Completeness via canonicity for distributive substructural logics: a coalgebraic perspective

Fredrik Dahlqvist and David Pym

University College London

Abstract. We prove strong completeness of a range of substructural logics with respect to their relational semantics by completeness-via-canonicity. Specifically, we use the topological theory of canonical (in) equations in distributive lattice expansions to show that distributive substructural logics are strongly complete with respect to their relational semantics. By formalizing the problem in the language of coalgebraic logics, we develop a modular theory which covers a wide variety of different logics under a single framework, and lends itself to further extensions.

1 Introduction

This work lies at the intersection of resource semantics/modelling, substructural logics, and the theory of canonical extensions and canonicity. These three areas respectively correspond to the semantic, proof-theoretic, and algebraic sides of the problem we tackle: to give a systematic, modular account of the relation between resource semantics and logical structure. We do not delve into the proof theory of substructural logics, but rather deal with the algebraic formulations of many such substructural proof systems ([29] summarizes the correspondence between classes of residuated lattices and substructural logics). A version of this work that includes detailed proofs can be found as a UCL Research Note [12].

Resource semantics and modelling. Resource interpretations of substructural logics — see, for example, [18,30,31,15,7] — are well-known and exemplified in the context of program verification and semantics by Ishtiaq and O'Hearn's pointer logic [23] and Reynolds' separation logic [32], each of which amounts to a model of a specific theory in Boolean BI. Resource semantics and modelling with resources has become an active field of investigation in itself (see, for example, [8]). Certain requirements, discussed below, seem natural (and useful in practice) in order to model naturally arising examples of resource.

1. We need to be able to compare at least some resources. Indeed, in a completely discrete model of resource (i.e., where no two resources are comparable) it is impossible to model key concepts such as 'having enough resources'. On the other hand, there is no reason to assume that *any two* resources be comparable (e.g., heaps). This suggests at least a preorder structure on models. In fact, we take the view that comparing two resources is fundamental,

and in particular, if two resources cannot be distinguished in this way then they can be identified. We thus add antisymmetry and work with posets.
2. We need to be able to combine (some) resources to form new resources (e.g., union of heaps with disjoint domains [23]). We denote the combination operation by $*$. An equivalent, but often more useful, point of view is to be able to specify how resources can be 'split up' into pairs of constituent resources. Moreover, since comparing resources is more important than establishing their equality, it makes sense to be able to list for a given resource r, the pairs (s_1, s_2) of resources which combine to form a resource $s_1 * s_2 \leq r$.
3. All reasonable examples of resources possess 'unit' resources with respect to the combination operation $*$; that is, special resources that leave other resources unchanged under the combination operation.
4. The last requirement is crucial, but slightly less intuitive. In the most well-behaved examples of resource models (e.g., $\mathbb{N}$), if we are given a resource r and a 'part' s of r, there exists a resource s' that 'completes' s to make r; that is, we can find a resource s' such that $s * s' = r$. More generally, given two resources r, s, we want to be able to find the the best s' such that $s*s' \leq r$. In a model of resource without this feature, it is impossible to provide an answer to legitimate questions such as 'how much additional resource is needed to make statement ϕ hold?'. Mathematically, this requirement says that the resource composition is a residuated mapping in both its arguments.

The literature on resource modelling, and on separation logic in particular, is vast, but two publications ([6] and [4]) are strongly related to this work. Both show completeness of 'resource logics' by using Sahlqvist formulas, which amounts to using completeness-via-canonicity ([3,24]).

Completeness-via-canonicity and substructural logics. The logical side of resource modelling is the world of substructural logics, such as BI, and of their algebraic formulations; that is, residuated lattices, residuated monoids, and related structures. The past decade has seen a fair amount of research into proving the completeness of relational semantics for these logics (for BI, for example, [31,15]), using, among other approaches, techniques from the duality theory of lattices. In [13], Dunn *et al.* prove completeness of the full Lambek calculus and several other well-known substructural logics with respect to a special type of Kripke semantics by using duality theory. This type of Kripke semantics, which is two-sorted in the non-distributive case, was studied in detail by Gerhke in [16]. The same techniques have been applied to prove Kripke completeness of fragments of linear logic in [5]. Finally, the work of Suzuki [33] explores in much detail completeness-via-canonicity for substructural logics. Our work follows in the same vein but with with some important differences. Firstly, we use a dual adjunction rather than a dual equivalence to connect syntax and semantics. This is akin to working with Kripke frames rather than descriptive general frames in modal logics: the models are simpler and more intuitive, but the tightness of the fit between syntax and semantics is not as strong. Secondly, we use the topological approach to canonicity of [17,21,34] because we feel it is the most flexible and modular approach to building canonical (in)equations. Thirdly, we only consider distributive structures. This is to some extent a matter a taste. Our choice is driven by the desire to keep the theory relatively simple (the non-distributive case is more involved), by the fact that from a resource modelling perspective the non-distributive case does not seem to occur 'in the wild', and finally because we place ourselves in the framework of coalgebraic logic, where the category of distributive lattices forms a particularly nice 'base category'.

The coalgebraic perspective. Coalgebraic methods bring many advantages to the study of completeness-via-canonicity. First, it greatly clarifies the connection between canonicity as an algebraic method and the existence of 'canonical models'; that is, strong completeness. Second, it provides a generic framework in which to prove completeness-via-canonicity for a vast range of logics ([11]). Third, it is intrinsically modular; that is, it provides theorems about complicated logics by combining results for simpler ones ([9,10]).

2 Substructural logics: a coalgebraic perspective

We use the 'abstract' version of coalgebraic logic developed in, for example, [27], [28] and [25]; that is, we require the following basic situation:

$$\begin{tikzcd} \mathscr{C} \arrow[loop left, "L"] \arrow[rr, bend left, "F"] \arrow[rr, phantom, "\perp"] && \mathscr{D}^{\mathrm{op}} \arrow[loop right, "T^{\mathrm{op}}"] \arrow[ll, bend left, "G"] \end{tikzcd} \tag{1}$$

The left hand-side of the diagram is the syntactic side, and the right-hand side the semantic one. The category $\mathscr{C}$ represents a choice of 'reasoning kernel'; that is, of logical operations which we consider to be fundamental, whilst L is a syntax constructing functor which builds terms over the reasoning kernel. Objects in $\mathscr{D}$ are the carriers of models and T specifies the coalgebras on these carriers in which the operations defined by L are interpreted. The functors F and G relate the syntax and the semantics, and F is left adjoint to G. We will denote such an adjunction by $F \dashv G : \mathscr{C} \to \mathscr{D}$. Note, as mentioned in the introduction, that we only need a dual adjunction, not a full duality.

