VARIOUS QUESTERS

THE HITCHIKER'S GUIDE TO C.ESPINDOLA'S PROOF

OF

SHELAH'S EVENTUAL CATEGORICITY CONJECTURE

Contents

1
Preliminary Notions

This chapter serves to set up the main concepts needed to begin reading the paper. (...)

1.1 Morley's Theorem

This section is heavily based on Chapter 7.1 of [CK12], the reader is encouraged to consult this for more information.

We start off by restating the target theorem, Morley's theorem.

Theorem 1.1.1 (Morley)**.** *Let $\mathcal{L}$ be a countable language, and let T be a first-order $\mathcal{L}$-theory that has infinite models, and is categorical in some uncountable power. Then T is categorical in every uncountable power.*

Here is a refresher on saturated models, which will be used in the proof:

Definition 1.1.2 (Saturated model)**.** Let $\mathfrak{M}$ be a structure, and $X \subseteq M$. We write $\mathfrak{M}_X$ to denote the model $\mathfrak{M}$ expanded to the language containing a new constant symbol, c_x, for every element $x \in X$, with each c_x interpreted as the element x.

Let $\mathfrak{M}$ be a structure, and κ an infinite cardinal. We say that $\mathfrak{M}$ is *κ-saturated* iff, for any $X \subseteq M$ with $|X| < \kappa$, every type that is finitely realised in $\mathfrak{M}_X$ is realised in $\mathfrak{M}$.

We say that $\mathfrak{M}$ is *saturated* iff $\mathfrak{M}$ is κ-saturated for $\kappa = |M|$.

Theorem 1.1.3. *Let $\mathfrak{M}$ be an infinite model, and let κ be an infinite cardinal. Then $\mathfrak{M}$ has a κ-saturated elementary extension $\mathfrak{N}$. (Note that $\mathfrak{N}$ might not be same cardinality of $\mathfrak{M}$! It might be much bigger)*

Proof. (Adapted from van den Berg (2018), 'Syllabus Model Theory 2018/2019')

We will find a model which is κ^+-saturated; we need a regular cardinal to make the proof work, and it is immediate from the definition of κ-saturation that such a model must also be κ-saturated.

We build the new model $\mathfrak{N}$ in κ^+ many stages. For successor ordinals $\alpha + 1$, let $\mathfrak{M}_\alpha$ be the model constructed so far. Let A denote the universe of $\mathfrak{M}_\alpha$. Let $(p_i(x))_{i \in I}$ be an enumeration of all types over $Th((\mathfrak{M}_\alpha)_A)$. Let $\{b_i \mid i \in I\}$ be a collection of fresh constant symbols. Consider the theory $T_{\alpha+1} := Th((\mathfrak{M}_\alpha)_A) \cup \{p_i(b_i) : i \in I\}$. This theory is finitely satisfiable by assumption, so by the Compactness theorem, it's satisfiable. Set $\mathfrak{M}_{\alpha+1}$ to be some model of $T_{\alpha+1}$- it will then be an elementary extension of $\mathfrak{M}_\alpha$, since $T_{\alpha+1}$ includes the elementary diagram of $\mathfrak{M}_\alpha$.

For limit ordinals γ, define $\mathfrak{M}_\gamma$ to be the union of the chain $(\mathfrak{M}_\alpha)_{\alpha<\gamma}$. Then $\mathfrak{M}_\gamma$ is an elementary extension of each of the $\mathfrak{M}_\alpha$.

Here we use the closure property of elementary classes under unions of chains. Recall that AECs also have this property!

Keep going until we reach $\mathfrak{M}_{\kappa^+}$. Now consider any $X \subseteq \mathfrak{M}_{\kappa^+}$ with $|X| \leqslant \kappa^+$, and any finitely realisable type $p(x)$ over $(\mathfrak{M}_{\kappa^+})_X$. By the regularity of κ^+, there must be some $\alpha < \kappa^+$ such that $X \subseteq \mathfrak{M}_\alpha$. But then $p(x)$ is realised in $\mathfrak{M}_{\alpha+1}$ by construction, and the same element will realise p in $\mathfrak{M}_{\kappa^+}$. □

To prove Morley's theorem, we will need many, many lemmas. Here's the first one:

Lemma 1.1.4. Let $\mathcal{L}$ be a countable language, and let T be a complete $\mathcal{L}$- theory such that every model of T of cardinality ω_1 is saturated. Then every uncountable model of T is saturated.

Proof. We proceed by contraposition: we suppose that T has an uncountable model which is not saturated, and find a model of T of size ω_1 which is not saturated.

Let $\mathfrak{A}$ be a model of T which is not saturated. Then let X be a subset of A, with $|X| < |A|$, and let $p(x)$ be a type over $\mathfrak{A}_X$, such that p is finitely satisfiable in $Th(\mathfrak{A}_X)$ but not realised in $\mathfrak{A}_X$. We now let $U \subseteq A$ be any subset of A such that $|U| = |p|$. (Observe that $|U| = |p| < |A|$, because $\mathcal{L}$ is countable and $|X| < |A|$.)

□

1.2 *Abstract Elementary Classes*

This section is modelled after [Gro02]; the reader is encouraged to consult this for more information.

Definition 1.2.1. Let $\langle K, \preceq_K\rangle$ be a pair consisting of a collection of structures K for some language $L(K)$, and a relation $\preceq_K$ holding between these structures, such that:

1. $\preceq_K$ is a partial order.
2. If $M \preceq_K N$ then M is a substructure of N.
3. (Isomorphism closure): K is closed under isomorphism, and if $M, N, M', N' \in K$, $f : M \cong M'$ and $g : N \cong N'$, $f \subseteq g$ and $M \preceq_K N$ then $M' \preceq_K N'$.
4. (Coherence): If $M \preceq_K N$ and $P \preceq_K N$, and $M \subseteq P$, then $M \preceq_K P$.
5. (Tarski-Vaught Axioms): If γ is an ordinal and $\{M_\alpha : \alpha \in \gamma\} \subseteq K$ is a chain under $\preceq_K$, then
   - $\bigcup_{\alpha\in\gamma} M_\alpha \in K$;
   - If $M_\alpha \preceq_K N$ for all $\alpha \in \gamma$ then $\bigcup_{\alpha\in\gamma} M_\alpha \preceq_K N$.

6. (Lowenheim-Skolem Axiom) : There exists a cardinal $\mu \geqslant |L(K)| + \aleph_0$, such that if A is a subset of $M \in K$, then there is N such that $A \subseteq N$, $|N| \leqslant |A| + \mu$ and $N \leq_K M$.

Given an AEC $\langle K, \leq_K \rangle$ we denote by $\mathsf{LS}(K)$ the least μ in the conditions of the LS-axiom, and call it the *Lowenheim-Skolem number*.

Given such an AEC, a map $f : M \to N$ where $M, N \in K$ is a K-*embedding* if $f[M] \leq_K N$, and f is an isomorphism from M onto $f[M]$.

Let us consider some examples:

Example 1.2.2. If K is an elementary class, i.e., $K = \mathsf{Mod}(T)$ for some theory T, then it is an abstract elementary class with the relation $\leq_K$ being given by elementary *substructure*. The two first axioms are trivial, and isomorphism closure, coherence, the Tarski-Vaught axioms and the Lowenheim-Skolem axiom are all properties known in classical model theory. The Lowenheim-Skolem number is $|T| + \aleph_0$.

Originally this paragraph claims when $\leq_K$ is given elementary *embedding*, the resulting class will be an AEC, which is *not* true. Again, axiom 2 means $M \leq_K N$ implies M is a substructure of N. Of course, none of these nitpicking things survive once we consider accessible categories, which generalise AECs (cf. next Section).

Of course we would not be interested in abstract elementary classes if this were the only example on hand. The key motivation of the theory lies in the fact that some classes are, from a certain point of view, very natural, and do not look too wild to be analysed through model-theoretic methods. For example we have

- Finitely generated groups;
- Archimedean fields;
- Connected graphs;
- Noetherian rings;
- The class of algebraically closed fields with infinite transcendence degree.

It seems like there should be some setting in which one could study these models that offered tools to classify them. But it is not immediately obvious what that would be. For the majority of the 20th century, it was judged that the way forward would be to construct languages which could "tame" these classes. Let us turn to some examples of this kind; for general references on infinitary logic, the reader can consult [Kar64; Mar02; Dic85]:

Definition 1.2.3. Let κ and λ be regular cardinals and $\lambda \leqslant \kappa$. Let τ be a first order vocabulary. We denote by $\mathcal{L}_{\kappa,\lambda}(\tau)$ the language constructed using $\bigvee_\kappa$, $\bigwedge_\kappa$ and $\exists_\lambda$ and $\forall_\lambda$:

Regularity of the cardinals is assumed to ensure that any formula ϕ has at most κ many subformulas; also $\lambda \leqslant \kappa$ is a sanity condition given we maintain our atomic formulas finite (which is not necessary, but very convenient).

(1) Terms and atomic formulas are as in first order logic;

(2) If ϕ is a formula then so is $\neg\phi$;

(3) If $(\phi_\alpha)_{\alpha\in\kappa}$ is a collection of less than κ formulas, then $\bigwedge_{\alpha\in\kappa} \phi_\alpha$ and $\bigvee_{\alpha\in\kappa} \phi_\alpha$ are formulas;

(4) For a formula ϕ and variables x_α for every $\alpha \in \lambda$, we have formulas $\exists_{\alpha\in\lambda}\phi(x_\alpha)$ and $\forall_{\alpha\in\lambda}\phi(x_\alpha)$.

We let:

$$\mathcal{L}_{\infty,\lambda} = \bigcup_{\kappa\in\mathsf{Ord}} \mathcal{L}_{\kappa,\lambda}$$

We briefly mention some notions which are relevant here. Namely, given two first-order theories M, N, we write:

$$M \simeq_{\infty,\omega} N,$$

If and only if the two structures satisfy the same formulas from $\mathcal{L}_{\infty,\omega}$. An equivalent, and more useful, characterization due to Karp uses the notion of a *partial isomorphism system*.

Definition 1.2.4. Let $\mathbf{M}, \mathbf{N}$ be two first order structures in a language τ. We say that a collection P of partial τ-embeddings $f : A \to \mathbf{N}$ where $A \subseteq M$ is a *partial isomorphism system* if:

- For all $f \in P$, and $a \in \mathbf{M}$, there is some $g \in P$ such that $f \subseteq g$ and $a \in \mathsf{dom}(g)$;
- For all $f \in P$ and $b \in \mathbf{N}$, there is $g \in P$ such that $f \subseteq g$ and $b \in \mathsf{img}(g)$

We write $\mathbf{M} \cong_{\mathbf{p}} \mathbf{N}$ if there is a partial isomorphism system between these two structures.

The following relates this concept to Ehrenfeucht-Fraisse games, to the above notion of infinitary equivalence, as well as to some set-theoretic notions.

Theorem 1.2.5. *The following are equivalent for* $\mathbf{M}, \mathbf{N}$*:*

1. $\mathbf{M} \simeq_{\infty,!} \mathbf{N}$*;*
2. $\mathbf{M} \cong_{\mathbf{p}} \mathbf{N}$*;*
3. *Player II has a winning strategy in the Ehrenfeucht-Fraisse game* $G(\mathbf{M}, \mathbf{N})$*;*
4. *There is a forcing extension* $V[G]$ *such that* $V[G] \models \mathbf{M} \cong \mathbf{N}$ *(i.e., the structures are isomorphic in a forcing extension).*

Proof. The equivalence of 1-3 is known as *Karp's theorem*; for a good proof see [Mar02, Theorem 2.1.4]. The equivalence with (4) is a known set-theoretic fact, mentioned in [Mar02, Exercise 2.1.8] □

This uses essentially the fact that the language $\mathcal{L}_{\infty,\omega}$ is absolute, and constructs the extension using a forcing poset consisting of a partial isomorphism system.

For the most part we will look at infinitary logics with finite quantifiers. We will also need the notion of a *fragment*:

Definition 1.2.6. Let $\mathbb{A} \subseteq \mathcal{L}_{\infty,\omega}$ be a set of formulas in the language τ such that there is an infinite set of variables V, such that if $\phi \in \mathbb{A}$ then all of its variables occur in V. We say that $\mathbb{A}$ is a *fragment* of τ if $\mathbb{A}$ satisfies the following closure properties:

1. All atomic formulas using only the constant symbols in the vocabulary τ and the variables in V are in $\mathbb{A}$;
2. $\mathbb{A}$ is closed under subformulas;
3. $\mathbb{A}$ is closed under substitution of terms assembled from V: if $\phi \in \mathbb{A}$ and v is free in ϕ and t is a term with all of its variables in V, then the formula obtained by replacing all instances of v in ϕ by t is in $\mathbb{A}$;

4. $\mathbb{A}$ is closed under formal/single negations;

5. $\mathbb{A}$ is closed under $\neg, \wedge, \vee, \exists v, \forall v$ for $v \in V$.

If $\mathbb{A} \subseteq \mathcal{L}_{\omega_1,\omega}$, and $|\mathbb{A}| \leqslant \aleph_0$, we say that it is a *countable fragment*.

The following definition is the crucial one:

Definition 1.2.7. Let M and N be structures in a language L, and $\mathbb{A}$ is an L-fragment. We write $M \subseteq_{TV,\mathbb{A}} N$ if and only if:

1. $M \subseteq N$ and,

2. For every $\bar{a} \in M$ and every formula $\phi(y, \bar{x}) \in \mathbb{A}$, if $N \models \exists y \phi(y, \bar{a})$, then there exists some $b \in M$, such that $N \models \phi(b, \bar{a})$.

Example 1.2.8 (Models of a countable infinitary theory)**.** Let T be a countable theory in a language $\mathcal{L}$, and let $\mathbb{A}$ be a fragment containing T. Let $K = \mathsf{Mod}(T)$. Let $M \preceq_K N$ if and only if $M \subseteq_{TV,\mathbb{A}} N$. Then $\langle K, \preceq_K \rangle$ is an abstract elementary class. The trickier parts to verify are the Lowenheim-Skolem and the union axiom; but both of these follow by the same proofs as their first-order correspondents.

The restriction to countable theories is sharp: it is not hard to find a theory T, in a countable language, of $\mathcal{L}_{\omega_1,\omega}$ which models are at least of size $2^{\aleph_0}$.

However AEC's are not at all limited to examples coming from logic. Let us see some preliminary examples, and then conclude with a wild, unexpected, example, which breathed new life to the field.

Example 1.2.9 (Noetherian Rings)**.** Let K be the class of noetherian rings. We define $R \preceq_K S$ if and only if R is a subring of S, and $R \simeq_{\infty,\omega} S$. Note that then R is noetherian if and only if S is noetherian. To see this, note that if we assume that R is noetherian and S is not, then (by an equivalent characterization), there is $f_1, f_2, ...,$ a sequence of elements such that for every integer n there is some f_i, such that f_i cannot be written in terms of the smaller elements. Then we claim that Player I has winning strategy in an unbounded Ehrenfeucht-Fraisse game: successively pick elements from that sequence. Once the game is played out, whatever Player II has chosen, say a sequence $g_1, g_2, ...,$ there must be an integer n such that each g_i is a linear combination of g_k for $k \leqslant n$. But then this sequence cannot be isomorphic to the former.

It is clear that if $R \preceq_K S$ then R is a substructure of S, and isomorphism closure and coherence are obvious. The Tarski-Vaught axiom follows from the fact that chains of models respect the $\simeq_{\infty,\omega}$ relation. I could not prove the Lowenheim-Skolem axiom, though Grossberg's notes claim it (Shrug).

Perhaps the most striking example – and one which in part revived the interest in this topic from the point of view of mainstream mathematics – is in the work of Boris Zil'ber's "Schanuel's Structures".

Definition 1.2.10. Let $\mathcal{K}_e$ be defined as:

$$\mathcal{K}_e := \{\langle F, +, \cdot, exp\rangle : F \text{ is an alg. closed field of characteristic zero, } \forall x \forall y (exp(x+y) = exp(x) \cdot exp(y))\}$$

also let

$$\mathcal{K}_{pexp} := \{\langle F, +, \cdot, exp\rangle \in \mathcal{K}_e : \mathsf{ker}(exp) = \pi\mathbb{Z}\}.$$

We consider the class of *Schanuel structures* to be the class $\mathcal{K}_{exp} \subseteq \mathcal{K}_{pexp}$ which satisfies some conditions, amongst them the *Schanuel condition*.

We are willfully vauge on the extra conditions, as they are not important, except for the Schanuel condition. For more information on this, check [Mar02, Chapter 8], and see also Will Boney's notes on this topic.

This essentially imposes that the so-called "Schanuel conjecture" be true:

Conjecture 1.2.11 (Schanuel,1960)**.** Assume that $x_0, ..., x_n \in F$ are linearly independent over $\mathbb{Q}$. Then $\mathbb{Q}(x_0, ..., x_n, exp(x_0), ..., exp(x_n))$ has transcendence degree at least n over $\mathbb{Q}$.

To understand the significance of this, it should be noted that the study of transcendental numbers is one of the most unexplored and difficult areas of number theory.

