Achieving completeness:
from constructive set theory to large cardinals

Christian Espíndola

Stockholm
University

Achieving completeness: from constructive set theory to large cardinals

Christian Espíndola

Address: Matematiska Institutionen, Stockholms Universitet, 106 91 Stockholm
E-mail address: espindola@math.su.se

ISBN 978-91-7649-458-5

Printed in Sweden by Holmbergs, Malmö, 2016

Distributor: Department of Mathematics, Stockholm University

Abstract

This thesis is an exploration of several completeness phenomena, both in the constructive and the classical settings. After some introductory chapters in the first part of the thesis where we outline the background used later on, the constructive part contains a categorical formulation of several constructive completeness theorems available in the literature, but presented here in an unified framework. We develop them within a constructive reverse mathematical viewpoint, highlighting the metatheory used in each case and the strength of the corresponding completeness theorems.

The classical part of the thesis focuses on infinitary intuitionistic propositional and predicate logic. We consider a propositional axiomatic system with a special distributivity rule that is enough to prove a completeness theorem, and we introduce weakly compact cardinals as the adequate metatheoretical assumption for this development. Finally, we return to the categorical formulation focusing this time on infinitary first-order intuitionistic logic. We propose a first-order system with a special rule, transfinite transitivity, that embodies both distributivity as well as a form of dependent choice, and study the extent to which completeness theorems can be established. We prove completeness using a weakly compact cardinal, and, like in the constructive part, we study disjunction-free fragments as well. The assumption of weak compactness is shown to be essential for the completeness theorems to hold.

Sammanfattning

Denna avhandling är en undersökning av flera fullständighetsfenomen, både i konstruktiva och klassiska versioner. Efter några inledande kapitel i den första delen av avhandlingen där vi beskriver bakgrunden som kommer att användas senare, innehåller den konstruktiva delen en kategoriteoretisk formulering av flera konstruktiva fullständighetssatser som finns i litteraturen, men här presenterade i ett enhetligt ramverk. Vi utvecklar dem från den konstruktiva "omvända matematikens" perspektiv, med fokus på metateorien som används i varje enskilt fall och styrkan hos motsvarande fullständighetssatser.

Den klassiska delen av avhandlingen fokuserar på infinitär intuitionistisk satslogik och predikatlogik. Vi betraktar ett axiomatiskt satslogiskt system med en speciell distributivitetsregel som är tillräcklig för att bevisa en fullständighetssats, och vi introducerar svagt kompakta kardinaltal som det adekvata metateoretiska antagande. Slutligen återvänder vi till den kategoriska formuleringen, fokuserande denna gång på infinitär första ordningens intuitionistisk logik. Vi föreslår ett predikatssystem med en särskild regel, transfinit transitivitet, som innehåller både distributivitet och en form av axiomet om beroende urval, och studerar i vilken utsträckning fullständighetssatser gäller. Vi fastlägger fullständighet under antagande om ett svagt kompakt kardinaltal, och studerar det disjunktionsfria fragmentet precis som i avhandlingens konstruktiva del. Antagandet om svag kompakthet bevisas vara väsentligt för att fullständighetssatserna ska gälla.

Contents

Acknowledgements

This thesis would not have been possible without the support of many people. First I would like to thank my supervisor Erik Palmgren and my co-supervisor Henrik Forssell for having so much patience along the years in which I have been a student, as well as the Swedish Research Council for financial support through the VR grant "Constructive and category-theoretic foundations of mathematics" (dnr 2008-5076) that allowed me to finish my studies. Likewise I would like to thank the members of the logic group Peter Lumsdaine, Håkon Gylterud, Jacopo Emmenegger and Johan Lindberg and extend my gratitude to all my colleagues at the Mathematics Department, especially those who accompanied me most of my time at Stockholm, some of whom are no longer here: Alessandro Oneto, Ivan Martino, Ornella Greco, Frédéric Gaunard, Stefano Marseglia, Christoph Neuner, Felix Wierstra, Olof Bergvall, Iara Gonçalves, Rune Suhr, Ketil Tveiten, Gleb Nenashev, Gabriele Balletti and Qimh Xantcha. Finally, I would like to thank all my friends and acquaintances that have made my stay in Sweden an interesting experience.

Part I

Introduction

Vi får inte veta.
Vi kommer aldrig att veta.

1. Introduction

The axiomatization of mathematics over first-order logic has led to the quest for semantically sufficient sets of axioms from which all valid statements can be proved. Stemming from this original idea, one analyzes carefully the metatheory and inspects the fragment of logic under consideration to produce the several completeness theorems achieving this. The aim of this thesis is to present evidence of the interplay between metatheory and expressive power of the logic for which one is able to prove a completeness theorem. In particular, how constructivist constraints affect the range of completeness results we can obtain, and how the addition of more powerful axioms to the metatheory becomes a natural tendency when strengthening the richness of the logic we use. It will, ultimately, be part of everyone's choice to find the appropriate balance, which is not the goal here; instead, we would like to raise awareness of what seems to be an intrinsic feature of our mathematical investigations, namely, the lack of a corresponding "theory of everything" that can join the two worlds we will explore: that of the constructive mathematician based on intuitionistic logic and that of the classical set theorist, whose preferred metatheory goes far beyond innocent uses of the axiom of choice, and actually delves into the upper attic of Cantor's paradise, adopting large cardinal axioms which seem more and more natural.

The attention of the reader should not be distracted with the (necessarily) technical details of the thesis, but should rather focus on the general aim of trying to search for the most convenient axiomatization of current mathematics, one which is able to keep all the essential features of every proposal, but gives a more unifying, combined picture. The significance of the results here obtained should not be judged by their usefulness/uselessness, but rather by the rôle in this overall example of the problem that supposes merging two different points of view regarding mathematical practice. Perhaps the best way to consider them is adopting a relativistic point of view, which can change with the circumstances, sometimes in an unnoticed way. In fact, the transition from constructivism to classical non-constructive settings has been presented in as smooth a manner as possible, to have at least a taste of what a unifying picture should look like.

One of the background themes around which this thesis evolves is Joyal's categorical completeness theorem, presented for the first time in the seventies, and appearing in [MR77]. Joyal's theorem is in itself a proof of the unifying power of categorical language, which allowed him to combine in a single statement three different completeness theorems: that of coherent logic, that of first-order classical logic and that of first-order intuitionistic logic. Our contribution here will be exploiting a bit further the key features of his proof to obtain a more diverse fan of completeness theorems, of different types of logics in different metatheories. This will allow to extend the unification gathering results from constructive model theory to infinitary languages, presenting them as aspects of this unique categorical construction, casting different shadows according to the metatheory or the logic one is using to shed some light.

We will describe a constructive version of Joyal's theorem, from which several known constructive completeness theorems arise, and study their strength in some broad constructive reverse mathematical sense, that is, as equivalences with certain canonical statements over a particular metatheory which serves as a foundation. This will allow to connect the existing intuitionistic completeness proofs for intuitionistic logic as well as to get new completeness theorems. We will also direct our attention to intuitionistic infinitary logics, obtaining as a product completeness theorems for this case as well. Unlike classical infinitary logics, whose related completeness results have been known for decades, the main difficulty in studying infinitary intuitionistic logics is the huge variety of non-equivalent formulas that one can obtain. This, as we will show, begs the introduction of large cardinal axioms appropriate to handle this unexpected richness that one gets by dropping the excluded middle.

1.1 Summary of results and open problems

The reverse mathematics programme started classically with the work of Friedman and Simpson, with the goal of identifying statements over a given axiomatic base which are equivalent in strength. Intuitionistically it had some developments, although less sharply defined in terms of formal theories, with the work of Brouwer and Bishop, followed by Richman, Bridges and Ishihara. It has been known since Henkin that the completeness theorem for classical first-order theories has, within ZF, the strength of the Boolean Prime Ideal theorem. In fact, it is not difficult to see (as will become apparent in the following sections) that the same is true if we consider intuitionistic first-order theories within the same metatheory. The Boolean Prime Ideal theorem is known in turn to be equivalent within ZF to the compactness of the generalized Cantor space 2^α for every set α. Our thesis will consist of studying these completeness results and their relation to the metatheory when we move in two different directions. First, when going down to the constructive metatheory IZF, in which case a statement of completeness for intuitionistic/classical enumerable theories will be shown to be equivalent to the FAN theorem, i.e., the compactness of the Cantor space 2^ω. Second, we will enhance our theories under study expressing its axioms in infinitary logic, and in the metatheory ZFC we will prove that a statement of completeness for intuitionistic/classical theories of cardinality at most κ is equivalent to the weak compactness of κ, which can be defined by saying that the space 2^κ with the topology generated by initial segments of length less than κ satisfies a generalized compactness property, namely, every open cover contains a subcover of cardinality less than κ. If we compare this to the constructive result, we see that there is an analogy, and both cases in which we vary the metatheory and the underlying theory of study behave in a somewhat similar way. This analogy will become even more explicit when we consider first-order object theories which do not contain disjunction as a connective. As we shall see, dropping disjunction will result in the first case dropping the FAN theorem from the metatheory when we establish the completeness theorem, and in the second case, dropping the need of the weak compactness of κ. To motivate the necessity of the weak compactness we will study, in between, the case of propositional infinitary intuitionistic logics, which does not require any categorical logic and will serve as a pivotal example that will allow us to smoothly introduce the last part of the thesis.

The results of the constructive part almost do not contain (apart from metatheoretical equivalents of some completeness theorems) new theorems, but it is rather an organized unified framework where old results all fit together. The classical part, dealing with infinitary logics, contains new completeness results, that of intuitionistic theories (the study of classical infinitary logics had already reached an advance stadium in the monograph [Kar64] of Karp). We prove that a valid sequent (with respect to Kripke semantics) is derivable from a theory of cardinality at most κ within infinitary intuitionistic logic in case κ is a weakly compact cardinal, but we leave open the exact strength within the large cardinal hierarchy of a cardinal for which one can prove completeness of theories of cardinality strictly less than κ (what in [Kar64] is called weak completeness).

1.2 Papers related to this thesis

The constructive part of the thesis is based on a joint paper with Henrik Forssell, [FE]. Although the direction in this thesis follows a somewhat different approach to that of the paper, the content of the first part has been made possible thanks to the contribution during conversations, and the ideas of the proofs have been inspired by this fruitful exchange. In fact, the results which belong completely to the scope of the paper have been quoted here with the corresponding mention of the author, while the propositions that are unnamed have been proved with a slightly different approach, but that essentially conveys all the essential features of the ideas introduced in that paper. The first lemma of the thesis uses a result of another paper of the author, [Esp16], which is intended from the very beginning to show the limitations of the constructive results that we should be expecting to get, paving the way for the development of the constructive part of the thesis. Another paper by the same author, [Esp13] is not explicitely discussed here, but its ideas are lurking behind some of the reasoning when considering the possibility of getting intuitionistic proofs of classically valid results.

Part II

Infinitary first-order categorical logic

2. Infinitary first-order logic

Systems for classical infinitary propositional and first-order logic have been described and studied extensively in the monograph [Kar64] with Hilbert-type systems (see also [MT61] for a related development with Gentzen's sequents). Infinitary languages $\mathcal{L}_{\kappa,\lambda}$ were defined according to the length of infinitary conjunctions/disjunctions as well as quantification it allows. In that way, assuming a supply of κ variables to be interpreted as ranging over a nonempty domain, one includes in the inductive definition of formulas an infinitary clause for conjunctions and disjunctions, namely, whenever the ordinal indexed sequence $A_0, ..., A_\delta, ...$ of formulas has length less than κ, one can form the infinitary conjunction/disjunction of them to produce a formula. Analogously, whenever an ordinal indexed sequence of variables has length less than λ, one can introduce one of the quantifiers $\forall$ or $\exists$ together with the sequence of variables in front of a formula to produce a new formula. One also stipulates that κ be a regular cardinal, so that the length of any well-formed formula is less than κ itself.

It is then a natural question to ask for which of these infinitary languages one can provide a notion of provability for which a form of completeness theorem can be proven, in terms, for example, of the obvious Tarskian semantics associated to them. In [Kar64], Karp proves completeness theorems for the classical logic $\mathcal{L}_{\kappa,\kappa}$ within a Hilbert-style system including the distributivity and the dependent choice axioms. These axioms consist of the following schemata:

1. $A \to [B \to A]$
2. $[A \to [B \to C] \to [[A \to B] \to [A \to C]]]$
3. $[\neg B \to \neg A] \to [A \to B]$
4. $[\bigwedge_{i<\alpha} [A \to A_i]] \to [A \to \bigwedge_{i<\alpha} A_i]$
5. $[\bigwedge_{i<\alpha} A_i] \to A_j$
6. $[\forall \mathbf{x} [A \to B] \to [A \to \forall \mathbf{x} B]]$
   provided no variable in $\mathbf{x}$ occurs free in A;

7. $\forall \mathbf{x} A \rightarrow S_f(A)$

   where $S_f(A)$ is a substitution based on a function f from $\mathbf{x}$ to the terms of the language;

8. Equality axioms:

   (a) $t = t$

   (b) $[\bigwedge_{i<\alpha} t_i = t'_i] \rightarrow [\phi(t_0, ..., t_\xi, ...) = \phi(t'_0, ..., t'_\xi, ...)]$

   (c) $[\bigwedge_{i<\alpha} t_i = t'_i] \rightarrow [P(t_0, ..., t_\xi, ...) \rightarrow P(t'_0, ..., t'_\xi, ...)]$
   for each $\alpha < \kappa$, where t, t_i are terms and ϕ is a function symbol of arity α and P a relation symbol of arity α;

9. Classical distributivity axiom[1]:

$$\bigwedge_{i<\gamma} \bigvee_{j<\gamma} \psi_{ij} \rightarrow \bigvee_{f \in \gamma^\gamma} \bigwedge_{i<\gamma} \psi_{if(i)}$$

10. Classical dependent choice axiom:

$$\bigwedge_{\alpha<\gamma} \forall_{\beta<\alpha} \mathbf{x}_\beta \exists \mathbf{x}_\alpha \psi_\alpha \rightarrow \exists_{\alpha<\gamma} \mathbf{x}_\alpha \bigwedge_{\alpha<\gamma} \psi_\alpha$$

   provided the sets $\mathbf{x}_\alpha$ are pairwise disjoint and no variable in $\mathbf{x}_\alpha$ is free in ψ_β for $\beta < \alpha$.

The inference rules are modus ponens, conjunction introduction and generalization.

In the same way that for finitary languages proofs are finitary objects, the right metatheory to study formal proofs of infinitary languages is that of sets hereditarily of cardinal less than κ. Similarly, Gödel numberings of finitary formulas can be generalized to the infinitary case if one uses (not necessarily finite) ordinal numbers (see [Kar64]), by considering one-to-one functions from the symbols of the language into κ. It is then possible to consider Gödel numbers of formulas and prove that they correspond to those sets hereditarily of cardinal less than κ that satisfy a precise ordinary predicate in a certain metalanguage. Moreover, Gödel numbers of provable formulas must satisfy a precise provability predicate in such metalanguage.

[1]Throughout this work the notation α^β for ordinals α, β will always denote the set of functions $f : \beta \rightarrow \alpha$, and should not be confused with ordinal exponentiation.

The development of [Kar64] is classical, that is, the infinitary systems considered formalize infinitary classical logic. Intuitionistic systems of infinitary propositional logic using countable many conjunctions and disjunctions was studied in [Nad78]. Our purpose here is to study systems for the intuitionistic general case, together with corresponding completeness theorems.

2.1 Infinitary first-order systems

Let κ be an inaccessible cardinal (we consider ω to be inaccessible as well, so that our account embodies in particular the finitary case). The syntax of intuitionistic κ-first-order logics $\mathcal{L}_{\kappa,\kappa}$ consists of a (well-ordered) set of sorts and a set of function and relation symbols, these latter together with the corresponding type, which is a subset with less than κ many sorts. Therefore, we assume that our signature may contain relation and function symbols on $\gamma < \kappa$ many variables, and we suppose there is a supply of κ many fresh variables of each sort. Terms and atomic formulas are defined as usual, and general formulas are defined inductively according to the following:

Definition 2.1.1. Let κ be an inaccessible cardinal. If $\phi, \psi, \{\phi_\alpha : \alpha < \gamma\}$ (for each $\gamma < \kappa$) are formulas of $\mathcal{L}_{\kappa,\kappa}$, the following are also formulas: $\bigwedge_{\alpha<\gamma} \phi_\alpha$, $\bigvee_{\alpha<\gamma} \phi_\alpha$, $\phi \to \psi$, $\forall_{\alpha<\gamma} x_\alpha \phi$ (also written $\forall \mathbf{x}_\gamma \phi$ if $\mathbf{x}_\gamma = \{x_\alpha : \alpha < \gamma\}$), $\exists_{\alpha<\gamma} x_\alpha \phi$ (also written $\exists \mathbf{x}_\gamma \phi$ if $\mathbf{x}_\gamma = \{x_\alpha : \alpha < \gamma\}$).

The inductive definition of formulas allows to place them in hierarchies or levels up to κ. Formulas in a successor level are built using the clauses of the definition from formulas in the previous level, while at limit levels one takes the union of all formulas in all levels so far defined. Proofs by induction on the complexity of the formula are proofs by transfinite induction on the least level of the formulas.

The infinitary systems that we will use in categorical logic for our purposes have all the rules of finitary first-order logic, except that in the case of $\mathcal{L}_{\kappa,\kappa}$ we allow infinite sets of variables as contexts of the sequents. Since the variables of each sort are assumed to in correspondence with (the elements of) κ, each subset of variables comes with an inherited well-order, which we will assume as given when we quantify over sets of variables. There are two special types of formulas that one can consider. One is the class of κ-*regular* formulas (see [Mak90]), which are those build of atomic formulas, κ-conjunctions and κ-existential quantification. Adding κ-disjunction results in the class of κ-*coherent* formulas, which we shall introduce in more detail later.

We use sequent style calculus to formulate the axioms of first-order logic, as can be found, e.g., in [Joh02], D1.3. The system for κ-first order logic is described below. Its key feature and the difference with the system of Karp essentially resides, besides being an intuitionistic system, in the introduction of the transfinite transitivity rule, which, as we shall see, is an intuitionistic way of merging the classical distributivity and dependent choice axioms. The intuitive meaning of this rule will be further explained after the following:

Definition 2.1.2. The system of axioms and rules for κ-first-order logic consists of

1. Structural rules:

   (a) Identity axiom:

   $$\phi \vdash_{\mathbf{x}} \phi$$

   (b) Substitution rule:

   $$\frac{\phi \vdash_{\mathbf{x}} \psi}{\phi[\mathbf{s}/\mathbf{x}] \vdash_{\mathbf{y}} \psi[\mathbf{s}/\mathbf{x}]}$$

   where $\mathbf{y}$ is a string of variables including all variables occurring in the string of terms $\mathbf{s}$.

   (c) Cut rule:

   $$\frac{\phi \vdash_{\mathbf{x}} \psi \qquad \psi \vdash_{\mathbf{x}} \theta}{\phi \vdash_{\mathbf{x}} \theta}$$

2. Equality axioms:

   (a)

   $$\top \vdash_x x = x$$

   (b)

   $$(\mathbf{x} = \mathbf{y}) \wedge \phi \vdash_{\mathbf{z}} \phi[\mathbf{y}/\mathbf{x}]$$

   where $\mathbf{x}$, $\mathbf{y}$ are contexts of the same length and type and $\mathbf{z}$ is any context containing $\mathbf{x}$, $\mathbf{y}$ and the free variables of ϕ.

3. Conjunction axioms and rules:

$$\bigwedge_{i<\gamma} \phi_i \vdash_{\mathbf{x}} \phi_j$$

$$\frac{\{\phi \vdash_{\mathbf{x}} \psi_i\}_{i<\gamma}}{\phi \vdash_{\mathbf{x}} \bigwedge_{i<\gamma} \psi_i}$$

for each cardinal $\gamma < \kappa$;

$$\frac{\phi \vdash_{\mathbf{x}} \psi \qquad \phi \vdash_{\mathbf{x}} \theta}{\phi \vdash_{\mathbf{x}} \psi \wedge \theta}$$

4. Disjunction axioms and rules:

$$\phi_j \vdash_{\mathbf{x}} \bigvee_{i<\gamma} \phi_i$$

$$\frac{\{\phi_i \vdash_{\mathbf{x}} \theta\}_{i<\gamma}}{\bigvee_{i<\gamma} \phi_i \vdash_{\mathbf{x}} \theta}$$

for each cardinal $\gamma < \kappa$;

5. Implication rule:

$$\frac{\phi \wedge \psi \vdash_{\mathbf{x}} \theta}{\overline{\phi \vdash_{\mathbf{x}} \psi \to \theta}}$$

6. Existential rule:

$$\frac{\phi \vdash_{\mathbf{xy}} \psi}{\overline{\exists \mathbf{y} \phi \vdash_{\mathbf{x}} \psi}}$$

where no variable in $\mathbf{y}$ is free in ψ.

7. Universal rule:

$$\frac{\phi \vdash_{\mathbf{xy}} \psi}{\phi \vdash_{\mathbf{x}} \forall \mathbf{y} \psi}$$

where no variable in $\mathbf{y}$ is free in ϕ.

8. Transfinite transitivity:

$$\frac{\begin{array}{ll} \phi_f \vdash_{\mathbf{y}_f} \bigvee_{g \in \gamma^{\beta+1}, g|_\beta = f} \exists \mathbf{x}_g \phi_g & \beta < \gamma, f \in \gamma^\beta \\ \phi_f \dashv\vdash_{\mathbf{y}_f} \bigwedge_{\alpha < \beta} \phi_{f|_\alpha} & \beta < \gamma, \text{ limit } \beta, f \in \gamma^\beta \end{array}}{\phi_\emptyset \vdash_{\mathbf{y}_\emptyset} \bigvee_{f \in \gamma^\gamma} \exists_{\beta < \gamma} \mathbf{x}_{f|_{\beta+1}} \bigwedge_{\beta < \gamma} \phi_{f|_\beta}}$$

for each cardinal $\gamma < \kappa$, where $\mathbf{y}_f$ is the canonical context of ϕ_f, provided that, for every $f \in \gamma^{\beta+1}$, $FV(\phi_f) = FV(\phi_{f|_\beta}) \cup \mathbf{x}_f$ and $\mathbf{x}_{f|_{\beta+1}} \cap FV(\phi_{f|_\beta}) = \emptyset$ for any $\beta < \gamma$, as well as $FV(\phi_f) = \bigcup_{\alpha<\beta} FV(\phi_{f|_\alpha})$ for limit β. Note that we assume that there is a fixed well-ordering of γ^γ for each $\gamma < \kappa$.

In this formulation the double line indicates a bidirectional rule. Note that in full infinitary first-order logic we can dispense with the use of sequents and treat $\phi \vdash_{\mathbf{x}} \psi$ as simply $\forall \mathbf{x}(\phi \to \psi)$. Conversely, any formula $\phi(\mathbf{x})$ can be interpreted as the sequent $\top \vdash_{\mathbf{x}} \phi$, thereby obtaining a translation with Hilbert style systems.

The transfinite transitivity rule can be understood as follows. Consider $\gamma^{\leq\gamma}$, the γ-branching tree of height γ, i.e., the poset of functions $f : \beta \to \gamma$ for $\beta \leq \gamma$ with the order given by inclusion. Suppose there is an assignment of formulas ϕ_f to each node f of $\gamma^{\leq\gamma}$. Then the rule expresses that if the assignment is done in a way that the formula assigned to each node entails the join of the formulas assigned to its immediate successors, and if the formula assigned to a node in a limit level is equivalent to the meet of the formulas assigned to its predecessors, then the formula assigned to the root entails the join of the formulas assigned to the nodes in level γ.

In full first-order logic the transfinite transitivity rule can be replaced by the axiom schema, for each $\gamma < \kappa$:

$$\bigwedge_{f\in\gamma^\beta,\beta<\gamma} \forall \mathbf{y}_f \left(\phi_f \to \bigvee_{g\in\gamma^{\beta+1}, g|_\beta = f} \exists \mathbf{x}_g \phi_g\right)$$

$$\wedge \bigwedge_{\beta<\gamma,\ \text{limit } \beta, f\in\gamma^\beta} \forall \mathbf{y}_f \left(\phi_f \leftrightarrow \bigwedge_{\alpha<\beta} \phi_{f|_\alpha}\right)$$

$$\vdash_{\mathbf{y}_\emptyset} \phi_\emptyset \to \bigvee_{f\in\gamma^\gamma} \exists_{\beta<\gamma} \mathbf{x}_{f|_{\beta+1}} \bigwedge_{\beta<\gamma} \phi_{f|_\beta}.$$

There are two particular cases of the transfinite transitivity rule which are of interest:

1. Distributivity rule:

$$\frac{\begin{array}{ll} \phi_f \vdash_{\mathbf{x}} \bigvee_{g\in\gamma^{\beta+1}, g|_\beta = f} \phi_g & \beta<\gamma, f\in\gamma^\beta \\ \phi_f \dashv\vdash_{\mathbf{x}} \bigwedge_{\alpha<\beta} \phi_{f|_\alpha} & \beta<\gamma,\ \text{limit } \beta, f\in\gamma^\beta \end{array}}{\phi_\emptyset \vdash_{\mathbf{x}} \bigvee_{f\in\gamma^\gamma} \bigwedge_{\beta<\gamma} \phi_{f|_\beta}}$$

   for each $\gamma<\kappa$ (we assume that there is a fixed well-ordering of γ^γ for each $\gamma<\kappa$).

2. Dependent choice:

$$\frac{\begin{array}{ll} \phi_\beta \vdash_{\mathbf{y}_\beta} \exists \mathbf{x}_{\beta+1} \phi_{\beta+1} & \beta<\gamma \\ \phi_\beta \dashv\vdash_{\mathbf{y}_\beta} \bigwedge_{\alpha<\beta} \phi_\alpha & \beta\leq\gamma,\ \text{limit } \beta \end{array}}{\phi_\emptyset \vdash_{\mathbf{y}_\emptyset} \exists_{\beta<\gamma} \mathbf{x}_{\beta+1} \phi_\gamma}$$

   for each $\gamma<\kappa$, where $\mathbf{y}_\beta$ is the canonical context of ϕ_β, provided that, for every $f\in\gamma^{\beta+1}$, $FV(\phi_f) = FV(\phi_{f|_\beta}) \cup \mathbf{x}_f$ and $\mathbf{x}_{f|_{\beta+1}} \cap FV(\phi_{f|_\beta}) = \emptyset$ for any $\beta<\gamma$, as well as $FV(\phi_f) = \bigcup_{\alpha<\beta} FV(\phi_{f|_\alpha})$ for limit β.

