

The points of canonical extensions of doctrines

Sam van Gool¹ and Joshua L. Wrigley²

¹LMF, ENS Paris-Saclay

²Department of Mathematics and Statistics, Faculty of Sciences, Masaryk University

June 27, 2026

Abstract

We analyse the space of points of the canonical extension of a coherent doctrine. We first give a full characterisation of doctrine morphisms that are extensible, and relate it to the existing notion of p -model of a coherent category. Through this characterisation, the extensible morphisms are shown to be exactly those which are ω -saturated in the sense of coherent first-order logic. Next, we answer the question: when does a presheaf of models fully describe the canonical extension? We prove a characterisation theorem via two conditions, which are again natural from the perspective of coherent logic, namely, homogeneity and the realisation of all prime types in a strict sense. The characterisation theorem allows us to deduce a reconstruction result for any coherent theory with the property that all prime types can be realised in a countable, saturated model. For instance, ω -stable coherent theories always have this property. We conclude by explaining how our results can be interpreted topos-theoretically, by relating them to the classifying topos and to the topos of types.

Contents

1	Introduction	2
2	Canonical extensions of coherent doctrines	4
2.1	Coherent doctrines	4
2.2	Canonical extensions	5
2.3	Canonical extensions of coherent doctrines	8
3	Extensible morphisms of coherent doctrines	10
4	Points	13
5	Presheaves of points	16
5.1	Presheaf-valued points and presheaves of points	17
5.2	Subpresheaves and orbits of elements	18
5.3	Types in a doctrine	19
5.4	Plentiful categories of models	20
6	Recovering the canonical extension for ω-stable theories	23
7	A tale of two topoi	26

1 Introduction

The duality theory of Stone [37] shows that a distributive lattice may be equivalently studied through its space of *points*, which can be viewed as its prime filters, its prime ideals, or its morphisms to the two-element lattice. A logical interpretation of Stone duality views a lattice as an algebra of equivalence classes of formulae modulo a propositional theory, and the points of its dual space as model-theoretic types of the theory. In the absence of classical negation, types may have non-trivial order relations between them. It is a key insight of Priestley’s seminal 1970 paper [34] that duality theory for distributive lattices is greatly simplified and enhanced if one considers this partial order on points as a primitive, rather than a derived, notion in the space. Under Priestley duality, lattice morphisms correspond to continuous order-preserving maps between their dual spaces in the opposite direction.

Gehrke and Jónsson [10], building on work of Jónsson and Tarski for Boolean algebras [22], introduce an algebraic counterpart to Priestley’s dual space: the *canonical extension* L^δ of a distributive lattice L . While, importantly, this extension can be constructed in a point-free manner without appeal to any choice principles (cf. Subsection 2.2), a useful spatial point of view on the elements of L^δ is that they are upwards closed subsets of points of L , with respect to Priestley’s partial order. In other words, the completion $L \hookrightarrow L^\delta$ is an algebraic rendering of the idea of equipping the topological space of points of L with a partial order.

In this way, Priestley duality can, *a posteriori*, be *derived* from the canonical extension construction: Given the canonical extension $L \hookrightarrow L^\delta$ of a distributive lattice L , the set $\text{Pt}(L)$ of points of L is in a natural bijection with the set of *complete* morphisms from L^δ to the two-element lattice 2 , or, equivalently, with the set $J^\infty L^\delta$ of *completely join-prime* elements of L^δ . The Priestley order on $\text{Pt}(L)$ can then be obtained from the restriction of the order of L^δ to $J^\infty L^\delta$, and the Priestley topology, viewed on $J^\infty L^\delta$, is generated by the sets $\downarrow a \cap J^\infty L^\delta$ and their complements, as a ranges over the elements of L .

Syntactically speaking, the canonical extension L^δ may be thought of as an enrichment of the propositional theory that L represents, by adding infinitary conjunction and disjunction, while *only taking into account the finitary part of L* . This means, for instance, that even if a sequence of elements $(\varphi_n)_{n \in \mathbb{N}}$ of L has an infimum φ , say, in L , then it will have a new infimum φ' in L^δ , where $\varphi' \notin L$. This new element φ' will be the unique successor of φ in L^δ . In the definition of canonical extensions, this is formalised as a *compactness* property of the completion. This strict separation of the finitary and the infinitary distinguishes the canonical extension from, for instance, the Dedekind-McNeille completion, and is crucial for the ensuing theory.

In her PhD work [4, 5] from 2013, Coumans extends the canonical extension construction to *coherent doctrines*, the algebraic structures encoding first-order theories, still in a negation-free setting. Specifically, given a first-order theory \mathbb{T} , one may associate with it a collection $D_{\mathbb{T}}$ of distributive lattices, indexed over finite sets of variables called *contexts*, as follows. For each context c , write $D_{\mathbb{T}c}$ for the lattice of \mathbb{T} -equivalence classes of formulae of which all free variables are in c . Substitutions between contexts then induce lattice morphisms between the fibers. This leads to Lawvere’s formulation [25] of *doctrines* as functors from a category of contexts to the category of distributive lattices, where the adjective *coherent* signifies that these functors are subject to axioms encoding the rules of (negation-free) first-order logic (called coherent logic by topos theorists; see [20, §D1]); we recall the precise definitions in Subsection 2.1 and discuss syntactic doctrines further in Section 4.

The *canonical extension* D^δ of a coherent doctrine D is obtained by post-composing D with the functor that assigns to a lattice L its canonical extension L^δ . The first key result of Coumans, recalled in Subsection 2.3, is that the resulting functor is again a

coherent doctrine, i.e., D^δ still satisfies the axioms encoding first-order logic. A further deep insight of Coumans’s work [5] is that the canonical extension construction, at the first-order level, is closely related to Makkai’s *topos of types* construction [28], which, as its name suggests, intends to capture the ‘theory of all model-theoretic types’ of a given first-order theory. This leads us to the question motivating this paper:

How do the type spaces of a first-order theory relate to the canonical extension of its doctrine?

A first answer to this question is essentially immediate from the definitions, generalising the Gehrke-Jónsson viewpoint on Priestley duality outlined above: the types of the theory are still the complete points, in the appropriate sense, of the canonical extension of its doctrine, and it is possible to assemble these into Priestley spaces of types called *polyadic Priestley spaces*, see for instance [15]. However, as already noticed by both Makkai [28] and Coumans [5], the appropriate definition of ‘complete point’ for doctrines is more subtle than it is for lattices, and our first contribution here clarifies this subtlety, as we will briefly describe now.

A *point* of a doctrine D is usually conceptualized as a morphism from D to the semantic *doctrine of predicates*, \mathcal{P} , i.e., the contravariant powerset functor. The key problem is that, contrary to the lattice setting, a doctrine D may have points which cannot be extended to complete points of D^δ ; we give a concrete example in Example 2.21 below. In Section 3, we then study the ensuing extensibility problem at the level of general morphisms between doctrines, and we arrive at a characterisation of extensible morphisms in Theorem 3.1. Our condition characterising extensible morphisms is similar to one previously identified in Makkai’s work [28] under the name ‘ p -model’, and this notion also plays a central role in Coumans’s work; for more on the comparison, see Remarks 3.5 and 4.11, and Section 7 below. Our new contribution in Section 3 is to derive this concept from first principles.

In Section 4, we instantiate our general concept of extensibility to points of a doctrine of the form D_{\top} , for \top a coherent theory. This allows us to show, in Theorem 4.8, that the extensible points of a syntactic doctrine precisely correspond to ω -saturated models, a correspondence that had only been implicit in the previous literature (cf. Remark 4.11). Here, as everywhere in the paper, the concept of ω -saturation is meant in a negation-free sense that takes into account the partial order between prime types. We discuss the relationship with classical ω -saturation in Remark 4.11.

In Section 5, we move from studying individual points of a doctrine to *families* of points. Such families will be indexed by a small category \mathcal{K} , and can be thought of either as presheaves on \mathcal{K} valued in (usual) points of D or as points of D valued in a presheaf category $\mathbf{Set}^{\mathcal{K}}$; this correspondence is made explicit in Subsection 5.1. The presheaf category $\mathbf{Set}^{\mathcal{K}}$ induces a semantic doctrine of *subobjects*, generalising the doctrine of predicates \mathcal{P} mentioned above.

In Proposition 5.6, we identify the completely join-prime elements of the subobject doctrine of $\mathbf{Set}^{\mathcal{K}}$ as the *orbits* in a presheaf. Using this characterisation of the completely join-prime elements, we obtain the main result of Section 5, Theorem 5.15, in which we give two model-theoretic conditions characterising exactly when a presheaf-valued point M ‘captures’ the canonical extension of a doctrine, in the sense that the family induces an isomorphism between D^δ and the doctrine of subobjects of $\mathbf{Set}^{\mathcal{K}}$, along a restricted domain. Thus, we recover another important result of Coumans [5, Theorem 40], as we explain in Section 7. The two model-theoretic conditions on M , which we call *homogeneity* and *strictly realising all (prime) types*, characterise the surjectivity and injectivity of the morphism, respectively. This result demonstrates how the canonical extension of doctrines allows for a lattice-theoretic approach to the model theory of first-order logic, in a way that is analogous to the use of canonical extensions of lattices for the model theory of propositional logics, see, e.g., [12].

In Section 6, we use Theorem 5.15 to derive an entirely new result: we present a ‘reconstruction theorem’ for the canonical extension of a coherent doctrine of the form $D_{\mathbb{T}}$, Theorem 6.6, which applies whenever every prime type of \mathbb{T} can be realised in a countable, saturated model. This theorem in particular applies to ω -stable coherent theories, as explained in Corollary 6.9. Given some algebraic structure A associated with a theory \mathbb{T} , a ‘reconstruction problem’ asks what further data is needed on the points of \mathbb{T} in order to be able to reconstruct A . For instance, Priestley duality can be understood as a reconstruction theorem for coherent propositional logic: the Lindenbaum-Tarski algebra of formulae for \mathbb{T} can be obtained from the points of \mathbb{T} , equipped with a topology. Our reconstruction results, Theorem 6.6 and Corollary 6.9, assert that, in the particular setting described above, the canonical extension of the syntactic doctrine of \mathbb{T} can be reconstructed from a certain subcategory of the models of \mathbb{T} endowed with ‘knowledge of the underlying sets of the models’.

In the concluding Section 7, we relate our results to the topos-theoretic literature.

2 Canonical extensions of coherent doctrines

In this section, we present a theory of canonical extension for coherent doctrines. This theory was first developed in [5, 4], combining doctrine theory [25, 29] and canonical extension theory [22, 23, 10, 11]. We will here recall the definitions necessary for understanding this theory, and prove some preliminary results about them.

2.1 Coherent doctrines

Distributive lattices¹ provide the algebraic semantics of *coherent propositional logic*. For the algebraic semantics of *coherent first-order logic*, we replace distributive lattices by *coherent doctrines* [25]. The concept of doctrine categorifies the operation which, to any context of variables, associates the distributive lattice of equivalences classes of first-order formulae whose free variables are in the context. First-order connectives, such as existential quantification and equality, are provided by structure between these lattices, namely the existence of left adjoints.

Definition 2.1. A *coherent doctrine*² is a functor $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ where \mathcal{C} , the *base category*, has all finite limits, and such that, for each arrow $f: c \rightarrow c'$ in \mathcal{C} , the lattice morphism $Pf: Pc' \rightarrow Pc$ has a left adjoint, denoted $\exists_f^P: Pc \rightarrow Pc'$, satisfying:

- (i) (Frobenius reciprocity) for all $\varphi \in Pc$ and $\psi \in Pc'$,

$$\exists_f^P(\varphi \wedge (Pf)(\psi)) = \exists_f^P(\varphi) \wedge \psi, \text{ and}$$

- (ii) (Beck-Chevalley condition) for any pullback square in \mathcal{C} as given on the left, the square in \mathbf{DL} on the right commutes:

$$\begin{array}{ccc} c_1 \times_c c_2 & \xrightarrow{\pi_2} & c_2 \\ \pi_1 \downarrow & & \downarrow f_2 \\ c_1 & \xrightarrow{f_1} & c \end{array} \quad \begin{array}{ccc} Pc_2 & \xrightarrow{P\pi_2} & Pc_1 \times_c c_2 \\ \exists_{f_2}^P \downarrow & & \downarrow \exists_{\pi_1}^P \\ Pc & \xrightarrow{Pf_1} & Pc_1. \end{array}$$

Let $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ and $Q: \mathcal{D}^{\text{op}} \rightarrow \mathbf{DL}$ be coherent doctrines. A *morphism* from P to Q is a pair (M, α) , where $M: \mathcal{C} \rightarrow \mathcal{D}$ is a finite-limit-preserving functor, and

¹Throughout this paper, ‘lattice’ means ‘bounded lattice’, i.e., a partially ordered set admitting finite suprema and finite infima, and lattice morphisms are required to preserve finite suprema and finite infima, unless explicitly mentioned otherwise.

²These are also known as ‘coherent hyperdoctrines’ in the literature, but we drop the prefix ‘hyper’ for brevity’s sake.

$\alpha: P \Rightarrow Q \circ M^{\text{op}}$ is a natural transformation such that, for any arrow $f: c \rightarrow c'$ in \mathcal{C} , the following *Beck-Chevalley square* commutes:

$$\begin{array}{ccc} Pc & \xrightarrow{\exists_f^P} & Pc' \\ \alpha_c \downarrow & & \downarrow \alpha_{c'} \\ QMc & \xrightarrow{\exists_{Mf}^Q} & QMc'. \end{array} \quad (1)$$

Given morphisms $(M, \alpha), (N, \beta): P \rightarrow Q$, a *transformation* from (M, α) to (N, β) is a natural transformation $\sigma: M \Rightarrow N$ such that $\alpha_c(\varphi) \leq Q(\sigma_c)\beta_c(\varphi)$ for all $c \in \mathcal{C}$ and $\varphi \in Pc$. We denote the 2-category of coherent doctrines by **CohDoc**.

Notation 2.2. When it can be inferred from the context what is the coherent doctrine $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ under consideration, and f is a morphism of \mathcal{C} , we write f^* for the morphism Pf and \exists_f for its left adjoint \exists_f^P .

Remark 2.3. If $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ is a coherent doctrine, and $f: c \rightarrow c'$ is a monomorphism in \mathcal{C} , then the square

$$\begin{array}{ccc} c & \xrightarrow{\text{id}_c} & c \\ \text{id}_c \downarrow & & \downarrow f \\ c & \xrightarrow{f} & c' \end{array}$$

is a pullback in \mathcal{C} . Thus, by the Beck-Chevalley condition, $Pf \circ \exists_f^P = \exists_{\text{id}_c}^P \circ P\text{id}_c = \text{id}_{Pc}$.

Example 2.4. Let \mathcal{C} be a category with finite limits which is *well-powered*, i.e., for each object c of \mathcal{C} , the collection $\text{Sub}_{\mathcal{C}}(c)$ of subobjects of c is a set. We denote by $\text{Sub}_{\mathcal{C}}: \mathcal{C}^{\text{op}} \rightarrow \mathbf{Poset}$ the functor that sends each object to its meet-semilattice of subobjects, and acts by pullback on morphisms. The functor $\text{Sub}_{\mathcal{C}}$ is a coherent doctrine precisely when \mathcal{C} is a *coherent category* (see, e.g., [20, §A.1]). When $\mathcal{C} = \mathbf{Set}$, the coherent doctrine $\text{Sub}_{\mathbf{Set}}$ is (isomorphic to) the power set doctrine \mathcal{P} , see Definition 2.19 below. In Section 5, we will consider $\text{Sub}_{\mathcal{C}}$ for \mathcal{C} a category of presheaves.

Remark 2.5. Let $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ and $Q: \mathcal{D}^{\text{op}} \rightarrow \mathbf{DL}$ be coherent doctrines. If $M: \mathcal{C} \rightarrow \mathcal{D}$ is a finite-limit-preserving functor and $\alpha: P \Rightarrow Q \circ M^{\text{op}}$ is a natural *isomorphism*, then (M, α) is always a morphism of coherent doctrines $(M, \alpha): P \rightarrow Q$. Indeed, we only need to check the commutativity of (1). The inverse α_c^{-1} defines a left adjoint to the component α_c , and so by taking left adjoints of the commuting naturality square

$$\begin{array}{ccc} Pc' & \xrightarrow{Pf} & Pc \\ \alpha_{c'} \downarrow & & \downarrow \alpha_c \\ QMc' & \xrightarrow{QMf} & QMc, \end{array}$$

we obtain the equation $\exists_f^P \circ \alpha_c^{-1} = \alpha_{c'}^{-1} \circ \exists_{Mf}^Q$, from which we derive $\alpha_{c'} \circ \exists_f^P = \exists_{Mf}^Q \circ \alpha_c$ as desired.

2.2 Canonical extensions

The canonical extension of a lattice [21, 10, 9, 11, 7] gives an algebraic, point-free presentation of Stone duality and its extensions, without any choice assumptions. We briefly recall here the definition and essential facts about the canonical extension.

Let L be a lattice. A *completion* of L is an injective lattice morphism $\eta: L \hookrightarrow C$ where C is a complete lattice. We will always identify L with its isomorphic image $\eta[L]$ in C , so that we may consider any completion as a sublattice inclusion $\eta[L] \hookrightarrow C$. In this setting, an element $x \in C$ is called a *filter element* when $x = \bigwedge \{a \in L \mid x \leq a\}$, and $y \in C$ is called an *ideal element* when $y = \bigvee \{b \in L \mid b \leq y\}$. The sets of filter and ideal elements of C are denoted $\mathcal{F}(C)$ and $\mathcal{I}(C)$, respectively.

