



A proof of Shelah's eventual categoricity conjecture and an extension to accessible categories with directed colimits

Christian Espíndola

► To cite this version:

Christian Espíndola. A proof of Shelah's eventual categoricity conjecture and an extension to accessible categories with directed colimits. 2023. hal-04353199

HAL Id: hal-04353199

<https://hal.science/hal-04353199>

Preprint submitted on 19 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A proof of Shelah’s eventual categoricity conjecture and an extension to accessible categories with directed colimits

Christian Espíndola

Abstract

We provide a proof, in *ZFC*, of Shelah’s eventual categoricity conjecture for abstract elementary classes (AEC’s). Moreover, assuming in addition the Singular Cardinal Hypothesis (*SCH*), we prove a direct generalization to the more general context of accessible categories with directed colimits. If \mathcal{K} is such a category, we show that there is a cardinal μ such that if \mathcal{K} is λ -categorical for some $\lambda \geq \mu$ (i.e., it has only one object of internal size λ up to isomorphism), then \mathcal{K} is eventually categorical (i.e., it is λ' -categorical for every $\lambda' \geq \mu$). When considering cardinalities of models of infinitary theories \mathbb{T} of $\mathcal{L}_{\kappa,\theta}$ that axiomatize \mathcal{K} , the result implies, under *SCH*, the following infinitary version of Morley’s categoricity theorem: let S be the class of cardinals λ which are of cofinality at least θ but are not successors of cardinals of cofinality less than θ . Then, if \mathbb{T} is a $\mathcal{L}_{\kappa,\theta}$ theory whose models have directed colimits and it is λ -categorical for some $\lambda \geq \mu$ in S , then it is λ' -categorical for every $\lambda' \geq \mu$ in S ; moreover, we also exhibit an example that shows that the exceptions in the class S are needed. Along the way we also prove Grossberg conjecture, according to which categoricity in a high enough cardinal implies eventual amalgamation. We establish this result in AEC’s and, assuming in addition *SCH*, in the more general context of accessible categories whose morphisms are monomorphisms.¹

1 Introduction

This paper belongs to the classification theory program initiated by Shelah, intended to determine dividing lines in the behaviour of the categoricity spectrum of theories. Continuing the work of Morley, who established that categoricity of a countable finitary theory in one uncountable cardinal entailed categoricity in all uncountable cardinals, Shelah intended to generalize this setup, to uncountable finitary theories and even to finite quantifier infinitary theories. He later pushed the classification to more general infinitary theories introducing abstract elementary classes. The main open problem in the area, Shelah’s eventual categoricity conjecture, asserts that for any abstract elementary class (AEC) there is a cardinal κ such that if the AEC is categorical in some $\lambda \geq \kappa$, it is categorical in *all* $\lambda' \geq \kappa$. This general conjecture was stated in [She09], while the version for the particular case of sentences in $\mathcal{L}_{\omega_1,\omega}$ was conjectured around 1977. Several approximations are known (see, e.g., the account presented in the introduction of [Vas18]), though the full conjecture has been open until now. When there is a proper class of strongly compact cardinals or when amalgamation and a suitable weakening of the Generalized Continuum Hypothesis (*GCH*) hold, the conjecture was proven to be true in [SS18]. We will provide here a full proof of the conjecture, in *ZFC*, and show as well that eventual amalgamation in fact follows from categoricity in a high enough cardinal, thus proving a conjecture of Grossberg from 1986. Moreover, our topos-theoretic argument is general enough so that, assuming in addition *SCH*, it applies to accessible categories with directed colimits, establishing

¹This research has been partially supported through the grant 19-00902S from the Grant Agency of the Czech Republic.

these conjectures in a much wider variety of categories of models. This also provides a version of Morley’s categoricity theorem for the corresponding infinitary theories \mathbb{T} in $\mathcal{L}_{\kappa,\theta}$ whose models have directed colimits; if S be the class of cardinals λ which are of cofinality at least θ but are not successors of cardinals of cofinality less than θ then, under SCH , when \mathbb{T} is λ -categorical for some $\lambda \geq \mu$ in S , it is λ' -categorical for every $\lambda' \geq \mu$ in S . The case $\theta = \omega$ is (a generalization of) Shelah’s original conjecture, since in that case the use of SCH can be eliminated²; also, the threshold cardinal we find is not explicit, but it is provably the least possible for which the conjecture is true. We prove as well that this infinitary version of Morley’s theorem is sharp in the sense that we build examples of theories (see Example 9.5) where the cardinals λ which are not in the class S fail in fact to have categoricity in the cardinality of the model, making our result best possible.

The main tool that we will use for these results is a completeness theorem for infinitary intuitionistic logics developed by the author in [Esp19] and [Esp20]. These are theorems valid for theories in infinite quantifier languages, for which a new sound rule of inference has been devised. Even when restricting to the particular case of classical theories, these completeness theorems are sharper than known completeness theorems for infinitary classical logic (e.g., those developed by Karp in [Kar64]), as unlike the latter, the former are able to be cast in the setting of categorical logic, making use of an infinitary generalization of syntactic categories that begun to be explored in the ’90’s for the infinitary regular fragment. As we will see, the categorical language is essential in this paper since the crucial arguments are topos-theoretic in nature. What lies below is the fact that any abstract elementary class is equivalent to the category of κ -points of some κ -topos (introduced in [Esp20], where they are called κ -geometric toposes). It is this precise observation which will allow for our generalization to accessible categories with directed colimits.