Schanuel's conjecture is a piece of machinery that would clarify many difficult conjectures in transcendental number theory. As a toy example, recall that it is widely assumed that $e + \pi$ is transcendental, though no proof of it is in sight; this would immediately fall off from the above result: if we set $x_0 = 1$ and $x_1 = \pi * i$, then $\mathbb{Q}(\pi, e)$ (the result of the field extension) would have transcendence degree at least 2, showing that there is no polynomial $f(x, y)$ such that $f(\pi, e) = 0$; this implies that $e + \pi$ is transcendental.

Now what Zil'ber did was note that $\mathcal{K}_{exp}$ can be given a relation $\preceq$, forming an abstract elementary class. Additionally, using some heavy model-theoretic and number-theoretic weaponry, he managed to prove that:

Theorem 1.2.12. *The theory $\mathcal{K}_{exp}$ has a unique model of cardinality $2^{\aleph_0}$.*

Thus, the only problem lies in proving that this model is indeed the model of the complex field $\mathbb{C}$, i.e., prove that the latter has the model-theoretically desirable properties. This is an active research area today.

Such examples motivate the idea that abstract elementary classes are indeed ubiquitous, and serve as a strong foundation for exploring non-trivial solutions to mathematical problems. However, as discussed in the logical dream, just like for first-order logic, this appears as a matter of finding the right "dividing lines". Hence we can encounter our appropriate generalization of Los' conjecture:

Conjecture 1.2.13 (Shelah)**.** Let $\mathcal{K}$ be an AEC. If there is a $\lambda \geqslant \beth_{2^{LS(\mathcal{K})})^+}$ such that $\mathcal{K}$ is categorical in λ, then $\mathcal{K}$ is categorical in all μ for $\mu \geqslant \beth_{2^{LS(\mathcal{K})})^+}$.

1.3 Accessible Categories

In the category of sets, we have the special property that the elements of a set X (its internal structure) correspond exactly to maps from the terminal object to X (its external structure), i.e. $\mathrm{Hom}_{\mathbf{Set}}(1, X) \cong X$. This is not the case for all categories in general. For a wider class of

A warning: we do not mention *sketches* in our presentation. While they play a central role in the books that introduced accessible categories, they are not as necessary for modern treatments, which is why we have chosen not to include them here. Similarly goes for *synthetic categories* as theories, but for different reasons.

categories, the internal structure of objects can still be reflected externally, but we require more than just one object to "probe" its internal structure.

Example 1.3.1. In the category of undirected graphs, the terminal object is the single node graph with a reflexive edge. However, graph homomorphisms from this terminal graph will never reveal anything about irreflexive nodes.

Figure 1.1: The terminal graph.

What we can do instead is to use two graphs: one graph with a single irreflexive node (targetting the nodes) and one graph with two nodes and an edge between them(targetting the edges). Any graph can be constructed by taking multiple copies of these diagrams (with some gluing homomorphisms f and g that designate a node to be the endpoint of some edge). Another way of viewing this is that any graph can be reconstructed as a colimit of a diagram inside the subcategory consisting of these two graphs and f, g.

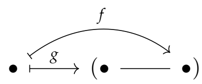

Figure 1.2: The probing subcategory.

In general the objects that we consider will have to be simple building blocks: a first candidate is for the object to be "finite".

Locally Finitely Presentable Categories

Definition 1.3.2. A diagram $D : \mathcal{I} \to \mathcal{C}$ is *directed* if $\mathcal{I}$ is a directed poset (considered as a category): every finite subset has an upper bound.

Definition 1.3.3 (Finitely Presentable Object [AR94])**.** An object K of a category $\mathcal{C}$ is *finitely presentable* if for each directed colimit $(D_i \xrightarrow{c_i} C)_{i \in I}$, any morphism $f : K \to C$ factorizes uniquely through some c_i, i.e. there exists a c_i (not necessarily unique) such that there exists a unique $g : K \to D_i$ with $f = c_i \circ g$.

More succintly, K is finitely presentable if the hom-functor $\mathrm{Hom}_{\mathcal{K}}(C, -)$ preserves directed colimits.

Example 1.3.4. The finitely presentable objects in **Set** are exactly the finite sets. We sketch why this is the case.

One can see this example as the motivating example for the above category-theoretic definition - after all, finiteness for an object is only well defined in categories of sets.

For any set X, the diagram D_X consisting of its finite subsets with inclusion functions is directed with X being the colimit of X. If X is finitely presentable, then the identity map id_X factors as $id_X = (D_X(i) \hookrightarrow X) \circ g$, but this means g has to be the identity map with $D_X(i) = X$, where $D_X(i)$ is finite.

On the other hand, suppose X is finite, and take any $f : X \to C$ and any directed diagram D with C as colimit. For each $x \in X$, $f(x)$ must be an element of $D(i_x)$ for some i_x. However, since D is a directed diagram and X is finite we can find a $D(i)$ which contains all the $D(i_x)$, and therefore all the $f(x)$. We can then factor f through $D(i)$.

$$\begin{tikzcd}
C & & & \\
& f[X] \subseteq D(i) \arrow[ul] & & \\
f(x_1) \in D(i_{x_1}) \arrow[ur] & f(x_2) \in D(i_{x_2}) \arrow[u] & \dots & f(x_k) \in D(i_{x_k}) \arrow[ull] \\
X \arrow[uuu, bend left, "f"] \arrow[uur] & & &
\end{tikzcd}$$

Since any set X is a colimit of a directed diagram of finitely presentable sets D_X, we say that **Set** is locally finitely presentable.

Definition 1.3.5. A category $\mathcal{C}$ is *locally finitely presentable* (LFP) if it is cocomplete and has a set of finitely presentable objects $\mathcal{A}$ s.t. every object is a directed colimit of objects from $\mathcal{A}$.

Intuition: $\mathcal{C}$ is locally finitely presentable if the finitely presentable objects essentially determine the rest of the category.

Many categories are LFP - in particular, the category of models of an equational theory (i.e. finitary varieties of algebras) is LFP. We illustrate this by way of the category of groups.

Example 1.3.6 ([AR94]). A *finitely presentable* group is a group which has a presentation $\langle A|R \rangle$ with A a finite set of generators and R a finite set of relations of said generators. The finitely presentable objects in **Grp** are exactly the finitely presentable groups[1].

[1] In fact the general terminology is derived from this.

Suppose G is a finitely presentable group. Then analogously to **Set**, we can consider the diagram D_G of subgroups generated by finitely many generators (with inclusion maps between subgroups), which has G as its colimit. Then the identity map id_G factors through some $D_G(i)$ and therefore $A = D_G(i)$.

For the relations, we again use the same technique of constructing a diagram and factoring the identity map. We have shown G can be finitely generated (let's call this finite set of generators X), so we can consider the canonical map $k : F(X) \to G$ from the free group generated by X to G. The kernel set

$$\ker k = \{(t, t') \in F(X) \mid k(t) = k(t')\}$$

contains all possible relations that hold between the terms when interpreted as elements of A. Hence, X generates G using the equivalence relation $\ker k$. Now, we can once again construct a directed diagram of the groups $D_G(i)$ generated by X and using finite subsets E_i of $\ker k$[2] - this diagram has G as colimit, and we can once again factor the identity map id_G through this diagram (see [AR94, p. 144] for details, as well as for the proof of the other direction).

We can think of k as the mapping from syntax (i.e. elements of the free group) to semantics (elements of G).

$$\begin{aligned} &k : F(X) \to G \\ &k(x \in X) := x \\ &k(e) := e_G \\ &k(t_1 \cdot t_2) := k(t_1) \cdot_G k(t_2) \\ &k(t^{-1}) := k(t_1)^{-1} \end{aligned}$$

Categorically, the kernel set can be constructed as the pullback

$$\begin{tikzcd} \ker k \arrow[r] \arrow[d] \arrow[dr, phantom, "\lrcorner", very near start] & X \arrow[d, "|k|"] \\ |F(X)| \arrow[r, "|k|"'] & |G| \end{tikzcd}$$

[2] i.e. take the congruence closure of E_i.

Any group is a filtered colimit of its finitely generated subgroups. Hence, the category **Grp** is locally finitely presentable. ■

In fact, even some non-equational theories have LFP categories of models. Consider the theory of partially ordered sets.

Example 1.3.7. In the same way as *Set* is a *LFP*, *Pos*, the category of posets and monotone maps, is an *LFP*. Since every poset is the union of all its finite subsets under the (restricted) ordering, the *FP*-objects in *Pos* are exactly the finite posets.

It can be enlightening to see what a locally finitely presentable poset.

Example 1.3.8. A poset, seen as a category, is *LFP* if and only if it is a complete algebraic lattice. Since algebraic lattices are those which are generated by joins of finite (compact) elements, it is exactly the finite elements that correspond to the finitely presentable objects in the category.

Example 1.3.9. A *non-example* is the category **FinSet** of finite sets and functions. We will come back to this example later.

Limit Theories

In general, we can characterise the locally finitely presentable categories as the categories of models of a finitary limit theory.

Definition 1.3.10. Let Σ be a first-order signature of function and relation symbols. We define the category $\mathbf{Str}\Sigma$ of structures interpreting Σ, with the morphisms being structure homomorphisms.

Definition 1.3.11 (Limit Theory)**.** A set of $\mathcal{L}_{\omega,\omega}$ sentences T is a *limit theory* if every sentence in T is of the form

$$\forall \vec{x}(\varphi(\vec{x}) \to \exists! \vec{y}\, \psi(\vec{x}, \vec{y}))$$

where φ and ψ are conjunctions of atoms.

Definition 1.3.12 (Model Category)**.** Given a theory T in signature Σ, the category $\mathbf{Mod}T$ of models of T is the full subcategory of $\mathbf{Str}\Sigma$ containing structures that satisfy T.

Theorem 1.3.13 ([AR94, p. 207])**.** *$\mathcal{C}$ is LFP iff it is equivalent to $\mathbf{Mod}T$ for some limit theory T in $\mathcal{L}_{\omega,\omega}$.*

Proof. (sketch) ($\Leftarrow$) The category $\mathbf{Str}\Sigma$ is complete, cocomplete and LFP [AR94, p. 201][3]. We show that $\mathbf{Mod}T$ is closed under limits and directed colimits when T is a limit theory. This is enough to show that $\mathbf{Mod}T$ is LFP. TODO

($\Rightarrow$) TODO □

[3] $\mathbf{Str}\Sigma$ inherits a lot of its structure from considering the algebraic structures interpreting its function symbols (forgetting the relations).

We use the theories of groups and posets as concrete examples of limit theories.

Example 1.3.14. We axiomatise the theory of groups with a constant symbol e, unary function symbol $-^{-1}$ and $- \cdot -$ in the following way:

$$\forall x(\top \Rightarrow x \cdot e = x \wedge e \cdot x = x)$$
$$\forall x(\top \Rightarrow \exists! y(x \cdot y = e \wedge y \cdot x = e))$$
$$\forall x_1, x_2, x_3(\top \Rightarrow ((x_1 \cdot x_2) \cdot x_3 = x_1 \cdot (x_2 \cdot x_3)))$$

Since each sentence is a limit sentence, the theory of groups is a limit theory. The category of models is exactly $\mathbf{Grp}$.

Example 1.3.15. We similarly axiomatise posets by limit sentences:

$$\forall x(\top \Rightarrow x \leqslant x)$$
$$\forall x_1, x_2(x_1 \leqslant x_2 \wedge x_2 \leqslant x_1 \Rightarrow x_1 = x_2)$$
$$\forall x_1, x_2, x_3(x_1 \leqslant x_2 \wedge x_2 \leqslant x_3 \Rightarrow x_1 \leqslant x_3)$$

Hence, the theory of posets is also a limit theory. The category of models is exactly $\mathbf{Pos}$.

Locally Presentable Categories

The correspondence between locally finitely presentable[4] categories and models of finite limit theories suggests that if we move to infinitary logic, we must also loosen the finite presentability[5].

[4] A more suggestive name in this context would be "compact categories"

[5] i.e. infinitary logic is no longer compact Is this true? idk I'm tired

Definition 1.3.16. Let λ be a regular cardinal. The diagram $D : \mathcal{I} \to \mathcal{C}$ is *λ-directed* if $\mathcal{I}$ is a poset where every subset of cardinality $< \lambda$ has an upper bound.

We consider only regular cardinals because otherwise if λ is singular, we may get an upper bound whose "size" is greater than or equal to λ. In particular, consider the example of the $\aleph_\omega$-directed diagram of subsets of $\aleph_\omega$.

Definition 1.3.17. Let λ be a regular cardinal. An object C of a category $\mathcal{C}$ is *λ-presentable* if $\mathrm{Hom}_{\mathcal{C}}(C, -)$ preserves λ-directed colimits.

$\mathcal{C}$ is *locally λ-presentable* if it is cocomplete and has a set of λ-presentable objects $\mathcal{A}$ s.t. every object is a directed colimit of objects from $\mathcal{A}$.

$\mathcal{C}$ is *locally presentable* if it is *locally λ-presentable* for some λ.

The less succint definition is analogously obtained by replacing "directed diagram" with "λ-directed diagram".

Going back to our example of a poset, seen as a category, we can see the difference between an *LFP* category and an *LP* category.

Example 1.3.18. A poset, seen as a category, is locally presentable if and only if it is a complete lattice. Notice here the lack of requirement of algebraicity, that the lattice is generated by finite (compact) elements.

There is an interesting theorem, which we will note for posterity's sake but not use in the sequel.

Theorem 1.3.19 (Gabriel-Ulmer). CITE HERE *If $\mathbb{C}$ is a locally presentable category (locally finitely presentable), $\mathbb{C}^{op}$ is never locally presentable (locally finitely presentable), unless it is a poset.*

We end with going back to our earlier non-example.

Example 1.3.20. **FinSet** is *not* locally presentable since it is not cocomplete. To prove this, take all singletons $\{x\}$ with $x \in \mathbb{N}$. Their directed colimit, seen as union, is all of $\mathbb{N}$, which clearly is not finite.

This shows us that while local presentability generalises local *finite* presentablility, it does not include some quite natural categories.

On the logical side, the corresponding language for FPs is $\mathcal{L}\infty$. Still allowing the same limit theories, we obtain an analogous representation theorem as for LFPs.

Theorem 1.3.21. *$\mathcal{C}$ is locally presentable iff it is equivalent to the category of models of some limit theory in $\mathcal{L}\infty$.*

In other words, the generalisation to LPs changes the logical language to an infinitary one, but the theories are the same. As noted above, there are many familiar theories that are not limit theories, such as the theory of finite sets and the theory of linear orders.

Accessible Categories

Accessible categories are locally presentable categories that aren't necessarily cocomplete.

Definition 1.3.22. A category $\mathbb{C}$ is *λ-accessible* for a regular cardinal λ if $\mathbb{C}$ has λ-directed colimits and $\mathbb{C}$ has a set of λ-presentable objects such that every object in $\mathbb{C}$ is a λ-directed colimit of objects from this set. A category $\mathbb{C}$ is *accessible* if it is *λ-accessible* for some regular cardinal λ.

Compare this with the definition of locally presentable categories. Note that the condition for cocompleteness has been removed.

Example 1.3.23. The category of linear orders, **LinOrd** is accessible. The finitely presentable objects are exactly the finite linear orders and every linear order is a directed colimit of these.

Note that the reason **LinOrd** is not LP is due to the fact that **LinOrd** is not cocomplete.

Example 1.3.24. **FinSet** is accessible. Every finite set is a presentable object and thus every object in **FinSet** is finitely presentable.

Since we no longer require **FinSet** to be cocomplete, it now fits in our definition.

Example 1.3.25. All the previous examples of LFPs and LPs are trivially accessible categories.

Basic Theories

Just as LP categories correspond to limit theories, accessible categories correspond as well to basic theories.

Definition 1.3.26 ([AR94, p. 227])**.** A formula in $\mathcal{L}_{\infty,\infty}$ is called

1. *positive-primitive* if it has the form $\exists Y \psi(X,Y)$[6] where $\psi(X,Y)$ is a conjunction of atomic formulas.

2. *positive-existential* if it is a disjunction of positive-primitive formulas.

3. *basic* if it has the form $\forall X(\phi(X) \to \psi(X))$ where ϕ and ψ are positive-existential formulas.

A set of $\mathcal{L}_{\infty,\infty}$ sentences T is a *basic theory* if every sentence in T is basic.

[6] We use capital letters to denote a set of variables being quantified over.

Note that basic theories indeed generalises limit theories because limit sentences of the form

$$\forall X(\phi(X) \to \exists! Y \psi(X,Y))$$

can be replaced by two basic sentences:

$$\forall X(\phi(X) \to \exists Y \psi(X,Y))$$

$$\forall X, Y, Z(\phi(X) \wedge \psi(X,Y) \wedge \psi(X,Z) \to Y = Z)$$

Example 1.3.27. **FinSet** is basic since we can axiomatise finite sets via the sentence

$$\forall x_0, x_1, \dots (\bigvee_{i \neq j \in \omega} x_i = x_j)$$

.

Example 1.3.28. **LinOrd** is basic - in particular the sentence

$$\forall x, y(x \leqslant y \vee y \leqslant x)$$

is basic.

