

Imperial College of Science, Technology and Medicine
Department of Computing

Completeness-via-canonicity for coalgebraic logics

Fredrik Paul Herbert Dahlqvist

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Computing of the University of London and
the Diploma of Imperial College, April 2014

Abstract

This thesis aims to provide a suite of techniques to generate completeness results for coalgebraic logics with axioms of arbitrary rank. We have chosen to investigate the possibility to generalize what is arguably one of the most successful methods to prove completeness results in ‘classical’ modal logic, namely *completeness-via-canonicity*. This technique is particularly well-suited to a coalgebraic generalization because of its clean and abstract algebraic formalism. In the case of classical modal logic, it can be summarized in two steps, first it isolates the purely algebraic problem of canonicity, i.e. of determining when a variety of boolean Algebras with Operators (BAOs) is closed under canonical extension (i.e. canonical). Secondly, it connects the notion of canonical varieties to that of canonical models to explicitly build models, thereby proving completeness.

The classical algebraic theory of canonicity is geared towards *normal* logics, or, in algebraic terms, BAOs (or generalizations thereof). Most coalgebraic logics are not normal, and we thus develop the algebraic theory of canonicity for Boolean Algebra with Expansions (BAEs), or more generally for Distributive Lattice Expansions (DLEs). We present new results about a class of expansions defined by weaker preservation properties than meet or join preservation, namely (anti)- k -additive and (anti)- k -multiplicative expansions. We show how canonical and Sahlqvist equations can be built from such operations.

In order to connect the theory of canonicity in DLEs and BAEs to coalgebraic logic, we choose to work in the abstract formulation of coalgebraic logic. An abstract coalgebraic logic is defined by a functor $L : \mathbf{BA} \rightarrow \mathbf{BA}$, and we can heuristically separate these logics in two classes. In the first class the functor L is relatively simple, and in particular can be interpreted as defining a BAE. This class includes the predicate lifting style of coalgebraic logics. In the second class the functor L can be very complicated and the whole theory requires a different approach. This class includes the nabla style of coalgebraic logics.

For simple functors, we develop results on strong completeness and then

prove strong completeness-via-canonicity in the presence of canonical frame conditions for strongly complete abstract coalgebraic logics. In particular we show coalgebraic completeness-via-canonicity for Graded Modal Logic, Intuitionistic Logic, the distributive full Lambek calculus, and the logic of trees of arbitrary branching degrees defined by the List functor. These results are to the best of our knowledge, new.

For a complex functor L we use an indirect approach via the notion of *functor presentation*. This allows us to represent L as the quotient of a much simpler polynomial functor. Polynomial functors define BAEs and can thus be treated as objects in the first class of functors, in particular we can apply all the above mentioned techniques to the logics defined by such functors. We develop techniques that ensure that results obtained for the simple *presenting* logic can be transferred back to the complicated *presented* logic. We can then prove strong-completeness-via-canonicity in the presence of canonical frame conditions for coalgebraic logics which do not define a BAE, such as the nabla coalgebraic logics.

Acknowledgements

This thesis is the fruit of a long discussion (four years!) with my supervisor Dirk Pattinson, to whom I am immensely grateful for a number of things: for taking me on as his student, for passing on a great deal of technical knowledge and teaching me how to write a paper, for taking the time to see me every week (at least when we were not at opposite ends of the world), for his amazing ability to think, talk and solve problems for hours on end without need for food or drink, and for always being up for a pint! I would also like to thank Ian Hodgkinson for stepping in after Dirk's move down under and his invaluable help when I was trying to understand canonicity. I am also very grateful to Alexander Kurz and Corina Cîrstea for accepting to examine my thesis, and for doing it so conscientiously.

My thanks of course to Tanja, firstly for always encouraging me, despite the massive pay cut, to pursue what I really wanted to do; secondly for bringing Ingrid to the world, by far the most wonderful thing to have happened during my PhD years; and finally for her patience and sacrifices during the long writing up period. I also thank my family, and my parents in particular, for their constant support, understanding and encouragements, in the face of what might have appeared as a strange career move.

I must also thank all the people who made going to work such a fun and interesting experience, in no particular order Björn, Ray, Florian, Victor, Tony, Rumi, Clemens, Nick, Faris and Rob.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Fredrik Dahlqvist)

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work

Contents

1	Prolegomenon	1
1.1	What this works sets out to achieve	1
1.1.1	Introductory remarks	1
1.1.2	Structure, summary and main contributions of the thesis.	4
1.2	Notation and Conventions	11
1.3	Reasoning structures	12
1.3.1	The category BA	12
1.3.2	The category DL	15
1.3.3	Boolean algebra and distributive lattice expansions	16
1.3.4	Algebras in BA and DL	20
1.4	Relations, categorically	21
1.4.1	Relations	21
1.4.2	Relation liftings	22
1.4.3	Equivalence relations and exactness	27
1.4.4	Fully invariant equivalence relations	34
1.5	Coalgebraic logics	37
1.5.1	Three styles of language	37
1.5.2	Three styles of semantics	41
1.5.3	Two styles of axiomatization	45
1.5.4	One style of algebraic semantics	53
2	Algebraic canonicity	67
2.1	Canonical Extensions	68
2.1.1	Canonical extension of distributive lattices and boolean algebras	68
2.1.2	Canonical extension of expansions	76
2.2	Algebraic Properties of the Canonical Extension	81
2.2.1	Properties preserved and strengthened by canonical extension	81
2.2.2	Canonical extension and L -algebras.	89
2.2.3	Canonical extension and function composition	90

2.3	Topology to the rescue	92
2.4	Canonicity	103
2.5	Sahlqvist identities	110
2.6	Applications	116
2.6.1	Graded Modal Logic	116
2.6.2	Intuitionistic Logic and Distributive Substructural Logics	119
3	Functor Presentations	123
3.1	The base transformation.	124
3.2	The concept of functor presentation	135
3.2.1	The category $\mathbf{El}(T)$	136
3.2.2	Colimits of representables	139
3.3	Some essential tools and concepts	141
3.3.1	Locally presentable and accessible categories	141
3.3.2	Kan extensions and cofinal diagrams	144
3.4	Size-bound functor presentations	148
3.4.1	λ -presentable and λ -presented functors	149
3.4.2	λ -generatable and λ -generated functors	160
3.4.3	Free cocompletions	167
3.5	Presentations of accessible functors	169
3.5.1	General presentations	169
3.5.2	Size-bound presentations	171
3.5.3	The canonical presentation.	175
3.5.4	Atomic objects of $\mathbf{El}(T)$	177
3.5.5	The minimal presentation	179
3.6	Lifting presentations using adjunctions	182
4	Translations	187
4.1	Syntax translations	187
4.2	Models and semantics translations	192
4.2.1	Semantic translations	192
4.2.2	Compatible semantics	194
4.2.3	The case of nabla logics	198
4.3	Proof-theoretic translations	205
5	Completeness-via-canonicity	211
5.1	Weak completeness	211
5.1.1	Completeness theorem	212
5.1.2	Completeness by filtration	218
5.1.3	A brief note on expressivity	219
5.2	Strong completeness	220
5.2.1	The coalgebraic Jónsson-Tarski theorem	220

5.2.2	Strong completeness for positive coalgebraic logics	223
5.2.3	Strong completeness for boolean coalgebraic logics	233
5.2.4	Strong completeness by semantic completion	250
5.3	Completeness-via-canonicity	266
5.3.1	Jónsson-Tarski vs canonical extensions	266
5.3.2	L -algebras defining DLEs with smooth extensions	269
5.3.3	General L -algebras	279

Chapter 1

Prolegomenon

1.1 What this works sets out to achieve

1.1.1 Introductory remarks

Coalgebras have gained popularity as an elegant and general framework to study and represent a wide variety of dynamical systems in computer science (see [Rut00]) and even in physics (see [Abr09]). The fact that they encompass so many of the most frequently used models of computation suggests that they provide the correct level of abstraction to formalize the dynamics of state based systems.

In parallel to the research on coalgebras *per se*, the field of coalgebraic logic has emerged as a unifying framework for the many types of (modal) logics used in knowledge representation and to reason about dynamical systems (see [KP11] for an overview). One of the great insights in the relationship between coalgebras and coalgebraic logics, is that the class of all T -coalgebras for a functor T can always be characterised logically by its one-step behaviour, i.e. axioms and rules with nesting depth of modal operators uniformly equal to 1 (see [Sch06]). However, once the transition type (i.e. the functor T) has been described logically in such a way, one may be interested in subclasses of T -coalgebras which are characterised by more complex axioms (such as transitivity for example) which we will refer to as *frame conditions*. The problem of logically characterising subclasses of the class of all T -coalgebras for an arbitrary functor T is by and large still open. To our knowledge, Pattinson and Schröder's [PS08] offer the only attempt at making progress in this direction and offer a solution for some of the standard frame conditions of classical modal logic.

This thesis aims to provide a suite of techniques to generate completeness results for coalgebraic logics with axioms of arbitrary rank. As practitioners of modal logic know all too well, there is no panacea for proving completeness in

the presence of frame conditions, and the generality of coalgebraic logic only makes this observation more true. Thus what we ambition to provide is not a general theory of completeness in the presence of frame conditions, but rather *some* techniques to deal with *some* frame conditions in *some* coalgebraic logics. In this spirit, we have chosen to investigate the possibility to generalize what is arguably one of the most successful methods to prove completeness results in ‘classical’ modal logic (i.e. the normal modal logic of e.g. [BdRV01]), namely *completeness-via-canonicity*. This technique is particularly well-suited to a coalgebraic generalization because of its clean and abstract algebraic formalism. In the case of classical modal logic, it can be summarized in two steps, first it isolates the purely algebraic problem of determining when a variety of boolean Algebras with Operators (BAOs) is canonical, i.e. closed under canonical extension. Secondly, it connects the notion of canonical varieties to that of canonical models, i.e. the algebraic and syntactic world to the coalgebraic and semantic world, and explicitly builds models proving completeness.

There are two challenges in trying to generalize classical completeness-via-canonicity to the coalgebraic framework. Firstly, the classical algebraic theory of canonicity (see for example in chronological order [JT51, Rib52, Sah75, Gol89, Jón94, Gol95, dRV95, GH01, GJ04, GNV05, Ven06]) is geared towards *normal* logics, or, in algebraic terms, BAOs (or generalizations thereof). Most coalgebraic logics are not normal, and we thus need to develop the algebraic theory of canonicity for boolean Algebra with Expansions (BAEs), i.e. boolean algebras with additional operators that have *a priori* no preservation properties. Our key findings on this topic are presented in Chapter 2 and can be summarized as follows:

- Some results for non-normal expansions can already be found in [GH01, GJ04, GNV05, Ven06], and we generalize them further to isotone and antitone expansions of arbitrary arities
- We study the canonical extension of expansions which are (anti-) k -additive (see [Hen70]) or (anti-) k -multiplicative. To the best of our knowledge, our results on this class of expansions are new. A very interesting application is Graded Modal Logic (GML) whose modalities are k -additive (and whose dual modalities are k -multiplicative).
- We show how we can define canonical identities via a generalization of the notion of Sahlqvist identity to non-normal modal logics whose expansions are (anti-) k -additive or (anti-) k -multiplicative, or (anti)-preserve up-directed joins, or (anti)-preserve down-directed meets.

The second challenge is to relate this purely algebraic theory to the coalgebraic semantic of coalgebraic logic. As we shall detail later in this Chapter,

there are three flavours of coalgebraic logic: the predicate lifting, the nabla, and the abstract flavour. We choose to place ourselves in the abstract framework for most of this thesis because it is at this level of abstraction that results are easiest to get and cleanest to present. Moreover, the abstract flavour subsumes the other two. An abstract coalgebraic logic is essentially defined by a functor $L : \mathbf{BA} \rightarrow \mathbf{BA}$, and we can heuristically separate these logics in two classes. In the first class the functor L is relatively simple, and in particular can be interpreted as defining a BAE; this case can be thought of as the abstraction of the predicate lifting flavour of coalgebraic logic. In the second class the functor L can be very complicated and the whole theory requires a different approach; this case can be thought of as the abstraction of the nabla flavour of coalgebraic logic.

The main results concerning the first class of abstract coalgebraic logics are

- In Theorem 5.16 we show strong completeness for a very large class of positive coalgebraic logic, which we call *relational logics* and which includes positive modal logic, intuitionistic logic and all distributive substructural logics. This results generalises results published in [DP15].
- In the case of boolean coalgebraic logics, we prove a suite of Theorems 5.24-5.27 which guarantee the existence of (quasi)-canonical models, and thus strong completeness. The key requirement is that the functor defining the semantics must weakly preserve cofiltered limits.
- For functors which do not weakly preserve cofiltered limits, we present a new technique which we call *semantic completion*, which extends the semantics in a generic way in order for Theorems 5.24-5.27 to be used.
- We clarify the connection between the syntactic notion of *canonical extension* and the semantic notion of *canonical model* in Theorems 5.44 and 5.45 in a coalgebraic way. To our knowledge this is the first time this question is addressed, let alone answered.
- With the connection between canonical extensions and canonical models clarified, we present coalgebraic completeness-via-canonicity results for functors defining DLEs and BAEs in Theorems 5.47 and 5.49.

The second class of abstract coalgebraic logic requires an indirect approach via the notion of *functor presentation*. If the functor L presenting the logic is very complex, as is typically the case in the nabla flavour of coalgebraic logic, the main technical tool for dealing with this complexity is to represent L as the quotient of a much simpler polynomial functor. Polynomial functors define BAEs and can thus be treated as objects in first class of functors, in particular we can apply all the above mentioned techniques to the logics defined by such

functors. We develop techniques that ensure that results obtained for the simple *presenting* logic can be transferred back to the complicated *presented* logic. The main results of this line of research are:

- We study functor presentation in great detail and show how it is related to the classical category theoretic construction presenting a **Set**-functor as a colimit of representable functors. We show a very close relationship between the notion of functor presentation and the notion of presentation by generators and relations for an algebraic variety. In particular, we define notions of λ -presented and λ -generated functors, and show that they are equivalent to being λ -presentable or λ -generatable in a category of functor. The latter concepts being precisely those which abstract the usual notion of being λ -presented or λ -generated for an algebraic variety.
- We show in Chapter 4 how natural transformations (such as functor presentations) can be used to (1) translate between abstract coalgebraic logics (Theorem 4.1), (2) translate between their semantic domains (Theorem 4.5), (3) translate syntax and semantics simultaneously whilst preserving satisfiability (Theorem 4.7), and finally (4) translate between proof systems in a way that preserve derivability in a precise way (Theorem 4.13).
- By using presentations by polynomial functors, all our translation results, and our completeness-via-canonicity results on the BAEs defined by polynomial functors, we show a step-by-step technique for proving strong-completeness-via-canonicity in the presence of canonical frame conditions for coalgebraic logics which do not define BAEs, such as the nabla coalgebraic logic for the finitary powerset functor (Section 5.3.3).

1.1.2 Structure, summary and main contributions of the thesis.

- **Chapter 1: Prolegomenon**
 - **Section 1.3: boolean structures**, in which all the basic constructs on boolean algebras are introduced concretely, as well as in the appropriate categorical setting. The question of when the category $\mathbf{Alg}_{\mathbf{BA}}(L)$ (for a variator $L : \mathbf{BA} \rightarrow \mathbf{BA}$) is monadic over **Set** is raised and the answer, which depends on results from the next Section, is presented.
 - **Section 1.4: Relations, categorically**: in which the notions of relation, relation lifting, equivalence relation and fully invariant relation are defined and developed in the context of regular categories. Although the results of Propositions 1.11 and 1.12 are known, the

direct proofs we present are, to our knowledge, new and provide a simple and yet purely categorical handle on the key properties of relation liftings which are used extensively in the nabla and abstract flavours of coalgebraic logic. The study of the (Barr-) exactness of categories of algebras is a relatively straightforward exercise on regular categories, but provides a nice new framework for the algebraic semantics of abstract coalgebraic logics developed in Section 1.5. Many ideas originate from the thesis of Hughes [Hug01] on categories of algebras and coalgebras. The notion of fully invariant relation was defined in [Hug01], but our construction of the fully invariant closure of a relation (Proposition 1.23) is, to the best of our knowledge, new.

- **Section 1.5: Coalgebraic logics**, in which coalgebraic languages, their semantics and their axiomatizations are defined. We pay particular attention to working out the details of the abstract flavour of coalgebraic logic and its connections with the other two flavours. We present an algebraic semantic for abstract coalgebraic logics in great detail and at what we believe to be the correct level of abstraction, i.e. working over the category **BA**. Most of this material is new, and generalizes/categorifies earlier work on the topic such as [KKP04]), relying heavily on the notions developed in the previous section, notably the notion of fully invariant closure.

• Chapter 2: Algebraic Canonicity

- **Section 2.1: Canonical Extensions**, in which basic concepts such as canonical extensions, compactness, closed and open elements, etc are introduced. Whilst all the results shown in this section are known, we have opted for a purely algebraic presentation, i.e. we make no reference to duality (Priestley or Stone spaces) as is usually the case in the literature. The basic properties of the canonical extensions of maps on distributive lattices or boolean algebras, developed in e.g. [GJ94, GH01, GJ04, Ven06], are generalised to n -ary monotone maps, i.e. maps f such that for each argument, f is either order preserving (isotone) or order reversing (antitone) in this argument.
- **Section 2.2: Algebraic properties of the canonical extension**, in which known results from [GJ94, GH01, GJ04, Ven06] are extended to the case of maps which are (anti-) k -additive or (anti-) k -multiplicative in Theorem 2.18. This Theorem generalises known results about join and meet-preserving maps. Lemma 2.23 which shows how canonical extension and function composition interact is of paramount importance for the rest of the Chapter.

- **Section 2.3: Topology to the rescue**, in which topological techniques developed in [GJ04] and [Ven06] are presented and extended to the case of n -ary maps which are isotone or antitone in each argument (Theorem 2.28). We also show that (anti-) k -additive and (anti-) k -multiplicative maps are smooth (Theorem 2.35); to the best of our knowledge this result is new. These topological results, and in particular the *principle of matching topologies* of [Ven06], will form the backbone of our study of canonicity. They provide us with conditions under which the converse of the inequality of Lemma 2.23 holds, i.e. conditions under which canonical extension and function composition commute. This property is key to building canonical terms.
- **Section 2.4: Canonicity**, in which the notions of canonical equation, stable and expanding terms are defined. Our presentation broadly follows that of [Jón94], defining terms as maps and studying their stability under canonical extension. The key result is the Principle of Matching Topology, an idea of [Ven06], which we use to generate a large number of conditions under which terms are stable under canonical extension. The results are presented in Table 2.1.
- **Section 2.5: Sahlqvist identities**, in which we attempt to clarify and formalise the notion of quasi-equation which is the key technical tool used in [Jón94] to define Sahlqvist identities. We suggest an abstract definition of Sahlqvist identities for a general BAE, based on algebraic properties of terms. We then present concrete - i.e. syntactically defined - Sahlqvist identities; first the classical Sahlqvist identities of modal logic, then a more general definition of Sahlqvist identities for BAEs whose operations satisfy one of the preservation properties listed in Table 2.1. We use Tables 2.1 and 2.2 as the basis of a method for building these general Sahlqvist identities. We illustrate the method by examples, including an example of a general Sahlqvist formula for the classical modal logic, which is not a classical Sahlqvist formula.
- **Section 2.6: Applications**. In this section we present interesting examples of DLEs and BAEs and define notions of Sahlqvist identities for them. We start with the case of Graded Modal Logic (GML) which is characterized by BAEs whose expansions are k -additive. Canonical equations and Sahlqvist identities can therefore be defined by using the results on k -additive maps developed in the Chapter. Next, we present Intuitionistic Logic (IL) and the distributive Lambek calculus as examples of positive modal logics with binary expansions whose arguments are all either isotone or antitone. The

logics provide examples of expansions which preserve joins or preserve meets or anti-preserve joins in their arguments. We then use results developed in the Chapter to define Sahlqvist identities in these logics.

- **Chapter 3: Functor Presentations**

- **Section 3.1: The base transformation**, where the base transformation, which is key to the nabla flavour of coalgebraic logic, is defined and discussed in a very general categorical setting, viz. for any **Set**-valued functor on a locally small, well-powered category with pullbacks of monos (i.e. intersections). Theorem 3.12 is a generalisation of results from [Gum05], and characterises functors for which the base transformation is natural.
- **Section 3.2: The Concept of Functor Presentation**, in which the most basic concepts of the theory of functor presentation are introduced: the category of elements of a functor and the associated expression of a functor as a colimit of representables (Theorem 3.19). The analogy between functor presentations and algebraic varieties, which will be the unifying thread of the Chapter, is first hinted at.
- **Section 3.3: Some Essential Tools and Concepts**, in which the concepts of locally presentable and accessible categories, left Kan extensions, canonical and cofinal diagrams are introduced. Proposition 3.36 shows how how a λ -accessible functor is determined uniquely, via the left Kan extension, by its restriction to λ -presentable objects.
- **Section 3.4: Size-Bound Functor Presentations**. In this large section we develop the analogy between **Set**-valued functors on a small category and elements of an algebraic variety. We define two notions of a functor being λ -presented and show that they are equivalent in Proposition 3.44. The key result of this section is Proposition 3.46 which shows that a functor is finitely presented iff it is a finitely presentable object in **Set**^{**C**}. We then define two notions of a functor being λ -generated and determine a condition on **C** that ensures that they agree. The notion of a ‘generatable’ object is our terminology for the notion of generated of [AR94]. Proposition 3.59 is the equivalent for generated functors of Proposition 3.46, i.e. it shows that a functor is finitely generated iff it is a finitely generatable object in **Set**^{**C**}. A summary of the analogy between functors in **Set**^{**C**} and elements of an algebraic variety is presented in Table 3.1. The section ends with a re-interpretation of some of its results in terms of various notions of cocompletion.

- **Section 3.5: Presentations of accessible functors**, where the theory developed for **Set**-valued functors over a small category is extended to the case of accessible **Set**-valued functors on accessible categories. We show how the left Kan extension construction and the accessibility assumption allows us to lift all result from the previous section to accessible categories. In particular we define and characterise presented/presentable and generated/generatable accessible functor and show that the category of λ -accessible functors is locally finitely presentable (Proposition 3.71). We also present a class of accessible categories over which finitely presentable and finitely generatable functors coincide (Proposition 3.73). We then introduce the so-called canonical presentation, and show a generic method for removing the redundancy of this presentation via the notion of atomic elements in the category of elements, itself related to the concept of base. We call the presentation obtained in this way the minimal presentation (already present in some form in [AT90]).
- **Section 3.6: Lifting presentations using adjunctions**, in which we show how we can extend the result of the previous section from the case of functors $\mathbf{C} \rightarrow \mathbf{Set}$ to the case of functors $\mathbf{C} \rightarrow \mathbf{C}$ when there exists an adjunction $F \dashv U : \mathbf{Set} \rightarrow \mathbf{C}$. We show how any finitary functor $T : \mathbf{C} \rightarrow \mathbf{C}$ can be given such a presentation if (1) the adjunction is finitary, (2) the adjunction is of descent type, (3) T preserves regular epis, and (4) the composition of regular epis is a regular epi in $\mathbf{C}$. All these conditions are met by our main application, namely endofunctors on **BA**.

• **Chapter 4: Translations**

- **Section 4.1 Syntax translation.** In this section we examine how a natural transformations between functors presenting two coalgebraic languages defines a translation map between these languages. We show how this translation, which is in fact a natural transformation, is defined, that it exists, and that it is a regular epi-transformation whenever the natural transformation between the functors is in Theorem 4.1
- **Section 4.2 Models and semantic translations**, in which we show how a natural transformation between two functors defining different coalgebraic semantics defines translation between the two types of models. We show in Theorem 4.5 how the cofree coalgebra for the two functors are related by this translation map, and in particular that if the natural transformation is epi, so is the translation map. Next, we show how the syntax and semantic translations can

be combined in a compatible way and prove that the truth predicate is preserved by such compatible translations. We then show how to implement the notion of compatible semantics in the case of the nabla logics associated with functors $S, T : \mathbf{Set} \rightarrow \mathbf{Set}$ for which there exists an epi-transformation $q : S \twoheadrightarrow T$. This involves presenting nabla logic in the abstract style and some detailed examination of the connection between derivability under the proof systems $\mathbf{KKV}(S)$ and $\mathbf{KKV}(T)$ associated with S and T .

- **Section 4.3: Proof-theoretic translation** in which proof systems are translated via the syntax translation defined in Section 4.1. Given two functors K, L presenting two abstract coalgebraic logics, an epi-transformation $q : K \twoheadrightarrow L$ and a set of equations for the L -logic, we show how this set of equations can be lifted to the K language. In this way we show how the abstract Hilbert systems defined in Section 1.5.4 can be lifted along an epi-transformation. The key result is Theorem 4.13 which shows that there is a very tight relationship between the variety defined by the L -equations and the variety defined by the lifted K -equations.

• Chapter 5

- **Section 5.1: Weak completeness.** We start this Chapter by introducing classical results about weak completeness in the abstract flavour of coalgebraic logic, and in particular the connection with the injective nature of the semantic natural transformation. For an abstract coalgebraic logic given by a triple $(L : \mathbf{BA} \rightarrow \mathbf{BA}, T : \mathbf{Set} \rightarrow \mathbf{Set}, \delta : L\mathcal{P} \rightarrow \mathcal{P}T)$, the precise connection between the free L -algebra over the boolean algebra $\mathbf{FV}$ and the cofree coalgebra over the set (of colours) $\mathbf{QV}$ is established in Theorem 5.25. This result is probably folklore, but left us perplexed for a long time until the connection between the adjunctions $\mathbf{Uf} \dashv \mathcal{P}$ and $\mathbf{F} \dashv \mathbf{U}$ became clear. This greatly clarifies the role of propositional variables and valuations. We also review the notion of filtration and of expressivity in the abstract setting.
- **Section 5.2: Strong completeness.** This Section starts by reformulating and developing some recent results from [SP09] and [KR12] about strong completeness in coalgebraic logic, most notably the coalgebraic Jónsson-Tarski Theorem (Theorem 5.10). We then give a list of Theorems under which strong completeness is guaranteed and illustrate them with important Examples which are developed in some detail. The case of positive logics (i.e. logics defined over $\mathbf{DL}$) is treated first and strong completeness is shown in Theorem 5.16 for

a very wide class of logics which generalise positive modal logics and which we call *relational logics*. This Theorem generalises a result from [DP15]. The case of boolean logics is treated next, and a particularity of **Set**, namely that all epis are split, can be exploited to provide conditions for strong completeness in Theorem 5.24-5.27. We study Examples of boolean logics in detail, with a particular emphasis on establishing the injectivity and surjectivity of the semantic transformation and of its transpose. One of our main tools for this is the construction of retractions and sections, which is a technique we have not seen being used elsewhere in this context, and our direct proofs of weak completeness for modal logic and graded modal logics are to our knowledge new. In the case of boolean logics, the preservation of cofiltered limits is a key requirement on the semantic functor in order to get strong completeness and we present a generic technique based on the right Kan extension construction to modify a functor in a way that ensures this preservation property. This new technique, which we call *semantic completion* is applied to the case of graded modal logic, where it is shown to yield strong completeness.

- **Section 5.3: Completeness-via-canonicity.** This section gathers all the results of the thesis, into proofs of completeness-via-canonicity in the presence of canonical frame conditions. It starts by showing when the extension obtained by the coalgebraic Jónsson-Tarski Theorem coincides with the canonical extension defined in Chapter 2. To our knowledge, this question was never addressed, let alone answered, previously and Theorems 5.44 and 5.45 provide some clarity about the exact relation between the *syntactic* notion of canonical extension, and the *semantic* notion of canonical model at an abstract, coalgebraic, level. We then give coalgebraic completeness-via-canonicity results in the case of logics defined by endofunctors in **DL** or **BA** which define DLEs or BAEs (Theorems 5.47-5.49). We illustrate the results with Examples which include classical and graded modal logics, as well as intuitionistic logic and the distributive Lambek calculus (these examples have been published in [DP15]). We finish the Chapter with the more complicated case of logics whose defining endofunctor must be presented as the quotient of a polynomial functor. We detail a procedure for proving completeness-via-canonicity in this case too by exploiting the results of Chapter 4. The method is illustrated in the case of the nabla logic for the finitary covariant powerset functor P_ω .

1.2 Notation and Conventions

All important terms will be written in bold font where they are first introduced, the corresponding entry in the index will usually link to this place in the text. In terms of notational convention, we have tried to adhere to the following guiding principles:

1. Categories are denoted in bold font, e.g. $\mathbf{C}$, $\mathbf{Set}$, $\mathbf{BA}$, ... and we reserve the letters $\mathbf{I}$ and $\mathbf{J}$ for filtered or directed (index) categories
2. Functors come in a variety of shapes depending on custom and requirements. However, functors defining diagram will be denoted by a calligraphic font, typically $\mathscr{D}$. Many functors will be denoted using a sans-serif font, in particular free and forgetful functors will be written as $\mathbf{F}$, $\mathbf{G}$ and $\mathbf{U}$, $\mathbf{V}$ respectively, and the identity functor on a category $\mathbf{C}$ as $\mathbf{Id}_{\mathbf{C}}$. We will use three types of powerset functors: $\mathcal{P} : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{BA}$, $\mathcal{Q} : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$ and $\mathbf{P} : \mathbf{Set} \rightarrow \mathbf{Set}$.
3. Natural transformations will be denoted by lower-case Greek letters, typically $\delta, \epsilon, \zeta, \eta, \theta$.
4. Sets will be denoted by X, Y, Z and subsets by U, V, W , whilst elements will be denoted by x, y, z
5. Objects in a general category will be denoted by A, B, C and morphisms will be written using lower case Roman letters, typically, f, g, h, k . The letters i, m will often be used for monomorphisms, the letters e, p, q for epimorphisms.
6. We follow the notation of [KKV12b] for anything related to the nabla-style of coalgebraic logic. In particular, formulas will be denoted by $a, b, c, \dots$ throughout. Propositional variables will be written as $p, q, r, \dots$
7. When the need arises to differentiate between a boolean algebra and its carrier, the former will typically be written as $\mathbf{A}$ and the latter by A .
8. The category of algebras for an endofunctor $L : \mathbf{C} \rightarrow \mathbf{C}$ will be denoted $\mathbf{Alg}_{\mathbf{C}}(L)$. Dually the category of coalgebras for an endofunctor $T : \mathbf{C} \rightarrow \mathbf{C}$ will be denoted $\mathbf{CoAlg}_{\mathbf{C}}(T)$. We will drop the subscript category if $\mathbf{C} = \mathbf{Set}$ and there is no ambiguity. In the case of algebra for a monad $T : \mathbf{C} \rightarrow \mathbf{C}$, we stick to conventional notation $\mathbf{C}^T$. Generic structure maps and transition maps will be written using lower case Greek letters.
9. We have opted for the lower case spelling of ‘boolean’, both because it will occur very frequently in the text, and because we feel that it is a greater honour to George Boole to have his name turned into a regular adjective rather than a unusual, capitalized, one.

1.3 Reasoning structures

One of the attractive features of coalgebraic logic is its relative parametricity in the choice of algebraic structure for the syntax of the logics. As long as a dual adjunction exists with the category which the coalgebraic models will inhabit, we are free to choose the ‘reasoning kernel’ of our logic. In other words, the framework of coalgebraic logic is parametric in the choice of a fundamental reasoning structure, as long as it comes with a dual adjunction to an appropriate ‘semantic’ category. Several such adjunctions are presented in [JS10]. In this work we will focus on two reasoning structures: the category of boolean algebras (**BA**) and the category of distributive lattices (**DL**). They will play a key role both as the main algebraic environment in which to study canonicity, and as the base category in which the ‘abstract’ version of coalgebraic logic is formulated. This latter use will require a fair amount of categorical information about **BA** and **DL**, and algebras defined over them. Throughout this work we use **A** to denote either **BA** or **DL**, thus a statement about **A** will be equally applicable to either categories.

1.3.1 The category **BA**

Let us briefly summarize some of the fundamental properties of the category **BA** whose objects are boolean algebras and whose morphisms are boolean homomorphisms. We assume all the basic facts about boolean algebras (BAs), for details we refer the reader to the excellent introduction by Givant and Halmos [GH09]. Perhaps the most important feature of **BA** is that, together with the familiar (faithful) forgetful functor $U : \mathbf{BA} \rightarrow \mathbf{Set}$, it forms a concrete category $(\mathbf{BA}, U)$ with all the attributes of a finitary algebraic category. Since these attributes are common to all such categories, we will simply list the following useful facts and refer the reader to e.g. [AHS04, AR94, ARV10, Bor94a, Bor94b, Mac98] for all the details

Fact 1 U has a left adjoint $F : \mathbf{Set} \rightarrow \mathbf{BA}$ which builds the free boolean algebra over a given set (of propositional variable).

Fact 2 U is therefore continuous, i.e. it preserves all limits. In fact, U *creates* all limits.

Fact 3 As a consequence of the above, **BA** is complete, since **Set** is.

Fact 4 All of the above facts are a consequence of the fact that U is **monadic** (see e.g. Proposition 20.12 of [AHS04]), i.e. that UF is a monad and that **BA** is isomorphic to the category of $\mathbf{Set}^{UF}$ of UF -algebras (in the sense of algebra for a monad).

- Fact 5 Dually, $\mathbf{F}$ is co-continuous, i.e. it preserves all the colimits.
- Fact 6 Since $\mathbf{F} \dashv \mathbf{U}$ is monadic, it is in particular of descent type (see [KP93, VK11]) and thus its counit $\epsilon : \mathbf{F}\mathbf{U} \rightarrow \mathbf{Id}_{\mathbf{BA}}$ is a regular epi-transformation.
- Fact 7 $\mathbf{U}$ creates filtered colimits (see e.g. [Mac98] p.213).
- Fact 8 $\mathbf{BA}$ is a **regular category**, i.e. it is finitely complete, it has coequalizers of kernel pairs and the pullback of a regular epimorphism along any morphism is a regular epimorphism.
- Fact 9 In particular, $\mathbf{BA}$ has regular-epi-mono factorisation.
- Fact 10 Regular epimorphisms are those whose underlying set function is surjective
- Fact 11 Like all finitary varieties (see Corollary 3.7 of [AR94]), $\mathbf{BA}$ is a locally finitely presentable category (see Chapter 3 for details), and in particular is cocomplete.
- Fact 12 The free boolean algebras over finite sets of variables are regular projectives, i.e. projective with respect to regular epis, and they form a dense collection, i.e. every boolean algebra is a canonical colimit of free boolean algebras
- Fact 13 Moreover, $\mathbf{BA}$ is a locally finite variety (see [Bez01]), i.e. finitely generated boolean algebras are finite (i.e. have a finite underlying set) (see [GH09] for a proof of this).

Apart from the last fact, we would like to emphasise that these facts apply to any finitary variety, in particular to the boolean algebras expansions (BAE) which we will soon consider.

Recall that, as is the case in any finitary variety, boolean algebras have a **presentation by generators and relations**. Formally, this means that for any A in $\mathbf{BA}$, we can find a set X , from which freely generated terms are built, and a pair of jointly monic arrows $e_1, e_2 : E \rightarrow \mathbf{U}FX$ which selects freely generated terms to be identified, such that A is the coequalizer of the adjoint transpose maps $\hat{e}_1, \hat{e}_2 : FE \rightarrow FX$ (i.e. the maps freely generated from e_1, e_2).

$$FE \begin{array}{c} \xrightarrow{\hat{e}_1} \\ \xrightarrow{\hat{e}_2} \end{array} FX \xrightarrow{q} \twoheadrightarrow A$$

If X is finite, then A is said to be **finitely generated**. If both X and E are finite, then A is said to be **finitely presented**. Note that q is a regular epi, and indeed we could abstract away from e_1, e_2 and write A as a regular quotient of a free algebra, however the information on the cardinality of E is then a

lot less transparent. Typically, an algebra will have numerous presentations by generators and relations, but to see that such a presentation always exists, note that from Fact 6, every boolean algebra A is a regular quotient of a free boolean algebra:

$$\ker \epsilon_A \begin{array}{c} \xrightarrow{p_1} \\ \xrightarrow{p_2} \end{array} \text{FUA} \xrightarrow{\epsilon_A} A$$

where p_1, p_2 are the projections of the kernel pair $\ker \epsilon_A$ (all regular epis are the coequalizers of their kernel pairs in a regular category). This presentation is trivial in the sense that if we consider the adjoint transpose $\hat{\epsilon}_A : \text{UA} \rightarrow \text{UA}$ of ϵ_A , it is easy to check that $\hat{\epsilon}_A = \text{Id}_{\text{UA}}$, which makes this presentation pretty poor in terms of information compression.

We conclude with a comment on colimits in **BA**. As we have just stated, **BA** is cocomplete, but unlike limits which are essentially built in **Set** (by Fact 2), colimits are less straightforward (unless, by Fact 7, they are filtered). Coproducts in **BA** are usually built as products in the dual category of Stone spaces, however, they can be built directly in **BA** as an example of ‘free product’ and this can be done particularly naturally if we use presentations by generators and relations. Consider A, B in **BA** and assume that they have presentations as coequalizers of $e_1, e_2 : FE \rightarrow FX$ and $e'_1, e'_2 : FE' \rightarrow FX'$ respectively. Now consider the following commutative diagram:

$$\begin{array}{ccccc} FE & \xrightarrow{i_E} & FE + FE' & \xleftarrow{i_{E'}} & FE' \\ \downarrow e_1 & & \downarrow i_E \circ e_1 + i_{E'} \circ e'_1 & & \downarrow e'_1 \\ & & \downarrow & & \downarrow e'_2 \\ FX & \xrightarrow{i_{FX}} & FX + FX' \simeq F(X + X') & \xleftarrow{i_{FX'}} & FX' \\ \downarrow q & & \downarrow i_A \circ q + i_B \circ q' & & \downarrow q' \\ A & \xrightarrow{i_A} & A + B & \xleftarrow{i_B} & A \end{array}$$

It is straightforward to check that $i_A \circ q + i_B \circ q'$ is indeed the coequalizer of the two arrows above it: note first that by the universal property of coproducts it is a cocone of the diagram defining the coequalizer, since it is a rival cocone for $FX + FX'$ and the diagram commutes. To see that it is a limiting cocone, assume that some C coequalizes $i_E \circ e_1 + i_{E'} \circ e'_1$ and $i_E \circ e_2 + i_{E'} \circ e'_2$, then it is easy to check that by commutativity of the diagram it would also coequalize e_1, e_2 and e'_1, e'_2 . There would therefore exist unique arrows $A \rightarrow C$ and $B \rightarrow C$, and this would in turn trigger the existence of a unique arrow $A + B \rightarrow C$ proving that $A + B$ is indeed the coequalizer. What this means, is that the coproduct $A + B$ can be described as the boolean algebra given by the presentation whose set of generators is $X + X'$, i.e. the disjoint union of the generators of A and B , and whose relations are given by the coproduct of the relations defining A and B .

Coequalizers are more straightforward: let $f, g : A \rightarrow B$ be two boolean homomorphisms. We define the equivalence relation $\equiv$ (more on equivalence relations in **BA** in the next section) as the smallest equivalence relation on the set UB generated by

$$b \equiv b' \iff \exists a \in A \text{ s.th. } b = f(a), b' = g(a)$$

for $b, b' \in UB$. The relation $\equiv$ is, by construction, an equivalence relation on the carrier of B , but it is in fact also an equivalence relation on the boolean algebra B , i.e. a sub-boolean algebra of $B \times B^1$. This is an easy consequence of the fact that f, g are boolean homomorphisms. Indeed if $b_1 \equiv b'_1$ and $b_2 \equiv b'_2$, then there exist a_1, a_2 such that $f(a_1) = b_1, g(a_1) = b'_1, f(a_2) = b_2, g(a_2) = b'_2$ and thus $f(a_1 \wedge a_2) = f(a_1) \wedge f(a_2) = b_1 \wedge b_2$ and $g(a_1 \wedge a_2) = g(a_1) \wedge g(a_2) = b'_1 \wedge b'_2$, i.e. $b_1 \wedge b_2 \equiv b'_1 \wedge b'_2$. A similar argument shows that $\equiv$ is closed under $\vee$ and $\neg$. The quotient $B/\equiv$ is defined as the set of equivalence classes $[b], b \in B$ of $\equiv$ together with the obvious operations, for example $[b_1] \wedge [b_2] = [b_1 \wedge b_2]$ etc. These operations are independent of the choice of representative precisely because $\equiv$ is a boolean sub-algebra of $B \times B$. The coequalizer of f, g is then the morphism $q : B \twoheadrightarrow B/\equiv$ mapping an element $b \in B$ to its equivalence class $[b]$. It is a boolean homomorphism by construction of $B/\equiv$.

We conclude this examination of **BA** by introducing the dual adjunction which lies at the heart of its use in coalgebraic logic, namely the adjunction $\mathbf{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}$, where $\mathbf{Uf}$ is the functor associating to each boolean algebra its set of ultrafilters and to each **BA**-morphism its inverse image, and $\mathcal{P}$ is the functor associating to each set the boolean algebra of its subsets (under set operations) and to each function its inverse image.

Note that for each $A \in \mathbf{BA}$, $\mathbf{Uf}A = \text{Hom}(A, \mathbf{2})$, where $\mathbf{2} = \{\top, \perp\}$ is the two element boolean algebra, and for each set $UP(X) = \text{Hom}(X, \mathbf{2})$, where $\mathbf{2}$ is the two elements set.

1.3.2 The category DL

The category of distributive lattices is also a finitary algebraic category and all the facts listed above for **BA** also apply to **DL** with the obvious free-forgetful adjunction $\mathbf{F} \dashv \mathbf{U} : \mathbf{DL} \rightarrow \mathbf{Set}$. The construction of colimits works just as in **BA**. It is interesting and important to note that Fact 13 also holds in **DL**, i.e. finitely generated distributive lattices are finite (see e.g. [BD75]). Local

¹We use the term *equivalence relation* where many authors use the term *congruence*. The reason for this choice of nomenclature is twofold. Firstly, we consider relations in a categorical setting, i.e. as subobjects in a given category, thus relations naturally inherit the structure of the ambient category. Secondly, the term *congruence* will be used heavily in the context of the congruence rule of coalgebraic logics. This rule is ultimately related to the notion of a subobject in a category of algebras, but we like to keep the two notions separate.

finiteness will prove to be extremely important in Chapters 3 and 5, and this justifies to a great extent our focus on **BA** and **DL**, but not for example on **HA**, the category of Heyting algebras.

The fundamental dual adjunction corresponding to $\mathbf{Uf} \dashv \mathcal{P}$ in the case of **BA** is the adjunction $\mathbf{Pf} \dashv \mathcal{U} : \mathbf{DL} \rightarrow \mathbf{Pos}$ where **Pos** is the category of posets and monotone maps, **Pf** is the functor sending a distributive lattice to its poset of prime filters and a **DL**-morphism to its inverse image (which is easily seen to be monotone), and $\mathcal{U}$ is the functor sending a poset to the distributive lattice of its downsets and a monotone map to its inverse image.

Note that in the case of boolean algebra, the two adjunctions, i.e. $\mathbf{F} \dashv \mathbf{U}$ and $\mathbf{Uf} \dashv \mathcal{P}$ involve the same categories (although one is covariant and the other contravariant). However, for **DL** the situation is different since the prime filter functor **Pf** takes values in **Pos**. However, since each distributive lattice is a poset, there exists a forgetful functor $\mathbf{U} : \mathbf{DL} \rightarrow \mathbf{Pos}$. Moreover, this functor has a left adjoint $\mathbf{F} : \mathbf{Pos} \rightarrow \mathbf{DL}$ which builds the free distributive lattice over a poset (see [Mon67]). As it turns out this functor yields a ‘finitary version’ of the canonical extension functor which we will describe in Chapter 2.

1.3.3 Boolean algebra and distributive lattice expansions

We now turn to the algebraic structures that underpin most of modal logic: BAOs and BAEs, as well as their distributive lattice counterparts. In this subsection, as in Chapter 2, we will use the convention of writing an object of **A** (i.e. **BA** or **DL**) as $\mathbb{A}$ and their underlying set as A if there is any danger of confusion between the two entities. Recall that, given a finitary signature $(\Sigma, \text{ar} : \Sigma \rightarrow \mathbb{N})$, we have:

Definition 1.1. A Σ -**Boolean Algebra Expansion** (BAE), or expanded boolean algebra, is a boolean algebra $\mathbb{A}$ together with maps $f_s : A^{\text{ar}(s)} \rightarrow A$ for each $s \in \Sigma$, where A is the underlying set of $\mathbb{A}$. In other words, a BAE is a pair $\mathcal{A} = (\mathbb{A}, (f_s)_{s \in \Sigma})$. A function $f_s : A^{\text{ar}(s)} \rightarrow A$ which preserves the bottom element $\perp \in A$ and joins in each of its arguments, is called an **operator**. A BAE whose set of extending functions are all operators is called a **Boolean Algebra with Operators** (BAO). We similarly define a Σ -**Distributive Lattice Expansion** (DLE) as a distributive lattice $\mathbb{A}$ together with maps $f_s : A^{\text{ar}(s)} \rightarrow A$ for each $s \in \Sigma$.

The theory of BAOs is well understood (see the seminal [JT51] and [Jón94, dRV95, Ven06]), but is tailored to the study of classical (relational) modal logics. In the context of coalgebraic logics we are interested in a larger class of interpretations and this is reflected in the fact that - in its predicate lifting flavour (see Section 1.5) - the ‘operators’ of coalgebraic modal logics are not,

well, operators; they do not in general preserve $\perp$ or joins in each of their argument. We will therefore consider the weaker structure of BAEs where the functions $f_s : A^{\text{ar}(s)} \rightarrow A$ have, in general, no preservation properties. One purpose of Chapter 2 will be to isolate some properties of these functions, other than ‘operator-ness’, which are relevant and useful to the study of canonicity.

For now, let us develop some lightweight categorical infrastructure. For a given signature Σ , the collection of all Σ -BAEs (resp. DLEs) forms a category which we will denote by **BAE**(Σ) (resp. **DLE**(Σ)), or simply **BAE** (resp. **DLE**) when there is no ambiguity about the signature. This category is defined as follows: the objects are of course the BAEs (resp. DLEs) and a morphism ϕ between two BAEs (resp. DLEs) $\mathcal{A} = (\mathbb{A}, (f_s)_{s \in \Sigma})$ and $\mathcal{B} = (\mathbb{B}, (g_s)_{s \in \Sigma})$ is a **BA**- (resp. **DL**-) morphism $\phi : \mathbb{A} \rightarrow \mathbb{B}$ such that for all $s \in \Sigma$,

$$\phi \circ f_s = g_s \circ \langle \phi, \dots \phi \rangle$$

where $\langle \phi, \dots \phi \rangle$ is the n -fold product of ϕ and $n = \text{ar}(f_s) = \text{ar}(g_s)$.

An important fact about the categories **BAE**(Σ) and **DLE**(Σ) is that they can be characterised as categories of algebras over **BA** and **DL** respectively. To describe this relation, let **A** denote either **BA** or **DL** and **AE**(Σ) (or **AE** if there is no ambiguity over the signature) denote either **BAE**(Σ) or **DLE**(Σ). We will talk about ‘AEs’ (expansions of objects of **A**) if we want to refer to either DLEs or BAEs. We now build a polynomial endofunctor $S_\Sigma : \mathbf{A} \rightarrow \mathbf{A}$:

$$S_\Sigma(\mathbb{A}) = F \left(\coprod_{s \in \Sigma} (A)^{\text{ar}(s)} \right) \quad (1.1)$$

(where $F \dashv U$ are the obvious free/forgetful functors) and we have:

Proposition 1.2. *The category $\mathbf{Alg}_{\mathbf{A}}(S_\Sigma)$ is equivalent to $\mathbf{AE}(\Sigma)$.*

Proof. Given an object $\mathbb{A}$ of **A**, the freeness of the construction of $S_\Sigma(\mathbb{A})$ means that the structure map $\alpha : S_\Sigma(\mathbb{A}) \rightarrow \mathbb{A}$ of any S_Σ -algebra in **A** is in unique correspondence with its adjoint transpose $\tilde{\alpha} : \coprod_{s \in \Sigma} (A)^{\text{ar}(s)} \rightarrow A$, where $A = U\mathbb{A}$, i.e. by a collection of maps $f_s : A^{\text{ar}(s)} \rightarrow A, s \in \Sigma$. This collection of maps together with $\mathbb{A}$ clearly define an object of **AE**(Σ). Conversely, given such an object $(\mathbb{A}, (f_s)_{s \in \Sigma})$, the coproduct morphism $\coprod_{s \in \Sigma} f_s : \coprod_{s \in \Sigma} (A)^{\text{ar}(s)} \rightarrow A$ can be freely extended to a **A**-morphism $\alpha : S_\Sigma(\mathbb{A}) \rightarrow \mathbb{A}$. These two operations define functors from $\mathbf{Alg}_{\mathbf{A}}(S_\Sigma)$ to **AE**(Σ) and back, and the compositions of these two functors are clearly isomorphic to the identities on $\mathbf{Alg}_{\mathbf{A}}(S_\Sigma)$ and **AE**(Σ) respectively. $\square$

Following the characterisation of the coproduct in **A** as a free product, the fact that free boolean algebras are presented without any relations, and that

products are constructed in **Set**, it is easy to see that Eq. (1.1) can be rewritten as $S_\Sigma \mathbb{A} = \coprod_{s \in \Sigma} (\mathbb{A})^{\text{ar}(s)}$, where the product and coproduct are now taken in **A**. In other words we can view S_Σ as the formal equivalent in **BA** or **DL** of its **Set** counterpart.

We will now examine how the categories **BAE**(Σ) and **DLE**(Σ) are related to the categories **BA** and **DL**. Using the same notation as above, we define $V : \mathbf{AE} \rightarrow \mathbf{A}$ as the obvious forgetful functor, i.e. we make **AE** concrete over **A**. Initial algebras for endofunctors on **A** will be important in the following Proposition and we cite without proof the following classic result, for details see for example [Awo10] 10.12.

Theorem 1.3. *Let $\mathbf{C}$ be a category with an initial element 0 and ω -colimits (i.e. colimits of diagrams whose index category is the totally ordered set of natural numbers), and let $T : \mathbf{C} \rightarrow \mathbf{C}$, then if T preserves ω -colimits it has an initial algebra.*

Corollary 1.4. *Let $T : \mathbf{C} \rightarrow \mathbf{C}$ be an finitary endofunctor on a locally finitely presentable category, then T has an initial algebra.*

Proof. Every locally presentable is cocomplete, and thus has an initial object. Moreover, a diagram indexed by ω is trivially directed, thus preserved by finitary functors. The conclusion then follows from Theorem 1.3 $\square$

Proposition 1.5. *The forgetful functor $V : \mathbf{AE} \rightarrow \mathbf{A}$ has a left adjoint $G : \mathbf{A} \rightarrow \mathbf{AE}$.*

Proof. It is enough to show that for any boolean algebra A (we will use the forgetful functor explicitly to denote the carrier in this proof), the functor $S_\Sigma(-) + A$ has an initial algebra, since this amounts to the existence of a free S_Σ -algebra over A , and by Proposition 1.2 this will mean that **AE** has free objects. We will use Corollary 1.4 above and show that S_Σ is finitary, i.e. preserves directed colimits. We are here using some notions which will be fully developed in Chapter 3, but we will only need one very basic concept, namely that finite sets are finitely presentable, i.e. their hom functors preserves directed colimits. Let $(I, \leq)$ be a directed set and let $\mathcal{D} : I \rightarrow \mathbf{A}$ be a directed diagram in **A**, we

then have

$$\begin{aligned}
S_\Sigma(\operatorname{colim}_i A_i) &= F \left(\prod_{s \in \Sigma} (\operatorname{U} \operatorname{colim}_i A_i)^{\operatorname{ar}(s)} \right) \\
&\stackrel{1}{=} F \left(\prod_{s \in \Sigma} (\operatorname{colim}_i \operatorname{U} A_i)^{\operatorname{ar}(s)} \right) \\
&\stackrel{2}{=} F \left(\prod_{s \in \Sigma} \operatorname{hom}(\operatorname{ar}(s), \operatorname{colim}_i \operatorname{U} A_i) \right) \\
&\stackrel{3}{=} F \left(\prod_{s \in \Sigma} \operatorname{colim}_i \operatorname{hom}(\operatorname{ar}(s), \operatorname{U} A_i) \right) \\
&\stackrel{4}{=} F \left(\operatorname{colim}_i \prod_{s \in \Sigma} \operatorname{hom}(\operatorname{ar}(s), \operatorname{U} A_i) \right) \\
&\stackrel{5}{=} \operatorname{colim}_i F \left(\prod_{s \in \Sigma} (\operatorname{U} A_i)^{\operatorname{ar}(s)} \right) \\
&= \operatorname{colim}_i S_\Sigma A_i
\end{aligned}$$

where (1) follows from Fact 7, (2) is a rewriting of the hom set in **Set**, (3) follows from the fact that since the signature is finitary $\operatorname{ar}(s)$ is finite and finite sets are finitely presentable, (4) is a special case of the general rule that colimits commute, and finally (5) follows from the fact that F is cocontinuous by virtue of being left-adjoint. Thus S_Σ is finitary, as is $S_\Sigma(-) + A$ for any boolean algebra A . By corollary 1.4, $S_\Sigma(-) + A$ has an initial algebra $\mu(S_\Sigma(-) + A)$ for every A .

So let us define $G : \mathbf{A} \rightarrow \mathbf{AE}$ by $G(A) = \mu(S_\Sigma(-) + A)$ and let us check that we do indeed have $G \dashv V$, i.e. that for any $\mathbf{A}$ -morphism $f : A \rightarrow V(B, \alpha)$ where $\alpha : S_\Sigma(B) \rightarrow B$ (i.e. any morphism $A \rightarrow B$), there exists a unique $\mathbf{Alg}_{\mathbf{A}}(S_\Sigma)$ -morphism $G(A) \rightarrow (B, \alpha)$. This is immediate by initiality of $G(A)$ as can be seen from the following diagram

$$\begin{array}{ccc}
S_\Sigma(GA) + A = S_\Sigma(\mu(S_\Sigma(-) + A)) & \xrightarrow{[S_\Sigma \hat{f} + \operatorname{Id}_A]} & S_\Sigma B + A \\
\downarrow & & \downarrow [\alpha + f] \\
GA = \mu(S_\Sigma(-) + A) & \xrightarrow{\hat{f}} & B
\end{array}$$

and $\hat{f}$ defines an $\mathbf{Alg}_{\mathbf{A}}(S_\Sigma)$ -morphism. □

Proposition 1.6. $G \circ F \dashv U \circ V$.

Proof. Let us consider A in **Set**, (B, α) in $\mathbf{Alg}_{\mathbf{A}}(S_{\Sigma})$ and $f : A \rightarrow U \circ V((B, \alpha))$. We need to show that f determines a unique morphism $g : G \circ F(A) \rightarrow (B, \alpha)$. But since $F \dashv U$, f determines a unique map $h : F(A) \rightarrow V((B, \alpha))$ and since $G \dashv V$, this h in turns determines a unique $g : G \circ F(A) \rightarrow (B, \alpha)$. $\square$

This is clearly a much more general result, i.e. adjunctions compose.

1.3.4 Algebras in **BA** and **DL**

We have just encountered a particular class of algebras in **BA** and **DL** which are defined by certain finitary polynomial endofunctors. In fact, algebras in **A** (where again **A** is either **BA** or **DL**) for general finitary functors are a very useful tool to study coalgebraic logic. However, they introduce a nesting of algebraic constructions (i.e. algebras defined over algebras) which is not completely straightforward. We have just seen in Proposition 1.5 that for finitary polynomial functors $S_{\Sigma} : \mathbf{A} \rightarrow \mathbf{A}$, free S_{Σ} -algebras exist. This makes finitary polynomial functors **variators**. A variator is an endofunctor $L : \mathbf{C} \rightarrow \mathbf{C}$ such that the forgetful functor $U : \mathbf{Alg}_{\mathbf{C}}(L) \rightarrow \mathbf{C}$, has a left adjoint $F : \mathbf{C} \rightarrow \mathbf{Alg}_{\mathbf{C}}(L)$, i.e. such that $\mathbf{Alg}_{\mathbf{C}}(L)$ has free objects. As we have seen from Corollary 1.4, every finitary functor on a locally presentable category is a variator. The following property of variators is crucial.

Theorem 1.7 ([AHS04] Theorem 20.56). *If $T : \mathbf{C} \rightarrow \mathbf{C}$ is a variator, then $\mathbf{Alg}_{\mathbf{C}}(T)$ is monadic over $\mathbf{C}$*

For any variator $L : \mathbf{A} \rightarrow \mathbf{A}$, we thus have two monadic functors: $U : \mathbf{A} \rightarrow \mathbf{Set}$ and $V : \mathbf{Alg}_{\mathbf{A}}(L) \rightarrow \mathbf{A}$. However, as shown in Example 20.45 of [AHS04], monadic functors do *not* in general compose, and in particular, $UV : \mathbf{Alg}_{\mathbf{A}}(L) \rightarrow \mathbf{Set}$ is not, in general, monadic. Since monadicity is a desirable feature, it is important to identify which additional feature on U and V will make their composition monadic. This feature is detailed in Exercise 23.F of [AHS04], and is the notion of **variational functor**. A regularly monadic functor $U : \mathbf{C} \rightarrow \mathbf{D}$ (i.e. U is monadic, $\mathbf{D}$ has regular epi-mono factorisations, and U preserves regular epis) is variational if $\mathbf{D}$ has pullbacks and U creates coequalizers of kernel pairs, also known as exact sequences. We will go through all these terms in more details in the next section; for now, let us just state the following three propositions. The first is a re-wording of Proposition 1.16, the second is a re-wording of Proposition 1.20.

Proposition 1.8. *Every monadic functor over **Set** is variational.*

Proposition 1.9. *If $L : \mathbf{C} \rightarrow \mathbf{C}$ is a variator over a regular category, and L preserves regular epis, then $U : \mathbf{Alg}_{\mathbf{C}}(L) \rightarrow \mathbf{C}$ is variational.*

Proposition 1.10 ([AHS04] 23.F (f)). *Varietal functors are monadic and compose.*

These Propositions provide an answer the question of knowing when $UV : \mathbf{Alg}_A(L) \rightarrow \mathbf{Set}$ is monadic: it is monadic if L preserves regular epimorphisms. Note that no assumption is needed for U to be varietal, it is essentially a feature of exact sequences in $\mathbf{Set}$ which gives this nice behaviour for free. Note also that preservation of regular epis has also been identified as a desirable feature of endofunctors on algebraic varieties in [Mye11], where this condition guarantees the existence of a notion of functor presentation.

1.4 Relations, categorically


This section will be fairly technical and aims to introduce and develop the familiar $\mathbf{Set}$ notions of relations, relation liftings, equivalence relations and fully invariant relations, in a setting which is abstract enough to capture the categories encountered in the previous section, namely $\mathbf{BA}$, $\mathbf{DL}$, and $\mathbf{Alg}_A(L)$ for a varietor $L : \mathbf{A} \rightarrow \mathbf{A}$ (where $\mathbf{A}$ is $\mathbf{BA}$ or $\mathbf{DL}$). The most natural and convenient environment to achieve this is to work in the framework of *regular categories*. The special case of the category of $\mathbf{Set}$ will of course be very useful in its own right, particularly in the context of the nabla and abstract flavours of coalgebraic logic.

1.4.1 Relations

There is a large literature on the topic of relations in categories other than $\mathbf{Set}$, see e.g. [Kel91], which considers various generalization of the concept and various classes of suitable categories, but the classical theory of relations in a category other than $\mathbf{Set}$ takes place in regular categories. Which suits us fine since $\mathbf{BA}$ and $\mathbf{DL}$ are regular category. The key features of regular categories which allow them to have a good notion of relations is that (1) they have regular epi/mono factorisation, (2) the class of regular epimorphisms is closed under pullback, i.e. the factorisation system is *stable*.


A relation $R \subseteq X \times Y$ between X, Y in $\mathbf{Set}$ can be seen as a subobject of $X \times Y$, and this notion can clearly be generalized to other categories (for more details on subobjects, see Section 3.1). For a well-powered regular category $\mathbf{C}$, let use write $\mathbf{Rel}(\mathbf{C})$ for the category whose objects are the same objects as $\mathbf{C}$ but whose morphism $R : A \rightarrow B$ are subobjects of $A \times B$. We need to show how morphisms in $\mathbf{Rel}(\mathbf{C})$ are composed. Let $R : A \rightarrow B$ and $S : B \rightarrow C$ be $\mathbf{Rel}(\mathbf{C})$ -morphisms, we compose $S \circ R$ in the same way as spans (since relations are spans), i.e. we build the pullback of $R \xrightarrow{p_2^R} B \xleftarrow{p_1^S} C$ where p_i^R, p_i^S , $i = 1, 2$ are the i^{th} projections of R and S respectively. We then take the regular epi-mono

factorisation of the map from the limiting cone of this pullback to $A \times C$ (which exists by universality of products) and define $S \circ R$ to be the subobject of $A \times C$ thus defined.



It can be shown that composition defined in this way is associative iff the class of regular epis is stable under pullback, which is the case by definition in a regular category. This explains why regular categories are a particularly natural setting in which to generalize relations.

There exists a very useful alternative description of a relation. In a finitely complete category, relations $R \rightharpoonup A \times A$ are in one-to-one correspondence with pairs of maps $e_1, e_2 : R \rightarrow A$ which are jointly monic, i.e. such that for any $f, g : B \rightarrow R$, if $e_1 \circ f = e_1 \circ g$ and $e_2 \circ f = e_2 \circ g$, then $f = g$. This correspondence is clear from the following diagrams



where it is not difficult to check that $e_1 \times e_2$ is monic if e_1, e_2 are jointly monic, and conversely, that $\pi_1 \circ m, \pi_2 \circ m$ are jointly monic if m is monic. In what follows, having these equivalent descriptions of what a relation is will prove very useful.

1.4.2 Relation liftings

With $\mathbf{Rel}(\mathbf{C})$ defined, it is natural to consider relation liftings for endofunctor $L : \mathbf{C} \rightarrow \mathbf{C}$. In fact, being in a regular category ensures that what is done in $\mathbf{Set}$, i.e. essentially taking a direct image, can be transferred to $\mathbf{C}$. Formally, given a relation $R : A \rightarrow B$ and a $\mathbf{C}$ -endofunctor L , we define its relation lifting $\bar{L}R$ as the subobject of $LA \times LB$ defined as the image of LR in $LA \times LB$, i.e.

we build the following regular epi-mono factorisation:

$$\begin{array}{ccccc}
 LA & \xleftarrow{Lp_1} & LR & \xrightarrow{Lp_2} & LB \\
 & & \downarrow \scriptstyle Lp_1 \times Lp_2 & & \\
 & & \overline{LR} & & \\
 & \swarrow \scriptstyle \pi_1 & \downarrow & \searrow \scriptstyle \pi_2 & \\
 & & LA \times LB & &
 \end{array}$$

A proof of the following Proposition can be found in [CKW91], but involves advanced concepts of category theory. We can in fact prove it using the simple machinery of regular categories.

Proposition 1.11. *Let $\mathbf{C}$ be a regular category, and let $R \rightharpoonup A \times B$, $S \rightharpoonup B \times C$ be relations in $\mathbf{C}$. Assume that $L : \mathbf{C} \rightarrow \mathbf{C}$ preserves regular epis and weak pullbacks, then $\overline{LS} \circ \overline{LR} = \overline{L(S \circ R)}$.*

Proof. We first show that $\overline{L(S \circ R)} \rightharpoonup \overline{LS} \circ \overline{LR}$. By definition, the relation $S \circ R$ is defined by building the pullback of $R \xrightarrow{p_2^R} B \xleftarrow{p_1^S} S$ where $p_i^R, i = 1, 2$ are the projection from R to A, B , and similarly for S . By the regular epi-mono factorisation of $Lp_i^R, Lp_i^S, i = 1, 2$ we first construct $\overline{LS}$ and $\overline{LR}$. The relation $\overline{LS} \circ \overline{LR}$ is then defined by building the pullback of $\overline{LR} \xrightarrow{r_2} LB \xleftarrow{s_1} \overline{LS}$ where r_1, r_2 are the projections from $\overline{LR}$ to LA and LB respectively, and s_1, s_2 are the projections from $\overline{LS}$ to LB and LC respectively.

It is easy to see that L applied to the first pullback forms a cone for the diagram defining the latter:

$$\begin{array}{ccccc}
 Lpb(p_2^R, p_1^S) & \xrightarrow{\quad} & LS & & \\
 \downarrow & \searrow \scriptstyle u & \downarrow \scriptstyle q^S & & \\
 & & pb(r_2, s_1) & \xrightarrow{\pi_1} & \overline{LS} \\
 & & \downarrow \scriptstyle \pi_2 & & \downarrow \scriptstyle s_1 \\
 LR & \xrightarrow{q^R} & \overline{LR} & \xrightarrow{r_2} & LB \\
 & \searrow \scriptstyle Lp_2^R & & &
 \end{array}$$

and there must therefore exist a unique morphism $u : Lpb(p_2^R, p_1^S) \rightarrow pb(r_2, s_1)$. Since L preserves regular epi, we have a regular epi

$$Lpb(p_2^R, p_1^S) \twoheadrightarrow L(S \circ R) \twoheadrightarrow \overline{L(S \circ R)}$$

where the first regular epi is just L applied to the regular epi factor defining $S \circ R$, and the second regular epi is by definition of relation liftings. Combining our observations so far we have a square:

$$\begin{array}{ccc}
 \text{Lpb}(p_2^R, p_1^S) & \twoheadrightarrow & \overline{L}(S \circ R) \\
 \downarrow u & & \downarrow \\
 \text{pb}(r_2, s_1) & & \\
 \downarrow & & \downarrow \\
 \overline{L}S \circ \overline{L}R & \twoheadrightarrow & LA \times LC
 \end{array}$$

and since regular epis are strong, there must exist a unique diagonal fill-in, which, moreover, must be monic, i.e. $\overline{L}(S \circ R) \rightarrow \overline{L}S \circ \overline{L}R$ as claimed.

For the converse direction, it is enough to show that u is a regular epi, and we can then use the same square as above to find a diagonal fill-in $\overline{L}S \circ \overline{L}R \rightarrow \overline{L}(S \circ R)$. To see that u is indeed a regular epi, we use the fact that L weakly preserves pullbacks. We build the following two pullbacks

$$\begin{array}{ccccc}
 \text{Lpb}(p_2^R, p_1^S) & & & & \\
 \swarrow h & & & & \searrow \\
 & \text{pb}(Lp_2^R, Lp_1^S) & \xrightarrow{\quad} & LS & \\
 \downarrow & \swarrow v & & \downarrow q^S & \\
 & \text{pb}(r_2, s_1) & \xrightarrow{\pi_1} & \overline{L}S & \\
 & \downarrow \pi_2 & & \downarrow s_1 & \\
 LR & \xrightarrow{q^R} & \overline{L}R & \xrightarrow{r_2} & LB \\
 & \searrow & & & \swarrow \\
 & & & & Lp_1^S \\
 & & & & \uparrow \\
 & & & & Lp_2^R
 \end{array}$$

it is not hard to verify by taking intermediary pullbacks and using the pullback lemma that v is a regular epi (see [Bor94b] Lemma 2.1.2). Moreover, since L weakly preserves pullbacks there must exist a (not necessarily unique) morphism $h : \text{Lpb}(p_2^R, p_1^S) \rightarrow \text{pb}(Lp_2^R, Lp_1^S)$. By the fact that $\text{pb}(Lp_2^R, Lp_1^S)$ is the actual pullback of the same diagram, there must exist a unique arrow $w : \text{pb}(Lp_2^R, Lp_1^S) \rightarrow \text{Lpb}(p_2^R, p_1^S)$, and by the universal property of the pullback it is clear that $h \circ w = \text{Id}$. It is easy to check that morphisms with a right inverse such as h are coequalizers (in this case of Id and $h \circ w$), and thus h is a regular epi. We can therefore conclude that $u = v \circ h$ is a regular epi, since

regular epis compose in regular categories ([Bor94b] Corollary 2.1.5), and this concludes the proof. $\square$

The following proposition is an application of the previous one, and is particularly important for the nabla and abstract versions of coalgebraic logics, since it allows well-defined inductive definitions.

Proposition 1.12. *Let $L : \mathbf{C} \rightarrow \mathbf{C}$ where $\mathbf{C}$ is regular, then*

- (i) *Every $\mathbf{C}$ -morphism $f : A \rightarrow B$ induces a relation $\text{Gr}(f) \rightharpoonup A \times B$ and $\overline{L}\text{Gr}(f) = \text{Gr}(Lf)$*
- (ii) *Every relation $R \rightharpoonup A \times B$ has a converse relation $R^{\text{op}} \rightharpoonup B \times A$ and $\overline{L}(R^{\text{op}}) = (\overline{L}R)^{\text{op}}$*
- (iii) *If L weakly preserves pullbacks and preserves regular epis, then for any $A_1 \rightharpoonup B_1$, $A_2 \rightharpoonup B_2$ and $R \rightharpoonup B_1 \times B_2$*

$$\overline{L}(R \cap (A_1 \times A_2)) = \overline{L}R \cap (LA_1 \times LA_2)$$

Proof. (i) The relation R_f defined by a morphism is obtained by the pullback

$$\begin{array}{ccc} R_f & \xrightarrow{p_1} & B \\ p_2 \downarrow & & \downarrow \text{Id}_B \\ A & \xrightarrow{f} & B \end{array}$$

or, alternatively and equivalently, as the equalizer of $\pi_2, f \circ \pi_1 : A \times B \rightarrow B$. Note that since the identity map is an isomorphism, and the pullback of an iso is an iso, p_2 must be an iso. If we apply L to the previous diagram we get:

$$\begin{array}{ccccc} LR_f & & \xrightarrow{Lp_1} & & LB \\ & \searrow u & & \searrow q_1 & \\ & & R_{Lf} & \xrightarrow{q_1} & LB \\ & \searrow Lp_2 & \downarrow q_2 & & \downarrow \text{Id}_{LB} = L\text{Id}_B \\ & & LA & \xrightarrow{Lf} & LB \end{array}$$

The morphism q_2 must be an iso, and since any functor preserves isos, Lp_2 is also an iso. Therefore u must be an iso, and its image is isomorphic to its domain, i.e. $\overline{L}R_f = R_{Lf}$.

(ii) The converse relation R^{op} of a relation $R \rightharpoonup A \times B$ is also obtained by a pullback

$$\begin{array}{ccc} R^{\text{op}} & \longrightarrow & R \\ \downarrow & & \downarrow \\ B \times A & \xrightarrow{\pi_2 \times \pi_1} & A \times B \end{array}$$

And since $\pi_1 \times \pi_2$ is an iso, the proof follows the same lines as (i).

For (iii), let $m : R \rightharpoonup B_1 \times B_2$, $i_1 : A_1 \rightharpoonup B_1$ and $i_2 : A_2 \rightharpoonup B_2$. We will first show that

$$R \cap (A_1 \times A_2) = \text{Gr}(i_2)^{\text{op}} \circ R \circ \text{Gr}(i_1)$$

To achieve this, we first show that $R \cap (A_1 \times B_2) = R \circ \text{Gr}(i_1)$. To see this let us first build the appropriate pullback:

It is not too hard to check that by definition of Gr , p_1 must be an iso, and that $p_2 \circ p_1^{-1} = i_1$. Let us now show that $(pb(p_2, q_1), r_1, r_2)$ is the pullback of $A_1 \times B_2 \rightharpoonup B_1 \times B_2 \leftarrow R$. It is clearly a cone by commutativity of the above diagram. Now let $g : A \rightarrow R, h_1 \times h_2 : A \rightarrow A_1 \times B_2$ such that $m \circ g = (i_1 \times \text{Id}_{B_2}) \circ (h_1 \times h_2)$ be another cone for the same diagram. We then have

$$\begin{aligned} q_1 \circ g &= \pi_1^B \circ m \circ g \\ &= \pi_1^B \circ (i_1 \times \text{Id}_{B_1}) \circ (h_1 \times h_2) \\ &= i_1 \circ \pi_1 \circ (h_1 \times h_2) \\ &= p_2 \circ p_1^{-1} \circ \pi_1 \circ (h_1 \times h_2) \end{aligned}$$

i.e. $(A, g, p_1^{-1} \circ \pi_1 \circ (h_1 \times h_2))$ forms a cone for the diagram defining $pb(p_2, q_1)$ and so there must exist a unique morphism $u : A \rightarrow pb(p_2, q_1)$ such that $r_2 \circ u = g$ and $r_1 \circ u = p_1^{-1} \circ \pi_1 \circ (h_1 \times h_2)$. To see that this cone is also a cone for the intersection $R \cap A_1 \times B_2$, we now just need to check that $(p_1 \circ r_1 \times q_2 \circ r_2) \circ u =$

$(h_1 \times h_2)$ (since $r_2 \circ u = g$ is already satisfied). We just calculate:

$$\begin{aligned}
 (p_1 \circ r_1 \times q_2 \circ r_2) \circ u &= p_1 \circ r_1 \circ u \times q_2 \circ r_2 \circ u \\
 &= p_1 \circ p_1^{-1} \circ \pi_1 \circ (h_1 \times h_2) \times q_2 \circ g \\
 &= \pi_1 \circ (h_1 \times h_2) \times \pi_2^B \circ m \circ g \\
 &= h_1 \times \pi_2^B \circ (i_1 \circ \text{Id}_{B_2}) \circ (h_1 \circ h_2) \\
 &= h_1 \times h_2
 \end{aligned}$$

as required. Thus $pb(p_2, q_1) = R \cap (A_1 \times B_2)$, and in particular $p_1 \circ r_1 \times q_2 \circ r_2$ must be a mono, and therefore there is no regular epi part to its factorization and we can conclude

$$R \circ \text{Gr}(i_1) = R \cap (A_1 \times B_2)$$

A completely analogous proof shows that $\text{Gr}(i_2)^{\text{op}} \circ R = R \cap B_1 \times A_2$. And thus the composition $\text{Gr}(i_2)^{\text{op}} \circ R \circ \text{Gr}(i_1)$ is defined by first building the pullback of $R \cap A_1 \times B_2$ with $R \cap B_1 \times A_2$, and then taking its image. But since both the morphisms defining this pullback are monos, we get

$$\text{Gr}(i_2)^{\text{op}} \circ R \circ \text{Gr}(i_1) = R \cap (A_1 \times A_2)$$

To complete the proof, we use Proposition 1.11 on the above composition, as well as (i) and (ii) to get

$$\begin{aligned}
 \overline{L}(R \cap (A_1 \times A_2)) &= \overline{L}(\text{Gr}(i_2)^{\text{op}} \circ R \circ \text{Gr}(i_1)) \\
 &= \overline{L}\text{Gr}(i_2)^{\text{op}} \circ \overline{L}R \circ \overline{L}\text{Gr}(i_1) \\
 &= \text{Gr}(Li_2)^{\text{op}} \circ \overline{L}R \circ \text{Gr}(Li_1)
 \end{aligned}$$

and the same argument as above shows that the latter expression is in fact $\overline{L}R \cap (LA_1 \times LA_2)$. $\square$

1.4.3 Equivalence relations and exactness

A particular type of relation is the notion of **equivalence relation**, sometimes called ‘congruence’, but we prefer equivalence relation as it avoids confusion with the congruence rule, which, as we shall see a bit later, is related to the idea of a relation in **BA**, but not specifically to that of an *equivalence* relation in **BA**. Categorically, an equivalence relation R on an object A in a regular category **C** is a subobject $E \xrightarrow{e_1 \times e_2} A \times A$ which is an internal equivalence relation on A , where the axioms defining an equivalence are categorified as follows:

- reflexivity: the diagonal Δ_A is a subobject of E
- symmetry: $e_1 \times e_2 = e_2 \times e_1$ as subobjects of E

- transitivity: if $E \xleftarrow{e'_1} E \times_A E \xrightarrow{e'_2} E$ is the pullback of $E \xrightarrow{e_1} A \xleftarrow{e_2} E$, then $(e_1 \circ e'_1) \times (e_2 \circ e'_2)$ is a subobject of $e_1 \times e_2$.

Kernel pairs provide the typical example of equivalence relation.

Lemma 1.13. *Let $\mathbf{C}$ be a regular category, and let $f : A \rightarrow B$ be a $\mathbf{C}$ -morphism, then $\ker f$ defined by the pullback*

$$\begin{array}{ccc} \ker f & \xrightarrow{p_1} & A \\ p_2 \downarrow & & \downarrow f \\ A & \xrightarrow{f} & B \end{array}$$

is an equivalence relation on A .

Proof. This is essentially an exercise on pullbacks. We first need to show that $\ker f$ is a relation on A . To see this note first that by universality of the product, there exists a unique arrow $p_1 \times p_2 : \ker f \rightarrow A \times A$. To see that this arrow is monic, consider two morphisms $h, g : C \rightarrow \ker f$ with the property that

$$(p_1 \times p_2) \circ g = (p_1 \times p_2) \circ h$$

Since $f \circ p_1 = f \circ p_2$, it is clear that

$$f \circ p_1 \circ g = f \circ p_2 \circ g$$

and similarly for h , i.e. the arrows $p_1 \circ g, p_2 \circ g$ and $p_1 \circ h, p_2 \circ h$ make C a cone for the limit defining the pullback, thus there must exist a single map $C \rightarrow \ker f$ making the obvious diagram commute. But g and h are also two such maps, thus we must have $g = h$.

We now show $\ker f$ is an equivalence relation. For the reflexivity note that A with two copies of the identity map Id_A forms a cone of the diagram defining the pullback $\ker f$, and there must therefore exist a unique map $h : A \rightarrow \ker f$. Moreover, by the same argument as above, it is easy to see that the unique map $\Delta_A : A \rightarrow A \times A$ is monic, and since $p_1 \times p_2$ is also monic, we must have h monic, and by unicity $\Delta_A = (p_1 \times p_2) \circ h$, which is to say that $\Delta_A \subset (p_1 \times p_2)$ as subobjects of $A \times A$. Symmetry is immediate, since $\ker f$ together with the projection maps p_1, p_2 in the reverse order forms a cone which must be isomorphic to $\ker f$. Finally for the transitivity, if we build the pullback $\ker f \xleftarrow{p'_1} \ker f \times_A \ker f \xrightarrow{p'_2} \ker f$ of $\ker f \xrightarrow{p_1} A \xleftarrow{p_2} \ker f$, then it is clear that $p_1 \circ p'_1$ and $p_2 \circ p'_2$ define a cone for the diagram defining the pullback $\ker f$, and thus there must exist a unique mediating morphism $\ker f \times_A \ker f \rightarrow \ker f$ and by the same argument as above this morphism must be a mono, and in consequence $(p_1 \circ p'_1 \times p_2 \circ p'_2)$ is a subobject of $p_1 \times p_2$ are required. $\square$

In fact, in many categories of interest, and all categories of interest to us, the concepts of kernel pairs and equivalence relation coincide.

Definition 1.14. Let $\mathbf{C}$ be a regular category, an equivalence relation $R \rightrightarrows A \times A$ in $\mathbf{C}$ is called **effective** if it is the kernel pair of some morphism. A regular category where all equivalence relations are effective is called **exact**. Finally, a diagram of the shape

$$\ker q \begin{array}{c} \xrightarrow{e_1} \\ \xrightarrow{e_2} \end{array} B \xrightarrow{q} Q$$

where q is the coequalizer of its kernel pair, is called an **exact sequence**.

Lemma 1.15. *Let $F \dashv U : \mathbf{C} \rightarrow \mathbf{D}$ be an adjunction between regular categories, if $R \rightrightarrows A \times A$ is an equivalence relation in $\mathbf{D}$, then $UR \rightrightarrows U(A \times A) = UA \times UA$ is an equivalence relation in $\mathbf{C}$. In other words, right adjoints preserve equivalence relations.*

Proof. This is a simple consequence of the fact that U preserves monos and limits. $\square$

Proposition 1.16. *If $(\mathbf{C}, U)$ is a concrete category over $\mathbf{Set}$ and U is monadic, then U creates exact sequences and $\mathbf{C}$ is exact.*

Proof. Notice first that $\mathbf{Set}$ is exact: if $R \subset X \times X$ is an equivalence relation in $\mathbf{Set}$, consider the map $q : X \rightarrow X/R$, where X/R is the quotient of X under R . It is easy to see that q is the coequalizer of the component maps $e_1, e_2 : R \rightarrow A$ of R and that $R = \ker q$. Note that the diagram defined by e_1, e_2, q is an exact sequence.

Next, we will show that the coequalizer of an equivalence relation in $\mathbf{Set}$ is absolute, i.e. preserved by any functor. To see this we need to find a section s of q and a section t of e_1 such $e_2 \circ t = s \circ q$. Since every epi in $\mathbf{Set}$ has a section, we can always find a section s of q . Given such a section s , we build $t : X \rightarrow R$ as follows:

$$t = (\text{Id}_X \times (s \circ q)) \circ \Delta_X$$

This map is well defined since by definition of q , $(x, s \circ q(x)) \in R$ for any choice of s , and it is clearly a section of e_1 such that $e_2 \circ t = s \circ q$. It follows that q is an absolute coequalizer.

Finally, since U is monadic, by Beck's theorem, U creates absolute coequalizers. In particular this means that if we take any equivalence relation $e_1 \times e_2 : R \rightrightarrows A \times A$ in $\mathbf{C}$, then by Lemma 1.15, $Ue_1 \times Ue_2 : UR \subset UA \times UA$ is an equivalence relation in $\mathbf{Set}$, and is thus equal to the kernel pair of its coequalizer. Moreover, since this coequalizer is absolute and U is monadic, there must exist Q in $\mathbf{C}$ and $q : A \rightarrow Q$ such that q is the coequalizer of e_1, e_2 and


$Uq : UA \rightarrow UQ$ is the coequalizer of Ue_1, Ue_2 . Finally, since U is monadic, it also *creates* all limits, and in particular since UR is the kernel pair of Uq , R is the kernel pair of q . $\square$

This proposition tells us that **BA** and **DL** are exact categories, as is and $\mathbf{Alg}_{\mathbf{Set}}(T)$ whenever T is a variator. In what follows, we will be confronted with relations in $\mathbf{Alg}_{\mathbf{A}}(L)$ for an endofunctor $L : \mathbf{A} \rightarrow \mathbf{A}$ where $\mathbf{A}$ is **BA** or **DL**. In this situation it will no longer be possible to use the very nice, but exceptional, property of coequalizers of equivalence relation in **Set** to conclude that $\mathbf{Alg}_{\mathbf{A}}(L)$ is exact. We can however guarantee that $\mathbf{Alg}_{\mathbf{A}}(L)$ is exact for a large class of variators $L : \mathbf{A} \rightarrow \mathbf{A}$, as we will now show.

Lemma 1.17. *Let $\mathbf{C}$ be a regular category. Then*

- (i) *Every regular epi is the coequalizer of its kernel pair.*
- (ii) *Let $f : A \rightarrow B$ and let f^* be the regular epi part of its regular epi-mono factorization, then $\ker f \simeq \ker f^*$.*
- (iii) *If $\mathbf{C}$ is exact, then any equivalence relation is the kernel pair of its coequalizer.*

Proof. To show (i), let $p : A \twoheadrightarrow B$ be a regular epi, i.e. there exists $g, h : C \rightarrow A$ such that p is the coequalizer of g, h . Now consider its kernel pair $p_1, p_2 : \ker p \rightarrow A$. We then have



Since C together with g, h constitutes a cone for the pullback diagram defining $\ker p$, there exists a unique mediating morphism $k : C \rightarrow \ker p$. Now, let $p' : A \rightarrow B'$ be the coequalizer of p_1, p_2 . Clearly since $p' \circ p_1 = p' \circ p_2$, then $p' \circ p_1 \circ k = p' \circ g = p' \circ p_1 \circ h = p' \circ h$ and thus there must exist a unique morphism $u : B \rightarrow B'$ since p is the coequalizer of g, h . Conversely, since $p \circ p_1 = p \circ p_2$ by construction, we must have a unique arrow $u' : B' \rightarrow B$ since p' is the coequalizer of p_1, p_2 . Since u, u' must be inverse of each other we have $B \simeq B'$.

For (ii), it is easy to see than an arrow is mono iff its kernel pair is the identity. With this in mind, consider the following diagram:

$$\begin{array}{ccccc}
 \ker f & \xrightarrow{p_1} & A & & \\
 \downarrow p_2 & & \downarrow f^* & & \\
 & \text{im } f & \xrightarrow{\text{Id}_{\text{im } f}} & \text{im } f & \\
 & \downarrow \text{Id}_{\text{im } f} & & \downarrow & \\
 A & \xrightarrow{f^*} & \text{im } f & \xrightarrow{\quad} & B
 \end{array}$$

and it is easy to see that $p_1, p_2 : \ker f \rightrightarrows A$ is the kernel pair of f^* .

For (iii), take any equivalence relation $e_1 \times e_2 : R \rightarrow A \times A$. Since $\mathbf{C}$ is exact, we are guaranteed that $R = \ker f$ for some $f : A \rightarrow B$. Now, take the regular epi-mono factorisation of f , and it is easy to see from (ii) that $R = \ker f^*$ where f^* is the regular epi part of the factorisation.

$$\begin{array}{ccccc}
 R = \ker f & \xrightarrow{e_1} & A & & \\
 \downarrow e_2 & & \downarrow f^* & & \downarrow f \\
 A & \xrightarrow{f^*} & B^* & & B \\
 & \searrow f & \swarrow & & \\
 & & & & B
 \end{array}$$

Since f^* is regular epi, by (i) it is the coequalizer of its kernel pair, i.e. f^* is the coequalizer of e_1, e_2 and R is its kernel pair by construction. $\square$

Proposition 1.18 ([Hug01]). *Let $\mathbf{C}$ be a regular category and let $L : \mathbf{C} \rightarrow \mathbf{C}$ preserve regular epis, then $\mathbf{Alg}_{\mathbf{C}}(L)$ has regular epi-mono factorisation preserved and reflected by $\mathbf{U} : \mathbf{Alg}_{\mathbf{C}}(L) \rightarrow \mathbf{C}$.*

Proof. Let $f : (A, \alpha) \rightarrow (B, \beta)$. Since $\mathbf{C}$ is regular we have a regular epi-mono factorization for $Uf : A \rightarrow B$ as $f = i \circ f^*$. Since L preserves regular epis, Lf^* is a regular epi, and is in particular a strong epi. If we now consider the following diagram

$$\begin{array}{ccc}
 LA & \xrightarrow{Lf^*} & \text{Lim } f \\
 \downarrow f^* \circ \alpha & \swarrow \text{dashed} & \downarrow \beta \circ Li \\
 \text{im } f & \xrightarrow{i} & B
 \end{array}$$

which provides us with a structure map for the mediating L -algebra $\text{Lim } f \rightarrow \text{im } f$. This shows that $\mathbf{Alg}_{\mathbf{C}}(L)$ has regular-epi factorisation reflected by $\mathbf{U}$. That it is preserved by $\mathbf{U}$ is immediate. $\square$

Proposition 1.19. *Let $\mathbf{C}$ be a regular category and $L : \mathbf{C} \rightarrow \mathbf{C}$ be a variator preserving regular epis, then $\mathbf{Alg}_{\mathbf{C}}(L)$ is regular.*

Proof. We need to check that $\mathbf{Alg}_{\mathbf{C}}(L)$ is (i) finitely complete, (ii) has coequalizers of kernel pairs, and (iii) that regular epis in $\mathbf{Alg}_{\mathbf{C}}(L)$ are stable under pullback.

- (i) Since L is a variator, $U : \mathbf{Alg}_{\mathbf{C}}(L) \rightarrow \mathbf{C}$ is monadic, and in particular creates all limits which exists in $\mathbf{C}$. Since $\mathbf{C}$ is regular, it is finitely complete and thus so is $\mathbf{Alg}_{\mathbf{C}}(L)$.
- (ii) By Proposition 1.18 we know that $\mathbf{Alg}_{\mathbf{C}}(L)$ has regular epi-mono factorisation, and by Lemma 1.17 (ii), we know that the kernel pair of a morphism is isomorphic to the kernel pair of its regular epi factor. We can therefore assume without loss of generality that the morphism $q : (A, \alpha) \rightarrow (B, \beta)$ whose kernel pair we will consider is a regular epi. By (i) we know that we can build its kernel pair $p_1, p_2 : (\ker q, \kappa) \rightrightarrows (A, \alpha)$ in $\mathbf{Alg}_{\mathbf{C}}(L)$. Let us now check that q is the coequalizer of p_1, p_2 . For this consider another map $c : (A, \alpha) \rightarrow (C, \gamma)$ coequalizing p_1, p_2 (in $\mathbf{Alg}_{\mathbf{C}}(L)$). Since $\mathbf{C}$ is regular, q coequalizes its kernel pair p_1, p_2 , and thus there must exist a unique arrow $u : B \rightarrow C$ with $u \circ q = c$. We then have the following diagram:

$$\begin{array}{ccccc}
 L \ker q & \xrightarrow{Lp_1} & LA & \xrightarrow{Lq} & LB \\
 \vdots & & \downarrow \alpha & \searrow Lc & \downarrow \beta \\
 & & & & LC \\
 & & & \swarrow Lu & \\
 \ker q & \xrightarrow{p_1} & A & \xrightarrow{q} & B \\
 & \xrightarrow{p_2} & & & \\
 & & \downarrow c & & \downarrow u \\
 & & C & &
 \end{array}$$

All we now need to check is that u defines a morphism in $\mathbf{Alg}_{\mathbf{C}}(L)$. This follows from the fact that L preserves regular epis, and that Lq is therefore an epi. Indeed, we necessarily have

$$\begin{aligned}
 \gamma \circ Lc &= \gamma \circ Lu \circ Lq \\
 &= c \circ \alpha \\
 &= u \circ q \circ \alpha \\
 &= u \circ \beta \circ Lq
 \end{aligned}$$

and thus $\gamma \circ Lu = u \circ \beta$ as required, since Lq is epi.

- (iii) Finally, we need to show that regular epis are stable under pullback, so let $p : (A, \alpha) \twoheadrightarrow (B, \beta)$ be a regular epi and let $f : (C, \gamma) \rightarrow (B, \beta)$ be an arbitrary morphism. By (i), we know that the pullback of p along f in $\mathbf{Alg}_{\mathbf{C}}(L)$ is created by the pullback of p along f in $\mathbf{C}$, and since $\mathbf{C}$ is regular, we know that it must be a regular epi in $\mathbf{C}$. Let us call it $p' : (A \times_B C, \delta) \twoheadrightarrow (A, \alpha)$. By assumption on L , Lp' is also a regular epi, and it remains to show that p' and Lp' combine to form a regular epi in $\mathbf{Alg}_{\mathbf{C}}(L)$, i.e. a coequalizer of some morphisms. To this effect, let us consider the kernel pair $k_1, k_2 : (\ker p', \kappa) \rightrightarrows (A \times_B C, \delta)$ of p' . We now proceed exactly as in (ii) to show that p' is the coequalizer of k_1, k_2 in $\mathbf{Alg}_{\mathbf{C}}(L)$, and thus p' is a regular epi in $\mathbf{Alg}_{\mathbf{C}}(L)$.

□

Proposition 1.20. *Let $\mathbf{C}$ be an exact category and let $L : \mathbf{C} \rightarrow \mathbf{C}$ be a variator which preserves regular epimorphisms, then $U : \mathbf{Alg}_{\mathbf{C}}(L) \rightarrow \mathbf{C}$ creates exact sequences and $\mathbf{Alg}_{\mathbf{C}}(L)$ is exact.*

Proof. Showing that U creates exact sequences in $\mathbf{Alg}_{\mathbf{C}}(L)$ uses the same proof as (ii) of Proposition 1.19, i.e. we lift an exact sequence

$$\ker q \xrightleftharpoons[p_2]{p_1} A \xrightarrow{q} B$$

in $\mathbf{C}$ by first using the fact that U creates the kernel pair in $\mathbf{Alg}_{\mathbf{C}}(L)$ and then using the fact L preserves regular epi to show that $q : (A, \alpha) \rightarrow (B, \beta)$ is the coequalizer of p_1, p_2 .

$$\begin{array}{ccccc}
 L \ker q & \xrightleftharpoons[Lp_2]{Lp_1} & LA & \xrightarrow{Lq} & LB \\
 \downarrow \kappa & & \downarrow \alpha & \searrow Lc & \swarrow Lu \\
 & & & LC & \\
 & & & \downarrow \gamma & \\
 \ker q & \xrightleftharpoons[p_2]{p_1} & A & \xrightarrow{q} & B \\
 & & \downarrow c & \swarrow u & \\
 & & C & &
 \end{array}$$

which shows that any exact sequence in $\mathbf{C}$ can be lifted to an exact sequence in $\mathbf{Alg}_{\mathbf{C}}(L)$.

To see that $\mathbf{Alg}_{\mathbf{C}}(L)$ is exact, notice first that by Proposition 1.19 it is regular, so we need only show that all equivalence relations are effective. Let $e_1 \times e_2 : (R, \rho) \twoheadrightarrow (A \times A, (\alpha \circ L\pi_1) \times (\alpha \circ L\pi_2))$ be an equivalence relation in

$\mathbf{Alg}_{\mathbf{C}}(L)$. By Lemma 1.15, $U(R, \rho) = R$ is an equivalence relation in $\mathbf{C}$, and since $\mathbf{C}$ is exact, R is the kernel pair of a morphism. Since L is a variator, U is monadic and thus creates limits in $\mathbf{Alg}_{\mathbf{C}}(L)$, and in particular the kernel pair $(\ker q, \kappa)$ where $\kappa : L \ker q \rightarrow \ker q$ defines the unique structure map of $\ker q$ as an L -algebra. Finally, since U is monadic it reflects isomorphism (see [AHS04] Proposition 20.12), and since $U(R, \rho) = U(\ker q, \kappa)$, we get $(R, \rho) = (\ker q, \kappa)$ as desired. $\square$

1.4.4 Fully invariant equivalence relations

Whilst, equivalence relations are ubiquitous in the whole of mathematics, the stronger notion of *fully invariant* equivalence relation is key in universal algebra, and will be very important in the development of an algebraic semantics for the abstract flavour of coalgebraic logic in the next section. The notion of a fully invariant relation is touched upon in [Hug01] (we essentially use his definition 3.5.2.) but our construction of the fully invariant closure of a relation in a categorical setting is, to the best of our knowledge, new. We start by the basic definitions.

Definition 1.21. Let A be an object of some well-powered category $\mathbf{C}$. We say that a subobject $m : B \rightharpoonup A$ in $\text{Sub}(A)$ is **fully invariant** if for any endomorphism $f : A \rightarrow A$, there exist a unique endomorphism $\tilde{f} : B \rightarrow B$ such that

$$\begin{array}{ccc} B & \xrightarrow{m} & A \\ \tilde{f} \downarrow & & \downarrow f \\ B & \xrightarrow{m} & A \end{array}$$

commutes. In particular, if $e_1 \times e_2 : R \rightharpoonup A \times A$ is a relation on A , we say that R is a **fully invariant relation** if for any endomorphism $f : A \rightarrow A$, there exists a unique morphism $\tilde{f} : R \rightarrow R$ making the following diagram commute:

$$\begin{array}{ccc} R & \xrightarrow{e_1 \times e_2} & A \times A \\ \tilde{f} \downarrow & & \downarrow f \times f \\ R & \xrightarrow{e_1 \times e_2} & A \times A \end{array}$$

Let us give some examples of this interesting concept.

Example 1.22. 1. Let $\mathbf{C} = \mathbf{Grp}$ be the category of group. In this case, the concept of a fully invariant subgroup H of a group G is well known. In a cyclic group $\mathbb{Z}_p$, every subgroup H is fully invariant. Indeed, since H must itself be cyclic, $H = \mathbb{Z}_m$ with m a divisor of p , and H can be viewed as the powers of m . An endomorphism f of $\mathbb{Z}_p$ is uniquely determined by

$f(1) = n$, and thus $f[H]$ are the powers of mn which is included or equal to H itself. Another example of fully invariant subgroup is the commutator subgroup G' of any group G as is easily seen: if $a, b \in G$ then

$$f([a, b]) = f(aba^{-1}b^{-1}) = f(a)f(b)f(a)^{-1}f(b)^{-1} = [f(a), f(b)]$$

As an example of fully invariant relation, consider $R \subset G \times G$ defined by $(a, b) \in R$ if $b = a^{-1}$ for $a, b \in G$. If G is abelian, this **Set** relation is a subgroup of $G \times G$, and thus a relation in **Grp**. It is easy to see that it is fully invariant.

2. Let $\mathbf{C} = \mathbf{Vect}$. It is clear that no proper subvector space U of a vector space V can be fully invariant. Indeed, consider the endomorphism f projecting everything onto a one dimensional subspace which is orthogonal to U ; clearly $f[U] \not\subseteq U$.
3. Let A be an object of **DL** and let us define a relation R on A by $(a, b) \in R$ iff $a \wedge b = a$, for $a, b \in A$ (i.e. $R = \leq_A$). Then it is easy to check by applying the distributivity rule that R defines a relation in **DL**, and that this relation is fully invariant. This does not hold for boolean algebras, since the relation is not closed under negation.
4. The kind of example which will interest us most is as follows. Let $\mathbf{C} = \mathbf{Alg}_{\mathbf{Set}}(\mathbf{S}_{\mathbf{BA}})$ where $\mathbf{S}_{\mathbf{BA}} : \mathbf{Set} \rightarrow \mathbf{Set}$ is a polynomial functor whose signature can be used to define boolean algebras, for example $\{\perp, \neg, \wedge, \vee\}$. Consider $\mathbf{F}_{\mathbf{BA}}V$, the free $\mathbf{S}_{\mathbf{BA}}$ -algebra over a set V , and define a relation R on $\mathbf{F}_{\mathbf{BA}}V$ as aRb for $a, b \in \mathbf{F}_{\mathbf{BA}}V$ if $\mathbf{BA} \vdash_{\text{EL}} a = b$, where $\mathbf{BA} \vdash_{\text{EL}}$ means 'is derivable using equational logic and the axioms of boolean algebras'. It is easy to see that R is an equivalence relation in $\mathbf{Alg}_{\mathbf{Set}}(\mathbf{S}_{\mathbf{BA}})$. Moreover, by adjunction, any endomorphism $f : \mathbf{F}_{\mathbf{BA}}V \rightarrow \mathbf{F}_{\mathbf{BA}}V$ is equivalent to a map $\hat{f} : V \rightarrow \mathbf{U}_{\mathbf{BA}}\mathbf{F}_{\mathbf{BA}}V$ (where $\mathbf{U}_{\mathbf{BA}}$ is the obvious forgetful functor), i.e. f is just a substitution instance. It is thus clear from the rules of equational logic that R is a fully invariant.

Let us now show how to build the fully invariant closure of a relation. Strictly speaking we will need two constraints on the ambient category (apart from being well-powered): it must be cocomplete and locally small. However, since we are dealing with relations, and relations are well behaved in regular category, it is not unreasonable to place ourselves in such a category. In particular, well-poweredness and local smallness then coincide.

Proposition 1.23. *Let $\mathbf{C}$ be a well-powered cocomplete regular category, A be an object of $\mathbf{C}$ and $e_1 \times e_2 : R \rightrightarrows A \times A$ be a relation on R . Then the kernel*

pair $E = \ker q$ of the following coequalizer is the fully invariant closure of R :

$$\coprod_{f \in \text{hom}(A,A)} R \begin{array}{c} \xrightarrow{\coprod_{f \in \text{hom}(A,A)} f \circ e_1} \\ \xrightarrow{\coprod_{f \in \text{hom}(A,A)} f \circ e_2} \end{array} A \xrightarrow{q} \twoheadrightarrow Q$$

Proof. Clearly, this coequalizer and the coproducts defining it exist by the co-completeness assumption, and $E = \ker q$ is a relation on A (in fact an equivalence relation). To see that it is a fully invariant relation, consider any endomorphism $g : A \rightarrow A$, and the diagram

$$\begin{array}{ccccc} \coprod_{f \in \text{hom}(A,A)} R & \begin{array}{c} \xrightarrow{\coprod_{f \in \text{hom}(A,A)} f \circ e_1} \\ \xrightarrow{\coprod_{f \in \text{hom}(A,A)} f \circ e_2} \end{array} & A & \xrightarrow{q} & Q \\ & & \downarrow g & & \downarrow u \\ \coprod_{f \in \text{hom}(A,A)} R & \begin{array}{c} \xrightarrow{\coprod_{f \in \text{hom}(A,A)} f \circ e_1} \\ \xrightarrow{\coprod_{f \in \text{hom}(A,A)} f \circ e_2} \end{array} & A & \xrightarrow{q} & Q \end{array}$$

We would like to show that there exists a unique arrow $u : Q \rightarrow Q$ making the diagram commute. We do this by showing that $q \circ g$ coequalizes the parallel pair of the top row. To see this, note that, by definition, q coequalizes any pair $f \circ e_1, f \circ e_2$ for $f \in \text{hom}(A, A)$. In particular it must coequalize all such pairs factoring through g , i.e. all the pairs of the type $g \circ f \circ e_1, g \circ f \circ e_2$, and thus q coequalizes $g \circ \coprod_{f \in \text{hom}(A,A)} f \circ e_i, i = 1, 2$. So $u : Q \rightarrow Q$ exists and is unique, which in turn implies that there must exist a unique $v : E \rightarrow E$ making the following diagram commute (since the top row of the diagram can be seen as a cone for the diagram defining $E = \ker q$).

$$\begin{array}{ccccc} E & \xrightarrow{p_1} & A & \xrightarrow{q} & Q \\ \downarrow v & \downarrow p_2 & \downarrow g & & \downarrow u \\ E & \xrightarrow{p_1} & A & \xrightarrow{q} & Q \\ & & \downarrow p_2 & & \end{array}$$

This shows that E is fully invariant.

We can easily see that E contains R (i.e. that there exists a mono from R to E). Note first that since $q \circ e_1 = q \circ e_2$ (since $\text{Id}_A \in \text{hom}(A, A)$), R is a cone for the diagram defining E , so there must exist a unique morphism $u : R \rightarrow E$. Moreover, this morphism must satisfy $p_1 \times p_2 \circ u = e_1 \times e_2$, which means that u must be a mono since $e_1 \times e_2$ is a mono. To see that E is the smallest such

fully invariant relation, simply notice that the coequalizer q' of any other fully invariant relation E' containing R would have to coequalize the pair $f \circ e_1, f \circ e_2$ for any $f \in \text{hom}(A, A)$ as is clear from the following diagram

$$\begin{array}{ccccc}
 R & & & & \\
 \downarrow & \searrow^{e_1 \times e_2} & & & \\
 E' & \xrightarrow{p'_1 \times p'_2} & A \times A & & \\
 \downarrow & & \downarrow f \times f & & \\
 E' & \xrightarrow{p'_1 \times p'_2} & A \times A & \xrightarrow{q'} & Q'
 \end{array}$$

where p'_1, p'_2 are the projections of the from the relation E' to A . In consequence, there must exist a unique morphism $Q \rightarrow Q'$, and thus a unique morphism $u : E = \ker q \rightarrow E' = \ker q'$ such that $(p'_1 \times p'_2) \circ u = p_1 \times p_2$. This morphism u is a mono since $p_1 \times p_2$ is a mono (Lemma 1.13). $\square$

1.5 Coalgebraic logics

We introduce coalgebraic logics by following the pedagogically helpful trichotomy of Kupke and Pattinson's review paper [KP11], namely that coalgebraic logics come in three flavours: the **predicate lifting**, **nabla**, and **abstract** flavour. The latter unifying the former two. The existing literature on coalgebraic logic is overwhelmingly boolean (see [JS10, KKV12a, BKV13] for some exceptions). The predicate lifting and nabla flavours of coalgebraic logic in particular can, for all intents and purposes, be considered to be boolean logics. The abstract flavour of coalgebraic logic however, explicitly parametrizes the choice of reasoning structure, and indeed [JS10] gives several interesting example of non-boolean coalgebraic logics. We will therefore proceed as follows: the predicate lifting and nabla styles will be presented as strictly boolean, whereas the abstract style will be presented in its full generality, i.e. parametric in a choice of reasoning structure.

1.5.1 Three styles of language

We start with the **predicate lifting** style of coalgebraic logic. Given a finitary signature (Σ, ar) , $\text{ar} : \Sigma \rightarrow \mathbb{N}$ and a set V of propositional variables, the **predicate lifting language** $\mathcal{L}_\Sigma(V)$ (or simply $\mathcal{L}_\Sigma$ if there is no ambiguity) has a syntax given by

$$a ::= p \mid \perp \mid \neg a \mid a \wedge a \mid \sigma(a, \dots, a)$$

where $p \in V$, and $\sigma \in \Sigma$. Note that we are using the notational convention of [KKV12b] where the lower case Roman letters a, b, c stand for formulae.

Algebraically, if we denote by Σ_{BA} the signature we've chosen for the propositional (i.e. boolean) part of the language (in this instance $\Sigma_{BA} = \{\perp, \neg, \wedge\}$ with the usual arities), then the language $\mathbf{L}_\Sigma$ is the free algebra over V for the **Set**-endofunctor $S_{\Sigma_{BA}+\Sigma}$ defined as

$$S_{\Sigma_{BA}+\Sigma}X = \coprod_{\sigma \in \Sigma_{BA}+\Sigma} X^{\text{ar}(\sigma)}$$

We know that $\mathbf{L}_\Sigma$ exists by Corollary 1.4 and the fact that $S_{\Sigma_{BA}+\Sigma}$ is a finitary polynomial functor.

The **∇ -style** of coalgebraic logic (also known as Moss style, or coalgebraic logic for the cover modality) has a very different flavour. We refer to [KKV12b] for a very good and very thorough overview of this quirky but very powerful logic. Since the language involves objects of many types, the notational conventions are crucial to avoid any confusion. We will scrupulously follow the conventions of [KKV12b]. We start by fixing a weak-pullback preserving and standard² **Set**-endofunctor T which will be used both to build the syntax of the logic and to define its coalgebraic semantics. We define $T_\omega X = \bigcup\{TY \mid Y \subseteq X, Y \text{ finite}\}$, the **finitary version** of T (more on finitary functors in Chapter 3 and on the finitary version in Chapter 5). The **nabla language** induced by T is written as $\mathcal{L}_T$ and is given by:

$$a ::= p \mid \neg a \mid \bigwedge \phi \mid \bigvee \phi \mid \nabla \alpha$$

where $p \in V$, $\phi \in \mathcal{P}_\omega \mathcal{L}_T$ and $\alpha \in T_\omega \mathcal{L}_T$. Elements of type $\mathcal{P}_\omega \mathcal{L}_T$ will always be denoted by $\phi, \psi, \dots$ whilst elements of type $T_\omega \mathcal{L}_T$ will always be denoted by $\alpha, \beta, \dots$. Note that $\bigwedge$ and $\bigvee$ are operators of type $\mathcal{P}_\omega \mathcal{L}_T \rightarrow \mathcal{L}_T$ and of course corresponds to the usual meet and join operations, the constant $\perp$ is defined as $\bigvee \emptyset$ and $\top$ as $\bigwedge \emptyset$. Once again, we can view the language as a free algebra over V for a **Set**-endofunctor, this time we use $L_T = \text{Id} + \mathcal{P}_\omega + \mathcal{P}_\omega + T_\omega$, and $\mathcal{L}_T$ is the free L_T -algebra over V . It is clear that since L_T is a coproduct of finitary functors, it is itself finitary, and thus $\mathcal{L}_T$ exists by Corollary 1.4. As is customary, the semantics of $\mathcal{L}_T$ will be defined inductively, which means that we need to define what subformulas are in this rather unusual language, which is not immediately obvious for ∇ -terms. For this we need the notion of **base** to which we will return at the beginning of Chapter 3. For any $\alpha \in T_\omega \mathcal{L}_T$ we define the base of α by $B_T(\alpha) = \bigcap\{U \subseteq \mathcal{L} \mid \alpha \in TU\}$. The base of α is the set of immediate subformulas of $\nabla \alpha$. We can now define the set $Sfml(a)$ of subformulas of a formula a recursively in the obvious way for formulas whose

²**Standard** means that the functor preserves monomorphisms. As shown in [Trn69] and [Gum05], for the purpose of studying coalgebras defined by a **Set**-endofunctor T we can always assume that T preserves monomorphisms.

outermost operators are boolean connectives, and by

$$Sfml(\nabla\alpha) = \{\nabla\alpha\} \cup \bigcup_{a \in \mathbf{B}_T(\alpha)} Sfml(a)$$

for ∇ -terms.

Finally, the **abstract style** of coalgebraic logic is both a generalization and a unification of the previous two styles. One of its key features is that it hides the fundamental structure of the formulas by defining formulas as elements of a free algebra for an endofunctor defined not on **Set**, but on a choice of reasoning structure, typically **BA** or **DL**. In other words all the terms are understood to be elements of an algebraic structure, and equivalence in this structure is hidden away. The basic setup, which can be traced back to [KKP04] and was subsequently developed in [KKP05] and [JS10], is as follows:

$$\begin{array}{ccc}
 & L & \\
 & \downarrow & \\
 \mathbf{C} & \xrightarrow{F} & \mathbf{D}^{\text{op}} \\
 & \uparrow & \\
 & T^{\text{op}} & \\
 \uparrow & \perp & \downarrow \\
 \mathbf{F} & \dashv & \mathbf{U} \\
 \downarrow & & \downarrow \\
 \mathbf{Set} & & \mathbf{Set}
 \end{array}
 \quad (1.2)$$

The left hand-side of the diagram is the syntactic side, and the right-hand side the semantic one. The category $\mathbf{C}$ represents a choice of ‘reasoning kernel’, i.e. of logical operations which we consider to be fundamental, whilst L is a syntax constructing functor which builds terms over the reasoning kernel. The category $\mathbf{C}$ is a finitary variety and there exist and free-forgetful adjunction $\mathbf{F} \dashv \mathbf{U} : \mathbf{C} \rightarrow \mathbf{Set}$ which is $\mathbf{U}$ is monadic and creates directed colimits. On the semantics side, objects in $\mathbf{D}$ are the carriers of models and T specifies the coalgebras on these carriers in which the operations defined by L will be interpreted. The functors F and G (where $F \dashv G$) relate the syntax and the semantics. Note that we only need a dual adjunction, not a full duality.

A typical example of this situation is the adjunction $\mathbf{Pf} \dashv \mathbf{U} : \mathbf{DL} \rightarrow \mathbf{Pos}^{\text{op}}$, where $\mathbf{Pf}$ is the functor sending a distributive lattice to its poset of prime filters, and $\mathbf{DL}$ -morphisms to their inverse images, and $\mathbf{U}$ is the functor sending a poset to the distributive lattice of its up-sets and monotone maps to their inverse images. In the case where a distributive lattice is a boolean algebra, it is well-known that prime filters are maximal, i.e. ultrafilters, and the partial order on the set of ultrafilter is thus discrete, i.e. ultrafilters are only related to themselves. Downsets are thus simply subsets and the adjunction $\mathbf{Pf} \dashv \mathbf{U}$ becomes the well-known adjunction $\mathbf{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}^{\text{op}}$ where $\mathcal{P}$ is the contravariant powerset functor and powersets are endowed with the usual boolean structure. More situations of the type of Diagram (1.2) can be found in [JS10].