Among the consequences of the completeness theorem that we will use, the omitting types theorem for infinite quantifier languages stands out. Whenever the category of models of an infinitary theory over an infinite quantifier language has directed colimits, an analogue of the well-known omitting types theorem for finitary logic (or $\mathcal{L}_{\omega_1,\omega}$) can be derived as a consequence of our completeness theorem. It was realized in the ’70’s that the omitting types theorem was an essential tool in model theory, having consequences for atomic and prime models, among other applications; this will be recast in our case to study infinitary theories which are categorical in some cardinal. The last ingredient of our proof is a precise characterisation of the κ -classifying topos for κ -saturated models, which will be very useful for deriving a downward categoricity transfer.

Our approach avoids large cardinals or amalgamation since these are essentially, as we will see, compactness properties of the theories in question, and we will be able to derive them directly from a categoricity assumption by using instead an approach that relies on infinitary regular logic, introduced by Makkai in [Mak90]. Although categoricity is not enough to make our theory axiomatizable in infinitary regular logic, it can be axiomatized in a slightly bigger fragment, which can then be studied through an equivalent infinitary regular theory on a different signature, an instance of the process known as Morleyization. κ -regular logic is known to be κ -compact, as proven already by Makkai in [Mak90]. The reason is essentially that while compactness arguments usually involve Łoś theorem and the consequent use of κ -complete ultrafilters extending a given κ -complete filter, the absence of disjunctions in the language of κ -regular logic allows to derive the corresponding version of Łoś theorem by using only an appropriate κ -complete filter.

²In fact, the precise amount of SCH that is required to hold for infinitary theories \mathbb{T} in $\mathcal{L}_{\kappa,\theta}$ is $SCH_{\theta, \geq 2^{<\theta}}$, defined as “for all $\mu \geq 2^{<\theta}$ there is a set of cardinals $\lambda_i \leq \mu$ unbounded below μ such that, for each i , $\nu^{<\theta} \leq \lambda_i$ for all $\nu < \lambda_i$ ”. This is discussed in Remark 2.3 of [LRV19].

The existence of this latter does not need to invoke any large cardinal assumptions. This together with the topos-theoretic characterisation of saturated models provides enough compactness to prove Grossberg conjecture.

Although there is a vast literature covering attempts to prove Shelah’s eventual categoricity conjecture (there are more than 2000 pages of approximations), our approach differs from them all in that we have a completeness theorem available, which makes possible, as we said, an omitting types theorem and its consequences. Nevertheless there are several partial results that can be found already in the literature and are derived here with fundamentally different proofs. The closest to our arguments seems to be [SM90], where the strategy followed has many points in common with ours. We can point, though, the following improvements (besides the fact that we avoid large cardinals):

- Several known facts about AEC’s, like the properties described in our Lemma 6.6, Corollary 6.2, Lemma 6.4 or Theorem 8.2, are here proved in the more general context of accessible categories with directed colimits. This shows that the notion of internal size is the correct framework in which to state these results.
- The use of the monster model is avoided altogether by using instead the syntactic category of the corresponding theory, whose universal property is enough to derive many of the statements in [SM90] without having to go through the semantic detour.
- Our Theorem 9.1 and its proof should correspond roughly to Proposition 4.27 of [SM90] (as explained in our Remark 9.2), but the heavy use of stability theory is here avoided through an argument relying instead on Kripke models and the completeness of infinitary intuitionistic logic, which simplifies the proof considerably.

As we mentioned, when restricting to an AEC we recover a proof, in *ZFC*, of Shelah’s eventual categoricity conjecture. Moreover, the threshold cardinal is the maximum of the Hanf numbers³ for categoricity and non-categoricity, which is in fact the least possible cardinal for which the conjecture is true. We avoid *GCH*, which is spared through a forcing argument that adapts well in this particular proof, circumventing thus the fact that categoricity is not absolute. Indeed, almost all uses of *GCH* in our paper derive from the equality $\kappa^{<\kappa} = \kappa$ for regular cardinals κ . This is in contrast with the use of it in previous literature as a device that, like uses of the diamond principle, makes the set-theoretical universe better behaved. In our case, this particular cardinal arithmetic equality will be easy to derive through the use of forcing (collapsing $\kappa^{<\kappa}$ to κ), in a way that does not modify the essential properties that we are proving (as this forcing does not change the category of models of size less than κ). This is what in the end will allow to eliminate many uses of *GCH* in the case of AEC’s, and to downgrade them to mere *SCH* in the case of accessible categories with directed colimits. If we restrict to those accessible categories in which all morphisms are monomorphisms (in particular, if we work with AEC’s), there is also an alternative argument, which we explain in the last section, that allows to achieve this same goal without recurring to forcing, by using instead the Löwenheim-Skolem property. Thus, the reader who is not acquainted with forcing techniques can easily skip to the very last portion of the last section, in which we explain how the version of the completeness theorem that we use in our arguments does not need *GCH* (see Remark 10.2).

Our main theorem shows that a classification theory for infinite quantifier theories is indeed possible, at least when the category of models has directed colimits. Thus, our results open up an entire new path so far unexplored in the model theory of non-elementary classes.

³The Hanf number for a certain property is the least cardinal μ such that if the property holds at a certain $\lambda \geq \mu$, then it holds in a proper class of cardinals.