Definition 2.6. A completion $L \hookrightarrow L^\delta$ is a *canonical extension* of L if it satisfies the following two properties:

- (i) (density) the filter elements join-generate and the ideal elements meet-generate, i.e., for every $u \in L^\delta$,

$$\bigvee\{x \in \mathcal{F}(L^\delta) \mid x \leq u\} = u = \bigwedge\{x \in \mathcal{I}(L^\delta) \mid u \leq y\},$$

- (ii) (compactness) for any subsets S, T of L , if $\bigwedge S \leq \bigvee T$, then there exist finite subsets S' of S and T' of T such that $\bigwedge S' \leq \bigvee T'$.

Every lattice L has a canonical extension, which is unique up to an isomorphism that preserves L . This can be proved without appeal to any choice principle [9, Section 2]. One such construction exhibits L^δ as the MacNeille completion of an *intermediate poset*, built from the free directed join completion and the free directed meet completion of L ; the original source for this abstract viewpoint on the construction is the paper [13] by Gehrke and Priestley. In more concrete terms, as was done in [9], L^δ can be built as the complete lattice of stable elements of a Galois connection induced by a relation between the filter lattice of L and the ideal lattice of L .

Let L be a distributive lattice. The canonical extension L^δ of L is a complete distributive lattice which, assuming a choice principle, can moreover be shown to be isomorphic to the lattice of up-sets of a poset.³ This (up to isomorphism unique) poset representing L^δ can be obtained by equipping the set $J^\infty L^\delta$ of *completely join-prime* elements of L^δ with the opposite of the order of L^δ , restricted to $J^\infty L^\delta$. Here, an element x in a complete lattice C is *completely join-prime* if, for any $S \subseteq C$, $x \leq \bigvee S$ implies that $x \leq s$ for some $s \in S$. In particular, assuming a choice principle, the completely join-prime elements of L^δ join-generate L^δ . For reasons to be explained further in Section 4, we will sometimes refer to completely join-prime elements of L^δ as *complete points*.

The poset of completely join-prime elements of L^δ is isomorphic to the poset of prime filters of L under inclusion, via the function which sends $x \in J^\infty L^\delta$ to $\{a \in L \mid x \leq a\}$. In summary, we obtain a different construction of L^δ as the complete lattice of up-sets of the poset underlying the Priestley dual space of L .

Recall that a lattice C is isomorphic to a lattice of up-sets of a poset if and only if C is complete and join-generated by its completely join-prime elements (and then, in particular, C is completely distributive); we will call such a lattice a DL^+ .

Definition 2.7. We write \mathbf{DL}^+ for the category whose objects are DL^+ s and whose morphisms are *complete morphisms*, i.e., functions that preserve arbitrary joins and arbitrary meets.

Note that any DL^+ is in particular a complete Heyting algebra, but the morphisms differ; also see Remark 2.16.

We will now recall how the object assignment $L \mapsto L^\delta$ extends to a *functor* from the category \mathbf{DL} , of distributive lattices with lattice morphisms, to \mathbf{DL}^+ . First, for any monotone function $f: L \rightarrow M$, we define, for any $x \in \mathcal{F}(L)$ and $y \in \mathcal{I}(L)$,

$$\bar{f}(x) := \bigwedge\{f(a) \mid a \in L, x \leq a\} \text{ and } \bar{f}(y) := \bigvee\{f(b) \mid b \in L, b \leq y\},$$

where we note that $\mathcal{F}(L) \cap \mathcal{I}(L) = L$, and the right-hand-sides of both formulae are equal to the value of f when x or y belongs to L . We then define, for any $u \in L^\delta$,

$$f^\sigma(u) := \bigvee\{\bar{f}(x) \mid x \in \mathcal{F}(L), x \leq u\} \text{ and } f^\pi(u) := \bigwedge\{\bar{f}(y) \mid y \in \mathcal{I}(L), u \leq y\}.$$

³For proofs of the claims in this paragraph and the next, see, e.g., [7, Corollary 3.43, Corollary 3.48 and Theorem 7.5]; note that [7] represents L^δ using down-sets, but we here use up-sets in order to fit with topos-theoretic tradition.

The extensions f^σ and f^π are different in general, and neither operation respects composition; see, e.g., [11, Sections 2.4 and 2.5]. However, when f is either finite-join-preserving or finite-meet-preserving, then $f^\sigma = f^\pi$ [11, Corollary 2.25], and in this case we write f^δ for the extension. In particular, for any lattice morphism $f: L \rightarrow M$ between distributive lattices, we have a unique extension $f^\delta: L^\delta \rightarrow M^\delta$. The following result, Proposition 2.8, is well-known in the theory of canonical extensions; see e.g. [11, Theorem 3.2], [14, Section 1.1], and [7, Theorem 8.21 and Corollary 8.22]. Note that distributivity is crucial for the adjointness statement in Proposition 2.8; also see [32].

Proposition 2.8. *The assignments $L \mapsto L^\delta$ and $f \mapsto f^\delta$ define a functor $\mathbf{DL} \rightarrow \mathbf{DL}^+$, which is left adjoint to the forgetful functor from \mathbf{DL}^+ to \mathbf{DL} .*

Proposition 2.8 expresses the universal property that, for every morphism of distributive lattices $f: L \rightarrow C$, with C a DL^+ , there exists a unique complete morphism $f^\sharp: L^\delta \rightarrow C$ such that the following diagram commutes:

$$\begin{array}{ccc} & & L^\delta \\ & \nearrow & \vdots \\ L & \xrightarrow{f} & C \\ & & \downarrow f^\sharp \end{array} \quad (2)$$

We call the complete morphism $f^\sharp: L^\delta \rightarrow C$ the *lifting* of f to L^δ .

Remark 2.9. If L is a *finite* distributive lattice, then $L^\delta = L$, and $f^\sharp = f$ for all f .

Since f^\sharp is a complete morphism, it in particular has a left adjoint. We now prove a technical lemma about this left adjoint which will be important when we consider types in Section 5; this lemma is a consequence of known facts in the canonical extensions literature, compare for instance [7, Lemma 7.9]. For the reader's convenience, we state and prove it in the form we will use it below.

Lemma 2.10. *Let L be a distributive lattice, C a DL^+ , and $\alpha: L \rightarrow C$ a morphism of distributive lattices. Let $\beta: C \rightarrow L^\delta$ be the left adjoint of the lifting $\alpha^\sharp: L^\delta \rightarrow C$. Then, for any completely join-prime element $j \in J^\infty C$, we have*

$$\beta(j) = \bigwedge \{a \in L \mid j \leq \alpha(a)\},$$

where the infimum is taken in L^δ . Moreover, $\beta(j) \in J^\infty(L^\delta)$.

Proof. The left-to-right inequality is immediate from the facts that β is left adjoint to α^\sharp , and that α^\sharp extends α . For the right-to-left-inequality, we use that the ideal elements are meet-dense in L^δ . Let $y \in \mathcal{I}(L^\delta)$ be arbitrary such that $\beta(j) \leq y$. By the adjunction $\beta \dashv \alpha^\sharp$, we have $j \leq \alpha^\sharp(y)$. By definition of ideal elements and the fact that α^\sharp preserves arbitrary joins, $\alpha^\sharp(y) = \bigvee \{\alpha(a) \mid a \in L, a \leq y\}$. Since j is completely join-prime, pick $a \in L$ such that $a \leq y$ and $j \leq \alpha(a)$. We conclude that $\bigwedge \{a \in L \mid j \leq \alpha(a)\} \leq y$.

We now turn to the ‘moreover’ part. Since α is finite-join-preserving, the set $\{a \in L \mid j \leq \alpha(a)\}$ is a prime filter of L . Therefore, its infimum is completely join-prime in L^δ by, e.g., [7, Corollary 3.48]. \square

Remark 2.11. Let $f: L \rightarrow K$ be a morphism between distributive lattices. Post-composing f with the embedding $K \hookrightarrow K^\delta$, we obtain a lattice morphism $f': L \rightarrow K^\delta$, and $(f')^\sharp$ coincides with the complete morphism $f^\delta: L^\delta \rightarrow K^\delta$, by uniqueness in (2). Lemma 2.10 implies that the left adjoint of $(f')^\sharp = f^\delta$ restricts to a function $g: J^\infty K^\delta \rightarrow J^\infty L^\delta$. As explained in Section 1, $J^\infty K^\delta$ and $J^\infty L^\delta$ can be seen as the Priestley dual spaces of K and L , respectively, and this function g is then the continuous order-preserving map dual to f under Priestley duality.

2.3 Canonical extensions of coherent doctrines

It is *a priori* not obvious that applying the canonical extension construction to a coherent doctrine would again yield a coherent doctrine. That this is true is an important result due to Coumans [5, Proposition 9], which we recall in Proposition 2.13 below.

Definition 2.12. The *canonical extension* of a coherent doctrine $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ is the functor $P^\delta: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}^+$, defined as the composite $(-)^\delta \circ P$. We denote by η^P the natural transformation $P \Rightarrow P^\delta$ whose component at an object c of \mathcal{C} is the inclusion $Pc \hookrightarrow (Pc)^\delta$.

Proposition 2.13. *For any coherent doctrine $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$, the canonical extension $P^\delta: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}^+$ is a coherent doctrine, and the pair $(\mathbf{id}_{\mathcal{C}}, \eta^P)$ is a morphism of coherent doctrines.*

Note that, in the statement of Proposition 2.13, strictly speaking, $(\mathbf{id}_{\mathcal{C}}, \eta^P)$ is a morphism from P to $U \circ P^\delta$, where U is the forgetful functor $\mathbf{DL}^+ \rightarrow \mathbf{DL}$. When no ambiguity can arise, we will omit such occurrences of U from our notation.

Remark 2.14. Adjunctions lift to the canonical extension: if $\exists_f: L \rightarrow M$ is left adjoint to $f^*: M \rightarrow L$, then $(\exists_f)^\sigma$ is left adjoint to $(f^*)^\pi$; see, e.g., [6, Proposition 3.6]. As a consequence, in the canonical extension P^δ of a coherent doctrine $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$, the left adjoint $\exists_f^{P^\delta}$ is the σ -extension of \exists_f^P . When P is clear from the context, we therefore denote this map by \exists_f^σ .

Definition 2.15. A DL^+ -doctrine is a coherent doctrine that takes values in the category \mathbf{DL}^+ . A *morphism* of DL^+ -doctrines is a morphism (M, α) of coherent doctrines such that each component of α is moreover a complete morphism. Transformations are as in Definition 2.1. We denote the 2-category of DL^+ -doctrines by \mathbf{CohDoc}^+ .

Remark 2.16. As pointed out earlier, any DL^+ is in particular a Heyting algebra, but morphisms in the category \mathbf{DL}^+ do not necessarily preserve the Heyting implication. However, if $Q: \mathcal{D}^{\text{op}} \rightarrow \mathbf{DL}^+$ is a DL^+ -doctrine, then Qf must preserve the Heyting implication for any arrow f in \mathcal{D} . This follows from a more general fact: if $Q: \mathcal{D}^{\text{op}} \rightarrow \mathbf{DL}$ is a functor such that Qd is a complete Heyting algebra for all objects d in \mathcal{D} , then the Frobenius condition is equivalent to the statement that Qf preserves the Heyting implication for all arrows f in \mathcal{D} ; see, e.g., [21, Proposition V.1.1].

Another consequence of the Frobenius condition, which we will use in Section 3, is the following.

Lemma 2.17. *Let $e: L \rightleftharpoons L' : h$ be an adjunction between complete lattices such that $h \circ e = \mathbf{id}_L$, and suppose that, for all $a \in L$ and $b \in L'$, the Frobenius condition $e(a) \wedge b = e(a \wedge h(b))$ is satisfied. Then e preserves all non-empty meets.*

Proof. Let $S \subseteq L$ be non-empty. Write $s_0 := \bigwedge S$ and $s'_0 := \bigwedge e[S]$. Fix $s \in S$. Since $s'_0 \leq e(s)$, we have

$$s'_0 = e(s) \wedge s'_0 = e(s \wedge h(s'_0)) .$$

Using that h preserves all meets, and that $h \circ e = \mathbf{id}_L$, we have

$$h(s'_0) = \bigwedge_{s \in S} h(e(s)) = s_0 .$$

Combining the two equalities, we see that $s'_0 = e(s \wedge s_0) = e(s_0)$, as required. \square

For any category \mathcal{C} , the adjunction $\mathbf{DL} \rightleftharpoons \mathbf{DL}^+$ yields an adjunction between functor categories $[\mathcal{C}^{\text{op}}, \mathbf{DL}] \rightleftharpoons [\mathcal{C}^{\text{op}}, \mathbf{DL}^+]$, where the unit and counit of the latter adjunction are given pointwise by the unit and counit of $\mathbf{DL} \rightleftharpoons \mathbf{DL}^+$. In particular, we have the following result.

Proposition 2.18. *Let $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine, let $Q: \mathcal{D}^{\text{op}} \rightarrow \mathbf{DL}^+$ be a DL^+ -doctrine, and let $(M, \alpha): P \rightarrow Q$ be a coherent morphism. There exists a*

unique natural transformation $\alpha^\sharp: P^\delta \Rightarrow Q \circ M^{\text{op}}$ such that, for each object c of \mathcal{C} , the restriction of α_c^\sharp to Pc is equal to α_c .

By Proposition 2.13, the object assignment $P \mapsto P^\delta$ associates to any coherent doctrine a DL^+ -doctrine, and this assignment readily extends to a functor from \mathbf{CohDoc} to \mathbf{CohDoc}^+ . However, as already observed in the work of Coumans [4, p. 150], this assignment does *not* satisfy a universal property analogous to (2), i.e., the functor $(-)^\delta$ is not left adjoint to the forgetful functor from \mathbf{CohDoc}^+ to \mathbf{CohDoc} .

Let $(M, \alpha): P \rightarrow Q$ be a morphism of coherent doctrines with Q a DL^+ -doctrine. Although, via Proposition 2.18, we always obtain a unique natural transformation $\alpha^\sharp: P^\delta \Rightarrow Q \circ M^{\text{op}}$ whose restriction is α , the pair (M, α^\sharp) can fail to be a morphism of DL^+ -doctrines because the square

$$\begin{array}{ccc} P^\delta c & \xrightarrow{\exists_f^\sigma} & P^\delta c' \\ \alpha_c^\sharp \downarrow & & \downarrow \alpha_{c'}^\sharp \\ QMc & \xrightarrow{\exists_{Mf}^Q} & QMc' \end{array}$$

may not commute. We will give an explicit instance of such a failure in Example 2.21, involving the following two doctrines, which will continue to play an important role in the later sections.

Definition 2.19. The *power set DL^+ -doctrine* is the functor $\mathcal{P}^+: \mathbf{Set}^{\text{op}} \rightarrow \mathbf{DL}^+$ which sends a set S to the complete lattice $\mathcal{P}S$ of subsets of S , ordered by inclusion, and a function $f: S \rightarrow T$ to the inverse image morphism $\mathcal{P}^+ f := f^{-1}: \mathcal{P}T \rightarrow \mathcal{P}S$.

The *power set coherent doctrine* is the functor $\mathcal{P} := U \circ \mathcal{P}^+$, where $U: \mathbf{DL}^+ \rightarrow \mathbf{DL}$ is the forgetful functor.

Remark 2.20. As the names suggest, \mathcal{P}^+ is a DL^+ -doctrine and \mathcal{P} is a coherent doctrine. When no ambiguity can arise, we sometimes also write \mathcal{P} for \mathcal{P}^+ and refer to either as ‘the power set doctrine’.

Example 2.21. Proposition 2.18 gives a natural transformation $u^\sharp: \mathcal{P}^\delta \Rightarrow \mathcal{P}^+$, where u denotes the identity transformation $\mathcal{P} = U \circ \mathcal{P}^+$. We will now show that the pair $(\mathbf{id}_{\mathbf{Set}}, u^\sharp)$ is *not* a morphism of coherent doctrines, because not all the squares for the left adjoints in the square (1) commute. Denote by $!$ the unique function from \mathbb{N} to 1. We claim that the following square does not commute:

$$\begin{array}{ccc} \mathcal{P}(\mathbb{N})^\delta & \xrightarrow{\exists_1^{\mathcal{P}^\delta}} & \mathcal{P}(1)^\delta \\ u_{\mathbb{N}}^\sharp \downarrow & & \downarrow u_1^\sharp \\ \mathcal{P}(\mathbb{N}) & \xrightarrow{\exists_1^{\mathcal{P}}} & \mathcal{P}(1). \end{array}$$

Let \mathfrak{F} be the collection of co-finite subsets of \mathbb{N} , and write $x := \bigwedge \mathfrak{F}$ for the infimum of \mathfrak{F} , computed in $\mathcal{P}(\mathbb{N})^\delta$. Now, using that $u_{\mathbb{N}}^\sharp$ is a complete morphism extending the identity on $\mathcal{P}(\mathbb{N})$, we have

$$u_{\mathbb{N}}^\sharp(x) = \bigcap_{a \in \mathfrak{F}} u_{\mathbb{N}}^\sharp(a) = \bigcap \mathfrak{F} = \emptyset.$$

Thus, $\exists_1^{\mathcal{P}} u_{\mathbb{N}}^\sharp(x) = \emptyset$, since $\exists_1^{\mathcal{P}}$, being a left adjoint, preserves the minimum.