Again, if implication is available in the fragment we are considering, we can instead replace the distributivity rule by an axiom schema expressible with a single sequent, for each $\gamma < \kappa$:

$$\bigwedge_{f \in \gamma^{\beta}, \beta < \gamma} \left(\phi_f \to \bigvee_{g \in \gamma^{\beta+1}, g|_{\beta} = f} \phi_g \right)$$

$$\wedge \bigwedge_{\beta < \gamma, \text{ limit } \beta, f \in \gamma^{\beta}} \left(\phi_f \leftrightarrow \bigwedge_{\alpha < \beta} \phi_{f|_{\alpha}} \right) \vdash_{\mathbf{x}} \phi_{\emptyset} \to \bigvee_{f \in \gamma^{\gamma}} \bigwedge_{\beta < \gamma} \phi_{f|_{\beta}}$$

In turn, the rule of dependent choice has as a particular case the rule of choice:

$$\frac{\phi \vdash_{\mathbf{x}} \bigwedge_{\beta < \gamma} \exists \mathbf{x}_{\beta} \phi_{\beta}}{\phi \vdash_{\mathbf{x}} \exists_{\beta < \gamma} \mathbf{x}_{\beta} \bigwedge_{\beta < \gamma} \phi_{\beta}}$$

where the $\mathbf{x}_{\beta}$ are disjoint canonical contexts of the ϕ_{β}. This can be seen by applying dependent choice to the formulas $\psi_{\beta} = \phi \wedge \bigwedge_{\alpha < \beta} \phi_{\alpha+1}$. From this rule one can also derive Frobenius axiom:

$$\phi \wedge \exists \mathbf{y} \psi \vdash_{\mathbf{x}} \exists \mathbf{y} (\phi \wedge \psi)$$

where no variable in $\mathbf{y}$ is in the context $\mathbf{x}$.

Lemma 2.1.3. *All instances of the classical distributivity axiom:*

$$\bigwedge_{i < \gamma} \bigvee_{j < \gamma} \psi_{ij} \vdash_{\mathbf{x}} \bigvee_{f \in \gamma^{\gamma}} \bigwedge_{i < \gamma} \psi_{i f(i)}$$

are derivable from those of the axiom schema:

$$\bigwedge_{f \in \gamma^{\beta}, \beta < \gamma} \left(\phi_f \to \bigvee_{g \in \gamma^{\beta+1}, g|_{\beta} = f} \phi_g \right)$$

$$\wedge \bigwedge_{\beta < \gamma, \text{ limit } \beta, f \in \gamma^{\beta}} \left(\phi_f \leftrightarrow \bigwedge_{\alpha < \beta} \phi_{f|_{\alpha}} \right) \vdash_{\mathbf{x}} \phi_{\emptyset} \to \bigvee_{f \in \gamma^{\gamma}} \bigwedge_{\beta < \gamma} \phi_{f|_{\beta}}$$

Proof. Assign to the nodes of the tree $\gamma^{<\gamma}$ the following formulas: to the immediate succesors of a node ϕ_f, for $f \in \gamma^\beta$, assign the formulas $\psi_{\beta j}$, then set $\phi_\emptyset = \top$, and $\phi_f = \bigwedge_{\alpha<\beta} \phi_{f|_\alpha}$ for $f \in \gamma^\beta$ and limit β. Then the antecedent of our premise derives the consequent $\bigvee_{g \in \gamma^{\beta+1}, g|_\beta = f} \phi_g$, and in particular $\phi_f \to \bigvee_{g \in \gamma^{\beta+1}, g|_\beta = f} \phi_g$, so that applying the distributivity rule we get $\bigvee_{f \in \gamma^\gamma} \bigwedge_{i<\gamma} \psi_{if(i)}$, as we wanted.

□

Remark 2.1.4. Note that it follows from the classical distributivity law not only that $\top \vdash_{\mathbf{x}} \phi \wedge \bigvee_{i<\gamma} \psi_i \to \bigvee_{i<\gamma} (\phi \wedge \psi_i)$, but also that we have $\top \vdash_{\mathbf{x}} \bigwedge_{i<\gamma} (\phi \vee \psi_i) \to \phi \vee \bigwedge_{i<\gamma} \psi_i$, which does not generally hold in a complete Heyting algebra. Indeed, we can write $\phi \vee \psi_i$ as $\phi \vee \psi_i \vee \bot \ldots$, and an application of classical distributivity shows that this can be rewritten as a disjunction of three types of disjuncts: $\bigwedge_{i<\gamma} \phi = \phi$, $\bigwedge_{i<\gamma} \psi_i$ and disjuncts of the form $\phi \wedge \bigwedge_{i \neq j} \psi_i$, and this latter type of disjuncts all imply ϕ.

Also, we prove below that within full first-order logic the dependent choice schema derives the form of dependent choice schema in [Kar64]:

$$\bigwedge_{\alpha<\gamma} \forall_{\beta<\alpha} \mathbf{x}_\beta \exists \mathbf{x}_\alpha \psi_\alpha \vdash_{\mathbf{x}} \exists_{\alpha<\gamma} \mathbf{x}_\alpha \bigwedge_{\alpha<\gamma} \psi_\alpha$$

for each $\gamma < \kappa$, provided the sets of variables $\mathbf{x}_\alpha$ are pairwise disjoint and no variable in $\mathbf{x}_\alpha$ appears in $\mathbf{x}_\beta$ for any $\beta < \alpha$.

Note that if universal quantification is available, we can instead replace the dependent choice rule by an axiom schema expressible with a single sequent for each $\gamma < \kappa$:

$$\bigwedge_{\beta<\gamma} \forall \mathbf{y}_\beta \left(\phi_\beta \to \exists \mathbf{x}_{\beta+1} \phi_{\beta+1} \right)$$

$$\wedge \bigwedge_{\beta \leq \gamma,\text{ limit } \beta} \forall \mathbf{y}_\beta \left(\phi_\beta \leftrightarrow \bigwedge_{\alpha<\beta} \phi_\alpha \right) \vdash_{\mathbf{y}_\emptyset} \phi_\emptyset \to \exists_{\alpha<\gamma} \mathbf{x}_\alpha \phi_\gamma.$$

Lemma 2.1.5. *All instances of the classical dependent choice axiom:*

$$\bigwedge_{\alpha<\gamma} \forall_{\beta<\alpha} \mathbf{x}_\beta \exists \mathbf{x}_\alpha \psi_\alpha \vdash_{\mathbf{x}} \exists_{\alpha<\gamma} \mathbf{x}_\alpha \bigwedge_{\alpha<\gamma} \psi_\alpha$$

are derivable from those of the axiom schema:

$$\bigwedge_{\beta<\gamma} \forall \mathbf{y}_\beta \left(\phi_\beta \to \exists \mathbf{x}_{\beta+1} \phi_{\beta+1} \right)$$

$$\wedge \bigwedge_{\beta \leq \gamma,\ limit\ \beta} \forall \mathbf{y}_\beta \left(\phi_\beta \leftrightarrow \bigwedge_{\alpha < \beta} \phi_\alpha \right) \vdash_{\mathbf{y}_\emptyset} \phi_\emptyset \to \exists_{\alpha<\gamma} \mathbf{x}_\alpha \phi_\gamma$$

Proof. Suppose that we have $\bigwedge_{\alpha<\gamma} \forall_{\beta<\alpha} \mathbf{x}_\beta \exists \mathbf{x}_\alpha \psi_\alpha$. Define by transfinite induction on γ the following formulas: set $\phi_\emptyset = \top$, $\phi_{\beta+1} = \psi_\alpha$ if $\alpha = \beta + 1$ is a successor ordinal, and $\phi_\alpha = \bigwedge_{\beta<\alpha} \psi_\beta$ if α is a limit ordinal. Then we have $\top \vdash_{\mathbf{y}_\beta} \exists \mathbf{x}_{\beta+1} \phi_{\beta+1}$, and in particular, $\top \vdash_{\mathbf{y}_\beta} \phi_\beta \to \exists \mathbf{x}_{\beta+1} \phi_{\beta+1}$. Applying generalization and the dependent choice rule we then get $\exists_{\alpha<\gamma} \mathbf{x}_\alpha \bigwedge_{\alpha<\gamma} \psi_\alpha$, as desired. □

There is also a version of the deduction theorem that holds here:

Lemma 2.1.6. *Let Σ be a set of sequents and let γ be a sentence. If the theory $\Sigma \cup \{\top \vdash \gamma\}$ derives the sequent $\phi \vdash_{\mathbf{x}} \psi$, then the theory Σ derives the sequent $\phi \wedge \gamma \vdash_{\mathbf{x}} \psi$.*

Proof. Straightforward induction on the length of the derivation. □

We are going to need three more fragments to work with:

Definition 2.1.7. The κ-$\mathcal{R}eg_\perp$ (resp. κ-pre-Heyting) fragment is the fragment of κ-coherent (resp. κ-first-order) logic that drops the disjunction $\vee$ from the language, and hence drops the rules involving it, but keeps the rule of dependent choice and the ex falso quodlibet axiom $\perp \vdash_{\mathbf{x}} \phi$). If $\perp$ and the ex falso quodlibet axiom are also dropped from the κ-$\mathcal{R}eg_\perp$ fragment, we get the κ-regular fragment.

2.2 κ-coherent categories

We will start now the study of the parallel between fragments of infinitary logics and the corresponding categorical notions. As it is customary in categorical logic, one considers several fragments of infinitary logic $\mathcal{L}_{\infty,\infty}$ for most of which one can prove completeness theorems in terms of the usual Tarski semantics. While classical theories over the finite quantifier fragment of $\mathcal{L}_{\infty,\infty}$ can be proved, assuming Grothendieck's universe axiom, to be complete (see, e.g., [Gre75]), one can also see, as presented, for example, in [Joh02], that the cartesian, regular and coherent fragments all enjoy such completeness theorems. One notable exception is geometric logic, the fragment obtained by adding to the coherent fragment arbitrary disjunctions. For instance, there are several examples of consistent geometric theories which happen to have no set-valued models. As we shall prove, the only obstacle to obtaining a derivation of $\perp$ in such theories is essentially the lack of rules handling infinitary conjunction and existential quantification, together with appropriate associated rules. For example, the following geometric theory:

$$\top \vdash \bigvee_{\alpha \in \omega_1} P_{n,\alpha}$$

for each $n \in \omega$,

$$\top \vdash \bigvee_{n \in \omega} P_{n,\alpha}$$

for each $\alpha \in \omega_1$,

$$P_{n,\alpha} \wedge P_{n,\beta} \vdash \bot$$

for each $n \in \omega$ and $\alpha \neq \beta$, cannot have set models and yet $\bot$ cannot be derived, as it can be proven to be consistent (see [Joh02]).

We will prove that the addition of these extra rules will be enough, under certain large cardinal assumptions which are in a sense unavoidable, to recover set-valued completeness in all cases. Hence, this brings to consideration a new fragment of infinitary logic which we shall call *infinitary coherent* (resp. *κ-coherent*).

Definition 2.2.1. The infinitary coherent (resp. κ-coherent) fragment of full intuitionistic infinitary first-order logic is the fragment of those sequents where formulas are infinitary coherent (resp. κ-coherent), i.e., only use $\bigwedge$, $\bigvee$, $\exists$, where the transfinite transitivity rule is restricted to instantiations on infinitary coherent (resp. κ-coherent) formulas only, and where disjunctions and conjunctions are indexed by arbitrary ordinals (resp. ordinals less than κ).

The κ-coherent fragment of first-order logic, which is an extension of the usual finitary coherent fragment, has a corresponding category which we are now going to define. Following [Mak90], consider a κ-chain in a category $\mathcal{C}$ with κ-limits, i.e., a diagram $\Gamma : \gamma^{op} \to \mathcal{C}$ specified by morphisms $(h_{\beta,\alpha} : C_\beta \to C_\alpha)_{\alpha \leq \beta < \gamma}$ such that the restriction $\Gamma|_\beta$ is a limit diagram for every limit ordinal β. We say that the morphisms $h_{\beta,\alpha}$ compose transfinitely, and take the limit projection $f_{\beta,0}$ to be the transfinite composite of $h_{\alpha+1,\alpha}$ for $\alpha < \beta$.

Given a cardinal $\gamma < \kappa$, consider the tree $T = \gamma^{<\gamma}$. We will consider diagrams $F : T^{op} \to \mathcal{C}$, which determine, for each node f, a family of arrows in $\mathcal{C}$, $\{h_{g,f} : C_g \to C_f | f \in \gamma^\beta, g \in \gamma^{\beta+1}, g|_\beta = f\}$. A κ-family of morphisms with the same codomain is said to be *jointly covering* if the union of the images of the morphisms is the whole codomain. We say that a diagram $F : T^{op} \to \mathcal{C}$ is *proper* if the $\{h_{g,f} : f \in T\}$ are jointly covering and, for limit β, $h_{f,\emptyset}$ is the transfinite composition of the $h_{f|_{\alpha+1},f|_\alpha}$ for $\alpha + 1 < \beta$. Given a proper diagram, we say that the families $\{h_{g,f} : f \in T\}$ compose transfinitely, and refer to the projections $\{h_{g,\emptyset} | g \in \gamma^\gamma\}$ as the transfinite composites of these families. If in a proper diagram the transfinite composites of the κ-families of morphisms form itself a jointly covering family, we will say that the diagram is completely proper.

Definition 2.2.2. A κ-coherent category is a κ-complete coherent category with κ-complete subobject lattices where unions of cardinality less than κ are stable under pullback, and where every proper diagram is completely proper, i.e., the transfinite composites of jointly covering κ-families of morphisms form a jointly covering family.

The latter property, which is the categorical analogue of the transfinite transitivity rule, can be considered as an exactness property of **Set**, generalizing the property in [Mak90] where the families consisted of single morphisms. The transfinite transitivity rule expresses that transfinite compositions of covering families (in the Grothendieck topology given by the jointly covering families of less than κ-morphisms) are again covering families; whence its name. It is easy to see that the rule holds in **Set**, and in fact in every presheaf category.

κ-coherent categories have an internal logic, in a signature containing one sort for each object, no relation symbols and one unary function symbol for each arrow, and axiomatized by the following sequents:

$$\top \vdash_x Id_X(x) = x$$

for all objects X (here x is a variable of sort X);

$$\top \vdash_x f(x) = h(g(x))$$

for all triples of arrows such that $f = h \circ g$ (here x is a variable whose sort is the domain of f);

$$\top \vdash_y \exists x f(x) = y$$

for all covers f (here x is a variable whose sort is the domain of f);

$$\top \vdash_x \bigvee_{i<\gamma} \exists y_i m_i(y_i) = x$$

whenever the sort A of x is the union of γ subobjects $m_i : A_i \rightarrowtail A$ (here y_i is a variable of sort A_i);

$$\bigwedge_{i:I \to J} \bar{i}(x_I) = x_J \vdash_{\{x_I : I \in \mathbf{I}\}} \exists x \bigwedge_{I \in \mathbf{I}} \pi_I(x) = x_I$$

$$\bigwedge_{I \in \mathbf{I}} \pi_I(x) = \pi_I(y) \vdash_{x,y} x = y$$

whenever there is a κ-small diagram $\Phi : \mathbf{I} \to \mathcal{C}$, $(\{C_I\}_{I \in \mathbf{I}}, \{\bar{i} : C_I \to C_J)\}_{i:I \to J})$ and a limit cone $\pi : \Delta_C \Rightarrow \Phi$, $(\pi_I : C \to C_I)_{I \in \mathbf{I}}$. Here x_I is a variable of type C_I, and x, y are variables of type C.

Functors preserving this logic, i.e., κ-coherent functors, are just coherent functors which preserve κ-limits and κ-unions of subobjects, and they can be easily seen to correspond to structures of the internal theory in a given κ-coherent category, where we use a straightforward generalization of categorical semantics, to be explained in the next section.

2.3 Categorical semantics

Categorical model theory techniques explore the study of models in arbitrary categories besides the usual category of sets. Unlike classical model theory, the logics one uses for this purpose formulate theories in terms of sequents; the type of theory studied depends on the type of formula one encounters in these sequents. The theories of the fragments mentioned so far all correspond to specific types of categories. We have the κ-regular categories, which are categories with κ-limits, regular epimorphism-monomorphism factorizations stable under pullback and where the transfinite composition of epimorphisms is an epimorphism. A κ-$\mathcal{R}eg_{\bot}$ category has in addition a strict initial object, to interpret $\bot$. The κ-coherent categories have, in addition to this, stable κ-unions of subobjects and satisfy the property that the transfinite composition of jointly covering families is jointly covering. Finally, the κ-Heyting categories have, in addition, right adjoint for pullback functors between subobject lattices, which makes interpreting universal quantification possible.

There is a categorical semantics that one can associate with each type of category and theory, which is usually defined according to some inductive clauses. Following [Joh02], D1.2, given a category $\mathcal{C}$, for each signature Σ of a first order language we can associate the so called Σ-structure within $\mathcal{C}$ in a way that generalizes the **Set**-valued interpretations to all κ-Heyting categories:

Definition 2.3.1. A Σ-structure in $\mathcal{C}$ consists of the following data:

1. for each sort A of variables in Σ there is a corresponding object $M(A)$;

2. for each γ-ary function symbol f there is a morphism $M(f)$: $M(A_1, ..., A_\alpha, ...) = \Pi_{i<\gamma} M(A_i) \to M(B)$;

3. for each γ-ary relation symbol R there is a subobject $M(R) \rightarrowtail M(A_1, ..., A_\alpha, ...)$, where A_i are the sorts corresponding to the individual variables corresponding to R (which will specify, by definition, the type of R).

The Σ-structure will serve as a setup for interpreting all formulas of the language considered. Due to the need of distinguishing the context in which the free variables of the formula occur, for the purpose of a correct interpretation, we shall adopt the notation $(\mathbf{x}, \phi)$ to represent a term/formula ϕ whose free variables occur within $\mathbf{x} = x_1, ..., x_\alpha, ...$. We now define the interpretation of such formulas by induction on their complexity:

Definition 2.3.2. Given a term in context $(\mathbf{x}, s)$ of a κ-first order theory, its interpretation $[\![\mathbf{x}, s]\!]$ within the κ-Heyting category $\mathcal{C}$ is a morphism of $\mathcal{C}$ defined in the following way:

1. If s is a variable, it is necessarily some x_i, and then the corresponding morphism is $[\![\mathbf{x}, x_i]\!] = \pi_i : M(A_0, ..., A_\alpha, ...) \to M(A_i)$, the i-th product projection.

2. If s is a term $f(t_0, ..., t_\alpha, ...)$, where each term t_α is of type C_α, its interpretation is the composite:

$$M(A_0, ..., A_\alpha, ...) \xrightarrow{([\![\mathbf{x}, t_0]\!], ..., [\![\mathbf{x}, t_\alpha]\!], ...)} M(C_0, ..., C_\alpha, ...) \xrightarrow{M(f)} M(B)$$

The interpretation in $\mathcal{C}$ of the formula in context $(\mathbf{x}, \phi)$, where $\mathbf{x} = x_0 ... x_\alpha ...$ and x_i is a variable of sort A_i, is defined as a subobject $[\![\mathbf{x}, \phi]\!] \rightarrowtail M(A_0, ..., A_\alpha, ...)$ in the following way:

1. If ϕ is the formula $R(t_0, ..., t_\alpha, ...)$, where R is a γ-ary relation symbol of type $B_0, ..., B_\alpha, ...$, then $[\![\mathbf{x}, \phi]\!]$ is given by the pullback:

$$\begin{array}{ccc} [\![\mathbf{x}, \phi]\!] & \longrightarrow & M(R) \\ \downarrow & & \downarrow \\ M(A_0, ..., A_\alpha, ...) & \xrightarrow{([\![\mathbf{x}, t_1]\!], ..., [\![\mathbf{x}, t_\alpha]\!], ...)} & M(B_0, ..., B_\alpha, ...) \end{array}$$

2. If ϕ is the formula $s = t$ where s, t are terms of sort B, then $[\![\mathbf{x}, \phi]\!]$ is the equalizer of the arrows:

$$M(A_0, ..., A_\alpha, ...) \underset{[\![\mathbf{x}, t]\!]}{\overset{[\![\mathbf{x}, s]\!]}{\rightrightarrows}} M(B)$$

   Equivalently, $[\![\mathbf{x}, \phi]\!]$ is the pullback of the diagonal $M(B) \rightarrowtail M(B) \times M(B)$ along the morphism $([\![\mathbf{x}, s]\!], [\![\mathbf{x}, t]\!])$.

3. If ϕ is the formula $\bigvee_{i<\gamma} \psi_i$, then $[\![\mathbf{x}, \phi]\!]$ is the union $\bigvee_{i<\gamma} [\![\mathbf{x}, \psi_i]\!]$ in $\mathcal{S}ub(M(A_0, ..., A_\alpha, ...))$. If ϕ is the formula $\bigwedge_{i<\gamma} \psi_i$, then $[\![\mathbf{x}, \phi]\!]$ is the intersection $\bigwedge_{i<\gamma} [\![\mathbf{x}, \psi_i]\!]$ in $\mathcal{S}ub(M(A_0, ..., A_\alpha, ...))$. Similarly, if ϕ is the formula $\neg\psi$, the corresponding subobject is $\neg[\![\mathbf{x}, \psi]\!]$.

4. If ϕ is the formula $(\exists y)\psi$, then $[\![\mathbf{x}, \phi]\!]$ is the image of the composite:

$$[\![\mathbf{x}y, \psi]\!] \rightarrowtail M(A_0, ..., A_\alpha, ..., B) \xrightarrow{\pi} M(A_0, ..., A_\alpha, ...)$$

   where π is the projection to the first γ coordinates. Equivalently, this amounts to applying the left adjoint to the pullback functor $\pi^{-1} : \mathcal{S}ub(M(A_0, ..., A_\alpha, ..., B)) \to \mathcal{S}ub(M(A_0, ..., A_\alpha, ...))$.

5. If ϕ is the formula $(\forall y)\psi$, then $[\![\mathbf{x}, \phi]\!]$ can be obtained by applying to $[\![\mathbf{x}y, \psi]\!]$ the right adjoint to the pullback functor $\pi^{-1} : \mathcal{S}ub(M(A_0, ..., A_\alpha, ..., B)) \to \mathcal{S}ub(M(A_0, ..., A_\alpha, ...))$, where π is the projection to the first γ coordinates. Implication can be seen as a particular case of this right adjoint, by considering in $\mathcal{S}ub(M(A_0, ..., A_\alpha, ...))$ the pullback functor $\phi \wedge - : \mathcal{S}ub(M(A_0, ..., A_\alpha, ...)) \to \mathcal{S}ub(M(A_0, ..., A_\alpha, ...))$.

Functors between the appropriate categories preserving the corresponding structure correpond to models in the codomain category of the internal theory of the domain category. Such functors are called conservative if they reflect isomorphisms, and hence they reflect also the validity of formulas in the corresponding models.

One then has:

Lemma 2.3.3. *κ-coherent logic is sound with respect to models in κ-coherent categories.*

Proof. This is straightforward for all axioms and rules, except for the rule of transfinite transitivity. But here the proof is the natural generalization of that of the soundness of dependent choice, presented in [Mak90] for κ-regular logic. Let $S_{\mathbf{y}_f}$ be the product of the sorts assigned to the variables in $\mathbf{y}_f$ in the structure within a κ-coherent category, and assume that the premises of the transfinite transitivity rule hold there. We must show that the conclusion holds. We can also assume, without loss of generality, that:

$$\phi_g \vdash_{\mathbf{y}_g} \phi_f$$

for each $g \in \gamma^{\beta+1}, g|_\beta = f$; otherwise we can take, for each $f \in \gamma^\beta$:

$$\psi_f = \bigwedge_{\alpha \leq \beta} \phi_{f|_\alpha}$$

which, using the law $\phi \wedge \bigvee_{i<\gamma} \psi_i \to \bigvee_{i<\gamma}(\phi \wedge \psi_i)$ as well as Frobenius axiom can be seen to satisfy the premises of the rule as well, and both this form of distributivity and Frobenius axiom hold in any κ-coherent category because κ-unions and covers are stable under pullback.

Let $m_\alpha : C_\alpha \to S_{\mathbf{y}_f}$ be the monomorphism representing the subobject $[\mathbf{y}_f, \phi_f]$. The assumption we have provides arrows:

$$h_{g,f} : C_g \to C_f$$

for $g \in \gamma^{\beta+1}$, $g|_\beta = f$, and by interpreting the premises of the rule it follows that the arrows:

$$\{h_{g,f} | g \in \gamma^{\beta+1}, g|_\alpha = f\}$$

form a jointly covering family. For a fixed $f \in \gamma^\gamma$ and limit β, the limit of the diagram formed by the $C_{f|_\alpha}$ for $\alpha < \beta$ is given by the intersection in the subobject lattice of $S_{\mathbf{y}_{f|_\beta}}$ of the pullbacks of each m_α along the projections $\pi_{f|_\beta, f|_\alpha} : S_{\mathbf{y}_{f|_\beta}} \to S_{\mathbf{y}_{f|_\alpha}}$. This intersection is in turn given by the subobject:

$$C_{f|_\beta} = \bigwedge_{\alpha<\beta} \phi_{f|_\alpha} \to S_{\mathbf{y}_{f|_\beta}}$$

By the property of the κ-coherent category, the arrows $C_{f|_\beta} \to C_\emptyset$ for $f \in \gamma^\gamma$ form a jointly covering family whenever β is a limit ordinal, and the interpretation of conclusion of the rule is precisely this statement for the case $\beta = \gamma$. This proves the soundness of the rule. $\square$

3. Categorical completeness

3.1 Syntactic categories

The development of syntactic categories for infinite quantifier logics follows precisely the same pattern as the finitary case, except that instead of finite contexts for the objects of the syntactic category of a theory over such logic, we allow arbitrary sets of variables of cardinality less than κ, following, e.g., [Mak90].

