METHODS FOR CANONICITY

MSc Thesis (*Afstudeerscriptie*)

written by

Samuel Jacob van Gool
(born November 1st, 1987 in Amsterdam, The Netherlands)

under the supervision of **Dr Alessandra Palmigiano**, and submitted to the Board of Examiners 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:**
November 30, 2009	Prof Dr Mai Gehrke
	Prof Dr Dick de Jongh
	Dr Alessandra Palmigiano
	Dr Yde Venema

INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

Abstract

In the first part of this thesis, we focus on the canonical extension of partially ordered sets, which was defined by algebraic means by Dunn, Gehrke and Palmigiano [9]. We show that it can be obtained alternatively via a generalization of Urquhart and Hartung's maximal filter-ideal pair construction ([43], [27]). We further give a first-order dual characterization of perfect lattice hemi- and homomorphisms, in the spirit of, but different from Gehrke [16], and make category-theoretic observations regarding the canonical extension.

The second part of the thesis concerns the algebraic canonicity proof of the Sahlqvist fragment for distributive modal logic by Gehrke, Nagahashi and Venema [22]. We pay particular attention to the additional operation **n**, which is crucial to that proof, and show that the proof can not be straightforwardly translated to an algebraic canonicity proof of the inductive fragment for distributive modal logic [7]. We extract requirements on a new version of the operation **n**, which would yield a proof of the canonicity of the inductive fragment, and finish by starting to explore two new perspectives on the magical nature of the operation **n**.

Contact information

Sam van Gool

samvg@mac.com

Contents

Acknowledgements

An earlier version of this document was submitted as my Master of Logic thesis at the University of Amsterdam, and successfully defended on 30 November, 2009. The main part of it was written between June and November 2009. Some parts, most notably section 1.4 of Chapter 1, needed to be revised following up questions and discussions at the occasion of and after my Thesis defense, and this was done somewhat later and sporadically in the first months of 2010 at the Radboud Universiteit Nijmegen and LIAFA, Université Paris-Diderot, France.

I would like to take this opportunity to thank my thesis committee, consisting of Mai Gehrke, Yde Venema, Dick de Jongh and Alessandra Palmigiano for their critical questions, from which I benefited a lot, and so did the thesis (hopefully).

Alessandra Palmigiano, as the supervisor of this thesis, deserves special mention: she inspired me to take up this project; I fondly remember many afternoons of working on the project together, which was sometimes hard, and even stressful work, but without exception a lot of fun.

Writing this thesis would also not have been possible without moral support and necessary distractions from many people. Thanks to: Jacobien Carstens, Chris van Dorp, Evelien, Floor and Sophie Frenkel, Lisa Fulford, Christian Geist, my parents Pim and Carolien, my sisters Sophie, Rosa and Joanne, Jan Loerakker, Niels van Maanen, Gover Meit, Joe Mier, Joep Pijpers, Hanna Schreuders, Jonathan Shaheen, Bouwien Smid.

Thanks also to all those whom I did not mention here; I do explicitly express my sincere hope that they are somewhat consoled by being mentioned implicitly in the current sentence.

Paris, April 21, 2010.

Introduction

1 The canonical extension via duality

Representation theorems are statements of the form: every element of the class of structures X is isomorphic to some element of the proper subclass Y of X. Representation theorems can be found in various fields of mathematics: among one of the first and most famous examples is *Cayley's theorem* that every group is isomorphic to a subgroup of a symmetric group, stated and partly proved by Cayley in 1854 [6]. In functional analysis, *Riesz's theorems*, to the effect that certain linear functionals on a space can be represented as measures, are important instances of representation theorems, dating from the first decade of the twentieth century [39], [40]. Taking a leap of roughly fifty years, in 1954 or '55, Nobuo Yoneda met Saunders MacLane in a café at the Gare du Nord in Paris, and told him about the category-theoretic fact, which was later baptized the *Yoneda Lemma* ([33], [32]). This fact greatly generalised and unified many of the previous representation theorems. For instance, when applied to a group regarded as a category, the Yoneda Lemma yields Cayley's theorem (*Cf.* [8], Example 2.7.5).
The conceptual importance of a representation theorem is proportional to the size and complexity of the proper subclass Y of the general class of structures X: if this subclass is much more transparent and easy to work with than the general class, then the representation theorem makes it easier both to conceive and to prove isomorphism-invariant properties of *all* structures in the (apparently) larger class X.
Representation theorems are of particular importance in logic because they pave the way to *completeness results*, to be made precise below in Section 2. In the field of mathematical logic, the following theorem, proved by Marshall H. Stone in 1936, is probably the best known representation theorem.

Stone's representation theorem (1936) [41]. *Every Boolean algebra $\mathbb{A}$ can be embedded in the complete and atomic Boolean algebra $\mathbb{A}^\sigma$, defined as the power set algebra of the set of ultrafilters of $\mathbb{A}$.*

This statement already looks very significant. For instance, it tells us a lot about the behaviour of the operations in a Boolean algebra and gives us a very powerful grasp on their *meaning*. However, it is fair to say that the abbreviated form in which the theorem is stated here discloses only a tiny fraction of the true significance and potential of this result, which is mostly concentrated in the special properties of the *embedding* of $\mathbb{A}$ into $\mathbb{A}^\sigma$.
These special properties are *abstractly* encoded in the definition of the central concept of this thesis: the *canonical extension* $\mathbb{A}^\sigma$. Moreover, they are also *concretely* rooted in the 'powerset'

and ‘ultrafilter’ functors, each of which is part of a *dual equivalence* or *duality* of categories. Loosely speaking, a duality is a one-to-one correspondence, or ‘back-and-forth translation’, between two categories that preserves and reflects the logical information stored in the two categories. It works in part as a dictionary and in part as a comparative grammar: a duality does not only translate words (e.g., homomorphism), but it also translates properties (e.g., injective) and constructions (e.g., product) from the one category into the *dual* properties (e.g., surjective) and constructions (e.g., disjoint sum) of the other category.
The following diagram is meant to express how Stone’s representation theorem relies on the existence of two dualities of categories.

$$
\begin{tikzcd}[column sep=huge, row sep=huge]
\mathsf{CABA} \arrow[r, shift left, "(\cdot)_+"] & \mathsf{Set} \arrow[l, shift left, "(\cdot)^+"] \\
\mathsf{BA} \arrow[u, dashed, "(\cdot)^\sigma"] \arrow[ur, dashed, "(\cdot)_\bullet"] \arrow[r, shift left, "(\cdot)_*"] & \mathsf{Stone} \arrow[u, "U"'] \arrow[l, shift left, "(\cdot)^*"]
\end{tikzcd}
$$

Figure 1: *Duality and Stone’s representation theorem*

Let us take some time to examine this diagram in detail, since its shape is a ‘blueprint’ for some more such diagrams that we will consider in this thesis. We first describe the categories and (object components of the) functors involved.
The category BA is that of Boolean algebras with homomorphisms, CABA is the category of complete atomic Boolean algebras with complete homomorphisms, Stone is the category of Stone spaces (i.e., compact, Hausdorff, totally disconnected topological spaces) with continuous maps, and Set the category of sets with functions.
The functor pair $(\cdot)_+$, $(\cdot)^+$ constitutes a dual equivalence between the categories CABA and Set, which we will refer to as the *discrete duality.* For a set S, S^+ is defined as the power set algebra of S, and for a complete atomic Boolean algebra $\mathbb{A}$, $\mathbb{A}_+$ is the set of atoms of $\mathbb{A}$.
The functors $(\cdot)_*$, $(\cdot)^*$ constitute a *topological duality* between BA and Stone, respectively. To any Boolean algebra $\mathbb{A}$, the functor $(\cdot)_*$ associates the *ultrafilter space* $\mathbb{A}_*$, and for any Stone space X, X^* is defined to be the Boolean algebra of clopen sets of that space.
Finally, the *canonical extension functor* $(\cdot)^\sigma$ is defined as $(\cdot)^\sigma := (\cdot)^+ \circ U \circ (\cdot)_*$, and the functor $(\cdot)_\bullet$ is defined to be the composition $U \circ (\cdot)_*$.
It is now easy to see that by combining the topological and discrete duality, we actually get two important results:

(i) Any Boolean algebra $\mathbb{A}$ can be *embedded* in the algebra $\mathbb{A}^\sigma$, and,

(ii) Any Boolean algebra homomorphism $f : \mathbb{A} \to \mathbb{B}$ can be *extended* to a complete Boolean algebra homomorphism $f^\sigma : \mathbb{A}^\sigma \to \mathbb{B}^\sigma$.

Putting the first item in a slightly different wording, we see that *any Boolean algebra is (up to isomorphism) a subalgebra of a power set algebra*: Stone's representation theorem.
The second item can be generalised: we can actually apply the functor $(\cdot)^\sigma$ not only to Boolean algebra homomorphisms, but to any (order-preserving) function. This construction will be studied in a more general context in Chapter 1 of this thesis, and is one of the important ingredients which enables us to use canonical extensions to study *modal* logic.
Using this construction, it is possible to devise a diagram, similar to that in Figure 1, in which Boolean algebras *with operators* (BAO's) occupy the lower left corner.

$$\begin{tikzcd}
\mathsf{CABACO} \arrow[rr, shift left, "(\cdot)_+"] & & \mathsf{KrFr} \arrow[ll, shift left, "(\cdot)^+"] \\
& & \\
\mathsf{BAO} \arrow[uu, dashed, "(\cdot)^\sigma"] \arrow[uurr, dashed, "(\cdot)^\bullet"] \arrow[rr, shift left, "(\cdot)_*"] & & \mathsf{DGF} \arrow[uu, "U"'] \arrow[ll, shift left, "(\cdot)_*"]
\end{tikzcd}$$

Figure 2: *Duality and the Jónsson-Tarski representation theorem*

A Boolean algebra with an operator (BAO) is a pair $\langle \mathbb{A}, \Diamond \rangle$ where $\mathbb{A}$ is a Boolean algebra and $\Diamond : \mathbb{A} \to \mathbb{A}$ is a function which preserves all finite joins. A unary operator is called *complete* if it preserves all joins, so a 'CABACO' is a complete atomic Boolean algebra with a complete operator. KrFr denotes the category of Kripke frames (sets with a relation on the set) and DGF the category of descriptive general frames, which can also be seen as Stone spaces, equipped with a relation which is point-closed.[1]
The following theorem is now the classical starting point for the use of canonical extensions in the study of modal logic.

Jónsson-Tarski representation theorem (1952) [31]. *Any BAO $\mathbb{A}$ can be embedded in its canonical extension $\mathbb{A}^\sigma$.*

In the next section, we will discuss the relevance of canonical extensions for logic in general and modal logic in particular.

2 Canonical extensions and completeness

From a logical point of view, representation theorems and canonical extensions are relevant because they yield frame completeness results. Probably the simplest example of this phenomenon is that Stone's representation theorem can be used to prove the completeness of classical propositional calculus with respect to subset semantics.

[1]For a more comprehensive discussion, the reader is referred to [5].

We focus on a little bit more involved example from *modal logic.* Using the Jónsson-Tarski theorem outlined in the previous section, we can easily prove the completeness of the modal logic **K**. Say that $\mathbf{K} \vdash \phi$ iff the equation $\phi = \top$ is valid in all Boolean algebras with operators.
The definition of Kripke semantics immediately gives us the following lemma.

Lemma 2.1. *For any Kripke frame $\langle S, R\rangle$ and any modal formula ϕ, $\langle S, R\rangle \models \phi$ if and only if $\langle S, R\rangle^+ \models \phi = \top$.*

Using the Jónsson-Tarski representation theorem, we then easily prove the completeness of **K** with respect to the class of Kripke frames.

Theorem 2.2. *Let ϕ be a modal formula. Then $\mathbf{K} \vdash \phi$ if and only if ϕ is valid on all Kripke frames.*

Proof. If $\mathbf{K} \vdash \phi$, then in particular $\phi = \top$ holds on all BAO's of the form $\langle S, R\rangle^+$, so, by the Lemma, $\langle S, R\rangle \models \phi$ for all Kripke frames $\langle S, R\rangle$.
Conversely, suppose $\mathbf{K} \not\vdash \phi$, and let $\mathbb{A}$ be a BAO which falsifies the equation $\phi = \top$. Then, by the Jónsson-Tarski representation theorem, the algebra $\mathbb{A}^\sigma$ also falsifies the equation $\phi = \top$, since validity of equations is preserved under taking subalgebras. Hence, by the Lemma, $\mathbb{A}_\bullet \cong (\mathbb{A}^\sigma)_+$ is a frame falsifying ϕ. □

We hope that this simple example already convinces the reader of the power and relevance to logic of representation theorems like Stone's and Jónsson and Tarski's. However, this example can be enhanced in several ways. We can ask, for example, if we can give a similiar proof of the completeness of the logic **K4** (that is, the logic of the class of BAO's in which the inequality $\Diamond\Diamond p \leq \Diamond p$ holds) with respect to the transitive frames. An easy inspection of the proof of Theorem 2.2 shows that the proof carries over, provided that we can make an inference of the form

$$\text{If } \mathbb{A} \models \alpha \leq \beta, \text{ then } \mathbb{A}^\sigma \models \alpha \leq \beta.$$

Inequalities $\alpha \leq \beta$ with this property are called *canonical.* The second part of this thesis, Chapter 2, will be devoted, broadly speaking, to finding canonical formulas, in a more general context than that of BAO's. Another generalisation of this example, obtained by moving away from classical logic, will be introduced in the next section.

3 Beyond classical logic

In recent years, there has been an increasing amount of interest in logics for which the algebras are not Boolean algebras, but rather more general. For example, [19], [24], [21], [22], and [7] treat the topics of canonical extension and canonicity in a context where Boolean algebras are replaced by *distributive lattices*, building on Priestley's duality theory [38]. Later, this approach was generalised to *bounded lattices* which are no longer required to be distributive ([17], [10]), for which Urquhart [43] and Hartung [27], among others, (*cf.* Chapter 1 and the Conclusion for more details) had developed a duality theory. Finally, in [9] and [16], a theory of duality and canonical extension is developed for *partially ordered sets*, and reported on and applied to the study of substructural logics in [14]. The latter is the most general approach in this line of research so far.

Importantly, just in the same way in which classical modal logic was treated algebraically as an expansion of the Boolean signature, a vast number of lattice- or poset-based logics can be accounted for in a completely analogous way in the lattice or poset setting. As a result, the algebraic study of *canonicity* of formulas that was revived by Jónsson for classical modal logic [30] could be generalised to the context of distributive lattices, as was first done by Ghilardi and Meloni [24] in a constructive setting for intermediate and intuitionistic modal logics. The work that we will focus on in Chapter 2 of this thesis is that of Gehrke, Nagahashi and Venema [22], who prove the canonicity of a Sahlqvist fragment for distributive modal logic (i.e., a modal logic which is based on distributive lattices). One of the main common points of the works mentioned in this paragraph is that canonicity is studied independently of correspondence. In Chapter 2, we will investigate the possibility of extending this approach of 'proving canonicity independently of correspondence' from the Sahlqvist fragment to the inductive fragment for distributive modal logic, which was introduced and proved to be canonical *via* correspondence by Palmigiano and Conradie [7]. We will come back to the links between canonicity and correspondence in Section 3 of the Conclusion of this thesis.

Studying canonical extensions and duality in more general contexts than that of Boolean algebras is interesting and important, for at least three reasons. The first reason is that we thus obtain a *uniform* approach to the theory of duality and canonical extensions. Rather than having to develop a whole new theory whenever confronted with a unknown logic which is associated with a certain class of algebras, one can rely on the existing general theory. An example of such an application of the general theory of canonical extensions is [13], where the canonical extension for so-called *bi-implicative algebras* is defined as a direct application the theory of canonical extensions for posets from [9].

Related to this, moving beyond Boolean algebras is interesting because it enables us to get a *modular* picture of the theory of canonical extensions. We then regard Boolean algebras, for example, as partially ordered sets with a number of additional operations (i.e., $\wedge$, $\vee$, $\neg$, $\top$) which are required to satisfy certain equations. In this view, it becomes apparent that the majority of the algebraic structures mentioned so far (Boolean algebras, BAO's, distributive lattices, bounded lattices) can all be seen as 'poset expansions' of a certain type. This leads one to conclude that the *order* on an algebra is actually all that is required to define the canonical extension. Adding 'privileged' operations like $\wedge$, $\vee$ and $\neg$ then slightly simplifies the appearance of the canonical extension, however at the expense of partially blurring our view of what the canonical extension actually is: a certain completion of a partially ordered set. This perspective has recently been investigated in [23] and [18].

The two reasons mentioned so far were of a more theoretical and methodological nature, but there is also an application-driven motivation behind the project of studying canonical extensions and duality in non-classical logic: the project can provide powerful tools to obtain relational semantics for non-classical logics, such as, among others, distributive modal logic, logics based on semilattice expansions, and substructural logics. For instance, as mentioned above, Chapter 2 of this thesis will be devoted to applying the canonical extension of distributive lattices to the study of canonicity of formulas. Many more examples can be found in [14].

4 Outline of the thesis

In Chapter 1, we study the construction of the canonical extension of a partially ordered set. In particular, we will try to 'complete the square diagram', using a generalisation of the ultrafilter construction for Boolean algebras, inspired by Urquhart's maximal filter-ideal pairs construction for bounded lattices. We will show that the canonical extension can be realized as the complex algebra of a polarity that naturally arises from the maximal filter-ideal pairs. More specifically, we will identify the 'atoms and coatoms' of the polarity associated with the canonical extension as the *optimal* filters and ideals of the original poset, thus generalising the result that the atoms of the canonical extension $\mathbb{A}^\sigma$ of a Boolean algebra $\mathbb{A}$ may be realized as the ultrafilters of the original Boolean algebra.

After studying the object part for the discrete duality that we need in in this construction, we will focus on the morphism part of the discrete duality in Section 4 of Chapter 1. Morphisms are of particular importance from a logical perspective when we want to *add operations* to the logics we study.

An example of such an added operation is the operation called $\mathbf{n}$, which will play an important role in Chapter 2, where we will extensively discuss the canonicity proof of the Sahlqvist fragment for distributive modal logic by Gehrke, Nagahashi and Venema [22]. We will study and explain in detail some of the mechanisms underlying that canonicity proof, which is special because it is mainly *algebraic* in nature, and does not directly rely on a correspondence result. In the last section of that chapter, we explain why the proof method does not directly transfer to other syntactic fragments, such as the *inductive fragment* for distributive modal logic, introduced by Conradie and Palmigiano [7]. This led us to consider the operation $\mathbf{n}$ in more detail, from which some interesting similarities between $\mathbf{n}$ and the Ackermann Lemma emerged, on which we will report at the end of Chapter 2.

In the Conclusion, we will reflect on the thesis, and collect the open problems and possible directions for further research that we encountered in it.

Chapter 1

Canonical extensions of partially ordered sets

Existence and uniqueness of the canonical extension of partially ordered sets were proved by algebraic methods in [9]. There, the algebras which are canonical extensions of partially ordered sets were characterized as *perfect lattices*. Gehrke [16] developed a 'discrete duality' between perfect lattices and two-sorted structures called *RS polarities*. However, an analogue of the topological duality for Boolean algebras, or even of the 'ultrafilter frame' functor, is missing. Thus, the current situation can be expressed by the following diagram.

$$
\begin{tikzcd}
\mathsf{PLat} \arrow[r, shift left, "(\cdot)_+"] & \mathsf{RSPol} \arrow[l, shift left, "(\cdot)^+"] \\
\mathsf{Poset} \arrow[u, dashed, "(\cdot)^\sigma"] \arrow[ur, dashed] \arrow[r, shift left, "?"] & ? \arrow[l, shift left, "?"] \arrow[u, "U\ ?"']
\end{tikzcd}
$$

Figure 1.1: *Partial duality square for posets*

In this chapter, we start by reviewing the algebraic construction of the canonical extension (Section 1) and the discrete duality (Section 2). After that, we move on to 'expand' this diagram in the following two ways:

First, in Section 3, we will develop an analogue of the ultrafilter frame functor (the diagonal arrow in the diagram), generalizing work done by Urquhart [43] and Hartonas and Dunn [26] for lattices, also drawing inspiration from the MSc Logic thesis by Haim [25].

Secondly, in section 4, we will re-examine the categories and the duality in the upper half of the diagram. In particular, we aim to answer the question: What should be the morphisms in the categories PLat and RSPol? In order to give a well-balanced answer to this question, we start developing a *correspondence theory* for perfect lattices and RS polarities, which can be seen as a relatively new attempt to uniformly generalize correspondence theory for modal logic to a much wider setting.

1 Algebraic construction

In this section, we briefly review the algebraic construction of the canonical extension of a poset as given in [9]. While doing so, we will also encounter the definition of the canonical extension, which we parametrize by a choice of filters and ideals, since different definitions of 'filter' are possible in a poset.

1.1 Filter definitions

Because there are many reasonable choices for the definition of 'filter' in a poset, each having different advantages and disadvantages, we here list the properties such a definition must have in order for us to be able to define a canonical extension, and show that two standard definitions of 'filter' in a poset satisfy these properties.

Definition 1.1.1. Given a poset $\mathbb{P} = (P, \leq_{\mathbb{P}})$, we say a collection $\mathcal{F}(\mathbb{P}) \subseteq \mathcal{P}(P)$ is a **filter system** in $\mathbb{P}$ if it has the following four properties:

(i) $\emptyset \notin \mathcal{F}(\mathbb{P})$,

(ii) every $F \in \mathcal{F}(\mathbb{P})$ is upward closed,

(iii) for every $p \in P$, $p\!\uparrow \in \mathcal{F}(\mathbb{P})$,

(iv) $\mathcal{F}(\mathbb{P})$ is closed under unions of chains, i.e.

$$\text{If } \{F_\alpha\}_{\alpha \in A} \text{ is a } \subseteq\text{-chain of elements in } \mathcal{F}(\mathbb{P}) \text{ then } \bigcup_{\alpha \in A} F_\alpha \in \mathcal{F}(\mathbb{P}).$$

An assignment $\mathcal{F}$ which sends every poset $\mathbb{P}$ to a filter system $\mathcal{F}(\mathbb{P})$ is called a **filter definition**. Given a filter definition $\mathcal{F}$, we let the **ideal definition $\mathcal{I}$ associated with $\mathcal{F}$** for every poset $\mathbb{P}$ be $\mathcal{I}(\mathbb{P}) := \{F \in \mathcal{P}(P) : F \in \mathcal{F}(\mathbb{P}^\partial)\}$.
After fixing a filter definition $\mathcal{F}$, elements of $\mathcal{F}(\mathbb{P})$ are called **filters of** $\mathbb{P}$ and elements of the associated ideal definition $\mathcal{I}(\mathbb{P})$ are called **ideals of** $\mathbb{P}$.

Let us look at two filter definitions in particular.

Examples 1.1.2. (i) For a poset $\mathbb{P}$, let $\mathcal{F}(\mathbb{P})$ be the collection of non-empty down-directed up-sets of $\mathbb{P}$. We show that $\mathcal{F}$ is a filter definition, that is, every $\mathcal{F}(\mathbb{P})$ is a filter system.