2.1 Syntax

Reasoning kernels. There are three choices for the category $\mathscr{C}$ which are particularly suited to our purpose, the category $\mathbf{DL}$ of distributive lattices, the category $\mathbf{BDL}$ of bounded distributive lattices, and the category $\mathbf{BA}$ of boolean algebras. The choice of $\mathbf{DL}$ as our most basic category was justified in the introduction, but we should also mention an important technical advantage of $\mathbf{DL}, \mathbf{BDL}$ and $\mathbf{BA}$ from the perspective of coalgebraic logic: each category is locally finite; that is, finitely generated objects are finite. This is a very desirable technical property for the presentation of endofunctors on this category and for coalgebraic strong completeness theorems. We denote by $\mathsf{F} \dashv \mathsf{U}$ the usual free-forgetful adjunction between $\mathbf{DL}$ (resp. $\mathbf{BDL}$, resp. $\mathbf{BA}$) and $\mathbf{Set}$.

True and false. The choice of including (or not) $\top$ and $\bot$ to the logic is clearly provided by the choice of reasoning kernel.

Algebras. Recall that an algebra for an endofunctor $L : \mathscr{C} \to \mathscr{C}$ is an object A of $\mathscr{C}$ together with a morphism $\alpha : LA \to A$. We refer to endofunctors $L : \mathscr{C} \to \mathscr{C}$ as *syntax constructors*.

Intuitionistic implication. We do not consider the intuitionistic implication as a fundamental operation; in particular, the category of Heyting algebras does not form a reasoning kernel. This choice is motivated by the fact that the semantics of intuitionistic logic can be given in terms of Kripke frames, that the intuitionistic implication is not usually part of the basic language of substructural logics, and that the category **HA** of Heyting algebras is not as well-behaved as our choices of reasoning kernels. We therefore add the implication as an additional (modal) operation on (bounded) distributive lattices via the syntax constructor:

$$L_{\mathrm{Hey}} : \mathbf{DL} \to \mathbf{DL}, \begin{cases} A \mapsto \mathsf{F}\{a \to b \mid a, b \in \mathsf{U}A\}/\equiv \\ L_{\mathrm{RL}} f : L_{\mathrm{Hey}} A \to L_{\mathrm{Hey}} B, [a]_\equiv \mapsto [f(a)]_\equiv, \end{cases}$$

where $\equiv$ is the fully invariant equivalence relation in **DL** generated by the following Heyting Distribution Laws for *finite* subsets X of A:

HDL1. $a \to \bigwedge X = \bigwedge[a \to X]$
HDL2. $\bigvee X \to a = \bigwedge[X \to a]$.

where we use the notation $\bigwedge[a \to X] := \bigwedge_{x \in X} a \to x$ and the convention that $\bigwedge \emptyset = \top$ and $\bigvee \emptyset = \bot$ when the objects of the reasoning kernel are bounded. The language defined by L_{Hey} for a set V of propositional variables is the free L_{Hey}-algebra over $\mathsf{F}V$; that is, the language of intuitionistic propositional logic quotiented by the axioms of 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). The remaining features of HAs will be captured in a second stage via *canonical frame conditions*. The reason for proceeding in this step-by-step way will become clear in the sequel and is similar in spirit to the approach of [1]. The main difference is that in [1], the axioms of Heyting algebras are separated into rank 1 and non-rank 1 axioms, leading to the notion of *weak Heyting algebras* which obey the axioms HDL1-2 and also $a \to a = \top$. In this work, we want to build a minimal 'pre-Heyting' logic with a strongly complete semantics and well-behaved (viz. smooth, see Section) operations, and L_{Hey}-algebras perform this role.

Resource operations. The operations on resources specified in the introduction; that is, a combination operation and its left and right residuals, are introduced via the following syntax constructor:

$$L_{\mathrm{RL}} : \mathscr{C} \to \mathscr{C}, \begin{cases} L_{\mathrm{RL}} A = \mathsf{F}\{I, a * b, a \backslash b, a/b \mid a, b \in \mathsf{U}A\}/\equiv \\ L_{\mathrm{RL}} f : L_{\mathrm{RL}} A \to L_{\mathrm{RL}} B, [a]_\equiv \mapsto [f(a)]_\equiv \,, \end{cases}$$

where $\equiv$ is the fully invariant equivalence relation in $\mathscr{C}$ generated by following the Distribution Laws for finite subsets X of A:

DL1. $\bigvee X * a = \bigvee[X * a]$
DL2. $a * \bigvee X = \bigvee[a * X]$
DL3. $a \backslash \bigwedge X = \bigwedge[a \backslash X]$
DL4. $\bigvee X \backslash a = \bigwedge[X \backslash a]$
DL5. $\bigwedge X / a = \bigwedge[X / a]$
DL6. $a / \bigvee X = \bigwedge[a / X]$.

The language defined by L_{RL} is the free L_{RL}-algebra over $\mathsf{F}V$, which is the language of the distributive full Lambek calculus (or residuated lattices) quotiented under the axioms of $\mathscr{C}$ and DL1-6. An L_{RL}-algebra is simply an object of $\mathscr{C}$ endowed with a nullary operation I and binary operations $*, \backslash$ and $/$ satisfying the distribution laws above. 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 subsequently as *canonical frame conditions*. Both L_{Hey}-algebras and L_{RL}-algebras are examples of *Distributive Lattice Expansions*, or DLEs; that is, distributive lattices endowed with a collection of maps of finite arities. When $\mathscr{C} = \mathbf{BA}$, L_{RL}-algebras are an example *Boolean Algebra Expansions*, or BAEs.

Modularity. The syntax developed above is completely modular. For example, if we wish to study boolean BI, it is natural to consider $L_{\mathrm{RL}} : \mathbf{BA} \to \mathbf{BA}$ as our syntax constructor. If we wish to study intuitionistic BI, then we should consider

$$L_{\mathrm{Hey}} + L_{\mathrm{RL}} : \mathbf{BDL} \to \mathbf{BDL}, \text{ where } (L_{\mathrm{Hey}} + L_{\mathrm{RL}})A = L_{\mathrm{Hey}}A + L_{\mathrm{RL}}A,$$

where the coproduct is taken in $\mathbf{BDL}$, and is thus a 'free product' generating precisely the expected language. Finally, we may wish to add modal operators to the language (see the 'relevant modal logic' in [33]), for example $\Diamond$. In this case, we can in the same way add the syntax constructor for modal logic, namely,

$$L_{\Diamond} : \mathscr{C} \to \mathscr{C}, A \mapsto \mathsf{F}\{\Diamond a \mid a \in \mathsf{U}A\} / \{\Diamond(\bigvee X) = \bigvee[\Diamond X]\}$$

2.2 Coalgebraic semantics

Semantic domain. As we mentioned in the introduction, it is reasonable to assume that a model of resources should be a poset, and thus taking $\mathscr{D} = \mathbf{Pos}$ is intuitively justified. This is a particularly attractive choice of 'semantic domain' given that the category $\mathbf{Pos}$ is related to $\mathbf{DL}$ by the dual adjunction $\mathsf{Pf} \dashv \mathcal{U} : \mathbf{DL} \to \mathbf{Pos}^{\mathrm{op}}$, where Pf is the functor sending a distributive lattice to its poset of prime filters, and $\mathbf{DL}$-morphisms to their inverse images, and $\mathcal{U}$ is the functor sending a poset to the distributive lattice of its up-sets and monotone maps to their inverse images. In the case in which 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; that is, ultrafilters are only related to themselves. Thus the dual adjunction $\mathsf{Pf} \dashv \mathcal{U}$ becomes the well-known adjunction $\mathsf{Uf} \dashv \mathsf{P}_c : \mathbf{BA} \to \mathbf{Set}^{\mathrm{op}}$.