The paper is structured as follows. In the section of preliminaries we introduce the notion of internal size from [LRV19] and establish the existence spectrum of accessible categories with directed colimits with this notion of size (which coincides in AEC's with the size of the underlying set). We then present a completely categorical account of the completeness theorem for intuitionistic logics, along with the downward Löwenheim-Skolem theorem, which will be used for the downward categoricity transfer later. We include also the representation of a κ -classifying topos as a topos of equivariant sheaves. This latter is only needed for the proof of the generalized conjecture for accessible categories with directed colimits, but plays no rôle in the case of AEC's; thus the reader interested only in understanding the proof of Shelah's eventual categoricity conjecture for AEC's can disregard this material. We subsequently move to prove a technical result regarding κ -classifying toposes, introduced by the author in [Esp20] as a generalization of the usual notion of classifying topos that provides the adequate context to study infinitary logics; the result established in this section is essential for the rest of the paper. The next immediate application is the omitting types theorem for infinitary theories whose categories of models have directed colimits; we provide two formulations, a syntactic one and a semantic one which does not presuppose any proof system. This theorem is in turn then used in the next section to establish a topos-theoretic characterization of $\mathcal{L}_{\infty, \kappa}$ -elementary equivalence. Next, we give a precise description of the κ -classifying topos for κ -saturated models, and then complement with results on amalgamation, establishing Grossberg's conjecture in accessible categories whose morphisms are monomorphisms. Once we have amalgamation, we then present the downward categoricity transfer and the main result of the paper, namely the generalization of Shelah's eventual categoricity conjecture to accessible categories with directed colimits. We postpone for the last section the presentation of the arguments used to eliminate many of the uses of GCH made thus far, which are based on a common technique described in detail in one particular case. This means that most uses of GCH made before are not essential, but we prefer to prove this later to make the reading more fluid.

To make the paper as self-contained as possible for readers not familiar with the general results of the author in [Esp19] and with those in the sequel [Esp20], we will include a brief description of κ -coherent theories (a particular case of κ -geometric from [Esp20]), together with the notion of κ -classifying topos. In particular, we will mention the transfinite transitivity rule, which makes the completeness theorem possible.

2 Preliminaries

We start with the categorical definition of an abstract elementary class:

Definition 2.1. *An abstract elementary class (AEC) \mathcal{K} is a category equivalent to an accessible category with directed colimits whose morphisms are monomorphisms, that admits an embedding $F : \mathcal{K} \longrightarrow \mathcal{A}$ into a finitely accessible category preserving directed colimits and monomorphisms which is, in addition:*

1. *Full on isomorphisms: for every isomorphism $h : F(A) \longrightarrow F(B)$ there is an isomorphism $m : A \longrightarrow B$ with $F(m) = h$.*
2. *Nearly full: for every commutative triangle:*

$$\begin{array}{ccc} F(A) & \xrightarrow{F(f)} & F(B) \\ & \searrow h & \nearrow F(g) \\ & F(C) & \end{array}$$

there is $m : A \longrightarrow C$ with $F(m) = h$.

A result of [BR12] gives us that an AEC automatically admits an iso-full, nearly full embedding $E : \mathcal{K} \longrightarrow \text{Emb}(\Sigma)$ (for some signature Σ , where $\text{Emb}(\Sigma)$ is the category of Σ -structures and embeddings) preserving directed colimits. It also has eventually a Löwenheim-Skolem number λ : for every substructure $i : A \longrightarrow F(B)$ there is $E(f) : F(C) \longrightarrow F(B)$ such that i factors through $E(f)$ and $|E(C)| \leq |A| + \lambda$. This gives the connection with the standard model-theoretic definition of an AEC. As we remarked earlier, an important fact in AEC's is that for each object A , $|E(A)|$ coincides with its internal size $|A|$. This means that our eventual categoricity result for accessible categories with directed colimits is really a direct generalization of the same result for AEC's.

When in the definition of AEC we allow for concrete μ -directed colimits instead of directed, we recover the notion of μ -AEC. As proven in [BGL⁺16], this concept admits a much simpler categorical description:

Definition 2.2. *A μ -abstract elementary class (μ -AEC) \mathcal{K} is a category equivalent to an accessible category whose morphisms are monomorphisms.*

It is well-known that accessible categories are axiomatizable with sequents in κ -coherent logic (introduced in section 3.1) for some κ ; moreover, one can choose the signature to contain only a binary relation symbol. This changes the signature (in fact it introduces multiple sorts), but since in AEC's the size of a model (i.e. the cardinal sum of the underlying sets of each sort) coincides with its internal size, and this latter is a categorical invariant, the change of signature does not affect the categoricity spectrum; in fact, the category of models with the new signature will be equivalent to the original AEC. We will see soon that this will determine, up to equivalence, the κ -classifying topos. If one allows an infinitary signature, other axiomatizations are possible as well, as we will see later. Due to the equivalence between an accessible category and the category of models of its axiomatization, we will usually identify the object of the accessible category with the corresponding model, and loosely speak of a model and make reference to its underlying set when we deal with objects of the accessible category.

Throughout the rest of the paper we will deal with a large accessible category \mathcal{K} with directed colimits equivalent to the category of models of some $\mathcal{L}_{\kappa,\theta}$ -theory. We will work with the notion of internal size $|A|$ from [LRV19], defined as follows:

Definition 2.3. *If $r(A)$ is the least regular cardinal λ such that A is λ -presentable, then the internal size $|A|$ is defined as:*

$$|A| = \begin{cases} \kappa & \text{if } r(A) = \kappa^+ \\ r(A) & \text{if } r(A) \text{ is limit} \end{cases}$$

We will also assume the Singular Cardinal Hypothesis (*SCH*), which states that for every infinite singular cardinal λ , we have that $\lambda^{cf(\lambda)} = 2^{cf(\lambda)} + \lambda$. This is a weakening of the Generalized Continuum Hypothesis (*GCH*), which holds, e.g., above a strongly compact cardinal, and is still powerful enough to imply well-behaved cardinal arithmetic properties. For example, it will imply (see [Jec03]) that $\kappa^{<\theta} = \kappa$ for all $\kappa > 2^\theta$ of cofinality at least θ . This allows to apply the downward Löwenheim-Skolem theorem to find submodels of size exactly κ .