Requirements (1) and (2) are met by definition of $\mathcal{F}(\mathbb{P})$.

Moreover, if $p \in P$, then $p\!\uparrow$ is clearly an up-set, and it is down-directed, for if $q_1, q_2 \in p\!\uparrow$, then $p \in p\!\uparrow$ and $p \leq q_1$ and $p \leq q_2$. So $p\!\uparrow \in \mathcal{F}(\mathbb{P})$, showing that requirement (3) holds.

To see that $\mathcal{F}(\mathbb{P})$ is closed under unions of chains, suppose $\{F_\alpha\}_{\alpha\in A}$ is a $\subseteq$-chain of down-directed up-sets, and let $F := \bigcup_{\alpha\in A} F_\alpha$.

Suppose $p \in F$ and $p \le q$, then $p \in F_\alpha$ for some $\alpha \in A$, so $q \in F_\alpha \subseteq F$. Hence, F is an up-set.

If $p_1, p_2 \in F$, then, because the F_α form a chain, there is some $\alpha \in A$ such that $p_i \in F_\alpha$ for both $i = 1, 2$. Since F_α is down-directed, there is $p \in F_\alpha \subseteq F$ such that $p \le p_1$ and $p \le p_2$. Hence, F is down-directed. We conclude that $F \in \mathcal{F}(\mathbb{P})$, as required.

(ii) For $\mathbb{P}$ a poset, let $\mathcal{F}(\mathbb{P})$ be the collection of non-empty up-sets F which are closed under existing binary meets, i.e., if $a, b \in F$ and $\{a, b\}$ has a greatest lower bound c in $\mathbb{P}$, then $c \in F$. We show again that $\mathcal{F}$ is a filter definition.

Again, (1) and (2) are met by definition.

Let $p \in P$. Then $p\uparrow$ is clearly a non-empty up-set. Suppose that $q_1, q_2 \in p\uparrow$ and their meet exists. Then $q_1 \wedge q_2 \ge p$ since both $q_i \ge p$, so $q_1 \wedge q_2 \in p\uparrow$. Therefore, $p\uparrow \in \mathcal{F}(\mathbb{P})$, which is requirement (3).

To see for requirement (4) that $\mathcal{F}(\mathbb{P})$ is closed under unions of chains, suppose $\{F_\alpha\}_{\alpha\in A}$ is a $\subseteq$-chain of down-directed up-sets, and let $F := \bigcup_{\alpha\in A} F_\alpha$.

By the same argument as in part (1), F is an up-set.

Let $p_1, p_2 \in F$ such that their meet exists in $\mathbb{P}$. Then, because the F_α form a chain, there is some $\alpha \in A$ such that $p_i \in F_\alpha$ for both $i = 1, 2$. Since F_α is closed under existing binary meets, we have $p_1 \wedge p_2 \in F_\alpha \subseteq F$. Hence, F is closed under existing meets, so we see that $F \in \mathcal{F}(\mathbb{P})$, as required.

From now on, we will always use the letter $\mathcal{F}$ to refer to an arbitrary, fixed, filter definition and the letter $\mathcal{I}$ to refer to its associated ideal definition.

1.2 Definition of the canonical extension

We will now follow the same route as in [34] and [13], and define the notion of an extension in the poset setting, parametrized by a filter definition $\mathcal{F}$. This generalizes the notion of canonical extension of a poset from [9], where the particular filter definition from Example 1.1.2(i) is used.
Let $\mathbb{P} = (P, \le_{\mathbb{P}})$ be a partially ordered set. Recall that a poset $\mathbb{Q} = (Q, \le_{\mathbb{Q}})$ is called an **extension** of $\mathbb{P}$ via η if there is an order embedding $\eta : \mathbb{P} \to \mathbb{Q}$. Furthermore, if $\mathbb{Q}$ is a complete lattice, then an extension $\mathbb{Q}$ of $\mathbb{P}$ is also called a **completion** of $\mathbb{P}$.
With the following definition, we put two (in general different) topologies on such an extension $\mathbb{Q}$.

Definition 1.2.1. Let $\mathbb{Q}$ be an extension of a poset $\mathbb{P}$ via η.
An element k in the extension $\mathbb{Q}$ is $\mathcal{F}(\mathbb{P})$**-closed** if there is a filter $F \in \mathcal{F}(\mathbb{P})$ such that $k = \bigwedge \eta[F]$. The set of closed elements of $\mathbb{Q}$ is denoted by $\mathcal{K}_{\mathcal{F}(\mathbb{P})}(\mathbb{Q})$, or simply $\mathcal{K}(\mathbb{Q})$, or even $\mathcal{K}$, when the other parameters are clear from the context.
Dually, an element j in $\mathbb{Q}$ is called $\mathcal{I}(\mathbb{P})$**-open** if there is an ideal $I \in \mathcal{I}(\mathbb{P})$ such that $j = \bigvee \eta[I]$, and the set of open elements is denoted by $\mathcal{O}_{\mathcal{I}(\mathbb{P})}(\mathbb{Q})$, $\mathcal{O}(\mathbb{Q})$ or simply $\mathcal{O}$.

We are now ready to state the two properties which will distinguish the canonical extension among other poset completions.

Definition 1.2.2. Let $\mathbb{P}$ be a poset and let $\mathbb{Q}$ be an extension of $\mathbb{P}$ via η.

(i) The extension $\mathbb{Q}$ of $\mathbb{P}$ is **dense** if, for every $u \in Q$:

$$\bigvee\{k \in \mathcal{K}(\mathbb{Q}) : k \leq u\} = u = \bigwedge\{j \in \mathcal{O}(\mathbb{Q}) : j \geq u\}.$$

(ii) The extension $\mathbb{Q}$ of $\mathbb{P}$ is **compact** if, for any $F \in \mathcal{F}(\mathbb{P})$ and $I \in \mathcal{I}(\mathbb{P})$, if $\bigwedge \eta[F] \leq \bigvee \eta[I]$, then $F \cap I \neq \emptyset$.

(iii) A **canonical extension** of $\mathbb{P}$ is a compact and dense completion of $\mathbb{P}$.

Note that these definitions very much depend on the chosen filter definition $\mathcal{F}$: different filter definitions give rise to different canonical extensions. For sake of readability, we suppress explicit reference to the filter definition, but we would like to note here that the concepts defined in Definition 1.2.2 should be properly named '$(\mathcal{F}, \mathcal{I})$-dense', '$(\mathcal{F}, \mathcal{I})$-compact' and '$(\mathcal{F}, \mathcal{I})$-canonical extension'.
In the next subsection, we will review the algebraic proof that, for any filter definition $\mathcal{F}$, the canonical extension exists uniquely.

1.3 Existence and uniqueness

We will merely sketch the algebraic existence and uniqueness proofs from [9] and [34], because one of the points of this chapter is that an alternative proof can be given by combining discrete duality and the optimal filter-construction. We start by sketching the existence proof.
First note that, by clause (iii) of the definition of a filter system (Definition 1.1.1), the assignment $p \mapsto p\!\uparrow$ defines an order embedding of $\mathbb{P}$ into the poset $\mathbb{F}(\mathbb{P}) = (\mathcal{F}(\mathbb{P}), \supseteq)$. So $\mathbb{F}$ is an extension of $\mathbb{P}$. Dually, $\mathbb{I}(\mathbb{P}) = (\mathcal{I}(\mathbb{P}), \subseteq)$ is an extension of $\mathbb{P}$ via the embedding $p \mapsto p\!\downarrow$.
The algebraic recipe for obtaining the canonical extension now runs as follows: first amalgamate the posets $\mathbb{F}$ and $\mathbb{I}$ into a compact and dense extension of $\mathbb{P}$, and then take the Dedekind-MacNeille (DM) completion of this amalgamation in order to obtain a compact and dense completion. It then remains to show that taking the DM-completion preserves the properties of compactness and denseness, which the amalgamation held.
As we will show in Section 1.4, the amalgamation could be simply defined as a pushout in the category Poset, of posets with monotone maps, and the DM completion as a left adjoint to a particular forgetful functor. Still, we think it is useful to outline the concrete, algebraic construction here.
Let $\mathcal{F}(\mathbb{P}) + \mathcal{I}(\mathbb{P})$ be the disjoint union of the sets $\mathcal{F}(\mathbb{P})$ and $\mathcal{I}(\mathbb{P})$. We define a preliminary relation R on $\mathcal{F}(\mathbb{P}) + \mathcal{I}(\mathbb{P})$ as the union

$$R := \leq_{\mathbb{F}(\mathbb{P})} \cup \leq_{\mathbb{I}(\mathbb{P})} \cup \{(F, I) : F \cap I \neq \emptyset\} \cup \{(I, F) : \forall p \in I, \forall p' \in F : p \leq_{\mathbb{P}} p'\}.$$

Lemma 1.3.1. *The relation R is a preorder on $\mathcal{F}(\mathbb{P}) + \mathcal{I}(\mathbb{P})$.*

Proof. R is clearly reflexive. The proof that R is transitive is elementary, and uses that every filter is upward closed, and every ideal is downward closed, which holds by clause (ii) of Definition 1.1.1. $\square$

Let $\sim$ be the equivalence relation on $\mathcal{F}(\mathbb{P}) + \mathcal{I}(\mathbb{P})$ defined by $x \sim y$ iff xRy and yRx, and let $\mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P}) := (\mathcal{F}(\mathbb{P}) + \mathcal{I}(\mathbb{P}))/\sim$ be the set of equivalence classes. We order this set by $[x] \leq_\oplus [y]$ iff there are $x' \in [x]$ and $y' \in [y]$ such that $x'Ry'$. This is an instance of the general construction to turn a preorder into a partial order. In the current case, we actually have a nice characterization of this equivalence relation.

Lemma 1.3.2. *$x \sim y$ if and only if $x = y$ or $\{x, y\} = \{p\!\uparrow, p\!\downarrow\}$ for some $p \in P$.*

Define the projected embeddings $j_F : \mathcal{F}(\mathbb{P}) \to \mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P})$ and $j_I : \mathcal{I}(\mathbb{P}) \to \mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P})$ by $f \mapsto [f]$ and $i \mapsto [i]$. The following lemma actually contains nothing more than the general fact that this quotient of a preorder is a partial order.

Lemma 1.3.3. *The relation $\leq_\oplus$ defines a partial order on $\mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P})$ such that j_F and j_I are order preserving maps.*

The amalgamated poset $\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P}) := \langle \mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P}), \leq_\oplus \rangle$ is close to being the canonical extension, except that it fails to be complete in general.

Proposition 1.3.4. *$\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P})$ is a compact and dense extension of $\mathbb{P}$.*

Proof. The proof rests on the fact that the closed elements of the extension $\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P})$ of $\mathbb{P}$ are precisely the classes of the form $[F]$, for $F \in \mathcal{F}(\mathbb{P})$, and the open elements are precisely the classes of the form $[I]$, for $I \in \mathcal{I}(\mathbb{P})$.
For denseness, it is then a matter of calculation, using only that filters are upward closed and that principal upsets are filters, to show that, any element of $\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P})$ is the join of the closed elements below it, and, dually, the meet of open elements above it. Compactness then follows almost directly from the definition of the order $\leq_\oplus$.
A detailed proof is given in Theorem 2.22 of [34]. $\square$

So, in order to obtain a canonical extension of $\mathbb{P}$, we need to make $\mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P})$ into a complete lattice. The easiest (and most economic) way to do this is by taking the Dedekind-MacNeille completion $\overline{\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P})}$ of $\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P})$. We then make use of the following fact.

Lemma 1.3.5. *If $\mathbb{Q}$ is a compact and dense extension of $\mathbb{P}$, then so is $\overline{\mathbb{Q}}$.*

The proof of this fact rests on the fact that the Dedekind-MacNeille completion preserves all joins and meets which already exist in $\mathbb{Q}$.
Finally, combining Proposition 1.3.4 and Lemma 1.3.5, we get to the end of our algebraic construction.

Theorem 1.3.6. *For any poset $\mathbb{P}$, $\overline{\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P})}$ is a canonical extension of $\mathbb{P}$.*

After successfully performing the algebraic construction to prove existence, the uniqueness proof consists of showing that any canonical extension of $\mathbb{P}$ will be isomorphic to the complete lattice $\overline{\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P})}$ just constructed. We only state the Theorem for completeness' sake, and refer the reader to Theorem 2.5 of [9] for the algebraic proof.

Theorem 1.3.7. *The canonical extension of* $\mathbb{P}$ *is unique up to an isomorphism fixing* $\mathbb{P}$.

We now also define the extension of *monotone maps* between posets, which we will be relevant for Chapter 2. This is precisely Definition 3.2 from [9].

Definition 1.3.8. Let $\mathbb{P}$ and $\mathbb{Q}$ be posets, and $f : \mathbb{P} \to \mathbb{Q}$ any function. Define f^σ and $f^\pi : \mathbb{P}^\sigma \to \mathbb{Q}^\sigma$, called the **sigma and pi extension of** f, by

$$f^\sigma(u) := \bigvee\{\bigwedge\{f(p) : x \le p, p \in P\} : x \le u, x \in \mathcal{K}(\mathbb{P}^\sigma)\},$$

$$f^\pi(u) := \bigwedge\{\bigvee\{f(p) : y \ge p, p \in P\} : y \ge u, y \in \mathcal{O}(\mathbb{P}^\sigma)\}.$$

1.4 A categorical perspective

[1] As mentioned above, both steps in the construction of Section 1.3, i.e., the amalgamation and the DM-completion, can be viewed from the perspective of category theory.
Regarding the amalgamation, although one might hope that it is a pushout in the category of posets with monotone maps, the following example shows that this is not the case in general.

Example 1.4.1. Consider the poset $\mathbb{P}$ whose underlying set is $\mathbb{N} \times \{0, 1\}$, with the order defined by $(n, 0) \le (m, 0)$ iff $n \le m$, $(n, 1) \le (m, 1)$ iff $m \le n$, and $(n, 0) \le (m, 1)$ for all n, m. We think of $\mathbb{P}$ as the partial order $\mathbb{N}$ with a mirrored copy above it.
Then the non-principal ideals of $\mathbb{P}$ are $\emptyset$ and $\{(n, 0) : n \in \mathbb{N}\}$, which we denote by ∞, and the non-principal filters of $\mathbb{P}$ are $\emptyset$ and $\{(n, 1) : n \in \mathbb{N}\}$, which we denote by ∞^*.
One can calculate that the amalgamation $\mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P})$ has as underlying set that of $\mathbb{P}$, with two additional points ∞ and ∞^*, and the order is that from $\mathbb{P}$, with the additional relations $(n, 0) \le \infty \le \infty^* \le (n, 1)$ for all n.
However, we can consider the poset $\mathbb{Q}$ which has the same underlying set and order, except that ∞ and ∞^* are not comparable in $\mathbb{Q}$, the 'legs' $\mathcal{F}(\mathbb{P}) \to \mathbb{Q}$ and $\mathcal{I}(\mathbb{P}) \to \mathbb{Q}$ are the obvious inclusion functions, which preserve the orders of $\mathcal{F}(\mathbb{P})$ and $\mathcal{I}(\mathbb{P})$.
Now there is clearly no order-preserving map from $\mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P}) \to \mathbb{Q}$ commuting with the legs of $\mathbb{Q}$, so $\mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{Q})$ is not the pushout.
It is not hard to see that the poset $\mathbb{Q}$ we defined above *is* the pushout in Poset. This is an instance of a general way to construct the pushout in Poset, that we do not discuss any further here.

However, the amalgamation does have another, rather delicate, universal property.

Proposition 1.4.2. *Let* $\mathbb{P}$ *be a poset,* $\mathbb{F}(\mathbb{P}) = \langle \mathcal{F}(\mathbb{P}), \supseteq \rangle$ *and* $\mathbb{I}(\mathbb{P}) = \langle \mathcal{I}(\mathbb{P}), \subseteq \rangle$ *the filter and ideal posets of* $\mathbb{P}$*, and* $\mathbb{F}(\mathbb{P}) \oplus \mathbb{I}(\mathbb{P})$ *the amalgamated poset constructed in Section 1.3. Suppose we have a diagram*

$$\begin{tikzcd} \mathbb{P} \arrow[r, "\iota_F"] \arrow[d, "\iota_I"'] & \mathbb{F}(\mathbb{P}) \arrow[d, "f"] \\ \mathbb{I}(\mathbb{P}) \arrow[r, "g"] & \mathbb{Q} \end{tikzcd}$$

[1]This section was heavily revised after the defense of this thesis, and benefited greatly from several discussions with Mai Gehrke and Dion Coumans.

where $\mathbb{Q}$ is a poset such that f preserves directed meets and g preserves directed joins (more precisely, for a down-directed set D in $\mathbb{F}(\mathbb{P})$, the directed meet of $f[D]$ exists and is equal to $f(\bigsqcap D)$, and similarly for g).
Then there exists a unique order-preserving function $h : \mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P}) \to \mathbb{Q}$ such that $h \circ j_F = f$ and $h \circ j_I = g$.

In the category Poset, the Dedekind-MacNeille completion of a poset $\mathbb{P}$ can be characterized categorically as the *injective hull* of the object $\mathbb{P}$ (*Cf.* [2] and 9.17 of [1] for details). More related to our work in Section 4 of this chapter, there is also a universal mapping characterization of the DM completion, provided by Bishop [4], which we will briefly review here[2].
Fix a poset $\mathbb{P}$. For $u \subseteq P$, let $\mathsf{u}_{\mathbb{P}}$ and $\mathsf{l}_{\mathbb{P}}$, defined by $\mathsf{u}_{\mathbb{P}}(u) := \{p' \in P : \forall p \in u\,(p \leq p')\}$ and $\mathsf{l}_{\mathbb{P}}(u) := \{p \in P : \forall p' \in u\,(p \leq p')\}$, be the **upper and lower maps** on $\mathcal{P}(P)$. These maps form a residuated pair, and their composition $\mathsf{c}_{\mathbb{P}} := \mathsf{l}_{\mathbb{P}}\mathsf{u}_{\mathbb{P}}$ is a closure operator on $\mathbb{P}$. We can use this closure operator to realize the DM completion of $\mathbb{P}$.

Theorem 1.4.3. *The complete lattice $\overline{\mathbb{P}}$, consisting of the subsets of $\mathbb{P}$ which are closed under $\mathsf{c}_{\mathbb{P}}$, is the Dedekind-MacNeille completion of $\mathbb{P}$ (up to an isomorphism fixing $\mathbb{P}$) with the order embedding $\eta_{\mathbb{P}} : \mathbb{P} \to \overline{\mathbb{P}}$ which sends p to $p{\downarrow}$. Moreover, $\eta_{\mathbb{P}}$ preserves any joins or meets which exist in $\mathbb{P}$.*

One may expect at first sight that $\eta_{\mathbb{P}}$ is a universal arrow in the category Poset, but this is not the case; we only have the following result.

Proposition 1.4.4. *For any monotone map $f : \mathbb{P} \to \mathbb{Q}$, where $\mathbb{P}$ is a poset and $\mathbb{Q}$ is a complete lattice, there exists a monotone map $\bar{f} : \overline{\mathbb{P}} \to \mathbb{Q}$ such that $\bar{f} \circ \eta_{\mathbb{P}} = f$. In general, $\bar{f}$ is not unique.*

If we want to get an actual universal arrow, we need to consider a subcategory of Poset by restricting the functions which we admit as morphisms in our category. The following definition is the key point in Bishop [4].

Definition 1.4.5. We say a map $f : \mathbb{P} \to \mathbb{Q}$ is **closed** if $f^{-1} : \mathcal{P}(Q) \to \mathcal{P}(P)$ sends $c_{\mathbb{Q}}$-closed sets to $c_{\mathbb{P}}$-closed sets.
More explicitly, f is closed if, for every subset $u \subseteq Q$ such that $c_{\mathbb{Q}}(u) = u$, we have that $c_{\mathbb{P}}(f^{-1}(u)) = f^{-1}(u)$.

It is not hard to see that residuated maps are closed, and that closed maps preserves arbitrary joins which exist in the domain. Then, Bishop has the following result.

Theorem 1.4.6 (Theorem 3, [4])**.** *For any closed map $f : \mathbb{P} \to \mathbb{Q}$ between posets, there exists a unique closed map $\bar{f} : \overline{\mathbb{P}} \to \overline{\mathbb{Q}}$ such that the diagram*

$$\begin{tikzcd} \mathbb{P} \arrow[r, "f"] \arrow[d, "\eta_{\mathbb{P}}"] & \mathbb{Q} \arrow[d, "\eta_{\mathbb{Q}}"] \\ \overline{\mathbb{P}} \arrow[r, "\bar{f}"] & \overline{\mathbb{Q}} \end{tikzcd}$$

[2]Although Bishop only considers completions of lattices, it is straightforward to see that the same results hold for posets, too.

commutes.

We observe the following immediate consequence of this theorem.

Corollary 1.4.7. *Let* Poset_C *be the category of posets with closed maps, and* $\mathsf{CLat}_\vee$ *the category of complete lattices with completely join-preserving functions. We have functors*

$$\begin{array}{rr} L : \mathsf{Poset}_C \to \mathsf{CLat}_\vee & U : \mathsf{CLat}_\vee \to \mathsf{Poset}_C \\ \mathbb{P} \mapsto \overline{\mathbb{P}} & \mathbb{C} \mapsto \mathbb{C} \\ (f : \mathbb{P} \to \mathbb{Q}) \mapsto (\bar{f} : \overline{\mathbb{P}} \to \overline{\mathbb{Q}}) & f \mapsto f. \end{array}$$

These functors $L : \mathsf{Poset}_C \leftrightarrows \mathsf{CLat}_\vee : U$ *form a reflection, i.e., an adjunction whose counit is an isomorphism.*

Proof. The unit of the adjunction $\eta_{\mathbb{P}}$ is the embedding of $\mathbb{P}$ into $\overline{\mathbb{P}}$, which is a closed map. The co-unit is the isomorphism between the complete lattices $\overline{U\mathbb{C}}$ and $\mathbb{C}$, which is given by sending a $c_{U\mathbb{C}}$-closed set u to $\bigvee_{\mathbb{C}} u$. $\square$

With this Corollary, we have represented the Dedekind-MacNeille completion categorically as the reflector of the the subcategory $\mathsf{CLat}_\vee$ of Poset_C. Erné [11] later showed that the category CLat of complete lattices with complete lattice homomorphisms is also reflective in a category whose objects are pre-ordered sets, and whose morphisms are so-called *cut-stable* maps. It is clear that the results from Gehrke [16] and our own results in Section 4 of this chapter are closely related to Erné's and Bishop's work, and we think that making these connections explicit would make an interesting and necessary further research project.
Furthermore, we would like to reconsider Proposition 1.4.2 in this context. It is not clear at this point in what category $\mathcal{F}(\mathbb{P}) \oplus \mathcal{I}(\mathbb{P})$ could be a pushout, if in any category at all.
We conclude that, in order to simultaneously capture the two parts of the canonical extension construction (amalgamation and completion) in a categorical framework, and hence give a coherent category-theoretic perspective on the canonical extension, we would need a subcategory of Poset for which both the Dedekind-MacNeille completion is a reflector, and the diagram from Proposition 1.4.2 is a pushout. As far as we know, such a subcategory is not presently known.