On the other hand, we will now prove that $u_1^\sharp(\exists_1^{\mathcal{P}^\delta} x) \neq \emptyset$. By Remark 2.9, $u_1^\sharp = u_1$, the identity on $\mathcal{P}(1)$, so it is equivalent to show that $\exists_1^{\mathcal{P}^\delta} x \neq \emptyset$. Since $\emptyset = \mathcal{P}^\delta!(\emptyset)$, applying the fact that $\exists_1^{\mathcal{P}^\delta}$ is left adjoint to $\mathcal{P}^\delta!$, we see that this is in turn equivalent to $x \neq \emptyset$ in $\mathcal{P}(\mathbb{N})^\delta$. The latter holds by compactness, since any finite subset of \mathfrak{F} has a non-empty intersection. Thus, $u_1^\sharp \circ \exists_1^{\mathcal{P}^\delta} \neq \exists_1^{\mathcal{P}} \circ u_{\mathbb{N}}^\sharp$ as desired.

Example 2.21 prompts us to pose the question: for *which* morphisms from a coherent doctrine to a DL^+ -doctrine is the lifting still a morphism? We will provide an answer in Section 3.

3 Extensible morphisms of coherent doctrines

The main result of this section is the following characterisation of morphisms from a coherent doctrine to a DL^+ -doctrine that admit liftings; we will call such morphisms *extensible* (see Definition 3.3 below).

Theorem 3.1. *Let $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine, let $Q: \mathcal{D}^{\text{op}} \rightarrow \mathbf{DL}^+$ be a DL^+ -doctrine, and let $(M, \alpha): P \rightarrow Q$ be a coherent morphism. The following are equivalent:*

- (i) *there exists a morphism $P^\delta \rightarrow Q$ in \mathbf{CohDoc}^+ whose restriction to P is (M, α) ;*
- (ii) *the pair $(M, \alpha^\sharp): P^\delta \rightarrow Q$ is a morphism in \mathbf{CohDoc}^+ ;*
- (iii) *for any objects c_1, c_2 of \mathcal{C} , and $\pi_1: c_1 \times c_2 \rightarrow c_1$ the left projection, the following diagram lax commutes:*

$$\begin{array}{ccc} P^\delta(c_1 \times c_2) & \xrightarrow{(\exists_{\pi_1}^P)^\sigma} & P^\delta c_1 \\ \alpha_{c_1 \times c_2}^\sharp \downarrow & \not\approx & \downarrow \alpha_{c_1}^\sharp \\ QM(c_1 \times c_2) & \xrightarrow[\exists_{M\pi_1}^Q]{} & QM c_1. \end{array}$$

In (i) of Theorem 3.1, by the *restriction* of a morphism $(N, \beta): P^\delta \rightarrow Q$ in \mathbf{CohDoc}^+ , we mean the pair $(N, \beta'): P \rightarrow Q$, where the component of β' at an object c of \mathcal{C} is the restriction of the complete morphism $\beta_c: P^\delta c \rightarrow QNc$ to the sublattice $Pc \subseteq P^\delta c$.

In the remainder of this section, let $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine, let $Q: \mathcal{D}^{\text{op}} \rightarrow \mathbf{DL}^+$ be a DL^+ -doctrine, and let $(M, \alpha): P \rightarrow Q$ be a coherent morphism. The crucial step towards proving Theorem 3.1 is the following lemma.

Lemma 3.2. *For any arrow $f: c \rightarrow c'$ in \mathcal{C} , the following diagram lax commutes:*

$$\begin{array}{ccc} P^\delta c & \xrightarrow{\exists_f^\sigma} & P^\delta c' \\ \alpha_c^\sharp \downarrow & \leq & \downarrow \alpha_{c'}^\sharp \\ QM c & \xrightarrow[\exists_{Mf}]{} & QM c'. \end{array}$$

If, moreover, f is a monomorphism, then the diagram strictly commutes.

Proof. Essentially, the first statement of the lemma is appealing to the categorical *mate*. Explicitly, since \exists_f^σ is left adjoint to $P^\delta f$ by Remark 2.14, we have $\mathbf{id}_{P^\delta c} \leq P^\delta f \circ \exists_f^\sigma$. Therefore, $\alpha_c^\sharp \leq \alpha_c^\sharp \circ P^\delta f \circ \exists_f^\sigma$. By naturality of α^\sharp (Proposition 2.18), $\alpha_c^\sharp \circ P^\delta f = QMf \circ \alpha_{c'}^\sharp$, so $\alpha_c^\sharp \leq QMf \circ \alpha_{c'}^\sharp \circ \exists_f^\sigma$. Since \exists_{Mf} is left adjoint to QMf , the desired inequality $\exists_{Mf} \circ \alpha_c^\sharp \leq \alpha_{c'}^\sharp \circ \exists_f^\sigma$ now follows.

If f is a monomorphism, then so is Mf , since M preserves finite limits. Thus, by Remark 2.3, we obtain $QMf \circ \exists_{Mf} = \mathbf{id}_{M c}$. By Lemma 2.17, \exists_{Mf} thus preserves all non-empty meets. Now note that it suffices to prove the equality $\alpha_{c'}^\sharp(\exists_f^\sigma x) = \exists_{Mf}(\alpha_c^\sharp x)$ for a filter element $x \in \mathcal{F}(Pc)$, since any element of $P^\delta c$ is a join of such elements, and both $\alpha_{c'}^\sharp \circ \exists_f^\sigma$ and $\exists_{Mf} \circ \alpha_c^\sharp$ preserve arbitrary joins.

Let $x \in \mathcal{F}(Pc)$ be arbitrary. We have $x = \bigwedge \{a \in Pc, x \leq_{P^\delta c} a\}$, and this meet is

non-empty since \top_{Pc} is in the set on the right hand side. Now

$$\begin{aligned}
\alpha_{c'}^\sharp(\exists_f^\sigma x) &= \alpha_{c'}^\sharp \left(\bigwedge \{ \exists_f a \mid a \in Pc, x \leq a \} \right) && \text{(definition of } \exists_f^\sigma) \\
&= \bigwedge \{ \alpha_{c'}(\exists_f a) \mid a \in Pc, x \leq a \} && (\alpha_{c'}^\sharp \text{ preserves meets and lifts } \alpha_{c'}) \\
&= \bigwedge \{ \exists_{Mf}(\alpha_{ca}) \mid a \in Pc, x \leq a \} && ((M, \alpha) \text{ is a morphism of doctrines)} \\
&= \exists_{Mf}(\alpha_c^\sharp x). && (\exists_{Mf} \circ \alpha_c^\sharp \text{ preserves non-empty meets}),
\end{aligned}$$

which is the required equality. \square

Proof of Theorem 3.1. (i) \iff (ii). If there exists a morphism of DL^+ -doctrines $(M, \beta): P^\delta \rightarrow Q \circ M^{\text{op}}$ whose restriction to P is (M, α) , then β must be natural, and thus equal to α^\sharp by Proposition 2.18. Conversely, if (M, α^\sharp) is a morphism, then its restriction is clearly (M, α) .

(ii) \iff (iii). Since α^\sharp is always a natural transformation from P^δ to $Q \circ M^{\text{op}}$ (Proposition 2.18), (ii) is equivalent to the condition that all squares as in (1) commute for α^\sharp . By Lemma 3.2, this is in turn equivalent to the statement that, for every $f: c \rightarrow c'$, the following square lax commutes:

$$\begin{array}{ccc}
P^\delta c & \xrightarrow{\exists_f^\sigma} & P^\delta c' \\
\alpha_c^\sharp \downarrow & \not\approx & \downarrow \alpha_{c'}^\sharp \\
QM c & \xrightarrow{\exists_{Mf}} & QM c'.
\end{array} \tag{3}$$

It is thus clear that (ii) \Rightarrow (iii). Assume (iii) holds, and let $f: c \rightarrow c'$ be an arbitrary arrow in \mathcal{C} . Since \mathcal{C} has products, we can take the *graph* G_f of f , i.e. the universally induced arrow

$$\begin{array}{ccccc}
& & c & & \\
& f \swarrow & \downarrow G_f & \searrow & \\
c' & \xleftarrow{\pi_1} & c' \times c & \xrightarrow{\quad} & c.
\end{array}$$

The graph of any arrow is a monomorphism, so the equality $\alpha_{c' \times c}^\sharp \circ \exists_{G_f}^\sigma = \exists_{M(G_f)} \circ \alpha_c^\sharp$ holds by Lemma 3.2. Moreover, the inequality $\alpha_{c'}^\sharp \circ \exists_{\pi_1}^\sigma \leq \exists_{M\pi_1} \circ \alpha_{c' \times c}^\sharp$ holds by assumption. Therefore, using that left adjoints are functorial, we obtain

$$\alpha_{c'}^\sharp \circ \exists_f^\sigma = \alpha_{c'}^\sharp \circ \exists_{\pi_1}^\sigma \circ \exists_{G_f}^\sigma \leq \exists_{M\pi_1} \circ \alpha_{c' \times c}^\sharp \circ \exists_{G_f}^\sigma = \exists_{M\pi_1} \circ \exists_{M(G_f)} \circ \alpha_c^\sharp = \exists_{Mf} \circ \alpha_c^\sharp,$$

which is the required inequality. \square

Definition 3.3. We say that a morphism $(M, \alpha): P \rightarrow Q$, with P a coherent doctrine and Q a DL^+ -doctrine, is *extensible* if it satisfies the equivalent properties of Theorem 3.1.

We record an immediate consequence of Theorem 3.1 for future reference.

Corollary 3.4. *Let $P: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine, $Q: \mathcal{D}^{\text{op}} \rightarrow \mathbf{DL}^+$ a DL^+ -doctrine, and $(M, \alpha): P \rightarrow Q$ a morphism. Suppose that, for each object c of \mathcal{C} , $X(c)$ is a subset of $P^\delta(c)$ which join-generates $P^\delta(c)$. Then (M, α) is extensible if and only if for any objects c_1, c_2 of \mathcal{C} and any $x \in X(c_1 \times c_2)$,*

$$\alpha_{c_1}^\sharp(\exists_{\pi_1}^\sigma(x)) \leq \exists_{M\pi_1}(\alpha_{c_1 \times c_2}^\sharp(x)). \tag{4}$$

Proof. Since all maps involved preserve arbitrary joins, the stated condition is equivalent to (iii) in Theorem 3.1. \square

Note that Corollary 3.4 applies when $X(c) \subseteq P^\delta(c)$ is taken to be the subset of filter elements $\mathcal{F}(Pc)$, or, assuming a choice principle, the subset of completely join-prime elements $J^\infty(P^\delta c)$.

Remark 3.5. We explain how our notion of extensibility is a doctrinal version of the notion of p -model that figures in [28] and [5]. In light of the equivalence between completely join-prime elements and prime filters and of the definitions of $(-)^{\sigma}$ and $(-)^{\sharp}$ given in Section 2, using Corollary 3.4, we see that (M, α) is extensible if and only if, for any objects c_1, c_2 of \mathcal{C} and any prime filter ρ in $P(c_1 \times c_2)$, we have the inequality

$$\bigwedge_{a \in \rho} \alpha_{c_1} \exists_{\pi_1}(a) \leq \exists_{M\pi_1} \left(\bigwedge_{a \in \rho} \alpha_{c_1 \times c_2}(a) \right).$$

In [4, Definition 5.2.7], following [28, Section 1], a p -model is defined as a coherent functor $F: \mathcal{C} \rightarrow \mathcal{D}$ between coherent categories such that, for any morphism $f: c \rightarrow c'$ in \mathcal{C} and any prime filter ρ in $\text{Sub}_{\mathcal{C}}(c)$, we have the equality

$$\bigwedge_{a \in \rho} \exists_{Ff}(Fa) = \exists_{Ff} \left(\bigwedge_{a \in \rho} Fa \right).$$

As also noted there, this equality is equivalent to the left-to-right inequality, as the other inequality always holds. Since F is coherent, the left-hand-side can be rewritten as $\bigwedge_{a \in \rho} F\exists_f(a)$. From this and our work earlier in this section, we can conclude that our notion of extensibility is a doctrinal version of the definition of p -model of [28] and [5].

Example 3.6. Let $\mathbf{DL}_f \subseteq \mathbf{DL}$ denote the subcategory of finite distributive lattices, and let \mathbf{Poset}_f denote the category whose objects are finite posets and whose morphisms are monotone functions. Consider the functor $[-, \mathbf{2}]: \mathbf{Poset}_f^{\text{op}} \rightarrow \mathbf{DL}_f \subseteq \mathbf{DL}$, where a finite poset P is sent to the distributive lattice of monotone functions from P into $\mathbf{2}$ equipped with the pointwise ordering (equivalently, the lattice of up-sets of P ; recall that this is the functor that witnesses Birkhoff duality). This functor describes a coherent doctrine, where the left adjoint to $- \circ f: [Q, \mathbf{2}] \rightarrow [P, \mathbf{2}]$, for a monotone map $f: P \rightarrow Q$, sends $g \in [P, \mathbf{2}]$ to the map $\exists_f(g): Q \rightarrow \mathbf{2}$ given by

$$\begin{aligned} \exists_f g(q) &= \bigvee \{ g(p) \mid p \in P \text{ such that } f(p) \leq q \}, \\ &= \begin{cases} \top & \text{if there is } p \in P \text{ with } f(p) \leq q \text{ and } g(p) = \top, \\ \perp & \text{otherwise.} \end{cases} \end{aligned}$$

In an analogous fashion, for any distributive lattice A , the functor $[-, A]: \mathbf{Poset}_f^{\text{op}} \rightarrow \mathbf{DL}$ is a coherent doctrine.

Let us consider morphisms of the form $(\mathbf{id}_{\mathbf{Poset}_f}, \alpha): [-, A] \rightarrow [-, \mathbf{2}]$. Note that $[-, \mathbf{2}]$ is a DL^+ -doctrine (indeed, each fibre is finite), and so it makes sense to ask what it means for $(\mathbf{id}_{\mathbf{Poset}_f}, \alpha)$ to be extensible. By Corollary 3.4, it suffices to check that, for any pair of finite posets P, Q and a filter element $x \in \mathcal{F}([P \times Q, A])$, we have that $\alpha_P^{\sharp}(\exists_{\pi_1}^{\sigma}(x)) \leq \exists_{\pi_1}(\alpha_{P \times Q}^{\sharp}(x))$, i.e. there is an inclusion of the corresponding up-sets. Recall, from the proof of Lemma 3.2, that

$$\alpha_P^{\sharp}(\exists_{\pi_1}^{\sigma}(x)) = \bigwedge \{ \exists_{\pi_1} \alpha_{P \times Q}(g) \mid g \in [P \times Q, A], x \leq g \}.$$

Thus, an element $p \in P$ is contained in the up-set $\alpha_P^{\sharp}(\exists_{\pi_1}^{\sigma}(x)) \subseteq P$ if and only if, for all $g \in [P \times Q, A]$ with $g \geq x$, there exists some $q_g \in Q$ such that the pair (p, q_g) is contained in the up-set $\alpha_{P \times Q}(g) \subseteq P \times Q$. Conversely, we readily calculate that p is contained in $\exists_{\pi_1}(\alpha_{P \times Q}^{\sharp}(x))$ if and only if there exists some $q \in Q$ such that, for all $g \geq x$, the pair (p, q) is contained in $\alpha_{P \times Q}(g)$. Thus, if we are given a family of q_g 's that depend on g for which $(p, q_g) \in \alpha_{P \times Q}(g)$, we must find a single $q' \in Q$ for which $(p, q') \in \alpha_{P \times Q}(g)$. In particular, if Q were a join-semilattice, then we can ensure the existence of such a q' by taking the join $q' := \bigvee_{g \geq x} q_g$ since, by monotonicity, we arrive

at

$$(p, q_g) \in \alpha_{P \times Q}(g) \implies (p, q') \in \alpha_{P \times Q}(g).$$

This pattern of seeking a uniform witness from a dependent choice prefigures the introduction of model-theoretic saturation in the next section.

4 Points

In this section, we give a model-theoretic interpretation to the notion of extensible morphism, and in Theorem 4.8 we establish its close connection with ω -saturation. To this end, we first recall the different notions of *point* of a doctrine, and how they correspond to *models* of a theory.

For L a distributive lattice, we call a morphism from L to the two-element distributive lattice $\mathbf{2}$ a *coherent point* of L . A coherent point $h: L \rightarrow \mathbf{2}$ uniquely determines a prime filter $x_h := h^{-1}(1)$ of L . For C a DL^+ , we call a complete morphism from C to $\mathbf{2}$ a *complete point* of C . A complete point $h: C \rightarrow \mathbf{2}$ uniquely determines a completely join prime element $x_h := \bigwedge h^{-1}(1)$ of C . By the universal property (2), the complete points of L^δ are in bijection with the coherent points of L .

When defining points in the context of coherent and DL^+ doctrines, the distributive lattice $\mathbf{2}$ is replaced by the power set doctrines \mathcal{P} and \mathcal{P}^+ (see Definition 2.19), respectively.

Definition 4.1. Let P be a coherent doctrine. A *coherent point* of P is a morphism $(M, \alpha): P \rightarrow \mathcal{P}$ in \mathbf{CohDoc} . We denote the category of points of P and their transformations by $\text{Pt}(P)$.

Let Q be a DL^+ -doctrine. A *complete point* of Q is a morphism $(M, \alpha): Q \rightarrow \mathcal{P}^+$ in \mathbf{CohDoc}^+ . We denote the category of complete points of Q and their transformations by $\text{Pt}^+(Q)$.