Coalgebras. Recall that a coalgebra for an endofunctor $T : \mathscr{D} \to \mathscr{D}$, is an object W of $\mathscr{D}$ together with a morphism $\gamma : W \to TW$. The endofunctors that we will consider are built from products and 'powersets' and will be referred to as *model constructors*. Note that **Pos** has products, which are simply the **Set** products with the obvious partial order on pairs of elements. The 'powerset' functor which we will consider is the *convex powerset* functor: $\mathsf{P}_c : \mathbf{Pos} \to \mathbf{Pos}$, sending a poset to its set of convex subsets, where a subset U of a poset $(X, \leq)$ is convex if $x, z \in U$ and $x \leq y \leq z$ implies $y \in U$. The set $\mathsf{P}_c X$ is given a poset structure via the *Egli-Milner* order (see [2]).

Coalgebras for the intuitionistic implication. We define the following model constructor, which will interpret $\to$:

$$T_{\mathrm{Hey}} : \mathbf{Pos} \to \mathbf{Pos}, \begin{cases} T_{\mathrm{Hey}} W = \mathsf{P}_c(W^{\mathrm{op}} \times W) \\ T_{\mathrm{Hey}} f : T_{\mathrm{Hey}} W \to T_{\mathrm{Hey}} W', U \mapsto (f \times f)[U]. \end{cases}$$

where W^{op} is the poset whose carrier is W and whose order is dual to that of W.

Coalgebras for the resource operations. We define the following model constructor, which is used to interpret $I, *, \backslash$ and $/$:

$$T_{\mathrm{RL}} : \mathscr{D} \to \mathscr{D}, \begin{cases} T_{\mathrm{RL}} W = \mathbb{2} \times \mathsf{P}_c(W \times W) \times \mathsf{P}_c(W^{\mathrm{op}} \times W) \times \mathsf{P}_c(W \times W^{\mathrm{op}}) \\ T_{\mathrm{RL}} f : T_{\mathrm{RL}} W \to T_{\mathrm{RL}} W', U \mapsto (\mathsf{Id}_{\mathbb{2}} \times (f \times f)^3)[U]. \end{cases}$$

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 $\backslash$ and $/$, respectively, and turn out to be very closely related to the second component. Note that if $\mathscr{D} = \mathbf{Pos}$, the structure map of coalgebras are monotone, intuitively this means bigger resources can be split up in more ways.

The semantic transformations. In the abstract flavour of coalgebraic logic, the semantics is provided by a natural transformation $\delta : LG \to GT^{\mathrm{op}}$ called the *semantic transformation*. We show below how this defines an interpretation map, but we first define our two semantic transformations. As already noted above, a $\mathscr{C}$-morphism $\delta_W^{\mathrm{Hey}} : L_{\mathrm{Hey}} GW \to GT_{\mathrm{Hey}} W$ is equivalent to a function over the set of generators $\{U \to V \mid U, V \in \mathsf{U}GW\}$ satisfying the distributivity laws HDL1-2, and similarly for $\delta_W^{\mathrm{RL}} : L_{\mathrm{RL}} GW \to GT_{\mathrm{RL}} W$ and the distributivity laws DL1-6. We now define

$$\delta_W^{\mathrm{Hey}}(U \to V) = \{(x, y) \in T_{\mathrm{Hey}} W \mid x \in U \Rightarrow y \in V\}$$

and similarly (by using the usual projections maps $\pi_i, 1 \leq i \leq 4$)

$$\begin{aligned} \delta_W^{\mathrm{RL}}(I) &= \{t \in T_{\mathrm{RL}} W \mid \pi_1(t) = 0 \in \mathbb{2}\} \\ \delta_W^{\mathrm{RL}}(u * v) &= \{t \in T_{\mathrm{RL}} W \mid \exists (x, y) \in \pi_2(t), x \in u, y \in v\} \\ \delta_W^{\mathrm{RL}}(u \backslash w) &= \{t \in T_{\mathrm{RL}} W \mid \forall (x, y) \in \pi_3(t), x \in u \Rightarrow y \in w\} \\ \delta_W^{\mathrm{RL}}(w / v) &= \{t \in T_{\mathrm{RL}} W \mid \forall (x, y) \in \pi_4(t), x \in v \Rightarrow y \in w\}. \end{aligned}$$

Proposition 1. *The natural transformations δ^{Hey} and δ^{RL} are well-defined, in particular each map δ_W^{Hey} satisfies the distributivity laws HDL1-2, and each map δ_W^{RL} satisfies the distributivity laws the distributivity laws DL1-6.*

The semantic transformations are thus well-defined. We now show how the interpretation map arises from the semantic transformation. Recall that, for a given syntax constructor $L : \mathscr{C} \to \mathscr{C}$, the language of L is the free L-algebra over $\mathsf{F}V$. This is equivalent to saying that it is the initial $L(-) + \mathsf{F}V$-algebra. We use initiality to define the interpretation map by putting an $L(-) + \mathsf{F}V$-algebra structure on the 'predicates' of a T-coalgebra $\gamma : W \to TW$; that is, on the carrier set GW. By definition of the coproduct, this means defining a morphism $LGW \to GW$ and a morphism $\mathsf{F}V \to GW$. By adjointness it is easy to see that the latter is simply a valuation $v : V \to \mathsf{U}GW$. For the former we simply use the semantic transformation and G applied to the coalgebra. The interpretation map $[\![-]\!]_W$ is thus given by the catamorphism:

$$
\begin{tikzcd}
L\mu(L(-)+\mathsf{F}V)+\mathsf{F}V \arrow[r, dashed, "L[\![-]\!]_W+\mathsf{Id}_{\mathsf{F}V}"] \arrow[dd] & LGW+\mathsf{F}V \arrow[d, "\delta_W+\mathsf{Id}_{\mathsf{F}V}"] \\
 & GTW+\mathsf{F}V \arrow[d, "G\gamma+v"] \\
\mu(L(-)+\mathsf{F}V) \arrow[r, dashed, "[\![-]\!]_W"'] & GW
\end{tikzcd}
$$