Let us for the moment focus only on the syntax builder L . It defines an **abstract coalgebraic language** which is defined for a set of propositional

variables V as the free L -algebra over FV . We therefore require that L be a variator. If we denote by $\mathbf{G} \dashv \mathbf{V}$ the adjunction $\mathbf{C} \rightarrow \mathbf{Alg}_{\mathbf{C}}(L)$, then the abstract language is given by $(\mathbf{GFV}, \langle - \rangle)$, where $\langle - \rangle$ denotes the structure map of the free algebra. Note that we can equally well consider the language as the initial algebra of the functor $L' = L(-) + FV$. Indeed, we can easily put an L' -algebra structure on $\mathbf{GFV}$ via the map $\langle - \rangle + \eta_{FV}^{\mathbf{G}}$, where $\eta^{\mathbf{G}}$ is the unit of $\mathbf{G} \dashv \mathbf{V}$. So the two descriptions are equivalent, and being able to use either will prove very useful indeed.

We can once again use the notion of *base* to describe the subformulas of a formula $\langle \alpha \rangle \in \mathbf{GFV}$, $\alpha \in L\mathbf{GFV}$. We use the same notational convention for abstract coalgebraic languages as for nabla languages. Some care must be taken, since we are no longer working in **Set**, however as we shall see in Chapter 3, the notion of base is well-behaved when L preserves all intersections. We can then define subformulas inductively, just as in the case of the nabla logic, with the usual clauses for propositional connectives, and with

$$Sfml(\langle \alpha \rangle) = \{\langle \alpha \rangle\} \cup \bigcup_{a \in B_L(\alpha)} Sfml(a)$$

This abstract language covers the previous two styles of languages in the following sense. Let us first look at the predicate lifting language $\mathbf{L}_{\Sigma}$ for a signature Σ . We define $L : \mathbf{BA} \rightarrow \mathbf{BA}$ as

$$L = \mathbf{F} \left(\prod_{\sigma \in \Sigma} (\mathbf{U}(-))^{\text{ar}(\sigma)} \right) = \mathbf{F} \mathbf{S}_{\Sigma} \mathbf{U}$$

In other words we use the adjunction $\mathbf{F} \dashv \mathbf{U}$ to lift $\mathbf{S}_{\Sigma}$ from **Set** $\rightarrow$ **Set** to $\mathbf{BA} \rightarrow \mathbf{BA}$. We can then see that

$$\mathbf{GFV} = \mathbf{L}_{\Sigma} / \equiv$$

where $a \equiv b$ if a and b are provably equivalent using equational logic and the axioms of boolean algebras. In other words, the abstract language for L is just the predicate lifting language quotiented by boolean equivalence. Note that $\mathbf{GFV}$ exists by Proposition 1.5.

Similarly, we can define an abstract language corresponding to the quotienting of the nabla language under boolean equivalence. For this we define $L = \mathbf{F} \mathbf{T} \mathbf{U}$ and we then have

$$\mathbf{GFV} = \mathcal{L}_T / \equiv$$

Again, we are in the presence of a functor in $\mathbf{BA}$ which is the ‘lifted’ version of a functor in **Set**. $\mathbf{GFV}$ exists when L is a variator, which is for example the case when T is finitary (since $\mathbf{F}$ is a left adjoint it preserves all colimits, and in particular filtered ones).

A functor L , defining an abstract coalgebraic language, does not have to be of the form FTU for some **Set** endofunctor T . It can be much more general, and encode much more than simply quotienting term building in **Set** under boolean equivalence. For example it can encode some extra axioms which we would like to take as standard for all the logics we consider. For example, we might want to encode the axiom **K** in any functor L describing classical modal logics. As a rule of thumb, everything we want as standard, should be encoded in L , whilst formulas we want to experiment with should be kept out of L . Alternatively, we can use the rule of thumb that L should contain all the axioms necessary to axiomatize the class of all coalgebras for a certain functor of interest, whilst extra axioms, kept out of the functor L , can be used to describe proper classes (covarieties) of such coalgebras. Time to move on to the semantic side.

1.5.2 Three styles of semantics

All three types of coalgebraic languages are, not surprisingly, interpreted in coalgebras. Given a standard **Set**-endofunctor T , a **coalgebra** is a pair (W, γ) where W is a set (of worlds) and $\gamma : W \rightarrow TW$ is a transition map (T defines the ‘transition type’).

We start with the predicate lifting style of coalgebraic logic which owes its name to the device used to interpret its operators. Let us fix once again a signature (Σ, ar) and let L_Σ be the associated language as defined above. Each modal operator $\sigma \in \Sigma$ is interpreted by a **predicate lifting**, i.e. a natural transformation $\llbracket \sigma \rrbracket : Q^n \rightarrow QT$ where $Q : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$ is the contravariant powerset functor. Intuitively, predicate liftings ‘lift’ n -tuples of predicates (i.e. subsets, hence the powerset functor Q) to a predicate on transitions (hence QT). A **T -coalgebraic model** - or T -model (or just model if there is no ambiguity about T) - is a triple $\mathcal{M} = (W, \gamma, \pi)$ where $\pi : W \rightarrow \mathcal{P}(V)$ is a valuation. The notion of **truth** of a formula a at a point $w \in W$ is defined inductively in the usual manner for propositional variables and boolean operators, and by

$$\mathcal{M}, w \models \sigma(a_1, \dots, a_n) \text{ iff } \gamma(w) \in \llbracket \sigma \rrbracket_W(\llbracket a_1 \rrbracket, \dots, \llbracket a_n \rrbracket)$$

for modal operators, where $\llbracket a_i \rrbracket$ is the interpretation of a_i in W . A formula a is **satisfiable** in $\mathcal{M}$ if there exists $w \in W$ such that $\mathcal{M}, w \models a$. A **coalgebraic frame** - or T -frame - is just a T -coalgebra (W, γ) and a formula a is **valid** on the frame if for any valuation π , a is true at every point in the model (W, γ, π) .

Let us now turn to the interpretation of the nabla style of coalgebraic logic. As we mentioned earlier, the same functor is used both for the syntax and the semantics. So let T be a weak pullback preserving **Set**-endofunctor, and let $\mathcal{L}_T$ be the language defined by T as above. In order to define the semantics of nabla formulas we need the notion of relation lifting, developed in the previous section for general regular categories, in **Set**. Note that the ubiquitous requirement that

a functor should preserve regular epis in order for liftings to be well-behaved is automatically satisfied in **Set** since every epi has a section (assuming the axiom of choice). Let us briefly specialize the definition of relation lifting from the previous section to the category **Set**. Let X_1, X_2 be sets and let $R \subseteq X_1 \times X_2$ be a relation between X_1 and X_2 . By construction, the lifting $\bar{TR}$ of R is $\bar{TR} = (T\pi_1 \times T\pi_2)[TR]$. Alternatively and equivalently, $\bar{TR}$ is the relation between TX_1 and TX_2 defined by

$$\bar{TR} = \{(\alpha, \beta) \in TX_1 \times TX_2 \mid \exists z \in TR \text{ s.th. } T\pi_1(z) = \alpha, T\pi_2(z) = \beta\}$$

where $\pi_i : R \rightarrow X_i, i = 1, 2$ is the projection onto the i^{th} component. If T preserves weak-pullbacks (regular epis is automatic), then T -liftings satisfy Propositions 1.11 and 1.12 and we can then define the semantics of $\mathcal{L}_T$. Given a T -model $\mathcal{M} = (W, \gamma, \pi)$, we define the validity relation $\models \subseteq W \times \mathcal{L}_T$ inductively for any $a \in \mathcal{L}_T$ and world $w \in W$ as usual for atomic propositions and propositional connectives, together with

$$\mathcal{M}, w \models \nabla \alpha \text{ iff } (\gamma(w), \alpha) \in \bar{T} \models$$

where $(\bar{T} \models) \subseteq TW \times T\mathcal{L}_T$ is the relation lifting of the truth-relation $\models \subseteq W \times \mathcal{L}_T$. Note that for this definition to indeed be inductive, it must be the case that we only need to know $\models$ on the subformulas of $\nabla \alpha$, i.e. on the base of α . This is where Propositions 1.11 and 1.12 come into play: since we're assuming that T preserves weak pullbacks we have

$$\begin{aligned} (\gamma(w), \alpha) \in \bar{T} \models & \text{ iff } (\gamma(w), \alpha) \in \bar{T} \models \uparrow (TW \times TB_T(\alpha)) \\ & \text{ iff } (\gamma(w), \alpha) \in \bar{T}(\models \uparrow (W \times B_T(\alpha))) \end{aligned}$$

i.e. we only need to know how $\models$ relates states to subformulas of $\nabla \alpha$, which is what was required for an inductive definition.

Finally, let us describe the semantics of the abstract version of coalgebraic logic. Again, this semantics generalises the two types of coalgebraic semantics we have defined above. Recall that the fundamental situation of Diagram (1.2)

$$\begin{array}{ccc} & & T^{\text{op}} \\ & \curvearrowright & \\ L & \begin{array}{c} \text{C} \end{array} & \begin{array}{c} \text{D}^{\text{op}} \end{array} \\ & \begin{array}{c} \text{F} \\ \perp \\ G \end{array} & \\ & \curvearrowleft & \end{array}$$

where L describes the syntax, and T describes the semantic domain, i.e. T -coalgebras. Intuitively, the semantic should be something which associates to an L -formula, a 'predicate' on the carrier of a T -frame (W, γ) , namely the denotation of the formula. In the abstract setting 'predicates' are given by the functor G . Indeed in practise the elements of GW can often be understood as

morphisms into a two-element structure $\mathbf{2}$. Down-sets on a poset W for example, are equivalent to monotone maps $W \rightarrow \mathbf{2}$, where $\mathbf{2}$ is the obvious two-element poset, and similarly subsets are equivalent to their characteristic function.

Since the semantic is given inductively, we can assume that to interpret an L -formula we need to interpret an L -construct of subformulas which have already been interpreted, i.e. an element of the type LGW . In this perspective it therefore natural to define an abstract semantic as a device which translates L -constructs of interpreted sub-formulas to their denotation, i.e. a natural transformation of the type

$$\delta : LG \rightarrow GT$$

We will call such a natural transformation a **semantic transformation**. Recall that the language is the free L -algebra over FV . Its carrier is equivalently given by the initial algebra $\mu L'$ where $L' = L(-) + FV$. The semantic natural transformation δ yields an interpretation maps $\llbracket - \rrbracket_W : \mu L' \rightarrow GW$ as follows. Given a coalgebra (W, γ) , if we can endow GW with an L' -algebra structure, then we get a unique interpretation map $\llbracket - \rrbracket_W : \mu L' \rightarrow GW$ by initiality. To define an L' -algebra structure on GW , we need δ and a map $FV \rightarrow GW$, or equivalently, a function $v : V \rightarrow UGW$. Such functions are abstract versions of the familiar valuations as is clear by considering the case where $F = \mathbf{Uf}$, $G = \mathcal{P}$. Given such a valuation, we get:

$$\begin{array}{ccc} L\mu L' + FV & \xrightarrow{L\llbracket - \rrbracket_W + \text{Id}_{FV}} & L\mathcal{P}W + FV \\ \downarrow & & \downarrow \delta_W + \text{Id}_{FV} \\ & & \mathcal{P}TW + FV \\ & & \downarrow \mathcal{P}\gamma + \hat{v} \\ \mu L' & \xrightarrow{\llbracket - \rrbracket_W} & \mathcal{P}W \end{array}$$

where $\hat{v}$ is the adjoint transpose of the valuation. The natural transformation δ inserted in this way does precisely what we had in mind, namely inductively define the semantics of formulas in $\mathbf{GFV}$.

Let us briefly show how this type of semantics covers the two cases above. In the case of the predicate lifting semantics interpreted over T -coalgebras, recall that we had defined L as $\mathbf{FS}_\Sigma \mathbf{U}$ where $\mathbf{S}_\Sigma$ is the polynomial functor on $\mathbf{Set}$ induced by the signature Σ . The abstract definition of the semantics requires a natural transformation

$$\delta : L\mathcal{P} = \mathbf{FS}_\Sigma \mathbf{U}\mathcal{P} = \mathbf{FS}_\Sigma \mathcal{Q} \rightarrow \mathcal{P}T$$

which, by adjunction, is equivalent to a natural transformation

$$\tilde{\delta} : \mathbf{S}_\Sigma \mathcal{Q} \rightarrow \mathbf{U}\mathcal{P}T = \mathcal{Q}T$$

which, in turn, is equivalent to a collection of natural transformations

$$\llbracket \sigma \rrbracket : \mathcal{Q}^{\text{ar}(\sigma)} \rightarrow \mathcal{Q}T$$

for each $\sigma \in \Sigma$. In other words, we recover precisely the concept of predicate lifting. The predicate lifting style of semantics is thus simply the adjoint transpose of the abstract style of semantics for ‘polynomial syntax’ functors.

Let us now show that the abstract semantics also subsumes the nabla-style of semantics. Recall first that in this case we had put $L : FTU$. Again we write $L' = FTU(-) + FV$, and recall that the interpretation maps should be given by:

$$\begin{array}{ccc} FTU\mu L' + FV & \xrightarrow{L[-]_w + \text{Id}_{FV}} & FT\mathcal{Q}W + FV \\ \downarrow & & \downarrow \delta_W + \text{Id}_{FV} \\ & & \mathcal{P}TW + FV \\ & & \downarrow \mathcal{P}\gamma + \hat{v} \\ \mu L' & \xrightarrow{\llbracket - \rrbracket_w} & \mathcal{P}W \end{array}$$

for a valuation $v : V \rightarrow \mathcal{Q}W$ and a well-chosen semantic transformation δ . By using the adjoint transpose maps where required we can re-write the diagram above as

$$\begin{array}{ccc} TU\mu L' + V & \xrightarrow{T[-]_w + \text{Id}_V} & T\mathcal{Q}W + V \\ \downarrow & & \downarrow \tilde{\delta}_W + \text{Id}_V \\ & & \mathcal{Q}TW + V \\ & & \downarrow \mathcal{Q}\gamma + v \\ U\mu L & \xrightarrow{\llbracket - \rrbracket_w} & \mathcal{Q}W \end{array}$$

where we have used the same notation for the interpretation maps and its adjoint transpose to clarify the notation. Notice now that for a purely propositional formula $p \in \mathcal{L}_T$ we have

$$w \models p \text{ iff } w \in v(p) \text{ iff } (w, p) \in (\in; \llbracket - \rrbracket_w^{\text{op}})$$

For a formula $\nabla\alpha \in \mathcal{L}_T$ of modal depth equals to one, we can use this equivalence and the nabla-style of semantics to conclude that

$$\begin{aligned} w \models \nabla\alpha &\Leftrightarrow (\gamma(w), \alpha) \in \bar{T}(\models \upharpoonright (W \times \mathbf{B}_T(\alpha))) \\ &\Leftrightarrow (\gamma(w), \alpha) \in \bar{T}(\in_W; \llbracket - \rrbracket_w^{\text{op}} \upharpoonright (W \times \mathbf{B}_T(\alpha))) && \text{Equiv. above} \\ &\Leftrightarrow (\gamma(w), \alpha) \in \bar{T} \in_W; (T[-]_w)^{\text{op}} \upharpoonright TW \times T\mathbf{B}_T(\alpha) && \text{Prop. 1.12} \\ &\Leftrightarrow \gamma(w) \bar{T} \in_W T\llbracket \alpha \rrbracket_w \end{aligned}$$

The argument can then be repeated for formulas of any modal depth. It therefore follows that if we define $\tilde{\delta}_W$ by

$$\tilde{\delta}_W : TQW \rightarrow QTW, A \mapsto \{t \in TW \mid t \bar{T} \in_W A\}$$

the diagram above will commute, and we thus have shown that the abstract style of semantics also covers the nabla case.

1.5.3 Two styles of axiomatization

One purpose of coalgebraic logics is to logically describe classes of coalgebras. In order to achieve this some axiomatization is needed. There are essentially two types of axiomatization: using *rules* or using *axioms*. For the predicate lifting style, either style can be used, and the choice is largely a matter of taste and convenience. The nabla flavour of coalgebraic logic is again rather different, and is arguably the most elegant type of coalgebraic logic in that it possesses a generic, rule-based, axiomatization which is fully parametric in the choice of functor defining the coalgebraic semantics. Finally, the abstract style of coalgebraic logic is rarely used in practise to axiomatize classes of coalgebras but could be given either an axiom based or a rule based axiomatization, although we will opt for axioms only.

Rule and axiom-based axiomatizations in the predicate lifting style

Let us once again start by examining the case of predicate-lifting style coalgebraic logic. Given a choice of semantics - i.e. the choice of a **Set**-endofunctor T - it has been shown in [Pat03, Sch06] that it is always possible to find an axiomatization of the class **CoAlg**(T) of all T -coalgebra in a Hilbert system using either rules or axioms. Moreover, the rules or axioms can be assumed to have a particularly simple form. We cite the key definitions and results, for more details see [Pat03, Sch06, KP11].

Definition 1.24. Let (Σ, ar) be a finitary signature, let V be a set of propositional variables and let L_Σ be the associated predicate-lifting style language. For any subset $U \subset L_\Sigma$ we define $\text{Up}(U)$ as

$$\text{Up}(U) = \{\sigma(a_1, \dots, a_n) \mid \sigma \in \Sigma, \text{ar}(\sigma) = n, a_i \in U, 1 \leq i \leq n\}$$

For any subset U of L_Σ we define $\text{Cl}(U)$ to be the set of clauses over U , i.e. the set of all formulas of the form $\bigvee_{i=1}^n \varepsilon_i b_i$ where ε_i is either nothing or $\neg$ and $b_i \in U$, $1 \leq i \leq n$.

A **one step rule** a/b is a pair $(a, b) \in FV \times \text{Cl}(\text{Up}(V))$. A **rank 1 clause** is a formula $a \in \text{Cl}(\text{Up}(FV))$.

Given a collection R of rules, we define a very simple Hilbert deduction system whose axioms are given by some choice of axioms of propositional logic (considered as rules without premise) together with the rules of R , and whose rules of inference are modus ponens, uniform substitution, and the **congruence rule**

$$(\text{Congruence}) \frac{a_i \leftrightarrow b_i, 1 \leq i \leq n}{\sigma(a_1, \dots, a_n) \leftrightarrow \sigma(b_1, \dots, b_n)} \quad (\sigma \in \Sigma)$$

This deduction system defines a provability predicate $R \vdash_{\text{ML}}$ on L_Σ , and, as usual, we define the **logic** L_R as the set of formulas $a \in L_\Sigma$ such that $R \vdash_{\text{ML}} a$. We use the notation $\vdash_{\text{ML}}$ because the Hilbert system is essentially that of modal logic, the only difference being that the necessitation rule is replaced by the congruence rule.

Similarly, given a set of axioms $\text{Ax} \subset L_\Sigma$, we define the Hilbert deduction system whose axioms are given by some choice of axioms of propositional logic together with the axioms of Ax , and whose inference rules are modus ponens, uniform substitution, and congruence. We denote the associated provability predicate on L_Σ as $\text{Ax} \vdash_{\text{ML}}$ and define the logic L_{Ax} as all the formulas $a \in L_\Sigma$ such that $\text{Ax} \vdash_{\text{ML}} a$.

The main technical tool in the study of completeness of such Hilbert systems with respect to a choice of coalgebraic semantics, are the notions of *one step soundness* and *one step completeness*.

Definition 1.25. Given a predicate lifting style language L_Σ and an interpretation of this language in the class of T -coalgebra, a one step rule a/b where $b = \bigvee_{i=1}^n \varepsilon_i b_i$ and each $b_i = \sigma_i(c_1^i, \dots, c_{k_i}^i)$ with $c_1^i, \dots, c_{k_i}^i \in V$, is **one step sound** if for any T -model $\mathcal{M} = (W, \gamma, \pi)$, whenever $\mathcal{M} \models a$

$$\bigcup_{i=1}^n \varepsilon_i \llbracket \sigma_i \rrbracket (\llbracket c_1^i \rrbracket, \dots, \llbracket c_{k_i}^i \rrbracket) = TW$$

where ε_i here stands for nothing or complementation. In other words, if a is true in the entire model, then any transition must make b true too.

Conversely, a rule a one step rule a/b is **one step complete** if for every T -model (W, γ, π) and every formula $d \in \text{Cl}(\text{Up}(V))$, $d = \bigvee_{i=1}^n \varepsilon_i \sigma_i(c_1^i, \dots, c_{k_i}^i)$,

$$\bigcup_{i=1}^n \varepsilon_i \llbracket \sigma_i \rrbracket (\llbracket c_1^i \rrbracket, \dots, \llbracket c_{k_i}^i \rrbracket) = TW$$

implies that

$$\{\sigma b \mid \sigma : \text{FV} \rightarrow \text{FV}, (W, \gamma) \models \sigma a\} \vdash_{\text{PL}} d$$

where $\vdash_{\text{PL}}$ means ‘is derivable using propositional calculus only’. In other words, if a formula in $\text{Cl}(\text{Up}(V))$ (the format of the conclusion of one step rules) is true for

any transition of the frame (W, γ) , then it must be derivable, using propositional logic, from the substitution instances of the conclusion of the rule which make the premise valid on (W, γ) . A rank 1 clause is one-step complete if it is one step complete when viewed as a one-step rule with empty premise.

Theorem 1.26 ([Sch06]). *Given a choice T of semantics, and a predicate lifting style language $\mathcal{L}_\Sigma$, the set of all one-step sound one-step rules is complete with respect to $\mathbf{CoAlg}(T)$.*

As was shown in Proposition 15 of [Sch06], for any one-step rule a/b , there exists a rank 1 clause c which is equivalent to a/b in the sense that $\{a, c\} \vdash_{\text{ML}} b$ iff $\{a/b\} \vdash_{\text{ML}} c$. We therefore have the corollary:

Corollary 1.27 ([Sch06]). *Given a choice T of semantics, and a predicate lifting style language $\mathcal{L}_\Sigma$, the set of all one-step sound rank-1 clauses is complete with respect to $\mathbf{CoAlg}(T)$.*

Note that this Theorem and its Corollary, only shows that complete axiomatizations exists in principle. Of course in practise, there exist more compact axiomatizations of many well-known endofunctor, and we refer the reader to e.g. Example 3.17 of [SP06] for a list of such Hilbert systems (typically given in terms of one-step rules).

The rule-based axiomatization of nabla logic

Let us now detail the generic axiomatization of the nabla style of coalgebraic logics. From its genericity and its simplicity it warrants to be studied separately, and this what we will do in Chapter 4 too. We once again refer the reader to [KKV12b] for all the details about this axiomatization, and the proof of its completeness. For our purpose it will be enough to state that for a weak pullback preserving standard functor T , the 2-dimensional Hilbert system, which we call $\text{KKV}(T)$, given by the axioms and rules of Table 1.1 is sound and weakly complete with respect to the nabla-style of semantics in the class $\mathbf{CoAlg}(T)$.

Table 1.1 requires some additional information, in particular the following typing information: $a, b \in \mathcal{L}_T$, $\phi, \psi \in \mathcal{P}_\omega \mathcal{L}_T$, $X \in \mathcal{P}_\omega \mathcal{P}_\omega \mathcal{L}_T$, $\alpha, \beta \in T_\omega \mathcal{L}_T$, $\Phi \in T_\omega \mathcal{P}_\omega \mathcal{L}_T$, $A \in \mathcal{P}_\omega T_\omega \mathcal{L}_T$. The set $\text{Choice}(X)$ is the set of choice functions on X , i.e. the maps $\gamma : X \rightarrow \mathcal{L}_T$ such that $\gamma(\phi) \in \phi$, and rng denotes the range of the function. $R \subseteq \mathcal{L}_T \times \mathcal{L}_T$ is any relation and $\bar{T}R$ is its lifting. Finally $\text{SRD}(A)$ is the set of so-called ‘slim redistributions’ of A . This last concept is important, and we therefore define it in extenso. A **redistribution** of $A \in \mathcal{P}_\omega T_\omega \mathcal{L}_T$ is an element Φ of $T_\omega \mathcal{P}_\omega \mathcal{L}_T$ which ‘contains’ all the elements of A as lifted members, i.e. $\alpha \bar{T} \in \Phi$ for all $\alpha \in A$. It is called **slim** if it is build from the direct subformulae of the elements of A , i.e. if $\Phi \in T_\omega \mathcal{P}_\omega (\bigcup_{\alpha \in A} \text{B}(\alpha))$.

$\frac{}{a \leq a}$	$(\text{Cut}) \frac{a \leq c \quad c \leq b}{a \leq b}$
$(\vee L) \frac{\{a \leq b \mid a \in \phi\}}{\vee \phi \leq b}$	$(\vee R) \frac{a \leq b}{a \leq \vee \phi} b \in \phi$
$(\wedge L) \frac{a \leq b}{\wedge \phi \leq b} a \in \phi$	$(\wedge R) \frac{\{a \leq b \mid b \in \phi\}}{a \leq \wedge \phi}$
$(\neg E) \frac{\wedge \{\phi \cup \{\neg a\}\} \leq \vee \psi}{\wedge \phi \leq \vee \{\psi \cup \{a\}\}}$	$(\neg I) \frac{\wedge \{\phi \cup \{a\}\} \leq \vee \psi}{\wedge \phi \leq \vee \{\psi \cup \{\neg a\}\}}$
$(\text{Distributivity}) \frac{}{\wedge \{\vee \phi \mid \phi \in X\} \leq \vee \{\wedge \text{rng}(\gamma) \mid \gamma \in \text{Choice}(X)\}}$	
$(\nabla 1) \frac{\{a \leq b \mid (a, b) \in R\}}{\nabla \alpha \leq \nabla \beta} (\alpha, \beta) \in \bar{T}R$	
$(\nabla 2) \frac{\{\nabla (T \wedge)(\Phi) \leq b \mid \Phi \in \text{SRD}(A)\}}{\wedge \{\nabla \alpha \mid \alpha \in A\} \leq b}$	$(\nabla 3) \frac{\{\nabla \alpha \leq b \mid \alpha \bar{T} \in \Phi\}}{\nabla (T \vee)(\Phi) \leq b}$

Table 1.1: The KKV system

Before we move on to axiomatizing the abstract flavour of coalgebraic logic, we would like to highlight the fact that $(\nabla 1)$ is a congruence rule, and that the use of an auxiliary relation R in its premise allows for the system to be well-defined in the sense that $\leq$ is then not defined in terms of itself. A particularly natural choice of relation R is to consider the restriction of $\leq$ to subformulas. We will return to this point in examining the abstract case.

Axiomatizations in the abstract setting

Let us now turn to the abstract style of coalgebraic logic. We would like to define what we mean by an axiomatization in this setting too, and show that it subsumes what was done in the predicate lifting and nabla case, just as we did for the definition of the language and of the semantics. We will however delay this last point to the end of the chapter because it will involve concepts which we first need to introduce.

Let us place ourselves once again in the fundamental situation of Diagram (1.2) and assume that the syntax constructor L is regular epi-preserving. Recall that the abstract coalgebraic language associated with L is defined as $\mathbf{GFV}$, the free L -algebra over $\mathbf{FV}$. We must now abandon some of the generality of the abstract setting if we want to define a style of Hilbert systems which generalises the systems for the predicate lifting and nabla style, and make some assumptions on the category $\mathbf{C}$. For reasons which will become clear in the rest of this work, we will consider the category $\mathbf{C} = \mathbf{DL}$ as a good ‘minimal reasoning structure’. In fact we will consider Hilbert systems for the following choices of $\mathbf{C}$: $\mathbf{DL}$, $\mathbf{BDL}$ and $\mathbf{BA}$.

For any of the above choices of fundamental reasoning structure we will define a relation $\vdash_{\mathbf{ML}} \subseteq \mathbf{UVGFV} \times \mathbf{UVGFV}$ whose intuitive meaning will be $a \vdash_{\mathbf{ML}} b$ iff ‘ b can be derived from a ’ in the abstract Hilbert system. The reasons for choosing a 2-dimensional Hilbert system is that in $\mathbf{DL}$ inequalities are the only meaningful statements, since there might not be a top or a bottom element. We start with axioms. In the 2-dimensional setting, axioms are of the form $a \vdash_{\mathbf{ML}} b$ for some $a, b \in \mathbf{UVGFV}$. The propositional part of the system is given, by increasing strength of the base category, by the following axioms (we use the same notation as in the $\mathbf{KKV}(T)$ -axiomatization):

For **DL** :

$$\begin{aligned}
(\vee L) & \frac{\{a \vdash_{\text{ML}} b \mid a \in \phi\}}{\vee \phi \vdash_{\text{ML}} b} \\
(\wedge L) & \frac{a \vdash_{\text{ML}} b}{\wedge \phi \vdash_{\text{ML}} b} \quad (a \in \phi) \\
(\vee R) & \frac{a \vdash_{\text{ML}} b}{a \vdash_{\text{ML}} \vee \phi} \quad (b \in \phi) \\
(\wedge R) & \frac{\{a \vdash_{\text{ML}} b \mid b \in \phi\}}{a \vdash_{\text{ML}} \wedge \phi} \\
(\text{Distributivity}) & \frac{}{\wedge \{\vee \phi \mid \phi \in X\} \vdash_{\text{ML}} \vee \{\wedge \text{rng}(\gamma) \mid \gamma \in \text{Choice}(X)\}}
\end{aligned}$$

For **BDL**, add

$$\begin{aligned}
(\top) & \frac{}{a \vdash_{\text{ML}} \top} \\
(\perp) & \frac{}{\perp \vdash_{\text{ML}} a}
\end{aligned}$$

For **BA**, add

$$\begin{aligned}
(\neg E) & \updownarrow \frac{\wedge \{\phi \cup \{\neg a\}\} \vdash_{\text{ML}} \vee \psi}{\wedge \phi \vdash_{\text{ML}} \vee \{\psi \cup \{a\}\}} \\
(\neg I) & \updownarrow \frac{\wedge \{\phi \cup \{a\}\} \vdash_{\text{ML}} \vee \psi}{\wedge \phi \vdash_{\text{ML}} \vee \{\psi \cup \{\neg a\}\}}
\end{aligned}$$

where $\updownarrow$ means that the rule can be used in either direction. We also include the following structural rules:

$$\begin{aligned}
(\text{Reflexivity}) & \frac{}{a \vdash_{\text{ML}} a} & (\text{Cut}) & \frac{a \vdash_{\text{ML}} c \quad c \vdash_{\text{ML}} b}{a \vdash_{\text{ML}} b}
\end{aligned}$$

Finally, we need to add **substitution** and **congruence**. The standard notion of substitution can readily be transported to the abstract setting unchanged: given a map $\sigma : V \rightarrow \text{UVGFV}$ in **Set**, by using the two adjunctions connecting **Set** to $\mathbf{Alg}_{\mathbf{C}}(L)$, we get an adjoint transpose $\hat{\sigma} : \mathbf{GFV} \rightarrow \mathbf{GFV}$, such a map is called a substitution, and given a term $a \in \mathbf{GFV}$, the application of $\hat{\sigma}$ to a is traditionally denoted as $\hat{\sigma}a$ or $a\hat{\sigma}$ (we will use the former notation). We therefore add the rule

$$(\text{Substitution}) \frac{a \vdash_{\text{ML}} b}{\hat{\sigma}a \vdash_{\text{ML}} \hat{\sigma}b} \quad \sigma : V \rightarrow \text{UVGFV}$$

The notion of congruence is a little trickier to adapt to the abstract framework since L is not necessarily polynomial, i.e. the terms are not necessarily built from n -ary operation symbols. Recall that the congruence rule in the predicate lifting setting was given by:

$$(\text{Congruence}) \frac{a_i \leftrightarrow b_i, 1 \leq i \leq n}{\sigma(a_1, \dots, a_n) \leftrightarrow \sigma(b_1, \dots, b_n)} (\sigma \in \Sigma)$$

The key to generalizing this rule is first to notice that the premise $a_i \leftrightarrow b_i, 1 \leq i \leq n$ can be re-written as $(a_i, b_i) \in R, 1 \leq i \leq n$ where $R \rightarrow \mathbf{L}_\Sigma \times \mathbf{L}_\Sigma$ is the relation (in **Set**) defined by $(a, b) \in R$ iff $\mathbf{Ax} \vdash a \leftrightarrow b$. It is then easy to realize that the premise can be understood as a witness of the lifting of this relation R by a polynomial functor. We can thus rewrite the congruence rule as

$$(\text{Congruence}) \frac{((a_1, \dots, a_n), (b_1, \dots, b_n)) \in \overline{\Sigma}R}{\sigma(a_1, \dots, a_n), \sigma(b_1, \dots, b_n) \in R}$$

Prefixing by σ is essentially the structure map of the language $\mathbf{L}_\Sigma$, viewed as an initial $\mathbf{S}_\Sigma(-) + \mathbf{FV}$ -algebra. The congruence rule written in this way is clearly ready to be generalized to our Hilbert system for abstract coalgebraic logic. We re-define the relation R as the relation on $\mathbf{UVGFV}$ defined by aRb iff $a \vdash_{\mathbf{ML}} b$ and $b \vdash_{\mathbf{ML}} a$. Using the notion of relation lifting in a regular category developed in the previous section, we can then consider the lifting, $\overline{LR}$ of R and put for any $\alpha, \beta \in \mathbf{LVGFV}$:

$$(\text{Congruence}) \frac{(\alpha, \beta) \in \overline{LR}}{(\langle \alpha \rangle, \langle \beta \rangle) \in R}$$

where $\langle - \rangle : \mathbf{LGFV} \rightarrow \mathbf{GFV}$ is the structure map of the language. We will call this rule the **abstract congruence rule**, and note its similarity with the rule $(\nabla 1)$ of the **KKV** system. The reader will be justifiably concerned that the derivability relation $\vdash_{\mathbf{ML}}$ is now defined in terms of itself via the relation R . However, by using the notion of base, mentioned earlier and developed in detail in Chapter 3, we can tighten the abstract congruence rule and make it more similar to the congruence rule in the predicate lifting set-up which only uses in its premise the ‘ingredients’ necessary to build the terms in its conclusion. Since limits in **C** are created in **Set** (recall that in Diagram (1.2) we assume $\mathbf{U}$ to be monadic), and since the base is essentially a limit, the base $\mathbf{B}_{\mathbf{GFV}}(\alpha)$ is created in **Set**. Assuming that L preserves weak pullbacks, then it also preserves monos by Proposition 3.5, and this property is clearly preserved by $\mathbf{U}$. Moreover, if we assume that $\mathbf{UL}$ is finitary, then by Proposition 3.83 it preserves all intersections. In consequence under the relatively ‘normal’ assumption of weak-pullback preservation and finitariness, we have for any $\alpha \in \mathbf{LGFV}$, that $\alpha \in Li[\mathbf{LB}_L(\alpha)]$, where $i : \mathbf{B}_L \alpha \rightarrow \mathbf{GFV}$ is the obvious injection, and $Li[\mathbf{LB}_L(\alpha)]$

is the image of Li (which we can always take in a regular category). Since L is assumed to preserve regular epis we get, by Proposition 1.12, the following congruence rule:

$$(\text{Congruence}) \frac{(\alpha, \beta) \in \overline{L}(R \upharpoonright B_L(\alpha) \times B_L(\beta))}{(\langle \alpha \rangle, \langle \beta \rangle) \in R}$$

In other words, we need only find a witness in the L -lifting of the much smaller relation $(R \upharpoonright B_L(\alpha) \times B_L(\beta))$, just as was the case in the predicate lifting style Hilbert system. The definition of $\mathbf{Ax} \vdash_{\mathbf{ML}}$ is thus an inductive definition. Formally, given a set of axioms $\mathbf{Ax} \subseteq (\mathbf{UVGFV})^2$, we define the abstract Hilbert system associated with it as the system whose axioms are those of $\mathbf{Ax}$, and whose rules are (1) the propositional rules defined above and corresponding to the appropriate base category (1) the reflexivity and cut rules, (3) closure under substitution instances, and (4) the abstract congruence rule we have just detailed.

The ‘propositional rules’ essentially enforce that the relation $\vdash_{\mathbf{ML}}$ is in fact a relation in $\mathbf{C}$, i.e. a subobject of $(\mathbf{VGFV})^2$. Similarly, the congruence rule guarantees that $\vdash_{\mathbf{ML}}$ is in fact a relation in $\mathbf{Alg}_{\mathbf{C}}(L)$, i.e. a subobject of $(\mathbf{GFV})^2$ as we will now show. Note first that since $\mathbf{C}$ is assumed to be monadic over $\mathbf{Set}$, it is regular. Moreover, since L is a regular epi preserving variator $\mathbf{Alg}_{\mathbf{C}}(L)$ is regular (by Proposition 1.19), and relations can thus be defined in $\mathbf{Alg}_{\mathbf{C}}(L)$. It is also monadic over $\mathbf{C}$ (by Proposition 1.9) and limits in $\mathbf{Alg}_{\mathbf{C}}(L)$ are therefore created in $\mathbf{C}$; and in particular the product of two L -algebras $\theta : LA \rightarrow A$ and $\xi : LB \rightarrow B$ is defined by the structure map $L(A \times B) \rightarrow A \times B$ via the unique map $\theta \circ L\pi_1 \times \xi \circ L\pi_2$ where π_1, π_2 are the obvious projections. By definition of the product in $\mathbf{Alg}_{\mathbf{C}}(L)$, and of relation lifting we have the following commutative diagram:

$$\begin{array}{ccccc}
 & & \overline{LR} & & \\
 & \nearrow & & \nwarrow & \\
 LR & \xrightarrow{\quad} & L(A \times A) & \xrightarrow{L\pi_1 \times L\pi_2} & LA \times LA \\
 \downarrow & & \downarrow \theta \circ L\pi_1 \times \theta \circ L\pi_2 & & \swarrow \theta \times \theta \\
 R & \xrightarrow{\quad} & A \times A & &
 \end{array} \tag{1.3}$$

Recall that the epi $LR \rightarrow \overline{LR}$ is regular i.e. the underlying $\mathbf{Set}$ map is surjective. The meaning of the congruence rule is now much clearer: it simply states that if $(\alpha, \beta) \in \overline{LR}$ for some $\alpha, \beta \in LA$, then $(\theta(\alpha), \theta(\beta)) \in R$, i.e. the congruence rule allows us to specify a relation on $L(A \times A) \rightarrow A \times A$ in $\mathbf{Alg}_{\mathbf{C}}(L)$ by using a relation on LA in $\mathbf{C}$. We have thus shown:

Proposition 1.28. *Let $L : \mathbf{C} \rightarrow \mathbf{C}$ be a regular epi-preserving variator. A relation $R \rightrightarrows A \times A$ in $\mathbf{C}$ is a relation in $\mathbf{Alg}_{\mathbf{C}}(L)$ on the L -algebra $\theta : LA \rightarrow A$*

iff the congruence rule is satisfied, i.e. iff $(\theta(\alpha), \theta(\beta)) \in R$ whenever $(\alpha, \beta) \in \overline{LR}$.

Thus, the congruence rule gives us a criterion on pairs of modal formulas (i.e. elements of $LGFV \times LGFV$) which enforces that the (equivalence) relation R in $\mathbf{C}$ is also an (equivalence) relation in $\mathbf{Alg}_{\mathbf{C}}(L)$.

1.5.4 One style of algebraic semantics

As in the case of (classical) modal logic (see Chapter 5 of [BdRV01]), it is very useful to define an algebraic semantics for the languages defined above. Fundamentally, the usefulness of an algebraic semantics lies in its close proximity with the languages themselves and with derivability. As we will now show, the algebraic semantics gives us an alternative description of derivability in the language of universal algebra, which is in many ways easier to manipulate than individual (modal) axiomatizations. This will be crucial in our discussion of translations in Chapter 4.

We saw earlier that the abstract description of coalgebraic languages and their semantics provides a unifying framework in which we don't have to worry about boolean reasoning. We therefore choose this level of abstraction to define an algebraic semantics (it is also interesting to look at varieties over a category different than $\mathbf{Set}$).

Let once again $L : \mathbf{C} \rightarrow \mathbf{C}$ be a regular epi-preserving variator and let us consider the language $\mathbf{GFV}$. We will use the fact that if we define $L' = L(-) + \mathbf{FV}$, the language can also be seen as $\mu L'$, the initial L' -algebra, to give a very simple and elegant definition of the algebraic semantic of $\mu L'$. Let $\theta : L'A \rightarrow A$ be an arbitrary L' -algebra, we define the **algebraic interpretation map** $\llbracket - \rrbracket_A$ as the unique morphism:

$$\begin{array}{ccc} L'\mu L' & \xrightarrow{L'\llbracket - \rrbracket_A} & L'A \\ \downarrow & & \downarrow \theta \\ \mu L' & \xrightarrow{\llbracket - \rrbracket_A} & A \end{array}$$

Observe that by definition of L' , structure maps of the type $\theta : L'A = LA + \mathbf{FV} \rightarrow A$ are in one to one correspondence with pairs of maps $\xi : LA \rightarrow A$ and $\hat{v}_A : \mathbf{FV} \rightarrow A$. By adjunction, the latter are in one-to-one correspondence with functions $v_A : V \rightarrow \mathbf{UA}$, i.e. valuations. In other words, by considering L' algebras rather than L algebras we automatically include the notion of valuation which is essential to the definition of an algebraic semantics. In particular, every choice of $v_A : V \rightarrow \mathbf{UA}$ for a fixed $\xi : LA \rightarrow A$ defines a new L' -algebra and a new algebraic interpretation map into A . We will say that a formula $a \in \mu L'$

is **algebraically satisfiable** if there exists an L' -algebra $\xi : L'A \rightarrow A$ such that $\llbracket a \rrbracket_A \neq \perp$. Similarly, we will say that a formula $a \in \mu L'$ is **algebraically valid over a class $\mathcal{C}$ of L -algebras** - denoted $\mathcal{C} \models a$ - if for every algebra $\xi : LA \rightarrow A$ in $\mathcal{C}$ and every valuation $v_A : V \rightarrow \mathbf{UA}$, $\llbracket a \rrbracket_A = \top$, where $\llbracket - \rrbracket_A$ is defined as above by the catamorphism into $\xi + \hat{v}_A : L'A \rightarrow A$. Similarly, we say that an *equation* $a = b$ between two abstract terms is algebraically valid over $\mathcal{C}$ - denoted $\mathcal{C} \models a = b$ if $\llbracket a \rrbracket_A = \llbracket b \rrbracket_B$ for all algebras A in $\mathcal{C}$ and every valuation. The purpose of this section is to prove that an axiomatization is always sound and complete with respect to the algebraic semantics over a class of algebra which is naturally associated with the axiomatization.

Axiomatizations are often given in terms of a set $\mathbf{Ax} \subset \mathbf{UVGFV}$ of *axioms*, but we will prefer to work with a set of *equations* i.e. a set $E \subset \mathbf{UVGFV} \times \mathbf{UVGFV}$. From E we will build two entities. First, on the semantic side, we will use E to build an equivalence relation (in $\mathbf{Alg}_{\mathbf{C}}(L)$) on $\mathbf{GFV}$. This equivalence relation will define, via its coequalizer, a class of L -algebras (more precisely a variety of L -algebras) which will form the semantic domain of our algebraic semantics. Secondly, on the proof theoretic side, we will use the set of equations to define an equational calculus, i.e. a set of rules to build derivations of equations. Finally, we will prove the adequacy of the semantics and proof-theoretic constructions.

Defining a variety from equations

We want to build an equivalence relation on $\mathbf{GFV}$ whose interpretation will be that two terms a, b are related if they can be proved to be equal in the equational calculus we will define in a moment. Having defined relations in regular categories, we now know that this means a subobject of $\mathbf{GFV} \times \mathbf{GFV}$ viewed as an L -algebra. But all we have so far is a *set* $\mathbf{Ax}$ of equations. In order to turn this set into a subobject of $\mathbf{GFV} \times \mathbf{GFV}$, and in particular into an L -algebra, we proceed as follows. The set $\mathbf{Ax}$ of equations naturally defines two functions

$$e_1, e_2 : \mathbf{Ax} \rightarrow \mathbf{UVGFV}$$

where e_1 picks the left-hand-side term of the equations and e_2 the right-hand-side term. It is natural to assume that e_1, e_2 are jointly monic in $\mathbf{Set}$, i.e. that if $e_1(a) = e_1(b)$ and $e_2(a) = e_2(b)$ then $a = b$. Indeed the premise $e_1(a) = e_1(b)$ and $e_2(a) = e_2(b)$ means that we have the same equation appearing twice in $\mathbf{Ax}$, and we can clearly eliminate the redundant copy without loss of generality. So we will assume e_1, e_2 jointly monic, i.e. $(e_1 \times e_2)$ monic as we saw earlier. We would like to turn these arrows in $\mathbf{Set}$ into an equivalence relations in $\mathbf{Alg}_{\mathbf{C}}(L)$, which suggests the following construction. First we use the adjunction to get two arrows in $\mathbf{Alg}_{\mathbf{C}}(L)$, i.e.

$$\hat{e}_1, \hat{e}_2 : \mathbf{GFAx} \rightarrow \mathbf{GFV}$$

Intuitively, $\hat{e}_1, \hat{e}_2$, create algebraic combinations of the original equations. Second, we take the kernel pair of their coequalizer, i.e. we define

$$\text{GFAx} \begin{array}{c} \xrightarrow{\hat{e}_1} \\ \xrightarrow{\hat{e}_2} \end{array} \text{GFV} \xrightarrow{q} Q$$

This coequalizer quotients GFV under the equations. Finally, we consider the kernel pair $\ker q$ of q

$$\begin{array}{ccc} \ker q & \xrightarrow{p_1} & \text{GFV} \\ p_2 \downarrow & & \downarrow q \\ \text{GFV} & \xrightarrow{q} & Q \end{array}$$

This kernel pair clearly defines an equivalence relation in $\mathbf{Alg}_{\mathbf{C}}(L)$ (see Lemma 1.13), and it captures all the pairs of terms which are quotiented by q . The regular quotient q defines a variety of L -algebras which we will denote $\mathbb{V}_q$. To define a variety we need the following concept:

Definition 1.29. Let C be an object and $f : A \rightarrow B$ be a morphism of a category $\mathbf{C}$. We say that C is **orthogonal** to f , or that A is in the **orthogonality class** of f if for every morphism $g : A \rightarrow C$ there exist a unique morphism $u : B \rightarrow C$ such that $g = u \circ f$. The orthogonality relation is denoted by $C \perp f$.

We refer the reader to [Bor94a] 5.4 or [AR94] 1.C for more details on this concept. We follow [Kur00, Hug01] and define a **variety** in a category $\mathbf{C}$ as the orthogonality class of a regular epimorphism whose domain is a projective object. In all our applications we will consider categories of the form $\mathbf{Alg}_{\mathbf{C}}(L)$ for a regular epi-preserving variator L and a regular category $\mathbf{C}$. In this case it is easy to check by using the obvious adjunction $\mathbf{G} \dashv \mathbf{V} : \mathbf{C} \rightarrow \mathbf{Alg}_{\mathbf{C}}(L)$, that any object in $\mathbf{Alg}_{\mathbf{C}}(L)$ of the form $\mathbf{F}A$ for some A in $\mathbf{C}$ is projective if A is projective in $\mathbf{C}$. Note that in the case of an adjunction $\mathbf{F} \dashv \mathbf{U} : \mathbf{Set} \rightarrow \mathbf{C}$ where $\mathbf{Set}$ is understood to satisfy the axiom of choice, $\mathbf{F}A$ is always projective, since all sets are projective objects. This means in particular that if $L : \mathbf{C} \rightarrow \mathbf{C}$ is a regular epi-preserving variator, then GFV is projective, so our very general definition of variety does apply to the case we are most interested in, i.e. if we define $\mathbb{V}_q$ as the orthogonality class of q , then $\mathbb{V}_q$ is a variety. We refer the reader to Chapter 3 of [Hug01] for a categorical version of Birkhoff's HSP theorem which uses the definition of variety that we have presented here.

Before moving on to the equational logic defined by Ax , we would like to make the following important comment which will simplify proofs further on. If $q : \text{GFV} \twoheadrightarrow Q$ is a regular epi, then q defines a variety, but Q is in general not a member of this variety. Equivalently (see Theorem 3.5.3. in [Hug01]), $\ker q$ is in general not fully invariant. If either of these conditions were satisfied, Q

would represent the ‘prototypical’ or ‘most general’ object in the variety. In particular, it would be the initial object in the category defined by the variety. Obtaining such a prototypical Q simplifies many proofs considerably, since it is then enough to reason only about Q .

Proposition 1.30. *Let $\mathbf{C}$ be a cocomplete well-powered regular category, let P be a projective object in $\mathbf{C}$ and $q : P \twoheadrightarrow Q$ be the coequalizer of two jointly monomorphic arrows $p_1, p_2 : E \rightarrow P$. If we denote by $q^* : P \twoheadrightarrow Q^*$ the regular epi defined in Proposition 1.23 and which constructs the fully invariant closure of $\ker q$, then*

$$\mathbb{V}_q = \mathbb{V}_{q^*}$$

Proof. By using the construction of Proposition 1.23 it is clear q^* also coequalizes p_1, p_2 (since it coequalizes $f \circ p_1, f \circ p_2$ for all $f \in \text{hom}(P, P)$) and we thus have a unique morphism u :

$$\begin{array}{ccccc} E & \xrightarrow{p_1} & P & \xrightarrow{q} & Q \\ & \xrightarrow{p_2} & & \searrow & \downarrow u \\ & & & q^* & Q^* \end{array}$$

Note that since $\mathbf{C}$ is regular and q^* is a regular epi, then u must be a regular epi ([Bor94b] 2.1.5.). It is now easy to check that if A is orthogonal to q^* , then it must also be orthogonal to q , i.e. for every $f : P \rightarrow A$

$$\begin{array}{ccc} & Q & \\ q \nearrow & & \searrow u \\ P & \xrightarrow{q^*} & Q^* \\ f \searrow & & \swarrow \text{---} \\ & A & \end{array}$$

For the converse implication, consider any $f : P \rightarrow A$, any $h : P \rightarrow A$, and assume that A is orthogonal to q , i.e. for any $g : P \rightarrow A$, there exists a unique $v : Q \rightarrow A$ such that $v \circ q = g$. In particular if we take $g = h \circ f$ for some $h : P \rightarrow A$. We then have that

$$h \circ f \circ e_1 = g \circ e_1 = v \circ q \circ e_1 = v \circ q \circ e_2 = g \circ e_2 = h \circ f \circ e_2$$

and thus $h : P \rightarrow A$ coequalizes the pair $f \circ e_1, f \circ e_2$. Since the choice of f is irrelevant, we actually get that h coequalizes $\coprod_{f \in \text{hom}(P, P)} f \circ e_i, i = 1, 2$, and there must therefore exist a unique morphism $u : Q^* \rightarrow A$ such that $u \circ q^* = h$, by definition of q^* and Q^* . In other words A is orthogonal to q^* which concludes the proof. $\square$

Equational logic in $\mathbf{C}$

We now define the equational logic associated with $\mathbf{Ax}$ by specifying the following derivation rules. For any $a, b, c, d \in \mathbf{GFV}$, $\alpha, \beta \in \mathbf{LGFV}$

EL(1) Axioms: $\mathbf{Ax} \vdash_{\mathbf{EL}} a = b$ if $(a, b) \in \mathbf{Ax}$

EL(2) Reflexivity: $\mathbf{Ax} \vdash_{\mathbf{EL}} a = a$

EL(3) Symmetry: if $\mathbf{Ax} \vdash_{\mathbf{EL}} a = b$ then $\vdash_{\mathbf{EL}} b = a$

EL(4) Transitivity: if $\mathbf{Ax} \vdash_{\mathbf{EL}} a = b$ and $\vdash_{\mathbf{EL}} b = c$ then $\vdash_{\mathbf{EL}} a = c$

EL(5) Substitutions: if $\mathbf{Ax} \vdash_{\mathbf{EL}} a = b$ and $\sigma : \mathbf{GFV} \rightarrow \mathbf{GFV}$ is a substitution, then $\mathbf{Ax} \vdash_{\mathbf{EL}} \sigma a = \sigma b$

EL(6) Modal congruence: if $(\alpha, \beta) \in \overline{L}(\mathbf{Ax} \vdash_{\mathbf{EL}})$ then $\mathbf{Ax} \vdash_{\mathbf{EL}} \langle \alpha \rangle = \langle \beta \rangle$

EL(7) $\mathbf{C}$ -congruence: $\vdash_{\mathbf{EL}}$ is a relation in $\mathbf{C}$. In particular, if $\mathbf{C} = \mathbf{DL}$ we have that $\mathbf{Ax} \vdash_{\mathbf{EL}} a = b$ and $\mathbf{Ax} \vdash_{\mathbf{EL}} c = d$ imply $\mathbf{Ax} \vdash_{\mathbf{EL}} a \wedge c = b \wedge d$, $\mathbf{Ax} \vdash_{\mathbf{EL}} a \vee c = b \vee d$. If $\mathbf{C} = \mathbf{BA}$ we also have $\mathbf{Ax} \vdash_{\mathbf{EL}} \neg a = \neg b$

Proposition 1.31. *Assuming that L preserves regular epimorphisms and weak pullbacks, the binary derivability predicate $\mathbf{Ax} \vdash_{\mathbf{EL}}$ defines the smallest fully invariant equivalence relation on $\mathbf{GFV}$ which contains $\mathbf{Ax}$.*

Proof. Let us first show that $\mathbf{Ax} \vdash_{\mathbf{EL}}$ does indeed define a fully invariant equivalence relation on $(\mathbf{GFV}, \langle - \rangle)$ in $\mathbf{Alg}_{\mathbf{C}}(L)$. Let K denote the relation defined by $(a, b) \in K$ iff $\mathbf{Ax} \vdash_{\mathbf{EL}} a = b$. It is immediate by virtue of EL7, that $m : K \rightarrow \mathbf{GFV} \times \mathbf{GFV}$ in $\mathbf{C}$. The congruence rule EL6 ensures that K is an L -algebra. Indeed, if we define

$$\xi = \langle - \rangle \times \langle - \rangle \circ L\pi_1 \times L\pi_2 \circ Lm$$

then the modal congruence rule guarantees that the image of that map lies in K (see Diagram 1.3). Hence $(K, \xi) \rightarrow (\mathbf{GFV} \times \mathbf{GFV}, \langle - \rangle \circ L\pi_1 \times \langle - \rangle \circ L\pi_2)$, i.e. K is a sub L -algebra of $\mathbf{GFV} \times \mathbf{GFV}$. Finally, the rules EL2-4 ensures that K is an equivalence relation, and rule 1 ensures that it contains $\mathbf{Ax}$. Thus K is an equivalence relation in $\mathbf{Alg}_{\mathbf{C}}(L)$ which contains $\mathbf{Ax}$. The fact that it is fully invariant is a direct consequence of EL5.

Now let us show that K is the *smallest* fully invariant equivalence relation on $\mathbf{GFV}$ containing $\mathbf{Ax}$. Assume K' is another fully invariant equivalence relation on $(\mathbf{GFV}, \langle - \rangle)$ containing $\mathbf{Ax}$, and let's show that $(a, b) \in K \Rightarrow (a, b) \in K'$. We proceed by induction on the length of derivation of (a, b) .

Base case: Clearly, if $(a, b) \in \mathbf{Ax}$, then $(a, b) \in K'$ by assumption. Similarly, if $(a, a) \in K$, then $(a, a) \in K'$ since K' is an equivalence relation.

Inductive case: We go through the possible last steps of the derivation $\text{Ax} \vdash_{\text{EL}} a = b$.

EL3: the last step of the derivation is an application of the symmetry rule, i.e. there is a shorter proof of $\text{Ax} \vdash_{\text{EL}} b = a$, and by induction hypothesis, this means that $(b, a) \in K'$ and thus $(a, b) \in K'$ since it's an equivalence relation.

EL4: the proof is the same as in the previous case.

EL5: the last step of the derivation is an application of the substitution rule, i.e. there exists a shorter proof of $\text{Ax} \vdash_{\text{EL}} c = d$ and a substitution $\sigma : \text{GFV} \rightarrow \text{GFV}$ such that $\sigma c = a, \sigma d = b$. By the induction hypothesis we have $(c, d) \in K'$, and since by assumption K' is fully invariant, we must also have $(\sigma c, \sigma d) \in K'$.

EL6: The congruence rule is quite different, and it is not immediately clear how the induction hypothesis works in this case. Since L is assumed to preserve regular epis and weak pullbacks we can assume that the last rule applied was in fact the congruence rule restricted to the bases. Thus, what we need to apply this rule, is to have established derivability between enough formulas in the bases of α and β - but not necessarily to have established the full relation $K \cap \mathbf{B}_L(\alpha) \times \mathbf{B}_L(\beta)$ - to be able to apply L to a subobject of $R \rightrightarrows K \cap \mathbf{B}_L(\alpha) \times \mathbf{B}_L(\beta)$ and get a witness in $LR \rightrightarrows L(K \cap \mathbf{B}_L(\alpha) \times \mathbf{B}_L(\beta))$ which projects to α, β , i.e. such that modulo some mono, we have $(\alpha, \beta) \in \overline{LR}$. In other words, if we apply this rule in a derivation, it means that we have shorter proofs of equality between terms in the bases of α, β , which we can collect in a subrelation $R \rightrightarrows K \cap \mathbf{B}_L(\alpha) \times \mathbf{B}_L(\beta)$ whose lifting contains the witness we need. By the induction hypothesis we have $R \rightrightarrows K' \cap \mathbf{B}_L(\alpha) \times \mathbf{B}_L(\beta)$. By using the properties of relation liftings from Propositions 1.11 and 1.12 we then get $(\alpha, \beta) \in \overline{LK'} \cap (LB_L(\alpha) \times LB_L(\beta))$ and thus $(\langle \alpha \rangle, \langle \beta \rangle) \in K'$ by the fact that K' is a relation in $\mathbf{Alg}_{\mathbf{C}}(L)$.

EL7: immediate since K' is an L -algebra, i.e. its carrier is a $\mathbf{C}$ -object. $\square$

Soundness and completeness results

We now prove that provability in the abstract Hilbert system is equivalent to provability in the equational logic derivation system. We first translate axioms between the two system. Assuming that $\mathbf{C} = \mathbf{DL}, \mathbf{BDL}$ or $\mathbf{BA}$, it makes sense to define the following translation of axioms in the Hilbert system into equations:

$$(a, b)^t = (a \wedge b, a)$$

And conversely from equations into Hilbert axioms:

$$(a, b)_t = \{(a, b), (b, a)\}$$

More generally, we will denote by $(\text{Ax})^t$ the set of equations obtained from a set of Hilbert-style axioms by applying the map $(-)^t$, and conversely for Ax_t . Note

that every axiom of $\mathbf{Ax}$ can be recovered from $((\mathbf{Ax})^t)_t$: assume $(a, b) \in \mathbf{Ax}$ then $(a \wedge b, a), (a, a \wedge b) \in ((\mathbf{Ax})^t)_t$ and we get

$$(\text{Cut}) \frac{a \vdash_{\text{ML}} a \wedge b \quad (\wedge L) \frac{b \vdash_{\text{ML}} b}{a \wedge b \vdash_{\text{ML}} b} \quad b \in \{a, b\}}{a \vdash_{\text{ML}} b}$$

Conversely, every equation in $\mathbf{Ax}$ can be recovered from $((\mathbf{Ax})_t)^t$: if $(a, b) \in \mathbf{Ax}$ then $(a \wedge b, a), (b \wedge a, b) \in ((\mathbf{Ax})_t)^t$, and the original equality can be recovered from the fact that $a \wedge b = b \wedge a$ modulo $\mathbf{C}$ -equivalence and EL3 and EL4.

Proposition 1.32. *Let $L : \mathbf{C} \rightarrow \mathbf{C}$ be a finitary, weak pullback and regular epi-preserving functor and let V be a set of variables. Consider $a, b \in \text{UVGFV}$ and a set of axioms $\mathbf{Ax} \subseteq (\text{UVGFV})^2$ in the Hilbert system of Section 1.5.3*

$$\text{if } a \vdash_{\text{ML}} b \text{ then } (\mathbf{Ax})^t \vdash_{\text{EL}} a \wedge b = a$$

Conversely, assuming a set of equations $\mathbf{Ax} \subseteq (\text{UVGFV})^2$,

$$\text{if } \mathbf{Ax} \vdash_{\text{EL}} a = b \text{ then } a \vdash_{\text{ML}} b \text{ and } b \vdash_{\text{ML}} a \text{ can be derived from } (\mathbf{Ax})_t$$

Proof. By induction on the size of the derivation trees.

Let us show the first implication. If $a \vdash_{\text{ML}} b$ is derived in 1-step, then $(a, b) \in \mathbf{Ax}$, i.e. $(a \wedge b, a) \in (\mathbf{Ax})^t$ by definition of $(\mathbf{Ax})^t$.

Now for the inductive step. If the last step of the proof is propositional it is not hard to check that since $\vdash_{\text{EL}}$ is defined on equivalence classes modulo equivalence in $\mathbf{C}$, we can translate the last step into an equational proof. We show the example of the rule $(\vee L)$ in the case where $\phi = \{a, b\}$ and $a \vee b \vdash_{\text{ML}}$ is the conclusion. By the induction hypothesis we can then build an equational derivation ending in

$$\text{EQ7} \frac{a \wedge c = a \quad b \wedge c = b}{(a \vee b) \wedge c = a \vee b}$$

where we've used the fact that $(a \wedge c) \vee (b \wedge c)$ and $(a \vee b) \wedge c$ are two representative of the same equivalence class under the axioms of $\mathbf{C}$. The same goes for all propositional rules. The case of the Reflexivity rule is trivial. Assume now that the last rule applied was the Cut rule, i.e. that we have shorter derivations of $a \vdash_{\text{ML}} b'$ and $b' \vdash_{\text{ML}} b$. By the induction hypothesis, this means that we have derivations of $\vdash_{\text{EL}} a \wedge b' = a$ and $\vdash_{\text{EL}} b' \wedge b = b'$. We can then proceed as follows:

$$\begin{array}{c} \text{EL3} \frac{a \wedge b' = a}{a = a \wedge b'} \quad b = b \quad \text{EL7} \frac{a = a \quad b' \wedge b = b'}{a \wedge b \wedge b' = a \wedge b'} \\ \text{EL7} \frac{a = a \wedge b' \quad b = b}{a \wedge b = a \wedge b' \wedge b} \\ \text{EL4} \frac{a \wedge b = a \wedge b' \wedge b \quad a \wedge b' = a}{a \wedge b = a \wedge b'} \quad \text{EL4} \frac{a \wedge b = a \wedge b' \quad a \wedge b' = a}{a \wedge b = a} \end{array}$$

The case of the substitution rule follows trivially from the fact that substitutions are L -algebra morphisms. Finally, for the congruence rule, assume that $\alpha, \beta \in \overline{L}(R \upharpoonright \mathbf{B}_L(\alpha) \times \mathbf{B}_L(\beta))$. By the induction hypothesis we have for any $(a, b) \in R \upharpoonright \mathbf{B}_L(\alpha) \times \mathbf{B}_L(\beta)$, $\vdash_{\text{EL}} a \wedge b = b$ and $\vdash_{\text{EL}} b \wedge a = a$, i.e. $\vdash_{\text{EL}} a = b$, and the result then follows from an application of EL6.

The second implication poses no difficulty. The case of axioms follows immediately from the definition of $(-)_t$ and if the equational proof ends in EL1-EL3 the I.H. immediately provides the corresponding $\vdash_{\text{ML}}$ proof. In the case of EL4, we easily get a $\vdash_{\text{ML}}$ proof from the I.H. and the Cut rule. The case of the substitution rule is trivial and Congruence is treated exactly as above but in the opposite direction. Finally it is not difficult to check that any equational proof ending in an application of EL7 can be translated into a pair of $\vdash_{\text{ML}}$ proof by using the propositional rules pertaining to $\mathbf{C}$. For example, if the last step of an equation derivation is

$$\text{EL7} \frac{a = c \quad b = d}{a \wedge b = c \wedge d}$$

then by the I.H., we have proofs of $a \vdash_{\text{ML}} c$, $c \vdash_{\text{ML}} a$, $b \vdash_{\text{ML}} d$ and $d \vdash_{\text{ML}} b$ and it follows that we have a proof ending in

$$\frac{(\wedge L) \frac{b \vdash_{\text{ML}} d}{a \wedge b \vdash_{\text{ML}} d} b \in \{a, b\} \quad (\wedge L) \frac{a \vdash_{\text{ML}} c}{a \wedge b \vdash_{\text{ML}} c} a \in \{a, b\}}{(\wedge R) \frac{a \wedge b \vdash_{\text{ML}} d \quad a \wedge b \vdash_{\text{ML}} c}{a \wedge b \vdash_{\text{ML}} c \wedge d}}$$

We similarly get a proof of $c \wedge d \vdash_{\text{ML}} a \wedge b$ and the result follows. $\square$

Let us now move to the algebraic semantic. Recall that the set of equations Ax can be used to define a regular quotient $q : \text{GFV} \rightarrow Q_{\text{Ax}}$ via the coequalizer (in $\mathbf{Alg}_{\mathbf{C}}(L)$) of the maps $\hat{e}_1, \hat{e}_2$ associated with the set Ax . Moreover, using the construction of Proposition 1.23 and the result of Proposition 1.30, we can assume that Q_{Ax} defines the smallest fully invariant equivalence relation containing the equations of Ax , and is thus orthogonal to q .

Proposition 1.33. *Using the above notation, we have*

$$\text{Ax} \vdash_{\text{EL}} a = b \iff Q_{\text{Ax}} \models a = b$$

Proof. We just need to show that the relation $K \subseteq \text{GFV} \times \text{GFV}$ defined by $(a, b) \in K$ if $Q_{\text{Ax}} \models a = b$ is the smallest fully invariant equivalence relation containing Ax , and the result will then follow from Proposition 1.31. For notational clarity we drop the Ax subscript for the remainder of the proof. The definition of the algebraic semantic means that for any valuation v_Q on Q ,

$$\llbracket a \rrbracket_Q = \llbracket b \rrbracket_Q$$

where $\llbracket - \rrbracket_Q : \mu L' \rightarrow Q$ is the unique L' -algebra morphism defined by initiality of $\mu L'$ and the fact that v_Q endows Q with an L' -algebra structure. We show that the equality

$$\llbracket a \rrbracket_Q = \llbracket b \rrbracket_Q$$

implies $q(a) = q(b)$. Since the equality holds for any valuation, it must, in particular, hold for the valuation defined by

$$v_Q : FV \rightarrow Q, p \mapsto q \circ \eta_{FV}^G(p)$$

where $\eta^G : \text{Id} \rightarrow \mathbf{V} \mathbf{G}$ is the unit of the adjunction $\mathbf{G} \dashv \mathbf{V} : \mathbf{C} \rightarrow \mathbf{Alg}_{\mathbf{C}}(L)$. This valuation turns the L -algebra morphism q into an L' -algebra morphism, $\llbracket a \rrbracket_Q = \llbracket b \rrbracket_Q$ then implies $q(a) = q(b)$ by definition. Conversely, it is immediate that if $q(a) = q(b)$, then the fact that $Q \perp q$ implies that $\llbracket a \rrbracket_Q = \llbracket b \rrbracket_Q$ for any valuation on Q .

Thus $(a, b) \in K$ iff $(a, b) \in \ker q$. But we know from the construction in Proposition 1.23, that $\ker q$ is the smallest fully invariant equivalence relation which contains Ax , which is what we needed to show. $\square$

Proposition 1.34. *Let $L : \mathbf{C} \rightarrow \mathbf{C}$ be a variator, and consider a regular quotient $q : \text{GFV} \rightarrow Q$ defining a variety $\mathbb{V}_q$ such that $Q \perp q$, then*

$$Q \models a = b \iff \text{for all } A \text{ in } \mathbb{V}_q, A \models a = b$$

Proof. By assumption, Q itself belongs to the variety $\mathbb{V}_q$ which is defined by q , thus the implication from right to left is trivial.

For the implication from left to right, it is enough to unravel the definitions. Consider any L -algebra A in the variety $\mathbb{V}_q$ defined by q , and assume that $Q \models a = b$, i.e. $\llbracket a \rrbracket_Q = \llbracket b \rrbracket_Q$. As we've just seen in the proof of Proposition 1.33, this in turn implies that $q(a) = q(b)$. By definition of the variety, the unique L -algebra morphism $\llbracket - \rrbracket_A : \mu L' \rightarrow A$ factors through q , and the result follows immediately. $\square$

Gathering our results, we have the following theorem summing up the completeness of the algebraic semantics of coalgebraic logics.

Theorem 1.35. *Under the conditions of Proposition 1.34 the following are equivalent:*

- (i) $a \vdash_{\text{ML}} b$ and $b \vdash_{\text{ML}} a$ can be derived from $(\text{Ax})_t$
- (ii) $\text{Ax} \vdash_{\text{EL}} a = b$
- (iii) $Q_{\text{Ax}} \models a = b$
- (iv) for all A in $\mathbb{V}_{q_{\text{Ax}}}$, $A \models a = b$

Alternatively, given a set $\mathbf{Ax} \subseteq (\mathbf{UVGFV})^2$, we could have written $a \vdash_{\mathbf{ML}} b$ can be derived from $\mathbf{Ax}$ iff $(\mathbf{Ax})^t \vdash_{\mathbf{EL}} a \wedge b = a$ iff $Q_{(\mathbf{Ax})^t} \models a \wedge b = a$ iff for all A in $\mathbb{V}_{q_{(\mathbf{Ax})^t}}$, $A \models a \wedge b = a$. Note however, that by working with equations, and using the universal constructions detailed above, the collection of valid/derivable equations forms a fully invariant equivalence relation in $\mathbf{C}$. In contrast, the set of derivable pairs of formula in a modal-style Hilbert system does not live in $\mathbf{Alg}_{\mathbf{C}}(L)$ when $\mathbf{C} = \mathbf{BA}$, and in particular it is not a subobject of the language (viewed as an L -algebra). In fact it does not even live in $\mathbf{C}$, since this would imply a system where $a \vdash_{\mathbf{ML}} b$ iff $\neg a \vdash_{\mathbf{ML}} \neg b$.

Comparing the systems

We finish this chapter by showing that the Hilbert systems which we have defined for the abstract version of coalgebraic logic are enough to reason about derivability in the predicate lifting and nabla flavour of coalgebraic logic. Combined with the results of Theorem 1.35, this will be very useful in Chapter 4. Let us first show that axiomatizations of the abstract style of coalgebraic logic subsume axiomatizations of the predicate lifting and nabla style. In order to achieve this, we will use some of the tools developed so far in order to understand precisely how the concrete and abstract languages are related. The intuitive relation is clear, but formalizing this relation is not totally trivial.

Let $\Sigma_{\mathbf{BA}}$ denote a signature which can accommodate a description of boolean algebras, for example $\Sigma_{\mathbf{BA}} = \{\perp, \neg, \wedge\}$ (as in the predicate lifting syntax), or $\Sigma_{\mathbf{BA}} = \{\neg, \wedge, \vee\}$ (as in the nabla logic syntax). Let $\mathbf{S}_{\mathbf{BA}}$ denote the **Set**-endofunctor associated with this signature, i.e. the obvious polynomial functor, or the functor $\text{Id} + \mathcal{P}_{\omega} + \mathcal{P}_{\omega}$ in the nabla case.

Let also T be a finitary functor. We can consider the nabla (or predicate lifting if T is polynomial) style language associate with T , which is just the free $T + \mathbf{S}_{\mathbf{BA}}$ -algebra over a chosen set V of propositional variables, as was discussed earlier. We wish to show precisely how this language is related to the abstract language defined by T , i.e. the free $\mathbf{FTU}$ -algebra (in $\mathbf{BA}$) over $\mathbf{FV}$. Let $\mathbf{F}_{T+\mathbf{BA}} : \mathbf{Set} \rightarrow \mathbf{Alg}_{\mathbf{Set}}(T + \mathbf{S}_{\mathbf{BA}})$ be the free functor defined by $T + \mathbf{S}_{\mathbf{BA}}$. Since both T and $\mathbf{S}_{\mathbf{BA}}$ are finitary, so is their coproduct, which is therefore a variator. Next, consider any set of equation axiomatizing boolean algebras, which we call $\mathbf{BA}$. We can assume without loss of generality, that the (three) variables used in these equations form a subset of V . We then proceed as above by considering the maps $e_1, e_2 : \mathbf{BA} \rightarrow \mathbf{UF}_{T+\mathbf{BA}}V$ defined by these equations, taking their adjoint transpose $\hat{e}_1, \hat{e}_2$, and then their coequalizer q . Finally, we form the fully invariant closure of $\ker q$, and get a regular epi $q_{\mathbf{BA}} : \mathbf{F}_{T+\mathbf{BA}}V \twoheadrightarrow Q_{\mathbf{BA}}$ such that $Q_{\mathbf{BA}} \perp q_{\mathbf{BA}}$.

Proposition 1.36. *The $T + \mathbf{S}_{\mathbf{BA}}$ -algebra $Q_{\mathbf{BA}}$ defined above is isomorphic to the free $\mathbf{FTU}$ -algebra over $\mathbf{FV}$.*

Proof. By ‘is isomorphic to’ we mean that by modifying its structure map, $Q_{\mathbf{BA}}$ (viewed as the carrier of a $T + \mathbf{S}_{\mathbf{BA}}$ -algebra structure) is the free $\mathbf{FTU}$ -algebra over $\mathbf{FV}$. To see this, we first need to show that it can be equipped with a $\mathbf{FTU}(-) + \mathbf{FV}$ structure. Next we’ll show that it is initial. By adjunction and the fact that $\mathbf{F}$ preserved coproducts, to be equipped with $\mathbf{FTU}(-) + \mathbf{FV}$ structure in $\mathbf{BA}$ is equivalent to being equipped with a $\mathbf{TU}(-) + \mathbf{V}$ structure in $\mathbf{Set}$. Since $Q_{\mathbf{BA}}$ is, by construction, a boolean algebra, we can consider its T -algebra structure map α (we forget the $\mathbf{S}_{\mathbf{BA}}$ -structure for now) as a map with signature $\mathbf{TU}Q_{\mathbf{BA}} \rightarrow \mathbf{U}Q_{\mathbf{BA}}$, and we can also endow it with a map $v : V \rightarrow \mathbf{U}Q_{\mathbf{BA}}$ defined as $v = q_{\mathbf{BA}} \circ \eta^{T+\mathbf{BA}}$ where $\eta^{T+\mathbf{BA}}$ is the unit of the adjunction $\mathbf{F}_{T+\mathbf{BA}} \dashv \mathbf{U}_{T+\mathbf{BA}} : \mathbf{Set} \rightarrow \mathbf{Alg}_{\mathbf{Set}}(T + \mathbf{BA})$. By using the adjunction $\mathbf{F} \dashv \mathbf{U}$, we then get that the dual transpose of $\alpha + v$ is a map

$$\widehat{\alpha + v} = \hat{\alpha} + \hat{v} : \mathbf{FTU}Q_{\mathbf{BA}} + \mathbf{FV} \rightarrow Q_{\mathbf{BA}}$$

where $Q_{\mathbf{BA}}$ is naturally viewed as a boolean algebra (since it satisfies all their equations). This shows that $Q_{\mathbf{BA}}$ is a $\mathbf{FTU}(-) + \mathbf{FV}$ -algebra.

To show that it is the initial such algebra, let us consider any other such algebra $\beta : \mathbf{FTU}B + \mathbf{FV} \rightarrow B$. Since B is a boolean algebra, we can define a trivial $\mathbf{S}_{\mathbf{BA}}$ -structure on $\mathbf{U}B$ by simply evaluating the boolean connectives, let us call this map $ev : \mathbf{S}_{\mathbf{BA}}\mathbf{U}B \rightarrow \mathbf{U}B$. By combining ev with the adjoint transpose $\tilde{\beta}$ of β , we can put a $T(-) + \mathbf{S}_{\mathbf{BA}}(-) + \mathbf{V}$ -algebra structure on $\mathbf{U}B$. We then get the following commutative diagram:

$$\begin{array}{ccccc}
 \mathbf{TF}_{T+\mathbf{BA}}V + \mathbf{S}_{\mathbf{BA}}V + V & \xrightarrow{Tq_{\mathbf{BA}} + \text{Id}_V} & \mathbf{TU}Q_{\mathbf{BA}} + \mathbf{S}_{\mathbf{BA}}\mathbf{U}Q_{\mathbf{BA}} + V & & \\
 \downarrow \langle - \rangle_V + \eta^{T+\mathbf{BA}} & \searrow T! & \swarrow & \downarrow \alpha + b + v & \\
 & & \mathbf{TU}B + \mathbf{S}_{\mathbf{BA}}B + V & & \\
 & & \downarrow ev + \tilde{\beta} & & \\
 \mathbf{F}_{T+\mathbf{BA}}V & \xrightarrow{q_{\mathbf{BA}}} & \mathbf{U}Q_{\mathbf{BA}} & & \\
 & \searrow ! & \swarrow & & \\
 & & \mathbf{U}B & &
 \end{array}$$

Where $!$ denotes the unique $T(-) + \mathbf{S}_{\mathbf{BA}} + \mathbf{V}$ -algebra morphism which exists by virtue of $\mathbf{F}_{T+\mathbf{BA}}$ being the initial such algebra, and b is the boolean part of the structure map of $Q_{\mathbf{BA}}$. Finally, since B is a boolean algebra, it is clear that $\mathbf{U}B$ must belong to the variety defined by the equations axiomatizing boolean algebras, and in particular we must have $\mathbf{U}B \perp q$, which proves the existence and uniqueness of the morphism from $\mathbf{U}Q_{\mathbf{BA}}$ to $\mathbf{U}B$, which in turns provides us by adjoint transpose with the unique $\mathbf{FTU}(-) + \mathbf{FV}$ -algebra morphism $Q_{\mathbf{BA}} \rightarrow B$. Note that the unique morphism $\mathbf{U}Q_{\mathbf{BA}} \rightarrow \mathbf{U}A$ is a boolean algebra morphism

(i.e. is of the form Uf for some f) by assumption on $q_{\mathbf{BA}}$. The U functor only really appears because the domain and codomain are viewed as carriers of algebra structures. $\square$

Note that the construction above also gives us an expression for the structure map of the abstract language, namely as the adjoint transpose of the T -algebra structure map of $Q_{\mathbf{BA}}$.

Before we can present the main result of this subsection, we need to briefly examine substitution instances in the concrete and abstract setting. In the concrete case, a substitution is defined uniquely by a map from the set V of propositional languages to the language, i.e.

$$\sigma : V \rightarrow U_{T+\mathbf{BA}} F_{T+\mathbf{BA}}(V)$$

or, equivalently, as an endomorphism of $\hat{\sigma} : F_{T+\mathbf{BA}}(V) \rightarrow F_{T+\mathbf{BA}}(V)$. By using the construction of Proposition 1.36 we can immediately extend this valuation to the abstract case.

Lemma 1.37. *For any endomorphism $\sigma : F_{T+\mathbf{BA}} V \rightarrow F_{T+\mathbf{BA}} V$, there exists an endomorphism $\tilde{\sigma} : Q_{\mathbf{BA}} = \mathbf{GF}V \rightarrow \mathbf{GF}V$ such that the following square commutes*

$$\begin{array}{ccc} F_{T+\mathbf{BA}} V & \xrightarrow{\sigma} & F_{T+\mathbf{BA}} V \\ q_{\mathbf{BA}} \downarrow & & \downarrow q_{\mathbf{BA}} \\ Q_{\mathbf{BA}} & \xrightarrow{\tilde{\sigma}} & Q_{\mathbf{BA}} \end{array}$$

Proof. We use the fact that $Q_{\mathbf{BA}} \perp q_{\mathbf{BA}}$ or, equivalently, the construction of $q_{\mathbf{BA}}$, to get the unique arrow by orthogonality:

$$\begin{array}{ccc} F_{T+\mathbf{BA}} V & \xrightarrow{q} & Q_{\mathbf{BA}} \\ & \searrow q \circ \sigma & \swarrow \tilde{\sigma} \\ & Q_{\mathbf{BA}} & \end{array}$$

$\square$

Let us denote by $\vdash_{\text{PL}}$ the derivability relation for predicate lifting style coalgebraic logics, and recall that $\vdash_{\text{ML}}$ is the derivability relation for abstract coalgebraic logics, as defined above. If $\mathbf{Ax}$ is a set of formulas in a predicate lifting language, let $[\mathbf{Ax}]$ denote the set $\{\top \vdash_{\text{ML}} [a]_{\sim_{\mathbf{BA}}} \mid a \in \mathbf{Ax}\}$ where $[a]_{\sim_{\mathbf{BA}}}$ is the class of terms which are $\mathbf{BA}$ -equivalent to a .

Theorem 1.38. *Let $T = S_{\Sigma+\mathbf{BA}}$ define a predicate lifting style language $\mathbf{L}_{\Sigma}$. If $\mathbf{Ax}$ defines a Hilbert system for $\mathbf{L}_{\Sigma}$ with axioms only, then for any $b \in \mathbf{L}_{\Sigma}$*

$$\mathbf{Ax} \vdash_{\text{PL}} b \quad \Leftrightarrow \quad \top \vdash_{\text{ML}} [b] \text{ can be derived from } [\mathbf{Ax}]$$

Proof. $\Rightarrow$: by induction on the size of the proof. For the base case, we must have that b is a substitution instance of an axiom in $\mathbf{Ax}$ or a tautology of propositional logic. If $b = \sigma a$ for some endomorphism $\sigma : \mathbf{L}_\Sigma \rightarrow \mathbf{L}_\Sigma$ and some $a \in A$, then it immediately follows from Lemma 1.37 that there exist an endomorphism $\tilde{\sigma}$ on $\mathbf{GFV}$ such that $\tilde{\sigma}([a]) = [\sigma(a)] = [b]$, i.e. $[b]$ is a substitution instance of an axiom in $[\mathbf{Ax}]$. If a is a tautology of propositional logic, then $[a] = [\top]$, which is an axiom in the abstract system.

For the inductive step, we go through the possible last steps of the derivation of a . We can ignore the rules of propositional calculus since they can assumed to be equivalent in both systems. If it was modus ponens, i.e.

$$(\text{M.P.}) \frac{c \quad c \rightarrow b}{b}$$

then by using the induction hypothesis, we have proof $\top \vdash_{\text{ML}} [c]$ and $\top \vdash_{\text{ML}} [c \rightarrow b]$, i.e. $\top \vdash_{\text{ML}} \neg[c] \vee [b]$, from which it follows by application of $(\neg I)$ that $[c] \vdash_{\text{ML}} [b]$, and thus $\top \vdash_{\text{ML}} [b]$ by an application of the Cut rule. If the last applied rule was the congruence rule, i.e.

$$(\text{Congruence}) \frac{a_i \leftrightarrow b_i, 1 \leq i \leq n}{\sigma(a_1, \dots, a_n) \leftrightarrow \sigma(b_1, \dots, b_n)} \quad (\sigma \in \Sigma)$$

then by using the induction hypothesis, we have a proof of $\top \vdash_{\text{ML}} [a_i] \leftrightarrow [b_i], 1 \leq i \leq n$. By application of $(\neg I), (\bigwedge L)$ and Cut it is easy to see that we have proofs of $[b_i] \vdash_{\text{ML}} [a_i]$, and symmetrically, we can derive proofs of $[a_i] \vdash_{\text{ML}} [b_i], 1 \leq i \leq n$. As we have described earlier, this can be viewed as a witness of the lifted relation

$$\overline{\text{FS}_\Sigma \mathbf{U}}(R \upharpoonright (\mathbf{B}_{\Sigma}([\sigma([a_1], \dots, [a_n])]) \times \mathbf{B}_{\Sigma}([\sigma([b_1], \dots, [b_n])])))$$

where $R \subset \mathbf{GFV} \times \mathbf{GFV}$ is defined by $[a]R[b] \Leftrightarrow [a] \vdash_{\text{ML}} [b]$, and we can then apply the abstract congruence rule of the $\vdash_{\text{ML}}$ system to $\alpha = ([a_1], \dots, [a_n]), \beta = ([b_1], \dots, [b_n])$

$$(\text{Congruence}) \frac{(\alpha, \beta) \in \overline{\text{FS}_\Sigma \mathbf{U}}(R \upharpoonright \mathbf{B}_L(\alpha) \times \mathbf{B}_L(\beta))}{(\langle \alpha \rangle, \langle \beta \rangle) \in R}$$

to conclude $\langle \alpha \rangle \vdash_{\text{ML}} [\langle \beta \rangle]$ and $[\langle \beta \rangle] \vdash_{\text{ML}} [\langle \alpha \rangle]$, where $\langle \alpha \rangle = [\sigma([a_1], \dots, [a_n])]$ and similarly for $\langle \beta \rangle$. From this $\top \vdash_{\text{ML}} \langle \alpha \rangle \leftrightarrow \langle \beta \rangle$ then follows easily from the propositional rules.

$\Rightarrow$: the opposite direction is shown in exactly the same way but in reverse and hinges on the fact that $\vdash_{\text{PL}}$ axiomatizes all propositionally valid formulas. In particular we are free to choose any representative a in an equivalence class $[a]$, since if we choose another representative a' then there must exist a proof $\vdash_{\text{PL}} a \leftrightarrow a'$. $\square$

We have thus shown that the two-dimensional $\vdash_{\text{ML}}$ - system for abstract coalgebraic system subsumes the more familiar proof system of coalgebraic logic in the predicate lifting format. It is also clear that it generalises the $\text{KKV}(T)$ proof system of the nabla style of predicate logic.

Chapter 2

Algebraic canonicity

This chapter will be almost purely algebraic, and will pursue a line of inquiry which dates back to Tarski and Jónsson's seminal work on boolean Algebras with Operators [JT51], namely the study of modal logics through algebraic means (although [JT51] does not explicitly mention modal logics). The idea behind this approach is simple: since propositional logic can be given an algebraic semantic in terms of boolean algebras, it is reasonable to give an algebraic semantic to modal logics - broadly construed as in the prolegomenon - in terms of boolean algebras extended with some additional maps to interpret the modal operators. This approach is very natural, and just as the language of propositional logic can be seen as the free BA over a set V of propositional variables, modal languages can also be seen as free extended boolean algebras over V . Thus from a universal algebra perspective, the algebraic study of modal logic is the study of varieties defined by adding equations to these free constructions. In this way, algebraic modal logic offers an algebraic handle on the main concepts and problems of modal logic, and the reader is referred to the classic [BdRV01] for an introduction and overview to the subject and to [Ven06] for more details.

The particular problem which we wish to tackle algebraically in this chapter, and to some extent throughout this thesis, is that of canonicity. A precursor to a systematic algebraic study of canonicity can clearly be found in [JT51]. This area of research then followed two complementary paths. Very crudely, the first path is concerned with the canonicity of *varieties* (through e.g. closure properties), whilst the second path investigates syntactic criteria for *formulas* to define canonical varieties. The first line of research was developed by Goldblatt who studied canonicity in the wider context of relation, cylindric and modal algebras as a way to understand the connection between first-order logic and equational logic. In [Gol89, Gol95], he investigated when an elementary relational structure (i.e. a relational structure defined by a first order formula)

gives rise to a canonical variety. This wide-ranging and powerful line of research uses dualities and many ideas for universal algebra in the spirit of the famous Golblatt-Thomason theorem of modal logic. The other track, which we follow here, can be traced back to the near simultaneous papers of Jónsson [Jón94] and de Rijke and Venema [dRV95]. The focus of both papers is to identify classes of formulas that define canonical varieties, and in particular to prove the classical Sahlqvist completeness result of [Sah75] algebraically (see [BdRV01] for the ‘traditional’ treatment of the subject).

Since then, the main focus of investigations into this algebraic theory of canonicity has been to weaken or modify the boolean core of the theory or to consider other types of extensions than the canonical extensions. The work of Gehrke for example has extended much of the theory to bounded distributive lattice expansions [GJ94, GJ04] and bounded lattices expansions [GH01]. Sahlqvist results in these settings can be found in [GNV05]. The thesis of Teheux [Teh08] examines canonicity in the context of multi-valued modal algebras, i.e. the boolean part of the theory is generalised to MV-algebras. Another strand of research can be found in Givant and Venema’s [GV99] where the Dedekind-Mac Neille completion is considered rather than the canonical extension, or in the work of Vosmaer who studies canonical extensions of various lattice structures in relation to profinite completions (see [Vos10] for an overview). Finally, canonicity for lattices has been used to prove completeness with respect to relational semantics for substructural logics, see for example [DGP05, Geh06, Suz11, CGvR11].

In this chapter we will study canonical extensions of BAEs and DLEs whose expansions satisfy a large range of algebraic properties, with the ambition of getting some new canonicity results for non-normal modal logics. The approach will be very similar to that taken in [GJ04], but with a focus on covering a wide range of types of expansions and providing methods for building canonical (in)equations. In particular we aim to present a theory that includes n -ary maps which are either isotone or antitone in their arguments and which includes the notion of k -additivity of [Hen70].

2.1 Canonical Extensions

2.1.1 Canonical extension of distributive lattices and boolean algebras

We now briefly present the salient facts about canonical extensions. For more details the reader is referred to [GH09] for BAs, [JT51] for BAOs, and [GJ94, GJ04] for DLEs. The main rationale for studying canonical extensions is to embed a lattice-based structure, typically a language quotiented by some axioms, into

a similar structure which is more ‘set-like’, i.e. whose elements can be viewed as parts of a set, or of a set with some additional structure. In this way, we can hope to establish a connection between the syntax and the semantics, i.e. build models from formulas. But what does being ‘set-like’ mean? Two criteria emerge as being fundamental: completeness and being generated from below (i.e. by joins) by something akin to ‘elements’. Canonical extensions satisfy these two conditions. Our approach here will be purely algebraic, i.e. we will not make use of the duality theory of distributive lattices or boolean algebras.

We start with distributive lattices. For a distributive lattice $\mathbb{A}$, the idea behind the construction of its canonical extension $\mathbb{A}^\sigma$ is to build a completion of $\mathbb{A}$ which is not ‘too big’ and not ‘too different’ from $\mathbb{A}$. Technically, we want $\mathbb{A}$ to be **dense** and **compact** in $\mathbb{A}^\sigma$.

Density. To build a completion of $\mathbb{A}$ it is natural to formally add to $\mathbb{A}$ all meets, all joins, all meets of all joins, all joins of all meets, etc. In the case of the canonical extension we require that this procedure stops after two iterations (i.e. we want a Δ_1 -completion, see [GJP13]). Intuitively, this prevents the completion from becoming ‘too big’. Based on this intuition we introduce the following terminology. Given a sub-lattice $\mathbb{A}$ of a complete distributive lattice $\mathbb{C}$, we define the meets in $\mathbb{C}$ of elements of $\mathbb{A}$ as the **closed elements** of $\mathbb{C}$ and denote the set of closed elements as $K(\mathbb{A})$ (or simply K when there is no ambiguity). The sub-lattice $\mathbb{A}$ is said to be **meet-dense** in $\mathbb{C}$ if $K(\mathbb{A}) = \mathbb{C}$, i.e. if every element is closed. Dually, we define the joins in $\mathbb{C}$ of elements of $\mathbb{A}$ as the **open elements** of $\mathbb{C}$ and denote the set of open elements as $O(\mathbb{A})$. The sub-lattice $\mathbb{A}$ is said to be **join-dense** in $\mathbb{C}$ if $O(\mathbb{A}) = \mathbb{C}$, i.e. if every element is open. Finally, we say that $\mathbb{A}$ is **dense** in $\mathbb{C}$ if

$$\mathbb{C} = O(K(\mathbb{A})) = K(O(\mathbb{A}))$$

It is natural to describe the elements of $\mathbb{A}$ as clopens in $\mathbb{C}$ since they are both open and closed.

Compactness. The canonical extension $\mathbb{A}^\sigma$ of $\mathbb{A}$ is also required not to be too different from $\mathbb{A}$ in the sense that facts about arbitrary meets and joins of elements of $\mathbb{A}$ in $\mathbb{A}^\sigma$ must already be ‘witnessed’ by finite meets and joins in $\mathbb{A}$. Formally, if $\mathbb{A}$ is a sub-lattice of $\mathbb{C}$, we say that $\mathbb{A}$ is **compact** in $\mathbb{C}$ if for every $X, Y \subseteq \mathbb{A}$ such that

$$\bigwedge X \leq \bigvee Y$$

there exists *finite* subsets $X_0 \subseteq X, Y_0 \subseteq Y$ such that

$$\bigwedge X_0 \leq \bigvee Y_0$$

An equivalent definition is that $\mathbb{A}$ is compact in $\mathbb{C}$ if for every closed element $p \in K(\mathbb{A})$ and open element $u \in O(\mathbb{A})$ such that $p \leq u$, there exists an element $a \in \mathbb{A}$ such that $p \leq a \leq u$.

We are now ready to define the **canonical extension** $\mathbb{A}^\sigma$ of a distributive lattice $\mathbb{A}$ as the complete distributive lattice such that $\mathbb{A}$ is dense and compact in $\mathbb{A}^\sigma$. We will show that canonical extensions exist and are unique up to isomorphism. We start with the following lemmas concerning infinite distributivity in canonical extensions.

Lemma 2.1. *Let $\mathbb{A}^\sigma$ be the canonical extension of a distributive lattice $\mathbb{A}$, $U \subseteq O$ be a set of open elements and $P \subseteq K$ be a set of closed elements, then for any $x \in \mathbb{A}^\sigma$*

$$1. x \wedge \bigvee U = \bigvee \{x \wedge u \mid u \in U\}$$

$$2. x \vee \bigwedge P = \bigwedge \{x \vee p \mid p \in P\}$$

Proof. We show the first distribution law, the second follows dually. Clearly, since $x \wedge u \leq a \wedge \bigvee U$ for each $u \in U$, we have $\bigvee \{x \wedge u \mid u \in U\} \leq x \wedge \bigvee U$. For the opposite inequality, we will show that if a closed element $p \leq x \wedge \bigvee U$, then $p \leq \bigvee \{x \wedge u \mid u \in U\}$, and the result will then follow by density. If $p \leq x \wedge \bigvee U$, then $p \leq x$ and $p \leq \bigvee U$, and since $\bigvee U$ is a set of open elements, compactness gives us a finite subset $U_0 \subseteq U$ such that $p \leq \bigvee U_0$. In consequence, by using the fact that $\mathbb{A}^\sigma$ is distributive

$$p \leq x \wedge \bigvee U_0 = \bigvee \{x \wedge u \mid u \in U_0\} \leq \bigvee \{x \wedge u \mid u \in U\}$$

which is what we wanted to show. $\square$

Using a very similar proof, we can also show complete distributivity for directed sets of closed and open elements.

Lemma 2.2. *Let $\mathbb{A}^\sigma$ be the canonical extension of a distributive lattice $\mathbb{A}$, $U_i, i \in I$ be collection of an upward-directed sets of open elements and $P_j, j \in J$ be a collection of downward-directed sets of closed elements, then*

$$1. \bigwedge_i \bigvee U_i = \bigvee \{\bigwedge \text{im} \phi \mid \phi : I \rightarrow O, \phi(i) \in U_i\}$$

$$2. \bigvee_j \bigwedge P_j = \bigwedge \{\bigvee \text{im} \phi \mid \phi : J \rightarrow K, \phi(j) \in P_j\}$$

Proof. See Lemma 3.2. of [GH01]. $\square$

We can now show the following important infinite distribution laws.

Lemma 2.3. *Let $\mathbb{A}^\sigma$ be the canonical extension of a distributive lattice $\mathbb{A}$, U be a set of open elements and P be a set of closed elements, then for any $q \in K, v \in O$*

$$1. q \wedge \bigvee P = \bigvee \{q \wedge p \mid p \in P\}$$

$$2. \ v \vee \bigwedge U = \bigwedge \{v \vee u \mid u \in U\}$$

Proof.

$$\begin{aligned} q \wedge \bigvee P &= q \wedge \bigvee \{ \bigwedge \{a \in A \mid a \geq p\} \mid p \in P \} && \text{Density} \\ &= q \wedge \bigwedge \{ \bigvee \text{im} \phi \mid \phi : P \rightarrow A, \phi(p) \geq p \} && \text{Lemma 2.2} \\ &= \bigwedge \{ q \wedge \bigvee \text{im} \phi \mid \phi : P \rightarrow A, \phi(p) \geq p \} \\ &= \bigwedge \{ \bigvee \{ q \wedge a \mid a \in \text{im} \phi \} \mid \phi : P \rightarrow A, \phi(p) \geq p \} && \text{Lemma 2.1} \end{aligned}$$

Note that we have used Lemma 2.2 because each $\{a \in A \mid a \geq p\}$ is a down-directed set of clopen - and thus closed - elements, and Lemma 2.1 because each $\text{im} \phi$ is open. We want to show that the last term of the derivation is equal to

$$\bigwedge \{ \bigvee \text{im} \psi \mid \psi : P \rightarrow A, \psi(p) \geq q \wedge p \}$$

For every ϕ , we can define the following $\psi : P \rightarrow A, p \mapsto q \wedge \phi(p)$ and we get

$$\bigvee \{ q \wedge a \mid a \in \text{im} \phi \} = \bigvee \text{im} \psi$$

Thus

$$\begin{aligned} \bigwedge \{ \bigvee \text{im} \psi \mid \psi : P \rightarrow A, \psi(p) \geq q \wedge p \} &\leq \\ \bigwedge \{ \bigvee \{ q \wedge a \mid a \in \text{im} \phi \} \mid \phi : P \rightarrow A, \phi(p) \geq p \} \end{aligned}$$

For the opposite inequality, we start with a ψ which picks a clopen above each element $q \wedge p$, and we want to define a ϕ which picks a clopen above each p and satisfies $\bigvee \{ q \wedge a \mid a \in \text{im} \phi \} = \bigvee \text{im} \psi$. To define such a ϕ , we choose for $p \in P$ any clopen a' above p such that $q \wedge a' \leq a = \psi(p)$. To see that such an element exists, we proceed by contraposition. Assume that no such element exist, i.e. that for all $a' \geq p$ it is the case that $a' \wedge q \not\leq a$, i.e. that for all $a' \geq p$, $a \vee (a' \wedge q) \not\geq a$. But then

$$\begin{aligned} &\bigwedge \{ a \vee (a' \wedge q) \mid a' \geq p \} \not\geq a \\ \Leftrightarrow &a \vee \bigwedge \{ a' \wedge q \mid a' \geq p \} \not\geq a && \text{Lemma 2.1 (each } a' \wedge q \text{ if closed)} \\ \Leftrightarrow &a \vee (q \wedge \bigwedge \{ a' \mid a' \geq p \}) \not\geq a \\ \Leftrightarrow &a \vee (q \wedge p) \not\geq a && p \text{ is closed} \end{aligned}$$

which contradicts the hypothesis that $a = \psi(p) \geq q \wedge p$. Thus for each p we can choose a clopen a' above it such that $q \wedge a' \leq a = \psi(p)$, and this defines a choice function $\phi : P \rightarrow A$ such that

$$\bigvee \text{im} \psi = \bigvee \{ q \wedge a' \mid a' \in \text{im} \phi \}$$

and we thus get inequality

$$\bigwedge \{ \bigvee \{ q \wedge a \mid a \in \text{im} \phi \} \mid \phi : P \rightarrow A, \phi(p) \geq p \} \leq \\ \bigwedge \{ \bigvee \text{im} \psi \mid \psi : P \rightarrow A, \psi(p) \leq q \wedge p \}$$

Thus we have

$$\begin{aligned} q \wedge \bigvee P &= \bigwedge \{ \bigvee \text{im} \psi \mid \psi : P \rightarrow A, \psi(p) \geq q \wedge p \} \\ &= \bigvee \{ \bigwedge \{ a \in A \mid a \geq q \wedge p \} \mid p \in P \} && \text{Lemma 2.2} \\ &= \bigvee \{ q \wedge p \mid p \in P \} && p \wedge q \text{ is closed} \end{aligned}$$

The other distribution rule follows dually. $\square$

There are two explicit representations of the canonical extension of a distributive lattice. The first one requires the Prime Filter Theorem which is a non-constructive principle, albeit strictly weaker than the axiom of choice (see [HL71]). The second is constructive and based on Galois connections (see [GH01]) but is not compatible with the adjunction $\mathbf{Pf} \dashv \mathcal{U} : \mathbf{DL} \rightarrow \mathbf{Pos}$ and the process of model building which we develop in Chapter 5. We will therefore follow the more traditional non-constructive approach. The key to representing canonical extensions is to show that certain classes of elements of the abstractly defined canonical extension $\mathbb{A}^\sigma$ can be realized as familiar constructions on the original lattice $\mathbb{A}$.

Lemma 2.4. *Let $\mathbb{A}^\sigma$ be the canonical extension of a distributive lattice $\mathbb{A}$, then the sub-lattice K of closed elements of $\mathbb{A}^\sigma$ is isomorphic to the lattice of filters of $\mathbb{A}$. Dually, the sub-lattice O of open elements of $\mathbb{A}^\sigma$ is isomorphic to the lattice of ideals of $\mathbb{A}$.*

Proof. For a full proof, see Lemma 3.3. of [GH01]. The isomorphism maps a closed element of $\mathbb{A}^\sigma$ to the set of elements of $\mathbb{A}$ (clopens) above it, and, in the opposite direction, a filter of $\mathbb{A}$ to its meet in $\mathbb{A}^\sigma$. Dually, open elements are mapped to the set of elements of $\mathbb{A}$ below it, and ideals of $\mathbb{A}$ to their joins. $\square$

Since every element of the canonical extension $\mathbb{A}^\sigma$ is a join of closed element, we could try to represent arbitrary elements as formal joins of filters of $\mathbb{A}$. However, using the Prime Filter Theorem, it is possible to isolate a subset of the closed elements which is join-dense in $\mathbb{A}^\sigma$ and has a much more natural logical interpretation.

Theorem 2.5 (Prime Filter Theorem for distributive lattices). *Suppose $\mathbb{A}$ is a distributive lattice, I is an ideal of $\mathbb{A}$, F is a filter of $\mathbb{A}$, and $I \cap F = \emptyset$. then there exists a prime filter G such that $F \subseteq G$ and $I \cap G = \emptyset$*

Proof. See Section III.4 Theorem 1 in [BD75]. $\square$

Recall that an element x of a lattice L is called **join irreducible** if it cannot be written as a (non-trivial) join of elements of L , i.e. if $x = y_1 \vee y_2, y_1, y_2 \in L$ implies $x = y_1$ or $x = y_2$. Similarly, x is said to be **join prime** if $x \leq y_1 \vee y_2$ implies that $x \leq y_1$ or $x \leq y_2$. Every join prime element is join irreducible and the converse can easily be seen to hold if the lattice is distributive. If a lattice L is complete, we can generalise these concepts as follows: we will say that $x \in L$ is **completely join irreducible** if $x = \bigvee Y$ implies $x = y$ for some $y \in Y$, and that x is **completely join prime** if $x \leq \bigvee Y$ implies $x \leq y$ for some $y \in Y$. It is conventional to denote the set of completely join irreducible elements of a complete lattice L as $J(L)$. The following corollary of Lemma 2.3 shows that the two concepts coincide in canonical extensions of distributive lattices.

Lemma 2.6. *Let $\mathbb{A}^\sigma$ be the canonical extension of a distributive lattice $\mathbb{A}$, the completely join irreducible elements of $\mathbb{A}^\sigma$ are completely join prime.*

Proof. Let p be completely join irreducible. Clearly, since every element of $\mathbb{A}^\sigma$ is a join of closed element, p must be closed. Assume now that $p \leq \bigvee X$ for some $X \subseteq \mathbb{A}^\sigma$. By density, we can assume w.l.o.g. that X is a set of closed elements. We now use Lemma 2.3 to conclude that

$$p = p \wedge \bigvee X = \bigvee \{p \wedge x \mid x \in X\}$$

and thus $p \leq x$ for some $x \in X$, since it is completely join irreducible. $\square$

Lemma 2.7 ([GH01]). *Let $\mathbb{A}^\sigma$ be the canonical extension of a distributive lattice $\mathbb{A}$, then the poset $J(\mathbb{A}^\sigma)$ of completely join irreducible elements of $\mathbb{A}^\sigma$ is isomorphic to $(\text{Pf}\mathbb{A})^{\text{op}}$, the poset of prime filters under reverse set-theoretic ordering. Moreover, every element of $\mathbb{A}^\sigma$ can be written as a join of elements of $J(\mathbb{A}^\sigma)$.*

Proof. The isomorphism in **Pos** works as follows: we define $\phi : J(\mathbb{A}^\sigma) \rightarrow (\text{Pf}\mathbb{A})^{\text{op}}, p \mapsto \{a \in \mathbb{A} \mid p \leq a\}$ and $\psi : (\text{Pf}\mathbb{A})^{\text{op}} \rightarrow J(\mathbb{A}^\sigma), F \mapsto \bigwedge F$. Clearly if $F \subseteq F'$ then $\psi(F) \geq \psi(F')$, and conversely if $p \leq p'$ then $\phi(p) \supseteq \phi(p')$, hence the $(-)^{\text{op}}$ order. It is not difficult to see that if both maps are well-defined, then $\psi \circ \phi(p) = p$. For the opposite direction we clearly have $F \subseteq \{a \in \mathbb{A} \mid a \geq \bigwedge F\}$, to show the equality we use compactness: since a is clopen it is in particular open and thus $\bigwedge F \leq a$ is witnessed by a finite set F_0 of elements of F , i.e. $\bigwedge F_0 \leq a$, but since F is a filter it follows that $\bigwedge F_0 \in F$ and thus $a \in F$. Thus ψ, ϕ are inverse of each other provided that they are well-defined. The fact that ϕ is well-defined is easy: $\phi(p)$ is clearly a filter, moreover it is prime because p is join irreducible. To see that $\psi(F) = \bigwedge F$ is completely join irreducible when F is prime, let $\bigwedge F = \bigvee M$, where we can assume without loss of generality that

M is a set of closed elements. Clearly $m \leq \bigwedge F$ each $m \in M$, thus if we want to proceed by contradiction we should assume that $m < \bigwedge F$ for each $m \in M$. But then $\{a \in \mathbb{A} \mid a \geq m\} \supsetneq F$ since m is closed, and for each $m \in M$ we can therefore select an $a_m \geq m$ not in F . It follows that

$$\bigwedge F \leq \bigvee M \leq \bigvee_{m \in M} a_m$$

which by compactness means that there must exist a finite set M_0 of a_m s such that $\bigvee_{m \in M_0} a_m \in F$, a contradiction since F is prime.

The second part of the lemma requires the Prime Filter Theorem for distributive lattices, and is central to the ‘representability’ of the canonical extension. Since every element of $\mathbb{A}^\sigma$ is a join of closed elements, it is enough to show the property for closed elements. To show the property we’ll show that if $k \neq k', k, k' \in K$ then there exists a completely join irreducible element p witnessing the non-equality from below, i.e. such that $p \leq k, p \not\leq k'$. Note that by density, if $k \neq k'$ there must exist an open set witnessing the inequality from above, i.e. $u \in O$ s.th. $k \not\leq u, k' \leq u$. Using Lemma 2.4, we know that k uniquely defines a filter $F(k) = \{a \in \mathbb{A} \mid k \leq a\}$ and u uniquely defines an ideal $I(u) = \{a \in \mathbb{A} \mid a \leq u\}$. If $a \in F(k) \cap I(k)$ then $k \leq a$ and $a \leq u$, and thus $k \leq u$ which contradicts the assumption that $k \not\leq u$. Thus $F(k) \cap I(u) = \emptyset$, and we can apply the Prime Filter Theorem to get a prime filter $P \supseteq F(k)$ such that $P \cap I(u) = \emptyset$. By the first part of the proof, and Lemma 2.4, we can map this prime filter back to $\mathbb{A}^\sigma$ to get a completely join irreducible element p such that

$$p = \bigwedge P \leq \bigwedge F(k) = k \not\leq \bigvee I(u) = u$$

as desired. $\square$

Theorem 2.8. *The canonical extension of a distributive lattice is given by*

$$\mathbb{A}^\sigma \simeq \mathcal{UPf}\mathbb{A}$$

and is unique up to isomorphism.

Proof. Using Lemma 2.7 it suffices to show that $\mathbb{A}^\sigma \simeq \mathcal{DJ}(\mathbb{A}^\sigma)$, where $\mathcal{D}$ is the down-set functor. We define

$$\phi : \mathbb{A}^\sigma \rightarrow \mathcal{DJ}(\mathbb{A}^\sigma) : x \mapsto \{p \in J(\mathbb{A}^\sigma) \mid p \leq x\}$$

Conversely, we define

$$\psi : \mathcal{DJ}(\mathbb{A}^\sigma) \rightarrow \mathbb{A}^\sigma, X \mapsto \bigvee X$$

It is clear that ϕ preserves meets, and it preserves joins by the fact that each p is join prime. The fact that ψ preserve joins is trivial and the fact that it

preserves meets follows from the fact that $\mathbb{A}^\sigma$ is distributive. Thus ϕ, ψ are **DL**-morphisms. From the fact that every element of $\mathbb{A}^\sigma$ is a join of element of $J(\mathbb{A}^\sigma)$ (by Lemma 2.7), it is clear that $\psi \circ \phi(x) = x$ for all $x \in \mathbb{A}^\sigma$. Conversely, since every join irreducible is join prime (by Lemma 2.6), we have

$$\begin{aligned} \phi \circ \psi(X) &= \phi(\bigvee X) \\ &= \{p \in J(\mathbb{A}^\sigma) \mid p \leq \bigvee X\} \\ &= \{p \in J(\mathbb{A}^\sigma) \mid p \leq x \text{ for some } x \in X\} \\ &= X \end{aligned}$$

where the last step uses the fact that X is a down-set.

The unicity follows from the fact that every element in the canonical extension can be written as a join of completely join irreducibles, and that these are in bijective correspondence with the prime filters of $\mathbb{A}$. This allows us to build a bijection between the sets of completely join irreducibles of any two canonical extensions, and thus between any two canonical extensions. $\square$

Both functors **Pf** and $\mathcal{U}$ arise from the fact that $\mathbb{2}$ is both a distributive lattice and a poset, and $\mathbf{Pf}\mathbb{A}$ is just the set $\mathbf{Hom}_{\mathbf{DL}}(\mathbb{A}, \mathbb{2})$ with the pointwise partial ordering, whilst $\mathcal{U}X$ is the natural lattice structure on $\mathbf{Hom}_{\mathbf{Pos}}(X, \mathbb{2})$. An alternative representation of $\mathbb{A}^\sigma$ is thus given by the set of monotone maps from $\mathbf{Pf}\mathbb{A}$ (or, equivalently $J(\mathbb{A}^\sigma)$ by Lemma 2.7) into $\mathbb{2}$ via the inverse image of 0. In the case of the boolean algebras to which we will now turn our attention, the down-sets functor becomes the power set functor which can also be characterized as inverse images under characteristic functions.

Each BA $\mathbb{A}$ is a distributive lattice, and can thus be embedded into a unique (up to isomorphism) complete boolean algebra $\mathbb{A}^\sigma$ in which it is both dense and compact. Since prime filters in a boolean algebra are always maximal, i.e. ultrafilters, it follows that the partial order on $\mathbf{Pf}\mathbb{A} = \mathbf{Uf}\mathbb{A}$ is discrete, i.e. each ultrafilter is only related to itself, and thus $\mathbf{Uf}\mathbb{A}$ is really just a set. For this discrete partial order any map into $\mathbb{2}$ is monotone, which explains why the adjunction $\mathbf{Pf} \dashv \mathcal{U} : \mathbf{DL} \rightarrow \mathbf{Pos}$ becomes $\mathbf{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}$.