Let P be a coherent doctrine. Instantiating Definition 3.3 with $Q := \mathcal{P}^+$, we see that the complete points of P^δ are in bijection with the *extensible* coherent points of P .

Let us now briefly recall the connection between coherent points and models of a coherent theory. For simplicity, we restrict ourselves to the case of coherent theories that are single-sorted and relational, i.e., having no function symbols. A *coherent formula* is built from atomic predicates using logical operators $\wedge, \vee, \perp, \top$ and \exists , and a *coherent theory* is then formulated using appropriate sequent derivation rules of the form $\varphi(x) \vdash \psi(x)$, which should be read as “for all x , if $\varphi(x)$ then $\psi(x)$ ” (for more on the syntax of coherent logic, see [20, §D1]). A *model* \mathcal{M} of a coherent theory \mathbb{T} is a set, together with an interpretation of each symbol of the signature of \mathbb{T} , such that all sequents in \mathbb{T} are interpreted in \mathcal{M} as true statements.

In Definition 4.2, we recall how to build a syntactic doctrine for coherent theory, which is a coherent doctrine on base category $\mathbf{FinSet}^{\text{op}}$, and in Remark 4.4 we recall how coherent points of this doctrine correspond to models of the coherent theory. For further details and proofs, see, e.g., [29, § 8] or [4, § 5].

Definition 4.2. Let \mathbb{T} be a (single-sorted, relational) coherent theory. The *syntactic doctrine* of \mathbb{T} is the functor $D_{\mathbb{T}}: \mathbf{FinSet} \rightarrow \mathbf{DL}$ which sends a finite set X to the distributive lattice of \mathbb{T} -equivalence classes of \mathbb{T} -formulae whose free variables are in X , and which sends a function $r: X \rightarrow Y$ to the homomorphism $D_{\mathbb{T}}r: D_{\mathbb{T}}X \rightarrow D_{\mathbb{T}}Y$ induced by the substitution function $\varphi(x_1, \dots, x_n) \mapsto \varphi(r(x_1), \dots, r(x_n))$.

In what follows, by a common abuse of notation, we will treat elements of $D_{\mathbb{T}}X$ as formulae, rather than as \mathbb{T} -equivalence classes of formulae. This abuse of notation will be harmless since all our arguments can be done modulo \mathbb{T} -equivalence.

Remark 4.3. The functor $D_{\mathbb{T}}: \mathbf{FinSet} \rightarrow \mathbf{DL}$ defines a coherent doctrine. For a function $r: X \rightarrow Y$, the left adjoint to $D_{\mathbb{T}}r$ is the map $\exists_r^{D_{\mathbb{T}}}: D_{\mathbb{T}}Y \rightarrow D_{\mathbb{T}}X$ given by

$$\psi(y_1, \dots, y_m) \mapsto \exists y_1, \dots, y_m. \psi(y_1, \dots, y_m) \wedge x_1 = y_{r(1)} \wedge \dots \wedge x_n = y_{r(n)},$$

where $y_{r(i)}$ denotes the image of $x_i \in X$ under r . In particular, for a coproduct inclusion $i_X: X \rightarrow X + Y$, we have $\exists_{i_X}^{D_{\mathbb{T}}}(\psi(x, y)) \equiv \exists y. \psi(x, y)$ (where x and y now denote tuples of variables), and for the co-diagonal map $\nabla_X: X + X \rightarrow X$, we have $\exists_{\nabla_X}^{D_{\mathbb{T}}}(\varphi(x)) \equiv \varphi(x) \wedge x = x'$. For further details that this satisfies the Beck-Chevalley condition, etc., see [36].

Remark 4.4. When \mathbb{T} is a coherent theory and $D_{\mathbb{T}}$ its syntactic doctrine, the coherent points of $D_{\mathbb{T}}$ correspond to the models of \mathbb{T} : if \mathcal{M} is a model with underlying set M , then we can define a morphism $(M, \llbracket - \rrbracket^{\mathcal{M}}): D_{\mathbb{T}} \rightarrow \mathcal{P}$ as follows. Let $M: \mathbf{FinSet}^{\text{op}} \rightarrow \mathbf{Set}$ be the functor defined on objects by sending a finite set X to the set M^X of X -tuples in M , and defined on morphisms by sending a function $r: X \rightarrow Y$ to the function $Mr: M^Y \rightarrow M^X$, which associates any Y -tuple $m \in M^Y$ with the X -tuple $m \circ r \in M^X$. For any formula φ with free variables in X , let $\llbracket \varphi \rrbracket_X^{\mathcal{M}} := \{m \in M^X \mid M \models \varphi(m)\}$, i.e. the interpretation of φ in \mathcal{M} . Since $\mathcal{M} \models \mathbb{T}$ is a model, this assignment is well-defined on \mathbb{T} -equivalence classes of formulae, and the pair $(M, \llbracket - \rrbracket^{\mathcal{M}})$ thus defined is a morphism in \mathbf{CohDoc} . We call $(M, \llbracket - \rrbracket^{\mathcal{M}})$ the *associated point* of \mathcal{M} . When \mathcal{M} is clear from the context, we will write $\llbracket - \rrbracket$ instead of $\llbracket - \rrbracket^{\mathcal{M}}$. Note that the model \mathcal{M} can be recovered from its associated point: the underlying set is the value of M at a singleton set and $\llbracket - \rrbracket^{\mathcal{M}}$ in particular records the interpretation of all predicate symbols. This assignment $\mathcal{M} \mapsto (M, \llbracket - \rrbracket^{\mathcal{M}})$ is the object part of an isomorphism between the category of models of \mathbb{T} and the category of points of $D_{\mathbb{T}}$ (where the morphisms of points are the transformations of Definition 2.1). See, e.g., [40, § 1] for more details.

In Theorem 4.8, we will instantiate Corollary 3.4 to the specific case of a point associated to a model of a coherent theory, in the sense of Remark 4.4. Working towards this, let us calculate what the two sides of the inequality (4) in Corollary 3.4 look like in this specific case.

Let \mathcal{M} be a model of a coherent theory \mathbb{T} , with associated point $(M, \llbracket - \rrbracket): D_{\mathbb{T}} \rightarrow \mathcal{P}$. For any finite sets X, Y , their product in $\mathbf{FinSet}^{\text{op}}$ can be realized as the disjoint union $X + Y$ and the left projection is then realized as the left inclusion $i_X: X \rightarrow X + Y$. The homomorphism $D_{\mathbb{T}}i_X: D_{\mathbb{T}}X \rightarrow D_{\mathbb{T}}(X + Y)$ includes the formulae with free variables in X as a sublattice of the formulae with free variables in $X + Y$. Recall from Remark 4.3 that its left adjoint $\exists_{i_X}: D_{\mathbb{T}}(X + Y) \rightarrow D_{\mathbb{T}}X$ sends a formula $\varphi(x, y)$ to $\exists y. \varphi(x, y)$, where ‘ x ’ denotes the tuple of variables in X and ‘ y ’ denotes the tuple of variables in Y . Let t be a filter element of $D_{\mathbb{T}}^{\delta}(X + Y)$, and write $\tau := \{\varphi \in D_{\mathbb{T}}(X + Y) \mid t \leq \varphi\}$. The σ -extension map $\exists_{i_X}^{\sigma}$ sends t to $\bigwedge \{\exists y. \varphi \mid \varphi \in \tau\}$, a filter element in $D_{\mathbb{T}}^{\delta}X$. Since $\llbracket - \rrbracket_X^{\#}$ preserves meets, we thus obtain the following expression for the left-hand-side of (4):

$$\llbracket \exists_{i_X}^{\sigma}(t) \rrbracket_X^{\#} = \left\{ m \in M^X \mid \text{for all } \varphi \in \tau, \mathcal{M} \models \exists y. \varphi(m, y) \right\}. \quad (5)$$

For the right-hand-side of (4), first note, using that $\llbracket - \rrbracket_{X+Y}^{\#}$ preserves meets,

$$\llbracket t \rrbracket_{X+Y}^{\#} = \left\{ (m, n) \in M^{X+Y} \mid \text{for all } \varphi \in \tau, \mathcal{M} \models \varphi(m, n) \right\}.$$

The function $Mi_X: M^{X+Y} \rightarrow M^X$ sends an $(X + Y)$ -tuple (m, n) to the X -tuple m , and $\exists_{Mi_X}: \mathcal{P}(M^{X+Y}) \rightarrow \mathcal{P}(M^X)$ takes the direct image under Mi_X . Therefore,

$$\exists_{Mi_X}(\llbracket t \rrbracket_{X+Y}^{\#}) = \left\{ m \in M^X \mid \text{there exists } n \in M^Y \text{ such that for all } \varphi \in \tau, \mathcal{M} \models \varphi(m, n) \right\}. \quad (6)$$

We will now use the equalities (5) and (6) to give a new perspective on the following

notions from positive model theory; cf. [24, §3.3] for instance.

Definition 4.5. Let \mathcal{M} be a model of a coherent theory \mathbb{T} and let X, Y be disjoint sets of variables. A Y -type with X -parameters in \mathcal{M} is a pair (τ, m) , where τ is a set of formulae with free variables in $X \cup Y$, and $m \in M^X$ is an X -tuple of elements of \mathcal{M} .

Let (τ, m) be a Y -type with X -parameters in \mathcal{M} . We say that (τ, m) is:

- (i) *consistent* if, for every $\varphi \in \tau$, there exists $n \in M^Y$ such that $\mathcal{M} \models \varphi(m, n)$;
- (ii) *realised* if there exists $n \in M^Y$ such that, for every $\varphi \in \tau$, $\mathcal{M} \models \varphi(m, n)$.

The model \mathcal{M} is ω -saturated if, for any finite X, Y , any consistent Y -type with X -parameters in \mathcal{M} is realised.

Remark 4.6. We have deliberately omitted the word ‘positively’ from our notion of ω -saturation, since it is clear by context. See Remark 4.11 for further discussion of our choice of terminology.

Remark 4.7. Note that our notion of type is often called a *partial type* by others (e.g., [24, Definition 2.2.3]). Classically, partial types in context Y are distinguished from *complete* types, the maximal (with respect to inclusion) consistent subsets of $D_{\mathbb{T}}(Y)$, which are more commonly abbreviated to just ‘types’. Note again that our types are with respect to positive logic, and so will in general be different from the types found in classical model theory (cf. [17]).

Later, in addition to partial and complete types, we will wish to speak of *prime* types. In classical logic, i.e., where every formula has a negation, we are motivated to consider complete types since a type τ is complete if and only if there is a model \mathcal{M} and some $m \in M^Y$ such that

$$\tau = \{ \varphi \text{ classical} \mid \mathcal{M} \models \varphi(m) \}.$$

However, in positive logic, this equivalence breaks down. Instead, for a filter τ of coherent formulae with free variables in Y , we have the following equivalence (also see Subsection 5.3 below):

- (i) τ is of the form $\text{tp}_{\mathcal{M}}(m) = \{ \varphi \mid \mathcal{M} \models \varphi(m) \}$ for some model \mathcal{M} and tuple $m \in M^Y$ (note that this is the *positive* type of m), which we call the *type of m* ;
- (ii) τ is *prime* in the sense that $\perp \notin \tau$ and if $\varphi \vee \psi \in \tau$, then either $\varphi \in \tau$ or $\psi \in \tau$.

For our purposes, it is more fitting to speak of partial types and prime types. Nonetheless, complete (positive) types play an important role in positive model theory – they are precisely the types realised by elements of a *positively closed* model ([24, Proposition 2.2.2]).

Theorem 4.8. *Let \mathcal{M} be a model of a coherent theory \mathbb{T} , with the associated point $(M, \llbracket - \rrbracket): D_{\mathbb{T}} \rightarrow \mathcal{P}$. The following are equivalent:*

- (i) *the point $(M, \llbracket - \rrbracket)$ is extensible;*
- (ii) *the model \mathcal{M} is ω -saturated.*

In particular, the ω -saturated models of \mathbb{T} correspond to the complete points of $D_{\mathbb{T}}^{\delta}$.

Proof. By Corollary 3.4 instantiated for the morphism $(M, \llbracket - \rrbracket)$, (i) is equivalent to the statement that $\llbracket \exists_{i_X}^{\sigma}(t) \rrbracket_X^{\sharp} \subseteq \exists_{Mi_X}(\llbracket t \rrbracket_{X+Y}^{\sharp})$ holds for any finite sets X, Y , with $i_X: X \rightarrow X + Y$ the left inclusion, and any $t \in \mathcal{F}(D_{\mathbb{T}}^{\delta}(X))$. By the equalities (5) and (6), this is in turn equivalent to the statement that \mathcal{M} is ω -saturated. The ‘in particular’ statement now follows in light of Remark 4.4. \square

Remark 4.9. For the reader familiar with the theory of polyadic spaces developed in [15], we explain how Theorem 4.8 relates to the definition of ‘ ω -saturated model’ given there.

Let \mathbb{T} be a single-sorted relational coherent theory⁴, with associated syntactic doctrine $D_{\mathbb{T}}: \mathbf{FinSet} \rightarrow \mathbf{DL}$. In [15], the authors consider the *polyadic Priestley space dual to $D_{\mathbb{T}}$* , which is defined as the functor $S_{\mathbb{T}}: \mathbf{FinSet}^{\text{op}} \rightarrow \mathbf{Priestley}$ obtained by post-composing $D_{\mathbb{T}}^{\text{op}}$ with the Priestley duality functor $\mathbf{DL}^{\text{op}} \rightarrow \mathbf{Priestley}$. Thus, for any finite set n , the poset underlying $S_{\mathbb{T}}(n)$ is $J^{\infty} D_{\mathbb{T}}^{\delta}(n)$, or alternatively, the prime filters of $D_{\mathbb{T}}(n)$ under inclusion. It follows that the canonical extension doctrine $D_{\mathbb{T}}^{\delta}$ is isomorphic to the doctrine $\text{Up} \circ S_{\mathbb{T}}$, where $\text{Up}: \mathbf{Priestley}^{\text{op}} \rightarrow \mathbf{DL}^+$ is the functor that sends a Priestley space to the DL^+ of up-sets of its underlying poset.

In this context, a model of \mathbb{T} , i.e., a point of $D_{\mathbb{T}}$, yields, via duality, a set X together with a natural transformation from $X^{(-)}$ to $S_{\mathbb{T}}$, where $X^{(-)}$ denotes the functor $\mathbf{FinSet}^{\text{op}} \rightarrow \mathbf{Set}$ that sends a finite set n to X^n and acts on morphisms by precomposition. Among all natural transformations $X^{(-)} \Rightarrow S_{\mathbb{T}}$, those which correspond to models are precisely those satisfying a certain exactness condition called the weak interpolation property [15, Proposition 5.6]. In this setting, ω -saturated models are then *defined* as those for which a stronger exactness property called the interpolation property holds [15, Definition 5.7].

Given a natural transformation $\tau: X^{(-)} \Rightarrow S_{\mathbb{T}}$, we can take its inverse image component-wise and obtain a natural transformation $\tau^{-1}: \text{Up} \circ S_{\mathbb{T}} \Rightarrow \mathcal{P} \circ X^{(-)}$. Under the isomorphism $D_{\mathbb{T}}^{\delta} \cong \text{Up} \circ S_{\mathbb{T}}$, τ^{-1} may be seen as a natural transformation $D_{\mathbb{T}}^{\delta} \Rightarrow \mathcal{P} \circ X^{(-)}$. Now, τ has the interpolation property if and only if the pair $(X^{(-)}, \tau^{-1})$ is a \mathbf{CohDoc}^+ -morphism $D_{\mathbb{T}}^{\delta} \rightarrow \mathcal{P}$. Thus, ω -saturated models in the sense of [15] correspond precisely to \mathbf{CohDoc}^+ -morphisms $D_{\mathbb{T}}^{\delta} \rightarrow \mathcal{P}$. By Theorem 3.1, these are precisely the extensible points of $D_{\mathbb{T}}$. In conclusion, the ω -saturated models in the sense of [15] are the same as ours, by Theorem 4.8.

Example 4.10. Suppose that $D_{\mathbb{T}}: \mathbf{FinSet} \rightarrow \mathbf{DL}_f \subseteq \mathbf{DL}$ factors through the subcategory of finite distributive lattices. Since the canonical extension is inert on finite distributive lattices, and so $D_{\mathbb{T}}^{\delta} = D_{\mathbb{T}}$, it follows that every point of $D_{\mathbb{T}}$ is extensible. Thus, if \mathbb{T} is a coherent theory such that, for any n , there are a finite number of formulae in n free variables up to \mathbb{T} -equivalence, every model is ω -saturated (cf. [18, Corollary 12.2.13]). Note that if \mathbb{T} is a classical theory, then there are finitely many formulae in n free variables up to equivalence, for any n , if and only if \mathbb{T} is *ω -categorical*, meaning that there is a unique (up to isomorphism) countable model of \mathbb{T} (see also Example 5.20).

Remark 4.11. We feel that to a select group of readers the correspondence between extensible points and ω -saturation may be already known (cf. [2]), but perhaps not as publicised as it ought to be, and so we conclude this section with a few bibliographic remarks. As discussed in Remark 3.5, the notion of extensible point we have been discussing is called *p-model* in Makkai's terminology ([28, §1]), which is also the terminology used by Coumans in [4, 5]. Makkai writes in [28] that the notion of *p-model* is a generalisation of the notion of (classical) ω -saturation, a sentiment that is echoed in [8, §6.3], and in [28, §1.4] it is proved that a classically ω -saturated model is a *p-model* (i.e., a positively ω -saturated model in the sense we have used here). The correspondence is more explicit with Butz's work [2], for instance in Proposition 7.10 *op. cit.*, which when combined with the observations from [5] expresses that $D_{\mathbb{T}}^{\delta}$ is (positively) ω -saturated in an internal sense.