2 Discrete duality: Objects

Whereas the construction outlined in the previous section took place on the algebraic (left) side of the square diagram 1.1, our first aim in this chapter is to 'reconstruct' the anti-clockwise route through the diagram, which was how the canonical extension was first defined in the context of Boolean algebras and BAO's (*cf.* Introduction).
As mentioned at the beginning of this chapter, the *discrete duality* in the upper half of the diagram was proposed in this context in [16], based on earlier work by Hartung [27] and Erné [11]. We review the object part of this duality in this section, and then show how to obtain the canonical extension via the discrete duality in section 3, deferring the discussion of morphisms to section 4.

2.1 Perfect Lattices and Polarities

In order to reconstruct the dual object of the canonical extension $\mathbb{P}^\sigma$ of a poset $\mathbb{P}$, it is a necessary first step to investigate the algebraic properties of the complete lattice $\mathbb{P}^\sigma$ in some more detail. Let us first try to motivate this investigation and the ensuing choice of a dual object for $\mathbb{P}^\sigma$ by considering the more specialized cases where $\mathbb{P}$ is a Boolean algebra, and where $\mathbb{P}$ is a distributive lattice.

Recall that the canonical extension $\mathbb{A}^\sigma$ of a Boolean algebra $\mathbb{A}$ is always *complete* and *atomic*. This observation then leads one to take the *atom set* of $\mathbb{A}^\sigma$ as the dual object of the algebra, because the entire structure of any complete and atomic Boolean algebra is already determined by its set of atoms.

A similar process can be seen at work for the canonical extension $\mathbb{D}^\sigma$ of a distributive lattice $\mathbb{D}$: here, one first observes that $\mathbb{D}^\sigma$ is always *complete*, and also join-generated by its completely join-prime elements and meet-generated by its completely meet-prime elements. Then the sets of completely join-prime and completely meet-prime elements turn out to be isomorphic because of the distributivity of $\mathbb{D}$, and it suffices to take the *poset*[3] *of completely join-primes* as the dual object of $\mathbb{D}^\sigma$. Note that, in Boolean algebras, the completely join-prime elements are precisely the same as the atoms, and the order of the Boolean algebra restricted to the atoms is always discrete. So, we can regard the canonical extension of a Boolean algebra as a special case of the construction for distributive lattices.

Moving back to the general context of the canonical extension $\mathbb{P}^\sigma$ of a partially ordered set $\mathbb{P}$, we find that the symmetry between the sets of completely join- and meet-irreducibles that was present in distributive lattices is lost, so that we will have to keep track of both sets on the dual side. The dual object of $\mathbb{P}^\sigma$ will thus become a two-sorted object. To encode the order of $\mathbb{P}^\sigma$ on the dual side, it suffices to remember only the order between the completely join- and meet-irreducibles, from which the rest of the order in $\mathbb{P}^\sigma$ can be reconstructed. Now, a remark analogous to the one at the end of the previous paragraph can be made: in a distributive lattice, the completely join-/meet-irreducible elements are precisely the same as the completely join-/meet-prime elements, so that the canonical extension of a distributive lattice can be regarded as a special case of the canonical extension of a poset.

We now formalise the ideas sketched here, closely following [16].

Definition 2.1.1. An element x in a complete lattice $\mathbb{L}$ is called **completely join-irreducible** if, for any $A \subseteq L$, $x = \bigvee A$ implies $x \in A$. The dual property is called **completely meet-irreducible**. The sets of completely join- and meet-irreducible elements of $\mathbb{L}$ are denoted by $J^\infty(\mathbb{L})$ and $M^\infty(\mathbb{L})$, respectively.

A complete lattice $\mathbb{L}$ is called **perfect** if it is join-generated by its completely join-irreducibles, and meet-generated by its completely meet-irreducibles. That is, $\mathbb{L}$ is perfect if for any $u \in L$, we have

$$\bigvee\{x \in J^\infty(\mathbb{L}) : x \leq u\} = u = \bigwedge\{a \in M^\infty(\mathbb{L}) : a \geq u\}.$$

Lemma 2.1.2. *If $\mathbb{P}$ is a poset, then $\mathbb{P}^\sigma$ is a perfect lattice.*

[3]It is interesting and relevant to duality theory that we get an *ordered* set as the dual of a distributive lattice, whereas the dual of a Boolean algebra is simply a set. However, we do not go into this any further here, because it is not central to the thread of our story.

Proof. By denseness of the closed elements in $\mathbb{P}^\sigma$, it suffices to show that any closed element of $\mathbb{P}^\sigma$ is equal to the join of the completely join-irreducible elements below it (and that the dual statement holds for the open elements). The proof of this fact uses the denseness of the open elements in $\mathbb{P}^\sigma$ and Zorn's Lemma. See the proof of Theorem 2.8(3) in [9] for the details. $\square$

We are now ready to set up our objects on the dual side. By our reasoning above, these objects should consist of two sets with a relation between them.

Definition 2.1.3. A **polarity** is a triple $\mathbb{F} = \langle X, A, R \rangle$, where X and A are sets, and $R \subseteq X \times A$ is a relation from X to A.[4]

Remark 2.1.4. The polarity object has been studied in full generality from several different perspectives, of which we want to mention a few here.
Firstly, polarities form the central object of study in Formal Concept Analysis, where a polarity is called a 'context'. The field of Formal Concept Analysis focuses on applications, of which there are many, ranging from linguistics to economics [15].
A polarity can also be regarded as a special case of a Chu space [37]: in that field, a polarity is called a $\{0, 1\}$-Chu space. We will refer to the latter observation when discussing morphisms in Section 4 of this chapter.

Since the structure of a perfect lattice is determined by its completely join- and meet-irreducibles, there is a natural way to encode a perfect lattice as a polarity:

Definition 2.1.5. Let $\mathbb{L}$ be a perfect lattice. The **polarity associated with** $\mathbb{L}$, denoted $\mathbb{L}_+$, is defined to be the triple $\langle J^\infty(\mathbb{L}), M^\infty(\mathbb{L}), \leq_\mathbb{L} \cap J^\infty(\mathbb{L}) \times M^\infty(\mathbb{L}) \rangle$.

A natural question to ask, now that we have defined a 'dual object' for a perfect lattice, is the following: how can we retrieve a perfect lattice $\mathbb{L}$ from its associated polarity $\mathbb{L}_+$? To answer this question, we need to recall how a polarity gives rise to a closure operator in a natural way.

Definition 2.1.6. If $\mathbb{F} = \langle X, A, R \rangle$ is a polarity, define functions u_R (upper) and l_R (lower) between the posets $\langle \mathcal{P}(X), \supseteq \rangle$ and $\langle \mathcal{P}(Y), \subseteq \rangle$, as follows:

$$\text{for } u \in \mathcal{P}(X),\ \text{let } \mathsf{u}_R(u) := \{a : \forall x \in u : xRa\}$$

$$\text{for } v \in \mathcal{P}(A),\ \text{let } \mathsf{l}_R(v) := \{x : \forall a \in v : xRa\}.^5$$

Then $(\mathsf{u}_R, \mathsf{l}_R)$ is a residuated pair, so $\mathsf{c}_R := \mathsf{l}_R \mathsf{u}_R$ is a closure operator on the set X, and we call a subset $u \subseteq X$ c_R**-closed**[6] if $\mathsf{c}_R(u) = u$. We write $\mathbb{F}^+$ for the complete lattice of subsets of X which are closed under the closure operator c_R. Recall that the lattice operations on $\mathbb{F}^+$ are given by $\bigwedge U = \bigcap U$ and $\bigvee U = \mathsf{c}_R(\bigcup U)$, for any set U of c_R-closed subsets.

[4] For notational convenience, especially in Section 4, we will denote variables which vary over the set X by x, x', and variables varying over the set A by a, a', and then sometimes abbreviate expressions such as "$\forall x \in X\ \exists a \in A$" as simply "$\forall x \exists a$".

[5] For sake of readability, if $x \in X$ and $a \in A$, we will often write $\mathsf{u}_R(x)$ instead of $\mathsf{u}_R(\{x\})$ and $\mathsf{l}_R(a)$ instead of $\mathsf{l}_R(\{a\})$.

[6] Note that, unfortunately, we need two completely different uses of the word 'closed': on the one hand 'closed' is the property of an element of the canonical extension that can be written as the meet of a filter, and on the other hand it is the word for a subset that is invariant under the closure operator c_R. To avoid confusion, we will always use the term 'c_R-closed' when we mean closed in the latter sense.

Remark 2.1.7. Of course, the definitions of the upper and lower maps and the associated closure operator $\mathsf{c}_{\mathbb{P}}$ that we considered in Section 1.4 are the particular case of this definition where $X = P = A$ and $R = \leq_{\mathbb{P}}$.

Lemma 2.1.8. *If* $\mathbb{L}$ *is a perfect lattice, then* $\mathbb{L} \cong (\mathbb{L}_+)^+$.

Proof. Since $\mathbb{L}$ is perfect, $\mathbb{L}$ is the DM completion of the poset $J^\infty(\mathbb{L}) \cup M^\infty(\mathbb{L})$. It is not hard to see that $(\mathbb{L}_+)^+$ is *also* the DM completion of the poset $J^\infty(\mathbb{L}) \cup M^\infty(\mathbb{L})$. The result then follows from the uniqueness part of Theorem 1.4.3 above. *Cf.* Proposition 4.7(1) in [9]. □

2.2 RS Polarities

Lemma 2.1.8 gives us half of the object correspondence between perfect lattices and polarities. We are now left with the question: which polarities arise as the dual objects of perfect lattices? Such polarities must have certain special properties, expressing the fact that $(\cdot)^+$ turns these polarities into perfect lattices. In this subsection, we will express these properties without any reference to the perfect lattices.
For any polarity $\mathbb{F} = \langle X, A, R\rangle$, we let $i_X : X \to \mathbb{F}^+$ be the map $x \mapsto \mathsf{c}_R(x)$ and $i_A : A \to \mathbb{F}^+$ the map $a \mapsto \mathsf{l}_R(a)$. For $\mathbb{F}$ to be of the form $\mathbb{L}_+$ for some perfect lattice $\mathbb{L}$, one first requirement is that both maps i_X and i_A are injective, so that we can regard both X and A as subsets of $\mathbb{F}^+$.

Lemma 2.2.1. *Let* $\mathbb{F} = \langle X, A, R\rangle$ *be a polarity.*

(i) The following are equivalent:

(a) i_X *is injective,*

(b) $\forall x_1, x_2 \in X \ (x_1 \neq x_2 \to \mathsf{u}_R(x_1) \neq \mathsf{u}_R(x_2))$. (S_X)

(ii) The following are equivalent:

(a) i_A *is injective,*

(b) $\forall a_1, a_2 \in A \ (a_1 \neq a_2 \to \mathsf{l}_R(a_1) \neq \mathsf{l}_R(a_2))$. (S_A)

Proof. (i) Suppose i_X is injective. Let $x_1, x_2 \in X$, $x_1 \neq x_2$. If we would have $\mathsf{u}_R(x_1) = \mathsf{u}_R(x_2)$, then clearly $i_X(x_1) = \mathsf{c}_R(x_1) = \mathsf{l}_R\mathsf{u}_R(x_1) = \mathsf{l}_R\mathsf{u}_R(x_2) = \mathsf{c}_R(x_2) = i_X(x_2)$, contrary to the assumption. So $\mathsf{u}_R(x_1) \neq \mathsf{u}_R(x_2)$.

Conversely, suppose (S_X) holds. Let $x_1, x_2 \in X$, $x_1 \neq x_2$. By (S_X) we get that $\mathsf{u}_R(x_1) \neq \mathsf{u}_R(x_2)$. Now $\mathsf{c}_R(x_1) = \mathsf{c}_R(x_2)$ would imply $\mathsf{u}_R(x_1) = \mathsf{u}_R\mathsf{c}_R(x_1) = \mathsf{u}_R\mathsf{c}_R(x_2) = \mathsf{u}_R(x_2)$, which is a contradiction. So $\mathsf{c}_R(x_1) \neq \mathsf{c}_R(x_2)$.

(ii) is the definition of injectivity. □

Definition 2.2.2. We say a polarity is **separating**, or an **S polarity**, if it satisfies conditions (S_X) and (S_A) of Lemma 2.2.1. In this case, we may thus regard X and A as subsets of $\mathbb{F}^+$, and we then write $\leq$ for the partial order on the set $X \cup A$ inherited from $\mathbb{F}^+$.

An important, well-known fact about polarities is that $\mathbb{F}^+$ is determined by the way in which X and A sit inside $\mathbb{F}^+$.

Proposition 2.2.3. *Let $\mathbb{F} = \langle X, A, R\rangle$ be an S polarity. Then X join-generates $\mathbb{F}^+$ and A meet-generates $\mathbb{F}^+$.*

Proof. In fact, for *any subset* u of X, we have the equality

$$\mathsf{c}_R(u) = \mathsf{c}_R\left(\bigcup\{\mathsf{c}_R(x) : x \in u\}\right), \tag{1.1}$$

which we will prove first.
Since $u \subseteq \{\mathsf{c}_R(x) : x \in u\}$ and c_R is monotone, the direction $\subseteq$ is clear.
For the direction $\supseteq$, it suffices to show $\mathsf{u}_R(u) \subseteq \mathsf{u}_R(\bigcup\{\mathsf{c}_R(x) : x \in u\})$, since l_R is antitone.
Let $a \in \mathsf{u}_R(u)$ be arbitrary. Take any $x' \in \bigcup\{\mathsf{c}_R(x) : x \in u\}$. Pick $x \in u$ such that $x' \in \mathsf{c}_R(x)$. Then xRa since $a \in \mathsf{u}_R(u)$. In particular, we then have $a \in \mathsf{u}_R(x)$, so $x'Ra$, since $x' \in \mathsf{c}_R(x)$. Since x' was arbitrary, this shows that $a \in \mathsf{u}_R(\bigcup\{\mathsf{c}_R(x) : x \in u\})$, as required.
Now, by definition of the operations in $\mathbb{F}^+$ and the embedding i_X of X in $\mathbb{F}^+$, for any c_R-*closed* subset u of X, (1.1) specializes to $u = \mathsf{c}_R(u) = \bigvee\{i_X(x) : x \leq u\}$. The proof that A meet-generates $\mathbb{F}^+$ is similar. □

We now formulate a second property of polarities, which isolates those polarities among the S polarities in which X really fulfills the role of 'set of completely join-irreducibles', and A the role of 'set of completely meet-irreducibles'.

Lemma 2.2.4. *Let $\mathbb{F} = \langle X, A, R\rangle$ be an S polarity.*

(i) The following are equivalent.

(a) The elements of X are completely join-irreducible in $\mathbb{F}^+$,

(b) $\forall x \in X\, \exists a \in A (x \not\leq a \wedge \forall x' \in X\, (x' < x \to x' \leq a)).$ $\qquad (R_X)$

(ii) The following are equivalent.

(a) The elements of A are completely meet-irreducible in $\mathbb{F}^+$,

(b) $\forall a \in A\, \exists x \in X (a \not\geq x \wedge \forall a' \in A\, (a' > a \to a' \geq x)).$ $\qquad (R_A)$

Proof. We only prove the equivalence in (i), the proof of (ii) being analogous.
Suppose that the elements of X are completely join-irreducible in $\mathbb{F}^+$. Let $x \in X$ be arbitrary. Consider the element $u := \bigvee\{x' : x' < x\}$ of $\mathbb{F}^+$, which is clearly less than or equal to x. If we had that $u = x$, then the join-irreducibility of x would imply that $x \in \{x' : x' < x\}$, which is clearly absurd. So we must have $u < x$. Since A meet-generates $\mathbb{F}^+$ (Proposition 2.2.3), pick an element $a \in A$ such that $x \not\leq a$ and $u \leq a$. If $x' < x$, then $x' \leq u \leq a$. So a is the element required to exist by (R_X).
Conversely, suppose that (R_X) holds. Let $x \in X$ be arbitrary. Suppose that $x = \bigvee U$ for some collection U of elements of $\mathbb{F}^+$. Assume, to obtain a contradiction, that $x \notin U$, so that $u < x$ for all $u \in U$. By (R_X), pick $a \in A$ such that $x \not\leq a$ and, for all $x' < x$, $x' \leq a$.
Let $u \in U$ be arbitrary. Then $u = \bigvee\{x' \in X : x' \leq u\}$, since X join-generates $\mathbb{F}^+$. Now, for any x' with $x' \leq u$, we have $x' < x$, so $x' \leq a$ by the choice of a. Since this holds for arbitrary x', we get $u \leq a$. However, this holds for arbitrary $u \in U$, so we conclude $x = \bigvee U \leq a$, contrary to the choice of a. □

Definition 2.2.5. An S polarity is called **reduced**, or an **RS polarity**, if it satisfies both conditions (R_X) and (R_A) of the Lemma.

The work we have done in this section is meant to motivate the following result.

Fact 2.2.6. *The assignments* $\mathbb{L} \mapsto \mathbb{L}_+$ *and* $\mathbb{F} \mapsto \mathbb{F}^+$ *define a bijective correspondence between perfect lattices and RS polarities.*

Strictly speaking, because we do not yet have a proper notion of 'isomorphism' between RS polarities, we can not prove a statement of the form 'If $\mathbb{F}$ is an RS polarity, then $\mathbb{F} \cong (\mathbb{F}^+)_+$.' We do hope to have convinced the reader, however, that such a statement, when made precise, is now within reach.

We have thus defined the objects of a category that will be dual to PLat and take up the place in the upper right corner of Diagram 1.1. In the next section, we will show how the obtain the canonical extension object via RS polarities. We will then return to the question of what are the morphisms in the category RSPol in Section 4.

3 Canonical extension via discrete duality

As indicated in the Introduction, Stone [41], Priestley [38], Urquhart [43] and Hartung [27] are concerned with developing topological dualities for increasingly general algebraic structures. Stone's duality establishes a correspondence between the categories of Boolean algebras and of Stone spaces. Priestley duality extends the correspondence to the categories of distributive lattices and of Priestley spaces, and Urquhart and Hartung again generalize this to the categories of bounded lattices and of L-spaces (Urquhart), or topological contexts (Hartung). In the last step of generalization, however, both Urquhart and Hartung needed to consider the subcategory of bounded lattices with *surjective* homomorphisms. More recent attempts by Hartonas and Dunn [26] and Jipsen and Moshier [28], [29] follow a different route to obtain a duality between bounded lattices with *all* homomorphisms and an appropriate subcategory of the category of topological spaces. However, this duality is not a generalisation of Stone duality. Our motivation here stems from logic rather than algebra, so that we want to remain consistent with Stone duality.

The algebraic objects needed for the mentioned dualities are of increasing complexity: Stone duality uses *ultrafilters*, Priestley duality needs *prime filters*, and Urquhart's bounded lattice duality rests on *maximal filter-ideal pairs.* A natural next step would be to extend this duality to the category of partially ordered sets. In this section, as a first part of this larger project, we show how to construct the functor $(\cdot)_\bullet$, which will be a straightforward generalization of the optimal filter-ideal functor for bounded lattices, in such a way that composing it with the discrete duality functor $(\cdot)^+$ introduced in the previous section yields the canonical extension.

More concretely, we will show that Urquhart and Hartung's optimal filter-ideal construction still works for partially ordered sets, no matter what filter definition we choose, and that it yields the canonical extension. We thus obtain a second construction of the canonical extension, uniformly generalizing the canonical-extension-via-duality approach for lattices to the setting of posets.

3.1 Definition and fundamental theorem of maximal pairs

Definition 3.1.1. Let F be a filter and I an ideal in a poset $\mathbb{P}$. We say F is **I-maximal** if it is a $\subseteq$-maximal element of the set $\{F' \in \mathcal{F}(\mathbb{P}) : F' \cap I = \emptyset\}$. Dually, we say I is **F-maximal** if it is a $\subseteq$-maximal element of the set $\{I' \in \mathcal{I}(\mathbb{P}) : F \cap I' = \emptyset\}$. The pair $\langle F, I\rangle$ is called a **maximal pair** if F is I-maximal and I is F-maximal.
An **optimal filter**[7] is a filter F for which there exists an ideal I such that $\langle F, I\rangle$ is a maximal pair. The set of optimal filters is denoted by $\mathcal{F}_*(\mathbb{P})$. Dually, we define an **optimal ideal** to be an ideal I for which there exists a filter F such that $\langle F, I\rangle$ is a maximal pair, and denote the set of optimal ideals by $\mathcal{I}_*(\mathbb{P})$.

We now generalize some of the theory from bounded lattices to the context of posets.

Theorem 3.1.2. *For any filter F and ideal I in a poset $\mathbb{P}$ such that $F \cap I = \emptyset$, there is a filter $F_* \supseteq F$ such that F_* is I-maximal, and, dually, there is an ideal $I_* \supseteq I$ such that I_* is F-maximal. Moreover, $\langle F_*, I_*\rangle$ is a maximal pair.*

Proof. Let F be a filter and I an ideal such that $F \cap I = \emptyset$. Consider the collection

$$\mathcal{X} := \{F' \in \mathcal{F}(P) : F \subseteq F' \text{ and } F' \cap I = \emptyset\}.$$

Since the union of a chain of filters is still a filter, by property (iv) of filter systems, we have that every chain of elements of $\mathcal{X}$ has an upper bound, namely, the union of the chain. Therefore, by Zorn's Lemma, there is a $\subseteq$-maximal element F_* in $\mathcal{X}$.
Now, if we would have $F_* \subsetneq F'$ for some $F' \in \mathcal{F}(\mathbb{P})$ such that $F' \cap I = \emptyset$, then also $F \subseteq F_* \subseteq F'$, so $F' \in \mathcal{X}$, contradicting that F_* is $\subseteq$-maximal in $\mathcal{X}$. Hence, F_* is I-maximal.
By a dual argument, we obtain an ideal $I_* \supseteq I$ which is F_*-maximal.
It remains to show that $\langle F_*, I_*\rangle$ is a maximal pair. We will only show that F_* is I_*-maximal, the other property being dual.
Suppose $F_* \subsetneq F'$ for some $F' \in \mathcal{F}(\mathbb{P})$ and $F' \cap I_* = \emptyset$. Then we also have $F' \cap I \subseteq F' \cap I_* = \emptyset$, and $F \subseteq F_* \subsetneq F'$, so $F' \in \mathcal{X}$, contradicting that F_* is $\subseteq$-maximal in $\mathcal{X}$. Hence, F_* is I_*-maximal. □

In particular, this theorem shows that *every disjoint filter-ideal pair can be extended to a maximal pair*, which can be viewed as a generalization of the well-known fact in Boolean algebras that *every filter can be extended to an ultrafilter*.
An important corollary, that we use more often than the theorem itself, is the following:

Corollary 3.1.3. *Let F be a filter and $p \notin F$. There is a maximal pair $\langle F_*, I_*\rangle$ such that $F \subseteq F_*$ and $p \in I_*$.*
Dually, for an ideal I and $q \notin I$, there is a maximal pair $\langle F_*, I_*\rangle$ such that $I \subseteq I_*$ and $q \in F_*$.