Modularity. Model constructors and semantic transformations can be assembled in a way that is dual to the the syntax constructors. For example, if we wish to interpret both the intuitionistic implication and the resource operations, we use a coalgebra of type $\gamma_1 \times \gamma_2 : W \to T_{\mathrm{Hey}}W \times T_{\mathrm{RL}}W$. The overall semantics is then inherited from that of the constituents via the following diagram:

$$
\begin{tikzcd}
(L_{\mathrm{Hey}}+L_{\mathrm{RL}})\mu(L_{\mathrm{Hey}}+L_{\mathrm{RL}}(-)+\mathsf{F}V)+\mathsf{F}V \arrow[r, dashed, "L_{\mathrm{Hey}}+L_{\mathrm{RL}}[\![-]\!]_W+\mathsf{Id}_{\mathsf{F}V}"] \arrow[dd] & L_{\mathrm{Hey}}GW+L_{\mathrm{RL}}GW+\mathsf{F}V \arrow[d, "\delta_W^{\mathrm{Hey}}+\delta_W^{\mathrm{RL}}+\mathsf{Id}_{\mathsf{F}V}"'] \\
 & GT_{\mathrm{Hey}}W+GT_{\mathrm{RL}}W+\mathsf{F}V \arrow[d, "G(\gamma_1\times\gamma_2)\circ(G\pi_1+G\pi_2)+v"'] \\
\mu(L_{\mathrm{Hey}}+L_{\mathrm{RL}}(-)+\mathsf{F}V) \arrow[r, dashed, "[\![-]\!]_W"'] & GW
\end{tikzcd}
$$

3 Canonicity

3.1 Canonical extension of distributive lattices

We now briefly present the salient facts about canonical extensions. For more details the reader is referred to [19] for BAs, [26,24] for BAOs, and [20,21] 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'; that is, whose elements can be viewed as parts of a set, or of a set with some additional structure. In this way, we can establish a connection between the syntax and the semantics; that is, 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. For a distributive lattice A, the idea behind the construction of its canonical extension A^σ is to build a completion of A which is not 'too big' and not 'too different' from A. Technically, we want A to be *dense* and *compact* in A^σ.

Density. To build a completion of A it is natural to formally add to 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 [22]). Intuitively, this prevents the completion from becoming 'too big'. Based on this intuition we introduce the following terminology: given a sub-lattice A of a complete distributive lattice C, we define the meets in C of elements of A as the *closed elements* of C and denote this set by $K(A)$ (or simply K when there is no ambiguity); dually, we define the joins in C of elements of A as the *open elements* of C and denote this set by $O(A)$. Finally, we say that A is *dense* in C if $C = O(K(A)) = K(O(A))$.

Compactness. The canonical extension A^σ is also required not to be too different from A in the sense that facts about arbitrary meets and joins of elements of A in A^σ must already be 'witnessed' by finite meets and joins in A. Formally, if A is a sub-lattice of C, A is *compact* in C if, for every $X, Y \subseteq A$ such that $\bigwedge X \leq \bigvee Y$, there exist *finite* subsets $X_0 \subseteq X, Y_0 \subseteq Y$ such that $\bigwedge X_0 \leq \bigvee Y_0$. An equivalent definition is that A is compact in C if for every closed element $p \in K(A)$ and open element $u \in O(A)$ such that $p \leq u$, there exists an element $a \in A$ such that $p \leq a \leq u$. The *canonical extension* A^σ of a distributive lattice A is the complete distributive lattice such that A is dense and compact in A^σ. We can summarize what we need to know about A^σ in the following theorem:

Theorem 1 ([20,17,21]). *The canonical extension A^σ of a distributive lattice A can be concretely represented as the lattice $A^\sigma \simeq \mathcal{U}\mathsf{Pf}A$; in particular, it is completely distributive.*

Note that this theorem requires the Prime Ideal Theorem for distributive lattices which is a non-constructive principle, albeit one that is strictly weaker than the axiom of choice. Note also that since the canonical extension of a BA is complete and completely distributive, it is also *atomic* (see [19] Ch. 14); that is, it is a complete atomic boolean algebra. It is concretely represented by $A^\sigma = \mathsf{P}_c\mathsf{Uf}A$, in which case it is not simply 'set-like', but an actual algebra of subsets.

3.2 Canonical extension of distributive lattice expansions

We now sketch the theory canonical extensions for Distributive Lattice Expansions (DLE) — for the details, see [20,21]. Each map $f : \mathsf{U}A^n \to \mathsf{U}A$ can be extended to a map $(\mathsf{U}A^\sigma)^n \to \mathsf{U}A^\sigma$ in two canonical ways:

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

where $f[d,u] = \{f(a) \mid a \in A^n, d \leq a \leq u\}$. Note that since A is compact in A^σ the intervals $[d,u]$ are never empty, which justifies these definitions. For a signature Σ, the *canonical extension* of a Σ-DLE $(A, (f_s : \mathsf{U}A^{\mathrm{ar}(n)} \to \mathsf{U}A)_{s\in\Sigma})$ is defined to be the Σ-DLE $(A^\sigma, (f_s^\sigma : \mathsf{U}(A^\sigma)^{\mathrm{ar}(n)} \to \mathsf{U}A^\sigma)_{s\in\Sigma})$, and similarly for BAEs. We summarize some important facts about canonical extensions of maps in the following proposition, proofs can be found in, for example, [17,21,34]:

Proposition 2. *Let A be a distributive lattice, and $f : \mathsf{U}A^n \to \mathsf{U}A$.*

1. *$f^\sigma \restriction A^n = f^\pi \restriction A^n = f$.*
2. *$f^\sigma \leq f^\pi$ under pointwise ordering.*
3. *If f is monotone in each argument, then $f^\sigma \restriction (K \cup O)^n = f^\pi \restriction (K \cup O)^n$.*

We call a monotone map $f : \mathsf{U}A^n \to \mathsf{U}A$ *smooth in its i^{th} argument* $(1 \leq i \leq n)$ if, for every $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n \in K \cup O$,

$$f^\sigma(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n) = f^\pi(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_n),$$

for every $x_i \in A^\sigma$. A map $f : \mathsf{U}A^n \to \mathsf{U}A$ is called *smooth* if it is smooth in each of its arguments.