5 Presheaves of points

Having studied individual models of a coherent theory \mathbb{T} , it is natural to next study categories of models, i.e., subcategories $\mathcal{K} \hookrightarrow \text{Mod}(\mathbb{T}) \cong \text{Pt}(D_{\mathbb{T}})$. In fact, it will be more convenient to consider arbitrary functors $\mathcal{K} \rightarrow \text{Pt}(D_{\mathbb{T}})$. This is because, as we shall see

⁴The definitions in [15] are given at a greater level of generality, but for simplicity we specialize to the setting of single-sorted relational coherent theories here.

in Construction 5.1, functors $\mathcal{K} \rightarrow \text{Pt}(D_{\mathbb{T}})$ bijectively correspond to presheaf-valued points of $D_{\mathbb{T}}$, i.e., morphisms of coherent doctrines $D_{\mathbb{T}} \rightarrow \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$, where $\text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ is the doctrine of *subpresheaves*. This correspondence can be summarised in the slogan ‘models internal to $\mathbf{Set}^{\mathcal{K}}$ are \mathcal{K} -indexed models’.

Inspired by [5, §6], we wish to identify conditions on a class of models \mathcal{K} so that \mathcal{K} fully describes ‘all the types’ of the theory \mathbb{T} , in a way that will be made precise in Theorem 5.15 below.

5.1 Presheaf-valued points and presheaves of points

Let \mathcal{K} be a small category. We denote by $\mathbf{Set}^{\mathcal{K}}$ the category of (covariant) presheaves on \mathcal{K} , i.e., the objects are functors $\mathcal{K} \rightarrow \mathbf{Set}$, and the morphisms are natural transformations between such functors. First, we give the necessary details to justify our slogan: ‘models internal to $\mathbf{Set}^{\mathcal{K}}$ are \mathcal{K} -indexed models’. In Definition 4.1, we defined a coherent point of a doctrine D to be a coherent doctrine morphism from D to \mathcal{P} . Generalizing this definition, a *presheaf-valued* coherent point of D is, by definition, a coherent doctrine morphism from D to $\text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$. Here, note that, since $\mathbf{Set}^{\mathcal{K}}$ is a coherent category (see, e.g., [20, p. 34]), $\text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ is a coherent doctrine by Example 2.4; we return to the structure of the doctrine $\text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ in more detail in Subsection 5.2, where we will show in particular that $\text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ is a DL^+ -doctrine.

The main interest of presheaf-valued coherent points of D is that they correspond to *presheaves of points*. Below, we explicitly show how to obtain a presheaf of points out of a presheaf-valued point. In essence, this is a currying argument. We will first give the explicit construction, and then explain (cf. Remark 5.2) how one might arrive at it in a more principled way using the adjunction between coherent doctrines and coherent categories.

Construction 5.1. Let $D: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine, and let \mathcal{K} be a small category. Let $(M, \alpha): D \rightarrow \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ be a coherent doctrine morphism.

The *associated presheaf of points* is the following functor $\mathcal{M}: \mathcal{K} \rightarrow \text{Pt}(D)$. For k an object of \mathcal{K} , we define a point $\mathcal{M}k := (M_k, \alpha_k): D \rightarrow \mathcal{P}$, as follows:

- the functor $M_k: \mathcal{C}^{\text{op}} \rightarrow \mathbf{Set}$ is defined by sending $c \in \text{ob } \mathcal{C}$ to the set $(Mc)k$, and a morphism $f: c \rightarrow c'$ of \mathcal{C} to the function $(Mf)_k: (Mc)k \rightarrow (M'c)k$;
- the natural transformation $\alpha_k: D \Rightarrow \mathcal{P} \circ (M_k)^{\text{op}}$ has its component at $c \in \text{ob } \mathcal{C}$ defined to be the function $\alpha_{k,c}: Dc \rightarrow \mathcal{P}(M_k c)$ that sends $\varphi \in Dc$ to

$$\alpha_{k,c}(\varphi) := (\alpha_c \varphi)k ,$$

i.e., $\alpha_{k,c}(\varphi)$ is the subset of $(Mc)k$ obtained by evaluating at the object k the subpresheaf $\alpha_c \varphi$ of Mc .

For $r: k \rightarrow k'$ a morphism of \mathcal{K} , we define $\mathcal{M}r: \mathcal{M}k \rightarrow \mathcal{M}k'$ to be the natural transformation from M_k to $M_{k'}$ whose component at $c \in \text{ob } \mathcal{C}$ is the function $(M_c)r: (M_c)k \rightarrow (M_c)k'$.

We omit the proof that \mathcal{M} is indeed a well-defined functor $\mathcal{M}: \mathcal{K} \rightarrow \text{Pt}(D)$, and that the data of a morphism (M, α) can be entirely recovered from its associated presheaf of points, so that we in fact have a bijective correspondence between presheaves of points of D and presheaf-valued points of D . The following remark justifies this omission.

Remark 5.2. We explain how the theory of coherent categories can be used to clarify and justify Construction 5.1. To do so, and only for the purposes of this remark, we will make use of the 2-category \mathbf{Coh} of coherent categories, coherent functors, and natural transformations, and the fact that the functor $\text{Sub}_{(-)}: \mathbf{Coh} \rightarrow \mathbf{CohDoc}$ has a fully faithful left pseudo-adjoint $\mathcal{A}: \mathbf{CohDoc} \rightarrow \mathbf{Coh}$, cf., e.g., [33, § 1] or [4, § 5.1.3]. In essence, the category $\mathcal{A}(D)$ is the ‘syntactic category’ of functional predicates in the internal language of D .

Let $D: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine, let \mathcal{K} be a small category, and write $A := \mathcal{A}(D)$. By the pseudo-adjunction $\mathcal{A} \dashv \text{Sub}_{(-)}$, we have an equivalence of categories,

$$\mathbf{CohDoc}(D, \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}) \simeq \mathbf{Coh}(A, \mathbf{Set}^{\mathcal{K}}). \quad (7)$$

The category $\mathbf{Coh}(A, \mathbf{Set}^{\mathcal{K}})$ is a subcategory of the functor category $[A, \mathbf{Set}^{\mathcal{K}}]$. The universal property of the functor category construction $\mathbf{Set}^{(-)}$ gives an isomorphism of categories

$$[A, \mathbf{Set}^{\mathcal{K}}] \cong [\mathcal{K} \times A, \mathbf{Set}] \cong [\mathcal{K}, \mathbf{Set}^A].$$

Since finite limits, existential quantification, and finite joins are all computed point-wise in the coherent category $\mathbf{Set}^{\mathcal{K}}$, the above isomorphism restricts to an isomorphism

$$\mathbf{Coh}(A, \mathbf{Set}^{\mathcal{K}}) \cong [\mathcal{K}, \mathbf{Coh}(A, \mathbf{Set})]. \quad (8)$$

Applying the pseudo-adjunction $\mathcal{A} \dashv \text{Sub}_{(-)}$ once more, and noting that $\text{Sub}_{\mathbf{Set}} = \mathcal{P}$, we have

$$\mathbf{Coh}(A, \mathbf{Set}) \simeq \mathbf{CohDoc}(D, \mathcal{P}) = \text{Pt}(D). \quad (9)$$

Putting (7), (8), (9) together, we obtain an equivalence

$$\mathbf{CohDoc}(D, \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}) \simeq [\mathcal{K}, \text{Pt}(D)],$$

and the object part of this equivalence functor was made explicit in Construction 5.1.

We will give another viewpoint on morphisms $D \rightarrow \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ in Section 7, where we will see that they also correspond to geometric morphisms from the presheaf topos $\mathbf{Set}^{\mathcal{K}}$ to a topos $\mathcal{E}_D^{\text{coh}}$ defined there.

Remark 5.3. Recall from Section 4 that, if D is the syntactic doctrine of a coherent theory \mathbb{T} , then coherent points of D can be seen as models of \mathbb{T} . In this case, if $(M, \alpha): D \rightarrow \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ is a coherent morphism, then Construction 5.1 yields a functor from \mathcal{K} to models of \mathbb{T} .

5.2 Subpresheaves and orbits of elements

We now examine the coherent doctrine $\text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ in more detail. Recall that subobjects in $\mathbf{Set}^{\mathcal{K}}$ can be described explicitly as follows: for any presheaf $X: \mathcal{K} \rightarrow \mathbf{Set}$, a *subpresheaf* of X is a functor $Y: \mathcal{K} \rightarrow \mathbf{Set}$ such that $Yk \subseteq Xk$ for every object k of \mathcal{K} , and, for any morphism $f: k \rightarrow k'$ of \mathcal{K} , the function Yf is the restriction of Xf to Yk . In other words, a subobject of the presheaf X is uniquely determined by a family of subsets $(Yk)_{k \in \text{ob } \mathcal{K}}$ of $(Xk)_{k \in \text{ob } \mathcal{K}}$ which has the property of being stable under the action of X on morphisms. From this description, it follows that $\text{Sub}_{\mathbf{Set}^{\mathcal{K}}}(X)$ is a complete lattice, in which joins and meets are computed by taking object-wise unions and intersections, respectively. It follows that $\text{Sub}_{\mathbf{Set}^{\mathcal{K}}}(X)$ is a completely distributive lattice.

For any presheaf $X: \mathcal{K} \rightarrow \mathbf{Set}$, an *element* of X is a pair (k, x) , where $k \in \text{ob } \mathcal{K}$ and $x \in Xk$. The *category of elements* of X , which we denote by $\text{el } X$, has as its objects the elements of X , and as its morphisms $(k, x) \rightarrow (k', x')$ those morphisms $f: k \rightarrow k'$ in \mathcal{K} such that $(Xf)(x) = x'$. The assignment $X \mapsto \text{el } X$ is the object part of an equivalence between $\mathbf{Set}^{\mathcal{K}}$ and a category of (discrete) *fibrations* over \mathcal{K}^{op} (see, e.g., [26, Theorem 2.1.2]), but we will only use the object part here.

Note that any element (k, x) of X defines a subpresheaf

$$\begin{aligned} \circ_{k,x}: \mathcal{K} &\rightarrow \mathbf{Set} \\ k' &\mapsto \{x' \in Xk' \mid \text{there exists } f: k \rightarrow k' \text{ such that } (Xf)x = x'\} \end{aligned}$$

of X , which we call the *orbit* of the element (k, x) in the presheaf X .

Example 5.4. Taking \mathcal{K} to be a group G , viewed as a one-object category, a presheaf is then simply a set X equipped with a G -action. In this setting, our notion of orbit is precisely the usual sense of the word.

Lemma 5.5. *For any element (k, x) of X , $\circ_{k,x}$ is the smallest subpresheaf Y of X such that $x \in Yk$.*

Proof. We have $x \in \circ_{k,x}k$, and if Y is any subpresheaf of X such that $x \in Yk$, then the definition of subpresheaves implies that $\circ_{k,x}k' \subseteq Yk'$ for any object k' of \mathcal{K} . \square

Lemma 5.5 in particular yields that, for any subpresheaf Y of X , we have

$$Y = \bigvee_{(k,x) \in \text{ob el } Y} \circ_{k,x} . \quad (10)$$

Proposition 5.6. *For any small category \mathcal{K} , and any presheaf $X: \mathcal{K} \rightarrow \mathbf{Set}$, the completely join-prime elements of $\text{Sub}_{\mathbf{Set}\mathcal{K}}(X)$ are precisely the orbits of elements of X . Moreover, the functor $\text{Sub}_{\mathbf{Set}\mathcal{K}}: \mathbf{Set}^{\mathcal{K}} \rightarrow \mathbf{DL}^+$ is a \mathbf{DL}^+ -doctrine.*

Proof. Let $X \in \mathbf{Set}^{\mathcal{K}}$. We show that the completely join-prime elements of $\text{Sub}_{\mathbf{Set}\mathcal{K}}(X)$ are precisely the orbits of elements of X . If $(Y_i)_{i \in I}$ is a family of subpresheaves of X and $\circ_{k,x} \leq \bigvee_{i \in I} Y_i$, then, since $x \in \circ_{k,x}k$, there exists $i \in I$ such that $x \in Y_i k$, and thus $\circ_{k,x} \leq Y_i$. Conversely, if J is completely join-prime in $\text{Sub}_{\mathbf{Set}\mathcal{K}}(X)$, then (10) implies $J = \circ_{k,x}$ for some element (k, x) of J . The equality (10) shows that $\text{Sub}_{\mathbf{Set}\mathcal{K}}(X)$ is join-generated by the orbits, which are completely join-prime elements, and hence is a \mathbf{DL}^+ for any presheaf X . \square

Remark 5.7. Let $X: \mathcal{K} \rightarrow \mathbf{Set}$ be a presheaf. The restriction of the partial order \leq on $\text{Sub}_{\mathbf{Set}\mathcal{K}}(X)$ to the set $J^\infty(\text{Sub}_{\mathbf{Set}\mathcal{K}}(X))$ of orbits of elements of X can be described as follows, using Lemma 5.5: for any elements (k, x) and (k', x') of X , we have

$$\circ_{k',x'} \leq \circ_{k,x} \iff x' \in \circ_{k,x}k' \iff \text{there exists } f: (k, x) \rightarrow (k', x') \text{ in } \text{el } X .$$

Thus, the poset $J^\infty(\text{Sub}_{\mathbf{Set}\mathcal{K}}(X))$ is the poset reflection of the category $(\text{el } X)^{\text{op}}$.

5.3 Types in a doctrine

As we already saw in Section 4 above, the extensibility of a point $(M, \llbracket - \rrbracket)$ of a doctrine $D_{\mathbb{T}}$ is closely related to the types realised by the model \mathcal{M} of \mathbb{T} that $(M, \llbracket - \rrbracket)$ corresponds to. We now study types in the setting of a general coherent doctrine.

Definition 5.8. Let $D: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine and $(M, \alpha): D \rightarrow \mathcal{P}$ a point of D . For $c \in \text{ob } \mathcal{C}$ and $x \in Mc$, the *type* of the element x in (M, α) is

$$\text{tp}_{M,\alpha}(x) := \{\varphi \in Dc \mid x \in \alpha_c \varphi\} .$$

Remark 5.9. Note that $\text{tp}_{M,\alpha}(x)$ is a prime filter of Dc . Although we will not expand upon this point here, it follows from the completeness theorem for coherent logic that every prime filter of Dc is of the form $\text{tp}_{M,\alpha}(x)$ for some point $(M, \alpha): D \rightarrow \mathcal{P}$ and $x \in Mc$.

We can characterise types of elements abstractly in terms of the canonical extension and the lifting α_c^\sharp . Indeed, let $\beta_c: \mathcal{P}(Mc) \rightarrow D^\delta c$ be the left adjoint of the lifting $\alpha_c^\sharp: D^\delta c \rightarrow \mathcal{P}(Mc)$. Then, applying Lemma 2.10 in the context of Definition 5.8, we see that, for any $x \in Mc$, $\beta_c(\{x\})$ is the infimum of $\text{tp}_{M,\alpha}(x)$ in $D^\delta c$. We will now generalize this observation to presheaf-valued points.

Notation 5.10. Throughout the rest of this section, let $D: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine, \mathcal{K} a small category, $(M, \alpha): D \rightarrow \text{Sub}_{\mathbf{Set}\mathcal{K}}$ a presheaf-valued point of D , and $\mathcal{M}: \mathcal{K} \rightarrow \text{Pt}(D)$ the associated presheaf of points in the sense of Construction 5.1. For any $c \in \text{ob } \mathcal{C}$, we write $\alpha_c^\sharp: D^\delta c \rightarrow \text{Sub}_{\mathbf{Set}\mathcal{K}}(Mc)$ for the lifting of α_c , and β_c for its left adjoint.

Lemma 5.11. *For any $c \in \text{ob } \mathcal{C}$ and any element (k, x) of M_c ,*

$$\beta_c(\mathfrak{o}_{k,x}) = \bigwedge \text{tp}_{M_k, \alpha_k}(x), \quad (11)$$

where the infimum is taken in $D^\delta c$. Moreover, $\beta_c(\mathfrak{o}_{k,x}) \in J^\infty(D^\delta c)$.

Proof. Let (k, x) be an element of M_c . By Proposition 5.6, $\mathfrak{o}_{k,x}$ is completely join-prime, and so, applying Lemma 2.10, we have

$$\beta_c(\mathfrak{o}_{k,x}) = \bigwedge \{\varphi \in Dc \mid \mathfrak{o}_{k,x} \leq \alpha_c \varphi\},$$

where the infimum is taken in $D^\delta c$, and this element is completely join-prime in $D^\delta c$. By Lemma 5.5, $\mathfrak{o}_{k,x} \leq \alpha_c \varphi$ if and only if $x \in (\alpha_c \varphi)k$, and the latter is equal to $\alpha_{k,c}(\varphi)$ by Construction 5.1. The desired equality (11) now follows from the definition of $\text{tp}_{M_k, \alpha_k}(x)$. \square

5.4 Plentiful categories of models

By Proposition 5.6, each presheaf valued point $(M, \alpha): D \rightarrow \text{Sub}_{\text{Set } \mathcal{K}}$ induces a natural transformation $\alpha^\sharp: D^\delta \Rightarrow \text{Sub}_{\text{Set } \mathcal{K}} \circ M^{\text{op}}$. Using Lemma 5.11, we can characterise when α^\sharp is pointwise injective and surjective in terms of model-theoretic properties, namely, realisation of types and homogeneity, respectively. In this way, we justify these properties in a purely categorical way. This will also immediately yield a streamlined proof of Theorem 5.15, the main result of this section. This result is analogous to [5, Theorem 40] due to Coumans, as we will explain in more depth in Section 7.