Proof. Let F be a filter and $p \notin F$. Then $p{\downarrow}$ is an ideal. Also, $F \cap p{\downarrow} = \emptyset$: if there was $p' \in P$ such that $p' \in F$ and $p' \leq p$, then we would have $p \in F$, because F is an up-set, contradicting the assumption that $p \notin F$.
By the Theorem, there is a maximal pair $\langle F_*, I_*\rangle$ such that $F \subseteq F_*$ and $p{\downarrow} \subseteq I_*$. In particular, $p \in I_*$.
The proof of the second part is dual. □

[7] This terminology was introduced, in the context of lattices, in the MSc Logic thesis of Haim [25].

3.2 The canonical extension via maximal pairs

There is an obvious way to put optimal filters and ideals together in a polarity.

Definition 3.2.1. Let $\mathbb{P}$ be a poset. Recall that $\mathcal{F}_*(\mathbb{P})$ and $\mathcal{I}_*(\mathbb{P})$ denote the sets of optimal filters and ideals of $\mathbb{P}$. We define the **incidence relation** $\perp \subseteq \mathcal{F}_*(\mathbb{P}) \times \mathcal{I}_*(\mathbb{P})$ by

$$F \perp I \iff F \cap I \neq \emptyset,$$

and we call $\mathbb{P}_\bullet := \langle \mathcal{F}_*(\mathbb{P}), \mathcal{I}_*(\mathbb{P}), \perp \rangle$ the **optimal polarity** of $\mathbb{P}$.

We now aim to show that applying the discrete duality functor $(\cdot)^+$ to the optimal polarity of $\mathbb{P}$ yields the canonical extension. In brief: we want to show that $\mathbb{P}^\sigma \cong (\mathbb{P}_\bullet)^+$.
We first show that $(\mathbb{P}_\bullet)^+$ is always a completion of $\mathbb{P}$ in a natural way. Recall that $(\mathbb{P}_\bullet)^+$, by definition, consists of sets of optimal filters which are closed under the closure operator $\mathsf{c}_\perp$ (this is Definition 2.1.6, applied to $R := \perp$).

Proposition 3.2.2. *The assignment* $p \mapsto \eta(p) := \{F \in \mathcal{F}_*(\mathbb{P}) : p \in F\}$ *defines an order embedding* $\eta : \mathbb{P} \to (\mathbb{P}_\bullet)^+$.

Proof. To see that η is well-defined, we need to show that $\eta(p)$ is a $\mathsf{c}_\perp$-closed set, i.e., that $\eta(p) = \mathsf{c}_\perp(\eta(p))$.
First note that $\eta(p) \subseteq \mathsf{c}_\perp(\eta(p))$ always holds, since $\mathsf{c}_\perp$ is a closure operator.
For the other inclusion, suppose that $F \in \mathcal{F}_*(\mathbb{P})$ is such that $F \notin \eta(p)$, i.e., $p \notin F$. By Corollary 3.1.3, there is a maximal pair $\langle F_*, I_* \rangle$ such that $F \subseteq F_*$ and $p \in I_*$. Note that $I_* \in \mathsf{u}_\perp(\eta(p))$: for any $F' \in \eta(p)$, we have $p \in F'$, so, since $p \in I_*$, we get $F' \perp I_*$. However, $F \cap I_* \subseteq F_* \cap I_* = \emptyset$, so $F \not\perp I_*$, so $F \notin \mathsf{c}_\perp(\eta(p))$.
So η is well-defined, and it remains to show that it is an order embedding. Clearly, if $p_1 \leq p_2$, then $\eta(p_1) \subseteq \eta(p_2)$, because any filter containing p_1 will then also contain p_2, since filters are upward closed.
For the other direction, suppose $p_1 \not\leq p_2$. Then $F := p_1\uparrow$ is a filter and $p_2 \notin F$ by assumption, so by Corollary 3.1.3, there is a maximal pair $\langle F_*, I_* \rangle$ such that $p_1\uparrow \subseteq F_*$ and $p_2 \in I_*$. In particular, $p_1 \in F_*$ and $p_2 \notin F_*$, so that F_* is in the set $\eta(p_1)$ but not in $\eta(p_2)$. Therefore, $\eta(p_1) \not\subseteq \eta(p_2)$, which is what we wanted to show. ☐

We are now ready to prove the main result of this section.

Theorem 3.2.3. *The complete lattice* $(\mathbb{P}_\bullet)^+$ *is the canonical extension of* $\mathbb{P}$.

Proof. We need to show that $\eta : \mathbb{P} \to (\mathbb{P}_\bullet)^+$ is a compact and dense extension.

- *Compactness.*

  Let $F \in \mathcal{F}(\mathbb{P})$ and $I \in \mathcal{I}(\mathbb{P})$ and suppose that $F \cap I = \emptyset$. We show that $\bigwedge \eta[F] \not\leq \bigvee \eta[I]$.

  By Theorem 3.1.2, there is a maximal pair $\langle F_*, I_* \rangle$ such that $F \subseteq F_*$ and $I \subseteq I_*$. We will show that $F_* \in \bigwedge \eta[F]$ and $F_* \notin \bigvee \eta[I]$. This will conclude the proof of the compactness property, since F_* is then an element of $\mathcal{F}_*(\mathbb{P})$ witnessing that $\bigwedge \eta[F] \not\leq \bigvee \eta[I]$.

  Since $F \subseteq F_*$, we have $f \in F_*$ for all $f \in F$, i.e., $F_* \in \eta(f)$ for all $f \in F$. Hence, $F_* \in \bigcap_{f \in F} \eta(f) = \bigwedge \eta[F]$.

Note that I_* is an element of the set $\mathsf{u}_\perp(\bigcup \eta[I])$: for any $F' \in \bigcup \eta[I]$, there is $i \in I$ such that $F' \in \eta(i)$, i.e., $i \in F'$. But then we also have $i \in I_*$, since $I \subseteq I_*$, so $F' \cap I_* \neq \emptyset$, i.e, $F' \perp I$, as claimed.

So, since $F_* \cap I_* = \emptyset$, we conclude that $F_* \notin \mathsf{l}_\perp \mathsf{u}_\perp(\bigcup \eta[I]) = \bigvee \eta[I]$.

- *Denseness.*

  Let $u \in (\mathbb{P}_\bullet)^+$ be a $\mathsf{c}_\perp$-closed set of optimal filters.

  - Let $K_u := \{k \in \mathcal{K}((\mathbb{P}_\bullet)^+) : k \leq u\}$. We want to show that $u = \bigvee K_u$.

    It is clear that u is an upper bound for K_u, so we need to show that it is a (the) least upper bound. Suppose u' is an upper bound in $(\mathbb{P}_\bullet)^+$ for K_u. It suffices to show that $\mathsf{u}_\perp(u') \subseteq \mathsf{u}_\perp(u)$: from this it follows by the residuation property of $(\mathsf{u}_\perp, \mathsf{l}_\perp)$ that $u \subseteq \mathsf{l}_\perp \mathsf{u}_\perp(u') = u'$, which is what we need to show.

    Suppose $I \notin \mathsf{u}_\perp(u)$. Pick $F \in u$ such that $F \cap I = \emptyset$. Since u is $\mathsf{c}_\perp$-closed, the set $k := \{F' \in \mathcal{F}_*(\mathbb{P}) : F \subseteq F'\}$ must be a subset of u. Note also that $k = \bigwedge \eta[F]$, and F is a filter in $\mathbb{P}$, so k is a closed element in $(\mathbb{P}_\bullet)^+$ that is below u, i.e., $k \in K_u$. Since u' is an upper bound for K_u, we have $k \subseteq u'$. In particular, $F \in u'$, so since $F \cap I = \emptyset$, we have $I \notin \mathsf{u}_\perp(u')$.

  - Let $O_u := \{j \in \mathcal{O}((\mathbb{P}_\bullet)^+) : u \leq j\}$. We want to show that $u = \bigwedge O_u = \bigcap O_u$. It is clear that u is contained in $\bigcap O_u$. For the converse inclusion, suppose that $F \notin u$ for some $F \in \mathcal{F}_*(\mathbb{P})$. We need to show that there is an open element above u such that $F \notin u$.

    Since u is $\mathsf{c}_\perp$-closed, we also have $F \notin \mathsf{l}_\perp \mathsf{u}_\perp(u)$. Pick $I \in \mathsf{u}_\perp(u)$ such that $F \cap I = \emptyset$. Let j be the open element $\bigvee \eta[I]$. The following calculation shows that $j = \mathsf{l}_\perp(\{I\})$:

$$
\begin{aligned}
j &= \mathsf{c}_\perp\left(\bigcup \eta[I]\right) \\
&= \mathsf{c}_\perp(\{F : \exists i \in I : F \in \eta(i)\}) \\
&= \mathsf{c}_\perp(\{F : F \cap I \neq \emptyset\}) \\
&= \mathsf{c}_\perp \mathsf{l}_\perp(\{I\}) = \mathsf{l}_\perp(\{I\}).
\end{aligned}
$$

    Now, since $\{I\} \subseteq \mathsf{u}_\perp(u)$ by assumption, we get $u \subseteq \mathsf{l}_\perp(\{I\}) = j$ by the residuation property of $(\mathsf{u}_\perp, \mathsf{l}_\perp)$. Furthermore, since $F \cap I = \emptyset$, we have $F \notin j$, so we have found an open element above u which does not contain F, as required. □

4 Discrete duality: Morphisms

In this section[8], we show how the object correspondence between RS polarities and perfect lattices from Section 2 can be extended to a duality. As indicated in the beginning of Section 3, finding appropriate duals for morphisms was already a problem in the category of bounded lattices, as can be seen in the development of duality by Urquhart [43] and Hartung [27]. There, it was only possible to find *functional* duals when the category of bounded lattices was restricted to contain only *surjective* morphisms.

In [9] and [16], the idea arose that allowing morphisms between polarities to be *relations* instead of requiring them to be functions largely widens the scope of possibilities. We pursue this approach and will consider categories, whose objects are RS polarities, in which the morphisms are relations, and even pairs of relations, satisfying certain properties.

Our approach can be distinguished from previous work in a few ways. First of all, an aspect which may be regarded as merely a cosmetic difference, but we believe is still worth mentioning, is that in [9] and [16], a morphism between RS polarities $\langle X, A, R\rangle$ and $\langle Y, B, S\rangle$ was always 'cross-wise': from X to B, and from Y to A. We have chosen to take our relations 'horizontally', i.e., from X to Y and from B to A.

Furthermore, we will show that our definition of morphisms between RS polarities can be seen to generalize the definition of Chu transform from the theory of Chu spaces, thus providing an independent reason why our definitions may be considered natural.

Finally, we try to be as impartial as possible about the question which are *the* morphisms in the category of perfect lattices, but rather treat several reasonable choices in a modular manner. We thus get a picture of increasingly more restricting choices of morphisms between perfect lattices, corresponding to increasingly more restricting conditions on the morphisms between RS polarities. For the reader's convenience, our results are informally summarized and indexed in the table below.

Perfect lattice	**RS polarity**	**Section**
Adjoint pair	Stable adjoint relation pair	4.1
$\bigvee/\bigwedge$-hemimorphism	1 relation satisfying a first-order condition (A)/(B)	4.4
Complete homomorphism	Pair of relations satisfying two first-order conditions: $\Diamond \leq \Box$ and $\Box \leq \Diamond$	4.5

Figure 1.2: Table of correspondence results

Furthermore, in Sections 4.2 and 4.3, we will define a composition on the morphisms between RS polarities, and indicate how, with this composition, the functors $(\cdot)^+$ and $(\cdot)_+$ constitute a duality.

[8] I would like to acknowledge explicitly the impact on this section of several discussions with Mai Gehrke and Adrian Pigors during TACL '09, and to mention the latter's ongoing research project [36], which is closely related to the matter treated in this section.

4.1 Adjoint pairs

As a first step, we will develop a duality between the category PLat_{ad} of perfect lattices with adjoint pairs between them and a category of RS polarities which we will denote by RSPol_{ad}.
We start by listing a number of possible notions of morphisms between perfect lattices.

Definition 4.1.1. Let $h : \mathbb{L}_1 \to \mathbb{L}_2$ be a function between two perfect lattices. We say

- h is a complete join-hemimorphism if h preserves arbitrary joins,
- h is a complete meet-hemimorphism if h preserves arbitrary meets,
- h is a complete homomorphism if h preserves arbitrary joins and meets.

Definition 4.1.2. If $f : \mathbb{L}_1 \to \mathbb{L}_2$ and $g : \mathbb{L}_2 \to \mathbb{L}_1$ are functions between perfect lattices, then (f, g) is an **adjoint pair**, written as $f \dashv g$, if

$$\forall u \in \mathbb{L}_1 \, \forall v \in \mathbb{L}_2 (f(u) \leq v \leftrightarrow u \leq g(v)).$$

We adopt the convention that the left adjoint determines the direction of an adjunction, i.e., in the above notation, (f, g) is an adjoint pair 'from' $\mathbb{L}_1$ 'to' $\mathbb{L}_2$. We thus get a category of complete lattices with adjoint pairs between them.

Note that right adjoints are unique: if $f \dashv g$ and $f \dashv g'$, then $g = g'$, and similarly for left adjoints. The following piece of lattice theory provides a useful description of these morphisms, and actually holds in general for complete lattices, but we will only be concerned with perfect lattices here.

Fact 4.1.3. *Let $h : \mathbb{L}_1 \to \mathbb{L}_2$ be a function between perfect lattices.*

(i) h is a complete join-hemimorphism if and only if there exists a g such that $h \dashv g$.

(ii) h is a complete meet-hemimorphism if and only if there exists a f such that $f \dashv h$.

It follows that the category of perfect lattices with adjoint pairs between them is equivalent to the category of perfect lattices with complete join-hemimorphisms, and dually equivalent to the category of perfect lattices with complete meet-hemimorphisms.
The most natural morphisms between perfect lattices are complete homomorphisms. In the end, we would like to find a duality between the category PLat of perfect lattices with complete homomorphisms between them and a category of RS polarities.
Let $f : \mathbb{L}_1 \leftrightarrows \mathbb{L}_2 : g$ be an adjoint pair from the perfect lattice $\mathbb{L}_1$ to the perfect lattice $\mathbb{L}_2$. Denote the RS polarity $(\mathbb{L}_1)_+$ by $\langle X, A, R \rangle$ and denote $(\mathbb{L}_2)_+$ by $\langle Y, B, S \rangle$.
Observe that f is completely join preserving, so that it is completely determined by its action on the completely join irreducibles X of $\mathbb{L}_1$: for any $u \in \mathbb{L}_1$, we have

$$f(u) = f\left(\bigvee\{x \in X : x \leq u\}\right) = \bigvee\{f(x) : x \leq u\}.$$

Also, since $\mathbb{L}_2$ is join-generated by Y, we have, for any $x \in X$, that

$$f(x) = \bigvee\{y \in Y : y \leq f(x)\}.$$

These two basic arguments show that f is fully determined by a relation $F \subseteq X \times Y$, defined by

$$xFy \iff y \leq f(x). \tag{1.2}$$

We then get that

$$f(x) = \bigvee F[x], \tag{1.3}$$
$$f(u) = \bigvee \{f(x) : x \leq u\}. \tag{1.4}$$

Note that equation (1.3) implies that the relation F has the following special property:

$$\forall x \forall y (y \leq \bigvee F[x] \to xFy). \tag{1.5}$$

We can describe the property $y \leq \bigvee F[x]$ without referring to the join in the perfect lattice by saying that y is in the c_S-closure of the set $F[x]$, using the fact that $\mathbb{L}_2$ is meet-generated by the set B:

$$\begin{aligned} y \leq \bigvee F[x] &\iff \forall b (b \geq \bigvee F[x] \to b \geq y) \\ &\iff \forall b (\forall y' (xFy' \to y'Sb) \to ySb) \\ &\iff y \in \mathsf{l}_S \mathsf{u}_S (F[x]). \end{aligned}$$

So, we can concisely describe the property (1.5) as

$$\forall x (\mathsf{c}_S(F[x]) \subseteq F[x]), \tag{S_F}$$

or, in words, for all $x \in X$, $F[x]$ is stable under the Galois connection induced by S. We call a relation F satisfying (S_F) **stable**.

Lemma 4.1.4. *If $F \subseteq X \times Y$ is a stable relation, then $y \leq \bigvee F[x]$ implies $y \in F[x]$.*

By a completely analogous, dual, argument for the completely meet preserving function g, if we define a relation $G \subseteq B \times A$ by

$$bGa \iff a \geq g(b), \tag{1.6}$$

then we get that

$$g(b) = \bigwedge G[b], \tag{1.7}$$
$$g(v) = \bigwedge \{g(b) : b \geq v\}, \tag{1.8}$$

and that

$$\forall b (\mathsf{c}_R(G[b]) \subseteq G[b]), \tag{S_G}$$

i.e., G is a **stable** relation.

Of course, since the functions f and g are 'connected' – they form an adjoint pair – we expect that the relations F and G are connected as well. The following proposition makes the nature of this connection precise.

Proposition 4.1.5. *If $f \dashv g$ and F and G are defined as in (1.2) and (1.6), then*

$$\forall x \forall b (R[x] \supseteq G[b] \leftrightarrow F[x] \subseteq S^{-1}[b]). \tag{ADJ}$$

Proof. Take any $x \in X$, $b \in B$. Using (1.3), (1.7) and the fact that (f, g) is an adjoint pair, we have

$$\begin{aligned} F[x] \subseteq S^{-1}[b] &\iff \bigvee F[x] \leq b \\ &\iff f(x) \leq b \\ &\iff x \leq g(b) \\ &\iff x \leq \bigwedge G[b] \\ &\iff R[x] \supseteq G[b]. \end{aligned}$$

□

Definition 4.1.6. We call a pair (F, G) of relations **stable** if it satisfies (S_F) and (S_G) and **adjoint** if it satisfies (ADJ).
Let RSPol_{ad} be the category of RS polarities, in which the morphisms from an RS polarity $\langle X, A, R\rangle$ to an RS polarity $\langle Y, B, S\rangle$ are the **stable adjoint pairs** of relations (F, G), where $F \subseteq X \times Y$ and $G \subseteq B \times A$.

Before defining the composition of two stable adjoint pairs of relations, we summarize the above arguments as follows.

Proposition 4.1.7. *If $f \dashv g$ is an adjoint pair from $\mathbb{L}_1$ to $\mathbb{L}_2$, then $(f, g)_+ := (F, G)$ is a stable adjoint pair from $(\mathbb{L}_1)_+$ to $(\mathbb{L}_2)_+$.*

Remark 4.1.8. The particular definitions of the relation pair (F, G), that we gave in this section, were inspired by the theory of *Chu spaces* (see, for example, [37], [3]).
To go into some detail: a *Chu space* over a *colour set* C is defined to be a triple $\langle X, A, t\rangle$, where $t : X \times A \to C$ is a function. It is clear that a Chu space over $\{0, 1\}$ is precisely what we have called a polarity above. The general definition of a *Chu morphism* applied to the case of $\{0, 1\}$-Chu spaces boils down to an 'adjoint function pair', i.e., a morphism from the $\{0, 1\}$-Chu space $\langle X, A, R\rangle$ to the $\{0, 1\}$-Chu space $\langle Y, B, S\rangle$ is a pair of functions (k, l), where $k : X \to Y$, $l : B \to A$, such that $k(x)Sb$ iff $xRl(b)$, for all $x \in X$, $b \in B$. In this light, an 'adjoint relation pair' as defined above may also be called a *relational Chu space morphism.*

4.2 Composition

A disadvantage of the relational morphisms considered in Gehrke [16] was that they did not compose in an obvious way. As a first step towards a natural composition, we defined the relations in a *prima facie* 'composable' manner, drawing inspiration from the definition of functions for Chu spaces. However, if we try to compose two stable adjoint relation pairs using the usual relational composition, we are not guaranteed to end up with a pair that is still stable.
There are now two possible ways to proceed: first of all, we know what the composition of stable adjoint pairs *should* be, because we want the assignment $(\cdot)_+$ on morphisms to be a functor between the categories PLat_{ad} and RSPol_{ad}. We will show the definition resulting from this constraint below.

Secondly, it will turn out that a more illuminating way to view this composition is as the *closure* of usual composition of relations. It then becomes much easier to show that this composition has all of the desired properties. We will state the result below.

The constraint that $(\cdot)_+$ should be functorial dictates our definition of the composition $*$ of two stable adjoint pairs (F_1, G_1) and (F_2, G_2), since we want that, if $f_1 \dashv g_1$ and $f_2 \dashv g_2$ are composable adjoint pairs, then

$$(f_2 \circ f_1, g_1 \circ g_2)_+ = (f_1, g_1)_+ * (f_2, g_2)_+.$$

A calculation shows that this requirement leads to the following definition.

Definition 4.2.1. Let $(F_1, G_1) : \langle X, A, R\rangle \to \langle Y, B, S\rangle$ and $(F_2, G_2) : \langle Y, B, S\rangle \to \langle Z, C, T\rangle$ be two stable adjoint pairs of relations. Define their **composition** $(F_1, G_1) * (F_2, G_2)$ to be the pair of relations $(F_1 * F_2, G_2 * G_1)$, where

$$(x, z) \in F_1 * F_2 \iff \forall c(\forall y(xF_1y \to F_2[y] \subseteq T^{-1}[c]) \to zTc),$$

$$(c, a) \in G_2 * G_1 \iff \forall x(\forall b(cG_2b \to G_1[b] \subseteq R[x]) \to xRa).$$

Note[9] that a more concise way to state these definitions is:

$$F_1 * F_2[x] = \mathsf{c}_T(F_1 \circ F_2[x]),$$

$$G_2 * G_1[c] = \mathsf{c}_R(G_2 \circ G_1[c]),$$

where $\circ$ is the usual relational composition.

Proposition 4.2.2. *With the composition $*$, the collection of RS polarities with stable relation pairs is a category, and the assignments $(\cdot)^+$ and $(\cdot)_+$ are functorial.*

4.3 Equivalence of PLat_{ad} and RSPol_{ad}

We can also go back, from stable adjoint pairs of relations to adjoint pairs of perfect lattice morphisms.

Let $(F, G) : \langle X, A, R\rangle \to \langle Y, B, S\rangle$ be an adjoint pair. Define $f : \langle X, A, R\rangle^+ \leftrightarrows \langle Y, B, S\rangle^+ : g$ by

$$f(x) := \bigvee F[x], \qquad g(b) := \bigwedge G[b], \tag{1.9}$$

$$f(u) := \bigvee\{f(x) : x \leq u\}, \qquad g(v) := \bigwedge\{g(b) : b \geq v\}. \tag{1.10}$$

Clearly, we need to show that f and g defined this way form an adjoint pair. We split this task up into a few reasonably simple lemmas.

