalence class $[t] \subseteq \mathbb{F}(a_0, \ldots, a_n)$, the map f_i is injective on $[t]$
2. For each path π through $\mathbb{F}(a_0, \ldots, a_n)$, the set $B_i \cup f_i[\pi]$ is guarded

In response, Duplicater must specify a map $f_{1-i} : \mathbb{F}(a_0, \ldots, a_n) \to M_{1-i}$, which also satisfies condition 1.

Next, Spoiler selects any path π_{1-i} through $\mathbb{F}(a_0, \ldots, a_n)$. In response, Duplicater must select a path π_i (which may but need not be distinct from π_{1-i}) through $\mathbb{F}(a_0, \ldots, a_n)$. Let B_{1-i} denote the elements of C_{1-i} linked with an element of B_i by g. Then the new configuration (C'_0, C'_1, g') is given by:

$$\begin{aligned}
&C'_0 = B_0 \cup f_0[\pi_0] \\
&C'_1 = B_1 \cup f_1[\pi_1] \\
&g'(b) = g(b) \text{ for any } b \in B_0, \text{ and } g'(f_0((\pi_0)_m)) = f_1((\pi_1)_m)
\end{aligned}$$

(If g' is not well-defined, then Duplicater loses the game).

Spoiler wins the game if the players ever reach a configuration (C_0, C_1, g) such that g is not a bijective local isomorphism. (So if the starting configuration is such that g fails to be a local isomorphism, then Spoiler wins instantly.) Otherwise, Duplicater wins.

We write $\mathsf{Bis}_{GF}(\mathfrak{M}_0, \mathfrak{M}_1, C_0, C_1, g)$ to denote the game played on $\mathfrak{M}_0, \mathfrak{M}_1$ with starting configuration (C_0, C_1, g). Where $\overline{m}_0, \overline{m}_1$ are ordered tuples and g is the map which sends each $(\overline{m}_0)_i$ to $(\overline{m}_1)_i$, we write $\mathfrak{M}_0, \overline{m}_0 \simeq_{GF} \mathfrak{M}_1, \overline{m}_1$ iff Duplicater has a winning strategy in $\mathsf{Bis}_{GF}(\mathfrak{M}_0, \mathfrak{M}_1, \overline{m}_0, \overline{m}_1, g)$. We write $\mathfrak{M}_0, \overline{m}_0 \simeq^k_{GF} \mathfrak{M}_1, \overline{m}_1$ iff Duplicater has a winning strategy in the corresponding k-round version of this game.

Naturally, we want our fragment GF^∞ to be invariant over this relation.

Proposition 5.4 (Bisimulation Invariance for GF^∞)**.** *Let $\mathfrak{M}, \mathfrak{N}$ be τ- structures for some relational structure τ, let $m_1, \ldots, m_k$ be a k-tuple of elements from $\mathfrak{M}$, and $n_1, \ldots, n_k$ a k-tuple of elements from $\mathfrak{N}$. Suppose $\mathfrak{M}, \overline{m} \simeq_{GF} \mathfrak{N}, \overline{n}$. Then for any $\phi \in \mathfrak{Form}^\infty_{GF}(\tau)$, $\mathfrak{M}, \overline{m} \models \phi$ iff $\mathfrak{N}, \overline{n} \models \phi$.*

Proof. We proceed by induction on the construction of ϕ.

Base case: ϕ is atomic. This is trivial if $\phi = \bot$. If $\phi = P\overline{x}$ for some predicate P, then if e.g. $m_1, \ldots, m_k$ satisfied ϕ but $n_1, \ldots, n_k$ did not satisfy ϕ, then $\overline{m}, \overline{n}$ would not be locally isomorphic, and Duplicater would lose the bisimulation game instantly.

The induction step for each of the Boolean connectives is routine, so we omit it.

Induction step: $\phi = Q_0 x_0 \ldots Q_n x_n(\alpha \wedge \psi)$. The induction hypothesis is that for any ordered tuples $\overline{m}' \subseteq M$ and $\overline{n}' \subseteq N$, if $\mathfrak{M}, \overline{m}' \simeq_{GF} \mathfrak{N}, \overline{n}'$ then $\mathfrak{M}, \overline{m}' \models \psi$ iff $\mathfrak{N}, \overline{n}' \models \psi$.

Suppose $\mathfrak{M}, \overline{m} \models Q_0 x_0 \ldots Q_n x_n(\alpha \wedge \psi)$. Let $(a_0, \ldots, a_n)$ be the index sequence for $Q_0, \ldots, Q_n$. Then by Observation 2.14, there is a function $f : \mathbb{F}(a_0, \ldots, a_n)$ which is injective on $\sim_{\mathbb{F}}$ equivalence classes, and is such that for every path π through $\mathbb{F}(a_0, \ldots, a_n)$, we have

$$\mathfrak{M}, \overline{m}, f((\pi)_0), \ldots, f((\pi)_n) \models \alpha \wedge \psi$$

Define

$$\mathsf{m} := \{m_i \in \overline{m} : \text{ some variable occurring in } \alpha \text{ corresponds to } m_i\}$$

Also, let $\mathsf{n} := \{n_i : m_i \in \mathsf{m}\}$. WLOG we may assume that each variable x_i appears in α. Because α is atomic, it follows that the set $\mathsf{m} \cup \{f((\pi)_0), \dots, f((\pi)_n)\}$ is guarded. This means that the set m, together with the function f, would be a legal move for Spoiler in $\mathsf{Bis}_{GF}(\mathfrak{M}, \overline{m}, \mathfrak{N}, \overline{n})$. Duplicater has a winning strategy in this game, so let $f_N : \mathbb{F}(a_0, \dots, a_n) \to \mathfrak{N}$ be the response for Duplicater prescribed by this strategy. Because the strategy is winning, we have that for any path π through the forest $\mathbb{F}(a_0, \dots, a_n)$, there is a path π' through $\mathbb{F}(a_0, \dots, a_n)$ such that

$$\mathfrak{M}, \mathsf{m}, f((\pi')_0), \dots, f((\pi')_n) \simeq_{GF} \mathfrak{N}, \mathsf{n}, f_N((\pi)_0), \dots, f_N((\pi)_n)$$

But for each such path π', we know that $\mathfrak{M}, \mathsf{m}, f((\pi')_0), \dots, f((\pi')_n) \models \alpha \wedge \psi$. Therefore by the Induction Hypothesis, for each path π,

$$\mathfrak{N}, \mathsf{n}, f_N((\pi)_0), \dots, f_N((\pi)_n) \models \alpha \wedge \psi$$

But this means that the map f_N satisfies the conditions given in Observation 2.14. So we can conclude that $\mathfrak{N}, \overline{n} \models Q_0 x_0 \dots Q_n x_n (\alpha \wedge \psi)$, as required.

Precisely the same reasoning, in reverse, shows that if $\mathfrak{N}, \overline{n} \models \phi$ then $\mathfrak{M}, \overline{m} \models \phi$.

This concludes the induction step for quantified formulae. So the induction is complete, and the proposition has been shown. $\square$

5.3 Characterisation of GF^∞

We can now ask of GF^∞ the same question as we asked of the other fragments of FOL^∞ considered in this thesis: is GF^∞ characterised by invariance under its associated notion of bisimulation? That is, is there any formula $\phi \in FOL^\infty$ that is invariant under $\simeq_{GF}$, but is not equivalent to any formula in GF^∞?

Just as in the case of the other fragments, we cannot use a compactness argument to answer this question, because FOL^∞ is not compact. So, the only other strategy which seems available is to use the 'upgrading' method, which we successfully applied to the other fragments. For this, we would need an upgrading result, stating that if $\mathfrak{M}, \overline{m} \simeq^k_{GF} \mathfrak{N}, \overline{n}$, then there are fully bisimilar companion structures $\mathfrak{M}', \overline{m}' \simeq_{GF} \mathfrak{M}, \overline{m}$ and $\mathfrak{N}', \overline{n}' \simeq_{GF} \mathfrak{N}, \overline{n}$ such that $\mathfrak{M}', \overline{m}' \equiv^k_{q,n} \mathfrak{N}', \overline{n}'$.

Unfortunately, proving this upgrading result seems to be much more difficult than in any of the cases of T^∞, G^∞ or GT^∞. In proving the previous upgrading results, the crucial step was to use unravelling to obtain structures whose Gaifman graphs were acyclic, and which were bisimilar (in the relevant sense) to the original structures. We made vital use of the fact that every structure was bisimilar (in the relevant sense) to some acyclic structure. (In the remainder of this chapter, we call a structure 'acyclic' iff its Gaifman graph is acyclic.)

In the GF^∞ setting, this is simply not the case: there are some structures for which we cannot find acyclic bisimilar companions. This happens for at least two reasons.