Proposition 5.12. *The following are equivalent:*

- (i) *the natural transformation $\alpha^\sharp: D^\delta \rightarrow \text{Sub}_{\text{Set } \mathcal{K}} \circ M^{\text{op}}$ is a monomorphism;*
- (ii) *for any $c \in \text{ob } \mathcal{C}$, $\beta_c: \text{Sub}_{\text{Set } \mathcal{K}}(M_c) \rightarrow D^\delta c$ is surjective;*
- (iii) *for any $c \in \text{ob } \mathcal{C}$ and $t \in J^\infty D^\delta c$, there exist $k \in \text{ob } \mathcal{K}$ and $x \in M_k c$ such that $t = \bigwedge \text{tp}_{M_k, \alpha_k}(x)$.*

Proof. The equivalence of (i) and (ii) follows from the fact that α^\sharp is monomorphic if and only if it is component-wise injective, if and only if the left adjoints β_c are all surjective.

In view of Lemma 5.11, property (iii) asserts that, for any $c \in \text{ob } \mathcal{C}$, the restriction of β_c to a function $J^\infty \text{Sub}_{\text{Set } \mathcal{K}}(M_c) \rightarrow J^\infty D^\delta c$ is surjective. Since completely join-prime elements join-generate $D^\delta c$ and β_c preserves joins, this yields (iii) \Rightarrow (ii). For (ii) \Rightarrow (iii), if β_c is surjective, then in particular $J^\infty D^\delta c$ is contained in the image of β_c . Then, if $t = \beta_c(Y)$ for some subpresheaf $Y \leq M_c$, writing Y as in (10) and using that β_c is join-preserving yields some element (k, x) of Y such that $t = \beta_c(\mathfrak{o}_{k,x})$, as required. \square

Proposition 5.13. *The following are equivalent:*

- (i) *the natural transformation $\alpha^\sharp: D^\delta \rightarrow \text{Sub}_{\text{Set } \mathcal{K}} \circ M^{\text{op}}$ is an epimorphism;*
- (ii) *for any $c \in \text{ob } \mathcal{C}$, $\beta_c: \text{Sub}_{\text{Set } \mathcal{K}}(M_c) \rightarrow D^\delta c$ is an order embedding;*
- (iii) *for any $c \in \text{ob } \mathcal{C}$, $k, k' \in \text{ob } \mathcal{K}$, $x \in M_k c$ and $x' \in M_{k'} c$, if*

$$\text{tp}_{M_k, \alpha_k}(x) \subseteq \text{tp}_{M_{k'}, \alpha_{k'}}(x'),$$

then there exists a morphism $r: k \rightarrow k'$ in \mathcal{K} such that $((M_c)r)(x) = x'$.

Proof. In view of Lemma 5.11, property (iii) asserts that, for any $c \in \text{ob } \mathcal{C}$, if $\beta_c(\mathfrak{o}_{k',x'}) \leq \beta_c(\mathfrak{o}_{k,x})$, then $x' \in \mathfrak{o}_{k,x}k'$. By Remark 5.7, the latter is equivalent to $\mathfrak{o}_{k',x'} \leq \mathfrak{o}_{k,x}$. Thus, property (iii) is equivalent to property (ii). The equivalence of (i) and (ii) follows from the fact that α^\sharp is epimorphic if and only if each α_c^\sharp is surjective, if and only if each β_c is an order embedding. \square

Definition 5.14. (i) We say that (M, α) *strictly realises all prime types* if the equivalent conditions of Proposition 5.12 hold.

(ii) We say that (M, α) is *homogeneous* if the equivalent conditions of Proposition 5.13 hold.

(iii) We say that (M, α) is *plentiful* if it is homogeneous and strictly realises all prime types.

We remark on the appropriateness of our terminology in Remark 5.16. Proposition 5.12 and 5.13 together yield the following theorem.

Theorem 5.15 (cf. [5, Theorem 40]). *Let $(M, \alpha): D \rightarrow \text{Sub}_{\mathbf{Set}^\mathcal{K}}$ be a presheaf-valued point of a doctrine D . Then $\alpha^\sharp: D^\delta \Rightarrow \text{Sub}_{\mathbf{Set}^\mathcal{K}} \circ M^{\text{op}}$ is a natural isomorphism if and only if (M, α) is plentiful.*

Remark 5.16. (i) Let \mathbb{T} be a coherent theory, and let τ be a prime X -type (without parameters). Recall from Definition 4.5 we say that τ is realised by a tuple $\vec{m} \in \mathcal{M}^{|X|}$, in a model $\mathcal{M} \models \mathbb{T}$, if $\mathcal{M} \models \varphi(\vec{m})$ for each $\varphi \in \tau$, or in other words $\tau \subseteq \text{tp}_{\mathcal{M}}(\vec{m})$. In Definition 5.14, to say that \vec{m} *strictly realises* τ , we instead ask for an *equality* $\tau = \text{tp}_{\mathcal{M}}(\vec{m})$.

(ii) Our terminology of ‘homogeneous’ agrees with (one of) the standard use(s) of the term in model theory when in the following special case. Let \mathbb{T} be a classical theory and consider the automorphism group $\text{Aut}(\mathcal{M})$ of a single model $\mathcal{M} \models \mathbb{T}$, which can be viewed as a presheaf of points $\text{Aut}(\mathcal{M}) \rightarrow \text{Pt}(D_{\mathbb{T}})$ in the obvious way. Recall that \mathcal{M} is called *homogeneous* if every partial isomorphism $\vec{m} \cong \vec{m}'$ between finite subsets $\vec{m}, \vec{m}' \subseteq \mathcal{M}$ admits an extension to a total isomorphism of \mathcal{M} , as in the diagram

$$\begin{array}{ccc} \vec{m} & \xrightarrow{\cong} & \vec{m}' \\ \downarrow & & \downarrow \\ M & \xrightarrow{\cong} & M. \end{array}$$

There is a partial isomorphism $\vec{m} \cong \vec{m}'$ if and only if $\text{tp}_{\mathcal{M}}(\vec{m}) = \text{tp}_{\mathcal{M}}(\vec{m}')$. Thus, since the types of a classical theory are discretely ordered, $\text{Aut}(\mathcal{M})$ is homogeneous as in Definition 5.14 if and only if \mathcal{M} is homogeneous in the above sense.

Remark 5.17. It follows that if $(M, \alpha): D \rightarrow \text{Sub}_{[\mathcal{K}, \mathbf{Set}]}$ is plentiful, then it is also extensible in the sense of Definition 3.3 (cf. [5, Lemma 39]). In [5, §6], Coumans’s version of the statement of Theorem 5.15 asked for the further condition that each evaluation (M_k, α_k) is extensible (equivalently, if D is taken as the syntactic doctrine of a coherent theory \mathbb{T} , each evaluation is ω -saturated). We have observed that this assumption is not necessary; the extensibility of each (M_k, α_k) follows automatically from the other two conditions.

Proposition 5.18. *Let $(M, \alpha): D \rightarrow \text{Sub}_{\mathbf{Set}^\mathcal{K}}$ be a plentiful presheaf-valued point of D . Then for each object $k \in \mathcal{K}$, the associated point $(M_k, \alpha_k): D \rightarrow \mathcal{P}$ (see Construction 5.1) is extensible.*

In particular, if $\mathcal{K} \subseteq \text{Mod}(\mathbb{T})$ is a subcategory of models for a coherent theory \mathbb{T} that strictly realises all prime types and is homogeneous, then each model in \mathcal{K} is ω -saturated.

Proof. Let $\text{ev}_k: \text{Sub}_{\mathbf{Set}^\mathcal{K}} \rightarrow \mathcal{P}$ be the morphism of doctrines that evaluates subpresheaves at the object $k \in \mathcal{K}$, i.e., the underlying functor $\mathbf{Set}^\mathcal{K} \rightarrow \mathbf{Set}$ of ev_k sends a presheaf X to the set Xk , and a subpresheaf $Y \subseteq X$ to the subset $Yk \subseteq Xk$. Thus, $\text{ev}_k \circ (M, \alpha) = (M_k, \alpha_k)$. Since arbitrary meets, joins and images are computed pointwise in $\mathbf{Set}^\mathcal{K}$, it follows that $\text{ev}_k: \text{Sub}_{\mathbf{Set}^\mathcal{K}} \rightarrow \mathcal{P}$ is a complete morphism of DL^+ -doctrines.

As α^\sharp is a natural isomorphism, by Remark 2.5, $(M, \alpha^\sharp): D^\delta \rightarrow \text{Sub}_{\mathbf{Set}^\mathcal{K}}$ is a morphism of DL^+ doctrines. By taking the composite $\text{ev}_k \circ (M, \alpha^\sharp)$, we obtain our desired extension of (M_k, α_k) . \square

Example 5.19 (Decidable objects). Consider the theory \mathbb{O}_{\neq} of *decidable objects*, i.e., the theory with one binary relation symbol \neq and the axioms

$$(x = y) \wedge (x \neq y) \vdash_{x,y} \perp \quad \text{and} \quad \top \vdash_{x,y} (x = y) \vee (x \neq y).$$

The category of models of this theory is $\mathbf{Set}_{\hookrightarrow}$, the category of sets and injections. We will show in this example that there are (full) subcategories $\mathcal{K} \subseteq \mathbf{Set}_{\hookrightarrow}$ that strictly realise all prime types but are not homogeneous, that are homogeneous but do not strictly realise all prime types, and conclude with a complete characterisation of the plentiful categories of models for \mathbb{O}_{\neq} .

Let us first describe the prime X -types of \mathbb{O}_{\neq} , i.e., the prime filters of $D_{\top}(X)$. For each natural number n , we use the shorthand E_n to denote the regular formula expressing that there are at least n elements in the model,

$$E_n := \exists y_1 \dots \exists y_n \left(\bigwedge_{k,\ell \leq n, k \neq \ell} y_k \neq y_\ell \right).$$

Using [20, Lemma D1.3.8] and the axioms of \mathbb{O}_{\neq} , it follows that every coherent formula in a context X is equivalent, modulo the theory \mathbb{O}_{\neq} , to a disjunction of formulae of the form

$$\left(\bigwedge_{(x_i, x_{i'}) \in I} x_i = x_{i'} \right) \wedge \left(\bigwedge_{(x_j, x_{j'}) \in J} x_j = x_{j'} \right) \wedge E_n,$$

for some $n \in \mathbb{N}$ and subsets $I, J \subseteq X^2$ (where I, J can moreover be chosen to be disjoint). For any equivalence relation \sim on X , we write

$$S_{\sim} := \left(\bigwedge_{x,y \in X, x \sim y} x = y \right) \wedge \left(\bigwedge_{x,y \in X, x \not\sim y} x \neq y \right).$$

The axioms of the theory \mathbb{O}_{\neq} imply that a prime X -type contains exactly one of $(x = y)$ and $(x \neq y)$, for each $x, y \in X$. Thus, a prime type τ contains the formula S_{\sim} for some equivalence relation \sim of the context X . By our above analysis of the coherent formulae of \mathbb{O}_{\neq} , it follows that τ can be of one of the two species:

$$\tau_{(\sim, n)} = \uparrow(S_{\sim} \wedge E_n) \quad \text{or} \quad \tau_{(\sim, \infty)} = \bigcup_{n \in \mathbb{N}} \uparrow(S_{\sim} \wedge E_n),$$

where, for the former species, we require that n is greater than or equal to the number of equivalence classes of \sim . That is to say, a prime X -type of \mathbb{O}_{\neq} expresses which variables in X are equal, which variables are not, and expresses that the underlying cardinality of the model has at least n elements or is infinite. Note that there is an inclusion of types $\tau_{(\sim, a)} \subseteq \tau_{(\sim', b)}$ if and only if the equivalence relations \sim and \sim' are identical and $a \leq b$, where a, b denote elements of $\mathbb{N} \cup \{\infty\}$ with the usual ordering.

Therefore, for a subcategory of models $\mathcal{K} \subseteq \mathbf{Set}_{\hookrightarrow}$ to *strictly* realise all prime types, \mathcal{K} must contain a finite set of cardinality n , for each $n \in \mathbb{N}$, and at least one infinite set.

Now let us consider homogeneity. Let $Y, Y' \in \mathcal{K} \subseteq \mathbf{Set}_{\hookrightarrow}$ be sets that are included in the chosen subcategory \mathcal{K} , and suppose that the type of some tuple $(y_1, \dots, y_n) \in Y^n$ is included in the type of $(y'_1, \dots, y'_n) \in Y'^n$, which, as we have just seen, entails that, for any $1 \leq i, j \leq n$,

- (i) $y_i = y_j$ if, and only if, $y'_i = y'_j$, and
- (ii) either Y is finite and $|Y'| \geq |Y|$, or both Y and Y' are infinite.

If $|Y| \leq |Y'|$, then there exists an injection $Y \hookrightarrow Y'$ that sends y_i to y'_i for each $1 \leq i \leq n$. However, if $|Y| > |Y'|$, then there can be no injection $Y \hookrightarrow Y'$. Thus, as

soon as \mathcal{K} contains two infinite sets of differing cardinality, \mathcal{K} is not homogeneous.

Let $\mathbf{Set}_{\rightarrow}^{\leq \kappa}$ (respectively, $\mathbf{Set}_{\rightarrow}^{< \kappa}$) denote the full subcategory of $\mathbf{Set}_{\rightarrow}$ consisting of all sets of cardinality less than or equal to κ (resp., strictly less than κ). By our above analysis, we conclude the following: the subcategory $\mathbf{Set}_{\rightarrow}^{\leq \omega}$ of finite sets is a homogeneous category of models for 0_{\neq} , but does not realise all prime types; $\mathbf{Set}_{\rightarrow}^{\leq 2^{\omega}}$ realises all types but is not homogeneous; the category $\mathbf{Set}_{\rightarrow}^{\leq \omega}$ of all countable sets is, by contrast, homogeneous and realises all types. Indeed, a full subcategory $\mathcal{K} \subseteq \mathbf{Set}_{\rightarrow}$ is plentiful if and only if \mathcal{K} contains sets of all finite cardinalities and sets of exactly one infinite cardinality.

Example 5.20 (Reconstructing ω -categorical theories). Recall that a classical (single-sorted) theory \mathbb{T} is said to be ω -categorical if one of the following equivalent conditions is satisfied:

- (i) for each natural number n , there are finitely many n -types;
- (ii) for each n , there are finitely many formulae with n free variables up to \mathbb{T} -provable equivalence, or equivalently the associated doctrine $D_{\mathbb{T}}$ factors through finite Boolean algebras;
- (iii) each type is *isolated*, i.e., each ultrafilter of $D_{\mathbb{T}}(n)$ is principal;
- (iv) the theory \mathbb{T} has, up to isomorphism, a unique countable model.

(Of course, the equivalence of conditions (i), (ii) and (iii) is a direct consequence of Stone duality for Boolean algebras, while it is the final condition (iv) that gives ω -categorical theories their name.) There is a well-established sense in which an ω -categorical (classical) theory can be reconstructed from the automorphism group of any countable model [1, §1], and more generally from any homogeneous model [3].

The salient steps of the assertion that an ω -categorical theory can be reconstructed from the automorphism group of homogeneous model can essentially be deduced from Theorem 5.15. Given a classical, ω -coherent theory \mathbb{T} , since $D_{\mathbb{T}}$ factors through finite Boolean algebras, we have that $D_{\mathbb{T}} = D_{\mathbb{T}}^{\delta}$, and as \mathbb{T} is a complete theory with isolated types, any model (strictly) realises all types. Thus, given a model $\mathcal{M} \models \mathbb{T}$, viewed as a functor $M: \text{Aut}(\mathcal{M}) \rightarrow \text{Pt}(D_{\mathbb{T}})$, the associated interpretation transformation $\llbracket - \rrbracket_M: D_{\mathbb{T}} \Rightarrow \text{Sub}_{\text{Aut}(\mathcal{M})} \circ M^{\text{op}}$ is a natural isomorphism if and only if \mathcal{M} is homogeneous. In particular, the unique countable model of an ω -categorical theory is homogeneous ([18, §10]), and so \mathbb{T} can be recovered from the automorphism group of its (unique) countable model (cf. [1, §1]). We mention, for interest, that recent work also explores homogeneity in the context of endomorphism monoids of models [35].

6 Recovering the canonical extension for ω -stable theories

In this final section, as an application of Theorem 5.15, we will show that the canonical extension $D_{\mathbb{T}}^{\delta}$ of the syntactic doctrine for an ω -stable theory can be recovered from a subcategory of models of \mathbb{T} equipped only with knowledge of the underlying sets of these models (but not, for instance, knowledge of the definable subsets, etc.). To express “knowledge of the underlying sets”, we employ the following formalism.

Definition 6.1. Let \mathcal{C} and \mathcal{D} be categories equipped with functors $A: \mathcal{C} \rightarrow \mathcal{E}$ and $B: \mathcal{D} \rightarrow \mathcal{E}$. We say that \mathcal{C} and \mathcal{D} are *equivalent over* \mathcal{E} , and write $\mathcal{C} \simeq_{\mathcal{E}} \mathcal{D}$, if there are functors $F: \mathcal{C} \rightleftarrows \mathcal{D}: G$ witnessing an equivalence $\mathcal{C} \simeq \mathcal{D}$ which moreover commute

with the functors A and B , as in the diagram:

$$\begin{array}{ccc} \mathcal{C} & \xrightleftharpoons[F]{G} & \mathcal{D} \\ & \searrow A & \swarrow B \\ & & \mathcal{E}. \end{array}$$

In other words, A and B are equivalent as objects of the slice category \mathbf{CAT}/\mathcal{E} .