Given a κ-coherent theory $\mathbb{T}$, we explain how to define, following [Joh02], D1.4, and [Esp], its syntatic category $\mathcal{C}_{\mathbb{T}}$ and a categorical model $M_{\mathbb{T}}$ inside it, in such a way that a formula in $\mathbb{T}$ will be provable if and only if its interpretation in $\mathcal{C}_{\mathbb{T}}$ is satisfied by the model $M_{\mathbb{T}}$. Formulas shall be considered in suitable contexts, which are (possibly empty) subsets of variables of cardinality less than κ containing the free variables of the formula. We will say that two formulas in context $(\mathbf{x}, \phi), (\mathbf{y}, \psi)$ are α-equivalent if the second has been obtained from the first after renaming the bound variables of ϕ and the variables in the context (some of them appearing as free variables in ϕ). We take the objects of $\mathcal{C}_{\mathbb{T}}$ to be the α-equivalence classes of formulas $(\mathbf{x}, \phi)$. To describe the morphisms, consider two objects $[\mathbf{x}, \phi], [\mathbf{y}, \psi]$, and assume, without loss of generality, that their set of variables $\mathbf{x}, \mathbf{y}$ are disjoint. Consider now a formula θ that satisfies the following conditions:

a) Its free variables are amongst $\mathbf{xy}$.
b) The following sequents are provable in $\mathbb{T}$:

$$\theta(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{xy}} \phi(\mathbf{x}) \wedge \psi(\mathbf{y})$$
$$\phi(\mathbf{x}) \vdash_{\mathbf{xy}} \exists \mathbf{y}(\theta(\mathbf{x}, \mathbf{y}))$$
$$\theta(\mathbf{x}, \mathbf{y}) \wedge \theta(\mathbf{x}, \mathbf{z}/\mathbf{y}) \vdash_{\mathbf{xyz}} (\mathbf{y} = \mathbf{z})$$

Define now the morphisms between $[\mathbf{x}, \phi]$ and $[\mathbf{y}, \psi]$ to be the provable-equivalence class of all those formulas of $\mathbb{T}$ that satisfy conditions a) and b) above.

The idea behind this definition is to allow only those morphisms that are exactly needed for our purposes. More precisely, the first formula in condition b) restricts the interpretation $[\![\theta(\mathbf{x},\mathbf{y})]\!]$ in any model to be a subobject of $[\![\phi(\mathbf{x}) \wedge \psi(\mathbf{y})]\!]$, while the last two formulas imply, if the category has finite limits, that it will be the graph of a morphism from $[\![\phi(\mathbf{x})]\!]$ to $[\![\psi(\mathbf{y})]\!]$. Because of the particular construction of the category $\mathcal{C}_{\mathbb{T}}$, this says exactly that the class $[\mathbf{xy}, \theta(\mathbf{x},\mathbf{y})]$ is a morphism from $[\mathbf{x}, \phi(\mathbf{x})]$ to $[\mathbf{y}, \psi(\mathbf{y})]$.

The composite of two morphisms:

$$[\mathbf{x}, \phi] \xrightarrow{[\mathbf{xy},\theta]} [\mathbf{y}, \psi] \xrightarrow{[\mathbf{yz},\delta]} [\mathbf{z}, \eta]$$

is defined to be the class $[\mathbf{xz}, \exists\mathbf{y}(\theta \wedge \delta)]$. It can be verified that this definition does not depend on the choice of representatives θ, δ and that this morphism so defined satisfies conditions a) and b) above. It can also be verified that composition of morphisms is associative. Finally, the identity morphism on an object $[\mathbf{x}, \phi]$ can be defined to be arrow:

$$[\mathbf{x}, \phi] \xrightarrow{[\mathbf{xy},\phi(\mathbf{x})\wedge(\mathbf{x}=\mathbf{y})]} [\mathbf{y}, \phi(\mathbf{y}/\mathbf{x})]$$

Again, it is easily checked that this morphism satisfies condition a) and b) and that it is the unity for composition. Also, note that these definitions do not depend on the choices of representatives in each class. This makes $\mathcal{C}_{\mathbb{T}}$ a small category.

Our goal is to relate syntactical provability in $\mathbb{T}$ with semantic validity in the categorical model $M_{\mathbb{T}}$ to be defined. One aspect of this relation is given by the following lemma, which highlights the syntactical properties of $\mathcal{C}_{\mathbb{T}}$:

Lemma 3.1.1. *1) A morphism $[\mathbf{xy}, \theta] : [\mathbf{x}, \phi] \to [\mathbf{y}, \psi]$ is an isomorphism if and only if $[\mathbf{yx}, \theta] : [\mathbf{y}, \psi] \to [\mathbf{x}, \phi]$ is a valid morphism in $\mathcal{C}_{\mathbb{T}}$ (i.e., it satisfies conditions a) and b) of the definition of morphism).*
2) A morphism $[\mathbf{xy}, \theta] : [\mathbf{x}, \phi] \to [\mathbf{y}, \psi]$ is a monomorphism if and only if the sequent $\theta(\mathbf{x},\mathbf{y}) \wedge \theta(\mathbf{z},\mathbf{y}) \vdash_{\mathbf{xyz}} \mathbf{x} = \mathbf{z})$ is provable in $\mathbb{T}$.
3) Every subobject of $[\mathbf{y}, \phi]$ is isomorphic to one of the form:

$$[\mathbf{x}, \psi] \rightarrowtail^{[\psi\wedge(\mathbf{x}=\mathbf{y})]} [\mathbf{y}, \phi]$$

where ψ is such that the sequent $\psi(\mathbf{y}) \vdash_{\mathbf{y}} \phi(\mathbf{y})$ is provable in $\mathbb{T}$. More-

over, any two subobjects $[\mathbf{y}, \psi], [\mathbf{y}, \eta]$ *in* $\mathcal{S}ub([\mathbf{y}, \phi])$ *satisfy* $[\mathbf{y}, \psi] \leq [\mathbf{y}, \eta]$ *if and only if the sequent* $\psi(\mathbf{y}) \vdash_{\mathbf{y}} \eta(\mathbf{y})$ *is provable in* $\mathbb{T}$.

Proof. To prove 1), suppose $[\mathbf{yx}, \theta]$ is a valid morphism from $[\mathbf{y}, \psi]$ to $[\mathbf{x}, \phi]$. Then it can be easily checked that $[\mathbf{yx}, \theta]$ itself is an inverse for $[\mathbf{xy}, \theta]$. Conversely, if $[\mathbf{xy}, \theta] : [\mathbf{x}, \phi] \to [\mathbf{y}, \psi]$ has an inverse $[\mathbf{yx}, \delta]$ (which is a valid morphism), then it can be verified that θ and δ are necessarily provable equivalent in $\mathbb{T}$, from which the result follows.

To prove 2), construct the kernel pair of $[\mathbf{xy}, \theta] : [\mathbf{x}, \phi] \to [\mathbf{y}, \psi]$, which, using the construction of products and equalizers, can be verified to be the class $[\mathbf{xz}, \exists \mathbf{y}(\theta(\mathbf{x}, \mathbf{y}) \wedge \theta(\mathbf{z}, \mathbf{y}))]$. Then, as can be easily checked, the provability of the stated sequent is equivalent, by 1), to the fact that the diagonal morphism from $[\mathbf{x}, \phi]$ to this kernel pair is an isomorphism, which is in turn equivalent to the fact that $[\mathbf{xy}, \theta]$ is a monomorphism.

Finally, suppose we have a monomorphism $[\mathbf{xy}, \theta] : [\mathbf{x}, \psi] \rightarrowtail [\mathbf{y}, \phi]$. By 1), the morphism $[\mathbf{xy}, \theta] : [\mathbf{x}, \psi] \to [\mathbf{y}, \exists \mathbf{x} \theta(\mathbf{x}, \mathbf{y})]$ is an isomorphism. Then, composing its inverse with the original monomorphism we have a subobject of the stated form, where $\psi(\mathbf{y})$ is the formula $\exists \mathbf{x} \theta(\mathbf{x}, \mathbf{y})$. Now, two subobjects $[\mathbf{y}, \psi], [\mathbf{y}, \eta]$ of $[\mathbf{y}, \phi]$ satisfy $[\mathbf{y}, \psi] \leq [\mathbf{y}, \eta]$ if and only if there exists a monomorphism $[\mathbf{y}, \psi] \rightarrowtail [\mathbf{y}, \eta]$, which by the previous argument must have the form $[\psi' \wedge (\mathbf{x} = \mathbf{y})] : [\mathbf{x}, \psi'] \rightarrowtail [\mathbf{y}, \eta]$ for some ψ'. But then, since ψ and ψ' must be provable equivalent, this is a valid morphism if and only if the sequent $\psi(\mathbf{y}) \vdash_{\mathbf{y}} \eta(\mathbf{y})$ is provable in $\mathbb{T}$. This completes the proof of 3). □

To construct the desired model $M_{\mathbb{T}}$ in the syntactic category of $\mathbb{T}$, note that there is a natural Σ-structure assigning to the sort A the formula $[x, \top]$ where x is a variable of sort A, and to the relation symbols R over variables $\mathbf{x} = x_1, ..., x_\alpha, ...$ of sorts $A_, ..., A_\alpha, ...$ respectively, the subobject $[\mathbf{x}, R(x_1, ..., x_\alpha, ...)] \rightarrowtail [\mathbf{x}, \top]$. We have now finally gotten to the important relationship between syntactic provability and semantic validity in $M_{\mathbb{T}}$:

Proposition 3.1.2. *The sequent* $\phi(\mathbf{x}) \vdash_{\mathbf{x}} \psi(\mathbf{x})$ *is satisfied by the* Σ*-structure* $M_{\mathbb{T}}$ *if and only if it is provable in* $\mathbb{T}$. *Consequently, a formula* $\eta(\mathbf{x})$ *has full extension in* $M_{\mathbb{T}}$ *if and only if it is provable in* $\mathbb{T}$.

Proof. By definition, the stated sequent is satisfied by $M_{\mathbb{T}}$ if and only if the corresponding subobjects in the interpretation satisfy $[\![\mathbf{x}, \phi]\!] \leq [\![\mathbf{x}, \psi]\!]$. By the construction of $M_{\mathbb{T}}$, a straightforward induction on the complexity of ϕ proves that the interpretation $[\![\mathbf{x}, \phi]\!]$ is the subobject

$[\mathbf{x}, \phi] \rightarrowtail [\mathbf{x}, \top]$. For example, the base of the induction corresponds to the verification of this property for atomic formulas. If $[\mathbf{x}, \phi]$ is the formula $[R(x_1, ..., x_\alpha, ...]$ (which in the described interpretation has a sort corresponding to $[x_1, ..., x_\alpha, ..., \top]$), the interpretation $[\![\mathbf{x}, \phi]\!]$ is by definition the pullback of $[R(x_1, ..., x_\alpha, ...)] \rightarrowtail [x_1, ..., x_\alpha, ..., \top]$ along $[x_1, ..., x_\alpha, ..., \top]$, that is, it is precisely the subobject $[R(x_1, ..., x_\alpha, ...)] \rightarrowtail [x_1, ..., x_\alpha, ..., \top]$. If $[\mathbf{x}, \phi]$ is the atomic formula $x = x'$, the sort of the variables x, x' correspond to $[y, \top]$ and hence, by definition, the interpretation $[\![\mathbf{x}, \phi]\!]$ is the equalizer of $[x, x], [x', x'] : [xx', \top] \rightarrowtail [y, \top]$, that is, the subobject $[xx', x = x'] \rightarrowtail [xx', \top]$. Similarly, the rest of the cases of the induction process can be carried out.

Therefore, the assertion $[\![\mathbf{x}, \phi]\!] \leq [\![\mathbf{x}, \psi]\!]$ is equivalent to the fact that the two subobjects $[\mathbf{x}, \phi], [\mathbf{x}, \psi]$ of $[\mathbf{x}, \top]$ satisfy $[\mathbf{x}, \phi] \leq [\mathbf{x}, \psi]$, which, by Lemma 3.1.1 3), is in turn equivalent to the fact that $\phi(\mathbf{x}) \vdash_{\mathbf{x}} \psi(\mathbf{x})$ is provable in $\mathbb{T}$. □

Proposition 3.1.2 says in a way that the model $M_{\mathbb{T}}$ reflects all syntactical relations in the theory $\mathbb{T}$; therefore, the analysis of categorical properties of $M_{\mathbb{T}}$ will reveal facts about provability in $\mathbb{T}$.

We now have:

Proposition 3.1.3. *If $\mathbb{T}$ is a κ-coherent (resp. κ-Heyting) theory, then $\mathcal{C}_{\mathbb{T}}$ is a κ-coherent (resp. κ-Heyting) category.*

Proof. To prove $\mathcal{C}_{\mathbb{T}}$ has κ-limits it suffices to prove it has κ-products and equalizers. As the product of γ-many objects $[\mathbf{x_i}, \phi_i]_{i<\gamma}$ (where the $\mathbf{x_i}$ are assumed to be disjoint) we can take the class $[\bigcup_{i<\gamma} \mathbf{x_i}, \bigwedge_{i<\gamma} \phi_i]$ together with the projections indicated below:

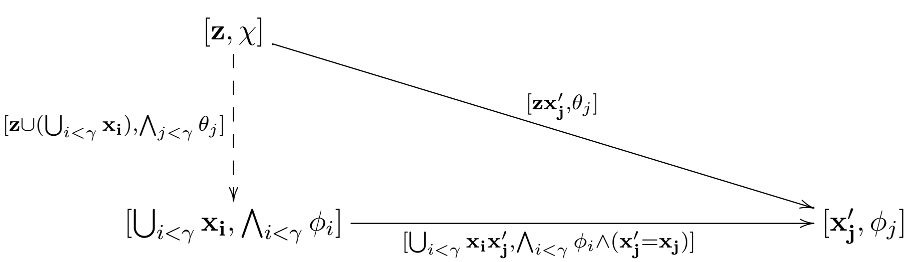

Given morphisms $[\mathbf{z}\mathbf{x'_j}, \theta_j]$, the induced morphism into the product is given by the class $[\mathbf{z} \cup (\bigcup_{i<\gamma} \mathbf{x_i}), \bigwedge_{j<\gamma} \theta_j]$, since it can be easily verified that this is the only morphism that makes the diagram commute.

For the equalizer of a parallel pair of morphisms $[\mathbf{xy}, \theta], [\mathbf{xy}, \delta]$, we take:

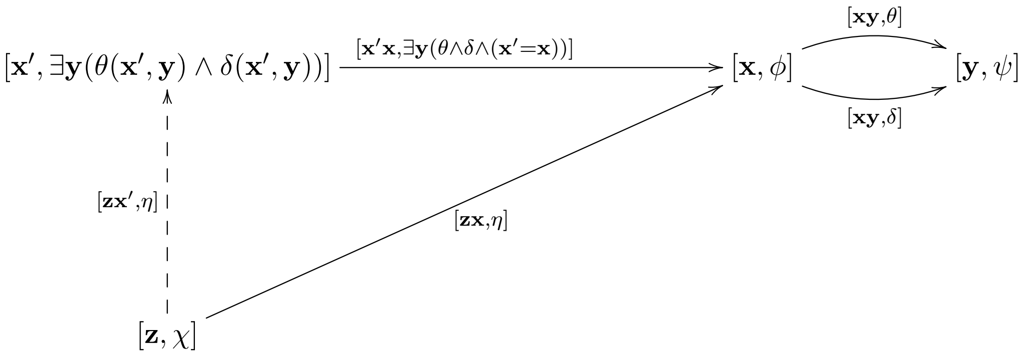

and the universal property is satisfied with the indicated induced morphism. This proves that $\mathcal{C}_T$ has κ-limits. Note as well that there is an initial object given by $[\{\}, \bot]$, and a terminal object given by $[\{\}, \top]$.

To prove that the category has image factorizations, given a morphism $[\mathbf{xy}, \theta] : [\mathbf{x}, \phi] \to [\mathbf{y}, \psi]$ we take its image as the subobject $[\mathbf{y}, \exists\mathbf{x}(\theta)] \rightarrowtail [\mathbf{y}, \psi]$. In particular, $[\mathbf{xy}, \theta]$ is a cover if and only if the sequent $\psi(\mathbf{y}) \vdash_{\mathbf{y}} \exists\mathbf{x}\theta(\mathbf{x},\mathbf{y})$ is provable in $\mathbb{T}$. Then, from the construction of limits above, it can be verified straightforwardly that covers are stable under pullbacks. To prove the category has unions, take subobjects $[\mathbf{x}, \phi_i]_{i<\gamma}$ of $[\mathbf{x}, \phi]$ and define their union to be $[\mathbf{x}, \bigvee_{i<\gamma} \phi_i]$.

The validity of the transfinite transitivity property is proven using the transfinite transitivity rule, the construction of limits and the fact that $([\mathbf{x}_i\mathbf{y}, \theta_i] : [\mathbf{x}_i, \phi_i] \to [\mathbf{y}, \psi])_{i<\gamma}$ are jointly covering if and only if the sequent $\psi(\mathbf{y}) \vdash_{\mathbf{y}} \bigvee_{i<\gamma} \exists\mathbf{x}_i\theta_i(\mathbf{x}_i, \mathbf{y})$ is provable in $\mathbb{T}$. It is necessary to compute the limit of a κ-chain $([\mathbf{x}_{\alpha+1}\mathbf{x}_\alpha, \theta_\alpha] : [\mathbf{x}_{\alpha+1}, \phi_{\alpha+1}] \to [\mathbf{x}_\alpha, \phi_\alpha])_{\alpha<\gamma}$. This can be computed using the construction of limits with products and equalizers; in this case, one can verify that the limit of such a chain reduces to compute the equalizer of the following diagram:

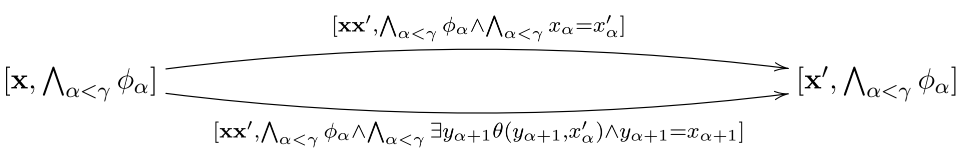

where $\mathbf{x} = x_0, ...x_\alpha, ...$. From this construction and the construction of equalizers in the syntactic category we can derive the sequent that expresses that the transfinite composition of jointly covering families are jointly covering, and verify that the sequent is provable within the theory making use of the transfinite transitivity rule.

Finally, if the theory is κ-Heyting, to construct universal quantification along a morphism $[\mathbf{xy}, \theta] : [\mathbf{x}, \phi] \to [\mathbf{y}, \psi]$, take a subobject $[\mathbf{x}, \eta]$ of its domain, in the canonical form given in Lemma 3.1.1 3). Then define $\forall_{[\mathbf{xy},\theta]}([\mathbf{x}, \eta])$ to be the subobject $[\mathbf{y}, \psi \wedge \forall \mathbf{x}(\theta \to \eta)]$ of $[\mathbf{y}, \psi]$. It follows from Lemma 3.1.1 3) that this works.

This concludes the proof. □

We now get:

Corollary 3.1.4. *If κ is any inaccessible cardinal, κ-coherent (resp. κ-Heyting) theories with respect to models in κ-coherent (resp. κ-Heyting) categories.*

Remark 3.1.5. We will see later than when establishing the completeness with respect to models in **Set**, we need to ask a large cardinal property (weak compactness) to hold for κ.

3.2 Morleyization and exploding models

The internal κ-coherent theory of a, say, κ-Heyting category can alternatively be described by a different axiomatization, which will be simpler for our purposes. Following [Joh02], where the process of rewriting of a classical first-order theory as an equivalent coherent theory is referred to as "Morleyization", we will also call "Morleyizing" a theory, in general, rewriting it into a theory in a less expressive fragment. From a categorical viewpoint (as opposed to the standard syntactic point of view), the syntactic category $\mathcal{C}_{\mathbb{T}}$ of, for example, an intuitionistic κ-first-order theory $\mathbb{T}$ which is a κ-Heyting category, is also a κ-coherent (resp. κ-regular) category, and thus $\mathcal{C}_{\mathbb{T}}$ has an internal κ-coherent theory (resp. internal κ-regular theory), which we refer to as "the theory of κ-coherent (resp. κ-regular) models of $\mathbb{T}$", (its "Morleyization" $\mathbb{T}^m$). As we will see below, the theory and its Morleyization have equivalent syntactic categories:

$$\mathcal{C}_{\mathbb{T}} \simeq \mathcal{C}_{\mathbb{T}^m}$$

Although for classical κ-first-order theories, the κ-coherent Morleyization will have the same models in all Boolean κ-coherent categories, in general when Morleyizing a κ-first-order theory to a κ-coherent one (or a κ-coherent theory to a κ-regular one), this is not the case, but there still some gain in considering the category of models of the Morleyized theory, as Joyal's theorem will show.

Definition 3.2.1. The theory of κ-regular ((i)–(iv) below) and the theory of κ-coherent ((i)–(v)) models $\mathbb{T}^m$ of a κ-first-order (resp. κ-pre-Heyting, κ-coherent) theory $\mathbb{T}$ over a signature Σ is defined as follows: its signature Σ^m extends Σ by adding for each κ-first-order (resp. κ-pre-Heyting, κ-coherent) formula ϕ over Σ with free variables $\mathbf{x}$ the relation symbol $P_\phi(\mathbf{x})$; then $\mathbb{T}^m$ is the theory axiomatized by the following axioms:

(i) $P_\phi \dashv\vdash_{\mathbf{x}} \phi$ for every atomic formula ϕ

(ii) $P_\phi \vdash_{\mathbf{x}} P_\psi$ for every sequent $\phi \vdash_{\mathbf{x}} \psi$ provable in $\mathbb{T}$;

(iii) $P_{\bigwedge_{i<\gamma} \phi_i} \dashv\vdash_{\mathbf{x}} \bigwedge_{i<\gamma} P_{\phi_i}$;

(iv) $P_{\exists \mathbf{y}. \phi} \dashv\vdash_{\mathbf{x}} \exists \mathbf{y}. P_\phi$;

(v) $P_{\bigvee_{i<\gamma} \phi_i} \dashv\vdash_{\mathbf{x}} \bigvee_{i<\gamma} P_{\phi_i}$.

The theory of κ-regular models of a κ-$\mathcal{R}eg_\perp$ theory is defined similarly; alternatively (since we are only discarding $\perp$) we could also treat $\perp$ as a propositional variable and add the axioms

$$\perp \vdash_{\mathbf{x}} \phi$$

for all formulas $[\mathbf{x}, \phi]$ in context. The case for the positive κ-coherent[1] Morleyization of a κ-coherent theory can be treated analogously.

Definition 3.2.2. We will say that a κ-regular (resp. positive κ-coherent) model of a κ-$\mathcal{R}eg_\perp$ (resp. κ-coherent) theory is *possibly exploding*, and make the convention that such a model is *exploding* if it assigns $\perp$ the value true.

Note that since $P_\phi \vdash_{\mathbf{x}} P_\psi$ in $\mathbb{T}^m$ if and only if $\phi \vdash_{\mathbf{x}} \psi$ in $\mathbb{T}$, if κ-regular theories are complete for **Set**-valued models, then κ-$\mathcal{R}eg_\perp$ theories will be complete for modified (i.e., possibly exploding) **Set**-valued models. Incidentally, any model of $\mathbb{T}^m$ that assigns $\perp$ the value true must be inhabited, since $(\perp \vdash \exists x. x = x) \in \mathbb{T}^m$.

In the case $\kappa = \omega$ we can constructively prove the following:

Lemma 3.2.3. *(IZF) The functor $F : \mathcal{C}_{\mathbb{T}} \to \mathcal{C}_{\mathbb{T}^m}$ sending a formula in context $[\mathbf{x}, \phi]$ to $[\mathbf{x}, P_\phi]$ has a pseudoinverse, that is, is part of an equivalence of categories.*

[1]The fragment which results after discarding $\perp$.

Proof. The axioms for the Morleyization allow to find for each formula in context $[\mathbf{x}, \psi]$ over Σ^m a provably equivalent formula of the form P_{ϕ_ψ} and over the same context. Indeed, those axioms are precisely the inductive clauses of a proof by induction on the complexity of the formula that every ψ is provable equivalent to a formula P_{ϕ_ψ} for a fixed formula ϕ_ψ. The base of the induction (the case of atomic ψ) is handled by item (i), which entails that $P_\psi \vdash_{\mathbf{x}} \psi$. If $\psi = \bigwedge_{i<\gamma} \phi_i$ or $\psi = \bigvee_{i<\gamma} \phi_i$, we know by inductive hypothesis that there are formulas ϕ_{ϕ_i} such that each $P_{\phi_{\phi_i}}$ is provable equivalent to ϕ_i, and then the use of clauses (iii)–(v) allow to prove that we can take $\phi_\psi = \bigwedge_{i<\gamma} \phi_{\phi_i}$ or $\phi_\psi = \bigvee_{i<\gamma} \phi_{\phi_i}$, respectively. Similarly, if $\psi = \exists \mathbf{y}. \phi$ we use clause (iv). This procedure provides hence a choice function from formulas over Σ^m to formulas over Σ.