In order to study effectively the canonical extension of maps, we need to define six topologies on A^σ. First, we define $\sigma^\uparrow = \{\uparrow p \mid p \in K\}$, $\sigma^\downarrow = \{\downarrow u \mid u \in O\}$ and $\sigma = \sigma^\uparrow \cup \sigma^\downarrow$; that is, the join of $\sigma^\uparrow$ and $\sigma^\downarrow$ in the lattice of topologies on A^σ. It is easy to check that the sets above do define topologies and that $\sigma = \{\uparrow p \cap \downarrow u \mid K \ni p \leq u \in O\}$. The next set of topologies is well-known to domain theorists: a *Scott open* in A^σ is a subset $U \subseteq A^\sigma$ such that (1) U is an upset and (2) for any up-directed set D such that $\bigvee D \in U$, $D \cap U \neq \emptyset$. The collection of Scott opens forms a topology called the *Scott topology*, which we denote $\gamma^\uparrow$. The dual topology will be denoted by $\gamma^\downarrow$, and their join by γ. It is not too hard to show (see [17,34]) that $\gamma^\uparrow \subseteq \sigma^\uparrow$, $\gamma^\downarrow \subseteq \sigma^\downarrow$, and $\gamma \subseteq \sigma$. We denote the product of topologies by $\times$, and the n-fold product of a topology τ by τ^n. The following result shows why these topologies are important: they essentially characterize the canonical extensions of maps:

Proposition 3 ([17]). *For any DL A and any map $f : \mathsf{U}A^n \to \mathsf{U}A$,*

1. *f^σ is the largest $(\sigma^n, \gamma^\uparrow)$-continuous extension of f,*
2. *f^π is the smallest $(\sigma^n, \gamma^\downarrow)$-continuous extension of f*
3. *f is smooth iff it has a unique (σ^n, γ)-continuous extension.*

From this important result, it is not hard to get the following key theorem, sometimes known as *Principle of Matching Topologies*, which underlies the basic 'algorithm' for canonicity:

Theorem 2 (Principle of Matching Topologies,[17,34]). *Let A be a distributive lattice, and $f : \mathsf{U}A^n \to \mathsf{U}A$ and $g_i : \mathsf{U}A^{m_i} \to \mathsf{U}A, 1 \leq i \leq n$ be arbitrary maps. Assume that there exist topologies τ_i on A, $1 \leq i \leq n$ such that each g_i^σ is (σ^{m_i}, τ_i)-continuous, then*

1. *if f^σ is $(\tau_1 \times \ldots \times \tau_n, \gamma^\uparrow)$-continuous, then $f^\sigma(g_1^\sigma, \ldots, g_n^\sigma) \leq (f(g_1, \ldots, g_n))^\sigma$,*
2. *if f^σ is $(\tau_1 \times \ldots \times \tau_n, \gamma^\downarrow)$-continuous, then $f^\sigma(g_1^\sigma, \ldots, g_n^\sigma) \geq (f(g_1, \ldots, g_n))^\sigma$*
3. *if f^σ is $(\tau_1 \times \ldots \times \tau_n, \gamma)$-continuous, then $f^\sigma(g_1^\sigma, \ldots, g_n^\sigma) = (f(g_1, \ldots, g_n))^\sigma$.*

The last piece of information we need to effectively use the Principle of Matching Topologies is to determine when maps are continuous for a certain topology, based on the distributivity laws they satisfy. For our purpose the following results will be sufficient:

Proposition 4 ([20,17,21,34]). *Let A be a distributive lattice, and let $f : \mathsf{U}A^n \to \mathsf{U}A$ be a map. For every $(n-1)$-tuple $(a_i)_{1 \leq i \leq n-1}$, we denote by $f_a^k : A \to A$ the map defined by $x \mapsto f(a_1, \ldots, a_{k-1}, x, a_k, \ldots, a_{n-1})$.*

1. *If f_a^k preserves binary joins, then $(f^\sigma)_a^k$ preserve all non-empty joins and is $(\sigma^\downarrow, \sigma^\downarrow)$-continuous.*
2. *If f_a^k preserves binary meets, then $(f^\sigma)_a^k$ preserve all non-empty meets and is $(\sigma^\uparrow, \sigma^\uparrow)$-continuous.*
3. *If f_a^k anti-preserves binary joins (i.e., turns them into meets), then $(f^\sigma)_a^k$ anti-preserve all non-empty joins and is $(\sigma^\downarrow, \sigma^\uparrow)$-continuous.*
4. *If f_a^k anti-preserves binary meets (i.e., turns them into joins), then $(f^\sigma)_a^k$ anti-preserve all non-empty meets and is $(\sigma^\uparrow, \sigma^\downarrow)$-continuous.*
5. *In each case f is is smooth in its k^{th} argument.*

3.3 Canonical (in)equations

To say anything about the canonicity of equations, we need to compare interpretations in A with interpretations in A^σ. It is natural to try to use the extension $(\cdot)^\sigma$ to mediate between these interpretations, but $(\cdot)^\sigma$ is defined on maps, not on terms. Moreover, not every valuation on A^σ originates from valuation on 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.

Term functions. The solution is to adopt the language of *term functions* (as first suggested in [24]). Given a signature Σ, let $\mathsf{T}(V)$ denote the language of Σ-DLEs (or Σ-BAEs) over a set V of propositional variables. We view each term $t \in \mathsf{T}(V)$ as defining, for each Σ-DLE A, a map $t^A : A^n \to A$. This allows us to consider its canonical extension $(t^A)^\sigma$, and also allows us to reason without having to worry about specifying valuations. Formally, given a signature Σ and a set V a propositional variables, we inductively define the term function associated with an element t built from variables $x_1, \ldots, x_n \in V$ as follows:

- $x_i^A = \pi_i^n : A^n \to A, 1 \leq i \leq n$;
- $(f(t_1, \ldots, t_m))^A = f^A \circ \langle t_1^A, \ldots, t_m^A \rangle$.

where π_i is the usual projection on the i^{th} component, f^A is the interpretation of the symbol f in A and $\langle t_1^A, \ldots, t_m^A \rangle$ 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 $\vee, \wedge$ and possibly $\neg$.

Proposition 5. *Let s, t be terms in the language defined by a signature Σ and A be a Σ-DLE,*

$$A \models s = t \text{ iff } s^A = t^A .$$

Canonical (in)equations. An equation $s = t$ where $s, t \in \mathsf{T}(V)$ is called *canonical* if $A \models s = t$ implies $A^\sigma \models s = t$, and similarly for inequations. Following [24], we say that $t \in \mathsf{T}(V)$ is *stable* if $(t^A)^\sigma = t^{A^\sigma}$, that t is *expanding* if $(t^A)^\sigma \leq t^{A^\sigma}$, and that t is *contracting* if $(t^A)^\sigma \geq t^{A^\sigma}$, for any A. The inequality between maps is taken pointwise. The following proposition illustrates the usefulness of these notions:

Proposition 6 ([24]). *If $s, t \in \mathsf{T}(V)$ are stable then the equation $s = t$ is canonical. Similarly, let $s, t \in \mathsf{T}(V)$ such that s is contracting and t is expanding, then the inequality $s \leq t$ is canonical.*

4 Completeness via-canonicity

4.1 Axiomatizing HAs and distributive residuated lattices

So far we have only captured part of the structure of Heyting algebras and distributive residuated lattices, namely we have enforced the distribution properties of $\to, *, \backslash$ and $/$ by our definition of the syntax constructors L_{Hey} and L_{RL}. In order to capture the rest of the structures we now add *frame conditions* to the coalgebraic models. To do this we need to find axioms which, when added to HDL1-2 and DL1-6 axiomatize HAs and distributive residuated lattices respectively. Due to the constraints that these axioms must be canonical, we choose the following Heyting Frame Conditions:

HFC1. $a \to a = \top$,
HFC2. $a \wedge (a \to b) = a \wedge b$
HFC3. $(a \to b) \wedge b = b$

and, for distributive lattices, the Frame Conditions:

FC1. $a * I = a$, $I * a = a$,
FC2. $I \leq a \backslash a$, $I \leq a/a$,
FC3. $a * (b \backslash c) \leq (a * b) \backslash c$,
FC4. $(c/b) * a \leq c/(a * b)$,
FC5. $(a/b) * b \leq a$, and
FC6. $b * (b \backslash a) \leq a$,

Proposition 7. *The axioms HDL1-2 and HFC1-3 axiomatize Heyting algebras, and similarly, the axioms DL1-6 and FC1-6 axiomatize distributive residuated lattices.*

We now show one of the crucial steps.