In our setting, $\mathcal{C} \subseteq \mathbf{Str}(\Sigma)$ and $\mathcal{D} \subseteq \mathbf{Str}(\Sigma')$ will be subcategories of structures, where Σ, Σ' are a pair of single-sorted signatures, and so are both equipped with forgetful functors $U: \mathcal{C} \rightarrow \mathbf{Set}, U': \mathcal{D} \rightarrow \mathbf{Set}$ that send a structure to its underlying set. Thus, \mathcal{C}, \mathcal{D} are equivalent over \mathbf{Set} if, in addition to there being an equivalence $\mathcal{C} \simeq \mathcal{D}$, this equivalence preserves the underlying set of each structure.

We are motivated to consider equivalences of categories over a third category since it is precisely the notion of equivalence that arises when considering isomorphisms of doctrines.

Lemma 6.2. *Let $D, D': \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be coherent doctrines (respectively, DL^+ -doctrines) over the same base category. If $D \cong D'$, then $\text{Pt}(D) \simeq_{\mathbf{Lex}[\mathcal{C}, \mathbf{Set}]} \text{Pt}(D')$ (resp., $\text{Pt}^+(D) \simeq_{\mathbf{Lex}[\mathcal{C}, \mathbf{Set}]} \text{Pt}^+(D')$).*

Proof. The category of points $\text{Pt}(D)$ of D lives over the category $\mathbf{Lex}[\mathcal{C}, \mathbf{Set}]$ by sending a morphism $(M, \alpha): D \rightarrow \mathcal{P}$ to the functor $M: \mathcal{C} \rightarrow \mathbf{Set}$. Using $\beta: D \Rightarrow D'$ to denote the natural isomorphism, the equivalence $\text{Pt}(D) \simeq_{\mathbf{Lex}[\mathcal{C}, \mathbf{Set}]} \text{Pt}(D')$ is witnessed by sending a morphism $(M, \alpha) \in \text{Pt}(D)$ to $(M, \alpha \circ \beta) \in \text{Pt}(D')$, and in reverse sending $(M, \gamma) \in \text{Pt}(D')$ to $(M, \gamma \circ \beta^{-1}) \in \text{Pt}(D)$. \square

Definition 6.3. Recall that a model \mathcal{M} is ω -saturated if, for any finite X, Y , every consistent Y -type with X -parameters in \mathcal{M} is realised. We say that \mathcal{M} is *countable and saturated* if the underlying set of \mathcal{M} is countable and \mathcal{M} is ω -saturated. We use $\mathbf{Sat}_\omega(\mathbb{T})$ to denote the category of countable, saturated models of \mathbb{T} and model homomorphisms between these.

Lemma 6.4 (cf. [18, Lemma 10.1.3]). *Let \mathcal{M} and \mathcal{M}' be models of a coherent theory \mathbb{T} with $\vec{m} \in M^n$ and $\vec{m}' \in M'^n$ such that $\text{tp}_{\mathcal{M}}(\vec{m}) \subseteq \text{tp}_{\mathcal{M}'}(\vec{m}')$. Suppose also that \mathcal{M} is countable and \mathcal{M}' is ω -saturated. Then there exists a homomorphism of models $f: \mathcal{M} \rightarrow \mathcal{M}'$ sending \vec{m} pointwise to \vec{m}' .*

Proof. Since M is countable, fix an enumeration of $M = (m_1, m_2, \dots)$ where the initial segment (m_1, \dots, m_n) is the given tuple $\vec{m} \in M^n$. We will construct the homomorphism $f: \mathcal{M} \rightarrow \mathcal{M}'$ inductively by building partial structure homomorphisms $f_k: \mathcal{M} \rightarrow \mathcal{M}'$ whose domain at each stage is given by (m_1, \dots, m_k) . As a base case, take $f_n: \mathcal{M} \rightarrow \mathcal{M}'$ to be the partial structure homomorphism that sends $\vec{m} = (m_1, \dots, m_n)$ pointwise to $\vec{m}' = (m'_1, \dots, m'_n) \subseteq M'$. This is indeed a partial structure homomorphism since $\text{tp}(\vec{m}) \subseteq \text{tp}(\vec{m}')$.

Suppose we have built $f_k: (m_1, \dots, m_k) \rightarrow \mathcal{M}'$ for $k \geq n$. We obtain f_{k+1} as follows. Firstly, we set $f_{k+1}(m_i) = f_k(m_i)$ for $i \leq k$. Turning to $f_{k+1}(m_{k+1})$, let $t(x_1, \dots, x_k, x_{k+1})$ denote the type of $(m_1, \dots, m_{k+1}) \subseteq M^{k+1}$. The type with (finitely many) parameters $t(f_k(m_1), \dots, f_k(m_k), x_{k+1})$ is consistent with \mathcal{M}' since, for each $\varphi \in t$, we have that $\exists x_{k+1} \varphi \in \text{tp}(m_1, \dots, m_k)$, from which we deduce that $\mathcal{M} \models \exists x_{k+1} \varphi(m_1, \dots, m_k)$ and so $\mathcal{M}' \models \exists x_{k+1} \varphi(f_k(m_1), \dots, f_k(m_k))$ by the fact that f_k is a structure homomorphism. Hence, since \mathcal{M}' is ω -saturated, there is some $m'_{k+1} \in M'$ that realises $t(f_k(m_1), \dots, f_k(m_k), x_{k+1})$, and we set $f_{k+1}(m_{k+1})$ as m'_{k+1} . The function $f_{k+1}: (m_1, \dots, m_{k+1}) \rightarrow \mathcal{M}'$ is a partial structure homomorphism since, by design, $\text{tp}(m_1, \dots, m_{k+1}) \subseteq \text{tp}(f(m_1), \dots, f(m_{k+1}))$.

The structure morphism $f: \mathcal{M} \rightarrow \mathcal{M}'$ is defined as the union $\bigcup_{k \leq \omega} f_k$, which is total since (m_1, m_2, \dots) enumerated the whole of M . This is clearly a structure homomorphism, i.e., if $\mathcal{M} \models \varphi(\vec{m}'')$ then $\mathcal{M}' \models \varphi(f(\vec{m}''))$, since we can always restrict to the partial structure homomorphism f_k for some k large enough that each $m'' \in \vec{m}''$ is contained in (m_1, \dots, m_k) . \square

Corollary 6.5. *The category $\mathbf{Sat}_\omega(\mathbb{T})$ of countable, saturated models of \mathbb{T} is a homogeneous category of models in the sense of Definition 5.14.*

Theorem 6.6. *Let \mathbb{T} and \mathbb{T}' be coherent theories. If every prime type of \mathbb{T} (respectively, \mathbb{T}') is strictly realised in a countable, saturated model, then the canonical extensions of their associated syntactic doctrines are isomorphic if and only if their countable, saturated models are equivalent over \mathbf{Set} , i.e.,*

$$D_{\mathbb{T}}^\delta \cong D_{\mathbb{T}'}^\delta \iff \mathbf{Sat}_\omega(\mathbb{T}) \simeq_{\mathbf{Set}} \mathbf{Sat}_\omega(\mathbb{T}').$$

Proof. Let us write $\mathcal{K} := \mathbf{Sat}_\omega(\mathbb{T})$ and $\mathcal{K}' := \mathbf{Sat}_\omega(\mathbb{T}')$. Suppose first that $\mathcal{K} \simeq_{\mathbf{Set}} \mathcal{K}'$. Then it is easily seen that there is a natural isomorphism $\mathbf{Sub}_{\mathbf{Set}\mathcal{K}} \circ U^{\text{op}} \cong \mathbf{Sub}_{\mathbf{Set}\mathcal{K}'} \circ U'^{\text{op}}$, where U and U' denote the respective forgetful functors $U: \mathcal{K} \rightarrow \mathbf{Set}$, $U': \mathcal{K}' \rightarrow \mathbf{Set}$. By assumption, \mathcal{K} strictly realises all prime types, and by Corollary 6.5 it is also homogeneous (and similarly for \mathcal{K}'). Thus, by Theorem 5.15, there is a chain of natural isomorphism

$$D_{\mathbb{T}}^\delta \cong \mathbf{Sub}_{\mathbf{Set}\mathcal{K}} \circ U^{\text{op}} \cong \mathbf{Sub}_{\mathbf{Set}\mathcal{K}'} \circ U'^{\text{op}} \cong D_{\mathbb{T}'}^\delta.$$

Conversely, suppose that $D_{\mathbb{T}}^\delta \cong D_{\mathbb{T}'}^\delta$. By Lemma 6.2, we have

$$\mathbf{Pt}^+(D_{\mathbb{T}}) \simeq_{\mathbf{Lex}[\mathbf{FinSet}, \mathbf{Set}]} \mathbf{Pt}^+(D_{\mathbb{T}'}).$$

Since the equivalence $\mathbf{Pt}^+(D_{\mathbb{T}}) \simeq \mathbf{Pt}^+(D_{\mathbb{T}'})$ occurs over $\mathbf{Lex}[\mathbf{FinSet}, \mathbf{Set}]$, we can restrict this equivalence to $\mathcal{C} \simeq_{\mathbf{Lex}[\mathbf{FinSet}, \mathbf{Set}]} \mathcal{C}'$, where $\mathcal{C} \subseteq \mathbf{Pt}^+(D_{\mathbb{T}})$ is the full subcategory of those complete points (M, α) where M takes values in countable sets, and similarly for $\mathcal{C}' \subseteq \mathbf{Pt}^+(D_{\mathbb{T}'})$. Now recall, from Theorem 4.8, that an object $(M, \alpha) \in \mathcal{C}$ is just the datum of a countable, ω -saturated model of \mathbb{T} , and moreover this equivalence commutes with the canonical forgetful functors in the sense that there is a commuting square

$$\begin{array}{ccc} \mathbf{Pt}^+(D_{\mathbb{T}}) & \xrightarrow{\sim} & \mathbf{Sat}_\omega(\mathbb{T}) \\ \downarrow & & \downarrow \\ \mathbf{Lex}[\mathbf{FinSet}, \mathbf{Set}] & \xrightarrow{\sim} & \mathbf{Set}. \end{array}$$

The same analysis holds for \mathbb{T}' , and so from $\mathcal{C} \simeq_{\mathbf{Lex}[\mathbf{FinSet}, \mathbf{Set}]} \mathcal{C}'$ we deduce that $\mathbf{Sat}_\omega(\mathbb{T}) \simeq_{\mathbf{Set}} \mathbf{Sat}_\omega(\mathbb{T}')$ as desired. \square

Thus, we obtain a reconstruction result for the canonical extension of the syntactic doctrine associated with any theory that has enough countable, saturated models. The existence of countable, saturated models is, in general, a hard problem, but a sufficient condition is found in the model-theoretic notion of ω -stability, which imposes a bound on the complexity of models as parameters are introduced.

Definition 6.7. A coherent theory \mathbb{T} is said to be ω -stable if, given any model $\mathcal{M} \models \mathbb{T}$ and a countable subset $A \subseteq M$, there are countably many prime types with parameters in A .

Proposition 6.8 ([18], Theorem 10.2.7). *If \mathbb{T} is ω -stable, every prime type of a coherent theory is strictly realised in a countable, saturated model.*

While the reader is directed to [18] for a detailed proof in the classical setting (from which a version for coherent logic can be derived), we give some intuition about why ω -stable theories have enough countable, saturated models: suppose we started with a

countable model \mathcal{M}_0 . If there are types with finitely many parameters from \mathcal{M}_0 that are not realised in \mathcal{M}_0 , we find, using standard model-theoretic techniques, a larger, countable model $\mathcal{M}_1 \supseteq \mathcal{M}_0$ that does realise those types. If \mathcal{M}_1 is not ω -saturated, we repeat to find $\mathcal{M}_2 \supseteq \mathcal{M}_1$, and so on. The bound on complexity of models imposed by ω -stability ensures that after ω -many steps, in the limit $\bigcup_{n < \omega} \mathcal{M}_n$, the process terminates and we are left with a countable, saturated model as desired.

Thus, by specialising Theorem 6.6 we arrive at our reconstruction result for ω -stable theories.

Corollary 6.9. *Let \mathbb{T} and \mathbb{T}' be ω -stable coherent theories. The canonical extensions of their associated syntactic doctrines are isomorphic if and only if their countable, saturated models are equivalent over \mathbf{Set} .*

That is to say, for ω -stable theories, the data of their (prime) types are entirely determined by the underlying sets and functions of their countable, saturated models and homomorphisms thereof. This hints that the correct notion of *Galois types* (cf., e.g. [16, §6]), in the case of a coherent theory, ought to involve homomorphisms of models, not only isomorphisms as in the classical case.

Example 6.10. The assumption that \mathbb{T} is ω -stable is not necessary in order to have enough countable, saturated models. For instance, the theory of the Rado (or random) graph gives a counterexample. To encode this theory in coherent logic, we take three binary relation symbols $\neq, \sim, \not\sim$, and the axioms

$$\begin{aligned} (x = y) \wedge (x \neq y) &\vdash \perp, & \mathbb{T} &\vdash (x = y) \vee (x \neq y), \\ (x \sim y) \wedge (x \not\sim y) &\vdash \perp, & \mathbb{T} &\vdash (x \sim y) \vee (x \not\sim y), \end{aligned}$$

expressing that $=, \neq$ and $\sim, \not\sim$ are complements, $x \sim y \vdash y \sim x$ making \sim symmetric, i.e., the edge relation of a graph, and for all $n, m \in \mathbb{N}$ the axiom

$$\mathbb{T} \vdash \exists y. \left(\bigwedge_{i \leq n+m} y \neq x_i \wedge \bigwedge_{i \leq n} y \sim x_i \wedge \bigwedge_{n < j \leq m} y \not\sim x_j \right)$$

expressing that for any two finite subsets of vertices, there is another distinct vertex that is joined to all the vertices in the first set, but none of the latter. This theory is not ω -stable ([30, Example 4]) but it is an ω -categorical theory ([18, §7.4]), and so every model is ω -saturated (via Example 4.10).

7 A tale of two topoi

Motivating Coumans's work [4, 5] on the canonical extension of doctrines was the connection to topos theory, which we have so far mostly omitted. In this concluding section, we describe the topos-theoretic counterparts to the results contained in this paper, which in particular lets us compare our results in more depth to those by Makkai [28] and Coumans. We will therefore assume some knowledge of topos theory for this discussion, as can be found in [27] or [20].

Given a coherent theory \mathbb{T} , there are two topoi⁵ we can naturally associate with the theory \mathbb{T} : its classifying topos and its topos of types. We begin by briefly recalling these constructions.

Classifying topos. The *classifying topos* of a coherent theory \mathbb{T} , which we denote by $\mathcal{E}_{\mathbb{T}}$, is characterised by the property that, for any topos \mathcal{F} , geometric morphisms $\mathcal{F} \rightarrow \mathcal{E}_{\mathbb{T}}$ correspond to models of \mathbb{T} internal to \mathcal{F} . Similarly, each coherent $D: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ has a 'classifying topos' too. Denote by $\mathcal{E}_D^{\text{coh}}$ the topos of sheaves on the site consisting

⁵Throughout this section, 'topos' is used as shorthand for 'Grothendieck topos'.

of the category $\mathcal{A}(D)$ equipped with the coherent topology, where \mathcal{A} is the ‘syntactic category’ functor from Remark 5.2. Then the topos $\mathcal{E}_D^{\text{coh}}$ classifies D , in the sense that, for any topos \mathcal{F} , we have an equivalence $\mathbf{Geom}(\mathcal{F}, \mathcal{E}_D^{\text{coh}}) \simeq \mathbf{CohDoc}(D, \text{Sub}_{\mathcal{F}})$. In particular, for \mathbb{T} a coherent theory, the classifying topos $\mathcal{E}_{D_{\mathbb{T}}}^{\text{coh}}$ of the syntactic doctrine $D_{\mathbb{T}}$ is the usual classifying topos of \mathbb{T} .

We have used the superscript *coh* to emphasise the fact that $\mathcal{E}_D^{\text{coh}}$ classifies the *coherent* morphisms $D \rightarrow \text{Sub}_{\mathcal{F}}$. If the coherent doctrine D takes values in the subcategory of frames $\mathbf{Frm} \subseteq \mathbf{DL}$, we may also wish to classify morphisms of coherent doctrines $(M, \alpha): D \rightarrow \text{Sub}_{\mathcal{F}}$ where $\alpha: D \Rightarrow \text{Sub}_{\mathcal{F}} \circ M^{\text{op}}$ is a natural transformation of \mathbf{Frm} -valued functors. We call D a *geometric doctrine* and (M, α) a morphism of *geometric doctrines* ([38, Definition III.43]). In this case, there is a topos $\mathcal{E}_D^{\text{geom}}$ for which $\mathbf{Geom}(\mathcal{F}, \mathcal{E}_D^{\text{geom}}) \simeq \mathbf{GeomDoc}(D, \text{Sub}_{\mathcal{F}})$, which can be obtained by taking sheaves on the geometric category $\mathcal{A}(D)$ with the geometric topology. Note that, in particular, the canonical extension D^δ of any coherent doctrine is a geometric doctrine.