There is a functor $G : \mathcal{C}_{\mathbb{T}^m} \to \mathcal{C}_{\mathbb{T}}$ which assigns to an object $[\mathbf{x}, \psi]$ the formula in context $[\mathbf{x}, \phi_\psi]$, and to a morphism $[\mathbf{xy}, \theta]$ the morphism $[\mathbf{xy}, \phi_\theta]$. This latter is a valid morphism because ϕ_θ is provably equivalent to θ, which is functional. The functor can also be seen to be the desired pseudoinverse. Indeed, we have on one hand $GF = Id_{\mathcal{C}_{\mathbb{T}}}$; on the other hand, we can see that there is a natural isomorphism $FG \sim Id_{\mathcal{C}_{\mathbb{T}^m}}$ which just assigns to $[\mathbf{x}, \psi]$ the object $[\mathbf{x}, \phi_\psi]$. Then the naturality conditions reduce to note that θ is provable equivalent to a formula P_{ϕ_θ} for any θ. □

3.3 Syntactic sites

The syntactic categories for fragments of κ-first-order logic can be equipped with appropriate Grothendieck topologies in such a way that the corresponding sheaf toposes are conservative models of the corresponding theories. Given a κ-regular category, we can define the κ-regular coverage, where the covering families are all singletons f where f is a cover. Similarly, for a κ-coherent category we can define the κ-coherent coverage, where the covering families are given by families of arrows $f_i : A_i \to A$ of cardinality less than κ such that the union of their images is the whole of A (in particular, the initial object 0 is covered by the empty family). We can also find (see [BJ98]) a conservative sheaf model given by Yoneda embedding into the sheaf topos obtained with the κ-coherent coverage. As proven in [BJ98], the embedding preserves κ-unions and κ-intersections, as well as any Heyting structure that might exist in $\mathcal{C}$. To highlight the fact that images and unions are stable under pullback is crucial, we prove the following lemma, which can be regarded as a generalization of the result corresponding to the finitary case:

Lemma 3.3.1. *Given a κ-coherent (resp. κ-Heyting) category $\mathcal{C}$ with the κ-coherent coverage τ, Yoneda embedding $\mathcal{C} \to \mathcal{S}h(\mathcal{C}, \tau)$ is a conservative κ-coherent (resp. κ-Heyting) functor and $\mathcal{S}h(\mathcal{C}, \tau)$ is a κ-Heyting category.*

Proof. By [MR77], Proposition 3.3.3, we know that all representable functors are sheaves for the κ-coherent coverage, since the fact that the union of the images of the arrows in a covering family over A is the whole of A is equivalent to the fact that the family is effective epimorphic, and this is precisely the sheaf condition on representable functors. Also, in case $\mathcal{C}$ is κ-Heyting, the embedding preserves universal quantification as shown in [BJ98], Lemma 3.1. Yoneda embedding preserves limits, and limits of sheaves are computed as in presheaves, so it remains to prove that it preserves images and κ-unions. Given a cover $f : A \to B$, we need to prove that $[-, A] \to [-, B]$ is a sheaf epimorphism, i.e., that it is locally surjective. For this it is enough to find a covering family over each object C that witnesses the local surjectivity. Given an element g in $[C, B]$, we can simply form the pullback of f along g, obtaining thus a covering family over consisting on the single arrow $g^*(f)$ which will clearly witness the local surjectivity.

The argument for the preservation of unions is similar: given the union $\bigvee_{i<\gamma} A_i$ of subobjects $f_i : A_i \to B$ we need to show that $[-, \bigvee_{i<\gamma} A_i]$ is the union of the sheaves $[-, A_i]$. Given an object C and an element g in $[C, \bigvee_{i<\gamma} A_i]$, the pullbacks along g of $f'_i : A_i \to \bigvee_{i<\gamma} A_i$ give a covering family $\{g^*(f'_i) : P_i \to C\}_{i<\gamma}$ with the property that $g.g^*(f'_i) \in [P_i, \bigvee_{i<\gamma} A_i]$ belongs to $[P_i, A_i]$, which is enough to guarantee that $[-, \bigvee_{i<\gamma} A_i]$ is indeed the union of the $[-, A_i]$.

Finally, we show that the sheaf topos is a κ-coherent category by proving that the transfinite transitivity property holds in $\mathcal{S}h(\mathcal{C}, \tau)$. To this end, suppose we have a family of sheaves $\{S_f : \beta < \gamma, f \in \gamma^\beta\}$ satisfying the premises of the transfinite transitivity property, that is, that $\{S_g \to S_f : g \in \gamma^{\beta+1}, g|_\beta = f\}$ form a jointly covering family and that $S_{f|_\beta} = \lim_{\alpha<\beta} S_{f|_\alpha}$ for limit β. Then given $c \in S_\emptyset(C)$ we define by transfinite recursion a covering family $\{C_f \to C : \beta < \gamma, f \in \gamma^\beta\}$ such that, given $f \in \gamma^\gamma$, $c.f \in \bigwedge_{\alpha<\gamma} S_{f_i|_\alpha}(C)$ for some $f_i \in \gamma^\gamma$, witnessing that $\{\bigwedge_{\alpha<\gamma} S_{f|_\alpha} \to S_\emptyset : f \in \gamma^\gamma\}$ is a jointly covering family. In fact, the covering family over C will be such that for any fixed $\beta < \gamma$ we will have that $\{C_f \to C : f \in \gamma^\beta\}$ is a witness of the joint covering of the sheaves $\{S_f : f \in \gamma^\beta\}$, that is, given $f \in \gamma^\beta$ we will have $c.f \in S_{f_i}(C)$ for some $f_i \in \gamma^\beta$.

Supposing that $\{C_f \to C : \beta < \mu, f \in \gamma^\beta\}$ has been defined, we show

how to define the family at level μ. If μ is a successor ordinal $\mu = \alpha + 1$, we have by inductive hypothesis a covering $\{C_f \to C : f \in \gamma^\alpha\}$ such that, given $f \in \gamma^\alpha$, $c.f \in S_{f_i}(C)$ for some $f_i \in \gamma^\alpha$. Then, because $\{S_g \to S_{f_i} : g \in \gamma^\mu, g|_\alpha = f_i\}$ is jointly covering, we can find a covering $\{h_{gf_i} : C_g \to C_{f_i} : g \in \gamma^\mu, g|_\alpha = f_i\}$ such that, given $g \in \gamma^\mu, g|_\alpha = f_i$, $c.g = (c.f_i).h_{gf_i} \in S_{g_j}(C)$ for some $g_j \in \gamma^\mu$. This extends, by transitivity, the definition of the covering family to level μ. If μ is a limit ordinal and $f \in \gamma^\mu$, we simply take C_f to be the limit of the diagram formed by $C_{f|_\alpha} : \alpha < \mu$. Then clearly, given $f \in \gamma^\mu$, $c.f \in \bigwedge_{\alpha<\mu} S_{f_k|_\alpha}$ for some $f_k \in \gamma^\mu$. This finishes the recursive construction of the family over C and proves the transfinite transitivity property for the sheaves.

□

We get immediately:

Corollary 3.3.2. *If κ is any inaccessible cardinal, κ-first-order theories are semantically complete for models in κ-Heyting Grothendieck toposes.*

3.4 Beth and Kripke models

Let Σ be a first-order signature. Assume without loss of generality that Σ is relational, that is, without function or constant symbols. From now on we will implicitely assume that our theories are single-sorted, for the sake of simplicity. For our purposes, the following definition of modified Beth model (i.e., one admitting nodes forcing $\bot$) will suffice:

Definition 3.4.1. A Beth model for pure κ-first-order logic over Σ is a quadruple $\mathcal{B} = (K, \leq, D, \Vdash)$, where $(K, \leq)$ is a tree with a set B of branches (i.e., maximal chains in the partial order), D is a set-valued functor on K and the forcing relation $\Vdash$ is a binary relation between elements of K and sentences of the language with constants from $\bigcup_{k \in K} D(k)$, defined for atomic formulas ϕ with the following condition: if B_k denotes the subset of branches containing the node k, then $k \Vdash \phi(\mathbf{d}) \iff \forall b \in B_k \exists l \in b, l \geq k (l \Vdash \phi(D_{kl}(\mathbf{d})))$ and $\mathbf{d} \subseteq D(k)$. This definition is recursively extended to arbitrary formulas as follows:

1. $k \Vdash \bigwedge_{i<\gamma} \phi_i(\mathbf{d}) \iff k \Vdash \phi_i(\mathbf{d})$ for every $i < \gamma$
2. $k \Vdash \bigvee_{i<\gamma} \phi_i(\mathbf{d}) \iff \forall b \in B_k \exists l \in b, l \geq k \quad (l \Vdash \phi_i(D_{kl}(\mathbf{d})$ for some $i < \gamma)$
3. $k \Vdash \phi(\mathbf{d}) \to \psi(\mathbf{d}') \iff \forall k' \geq k(k' \Vdash \phi(D_{kk'}(\mathbf{d})) \implies k' \Vdash \psi(D_{kk'}(\mathbf{d}'))$

4. $k \Vdash \exists \mathbf{x}\phi(\mathbf{x}, \mathbf{d}) \iff \forall b \in B_k \exists l \in b, l \geq k \quad \exists \mathbf{e} \subseteq D(l)(l \Vdash \phi(\mathbf{e}, D_{kl}(\mathbf{d}))$

5. $k \Vdash \forall \mathbf{x}\phi(\mathbf{x}, \mathbf{d}) \iff \forall k' \geq k \forall \mathbf{e} \subseteq D_{k'}(k' \Vdash \phi(\mathbf{e}, D_{kk'}(\mathbf{d})))$

A Beth model for a theory $\mathbb{T}$ is a Beth model forcing all the axioms of the theory. If the clauses for atomic formulas, disjunction and existential quantification are strengthened by requiring that $level(l) = level(k) + \alpha$ for a fixed $\alpha < \kappa$, the Beth model will be called *weak*.

A Kripke model is a special kind of Beth model none of whose nodes forces $\perp$ and where the forcing relation for atomic formulas, disjunction and existential quantification satisfies the stronger condition $level(l) = level(k)$:

Definition 3.4.2. A Kripke model for pure first-order logic over Σ is a quadruple $\mathcal{B} = (K, \leq, D, \Vdash)$, where $(K, \leq)$ is a tree, D is a set-valued functor on K and the forcing relation $\Vdash$ is a binary relation between elements of K and sentences of the language with constants from $\bigcup_{k \in K} D(k)$, defined for atomic formulas ϕ with the conditions that $k \nVdash \perp$ and that $k \Vdash \phi(\mathbf{d}) \implies l \Vdash \phi(D_{kl}(\mathbf{d})))$ for $\mathbf{d} \subseteq D(k)$, and recursively extended to arbitrary formulas as follows:

1. $k \Vdash \bigwedge_{i<\gamma} \phi_i(\mathbf{d}) \iff k \Vdash \phi_i(\mathbf{d})$ for every $i < \gamma$

2. $k \Vdash \bigvee_{i<\gamma} \phi_i(\mathbf{d}) \iff k \Vdash \phi_i(\mathbf{d})$ for some $i < \gamma$)

3. $k \Vdash \phi(\mathbf{d}) \to \psi(\mathbf{d}') \iff \forall k' \geq k(k' \Vdash \phi(D_{kk'}(\mathbf{d})) \implies k' \Vdash \psi(D_{kk'}(\mathbf{d}')))$

4. $k \Vdash \exists \mathbf{x}\phi(\mathbf{x}, \mathbf{d}) \iff \exists \mathbf{e} \subseteq D(k)(k \Vdash \phi(\mathbf{e}, \mathbf{d})$

5. $k \Vdash \forall \mathbf{x}\phi(\mathbf{x}, \mathbf{d}) \iff \forall k' \geq k \forall \mathbf{e} \subseteq D_{k'}(k' \Vdash \phi(\mathbf{e}, D_{kk'}(\mathbf{d})))$

A Kripke model for a theory $\mathbb{T}$ is a Kripke model forcing all the axioms of the theory.

A Kripke model can also be seen categorically as a model on a presheaf category. That is, if $\mathcal{C}_{\mathbb{T}}$ is the syntactic category of the theory, a Kripke model on $(K, \leq)$ is nothing but a κ-Heyting functor $F : \mathcal{C}_{\mathbb{T}} \to \mathbf{Set}^K$, since such a functor determines the set-valued functor $D = F([x, \top]) : K \to \mathbf{Set}$ which specifies the underlying domains of the nodes. In this case the forcing relation is given by $k \Vdash \phi(\mathbf{d})$ for $\mathbf{d} \in D(k)^n$ if and only if $\mathbf{d} : [-, k] \to D^n = F([\mathbf{x}, \top])$ factors through the subobject $F([\mathbf{x}, \phi]) \rightarrowtail F([\mathbf{x}, \top])$, where we use Yoneda lemma to identify elements of $D(k)$ with natural transformations $[-, k] \to D$. This definition is precisely the forcing relation for the Kripke-Joyal semantics in the topos $\mathbf{Set}^K$, whence the name Kripke associated to it.

More generally, one can consider Kripke models on arbitrary categories $\mathcal{M}$ instead of the tree K, and it turns out that the semantics of the Kripke model over $\mathcal{M}$ can be recovered in terms of Kripke semantics over a certain collection of trees. To do that, consider first the poset P which consists of finite composable sequences of morphisms of $\mathcal{M}$, i.e., chains $A_0 \to \ldots \to A_n$ in $\mathcal{M}$. One such sequence is below another in P if the former is an initial segment of the latter. There is a functor $E : P \to \mathcal{M}$ sending each chain to the last object in it and sending any morphism f of P to the composite of the morphisms of $\mathcal{M}$ that are in the codomain minus the domain of f. Now, given a Kripke model $F : \mathcal{C}_{\mathbb{T}} \to \mathbf{Set}^{\mathcal{M}}$, we can compose F with the transpose $E^* : \mathbf{Set}^{\mathcal{M}} \to \mathbf{Set}^P$, and if this latter is a conservative κ-Heyting functor, this will provide a Kripke model on P forcing precisely the same formulas as the original model. Finally, the Kripke model on P can be regarded as a collection of Kripke models on trees, where the roots of the trees are given by one-element chains. This construction amounts to build the Diaconescu cover of the topos $\mathbf{Set}^{\mathcal{M}}$ (see e.g. [MM94]). In our case the discussion above shows that for our purposes it is enough to prove the following, which is a particular case of section 1.744 of [FS90]:

Lemma 3.4.3. *The functor* $E^* : \mathbf{Set}^{\mathcal{M}} \to \mathbf{Set}^P$ *is conservative and* κ*-Heyting.*

Proof. The conservativity of E^* follows from the fact that E is surjective on objects and arrows. To prove that it is κ-Heyting, the non-trivial part is proving that it preserves $\forall$. For a natural transformation $f : F \to G$ in $\mathcal{S}et^{\mathcal{M}}$ and a subfunctor A of F, we need to show that $E^*(\forall_f A)$ is the same subfunctor of $E^*(G)$ as $\forall_{E^*(f)} E^*(A)$. By definition, for any object p in P and $y \in E^*(G)(p) = G(E(p))$, we have $y \in \forall_{E^*(f)} E^*(A)(p)$ if and only if for all arrows $l : p \to q$ in P one has:

$$E^*(f)_q^{-1}(G(E(l))(y)) \subseteq E^*(A)(q)$$

$$\iff \forall x(E^*(f)_q(x) = G(E(l))(y) \implies x \in E^*(A)(q))$$

$$\iff \forall x(f_{E(q)}(x) = G(E(l))(y) \implies x \in A(E(q))) \qquad (1)$$

On the other hand, also by definition, for $y \in G(E(p))$ one has $y \in \forall_f A(E(p))$ if and only if for all arrows $t : E(p) \to r$ in $\mathcal{M}$ one has:

$$f_r^{-1}(G(t)(y)) \subseteq A(r)$$

$$\iff \forall x(f_r(x) = G(t)(y) \implies x \in A(r)) \qquad (2)$$

But because the functor E is surjective (both on objects and arrows), we can find $q, l \in P$ such that $r = E(q)$ and $t = E(l)$, from which we deduce that (1) and (2) above are equivalent. Hence, $E^*\forall = \forall E^*$, as we wanted.

□

Part III

Completeness in constructive settings

4. Constructive completeness

It has been known since 1957 that completeness theorems for intuitionistic or classical theories imply non-constructive principles (see this and related ideas in [Kre62], [McC94], [McC08]). This motivated the question of what the correct notion of constructive model is. In a classical metatheory, one works with Tarskian models for classical logic and Kripke and Beth models for intuitionistic logic. If we wish to work in a strictly constructive metatheory, however, these semantics cannot be proven to be complete, and one has to appeal to more general semantics like locale and topoi models, which seem more suitable. In 1976, Veldman proved in [Vel76] that if one restricts to a countable signature, it is possible to prove the completeness of intuitionistic logic with respect to a modified notion of Kripke semantics which is classically equivalent to the usual one, while intuitionistically lies very close to it. As he explained, to make the completeness theorem constructive, one needs to redefine the notion of Kripke models and allow for the so called exploding nodes, in which $\bot$ (and hence every formula) is forced. In that way one sacrifices the decidability of a node being exploding, but the resulting completeness theorem is intuitionistically provable.

More generally, some intuitionistic completeness theorems for both classical and intuitionistic logic have been proposed later. For the classical case, one has, e.g., Krivine's results in [Kri96], which works with a notion of Tarskian model which is allowed to be exploding in the aforementioned sense. For the intuitionistic case, Friedman has shown that the introduction of exploding nodes in the Beth models (related to Kripke semantics) gives as well a fully constructive proof of completeness, with a metatheory weaker than that of Veldman's proof. One of the main purposes of this chapter is to present a unifying categorical approach that proves constructive completeness theorems for classical and intuitionistic first-order logic, as well as for the coherent fragment (see [Joh02]). We will see that Veldman's modified semantics is actually a particular case of sheaf semantics, which confirms the adequacy of sheaf models in constructive mathematics. On the other hand, our proof of Kripke completeness, as that of Veldman's, makes use of the FAN theorem, which raises the question of how far one can get in proving completeness of Kripke semantics constructively and without its use. We present here a fragment of first-order logic for which the appeal to the FAN theorem is eliminated. These and related ideas appear in [FE]. Finally, we show that the size of the signature is a fundamental constraint in Veldman's proof.

When we work in a constructive metatheory such as intuitionistic Zermelo Fraenkel set theory, IZF (an account of this can be found in [Myh73]), sheaf semantics is a candidate for the "correct" constructive notion of model. For example, it is shown in [Pal97] that there is a conservative model of any first-order theory in the category of coherent sheaves on its syntactic category. In a category-theoretic formulation, if $\mathcal{C}$ is a small Heyting category, then the Yoneda embedding

$$y : \mathcal{C} \to \mathrm{Sh}(\mathcal{C}, K)\,,$$

where K is the coherent coverage (finite covering families), is a conservative Heyting functor. As we shall see, this sheaf topos is essentially the same as a particular modified Beth model which encodes all the information of the topos into a tree structure. But the sheaf topos can also be made more reminiscent of Kripke semantics if one considers instead an embedding into a presheaf category, which is essentially the same as working with Kripke models. However, the passing from sheaves to presheaves is done at the expenses that the proof becomes non-constructive.

As explained in [MR77], Joyal's completeness theorem provides this embedding in the form of an evaluation functor:

$$ev : \mathcal{C} \to \mathbf{Set}^{\mathrm{Mod}_c(\mathcal{C})^{\mathrm{op}}},$$

where $\mathrm{Mod}_c(\mathcal{C})$ is the category of coherent models of $\mathcal{C}$, and this new functor ev factors through y. As we shall prove, when the signature is countable and the theory is semi-decidable, one can make the completeness proof constructive by replacing the presheaf category with a sheaf model $\mathrm{Sh}\big(\mathrm{Mod}_c^E(\mathcal{C}), E\big)$, where the underlying category of the site differs from $\mathrm{Mod}_c(\mathcal{C})$ just in that it also includes exploding models, and where E is an appropriate "exploding" topology. Classically, $\mathrm{Sh}\big(\mathrm{Mod}_c^E(\mathcal{C}), E\big)$ and $\mathbf{Set}^{\mathrm{Mod}_c(\mathcal{C})^{\mathrm{op}}}$ will be equivalent categories. Constructively, though, the topos $\mathrm{Sh}\big(\mathrm{Mod}_c^E(\mathcal{C}), E\big)$ will be essentially the same as Veldman's universal modified Kripke model (see [Vel76]).

We cannot relax the hypothesis on the size of the signature while remaining constructive. Indeed, we have:

Lemma 4.0.4. *The completeness theorem for modified Kripke semantics for full first-order logic is unprovable in IZF.*

Proof. Since classically modified Kripke semantics is just the usual Kripke semantics, the completeness of the former (MKC) is classically equivalent to the completeness of the latter (KC). On the other hand, adding to IZF the law of excluded middle (LEM) one gets ZF. Hence, if MKC was derivable in IZF, we would have ZF=IZF+LEM=IZF+MKC+LEM=ZF+KC, that is absurd since ZF+KC implies BPI (see [Esp16]), which is known to be independent of ZF (see, e.g., [Jec73]). □

4.1 Joyal's theorem

The construction of the syntactic category is an aspect of the philosophy of theories as categories, which is supplemented by the concept of internal theory of a given category and the functorial semantics associated with it. For, say, a coherent category $\mathcal{C}$ there is a canonical signature and coherent axioms associated to the category in such a way that coherent models of this theory correspond to coherent functors having the category as a domain. That is, functors which preserve the categorical properties are seen as models of the internal theory of the categories in the codomain categories. Moreover, model homomorphisms correspond in this view to natural transformations of functors. This allows us to think, for example, of the category $\mathcal{M}$ of set-valued coherent models of a theory as corresponding functors from the syntactic category of the theory to the category **Set** of sets. Consider now the further functor category $\mathbf{Set}^{\mathcal{M}}$. To each coherent formula in context we can assign its extension in each of the models of $\mathcal{M}$, or equivalently, evaluate the models, seen as functors, on the corresponding object represented by the formula. This assignment is in fact functorial, and thus each coherent formula in context gives rise to a functor in $\mathbf{Set}^{\mathcal{M}}$, which we call the evaluation functor at the corresponding formula. If we do this for every coherent formula in context, the assignment of evaluation functors at formulas is itself functorial, and gives rise to a functor $ev: \mathcal{C}_{\mathbb{T}} \to \mathbf{Set}^{\mathcal{M}}$.

In its original version, Joyal's theorem is a statement over ZFC which could be described as follows:

Theorem 4.1.1. *(Joyal) Let $\mathbb{T}$ be a coherent theory and let $\mathcal{M}$ be the category of coherent models of $\mathbb{T}$. Then the functor*

$$ev: \mathcal{C}_{\mathbb{T}} \to \mathbf{Set}^{\mathcal{M}}$$

is conservative and Heyting.

For the proof of Joyal's theorem we refer to [MR77], Ch. 6, pp. 189, since we will later study and prove some variations. The significance of the theorem resides in that it encapsulates three different completeness theorems. The conservativity of ev is a categorical way of saying that models in $\mathcal{M}$ are semantically complete for coherent logic. In the particular case in which the logic is classical, this is precisely Gödel's completeness theorem for first order logic. But even when we consider intuitionistic logic, the preservation of the right adjoint entails that ev preserves the first-order structure of $\mathcal{C}_{\mathbb{T}}$, and through categorical semantics in the presheaf category $\mathbf{Set}^{\mathcal{M}}$ we can see that the conservative embedding provides a universal Kripke model of the theory, resulting thus in Kripke completeness theorem for first-order intuitionistic logic. We shall go some steps further and consider variations that provide new completeness theorems, both in the constructive and in the infinitary case. Our aim here is to adapt Joyal's theorem to different circumstances, according to the type of logic and metatheory we have in each case.

For the constructive part, we will employ the methods of categorical logic understood as being formalized in IZF. This means that all of our constructions have to avoid any use of the principle of excluded middle or the axiom of choice. The developments on constructive sheaf semantics, like the ones in [Pal97] or [TVD88] for instance, will be used here. There are various notions of countability which are discussed in [AR01]; for the sake of clarity, we give the explicit definitions of some of the concepts we will make use of:

Definition 4.1.2. A subset X of a set A is called:

1. discrete if for any $x, y \in X$ it is the case $x = y \vee x \neq y$;
2. decidable if for any element $x \in A$ it is the case $x \in X \vee x \notin X$;
3. enumerable if there exists a surjection from $\mathbf{N}$ onto $X + 1$;
4. semi-decidable if there is a function $f : \mathbf{N} \times X \to 2$ such that $x \in X \leftrightarrow \exists n \in \mathbf{N}(f(n, x) = 0)$.

We will assume that the signature is a discrete, countable set, so that we can have a robust notion of formula, and that our theories are single sorted. Let Σ be a first-order signature, and assume without loss of generality that Σ is relational, that is, without function or constant symbols. We consider here enumerable theories over Σ. Note that theories whose axioms are semi-decidable subsets of the countable set of formulas are also enumerable.

The constructive version of Joyal's theorem, as presented in [FE] and to be proved in the next section, is as follows:

Theorem 4.1.3. *Let $\mathbb{T}$ be a semi-decidable $\mathcal{R}eg_\perp$ (resp. coherent) theory and let $\mathcal{M}$ be a full subcategory of the category of regular (resp. coherent) models of $\mathbb{T}$ such that:*

i) $\mathbb{T}$ is complete with respect to $\mathcal{M}$, and

ii) for every $\mathbf{M} \in \mathcal{M}$ the theory $Th(\mathbf{M})$ is complete with respect to models (the reducts of which are) in $\mathcal{M}$.

Then the functor

$$ev : \mathcal{C}_\mathbb{T} \to \mathbf{Set}^\mathcal{M}$$

is a) conservative and b) whenever the pullback functor $f^* : \mathrm{Sub}_{\mathcal{C}_\mathbb{T}}(B) \to \mathrm{Sub}_{\mathcal{C}_\mathbb{T}}(A)$ induced by a morphism $f : A \to B$ in $\mathcal{C}_\mathbb{T}$ has a right adjoint $\forall_f$, we have for all $S \in \mathrm{Sub}_{\mathcal{C}_\mathbb{T}}(A)$ that $ev(\forall_f(S)) = \forall_{ev(f)}(ev(S))$.

This version will be possible through Morleyization, because that process does not affect the enumerability of theories. We state this in the following:

Lemma 4.1.4. *If $\mathbb{T}$ is enumerable, so is $\mathbb{T}^m$.*

Proof. Straightforward from the definition of $\mathbb{T}^m$. □

4.2 Friedman's theorem

There exists a constructive completeness proof for pure intuitionistic logic with respect to weak Beth models due to Friedman and presented, e.g. in [TVD88]; we will briefly describe the construction.