Assuming *SCH*, we can identify in an accessible category of models of some theory in $\mathcal{L}_{\kappa,\theta}$ for which cardinals the notion of internal size and that of external size (i.e., the cardinality of the underlying set of the model) coincide. Indeed, by Theorem 4.13 of [LRV19], that will be the case for all regular cardinals λ which are not successors of cardinals of cofinality less than θ . This motivates the following:

Definition 2.4. *S is the class of cardinals which are of cofinality at least θ but are not successors of cardinals of cofinality less than θ .*

By Theorem 4.12 of [LRV19] we have:

Remark 2.5. *Internal and external sizes coincide for all high enough $\lambda \in S$.*

By the results of Beke and Rosicky ([BR12]), any large accessible category with directed colimits is eventually λ -accessible for every λ . We will see in Lemma 6.3 below that under a categoricity assumption the subcategory where we consider only monomorphisms is closed under directed colimits above the categoricity cardinal. It follows from Corollary 4.15 of [LRV19] that in those conditions, and under *SCH*, there is an object of every high enough internal size whose cofinality is at least θ . The following lemma shows that more can be said:

Lemma 2.6. *Assume *SCH* and that all morphisms are monomorphisms. If κ is a singular cardinal of cofinality less than θ , models of internal size κ are precisely the directed colimits of models of internal size $\nu < \kappa$ in S .*

Proof. It is easy to see that directed colimits of models of internal size $\nu < \kappa$ in S must be κ^+ -presentable. On the other hand, by Löwenheim-Skolem theorem, for any subset A of size $\kappa = \bigcup_{\nu < \kappa, \nu \in S} \nu$ of a model M of internal size κ there is a submodel of internal size κ containing A . Indeed, writing A as a union of a chain of κ subsets of size less than κ , we can find submodels containing these subsets, that are of the same internal size and form moreover a chain under inclusion. Then the directed colimit of this chain will be the desired submodel. This argument also shows, as is not difficult to see, that M is a κ^+ -directed colimit of submodels which are themselves directed colimits of models of internal size less than κ . Since M is κ^+ -presentable, it follows that it must coincide with one of these submodels, as we wanted. \square

Corollary 2.7. *Assume *SCH*. Then in any large accessible category with directed colimits, categorical in some cardinal, there is eventually an object of every high enough internal size.*

Proof. It is enough to note that the internal size remains unchanged when one restricts to the subcategory where all morphisms are monomorphisms. \square

In the next sections we will assume that *SCH* holds, at least. \mathcal{K}_κ will denote the subcategory of objects of internal size κ and $\mathcal{K}_{\geq \kappa}$ that of objects of internal size at least κ .

3 The completeness theorem

3.1 κ -coherent logic

Let κ be a regular cardinal such that $\kappa^{<\kappa} = \kappa$. The syntax of κ -coherent logic 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 3.1. *If $\phi, \psi, \{\phi_\alpha : \alpha < \gamma\}$ (for each $\gamma < \kappa$) and $\{\psi_\alpha : \alpha < \delta\}$ (for each $\delta < \kappa^+$) are κ -coherent formulas, the following are also formulas: $\bigwedge_{\alpha < \gamma} \phi_\alpha$, $\exists_{\alpha < \gamma} x_\alpha \phi$ (also written $\exists_{\mathbf{x}_\gamma} \phi$ if $\mathbf{x}_\gamma = \{x_\alpha : \alpha < \gamma\}$) and $\bigvee_{\alpha < \delta} \psi_\alpha$, this latter provided that $\bigcup_{\alpha < \delta} FV(\psi_\alpha)$, the set of free variables of all ψ_α , has cardinality less than κ .*

We use sequent style calculus to formulate the axioms of κ -coherent logic, as explained in [Esp20], where the system for κ -coherent logic is described. The only substantial difference with usual geometric logic (which can be seen to be exactly ω -geometric logic) is the introduction of the transfinite transitivity rule (called “rule T ” in [Esp20]):

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

for each cardinal $\gamma \leq \kappa$, 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 . Here $B \subseteq \gamma^{<\kappa}$ consists of the minimal elements of a given bar⁴ over the tree $\gamma^{<\kappa}$, and the δ_f are the levels of the corresponding $f \in B$.

Basically, what it expresses is that given the tree $\gamma^{<\kappa}$ (i.e., the poset of functions $f : \beta \rightarrow \gamma$ for $\beta \leq \gamma$ with the order given by inclusion), and assuming there is an assignment of formulas ϕ_f to each node of the tree in such a way that the formula assigned to each node is “covered” by the formulas assigned to its immediate successors, and 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 must be “covered” by the formulas assigned to the nodes ranging among the minimal elements of a given bar over the tree. This rule is the syntactic counterpart of the corresponding exactness property of the category of sets, in which we identify each formula in a node with the set of elements where the formula holds, and where each “cover” is really a jointly epic family of functions.