Topoi of types. In addition to the classifying topos, a coherent theory \mathbb{T} also has a corresponding *topos of types* $\mathcal{T}_{\mathbb{T}}$, as introduced by Makkai [28]. A key contribution of Coumans ([5, Proposition 24]) was to show that there is an equivalence

$$\mathcal{T}_{\mathbb{T}} \simeq \mathcal{E}_{D_{\mathbb{T}}}^{\text{geom}}.$$

Generalising, we write \mathcal{T}_D for $\mathcal{E}_{D^\delta}^{\text{geom}}$ and refer to it as the topos of types for a coherent doctrine D . Thus, for any topos \mathcal{F} , geometric morphisms $\mathcal{F} \rightarrow \mathcal{T}_D$ correspond to morphisms of geometric doctrines $D^\delta \rightarrow \text{Sub}_{\mathcal{F}}$.

If \mathcal{F} is a topos such that each subobject lattice of \mathcal{F} is a DL^+ , i.e., $\text{Sub}_{\mathcal{F}}$ is a DL^+ -doctrine, then we can restrict $\mathbf{Geom}(\mathcal{F}, \mathcal{T}_D) \simeq \mathbf{Geom}(D^\delta, \text{Sub}_{\mathcal{F}})$ to an equivalence

$$\mathbf{Geom}^+(\mathcal{F}, \mathcal{T}_D) \simeq \mathbf{CohDoc}^+(D^\delta, \text{Sub}_{\mathcal{F}}),$$

where $\mathbf{Geom}^+(\mathcal{F}, \mathcal{T}_D)$ is defined as the class of geometric morphisms $f: \mathcal{F} \rightarrow \mathcal{T}_D$ such that the inverse image f^* preserves arbitrary meets of subobjects. This is the original universal property of \mathcal{T}_D that Makkai used to characterise the topos of types ([28, Theorem 1.1]). Thus, our Theorem 4.8 is equivalent to the following corollary:

Corollary 7.1. *For \mathbb{T} a coherent theory, $\mathbf{Geom}^+(\mathbf{Set}, \mathcal{T}_{\mathbb{T}})$ is equivalent to the category of (set-based) ω -saturated models of \mathbb{T} .*

Note that in general $\mathcal{E}_{D^\delta}^{\text{geom}} \not\cong \mathcal{E}_{D^\delta}^{\text{coh}}$, and also $\mathbf{Geom}^+(\mathcal{F}, \mathcal{T}_D)$ is strictly contained in $\mathbf{Geom}(\mathcal{F}, \mathcal{T}_D)$ – we must specify with respect to which syntax we are classifying D^δ !

Internal locales. Here is another perspective on the topos of types: the geometric doctrine $D^\delta: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}^+$ describes an *internal frame*, or equivalently an *internal locale*, of the presheaf topos $\mathbf{Set}^{\mathcal{C}^{\text{op}}}$ (see [21, §VI.2]). As an internal locale, we can take the topos $\mathbf{Sh}_{\mathbf{Set}^{\mathcal{C}^{\text{op}}}}(D^\delta)$ of *internal sheaves* on D^δ , generalising the usual notion of sheaves on a locale. The topos $\mathbf{Sh}_{\mathbf{Set}^{\mathcal{C}^{\text{op}}}}(D^\delta)$ is equivalent to \mathcal{T}_D (compare the constructions of $\mathbf{Sh}_{\mathbf{Set}^{\mathcal{C}^{\text{op}}}}(D)$ in [31, Example 1.3(c)] and $\mathcal{E}_{D^\delta}^{\text{geom}}$ in [38, §III]). As a consequence, there is a canonical *localic* geometric morphism $v: \mathcal{T}_D \simeq \mathbf{Sh}_{\mathbf{Set}^{\mathcal{C}^{\text{op}}}}(D^\delta) \rightarrow \mathbf{Set}^{\mathcal{C}^{\text{op}}}$ (see [19, Lemma 1.2]).

In fact, for any coherent doctrine $D: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$, there is a corresponding canonical localic geometric morphism $u_D: \mathcal{E}_D^{\text{coh}} \rightarrow \mathbf{Set}^{\mathcal{C}^{\text{op}}}$ ([38, §IV.1]).

Hyperconnected-localic factorisations. Every geometric morphism $f: \mathcal{F} \rightarrow \mathcal{G}$ between toposes can be factorised uniquely up to equivalence as

$$\mathcal{F} \xrightarrow{f_{\text{hyp}}} \mathcal{L}(f) \xrightarrow{f_{\text{loc}}} \mathcal{G},$$

where f_{hyp} is *hyperconnected* and f_{loc} is *localic* (and thus the factoring topos $\mathcal{L}(f)$ is the topos of sheaves on an internal locale of \mathcal{G} ; see [20, §A.4.6.5]).

Let $D: \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a coherent doctrine, \mathcal{K} a small category, and let $p: \mathbf{Set}^{\mathcal{K}} \rightarrow \mathcal{E}_D^{\text{coh}}$ be a geometric morphism. By the characteristic property of $\mathcal{E}_D^{\text{coh}}$, p is uniquely induced by some morphism of coherent doctrines $(M, \alpha): D \rightarrow \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$. Consider the hyperconnected-localic factorisation $(p_{\text{hyp}}, p_{\text{loc}})$. Since the geometric morphism $u_D: \mathcal{E}_D^{\text{coh}} \rightarrow \mathbf{Set}^{\mathcal{C}^{\text{op}}}$ is localic, the factorisation of the composite $u_D \circ p$ is given by $((u_D \circ p)_{\text{hyp}}, (u_D \circ p)_{\text{loc}}) = (p_{\text{hyp}}, u_D \circ p_{\text{loc}})$, as in the following diagram:

$$\begin{array}{ccc} \mathbf{Set}^{\mathcal{K}} & \xrightarrow{p} & \mathcal{E}_D^{\text{coh}} \xrightarrow{u_D} \mathbf{Set}^{\mathcal{C}^{\text{op}}} \\ & \searrow p_{\text{hyp}} & \nearrow p_{\text{loc}} \\ & & \mathcal{L}(p) \end{array}$$

The internal locale of $\mathbf{Set}^{\mathcal{C}^{\text{op}}}$ corresponding to the localic morphism $u_D \circ p_{\text{loc}}$ is easy to calculate: it is given by the (geometric) doctrine $\text{Sub}_{\mathcal{K}} \circ M^{\text{op}}: \mathcal{C}^{\text{op}} \rightarrow \mathbf{Frm}$ (see, for instance, [39, §3]). Thus, since there are equivalences $\mathcal{T}_D \simeq \mathbf{Sh}_{\mathbf{Set}^{\mathcal{C}^{\text{op}}}}(D^\delta)$ and $\mathcal{L}(p) \simeq \mathbf{Sh}_{\mathbf{Set}^{\mathcal{C}^{\text{op}}}}(\text{Sub}_{\mathcal{K}} \circ M^{\text{op}})$, Theorem 5.15 can also be expressed as follows:

Corollary 7.2. *The morphism $(M, \alpha): D \rightarrow \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ is plentiful if and only if there is an equivalence of topoi $\mathcal{T}_D \simeq \mathcal{L}(p)$.*

The left-to-right direction of this corollary is how Coumans originally expressed [5, Theorem 40].

Comparing internal locales. We end by describing the topos-theoretic intuition behind the constructions in Section 6.

Let $D, D': \mathcal{C}^{\text{op}} \rightarrow \mathbf{DL}$ be a pair of coherent doctrines over the same base category, and let $\mathcal{K} \subseteq \text{Pt}(D)$ and $\mathcal{K}' \subseteq \text{Pt}(D')$ be subcategories of points, which we recall from Remark 5.2 can be alternatively viewed as a pair $(M, \alpha): D \rightarrow \text{Sub}_{\mathbf{Set}^{\mathcal{K}}}$ and $(M', \alpha'): D' \rightarrow \text{Sub}_{\mathbf{Set}^{\mathcal{K}'}}$ of presheaf-valued points, or, as above, as a pair of geometric morphisms $p: \mathbf{Set}^{\mathcal{K}} \rightarrow \mathcal{E}_D^{\text{coh}}$ and $p': \mathbf{Set}^{\mathcal{K}'} \rightarrow \mathcal{E}_{D'}^{\text{coh}}$. Now suppose there is an equivalence $\mathcal{K} \simeq_{\mathbf{Lex}[\mathcal{C}, \mathbf{Set}]} \mathcal{K}'$ over $\mathbf{Lex}[\mathcal{C}, \mathbf{Set}]$, where the functor $\mathcal{K} \rightarrow \mathbf{Lex}[\mathcal{C}, \mathbf{Set}]$ is the restriction of the canonical forgetful functor $\text{Pt}(D) \rightarrow \mathbf{Lex}[\mathcal{C}, \mathbf{Set}]$, and similarly for \mathcal{K}' . It follows that there is an equivalence of topoi $\mathbf{Set}^{\mathcal{K}} \simeq_{\mathbf{Set}^{\mathcal{C}^{\text{op}}}} \mathbf{Set}^{\mathcal{K}'}$ over $\mathbf{Set}^{\mathcal{C}^{\text{op}}}$, i.e., there is an equivalence $\mathbf{Set}^{\mathcal{K}} \simeq \mathbf{Set}^{\mathcal{K}'}$ that commutes with the geometric morphisms $u_D \circ p$ and $u_{D'} \circ p'$, as in the diagram

$$\begin{array}{ccc} \mathbf{Set}^{\mathcal{K}} \simeq \mathbf{Set}^{\mathcal{K}'} & & \\ p \downarrow & & \downarrow p' \\ \mathcal{E}_D^{\text{coh}} & & \mathcal{E}_{D'}^{\text{coh}} \\ u_D \swarrow & & \swarrow u_{D'} \\ & \mathbf{Set}^{\mathcal{C}^{\text{op}}} & \end{array}$$

Hence, by the above analysis, if there is an equivalence $\mathcal{K} \simeq_{\mathbf{Lex}[\mathcal{C}, \mathbf{Set}]} \mathcal{K}'$, then there is a chain of equivalences of topoi $\mathcal{L}(p) \simeq \mathcal{L}(u_D \circ p) \simeq \mathcal{L}(u_{D'} \circ p') \simeq \mathcal{L}(p')$. Thus, we obtain the following corollary:

Corollary 7.3. *Suppose that \mathcal{K} and \mathcal{K}' are plentiful for \mathbb{T} , respectively, \mathbb{T}' , and that $\mathcal{K} \simeq_{\mathbf{Lex}[\mathcal{C}, \mathbf{Set}]} \mathcal{K}'$. Then $\mathcal{T}_D \simeq_{\mathbf{Set}^{\mathcal{C}^{\text{op}}}} \mathcal{T}_{D'}$.*

Acknowledgements. We thank Mark Kamsma for patiently answering our questions on model theory, no matter how basic. Both authors acknowledge the financial support from Agence Nationale de la Recherche project *Topology for types and terms* ANR-23-CE48-0012-01. The second author also acknowledges the support from Marie Skłodowska-Curie Grant No. 101273434.



Funded by
the European Union

References

- [1] Ahlbrandt, G., Ziegler, M.: Quasi finitely axiomatizable totally categorical theories. *Annals of Pure and Applied Logic* **30**(1), 63–82 (1986)
- [2] Butz, C.: Saturated models of intuitionistic theories. *Annals of Pure and Applied Logic* **129**(1), 245–275 (2004)
- [3] Caramello, O.: Topological galois theory. *Advances in Mathematics* **291**, 646–695 (2016)
- [4] Coumans, D.: Canonical extensions in logic. Ph.D. thesis, Radboud University Nijmegen (2012). URL <http://hdl.handle.net/2066/98663>
- [5] Coumans, D.: Generalising canonical extension to the categorical setting. *Annals of Pure and Applied Logic* **163**(12), 1940–1961 (2012)
- [6] Dunn, J.M., Gehrke, M., Palmigiano, A.: Canonical extensions and relational completeness of some substructural logics. *Journal of Symbolic Logic* **70**(3), 713–740 (2005)
- [7] Fussner, W., Gehrke, M.: Canonical Extensions Quickly (2025). URL <https://hal.science/hal-05231128>. Preprint
- [8] Garner, R.: Ultrafilters, finite coproducts and locally connected classifying toposes. *Annals of Pure and Applied Logic* **171**(10), 102831 (2020)
- [9] Gehrke, M., Harding, J.: Bounded lattice expansions. *Journal of Algebra* **238**(1), 345–371 (2001)
- [10] Gehrke, M., Jónsson, B.: Bounded distributive lattice expansions. *Mathematica Japonica* **40**, 207–15 (1994)
- [11] Gehrke, M., Jónsson, B.: Bounded distributive lattice expansions. *Math. Scand.* **94**, 13–45 (2004)
- [12] Gehrke, M., Priestley, H.: Canonical extensions of double quasioperator algebras: An algebraic perspective on duality for certain algebras with binary operations. *Journal of Pure and Applied Algebra* **209**(1), 269–290 (2007)
- [13] Gehrke, M., Priestley, H.A.: Canonical extensions and completions of posets and lattices. *Reports Math. Log.* **43**, 133–152 (2008)
- [14] Gehrke, M., Vosmaer, J.: Canonical extensions and canonicity via dcpo presentations. *Theoretical Computer Science* **412**(25), 2714–2723 (2011)
- [15] Gool, S.v., Marquès, J.: On duality and model theory for polyadic spaces. *Annals of Pure and Applied Logic* **175**(2), 103388 (2024)
- [16] Grossberg, R.: Classification theory for abstract elementary classes. In: Y. Zhang (ed.) *Logic and Algebra*, *Contemporary Mathematics*, vol. 302, pp. 165–204. American Mathematical Society (2003)
- [17] Haykazyan, L.: Spaces of types in positive model theory. *The Journal of Symbolic Logic* **84**, 833–848 (2019)
- [18] Hodges, W.: *Model theory*. Cambridge University Press (1993)
- [19] Johnstone, P.T.: Factorization theorems for geometric morphisms, I. *Cahiers de topologie et géométrie différentielle* **22**(1), 3–17 (1981)

- [20] Johnstone, P.T.: Sketches of an Elephant: A Topos Theory Compendium. Oxford Logic Guides 43 & 44. Oxford University Press (2002)
- [21] Joyal, A., Tierney, M.: An extension of the Galois theory of Grothendieck. *Memiors of the American Mathematical Society* **51**(309) (1984)
- [22] Jónsson, B., Tarski, A.: Boolean algebras with operators. Part I. *American Journal of Mathematics* **73**(4), 891–939 (1951)
- [23] Jónsson, B., Tarski, A.: Boolean algebras with operators. Part II. *American Journal of Mathematics* **74**(1), 127–162 (1952)
- [24] Kamsma, M.: Positive logic: An introduction for model theorists (2025). URL <https://arxiv.org/abs/2511.10167>
- [25] Lawvere, F.W.: Adjointness in foundations. *Dialectica* **23**(3/4), 281–296 (1969)
- [26] Loregian, F., Riehl, E.: Categorical notions of fibration. *Expositiones Mathematicae* **38**(4), 496–514 (2020)
- [27] Mac Lane, S., Moerdijk, I.: *Sheaves in geometry and logic: a first introduction to topos theory*. Springer (1994)
- [28] Makkai, M.: The topos of types. In: M. Lerman, J.H. Schmerl, R.I. Soare (eds.) *Logic Year 1979–80*, pp. 157–201. Springer Berlin Heidelberg, Berlin, Heidelberg (1981)
- [29] Makkai, M., Reyes, G.: *First Order Categorical Logic: Model-Theoretical Methods in the Theory of Topoi and Related Categories*. Lecture Notes in Mathematics. Springer (1977)
- [30] Malliaris, M., Shelah, S.: General topology meets model theory, on \mathfrak{p} and \mathfrak{t} . *Proceedings of the National Academy of Sciences* **110**(33), 13300–13305 (2013)
- [31] Moerdijk, I.: Continuous fibrations and inverse limits of toposes. *Compositio Mathematica* **58**(1), 45–72 (1986)
- [32] Morton, W., van Alten, C.J.: Distributive and completely distributive lattice extensions of ordered sets. *International Journal of Algebra and Computation* **28**(3), 521–541 (2018)
- [33] Pitts, A.: An application of open maps to categorical logic. *Journal of Pure and Applied Algebra* **29**(3), 313–326 (1983)
- [34] Priestley, H.A.: Representation of distributive lattices by means of ordered Stone spaces. *Bull. London Math. Soc.* **2**, 186–190 (1970)
- [35] Rogers, M.: Topological endomorphism monoids of models of geometric theories. *Theory and Applications of Categories* **42**, 41–58 (2024)
- [36] Seely, R.A.G.: Hyperdoctrines, natural deduction and the Beck condition. *Mathematical Logic Quarterly* **29**(7), 97–173 (1983)
- [37] Stone, M.H.: Topological representations of distributive lattices and Brouwerian logics. *Časopis pro Peštování Matematiky a Fysiky* **67**, 1–25 (1938)
- [38] Wrigley, J.L.: Doctrinal and groupoidal representations of classifying topoi. Ph.D. thesis, Università degli Studi dell’Insubria (2024). URL https://jlrwrigley.github.io/thesis/thesis_wrigley.pdf

- [39] Wrigley, J.L.: Some properties of internal locale morphisms externalised. *Theory and Applications of Categories* **41**(35), 1160–1202 (2024)
- [40] Wrigley, J.L.: Existential completions and Herbrand’s theorem. *Theory and Applications of Categories* **45**(23), 924–950 (2026)