Reason 1. Consider the structure $\mathfrak{N}_1$ shown in Figure 5.2 below. Let P be interpreted as the set $\{a, b, c\}$ in $\mathfrak{N}_1$, and R as the set of pairs $\{(a, b), (b, c), (c, a)\}$ in $\mathfrak{N}_1$.

$$\begin{tikzcd}
 & & P & & \\
 & & c \arrow[dl] & & \\
P & a \arrow[rr] & & b \arrow[ul] & P
\end{tikzcd}$$

Figure 5.2: $\mathfrak{N}_1$

A natural candidate for an acyclic $\simeq_{GF}$-bisimilar companion structure to $\mathfrak{N}_1$ is shown in Figure 5.3 below.

$$\begin{tikzcd}
P & P & P & P & P \\
\dots \arrow[r] & a' \arrow[r] & b' \arrow[r] & c' \arrow[r] & \dots
\end{tikzcd}$$

Figure 5.3: $\mathfrak{N}'_1$

But in fact, $\mathfrak{N}'_1 \simeq_{GF} \mathfrak{N}_1$ does not hold. To see this, simply observe that the structures are not even GF^∞-equivalent: $\mathfrak{N}_1$ does not satisfy $\exists^\infty x(Px \wedge \top)$, while $\mathfrak{N}'_1$ does. Further, this is not merely a problem with the particular choice of $\mathfrak{N}'_1$. The structure $\mathfrak{N}_1$ has an infinite path, in which every point satisfies P. So, any bisimilar companion structure must also have an infinite path, in which every point satisfies P. However, $\mathfrak{N}_1$ only has a finite number of points satisfying P, hence any bisimilar companion must also have only finitely many P-elements. This implies that the infinite path must involve a cycle. So $\mathfrak{N}_1$ cannot have a $\simeq_{GF}$-bisimilar acyclic companion.

A similar problem to this occurs in finite model theory. In this setting, $(\mathfrak{N}_1, a)$ would not even have a classically bisimilar acyclic companion structure, because infinite structures are banned. The best one can achieve is a bisimilar companion structure that is *locally* acyclic, in that it avoids 'short' cycles (all cycles are longer than some fixed k). In many cases, this type of companion structure is enough to prove the desired upgrading result; see [15] for an example of this method. So, it might be hoped that locally acyclic companions can be found in the $\simeq_{GF}$ case as well, and that we can use them to prove the desired upgrading result.

Indeed, for any finite k, we can find bisimilar companions to $\mathfrak{N}_1$ which avoid cycles of length $\leq k$: just take the R-cycle of length $k+1$, with all points satisfying P. However, the next example shows that we cannot do this in general, and this brings us onto:

Reason 2. Consider the structure $\mathfrak{N}_2$ shown below, in which P is interpreted as the set $\{c\}$, and R as the set of pairs $\{(a, b_j) : j \in \mathbb{Z}\} \cup \{(b_j, c) : j \in \mathbb{Z}\}$.

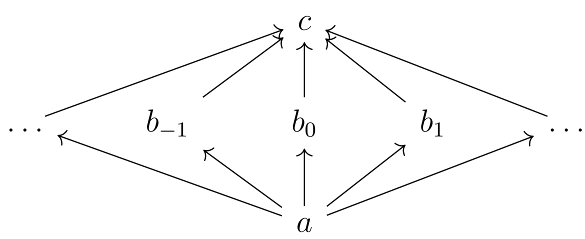

Figure 5.4: $\mathfrak{N}_2$

$\mathfrak{N}_2$ is evidently not acyclic: its Gaifman graph contains the cycle, e.g., (a, b_0, c, b_1). Now let $\mathfrak{N}_2'$ be any structure $\simeq_{GF}$-bisimilar to $\mathfrak{N}_2$. Then by Proposition 5.4, $\mathfrak{N}_2'$ and $\mathfrak{N}_2$ agree on all GF^{∞} formulas. Observe that $\mathfrak{N}_2$ satisfies the formula

$$\phi := \exists x \exists^{\infty} y (Rxy \wedge \exists z (Ryz \wedge Pz)) \wedge \neg \exists^{\infty} z (Pz \wedge \top)$$

So $\mathfrak{N}_2'$ must also satisfy this formula.

This means that $\mathfrak{N}_2'$ has to contain an element, a', which has infinitely many successors b', all of which see a P-element- but $\mathfrak{N}_2'$ can only contain finitely many P-elements. If every b' saw a *distinct* P- element c', then there would be infinitely many P-elements, which is an immediate contradiction! Therefore, there must be at least two R-successors of a', call them b_1' and b_2', which see the *same* P-element c'. This means that $\mathfrak{N}_2'$ cannot be acyclic. Indeed, this argument shows more: it shows that the Gaifman graph of $\mathfrak{N}_2'$ has a cycle of length 4 (namely a', b_0', b_1', c'). So, unlike in the finite model theory case, we cannot even find a bisimilar companion to $\mathfrak{N}_2$ which is locally acyclic.

Harking back to Definition 2.5, we may now finally explain why we did not define G^{∞} to include a global infinity modality as well as a global modality. In a nutshell, adding this modality would have made G^∞ strong enough for the problem discussed above to arise. Letting E^∞ denote the global infinity modality, any model of the formula:

$$E\Diamond^\infty\Diamond p \wedge \neg E^\infty p$$

would fail to have an acyclic bisimilar companion.

Another point to make about the above example is that it is closely connected to the natural language inference, **I**, discussed at the beginning of Chapter 1. Observe that GF^∞ is strong enough to express this inference- indeed, it can express this inference within a vocabulary containing only unary predicate symbols, and a single binary relation symbol. Using the unary predicates 'D' for the property of being a dog and 'C' for the property of being a cat, and the binary predicate 'R' for the relation 'x sees y', we can formalise the inference as:

- **P1**: $\exists^\infty x \exists y(Rxy \wedge Dx \wedge Cy)$
- **P2**: $\neg\exists y \exists^\infty x(Rxy \wedge Dx \wedge Cy)$
- **C**: $\exists^\infty y Cy$

This inference turns on the combinatorial principle that if there are infinitely many dogs which each see a cat, but only finitely many cats, then some cat would have to be seen by infinitely many different dogs. This is closely analogous to the above argument that, in any bisimilar companion $\mathfrak{N}_2'$ to $\mathfrak{N}_2$, there would have to be two distinct b_1', b_2' seeing the same element c'. So, although it is a nice feature of GF^∞ that it is able to express inferences such as **I**, this expressivity may also be viewed as the main obstacle to characterising GF^∞ by invariance under $\simeq_{GF}$.

Of course, none of this shows that GF^∞ is *not* characterised by being the $\simeq_{GF}$-invariant fragment of FOL^∞. It does not even show that the Upgrading statement corresponding to $\simeq_{GF}$ does not hold. It might be that the Upgrading statement can be proved using some more complicated construction for obtaining bisimilar companions, or that the Upgrading statement fails, but the characterisation still holds. However, the unavailability of locally ayclic $\simeq_{GF}$-bisimilar companion structures seems like a serious obstacle to proving the $\simeq_{GF}$-invariance characterisation claim. We therefore leave this question for future work.

Chapter 6

Decidability

The final chapter of this thesis will discuss the decidability of some of the logics considered so far. We begin by giving a complete axiomatisation of T^∞, and showing that this logic is decidable. Our strategy will be to translate formulas of T^∞ into formulas of a polymodal logic, and then to argue that the latter is decidable and the translation preserves satisfiability. We then turn to the fragment GF^∞ introduced in Chapter 5. We define a similar translation, from formulas of GF^∞ into formulas of GF, and we conjecture that this translation also preserves satisfiability.

6.1 Decidability of T^∞

6.1.1 Preliminaries

For the remainder of this chapter we fix a countable set of proposition letters, P.

Definition 6.1 (Syntax of T^2)**.** We extend the basic temporal language $\{\Diamond_\rightarrow, \Diamond_\leftarrow\}$ with two new diamonds: $\Diamond^2_\rightarrow$ and $\Diamond^2_\leftarrow$. We let T^2 denote the set of all modal formulas that can be formed out of the modalities $\{\Diamond_\rightarrow, \Diamond_\leftarrow, \Diamond^2_\rightarrow, \Diamond^2_\leftarrow\}$, according to the standard inductive definitions. We regard $\Box^2_\rightarrow$ and $\Box^2_\leftarrow$ as duals of $\Diamond^2_\rightarrow$ and $\Diamond^2_\leftarrow$, respectively, definable in the usual way.

Definition 6.2 (Semantics of T^2)**.** A T^2 frame $\mathcal{F}$ is an ordered tuple $(S, R, R^2_\rightarrow, R^2_\leftarrow)$, with S a set and each $R, R^2_\rightarrow, R^2_\leftarrow$ a binary relation on S. We say that $\mathcal{F}$ is a $\mathbf{T}^2$ frame iff it satisfies the additional conditions: $R^2_\rightarrow \subseteq R$, $R^2_\leftarrow \subseteq \breve{R}$.

A T^2 model $\mathcal{M}$ is a T^2 frame $\mathcal{F}$, together with a valuation $V : \mathsf{P} \to \wp(S)$.

We can give a natural semantics for T^2 on the class of T^2 models. The recursive clauses for the new modalities are given by:

- $\mathcal{M}, s \Vdash \Diamond^2_\rightarrow \phi$ iff there is some $t \in S$ with $sR^2_\rightarrow t$ and $\mathcal{M}, t \Vdash \phi$

- $\mathcal{M}, s \Vdash \Diamond^2_\leftarrow \phi$ iff there is some $t \in S$ with $sR^2_\leftarrow t$ and $\mathcal{M}, t \Vdash \phi$

One important thing to note about the above definition is that on T^2 models, we do *not* require that $R^2_\leftarrow$ be the converse of $R^2_\rightarrow$. So, we do not necessarily have $\mathcal{M}, s \Vdash \Diamond^2_\leftarrow \phi$ iff there is some $t \in S$ such that $tR^2_\rightarrow s$ and $\mathcal{M}, t \Vdash \phi$.