A κ -coherent theory corresponds precisely to a κ -coherent category, which is a coherent category with κ -small limits, stable unions of at most κ many subobjects and satisfying that the transfinite composites of jointly epic families of arrows form a jointly epic family. Indeed, each such theory \mathbb{T} gives rise to its κ -coherent syntactic category $\mathcal{C}_{\mathbb{T}}$ containing a generic model M_0 of \mathbb{T} , characterized by the universal property that models M in any other κ -coherent category \mathcal{D} correspond precisely to functors $\mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{D}$ preserving the κ -coherent structure, i.e., to κ -coherent functors:

$$\begin{array}{ccc} \mathcal{C}_{\mathbb{T}} & \xrightarrow{\kappa\text{-coherent}} & \mathcal{D} \\ \uparrow & & \uparrow \\ M_0 & \vdash & M \end{array}$$

Presheaves inherit the transfinite transitivity property of **Set**, as do sheaves over a κ -coherent category with a κ -topology (i.e., a Grothendieck topology in which transfinite composites of covers are still covers). Such toposes will be here called κ -toposes (κ -geometric toposes in [Esp20]). The (2-)category of κ -toposes has as 1-cells the κ -geometric morphisms, which are geometric morphisms whose inverse image preserve, in addition, all κ -small limits.

Each κ -coherent theory admits a κ -classifying topos, introduced in [Esp20] as the κ -topos containing a generic model of the theory which is universal among models in other κ -toposes, in the sense that the diagram:

⁴A bar over the tree $\gamma^{<\kappa}$ is an upward closed subset of nodes intersecting every branch of the tree.

$$\begin{array}{ccc}
\mathcal{C}_{\mathbb{T}} & \xrightarrow{Y} & \mathbf{Set}[\mathbb{T}]_{\kappa} \\
& \searrow M & \swarrow \text{\scriptsize κ-small limit preserving} \\
& & \mathcal{E}
\end{array}
\quad \cong$$

where \mathcal{E} is a κ -topos and M is a κ -coherent functor from the syntactic category $\mathcal{C}_{\mathbb{T}}$, commutes up to invertible 2-cell for an essentially unique geometric morphism (the dashed inverse image). To simplify the notation, we write throughout $(\mathcal{C}_{\mathbb{T}})_{\lambda}$, with regular λ , for the syntactic category of the theory with the axioms of \mathbb{T} but in λ -coherent logic $\mathcal{L}_{\lambda^+, \lambda}$. When λ is limit, the we understand by the notation the colimit in \mathbf{Cat} of $(\mathcal{C}_{\mathbb{T}})_{\kappa}$ for regular $\kappa < \lambda$. Likewise, the λ -classifying topos of \mathbb{T} is written as $\mathbf{Set}[\mathbb{T}]_{\lambda}$, understanding that this latter is, for limit λ , defined as the pseudolimit of the toposes $\mathbf{Set}[\mathbb{T}]_{\kappa}$ for regular $\kappa < \lambda$.

Finally, we briefly mention the κ -regular fragment as defined in [Mak90], which is the subfragment of κ -coherent logic in which we remove disjunctions from the language. In this case, the transfinite transitivity rule is only applied to linear trees, where each node has just one successor, and thus reduces to the rule of dependent choices.

3.2 Completeness of κ -coherent logic

The main tool in the analysis of infinitary logic will be the completeness theorem proved by the author in [Esp19] and [Esp20]. Here we will present the same proof, but cast entirely in categorical language to better suit our purposes. In topos-theoretic language, it asserts that every κ -separable⁵ κ -topos has enough κ -points. The κ -separability condition puts a restriction on the cardinality of formulas of the internal language of the site, which is achieved by working in the so called κ -fragments.

In the classical case, a κ -fragment of $\mathcal{L}_{\kappa^+, \kappa}$ will be a subset of formulas formed in a language with a signature of cardinality at most κ and a supply of κ many fresh variables, that in addition is closed under κ -small conjunctions, disjunctions and quantification, negation and formal negation, subformulas and substitution. It follows that for any sentence ϕ of $\mathcal{L}_{\kappa^+, \kappa}$ there is a smallest κ -fragment containing ϕ , and it has κ many formulas. The same is true for any theory with at most κ many axioms.


An analogous definition for κ -coherent logic yields that the syntactic category of a theory with at most κ axioms has a subcategory given by the formulas in context belonging to the κ -fragment generated by the theory. This subcategory is itself a κ -coherent category of size κ , for which we will now prove a completeness theorem.

Given a κ -coherent category \mathcal{C} of size κ constructed as above, we will find a jointly conservative set of κ -coherent functors $F_i : \mathcal{C} \rightarrow \mathbf{Set}$. The definition of these functors is as follows. First, for each proper monomorphism $n : Y \rightarrow X$ in \mathcal{C} we consider the slice \mathcal{C}/X . In this slice, we define for each object B a set of κ many jointly epic families of at most κ many arrows each, in the following way. For each A in $\mathcal{C}' := \mathcal{C}/X$ we consider all jointly epic families of at most κ many arrows each generated, as a κ -Grothendieck topology, by the axioms of the theory. We then pullback these covers of A along the morphism $B \rightarrow 1$, which gives us covers of $A' := A \times B$. Now we consider all sections $s : B \rightarrow A'$ of the projection $\pi_B : A' \rightarrow B$, and we pullback the κ many covers over A' along each of these sections. This defines a sequence of κ many covers

⁵ A κ -separable κ -topos is one for which there is a site of size at most κ where the κ -Grothendieck topology is generated by at most κ many covering families.