Lemma 4.3.1. *If (F, G) is an adjoint pair, then f and g defined above satisfy*

$$\forall x \forall b(f(x) \leq b \leftrightarrow x \leq g(b)). \tag{adj}$$

Proof. Similar to the proof of Proposition 4.1.5. □

[9] Adrian Pigors (personal communication), 05-07-2009

Lemma 4.3.2. *If (F, G) is an adjoint pair, then f is a complete join-hemimorphism and g is a complete meet-hemimorphism.*

Proof. We only show that f is a complete join-hemimorphism, the proof that g is a complete meet-hemimorphism is dual.
Take an arbitrary $U \subseteq L_1$ and write $u_0 := \bigvee U$. We need to show that $f(u_0) = \bigvee f[U]$.
If $u \in U$, then $u \leq u_0$, so $f(u) \leq f(u_0)$ since f is order preserving. So we get that $f(u_0) \geq \bigvee f[U]$.
For the inequality $f(u_0) \leq \bigvee f[U]$, take an arbitrary $b \in B$ such that $b \geq \bigvee f[U]$. It suffices to show that $b \geq f(u_0)$, since B meet-generates $\mathbb{L}_2$. Since $b \geq \bigvee f[U]$, we have $b \geq f(u)$ for all $u \in U$. By definition, $f(u) = \bigvee\{f(x) : x \leq u\}$, so we get that $b \geq f(x)$ for all $x \leq u$, for all $u \in U$. By (adj), we get $g(b) \geq x$ for all $x \leq u$, for all $u \in U$. But $u = \bigvee\{x : x \leq u\}$, since X join-generates $\mathbb{L}_1$, so we get that $g(b) \geq u$, for all $u \in U$. In other words, $g(b)$ is an upper bound for U, so $g(b) \geq u_0$. Applying f to both sides, we get $f(g(b)) \geq f(u_0)$. Furthermore, by (adj) applied to $x := g(b)$, we get that $b \geq f(g(b))$. So we conclude that $b \geq f(u_0)$, as required. □

Lemma 4.3.3. *If $f : \mathbb{L}_1 \leftrightarrows \mathbb{L}_2 : g$ are functions between perfect lattices such that (adj) holds, f is a complete join-hemimorphism and g is a complete meet-hemimorphism, then $f \dashv g$.*

Proof. Take $u \in \mathbb{L}_1$ and $v \in \mathbb{L}_2$ arbitrary.
Suppose $f(u) \leq v$. To show that $u \leq g(v)$, note that $u = \bigvee\{x : x \leq u\}$ and $g(v) = g(\bigwedge\{b : b \geq v\}) = \bigwedge\{g(b) : b \geq v\}$ (because g is a complete meet-hemimorphism).
Let $x \in X$ and $b \in B$ be arbitrary such that $x \leq u$ and $v \leq b$. Since $x \leq u$ we have $f(x) \leq f(u)$, so we get that $f(x) \leq f(u) \leq v \leq b$. Now by (adj), we have $x \leq g(b)$, as required.
The proof of the other implication is dual. □

We conclude:

Proposition 4.3.4. *If $(F, G) : \langle X, A, R\rangle \to \langle Y, B, S\rangle$ is an adjoint pair, then (f, g) defined by (1.9) and (1.10) is an adjoint pair $\langle X, A, R\rangle^+ \to \langle Y, B, S\rangle^+$.*

In particular, $(\cdot)^+$ maps adjoint pairs of relations which are stable to adjoint pairs of functions. This stability requirement will guarantee that $(\cdot)_+$ defined in Proposition 4.1.7 is a left inverse to $(\cdot)^+$:

Proposition 4.3.5. *If (F, G) are stable relations then $((F, G)^+)_+ = (F, G)$.*

Proof. Let f and g be the functions defined in (1.9) and (1.10).
Let F' be the relation defined by $xF'y$ iff $y \leq f(x)$. We show that $F' = F$.
Clearly, if xFy, then $y \in F[x]$ so $y \leq \bigvee F[x] = f(x)$, so $xF'y$.
Conversely, if $y \leq f(x) = \bigvee F[x]$, then $y \in F[x]$ by Lemma 4.1.4.
Similarly, the stability of G guarantees that the relation defined by $g(b) \leq a$ is equal to G. □

The following result now follows by combining the above Lemmas and Propositions.

Theorem 4.3.6. *$(\cdot)^+$ and $(\cdot)_+$ form an equivalence between the categories PLat_{ad} and RSPol_{ad}.*

4.4 Hemimorphisms

Although the route via adjoints described in the previous section works, it would be useful to have necessary and sufficient conditions on just one relation $F \subseteq X \times Y$ to be the dual of a complete join-hemimorphism, without reference to its adjoint relation. The aim of this section is to find such first-order conditions on the relation F.

From lattice theory, we know that if a function f has a right adjoint, then its value at a point v is the greatest element which f sends to a point below v. Since we are working with perfect lattices, it suffices to look at the join- and meet-irreducibles.

Proposition 4.4.1. *Let $f : \mathbb{L}_1 \leftrightarrows \mathbb{L}_2 : g$ be functions between two perfect lattices and $\langle X, A, R\rangle$ and $\langle Y, B, S\rangle$ their associated RS polarities, respectively.*
The following are equivalent:

(i) $f \dashv g$,

(ii) f is a complete join-hemimorphism, g is a complete meet-hemimorphism, and for all $b \in B$,

$$g(b) = \bigvee\{x \in X : f(x) \leq b\},$$

(iii) f is a complete join-hemimorphism, g is a complete meet-hemimorphism, and for all $x \in X$,

$$f(x) = \bigwedge\{b \in B : x \leq g(b)\}.$$

From part (ii) of this proposition, we see that if $f \dashv g$, then the relation G is determined by the relation F:

$$\begin{aligned} bGa &\iff a \geq g(b) \\ &\iff a \geq \bigvee\{x \in X : f(x) \leq b\} \\ &\iff \forall x(f(x) \leq b \to x \leq a) \\ &\iff \forall x(F[x] \subseteq S^{-1}[b] \to xRa). \end{aligned}$$

Similarly, part (iii) of the proposition yields that F can be determined from the relation G:

$$xFy \iff \forall b(G[b] \subseteq R[x] \to ySb).$$

Conversely, if a pair of relations (F, G) is 'intertwined' in this way, then their duals f and g are adjoints.

Lemma 4.4.2. *Let $F \subseteq X \times Y$ and $G \subseteq B \times A$ be relations.*
The following are equivalent:

(i) (F, G) is a stable adjoint pair.

(ii) F and G satisfy the conditions $(F \dashv G)$ and $(G \vdash F)$:

$$xFy \iff \forall b(G[b] \subseteq R[x] \to ySb). \qquad (G \vdash F)$$

$$bGa \iff \forall x(F[x] \subseteq S^{-1}[b] \to xRa). \qquad (F \dashv G)$$

Substituting one of the two equivalences in the other, we get separate first-order conditions for both F and G to be part of a stable adjoint pair:

Proposition 4.4.3.

(i) Let $F \subseteq X \times Y$ be a relation.

The following are equivalent:

(a) There exists $G \subseteq B \times A$ such that (F, G) is a stable adjoint pair.

(b) F satisfies

$$\forall x \forall y \left(xFy \leftrightarrow \forall b(\forall a(\forall x'(F[x'] \subseteq S^{-1}[b] \to x'Ra) \to xRa) \to ySb)\right) \qquad (\textstyle\bigvee_F)$$

(ii) Let $G \subseteq B \times A$ be a relation.

The following are equivalent:

(a) There exists $F \subseteq X \times Y$ such that (F, G) is a stable adjoint pair.

(b) G satisfies

$$\forall b \forall a \left(bGa \leftrightarrow \forall x(\forall y(\forall b'(G[b'] \subseteq R[x] \to ySb') \to ySb) \to xRa)\right). \qquad (\textstyle\bigwedge_G)$$

Corollary 4.4.4. *(i) There is a one-to-one correspondence between complete join-hemimorphisms $\mathbb{L}_1 \to \mathbb{L}_2$ and relations $F \subseteq X \times Y$ which satisfy $\bigvee_F$.*

(ii) There is a one-to-one correspondence between complete meet-hemimorphisms $\mathbb{L}_1 \to \mathbb{L}_2$ and relations $G \subseteq B \times A$ which satisfy $\bigwedge_G$.

4.5 Complete homomorphisms

We now regard a complete homomorphism $h : \mathbb{L}_2 \to \mathbb{L}_1$ as having two 'faces': its join-preserving face yields a right adjoint g, whereas its meet-preserving face yields a left adjoint f. Since (f, h) and (h, g) are adjoint pairs, they correspond to stable adjoint relation pairs $(f, h)_+ = (F, H_m)$ and $(h, g)_+ = (H_j, G)$. The names H_m and H_j reflect the idea that H_m is the relation representing the meet-preserving face of h and H_j is the relation representing the join-preserving face of h.
Our aim is now to describe conditions on the relations H_m and H_j which correspond to the fact that H_m and H_j stem from the same homomorphism h. If we succeed in this, we can take the relation pair (F, G) to be a full dual description of the complete homomorphism h.
Let $h_m : \mathbb{L}_2 \to \mathbb{L}_1$ and $h_j : \mathbb{L}_2 \to \mathbb{L}_1$ be the functions defined from H_m and H_j, respectively. So

$$h_j(y) := \bigvee H_j[y], \qquad h_m(b) := \bigwedge H_m[b],$$
$$h_j(v) := \bigvee \{h_j(y) : y \leq v\}, \qquad h_m(v) := \bigwedge \{h_m(b) : b \geq v\}.$$

We want to find conditions on H_m and H_j such that $h_m = h_j$. We regard this equality as two inequalities:

- $\forall v(h_j(v) \leq h_m(v))$,
- $\forall v(h_m(v) \leq h_j(v))$.

Our aim is now to find first-order conditions on H_m and H_j which guarantee the inequalities to hold. As is already noted in [9] (above Proposition 4.21), these inequalities are reminiscent of Sahlqvist correspondence theory: regarding h_j as a $\Diamond$ and h_m as a $\Box$, our aim is to find first-order correspondents of the inequalities $\Diamond \leq \Box$ and $\Box \leq \Diamond$. In the terminology of [5] and [35], the first of these inequalities is the analogue of a *very simple* Sahlqvist implication, and the second of a *simple* Sahlqvist implication.

Proposition 4.5.1. *The following are equivalent:*

(i) $\forall v(h_j(v) \leq h_m(v))$,

(ii) $\forall v \forall y \forall b(y \leq v \leq b \rightarrow h_j(y) \leq h_m(b))$,

(iii) $\forall y \forall b(y \leq b \rightarrow h_j(y) \leq h_m(b))$,

(iv) $\forall y \forall b(ySb \rightarrow \forall x \forall a(yH_jx \wedge bH_ma \rightarrow xRa))$.

(v) $H_j^{-1} \circ S \circ H_m \subseteq R$.

Proof. (i) and (ii) are equivalent by the definitions of $h_j(v)$ and $h_m(v)$.
(ii) and (iii) are equivalent by transitivity of $\leq$.
(iii) and (iv) are equivalent because the inequality $h_j(y) \leq h_m(b)$ can be rewritten as

$$\forall x \forall a(yH_jx \wedge bH_ma \rightarrow xRa),$$

using the definitions of $h_j(y)$ and $h_m(b)$, as we have already done a couple of times above.
(iv) and (v) are equivalent by shuffling quantifiers and the definition of relational composition. $\Box$

For the other inequality, we will need the fact that H_m and H_j are part of adjoint pairs (F, H_m) and (H_j, G), so that the properties $(H_m \vdash F)$ and $(H_j \dashv G)$ from Lemma 4.4.2 hold. That we need to make use of the residuation properties at this point may not come as a surprise. Indeed, as already noted above, the inequality $h_m \leq h_j$ can be seen as an example of (the poset version of) a so-called *simple* Sahlqvist implication, and in classical modal logic, one exactly needs to use the residuation properties for finding first-order correspondents of such *simple* Sahlqvist implications ([5], [35]).

Proposition 4.5.2. *The following are equivalent:*

(i) $\forall v(h_m(v) \leq h_j(v))$,

(ii) $\forall v \forall x \forall a(x \leq h_m(v) \wedge h_j(v) \leq a \rightarrow x \leq a)$,

(iii) $\forall v \forall x \forall a(f(x) \leq v \leq g(a) \rightarrow x \leq a)$,

(iv) $\forall x \forall a(f(x) \leq g(a) \rightarrow x \leq a)$,

(v) $\forall x \forall a(\forall y \forall b(xFy \wedge aGb \rightarrow ySb) \rightarrow xRa)$,

(vi) $\forall x \forall a \big(\forall y \forall b \big(\forall b'(H_m[b'] \subseteq R[x] \to ySb') \wedge \forall y'(H_j[y'] \subseteq R^{-1}[a] \to y'Sb) \to ySb\big) \to xRa\big)$.

Proof. (i) and (ii) are equivalent because $h_m(v) = \bigvee\{x : x \leq h_m(v)\}$ and $h_j(v) = \bigwedge\{a : a \geq h_j(v)\}$.
(ii) and (iii) are equivalent because (f, h_m) and (h_j, g) are adjoint pairs.
(iii) and (iv) are equivalent by transitivity of $\leq$.
(iv) and (v) are equivalent by the facts that $f(x) = \bigvee F[x]$ and $g(a) = \bigwedge G[a]$.
(v) and (vi) are equivalent because (F, H_m) and (H_j, G) are adjoint pairs of relations, so we may write $\forall b'(H_m[b'] \subseteq R[x] \to ySb')$ instead of xFy by the property $(H_m \vdash F)$, and similarly we may rewrite aGb using the property $(G \dashv H_j)$. □

Proposition 4.5.3. *Let $H_j \subseteq Y \times X$ and $H_m \subseteq B \times A$ be relations, and h_j and h_m their dual functions as defined above.*
Then h_j and h_m are equal to the same complete homomorphism h if and only if the following four conditions hold:

(i) $\bigvee_{H_j}$

(ii) $\bigwedge_{H_m}$

(iii) $H_j^{-1} \circ S \circ H_m \subseteq R$.

(iv) $\forall x \forall a(\forall y \forall b((\forall b'(H_m[b'] \subseteq R[x] \to ySb') \wedge \forall y'(H_j[y'] \subseteq R^{-1}[a]) \to y'Sb) \to ySb) \to xRa)$.

Proof. If (i)-(iv) hold, then by (i) and (ii) and Proposition 4.4.3, there are relations $F \subseteq X \times Y$ and $G \subseteq A \times B$ such that (F, H_m) and (H_j, G) are stable adjoint pairs. Then h_j is a complete join-hemimorphism and h_m is a complete meet-hemimorphism. By (iii) and Proposition 4.5.1, $h_j \leq h_m$. By (iv) and Proposition 4.5.2, $h_m \leq h_j$. So $h_j = h_m =: h$ is a complete homomorphism. Similar reasoning applies for the other direction, since all propositions we referred to were equivalences. □

Corollary 4.5.4. *There is a one-to-one correspondence between complete homomorphisms from $\mathbb{L}_1$ to $\mathbb{L}_2$ and pairs of relations (H_j, H_m) satisfying the conditions (i)-(iv) from Proposition 4.5.3 above.*

To briefly summarize this chapter: after reviewing the algebraic construction of the canonical extension in Sections 1.1–1.3, we made some remarks about the categorical properties of the canonical extension in Section 1.4. In Section 2, we reviewed the object part of the duality between perfect lattices and two-sorted, frame-like structures called RS polarities. We then gave, in Section 3, an alternative way to obtain the canonical extension via a generalisation of the maximal filter-ideal pair construction from lattices to posets. Finally, in Section 4, we gave a modular account of the several possibilities to define morphisms between RS polarities, and found first-order correspondents for natural properties of maps between perfect lattices.
In the next chapter, we will discuss how the canonical extension can be applied to distributive lattices, in order to show the canonicity of a Sahlqvist fragment for distributive modal logic by algebraic methods.

Chapter 2

Canonicity of inequalities

In the Introduction, after the proof of Theorem 2.2, we remarked that *canonicity* of a formula, or inequality, is very useful for, among other things, proving strong frame completeness of the logic axiomatized by that formula or inequality. An entire research field, which can be broadly described as 'the algebraic study of canonicity', originated from this. The question underlying this research field can be stated algebraically as follows:

Main question. Under what conditions on terms α and β do we have, for any algebra $\mathbb{A}$, that

$$\mathbb{A} \models \alpha \leq \beta \implies \mathbb{A}^\sigma \models \alpha \leq \beta \qquad ?$$

Note that the question, formulated as such, is (deliberately) vague: we are unclear about the formal language in which the terms α and β are formulated, what we mean by 'conditions' on α and β, and about what are the algebras $\mathbb{A}$ interpreting the inequality $\alpha \leq \beta$. We will be more precise about the specific question that we are interested in in this chapter below.

A second important thing to note is that we phrased the *Main question* only referring to algebras, without mentioning their dual frames. Originally, however, canonicity was developed mainly using *correspondence* ideas, similar to those outlined in the previous chapter of this thesis, with the exception of [31], which in fact started algebraic canonical extension theory as early as 1951, but remained relatively unnoticed for a long time. Only relatively recently, in Jónsson's paper [30], the first proof that the Sahlqvist fragment of Boolean modal logic is canonical was given that used *purely algebraic methods*. The paper [30] instigated a revival of the algebraic study of canonicity, taken up in many papers, of which we name a few here, by no means, however, pretending to give an extensive list: [7], [9], [14], [20], [22], [42]. More detailed historical surveys of the research field of canonicity can be found in the notes to Chapter 5 of [5] and in Section 7 of [44].

In this chapter, we focus on the algebraic canonicity proof for the 'Sahlqvist' fragment of distributive modal logic that was given in [22]. In particular, we revisit that proof, filling in some details, and we pay a considerable amount of attention to the binary operation called "**n**" (for **n**ot less than or equal), which was used in the canonicity proof of [22] in order to broaden the 'Sahlqvist' fragment considerably. We isolate this "**n**-trick"[9], and indicate how we believe it may be related to some more general phenomena, in particular Ackermann's Lemma, which plays an important

[9]For the sake of brevity and in lack of a better name, we will refer to the syntactic step which uses the operation **n** as "the **n**-trick", because of its apparently magical and at the same time mysterious properties, which will be a main topic of discussion in this chapter.

role in the canonicity theory of the *inductive* fragments of Boolean and distributive modal logic, surveyed and developed in [45], [46], and [7].
The organisation of this chapter is as follows. In Section 1, we review the definitions of bounded lattice logic, generation trees, and the Sahlqvist fragment, and in the process we introduce some useful, unifying notation. We then study one particular part of the canonicity proof of the Sahlqvist fragment for distributive modal logic in considerable detail in Section 2, and we sketch how the proof is completed in Sections 3 and 4. In Section 5, we will discuss the possibility of generalising this proof to the inductive fragment.

1 Terms and trees

1.1 Terms

Although [22] talks only about *distributive* modal logic, part of the theory works for more general logics whose algebraic semantics are given by bounded, but not necessarily distributive, lattices, semilattices, or even posets. We define our language as follows.

Definition 1.1.1. An **order type** is a vector $\epsilon \in \{1, \partial\}^n$, for some n. We denote by $\#\epsilon$ the **length** of the vector ϵ and by ϵ^∂ the **dual order type** of ϵ, which contains the precise opposite signs from ϵ.
An **order similarity type** is a list $\tau = (\epsilon^{(1)}, \dots, \epsilon^{(k)})$ of order types.
Given an order similarity type τ, the language $\mathcal{L}_{\text{BLM}_\tau}$ (for *bounded lattice with monotone operations*) consists of the formulas (also called BLM_τ-terms) defined inductively by

$$\phi ::= x \mid \top \mid \bot \mid \phi \vee \phi \mid \phi \wedge \phi \mid \Diamond_\epsilon(\phi_1, \dots, \phi_{\#\epsilon}) \quad (\epsilon \in \tau),$$

where x is a proposition letter from a fixed set X of propositional variables.
In what follows, we will also refer to "the order type of $\vee$" and "the order type of $\wedge$", both meaning the order type $(1, 1)$.

The algebraic semantics for $\mathcal{L}_{\text{BLM}_\tau}$ are readily given by the following definition.

Definition 1.1.2. A **bounded lattice with monotone operations**, or BLM_τ, is a tuple $\mathbb{A} = (\mathbb{L}, f_\epsilon)_{\epsilon \in \tau}$, where $\mathbb{L}$ is a bounded lattice and, for each $\epsilon \in \tau$, $f_\epsilon : \mathbb{A}^\epsilon \to \mathbb{A}$ is a monotone function (i.e., $f_\epsilon : L^{\#\epsilon} \to L$ is monotone in those coordinates in which the corresponding coordinate of ϵ is 1, and antitone in the coordinates where the corresponding coordinate of ϵ is ∂).
Let ϕ be an arbitrary formula of $\mathcal{L}_{\text{BLM}_\tau}$ with variables $x_1, \dots, x_n$ occurring in it. The **term function** $\phi^\mathbb{A}$ is defined inductively by interpreting the symbols $\top$, $\bot$, $\wedge$ and $\vee$ as the operations of the underlying lattice $\mathbb{L}$ of $\mathbb{A}$, and each operation symbol $\Diamond_\epsilon$ as f_ϵ.
Given two BLM_τ-terms α, β, let $x_1, \dots, x_n$ be all variables occurring in α and β. We say that **the inequality** $\alpha \leq \beta$ **holds in** $\mathbb{A}$ if the inequality $\alpha^\mathbb{A} \leq \beta^\mathbb{A}$ holds pointwise for the term functions. If this is the case, we write $\mathbb{A} \models \alpha \leq \beta$. We say two inequalities are **equivalent** if they hold in precisely the same algebras.

Clearly, the definitions in [22] are a special case of this definition. There, the **distributive similarity type** $\tau_D = ((1), (1), (\partial), (\partial))$ is used, the operation symbols are denoted by $\Diamond, \Box, \lhd, \rhd$, respectively, and some more restrictions on the interpreting algebras are imposed: $\mathbb{L}$ must be a *distributive* lattice, each f_ϵ preserves finite joins or meets, or turns finite meets into joins, or vice versa. Such a structure is called a DMA, for Distributive Modal Algebra. We refer the reader to Definition 2.9 of [22] for details.

1.2 Trees

To define the conditions on the inequalities, we need the notion of a **signed generation tree**. Because formally defining such trees seems inevitably cumbersome and not very enlightening, we first give a hopefully clarifying example.

Example 1.2.1. In BLM_{τ_D}, let α be the formula $\lhd x \vee \Box(x \wedge y)$. Then the *positive generation tree* for α looks as follows:

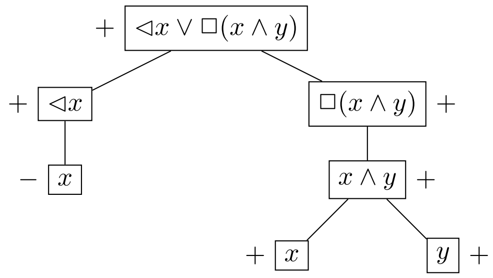

The adjective *positive* refers to the fact that in the construction of the tree, the top node (the *root* of the tree) was labeled '+'. If we put a $-$ at the top node, we get the *negative* version of the tree. The rule for labelling subsequent nodes with signs is simple: in the coordinates where the order type of the operation at hand is 1, we label the node with the same sign as the node directly above it, and where the order type is ∂, we put the opposite sign.
One more example: take $\tau = ((\partial, 1), (1, \partial))$, and denote the operations by $\circ$ and $\bullet$, respectively. Then the negative generation tree of $\beta = x \circ (x \bullet y)$ looks as follows:

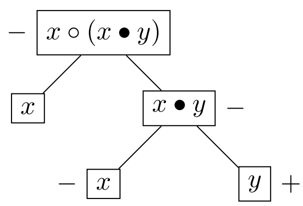

$\Box$

Now for the formal definition.