Definition 6.3 (Translation of T^∞ into T^2)**.** We now give a very simple translation, $\psi \mapsto tr(\psi)$, from T^∞ into T^2.

- $tr(p) = p$ for all proposition letters p
- $tr(\bot) = \bot$
- $tr(\psi \wedge \chi) = tr(\psi) \wedge tr(\chi), tr(\neg\psi) = \neg tr(\psi)$
- $tr(\Diamond_\rightarrow \psi) = \Diamond_\rightarrow tr(\psi), tr(\Diamond_\leftarrow \psi) = \Diamond_\leftarrow tr(\psi)$
- $tr(\Diamond^\infty_\rightarrow \psi) = \Diamond^2_\rightarrow tr(\psi), tr(\Diamond^\infty_\leftarrow) = \Diamond^2_\leftarrow tr(\psi)$

It is easy to see that, syntactically speaking, tr is an isomorphism from T^∞ to T^2. However, T^∞ and T^2 have very different semantics, so it is far from obvious that tr also respects semantic properties such as satisfiability. Our goal will be to show that, indeed, $\phi \in T^\infty$ is satisfiable iff $tr(\phi)$ is satisfiable on a $\mathbf{T}^2$ frame.

6.1.2 Satisfiability: From ϕ to $tr(\phi)$

We begin by showing that if ϕ is satisfiable, then $tr(\phi)$ is satisfiable on a $\mathbf{T}^2$ frame.

Definition 6.4 ($\mathbf{T}^\infty$)**.** We let $\mathsf{Ax}(T^\infty)$ denote the following set of T^∞ formulas:

- All propositional tautologies
- $\Box_\rightarrow(p \to q) \to (\Box_\rightarrow p \to \Box_\rightarrow q)$
- $\Box_\leftarrow(p \to q) \to (\Box_\leftarrow p \to \Box_\leftarrow q)$
- $\Box^\infty_\rightarrow(p \to q) \to (\Box^\infty_\rightarrow p \to \Box^\infty_\rightarrow q)$
- $\Box^\infty_\leftarrow(p \to q) \to (\Box^\infty_\leftarrow p \to \Box^\infty_\leftarrow q)$
- $\Diamond^\infty_\rightarrow p \to \Diamond_\rightarrow p$
- $\Diamond^\infty_\leftarrow p \to \Diamond_\leftarrow p$
- $p \to \Box_\leftarrow \Diamond_\rightarrow p$
- $p \to \Box_\rightarrow \Diamond_\leftarrow p$

We let $\mathbf{T}^\infty$ denote the least subset of T^∞ which extends $\mathsf{Ax}(T^\infty)$, and is closed under the following rules of inference:

- Modus ponens
- Necessitation: if $\psi \in \mathbf{T}^\infty$ then $\Box_\to\psi, \Box_\leftarrow\psi, \Box_\to^\infty\psi, \Box_\leftarrow^\infty\psi \in \mathbf{T}^\infty$
- Uniform substitution: If $\phi \in \mathbf{T}^\infty$, p is a proposition letter and ψ is a well-formed formula of T^∞, then $\phi[\psi/p] \in \mathbf{T}^\infty$

Proposition 6.5 (Soundness)**.** *For any formula $\psi \in T^\infty$, if $\psi \in \mathbf{T}^\infty$ then ψ is valid on all Kripke structures.*

Proof. Clearly the inference rules are sound, and it is straightforward to verify that each axiom is sound. □

Definition 6.6 ($\mathbf{T}^2$)**.** We let $\mathsf{Ax}(T^2)$ denote the following set of T^2 formulas:

- All propositional tautologies
- $\Box_\to(p \to q) \to (\Box_\to p \to \Box_\to q)$
- $\Box_\leftarrow(p \to q) \to (\Box_\leftarrow p \to \Box_\leftarrow q)$
- $\Box_\to^2(p \to q) \to (\Box_\to^2 p \to \Box_\to^2 q)$
- $\Box_\leftarrow^2(p \to q) \to (\Box_\leftarrow^2 p \to \Box_\leftarrow^2 q)$
- $\Diamond_\to^2 p \to \Diamond_\to p$
- $\Diamond_\leftarrow^2 p \to \Diamond_\leftarrow p$
- $p \to \Box_\leftarrow\Diamond_\to p$
- $p \to \Box_\to\Diamond_\leftarrow p$

We let $\mathbf{T}^2$ denote the least subset of T^∞ which extends $\mathsf{Ax}(T^\infty)$, and is closed under the inference rules of Modus ponens, Necessitation and Uniform substitution.

Proposition 6.7. *For any formula $\phi \in T^\infty$, it holds that $\phi \in \mathbf{T}^\infty$ iff $tr(\phi) \in \mathbf{T}^2$.*

Proof. It is easy to see that any derivation of ϕ in $\mathbf{T}^\infty$ may be turned into a derivation of $tr(\phi)$ in $\mathbf{T}^2$ by translating every formula ψ occurring in the derivation to $tr(\psi)$. Similarly, a derivation of $tr(\phi)$ in $\mathbf{T}^2$ may be transformed into a derivation of ϕ in $\mathbf{T}^\infty$ by replacing every formula ψ, occurring in the derivation, with $tr^{-1}(\psi)$. □

Proposition 6.8. *The axiomatisation* $\mathbf{T}^2$ *is sound and complete for the class of* $\mathbf{T}^2$*-frames.*

Proof. First of all, we claim that the class of $\mathbf{T}^2$-frames is precisely the class of frames which validate all formulas in $\mathbf{T}^2$. It is easy to see that if a frame $\mathcal{F}$ is a $\mathbf{T}^2$ frame, then every formula $\phi \in \mathbf{T}^2$ is valid on $\mathcal{F}$.

Conversely, suppose $\mathcal{F}$ is not a $\mathbf{T}^2$ frame. Then one of $R^2_{\rightarrow} \subseteq R, R^2_{\leftarrow} \subseteq \check{R}$ must fail. Suppose, for example, that there are $s, t \in S$ with $sR^2_{\rightarrow}t$ but not sRt. Then if V is a valuation with $V(p) = \{t\}$, then $\mathcal{M}, s$ refutes the axiom $\Diamond^2_{\rightarrow} p \rightarrow \Diamond p$. The other case is analogous.

We may conclude that $\mathbf{T}^2$ is valid on precisely the class of $\mathbf{T}^2$-frames. Now, observe that every axiom in $\mathsf{Ax}(T^2)$ is a Sahlqvist formula. So by the Sahlqvist Completeness Theorem (see e.g. [5]), $\mathbf{T}^2$ is complete with respect to the class of $\mathbf{T}^2$-frames. $\square$

Proposition 6.9. *Let ϕ be a satisfiable formula of T^∞. If ϕ is satisfiable on a Kripke frame, then $tr(\phi)$ is satisfiable on a* $\mathbf{T}^2$ *frame.*

Proof. Suppose $\phi \in T^\infty$ is satisfiable on a Kripke frame. Then since the axiomatisation $\mathbf{T}^\infty$ is sound, we can infer $\neg\phi \notin \mathbf{T}^\infty$. By Proposition 6.7, it follows that $tr(\neg\phi) \notin \mathbf{T}^2$. Now, $tr(\neg\phi)$ is just $\neg tr(\phi)$, and since $\mathbf{T}^2$ is complete with respect to $\mathbf{T}^2$ frames, it follows that $tr(\phi)$ is satisfiable on a $\mathbf{T}^2$ frame. $\square$

6.1.3 Satisfiability: From $tr(\phi)$ to ϕ

We now want to go in the opposite direction- given a formula ϕ which is satisfiable on some $\mathbf{T}^2$ model, we need to show that $tr^{-1}(\phi)$ is satisfiable. The first step is to show that T^2 has the finite model property.

Proposition 6.10. *Let ϕ be a satisfiable formula of T^2. Then ϕ is satisfiable on a finite model.*

Proof. If ϕ is satisfiable on a $\mathbf{T}^2$ model, then observe that $\phi' := ST(x, \phi) \wedge \forall xy(R^2_{\rightarrow}xy \rightarrow Rxy) \wedge \forall xy(R^2_{\leftarrow}yx \rightarrow Rxy)$ is also satisfied on this model. But ϕ' is in the guarded fragment of FOL, and the guarded fragment has the finite model property (see [10]). Therefore, ϕ' is satisfiable on some finite model $\mathcal{M}'$. Clearly $\mathcal{M}'$ is a $\mathbf{T}^2$ model. $\square$

Our strategy will be to define a model construction which turns any finite $\mathbf{T}^2$ model $\mathcal{M}$ into a Kripke model $\mathbb{B}(\mathcal{M})$ (the 'blooming' of $\mathcal{M}$), such that $\mathbb{B}(\mathcal{M}) \Vdash \psi$ iff $\mathcal{M} \Vdash tr(\psi)$, for *any* formula ψ. Before we can define this construction, we need the following lemma.

Lemma 6.11. *There exists a function $g : \mathbb{N} \rightarrow \mathbb{N}$ such that:*

- *g is a surjection*
- *For every $n \in \omega$, the preimage $g^{-1}[n]$ of n is infinite*

Proof. We build g in stages. At stage 0, let S_0 be some arbitrary infinite subset of $\mathbb{N}$, such that $\mathbb{N} \setminus S_0 := S_0'$ is also infinite. Let $g_0 := \{(n, 0) : n \in S_0\} \cup \{(0, 0)\}$, Finally, set $T_0 := S_0' \setminus \{0\}$.