Proposition 8. *The axioms HFC1-3 and FC1-6 are canonical.*

Proof. The proof is an application of Theorem 2 and Proposition 6.

FC1: Since $*$ preserves binary joins in each argument, it is smooth by Prop. 4, and it follows that it is (σ^2, γ)-continuous. Since π_1^σ and I^σ are trivially (σ, σ)-continuous, it follows from Theorem 2 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 Prop. 6.

FC2: I is stable and thus contracting, and $(\backslash \circ \langle \pi_1, \pi_1 \rangle)^\sigma = \backslash^\sigma \circ \langle \pi_1, \pi_1 \rangle^\sigma$, since π_1^σ is $(\sigma, \sigma)-$continuous and $\backslash^\sigma$ 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 $\backslash^\sigma$ is smooth it is in particular $(\sigma^2, \gamma^\uparrow)$-continuous. Since π_1^σ is $(\sigma, \gamma^\uparrow)$-continuous, we get that $*^\sigma \circ \langle \pi_1^\sigma, \backslash^\sigma \circ \langle \pi_2^\sigma, \pi_3^\sigma \rangle \rangle$ is $(\sigma^3, \gamma^\uparrow)$-continuous and thus contracting. For the RHS, note that since $\backslash^\sigma$ preserves meets in its first argument, it must in particular preserve down-directed ones, thus $\backslash^\sigma$ is $(\gamma^\downarrow, \gamma^\downarrow)$-continuous in its first argument. Similarly, since $\backslash^\sigma$ 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 $\backslash^\sigma$ 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.

4.2 Strong completeness results

The Jónsson-Tarski theorem. We first establish the strong completeness of the logics defined by our syntax constructors L_{Hey} and L_{RL} with respect to their T_{Hey}- and T_{RL}-coalgebraic models. The proof is an application of the coalgebraic Jónsson-Tarksi theorem. Recall from Theorem 1 and Diagram (1), that the canonical extension of an object A in any of our reasoning kernels $\mathscr{C}$ is given by GFA. This justifies the following:

Theorem 3 (Coalgebraic Jónsson-Tarksi theorem, [28]). *Assuming the basic situation of Diagram (1) and a semantic transformation $\delta : LG \to GT$, if its adjoint transpose $\hat{\delta} : TF \to FL$ has a right-inverse $\hat{\delta}^{-1} : FL \to TF$, then for every L-algebra $\alpha : LA \to A$, the embedding $\eta_A : A \to GFA$ of A into its canonical extension can be lifted to the following L-algebra embedding:*

$$\begin{tikzcd}
LA \arrow[rrr, "\alpha"] \arrow[d, "L\eta_A"'] & & & A \arrow[d, "\eta_A"] \\
LGFA \arrow[r, "\delta_{FA}"'] & GTFA \arrow[r, "G\hat{\delta}_A^{-1}"'] & GFLA \arrow[r, "GF\alpha"'] & GFA
\end{tikzcd} \tag{2}$$

We call the coalgebra $\hat{\delta}^{-1} \circ F\alpha : FA \to TFA$ a *canonical model* of (the L-algebra) A. If A is the free L-algebra over $\mathsf{F}V$ we recover the usual notion of canonical model. The 'truth lemma' follows from the definition of η.

We now prove the existence of canonical models for the logics defined by L_{Hey} and L_{RL}. The result generalizes lemma 5.1 of [14], which builds canonical models for countable DLs with a unary operator, and lemma 4.26 of [3], which builds canonical models for countable BAs with n-ary operators. We essentially show how to build canonical models for arbitrary DLs with n-ary expansions all of whose arguments either (1) preserve joins or anti-preserve meets, or (2) preserve meets or anti-preserve joins.

Theorem 4. *The logic defined by L_{Hey} (resp. L_{RL}) is sound and strongly complete with respect to the class of all T_{Hey}- (resp. T_{RL}-) coalgebras.*

Proof (Sketch). The proof follows a Prime Ideal Theorem argument. To interpret $*$ on $\mathsf{Pf}A$ for some A in **DL** we define $\gamma^*_A : \mathsf{Pf}A \to \mathsf{P}_c(\mathsf{Pf}A \times \mathsf{Pf}A), F \mapsto \{(F_1, F_2) \mid a \in F_1, b \in F_2 \Rightarrow a * b \in F\}$. It is easy to check that if $\exists F_1, F_2$ s.th. $(F_1, F_2) \in \gamma_A(F)$ and $a \in F_1, b \in F_2$, then $a * b \in F$ and $F \models a * b$. The converse is harder: given $a * b \in F$, we must build prime filters F_1, F_2 s.th. $a \in F_1, b \in F_2$ and $c * d \notin F \Rightarrow c \notin F_1$ or $d \notin F_2$. We consider the set $\mathscr{P}(a, b)$ of pairs of filter-ideal pairs $((F_1, I_1), (F_2, I_2))$ s.th.

1. $\uparrow a \subseteq F_1 \subseteq \{c \mid \forall d \in F_2, c * d \in F\}$
2. $\uparrow b \subseteq F_2 \subseteq \{d \mid \forall c \in F_1, c * d \in F\}$
3. $I_1 = \{c \mid \exists d \in F_2 \text{ s.th. } c * d \notin F\}$
4. $I_2 = \{d \mid \exists c \in F_1 \text{ s.th. } c * d \notin F\}$

It can be shown that $\mathscr{P}(a, b)$ is not-empty, forms a poset, has the property that I_1, I_2 are ideals such that $F_1 \cap I_1 = F_2 \cap I_2 = \emptyset$, and is closed under union of chains. Zorn's lemma then yields a maximum element which provides the desired prime filters. The same technique can be applied to define $\gamma^{\backslash}_A, \gamma^{/}_A$ interpreting $\backslash, /$, and it is easy to check that $\langle 0, \gamma^*_A, \gamma^{\backslash}_A, \gamma^{/}_A \rangle$ is a right inverse of $\hat{\delta}^{\mathrm{DL}}_A$.