One starts with a language containing at least one predicate symbol and adjoins a countable set C of new constants to the language $\mathcal{L}$. Then one takes an enumeration $\{A_n : n \in \mathbf{N}\}$ with infinite repetitions of the formulas in this new language $\mathcal{L}'$; finally, one defines $\mathcal{L}(\Gamma)$ to be the language containing constants from C that appear in the set of formulas Γ. The weak Beth model has an underlying binary tree given by finite $0-1$ sequences k ordered by prolongation, and to each node k one assigns a finite set Γ_k of formulas of $\mathcal{L}'$. Denote $\Gamma \vdash_m \phi$ to mean that there exists a proof of ϕ from Γ with code number less than m. The assignment is defined inductively as follows: let $length(k) = u$, $\Gamma_{\{\}} = \emptyset$ and define Γ_{k*i} by cases:

1. if $A_u \notin \mathcal{L}(\Gamma_k)$, take $\Gamma_{k*0} = \Gamma_{k*1} = \Gamma_k$,

2. if $A_u \in \mathcal{L}(\Gamma_k)$, $A_u = B \vee C$ and $\Gamma_k \vdash_u B \vee C$, take $\Gamma_{k*0} = \Gamma_k \cup \{B\}$ and $\Gamma_{k*1} = \Gamma_k \cup \{C\}$;

3. if $A_u \in \mathcal{L}(\Gamma_k)$, $A_u = \exists x B(x)$, $\Gamma_k \vdash_u \exists x B(x)$, let c_i be the first constant of C not in $\Gamma_k \cup \{B(x)\}$ and take $\Gamma_{k*0} = \Gamma_{k*1} = \Gamma_k \cup \{B(c_i)\}$;

4. in any other case, take $\Gamma_{k*0} = \Gamma_{k*1} = \Gamma_k \cup \{A_u\}$.

One completes the definition of the Beth model by assigning the constant domain C to the nodes, and setting $k \Vdash P$ if and only if $\vdash \bigwedge \Gamma_k \to P$ for all atomic sentences P. By an inductive proof one may then verify that this equivalence holds in fact for every sentence P, from which completeness follows if one considers the root node. Note that the given forcing relation is semi-decidable, since one has an enumeration of all proofs.

We shall now restrict our attention to Beth models the underlying domain of whose nodes are enumerable, and where the forcing relation is semi-decidable. Likewise, while relaxing the hypothesis of constant domains we will allow for arbitrary functions connecting the underlying set of a node with that of its successors.

In [Pal97], Palmgren developed a constructive proof that the classifying topos for the coherent fragment of a first-order theory is a universal model of the theory. The topos is constructed by taking sheaves on the syntactic category of the theory with respect to the coherent coverage, and the semantics considered is the usual Kripke-Joyal forcing on the objects. Because the category has a terminal object and we work with first-order logic, it is enough to consider the sheaf forcing at such object. We prove next that this forcing is equivalent to the forcing on a weak Beth model constructed on a binary tree.

Theorem 4.2.1. *Any enumerable first-order theory has a universal*[1] *weak Beth model on a binary tree.*

Proof. Consider the syntactic category $\mathcal{C}_{\mathbb{T}}$ of the theory and its conservative embedding in the topos of sheaves with the coherent coverage, $\mathcal{C}_{\mathbb{T}} \to \mathcal{S}h(\mathcal{C}_{\mathbb{T}}, \tau_C)$. Consider an enumeration of all possible proofs; it is easy to see that the set of basic covering families of cardinality 2 over

[1]Throughout this work, by an universal model we mean a conservative one, that is, one where the sentences that are valid/forced are precisely the provable ones.

any object, which are given by jointly covering pairs of arrows, is then a semi-decidable subset of the enumerable set of all pairs of arrows (indeed, if the object is $[\mathbf{x}, \phi]$, it consists of those pairs $\theta(\mathbf{y}, \mathbf{x})$, $\eta(\mathbf{z}, \mathbf{x})$ of provably functional formulas such that $\phi \vdash_{\mathbf{x}} \exists \mathbf{y} \theta \vee \exists \mathbf{z} \eta)$. Hence, the subset is also enumerable. Construct a functor from the binary tree to the syntactic category, defined recursively on the levels of the tree. Start with an denumeration $n : \mathbf{N} \times \mathbf{N} \to \mathbf{N}$ of $\mathbf{N} \times \mathbf{N}$ with the property that $n(k, l) \geq l$ (for example, the usual canonical denumeration $(0,0), (1,0), (0,1), (1,1), (0,2), ...$ will do). We describe by an inductive definition how the tree obtained as the image of the functor is constructed.

The root of the tree is the terminal object. Suppose now that the tree is defined up to level $n = n(k, l)$. In particular, the nodes $\{p_i\}_{i=1}^{m_l}$ at level l are defined because of the hypothesis $n(k, l) \geq l$. Consider the morphisms g_{ij}^n over p_i assigned to the paths from each of the nodes p_i to the nodes of level n. To define the nodes at level $n + 1$, take then the $k-th$ covering pair over each p_i and pull it back along the morphisms g_{ij}^n. This produces covering pairs over each node at level n, whose domains are then the nodes of level $n + 1$.

Clearly, the morphisms assigned to the paths from any node p till the nodes of level m in the subtree over p form a basic covering family of p. Define now a Beth model B over this binary tree by defining as the underlying set of a node q the set of arrows from q to the object $[x, \top]$ in the syntactic category, and where the function between the underlying set of a node and its successor is given by composition with the corresponding arrow. We set by definition $q \Vdash_B R(\alpha)$ if and only if q forces $R(\alpha)$ in the sheaf semantics of the topos (we identify the category with its image through Yoneda embedding). That p forces $\phi(\alpha)$ in the sheaf semantics of the topos will be denoted by $p \Vdash \phi(\alpha)$. We shall now prove the following:

Claim : For every node p and every tuple α, $p \Vdash \phi(\alpha)$ if and only if $p \Vdash_B \phi(\alpha)$.

The proof goes by induction on ϕ.

1. If ϕ is atomic, the result is immediate by definition of the underlying structures on each node.

2. If $\phi = \psi \wedge \theta$, the result follows easily from the inductive hypothesis, since we have $p \Vdash \psi(\alpha) \wedge \theta(\alpha)$ if and only if $p \Vdash \psi(\alpha)$ and $p \Vdash \theta(\alpha)$,

if and only if $p \Vdash_B \psi(\alpha)$ and $p \Vdash_B \theta(\alpha)$, if and only if $p \Vdash_B \psi(\alpha) \wedge \theta(\alpha)$.

3. Suppose $\phi = \psi \vee \theta$. If $p \Vdash \psi(\alpha) \vee \theta(\alpha)$, then there is a basic covering family $\{f_i : A_i \to p\}_{i=1}^n$ such that for each i, $A_i \Vdash \psi(\alpha f_i)$ or $A_i \Vdash \theta(\alpha f_i)$. By taking the joint image of the arrows in each of the two groups of arrows, we get a covering pair $\{f, g\}$ with the same property. Since this covering pair appears at some point in the enumeration, it is pulled back along all paths g_j of a finite subtree to create the nodes of a certain level of the subtree over p. Hence, every node m_k in such a level satisfies $m_k \Vdash \psi(\alpha f g'_j)$ or $m_k \Vdash \theta(\alpha g g'_j)$. By inductive hypothesis, $m_k \Vdash_B \psi(\alpha f g'_j)$ or $m_k \Vdash_B \theta(\alpha g g'_j)$, and hence we have $p \Vdash_B \psi(\alpha) \vee \theta(\alpha)$.

   Conversely, if $p \Vdash_B \psi(\alpha) \vee \theta(\alpha)$, there is a level in the subtree over p such that for every node m_k there one has $m_k \Vdash_B \psi(\alpha f_k)$ or $m_k \Vdash_B \theta(\alpha f_k)$, so by inductive hypothesis $m_k \Vdash \psi(\alpha f_k)$ or $m_k \Vdash \theta(\alpha f_k)$. Since $\{f_k : m_k \to p\}$ is, by construction, a basic covering family, we must have $p \Vdash \psi(\alpha) \vee \theta(\alpha)$.

4. Suppose $\phi = \psi \to \theta$. If $p \Vdash \psi(\alpha) \to \theta(\alpha)$, for every $f : c \to p$ in the category one has $c \Vdash \psi(\alpha f) \implies c \Vdash \theta(\alpha f)$. In particular, this holds when c is any node q in the tree above p, and by inductive hypothesis one has $q \Vdash_B \psi(\alpha f) \implies q \Vdash_B \theta(\alpha f)$ for all such nodes. Therefore, $p \Vdash_B \psi(\alpha) \to \theta(\alpha)$.

   Conversely, suppose that $p \Vdash_B \psi(\alpha) \to \theta(\alpha)$ and consider an arrow $f : c \to p$. Together with the identity, this arrow forms a covering pair which appears at some point in the enumeration and is hence pulled back along paths g_j of a finite subtree to build the next level of the subtree over p. Suppose that $c \Vdash \psi(\alpha)$; then $g_j^*(c) \Vdash \psi(\alpha g'_j)$, so by inductive hypothesis one has $g_j^*(c) \Vdash_B \psi(\alpha g'_j)$. Therefore, we get $g_j^*(c) \Vdash_B \theta(\alpha g'_j)$, and using once more the inductive hypothesis, $g_j^*(c) \Vdash \theta(\alpha g'_j)$. But $g'_j = f^*(g_j) : g_j^*(c) \to c$ is a basic cover of c (since the g_j form a basic cover of p), and hence we will have $c \Vdash \theta(\alpha)$. We have, thus, proved that $p \Vdash \psi(\alpha) \to \theta(\alpha)$.

5. Suppose $\phi = \exists x \psi(x)$. If $p \Vdash \exists x \psi(x, \alpha)$, then there is a basic covering family $\{f_i : A_i \to p\}_{i=1}^n$ such that for each i one has $A_i \Vdash \psi(\beta_i, \alpha f_i)$ for some $\beta_i : A_i \to [x, \top]$ (even if p is covered by the empty family, it is isomorphic to the initial object and then we can take the identity to be f_1). This basic cover can be decomposed in a covering pair $\{f_1, f'_1\}$ taking the first arrow together with the subobject $s \rightarrowtail p$ given by the union of the images of the rest of

the arrows. This subobject has hence a covering with a similar property to that of p, but with one arrow less, which allows us to give a recursive argument as follows. The binary cover of p appears at some point in the enumeration and is hence pulled back along all paths g_j of a finite subtree to create the nodes of a certain level of the subtree over p. Half of the nodes m_k in this level (namely, the ones corresponding to an odd k) will have the property that $m_k \Vdash \psi(\beta_i g'_j, \alpha f_1 g'_j)$, and hence, by inductive hypothesis, that $m_k \Vdash_B \psi(\beta_i g'_j, \alpha f_1 g'_j)$. The covering over the subobject s can now be pulled back to create coverings for each m_k with even k, and hence we can also prove for them (in view of the recursion) that $m_k \Vdash_B \psi(\beta_i g'_j, \alpha f'_1 g'_j)$. By definition, we get thus $p \Vdash_B \exists x\psi(x, \alpha)$.

Conversely, suppose that $p \Vdash_B \exists x\psi(x, \alpha)$. Then there is a level in the subtree over p such that for every node m_k there one has $m_k \Vdash_B \psi(\beta_k, \alpha f_k)$ for some $\beta_k : m_k \to [x, \top]$, and hence, by inductive hypothesis, such that $m_k \Vdash \psi(\beta_k, \alpha f_k)$. Since the arrows $f_k : m_k \to p$ form a basic cover of p, we must have $p \Vdash \exists x\psi(x, \alpha)$.

6. Suppose $\phi = \forall x\psi(x)$. If $p \Vdash \forall x\psi(x, \alpha)$, for every $f : c \to p$ in the category and every $\beta : c \to [x, \top]$ one has $c \Vdash \psi(\beta, \alpha)$. In particular, this holds when c is any node q in the tree above p, and by inductive hypothesis one has $q \Vdash_B \psi(\beta, \alpha)$ for all such nodes. Therefore, $p \Vdash_B \forall x\psi(x, \alpha)$.

   Conversely, suppose that $p \Vdash_B \forall x\psi(x, \alpha)$ and consider an arrow $f : c \to p$. Together with the identity, this arrow forms a covering pair which appears at some point in the enumeration and is hence pulled back along paths g_j of a finite subtree to build the next level of the subtree over p. Suppose we have some $\beta : c \to [x, \top]$; then we have arrows $\beta f^*(g_j) : g_j^*(c) \to [x, \top]$, and by definition we must have $g_j^*(c) \Vdash_B \psi(\beta f^*(g_j), \alpha f g'_j)$, so by inductive hypothesis one has $g_j^*(c) \Vdash \psi(\beta f^*(g_j), \alpha f g'_j)$. But $f^*(g_j) : g_j^*(c) \to c$ is a basic cover of c (since the g_j form a basic cover of p), and hence we will have $c \Vdash \psi(\beta, \alpha)$. We have thus proved that $p \Vdash \forall x\psi(x, \alpha)$.

□

Remark 4.2.2. It is possible to prove that at each node q, the forcing relation of each formula is semi-decidable. Indeed, that q forces $\phi(\alpha)$ in the sheaf semantics is equivalent to stating that in the syntactic category the arrow $\alpha : q \to [\mathbf{x}, \top]$ factors through the subobject $[\mathbf{x}, \phi(\bar{x})] \rightarrowtail [\mathbf{x}, \top]$, which in turns involves the provability of certain sequents. More especifically, if $q = [\mathbf{y}, \psi]$, the procedure reduces to check whether some formula $\theta(\mathbf{y}, \mathbf{x})$ amongst all formulas that are provable functional (which are enumerable) satisfies $\psi(\mathbf{y}) \iff \exists \mathbf{x}(\theta \wedge \phi)$.

Proposition 4.2.3. *(IZF+FAN) Enumerable coherent theories are constructively complete for (possibly exploding) coherent models*[1] *over enumerable domains and with semi-decidable satisfaction relations.*

Proof. It is enough to prove that every object in the sheaf model forcing the antecedent $\phi(\alpha)$ of a sequent $\phi \vdash_x \psi$ also forces the consequent $\psi(\alpha)$ for every tuple α in the domain. Construct a weak Beth model over a binary tree as above but taking as the root of the tree a given object forcing $\phi(\alpha)$. For each branch $\mathbf{b}$ of the tree, consider the directed colimit $\mathbf{D_b}$ of all the underlying structures in the nodes of the branch, with the corresponding functions between them. Such a directed colimit is a structure under the definition $R(\bar{x_1}, ..., \bar{x_n}) \iff R(x_1, ..., x_n)$ for some representatives x_i of $\bar{x_i}$. We will show that such a structure is a (possibly exploding) coherent model of the theory satisfying $\phi(\bar{\alpha})$. Indeed, we have the following:

Claim : Given any coherent formula $\phi(\bar{x_1}, ..., \bar{x_n})$, we have $D_b \vDash \phi(\bar{\alpha_1}, ..., \bar{\alpha_n})$ if and only if for some node n in the path $\mathbf{b}$, the underlying structure C_n satisfies $C_n \Vdash \phi(\alpha_1, ..., \alpha_n)$ for some representatives α_i of $\bar{\alpha_i}$.

The proof of the claim is by induction on the complexity of ϕ.

1. If ϕ is $R(x_1, ..., x_s)$, the result follows by definition of the structure.

2. If ϕ is of the form $\theta \wedge \eta$ the result follows from the inductive hypothesis.

3. If ϕ is of the form $\theta \vee \eta$ and $\mathbf{D_b} \vDash \phi(\bar{\alpha_1}, ..., \bar{\alpha_s})$, then we can assume that $\mathbf{D_b} \vDash \theta(\bar{\alpha_1}, ..., \bar{\alpha_s})$ or that $\mathbf{D_b} \vDash \eta(\bar{\alpha_1}, ..., \bar{\alpha_s})$, so that by inductive hypothesis we get $C_n \Vdash \phi(\alpha_1, ..., \alpha_s)$ for some node n in $\mathbf{b}$ in either case. Conversely, if $C_n \Vdash \phi(\alpha_1, ..., \alpha_s)$ for some node n in $\mathbf{b}$, by definition of the forcing there is a node m above n in $\mathbf{b}$ and a function $f_{nm} : D_n \to D_m$ for which either $C_m \Vdash \theta(f_{nm}(\alpha_1), ..., f_{nm}(\alpha_s))$ or $C_m \Vdash \eta(f_{nm}(\alpha_1), ..., f_{nm}(\alpha_s))$, so that by inductive hypothesis we get $\mathbf{D_b} \vDash \phi(\bar{\alpha_1}, ..., \bar{\alpha_s})$ in either case.

4. Finally, if ϕ is of the form $\exists x \psi(x, \bar{\alpha_1}, ..., \bar{\alpha_s})$ and $\mathbf{D_b} \vDash \phi(\bar{\alpha_1}, ..., \bar{\alpha_s})$, then $\mathbf{D_b} \vDash \psi(\bar{\alpha}, \bar{\alpha_1}, ..., \bar{\alpha_s})$ for some $\bar{\alpha}$, and then $C_n \Vdash \psi(\alpha, \alpha_1, ..., \alpha_s)$

[1]Technically the models are positive coherent, since they might be exploding, but we trust that the context prevents confusion.

for some node n by inductive hypothesis. Conversely, if $C_n \Vdash \phi(\alpha_1, ..., \alpha_s)$ for some node n in $\mathbf{b}$, then by definition of the forcing there is a node m above n in $\mathbf{b}$ and a function $f_{nm} : D_n \to D_m$ for which $C_m \Vdash \psi(f_{nm}(\alpha), f_{nm}(\alpha_1), ..., f_{nm}(\alpha_s))$, which implies that $\mathbf{D_b} \vDash \psi(\bar{\alpha}, \bar{\alpha_1}, ..., \bar{\alpha_s})$ and hence $\mathbf{D_b} \vDash \phi(\bar{\alpha_1}, ..., \bar{\alpha_s})$.

The (possibly exploding) coherent model has an enumerable domain (since the underlying sets of the structures in the colimit are themselves enumerable) and the satisfaction relation there is semi-decidable (since by the claim it can be stated in terms of the forcing relation of the nodes in the path, which is itself semi-decidable by Remark 4.2.2). We can now define a bar on the tree by saying that a sequence is in the bar if the leaf node, of level, say, n, forces the sentence, and the proof of this uses the first n axioms of the theory in the enumeration. Then this bar is decidable. Therefore, since $\psi(\bar{\alpha})$ is satisfied in all (possibly exploding) coherent models of the theory, $\psi(\alpha)$ is forced at a certain node of every branch of the tree. By the FAN theorem, there is a uniform level of the tree each of whose nodes forces $\psi(\alpha)$. Therefore, $\psi(\alpha)$ is also forced at the root. □

A simplified version of the ideas above can be also used to obtain the completeness theorem for $Reg_\perp$ theories. Consider the fragment $Reg_\perp$ of first-order logic, the syntactic category of a theory over this fragment is a regular category with a strict initial object. If we consider the topos of sheaves over this category with the $Reg_\perp$ coverage given by finite epimorphic families of at most one arrow (so a cover is either empty or a single epimorphism), the coverage is subcanonical and the topos is a conservative sheaf model for the theory. We have now:

Theorem 4.2.4. *Any enumerable $Reg_\perp$ theory has a universal linear weak Beth model.*

Proof. Consider the syntactic category $\mathcal{C}_{\mathbb{T}}$ of the theory and its conservative embedding in the topos of sheaves with the $Reg_\perp$ coverage, $\mathcal{C}_{\mathbb{T}} \to \mathcal{S}h(\mathcal{C}_{\mathbb{T}}, \tau_R)$. Consider an enumeration of all possible proofs; it is easy to see that the set of epimorphisms over any object is then a semi-decidable subset of the enumerable set of all arrows (indeed, if the object is $[\mathbf{x}, \phi]$, it consists of those provably functional formulas $\theta(\mathbf{y}, \mathbf{x})$ such that $\phi \vdash_{\mathbf{x}} \exists \mathbf{y}\theta)$. Hence, the subset is also enumerable. Construct a functor from the linear tree to the syntactic category, defined recursively on the levels of the tree. Start with a denumeration $n : \mathbf{N} \times \mathbf{N} \to \mathbf{N}$

of $\mathbf{N} \times \mathbf{N}$ with the property that $n(k,l) \geq l$. We describe by an inductive definition how the tree obtained as the image of the functor is constructed.

The root of the tree is the terminal object. Suppose now that the tree is defined up to level $n = n(k,l)$. In particular, the node p_l at level l is defined because of the hypothesis $n(k,l) \geq l$. Consider the morphisms g_l^n over p_l assigned to the path from the node p_l to the node at level n. To define the node at level $n+1$, take then the $k-th$ epimorphism over p_l and pull it back along the morphism g_l^n. This produces an epimorphism over the node at level n, whose domain is then the node of level $n+1$.

Clearly, the morphism assigned to the path from any node p till the node of level m in the subtree over p is itself an epimorphism. Define now a Beth model B over this linear tree by defining as the underlying set of a node q the set of arrows from q to the object $[x, \top]$ in the syntactic category, and where the function between the underlying set of a node and its successor is given by composition with the corresponding arrow. By definition we set $q \Vdash_B R(\alpha)$ if and only if q forces $R(\alpha)$ in the sheaf semantics of the topos (we identify the category with its image through Yoneda embedding). That p forces $\phi(\alpha)$ in the sheaf semantics of the topos will be denoted by $p \Vdash \phi(\alpha)$. We can then prove the following:

Claim : For every node p, every $Reg_{\perp}$ formula ϕ and every tuple α, $p \Vdash \phi(\alpha)$ if and only if $p \Vdash_B \phi(\alpha)$.

The proof goes by induction on ϕ similarly to the proof in the case of coherent theories.

1. If ϕ is atomic, the result is immediate by definition of the underlying structures on each node.

2. If $\phi = \psi \wedge \theta$, the result follows easily from the inductive hypothesis, since we have $p \Vdash \psi(\alpha) \wedge \theta(\alpha)$ if and only if $p \Vdash \psi(\alpha)$ and $p \Vdash \theta(\alpha)$, if and only if $p \Vdash_B \psi(\alpha)$ and $p \Vdash_B \theta(\alpha)$, if and only if $p \Vdash_B \psi(\alpha) \wedge \theta(\alpha)$.

3. Suppose $\phi = \exists x \psi(x)$. If $p \Vdash \exists x \psi(x, \alpha)$, then there is a basic cover $f : A \to p$ such that one has $A \Vdash \psi(\beta, \alpha f)$ for some $\beta : A \to [x, \top]$ (as in the coherent case, if p is covered by the empty cover, p is isomorphic to the initial object and then we can take f to be the identity map). This cover of p appears at some point in the enumeration and is hence pulled back along the path g of a finite subtree to create the node of a certain level of the subtree over p. The

node m in this level will have the property that $m \Vdash \psi(\beta g', \alpha f g')$, and hence, by inductive hypothesis, that $m \Vdash_B \psi(\beta g', \alpha f g')$. By definition, we get thus $p \Vdash_B \exists x \psi(x, \alpha)$.

Conversely, suppose that $p \Vdash_B \exists x \psi(x, \alpha)$. Then there is a level in the subtree over p such that the node m there satisfies $m \Vdash_B \psi(\beta, \alpha f)$ for some $\beta : m \to [x, \top]$, and hence, by inductive hypothesis, such that $m \Vdash \psi(\beta, \alpha f)$. Since the arrow $f : m \to p$ is a basic cover of p, we must have $p \Vdash \exists x \psi(x, \alpha)$.

□

As a consequence, we immediately get:

Proposition 4.2.5. *(IZF) Enumerable $Reg_\perp$ theories are complete with respect to (possibly exploding) regular models over enumerable domains and with semi-decidable satisfaction relations.*

Proof. It is enough to prove that every object in the sheaf model forcing the antecedent $\phi(\alpha)$ of a sequent $\phi \vdash_x \psi$ also forces the consequent $\psi(\alpha)$ for every tuple α in the domain. We can thus consider a weak Beth model over a linear tree constructed as above but taking instead as root of the tree a given object forcing $\phi(\alpha)$, and the directed colimit $\mathbf{D}$ of all the underlying structures in the nodes of the tree. We then make it into a structure with the expected definition and prove the following:

Claim : Given any $Reg_\perp$ formula $\phi(\bar{x_1}, ..., \bar{x_n})$, we have $D \vDash \phi(\bar{\alpha_1}, ..., \bar{\alpha_n})$ if and only if for some node n in the tree, the underlying structure C_n satisfies $C_n \Vdash \phi(\alpha_1, ..., \alpha_n)$ for some representatives α_i of $\bar{\alpha_i}$.

The proof is similar to the coherent case.

1. If ϕ is $R(x_1, ..., x_s)$, the result follows by definition of the structure.

2. If ϕ is of the form $\theta \wedge \eta$ the result follows from the inductive hypothesis.

3. Finally, if ϕ is of the form $\exists x \psi(x, \bar{\alpha_1}, ..., \bar{\alpha_s})$ and $\mathbf{D} \vDash \phi(\bar{\alpha_1}, ..., \bar{\alpha_s})$, then $\mathbf{D} \vDash \psi(\bar{\alpha}, \bar{\alpha_1}, ..., \bar{\alpha_s})$ for some $\bar{\alpha}$, and then $C_n \Vdash \psi(\alpha, \alpha_1, ..., \alpha_s)$ for some node n by inductive hypothesis. Conversely, if $C_n \Vdash \phi(\alpha_1, ..., \alpha_s)$ for some node n, then by definition of the forcing there is a node m above n and a function $f_{nm} : D_n \to D_m$ for which $C_m \Vdash \psi(f_{nm}(\alpha), f_{nm}(\alpha_1), ..., f_{nm}(\alpha_s))$, which implies that $\mathbf{D} \vDash \psi(\bar{\alpha}, \bar{\alpha_1}, ..., \bar{\alpha_s})$ and hence $\mathbf{D} \vDash \phi(\bar{\alpha_1}, ..., \bar{\alpha_s})$.

Therefore, since any $Reg_\perp$ formula satisfied in the models given by the directed colimits of the underlying structures of the nodes in the linear trees is forced at their roots, $\psi(\alpha)$ is forced at the roots, as we wanted to prove. □

4.3 Veldman's theorem and the rôle of FAN

In what follows it will be convenient to consider logics in which the disjunction $\vee$ is not part of the group of connectives considered. Classically, of course, this makes no change, since the effect of a disjunct can be recovered in terms of conjunctions and negations. But intuitionistically there is an important difference, that will become clearer when we present the completeness theorems (when removing disjunction, the metatheory one needs to prove completeness is considerably simplified). We start with the following:

Theorem 4.3.1. *(IZF) Every enumerable pre-Heyting theory has a universal modified*[1] *presheaf model with a semi-decidable forcing relation.*

Proof. Consider the syntactic category $\mathcal{C}$ of the $\mathcal{R}eg_\perp$-Morleyization $\mathbb{T}^m$ of a pre-Heyting theory $\mathbb{T}$, which is a regular category with a strict initial object and right adjoints for every pullback functor between subobject categories. Let $\mathcal{R}eg(\mathcal{C})$ be a suitable sized restriction[2] of the category of (possibly exploding) regular models of $\mathbb{T}^m$ with enumerable domains and semi-decidable satisfaction relations, and where arrows are model homomorphisms. We have a functor $ev : \mathcal{C} \to \mathbf{Set}^{\mathcal{R}eg(\mathcal{C})}$ sending an object A to the evaluation functor $ev(A)$. It is clear that this functor is regular, and by Proposition 4.2.5, it is also conservative. Classically, we can see that it also sends the initial object to the presheaf that takes the value 0 at every model except at an exploding model, where it takes the value 1. We prove, constructively, that ev also preserves $\forall$.