over B , $\mathcal{F}(B)$, which without loss of generality we can assume well-ordered and of order type κ . We then build a tree of height κ each of whose nodes is an object of \mathcal{C}' such that its immediate successors are the domains of a certain jointly epic family of arrows over the object. At step 0, we start with the terminal object in \mathcal{C}' . We fix a well-ordering $w : \kappa \times \kappa \rightarrow \kappa$ such that $w(\alpha, \beta) \geq \beta$ (e.g., the canonical well-ordering of $\kappa \times \kappa$). At step $\gamma + 1$ we consider $w^{-1}(\gamma) = (\alpha, \beta)$ and we pull back the α -cover in each $\mathcal{F}(m)$, for m a node in the β -level of the tree, along all the branches of height γ over each such node m , to form the objects in the level $\gamma + 1$ of the tree. At limit ordinals δ we take the limit of the chain of objects so defined along each branch of height δ . This defines a tree of height κ whose branches are transfinite chains of morphisms $b : 1 \leftarrow B_0 \leftarrow B_1 \leftarrow \dots$. To finish the definition of the jointly conservative family of functors, we compute the corresponding slices $\mathcal{C}' \rightarrow \mathcal{C}'/B_0 \rightarrow \mathcal{C}'/B_1 \rightarrow \dots$ with (some choice of) pullback functors between them, and take their (pseudo-)colimit in \mathcal{Cat} , \mathcal{C}'_b . Our family of functors will then be the composites $\mathcal{C} \rightarrow \mathcal{C}' \rightarrow \mathcal{C}'_b \rightarrow \mathbf{Set}$, where we vary over all proper monomorphisms $n : Y \rightarrow X$ and all branches b , and where each $\mathcal{C}'_b \rightarrow \mathbf{Set}$ is given by the term model functor $[1, -]$.

We need to prove that each functor is κ -coherent, for which it is enough to prove that each one of these composites $[1, -] \circ F$ preserves jointly epic families of at most κ many arrows. Suppose $\{C_i \twoheadrightarrow M_i \twoheadrightarrow A\}_{i < \kappa}$ is such a family and we have in \mathcal{C}'_b an arrow $s : 1 \rightarrow \bigvee_{i < \kappa} F(M_i) = F(A)$, which is a section of the map $F(A) \rightarrow 1$. By construction, such an arrow has a representative $s_B : B \rightarrow A'$ in some slice \mathcal{C}'/B which is a section of a map $h : A' \rightarrow B$, so that, by definition, the image $\{C'_i \twoheadrightarrow M'_i \twoheadrightarrow A'\}_{i < \kappa}$ in \mathcal{C}'/B of the family $\{C_i \twoheadrightarrow M_i \twoheadrightarrow A\}_{i < \kappa}$, when pulled back along the section $s_B : B \rightarrow A'$, provides a jointly epic family $\{P_i \twoheadrightarrow B\}_{i < \kappa}$ that belongs to $\mathcal{F}(B)$. Such a jointly epic family is then, in turn, pulled back at some step in the transfinite construction to form the successors of some node C , so that the branch b chosen for our functor F will have some object $P'_i := P_i \times_B C$ as a successor of C . It follows that in the slice \mathcal{C}'/P'_i the section $s_{P'_i} : P'_i \rightarrow Q$, which represents s , will factor through $i : P \rightarrow Q$ via the induced morphism $(f, Id_{P'_i}) : P'_i \rightarrow P$ to the pullback P (which is the image of C_i in the slice \mathcal{C}'/P'_i). That this is the case can be seen by noticing that $s_{P'_i} : P'_i \rightarrow Q$ and $i \circ (f, Id_{P'_i}) : P'_i \rightarrow Q$ give the same morphism when composed with a and b (see the diagram below).



Therefore, in \mathcal{C}'_b , $s : 1 \rightarrow F(\bigvee_{i < \kappa} M_i) = F(A)$ factors through some $F(C_i)$, as we wanted to prove.

Finally, it should be clear now that this family of κ -coherent functors is jointly conservative, since for every proper monomorphism $m : Y \rightarrow X$, its image in the slice $\mathcal{C}' := \mathcal{C}/X$ is a proper subterminal object. If all functors in our family sent this to an isomorphism, it would be an

isomorphism in each \mathcal{C}'_b and hence also in some slice \mathcal{C}'/B_b for some B_b in each branch b . Now the transfinite transitivity property implies that the family of functors $\{\mathcal{C}' \rightarrow \mathcal{C}'/B_b\}_b$ is jointly conservative, getting that the original monomorphism was not proper after all. Whence, the family really is jointly conservative, as we wanted.

Note that the cardinal arithmetic assumption on κ that we need to put is $\kappa^{<\kappa} = \kappa$ (in particular, κ must be regular). This guarantees that the set of κ -coherent formulas in a κ -fragment containing any set of κ many of them will have cardinality κ , so that each set of covers $\mathcal{F}(B)$ has size κ and we can perform the transfinite construction up to κ . There is however the possibility to extend the theorem even if κ is singular, provided the fragment in which the axioms of our theory are expressed is the union of all ν -fragments of size ν for regular $\nu < \kappa$ and that we restrict to those morphisms which are monomorphisms. This happens because then the well-ordering of $\kappa \times \kappa$ can be chosen to be the union of the canonical well-orderings of each $\nu \times \nu$. This observation will be used later on when we prove the downward categoricity transfer.