Definition 1.2.2. Let α be a BLM_τ-term. We inductively define the **generation tree** $T(\alpha)$ of α, as follows:

- If $\alpha = x$ for some variable x, $T(\alpha)$ is the tree consisting of one node called 'x' and no edges.
- If $\alpha = \heartsuit(\alpha_1, \ldots, \alpha_n)$ for some n-ary operation symbol $\heartsuit$ (which may also be one of the binary operations $\vee$ and $\wedge$, or the nullary operations $\bot$ and $\top$), then $T(\alpha)$ is the tree consisting of a node called 'α' and n disjoint subtrees $T(\alpha_1), \ldots, T(\alpha_n)$, with edges from the node α to

the root of each $T(\alpha_i)$. By **the operation at the node** α **in** $T(\alpha)$ we mean $\heartsuit$, and the root of $T(\alpha_i)$ is called the i^{th} **direct successor of the node** α.

We now define the **positive** and **negative** generation tree, $T^+(\alpha)$ and $T^-(\alpha)$, respectively, as follows. Let $T^+(\alpha)$ be the tree $T(\alpha)$ with a label $+$ added to the root, and $T^-(\alpha)$ the tree with a label $-$ added to the root. Then, walking down all the paths in the tree from the root to the leaves, label all nodes in $T^+(\alpha)$ and $T^-(\alpha)$ with either a $+$ or a $-$ according to the following rule:

To the i^{th} direct successor of n, assign the *same* or *opposite* label as the node n, according to whether the i^{th} coordinate of the order type of the operation at the node n is 1 or ∂, respectively.

After this process, all nodes in the trees $T^+(\alpha)$ and $T^-(\alpha)$ will be labeled with a sign.

We now introduce some useful notation from [7].
A **subterm** s of a term α is the content of any node in the generation tree of α. By a subterm we will always mean *a specific occurrence* of that subterm. (That is, in Example 1.2.1 above, there are two *different* subterms of β which have 'x' as their content. It is useful to regard them as different, because these two nodes get different signs in the signed generation tree.)
Note that a subterm s generates a **subtree** of α, namely the one which has s as its root. We will identify subterms with their subtrees, and write $s \prec \alpha$ for "s is a subterm of α".
A useful expansion of this notation is to write, for example, "$+s \prec -\alpha$" if s has the sign $+$ at its root in the negative generation tree of α. We will also write "$\epsilon_i s \prec +\alpha$" to mean that s has the sign that is dictated by ϵ_i in the positive tree of α.
Given an order type $\epsilon = (\epsilon_1, \ldots, \epsilon_n)$ for all of the variables $x_1, \ldots, x_n$ occurring (possibly more than once) in α, we say that an occurrence $x_i \prec \alpha$ **agrees with** ϵ if $\epsilon_i x_i \prec +\alpha$. We say that the term α **agrees with** ϵ if every occurrence of a variable agrees with ϵ, and we abbreviate this as '$\epsilon(\alpha)$'. A term α is called **uniform** if there exists some ϵ such that $\epsilon(\alpha)$.
If $s \prec \alpha$ and ϕ is any BLM_τ-formula, we write $\alpha[\phi/s]$ for the result of substituting the subterm s by ϕ. (Note, again, that this substitution occurs only once, at the location of the particular subterm s in the tree of α.)

Example 1.2.3 (Example 1.2.1, continued)**.** In the example formula $\beta = x \circ (x \bullet y)$ above, we have $+x \prec -\beta$, $-x \prec -\beta$ and $+y \prec -\beta$, but *not* $-y \prec -\beta$.
For an example of the substitution notation, we have $\beta[z/(x \bullet y)] = x \circ z$. We cannot unambiguously write $\beta[z/x]$, because substitution can happen only once, so we would need to specify explicitly *which* occurence of x we want to replace with z.
Neither α nor β is uniform, but the formula $\alpha' := \alpha[z/\lhd x]$ *is* uniform, because α' agrees with the order type $\epsilon = (1, 1, 1)$, that is, $\epsilon(\alpha')$ holds.

The following lemma is an obvious consequence of these definitions.

Lemma 1.2.4. *If a term α agrees with an order type ϵ, then the term function $\alpha^{\mathbb{A}} : \mathbb{A}^\epsilon \to \mathbb{A}$ is monotone.*

Proof. By induction on the complexity of α. The base step is the observation that all variables have monotone term functions, and the inductive step holds because each operation symbol $\diamondsuit_\epsilon$ is interpreted as a monotone function $f_\epsilon : \mathbb{A}^\epsilon \to \mathbb{A}$, and monotonicity is a property that is preserved by functional composition. $\square$

To finish this section, we give a small glossary which translates the the terminology of the paper [22] to the terms we have defined above.

Term in [22]	*Corresponding term here*
positive	agrees with $(1, \ldots, 1)$
negative	agrees with $(\partial, \ldots, \partial)$
ϵ-positive	agrees with ϵ
ϵ-negative	agrees with ϵ^∂
uniform	agrees with some ϵ

It is apparent from this table that the "agreeing with"-relation unifies and simplifies some of the 'old' terminology.

1.3 The Sahlqvist fragment

In this section, we define a syntactic fragment of the language BLM_{τ_D}, in which all formulas will turn out to be canonical: the Sahlqvist fragment [22]. The definitions in this section have been generalised in two ways: firstly, the Sahlqvist fragment has been generalised to the context of partially ordered sets by Suzuki [42], and secondly, in order to include more formulas, Conradie and Palmigiano define a version of the *inductive fragment* for distributive modal logic [7]. We will say more about the inductive fragment in Section 5.
We need the following preliminary definitions regarding the generation trees.

Definition 1.3.1. Given a term $\alpha(x_1, \ldots, x_n)$ and an order type $\epsilon = (\epsilon_1, \ldots, \epsilon_n)$, an occurrence of the variable x_i in a signed tree $T^\delta(\alpha)$ (where $\delta \in \{+, -\}$) is **ϵ-critical** if $\epsilon_i x_i \prec \delta\alpha$. An **$\epsilon$-critical path** is a path[10] in the tree $T^\delta(\alpha)$ from an ϵ-critical occurrence of a variable to the root of the tree.

So, we now use the order type ϵ to pinpoint certain occurrences of variables in a term as 'critical'. The definition of our syntactic fragment is of the general form "no forbidden configurations in critical paths". To say what configurations are 'forbidden', we need to categorize occurrences of connectives, in the following way.

Definition 1.3.2. Let α be a term and $T^\delta(\alpha)$ a signed tree of α (where $\delta \in \{+, -\}$).

- A node with the sign $+$ and the operation $\vee$, $\Diamond$, or $\lhd$, or with the sign $-$ and the operation $\wedge$, $\Box$, or $\rhd$, is called a **choice node**.
- A node with the sign $+$ and the operation $\Box$ or $\rhd$, or with the sign $-$ and the operation $\Diamond$ or $\lhd$ is called a **universal node**.

We now have a concise way to phrase the definitions given in [22]:

[10] We only consider paths from a leaf (which will always be a signed variable) to the root. So, for us, a 'path' will always be what is sometimes called a 'maximal path'.

Definition 1.3.3. A path in a generation tree is **Sahlqvist-harmless** if there is no universal node above a choice node in the path. If a path is not Sahlqvist-harmless, we call it (of course) **Sahlqvist-harmful**.
A term α is **ϵ-left Sahlqvist** if all ϵ-critical paths in $T^+(\alpha)$ are Sahlqvist-harmless.
A term β is **ϵ-right Sahlqvist** if all ϵ-critical paths in $T^-(\beta)$ are Sahlqvist-harmless.
An inequality $\alpha \leq \beta$ is a **Sahlqvist inequality** if there is some ϵ such that α is ϵ-left Sahlqvist and β is ϵ-right Sahlqvist.

In this chapter, we discuss the proof from [22] that all Sahlqvist inequalities are canonical in the distributive fragment.
We now also define the "proper Sahlqvist" fragment. It will turn out that we are able to "almost" reduce general Sahlqvist inequalities as defined above to *proper* Sahlqvist inequalities by means of the aforementioned trick with the operation $\mathbf{n}$. This "almost" will be made precise in the explanation of this trick, which will take up Section 2 of this chapter.

Definition 1.3.4. A term α is **proper ϵ-left Sahlqvist** if α is ϵ-left Sahlqvist and agrees with ϵ. Similarly, a term β is **proper ϵ-right Sahlqvist** if β is ϵ-right Sahlqvist and agrees with ϵ^∂.
A **proper Sahlqivst inequality** is of the form $\alpha \leq \beta$, where α is proper ϵ-left Sahlqvist and β is proper ϵ-right Sahlqvist, for some ϵ.

What we call *proper left and right Sahlqvist* here was simply called *left and right Sahlqvist* in [22], but we chose to add the word "proper" to emphasize the special nature of such terms: it is much harder for an inequality to be *proper* Sahlqvist than 'just' Sahlqvist. In particular, note the following property, which follows almost directly from the definitions:

Lemma 1.3.5. *Let α and β be BLM_τ-terms.*

(i) α is proper left Sahlqivst if and only if it agrees with some ϵ and all paths in $T^+(\alpha)$ are Sahlqvist-harmless.

(ii) β is proper right Sahlqivst if and only if it agrees with some ϵ and all paths in $T^-(\beta)$ are Sahlqvist-harmless.

Proof. The following chain of equivalences gives the proof of the first part.

$$\begin{aligned}
\alpha \text{ is proper left Sahlqvist} &\iff \text{there is } \epsilon \text{ s.t. } [\epsilon(\alpha) \text{ and } \alpha \text{ is } \epsilon\text{-left Sahlqvist}] \\
&\iff \text{there is } \epsilon \text{ s.t. } [\epsilon(\alpha) \text{ and every } \epsilon\text{-critical path in } T^+(\alpha) \text{ is harmless}] \\
&\iff [\text{there is } \epsilon \text{ s.t. } \epsilon(\alpha)] \text{ and every path in } T^+(\alpha) \text{ is harmless.}
\end{aligned}$$

Here, the last $\iff$ holds because if $\epsilon(\alpha)$ holds, then in fact every path in $T^+(\alpha)$ is ϵ-critical, by definition.
The proof of the second part is dual.

□

Remark 1.3.6. In [42], the definitions of this section are generalised to a *non-distributive* environment. There, canonicity is proved for a slightly larger fragment than the Sahlqvist fragment, which does allow for a binary choice node in the scope of a universal node, as long as one of the branches emerging from the choice node contains no variables, but only constants. In Section 5, we will encounter this fragment as the “proper inductive” fragment.

2 Collapsing terms

This section forms the heart of the chapter. We aim to expand on the proof of Lemma 5.14 in [22], and to show how the proof of that Lemma can be seen as providing a general method for rewriting inequalities.

In Subsection 2.1, we will describe “the **n**-trick”, thanks to which we are able to rewrite *general* well-behaved inequalities into equivalent *proper* well-behaved inequalities. The difference between *general* and *proper* inequalities is, as we saw in Lemma 1.3.5 of this chapter, that general inequalities only need to be well-behaved in certain critical paths, whereas proper inequalities need to be well-behaved everywhere. Then we show how to apply the **n**-trick to rewrite inequalities in a shape that we dub the ‘minimal collapse’ (Section 2.2).

This is essentially how, in [22], the ‘proper’ Sahlqvist fragment was considerably enlarged, while maintaining canonicity of the formulas in the fragment. Using the set-up given above, we try to give a proof that is a bit more intuitive than (although essentially the same as) the proof of the corresponding Lemma 5.14 in [22]. The main difference between our proof and the one given in [22] is presentational: we separate the general syntactic step from the specific reduction given there for Sahlqvist inequalities, which will hopefully stimulate attempts to apply this general method to other syntactic fragments of interest as well.

2.1 The binary operation n

In this subsection, we isolate the main technical step, introduced in [22], which is used there for the reduction of Sahlqvist inequalities to proper Sahlqvist inequalities. We believe that there is an advantage in discussing this step and the ensuing algorithm ‘outside’ of the Sahlqvist setting: the following discussion hopefully makes it apparent that it actually yields a very general reduction method, which could be of use in other proofs ‘by reduction’.

The **n**-trick relies on the definition of a new binary operation, called **n** (for **n**ot less than or equal).

Definition 2.1.1. In any BLM_τ $\mathbb{A}$, define the binary operation $\mathbf{n}^{\mathbb{A}} : \mathbb{A}^2 \to \mathbb{A}$ by

$$\mathbf{n}^{\mathbb{A}}(a,b) := \begin{cases} \bot & \text{if } a \leq b \\ \top & \text{if } a \not\leq b. \end{cases}$$

Remark 2.1.2. If we regard the poset reduct $\mathbb{P}_A := \langle A, \leq \rangle$ of the algebra $\mathbb{A}$ as a category, we observe that the bifunctor $\mathsf{Hom} : (\mathbb{P}_A)^{\mathrm{op}} \times \mathbb{P}_A \to \mathsf{Set}$ has the following value on objects:

$$\mathsf{Hom}_{\mathbb{P}_A}(a,b) = \begin{cases} \{*\} & \text{if } a \leq b \\ \emptyset & \text{if } a \not\leq b. \end{cases}$$

Noting that $\{*\}$ and $\bot$ are the initial objects of the categories $\mathsf{Set}^{\mathrm{op}}$ and $\mathbb{P}_A$, respectively, and that $\emptyset$ and $\top$ are the terminal objects of these categories, the similarity with the definition of the operation $\mathbf{n}$ becomes striking. This suggests that the operation $\mathbf{n}$ as defined above is a "Hom-functor in disguise". We do not develop this idea any further at this point, but we think it is an interesting observation which may be non-accidental and useful, especially with respect to future extensions of the theory of canonical extenstions from posets to categories.

The following basic observation about $\mathbf{n}$, which will be of importance later, is stated but not proven in [22]. We give the proof here.

Lemma 2.1.3. *Let $\mathbb{A}$ be any BLM_τ. Then $\mathbf{n}$ agrees with $(1, \partial)$, i.e. $\mathbf{n}^{\mathbb{A}} : \mathbb{A}^{(1,\partial)} \to \mathbb{A}$ is order-preserving.*

Proof. If $a, a', b \in \mathbb{A}$ and $a \leq a'$, then either $\mathbf{n}(a, b) = \bot$, in which case automatically $\mathbf{n}(a, b) \leq \mathbf{n}(a', b)$, or $\mathbf{n}(a, b) = \top$, which means $a \not\leq b$. But then also $a' \not\leq b$, since $a \leq a'$ by assumption. Hence, in this case we also have $\top = n(a, b) \leq n(a', b) = \top$.
If $a, b, b' \in \mathbb{A}$ and $b \leq b'$, then either $\mathbf{n}(a, b) = \top$, in which case automatically $\mathbf{n}(a, b) \geq \mathbf{n}(a, b')$, or $\mathbf{n}(a, b) = \bot$, which means $a \leq b$. But then also $a \leq b'$, since $b \leq b'$ by assumption. Hence, in this case we also have $\bot = n(a, b) \geq n(a, b') = \bot$. □

Another important property of the operation $\mathbf{n}$ is that, like all standard operations of the modal similarity type, it is itself stable under taking canonical extensions.

Lemma 2.1.4. *For any BLM_τ $\mathbb{A}$, we have $(\mathbf{n}^{\mathbb{A}})^\sigma = \mathbf{n}^{\mathbb{A}^\sigma}$.*

Proof. See Lemma 5.15 of [22]. □

The use of the operation $\mathbf{n}$ can be explained intuitively by saying that $\mathbf{n}$ enables us to 'extract subterms' which occur in a BLM_τ-inequality. The following lemma, which corresponds to Lemma 5.13(1) of [22], shows how such an 'extraction' works in detail.

Lemma 2.1.5 ("**n**-trick")**.** *Let α and β be BLM_τ-terms. Let s be a subterm of α such that $+s \prec +\alpha$. Then, for any BLM_τ $\mathbb{A}$, we have*

$$\mathbb{A} \models \alpha \leq \beta \iff \mathbb{A} \models \alpha[z/s] \leq \beta \vee \mathbf{n}(z, s)$$

where z is a new variable which does not occur in α or β.

Proof. Suppose that $x_1, \ldots, x_n$ are the variables occurring in α and β.
For the 'if' direction, let $a_i \in \mathbb{A}$ $(1 \leq i \leq n)$ be arbitrary. Then

$$\alpha^{\mathbb{A}}(a_1, \ldots, a_n) = (\alpha[z/s])^{\mathbb{A}}(a_1, \ldots, a_n, s^{\mathbb{A}}(a_1, \ldots, a_n)).$$

By assumption, $\mathbb{A} \models \alpha[z/s] \leq \beta \vee \mathbf{n}(z, s)$, so in particular

$$(\alpha[z/s])^{\mathbb{A}}(a_1, \ldots, a_n, s^{\mathbb{A}}(a_1, \ldots, a_n)) \leq \beta^{\mathbb{A}}(a_1, \ldots, a_n) \vee \mathbf{n}^{\mathbb{A}}(s^{\mathbb{A}}(a_1, \ldots, a_n), s^{\mathbb{A}}(a_1, \ldots, a_n)).$$

Now, $\mathbf{n}^{\mathbb{A}}$ is defined to take the value $\bot$ at the pair $(s^{\mathbb{A}}(a_1, \ldots, a_n), s^{\mathbb{A}}(a_1, \ldots, a_n))$, so we conclude

$$\alpha^{\mathbb{A}}(a_1, \ldots, a_n) \leq \beta^{\mathbb{A}}(a_1, \ldots, a_n).$$

For the 'only if' direction, take arbitrary $a_i \in \mathbb{A}$ $(1 \leq i \leq n)$ and $b \in \mathbb{A}$. We want to show that

$$(\alpha[z/s])^{\mathbb{A}}(a_1, \ldots, a_n, b) \leq \beta^{\mathbb{A}}(a_1, \ldots, a_n) \vee \mathbf{n}^{\mathbb{A}}(b, s^{\mathbb{A}}(a_1, \ldots, a_n, b)).$$

We distinguish two cases:

- $b \leq s^{\mathbb{A}}(a_1, \ldots, a_n)$. Then we have

$$\begin{aligned}(\alpha[z/s])^{\mathbb{A}}(a_1, \ldots, a_n, b) &\leq (\alpha[z/s])^{\mathbb{A}}(a_1, \ldots, a_n, s^{\mathbb{A}}(a_1, \ldots, a_n)) \\ &= \alpha^{\mathbb{A}}(a_1, \ldots, a_n) \\ &\leq \beta^{\mathbb{A}}(a_1, \ldots, a_n) \\ &= \beta^{\mathbb{A}}(a_1, \ldots, a_n) \vee \mathbf{n}^{\mathbb{A}}(b, s^{\mathbb{A}}(a_1, \ldots, a_n)).\end{aligned}$$

  Here, the first inequality holds because $b \leq s^{\mathbb{A}}(a_1, \ldots, a_n)$ and $+s \prec +\alpha$ (so $+z \prec +\alpha[z/s]$), the second equality holds by the inductive definition of term functions, the third inequality holds by assumption, and the last equality holds since $\mathbf{n}^{\mathbb{A}}(b, s^{\mathbb{A}}(a_1, \ldots, a_n)) = \bot$ by definition of $\mathbf{n}$.

- $b \not\leq s^{\mathbb{A}}(a_1, \ldots, a_n)$. Then $\mathbf{n}^{\mathbb{A}}(b, s^{\mathbb{A}}(a_1, \ldots, a_n)) = \top$ by definition, so we immediately have that

$$(\alpha[z/s])^{\mathbb{A}}(a_1, \ldots, a_n, b) \leq \beta^{\mathbb{A}}(a_1, \ldots, a_n) \vee \mathbf{n}^{\mathbb{A}}(b, s^{\mathbb{A}}(a_1, \ldots, a_n)) = \top. \qquad \square$$

We observe next that the proof of the above lemma can be dualized to the cases $-s \prec +\alpha$, $+s \prec -\beta$ and $-s \prec -\beta$.

Corollary 2.1.6. *Let α and β be BLM_τ-terms. Let s be a BLM_τ-term and z a new variable not occuring in α or β.*

(i) If $+s \prec +\alpha$, then $\alpha \leq \beta \iff \alpha[z/s] \leq \beta \vee \mathbf{n}(z, s)$.

(ii) If $-s \prec +\alpha$, then $\alpha \leq \beta \iff \alpha[z/s] \leq \beta \vee \mathbf{n}(s, z)$.

(iii) If $+s \prec -\beta$, then $\alpha \leq \beta \iff \alpha \leq \beta[z/s] \vee \mathbf{n}(z, s)$.

(iv) If $-s \prec -\beta$, then $\alpha \leq \beta \iff \alpha \leq \beta[z/s] \vee \mathbf{n}(s, z)$.

Remark 2.1.7. In Section 5.2 of this chapter, we will observe an interesting connection between Lemma 2.1.5 and the Ackermann Lemma on substitution of terms.

2.2 The minimal collapses

In this subsection, we show how to use the result from the previous subsection to rewrite general inequalities into equivalent *proper* inequalities.

Proposition 2.2.1. *Let α and β be BLM_τ terms. For any order type ϵ, there exist an order type $\bar{\epsilon}$ extending ϵ, and terms α', β' and γ such that:*

(i) $\alpha \leq \beta \iff \alpha' \leq \beta' \vee \gamma$,

(ii) $\bar{\epsilon}(\alpha')$, $\bar{\epsilon}^{\partial}(\beta')$ and $\bar{\epsilon}(\gamma)$,

(iii) α' and β' are obtained by substituting subterms of α and β, respectively, with new variables.

Proof. The idea of the proof is to substitute parts of the terms α and β which are badly behaved ('junk') by fresh variables, moving all the 'junk' into a long, uniform disjunction γ of $\mathbf{n}$'s on the right side of the inequality. We will call α' and β' the *ϵ-minimal collapses* of α and β, respectively. The proof will occupy the rest of this section, and naturally falls apart into three parts:

- The algorithm which defines the minimal collapses of α and β;
- A relevant concrete example of how the algorithm works;
- The formal proof that the algorithm always works.