At stage $k + 1$, let S_{k+1} be some arbitrary infinite subset of T_k such that $S_{k+1}' := T_k \setminus S_{k+1}$ is also infinite. If $k + 1 \notin dom(g_k)$, then set

$$g_{k+1} := g_k \cup \{(n, k + 1) : n \in S_{k+1}\} \cup \{(k + 1, k + 1)\}$$

Otherwise, set $g_{k+1} := g_k \cup \{(n, k + 1) : n \in S_{k+1}\}$. In addition, set $T_{k+1} := S_{k+1}' \setminus \{k + 1\}$.

Finally, define $g := \bigcup_{k \in \mathbb{N}} g_k$. Then g satisfies the desired conditions. □

For the remainder of this section, let g be some fixed function satisfying the conditions of Lemma 6.11. We now define the blooming construction mentioned above.

Definition 6.12 ($\mathbb{B}(\mathcal{M})$)**.** Let $\mathcal{M} = (S, R, R^2_{\rightarrow}, R^2_{\leftarrow}, V)$ be a $\mathbf{T}^2$ model. The universe S^B of $\mathbb{B}(\mathcal{M})$ is given by: $S^B := S \times \mathbb{N}$. The valuation V^B of $\mathbb{B}(\mathcal{M})$ is given by: $(s, i) \in V^B(p)$ iff $s \in V(p)$, for any proposition letter p.

To define whether the relation R^B holds between any two points $(s, i), (t, j) \in S^B$, we need to distinguish some cases.

- sRt does not hold. Then set $(s, i) \not R^B (t, j)$.
- sRt holds, but neither $sR^2_{\rightarrow}t$ nor $tR^2_{\leftarrow}s$ holds. Then set $(s, i)R^B(t, j)$ iff $i = j$.
- sRt and $sR^2_{\rightarrow}t$ hold, but $tR^2_{\leftarrow}s$ does not hold. Then set $(s, i)R^B(t, j)$ iff $i = g(j)$.
- sRt and $tR^2_{\leftarrow}s$ hold, but $sR^2_{\rightarrow}t$ does not hold. Then set $(s, i)R^B(t, j)$ iff $j = g(i)$.
- $sRt, sR^2_{\rightarrow}t$ and $tR^2_{\leftarrow}s$ all hold. Then set $(s, i)R^B(t, j)$, whatever the values of i, j.

Observe that these are the only cases that can occur, because $\mathcal{M}$ is a $\mathbf{T}^2$ model.

For any $s \in S$ and any $(t, i) \in S^B$, we say that (t, i) is a *representative* of s iff $t = s$.

We include some diagrams to give a feel for this construction.

$$
\begin{tikzcd}
\dots \arrow[r] & \dots \\
(s,4) \arrow[r] & (t,4) \\
(s,2) \arrow[r] & (t,2) \\
(s,0) \arrow[r] & (t,0) \\
(s,1) \arrow[r] & (t,1) \\
(s,3) \arrow[r] & (t,3) \\
\dots \arrow[r] & \dots
\end{tikzcd}
$$

Figure 6.1: The case where Rst, but neither $R^2_{\rightarrow}st$ nor $R^2_{\leftarrow}ts$

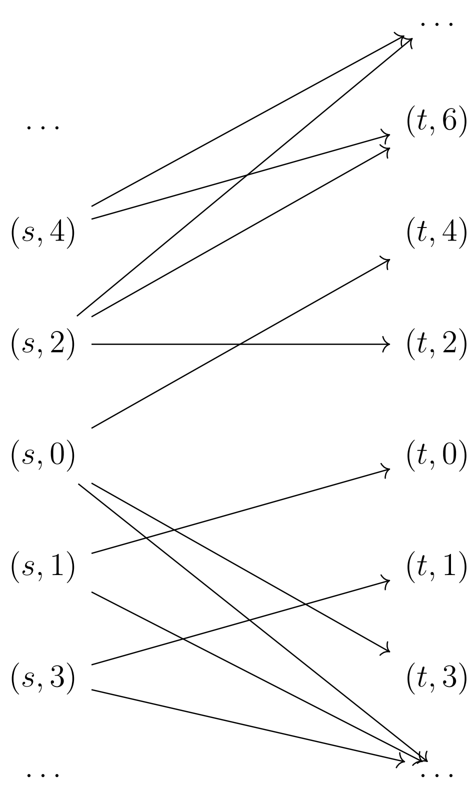

Figure 6.2: The case where sRt and $sR^2_{\rightarrow}t$, but not $tR^2_{\leftarrow}s$

Lemma 6.13. *Let $\mathcal{M}$ be a $\mathbf{T}^2$ model, $s, t \in S$. Then the following hold:*

1. *If sRt, then for every representative (s,i) of s, there is a representative (t,j) of t such that $(s,i)R^B(t,j)$*
2. *If sRt, then for every representative (t,j) of t, there is a representative (s,i) of s such that $(s,i)R^B(t,j)$*
3. *If $sR^2_\rightarrow t$, then for every representative (s,i) of s, there are infinitely many representatives (t,j) of t such that $(s,i)R^B(t,j)$*
4. *If $tR^2_\leftarrow s$, then for every representative (t,j) of t, there are infinitely many representatives (s,i) of s such that $(s,i)R^B(t,j)$*

Proof.

1. Suppose sRt. Let (s,i) be a representative of s. In the case that sRt and $tR^2_\leftarrow s$ but not $sR^2_\rightarrow t$, we have $(s,i)R^B(t,g(i))$. If $sR^2_\rightarrow t$ but not $tR^2_\leftarrow s$, then take any $j \in g^{-1}[i]$; then $(s,i)R^B(t,j)$. In both other cases, we are guaranteed to have $(s,i)R^B(t,i)$.
2. Suppose sRt. Let (t,j) be a representative of t. In the case that sRt and $sR^2_\rightarrow t$ but not $tR^2_\leftarrow s$, we have $(s,g(j))R^B(t,j)$. If $tR^2_\leftarrow s$ but not $sR^2_\rightarrow t$, then take any $i \in g^{-1}[j]$; then $(s,i)R^B(t,j)$. In both other cases, we are guaranteed to have $(s,j)R^B(t,j)$.
3. Suppose $sR^2_\rightarrow t$. Let (s,i) be a representative of s. If $tR^2_\leftarrow s$ does not hold, then we have $(s,i)R^B(t,j)$ for every $j \in g^{-1}[i]$, which is an infinite set. If $tR^2_\leftarrow s$ does hold, then we have $(s,i)R^B(t,j)$ for every $j \in \mathbb{N}$.
4. Suppose $tR^2_\leftarrow s$. Let (t,j) be a representative of t. If $sR^2_\rightarrow t$ does not hold, then we have $(s,i)R^B(t,j)$ for every $i \in g^{-1}[j]$, which is an infinite set. If $sR^2_\rightarrow t$ does hold, then we have $(s,i)R^B(t,j)$ for every $i \in \mathbb{N}$.

□

Lemma 6.14. *Let $\mathcal{M}$ be a $\mathbf{T}^2$ model, $s, t \in S$. Then the following hold:*

1. *If there are representatives (s,i) of s and (t,j) of t such that $(s,i)R^B(t,j)$, then sRt*
2. *If there is a representative (s,i) of s such that there are infinitely many representatives (t,j) of t such that $(s,i)R^B(t,j)$, then $sR^2_\rightarrow t$*
3. *If there is a representative (t,j) of t such that there are infinitely many representatives (s,i) of s such that $(s,i)R^B(t,j)$, then $tR^2_\leftarrow s$*

Proof. For each item, we proceed by contraposition.

1. Suppose $s \not R t$. Then by definition of R^B, we do not have $(s,i)R(t,j)$ for any $i,j \in \mathbb{N}$.

2. Suppose $sR^2_{\rightarrow}t$ does not hold. Consider any representative (s,i) of s. If $s \not R t$ then there is no j whatsoever such that $(s,i)R^B(t,j)$. If sRt but not $tR^2_{\leftarrow}s$, then $(s,i)R(t,j)$ only if $j = i$. If sRt and $tR^2_{\leftarrow}s$, then $(s,i)R(t,j)$ iff $j = g(i)$. In each case, we only have finitely many representatives (t,j) of t such that $(s,i)R^B(t,j)$.

3. Suppose $tR^2_{\leftarrow}s$ does not hold. Consider any representative (t,j) of t. If $s \not R t$ then there is no i whatsoever such that $(s,i)R^B(t,j)$. If sRt but not $sR^2_{\rightarrow}t$, then $(s,i)R(t,j)$ only if $i = j$. If sRt and $sR^2_{\rightarrow}t$, then $(s,i)R(t,j)$ iff $i = g(j)$. In each case, we only have finitely many representatives (s,i) of s such that $(s,i)R^B(t,j)$.

$\square$

Proposition 6.15. *Let $\mathcal{M}$ be a finite* $\mathbf{T}^2$ *model, $\phi \in T^\infty$, $s \in S$, $i \in \mathbb{N}$. Then $\mathcal{M}, s \Vdash tr(\phi)$ iff $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \phi$.*

Proof. We argue by induction on the construction of ϕ.