The following results specialise the canonical extension construction to the case of boolean algebras.

Proposition 2.9. *An element of a boolean algebra is an atom iff it is completely join irreducible.*

Proof. See Theorem 8.10.2. of [DH01]. $\square$

Since every element of the canonical extension can be written as a join of completely join irreducibles (Lemma 2.7), the previous Proposition shows that the canonical extension of a boolean algebra is *atomic*. Thus the canonical

extension of a boolean algebra is a Complete Atomic Boolean Algebra (CABA). Moreover, we have:

Theorem 2.10 ([GH09] Ch. 14). *A complete boolean algebra is atomic iff it is completely distributive.*

In particular, the canonical extension of a boolean algebra is always completely distributive. In fact, this holds for any distributive lattice since the canonical extension $\mathbb{A}^\sigma$ of a distributive lattice $\mathbb{A}$ is isomorphic to $\mathcal{UPf}\mathbb{A}$ which is completely distributive.

2.1.2 Canonical extension of expansions

We now turn our attention to maps defined on distributive lattices and boolean algebras, fix some terminology and define their canonical extensions. We will use the same convention as in Section 1.3 and use $\mathbf{A}$ to denote either of the categories **DL**, **BDL** or **BA**. For any object $\mathbb{A}$ in $\mathbf{A}$ with underlying set A , will say that a function $f : A^m \rightarrow A$ is **isotone in its i^{th} argument** if whenever $a_i \leq b_i$

$$f(a_1, \dots, a_{i-1}, a_i, a_{i+1}, \dots, a_m) \leq f(a_1, \dots, a_{i-1}, b_i, a_{i+1}, \dots, a_m)$$

Similarly, we will say that f is **antitone in its i^{th} argument** if whenever $a_i \leq b_i$

$$f(a_1, \dots, a_{i-1}, a_i, a_{i+1}, \dots, a_m) \geq f(a_1, \dots, a_{i-1}, b_i, a_{i+1}, \dots, a_m)$$

A map f is **monotone in its i^{th} argument** if it is either isotone or antitone in this argument. Finally, we will say that f is isotone (resp. antitone, resp monotone) if it is isotone (resp. antitone, resp. monotone) in *all* of its arguments, and we will say that f is **mixed** if it has both isotone and antitone arguments, i.e. monotone but neither isotone nor antitone.

The existence and unicity of canonical extensions for BAs was extended in [JT51] to any BAO $\mathcal{A}$ by defining the **extension of an operator** f of arity n as either of the following alternatives

$$f^\sigma(x) = \bigvee_{x \geq y \in K^n} \bigwedge_{y \leq z \in A^n} f(z) \quad (2.1)$$

$$f^\pi(x) = \bigwedge_{x \leq u \in O^n} \bigvee_{u \geq z \in A^n} f(z) \quad (2.2)$$

Note that the original definition in [JT51, Jón94] was Eq. (2.1), the dual possibility of Eq. (2.2) is defined for example in [GH01, Ven06]. A map f is called **smooth** if $f^\sigma = f^\pi$, we will return to this notion shortly. The canonical extension of a BAO $\mathcal{A}$ is traditionally defined as $\mathcal{A}^\sigma = (\mathbb{A}^\sigma, (f_s^\sigma)_{s \in \Sigma})$. In fact,

the definitions of f^σ and f^π make sense not just for operators, but more generally for any map which is isotone. However, for a completely general function $f : A^n \rightarrow A$, or even for an antitone map, the definitions Eqs. (2.1) and (2.2) do not make sense. We therefore need a more powerful definition of map extension. Such a definition was developed for the study of canonical extensions of various classes of lattices, see e.g. [GJ94, GH01, GJ04, Ven06]. Given an object in $\mathbf{A}$ with underlying set A , the σ - and π - extensions of a map $f : A^n \rightarrow A$ are the maps $(A^\sigma)^n \rightarrow A^\sigma$ defined by:

$$f^\sigma(x) = \bigvee \{ \bigwedge f[d, u] \mid K^n \ni d \leq x \leq u \in O^n \} \quad (2.3)$$

$$f^\pi(x) = \bigwedge \{ \bigvee f[d, u] \mid K^n \ni d \leq x \leq u \in O^n \} \quad (2.4)$$

where A^σ is the set underlying $\mathbb{A}^\sigma$ and $f[d, u] = \{f(a) \mid a \in A^n, d \leq a \leq u\}$. Note that $f[d, u]$ is always non-empty by compactness, which is the reason for choosing this type of ‘intervals’. It is straightforward to check that if f is isotone Eqs. (2.1) and (2.3) coincide, and similarly for the dual definitions of Eqs. (2.2) and (2.4). We can now define the canonical extension of a DLE (resp. BAE) $\mathcal{A} = (\mathbb{A}, (f_s)_{s \in \Sigma})$ as the DLE (resp. BAE) $\mathcal{A}^\sigma = (\mathbb{A}^\sigma, (f_s^\sigma)_{s \in \Sigma})$. This choice is purely arbitrary since using the $(\cdot)^\pi$ extension would work too, but it is conventional. In fact the choice of which extension to use can be guided, as we shall see later, by the algebraic properties of the maps which we want to extend, since some properties are preserved by the $(\cdot)^\sigma$ extension only, and others by the $(\cdot)^\pi$ extension, so having two possible definitions can be turned into an advantage. Let us start by examining some of the basic properties of these extensions.

Proposition 2.11. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A , and let $f : A^n \rightarrow A$ be any map, then*

$$(i) \quad f^\sigma \upharpoonright A^n = f^\pi \upharpoonright A^n = f$$

$$(ii) \quad f^\sigma \leq f^\pi \text{ under the pointwise order}$$

Proof. The proof can be found in for example [GH01] or [Ven06] for the case of unary operators and is readily adapted to the n -ary case by considering tuples of arguments. $\square$

If f is monotone we can say a bit more.

Proposition 2.12. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ and let $f : \mathbb{A}^n \rightarrow \mathbb{A}$ be a monotone map, then*

$$f^\sigma \upharpoonright (K \cup O)^n = f^\pi \upharpoonright (K \cup O)^n$$

Proof. The case where f is isotone is shown in for example [GH01, Ven06], so we just show the antitone case. Let us start by assuming that f is unary. We

know from Proposition 2.11 that $f^\sigma(x) \leq f^\pi(x)$ for any $x \in A^\sigma$, so we only need to show the opposite inequality. Start with $p \in K$, then for all $p \leq a \in A$ we have by antitonicity $f(a) = f^\sigma(a) \leq f^\sigma(p)$ and since this holds for any such a we have

$$\bigvee \{f(a) \mid p \leq a \in A\} \leq f^\sigma(p)$$

i.e. $f^\pi(p) \leq f^\sigma(p)$ by Eq. (2.4) applied to an antitone map. Similarly, for $u \in O$ we have for any $u \geq a \in A$ that $f^\pi(u) \leq f^\pi(a) = f(a)$ and thus

$$f^\pi(u) \leq \bigwedge \{f(a) \mid u \geq a \in A\}$$

i.e. $f^\pi(u) \leq f^\sigma(u)$ by Eq. (2.3) applied to an antitone map.

The case where f is n -ary is treated by first considering tuples consisting of open elements for the isotone arguments and closed elements for the antitone arguments. From there the general case is simply computed. We show how this is done in the case where $f : A^2 \rightarrow A$ is binary, isotone in the first argument and antitone in the second. By Proposition 2.11, we know that for any $x, y \in A^\sigma$, $f^\sigma(x, y) \leq f^\pi(x, y)$. We briefly show the opposite inequality for $u \in O, p \in K$. Let $a_1, a_2 \in A$ such that $a_1 \leq u, p \leq a_2$, then we have

$$f(a_1, a_2) = f^\sigma(a_1, a_2) \leq f^\sigma(u, p)$$

and since this holds for any such a_1, a_2 we get

$$f^\pi(u, p) = \bigvee_{a_1 \leq u, p \leq a_2} f(a_1, a_2) \leq f^\sigma(u, p)$$

A completely analogous dual proof shows that $f^\pi(p, u) \leq f^\sigma(p, u)$. Consider now, $p, q \in K$, from what we have just shown and the definitions of Eqs. 2.3 and 2.4, we get

$$\begin{aligned} f^\pi(p, q) &= \bigwedge_{p \leq u_1 \in O} \bigvee_{\substack{u \geq a_1 \in A \\ q \leq a_2 \in A}} f(a_1, a_2) \\ &= \bigwedge_{p \leq u \in O} f^\pi(u, q) \\ &= \bigwedge_{p \leq u \in O} f^\sigma(u, q) \\ &= \bigwedge_{p \leq u \in O} \bigvee_{\substack{u \geq d \in K \\ q \leq v \in O}} \bigwedge_{\substack{d \leq a_1 \in A \\ v \geq a_2 \in A}} f(a_1, a_2) \\ &= \bigvee_{\substack{q \leq v \in O \\ v \geq a_2 \in A}} \bigwedge_{\substack{p \leq a_1 \in A \\ v \geq a_2 \in A}} f(a_1, a_2) \\ &= f^\sigma(p, q) \end{aligned}$$

where the penultimate step follows from the fact that taking the meet over all $u \in O$ such that $p \leq u$, of the join over all $d_1 \in K$ such that $d_1 \leq u$, amounts to doing nothing since every element of the canonical extension can be written in this way. An analogous dual proof shows that $f^\pi = f^\sigma$ for pairs of open elements too. $\square$

The natural question is now: when do f^σ and f^π agree on *any* $x \in A^\sigma$, i.e. when is f smooth? We will present a necessary and sufficient condition for a map to be smooth in Section 2.3, but we can already show that some frequent algebraic conditions on f^σ and f^π (which can themselves be inherited from conditions on f as we shall see in a short while) guarantee smoothness. We will say that a map $f : A^n \rightarrow A$ is **smooth in its i^{th} argument** ($1 \leq i \leq n$) if for any n -tuple $(x_1, \dots, x_i, \dots, x_n)$ such that $x_i \in A^\sigma$ and $x_j \in K \cup O$ for every $1 \leq j \leq n, j \neq i$ we have

$$f^\sigma(x_1, \dots, x_i, \dots, x_n) = f^\pi(x_1, \dots, x_i, \dots, x_n)$$

Proposition 2.13. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A^n \rightarrow A$ be an isotone map.*

- (i) *If f^σ preserves down-directed meets or non-empty joins in an argument, then f is smooth in that argument.*
- (ii) *If f^π preserves upwardly directed joins or non-empty meets in an argument, then f is smooth in that argument.*

Proof. The proof for n -ary functions is exactly the same as the proof for unary functions, so for notational clarity we stick to unary functions. If f^σ preserves down-directed meets then the proof is essentially in [GH01] and goes as follows: let $x \in A^\sigma$, then by definition of the canonical extension we can write

$$x = \bigwedge \{u \mid x \leq u \in O\}$$

Since the set $\{u \mid x \leq u \in O\}$ is down-directed we have by our hypothesis that $f^\sigma(x) = \bigwedge \{f^\sigma(u) \mid x \leq u \in O\}$ but since $f^\sigma = f^\pi$ on O we have

$$f^\sigma(x) = \bigwedge \{f^\pi(u) \mid x \leq u \in O\} = f^\pi(x)$$

since f^σ and f^π agree on $K \cup O$. The proof for non-empty join preservation can be found in [GH01]. The proof for f^π follows by duality. $\square$

We will say that $f : A \rightarrow A$ **anti-preserves (down-directed) meets** if for any (down-directed) subset $X \subseteq A$

$$f(\bigwedge X) = \bigvee f[X]$$

and dually for joins. Note that if f anti-preserves meets $a \leq b$ in A , i.e. $a \wedge b = a$ then $f(a) = f(a \wedge b) = f(a) \vee f(b)$ and f is automatically antitone, and

dually for joins. Similarly, by considering the down-directed closure of $\{a, b\}$, i.e. $\{a, b, a \wedge b\}$ the anti-preservation of down-directed meets gives

$$\begin{aligned} f(a) &= f(\bigwedge \{a, b, a \wedge b\}) \\ &= \bigvee f[\{a, b, a \wedge b\}] \\ &= f(a) \vee f(b) \vee f(a) = f(a) \vee f(b) \end{aligned}$$

and we again automatically get that f is antitone. This of course dually holds for the anti-preservation of up-directed joins too.

Corollary 2.14. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A^n \rightarrow A$.*

- (i) *If f^σ anti-preserves up-directed joins or non-empty meets in one argument, then f is smooth in that argument.*
- (ii) *If f^π anti-preserves down-directed meets or non-empty joins in one argument, then f is smooth in that argument.*

Proof. The proof is very similar to 2.13, but we show (i) as an illustration of working with unary antitone maps. The case of arbitrary n -ary maps follows easily from the unary cases by working component-wise. The proof of (ii) follows by duality.

Let us first assume that f^σ anti-preserves up-directed joins. As we just saw this means that f^σ is antitone, and since $f^\sigma = f$ on A , this means that f must be antitone too. We now write $x \in A^\sigma$ as

$$x = \bigvee \{p \mid x \geq p \in K\}$$

which is an up-directed join. We thus have

$$\begin{aligned} f^\sigma(x) &= f^\sigma(\bigvee \{p \mid x \geq p \in K\}) \\ &= \bigwedge \{f^\sigma(p) \mid x \geq p \in K\} \\ &= \bigwedge \{f^\pi(p) \mid x \geq p \in K\} = f^\pi(x) \end{aligned}$$

by the fact that f is antitone, f^σ anti-preserves down-directed meets, Proposition 2.12, and the definition of f^π .

Let us now assume that f^σ anti-preserves non-empty meets. In particular this

means that f must be antitone, and thus for any $x \in A^\sigma$ we have

$$\begin{aligned}
f^\sigma(x) &= \bigvee \{ \bigwedge \{ f(a) \mid u \geq a \in A \} \mid x \leq u \in O \} \\
&= \bigwedge \{ \bigvee \{ f(a) \mid a \in \text{im}(\gamma) \} \mid \gamma : \uparrow x \cap O \rightarrow A \text{ s.th. } \gamma(u) \leq u \} \\
&= \bigwedge \{ \bigvee \{ f^\sigma(a) \mid a \in \text{im}(\gamma) \} \mid \gamma : \uparrow x \cap O \rightarrow A \text{ s.th. } \gamma(u) \leq u \} \\
&= \bigwedge \{ f^\sigma(\bigwedge \text{im}(\gamma)) \mid \gamma : \uparrow x \cap O \rightarrow A \text{ s.th. } \gamma(u) \leq u \} \\
&= \bigwedge \{ f^\pi(\bigwedge \text{im}(\gamma)) \mid \gamma : \uparrow x \cap O \rightarrow A \text{ s.th. } \gamma(u) \leq u \} \\
&\geq \bigwedge \{ f^\pi(p) \mid x \geq p \in K \} = f^\pi(x)
\end{aligned}$$

where the first step uses general distributivity (Theorem 2.10) and the second step follows from the anti-preservation property of f^σ . The second to last step comes from $\bigwedge \text{im}(\gamma) \in K$ and $f^\sigma = f^\pi$ on $K \cup O$. The last step of the derivation follows from the fact that each $\bigwedge \text{im}(\gamma) \leq x$ since each x is the meet of all the opens u above it and each γ picks elements of A below each u . $\square$

2.2 Algebraic Properties of the Canonical Extension

In this section we will examine how the canonical extension construction interacts with the algebraic properties of maps. Assume a map $f : A^n \rightarrow A$ satisfies a certain algebraic property, for example meet-preservation in its i^{th} argument, does $f^\sigma : (A^\sigma)^n \rightarrow A^\sigma$ satisfy a similar property? We will examine several such properties and see that the canonical extension of a map often satisfies a ‘complete version’ of the algebraic properties of the map it extends. We will show how this allows for the integration of the notion of canonical extension of DLEs and BAEs with their abstract presentation as L -algebras (see Section 1.3). Finally, we will examine how canonical extensions interact with function composition.

2.2.1 Properties preserved and strengthened by canonical extension

We start with an algebraic property which is not often considered but which will be crucial for the applications at the end of this Chapter. The following concepts are due to Henkin in [Hen70].

Definition 2.15. Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A \rightarrow A$. We will say that f is **k -additive** if for any finite set $U \subseteq A$

$$f(\bigvee U) = \bigvee \{ f(\bigvee V) \mid V \in \mathcal{P}_k(U) \}$$

where $\mathcal{P}_k(U)$ is the set of subsets of U with at most k elements. A map $f : A \rightarrow A$ will be called **completely k -additive** if for any set $U \subseteq A$, the join $\bigvee \{f(\bigvee V) \mid V \in \mathcal{P}_k(U)\}$ exists and equals $f(\bigvee U)$. We will say that a map is (completely) ω -additive if it is (completely) k -additive for some $k < \omega$.

The following two remarks illustrate the importance of the concept of k -additivity:

1. 1-additivity just means additivity, i.e. distribution over joins.
2. A completely k -additive map on a complete lattice preserves up-directed joins: indeed if $U \subseteq A$ is up-directed, then for any $V \in \mathcal{P}_k(U)$, $\bigvee V \in U$ and thus

$$\bigvee \{f(\bigvee V) \mid V \in \mathcal{P}_k(U)\} = \bigvee f[U] = f(\bigvee U)$$

It is natural to generalize the concept of k -additivity as follows.

Definition 2.16. Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A \rightarrow A$. We will say that f is **k -multiplicative** if for any finite set $U \subseteq A$

$$f(\bigwedge U) = \bigwedge \{f(\bigwedge V) \mid V \in \mathcal{P}_k(U)\}$$

A map $f : A \rightarrow A$ will be called **completely k -multiplicative** if for any set $U \subseteq A$, the meet $\bigwedge \{f(\bigwedge V) \mid V \in \mathcal{P}_k(U)\}$ exists and equals $f(\bigwedge U)$. We will say that a map is (completely) ω -multiplicative if it is (completely) k -multiplicative for some $k < \omega$. We will say that a map $f : A \rightarrow A$ is **anti- k -additive** if

$$f(\bigvee U) = \bigwedge \{f(\bigwedge V) \mid V \in \mathcal{P}_k(U)\}$$

and similarly that it is **anti- k -multiplicative** if

$$f(\bigwedge U) = \bigvee \{f(\bigvee V) \mid V \in \mathcal{P}_k(U)\}$$

Complete anti- k -additivity and multiplicativity are defined analogously.

The following is an easy consequence of the definitions.

Lemma 2.17. *For any $k > 0$, k -additive and k -multiplicative operators are isotone. Anti- k -additive and anti- k -multiplicative operators are antitone.*

Proof. Let $f : A \rightarrow A$ be a k -additive operator and let $a, b \in A$ s.th. $a \leq b$, i.e. $a \vee b = b$. If $k = 1$ we obviously get $f(a \vee b) = f(b) = f(a) \vee f(b)$ whence $f(a) \leq f(b)$. If $k \geq 2$ we get $f(a \vee b) = f(b) = f(a) \vee f(b) \vee f(a \vee b) = f(a) \vee f(b) \vee f(b) = f(a) \vee f(b)$ whence $f(a) \leq f(b)$ also. The proof is similar for the other types of operators. $\square$

The following theorem will be used heavily at the end of this chapter, and since we will only need it for unary operators, this is how we will present the result.

Theorem 2.18. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A , and let $f : A \rightarrow A$ be a function.*

- (i) *If f is k -additive, then $f^\sigma : A^\sigma \rightarrow A^\sigma$ is completely k -additive. In particular, f^σ preserves up-directed joins.*
- (ii) *If f is k -multiplicative, then $f^\pi : A^\sigma \rightarrow A^\sigma$ is completely k -multiplicative. In particular, f^π preserves down-directed meets.*
- (iii) *If f is anti- k -additive, then $f^\pi : A^\sigma \rightarrow A^\sigma$ is completely anti- k -additive. In particular, f^π anti-preserved up-directed joins.*
- (iv) *If f is anti- k -multiplicative, then $f^\sigma : A^\sigma \rightarrow A^\sigma$ is completely anti- k -multiplicative. In particular, f^σ anti-preserved down-directed meets.*

Proof. We show (i), (ii) then follows by duality, (iii) is shown in a very similar way to (i) and (iv) follows from (iii) by duality. Let $f : A \rightarrow A$ be k -additive, and let us first show that f^σ is completely k -additive on sets of closed elements. Let $P \subseteq K$, we need to show $f^\sigma(\bigvee P) = \bigvee \{f(\bigvee V) \mid V \in \mathcal{P}_k(P)\}$. It follows immediately from Lemma 2.17 that

$$\bigvee \{f^\sigma(\bigvee V) \mid V \in \mathcal{P}_k(P)\} \leq f^\sigma(\bigvee P)$$

So we need only show

$$f^\sigma(\bigvee P) \leq \bigvee \{f^\sigma(\bigvee V) \mid V \in \mathcal{P}_k(P)\} \quad (2.5)$$

We will now rewrite this inequation in a $\bigvee \leq \bigwedge$ format. For the left hand-side we simply use the definition of canonical extension:

$$f^\sigma(\bigvee P) = \bigvee \{f^\sigma(q) \mid \bigvee P \geq q \in K\}$$

For the left-hand side, we use the definition of canonical extension, the fact that a finite join of closed sets is closed, and the complete distributivity property of Theorem 2.10 to get:

$$\begin{aligned} & \bigvee \{f^\sigma(\bigvee V) \mid V \in \mathcal{P}_k(P)\} \\ &= \bigvee \{ \bigwedge \{f(a) \mid A \ni a \geq \bigvee V\} \mid V \in \mathcal{P}_k(P) \} \\ &= \bigwedge \{ \bigvee f[\text{im } \gamma] \mid \gamma : \mathcal{P}_k(P) \rightarrow A \text{ s.th. } \gamma(V) \geq \bigvee V \} \end{aligned}$$

Thus to show (2.5), we need to show that for any choice of closed element $q \leq \bigvee P$, and any choice function γ as above, $f^\sigma(q) \leq \bigvee f[\text{im } \gamma]$. To see this,

note first that by definition $q \leq \bigvee P \leq \bigvee \text{im} \gamma$, since γ picks for a set of at most k elements of P a clopen a above their join. Then note that $\bigvee \text{im} \gamma$ is open (as a join of elements of A) and that q is assumed to be closed. So by compactness, there exists a finite subset F of the domain of γ (i.e. $\mathcal{P}_k(P)$) such that

$$q \leq \bigvee \text{im}(\gamma \upharpoonright F)$$

We now use a trick which we shall use again in Theorem 2.35 which is to make γ ‘locally join preserving’. We can assume, modulo the addition of a finite number of elements to F , that $F = \mathcal{P}_K(I_q)$ for a finite set $I_q \subseteq P$ such that $q \in I_q$. We now define the ‘locally join preserving’ version of γ as the map $\gamma_q : \mathcal{P}_k(I_q) \rightarrow A$ defined on singletons as

$$\gamma_q(\{p\}) = \bigwedge \{\gamma(U) \mid p \in U \in \mathcal{P}_k(I_q)\}$$

and on the other subsets $U \subseteq I_q$ as

$$\gamma_q(U) = \bigvee \{\gamma_q(p) \mid p \in U\}$$

Since I_q is finite, all meets and all joins are finite, and γ_q does indeed take its values in A . For notational clarity we also define the map $\psi_q : I_q \rightarrow A, p \mapsto \gamma_q(\{p\})$. We now have

$$\begin{aligned} \bigvee f[\text{im} \gamma_q] &= \bigvee \{f(\gamma_q(U)) \mid U \in \mathcal{P}_k(I_q)\} \\ &= \bigvee \{f(\bigvee \{\gamma_q(p) \mid p \in U\}) \mid U \in \mathcal{P}_k(I_q)\} \\ &= \bigvee \{f(\bigvee V) \mid V \in \mathcal{P}_k(\text{im} \psi_q)\} \\ &= f(\bigvee \text{im} \psi_q) \\ &= f(\bigvee \{\gamma_q(\{p\}) \mid p \in I_q\}) \\ &= f(\bigvee \text{im} \gamma_q) \end{aligned}$$

where we have used the k -additivity of f at the fourth step. Moreover, since $q \leq \bigvee U \leq \gamma(U)$ for each $q \in U \in \mathcal{P}_k(I_q)$ we have $q \leq \gamma_q(\{q\}) \leq \bigvee \text{im} \gamma_q$. We also have $\gamma_q(\{p\}) \leq \gamma(U)$ for each $p \in U$ and thus $\gamma_q(U) \leq \gamma(U)$. It now follows that

$$f^\sigma(q) \leq f^\sigma(\bigvee \text{im} \gamma_q) = \bigvee f^\sigma[\text{im} \gamma_q] \leq f^\sigma[\text{im} \gamma \upharpoonright \mathcal{P}_k(I_q)] \leq \bigvee f^\sigma[\text{im} \gamma]$$

For the case of a general set $U \subseteq A^\sigma$, we will show that

$$f^\sigma(\bigvee U) \leq \bigvee \{f(\bigvee V) \mid V \in \mathcal{P}_k(\overline{U})\}$$

where $\overline{U}$ is the set of closed elements below an element of U . The result will then follow by the monotonicity result of Lemma 2.17 and the definition of canonical extension. By definition we have

$$f^\sigma(\bigvee U) = \bigvee \{f^\sigma(p) \mid \bigvee U \geq p \in K\}$$

and any such p has the property that $p \leq \bigvee \overline{U}$ by definition of $\overline{U}$. By using monotonicity and the fact that we have established the result for sets of closed elements we can now conclude that for any $p \leq \bigvee U$ we have

$$\begin{aligned} f^\sigma(p) &\leq f^\sigma\left(\bigvee \overline{U}\right) \\ &\leq \bigvee \{f^\sigma\left(\bigvee V\right) \mid V \in \mathcal{P}_k(\overline{U})\} \\ &\leq \bigvee \{f^\sigma\left(\bigvee V\right) \mid V \in \mathcal{P}_k(U)\} \end{aligned}$$

which concludes the proof. $\square$

We now focus on meet and join preservation. We prove the following for binary monotone operators, but the results clearly generalises to arbitrary finite arities.

Lemma 2.19. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A \times A \rightarrow A$ be monotone in both arguments.*

- (i) *If f preserves finite joins in one argument then f^σ preserves non-empty joins in that argument*
- (ii) *If f anti-preserved finite meets in one argument then f^σ anti-preserved arbitrary meets in that argument*
- (iii) *If f preserves finite meets in one argument then f^π preserves non-empty meets in that argument*
- (iv) *If f anti-preserved finite joins in one argument then f^π anti-preserved arbitrary joins in that argument*

Proof. We will assume without loss of generality that f has the preservation property in the first argument. The proofs of (i) can be found e.g. [GH01] and [Ven06] in the case where f is isotone in both arguments, our proof of (ii) will show how to generalise this proof to maps which have both isotone and antitone arguments. Note also that the proof of 2.18 also implies (i) (in fact it is an adaptation of the proof from [GH01]).

Let us now show (ii) when f is antitone in the first and isotone in the second argument. We start with showing $f^\sigma(\bigwedge U, p) \leq \bigvee \{f^\sigma(u, p) \mid u \in U\}$ where $U \subseteq O$ and $p \in K$ (since $\bigwedge U \leq u$ for all $u \in U$ and f is antitone in the first argument $f(\bigwedge U, p) \geq f(u, p)$ for each $u \in U$ and the opposite inequality follows). Again by definition

$$f^\sigma(\bigwedge U, p) = \bigvee \{f^\sigma(v, p) \mid \bigwedge U \leq v \in O\}$$

We also have by generalised distributivity that:

$$\begin{aligned} & \bigvee \{f^\sigma(u, p) \mid u \in U\} \\ &= \bigvee \{ \bigwedge \{f(a, b) \mid u \geq a \in A, p \leq b \in A\} \mid u \in U \} \\ &= \bigwedge \{ \bigvee \{f(a, b) \mid (a, b) \in \text{im}(\gamma)\} \mid \gamma : U \rightarrow A \times (p \uparrow \cap A) \text{ s.th. } \pi_1(\gamma(u)) \leq u \} \end{aligned}$$

So we need to show that for any $v \in O, \bigwedge U \leq v$ and any choice function γ as defined above $f^\sigma(v, p) \leq \bigvee f[\text{im}(\gamma)]$. For this note first that by definition of O

$$\begin{aligned} \bigwedge U &= \bigwedge \{ \bigvee \{a \mid u \geq a \in A\} \mid u \in U \} \\ &= \bigvee \{ \bigwedge \text{im}(\delta) \mid \delta : U \rightarrow A \text{ s.th. } \delta(u) \leq u \} \end{aligned}$$

Thus if $\bigwedge U \leq v$, then for any δ as above $\bigwedge \text{im}(\delta) \leq v$, and since $\bigwedge(\text{im}(\delta))$ is closed and v is open, there must exist by compactness a finite set $V_\delta \subseteq_\omega A$ such that $\bigwedge V_\delta \leq v$. Notice now that for any of the functions γ defined above, $\pi_1 \circ \gamma$ is one of these maps δ . Thus for any γ we have a finite set $V_\gamma \subseteq \text{im}(\pi_1 \circ \gamma)$ such that $\bigwedge U \leq \bigwedge V_\gamma \leq v$. This finite set also determines a finite subset $B_\gamma \subseteq \text{im}(\pi_2 \circ \gamma)$ by simply taking for each element of V_γ its π_2 correspondent. Each element of B_γ is, by definition of γ , a clopen above p . But since f anti-preserves finite meets and is antitone in the first argument we then have

$$\begin{aligned} f^\sigma(v, p) &\leq f^\sigma(\bigwedge V_\gamma, p) \\ &= f^\sigma(\bigwedge V_\gamma, \bigwedge B_\gamma) \\ &= f(\bigwedge V_\gamma, \bigwedge B_\gamma) \\ &= \bigvee \{f(a, \bigwedge B_\gamma) \mid a \in V_\gamma\} \\ &\leq \bigvee \{f^\sigma(a, b) \mid (a, b) \in \text{im}(\gamma)\} \end{aligned}$$

We now need to show that the inequality holds for an arbitrary set $X \subseteq A^\sigma$ and element $y \in A^\sigma$. Again, we proceed similarly to (i) by showing that $f^\sigma(\bigwedge X, y) \leq \bigvee \{f^\sigma(\tilde{x}, y) \mid \tilde{x} \in \tilde{X}\}$ where $\tilde{X}$ is the set of open elements above an element of X . By definition

$$f^\sigma(\bigwedge X, y) = \bigvee \{f^\sigma(u, p) \mid \bigwedge X \leq u \in O, y \geq p \in K\}$$

Since for any such (u, p) we have $\bigwedge \tilde{X} \leq u$

$$\begin{aligned} f^\sigma(u, p) &\leq f^\sigma(\bigwedge \tilde{X}, p) \\ &= \bigvee \{f(\tilde{x}, y) \mid \tilde{x} \in \tilde{X}\} \end{aligned}$$

which completes the proof of this case. The case for (ii) when f is antitone in the second argument is easy to prove by combining the two cases we have already shown.

The proof for (iii) and (iv) follow from (i) and (ii) by duality. $\square$

Corollary 2.20. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A \times A \rightarrow A$.*

- (i) *If f preserves finite joins in one argument then f^π preserves non-empty joins in that argument.*
- (ii) *If f preserves finite meets in one argument, then f^σ preserves non-empty meets in that argument.*
- (iii) *If f anti-preserves finite meets in one argument then f^π anti-preserves non-empty meets in that argument.*
- (iv) *If f anti-preserves finite joins in one argument then f^σ anti-preserves non-empty joins in that argument.*

Proof. For (i) we know by Lemma 2.19 that if f preserves finite joins, then f^σ preserves non-empty-joins. Moreover, we know from Proposition 2.13 that if f^σ preserves non-empty joins, then $f^\sigma = f^\pi$ and thus f^π preserves non-empty joins, and dually for (ii).

Similarly for (iii), by Lemma 2.19 if f anti-preserves finite meets, then f^σ anti-preserves non-empty meets and by Corollary 2.14 this means that $f^\sigma = f^\pi$, thus f^π anti-preserves non-empty meets, and dually for (iv). $\square$

The list of algebraic properties which we may wish to prove to be preserved by canonical extension is clearly infinite, and the choice will greatly depend on the application, i.e. the logic, at hand. We can show some simple preservation results for frequent algebraic properties such as being a expanding, contracting, or idempotent.

Proposition 2.21. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A \rightarrow A$ be expanding, then f^σ and f^π are expanding. Similarly, if f is contracting, then so are f^σ and f^π .*

Proof. (1) Expanding maps. We start by showing the property on closed elements $p \in K$, i.e. we want to show that

$$p \leq f^\sigma(p) = \bigwedge \{f(a) \mid p \leq a \in A\}$$

But recall that p is closed precisely when $p = \bigwedge \{a \mid p \leq a \in A\}$, thus we need to show that

$$\bigwedge \{a \mid p \leq a \in A\} \leq \bigwedge \{f(a) \mid p \leq a \in A\}$$

which is immediate from the assumption since for every term in the RHS meet there exist a smaller term on the LHS meet. Similarly, for a general element $x \in A^\sigma$ we have

$$x = \bigvee \{p \mid x \geq p \in K\} \leq \bigvee \{f^\sigma(p) \mid x \geq p \in K\} = f^\sigma(x)$$

The proof for f^π uses similar ideas: starting with elements $u \in O$ we get

$$u = \bigvee \{a \mid u \geq a \in A\} \leq \bigvee \{f(a) \mid u \geq a \in A\} = f^\pi(u)$$

And for a general $x \in A^\sigma$ we thus have

$$x = \bigwedge \{u \mid x \leq u \in O\} \leq \bigwedge \{f^\sigma(u) \mid x \leq u \in O\} = f^\pi(x)$$

(2) Contracting maps. For the contracting case, consider again $p \in K$, then by definition and by assumption it is easy to see that

$$f^\sigma(p) = \bigwedge \{f(a) \mid p \leq a \in A\} \leq \bigwedge \{a \mid p \leq a \in A\} = p$$

And for a general $x \in A^\sigma$ we then have

$$f^\sigma(x) = \bigvee \{f^\sigma(p) \mid x \geq p \in K\} \leq \bigvee \{p \mid x \geq p \in K\} = x$$

The proof for f^π follows exactly the same pattern. $\square$

Proposition 2.22. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A \rightarrow A$ be isotone and idempotent, then f^σ is idempotent.*

Proof. Since f is assumed to be monotone, we have by Lemma 2.23 that

$$f^\sigma = (ff)^\sigma \leq f^\sigma f^\sigma$$

So we need only to show $f^\sigma f^\sigma \leq f^\sigma$. We start by showing that this property holds on closed elements $p \in K$. By definition and the fact that f^σ sends closed elements to closed elements, we must show that for any $p \in K$

$$\bigwedge \{f(a) \mid f^\sigma(p) \leq a \in A\} \leq \bigwedge \{f(b) \mid p \leq b \in A\}$$

i.e. we must show that for any $b \in A$ such that $p \leq b$

$$\bigwedge \{f(a) \mid f^\sigma(p) \leq a \in A\} \leq f(b)$$

This can be shown by finding an $a \in A$ such that $f^\sigma(p) \leq a$ and $f(a) \leq f(b)$. We get such an a by putting $a = f(b)$. Since $p \leq b$ we then indeed have $f^\sigma(p) \leq f^\sigma(b) = f(b) = a$ and by idempotency we have $f(a) = f(f(b)) = f(b)$ and in particular $f(a) \leq f(b)$. Thus $f^\sigma f^\sigma = f^\sigma$ on K .

It is now immediate that for an arbitrary $x \in A^\sigma$ we have

$$(f^\sigma f^\sigma)(x) = \bigvee \{(f^\sigma f^\sigma)(p) \mid x \geq p \in K\} = \bigvee \{f^\sigma(p) \mid x \geq p \in K\} = f^\sigma(x)$$

$\square$

It is easy to check that an antitone idempotent map is constant, so this case need not be considered.

2.2.2 Canonical extension and L -algebras.

Let us briefly explore what the preservation results of this section mean in terms of the L -algebras in $\mathbf{A}$ described in Section 1.3.4. As we have seen in Section 1.5 and as we will see in Chapter 5, L -algebras of interest will often be defined by functors of the type

$$L\mathbb{A} = (\mathbf{F}S_\Sigma \mathbf{U}\mathbb{A}) / \simeq$$

where S_Σ is a polynomial **Set**-functor and $\simeq$ an equivalence relation generated by some axioms. These axioms are typically used to enforce algebraic conditions on the distributive lattice or boolean algebra expansions. The functor $L : \mathbf{BA} \rightarrow \mathbf{BA}$ used to define classical modal logic for example is defined by

$$L\mathbb{A} = \mathbf{F}\{\Diamond a \mid a \in A\} / (\Diamond(a \vee b) = \Diamond a \vee \Diamond b, \Diamond \perp = \perp)$$

The axioms ensure that the structure map of an L -algebra $L\mathbb{A} \rightarrow \mathbb{A}$ is equivalent to a map on A which preserves binary joins and the bottom element. More generally, if $\simeq$ defines some preservation property, then we have the following adjunction:

$$\frac{\mathbf{A}\text{-morphisms } (\mathbf{F}S_\Sigma \mathbf{U}\mathbb{A}) / \simeq \rightarrow \mathbb{A}}{\mathbf{Set}\text{-morphisms } S_\Sigma A \rightarrow A \text{ with the algebraic properties defined by } \simeq}$$

The preservation results developed in this section can now be understood in a new light. Let us consider a functor $L\mathbb{A} = \mathbf{F}(\mathbf{U}\mathbb{A})^n / \simeq$ where $\simeq$ describes some algebraic property of an n -ary expansion $f : (A)^n \rightarrow A$ (for example distribution over joins in its first argument). If this algebraic property is preserved under canonical extension, i.e. $f^\sigma : (A^\sigma)^n \rightarrow A^\sigma$ also satisfies this property, then by the general adjunction sketched above, the canonical extension construction can be lifted from $\mathbf{A}$ to $\mathbf{Alg}_{\mathbf{A}}(L)$, i.e. we have a map $f^\sigma : L(A^\sigma) \rightarrow A^\sigma$. Moreover, since $f^\sigma \upharpoonright A^n = f$ we clearly have

$$\begin{array}{ccc} L\mathbb{A} & \xrightarrow{L\eta_A} & L(\mathbb{A}^\sigma) \\ f \downarrow & & \downarrow f^\sigma \\ \mathbb{A} & \xrightarrow{\eta_A} & \mathbb{A}^\sigma \end{array}$$

This is a version of the Jónsson-Tarski theorem to which we will return in Chapter 5. Note however that for a general varietor L , there is no reason for the canonical extension construction to lift from $\mathbf{A}$ to $\mathbf{Alg}_{\mathbf{A}}(L)$, but this section has only shown a set of situations, which occur very frequently in practise, in which such a lifting is possible.

2.2.3 Canonical extension and function composition

We now move on to studying how the canonical extension procedure interacts with the composition of maps. Understanding this interaction will be crucial to understanding the canonicity of terms built by nesting modal operators, since this essentially amounts to functional composition. For clarity's sake we will denote the functional composition of $f, g : A \rightarrow A$ as fg rather than the more conventional $f \circ g$. The following result will be crucial and dates back to [JT51] and [Rib52] in the case of purely isotone maps. We extend it to monotone functions.

Lemma 2.23. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A , $f_i : A^{n_i} \rightarrow A$, $1 \leq i \leq m$ be n_i -ary monotone maps and $g : A^m \rightarrow A$ be an m -ary monotone map, then*

$$(g(f_1, \dots, f_m))^\sigma \leq g^\sigma(f_1^\sigma, \dots, f_m^\sigma)$$

Proof. We can assume w.l.o.g. that the arguments of the functions are gathered into isotone and antitone groups, starting with isotone arguments. For notational convenience, we show the case where all maps are binary with the first argument isotone and the second argument antitone, but the proof shown below is readily generalized to the arbitrary case. So let us consider $g, f_1, f_2 : A^2 \rightarrow A$ isotone in the first and antitone in the second argument. We start by showing the result for tuples of arguments $(p, u, v, q) \in K \times O \times O \times K$. By (2.3) we have

$$f_1^\sigma(p, u) = \bigwedge \{f_1(a_1, a_2) \mid p \leq a_1 \in A \ni a_2 \leq u\} \quad (2.6)$$

and by Proposition 2.12 and (2.4) we know that

$$f_2^\sigma(v, q) = f_2^\pi(v, q) = \bigvee \{f_2(b_1, b_2) \mid q \leq b_2 \in A \ni b_1 \leq v\} \quad (2.7)$$

Note in particular that $f_1^\sigma(p, u)$ is thus closed, whilst $f_2^\sigma(v, q)$ is open. From this it follows that

$$g^\sigma(f_1^\sigma(p, u), f_2^\sigma(v, q)) = \bigwedge \{g(c_1, c_2) \mid f_1^\sigma(p, u) \leq c_1 \in A \ni c_2 \leq f_2^\sigma(v, q)\}$$

Take any c_1, c_2 as above, i.e. such that

$$\bigwedge \{f_1(a_1, a_2) \mid p \leq a_1 \text{ and } a_2 \leq u\} \leq c_1 \quad c_2 \leq \bigvee \{f_2(b_1, b_2) \mid q \leq b_2 \text{ and } b_1 \leq v\}$$

since c_1 (resp. c_2) is clopen, it is open (resp. closed). By compactness and the fact that the sets above are down-directed and up-directed respectively, there must exist $\tilde{a}_1, \tilde{a}_2$ and $\tilde{b}_1, \tilde{b}_2$ in A such that $p \leq \tilde{a}_1, \tilde{a}_2 \leq u, \tilde{b}_1 \leq v, q \leq \tilde{b}_2$ and

$$f_1(\tilde{a}_1, \tilde{a}_2) \leq c_1 \quad c_2 \leq f_2(\tilde{b}_1, \tilde{b}_2)$$

Since $g(f_1, f_2)$ is isotone in its first and last parameters and antitone in its second and third parameters, we get that

$$\begin{aligned} (g(f_1, f_2))^\sigma(p, u, v, q) &\leq (g(f_1, f_2))^\sigma(\tilde{a}_1, \tilde{a}_2, \tilde{b}_1, \tilde{b}_2) \\ &= (g(f_1, f_2))(\tilde{a}_1, \tilde{a}_2, \tilde{b}_1, \tilde{b}_2) \\ &\leq g(c_1, c_2) \end{aligned}$$

Since this holds for any c_1, c_2 such that $f_1^\sigma(p, u) \leq c_1$ and $c_2 \leq f_2^\sigma(v, q)$, we must have

$$\begin{aligned} (g(f_1, f_2))^\sigma(p, u, v, q) &\leq \bigwedge \{g(c_1, c_2) \mid f_1^\sigma(p, u) \leq c_1 \in A \ni c_2 \leq f_2^\sigma(v, q)\} \\ &= g^\sigma(f_1^\sigma(p, u), f_2^\sigma(v, q)) \end{aligned}$$

From this we recover the general case as follows. Let $x_1, x_2, y_1, y_2 \in A^\sigma$, then we can re-write Eq. (2.3) as

$$\begin{aligned} (g(f_1, f_2))^\sigma(x_1, x_2, y_1, y_2) &= \\ \bigvee \{(g(f_1, f_2))^\sigma(p, u, v, q) \mid x_1 \geq p \in K, x_2 \leq u \in O, y_1 \leq v \in O, y_2 \geq q \in K\} \end{aligned} \quad (2.8)$$

and

$$\begin{aligned} g^\sigma(f_1^\sigma(x_1, x_2), f_2^\sigma(y_1, y_2)) &= \\ \bigvee \{g^\sigma(r, w) \mid f_1^\sigma(x_1, x_2) \geq r \in K, f_2^\sigma(y_1, y_2) \leq w \in O\} \end{aligned} \quad (2.9)$$

To prove the inequality we show that for every $(g(f_1, f_2))^\sigma(p, u, v, q)$ in the set of Eq. (2.8), there exist a $g^\sigma(r, w)$ in the set of Eq. (2.9) such that

$$(g(f_1, f_2))^\sigma(p, u, v, q) \leq g^\sigma(r, w)$$

But this follows immediately from what we have shown above, we just need to take $r = f_1^\sigma(p, u)$ which is indeed closed as was shown earlier, and $w = f_2^\sigma(v, q)$ which is open. The result then follows from the special case we started off with. $\square$

Corollary 2.24. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A , $f_i : A^{n_i} \rightarrow A$, $1 \leq i \leq m$ be n_i -ary monotone maps and $g : A^m \rightarrow A$ be an m -ary monotone map, then*

$$(g(f_1, \dots, f_m))^\sigma \upharpoonright (K \cup O)^N = g^\sigma(f_1^\sigma, \dots, f_m^\sigma) \upharpoonright (K \cup O)^N$$

where $N = \sum_i n_i$

Proof. The proof can be found in [GH01] for unary isotone maps. For general monotone maps, we get by duality with the preceding Lemma 2.23 that

$$g^\pi(f_1^\pi, \dots, f_m^\pi) \leq (g(f_1, \dots, f_m))^\pi$$

By Proposition 2.12 we know that for every $1 \leq i \leq m$ we have $f_i^\sigma \upharpoonright (K \cup O)^{n_i} = f_i^\pi \upharpoonright (K \cup O)^{n_i}$, and from Proposition 2.11 we know that $g^\sigma \leq g^\pi$, we therefore get :

$$\begin{aligned} (g(f_1, \dots, f_m))^\sigma &\leq g^\sigma(f_1^\sigma, \dots, f_m^\sigma) \\ &= g^\sigma(f_1^\pi, \dots, f_m^\pi) \\ &\leq g^\pi(f_1^\pi, \dots, f_m^\pi) \\ &\leq (g(f_1, \dots, f_m))^\pi \end{aligned}$$

on $(K \cup O)^N$. By applying Proposition 2.12 one more time, we see that

$$(g(f_1, \dots, f_m))^\sigma \upharpoonright (K \cup O)^N = (g(f_1, \dots, f_m))^\pi \upharpoonright (K \cup O)^N$$

and as a consequence, all the terms in the above sequence of inequalities are equal on $(K \cup O)^N$, and the result follows. $\square$

Ideally, we would like to strengthen the inequality which we have just proved in Lemma 2.23 to a full equality, i.e. we would like

$$g^\sigma(f_1^\sigma, \dots, f_m^\sigma) = (g(f_1, \dots, f_m))^\sigma$$

As it turns out, several of the algebraic properties which we have come across guarantee this equality. For example, in the case of unary operators, if g preserves joins, then $g^\sigma f^\sigma = (gf)^\sigma$ for any monotone g . If f preserves meets, then $g^\sigma f^\sigma = (gf)^\sigma$ for any isotone f . These results can be proven directly by standard canonical extension algebra (see e.g. [GH01] for several such results). However, there is a more compact and modular method for proving this type of results. This method was developed in [GJ04] and [Ven06] and relies on some topology.

2.3 Topology to the rescue

In this section we will make heavy use of the notation and presentation of [Ven06] and in particular of the *principle of matching continuities*.

Proposition and Definition 2.25. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with canonical extension $\mathbb{A}^\sigma$ and let us define the collections of subsets of A^σ .*

$$- \sigma^\uparrow = \{ \uparrow p \mid p \in K \}$$

- $\sigma^\downarrow = \{\downarrow u \mid u \in O\}$
- $\sigma = \{\uparrow p \cap \downarrow u \mid K \ni p \leq u \in O\}$

then $\sigma^\uparrow, \sigma^\downarrow$ and σ (together with the empty set) constitute topologies on A^σ and σ is the join of $\sigma^\uparrow$ and $\sigma^\downarrow$ in the lattice of topologies on A^σ .

Proof. Let's first show that $\sigma^\uparrow$ forms a topology. Clearly $A^\sigma \in \sigma^\uparrow$ as $A^\sigma = \uparrow \perp$ and $\perp \in K$ (since it is clopen). If $p_1, \dots, p_n \in K$, then

$$\begin{aligned} \bigcap_{i=1}^n (\uparrow p_i) &= \{x \in A^\sigma \mid p_i \leq x, 1 \leq i \leq n\} \\ &= \{x \in A^\sigma \mid \bigvee_{i=1}^n p_i \leq x\} = \uparrow(\bigvee_{i=1}^n p_i) \end{aligned}$$

and $(\bigvee_{i=1}^n p_i) \in K$. Similarly, if $(p_i)_{i \in I} \subseteq K$, then

$$\begin{aligned} \bigcup_{i \in I} (\uparrow p_i) &= \{x \in A^\sigma \mid \exists i \in I : p_i \leq x\} \\ &= \{x \in A^\sigma \mid \bigwedge_{i \in I} p_i \leq x\} = \uparrow(\bigwedge_{i \in I} p_i) \end{aligned}$$

and by definition of K , $\bigwedge_{i \in I} p_i \in K$.

The proof that $\sigma^\downarrow$ is a topology is dual to that of $\sigma^\uparrow$.

The join of $\sigma^\uparrow$ and $\sigma^\downarrow$ in the lattice of topologies over A^σ is the topology generated by taking $\sigma^\uparrow \cup \sigma^\downarrow$ as a sub-base, i.e. it is the topology whose opens are unions of finite intersections of elements of $\sigma^\uparrow \cup \sigma^\downarrow$. These finite intersections are of the form $\uparrow p \cap \downarrow u$ for some $K \ni p \leq u \in O$, and it is easy to verify from the case of $\sigma^\uparrow$ and $\sigma^\downarrow$ that the set of all such finite intersections - i.e. σ - is closed under arbitrary unions. $\square$

The next three topologies are well known to domain theorists and are defined as follows. A **Scott open** subset of A^σ is defined as a subset $U \subseteq A^\sigma$ which has the properties that it is (1) an up-set, and (2) $U \cap D \neq \emptyset$ for any up-directed set D such that $\bigvee D \in U$. The **Scott closed** sets corresponding to the Scott opens can be defined as the down-sets C such that for any up-directed set $D \subseteq C$, $\bigvee D \in U$. The **Scott topology** on A^σ is defined as the collection of all Scott open subsets of A^σ and is denoted by $\gamma^\uparrow$. By order duality, one can also define a collection $\gamma^\downarrow$ of dual Scott opens, i.e. down-sets U with the property that for any down-directed sets D such that $\bigwedge D \in U$, $U \cap D \neq \emptyset$. Finally as in the previous case we can define $\gamma = \{U \cap V \mid U \in \gamma^\uparrow, V \in \gamma^\downarrow\}$. This set of topologies is related to the previous one in the lattice of topologies on A^σ by:

Proposition 2.26. *Using the above notation, $\gamma^\uparrow \subseteq \sigma^\uparrow$, $\gamma^\downarrow \subseteq \sigma^\downarrow$, $\gamma \subseteq \sigma$*

Proof. We show that $\gamma^\uparrow \subseteq \sigma^\uparrow$. Take an Scott-open set $U \in \gamma^\uparrow$. By definition of the canonical extension operation, any $x \in U$ can be written as $x = \bigvee \{p \mid x \geq p \in K\}$, and since $\{p \mid x \geq p \in K\}$ is clearly up-directed, by definition of the Scott-open we must have $U \cap \{p \mid x \geq p \in K\} \neq \emptyset$. This means that there exist at least one $p \in K$ such that $x \in \uparrow p$. We therefore have

$$U \subseteq \bigcup \{\uparrow p \mid p \in U \cap K\}$$

Conversely, if $x \in \bigcup \{\uparrow p \mid p \in U \cap K\}$, there must exist a $p \in U \cap K$ such that $p \leq x$ and since U is an up-set this means that $x \in U$. Thus

$$U = \bigcup \{\uparrow p \mid p \in U \cap K\} \in \sigma^\uparrow$$

□

Note that whilst the topologies $\sigma^\uparrow, \sigma^\downarrow$ and σ only make sense for objects in $\mathbf{A}$ which are canonical extensions (or at least which have a dense sub-object), the topologies $\gamma^\uparrow, \gamma^\downarrow$ and γ are legitimate topologies on any object in $\mathbf{A}$.

Lemma 2.27 ([GJ04, Gie80]). *In the canonical extension of a distributive lattice A , the Scott topology has a basis of sets $\uparrow j$ where $j \in J_\omega(A^\sigma)$, the set of finite joins of completely join irreducible elements.*

Proof. Let $U \in \gamma^\uparrow$ and $x \in U$. By Lemma 2.7 we can write $x = \bigvee \{j \in J(A^\sigma) \mid p \leq x\}$ which is the join of an up-directed set. It follows immediately from the definition of $\gamma^\uparrow$ that there exists $j \in J(A^\sigma)$ such that $\uparrow p \subseteq U$ and $x \in \uparrow j$. Taking finite joins of elements of $J(A^\sigma)$ allows the collection $\{\uparrow j \mid j \in J_\omega(A^\sigma)\}$ to be a basis (since it is then closed under finite intersections). □

With these topologies in place we can talk about continuous maps from A^σ to itself, but there are of course as many notions of continuity as there are choices of topologies for the domain and codomain. Formally, we will say that $f : A^\sigma \rightarrow A^\sigma$ is (τ_1, τ_2) -continuous if f is a continuous map between the topological spaces (A^σ, τ_1) and (A^σ, τ_2) . For n -ary maps $f : (A^\sigma)^n$, given a topology τ on A^σ , we will also denote by τ the n -fold product topology on A^σ , i.e. the topology defined by the subbase of elements $\pi_i^{-1}(U), U \in \tau, 1 \leq i \leq n$.

The following Theorem highlights the importance of the topological approach. It shows that the topologies can be used to characterise the $(\cdot)^\sigma$ operation on maps.

Theorem 2.28. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A^n \rightarrow A$, be a monotone map which for notational convenience we assume to be isotone in its first m arguments ($m \leq n$) and antitone in the subsequent $m' = n - m$ arguments. Then $f^\sigma : (A^\sigma)^n \rightarrow A^\sigma$ is the largest $((\sigma^\uparrow)^m \times (\sigma^\downarrow)^{m'}, \gamma^\uparrow)$ -continuous extension of f .*

Moreover, for any $f : A^n \rightarrow A$ (i.e. not necessarily monotone), f^σ is the largest $(\sigma^\uparrow, \gamma^\uparrow)$ -continuous extension of f .

Proof. The proof is adapted from Proposition 7.14 of [Ven06]. To see that f^σ is $((\sigma^\uparrow)^m \times (\sigma^\downarrow)^{m'}, \gamma^\uparrow)$ -continuous, let U be a Scott-open set and let $(x_1, \dots, x_n) \in (A^\sigma)^n$ be such that $f^\sigma(x_1, \dots, x_n) \in U$. By definition of $(\cdot)^\sigma$ we know that

$$f^\sigma(x_1, \dots, x_n) = \bigvee \{f^\sigma(p_1, \dots, p_m, u_1, \dots, u_{m'}) \mid K \ni p_i \leq x_i, 1 \leq i \leq m, \\ x_{m+i} \leq u_i \in O, 1 \leq i \leq m'\}$$

It is easy to see that that this join is taken over an up-directed set of tuples in $K^m \times O^{m'}$ (by taking joins of closed sets and meets of open sets). By the definition of Scott opens, since $f^\sigma(x_1, \dots, x_n) \in U$ there must exist an n -tuple $(p_1, \dots, p_m, u_1, \dots, u_{m'}) \in K^m \times O^{m'}$ such that $p_i \leq x_i, 1 \leq i \leq m, x_{m+i} \leq u_i, 1 \leq i \leq m'$ and $f^\sigma(p_1, \dots, p_m, u_1, \dots, u_{m'}) \in U$, and thus

$$(x_1, \dots, x_n) \in \bigcap_{1 \leq i \leq m} \pi_i^{-1}(\uparrow p_i) \cap \bigcap_{1 \leq i \leq m'} \pi_{m+i}^{-1}(\downarrow u_i)$$

where $\pi_i : (A^\sigma)^n \rightarrow A^\sigma$ is the usual projection map. Thus $(x_1, \dots, x_n) \in (\sigma^\uparrow)^m \times (\sigma^\downarrow)^{m'}$ by definition of the product topology.

Let us now show that f^σ is the largest such extension. Assume that $g : (A^\sigma)^n \rightarrow A$ is another $((\sigma^\uparrow)^m \times (\sigma^\downarrow)^{m'}, \gamma^\uparrow)$ -continuous extension of f . To show that $g(x) \leq f^\sigma(x)$ for every $x \in (A^\sigma)^n$, we will use Lemma 2.7 and show that if p is completely join irreducible element of A^σ and $p \leq g(x)$, then $p \leq f^\sigma(x)$ too. Note that if $p \in J(A^\sigma)$ then $\uparrow p \in \gamma^\uparrow$. Indeed, if D is a directed set and $\bigvee D \in \uparrow p$, then $p \leq \bigvee D$, and thus $p \leq d$ for some $d \in D$ by Lemma 2.6. It then follows that $g(x) \in \uparrow p \in \gamma^\uparrow$, and since g is assumed to be $((\sigma^\uparrow)^m \times (\sigma^\downarrow)^{m'}, \gamma^\uparrow)$ -continuous there must exist $(p_1, \dots, p_m, u_1, \dots, u_{m'}) \in K^m \times O^{m'}$ such that $p_i \leq x_i, 1 \leq i \leq m, x_{m+i} \leq u_i, 1 \leq i \leq m'$ and $p \leq g(p_1, \dots, p_m, u_1, \dots, u_{m'})$. It then follows that for any $(a_1, \dots, a_n) \in A^n$ such that $a_i \geq p_i, 1 \leq i \leq m$ and $a_{m+i} \leq u_i, 1 \leq i \leq m'$

$$p \leq g(p_1, \dots, p_m, u_1, \dots, u_{m'}) \leq g(a_1, \dots, a_n) = f(a_1, \dots, a_n)$$

Since this holds for any such n -tuples of elements of A it follows that

$$p \leq \bigvee \{f(a_1, \dots, a_n) \mid a_i \geq p_i, 1 \leq i \leq m, a_{m+i} \leq u_i, 1 \leq i \leq m'\} \\ = f^\sigma(p_1, \dots, p_m, u_1, \dots, u_{m'}) \leq f^\sigma(x)$$

as desired.

The proof for arbitrary maps is similar. If $f^\sigma(x_1, \dots, x_n) \in U \in \gamma^\uparrow$, then by Eq.(2.3),

$$f^\sigma(x_1, \dots, x_n) = \bigvee \{ \bigwedge f[d, u] \mid K^n \ni d \leq x \leq u \in O^n \}$$

and it is not hard to see that the join is over an up-directed set: if $\bigwedge f[d_1, u_1]$ and $\bigwedge f[d_2, u_2]$ are in this set, then $\bigwedge f[d_1 \vee d_2, u_1 \wedge u_2]$ also lies in this set and $\bigwedge f[d_1, u_1] \leq \bigwedge f[d_1 \vee d_2, u_1 \wedge u_2]$ and $\bigwedge f[d_2, u_2] \leq \bigwedge f[d_1 \vee d_2, u_1 \wedge u_2]$. Thus there exists an interval $[d, u]$ such that $d \leq x \leq u$ and $\bigwedge f[d, u] \in U$, and therefore $f[d, u] \subseteq U$ since U is an up-set, and $(\sigma, \gamma^\uparrow)$ -continuity then follows from the definition of σ . The proof that f^σ is the greatest such extension is identical to the monotone case by using meets of intervals rather than closed and open elements. $\square$

By duality we then immediately get the following two corollaries.

Corollary 2.29. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A^n \rightarrow A$, be a monotone map which for notational convenience we assume to be isotone in its first m arguments ($m \leq n$) and antitone in the subsequent $m' = n - m$ arguments. Then $f^\pi : (A^\sigma)^n \rightarrow A^\sigma$ is the smallest $(\sigma^\downarrow)^m \times (\sigma^\uparrow)^{m'}, \gamma^\downarrow$ -continuous extension of f . Moreover, for any $f : A^n \rightarrow A$ (i.e. not necessarily monotone), f^π is the smallest $(\sigma^n, \gamma^\downarrow)$ -continuous extension of f .*

Corollary 2.30. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A^n \rightarrow A$. Then f is smooth iff it has a unique (σ, γ) extension.*

Proof. If f is smooth, then $f^\sigma = f^\pi$ is both $(\sigma^\uparrow, \gamma^\uparrow)$ - and $(\sigma^\downarrow, \gamma^\downarrow)$ -continuous by Proposition 2.28 and Corollary 2.29 and thus (σ, γ) -continuous by definition of γ . Any other (σ, γ) extension will be both smaller and greater, and thus equal to f^σ , proving unicity.

Conversely, if f has a (σ, γ) -extension g , then by the previous results we know that $g \leq f^\sigma$ and $f^\pi \leq g$. Moreover, by Proposition 2.11, $f^\sigma \leq f^\pi$, and smoothness follows. $\square$

As the proof of Proposition 2.28 shows there is no difficulty in dealing with n -ary maps, one just needs to consider n -tuples, pointwise operations, and use the product topology. For notational convenience, we will therefore proceed with unary maps only, with the understanding that every result in the remainder of this section can easily be extended to the n -ary case.

Proposition 2.31. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A , then $f : A^\sigma \rightarrow A^\sigma$ is:*

- (i) $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous iff it preserves up-directed joins,
- (ii) $(\gamma^\downarrow, \gamma^\downarrow)$ -continuous iff it preserves down-directed meets,
- (iii) $(\gamma^\downarrow, \gamma^\uparrow)$ -continuous iff it anti-preserved down-directed meets,
- (iv) $(\gamma^\uparrow, \gamma^\downarrow)$ -continuous iff it anti-preserved up-directed joins

Proof. The proof of (i) is a standard domain theoretical result and (ii) is the dual result. (iii) and (iv) follow the same lines and we will show (iii) as an illustration. Let us first assume that f anti-preserves down-directed meets, and let U be a Scott-open set in $\gamma^\uparrow$. We need to show that $f^{-1}(U) \in \gamma^\downarrow$. Note first that if $x \in f^{-1}(U)$, and $z \leq x$, since f is antitone we have $f(z) \geq f(x) \in U$ and since U is an up-set we have $f(z) \in U$, i.e. $z \in f^{-1}(U)$. So $f^{-1}(U)$ is a down set as required. Let us now take D a down-directed set such that $\bigwedge D \in f^{-1}(U)$, we need to show that $D \cap f^{-1}(U) \neq \emptyset$. Since f anti-preserves down-directed meets we have $f(\bigwedge D) = \bigvee f[D] \in U$ and since f is antitone we have that $f[D]$ is up-directed. Indeed if $d_1, d_2 \in f[D]$, then there exist $c_1, c_2 \in D$ such that $f(c_1) = d_1$ and $f(c_2) = d_2$ and since D is down-directed, there exist $r \in D, r \leq c_1$ and $r \leq c_2$ and thus $d_1 \leq f(r), d_2 \leq f(r)$ with $f(r) \in f[D]$. But if $f[D]$ is directed and $\bigvee f[D] \in U$, then $f[D] \cap U \neq \emptyset$, since U is a Scott-open, and thus there exists $d \in D$ such that $f(d) \in U$, i.e. $D \cap f^{-1}(U) \neq \emptyset$ as required.

For the opposite direction, assume that f is $(\gamma^\downarrow, \gamma^\uparrow)$ -continuous and let D be a down-directed set. We need to show that $f(\bigwedge D) = \bigvee f[D]$. Let us first show that f is antitone. For any $x, y \in A^\sigma$ such that $x \leq y$, consider the set $\downarrow f(x)$. It is straightforward to check that $\downarrow f(x)$ is a Scott closed set and thus $f^{-1}(\downarrow f(x))$ is a dual Scott closed set, and in particular it is an up-set, hence $y \in f^{-1}(\downarrow f(x))$, i.e. $f(y) \leq f(x)$. By antitonicity we immediately get

$$\bigvee f[D] \leq f(\bigwedge D)$$

Now, consider the Scott closed set $\downarrow(\bigvee f[D])$. By continuity $f^{-1}(\downarrow(\bigvee f[D]))$ is a dual Scott closed set and in particular it has the property that it must be closed under taking meets of down-directed subsets. Since for any $d \in D$, $f(d) \leq \bigvee f[D]$ we have $D \subseteq f^{-1}(\downarrow(\bigvee f[D]))$, and since D is assumed to be down-directed we must have $\bigwedge D \in f^{-1}(\downarrow(\bigvee f[D]))$, i.e.

$$f(\bigwedge D) \leq \bigvee f[D]$$

which concludes the proof. $\square$

The following two seemingly technical lemmas have in fact very important topological consequences which will be detailed in Proposition 2.34.

Lemma 2.32. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f : A \rightarrow A$ preserve finite meets. For any $p \in K$ and $x \in A^\sigma$, if $p \leq f^\sigma(x)$, there exists $p' \in K$ such that $p' \leq x$ and $p \leq f^\sigma(p') \leq f^\sigma(x)$.*

Proof. If f preserves finite meets then by Lemma 2.19 f^π preserves non-empty meets and by Proposition 2.13 f is then smooth. Hence

$$p \leq f^\sigma(x) = f^\pi(x) = \bigwedge \{f^\pi(u) \mid x \leq u \in O\}$$

i.e. for any $u \geq x$ we must have $p \leq f^\pi(u)$ and since

$$f^\pi(u) = \bigvee \{f(a) \mid u \geq a \in A\}$$

and $\{f(a) \mid u \geq a \in A\}$ is up-directed, there must exist by compactness an element $a_u \in A$ with $a_u \leq u$ such that $p \leq f(a_u)$. Consider now the closed element defined as the meet of all these elements of A , namely $p' = \bigwedge \{a_u \mid x \leq u, a_u \in A\}$. Since x is the meet of all opens above it and since each a_u lies below one of these opens we must have $p' \leq x$. Moreover, since f^σ preserve meets we have that $f^\sigma(p') = \bigwedge \{f^\sigma(a_u) \mid x \leq u, a_u \in A\}$ and thus $p \leq f^\sigma(p')$ by construction of the elements a_u . $\square$

Lemma 2.33. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and let $f^\sigma : A^\sigma \rightarrow A^\sigma$ anti-preserve joins then for any $p \in K$ and $x \in A^\sigma$, if $p \leq f^\sigma(x)$, there exists $u \in O$ such that $u \geq x$ and $p \leq f^\sigma(u) \leq f^\pi(x)$.*

Proof. The proof is an antitone version of the preceding Lemma 2.32. By Lemma 2.19 f^π must anti-preserve non-empty joins and by Corollary 2.14 f must be smooth. Hence

$$p \leq f^\sigma(x) = f^\pi(x) = \bigwedge \{f^\pi(q) \mid x \geq q \in K\}$$

i.e. for any $q \in K, q \leq x$ we must have $p \leq f^\pi(q)$. By definition this means

$$p \leq \bigvee \{f(a) \mid q \leq a \in A\}$$

and again by compactness and the fact that $\{f(a) \mid q \leq a \in A\}$ is up-directed (by antitonicity of f) this means that there exist an element $a_q \in A$ such that $p \leq f(a_q)$. Consider the open element defined as the join of all these elements, i.e. $u = \bigvee \{a_q \mid q \in K, x \geq q\}$. Since x is the join of all closed elements below it and since each a_q lies above one such closed element $u \geq x$. It remains to check that $p \leq f^\sigma(u)$. By definition and our assumption on f we have

$$p \leq \bigwedge \{f^\sigma(a_q) \mid q \in K, x \geq q\} = f^\sigma(\bigvee \{a_q \mid q \in K, x \geq q\}) = f^\sigma(u)$$

$\square$

Using the second set of topologies, we can restate and generalise the technical Lemmas 2.32 and 2.33 topologically as:

Proposition 2.34. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A , a map $f : A^\sigma \rightarrow A^\sigma$ is*

- (i) $(\sigma^\uparrow, \sigma^\uparrow)$ -continuous if it preserves non-empty meets
- (ii) $(\sigma^\downarrow, \sigma^\downarrow)$ -continuous if it preserves non-empty joins

(iii) $(\sigma^\downarrow, \sigma^\uparrow)$ -continuous if it anti-preserves non-empty joins

(iv) $(\sigma^\uparrow, \sigma^\downarrow)$ -continuous if it anti-preserves non-empty meets

Proof. We will show (i) and (iii), (ii) and (iv) follow by duality. For (i) we need to show that for any $p \in K$, $f^{-1}(\uparrow p) \in \sigma^\uparrow$. Take $x \in f^{-1}(\uparrow p)$, by lemma 2.32 we know that there exists $p'(x) \in K$ such that $p'(x) \leq x$ and $p \leq f(p'(x))$. We now define the closed element $q = \bigwedge \{p'(x) \mid x \in f^{-1}(\uparrow p)\}$ and claim that $\uparrow q = f^{-1}(\uparrow p)$. Since f preserves non-empty meets we have $f(q) = \bigwedge \{f(p'(x)) \mid x \in f^{-1}(\uparrow p)\}$ and $p \leq f(q)$, thus

$$\uparrow q \subseteq f^{-1}(\uparrow p)$$

Conversely, since for any $x \in f^{-1}(\uparrow p)$ there exist $p'(x) \leq x$ we have by definition of q that $x \in \uparrow q$ and thus

$$f^{-1}(\uparrow p) \subseteq \uparrow q$$

For (iii) we proceed in a similar way: take $x \in f^{-1}(\uparrow p)$, by lemma 2.33 we know that there exists $u(x) \in O$ such that $x \leq u(x)$ and $p \leq f(u)$. Define an open element $v = \bigvee \{u(x) \mid x \in f^{-1}(\uparrow p)\}$, then by assumption on f we have $f(v) = \bigwedge \{f(u(x)) \mid x \in f^{-1}(\uparrow p)\}$ and thus $\downarrow v \subseteq f^{-1}(\uparrow p)$ and conversely if $x \in f^{-1}(\uparrow p)$, there exists $x \leq u(x)$ and $x \leq v$ and thus $f^{-1}(\uparrow p) \subseteq \downarrow v$. $\square$

Since preservation of all non-empty joins in particular implies that of directed ones, it is clear that join preserving maps are also $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous, and similarly for the other combinations of Propositions 2.31 and 2.34.

We can also use the topological formalism to show an important property of (anti-)k-additive and (anti-)k-multiplicative maps.

Theorem 2.35. *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A and if $f : A \rightarrow A$ is k -additive, k -multiplicative, anti- k -additive, or anti- k -multiplicative, then its extension is (σ, γ) -continuous and thus smooth.*

Proof. We show the result for k -additive maps, the other cases are treated similarly/dually. Recall first that by Theorem 2.18 if f is k -additive, then f^σ is completely k -additive, and in particular it preserves up-directed joins. It follows from Theorem 2.31 that it is $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous. We now show that it is $(\sigma^\downarrow, \gamma^\downarrow)$ -continuous too. Let $\Gamma \in \gamma^\downarrow$, i.e. Γ is a down-set such that if a down directed set D is such that $\bigwedge D \in \Gamma$, then $\Gamma \cap D \neq \emptyset$. Assume now that $f^\sigma(x) \in \Gamma$ for some x , we will show that there exist $x \leq u \in O$ such that $f^\sigma(u) \in \Gamma$. From Lemma 2.17, we know that f is isotone and by Theorem 2.18 we know that it

is completely k -additive and we thus have:

$$\begin{aligned}
 f^\sigma(x) &= f^\sigma\left(\bigvee K \cap \downarrow x\right) \\
 &\stackrel{(1)}{=} \bigvee \{f^\sigma(\bigvee U) \mid U \in \mathcal{P}_k(K \cap \downarrow x)\} \\
 &\stackrel{(2)}{=} \bigvee \{\bigwedge \{f^\sigma(a) \mid \bigvee U \leq a \in A\} \mid U \in \mathcal{P}_k(K \cap \downarrow x)\} \\
 &\stackrel{(3)}{=} \bigwedge \{\bigvee f^\sigma[\text{im}\phi] \mid \phi : \mathcal{P}_k(K \cap \downarrow x) \rightarrow A, \bigvee U \leq \phi(U)\}
 \end{aligned}$$

where (1) follows by definition of complete k -additivity, (2) follows from the fact that each $\bigvee U$ being a finite join of closed elements is closed, and (3) follows from the complete distributivity property of Theorem 2.10. It is easy to see by taking pointwise meets of choice maps that $\{\bigvee f^\sigma[\text{im}\phi] \mid \phi : \mathcal{P}_k(K \cap \downarrow x) \rightarrow A, \bigvee U \leq \phi(U)\}$ is down-directed. Since $\Gamma \in \gamma^\downarrow$, there must exist $\phi_0 : \mathcal{P}_k(K \cap \downarrow x) \rightarrow A, \bigvee U \leq \phi_0(U)$ such that $\bigvee f^\sigma[\text{im}\phi_0] \in \Gamma$. Recall that $f^\sigma(x) = \bigvee \{f^\sigma(p) \mid x \geq p \in K\}$, and for the equality (3) above to hold, it must be the case that for every closed element $p \leq x$ we must have

$$f^\sigma(p) \leq \bigvee f^\sigma[\text{im}\phi]$$

for any choice function ϕ , and a fortiori for ϕ_0 . So let $p \leq x$ be closed, then $f^\sigma(p)$ is closed as well, and by compactness there must exist a finite subset $F \subseteq_\omega \mathcal{P}_k(K \cap \downarrow x)$ such that

$$f^\sigma(p) \leq \bigvee f^\sigma[\text{im}(\phi_0 \upharpoonright F)]$$

Modulo the addition of a finite number of elements to F we can assume that $F = \mathcal{P}_k(I_p)$ for I_p a finite subset of $K \cap \downarrow x$ such that $p \in I_p$. We now define a map $\phi_p : \mathcal{P}_k(I_p) \rightarrow A$ by first defining it on singletons $\{q\}, q \in I_p$

$$\phi_p(\{q\}) = \bigwedge \{\phi_0(U) \mid q \in U \in F\}$$

which is an element of A since it is a finite meet of elements of A . We then extend ϕ_p to the whole of $\mathcal{P}_k(I_p)$ by

$$\phi_p(U) = \bigvee \{\phi_p(\{q\}) \mid q \in U\}$$

It is clear that ϕ_p preserves all joins of at most k -singletons by construction. It

follows that if we define $\psi_p : I_p \rightarrow A$ by $\psi(q) = \phi_p\{q\}$ for each $q \in I_p$, we have

$$\begin{aligned}
 \bigvee f^\sigma[\text{im}\phi_p] &= \bigvee \{f^\sigma(\phi_p(U)) \mid U \in \mathcal{P}_k(I_p)\} \\
 &= \bigvee \{f^\sigma(\bigvee \{\phi_p(\{q\}) \mid q \in U\}) \mid U \in \mathcal{P}_k(I_p)\} \\
 &= \bigvee \{f^\sigma(\bigvee V) \mid V \in \mathcal{P}_k(\text{im}\psi_p)\} \\
 &= f^\sigma(\bigvee \text{im}\psi_p) \\
 &= f^\sigma(\bigvee \{\phi_p(\{q\}) \mid q \in I_p\}) \\
 &= f^\sigma(\bigvee \text{im}\phi_p)
 \end{aligned}$$

where we have used the k -additivity of f^σ at the fourth step. Moreover, given $U \in \mathcal{P}_k(I_p)$ we have by construction that $\phi_p(q) \leq \phi(U)$ for every $q \in U$ and thus $\phi_p(U) \leq \phi_0(U)$. In particular

$$f^\sigma(\bigvee \text{im}\phi_p) = \bigvee f^\sigma[\text{im}\phi_p] \leq \bigvee f^\sigma(\text{im}\phi_0 \upharpoonright F) \leq \bigvee f^\sigma(\text{im}\phi_0)$$

and thus $\bigvee f^\sigma(\text{im}\phi_p) \in \Gamma$. Note also that since $q \leq \bigvee U \leq \phi_0(U)$ whenever $q \in U$ we have $q \leq \phi_p(U)$ whenever $q \in U$ and in particular since we're assuming $p \in I_p$ we have

$$p \leq \bigvee \text{im}\phi_p$$

and $\text{im}\phi_p$ is an element of A . Let us write $a_p = \text{im}\phi_p$. For every $p \leq x$ we can thus find a clopen element a_p such that $p \leq a_p$ and $f^\sigma(a_p) \in \Gamma$. We claim that if we define $u = \bigvee_p a_p$, then u is the open element we are looking for, i.e. $x \leq u$ and $f^\sigma(u) \in \Gamma$. The fact that $x \leq u$ is easy: x is the join of all closed elements p below it, and for each such closed element $p \leq a_p$. Now let us show that $f^\sigma(u) \in \Gamma$. By complete k -additivity we have

$$f^\sigma(u) = f^\sigma(\bigvee_p a_p) = \bigvee \{f^\sigma(\bigvee U) \mid U \in \mathcal{P}_k\{a_p \mid p \in K \cap \downarrow x\}\}$$

The result will follow immediately if we can show that each $f^\sigma(\bigvee U) \leq \bigvee f^\sigma[\text{im}\phi]$, since $\bigvee f^\sigma[\text{im}\phi] \in \Gamma$. So let us consider $U \in \mathcal{P}_k\{a_p \mid p \in K \cap \downarrow x\}$, by definition it is of the form $\bigvee_i \bigvee \text{im}\phi_{p_i}$ where each ϕ_{p_i} is constructed as above. Based on these maps we define $I_U = \bigcup_i I_{p_i}$ and $\phi_U : \mathcal{P}_k(I_U) \rightarrow A$ in the same way as we defined ϕ_{p_i} : we start on singleton sets

$$\phi_U(\{q\}) = \bigvee \{\phi_{p_i}(q) \mid q \in I_{p_i}\}$$

and then on the entire domain:

$$\phi_U(V) = \bigvee \{\phi_U(\{q\}) \mid q \in V\}$$

Note that since all joins are finite, the map ϕ_U does indeed take its values in A . We now proceed as above: we define $\psi_U : I_U \rightarrow A, q \mapsto \phi_U(\{q\})$ and compute

$$\begin{aligned}
 \bigvee f^\sigma[\text{im}\phi_U] &= \bigvee \{f^\sigma(\phi_U(V)) \mid V \in \mathcal{P}_k(I_U)\} \\
 &= \bigvee \{f^\sigma(\bigvee \{\phi_U(\{q\}) \mid q \in U\}) \mid U \in \mathcal{P}_k(I_U)\} \\
 &= \bigvee \{f^\sigma(\bigvee V) \mid V \in \mathcal{P}_k(\text{im}\psi_U)\} \\
 &= f^\sigma(\bigvee \text{im}\psi_U) \\
 &= f^\sigma(\bigvee \{\phi_U(\{q\}) \mid q \in I_U\}) \\
 &= f^\sigma(\bigvee \text{im}\phi_U)
 \end{aligned}$$

For an given $q \in I_{p_i}$, we already know that $\phi_{p_i}(q) \leq \phi_0(q)$, and it follows that $\phi_U(q) \leq \phi_0(p)$ too. Similarly, given $V \in \mathcal{P}_k(I_U)$, each $\phi_{p_i}(q) \leq \phi_0(V)$ if $q \in V$ and $q \in I_{p_i}$, and thus $\phi_U(q) \leq \phi_0(V)$, whence $\phi_U(V) \leq \phi_0(V)$. It follows that

$$f^\sigma(\bigvee \text{im}\phi_U) = \bigvee f^\sigma[\text{im}\phi_U] \leq \bigvee f^\sigma(\text{im}\phi_0 \upharpoonright \mathcal{P}_k(I_U)) \leq \bigvee f^\sigma(\text{im}\phi_0)$$

which concludes the proof. $\square$

The following Theorem shows why these topological results are useful to study the interaction of canonical extension and functional composition. The result is essentially Theorem 2.30. of [GJ04].

Theorem 2.36 (Principle of matching topologies). *Let $\mathbb{A}$ be an object in $\mathbf{A}$ with underlying set A , and $f : A^n \rightarrow A$ and $g_i : A^{m_i} \rightarrow A, 1 \leq i \leq n$ be arbitrary maps, then*

- (i) *If there exist topologies τ_i on $A, i \leq i \leq n$ such that each g_i^σ is (σ^{m_i}, τ_i) -continuous and f^σ is $(\tau_1 \times \dots \times \tau_n, \gamma^\uparrow)$ -continuous, then*

$$f^\sigma(g_1^\sigma, \dots, g_n^\sigma) \leq (f(g_1, \dots, g_n))^\sigma$$

- (ii) *If there exist topologies τ_i on $A, i \leq i \leq n$ such that each g_i^σ is (σ^{m_i}, τ_i) -continuous and f^σ is $(\tau_1 \times \dots \times \tau_n, \gamma^\downarrow)$ -continuous, then*

$$f^\sigma(g_1^\sigma, \dots, g_n^\sigma) \geq (f(g_1, \dots, g_n))^\sigma$$

- (iii) *If there exist topologies τ_i on $A, i \leq i \leq n$ such that each g_i^σ is (σ^{m_i}, τ_i) -continuous and f^σ is $(\tau_1 \times \dots \times \tau_n, \gamma)$ -continuous, then*

$$f^\sigma(g_1^\sigma, \dots, g_n^\sigma) = (f(g_1, \dots, g_n))^\sigma$$

Proof. For (i) we use Theorem 2.28: the existence of the matching set of topologies $\tau_1, \dots, \tau_n$ means that $f^\sigma(g_1^\sigma, \dots, g_n^\sigma)$ is $(\sigma^{m_1} \times \dots \times \sigma^{m_n}, \gamma^\uparrow)$ -continuous, and since $(f(g_1, \dots, g_n))^\sigma$ is the largest such map the result follows. For (ii), we use Corollary 2.29 and the existence of the set of matching topologies to conclude that $f^\sigma(g_1^\sigma, \dots, g_n^\sigma) \geq (f(g_1, \dots, g_n))^\pi$ and the result then follows from Theorem 2.11. Claim (iii) is simply the conjunction of (i) and (ii). $\square$

This *principle of matching topologies* will be pivotal in our understanding of canonicity in the next sections.

2.4 Canonicity

We are now ready to define and study canonicity algebraically. Recall from Eq. (1.1) that given a signature Σ we can define the free $\mathbf{S}_\Sigma$ -algebra over $\mathbf{F}(V)$, i.e. $\mathbf{G} \circ \mathbf{F}(V)$. This Σ -AE will be denoted henceforth by $\mathbf{T}(V)$ and is the set of terms in the language of Σ -AEs modulo equivalence in $\mathbf{A}$. From Propositions 1.2, 1.6 and the universal property of free objects, it is clear that for every $\mathcal{A} = (\mathbb{A}, (f_s)_{s \in \Sigma})$ in $\mathbf{AE}(\Sigma)$ and every valuation $v : V \rightarrow \mathcal{A}$, there exist a unique interpretation map $\llbracket \cdot \rrbracket_v^{\mathcal{A}} : \mathbf{T}(V) \rightarrow \mathcal{A}$ such that the following triangle commutes (in **Set**):

$$\begin{array}{ccc} & & \mathbf{T}(V) \\ & \nearrow \eta_V & \downarrow \llbracket \cdot \rrbracket_v^{\mathcal{A}} \\ V & & \mathcal{A} \\ & \searrow \pi & \end{array}$$

An equation between two terms $s, t \in \mathbf{T}(V)$ will be called **canonical** if $\llbracket s \rrbracket_v^{\mathcal{A}} = \llbracket t \rrbracket_v^{\mathcal{A}}$ for any valuation $v : V \rightarrow \mathcal{A}$ implies that $\llbracket s \rrbracket_v^{\mathcal{A}^\sigma} = \llbracket t \rrbracket_v^{\mathcal{A}^\sigma}$ for any valuation $v : V \rightarrow \mathcal{A}^\sigma$. Using the notation introduced in Section 1.5.4, $s = t$ is canonical if

$$\mathcal{A} \models s = t \quad \Rightarrow \quad \mathcal{A}^\sigma \models s = t$$

Recall from the previous Chapter that equations between terms of the free algebra define varieties. From this perspective, an equation is canonical if the variety it defines is closed under taking canonical extensions. Similarly, we will call an inequation $s \leq t$ canonical if

$$\mathcal{A} \models s \leq t \quad \Rightarrow \quad \mathcal{A}^\sigma \models s \leq t$$

In order to be able to say anything about the canonicity of equations, we somehow need to compare the interpretations in $\mathcal{A}$ with interpretations in $\mathcal{A}^\sigma$. It is natural to try to use the map $(\cdot)^\sigma$ to mediate between these interpretations, but $(\cdot)^\sigma$ is defined on maps, not on terms. Moreover, not every valuation on $\mathcal{A}^\sigma$

originates from valuation on $\mathcal{A}$. We would therefore like to recast the problem in such a way that (1) terms are viewed as maps, and (2) we do not need to worry about valuations.

We will therefore adopt the language of *term functions*, i.e. we will view each term $t \in \mathsf{T}(V)$ built using n variables of V as defining for each AE $\mathcal{A}$, a map $t^{\mathcal{A}} : A^n \rightarrow A$. This will allow us to consider its canonical extension $(t^{\mathcal{A}})^{\sigma}$, and it will also allow us to reason without having to worry about specifying valuations. Formally, given a signature Σ and a set V of propositional variables, we inductively define the term function associated with an element $t \in \mathsf{T}(V)$ built from variables $x_1, \dots, x_n \in V$ as follows:

- $x_i^{\mathcal{A}} = \pi_i^n : A^n \rightarrow A, 1 \leq i \leq n$
- $(f(t_1, \dots, t_m))^{\mathcal{A}} = f^{\mathcal{A}} \circ \langle t_1^{\mathcal{A}}, \dots, t_m^{\mathcal{A}} \rangle$

where π_i is the usual projection on the i^{th} component, $f^{\mathcal{A}}$ is the interpretation of the symbol f in $\mathcal{A}$ and $\langle t_1^{\mathcal{A}}, \dots, t_m^{\mathcal{A}} \rangle$ is usual the product of m maps. Note that in this definition we work in **Set**, and the building blocks of term functions are thus the variables in V (interpreted as projections) and all operation symbols, including boolean operations. However, we could equally well work directly in **A** in which case we would consider FV as the set of variables and only interpret the symbols defined by the signature Σ .

We now show how term functions allow us to reason without explicit reference to valuations. Note first that when we enforce an equation $s = t$, the set of variables $\text{Var}(s)$ and $\text{Var}(t)$ of s and t need not be equal. However for the equation $s^{\mathcal{A}} = t^{\mathcal{A}}$ to be well-typed, the term functions must have the same domain, i.e. be built on the same number of variables. This presents no problem, since we can always consider a term on m variables to be a term on $n > m$ variables with $n - m$ ‘silent’ variables which are simply ignored. We can thus assume without loss of generality that $s^{\mathcal{A}} = t^{\mathcal{A}}$ is well-typed by simply adding silent variables where needed. In fact we can assume that $\text{Var}(s) = \text{Var}(t)$ by using this procedure.

Proposition 2.37. *Let $s, t \in \mathsf{T}(V)$ and $\mathcal{A}$ be an object in $\mathbf{AE}(\Sigma)$ then*

$$\mathcal{A} \models s = t \text{ iff } s^{\mathcal{A}} = t^{\mathcal{A}}$$

Proof. We show that if $\text{Var}(t) = \{x_1, \dots, x_n\} \subseteq V$, and $v : V \rightarrow \mathcal{A}$ is a valuation

$$\llbracket t \rrbracket_v^{\mathcal{A}} = t^{\mathcal{A}}(v(x_1), \dots, v(x_n))$$

The result then follows: since we can assume $\text{Var}(s) = \text{Var}(t)$ whenever $s^{\mathcal{A}} = t^{\mathcal{A}}$, then for any valuation

$$\llbracket s \rrbracket_v^{\mathcal{A}} = s^{\mathcal{A}}(v(x_1), \dots, v(x_n)) = t^{\mathcal{A}}(v(x_1), \dots, v(x_n)) = \llbracket t \rrbracket_v^{\mathcal{A}}$$

and thus $\mathcal{A} \models s = t$. Conversely, to show $s^{\mathcal{A}} = t^{\mathcal{A}}$ we need to show that $s^{\mathcal{A}}(a_1, \dots, a_n) = t^{\mathcal{A}}(a_1, \dots, a_n)$ for any $(a_1, \dots, a_n) \in A^n$. But any such choice can be used to build a valuation $v : V \rightarrow A$ such that $v(x_i) = a_i, 1 \leq i \leq n$. If $\mathcal{A} \models s = t$, then in particular

$$\begin{aligned} s^{\mathcal{A}}(a_1, \dots, a_n) &= s^{\mathcal{A}}(v(x_1), \dots, v(x_n)) \\ &= \llbracket s \rrbracket_v^{\mathcal{A}} = \llbracket t \rrbracket_v^{\mathcal{A}} \\ &= t^{\mathcal{A}}(v(x_1), \dots, v(x_n)) \\ &= t^{\mathcal{A}}(a_1, \dots, a_n) \end{aligned}$$

To show that $\llbracket t \rrbracket_v^{\mathcal{A}} = t^{\mathcal{A}}(v(x_1), \dots, v(x_n))$ we proceed by induction on the complexity of t . If $t = x$ is just a propositional variable, then

$$\llbracket t \rrbracket_v = v(x) = \pi_1^1(v(x)) = x^{\mathcal{A}}(v(x)) = t^{\mathcal{A}}(v(x))$$

If $t = f(t_1, \dots, t_n)$ where each t_i is m_i -ary and such that $\bigcup_i \text{Var}(t_i) = \text{Var}(t)$, then

$$\begin{aligned} \llbracket t \rrbracket_v^{\mathcal{A}} &= f^{\mathcal{A}}(\llbracket t_1 \rrbracket_v^{\mathcal{A}}, \dots, \llbracket t_n \rrbracket_v^{\mathcal{A}}) \\ &= f^{\mathcal{A}}(t_1^{\mathcal{A}}(v(x_1^1), \dots, v(x_{m_1}^1)), \dots, t_n^{\mathcal{A}}(v(x_1^n), \dots, v(x_{m_n}^n))) \\ &= t^{\mathcal{A}}(v(x_1), \dots, v(x_n)) \end{aligned}$$

□

A completely analogous result holds for inequations.

We now consider the canonical extensions $(t^{\mathcal{A}})^{\sigma}$ of the term functions $(t)^{\mathcal{A}}$. For example, for a signature containing a single unary operator $\Diamond$, the term $t = \neg \Diamond p \in \mathsf{T}(V)$ defines the maps $\lambda p. (\neg \Diamond p)^{\mathcal{A}} : A \rightarrow A$ and $\lambda p. (\neg \Diamond p)^{\mathcal{A}\sigma} : A^{\sigma} \rightarrow A^{\sigma}$. Note that unlike [Jón94], the terms $t \in \mathsf{T}(V)$ need not define an isotone map since the definition of canonical extension of Eq. (2.3) is suitable for any map. Note also that the idea of interpreting all propositional variables as the identity function makes sense from the point of view of quantification over valuations: a propositional variable can be mapped to any element of A and in this sense propositional variables are indistinguishable.

Following [Jón94] we say for any $t \in \mathsf{T}(V)$ that:

- t is **stable** if $(t^{\mathcal{A}})^{\sigma} = t^{\mathcal{A}\sigma}$
- t is **expanding** if $(t^{\mathcal{A}})^{\sigma} \leq t^{\mathcal{A}\sigma}$
- t is **contracting** if $(t^{\mathcal{A}})^{\sigma} \geq t^{\mathcal{A}\sigma}$

for any $\mathcal{A}$ (and valuation $\pi : V \rightarrow \mathbb{A}$). The inequality between maps is taken pointwise. The following trivial propositions illustrate the usefulness of these notions.

Proposition 2.38. *If $s, t \in \mathsf{T}(V)$ are stable then the equation $s = t$ is canonical. Similarly, if s is contracting and t is expanding, then the inequality $s \leq t$ is canonical.*

Proof. Let $\mathcal{A}$ be an arbitrary object in $\mathbf{AE}(\Sigma)$. If $\mathcal{A} \models s = t$ then $s^{\mathcal{A}} = t^{\mathcal{A}}$ by Proposition 2.37, therefore $(s^{\mathcal{A}})^{\sigma} = (t^{\mathcal{A}})^{\sigma}$ and thus $s^{\mathcal{A}^{\sigma}} = t^{\mathcal{A}^{\sigma}}$ by stability, and it follows that $\mathcal{A}^{\sigma} \models s = t$ by Proposition 2.37.

Similarly, if $\mathcal{A} \models s \leq t$ then $s^{\mathcal{A}} \leq t^{\mathcal{A}}$ by Proposition 2.37 and thus $(s^{\mathcal{A}})^{\sigma} \leq (t^{\mathcal{A}})^{\sigma}$. By the assumptions on s and t , this means that we also have $s^{\mathcal{A}^{\sigma}} \leq t^{\mathcal{A}^{\sigma}}$, and thus $\mathcal{A}^{\sigma} \models s \leq t$ by Proposition 2.37. $\square$

If we can determine if the terms of an (in)equation are stable, expanding or contracting, we can thus also determine if it is canonical. It is easy to see that stable terms always exist.

Lemma 2.39. *The meet and join operations are stable.*

Proof. Take $t = p \wedge q \in \mathsf{T}(V)$, then consider $t^{\mathcal{A}} = \lambda p \lambda q. (p \wedge q)^{\mathcal{A}}$. Since $\wedge^{\mathcal{A}}$ is isotone in both argument, we have by definition

$$\begin{aligned} (t^{\mathcal{A}})^{\sigma}(x, y) &= \bigvee \{ \bigwedge \{ t^{\mathcal{A}}(a, b) \mid p \leq a \in A, q \leq b \in A \} \mid x \geq p \in K, y \geq q \in K \} \\ &= \bigvee \{ \bigwedge \{ a \wedge^{\mathcal{A}} b \mid p \leq a \in A, q \leq b \in A \} \mid x \geq p \in K, y \geq q \in K \} \\ &= \bigvee \{ \bigwedge \{ a \wedge^{\mathcal{A}^{\sigma}} b \mid p \leq a \in A, q \leq b \in A \} \mid x \geq p \in K, y \geq q \in K \} \\ &= \bigvee \{ \bigwedge \{ a \mid p \leq a \in A \} \wedge^{\mathcal{A}^{\sigma}} \bigwedge \{ b \mid q \leq b \in A \} \mid x \geq p \in K, y \geq q \in K \} \\ &= \bigvee \{ p \wedge^{\mathcal{A}^{\sigma}} q \mid x \geq p \in K, y \geq q \in K \} \\ &= \bigvee \{ p \mid x \geq p \in K \} \wedge^{\mathcal{A}^{\sigma}} \bigvee \{ q \mid x \geq q \in K \} \\ &= x \wedge^{\mathcal{A}^{\sigma}} y = t^{\mathcal{A}^{\sigma}}(x, y) \end{aligned}$$

where we've used infinite distributivity (which holds for all canonical extension as was noted in Section 2.1) and the fact that $\mathbb{A}$ is a sub-object of $\mathbb{A}^{\sigma}$ and thus $\wedge^{\mathcal{A}} = \wedge^{\mathcal{A}^{\sigma}}$ on elements of A . The proof for joins is almost identical. $\square$

By using Definition (2.3) for the extension of a map we can also show:

Lemma 2.40. *In the case of boolean algebras, the complementation operation is stable.*

Proof. Let $t = \neg p \in \mathsf{T}(V)$, then $t^{\mathcal{A}} = \lambda p. (\neg p)^{\mathcal{A}}$ and since $\neg^{\mathcal{A}}$ is antitone we

have

$$\begin{aligned}
(t^{\mathcal{A}})^{\sigma}(x) &= (\neg^{\mathcal{A}})^{\sigma}(x) \\
&= \bigvee \{ \bigwedge \{ \neg^{\mathcal{A}}(a) \mid u \geq a \in A \} \mid x \leq u \in O \} \\
&= \bigvee \{ \bigwedge \{ \neg^{\mathcal{A}^{\sigma}}(a) \mid u \geq a \in A \} \mid x \leq u \in O \} \\
&= \bigvee \{ \neg^{\mathcal{A}^{\sigma}} \bigvee \{ a \mid u \geq a \in A \} \mid x \leq u \in O \} \\
&= \bigvee \{ \neg^{\mathcal{A}^{\sigma}} u \mid x \leq u \in O \} \\
&= \neg^{\mathcal{A}^{\sigma}} \bigwedge \{ u \mid x \leq u \in O \} \\
&= \neg^{\mathcal{A}^{\sigma}} x = (t)^{\mathcal{A}^{\sigma}}(x)
\end{aligned}$$

where we have used the infinite version of the de Morgan laws in a CABA and the fact that since $\mathbb{A}$ is a sub boolean algebra of $\mathbb{A}^{\sigma}$, $\neg^{\mathcal{A}^{\sigma}} = \neg^{\mathcal{A}}$ on $\mathbb{A}$. $\square$

Moreover, for any signature Σ , all the functions $f_s, s \in \Sigma$ define stable terms in $\mathsf{T}(V)$. Indeed, let f_s be n -ary and $t = f_s(p_1, \dots, p_n) \in \mathsf{T}(V)$, then if we consider $t^{\mathcal{A}} = \lambda p_1 \dots \lambda p_n. f_s(p_1, \dots, p_n) : A^n \rightarrow A$ we have

$$(t^{\mathcal{A}})^{\sigma}(x_1, \dots, x_n) = f_s^{\sigma}(x_1, \dots, x_n) = t^{\mathcal{A}^{\sigma}}(x_1, \dots, x_n)$$

by definition of the canonical extension of a AE $\mathcal{A} = (\mathbb{A}, (f_s)_{s \in \Sigma})$ as the AE $\mathcal{A}^{\sigma} = (\mathbb{A}, (f_s^{\sigma})_{s \in \Sigma})$. We have thus shown that all the building blocks of a modal language are stable. The key to studying canonicity is thus to understand how strings of these building blocks behave, i.e. to understand the interaction between function composition and canonical extension. This explains why we dedicated a large part of the previous sections to this very topic, and our main tool to establish canonicity will be the *Principle of Matching Topologies* of Theorem 2.36.

Armed with this principle we build Table 2.1 which gives a list of sufficient conditions on two unary terms to guarantee the preservation of stability under functional composition. Note that two columns are particularly well-behaved: if f^{σ} preserves non-empty or just up-directed joins, then canonical extension always commutes with functional composition if g^{σ} is monotone. Terms with this property (i.e. whose composition with any monotone term is stable) have been called **conservative** in [Jón94]. We will rely on the results summarized in this table to define a notion of Sahlqvist formulas. Table 2.2 summarizes how the preservation properties we have examined combine under functional composition, together with Table 2.1 it will allow us to build well-behaved terms in the coming section.

Table 2.1: Stability inheritance of unary terms under composition ($f^\sigma \circ g^\sigma$)

f^σ g^σ		Preservation					Anti-preservation				
		$\leq$	$\wedge$	$\vee$	d.d. $\wedge$	u.d. $\vee$	$\leq$	$\wedge$	$\vee$	d.d. $\wedge$	u.d. $\vee$
Preservation	$\leq$	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗
	$\wedge$	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗
	$\vee$	✗	✗	✓	✗	✓	✓	✓	✓	✓	✓
	d.d. $\wedge$	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗
	u.d. $\vee$	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗
Anti-preservation	$\leq$	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗
	$\wedge$	✗	✗	✓	✗	✓	✓	✓	✓	✓	✓
	$\vee$	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗
	d.d. $\wedge$	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗
	u.d. $\vee$	✗	✗	✓	✗	✓	✗	✗	✗	✗	✗

Table 2.2: Property inheritance under composition ($f^\sigma \circ g^\sigma$)

f^σ	g^σ	Preservation					Anti-preservation				
		$\leq$	$\wedge$	$\vee$	$\text{d.d.} \wedge$	$\text{u.d.} \vee$	$\leq$	$\wedge$	$\vee$	$\text{d.d.} \wedge$	$\text{u.d.} \vee$
Preservation	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$
	$\wedge$	$\leq$	$\wedge$	$\leq$	$\text{d.d.} \wedge$	$\leq$	$\leq$	$\text{a-}\wedge$	$\leq$	$\text{a-d.d.} \wedge$	$\leq$
	$\vee$	$\leq$	$\leq$	$\vee$	$\leq$	$\text{u.d.} \vee$	$\leq$	$\leq$	$\text{a-}\vee$	$\leq$	$\text{a-u.d.} \vee$
	$\text{d.d.} \wedge$	$\leq$	$\text{d.d.} \wedge$	$\leq$	$\text{d.d.} \wedge$	$\leq$	$\leq$	$\text{a-d.d.} \wedge$	$\leq$	$\text{a-d.d.} \wedge$	$\leq$
	$\text{u.d.} \vee$	$\leq$	$\leq$	$\text{u.d.} \vee$	$\leq$	$\text{u.d.} \vee$	$\leq$	$\leq$	$\text{a-u.d.} \vee$	$\leq$	$\text{a-u.d.} \vee$
Anti-preservation	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$	$\leq$
	$\wedge$	$\leq$	$\leq$	$\text{a-}\wedge$	$\leq$	$\text{a-d.d.} \wedge$	$\leq$	$\leq$	$\vee$	$\leq$	$\text{d.d.} \wedge$
	$\vee$	$\leq$	$\text{a-}\vee$	$\leq$	$\text{a-u.d.} \vee$	$\leq$	$\leq$	$\vee$	$\leq$	$\text{u.d.} \vee$	$\leq$
	$\text{d.d.} \wedge$	$\leq$	$\leq$	$\text{a-d.d.} \wedge$	$\leq$	$\text{a-d.d.} \wedge$	$\leq$	$\leq$	$\text{d.d.} \wedge$	$\leq$	$\text{d.d.} \wedge$
	$\text{u.d.} \vee$	$\leq$	$\text{a-u.d.} \vee$	$\leq$	$\text{a-u.d.} \vee$	$\leq$	$\leq$	$\text{u.d.} \vee$	$\leq$	$\text{u.d.} \vee$	$\leq$

$\leq$ stands for isotonicity, $\geq$ for antitonicity, 'u.d.' for up-directed, 'd.d.' for down-directed and 'a-' for anti-preservation.

2.5 Sahlqvist identities

We've seen in the previous section how to build stable terms and thus canonical equations. However, there exists a more general technique to build canonical equations, where not all terms need to be stable: the theory of Sahlqvist formulas. This theory was first laid out in 1975 by Henrik Sahlqvist in the now legendary [Sah75]. This theory was developed in the context of classical 'box-diamond style' modal logic. The theory was then generalised to the context of arbitrary boolean Algebras with Operators in [Jón94] and [dRV95], and subsequently to more general structures [GNV05, Teh08] and to other completions of boolean algebras [GV99]. Here our aim is to generalize the theory to boolean Algebras Expansions, with the hope of dealing with some non-normal modal logics. Our path is straightforward: we will follow the algebraic treatment of [Jón94, Ven06] but without assuming that the maps in the signature are operators.

The algebraic treatment of Sahlqvist formulas hinges on the notion of quasi-equation and how they can be reduced to equations in the context of AEs. Let $\mathsf{T}(V)$ be the free AE for a signature Σ . A **quasi-equation** is a set of $n + 1$ pairs of terms $\{(s_1, t_1), \dots, (s_n, t_n), (u, v)\} \subset \mathsf{T}(V) \times \mathsf{T}(V)$, denoted by $s_1 = t_1, \dots, s_n = t_n \Rightarrow u = v$. A quasi-equation is interpreted as being valid in a AE $\mathcal{A}$ if for any interpretation map $(\cdot)^{\mathcal{A}} : \mathsf{T}(V) \rightarrow \mathcal{A}$,

$$s_1^{\mathcal{A}} = t_1^{\mathcal{A}}, \dots, s_n^{\mathcal{A}} = t_n^{\mathcal{A}} \text{ implies } u^{\mathcal{A}} = v^{\mathcal{A}}$$

in which case we write $s_1^{\mathcal{A}} = t_1^{\mathcal{A}}, \dots, s_n^{\mathcal{A}} = t_n^{\mathcal{A}} \Rightarrow u^{\mathcal{A}} = v^{\mathcal{A}}$. Just as an equation, a quasi-equation is canonical if

$$s_1^{\mathcal{A}^\sigma} = t_1^{\mathcal{A}^\sigma}, \dots, s_n^{\mathcal{A}^\sigma} = t_n^{\mathcal{A}^\sigma} \Rightarrow u^{\mathcal{A}^\sigma} = v^{\mathcal{A}^\sigma}$$

whenever

$$s_1^{\mathcal{A}} = t_1^{\mathcal{A}}, \dots, s_n^{\mathcal{A}} = t_n^{\mathcal{A}} \Rightarrow u^{\mathcal{A}} = v^{\mathcal{A}}$$

We will only need quasi-equations of the shape $s = \perp \Rightarrow t = u$.

We now show how the validity of a quasi-equation over some variety is equivalent to the validity of a certain equation over a certain collection of AEs. This makes precise Proposition 1.2 of [Jón94] which also appears, in a slightly different format, in [Ven06]. Let Σ be a signature and $\mathsf{T}(V)$ be the corresponding free Σ -AE and let $s, t, u \in \mathsf{T}(V)$. We are concerned with the validity over a variety $\mathbb{V}$ of Σ -AEs of the quasi-equation $s = \perp \Rightarrow t = u$. Consider now the signature $\Sigma \cup \{\mathbf{g}\}$ with $\mathbf{g}$ unary and the corresponding free $\Sigma \cup \{\mathbf{g}\}$ -AE $\mathsf{T}'(V)$. We first build a variety of $\Sigma \cup \{\mathbf{g}\}$ -AE by quotienting $\mathsf{T}'(V)$ with all the equations defining the variety $\mathbb{V}$, let us call this variety $\mathbb{V}'$. We then consider the collection of AEs in $\mathbb{V}'$ which satisfy:

$$\mathbf{g}(\perp) = \perp \text{ and } \mathbf{g}(x) = \top \text{ where } \perp \neq x \in \mathsf{T}'(V)$$

Note that this collection is *not* a variety, since $x \neq \perp$ cannot be expressed equationally. We will write $\mathcal{C}$ for this collection. Note also that every Σ -AE $\mathcal{A} \in \mathbb{V}$ trivially gives rise to an element $\mathcal{A}_{\mathbf{g}} \in \mathcal{C}$ by simply appending the operator $\mathbf{g}$ to $\mathcal{A}$ and declaring that it satisfies the condition above. Similarly any valuation $\pi : V \rightarrow \mathcal{A}$ trivially gives rise to an interpretation map

$$(\cdot)^{\mathcal{A}_{\mathbf{g}}} : \mathsf{T}'(V) \rightarrow \mathcal{A}_{\mathbf{g}}$$

by using $(\cdot)^{\mathcal{A}}$ and defining $(\mathbf{g})^{\mathcal{A}_{\mathbf{g}}}$ in the obvious way. It is now straightforward to check from the construction that for any $\mathcal{A} \in \mathbb{V}$ we have

$$s^{\mathcal{A}} = \perp \implies t^{\mathcal{A}} = u^{\mathcal{A}} \quad \text{iff} \quad t^{\mathcal{A}_{\mathbf{g}}} \vee (\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}} = u^{\mathcal{A}_{\mathbf{g}}} \vee (\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}}$$

We can now use this relation between quasi-equation and equation to investigate the canonicity of quasi-equations.

Lemma 2.41. *Let $s \in \mathsf{T}(V)$ be an expanding term, and let $t \in \mathsf{T}(V)$ be a stable term, then the quasi-equation $s = \perp \Rightarrow t = \perp$ is canonical.*

Proof. Assume that $s^{\mathcal{A}} = \perp \Rightarrow t^{\mathcal{A}} = \perp$. By the preceding construction this is equivalent to

$$t^{\mathcal{A}_{\mathbf{g}}} \vee (\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}} = (\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}}$$

or in other words,

$$t^{\mathcal{A}_{\mathbf{g}}} \leq (\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}}$$

We need to check that $t^{\mathcal{A}_{\mathbf{g}}^{\sigma}} \leq (\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}^{\sigma}}$. Since t is stable we have $(t^{\mathcal{A}_{\mathbf{g}}})^{\sigma} = t^{\mathcal{A}_{\mathbf{g}}^{\sigma}}$ and thus

$$t^{\mathcal{A}_{\mathbf{g}}^{\sigma}} \leq ((\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}})^{\sigma}$$

Notice now that $\mathbf{g}$ trivially preserves finite joins, and thus by the results of the previous section, and Table 2.1 in particular, it is clear that irrespective of s ,

$$((\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}})^{\sigma} = (\mathbf{g}^{\mathcal{A}_{\mathbf{g}}})^{\sigma} (s^{\mathcal{A}_{\mathbf{g}}})^{\sigma}$$

Since $\mathbf{g}$ is one of the extension maps, it is stable, and since s is assumed to be expanding we thus get

$$t^{\mathcal{A}_{\mathbf{g}}^{\sigma}} \leq ((\mathbf{g}(s))^{\mathcal{A}_{\mathbf{g}}})^{\sigma} \leq \mathbf{g}^{\mathcal{A}_{\mathbf{g}}^{\sigma}} (s^{\mathcal{A}_{\mathbf{g}}})^{\sigma} \leq \mathbf{g}^{\mathcal{A}_{\mathbf{g}}^{\sigma}} (s^{\mathcal{A}_{\mathbf{g}}^{\sigma}})$$

which is equivalent to $s^{\mathcal{A}^{\sigma}} = \perp \Rightarrow t^{\mathcal{A}^{\sigma}} = \perp$. $\square$

The following result captures the key idea behind Sahlqvist identity. In a AE defined by an arbitrary signature Σ , we will take this Proposition as the definition of a Sahlqvist identity. However, since we would really like Sahlqvist identities to be syntactically defined, rather than through algebraic properties of their subterms, we will use the terminology of **abstract Sahlqvist identity** in the following result and of **concrete Sahlqvist identity** when it is instantiated purely syntactically in a specific modal language.

Proposition and Definition 2.42 ([Jón94]). *Let $\mathsf{T}(V)$ be the modal language defined by a signature Σ . Every identity in $\mathsf{T}(V)$ of the type*

$$s(t_1, \dots, t_n, \neg u_1, \dots, \neg u_m) = \perp \quad (2.10)$$

*where $t_1, \dots, t_n$ are stable terms, $u_1, \dots, u_m$ are expanding terms, and $s \in$ is a conservative term will be called an **abstract Sahlqvist identity**. Every abstract Sahlqvist identity is canonical.*

Proof. Let us define

$$\begin{aligned} v &= (u_1 \wedge p_1) \vee \dots \vee (u_n \wedge p_m) \\ w &= s(t_1, \dots, t_n, p_1, \dots, p_m) \end{aligned}$$

where $p_1, \dots, p_m \in V$ are fresh variables not occurring in any of the terms. Then the identity (2.10) is equivalent to the quasi-identity $v = \perp \Rightarrow w = \perp$. Theorem 5.5 in [Jón94] shows that v must be expanding and, by definition of stability and of being conservative, w must be stable. The results then follows from Lemma 2.41. $\square$

The key to a practical use of Proposition 2.42 is to isolate, for a given logic, syntactically defined classes of stable, expanding and conservative terms. Of course, going from abstract to concrete Sahlqvist identities will always be specific to the language at hand. Here we have gathered results for languages whose operations satisfy some of the preservation properties listed in Table 2.1, with the case of Graded Modal Logic and Probability Logic, detailed later in this Chapter, in mind. As we shall detail shortly, Tables 2.1 and 2.2 provide a tool for building stable and conservative terms in such languages. However, we do not yet have a general technique for building expanding terms. For this we will need the following nomenclature. Let us restrict our attention to the class of AE whose endo-maps $\{f_s\}_{s \in \Sigma}$ are all monotone. We introduce the following dual notions: let $t \in \mathsf{T}(V)$ and $p \in V$ be a (propositional) variable, we will say that a term t in the language associated with the signature Σ is **positive in p** if every occurrence of p in t is in the scope of an even number of antitone operations (including the complementation). Dually, a term $t \in \mathsf{T}(V)$ will be called **negative in p** if every occurrence of p in t is in the scope an odd number of antitone operations. For example, if $f : A^2 \rightarrow A$ is isotone in the first and antitone in the second parameter, then the term $f(p, p)$ is neither positive nor negative in p since the first occurrence of p is under the scope of no antitone operations, whereas the second occurrence is under the scope of one antitone operation. The terms $f(p, \neg p)$ or $f(p, g(p))$ where g is unary and antitone, are both positive in p . A term is **positive** if it is positive in all its propositional variables, and dually for **negative** terms. Positivity and negativity as we have just defined it are interesting concepts because they provide a purely syntactic

description of isotone and antitone terms as the following well-known result shows.