The algorithm

Let $x_1, \dots, x_n$ be the variables occurring in α and β. Let $\bar{\epsilon}_i := \epsilon_i$ for $1 \le i \le n$. We now explicitly construct the rest of the order type $\bar{\epsilon}$ and the terms α', β' and γ, as follows.
In the tree $T^+(\alpha)$, first colour all the nodes which are on at least one ϵ-critical path *red*, say, and then colour all the remaining nodes *green*.
Let $p_1, \dots, p_m$ be the paths in this tree which contain some *green* node. Now, for any $1 \le i \le m$, let s_i be the maximal *green* node on the path p_i. Let $\bar{\epsilon}_{n+i} \in \{1, \partial\}$ be the sign of the node s_i in the tree $T^+(\alpha)$, and let $z_1, \dots, z_m$ be m new variables, which are not among the variables x_i. Now put

$$\alpha'(x_1, \dots, x_n, z_1, \dots, z_m) := \alpha[z_1/s_1, \dots, z_m/s_m].$$

We do the exact analogous thing for the tree $T^-(\beta)$: after colouring the ϵ-critical paths *red*, we have m' paths $q_1, \dots, q_{m'}$, each of which contains some *green* node, and their maximal nodes $t_1, \dots, t_{m'}$. We let $\bar{\epsilon}_{n+m+i}$ be the sign of the node t_i in the tree $T^-(\beta)$, let $w_1, \dots, w_{m'}$ be m' new variables, and put

$$\beta'(x_1, \dots, x_n, w_1, \dots, w_{m'}) := \beta[w_1/t_1, \dots, w_{m'}/t_{m'}].$$

Now let

$$\gamma(x_1, \dots, x_n, z_1, \dots, z_m, w_1, \dots, w_{m'}) := \bigvee_{i=1}^{m} \mathbf{n}((z_i, s_i)^{\bar{\epsilon}_{n+i}}) \vee \bigvee_{j=1}^{m'} \mathbf{n}((w_j, t_j)^{\bar{\epsilon}_{n+m+j}}),$$

where we use the notational convention that $(a, b)^1 := (a, b)$ and $(a, b)^\partial = (b, a)$.
Before completing the proof by showing that α', β' and γ thus constructed satisfy the required properties, we think it is useful to give an example of how the described algorithm works.

An example

Example 2.2.2. Let $\alpha = \lhd(x \wedge y) \vee x$ and $\beta = \Diamond((x \wedge y) \vee \lhd y)$, and let $\epsilon = (1, 1)$. The coloured trees $T^+(\alpha)$ and $T^-(\beta)$ are depicted below.

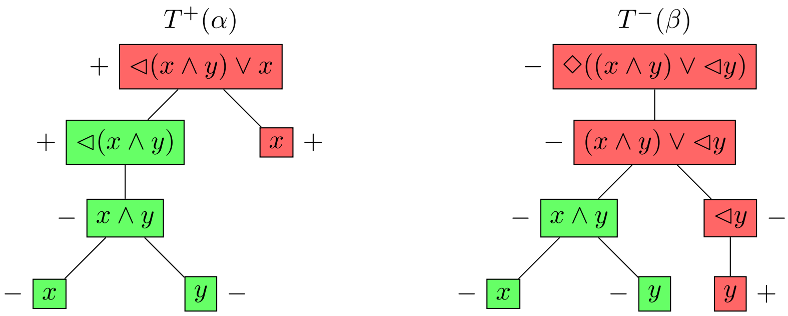

In this case, both α and β have one maximal green node, so we put $s := \lhd(x \wedge y)$ and $t := x \wedge y$. Then $+s \prec +\alpha$ and $-t \prec -\beta$. So, the algorithm tells us to introduce two new variables, z and w, and to extend ϵ to $\bar{\epsilon}$ by putting $\bar{\epsilon}_z := 1$ and $\bar{\epsilon}_w := \partial$. Now α' and β' are defined as $\alpha[z/s]$ and $\beta[w/t]$, respectively. Their generation trees are depicted below.

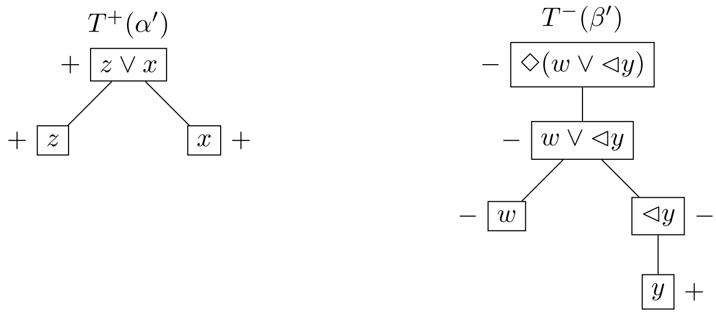

In these trees, we can see clearly that $\bar{\epsilon}(\alpha')$ and $\bar{\epsilon}^\partial(\beta')$.
According to the algorithm, γ is defined as

$$\gamma := \mathbf{n}(z, \lhd(x \wedge y)) \vee \mathbf{n}(x \wedge y, w).$$

The tree $T^+(\gamma)$ looks as follows:

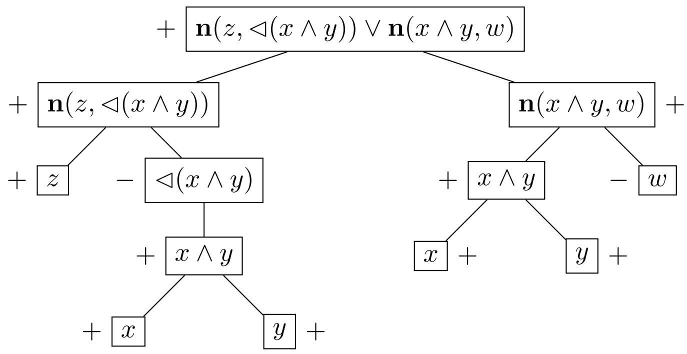

Note that γ agrees with $\bar{\epsilon} = (1, 1, 1, \partial)$. This is not a coincidence, of course, but it is rather subtle to see why precisely this is the case. The answer to this will be given in the proof below.

Proof of correctness of the algorithm

We now show in detail that the algorithm indeed gives the α', β' and γ which satisfy the properties (i)-(iii). The details are unfortunately rather involved, and we suggest the reader checks back to see what the details amount to in the above example.

(i) The fact that $\alpha \leq \beta$ is equivalent to $\alpha' \leq \beta' \vee \gamma$ is immediate by applying Corollary 2.1.6 $m + m'$ times: first, for the s_i, we are in case (i) or (ii), depending on the sign $\bar{\epsilon}_{n+i}$, and then, for the t_j, we are in case (iii) or (iv), depending on the sign $\bar{\epsilon}_{n+m+j}$.

(ii) We first show that $\bar{\epsilon}(\alpha')$.

If some 'old' variable x_i occurs in α', then it must be coloured *red* in $T^+(\alpha)$, because otherwise it would be the leaf of a partially green path in $T^+(\alpha)$, but all green paths are replaced by a fresh variable. The fact that x_i is coloured *red* means precisely that it is an ϵ-critical occurrence, so it agrees with ϵ.

For a 'new' variable z_i in α, the sign $\bar{\epsilon}_{n+i}$ was precisely chosen such that z_i agrees with $\bar{\epsilon}$.

The proof that $\bar{\epsilon}^\partial(\beta')$ is dual, because we coloured the *negative* generation tree of β.

Finally, we show that $\bar{\epsilon}(\gamma)$.

For an occurrence of an old variable x_i in γ, it must be in one of the s_i or t_j.

- If the occurrence of x_i is in one of the s_i, there are two possibilities:
  (a) $+s_i \prec +\alpha$, so $\bar{\epsilon}_{n+i} = 1$ by definition. Since s_i was coloured *green* in $T^+(\alpha)$, we have in particular that the occurrence of x_i was not critical. So the sign of this occurrence of x_i in $T^+(\alpha)$ was *opposite* to the sign ϵ_i, that is, $\epsilon_i^\partial x_i \prec +\alpha$. So $\epsilon_i^\partial x_i \prec +s_i$, and since $\mathbf{n}$ agrees with $(1, \partial)$ we get $\epsilon_i x_i \prec \mathbf{n}(z_i, s_i)$.
  (b) $-s_i \prec +\alpha$, so $\bar{\epsilon}_{n+i} = \partial$ by definition. The reasoning is dual to that in the previous item and leads to the conclusion that $\epsilon_i x_i \prec \mathbf{n}((z_i, s_i)^\partial)$.
- If the occurrence of x_i is in one of the t_j, there are two analogous possibilities, and we can again conclude either $\epsilon_i x_i \prec \mathbf{n}(w_i, t_i)$ or $\epsilon_i x_i \prec \mathbf{n}((w_i, t_i)^\partial)$, depending on the sign of $\bar{\epsilon}_{n+m+j}$.

For an occurrence of a new variable z_i or w_j in γ, it occurs as the first coordinate of $\mathbf{n}$ precisely when the corresponding sign in $\bar{\epsilon}$ is 1, and it occurs as the second coordinate of $\mathbf{n}$ when the corresponding sign in $\bar{\epsilon}$ is ∂, by the definition of γ. This means, again because $\mathbf{n}$ agrees with $(1, \partial)$, that $\bar{\epsilon}_{n+i} z_i \prec +\gamma$ and $\bar{\epsilon}_{n+m+j} w_j \prec +\gamma$, for all $1 \leq i \leq m$ and $1 \leq j \leq m'$, as required.

(iii) is immediate from the construction of α' and β'.

Related to this point, we further note that the algorithm in fact gives us slightly more: the tree $T^+(\alpha')$ is nothing more than a *truncated version* of the tree $T^+(\alpha)$ in which only *green* nodes are eliminated. The same assertion holds for $T^-(\beta')$ and $T^-(\beta)$. □

3 After the collapse

The motivation for the minimal collapse algorithm, outlined in the previous section, was to rewrite any general inequality as a proper inequality. One could argue that, until now, we only succeeded in doing this partially. After rewriting an inequality $\alpha \leq \beta$ as $\alpha' \leq \beta' \vee \gamma$, where α' and β' are the minimal collapses of α and β, there are now two main hurdles left:

- Are α' and β' really *proper* (Sahlqvist, inductive) terms?
- Although we wanted to rewrite $\alpha \leq \beta$ as a proper inequality $\alpha' \leq \beta'$, we ended up with something slightly worse, namely an equivalent inequality of the form $\alpha' \leq \beta' \vee \gamma$, where γ is a formula which agrees with ϵ. Can we maintain canonicity of the inequality with such an added term on the right?

The second of these questions is dealt with in detail in Lemma 5.11 of [22], which we will briefly report in Section 3.2. Regarding the first question, however, it is only stated in the proof of Lemma 5.14 of that paper that the answer is affirmative, but a detailed proof is omitted. We will give a detailed proof in this section.
Spelling out this proof is important, in our view, because the fact that the answer to the first question is 'yes' is the crucial observation which makes the minimal collapse algorithm useful: it is the only place in the proof where we need to use that the collapse we are taking is the *minimal* one (i.e., we collapse at the maximal green nodes). Indeed, the crucial observation here is that *only* the *minimal* collapses guarantee that when Sahlqvist inequalities are fed to this algorithm, the outcome is still Sahlqvist (and proper, as it is meant to be by construction). As we will see in detail in Section 5, this delicate step will fail for the inductive formulas, and this is the reason why the automatic application of the 'minimal collapse' methodology fails when trying to prove canonicity of the inductive fragment by algebraic means.

3.1 Minimal collapses are proper Sahlqvist

Lemma 3.1.1. *If $\alpha \leq \beta$ is a Sahlqvist inequality, and this is witnessed by an order type ϵ, then the ϵ-minimal collapses α' and β' are proper left and right Sahlqvist, respectively.*

Proof. We show that α' is proper left Sahlqvist. By assumption, α is ϵ-left Sahlqvist.
By Lemma 1.3.5, it suffices to show that α' is uniform and all paths in $T^+(\alpha')$ are Sahlqvist-harmless. By Proposition 2.2.1(ii), we have $\bar{\epsilon}(\alpha')$, so α' is uniform, since it agrees with $\bar{\epsilon}$.
Let y be an occurrence of a variable in α'. There are two possibilities:

- y is an 'old' variable x_i which already occurs in α.

  Then $\epsilon x_i \prec +\alpha$ since we have that $\bar{\epsilon}(\alpha')$, $T^+(\alpha')$ is a subtree of $T^+(\alpha)$, and $\bar{\epsilon}$ extends ϵ. So, the occurrence of x_i is ϵ-critical in α. Therefore, its path to the root in $T^+(\alpha)$ must be Sahlqvist-harmless, because any ϵ-critical path in $T^+(\alpha)$ is Sahlqvist-harmless. Also, the path from x_i to the root in $T^+(\alpha)$ was coloured *red* entirely by the algorithm, so that it remains intact in the truncated version $T^+(\alpha')$ of $T^+(\alpha)$, proving that this path is Sahlqvist-harmless in $T^+(\alpha')$ as well.

- y is a 'new' variable z_i which replaces a subterm s_i of α.

  Suppose, to obtain a contradiction, that the path from z_i to the root of $T^+(\alpha')$ is not Sahlqvist-harmless, so there is a choice node n'_c below a universal node n'_u somewhere in the path. Since $T^+(\alpha')$ is a truncated version of $T^+(\alpha)$, there are a corresponding choice node n_c and a corresponding universal node n_u in the tree $T^+(\alpha)$. The node n_c must have been coloured *red* by the algorithm, because the node s_i was the *maximal* green node on the path. The fact that n_c was coloured *red* means that there is an ϵ-critical path going through n_c in $T^+(\alpha)$, but this ϵ-critical path would then fail to be Sahlqvist-harmless, contradicting that α is ϵ-left Sahlqvist.

The proof that β' is proper right Sahlqvist is completely analogous. □

3.2 Positive terms on the right

The second question posed in the introduction of this section was: what to do with the ϵ-agreeing term γ that the minimal collapse algorithm adds to the right side of the inequality? It turns out that we are able to make one useful interpolating step when taking the canonical extension of a term of the form $\beta \vee \gamma$, and then (see Lemma 4.1.2 below) we can use the fact that terms which agree with ϵ are σ-contracting. The details are in the following lemma and corollary, a restatement of Lemma 5.11 in [22]. Note that it is based on a general fact about monotone maps between distributive lattices and their canonical extensions, and that it is not specific for Sahlqvist inequalities at all.

Lemma 3.2.1. *If $f, g : \mathbb{A} \to \mathbb{B}$ are maps between distributive lattices such that f is antitone and g is monotone, then $(f \vee g)^\sigma \leq f^\pi \vee g^\sigma$.*

Corollary 3.2.2. *If β' and γ are* BLM$_\tau$*-terms, β' agrees with ϵ^∂, and γ agrees with ϵ, then, for any* DMA *$\mathbb{A}$, we have $(\beta'^{\mathbb{A}} \vee \gamma^{\mathbb{A}})^\sigma \leq (\beta'^{\mathbb{A}})^\pi \vee (\gamma^{\mathbb{A}})^\sigma$.*

4 Canonicity of the Sahlqvist fragment

We are now ready to prove the canonicity of the Sahlqvist fragment, combining the stability of proper Sahlqvist inequalities (Theorem 4.1.1 below) from [22] and the minimal collapse algorithm that we discussed in Section 2.

4.1 Stability and contraction of certain terms

The main result needed to prove the canonicity of the proper Sahlqvist fragment is

Theorem 4.1.1. *Let α and β be* DMA*-terms and $\mathbb{A}$ any* DMA*.*

(i) If α is proper left Sahlqvist, then $\alpha^{\mathbb{A}}$ is σ-stable.

(ii) If β is proper right Sahlqvist, then $\beta^{\mathbb{A}}$ is π-stable.

Proof. This is Lemma 5.10 from [22]. □

Finally, a last lemma that we need, in order to deal with the extra term γ, is the following.

Lemma 4.1.2. *If a* BLM$_\tau$*-term* γ *agrees with* ϵ*, then* γ *is* σ*-contracting, i.e.* $(\gamma^{\mathbb{A}})^\sigma \leq \gamma^{\mathbb{A}^\sigma}$ *for any* $\mathbb{A}$.

Proof. This is stated as Lemma 5.5 for distributive lattices in [22], but the proof never uses the fact that $\mathbb{A}$ is a distributive lattice. □

4.2 Canonicity of general Sahlqvist terms

Theorem 4.2.1. *Every Sahlqvist inequality* $\alpha \leq \beta$ *is canonical.*

Proof. Let $\alpha \leq \beta$ be a Sahlqivst inequality and let $\mathbb{A}$ be any DMA such that $\mathbb{A} \models \alpha \leq \beta$. Then $\alpha^{\mathbb{A}} \leq \beta^{\mathbb{A}}$, and we need to show that $\alpha^{\mathbb{A}^\sigma} \leq \beta^{\mathbb{A}^\sigma}$.
Let $\alpha' \leq \beta' \vee \gamma$ be the inequality equivalent to $\alpha \leq \beta$, obtained by the minimal collapse algorithm (Proposition 2.2.1). Then $\mathbb{A} \models \alpha' \leq \beta' \vee \gamma$.
By Theorem 4.1.1, $(\alpha')^{\mathbb{A}}$ and $(\beta')^{\mathbb{A}}$ are σ- and π-stable, respectively. Also, γ is σ-contracting since it agrees with ϵ (Lemma 4.1.2). Also, using Corollary 3.2.2, we have $(\beta'^{\mathbb{A}} \vee \gamma^{\mathbb{A}})^\sigma \leq (\beta'^{\mathbb{A}})^\pi \vee (\gamma^{\mathbb{A}})^\sigma$.
Combining these facts gives us:

$$(\alpha')^{\mathbb{A}^\sigma} = (\alpha'^{\mathbb{A}})^\sigma \leq (\beta'^{\mathbb{A}} \vee \gamma^{\mathbb{A}})^\sigma \leq (\beta'^{\mathbb{A}})^\pi \vee (\gamma^{\mathbb{A}})^\sigma = \beta'^{\mathbb{A}^\sigma} \vee (\gamma^{\mathbb{A}})^\sigma \leq \beta'^{\mathbb{A}^\sigma} \vee \gamma^{\mathbb{A}^\sigma},$$

that is, $\mathbb{A}^\sigma \models \alpha' \leq \beta' \vee \gamma$. Since $\alpha' \leq \beta' \vee \gamma$ is equivalent to $\alpha \leq \beta$, we conclude that $\mathbb{A}^\sigma \models \alpha \leq \beta$, as we needed to show. □

This concludes our discussion of the algebraic canonicity of the Sahlqvist fragment. Before we move on to discuss the inductive fragment in the next section, we now summarize the global form of this proof. In the above discussion, we have identified the following crucial steps:

(i) Reduce *any* BLM$_\tau$ inequality (independently of its being Sahlqvist ot not) to one which agrees with some ϵ, using the minimal collapse algorithm (Section 2.2);

(ii) Show that this reduction 'preserves Sahlqvist', i.e., that when the algorithm is applied to Sahlqvist inequalities, the 'minimal collapse' algorithm guarantees that no Sahlqvist-harmful paths are anywhere in the output inequality, even though there are potentially more critical paths than in the original formula (Lemma 3.1.1);

(iii) For the resulting *proper* Sahlqvist inequality, we have a direct canonicity proof (Corollary 3.2.2, Theorem 4.1.1, Lemma 4.1.2).

Our reason for summarizing the proof in this way, emphasizing the syntactic steps rather than the actual canonicity proof, is that these are the steps which can not be automatically translated when trying to prove canonicity of the inductive fragment algebraically. In the next section, we will discuss the difficulties in trying to follow this proof method and we state a negative result.

5 The inductive fragment

The analysis of the canonicity mechanism that we gave at the end of the previous section raises the question whether the proof carries over to a canonicity proof of more general syntactic fragments of distributive modal logic. Conradie and Palmigiano [7] prove canonicity of the *inductive fragment* of distributive modal logic via correspondence. In this section, we discuss the possibility of obtaining an *algebraic* proof of that result.

5.1 Definitions and comparison with Sahlqvist fragment

Let us first define the inductive fragment.

Definition 5.1.1. Let α be a DMA-term. Let ϵ be an order type and Ω a strict partial order, i.e., an irreflexive and transitive relation, on the variables occurring in α.
We say that the pair (Ω, ϵ) **solves a binary choice node** $c(s,t)$ **for the variable** p_i in a certain tree if the following conditions hold:

- s agrees with ϵ^∂,
- $\epsilon_i p_i \prec +t$,
- for every variable p_j occurring in s we have $p_j <_\Omega p_i$.

We say that a path is (Ω, ϵ)**-harmless** if (Ω, ϵ) solves every choice node below a universal node in the path for the variable at the leaf.
A term α is (Ω, ϵ)**-left inductive** if all ϵ-critical paths in $T^+(\alpha)$ are (Ω, ϵ)-harmless.
A term β is (Ω, ϵ)**-right inductive** if all ϵ-critical paths in $T^-(\beta)$ are (Ω, ϵ)-harmless.
An **inductive inequality** is of the form $\alpha \leq \beta$ such that α is (Ω, ϵ)-left inductive and β is (Ω, ϵ)-right inductive, for some Ω and ϵ.

Comparing this definition with Definition 1.3.3, of the Sahlqvist fragment for distributive modal logic, it becomes apparent that it is of the same general form of "no forbidden configurations in critical paths", only the definition of 'forbidden configurations' is loosened considerably.
If we try to apply the algebraic canonicity proof for the Sahlqvist fragment to the inductive fragment, we see immediately that step (i) of the outline at the end of Section 4 still works: in fact, it does not use the fact that the inequality we are dealing with is Sahlqvist, at all.
Let us now first look at step (iii). To be able to say if this step carries over, we need to know what 'proper inductive inequalities' are. The following definition naturally generalises the definition for the Sahlqvist fragment.

Definition 5.1.2. An inequality $\alpha \leq \beta$ is **proper inductive** if it is inductive for some (Ω, ϵ), α agrees with ϵ, and β agrees with ϵ^∂.

Incidentally, this definition coincides with Suzuki's small enlargement of the Sahlqvist fragment that we discussed in Remark 1.3.6. So, we indeed have a proof that 'proper inductive' formulas are canonical in [42].
However, what about step (ii)? Does the minimal collapse algorithm preserve inductivity of formulas? It turns out that the answer to this question is negative, and this is the hurdle in

directly applying the proof method from [22] to the inductive fragment. The following example shows that the minimal collapse algorithm does not preserve inductivity of formulas.

Example 5.1.3. Let α be the formula $\Box(\lhd x \vee y)$. This formula is (Ω, ϵ)-left inductive for $\epsilon = (1,1)$ and $x <_\Omega y$. However, the result of applying the minimal collapse algorithm to α is $\alpha' = \Box(z \vee y)$, and $\bar{\epsilon} = (1,1,1)$. Now, there is no $\overline{\Omega}$ such that α' is $(\overline{\Omega}, \bar{\epsilon})$-left inductive: for this, we would need a partial order Ω such that both $z <_\Omega y$ (since the path from y is critical) and $y <_\Omega z$ (since the path from z is critical), which is clearly impossible.

We have thus proved

Proposition 5.1.4. *There exists an inductive inequality $\alpha \leq \beta$ such that the inequality $\alpha' \leq \beta'$, where α' an β' are obtained from applying the minimal collapse algorithm (Proposition 2.2.1) is not inductive.*

5.2 Connections with Ackermann's Lemma

From Proposition 5.1.4, we conclude that, interestingly, to prove canonicity of the inductive fragment by algebraic means, we will need a proof strategy that substantially improves on some aspects of the algebraic canonicity proof of the Sahlqvist fragment. In particular, we expect that we need an enhanced version of the canonical operation $\mathbf{n}$, which would enable us to collapse more nodes, while maintaining equivalence.