In some sense[7], working with basic theories only causes no loss in generality as compared to working with all theories.

[7] I'm not really sure in what "sense" this causes no loss in generality, but this seems like an important thing to know.

Definition 1.3.29. $\mathcal{F} \subseteq \mathcal{L}_{\infty,\infty}$ is a *fragment* if

1. all atomic formulas are in $\mathcal{F}$.

2. $\mathcal{F}$ is closed under substitution of terms for free variables.

3. $\mathcal{F}$ is subformula closed.

4. $\forall X \phi(X) \in \mathcal{F}$ implies $\neg \exists X \neg \phi(X) \in \mathcal{F}$ and if $\phi \to \psi \in \mathcal{F}$ then $\neg \phi \vee \psi \in \mathcal{F}$.

The $\mathcal{F}$-basic formulas are defined just like basic formulas, except atomic formulas are replaced by formulas of $\mathcal{F}$

Definition 1.3.30. Given a fragment $\mathcal{F}$, a structure map $h : M \to N$ is an *$\mathcal{F}$-elementary map* if h preserves the meaning of all formulas in $\mathcal{F}$, i.e. $M \models \varphi[A \subseteq M]$ iff $N \models \varphi[h[A]]$.

Define the category $\mathbf{Mod}^{\mathcal{F}}T$ as models of T but where the morphisms are the $\mathcal{F}$-elementary maps.

Theorem 1.3.31 ([MP89, p. 52])**.** *Given any small fragment $\mathcal{F}$ and a $\mathcal{F}$-basic theory T, the category $\mathbf{Mod}^{\mathcal{F}}T$ is equivalent to $\mathbf{Mod}T'$ for some other basic theory T' in some other language L'.*

AECs are Accessible Categories

This whole subsection follows from CITE HERE, Lieberman

We can see an AEC as a category by letting its models be the objects and morphisms the strong embeddings between these. The following theorem connects the story of accessible categories to that of AECs.

Theorem 1.3.32. *CITE HERE Let K be an AEC and μ its cardinal defined in the Löwenheim-Skolem Axiom. Then, K is μ^+-accessible. More generally, K is λ-accessible for all regular cardinals $\lambda > \mu$.*

It follows that every AEC is accessible. Moreover, we have a characterisation theorem which needs just a little preamble.

Given a signature L over an AEC K, denoted $L(K)$, we can define the category of L-structures as the category with objects being the L-structures and morphisms are injective L-morphisms that preserve and reflect relations in L. We will denote this category by L-**Struct**.

We now need two definitions to obtain our characterisation. That of *repleteness* and being *almost full*.

Definition 1.3.33. A subcategory $\mathbb{D}$ of some category $\mathbb{C}$ is *replete* if there are no isomorphic objects $X \cong Y$ such that $X \in \mathbb{D}_0$ while $Y \notin \mathbb{D}_0$.

In other words, replete subcategories respect isomorphisms. Examples of these abound: take topological spaces and continuous maps. Any combination of topological properties forms a replete subcategory, with *Haus* for Hausdorff spaces, *Stone* for Stone spaces, and Sob_ω for Sober spaces of cardinality ω.

Definition 1.3.34. A subcategory $\mathbb{D}$ of some category $\mathbb{C}$ is *almost full* if for any two objects X, Y, Z in $\mathbb{D}$, with $f : X \to Z$ and $g : Y \to Z$, if there is a $h : X \to Y$ in $\mathbb{C}$ such that $g \circ h = f$, then h is in $\mathbb{D}$ as well.

So, $\mathbb{D}$ is almost full if it has the base of every commutative triangle of which it has the other sides. Diagrammatically:

$$\begin{tikzcd} X \arrow[rr, dashed, "h"] \arrow[dr, "f"'] & & Y \arrow[dl, "g"] \\ & Z & \end{tikzcd}$$

This property corresponds to the coherency axiom for AECs.

Now, we can characterise AECs more precisely as a special form of subcategory of L-**Struct**.

Theorem 1.3.35. *Let K be an AEC. Then K is a nearly full, replete subcategory of L-**Struct** which is λ-accessible for every $\lambda \gneq \mu$ (where μ is the cardinal from the Löwenheim-Skolem axiom) and which has all directed colimits. (The directed colimits are computed as in L-**Struct**.)*

We have a converse result as well.

Theorem 1.3.36. *Any nearly full, replete subcategory of L-**Struct** which is λ-accessible for every $\lambda \gneq \mu$, for some cardinal μ, and which has all directed colimits which are computed as in L-**Struct** can be seen as an AEC.*

Give the specific construction here.

1.4 Categorical Logic

The essential theme of categorical logic can be formulated as follows:

> Theories *are* (certain types of) categories; Models *are* (certain types of functors.

In fact, we have a list of correspondence between different fragments of logic and various types of categories as follows:[8]

Logic	Category
algebraic $(=)$	fin. product
cartesian $(\top, \wedge, =^*)$	fin. limit
regular (cat. + $\exists$)	regular
coherent (reg. + $\perp, \vee$)	coherent
boolean (coh. + $\neg$)	boolean
geometric (coh. + $\bigvee$)	geometric

For each fragment of logic, we require certain categorical structures to interpret the logical connectives within the logic. The basic ideal is as follows. Suppose we have a category $\mathcal{C}$, for any object A within, we view $\mathrm{Sub}(A)$ as the *predicates* in $\mathcal{C}$ over A.[9] Now the category $\mathcal{C}$ can interpret the logic connective $\wedge$, iff for any object A in $\mathcal{C}$, $\mathrm{Sub}(A)$ admits finite meets. It is slightly more involved for other logical connectives, but the essential idea is the same. We refer the readers to [Car18] for more details.

[8] Here $=^*$ means it allows partial functions, and equality can compare values of potentially undefined terms.

[9] $\mathrm{Sub}(A)$ is the partial order of monic maps into A modulo isomorphism. For two monic maps, $i : X \hookrightarrow A$ and $j : Y \hookrightarrow A$, $i \leqslant j$ iff there exists map $k : X \to Y$ that commutes with i, j.

Theories as Categories

Now suppose we have a theory T within some fragment of logic. For simplicity and sufficient generality, we suppose T is *regular*, so that it allows $\top, \wedge, \exists$ in its language. Since T potentially lacks implication, its axiomatisation and derivation all involve *sequents* of formulas of the following form,

$$\varphi \vdash_{\vec{x}} \psi.$$

In the above sequent, φ and ψ are required to be formulas within our fragment of logic. In our example of a regular theory T, they can only contain $\top, \wedge, \exists$ as logical connectives. The subscript $\vec{x}$ is a list of variable names called the *context* of the sequent, and it is required that all free variables of φ and ψ are contained in this context.

Given such a theory T, we can construct a corresponding category $\mathcal{C}_T$, called its *syntactic category*. It is the first-order analogue of the Lindenbaum-Tarski algebra of a propositional theory. Its concrete construction is as follows:

- Objects are of the form $(\vec{x}.\varphi)$, where $\vec{x}$ is a context and φ is a formula, whose free variables are contained in the context $\vec{x}$.[10]
- Morphism from $(\vec{x}.\varphi)$ to $(\vec{y}.\psi)$ are T-provably functional formulas $\theta(\vec{x},\vec{y})$ between them, identified upto T-provable equivalence.[11]

[10] For simplicity, identify α-equivalent formulas. Two objects $(\vec{x}.\varphi)$ and $(\vec{y}.\psi)$ are α-equivalent iff there is a bijection between $\vec{x}$ and $\vec{y}$, such that ψ is obtained from φ by substituting its free variables along this map from $\vec{x}$ to $\vec{y}$ (with possibly renaming bounded variables). The upshot is that for any two objects $(\vec{x}.\varphi)$ and $(\vec{y}.\psi)$, we may assume their contexts are disjoint.

[11] From now on, whenever we write $\vec{x},\vec{y},\vec{z},\cdots$, we always assume these contexts are disjoint; see the above footnote.

Since our theory T may lack implication and universal quantifier, a formula $\theta(\vec{x},\vec{y})$ is T-provably functional from $(\vec{x}.\varphi)$ to $(\vec{y}.\psi)$ iff the following sequents are provable in T:[12]

$$\theta(\vec{x},\vec{y}) \vdash_{\vec{x},\vec{y}} \varphi(\vec{x}) \wedge \psi(\vec{y})$$
$$\varphi(\vec{x}) \vdash_{\vec{x}} \exists\vec{y}\theta(\vec{x},\vec{y})$$
$$\theta(\vec{x},\vec{y}) \wedge \theta(\vec{x},\vec{z}) \vdash_{\vec{x},\vec{y},\vec{z}} \vec{y} = \vec{z}$$

[12] If T even lacks $\exists$, then T provably functional formulas are interpreted as terms in T.

Theorem 1.4.1. *For any theory T in some fragment of logic, its syntactic category $\mathcal{C}_T$ will be a category of the corresponding type of the logic. Moreover, for any (small) category $\mathcal{C}$ of some type, there exists some theory T in its corresponding fragment of logic such that $\mathcal{C} \simeq \mathcal{C}_T$.*

The above theorem justifies the identification of (particular fragments of) theories with (certain types of) categories.

1.5 Models as Functors

The importance of the syntactic category lies in the fact that it is the representable object of the (2-)functor T-mod, which takes a category $\mathcal{E}$ to its category of T-models within $\mathcal{E}$. This sounds like a mouthful, but let us see closely what this sentence really says.

Again, consider a, say regular, theory T. Suppose M is a *model* of T, then we can naturally associate a *functor* from $F_M : \mathcal{C}_T \to \mathbf{Set}$ as follows,

- For object $(\vec{x}.\varphi)$, F_M sends it to the following definable subset,[13]
$$F_M(\vec{x}.\varphi) = \{\, \vec{a} \in M^n \mid M \models \varphi[\vec{a}] \,\}.$$
- For each T-provably functional formula θ from $(\vec{x}.\varphi)$ to $(\vec{y}.\psi)$, F_M sends it to the function whose graph is defined by θ.[14]

[13] Notice that α-equivalent formulas define the same definable subset, hence this is well-defined.

[14] Again, notice that T-provably equivalent formulas must define the same graph, hence this is well-defined.

The intuition is that since $\mathcal{C}$ is built out of syntactic objects, any model M can interpret these syntactic objects, and produce corresponding definable subsets of M. Furthermore, such a functor F_M must preserve the corresponding logical/categorical structure within $\mathcal{C}_T$. For instance, we must have

$$F_M(\vec{x}.\varphi \wedge \psi) = F_M(\vec{x}.\varphi) \cap F_M(\vec{x}.\psi),$$

and similarly for other logical connectives (that exists in T).

We turn the above observation into a *definition*:

Definition 1.5.1. Let T be some theory in a fragment of logic, and let $\mathcal{E}$ be a category of its corresponding type. A *model* of T in $\mathcal{E}$ is defined to be a *functor* from $\mathcal{C}_T$ to $\mathcal{E}$ that preserves the logical/categorical structures within this fragment.

More precisely, we have the following definition,

$$T\text{-mod}(\mathcal{E}) := \mathrm{Fun}^*(\mathcal{C}_T, \mathcal{E}),$$

where the right hand side is the category of functors from $\mathcal{C}_T$ to $\mathcal{E}$ that preserves the corresponding logical/categorical structures, with morphisms being natural transformations. This way, we can talk about models of T in *any* (suitable) category $\mathcal{E}$.

2
Introduction to Topos Theory

In this chapter we introduce the basic philosophy of topos theory.

2.1 What is Topos Theory About?

We start with an explanation of a common theme in mathematics. Suppose now X is some *space*. For concreteness, let us assume X to be a "nice" topological space. Associating to X there are two important aspects we can consider:

- *Geometry*: *points* on X: $\{ x : X \}$, which will be equipped with a continuous structure, viz. *topology*;
- *Algebra*: continuous *functions* on X: $C(X) = \{ f : X \to \mathbb{R} \}$, which will be equipped with an algebraic structure, viz. an $\mathbb{R}$-*algebra*.[1]

Notice that our notation suggests we really want to distinguish a space X with its set of points $\{ x : X \}$ equipped with a topology. The reason for this will become more clear later. The two aspects are related to each other in the following manner:

1. Any $f \in C(X)$ will induce a function from points $\{ x : X \}$ to $\mathbb{R}$

$$x \mapsto f(x),$$

   that respects the geometric structure, i.e. it will be a continuous map between topological spaces (simply by definition).

2. Any $x \in \{ x : X \}$ will induce a function ev_x from $C(X)$ to $\mathbb{R}$,

$$f \mapsto f(x),$$

   that respects the algebraic structure, i.e. it will be a homomorphism between $\mathbb{R}$-algebras.

All things we have said so far seem very trivial, but they embodies an absolutely *essential* philosophy of modern mathematics: the *duality* between geometry and algebra. From a more categorical perspective, we actually have the following picture:

$$\begin{tikzcd}
\{ y : Y \} \arrow[r, dashed, no head] \arrow[d, "\alpha"'] & C(Y) \\
\{ x : X \} \arrow[r, dashed, no head] & C(X) \arrow[u, "\alpha^*"']
\end{tikzcd}$$

[1] An $\mathbb{R}$-algebra is a commutative ring equipped with a compatible $\mathbb{R}$-action. Here the ring structure on $C(X)$ is defined point-wise using the ring structure on $\mathbb{R}$, and the $\mathbb{R}$-action is simply given by multiplying a constant to a function.

Here α denotes a continuous morphism between topological space. It induces a homomorphism between $\mathbb{R}$-algebras,

$$\alpha^* : C(X) \to C(Y),$$

essentially by *precomposing* with α,

$$\begin{tikzcd} Y \arrow[r, "\alpha"] \arrow[dr, "\alpha^*(f)"'] & X \arrow[d, "f"] \\ & \mathbb{R} \end{tikzcd}$$

From this perspective, the (1) and (2) above can actually be realised as special cases when $\{ y : Y \}$ is a *one-point space*: A point in X is the same as a continuous map from the one-point space to X that selects this point, and ev_x is exactly given by precomposition with this map.[2]

The deep fact is that, at least in good cases, the two perspectives are *equivalent*, and these results are usual denoted as *duality* result. For compact Hausdorff spaces, this is usually called *Gelfand duality*; for Boolean algebras, this is called *Stone duality*.

Notice that in these types of duality, we can always find a special object that serve as *both* a geometric object *and* an algebraic object. In the above example, this object is $\mathbb{R}$. We have secretly used the fact that $\mathbb{R}$ can be both viewed as a space and an algebra; we already require both to even define $C(X)$ and show it has an $\mathbb{R}$-algebra structure.

Now we argue that the very *same* set of ideas works for *theories* as well. Let T be any theory.[3] There are again the following two perspectives on T:

- *Algebra*: The syntactic category $\mathcal{C}_T$ can be thought of *the algebra of* **Set**-*valued functions* on T.
- *Geometry*: Any *model* M of T induces an evaluation functor

2.2 Grothendieck toposes

In this section we will define the notion of a Grothendieck topos over **Set**. Almost all of the material comes from [Car18]. For the sake of simplicity we bluntly ignore all size issues. We begin by recalling the following fundamental construction in category theory.

Definition 2.2.1. Let $\mathcal{C}$ be any category. We denote by $\widehat{\mathcal{C}}$ the *category of presheaves* on $\mathcal{C}$. Its objects are *presheaves*, *i.e.* functors from $\mathcal{C}^{op} \to \mathbf{Set}$, and its morphisms are natural transformations.

A Grothendieck topos will be a certain kind of full subcategory of a category of presheaves. The presheaves that belong to this subcategory will be called *sheaves*, but exactly which presheaves those are depends on additional data.

Definition 2.2.2. A *sieve*[4] S on an object C of some category $\mathcal{C}$ is a collection of arrows with codomain C that is closed under precompositon. That is, if $f \in S$ and $\mathrm{dom}(f) = D$, then $gf \in S$ for every arrow g in $\mathcal{C}$ with $\mathrm{cod}(g) = D$.

[2] Notice that the algebra of continuous functions on a one-point space is exactly given by $\mathbb{R}$ itself.

[3] At this point we do not impose any restriction on T except insisting it is first-order.

[4] Die-hard category theorists equivalently define a *sieve* on C to be any subfunctor of $\mathcal{Y}C$.

If S is a sieve on C and $f : C' \to C$, we define the *pullback of S along f* by $f^* := \{g : D \to C' : fg \in S\}$. The *maximal sieve* on C is the collection of all arrows with codomain C.

Definition 2.2.3. A *Grothendieck topology* on a category $\mathcal{C}$ is a function J assigning to each object C of $\mathcal{C}$ a collection $J(C)$ of *covering sieves* on C, such that for each object C:

1. $J(C)$ contains the the maximal sieve on C. (Maximality)
2. If $S \in J(C)$ and $f : C' \to C$, then $f^*S \in J(C')$. (Pullback stability)
3. If R is a sieve on C and $S \in J(C)$ such that for every $f \in S$ it holds that $f^*R \in J(\mathrm{dom}(f))$, then $R \in J(C)$. (Transitivity)

A pair $(\mathcal{C}, J)$ of a category equipped with a Grothendieck topology is called a *site*.