Proposition 2.43. *Let Σ be a signature containing only functions which are monotone in each of their arguments, let $\mathsf{T}(V)$ be the free Σ -AE and let $t \in \mathsf{TV}$. If t is positive in p , then it is isotone in p , if it is negative in p , then it is antitone in p .*

Proof. We prove both parts simultaneously by induction on the complexity of t . The base case for terms positive in p is easy since t is then the propositional variable p , which is trivially isotone since the identity map is isotone. The base case for terms negative in p is either $\neg p$ or a term $f_s(q_1, \dots, q_{i-1}, p, q_{i+1}, \dots, q_n)$ where $f_s, s \in \Sigma$ is antitone in its i^{th} argument. Both cases are clearly antitone.

Now for the inductive step. We can treat all cases in one go. Assume $t = f(s_1, \dots, s_{n_i}, s_{n_i+1}, \dots, s_{n_i+n_a})$ where n_i is the number of isotone parameters of f , and n_a is the number of antitone parameters. For example if $f = \wedge$, then $n_i = 2, n_a = 0$ and if $f = \neg$, then $n_i = 0, n_a = 1$, and similarly for the monotone maps $f_s, s \in \Sigma$.

If t is positive in p , then each $s_i, 1 \leq i \leq n_i$ must be positive in p , and thus by induction isotone in p . Dually each $s_i, n_i + 1 \leq i \leq n_i + n_a$ must be negative in p and thus antitone in p by the induction hypothesis. The term t is then clearly isotone in p . The case where t is negative in p works in exactly the same way. $\square$

The following key result is Theorem 5.5 of [Jón94] (Theorem 7.20 of [Ven06]).

Theorem 2.44. *If a signature Σ contains only isotone operations, then every positive term is expanding.*

In this instance, we were not able to generalise this important theorem to include antitone operations. The reason is that Lemma 2.23 induces a profound asymmetry in the algebraic theory of canonicity. By working uniformly for all monotone maps, it gives a preferred direction to all inequalities which quickly becomes incompatible with order reversal, i.e. with antitone maps. We illustrate this phenomenon with a simple example: let a signature contain a binary map g which is isotone in its first and antitone in its second argument, unary isotone maps f_1, f_2 and a unary antitone map h . Then the term

$$g(f_1(p), f_2(h(q)))$$

is positive but if we try to apply Lemma 2.23 repeatedly, we get

$$(g(f_1, f_2(h)))^\sigma \leq g^\sigma(f_1^\sigma, (f_2(h))^\sigma) \geq g^\sigma(f_1^\sigma, (f_2^\sigma(h^\sigma)))$$

and it is thus impossible to use the expanding/contracting properties of the subterms to infer this property for the entire term, despite the term being

positive. Because of its full generality, Lemma 2.23 cannot be used for nested sequences of mixed or antitone terms. In consequence, for a signature $\Sigma = \Sigma_{\text{iso}} + \Sigma_{\text{mix+anti}}$ containing antitone and/or mixed operations, our only source of expanding terms will be those terms which are positive for the sub-signature Σ_{iso} .

Having built this (not entirely satisfying) syntactically defined class of expanding terms, let us now show how to build stable and conservative terms. We start with the well-known case of classical modal logic defined by the signature $\Sigma = \{\Diamond, \Box\}$. Since all operators are either join- (the diamond operator $\Diamond$) or meet- (the box operator $\Box$) preserving, we can use Table 2.1 to notice that terms involving exclusively $\Diamond$, meets and joins are stable. These terms are usually defined as the **strictly positive** terms since $\Diamond$ is often considered as the fundamental operator and $\Box$ the derived one. Moreover, by using the distributivity rules of boolean algebras, we have that:

Lemma 2.45. *In the signature $\Sigma = \{\Diamond\}$ where $\Diamond$ is unary and join preserving, the strictly positive terms preserve joins. In particular, the strictly positive terms are conservative*

Proof. By induction on the complexity of the terms. For the base case, note that $\Diamond$ preserves joins, $\wedge$ considered as a binary operation also preserves joins by virtue of the distributivity of $\wedge$ over $\vee$, finally $\vee$ considered as binary operation trivially preserves $\vee$. The inductive step follows for the same reason. The fact that terms preserving joins are conservative follows from Lemma 2.19 and Table 2.1. $\square$

Another conclusion of Table 2.1 is that nested sequences of meet-preserving operators, i.e. nested sequences of $\Box$, are also stable, but are not conservative, so they can play the role of the terms $t_1, \dots, t_n$ in Proposition 2.42. This provides us with the classical definition of Sahlqvist identities.

Proposition and Definition 2.46. *A classical Sahlqvist identity is an equation of the form*

$$s(t_1, \dots, t_n, \neg u_1, \dots, \neg u_m) = \perp$$

where s is a strictly positive term, $t_1, \dots, t_n$ are nested sequences of $\Box$, and the terms $u_1, \dots, u_m$ are positive terms. Classical Sahlqvist identities are canonical.

Proof. This is just a special case of Proposition 2.42. $\square$

Note that we now have a purely syntactic implementation of the abstract notion of Sahlqvist identities, i.e. we now have what we have defined as ‘concrete Sahlqvist identities’. Note also that other stable combinations suggests themselves from Tables 2.1 and 2.2. In fact we can use the two tables as the

basis of an algorithm to generate stable terms: pick two stable terms s and t with one of the preservation or anti-preservation properties, then from Table 2.1 check whether or not the composition st is also stable. If it is, use Table 2.2 update the preservation or anti-preservation property of st , select a new term r and repeat the process. This leads us to the following revised definition of Sahlqvist identities:

Proposition and Definition 2.47. *Given a signature Σ whose operations satisfy any of the preservation properties of Table 2.1, a **general Sahlqvist identity** is an equation between terms in the language defined by Σ of the form*

$$s(t_1, \dots, t_n, \neg u_1, \dots, \neg u_m) = \perp$$

where s is a conservative term built from the Table 2.1, $t_1, \dots, t_n$ are stable terms built from the Table 2.1, and the terms $u_1, \dots, u_m$ are terms which are positive for the sub-signature of Σ consisting of isotone operations only.

Example 2.48. Let us illustrate this definition in the case of BAOs defined by the signature $\{\Diamond\}$.

- As an example of stable terms which are not nested boxes, consider the following formulas: $\Diamond\Box p$ and $\Diamond\Box(p \wedge \neg q)$. From table 2.1 it is immediate that $\Diamond\Box p$ is stable. It is clear by Table 2.1 that since $\neg$ anti-preserves joins and $\wedge$ trivially preserves meets in both arguments that $\wedge(\cdot, \neg(\cdot))$ is a stable term which anti-preserves joins in its second argument. Thus by Table 2.1 $\Box(\wedge(\cdot, \neg(\cdot)))$ is also stable and by Table 2.2 it preserves meets in its first argument and anti-preserves joins in its second. Finally, since $\Diamond$ is conservative and stable, we get that $\Diamond\Box(p \wedge \neg q)$ is stable.
- As an example of general but not classical Sahlqvist formula, consider the term

$$t' = \Diamond p \wedge \Diamond(q \wedge \neg p) \wedge \Diamond(r \wedge \neg p \wedge \neg q)$$

The equation $t' = \top$ clearly enforces the existence of at least three distinct successors in a Kripke frame. In the formalism of Definition 2.47, this equation can be re-written as

$$t = \neg t' = \Box\neg p \vee \Box(\neg q \vee p) \vee \Box(\neg r \vee p \vee q) = \perp$$

It cannot be a classical Sahlqvist term since $\Box(\neg q \vee p)$ for example is neither a nested box nor the negation of a positive formula. However, by using De Morgan's laws as a proof that $\neg$ anti-preserves joins, it is easy to check directly from Table 2.1 that each of the terms $t_1 = \Box\neg p$, $t_2 = \Box(\neg q \vee p)$ and $t_3 = \Box(\neg r \vee p \vee q)$ is stable, and $t = \vee(t_1, t_2, t_3)$ with each t_i stable, is thus a general Sahlqvist formula. Thus the logic $K + (t' = \top)$ is sound and

strongly complete w.r.t. Kripke frames whose elements all have at least three distinct successors. This example clearly generalizes to n -successors for any n .

- As an example of non general Sahlqvist and non-canonical formula, consider McKinsey's axiom

$$\Box \Diamond p \rightarrow \Diamond \Box p$$

It can be re-written as $\vee(\Box \Diamond p, \neg \Diamond \Box p)$ and we immediately see that one subterm is positive in p and the other negative, so the only way this formula could be a general Sahlqvist is if one of the two terms was stable. From Table 2.1 we can easily see that $\Box \Diamond$ is a problematic combination which is not stable, and whilst $\Diamond \Box p$ is stable, it has no preservation property (apart from being isotone) and thus $\neg \Diamond \Box p$ is not in general stable either.

For non-normal logics, we propose using the other rows and columns of Table 2.1 to isolate classes of stable and conservative terms. The next section will provide us with concrete case studies involving most of the types of extension functions studied so far.

2.6 Applications

2.6.1 Graded Modal Logic

In this section we consider Graded Modal Logic (see e.g. [Fin72, FBC85]), or GML for short, as an example of a logic whose modal operators are not operators in the technical sense, i.e. do not preserve joins. This logic is based on the choice of **BA** as a fundamental reasoning structure, and on the signature $\Sigma = \{\langle k \rangle \mid k \in \mathbb{N}\}$ where all the function symbols are unary. In the BNF formalism, GML formulas are thus given by:

$$\phi ::= \perp \mid p \mid \neg \phi \mid \phi \wedge \phi \mid \langle k \rangle \phi, k \in \mathbb{N}$$

The intended interpretation of $\langle k \rangle p$ in a Kripke frame is that it should be true at a point w if w has *at least* k successors where p holds. The function symbol $\langle 1 \rangle$ is therefore equivalent to the $\Diamond$ operator of classical modal logic. We will not use the dual operators, but it is frequent to define $[k] = \neg \langle k \rangle \neg a$ whose interpretation is thus ‘there are *strictly fewer* than k successors where a does not hold’. These operators can thus be understood as counting ‘exceptions’. Clearly $[1]$ is equivalent to $\Box$ in classical modal logic.

Formally the semantics of GML can be given either in terms of Kripke frames or in terms of coalgebras for the so-called ‘bag’ or multiset functor $\mathcal{B} : \mathbf{Set} \rightarrow \mathbf{Set}$. For a Kripke frame (W, R) , a valuation $\pi : V \rightarrow W$ and a point $w \in W$,

the semantics of the terms $\mathsf{T}(V)$ is defined inductively in the usual way for propositional variables and boolean connectives and by:

$$(W, R, w) \models \langle k \rangle p \text{ iff } |\{w' \in W \mid wRw' \text{ and } w' \models p\}| \geq k$$

The coalgebraic semantics of GML is given by the bag functor:

$$\mathcal{B}(W) = \{f : W \rightarrow \mathbb{N} \mid \text{supp}(f) \text{ is finite}\}$$

The modal operators $\langle k \rangle$ are then interpreted as the predicate liftings

$$[\![\langle k \rangle]\!] : \mathcal{P}(W) \rightarrow \mathcal{P}(\mathcal{B}W), U \mapsto \{f \in \mathcal{B}(W) \mid \sum_{x \in U} f(x) \geq k\}$$

and the coalgebraic semantics of GML terms is then defined in the usual inductive way.

GML can be axiomatized in several ways, but we follow the axiomatization of [Fin72] which provides an elegant strong completeness proof.

$$\text{GML1 } \langle 1 \rangle (a \vee b) = \langle 1 \rangle (a) \vee \langle 1 \rangle (b)$$

$$\text{GML2 } \langle 1 \rangle \perp = \perp$$

$$\text{GML3 } \langle k \rangle a \rightarrow \langle l \rangle a, \quad l < k$$

$$\text{GML4 } \langle k \rangle a \leftrightarrow \bigvee_{i=0}^k \langle i \rangle (a \wedge b) \wedge \langle k-i \rangle (a \wedge \neg b)$$

$$\text{GML5 } \neg \langle 1 \rangle (a \wedge \neg b) \rightarrow (\langle k \rangle a \rightarrow \langle k \rangle b)$$

The axiomatization of GML defines a quotient of the free Σ -BAE $\mathsf{T}(V)$, and thus a variety of Σ -BAEs, which we denote by $\mathbf{Q}_{GML}$.

Theorem 2.49 ([Fin72, FBC85]). *GML is sound and strongly complete w.r.t. to its Kripke frame semantics, in other words for any term $t \in \mathsf{T}(V)$, $(t)^{\mathbf{Q}_{GML}} = \top$ iff $(W, R) \models t$ for any Kripke frame (W, R)*

It is easy to verify from the above theorem that the expansions $\langle k \rangle$ are isotone for any $k \in \mathbb{N}$. Moreover, the following proposition captures a fundamental distributivity rule satisfied by each $\langle k \rangle$. This distributivity rule will be, via Theorem 2.18, key to building Sahlqvist GML terms via Table 2.1.

Proposition 2.50 (GML distributivity). *For each k , $\langle k \rangle$ is k -additive.*

Proof. We use a semantic argument in conjunction with the soundness and completeness result of Theorem 2.49. To show

$$\langle k \rangle \bigvee X \leq \bigvee \{\langle k \rangle U \mid U \in \mathcal{P}_k(X)\}$$

is easy: let (W, R) be a Kripke frame and $w \in W$, then if $(w, W, R) \models \langle k \rangle \bigvee X$, then there are at least k successors of w which land in $\bigvee X$. Pick k of them, each of them must land in at least one of the elements $a \in X$. By choosing one such element $a(i)$ for each $1 \leq i \leq k$ we define subset U of X of cardinality at most k and it is clear that $x \models \langle k \rangle \bigvee U$. The opposite direction follows from the monotonicity of $\langle k \rangle$ and the observation that $\bigvee U \leq \bigvee X$ for any $U \in \mathcal{P}_k(X)$. $\square$

Note that for $k = 1$, the distributivity law of Proposition 2.50 is the usual join-preservation property of $\diamond$. Note also that if $k = |X|$, we would just have the triviality $\langle k \rangle \bigvee X = \langle k \rangle \bigvee X$ since one of the functions ϕ can then pick every element of X .

Recall from the comments made before Theorem 2.18 that if $\langle k \rangle^\sigma$ is completely k -additive, then in particular it preserves up-directed joins, i.e. it is $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous. From Table 2.1, it is clear that, for all intent and purposes, functions preserving up-directed joins are just as good as functions preserve non-empty joins. In particular terms built from meets, joins and up-directed join preserving functions are stable and conservative. Note that GML has modal operators satisfying each of the preservation properties detailed in Table 2.1: $\langle 1 \rangle$ preserves joins, $\langle k \rangle, k > 1$ preserve up-directed joins, $[0]$ preserves meets and $[k], k > 1$ preserves down-directed meets. We can now define Sahlqvist formulas for GML.

Proposition and Definition 2.51. *A general Sahlqvist identity for GML is an equation of the form*

$$s[t_1, \dots, t_n, \neg u_1, \dots, \neg u_m] = \perp$$

where s is strictly positive (i.e. it is built from $\vee, \wedge, \langle k \rangle$ and propositional variables), $t_1, \dots, t_n$ are stable terms obtained from Table 2.1 and $u_1, \dots, u_m$ are positive terms. Sahlqvist identities for GML are canonical.

Proof. The proof follows immediately from Proposition 2.42 and the fact that since the functions $\langle k \rangle^\sigma$ preserve up-directed joins by Theorem 2.18, strictly positive terms are both stable and conservative. $\square$

Example 2.52. Let us consider a scenario with a Description Logic flavour. Assume we have a GML with several classes of graded operators intended to represent genealogical relations with number restrictions. Concretely we have a $\langle k \rangle^{\text{Child}}$ operator, with the intended interpretation of $\langle k \rangle^{\text{Child}} p$ being ‘has at most k children satisfying p ’, and an operator $\langle k \rangle^{\text{G-Child}}$ whose interpretation is similar but with grand-children instead. Now it would be natural that any model for such a logic should have cardinality restrictions of the type:

$$\langle k_1 \rangle^{\text{Child}} \langle k_2 \rangle^{\text{Child}} p \rightarrow \langle k_1 \cdot k_2 \rangle^{\text{G-Child}} p$$

Such a formula is easily seen to belong to the class of general Sahlqvist formula for GML defined above since $\langle k_1 \rangle^{\text{Child}} \langle k_2 \rangle^{\text{Child}} p$ is stable and $\neg \langle k_1 \cdot k_2 \rangle^{\text{G-Child}} p$ is negative in p .

2.6.2 Intuitionistic Logic and Distributive Substructural Logics

In this section we will provide examples of logics build on **BDL** and **DL** and containing non-isotone operators. We will also use these as examples of logics presented in the abstract flavour of coalgebraic logic.

We start with the well-know **intuitionistic logic**. Algebraically, intuitionistic logic over a set V of propositional variables is the free **Heyting algebra** over V . However, it can also be viewed as a positive modal logic with a binary operator $\rightarrow$, i.e. of modal logic based on the free distributive lattice over V (since Heyting algebras are necessarily distributive). Algebraically, this amounts to viewing a Heyting algebra as a DLE of a particular kind. This is the point of view which we adopt here and we therefore define the following syntax constructor for any $f; A \rightarrow B$ in **BDL**:

$$L_{\text{Hey}} : \mathbf{BDL} \rightarrow \mathbf{BDL}, \begin{cases} A \mapsto \mathbf{F}\{a \rightarrow b \mid a, b \in \mathbf{UA}\} / \equiv \\ L_{\text{RL}} f : L_{\text{Hey}} A \rightarrow L_{\text{Hey}} B, [a]_{\equiv} \mapsto [f(a)]_{\equiv} \end{cases}$$

where $\equiv$ is the fully invariant equivalence relation in **BDL** (see Section 1.4.4) generated by the following Heyting Distribution Laws:

$$\text{HDL1 } a \rightarrow (b \wedge c) = (a \rightarrow b) \wedge (a \rightarrow c)$$

$$\text{HDL2 } (a \vee b) \rightarrow c = (a \rightarrow c) \wedge (b \rightarrow c)$$

The language defined by L_{Hey} for a set V of propositional variables is the free L_{Hey} -algebra over $\mathbf{FV}$, i.e. the language of intuitionistic propositional logic quotiented by the axioms of bounded distributive lattices and HDL1-2. Note that an L_{Hey} -algebra is *not* a Heyting algebra, the axioms HDL1-2 only capture some of the Heyting algebra structure. Instead, an L_{Hey} -algebra is simply a distributive lattice with a binary map satisfying the distribution laws above (which happen to be valid in HAs), i.e. a DLE. The remaining features of HAs will be captured in a second stage in Chapter 5 via *canonical frame conditions*, for reason which will become clear. The reason for choosing **BDL** rather than **DL** is that one of the additional axioms needed to fully capture intuitionistic logic involve $\top$.

Since the (modal) operator $\rightarrow$ preserves meet in one argument and anti-preserves joins in the other, it is isotone in its first argument and antitone in its second, and L_{Hey} -algebras fall within the class of DLEs for which we can

use the results of Sections 2.1-2.5 to define a notion of canonical (in)equation. Indeed this is the reason we have chosen to define L_{Hey} as we did. On the other hand, the notion of Sahlqvist identity for L_{Hey} -algebra is problematic. By the comment following Theorem 2.44, *positive terms* cannot be built from non-isotone expansions. In particular no L_{Hey} -term is positive. Moreover, no L_{Hey} -term can be strictly positive (even though we are working without negation!), since we do not have any conservative operation. Thus a general Sahlqvist identity in L_{Hey} simply amount to an equation

$$t = \perp$$

where t is stable. However, it is perfectly possible to define useful and interesting *canonical (in)equalities* in the language of L_{Hey} -algebra, as we will see in Chapter 5. The format of Sahlqvist identities is simply not very well-suited to L_{Hey} -algebras.

We proceed in an analogous way to define a basic language for **distributive substructural logics** (see [Ono03] for an overview of the topic). In a nutshell, substructural logic disallow some or all of the structural rule of propositional logic, viz. weakening, contraction and exchange. In this setting an additional operator, which we denote by $*$, is added to the language as a kind of $\wedge$ which does not satisfy the structural rule. This operator is required to be left and right-residuated, i.e. there must exist $\backslash$ and $/$ operators such that $a * b \leq c$ iff $b \leq a \backslash c$ iff $a \leq c / a$ and posses a unit I . Algebraically, such a structure corresponds to a **residuated lattice**, i.e. a lattice with a monoidal operation which is left and right-residuated (see [Ono03, GJKO07]. Just as Heyting algebra can be seen as a DLE, a residuated lattice can be seen as a lattice with expansions. We will in fact restrict ourselves to distributive residuated lattices in order to make use of the theory of canonicity developed in the previous Sections, which means that we will view distributive residuated lattices as DLEs. We define the following syntax constructor:

$$L_{\text{RL}} : \mathbf{DL} \rightarrow \mathbf{DL}, \begin{cases} L_{\text{RL}}A = F\{I, a * b, a \backslash b, a / b \mid a, b \in \mathbf{UA}\} / \equiv \\ L_{\text{RL}}f : L_{\text{RL}}A \rightarrow L_{\text{RL}}B, [a]_{\equiv} \mapsto [f(a)]_{\equiv} \end{cases}$$

where $\equiv$ is the fully invariant equivalence relation in $\mathbf{DL}$ (see Section 1.4.4) generated by the Distribution Laws:

$$\text{DL1 } (a \vee b) * c = (a * c) \vee (b * c)$$

$$\text{DL2 } a * (b \vee c) = (a * b) \vee (a * c)$$

$$\text{DL3 } a \backslash (b \wedge c) = (a \backslash b) \wedge (a \backslash c)$$

$$\text{DL4 } (a \vee b) \backslash c = (a \backslash c) \wedge (b \backslash c)$$

$$\text{DL5 } (a \wedge b)/c = (a/c) \wedge (b/c)$$

$$\text{DL6 } a/(b \vee c) = (a/b) \wedge (a/c)$$

The language defined by L_{RL} is the free L_{RL} -algebra over $\mathbf{FV}$, which is the language of the full Lambek calculus (or residuated lattices) quotiented under the axioms of **DL** and DL1-6. An L_{RL} -algebra is simply an object of **DL** endowed with a nullary operation I and binary operations $*$, $\backslash$ and $/$ satisfying the distribution laws above, i.e. a DLE whose expansions are amenable to the techniques developed earlier in the Chapter. Again, note that an L_{RL} -algebra is *not* a distributive residuated lattice. Only some features of this structure have been captured by the axioms above. But several are still missing, and will be added in Chapter 5 as *canonical frame conditions*. Note again that $\backslash$ and $/$ are isotone in one argument but antitone in the other.

We now proceed to define a notion of general Sahlqvist identity for the language L_{RL} -algebras. The positive L_{RL} -terms are those built using $I, \wedge, \vee$ and $*$ only, since $\backslash$ and $/$ are not isotone. But since $*$ is like a binary $\diamond$ operator (it preserves joins in each argument), and negation is not part of the language, the positive terms are also the strictly positive ones. Consequently we have

Proposition and Definition 2.53. *A **general Sahlqvist formula** for distributive substructural logics is an equation of the form*

$$s[t_1, \dots, t_n, u_1, \dots, u_m] = \perp$$

where $s, u_1, \dots, u_m$ are positive terms, and $t_1, \dots, t_n$ are stable terms obtained from Table 2.2. *Sahlqvist identities for distributive substructural logics are canonical.*

As we will see in Chapter 5, the notions of canonical equation and inequation are in fact more useful and better suited to the language of L_{RL} -algebras.

Chapter 3

Functor Presentations

The power of coalgebraic logic comes from its high level of abstraction, both at the syntactic and at the semantic level. In its most abstract formalism (promoted by Alexander Kurz, see e.g. [KKP05, KR07, KP11]) a coalgebraic logic and its interpretation are determined by an endofunctor $L : \mathbf{BA} \rightarrow \mathbf{BA}$, an endofunctor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ and a natural transformation $\lambda : LP \rightarrow PT$. In itself a functor is a very abstract entity, with very simple specifications (viz. it needs to preserve identities and commute with functional composition). Faced with such abstraction the logician may yearn for something concrete to manipulate, guide his/her intuition and prove things with. This is where the idea of representing abstract functors using simple, well known functors often proves useful. This idea of representing abstract entities by using simpler and more well-known ones is pervasive in much of mathematics and it is therefore not surprising that the theory of functor representation should follow very similar lines. We will return to this point in the third section of this chapter, but for now let us point out that whilst we use the term ‘representation’ loosely and intuitively, the main type of functor representation we will study is the notion of functor *presentation* which has a strict technical meaning.

For functors $F : \mathbf{C} \rightarrow \mathbf{Set}$, a natural candidate for the role of simple functors which can be used to represent arbitrary ones is the class of covariant hom functors $\text{hom}(A, -)$, A in $\mathbf{C}$. More generally for a category $\mathbf{C}$ enriched over a category $\mathbf{V}$, it is natural to try to represent functors $F : \mathbf{C} \rightarrow \mathbf{V}$ in terms of the ‘enriched hom functors’ in $\mathbf{V}$. The theory of enriched functors representation has been studied in e.g. [KP93, VK11, AMM12], although from a slightly different perspective to the one developed in what follows. For our purpose we only need enrichment over $\mathbf{Set}$, i.e. no enrichment at all, but the enriched reader might find comfort in knowing that most of what follows can be extended to the enriched setting.

Underlying every type of functor representations we will define and study, is the ability to express a functor as a certain colimit of hom functors. A representation of a functor T in these terms will be called a *presentation* of T . When $\mathbf{C}$ is small, the basic construction is a textbook category theory construction (although it is usually carried out for presheaves, see e.g. [LM92]), and is essentially a consequence of Yoneda’s lemma. For sufficiently well-behaved functors $F : \mathbf{C} \rightarrow \mathbf{Set}$, this construction can be generalized to a large class of categories called accessible categories which are characterised by the fact that they contain a small subcategory which can be used to approximate all its objects. In this case, the notion of functor presentation is also a standard construction which is detailed for example in [AT90, AR94, AGT10]. A general theory of **Set**-endofunctor presentations which includes a characterization of properties of functors in terms of properties of the presentations can be found in [AGT10].

Our small contribution to this already well-developed theory is twofold.

- We clarify the notions of functor presentation which exist in the literature (e.g. [Fre71, AT90, AGT10, KRV05, VK11, AMM12]) and relate them to standard notions from the theory of locally presentable and accessible categories (see [AR94, MP89]) by considering functors as objects in functor categories. From this, we extract the notions of *functor presentation*, *λ -presentable* functors, *λ -presented* functor, *λ -generatable* functor and *λ -generated* functor. We highlight and detail a striking similarity between **Set**-functors and algebraic varieties (already mentioned in passing in [Fre71, Kel82]).
- We isolate and characterize a ‘minimal presentation’ (already introduced, although less systematically, in [AT90]) and thereby study the interaction between presentations and the notion of *base* which is introduced in the first section of this chapter. This notion is key to the formulation of the nabla-style of coalgebraic logic, but is interesting in its own right to understand **Set**-functors.

3.1 The base transformation.

We start this chapter with a slight detour by looking at the notion of base. For **Set**-endofunctors this notion is well understood and underpins the entire edifice of coalgebraic logic in the nabla-style. Definitions and basic properties can thus be found in treatments of the nabla-logics, for example the excellent [KKV12b]. The notion of base for a **Set**-endofunctor is also discussed in some details, although not under the name ‘base’, in [Gum05] which has inspired this section significantly. Intuitively, given a functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$, a set X and an element $\alpha \in TX$, the base of α , which we will denote as $B_X(\alpha)$, is the smallest

subset of X that is required to ‘support’ α , i.e. the collection of elements of X that are ‘ingredients’ in the construction of α . In this section we will formalize and discuss this intuition in as general a context as possible.

More generally, given an arbitrary category $\mathbf{C}$ and a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$, what do we need in order to be able to talk about the ‘support’ of $\alpha \in TC$? Well, we clearly need to be able to talk about subobjects of C , since $B(\alpha)$ is going to be such a subobject. We will therefore require that $\mathbf{C}$ be **well-powered**, i.e. that every object C of $\mathbf{C}$ should possess a small (i.e. set-sized) poset of subobjects denoted by $\text{Sub}(C)$. We refer the reader to e.g. [Mac98] or [Gol06] for more details on the precise construction of the poset $\text{Sub}(C)$. We will denote a subobject U of C either by the object U or by the injection map $i_U^C : U \rightarrowtail C$ depending on the context and we will denote the partial order on $\text{Sub}(C)$ by $\subseteq$. The fact that a subobject i_U^C ‘supports’ α , can now be formalized as

$$\alpha \in Ti_U^C[TU]$$

where $Ti_U^C[TU]$ is the usual direct image (in $\mathbf{Set}$) of TU under Ti_U^C .

Secondly, we would like to identify the subobject of C which *only* contains the ingredients that are *strictly* necessary to build α , in other words we want the smallest possible subobject supporting α . We therefore require that $\mathbf{C}$ should have pullbacks of monos (i.e. intersections when $\mathbf{C} = \mathbf{Set}$). In fact, requiring $\mathbf{C}$ to have (all) pullbacks is probably a more ‘natural’ requirement, and in fact the categories we will be looking at in what follows (viz. locally presentable categories) will even be complete. The definition of the base of α is then given by:

Definition 3.1. Let $\mathbf{C}$ be a well-powered category with pullbacks of monos, let $T : \mathbf{C} \rightarrow \mathbf{Set}$ and let C be an object of $\mathbf{C}$, then the base of $\alpha \in TC$ is the subobject of C defined as

$$B_C(\alpha) = \bigcap \{U \xrightarrow{i} C \mid \alpha \in Ti_U^C[TU]\}$$

where the $\bigcap$ symbol signifies the (limiting object of the) pullback of all the monos in the set, which of course reduces to the usual intersection when $\mathbf{C} = \mathbf{Set}$.

Whilst the object $B_C(\alpha)$ is well-defined in $\mathbf{C}$ as soon as $\mathbf{C}$ has small intersections, it does not necessarily mean that it constitutes a meaningful representation of the aforementioned vision of a smallest subobject from which α can be built, as the following important example shows.

Example 3.2. Let $\mathcal{Q} : \mathbf{Set} \rightarrow \mathbf{Set}^{\text{op}}$ be the contravariant powerset functor, then $\mathcal{Q}^2 : \mathbf{Set} \rightarrow \mathbf{Set}$ is a covariant functor sometimes called the neighbourhood functor. Let us now define the subfunctor of $\mathcal{Q}^2$ defined on sets X as

$$\mathcal{F}X = \{\mathcal{F} \subseteq \mathcal{P}X \mid \mathcal{F} \text{ is a filter}\}$$

and on functions $f : X \rightarrow Y$ as

$$\mathcal{F}f = (f^{-1})^{-1} : \mathcal{F}X \rightarrow \mathcal{F}Y, \mathcal{F} \mapsto \{U \subseteq Y \mid f^{-1}(U) \in \mathcal{F}\}$$

It is easy to verify that $\mathcal{F}f(\mathcal{F})$ is indeed a filter on $\mathcal{P}Y$ and that $\mathcal{F}$ is thus a functor. It is interesting for our purpose to understand how $\mathcal{F}$ acts on inclusion maps. Let $i_U^X : U \hookrightarrow X$ and consider the trivial filter $\mathcal{F}_U = \{\{U\}\} \in \mathcal{F}U$, it is easy to see that

$$\mathcal{F}i_U^X(\mathcal{F}_U) = \{V \subseteq X \mid X \cap U \in \mathcal{F}_U\} = \uparrow U$$

Thus if we apply our definition of B to $\uparrow U$ we get

$$B(\uparrow U) = \bigcap \{V \subseteq X \mid \uparrow U \in \mathcal{F}i_V^X[\mathcal{F}V]\} = U$$

which is in line with the intuition of what the base should be, namely the least amount of information required to build the element $\uparrow U$; and in particular $\uparrow U \in \mathcal{F}i_U^X[\mathcal{F}U]$. However, assume now that X is infinite, and consider the cofinal (or Fréchet) filter $\mathcal{F}_X^c \in \mathcal{F}X$, and let $x \in X$ and $X_x = X \setminus \{x\}$. It is easy to check that if $\mathcal{F}_{X_x}^c$ is the cofinal filter on X_x , then

$$\mathcal{F}i_{X_x}^X(\mathcal{F}_{X_x}^c) = \{V \subseteq X \mid V \cap X_x \in \mathcal{F}_{X_x}^c\} = \mathcal{F}_X^c$$

Note that the same holds for any subset of X to which finitely many elements have been removed. In particular, there is no smallest such subset. Note also that

$$B_X(\mathcal{F}_X^c) = \bigcap \{U \subseteq X \mid \mathcal{F}_X^c \in \mathcal{F}i_U^X[\mathcal{F}U]\} \subseteq \bigcap_{x \in X} X_x = \emptyset$$

This clearly violates the intuition of what the base of an element should be, in particular we have the problematic situation that $\mathcal{F}_X^c \notin \mathcal{F}i_{B_X(\mathcal{F}_X^c)}^X[\mathcal{F}B_X(\mathcal{F}_X^c)]$.

To remedy the pathology illustrated by the previous example we need to add some criterion to the functors for which we want to define a notion of base. The criterion we are interested is very natural in light of the definition of the base, namely the preservation of small intersections. However, let us take this opportunity to define, more generally, the notion of preservation of classes of limits.

Definition 3.3. Let $\mathbf{J}$ be a small category and $D : \mathbf{J} \rightarrow \mathbf{C}$ be a diagram in $\mathbf{C}$. We will say that a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ **preserves D -limits** if every limiting cone of D in $\mathbf{C}$ gets mapped by T to a limiting cone of $T \circ D$ in $\mathbf{Set}$. A functor preserving all small limits is called continuous.

Similarly, we will say that a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ **weakly preserves D -limits** if every limiting cone of D in $\mathbf{C}$ gets mapped by T to a weakly limiting cone of $T \circ D$ in $\mathbf{Set}$. Preservation of D -colimits, co-continuity and weak-preservation of D -colimits are defined similarly.

For our purpose, pullbacks will be the most important type of limits, and in particular we will be most interested in the (weak) preservation pullbacks involving monic arrows, viz. inverse images and small intersections. We will say that a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ **preserves monos** if for any monic arrow $m : U \rightarrowtail A$, $Tm : TU \rightarrow TA$ is also monic, and dually, T **preserves epis** if for any epic arrow: $e : B \twoheadrightarrow A$, $Te : TB \rightarrow TA$ is also epi. The following two sets of dual criteria can be used to determine when a functor preserves monos and or epis.

Proposition 3.4. *If $\mathbf{C}$ is a category where all monos are sections (resp. all epis are retractions), then all functors $T : \mathbf{C} \rightarrow \mathbf{C}'$ preserve monos (resp. epis).*

Proof. Every section is a mono (and every retraction is an epi) and being a section (or a retraction) is clearly preserved by functoriality. $\square$

Proposition 3.5. *If $\mathbf{C}$ is a category that has all kernel pairs (resp. co-kernel pairs) and $T : \mathbf{C} \rightarrow \mathbf{Set}$ is a functor weakly preserving these pairs, then T preserves monos (resp. epis).*

Proof. It is easy to check that a morphism $f : A \rightarrow B$ is a monomorphism iff the following commutative square is a pullback square

$$\begin{array}{ccc} A & \xrightarrow{\text{Id}_A} & A \\ \text{Id}_A \downarrow & & \downarrow f \\ A & \xrightarrow{f} & B \end{array}$$

(and dually for epimorphisms), and the result follows (almost) immediately from the assumption on T . $\square$

Note in particular that weak-pullback preserving functors which are the bread and butter of coalgebraic logic must therefore preserve monos.

Given an arrow $f : A \rightarrow B$ in $\mathbf{C}$ and an element $i_V^B \in \text{Sub}(B)$, the **inverse image** of i_V^B along f is defined as the pullback:

$$\begin{array}{ccc} f^{-1}(V) & \xrightarrow{f_*^V} & V \\ i_*^V \downarrow & & \downarrow i_V^B \\ A & \xrightarrow{f} & B \end{array}$$

A mono-preserving functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ will thus be said to **(weakly) preserve inverse images** if applying T to the diagram above yields a (weak) pullback in $\mathbf{Set}$. As we have already seen, an **intersection** is a (small) pullback of monomorphisms. Thus, we will say that $T : \mathbf{C} \rightarrow \mathbf{Set}$ **(weakly) preserves**

intersections if for any set of monos $\{i_x : U_x \rightarrowtail A \mid x \in X\}$ the cone of the diagram $\{Ti_x : U_x \rightarrowtail A \mid x \in X\}$ given by

$$\left(T \bigcap \{i_x \mid x \in X\}, \{Tj_x : T \bigcap \{i_x \mid x \in X\} \rightarrowtail TU_x\}_{x \in X} \right)$$

is a (weak) pullback. If $T : \mathbf{C} \rightarrow \mathbf{Set}$ preserves intersections we have the following very natural property.

Lemma 3.6. *Let $\mathbf{C}$ be well-powered with pullbacks of monos let $T : \mathbf{C} \rightarrow \mathbf{Set}$ preserve intersections and monos, then for any A in $\mathbf{C}$ and any $\alpha \in TA$:*

$$\alpha \in Ti_{\mathbf{B}_A(\alpha)}^A[TB_A(\alpha)]$$

Proof. Let us write

$$\mathcal{B}(\alpha) = \{Ti_U^A[TU] \mid i_U^A \in \text{Sub}(A), \alpha \in Ti_U^A[TU]\}$$

By definition, for each $Ti_U^A[TU] \in \mathcal{B}(\alpha)$ there exists an element $\beta_{TU} \in Ti_U^A[TU]$ which gets mapped to α along the corresponding maps Ti_U^A . Since T preserve monos, each Ti_U^A is a mono, and they therefore define an intersection $\bigcap \mathcal{B}(\alpha)$. From the fact that this intersection $\bigcap \mathcal{B}(\alpha)$ is a pullback, any collection of elements $\{\beta_{TU}\}_{TU \in \mathcal{B}(\alpha)}$ which all get mapped to α (and there exists at least one) determines a unique element

$$\beta_0 \in \bigcap \mathcal{B}(\alpha)$$

which gets mapped to the various $\beta_{TU} \in TU$ and then to α along the various monos defining the pullback. Since T preserves intersections, we have

$$\beta_0 \in T \left(\bigcap \{i_U^A \in \text{Sub}(A) \mid \alpha \in Ti_U^A[TU]\} \right) = TB_A(\alpha)$$

and thus $Ti_{\mathbf{B}_A(\alpha)}^A(\beta_0) = \alpha$. □

Note that weak-preservation of intersections is actually sufficient.

We will spend the rest of this section showing that, for a large class of functors, the base construction actually defines a natural transformation between T and the Sub functor. To show this we must first turn Sub into a (covariant) functor. This is done by specifying a generalized notion of ‘direct image’ for the category $\mathbf{C}$. Given a morphism $f : A \rightarrow B$ in $\mathbf{C}$, we will define the **direct image** $f[A]$ as the smallest subobject of B through which f factors. Since we’re assuming that $\mathbf{C}$ has pullbacks of monos to define bases, we can build $f[A]$ explicitly as

$$f[A] = \bigcap \{i_V^B : V \hookrightarrow B \mid f \text{ factors through } i_V^B\}$$

Note that the preservation of intersections does not imply the preservation of direct images, since for $f : A \rightarrow B$, there are generally subsets of TB through which Tf factors but which are not of the form TU for $U \in \text{Sub}(B)$. Note also that whilst the direct image $f[A]$ is, strictly speaking, a subobject of B , i.e. an equivalence class of monomorphisms into B , for notational convenience we will often denote it as an object, i.e. $f[A]$, with the understanding that we are really looking at (the equivalence class of) $i_{f[A]}^A$. In order for f to factor through $f[A]$, we need a map $A \rightarrow f[A]$ which we will denote f^* , i.e. f^* will from now on denote the restriction of a morphism to its (direct) image. This construction endows $\mathbf{C}$ with an **orthogonal factorization system**

Proposition 3.7. *Let $\mathbf{C}$ be well-powered and with small pullbacks and let $f : A \rightarrow B$ be a $\mathbf{C}$ -morphism. The decomposition*

$$A \xrightarrow{f} B = A \xrightarrow{f^*} f[A] \xrightarrow{i^*} B$$

has the property that f^ is orthogonal to the class of all monomorphism in $\mathbf{C}$, i.e. for any monomorphism $m : C \rightarrow D$ and arrows $g : A \rightarrow C, h : f[A] \rightarrow D$ such that the following square commutes*

$$\begin{array}{ccc} A & \xrightarrow{f^*} & f[A] \\ g \downarrow & & \downarrow h \\ C & \xrightarrow{m} & D \end{array}$$

there exist a unique diagonal fill-in morphism $e : f[A] \rightarrow C$ such that $g = e \circ f^$ and $h = m \circ e$.*

Proof. Recall we have defined a decomposition $f = i^* \circ f^*$. Let us first show that if $f^* = n \circ k$ for a certain monomorphism n and a certain morphism k , then n is an isomorphism. Note first that $i^* \circ n \circ k = i^* \circ f^* = f$ and since $i^* \circ n$ is the composition of two monos, it is itself a mono through which f factors. But i^* is the smallest such mono, so i^* must factor through $i^* \circ n$, i.e. there must exist p such that $i^* \circ n \circ p = i^*$, and since i^* is mono, this means $n \circ p = \text{Id}$. As a consequence, we also have that $(i^* \circ n) \circ (p \circ n) = i^* \circ (n \circ p) \circ n = i^* \circ n$ and thus $(p \circ n) = \text{Id}$ too, by virtue of $i^* \circ n$ being mono. Thus n is iso.

Now staring from the commutative square above, since $\mathbf{C}$ has pullbacks we

can build the pullback of h along m :

$$\begin{array}{ccccc}
 A & & & & \\
 \swarrow \scriptstyle f^* & & & & \\
 & h^{-1}(C) & \xrightarrow{n} & f[A] & \\
 \searrow \scriptstyle g & \downarrow & & \downarrow \scriptstyle h & \\
 & C & \xrightarrow{m} & D &
 \end{array}$$

By the universal property of the pullback there exist a unique morphism $A \rightarrow h^{-1}(C)$, and since pulling back along a mono produces a mono, we have that f^* factors through a mono n and, by our previous point, this means that this mono n is actually an iso. It is then easy to see that we have an arrow from $f[A] \rightarrow C$. The unicity follows from the fact that inverses are unique up-to-isomorphism and so are pullbacks. $\square$

Note that algebraic theories, and in particular **Set** and **BA**, have a regular epi-mono factorisation. More generally, every locally presentable category (see next section) has strong-epi mono factorisation. Both are orthogonal factorization systems.

If $\mathbf{C}$ is well-powered and has pullbacks we can therefore turn Sub into a covariant functor $\text{Sub} : \mathbf{C} \rightarrow \mathbf{Set}$ by defining it as usual on objects and by setting for any $f : A \rightarrow B$:

$$\text{Sub}(f) : \text{Sub}(A) \rightarrow \text{Sub}(B), \quad i_U^A \mapsto \bigcap \{i_V^B \mid f \circ i_U^A \text{ factors through } i_V^B\}$$

For clarity's sake we need to make the following notational point precise. It is customary mathematical notation to write $\text{Sub}(f)(i_U^A)$ as $f[U]$, but in our very general categorical context this notation can be misleading. Consider again the $\mathbf{C}$ -morphism $f : A \rightarrow B$ and the subobject $i_U^A : U \hookrightarrow A$. Whilst $f[A]$ is the direct image of f , $f[U]$ is the direct image of $f \circ i_U^A$. Thus $f[U]$ is an abuse of notation since U is not the domain of f . In fact we can make this notation both convenient and precise as follows: since $\text{Sub}(f)(i_U^A)$ is the smallest subobject of B through which $f \circ i_U^A$ factors, there must exist a morphism f_U^* such that $\text{Sub}(f)(i_U^A) \circ f_U^* = f \circ i_U^A$. For brevity's sake we will often denote $\text{Sub}(f)(i_U^A)$ by i_U^* . Notice that we are using a notation for the maps defining the limiting cone which is dual to that in the inverse image case, e.g. f_*^U versus f_U^* . With this notation in place, it is clear that $f_U^* = (f \circ i_U^A)^*$, i.e. in **Set**-language f_U^* is the image restriction of f restricted to U .

Let us now check that Sub indeed defines a functor. Consider first the

diagram:

$$\begin{array}{ccc} U & \hookrightarrow & V \\ \downarrow i & & \downarrow \\ A & \xrightarrow{\text{Id}_A} & A \end{array}$$

Clearly, if V is a subobject of A through which $\text{Id}_A \circ i = i$ must factor, then $U \hookrightarrow V$, and the smallest such V is U itself. Thus $\text{Sub}(\text{Id}_A) = \text{Id}_{\text{Sub}(A)}$. Now for composition consider the following diagram:

$$\begin{array}{ccccc} U & \xrightarrow{f_U^*} & f_U^*[U] & \xrightarrow{g_{f_U^*[U]}^*} & g_{f_U^*[U]}^*[f_U^*[U]] \\ \downarrow i_U^A & & \downarrow i_U^* & & \downarrow i_{f_U^*[U]}^* \\ A & \xrightarrow{f} & B & \xrightarrow{g} & C \end{array}$$

In light of the above diagram, the abuse of notation consisting in writing $f[U]$ for $f_U^*[U]$ is here fully justified, and we shall thus use this convention for what follows. By definition, $(g \circ f)[U]$ is the smallest subobject of C through which $g \circ f \circ i_U^A$ factors. On the other hand $g[f[U]]$ is the smallest subobject of C through which $g \circ i_{f[U]}^B$ factors. But that means that $g \circ i_{f[U]}^B \circ f^*$ also factors through $g[f[U]]$, and by commutativity of the above diagram this means that $g \circ f \circ i_U^A$ factors through $g[f[U]]$, and thus $(g \circ f)[U] \hookrightarrow g[f[U]]$. For the converse direction, note that from the diagram above and the definition of $(g \circ f)[U]$ we can extract the following diagram:

$$\begin{array}{ccc} U & \xrightarrow{f_U^*} & f[U] \\ (g \circ f)_U^* \downarrow & & \downarrow g \circ i_U^* \\ (g \circ f)[U] & \hookrightarrow & C \end{array}$$

By Proposition 3.1 we know that f^* is orthogonal to the class of all monomorphisms and thus there must exist a unique diagonal fill-in morphism $f[U] \rightarrow (g \circ f)[U]$. This in turns means that $g \circ i_{f[U]}^B$ factors through $(g \circ f)[U]$ and since $g[f[U]]$ is the smallest subobject of C with this property we have $g[f[U]] \hookrightarrow (g \circ f)[U]$.

The notion of base can now be understood as a transformation $\mathbf{B} : T \rightarrow \text{Sub}$ between the functors $T : \mathbf{C} \rightarrow \mathbf{Set}$ and $\text{Sub} : \mathbf{C} \rightarrow \mathbf{Set}$ defined at each stage A in $\mathbf{C}$ by:

$$\mathbf{B}_A : TA \rightarrow \text{Sub}(A), \quad \alpha \mapsto i_{\mathbf{B}_A(\alpha)}^A = \bigcap \{i_U^A : U \hookrightarrow A \mid \alpha \in Ti_U^A[TU]\}$$

This transformation is in general not natural, but we will now isolate preservation properties for T which guarantee the naturality of $\mathbf{B}$. We first need the following lemmas.

Lemma 3.8. Consider $U \xrightarrow{i_U^V} V \xrightarrow{i_V^A} A$, then $\text{Sub}(i_V^A)(i_U^V) = i_U^A$.

Proof. If $i_W^A \in \text{Sub}(A)$ is such that $i_V^A \circ i_U^V$ factors through i_W^A , then since $i_U^A = i_V^A \circ i_U^V$, it is trivially the case that i_U^A factors through i_W^A and thus $i_U^A \leq i_W^A$ which shows that i_U^A is the smallest such object and the result follows. $\square$

Lemma 3.9. The functor Sub preserves monos.

Proof. Let $i : A \rightarrow B$ be a $\mathbf{C}$ -monomorphism. Since $\text{Sub}(i) : \text{Sub}(A) \rightarrow \text{Sub}(B)$ is an arrow in $\mathbf{Set}$, it will be equivalent to show that it is injective. By the preceding Lemma 3.8, if $i_U^A : U \hookrightarrow A$ is a subobject of A , then $\text{Sub}(i)(i_U^A) = i \circ i_U^A$. Thus if we have two subobjects $i_U^A, i_V^A \in \text{Sub}(A)$ such that $\text{Sub}(i)(i_U^A) = \text{Sub}(i)(i_V^A)$, then clearly $i \circ i_U^A = i \circ i_V^A$, and since i is mono, $i_U^A = i_V^A$. $\square$

Lemma 3.10. The functor Sub preserves intersections.

Proof. Let us consider an arbitrary collection of monos $\{i_x : A_x \rightarrow B \mid x \in X\}$ in $\mathbf{C}$. By definition, any collection of monos $\{j_x : U_x \hookrightarrow A_x \in \text{Sub}(A_x) \mid x \in X\}$ such that $\text{Sub}(i_x)(j_x) = k \in \text{Sub}(B)$ for all $x \in X$ defines an element of $\bigcap \{\text{Sub}(i_x) \mid x \in X\}$, we will show that it also defines a unique element of $\text{Sub}(\bigcap \{i_x \mid x \in X\})$, providing us with a map

$$\bigcap \{\text{Sub}(i_x) \mid x \in X\} \rightarrow \text{Sub}\left(\bigcap \{i_x \mid x \in X\}\right)$$

From which the result follows immediately. So let us consider such a collection of monos $j_x \in \text{Sub}(A_x), x \in X$, from Lemma 3.8 we know that $\text{Sub}(i_x)(j_x) = i_x \circ j_x$, which is a mono into B . Since $\mathbf{C}$ has intersections we can form the intersection $\bigcap \{i_x \circ j_x : U_x \hookrightarrow B \mid x \in X\}$. Since this intersection is a cone for $\{i_x : A_x \rightarrow B \mid x \in X\}$, there must exist a unique monomorphism

$$\bigcap \{i_x \circ j_x : U_x \hookrightarrow B \mid x \in X\} \rightarrow \bigcap \{i_x : A_x \rightarrow B\}$$

But this precisely defines an element of $\text{Sub}(\bigcap \{i_x \mid x \in X\})$, as desired. $\square$

Lemma 3.11. The functor Sub preserves inverse images.

Proof. Let $f : A \rightarrow B$ and $i_V^B : V \hookrightarrow B$ in $\mathbf{C}$. We need to show that $\text{Sub}(f^{-1}(V))$ together with the arrows $\text{Sub}(i_*^V)$ and $\text{Sub}(f_*^V)$ is a pullback for $\text{Sub}(A) \xrightarrow{\text{Sub}(f)} \text{Sub}(B) \xleftarrow{\text{Sub}(i_V^B)} \text{Sub}(V)$. As usual, we pick $m \in \text{Sub}(A)$ and $j_W^V \in \text{Sub}(V)$ such that $\text{Sub}(f)(m) = \text{Sub}(i_V^B)(j_W^V) = i_V^B \circ j_W^V$ and we need to find a unique element of $\text{Sub}(f^{-1}(V))$ which gets mapped to m and j_W^V .

Consider the following commutative diagram:

$$\begin{array}{ccc}
 U & \xrightarrow{(f \circ m)_U^*} & W = (f \circ m)_U^*[U] \\
 & \searrow \text{dashed} & \downarrow j_W^V \\
 & f^{-1}(V) & \xrightarrow{f_*^V} V \\
 & \downarrow i_*^V & \downarrow i_V^B \\
 A & \xrightarrow{f} & B
 \end{array}$$

m (curved arrow from U to A)

Since the bottom square is a pullback and since $(U, \{j_W^V \circ (f \circ m)_U^*, m\})$ forms a cone for the same diagram, there must exist a unique arrow $n : U \rightarrow f^{-1}(V)$. Moreover, this arrow needs to be a monomorphism since i_*^V and m are monomorphisms. We therefore have a unique element of $n \in \text{Sub}(f^{-1}(V))$. It remains to check that it gets mapped to m and j_W^V by $\text{Sub}(i_*^V)$ and $\text{Sub}(f_*^V)$ respectively. By Lemma 3.8 we know that $\text{Sub}(i_*^V)(n) = i_*^V \circ n = m$. Moreover, if $\text{Sub}(f_*^V)(n)$ was a subobject of W , then $f \circ m$ would factor through it and since W is the smallest subobject of B with this property we must have $\text{Sub}(f_*^V)(n) = j_W^V$. $\square$

Theorem 3.12. *Let $\mathbf{C}$ be well-powered with pullbacks and direct images, then $T : \mathbf{C} \rightarrow \mathbf{Set}$ preserves monos, intersections and inverse images iff $\mathbf{B} : T \rightarrow \mathbf{Sub}$ is natural and sub-cartesian.*

Proof. Let us start by assuming that T preserves monos, intersections and inverse images. We must then show that for any $f : A \rightarrow B$ in $\mathbf{C}$ the following diagram commutes.

$$\begin{array}{ccc}
 TA & \xrightarrow{Tf} & TB \\
 \mathbf{B}_A \downarrow & & \downarrow \mathbf{B}_B \\
 \text{Sub}(A) & \xrightarrow{\text{Sub}(f)} & \text{Sub}(B)
 \end{array}$$

Let us pick an $\alpha \in TA$. In order to show that $\text{Sub}(f)(\mathbf{B}_A(\alpha)) \subseteq \mathbf{B}_B(Tf(\alpha))$. First note that by Lemma 3.10, Sub preserves intersections and thus

$$\begin{aligned}
 \text{Sub}(f)(\mathbf{B}_A(\alpha)) &= \text{Sub}(f) \left(\bigcap \{i_U^A \mid \alpha \in Ti_U^A[TU]\} \right) \\
 &= \bigcap \{\text{Sub}(f)(i_U^A) \mid \alpha \in Ti_U^A[TU]\}
 \end{aligned}$$

Now starting with a mono $i_V^B : V \rightarrow B$ such that $Tf(\alpha) \in Ti_V^B[TV]$ let us build the preimage of TV under Tf . Since T preserves preimages we have the

following pullback square:

$$\begin{array}{ccc} Tf^{-1}(TV) = T(f^{-1}(V)) & \xrightarrow{Tf_{Tf^{-1}(TV)} = T(f_{f^{-1}(V)})} & TV \\ \downarrow Ti_{f^{-1}(V)}^A & & \downarrow Ti_V^B \\ TA & \xrightarrow{Tf} & TB \end{array}$$

And since this square is a pullback, the fact that there exist $\beta \in TV$ such that $Ti_V^B(\beta) = Tf(\alpha)$ means that there must exist a $\beta' \in T(f^{-1}(V))$ such that $T(f_{f^{-1}(V)})(\beta') = \beta$ and $Ti_{f^{-1}(V)}^A(\beta') = \alpha$. If we now set $U = f^{-1}(V)$, then $\alpha \in Ti_U^A[TU]$ and $\text{Sub}(f)(i_U^A)$ is the smallest subobject of B through which $f \circ i_U^A$ factors. Clearly V is such a subobject since $i_V^B \circ f_{f^{-1}(V)} = f \circ i_U^A$, so $\text{Sub}(f)(i_U^A) \hookrightarrow V$. Since we can find such an U for any V with $Tf(\alpha) \in Ti_V^B[TV]$ we must have

$$\bigcap \{ \text{Sub}(f)(i_U^A) \mid \alpha \in Ti_U^A[TU] \} \subseteq \bigcap \{ i_V^B \mid Tf(\alpha) \in Ti_V^B[TV] \}$$

i.e. $\text{Sub}(f)(B_A(\alpha)) \subseteq B_B(Tf(\alpha))$.

Now for the opposite direction, i.e. $B_B(Tf(\alpha)) \subseteq \text{Sub}(f)(B_A(\alpha))$. For notational convenience let us write i^* for the mono $i^* : \text{Sub}(f)(B_A(\alpha)) \hookrightarrow B$. It is enough to show that $Tf(\alpha) \in Ti^*[T(\text{Sub}(f)(B_A(\alpha)))]$. But this follows immediately from Lemma 3.6: since $\alpha \in Ti_{B_A(\alpha)}^A[TB_A(\alpha)]$, then there must exist $\beta \in TB_A(\alpha)$ with $Ti_{B_A(\alpha)}^A(\beta) = \alpha$. Since $\text{Sub}(f)(B_A(\alpha))$ is the smallest subobject of B through which $f \circ i_{B_A(\alpha)}^A$ factors, $Tf \circ Ti_{B_A(\alpha)}^A$ factors through $T(\text{Sub}(f)(B_A(\alpha)))$ and so there must exist a $\beta' \in T(\text{Sub}(f)(B_A(\alpha)))$ such that $Ti^*(\beta') = Tf(\alpha)$, which is what we needed to show.

Let us conclude the ‘if’ direction by showing that B is sub-cartesian. Let $m : A \rightarrow B$ be a monic arrow in $\mathbf{C}$, we need to show that the following square is a pullback:

$$\begin{array}{ccc} TA & \xrightarrow{Tm} & TB \\ \downarrow B_A & & \downarrow B_B \\ \text{Sub}(A) & \xrightarrow{\text{Sub}(m)} & \text{Sub}(B) \end{array}$$

Pick an element $\beta \in TB$ and an element $i_U^A \in \text{Sub}(A)$ such that $\text{Sub}(m)(i_U^A) = B_B(\beta)$. From Lemma 3.9 we know that $\text{Sub}(m)$ is injective and from Lemma 3.8 we know that $\text{Sub}(m)(i_U^A) = m \circ i_U^A$. Since T preserves intersections we know from Lemma 3.6 that $\beta \in Ti_{B_B(\beta)}^B[TB_B(\beta)] = T(m \circ i_U^A)[TU]$. Thus there must exist $\alpha \in TU$ such that $Tm \circ Ti_U^A(\alpha) = \beta$, and since T preserves monos, this α must be unique. If we now choose $\alpha' = Ti_U^A(\alpha)$ we get a unique element of TA such that $Tm(\alpha') = \beta$ and $B_A(\alpha') = U$.

3.2.1 The category $\mathbf{El}(T)$.

The Yoneda lemma has been described as the deepest triviality known to humanity. The following subsection adds further weight to this observation by showing how it underlies the theory of functor presentation. We start by recalling the Lemma.

Lemma 3.13 (Yoneda Lemma). *Let $\mathbf{C}$ be a locally small category and $T : \mathbf{C} \rightarrow \mathbf{Set}$ be a functor, then for any object A of $\mathbf{C}$, there exist a bijection*

$$\mathrm{Nat}(\mathrm{hom}(A, -), T) \simeq TA$$

Proof. See for example [Mac98]. □

In the special case where T is also a hom functor, say $\mathrm{hom}(B, -)$, then the Yoneda lemma tells us that

$$(\mathrm{hom}(A, -), \mathrm{hom}(B, -)) \simeq \mathrm{hom}(B, A) \quad (3.1)$$

i.e. to each $\mathbf{C}$ -arrow $f : B \rightarrow A$ corresponds a $\mathbf{Set}^{\mathbf{C}}$ -arrow (i.e. a natural transformation) $\eta_f : \mathrm{hom}(A, -) \rightarrow \mathrm{hom}(B, -)$. Moreover this correspondence is well-behaved with respect to arrow composition and identities. This suggests that the assignment $A \mapsto \mathrm{hom}(A, -)$ in fact defines a contravariant functor. This is indeed the case and this functor is known as the Yoneda embedding. It is defined formally as the functor $\mathbf{Y} : \mathbf{C}^{\mathrm{op}} \rightarrow \mathbf{Set}^{\mathbf{C}}$ defined on objects as $\mathbf{Y}(A) = \mathrm{hom}(A, -)$ and on $\mathbf{C}^{\mathrm{op}}$ -morphisms $f : B \rightarrow A$ by the natural transformation $\mathbf{Y}f : \mathrm{hom}(A, -) \rightarrow \mathrm{hom}(B, -)$ defined by:

$$\mathbf{Y}f_C : h \in \mathrm{hom}(A, C) \mapsto h \circ f \in \mathrm{hom}(B, C)$$

Corollary 3.14. *The Yoneda embedding $\mathbf{Y}$ defined above is a fully faithful functor.*

Proof. This is the content of the isomorphism of Eq. (3.1). □

Given a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ and an object A of $\mathbf{C}$, the Yoneda lemma suggests viewing elements of TA as natural transformations. Moreover, the existence of the Yoneda embedding suggests that since we have morphisms between hom functors, it may also be possible to define morphisms between these natural transformations. As we shall shortly show, this is indeed the case and allows us to view the *elements* of TA as *objects* in a category which is called the **category of elements** of T and denoted $\mathbf{El}(T)$ or $\int T$ (we will use the former notation throughout). The formal construction is as follows: the objects of $\mathbf{El}(T)$ are the natural transformations $\eta : \mathrm{hom}(A, -) \rightarrow T$ for some A in $\mathbf{C}$ and there is a morphism between $\eta : \mathrm{hom}(A, -) \rightarrow T$ and $\eta' : \mathrm{hom}(B, -) \rightarrow T$

whenever there is a $\mathbf{C}$ -morphism $f : B \rightarrow A$, such that the following diagram commute:

$$\begin{array}{ccc} \mathrm{hom}(A, -) & \xrightarrow{\Upsilon f} & \mathrm{hom}(B, -) \\ & \searrow \eta \quad \swarrow \eta' & \\ & T & \end{array}$$

Thus $\mathbf{El}(T)$ is the comma category $\Upsilon \downarrow T$. There is an obvious forgetful functor $\mathbf{U}_T : \mathbf{El}(T) \rightarrow \mathbf{C}^{\mathrm{op}}$ which is defined on objects by:

$$\left(\mathrm{hom}(A, -) \xrightarrow{\eta} T \right) \mapsto A$$

and on arrows by $(\eta \rightarrow \eta') \mapsto f$ where $f : B \rightarrow A$ is the $\mathbf{C}$ -arrow making $\eta' \circ \Upsilon(f) = \eta$.

We can use Yoneda's lemma to give a more intuitive description of the category $\mathbf{El}(T)$ which is often easier to work with. Formally, $\mathbf{El}(T)$ can equivalently be defined as follows: to each $\eta : \mathrm{hom}(A, -) \rightarrow T$ we associate the pair (A, α) where $\alpha \in TA$ is the element corresponding to η by Yoneda, i.e. $\eta_A(\mathrm{Id}_A)$. This assignment is clearly bijective. For arrows, we need to define an arrow $(A, \alpha) \rightarrow (B, \beta)$ for $\beta \in TB$ whenever the natural transformation $\eta' : \mathrm{hom}(B, -) \rightarrow T$ satisfies a commutative triangle as above, i.e. $\eta' \circ \Upsilon(f) = \eta$. In other words we want

$$\alpha = (\eta' \circ \Upsilon f)_A(\mathrm{Id}_A) = \eta'_A(f) = \eta'_A \circ \mathrm{hom}(f)(\mathrm{Id}_B) = Tf \circ \eta'_B(\mathrm{Id}_B) = Tf(\beta)$$

It is then clear that the category whose objects are pairs (A, α) where A is in $\mathbf{C}$ and $\alpha \in TA$ and where there exists an arrow $(A, \alpha) \rightarrow (B, \beta)$ whenever there exists a map $f : B \rightarrow A$ such that $\alpha = Tf(\beta)$ is another implementation of $\mathbf{El}(T)$. Note that the reversal of arrows from $\mathbf{C}$ to $\mathbf{El}(T)$ is conventional and is not universal in the literature, but is convenient for our purpose, since we are interested in covariant functors. But in order to reason about the category $\mathbf{El}(T)$ itself it often is conceptually and notationally convenient to forget about the arrow reversal and consider $\mathbf{El}(T)^{\mathrm{op}}$ instead, since we then have an arrow $(A, \alpha) \rightarrow (B, \beta)$ whenever there exists a map $f : A \rightarrow B$ such that $\beta = Tf(\alpha)$, and then proceed by duality. The category $\mathbf{El}(T)^{\mathrm{op}}$ is also sometimes called the Grothendieck construction, although this terminology usually has a broader meaning in the context of fibrations. As the following examples will illustrate, the category $\mathbf{El}(T)$ can often be seen as a generalisation of the concept of 'pointed space'.

Example 3.15. 1. If T is taken to be the identity functor on $\mathbf{Set}$, then $\mathbf{El}(\mathrm{Id})^{\mathrm{op}}$ is just the category whose objects are pairs (X, x) with a distinguished $x \in X$ and morphisms in $\mathbf{El}(T)^{\mathrm{op}}$ send distinguished elements to

distinguished elements. In other words, $\mathbf{El}(T)$ is the category of pointed sets.

2. Consider the finitary covariant powerset functor $\mathcal{P}_\omega : \mathbf{Set} \rightarrow \mathbf{Set}$. The objects of $\mathbf{El}(\mathcal{P}_\omega)^{\text{op}}$ are pairs $(X, \{x_1, \dots, x_n\})$ where $x_i \in X, 1 \leq i \leq n$ and there is a morphism between $(X, \{x_1, \dots, x_n\})$ and $(Y, \{y_1, \dots, y_m\})$ only if there exists a function $f : X \rightarrow Y$ such that $\mathcal{P}_\omega f(\{x_1, \dots, x_n\}) = f[\{x_1, \dots, x_n\}] = \{y_1, \dots, y_m\}$. This is a generalisation of the previous example in the sense that a set is equipped with a distinguished element, but this time the set is of the form $TX = \mathcal{P}_\omega X$.
3. Next, consider the forgetful functor $U : \mathbf{BA} \rightarrow \mathbf{Set}$. An object of $\mathbf{El}(U)^{\text{op}}$ is a pair (A, a) where A is a boolean algebra and $a \in A$. There will exist a morphism between (A, a) and (B, b) in $\mathbf{El}(U)^{\text{op}}$ if there exists a boolean homomorphism $f : A \rightarrow B$ such that $f(a) = b$. So $\mathbf{El}(U)^{\text{op}}$ is the category of ‘pointed boolean algebras’.
4. For a slightly different type of example, recall that a group G can be seen as a category with a single object and arrows for every element $g \in G$ (encoding ‘multiplication’ by $g \in G$). A functor $T : G \rightarrow \mathbf{Set}$ defines a set S with a group action, or G -set. In this case, $\mathbf{El}(T) \simeq S$.

As we shall now see, the category $\mathbf{El}(T)^{\text{op}}$ (and therefore $\mathbf{El}(T)$ too) is not very well behaved when it comes to the existence of limits and co-limits.

Proposition 3.16. *The category $\mathbf{El}(T)^{\text{op}}$ has an initial object iff $T = \text{hom}(A, -)$ for some A in $\mathbf{C}$.*

Proof. If $T = \text{hom}(A, -)$, then (A, Id_A) is the initial object of $\mathbf{El}(T)^{\text{op}}$. Indeed, take any (B, f) in $\mathbf{El}(T)^{\text{op}}$ (i.e. $f \in \text{hom}(A, B)$), we need to show that there exists a unique arrow $(A, \text{Id}_A) \rightarrow (B, f)$, i.e. that there must exist a unique arrow $! : A \rightarrow B$ such that $\text{hom}(A, -)(!)(\text{Id}_A) = f$, where $\text{hom}(A, -)(!) : \text{hom}(A, A) \rightarrow \text{hom}(A, B)$ is the usual application of the covariant $\text{hom}(A, -)$ functor to $!$, i.e. $\text{hom}(A, -)(!)(\text{Id}_A) = ! \circ \text{Id}_A$. Clearly, taking $! = f$ does the job.

Conversely, assume that (A, α) is an initial object in $\mathbf{El}(T)^{\text{op}}$, then we show that this provides a bijection $\phi : \text{hom}(A, B) \rightarrow TB$ for any object B in $\mathbf{C}$. Start with an element $f \in \text{hom}(A, B)$ and define $\phi(f) = Tf(\alpha) \in TB$. This assignment is injective because it defines an arrow $(A, \alpha) \rightarrow (B, Tf(\alpha))$ in $\mathbf{El}(T)^{\text{op}}$ which must be unique since (A, α) is assumed to be initial. To see that ϕ is surjective, start with an element $\beta \in TB$, then (B, β) is an object of $\mathbf{El}(T)^{\text{op}}$ and since (A, α) is initial, there must exist a unique arrow $(A, \alpha) \rightarrow (B, \beta)$, and thus an arrow $f : A \rightarrow B$ in $\mathbf{C}$ such that $Tf(\alpha) = \beta$, i.e. an element of $\text{hom}(A, B)$ such that $\phi(f) = \beta$. $\square$

Thus $\mathbf{El}(T)$ has a terminal object iff T is a hom functor. As we shall see later, having a terminal object means that $\mathbf{El}(T)$ is filtered which will be an important feature of a category in what follows. More generally we have the following result.

Proposition 3.17. *If $\mathbf{C}$ has limits of some class and T preserves them, then $\mathbf{El}(T)^{\text{op}}$ has these limits too. In particular, if $\mathbf{C}$ is complete and T is continuous, then $\mathbf{El}(T)^{\text{op}}$ is complete.*

Proof. See Proposition 4.87 in [Kel82]. $\square$

We conclude this brief description of the category $\mathbf{El}(T)$ with the following simple observation.

Lemma 3.18. *If $\mathbf{C}$ is small, then for any $T : \mathbf{C} \rightarrow \mathbf{Set}$, $\mathbf{El}(T)$ is small.*

Proof. By smallness of $\mathbf{C}$ both $\mathbf{C}$ and $\bigcup_{A \in \mathbf{C}} TA$ are sets, and thus the objects $(A, \alpha), \alpha \in TA$ of $\mathbf{El}(T)$ all lie in a set. $\square$

3.2.2 Colimits of representables

We now use what is sometimes known as the co-Yoneda lemma (particularly in the context of enriched category theory) to build a representation of functors in $\mathbf{Set}^{\mathbf{C}}$ when $\mathbf{C}$ is small. This construction is a classic construction and is detailed e.g. in Proposition 1 of [LM92] for the case of contravariant functors.

Theorem 3.19. *Let $\mathbf{C}$ be a small category and $T : \mathbf{C} \rightarrow \mathbf{Set}$ be a functor, then T is isomorphic to a colimit of hom-functors.*

Proof. We need to define a diagram $\mathcal{D}_T : \mathbf{D} \rightarrow \mathbf{Set}^{\mathbf{C}}$ such that $T \simeq \text{colim } \mathcal{D}_T$. We take as index category $\mathbf{D}$, the category of elements of T , and we use the Yoneda embedding $\mathbf{Y} : \mathbf{C}^{\text{op}} \rightarrow \mathbf{Set}^{\mathbf{C}}$ together with the forgetful functor $\mathbf{U}_T$ to define a diagram $\mathcal{D}_T$

$$\mathcal{D}_T : \mathbf{El}(T) \xrightarrow{\mathbf{U}_T} \mathbf{C}^{\text{op}} \xrightarrow{\mathbf{Y}} \mathbf{Set}^{\mathbf{C}}$$

We claim that T is its colimit, i.e.

$$T \simeq \text{colim } \mathcal{D}_T$$

To see that $\text{colim } \mathcal{D}_T$ is well-defined, note that since $\mathbf{Set}$ is co-complete, so is $\mathbf{Set}^{\mathbf{C}}$, and thus since $\mathbf{C}$ is small, $\text{colim } \mathcal{D}_T$ exists by Lemma 3.18. To check that T is the colimit, note first that by construction T is a co-cone of the diagram, the arrows defining this co-cone being precisely those forgotten by $\mathbf{U}_T$ (when $\mathbf{El}(T)$ is viewed in terms of natural transformations). Let's now check that it is a colimit, i.e. we consider any other co-cone G of $\mathcal{D}_T$ and show that there exists a

unique natural transformation $\theta : T \rightarrow G$. Consider the $\mathbf{El}(T)$ -morphism $\eta \rightarrow \eta'$ where $\eta : \text{Hom}(A, -) \rightarrow T$ and $\eta' : \text{Hom}(B, -) \rightarrow T$ are such that $\eta = \eta' \circ Y(f)$ for an $f : B \rightarrow A$ in $\mathbf{C}$.

$$\begin{array}{ccc}
 \mathcal{D}_T(\eta) = \text{Hom}(A, -) & \xrightarrow{Y(f)} & \mathcal{D}_T(\eta') = \text{Hom}(B, -) \\
 & \searrow \eta \quad \swarrow \eta' & \\
 & T & \\
 & \downarrow \theta & \\
 & G & \\
 & \swarrow \zeta \quad \searrow \zeta' &
 \end{array}$$

Following our alternative characterisation of $\mathbf{El}(T)$ and Yoneda, $\eta, \eta', \zeta, \zeta'$ correspond to elements $x(\eta) \in TA, x'(\eta') \in TB, y(\zeta) \in GA$ and $y'(\zeta') \in GB$ respectively which are related by $x(\eta) = Tf(x'(\eta'))$ and $y(\zeta) = Gf(y'(\zeta'))$. Define the transformation $\theta : T \rightarrow G$ by

$$\theta_A : TA \rightarrow GA, x(\eta) \mapsto y(\zeta)$$

If θ is well-defined, then the above diagram is commutative since $\theta \circ \eta = \zeta$ iff they correspond by Yoneda to the same element of GA . To see that this transformation is well-defined note since *every* element of TA , for an A in $\mathbf{C}$, corresponds to a natural transformation η , and since these are the objects of $\mathbf{El}(T)$ and G is a co-cone there must exist a ζ corresponding to every η . Thus θ_A is defined everywhere on TA . The fact that θ_A is a functional follows from the fact that G is a co-cone of $\mathcal{D}_T$.

Now we need to check that the transformation is natural, i.e. that for $f : B \rightarrow A$, $\theta_A \circ Tf = Gf \circ \theta_B$. Starting with $y(\eta') \in TB$, let $x(\eta) = Tf(y)$. This means that we have an arrow $\eta \rightarrow \eta'$ in $\mathbf{El}(T)$ as in the diagram above and we therefore have

$$\begin{aligned}
 \theta_A \circ Tf(y(\eta')) &= \theta_A(x(\eta)) \\
 &= y(\zeta) \\
 &= Gf(y'(\zeta')) \\
 &= Gf \circ \theta_B(\eta')
 \end{aligned}$$

□

Thus any functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ can be expressed as a colimit of representables when $\mathbf{C}$ is small. More generally, we propose the following definition.

Definition 3.20 (Functor presentation). A **presentation** for a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$, $\mathbf{C}$ small, is a diagram $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ which factors through the Yoneda embedding $Y : \mathbf{C} \rightarrow \mathbf{Set}^{\mathbf{C}}$ such that $T = \text{colim } \mathcal{D}$.

Remark 3.21. This is a good place to start describing an analogy between **Set**-functors on a small category and algebraic varieties. For a presentation $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ of a functor T , we can think of the objects of $\mathbf{J}$ as *generators* and the morphisms of $\mathbf{J}$ as *relations*. In particular, $\mathbf{El}(T)$ can be seen as a canonical category of generators and relations for T . We will make this analogy more precise further on.

3.3 Some essential tools and concepts

Before we move on to describe special types of functor presentations, we need some essential concepts from the theory of locally presentable and accessible categories, and some essential tools from general category theory.

3.3.1 Locally presentable and accessible categories

The constraint that $\mathbf{C}$ be small is too strong for many practical applications, not least the study of **Set** endofunctors. However, the construction of Theorem 3.19 seems to require *all* the objects of $\mathbf{C}$. So what are we to do? A natural solution is to consider categories which are ‘spanned’ by a small subcategory. Given a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ we can then apply Theorem 3.19 to the restriction of T to this small subcategory and find criteria under which this is enough to present the full functor T . To make this strategy precise we must first settle on a notion of ‘spanning a category’. Such a notion is provided by the theory of accessible and locally presentable categories, developed in great detail in [MP89] and [AR94] which we will follow in this subsection. We will only mention the essential definitions here and will discuss the salient facts as they become necessary in the remainder of the chapter. We refer the reader to [AR94] for further details. We assume throughout that λ is a regular cardinal and that all categories are locally small.

Definition 3.22. A non-empty partially ordered set J is λ -**directed** if every subset $J_0 \subseteq J$ such that $|J_0| \leq \lambda$ has an upper bound. If J is λ -directed, the colimit of a diagram $D : J \rightarrow \mathbf{C}$ is called a λ -**directed colimit**. An object A of a locally small category $\mathbf{C}$ is said to be λ -**presentable** if $\mathrm{hom}(A, -)$ preserves all λ -directed colimits in $\mathbf{C}$. The $\aleph_0$ -presentable objects are called **finitely presentable** and an object will be simply called **presentable** if it is λ -presentable for some λ .

An equivalent characterisation of λ -presentable objects is as follows. Assume $\mathcal{D} : J \rightarrow \mathbf{C}$ is a λ -directed diagram and $B = \mathrm{colim} \mathcal{D} = \mathrm{colim}_{j \in J} B_j$ is its λ -directed colimit. If there exists a morphism $f : A \rightarrow B$, then A is λ -presentable iff f factorises essentially uniquely through one of the morphisms $c_j : B_j \rightarrow B$

of the colimiting co-cone, where ‘essentially uniquely’ means that if f factors through c_j in two different ways, i.e. $f = c_j \circ g = c_j \circ g'$, then we can find an $i \geq j$ where these two different ways are unified, i.e. such that $\mathcal{D}(j \rightarrow i) \circ g = \mathcal{D}(j \rightarrow i) \circ g'$. The λ -presentable objects will form the basis or generators of the notion of ‘spanning a category’ which we will use. We now need to specify how to assemble these basic objects.

Definition 3.23. A locally small category $\mathbf{C}$ will be called λ -**accessible** if it has λ -directed colimits and a small subcategory $\mathbf{C}_\lambda$ of λ -presentable objects such that every object in $\mathbf{C}$ is a λ -directed colimit of objects in $\mathbf{C}_\lambda$. If in addition $\mathbf{C}$ is cocomplete (i.e. $\mathbf{C}$ has *all* colimits, not just λ -directed ones), then $\mathbf{C}$ is said to be **locally λ -presentable**. Again, for $\lambda = \aleph_0$ we will say finitely accessible and locally finitely presentable, whilst accessible and locally presentable will mean λ -accessible and λ -locally presentable for some λ .

- Example 3.24.**
1. The standard example is that of the category **Set** which is locally finitely presentable. A set X is finitely presentable iff it is finite, and since any set is the directed colimit (i.e. union in this instance) of its finite subsets, it is finitely accessible. Since **Set** is also cocomplete, it is locally finitely presentable.
 2. The category **BA** is also locally finitely presentable. A boolean algebra A is finitely presentable iff it can be presented by finitely many generators and finitely many relations. This is in fact a general property of all varieties with a finitary signature. For details of why this is the case and why **BA** and all finitary varieties are locally finitely presentable, we refer the reader to chapter 3 of [AR94].
 3. As we will prove later, the category of functors $\mathbf{Set}^{\mathbf{C}}$ is also locally finitely presentable when $\mathbf{C}$ is small.

In light of these examples, let us clarify the status of $\mathbf{C}_\lambda$. There are two ways of looking at the smallness of $\mathbf{C}_\lambda$. The first is to consider $\mathbf{C}_\lambda$ as the entire subcategory of λ -presentable objects. The correct notion of smallness is then to say that $\mathbf{C}_\lambda$ is *essentially small*, i.e. that $\mathbf{C}_\lambda$ has a small skeleton, i.e. that it is locally small and has a small number of isomorphism classes of objects. This is the approach taken in [MP89]. The other approach, taken in [AR94], is to assume that a choice of representative of each isomorphism class of λ -presentable object has been made, such as taking the finite ordinals in the case of **Set**, in which case $\mathbf{C}_\lambda$ is small in the usual sense. This is the approach we will choose throughout.

It is often convenient to work with a class of colimits which is equivalent to directed colimits. A category $\mathbf{J}$ is called λ -**filtered** if it satisfies the following conditions:

1. it is non-empty
2. if I is a set of cardinality strictly smaller than λ and $\{A_i\}_{i \in I}$ are objects in $\mathbf{J}$, then there exists an object B and $\mathbf{J}$ -morphisms $f_i : A_i \rightarrow B$
3. if I is a set of cardinality strictly smaller than λ and $\{f_i : A \rightarrow B\}_{i \in I}$ are $\mathbf{J}$ -morphisms, there exists a $\mathbf{J}$ -morphism $h : B \rightarrow C$ which simultaneously equalizes all $f_i, i \in I$

We will say ‘filtered’ for $\aleph_0$ -filtered. A λ -filtered diagram is a diagram $D : \mathbf{J} \rightarrow \mathbf{C}$ where $\mathbf{J}$ is λ -filtered. A λ -**filtered colimit** is the colimit of such a diagram. We cite the following two useful results without proof.

Lemma 3.25. *A category has λ -filtered colimits iff it has λ -directed colimits. A functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ from such a category preserves λ -filtered colimits iff it preserves λ -directed colimits*

Proof. See [AR94] 1.7. □

Lemma 3.26. *A colimit of a λ -small diagram of λ -presentable objects is λ -presentable.*

Proof. See [AR94] 1.16. □

This result generalises as follows.

Lemma 3.27. *Let λ, μ be two regular cardinals, then λ is said to be sharply smaller than μ , written $\lambda \triangleleft \mu$, if $\lambda < \mu$ and for any set X such that $|X| < \mu$, $\mathcal{P}_\lambda(X)$ has a cofinal set of cardinality strictly smaller than μ . Assume now that $\mathbf{C}$ is a λ -accessible category and $\lambda \triangleleft \mu$, then an object A in $\mathbf{C}$ is μ -presentable iff it is a split subobject of a μ -small λ -directed colimit of λ -presentable objects.*

Proof. See [MP89] 2.3.11. □

The relation $\triangleleft$ turns out to be very important in the study of accessible categories, for example if a category is λ -accessible, and $\lambda \triangleleft \mu$, then it is also μ -accessible. As a special example of such a $\lambda \triangleleft \mu$ relation, note that $\lambda \triangleleft \lambda^+$, where λ^+ is the successor cardinal of λ . It is also worth mentioning that $\triangleleft$ is transitive.

To conclude this brief introduction to locally presentable and accessible categories we need to define the following crucial class of functors.

Definition 3.28. Let $\mathbf{C}$ be a category with λ -directed colimits and let $\mathbf{D}$ be a λ -accessible category. We will say that a functor $T : \mathbf{C} \rightarrow \mathbf{D}$ is λ -**accessible** if it preserves λ -directed colimits.

Given an accessible category $\mathbf{C}$, this property on a functor $\mathbf{C} \rightarrow \mathbf{Set}$ will be the key to lifting Theorem 3.19 from the small subcategory of presentable objects of $\mathbf{C}$ to the full category $\mathbf{C}$.

3.3.2 Kan extensions and cofinal diagrams

We now present two extremely useful technical concepts on which we will rely in what follows. The notion of Kan extension is very useful in the context of extending a functor defined on a subcategory to the entire category. Here we will merely recall the definition and a few easy but salient facts, for more details the reader is referred to Chapter X of [Mac98]. Since we are dealing with colimits in this Chapter, the useful notion is that of *left* Kan extension (in Section 5.2.4 we will deal with limits and use right Kan extensions).

Definition 3.29. Given two functors $\mathbf{B} \xleftarrow{F} \mathbf{A} \xrightarrow{G} \mathbf{C}$, a left Kan extension $\text{Lan}_G F$ of F along G is a functor $\text{Lan}_G F : \mathbf{C} \rightarrow \mathbf{B}$ together with a natural transformation $\eta : F \rightarrow \text{Lan}_G F \circ G$. The pair $(\text{Lan}_G F, \eta)$ is universal in the sense that for any other pair (L', η') with $L' : \mathbf{C} \rightarrow \mathbf{B}$ and $\eta' : F \rightarrow L' \circ G$, there must exist a natural transformation $\zeta : \text{Lan}_G F \rightarrow L'$ such that $\eta' = \zeta_G \circ \eta$ where ζ_G is the obvious natural transformation $\zeta_G : \text{Lan}_G F \circ G \rightarrow L' \circ G$.

In light of the definition of accessible and presentable categories, let

$$I_\lambda : \mathbf{C}_\lambda \rightarrow \mathbf{C}$$

be the inclusion functor of the subcategory of λ -presentable objects of $\mathbf{C}$. Let also $T : \mathbf{C} \rightarrow \mathbf{Set}$ be an arbitrary functor and $T_\lambda = T I_\lambda : \mathbf{C}_\lambda \rightarrow \mathbf{C}$ be its restriction to the subcategory $\mathbf{C}_\lambda$. We will examine the left Kan extension $\text{Lan}_{I_\lambda} T_\lambda$. First recall,

Proposition 3.30. *Let $\mathbf{B} \xleftarrow{F} \mathbf{A} \xrightarrow{G} \mathbf{C}$ such that $\mathbf{A}$ is small and $\mathbf{B}$ is cocomplete, then F has a left Kan extension $\text{Lan}_G F$ along any G .*

Proof. See [Mac98] X.3.2. □

Thus $\text{Lan}_{I_\lambda} T_\lambda$ always exists. Its explicit construction is given by

$$\text{Lan}_{I_\lambda} T_\lambda A = \text{colim} \left((\mathbf{C}_\lambda \downarrow A) \xrightarrow{\mathcal{D}_c^A} \mathbf{C}_\lambda \xrightarrow{T_\lambda} \mathbf{Set} \right)$$

where $\mathcal{D}_c^A$ is the functor which projects an object of $\mathbf{C}_\lambda \downarrow A$ to its domain (and is called the canonical diagram, see below). Moreover, we have the following:


Proposition 3.31. *Let $\mathbf{B} \xleftarrow{F_0} \mathbf{A}_0 \xrightarrow{I} \mathbf{A}$ such that $\mathbf{A}_0$ is a small full subcategory of $\mathbf{A}$ and $\mathbf{B}$ is cocomplete, then F_0 has a left Kan extension $\text{Lan}_I F_0$ whose associated natural transformation $\eta : F_0 \rightarrow \text{Lan}_I F_0 \circ I$ is the identity natural transformation.*

Proof. See [Mac98] X.3.3. □

The fact that $\text{Lan}_{I_\lambda} T_\lambda$ can be written as a colimit suggests that if this colimit is filtered and if the functor F is accessible we should be able to recover T from T_λ . This is indeed the case, but to show it we will need the following concept, which will prove very useful later on in this chapter.

Definition 3.32. A functor $G : \mathbf{A} \rightarrow \mathbf{B}$ is called **cofinal** (or sometimes ‘final’, see e.g. [Mac98]) if for any object B of $\mathbf{B}$ the comma category $B \downarrow G$ is non-empty and connected, i.e. if the following two conditions hold for any B

- (i) there exists an object A of $\mathbf{A}$ and a morphism $B \rightarrow GA$
- (ii) any two objects in $B \downarrow G$, i.e. any two morphisms $B \rightarrow GA$ and $B \rightarrow GA'$, are related by a zigzag of morphisms:



(or the same pattern with the arrows between G -terms reversed).

Cofinal functors are very useful for computing colimits. Indeed, if a functor $G : \mathbf{A} \rightarrow \mathbf{B}$ is cofinal, then we can restrict diagrams $\mathcal{D}$ on $\mathbf{B}$ to diagrams on $\mathbf{A}$ along G without altering their colimits. In this context we will speak of a **cofinal diagram** rather than cofinal functor.

Theorem 3.33. If $G : \mathbf{A} \rightarrow \mathbf{B}$ is cofinal and $\mathcal{D} : \mathbf{B} \rightarrow \mathbf{C}$ is a diagram such that $\text{colim } \mathcal{D}G$ exists, then $\text{colim } \mathcal{D} \simeq \text{colim } \mathcal{D}G$.

Proof. See [Mac98], Chapter IX, Section 3, Theorem 1. □


We will be particularly interested in the case where G is an inclusion functor and $\mathbf{A}$ and $\mathbf{B}$ are directed or filtered categories. Note that if G is full and $\mathbf{A}$ is filtered, then the existence of a zigzag of morphisms between any two objects $B \rightarrow GA$ and $B \rightarrow GA'$ of $B \downarrow G$ implies that the span $GA \leftarrow B \rightarrow GA'$ can be completed into a commutative square

$$\begin{array}{ccc} B & \longrightarrow & GA \\ \downarrow & & \downarrow \\ GA' & \longrightarrow & GA'' \end{array}$$

which, together with the non-emptiness of $B \downarrow G$, is the - slightly non-standard but perhaps more *cofinal* - definition of cofinality given in 0.11 of [AR94].

Lemma 3.34. *Let $\mathbf{C}$ be a λ -accessible category whose subcategory of λ -presentable objects is $\mathbf{C}_\lambda$, and let A be an object of $\mathbf{C}$. Define the **canonical diagram** to be the projection functor $\mathcal{D}_c^A : \mathbf{C}_\lambda \downarrow A \rightarrow \mathbf{C}$ defined above, then*

- Proof.* We start by showing (i). By definition of a λ -accessible category, A must be expressible as a λ -directed colimit $A = \operatorname{colim}_{j \in J} A_j$ of λ -presentable objects. We will use this colimit to show that $\mathbf{C}_\lambda \downarrow A$ is λ -filtered. Note first that $\mathbf{C}_\lambda \downarrow A$ is clearly non-empty, or else J would have to be empty which would contradict the fact that it is λ -filtered. Note also that if we take any object of $\mathbf{C}_\lambda \downarrow A$, i.e. any arrow $B \rightarrow A$ with B λ -presentable, then by definition of λ -presentability $B \rightarrow A = \operatorname{colim}_{j \in J} A_j$ must factor through an essentially unique $c_j : A_j \rightarrow A$. If we select strictly fewer than λ such objects $f_\kappa : B_\kappa \rightarrow A, \kappa < \lambda$, then each f_κ must factor through some map $c_{j_\kappa} : A_{j_\kappa} \rightarrow A$ and since J is λ -directed, there must exist $c_j : A_j \rightarrow A$ such that $c_j \circ (A_{j_\kappa} \rightarrow A_j) = c_{j_\kappa}$ for each κ and thus such that $f_\kappa = c_j \circ (B_\kappa \rightarrow A_{j_\kappa} \rightarrow A_j)$ for each κ , i.e. we have arrows in $\mathbf{C}_\lambda \downarrow A$ from every $B_\kappa \rightarrow A$ to a common codomain. The situation is illustrated by the following commutative diagram.



$$A = \operatorname{colim} \mathcal{D}^A = \operatorname{colim} \mathcal{D}_c^A \circ I = \operatorname{colim} \mathcal{D}_c^A$$
☐

Corollary 3.35. *Let $\mathbf{C}$ be a non-small λ -accessible category, then the set of objects in $\mathbf{C}_\lambda$ has cardinality at least λ .*

Proof. Assume, for the sake of contradiction that $|(\mathbf{C}_\lambda)^{\text{ob}}| < \lambda$. Take any object A in $\mathbf{C}$. From Lemma 3.34 we know that $A = \text{colim } \mathcal{D}_c^A$. But since $|(\mathbf{C}_\lambda)^{\text{ob}}| < \lambda$ and a λ -small colimit of λ -presentable objects is λ -presentable by Lemma 3.26, this means that A is λ -presentable, and thus $\mathbf{C} = \mathbf{C}_\lambda$ which is a contradiction since we're assuming that $\mathbf{C}$ is not small. $\square$

We can now state the main result of this subsection.

Proposition 3.36. *Let $\mathbf{C}$ be a λ -accessible category whose subcategory of λ -presentable objects is $\mathbf{C}_\lambda$ and let $l_\lambda : \mathbf{C}_\lambda \rightarrow \mathbf{C}$ be the injection functor. For any λ -accessible functor $T : \mathbf{C} \rightarrow \mathbf{Set}$, with restriction $T_\lambda : \mathbf{C}_\lambda \rightarrow \mathbf{Set}$, it is the case that*

$$\text{Lan}_{l_\lambda} T_\lambda = T$$

Proof. So far we have used the notation $\mathcal{D}_c^A$ to denote the projection functors with codomain $\mathbf{C}_\lambda$ or $\mathbf{C}$ depending on the context. Here we will need to draw the distinction, and the two variants are introduced in the following trivially commutative diagram:

$$\begin{array}{ccc} & \mathbf{C}_\lambda \downarrow A & \\ \mathcal{D}_{c_\lambda}^A \swarrow & & \searrow \mathcal{D}_c^A \\ \mathbf{C}_\lambda & \xrightarrow{l_\lambda} & \mathbf{C} \\ T_\lambda \searrow & & \swarrow T \\ & \mathbf{Set} & \end{array}$$

By construction we have

$$\begin{aligned} \text{Lan}_{l_\lambda} T_\lambda A &= \text{colim} \left((\mathbf{C}_\lambda \downarrow A) \xrightarrow{\mathcal{D}_{c_\lambda}^A} \mathbf{C}_\lambda \xrightarrow{T_\lambda} \mathbf{Set} \right) \\ &\stackrel{1}{=} \text{colim} \left((\mathbf{C}_0 \downarrow A) \xrightarrow{\mathcal{D}_c^A} \mathbf{C} \xrightarrow{T} \mathbf{Set} \right) \\ &\stackrel{2}{=} T \text{ colim } \mathcal{D}_c^A \\ &\stackrel{3}{=} TA \end{aligned}$$

where step 1 is by definition of T_λ and the commutativity of the diagram above, step 2 is by the fact that $\mathcal{D}_c^A$ is a λ -filtered diagram (by Lemma 3.34) and that T is λ -accessible, and step 3 is by the fact that $\text{colim } \mathcal{D}_c^A = A$ by Lemma 3.34. $\square$

3.4 Size-bound functor presentations

We now return to the functors $T : \mathbf{C} \rightarrow \mathbf{Set}$ where $\mathbf{C}$ is small. As was shown in Theorem 3.19, any such functor has a presentation arising from its category

of elements $\mathbf{El}(T)$. In this section we will consider cases where simpler presentations of T are possible, and relate the existence of such presentations to the status of T in $\mathbf{Set}^{\mathbf{C}}$.

3.4.1 λ -presentable and λ -presented functors

We start with the following obvious definition.

Definition 3.37 (Presentable functor). We will say that a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ ($\mathbf{C}$ arbitrary) is **λ -presentable** if it is a λ -presentable object in $\mathbf{Set}^{\mathbf{C}}$, i.e. if $\text{Nat}(T, -)$ preserves λ -directed colimits. As usual if $\lambda = \aleph_0$ we will say *finitely* presentable.

We have the following easy example of presentable functors.

Proposition 3.38. *If $\mathbf{C}$ is a small category, then all hom functors $\text{hom}(A, -) : \mathbf{C} \rightarrow \mathbf{Set}$ are λ -presentable.*

Proof. In fact we will show a stronger result, namely that hom functors are *absolutely presentable* (see [KRV05]) in $\mathbf{Set}^{\mathbf{C}}$, i.e. that for any object A of $\mathbf{C}$, $\text{Nat}(\text{hom}(A, -), -)$ preserves *all* colimits, and a fortiori the λ -directed ones. Moreover, the smallness of $\mathbf{C}$ is not necessary, and the same proof works for $\mathbf{C}$ locally small. This latter requirement cannot be relaxed as the Yoneda lemma is key.

Let A be an object of $\mathbf{C}$, and let $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ be a diagram whose colimit is G . We need to show that if there exists a natural transformation $\eta : \text{hom}(A, -) \rightarrow G$, then it must factor through an essentially unique object $\mathcal{D}(j)$ for a certain object j of $\mathbf{J}$. For notational clarity, let us write $G_j = \mathcal{D}(j)$, i.e. $G = \text{colim}_j G_j$. By Yoneda's lemma any morphism $\eta : \text{hom}(A, -) \rightarrow G$ corresponds to a unique element $\alpha \in GA$ where $\alpha = \eta_A(\text{Id}_A)$. Since colimits in $\mathbf{Set}^{\mathbf{C}}$ are computed pointwise, GA is the colimit of the diagram $\{G_j A\}_{j \in \mathbf{J}}$ in $\mathbf{Set}$ and thus, by the general construction of colimits in $\mathbf{Set}$, $\alpha \in GA$ must originate from an essentially unique $\beta \in G_j A$. By using Yoneda in the opposite direction, it is clear that this β corresponds to a natural transformation $\theta : \text{hom}(A, -) \rightarrow G_j$ defined by $\theta_A(\text{Id}_A) = \beta$, and for any $f : A \rightarrow B$ by $\theta_B(f) = Tf(\beta)$. If we denote by c_j the natural transformation $c_j : G_j \rightarrow G$, we can then check that

$\eta = c_j \circ \theta$. For every B in $\mathbf{C}$ we have:

$$\begin{aligned}
 \eta_B(f) &= \eta_B \circ \text{hom}(A, -)(f)(\text{Id}_A) \\
 &= Gf \circ \eta_A(\text{Id}_A) && \eta \text{ natural} \\
 &= Gf(\alpha) && \text{Definition of } \alpha \\
 &= Gf((c_j)_A(\beta)) && \text{Definition of } \beta \\
 &= Gf \circ (c_j)_A \circ \theta_A(\text{Id}_A) && \text{Definition of } \theta \\
 &= (c_j)_B \circ \theta_B \circ \text{hom}(A, -)(f)(\text{Id}_A) && c_j, \theta \text{ natural} \\
 &= (c_j)_B \circ \theta_B(f)
 \end{aligned}$$

i.e. η factors through θ as desired. The essential unicity of the factorisation is a consequence of the essential unicity of β and of the corresponding θ by Yoneda's lemma. $\square$

We can in fact strengthen the previous Proposition 3.38 as follows:

Proposition 3.39. *Let $\mathbf{C}$ be a non-empty small finitely complete category and $T : \mathbf{C} \rightarrow \mathbf{Set}$. T is left-exact (i.e. preserves finite limits) and finitely presentable in $\mathbf{Set}^{\mathbf{C}}$ iff it is a hom functor.*

Proof. The ‘if’ direction follows from the previous Proposition 3.38 and the fact that (covariant) hom functors preserve limits. For the ‘only if’ direction we will first show that $\mathbf{El}(T)$ is filtered. Since $\mathbf{C}$ is non empty, $\mathbf{El}(T)$ must be non-empty too. Moreover, since $\mathbf{C}$ has finite limits it must have products and equalizers (in fact these two requirements are equivalent). Thus for any $(A, \alpha), (B, \beta)$ in $\mathbf{El}(T)$ we can build the object $(A \times B, (\alpha, \beta))$ in $\mathbf{El}(T)$ and since T is left-exact the projections give us arrows $T\pi_1 : T(A \times B) = TA \times TB \rightarrow TA$ and $T\pi_2 : T(A \times B) = TA \times TB \rightarrow TB$ such that $T\pi_1(\alpha, \beta) = \alpha$ and $T\pi_2(\alpha, \beta) = \beta$, i.e. arrows in $(A, \alpha) \rightarrow (A \times B, (\alpha, \beta))$ and $(B, \beta) \rightarrow (A \times B, (\alpha, \beta))$ in $\mathbf{El}(T)$. Similarly, given two arrows $(A, \alpha) \rightrightarrows (B, \beta)$, i.e. arrows $f, g : B \rightarrow A$ such that $Tf(\beta) = Tg(\beta) = \alpha$, we can build the equalizer $e : E \rightarrow B$ of f, g in $\mathbf{C}$. Since T preserves equalizers, we know that Te equalizes Tf and Tg and thus since $Tf(\beta) = Tg(\beta)$, there must exist an element γ in the equalizer TE such that $Te(\gamma) = \beta$. We therefore have an $\mathbf{El}(T)$ -arrow $(B, \beta) \rightarrow (E, \gamma)$ which composes with the two arrows $(A, \alpha) \rightrightarrows (B, \beta)$ to give a unique arrow $(A, \alpha) \rightarrow (E, \gamma)$.

Since $\mathbf{El}(T)$ is filtered, it is obvious that if we write T as a colimit of representables following the construction of Theorem 3.19, T can in fact be written as a *filtered* colimit of representables. Since it is assumed to be finitely presentable in $\mathbf{Set}^{\mathbf{C}}$ it means that the identity natural transformation $\text{Id}_T : T \rightarrow T$ must factor uniquely through one of the hom functors of the filtered colimit representing T . This in turn, means that T must be equivalent to this hom functor. $\square$

In fact, we can demand slightly less than left-exactness. If we instead require that T *weakly* preserve limits, then the proof goes through unchanged. Note also that weak preservation of limits is equivalent to weak preservation of pullbacks and terminal objects. The proof can be generalised to higher cardinalities, in which case left exactness must be generalised to λ -left exactness, i.e. the preservation of λ -small limits. We can also characterize left-exact functors.

Proposition 3.40. *Let $\mathbf{C}$ be a small category, then $T : \mathbf{C} \rightarrow \mathbf{Set}$ is left exact iff $\mathbf{El}(T)$ is filtered.*

Proof. Assume first that T is left exact, then as we saw in the proof of the previous Proposition 3.39, $\mathbf{El}(T)$ is filtered. Let us now assume that $\mathbf{El}(T)$ is filtered, and let $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{C}$ be a finite diagram, we then have

$$\begin{aligned}
 T(\lim \mathcal{D}) &= \operatorname{colim}(\mathbf{Y}\mathbf{U}_T)(\lim \mathcal{D}) \\
 &\stackrel{1}{=} \operatorname{colim}(ev_{\lim \mathcal{D}} \mathbf{Y}\mathbf{U}_T) \\
 &\stackrel{2}{=} \operatorname{colim}(\lim((ev_{(-)} \mathcal{D}) \mathbf{Y}\mathbf{U}_T)) \\
 &\stackrel{3}{=} \lim(\operatorname{colim}((ev_{(-)} \mathcal{D}) \mathbf{Y}\mathbf{U}_T)) \\
 &\stackrel{4}{=} \lim(\operatorname{colim}(\mathbf{Y}\mathbf{U}_T) \mathcal{D}) \\
 &= \lim(T\mathcal{D})
 \end{aligned}$$

where (1) is because colimits are evaluated pointwise in $\mathbf{Set}^{\mathbf{C}}$, (2) follows from the fact that both $\mathbf{Y}$ and $\mathbf{U}_T$ preserves limits, (3) is the well-known fact that finite limits commute with filtered colimits (see [Mac98] IX.2.), and (4) is because each evaluation functor $ev_A : \mathbf{Set}^{\mathbf{C}} \rightarrow \mathbf{Set}$ preserves colimits (see [Mac98] V.3.). $\square$

One of the motivations behind the concept of finitely presentable object in a category, is that in the case of algebraic varieties, it provides an abstract characterisation of the objects which can be given a finite presentation by generators and relations, i.e. finitely presented objects. And in fact, as shown in Theorem 3.12. of [AR94], the two concepts are equivalent. We would like to show the same sort of correspondence for functors $\mathbf{Set} \rightarrow \mathbf{C}$ when $\mathbf{C}$ is small. The purpose of this section is to isolate classes of functors for which the cardinality of the set of relations and/or generators is bounded. There are two natural ways to specify the cardinality of generators and relations, and two associated notions of finitely-presented and finitely-generated functors. Recall from Remark 3.21 that for a presentation $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ of a functor T , we drew a parallel between the objects of $\mathbf{J}$ and generators and between the morphisms of $\mathbf{J}$ and relations.

Definition 3.41 (Strongly λ -presented functor). A functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ is **strongly λ -presented** if there exists a λ -small presentation $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ of T , i.e. a presentation such that $\mathbf{J}$ has fewer than λ objects and fewer than λ

morphisms. As usual, we will say that T is **finitely strongly presented** if it is strongly $\aleph_0$ -presentable, i.e. if it can be written as a finite colimit of hom functors.

The intuition behind the alternative notion of finite presentation is the following. We have seen that we can write any functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ as a colimit of representable functors. By the general construction of colimits as regular quotients of coproducts, this means that any functor T can be viewed as a regular quotient of a coproduct of hom functors. From this perspective it is natural to consider the hom functors as ‘generators’, and the following definition is then natural.

Definition 3.42 (λ -presented functor). A functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ is **λ -presented** if there exist a set I of cardinality $|I| < \lambda$, and a regular epimorphism (coequalizer) e

$$S \rightrightarrows \coprod_{i \in I} \text{hom}(A_i, -) \xrightarrow{e} T$$

such that S can itself be expressed as a regular quotient of a coproduct of fewer than λ representable functors.

In this definition, the only measure of ‘size’ is the cardinality of the sets indexing coproducts of homs. The number of relations necessary to encode T is therefore measured by the size of the functor encoding the relations, i.e. by the ‘size’ of the domain of the pair of arrows coequalized by e . As suggested by the terminology, we have

Proposition 3.43. *If a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ is strongly λ -presented, then it is λ -presented.*

Proof. If T is strongly finitely presented, then $T = \text{colim } Y \circ \mathcal{D}$ with $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{C}^{\text{op}}$ such that $\mathbf{J}$ has strictly fewer than λ objects and fewer than λ morphisms. The result follows immediately by expressing the colimit as a regular quotient of a coproduct in the canonical way, i.e.

$$\coprod_{f \in \mathbf{J}^{\text{mor}}} \text{hom}(\text{cod } \mathcal{D}(f), -) \xrightarrow[\psi]{\phi} \coprod_{A \in \mathbf{J}^{\text{ob}}} \text{hom}(\mathcal{D}(A), -) \xrightarrow{e} T$$

where $\phi \circ \text{in}_f = \text{in}_{\text{cod } f}$, $\psi \circ \text{in}_f = \text{in}_{\text{dom } f} \circ ((-) \circ \mathcal{D}(f))$ and $\lambda < |\mathbf{J}^{\text{mor}}|, |\mathbf{J}^{\text{ob}}|$. We can then take $S = \coprod_{f \in \mathbf{J}^{\text{mor}}} \text{hom}(\text{cod } \mathcal{D}(f), -)$ and the identity as a trivial regular quotient in the definition of λ -presented. $\square$

In fact, we can also show the converse.

Proposition 3.44. *If a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ is λ -presented, then it is also strongly λ -presented.*

Proof. By unravelling the definition of being λ -presented, it is immediate that there exist sets I, J of cardinality strictly less than λ such that

$$\coprod_{j \in J} \text{hom}(B_j, -) \xrightarrow{q} S \xrightleftharpoons[\phi_2]{\phi_1} \coprod_{i \in I} \text{hom}(A_i, -) \xrightarrow{e} T$$

It is easy to check that since q is epi, e is also the coequalizer of $\phi_1 \circ q, \phi_2 \circ q$. So we can assume w.l.o.g. that we in fact have a coequalizer

$$\coprod_{j \in J} \text{hom}(B_j, -) \xrightleftharpoons[\phi_2]{\phi_1} \coprod_{i \in I} \text{hom}(A_i, -) \xrightarrow{e} T$$

By the Yoneda lemma, the transformations ϕ_1, ϕ_2 are equivalent to the data for each $j \in J$ of pair maps $f_j^1 : A_i \rightarrow B_j, f_j^2 : A_k \rightarrow B_j$ for some indices i, k . We now define the category $\mathbf{J}$ as the subcategory of $\mathbf{C}$ such that

$$\mathbf{J}^{\text{ob}} = \{A_i \mid i \in I\} \cup \{B_j \mid j \in J\}$$

and

$$\mathbf{J}^{\text{mor}} = \text{closure of } \{f_j^1, f_j^2 \mid j \in J\} \text{ under composition and identities}$$

We claim that if we write $\mathbf{Y} : \mathbf{J}^{\text{op}} \rightarrow \mathbf{Set}^{\mathbf{C}}$ for the Yoneda transformation, then $T = \text{colim } \mathbf{Y}$. Since λ is regular, we have that $|\mathbf{J}^{\text{ob}}|, |\mathbf{J}^{\text{mor}}| < \lambda$, and if $T = \text{colim } \mathbf{Y}$ then it is strongly λ -presentable, as desired.

Let $T' = \text{colim } \mathbf{Y}$, we now work pointwise to establish $T' = T$. Let X be any object of $\mathbf{C}$, by assumption and the description of ϕ_1, ϕ_2 from the Yoneda lemma, TX is the quotient of $\coprod_{i \in I} \text{hom}(A_i, X)$ under the smallest equivalence relation generated by

$$h_i : A_i \rightarrow X \sim_X h_k : A_k \rightarrow X$$

if there exist $j \in J$ and $g : B_j \rightarrow X$ such that

$$\begin{array}{ccc} A_i & \xrightarrow{h_i} & X \\ f_j^1 \searrow & & \nearrow g \\ & B_j & \\ f_j^2 \nearrow & & \nwarrow h_k \\ A_k & \xrightarrow{h_k} & X \end{array} \quad (3.2)$$

commutes. If we build $T' = \text{colim } \mathbf{Y}$ as a regular quotient of a coproduct we get:

$$\coprod_{f \in \mathbf{J}^{\text{mor}}} \text{hom}(\text{cod } f, X) \xrightleftharpoons[\psi]{\phi} \coprod_{C \in \mathbf{J}^{\text{ob}}} \text{hom}(C, X) \longrightarrow T'X$$

where $\phi \circ \text{in}_f = \text{in}_{\text{cod}f}, \psi \circ \text{in}_f = \text{in}_{\text{dom}f} \circ ((-) \circ f)$. $T'X$ is the quotient of $\coprod_{C \in \mathbf{J}^{\text{ob}}} \text{hom}(C, X)$ under the smallest equivalence relation generated by

$$h_j : B_j \rightarrow X \approx_X h_k : A_k \rightarrow X$$

if there exists $f_j : A_k \rightarrow B_j \in \mathbf{J}^{\text{mor}}$ such that

$$\begin{array}{ccc} B_j & \xrightarrow{h_j} & X \\ \uparrow f_j & \nearrow h_k & \\ A_k & & \end{array} \quad (3.3)$$

commutes. We can now see that $T'X \simeq TX$ by defining a bijection between the two sets. We simply put

$$\theta_X : TX \rightarrow T'X, [A_i \xrightarrow{h_i} X]_{\sim_X} \mapsto [A_i \xrightarrow{h_i} X]_{\approx_X}$$

where the square brackets simply denote the equivalence classes under the respective equivalence relations. To see that the map θ_X is well-defined notice that from diagram (3.2) we get two copies of diagram (3.3), i.e.

$$\begin{aligned} & \text{if } A \xrightarrow{h_k} X \sim_X A \xrightarrow{h_i} X \text{ then} \\ & A_i \xrightarrow{h_i} X \approx_X B_j \xrightarrow{g} X \text{ and } B_j \xrightarrow{g} X \approx_X A_i \xrightarrow{h_i} X \end{aligned}$$

From this it is clear that if we pick two representatives $A_i \xrightarrow{h_i} X, A_j \xrightarrow{h_j} X$ of the same equivalence class under $\sim_X$, then they will belong to the same equivalence class under $\approx_X$. This also shows that θ_X is injective: if $\theta_X([A_i \xrightarrow{h_i} X]_{\sim_X}) = \theta_X([A_k \xrightarrow{h_k} X]_{\sim_X})$, i.e. $[A_i \xrightarrow{h_i} X]_{\approx_X} = [A_k \xrightarrow{h_k} X]_{\approx_X}$, then we must have $[A_i \xrightarrow{h_i} X]_{\sim_X} = [A_k \xrightarrow{h_k} X]_{\sim_X}$. Finally to see that θ_X is surjective, we only need to check that an equivalence class with representative $[B_j \xrightarrow{g} X]_{\approx_X}$ has a pre-image under θ_X . This follows from the fact that we have morphisms of the type $f_j : A_i \rightarrow B_j$ by construction of $\mathbf{J}$, and thus $[B_j \xrightarrow{g} X]_{\approx_X} = [A_i \xrightarrow{gf_j} X]_{\approx_X}$ which has pre-image $[A_i \xrightarrow{gf_j} X]_{\sim_X}$ under θ_X . $\square$

The reader would be justified in asking why we have defined two notions of λ -presentability if they are equivalent. The reason will become clear when we consider the associated notions of λ -generated and strongly λ -generated functors which are, in general, not equivalent. First though, let us prove the promised results that just as in the case of algebraic varieties, being finitely (strongly) presentable is equivalent to being finitely presented. We start with the following technical lemma.