3.3 The Löwenheim-Skolem theorem

Recall that the universal property of the slice category (which has 2-categorical aspects) says that the pullback functor to the slice is the universal morphism into a κ -coherent category which has a section (the diagonal $\Delta : (\mathbf{x}, \phi) \rightarrow (\mathbf{x}, \phi) \times (\mathbf{x}, \phi)$) for the morphism $(\mathbf{x}, \phi) \rightarrow 1$:

$$\begin{array}{ccc}
 \mathcal{C}_{\mathbb{T}} & \xrightarrow{\quad} & \mathcal{C}_{\mathbb{T}}/(\mathbf{x}, \phi), \Delta \\
 & \searrow F & \swarrow \bar{F} \\
 & \mathcal{D}, \mathbf{c} &
 \end{array}
 \quad \begin{array}{c}
 \cong \\
 \bar{F}(\Delta) = \mathbf{c}
 \end{array}$$

This universal property allows us to have an entirely categorical understanding of the completeness theorem just proven, since starting with the syntactic category of a κ -coherent theory \mathbb{T} , each slice over (\mathbf{x}, ϕ) is equivalent to the syntactic category of the theory $\mathbb{T} \cup \phi(\mathbf{c})$ where \mathbf{c} is a tuple of fresh constants.

Since the construction explained in the proof of the completeness theorem proceeds by successively taking slices over the syntactic category \mathcal{C} of \mathbb{T} , an easy application of the universal property shows that for any κ -coherent model $M : \mathcal{C} \rightarrow \mathbf{Set}$ and any tuple of less than κ elements $\mathbf{c} \subseteq M$ (where we identify the functor with the underlying set of $M((x, \top))$), factors, up to invertible 2-cell, through some \mathcal{C}'_b for some branch b . Indeed, we know it factors through the slice $\mathcal{C}' := \mathcal{C}/(\mathbf{x}, \top)$; since M is κ -coherent, it sends the tree of height κ built by transfinite recursion to a tree of sets each of whose nodes has a jointly epic family of morphisms from its immediate successors. This allows to define successively a compatible chain of elements $\mathbf{d}_B \in M(B)$ which, by the universal property of the slice, induce morphisms $M_B : \mathcal{C}'/B \rightarrow \mathbf{Set}$ along objects B in some branch b , with invertible 2-cells between them that form a pseudococone. It follows, as we claimed, that then there is an induced κ -coherent functor $\bar{M} : \mathcal{C}'_b \rightarrow \mathbf{Set}$ that extends M . Since M must be a colimit of representable functors in whose diagram there is always $[1, -]$, this latter representable provides, as we have proven above, a κ -coherent submodel M' of M containing the tuple \mathbf{c} . In other words, we have proven a version of the downward Löwenheim-Skolem theorem which will become essential for the downward categoricity transfer, as we will see.

Note that the submodel $[1, -] \circ F$ is a κ -coherent model of each theory $\mathbb{T} \cup \phi(\mathbf{c})$, where $F((\mathbf{x}, \phi)) = B_i$ is the i -th object in the branch b . This readily implies that it must also be a model of the subcategory of the syntactic category of the κ -coherent theory axiomatized by

$\mathbb{T} \cup \{\phi_i(\mathbf{c}_i)\}_{i < \kappa}$ which contains formulas in the κ -fragment generated by these axioms. In turn, such a subcategory has an induced canonical morphism from \mathcal{C}'_b due to the universal property of this latter. This will be used for the Kripke completeness theorem.

3.4 Kripke completeness for κ -first-order theories

The categorical proof of the completeness theorem can be applied in particular to any κ -Heyting category, e.g. the syntactic category of a κ -Heyting theory \mathbb{T} of size κ . The κ -Heyting structure is preserved under slicing; moreover, the colimit \mathcal{C}'_b has a functor $[1, -]$ such that the composite $[1, -] \circ F$ is a model of the subcategory of the κ -Heyting syntactic category of the theory $\mathbb{T} \cup \{\phi_i(\mathbf{c}_i)\}_{i < \kappa}$ which contains formulas in the κ -fragment generated by these axioms (see the last paragraph of section 3.3). In turn, the syntactic category of this latter theory has, by the universal property of \mathcal{C}'_b , an induced morphism from it. Hence, the construction really gives us a set of κ -Heyting prime theories⁶ over the language in which we add κ many fresh constants (namely, the κ -Heyting theories $\mathbb{T}_b := \mathbb{T} \cup \{\phi_i(\mathbf{c}_i)\}_{i < \kappa}$). These have, in addition, a conservativity property: given any tuple of less than κ constants \mathbf{c} in the new language such that $\mathbb{T}, \phi(\mathbf{c}) \not\vdash \psi(\mathbf{c})$, there is one theory \mathbb{T}_b in our set of new κ -Heyting theories such that $\mathbb{T}_b, \phi(\mathbf{c}) \not\vdash \psi(\mathbf{c})$. As it is well-known, at least in the finitary case, this property is essentially what allows to build a set of jointly conservative Kripke models⁷ of \mathbb{T} over a tree of height ω . Indeed, we can iterate the construction that builds the set of prime theories over the extended language ω many times, continuing with each $\mathcal{C}_{\mathbb{T}_b}$ at limit levels, forming the said tree. Each node is identified with the term model of the corresponding \mathcal{C}'_b , with the induced homomorphisms between them arising from the Löwenheim-Skolem theorem (cf. section 3.3). It is not hard to see that the theory \mathbb{T}_b , in the appropriate language, represents the set of all κ -Heyting sentences that are forced in the corresponding node. To prove this, we can make a straightforward induction on the complexity of the formula using the property of the set of prime theories (extending the theory of each node) that we have mentioned above. Since the functors between each category $\mathcal{C}_{\mathbb{T}_b}$ and each of its extensions is κ -Heyting, a formula proved by \mathbb{T}_b must be forced in the corresponding node. On the other hand, the theories \mathbb{T}_b are prime theories for the following reason: if \mathbb{T}_b proves $\bigvee_i \exists \mathbf{x} \phi_i$, its κ -coherent Morleyization⁸ $\bigvee_i \exists \mathbf{x} P_{\phi_i}$ holds in the κ -coherent model corresponding to its node, whence, some $P_{\phi_i}(\mathbf{c})$ must hold there and therefore also holds in the successor nodes. Thus, by the construction of section 3.2, it is provable in $\mathcal{C}_{\mathbb{T}_b}$. To see that forced formulas are provable in the corresponding theory, the non trivial case is when ϕ is of the form $\forall \mathbf{x}(\psi \longrightarrow \eta)$. Then if a node forces ϕ , the sequent $\psi \vdash_{\mathbf{x}} \eta$ must be forced in all successor nodes, being thus proved in each successor theory by inductive hypothesis. Therefore, its κ -coherent Morleyization holds in all successor nodes, so that by the construction of section 3.2, it is provable in $\mathcal{C}_{\mathbb{T}_b}$.