Example 2.2.4. • The *trivial topology* on a category $\mathcal{C}$ is the Grothendieck topology whose covering sieves are precisely the maximal sieves.

- Given a topological space X, consider the category $\mathcal{O}(X)$ of opens of X. There is a natural Grothendieck topology $J_{\mathcal{O}(X)}$ on $\mathcal{O}(X)$ where the covering sieves of an open U are precisely the open covers of U.

A site contains enough information to determine which presheaves are sheaves. We first need two more definitions.

Definition 2.2.5. Let P be a presheaf and let S be a sieve. A *matching family* assigns to each arrow $f : D \to C$ in S an element $x_f \in P(D)$ such that $P(g)(x_f) = x_{fg}$ for every arrow $g : E \to D$. An *amalgamation* for such a family is an $x \in P(C)$ such that $P(f)(x) = x_f$ for every $f \in S$.

Definition 2.2.6. A *sheaf* on a site $(\mathcal{C}, J)$ is a presheaf P in $\widehat{\mathcal{C}}$ such that every matching family on a covering sieve S in $J(C)$ for some object C of $\mathcal{C}$, has a unique amalgamation. The *category of sheaves* on $(\mathcal{C}, J)$ is the full subcategory of $\widehat{\mathcal{C}}$ given by the presheaves that are sheaves.

Finally, we are ready to define the our notion of topos.

Definition 2.2.7. A *Grothendieck topos* is any category of sheaves on some site.

Example 2.2.8. • The category **Set** is the the category of sheaves of the one point topological space under some topology.

- The category of presheaves on $\mathcal{C}$ is the category of sheaves on $\mathcal{C}$ equipped with the trivial topology.
- If X is a topological space, the category of sheaves on $(\mathcal{O}(X), J_{\mathcal{O}(X)})$ corresponds to the ordinary notion of sheaves of sets on the topological space X.

We shall use the following notion of morphism between toposes.

Definition 2.2.9. A *geometric morphism* $f : \mathcal{E} \to \mathcal{F}$ between toposes is a pair $f^* \dashv f_*$ of adjoint functors $f_* : \mathcal{E} \to \mathcal{F}$ and $f^* : \mathcal{F} \to \mathcal{E}$, such that f^* preserves limits.[5] The functors f_* and f^* are, respectively, called the *direct image* and the *inverse image* of f.

[5] Note that, as adjoints, f_* preserves all limits and f^* preserves all colimits.

The category **Topos** has as objects toposes and as morphisms geometric morphisms. If $\mathcal{E}$ and $\mathcal{F}$ are toposes, we denote by $\mathbf{Topos}(\mathcal{E}, \mathcal{F})$ the category with as objects geometric morphisms $\mathcal{E} \to \mathcal{F}$ and as morphisms from $f \to g$ natural transformations $f^* \to g^*$.

Recall the notion of a *geometric category* from Section 1.4. The following theorems show that we can interpret a geometric theory T in any Grothendieck toposes and, moreover, that models of T are preserved along the inverse image of a geometric functor.

Theorem 2.2.10. *Every topos is a geometric category.*

As a result we have for each geometric theory T and topos $\mathcal{E}$ a category $T\text{-mod}(\mathcal{E})$ of models of T in $\mathcal{E}$.

Theorem 2.2.11. *The inverse image of a geometric morphism is geometric.*

Hence for every geometric morphism $f : \mathcal{E} \to \mathcal{F}$, we get a geometric functor $f^* : T\text{-mod}(\mathcal{F}) \to T\text{-mod}(\mathcal{E})$ defined by composing a model $M : \mathcal{C}_T \to \mathcal{F}$ with the inverse image of f.

2.3 Classifying toposes

Definition 2.3.1. The *classifying topos* of a geometric theory T is a topos $\mathbf{Set}[T]$ such that there is an equivalence of categories

$$\mathbf{Topos}(\mathcal{E}, \mathbf{Set}[T]) \simeq T\text{-mod}(\mathcal{E}),$$

natural in $\mathcal{E}$.

The above naturality condition means that for any geometric morphism $\mathcal{E} \to \mathcal{F}$, the square

$$\begin{array}{ccc} \mathbf{Topos}(\mathcal{F}, \mathbf{Set}[T]) & \xrightarrow{\simeq} & T\text{-mod}(\mathcal{F}) \\ {\scriptstyle -\circ f}\big\downarrow & & \big\downarrow{\scriptstyle f^*} \\ \mathbf{Topos}(\mathcal{E}, \mathbf{Set}[T]) & \xrightarrow{\simeq} & T\text{-mod}(\mathcal{E}) \end{array}$$

commutes up to isomorphism.

The model U_T of T in $\mathbf{Set}[T]$ corresponding to $\mathrm{id}_{\mathbf{Set}[T]}$ under this equivalence is called the *universal model* of T. By taking $\mathcal{F}$ to be $\mathbf{Set}[T]$ in the above square, it can be seen that every geometric morphism $f : \mathcal{E} \to \mathbf{Set}[T]$ corresponds to the model $f^* U_T$ of T in $\mathcal{E}$ (up to isomorphism).

The classifying topos $\mathbf{Set}[T]$ is a representable object for the functor T-mod which sends a topos $\mathcal{E}$ to the category of models $T\text{-mod}(\mathcal{E})$. As a corollary of the (2-categorical) Yoneda Lemma, classifying toposes are unique up to equivalence.[6]

[6] Can this be seen in a more simple way?

Theorem 2.3.2. *Every geometric theory has a classifying topos.*

Proof sketch. The idea is to equip the syntactic category $\mathcal{C}_T$ with a Grothendieck topology J_T. Given an object $(\vec{y}.\psi)$, we let the covering sieves of J_T to be those sieves $\{\theta_i(\vec{x},\vec{y}) : (\vec{x}_i.\varphi_i) \to (\vec{y}.\psi) \mid i \in I\}$ such that

$$\psi \vdash_{\vec{y}} \bigvee_{i \in I} \exists \vec{x}_i.\theta_i$$

is provable in T. It can be proven that this $(\mathcal{C}_T, J_T)$ is a site and, moreover, that the topos of sheaves over $(\mathcal{C}_T, J_T)$ is the classifiyig topos for T. □

Theorem 2.3.3. *Every topos is the classifying topos of some geometric theory.*

Proof sketch. Let $\mathbf{Sh}(\mathcal{C}, J)$ be the category of sheaves over the site $(\mathcal{C}, J)$. The idea is to explicitly construct a geometric theory T for which $\mathbf{Sh}(\mathcal{C}, J)$ is the classifying topos. The signature for T consists of a type $\ulcorner C \urcorner$ for each object C of $\mathcal{C}$, a function symbol $\ulcorner f \urcorner$ for each arrow f of $\mathcal{C}$, and a relation symbol $\ulcorner R \urcorner$ for each subobject in $\mathcal{C}$.

In this signature, we let T be the theory of so-called *J-continuous flat functors on $\mathcal{C}$*. For any topos $\mathcal{E}$, a model of T in $\mathcal{E}$ corresponds to a J-continuous flat functor from $\mathcal{C}$ to $\mathcal{E}$. This correspondence lifts to an equivalence of categories, *i.e.* we have

$$\mathbf{Flat}_J(\mathcal{C}, \mathcal{E}) \simeq T\text{-mod}(\mathcal{E}).$$

By a result known a Diaconescu's theorem, it holds that

$$\mathbf{Topos}(\mathcal{E}, \mathbf{Sh}(\mathcal{C}, J)) \simeq \mathbf{Flat}_J(\mathcal{C}, \mathcal{F}),$$

hence we obtain the required result. □

Classifying toposes often have a more familiar shape, as illustrated by the following example.

Example 2.3.4. The theory T_{int} of intervals is formulated in the signature with one-sort, a single binary relation $\leqslant$, and two constants b and t. Its axioms are:

$$\begin{gathered}
\top \vdash_x x \leqslant x \\
x \leqslant y \wedge y \leqslant z \vdash_{x,y,z} x \leqslant z \\
x \leqslant y \wedge y \leqslant x \vdash_{x,y} x = y \\
\top \vdash_x x \leqslant t \wedge b \leqslant x \\
b = t \vdash \bot \\
\top \vdash_{x,y} x \leqslant y \vee y \leqslant x
\end{gathered}$$

The classifying topos of T_{int} is the category of simplicial sets.

3 Infinitary Logic

Christian Espindola's proof of an infinitary generalisation of Deligne's theorem [ESP20] has a number of moving parts: it is intuitionistic, infinitary, done using a category theoretic language and apparatus, and involves some cardinal assumptions (in [Esp19], this involves large cardinals; in [ESP20] this involves assumptions like $\kappa^{<\kappa} = \kappa$). Speaking personally, this can make it difficult for me to understand what is going on in the proof, even if the core idea is more or less clear.

3.1 A Tour in Known Lands

So I propose we take a number of steps back. Let us start somewhere really far back from this setting, and work our way up to it: *classical propositional first order-logic*. We will set up a blueprint for the kinds of questions we are interested in. Throughout I assume familiarity with the basics of this theory, as well as with Stone duality and the basic algebraic completeness one finds in this.

There are in general two ways to obtain completeness theorems for various logics. One, is the algebraic, or more or less syntactic way: we construct an *algebraic model* (of whatever kind) for our logical theory, and we use something like a free algebra, a term model, or a syntactic category, to do the job. Usually there is a clear and obvious choice for such models, which derives immediately from the axioms of our logic. In the case of the logic $\mathcal{L}(\omega)$, basic propositional logic with an infinite number of propositions, we will arrive at the concept of a *Boolean algebra*. We use them as models by considering valuations $v : \mathsf{Prop} \to \mathbf{B}$, which are lifted to the whole term algebra in the usual way, and write:

$$\mathbf{B} \vDash \phi$$

if and only if for all valuations v, $v(\phi) = 1$. To show completeness, we construct a Lindenbaum-Tarski algebra $\mathbb{F}(\omega)$ (called the *free Boolean algebra on ω-generators*, by taking the term algebra $\mathbf{Tm}(\omega)$ and quotienting it by our axioms, i.e., saying that the pair $(\phi \wedge \phi \to \psi, \phi \wedge \psi)$, interpreted as meaning $\phi \wedge \phi \to \psi \vdash \phi \wedge \psi$, should be equal. This happens to be the unique countable atomless Boolean algebra). So far all good. We will refer to these as *Algebraic completeness theorems*.

However, despite the best efforts of algebraic logicians and set theorists alike, human beings seemingly do not cope very well with Boolean valued semantics. So we are normally also interested in a second kind of completeness theorem, which for reasons that will be explained later, I will refer to as a *Relational Completeness Theorem*.

In the propositional case this is very simple: a model is a valuation $v : \mathsf{Prop} \to \{0,1\}$, which we lift through the usual inductive clauses to all formulas, generating $\overline{v} : \mathsf{Form} \to \{0,1\}$. Surely our logic for these models is sound as well; so we only have to deal with completeness. Let us see how this is done in detail:

Proposition 3.1.1. Propositional logic is relationally complete.

Proof. The proof follows the following steps:

- We assume that $\Gamma \not\vdash \phi$, and hence $\Gamma \cup \{\neg\phi\}$ is a consistent set of formulas.
- Taking the equivalence classes of these formulas, we have that $[\Gamma] \cup \{[\neg\phi]\}$ is then a subset of $\mathbf{B}$, with the property that it generates a *proper filter*.
- By the *prime filter theorem*, this can be extended to a prime filter, which on a Boolean algebra is an ultrafilter, i.e., a Boolean homomorphism $v : \mathbb{F}(\omega) \to \{0,1\}$. Composing this map with the identity map from $\mathsf{Prop} \to \mathbb{F}(\omega)$ assigning to each variable its value, provides the valuation we wanted.

□

Essentially, the key to going from an *algebraic* to a *relational* completeness theorem is, as Kristoff pointed out last week, to be able to extract 2-valued semantics from Boolean valued semantics. For that it is useful to think about what kinds of Boolean algebras we are working with.

Definition 3.1.2. Let $\mathbf{B}$ be a Boolean algebra. We say that $\mathbf{B}$ is *complete* if it is complete as a lattice. We say that it is κ-complete if for I such that $|I| < \kappa$ and $(a_i)_{i\in I}$ then $\bigwedge_{i\in I} a_i$ exists. We say that it is (κ,λ)-*distributive* if it satisfies the following: for any sets I and J such that $|I| \leqslant \kappa$ and $|J| \leqslant \lambda$ and for any family $(a_{i,j})_{i,j}$ of elements in $\mathbf{B}$ we have:

$$\bigwedge_{i\in I}\bigvee_{j\in J} a_{i,j} = \bigvee\{\bigwedge_{i\in I} a_{i,f(i)} : f \in J^I\}$$

We say that $\mathbf{B}$ is *atomic* if every elements lies above an atom.

Complete and atomic Boolean algebras are simply those of the form $\mathcal{P}(X)$ for a given set X. Whilst they form our prototype of a Boolean algebra, not all Boolean algebras are of this form, as for instance $\mathbb{F}(\omega)$ is not atomic, and also not complete.

A consequence of the proof of the relational consequence theorem we gave above is that we get a sharper algebraic completeness theorem: *classical logic is sound and complete with respect to CABAs*. To see why, note that Stone duality, which was implicitly used above, gives us a representation of our Boolean algebra as a subalgebra of a complete and atomic Boolean algebra, (namely, the algebra $\mathcal{P}(X)$ where X is the Stone space dual). Since subalgebras preserve validity of equations, this yields completeness.

Hence, the above blueprint gives us completeness results with respect to complete and atomic Boolean algebras. Interestingly for our purposes we have the following:

Proposition 3.1.3. A complete Boolean algebra is atomic if and only if it is completely distributive.

To see why this matters, recall that by a classic result, CABA's are dual to the category of set, a mapping which identifies the CABA with its set of atoms. Atoms are, from this point of view the *complete ultrafilters*: ultrafilters which are closed under all meets. Atomicity, in other words, says that if b is an element of a CABA, it is contained in such a complete ultrafilter. Hence we have an interesting relationship between two kinds of representation theorems; we also add the case of ortholattices, which is interesting on the other extreme:

- Stone's theorem says that in all Boolean algebras, every element is contained in some ultrafilter;
- In CABA's this can be strengthened to a complete ultrafilter.
- In Ortholattices (i.e., Boolean algebras without distributivity), the lack of distributivity means the prime filter theorem does not go through. Hence one needs to either use *all filters* [Gol74; Gol75], or a cleverly selected collection of such filters.

Now let us think about infinitary propositional logic, and the question of how we should generalise this. In light of the above, the answer seems clear: we need a notion of a κ-complete ultrafilter, and to prove an analogue of Stone's theorem which works for larger κ. However it is clear that this cannot work in general; Keisler and Tarski [KT64] showed that these are large cardinal assumptions. This certainly looks bad, but it might not be so damning. After all, we do not need all Boolean algebras to be embeddable in a CABA; only the ones we are specifically interested in. Let us introduce some terminology:

Definition 3.1.4. Let $\mathbf{A}$ be a Boolean algebra. We say that this is an *κ-algebra of sets* if it is a *κ-subalgebra* of a power set algebra.

Stone's theorem then says:

Theorem 3.1.5. *Every Boolean algebra is an ω-algebra of sets.*

However, there are well-known examples of ω_1-Boolean algebras (normally denoted as σ-algebras) which are not σ-algebras of sets: an example is the σ-algebra $[0,1]$ modulo the ideal generated by the sets of Lebesgue measure zero. However, this is again not the center of the problem. To see why let us see now the proof of completeness for $\mathcal{L}_{\omega_1}(\omega)$; this is the logic which only extends propositional logic by the obvious infinitary discharge and introduction rules, as well as ω_1

many proposition letters. It is here that we see the use of the famous Rasiowa-Sikorski lemma; we will need it in the following form:

Lemma 3.1.6. Assume that **B** is a Boolean algebra, and $Q = (\{X_n\})_{n\in\omega}$ is a countable collection of countable subsets of **B** such that $\bigwedge X_n \in B$. Then for each $a \in B$ there is an ultrafilter P, called a *Q-filter* such that:

1. $a \in P$.

2. For each n, $X_n \subseteq P$ if and only if $\bigwedge X_n \in P$.

Proposition 3.1.7. The logic $\mathcal{L}_{\omega_1}(\omega)$ is algebraically and relationally complete.

Proof. The algebraic proof of completeness proceeds in exactly the same way as before, except we now look at σ-complete Boolean algebras, and generate a free σ-complete Boolean algebra (the same method applies regardless), call it $\mathbb{F}()\sigma(\omega)$.