Lemma 3.45. *Let $\mathbf{C}$ be a category and let $\mathbf{l} : \mathbf{C}_0 \hookrightarrow \mathbf{C}$ be the embedding of a full subcategory. Assume that $T : \mathbf{C} \rightarrow \mathbf{Set}$ is a functor and that T_0 is its restriction to $\mathbf{C}_0$, assume further that T and T_0 have the property that $T = \text{Lan}_{\mathbf{l}} T_0$, then if T is a λ -presentable object in $\mathbf{Set}^{\mathbf{C}}$ so is T_0 in $\mathbf{Set}^{\mathbf{C}_0}$.*

Proof. Assume that $\mathcal{D}_0 : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}_0}$ is a λ -filtered diagram and that there exists a natural transformation $\sigma_0 : T_0 \rightarrow \text{colim } \mathcal{D}_0$. We need to show that σ_0 factors through an essentially unique object in the diagram $\mathcal{D}_0$. To show this recall first that we can turn $\mathcal{D}_0$ into a λ -filtered diagram in $\mathbf{Set}^{\mathbf{C}}$ by post-composing with the functor $\text{Lan}_I(-) : \mathbf{Set}^{\mathbf{C}_0} \rightarrow \mathbf{Set}^{\mathbf{C}}$, i.e. we define $\mathcal{D} = \text{Lan}_I(-) \circ \mathcal{D}_0$. Since the left Kan extension construction is a left adjoint, i.e.

$$\text{Lan}_I(-) \dashv (-) \circ I$$

it preserves all colimits, and thus

$$\text{colim } \mathcal{D} = \text{colim } \text{Lan}_I(-) \circ \mathcal{D}_0 = \text{Lan}_I(\text{colim } \mathcal{D}_0)$$

Since $T = \text{Lan}_I T_0$, by functoriality of $\text{Lan}_I(-)$ there must exist a natural transformation $\sigma = \text{Lan}_I(\sigma_0) : T \rightarrow \text{Lan}_I(\text{colim } \mathcal{D}_0)$, i.e. a natural transformation $\sigma : T \rightarrow \text{colim } \mathcal{D}$ where $\mathcal{D}$ is λ -filtered. Since T is λ -presentable σ must therefore factor essentially uniquely through an object $\text{Lan}_I(\mathcal{D}_0(j))$ for some j of $\mathbf{J}$. Finally, since we are taking left Kan extensions along a full inclusion functor, we are under the conditions of Proposition 3.31, i.e. $\text{Lan}_I(F) \circ I = F$ for any F in $\mathbf{Set}^{\mathbf{C}_0}$, and thus by applying the right adjoint $(-) \circ I$ of $\text{Lan}_I(-)$ to the commutative triangle

$$\begin{array}{ccc} & \text{Lan}_I(\mathcal{D}_0(j)) & \\ \nearrow & & \searrow \\ T = \text{Lan}_I(T_0) & \xrightarrow{\sigma} & \text{colim } \mathcal{D} = \text{Lan}_I(\text{colim } \mathcal{D}_0) \end{array}$$

we get that T_0 factors through $\mathcal{D}_0(j)$. The essential unicity of this factorisation is likewise inherited from the essential unicity of the factorisation for T . $\square$

We can now prove that finitely presentable and finitely presented functors coincide. This proof generalises, and is to some extent adapted from, Theorem 2.16 of [AMV11].

Proposition 3.46. *Let $\mathbf{C}$ be a small category and let $T : \mathbf{C} \rightarrow \mathbf{Set}$, then T is finitely presentable in $\mathbf{Set}^{\mathbf{C}}$ iff it is strongly finitely presented.*

Proof. The ‘if’ direction follows from the fact that hom functors are absolutely presentable in $\mathbf{Set}^{\mathbf{C}}$ (see the proof of Lemma 3.38) and thus finitely presentable objects of $\mathbf{Set}^{\mathbf{C}}$. Therefore, by Lemma 3.26, so is a finite colimit of hom functors.

For the ‘only if’ direction, assume that T is finitely presentable in $\mathbf{Set}^{\mathbf{C}}$. We need to show that T is a finite colimit of hom functors. In order to show this we need to prove two things: (1) that T can be written as a the colimit of a diagram defined on a category with finitely many *objects*, and (2) that this category also has finitely many *arrows*.

Let us first show (1). We use the fact that the objects of $\mathbf{C}$ form a set, and that all sets can be expressed as the colimit of their finite subsets, in order to express T as a filtered colimit. Let $\mathbf{J}$ denote the category whose objects are the full subcategories of $\mathbf{C}$ with finitely many objects and whose arrows are the inclusion functors between these categories. For any $\mathbf{C}_0$ of $\mathbf{J}$, let $T_0 : \mathbf{C}_0 \rightarrow \mathbf{Set}$ denote the restriction of T to $\mathbf{C}_0$ and let l_0 denote the inclusion functor $\mathbf{C}_0 \hookrightarrow \mathbf{C}$. We now define a diagram $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ whose colimit will be T : for any $\mathbf{C}_0$ in $\mathbf{J}$ let

$$\mathcal{D}(\mathbf{C}_0) = \text{Lan}_{l_0} T_0$$

The definition of $\mathcal{D}$ on inclusion functors follows automatically: let $l_0^1 : \mathbf{C}_0 \hookrightarrow \mathbf{C}_1$ in $\mathbf{J}$, then by construction of the left Kan extension we have for any object A of $\mathbf{C}$ that

$$\text{Lan}_{l_0} T_0 A = \text{colim} \left(\mathbf{C}_0 \downarrow A \xrightarrow{P_0} \mathbf{C}_0 \xrightarrow{T_0} \mathbf{Set} \right)$$

and since $\mathbf{C}_0 \downarrow A$ is a subcategory of $\mathbf{C}_1 \downarrow A$, the colimit defining $\text{Lan}_{l_1} T_1 A$ is a cocone for the diagram defining $\text{Lan}_{l_0} T_0 A$ and so there must exist a function

$$\mathcal{D}(l_0^1)_A : \text{Lan}_{l_0} T_0 A \rightarrow \text{Lan}_{l_1} T_1 A$$

These functions define a natural transformation $\mathcal{D}(l_0^1) : \text{Lan}_{l_0} T_0 \rightarrow \text{Lan}_{l_1} T_1$. Note that whilst all the morphisms in $\mathbf{J}$ are monos, their images under $\mathcal{D}$ typically are not, as the following example shows: let $\mathbf{C}_0^{\text{ob}} = \{A_1, A_2\}$ and $\mathbf{C}_0^{\text{mor}} = \emptyset$ with unique morphisms $f : A_1 \rightarrow B$ and $g : A_2 \rightarrow B$ in $\mathbf{C}$, then $\mathcal{D}(\mathbf{C}_0)B = \text{Lan}_{l_0} T_0 B = TA_1 + TA_2$. Let now $\mathbf{C}_1^{\text{ob}} = \{A_1, A_2, B\}$ and $\mathbf{C}_1^{\text{mor}} = \{f, g\}$, then $\mathcal{D}(\mathbf{C}_1)B = TB$ which in general is not a subset of $TA_1 + TA_2$.

Let us now show that $T = \text{colim } \mathcal{D}$. Since colimits of functors are computed pointwise, it is enough to show that $TA = \text{colim}(ev_A \circ \mathcal{D})$. To see that TA is a cocone, it is enough to see that we have maps $\mathcal{D}(\mathbf{C}_0)(A) = \text{Lan}_{l_0} T_0(A) \rightarrow TA$ for every $\mathbf{C}_0$. From the construction of the left Kan extension as a colimit in $\mathbf{Set}$, an element of $\text{Lan}_{l_0}(A)$ is an equivalence class of elements of the type $\alpha_0 \in TA_0$ where $f : A_0 \rightarrow A$ is in $\mathbf{C}_0 \downarrow A$. So let us pick α_0 as a representative of this equivalence class, then we can associate with it

$$(d_{\mathbf{C}_0})_A : \text{Lan}_{l_0} T_0(A) \rightarrow TA, [\alpha_0] \mapsto Tf(\alpha_0)$$

To check that this assignment is well-defined, let us choose another representative of the equivalence class of α_0 , say $\alpha_1 \in TA_1$ with $g : A_1 \rightarrow A$ in $\mathbf{C}_0 \downarrow A$. Because α_0 and α_1 are in the same equivalence class, there must exist a zigzag of morphisms connecting α_0 and α_1 . By definition of the morphisms in $\mathbf{C}_0 \downarrow A$, it is easy to see that this implies $Tg(\alpha_1) = Tf(\alpha_0)$, and thus the choice of representative does not matter, and our assignment is well-defined. This shows TA is a cocone for $ev_A \circ \mathcal{D}$ and thus T is a cocone for $\mathcal{D}$. We now check that

$TA = \text{colim}(ev_A \circ \mathcal{D})$, i.e. that for every $\alpha \in TA$, there exist an object $\mathbf{C}_0$ of $\mathbf{J}$ and an element $\alpha_0 \in \text{Lan}_{\mathbf{l}_0} T_0(A)$ such that $(d_{\mathbf{C}_0})_A(\alpha_0) = \alpha$ and that any other choice is related to this one by a zigzag. Clearly, we can always choose $\mathbf{C}_0$ as the full subcategory of $\mathbf{C}$ with only one object: A , in which case we get $\text{Lan}_{\mathbf{l}_0} T_0(A) = TA$ and we can take $\alpha_0 = \alpha$ since $(d_{\mathbf{C}_0})_A = \text{Id}_{TA}$ in this case. For the essential unicity, assume that there exist another object $\mathbf{C}_1$ in $\mathbf{J}$ and another element $[\alpha_1] \in \text{Lan}_{\mathbf{l}_1} T_1 A$ with $(d_{\mathbf{C}_1})_A([\alpha_1]) = \alpha$. By definition of $d_{\mathbf{C}_1}$, this means that there exists an object $f : A_1 \rightarrow A$ in $\mathbf{C}_1 \downarrow A$ such that $Tf(\alpha_1) = \alpha$. If A is an object of $\mathbf{C}_1$ we can take $\alpha_1 = \alpha$ and $\mathcal{D}(\mathbf{l}_0^1)_A = \text{Id}_{TA}$ gives us a (very simple) zigzag. If A is not in $\mathbf{C}_1$ we can embed $\mathbf{C}_1$ and $\mathbf{C}_0$ in the full subcategory $\mathbf{C}_2$ whose objects are those of $\mathbf{C}_1$ plus A . It is easy to see that $\mathcal{D}(\mathbf{l}_0^2)_A = \text{Id}_{TA}$ and $\mathcal{D}(\mathbf{l}_1^2)_A(\alpha_1)$ is precisely $Tf(\alpha_1) = \alpha$, and this provides us with a zigzag

$$\text{Lan}_{\mathbf{l}_0} T_0 A = TA \xrightarrow{\mathcal{D}(\mathbf{l}_0^2)_A} \text{Lan}_{\mathbf{l}_2} T_2 A = TA \xleftarrow{\mathcal{D}(\mathbf{l}_1^2)_A} \text{Lan}_{\mathbf{l}_1} T_1 A$$

connecting α_0 and α_1 which concludes the proof that $T = \text{colim } \mathcal{D}$.

Since $\mathcal{D}$ is directed, and T is assumed to be finitely presentable, Id_T must factor through a unique object $\mathcal{D}(\mathbf{C}_0)$, i.e.

$$T = \text{Lan}_{\mathbf{l}_0} T_0$$

for a certain full subcategory $\mathbf{C}_0$ with finitely many object. We are now going to show that (a) T can be written as a certain colimit of hom functors (b) that this colimit arises from a diagram over a category with only finitely many objects. This will conclude the proof of (1).

For (a), recall by Theorem 3.19 that $T_0 = \text{colim}(\mathbf{Y}_0 \mathbf{U}_{T_0})$ where $\mathbf{Y}_0 : \mathbf{C}_0 \rightarrow \mathbf{Set}^{\mathbf{C}_0}$ is the Yoneda transformation. We will now show that we can extend this result and that

$$T = \text{colim} \left(\mathbf{El}(T_0) \xrightarrow{\mathbf{l}_0 \mathbf{U}_{T_0}} \mathbf{C}^{\text{op}} \xrightarrow{\mathbf{Y}} \mathbf{Set}^{\mathbf{C}} \right)$$

To see that T is a cocone of the diagram $\mathcal{D}_0 = \mathbf{Y}_0 \mathbf{U}_{T_0}$, we use the fact that by Yoneda's lemma for any (A, α) in $\mathbf{El}(T_0)$, α determines a unique natural transformation $c^\alpha : \text{hom}(A, -) \rightarrow T$, all these transformations are compatible by definition of $\mathbf{El}(T_0)$, and T is thus indeed a cocone. To see that $T = \text{colim } \mathcal{D}_0$ we can again check that $TB = \text{colim}(ev_B \circ \mathcal{D}_0)$ at every stage B of $\mathbf{C}$. So let $\beta \in TB$, since $TB = \text{Lan}_{\mathbf{l}_0} T_0 B$ there must exist $f : A \rightarrow B$ in $\mathbf{C}_0 \downarrow B$ and $\alpha \in TA$ such that $Tf(\alpha) = \beta$, but this data provides us precisely with an $f \in \mathcal{D}_0(A, \alpha)B = \text{hom}(A, B)$ such that $c_B^\alpha(f) = \beta$. Now we just need to check that it is essentially unique. Assume that there exists $f' \in \mathcal{D}_0(A', \alpha')$ such that $c_B^{\alpha'}(f') = Tf'(\alpha') = \beta$, i.e. there exists another object of $\mathbf{C}_0 \downarrow B$ and another element $\alpha' \in TA'$ which gets mapped to $\beta \in TB = \text{Lan}_{\mathbf{l}_0} T_0 B$. But by

construction of the left Kan extension as a colimit, this means that there must exist a zigzag in $\mathbf{C}_0 \downarrow B$ connecting f and f' , and this immediately provides us with a zigzag (identical but with all the arrows reversed) connecting $\text{hom}(A, B)$ and $\text{hom}(A', B)$. Thus $T = \text{colim } \mathcal{D}_0$.

For (b), recall that the objects of $\mathbf{El}(T_0)$ are given by the pairs (A, α) with A an object of $\mathbf{C}_0$ and $\alpha \in T_0 A$. We know that $\mathbf{C}_0$ has finitely many objects, thus to show that $\mathbf{El}(T_0)$ has finitely many objects we need to show that $T_0 A$ is always finite. To achieve this let us write $\mathbf{C}_0^{\text{ob}} = \{A_1, \dots, A_n\}$ and consider the collection

$$\mathcal{U}(\mathbf{C}_0) = \{(U_1, \dots, U_n) \mid U_i \subseteq T_0 A_i, |U_i| < \omega, 1 \leq i \leq n\}$$

i.e. the n -tuples of finite subsets of the images of $\mathbf{C}_0$ -objects under T_0 . We now define for any $\vec{U} \in \mathcal{U}(\mathbf{C}_0)$ the functor $T_0^{\vec{U}} : \mathbf{C}_0 \rightarrow \mathbf{Set}$ on objects by

$$T_0^{\vec{U}} A = \bigcup_{f \in (\mathbf{C}_0 \downarrow A)^{\text{ob}}} T f[U_{\text{dom} f}]$$

where $U_{\text{dom} f}$ is the component of $\vec{U}$ corresponding to $\text{dom} f$. On morphisms $g : A \rightarrow B$ the functor is defined as

$$T_0^{\vec{U}} g = T_0 g \upharpoonright T_0^{\vec{U}} A$$

Note the following: (i) $\mathcal{U}$ can be partially ordered by taking the usual set inclusion component-wise and if $\vec{U}, \vec{V} \in \mathcal{U}$ are such that $\vec{U} \subseteq \vec{V}$, then $T_0^{\vec{U}} \hookrightarrow T_0^{\vec{V}}$; (ii) $\mathcal{U}(\mathbf{C}_0)$ is directed (by taking the union component-wise); and (iii) if we define $\mathcal{D}_{\mathcal{U}} : \mathcal{U} \rightarrow \mathbf{Set}^{\mathbf{C}_0}$ (which defines a diagram of monos by (i)) it is clear that $T_0 = \text{colim } \mathcal{D}_{\mathcal{U}}$. We know from Lemma 3.45 that since $T = \text{Lan}_{\mathbf{l}_0} T_0$, $\mathbf{C}_0$ is a full subcategory, and T is finitely presentable, T_0 must also be finitely presentable. Therefore, T_0 must be isomorphic to one of the $T_0^{\vec{U}}$, $\vec{U} \in \mathcal{U}$, i.e. image finite (or ‘super finitary’, see [AMV11]). Therefore $\mathbf{El}(T_0)^{\text{ob}}$ is indeed finite and T is a colimit of hom functors whose diagram is defined on a category with finitely many objects. This concludes the proof of (1).

Now we need to show (2), i.e. that $\mathbf{El}(T_0)$ also has finitely many *morphisms*. To prove this, let us consider the set of all finite subcategories of $\mathbf{C}_0$, i.e. the subcategories of $\mathbf{C}_0$ with only finitely many morphisms. These form a set $\mathcal{V}$ which is partially ordered by the usual inclusion. Note also that by taking the union of two subcategories (i.e. the union of their sets of objects and morphisms), we easily get that $\mathcal{V}$ is directed. We now define the diagram $\mathcal{D}_m : \mathcal{V} \rightarrow \mathbf{Set}^{\mathbf{C}_0}$ on objects by

$$\mathcal{D}_m(\mathbf{C}_i) = \text{Lan}_{\mathbf{l}_i}(T_0 \circ \mathbf{l}_i) : \mathbf{C}_0 \rightarrow \mathbf{Set}$$

where $\mathbf{l}_i : \mathbf{C}_i \hookrightarrow \mathbf{C}_0$. The definition of $\mathcal{D}_m$ on morphisms follows automatically: suppose $\mathbf{l}_{ik} : \mathbf{C}_i \hookrightarrow \mathbf{C}_k$ is an inclusion between two finite subcategories of $\mathbf{C}_0$,

and let us write T_i for $T_0 \circ l_i$; then at any stage A of $\mathbf{C}_0$ we have by construction of the left Kan extension that

$$\mathrm{Lan}_{l_i} T_i A = \mathrm{colim} \left(\mathbf{C}_i \downarrow A \xrightarrow{P_i} \mathbf{C}_i \xrightarrow{T_i} \mathbf{Set} \right)$$

where $P_i(f : l_i(B) \rightarrow A) = B$. If $\mathbf{C}_i \hookrightarrow \mathbf{C}_k$, the colimit defining $\mathrm{Lan}_{l_k} T_k A$ is a cocone for the diagram $T_i P_i$ and therefore there must exist a unique function

$$\mathcal{D}_m(l_{ik})_A : \mathrm{Lan}_{l_i} T_i A \rightarrow \mathrm{Lan}_{l_k} T_k A$$

These functions define a natural transformation $\mathcal{D}_m(l_{ik}) : \mathcal{D}_m(\mathbf{C}_i) \rightarrow \mathcal{D}_m(\mathbf{C}_k)$.

The proof that $T_0 = \mathrm{colim} \mathcal{D}_m$ is exactly the same as the proof above that T was colimit for $\mathcal{D}$; the fact that the subcategories $\mathbf{C}_i$ are not full does not play any role in the construction (however it does imply that the restriction functors T_i are *not* finitely presentable). Therefore, since T_0 is finitely presentable we must have $T_0 = \mathrm{Lan}_{l_i} T_i$ for a subcategory $\mathbf{C}_i$ of $\mathbf{C}_0$ with finitely many morphisms. As was shown in Proposition 1.6, the composition of two adjunctions is an adjunction and thus we have that $\mathrm{Lan}_{l_0}(-) \circ \mathrm{Lan}_{l_i}(-) \dashv ((-)\circ l_i) \circ ((-)\circ l_0)$, i.e.

$$T = \mathrm{Lan}_{l_0}(\mathrm{Lan}_{l_i} T_i) = \mathrm{Lan}_{l_0 l_i} T_i$$

We can then use this fact to show that

$$T = \mathrm{colim}(\mathbf{El}(T_i) \xrightarrow{l_0 l_i \cup T_i} \mathbf{C}^{\mathrm{op}} \xrightarrow{\mathbf{Y}} \mathbf{Set}^{\mathbf{C}})$$

in exactly the same way as we showed that $T = \mathrm{colim} \mathbf{Y} l_0 \cup T_0$ earlier, and $\mathbf{El}(T_i)$ is a category with finitely many objects and finitely many morphisms. $\square$

Remark 3.47. We have voluntarily decomposed the proof above into two steps: first show that T can be written as a colimit of a diagram of hom functors defined on an index category with finitely many *objects*, and second that this index category can also be assumed to have finitely many *morphisms*. The reader might wonder if this decomposition is necessary, and why we didn't immediately write T as a filtered colimit of left Kan extensions of restrictions of T to finite categories. The problem with this approach is that by taking non full subcategories, the restriction functors do not inherit presentability (i.e. we cannot use Lemma 3.45), which means that we cannot prove image finiteness. Thus, though it looks like we are performing the same trick twice, there is a good reason to proceed in this way. Moreover, as we shall see when we reach the notion of λ -generated functor, this distinction between the number of objects and the number of morphisms in a diagram will become very important.

Corollary 3.48. *For any small category $\mathbf{C}$, the category $\mathbf{Set}^{\mathbf{C}}$ is locally finitely presentable.*

Proof. Using exactly the same construction as in the previous proof of Proposition 3.46, it is clear that any T in $\mathbf{Set}^{\mathbf{C}}$ can be expressed as a filtered colimit of functors of the type $T_{\mathbf{C}_0}$ where $\mathbf{C}_0$ is a finite subcategory of $\mathbf{C}$. Since by the same proof each of these $T_{\mathbf{C}_0}$ is a finitely presentable object of $\mathbf{Set}^{\mathbf{C}}$, it means that every object in $\mathbf{Set}^{\mathbf{C}}$ is a filtered colimit of presentable objects. And since $\mathbf{Set}^{\mathbf{C}}$ is also cocomplete, it is locally finitely presentable. $\square$

Corollary 3.49. *Let $\mathbf{C}$ be a small category, and let λ be a regular cardinal, then an object T in $\mathbf{Set}^{\mathbf{C}}$ is λ -presentable iff it is a λ -small colimit of hom functors.*

Proof. For the ‘if’ direction, notice first that since hom functors are absolutely presentable, they are λ -presentable, and thus so is a λ -small colimit of them by Lemma 3.26. For the ‘only if’ direction we proceed as in the proof of Proposition 3.46, but this time we use the fact that the set $\mathbf{C}^{\mathbf{ob}}$ can be expressed as the colimit of its subsets of cardinality smaller than λ . Since λ is regular, this colimit is λ -directed. In consequence, we can write T as a λ -directed colimit of subfunctors in the same fashion as in the proof of Proposition 3.46 and the rest of the proof follows. $\square$

Remark 3.50. Similarly to Theorem 3.12 of [AR94], Proposition 3.46 is not trivial, but it tightens the analogy, started in Remark 3.21, between functors $\mathbf{C} \rightarrow \mathbf{Set}$, $\mathbf{C}$ small, and algebraic varieties. A presentation defined on a small index category $\mathbf{J}$ is like the presentation of a variety, with $\mathbf{J}^{\mathbf{ob}}$ as its set of generators and $\mathbf{J}^{\mathbf{mor}}$ as its sets of relations. If $\mathbf{J}$ is a λ -small category, then the functor is λ -presented, i.e. presented by fewer than λ ‘generators’ and fewer than λ ‘relations’. In particular coproducts of hom functors (i.e. when $\mathbf{J}$ is discrete) can be thought of as the analogue of free algebraic structures.

3.4.2 λ -generatable and λ -generated functors

The standard notion of functor presentation in the literature, is to express a functor T as an epimorphic image of a coproduct of representables whose size is typically bounded by some cardinal (see e.g. [AGT10]). This notion probably dates back to [Fre71] where a functor with such a representation is called ‘petty’. For reasons that will soon be clear we prefer the terminology of [KRV05]:

Definition 3.51 (λ -generated functor, [KRV05]). A functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ will be called **λ -generated** if there exists a set I of cardinality $|I| < \lambda$, and regular epimorphism

$$e : \coprod_{i \in I} \mathrm{Hom}(A_i, -) \twoheadrightarrow T$$

The associated ‘strong’ notion is defined as in the case of strongly presentable functor:

Definition 3.52 (Strongly λ -generated functor). A functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ will be called **strongly λ -generated**, if there exists a presentation $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ of T (i.e. $T = \text{colim } \mathcal{D}$) such that $\mathbf{J}$ has fewer than λ objects, there are no restrictions on the number of morphisms in $\mathbf{J}$.

Proposition 3.53. *A functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ is λ -generated (in the sense of [KRV05]) if it is strongly λ -generated.*

Proof. The proof is the same as the proof of Proposition 3.43. $\square$

Unlike Proposition 3.44, we cannot in general show the converse implication, i.e. that λ -generated functors are strongly λ -generated. However, in the case where the base category $\mathbf{C}$ is sufficiently well-behaved we do have a converse result which we will return to when discussing the canonical presentation in Section 3.5.3. A necessary technical condition on the base category for the following Proposition is that the cardinality of the set of quotients of any object should be bounded. We will call a category **λ -co-well powered** if for every object A of $\mathbf{C}$, the collection of (equivalence classes of) epimorphisms with domain A is a set of cardinality strictly less than λ .

Proposition 3.54. *Let $\mathbf{C}$ be a λ -co-well powered small category with an epi-split-mono factorisation system, then every λ -generated functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ is strongly λ -generated.*

Proof. By assumption, we have a set I of cardinality less than λ , a functor S and a natural transformation e

$$S \xrightarrow[\phi_2]{\phi_1} \coprod_{i \in I} \text{hom}(A_i, -) \xrightarrow{e} T$$

which is the coequalizer of ϕ_1, ϕ_2 . By Yoneda, the transformation e uniquely determines (and is uniquely determined by) a set $\{x^{A_i} \in TA_i \mid i \in I\}$, and for any object X in $\mathbf{C}$, two morphisms $f_i : A_i \rightarrow X$ and $f_j : A_j \rightarrow X$ in the coproduct $\coprod_i \text{hom}(A_i, X)$ are identified precisely when

$$Tf_i(x^{A_i}) = Tf_j(x^{A_j})$$

We use this characterization of e to define the following category $\mathbf{J}$ in two steps. In the first step, the objects of $\mathbf{J}$ will be all the pairs $(A_i, x^{A_i}), i \in I$, and there will be morphism $(A_i, x^{A_i}) \rightarrow (A_j, x^{A_j})$ whenever there exist $f : A_j \rightarrow A_i$ such that $Tf(x^{A_j}) = x^{A_i}$. In a second step we complete $\mathbf{J}$ by adding for every morphism as defined above its epi-split mono factorisation (which exists by assumption on $\mathbf{C}$), i.e. for every $f : A_j \rightarrow A_i$ such that $Tf(x^{A_j}) = x^{A_i}$ we consider the decomposition $f = f^m \circ f^e$ where f^e is epi and f^m is split mono. We then add $f[A_j]$ to $\mathbf{J}^{\text{ob}}$ and we add the morphisms induced by f^e and f^m

to $\mathbf{J}^{\mathbf{mor}}$: f^e defines a morphism $(f[A_j], Tf^e(x^{A_j})) \rightarrow (A_j, x^{A_j})$, and f^m defines a morphism $(A_i, x^{A_i}) \rightarrow (f[A_j], Tf^e(x^{A_j}))$. Finally, we close this collection of morphisms under composition and identity maps. Note that since $\mathbf{C}$ is assumed to be λ -co-well powered and since λ is regular $|\mathbf{J}^{\mathbf{ob}}| < \lambda$.

We now define

$$T' = \text{colim}(\mathbf{J} \xrightarrow{\mathbf{U}} \mathbf{C}^{\text{op}} \xrightarrow{\mathbf{Y}} \mathbf{Set}^{\mathbf{C}})$$

where $\mathbf{U}$ is the obvious forgetful functor. We want to show that $T' = T$. By working pointwise, this means that for any X in $\mathbf{C}$, the following coequalizers

$$SX \rightrightarrows \coprod_{i \in I} \text{hom}(A_i, X) \xrightarrow{e_X} TX$$

and

$$\coprod_{f \in \mathbf{J}^{\mathbf{mor}}} \text{hom}(\text{dom } f, X) \xrightarrow[\psi]{\phi} \coprod_{A \in \mathbf{J}^{\mathbf{ob}}} \text{hom}(A, X) \xrightarrow{q_X} T'X$$

where $\text{in}_f \circ \phi = \text{in}_{\text{dom } f}$ and $\text{in}_f \circ \psi = \text{in}_{\text{cod } f} \circ ((-) \circ f)$, should be the same object. To check that this is the case, note first that we can characterise $T'X$ as the quotient of the coproduct under the smallest equivalence class generated by

$$f : A \rightarrow X \approx_X f' : A' \rightarrow X$$

if there exist $h : A' \rightarrow A$ such that $Th(x^{A'}) = x^{A'}$, where x^A is the element of TA associated with A in $\mathbf{J}$, and similarly for $x^{A'}$. Note that by construction, the objects A of $\mathbf{J}$ are either of the shape $A_i, i \in I$ or $f[A_i]$ for an $f \in \mathbf{J}^{\mathbf{mor}}$. If A is of the shape $f[A_i]$, we can immediately see that $h : A \rightarrow X$ will be in the same equivalence class under $\approx_X$ as $A_i \xrightarrow{f^e} f[A_i] \xrightarrow{h} X$, by construction of $\mathbf{J}$. Thus we can always reason w.l.o.g. by using representatives of the shape $f : A_i \rightarrow X$ for an $i \in I$. From this it follows that we can define a map

$$\theta_X : T'X \rightarrow TX, [A_i \xrightarrow{f} X]_{\approx} \mapsto [A_i \xrightarrow{f} X]_{e_X}$$

where $[A_i \xrightarrow{f} X]_{\approx}$ is the equivalence class of f under $\approx_X$, and $[A_i \xrightarrow{f} X]_{e_X}$ is the equivalence class of maps identified with f by e_X . It is easy to check that θ_X is well-defined: if we start with maps $f_i : A_i \rightarrow X$ and $f_j : A_j \rightarrow X$ such that $f_i \approx_X f_j$, then there exists $h : A_j \rightarrow A_i$ with $Th(x^{A_j}) = x^{A_i}$, and thus $Tf_j(x^{A_j}) = Tf_i(Th(x^{A_j})) = Tf_i(x^{A_i})$, i.e. $e_X(f_i) = e_X(f_j)$. Moreover, θ_X is trivially surjective. It remains to show that θ_X is injective. For this assume that $e_X(f_i) = e_X(f_j)$, and build the epi-split mono decomposition $f_i = f_i^m \circ f_i^e$. We then construct the following diagram

$$\begin{array}{ccccc} & & & & TA_j \\ & & & & \downarrow f_j \\ TA_i & \xrightarrow{Tf_i^e} & Tf[A_i] & \xrightleftharpoons[r]{Tf_i^m} & TX \end{array}$$

where r is a retraction of f_i^m , i.e. $r \circ f_i^m = \text{Id}_{f[A_i]}$. Since $Tf_j(x^{A_j}) = Tf_i(x^{A_i})$ by assumption, it follows that $Tr \circ Tf_j(x^{A_j}) = Tr \circ Tf_i^m \circ Tf_i^e(x^{A_i}) = Tf_i^e(x^{A_i})$, and thus

$$[A_j \xrightarrow{f_j} X]_{\approx_X} = [f[A_i] \xrightarrow{f_i^m} X]_{\approx_X}$$

and since we have $[f[A_i] \xrightarrow{f_i^m} X]_{\approx_X} = [A_i \xrightarrow{f_i} X]_{\approx_X}$ by construction, we can conclude by transitivity that

$$[A_j \xrightarrow{f_j} X]_{\approx_X} = [A_i \xrightarrow{f_i} X]_{\approx_X}$$

and thus θ_X is also injective, which concludes the proof. $\square$

The assumptions on $\mathbf{C}$ under which strongly λ -generated and λ -generated functors coincide are identical (modulo the λ bound) to the conditions isolated in Proposition 4.6 of [Fre71]. This Proposition establishes under which conditions functors called ‘lucid’ are ‘petty’, i.e. generated. In the analogy between functors and elements of a variety, ‘ λ -lucid’ functors are λ -coherent elements, i.e. λ -generated functor whose λ -generated subfunctors are λ -presented (they are therefore λ -presented themselves).

We would now like to have an abstract characterization of functors which are λ -generated, in the same way that being λ -presentable is an abstract characterization of being λ -presented. As it turns out, this characterization exists, and is detailed in section 1.E. of [AR94], but unfortunately, an object satisfying this characterization is also called λ -generated. In order to both avoid confusion and make a parallel with presentable and presented objects we will use the following terminology.

Definition 3.55 (λ -generatable object). An object A in a category will be called **λ -generatable** if $\text{hom}(A, -)$ preserves λ -directed colimits of monomorphisms. In particular a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ will be called λ -generatable if it is a λ -generatable object in $\mathbf{Set}^{\mathbf{C}}$, i.e. if $\text{Nat}(T, -)$ preserves λ -directed colimits of monomorphisms.

We cite without proof the following relevant results from [AR94].

Proposition 3.56 ([AR94] 1.69). *In a locally λ -presentable category, an object is λ -generatable iff it is a strong quotient of a λ -presentable object.*

Theorem 3.57 ([AR94] 1.70). *A category $\mathbf{C}$ is locally-presentable iff it is co-complete and there exist a regular cardinal λ and a small subcategory $\mathbf{G}_\lambda$ of λ -generatable objects such that every object of $\mathbf{C}$ is a λ -directed colimit of its subobjects from $\mathbf{G}_\lambda$.*

Let us now show that as expected, λ -generated functors $\mathbf{C} \rightarrow \mathbf{Set}$, with $\mathbf{C}$ small, are precisely the λ -generatable objects in $\mathbf{Set}^{\mathbf{C}}$. We start with the analogue of Lemma 3.45 for λ -generated functors.

Lemma 3.58. *Let $\mathbf{C}$ be a category and let $\mathbf{l} : \mathbf{C}_0 \hookrightarrow \mathbf{C}$ be the embedding of a full subcategory. Assume that $T : \mathbf{C} \rightarrow \mathbf{Set}$ is a functor and that T_0 is its restriction to $\mathbf{C}_0$, assume further that T and T_0 have the property that $T = \text{Lan}_{\mathbf{l}} T_0$, then if T is a λ -generatable object in $\mathbf{Set}^{\mathbf{C}}$ so is T_0 in $\mathbf{Set}^{\mathbf{C}_0}$.*

Proof. The proof is identical to that of Lemma 3.45 once it is established that if $\mathcal{D}_0 : \mathbf{C}_0 \rightarrow \mathbf{Set}^{\mathbf{C}_0}$ is a diagram of monomorphisms, then so is $\text{Lan}_{\mathbf{l}}(-) \circ \mathcal{D}_0$, i.e. $\text{Lan}_{\mathbf{l}}(-)$ preserves monos. This is rather obvious from the construction of $\text{Lan}_{\mathbf{l}}$: if $F_i \hookrightarrow F_j$ are two functors in $\mathbf{Set}^{\mathbf{C}_0}$ then clearly for all A in $\mathbf{C}_0$

$$\text{colim} \left(\mathbf{C}_0 \downarrow A \xrightarrow{P} \mathbf{C}_0 \xrightarrow{F_i} \mathbf{Set} \right) \hookrightarrow \text{colim} \left(\mathbf{C}_0 \downarrow A \xrightarrow{P} \mathbf{C}_0 \xrightarrow{F_j} \mathbf{Set} \right)$$

i.e. $\text{Lan}_{\mathbf{l}} F_i A \hookrightarrow \text{Lan}_{\mathbf{l}} F_j A$. □

Proposition 3.59. *A functor $T : \mathbf{C} \rightarrow \mathbf{Set}$, $\mathbf{C}$ small, is λ -generated iff it is a λ -generatable object of $\mathbf{Set}^{\mathbf{C}}$.*

Proof. For the ‘if’ direction we proceed as in Proposition 3.46 and Corollary 3.49 and we write T as the following λ -filtered colimit of functors. We define $\mathbf{J}$ to be the category whose objects are the full subcategories of $\mathbf{C}$ with strictly fewer than λ objects and whose arrows are the inclusions between these subcategories. Then we can write define $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ by $\mathcal{D}(\mathbf{C}_0) = \text{Lan}_{\mathbf{l}_0} T_0$ where T_0 is the restriction of T to the full subcategory $\mathbf{C}_0$. The definition of $\mathcal{D}$ on inclusions follows automatically as in Proposition 3.46 and we once again get that

$$T = \text{colim } \mathcal{D}$$

Note once again that whilst $\mathbf{J}$ is a λ -directed category all of whose arrows are monos, the colimit is *not* a colimit of monos. To remedy this we need to take an epi-mono factorisation of all the morphisms $d_{\mathbf{C}_0} : \text{Lan}_{\mathbf{l}_0} T_0 \rightarrow T$ making T a cocone for $\mathcal{D}$ which is defined for every A of $\mathbf{C}$ as

$$(d_{\mathbf{C}_0})_A : \text{Lan}_{\mathbf{l}_0} T_0 A \rightarrow TA, [\alpha_0] \rightarrow Tf(\alpha_0)$$

where $[\alpha_0]$ is an equivalence class with representative $\alpha_0 \in TA_0$ for a certain $f : A_0 \rightarrow A$ in $\mathbf{C}_0 \downarrow A$.

Taking the epi-mono factorization of $d_{\mathbf{C}_0}$ is always possible since $\mathbf{Set}^{\mathbf{C}}$ is a topos. We can therefore define the following diagram

$$\mathcal{D}_m : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}, \mathcal{D}_m(\mathbf{C}_0) = d_{\mathbf{C}_0}[\text{Lan}_{\mathbf{l}_0} T_0]$$

For any inclusion $\mathbf{l}_0^1 : \mathbf{C}_0 \hookrightarrow \mathbf{C}_1$ in $\mathbf{J}$, $\mathcal{D}_m(\mathbf{l}_0^1)$ is the obvious inclusion of $d_{\mathbf{C}_0}[\text{Lan}_{\mathbf{l}_0} T_0]$ into $d_{\mathbf{C}_1}[\text{Lan}_{\mathbf{l}_1} T_1]$. For A in $\mathbf{C}$ we then have that T is the colimit over all $\mathbf{C}_0$ in $\mathbf{J}$ of the direct images

$$(d_{\mathbf{C}_0})_A[\text{Lan}_{\mathbf{l}_0} T_0 A] = \bigcup_{f \in (\mathbf{C}_0 \downarrow A)^{\text{ob}}} Tf[T \text{dom} f]$$

It is clear that $T = \text{colim } \mathcal{D}_m$ and that $\mathcal{D}_m$ is a λ -filtered diagram of monos since $(d_{\mathbf{C}_0})_A[\text{Lan}_{\mathbf{I}_0} T_0 A] \subseteq TA$ and for an inclusion $\mathbf{l}_0^1 : \mathbf{C}_0 \hookrightarrow \mathbf{C}_1$ in $\mathbf{J}$, we clearly have $(d_{\mathbf{C}_0})_A[\text{Lan}_{\mathbf{I}_0} T_0 A] \subseteq (d_{\mathbf{C}_0})_A[\text{Lan}_{\mathbf{I}_1} T_0 A]$. Note that this colimit of monos is the ‘ λ -generalization’ of the colimit of monos in Theorem 2.16 of [AMV11] which is the basis of the notion of *superfinitary functor*.

Since T is λ -generatable, Id_T must factor through one of these monomorphisms and thus $T = d_{\mathbf{C}_0}[\text{Lan}_{\mathbf{I}_0} T_0]$ for a full subcategory $\mathbf{C}_0$ with less than λ -objects. From Proposition 3.46 we know that

$$\text{Lan}_{\mathbf{I}} T_0 = \text{colim} \left(\mathbf{El}(T_0) \xrightarrow{\text{Id}_{T_0}} \mathbf{C}_0^{\text{op}} \xrightarrow{\gamma} \mathbf{Set}^{\mathbf{C}} \right)$$

By the previous Lemma 3.58 we know that T_0 itself is λ -generatable. Therefore, if we proceed as in Proposition 3.46, and express T_0 as a colimit of monos by restricting the cardinality of $T_0 A$ for all A_0 in $\mathbf{C}_0$, we can show that $|T_0 A| < \lambda$ for all A in $\mathbf{C}_0$, and thus $\mathbf{El}(T_0)$ has less than λ objects. In other words $\text{Lan}_{\mathbf{I}} T_0$ can be expressed a quotient of a coproduct of less than λ hom functors, and since T is a quotient of $\text{Lan}_{\mathbf{I}} T_0$, we immediately get that T is a quotient of a coproduct of less than λ hom functors (all epimorphisms in $\mathbf{Set}^{\mathbf{C}}$ are regular).

For the ‘only if’ direction, we assume that T can be written as

$$e : \coprod_{i \in I} \text{Hom}(A_i, -) \twoheadrightarrow T$$

where $|I| < \lambda$. Note that $\coprod_{i \in I} \text{Hom}(A_i, -)$ is a λ -small colimit of hom functors for the diagram $\mathcal{D} = \mathbf{Y} \circ A_{(-)} : I \rightarrow \mathbf{Set}^{\mathbf{C}}$ where $A_{(-)}(i) = A_i$ and I is viewed as a discrete category. By Corollary 3.49, this means that $\coprod_{i \in I} \text{Hom}(A_i, -)$ is a λ -presentable object in $\mathbf{Set}^{\mathbf{C}}$. Since T is a regular (and therefore strong) quotient of a λ -presentable object, it must be a λ -generated object of $\mathbf{Set}^{\mathbf{C}}$ by Proposition 3.56. $\square$

Corollary 3.60. *Let $\mathbf{C}$ be a small category. Every λ -presentable functor in $\mathbf{Set}^{\mathbf{C}}$ is λ -generated.*

Remark 3.61. We can now extend the analogy of Remarks 3.21 and 3.50 further by noting the similarity between λ -generated functors and λ -generated algebraic varieties. The analogy is detailed in Table 3.1. As indicated in this table, we do not know if there exist an abstract characterisation of strongly finitely generated functors of the type ‘preservation of a subclass of all cofiltered limits’. We conjecture that it may be possible to take this analogy further, that a theory of ‘varieties of functor’ might exist in which classes of functors (e.g. left-exact functors) may be described by small categories of generators and relations with a specified form, and that a Birkhoff style theorem for $\mathbf{Set}$ -valued functors might also be found for an appropriate set of operation on functors.

Table 3.1: Analogy between varieties and functors

Algebraic variety $\mathbf{V}$	Set-functors on small category $\mathbf{C}$	
	Strong version	Normal version
Presentation $\langle G, R \rangle$ by sets of generators G and relations R	Diagram $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ factoring through $\mathbf{Y} : \mathbf{C}^{\text{op}} \rightarrow \mathbf{Set}^{\mathbf{C}}$ s. th. $T = \text{colim } \mathcal{D}$	Regular epimorphism: $S \rightrightarrows \prod_{i \in I} \text{hom}(A_i, -) \twoheadrightarrow T$
Generators in G	$\mathbf{j}_{\text{ob}}$	$\{\text{hom}(A_i, -) \mid i \in I\}$
Relations in R	$\mathbf{j}_{\text{mor}}$	S
Finitely presented objects: G and R finite	Strongly finitely presented functor: $\mathbf{J}$ a finite category	Finitely presented functor I finite, S finitely generated
Finitely presentable object in $\mathbf{V}$	Finitely presentable object in $\mathbf{Set}^{\mathbf{C}}$	
Free object in $\mathbf{V}$	$\mathbf{J}$ discrete, i.e. coproduct of hom	
Finitely generated objects: G finite	Strongly finitely generated functor: $\mathbf{j}_{\text{ob}}$ finite	Finitely generated functor: I finite
Finitely generatable object in $\mathbf{V}$	?	Finitely generatable object in $\mathbf{Set}^{\mathbf{C}}$

3.4.3 Free cocompletions

We end this section with a different way of looking at the functor presentations. As we have just seen, when $\mathbf{C}$ is small, the finite colimits of representables are precisely the finitely presentable objects in $\mathbf{Set}^{\mathbf{C}}$ (Proposition 3.46). The filtered colimits of representable functors are precisely the objects of $\mathbf{Set}^{\mathbf{C}}$ which are left exact functors 3.40. Finally, by Theorem 3.19, general colimits of representable functors are precisely the functors in $\mathbf{Set}^{\mathbf{C}}$. These three facts can be interpreted as three different ‘free cocompletions’ of the small category $\mathbf{C}$.

Definition 3.62. Let $\mathbf{C}$ be a category. A **free finite cocompletion** of $\mathbf{C}$ is a full embedding $\mathbf{l}_r : \mathbf{C} \rightarrow \mathbf{Rex}(\mathbf{C})$ such that

- (i) $\mathbf{Rex}(\mathbf{C})$ is finitely cocomplete (in other words right exact or ‘rex’)
- (ii) $\mathbf{l}_r$ is universal amongst such embeddings in the sense that for any other functor $F : \mathbf{C} \rightarrow \mathbf{A}$ with $\mathbf{A}$ finitely cocomplete, there must exist a unique (up to isomorphism) right exact functor $F^* : \mathbf{Rex}(\mathbf{C}) \rightarrow \mathbf{A}$ such that $F^*\mathbf{l}_r = F$.

A **free filtered cocompletion** of $\mathbf{C}$ is a full embedding $\mathbf{l}_f : \mathbf{C} \rightarrow \mathbf{Acc}(\mathbf{C})$ such that

- 1. $\mathbf{Acc}(\mathbf{C})$ has all filtered colimits (and a small subcategory generating $\mathbf{Acc}$ by taking filtered colimits, i.e. $\mathbf{Acc}(\mathbf{C})$ is finitely accessible)
- 2. $\mathbf{l}_f$ is universal amongst such embeddings, i.e. for any other functor $F : \mathbf{C} \rightarrow \mathbf{A}$ where $\mathbf{A}$ has all filtered colimits, there exists an accessible functor $\tilde{F} : \mathbf{Acc}(\mathbf{C}) \rightarrow \mathbf{A}$ such that $\tilde{F}\mathbf{l}_f = F$

A **free cocompletion** of $\mathbf{C}$ is a full embedding $\mathbf{l}_c : \mathbf{C} \rightarrow \mathbf{CoCo}(\mathbf{C})$ such that

- 1. $\mathbf{CoCo}(\mathbf{C})$ has all colimits
- 2. $\mathbf{l}_c$ is universal amongst such embeddings, i.e. for any other functor $F : \mathbf{C} \rightarrow \mathbf{A}$ where $\mathbf{A}$ has all colimits, there exists an accessible functor $\tilde{F} : \mathbf{CoCo}(\mathbf{C}) \rightarrow \mathbf{A}$ such that $\tilde{F}\mathbf{l}_c = F$

As it turns out, each embedding can be taken to be the Yoneda embedding, but we must be careful in our choice of Yoneda embedding. In what preceded, we were interested in studying the structure of the category $\mathbf{Set}^{\mathbf{C}}$, and to this end we relied heavily on the Yoneda transformation $\mathbf{Y} : \mathbf{C}^{\text{op}} \rightarrow \mathbf{Set}^{\mathbf{C}}, A \mapsto \text{hom}(A, -)$. In the present context, we are interested in freely completing $\mathbf{C}$, so the source of the functor must be $\mathbf{C}$, not $\mathbf{C}^{\text{op}}$, i.e. we consider the Yoneda embedding $\mathbf{Y} : \mathbf{C} \rightarrow \mathbf{Set}^{\mathbf{C}^{\text{op}}}, A \mapsto \text{hom}(-, A)$, i.e. the embedding into the category $\hat{\mathbf{C}} = \mathbf{Set}^{\mathbf{C}^{\text{op}}}$ of presheaves. This change of convention does not modify

anything: everything we have done so far can be re-phrased using this embedding. For example, all presheaves can be written as a colimit of *contravariant* hom functors, the proof is exactly the same as that of Theorem 3.19 with the following small difference. If $T : \mathbf{C} \rightarrow \mathbf{Set}$ is contravariant, then there exists a morphism between (A, α) and (B, β) in $\mathbf{El}(T)$ if there exists a morphism f in $\mathbf{C}$ such that $Tf(\beta) = \alpha$. Since T is contravariant, this means that f must be a morphism from $A \rightarrow B$. The forgetful functor therefore has type $\mathbf{El}(T) \rightarrow \mathbf{C}$.

Proposition 3.63. *For a small category $\mathbf{C}$, a free finite cocompletion is given by the category $\mathbf{f.p.}(\widehat{\mathbf{C}})$ of finitely presentable objects in $\widehat{\mathbf{C}}$, together with the Yoneda embedding $\mathbf{Y} : \mathbf{C} \rightarrow \mathbf{f.p.}(\widehat{\mathbf{C}})$.*

Proof. First, we need to show that $\mathbf{Y}$ indeed maps $\mathbf{C}$ into $\mathbf{f.p.}(\widehat{\mathbf{C}})$, but this follows immediately by Proposition 3.38. Moreover, $\mathbf{f.p.}(\widehat{\mathbf{C}})$ is right exact by Lemma 3.26, and $\mathbf{Y}$ is a full embedding, so we need only check the universal property. This is where Kan extensions come into play. Let $F : \mathbf{C} \rightarrow \mathbf{A}$ be a functor and $\mathbf{A}$ be finitely cocomplete, then the span

$$\mathbf{A} \xleftarrow{F} \mathbf{C} \xrightarrow{\mathbf{Y}} \mathbf{f.p.}(\widehat{\mathbf{C}})$$

satisfies the conditions of Proposition 5.36 and there must therefore exist a functor $F^* = \mathbf{Lan}_{\mathbf{Y}} F : \mathbf{f.p.}(\widehat{\mathbf{C}}) \rightarrow \mathbf{A}$ such that $F^* \mathbf{Y} = F$. So all we need to check is that $F^* = \mathbf{Lan}_{\mathbf{Y}} F$ is right exact. The functor F^* is sometimes known as the ‘Yoneda extension’ of F and is known to be a left adjoint, and thus cocontinuous, and a fortiori right exact (see for example [AR94] 1.45. (i)). $\square$

Proposition 3.64. *For a small category $\mathbf{C}$, a free filtered cocompletion is given by the category $\widetilde{\mathbf{C}}$ of presheaves whose category of elements is filtered, i.e. by Proposition 3.40 the category of left-exact functors, together with the Yoneda embedding $\mathbf{Y} : \mathbf{C} \rightarrow \widetilde{\mathbf{C}}$.*

Proof. Note first that since the category of elements of a hom functor has a terminal element by Proposition 3.16, it is trivially filtered, and the Yoneda embedding is indeed defined into $\widetilde{\mathbf{C}}$. The rest of the proof is exactly as in the proof of Proposition 3.63, with Yoneda extensions preserving all colimit, and thus in particular filtered ones. $\square$

Proposition 3.65. *For a small category $\mathbf{C}$, a free cocompletion is given by the category $\mathbf{Set}^{\mathbf{C}^{\text{op}}}$ of presheaves together with the Yoneda embedding.*

Proof. The only thing to show is the universality of $\mathbf{l}_c$, and the proof is identical to Proposition 3.63. $\square$

Note that the embedding $\mathbf{l}_r$ into $\mathbf{Rex}(\mathbf{C})$ does not in general preserve *existing* colimits in $\mathbf{C}$, i.e. $\mathbf{l}_r$ is not cocontinuous. However, this can be remedied by

restricting the embedding to the subcategory of left exact presheaves on $\mathbf{C}$ (which makes sense since the covariant hom preserves limits, see Proposition 1.45 (ii) of [AR94]). This gives rise to the stronger notion of free cocompletion, which is to the free finite completion described above as the Dedekind MacNeille completion is to the canonical extension. Note also that the notion of free finite cocompletion above can be generalized to higher cardinalities by essentially appending a λ in front of all the relevant concepts, i.e. the λ -presentable functors provide a free λ -small cocompletion.

3.5 Presentations of accessible functors

In this section we will extend the results of the previous section to the case where $\mathbf{C}$ is no longer small, but can be approximated by a small subcategory, i.e. we will consider **Set**-functors on accessible categories. As we will see, all the results obtained so far can essentially be transferred to this case, provided that the functor is accessible.

3.5.1 General presentations

We will start by generalizing the construction of Theorem 3.19 to λ -accessible functors. Before we can present the main construction we need a few technical lemmas.

Lemma 3.66. *Let $\mathbf{C}$ be a λ -accessible category, let $\mathbf{C}_\lambda$ be its subcategory of λ -presentable objects and let A be such a λ -presentable object. If we denote by $\mathbf{l}$ the inclusion functor $\mathbf{C}_\lambda \hookrightarrow \mathbf{C}$, then*

$$\mathrm{Lan}_{\mathbf{l}}(\mathrm{hom}_{\mathbf{C}_\lambda}(A, -)) = \mathrm{hom}_{\mathbf{C}}(A, -)$$

Proof. By construction of the left Kan extension, and using the notation of Proposition 3.36, we have for any object B of $\mathbf{C}$

$$\begin{aligned} \mathrm{Lan}_{\mathbf{l}}(\mathrm{hom}_{\mathbf{C}_\lambda}(A, B)) &= \mathrm{colim}(\mathrm{hom}_{\mathbf{C}_\lambda}(A, -) \circ \mathcal{D}_{c_\lambda}^B) \\ &\stackrel{1}{=} \mathrm{colim}(\mathrm{hom}_{\mathbf{C}}(A, -) \circ \mathcal{D}_c^B) \\ &\stackrel{2}{=} \mathrm{hom}_{\mathbf{C}}(A, \mathrm{colim} \mathcal{D}_c^B) \\ &\stackrel{3}{=} \mathrm{hom}_{\mathbf{C}}(A, B) \end{aligned}$$

where (1) follows from the trivial commutative diagram of Proposition 3.36, (2) follows from the fact that $\mathcal{D}_c^B$ is λ -filtered (see (i) of Lemma 3.34) and that A is λ -presentable, and (3) follows from (ii) of Lemma 3.34. $\square$

Lemma 3.66 establishes the commutativity of the following diagram

$$\begin{array}{ccc}
 \mathbf{C}_\lambda^{\text{op}} & \xrightarrow{Y_\lambda} & \mathbf{Set}^{\mathbf{C}_\lambda} \\
 \downarrow \text{I}^{\text{op}} & & \downarrow \text{Lan}_\text{I}(-) \\
 \mathbf{C}^{\text{op}} & \xrightarrow{Y} & \mathbf{Set}^{\mathbf{C}}
 \end{array} \tag{3.4}$$

where Y_λ is the Yoneda transformation $\mathbf{C}_\lambda \rightarrow \mathbf{Set}^{\mathbf{C}_\lambda}$ and Y is the Yoneda transformation from the entire category $\mathbf{C}$. Note also that the λ -presentability of the object defining the hom functor is essential, for a general A in $\mathbf{C}$ we would *not* have $\text{Lan}_\text{I} \text{hom}_{\mathbf{C}_\lambda}(A, -) = \text{hom}_{\mathbf{C}}(A, -)$.

Lemma 3.67. *Let $\mathbf{C}$ be a λ -accessible category and let $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}}$ be a small diagram such that $\mathcal{D}(j)$ is a functor of the form $\text{hom}(A, -)$ where A is λ -presentable for each object j of $\mathbf{J}$, then $\text{colim } \mathcal{D}$ preserves λ -filtered colimits.*

Proof. Let $\mathcal{F} : \mathbf{K} \rightarrow \mathbf{C}$ be a λ -filtered diagram in $\mathbf{C}$. Note first that

$$\begin{aligned}
 \text{colim } \mathcal{D}(\text{colim } \mathcal{F}) &= \text{ev}_{\text{colim } \mathcal{F}} \text{colim } \mathcal{D} \\
 &= \text{colim } \text{ev}_{\text{colim } \mathcal{F}} \circ \mathcal{D}
 \end{aligned}$$

where $\text{ev}_A : \mathbf{Set}^{\mathbf{C}} \rightarrow \mathbf{Set}$ is the ‘evaluation at A ’ functor which preserves colimits (see [Mac98] V.3.). Consider now the functor $\text{ev}_{\text{colim } \mathcal{F}} \circ \mathcal{D}$: for any j in $\mathbf{J}$ we have by assumption that $\mathcal{D}(j)$ is of the form $\text{hom}(A_j, -)$ with A_j presentable. Thus

$$\text{ev}_{\text{colim } \mathcal{F}} \circ \mathcal{D}(j) = \mathcal{D}(j)(\text{colim } \mathcal{F}) = \text{colim}(\mathcal{D}(j) \circ \mathcal{F})$$

i.e. viewed as functors $\mathbf{J} \rightarrow \mathbf{Set}$ we have $\text{ev}_{\text{colim } \mathcal{F}} \circ \mathcal{D} = \text{colim}(\mathcal{D}(-) \circ \mathcal{F})$. Finally, the formula of interchange of colimits (see [Mac98] IX.2.) gives us

$$\text{colim}(\text{colim}(\mathcal{D}(-) \circ \mathcal{F})) = \text{colim}((\text{colim } \mathcal{D}) \circ \mathcal{F})$$

and we can conclude

$$\text{colim } \mathcal{D}(\text{colim } \mathcal{F}) = \text{colim}((\text{colim } \mathcal{D}) \circ \mathcal{F})$$

□

Lemma 3.68. *Let $\mathbf{C}$ be a λ -accessible category, $\text{I} : \mathbf{C}_\lambda \hookrightarrow \mathbf{C}$ be the inclusion of the subcategory of λ -presentable objects, and let $T : \mathbf{C} \rightarrow \mathbf{Set}$ be a λ -accessible functor and $T_\lambda = T \circ \text{I}$, under those assumptions, T inherits any presentation of T_λ . Formally, if*

$$T_\lambda = \text{colim } Y_\lambda \circ \mathcal{D}$$

for a functor $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{C}_\lambda^{\text{op}}$, then

$$T = \text{colim } Y \circ \text{I}^{\text{op}} \circ \mathcal{D}$$

Proof. For every A in $\mathbf{C}$ we have

$$\begin{aligned}
 TA &= \text{Lan}_{\mathbf{I}} T_{\lambda} A && \text{Proposition 3.36} \\
 &= \text{Lan}_{\mathbf{I}} \text{colim}(\mathbf{Y}_{\lambda} \circ \mathcal{D})(A) \\
 &= \text{colim} \text{Lan}_{\mathbf{I}}(\mathbf{Y}_{\lambda} \circ \mathcal{D})(A) && \text{Lan}_{\mathbf{I}}(-) \text{ is a left adjoint} \\
 &= \text{colim}(\mathbf{Y} \circ \mathbf{I}^{\text{op}} \circ \mathcal{D})(A) && \text{Diagram (3.4)}
 \end{aligned}$$

□

We are now ready to present the main theorem on functor presentation.

Theorem 3.69. *Let $\mathbf{C}$ be a λ -accessible category and $T : \mathbf{C} \rightarrow \mathbf{Set}$. T is accessible iff it is a small colimit of hom-functors $\text{hom}(A, -)$ where A is λ -presentable.*

Proof. The ‘if’ direction was shown in the previous Lemma 3.67.

For the ‘only if direction’, we once again write $\mathbf{I} : \mathbf{C}_{\lambda} \hookrightarrow \mathbf{C}$ and we start by considering the restriction $T_{\lambda} = T \circ \mathbf{I}$. Since this $\mathbf{C}_{\lambda}$ is small we can perform the construction of Theorem 3.19 and express T_{λ} as

$$T_{\lambda} = \text{colim} \left(\mathbf{El}(T_{\lambda}) \xrightarrow{\mathbf{U}_{T_{\lambda}}} \mathbf{C}_{\lambda}^{\text{op}} \xrightarrow{\mathbf{Y}_{\lambda}} \mathbf{Set}^{\mathbf{C}_{\lambda}} \right)$$

and where $\mathbf{Y}_{\lambda}$ is used as above. It follows that $TA = \text{colim}(\mathbf{Y}_{\lambda} \circ \mathbf{U}_{T_{\lambda}})(A)$ by Lemma 3.68.

□

3.5.2 Size-bound presentations

The proof above hints at the fact that accessible functors are those for which the results developed in the case where $\mathbf{C}$ was small, can be extended to the case where $\mathbf{C}$ is accessible. As the following results will illustrate, this is indeed the case. We start with the accessible version of Proposition 3.46 and Corollary 3.49 wrapped together. The general strategy of all these proofs is to reduce the study of T to that of T_{λ} , its restriction to the small subcategory of λ -presentable objects. We start with the following useful lemma.

Proposition 3.70. *Let $\mathbf{C}$ a λ -accessible category and let $\mu \leq \lambda$ be a regular cardinal, then an object T in the category $\mathbf{Acc}_{\lambda}$ of λ -accessible functors $\mathbf{C} \rightarrow \mathbf{Set}$ is μ -presentable iff it is a μ -small colimit of functors of the shape $\text{hom}(A, -)$, A λ -presentable.*

Proof. The ‘if’ direction is clear: if A is λ -presentable, then $\text{hom}(A, -)$ is accessible. From the proof of Proposition 3.38 it is clear that all hom functors

are λ -presentable since they are absolutely presentable and by Lemma 3.26, a λ -small colimit of λ -presentable objects is λ -presentable and $\mu \leq \lambda$.

For the ‘only if’ direction, note first that since T is λ -accessible, we have by Proposition 3.36 that $T = \text{Lan}_! T_\lambda$ where $! : \mathbf{C}_\lambda \rightarrow \mathbf{C}$ is the inclusion of the small subcategory λ -presentable objects and T_λ is the restriction of T to $\mathbf{C}_\lambda$. By Lemma 3.45 we therefore know that T_λ is μ -presentable because T is. We now apply Proposition 3.46 to T_λ , and we get that T_λ is a μ -small colimit of hom functors in $\mathbf{Set}^{\mathbf{C}_\lambda}$. We can now follow the same derivation as in the proof of Theorem 3.69 with the difference that the colimit of the diagram $\mathcal{D}_{T_\lambda} : \mathbf{El}(T_\lambda) \rightarrow \mathbf{Set}^{\mathbf{C}_\lambda}$ expressing T_λ is now a μ -small colimit of hom functors, and thus we finally get that T itself is a μ -small colimit of $\text{hom}(A, -)$ functors with all the A s λ -presentable as desired. $\square$

Note that the ‘if’ direction of the previous Proposition can be strengthened in the following way: from Proposition 3.38 we know that *any* hom functor is λ -presentable in $\mathbf{Set}^{\mathbf{C}}$ for any λ , and thus a μ -small colimit of hom functors is μ -presentable in $\mathbf{Set}^{\mathbf{C}}$. However, $\text{hom}(A, -)$ can only be μ -presentable in $\mathbf{Acc}_\lambda$ if A is λ -presentable.

Proposition 3.71. *Let $\mathbf{C}$ be a λ -accessible category and let $\mathbf{Acc}_\lambda$ be the subcategory of $\mathbf{Set}^{\mathbf{C}}$ consisting of λ -accessible functors. $\mathbf{Acc}_\lambda$ is locally finitely presentable.*

Proof. Let us first show that $\mathbf{Acc}_\lambda$ is cocomplete. Let us consider a diagram $\mathcal{D}_F : \mathbf{I} \rightarrow \mathbf{Acc}_\lambda$ and let $F = \text{colim } \mathcal{D}_F$. Of course, since $\mathbf{Acc}_\lambda$ is a subcategory of $\mathbf{Set}^{\mathbf{C}}$ which is cocomplete, we can build F in $\mathbf{Set}^{\mathbf{C}}$, but to check that F is an object of $\mathbf{Acc}_\lambda$ we need to check that it preserves λ -filtered colimits. So let $\mathcal{F} : \mathbf{J} \rightarrow \mathbf{C}$ be a λ -filtered diagram, by using the interchange of colimits we have:

$$\begin{aligned} F(\text{colim } \mathcal{F}) &= \text{colim } \mathcal{D}_F(\text{colim } \mathcal{F}) \\ &= \text{colim}(ev_{\text{colim } \mathcal{F}} \circ \mathcal{D}_F) \\ &= \text{colim}(\text{colim}(\mathcal{D}_F(-) \circ \mathcal{F})) \\ &= \text{colim}((\text{colim } \mathcal{D}_F) \circ \mathcal{F}) \\ &= \text{colim}(F \circ \mathcal{F}) \end{aligned}$$

and thus $\mathbf{Acc}_\lambda$ is cocomplete.

Let us now show that all λ -accessible functors are filtered colimits of finitely presentable objects in $\mathbf{Acc}_\lambda$. Again, we use the fact that $T = \text{Lan}_! T_\lambda$, and apply known results to T_λ . In this instance we apply Corollary 3.48 and get that T_λ is a filtered colimit of finitely presentable objects in $\mathbf{Set}^{\mathbf{C}_\lambda}$, i.e. a filtered colimit of finite colimits of functors of the shape $\text{hom}(A, -)$ with A λ -presentable (in

C). Lemma 3.68 then allows us to extend this result to T itself, i.e. T is also a filtered colimit of finite colimits of functors of the shape $\text{hom}(A, -)$ with A λ -presentable. By Proposition 3.70 this means that T is a filtered colimit of finitely presentable objects in $\mathbf{Acc}_\lambda$, which concludes the proof. $\square$

It is interesting to note from the previous Proposition that the ‘ λ ’ of ‘ λ -accessible’ controls the type of hom functors used to present a λ -accessible functor, i.e. homs of the shape $\text{hom}(A, -)$ for A λ -presentable in the base category, but is independent of the ‘size’ of the presentation which is controlled by size of the small category $\mathbf{C}_\lambda$. For a λ -accessible functor $T : \mathbf{C} \rightarrow \mathbf{Set}$, any bounds on the size of the presentation are therefore also inherited from bounds on the size of the presentation of its restriction T_λ . The notion of strongly λ -presented and strongly λ -generated functors, can therefore be lifted directly to accessible functor.

Definition 3.72. A λ -accessible functor on a λ -accessible category $\mathbf{C}$ will be called **strongly μ -presented** (resp. generated) if its restriction T_λ to the subcategory of λ -presentable objects is strongly μ -presented (resp. generated).

Similarly, we can use the left Kan extension functor to lift the notion of presented and generated functors to the accessible case. Using the same notation as above, if we assume that T_λ is μ -generated, then T_λ can be expressed as a regular epi $\coprod_{i \in I} \text{Hom}(A_i, -) \xrightarrow{e} T_\lambda$ for a set $|I| < \mu$. We now apply the left Kan extension functor $\text{Lan}_1(-)$ to get

$$\text{Lan}_1\left(\coprod_{i \in I} \text{Hom}(A_i, -)\right) = \coprod_{i \in I} \text{Lan}_1 \text{Hom}(A_i, -) = \coprod_{i \in I} \text{Hom}(A_i, -) \twoheadrightarrow \text{Lan}_1(T_\lambda) = T$$

i.e. $\text{Lan}_1(-)(e) : \coprod_{i \in I} \text{Hom}(A_i, -) \rightarrow T$, which is a regular epi too since $\text{Lan}_1(-)$ being a left adjoint preserves all colimits. The same construction clearly defines the notion of being μ -presented for a λ -accessible functor. Finally, the lifting by $\text{Lan}_1(-)$, together with Proposition 3.70 also allows us to transfer Propositions 3.46 and 3.59 to accessible functors, i.e. a λ -accessible functor is μ -presented (resp. generated) iff it is a μ -presentable (resp. μ -generatable) object in $\mathbf{Acc}_\lambda$.

The following proposition describes a class of base categories for which finitely-generated and finitely-presented functors coincide. This class includes the categories of distributive lattices and of boolean algebras.

Proposition 3.73. *Let $\mathbf{C}$ be a locally finite variety (i.e. a variety whose finitely generated objects are finite, see [Bez01]) and let $T : \mathbf{C} \rightarrow \mathbf{Set}$ be a finitary functor preserving monos. Then T is finitely presentable iff it is finitely generatable.*

Proof. We only need to show that finitely generatable functors are finitely presentable. So let T be finitely generatable and let T_ω be its restriction to the subcategory $\mathbf{C}_\omega \xrightarrow{l_\omega} \mathbf{C}$ of finitely presented objects. Since T is finitary, $T = \text{Lan}_{l_\omega} T_\omega$

by 3.36, and thus if we can express T_ω as a finite colimit of hom functors, we are done by Lemma 3.68.

The idea of the proof is as follows. Just as we did in the proof of Propositions 3.46 and 3.59, we start by considering the category whose objects are the subcategories $\mathbf{C}_0$ of $\mathbf{C}_\omega$ with finitely many objects and whose arrows are the obvious inclusions. We can immediately note that since $\mathbf{C}$ is a locally finite variety, (i.e. finitely generated objects are finite), each such subcategory has finitely many morphisms. We will now modify this category slightly: for every finite subcategory $\mathbf{C}_0$ of $\mathbf{C}_\omega$, we define the category $\overline{\mathbf{C}}_0$ as the full subcategory of $\mathbf{C}$ whose objects are those of $\mathbf{C}_0$ plus all their strong quotients. Since the strong quotient of a finitely presented variety is a finitely generated one, and since those are finite by assumption, $\overline{\mathbf{C}}_0$ also has finitely many objects and finitely many morphisms. We now consider the category $\mathbf{J}$ whose objects are all the completions $\overline{\mathbf{C}}_0$ of finite subcategories $\mathbf{C}_0$, and whose morphisms are the inclusions, and we define a diagram $\mathcal{D} : \mathbf{J} \rightarrow \mathbf{Set}^{\mathbf{C}_\omega}$ by $\mathcal{D}(\overline{\mathbf{C}}_0) = \text{Lan}_{\mathbf{l}_0}(T_\omega \circ \mathbf{l}_0)$ where $\mathbf{l}_0$ is the inclusion $\mathbf{l}_0 : \overline{\mathbf{C}}_0 \hookrightarrow \mathbf{C}$. It is easy to see that $\mathbf{J}$ is directed and that

$$T_\omega = \text{colim } \mathcal{D}$$

So far, the construction looks a lot like that of Propositions 3.46 and 3.59, but the key difference is that each $\overline{\mathbf{C}}_0$ has all quotients. This means that transformations $d_{\overline{\mathbf{C}}_0} : \text{Lan}_{\mathbf{l}_0}(T_\omega \circ \mathbf{l}_0) \rightarrow T_\omega$ are monos. Indeed, by construction of the left Kan extension as a colimit, for every $A \in \mathbf{C}_\omega$, $\text{Lan}_{\mathbf{l}_0}(T_\omega \circ \mathbf{l}_0)A$ is taken as a colimit over the objects of $\overline{\mathbf{C}}_0 \downarrow A$, but any such $f : C \rightarrow A$ for C in $\overline{\mathbf{C}}_0$ can be factored as a strong epi (which lives in $\overline{\mathbf{C}}_0$) followed by a mono, and since T is assumed to preserve monos, the colimit is in fact taken over a collection of monos pointing into TA . Formally,

$$\text{colim} \left(\overline{\mathbf{C}}_0 \downarrow A \xrightarrow{\mathbf{U}_0} \mathbf{C}_\omega \xrightarrow{T_\omega} \mathbf{Set} \right) = \text{colim} \left((\overline{\mathbf{C}}_0 \twoheadrightarrow A) \xrightarrow{\mathbf{V}_0} \mathbf{C}_\omega \xrightarrow{T_\omega} \mathbf{Set} \right)$$

where $(\overline{\mathbf{C}}_0 \twoheadrightarrow A)$ is the category whose objects are all the monos from $\overline{\mathbf{C}}_0$ into A , and whose morphisms are defined in the obvious fashion. The functors $\mathbf{U}_0, \mathbf{V}_0$ are the obvious forgetful functors. From this it follows that for every $\overline{\mathbf{C}}_0$ the map $d_{\overline{\mathbf{C}}_0} : \text{Lan}_{\mathbf{l}_0}(T_\omega \circ \mathbf{l}_0) \twoheadrightarrow T_\omega$ is a mono, and thus $\mathcal{D}$ in fact defines a filtered diagram of monos. Since T is finitely generatable, T_ω is also finitely generatable by Lemma 3.58, and it must therefore be equal to one of the elements of this diagram, i.e.

$$T_\omega = \text{Lan}_{\mathbf{l}_0}(T_\omega \circ \mathbf{l}_0) \tag{3.5}$$

for a certain finite full subcategory $\overline{\mathbf{C}}_0 \xrightarrow{\mathbf{l}_0} \mathbf{C}_\omega$.

Now we can use the exact same technique as in the proof of Proposition 3.46 and show that $T_0 = T_\omega \circ \mathbf{l}_0$ is image finite. The key is that this technique allows us to write T_0 as a filtered colimit of monos (all of which are image finite), and

we can thus use the fact that T_ω is finitely generatable together with Lemma 3.58 and Eq. (3.5) to conclude that T_0 is finitely generatable and must therefore be equal to one of the image finite functors. From there it follows easily that the category $\mathbf{El}(T_0)$ has finitely many objects and finitely many morphisms (since these arise from the finitely many morphisms of $\overline{\mathbf{C}}_0$). Two applications of 3.68 then allow us to conclude that T is strongly finitely presented, and thus finitely presented and thus finitely presentable (by Proposition 3.46). $\square$

Remark 3.74. The construction above (viz. extending the subcategory defining a λ -generated functor by adding all the quotients) can be carried out in general, but the result is that the enlarged subcategory contains objects which are no longer λ -presentable, but λ -generated. So we can express a λ -generated functor either as a direct image of a left Kan extension over a subcategory of λ -presentable objects, or, as a Kan extension over a subcategory of λ -generated objects, in which case the direct image is taken in the base category, and we need preservation of monos.

3.5.3 The canonical presentation.

Let again $\mathbf{C}$ be a λ -accessible category, $\mathbf{l} : \mathbf{C}_\lambda \hookrightarrow \mathbf{C}$ be the inclusion of the subcategory of λ -presentable objects, $T : \mathbf{C} \rightarrow \mathbf{Set}$ be a λ -accessible functor, and $T_\lambda = T \circ \mathbf{l}$. By Theorem 3.69 we have

$$T = \operatorname{colim}(\mathbf{El}(T_\lambda) \xrightarrow{\mathbf{U}} \mathbf{C}_\lambda^{\operatorname{op}} \xrightarrow{\mathbf{Y} \circ \mathbf{l}} \mathbf{Set}^{\mathbf{C}})$$

This presentation of T is what we will call the **canonical presentation**. The purpose of this section is to show what this colimit actually looks like. Re-writing this colimit as a coequalizer of a coproduct gives us

$$\coprod_{f \in \mathbf{El}(T_\lambda)^{\operatorname{mor}}} \mathbf{Y} \circ I \circ \mathbf{U}(\operatorname{dom}(f)) \xrightleftharpoons[\psi]{\phi} \coprod_{(A, \alpha) \in \mathbf{El}(T_\lambda)^{\operatorname{ob}}} \mathbf{Y} \circ I \circ \mathbf{U}((A, \alpha)) \longrightarrow T$$

where ϕ and ψ are defined uniquely by their action on the summands of the coproduct (by the universal property of the coproduct) as follows: $\phi \circ \operatorname{in}_f = \operatorname{in}_{\operatorname{dom}(f)}$ and $\psi \circ \operatorname{in}_f = \operatorname{in}_{\operatorname{cod}(f)} \circ \mathbf{Y} I \mathbf{U}(f)$. By using the definitions and Yoneda's lemma, this can be re-written as

$$\coprod_{f \in \mathbf{C}_\lambda \downarrow A, A \in \mathbf{C}_\lambda^{\operatorname{ob}}} T \operatorname{dom}(f) \times \operatorname{hom}(A, -) \xrightleftharpoons[\psi]{\phi} \coprod_{A \in \mathbf{C}_\lambda^{\operatorname{ob}}} T A \times \operatorname{hom}(A, -) \longrightarrow T$$

since by definition of $\mathbf{El}(T_\lambda)$ there exists a morphism in $\mathbf{El}(T_\lambda)$ for each choice of an object A in $\mathbf{C}_\lambda$, a morphism $f : B \rightarrow A$ and an element $\beta \in TB = T \operatorname{dom}(f)$. The natural transformations ϕ and ψ are now defined as follows:

$\phi \circ \text{in}_f = \text{in}_{\text{cod}(f)} \circ \langle Tf \times Id \rangle$ and $\psi \circ \text{in}_f = \text{in}_{\text{dom}(f)} \circ \langle Id \times ((-) \circ f) \rangle$. For any C of $\mathbf{C}$, we get that

$$\coprod_{f \in \mathbf{C}_\lambda \downarrow A, A \in \mathbf{C}_\lambda^{\text{ob}}} T\text{dom}(f) \times \text{hom}(A, C) \xrightleftharpoons[\psi]{\phi} \coprod_{A \in \mathbf{C}_\lambda^{\text{ob}}} TA \times \text{hom}(A, C) \twoheadrightarrow TC$$

i.e. TC is the quotient of the set $\coprod_{A \in \mathbf{C}_\lambda^{\text{ob}}} TA \times \text{hom}(A, C)$ by the smallest equivalence relation generated for the relation $\sim_C$ defined for any $(\alpha, g) \in TA \times \text{hom}(A, C)$ and $(\alpha', g') \in TA' \times \text{hom}(A', C)$ by

$$(\alpha, g) \sim_C (\alpha', g') \text{ if there exists } f : B \rightarrow B' \text{ and } (\beta, f') \in TB \times \text{hom}(B', C) \\ \text{such that } \phi(\beta, f') = (\alpha, g), \psi(\beta, f) = (\alpha', g')$$

i.e. if there exist $f : B \rightarrow B'$ and $(\beta, f') \in TB \times \text{hom}(B', C)$ such that

$$Tf(\beta) = \alpha, f' = g \text{ (and thus in particular } A = B')$$

and

$$\beta = \alpha', g' = f' \circ f \text{ (and thus in particular } A' = B)$$

We thus get the following characterisation:

$$(\alpha, g) \sim_C (\alpha', g') \text{ if there exists } f : A' \rightarrow A \text{ such that } g' = g \circ f \text{ and } \alpha = Tf(\alpha')$$

The **canonical presentation** can thus equivalently be defined as the regular epi transformation

$$q^c : \coprod_{A \in \mathbf{C}_\lambda^{\text{ob}}} TA \times \text{hom}(A, -) \twoheadrightarrow T$$

obtained by quotienting the functor $\coprod_{A \in \mathbf{C}_\lambda^{\text{ob}}} TA \times \text{hom}(A, -)$ at each stage C of $\mathbf{C}$ by the smallest equivalence relation generated by $\sim_C$.

In the general case where $\mathbf{C}$ is an arbitrary λ -accessible category, we know that the cardinality of $\mathbf{C}_\lambda^{\text{ob}}$ must be at least λ by Corollary 3.35, and thus the size of the coproduct must also be at least λ . However, the precise cardinality of $\mathbf{C}_\lambda^{\text{ob}}$ seems to be a largely open question. In the case where $\mathbf{C}$ is the locally finitely presentable category **Grp**, it is known that the set of finitely presentable objects is countable (see [Rut97] Chapter 9) but non-recursively enumerable (see [Cal03]). In this latter case the canonical presentation q^c of a finitary functor takes the shape T

$$q_X^c : \coprod_{n \in \mathbb{N}} Tn \times \text{hom}(n, X) \twoheadrightarrow TX$$

Our notion of canonical presentation (i.e. via the equivalence relation generated by $\sim_X$) can also be found in e.g. [VK11] but is not the ‘standard’ notion found in most of the literature (e.g. [AGT10, AT90, AMM12, Lea11, KL09]) where the focus is to present functors as quotients of polynomial functors rather than as a colimit of representables. Nevertheless, in the case of accessible **Set**-endofunctors the two notions coincide, although this is not immediately obvious.

Proposition 3.75. *Let $\sim_X$ be the relation on $\coprod_{n \in \mathbb{N}} Tn \times \text{hom}(n, X)$ defined above and let $\approx_X$ be the relation on $\coprod_{n \in \mathbb{N}} Tn \times \text{hom}(n, X)$ defined for any $(\sigma, f) \in Tm \times \text{hom}(m, X)$ and $(\tau, g) \in Tn \times \text{hom}(n, X)$ as in 3.9. of [AGT10], namely*

$$(\sigma, f) \approx_X (\tau, g) \text{ if } Tf(\sigma) = Tg(\tau)$$

The relations $\sim_X$ and $\approx_X$ generate the same equivalence relation on X .

Proof. The reader might recognise the relations $\sim_X$ and $\approx_X$, since they are the same relations as in the proof of Proposition 3.54. In fact the proof that $\sim_X$ and $\approx_X$ define the same quotient on a set X is identical to the proof that the quotients defined by the corresponding relations in Proposition 3.54 coincide. $\square$

As we have seen in Section 1.5.4, a regular epimorphism can be seen as a quotient under some equations. In particular a regular epi from a polynomial functor can alternatively be seen as an pair (Σ, E) where Σ is the signature defining the polynomial functor - in the case of a λ -accessible functor this would be the set of λ -presentable objects - and E is a set of equations - i.e. the set of pairs $((\sigma, f), (\tau, g))$ originating from a common element through the maps ϕ, ψ defined above. It is usually in terms of such a pair (Σ, E) that the notion of *presentation of a functor* is defined (see for example [AGT10]).

3.5.4 Atomic objects of $\text{El}(T)$.

In this section we will show that the canonical presentation which tends to have a lot a redundant information about the functor, can often be simplified. To show this we will need the notion of atomic object in a category. Let $\mathbf{C}$ be a well-powered category, and let us firstly assume that $\mathbf{C}$ has not got an initial object, we will then say that an object A in $\mathbf{C}$ is **atomic** if it has no proper subobjects, i.e. if all the monos with codomain C are isos, or using the Sub functor if $\text{Sub}(A) = \{\text{Id}_A\}$. If $\mathbf{C}$ has an initial object, we will define its atomic objects as the atomic objects of $\mathbf{C}_0$, the subcategory of $\mathbf{C}$ where the initial objects and all arrows from it have been removed. The atomic elements of $\mathbf{C}$ form a fully faithful subcategory $\mathbf{Atom}(\mathbf{C})$ of $\mathbf{C}$.

Example 3.76. 1. If we consider a poset $\mathbf{P}$ as a category, then since all arrows are monic, an element of $\mathbf{P}$ is atomic precisely when it's an atom in the order-theoretic sense.

2. In **Set**, an object is atomic iff it is a singleton set, so there really is only one atom in **Set**, and it is 1.
3. The atomic objects in **Grp**, i.e. the groups without proper subgroups are the cyclic groups $\mathbb{Z}_n$ with n prime. Indeed, if G is an atomic object in

Grp, then for any $g \in G, g \neq e$ the subgroup $\langle g \rangle$ generated by g must be the entire group G . Thus G must be cyclic. If G is infinite, then it must be isomorphic to $\mathbb{Z}$, but $\mathbb{Z}$ has plenty of proper subgroups (e.g. the even integers), so G must be finite and cyclic. Since the subgroups of cyclic groups are cyclic and since for a finite cyclic group of order n its subgroups' order must be a divisor of n , it is clear that $\mathbb{Z}_n$ cannot have a proper subgroup if n is prime.

4. A similar argument as in the previous case, or a duality argument from example 2, shows that the only atomic object in **BA** is the free boolean algebra on one element, due to the idempotency of all the operators defining a BA.
5. The category which will interest us is $\mathbf{El}(T)^{\text{op}}$, for a functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ where $\mathbf{C}$ is small or λ -accessible. Intuitively, an object (A, α) of $\mathbf{El}(T)^{\text{op}}$ is atomic if α is built using the entirety of the object A , i.e. there cannot exist $i : A' \hookrightarrow A$ and $\alpha' \in TA'$ such that $\alpha = Ti(\alpha')$. For example, if $T = \mathcal{P}$ and $A = 3 = \{0, 1, 2\}$ then $(3, \{0, 1\})$ is not atomic since it comes from the injection of $(2, \{0, 1\})$, but $(2, \{0, 1\})$ itself is atomic. We will denote the subcategory of atomic objects of $\mathbf{El}(T)^{\text{op}}$ by $\mathbf{Atom}(T)^{\text{op}}$.

From the example above it is clear that atomicity in $\mathbf{El}(T)^{\text{op}}$ is related to the notion of base.

Proposition 3.77. *Let $\mathbf{C}$ be well-powered and have pullbacks of monos, and let $T : \mathbf{C} \rightarrow \mathbf{Set}$ preserve intersections, then (A, α) is atomic in $\mathbf{El}(T)^{\text{op}}$ iff $A = B(\alpha)$.*

Proof. If $A = B(\alpha)$, then by definition of $\mathbf{El}(T)$ we have $\alpha \in TB(\alpha)$. Clearly if there existed a strict mono $i : A' \hookrightarrow A$ and an element $\alpha' \in TA'$ such that $Ti_{A'}^A(\alpha') = \alpha$, then A' would be a strict subobject of the intersection of all such subobjects which is a contradiction, so (A, α) is indeed atomic.

For the converse, assume that (A, α) is atomic. Since $i_{B(\alpha)}^A : B(\alpha) \hookrightarrow A$, if we can show that $\alpha \in Ti_{B(\alpha)}^A[TB(\alpha)]$, then we're done since (A, α) is atomic. But this is exactly the content of Lemma 3.6. $\square$

Corollary 3.78. *Let $\mathbf{C}$ be well-powered and have pullbacks of monos, and let $T : \mathbf{C} \rightarrow \mathbf{Set}$ preserve intersections, then for every object (A, α) in $\mathbf{El}(T)^{\text{op}}$, there exists an atomic object (B, β) and a monomorphism $i : B \rightarrow A$ such that $Ti(\beta) = \alpha$.*


Note that since all locally presentable categories are complete and well-powered (see [AR94] 1.56), Corollary 3.78 applies to all **Set**-valued functors on locally presentable categories which preserve intersections, in particular all functors defined on sets or algebraic varieties which preserve intersections.

3.5.5 The minimal presentation


Theorem 3.79. *Let $\mathbf{C}$ be a small category with pullbacks of monos and let $T : \mathbf{C} \rightarrow \mathbf{Set}$ preserve intersections, then $\downarrow : \mathbf{Atom}(T) \hookrightarrow \mathbf{El}(T)$ is cofinal.*

Proof. Let (A, α) be an object in $\mathbf{El}(T)$, we need to find an atomic object (B, β) and an $\mathbf{El}(T)$ -morphism $(A, \alpha) \rightarrow (B, \beta)$. The existence of such an object and such a morphism is precisely the content of Corollary 3.78 (since small categories are automatically well-powered).

Next consider two $\mathbf{El}(T)$ -morphisms $f_1 : (A, \alpha) \rightarrow (B_1, \beta_1)$ and $f_2 : (A, \alpha) \rightarrow (B_2, \beta_2)$ with (B_1, β_1) and (B_2, β_2) atomic. Then we need to find a zigzag of morphisms in $(A, \alpha) \downarrow \downarrow$ between f_1 and f_2 . Since $\mathbf{C}$ has all intersections, we can factorize f_1, f_2 through their direct images $f_1[B_1]$ and $f_2[B_2]$ and we get the following commutative diagram



by definition of the base of α and the fact that m_1 and m_2 are monomorphisms. It is then easy to check that we have the following zigzag of morphisms in $(A, \alpha) \downarrow \downarrow$:



Notice the reversal of arrows between the two diagrams which is due to the definition of $\mathbf{El}(T)$. $\square$

As an immediate corollary of the previous result we get the following representation theorem.

Theorem 3.80. *Let $\mathbf{C}$ be a small category closed under intersections and let $T : \mathbf{C} \rightarrow \mathbf{Set}$ preserve intersections, then*

$$T = \operatorname{colim} \left(\mathbf{Atom}(T) \xrightarrow{!} \mathbf{El}(T) \xrightarrow{U_T} \mathbf{C}^{\operatorname{op}} \xrightarrow{Y} \mathbf{Set}^{\mathbf{C}} \right)$$

Proof. Immediate by Theorems 3.79 and 3.33. $\square$

We would now like to proceed as we did for the previous representation theorem, namely by extending the result for small category to a result for λ -accessible categories and λ -accessible functors. However, compared with Theorem 3.19, Theorem 3.80 has an additional condition placed on the small category, and this condition might not in general hold in the small subcategory $\mathbf{C}_\lambda$ of λ -presentable objects in a λ -accessible category. The following example from [Mol68] shows that this is indeed the case.

Example 3.81. Let $F : \mathbf{Set} \rightarrow \mathbf{Grp}$ be the free group functor, and consider the direct product $F2 \times F1$, i.e. $F2 \times \mathbb{Z}$. This group can be finitely presented by $\langle a, b, c \mid ac = ca, bc = cb \rangle$, in particular it is finitely presented. Now consider the finitely presented subgroups $G = F2 = \langle a, b \rangle$ and $H = \langle a, bc \rangle$. As the following argument will show, the intersection $G \cap H = \langle b^i ab^{-i}, i \in \mathbb{N} \rangle$, which is clearly *not* finitely presentable. Using the commutation rules of the presentation of $F2 \times \mathbb{Z}$, it is clear that for any $i \in \mathbb{N}$ we have

$$b^i ab^{-i} = b^i c^i c^{-i} ab^{-i} = b^i c^i ac^{-i} b^{-i} = (bc)^i a (bc)^{-i}$$

and thus $\langle b^i ab^{-i}, i \in \mathbb{N} \rangle$ is clearly a subgroup both of G and H . Conversely, let $w(a, bc)$ be a word in $F2 \times \mathbb{Z}$ built from a and bc only, then using the commutation rules we can re-group all the c 's and re-write $w(a, bc) = w'(a, b)c^k$ where k is the sum of the exponents of the b 's in $w'(a, b)$. If $w(a, bc)$ lies in the intersection of G and H , we must clearly have $k = 0$, i.e. $w'(a, b) \in \langle b^i ab^{-i}, i \in \mathbb{N} \rangle$.

To our knowledge, no characterisation of the λ -accessible categories whose λ -presentable objects are closed under intersections exists. Indeed, as the previous example has shown us, even finite intersections are in general problematic. So the best we can hope for is to isolate some useful classes of accessible categories that do exhibit this property. The following Proposition aims at doing just that.

Proposition 3.82. *Let $\mathbf{C}$ be a locally finite variety, then the class of finitely presentable objects of $\mathbf{C}$ is closed under intersections.*

Proof. Let $\{A_i \hookrightarrow B\}_{i \in I}$ be a set of finitely presented subobjects of an algebra B . Since each A_i is finite, and since limits of algebras are created in $\mathbf{Set}$, $A = \bigcap_i A_i$ is also finite. To show that it is finitely generated, let $F \dashv U : \mathbf{Set} \rightarrow \mathbf{C}$ denote the usual free/forgetful adjunction. The counit of the adjunction ϵ_A provides a regular epi $\epsilon_A : FUA \twoheadrightarrow A$ since the adjunction $F \dashv U$ is monadic (see

[KP93]). We thus have a presentation of A using finitely many generators given by $\langle \mathbf{U}A \mid \mathbf{U}\ker(\epsilon_A) \rangle$. $\square$

Examples of locally finite varieties include **DL**, **BA** and **Set** which should convince the reader that requiring the closure of the set of presentable objects under intersection is not too strict a requirement in the context of coalgebraic logics. Moreover, for **Set**-endofunctors we have the following useful fact.

Proposition 3.83 ([KKV12b]). *Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be finitary, then T preserves intersections.*

Theorem 3.84. *Let $\mathbf{C}$ be a λ -accessible category whose small subcategory $\mathbf{l}_\lambda : \mathbf{C}_\lambda \hookrightarrow \mathbf{C}$ of λ -presented objects is closed under intersections. For any intersection-preserving λ -accessible functor $T : \mathbf{C} \rightarrow \mathbf{Set}$ let T_λ be its restriction to $\mathbf{C}_\lambda$, we then have*

$$T = \operatorname{colim} \left(\mathbf{Atom}(T_\lambda) \xrightarrow{\mathbf{l}} \mathbf{El}(T_\lambda) \xrightarrow{\mathbf{l}_\lambda^{\operatorname{opp}} \mathbf{U}_{T_\lambda}} \mathbf{C}^{\operatorname{op}} \xrightarrow{\mathbf{Y}} \mathbf{Set}^{\mathbf{C}} \right)$$

Proof. The proof is identical to that of Theorem 3.69, with the difference that Theorem 3.80 allows us to represent T_λ as the colimit

$$T_\lambda = \operatorname{colim} \left(\mathbf{Atom}(T_\lambda) \xrightarrow{\mathbf{l}} \mathbf{El}(T_\lambda) \xrightarrow{\mathbf{U}_{T_\lambda}} \mathbf{C}_\lambda^{\operatorname{op}} \xrightarrow{\mathbf{Y}} \mathbf{Set}^{\mathbf{C}_\lambda} \right)$$

and the conclusion follows from Lemma 3.68. $\square$

As we saw earlier, colimits can be expressed as coequalizers whose domains are coproducts. Under the hypotheses of Theorem 3.84, a functor T can be represented via the regular epi-transformation:

$$q^m : \coprod_{A \in \mathbf{C}_\lambda^{\operatorname{ob}}} \mathbf{B}_A^{-1} \times \operatorname{Hom}(A, -) \rightrightarrows T$$

where $\mathbf{B}_A^{-1} = \{\alpha \in TA \mid \mathbf{B}_A(\alpha) = A\}$ and where q_C^m identifies $(\sigma, f) \in \mathbf{B}_A^{-1} \times \operatorname{Hom}(A, C)$ and $(\tau, g) \in \mathbf{B}_B^{-1} \times \operatorname{Hom}(B, C)$ if there exists $h : A \rightarrow B$ such that $f = g \circ h$ and $\tau = Th(\sigma)$. The transformation q^m will be called the **minimal presentation** of T .

Example 3.85. We illustrate the construction in the case of the finite powerset functor $\mathcal{P}_\omega$. The canonical presentation of $\mathcal{P}_\omega$ is given by

$$q_X^c : \coprod_{n \in \mathbb{N}} \mathcal{P}_\omega n \times \operatorname{hom}(n, X) \rightrightarrows \mathcal{P}_\omega X$$

where q_X^c identifies (U, f) and (V, g) where $U \subset \{1, \dots, m\}$ and $V \subset \{1, \dots, n\}$ when there exist a map $h : m \rightarrow n$ such that $f = g \circ h$ and $\mathcal{P}_\omega h(U) = h[U] = V$,

or, equivalently by Proposition 3.75, when $\mathcal{P}_\omega f(U) = f[U] = \mathcal{P}_\omega g(V) = g[V]$. The minimal presentation is given by

$$q_X^m : \coprod_{n \in \mathbb{N}} n \times \text{hom}(n, X) \twoheadrightarrow \mathcal{P}_\omega X$$

since for each n there is only one atomic element, namely $(n, \{n\})$. The regular epi q_X^m identifies (U, f) and (V, g) under the same circumstances as for q_C^m . Notice how much redundant information was removed when going from the canonical to the minimal presentation.

3.6 Lifting presentations using adjunctions

Let $\mathbf{C}$ be a λ -accessible category for which there exists an adjunction $\mathbf{F} \dashv \mathbf{U}$ where $\mathbf{F} : \mathbf{Set} \rightarrow \mathbf{C}$ is a ‘free’ functor and $\mathbf{U} : \mathbf{C} \rightarrow \mathbf{Set}$ is a forgetful functor. In this section we will show how we can use the results developed so far together with the adjunction $\mathbf{F} \dashv \mathbf{U}$ to give a representation of a large class of functors $T : \mathbf{C} \rightarrow \mathbf{C}$. We restrict ourselves to finitary functors, but all the results that follow can be adapted easily to the general case of λ -accessible functors.

One way to represent T might be the following: we know how to represent accessible functors $\mathbf{C} \rightarrow \mathbf{Set}$ as colimits of representable functors, and from there as regular quotients of coproducts of representables. So, for an endofunctor T on $\mathbf{C}$ we could try to represent $\mathbf{U}T : \mathbf{C} \rightarrow \mathbf{Set}$. However, as we have illustrated earlier in the case of $\mathbf{C} = \mathbf{Grp}$, representing a λ -accessible functor $\mathbf{C} \rightarrow \mathbf{Set}$ can be difficult practically since it will involve a coproduct over all λ -presentable objects, which is already an incredibly complicated set for the case of finitely presented groups. The situation is simpler in the case of $\mathbf{BA}$, where the finitely presented boolean algebras are those of the form $\mathcal{P}n$, for all n finite. However, apart from the complexity of listing all presentable objects in a category, there is another reason for which we choose another path, namely that in the case of $\mathbf{C} = \mathbf{BA}$ we would like BAEs to emerge from our presentations, i.e. we would like $\mathbf{BA}$ -morphisms between freely generated BAs, or equivalently $\mathbf{Set}$ -morphisms between their sets of generators. Ideally, we would therefore like to involve the free functor $\mathbf{F}$.

With these two objections to representing $\mathbf{U}T$ in mind, a better way to proceed is to find a representation of the functor $\mathbf{U}T\mathbf{F} : \mathbf{Set} \rightarrow \mathbf{Set}$, since presentable objects in $\mathbf{Set}$ are easy to deal with, and with the free functor in the mix we will hopefully get to BAEs in the case of $\mathbf{C} = \mathbf{BA}$.

The first requirement to build a representation of a finitary $T : \mathbf{C} \rightarrow \mathbf{C}$ by using a representation of $\mathbf{U}T\mathbf{F}$ is naturally that $\mathbf{U}T\mathbf{F}$ should be finitary. Since $\mathbf{F}$ is cocontinuous and T is finitary, this means that $\mathbf{U}T\mathbf{F}$ is finitary iff $\mathbf{U}$ is finitary. A adjunction $\mathbf{F} \dashv \mathbf{U}$ with this property is called a **finitary adjunction**, and

the associated monad is called a **finitary monad**. These concepts are readily extended to λ -accessible adjunctions and monads. A well-known class of finitary adjunctions is given by the following well-known result:

Lemma 3.86. *If $\mathbf{C}$ is a finitary algebraic variety, then $\mathbf{U} : \mathbf{C} \rightarrow \mathbf{Set}$ creates filtered colimits*

Proof. See e.g. [Mac98] IX. 2, or [Joh86]. $\square$

Assuming $\mathbf{U}$ and T finitary, we can represent $\mathbf{UTF}$ by using Theorem 3.69:

$$\mathbf{UTF} = \text{colim} \left(\mathbf{El}((\mathbf{UTF})_\omega) \xrightarrow{\text{loV}} \mathbf{Set} \xrightarrow{\mathbf{Y}} \mathbf{Set}^{\mathbf{Set}} \right)$$

where $\mathbf{l} : \mathbf{Set}_f \hookrightarrow \mathbf{Set}$ is the inclusion of the full subcategory of finite sets, $(\mathbf{UTF})_\omega = \mathbf{UTFI}$ and $\mathbf{V} : \mathbf{El}((\mathbf{UTF})_\omega) \rightarrow \mathbf{Set}_f$ is the obvious forgetful functor. This yields the canonical presentation

$$q^m : \coprod_{n \in \mathbb{N}} \mathbf{UTF}(n) \times \text{hom}(n, -) \rightarrow \mathbf{UTF} \quad (3.6)$$

However, it is ultimately T we want a representation for, and to go from a representation of $\mathbf{UTF}$ to a representation of T we use the adjunction as follows. We pre-compose with $\mathbf{U}$ and get a presentation

$$q_U^m : \coprod_{n \in \mathbb{N}} \mathbf{UTF}(n) \times \text{hom}(n, \mathbf{U}(-)) \rightarrow \mathbf{UTFU} \quad (3.7)$$

which is regular and has a polynomial functor as its domain. We will now ‘kill’ the $\mathbf{FU}$ term by using the counit of the adjunction, or more precisely $\mathbf{UT}\epsilon$. This will give us a regular transformation onto $\mathbf{UT}$, which by the adjunction will give us a transformation onto T . Since we want a presentation, i.e. a regular epi from a coproduct of hom functors, we want $\mathbf{UT}\epsilon$ to be a regular epi. For a general adjunction the counit is in general *not* a regular epimorphism. The class of adjunctions for which the counit is a regular epi is in fact well-known to category theorists.

Definition 3.87 ([KP93]). An adjunction $\mathbf{F} \dashv \mathbf{U} : \mathbf{Set} \rightarrow \mathbf{C}$ is said to be of **descent type** if any of the following equivalent conditions is satisfied:

1. the comparison functor $\mathbf{K} : \mathbf{C} \rightarrow \mathbf{Set}^{\mathbf{U}\mathbf{F}}$ is full and faithful
2. ϵ is a regular epi-transformation

This class of adjunctions has been used in [KP93] and [VK11] to prove a slightly different type of presentation results (i.e. not directly related to the canonical colimit description of accessible functors) for finitary functors between enriched categories. Note that the more frequently studied class of *monadic* adjunctions, i.e. those for which the comparison functor is *an equivalence* of categories, is a subclass of the class of adjunctions of descent type.

Lemma 3.88. *The counit of an adjunction $F \dashv U : \mathbf{C} \rightarrow \mathbf{D}$ of monadic type is a regular epi.*

Proof. Consider the following fork:

$$UFUFUA \begin{array}{c} \xrightarrow{UFU\epsilon_A} \\ \xrightarrow{U\epsilon_{FUA}} \end{array} UFUA \xrightarrow{U\epsilon_A} UA$$

and let us show that it is a split coequalizer, the result will then follow from Beck's Monadicity Theorem. First, notice that η_{UA} is a section of $U\epsilon_A$ by definition of the interaction between the unit and the counit in an adjunction. Similarly, η_{UFUA} is a section of $U\epsilon_{FUA}$. So we just need to show that

$$\eta_{UA} \circ U\epsilon_A = UFU\epsilon_A \circ \eta_{UFUA}$$

but this is just a consequence of the naturality of η . $\square$

For $T\epsilon$ to be regular, there is unfortunately no choice: we need to impose that T preserve regular epimorphisms. This requirement was already shown to be particularly helpful in Sections 1.3 and 1.4, and also, albeit in a slightly different context, in Chapter 3 of [Mye11]. As it turns out, this is all we need to ask from our functors as the following result indicates.

Proposition 3.89 ([Bor94a]). *Let $F \dashv U : \mathbf{C} \rightarrow \mathbf{D}$ be an adjunction, then U preserves regular epimorphisms.*

Proof. Proposition 4.3.9. of [Bor94a]. $\square$

In particular if ϵ is a regular epi-transformation and T preserves regular epis, then $UT\epsilon$ is a regular epi. There is one last step before we have shown that UT can be written as a regular quotient of a polynomial functor. Recall the regular epi q_U^m of Eq. (3.7), we need to combine it with $UT\epsilon$ to get a presentation

$$UT\epsilon \circ q_U^m : \coprod_{n \in \mathbb{N}} UTF(n) \times \text{hom}(n, U(-)) \twoheadrightarrow UTFU \twoheadrightarrow UT$$

To combine these two regular epis we need regular epis to compose. This is not always the case as is shown in [Ple00], but we have discussed a natural framework in which this is the case in Section 1.4, namely regular categories. Since the categories we are interested in are all regular, this will not cause us any restriction. We summarize our results as follows.

Theorem 3.90. *Let $\mathbf{C}$ be an accessible category in which the composition of two regular epis is a regular epi, $F \dashv U : \mathbf{C} \rightarrow \mathbf{Set}$ be a finitary adjunction of descent type, and let $T : \mathbf{C} \rightarrow \mathbf{C}$ be a finitary functor which preserves regular epis, then T has a canonical presentation by the regular epi:*

$$F \left(\coprod_{n \in \mathbb{N}} UTF(n) \times \text{hom}(n, U(-)) \right) \twoheadrightarrow T$$

Proof. From the presentation $UT\epsilon \circ q_U^m$ defined above, we get the adjoint transpose transformation described in the statement, and we just need to show that it is a regular epi. This follows immediately from the fact that F is cocontinuous and that ϵ is a regular epi. $\square$

In particular any finitary endofunctor $T : \mathbf{BA} \rightarrow \mathbf{BA}$ or $T : \mathbf{DL} \rightarrow \mathbf{DL}$ which preserves regular epis can be given a presentation of the type described in the Theorem above. As we will now see in Chapter 4, such a transformation allows us to present T -algebras on $\mathbf{BA}$ as quotiented $F(\coprod_{n \in \mathbb{N}} UTF(n) \times \text{hom}(n, U(-)))$ -algebras, i.e. a BAE as we have shown in Chapter 1.

A very similar presentation of endofunctors on $\mathbf{BA}$ can be found in [BK06] and in [VK11]. In [BK06] the notion of presentation is introduced in an *ad-hoc* fashion, motivated by the example of modal algebras. Our presentation here is closer to that of [VK11], although their starting point lies with adjunctions $F \vdash U$ of descent type (in an enriched setting), whilst our presentation starts with the notion of presentable object in $\mathbf{Set}^{\mathbf{C}}$, with the adjunction used to lift presentations to $\mathbf{C}^{\mathbf{C}}$.

Chapter 4

Translations

This chapter will be of a relatively technical nature. Its primary use will be to develop the tools necessary to formulate a theory of completeness-via-canonicity for the ∇ flavour of coalgebraic logic. This version of coalgebraic logic is based on the category **BA**, its positive fragment (which would be defined over **DL**) has not yet been studied. This chapter will thus be concerned entirely with algebras in **BA** on the syntax side and coalgebras in **Set** on the semantics side. We note also that Theorem 4.5 makes essential use of a property of **Set**, and it is not clear how this could be adapted to a more general category such as **Pos**.

In the previous chapter we have seen how well-behaved **BA**-endofunctors can be given a presentation in terms of polynomial functors. A natural question to ask in this context is: given two such functors S, T and a natural transformation $q : S \rightarrow T$ (where S might define a presentation of T), what can we say about the relationship between the coalgebraic logics associated with S and T ? Is there a syntactic relationship? And what happens at the semantic level? These questions seem natural but, as far as we know, have not really been studied systematically in the literature.

4.1 Syntax translations

Recall from Chapter 1, that we can consider coalgebraic logics as coming in three flavours. In the predicate lifting style, the syntax is determined by a signature Σ of operators and associated arities. The result is that we can view the language of such a logic as the initial algebra for the functor

$$S : \mathbf{BA} \rightarrow \mathbf{BA}, A \mapsto F \left(\prod_{\sigma \in \Sigma} (UA)^{\text{ar}(\sigma)} \right) = FS_{\Sigma}U$$

where S_Σ is the polynomial **Set** functor associated with the signature Σ . In light of the previous chapter, we can consider this functor as being trivially presented by itself, and we can leave it at that. It is clear that the structure map of any S -algebra $s : SA \rightarrow A$ is in bijective correspondence with a **Set**-morphism $\coprod_{\sigma \in \Sigma} (UA)^{\text{ar}(\sigma)} \rightarrow UA$, i.e. with the data of a Σ -BAE with reduct A .

In the nabla formalism, the syntax is determined by the **Set**-endofunctor T defining the semantics, and the language $\mathcal{L}_T$ defined in Chapter 1 can be viewed as the initial algebra for the functor

$$FTU : \mathbf{BA} \rightarrow \mathbf{BA}$$

In this case, the functor defining the syntax is therefore not presented. To get a representation of FTU , note that by the adjunction $F \dashv U$, any presentation of TU will lead to a presentation of FTU , and any presentation of T trivially gives a presentation of TU . Thus we need to present T . Since T is assumed to be finitary and weak pullback preserving, we have two canonical choices: we can use the canonical or the minimal presentation, since, by Proposition 3.83, T will then preserve all intersections. As it turns out, the minimal presentation has some very nice properties, and we will therefore often choose this presentation for T , i.e. we have a regular epi

$$SX = \coprod_{n \in \mathbb{N}} \text{hom}(n, X) \times B_n^{-1} \twoheadrightarrow T$$

The structure map $s : FSUA \rightarrow A$ of any FSU-algebra is therefore once again in bijective correspondence with a **Set**-morphism $\coprod_{n \in \mathbb{N}} \text{hom}(n, UA) \times B_n^{-1} \rightarrow UA$, i.e. with the data of a $(B_n^{-1})_{n \in \mathbb{N}}$ -BAE with reduct A .

Finally in the abstract formalism, the syntax is determined by a **BA**-functor L . As was mentioned in Chapter 1, we need to make some assumptions on L , namely that it preserves regular epis and weak pullbacks and that it is a variator. Again, the language defined by L can be thought of as the initial L algebra which exists by 1.4. If L is finitary (or just accessible), we can use the results from the last section of the previous chapter to give a presentation to L .

In the last two cases (and trivially in the first), the presentation of the syntax functor provides us with a new syntax functor whose algebras can be viewed as BAEs, which is precisely the reason for which we consider these presentations in the first place. The task is now to understand how we can go from the language of the presentation to the original language.

Let K, L be two regular epi preserving variators, and let $q : K \twoheadrightarrow L$ be a regular epi-transformation. Let $Q : \mathbf{Alg}_{\mathbf{BA}}(L) \rightarrow \mathbf{Alg}_{\mathbf{BA}}(K)$ denote the functor defined by $Q(A, \alpha) = (A, \alpha \circ q_A)$ and for $f : (A, \alpha) \rightarrow (B, \beta)$ (i.e. $f \circ \alpha = \beta \circ Lf$) by $Qf = f$ which defines a K -algebra morphism by naturality of q .

Theorem 4.1. *Let $K, L : \mathbf{BA} \rightarrow \mathbf{BA}$ be two finitary functors such that K preserves regular epimorphisms, and let $q : K \rightarrow L$ be a natural transformation. Let $F_K \dashv U_K : \mathbf{Set} \rightarrow \mathbf{Alg}_{\mathbf{BA}}(K)$ denote the obvious free/forgetful adjunction associated with K , and similarly $F_L \dashv U_L$ for L . There exist a unique natural transformation*

$$\xi : F_K \rightarrow QF_L$$

Moreover, ξ is a regular epi-transformation whenever q is a regular epi-transformation.