Given an arrow $f : A \to B$, a subobject $C \rightarrowtail A$ and the subobject $Y = \forall_f(C) \rightarrowtail B$, we need to show that $ev(Y) = \forall_{ev(f)}(ev(C))$ as subobject of $ev(B)$. By the definition of $\forall$ in the Heyting category $\mathbf{Set}^{\mathcal{R}eg(\mathcal{C})}$, this reduces to proving the following equivalence, for every $\mathbf{y} \in ev(B)(M) = M(B)$:

$\mathbf{y} \in ev(Y)(M) \iff$ For every model N, for every model homomorphism

[1] The modification consists of the fact that $\perp$ may be forced by some nodes.

[2] It is enough to consider a full subcategory satisfying conditions i) and ii) in Theorem 4.1.3.

$$\phi : M \to N,$$

$$(ev(f)_N)^{-1}(\phi_B(\mathbf{y})) \subseteq ev(C)(N)$$

that is:

$$\mathbf{y} \in M(Y) \iff \text{ For every model } N, \text{ for every model homomorphism}$$

$$\phi : M \to N,$$

$$(N(f))^{-1}(\phi_B(\mathbf{y})) \subseteq N(C)$$

The implication $\implies$ can be proven as follows: if $\mathbf{y} \in M(Y)$, then $\phi_B(\mathbf{y}) \in N(Y)$, and so, since N is regular, $\phi_B(\mathbf{y}) = N(f)(\mathbf{x})$ gives $\mathbf{x} \in N(f)^{-1}(N(\forall_f(C))) = N(f^{-1}\forall_f(C)) \subseteq N(C)$.

Let us focus on the other implication. Consider the following diagram in $\mathcal{C}$:

$$\begin{array}{ccc}
C = [\mathbf{x}, \theta] & & \forall_f(C) = [\mathbf{y}, \gamma] \\
\downarrow & & \downarrow \\
A = [\mathbf{x}, \phi] & \xrightarrow[f=[\mathbf{xy},\lambda]]{} & B = [\mathbf{y}, \psi]
\end{array}$$

Applying the functor ev and evaluating at a model M gives the diagram:

$$\begin{array}{ccc}
\{\mathbf{d}|M \vDash \theta(\mathbf{d})\} & & \{\mathbf{c}|M \vDash \gamma(\mathbf{c})\} \\
\downarrow & & \downarrow \\
\{\mathbf{d}|M \vDash \phi(\mathbf{d})\} & \xrightarrow[\{\mathbf{d},\mathbf{c}|M\vDash\lambda(\mathbf{d},\mathbf{c})\}]{} & \{\mathbf{c}|M \vDash \psi(\mathbf{c})\}
\end{array}$$

Given $\mathbf{c} \in \forall_{ev(f)}(ev(C))$, we need to prove that $M \vDash \gamma(\mathbf{c})$. Consider the positive diagram of M, $Diag_+(M)$, which, in a language extended with constants c for every element c of the underlying set of M, consists of all sequents of the form $\top \vdash \psi(c_1, ..., c_m)$ for every positive atomic ψ such that $M \vDash \psi(c_1, ..., c_m)$ (we identify the constants symbols with the elements of M, to simplify the exposition). If N' is a model of $Th(M)$, then, defining N as the reduct of N' with respect to the elements $\{c^{N'} : c \in M\}$ we can define $\phi : M \to N$ by $\phi(c) = c^{N'}$, which is a well defined model homomorphism. But we know that for all $\phi : M \to N$ one has $N(f)^{-1}(\phi_B(\mathbf{c})) \subseteq N(C)$. This implies that for all models N' of $Th(M)$, the sequent $\lambda(\mathbf{x}, \mathbf{c}/\mathbf{y}) \vdash_{\mathbf{x}} \theta(\mathbf{x})$ holds, and therefore, the sequent $\psi(\mathbf{c}) \wedge \lambda(\mathbf{x}, \mathbf{c}/\mathbf{y}) \vdash_{\mathbf{x}} \theta(\mathbf{x})$ also holds.

By the assumption on completeness, this means that such a sequent is provable in $Th(M)$. Besides sequents in $\mathbb{T}^m$, this proof uses a finite number of sequents of the general form $\top \vdash \phi_i(\mathbf{c}, \mathbf{c_1}, ..., \mathbf{c_n})$, where the ϕ_i are positive atomic sentences corresponding to the diagram of M and the $\mathbf{c_i}$ are elements of M. Considering the conjunction ξ of the ϕ_i, we see that there is a proof in $\mathbb{T}^m$ from:

$$\top \vdash \xi(\mathbf{c}, \mathbf{c_1}, ..., \mathbf{c_n})$$

to

$$\psi(\mathbf{c}) \wedge \lambda(\mathbf{x}, \mathbf{c}/\mathbf{y}) \vdash_{\mathbf{x}} \theta(\mathbf{x})$$

By the deduction theorem (Lemma 2.1.6), since $\xi(\mathbf{c}, \mathbf{c_1}, ..., \mathbf{c_n})$ is a sentence, we obtain in $\mathbb{T}^m$ a derivation of:

$$\xi(\mathbf{c}, \mathbf{c_1}, ..., \mathbf{c_n}) \wedge \psi(\mathbf{c}) \wedge \lambda(\mathbf{x}, \mathbf{c}/\mathbf{y}) \vdash_{\mathbf{x}} \theta(\mathbf{x})$$

But it is always possible to replace the constants by variables as long as they are added to the contexts of the sequents, so using the existential rule, we have also a derivation of:

$$\exists \mathbf{x_1} ... \mathbf{x_n} \xi(\mathbf{y}, \mathbf{x_1}, ..., \mathbf{x_n}) \wedge \psi(\mathbf{y}) \wedge \lambda(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{xy}} \theta(\mathbf{x})$$

Calling $Y' = [\mathbf{y}, \Phi(\mathbf{y})]$ the subobject of B given by the interpretation in $\mathcal{C}$ of the formula:

$$\exists \mathbf{x_1} ... \mathbf{x_n} \xi(\mathbf{y}, \mathbf{x_1}, ..., \mathbf{x_n}) \wedge \psi(\mathbf{y})$$

we have a proof of the sequent:

$$\Phi(\mathbf{y}) \wedge \lambda(\mathbf{x}, \mathbf{y})) \vdash_{\mathbf{xy}} \theta(\mathbf{x})$$

and hence also of the sequent:

$$\exists \mathbf{y}(\Phi(\mathbf{y}) \wedge \lambda(\mathbf{x}, \mathbf{y})) \vdash_{\mathbf{x}} \theta(\mathbf{x})$$

Now the antecedent is precisely the pullback of the subobject $\Phi(\mathbf{y})$ of B along f, so by adjunction we have $Y' \leq \forall_f(C) = [\mathbf{y}, \gamma]$, i.e., the sequent $\Phi(\mathbf{y}) \vdash_{\mathbf{y}} \gamma(\mathbf{y})$ is provable. Therefore, since $M \vDash \Phi(\mathbf{c})$, it follows that $M \vDash \gamma(\mathbf{c})$, as we wanted to prove.

□

We have seen that the evaluation functor is regular, but it does not preserve the initial object, as $ev(0)$ is not the constant functor 0. Thus it can be viewed as a modified presheaf model of $\mathbb{T}$. A standard (sheaf) model is obtained by sheafifying with the least coverage so that $ev(0)$ is identified with 0, obtaining thus a model of the theory in a subtopos $\mathrm{Sh}(\mathcal{R}eg(\mathcal{C}), E)$ (see [FE]). Note that classically the standard (i.e., non-exploding) models in $\mathcal{M}$ are dense (in the sense of [Joh02]), so that then $\mathrm{Sh}(\mathcal{R}eg(\mathcal{C}), E) \simeq \mathbf{Set}^{\mathcal{M}^s}$ where $\mathcal{M}^s$ is the full subcategory of standard models. We state the preceeding in the following:

Theorem 4.3.2. *(Forssell) If $\mathcal{C}$ is the syntactic category of an enumerable pre-Heyting theory, the functor*

$$ev : \mathcal{C} \to \mathrm{Sh}(\mathcal{R}eg(\mathcal{C}), E)$$

is a conservative pre-Heyting embedding.

Corollary 4.3.3. *(IZF) Every enumerable pre-Heyting theory has a universal modified Kripke model.*

Proof. As explained, e.g., in [MM94], there is a conservative Heyting embedding of any sheaf topos into a localic topos, where the locale is the Diaconescu cover of the topos. By using the explicit construction of such embedding for the topos $\mathrm{Sh}(\mathcal{R}eg(\mathcal{C}), E)$, it is straightforward to check that Kripke-Joyal semantics there is equivalent to that of the localic topos, and that this latter semantics is precisely modified Kripke semantics (the underlying poset of the modified Kripke model being precisely the Diaconescu cover). □

Remark 4.3.4. Theorem 4.3.2 can be strengthened by removing the restriction on the cardinality of the signature. This is worked out in [FE].

It is possible to recast Joyal's completeness theorem using the completeness of coherent enumerable theories, which leads us finally to the following:

Theorem 4.3.5. *For enumerable theories, the completeness of modified Kripke semantics is provable in IZF+FAN.*

Proof. It is enough to reproduce the proofs of Theorem 4.3.1 and Corollary 4.3.3, replacing the category $\mathcal{R}eg(\mathcal{C})$ of regular models with the category $\mathcal{C}oh(\mathcal{C})$ of coherent models (where the domain is enumerable and the satisfaction relation semi-decidable) to get the universal sheaf model, and then consider the Diaconescu cover over it. This localic topos is essentially Veldman's universal modified Kripke model (see [Vel76]). In fact, the site of such topos is a forest with the natural topology induced by the exploding topology on $\mathcal{C}oh(\mathcal{C})$, and the Kripke-Joyal semantics in there corresponds precisely to the exploding Kripke semantics in each of the trees of the forest. □

It is also possible to consider classical theories in an intuitionistic meta-theory. If we add the law of excluded middle to a first-order theory and define a modified (i.e., possibly exploding) Tarski model as a set in which the extension of every formula over the language of the theory is complemented, then this modified Tarski semantics is not only sound, but also complete, as shown in the following result (compare with the completeness result in [CLR01]):

Corollary 4.3.6. *For classical enumerable theories, the completeness of modified Tarski semantics is provable in IZF+FAN.*

Proof. It is enough to note that Morleyization does not affect enumerability and that (see [Joh02]) a proof in the Morleyized theory translates into a classical proof of the original theory. □

The use of the FAN theorem we have made can be seen to be unavoidable, in the sense that the completeness theorems of enumerable theories for any of the modified semantics considered (Beth, Kripke or Tarski) is in fact equivalent to it. We have:

Theorem 4.3.7. *(Forssell-Gylterud) The completeness theorem of enumerable theories with respect to enumerable modified coherent models in which the satisfaction relation is semi-decidable, entails the FAN theorem.*

Proof. Let F be a fan (finitely branching) containing a decidable bar B. Consider the theory $\mathbb{T}$ over a language including a constant e, constants α for each element $\alpha \in F$ (we can identify each $\alpha \in F$ with the corresponding constant in the language), and a unary relation symbol P, and whose axioms are:

1. $\top \vdash P(r)$, where r is the root of the Fan;
2. $P(\alpha) \vdash \bigvee_{\alpha n \in F} P(\alpha n)$, for each $\alpha \in F$
3. $P(\alpha) \vdash P(e)$, for each $\alpha \in B$.

If N is any enumerable model of $\mathbb{T}$ with a semi-decidable satisfaction relation, it is not hard to define by recursion a maximal chain σ_N such that for each $\alpha \in \sigma_N$ we have $N \vDash P(\alpha)$. Since B is a bar, it follows that $N \vDash P(e)$ for every enumerable model N with a semi-decidable satisfaction relation. By completeness, $\top \vdash P(e)$ is provable in $\mathbb{T}$. Because this proof contains finitely many axioms, it follows that B must be uniform. □

Corollary 4.3.8. *The completeness theorem of enumerable theories with respect to countable modified Kripke (resp. Beth, Tarski) models in which the underlying set of each node is enumerable and the forcing relation is semidecidable, entails the FAN theorem.*

Proof. By the previous theorem, it is enough to show that the completeness of each of these semantics entails the modified coherent completeness. Consider first the case of modified Kripke semantics. Given an enumerable coherent theory, suppose a coherent sequent is valid in all countable modified coherent models with a semi-decidable satisfaction relation. Then it is necessarily forced at every node of every modified Kripke model, and therefore provable from the axioms of the theory in first-order logic. By conservativity, it has to be provable already in coherent logic.

For the case of Beth semantics, it is not in general true that the structure at each node is a modified coherent model, but the directed colimit of the underlying structures in the nodes of a maximal chain is; moreover, any coherent sequent valid in the structure given by such a colimit has to be forced at some node of the chain, from which we deduce it has to be provable by using Beth completeness. Again by conservativity, there is a proof in coherent logic.

Finally, the case of modified Tarski semantics is immediate since every modified Tarski model of a coherent theory is also a modified coherent model, and classical logic is constructively conservative over coherent logic. □

Part IV

Completeness in classical settings

5. Infinitary intuitionistic propositional logic

5.1 Weakly compact cardinals

It is known that even when there are only two propositional variables, there is a proper class of non-equivalent propositional formulas in $\mathcal{L}_\infty$ (see [dJ80]). Therefore, when studying full propositional logic $\mathcal{L}_\kappa$, we are required to appeal to large cardinal axioms for κ in order to continue with our study of completeness theorems. It turns out that the notion of weakly compact cardinal is needed for our purposes here. We start by stating, in this respect, the following:

Definition 5.1.1. A cardinal κ is weakly compact if it is inaccessible and has the tree property, i.e., every tree of height κ whose levels have size strictly less than κ has a cofinal branch.

Weakly compact cardinal were first studied in relation to the compactness theorems for infinitary classical logics. One can prove that the definition above is, for an inaccessible κ, equivalent to stating that a generalized version of the compactness theorem holds for every κ-first-order classical theory with at most κ-many axioms. More precisely, if every subtheory with strictly less than κ axioms has a Tarski model, then the whole theory has a Tarski model. In the axiomatization of [Kar64], the completeness theorem also holds for these type of theories, and in fact from the equivalence of the completeness theorem with model existence theorem, one can see that the completeness of κ-first-order classical theory with at most κ-many axioms implies that κ is weakly compact. For other equivalences of weak compactness of a combinatorial character, we refer to [Dra74] and [Kan09].

The weak compactness of κ is a large cardinal notion. It is relatively mild in strength within the large cardinal hierarchy, but is much stronger than just inaccessibility. As with any other large cardinal, the existence of a weakly compact cardinal is unprovable from ZFC even if one includes as extra axioms the existence of inaccessible cardinals.

Definition 5.1.2. A κ-complete lattice will be called κ-distributive if it satisfies the intuitionistic distributivity law, i.e., if for every $\gamma < \kappa$ and all elements $\{a_f : f \in \gamma^\beta, \beta < \gamma\}$ such that

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

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

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

for all limit β, $f \in \gamma^\beta, \beta < \gamma$, we have that

$$a_\emptyset \leq \bigvee_{f \in \gamma^\gamma} \bigwedge_{\beta < \gamma} a_{f|_\beta}.$$

A κ-complete filter in the lattice is a filter such that whenever $a_i \in \mathcal{F}$ for every $i \in I$, $|I| < \kappa$, then $\bigwedge_{i \in I} a_i \in \mathcal{F}$. A κ-prime filter in the lattice is a filter $\mathcal{F}$ such that whenever $\bigvee_{i \in I} a_i$ is in $\mathcal{F}$ for $|I| < \kappa$ then $a_i \in \mathcal{F}$ for some $i \in I$.

In the next section we study two-valued completeness. Although this is really a particular case of the completeness of κ-first-order theories, we prefer to present this in detail as it does not rely on categorical logic in an essential way.

5.2 Completeness of infinitary intuitionistic propositional logic

We will need the following technical lemma, which is an infinitary generalization of the canonical well-ordering we used in the proof of Theorem 4.2.1, and corresponds to the canonical well-ordering of $\kappa \times \kappa$ from [Jec03]:

Lemma 5.2.1. *For every regular cardinal κ there is a well-ordering $f : \kappa \times \kappa \to \kappa$ with the property that $f(\beta, \gamma) \geq \gamma$*

Proof. We define f by induction on $\max(\beta, \gamma)$ as follows:

$$f(\beta, \gamma) = \begin{cases} \sup\{f(\beta', \gamma') + 1 : \beta', \gamma' < \gamma\} + \beta & \text{if } \beta < \gamma \\ \sup\{f(\beta', \gamma') + 1 : \beta', \gamma' < \beta\} + \beta + \gamma & \text{if } \gamma \leq \beta \end{cases}$$

which clearly satisfy the required property. □

Proposition 5.2.2. *For a weakly compact cardinal κ, any non-trivial κ-complete, κ-distributive lattice of cardinality at most κ has a κ-complete, κ-prime filter. Moreover, the intersection of such filters is* 1.

Proof. Let $\mathcal{H}$ be a κ-complete, κ-distributive lattice of cardinality at most κ. We will first show how to build a κ-complete, κ-prime filter $\mathcal{F}$ in $\mathcal{H}$.

For each $a \in \mathcal{H}$, let $\mathcal{C}(a)$ be the set of all tuples $(b_\alpha)_{\alpha<\lambda}$ of elements of $\mathcal{H}$ such that $a = \bigvee_{\alpha<\lambda} b_\alpha$. Assume without loss of generality that $\mathcal{C}(a)$ is well-ordered and has order type κ (repeating tuples, if needed). Then we can build a tree of height κ whose nodes are elements of $\mathcal{H}$ and whose partial order corresponds to the reverse partial order on $\mathcal{H}$, through the following transfinite recursive procedure. The root of the tree is the top element of $\mathcal{H}$. Assuming that the tree is defined for all levels $\lambda < \mu$; we show how to define the nodes of level μ. Suppose first that μ is a successor ordinal $\mu = \alpha + 1$, and let $\alpha = f(\beta, \gamma)$. Since by hypothesis $f(\beta, \gamma) \geq \gamma$, the nodes $\{p_i\}_{i<m_\gamma}$ at level γ are defined. To define the nodes at level $\alpha+1$, we need to define the successors of a node n there; for this purpose, take then the $\beta - th$ tuple $(b_\alpha)_{\alpha<\lambda} \in \mathcal{C}(p)$ over the predecessor p of n at level γ, and define the successors of n to be $(n \wedge b_\alpha)_{\alpha<\lambda}$. Suppose now that μ is a limit ordinal. Then define every node at level μ to be the conjunction of the predecessors.

By construction, and because of the distributivity property, the join of all elements in any given level of a subtree generated by any given node n is equivalent to n; in particular, the join of all elements corresponding to nodes in any given level of the tree is equivalent to 1, and hence at least one node in that level is not 0. Since κ is weakly compact, we see that the subtree of all nodes that are not 0 must then have an infinite branch B. Then we can define $\mathcal{F}$ by stipulating $a \in \mathcal{F}$ if and only if there is some node b in B such that $b \leq a$.

Claim : $\mathcal{F}$ is a κ-complete, κ-prime filter.

1. It is clearly seen to contain 1 and not to contain 0.

2. If $a \leq b$ and $a \in \mathcal{F}$ then clearly $b \in \mathcal{F}$.

3. It is closed under $\gamma < \kappa$ conjunctions. Indeed, if for $\alpha < \nu$ we have that b_α at level $l(\alpha)$ witnesses that $a_\alpha \in \mathcal{F}$, then $\bigwedge_{\alpha<\nu} b_\alpha$ will be the node of the branch at level $\sup_{\alpha<\nu} l(\alpha)$, and will witness that $\bigwedge_{\alpha<\nu} a_\alpha \in \mathcal{F}$.

4. For the primeness property, suppose that b witnesses that $\bigvee_{\alpha<\gamma} a_\alpha$ is in $\mathcal{F}$. Then we have $(b \wedge a_\alpha)_{\alpha<\gamma} \in \mathcal{C}(b)$, so that by construction,

for some level in the subtree over the node b, every element c in that level will have successors of the form $b \wedge a_\alpha \wedge c$, which implies, by definition, that some a_α will be in $\mathcal{F}$.

To prove that the intersection of all κ-prime filters is the top element 1, consider an element b in that intersection, and suppose it is not 1. It follows that at least one of the successors nodes c of the root is such that $c \leq b$ does not hold, and by construction, the same is true for at least one node in each level of the tree. By weak compactness, there is a branch B composed of such nodes, providing a κ-complete, κ-prime filter not containing b, which is absurd. □

We have now:

Corollary 5.2.3. *For a weakly compact cardinal κ, theories of cardinality at most κ over infinitary intuitionistic propositional logic $\mathcal{L}_\kappa$ are complete for infinitary Kripke semantics.*

Proof. Given the κ-complete Heyting algebra of provable equivalence classes of propositional formulas in $\mathcal{L}_\kappa$, consider the Kripke model whose frame is given by the poset of κ-complete, κ-prime filters ordered by inclusion. The forcing relation given by $P \Vdash \phi \iff \phi \in P$ yields, according to Proposition 5.2.2, a universal infinitary Kripke model. □

Remark 5.2.4. It is also possible to get the completeness of propositional infinitary intuitionistic logic for infinitary Kripke semantics over trees, by considering in the universal Kripke model the posets of finite chains ordered by prolongation. These become Kripke models over trees with the forcing at a chain being induced by the forcing relation in the original model at the last element of the chain.

Remark 5.2.5. The completeness of propositional infinitary logic derived from Corollary 5.2.3 entails that, assuming a proper class of weakly compact cardinals, propositional infinitary intuitionistic logic $\mathcal{L}_\infty$ is complete for infinitary Kripke semantics.

As applications of these completeness theorems, we obtain a generalization of known results from the finitary case:

Corollary 5.2.6. *If κ is a weakly compact cardinal, κ-propositional intuitionistic logic $\mathcal{L}_\kappa$ has the infinitary disjunction property. That is, if $\top \vdash \bigvee_{i \in I} \phi_i$ is provable in the empty theory, then, for some $i \in I$, $\top \vdash \phi_i$ is already provable.*

Proof. This is a straightforward generalization of the usual semantic proof in the finitary case, based on the completeness with respect to Kripke models over trees. If no sequent $\top \vdash \phi_i$ was provable, there would be a countermodel for each. Then we can build a new Kripke tree appending to these countermodels a bottom node forcing no atoms. Such a Kripke tree would then be a countermodel for $\top \vdash \bigvee_{i \in I} \phi_i$. $\square$

Remark 5.2.7. Note that while in finitary propositional logic one has the law $\neg\neg\phi \wedge \neg\neg\psi \to \neg\neg(\phi \wedge \psi)$, the infinitary generalization does not generally hold. Indeed, for the formula $\bigwedge_{i<\omega} \neg\neg p_i \to \neg\neg \bigwedge_{i<\omega} p_i$ one can construct a Kripke countermodel on the linear tree over the natural numbers, by setting node k to force $p_0, ..., p_{k-1}$ while node 0 forces no atoms.

To justify the large cardinal hypothesis made above, we prove in the following result that the assumption of a weakly compact cardinal is essential for obtaining the previous propositions, which exhibits an equivalent form of weak compactness:

Proposition 5.2.8. *(ZFC) Given an inaccessible cardinal κ, the existence of a κ-complete, κ-prime filter in any κ-complete, κ-distributive lattice of cardinality at most κ implies that κ is weakly compact.*

Proof. Consider the κ-coherent fragment of full intuitionistic propositional logic, that is the fragment of those sequents where formulas are κ-coherent (i.e., only use $\bigwedge$, $\bigvee$ indexed by cardinals less than κ) and the distributivity rule is restricted to instantiations on κ-coherent formulas only. Given a tree of height κ and levels of size less than κ, consider the theory of a branch, over a language containing one propositional variable P_a for every node a in the tree and axiomatized as follows:

$$\top \vdash \bigvee_{a \in L_\alpha} P_a$$

for each $\alpha < \kappa$, where L_α is set of all nodes at level α;

$$P_a \wedge P_b \vdash \bot$$

for each pair $a \neq b \in L_\alpha$ and each $\alpha < \kappa$;

$$P_a \vdash P_b$$

for each pair a, b such that a is a successor of b.

Then this is a κ-coherent theory that is certainly consistent, as every subtheory of cardinality less than κ has a model, so the κ-complete, κ-distributive lattice of provable-equivalent classes of κ-coherent formulas is non-trivial. Then the κ-complete, κ-prime filter there corresponds to a model of the whole theory that yields a cofinal branch of the tree. □

5.3 Strongly compact cardinals

The restriction on the cardinalities of the theories considered in the completeness theorems can be removed if we use a stronger large cardinal notion. The key concept is that of a strongly compact cardinal:

Definition 5.3.1. A cardinal κ is strongly compact if and only if every κ-complete filter on a set can be extended to a κ-complete ultrafilter.

Using an infinitary version of Łoś theorem one can prove that the definition above implies (and in fact, is equivalent to) the alternate description of a strongly compact cardinal κ as follows: if every subtheory with less than κ many axioms of a given arbitrary theory has a model, the whole theory has a model. For details see [Dra74]. This property will be quite useful in extending Proposition 5.2.2 to provide yet another characterization of strong compactness:

Proposition 5.3.2. *For a strongly compact cardinal κ, any non-trivial κ-complete, κ-distributive lattice has a κ-complete, κ-prime filter. Moreover, the intersection of such filters is* 1.