As our relational models, we consider valuations $v : Prop_{\omega_1} \to \{0,1\}$ lifted to the algebra of formulas. It is trivial to show soundness. For completeness, suppose that $\nvdash_{\omega_1} \phi$; then consider a countably infinitary term algebra $\mathbf{Tm}_\phi$, constructed from ϕ, including all subformulas of ϕ, and closed *only under finitary Boolean operations*. Take the quotient under derivability in the same way, and note that if $\bigwedge_{n\in\omega} \psi_n$ is a subformula of ϕ, we have:

$$[\bigwedge_{n\in\omega} \psi_n] = \bigwedge_{n\in\omega} [\psi_n].$$

Now by construction ϕ may contain only countably many subformulas, so we can construct a collection $Q = (\{X_n\})_{n\in\omega}$ of all sequences of formulas appearing in ψ_n; so by the Rasiowa-Sikorski lemma, there is a Q-filter containing ϕ. By the same arguments as before, this generates then a valuation $v : Prop_{\omega_1}(\phi) \to \{0,1\}$, which we can extend arbitrarily to propositions not ocurring in ϕ, and this gives us the desired model. □

The proof also yields:

Corollary 3.1.8. $\mathcal{L}_{\omega_1}$ *is sound and complete with respect to CABA.*

Proof. Soundness is obvious. As for completeness, given the algebra $\mathfrak{F}(\phi)$ as above, let X' be the set:

$$\{x : x \text{ is a } Q\text{-filter }\}. \tag{3.1}$$

The fact that each such filter preserves the collection Q means that $\mathfrak{F}(\phi)$ embeds into $\mathcal{P}(X')$ which preserves all meets and joins ocurring in Q. By induction we can then show that this implies that $\mathcal{P}(X')$ refutes the formula ϕ, as desired. □

Hence the Rasiowa and Sikorski lemma allows us *just enough primeness in our filters* to prove the same kind of completeness theorem. As is well-known though, the Rasiowa and Sikorski lemma is intimately related to Martin's Axiom, so one could be reasonably skeptical about how this can be obtained for larger cardinals without invoking either large cardinals or a forcing axiom. This brings us to the concept of *representability*:

Definition 3.1.9. Let **A** be a κ-complete Boolean algebra. We say that **A** is κ-representable if there is an algebra of sets **B** and a κ-surjective homomorphism (i.e., preserving κ-infinitary operations) $f : \mathbf{B} \to \mathbf{A}$.

The key property of representability lies in the following, originally showed by Chang [Cha57]:

Proposition 3.1.10. Let **B** be a κ-representable Boolean algebra. Then whenever $a \in B$, and $Q = (\{X_\alpha\})_{\alpha<\lambda}$ is a collection of $\lambda \leqslant \kappa$ sets of elements which meet belongs to **B**, then there exists an ultrafilter of **B** containing a and preserving the meets in Q.

Proof. See the Appendix. □

By the above proof, the generalisation of completeness now reveals itself obvious, *if we can find some logical property implying representability*. The following laws were found by C.C. Chang [Cha57] who credits them in part to Tarski, and indeed do the job:

Definition 3.1.11. Let γ be an infinite cardinal number. We denote by Π_γ the γ-*Chang law*, for each family of formulas $\{A_\varepsilon : \varepsilon < \gamma\}$:

$$\bigvee_{\mu<\gamma} \bigwedge_{\eta<\gamma} A_{\mu,\eta}$$

where $\{A_{\mu,\eta} : \mu, \eta < \gamma\}$ is a family of formulas such that for each μ, η there is some $\varepsilon < \gamma$ such that $A_{\mu,\eta} = A_\varepsilon$ or $A_{\mu,\eta} = \neg A_\varepsilon$ and for all $g \in \gamma^\gamma$ there is a $\varepsilon < \gamma$ such that $\{A_\varepsilon, \neg A_\varepsilon\} \subseteq \{A_{\mu,g(\mu)} : \mu < \gamma\}$.

To see why this is needed, consider an arbitrary model. If the law were false under this model, then for each $\mu < \gamma$, there would be $\eta < \gamma$ such that $\neg A_{\mu,\eta}$; this defines a function $g \in \gamma^\gamma$, and such formulas now appear as $\bigwedge \neg A_{\mu,g(\mu)}$; but by assumption, there is a contradictory pair here, so this leads to the model satisfying a contradiction.

The following establishes the sanity of this condition (see Karp [Kar64, Theorem 6.4.4]):

Proposition 3.1.12. A Boolean algebra **B** is κ-representable if and only if it validates the κ-Chang's law.

Using the former we can derive a completeness theorem for any regular cardinal κ. Hence let $\mathcal{L}_{\kappa^+}$ denote the infinitary propositional logic which includes, in addition to the usual rules and axioms, also the κ-Chang law. Then:

Corollary 3.1.13. *The logic $\mathcal{L}_{\kappa^+}$ is algebraic and relationally sound and complete.*

Proof. We provide algebraic completeness with respect to κ-complete and κ-representable Boolean algebras. To prove relational completeness, assume that $\nvdash \phi$; then construct the free Boolean algebra generated by the at most κ-many infinite formulas occurring in ϕ, and otherwise closed for Boolean operations, just as before. Note that since ϕ

must contain at most κ formulas, this ensures that this Boolean algebra will have size at most κ, and hence, that we can collect all infinite meets $(\{X_\alpha\}_{\alpha<\kappa})$ ocurring in the algebra. Then proceed exactly as before, using κ-representability instead of the Rasiowa-Sikorski lemma. □

There is a deep connection between representability and distributivity. Indeed, in sufficient amounts, they each imply the other. One particularly sharp relationship is the following:

Proposition 3.1.14. Let $\mathbf{B}$ be a Boolean algebra. Then if $\mathbf{B}$ is κ^+-complete Boolean algebra which is $(\kappa, 2)$-distributive. Then $\mathbf{B}$ is κ-representable. In turn if $\mathbf{B}$ is $(2^\kappa)^+$-representable, then it is $(\kappa, 2)$-distributive.

Proof. See for example [KMB89, Proposition 14.12]. □

The key problem of working with the distributivity law, rather than Chang's laws, is that the $(\kappa, 2)$-distributivity law, in order to be included in a calculus, requires a conjunction of size κ^+. This is why in general, the distributivity laws do not suffice for obtaining completeness. This is also a hint for why very complicated distributivity-like laws can come in handy.

One final "logic" one might consider, and which is relevant for our discussion of geometric logic, is the system $\mathcal{L}_\infty$, consisting of a proper class of propositional variables, axioms and together with all the Chang distributivity laws (or equivalently all the distributivity laws). Given what we showed before, it follows that:

Proposition 3.1.15. $\mathcal{L}_\infty$ is algebraically sound and complete with respect to CABA's.

With this preamble in mind, we are ready to take an upgrading move to the setting of *first-order logic*.

3.2 *Infinitary First-Order Classical Logic*

When moving to first-order logic, we are suddenly faced with many complications, and we will seek to avoid these as much as possible. For a more in-depth coverage of this see [Mar02; Dic85; Kar64].

In this setting, our *relational* models are no longer simply valuations, but rather entire classes of structures, and their semantics is the expected one. One thing which we should note is that in this setting we can have infinitely long terms and relational formulas. For simplicity we will only tackle the case of $\mathcal{L}_{\kappa^+,\kappa}$ where terms and such formulas are of size less than κ (this is also Espíndola's setting in his papers).

In addition to this, we of course add quantifiers of the form:

$$\exists_{\alpha<\lambda} x_\alpha \text{ and } \forall_{\alpha<\lambda} x_\alpha.$$

The derivation system for this logic includes all instances of axioms and rules for $\mathcal{L}_{\kappa^+}$, introduction and elimination rules for existential and universal quantifiers. Additionally, we require a specific rule, which allows us to cope with the loss of compactness with respect to quantifiers. This rule is often called the rule of "Dependent choices", and it goes as follows: for each collection of formulas $(\psi_\alpha(v_\alpha)_{\alpha<\kappa}$

$$\text{If } \vdash \bigwedge_{\alpha<\kappa} \exists v_0 \psi_0(v_0) \wedge ... \wedge \forall_{\eta<\lambda} v_\eta \exists v_\lambda \psi_\lambda(v_\lambda) \text{ then } \vdash \exists_{\alpha<\kappa} v_\alpha \psi(v_\alpha).$$

Provided that the sequences of variables have disjoint range, and variables in v_α do not appear free in ψ_β for $\beta < \alpha$. The intuition behind this rule lies in the set theoretic axiom: if we can pick some elements to satisfy a formula, and from those elements pick some more elements to validate another formula, then we should be able to construct a single sequence satisfying the whole sequence of formulas. The contrapositive of the rule, however, will be more useful below:

$$\text{If } \vdash \forall x_\alpha \bigvee_{\alpha<\kappa} \psi(v_\alpha) \text{ then } \vdash \bigvee_{\alpha<\kappa} \forall v_0 \psi_0(v_0) \vee ... \vee \exists_{\eta<\lambda} v_\eta \forall v_\lambda \psi_\lambda(v_\lambda).$$

Let us now move on to the issue of completeness of these systems. In the case of first-order logic what we have called "algebraic completeness" becomes somewhat of a misnomer; if one wanted a genuinely "algebraic" proof this would be the domain of so-called "cylindric algebras". However, the theory of cylindric algebras is very different from the usual Boolean, Intuitionistic, and even Modal, logical setting, as it is vastly more complex and filled with subtleties. The approach taken here thus focuses more on the Henkin-Tarski relational completeness.

As such, the natural option is to assemble a "term model", and use this to serve as our generic model for a given theory. But this reveals us the need for set-theoretic assumptions. To see why, note that the way this is done for the logic $\mathcal{L}_{\kappa^+,\kappa,}$ is as follows: given a formula ϕ, let X be a collection of κ^+ many symbols not ocurring in ϕ, and let $T_0(\phi)$ be the collection of all symbols in X and all constants appearing in ϕ. Let Δ_0 be the collection of all substitutions of all subformulas of ϕ for terms in T_0. Note that ϕ can have at most κ many subformulas; each subformula must have fewer than κ free variables, say δ, and the set of terms in this language is of size κ^+ (given it includes the whole of X), hence there are $(\kappa^+)^{<\kappa}$ many possible substitutions. Since by assumption, $\kappa^{<\kappa} = \kappa^+$, then the above is surely also κ^+. Then we construct, by mutual induction the sets:

- $T_\gamma(\phi)$, consisting of all terms in atomic formulas of $\bigcup_{\xi<\gamma} \Delta_\xi$.
- Δ_γ, consisting of all substitutions of subformulas of A for terms in $\bigcup_{\eta<\gamma} T_\eta$.

We then construct $T(\phi)$ as the union of these sets over κ, and let Δ be the union over Δ_γ together with the set $\{f = g : f, g \in T(\phi)\}$. Note that the set-theoretic assumption gives us that all of these sets remain firmly of cardinality at most κ, and all formulas involved contain fewer than κ many variables.

The former is then why the assumption:

$$\kappa^{<\kappa} = \kappa$$

Appears naturally in the setting of first-order logic. With it, using Karp's result, one can prove the following:

Theorem 3.2.1. *Assume that $\kappa^{<\kappa} = \kappa$. Then the calculus $\mathcal{L}_{\kappa^+,\kappa}$ is relationally sound and complete.*

Proof. See Appendix. □

Additionally, we remark a few facts about this kind of completeness result:

- If we were looking instead at $\mathcal{L}_{\kappa^+,\omega}$, the hypothesis that $\kappa^{<\kappa} = \kappa$ is not necessary;
- If κ is strongly inaccessible, one can then show that $\mathcal{L}_{\kappa,\kappa}$ is sound and complete;
- However, we mention that, for example, $\mathcal{L}_{\omega_2,\omega_2}$ was shown to be incomplete [Kar64, Chapter 12], and like it all logics of the form $\mathcal{L}_{\kappa,\kappa}$ when κ is a successor.
- The logic $\mathcal{L}_{\infty,\kappa}$ is sound and complete, without any extra assumptions, by picking a sufficiently large first coordinate, and by using similar arguments as those sketched here.

We stress then that apart from the need for cardinal assumptions for some cases, and the rule of dependent choices, this setting still very much seems quite tamable. In the next section we will see how this changes when we move the setting to intuitionistic logic.

Thus far, the systems we have looked at have remained wholly classical. We will move to a different set of logics, and thus we need to adopt a few changes:

- We lose interdefinability of the connectives, so we are forced to work with all connectives. Additionally, for the majority of cases *we do not use the implication*. This is important, since the intuitionistic implication, unlike all other connectives, has a highly modal flavour, and hence, tends to require a strong form of prime filter separation to be validated (we will later see how this can be dropped in the special case of completeness we will be interested in).
- Our relational models – even in the propositional case – have to be more complex than in the classical case.
- Heyting algebras are less symmetric structures, which means we lose the ability to use structures like ideals with ease.

Additionally, as we will see, to make the results go through we often need to restrict to fragments containing different sizes of conjunction and disjunction, and even allowing arbitrarily large disjunctions. This introduces further degrees of freedom which we will encounter in the next few pages.

Definition 3.2.2. The language of propositional *κ-coherent logic*, denoted $\mathcal{L}_\kappa^{coh}$ is composed of proposition symbols, as well as conjunctions and disjunctions of size less than κ. The language of propositional *κ-geometric logic*, denoted $\mathcal{L}_\kappa^{g}$ is composed of arbitrary disjunctions and conunctions of size less than κ.

3.3 Algebraic and Relational Models

The algebraic models we will be mostly concerned with are *Distributive Lattices* and *Heyting algebras*. For general references see for instance [CZ97; Dav79; Esa19]. In the finitary case, ω-coherent logic is sometimes called *positive logic*, due to the lack of any form of negation (whether classical or otherwise) present in the language. The usual proof of algebraic completeness is routine, using the same technique of generating the free algebra. As for the relational completeness, we consider *Kripke semantics*[1]:

[1] Note that we use here the more usual semantics for positive and intuitionistic logic – including all posets, rather than just trees. This difference is immaterial for most purposes, and we think it helps connect the infinitary cases with the more usual duality-laden finitary cases.

Definition 3.3.1. (Propositional Kripke semantics) A positive *Kripke frame* consists of a tuple $\mathfrak{F} = (K, \leqslant)$ where $\leqslant$ is a partial order. A *propositional Kripke model* consists of a pair $\mathfrak{M} = (\mathfrak{F}, V)$ where $\mathfrak{F}$ is a Kripke frame and V is a function $V : \mathsf{Prop} \to \mathsf{Up}(K)$. We call a pair $(\mathfrak{M}, w)$ of a Kripke model and a world $w \in K$ a *Pointed Kripke model*. We define a forcing relation on this structure, $\Vdash$, as follows:

- $\mathfrak{M}, w \nVdash \bot$.
- $\mathfrak{M}, w \Vdash p$ if and only if $w \in V(p)$;
- $\mathfrak{M}, w \Vdash \bigwedge_{i \in I} \phi_i$ if and only if $\mathfrak{M}, w \Vdash \phi_i$ for each i;
- $\mathfrak{M}, w \Vdash \bigvee_{i \in I} \phi_i$ if and only if $\mathfrak{M}, w \Vdash \phi_i$ for some i.

We define propositional *intuitionistic* models by extending the clause for implication as:

- $\mathfrak{M}, w \Vdash \phi \to \psi$ if and only if whenever $w \leqslant v$ and $\mathfrak{M}, v \Vdash \phi$ then $\mathfrak{M}, v \Vdash \psi$.

Definition 3.3.2. (First-order Kripke semantics) A *first-order Kripke model* is a quadruple $\mathcal{B} = (K, \leqslant, D, V)$ where $(K, \leqslant)$ is a Kripke frame frame, D is a functor from K to **Set**, and V is a valuation on each of the models, over a fixed collection of relation and function symbols, and with added constants from $\bigcup_{k \in K} D_k$, which is required to be *persistent*, i.e., the following diagram should commute for each n-ary function symbol and each relation symbol, as described in Figure 3.1:

$$\begin{tikzcd} D_k^n \arrow[r, "g_{D_k}"] \arrow[d] & D_k \arrow[d, "D(k \leqslant m)"] \\ D_m^n \arrow[r, "g_{D_m}"'] & D_m \end{tikzcd}$$

Figure 3.1: Persistence of Models

For each model $(D(k), V)$ we define the $\vDash$-semantic satisfaction relation as usual. We define the positive satisfaction relation on $\mathcal{B}$, $\Vdash$, as follows:

- For $\phi(\overline{x})$ a closed atomic formula with constants in the language, and $k \in K$, $\mathfrak{M}, k \Vdash \phi(\overline{x})$ if and only if $D(k) \vDash \phi(\overline{x})$.

- The propositional clauses of satisfaction as before;

- $\mathfrak{M}, k \Vdash \exists_{\lambda<\alpha} x_\alpha \phi(\overline{x}_{\alpha<\lambda}, \overline{y})$ if and only if $\mathfrak{M}, k \Vdash \phi(c_{\alpha<\lambda}, \overline{y})$ for some sequence of elements in $D(k)$.

We further consider *intuitionistic Kripke models* by adding the clause for the universal:

- $\mathfrak{M}, k \Vdash \forall_{\alpha<\lambda} x_\alpha \phi(\overline{x}_{\alpha<\lambda}, \overline{y})$ if and only if if $k \leqslant t$ then for all $(c_\alpha)_{\alpha<\lambda}$ in $D(t)$, $\mathfrak{M}, t \Vdash \phi(c_\alpha, \overline{y})$.