Proof. The existence of ξ is not difficult to show. Since the free K -algebra over an object A of $\mathbf{BA}$ is the initial $K(-) + A$ algebra, we get by initiality the existence of a unique morphism

$$\begin{array}{ccc} KF_K(A) + A & \xrightarrow{K\xi_A + \text{Id}_A} & KF_L(A) + A \\ \downarrow & & \downarrow q_{F_L(A)} + \text{Id}_A \\ & & LF_L(A) + A \\ \downarrow & & \downarrow \\ F_K(A) & \xrightarrow{\xi_A} & F_L(A) \end{array} \quad (4.1)$$

It is not difficult to see that ξ is natural, and the first statement follows.

Now, let us show that ξ_A is a regular epi. Let us write $K' = K(-) + A$ and $L' = L(-) + A$. By construction of the initial algebras, we have the following diagram where $U : \mathbf{BA} \rightarrow \mathbf{Set}$ is the forgetful functor and 0 is the initial object of $\mathbf{BA}$

$$\begin{array}{ccccc} & UK'0 & \xrightarrow{UK'^1!} & UK'^20 & \xrightarrow{UK'^2!} \dots \\ & \swarrow & \downarrow Uq_0 & \downarrow U(q_{L'0} \circ K'q_0) & \\ \text{colim}_i UK'^i0 & & UL'0 & \xrightarrow{UL'^1!} & UL'^20 \xrightarrow{UL'^2!} \dots \\ \downarrow \phi & \swarrow & \downarrow & \downarrow & \\ & & \text{colim}_j UL'^j0 & & \end{array}$$

where ϕ is the unique arrow arising from the fact that $\text{colim}_j UL'^j0$ is a co-cone for the diagram defining $\text{colim}_i UK'^i0$. It is easy to check that this is a commutative diagram. Note also that since we're assuming that q is a regular epi-transformation, then Uq is also a regular epi by Proposition 3.89. Moreover, since K is assumed to preserve regular epis, we in fact have that every vertical arrow in the diagram is epi in $\mathbf{Set}$. By construction of colimits in $\mathbf{Set}$, we know that any element $x \in \text{colim}_j UL'^j0$ can be traced back to an element $y \in UL'^n0$

for a certain n , and thus by the fact that the vertical arrow landing in UL'^n0 is surjective, there must exist an element $y' \in UK'^n0$ which gets mapped to an element $x' \in \text{colim}_i UK'^i0$, and by commutativity of the diagram $\phi(x') = x$. Thus ϕ is surjective.

Next, notice that since U preserves (and creates) filtered colimits, we actually have

$$\text{colim}_i UK'^i0 = U \text{colim}_i K'^i0 = U\mu K'$$

and similarly for L' . By unicity of the the arrow ϕ we must also have $\phi = U\xi_A$. So we have that $U\xi_A$ is regular epi, and in particular it is the coequalizer of its kernel pair. Since U is monadic over **Set**, it reflects exact sequences (see Proposition 1.16), and ξ_A must be therefore be a regular epi. $\square$

The following trivial fact will be useful: it is easy to see that

$$U_K Q = U_L$$

This means in particular that

$$U_K \xi_{FV} : U_K F_K \rightrightarrows U_K Q F_L = U_L F_L$$

It is clear that since $U_K \xi_{FV}$ is just the underlying **Set** map of ξ_{FV} , it is (regular) epi. We will sometimes write $U_K \xi_{FV}$ as $(-)^q : U_K F_K \rightarrow U_L F_L$ and call it the **syntax translation** or **translation map** between the languages.

The natural transformation ξ also interacts in an interesting and useful way with the counits of the adjunctions $F_K \dashv U_K$ and $F_L \dashv U_L$ as the following Proposition shows:

Proposition 4.2. *Let K, L, q, Q and ξ be as above, and let ϵ^K and ϵ^L be the counits of the adjunctions $F_K \dashv U_K$ and $F_L \dashv U_L$ respectively, we then have for any A in $\mathbf{Alg}_{\mathbf{BA}}(L)$ that*

$$\epsilon_{QA}^K = Q\epsilon_A^L \circ \xi_{U_L A} \quad (4.2)$$

Proof. The proof is a consequence of the construction of free algebras. Recall that the carrier $\mu(L(-) + U_L A)$ of $F_L U_L A$ is the colimit of the initial sequence (in **BA**):

$$0 \xrightarrow{!} L(0) + U_L A \xrightarrow{L! + \text{Id}_{U_L A}} L(L(0) + U_L A) + U_L A \rightarrow \dots$$

It is easy to see that $U_L A$ is a cocone for this sequence and there must therefore exist a unique morphism

$$\epsilon_{U_L A}^L : \mu L(-) + U_L A \rightarrow U_L A$$

Our choice of notation is not innocent, since $\epsilon_{U_L A}$ is indeed the counit of $F_L \dashv U_L$ (on carriers). The same construction goes for $\epsilon_{U_L A}^K$. Since the initial sequences for L and K are connected by the natural transformation q we get the following commutative diagram

$$\begin{array}{c}
 \mu(K(-) + U_L A) \\
 \vdots \\
 \epsilon_{U_L A}^K \\
 \vdots \\
 U_K \xi_{U_L A} \quad U_L A \\
 \vdots \\
 \epsilon_{U_L A}^L \\
 \vdots \\
 \mu(L(-) + U_L A)
 \end{array}
 \begin{array}{c}
 \xleftarrow{\quad} 0 \xrightarrow{\quad} K(0) + U_L A \xrightarrow{\quad} K(K(0) + U_L A) + U_L A \xrightarrow{\quad} \dots \\
 \downarrow \text{Id}_0 \quad \downarrow u_0 = q_0 + \text{Id}_{U_L A} \quad \downarrow q_{L(0) + U_L A} \circ K u_0 + \text{Id}_{U_L A} \\
 \xleftarrow{\quad} 0 \xrightarrow{\quad} L(0) + U_L A \xrightarrow{\quad} L(L(0) + U_L A) + U_L A \xrightarrow{\quad} \dots
 \end{array}$$

The result follows by the unicity of the three dashed arrows and the naturality of q . □

Theorem 4.1 shows us how a natural transformation induces a translation between the languages associated with two functors. As an illustration, we show how this works in the case of the nabla flavour of coalgebraic logic. Since the nabla languages are defined by **Set**-endofunctors, let $S, T : \mathbf{Set} \rightarrow \mathbf{Set}$ be two weak-pullback preserving functors (note that **Set** functors always preserve regular epis) and let $q : S \rightarrow T$ be any natural transformation. We can read off the effect of the translation transformation

$$\xi_{FV} : \mu(FSU(-) + FV) \rightarrow \mu(FTU(-) + FV)$$

from Diagram 4.1 of Theorem 4.1. By defining $(-)^q = U_{FSU} \xi_{FV}$ as above, we get the recursive definition for every $\alpha \in \mu(FSU(-) + FV)$:

$$(\nabla \alpha)^q = \nabla(q_{\mu(FTU(-) + FV)} \circ S(-)^q)(\alpha)$$

If we use the nabla languages $\mathcal{L}_S$ and $\mathcal{L}_T$ (see Section 1.5.1) rather than the abstract language associated with FSU or FTU , we can rewrite this translation map a bit more pithily for every $\alpha \in \mathcal{L}_S$ as

$$(\nabla \alpha)^q = \nabla(q_{\mathcal{L}_T} \circ S(-)^q)(\alpha)$$

4.2 Models and semantics translations

This section starts with the dual counterpart of the previous section. We will then examine how the connections between syntax and models, i.e. the semantics, interacts with natural transformations, and in particular functor presentations. We end with an application to the case of the nabla flavour of coalgebraic logics.

4.2.1 Semantic translations

Since models are described in exactly the same way for the predicate-lifting, nabla and abstract descriptions of coalgebraic logic, i.e. by coalgebras, we just need to look at the impact of natural transformations on coalgebras. Note first that a natural transformation $q : S \rightarrow T$ between two **Set**-endofunctors induces a functor $Q : \mathbf{CoAlg}(S) \rightarrow \mathbf{CoAlg}(T)$ on the corresponding categories of coalgebras, given for any (W, γ) in $\mathbf{CoAlg}(S)$ and any $f : (W, \gamma) \rightarrow (W', \delta)$ (i.e. $Tf \circ \gamma = \delta \circ f$) by

$$Q(W, \gamma) = (W, q_W \circ \gamma), \quad Qf = f$$

and Qf is a coalgebra morphism by naturality of q . We will now show how we can use Q to define a translation from S -behaviours to T -behaviours, in a way that is dual to the syntactic translation of Theorem 4.1. First recall the classic result.

Theorem 4.3 (Adámek). *Let $T : \mathbf{C} \rightarrow \mathbf{C}$ be an endofunctor on a category with a terminal object 1. Assume further, that the limit of the following diagram exists in $\mathbf{C}$*

$$\dots \longrightarrow T^3 1 \xrightarrow{T^2!} T^2 1 \xrightarrow{T!} T 1 \xrightarrow{!} 1$$

and that T preserves it, then T has a terminal coalgebra.

Note that if $\mathbf{C}$ is locally presentable, then it is complete and cocomplete and in particular $\mathbf{C}$ has a terminal object and the limit above always exists, so the only requirement left is that T should preserve it. This is always the case for polynomial functors:

Proposition 4.4. *Coproducts of $\text{hom} : \mathbf{C} \rightarrow \mathbf{Set}$ functors preserve ω^{op} -limits.*

Proof. This is a consequence of the fact that hom functors preserve limits and that small coproducts in **Set** commute with connected limits, and in particular with ω^{op} -limits. $\square$

In particular polynomial functors always have terminal coalgebras. So if T is a finitary **Set**-endofunctor and S is a polynomial functor arising from a

presentation of T , i.e. $q : S \twoheadrightarrow T$, then we can always build the terminal coalgebra of S . Intuitively, the terminal coalgebra of T should be a quotient of the terminal coalgebra for S , since the latter is a more fine-grained specification of behaviours. This is indeed the case and is at the core of the method developed by Adámek and Milius in [AM06] to compute terminal coalgebras for functors which do not preserve the ω^{op} -limit of Theorem 4.3. The procedure is in two steps. First, they show that if νS is the terminal coalgebra of S , then the coalgebra

$$\nu S \longrightarrow S\nu S \xrightarrow{q\nu S} T\nu S$$

is a *weakly* terminal T -coalgebra. They then quotient this coalgebra to turn it into the terminal T -coalgebra by identifying behaviours that T cannot tell apart. In this way it can be shown that every finitary functor on **Set** has a terminal coalgebra. In fact it is the case that every accessible functor on a locally finitely presentable category has a terminal coalgebra (see Theorem 4.2.12. [AMM10]). With this fact established, we can easily quotient S -behaviours onto T -behaviours as follows.

Theorem 4.5. *Let $S, T : \mathbf{Set} \rightarrow \mathbf{Set}$ be two finitary functors, and let $q : S \twoheadrightarrow T$ be an epi-transformation. Let $C_S : \mathbf{Set} \rightarrow \mathbf{CoAlg}(S)$ denote the obvious cofree functor associated with S , and similarly with C_T for T . There exist a natural transformation*

$$\chi : C_S \rightarrow C_T$$

which is an epi-transformation.

Proof. The existence proof is dual to that of Theorem 4.1. Let V be any set, by using the fact that $C_T(V)$ is the terminal $T(-) \times V$ -coalgebra, we get:

$$\begin{array}{ccc} TC_S V \times V & \xrightarrow{T\chi_V \times \text{Id}_V} & TC_T V \times V \\ \uparrow q_{C_S V} & & \uparrow \\ SC_S V \times V & & \\ \uparrow & & \\ C_S V & \xrightarrow{\chi_V} & C_T V \end{array} \quad (4.3)$$

To prove that χ_V is surjective when q is an epi-transformation, we use the fact, in **Set**, every epimorphism has a section. In particular, there exist a section

$$s : TC_T V \rightarrow SC_T V, \text{ such that } q_{C_T V} \circ s = \text{Id}_{TC_T V}$$

Using s we can build for every set V the following commutative diagram (where

we drop all the $- \times V$ terms and $- \times \text{Id}_V$ maps for legibility):

$$\begin{array}{ccc}
 T\mathbf{C}_TV & \xleftarrow{T\chi_V} & T\mathbf{C}_SV \\
 \uparrow q_{\mathbf{C}_TV} & \xrightarrow{T\sigma} & \uparrow q_{\mathbf{C}_SV} \\
 \text{Id}_{T\mathbf{C}_TV} \curvearrowleft & & \\
 SC_TV & \xleftarrow{S\chi_V} & SC_SV \\
 \uparrow s & \xrightarrow{S\sigma} & \uparrow \\
 TC_TV & & \\
 \uparrow & & \uparrow \\
 \mathbf{C}_TV & \xleftarrow{\chi_V} & \mathbf{C}_SV \\
 & \xrightarrow{\sigma} &
 \end{array}$$

Where σ exists by terminality of $\mathbf{C}_SV$ and χ exists by terminality of $\mathbf{C}_TV$. It follows that $\chi_V \circ \sigma = \text{Id}_{\mathbf{C}_TV}$, i.e. χ_V is a split epi, and in particular a surjection. $\square$

If there is no ambiguity about the set V of colours, we will denote the **Set**-morphisms $\chi_V : \mathbf{C}_TV \rightarrow \mathbf{C}_SV$ by $(-)_q$, in duality to the syntax translation map $(-)^q$ defined in the previous section. We will call this map the **semantic translation**.

4.2.2 Compatible semantics

We now know how functor presentations (or more generally natural transformations) give rise to translations between languages and between models. The next natural question is to investigate how presentations interact with the relationship between language and model, i.e. the semantics.

Recall that for the abstract description of coalgebraic logic given by the diagram

$$\begin{array}{ccc}
 L & & T^{\text{op}} \\
 \downarrow & \xrightarrow{\text{Uf}} & \downarrow \\
 \mathbf{BA} & & \mathbf{Set}^{\text{op}} \\
 & \xleftarrow{\mathcal{P}} &
 \end{array}$$

the semantics is given by a natural transformation $\delta : L\mathcal{P} \rightarrow \mathcal{P}T^{\text{op}}$. Now, let $q : K \rightarrow L$ be a transformation arising from a presentation of L (recall that K is then of the form FPU , where P is polynomial in **Set**) and $r : S \rightarrow T$ be a transformation arising from a presentation of T (e.g. the minimal presentation). Since $\mathcal{P}$ is contravariant, we have a canonical choice of natural transformation

$\lambda : K\mathcal{P} \rightarrow \mathcal{P}S$ given by:

$$\begin{array}{ccc} K\mathcal{P} & \xrightarrow{\lambda} & \mathcal{P}S \\ q\mathcal{P} \downarrow & & \uparrow \mathcal{P}r \\ L\mathcal{P} & \xrightarrow{\delta} & \mathcal{P}T \end{array} \quad (4.4)$$

which is clearly natural. However, λ defined in this way has some unpleasant properties. As was mentioned in Chapter 1, the properties of δ (or its adjoint transpose) reflect important properties of the logic described by the diagram above, for example expressivity and completeness. If we choose λ as above, then λ will be neither mono (since $q\mathcal{P}$ is epi) nor epi (since $\mathcal{P}r$ generally isn't). So no matter how well-behaved δ is, it is all lost by building λ in this way. As will soon become clear, we need a well behaved logic at the polynomial level, in particular we will need completeness in order to build models. We therefore need to proceed differently, and the case of the nabla logic will guide us in finding a good notion of 'lifted semantics' for our purpose. Recall that in the case of the nabla style of coalgebraic logic, the abstract semantics takes the form:

$$\delta : FTUP \rightarrow \mathcal{P}T$$

which is freely generated by the (adjoint) maps

$$\tilde{\delta}_X : TUPX \rightarrow UPTX, \quad A \mapsto \{\alpha \in TX \mid \alpha \tilde{T} \in_X A\}$$

Clearly, this assignment is fully generic in the functor T , and we can similarly define for a polynomial functor S such that $q : S \twoheadrightarrow T$, the natural transformation

$$\lambda : FSUP \rightarrow \mathcal{P}S$$

which is freely generated by the (adjoint) maps

$$\tilde{\lambda}_X : SUPX \rightarrow UPSX, \quad A \mapsto \{\alpha \in SX \mid \alpha \tilde{S} \in_X A\}$$

The crucial observation to be made about this choice of semantics is that if $\alpha \tilde{S} \in_X A$, then there must exist a witness $w \in S(\in_X)$ which projects to α and A . By naturality of q we therefore get a witness $q_{\in_X}(w) \in T(\in_X)$ which projects to $q_X(\alpha)$ and $q_{UPX}(A)$. This means that we have the following weaker condition than Diagram (4.4):

$$\hat{\lambda}_X(A) \subseteq \mathcal{P}q_{UPX} \circ \hat{\delta}_X \circ q_{UPX}(A)$$

This relation can be understood as the statement that the λ -interpretation of A agrees with the δ -interpretation of $q_{UPX}(A)$. We formalize and generalize this notion from the nabla style of coalgebraic logic in the following definition.

Definition 4.6 (Compatible semantics). Let $(K, S, \lambda : K\mathcal{P} \rightarrow \mathcal{P}S)$ and $(L, T, \delta : L\mathcal{P} \rightarrow \mathcal{P}T)$ define two abstract coalgebraic logics and their semantics, and let $q : K \rightarrow L$ and $r : S \rightarrow T$ be two natural transformations. We will say that λ and δ are **compatible semantics** w.r.t. to q, r whenever for every set X , if we view $K\mathcal{P}X$ and $\mathcal{P}SX$ as skeletal categories, then there exists a transformation

$$\mu_X : \lambda_X \rightarrow \mathcal{P}r_X \circ \delta_X \circ q_{\mathcal{P}X}$$

Or using ‘evil’ (i.e. mentioning elements) terminology, for every $a \in K\mathcal{P}X$

$$\lambda_X(a) \subseteq \mathcal{P}r_X \circ \delta_X \circ q_{\mathcal{P}X}(a)$$

Let us now show that the term ‘compatible’ is indeed appropriate.

Theorem 4.7. *Let $(K, S, \lambda : K\mathcal{P} \rightarrow \mathcal{P}S)$ and $(L, T, \delta : L\mathcal{P} \rightarrow \mathcal{P}T)$ define two abstract coalgebraic logics and their semantics, and let $q : K \rightarrow L$ and $r : S \rightarrow T$ be two natural transformations making λ compatible with δ . Let also V be a set (of propositional variables) and let $a \in F_K V$, if for some coalgebra $W \xrightarrow{\gamma} TW$, some valuation $v : V \rightarrow \mathcal{Q}W$ and some $w \in W$*

$$w, \gamma, v \models a$$

then

$$w, r_W \circ \gamma, v \models (a)^q$$

Moreover, if $w \in C_S \mathcal{Q}V$ and


$$w \models a$$

then

$$(w)_q \models (a)^q$$

Proof. The interpretation maps are defined by initiality via the following commutative diagram (where we have dropped all the $(-)+FV$ and $(-)+Id_V$ terms

for clarity's sake)



Starting from $a = \langle \alpha \rangle_K, \alpha \in KF_KV$, we know from the compatibility assumption that


$$\lambda_W(K[-]_\gamma^K(\alpha)) \subseteq Pr_W(\delta_W(q_{PW}(K[-]_\gamma^K(\alpha))))$$

from which it follows by commutativity of the diagram that

$$[[a]]_\gamma^K \subseteq [(a)^q]_{r_W \circ \gamma}^L$$

which is what we needed to show. In terms of the terminal coalgebra semantics,

we adapt the diagram above as follows



Starting from $a = \langle \alpha \rangle_K, \alpha \in KF_K V$, the compatibility assumption once again entails that

$$\lambda_{c_S QV}(K[-]^K(\alpha)) \subseteq pr_{c_S QV}(\delta_{c_S QV}(q_{PC_S QV}(K[-]^K(\alpha))))$$

from which it follows by commutativity of the diagram that

$$[[a]]^K \subseteq \mathcal{P}(-)_q([[(a)^q]]^L)$$

i.e.

$$\text{if } w \models a \text{ then } (w)_q \models (a)^q$$

as desired. $\square$

4.2.3 The case of nabla logics

We illustrate the notion of compatible semantics with the case that inspired its definition, namely the nabla flavour of coalgebraic logic. However, we will adopt an abstract perspective on nabla logic as follows. Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be a weak pullback preserving finitary functor, we saw in Section 1.5 that we can view the nabla language modulo boolean equivalence as the free L -algebra for $L = FTU$, that is we can use L to define a simple logic where the only inference rules are those of propositional calculus. Here we add more structure to L by defining for any finitary $\mathbf{Set}$ -functor T , any boolean algebra A and morphism $f : A \rightarrow B$

$$\begin{cases} L^T A = F(\{\nabla \alpha \mid \alpha \in TUA\}) / \simeq_{\text{KKV}(T)} \\ L^T f : L^T A \rightarrow L^T B, [\nabla \alpha] \mapsto [\nabla T f \alpha] \end{cases}$$

where the quotient is taken under provability in $\mathbf{KKV}(T)$, i.e. $a \simeq_{\mathbf{KKV}(T)} b$ iff $\mathbf{KKV}(T) \vdash a \leq b$ and $\mathbf{KKV}(T) \vdash b \leq a$.

We now define the semantic transformation $\delta : L\mathcal{P} \rightarrow \mathcal{P}T$. The kind of definition of L given above will occur repeatedly in all the Examples we will examine, particularly in Chapter 5. In every case L will be of the form $L = (\mathbf{FTU})/\simeq$, so it is worth detailing exactly how we define a semantic natural transformation on such an L . The procedure will always be the same: for every set X , we will define the map $\hat{\delta}_X : T\mathcal{UP}X \rightarrow \mathcal{UP}TX$, which uniquely determines a **BA**-morphism $\delta_X : \mathbf{FTUP}X \rightarrow \mathcal{P}TX$ (via the adjunction). If δ_X can be shown to respect the equivalence relation, i.e. if $\delta_X(a) = \delta_X(b)$ whenever $a \simeq b$, then δ_X defines a **BA**-morphism $(\mathbf{FTUP}X)/\simeq \rightarrow \mathcal{P}TX$. We usually commit the abuse of terminology consisting in saying that we define δ_X on generators, rather than saying we define its adjoint transpose. For notational clarity we will also drop the $[-]$ brackets and simply remember that elements of LA are equivalence classes, of which we pick representatives.

In the case at hand we proceed as follows: we define δ_X on generators $\nabla\alpha, \alpha \in T\mathcal{UP}X$ by

$$\delta_X(\nabla\alpha) = \{t \in TX \mid t \bar{T} \in_X \alpha\}$$

Lemma 4.8. δ_X respects the equivalence relation $\simeq_{\mathbf{KKV}(T)}$.

Proof. We proceed by induction on the depth of the proof of $\nabla\Phi \leq \nabla\Psi$. The base case is trivial: if $\nabla\Phi \leq \nabla\Phi$, then $\delta_X(\nabla\Phi) \subseteq \delta_X(\nabla\Phi)$. For the inductive case we need only check the case where the last step of the proof is $(\nabla 1), (\nabla 2), (\nabla 3)$, the propositional rules being automatically taken care of by being in **BA**.

($\nabla 1$) Assume that the last applied rule was

$$(\nabla 1) \frac{\{U \subseteq V \mid (U, V) \in R\}}{\nabla\alpha \leq \nabla\beta} (\alpha, \beta) \in \bar{T}R$$

From the premise we can conclude that $R \subseteq \subseteq_{\mathcal{P}X}$, and thus $\bar{T}R \subseteq (\bar{T} \subseteq_{\mathcal{P}X})$. Moreover, it is clear that $(\in_X; \subseteq_{\mathcal{P}X}) \subseteq \in_X$, and thus since T preserves weak pullbacks we get $(\bar{T} \in_X; \bar{T} \subseteq_X) \subseteq \bar{T} \in_X$. We can then conclude that

$$\begin{aligned} t \bar{T} \in_X \alpha \bar{T} R \beta &\Rightarrow t \bar{T} \in_X \alpha \bar{T} \subseteq_X \beta \\ &\Rightarrow t \bar{T} \in_X \beta \end{aligned}$$

i.e. $\delta_X(\nabla\alpha) \subseteq \delta_X(\nabla\beta)$ as desired.

($\nabla 2$) Assume that the last applied rule was

$$(\nabla 2) \frac{\{\nabla(T \cap)\Phi \leq b \mid \Phi \in \mathbf{SRD}(A)\}}{\bigcap \{\nabla\alpha \mid \alpha \in A\} \leq b}$$

We need to show that if for all $\alpha \in A$, $t\bar{T} \in_X \alpha$, then there exists a $\Phi \in SRD(A)$ such that $t\bar{T} \in_X (T \cap) \Phi$, and we can then use the Induction Hypothesis on shorter proofs. We proceed as follows: A is a finite set, say of cardinality n , i.e. $A = \{\alpha_1, \dots, \alpha_n\}$, and we can therefore consider the kernel n -tuple $\ker_n(\pi_1)$ of the map $\pi_1 : \in_X \rightarrow X$, i.e. the pullback of n copies of π_1 . Intuitively an element of $\ker_n(\pi_1)$ is an n -tuple of subsets of X which all contain a common element. Since T weakly preserves pullbacks, it weakly preserves kernel n -tuples, and the data of $t\bar{T} \in_X \alpha$ for all $\alpha \in A$ defines an element of $x \in T\ker_n(\pi_1)$.

Next, note that $\ker_n$ defines a relation K_n on $X \times \mathcal{P}_\omega \mathcal{P}X$ defined by

$$xK_n\mathcal{V} \text{ if } |\mathcal{V}| = n \text{ and } x \in U \text{ for all } U \in \mathcal{V}$$

and that there exist a map $h : \ker_n(\pi_1) \rightarrow K_n$ defined by

$$((x, U_1), \dots, (x, U_n)) \mapsto (x, \{U_1, \dots, U_n\})$$

If we denote by $p_i, 1 \leq i \leq n$ the maps from the kernel n -tuple to the various copies of $\in_X$, it is clear that by construction

$$p_i; \pi_2; \in_{\mathcal{P}X} = h; \pi_2$$

i.e. the $\mathcal{P}X$ -component of the i^{th} projection of an n -tuple belongs to a finite collection of subsets iff this finite collection of subsets is the one built by h . Note also that $K_n \subseteq (\in_X; \cap)$, and that since T preserves weak pullbacks we therefore have $\bar{T}K_n \subseteq \bar{T}\in_X; T \cap$.

We can now gather our results and see that (1) by the action of the map h , our witness x gives us an element $Th(x) \in TK_n$, (2) this element also belongs to $\bar{T}\in_X; T \cap$, i.e. $t\bar{T} \in_X T \cap \Phi$ for some $\Phi \in T\mathcal{P}_\omega \mathcal{P}X$, and finally (3) this Φ is in $SRD(A)$ since $T\pi_2 \circ T\pi_i(x) = \alpha_i$ and $T\pi_2 \circ Th(x) = \Phi$, thus $\alpha_i \bar{T} \in_{\mathcal{P}X} \Phi$ for $1 \leq i \leq n$ as desired.

($\nabla 3$) Assume that the last applied rule was

$$(\nabla 3) \frac{\{\nabla \alpha \leq b \mid \alpha \bar{T} \in_X \Phi\}}{\nabla(T \cup) \Phi \leq b}$$

We need to show that if $t\bar{T} \in_X (T \cup) \Phi$ then there exists $\alpha \bar{T} \in_{\mathcal{P}X} \Phi$ such that $t\bar{T} \in_X \alpha$, and we can then use the I.H. on a shorter proof. To show this we first notice that if $x \in \bigcup_i U_i$ for a certain collection $(U_i)_{i \in I}$ of subsets of X , then x must belong to at least one U_i . We choose a selection map for any such pair $(x, \bigcup_i U_i)$, i.e. we define a map $h : (\in_X; \bigcup) \rightarrow (\in_X)$ which chooses a subset to which a given element of a union of subsets belongs.

Note that this does not require the axiom of choice since all unions are over finite sets. This map has the pleasant property that $\pi_1 \circ h = \pi_1$ (i.e. we keep the element), and that $h; \pi_2; \in_{\mathcal{P}X} = \pi_2$ (i.e. the subset we choose is an element of the collection of subset we started off with). From this it follows that since T weakly preserve pullbacks, we can send the witness w of $t\bar{T} \in_X (T \bigcup) \Phi$ to a witness $Th(w) \in T \in_X$ of $t\bar{T} \in_X \alpha$ where $\alpha = T\pi_2 Th(w)$ is such that $\alpha \bar{T} \in_{\mathcal{P}X} \Phi$ as desired.

□

Having shown that δ_X defines an appropriate semantic transformation for L^T we conclude by showing that for every epi-transformation r between two weak pullback preserving functors S, T if $\lambda : L^S \rightarrow \mathcal{P}S$ and $\delta : L^T \rightarrow \mathcal{P}T$ are the semantic transformations we have just defined, then there exists a natural transformation $q : L^S \rightarrow L^T$ making the semantics compatible in the sense of definition 4.6. This transformation is uniquely determined by $r : S \rightarrow T$, but is a little messy to formulate.

We temporarily define $q : L^S \rightarrow L^T$ by its action on generators

$$q_A : (\mathbf{FSUA}) / \simeq_{\mathbf{KKV}(S)} \rightarrow (\mathbf{FTUA}) / \simeq_{\mathbf{KKV}(T)}, \nabla \alpha \mapsto \nabla(r_{\mathbf{UA}}(\alpha))$$

For q_A to be well-defined in must be compatible with the $\mathbf{KKV}(T)$ axiomatization encoded in the target functor L^T . Specifically, for any $a, b \in L^S A$ if $\mathbf{KKV}(S) \vdash a \leq b$, then we need $\mathbf{KKV}(T) \vdash q_A(a) \leq q_A(b)$. As it turns out q defined as above is *not* well-defined. Let us try to show why this is, and how we can fix the problem. We proceed by induction on depth of the proof of $\mathbf{KKV}(S) \vdash a \leq b$, as in the Lemma above, then the base case is again trivial. For the inductive step let us first consider the $(\nabla 1)$ rule i.e.

$$(\nabla 1) \frac{\{a \leq b \mid (a, b) \in R\}}{\nabla \alpha \leq \nabla \beta} (\alpha, \beta) \in \bar{S}R$$

By definition, if $(\alpha, \beta) \in \bar{S}R$, then there exists a witness $w \in SR$ which we can turn into a witness $r_R(w) \in TW$, from which it follows immediately that $\mathbf{KKV}(T) \vdash \nabla(\mathbf{F}r_{\mathbf{UA}}(\alpha)) \leq \nabla(\mathbf{F}r_{\mathbf{UA}}(\beta))$ by the induction hypothesis and an application of the $(\nabla 1)$ rule.

However, the two other rules do not behave so well, in particular the induction hypothesis does *not* allow us to turn $\mathbf{KKV}(S)$ -proof into a $\mathbf{KKV}(T)$ -proof. Let us see what happens with the $(\nabla 2)$ rule. Assume that we have a proof $\mathbf{KKV}(S) \vdash a \leq b$ whose last step is

$$\nabla 2 \frac{\{\nabla(S \wedge) \Phi \leq b \mid \Phi \in \mathbf{SRD}(U)\}}{\bigwedge \{\nabla \alpha \mid \alpha \in U\} \leq b}$$

for some $U \in \mathcal{P}_\omega \text{SUA}$. If we wanted to show $\bigwedge \{\nabla \alpha \mid \alpha \in q_A[A]\}$ (i.e. the image of the conclusion under the transformation q_A defined above) by using the $(\nabla 2)$ rule of $\text{KKV}(T)$, then the following lemma shows us what would be needed in the premise.

Lemma 4.9. *Let S, T be two weak-pullback preserving functors on $\mathbf{Set}$, let $q : S \rightarrow T$ be a epi natural transformation. For any boolean algebra A and $U \in \mathcal{P}_\omega \text{SUA}$, the two following conditions are equivalent*

- (1) $\Phi \in \text{SRD}(q_{\text{UA}}[U])$
- (2) *there exist $\Phi' \in \text{SP}_\omega \text{UA}$ such that $q_{\mathcal{P}_\omega \text{UA}} \Phi' = \Phi$ and for all $\alpha \in U$ there exist α' such that $\alpha' \bar{S} \in_{\text{UA}} \Phi'$ and $q_{\text{UA}}(\alpha') = q_{\text{UA}}(\alpha)$*

Proof. For notational convenience we will use the following abbreviations:

- $(-)^q : \text{SUA} \rightarrow \text{SUA}, \alpha \mapsto q_{\text{UA}}(\alpha)$
- $[-]^q : \mathcal{P}_\omega \text{SUA} \rightarrow \mathcal{P}_\omega \text{TUA}, U \mapsto q_{\text{UA}}[U]$
- $\{-\}^q : \text{SP}_\omega A \rightarrow \text{TP}_\omega \text{UA}, \Phi \mapsto q_{\mathcal{P}_\omega \text{UA}}(\Phi)$

From (1) to (2). By definition of redistributions, we know that all elements α of $[U]^q$ are T -lifted members of Φ , i.e. we can always find a witness w in $T \in_{\text{UA}}$ such that $T\pi_1(w) = \alpha$ and $T\pi_2(w) = \Phi$ where $\pi_1 : \in_{\text{UA}} \rightarrow \text{UA}, \pi_2 : \in_{\text{UA}} \rightarrow \mathcal{P}_\omega \text{UA}$. Since $[U]^q$ is finite, we can encode this property categorically by considering the kernel n -tuple of $T\pi_2$ where $n = |[U]^q|$, i.e. the limit of diagram

$$\begin{array}{ccccc}
 T \in_{\text{UA}} & & T \in_{\text{UA}} & & \dots & & T \in_{\text{UA}} & & T \in_{\text{UA}} \\
 & \searrow & \downarrow T\pi_2 & & \downarrow T\pi_2 & & \downarrow T\pi_2 & & \downarrow T\pi_2 \\
 & & & & T\pi_2 & & & & \\
 & & & & T\mathcal{P}_\omega \text{UA} & & & &
 \end{array}$$

with one arrow $T\pi_2$ for each element $\alpha \in [U]^q$. Let us denote the limit of this diagram by $\ker_n(T\pi_2)$. This set contains an n -tuple of witnesses each of whom witnesses a relation $\alpha \bar{T} \in_{\text{UA}} \Phi$ for all $\alpha \in [U]^q$.

Since T preserves weak pullbacks, it also weakly preserves kernel pairs. Moreover, it is not difficult to see that the weak preservation of kernel pairs implies the weak preservation of kernel n -tuples when n is finite. We therefore have that the set of witnesses w_α of each relation $\alpha \bar{T} \in \Phi$ determines (non-uniquely) a ‘global witness’ $w_\Phi \in T \ker_n(\pi_2)$. Since we have an epi natural transformation $q : S \rightarrow T$, we have a surjective map

$$q_{\ker_n(\pi_2)} : S \ker_n(\pi_2) \rightarrow T \ker_n(\pi_2)$$

and in particular our ‘global’ witness w_Φ has a pre-image $w'_\Phi \in S \ker_n(\pi_2)$ under this map. It is straightforward to check that by naturality of q , this w'_Φ is an

n -tuple of witnesses, all of whom witness a relation $\alpha' \bar{S}_{\in \mathbf{U}A} \Phi'$ for a common element Φ' such that $\{\Phi'\}^q = \Phi$. Moreover, each α' gets mapped to a different element $(\alpha')^q$ of $[U]^q$.

From (2) to (1). If for any $\alpha \in U$ we have an α' such that $(\alpha)^q = (\alpha')^q$ and $\alpha' \bar{S}_{\in \mathbf{U}A} \Phi'$ then there must exist a witness $w_{\alpha'} \in S_{\in \mathbf{U}A}$ such that $S\pi_1(w_{\alpha'}) = \alpha', S\pi_2(w_{\alpha'}) = \Phi'$. Using the map

$$q_{S \in \mathbf{U}A} : S \in \mathbf{U}A \rightarrow T \in \mathbf{U}A$$

we can turn $w_{\alpha'}$ into an element $w_{(\alpha')^q} \in T \in \mathbf{U}A$ witnessing $(\alpha')^q \bar{T} \in \mathbf{U}A \Phi$ by naturality of q . Since this holds for any $\alpha \in U$ and since $(\alpha)^q = (\alpha')^q$ we clearly have that every element of $[U]^q$ is a T -lifted member of Φ , i.e. $\Phi \in \text{SRD}([U]^q)$. $\square$

Using the result and the notation of the lemma, it is clear that if we wanted to conclude that $\bigwedge \{\nabla \alpha \mid \alpha \in [U]^r\}$ by using the $(\nabla 2)$ rule, we would need proofs of $\{\nabla(S \bigwedge) \Phi \leq b \mid \Phi \in \text{SRD}(U')\}$ for *every* U' such that $[U']^r = [U]^r$. To remedy this problem, we will need q to identify more formulas than those identified by $r_{\mathbf{U}A}$. The easiest solution is to identify formulas $\bigwedge \{\nabla \alpha \mid \alpha \in U\}$ and $\bigwedge \{\nabla \alpha \mid \alpha \in V\}$ whenever $[U]^r = [V]^r$. Since the converse of the $(\nabla 2)$ rule is derivable using the $(\nabla 1)$ rule (see Remark 6.4 of [KKV12b]), identifying all these meets will automatically guarantee that the corresponding slim redistributions will also be identified. Formally we define the set of equations as follows. Let $\mathcal{S}_r(X)$ denote the set of all sections of the epi $r_X : SX \rightarrow TX$, using this notation we define

$$E_A^\wedge = \mathcal{S}_r(\mathbf{U}A))^2 \times \mathcal{P}_\omega T\mathbf{U}A$$

and the maps

$$e_i^\wedge E_A^\wedge \rightarrow \mathbf{U}L^S A, \quad (s_1, s_2, A) \mapsto \bigwedge s_i[A], \quad i = 1, 2$$

this set $E_A^\wedge$ together with the maps $e_1^\wedge, e_2^\wedge$ define the equations under which we will quotient $L^S A$ before applying the transformation $r_{\mathbf{U}A}$ in order to be compatible with $(\nabla 2)$.

Let us now turn our attention to the $(\nabla 3)$ rule by assuming that we have a $\text{KKV}(S)$ proof of $a \leq b$ whose last step is

$$(\nabla 3) \frac{\{\nabla \alpha \leq b \mid \alpha \bar{S}_{\in \mathbf{U}A} \Phi\}}{\nabla(S \bigvee) \Phi \leq b}$$

The following lemma shows what would be needed as a premise of a $(\nabla 3)$ rule to conclude $\nabla(T \bigvee) \{\Phi\}^r \leq b$.

Lemma 4.10. *Let S, T be two weak-pullback preserving functors, let $q : S \rightarrow T$ be an epi natural transformation. For any boolean algebra A and $\Phi \in T\mathcal{P}_\omega \mathbf{U}A$ and using the same conventions as in Lemma 4.9, the following two conditions are equivalent*

(1) $\alpha \bar{T} \in_{\mathcal{U}A} \Phi$

(2) *there exist $\Phi' \in S\mathcal{P}_\omega \mathcal{U}A$ such that $\{\Phi'\}^q = \Phi$ and $\alpha' \in S\mathcal{U}A$ such that $(\alpha')^q = \alpha$ and $\alpha' \bar{S} \in_{\mathcal{U}A} \Phi'$*

Proof. From (1) to (2). Assume $\alpha \bar{T} \in_{\mathcal{U}A} \Phi$, by definition, there must exist a witness $w \in T(\in_{\mathcal{U}A})$ such that $T\pi_1(w) = \alpha, T\pi_2(w) = \Phi$. Since q is epi, there must exist a pre-image $w' \in S(\in_{\mathcal{U}A})$ such that $q_{\in_{\mathcal{U}A}}(w') = w$ which witnesses that $\alpha' \bar{S} \in_{\mathcal{U}A} \Phi'$ for a certain $\alpha' \in S\mathcal{U}A$ and $\Phi' \in S\mathcal{P}_\omega \mathcal{U}A$. By naturality of q we have $(\alpha')^q = \alpha$ and $\{\Phi'\}^q = \Phi$.

From (2) to (1). Immediate by pushing the witness of $\alpha' \bar{S} \in_X \Phi'$ down to $T(\in_{\mathcal{U}A})$ by the maps defined above. $\square$

From this lemma we can see that if we wanted to conclude $\nabla(T \vee)\{\Phi\}^r \leq b$ by using the $(\nabla 3)$ rule, we would need to have a proof of $\nabla\alpha' \leq b$ for every lifted member α' of every Φ' which gets mapped to $\{\Phi\}^q$. Again, we have no choice but to identify more formulas than those identified by $r_{\mathcal{U}A}$. We proceed as in the $(\nabla 2)$ case and identify $\nabla(S \vee)\Phi$ and $\nabla(S \vee)\Psi$ whenever $\{\Phi\}^r = \{\Psi\}^r$. Since the converse of rule $(\nabla 3)$ is derivable using $(\nabla 1)$ (see Remark 6.6 of [KKV12b]) we will then automatically get the relevant premise for the $(\nabla 3)$ rule. We proceed formally as in the case of the $(\nabla 2)$ rule and define for every boolean algebra A the set (of equations)

$$E_A^\vee : \mathcal{S}(\mathcal{P}_\omega \mathcal{U}A)^2 \times T\mathcal{P}_\omega \mathcal{U}A$$

and the maps

$$e_i^\vee : E_A^\vee \rightarrow \mathcal{U}L^S A, \quad (s_1, s_2, \Phi) \mapsto s_i(\Phi), \quad i = 1, 2$$

We can finally define $q : L^S \rightarrow L^T$. We proceed in two steps. Let A be any boolean algebra, we first construct the coequalizer of the pair of morphisms

$$F(E_A^\wedge + E_A^\vee) \xrightarrow[\text{F}(e_2^\wedge + e_2^\vee)]{\text{F}(e_1^\wedge + e_1^\vee)} L^S A \xrightarrow{q_{E_A}} Q_{E_A}$$

The second step is defined on ‘generators’ of Q_{E_A} as

$$[\nabla\alpha] \mapsto [\nabla r_{\mathcal{U}A}(\alpha)]$$

By combining the two steps we get an epi-transformation

$$q_A : L^S A \rightarrow L^T A$$

which is a well-defined boolean algebra morphism. It is easy to see from the definition of the first quotient that $r_{\mathcal{U}A}$ is compatible with it. For example since

we have identified formulas $\nabla(S \vee) \Phi$ and $\nabla(S \vee) \Psi$ precisely when $\{\Phi\}^r = \{\Psi\}^r$, this means that if $\nabla(S \vee) \Phi$ and $\nabla(S \vee) \Psi$ are two representatives of an equivalence class under the equations of $E_A^\wedge + E_A^\vee$, then they are mapped to the same formula $\nabla T \vee \{\Phi\}^r = \nabla T \vee \{\Psi\}^r$. From this perspective, it looks like the quotient under the equations of $E_A^\wedge + E_A^\vee$ is not doing anything since we end up with the same elements as if we had only used $(-)^r$. This is true on formulas, but what the quotient achieves is to ensure that the transformation q_A preserve the $\leq$ relation. Put differently, let Φ, Ψ such that $\{\Phi\}^r = \{\Psi\}^r$ and assume that $\nabla(S \vee) \Phi \leq b$ in $L^S A$ and the proof of this ends in $(\nabla 3)$. For $\Psi \neq \Phi$, there is no reason to have $\nabla(S \vee) \Psi \leq b'$ for a b' such that $(b')^r = (b)^r$, i.e. there is no reason for $(-)^r$ to be compatible with the $\leq$ structure on $L^S A$. By quotienting under $E_A^\wedge + E_A^\vee$ we force $\nabla(S \vee) \Phi$ and $\nabla(S \vee) \Psi$ to be the same object, and in particular we then have $[\nabla \alpha] \leq [\nabla(S \vee) \Phi] = [\nabla(S \vee) \Psi] \leq b$ for any lifted member of Φ or Ψ . In this way we make it possible to prove by induction that if $\nabla(S \vee) \Phi \leq b$ and the proof of this ends in $(\nabla 3)$, then $q_A(\nabla(T \vee) \{\Phi\}) \leq q_A(b^q)$ can be shown in $\text{KKV}(T)$ using a proof ending in $(\nabla 3)$.

Let us finally show that with q thus defined, $\lambda : L^S \mathcal{P} \rightarrow \mathcal{P}S$ and $\delta : L^T \mathcal{P} \rightarrow \mathcal{P}T$ are compatible semantics in the sense of Definition 4.6. We only need to show the inequality defining compatible semantics on generators, i.e. for $\alpha \in \text{SUP}X$ we need to check that

$$\lambda_X(\nabla \alpha) \subseteq \mathcal{P}r_X \circ \delta_X \circ q_{\mathcal{P}X}(\nabla \alpha)$$

But this follows trivially by definition and naturality of r , i.e. if $t \bar{S} \in_X \alpha$, then by pushing the witness from $S \in_X$ to $T \in_X$ with $r_{\in_X}$ we get a witness of $r_X(t) \bar{T} \in_X r_{\text{UP}X}(\alpha)$ as desired.

4.3 Proof-theoretic translations

In this section we will deal with purely syntactic questions, i.e. no models will be involved. We will study the link between provability in a logic for a functor $K : \mathbf{BA} \rightarrow \mathbf{BA}$ and a functor $L : \mathbf{BA} \rightarrow \mathbf{BA}$ when they are related by a regular epi-transformation $q : K \twoheadrightarrow L$.

We will use the notation and constructions from our discussion on the algebraic semantics of coalgebraic logics in Section 1.5.4. Let $K, L : \mathbf{BA} \rightarrow \mathbf{BA}$ be weak-pullback and regular epi preserving finitary functors, in particular K, L are variators. Let V be a set of proposition variables, let $F_L \dashv U_L : \mathbf{Set} \rightarrow \mathbf{Alg}_{\mathbf{BA}}(L)$ be the obvious free and forgetful functors and similarly for $F_K \dashv U_K : \mathbf{Set} \rightarrow \mathbf{Alg}_{\mathbf{BA}}(K)$. By definition, the abstract languages over a set of propositional variables V , associated with L and K , are simply $F_L V$ and $F_K V$. Let also $e_1, e_2 : E \rightrightarrows U_L F_L V$ define a set E of equations in the

abstract coalgebraic language $F_L V$. These equations define a variety $\mathbb{V}_E$ of $L(-) + FV$ -algebras defined by the coequalizer of the adjoint transpose morphisms $\hat{e}_1, \hat{e}_2 : F_L E \rightrightarrows F_L V$. Given a natural transformation $q : K \rightarrow L$, we will now canonically lift these equations to define a variety of K -algebras. We follow the technique developed in Section 1.5.4, i.e. we start by building the smallest equivalence relation generated by the equations defined by e_1, e_2 , i.e. we take the kernel pair of the coequalizer of the adjoint transpose morphisms $\hat{e}_1, \hat{e}_2 : F_L E \rightarrow F_L V$. For notational simplicity, and because this kernel pair can still be thought of as ‘equations’ we will denote the kernel pair as E and the projections to $F_L V$ as e_1, e_2 since these also pick the left and right-hand side of an ‘equation’. We denote the equalizer as $q_E : F_L V \rightarrow Q_E$, and we thus have the following exact sequence:

$$E \begin{array}{c} \xrightarrow{e_1} \\ \xrightarrow{e_2} \end{array} F_L V \xrightarrow{q_E} Q_E \quad (4.5)$$

We now lift this exact sequence to a K -algebra by using the transformation $Q : \mathbf{Alg}_{\mathbf{BA}}(L) \rightarrow \mathbf{Alg}_{\mathbf{BA}}(K)$ defined at the beginning of this chapter as well the following easy Proposition:

Proposition 4.11. *Let K, L be two functors on $\mathbf{BA}$ such that K is a variator preserving regular epis, let $q : K \rightarrow L$ be a natural transformation and let $Q : \mathbf{Alg}_{\mathbf{BA}}(L) \rightarrow \mathbf{Alg}_{\mathbf{BA}}(K)$ be the associated functor, then Q preserves coequalizers.*

Proof. The proof is exactly the same as Proposition 1.19 (ii) and Proposition 1.20, and hinges on K preserving regular epis. $\square$

Proposition 4.12. *Let K, L be two regular epi-preserving functors on $\mathbf{BA}$, let $q : K \rightarrow L$ be a natural transformation and let $Q : \mathbf{Alg}_{\mathbf{BA}}(L) \rightarrow \mathbf{Alg}_{\mathbf{BA}}(K)$ be the associated functor, then Q preserves kernel pairs and exact sequences.*

Proof. Recall from Section 1.4 that a kernel pair $p_1, p_2 : (\ker f, \phi) \rightrightarrows (A, \alpha)$ of an L -algebra morphism $f : (A, \alpha) \rightarrow (B, \beta)$ defines a relation, i.e. a monomorphism $p_1 \times p_2 : (\ker f, \phi) \rightarrow (A, \alpha) \times (A, \alpha)$ in $\mathbf{Alg}_{\mathbf{BA}}(L)$. Let $r_1, r_2 : (C, \gamma) \rightrightarrows Q(A, \alpha)$ be a cone for the diagram $Q(A, \alpha) \xrightarrow{Qf} Q(B, \beta) \xleftarrow{Qf} (A, \alpha)$, we need to show that there exists unique K -algebra morphism $(C, \gamma) \rightarrow Q(\ker f, \delta)$. Since L preserves regular epis, $U_L : \mathbf{Alg}_{\mathbf{BA}}(L) \rightarrow \mathbf{BA}$ is monadic by Proposition 1.9, and thus creates all limits in $\mathbf{Alg}_{\mathbf{BA}}(L)$. In particular, $\ker f$ is the kernel pair of $f : A \rightarrow B$ in $\mathbf{BA}$. There must therefore exist a unique $\mathbf{BA}$ -morphism $u : C \rightarrow \ker f$ such that $p_1 \circ u = r_1$ and $p_2 \circ u = r_2$. Let us show that u defines a K -algebra morphism, i.e. that $u \circ \gamma = \delta \circ Ku$. The cone (C, γ) defines a morphism $r_1 \times r_2 : (C, \gamma) \rightarrow Q((A, \alpha) \times (A, \alpha))$ in $\mathbf{Alg}_{\mathbf{BA}}(K)$, and have the

following commutative diagram:

$$\begin{array}{ccccc}
 & & KC & & \\
 & \swarrow Ku & \downarrow \gamma & \searrow K(r_1 \times r_2) & \\
 K \ker f & \xrightarrow{K(p_1 \times p_2)} & K(A \times A) & & \\
 \downarrow \phi \circ q_{\ker f} & & \downarrow (L\pi_1 \circ \alpha \times L\pi_2 \circ \alpha) \circ q_{A \times A} & & \\
 \ker f & \xrightarrow{(p_1 \times p_2)} & A & & \\
 & \nwarrow u & \nearrow r_1 \times r_2 & & \\
 & & C & &
 \end{array}$$

from which it is clear that

$$\begin{aligned}
 r_1 \times r_2 \circ \gamma &= (p_1 \times p_2) \circ u \circ \gamma \\
 &= (L\pi_1 \circ \alpha \times L\pi_2 \circ \alpha) \circ q_{A \times A} \circ K(r_1 \times r_2) \\
 &= (L\pi_1 \circ \alpha \times L\pi_2 \circ \alpha) \circ q_{A \times A} \circ K(p_1 \times p_2) \circ Ku \\
 &= (p_1 \times p_2) \circ \phi \circ q_{\ker f} \circ Ku
 \end{aligned}$$

and thus $\phi \circ q_{\ker f} \circ Ku = u \circ \gamma$ as required since $p_1 \times p_2$ is a monomorphism. The preservation of exact sequences is now a consequence of Proposition 4.11. $\square$

Thus, as long as K, L preserve regular epis, the functor Q associated with the natural transformation q turns the exact sequence (4.5) in $\mathbf{Alg}_{\mathbf{BA}}(L)$ into an exact sequence

$$QE \xrightarrow[\text{Q}e_2]{\text{Q}e_1} QF_L V \xrightarrow{\text{Q}q_E} QQ_E$$

in $\mathbf{Alg}_{\mathbf{BA}}(K)$. We are now ready to define the lifting of the equations defined by E, e_1, e_2 into the syntax functor defined by K : we take the pullback of $e_1 \times e_2$ along the translation transformation $\xi_V : F_K V \rightarrow QF_L V$ and we call this K -algebra of equations E^* :

$$\begin{array}{ccc}
 E^* & \xrightarrow{e_1^* \times e_2^*} & F_K V \times F_K V \\
 \downarrow r & & \downarrow \xi_V \times \xi_V \\
 QE & \xrightarrow{\text{Q}e_1 \times \text{Q}e_2} & QF_L V \times QF_L V
 \end{array} \tag{4.6}$$

Note that pulling back a mono always yields a mono, and thus $e_1^* \times e_2^*$ is a mono, and as was shown in Section 1.4 this mono defines a pair of jointly monic morphisms $e_1^*, e_2^* : E^* \rightarrow F_K V$, which is what we expect from a pair of morphisms defining ‘equations’. Moreover since we are assuming K to be a variator

preserving regular epis, $\mathbf{Alg}_{\mathbf{BA}}(K)$ is a regular category by Proposition 1.19. By Theorem 4.1, ξ_V is a regular epi, and since pulling back a regular epi gives a regular epi in a regular category, r must be a regular epi.

We will now show that the maps e_1^*, e_2^* form the kernel pair of their coequalizer, this will in particular show that E^* is an equivalence relation, and intuitively shows that the procedure above gives us the ‘tightest’ lifting of the equations of E into the K -language, i.e. that we in fact have lifted the entire exact sequence of Diagram (4.5). So let $q_{E^*} : F_K V \rightarrow Q_{E^*}$ be the coequalizer of e_1^*, e_2^* , and notice first that there exist a unique morphism between Q_{E^*} and Q_{Q_E} , i.e. we have the following commutative diagram:

$$\begin{array}{ccccc}
 E^* & \xrightarrow{e_1^*} & F_K V & \xrightarrow{q_{E^*}} & Q_{E^*} \\
 \downarrow r & \searrow e_2^* & \downarrow \xi_V & & \downarrow q_{E^*}^{E^*} \\
 QE & \xrightarrow{Qe_1} & QF_L V & \xrightarrow{Qq_E} & QQ_E \\
 & \searrow Qe_2 & & &
 \end{array} \quad (4.7)$$

The reason is simply that by commutativity of the left-hand side square, $Qq_E \circ \xi_V$ also coequalizes e_1^* and e_2^* . This being established, let $p_1, p_2 : C \rightarrow F_K V$ be a cone for the diagram $F_K V \xrightarrow{q_{E^*}} Q_{E^*} \xleftarrow{q_{E^*}} F_K V$, if E^* is the kernel pair of q_{E^*} , there must exist a unique K -algebra morphism $u : C \rightarrow E^*$ such that $e_1^* \circ u = p_1$ and $e_2^* \circ u = p_2$. By the same argument as the one leading to Diagram (4.7), if q_{E^*} coequalizes p_1, p_2 , then so does $Qq_E \circ \xi_V$, and therefore $C, \xi_V \circ p_1, \xi_V \circ p_2$ is a cone for the diagram $QF_L V \xrightarrow{Qq_E} QQ_E \xleftarrow{Qq_E} QF_L V$, whose limit is QE . There must therefore exist a unique K -algebra morphism $v : C \rightarrow QE$ such that $e_1 \circ v = \xi_V \circ p_1$ and $e_2 \circ v = \xi_V \circ p_2$. Consequently,

$$\begin{aligned}
 e_1 \times e_2 \circ v &= e_1 \circ v \times e_2 \circ v \\
 &= \xi_V \circ p_1 \times \xi_V \circ p_2 \\
 &= \xi_V \times \xi_V \circ p_1 \times p_2
 \end{aligned}$$

and C is a cone for Diagram (4.6) which defines E^* as a pullback, so there must exist a unique morphism $u : C \rightarrow E^*$ such that $e_1^* \times e_2^* \circ u = p_1 \times p_2$, and the result follows.

We have thus shown three things: (1) we can lift equations in the logic defined by the functor L to equations in the logic defined by the functor K in a canonical manner, (2) the two algebras of equations and the two quotients obtained in this way are related by the nice commutative Diagram (4.7), (3) both original algebra of equations E and its lifted counterpart are equivalence relations on their respective free algebras. We can now state and prove the key result of this Section. Recall from Section 1.5.4, that an object A in a category is said to be *orthogonal* to a morphism q (notation $A \perp q$) is for every morphism $f : \text{dom} q \rightarrow A$ there exists a unique morphism $h : \text{cod} q \rightarrow A$ such that $f = h \circ q$.

Theorem 4.13. *For any L -algebra A , $A \perp q_E$ iff $QA \perp q_{E^*}$.*

Proof. Let us first assume that $A \perp q_E$. To show that $QA \perp q_{E^*}$, we take any K -algebra morphism $h : F_K A \rightarrow QA$ and we will show that h must factor through $QF_L V$. We proceed by using the usual adjunctions: $h : F_K V \rightarrow QA$ defines a unique morphism $\hat{h} : V \rightarrow U_K QA = U_L A$ which in turns define a unique morphism $\hat{h} : F_L V \rightarrow A$. Let us now show that $Q\hat{h} \circ \xi_V = h$, we do this by unravelling the construction of the two adjoint transpose maps and we get that $\hat{h} = \epsilon_A^L \circ F_L U_K h \circ F_L \eta_V^K$. For $Q\hat{h}$ we have:

$$\begin{array}{ccccc}
 & & & Q\hat{h} & \\
 & & & \curvearrowright & \\
 QF_L V & \xrightarrow{QF_L \eta_V^K} & QF_L U_K F_K V & \xrightarrow{QF_L U_K h} & QF_L U_L A \\
 \uparrow \xi_V & & \uparrow \xi_{U_K F_K V} & & \uparrow \xi_{U_L A} \\
 F_K V & \xrightarrow{F_K \eta_V^K} & F_K U_K F_K V & \xrightarrow{F_K U_K h} & F_K U_K QA \\
 & & & \curvearrowright h & \\
 & & & & QA
 \end{array}$$

$\begin{array}{l} \nearrow Q\epsilon_A^L \\ \nearrow \epsilon_{QA}^K \end{array}$

All the squares commute by naturality of ξ , so all we need to check is the commutativity $Q\epsilon_A^L \circ \xi_{U_L A} = \epsilon_{QA}^K$, but this precisely what we have shown in Proposition 4.2. Thus we have shown that $h = Q\hat{h} \circ \xi_V$. But since $A \perp q_E$, there must exist a unique morphism $\hat{u} : Q_E \rightarrow A$ such that $\hat{u} \circ q_E = \hat{h}$. If we now define $u = Q\hat{u} \circ q_E^{E^*}$, where $q_E^{E^*} : Q_{E^*} \rightarrow QQ_E$ is the unique morphism define in Diagram (4.7), we get a unique (since $\hat{u}$ and $q_E^{E^*}$ are unique) K -algebra morphism such that

$$u \circ q_{E^*} = Q\hat{u} \circ Qq_E \circ \xi_V = Q\hat{h} \circ \xi_V = h$$

i.e. $QA \perp q_{E^*}$ as desired.

For the converse, we start in a completely analogous fashion. Assume $QA \perp q_{E^*}$, and let $h : F_L V \rightarrow A$. By using the adjunctions, h defines a function $\hat{h} : V \rightarrow U_L A = U_K QA$, which in turns defines a map $\tilde{h} : F_K V \rightarrow QA$. By using the same argument as above and Proposition 4.2, it is easy to show that $\tilde{h} = h \circ \xi_V$. Since we're assuming that $QA \perp q_{E^*}$, there must exist a unique K -algebra morphism $u : Q_{E^*} \rightarrow QA$ such that $u \circ q_{E^*} = \tilde{h}$. By using the notation of Diagram (4.7) we then get that

$$\tilde{h} \circ e_1^* = u \circ q_{E^*} \circ e_1^* = u \circ q_{E^*} \circ e_2^* = \tilde{h} \circ e_2^*$$

From this it follows that

$$Qh \circ \xi_V \circ e_1^* = Qh \circ \xi_V \circ e_2^*$$

and thus

$$Qh \circ Qe_1 \circ r = Qh \circ Qe_2 \circ r$$

i.e. $Qh \circ Qe_1 = Qh \circ Qe_2$ since r is epi. But this means that h coequalizes e_1, e_2 , and since q_E is the coequalizer of e_1, e_2 , there must exist a unique map $v : Q_E \rightarrow A$ such that $v \circ q_E = h$, i.e. $A \perp q_E$ as desired. $\square$

We have thus show that not only can we lift equations from the L -language to the K -language, but that by doing so we define varieties which are in extremely tight correspondence. From our soundness and completeness results for the abstract algebraic semantic of coalgebraic logics developed in Section 1.5.4 (in particular Theorem 1.35), this means that the derivability of a formula a in the abstract Hilbert system associated with the set of equations E , is equivalent to the derivability of some pre-image a' such that $(a')^q = a$ in the abstract Hilbert system associated with E^* . And similarly for equations and equational systems. In particular, if a formula a in the logic defined by L and the axioms in E is consistent, then it must have an inverse image a' under the syntax translation which is consistent in the logic defined by K and the axioms of E^* .

Chapter 5

Completeness-via-canonicity


In this Chapter we will weave together the strands developed in all the previous Chapters. The first section is an overview of known results on weak-completeness (see for example [KKP05, KP10]). The second section will develop the notion of strong completeness via canonicity in the absence of frame conditions (i.e. of axioms) which is key to developing a theory of strong completeness in the presence of canonical axioms. This section uses the work of [KKP05, KR12, SP09], although our presentation is slightly different. Finally the third section collects all the results of the thesis into completeness-via-canonicity theorems for coalgebraic logics both in their abstract and in their nabla formulation.

5.1 Weak completeness

There are two ‘natural’ candidates for building coalgebraic models and proving things like completeness. The first is the terminal coalgebra (if it exists), or finitary version thereof (if it does not). It is the key to understanding weak completeness and is intimately connected with the properties of the semantic transformation δ . However, this type of construction is not amenable to techniques exploiting canonicity, and for this we must use the second natural candidate which is the notion of ‘canonical model’. The idea there is, in a nutshell, to put a coalgebraic structure on the set of ultrafilters of the initial L -algebra defining the logic. It is intimately related to the *inverse* of the transpose of the semantic transformation ($\hat{\delta}$), if it exists. Thus the two techniques work, in a certain sense, in opposite directions.

5.1.1 Completeness theorem

The following theorem captures the essential features of terminal coalgebra semantics and of weak completeness, it is similar in some respects to Theorem 6.15 of [KP10] and many ideas date back to [KKP04]. We place ourselves once again in the fundamental situation of Diagram (1.2), i.e.



Recall that $\mathbf{C}$ is assumed to be a finitary variety. In particular U creates limits and directed colimits.

Theorem 5.1. *In the fundamental situation described above, let $L : \mathbf{C} \rightarrow \mathbf{C}$ be a finitary variator, let $T : \mathbf{D} \rightarrow \mathbf{D}$ have a terminal coalgebra νT , and let $\delta : LG \rightarrow GT$ be a semantic transformation. For any set V , if $L_V = L(-) + FV$ and $T_V = T(-) \times FV$, then:*

- *There exists a unique interpretation morphism $\llbracket - \rrbracket_V : \mu L_V \rightarrow G\nu T_V$.*
- *This map is injective whenever (i) the morphism $0_{\mathbf{C}} \rightarrow G1_{\mathbf{D}}$ is injective, (ii) δ is injective at every object and (iii) L preserves injective maps. In particular any two formulas $a, b \in \mu L_V$ such that $a \neq b$ have different denotation, i.e. L is weakly complete w.r.t. $\mathbf{CoAlg}(T)$*
- *This map is surjective whenever (i) there exist an regular epi $0_{\mathbf{C}} \rightarrow G1_{\mathbf{D}}$ (ii) δ is a regular epi-transformation and (iii) L preserve regular epis.*

Proof. **Existence of $\llbracket - \rrbracket_V$:** The structure map of νT_V is of the form:

$$\gamma \times v : \nu T_V \rightarrow T(\nu T_V) \times FV$$

By the fact that $F \dashv G$, v has an adjoint transpose $\hat{v} : FV \rightarrow G\nu T_V$, and thus the structure map of the cofree coalgebra νT_V determines uniquely the map:

$$G\gamma + \hat{v} : GT\nu T_V + FV \rightarrow G\nu T_V$$

This having been established, we proceed as usual and use the natural transformation δ and the initiality of μL_V to define the interpretation map as the

catamorphism:

$$\begin{array}{ccc}
 L\mu L_V + FV & \xrightarrow{L[-]_V + \text{Id}_{FV}} & LG\nu T_V + FV \\
 \downarrow & & \downarrow \delta_{\nu T_V} + \text{Id}_{FV} \\
 & & GT\nu T_V + FV \\
 & & \downarrow G\gamma + \hat{v} \\
 \mu L_V & \xrightarrow{[-]_V} & G\nu T_V
 \end{array}$$

Note that we do not need to specify any valuation, instead the map $\hat{v}$ is purely endogenous to the cofree coalgebra, in fact it is merely a rewriting of its canonical colouring map.

Construction of $[-]_V$ from the initial sequence. Let us now give an alternative description of $[-]_V$ which will be essential in what follows. We use both the initial sequence of L_V and the terminal sequence of T_V and relate them as follows.

$$\begin{array}{ccccccc}
 & & & \mu L_V & & & \\
 & \nearrow & & \uparrow & \nwarrow & & \\
 0 & \xrightarrow{i} & L_V 0 & \xrightarrow{L_V i} & L_V^2 0 & \xrightarrow{L_V^2 i} & \dots \xrightarrow{} L_V^n 0 \dots \\
 \downarrow [-]_0 & \downarrow [-]_1 = \delta_1 \circ L_V [-]_0 & \downarrow [-]_2 = \delta_{T_V 1} \circ L_V [-]_1 & & & & \downarrow [-]_n = \delta_{T_V^{n-1} 1} \circ L_V [-]_{n-1} \\
 G1 & \xrightarrow{Gt} & GT_V 1 & \xrightarrow{GT_V t} & GT_V^2 1 & \xrightarrow{GT_V^2 t} & \dots \xrightarrow{} GT_V^n 1 \dots \\
 & \searrow & \searrow & \downarrow & \nwarrow & & \nwarrow \\
 & & & G\nu T_V & & &
 \end{array}$$

where 0 is the initial object in $\mathbf{C}$ (which always exists since $\mathbf{C}$ is a finitary variety) and $i : 0 \rightarrow L_V 0$ is the unique morphism arising from initiality of 0 . Dually, 1 is the terminal object in $\mathbf{D}$ and $t : T_V 1 \rightarrow 1$ is the unique morphism arising from the fact that 1 is terminal. Since $G \lim_i T_V^i 1 = G\nu T_V$ forms a cocone for the cochain $G1 \rightarrow GT_V 1 \rightarrow \dots$, it clearly also forms a cocone for the initial sequence $0 \rightarrow L_V 0 \rightarrow \dots$, and thus there must exist a unique morphism

$$[-] : \mu L_V = \text{colim}_i L_V^i 0 \rightarrow G\nu T_V$$

It is quite clear by construction that in fact $[-] = [-]_V$. In order to determine the set theoretic properties of $[-]_V$ (viz. injectivity and surjectivity), we now apply the forgetful functor $U : \mathbf{C} \rightarrow \mathbf{Set}$ to the above diagram. Recall that U creates filtered colimits (and a fortiori ω -cochains) and thus preserves them.

Hence we have that:

$$\mathsf{U}\mu L_V = \mathsf{U} \operatorname{colim}_i L_V^i 0 = \operatorname{colim}_i \mathsf{U} L_V^i 0$$

Moreover, since $UG\nu T_V$ forms a cocone for the cochain $UG1 \rightarrow UGT_V 1 \rightarrow \dots$ - and thus for the cochain $\mathsf{U}0 \rightarrow \mathsf{U}L_V 0 \rightarrow \dots$ - there must exist a unique morphism $\mathsf{U}\mu L_V \rightarrow UG\nu T_V$, which by unicity must be $\mathsf{U}[-]_V$ itself.

Injectivity of $\mathsf{U}[-]_V$. Let $a, b \in \mathsf{U}\mu L_V$ such that $\mathsf{U}[a]_V = \mathsf{U}[b]_V$. By construction of colimits in **Set** as coproducts quotiented by an equivalence relation, we know that there must exist i, j and pre-images $a' \in L_V^i 0, b' \in \mathsf{U}L_V^j 0$ which map to a, b respectively. By setting $k = \max\{i, j\}$ and choosing the appropriate representative in the equivalence class we can in fact choose $a', b' \in \mathsf{U}L_V^k 0$. By the fact that U creates limits and the construction of limits in **Set**, any element of $UG\nu T_V = UG \lim_i T_V^i 1$ is a set of coherent sequences $(u_0, u_1, u_2, \dots) \in \prod_i T_V^i 1$, where coherent means that $u_i = T_V^i t(u_{i+1})$. In particular $X = \mathsf{U}[a]_V = \mathsf{U}[b]_V$ is such a set.

Consider now $X_a = \mathsf{U}[a']_k$ and $X_b = \mathsf{U}[b']_k$, by commutativity of the diagram, they must get mapped to X by the inverse image of the projection map $\pi_k : \lim_i T_V^i 1 \rightarrow T_V^k 1$. In other words a sequence $(x_i) \in X$ ‘goes through’ X_a iff it ‘goes through’ X_b . Thus $X_a = X_b$. It is easy to check by induction that each $\mathsf{U}[-]_i$ is injective: $\mathsf{U}[-]_0$ is injective by assumption and since $\mathsf{U}[-]_{n+1} = \delta_{T_V^n 1} \circ L_V \mathsf{U}[-]_n$, $\mathsf{U}[-]_{n+1}$ is injective since $\delta_{T_V^n 1}$ is injective by assumption and L_V preserves injective maps. Thus $X_a = X_b$ implies that $a' = b'$, which in turns implies that $a = b$.

Surjectivity of $\mathsf{U}[-]_V$ The proof is very similar to the one for injectivity. Since $X \in UG\nu T_V$ can be represented as a set of coherent sequences, we can take its inverse image under the various projections maps and get a coherent sequence of elements $X_i = \pi_i^{-1}(X) \in UGT_V^i 1$, where coherent now means that $GT_V^i(X_i) = (T_V^i)^{-1}(X_i) = X_{i+1}$. Since $\mathsf{U}[-]_0$ is trivially a regular epi, since δ is assumed to be a regular epi, and since L is assumed to preserve them, it is clear that by induction each c_n is a regular epi, and thus each $\mathsf{U}c_n$ is surjective. Any set X_i therefore has a preimage $a_i \in \mathsf{U}L_V^i 0$ which determines an element of $\mathsf{U}\mu L_V$. $\square$

Remark 5.2. The condition that the initial object $0_{\mathbf{C}}$ of $\mathbf{C}$ should map injectively (resp. surjectively) into (resp. onto) $G1_{\mathbf{D}}$ (where $1_{\mathbf{D}}$ is the terminal object of $\mathbf{D}$) may at first seem ad hoc and unnatural. However, in most practical cases the categories $\mathbf{C}$ and $\mathbf{D}$ are connected in a very special way, namely the subcategories $\mathbf{C}_f$ and $\mathbf{D}_f$ of finite objects are dually equivalent and $\mathbf{C} = \operatorname{Ind}(\mathbf{C}_f), \mathbf{D} = \operatorname{Pro}(\mathbf{D}_f)$ (see [KR12]). Since initial and terminal objects tend to be finite, it is actually generally the case that G sends $1_{\mathbf{D}}$ to $0_{\mathbf{C}}$, i.e. we have an isomorphism between $0_{\mathbf{C}}$ and $G1_{\mathbf{D}}$ and both conditions are satisfied.

This is what happens in the case of $\text{Pf} \dashv \mathcal{U} : \mathbf{DL} \rightarrow \mathbf{Pos}^{\text{op}}$ and in the case of $\text{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}^{\text{op}}$.

Remark 5.3. Note that if $\mathbf{C} = \mathbf{BA}$, then any finitary functor preserves injective maps as is shown in Lemma 6.14. of [KP10]. We do not know if this holds for $\mathbf{DL}$, but the proof in [KP10] does not seem to be directly adaptable to $\mathbf{DL}$ as it relies on a very special property of $\mathbf{BA}$.

In the cases of interest to us, i.e. when $\mathbf{C} = \mathbf{DL}$ or $\mathbf{BA}$, we have the following representation of the ‘colours’ FFV of the cofree coalgebra.

Lemma 5.4. *There is a natural isomorphism $\text{PfF} \simeq \mathcal{Q}$, where $\mathcal{Q}$ is the contravariant powerset functor $\mathbf{Set} \rightarrow \mathbf{Pos}^{\text{op}}$. Similarly, $\text{UfF} \simeq \mathcal{Q}$ where $\mathcal{Q}$ is the contravariant powerset functor $\mathbf{Set} \rightarrow \mathbf{Set}^{\text{op}}$.*

Proof. We use the well-known fact that for any A in $\mathbf{DL}$ a prime filter $u \in \text{Pf}A$ is equivalent to a $\mathbf{DL}$ -morphism $\chi_u : A \rightarrow \mathbb{2}$, where $\mathbb{2}$ is the two element distributive lattice. In particular, prime filters $u \in \text{Pf}FV$ are in one-to-one correspondence with $\mathbf{DL}$ -morphisms $FV \rightarrow \mathbb{2}$, which are themselves, by adjunction, in one-to-one correspondence with maps $V \rightarrow \text{U}\mathbb{2} = \mathbb{2}$ where $\mathbb{2}$ is the set with two elements. In other words, prime filters $u \in \text{Pf}FV$ are in one-to-one correspondence with subsets of V . $\square$

The terminal coalgebra semantics in these cases is thus given by the terminal coalgebra for the functor $T(-) \times \mathcal{Q}V$, i.e. ‘ T -behaviours’ are simply coloured by sets of propositional variables. Using terminal coalgebra semantics we can show weak completeness of the abstract Hilbert system developed in Section 1.5.3. We show weak completeness in the case of $\mathbf{C} = \mathbf{DL}$ since the case of $\mathbf{BDL}$ and $\mathbf{BA}$ are clearly special sub-cases.

Corollary 5.5. *Let $L : \mathbf{DL} \rightarrow \mathbf{DL}$ be a variator which preserves injective maps, $T : \mathbf{Pos} \rightarrow \mathbf{Pos}$ and that $\delta : LU \rightarrow UT$ be injective at every stage, then the Hilbert system $\vdash_{\text{ML}}$ defined by L for a set V of propositional variables is weakly complete w.r.t. to $\mathbf{CoAlg}(T)$.*

Proof. Given a formulas $a, b \in \mu(L(-) + FV)$ such that $a \not\vdash_{\text{ML}} b$, we need to find a T -model which satisfies a but not b . Recall from Proposition 1.32 that $a \vdash_{\text{ML}} b$ iff $\vdash_{\text{EL}} a \wedge b = a$. Our starting assumption is thus that $\not\vdash_{\text{EL}} a \wedge b = a$, i.e. $a \wedge b \neq a$ in $\mu(L(-) + FV)$ by Theorem 1.35. We can now use Theorem 5.1 and the assumption that δ is injective to get that the interpretations of $a \wedge b$ and a in the terminal coalgebra do not coincide, i.e.

$$\llbracket a \wedge b \rrbracket_V \neq \llbracket a \rrbracket_V$$

In particular, there exists $t \in \llbracket a \rrbracket_V$ such that $t \notin \llbracket a \wedge b \rrbracket_V$, which by definition of the coalgebraic semantics in the abstract style means exactly that

$$t \models a \text{ and } t \not\models a \wedge b$$

and thus $t \not\models b$ as required. $\square$

If the terminal coalgebra does not exist we have the following result from [KP10], whose proof is very similar the proof of Theorem 5.1, but where the coalgebraic semantics is not given by the terminal coalgebra (which may not exist), but rather built from an arbitrary section $s : 1 \rightarrow T_V 1 \times FV$ of the terminal morphism $T_V 1 \rightarrow 1$ (which is always possible when $\mathbf{D} = \mathbf{Pos}$ or $\mathbf{Set}$ since it simply amounts to choosing an element in $T_V 1 \times FV$).

Theorem 5.6. *In the fundamental situation described above, let $L : \mathbf{C} \rightarrow \mathbf{C}$ be a finitary variator, let $T : \mathbf{D} \rightarrow \mathbf{D}$ be such that $\mathbf{D}$ has a terminal object $1_{\mathbf{D}}$ and that there exists a right-inverse $s : 1_{\mathbf{D}} \rightarrow T1_{\mathbf{D}} \times FV$ of the unique morphism to the terminal object. For any set V , if $L_V = L(-) + FV$ and $T_V = T(-) \times FV$ then:*

- *For any n , there exist a unique morphism $\llbracket - \rrbracket_n : L_V^n(0_{\mathbf{C}}) \rightarrow GT^n 1_{\mathbf{D}}$ interpreting formulas of modal depth at most n*
- *This map is injective if (i) the morphism $0_{\mathbf{C}} \rightarrow G1_{\mathbf{D}}$ is injective (ii) δ is injective, and (iii) L preserves injective maps. In particular for any two formulas $a, b \in \mu L_V$ such that $a \neq b$ there exist an n which provide a, b with different denotations, i.e. L is weakly complete w.r.t. $\mathbf{CoAlg}$.*
- *This map is surjective if (i) there exists a regular epi $0_{\mathbf{C}} \rightarrow G1_{\mathbf{D}}$ (ii) δ is a regular epi-transformation, and (iii) L preserves regular epis.*

Proof. See Theorem 6.15 of [KP10]. $\square$

In many, but not all, practical cases it is possible to prove the injectivity of δ from its injectivity on finite objects. The following result bears some similarity to the construction of Section 6 of [KR12], but the strict connection between L and T is abandoned for the assumption that L be finitary.

Proposition 5.7. *Let the adjunction $F \dashv G : \mathbf{C} \rightarrow \mathbf{D}^{\text{op}}$ be either the adjunction $\text{Pf} \dashv \mathcal{U} : \mathbf{DL} \rightarrow \mathbf{Pos}^{\text{op}}$ or $\text{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}^{\text{op}}$, let $L : \mathbf{C} \rightarrow \mathbf{C}$ be finitary, and let $\delta : LG \rightarrow GT$. If δ_X is injective for any finite $\mathbf{D}$ -object X , then δ is injective.*

Proof. This is an approximation result using the facts that (1) every object of $\mathbf{DL}$ or $\mathbf{BA}$ can be represented as a filtered colimit of finitely presentable objects, (2) that these are all finite, i.e. $\mathbf{DL}$ and $\mathbf{BA}$ are locally finite, and (3) that by Birkhoff's representation theorem every finite distributive lattice can be represented as the downsets of its poset of join irreducibles, and similarly every finite boolean algebra can be represented as the powerset of its atoms. We show the result for the adjunction $\text{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}^{\text{op}}$, the proof is exactly the same in the distributive lattice case.

So let X be any set, $\mathcal{P}X$ is a boolean algebra which can be written as

$$\mathcal{P}X = \operatorname{colim}_j A_j$$

where each A_j is a finitely presentable boolean algebra. Since such algebras are finite and since finite boolean algebras are isomorphic to the powerset algebra on their atoms, we can assume that each A_j is of the form $A_j = \mathcal{P}X_j$ for a finite set X_j , i.e. $\mathcal{P}X = \operatorname{colim}_j \mathcal{P}X_j$. From this, it is intuitively quite clear that $X = \operatorname{colim}_j X_j$. Formally, by the fact that the diagram X_j is a diagram of monos and by the construction of colimits in **Set**, it is clear that the maps $\operatorname{in}_j : \mathcal{P}X_j \rightarrow \mathcal{P}X$ give, by restriction to singletons, a collection of maps $\hat{\operatorname{in}}_j : X_j \rightarrow X$, and X is thus a cocone for the diagram X_j . To see that it is a colimiting cocone, let $(Y, h_j : X_j \rightarrow Y)$ be another cocone on the same diagram, then the direct images $h_j[-] : \mathcal{P}X_j \rightarrow \mathcal{P}Y$ turn $\mathcal{P}Y$ into a cocone (in **Set**!) for the diagram $\mathcal{P}X_j$ whose colimit is $\mathcal{P}X$. There must therefore exist a unique morphism $u : \mathcal{P}X \rightarrow \mathcal{P}Y$ making the obvious diagram commute. Since in_j and $h_j[-]$ send singletons to singletons, we can also restrict u to singletons and get a unique map $\hat{u} : Y \rightarrow X$, proving that X is the claimed colimit. Let us denote by i_j the injection $X_j \rightarrow X$ and by $\tilde{i}_j$ the injection $\mathcal{P}X_j \rightarrow \mathcal{P}X$. It is not difficult to see that $\mathcal{P}i_j \circ \tilde{i}_j = \operatorname{Id}_{\mathcal{P}X_j}$. By naturality of δ we have the following commutative diagrams:

$$\begin{array}{ccc} L\mathcal{P}X & \xrightarrow{\delta_X} & \mathcal{P}TX \\ L\mathcal{P}i_j \downarrow & & \downarrow \mathcal{P}Ti_j \\ L\mathcal{P}X_j & \xrightarrow{\delta_{X_j}} & \mathcal{P}TX_j \end{array}$$

Now let $x, y \in L\mathcal{P}X$ and assume that $\delta_X(x) = \delta_X(y)$. Since L is finitary we have

$$L\mathcal{P}X = L \operatorname{colim}_j \mathcal{P}X_j = \operatorname{colim}_j L\mathcal{P}X_j$$

and by construction of colimits in **Set**, there must exist $x_j \in L\mathcal{P}X_j, y_k \in L\mathcal{P}X_k$ which map to x, y by $L\tilde{i}_j, L\tilde{i}_k$. Since $\mathcal{P}X_j$ is a filtered diagram, we can assume w.l.o.g. that x_j and y_k actually belong to the same set (by pushing them up to a common superset if necessary), i.e. we take $x_j, y_j \in L\mathcal{P}X_j$. We now have

$$\begin{aligned} \mathcal{P}Ti_j \circ \delta_X(x) &= \mathcal{P}Ti_j \circ \delta_X(y) && \Leftrightarrow \\ \mathcal{P}Ti_j \circ \delta_X \circ L\tilde{i}_j(x_j) &= \mathcal{P}Ti_j \circ \delta_X \circ L\tilde{i}_j(y_j) && \Leftrightarrow \\ \delta_{X_j} \circ L\mathcal{P}i_j \circ L\tilde{i}_j(x_j) &= \delta_{X_j} \circ L\mathcal{P}i_j \circ L\tilde{i}_j(y_j) && \Leftrightarrow \\ \delta_{X_j}(x_j) &= \delta_{X_j}(y_j) \end{aligned}$$

and thus $x_j = y_j$ since δ_{X_j} is assumed to be injective. It follows immediately that $x = y$, and thus δ_X is injective too. $\square$

It is clear from the proof above that if T preserves finite objects (which is quite often the case in practise), then for δ_X to be injective on finite objects, it must also be the case that L should preserve finite objects. This can never be the case if L represents a logic with infinitely many modal operators such as GML (see Section 2.6.1). This problem is addressed in the next section.

5.1.2 Completeness by filtration

When a language includes infinitely many operators (and thus L does not preserve finite sets), it is often easier to prove completeness (or, the injectivity of the semantic transformation) on finite sublanguages. This is the idea behind the filtration method which is well-known to modal logicians. Here we follow the presentation of this method by Kurz and Petrişan in [KP10]. We will present examples in subsection ??.

The idea is to express the logic's functor L by a cofiltered limit of functors which preserve finite boolean algebras. As we shall see in the next section, the preservation of finite algebras is a particularly useful property for reasoning about the semantic transformation and its transpose. For now however, we will state the following more general approximation result.

Theorem 5.8 (Weak completeness by filtration). *In the fundamental situation of Diagram (1.2), let $L : \mathbf{C} \rightarrow \mathbf{C}$ be a finitary functor, let $T : \mathbf{D} \rightarrow \mathbf{D}$, and assume that*

$$L = \operatorname{colim}_k L_k$$

where the index k runs over a cofiltered category. Assume further that for each k there exists a natural transformation

$$\delta_k : L_k G \rightarrow GT$$

such that for every natural transformation $\zeta_{k,k'} : L_k \rightarrow L_{k'}$ in the diagram defining L , $\delta_k = \delta_{k'} \circ \zeta_{k,k'}$, then there exists a unique natural transformation $\delta : LG \rightarrow GT$ which agrees with δ_k for each k and is injective if each δ_k is injective.

Proof. Since $\mathbf{C}$ is a finitary variety, $U : \mathbf{C} \rightarrow \mathbf{Set}$ creates filtered colimits, and in particular

$$ULA = U \operatorname{colim}_k L_k A = \operatorname{colim}_k UL_k A$$

By construction of colimits in $\mathbf{Set}$, if $a \in LGX$ for some set X , then there must exist k in the index category and an element $a' \in L_k GX$ such that $\xi_{k,X}(a') = a$, where $\xi_k : L_k G \rightarrow LG$ is the obvious natural transformation. We now define

$$\delta_X(x) : LGX \rightarrow GTX, a \mapsto \delta_{k,X}(a')$$

To see that this is well-defined, assume that there exist k' and a'' such that $\xi_{k',X}(a'') = a$. By construction of colimits in **Set**, this means that there must exist a zigzag of natural transformations between L_k and $L_{k'}$ which connects the elements a' and a'' . By the assumption that $\delta_k = \delta_{k'} \circ \zeta_{k,k'}$ for any natural transformation $\zeta_{k,k'} : L_k \rightarrow L_{k'}$ in the diagram, it follows that $\delta_{k',X}(a'') = \delta_{k,X}(a')$, and δ is thus well-defined. If each δ_k is injective, it is clear from the construction we have just described that so is δ . $\square$

Weak completeness by filtration is then an immediate consequence of Theorems 5.8 and 5.6.

5.1.3 A brief note on expressivity

Given a T -coalgebra $\gamma : W \rightarrow TW$ and a valuation $v : V \rightarrow GW$, the interpretation map $\llbracket - \rrbracket_W : \mu L_V \rightarrow GW$ has an adjoint transpose:

$$\theta_W : W \rightarrow F\mu L$$

called the theory map and which associates to each point $w \in W$ the ultrafilter of formulas which are satisfied at w . The following result was elegantly shown in [JS10]:

Theorem 5.9 ([JS10]). *Let $L : \mathbf{C} \rightarrow \mathbf{C}$ be a finitary functor, let $T : \mathbf{D} \text{ to } \mathbf{D}$ where $\mathbf{D}$ has an strong-epi mono factorisation system, and let $\delta : LG \rightarrow GT$ be a semantic transformation. If the adjoint transpose $\hat{\delta} : TF \rightarrow FL$ is injective then logically indistinguishable points are behaviourally equivalent, i.e. if two points have the same theory, then there exists a coalgebra homomorphism identifying them.*

Proof. Let $\gamma : W \rightarrow TW$ be a coalgebra, and let $v : V \rightarrow GW$ be a valuation. This data defines the interpretation map $\llbracket - \rrbracket : \mu L_V \rightarrow GW$, where $L_V = L(-) + FV$. The idea is to factor the adjoint transpose of $\llbracket - \rrbracket$, i.e. the theory map $\theta : W \rightarrow F\mu L_V$, into a strong-epi mono factorisation, which is always possible in $\mathbf{D}$ by assumption. We have

$$\begin{array}{ccccc}
 W & \xrightarrow{q} & Q & \xrightarrow{m} & F\mu L_V \\
 \gamma \downarrow & & \downarrow & & \downarrow \simeq \\
 TW & \xrightarrow{Tq} & TQ & \xrightarrow{Tm} & TF\mu L_V \xrightarrow{\hat{\delta}_{\mu L_V}} FL\mu L_V
 \end{array}$$

θ (curved arrow from W to $F\mu L_V$)

where the unique diagonal fill-in morphism exists by virtue of q being strong. If $\theta(x) = \theta(y)$, then since m is injective, we must have $q(x) = q(y)$, and the two states are thus behaviourally equivalent. $\square$

5.2 Strong completeness

In this section we will present the basic set-up for the study of coalgebraic completeness-via-canonicity. In the language of modal logic we will investigate ‘canonical’ models in the coalgebraic framework, i.e. coalgebraic models built from maximally consistent sets of formulas.

5.2.1 The coalgebraic Jónsson-Tarski theorem

The construction of canonical models for coalgebraic logics was first studied in [KKP05] via the notion of ultrafilter extension, and subsequently in [KR12] and [SP09] where general conditions are presented for the existence of ‘canonical’ models (more on this terminology shortly). The idea of building models from formulas is pervasive in logics, and is a fundamental tool in modal logic, but the Jónsson-Tarski theorem actually shows more than this. Although it is not usually stated explicitly, the classical Jónsson-Tarski theorem can in fact be split into two independent components:

1. the first shows how to build Kripke models of normal modal logics from the algebra of formulas using the usual notion of maximally consistent sets (also called theory),
2. the second shows how validity in this Kripke model is equivalent to (algebraic) validity in the *canonical extension* of the BAO defined by the logic.


The first component of the theorem is used to show strong completeness of normal modal logics, whilst the second is used to show strong completeness in the presence of canonical equations. This section and [KKP05, KR12, SP09] focus solely on the first part of the theorem, we will examine the second part in the next section.

We start by making an important comment on the term ‘canonical’, which is related to the dual role of the Jónsson-Tarski theorem. In its first role, it builds ‘canonical models’, which in the coalgebraic framework correspond to T -coalgebra over $F\mu(L(-) + FV)$. However, as pointed out in [SP09, KR12] for example, these ‘canonical’ models are usually not defined uniquely. In fact an element of choice is always required, for example when a transformation $FL \rightarrow TF$ is used to define them (see [KR12]), it is in general not natural as we shall soon see, making canonical models incompatible with one another. The term ‘canonical’ is thus rather imprecise in this context, prompting Pattinson and Schröder to suggest the term ‘quasi-canonical’ in [SP09]. In the context of models we will adopt the following convention: we will use the term ‘canonical’ as meaning ‘governed by a law’ and thus call a model canonical if it can be

explicitly constructed. Models which can merely be shown to exist will be call ‘quasi-canonical’.

In its second role, the Jónsson-Tarski theorem connects the building of models from formulas to the notion of *canonical* extension. As we have seen in Chapter 2, every DLE can be embedded into a better behaved DLE, called its canonical extension, whose underlying distributive lattice structure is complete and whose join irreducibles are join-dense. Moreover, the canonical extension is *uniquely defined* (up to isomorphism, see e.g. [GH09]) by the requirement that that $\mathbb{A}$ must be dense and compact in $\mathbb{A}^\sigma$. The term ‘canonical’ is thus completely unambiguous in the context of DLEs (and BAEs). There is therefore a fundamental contradiction between the two usages of the term ‘canonical’, but as we shall see in the next section, this contradiction often resolve itself spontaneously thanks to the notion of *smoothness* introduced in Chapter 2, viz. we can get a uniquely defined canonical extension through the use of a non-uniquely defined canonical model.

Our criterion for the Jónsson-Tarski theorem is essentially (6.11) of [KR12], which is slightly different from that of [KKP05, SP09], but is easier to state. We will show how it is related to the conditions of Theorem 3 of [KKP05], and Theorem 2.5 of [SP09]. We place ourselves once again in the following situation



where $\mathbf{C}$ is a finitary algebraic variety.

Theorem 5.10 (Coalgebraic Jónsson-Tarski Theorem). *Let $\delta : LG \rightarrow GT$ be a semantic natural transformation and let η, ϵ denote the unit and counit of the adjunction $F \dashv G$. If the adjoint semantic transformation $\hat{\delta} : TF \rightarrow FL$ given by $\hat{\delta} = FL\eta \circ F\delta_F \circ \epsilon_{TF}$ has (not necessarily natural) right inverse $\hat{\delta}^{-1} : FL \rightarrow TF$, then for any A in $\mathbf{Alg}_{\mathbf{C}}(L)$, the morphism $\eta_A : A \rightarrow GFA$ lifts to an L -algebra morphism.*

Proof. Starting with an algebra $\alpha : LA \rightarrow A$, we show that the following dia-

gram commutes:

$$\begin{array}{ccc}
 LA & \xrightarrow{L\eta_A} & LGFA \\
 \downarrow \alpha & \searrow \eta_{LA} & \downarrow \delta_{FA} \\
 & & GTFA \\
 & & \downarrow G(\hat{\delta}_A^{-1}) \\
 & & GFLA \\
 & & \downarrow GF\alpha \\
 A & \xrightarrow{\eta_A} & GFA
 \end{array} \tag{5.1}$$

To see that it commutes, notice first that the trapezium at the bottom of the diagram defined by η_{LA} commutes by naturality of η . So we need only to show that the triangle of the top of the diagram commutes. To see this, consider the following diagram:

$$\begin{array}{ccccc}
 LA & \xrightarrow{L\eta_A} & LGFA & \xrightarrow{\delta_{FA}} & GTFA \\
 \eta_{LA} \downarrow & & & & \downarrow \eta_{GTFA} \\
 GFLA & \xrightarrow{GFL\eta_A} & GFLGFA & \xrightarrow{GF\delta_{FA}} & GF GTFA
 \end{array}$$

$\curvearrowright G\epsilon_{TFA}$

The naturality square clearly commutes, and moreover, by the fact that $F \dashv G$ it is also the case that $G\epsilon_{TFA} \circ \eta_{GTFA} = \text{Id}_{GTFA}$, and thus

$$\begin{aligned}
 G\hat{\delta}_A \circ \eta_{LA} &= G\epsilon_{TFA} \circ GF\delta_{FA} \circ GFL\eta_A \circ \eta_{LA} \\
 &= G\epsilon_{TFA} \circ \eta_{GTFA} \circ \delta_{FA} \circ L\eta_A \\
 &= \delta_{GA} \circ L\eta_A
 \end{aligned}$$

From this we get:

$$G(\hat{\delta}_A^{-1}) \circ \delta_{FA} \circ L\eta_A = G(\hat{\delta}_A^{-1}) \circ G\hat{\delta}_A \circ \eta_{LA} = G(\hat{\delta}_A \circ \hat{\delta}_A^{-1}) \circ \eta_{LA} = \eta_{LA}$$

By definition of $\hat{\delta}_A^{-1}$ as a right inverse. □

In the case where $A = \mathbf{GFV}$ is the free L -algebra over V , i.e. the abstract language, we will call the coalgebra

$$\hat{\delta}_{\mathbf{GFV}}^{-1} \circ F\langle - \rangle_V : F(\mathbf{GFV}) \rightarrow TF(\mathbf{GFV})$$

(where $\langle - \rangle_V$ is the structure map of the language), the **(pseudo)-canonical T -model** (depending on whether $\hat{\delta}^{-1}$ can be given explicitly) associated with the semantic transformation δ . More generally, for any object $\alpha : LA \rightarrow A$ in $\mathbf{Alg}_C(L)$, we will call a coalgebra $\hat{\delta}_A^{-1} \circ F\alpha : FA \rightarrow TFA$ a **(pseudo)canonical**

model on A . We will call the L -algebra $GF\alpha \circ G\hat{\delta}_A^{-1} \circ \delta_{FA} : LGFA \rightarrow GFA$ the **Jónsson-Tarski extension** of the L -algebra $\alpha : LA \rightarrow A$. We will show in that $\hat{\delta}^{-1}$ can in practise often be given explicitly, modulo a non-constructive choice principle.

Remark 5.11. We could recast the coalgebraic Jónsson-Tarski Theorem without any reference to the semantics of the language defined by L in terms of the **distributive law of a monad over a functor**. As is well-known, every adjunction gives rise to a monad, and in particular $GF : \mathbf{C} \rightarrow \mathbf{C}$ is a monad. A distributive law of GF over a functor $L : \mathbf{C} \rightarrow \mathbf{C}$ is a natural transformation $\lambda : LGF \rightarrow GFL$ which interacts well with the unit and multiplication of the monad, in particular the following triangle must commute

$$\begin{array}{ccc} LA & \xrightarrow{L\eta_A} & LGFA \\ & \searrow \eta_{LA} & \downarrow \lambda_A \\ & & GFLA \end{array}$$

But this is exactly the triangle from the proof of Theorem 5.10. In other words, given a distributive law of the monad GF over L , we immediately get an injective embedding of an L algebra into its canonical extension. Notwithstanding the elegance of this approach, our purpose is, at the end of the day, to build models, and factoring through GFA is therefore essential.

5.2.2 Strong completeness for positive coalgebraic logics

In this section we will apply our coalgebraic Jónsson-Tarski Theorem 5.10 to L -algebras in **DL** interpreted by T -coalgebras in **Pos**, i.e. to positive coalgebraic logics (see [Dun95, KKV12a, BKV13]). Clearly, weak completeness follows from the Jónsson-Tarski theorem if we can show that the unit η_{GFV} of the adjunction $F \dashv G$, evaluated at the ‘language object’ GFV , is injective; in which case different formulas will have different denotations. In the cases which interest us here, i.e. when $\mathbf{C} = \mathbf{DL}$ and $\mathbf{D} = \mathbf{Pos}$, η is given by

$$\eta_A : A \rightarrow \mathcal{UPf}A, a \mapsto \{p \in \mathcal{P}fA \mid a \in p\}$$

and we must distinguish two problems: (1) proving that η_{GFV} is injective and (2) proving that η is injective in general. The problems are different because V is traditionally assumed to be countable, and thus so is the logic GFV . In this case it can easily be shown that η_{GFV} is injective via a Lindenbaum lemma argument.

Lemma 5.12 (Lindenbaum’s Lemma, [Dun95]). *If $L : \mathbf{DL} \rightarrow \mathbf{DL}$ preserves countable sets and V is countable, then η_{GFV} is injective.*

Proof. Let $a, b \in \mathbf{GFV}$ such that $a \neq b$, we need to build a prime filter which contains one formula and not the other. If $a \neq b$ then either $a \not\leq b$ or $b \not\leq a$, we assume the former and it follows easily that $\uparrow a \cap \downarrow b = \emptyset$. Since V is countable and L preserves countable sets, so is $\mathbf{GFV}$ and we can enumerate all the formulas $a_0, a_1, \dots \in \mathbf{GFV}$. The strategy is to add a_0 to $\uparrow a$ if $\uparrow (a \wedge a_0) \cap \downarrow b = \emptyset$ or else add it to $\downarrow b$, and then iterate the construction. It can be shown that each formula can in this way be added to the left or to the right whilst preserving disjointness (see Lemma 3.5 of [Dun95]) and that the limit of the procedure provides a prime filter which contains a but not b . $\square$

In fact this kind of procedure gives more than weak completeness since it can be carried out for arbitrary pairs of sets forming a non-intersecting ideal-filter pair. This provides us with a strong completeness result for the abstract Hilbert system of Section 1.5.3.

Theorem 5.13 (Strong completeness for countable L -algebras over **DL**). *If $L : \mathbf{DL} \rightarrow \mathbf{DL}$ preserves countable sets and satisfies the conditions of the Jónsson-Tarski Theorem 5.10 for a choice of semantics $\delta : \mathcal{LU} \rightarrow \mathcal{UT}$, and if V is countable, then the Hilbert system $\vdash_{\mathbf{ML}}$ associated with L is strongly complete w.r.t. $\mathbf{CoAlg}(T)$*

Proof. Given arbitrary sets $\Phi, \Psi \subseteq \mathbf{GFV}$ of formulas such that $\Phi \not\vdash_{\mathbf{ML}} \Psi$ we must show that there exists a model satisfying all the formulas of Φ and none of the formulas of Ψ . It is not difficult to show that $\Phi \not\vdash_{\mathbf{ML}} \Psi$ iff for every *finite* sets of formulas $\phi_1, \dots, \phi_m \in \Phi$ and $\psi_1, \dots, \psi_n \in \Psi$

$$\phi_1 \wedge \dots \wedge \phi_m \not\leq \psi_1 \vee \dots \vee \psi_n$$

i.e. the filter $\langle \Phi \rangle^\uparrow$ generated by Φ has an empty intersection with the ideal $\langle \Psi \rangle^\downarrow$ generated by Ψ (in particular, if we work in **BDL** then no meet of elements of Φ can equal $\perp$, i.e. we have a ‘consistent’ set of formulas). By the same procedure as above, we can enumerate all formulas in $\mathbf{GFV}$ and successively add them either to $\langle \Phi \rangle^\uparrow$ or to $\langle \Psi \rangle^\downarrow$ whilst maintaining the intersection between the sets empty. The process yields a prime filter p_Φ containing every formula in Φ , and thus

$$p_\Phi \in \eta_{\mathbf{GFV}}(\phi), \text{ i.e. } p_\Phi \in \llbracket \phi \rrbracket_{\mathbf{PGFV}}, \text{ i.e. } p_\Phi \models \phi, \text{ for each } \phi \in \Phi$$

and since p_Φ is disjoint from Ψ it follows that

$$p_\Phi \notin \eta_{\mathbf{GFV}}(\psi), \text{ i.e. } p_\Phi \notin \llbracket \psi \rrbracket_{\mathbf{PGFV}}, \text{ i.e. } p_\Phi \not\models \psi, \text{ for each } \psi \in \Psi$$

which is what we wanted to show. $\square$

If we wish a more categorically pleasing result, and in particular include cases where V is not countable, then we must prove the injectivity of η in general. This requires an additional ingredient, namely the Prime Filter Theorem (see Theorem 2.5). Indeed, to extend a filter $\uparrow a$ such that $\uparrow a \cap \downarrow b = \emptyset$ to a prime filter F such that $F \cap \downarrow b \in \emptyset$ (as in the proof of Lemma 5.12) in a general distributive lattice requires precisely the Prime Filter Theorem. Note that this theorem is independent of ZF, but is strictly weaker than ZFC, and we must thus assume that we are doing mathematics in which the Prime Filter Theorem holds. Modulo this hypothesis, we can adapt Theorem 5.13 and show the following result:

Theorem 5.14 (Strong completeness for arbitrary L -algebras over **DL**). *Let $L : \mathbf{DL} \rightarrow \mathbf{DL}$, $T : \mathbf{Pos} \rightarrow \mathbf{Pos}$ and $\delta : LU \rightarrow UT$ satisfy the conditions of the Jónsson-Tarski Theorem 5.10. For any distributive lattice A and any sets $\Phi, \Psi \subseteq A$ such that no finite meet of elements of Φ lies below a finite join of element of Ψ , there exist a canonical T -model of A satisfying all formulas of Φ and no formula of Ψ .*

For an important class of problems it is in fact possible to explicitly build a right inverse $\hat{\delta}^{-1}$. This class of problem includes positive modal logic (see [Dun95]) and our description of intuitionistic and sub-structural logics of Section 2.6.2 (see [DP15]). It is characterised by n -ary modalities which either (1) preserve joins or anti-preserve meets in each of their arguments, or (2) preserve meets or anti-preserve joins in each of their arguments. Formally, for a given finitary signature Σ we define the syntax building functor:

$$L_\Sigma : \mathbf{DL} \rightarrow \mathbf{DL} \begin{cases} L_\Sigma A = \text{FS}_\Sigma \text{UA} / \equiv \\ L_\Sigma f : L_\Sigma A \rightarrow L_\Sigma B, [a]_\equiv \mapsto [\text{S}_\Sigma \text{U}f(a)]_\equiv \end{cases}$$

where $f : A \rightarrow B$ is a **DL**-morphism and $\equiv$ is the fully invariant relation in **DL** generated by equations of the type

$$\heartsuit(a_1 \vee b_1, \dots, a_k \vee b_k, a_{k+1} \wedge b_{k+1}, \dots, a_n \wedge b_n) = \heartsuit(a_1, \dots, a_n) \vee \heartsuit(b_1, \dots, b_n) \quad (5.2)$$

or

$$\heartsuit(a_1 \wedge b_1, \dots, a_k \wedge b_k, a_{k+1} \vee b_{k+1}, \dots, a_n \vee b_n) = \heartsuit(a_1, \dots, a_n) \wedge \heartsuit(b_1, \dots, b_n) \quad (5.3)$$

for every n -ary symbol $\heartsuit$ in the signature Σ . We will always group the arguments of each n -ary symbol according to whether they are isotone or antitone, starting with k isotone arguments and followed $n - k$ antitone arguments. Note that the syntax builders L_{Hey} and L_{RL} defined in Section 2.6.2 do indeed fall into this category. We will call logics defined by this kind of functor **positive relational logics**. The reason for this terminology is that we are going to interpret

these logics in coalgebras defined by the **Pos** equivalent of the powerset functor. Following [VK11], we define the convex powerset functor $\mathbf{P}_c : \mathbf{Pos} \rightarrow \mathbf{Pos}$ as the functor associating with each poset X its set of convex subsets (i.e. subsets Y with the property that if $x, z \in Y$ then $y \in Y$ whenever $x \leq y \leq z$) ordered by the Egli-Milner order, i.e. $Y \leq Y'$ in $\mathbf{P}_c$ if

$$\begin{aligned} &(\forall y \in Y)(\exists y' \in Y') \text{ s.th. } y \leq y', \text{ and} \\ &(\forall y' \in Y')(\exists y \in Y) \text{ s.th. } y \leq y' \end{aligned}$$

Note that upsets are particular examples of convex sets. The convex powerset functor associates with a monotone map the convex closure of its direct image, i.e. if $f : X \rightarrow Y$ is a monotone map, then $\mathbf{P}_c f : \mathbf{P}_c X \rightarrow \mathbf{P}_c Y, U \mapsto \overline{f[U]}$, where $\overline{(-)}$ denotes the convex closure. We now define for a signature Σ the functor:

$$T_\Sigma : \mathbf{Pos} \rightarrow \mathbf{Pos} \begin{cases} T_\Sigma X = \prod_{\heartsuit \in \Sigma} \mathbf{P}_c(X^k \times (X^{\text{op}})^{n-k}) \\ T_\Sigma f = \prod_{\heartsuit \in \Sigma} \mathbf{P}_c f^n \end{cases}$$

where k is the number of isotone arguments of $\heartsuit$, n its total arity, X^{op} is the dual poset of X (i.e. same carrier, order reversed), and f^n is the usual n -ary product of maps (it is easy to see that **Pos** has products which are created in **Set**).

For better readability we will follow the following convention. For any symbol $\heartsuit$ in the signature Σ such that $n = \text{ar}(\heartsuit)$ and k denotes the number of isotone arguments (regrouped as the first arguments), i will always denote an index value between 1 and k , and such an index will be called an isotone index. Similarly, we assume from now on that j will always denote an index value between $k+1$ and n , and such an index will be called an antitone index. Using this convention, we interpret the logic defined by L_Σ in T_Σ -coalgebras via the following semantic transformation $\delta_X^\Sigma : L_\Sigma \mathcal{U}X \rightarrow \mathcal{U}T_\Sigma X$:

$$\begin{aligned} &[\heartsuit(u_1, \dots, u_k, u_{k+1}, \dots, u_n)] \mapsto \\ &\{t \in T_\Sigma X \mid \exists (a_1, \dots, a_n) \in \pi_\heartsuit(t), (a_i \in u_i \text{ for all } i) \text{ and } (a_j \notin u_j \text{ for all } j)\} \\ &\text{if } \heartsuit \text{ satisfies Eq. (5.2)} \\ &[\heartsuit(u_1, \dots, u_k, u_{k+1}, \dots, u_n)] \mapsto \\ &\{t \in T_\Sigma X \mid \forall (a_1, \dots, a_n) \in \pi_\heartsuit(t), (a_j \in u_j \text{ for all } j) \Rightarrow (a_i \in u_i \text{ for all } i)\} \\ &\text{if } \heartsuit \text{ satisfies Eq. (5.3)} \end{aligned}$$

where $\pi_\heartsuit$ is the projection on the $\heartsuit$ -component of the product defining T_Σ .

Proposition 5.15. δ^Σ is well-defined.

Proof. For δ^Σ to be well defined, the image of each ‘generator’ must form an upset, and δ^Σ must preserve the equations (5.2) and (5.3). Let us fix an arbitrary

poset X , a symbol $\heartsuit \in \Sigma$ with $n = \text{ar}(\heartsuit)$, and let $(u_1, \dots, u_n), (v_1, \dots, v_n) \in (\mathcal{U}X)^n$. We first show that if $t \in \delta^\Sigma(\heartsuit(u_1, \dots, u_n))$ and $t \leq t' \in T_\Sigma X$, then $t' \in \delta^\Sigma(\heartsuit(u_1, \dots, u_n))$ too. By definition of the partial order on $T_\Sigma X$, this means that $\pi_\heartsuit(t) \leq \pi_\heartsuit(t')$ for the (component-wise) Egli-Milner order, i.e. for each $(a_1, \dots, a_k, a_{k+1}, \dots, a_n) \in \pi_\heartsuit(t)$ there exist $(a'_1, \dots, a'_k, a'_{k+1}, \dots, a'_n) \in \pi_\heartsuit(t')$ such that $a_i \leq a'_i$ and $a'_j \leq a_j$ (note the reversal of direction due to the presence of $(-)^{\text{op}}$ in the definition of T_Σ), and conversely. We need to distinguish two cases:

1. $\heartsuit$ satisfies (5.2). By definition there exists $(a_1, \dots, a_n) \in \pi_\heartsuit(t)$ such that $a_i \in u_i$ and $a_j \notin u_j$. We want to show that there exist $(a'_1, \dots, a'_n) \in \pi_\heartsuit(t')$ with the same property. By definition of the Egli-Milner order we know that there exist $(a'_1, \dots, a'_n) \in \pi_\heartsuit(t')$ s.th. $a_i \leq a'_i$ and $a'_j \leq a_j$. Since u_i is an upset, $a'_i \in u_i$ and since u_j is an upset it follows that $a'_j \notin u_j$.
2. $\heartsuit$ satisfies (5.3). By definition, for all $(a_1, \dots, a_n) \in \pi_\heartsuit(t)$, if $a_j \in u_j$ for every antitone index j then $a_i \in u_i$ for every isotone index i . We want to show that the same holds for every $(a'_1, \dots, a'_n) \in \pi_\heartsuit(t')$. So let us assume that each $a'_j \in u_j$. By definition of the Egli-Milner order, we know that there exists $(a_1, \dots, a_n) \in \pi_\heartsuit(t)$ s.th. $a_i \leq a'_i$ and $a'_j \leq a_j$. Since u_j is an upset it follows that $a_j \in u_j$ for each antitone j , and thus $a_i \in u_i$ for every isotone i , and thus $a'_i \in u_i$ since u_i is an upset.

Finally, we need to show that δ^Σ preserves equations (5.2) and (5.3). We have (dropping the square brackets for notational clarity):

$$\begin{aligned}
& \delta_X^\Sigma(\heartsuit(u_1 \vee v_1, \dots, u_k \vee v_k, u_{k+1} \wedge v_{k+1}, \dots, u_n \wedge v_n)) \\
&= \{t \in T_\Sigma X \mid \exists (a_1, \dots, a_n) \in \pi_\heartsuit(t), (a_i \in u_i \vee v_i) \text{ and } (a_j \notin u_j \wedge v_j)\} \\
&= \{t \in T_\Sigma X \mid \exists (a_1, \dots, a_n) \in \pi_\heartsuit(t), (a_i \in u_i) \text{ and } (a_j \notin u_j)\} \cup \\
&\quad \{t \in T_\Sigma X \mid \exists (a_1, \dots, a_n) \in \pi_\heartsuit(t), (a_i \in v_i) \text{ and } (a_j \notin v_j)\} \\
&= \delta_X^\Sigma(\heartsuit(u_1, \dots, u_n)) \vee \delta_X^\Sigma(\heartsuit(v_1, \dots, v_n))
\end{aligned}$$

A similar proof shows that δ^Σ preserves equation (5.3). $\square$

We will now show how a right inverse to $\hat{\delta}^\Sigma$ can be defined explicitly, providing strong completeness of the logic defined by L_Σ w.r.t. T_Σ -coalgebras. To keep the proof readable we will need the following additional notation: if $\vec{x}$ is a $k-1$ -tuple of isotone arguments (resp. an $n-k-1$ -tuple of antitone arguments) and y_i (resp. y_j) is a particular isotone (resp. antitone) argument, then $\vec{x}/y_i$ (resp. $\vec{x}/y_j$) denotes the k -tuple (resp. $n-k$ -tuple) formed by inserting y_i (resp. y_j) in i^{th} (resp. j^{th}) position in $\vec{x}$.