Although we were not able to find such an enhanced version of $\mathbf{n}$ yet, in the quest for it, we discovered some interesting parallels of the original operation $\mathbf{n}$ with Ackermann's Lemma, which is used for the correspondence and canonicity proofs for the inductive fragment in [7].

We first state Ackermann's Lemma, in a slightly different form from the way it is stated in [7]. We do this mainly for presentational purposes, in order to make the connection with the key lemma for the operation $\mathbf{n}$ (Lemma 2.1.5) explicit.

Lemma 5.2.1 (Ackermann's Lemma). *Let α and β be BLM_τ-terms. Let z be a variable, occurring only in α, such that $+z \prec +\alpha$. Then, for any BLM_τ $\mathbb{A}$ and any valuation V of the variables occurring in α and β, the following are equivalent:*

(i) $\langle \mathbb{A}, V \rangle \models \alpha \leq \beta$

(ii) *There exists a BLM_τ-term γ, using only variables from α and β, such that $\langle \mathbb{A}, V \rangle \models z \leq \gamma$ and $\langle \mathbb{A}, V \rangle \models \alpha[\gamma/z] \leq \beta$.*

Proof. The direction (i) $\Rightarrow$ (ii) holds trivially by putting $\gamma := z$.

For the direction (ii) $\Rightarrow$ (i), suppose γ is a BLM_τ-term such that both $\langle \mathbb{A}, V \rangle \models \alpha[\gamma/z] \leq \beta$ and $\langle \mathbb{A}, V \rangle \models z \leq \gamma$.

Let $x_1, \ldots, x_n, z$ be the variables occurring in α and β. Then, since $\langle \mathbb{A}, V \rangle \models z \leq \gamma$, we have $V(z) \leq \gamma^{\mathbb{A}}(V(x_1), \ldots, V(x_n), V(z))$. By Lemma 1.2.4, $\alpha^{\mathbb{A}}$ is monotone in the last coordinate, and by assumption, $(\alpha[\gamma/z])^{\mathbb{A}} \leq \beta^{\mathbb{A}}$, so

$$\begin{aligned}
\alpha^{\mathbb{A}}(V(x_1), \ldots, V(x_n), V(z)) &\leq \alpha^{\mathbb{A}}(V(x_1), \ldots, V(x_n), \gamma^{\mathbb{A}}(V(x_1), \ldots, V(x_n), V(z))) \\
&= (\alpha[\gamma/z])^{\mathbb{A}}(V(x_1), \ldots, V(x_n)) \\
&\leq \beta^{\mathbb{A}}(V(x_1), \ldots, V(x_n)).
\end{aligned}$$
$\square$

Remark 5.2.2. We note the similarity of the proof of Lemma 5.2.1 with parts of the proof of Lemma 2.1.5. In a sense, it seems that the **n**-trick is a way of encoding Ackermann's Lemma *inside the logic*, whereas Ackermann's Lemma as it is stated here is rather a result *about the logic*, containing, for example, existential quantification over the term γ. One could then summarize this observation as: "**n**-trick = Ackermann + encoding".

We think that making this similarity precise will throw more light on the reasons why the proof method, that we have dubbed "the **n**-trick" in this chapter, works precisely for the canonicity proof of the Sahlqvist fragment, but fails for the inductive fragment.

Moreover, the correspondence and canonicity proof for the inductive fragment that is given in [7] via general frames heavily relies on Ackermann's Lemma as one of its cornerstones. Therefore, we expect that spelling out the details of these parallels will also provide the key to an algebraic canonicity proof for the inductive fragment, in the spirit of the algebraic canonicity proof of the Sahlqvist fragment which we examined in this chapter.

Conclusion

The central theme and red thread of this thesis is the notion of *canonical extensions*. We now collect the results and open questions emerging from our investigation of canonical extensions that we reported on in this thesis.

1 Reconstructing the lost duality

In the first chapter, we 'reconstructed' part of the duality square for the canonical extension of partially ordered sets, so that we now have the following diagram, which is a bit more complete than Diagram 1.1 of Chapter 1.

$$\begin{tikzcd}
\mathsf{PLat} \arrow[rr, shift left, "(\cdot)_+"] & & \mathsf{RSPol} \arrow[ll, shift left, "(\cdot)^+"] \\
& & \\
\mathsf{Poset} \arrow[uu, dashed, "(\cdot)^\sigma"] \arrow[uurr, dashed, "(\cdot)_\bullet"] \arrow[rr, shift left, "?"] & & ? \arrow[ll, shift left, "?"] \arrow[uu, "U\ ?"']
\end{tikzcd}$$

Figure 1: *Partial duality square for posets, with* $(\cdot)_\bullet$

Furthermore, we gave dual first-order characterisations of adjoint pairs, hemimorphisms and homomorphisms between perfect lattices, following an alternative route to [16].
From a methodological point of view, we studied duality from the perspective of a hierarchy of increasingly general dualities. We regard this hierarchy as 'telescopic', in the sense that each duality encompasses and generalises the previous one: Stone duality for Boolean algebras can be viewed as a special case of Priestley duality for distributive lattices, part of which can be viewed as a special case of Urquhart duality for bounded lattices. However, Urquhart duality has a narrower scope, because it only provides a duality for the category of bounded lattices with *surjective* morphisms.
An important feature of our methodology is that we embraced a modal logic perspective on these

dualities, in the following three ways:

- Rather than lattice homomorphisms, we regarded *join- and meet-hemimorphisms* as the first-class citizens in the category of perfect lattices. This idea agrees with the earlier work of Jónsson and Tarski on canonical extensions [31], as applied to modal logic, where operators and dual operators are the join- and meet-hemimorphisms.
- We derived our first-order dual characterisations using methods inspired by the correspondence theory for modal logic. For instance, we regarded a homomorphism as a two-faced object, whose two personalities are '$\Diamond$' and '$\Box$', satisfying the two Sahlqvist inequalities $\Diamond \leq \Box$ and $\Box \leq \Diamond$.
- Our generalised poset perspective dictated that the operations in any ordered algebras be treated as part of the expansion of the primitive (empty) signature. We do not regard the lattice operations (e.g., $\wedge$, $\vee$, $\neg$) as privileged, but as if they were a kind of 'modal operation' (and, in fact, operators in case the lattice is distributive). This has the advantage of a neat separation of roles between the consequence relation (encoded in the order) and the algebraic behaviour of the logic (expressed by the operations).

An obvious question that remains open is how to complete the square diagram for posets: is there a topological duality in the lower half of Diagram 1 which generalises the topological Stone and Priestley dualities satisfying our 'telescopic' requirement?
As mentioned in the introduction to Section 3 of Chapter 1, there are many proposals for such a duality in the case of lattices: Urquhart, Hartung and Haim ([43], [27], [25]), Hartonas and Dunn ([26]), and Jipsen and Moshier ([28], [29]). All of these dualities cross the boundary from distributive to non-distributive lattices, which turns out to be a significantly harder step than the step from Boolean algebras to distributive lattices.
In particular, all of the dualities for general bounded lattices that we mentioned above are two-sorted instead of one-sorted. Moreover, all of these dualities, that were designed for lattices, turn out to have features that naturally apply to the category of posets, in a similar way as the Boolean and distributive lattice settings are very much alike. Although limited, Urquhart/Hartung dualities, contrary to the dualities of Hartonas and Dunn, and of Jipsen and Moshier, are structurally compatible with Stone duality. We therefore try to broaden the scope of Urquhart/Hartung dualities to a duality which encompasses not only the surjective, but all morphisms.
In accordance with [16], we take duals of morphisms to be pairs of relations, rather than functions, and we view homomorphisms as two-faced objects (as explained in the second item in the list above). However, differently from [16], in order to satisfy the telescopic requirement, we want the relations to go in the same direction, i.e. 'horizontal' rather than 'cross-wise', as the functions in the dual categories of Stone and Priestley dualities.
Earlier on, we remarked that the similarity between lattices and posets is analogous to the similarity between Boolean and distributive lattices. Having said that, there are differences when moving from lattices to posets: these differences essentially don't show up at the level of the discrete duality (the upper half of Diagram 1), but do show up at the level of the topological duality (the lower half of Diagram 1). In this respect, we sketch a crucial example, omitting the proofs.[11]

[11] We thank Mai Gehrke for drawing our attention to this example.

Example 1. For any poset $\mathbb{P}$, let $\mathcal{F}(\mathbb{P}) = \{p\!\uparrow :\ p \in P\}$ and $\mathcal{I}(\mathbb{P}) = \{p\!\downarrow :\ p \in P\}$ be the filter and ideal systems consisting of just the principal up- and down-sets. Consider the 'butterfly' poset $\mathbb{P}$ and its Dedekind-MacNeille completion $\mathbb{Q}$:

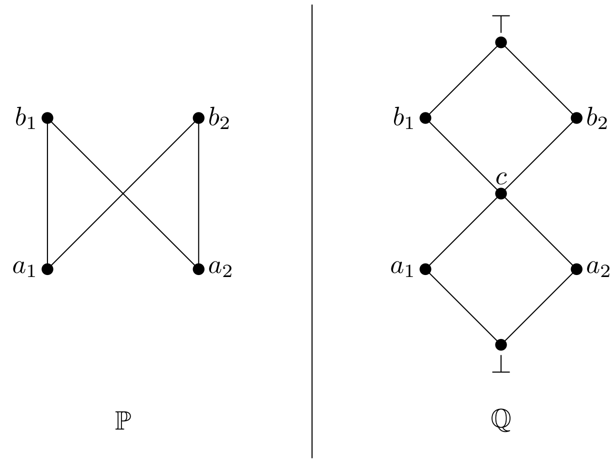

Then the topologies on the optimal filter sets $\mathcal{F}_*(\mathbb{P})$ and $\mathcal{F}_*(\mathbb{Q})$ which are generated by the subbases $\{\eta_{\mathbb{P}}(p) : p \in \mathbb{P}\}$ and $\{\eta_{\mathbb{Q}}(q) : q \in \mathbb{Q}\}$ are isomorphic, although the subbases themselves are different.

The last observation in this example could be the key to the problem of finding a duality in the lower half of Diagram 1, drawing once more from the modal logic methodology. Indeed, we expect that the most fruitful direction to an answer would be to try to generalise the *descriptive general frames* from classical modal logic. In that setting, the adjective *general* means that an ordinary Kripke frame is expanded with a collection of 'admissible' subsets of the set of worlds. The frame being *descriptive* means, in the classical case, that the collection of admissible subsets generates a Stone topology as a base (*Cf.* the end of Section 1 of the Introduction, and [5] for details).
Analogously, we expect that, in the case of posets, the category in the lower right corner of Diagram 1 would be the category of 'descriptive general RS polarities'. A descriptive general RS polarity would be defined as an RS polarity, expanded with two collections of admissible subsets for the two sorts, satisfying certain additional properties which make it 'descriptive'. However, when defining the notion of 'descriptive' in this context, we need to take care. In particular, because of Example 1, we can not have a definition which would yield an equivalence between the category of descriptive general RS polarities and a category of topologized RS polarities. Note that we still have such an equivalence in the lattice case.
Therefore, the following interesting conclusion can be drawn from these considerations. We saw in Example 1 that a topological approach to a duality for posets does not work. Posets form the first layer in the hierarchy where descriptive general frames are essential, showing that they are really more than a mathematically unorthodox reformulation of the topological approach, for which they are often criticized. Just as we 'lose' operations when moving from lattices to posets, we 'lose' the possibility of using topology when moving from RS polarities for lattices to RS polarities for posets. Thus, the considerations in this section suggest in what directon to look for a duality in the lower half of Diagram 1.

2 Category-theoretic perspective on canonical extensions

Another question that arose in Chapter 1 was: can we give a categorical characterisation of the canonical extension construction? In particular, by appropriately restricting the notion of morphism between posets, is it possible to find a subcategory of Poset for which both the Dedekind-MacNeille completion is a reflector, and the amalgamation is a pushout? Finding, and subsequently studying the properties of this subcategory could pave the way to generalisations of the canonical extension construction to other categories of possibly non-ordered structures, such as groups, groupoids, topological spaces, and so forth.

A related, but different, point is that we can regard any partially ordered set itself as a category in a straightforward way. This observation, as well as the fact that all of the additional operations in lattices and Boolean algebras have natural categorical interpretations (e.g., meets and joins are *limits* and *colimits* in category-theoretic terms, and top and bottom elements are *initial* and *terminal* objects), leads one to suspect that there may be a construction of 'the canonical extension of a category'. We must admit that this idea should be motivated and clarified further, but here, we just wanted to make note of this possibly interesting direction.

Moving on to another point, the discussion of morphisms in Section 4 of Chapter 1 would benefit from a more structural study of the categorical properties of the category of (general) RS polarities. Here, we are thinking of characterizing important constructions such as products, coproducts, subobjects and limits and colimits in general. As for RS polarities, we can use the discrete duality that we discussed in Chapter 1. Still, we would like to find independent characterizations of these constructions in the category of RS polarities. For general RS polarities, a uniform way to treat these questions is to study the adjoint behaviour of the forgetful functor from general RS polarities to RS polarities. This would be a two-sorted version of the well-known category-theoretic methodology, where limits and colimits in a more complicated category (e.g., Top) can be uniformly lifted from a simpler, better-known category (e.g., Set) [1].

3 Canonicity and correspondence

In Chapter 2, we focused on one particular application of the canonical extension: we revisited the algebraic canonicity proof of the Sahlqvist fragment for distributive modal logic [22], in the quest of an analogous algebraic canonicity proof of the inductive fragment. We isolated three important steps in the proof, and described which were essentially order-theoretic, and which pertained particularly to Sahlqvist, paying attention to how these different steps interact. In particular, we focused on the properties of the additional operation $\mathbf{n}$, which played a crucial role in the first part of the algebraic canonicity proof.

Building on this analysis, we proved that the Sahlqvist proof can not be straightforwardly translated to a proof for the inductive fragment. The first open problem is therefore to find an algebraic proof of the canonicity of the inductive fragment. We believe that finding an enhanced version of the operation $\mathbf{n}$ would be an important step towards such a result. There are some open questions which are related to the operation $\mathbf{n}$ itself, so that we believe investigating the operation $\mathbf{n}$ in its own right could be benificial to finding 'the enhanced $\mathbf{n}$'.

We have two perspectives on $\mathbf{n}$ that we would like to investigate further. Firstly, we observed that the operation $\mathbf{n}$ has many similarities with the Hom functor from a lattice, regarded as a category, to Set. Moreover, we discovered some interesting parallels between (what we called) the '$\mathbf{n}$-trick' and the Ackermann Lemma, which is central to the correspondence for the inductive fragment. The second perspective raises another question, pertaining to the interaction between canonicity and correspondence.

One of the widely recognized features of the algebraic canonicity approach is its providing proofs of canonicity independent of correspondence. It was already observed by Fine [12] that in classical modal logic, there are formulas which are canonical but do not have a first-order correspondent. From this result, we know that having a first-order correspondent is a strictly stronger property than canonicity. Jónsson ([30], section 6) proved the algebraic canonicity of certain isolated formulas (such as Fine's formula) which do not have any first-order frame correspondent, showing that algebraic canonicity proofs really are more than an algebraic reformulation of the canonicity-via-correspondence method.

On the other hand, the observed parallels between the Ackermann Lemma and the $\mathbf{n}$-trick suggest that the operation $\mathbf{n}$ may partly bring a 'correspondence-type' ingredient into the algebraic canonicity proof. Could the correspondence that is used for the proof of canonicity which is explicit in proofs on the frame side, be actually present, but hidden, in algebraic canonicity proofs? In what sense precisely does the operation $\mathbf{n}$ provide a substitute, or an enhancement, of correspondence methods? We do not give precise formulations of, nor definite answers to, these questions, but think they suggest another interesting direction for further investigations.

Bibliography

[1] Jirí Adámek, Horst Herrlich, and George E. Strecker, *Abstract and Concrete Categories: The Joy of Cats*, published under GNU Free Documentation License at http://katmat.math.uni-bremen.de/acc/, 2004.

[2] Bernhard Banaschewski and Günter Bruns, *Categorical Characterization of the MacNeille Completion*, Archiv der Mathematik **18** (1967), 369–377.

[3] Johan van Benthem, *Information Transfer across Chu Spaces*, Logic Journal of the Interest Group in Pure and Applied Logic **8** (2000), no. 6, 719–731.

[4] Alan A. Bishop, *A universal mapping characterization of the completion by cuts*, Algebra Universalis **8** (1978), 349–353.

[5] Patrick Blackburn, Maarten de Rijke, and Yde Venema, *Modal logic*, Cambridge Tracts in Theoretical Computer Science, vol. 53, Cambridge University Press, 2001.

[6] Arthur Cayley, *On the theory of groups as depending on the symbolic equation* $\theta^n = 1$, Philosophical Magazine **VII** (1854).

[7] Willem Conradie and Alessandra Palmigiano, *Algorithmic Correspondence and Canonicity for Distributive Modal Logic*, to appear (2009).

[8] Roy L. Crole, *Categories for Types*, Cambridge Mathematical Textbooks, Cambridge University Press, 1993.

[9] J. Michael Dunn, Mai Gehrke, and Alessandra Palmigiano, *Canonical extensions of ordered algebraic structures and relational completeness of some substructural logics*, Journal of Symbolic Logic **70** (2005), no. 3, 713–740.

[10] J. Michael Dunn and Gary Hardegree, *Algebraic Methods in Philosophical Logic*, Oxford Logic Guides, vol. 41, Oxford University Press, 2001.

[11] Marcel Erné, *The Dedekind-MacNeille Completion as a Reflector*, Order **8** (1991), 159–173.

[12] Kit Fine, *Some connections between elementary and modal logic*, Proceedings of the 3rd Scandinavian Logic Symposium, Uppsala 1973 (Amsterdam) (S. Kanger, ed.), Springer-Verlag, 1975, pp. 110–143.

[13] Lisa Fulford, *A study of Canonicity for Bi-Implicative Algebras*, (2009), Master of Logic Thesis, to appear.

[14] Nikolaos Galatos, Peter Jipsen, Tomasz Kowalski, and Hiroakira Ono, *Residuated Lattices: an algebraic glimpse at substructural logics*, Studies in Logic and the Foundations of Mathematics, vol. 151, Elsevier, 2007.

[15] Bernhard Ganter, Gerd Stumme, and Rudolf Wille (eds.), *Formal concept analysis, foundations and applications*, Lecture Notes in Computer Science, vol. 3626, Springer, 2005.

[16] Mai Gehrke, *Generalized Kripke Frames*, Studia Logica **84** (2006), no. 2, 241–275.

[17] Mai Gehrke and John Harding, *Bounded lattice expansions*, Journal of Algebra **238** (2001), no. 1, 345–371.

[18] Mai Gehrke, Ramon Jansana, and Alessandra Palmigiano, *Some δ_1-completions of a poset and the canonical extension*, (2008), Preprint.

[19] Mai Gehrke and Bjarni Jónsson, *Bounded distributive lattices with operators*, Mathematica Japonica **40** (1994), no. 2, 207–215.

[20] ———, *Monotone bounded distributive lattice expansions*, Mathematica Japonica **52** (2000), no. 2, 197–213.

[21] ———, *Bounded Distributive Lattice Expansions*, Mathematica Scandinavica **94** (2004), 13–45.

[22] Mai Gehrke, Hideo Nagahashi, and Yde Venema, *A Sahlqvist Theorem for Distributive Modal Logic*, Annals of Pure and Applied Logic **131** (2005), no. 1-3, 65–102.

[23] Mai Gehrke and Hilary A. Priestley, *Canonical extensions and completions of posets and lattices*, Reports on Mathematical Logic **43** (2008), 133–152.

[24] Silvio Ghilardi and Giancarlo Meloni, *Constructive Canonicity in Non-Classical Logics*, Annals of Pure and Applied Logic **86** (1997), no. 1, 1–32.

[25] Marianna Haim, *Duality for Lattices with Operators: A Modal Logic Approach*, (2000), Master of Logic Thesis, MoL 2000-02.

[26] C. Hartonas and J. M. Dunn, *Stone duality for lattices*, Algebra Universalis **37** (1997), 391–401.

[27] Gerd Hartung, *A topological representation of lattices*, Algebra Universalis **29** (1992), 273–299.

[28] Peter Jipsen and Andrew Moshier, *Topological Duality and Lattice Expansions Part I: A Topological Construction of Canonical Extensions*, (2009), preprint.

[29] ———, *Topological Duality and Lattice Expansions Part II: Lattice Expansions with Quasi-operators*, (2009), preprint.

[30] Bjarni Jónsson, *On the canonicity of Sahlqvist identities*, Studia Logica **4** (1994), 473–491.

[31] Bjarni Jónsson and Alfred Tarski, *Boolean Algebras with Operators*, American Journal of Mathematics **74** (1952), no. 1, 127–162.

[32] Saunders MacLane, *Categories for the working mathematician, Second edition*, Graduate Texts in Mathematics, vol. 5, Springer, 1998.

[33] ______, *The Yoneda Lemma*, Mathematica Japonica **47** (1998), no. 1, 156.

[34] Alessandra Palmigiano, *Notes on Canonical Extensions of Partial Orders*, (2004), unpublished.

[35] ______, *Notes on Sahlqvist Correspondence*, (2008), Notes accompanying the MSc Logic course 'Introduction to Modal Logic', ILLC, Universiteit van Amsterdam.

[36] Adrian Pigors, *PhD Thesis*, (2010), forthcoming.

[37] Vaughan Pratt, *Chu Spaces*, Course notes for the School in Category Theory and Applications (1999).

[38] Hilary Priestley, *Representation of distributive lattices by means of ordered Stone spaces*, Bulletin of the London Mathematical Society **2** (1970), 186–190.

[39] Frigyes Riesz, *Sur une espèce de géométrie analytiques des systèmes de fonctions sommables*, Comptes rendus hebdomadaires des séances de l'Académie des Sciences **144** (1907), 1414–1416.

[40] ______, *Sur les opérations fonctionelles linéaires*, Comptes rendus hebdomadaires des séances de l'Académie des Sciences **149** (1909), 974–977.

[41] M. H. Stone, *The Theory of Representation for Boolean Algebras*, Transactions of the American Mathematical Society **74** (1936), no. 1, 37–111.

[42] Tomoyuki Suzuki, *Canonicity Results of Substructural and Lattice-based Logics*, The Review of Symbolic Logic (2009), to appear.

[43] Alasdair Urquhart, *A topological representation theory for lattices*, Algebra Universalis **8** (1978), no. 1, 45–58.

[44] Yde Venema, *Algebras and Coalgebras*, Handbook of Modal Logic (Patrick Blackburn, Johan van Benthem, and Frank Wolter, eds.), Elsevier, 2007, pp. 331–426.

[45] Willem Conradie and Valentin Goranko and Dimiter Vakarelov, *Elementary canonical formulae: a survey on syntactic, algorithmic, and model-theoretic aspects*, Advances in Modal Logic (R. Schmidt and I. Pratt-Hartmann and M. Reynolds and H. Wansing., ed.), vol. 5, 2005, pp. 17–51.

[46] ______, *Algorithmic Correspondence and Completeness in Modal Logic: I. The Core Algorithm SQEMA*, Logical Methods in Computer Science **2** (2006), 1–26.