In both the propositional and first-order cases, the intuition behind the persistence condition is as follows: whereas classical logic imagines truth as being modelled absolutely, positive and intuitionistic logic understand knowledge as something which can be acquired as more facts are discovered, i.e., as one progresses along a Kripke model. It is a well-known fact in the theory of intuitionistic and positive logic that nothing is lost when considering *complete partial orders*; in that case, we can consider the leaf-nodes of the model as eventual classical models, which always exist above any non-classical model.

Like before, we will begin by analysing the propositional case, and then return to the first-order case once the propositional complexities have been tamed.

The completeness of basic positive logic follows a very similar strategy as the classical case before it. For the relational completeness, we note that the ultrafilter lemma admits an extension to a *prime filter lemma*, saying that whenever F is a filter and I an ideal such that $F \cap I = \varnothing$, then there is a prime filter $F' \supseteq F$ such that $F' \cap I = \varnothing$. We have the following analogue of a Stone space that makes the situation very close:

Definition 3.3.3. Let $(X, \leqslant, \tau)$ be a partially ordered topological space. We say that this is a *Priestley space* if:

1. (X, τ) is a compact topological space;

2. (Priestley separation condition) Whenever $x \nleqslant y$ there is a clopen upwards closed subset $U \subseteq X$ such that $x \in U$ and $y \notin U$;

The following is shown for instance in [DP02].

Proposition 3.3.4. The category of distributive lattices is dually equivalent to the category of Priestley spaces. In particular, each distributive lattice $\mathbf{H}$ is isomorphic to a distributive lattice of sets of the form $\mathsf{ClopUp}(X_H)$ where X_H is its dual Priestley space.

Proposition 3.3.5. Propositional logic $\mathcal{L}^{int}_{\omega}$ is algebraically and relationally complete.

Proof. The proof is identical to the classical case, except we now use the prime filter theorem in the form specified above. □

Many notions defined before generalise with obvious modifications – for instance the (κ, μ)-distributive laws, or the notion of κ-representability, given these only involve the lattice language. One aspect which has a less obvious generalisation is *atomicity*. Indeed, whilst the concept of "atomic distributive lattice" is certainly sensible, complete and atomic Distributive lattices do not play the same role as their Boolean counterparts. The relevant notion is the following:

Definition 3.3.6. Let $\mathbf{H}$ be a distributive lattice. We say that a filter F on $\mathbf{H}$ is *completely join-prime* if $\bigvee_{i \in I} a_i \in F$ if and only if there is some $i \in I$ and $a_i \in F$. We say that an element $a \in H$ is completely prime if $\uparrow a$ is a completely join-prime element.

We say that an element $a \in H$ is *completely meet-prime* if whenever $\bigwedge_{i \in I} b_i \leqslant a$ then for some $i \in I$, $b_i \leqslant a$. We say that a pair of elements (p, q) where p is completely join-prime and q is completely meet-prime is a *splitting pair* if $\uparrow p \cap \downarrow q = \varnothing$ and $\uparrow p \cup \downarrow q = H$.

We say that $\mathbf{H}$ is *completely join-prime generated* if every element is the join of completely prime elements.

We note the following well-known characterization:

Proposition 3.3.7. Let $\mathbf{H}$ be a complete distributive lattice. Then the following are equivalent:

1. $\mathbf{H}$ is completely join-prime generated.
2. $\mathbf{H}$ is isomorphic to $\mathsf{Up}(X)$, the set of upwards closed subsets of a partially ordered set $(X, \leqslant)$.
3. For each $a, b \in H$, if $a \nleqslant b$, there is a splitting pair (p, q) such that $p \leqslant a$ and $b \leqslant q$.

Proof. (1) implies (2). If $\mathbf{H}$ is completely join-prime generated, consider P_∞ the set of completely join prime elements, ordered by inclusion. Let $\mathsf{Up}(P_\infty)$ denote the upwards closed subsets of P_∞. Define:

$$\begin{aligned} \phi : H &\to \mathsf{Up}(P_\infty) \\ a &\mapsto \{x : x \leqslant a\} \end{aligned}$$

The hypothesis of complete generation ensures that this is a complete embedding. It is easy to see that it is onto since the algebras are both complete.

(2) implies (3): assume that U and V are upwards closed subsets and $U \nsubseteq V$; then there is some $x \in U$ such that $x \notin V$. Note that $\uparrow x$ is an element of the lattice, and it is completely join prime. Similarly, $X - \downarrow x$ is a completely meet prime element, and $V \subseteq X - \downarrow x$ (since V is upwards closed). So $(\uparrow x, X - \downarrow x)$ is a splitting pair: if $\uparrow x \subseteq U \subseteq X - \downarrow x$ we surely get a contradiction, and given any W, either $x \in W$, in which case, $\uparrow x \subseteq W$, or $x \notin W$, hence $W \subseteq X - \downarrow x$ (given otherwise $x \in W$ by upwards closure). This gives us the desired splitting pair.

(3) implies (1): Given an arbitrary element $a \in H$, we claim that it is the join of all the completely join-prime elements below it. Indeed,

assume not; then $a \nleqslant \bigvee_{c \leqslant a} c$, where the latter is the join of all completely join primes. Then by assumption, let (p, q) be a splitting pair. It is easy to see that p is a completely join-prime element; but then $\bigvee_{c \leqslant a} c$ is contained in $\uparrow p \cap \downarrow q$, a contradiction. □

Hence, the algebras of the form $\mathsf{Up}(X)$ are, in the setting of positive logic, the correct generalisation of the power set algebras. For ease of reference, and in light of the result, we refer to the above as *Splitting Algebras*. One can note that the above completeness theorem also showed completeness with respect to these kinds of algebras, and in general, we will want this from our completeness theorems in analogy with the classical case. On the other hand, we can see that we cannot obtain this through the usual distributive laws, since complete distributivity is not enough to guarantee one is a splitting Heyting algebra:

Example 3.3.8. Consider $[0, 1]$ with complete meet given by infimum and complete join given by supremum. Notice that it is completely distributive: if $(x_{i,j})_{i \in I, j \in J}$ is a collection of elements, then note that if z is equal to $\bigvee \{\bigwedge_{i \in I} x_{i,f(i)} : f \in J^I\}$, then z is the supremum element amongst the f, of the infima of $x_{i,f(i)}$. Let i be arbitrary. Then note that for each f such that j is in the range of f, $\bigwedge_{i \in I} x_{i,f(i)} \leqslant x_{i,j}$; hence the supremum of the former is less than or equal to $\bigvee_{j \in J} x_{i,j}$, since all j will be in the range of some function. Hence $z \leqslant \bigvee_{j \in J} x_{i,j}$. This suffices to show complete distributivity.

To see that this algebra is not splitting, note that it contains *no* completely join prime elements other than 0; indeed, given any element, we can consider it as the supremum of the elements coming below it.

Indeed, in this setting, complete distributivity is rather equivalent to *complete representability*. This follows by a result due to Raney [Ran52]:

Lemma 3.3.9. Let $\mathbf{H}$ be a complete distributive lattice. Then the following are equivalent:

1. $\mathbf{H}$ is a complete homomorphic image of a completely join-prime generated algebra $\mathbf{H}'$;
2. $\mathbf{H}$ is completely distributive.

Proof. See Appendix. □

Hence in this respect the situation is distinct from the classical case. κ-distributive laws cannot ensure us completeness with respect to algebras of sets, given if that were the case, the above counterexample would not exist.

The notion of a Q-set adapts to our setting, with minimal changes:

Definition 3.3.10. Let $\mathbf{H}$ be a κ^+-complete distributive lattice, and $Q = (\{X_\alpha\}, \{Y_\alpha\})_{\alpha < \kappa}$ be subsets of size at most κ, such that $\bigwedge X_\alpha$ and $\bigvee Y_\alpha$ exist for each α. We say that a subset $F \subseteq H$ is a Q-filter if it is a prime filter and additionally:

- $X_\alpha \subseteq F$ if and only if $\bigwedge X_\alpha \in F$;
- $\bigvee Y_\alpha \in F$ if and only if $Y_\alpha \cap F \neq \emptyset$.

However, again there is a priori no device that ensures that enough Q-filters exist. What we would need is some rule or law that ensured the existence of enough such points. It is here that we encounter Espindola's "Transfinite Transitivity Rule".

3.4 *Propositional Transfinite Transitivity Rule*

We have so far seen a number of distributivity properties that a given lattice might enjoy. The last section gave us a taste for the difficulties of dealing with this setting. In this section we discuss C. Espindola's Transfinite Transitivity rule, as a device enforcing strong distributivity.

Definition 3.4.1. Let $\mathbf{H}$ be a distributive lattice. We say that $\mathbf{H}$ is TT_κ-distributive (or satisfies the *Propositional TT_κ rule*) if it is κ-complete, possibly with a collection of κ-many κ-joins, if for each $\gamma < \kappa^+$ and all elements $\{a_f : f \in \gamma^{<\kappa}\}$ such that:

$$a_f = \bigvee_{g \in \gamma^{\beta+1}, g\restriction_\beta = f} a_g$$

for all $f \in \gamma^\beta$, $\beta < \kappa$ and:

$$a_f = \bigwedge_{\alpha<\beta} a_{f\restriction_\alpha}$$

for all limit β, we have that $\bigvee_{f \in B} \bigwedge_{\beta < \delta_f} a_f$ exists and is equal to $a_\emptyset$ where B is the set of minimal elements of a bar in $\gamma^{<\kappa}$.

We denote the TT_κ-rule as follows:

$$\frac{\begin{array}{c}\phi_f \vdash \bigvee_{g\in\gamma^{\beta+1}, g\restriction_\beta = f} \phi_g,\ \beta < \kappa, f \in \gamma^\beta \\ \phi_f \dashv\vdash \bigwedge_{\alpha<\beta} \phi_{f\restriction_\alpha}, \beta < \kappa,\ \text{limit } \beta,\ f \in \gamma^\beta\end{array}}{\phi_\emptyset \vdash \bigvee_{f\in B} \bigwedge_{\beta<\delta_f} \phi_{f\restriction_\beta}}$$

where $\gamma \leqslant \kappa$, $B \subseteq \gamma^{<\kappa}$ is a collection of the minimal elements of a given bar, and the $\delta_f < \kappa$ are the levels of $f \in B$, and all meets and joins are assumed to exist.

We say that a complete distributive lattice $\mathbf{H}$ satisfies the *Complete Propositional TT* rule if it is TT_κ-distributive for all κ.

We note some aspects of the former rule: first of all, note that it involves only conjunctions of size less than κ, and disjunctions of size at most κ. Moreover, the conclusion seems somewhat similar to the conclusion of the (κ, κ)-distributivity law, if only visually.

In his papers, C. Espindola's motivates the the former rule as something one ought to have if completeness is to be possible, by taking the situation where we have been able to establish (somehow) a representation using cleverly chosen points, and demonstrating how this rule needs to be satisfied. Another demonstration of this sort (proven by him in a MathOverflow question) shows that the TT-rule is enough to ensure a strong form of representability; we reproduce here the argument as this helps to illustrate the uses of the rule:

Proposition 3.4.2. Let **H** be a complete distributive lattice. Then the following are equivalent:

1. **H** is a Splitting algebra;

2. **H** satisfies the Complete Propositional TT-rule.

Proof. First assume that **H** does not satisfy the complete Propositional TT-rule; then there is a tree of elements in the conditions of the antecedent, and a root $\phi_\varnothing$ which is not below $\bigvee_{f \in B} \bigwedge_{\beta < \delta_f} \phi_{f \restriction_\beta}$ for any bar B. Pick an arbitrary bar, and use the splitting property to find a completely join prime element $x \leqslant \phi_\varnothing$ which is not below the given join; then using the complete join-primeness and the construction of the tree, we can construct a branch, ensuring that indeed x is eventually below $\bigwedge_{\beta < \delta_f} \phi_{f \restriction_\beta}$, a contradiction.

Conversely, assume that the TT-rule is valid. Let $\kappa = (2^\delta)^+$ where $|H| = \delta$; for each $a \in H$, let $C(a)$ denote the collection of all sequences $(b_\alpha)_{\alpha < \lambda}$ for $\lambda \leqslant \kappa$ where $a = \bigvee_{\alpha < \lambda} b_\alpha$. Now let $f : \kappa \times \kappa \to \kappa$ be the canonical well-ordering of κ, with the property that $f(\alpha, \gamma) \geqslant \gamma$.

Now let c be an arbitrary element. We will construct a tree of height κ. Let $\phi_\varnothing = c$. For each β, assume that the tree has been defined on level β, and we specify it at level $\beta + 1$, by considering a node p at level β, such that we will outline how to construct its successors. Indeed, if $f(\alpha, \gamma) = \beta$, note that the predecessors at level γ will have been defined for p; so take the α-th tuple $(b_{\eta < \lambda}) \in C(m)$ where m is the (unique) predecessor of p in the level γ, and note that then:

$$p \leqslant p \wedge m = p \wedge \bigvee_{\eta < \lambda} b_\eta = \bigvee_{\eta < \lambda} p \wedge b_\eta$$

Hence, we let the successors at level $\beta + 1$ be exactly $p \wedge b_\eta$. At limit levels we take the conjunction of all predecessors.

Note that since $\kappa > \delta$, and the nodes along a branch are decreasing, for each branch there is some β where the tree eventually stabilises. Hence if p is such a node where all of its successors are equal to it, then p must be completely join-prime: if $p = \bigvee_{i \in I} c_i$, then $(c_i)_{i \in I} \in C(p)$, so at a sufficiently large cardinal this will ocurr in the branch.

Now pick a bar B consisting of the nodes where the branch stabilises. By the TT-rule, we have that $c \leqslant \bigvee_{f \in B} \bigwedge_{\beta < \delta_f} \phi_{f \restriction_\beta}$. Each of these elements is below c, so this is equality, and by what we just showed they are completely join-prime, so c is the join of completely join-prime elements. By Proposition 3.3.7, then **H** is a Splitting algebra. □

The proof technique employed in the previous proposition illustrates the essence of the TT-rule: like all distributivity laws, its goal is to eliminate specific combinatorial structures living in our lattices, and the specific point here is to control the structure of trees of elements to ensure that the characteristic behaviour of distributivity – forcing a given top node to be below a join – is preserved at limit steps.

The key use of the TT-rule for our purposes lies in the following proposition, which uses the same ideas as above, and is proven in [Esp18]:

Proposition 3.4.3. Assume that $\kappa^{<\kappa}$. Let $\mathbf{H}$ be a distributive lattice such that:

- $\mathbf{H}$ has cardinality at most κ;
- $\mathbf{H}$ is closed under meets of less than κ many elements, and joins of at most κ many elements;
- $\mathbf{H}$ satisfies the Propositional TT_κ rule.
- $\mathbf{H}$ contains a collection Q of at most κ many joins of at most κ many elements, closed under the following distributivity requirement: if $(b_\alpha)_{\alpha<\lambda}$ is such that $\bigvee_{\alpha<\lambda} b_\alpha$ and $x \in \mathbf{H}$ then $\bigvee_{\alpha<\lambda} x \wedge b_\alpha \in Q$.

Then whenever $a \nleqslant b$, there is a Q-filter F which is κ-complete and κ-prime such that $a \in F$ and $b \notin F$.

Proof. Consider again the canonical well-ordering $f : \kappa \times \kappa \to \kappa$, and for each $c \in \mathbf{H}$ let $C(c)$ be the collection of tuples $(d_{\alpha<\lambda})$ such that either $\lambda < \kappa$, or if $\lambda = \kappa$, then $\bigvee d_\alpha \in Q$, and the join is equal to c in both cases. Define the tree in a similar way to above, using the new sets $C(c)$, and letting $\phi_\varnothing = c$. By the TT-rule, we have that c is below the join of the elements at a given fixed (but arbitrary) bar. Now assume towards a contradiction that along each branch of this tree, we eventually arrived at some $e \leqslant b$; hence for each such branch we could take the least element e in these conditions, and this would form a bar on the tree. Then we would have that $a \leqslant \bigvee_{e \in B} e \leqslant b$, a contradiction. So there must exist a branch of elements e such that $e \nleqslant b$ for all the elements in that branch.

Now define F by letting $c \in F$ if and only if for some $x \in \mathbb{B}$, $x \leqslant c$. Then note that F will be closed under meets of size smaller than κ, since the branch is of size κ; $a \in F$, and $b \notin F$; and if $\bigvee_{\alpha<\lambda} b_\alpha \in F$, then for some x in the branch we have that $x = \bigvee_{\alpha<\lambda} x \wedge b_\alpha$; hence each successor of x will be of the form $x \wedge b_\alpha$, and hence, $b_\alpha \in F$ for some α. This shows that F is a Q-filter as desired. □

The former now allows us to prove the completeness theorem for propositional κ-coherent logic, under the assumption that we only include conjunctions of size at most κ.

Definition 3.4.4. Let $\mathcal{L}^{coh-}_{\kappa^+}$ be κ^+-coherent logic with the TT_κ-rule, and with conjunctions limited to size less than κ.

Proposition 3.4.5. The logic $\mathcal{L}^{coh-}_{\kappa^+}$ is algebraic and relationally complete.