The Jónsson-Tarski embedding and canonical extensions. We now apply the theory of canonicity to show that HAs and distributive residuated lattices are strongly complete with respect to the (proper) classes of T_{Hey}- and T_{RL}-coalgebras validating HFC1-3 and FC1-6 respectively. We need one important technical result, which shows that the Jónsson-Tarski embedding of Theorem 3 is the canonical extension defined in Section 3.2.

Proposition 9. *The structure map of the Jónsson-Tarski extension of an L_{Hey}- or L_{RL}-algebra is equal to the canonical extension of its structure map (in the sense of Section 3.2).*

Proof (Sketch). Recall Diagram (2) and that a DL-morphism $\alpha : L_{\mathrm{RL}}A \to A$ is equivalent to being given a constant and binary operations $\alpha_*, \alpha_{\backslash}, \alpha_{/}$ on $\mathsf{U}A$ satisfying DL1-DL6. Similarly, $\mathcal{U}\mathsf{Pf}\alpha \circ \mathcal{U}\gamma_A \circ \delta^{\mathrm{RL}}_{\mathsf{Pf}A}$ is equivalent to a constant and three binary operations $\mathcal{U}\mathsf{Pf}\alpha \circ \mathcal{U}\gamma_A \circ \delta_*, \mathcal{U}\mathsf{Pf}\alpha \circ \mathcal{U}\gamma_A \circ \delta_{\backslash}, \mathcal{U}\mathsf{Pf}\alpha \circ \mathcal{U}\gamma_A \circ \delta_{/}$ on A^σ. By commutativity of (2), the latter are extensions of the former. It is not hard to show that if an extension of a map on $\mathsf{U}A$ preserves or anti-preserves all non-empty meets or joins, then it is smooth and thus unique by Proposition 6. Direct calculation shows that $\delta_*, \delta_\backslash$ and $\delta_/$ all have such preservation properties in each argument. Moreover, $\mathcal{U}\mathsf{Pf}\alpha$ and $\mathcal{U}\gamma_A$ being inverse images preserve any meet or join. We thus get that $\mathcal{U}\mathsf{Pf}\alpha \circ \mathcal{U}\gamma_A \circ \delta_*$ is smooth and thus equal to α^σ as desired, and similarly for the other operations.

Strong completeness. We are now ready to state our main result.

Theorem 5 (Strong completeness theorem). *Intuitionistic logic is strongly complete with respect to the class of T_{Hey}-coalgebras validating HFC1-3. The Distributive Full Lambek Calculus is strongly complete with respect to the class of T_{RL}-coalgebras validating FC1-6.*

Proof (Sketch). We treat the case of the distributive full Lambek calculus; intuitionistic logic is treated similarly. Let Φ, Ψ be sets of L_{RL}-formulas such that FC1-6+$\Phi \nvdash \Psi$. We need to find a model in which FC1-6 are valid, and which satisfies all formulas of Φ and no formula of Ψ at a certain point. Consider the Lindenbaum-Tarski L_{RL}-algebra $\mathcal{L}$ defined by FC1-6. These axioms are clearly valid in $\mathcal{L}$, and since they are canonical by Prop. 8, they are also valid in $\mathcal{L}^\sigma$, which by Prop. 9 is just its coalgebraic Jónsson-Tarski extension. It follows that FC1-6 are valid on the *model* $\mathsf{Pf}\mathcal{L} \to T_{\mathrm{RL}}\mathsf{Pf}\mathcal{L}$. To find the desired point, note that the filter generated by Φ in $\mathcal{L}$ is proper and does not intersect the ideal $\langle\Psi\rangle$ generated by Ψ, or else our staring assumption would be contradicted. We can thus find $\mathsf{Pf}\mathcal{L} \ni p_\Phi \supseteq \Phi$ s.th. $p_\Phi \cap \langle\Psi\rangle = \emptyset$, and $p_\Phi \models \Phi$, $p_\Phi \not\models \Psi$ follows.

Describing T_{Hey}-coalgebras validating the Heyting Frame Conditions. Let us examine what T_{Hey}-coalgebras validating HFC1-3 look like. For every $\gamma : W \to T_{\mathrm{Hey}}W$ in this class, every $w \in W$ and every valuation, $w \models a \to a$. By considering a formula satisfied at a single point in the model is easy to see that $(x, y) \in \gamma(w) \Rightarrow x = y$; that is, the structure map of the coalgebra only really defines a binary relation to interpret $\to$. Thus T_{Hey}-coalgebras validating HFC1 are equivalent to P_c-coalgebras where $w \models a \to b$ iff $\forall x \in \gamma(w), x \models a \Rightarrow x \models b$. The distributivity laws of $\to$ together with HFC2-3 encode the well-known residuation property of $\to$ 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. It also follows that $x \in \gamma(x)$; that is, the relation is reflexive. Finally, from HFC3 we get that $a \wedge b \leq c$ iff $b \leq a \to c$ iff $b \leq a \to (c \wedge (a \to c))$. By unravelling the interpretation of this last inequality, we get that the relation interpreting $\to$ 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.

Describing 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, [6]. 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 $T'_{\mathrm{RL}} : \mathscr{D} \to \mathscr{D}, T'_{\mathrm{RL}} W = 2 \times \mathsf{P}_c(W \times W)$ 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 $x \models a$.

Modularity. The coalgebraic setting allows us to combine completeness-via-cano-nicity results from simple logics to get results for more complicated logics. It can be shown that the coalgebraic Jónsson-Tarski theorem is modular in the sense that if logics defined by syntax constructors L_1 and L_2 and interpreted in T_1- and T_2-coalgebras respectively via semantic transformations δ_1 and δ_2 whose adjoint transposes have right-inverses, then the logic defined by $(L_1 + L_2)$ is strongly complete w.r.t. $(T_1 \times T_2)$-coalgebras.

Theorem 6 (Strong completeness of intuitionistic BI). *Intuitionistic BI is strongly complete w.r.t. the class of $T_{\mathrm{Hey}} \times T_{\mathrm{RL}}$-coalgebra satisfying HFC1-3 and FC1-6.*

Additional frame conditions. We can consider more axioms to restrict further the classes of models we might be interested in. The following (in)equations can all easily be verified to be canonical and each corresponds to admitting a structural rule to the full distributive Lambek calculus: (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). More generally, we have presented a general methodology to get completeness results for axioms that could capture the behaviour of certain sub-classes of resources (e.g., heaps in separation logic).

5 Conclusion and future work

We have shown how distributive substructural logics can be formalized and given a semantics in the framework of coalgebraic logic, and highlighted the modularity of this approach. By choosing a syntax whose operators explicitly follow distribution rules, we can use the elegant topological theory of canonicity for DLs, and in particular the notion of smoothness and of topology matching, to build a set of canonical (in)equation capturing intuitionistic logic and the distributive full Lambek calculus. The coalgebraic approach makes the connection between algebraic canonicity and canonical models explicit, categorical and generalizable.