Base case: $\phi = p$. Then by definition of $\mathbb{B}(\mathcal{M})$, $(s,i) \in V^B(p)$ iff $s \in V(p)$.

Base case: $\phi = \bot$. Neither $\mathcal{M}, s \Vdash \bot$ nor $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \bot$.

Induction step: $\phi = \chi \wedge \psi$. Suppose $\mathcal{M}, s \Vdash tr(\chi \wedge \psi)$. Then $\mathcal{M}, s \Vdash tr(\psi) \wedge tr(\chi)$ by definition of tr. Hence, by IH, $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \psi$ and $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \chi$. Conversely, suppose $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \psi \wedge \chi$. Then $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \psi$ and $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \chi$, so by IH, $\mathcal{M}, s \Vdash tr(\chi)$ and $\mathcal{M}, s \Vdash tr(\psi)$. It follows that $\mathcal{M}, s \Vdash tr(\psi) \wedge tr(\chi)$ i.e. $\mathcal{M}, s \Vdash tr(\psi \wedge \chi)$.

Induction step: $\phi = \neg\psi$. If $\mathcal{M}, s \Vdash tr(\neg\psi)$ then $\mathcal{M}, s \nVdash tr(\psi)$, so by IH, $\mathbb{B}(\mathcal{M}), (s,i) \nVdash \psi$. Conversely, if $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \neg\psi$, then $\mathbb{B}(\mathcal{M}), (s,i) \nVdash \psi$, so by IH, $\mathcal{M}, s \nVdash tr(\psi)$. It follows that $\mathcal{M}, s \Vdash tr(\neg\psi)$.

Induction step: $\phi = \Diamond\psi$. The induction hypothesis is that for any $t \in S$, for any $j \in \mathbb{N}$, $\mathcal{M}, t \Vdash tr(\psi)$ iff $\mathbb{B}(\mathcal{M}), (t,j) \Vdash \psi$.

Suppose $\mathcal{M}, s \Vdash \Diamond tr(\psi)$. Then there is some $t \in S$ such that sRt and $\mathcal{M}, t \Vdash tr(\psi)$. By item 1 of Lemma 6.13, there is some $j \in \mathbb{N}$ such that $(s,i)R^B(t,j)$. By IH, $\mathbb{B}(\mathcal{M}), (t,j) \Vdash \psi$, and therefore $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \Diamond\psi$.

Conversely, suppose $\mathbb{B}(\mathcal{M}), (s,i) \Vdash \Diamond\psi$. Then there is some (t,j) such that $(s,i)R^B(t,j)$ and $\mathbb{B}(\mathcal{M}), (t,j) \Vdash \psi$. By item 1 of Lemma 6.14, sRt, and by IH,

$\mathcal{M}, t \Vdash tr(\psi)$. Hence $\mathcal{M}, s \Vdash \Diamond tr(\psi)$, which was to show.

Induction step: $\phi = \Diamond_{\leftarrow}\psi$. The inductive hypothesis is that for any $t \in S$, for any $j \in \mathbb{N}$, $\mathcal{M}, t \Vdash tr(\psi)$ iff $\mathbb{B}(\mathcal{M}), (t, j) \Vdash \psi$.

Suppose $\mathcal{M}, s \Vdash \Diamond_{\leftarrow} tr(\psi)$. Then there is some $t \in S$ such that tRs and $\mathcal{M}, t \Vdash tr(\psi)$. By item 2 of Lemma 6.13, there is some $j \in \mathbb{N}$ such that $(t, j)R^B(s, i)$. By IH, $\mathbb{B}(\mathcal{M}), (t, j) \Vdash \psi$, hence $\mathbb{B}(\mathcal{M}), (s, i) \Vdash \Diamond_{\leftarrow}\psi$.

Conversely, suppose $\mathbb{B}(\mathcal{M}), (s, i) \Vdash \Diamond_{\leftarrow}\psi$. Then there is some (t, j) with $(t, j)R^B(s, i)$ and $\mathbb{B}(\mathcal{M}), (t, j) \Vdash \psi$. By item 1 of Lemma 6.14, tRs, and by IH, $\mathcal{M}, t \Vdash tr(\psi)$. Hence $\mathcal{M}, s \Vdash \Diamond_{\leftarrow} tr(\psi)$.

Induction step: $\phi = \Diamond^{\infty}_{\rightarrow}\psi$. The inductive hypothesis is that for any $t \in S$, for any $j \in \mathbb{N}$, $\mathcal{M}, t \Vdash tr(\psi)$ iff $\mathbb{B}(\mathcal{M}), (t, j) \Vdash \psi$. Recall that $tr(\Diamond^{\infty}_{\rightarrow}\psi) = \Diamond^{2}_{\rightarrow} tr(\psi)$.

Suppose $\mathcal{M}, s \Vdash \Diamond^{2}_{\rightarrow} tr(\psi)$. Then there is some $t \in S$ such that $sR^2_{\rightarrow}t$ and $\mathcal{M}, t \Vdash tr(\psi)$. By item 3 of Lemma 6.13, there are infinitely many $j \in \mathbb{N}$ such that $(s, i)R^B(t, j)$, and by IH, for every such j we have $\mathbb{B}(\mathcal{M}), (t, j) \Vdash \psi$. Hence $\mathbb{B}(\mathcal{M}), (s, i) \Vdash \Diamond^{\infty}_{\rightarrow}\psi$.

Conversely, suppose $\mathbb{B}(\mathcal{M}), (s, i) \Vdash \Diamond^{\infty}_{\rightarrow}\psi$. Then there are infinitely many objects $x \in S^B$ such that $(s, i)R^B x$ and $\mathbb{B}(\mathcal{M}), x \Vdash \psi$. Further, by assumption, $\mathcal{M}$ is finite. It follows by the pigeonhole principle that there is some $b \in S$ such that there are infinitely many $j \in \mathbb{N}$ such that $(s, i)R^B(t, j)$ and $\mathbb{B}(\mathcal{M}), (t, j) \Vdash \psi$. By item 2 of Lemma 6.14, $sR^2_{\rightarrow}t$, and by IH, $\mathcal{M}, t \Vdash tr(\psi)$. Hence $\mathcal{M}, s \Vdash \Diamond^{2}_{\rightarrow} tr(\psi)$, which means $\mathcal{M}, s \Vdash tr(\Diamond^{\infty}_{\rightarrow}\psi)$.

Induction step: $\phi = \Diamond^{\infty}_{\leftarrow}\psi$. The inductive hypothesis is that for any $t \in S$, for any $j \in \mathbb{N}$, $\mathcal{M}, t \Vdash tr(\psi)$ iff $\mathbb{B}(\mathcal{M}), (t, j) \Vdash \psi$. Recall that $tr(\Diamond^{\infty}_{\leftarrow}\psi) = \Diamond^{2}_{\leftarrow} tr(\psi)$.

Suppose $\mathcal{M}, s \Vdash \Diamond_{\leftarrow} tr(\psi)$. Then there is some $t \in S$ such that $sR^2_{\leftarrow}t$ and $\mathcal{M}, t \Vdash tr(\psi)$. By item 4 of Lemma 6.13, there are infinitely many $j \in \mathbb{N}$ such that $(t, j)R^B(s, i)$, and by IH, for every such j, we have $\mathbb{B}(\mathcal{M}), (t, j) \Vdash \psi$. Hence $\mathbb{B}(\mathcal{M}), (s, i) \Vdash \Diamond^{\infty}_{\leftarrow}\psi$.

Conversely, suppose $\mathbb{B}(\mathcal{M}), (s, i) \Vdash \Diamond^{\infty}_{\leftarrow}\psi$. Then there are infinitely many $x \in S^B$ such that $xR^B(s, i)$ and $\mathbb{B}(\mathcal{M}), x \Vdash \psi$. Again, we use the finiteness of $\mathcal{M}$ and the pigeonhole principle to infer that there is some $b \in S$ such that there are infinitely many $j \in \mathbb{N}$ such that $(b, j)R^B(s, i)$ and $\mathbb{B}(\mathcal{M}), (b, j) \Vdash \psi$. By item 3 of Lemma 6.14, $sR^2_{\leftarrow}t$, and by IH, $\mathcal{M}, b \Vdash tr(\psi)$. Hence $\mathcal{M}, s \Vdash \Diamond^{2}_{\leftarrow} tr(\psi)$, which means $\mathcal{M}, s \Vdash tr(\Diamond^{\infty}_{\leftarrow}\psi)$.

This concludes the induction. $\square$

Corollary 6.16. *Let $\phi \in T^2$ be satisfiable on a $\mathbf{T}^2$ model. Then $tr^{-1}(\phi)$ is satisfiable on a Kripke model.*

Proof. Let ϕ be satisfied on the $\mathbf{T}^2$ model $\mathcal{M}$. Then by Proposition 6.10, ϕ is satisfied on some finite $\mathbf{T}^2$ model $\mathcal{M}'$. Hence, by Proposition 6.15, $tr^{-1}(\phi)$ is satisfied on $\mathbb{B}(\mathcal{M}')$. $\square$

6.1.4 Conclusions

We can now put the previous two sections together to obtain the following results:

Theorem 6.17 (Decidability of T^∞). *The satisfiability problem for T^∞ is decidable.*

Proof. By Proposition 6.9 and Corollary 6.16, for any $\phi \in T^\infty$, ϕ is satisfiable iff $tr(\phi)$ is satisfiable. Clearly the translation $tr(\cdot)$ is computable. Therefore, tr is an effective reduction of the satisfiability problem for T^∞ to the satisfiability problem for T^2. But the satisfiability problem for T^2 is decidable, because $\mathbf{T}^2$ is finitely axiomatisable by Sahlqvist formulas.[1] $\square$

Theorem 6.18 (Weak completeness of T^∞). $\mathbf{T}^\infty$ *is a complete axiomatisation of* T^∞.