A consequence of the previous result is that κ-geometric logic is sound and complete, both algebraic and relationally. We delay the proof to include the first-order case.

One aspect which we might note now is that the above would not be (necessarily) enough if we also wanted to represent an *implication* connective. This is because to model the latter using Kripke semantics, one needs to represent it somewhat as:

$$\phi(a \to b) = X - \downarrow(\phi(a) - \phi(b)),$$

and to prove such an equality, we at some point must assume that if $x \notin \phi(a \to b)$, then there is some extension of x which contains a and does not contain b. In the finitary case this is ensured by the strong form of the prime filter lemma we mentioned above, and in the countably infinitary case, by the Rasiowa-Sikorski lemma on Heyting algebras (see for instance [Gol12]). Moreover, for weakly compact cardinals κ, C. Espindola has a proof that the strong prime filter lemma generalises in the relevant way. However, this does not seem necessary, as we will not need this to prove completeness with respect to intuitionistic logic.

We conclude by noting the first-order version of this rule:

Definition 3.4.6. The *(full) TT_κ rule* is the following rule:

$$\frac{\begin{array}{c}\phi_f \vdash_{\mathbf{y}_f} \bigvee_{g \in \gamma^{\beta+1}, g\restriction_\beta = f} \exists \mathbf{x}_g \phi_g,\ \beta < \kappa, f \in \gamma^\beta \\ \phi_f \dashv\vdash_{\mathbf{y}_f} \bigwedge_{\alpha<\beta} \phi_{f\restriction_\alpha}, \beta < \kappa,\ \text{limit}\ \beta,\ f \in \gamma^\beta\end{array}}{\phi_\varnothing \vdash_{\mathbf{y}_\varnothing} \bigvee_{f \in B} \exists_{\beta<\delta_f} \mathbf{x}_{f\restriction_{\beta+1}} \bigwedge_{\beta<\delta_f} \phi_{f\restriction_\beta}}$$

where $\gamma \leqslant \kappa$, where $\mathbf{y}_f$ is the context of ϕ_f, assuming that for each $f \in \gamma^\beta$, $FV(\phi_f) = FV(\phi_{f\restriction_\beta}) \cup \mathbf{x}_f$, where $\mathbf{x}_f \cap FV(\phi_{f\restriction_\beta}) = \varnothing$, and $FV(\phi_f) = \bigcup_{\alpha<\beta} FV(\phi_{f\restriction_\alpha})$ for limit β; where $B \subseteq \gamma^{<\kappa}$ is a collection of the minimal elements of a given bar, and the $\delta_f < \kappa$ are the levels of $f \in B$, and all meets and joins are assumed to exist, and additionally

As noted by C.Espindola, in the classical case, the former allow us to derive both the (κ, κ)-distributivity rules, and also the axiom of κ-dependent choices. Hence, this consists of the appropriate first-order generalisation of Karp's system. Hence, let us see this in action:

3.5 *First-order Infinitary Coherent Logic*

We return to the first-order case. Once again the idea for proving the soundness and completeness is to use a term model construction.

Definition 3.5.1. We denote the first-order κ-coherent logic with restricted conjunction, $\mathcal{L}^{coh-}_{\kappa^+,\kappa}$, the logic containing κ-disjunctions (provided the resulting set has less than κ many free variables), conjunctions of size less than κ, and existential quantification of formulas with existentials of size less than κ. We include in this logic the basic axioms and the full TT_κ-rule.

The key ideas here are the following:

- As before, we construct for each formula an associated collection of terms $T_\gamma(\phi)$ and a set of formulas $\Delta(\phi)$ containing all relevant substitution instances from a pool X of fresh variables.

- We assume that $\phi \nvdash \psi$, and construct a term algebra containing at most the κ-large joins of ϕ and ψ, including all substitution instances of the formulas obtained for variables, as before; the assumption that $\kappa^{<\kappa} = \kappa$ ensures this can be done in such a way that the resulting algebra has size at most κ. Given a formula of the form $\exists \mathbf{x}\phi(\mathbf{x})$ we add also the join $\bigvee \phi(\mathbf{y})$ ranging over all formulas (where again $\kappa^{<\kappa} = \kappa$ and the assumption on size of the formulas ensures this is well-defined). We take the quotient over derivability as usual.

- To obtain a condition for consistency, we require something analogous to Lemma 3.9.1 (see Appendix), where the requirement of maximality (i.e., $\phi \in \Gamma$ if and only if $\neg\phi \notin \Gamma$), given the absence of negations.

Indeed, the last step provides us a criterion for satisfiability, which allows the construction of a term model.

Hence, suppose that $\phi \nvdash_{\mathbf{x}} \psi$. All that is needed is to extract a prime theory from the Lindenbaum-Tarski algebra. One fact which is necessary is that the resulting algebra is indeed TT_κ-distributive.

Lemma 3.5.2. The algebra $\mathbf{H}$ as constructed above is TT_κ-distributive.

Proof. Assume that $\{a_f : f \in \gamma^{<\kappa}\}$ is a tree as defined. For each such formula, we can assume that $a_f = \phi_f$; by hypothesis, $a_g = \exists \mathbf{x}\theta_g(\mathbf{x})$ for each a_g such that $a_f = \bigvee_g a_g$. Hence, without loss of generality, we can consider $\phi_f = \bigvee_{g \in \gamma^{\beta+1}, g\restriction_\beta = f} \exists \mathbf{x}\theta_g(\mathbf{x})$, for each such f. In this way we construct a tree for which we can apply the TT_κ-rule.

Now notice that by the above rule, $\psi_\varnothing$ implies $\bigvee_{f\in B} \exists_{\beta<\delta_f} \mathbf{x}_{f\restriction_{\beta+1}} \bigwedge_{\beta<\delta_f} \phi_{f\restriction_\beta}$ in the context $\mathbf{y}_\varnothing$. Additionally we have that:

$$\exists_{\beta<\delta_f} \mathbf{x}_{f\restriction_{\beta+1}} \bigwedge_{\beta<\delta_f} \psi_{f\restriction_\beta} \vdash \phi_\varnothing$$

where the consequence is taken in the shared context; the latter fact follows from the fact that we are given all the witnesses outside of the conjunction, ensuring there are no clashes of variables, and indeed, the witnesses are as desired.

Hence by construction of the algebra this ensures that $a_\varnothing$ is the desired join of elements. □

Definition 3.5.3. Let $\mathbf{H}$ be an algebra of formulas as constructed in all the previous sections. If F is a prime filter over $\mathbf{H}$, let:

$$\Gamma_F = \{\phi : [\phi] \in F\}.$$

We call such a collection of formulas a (κ-)*prime theory*.

With the previous lemma in mind we can now prove:

Proposition 3.5.4. The logic $\mathcal{L}^{coh-}_{\kappa^+,\kappa}$ is algebraic and relationally complete.

Proof. Soundness is trivial. Now assume that $\phi \nvdash \psi$ in a specific context. Consider the algebra $\mathbf{H}$ as above. It is of size κ, has a collection of at most κ joins of size κ, is κ-complete, and TT_κ-distributive by Lemma 3.5.2. Hence, by Proposition 3.4.3, find a Q-filter P containing ϕ and not containing ψ. Then the term model $T(\phi)$, modulo the theory Γ_P, is the model we want (see the Appendix for the arguments for the classical case, the same facts apply here). □

Hence, the first-order case offers no difficulties when we are only concerned with coherent theories. It is then clear that:

Corollary 3.5.5. *First order κ-geometric logic is sound and complete.*

However, you will have noted that at no point did we talk about Kripke models. In the next section we will show where these become relevant by extending the previous results to the full first-order setting.

3.6 First-Order Infinitary Intuitionistic Logic

Despite the problems we outlined in Section 3.4, it is possible to extend the results to deal with the implication and the universal quantifier. The key to handle this is the following technical lemma:

Lemma 3.6.1. Let ϕ be a formula in first-order κ-intuitionistic logic. Assume that $\mathbf{H}$ is the Lindenbaum-Tarski algebra of size at most κ constructed in the same way as described in the previous section. Suppose that F is a Q-filter over $\mathbf{H}$. If $[a \to b] \notin F$, then there is a Q-filter $F' \supseteq F$ such that $[a] \in F'$ and $[b] \notin F'$. Similarly, if $[\forall \mathbf{x}\phi(\mathbf{x})] \notin F$, there is a Q-filter $F' \supseteq F$ such that for some collection of variables $\phi(\mathbf{y}) \notin F$.

Proof. Let Γ_F be the prime theory corresponding to F. Let T be the term algebra, such that $\mathbf{H}$ is a quotient of T. Let $\mathbf{H}'$ be the Lindenbaum-Tarski algebra modulo Γ_F, i.e., we take the quotient under derivability, adding the formulas in Γ_F as axioms. Note that then $\mathbf{H}'$ will be again an algebra of size at most κ, κ-complete, TT_κ-distributive (by the same argument as in Lemma 3.5.2). Note that if $[a] \to [b] \notin F$, then $a \to b \notin \Gamma_F$, and so $[a]_F \nleqslant [b]_F$, where these represent the equivalence classes in this algebra. Hence by Lemma 3.4.3, there is a Q-prime filter containing $[a]_F$ and not containing $[b]_F$, say G. Now define the following:

$$F' := \uparrow\{[\psi] \in \mathbf{H} : [\psi]_F \in \mathbf{H}'\}.$$

We will show that this is a Q-filter:

- Closure under $< \kappa$-meets is straightforward to verify; now assume that $[\psi] \leqslant \bigvee_{\alpha<\kappa}[\mu_\alpha]$; hence $\psi \vdash \bigvee_{\alpha<\kappa}\mu_\alpha$, so in $\mathbf{H}'$, $[\mu_\alpha]_F \in G$ for some α. Hence $[\mu_\alpha] \in F'$ as desired.

- $[a] \in F'$ and $[b] \notin F'$; to see the latter, assume that $[\psi] \leqslant [b]$ where $[\psi]_F \in G$; the former implies that $\psi \vdash b$, which would then force that $[b]_F \in G$.
- $F \subseteq F'$: if $[\phi] \in F$, then by construction $[\phi]_F = [\top]_F$, so $[\phi]_F \in G$, and hence $[\phi] \in F'$.

Thus, F' is indeed the desired theory. The case where $[\forall \mathbf{x} \phi(\mathbf{x})] \notin F$ is wholly similar. Consider again the quotient algebra under the theory Γ_F. Since $[\forall \mathbf{x} \phi(\mathbf{x})]_F \neq [1]_F$, and this is the meet of all the formulas $\phi(\mathbf{y})$ for $\mathbf{y}$ not free in ϕ, we have that there must be some $\mathbf{y}$ for which $[\phi(\mathbf{y})]_F \neq [1]$. Then proceed as above, obtaining a filter $F' \supseteq F$ such that $[\phi(\mathbf{y})] \notin F'$. □

Definition 3.6.2. Let $\mathcal{L}^{int-}_{\kappa^+,\kappa}$ denote the first-order intuitionistic logic where conjunctions are of size $<\kappa$.

Corollary 3.6.3. *The logic $\mathcal{L}^{int-}_{\kappa^+,\kappa}$ is sound and complete with respect to Kripke models.*

Proof. Soundness is easy. To see completeness, again assume that $\phi \nvdash \psi$. Construct the Lindenbaum-Tarski algebra $\mathbf{H}$ as before, which is a Heyting algebra saturated with witnesses. Let

$$Pr(\mathbf{H}) = \{\Gamma_F : F \text{ is a } Q\text{-filter over } \mathbf{H}\}$$

be the collection of prime theories over this language, ordered by inclusion. Note that since $\phi \nvdash \psi$, then there is a theory T_0 containing ϕ and not containing ψ, which we take as our root. Now, for each T' such that $T_0 \subseteq T'$, let $\mathbb{T}'$ be the term model containing the witnesses from T', and quotiented under T'. Let **Term** denote the collection of all such models. Note that these are all first-order models, which additionally satisfy persistence, given the theories are ordered by inclusion. This forms a Kripke model $\mathfrak{M}$, and we claim that for all $\mathbb{T} \in \mathfrak{M}$:

$$\mathbb{T} \Vdash \phi \iff \phi \in T$$

This is straightforward for the "local" clauses, since the theories are obtained from Q-filters, and hence are closed under disjunctions and conjunctions of appropriate size, and under existential quantifiers. As for the implications and universal quantifiers, we use Lemma 3.6.1, together with the induction hypothesis, to ensure the result. □

3.7 Set-Theoretic and Algebraic Rasiowa-Sikorskis

We include here, since I could not find the proof anywhere, a brief discussion on how the version of Rasiowa-Sikorski that we are using (which can be called the "Algebraic" Rasiowa-Sikorski) relates to the more usual version in set theory.

Definition 3.7.1. The "Set-Theoretic Rasiowa-Sikorski Lemma" refers to the following statement:

- For every poset P, and every countable family $\mathcal{D} = \{D_n : n \in \omega\}$ of dense subsets of P, there is a $\mathcal{D}$-generic filter F, i.e., for every n, $F \cap D_n \neq \varnothing$.

Lemma 3.7.2. The Algebraic form of Rasiowa-Sikorski follows from the Set-Theoretic one.

Proof. Let B be a Boolean algebra, and consider $Q = (\{X_n\})$ a collection of subsets which have an infinite join in the algebra. Then for each n look at:

$$D_n = \{\bigwedge X_n\} \cup \{\neg a_m : a_m \in X_n, m \in \omega\}$$

Indeed, we can show that this is a dense subset, since if x is arbitrary, and $x \wedge \bigwedge X_n = 0$, then $x \leqslant \bigvee_{n\in\omega} \neg a_n$, so $x = \bigvee_{n\in\omega} x \wedge \neg a_n$, so there must be some m such that $x \wedge \neg a_m \neq \varnothing$. Hence let G be a generic filter, intersecting all these filters and containing p, and let U be an ultrafilter extending G. Then U is a Q-filter, as desired. $\square$

3.8 *Proof that κ-representability implies completeness*

The following is our topological adaptation of Chang's proof that κ-representability implies the existence of a Q-filter (the original proof was formulated in a purely algebraic fashion):

Proposition 3.8.1. Let κ be a regular cardinal. Then $\mathbf{A}$ is a κ-representable Boolean algebra if and only if in the dual space X, every intersection of less than κ many dense open subsets, each a union of less than κ many clopen sets, is non-empty.

Proof. Consider X the dual Stone space of A, and identify A with $\mathsf{Clop}(X)$. Let $g : B \to A$ be a surjective κ-complete epimorphism from $B \leqslant \mathcal{P}(X)$ onto A, the former of which is a κ-algebra of sets. Let $h : A \to B$ be a map defined as follows: pick an ultrafilter U containing a, and whenever $c \in U$, let $h(c)$ be some element such that $g(h(c)) = c$ and when $c \notin U$, let $h(c) = \neg h(\neg c)$. Assume that $(\phi(a_{i,j}))_{i,j}$ is a collection of clopen sets such that for each i, $\bigcup_{j\in J} \phi(a_{i,j})$ is dense. Assume that for some c in $\mathbf{A}$

$$\phi(c) \subseteq \bigcup_{i\in I} \bigcap_{j\in J} \phi(\neg a_{i,j}).$$

Then in particular $c \subseteq \bigvee_{i\in I} \phi(a_{i,f(i)})$ for each $f \in J^I$.

Note that $\bigwedge_{j\in J} \neg\phi(a_{i,j}) = int(\bigcap_{j\in J} \neg\phi(a_{i,j})) = \varnothing$. Then note that:

$$\bigcup_{i\in I} \bigcap_{j\in J} \neg h(a_{i,j}) \neq 0_B$$

Otherwise, we would have that:

$$g(\bigcup_{i\in I} \bigcap_{j\in J} \neg h(a_{i,j})) = \bigvee_{i\in I} \bigwedge_{j\in J} \neg\phi(a_{i,j}) = \varnothing$$

$\square$

Proposition 3.8.2. Let **B** be a κ^+-complete and κ-representable Boolean algebra for regular κ. Then whenever $a \in B$, and $Q = (\{X_\alpha\})_{\alpha<\lambda}$ is a collection of $\lambda \leqslant \kappa$ sets of elements which meet belongs to **B**, then there exists an ultrafilter of **B** containing a and preserving the meets in Q.

Proof. Consider X the dual Stone space of A, and identify A with $\mathsf{Clop}(X)$. Let $g : B \to A$ be a surjective κ-complete epimorphism from $B \leqslant \mathcal{P}(X)$ onto A, the former of which is a κ-algebra of sets. Since $a \neq 0$, let U be an ultrafilter containing a. Let $h : B \to A$ be a function defined as follows: if $c \in U$, then $h(c)$ is an arbitrary element of B such that $g(h(c))$, and if $c \notin U$, then $h(c) = \neg h(c)$. Consider the family $\mathcal{H}$ of subsets of **A** as follows:

- $\{a\}$;
- $\neg \bigwedge X_\alpha \cup \bigcup_{\beta<\kappa} x_\beta$ for each α in the Q-set.
- $\{\neg c_0, ..., \neg c_n, c_0 \wedge ... \wedge c_n\}$ for each finite subset of elements from **A**.