The modularity provided by our approach is twofold. Firstly, we have a generic method for building canonical (in)equations by using the Principle of Matching Topologies. Getting completeness results with respect to simple Kripke models for variations of the distributive full Lambek calculus (e.g., distributive affine logic) becomes very straightforward. Secondly, adding more operators to the fundamental language simply amounts to taking a *coproduct* of syntax constructors (e.g., $L_{\mathrm{RL}} + L_{\mathrm{Hey}}$ to define intuitionistic BI) and interpreting it with a *product* of model constructors (e.g., $T_{\mathrm{RL}} \times T_{\mathrm{Hey}}$). This seems particularly suited to logics which build on BI such as the bi-intuitionistic boolean BI of [4].

The operators $*, \backslash, /$, and $\to$ all satisfy simple distribution laws, but our approach can also accommodate operators with more complicated distribution laws and non-relational semantics. For example, the theory presented in this work could be extended to cover a graded version of $*$, say $*_k$, whose interpretation would be 'there are at least k ways to separate a resource such that...', the semantics would be given by coalgebras of the type $\mathbb{2} \times \mathcal{B}(- \times -)$ where $\mathcal{B}$ is the 'bag' or multiset functor. Similarly, a graded version $\to_k$ of the intuitionistic implication whose meaning would be '... implies ... apart from at most k exceptions' and interpreted by $\mathcal{B}(- \times -)$-coalgebras could also be covered by our approach. Crucially, such operators do satisfy (more complicated) distribution laws which lead to generalizations of the results in Section 3.2, and the possibility of building canonical (in)equations. The coalgebraic infrastructure then allows the rest of the theory to stay essentially unchanged.

References

1. N. Bezhanishvili, M. Gehrke. Free Heyting Algebras: Revisited. LNCS 5728, 251–266, 2009.
2. M. Bilkova, A. Kurz, D. Petrisan, J. Velebil Relation Liftings on Preorders and Posets. LNCS 6859, 115–129, 2011.
3. P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic.* Cambridge Tracts in Theoretical Computer Science 53. CUP, 2001.
4. J. Brotherston and J. Villard. Bi-intuitionistic boolean bunched logic. Research Note RN/14/06, University College London, 2014.
5. D. Coumans, M. Gehrke, and L. van Rooijen. Relational semantics for a fragment of linear logic. *Proceedings of PhDs in Logic III*, 2011.
6. C. Calcagno, P. Gardner, and U. Zarfaty. Context logic as modal logic: completeness and parametric inexpressivity. In *ACM SIGPLAN Not.* 42.1, 123–134, 2007.
7. M. Collinson and D. Pym. Algebra and logic for resource-based systems modelling. *Mathematical Structures in Computer Science*, 19(5):959–1027, 2009.
8. M. Collinson, B. Monahan, and D. Pym. *A Discipline of Mathematical Systems Modelling.* College Publications, 2012.
9. C. Cîrstea and D. Pattinson. Modular construction of complete coalgebraic logics. *Theoret. Comp. Sci.* 388(1-3), 83–108, 2007.
10. F. Dahlqvist and D. Pattinson. On the fusion of coalgebraic logics. LNCS 6859, 161–175, 2011.

11. F. Dahlqvist and D. Pattinson. Some sahlqvist completeness results for coalgebraic logics. LNCS 7794, 193–208, 2013.
12. F. Dahlqvist and D. Pym. Completeness via canonicity for distributive substructural logics: a coalgebraic perspective. Research Note RN/15/04, Department of Computer Science, UCL, 2015. `http://www.cs.ucl.ac.uk/fileadmin/UCL-CS/research/Research_Notes/rn-15-04.pdf` .
13. J. M. Dunn, M. Gehrke, and A. Palmigiano. Canonical Extensions and Relational Completeness of Some Substructural Logics. *J. Symb. Logic*, 70(3),713–740, 2005.
14. J. M. Dunn. Positive modal logic. *Studia Logica*, 55(2):301–317, 1995.
15. D. Galmiche, D. Méry, and D. Pym. The Semantics of BI and Resource Tableaux. *Mathematical Structures in Computer Science* (2005) 15, 1033–1088.
16. M. Gehrke. Generalized Kripke frames. *Studia Logica*, 84(2):241–275, 2006.
17. M. Gehrke and J. Harding. Bounded lattice expansions. *Journal of Algebra*, 238(1):345 – 371, 2001.
18. J.-Y. Girard. Linear logic. *Theoret. Comp. Sci.*, 50:1–102, 1987.
19. S. Givant and P. Halmos. *Introduction to Boolean Algebras.* Springer, 2009.
20. M. Gehrke and B. Jónsson. Bounded distributive lattices with operators. *Mathematica Japonica*, 40(2):207–215, 1994.
21. M. Gehrke and B. Jónsson. Bounded distributive lattice expansions. *Mathematica Scandinavica*, 94:13–45, 2004.
22. M. Gehrke, R. Jansana, and A. Palmigiano. Δ_1 -completions of a poset. *Order*, 30(1):39–64, 2013.
23. S. S. Ishtiaq and P. W. O'Hearn. BI as an assertion language for mutable data structures. In *ACM SIGPLAN Not.* 36.3, 14–26, 2001.
24. B. Jónsson. On the canonicity of Sahlqvist identities. *Stud. Log.*, 53(4):473–492,1994.
25. B. Jacobs and A. Sokolova. Exemplaric expressivity of modal logics. *J. Log. Compuat.*, 20:1041–1068, 2010.
26. B. Jónsson and A. Tarski. Boolean algebras with operators. part 1. *Amer. J. Math.*, 33:891–937, 1951.
27. C. Kupke, A. Kurz, and D. Pattinson. Algebraic semantics for coalgebraic logics. ENTCS 106, 219–241, 2004.
28. C. Kupke, A. Kurz, and D. Pattinson. Ultrafilter Extensions for Coalgebras. LNCS 3629, 263–277, 2005.
29. H. Ono. Substructural logics and residuated lattices — an introduction. *Trends in Logic*, 20:177–212, 2003.
30. P. W. O'Hearn and D. J. Pym. The logic of bunched implications. *Bulletin of Symbolic Logic*, 5(2):215–244, 1999.
31. D. Pym, P. O'Hearn, and H. Yang. Possible Worlds and Resources: The Semantics of BI. *Theoret. Comp. Sci.* 315(1): 257–305, 2002. Erratum: p. 285, l. -12: ', for some $P', Q = P; P'$' should be '$P \vdash Q$'.
32. J. C Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. 17th LICS*, 55–74. IEEE, 2002.
33. T. Suzuki. Canonicity results of substructural and lattice-based logics. *The Review of Symbolic Logic*, 4:1–42, 3 2011.
34. Y. Venema. Algebras and coalgebras. In J. van Benthem, P. Blackburn, and F. Wolter, editors, *Handbook of Modal Logic.* Elsevier, 2006.