Proof. Consider a non-trivial κ-complete, κ-distributive lattice $\mathcal{L}$, and an element c different from 1. It is enough to prove that there is a κ-complete, κ-prime filter not containing c. In a language containing one propositional variable P_a for every element of $\mathcal{L}$, consider the theory of a κ-complete, κ-prime filter, axiomatized as follows:

1. $P_a \vdash P_b$ for every pair $a \leq b$ in $\mathcal{L}$
2. $\bigwedge_{i<\gamma} P_{a_i} \vdash P_{\bigwedge_{i<\gamma} a_i}$ for all families $\{a_i\}_{i<\gamma}$ such that $\gamma < \kappa$
3. $P_{\bigvee_{i<\gamma} a_i} \vdash \bigvee_{i<\gamma} P_{a_i}$ for all tuples $\{a_i\}_{i<\gamma}$ such that $\gamma < \kappa$
4. $P_c \vdash \bot$

Here P_a is thought of as the assertion "a is in the filter". Every subtheory of cardinality less than κ contains less than κ many propositional variables P_{a_i}, and since the corresponding elements a_i generate a κ-complete, κ-distributive sublattice of cardinality at most κ, it must have a model according to Proposition 5.2.2. Since κ is strongly compact, the whole theory has a model, which corresponds to a κ-complete, κ-prime filter not containing c. □

We immediately deduce the following:

Corollary 5.3.3. *For a strongly compact cardinal κ, theories over infinitary intuitionistic propositional logic $\mathcal{L}_\kappa$ are complete for infinitary Kripke semantics.*

Proof. Analogous to the proof of Corollary 5.2.3. □

The assumption of strong compactness in Proposition 5.3.2 is necessary:

Proposition 5.3.4. *If every non-trivial κ-complete, κ-distributive lattice has a κ-complete, κ-prime filter, then κ is strongly compact.*

Proof. Consider a κ-complete filter $\mathcal{F}$ in the lattice $\mathcal{L}$ of subsets of a set. Form the quotient $\mathcal{L}/\mathcal{F}$, which will be a κ-complete, κ-distributive lattice. Any κ-complete, κ-prime filter in $\mathcal{L}/\mathcal{F}$ will be an ultrafilter (since $\mathcal{L}/\mathcal{F}$ is Boolean) and hence its preimage along the quotient map provides a κ-complete ultrafilter extending $\mathcal{F}$. Therefore, κ is strongly compact. □

6. Infinitary intuitionistic first-order logic

6.1 Completeness of infinitary coherent logic

The completeness of infinitary coherent logic is the straightforward infinitary generalization of the corresponding result for the finitary case (4.2.1, 4.2.3), and it can be proved by the same techniques. In our case the use of the canonical enumeration of $\mathbf{N} \times \mathbf{N}$ is replaced by Lemma 5.2.1, which gives a canonical well-ordering of $\kappa \times \kappa$.

Recall that we have a conservative κ-coherent embedding $\mathcal{C}_{\mathbb{T}} \to \mathcal{S}h(\mathcal{C}_{\mathbb{T}}, \tau)$ from the syntactic category into the topos of sheaves with the κ-coherent coverage consisting of jointly covering families of cardinality less than κ. We have now:

Theorem 6.1.1. *Let κ be an inaccessible cardinal. Then any κ-first-order theory of cardinality at most κ has a universal weak Beth model on a tree of height κ and levels of size less than κ.*

Proof. Consider the syntactic category $\mathcal{C}_{\mathbb{T}}$ of the theory and its conservative embedding in the topos of sheaves with the κ-coherent coverage, $\mathcal{C}_{\mathbb{T}} \to \mathcal{S}h(\mathcal{C}_{\mathbb{T}}, \tau)$. Consider for any object a well-ordering of the set of basic covering families over the object, which are given by jointly covering sets of arrows of cardinality less than κ. Because κ is inaccessible, this set has cardinality at most κ; repeating sets if needed we can assume this set has order type exactly κ. Construct a functor from a tree of height κ and levels of size less than κ to the syntactic category, defined recursively on the levels of the tree. Start with a well-ordering $f : \kappa \times \kappa \to \kappa$ as in Lemma 5.2.1, i.e., with the property that $f(\beta, \gamma) \geq \gamma$. We describe by an inductive definition how the tree obtained as the image of the functor is constructed.

The root of that tree is the terminal object. Suppose now that the tree is defined for all levels $\lambda < \mu$; we show how to define the nodes of level μ. Suppose first that μ is a successor ordinal $\mu = \alpha + 1$, and let $\alpha = f(\beta, \gamma)$. Since by hypothesis $f(\beta, \gamma) \geq \gamma$, the nodes $\{p_i\}_{i<m_\gamma}$ at level γ are defined. Consider the morphisms g_{ij}^{α} over p_i assigned to

the paths from each of the nodes p_i to the nodes of level α. To define the nodes at level $\alpha + 1$, take then the $\beta - th$ covering family over each p_i and pull it back along the morphisms g_{ij}^{α}. This produces covering families over each node at level α, whose domains are then the nodes of level $\alpha + 1$. Suppose now that μ is a limit ordinal. Then each branch of the tree of height μ already defined determines a diagram, whose limit is defined to be the node at level μ corresponding to that branch.

The tree has size κ and, because κ is inaccessible, has levels of size less than κ. Clearly, the morphisms assigned to the paths from any node p till the nodes of level α in the subtree over p form a basic covering family of p because of the transfinite transitivity property. Define now a Beth model B over this tree by defining as the underlying set of a node q the set of arrows from q to the object $[x, \top]$ in the syntactic category, and where the function between the underlying set of a node and its successor is given by composition with the corresponding arrow. We set by definition $q \Vdash_B R(\alpha)$ if and only if q forces $R(\alpha)$ in the sheaf semantics of the topos (we identify the category with its image through Yoneda embedding). That p forces $\phi(\alpha)$ in the sheaf semantics of the topos will be denoted as $p \Vdash \phi(\alpha)$. We shall now prove the following:

Claim : For every node p and every tuple α, $p \Vdash \phi(\alpha)$ if and only if $p \Vdash_B \phi(\alpha)$.

The proof goes by induction on ϕ.

1. If ϕ is atomic, the result is immediate by definition of the underlying structures on each node.

2. If $\phi = \bigwedge_{i<\gamma} \psi_i$, the result follows easily from the inductive hypothesis, since we have $p \Vdash \bigwedge_{i<\gamma} \psi_i(\alpha)$ if and only if $p \Vdash \psi_i(\alpha)$ for each $i < \gamma$, if and only if $p \Vdash_B \psi_i(\alpha)$ for each $i < \gamma$, if and only if $p \Vdash_B \bigwedge_{i<\gamma} \psi_i(\alpha)$.

3. Suppose $\phi = \bigvee_{i<\gamma} \psi_i$. If $p \Vdash \bigvee_{i<\gamma} \psi_i$, then there is a basic covering family $\{f_i : A_i \to p\}_{i<\lambda}$ such that for each $i < \lambda$, $A_i \Vdash \psi_{k_i}(\alpha f_i)$ for some $k_i < \gamma$. Since this covering family appears at some point in the well-ordering, it is pulled back along all paths g_j of a subtree to create the nodes of a certain level of the subtree over p. Hence, every node m_j in such a level satisfies $m_j \Vdash \psi_{k_j}(\alpha f g'_j)$ for some $k_j < \gamma$. By inductive hypothesis, $m_j \Vdash_B \psi_{k_j}(\alpha f g'_j)$, and hence we have $p \Vdash_B \bigvee_{i<\gamma} \psi_i$.

Conversely, if $p \Vdash_B \bigvee_{i<\gamma} \psi_i$, there is a level in the subtree over p such that for every node m_j there one has $m_j \Vdash_B \psi_{k_j}(\alpha f_j)$ for some $k_j < \gamma$, so by inductive hypothesis $m_j \Vdash \psi_{k_j}(\alpha f_j)$. Since $\{f_k : m_k \to p\}$ is, by construction, a basic covering family, we must have $p \Vdash \bigvee_{i<\gamma} \psi_i$.

4. Suppose $\phi = \psi \to \theta$. If $p \Vdash \psi(\alpha) \to \theta(\alpha)$, for every $f : c \to p$ in the category one has $c \Vdash \psi(\alpha f) \implies c \Vdash \theta(\alpha f)$. In particular, this holds when c is any node q in the tree above p, and by inductive hypothesis one has $q \Vdash_B \psi(\alpha f) \implies q \Vdash_B \theta(\alpha f)$ for all such nodes. Therefore, $p \Vdash_B \psi(\alpha) \to \theta(\alpha)$.

   Conversely, suppose that $p \Vdash_B \psi(\alpha) \to \theta(\alpha)$ and consider an arrow $f : c \to p$. Together with the identity, this arrow forms a covering family which appears at some point of the well-ordering and is hence pulled back along paths g_j of a subtree to build the next level of the subtree over p. Suppose that $c \Vdash \psi(\alpha)$; then $g_j^*(c) \Vdash \psi(\alpha g_j')$, so by inductive hypothesis one has $g_j^*(c) \Vdash_B \psi(\alpha g_j')$. Therefore, we get $g_j^*(c) \Vdash_B \theta(\alpha g_j')$, and using once more the inductive hypothesis, $g_j^*(c) \Vdash \theta(\alpha g_j')$. But $g_j' = f^*(g_j) : g_j^*(c) \to c$ is a basic cover of c (since the g_j form a basic cover of p), and hence we will have $c \Vdash \theta(\alpha)$. We have, thus, proved that $p \Vdash \psi(\alpha) \to \theta(\alpha)$

5. Suppose $\phi = \exists \mathbf{x} \psi(\mathbf{x})$. If $p \Vdash \exists \mathbf{x} \psi(\mathbf{x}, \alpha)$, then there is a basic covering family $\{f_i : A_i \to p\}_{i<\lambda}$ such that for each i one has $A_i \Vdash \psi(\boldsymbol{\beta_i}, \alpha f_i)$ for some $\boldsymbol{\beta_i} : A_i \to [\mathbf{x}, \top]$. This basic cover appears at some point in the well-ordering and is hence pulled back along all paths g_j of a subtree to create the nodes of a certain level of the subtree over p. The nodes m_{ij} in this level will have the property that $m_{ij} \Vdash \psi(\boldsymbol{\beta_i} g_j', \alpha f_i g_j')$, and hence, by inductive hypothesis, that $m_{ij} \Vdash_B \psi(\boldsymbol{\beta_i} g_j', \alpha f_i g_j')$. By definition, we get thus $p \Vdash_B \exists \mathbf{x} \psi(\mathbf{x}, \alpha)$.

   Conversely, suppose that $p \Vdash_B \exists \mathbf{x} \psi(\mathbf{x}, \alpha)$. Then there is a level in the subtree over p such that for every node m_k there one has $m_k \Vdash_B \psi(\boldsymbol{\beta_k}, \alpha f_k)$ for some $\boldsymbol{\beta_k} : m_k \to [\mathbf{x}, \top]$, and hence, by inductive hypothesis, such that $m_k \Vdash \psi(\boldsymbol{\beta_k}, \alpha f_k)$. Since the arrows $f_k : m_k \to p$ form a basic cover of p, we must have $p \Vdash \exists \mathbf{x} \psi(\mathbf{x}, \alpha)$.

6. Suppose $\phi = \forall \mathbf{x} \psi(\mathbf{x})$. If $p \Vdash \forall \mathbf{x} \psi(\mathbf{x}, \alpha)$, for every $f : c \to p$ in the category and every $\boldsymbol{\beta} : c \to [\mathbf{x}, \top]$ one has $c \Vdash \psi(\boldsymbol{\beta}, \alpha)$. In particular, this holds when c is any node q in the tree above p, and by inductive hypothesis one has $q \Vdash_B \psi(\boldsymbol{\beta}, \alpha)$ for all such nodes. Therefore, $p \Vdash_B \forall \mathbf{x} \psi(\mathbf{x}, \alpha)$.

Conversely, suppose that $p \Vdash_B \forall \mathbf{x}\psi(\mathbf{x}, \alpha)$ and consider an arrow $f : c \to p$. Together with the identity, this arrow forms a covering family which appears at some point in the well-ordering and is hence pulled back along the paths g_j of a subtree to build the next level of the subtree over p. Suppose we have some $\boldsymbol{\beta} : c \to [\mathbf{x}, \top]$; then we have arrows $\boldsymbol{\beta} f^*(g_j) : g_j^*(c) \to [\mathbf{x}, \top]$, and by definition we must have $g_j^*(c) \Vdash_B \psi(\boldsymbol{\beta} f^*(g_j), \alpha f g_j')$, so by inductive hypothesis one has $g_j^*(c) \Vdash \psi(\boldsymbol{\beta} f^*(g_j), \alpha f g_j')$. But $f^*(g_j) : g_j^*(c) \to c$ is a basic cover of c (since the g_j form a basic cover of p), and hence we will have $c \Vdash \psi(\boldsymbol{\beta}, \alpha)$. We have thus proved that $p \Vdash \forall \mathbf{x}\psi(\mathbf{x}, \alpha)$.

□

Proposition 6.1.2. *If κ is a weakly compact cardinal, κ-coherent theories of cardinality at most κ are complete for κ-coherent models.*

Proof. It is enough to prove that every object in the sheaf model forcing the antecedent $\phi(\alpha)$ of a sequent $\phi \vdash_x \psi$ also forces the consequent $\psi(\alpha)$ for every tuple α in the domain. Construct a weak Beth model over a tree as above but taking as the root of the tree a given object forcing $\phi(\alpha)$. For each branch $\mathbf{b}$ of the tree, consider the directed colimit $\mathbf{D_b}$ of all the underlying structures in the nodes of the branch, with the corresponding functions between them. Such a directed colimit is a structure under the definition $R(\bar{x_0}, ..., \bar{x_\lambda}, ...) \iff R(x_0, ..., x_\lambda, ...)$ for some representatives x_i of $\bar{x_i}$. We will show that such a structure is a (possible exploding) κ-coherent model of the theory satisfying $\phi(\bar{\alpha})$. Indeed, we have the following:

Claim : Given any κ-coherent formula $\phi(\bar{x_0}, ..., \bar{x_\lambda}, ...)$, we have $\mathbf{D_b} \vDash \phi(\bar{\alpha_0}, ..., \bar{\alpha_\lambda}, ...)$ if and only if for some node n in the path $\mathbf{b}$, the underlying structure C_n satisfies $C_n \Vdash \phi(\alpha_0, ..., \alpha_\lambda, ...)$ for some representatives α_i of $\bar{\alpha_i}$.

The proof of the claim is by induction on the complexity of ϕ.

1. If ϕ is $R(x_0, ..., x_s, ...)$, the result follows by definition of the structure.

2. If ϕ is of the form $\bigwedge_{i<\gamma} \theta_i$ the result follows from the inductive hypothesis.

3. If ϕ is of the form $\bigvee_{i<\gamma} \theta_i$ and $\mathbf{D_b} \vDash \phi(\bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$, then we can assume that $\mathbf{D_b} \vDash \theta_i(\bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$ for some $i < \gamma$, so that by inductive hypothesis we get $C_n \Vdash \phi(\alpha_1, ..., \alpha_s, ...)$ for some node n in $\mathbf{b}$. Conversely, if $C_n \Vdash \phi(\alpha_0, ..., \alpha_s, ...)$ for some node n in $\mathbf{b}$, by

definition of the forcing there is a node m above n in $\mathbf{b}$ and a function $f_{nm} : D_n \to D_m$ for which $C_m \Vdash \theta_i(f_{nm}(\alpha_0), ..., f_{nm}(\alpha_s), ...)$ for some $i < \gamma$, so that by inductive hypothesis we get $\mathbf{D_b} \vDash \phi(\bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$.

4. Finally, if ϕ is of the form $\exists \mathbf{x} \psi(\mathbf{x}, \bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$ and $\mathbf{D_b} \vDash \phi(\bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$, then $\mathbf{D_b} \vDash \psi(\bar{\boldsymbol{\alpha}}, \bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$ for some $\bar{\boldsymbol{\alpha}}$, and then $C_n \Vdash \psi(\boldsymbol{\alpha}, \alpha_0, ..., \alpha_s, ...)$ for some node n by inductive hypothesis. Conversely, if $C_n \Vdash \phi(\alpha_0, ..., \alpha_s, ...)$ for some node n in $\mathbf{b}$, then by definition of the forcing there is a node m above n in $\mathbf{b}$ and a function $f_{nm} : D_n \to D_m$ for which $C_m \Vdash \psi(f_{nm}(\boldsymbol{\alpha}), f_{nm}(\alpha_0), ..., f_{nm}(\alpha_s), ...)$, which implies that $\mathbf{D_b} \vDash \psi(\bar{\boldsymbol{\alpha}}, \bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$ and hence $\mathbf{D_b} \vDash \phi(\bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$.

Since $\psi(\bar{\alpha})$ is satisfied in all κ-coherent models of the theory, it is satisfied in all models of the form $\mathbf{D_b}$ (since we are in a classical metatheory, we can treat separately the case when the structure $\mathbf{D_b}$ is exploding). Hence, $\psi(\alpha)$ is forced at a certain node of every branch of the tree. By weak compactness, we can find a level of the tree where every node forces $\psi(\alpha)$. Indeed, suppose that is not the case; then every level of the tree would contain a node not forcing $\psi(\alpha)$, and hence, because the levels have size less than κ, there would be a cofinal branch composed of such nodes, which is absurd. Therefore, there must exist a level each of whose nodes forces $\psi(\alpha)$. Because these nodes form a basic covering family, $\psi(\alpha)$ is therefore forced at the root, as we wanted to prove. □

Remark 6.1.3. In the classical case, when the syntactic category is Boolean, Proposition 6.1.2 reduces to the completeness theorem of Karp in [Kar64]; the transfinite transitivity property can be rewritten into the classical distributivity and dependent choice axiom schemata.

As in the propositional case, one can remove the restriction on the cardinality of the theory if one assumes instead that κ is a strongly compact cardinal:

Proposition 6.1.4. *If κ is a strongly compact cardinal, κ-coherent theories are complete for κ-coherent models.*

Proof. Suppose that the sequent $\phi \vdash_{\mathbf{x}} \psi$ is valid in every model of a certain theory but not provable. Then it is not provable in any subtheory of cardinality less than κ. therefore, if we add to the language a new constant c and axioms $\top \vdash \phi(c)$ and $\psi(c) \vdash \bot$, any subtheory of cardinality less than κ together with these two new axioms has, by

Proposition 6.1.2, a model. Since κ is strongly compact, the whole theory has a model, which provides a model for the original theory where $\phi \vdash_{\mathbf{x}} \psi$ is not valid. □

A simplified version of the ideas above can be also used to obtain the completeness theorem for infinitary $Reg_\perp$ theories. Consider the fragment κ-$Reg_\perp$ of first-order logic; the syntactic category of a theory over this fragment is a regular category with a strict initial object. If we consider the topos of sheaves over this category with the κ-$Reg_\perp$ coverage given by finite epimorphic families of at most one arrow (so a cover is either empty or a single epimorphism), the coverage is subcanonical and the topos is a conservative sheaf model for the theory. We have now:

Theorem 6.1.5. *Let κ be an inaccessible cardinal. Then any κ-$Reg_\perp$ theory of cardinality at most κ has a universal linear weak Beth model of height κ.*

Proof. Consider the syntactic category $\mathcal{C}_\mathbb{T}$ of the theory and its conservative embedding in the topos of sheaves with the κ-$Reg_\perp$ coverage, $\mathcal{C}_\mathbb{T} \to \mathcal{S}h(\mathcal{C}_\mathbb{T}, \tau')$. Consider for any object the set of epimorphisms over the object. Because κ is inaccessible, this set has cardinality at most κ. Repeating elements if needed, we can assume that the set has order-type precisely κ. Construct a functor from a linear tree of height κ to the syntactic category, defined recursively on the levels of the tree. Start with a well-ordering $f : \kappa \times \kappa \to \kappa$ as in Lemma 5.2.1, i.e., with the property that $f(\beta, \gamma) \geq \gamma$. We describe by an inductive definition how the tree obtained as the image of the functor is constructed.

The root of that tree is the terminal object. Suppose now that the tree is defined for all levels $\lambda < \mu$; we show how to define the node of level μ. Suppose first that μ is a successor ordinal $\mu = \alpha + 1$, and let $\alpha = f(\beta, \gamma)$. Since by hypothesis $f(\beta, \gamma) \geq \gamma$, the node p_γ at level γ is defined. Consider the morphism g_γ^α over p_γ assigned to the path from the node p_γ to the node of level α. To define the node at level $\alpha + 1$, take then the $\beta - th$ epimorphism over p_γ and pull it back along the morphism g_γ^α. This produces an epimorphism over the node at level α, whose domain is then the node of level $\alpha + 1$. Suppose now that μ is a limit ordinal. Then the tree of height μ already defined determines a diagram, whose limit is defined to be the node at level μ corresponding to that branch.

Clearly, the morphism assigned to the path from any node p till the node of level α in the linear subtree over p is an epimorphism because of the dependent choice property. Define now a Beth model B over this tree by defining as the underlying set of a node q the set of arrows from q to the object $[x, \top]$ in the syntactic category, and where the function between the underlying set of a node and its successor is given by composition with the corresponding arrow. We set by definition $q \Vdash_B R(\alpha)$ if and only if q forces $R(\alpha)$ in the sheaf semantics of the topos (we identify the category with its image through Yoneda embedding). That p forces $\phi(\alpha)$ in the sheaf semantics of the topos will be denoted as $p \Vdash \phi(\alpha)$. We shall now prove the following:

Claim : For every node p, every κ-$Reg_\perp$ formula ϕ and every tuple α, $p \Vdash \phi(\alpha)$ if and only if $p \Vdash_B \phi(\alpha)$.

The proof goes by induction on ϕ similarly to the proof in the case of κ-coherent theories.

1. If ϕ is atomic, the result is immediate by definition of the underlying structures on each node.

2. If $\phi = \bigwedge_{i<\gamma} \psi_i$, the result follows easily from the inductive hypothesis, since we have $p \Vdash \bigwedge_{i<\gamma} \psi_i(\alpha)$ if and only if $p \Vdash \psi_i(\alpha)$ for each $i < \gamma$, if and only if $p \Vdash_B \psi_i(\alpha)$ for each $i < \gamma$, if and only if $p \Vdash_B \bigwedge_{i<\gamma} \psi_i(\alpha)$.

3. Suppose $\phi = \exists \mathbf{x} \psi(\mathbf{x})$. If $p \Vdash \exists \mathbf{x} \psi(\mathbf{x}, \alpha)$, then there is a basic cover $f : A \to p$ such that one has $A \Vdash \psi(\boldsymbol{\beta}, \alpha f)$ for some $\boldsymbol{\beta} : A \to [\mathbf{x}, \top]$. This cover of p appears at some point of the well-ordering and is hence pulled back along the path g of a subtree to create the node of a certain level of the subtree over p. The node m in this level will have the property that $m \Vdash \psi(\boldsymbol{\beta} g', \alpha f g')$, and hence, by inductive hypothesis, that $m \Vdash_B \psi(\boldsymbol{\beta} g', \alpha f g')$. By definition, we get thus $p \Vdash_B \exists \mathbf{x} \psi(\mathbf{x}, \alpha)$.

   Conversely, suppose that $p \Vdash_B \exists \mathbf{x} \psi(\mathbf{x}, \alpha)$. Then there is a level in the subtree over p such that the node m there satisfies $m \Vdash_B \psi(\boldsymbol{\beta}, \alpha f)$ for some $\boldsymbol{\beta} : m \to [\mathbf{x}, \top]$, and hence, by inductive hypothesis, such that $m \Vdash \psi(\boldsymbol{\beta}, \alpha f)$. Since the arrow $f : m \to p$ is a basic cover of p, we must have $p \Vdash \exists \mathbf{x} \psi(\mathbf{x}, \alpha)$.

□

As a consequence, we immediately get:

Proposition 6.1.6. *(ZFC) If κ is an inaccessible cardinal, κ-$Reg_\perp$ theories of cardinality at most κ are complete with respect to κ-regular models.*

Proof. It is enough to prove that every object in the sheaf model forcing the antecedent $\phi(\alpha)$ of a sequent $\phi \vdash_x \psi$ also forces the consequent $\psi(\alpha)$ for every tuple α in the domain. We can thus consider a weak Beth model over a linear tree as above but taking instead as the root of the tree an object forcing $\phi(\alpha)$, and the directed colimit $\mathbf{D}$ of all the underlying structures in the nodes of the tree. We then make it into a structure with the expected definition and prove the following:

Claim : Given any κ-*Reg*$_\perp$ formula $\phi(\bar{x_0}, ..., \bar{x_\lambda}, ...)$, we have $\mathbf{D} \vDash \phi(\bar{\alpha_0}, ..., \bar{\alpha_\lambda}, ...)$ if and only if for some node n in the tree, the underlying structure C_n satisfies $C_n \Vdash \phi(\alpha_0, ..., \alpha_\lambda, ...)$ for some representatives α_i of $\bar{\alpha_i}$.

The proof is similar to the κ-coherent case.

1. If ϕ is $R(x_0, ..., x_s)$, the result follows by definition of the structure.

2. If ϕ is of the form $\bigwedge_{i<\gamma} \theta_i$ the result follows from the inductive hypothesis.

3. Finally, if ϕ is of the form $\exists \mathbf{x} \psi(\mathbf{x}, \bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$ and $\mathbf{D} \vDash \phi(\bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$, then $\mathbf{D} \vDash \psi(\bar{\boldsymbol{\alpha}}, \bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$ for some $\bar{\boldsymbol{\alpha}}$, and then $C_n \Vdash \psi(\boldsymbol{\alpha}, \alpha_0, ..., \alpha_s, ...)$ for some node n by inductive hypothesis. Conversely, if $C_n \Vdash \phi(\alpha_0, ..., \alpha_s, ...)$ for some node n, then by definition of the forcing there is a node m above n and a function $f_{nm} : D_n \to D_m$ for which $C_m \Vdash \psi(f_{nm}(\boldsymbol{\alpha}), f_{nm}(\alpha_0), ..., f_{nm}(\alpha_s), ...)$, which implies that $\mathbf{D} \vDash \psi(\bar{\boldsymbol{\alpha}}, \bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$ and hence $\mathbf{D} \vDash \phi(\bar{\alpha_0}, ..., \bar{\alpha_s}, ...)$.