Note that except for the first case, then whenever D is such a set, $\bigwedge D = int(\bigcap D)$ is empty.

For each $D \in \mathcal{H}$, let $h_D = \{h(b) : b \in D\} \subseteq B$. Then note that:

$$\bigcup\{\bigcap h_D : D \in \mathcal{H}\} \neq 1_B$$

for otherwise, because g is sectioned by h and preserves κ-complete meets and joins:

$$\begin{aligned} g(\bigcup\{\bigcap h_D : D \in \mathcal{H}\}) &= g(\bigcup\{\bigcap\{h(b) : b \in D\} : D \in \mathcal{H}\} \\ &= \bigvee\{\bigwedge\{g(h(b)) : b \in D\} : D \in \mathcal{H}\} \\ &= \bigvee\{\bigwedge D : D \in \mathcal{H}\} \\ &= \{\neg\phi(a)\}. \end{aligned}$$

Since by assumption $a \neq 0$, then $\neg a \neq 1$, so this is a contradiction. Now, for each choice function on $\mathcal{H}$, f, let $c_{D,f(D)}$ denote $f(D)$. Hence using distributivity over B we have:

$$\bigcup\{\bigcap_{D\in\mathcal{H}} c_{D,f(D)} : D \in \mathcal{H}\} \neq \varnothing.$$

So let $T \in B$ be an element that belongs here. Thus for each α, either $T \subseteq \phi(\neg \bigwedge X_\alpha))$ or $T \subseteq \phi(x_\beta)$ for some $x_\beta \in X_\alpha$. For each α denote the relevant element by d_α Hence consider:

$$S = \{a\} \cup \{d_\alpha : \alpha < \kappa\}.$$

Note that this forms a filter basis: if $a, d_0, ..., d_n$ are arbitrary elements, then $T \subseteq h(d_k)$ for each k, hence look at the set $D_l = \{h(\neg a), h(\neg d_0), ..., h(\neg d_n), h(a \wedge d_0 \wedge ... \wedge d_n)\}$; if $T \subseteq h(\neg d_k)$, then because h preserves complements, $T \subseteq \neg h(d_k)$, so $T = \varnothing$, a contradiction. Hence the only option is that $T \subseteq h(a \wedge d_0 \wedge ... \wedge d_n)$, i.e. $a \wedge d_0 \wedge ... \wedge d_n \in S$; so extend S to an ultrafilter F. Then F is the desired Q-filter. □

3.9 *Proof of completeness of First-order Infinitary Calculus*

The following is due to Carol Karp [Kar64]. We begin by proving the following lemma, which gives us a criterion for satisfiability:

Lemma 3.9.1. Assume that ϕ is an arbitrary formula, and suppose that $T(\phi)$ and Δ are as previously outlined. Then ϕ is satisfiable if there is a set of formulas Γ containing ϕ and all formulas of the form $g = g$ for $g \in T(\phi)$, and satisfying the following conditions:

- If terms $g(t_\alpha)_{\alpha<\delta}$ and $g(t'_\alpha)_{\alpha<\delta}$ are in $T(\phi)$ then if $t_\alpha = t'_\alpha \in \Gamma$ for all $\alpha < \delta$ then $g(t_\alpha)_{\alpha<\delta} = g(t'_\alpha)_{\alpha<\delta} \in \Gamma$.
- If $R(t_\alpha)_{\alpha<\delta}$ and $R(t'_\alpha)_{\alpha<\delta}$ are in Δ then if $t_\alpha = t'_\alpha \in \Gamma$ for all $\alpha < \delta$ and $R(t_\alpha)_{\alpha<\delta} \in \Gamma$, then $R(t'_\alpha)_{\alpha<\delta} \in \Gamma$.
- If $\psi \in \Delta$ then $\psi \in \Gamma$ if and only if $\neg\psi \notin \Gamma$.
- If $\bigwedge_{\eta<\kappa} \psi_\eta \in \Delta$, then $\bigwedge_{\eta<\kappa} \psi_\eta \in \Gamma$ iff all the $\psi_\eta \in \Gamma$.
- If $\exists_{i\in I} v_i \psi(v_i) \in \Delta$, then $\exists_{i\in I} v_i \psi(v_i) \in \Gamma$ if and only if there is a substitution: $\psi(t_i) \in \Gamma$ for t_i a collection of I many terms from T.

Proof. This proceeds as in the finitary case: let $T(\phi)$ be the set of terms, and form the term model by taking the set of equivalence classes modulo Γ. The above conditions ensure that this forms a well-defined equivalence relation, and that interpreting function symbols by letting them name themselves is well-defined. We define relation symbols in the usual way: $R([t_\alpha])_{\alpha<\delta}$ if and only if there are $t'_\alpha \in [t_\alpha]$ such that $R(t_\alpha)_{\alpha<\delta} \in \Gamma$. Call this model $\mathbb{T}$. Then the remaining clauses ensure that for each formula $\psi \in \Delta$:

$$\mathbb{T} \models \psi \iff \psi \in \Gamma.$$

This is done by induction on complexity of formulas. For equality this is by definition; for relation symbols this is given, and all clauses except the existential follow immediately. Finally we look at the existential case. Assume that $\psi = \exists_{i\in I} v_i \chi(v_i)$. Indeed first suppose that $\psi \in \Gamma$. Then by the hypothesis, there is a substitution for a term in $T(\phi)$, t_i such that $\chi(v_i) \in \Gamma$. Hence by induction hypothesis, $\mathbb{T} \models \chi(v_i)$ which means by hypothesis that $\mathbb{T} \models \exists_{i\in I} v_i \chi(v_i)$. The converse is immediate. □

Theorem 3.9.2. *Assume that $\kappa^{<\kappa} = \kappa$. Then the calculus $\mathcal{L}_{\kappa^+,\kappa}$ is relationally sound and complete.*

Proof. Soundness as usual is obvious in all cases except perhaps the rule of Dependent Choices; this follows by the set theoretic assumptions of our meta-theory (namely, the fact that dependent choice holds in the outside universe). Now assume that $\nvdash_{\mathcal{L}_{\kappa^+,\kappa}} \phi$. Form Δ and $T(\phi)$ as before. Let the following list all formulas which are existentially quantified in Δ:

$$S = \{\exists_{i\in I} v_i \psi(v_i)_\gamma : |I| < \kappa, \gamma < \kappa\}$$

For each such formula we can find a fresh collection of symbols c_i for $i \in I$, and we consider the formulas $W_\eta := \exists_{i \in I} v_i \psi(v_i)_\gamma \to \psi(c_i)$. Note that there are at most κ many such formulas, so this adds at most κ many constants. We let:

$$\mathbb{F}_\phi$$

be the free Boolean algebra obtained by closing Δ under $< \kappa$-operations, as well as adding the meet $\neg\phi \wedge \bigwedge_{\alpha<\kappa} W_\alpha$. Note the resulting algebra is still of size at most κ. Then we claim that:

$$[\neg\phi \wedge \bigwedge_{\alpha<\kappa} W_\alpha] \neq [\bot]$$

Indeed, suppose that it was. Then by definition, $\vdash \neg(\neg\phi \wedge \bigwedge_{\alpha<\kappa} W_\alpha)$, so $\vdash \neg\phi \to \bigvee_{\alpha<\kappa} \neg W_\alpha$. Unfolding this means that $\vdash \neg\phi \to \bigvee_{\alpha<\kappa} \exists_{i \in I} v_i \psi(v_i) \wedge \neg\psi(c_i)_\alpha$. From this we can infer, by propositional reasoning, that $\vdash \bigvee_{\alpha<\kappa} \phi \vee \exists_{i \in I} v_i \psi(v_i) \wedge \neg\psi(c_i)_\alpha$, hence by the law of dependent choices, and given that the variables are all fresh where they must, we infer $\vdash \forall_{i \in I_0} c_i \phi \vee \forall x_0 W_0 \vee ... \vee \exists_{\eta<\lambda} v_\eta \forall v_\lambda W_\lambda ...$; distributing the universal quantifiers, since there is no clash of variables, we obtain $\exists_{i \in I} v_i \psi(v_i) \wedge \forall_{i \in I} v_i \neg\psi(v_i)$ for each such clause. Hence, we conclude that $\vdash \phi$, which is a contradiction. Hence by reductio, we have that $[\neg\phi \wedge \bigwedge_{\alpha<\kappa} W_\alpha] \neq [\bot]$.

With this in place, we now let Q consist of all infinitary meets in $\mathbb{F}_\phi$, and by hypothesis on κ-representability, obtain a Q-filter on the algebra containing $\neg\phi \wedge \bigwedge_{\alpha<\kappa} W_\alpha$. If P is such a Q-filter, we can consider $P' = \{\psi : [\psi] \in P\}$, and we can show that this satisfies the conditions of Lemma 3.9.1; we only check the last condition. Indeed if there is a substitution $\psi(t_i) \in P$, then because $\psi(t_i) \to \exists v_i \psi(v_i)$, $\exists v_i \psi(v_i) \in P'$. Otherwise, assume that $\exists v_i \psi_\alpha(v_i) \in P'$; then since $W_\alpha \in P'$, by deductive closure, $\psi_\alpha(c_i) \in P'$.

Now by Lemma 3.9.1, we have that $\neg\phi$ is satisfiable in a model $\mathbb{T}$. This shows completeness. □

3.10 *Distributivity and Representability in Heyting algebras*

Lemma 3.10.1. Let **H** be a complete Heyting algebra. Then the following are equivalent:

1. **H** is a complete homomorphic image of a completely join-prime generated algebra **H′**;

2. **H** is completely distributive.

Proof. First assume that (1) holds. Let $(a_{i,j})_{i \in I, j \in J}$ be a doubly indexed family of elements. Let $f(c_{i,j}) = a_{i,j}$ be elements in $\mathbf{H}'$. Then note that:

$$\begin{aligned}\bigwedge_{i \in I} \bigvee_{j \in J} a_{i,j} &= f(\bigwedge_{i \in I} \bigvee_{j \in J} c_{i,j}) \\ &= f(\bigvee \{\bigwedge_{i \in I} c_{i,g(i)} : g \in J^I\}) \\ &= \bigvee \{\bigwedge_{i \in I} a_{i,f(i)} : g \in J^I\}\end{aligned}$$

Now assume that (2) holds. Let $\mathbf{H}'$ be the lattice of downwards closed subsets of $\mathbf{H}$. Note that this is a completely join-prime generated Heyting algebra when we consider arbitrary unions and intersections. Now assume that $\{S_i : i \in I\}$ is an indexed collection of downwards closed subsets. Denote by $M(I)$ the collection of functions g from $I \to L$, such that $g(i) \in S_i$. Then note that:

$$\bigwedge\{S_i : i \in I\} = \{\bigcap g[I] : g \in M(I)\}$$

Indeed, if $x \in S_i$ for each i, then let g be a function mapping constantly to x; then this belongs to the set $\bigcap g[I]$. Conversely, let S be a set of the form $\bigcap g[I]$, then x must belong to each of these subsets.

Now define the map $f : \mathbf{H}' \to \mathbf{H}$ as follows: $f(S) = \bigvee S$. By completeness of $\mathbf{H}$ this is well-defined. We check that f preserves the operations: it is easy to see that it will preserve the complete union; if $(S_i)_{i \in I}$ is a collection of downwards closed subsets, then:

$$\begin{aligned} f(\bigwedge\{S_i : i \in I\}) &= \bigvee\{\bigcap_{i \in I} g(i) : g \in M(I)\} \\ &= \bigwedge_{i \in I} \bigvee_{l \in L} S_i \\ &= \bigwedge_{i \in I} f(S_i) \end{aligned}$$

which shows preservation of complete meets. Finally, assume that $V = U \Rightarrow W$. Then assume that $c \wedge f(U) \subseteq f(W)$. Consider $\downarrow c$, and note that then $\downarrow c \cap U \subseteq W$: indeed, if $x \leqslant c$ and $x \leqslant \bigcup U$, then $x \leqslant c \wedge f(U) \leqslant f(W)$, so since W is downwards closed, $x \in W$. Hence $\downarrow c \subseteq V$, hence $c \leqslant f(V)$, as intended. Hence we have that f is a complete homomorphism, and it is clearly surjective, as desired. □

Bibliography

[AR94] J. Adamek and J. Rosicky. *Locally Presentable and Accessible Categories*. Lecture note series / London mathematical society. Cambridge University Press, 1994. ISBN: 9780521422611. URL: https://books.google.nl/books?id=iXh6rOd7of0C.

[Car18] O. Caramello. *Theories, Sites, Toposes: Relating and studying mathematical theories through topos-theoretic'bridges'*. Oxford University Press, 2018.

[Cha57] C. C. Chang. "On the Representation of α-Complete Boolean Algebras". In: *Transactions of the American Mathematical Society* 85.1 (May 1957), p. 208. DOI: 10.2307/1992970. URL: https://doi.org/10.2307/1992970.

[CK12] C. Chang and H. Keisler. *Model Theory*. Mineola, New York: Dover, 2012. ISBN: 9780486488219.

[CZ97] A. Chagrov and M. Zakharyaschev. *Modal Logic*. Oxford Logic Guides. Oxford, England: Clarendon Press, 1997.

[Dav79] B. Davey. "On the lattice of subvarieties". In: *Houston journal of mathematics* 5 (Jan. 1979), pp. 183–192.

[Dic85] M. Dickman. "Large Infinitary Languages". In: *Model Theoretic Logics*. Ed. by J. Barwise and S. Feferman. New York, Berlin, Heidelberg, Tokyo: Springer-Verlag, 1985. Chap. IX, pp. 317–348.

[DP02] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge: Cambridge University Press, 2002, pp. 175–200.

[Esa19] L. Esakia. *Heyting Algebras: Duality Theory*. Ed. by G. Bezhanishvili and W. Holliday. English translation of the original 1985 book. Springer, 2019.

[Esp18] C. Espíndola. *A complete axiomatization of infinitary first-order intuitionistic logic over $\mathcal{L}_{\kappa^+,\kappa}$*. 2018. DOI: 10.48550/ARXIV.1806.06714. URL: https://arxiv.org/abs/1806.06714.

[Esp19] C. Espındola. "Infinitary first-order categorical logic". In: *Annals of Pure and Applied Logic* 170.2 (Feb. 2019), pp. 137–162. DOI: 10.1016/j.apal.2018.09.002. URL: https://doi.org/10.1016/j.apal.2018.09.002.

[ESP20] C. ESPÍNDOLA. "INFINITARY GENERALIZATIONS OF DELIGNE'S COMPLETENESS THEOREM". In: *The Journal of Symbolic Logic* 85.3 (Sept. 2020), pp. 1147–1162. DOI: 10.1017/jsl.2020.27. URL: https://doi.org/10.1017/jsl.2020.27.

[Gol12] R. Goldblatt. "Topological Proofs of Some Rasiowa-Sikorski Lemmas". In: *Studia Logica: An International Journal for Symbolic Logic* 100.1/2 (2012), pp. 175–191. ISSN: 00393215, 15728730. URL: http://www.jstor.org/stable/41475222 (visited on 11/26/2022).

[Gol74] R. Goldblatt. "Semantic analysis of orthologic". In: *Journal of Philosophical Logic* 3.1-2 (1974), pp. 19–35. DOI: 10.1007/bf00652069.

[Gol75] R. Goldblatt. "The Stone Space of an Ortholattice". In: *Bulletin of the London Mathematical Society* 7.1 (Mar. 1975), pp. 45–48. DOI: 10.1112/blms/7.1.45. URL: https://doi.org/10.1112/blms/7.1.45.

[Gro02] R. Grossberg. *Classification theory for abstract elementary classes*. 2002. DOI: 10.1090/conm/302/05080. URL: https://doi.org/10.1090/conm/302/05080.

[Kar64] C. R. Karp. *Languages with Expressions of Infinite Length*. Amsterdam: North-Holland Publishing Company, 1964.

[KMB89] S. Koppelberg, D. Monk, and R. Bonnet, eds. *Handbook of Boolean algebra: V. 1-3*. en. London, England: Elsevier Science, Feb. 1989.

[KT64] H. Keisler and A. Tarski. "From accessible to inaccessible cardinals (Results holding for all accessible cardinal numbers and the problem of their extension to inaccessible ones)". eng. In: *Fundamenta Mathematicae* 53.3 (1964), pp. 225–308. URL: http://eudml.org/doc/213737.

[Mar02] D. Marker. *Model Theory*. Springer-Verlag, 2002. DOI: 10.1007/b98860. URL: https://doi.org/10.1007/b98860.

[MP89] M. Makkai and R. Paré. *Accessible Categories: The Foundations of Categorical Model Theory: The Foundations of Categorical Model Theory*. Contemporary mathematics - American Mathematical Society. American Mathematical Society, 1989. ISBN: 9780821851111. URL: https://books.google.nl/books?id=ExIcCAAAQBAJ.

[Ran52] G. N. Raney. "Completely distributive complete lattices". In: *Proceedings of the American Mathematical Society* 3.5 (1952), pp. 677–680. DOI: 10.1090/s0002-9939-1952-0052392-3. URL: https://doi.org/10.1090/s0002-9939-1952-0052392-3.