Proof. Suppose $\phi \notin \mathbf{T}^\infty$. Then $tr(\phi) \notin \mathbf{T}^2$, which implies $\neg tr(\phi)$ is satisfiable on a $\mathbf{T}^2$ model. But then by Corollary 6.16, $\neg\phi$ is satisfiable on a Kripke model. By contraposition, we may conclude that if $\phi \in T^2$ is valid (i.e. $\neg\phi$ is unsatisfiable), then $\phi \in \mathbf{T}^\infty$. $\square$

We conclude this section with some remarks on the above proof. The strategy is closely connected to the strategy employed by Bellas Acosta in [1], to show that the class of ML^∞ formulas valid on all Kripke frames is axiomatisable. Bellas Acosta's approach is to give an alternative semantics directly for the logic ML^∞, in terms of bimodal frames with an extra accessibility relation R^∞ on which the infinity diamond $\Diamond^\infty$ is interpreted. We point out two differences between the approaches, one technical and one conceptual.

The technical difference is that the blooming construction given by Definition 6.12 is an improvement on Bellas Acosta's blooming construction (Definition 2.25 of [1]). As Bellas Acosta notes, his construction does not always preserve the truth value of ML^∞ formulas from bimodal $\{R, R^\infty\}$ models to Kripke models, although it *does* preserve truth values on tree models, which is sufficient for many purposes. Nevertheless, Proposition 6.15 shows that the construction given by Definition 6.12 above always preserves truth value from $tr(\phi)$ to ϕ. In this regard, it is more general and a clear improvement on the method introduced in [1].

[1] This reasoning is mistaken, because finite axiomatisation by Sahlqvist formulas is not sufficient to ensure decidability. Nevertheless, it is true that T^2 is decidable. This follows from Proposition 6.8 (finite axiomatisability) and Proposition 6.10 (finite model property).

The conceptual difference is that in the above proof, we first translated all formulas of T^{∞} into T^2 formulas, before interpreting them on T^2 frames. As noted earlier, the translation is syntactically an isomorphism, which means that the translation wasn't really necessary at all- we could have viewed T^2 as the very same logic as T^{∞}, and simply regarded T^2 frames as an alternative semantics for T^{∞}. Indeed, this is the approach taken by Bellas Acosta. However, an advantage of the translation approach is that it makes it clearer that T^{∞} *can* simply be viewed, syntactically, as a polymodal logic. There is less risk of being distracted by the seemingly special nature of the infinity diamond, making it easier to see that standard methods of modal logic may be applied to show that $\mathbf{T}^{\infty}$-consistent formulas are satisfiable on $\mathbf{T}^2$ frames. A further advantage is that it is very difficult to see how the alternative semantics approach might be generalised to the larger and more syntactically complex fragment, GF^{∞}, introduced in Chapter 5. By contrast, in the following section we will indicate a promising way in which the translation approach might be generalised to GF^{∞} (although we will not provide a conclusive proof that the approach works).

Finally, we briefly mention the question of the decidability of the fragments G^{∞} and GT^{∞}. It seems very likely that these fragments may also be axiomatised, and shown to be decidable, using similar methods to the above. However, we will not attempt to give a proof of this here. Instead, we prefer to discuss the decidability of GF^{∞}, since this would be a more interesting result and would imply the decidability of G^{∞} and GT^{∞} (since both are fragments of GF^{∞}). We turn to this question now.

6.2 Translation Scheme for GF^{∞}

In this section, we fix τ to be the signature containing countably many unary predicates $P_0, P_1, \ldots$, and a single binary predicate, R. We let GF^{∞}_{τ} denote the restriction of GF^{∞} to formulas in the signature τ. The following observation about GF^{∞}_{τ} is simple, but important:

Proposition 6.19. *Let ϕ be a sentence of GF^{∞}_{τ}. Then every subformula ψ of ϕ has at most two variables free.*

Proof. Suppose, for contradiction, that there is some subformula ψ of ϕ with three distinct free variables x, y, z. Since ϕ is a sentence, each of these variables must be bound by some quantifier; without loss of generality, suppose the quantifier which binds z is within the scope of the quantifiers binding the other two variables. Then the quantifier which binds z needs to be 'guarded' by an atomic formula containing all three variables x, y, z, which is impossible, because the signature τ contains no relation symbols of arity greater than 2. □

Clearly, for every formula $\phi \in GT^{\infty}$, we have $ST(x_i, \phi) \in GF^{\infty}_{\tau}$. So we may view GT^{∞} (and each of its sub-fragments) as fragments of GF^{∞}_{τ}. Further, our discussion in Chapter 5 showed that the natural language inference **I**, with which we began this thesis, may be formalised in τ. Reworking this idea, we may note that GF^{∞}_{τ} has validities such as:

$$\beta := \exists^{\infty} x \exists y (Rxy \wedge P_0 x \wedge P_1 y) \rightarrow (\exists y \exists^{\infty} x (Rxy \wedge P_0 x \wedge P_1 y) \vee \exists^{\infty} y P_1 y)$$

which also turns on the principle that a finite union of finite sets is finite. It is instructive to compare the validity β with the fourth axiom schema of the axiomatisation for $\mathcal{L}(Q_1)$ given by Keisler [11]. Keisler's axiom schema is:

$$Q_1 x \exists y \phi(x, y) \rightarrow (\exists y Q_1 x \phi(x, y) \vee Q_1 y \exists x \phi(x, y))$$

the validity of which turns on the principle that a countable union of countable sets is countable (as opposed to: a finite union of finite sets is finite). Further, the presence of such non-obvious validities as β in GF^{∞}_{τ} indicates that the validity problem for GF^{∞}_{τ} is significantly more difficult than that for T^{∞}, which doesn't seem to have any validities of comparable interest.

Nevertheless, we conjecture that the logic GF^{∞}_{τ} *is* decidable. We do not yet have a full proof of this claim. However, we will indicate a promising strategy for showing decidability of this logic. The strategy involves a translation of the sentences of GF^{∞}, analogous to the translation used to show decidability for T^{∞} in the previous section. (Cf. also Keisler's [11] completeness proof for $\mathcal{L}(Q_1)$, which also employs the idea of reducing formulas of the logic $\mathcal{L}(Q_1)$ to formulas of *FOL*.)

Definition 6.20 (τ^+)**.** We let τ^+ denote the signature τ, extended by a fresh unary predicate P^{∞} for every unary predicate $P \in \tau$, and also extended by the following binary relation symbols:

- $R^{1,\infty}_{1,2}$
- $R^{1,\infty}_{2,1}$
- $R^{\infty,1}_{1,2}$
- $R^{\infty,1}_{2,1}$
- $R^{\infty,\infty}_{1,2}$
- $R^{\infty,\infty}_{2,1}$
- $=^{1,\infty}_{1,2}$

- $=^{1,\infty}_{2,1}$
- $=^{\infty,1}_{1,2}$
- $=^{\infty,1}_{2,1}$
- $=^{\infty,\infty}_{1,2}$
- $=^{\infty,\infty}_{2,1}$

We also introduce the convention that P^1 is syntactically identical to P for each unary P, that $R^{1,1}_{1,2}$ and $R^{1,1}_{2,1}$ are both syntactically identical to the relation symbol R, and that $=^{1,1}_{1,2}, =^{1,1}_{2,1}$ are both syntactically identical to the relation symbol $=$.

The point of introducing all these new predicate symbols is in order to give a translation from GF^{∞}_{τ} formulas, into τ^{+} formulas of GF. Our translation will replace all infinity quantifiers with regular existential quantifiers. But we also want to keep track of which variables were originally bound by an infinity quantifier, and which occurred within the scope of which. The point of the superscripts and subscripts of the new predicates is to store this information. For example, we will translate the formula $\theta := \exists x \exists^{\infty} y(Rxy)$ to:

$$\exists x \exists y(R^{1,\infty}_{1,2} xy)$$

To see the importance of the scope markers, observe that the formula $\theta' := \exists^{\infty} y \exists x(Rxy)$ is *not* equivalent to θ (it is strictly weaker). Therefore, we do not want to translate θ' to $\exists y \exists x(R^{1,\infty}_{1,2} xy)$, as this would be equivalent to the translation of θ. Instead, we translate θ' to:

$$\exists y \exists x(R^{1,\infty}_{2,1} xy)$$

The subscript $2, 1$ indicates that in the original formula θ', the quantifier which binds x occurs *within* the scope of the quantifier which binds y (unlike in the formula θ).

This is the high-level intuition behind the translation; we must now give a precise definition of it. For simplicity, we restrict attention to the satisfiability problem for *sentences* of GF^{∞}_{τ}. We still need to provide a translation for formulas of GF^{∞}_{τ}, in order that the translation of quantified sentences may be inductively built up from the translations of their subformulas. However, in view of Proposition 6.19, we may confine our attention to formulas with at most two free variables, and we may further assume that a formula of the form $Qu\psi$, where Q is a quantifier, has at most one variable free. (Otherwise, ψ would have more than two variables free.) We in fact define a family of different translations; these different translations may disagree on formulas with free variables, but they are guaranteed to agree on sentences, as the reader may verify.