Theorem 5.16. *Any positive relational logic defined by a signature Σ is strongly complete with respect to T_Σ -coalgebras.*

Proof. By the coalgebraic Jónsson-Tarski Theorem 5.10 and Theorem 5.14 it suffices to define a right-inverse to the transformation $\hat{\delta} : T_\Sigma \text{Pf} \rightarrow \text{Pf} L_\Sigma$. For the purpose of finding a right inverse to $\hat{\delta}$, we only need to specify when a generator of $L_\Sigma A$ belongs to $\hat{\delta}_A(t)$ for some $t \in T_\Sigma \text{Pf} A$.

if $\heartsuit$ satisfies Eq. (5.2):

$\heartsuit(a_1, \dots, a_n) \in \hat{\delta}_A(t)$ iff

$\exists(F_1, \dots, F_n) \in \pi_{\heartsuit}(t)$ s.th. $(a_i \in F_i \text{ for all } i) \text{ and } (a_j \notin F_j \text{ for all } j)$

if $\heartsuit$ satisfies Eq. (5.3):

$\heartsuit(a_1, \dots, a_n) \in \hat{\delta}_A(t)$ iff

$\forall(F_1, \dots, F_n) \in \pi_{\heartsuit}(t) (a_j \in F_j \text{ for all } j) \Rightarrow (a_i \in F_i \text{ for all } i)$

We now define a morphism $\gamma_A : \text{Pf} L_\Sigma A \rightarrow T_\Sigma \text{Pf} A$ which we will show is a right inverse. It maps $F \in \text{Pf} L_\Sigma A$ to the $|\Sigma|$ -tuple of sets of n -tuples $(F_1, \dots, F_n), n = \text{ar}(\heartsuit), \heartsuit \in \Sigma$ of prime filters of A defined by:

$$\begin{aligned} \pi_{\heartsuit}(\gamma_A(F)) = & \{(F_1, \dots, F_n) \mid \\ & (a_i \in F_i) \text{ and } (a_j \notin F_j) \Rightarrow \heartsuit(a_1, \dots, a_n) \in F \text{ if } \heartsuit \text{ satisfies (5.2)} \\ & \heartsuit(a_1, \dots, a_n) \in F \Rightarrow \text{if } (a_j \in F_j) \text{ then } (a_i \in F_i) \text{ if } \heartsuit \text{ satisfies (5.3)}\} \end{aligned}$$

To see that γ_A is well-defined, we need to check that for each symbol $\heartsuit \in \Sigma$, $\pi_{\heartsuit}(\gamma_A(F))$ is a convex subset of $\text{P}_c(\text{Pf} A^k \times (\text{Pf} A^{\text{op}})^{n-k})$. So let $(F_1, \dots, F_n) \leq (F'_1, \dots, F'_n) \leq (F''_1, \dots, F''_n)$ with $(F_1, \dots, F_n), (F'_1, \dots, F'_n) \in \pi_{\heartsuit}(\gamma_A(F))$. We need to show that $(F'_1, \dots, F'_n) \in \pi_{\heartsuit}(\gamma_A(F))$. From the presence of $(-)^{\text{op}}$ in the definition, we get that $F_i \subseteq F'_i \subseteq F''_i$ for isotone indices and $F''_j \subseteq F'_j \subseteq F_j$ for antitone indices. Again we distinguish two cases

1. $\heartsuit$ satisfies (5.2). Assume $(a_i \in F'_i)$ and $(a_j \notin F'_j)$. From the inequalities above, it follows that $a_i \in F''_i$ and $a_j \notin F''_j$ and thus $\heartsuit(a_1, \dots, a_n) \in F$ since $(F''_1, \dots, F''_n) \in \pi_{\heartsuit}(\gamma_A(F))$.
2. $\heartsuit$ satisfies (5.3). Assume that $\heartsuit(a_1, \dots, a_n) \in F$ and that $a_j \in F'_j$. It follows that $a_j \in F_j$ and thus $a_i \in F_i$ and thus $a_i \in F'_i$ as desired.

With all these checks in place, we now move to the difficult part of the proof: showing that γ_A is a right inverse of $\hat{\delta}_A$. For this it must be the case that $\heartsuit(a_1, \dots, a_n) \in F \Leftrightarrow \heartsuit(a_1, \dots, a_n) \in \hat{\delta}_A(\gamma_A(F))$, i.e. by unravelling the

definitions:

$$\heartsuit(a_1, \dots, a_n) \in F \Leftrightarrow \exists (F_1, \dots, F_n) \in \pi_{\heartsuit}(\gamma_A(F)) \text{ s.th.} \quad (5.4)$$

$$(a_i \in F_i \text{ for all } i) \text{ and } (a_j \notin F_j, \text{ for all } j)$$

$$\text{if } \heartsuit \text{ satisfies Eq. (5.2)}$$

$$\heartsuit(a_1, \dots, a_n) \in F \Leftrightarrow \forall (F_1, \dots, F_n) \in \pi_{\heartsuit}(\gamma_A(F)) \text{ s.th.} \quad (5.5)$$

$$(a_j \in F_j, \text{ for all } j) \Rightarrow (a_i \in F_i \text{ for all } i)$$

$$\text{if } \heartsuit \text{ satisfies Eq. (5.3)}$$

It is easy to see that if $\heartsuit$ satisfies Eq. (5.2), then the right-to-left direction of Eq. (5.4) follows immediately from the definition of γ_A . Similarly, if $\heartsuit$ satisfies Eq. (5.3), then the left-to-right direction of Eq. (5.5) follows from the definition. The difficult part of the proof are the converse implications.

Left-to-right of Eq. (5.4) Assume that $\heartsuit(a_1, \dots, a_n) \in F$, we must build an n -tuple $(F_1, \dots, F_n) \in \pi_{\heartsuit}(\gamma_A(F))$ of prime filters such that $a_i \in F_i$ for every isotone i , and $a_j \notin F_j$ for every antitone j . Since this tuple must belong to $\pi_{\heartsuit}(\gamma_A(F))$ it should also have the property that if $(a'_i \in F_i)$ for all i and $(a'_j \notin F_j)$ for all j then $\heartsuit(a'_1, \dots, a'_n) \in F$, or equivalently that if $\heartsuit(a'_1, \dots, a'_n) \notin F$ then there must exist i with $a'_i \notin F_i$ or j with $a'_j \in F_j$. To achieve this we use a Prime Filter Theorem argument on a set of ideal-filter pairs $((F_1, I_1), \dots, (F_n, I_n))$ which we will call $\mathcal{P}(a_1, \dots, a_n)$ and define as follows. For $1 \leq l \leq k$ we put:

$$(\text{iso } 1) \uparrow a_l \subseteq F_l \subseteq \{c_l \mid (\forall \vec{c} \in \prod_{i \neq l} F_i)(\forall \vec{d} \in \prod_j I_j), \heartsuit(\vec{c}/c_l, \vec{d}) \in F\}$$

$$(\text{iso } 2) I_l = \{c_l \mid (\exists \vec{c} \in \prod_{i \neq l} F_i)(\exists \vec{d} \in \prod_j I_j) \text{ s.th. } \heartsuit(\vec{c}/c_l, \vec{d}) \notin F\}$$

For $k+1 \leq l \leq n$ we put:

$$(\text{anti } 1) \downarrow a_j \subseteq I_j \subseteq \{d_l \mid (\forall \vec{c} \in \prod_i F_i)(\forall \vec{d} \in \prod_{j \neq l} I_j), \heartsuit(\vec{c}, \vec{d}/d_l) \in F\}$$

$$(\text{anti } 2) F_l = \{d_l \mid (\exists \vec{c} \in \prod_i F_i)(\exists \vec{d} \in \prod_{j \neq l} I_j) \text{ s.th. } \heartsuit(\vec{c}, \vec{d}/d_l) \notin F\}$$

We show the following facts about $\mathcal{P}(a_1, \dots, a_n)$

- $((\uparrow a_1, I_1), \dots, (F_{k+1}, \downarrow a_{k+1}), \dots)$ belongs to $\mathcal{P}(a_1, \dots, a_n)$, which is therefore not empty. Since the I_i s and F_j s are defined from the filters $\uparrow a_i$ and the ideals $\downarrow a_j$, we need only check the inclusions in (iso 1) and (anti 1). For this let $1 \leq l \leq k$, $a_l \leq c_l$ and assume for the sake of contradiction that there exists $\vec{c} \in \prod_{i \neq l} \downarrow a_i$ and $\vec{d} \in \prod_j \downarrow a_j$ such that $\heartsuit(\vec{c}/c_l, \vec{d}) \notin F$, but then $\heartsuit(\vec{c}/c_l, \vec{d}) \geq \heartsuit(a_1, \dots, a_n) \notin F$, a contradiction with the starting assumption. Thus (iso 1) holds, and the proof that (anti 1) holds is dual.
- $\mathcal{P}(a_1, \dots, a_n)$ forms a poset under component-wise set inclusion

- Each $I_l, 1 \leq l \leq k$ is an ideal. Assume $c_l \in I_l$ and $c'_l \leq c_l$. By definition, there exists $\vec{c} \in \prod_{i \neq l} F_i$ and $\vec{d} \in \prod_j I_j$ s.th. $\heartsuit(\vec{c}/c_l, \vec{d}) \notin F$. It follows immediately from the fact that l is an isotone argument and that F is a filter that $\heartsuit(\vec{c}/c'_l, \vec{d}) \notin F$ and thus $c_{l'} \in I_l$. Now assume that $c_l, c'_l \in I_l$. By definition there exist $\vec{c}, \vec{c}' \in \prod_{i \neq l} F_i$ and $\vec{d}, \vec{d}' \in \prod_j I_j$ s.th. $\heartsuit(\vec{c}/c_l, \vec{d}), \heartsuit(\vec{c}'/c'_l, \vec{d}') \notin F$. It follows that $\heartsuit(\vec{c} \wedge \vec{c}'/c_l, \vec{d} \vee \vec{d}'), \heartsuit(\vec{c} \wedge \vec{c}'/c'_l, \vec{d} \vee \vec{d}') \notin F$. Since F is prime and $\heartsuit$ satisfies (5.4) it follows that

$$\begin{aligned} & \heartsuit(\vec{c} \wedge \vec{c}'/c_l, \vec{d} \vee \vec{d}') \vee \heartsuit(\vec{c} \wedge \vec{c}'/c'_l, \vec{d} \vee \vec{d}') \\ &= \heartsuit(\vec{c} \wedge \vec{c}'/c_l \vee c'_l, \vec{d} \vee \vec{d}') \notin F \end{aligned}$$

and since F_i is a filter and each I_j is an ideal, it follows that $\vec{c} \wedge \vec{c}' \in \prod_{i \neq l} F_i$ and $\vec{d} \vee \vec{d}' \in \prod_j I_j$ witness that $c_l \wedge c'_l \in I_l$.

- For totally dual reasons, $F_l, k+1 \leq l \leq n$ is a filter.
- $F_i \cap I_i = \emptyset = F_j \cap I_j$. For an isotone index l assume that there exist $x_l \in F_l, y_l \in I_l$ such that $x_l \leq y_l$. By definition of I_l there exists $\vec{c} \in \prod_{i \neq l} F_i$ and $\vec{d} \in \prod_j I_j$ s.th. $\heartsuit(\vec{c}/y_l, \vec{d}) \notin F$. It follows that $\heartsuit(\vec{c}/x_l, \vec{d}) \notin F$, contradicting the property (iso 1) of F_l . The proof for antitone arguments is identical.

Let us now check that $\mathcal{P}(a_1, \dots, a_n)$ has upper bound of chains. Assume $((F_1^k, I_1^k), \dots, (F_n^k, I_n^k))_{k \in \omega}$ is a chain of elements of $\mathcal{P}(a_1, \dots, a_n)$, and define

$$F_l^\infty = \bigcup_k F_l^k \quad I_l^\infty = \bigcup_k I_l^k, \quad k = 1, \dots, n$$

It is well-known and easy to check that the union of a chain of filters (resp. ideals) is a filter (esp. ideal). Let us now check that conditions (iso 1)-(iso 2) and (anti 1)-(anti 2) are satisfied. The first inclusions of (iso 1) and (anti 1) are trivially satisfied. For the second inclusion, let $1 \leq l \leq k$, $c_l \in F_l^\infty$, and let $\vec{c} \in \prod_{i \neq l} F_i^\infty$ and $\vec{d} \in \prod_j I_j^\infty$. By definition each component c_i of $\vec{c}$ can be traced back to F_i^k for a certain k , and similarly for each component d_j of $\vec{d}$. By taking the maximum of these k s we get an integer $\tilde{k}$ s.th. $\vec{c} \in \prod_{i \neq l} F_i^{\tilde{k}}, \vec{d} \in \prod_j I_j^{\tilde{k}}$. It follows immediately by the condition (iso 1) on $F_l^{\tilde{k}}$ that $\heartsuit(\vec{c}/c_l, \vec{d}) \in F$ as desired. The proof for l antitone is identical. Finally we need to show that (iso 2) and (anti 2) are satisfied. Let $c_l \in I_l^\infty$ for l isotone. By construction, there must exist k such that $c_l \in I_l^k$, from which it follows that there exist $\vec{c} \in \prod_{i \neq l} F_i^k$ and $\vec{d} \in \prod_j I_j^k$ s.th. $\heartsuit(\vec{c}/c_l, \vec{d}) \notin F$. By definition it follows immediately that there exist $\vec{c} \in \prod_{i \neq l} F_i^\infty$ and $\vec{d} \in \prod_j I_j^\infty$ s.th. $\heartsuit(\vec{c}/c_l, \vec{d}) \notin F$, i.e. we have

shown the inclusion from left to right in (iso 2). For the opposite direction, assume that there exist $\vec{c} \in \prod_{i \neq l} F_i^\infty$ and $\vec{d} \in \prod_j I_j^\infty$ s.th. $\heartsuit(\vec{c}/\vec{c}_l, \vec{d}) \notin F$. Once again, each component of $\vec{c}$ and of $\vec{d}$ can be traced back to a filter or an ideal indexed by a finite integer k . By taking $\tilde{k}$ to be the maximum over all these integers, we can assume that $\vec{c} \in \prod_{i \neq l} F_i^{\tilde{k}}$ and $\vec{d} \in \prod_j I_j^{\tilde{k}}$, and thus $c_l \in I_l^{\tilde{k}}$, and thus $c_l \in I_l^\infty$, which concludes the proof that (iso 2) holds. The proof that (anti 2) holds is identical.

Since $\mathcal{P}(a_1, \dots, a_n)$ is closed upper bound of chains, we can apply Zorn's lemma and get the existence of a maximal element of $\mathcal{P}(a_1, \dots, a_n)$ which we call

$$((G_1, J_1), \dots, (G_n, J_n))$$

We claim that $G_1, \dots, G_n$ are the prime filters we are looking for. Let us first check that they are prime. For isotone arguments, let assume that $c_l \vee c'_l \in G_l$ but that $c_l, c'_l \notin G_l$. It follows that

$$\begin{aligned} & ((G_1, J_1), \dots, (G_l, J_l), \dots, (G_n, J_n)) \subsetneq \\ & ((G_1, J'_1), \dots, (\langle G_l \cup \{c_l\} \rangle, J_l), (G_{l+1}, J'_{l+1}), \dots, (G'_{k+1}, J_{k+1}), \dots, (G'_n, J_n)) \end{aligned}$$

where $\langle G_l \cup \{c_l\} \rangle$, the filter generated by $G_l \cup \{c_l\}$, and for each isotone index i' , $J'_{i'}$ is defined by (iso 2) by allowing components in this filter. Explicitly:

$$J'_{i'} = \{c_{i'} \mid (\exists \vec{c} \in \prod_{i < l, i \neq i'} G_i \times \langle G_l \cup \{c_l\} \rangle \times \prod_{i > l, i \neq i'} G_i) (\exists \vec{d} \in \prod_j J_j) \text{ s.th. } \heartsuit(\vec{c}/c_{i'}, \vec{d}) \notin F\}$$

Similarly, for j' antitone, we make sure that (anti 2) is respected and define:

$$G_{j'} = \{d_{j'} \mid (\exists \vec{c} \in \prod_{i < l} G_i \times \langle G_l \cup \{c_l\} \rangle \times \prod_{i > l} G_i) (\exists \vec{d} \in \prod_{j \neq j'} J_j) \text{ s.th. } \heartsuit(\vec{c}, \vec{d}/d_{j'}) \notin F\}$$

Since $(G_1, J_1, \dots, (G_n, J_n))$ is maximal in $\mathcal{P}(a_1, \dots, a_n)$, the new n -tuple of filter-ideal pairs cannot be in $\mathcal{P}(a_1, \dots, a_n)$. Since (iso 2) and (anti 2) are satisfied by definition, and since the first inclusion of (iso 1) and (anti 1) are also necessarily satisfied, for the new n -tuple not to be in $\mathcal{P}(a_1, \dots, a_n)$, it must be the case that the second inclusion of (iso 1) or of (anti 1) must be violated. Since all G_i s are unchanged apart from G_l , and since every J_j is unchanged, it is not too hard to see that the violations of either inclusions are equivalent and happen precisely if there exists $(x_1, \dots, x_l \wedge c_l, \dots, x_k) \in \prod_{i < l} G_i \times \langle G_l \cup \{c_l\} \rangle \times \prod_{i > l} G_i$ and a $\vec{d} \in \prod_j J_j$ such that

$$\heartsuit(x_1, \dots, x_l \wedge c_l, \dots, x_k, \vec{d}) \notin F$$

By definition, this means that $x_l \wedge c_l \in J'_l$, i.e. there exist $y_l \in J'_l$ s.th. $x_l \wedge c_l \leq y_l$. A completely analogous argument shows that there must exist $(x'_1, \dots, x'_l \wedge$

$c'_l, \dots, x'_k) \in \prod_{i < l} G_i \times \langle G_l \cup \{c_l\} \rangle \times \prod_{i > l} G_i$ and $y'_l \in J'_l$ s.th. $x'_l \wedge c'_l \leq y'_l$. By combining the two inequalities we get

$$(x_l \wedge c_l) \vee (x'_l \wedge c'_l) \leq y_l \vee y'_l$$

i.e.

$$(x_l \vee x'_l) \wedge (x_l \vee c'_l) \wedge (c_l \vee x'_l) \wedge (c_l \vee c'_l) \leq y_l \vee y'_l$$

since J'_l is an ideal, G_l is a filter, $x_l, x'_l \in G_l$ and $c_l \vee c'_l \in G_l$ is assumed, we get $G_l \cap J'_l \neq \emptyset$ which is easily seen to be a contradiction by construction of J'_l . Thus we cannot have both $c_l \notin G_l$ and $c'_l \notin G_l$, and G_l is thus prime. A completely dual proof shows that each J_j is prime. It follows that J_i and G_j are prime too. We show this for J_l , l isotone, and a completely dual argument shows that each G_j is prime too. If $c_l \wedge c'_l \in J_l$, then by definition there exists $\vec{c} \in \prod_{i \neq l} G_i$, $\vec{d} \in \prod_j J_j$ s.th. $\heartsuit(c/c_l \wedge c'_l, \vec{d}) \notin F$. Assume now for the sake of contradiction that $c_l, c'_l \notin J_l$. Then they must both belong to G_l since G_l is maximal, and not being in J_l precisely ensures that the second inequality of (iso 1) is respected. But if $c_l, c'_l \in G_l$, then $c_l \wedge c'_l \in G_l$ since it's a filter, which contradicts the property established above that $G_l \cap J_l = \emptyset$. Thus either c_l or c'_l belongs to J_l which is thus prime.

Finally, we need to check that $((G_1, J_1), \dots, (G_n, J_n)) \in \pi_\heartsuit(\gamma_A(F))$, i.e. that if $\heartsuit(a'_1, \dots, a'_n) \notin F$ then there must exist either an isotone i with $a'_i \notin G_i$ or an antitone j with $a'_j \in G_j$. Assume that $a'_i \in G_i$ apart from a single isotone index l and that $a'_j \notin G_j$ for every antitone index j , we must show that $a_l \notin G_l$. By definition (iso 2), the k -tuple $(a'_i)_{i \neq l}$ and the $n - k$ -tuple (a_j) clearly witness the fact that $a_l \in J_l$, and since $G_l \cap J_l = \emptyset$ it follows that $a_l \notin G_l$ as desired. Similarly, if $a'_i \in G_i$ for each isotone index and $a'_j \notin G_j$ apart from a single antitone index l , then $a_l \in G_l$ by definition (anti 2).

Right-to-left of Eq. (5.5). Is shown in a very similar way as the left-to-right direction of Eq. (5.4) above. We show the contrapositive: if $\heartsuit(a_1, \dots, a_n) \notin F$, then there must exist $(F_1, \dots, F_n) \in \pi_\heartsuit(\gamma_A(F))$ s.th. $a_j \in F_j$ for all antitone j but $a_i \notin F_i$ for all isotone i (and thus for a particular i). Since $(F_1, \dots, F_n) \in \pi_\heartsuit(\gamma_A(F))$ we must ensure that if there exist $a'_1, \dots, a'_n$ s.th. $a'_j \in F_j$ for each j but $a'_i \notin F_i$ for a certain i , then $\heartsuit(a'_1, \dots, a'_n) \notin F$. We proceed as above and define a set $\mathcal{P}(a_1, \dots, a_n)$ of filter-ideal pairs $((F_1, I_1), \dots, (F_n, I_n))$ as follows. For $1 \leq i \leq k$ we put:

1. $F_i = \{c_i \mid \exists \vec{d} \in \prod_j F_j \text{ s.th. } \heartsuit(c_i, \vec{d}) \in F\}$
2. $\downarrow a_i \subseteq I_i \subseteq \{c_i \mid \forall \vec{d} \in \prod_j F_j, \heartsuit(c_i, \vec{d}) \notin F\}$

For $k + 1 \leq j \leq n$ we put:

1. $\uparrow a_j \subseteq F_j \subseteq \{d_j \mid \forall \vec{c} \in \prod_i I_i, \heartsuit(\vec{c}, d_j) \notin F\}$

$$2. I_j = \{d_j \mid \exists \vec{c} \in \prod_i I_i \text{ s.th. } \heartsuit(\vec{c}, d_j) \in F\}$$

The desired tuple of prime filters is obtained in the same way as above by dualization, i.e. by showing that the collection $\mathcal{P}(a_1, \dots, a_n)$ forms a non-empty poset, that each F_i (resp. I_j) is a filter (resp. an ideal), that the intersection of each filter-ideal pair is empty, and that $\mathcal{P}(a_1, \dots, a_n)$ is closed under union of chains. By applying Zorn's lemma we once again find a maximal element consisting of prime filter-prime ideal pairs which fulfill the desired requirement. $\square$

5.2.3 Strong completeness for boolean coalgebraic logics

We now turn our attention to the case where $\mathbf{C} = \mathbf{BA}$ and $\mathbf{D} = \mathbf{Set}$. The category of sets has a very remarkable property which allows for a reformulation of the coalgebraic Jónsson-Tarski theorem. In Theorem 3 of [KKP05], the criterion for the existence of a Jónsson-Tarski embedding an L -algebra A is essentially that of our Theorem 5.10, but is formulated in the following way: there must exist a map $h_A : \mathbf{Uf}LA \rightarrow T\mathbf{Uf}A$ such that for any $u \in \mathbf{Uf}LA$,

$$b \in u \Leftrightarrow h_A(u) \in \delta_{\mathbf{Uf}A} \circ L\eta_A \circ h_A(u) \quad (5.6)$$

If such an h_A exists then it is easy to see that $\hat{\delta}_A^{-1}$ can be substituted by h_A^{-1} in the Diagram (5.1) of the proof above and preserve commutativity, i.e. h_A is a right inverse to $\hat{\delta}_A$: if we unravel the definition of $\hat{\delta}$ we get for any $u \in \mathbf{Uf}LA$:

$$\begin{aligned} \hat{\delta}_A \circ h_A(u) &= \mathbf{Uf}L\eta_A \circ \mathbf{Uf}\delta_{\mathbf{Uf}A} \circ \epsilon_{T\mathbf{Uf}A} \circ h_A(u) \\ &= (\delta_{\mathbf{Uf}A} \circ L\eta_A)^{-1} \circ \epsilon_{T\mathbf{Uf}A}(h_A(u)) \\ &= (\delta_{\mathbf{Uf}A} \circ L\eta_A)^{-1}(\{U \in \mathcal{PT}\mathbf{Uf}A \mid h_A(u) \in V\}) \end{aligned}$$

If $b \in u$, then by (5.6) we know that $h_A(u) \in \delta_{\mathbf{Uf}A} \circ L\eta_A(b)$, and thus $\delta_{\mathbf{Uf}A} \circ L\eta_A(b)$ is one of these U s and in particular $b \in \hat{\delta}_A \circ h_A(u)$. Conversely, if $b \in \hat{\delta}_A \circ h_A(u)$, then $\delta_{\mathbf{Uf}A} \circ L\eta_A(b) \in \epsilon_{T\mathbf{Uf}A}(h_A(u))$, i.e. $h_A(u) \in \delta_{\mathbf{Uf}A} \circ L\eta_A(b)$ and thus $b \in u$ by (5.6). So we have established that $\hat{\delta}_A \circ h_A = \text{Id}_{\mathbf{Uf}LA}$, i.e. that h_A is a section of $\hat{\delta}_A$, which is thus a split epi.

The interesting property of $\mathbf{Set}$ which we will now exploit is that all epimorphisms are split (modulo the axiom of choice). In other words, if $\hat{\delta}_A$ is an epimorphism, then it must have a section, i.e. a right-inverse, and *any* section of $\hat{\delta}_A$ will act as a good h_A and define a Jónsson-Tarski embedding. Indeed, assume that $s_A : \mathbf{Uf}LA \rightarrow T\mathbf{Uf}A$ is such a section, by reusing the last steps of the proof of the coalgebraic Jónsson-Tarski Theorem we get:

$$\begin{aligned} \mathcal{P}s_A \circ \delta_{\mathbf{Uf}A} \circ L\eta_A &= \mathcal{P}s_A \circ \mathcal{P}\hat{\delta}_A \circ \eta_{LA} \\ &= \mathcal{P}(\hat{\delta}_A \circ s_A) \circ \eta_{LA} = \eta_{LA} \end{aligned}$$

as was shown above and as required for Diagram (5.1) to commute.

It is thus clear that the requirement that $\hat{\delta}$ be an epi-transformation is enough to embed an L -algebra A into its canonical extension $\mathcal{P}\mathbf{Uf}A$ equipped with an L -algebra structure. However, this embedding is not uniquely defined and relies completely on the axiom of choice. We therefore call a model obtained in this way ‘**quasi-canonical**’, following the terminology of [SP09], and reserve the term ‘canonical’ for the case where the right inverse to $\hat{\delta}$ can be built explicitly, such as in the case of (positive) relational logics described above. Moreover, the surjectivity of $\hat{\delta}$ only guarantees the existence of pseudo-canonical models for the $\mathbf{Uf} \dashv \mathcal{P}$ adjunction, i.e. for boolean coalgebraic logic. In general not all surjective morphisms are split. Indeed, even in the case of **Pos** it is easy to find surjective morphisms which do not split. Thus whilst useful and simple in the boolean case, showing that $\hat{\delta}$ is surjective does not generalize, even to the case of coalgebraic logics based on distributive lattices.

Following the discussion above and the prevalence of boolean coalgebraic logic, we will spend the remainder of this Section investigating conditions on L, T and δ which guarantee that $\hat{\delta}$ is an epi-transformation in the context of the adjunction $\mathbf{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}$. As we will also show, we can in many cases explicitly build the right inverse of $\hat{\delta}$ when it is also assumed to be a mono-transformation, i.e. when the logic is assumed to be expressive. We first need a few easy technical results.

Lemma 5.17. *The functors $\mathbf{Uf}$ and $\mathcal{P}$ turn monos into epis and vice-versa.*

Proof. To see that $\mathbf{Uf}$ turns epi into monos is easy: since it is a left adjoint it preserves epis, i.e. it sends an epi in $\mathbf{BA}$ to an epi in $\mathbf{Set}^{\text{op}}$, i.e. to a mono in $\mathbf{Set}$. The dual argument shows that $\mathcal{P}$ sends epis to monos.

To see that $\mathcal{P}$ sends monos to epis, let $f : X \rightarrow Y$ be a mono in $\mathbf{Set}$, and let $U \in \mathcal{P}A$, we need to find a $V \in \mathcal{P}B$ such that $U = \mathcal{P}fV = f^{-1}V$. Since f is injective we can simply take $V = f[U]$ and we then get $f^{-1}f[U] = U$ by injectivity.

Let us finally check that $\mathbf{Uf}$ also sends monos to epis. Let $f : A \rightarrow B$ be a monomorphism in $\mathbf{BA}$, then f is injective on the carriers. To check that $\mathbf{Uf}f$ is epi, we need to show that for every $u \in \mathbf{Uf}A$, there exists a $v \in \mathbf{Uf}B$ such that $u = \mathbf{Uf}f(v) = f^{-1}(v)$. We choose v as the ultrafilter generated by $f[u]$ (this means that we need to assume the axiom of choice) and show that $u = f^{-1}(v)$. Clearly, if $a \in u$ then $f(a) \in f[u] \subseteq v$, i.e. $u \subseteq f^{-1}(v)$. Conversely, if we choose $a \in f^{-1}(v)$, then either a or $\neg a$ must belong to u . If $\neg a \in u$, then $f(\neg a) \in v$, i.e. $\neg f(a) \in v$. But since $f(a)$ is in v by assumption that $a \in f^{-1}(v)$, and since v is an ultrafilter, we must have $f(a) = f(\neg a)$ to avoid a contradiction. But since f is injective this means that $a = \neg a$ which is impossible. Thus it must be the case that $a \in u$ and we thus have $f^{-1}(v) \subseteq u$, and therefore $u = f^{-1}(v)$. $\square$

Lemma 5.18. *The unit and counit η, ϵ and of the adjunction $\mathbf{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}$ are isomorphisms on finite boolean algebras and finite sets respectively.*

Proof. Let A be a finite boolean algebra, then $\mathcal{P}\mathbf{Uf}A$ is also finite. Recall that $\eta_A : A \rightarrow \mathcal{P}\mathbf{Uf}A, a \mapsto \{u \in \mathbf{Uf}A \mid a \in u\}$ is a monomorphism and $\mathbf{U}\eta_A$ is injective. Since every ultrafilter in a finite boolean algebra is principal, an element U of $\mathcal{P}\mathbf{Uf}A$ is a finite set of principal ultrafilter and is thus in one-to-one correspondence with a finite set $\{a_1, \dots, a_n\}$ of elements of A . We can therefore define for any finite A

$$h_A : \mathcal{P}\mathbf{Uf}A \rightarrow A, \{\uparrow a_1, \dots, \uparrow a_n\} \mapsto a_1 \vee \dots \vee a_n$$

It is then easy to see that

$$\eta_A \circ h_A(U) = \eta_A(a_1 \vee \dots \vee a_n) = \bigcup_{i=1}^n \eta_A(a_i) = U$$

and h_A is thus a right inverse of η_A . Similarly, it is easy to check that $h_A \circ \eta_A(a) = a$ for any $a \in A$, and thus h_A is also a left inverse and η_A is therefore iso.

For the counit, let X be a finite set, and recall that $\epsilon_X : X \rightarrow \mathbf{Uf}\mathcal{P}X, x \mapsto \{U \subseteq X \mid x \in U\} = \uparrow \{x\}$. This map is clearly injective. To see that it is surjective, pick $u \in \mathbf{Uf}\mathcal{P}X$, and consider the set $U = \bigcap u$. By the properties of ultrafilters we automatically get that $U \neq \emptyset$ and $u = \uparrow U$. To see that U is in fact a singleton, assume $\{x, y\} \subseteq U, x \neq y$. Since u is an ultrafilter, either $\{x\} \in u$ or $\{x\}^c \in u$. If $\{x\}^c \in u$, then clearly x cannot be an element of U , so it must in fact be the case that $\{x\} \in u$, but then we clearly cannot have $\{x, y\} \subseteq U$ and thus it must be the case that $U = \{x\}$. This shows that $u = \uparrow \{x\} = \epsilon_X(x)$ for some x . $\square$

From these lemmas, we can immediately show that if $\hat{\delta}$ is surjective at each stage A , then δ is injective and the logic is (weakly) complete as we have seen in Section 5.1.

Proposition 5.19. *If $L : \mathbf{BA} \rightarrow \mathbf{BA}$ preserves monos (in particular if it preserves weak pullbacks) and $\hat{\delta} : T\mathbf{Uf} \rightarrow \mathbf{Uf}L$ is surjective, then $\delta : L\mathcal{P} \rightarrow \mathcal{P}T$ is injective.*

Proof. By using the same argument as in the proof of Theorem 5.10, we know that the following diagram commutes:

$$\begin{array}{ccccc} LA & \xrightarrow{L\eta_A} & L\mathcal{P}\mathbf{Uf}A & \xrightarrow{\delta_{\mathbf{Uf}A}} & \mathcal{P}T\mathbf{Uf}A \\ \eta_{LA} \downarrow & & \nearrow \mathcal{P}\hat{\delta}_A & & \uparrow \mathcal{P}\epsilon_{T\mathbf{Uf}A} \\ \mathcal{P}\mathbf{Uf}LA & \xrightarrow{\mathcal{P}\mathbf{Uf}L\eta_A} & \mathcal{P}\mathbf{Uf}L\mathcal{P}\mathbf{Uf}A & \xrightarrow{\mathcal{P}\mathbf{Uf}\delta_{\mathbf{Uf}A}} & \mathcal{P}\mathbf{Uf}\mathcal{P}T\mathbf{Uf}A \end{array}$$

By Lemma 5.17, since $\hat{\delta}_A$ is epi (in **Set**) $\mathcal{P}\hat{\delta}_A$ must be mono (in **BA**). Moreover, since the unit $\eta : \text{Id} \rightarrow \mathcal{P}\text{Uf}$ is a mono-transformation and L preserves monos, we must have $\delta_{\text{Uf}A}$ mono. $\square$

Lemma 5.20. *Let $L : \mathbf{BA} \rightarrow \mathbf{BA}$ and let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be two functors related by a semantic transformation $\delta : L\mathcal{P} \rightarrow \mathcal{P}T$, and assume further that T preserve finite sets. The adjoint transformation $\hat{\delta} : T\text{Uf} \rightarrow \text{Uf}L$ is injective (resp. surjective) on finite objects iff δ is surjective (resp. injective) on finite objects.*

Proof. This follows by the definition of the adjoint transformation by

$$\hat{\delta}_A = \text{Uf}L\eta_A \circ \text{Uf}\delta_{\text{Uf}A} \circ \epsilon_{T\text{Uf}A}$$

Clearly, if A is a finite boolean algebra, then by Lemma 5.18 and the assumption on T , $\text{Uf}L\eta_A$ and $\epsilon_{T\text{Uf}A}$ are isomorphisms. The result then follows easily from Lemma 5.17. $\square$

The following technical lemma is not strictly necessary, but it makes the proof of Theorem 5.24 more symmetric, and it also provides an important example of functor weakly preserving cofiltered limits.

Lemma 5.21. *The covariant powerset functor $\mathbf{P} : \mathbf{Set} \rightarrow \mathbf{Set}$ weakly preserves cofiltered limits.*

Proof. Let $(X_i)_{i \in I}$ be a cofiltered diagram in **Set**, let $X = \lim_i X_i$ and let $\pi_i : X \rightarrow X_i$. The limit X is the set of coherent collection of elements $(x_i)_{i \in I}$, where coherent means that if $x_i \in X_i$ and $f_{ij} : X_i \rightarrow X_j$, then $x_j = f_{ij}(x_i)$.

We need to show that $\mathbf{P}\lim_i X_i$ is a weak limit for the diagram $(\mathbf{P}X_i)_{i \in I}$. It is clear that by functoriality of $\mathcal{Q}$, it is a cone for this diagram. Given a coherent collection of subsets $U_i \in \mathbf{P}X_i$, i.e. a collection of subsets such that $f_{ij}[U_i] = U_j$ for any $f_{ij} : X_i \rightarrow X_j$, we need to find an element of $\mathbf{P}X = \mathbf{P}\lim_i X_i$ mapping onto the various elements of the collection. For this we can simply take $U = \{(x_i)_{i \in I} \in \lim_i X_i \mid x_i \in U_i\}$. For U to be a good choice, it must be the case that $\mathbf{P}\pi_i(U) = \pi_i[U] = U_i$. Consider any $y_i \in U_i$, then we must find in any $U_j \in \mathbf{P}X_j$, an element y_j belonging to a coherent collection containing y_i . Since the diagram is cofiltered, there exist an X_k , a $U_k \in \mathbf{P}X_k$, and morphisms $f_{ki} : X_k \rightarrow X_i$ and $f_{kj} : X_k \rightarrow X_j$ such that $f_{ki}[U_k] = U_i$ and $f_{kj}[U_k] = U_j$. In particular, there must exist y_k such that $f_{ki}(y_k) = y_i$ and $f_{kj}(y_k) = y_j$, and since this holds for any U_j , y_i must belong to a coherent family in U , i.e. $\pi_i[U] = U_i$ as required. $\square$

Another important example, which will be crucial in the next subsection will be the following (the proof is very similar to the one just given):

Lemma 5.22. *Let n be a finite set, then contravariant hom functor $\text{hom}(-, n) : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$ weakly preserves cofiltered limits (i.e. turns cofiltered limits into weak filtered colimits).*

Proof. Let $X = \lim_i X_i$ be the limit of a cofiltered diagram, and let us consider a collection of maps $f_i \in \text{hom}(X_i, n)$ which is coherent in the sense that if $h_{ij} : X_i \rightarrow X_j$ is the image of a morphism from the index category, then

$$\text{hom}(-, n)(f_j) = f_j \circ h_{ij} = f_i$$

We need to show that we can find a (not necessarily unique) element $f \in \text{hom}(X, n)$ such that $\text{hom}(-, n)(\pi_i)(f_i) = f$ for all the projection maps $\pi_i : X \rightarrow X_i$. Let us define

$$f : X \rightarrow n, x \mapsto \begin{cases} y & \text{if } \forall i, f_i(\pi_i(x)) = y \\ 0 & \text{else} \end{cases}$$

Let us show that for any $x \in X$ we do have $f(x) = f_i(\pi_i(x))$. For this we need to show that if $f_i(\pi_i(x)) = y$ for a particular i , then $f_j(\pi_j(x)) = y$ for any other j of the index category. For this we use the fact that the diagram is cofiltered, i.e. given i, j , there must exist a k and morphisms $h_{ki} : X_k \rightarrow X_i$ and $h_{kj} : X_k \rightarrow X_j$. By a proof completely analogous to that of Proposition 5.37, we can assume without loss of generality that all morphisms h_{ki} are surjective, i.e. $x_i \in X_i$ must have a pre-image x_k under the map h_{ki} . By the fact that $f_k = f_i \circ h_{ki}$ we have

$$f_k(x_k) = f_i \circ h_{ki}(x_k) = f_i(x_i) = y$$

and thus

$$y = f_k(x_k) = f_j \circ h_{kj}(x_k) = f_j(x_j)$$

Thus we indeed have that if $f_i(\pi_i(x)) = y$ then $f_j(\pi_j(x)) = y$ too, for any index object j . This establishes $f(x) = f_i(\pi_i(x))$, and proves that x defined as above does indeed define a weak colimit. $\square$

Note the duality with the notion of finitely presentable object: in $\mathbf{Set}$ an object X is finitely presentable iff $\text{hom}(X, -)$ preserves filtered colimits iff X is finite. Here we have that finite sets are also ‘finitely co-presentable’.

Corollary 5.23. *The contravariant powerset functors $\mathcal{Q} : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Set}$ weakly preserve cofiltered limits.*

Proof. This is a simple consequence of the fact that $\mathcal{Q} = \text{hom}(-, 2)$. $\square$

The following characterisation of situations where completeness implies strong completeness is due to [KR12], but here we decouple L from T , whereas *op.cit.*

considers the case where the syntax functor L is defined as $L_T = \mathcal{PTUf}$ (which is similar in spirit to the nabla logic case where $L_T = \mathcal{FTU}$). We start by focusing on the epi nature of $\hat{\delta}$. Recall from the remark above that this will imply the existence of *quasi*-canonical models.

Theorem 5.24. *Let $L : \mathbf{BA} \rightarrow \mathbf{BA}$ and let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ preserve finite sets (i.e. if X is finite, so is TX), and assume further that L is finitary and that T weakly preserves cofiltered limits. If the semantic transformation $\delta : LP \rightarrow PT$ is a mono-transformation, then its transpose $\hat{\delta} : TUf \rightarrow UfL$ is an epi-transformation.*

Proof. The idea of the proof comes from Section 6 of [KR12]. Since $\mathbf{BA}$ is locally finitely presentable, we can write any object A of $\mathbf{BA}$ as a filtered colimit $\text{colim}_{i \in I} A_i$ where each A_i is finitely presentable. Moreover, since $\mathbf{BA}$ is a locally finite variety and finitely generated boolean algebras are finite, each A_i is *finite*.

Recall that $\hat{\delta} = UfL\eta \circ Uf\delta_{Uf} \circ \epsilon_{TUf}$. We examine each step of this composition starting from the last, in order to find an inverse image to an arbitrary element $u \in UfLA$. Consider first:

$$\begin{array}{ccc}
 & UfLPufA_i & \xrightarrow[\simeq]{UfL\eta_{A_i}} UfLA_i \\
 & \uparrow & \nearrow \\
 UfLPufA & \xrightarrow{UfL\eta_A} & UfLA \\
 & \downarrow & \searrow \\
 & UfLPufA_j & \xrightarrow[\simeq]{UfL\eta_{A_j}} UfLA_j \\
 & & \downarrow
 \end{array} \tag{5.7}$$

Since $A = \lim_i A_i$, L is finitary and Uf is left adjoint, we have $UfLA = UfL \text{colim}_i A_i = \lim_i UfLA_i$. Thus $u \in UfLA$ is uniquely associated with a coherent collection $u_i \in UfLA_i, i \in I$. Since L and Uf preserve finite objects, it follows from Lemma 5.18 that each $UfL\eta_{A_i}$ is an iso. We can therefore map the collection $(u_i)_{i \in I}$ to a coherent collection $(v_i)_{i \in I}, v_i \in UfLPufA_i$. By using the aforementioned properties of Uf and L and Corollary 5.23, we get that $UfLPufA$ is a weak limit of the diagram $(UfLPufA_i)_{i \in I}$, and thus the collection $(v_i)_{i \in I}$ defines a (not necessarily unique) element $v \in UfLPufA$. Since each diagram of the shape of (5.7) commutes it is clear that $UfL\eta_A(v) = u$.

For the second step we have:

$$\begin{array}{ccccc}
 & \text{UfPTUf}A_i & \xrightarrow{\text{Uf}\delta_{\text{Uf}A_i}} & \text{UfLPUf}A_i & \\
 & \uparrow & & \uparrow & \\
 \text{UfPTUf}A & \xrightarrow{\text{Uf}\delta_{\text{Uf}A}} & \text{UfLPUf}A & & \\
 & \downarrow & & \downarrow & \\
 & \text{UfPTUf}A_i & \xrightarrow{\text{Uf}\delta_{\text{Uf}A_i}} & \text{UfLPUf}A_j &
 \end{array} \quad (5.8)$$

By Lemma 5.17, since $\delta_{\text{Uf}A}$ is assumed to be mono, $\text{Uf}\delta_{\text{Uf}A}$ is epi for any A , and we can thus find an inverse image w such that $\text{Uf}\delta_{\text{Uf}A}(w) = v$. By projecting onto the various $\text{UfPTUf}A_i$ this w defines a coherent family $(w_i)_{i \in I}, w_i \in \text{UfPTUf}A_i$.

Finally, at the last step we have:

$$\begin{array}{ccccc}
 & T\text{Uf}A_i & \xrightarrow[\simeq]{\epsilon_{T\text{Uf}A_i}} & \text{UfPTUf}A_i & \\
 & \uparrow & & \uparrow & \\
 T\text{Uf}A & \xrightarrow{\epsilon_{T\text{Uf}A}} & \text{UfPTUf}A & & \\
 & \downarrow & & \downarrow & \\
 & T\text{Uf}A_j & \xrightarrow[\simeq]{\epsilon_{T\text{Uf}A_j}} & \text{UfPTUf}A_j &
 \end{array} \quad (5.9)$$

Since Uf and T preserve finite objects, each $T\text{Uf}A_i$ is finite, and by Lemma 5.18, each $\epsilon_{T\text{Uf}A_i}$ is an isomorphism. The coherent family $(w_i)_{i \in I}$ thus defines a coherent family $x_i \in T\text{Uf}A_i, i \in I$. Since T weakly preserves cofiltered limits and Uf preserves filtered colimits, $T\text{Uf}A$ is a weak limit of the diagram $(T\text{Uf}A_i)_{i \in I}$, and the coherent collection $(x_i)_{i \in I}$ therefore defines a (not necessarily unique) element $x \in \text{Uf}TA$. Since each diagram of the shape of (5.7) commutes it is clear that $\epsilon_{T\text{Uf}A}(x) = w$. By combining our results, it is clear that

$$\hat{\delta}_A(x) = \text{Uf}L\eta_A(\text{Uf}\delta_{\text{Uf}A}(\epsilon_{T\text{Uf}A}(x))) = u$$

i.e. we have shown that $\hat{\delta}_A$ is surjective. $\square$

We combine the previous Proposition 5.24 with the assumption of expressivity to build a natural right inverse to $\hat{\delta}$, and get the Jónsson-Tarski theorem 5.10 (as well as strong completeness).

Theorem 5.25. *Let $L : \mathbf{BA} \rightarrow \mathbf{BA}$ be finitary, let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ preserve finite sets, and assume further that T weakly preserves cofiltered limits. If the semantic transformation $\delta : \mathbf{LP} \rightarrow \mathbf{PT}$ is injective, and the logic defined by L is expressive w.r.t. $\mathbf{CoAlg}(T)$, then the Jónsson-Tarski Theorem 5.10 holds.*

Proof. We want to show that the construction detailed in the proof of Proposition 5.24 defines a right inverse to $\hat{\delta}$.

We use the same idea as in the previous proof, but we choose a particular representation of boolean algebras. For this let $f : A \rightarrow B$ be a morphism in $\mathbf{BA}$, and let $\mathcal{D}_c^A$ and $\mathcal{D}_c^B$ be the canonical diagrams of A and B respectively. Recall from Chapter 3, that the canonical diagram for an object A in a locally presentable category (such as $\mathbf{BA}$) is the diagram defined by the projection functor $\mathcal{D}_c^A : \mathbf{C}_\lambda \downarrow A \rightarrow \mathbf{C}$ (where $\mathbf{C}_\lambda$ is the subcategory of λ -presentable objects), recall also from Lemma 3.34 that $A = \text{colim } \mathcal{D}_c^A$. Note that if $f : A \rightarrow B$, then any object A_i of the canonical diagram $\mathcal{D}_c^A$ is also an object in the canonical diagram $\mathcal{D}_c^B$ satisfying

$$\begin{array}{ccc} A_i & & \\ h_i \downarrow & \searrow g_i = f \circ h_i & \\ A & \xrightarrow{f} & B \end{array}$$

In consequence, by applying $\text{Uf}L$, we have commutative diagrams between the corresponding cofiltered limits of the type:

$$\begin{array}{ccccc} \text{Uf}LA & \xleftarrow{\text{Uf}Lf} & & \text{Uf}LB & \\ \text{Uf}Lh_i \downarrow & & \text{Uf}Lg_i & \swarrow & \downarrow \text{Uf}Lg_k \\ \text{Uf}LA_i & & \text{Uf}LA_j & & \text{Uf}LA_k \end{array}$$

and it is clear that every coherent family in $\text{Uf}LB = \lim \text{Uf}L\mathcal{D}_c^B$ gets mapped to a coherent family in $\text{Uf}LA = \lim \text{Uf}L\mathcal{D}_c^A$.

We have shown in Theorem 5.24 that by lifting each element of a coherent family $v = (v_i)_{i \in \{\mathbf{BA}_\omega \downarrow B\}^{\text{ob}}}$ (with $v_i \in \text{Uf}LA_i$) in $\text{Uf}LB$ against the corresponding maps $\text{Uf}L\eta_{A_i} \circ \text{Uf}\delta_{\text{Uf}A_i} \circ \epsilon_{T\text{Uf}A_i}$, we can find a coherent family $(t_i)_{i \in \{\mathbf{BA}_\omega \downarrow B\}^{\text{ob}}}$ with $t_i \in T\text{Uf}A_i$ and thus an element $t_v \in T\text{Uf}B$ such that $\hat{\delta}_B(t_v) = v$. Since we are assuming that $\hat{\delta}$ is injective, this preimage is unique. The naturality of the maps $v \mapsto t_v$ then follows from the connection between $\mathcal{D}_c^A$ and $\mathcal{D}_c^B$ and thus between the unique coherent families representing the preimages. $\square$

The same proof can be used when the logic is not necessarily expressive, i.e. when $\hat{\delta}$ is not necessarily injective, but we must then require the surjectivity of δ on finite structures.

Theorem 5.26. *Let $L : \mathbf{BA} \rightarrow \mathbf{BA}$ be finitary, let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ preserve finite sets, and assume further that T weakly preserves cofiltered limits. If the semantic transformation $\delta : LP \rightarrow PT$ is injective and surjective on finite sets, then the Jónsson-Tarski Theorem 5.10 holds.*

Proof. The proof is similar to the proof of the preceding Theorem 5.25. Since T preserves finite sets, and since δ, η and ϵ are iso on finite sets, it is easy to see that $\hat{\delta}$ is iso on finite sets. By fixing a filtered diagram $(A_i)_{i \in I}$ representing A and using the fact that L is finitary and that $\mathbf{BA}$ is locally finite, we construct the following diagrams, simplifying the diagrams (5.9)-(5.7):

$$\begin{array}{ccccc}
 & & TUfA_i & \xrightarrow[\simeq]{\hat{\delta}_{A_i}} & UfLA_i \\
 & \nearrow & \downarrow & & \downarrow \\
 TUfA & \xrightarrow{\quad \delta_A \quad} & & & UfLA \\
 & \searrow & \downarrow & & \downarrow \\
 & & TUfA_j & \xrightarrow[\hat{\delta}_{A_j}]{\simeq} & UfLA_j
 \end{array}$$

with each $\hat{\delta}_{A_i}$ an isomorphism. Starting from an element t of $UfLA$, which is the limit of the cofiltered diagram $UfLA_i \rightarrow UfLA_j$, and using the fact that T weakly preserves cofiltered limits, it is clear that we can find an inverse image under $\hat{\delta}_A$ of t . $\square$

Unfortunately, there are many interesting examples of functors T which do *not* preserve finite sets, such as the bag functor or the finite probability distribution functor. This question is addressed in [SP09] via the notion of **strong completeness over finite algebras**. Formally, given a semantic transformation $\delta : LP \rightarrow PT$, the logic defined by L is strongly complete over finite algebras if $\hat{\delta}_A : TUfA \rightarrow UfLA$ is an isomorphism whenever A is finite. Note that for functors T *not* preserving finite sets, this is not equivalent to asking for δ to be an iso on finite *sets* (which is what we required in the previous Theorem). Intuitively, strong completeness over finite algebras allows to approximate the full logic by considering its finite fragments.

Theorem 5.27. *Let $L : \mathbf{BA} \rightarrow \mathbf{BA}$ be finitary, let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ weakly preserve cofiltered limits, and let $\delta : LP \rightarrow PT$ be a semantic transformation such $\hat{\delta}$ is an iso on finite boolean algebras, i.e. the logic defined by L is strongly complete over finite algebras then the Jónsson-Tarski 5.10 Theorem holds.*

Proof. We again proceed as in Proposition 5.24, i.e. we use the fact that L is finitary and that $\mathbf{BA}$ is locally finitely presentable and write any A in $\mathbf{BA}$ as a filtered colimit $A = \text{colim}_{i \in I} A_i$. We then have the following diagram, simplifying the diagrams (5.9)-(5.7), for any $A_i \rightarrow A_j$ in the filtered diagram

representing A :

$$\begin{array}{ccccc}
 & & T\mathbf{Uf}A_i & \xrightarrow[\cong]{\hat{\delta}_{A_i}} & \mathbf{Uf}LA_i \\
 & \nearrow & \downarrow & & \downarrow \\
 T\mathbf{Uf}A & & & \xrightarrow{\hat{\delta}_A} & \mathbf{Uf}LA \\
 & \searrow & \downarrow & & \downarrow \\
 & & T\mathbf{Uf}A_j & \xrightarrow[\cong]{\hat{\delta}_{A_j}} & \mathbf{Uf}LA_j
 \end{array}$$

Since $\mathbf{Uf}LA = \lim_i \mathbf{Uf}LA_i$, every element of $u \in \mathbf{Uf}LA$ defines a coherent family $(u_i)_{i \in I}$, $u_i \in \mathbf{Uf}LA_i$. By the fact that $\hat{\delta}_{A_i}$ is iso (since each A_i is finite) and that the above diagram commutes, this coherent family defines a coherent family $(t_i)_{i \in I}$ with $t_i \in T\mathbf{Uf}A_i$, and thence an element $t \in \mathbf{Uf}TA$ as it is a weak limit of the diagram $(\mathbf{Uf}TA_i)_{i \in I}$. By commutativity of the diagram $\hat{\delta}(t) = u$, and thus $\hat{\delta}$ is surjective. $\square$

Let us show a few examples where the Theorems developed above prove strong completeness.

Example 5.28 (A modal logic for trees with at most n successors). Consider the functor $S_n : \mathbf{Set} \rightarrow \mathbf{Set}$ defined for any $f : X \rightarrow Y$ in $\mathbf{Set}$ by

$$\begin{aligned}
 S_n X &= \coprod_{i=0}^n X^i \\
 S_n f &= [! + f + \dots + f^n]
 \end{aligned}$$

where $X^0 = 1$ the terminal object in $\mathbf{Set}$, $! : 1 \rightarrow 1$, f^k is the k -fold product of f , and $[. + \dots +.]$ denotes the obvious coproduct of maps. We will write in_k for the injections $X^k \rightarrow S_n X$, $1 \leq k \leq n$.

We consider the nabla language $\mathcal{L}_{S_n}$ for this functor which is the initial algebra for the functor $\mathbf{FS}_n \mathbf{U}(-) + \mathbf{FV}$ (see Section 1.5.1 and [DP13] for more details), and the language is thus equivalent to the free BAE defined by a single k -ary operator $\langle k \rangle$ for each $1 \leq k \leq n$. Due to the particularly simple shape of the functor, the axioms of the nabla logic take the form:

$$(\text{Tree 1}) \quad \bigwedge_i \{ \langle k_i \rangle (a_1^i, \dots, a_{k_i}^i) \} = \begin{cases} \langle k \rangle (\bigwedge_i a_1^i, \dots, \bigwedge_i a_k^i) & \text{if } k_i = k \text{ for all } i \\ \perp & \text{else} \end{cases}$$

$$(\text{Tree 2}) \quad \langle k \rangle (\bigvee \phi_1, \dots, \bigvee \phi_k) = \bigvee \{ \langle k \rangle (a_1, \dots, a_k) \mid a_i \in \phi_i, 1 \leq i \leq k \}$$

where $a_j^i \in \mathcal{L}_{S_n}$, $\phi_i \in \mathcal{P}_\omega \mathcal{L}_{S_n}$, $1 \leq i \leq k$. In other words, operator $\langle k \rangle$ distributes over meets and joins in each of its arguments. This last feature is very desirable

as we have seen in Chapter 2, and we will return to the significance of this later on. Using the discussion on negation from [KKV12b] (Section 6), we can also write down the following axiom which is derivable but which is convenient for the calculations that will follow.

$$(\text{Tree 3}) \neg\langle k \rangle(a_1, \dots, a_k) = \bigvee_{1 \leq j \neq k \leq n} \langle j \rangle(\top, \dots, \top) \vee \bigvee_{1 \leq i \leq k} \langle k \rangle(\epsilon_{1i}a_1, \dots, \epsilon_{ki}a_k)$$

where $\epsilon_{ij} = \neg$ if $i = j$ and is the empty symbol otherwise. Using these axioms we can define the functor L defining our logic of trees of degree at most n as the functor $\mathbf{BA} \rightarrow \mathbf{BA}$ defined for any object A and morphism $f : A \rightarrow B$ as

$$\begin{aligned} LA &= (\mathbf{FS}_n \mathbf{U}) / (\text{Tree 1-3}) \\ Lf : LA &\rightarrow LB, [\langle k \rangle(a_1, \dots, a_k)] \mapsto [\langle k \rangle(f(a_1), \dots, f(a_k))] \end{aligned} \quad (5.10)$$

where the quotient is taken under the smallest equivalence relation (in $\mathbf{BA}$) containing the axioms Tree 1-3. This equivalence relation can be explicitly constructed using the method described in Section 1.5.4 via the kernel pair of the coequalizer associated with the set of axioms. For notational clarity we drop the square brackets $[-]$ and simply remember that we are dealing with equivalence classes under the axioms Tree 1-3.

Before moving on to the semantics, we make the following important observation about terms in the language: every term can be expressed in the shape $\bigvee_i \langle k_i \rangle(a_1^i, \dots, a_{k_i}^i)$ where each k_i is different (i.e. $k_i \neq k_j$ when $i \neq j$). To see that this is the case, start with an arbitrary term and re-write it in conjunctive (or disjunctive) normal form. By (Tree 3), we can remove all the innermost negation, modulo the addition of some joins. Now by distributivity, we can re-write the term thus obtained as a join of meets. But by (Tree 1), we can also see that each meet can only be a meet of operators of the same arity, or else we have a bottom element which we can get rid of. Finally, we can use (Tree 1) to push the meets of subterms of the same arity inside the operator, leaving us with a join of the shape $\bigvee_i \langle k_i \rangle(a_1^i, \dots, a_{k_i}^i)$. Finally, we can use (Tree 2) to push the joins of identical k_i s inside the modal operator, leaving us with a join over different k_i s.

The semantic transformation $\delta : L\mathcal{P} \rightarrow \mathcal{PS}_n$ is given by the transpose maps:

$$\begin{aligned} \tilde{\delta}_X : \mathbf{S}_n \mathbf{U} \mathcal{P} X &\rightarrow \mathbf{U} \mathcal{PS}_n X, \\ (U_1, \dots, U_k) &\mapsto \{(x_1, \dots, x_k) \in \text{in}_k[X] \mid x_i \in U_i, 1 \leq i \leq k\} \end{aligned}$$

We can easily check that δ_X is well-defined, i.e. that it is independent of the choice of representative. As an illustration, let us show that it is the case for

(Tree 1), i.e. the distribution of meets.

$$\begin{aligned}
\delta_X(\bigwedge_i \langle k \rangle(U_1^i, \dots, U_k^i)) &= \bigcap_i \delta_X(\langle k \rangle(U_1^i, \dots, U_k^i)) \\
&= \bigcap_i \{ \langle k \rangle(x_1^i, \dots, x_n^i) \mid x_j^i \in U_j^i \} \\
&= \{ \langle k \rangle(x_1, \dots, x_n) \mid x_j \in \bigcap_i U_j^i \} \\
&= \delta_X(\langle k \rangle(\bigwedge_i U_1^i, \dots, \bigwedge_i U_n^i))
\end{aligned}$$

By using the general result that nabla logics are weakly complete, we get that this logic is sound and weakly complete with respect to infinite n -ary trees. But we can show this directly. We introduce a technique which we will use repeatedly in the following examples. The idea is to use the fact that L is finitary and Proposition 5.7 and show that δ is injective on *finite* sets by exhibiting a retraction for δ_X with X finite. We define $r_X : \mathcal{PS}_n X \rightarrow L\mathcal{P}X$ by

$$r_X(\{ \langle k_i \rangle(x_1^i, \dots, x_{k_i}^i) \}_{i \in I}) = \bigvee_k \langle k_i \rangle(\{x_1^i\}_{i|k_i=k}, \dots, \{x_{k_i}^i\}_{i|k_i=k})$$

where the join is taken over all the k s that occur in $\{ \langle k_i \rangle(x_1^i, \dots, x_{k_i}^i) \}_{i \in I}$, this set is finite since $\mathcal{S}_n X$ is finite.

It is easy to check that r_X is a right inverse of δ_X . By the comment above, every term in $LX = \mathcal{FS}_n \mathcal{UP}X$ can be written as $\bigvee_i \langle k_i \rangle(U_1^i, \dots, U_{k_i}^i)$ with every k_i different and with each $U_j^i \in \mathcal{P}X$. It then follows that

$$\begin{aligned}
r_X(\delta_X(\bigvee_i \langle k_i \rangle(U_1^i, \dots, U_{k_i}^i))) &= r_X(\bigcup_i \delta_X(\langle k_i \rangle(U_1^i, \dots, U_{k_i}^i))) \\
&= r_X(\bigcup_i \{ (x_1, \dots, x_{k_i}) \mid x_j \in U_j^i, 1 \leq j \leq k_i \}) \\
&= \bigvee_i \langle k_i \rangle(\{x_1 \mid x_1 \in U_1^i\}, \dots, \{x_{k_i} \mid x_{k_i} \in U_{k_i}^i\}) \\
&= \bigvee_i \langle k_i \rangle(U_1^i, \dots, U_{k_i}^i)
\end{aligned}$$

Thus $r_X \circ \delta_X = \text{Id}_{L\mathcal{P}X}$, and we can conclude that δ_X is injective on finite sets, and thus everywhere from Proposition 5.7 since L is finitary. Weak completeness follows by Theorem 5.6. The adjoint transformation $\hat{\delta} : \mathcal{S}_n \mathcal{Uf} \rightarrow \mathcal{Uf}L$ is given by:

$$\hat{\delta}_A((u_1, \dots, u_k)) = \{ \langle k \rangle(a_1, \dots, a_k) \mid a_i \in u_i, 1 \leq i \leq k \}$$

It is easy to check that $\hat{\delta}_A((u_1, \dots, u_k))$ is an ultrafilter, and that $\hat{\delta}_A$ is injective. The logic is thus expressive.

Let us check that we can use our theorems. The functor $\mathbf{FS}_n\mathbf{U}$ is finitary since $\mathbf{U}$ is monadic, and thus preserves filtered colimits, $\mathbf{S}_n$ is finitary, and $\mathbf{F}$ being a left adjoint preserve all colimits, and in particular filtered ones. To see that $\mathbf{S}_n$ preserves cofiltered limits, note first that each summand in the coproduct is essentially a hom functor $\text{hom}(k, -)$ which preserves all limits, and thus cofiltered ones. As for the coproduct, it is a standard result that coproducts commute with connected limits in **Set**, and since cofiltered limits are connected, we can conclude that $\mathbf{S}_n$ preserves cofiltered limits. Finally, since δ is injective, $\hat{\delta}$ is surjective by Theorem 5.24. We can thus use either Theorem 5.25 or Theorem 5.26 since we have both expressivity and a cofiltered limit preserving functor to conclude that the logic is strongly complete with respect to its semantics. This example exhibits the best possible case: we have an expressive and strongly complete logic. As we shall see, combining these two features is rarely possible.

Example 5.29 (Classical modal logic). Consider the functor $LA : \mathbf{BA} \rightarrow \mathbf{BA}$ defined by

$$LA = \mathbf{F}(\{\Diamond a \mid a \in A\})/(\Diamond \perp = \perp, \Diamond(a \vee b) = \Diamond a \vee \Diamond b) \\ Lf : LA \rightarrow LB, [\Diamond a] \mapsto [\Diamond f(a)]$$

where the quotient is under the smallest equivalence relation (in **BA**) generated by the two equations. The map Lf is well defined by virtue of f being a **BA**-morphism. Note that this is a particular instance of a relational logic, as defined in the previous section, with $\Diamond$ being a unary operator satisfying (Eq. 5.2). For notational clarity, we will from now on drop the square brackets of the equivalence class $[\Diamond a]$ and simply write $\Diamond a$. The functor L defines **classical modal logic**, and is interpreted in P-coalgebras, i.e. Kripke frames, via the semantic transformation $\delta : L\mathcal{P} \rightarrow \mathcal{P}\mathcal{P}$ defined inductively via its action on the generators $\Diamond U$:

$$\delta_X(\Diamond U) = \{V \subseteq X \mid V \cap U \neq \emptyset\}$$

The fact that δ is well-defined is a direct application of Proposition 5.15. We will now show that classical modal logic is strongly complete with respect to its Kripke frame semantics. Of course, we already know this to be true from Theorem 5.16, but the purpose of this example is to illustrate how the theorems developed in this section can be used in practise. An important requirement in the application of Theorems 5.24, 5.25, 5.26 and 5.27 proving strong completeness is that δ should be injective, and in particular the logic should be *weakly* complete (Theorems 5.1 or 5.6). To show that δ is injective, we show that it is injective on *finite* sets, and the general case then follow from Proposition 5.7 since L is finitary and the fundamental adjunction is $\mathbf{U}f \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}^{\text{op}}$. To show that δ_X is injective when X is finite, we will exhibit a retraction r_X , i.e. a left inverse, to δ_X . To define this retraction we follow the following steps.

(Step 1) For every finite set X , any element of LPX can be expressed as

$$\bigvee_i \left(\bigwedge_j \Diamond\{x_j^i\} \wedge \neg\Diamond V_i \right)$$

where each $V_i \subseteq X$ and each $x_j^i \in X$ is such that $x_j^i \notin V_i$. To see that this is the case, first note that for any finite set U the term $\Diamond U$ can be written as $\bigvee_{x \in U} \Diamond\{x\}$ by using the distributivity over joins and the finiteness of U . Thus starting with a term $t \in LPX$, we can rewrite it in such a way that every $\Diamond$ operator is applied to a singleton set. Next, we put this re-written term in disjunctive normal form. If in a clause (i.e. a meet) we find a literal and its negation, then this clause evaluates to $\perp$ and we can get rid of it. Finally, by using the de Morgan law and distributivity once more, we can regroup the negative literals to form the sets V_i and obtain a term of the form above. The fact that $x_j^i \notin V_i$ follows from the previous step.

(Step 2) Let us now consider the image under δ_X of a term without joins, i.e. a set of the form $\delta_X(\bigwedge_{x \in U} \Diamond\{x\} \wedge \neg\Diamond V)$. It is the collection of subsets of X which include U but avoids V , i.e.

$$\delta_X\left(\bigwedge_{x \in U} \Diamond\{x\} \wedge \neg\Diamond V\right) = \{W \subseteq X \mid U \subseteq W \subseteq V^c\}$$

The two inclusions are compatible since, as was shown in the previous step, we can assume $U \cap V = \emptyset$.

(Step 3) To recover U and V from the set above we define for any element $\mathcal{V} \in \mathcal{PPX}$ the following set of pairs of subsets:

$$\mathcal{I}(\mathcal{V}) = \{(Y, Z) \mid Y \subseteq W \subseteq Z^c \Rightarrow W \in \mathcal{V}\}$$

This set forms a poset under component-wise inclusion, and it is easy to check that $(\top, \top) \in \mathcal{I}(\mathcal{V})$ trivially.

(Step 4) Let us write for notational convenience $t = \bigwedge_{x \in U} \Diamond\{x\} \wedge \neg\Diamond V$. We now show that $\mathcal{I}(\delta_X(t))$ has (U, V) as a minimal element. It follows immediately from the definitions that $(U, V) \in \mathcal{I}(\delta_X(t))$. Now assume $(Y, Z) \in \mathcal{I}(\delta_X(t))$, we need to show that $U \subseteq Y$ and $V \subseteq Z$. Assume for the sake of contradiction that $U \not\subseteq Y$ and pick $W = Y$. It is clear that $Y \subseteq W \subseteq Z^c$, but $W \notin \delta_X(t)$ since $U \not\subseteq Y$. Similarly, if we assume that $V \not\subseteq Z$ i.e. $Z^c \not\subseteq V^c$, and pick $W = Z^c$, it is clear that $Y \subseteq W \subseteq Z^c$, but that $W \notin \delta_X(t)$ since $W \not\subseteq V^c$. We can thus recover a term $t = \bigwedge_{x \in U} \Diamond\{x\} \wedge \neg\Diamond V$ from its interpretation $\delta_X(t)$ by taking the minimal element of $\mathcal{I}(\delta_X(t))$.

(Step 5) Let us now show that we can recover a general term of the form detailed above by finding all the minima of $\mathcal{I}(\delta_X(t))$. Assume

$$t = \bigvee_i \left(\bigwedge_{x \in U_i} \Diamond\{x\} \wedge \neg\Diamond V_i \right)$$

We claim that the minima of $\mathcal{I}(\delta_X(t))$ are precisely the pairs (U_i, V_i) . First note that each $(U_i, V_i) \in \mathcal{I}(\delta_X(t))$ by definition. Now let us take an arbitrary $(Y, Z) \in \mathcal{I}(\delta_X(t))$, we need to show that $U_i \subseteq Y$ and $V_i \subseteq Z$ for some i . If $(Y, Z) \in \mathcal{I}(\delta_X(t))$, then in particular we must have $Y \in \delta_X(t)$, and thus there must exist i such that $U_i \subseteq Y \subseteq V_i^c$. Similarly, we must have $Z^c \in \delta_X(t)$, and thus there must exist j such that $U_j \subseteq Z^c \subseteq V_j^c$. If $i = j$ we are done, so let us assume that $i \neq j$, i.e. that we cannot find i such that $U_i \subseteq Y$ and $Z^c \subseteq V_i^c$. In particular this means that we are assuming that $U_j \not\subseteq Y$, i.e. there exists $x \in U_j, x \notin Y$, and $Z^c \not\subseteq V_i^c$, i.e. there exists $y \in Z^c, y \notin V_i^c$. We now have three possibilities:

- $y \in U_j^c$: We then consider $W = Y \cup \{y\}$: it is obvious that $Y \subseteq W$, but since $y \in Z^c$ it is also clear that $W \subseteq Z^c$, since $Y \subseteq Z^c$. Let us now show that $W \notin \delta_X(t)$, which will contradict the assumption that $(Y, Z) \in \mathcal{I}(\delta_X(t))$. It is enough to show that $U_j \not\subseteq W$ and $W \not\subseteq V_i^c$. Since we are assuming both $U_j \not\subseteq Y$ and $y \in U_j^c$ it is clear that $U_j \not\subseteq W$ holds. Similarly, since $y \notin V_i^c$, and $y \in W$, it follows that $W \not\subseteq V_i^c$. Thus $W \notin \delta_X(t)$ and we get the desired contradiction.
- $y \in U_j$ and $x \in V_i^c$: We once again put $W = Y \cup \{y\}$ and show the same contradiction as above. We again have $Y \subseteq W \subseteq Z^c$. Since $y \notin V_i^c$, and $y \in W$, it follows that $W \not\subseteq V_i^c$. To show that $U_j \not\subseteq W$, note first that since $x \in V_i^c$ and $y \notin V_i^c$, we must have $x \neq y$. Since $x \notin Y$ and $x \neq y$ it follows that $x \notin W$, and since $x \in U_j$ it follows that $U_j \not\subseteq W$.
- $y \in U_j$ and $x \in V_i$: This is the difficult case because we cannot immediately discount $x = y$. If we can find two distinct points $x \neq y$ satisfying the conditions above then $W = Y \cup \{y\}$ will lead to the same contradiction as above. For there to be no two distinct points y, x such that $y \in U_j$ and $x \in V_i$ it must be the case that $U_j = V_i = \{x\}$. We now show that this situation can be discounted. For notational convenience we will write $\bigwedge \Diamond U$ for $\bigwedge_{u \in U} \Diamond\{u\}$ and $\Diamond x$ for $\Diamond\{x\}$. We have the following join of

clauses in the expression of t :

$$\begin{aligned} & (\bigwedge \Diamond U_i \wedge \neg \Diamond x) \vee (\Diamond x \wedge \neg \Diamond V_j) = \\ & (\bigwedge \Diamond U_i \vee \Diamond x) \wedge (\bigwedge \Diamond U_i \vee \neg \Diamond V_j) \wedge (\neg \Diamond x \vee \neg \Diamond V_j) = \\ & (\bigwedge \Diamond U_i \wedge \neg \Diamond x) \vee (\bigwedge \Diamond U_i \wedge \neg \Diamond V_j) \vee (\Diamond x \wedge \neg \Diamond V_j) \end{aligned}$$

where the first step is by distributing meets over the join and the second step is by distributing the joins over the meets. Recall that we had assumed that $U_i \subseteq Y \subseteq V_i^c$, we show that $U_i \subseteq Y \subseteq V_j^c$. Assume that there exist $y \in Y, y \notin V_j^c$, i.e. $y \in V_j$. Since we have also assumed that $V_j \subseteq Z$ and Z is disjoint from Y , we get a contradiction. Thus $U_i \subseteq Y \subseteq V_j^c$, which means that modulo the re-writing of t described above we can find pairs $(\tilde{U}_i, \tilde{V}_i^c)$ and $(\tilde{U}_j, \tilde{V}_j^c)$ sandwiching Y and Z^c respectively but without the problematic property that $\tilde{V}_i = \tilde{U}_j = \{x\}$, and we can then proceed as above to find a contradictory W .

(Step 6) We have thus shown that by taking the minimum elements of $\mathcal{I}(\delta_X(t))$ we can (modulo re-writing, which makes no difference since $L\mathcal{P}X$ is in **BA**) always recover t , via the map $r_X : \mathcal{P}\mathcal{P}X \rightarrow L\mathcal{P}X$

$$\mathcal{V} \mapsto \bigvee_{\text{minima } (U,V) \text{ of } \mathcal{I}(\mathcal{V})} \left(\bigwedge_{x \in U} \Diamond \{x\} \wedge \neg \Diamond V \right)$$

which as we have shown is a retraction of δ_X , i.e. δ_X is injective on finite sets.

Since L is finitary can now apply Proposition 5.7 to conclude that δ is actually a mono-transformation, and modal logic is thus weakly complete with respect to its $\mathcal{P}$ -coalgebra, i.e. Kripke frame semantics by Theorem 5.6. Alternatively, we could use the fact that $\mathbf{P}_\omega$ has a terminal coalgebra which in particular is a $\mathbf{P}$ -coalgebra, and use Theorem 5.1. Note that unlike the usual weak completeness proofs (e.g. via tableaux systems, see for example [F⁺72]), we have in effect shown ‘categorical weak completeness’, i.e. a result which is independent of the choice of a base object in **Set**, i.e. of a model carrier. To show strong completeness of modal logic we can use the fact that $\mathbf{P}$ preserves finite sets, that L is finitary and that $\mathbf{P}$ weakly preserves cofiltered limits (see Lemma 5.21) to conclude that modal logic is strongly complete with respect to $\mathbf{P}$ -coalgebra by Theorem 5.24. Note that this only ensures the existence of *quasi*-canonical models.

We can also show that δ is surjective on finite sets. Let X be a finite set and let $\{U_i\}_{1 \leq i \leq n}$ be a collection of subsets of X . For each i , since X is finite,

we can write

$$\{U_i\} = \delta_X \left(\bigwedge_{x_i \in U_i} [\Diamond\{x_i\}] \wedge \neg \bigwedge_{y_i \in \neg U_i} [\Diamond\{y_i\}] \right)$$

i.e. each U_i is the intersection of the collection of subsets which intersect each of its elements and none of the elements of its complement. It follows that

$$\{U_i\}_{1 \leq i \leq n} = \delta_X \left(\bigvee_{i \leq i \leq n} \left[\bigwedge_{x_i \in U_i} [\Diamond\{x_i\}] \wedge \neg \bigwedge_{y_i \in \neg U_i} [\Diamond\{y_i\}] \right] \right)$$

and thus δ_X is surjective. Since it is always injective, we have that δ is bijective on finite sets. We can therefore also prove strong completeness of classical modal logic via Theorem 5.26. Moreover, since $\mathbf{P}$ preserves finite sets it follows from 5.20 that $\hat{\delta}$ is an iso on finite algebras, and the logic is thus strongly complete over finite algebras and strong completeness can be obtained from Theorem 5.27 as well.

It is difficult to give an explicit definition of the adjoint transformation, but the following observation provides us with a way around this difficulty. By definition, LA is the quotient of $F(\{\Diamond a \mid a \in A\})$ under an equivalence relation, i.e. there exist a (regular) epi $q_A : F(\{\Diamond a \mid a \in A\}) \rightarrow LA$. There must therefore exist a mono transformation $\text{Uf}q_A : \text{Uf}LA \rightarrow \text{Uf}F(\{\Diamond a \mid a \in A\})$. But by Lemma 5.4, this is to say that we have a monomorphism $\text{Uf}LA \rightarrow \text{UP}(\{\Diamond a \mid a \in A\})$, i.e. we can view ultrafilters on LA as collections of elements $\Diamond a, a \in A$. Note that not all sets of generators define valid ultrafilters, since the axioms encoded in L must be satisfied. For example, if $\Diamond a$ and $\Diamond b$ belong to an ultrafilter, then so does $\Diamond a \vee \Diamond b$ and thus $\Diamond(a \vee b)$ by the axioms defining L . Modulo this remark, viewing an ultrafilter in terms of the generators it contains allow us to write:

$$\begin{aligned} \hat{\delta}_A(U) &\simeq \{\Diamond a \in LA \mid U \in \delta_{\text{Uf}A}(L\eta_A)(\Diamond a)\} \\ &= \{\Diamond a \in LA \mid U \in \delta_{\text{Uf}A}(\Diamond\eta_A(a))\} \\ &= \{\Diamond a \in LA \mid \exists u \in U. a \in u\} \end{aligned}$$

and it is easy to check that this collection of generators does indeed satisfy the axioms defining L , and does therefore define an element of $\text{Uf}LA$.

The adjoint transformation $\hat{\delta}$ is not injective in general, reflecting the well-known fact that classical modal logic is not expressive (i.e. has not got the Hennessy-Milner property) for arbitrary Kripke frames (see [BdRV01] Section 2.2). However, $\hat{\delta}$ is injective on finite subsets. This can be shown from the surjectivity of δ on finite sets and Lemma 5.20, or directly by using an argument dual to the proof of Theorem 9 in [JS10]. Let U and V be finite sets of ultrafilters

of A such that $U \neq V$. Since V is assumed to be finite, we can write it as $V = \{v_1, \dots, v_n\}$ with $v_i \in \text{Uf}A$, $1 \leq i \leq n$, and since $U \neq V$, there must exist an $u \in U$ such that $u \notin V$, i.e. $\forall i \exists b_i \in u$ such that $b_i \notin v_i$. Consider now the element $a = \Diamond \bigwedge b_i \in LA$. It is easy to check that:

$$\begin{aligned} a \in \hat{\delta}_A(W) &\Leftrightarrow W \in \delta(L\eta_A(\Diamond \bigwedge b_i)) \\ &\Leftrightarrow W \cap L\eta_A(\Diamond \bigwedge b_i) \neq \emptyset \\ &\Leftrightarrow \exists w \in W \forall i, b_i \in w \end{aligned}$$

and thus we have $a \in \hat{\delta}_A(U)$ and $a \notin \hat{\delta}_A(V)$ by construction of a . The fact that the logic is not expressive means that we cannot use Theorem 5.25.

Of course we can also define a right inverse to $\hat{\delta}$ as we did in Theorem 5.16 (this is essentially the content of the traditional canonical model construction) by:

$$\hat{\delta}_A^{-1} : \text{Uf}LA \rightarrow \text{PUf}A, u \mapsto \{v \in \text{Uf}A \mid a \in v \text{ whenever } \Diamond a \in u\}$$

Strong completeness then follows immediately from the Jónsson-Tarski Theorem 5.10.

5.2.4 Strong completeness by semantic completion

We start this subsection by presenting three important cases for which the theorems developed in the previous section do not hold. We will then suggest a general recipe for recovering strong completeness in such cases. The examples are taken from [SP09], but the general method to recover strong completeness is, to the best of our knowledge, new.

Example 5.30 (Modal logic for trees with unbounded branching degree). Consider the functor $S_\omega : \mathbf{Set} \rightarrow \mathbf{Set}$ defined for any $f : X \rightarrow Y$ in $\mathbf{Set}$ by

$$\begin{aligned} S_\omega X &= \coprod_{i \in \omega} X^i \\ S_\omega f &= [! + f + \dots] \end{aligned}$$

where $X^0 = 1$ is the terminal object in $\mathbf{Set}$ and $! : 1 \rightarrow 1$. We consider the unbounded version of the language detailed in Example 5.28, i.e. the free BAE defined by k -ary expansions $\langle k \rangle$ for each $k \in \omega$. The axioms and the semantics are given in the same way as in Example 5.28, but the formula (Tree 3) does not hold since it would involve an infinitary disjunction. The semantic transformation δ is shown to be injective in exactly the same way as in the bounded branching case of Example 5.28, the only subtlety in defining the retraction r_X is to define it on the image of δ_X , since $\mathcal{PS}_\omega X$ is not finite anymore. Elements in $\delta_X[L\mathcal{P}X]$ for X finite are however finite.

The adjoint transpose of the semantic transformation δ is given as in Example 5.28 by:

$$\hat{\delta}_A : S_\omega \text{Uf} A \rightarrow \text{Uf} LA, (u_1, \dots, u_n) \mapsto \{\langle k \rangle(a_1, \dots, a_k) \mid a_i \in u_i, 1 \leq i \leq k\}$$

To see that $\hat{\delta}_A$ cannot be surjective, consider the filter generated by

$$\{\neg \langle k \rangle(\top, \dots, \top) \mid k \in \omega\}$$

This set is consistent, and thus generates an ultrafilter in $\text{Uf} LA$, but there can clearly not be any tuple $(u_1, \dots, u_k)$ mapped to it by $\hat{\delta}_A$.

Example 5.31 (Image finite semantic of classical modal logic). We proceed exactly as in Example 5.29, but we replace the semantic functor $\mathbf{P}$ by its finitary version $\mathbf{P}_\omega$ and the semantic transformation is then defined on generators by the maps:

$$\delta_X : L\mathcal{P}X \rightarrow \mathcal{P}\mathbf{P}_\omega X, [\Diamond U] \mapsto \{V \subseteq_\omega X \mid V \cap U \neq \emptyset\}$$

The properties of δ_X which we showed in Example 5.29 remain true, i.e. δ_X remains injective for each X , and it remains surjective on finite sets. Moreover, as is well known, the logic is now expressive, i.e. $\hat{\delta}$ is injective (see [JS10]). However, $\mathbf{P}_\omega$ does not weakly preserve cofiltered limits, as the following simple example immediately shows. Consider the sets $X_n = \{1, \dots, n\}$ together with the surjective maps

$$p_n : X_{n+1} \rightarrow X_n, i \mapsto \begin{cases} i & \text{if } i < n+1 \\ n & \text{else} \end{cases}$$

It is clear that the X_n and p_n form a cofiltered diagram and that $\lim_n X_n = \mathbb{N}$. A coherent family $(U_n)_{n \in \mathbb{N}}$ of elements $U_n \in \mathbf{P}_\omega X_n$ does not in general define an element of $\mathbf{P}_\omega \lim_n X_n = \mathbf{P}_\omega \mathbb{N}$ since the size of the U_n need not be bounded.

From this failure to preserve cofiltered limits we can show that $\hat{\delta}_A$ cannot be surjective for a general algebra A . Indeed, from the proof of Theorem 5.27, it is clear that if $\hat{\delta}_A$ was surjective, we could find for every coherent family of the cofiltered diagram $(\mathbf{P}_\omega \text{Uf} A_i)_{i \in I}$ (where $\text{colim}_i A_i = A$ is a representation of A as a filtered colimit of finitely presentable boolean algebras) an element of $\mathbf{P}_\omega \text{Uf} A$ projecting onto its components in $\mathbf{P}_\omega \text{Uf} A_i$, i.e. $\mathbf{P}_\omega$ would weakly preserve cofiltered limits. Thus the failure of $\mathbf{P}_\omega$ to weakly preserve cofiltered limits in fact forbids any Tarski-Jónsson embedding, and therefore strong completeness by canonicity or quasi-canonicity. Thus while we recover expressivity (which we did not have in Example 5.29), we lose strong completeness. We will return to this tension between the two concepts.

Example 5.32 (Graded modal logic). We have introduced graded modal logic in Chapter 2, and explained its semantics in terms of the bag functor. Graded modal logic can also be treated in its nabla flavour, as was done in [DP13]. In the abstract style, the syntax is given by the functor $L : \mathbf{BA} \rightarrow \mathbf{BA}$ defined by

$$\begin{aligned} LA &= F(\{\langle k \rangle a \mid k \in \mathbb{N}, a \in A\}) / \mathbf{Ax}_{\text{GML}} \\ Lf : LA &\rightarrow LB, [\langle k \rangle a] \mapsto [\langle k \rangle f(a)] \end{aligned}$$

The quotient in the definition of L is taken under the smallest equivalence relation (in $\mathbf{BA}$) containing the axioms $\mathbf{Ax}_{\text{GML}}$ of graded modal logic. Recall from Chapter 2 that these are given by:

$$\text{GML1 } \langle 1 \rangle (a \vee b) = \langle 1 \rangle (a) \vee \langle 1 \rangle (b)$$

$$\text{GML2 } \langle 1 \rangle \perp = \perp$$

$$\text{GML3 } \Diamond_k a \rightarrow \Diamond_l a, \quad l < k$$

$$\text{GML4 } \Diamond_k a \leftrightarrow \bigvee_{i=0}^k \Diamond_i (a \wedge b) \wedge \Diamond_{k-i} (a \wedge \neg b)$$

$$\text{GML5 } \neg \Diamond_1 (a \wedge \neg b) \rightarrow (\Diamond_k a \rightarrow \Diamond_k b)$$

Once again we will drop the square brackets and simply write the equivalence classes of the generators under $\mathbf{Ax}_{\text{GML}}$ as $\langle k \rangle a$. The semantic transformation $\delta : L\mathcal{P} \rightarrow \mathcal{PB}$ is given by its action on generators:

$$\delta_X : L\mathcal{P}X \rightarrow \mathcal{PB}X, \langle k \rangle U \mapsto \{f \in \mathcal{B}(X) \mid \sum_{x \in U} f(x) \geq k\}$$

which is precisely the definition of the predicate lifting we defined in Chapter 2. It is not very hard to check that δ_X preserves the axioms (GML 1-GML 5). To see that δ_X is injective, we proceed as in Example 5.29 and show that δ_X is injective when X is finite. Since L is finitary and we are working with the adjunction $\text{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}^{\text{op}}$, it will follow that δ_X is injective in general. We exhibit a retraction r_X for X finite, just as we did in Example 5.29, which we construct in several steps:

(Step 1) For every finite set X , any element $t \in L\mathcal{P}X$ can be expressed as:

$$t = \bigvee_i \left(\bigwedge_{j_i} \langle k_{j_i} \rangle \{x_{j_i}\} \wedge \bigwedge_{j'_i} \neg \langle k_{j'_i} \rangle \{y_{j'_i}\} \right)$$

where for a given i we can only have $\langle k_{j_i} \rangle \{x\}$ and $\neg \langle k_{j'_i} \rangle \{x\}$ in the corresponding meets if $k_{j_i} < k_{j'_i}$. To see that this is the case, consider first a term $\langle k \rangle U$ for $U \subseteq X$. Since X is finite, we can write U as

$\{x_1\} \vee \dots \vee \{x_n\}$, which for notational clarity we will abbreviate as $x_1 \vee \dots \vee x_n$. We will now repeatedly use the axiom (GML 4) as follows:

$$\begin{aligned}
\langle k \rangle U &= \langle k \rangle (x_1 \vee \dots \vee x_n) \\
&= \bigvee_{i=0}^k (\langle i \rangle ((x_1 \vee \dots \vee x_n) \wedge x_1^c) \wedge \langle k-i \rangle ((x_1 \vee \dots \vee x_n) \wedge x_1)) \\
&= \bigvee_{i=0}^k (\langle i \rangle (x_2 \vee \dots \vee x_n) \wedge \langle k-i \rangle (x_1)) \\
&= \bigvee_{i=0}^k \left(\bigvee_{j=0}^i \langle j \rangle ((x_2 \vee \dots \vee x_n) \wedge x_2^c) \wedge \langle i-j \rangle ((x_2 \vee \dots \vee x_n) \wedge x_2) \right) \\
&\quad \wedge \langle k-i \rangle (x_1) \\
&= \bigvee_{i=0}^k \left(\bigvee_{j=0}^i \langle j \rangle (x_3 \vee \dots \vee x_n) \wedge \langle i-j \rangle (x_2) \right) \wedge \langle k-i \rangle (x_1) \\
&= \dots
\end{aligned}$$

We keep on proceeding in this way for modalities $\langle k \rangle$ with $k > 1$. As soon as we reach a term $\langle 1 \rangle (x_1, \dots, x_m)$ or $\langle 1 \rangle \emptyset$ we use (GML 1) and (GML 2) to rewrite it as a join a modalities applied to singletons, or as $\perp$. It is clear that in this way we can re-write $\langle k \rangle U$ as a (much more complicated) term involving only singleton sets after the graded modalities. Thus we can assume without loss of generality that an arbitrary term has graded modalities applied to singleton sets only. We then put such a term in disjunctive normal form to get a term of the form announced above. The fact that for a given i we can only have $\langle k_{j_i} \rangle \{x\}$ and $\neg \langle k_{j'_i} \rangle \{x\}$ in the corresponding meets if $k_{j_i} < k_{j'_i}$ is an easy consequence of axiom (GML 3).

(Step 2) We can make an extra assumption on the form of the terms above which will prove useful in the last step of this proof. For a given $x \in X$, we can assume that all the modal operators appearing positively (resp. negatively) in front of x are identical. It is in fact a simple consequence of distributivity. Assume that we have a clause $\langle k_1 \rangle x \wedge \neg \langle k_2 \rangle x \wedge S$ where S stands for the sub-clause not involving x , and that we have another clause $\langle k'_1 \rangle x \wedge \neg \langle k'_2 \rangle x \wedge S'$. Part of the join defining t will thus be

$$(\langle k_1 \rangle x \wedge \neg \langle k_2 \rangle x \wedge S) \vee \langle k'_1 \rangle x \wedge \neg \langle k'_2 \rangle x \wedge S'$$

By applying distributivity back and forth using (GML 3) it is not

difficult to see that this join can be re-written as

$$\begin{aligned} & \langle (\min(k_1, k'_1))x \wedge \neg(\max(k_2, k'_2))x \wedge S \rangle \vee \\ & \langle (\min(k_1, k'_1))x \wedge \neg(\max(k_2, k'_2))x \wedge S' \rangle \end{aligned}$$

i.e. we do have the same modal operators in front of x in both clauses. This clearly generalizes to all clauses, by each time widening the interval of values which x is allowed to take.

(Step 3) Assume a term $t = \bigwedge_j \langle k_j \rangle \{x_j\} \wedge \bigwedge_{j'} \neg \langle k_{j'} \rangle \{y_{j'}\}$. The set $\delta_X(t)$ is the collection of maps $f : X \rightarrow \mathbb{N}$ such that $f(x_j) \geq k_j$ and $f(y_{j'}) < k_{j'}$. To recover the sets $\{x_j\}$ and $\{y_{j'}\}$ and all the integers k_j and $k_{j'}$ for $\delta_X(t)$, we define for any element $\mathcal{V} \in \mathcal{PB}(X)$, the following set of pairs of maps $X \rightarrow \mathbb{N}$:

$$\mathcal{I}(\mathcal{V}) = \{(f, g) \mid f \upharpoonright (\text{supp}(f) \cap \text{supp}(g)) < g \upharpoonright (\text{supp}(f) \cap \text{supp}(g)) \text{ and} \\ \text{if } f \leq h \upharpoonright \text{supp}(f) \text{ and } h \upharpoonright \text{supp}(g) \leq g \text{ then } h \in \mathcal{V}\}$$

where the inequalities are of course defined pointwise. The first condition in the definition ensures that the second is not always vacuously satisfied. We define a partial order on $\mathcal{I}(\mathcal{V})$ as follows:

$$(f_1, g_1) \leq (f_2, g_2) \text{ iff } f_1 \leq f_2 \upharpoonright \text{supp}(f_1) \text{ and } g_1 \geq g_2 \upharpoonright \text{supp}(g_1)$$

Note the reversal of order and the presence of the support in the second argument.

(Step 4) Let us define the maps $f_t, g_t : X \rightarrow \mathbb{N}$ by

$$f_t(x) = \begin{cases} k_j & \text{if } \exists j \text{ s.t. } x = x_j \\ 0 & \text{else} \end{cases} \quad g_t(x) = \begin{cases} k_{j'} & \text{if } \exists j' \text{ s.t. } x = y_{j'} \\ 0 & \text{else} \end{cases}$$

We show that $\mathcal{I}(\delta_X(t))$ has (f_t, g_t) as minimum. Clearly $(f_t, g_t) \in \mathcal{I}(\delta_X(t))$ by definition. Assume that $(f, g) \in \mathcal{I}(\delta_X(t))$ and define

$$h : X \rightarrow \mathbb{N}, x \mapsto \begin{cases} f(x) & \text{if } x \in \text{supp}(f) \\ g(x) & \text{if } x \in \text{supp}(g) \cap \text{supp}(f)^c \end{cases}$$

and

$$\tilde{h} : X \rightarrow \mathbb{N}, x \mapsto \begin{cases} f(x) & \text{if } x \in \text{supp}(f) \cap \text{supp}(g)^c \\ g(x) & \text{if } x \in \text{supp}(g) \end{cases}$$

Note that the conditions $x \in \text{supp}(f) \cap \text{supp}(g)^c$ and $x \in \text{supp}(g) \cap \text{supp}(f)^c$ ensures that h and $\tilde{h}$ are well defined, since the two supports might have a non-empty intersection. It also follows that h coincides

with f on $\text{supp}(f)$, whilst $\tilde{h}$ coincides with g on $\text{supp}(g)$. It follows from these definitions that $f \leq h \upharpoonright \text{supp}(f)$, and $h \upharpoonright \text{supp}(g) \leq g$, and similarly $f \leq \tilde{h} \upharpoonright \text{supp}(f)$, and $\tilde{h} \upharpoonright \text{supp}(g) \leq g$, since f is strictly below g on $\text{supp}(f) \cap \text{supp}(g)$. Since we're assuming that $(f, g) \in \mathcal{I}(\delta_X(t))$, it must be the case that $h, \tilde{h} \in \delta_X(t)$, which by definition means in particular that $f_t \leq h \upharpoonright \text{supp}(f) = f$, and that $g_t \geq \tilde{h} \upharpoonright \text{supp}(g) = g$. We have thus shown that $(f_t, g_t) \leq (f, g)$ for the partial order defined above, and it follows that (f_t, g_t) is indeed the minimum element of $\mathcal{I}(\delta_X(t))$ as claimed.

(Step 5) Let us now consider a general term t as described above, and for each i define f_i and g_i as above, i.e.

$$f_i(x) = \begin{cases} k_j^i & \text{if } \exists j \text{ s.th. } x = x_j^i \\ 0 & \text{else} \end{cases} \quad g_i(x) = \begin{cases} k_{j'}^i & \text{if } \exists j' \text{ s.th. } x = y_{j'}^i \\ 0 & \text{else} \end{cases}$$

We claim that the minima of $\mathcal{I}(\delta_X(t))$ are the pairs (f_i, g_i) . Clearly, each such pair belongs to $\mathcal{I}(\delta_X(t))$ by definition. Now, assume $(f, g) \in \mathcal{I}(\delta_X(t))$ and define $h^{(0)}$ as we defined h in the previous step. By definition $f \leq h^{(0)} \upharpoonright \text{supp}(f)$ and $h^{(0)} \upharpoonright \text{supp}(g) \leq g$, and thus we must have $h^{(0)} \in \delta_X(t)$, i.e. there must exist i_0 s.th. $f_{i_0} \leq h^{(0)} \upharpoonright \text{supp}(f_{i_0})$ and $h^{(0)} \upharpoonright \text{supp}(g_{i_0}) \leq g_{i_0}$. If we also have $g \upharpoonright \text{supp}(g_{i_0}) \leq g_{i_0}$, then we are done, since we will have shown $(f_{i_0}, g_{i_0}) \leq (f, g)$. If $g \upharpoonright \text{supp}(g_{i_0}) \not\leq g_{i_0}$, then there must exist $x_{i_0} \in \text{supp}(f) \cap \text{supp}(g)$ such that $g_{i_0}(x_{i_0}) < g(x_{i_0})$ (x_{i_0} cannot belong to $\text{supp}(f)^c \cap \text{supp}(g)$, or we would contradict that $h^{(0)} \upharpoonright \text{supp}(g_{i_0}) \leq g_{i_0}$). Now define

$$h^{(1)}(x) = \begin{cases} h^{(0)}(x) & \text{if } x \neq x_{i_0} \\ g(x_{i_0}) & \text{else} \end{cases}$$

By construction $h^{(1)} \upharpoonright \text{supp}(g_{i_0}) \not\leq g_{i_0}$, and moreover $f \leq h^{(1)} \upharpoonright \text{supp}(f)$ and $h^{(1)} \upharpoonright \text{supp}(g) \leq g$ (since the only change from $h^{(0)}$ is $g(x_{i_0})$). Thus if $(f, g) \in \mathcal{I}(\delta_X(t))$ we must have $h^{(1)} \in \delta_X(t)$, and thus there must exist i_1 such that $f_{i_1} \leq h^{(1)} \upharpoonright \text{supp}(f_{i_1})$ and $h^{(1)} \upharpoonright \text{supp}(g_{i_1}) \leq g_{i_1}$. We have established earlier that we can assume without loss of generality that f_{i_0} and f_{i_1} coincide on the intersection of their supports, and similarly for g_{i_0} and g_{i_1} . It follows that $f_{i_1} \leq f \upharpoonright \text{supp}(f_{i_1})$, indeed on $\text{supp}(f)$ $h^{(1)}$ only differs from f at x_{i_0} and if $x_{i_0} \in \text{supp}(f_{i_1})$ then

$$f_{i_1}(x_{i_0}) = f_{i_0}(x_{i_0}) \leq h^{(0)}(x_{i_0}) = f(x_{i_0})$$

Thus if $g \upharpoonright \text{supp}(g_{i_1}) \leq g_{i_1}$ we are done, and we will have shown $(f_{i_1}, g_{i_1}) \leq (f, g)$. Otherwise, we iterate our construction: assume

that $g \upharpoonright \text{supp}(g_{i_1}) \not\leq g_{i_1}$, then there exist $x_{i_1} \in \text{supp}(f) \cap \text{supp}(g)$ such that $g_{i_1}(x_{i_1}) < g(x_{i_1})$. We now define

$$h^{(2)}(x) = \begin{cases} h^{(1)}(x) & \text{if } x \neq x_{i_1} \\ g(x_{i_1}) & \text{else} \end{cases}$$

This map has the property that $h^{(2)} \upharpoonright \text{supp}(g_{i_0}) \not\leq g_{i_0}$ and $h^{(2)} \upharpoonright \text{supp}(g_{i_1}) \not\leq g_{i_1}$ (since we include two counter-example points in its image), but also $f \leq h^{(2)} \upharpoonright \text{supp}(f)$ and $h^{(2)} \upharpoonright \text{supp}(g) \leq g$. Thus if $(f, g) \in \mathcal{I}(\delta_X(t))$ we must have $h^{(2)} \in \delta_X(t)$. We then repeat the procedure detailed above. It is clear that at each step n we either find a pair $(f_{i_n}, g_{i_n}) \leq (f, g)$ and we have shown what we wanted, or else we construct a new map $h^{(n)}$ with the property that $h^{(k)} \upharpoonright \text{supp}(g_{i_k}) \not\leq g_{i_k}$, for every $1 \leq k \leq n$, i.e. $(f_{i_k}, g_{i_k}) \not\leq (f, g)$ for every $1 \leq k \leq n$. Since there is a finite number p of clauses in our term, this procedure must stop with either a pair $(f_{i_k}, g_{i_k}) \leq (f, g)$ (where $0 \leq k \leq p$) or with a map $h^{(p)}$ partially sandwiched between f and g with the property that $(f_{i_k}, g_{i_k}) \not\leq (f, g)$ for every $1 \leq k \leq p$, and thus $h^{(p)} \notin \delta_X(t)$, which contradicts $(f, g) \in \mathcal{I}(\delta_X(t))$.

(Step 6) We have thus exhibited the following retraction $r_X : \mathcal{PB}(X) \rightarrow L\mathcal{P}X$, defined by:

$$\mathcal{V} \mapsto \bigvee_{(f,g) \text{ minima of } \mathcal{I}(\mathcal{V})} \left(\bigwedge_{x \in \text{supp } f} \langle f(x) \rangle \{x\} \wedge \bigwedge_{y \in \text{supp } g} \neg \langle g(y) \rangle \{y\} \right)$$

and thus shown that δ_X is injective on finite sets, and thus everywhere by 5.7.

Since $\mathcal{B}$ is finitary and **Set** is locally presentable, we know from (see Theorem 4.2.12. of [AMM10]) that it has a terminal coalgebra, and we thus get weak completeness of GML with respect to its multiset semantics via Theorem 5.1. However, since $\mathcal{B}$ does *not* preserve finite sets, we cannot hope to prove strong completeness from weak-completeness via Theorems 5.24-5.26. Let us now see if Theorem 5.27 can be used. To explicitly characterise the adjoint transformation $\hat{\delta} : \mathcal{BUf} \rightarrow \mathcal{Uf}L$, we use the same observation as in Example 5.29, i.e. that LA is a quotient of a free algebra, and thus ultrafilters on LA can be seen as sets of generators $\langle k \rangle a, a \in A$ (again not all sets of generators define an ultrafilter since the axioms encoded by L must hold). We can then write

$$\begin{aligned} \hat{\delta}(f) &\simeq \{ \langle k \rangle a \in LA \mid f \in \delta_{\mathcal{Uf}A}((L\eta_A)(\langle k \rangle a)) \} \\ &= \{ \langle k \rangle a \in LA \mid f \in \delta_{\mathcal{Uf}A}(\langle k \rangle \eta_A(a)) \} \\ &= \{ \langle k \rangle a \in LA \mid \sum_{u|a \in u} f(u) \geq k \} \end{aligned}$$

Note that since $\mathcal{B}$ does *not* preserve finite sets, we cannot use Lemma 5.17 to infer properties of $\hat{\delta}$ on finite algebras from properties of δ on finite sets. In particular, the fact that δ is injective, does *not* imply that $\hat{\delta}$ should be surjective on finite algebras. We can in fact explicitly check that $\hat{\delta}$ is *not* surjective, even on finite algebras. To see this, let A be a finite boolean algebra and consider any ultrafilter u containing the set of generators $\{\langle k \rangle \top \mid k \in \mathbb{N}\}$. It is clear from the above characterisation that any such u cannot have a pre-image since such a map $f : \mathcal{B} \text{Uf} A \rightarrow \mathbb{N}$ would then have to satisfy $\sum_{u \in \text{supp}(f)} f(u) \geq k$ for each k . Since the support of each f is bounded we get a contradiction, thus $\hat{\delta}$ fails to be surjective, and in particular we cannot use Theorem 5.27 to strengthen weak completeness to strong completeness.

Note also $\mathcal{B}$ does not preserve cofiltered limits. Consider once again the sets X_n and the maps p_n defined in the previous Example 5.31. A coherent family $f_n \in \mathcal{B}(X_n)$ does not in general define an element of $\mathcal{B}(\lim_n X_n) = \mathcal{B}(\mathbb{N})$ for the same reason as above: whilst every element of $\mathcal{B}(\mathbb{N})$ must have a finite support, the family f_n might have unbounded supports.

What characterises these three examples is that the functor T providing the semantic domain is ‘too small’ to accommodate the logic. Alternatively and equivalently, the functor L defining the logics is too large in that it allows total logical descriptions of states (i.e. ultrafilters) which cannot be realized by T -models. We present a generic technique to try to solve this problem which essentially consists of enlarging the semantic domain by ‘completing’ the functor T in an appropriate fashion. We therefore call this procedure **semantic completion**. This technique seems to work well in cases where the only obstacle to applying Theorem 5.27 is the non weak preservation of cofiltered limits by the semantic functor.

The finitary version of a functor, revisited. Recall that the finitary version of a set functor $T : \mathbf{Set} \rightarrow \mathbf{Set}$ was defined as

$$T_\omega : \mathbf{Set} \rightarrow \mathbf{Set}, X \mapsto \bigcup_{U \subseteq_\omega X} TU$$

Following our work in Chapter 3, we know from Proposition 3.36, that if a functor T is finitary and if we denote by I the inclusion functor $\mathbf{l} : \mathbf{Set}_f \rightarrow \mathbf{Set}$, we have

$$T = \text{Lan}_I(T \circ \mathbf{l})$$

More generally, as was shown in Propositions 3.46 and 3.70, if T is any functor then T_ω is just the left Kan extension of $T \circ I$ w.r.t. the inclusion functor $\mathbf{l} : \mathbf{Set}_f \rightarrow \mathbf{Set}$ of the category of finite sets into the category of sets. Formally

$$T_\omega = \text{Lan}_I(T \circ \mathbf{l})$$

The functors $\mathcal{B}$ and $\mathcal{P}_\omega$ which lead to the failures of strong completeness in the examples above can be seen either as finitary functors or as finitary versions of more general functors. The finitary version of a functor clearly preserves filtered colimits and we present here a very simple proof.

Proposition 5.33. *Let $\mathbf{C}$ be a locally λ -presentable category, let $\mathbf{l} : \mathbf{C}_\lambda \rightarrow \mathbf{C}$ be the inclusion of the subcategory of λ -presentable objects, and let $T : \mathbf{C} \rightarrow \mathbf{C}$ be any functor. The λ -ary version of the functor given by*

$$T_\lambda = \text{Lan}_{\mathbf{l}}(T \circ \mathbf{l})$$

preserves λ -filtered colimits.

Proof. Since $\mathbf{C}$ is locally λ -presentable, it is in particular cocomplete, and the left Kan extension does indeed exist. Moreover it can be computed using the following coend formula (see [Mac98] X.4):

$$\text{Lan}_{\mathbf{l}}(T \circ \mathbf{l})A = \int^B \text{hom}(\mathbf{l}B, A) \cdot (T \circ \mathbf{l})B$$

where $-\cdot (T \circ \mathbf{l})B$ denotes the copower functor. Since a coend is a colimit, it will preserve the colimits preserved by the composition of $\text{hom}(\mathbf{l}B, -)$ and $-\cdot (T \circ \mathbf{l})B$. From the fact that

$$-\cdot (T \circ \mathbf{l})B \dashv \text{hom}((T \circ \mathbf{l})B, -)$$

it is clear that the copower functor preserves all colimits, and in particular cofiltered ones. So we need only show that $\text{hom}(\mathbf{l}B, -)$ preserves λ -filtered colimits, but this follows immediately from the fact that B is a λ -presentable object. $\square$

The co-finitary version of a functor. By dualising the above procedure, we can create a version of any **Set**-endofunctor which weakly preserves cofiltered limits, which we will call the **co-finitary version** of the functor, for lack of a better name. Note that what we present now is *not* the dual version of the finitary functor, it simply uses dual tools, namely the right Kan extension. The best property we will be able to show for the co-finitary version of a functor is *weak* preservation of cofiltered limits in the category of **Set**, which is a much more restricted result than Proposition 5.33, but sufficient for our purpose.

Definition 5.34. A functor $G : \mathbf{A} \rightarrow \mathbf{B}$ is called **initial** if for any object B of $\mathbf{B}$ the comma category $G \downarrow B$ is non-empty and connected, i.e. if the following two conditions holds for any B

- (i) there exists an object A of $\mathbf{A}$ and a morphism $GA \rightarrow B$

- (ii) any two objects in $B \downarrow G$, i.e. any two morphisms $GA \rightarrow B$ and $GA' \rightarrow B$, are related by a zigzag of morphisms.

Initial functors are used to compute limits as cofinal functors are used to compute colimits.

Theorem 5.35. *If $G : \mathbf{A} \rightarrow \mathbf{B}$ is initial and $\mathcal{D} : \mathbf{B} \rightarrow \mathbf{C}$ is a diagram such that $\lim \mathcal{D}G$ exists, then $\lim \mathcal{D} \simeq \lim \mathcal{D}G$.*

Proof. Dual to the result on cofinal functors (see [Mac98], Chapter IX, section 3, Theorem 1). $\square$

The notion of **right Kan extension** is defined dually to that of left Kan extension and we refer the reader to Section 3.3 or to see [Mac98] X.3. for all the definitions. For our purpose, the following result will be sufficient.

Proposition 5.36. *Let $\mathbf{B} \xleftarrow{F} \mathbf{A} \xrightarrow{G} \mathbf{C}$ such that $\mathbf{A}$ is small and $\mathbf{B}$ is complete, then F has a right Kan extension $\text{Ran}_G F$ along any G .*

Proof. See [Mac98] X.3.2. $\square$

Thus if we consider the inclusion $\mathbf{l} : \mathbf{Set}_f \rightarrow \mathbf{Set}$, the right Kan extension $\text{Ran}_{\mathbf{l}}(T \circ \mathbf{l})$ for any $T : \mathbf{Set} \rightarrow \mathbf{Set}$ always exists. Its explicit construction is given by

$$\text{Ran}_{\mathbf{l}}(T \circ \mathbf{l})X = \lim \left((X \downarrow \mathbf{l}) \xrightarrow{Q} \mathbf{Set}_f \xrightarrow{T \circ \mathbf{l}} \mathbf{Set} \right)$$

where Q is the obvious projection functor. We will call $\text{Ran}_{\mathbf{l}}(T \circ \mathbf{l})$ the co-finitary version of T and denote it by T^ω . We can in fact restrict the index category of this limit as the following Proposition shows.