Therefore, since any κ-*Reg*$_\perp$ formula satisfied in the models given by the directed colimits of of the underlying structures of the nodes in the linear trees, is forced at their roots, $\psi(\alpha)$ is forced at the roots, as we wanted to prove. □

6.2 Completeness of infinitary intuitionistic first-order logic

Having now at hand a completeness theorem for κ-coherent theories, we can adapt the proof of Joyal's theorem by replacing the category of coherent models with that of κ-coherent models. As a result, we get:

Theorem 6.2.1. *If κ is a weakly compact cardinal, every κ-first-order theory of cardinality at most κ has a universal Kripke model.*

Proof. Consider the syntactic category $\mathcal{C}$ of the κ-coherent Morleyization $\mathbb{T}^m$ of the theory $\mathbb{T}$. Let $\mathcal{C}oh(\mathcal{C})$ be the category of κ-coherent models of $\mathbb{T}^m$ of size at most κ, and where arrows are model homomorphisms. We have a functor $ev : \mathcal{C} \to \mathbf{Set}^{\mathcal{C}oh(\mathcal{C})}$ sending an object A to the evaluation functor $ev(A)$. It is clear that this functor is κ-coherent, and by Proposition 6.1.2, it is also conservative. We must prove that ev also preserves $\forall$.

Given an arrow $f : A \to B$, a subobject $C \rightarrowtail A$ and the subobject $Y = \forall_f(C) \rightarrowtail B$, we need to show that $ev(Y) = \forall_{ev(f)}(ev(C))$ as subobject of $ev(B)$. By the definition of $\forall$ in the Heyting category $\mathbf{Set}^{\mathcal{C}oh(\mathcal{C})}$, this reduces to proving the following equivalence, for every $\mathbf{y} \in ev(B)(M) = M(B)$:

$$\mathbf{y} \in ev(Y)(M) \iff \text{For every model } N, \text{ for every model homomorphism}$$

$$\phi : M \to N,$$

$$(ev(f)_N)^{-1}(\phi_B(\mathbf{y})) \subseteq ev(C)(N)$$

that is:

$$\mathbf{y} \in M(Y) \iff \text{For every model } N, \text{ for every model homomorphism}$$

$$\phi : M \to N,$$

$$(N(f))^{-1}(\phi_B(\mathbf{y})) \subseteq N(C)$$

The implication $\implies$ can be proven as follows: if $\mathbf{y} \in M(Y)$, then $\phi_B(\mathbf{y}) \in N(Y)$, and so, since N is κ-coherent, $\phi_B(\mathbf{y}) = N(f)(\mathbf{x})$ gives $\mathbf{x} \in N(f)^{-1}(N(\forall_f(C))) = N(f^{-1}\forall_f(C)) \subseteq N(C)$.

Let us focus on the other implication. Consider the following diagram in $\mathcal{C}$:

$$\begin{array}{ccc}
C = [\mathbf{x}, \theta] & & \forall_f(C) = [\mathbf{y}, \gamma] \\
\downarrow & & \downarrow \\
A = [\mathbf{x}, \phi] & \xrightarrow[f=[\mathbf{xy},\lambda]]{} & B = [\mathbf{y}, \psi]
\end{array}$$

Applying the functor ev and evaluating at a model M gives the diagram:

$$\begin{array}{ccc}
\{\mathbf{d}|M \vDash \theta(\mathbf{d})\} & & \{\mathbf{c}|M \vDash \gamma(\mathbf{c})\} \\
\big\downarrow & & \big\downarrow \\
\{\mathbf{d}|M \vDash \phi(\mathbf{d})\} & \xrightarrow[\{\mathbf{d},\mathbf{c}|M\vDash\lambda(\mathbf{d},\mathbf{c})\}]{} & \{\mathbf{c}|M \vDash \psi(\mathbf{c})\}
\end{array}$$

Given $\mathbf{c} \in \forall_{ev(f)}(ev(C))$, we need to prove that $M \vDash \gamma(\mathbf{c})$. Consider the positive diagram of M, $Diag_+(M)$, which, in a language extended with constants c for every element c of the underlying set of M, consists of all sequents of the form $\top \vdash \psi(c_0, ..., c_\alpha, ...)$ for every positive atomic ψ such that $M \vDash \psi(c_0, ..., c_\alpha, ...)$ (we identify the constants symbols with the elements of M, to simplify the exposition). If N' is a model of $Th(M)$, then, defining N as the reduct of N' with respect to the elements $\{c^{N'} : c \in M\}$ we can define $\phi : M \to N$ by $\phi(c) = c^{N'}$, which is a well defined model homomorphism. But we know that for all $\phi : M \to N$ one has $N(f)^{-1}(\phi_B(\mathbf{c})) \subseteq N(C)$. This implies that for all models N' of $Th(M)$, the sequent $\lambda(\mathbf{x}, \mathbf{c}/\mathbf{y}) \vdash_{\mathbf{x}} \theta(\mathbf{x})$ holds, and therefore, the sequent $\psi(\mathbf{c}) \wedge \lambda(\mathbf{x}, \mathbf{c}/\mathbf{y}) \vdash_{\mathbf{x}} \theta(\mathbf{x})$ also holds.

By Proposition 6.1.2, this means that such a sequent is provable in $Th(M)$. Besides sequents in $\mathbb{T}^m$, this proof uses less than κ sequents of the general form $\top \vdash \phi_i(\mathbf{c}, \mathbf{c_0}, ..., \mathbf{c}_\alpha, ...)$, where the ϕ_i are positive atomic sentences corresponding to the diagram of M and the $\mathbf{c_i}$ are elements of M. Considering the conjunction ξ of the ϕ_i, we see that there is a proof in $\mathbb{T}^m$ from:

$$\top \vdash \xi(\mathbf{c}, \mathbf{c_0}, ..., \mathbf{c}_\alpha, ...)$$

to

$$\psi(\mathbf{c}) \wedge \lambda(\mathbf{x}, \mathbf{c}/\mathbf{y}) \vdash_{\mathbf{x}} \theta(\mathbf{x})$$

By the deduction theorem (Lemma 2.1.6), since $\xi(\mathbf{c}, \mathbf{c_0}, ..., \mathbf{c}_\alpha, ...)$ is a sentence, we obtain in $\mathbb{T}^m$ a derivation of:

$$\xi(\mathbf{c}, \mathbf{c_0}, ..., \mathbf{c}_\alpha, ...) \wedge \psi(\mathbf{c}) \wedge \lambda(\mathbf{x}, \mathbf{c}/\mathbf{y}) \vdash_{\mathbf{x}} \theta(\mathbf{x})$$

But it is always possible to replace the constants by variables as long as they are added to the contexts of the sequents, so using the existential rule, we have also a derivation of:

$$\exists \mathbf{x_0} ... \mathbf{x}_\alpha ... \xi(\mathbf{y}, \mathbf{x_0}, ..., \mathbf{x}_\alpha, ...) \wedge \psi(\mathbf{y}) \wedge \lambda(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{xy}} \theta(\mathbf{x})$$

Calling $Y' = [\mathbf{y}, \Phi(\mathbf{y})]$ the subobject of B given by the interpretation in $\mathcal{C}$ of the formula:

$$\exists \mathbf{x_0} ... \mathbf{x}_\alpha ... \xi(\mathbf{y}, \mathbf{x_0}, ..., \mathbf{x}_\alpha, ...) \wedge \psi(\mathbf{y})$$

we have a proof of the sequent:

$$\Phi(\mathbf{y}) \wedge \lambda(\mathbf{x}, \mathbf{y}) \vdash_{\mathbf{xy}} \theta(\mathbf{x})$$

and hence also of the sequent:

$$\exists \mathbf{y}(\Phi(\mathbf{y}) \wedge \lambda(\mathbf{x}, \mathbf{y})) \vdash_{\mathbf{x}} \theta(\mathbf{x})$$

Now the antecedent is precisely the pullback of the subobject $\Phi(\mathbf{y})$ of B along f, so by adjunction we have $Y' \leq \forall_f(C) = [\mathbf{y}, \gamma]$, i.e., the sequent $\Phi(\mathbf{y}) \vdash_{\mathbf{y}} \gamma(\mathbf{y})$ is provable. Therefore, since $M \vDash \Phi(\mathbf{c})$, it follows that $M \vDash \gamma(\mathbf{c})$, as we wanted to prove.

□

Once more, it is possible to remove the restriction on the cardinality of the signature in Theorem 6.2.1 by assuming κ to be strongly compact:

Theorem 6.2.2. *If κ is a strongly compact cardinal, every κ-first-order theory has a universal Kripke model.*

Proof. It suffices to rewrite the proof of Theorem 6.2.1 using Proposition 6.1.4 in place of Proposition 6.1.2. □

As applications of these completeness theorems, we obtain not only the disjunction property over a language without function symbols (with a proof similar to 5.2.6) but also the existence property:

Corollary 6.2.3. *If κ is a weakly compact cardinal, κ-first-order intuitionistic logic $\mathcal{L}_\kappa$ over a language without function symbols and with at least one constant symbol has the infinitary existence property. That is, if $\top \vdash \exists \mathbf{x}\phi(\mathbf{x})$ is provable in the empty theory, then, for some constants $\mathbf{c}$, $\top \vdash \phi(\mathbf{c})$ is already provable.*

Proof. This is a straightforward generalization of the usual semantic proof in the finitary case, based on the completeness with respect to Kripke models over trees. If no sequent $\top \vdash \phi(\mathbf{c})$ was provable, there would be a countermodel for each choice of $\mathbf{c}$. Then we can build a new Kripke tree appending to these countermodels a bottom node forcing no atoms, whose underlying domain contains just the constants of the language, with the obvious injections into the roots of the countermodels. Such a Kripke tree would then be a countermodel for $\top \vdash \exists \mathbf{x}\phi(\mathbf{x})$. □

Having at hand a completeness theorem for κ-$\mathcal{R}eg_\perp$ theories, it is possible to adapt Joyal's proof of the completeness theorem for first-order intuitionistic logic over infinitary logic, by replacing coherent models with κ-regular models. This has as a consequence, like in the finitary pre-Heyting theories, a completeness result for κ-pre-Heyting theories.

Theorem 6.2.4. *(ZFC) If κ is inaccessible, every κ-pre-Heyting theory of cardinality at most κ has a universal Kripke model.*

Proof. We go along the lines of Joyal's proof, adapted for this case using Proposition 6.1.6. Consider the category $\mathcal{R}eg(\mathcal{C})$ of κ-regular models of size κ of our syntactic category (i.e., κ-regular functors from the κ-$\mathcal{R}eg_\perp$ Morleyization of our theory), with the corresponding κ-regular homomorphisms given by natural transformations. We have a functor $ev : \mathcal{C} \to \mathbf{Set}^{\mathcal{R}eg(\mathcal{C})}$ sending an object A to the evaluation functor $ev(A)$. It is clear that this functor is κ-regular, and by Proposition 6.1.6, it is also conservative. The proof that ev also preserves $\forall$ follows the same lines as the proof of the infinitary version of Joyal's theorem, where we replace the category of κ-coherent models by that of κ-regular models. Although some of these models could be exploding, generating thus exploding nodes in the Kripke model, we are working classically, so that we can simply eliminate all exploding nodes without affecting the forcing at the rest of the nodes. □

Remark 6.2.5. It is possible to weaken the hypothesis of Theorem 6.2.4 by allowing κ to be any regular cardinal and the theories to have arbitrary cardinality. Indeed, in [Mak90] Makkai presents a syntactic proof that κ-regular theories are complete for **Set**-valued models, from which we can derive the completeness of κ-$\mathcal{R}eg_{\perp}$ theories by Morleyization. From this we proceed as in Joyal's theorem, which makes no use of the stronger hypothesis.

Finally, we can also see that the statement of completeness of κ-first-order theories requires in an essential way the hypothesis of weak compactness:

Proposition 6.2.6. *(ZFC) Given an inaccessible cardinal κ, the Kripke (resp. Beth, Tarski) completeness of κ-first-order theories of cardinality at most κ implies that κ is weakly compact.*

Proof. This is a generalization of Corollary 4.3.8 and is proved in an analogous way. Beth completeness and Kripke completeness imply κ-coherent completeness of theories of cardinality at most κ. To prove that this latter implies weak compactness, given a tree of height κ and levels of size less than κ, consider the theory of a branch, over a language containing a unary relation symbol P and one constant a for every node in the tree and axiomatized as follows:

$$\top \vdash \bigvee_{a \in L_\alpha} P(a)$$

for each $\alpha < \kappa$, where L_α is the level of height α;

$$P(a) \wedge P(b) \vdash \perp$$

for each pair $a \neq b \in L_\alpha$ and each $\alpha < \kappa$;

$$P(a) \vdash P(b)$$

for each pair a, b such that a is a successor of b.

Then the theory is certainly consistent within $\mathcal{L}_{\kappa,\kappa}$, as every subtheory of cardinality less than κ has a Tarski model, so by completeness it follows that the whole theory has a Tarski model, corresponding to a cofinal branch. $\square$

Also, we have as well:

Proposition 6.2.7. *(ZFC) Given an inaccessible cardinal κ, the Kripke (resp. Beth, Tarski) completeness of κ-first-order theories implies that κ is strongly compact.*

Proof. This is another generalization of Corollary 4.3.8. It is enough to show that κ-coherent completeness implies strong compactness. Consider a κ-complete filter $\mathcal{F}$ in the lattice $\mathcal{L}$ of subsets of a set. In a language containing a constant a for every $a \in \mathcal{L}/\mathcal{F}$ and a unary relation symbol P, consider the following theory of a κ-complete ultrafilter:

1. $P(a) \vdash P(b)$ for every pair $a \leq b$ in $\mathcal{L}$
2. $\bigwedge_{i<\gamma} P(a_i) \vdash P(\bigwedge_{i<\gamma} a_i)$ for all families $\{a_i\}_{i<\gamma}$ such that $\gamma < \kappa$
3. $\top \vdash P(a) \vee P(\neg a)$ for every $a \in \mathcal{L}$

Since the theory is consistent (as it has a model in $\mathcal{L}/\mathcal{F}$ itself), by κ-coherent completeness it has a Tarski model, which provides a κ-complete ultrafilter in $\mathcal{L}/\mathcal{F}$ whose preimage along the quotient map yields a κ-complete ultrafilter in $\mathcal{L}$ extending $\mathcal{F}$. Therefore, κ is strongly compact. □

Part V

Conclusion

7. Conclusion

7.1 Future work

There are several lines of further exploration of the ideas here presented that can be interesting to pursue. The first is the question about conceptual completeness, or, more generally, the question about to which extent the category of κ-coherent models of a certain theory characterizes the theory, and whether some κ-coherent form of the κ-pretopos completion (see [MR77]) plays some rôle in it. Positive answers in the case $\kappa = \omega$ have been given, and naturally one would like to determine if such characterizations could hold in the general case.

The second question is the determination of the right consistency strength of κ that is needed to prove the completeness theorem for theories of cardinality strictly smaller than κ. More especifically, call κ a *Heyting* cardinal if it is inaccessible and κ-first-order theories are complete (with respect to Kripke semantics) for theories of cardinality strictly less than κ. Such a cardinal lies somewhere between inaccessible and weakly compact cardinals, and for classical theories is known that inaccessibility is enough. We leave for the future the determination, in the intuitionistic case, of the exact consistency strength of this property within the large cardinal hierarchy.

Finally, it remains to determine whether theories over finite quantifier languages $\mathcal{L}_{\kappa,\omega}$ satisfy a completeness theorem. The classical case has been handled in [Gre75], but the intuitionistic case seems as usual more difficult to analyze.

7.2 Philosophical (in)completeness

There is a different conception of completeness of a more philosophical flavour, concerning how much we are able to grasp through completeness theorems the universe of sets. Despite the many advocates of ZFC plus large cardinal axioms as the preferred axiomatic system, constructive mathematicians are more confortable with IZF or its predicative variant, CZF (constructive Zermelo-Fraenkel). We present here a third proposal that combines aspects of those two: κ-first-order logic axiomatized with purely logical axioms, in the metatheory ZFC plus the axiom "κ is weakly compact". This has the effect of sweeping the law of excluded middle to the metatheory, thereby guaranteeing desirable properties of the object theory (like the disjunction and the existence property), equipped with a semantics for which a completeness theorem holds. But the ultimate thesis here is that beyond all these proposals, the philosophical completeness remains, without remedy, utopic, and that this has to be the case due to the very nature of existence: an amphibious, dynamic and incomplete path to being from nothingness.

We have tried to give instances of what seems to be the general phenomenon where this systematic incompleteness manifests: the more expressive the logic is, the stronger the metatheory must be to support a proof of completeness. In the same way that, quite probably, the current inconsistencies found at the very top of the large cardinal hierarchy could be atributed to some of the axioms of the metatheory, also the search for the ideal completeness theorem might be doomed to fail: we would presumably only be able to achieve a very strong form of completeness of some infinitary intuitionistic logic by using an inconsistent metatheory. In this sense, although we have only presented completeness theorems here, the remaining scent that these results should leave behind in our attempt to achieve completeness is that of an inherent incompleteness. Perhaps frustrating, but necessary to avoid the calamity of inconsistency. *Vi får inte veta. Vi kommer aldrig att veta.*

References

[AR01] Peter Aczel and Michael Rathjen, *Notes on constructive set theory*, 2001. 60

[BJ98] Carsten Butz and Peter Johnstone, *Classifying toposes for first-order theories*, Annals of Pure and Applied Logic **91** (1998), no. 1, 33–58. 48, 49

[CLR01] Michel Coste, Henri Lombardi, and Marie-Françoise Roy, *Dynamical method in algebra: effective Nullstellensätze*, Annals of Pure and Applied Logic **111** (2001), no. 3, 203–256. 74

[dJ80] Dick de Jongh, *A class of intuitionistic connectives*, Studies in Logic and the Foundations of Mathematics **101** (1980), 103–111. 79

[Dra74] Frank Drake, *Set theory. An introduction to large cardinals*, North-Holland Publishing Company, 1974. 79, 84

[Esp] Christian Espíndola, *Métodos categóricos en teoría de modelos*, Tesis de licenciatura (2011). 41

[Esp13] ———, *A short proof of Glivenko theorems for intermediate predicate logics*, Archive for mathematical logic **52** (2013), no. 7-8, 823–826. 20

[Esp16] ———, *Semantic completeness of first-order theories in constructive reverse mathematics*, Notre Dame Journal of Formal Logic **57** (2016), no. 2, 281–286. 20, 59

[FE] Henrik Forssell and Christian Espíndola, *Constructive completeness and non discrete languages*, Unpublished. 20, 58, 61, 73

[FS90] Peter J Freyd and Andre Scedrov, *Categories, allegories*, vol. 39, Elsevier, 1990. 52

[Gre75] Judy Green, *Consistency properties for finite quantifier languages*, Infinitary Logic: In Memoriam Carol Karp, Springer, 1975, pp. 73–123. 32, 103

[Jec73] Thomas Jech, *The axiom of choice*, North-Holland Publishing Company, 1973. 59

[Jec03] ______, *Set theory. The Third Millenium Edition, revised and expanded*, Springer Monographs in Mathematics. Springer-Verlag, Berlin (2003). 80

[Joh02] Peter Johnstone, *Sketches of an elephant (volume 2). A topos theory compendium*, Oxford University Press, New York, 2002. 26, 32, 33, 36, 41, 46, 57, 73, 74

[Kan09] Akihiro Kanamori, *The higher infinite: large cardinals in set theory from their beginnings*, Springer-Verlag, 2009. 79

[Kar64] Carol Karp, *Languages with expressions of infinite length*, North-Holland Publishing Company, 1964. 19, 23, 24, 25, 31, 79, 91

[Kre62] Georg Kreisel, *On weak completeness of intuitionistic predicate logic*, The Journal of Symbolic Logic **27** (1962), no. 02, 139–158. 57

[Kri96] Jean-Louis Krivine, *Une preuve formelle et intuitionniste du théoreme de complétude de la logique classique*, Bulletin of Symbolic Logic **2** (1996), no. 04, 405–421. 57

[Mak90] Michael Makkai, *A theorem on Barr-exact categories, with an infinitary generalization*, Annals of pure and applied logic **47** (1990), no. 3, 225–268. 25, 33, 34, 38, 41, 98

[McC94] David McCarty, *On theorems of Gödel and Kreisel: completeness and Markov's principle*, Notre Dame Journal of Formal Logic **35** (1994), no. 1, 99–107. 57

[McC08] ______, *Completeness and incompleteness for intuitionistic logic*, The Journal of Symbolic Logic **73** (2008), no. 04, 1315–1327. 57

[MM94] Saunders Maclane and Ieke Moerdijk, *Sheaves in geometry and logic*, Springer-Verlag, New York, 1994. 52, 73

[MR77] Michael Makkai and Gonzalo Reyes, *First-order categorical logic. Model-theoretical methods in the theory of topoi and related categories*, Springer, 1977. 18, 49, 58, 60, 103

[MT61] Shoji Maehara and Gaisi Takeuti, *A formal system of first-order predicate calculus with infinitely long expressions*, Journal of the Mathematical Society of Japan **13** (1961), no. 4, 357–370. 23

[Myh73] John Myhill, *Some properties of intuitionistic Zermelo-Frankel set theory*, Cambridge Summer School in Mathematical Logic, Springer, 1973, pp. 206–231. 58

[Nad78] Mark Nadel, *Infinitary intuitionistic logic from a classical point of view*, Annals of Mathematical Logic **14** (1978), no. 2, 159–191. 25

[Pal97] Erik Palmgren, *Constructive sheaf semantics*, Mathematical Logic Quarterly **43** (1997), no. 3, 321–327. 58, 60, 62

[TVD88] Anne Troelstra and Dirk Van Dalen, *Constructivism in mathematics*, vol. 2, Elsevier Science Publishers, 1988. 60, 61

[Vel76] Wim Veldman, *An intuitionistic completeness theorem for intuitionistic predicate logic*, The Journal of Symbolic Logic **41** (1976), no. 1, 159–166. 57, 58, 74

Matematiska Institutionen, Stockholms Universitet, 106 91 Stockholm
E-mail address: espindola@math.su.se

STOCKHOLM DISSERTATIONS IN MATHEMATICS

Dissertations at the Department of Mathematics Stockholm University

1-22. 2005-2008.

23. Strohmayer, Henrik: *Prop profiles of compatible Poisson and Nijenhuis structures.* 2009.
24. Rökaeus, Karl: *Grothendieck rings and motivic integration.* 2009.
25. Nilsson, Lisa: *Amoebas, discriminants, and hypergeometric functions.* 2009.
26. Lundqvist, Samuel: *Computational algorithms for algebras.* 2009.
27. Lindell, Andreas: *Theoretical and practical applications of probability.* 2009.
28. Geli Rolfhamre, Patricia: *From penicillin binding proteins to community interventions: mathematical and statistical models related to antibiotic resistance.* 2009.
29. Zusmanovich, Pasha: *Low-dimensional cohomology of current Lie algebras.* 2010.
30. Bojarova, Jelena: *Toward sequential data assimilation for numerical weather prediction (NWP) models using Kalman filter tools.* 2010.
31. Granåker, Johan: *Wheeled operads in algebra, geometry, and quantization.* 2010.
32. Xantcha, Qimh: *The theory of polynomial functors.* 2010.
33. Emtander, Eric: *Chordal and complete structures in combinatorics and commutative algebra.* 2011.
34. Grünewald, Maria: *Design and analysis of response selective samples in observational studies.* 2011.
35. Samieinia, Shiva: *Digital geometry, combinatorics, and discrete optimization.* 2011.
36. Björkwall, Susanna: *Stochastic claims reserving in non-life insurance: Bootstrap and smoothing models.* 2011.
37. Waliullah, Shoyeb: *Topics in nonlinear elliptic differential equations.* 2011.
38. Andersson, Patrik: *Four applications of stochastic processes: Contagious disease, credit risk, gambling and bond portfolios.* 2011.
39. Tadesse, Yohannes: *Tangential derivations, Hilbert series and modules over Lie algebroids.* 2011.
40. Al Abdulaali, Ahmad Khalid: *On the extension and wedge product of positive currents.* 2012.
41. Lindenstrand, David: *Stochastic epidemic models in heterogeneous communities.* 2012.
42. Lundqvist, Johannes: *On amoebas and multidimensional residues.* 2013.
43. Jost, Christine: *Topics in computational algebraic geometry and deformation quantization.* 2013.
44. Alexandersson, Per: *Combinatorial methods in complex analysis.* 2013.
45. Höhna, Sebastian: *Bayesian phylogenetic inference: Estimating diversification rates from reconstructed phylogenies.* 2013.
46. Aermark, Lior: *Hardy and spectral inequalities for a class of partial differential operators.* 2014.
47. Alm, Johan: *Universal algebraic structures on polyvector fields.* 2014.
48. Jafari-Mamaghani, Mehrdad: *A treatise on measuring Wiener–Granger causality.* 2014.
49. Martino, Ivan: *Ekedahl invariants, Veronese modules and linear recurrence varieties.* 2014.
50. Ekheden, Erland: *Catastrophe, ruin and death: Some perspectives on insurance mathematics.* 2014.
51. Johansson, Petter: *On the topology of the coamoeba.* 2014.

52. Lopes, Fabio: *Spatial marriage problems and epidemics.* 2014.
53. Bergh, Daniel: *Destackification and motivic classes of stacks.* 2014.
54. Olsson, Fredrik: *Inbreeding, effective population sizes and genetic differentiation: A mathematical analysis of structured populations.* 2015.
55. Forsgård, Jens: *Tropical aspects of real polynomials and hypergeometric functions.* 2015.
56. Tveiten, Ketil: *Period integrals and other direct images of D-modules.* 2015.
57. Leander, Madeleine: *Combinatorics of stable polynomials and correlation inequalities.* 2016.
58. Petersson, Mikael: *Perturbed discrete time stochastic models.* 2016.
59. Oneto, Alessandro: *Waring-type problems for polynomials: Algebra meets Geometry.* 2016.
60. Malmros, Jens: *Studies in respondent-driven sampling: Directed networks, epidemics, and random walks.* 2016.
61. Espíndola, Christian: *Achieving completeness: from constructive set theory to large cardinals.* 2016

Distributor
Department of Mathematics
Stockholm University
SE-106 91 Stockholm, Sweden