Definition 6.21. We define a series of operations, $\{\psi(u,v) \mapsto \psi^{q,r}_{s,3-s}(u,v) : q,r \in \{1,\infty\}, s \in \{1,2\}\}$. Each operation is defined inductively on the set $\{\psi \in GF^{\infty}(\tau) : \psi$ has at most two free variables $\}$. Note that in the following, u, v and w are used schematically to range over variables, and they need not denote distinct variables.

We give one particular example, to give the reader a feeling for the definition. $\psi(u,v) \mapsto \psi^{1,\infty}_{1,2}(u,v)$ is defined by:

- $(Pu)^{1,\infty}_{1,2}(u,v) = Pu$
- $(Ruv)^{1,\infty}_{1,2}(u,v) = R^{1,\infty}_{1,2}uv$
- $(\psi \wedge \chi)^{1,\infty}_{1,2}(u,v) = \psi^{1,\infty}_{1,2}(u,v) \wedge \chi^{1,\infty}_{1,2}(u,v)$
- $(\neg\psi)^{1,\infty}_{1,2}(u,v) = \neg\psi^{1,\infty}_{1,2}(u,v)$
- $(\exists u\psi)^{1,\infty}_{1,2}(v) = \exists u(\psi^{1,\infty}_{2,1}(u,v))$
- $(\exists^{\infty} u\psi)^{1,\infty}_{1,2}(v) = \exists u(\psi^{\infty,\infty}_{2,1}(u,v))$

We now give the general definition. $\psi(u,v) \mapsto \psi^{q,r}_{s,3-s}(u,v)$ is given by:

- $(Pu)^{q,r}_{s,3-s}(u,v) = P^{q}u$
- $(Ruv)^{q,r}_{s,3-s}(u,v) = R^{q,r}_{s,3-s}uv$
- $(\bot)^{q,r}_{s,3-s}(u,v) = \bot$
- $(\psi \wedge \chi)^{q,r}_{s,3-s}(u,v) = \psi^{q,r}_{s,3-s}(u,v) \wedge \chi^{q,r}_{s,3-s}(u,v)$
- $(\neg\psi)^{q,r}_{s,3-s}(u,v) = \neg\psi^{q,r}_{s,3-s}(u,v)$
- $(\exists u\psi)^{q,r}_{s,3-s}(v) = \exists u(\psi^{1,r}_{2,1}(u,v))$
- $(\exists^{\infty} u\psi)^{q,r}_{s,3-s}(v) = \exists u(\psi^{\infty,r}_{2,1}(u,v))$

Since these different translations all agree on sentences, we are justified in using a functional notation, $\phi \mapsto \phi^*$, to denote 'the' translation of ϕ. For example, we may simply define $\phi^* := \phi^{1,1}_{1,2}(x,y)$.

Now, as it stands, the translation we have given will not preserve satisfiability. To see this, observe that the FOL^{∞} formula $\exists^{\infty}xPx \wedge \neg\exists xPx$ is unsatisfiable. However, the translation does not respect this: $\exists xP^{\infty}x \wedge \neg\exists xPx$ *is* satisfiable. Intuitively, we want the P^{∞} predicate to be stronger than the P predicate. In order to guarantee this, and other nice behaviour from our new predicate symbols, we introduce a series of *axioms*.

Definition 6.22 (Axioms for new predicates)**.** Let τ' be any signature such that $\tau' \subseteq \tau$. We define $\mathsf{Ax}(\tau')$ to be the set containing all instances of the following axiom schemata, where 'P' may be replaced by any unary predicate letter from τ'.

- $\forall x(P^{\infty}x \to (Px \wedge \mathit{Inf}x))$
- $\forall xy(R^{\infty,1}_{1,2}xy \to (Rxy \wedge \mathit{Inf}x))$
- $\forall xy(=^{\infty,1}_{1,2} xy \to (= xy \wedge \mathit{Inf}x))$
- $\forall xy(R^{1,\infty}_{2,1}xy \to (Rxy \wedge \mathit{Inf}y))$
- $\forall xy(=^{1,\infty}_{2,1} xy \to (= xy \wedge \mathit{Inf}y))$
- $\forall x(\mathit{Inf}x \to (Px \to P^{\infty}x))$
- $\forall xy(Rxy \to (\mathit{Inf}x \to R^{\infty,1}_{1,2}xy))$
- $\forall xy(x = y \to (\mathit{Inf}x \to =^{\infty,1}_{1,2} xy))$
- $\forall xy(Rxy \to (\mathit{Inf}y \to R^{1,\infty}_{2,1}xy))$
- $\forall xy(x = y \to (\mathit{Inf}y \to =^{1,\infty}_{2,1} xy))$
- $\forall xy(R^{1,\infty}_{1,2}xy \to R^{1,\infty}_{2,1}xy)$ (**Scope**)
- $\forall xy(R^{\infty,1}_{2,1}xy \to R^{\infty,1}_{1,2}xy)$ (**Scope**)
- $\forall xy(R^{\infty,1}_{1,2}xy \to (\neg R^{\infty,1}_{2,1}xy \to \mathit{Inf}y))$ (**Pigeonhole**)
- $\forall xy(R^{1,\infty}_{2,1}xy \to (\neg R^{1,\infty}_{1,2}xy \to \mathit{Inf}x))$ (**Pigeonhole**)
- $\forall x(R^{1,\infty}_{1,2}xx \to \bot)$ (**Non-refl**)
- $\forall x(R^{\infty,1}_{2,1}xx \to \bot)$ (**Non-refl**)
- $\forall xy(x =^{1,\infty}_{1,2} y \to \bot)$ (**Identity**)
- $\forall xy(x =^{\infty,1}_{2,1} y \to \bot)$ (**Identity**)

Observe the crucial fact that every one of these axioms is within the guarded fragment of *FOL*. Observe also that, given any formula $\phi \in GF^{\infty}_{\tau}$, the signature $\tau_{\phi} \subseteq \tau$, containing all predicate symbols occurring in ϕ, is finite, and therefore $\mathsf{Ax}(\tau_{\phi})$ is a finite set. Combining these facts, we can see that $\bigwedge \mathsf{Ax}(\tau_{\phi})$ is a well-formed sentence of *GF*. We may now give the final definition of our translation from GF^{∞}_{τ} to *GF*.

Definition 6.23. Let ϕ be a sentence of GF^{∞}_{τ}. We define

$$f(\phi) := \phi^* \wedge \bigwedge(\mathsf{Ax}(\tau_\phi))$$

Now, we would like to show that if $f(\phi)$ is satisfiable, then so is ϕ. One way to achieve this would be to give a model construction, transforming any model of $f(\phi)$ into a model of ϕ. A natural approach to this, along the same lines as the strategy employed in Section 6.1, is to turn any pair of points (a, b) satisfying, e.g., the predicate $R^{1;\infty}_{1,2}$, into witnesses for the FOL^{∞} formula $\exists x \exists^{\infty} y Rxy$. So, we could create infinitely many representatives of b, and have each representative of a stand in R to infinitely many of these representatives of b. We could do something analogous for each fresh predicate symbol, and associated quantifier combination. Keeping this in mind, we can explain the motivation for the axioms of $\mathsf{Ax}(\tau_\phi)$ a little more clearly.

Suppose we have two points (a, b) satisfying $R^{1;\infty}_{1,2}$ in a model of $f(\phi)$. Then in the model we would construct, each representative of a will see infinitely many representatives of b. That means there will certainly be infinitely many representatives of b which are each seen by some a. Hence, we want the pair (a, b) to also satisfy the predicate $R^{1;\infty}_{2,1}$. This is why the 'scope' axioms are included among $\mathsf{Ax}(\tau')$.

Now suppose there are two points a, b satisfying $R^{1;\infty}_{2,1}$ in a model of $f(\phi)$. Suppose further that these points do *not* satisfy $R^{1;\infty}_{1,2}$. Then in the model we construct, there will be infinitely many representatives of b which are each seen by some a. Now suppose the constructed model only contains finitely many representatives of a. Then a simple pigeonhole argument implies that some representative of a sees infinitely many representatives of b. Contraposing this reasoning, if no representative of a sees infinitely many representatives of b, then there must be infinitely many representatives of a; so we want a, b to satisfy $\neg R^{1;\infty}_{1,2} xy \to Inf x$. This is why the 'pigeonhole' axioms are included among $\mathsf{Ax}(\tau')$.

Finally, observe that if two points a, b satisfy $=^{1;\infty}_{1,2}$ in a model of $f(\phi)$, then in the constructed model, each representative of a should be identical to infinitely many *distinct* representatives of b. But this is clearly impossible. (Observe also that any FOL^{∞}- formula of the form $\exists x \exists^{\infty} y(\ldots \wedge x = y)$ is unsatisfiable, so its f-image should also be unsatisfiable.) This is why the 'identity' axioms are included among $\mathsf{Ax}(\tau')$.

We conclude this chapter with the following conjecture:

Conjecture 6.24. Let ϕ be a sentence of GF^{∞}_{τ}. Then ϕ is satisfiable if and only if $f(\phi)$ is satisfiable.