Proposition 5.37. *Let X be a set, and let us consider the subcategory of $X \downarrow \mathbf{Set}_f$ whose objects consists of surjections $X \twoheadrightarrow n, n \in \mathbb{N}$ and whose morphisms between $g : X \twoheadrightarrow n, h : X \twoheadrightarrow p$ are surjections $f : n \twoheadrightarrow p$ such that $h = f \circ g$. Let us call this category $\mathbf{J}$, and let I denote the obvious inclusion $\mathbf{J} \rightarrow X \downarrow \mathbf{Set}_f$, then I is initial.*

Proof. We need to show that for any $g : X \rightarrow n$ there exist an arrow $I(h) \rightarrow g$ (in $X \downarrow \mathbf{Set}_f$) for some object h of $\mathbf{J}$. For this, we consider the strong epi-mono factorisation of g as $g = m \circ e$ and get the morphism $m : I(e) \rightarrow g$ in $X \downarrow \mathbf{Set}_f$:

$$\begin{array}{ccc} & X & \\ e \swarrow & & \searrow g \\ n' & \xrightarrow{m} & n \end{array}$$

For the second condition, we need to show that given $e : X \twoheadrightarrow p$ and $e' : X \twoheadrightarrow p'$ such that there exist $g : X \rightarrow n$ and morphisms $h : p \rightarrow n, h' : p' \rightarrow n$ such that

$h \circ e = h' \circ e' = g$, we have a zigzag of morphism connecting e and e' . For this we take an strong epi-mono factorisation of g as $g = m \circ q$. Using the fact that the epimorphisms e and e' are strong we get fill-in morphisms u, u' which make the following diagram commute and provide us with the zigzag of morphisms:

$$\begin{array}{ccccc}
 & & X & & \\
 & e \swarrow & \downarrow q & \searrow e' & \\
 p & \xrightarrow{u} & n' & \xleftarrow{u'} & p' \\
 & h \searrow & \downarrow m & \swarrow h' & \\
 & & n & &
 \end{array}$$

Since strong epimorphisms have the property that if $u \circ e$ is a strong epimorphism, then u must be a strong epimorphism, u, u' indeed provide the zigzag of morphism in **J** connecting e and e' . $\square$

Let us now show that the co-finitary version of a functor does what we want it to do, namely weakly preserve cofiltered limits.

Proposition 5.38. *Let $T : \mathbf{Set} \rightarrow \mathbf{Set}$ be an endofunctor, let $\mathbf{l} : \mathbf{Set}_f \rightarrow \mathbf{Set}$ be the inclusion functor of the subcategory of finite sets, and let T^ω be defined as $\text{Ran}_{\mathbf{l}}(T \circ \mathbf{l})$, then T^ω weakly preserves cofiltered limits.*

Proof. We write $\text{Ran}_{\mathbf{l}}(T \circ \mathbf{l})A$ as the end (see [Mac98] X.4)

$$\text{Ran}_{\mathbf{l}}(T \circ \mathbf{l})X = \int_n (T \circ \mathbf{l})n^{\text{hom}(X, \mathbf{l}n)}$$

and use an argument totally analogous to Proposition 5.33: since an end is a limit, it weakly preserves the limits preserved by the composite of the contravariant functors $\text{hom}(-, \mathbf{l}n)$ and $(T \circ \mathbf{l})n^{(-)}$ where the latter is the power functor. Since

$$(T \circ \mathbf{l})n^{(-)} \vdash \text{hom}(-, (T \circ \mathbf{l})n)$$

the power functor (which is contravariant) turns all colimits into limits, and in particular it turns (weak) filtered colimits into (weak) cofiltered limits. Thus we need to show that $\text{hom}(-, \mathbf{l}n)$ weakly preserves cofiltered limits, i.e. that it turns cofiltered limits into weak filtered colimits. This is precisely what we have shown in Lemma 5.22. $\square$

We now have a uniform way to turn a **Set**-endofunctor into one that weakly preserves cofiltered limits, and is thus amenable to Theorems 5.25, 5.26 and 5.27. Note that in general, this will result in loosing expressivity.

Example 5.39 (Modal logic for trees with unbounded branching degree). It is easy to see that the functor S_ω does not weakly preserve cofiltered limits by considering a sequence of increasingly ‘large’ n -tuples of elements of X , whose limit cannot be an n -tuple, but a map $\mathbb{N} \rightarrow X$. Indeed this is what makes $\hat{\delta}$ non-surjective, even on finite algebras. We apply the construction above to remedy this situation:

$$\begin{aligned} S_\omega^\omega(X) &= \text{Ran}_1(S_\omega)(X) \\ &= \lim \left((X \downarrow 1) \xrightarrow{Q} \mathbf{Set}_f \xrightarrow{S_\omega} \mathbf{Set} \right) \\ &\simeq \{f : \mathbb{N} \rightarrow X\} = \text{hom}(\mathbb{N}, X) \end{aligned}$$

where finite sequences of elements of X are identified with functions $f : \mathbb{N} \rightarrow X$ which eventually become constant. We can interpret the language defined in Example 5.30 in S_ω^ω -coalgebras in exactly the same way as we did in S_ω -coalgebras, and the semantic transformation remains injective. The key difference is that the adjoint transpose $\hat{\delta} : S_\omega^\omega \mathbf{Uf} \rightarrow \mathbf{Uf} L$ is now surjective. For any A in $\mathbf{BA}$, define the following section, $s_A : \mathbf{Uf} L A \rightarrow S_\omega^\omega \mathbf{Uf} A$

$$s_A(t) = \begin{cases} \langle k \rangle(u_1, \dots, u_k) \text{ s.th. } a_i \in u_i, 1 \leq i \leq k \text{ iff } \langle k \rangle(a_1, \dots, a_k) \in t & \text{if } \exists k, \langle k \rangle(\top, \dots, \top) \in t \\ ((\{\neg \langle k \rangle(\top, \dots, \top) \mid k \in \omega\})^\uparrow, \dots) & \text{else} \end{cases}$$

where $\langle - \rangle^\uparrow$ denotes the ultrafilter generated by the set. Note that from the axiom (Tree 1), it is clear that if there exist a k with $\langle k \rangle(\top, \dots, \top) \in t$, then it is unique. It is easy to see that s_A is indeed a section of $\hat{\delta}_A$ (see Example 5.28 for a definition of $\hat{\delta}$), which is thus surjective. Since it is surjective on any boolean algebra, we can apply the Jónsson-Tarski Theorem 5.10 directly to conclude that the logic defined by $\mathbf{FS}_\omega \mathbf{U}$ and the axioms (Tree 1), (Tree 2) is strongly complete with respect to S_ω^ω -coalgebras.

However, the definition of s_A above suggests that S_ω^ω is ‘too big’ a functor. Indeed, most of the elements gained by enlarging S_ω to S_ω^ω are not really necessary. In fact a much more economical alternative (and much more fruitful as we shall see at the end of the Chapter) is to extend S_ω by a singleton, i.e. we define

$$S_\omega^* : \mathbf{Set} \rightarrow \mathbf{Set}, S_\omega^* X = S_\omega X + 1$$

We can interpret L -algebras in S_ω^* -coalgebras exactly as in Examples 5.28 and 5.30, and δ remains injective. Moreover, by mimicking the definition above, we can show that the adjoint transpose $\hat{\delta} : S_\omega^* \mathbf{Uf} \rightarrow \mathbf{Uf} L$ is surjective by defining

the section

$$s_A(t) = \begin{cases} \langle k \rangle(u_1, \dots, u_k) \text{ s.th. } a_i \in u_i, 1 \leq i \leq k & \text{iff } \langle k \rangle(a_1, \dots, a_k) \in t \\ \text{if } \exists k, \langle k \rangle(\top, \dots, \top) \in t & \\ * & \text{else} \end{cases}$$

where $*$ is the new singleton in the coproduct defining S_ω^* . Since there exists only one problematic ultrafilter, namely the one generated by $\{\neg \langle k \rangle(\top, \dots, \top) \mid k \in \omega\}$, it is clear that $\hat{\delta}_A \circ s_A = \text{Id}_{\text{Uf}LA}$. It follows from the Jónsson-Tarski Theorem 5.10 that the logic defined by L is strongly complete with respect to S_ω^* -coalgebras.

Example 5.40. (Countable semantics of classical modal logic) Recall from Examples 5.29 and 5.31 that we have the traditional semantic of classical modal logic which is strongly complete but not expressive and given in terms of the covariant powerset functor P , and the image finite semantics which is expressive but not strongly complete and given in terms of the finite powerset functor P_ω . If we apply the construction above to this functor we get (by using Proposition 5.37) that

$$\begin{aligned} P^\omega(X) &= \text{Ran}_I(P_\omega)(X) \\ &= \lim \left((X \downarrow I) \xrightarrow{Q} \mathbf{Set}_f \xrightarrow{P_\omega} \mathbf{Set} \right) \\ &= P_c(X) \end{aligned}$$

where $P_c(X)$ is the set of countable subsets of X . We can keep the definition of the semantic transformation δ from Example 5.29 since $P_\omega(X) \subseteq P_c(X)$ for all sets X . Moreover, since P_c preserves cofiltered limits by Proposition 5.38, and since the properties of the semantic transformation and its adjoint are the same as in Example 5.29, we can also use Theorem 5.27 to conclude that P_c provide a strongly complete semantic to classical modal logic. Again, the logic fails to be expressive.

Example 5.41. (Graded Modal Logic) We apply the co-finitary construction to the finitary bag functor which has provided the semantic for Graded Modal Logic in Example 5.32. There were two obstacles to strong completeness in this example: (1) the finitary bag functor $\mathcal{B}$ does not preserve cofiltered limits, and (2) the adjoint semantic transformation $\hat{\delta}$ is not surjective on finite sets. We will now show that by fixing the first problem via the co-finitary construction, we also fix the second. Let us first compute $\mathcal{B}^\omega$:

$$\begin{aligned} \mathcal{B}^\omega(X) &= \text{Ran}_I(\mathcal{B})(X) \\ &= \lim \left((X \downarrow I) \xrightarrow{Q} \mathbf{Set}_f \xrightarrow{\mathcal{B}} \mathbf{Set} \right) \\ &= \{f : X \rightarrow \mathbb{N} \cup \{\infty\} \mid \text{supp}(f) \subseteq_c X\} \end{aligned}$$

where $\subseteq_c$ means ‘is a countable subset of’. Since $\mathcal{B}(X) \subseteq \mathcal{B}^\omega(X)$ for any set X , we can keep the definition of the semantic transformation δ from Example 5.32. To see that $\mathcal{B}^\omega$ addresses the second problem raised above, we show that $\hat{\delta}$ is an isomorphism on finite boolean algebras. Since δ is essentially unchanged, we can recycle $\hat{\delta} : \mathcal{B}^\omega \mathbf{Uf}A \rightarrow \mathbf{Uf}LA$ too, which as we have shown above is defined by:

$$\hat{\delta}(f) = \{ \langle k \rangle b \mid \sum_{u \in \eta_A(b)} f(u) \geq k \}$$

where η is the unit of the adjunction $\mathbf{Uf} \dashv \mathcal{P}$. For a finite boolean algebra A , every ultrafilter on A is principal, so let us denote by a_u the elements of A which define an ultrafilter $u = \uparrow a_u$. A map $f : \mathbf{Uf}A \rightarrow \mathbb{N} \cup \{\infty\}$ can be seen as a map $f : \{a_u\}_{u \in \mathbf{Uf}A} \rightarrow \mathbb{N} \cup \{\infty\}$ and we then have:

$$\hat{\delta}_A(f) = \{ \langle k \rangle b \mid \sum_{a_u \leq b} f(a_u) \geq k \}$$

It is straightforward to check that this map is injective. To see that it is surjective, we define a section $s : \mathbf{Uf}LA \rightarrow \mathcal{B}^\omega \mathbf{Uf}A$ of $\hat{\delta}_A$ as follows. Recall that an ultrafilter of LA is determined uniquely by the generators $\langle k \rangle b, b \in A$ it contains. Thus we will alternatively write $w \in \mathbf{Uf}LA$ as a set $\{\langle k_i \rangle b_i \mid i \in I\}$ of generators of LA . We now define

$$s(w) = \lambda a_u. \begin{cases} \max(\{\langle k_i \rangle \mid \exists i \in I. b_i = a_u\} \cup \{0\}) & \text{if it exists} \\ \infty & \text{else} \end{cases}$$

Note that since the k_i need not be bounded, we do need the $\{\infty\}$ option provided by $\mathcal{B}^\omega$ for this map to be well defined. Let us check that it is indeed a section. We start with a collection $\{\langle k_i \rangle a_i\}_{i \in I}$ defining an ultrafilter in LA and for notational clarity we write $f_w = s(w)$ with w described by the set $\{\langle k_i \rangle a_i \mid i \in I\}$. We need to show that $\hat{\delta}_A(f_w) = w$ i.e. that the sets of generators describing $\hat{\delta}_A(f_w)$ and w are equal, in other words that

$$\{ \langle l \rangle b \mid \sum_{a_u \leq b} f_w(a_u) \geq l \} = \{ \langle k_i \rangle a_i \}_{i \in I}$$

Let us first show the inclusion from right to left, i.e. that each $\langle k_i \rangle a_i \in \hat{\delta}_A(f_w)$. We need the following lemma.

Lemma 5.42. *In the situation described above, if $\langle k \rangle a$ belongs to an ultrafilter of $\mathbf{Uf}LA$, then this ultrafilter must also contain a finite set of terms $\{\langle k_j \rangle a_{u_j}\}_{1 \leq j \leq m}$, $u_j \in \mathbf{Uf}A$, such that*

$$\sum_{j=1}^m k_j \geq k$$

Proof. The first step is to show that there exist elements of the type a_u below a . So let us assume that $\uparrow a$ is not a principal ultrafilter, i.e. there must exist a finite set $\{a_i \mid 1 \leq i \leq n\}$ of elements of A such that neither a_i nor $\neg a_i$ belong to $\uparrow a$. Each selection function

$$h : n \mapsto \{a_i\}_{1 \leq i \leq n} \cup \{\neg a_i\}_{1 \leq i \leq n}$$

of either a_i or $\neg a_i$ for each $1 \leq i \leq n$ completes the proper filter $\uparrow a$ into an ultrafilter $u_h = \uparrow(a \wedge \bigwedge_i h(i)) = \uparrow a_{u_h}$, and $a_{u_h} \leq a$. Let us write H for the set of all such selection functions. We have,

$$\bigvee_{h \in H} a_{u_h} = \bigvee_{h \in H} (a \wedge \bigwedge_i h(i)) = a \wedge \bigvee_{h \in H} \bigwedge_i h(i) = a \wedge \top = a$$

Let us now focus on the operators. There are two possibilities, either $\langle k \rangle a$ is the ‘largest’ such formula in the ultrafilter, or there is a larger integer $k' > k$ such that $\langle k' \rangle a$ is in the ultrafilter. In the first case, since $a = \bigvee_h a_{u_h}$ and since $a_{u_h} \wedge a_{u_{h'}} = \perp$ whenever $h \neq h'$ we can use an argument totally analogous to Step 1 of Example 5.32 (using axioms GML 1-2, 4) to show that there must exist a partition of k into a collection $\{k_h\}_{h \in H_0}$ (i.e. the k_h s add up to k) where $H_0 \subseteq H$ such that for every $h \in H_0$, $\langle k_h \rangle a_{u_h}$ is in the ultrafilter, and the result follows.

If there exists $k' > k$ such that $\langle k' \rangle a$ is in the ultrafilter then there are two possibilities: either there exists a largest such k' , in which case we can apply the same reasoning as above to $\langle k' \rangle a$, or there is no largest such k' in which case there are also terms of the type $\langle k_h \rangle a_{u_h}$ for arbitrary large k_h s by the same argument as above instantiated at each k' , in which case we will have $\infty \geq k$. $\square$

From this Lemma it follows that to any $\langle k_i \rangle a_i$ there must exist a collection of formulas $\{\langle k_{i_j} \rangle a_{u_j}\}_{1 \leq j \leq m}$ in the ultrafilter, where each a_{u_j} defines a principal ultrafilter. From the definition of f_w , it then follows that

$$\sum_{a_u \leq a_i} f_w(a_u) \geq \sum_{j=1}^m k_{i_j} \geq k_i$$

We therefore have $\langle k_i \rangle a_i \in \hat{\delta}_A(f_w)$.

Let us now show the opposite inclusion. If $\langle l \rangle b \in \hat{\delta}_A(f_w)$, then by definition there exist a (finite) collection of elements $a_{u_k} \leq b, 1 \leq k \leq p$ such that $f_w(a_{u_k}) > 0$ and $\sum_k f_w(a_{u_k}) \geq l$. By definition of f_w , this means that for each k , $a_{u_k} = a_i$ for some $i \in I$, and since f_w picks the largest integer l_k associated with a_{u_k} , this means that $\langle l_k \rangle a_{u_k}$ is in the ultrafilter. The condition that $\langle l \rangle b \in \hat{\delta}_A(f_w)$ takes the form $\sum_k l_k \geq l$. It is easy to see that for any ultrafilter $u, v \in \text{Uf}A$, $a_u \wedge a_v = \perp$. We use this fact and the axiom (GML 4)

to show that $\langle l \rangle b \in w$ as follows. To fix ideas let consider three elements $\langle l_1 \rangle a_{u_1}, \langle l_2 \rangle a_{u_2}, \langle l_3 \rangle a_{u_3}$, the result can straightforwardly be generalized to any number n of such elements. We have

$$\begin{aligned}
& \langle l_1 + l_2 + l_3 \rangle (a_{u_1} \vee a_{u_2} \vee a_{u_3}) \\
&= \bigvee_{j=0}^{l_1+l_2+l_3} \langle j \rangle ((a_{u_1} \vee a_{u_2} \vee a_{u_3}) \wedge a_{u_1}) \wedge \langle l_1 + l_2 + l_3 - j \rangle ((a_{u_1} \vee a_{u_2} \vee a_{u_3}) \wedge a_{u_1}^c) \\
&= \bigvee_{j=0}^{l_1+l_2+l_3} \langle j \rangle a_{u_1} \wedge \langle l_1 + l_2 + l_3 - j \rangle (a_{u_2} \vee a_{u_3}) \\
&= \bigvee_{j=0}^{l_1+l_2+l_3} \langle j \rangle a_{u_1} \wedge \left(\bigvee_{j'=0}^{l_1+l_2+l_3-j} \langle j' \rangle a_{u_2} \wedge \langle l_1 + l_2 + l_3 - j - j' \rangle a_{u_3} \right) \\
&= \bigvee_{j=0}^{l_1+l_2+l_3} \bigvee_{j'=0}^{l_1+l_2+l_3-j} (\langle j \rangle a_{u_1} \wedge \langle j' \rangle a_{u_2} \wedge \langle l_1 + l_2 + l_3 - j - j' \rangle a_{u_3})
\end{aligned}$$

In particular, we therefore have

$$\langle l_1 \rangle a_{u_1} \wedge \langle l_2 \rangle a_{u_2} \wedge \langle l_3 \rangle a_{u_3} \leq \langle l_1 + l_2 + l_3 \rangle (a_{u_1} + a_{u_2} + a_{u_3})$$

More generally, using exactly the same argument we have

$$\bigwedge_k \langle l_k \rangle a_{u_k} \leq \left\langle \sum_k l_k \right\rangle \left(\bigvee_k a_{u_k} \right)$$

Since each $\langle l_k \rangle a_{u_k}$ is in the ultrafilter and that there are finitely many of them, their meet is in the ultrafilter and we then have that

$$\bigwedge_k \langle l_k \rangle a_{u_k} \leq \left\langle \sum_k l_k \right\rangle \left(\bigvee_k a_{u_k} \right) \leq \langle l \rangle \left(\bigvee_k a_{u_k} \right) \leq \langle l \rangle b$$

and thus $\langle l \rangle b$ belongs to the ultrafilter w defined by $\{\langle k_i \rangle a_i \mid i \in I\}$. We have thus proved that

$$\hat{\delta}_A(s(w)) = w$$

i.e. that s is indeed a section and that in consequence $\hat{\delta}$ is surjective on finite boolean algebras.

We have now shown that $\hat{\delta}$ is an iso on finite algebra, i.e. that GML is strongly complete over finite algebras. Moreover, since $\mathcal{B}^\omega$ weakly preserves filtered colimits by construction, it follows that we can now apply Theorem 5.27 and conclude that GML is strongly complete with respect to $\mathbf{CoAlg}(\mathcal{B}^\omega)$.

5.3 Completeness-via-canonicity

5.3.1 Jónsson-Tarski vs canonical extensions

Strong completeness via the coalgebraic Jónsson-Tarski Theorem 5.10 is the bedrock of coalgebraic completeness-via-canonicity, but so far we have made no use of canonicity as it was developed in Chapter 2. To do this, we need first of all to restrict our attention to coalgebraic logics over **DL** or **BA**, which we will now assume until the end of this Chapter. Secondly, if we are to make use of the algebraic theory of canonicity developed in Chapter 2 in conjunction with the theory of strong completeness developed in the previous Section 5.2, we need to relate two constructions which are a priori unrelated: (1) the canonical extension of a DLE (or BAE) and (2) the Jónsson-Tarski extension defined by Theorem 5.10. Whilst these two entities share the same carrier, namely the canonical extension of a distributive lattice or of a boolean algebra, there is no a priori reason for their structure maps to be the same. Formally, given a functor $L : \mathbf{DL} \rightarrow \mathbf{DL}$ defining DLEs, if the logic defined by L is interpreted in T -coalgebras for $T : \mathbf{Pos} \rightarrow \mathbf{Pos}$ via a semantic transformation $\delta : \mathcal{LU} \rightarrow \mathcal{UT}$ whose adjoint transpose has right inverses, then for a given $\alpha : LA \rightarrow A$ in $\mathbf{Alg}_{\mathbf{DL}}(L)$ what is the relationship between

$$LA^\sigma \xrightarrow{\alpha^\sigma} A^\sigma$$

the canonical extension of the DLE (A, α) (defined in Section 2.2.2), and the Jónsson-Tarski extension given by

$$LA^\sigma = L\mathcal{UPf}A \xrightarrow{\delta_{\mathcal{P}fA}} \mathcal{UTP}fA \xrightarrow{\mathcal{U}\delta_A^{-1}} \mathcal{UPf}LA \xrightarrow{\mathcal{UPf}\alpha} \mathcal{UPf}A = A^\sigma$$

In particular, when can we say that these two structure maps define the same L -algebra? Answering this question is the purpose of this section and the last technical obstacle to a coalgebraic theory of completeness-via-canonicity. For notational convenience we will denote the Jónsson-Tarski extension of a structure map $\alpha : LA \rightarrow A$ by $\tilde{\alpha}^\sigma$, i.e. $\tilde{\alpha}^\sigma = \mathcal{UPf}\alpha \circ \mathcal{U}\delta_A^{-1} \circ \delta_{\mathcal{P}fA}$ and we want to understand the relationship between α^σ and $\tilde{\alpha}^\sigma$.

We can answer this question in two different ways depending on whether we are dealing with DLEs or BAEs. Let us first make the nomenclature absolutely precise. By definition, if L defines a variety of DLEs or BAEs whose expansion symbols belong to a signature Σ , then a structure morphism (in **DL** or **BA**) $\alpha : LA \rightarrow A$ can be seen as a collection of *functions* $\alpha_\heartsuit : A^n \rightarrow A$ for each $\heartsuit \in \Sigma$ of arity n . Thus α^σ defines a collection of maps $\alpha_\heartsuit^\sigma : (A^\sigma)^n \rightarrow A^\sigma$ whose are by definition the canonical extensions of each $\alpha_\heartsuit$. We will say that α is **smooth** if each $\alpha_\heartsuit, \heartsuit \in \Sigma$ is smooth. Similarly, we will say that the map α^σ or the map $\tilde{\alpha}^\sigma$ are (τ_1, τ_2) -continuous if each $\alpha_\heartsuit^\sigma$ or $\tilde{\alpha}_\heartsuit^\sigma$ is (τ_1, τ_2) -continuous in each argument.

In the case of positive logics (i.e. over **DL**) we will use the following result.

Lemma 5.43. *Let $L : \mathbf{DL} \rightarrow \mathbf{DL}$ define a variety of DLEs whose expansions are (anti-)k-additive or (anti-)k-multiplicative, and assume that L -algebras can be interpreted in T -coalgebras for $T : \mathbf{Pos} \rightarrow \mathbf{Pos}$ via a semantic transformation $\delta : L\mathcal{U} \rightarrow \mathcal{U}T$ whose adjoint transpose has right inverses. If δ defines maps which are completely (anti-)k-additive or completely (anti-)k-multiplicative in every argument, then for any L -algebra (A, α) , $\tilde{\alpha}^\sigma$ is (σ, γ) -continuous.*

Proof. By definition of the Jónsson-Tarski extension, since $\mathcal{U}\text{Pf}\alpha$ and $\mathcal{U}\hat{\delta}_A^{-1}$ are just inverse images and therefore preserve all meets and all joins, the preservation properties of $\tilde{\alpha}^\sigma$ are determined by $\delta_{\text{Pf}A}^\heartsuit$. For the sake of simplicity we assume that the extension symbol is unary, but the proof clearly extends to n -ary operators in a straightforward way. Assume that $\delta_{\text{Pf}A}^\heartsuit$ is completely k -additive, then by the comment we just made it follows that $\tilde{\alpha}^\sigma$ is also completely k -additive. It follows immediately that $\tilde{\alpha}^\sigma$ preserves up-directed joins, and is therefore $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous.

Let us now show that $\tilde{\alpha}^\sigma$ is also $(\sigma^\downarrow, \gamma^\downarrow)$ -continuous. Since it is completely k -additive we have for every $u \in O(A)$

$$\begin{aligned} \tilde{\alpha}^\sigma(u) &\stackrel{1}{=} \tilde{\alpha}^\sigma(\bigvee A \cap \downarrow u) \\ &\stackrel{2}{=} \bigvee \{\tilde{\alpha}^\sigma(\bigvee U) \mid U \in \mathcal{P}_k(A \cap \downarrow u)\} \\ &\stackrel{3}{=} \bigvee \{\tilde{\alpha}^\sigma a \mid u \leq a \in A\} \\ &\stackrel{4}{=} \bigvee \{\alpha^\sigma a \mid u \leq a \in A\} \\ &= \alpha^\sigma(u) \end{aligned}$$

where (1) is by definition of open elements, (2) is by definition of k -additivity, (3) is because the set $A \cap \downarrow u$ is up-directed, and (4) is because both the Jónsson-Tarski and the canonical extensions are extensions, i.e. they agree with α on A . Thus $\tilde{\alpha}^\sigma$ and α^σ agree on open elements. Now we use the fact that since α^σ is assumed to be k -additive, α^σ is completely k -additive by Theorem 2.18, and thus $(\sigma^\downarrow, \gamma^\downarrow)$ -continuous by Theorem 2.35. By using the dual version of Lemma 2.27, we find that the topology $\gamma^\downarrow$ has a basis given by $\{\downarrow m \mid m \in M_\omega(A^\sigma)\}$ where $M_\omega(A^\sigma)$ is the set of finite meets of completely meet irreducible elements of A^σ . Since α^σ is $(\sigma^\downarrow, \gamma^\downarrow)$ -continuous, it means that for every $m \in M_\omega(A^\sigma)$ there exist $u \in O(A)$ such that $\alpha^\sigma(u) = m$. But since α^σ and $\tilde{\alpha}^\sigma$ agree on open elements this means that for every $m \in M_\omega(A^\sigma)$ there exist $u \in O(A)$ such that $\tilde{\alpha}^\sigma(u) = m$. In other words the inverse image under $\tilde{\alpha}^\sigma$ of any element in a basis of $\gamma^\downarrow$ is in $\sigma^\downarrow$, and it follows that $\tilde{\alpha}^\sigma$ must be $(\sigma^\downarrow, \gamma^\downarrow)$ -continuous. By gathering our results, it is clear that $\tilde{\alpha}^\sigma$ is (σ, γ) -continuous. The proof for anti- k -additivity and (anti-)k-multiplicativity follow along exactly the same lines. $\square$

We can now give a very general criteria for the canonical and the Jónsson-Tarski extensions to coincide:

Theorem 5.44 (Completeness-via-canonicity for DLEs). *Let $L : \mathbf{DL} \rightarrow \mathbf{DL}$ define a variety of DLEs whose expansions are (anti-)k-additive or (anti-)k-multiplicative, and assume that L -algebras can be interpreted in T -coalgebras for $T : \mathbf{Pos} \rightarrow \mathbf{Pos}$ via a semantic transformation $\delta : \mathcal{LU} \rightarrow \mathcal{UT}$ whose adjoint transpose has right inverses. If δ defines maps which are completely (anti-)k-additive or completely (anti-)k-multiplicative in every argument, then for any L -algebra (A, α) :*

$$(A^\sigma, \alpha^\sigma) = (A^\sigma, \tilde{\alpha}^\sigma)$$

Proof. The result follows from the preceding Lemma 5.43 and Corollary 2.30 since the assumptions of the Theorem guarantee that both α^σ and $\tilde{\alpha}^\sigma$ will be smooth, and thus equal. $\square$

In the case of BAEs, i.e. if we assume that the base category on which the syntax building functors are defined is $\mathbf{BA}$, we can in fact relax the smoothness condition, i.e. the (σ, γ) -continuity of δ , significantly. In fact, we only require $(\sigma, \gamma^\uparrow)$ - or $(\sigma, \gamma^\downarrow)$ -continuity. The following theorem also shows that we do not necessarily need the canonical extension and the Jónsson-Tarski extension to coincide.

Theorem 5.45 (Completeness-via-canonicity for BAEs). *Let $L : \mathbf{BA} \rightarrow \mathbf{BA}$ define a variety of BAEs whose expansions are isotone, and assume that L -algebras can be interpreted in T -coalgebras for $T : \mathbf{Set} \rightarrow \mathbf{Set}$ via a semantic transformation $\delta : LP \rightarrow PT$ whose adjoint transpose has right inverses. If δ defines maps which are all $(\sigma, \gamma^\uparrow)$ -continuous or all $(\sigma, \gamma^\downarrow)$ -continuous, then if the canonical extension of (A, α) belongs to a variety, so does its Jónsson-Tarski extension.*

Proof. Let V be a set of variables, and $q : F_L FV = \mu(L(-) + FV) \rightarrow Q$ be a regular epi (in $\mathbf{Alg}_{\mathbf{BA}}(L)$) which defines a variety by orthogonality (see Definition 1.29). By definition, it coequalizes a pair $e_1, e_2 : E \rightarrow F_L FV$ of ‘equations’. If the canonical extension $(A^\sigma, \alpha^\sigma)$ belongs to this variety, then it too must coequalize e_1, e_2 , i.e. $\llbracket e_1(x) \rrbracket_{\alpha^\sigma} = \llbracket e_2(x) \rrbracket_{\alpha^\sigma}$ for every ‘equation’ $x \in E$, where $\llbracket - \rrbracket_{\alpha^\sigma}$ denotes the catamorphism into the (A, α) . Since we’re in $\mathbf{BA}$ this can be equivalently restated as

$$\llbracket \neg(e_1(x) \leftrightarrow e_2(x)) \rrbracket_{\alpha^\sigma} = \perp$$

From the fact δ defines $(\sigma, \gamma^\uparrow)$ maps and that $\mathcal{PU}f\alpha$ and $\mathcal{P}\hat{\delta}_A^{-1}$ are $(\gamma^\uparrow, \gamma^\uparrow)$ -continuous, it follows that we $\tilde{\alpha}^\sigma$ is $(\sigma, \gamma^\uparrow)$ -continuous and from Theorem 2.28

it follows that $\tilde{\alpha}^\sigma \leq \alpha^\sigma$ pointwise, since α^σ defines the largest such extension of α . It follows by a simple induction on the modal depth of terms in $\mathbf{F}_L\mathbf{FV}$ that

$$\llbracket - \rrbracket_{\tilde{\alpha}^\sigma} \leq \llbracket - \rrbracket_{\alpha^\sigma}$$

pointwise, and in particular we have

$$\llbracket \neg(e_1(x) \leftrightarrow e_2(x)) \rrbracket_{\tilde{\alpha}^\sigma} \leq \llbracket \neg(e_1(x) \leftrightarrow e_2(x)) \rrbracket_{\alpha^\sigma} = \perp$$

Hence $\llbracket - \rrbracket_{\tilde{\alpha}^\sigma}$ also coequalizes e_1, e_2 and thus $(A^\sigma, \tilde{\alpha}^\sigma)$ also belongs to the variety defined by q .

The case of $(\sigma, \gamma^\downarrow)$ -continuous maps is treated dually by using Corollary 2.29 and $e_1(x) = e_2(x)$ iff $e_1(x) \leftrightarrow e_2(x) = \top$. $\square$

Remark 5.46. Whilst the previous Theorem 5.45 is much stronger topologically than Theorem 5.44, we have not found any natural preservation property on δ which implies $(\sigma, \gamma^\uparrow)$ or $(\sigma, \gamma^\downarrow)$ -continuity. Of course, any δ which is completely (anti-)k-additive or completely (anti-)k-multiplicative maps in each argument satisfy this requirement, but it is a priori a much weaker constraint to satisfy.

For L -algebras defining DLEs with smooth extensions or BAEs we will rely on the Theorems above to directly prove completeness-via-canonicity for coalgebraic logics. For more complicated functors, such as those defined by the nabla style of coalgebraic logics, we will proceed indirectly by using the translation techniques developed in Chapter 4. We start with the ‘straightforward’ case.

5.3.2 L -algebras defining DLEs with smooth extensions

We work in $\mathbf{DL}$ but everything that follows can easily be specialized to the case of endofunctors on $\mathbf{BA}$. Let $L : \mathbf{DL} \rightarrow \mathbf{DL}$ be a functor defining a variety of DLEs with smooth extensions. In all concrete examples we will use DLEs whose extensions obey the conditions of Theorem 2.35 in each of their arguments. In other words for all practical purposes we can assume L to define only n -ary modalities which are (anti-)k-additive or (anti-)k-multiplicative in each of their arguments, and thus smooth. We will write $F_L \dashv U_L$ for the adjunction $\mathbf{DL} \rightarrow \mathbf{Alg}_{\mathbf{DL}}(L)$ and η_L for the associated unit.

Theorem 5.47 (Coalgebraic completeness-via-canonicity for smooth DLEs). *Let $L : \mathbf{DL} \rightarrow \mathbf{DL}$ be a regular epi-preserving finitary functor defining a variety of DLEs with smooth extensions, $T : \mathbf{Pos} \rightarrow \mathbf{Pos}$ and $\delta : LU \rightarrow UT$ be such that the adjoint transpose $\hat{\delta} : TPf \rightarrow PfL$ has right-inverses, and let $\mathbf{Ax} \subseteq \mathbf{UF}_L\mathbf{FV} \times \mathbf{UF}_L\mathbf{FV}$ be any set of canonical equations defined over a set V of variables. If δ_{PfA} defines maps which are completely (anti-)k-additive or completely (anti-)k-multiplicative, then logic defined by L together with the axioms of $\mathbf{Ax}$ is strongly*

complete with respect to the class of T -coalgebras on which the equations of $\mathbf{Ax}$ are valid.

Proof. Given sets of formulas Φ, Ψ such that $\Phi \not\vdash_{\mathbf{ML}+\mathbf{Ax}} \Psi$, we need to find a model satisfying all the formulas of Φ and none of the formulas of Ψ on which the equations of $\mathbf{Ax}$ are valid. We start the construction as in Sections 1.5.4 and 4.3 by taking the adjoint transpose of the maps e_1, e_2 associated with the set of axioms $\mathbf{Ax}$, and then taking their coequalizer.

$$F\mathbf{Ax} \begin{array}{c} \xrightarrow{\hat{e}_1} \\ \xrightarrow{\hat{e}_2} \end{array} F_L FV \xrightarrow{q_{\mathbf{Ax}}} Q_{\mathbf{Ax}}$$

This defines the quotient under the smallest equivalence class (in **DL**) generated by the equations of $\mathbf{Ax}$. We can explicitly recover this equivalence relation by taking the kernel pair of $q_{\mathbf{Ax}}$. By Proposition 1.23 we know that we can get the fully invariant closure of this equivalence relation. So let us assume that we have the following exact sequence

$$F_{\mathbf{Ax}} \begin{array}{c} \xrightarrow{e'_1} \\ \xrightarrow{e'_2} \end{array} F_L FV \xrightarrow{q_{F_{\mathbf{Ax}}}} \mathcal{L}_{\mathbf{Ax}}$$

where $F_{\mathbf{Ax}}$ is the fully invariant closure of the equivalence relation defined by $q_{\mathbf{Ax}}$ (and the kernel pair of $q_{F_{\mathbf{Ax}}}$). The algebra $\mathcal{L}_{\mathbf{Ax}}$ is usually known as the **Lindenbaum-Tarski algebra** for the logic defined by L with additional axioms in $\mathbf{Ax}$. It is the quotient of the logic $F_L V$ under provable equivalence using the axioms of $\mathbf{Ax}$, and $F_{\mathbf{Ax}}$ can be thought of as the algebra of pairs of terms which are provably equivalent using equational reasoning and $\mathbf{Ax}$. By construction, $\mathcal{L}_{\mathbf{Ax}}$ is in the variety defined by the axioms in $\mathbf{Ax}$ (in fact it completely characterizes this variety, see Proposition 1.30), and since all the axioms are assumed to be canonical, this means that the canonical extension of $\mathcal{L}_{\mathbf{Ax}}$ also belongs to this variety. Since we're assuming that $\hat{\delta}_{\mathcal{L}_{\mathbf{Ax}}}$ has a right-inverse $h : \mathbf{Pf} L \mathcal{L}_{\mathbf{Ax}} \rightarrow T\mathbf{Pf} \mathcal{L}_{\mathbf{Ax}}$, the Jónsson-Tarski extension of $\mathcal{L}_{\mathbf{Ax}}$ exists by Theorem 5.10. Since $\mathcal{L}_{\mathbf{Ax}}$ inherits the canonical valuation from $F_L V$, we can consider the algebraic

Theorem turns out not to be too restrictive in practise as the following result illustrates.

Proposition 5.48. *For any finitary signature Σ , and any distributive lattice A , the semantic transformation $\delta_{\text{Pf}A}^\Sigma : L_\Sigma \text{UPf}A \rightarrow \mathcal{UT}_\Sigma \text{Pf}A$ interpreting the positive relational logic defined by Σ (see Section 5.2.2) defines maps which either (1) preserves all joins or anti-preserve all meets in each argument or (2) preserve all meets or anti-preserve all joins in each argument.*

Proof. Recall that the modalities $\heartsuit$ in the signature Σ come in two flavours, satisfying the distribution law of Eqs. (5.2) or (5.3) respectively, and that the δ^Σ accordingly defines the maps $\delta^\heartsuit : (\mathcal{UPf}A)^n \rightarrow \mathcal{UT}_\Sigma \text{Pf}A$

$$\begin{aligned} [\heartsuit(u_1, \dots, u_k, u_{k+1}, \dots, u_n)] &\mapsto \\ \{t \in T_\Sigma X \mid \exists (a_1, \dots, a_n) \in \pi_\heartsuit(t), (a_i \in u_i \text{ for all } i) \text{ and } (a_j \notin u_j \text{ for all } j)\} \\ &\text{if } \heartsuit \text{ satisfies Eq. (5.2)} \\ [\heartsuit(u_1, \dots, u_k, u_{k+1}, \dots, u_n)] &\mapsto \\ \{t \in T_\Sigma X \mid \forall (a_1, \dots, a_n) \in \pi_\heartsuit(t), (a_j \in u_j \text{ for all } j) \Rightarrow (a_i \in u_i \text{ for all } i)\} \\ &\text{if } \heartsuit \text{ satisfies Eq. (5.3)} \end{aligned}$$

The fact that these maps preserves all joins or anti-preserve all meets in each argument for expansions satisfying Eq. (5.2) and preserve all meets or anti-preserve all joins in each argument for expansions satisfying Eq. (5.3) follows from elementary set-theoretic considerations. $\square$

In the case of BAEs we can state the following result.

Theorem 5.49. *Let $L : \mathbf{BA} \rightarrow \mathbf{BA}$ be a regular epi-preserving finitary functor defining a variety of BAEs, $T : \mathbf{Set} \rightarrow \mathbf{Set}$ and $\delta : LP \rightarrow PT$ be such that the adjoint transpose $\hat{\delta} : T\text{Uf} \rightarrow \text{Uf}L$ has right-inverses, and let $\mathbf{Ax} \subseteq \text{UF}_L \text{FV}$ be any set of canonical formulas defined over a set V of variables. If $\delta_{\text{Uf}A}$ defines maps which are all $(\sigma, \gamma^\uparrow)$ -continuous or all $(\sigma, \gamma^\downarrow)$ -continuous then the logic defined by L together with the axioms of $\mathbf{Ax}$ is strongly complete with respect to the class of T -coalgebras on which the axioms of $\mathbf{Ax}$ are valid.*

Proof. The result follows from Theorem 5.45 in exactly the same way as the Theorem 5.47 above follows from Theorem 5.44. $\square$

Example 5.50 (Logic for trees with bounded degrees). We return to our Example 5.28 of a logic for trees with bounded degrees, say n . Recall that the operators of the logic defined by Eq. (5.10) all preserve both meets and joins in each of their argument. They are thus particularly well-behaved with regards to Table 2.1. Crucially, recall that the logic has the following axiom

$$(\text{Tree } 3) \neg \langle k \rangle (a_1, \dots, a_k) = \bigvee_{1 \leq j \neq k \leq n} \langle j \rangle (\top, \dots, \top) \vee \bigvee_{1 \leq i \leq k} \langle k \rangle (\epsilon_{1i} a_1, \dots, \epsilon_{ni} a_1)$$

which, by repeated use, allows us to push negations down to the propositional level. The consequence of this is particularly nice: since join-preserving operators are conservative (i.e. they make stable combinations when applied to any monotone term), since all operators are join preserving, and since we can always push the negations in front on propositional variables, we get that *every term in the language is stable, and thus canonical*.

As we saw in Example 5.28, the adjoint transpose $\hat{\delta}$ is surjective, and thus has right-inverses in **Set**. Moreover, we can easily check that the semantic transformation δ defines maps which are $(\sigma^\uparrow, \gamma^\uparrow)$ -continuous in each argument. By definition we have for each set X

$$\delta_X(U_1, \dots, U_k) = \{(x_1, \dots, x_k) \in \text{in}_k[X] \mid x_i \in U_i, 1 \leq i \leq k\}$$

Now let $(U_i^j)_{j \in J}$ be a collection of subsets of X . It is immediate that

$$\delta_X(U_1, \dots, \bigcup_{j \in J} U_i^j, \dots, U_k) = \bigcup_{j \in J} \delta_X(U_1, \dots, U_i^j, \dots, U_k)$$

i.e. δ_X preserves all joins in each of its arguments, and thus in particular all up-directed ones. It follows that δ_X is $((\gamma^\uparrow)^k, \gamma^\uparrow)$ -continuous, and we can therefore apply Theorem 5.49 to *any* set of formulas in the logic and get strong completeness for any collection of frame conditions (note that if the set of frame conditions is not consistent we get strong completeness vacuously).

Example 5.51 (Classical modal logic). We continue from Example 5.29, but we could also consider the countable version of classical modal logic developed in Example 5.40, since it is also strongly complete. We have described in Section 2.5 a general class of Sahlqvist identities (see Definition 2.47). We have seen that the adjoint transpose $\hat{\delta}$ of the semantic transformation has a right inverse (essentially as a special case of Theorem 5.14), and using the same argument as in the previous example it is easy to see that δ_X is $(\sigma^\uparrow, \gamma^\uparrow)$ -continuous, and we can therefore apply Theorem 5.49 to any set canonical axioms and get strong completeness-via-canonicity. Consider for example the general but not classical Sahlqvist formula described in Chapter 2:

$$(3) \quad \Diamond p \wedge \Diamond(q \wedge \neg p) \wedge \Diamond(r \wedge \neg p \wedge \neg q)$$

We can conclude that the classical modal logic $K + (3)$ is sound and strongly complete with respect to the class of frames where every point has at least three successors.

Example 5.52 (GML). Recall from Example 5.41 that GML is strongly complete w.r.t. the class of $\mathcal{B}^\omega$ -coalgebra, where $\mathcal{B}^\omega$ is the semantic completion of the usual bag functor, as described in Section 5.2.4. Recall also from 2.50 that the expansions $\langle k \rangle$ of GML are all k -additive. From Theorem 2.35 we know

that they are therefore smooth and the L -algebras of GML thus satisfies the conditions of Theorem 5.47. We need only check that the semantic transformation map $\delta : L\mathcal{P} \rightarrow \mathcal{PB}^\omega$ define completely k -additive maps.

$$\begin{aligned} \delta_X(\bigcup_i U_{i \in I}) &= \{f \in \mathcal{B}^\omega(X) \mid \sum_{x \in \text{supp}(f) \cap \bigcup_i U_i} f(x) \geq k\} \\ &= \bigcup_{\{i_1, \dots, i_p\} \subseteq I} \{f \in \mathcal{B}^\omega(X) \mid \sum_{x \in \text{supp}(f) \cap \bigcup_{j=1}^p U_{i_j}} f(x) \geq k\} \\ &= \bigcup_{\{i_1, \dots, i_k\} \subseteq I} (\delta(\bigcup_{j=1}^p U_{i_j})) \end{aligned}$$

Where the second step arises from the fact that any map f assigning a weight of at least k to $\bigcup_i U_i$ it is necessary and sufficient to find a partition $k = k_1 + \dots + k_p$ (in at most k integers) and a corresponding list of sets $U_{i_1}, \dots, U_{i_p}$ with $\{i_1, \dots, i_p\} \subseteq I$ such that $\sum_{x \in \text{supp}(f) \cap U_{i_j}} f(x) \geq k_j, 1 \leq j \leq p$, and thus

$$\sum_{x \in \text{supp}(f) \cap \bigcup_{j=1}^k U_{i_j}} f(x) \geq k$$

The semantic transformation therefore also satisfies the conditions of Theorem 5.47 and canonical equations in GML therefore define extensions of GML which are strongly complete with respect to the class of $\mathcal{B}^\omega$ -coalgebras which validate these equations. As shown in Section 2.6.1 Sahlqvist identities can be defined for GML (see Proposition and Definition 2.51), generating a large class of canonical formulas. Some examples of Sahlqvist frame conditions could for example be

$$\neg \langle k \rangle (\top)$$

i.e. ‘there can never be k or more successors’, which is Sahlqvist since it is the negation of a positive term. GML plus this axiom are thus sound and strongly complete w.r.t. frames in which every point has no more than k successors (including multiplicities). Similarly, we could consider the infinite set of axioms

$$\{\langle 2k \rangle \top \wedge \neg \langle 2k + 1 \rangle \top\}_{k \in \mathbb{N}}$$

i.e. ‘every point must have an even number of successors (including multiplicities)’, which are all Sahlqvist since the conjunction of a stable term and the negation of a positive term. GML plus this collection of axioms is thus sound and strongly complete w.r.t. frames in which every point has an even number of successors (including multiplicities).

Example 5.53 (Intuitionistic Logic and Distributive Substructural Logics). We expand on the examples of Section 2.6.2, which are also detailed in [DP15].

Recall that the language of intuitionistic logic (IL) was defined by the functor

$$L_{\text{Hey}} : \mathbf{BDL} \rightarrow \mathbf{BDL}, \begin{cases} A \mapsto F\{a \rightarrow b \mid a, b \in \mathcal{U}A\} / \equiv \\ L_{\text{RL}}f : L_{\text{Hey}}A \rightarrow L_{\text{Hey}}B, [a]_{\equiv} \mapsto [f(a)]_{\equiv} \end{cases}$$

where $\equiv$ is the fully invariant equivalence relation in $\mathbf{BDL}$ (see Section 1.4.4) generated by the following Heyting Distribution Laws:

$$\text{HDL1 } a \rightarrow (b \wedge c) = (a \rightarrow b) \wedge (a \rightarrow c)$$

$$\text{HDL2 } (a \vee b) \rightarrow c = (a \rightarrow c) \wedge (b \rightarrow c)$$

The language of distributive substructural logics was defined by the functor

$$L_{\text{RL}} : \mathbf{DL} \rightarrow \mathbf{DL}, \begin{cases} L_{\text{RL}}A = F\{I, a * b, a \setminus b, a / b \mid a, b \in \mathcal{U}A\} / \equiv \\ L_{\text{RL}}f : L_{\text{RL}}A \rightarrow L_{\text{RL}}B, [a]_{\equiv} \mapsto [f(a)]_{\equiv} \end{cases}$$

where $\equiv$ is the fully invariant equivalence relation in $\mathbf{DL}$ (see Section 1.4.4) generated by the Distribution Laws:

$$\text{DL1 } (a \vee b) * c = (a * c) \vee (b * c)$$

$$\text{DL2 } a * (b \vee c) = (a * b) \vee (a * c)$$

$$\text{DL3 } a \setminus (b \wedge c) = (a \setminus b) \wedge (a \setminus c)$$

$$\text{DL4 } (a \vee b) \setminus c = (a \setminus c) \wedge (b \setminus c)$$

$$\text{DL5 } (a \wedge b) / c = (a / c) \wedge (b / c)$$

$$\text{DL6 } a / (b \vee c) = (a / b) \wedge (a / c)$$

The functors L_{Hey} and L_{RL} are examples of the class of functors described in Section 5.2.2 defining *relational logics*. We therefore interpret L_{Hey} -algebras in coalgebras for the functor:

$$T_{\text{Hey}} : \mathbf{Pos} \rightarrow \mathbf{Pos}, \begin{cases} T_{\text{Hey}}W = \mathbf{P}_c(W \times W) \\ T_{\text{Hey}}f : T_{\text{Hey}}W \rightarrow T_{\text{Hey}}W', U \mapsto (f \times f)[U]. \end{cases}$$

via the semantic transformation $\delta^{\text{Hey}} : L_{\text{Hey}}\mathcal{U} \rightarrow \mathcal{U}T_{\text{Hey}}$

$$\delta_W^{\text{Hey}}(U \rightarrow V) = \{(x, y) \in T_{\text{Hey}}W \mid x \in U \Rightarrow y \in V\}$$

Similarly we interpret L_{RL} -algebras in coalgebras for the functor:

$$T_{\text{RL}} : \mathbf{Pos} \rightarrow \mathbf{Pos}, \begin{cases} T_{\text{RL}}W = \mathbf{2} \times (\mathbf{P}_c(W \times W))^3 \\ T_{\text{RL}}f : T_{\text{RL}}W \rightarrow T_{\text{RL}}W', U \mapsto (\text{Id}_2 \times (f \times f)^3)[U]. \end{cases}$$

via the semantic transformation $\delta^{\text{RL}} : L_{\text{RL}}\mathcal{U} \rightarrow \mathcal{U}T_{\text{RL}}$

$$\begin{aligned}\delta_W^{\text{RL}}(I) &= \{t \in T_{\text{RL}}W \mid \pi_1(t) = 0 \in \mathbb{2}\} \\ \delta_W^{\text{RL}}(u * v) &= \{t \in T_{\text{RL}}W \mid \exists(x, y) \in \pi_2(t), x \in u, y \in v\} \\ \delta_W^{\text{RL}}(u \setminus w) &= \{t \in T_{\text{RL}}W \mid \forall(x, y) \in \pi_3(t), x \in u \Rightarrow y \in w\} \\ \delta_W^{\text{RL}}(w / v) &= \{t \in T_{\text{RL}}W \mid \forall(x, y) \in \pi_4(t), x \in v \Rightarrow y \in w\}.\end{aligned}$$

where $\pi_i, 1 \leq i \leq 4$ are the usual projections maps. The intuition is that the first component of the structure map of a T_{RL} -coalgebra (to the (po)set $\mathbb{2}$) separates states into units and non-units. The second component sends each ‘state’ $w \in W$ to the pairs of states which it ‘contains’, the next two components are used to interpret $\setminus$ and $/$, respectively, and turn out to be very closely related to the second component.

Since these two situations are particular instances of relational logics, as developed in Section 5.2.2, we know by Theorem 5.16 that both δ^{Hey} and δ^{RL} have right inverses, and that the Jónsson-Tarski extension of any L_{Hey} or L_{RL} -algebra exists by Theorem 5.10. Moreover, we know from Proposition 2.13 that L_{Hey} and L_{RL} define varieties of DLEs with smooth expansions, and from Proposition 5.48 that the semantic transformation δ^{Hey} and δ^{RL} either (1) preserves all joins or anti-preserved all meets in every argument, or (2) preserves all meets or anti-preserved all joins in every argument. It follows that L_{Hey} -algebras and δ^{Hey} satisfy the conditions of Theorem 5.47, and similarly for L_{RL} -algebras and δ^{RL} . We therefore have strong completeness of L_{Hey} or L_{RL} -algebras validating canonical axioms with respect to the class of T_{Hey} or T_{RL} -coalgebras validating these axioms. We will apply this fact to get strongly complete semantics of IL and of the full distributive Lambek calculus.

Recall that the axioms HDL1-HDL2 are not sufficient to completely capture IL. We also need to specify that $\rightarrow$ is the residuum of $\wedge$ and that $a \rightarrow a$ is a theorem of IL. This can be done by using the following axioms (Heyting Frame Conditions):

$$\text{HFC1 } a \rightarrow a = \top,$$

$$\text{HFC2 } a \wedge (a \rightarrow b) = a \wedge b$$

$$\text{HFC3 } (a \rightarrow b) \wedge b = b$$

It is not difficult to check that HDL1-HDL2 together with HFC1-HFC3 axiomatize IL.

Similarly, DL1-DL6 does not fully axiomatize the full distributive Lambek calculus, i.e. distributive residuated lattices. To achieve this we need to specify that I is a unit for $*$ and that $\setminus$ and $/$ are the left and right residuals of $*$. This can be achieved via the following axioms (Frame Conditions):

FC1 $a * I = a, I * a = a,$

FC2 $I \leq a \setminus a, I \leq a/a,$

FC3 $a * (b \setminus c) \leq (a * b) \setminus c,$

FC4 $(c/b) * a \leq c/(a * b),$

FC5 $(a/b) * b \leq a,$ and

FC6 $b * (b \setminus a) \leq a,$

Again, it is not difficult to check that DL1-DL6 together with FC1-FC6 axiomatize distributive residuated lattices.

Proposition 5.54. *The axioms HFC1-3 and FC1-6 are canonical.*

Proof. The proof is an application the results of Chapter 2. We show the result for FC1-6, the result for HFC1-3 follow very similar lines.

FC1: Since $*$ preserves binary joins in each argument, it is smooth by Proposition 2.13, and it follows that it is (σ^2, γ) -continuous by Corollary 2.14. Since π_1^σ and I^σ are trivially (σ, σ) -continuous, it follows from Theorem 2.36 that $(* \circ \langle \pi_1, I \rangle)^\sigma = *^\sigma \circ \langle \pi_1, 1 \rangle^\sigma$. Each side of the equation is thus stable and the result follows from Proposition 2.38.

FC2: I is stable and thus contracting, and $(\setminus \circ \langle \pi_1, \pi_1 \rangle)^\sigma = \setminus^\sigma \circ \langle \pi_1, \pi_1 \rangle^\sigma$, since π_1^σ is (σ, σ) -continuous and $\setminus^\sigma$ is smooth. The RHS of the inequality is thus stable, and a fortiori expanding, and the inequality is thus canonical.

FC3-4: Since $*^\sigma$ preserve joins in each argument, it preserves up-directed ones, and is thus $((\gamma^\uparrow)^2, \gamma^\uparrow)$ -continuous. Since $\setminus^\sigma$ is smooth it is in particular $(\sigma^2, \gamma^\uparrow)$ -continuous. Since π_1^σ is $(\sigma, \gamma^\uparrow)$ -continuous, we get that $*^\sigma \circ \langle \pi_1^\sigma, \setminus^\sigma \circ \langle \pi_2^\sigma, \pi_3^\sigma \rangle \rangle$ is $(\sigma^3, \gamma^\uparrow)$ -continuous and thus contracting. For the RHS, note that since $\setminus^\sigma$ preserves meets in its first argument, it must in particular preserve down-directed ones, thus $\setminus^\sigma$ is $(\gamma^\downarrow, \gamma^\downarrow)$ -continuous in its first argument. Similarly, since $\setminus^\sigma$ anti-preserve joins in its second argument, it must in particular anti-preserve up-directed ones, and is thus $(\gamma^\uparrow, \gamma^\downarrow)$ -continuous in its second argument. This means that $\setminus^\sigma$ is $(\gamma^2, \gamma^\downarrow)$ -continuous. We thus have that the full term is $(\sigma^3, \gamma^\downarrow)$ continuous, and thus expanding. The inequation is therefore canonical.

FC5-6: The LHS is contracting by the same reasoning as above, and the RHS is stable and thus expanding. □

It follows from Proposition 5.54 and Theorem 5.47 that

- The logic defined by L_{Hey} and the Heyting frame conditions HFC1-3 is strongly complete with respect to T_{Hey} -coalgebras validating HFC1-3.

- The logic defined by L_{RL} and the frame conditions FC1-6 is strongly complete with respect to T_{RL} -coalgebras validating FC1-6.

To conclude this example, let us examine what these classes of T_{Hey} and T_{RL} -coalgebras look like. Let us first examine what T_{Hey} -coalgebras validating HFC1-3 look like. For every $\gamma : W \rightarrow T_{Hey}W$ in this class, every $w \in W$ and every valuation, $w \models a \rightarrow a$. By considering a formula satisfied at a single point in the model is easy to see that $(x, y) \in \gamma(w) \Rightarrow x = y$, i.e. the structure map of the coalgebra only really defines a binary relation to interpret $\rightarrow$. Thus T_{Hey} -coalgebras validating HFC1 are equivalent to P_c -coalgebras where $w \models a \rightarrow b$ iff $\forall x \in \gamma(w), x \models a \Rightarrow x \models b$. The distributivity laws of $\rightarrow$ together with HFC2-3 encode the well-known residuation property of $\rightarrow$ with respect to $\wedge$. Combined with HFC1 and the associated reformulation in terms of P_c -coalgebra, the residuation property states that:

$$w \models a \wedge b \Rightarrow w \models c \quad \text{iff} \quad w \models b \Rightarrow (\forall x \in \gamma(w) (x \models a \Rightarrow x \models c))$$

Assuming the left-hand side, for the right-hand side to hold it is necessary that if $w \models b$, then $\forall x \in \gamma(w), x \models b$; that is, successor states satisfy the so-called ‘persistency’ condition. Assuming the right-hand side, for the left hand side to hold in a model where $w \models a, b$ it is necessary that $x \in \gamma(x)$; that is, the relation is reflexive. Finally, from HFC3 we get that $a \wedge b \leq c$ iff $b \leq a \rightarrow c$ iff $b \leq a \rightarrow (c \wedge (a \rightarrow c))$. By unravelling the interpretation of this last inequality, we get that the relation interpreting $\rightarrow$ must also be transitive. Thus we have recovered the traditional Kripke semantics of intuitionistic logic via a pre-order and persistent valuations by using the theory of canonicity for distributive lattices. Recall that γ takes its values in the *convex* powerset of the carrier $(W, \leq)$ of the model. Thus the order defined by γ and the order of the carrier are closely related: if $y \in \gamma(x)$, then for every $z \in W$ such that $x \leq z \leq y$, $z \in \gamma(x)$ too. Indeed, since each poset is in particular a pre-order, each element of **Pos** comes with a ‘canonical’ T_{Hey} -coalgebra structure validating HFC1-3 defined by $\gamma(x) = \uparrow x$ (upsets are convex).

Let us now describe T_{RL} -coalgebras validating FC1-6. Axiom FC1 means that at every w in a T_{RL} -coalgebra, amongst all the pairs of states into which w can be ‘separated’ there must exist a *unit* state i , viz. $\pi_1(\gamma(i)) = 0$, such that $(w, i) \in \pi_2(\gamma(w))$. Similarly, there must exist a unit state i' such that $(i', w) \in \pi_2(\gamma(w))$. This condition can be found in this form in, for example, [CGZ07]. The other axioms are simply designed to capture the residuation condition in such a way that canonicity can be used, so a model in which FC2-6 are valid is simply a model in which the residuation conditions hold. By considering models with only three points it is easy to see that these conditions imply that

$$(y, z) \in \pi_1(\gamma(x)) \text{ iff } (x, z) \in \pi_2(\gamma(y)) \text{ iff } (y, x) \in \pi_3(\gamma(z)),$$

that is, the last three components of a T_{RL} -coalgebra's structure map are determined by any one of them. If we choose the second as defining the last two, a T_{RL} -coalgebra validating FC1-6, really is a coalgebra for the functor

$$T'_{\text{RL}} : \mathbf{Pos} \rightarrow \mathbf{Pos}, T'_{\text{RL}} W = \mathbb{2} \times \mathbf{P}_c(W \times W)$$

in which the interpretation of the operators is given by:

1. $w \models a * b$ iff $\exists(x, y) \in \gamma(w)$ s.t. $x \models a$ and $y \models b$
2. $w \models a/b$ iff $\forall(x, y)$ s.th $(w, y) \in \gamma(x)$ if $y \models b$ then $x \models a$
3. $w \models b \backslash a$ iff $\forall(x, y)$ s.th $(y, w) \in \gamma(x)$ if $y \models b$ then $x \models a$.

It is frequent to consider variations of the full distributive Lambek calculus where one or several of the structural rules are allowed. This can be achieved by adding yet more canonical frame conditions. Specifically, it is very easy to check that the following (in)equations, each corresponding to admitting a structural rule, are canonical: (1) Commutativity: $a * b = b * a$; (2) Increasing idempotence: $a \leq a * a$ (defines relevant logic); and (3) Integrality: $a \leq I$ (defines affine logic).

5.3.3 General L -algebras

We now define a method for proving completeness-via-canonicity in the presence of frame conditions for logics defined by a functor $L : \mathbf{BA} \rightarrow \mathbf{BA}$ which does not define a DLE. The main application of this method is the case of coalgebraic logics with the cover modality, but in principle, this method applies to any functor L which is too complicated to deal with using the direct method detailed in the previous Section. Since the target usage of the method in this Section is the case of nabla logics, we place ourselves in the $\mathbf{Uf} \dashv \mathcal{P} : \mathbf{BA} \rightarrow \mathbf{Set}^{\text{op}}$ situation throughout.

For this method, we continue where we left off at the end of Chapter 4. We briefly recall the basic setting. We have two finitary regular epi-preserving functors $K, L : \mathbf{BA} \rightarrow \mathbf{BA}$, two functors $S, T : \mathbf{Set} \rightarrow \mathbf{Set}$ and two epi-transformations $q : K \twoheadrightarrow L$ and $r : S \twoheadrightarrow T$. We also have two semantic natural transformations $\lambda : K\mathcal{P} \rightarrow \mathcal{P}S$ and $\delta : L\mathcal{P} \rightarrow \mathcal{P}T$, and these define *compatible semantics*, i.e. for every set X and every $a \in L\mathcal{P}X$ we must have

$$\lambda_X(a) \subseteq \mathcal{P}r_X \circ \delta_X \circ q_{\mathcal{P}X}(a)$$

For the method to work, or even make sense, we need K to define a BAE. Typically, K would arise from a presentation of L , as discussed in Section 3.6. In this case K of the form $(\mathbf{FS}_{\Sigma}\mathbf{U})/\simeq$, for a polynomial $\mathbf{Set}$ -functor $\mathbf{S}_{\Sigma}$, and possibly some quotient $\simeq$. In other words, we can assume that whilst L might

not define a variety of BAEs, K does. We fix a set of propositional variables V throughout.

The problem is the following: given a set E of axioms (frame conditions) in the logic defined by L , and a (possibly infinite) set $\Phi \subseteq \mu(L(-) + \mathbf{FV})$ of L -formulas which are consistent with the axioms of E , we want to build an T -model in which every formula $a \in \Phi$ is satisfiable. We proceed in five steps.

Build the Lindenbaum-Tarski algebra.

Let us write $F_L \dashv U_L$ for the adjunction $\mathbf{BA} \rightarrow \mathbf{Alg}_{\mathbf{BA}}(L)$, and similarly for K . We proceed as in the direct method by building the Lindenbaum-Tarski algebra associated with E . As usual, we start by taking the adjoint transpose of the maps e_1, e_2 associated with the set of equations E , and then taking their coequalizer.

$$FE \xrightleftharpoons[e_2]{e_1} F_L FV \xrightarrow{q_E} Q_E$$

This defines the quotient under the smallest equivalence class generated by the equations of E . By Proposition 1.23 we know that we can get the fully invariant closure of this equivalence relation. So let us assume that we have the following exact sequence

$$F_E \xrightleftharpoons[e'_2]{e'_1} F_L FV \xrightarrow{q_{F_E}} \mathcal{L}_E$$

where F_E is the fully invariant closure of the equivalence relation defined by q_E (and the kernel pair of q_{F_E}). The algebra $\mathcal{L}_E$ is the Lindenbaum-Tarski algebra for the logic defined by L with additional axioms in E . As we saw in Section 1.5.4, $\mathcal{L}_E$ has the nice property that it is itself a member of the variety it defines, or put succinctly, $\mathcal{L}_E \perp q_{F_E}$ (see Proposition 1.30). Note that if $a \in \Phi$ is consistent with the axioms of E , then $\llbracket a \rrbracket_{\mathcal{L}_E} \neq \perp$. Since $\mathcal{L}_E$ inherits the canonical valuation of $F_L FV$, we can actually think of the formulas $a \in \Phi$ as elements of $\mathcal{L}_E$ which are not equal to $\perp$.

Lift the axioms.

We now apply the construction leading to Theorem 4.13, and lift the equations in F_E to the K -logic. Recall that we generated Diagram (4.7):

$$\begin{array}{ccccc} E^* & \xrightleftharpoons[e_2^*]{e_1^*} & F_K FV & \xrightarrow{q_{E^*}} & Q_{E^*} \\ r \downarrow & & \downarrow \xi_V & & \downarrow q_{E^*}' \\ QF_E & \xrightleftharpoons[Qe'_2]{Qe'_1} & QF_L FV & \xrightarrow{Qq_{F_E}} & Q\mathcal{L}_E \end{array}$$

where $Q : \mathbf{Alg}(L) \rightarrow \mathbf{Alg}(K)$ pre-composes with $q : K \rightarrow L$.

We can already point out that since all the vertical arrows are epis, every $a \in \Phi$ must have an inverse image $a' \neq \perp$ under the syntax translation map.

Check for canonicity and that validity in the canonical extension implies validity in the Jónsson-Tarski extension.

Since K is assumed to define a BAE, it makes sense to ask whether the variety defined by q_{E^*} is canonical. This will in particular be the case if E^* is generated by canonical equations, for example Sahlqvist identities. This leads us to the following natural definition.

Definition 5.55. Let $q : K \rightarrow L$ be a regular epi-transformation between two regular epi-preserving finitary functors on $\mathbf{BA}$. We will say that $a \in F_L FV$ is a **abstract coalgebraic Sahlqvist formula** for the transformation q if every pre-image of a under the translation map $(-)^q : U_K F_K FV \rightarrow U_L F_L FV$ is an abstract Sahlqvist formula in the sense of Definition 2.42. We will say that a is a **concrete coalgebraic Sahlqvist formula** if we have a notion of concrete Sahlqvist formula in the BAEs defined by K and every pre-image of a under the translation map $(-)^q$ is a concrete Sahlqvist formula, for example in the sense of Definition 2.47.

If E^* is generated by a set of Sahlqvist identities, then the variety defined by q_{E^*} is canonical, and thus if $A \perp q_{E^*}$, then $\mathcal{P}UfA \perp q_{E^*}$. Recall from Theorem 4.13 that in the context of the lifting of equations described by the diagram above, we have for any L -algebra A that $A \perp q_{F_E}$ iff $QA \perp q_{E^*}$. Since the Lindenbaum-Tarski algebra $\mathcal{L}_E$ has the property that $\mathcal{L}_E \perp q_{F_E}$, Theorem 4.13 tells us that $Q\mathcal{L}_E \perp q_{E^*}$. Thus if the variety defined by q_{E^*} is canonical we will have

$$\mathcal{P}UfQ\mathcal{L}_E \perp q_{E^*}$$

Since we are working over $\mathbf{BA}$, the most general result to verify that equations that hold in the canonical extension also hold in the Jónsson-Tarski extension is Theorem 5.45, but in practise Theorem 5.44 is more useful since one can often exhibit preservation properties of the expansions defined by the S -algebras.

Build an S -model.

If the adjoint transpose $\hat{\lambda} : T\mathcal{U}f \rightarrow \mathcal{U}fK$ has a right inverse h at $\mathcal{L}_E$, then we can build a coalgebra

$$\gamma : \mathcal{U}f\mathcal{L}_E \rightarrow \mathcal{U}fL\mathcal{L}_E \xrightarrow{\mathcal{U}fq_{\mathcal{L}_E}} \mathcal{U}fK\mathcal{L}_E \xrightarrow{h} S\mathcal{U}f\mathcal{L}_E$$

which provides a Jónsson-Tarski embedding of $\mathcal{L}_E$ into $\mathcal{PUf}\mathcal{L}_E$ as K -algebras (see Theorem 5.10 and remarks thereafter). Since this embedding $\mathcal{L}_E \rightarrow \mathcal{PUf}\mathcal{L}_E$ is injective and since we know that for every $a \in \Phi$, $a \neq \perp$, we get that $\llbracket a \rrbracket_{\mathcal{PUf}\mathcal{L}_E} \neq \perp$ for the valuation on $\mathcal{PUf}\mathcal{L}_E$ inherited from the canonical valuation on $\mathcal{L}_E$ as follows

$$\begin{array}{ccccc}
 & & K\llbracket - \rrbracket_{\mathcal{PUf}\mathcal{L}_E} + \text{Id}_{FV} & & \\
 & \nearrow & & \searrow & \\
 KF_KFV + FV & \xrightarrow{K(q_{FE} \circ (-)^q) + \text{Id}_{FV}} & K\mathcal{L}_E + FV & \xrightarrow{K\eta_{\mathcal{L}_E} + \text{Id}_{FV}} & K\mathcal{PUf}\mathcal{L}_E + FV \\
 \downarrow \langle - \rangle_V + \eta_V^K & & \downarrow \alpha \circ q_{\mathcal{L}_E} + \eta_V^L & \searrow \eta_{K\mathcal{L}_E} & \downarrow \delta_{\mathcal{Uf}\mathcal{L}_E} + \text{Id}_{FV} \\
 & & & & \mathcal{PSUf}\mathcal{L}_E + FV \\
 & & & & \downarrow \mathcal{P}h + \text{Id}_{FV} \\
 & & & & \mathcal{PUf}K\mathcal{L}_E + FV \\
 & & & \searrow \mathcal{PUf}(\alpha \circ q_{\mathcal{L}_E}) + \mathcal{PUf}\eta_V^L \circ \eta_{FV} & \downarrow \\
 F_KFV & \xrightarrow{q_{FE} \circ (-)^q} & \mathcal{L}_E & \xrightarrow{\eta_{\mathcal{L}_E}} & \mathcal{PUf}\mathcal{L}_E \\
 & \searrow & & \nearrow & \\
 & & \llbracket - \rrbracket_{\mathcal{PUf}\mathcal{L}_E} & &
 \end{array}$$

where η^K is the unit of the adjunction $F_K \dashv U_K$, η^L that of $F_L \dashv U_L$, η is the unit of the adjunction $\mathcal{Uf} \dashv \mathcal{P}$ and $\langle - \rangle_V$ and α are the structure morphisms of F_K and $\mathcal{L}_E$ respectively.

As noted above, there exists for every $a \in \Phi$ an $a' \in F_KFV$ such that $a' \neq \perp$ and $(a')^q = a \neq \perp$. Let us gather such a choice of pre-images in a set $\Phi' \subseteq F_KFV$. By assumption that Φ is consistent with the axioms of E we also have that for each $a \in \Phi$, $q_{FE}(a) \neq \perp$. Finally, since the Jónsson-Tarski embedding is injective we can conclude that $\llbracket a' \rrbracket_{\mathcal{PUf}\mathcal{L}_E} \neq \perp$ for every $a' \in \Phi'$. Note that the semantic map $\llbracket - \rrbracket_{\mathcal{PUf}\mathcal{L}_E}$ is explicitly given by $\eta_{\mathcal{L}_E} \circ q_{FE} \circ (-)^q$. This is just a reformulation of the usual *truth lemma*: the interpretation of a formula a' is simply the set of ultrafilters of $\mathcal{L}_E$ which contains its translation a . Since Φ is consistent, we can extend it to an ultrafilter by Lindenbaum's lemma, i.e. to an element $u \in \mathcal{PUf}\mathcal{L}_E$, and for every $a' \in \Phi'$, $u \in \llbracket a' \rrbracket_{\mathcal{L}_E}$. By definition of the semantic map (see Section 1.5.2), if $u \in \llbracket a' \rrbracket_{\mathcal{PUf}\mathcal{L}_E}$ for every $a' \in \Phi'$, then

$$u, \gamma \models a'$$

where the valuation is simply the valuation given by the map $FV \rightarrow \mathcal{PUf}\mathcal{L}_E$ defined in the diagram above. By the sequence of adjunctions described in the proof of Theorem 5.1, it is equivalent to a map $\mathcal{Uf}\mathcal{L}_E \rightarrow \mathcal{QV}$, i.e. a valuation on the coalgebra. Thus we have found an S -model in which every $a' \in \Phi$ is satisfiable (simultaneously).

We now need to check that the lifted axioms of E^* are valid over our S -model. Since the lifted axioms are assumed to be canonical and are valid in the K -algebra $\mathbf{Q}\mathcal{L}_E$ by construction, they are also valid over the canonical extension of $\mathbf{Q}\mathcal{L}_E$. By assumption we therefore know that they also hold on the Jónsson-Tarski extension of $\mathbf{Q}\mathcal{L}_E$, and all the axioms of E^* are automatically valid on the coalgebra $\gamma : \mathcal{P}\mathbf{Uf}\mathcal{L}_E \rightarrow S\mathcal{P}\mathbf{Uf}\mathcal{L}_E$ by definition of the coalgebraic semantics.

Build a T -model.

To conclude, we use Theorem 4.7 which tells us that if the semantics are compatible, then if

$$u, \gamma, v \models a'$$

then

$$u, q_{\mathbf{Uf}\mathcal{L}_E} \circ \gamma, v \models (a')^q = a$$

Since this holds for every $a' \in \Phi'$, the model above is the model we were looking for, i.e. a model in which every formula in Φ is satisfiable. We now just need to check that the axioms of E are valid on the coalgebra $q_{\mathbf{Uf}\mathcal{L}_E} \circ \gamma : \mathcal{P}\mathbf{Uf}\mathcal{L}_E \rightarrow T\mathcal{P}\mathbf{Uf}\mathcal{L}_E$. It is quite clear that this should be the case, by construction of E^* . For every axiom $b = \top$ in E , the axiom $b' = \top$ is in E^* for every pre-image b' of b (under $(-)^q$). Thus $\gamma \models b' = \top$, and thus by Theorem 4.7, $q_{\mathbf{Uf}\mathcal{L}_E} \circ \gamma \models (b')^q = b = \top$. Thus every axiom of E holds everywhere on our coalgebra $q_{\mathbf{Uf}\mathcal{L}_E} \circ \gamma : \mathcal{P}\mathbf{Uf}\mathcal{L}_E \rightarrow T\mathcal{P}\mathbf{Uf}\mathcal{L}_E$. as desired.

Remark 5.56. There are actually two ways to get a K -algebra embedding of $\mathcal{L}_E$ into $\mathcal{P}\mathbf{Uf}\mathcal{L}_E$. The strategy described above uses the assumption that $\hat{\lambda}$ has a right inverse at $\mathcal{L}_E$. This allows us to embed $\mathbf{Q}\mathcal{L}_E$ into $\mathcal{P}\mathbf{Uf}\mathbf{Q}\mathcal{L}_E$. The other strategy would be to assume that δ has a right inverse at $\mathcal{L}_E$, this would allow us to first embed $\mathcal{L}_E$ into its L -algebra canonical extension $\mathcal{P}\mathbf{Uf}\mathcal{L}_E$, and then use the transformation $\mathbf{Q}$ to get an embedding $\mathbf{Q}\mathcal{L}_E$ into $\mathbf{Q}\mathcal{P}\mathbf{Uf}\mathcal{L}_E$. However, we cannot use canonicity in the later strategy since $\mathbf{Q}\mathcal{P}\mathbf{Uf}\mathcal{L}_E$ is not the canonical extension of $\mathbf{Q}\mathcal{L}_E$.

Example 5.57 (Nabla logic for the finitary powerset functor $\mathbf{P}_f$). Recall that the logic is defined by $L : \mathbf{BA} \rightarrow \mathbf{BA}$ defined by $L = \mathbf{FP}_f\mathbf{U}(-) / \simeq_{\mathbf{KKV}(\mathbf{P}_f)}$ and is interpreted in $\mathbf{P}_\omega$ coalgebras via $\delta : L\mathcal{P} \rightarrow \mathcal{P}\mathbf{P}_\omega$ which takes the ‘generators’ of $L\mathcal{P}X$ to the set of their $\mathbf{P}_\omega$ -lifted members at each X . A presentation of $\mathbf{P}_\omega$ is given by the functor

$$S_\omega X = \left(\coprod_{n \in \omega} X^n \right)$$

and a regular quotient $q : S_\omega \twoheadrightarrow \mathbf{P}_\omega$ given by

$$q_X(x_1, \dots, x_n) = q_X(y_1, \dots, y_p) \text{ iff } \{x_1, \dots, x_n\} = \{y_1, \dots, y_p\}$$

the tuples with the same underlying set of elements are identified. the functor S_ω defines a nabla-style logic defined by the functor $K : \mathbf{BA} \rightarrow \mathbf{BA}$ given by $K = \mathbf{FS}_\omega \mathbf{U}(-) / \simeq_{\mathbf{KKV}(S_\omega)}$ and interpreted via a semantic transformation $\lambda : K\mathcal{P} \rightarrow \mathcal{PS}_\omega$ which takes the ‘generators’ of $K\mathcal{P}X$ to the set of their S_ω -lifted members at each X . We have shown in Chapter 4 how the natural transformation q can be used to construct a natural transformation $r : K \rightarrow L$, in such a way that the semantics λ and δ associated with K and L are *compatible* with respect to r and q in the sense of Definition 4.6. We showed in Theorem 4.1 that the transformation $r : K \rightarrow L$ defines a translation morphism $\xi_V : \mathbf{F}_K \mathbf{FV} \rightarrow \mathbf{QF}_L \mathbf{FV}$, where $\mathbf{Q} : \mathbf{Alg}_{\mathbf{BA}}(L) \rightarrow \mathbf{Alg}_{\mathbf{BA}}(K)$ is the functor defined by pre-composing with r .

As an illustration of the construction above, let us consider the axiom for transitivity $(4) = \nabla\{\nabla\{p\}\} \rightarrow \nabla\{p\}$, and let us consider an arbitrary set Φ of L -formulas which are $\mathbf{KKV}(\mathbf{P}_\omega) + (4)$ -consistent. We follow the steps detailed above.

- (Step 1) We start by building the Lindenbaum-Tarski algebra as the coequalizer of the equation $(4) = \top$, i.e. $q_{(4)} : \mathbf{F}_L \mathbf{FV} \rightarrow \mathcal{L}_4 = \mathbf{F}_L \mathbf{FV} / (4)$. It is an L -algebra, and we will denote its structure map by $\alpha : L\mathcal{L}_4 \rightarrow \mathcal{L}_4$. Note that since Φ is a $\mathbf{KKV}(\mathbf{P}_\omega) + (4)$ -consistent set of formulas, Φ can be extended to an ultrafilter $\langle \Phi \rangle^\uparrow$ in $\mathcal{L}_4$. Note also that $\mathcal{L}_4$ comes with a canonical valuation $\eta_V^L : \mathbf{FV} \rightarrow \mathcal{L}_4$ which simply injects propositional formulas into $\mathcal{L}_4$.
- (Step 2) We then lift the axiom (4) to the language of the functor K as was shown in in Section 4.3 by taking the pre-images under ξ_V of (4) (and boolean combinations thereof), for example $\nabla(\nabla(p, p), \nabla(p)) \rightarrow \nabla(p, p, p)$. Let us denote this set by $(4)^*$. Theorem 4.13 tells us that $\mathbf{Q}\mathcal{L}_4$ belongs to the variety of defined by $(4)^*$, since $\mathcal{L}_4$ belongs to the variety defined by (4) .
- (Step 3) Let us check that these inverse images are canonical. As was shown in [DP13] and Examples 5.28 and 5.30, the axioms of the logic for trees of unbounded branching degrees enforce that the operators defined by $\mathbf{FS}_\omega \mathbf{U}$ preserve meets and joins in each argument, and are thus in particular smooth. All elements of $(4)^*$ can be written from elements of the general form:

$$\vee(\neg\nabla(\nabla(p, \dots, p), \dots, \nabla(p, \dots, p)), \nabla(p, \dots, p))$$

which is always Sahlqvist since $\vee$ is conservative, $\nabla(p, \dots, p)$ is stable, and $\neg\nabla(\nabla(p, \dots, p), \dots, \nabla(p, \dots, p))$ is the negation of a positive term. In other words each element of $(4)^*$ is canonical. By canonicity of the elements of $(4)^*$ it follows that since $\mathbf{Q}\mathcal{L}_4$ belongs to the variety defined by $(4)^*$, so does its canonical extension.

Let us now check that validity on the canonical extension implies validity on the Jónsson-Tarski extension. Recall from Example 5.39 that the logic defined by K is strongly complete with respect to S_ω^* -coalgebra. The Jónsson-Tarski extension of $Q\mathcal{L}_4$ is thus given by

$$LP\mathcal{Uf}\mathcal{L}_4 \xrightarrow{\delta_{\mathcal{Uf}\mathcal{L}_4}} \mathcal{PS}_\omega^*\mathcal{Uf}\mathcal{L}_4 \xrightarrow{\mathcal{P}\hat{\delta}_{\mathcal{L}_4}^{-1}} \mathcal{P}\mathcal{Uf}K\mathcal{L}_4 \xrightarrow{\mathcal{P}\mathcal{Uf}(\alpha \circ r_{\mathcal{L}_4})} \mathcal{P}\mathcal{Uf}\mathcal{L}_4$$

together with the valuation $\tilde{v} : \mathcal{P}\mathcal{Uf}\eta_V^L \circ \eta_{FV} : FV \rightarrow \mathcal{P}\mathcal{Uf}\mathcal{L}_4$. It is easy to check that $\delta_{\mathcal{Uf}\mathcal{L}_4}$ defines maps which preserve all joins in each argument, and by Theorem 5.44 it follows that formulas which are valid on the canonical extension of $Q\mathcal{L}_4$ are also valid on its Jónsson-Tarski-extension.

(Step 4) Let us gather our results in the following diagram:

$$\begin{array}{ccccc}
 & & K[-]_{\mathcal{P}\mathcal{Uf}\mathcal{L}_E} + \text{Id}_{FV} & & \\
 & \swarrow & & \searrow & \\
 KFKFV + FV & \xrightarrow{K(q_{(4)} \circ \xi_V) + \text{Id}_{FV}} & K\mathcal{L}_4 + FV & \xrightarrow{K\eta_{\mathcal{L}_4} + \text{Id}_{FV}} & K\mathcal{P}\mathcal{Uf}\mathcal{L}_4 + FV \\
 \downarrow \langle - \rangle_V + \eta_V^K & & \downarrow \alpha \circ r_{\mathcal{L}_4} + \eta_V^L & \searrow \eta_{\mathcal{L}_4}^K + \text{Id}_{FV} & \downarrow \delta_{\mathcal{Uf}\mathcal{L}_4} + \text{Id}_{FV} \\
 & & & & \mathcal{PS}_\omega^*\mathcal{Uf}\mathcal{L}_4 + FV \\
 & & & & \downarrow \mathcal{P}\hat{\delta}_{\mathcal{L}_4}^{-1} + \text{Id}_{FV} \\
 & & & & \mathcal{P}\mathcal{Uf}K\mathcal{L}_4 + FV \\
 & & & \searrow \mathcal{P}\mathcal{Uf}(\alpha \circ r_{\mathcal{L}_4}) + \mathcal{P}\mathcal{Uf}\eta_V^L \circ \eta_{FV} & \downarrow \\
 F_KFV & \xrightarrow{q_{(4)} \circ \xi_V} & \mathcal{L}_4 & \xrightarrow{\eta_{\mathcal{L}_4}} & \mathcal{P}\mathcal{Uf}\mathcal{L}_4 \\
 & \searrow & & \swarrow & \\
 & & [-]_{\mathcal{P}\mathcal{Uf}\mathcal{L}_4} & &
 \end{array}$$

We know from the previous step and from the commutativity of the diagram above that all the formulas of $(4)^*$ are valid in the coalgebra

$$\mathcal{Uf}\mathcal{L}_4 \xrightarrow{\mathcal{Uf}(\alpha \circ r_{\mathcal{L}_4})} \mathcal{Uf}K\mathcal{L}_4 \xrightarrow{\hat{\delta}_{\mathcal{L}_4}^{-1}} S_\omega^*\mathcal{Uf}\mathcal{L}_4$$

together with the canonical valuation arising from $\tilde{v}$ by adjointness: $\hat{v} : \mathcal{Uf}\mathcal{L}_4 \rightarrow \mathcal{Uf}FV = QV$. For $a \in \Phi$, since r is epi there exists a pre-image a' under ξ_V . It is clear from the diagram above that

$$\begin{aligned}
 \llbracket a' \rrbracket_{\mathcal{P}\mathcal{Uf}\mathcal{L}_4} &= \eta_{\mathcal{L}_4} \circ q_{(4)} \circ \xi_V(a') \\
 &= \{u \in \mathcal{Uf}\mathcal{L}_4 \mid a \in u\}
 \end{aligned}$$

and thus $\langle \Phi \rangle^\dagger \in \llbracket a' \rrbracket_{\mathcal{P}\mathcal{Uf}\mathcal{L}_4}$. Equivalently, for every $a \in \Phi$ we can find a' such that $\xi_V(a') = a$ and such that

$$\langle \Phi \rangle^\dagger, \hat{\delta}_{\mathcal{L}_4}^{-1} \circ \mathcal{Uf}(\alpha \circ r_{\mathcal{L}_4}), \hat{v} \models a'$$

(Step 5) We now need to build a P_ω -model with the same property. This is the most delicate step. Notice that from the previous steps we know how to build an S_ω^* -model, but since we want to quotient our model into a P_ω model, we would have wanted a S_ω -model and used q to turn in into a P_ω -model.

We use a different trick from that developed in [DP13] to turn our S_ω^* -model into a P_ω -model. Here we simply define an extension of the epi-transformation $q : S_\omega \rightarrow P_\omega$ as follows

$$q_X^* : S_\omega^* X \rightarrow P_\omega X, x \mapsto \begin{cases} q(x) & \text{if } x \neq * \\ \emptyset & \text{else} \end{cases}$$

We now check that the triples defined by $(K, S_\omega^*, \lambda^* : K\mathcal{P} \rightarrow \mathcal{P}S_\omega^*)$ and $(L, P_\omega, \delta : L\mathcal{P} \rightarrow \mathcal{P}P_\omega)$ - where λ is the obvious extension of $\lambda : K\mathcal{P} \rightarrow \mathcal{P}S_\omega$ as defined in Examples 5.28 and 5.30 - are compatible in the sense of Definition 4.6, i.e. that

$$\lambda_X^*(a) \subseteq \mathcal{P}q_X^* \circ \delta_X \circ r_{\mathcal{P}X}(a)$$

But as was shown in Section 4.2.3, we already know that

$$\lambda_X(a) \subseteq \mathcal{P}q_X \circ \delta_X \circ r_{\mathcal{P}X}(a)$$

It is clear that $\lambda_X(a) = \lambda_X^*(a)$ for any $a \in K\mathcal{P}X$, since the only difference is that λ_X^* has an additional point $*$ in its codomain (which is never reached by λ_X^*). Moreover, by definition of q_X^* it is clear that $\mathcal{P}q_X(U) \subseteq \mathcal{P}q_X^*(U)$, with the inequality being strict when $\emptyset \in U$. We thus have

$$\lambda_X^*(a) = \lambda_X(a) \subseteq \mathcal{P}q_X \circ \delta_X \circ r_{\mathcal{P}X}(a) \subseteq \mathcal{P}q_X^* \circ \delta_X \circ r_{\mathcal{P}X}(a)$$

and the triples above therefore define compatible semantics. We can now use Theorem 4.7 to show that since for each $a \in \Phi$ there exists $a' \in F_K FV$ such that $\xi_V(a') = a$ and

$$\langle \Phi \rangle^\uparrow, \hat{\delta}_{\mathcal{L}_4^{-1}} \circ \text{Uf}(\alpha \circ q_{\mathcal{L}_4}), \hat{v} \models a',$$

it must also be the case that

$$\langle \Phi \rangle^\uparrow, q_{\text{Uf}\mathcal{L}_4}^* \circ \hat{\delta}_{\mathcal{L}_4^{-1}} \circ \text{Uf}(\alpha \circ r_{\mathcal{L}_4}), \hat{v} \models \xi_V(a') = a$$

We have thus found a P_ω -model satisfying every formula of Φ , moreover, since every formula of $(4)^*$ is valid on the frame defined by $\hat{\delta}_{\mathcal{L}_4^{-1}} \circ \text{Uf}(\alpha \circ q_{\mathcal{L}_4}) : \text{Uf}\mathcal{L}_4 \rightarrow S_\omega^* \text{Uf}\mathcal{L}_4$ and $\hat{v}$, it follows from Theorem 4.7 that (4) is valid on the frame defined by $q_{\text{Uf}\mathcal{L}_4}^* \circ \hat{\delta}_{\mathcal{L}_4^{-1}} \circ \text{Uf}(\alpha \circ q_{\mathcal{L}_4}) : \text{Uf}\mathcal{L}_4 \rightarrow P_\omega \text{Uf}\mathcal{L}_4$ and $\hat{v}$.

Index

- k -additive, 81
 - Anti-, 82
 - Completely, 82
- k -additive, k -multiplicative, 99
- k -multiplicative
 - Anti-, 82
 - Completely, 82
- Absolutely presentable, 149
- Abstract
 - language, 39
 - style, 39
- Accessible category, 142
- Accessible functor, 143
- Algebraic interpretation map, 53
- Anti-preservation
 - of down-directed meets, 79
 - of joins, 79
 - of meets, 79
 - of up-directed joins, 79
- Antitone map, 76
- Atomic object, 177
- Axiom-based axiomatization, 46
- Axiomatization of nabla logics, 47
- BAE, *see* boolean algebra expansion,
 - see* Boolean algebra with operators
- Base, 38, 125
- Boolean algebra expansion, 16
- Boolean algebra with operators, 16
- Canonical
 - equation, 103
 - inequation, 103
- Canonical diagram, 144
- Canonical extension, 70
 - of a boolean algebra, 75
- Canonical model, 222
 - on an algebra, 223
- Canonical presentation, 175, 176
- Category of elements, 136
- Classical modal logic, 245
- Closed element, 69
- Co-continuous, 13
- Co-continuous functor, 126
- Co-finitary version, 258
- Coalgebra, 41
- Coalgebraic Sahlqvist formula
 - abstract, 281
 - concrete, 281
- Cofinal functor, 145
- Compact, 69
- Compatible semantics, 196
- Congruence, 51
- Congruence rule, 46, 51
- Continuous functor, 12, 126
- Dense, 69
 - join-dense, 69
 - meet-dense, 69
- Direct image, 128
- Directed colimit, 141
- Directed set, 141
- Distributive lattice expansion, 16
- Distributive law of a monad, 223
- DLE, *see* Distributive lattice expansion

- Effective equivalence relation, 29
- Egli-Milner order, 226
- Equational logic, 57
- Equivalence relation, 27
- Exact category, 29
- Exact sequence, 29
- Extension of map, 76
- Factorization system, 129
- Filter, 125
- Filter functor, 125
- Filtered category, 142
- Filtered colimit, 143
- Finitary adjunction, 183
- Finitary monad, 183
- Finitary version, 38
- Finitely generated
 - algebraic object, 13
- Finitely presentable
 - functor, 149
- Finitely presented
 - algebraic object, 13
- Finitely strongly presented
 - functor, 152
- Frame
 - coalgebraic, 41
- Free cocompletion, 167
 - filtered, 167
 - finite, 167
- Fully invariant
 - relation, 34
 - subobject, 34
- Generatable, 163
- Generated functor, 160, 173
- Grothendieck construction, 137
- Heyting algebra, 119
- Initial functor, 258
- Intersection, 127
- Intuitionistic logic, 119
- Inverse image, 127
- Isotone map, 76
- Jónsson-Tarski extension, 223
- Jónsson-Tarski theorem, 221
- Join irreducible, 73
 - Completely, 73
- Join prime, 73
 - Completely, 73
- k-multiplicative, 82
- Kan extension, 144, 259
- Lindenbaum-Tarski algebra, 270, 280
- Locally presentable category, 142
- Logic, 46
- Model
 - coalgebraic, 41
- Monadic, 12
- Monotone map, 76
- Nabla
 - language, 38
 - operator, 38
 - style, 38
- Negative term, 112
- Neighbourhood functor, 125
- One step completeness, 46
- One step rule, 45
- One step soundness, 46
- Open element, 69
- Operator, 16
- Pointed set, 138
- Positive term, 112
- Powerset
 - contravariant functor, 39
- Predicate Lifting
 - language, 37
 - style, 37
- Predicate lifting, 41
- Presentable, 141
- Presentable functor, 149

- Presentation, 140
- Presentation by generators and relations, 13
- Presented functor, 152, 173
- Preservation
 - of colimits, 126
 - of limits, 126
- Prime filter
 - functor, 39
- Quasi-canonical
 - embedding, 234
- Quasi-equation, 110
- Rank 1 axiom, 45
- Redistribution, 47
- Regular category, 13
- Regular projective, 13
- Regularly monadic, 20
- Relational logic
 - positive, 225
- Residuated lattice, 120
- Rule-based axiomatization, 46
- Satisfiability
 - in the predicate lifting style, 41
- Satisfiable
 - algebraically, 54
- Scott topology, 93
- Semantic transformation, 43
- Semantics
 - abstract style, 43
 - algebraic, 53
 - nabla style, 42
 - predicate lifting style, 41
- Sharply smaller, 143
- Silm redistribution, 47
- Smooth, 266
- Smooth map, 76, 79
- Standard functor, 38
- Strong completeness
 - over finite algebras, 241
- Strongly generated functor, 161, 173
- Strongly presented functor, 151, 173
- Subformula, 38
- Subobject functor, 130
- Substitution, 50
- Substructural logics, 120
- Term function, 104
- Translation
 - semantics, 193
 - syntax, 189
- Ultrafilter
 - functor, 39
- Up-set
 - functor, 39
- Validity, 41
 - algebraic, 54
- Varietal, 20
- Variator, 20
- Variety, 55
- Weak preservation
 - of colimits, 126
 - of limits, 126
- Well-powered, 125
- Yoneda Lemma, 136

Bibliography

- [Abr09] Samson Abramsky. Coalgebras, Chu Spaces, and Representations of Physical Systems. *CoRR*, abs/0910.3959, 2009.
- [AGT10] Jirí Adámek, Hans Peter Gumm, and Vera Trnková. Presentation of set functors: A coalgebraic perspective. *J. Log. Comput.*, 20(5):991–1015, 2010.
- [AHS04] Jirí Adámek, Horst Herrlich, and George E. Strecker. *Abstract and Concrete Categories The Joy of Cats*. GNU Free Documentation License, Version 1.2, 2004.
- [AM06] Jirí Adámek and Stefan Milius. Terminal coalgebras and free iterative theories. *Inf. Comput.*, 204(7):1139–1172, 2006.
- [AMM10] Jirí Adámek, Stefan Milius, and Lawrence Moss. Initial Algebras and Terminal Coalgebras: a survey. Draft, available from: <https://www.tu-braunschweig.de/iti/mitarbeiter/adamek/artikel>, 2010.
- [AMM12] Jirí Adámek, Stefan Milius, and Lawrence S. Moss. On finitary functors and their presentations. In Dirk Pattinson and Lutz Schröder, editors, *CMCS*, volume 7399 of *Lecture Notes in Computer Science*, pages 51–70. Springer, 2012.
- [AMV11] Jirí Adámek, Stefan Milius, and Jirí Velebil. Semantics of higher-order recursion schemes. *Logical Methods in Computer Science*, 7(1), 2011.
- [AR94] Jirí Adámek and Jirí Rosický. *Locally Presentable and Accessible Categories*. Number 189 in London Mathematical Society Lecture Note Series. Cambridge University Press, 1994.
- [ARV10] Jirí Adámek, Jirí Rosický, and Enrico Maria Vitale. *Algebraic Theories: A Categorical Introduction to General Algebra*. Number 184

- in Cambridge Tracts in Mathematics. Cambridge University Press, 210.
- [AT90] Jirí Adámek and Vera Trnková. *Automata and Algebras in Categories*. Mathematics and Its Applications. Kluwer Academic Publishers, 1990.
- [Awo10] Steve Awodey. *Category Theory*. Oxford Logic Guides. Oxford University Press, second edition, 2010.
- [BD75] Raymond Balbes and Philip Dwinger. *Distributive lattices*. University of Missouri Press, 1975.
- [BdRV01] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.
- [Bez01] Guram Bezhanishvili. Locally finite varieties. *Algebra Universalis*, 46:531–548, 2001.
- [BK06] Marcello M Bonsangue and Alexander Kurz. Presenting functors by operations and equations. In *Foundations of Software Science and Computation Structures*, pages 172–186. Springer, 2006.
- [BKV13] Adriana Balan, Alexander Kurz, and Jirí Velebil. Positive fragments of coalgebraic logics. In Reiko Heckel and Stefan Milius, editors, *Algebra and Coalgebra in Computer Science - 5th International Conference, CALCO 2013, Warsaw, Poland, September 3-6, 2013. Proceedings*, volume 8089 of *Lecture Notes in Computer Science*, pages 51–65. Springer, 2013.
- [Bor94a] Francis Borceux. *Handbook of Categorical Algebra: Volume 1, Basic Category Theory*. Cambridge University Press, 1994.
- [Bor94b] Francis Borceux. *Handbook of Categorical Algebra: Volume 2, Categories and Structures*. Cambridge University Press, 1994.
- [Cal03] Danny Calegari. Problems in foliations and laminations of 3-manifolds. In Clint McCrory Gordana Matic, editor, *Topology and Geometry of Manifolds: 2001 Georgia International Topology Conference, May 21-June 2, 2001, University of Georgia, Athens, Georgia*, volume 71 of *Proceedings of symposia in pure mathematics*, page 357. American Mathematical Soc., 2003.
- [CGvR11] Dion Coumans, Mai Gehrke, and Lrijn van Rooijen. Relational semantics for a fragment of linear logic. *Proceedings of PhDs in Logic III*, 2011.

- [CGZ07] Cristiano Calcagno, Philippa Gardner, and Uri Zarfaty. Context logic as modal logic: Completeness and parametric inexpressivity. In *Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '07, pages 123–134, New York, NY, USA, 2007. ACM.
- [CKW91] Aurelio Carboni, Gregory Maxwell Kelly, and Richard J. Wood. A 2-categorical approach to change of base and geometric morphisms i. *Cahiers de Topologie et Geometrie Differentielle Catgoriques*, 32(1):47–95, 1991.
- [DGP05] Jon Michael Dunn, Mai Gehrke, and Alessandra Palmigiano. Canonical Extensions and Relational Completeness of Some Substructural Logics. *The Journal of Symbolic Logic*, 70(3):pp. 713–740, 2005.
- [DH01] Jon Michael Dunn and Gary Hardegree. *Algebraic Methods in Philosophical Logic*. OUP Oxford, 2001.
- [DP13] Fredrik Dahlqvist and Dirk Pattinson. Some Sahlqvist completeness results for coalgebraic logics. In Frank Pfenning, editor, *FoSSaCS*, volume 7794 of *Lecture Notes in Computer Science*, pages 193–208. Springer, 2013.
- [DP15] Fredrik Dahlqvist and David Pym. Completeness-via-canonicity for distributive substructural logics, a coalgebraic perspective. In *Relational and Algebraic Methods in Computer Science*, 2015.
- [dRV95] Maarten de Rijke and Yde Venema. Sahlqvist’s Theorem For Boolean Algebras With Operators With An Application To Cylindric Algebras. *Studia Logica*, 1995.
- [Dun95] Jon Michael Dunn. Positive modal logic. *Studia Logica*, 55(2):301–317, 1995.
- [F⁺72] Melvin Fitting et al. Tableau methods of proof for modal logics. *Notre Dame Journal of Formal Logic*, 13(2):237–247, 1972.
- [FBC85] Maurizio Fattorosi-Barnaba and Francesco De Caro. Graded Modalities (I). *Studia Logica*, 44(2), 1985.
- [Fin72] Kit Fine. In so many possible worlds. *Notre Dame J. Formal Logic*, 13(4):516–520, 10 1972.
- [Fre71] Peter Freyd. Several new concepts: lucid and concordant functors, pre-limits, pre-completeness, the continuous and concordant completions of categories. In *Lecture notes in mathematics*, volume 99, pages 196–241. Springer, Berlin, 1971.

- [Geh06] Mai Gehrke. Generalized Kripke frames. *Studia Logica*, 84(2):241–275, 2006.
- [GH01] Mai Gehrke and John Harding. Bounded lattice expansions. *Journal of Algebra*, 238(1):345 – 371, 2001.
- [GH09] Steven Givant and Paul Halmos. *Introduction to Boolean Algebras*. Springer, 2009.
- [Gie80] G. Gierz. *A Compendium of continuous lattices*. Springer, 1980.
- [GJ94] Mai Gehrke and Bjarni Jónsson. Bounded distributive lattices with operators. *Mathematica Japonica*, 40, No. 2:207–215, 1994.
- [GJ04] Mai Gehrke and Bjarni Jónsson. Bounded distributive lattice expansions. *Mathematica Scandinavica*, 94:13–45, 2004.
- [GJKO07] Nikolaos Galatos, Peter Jipsen, Tomasz Kowalski, and Hiroakira Ono. *Residuated Lattices: An Algebraic Glimpse at Substructural Logics: An Algebraic Glimpse at Substructural Logics*, volume 151. Elsevier, 2007.
- [GJP13] Mai Gehrke, Ramon Jansana, and Alessandra Palmigiano. Δ_1 - completions of a poset. *Order*, 30(1):39–64, 2013.
- [GNV05] Mai Gehrke, Hideo Nagahashi, and Yde Venema. A Sahlqvist theorem for distributive modal logic. *Ann. Pure Appl. Logic*, 131(1-3):65–102, 2005.
- [Gol89] Robert Goldblatt. Varieties of complex algebras. *Annals of pure and applied logic*, 44:173–242, 1989.
- [Gol95] Robert Goldblatt. Elementary generation and canonicity for varieties of boolean algebras with operators. *Algebra Universalis*, 34:551–607, 1995.
- [Gol06] Robert Goldblatt. *Topoi, The Categorical Analysis of Logic*. Dover, 2nd edition, 2006.
- [Gum05] Hans Peter Gumm. From T-coalgebras to Filter Structures and Transition Systems. In *Algebra and Coalgebra in Computer Science*, pages 194–212. Springer, 2005.
- [GV99] Steven Givant and Yde Venema. The preservation of Sahlqvist equations in completions of boolean algebras with operators. *Algebra Universalis*, 41:47–84, 1999.

- [Hen70] Leon Henkin. Extending Boolean Operations. *Pacific Journal of Mathematics*, 32(3), 1970.
- [HL71] Joseph D. Halpern and Azriel Levy. The boolean prime ideal theorem does not imply the axiom of choice. *Axiomatic Set Theory*, 13(Part 1):83–134, 1971.
- [Hug01] Jesse Hughes. *A Study of Categories of Algebras and Coalgebras*. PhD thesis, Carnegie Mellon University, May 2001.
- [Joh86] Peter Johnstone. *Stone Spaces*. Cambridge studies in advanced mathematics 3. Cambridge University Press, 1986.
- [Jón94] Bjarni Jónsson. On the canonicity of Sahlqvist identities. *Studia Logica*, 53(4):473–492, 1994.
- [JS10] Bart Jacobs and Ana Sokolova. Exemplaric expressivity of modal logics. *J. Log. and Comput.*, 20:1041–1068, October 2010.
- [JT51] Bjarni Jónsson and Alfred Tarski. Boolean algebras with operators. part 1. *Amer. J. Math.*, 33:891–937, 1951.
- [Kel82] Gregory Maxwell Kelly. *Basic Concepts of Enriched Category Theory*. Number 64 in Lecture Notes in Mathematics. Cambridge University Press, 1982.
- [Kel91] Gregory Maxwell Kelly. A note on relations relative to a factorization system. In *Category Theory, Proceedings of the International Conference held in Como, Italy, July 2228, 1990*, volume 1488, pages pp 249–261. Springer Berlin Heidelberg, 1991.
- [KKP04] Clemens Kupke, Alexander Kurz, and Dirk Pattinson. Algebraic semantics for coalgebraic logics. In J. Adámek and S. Milius, editors, *Coalgebraic Methods in Computer Science (CMCS 2004)*, volume 106 of *Electr. Notes in Theoret. Comp. Sci.*, pages 219–241, 2004.
- [KKP05] Clemens Kupke, Alexander Kurz, and Dirk Pattinson. Ultrafilter Extensions for Coalgebras. In José Luiz Fiadeiro, Neil Harman, Markus Roggenbach, and Jan J. M. M. Rutten, editors, *CALCO*, volume 3629 of *Lecture Notes in Computer Science*, pages 263–277. Springer, 2005.
- [KKV12a] Krzysztof Kapulkin, Alexander Kurz, and Jirí Velebil. Expressiveness of positive coalgebraic logic. In Thomas Bolander, Torben Braüner, Silvio Ghilardi, and Lawrence S. Moss, editors, *Advances in Modal Logic 9, papers from the ninth conference on "Advances*

- in Modal Logic,” held in Copenhagen, Denmark, 22-25 August 2012*, pages 368–385. College Publications, 2012.
- [KKV12b] Clemens Kupke, Alexander Kurz, and Yde Venema. Completeness for the coalgebraic cover modality. *Logical Methods in Computer Science*, 8(3), 2012.
- [KL09] Alexander Kurz and Raul Leal. Modalities in the Stone age: A comparison of coalgebraic logics. In *MFPS XXV*, Oxford, 2009.
- [KP93] Gregory Maxwell Kelly and John Power. Adjunctions whose counits are coequalizers, and presentations of finitary enriched monads. *J Pure Appl. Algebra*, 89:163–179, 1993.
- [KP10] Alexander Kurz and Daniela Petrişan. Presenting functors on many-sorted varieties and applications. *Inf. Comput.*, 208(12):1421–1446, December 2010.
- [KP11] Clemens Kupke and Dirk Pattinson. Coalgebraic semantics of modal logics: an overview. *Theoretical Computer Science*, 412(38):5070–5094, 2011. Special issue CMCS 2010.
- [KR07] Alexander Kurz and Jirí Rosický. The Goldblatt-Thomason Theorem for Coalgebras. In Till Mossakowski, Ugo Montanari, and Magne Haveraaen, editors, *CALCO*, volume 4624 of *Lecture Notes in Computer Science*, pages 342–355. Springer, 2007.
- [KR12] Alexander Kurz and Jirí Rosický. Strongly complete logics for coalgebras. *Logical Methods in Computer Science*, 8(3), 2012.
- [KRV05] Panagis Karazeris, Jirí Rosický, and Jirí Velebil. Completeness of co-completions. *Journal of Pure and Applied Algebra*, 196(2-3):229250, April 2005.
- [Kur00] Alexander Kurz. *Logics for Coalgebras and Applications to Computer Science*. PhD thesis, Ludwig Maximilian University of Munich, 2000.
- [Lea11] Raul Leal. *Modalities Through the Looking Glass*. PhD thesis, Institute for Logic, Language and Computation, 2011.
- [LM92] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in Geometry and Logic. A First Introduction to Topos Theory*. Springer, 1992.
- [Mac98] Saunders MacLane. *Categories for the Working Mathematician (Graduate Texts in Mathematics)*. Springer, 2nd edition, 1998.

- [Mol68] D. I. Moldavanskii. Intersection of finitely generated subgroups. *Siberian Mathematical Journal*, 9(6):1066–1069, 1968.
- [Mon67] Luiz Monteiro. Une construction du réticulé distributif libre sur un ensemble ordonné. *Colloquium Mathematicum*, 17:23–27, 1967.
- [MP89] Michael Makkai and Robert Paré. *Accessible Categories: The Foundations of Categorical Model Theory*. Number 104 in Contemporary Mathematics. American Mathematical Soc., 1989.
- [Mye11] Robert Myers. *Rational Coalgebraic Machines in Varieties: Languages, Completeness and Automatic Proofs*. PhD thesis, Imperial College London, 2011.
- [Ono03] Hiroakira Ono. Substructural logics and residuated lattices – an introduction. *Trends in Logic*, 20:177–212, 2003.
- [Pat03] Dirk Pattinson. Coalgebraic modal logic: Soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309(1–3):177–193, 2003.
- [Ple00] Till Plewe. Quotient maps of locales. *Applied Categorical Structures*, 8(1-2):17–44, 2000.
- [PS08] Dirk Pattinson and Lutz Schröder. Beyond rank 1: Algebraic semantics and finite models for coalgebraic logics. In R. Amadio, editor, *Proc. FoSSaCS 2008*, number 4962 in LNCS, pages 66–80, 2008.
- [Rib52] Hugo Ribeiro. A remark on boolean algebras with operators. *Amer. J. Math.*, 74:162–167, 1952.
- [Rut97] John W. Rutter. *Spaces of homotopy self-equivalences: a survey*. Number 1662 in Lecture Notes in Mathematics. Springer, 1997.
- [Rut00] Jan J. M. M. Rutten. Universal coalgebra: a theory of systems. *Theor. Comput. Sci.*, 249(1):3–80, 2000.
- [Sah75] Henrik Sahlqvist. Completeness and Correspondence in the First and Second Order Semantics for Modal Logic. In S. Kanger, editor, *Proceedings of the Third Scandinavian Logic Symposium*, volume 82 of *Studies in Logic and the Foundations of Mathematics*, pages 110–143. Elsevier, 1975.
- [Sch06] Lutz Schröder. A finite model construction for coalgebraic modal logic. In Luca Aceto and Anna Ingólfssdóttir, editors, *Foundations Of Software Science And Computation Structures*, volume 3921 of *Lecture Notes in Computer Science*, pages 157–171. Springer; Berlin; <http://www.springer.de>, 2006.

- [SP06] Lutz Schröder and Dirk Pattinson. PSPACE bounds for rank-1 modal logics. In *Proc. 21st IEEE Symposium on Logic in Computer Science (LICS 2006)*, pages 231–242, 2006.
- [SP09] Lutz Schröder and Dirk Pattinson. Strong completeness of coalgebraic modal logics. In S. Albers and J.-Y. Marion, editors, *Proc. STACS 2009*, volume 09001 of *Dagstuhl Seminar Proceedings*, pages 673–684. Schloss Dagstuhl, 2009.
- [Suz11] Tomoyuki Suzuki. Canonicity results of substructural and lattice-based logics. *The Review of Symbolic Logic*, 4:1–42, 3 2011.
- [Teh08] Bruno Teheux. *Algebraic Approach to Modal Extensions of Lukasiewicz Logics*. PhD thesis, Université de Liège, 2008.
- [Trn69] Vera Trnková. Some properties of set functors. *Comment. Math. Univ. Carolinae*, 10:323–352, 1969.
- [Ven06] Yde Venema. Algebras and coalgebras. In Johan van Benthem, Patrick Blackburn, and Frank Wolter, editors, *Handbook of Modal Logic*. Elsevier, 2006.
- [VK11] Jirí Velebil and Alexander Kurz. Equational presentations of functors and monads. *Mathematical Structures in Computer Science*, 21(2):363–381, 2011.
- [Vos10] Jacob Vosmaer. *Logic, Algebra and Topology. Investigations into canonical extensions, duality theory and point-free topology*. PhD thesis, ILLC Dissertation Series, 2010.