An immediate corollary of Conjecture 6.24 would be that the satisfiability problem for GF^{∞}_{τ} is decidable. To decide whether any formula $\phi \in GF^{\infty}_{\tau}$ is satisfiable, one could compute the translation $f(\phi)$ of ϕ (the translation f is clearly computable), then check whether $f(\phi)$ is satisfiable- using the fact that GF is decidable, and $f(\phi) \in GF$. A further corollary would be that GT^∞ and G^∞ are both decidable.

Chapter 7

Conclusion and Future Work

In this thesis, we have made some progress towards providing a nice model theory for fragments of FOL^∞. In Chapter 3, we established our first substantive result, the Gaifman theorem for FOL^∞, which shows that FOL^∞ is 'essentially local', in the same sense in which FOL may be said to be essentially local ([15], p. 188). We then used this result, and other model construction methods, to give a semantic characterisation of various syntactically given fragments of FOL^∞. Specifically, we showed that to each of T^∞, G^∞ and GT^∞, there corresponds a suitable notion of bisimulation for which each fragment is *the* bisimulation invariant fragment of FOL^∞. (In the cases of G^∞ and GT^∞ this requires some qualification, since our result was restricted to sentences with no free variables.)

On the other hand, we have had less success with the stronger *guarded fragment*, GF^∞, of FOL^∞. This logic has a much richer inferential structure, including inferences that fundamentally rely on facts about cardinality such as 'a finite union of finite sets is finite'. This increase in expressive power comes hand in hand with a loss of model-theoretic tractability: structures do not in general have acyclic bisimilar companions. Nevertheless, we have at least been able to define a natural notion of bisimulation, $\simeq_{GF}$, under which all formulas of GF^∞ are invariant. We have also left open the question of whether GF^∞ is decidable, although we briefly sketched a strategy for attempting to show that it is decidable.

The thesis leaves plenty of questions open for future investigation. We conclude by discussing a few of these.

Gaifman Theorems for Abstract Logics

As we observed at the end of Chapter 3, it seems plausible that an analogue of the Gaifman theorem for FOL^∞ (Theorem 3.8) would hold for other logics of the form $\mathcal{L}_{\omega,\omega}(Q_\alpha)$. Indeed, it seems plausible that the very same proof as was given for Theorem 3.8 could be adapted to the case of other logics. A natural next step is to investigate whether this is in fact the case. However, a more fundamental question also suggests itself. In Chapter 3, we observed that for any given abstract logic, a problem arises as to what should count as a 'basic local sentence' for that logic. We also observed that, if we presuppose an answer to this question, in the form of some map which assigns to each logic the set of sentences which are to count as its basic local sentences, then it makes sense to define a logic to *have the Gaifman property* iff every one of its formulas can be expressed in terms of local formulas, and its basic local sentences. This raises the question: which logics have the Gaifman property? Are there any more fundamental properties of logics which determine whether they have the Gaifman property, and does the Gaifman property in turn imply other important properties for abstract logics?

Bisimulation Characterisation for GF^∞

A major question which has been left open is whether GF^∞ is characterised by invariance under $\simeq_{GF}$. Specifically, is it the case that every formula of FOL^∞ which is invariant under $\simeq_{GF}$ is logically equivalent to a formula of GF^∞? The proof method used to characterise G^∞ and GT^∞ breaks down at the Upgrading stage, where we would like to show that for any structures $\mathfrak{A}, \mathfrak{B}$, if $\mathfrak{A} \simeq^k_{GF} \mathfrak{B}$, then there are $\simeq_{GF}$-bisimilar companions $\mathfrak{A}', \mathfrak{B}'$ with $\mathfrak{A}' \equiv^{k,q}_{1,\exists} \mathfrak{B}'$. In the $\simeq_{GF}$ setting, the uniform strategy of finding acyclic bismilar companions is not available, since acyclic bisimilar companions need not exist. For any two particular structures $\mathfrak{A} \simeq^k_{GF} \mathfrak{B}$, it is generally not difficult to find suitable companion structures $\mathfrak{A}', \mathfrak{B}'$, but the process may differ significantly from case to case: sometimes we can, and should, choose acyclic companions $\mathfrak{A}', \mathfrak{B}'$, whereas sometimes, we can't choose acyclic companions. Therefore, the challenge seems to be to find a *uniform* model construction, which is guaranteed to yield suitably indistinguishable bisimilar companions for any two $\mathfrak{A} \simeq^k_{GF} \mathfrak{B}$.

Decidability of GF^∞

The other major question left open in this thesis is whether the satisfiability problem for GF^∞ is decidable. In chapter 6, we stated a conjecture (Conjecture 6.24) which would immediately imply that this problem is indeed decidable. If this conjecture could be established, this would be a result of some significance. Not only would it show GF^∞ to be decidable, it would also vindicate the proof idea of reducing FOL^∞ formulas to first-order formulas by replacing infinity quantifiers with regular existential quantifiers, and adding fresh predicate symbols to store the information originally carried by the quantifiers. This would open up the possibility that the same method could be applied to other abstract logics, which make use of generalised quantifiers other than $\exists^\infty$. For this reason, the status of Conjecture 6.24 is an exciting question for future work.

Index of Notation

Bibliography

[1] Ignacio Bellas Acosta. *Studies in the Extension of Standard Modal Logic with an Infinite Modality*. 2020. URL: https://5dok.net/document/oz1e1r8y-studies-extension-standard-modal-logic-infinite-modality.html.

[2] Ignacio Bellas Acosta and Yde Venema. "Counting to Infinity: Graded Modal Logic with an Infinity Diamond". In: *Review of Symbolic Logic* (2022), pp. 1–36. DOI: 10.1017/S1755020322000247.

[3] Hajnal Andréka, Johan van Benthem, and Istvan Németi. "Modal Languages and Bounded Fragments of Predicate Logic". In: *Journal of Philosophical Logic* 27.3 (1998), pp. 217–274. DOI: https://www.jstor.org/stable/30226645.

[4] Jon Barwise. "Model-Theoretic Logics: Background and Aims". In: *Model-Theoretic Logics*. Ed. by Jon Barwise and Solomon Feferman. Perspectives on Mathematical Logic. New York: Springer-Verlag, 1985, pp. 3–24.

[5] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. New York: Cambridge University Press, 2001.

[6] Facundo Carreiro et al. "Model Theory of Monadic Predicate Logic with the Infinity Quantifier". In: *Archive of Mathematical Logic* 61.3-4 (2021), pp. 465–502. DOI: https://doi.org/10.1007/s00153-021-00797-0.

[7] Balder ten Cate, Johan van Benthem, and Jouko Väänänen. "Lindström theorems for fragments of first-order logic". In: *Logical Methods in Computer Science* 5.3 (2009), pp. 1–27. DOI: https://doi.org/10.2168/LMCS-5(3:3)2009.

[8] Gebhard Fuhrken. "Languages with added quantifier "There exist at least $\aleph_\alpha$"". In: *The Theory of Models*. Ed. by J.W. Addison, Leon Henkin, and Alfred Tarski. Proceedings of the 1963 International Symposium at Berkeley. Amsterdam: North-Holland, 1965, pp. 121–131.

[9] Haim Gaifman. "On Local and Non-Local Properties". In: *Studies in Logic and the Foundations of Mathematics* 107 (1982), pp. 105–135. DOI: https://doi.org/10.1016/S0049-237X(08)71879-2.

[10] Erich Grädel. “On the Restraining Power of Guards”. In: *Journal of Symbolic Logic* 64.4 (1999), pp. 1719–1742. DOI: https://doi.org/10.2307/2586808.

[11] H. Jerome Keisler. “Logic with the quantifier “there exist uncountably many””. In: *Annals of Mathematical Logic* 1.1 (1970), pp. 1–93. DOI: https://doi.org/10.1016/S0003-4843(70)80005-5.

[12] H. Jerome Keisler and Wafik Boulos Lotfallah. “A local normal form theorem for infinitary logic with unary quantifiers”. In: *Mathematical Logic Quarterly* 51.2 (2005), pp. 137–144. DOI: https://doi.org/10.1002/malq.200410013.

[13] Andrzej Mostowski. “Craig’s Interpolation Theorem in Some Extended Systems of Logic”. In: *Logic, Methodology and Philosophy of Science III*. Ed. by B. Van Rootselaar and J.F. Staal. Studies in Logic and the Foundations of Mathematics. North-Holland, 1968, pp. 87–103.

[14] Andrzej Mostowski. “On a Generalisation of Quantifiers”. In: *Fundamenta Mathematicae* 44.1 (1957), pp. 12–36. DOI: https://doi.org/10.4064/fm-44-1-12-36.

[15] Martin Otto. “Modal and guarded characterisation theorems over finite transition systems”. In: *Annals of Pure and Applied Logic* 130.1-3 (2004), pp. 173–205. DOI: https://doi.org/10.1016/j.apal.2004.04.003.

[16] Eric Rosen. “Modal Logic over Finite Structures”. In: *Journal of Logic, Language and Information* 6.4 (1997), pp. 427–439. DOI: https://doi.org/10.1023/A:1008275906015.

[17] Jouko Väänänen. *Models and Games*. New York: Cambridge University Press, 2011.

[18] Robert Vaught. “The completeness of logic with the added quantifier “there are uncountable many””. In: *Fundamenta Mathematicae* 54.3 (1964), pp. 303–304. DOI: https://doi.org/10.4064/fm-54-3-303-304.