If we start with the κ -Heyting theory \mathbb{T} , this procedure provides a tree of prime theories over each immediate successor which will be a jointly conservative set of Kripke models for \mathbb{T} . We will make use of this observation for the downward categoricity transfer.

3.5 Representation theorem for κ -toposes with enough κ -points

Awodey and Forssell provided in [AF13] a logical approach to the representation theorem of Butz and Moerdijk for toposes with enough points, providing actually a reconstruction result

⁶A prime theory is a deductively closed theory with the disjunction and existence properties.

⁷A Kripke model of \mathbb{T} over P is a κ -Heyting functor $\mathcal{C}_{\mathbb{T}} \longrightarrow \mathbf{Set}^P$ for a given poset P .

⁸The κ -coherent Morleyization of \mathbb{T}_b is the theory of κ -coherent models of $\mathcal{C}_{\mathbb{T}_b}$. Each formula ϕ in the original signature is then equivalent, in all κ -coherent models of $\mathcal{C}_{\mathbb{T}_b}$, to a formula P_{ϕ} in the new signature (its κ -coherent Morleyization).

that allows to recover the theory (up to pretopos completion) from its category of models. Unlike previous reconstruction results, which relied on using the ultraproduct structure in the category of models, this result puts instead a topological structure that allows to recover the classifying topos of the theory as the topos of equivariant sheaves on a topological groupoid. This is possible because the category of models carries a natural topology which can be presented in terms of the formulas of theory, to which one can associate the basic opens.

The fact that the set of opens satisfy the axioms of a topology is then a consequence of the formula construction, and in particular, the fact that a finite intersection of opens is open is directly linked to the fact that a finite conjunction of coherent formulas is coherent. When we work in infinitary logic like, e.g., $\mathcal{L}_{\kappa^+, \kappa}$, the formula construction allows for infinitary conjunctions, which in turn gives the associated topology on the category of models a distinctive property: the intersection of less than κ many open sets is still open. We refer to it as a κ -topology. It turns out that many of the properties and results on topological spaces have also variants for κ -topologies. As a result, almost all the setup available in [AF13] is readily generalizable to the infinitary case as soon as we apply our completeness theorem relating the syntax and the semantics. Such a completeness theorem is crucial since the classifying topos is obtained through the representation theorem of Butz and Moerdijk, which requires the topos to have enough points.

A similar representation theorem for κ -classifying toposes can now be shown. In the case of κ -Grothendieck toposes, it turns out that their main characteristic (the transfinite transitivity property) is shared by toposes of equivariant sheaves on topological groupoids where we use κ -topologies. This is what allows, in essence, the transfer of Butz-Moerdijk result to the infinitary case; we briefly mention the steps to do it:

1. The logical topology on the set of models corresponding to a κ -coherent decidable theory is a κ -topology, meaning that intersection of less than κ open sets is open. This is evident from the fact that intersections of basic opens corresponding to κ -coherent formulas corresponds in turn to the conjunction of the formulas, which is κ -coherent.
2. Sheaves on a κ -topological space X (étale bundles over X) satisfy the transfinite transitivity property. This is a direct consequence of a result proven in [Esp20], namely, that when the underlying category of the site satisfies the transfinite transitivity property (as is the case with the lattice of open sets of a κ -topology), so does the sheaf topos on it.
3. Continuous functions between κ -topological spaces give rise to κ -geometric morphisms between the corresponding sheaf toposes, i.e., geometric morphisms whose inverse images preserve κ -limits. Moreover, morphisms of sites that satisfy the transfinite transitivity property induce κ -geometric morphisms. This is essentially contained in [Esp20]. With a proof similar to that of the universal property of a κ -classifying topos we can see that morphisms of sites satisfying the transfinite transitivity property give rise to κ -geometric morphisms between their corresponding κ -Grothendieck toposes on them.
4. Given a κ -topological groupoid, the topos of equivariant sheaves, defined as a (pseudo-) colimit of toposes, satisfies the transfinite transitivity property. This is a consequence of the construction of pseudo-colimits of κ -Grothendieck toposes, as explained, e.g., in [Moe88] or [Kel89]: it suffices to prove that the transfinite transitivity property is preserved during the construction of coproducts, tensor products, iso-coinserters and iso-coequifiers. Alternatively, we can use the explicit description of the topos of equivariant sheaves in terms of local homeomorphisms.
5. The theory is recovered up to κ -pretopos completion as the full subcategory of the κ^+ -compact decidable subobjects. (A κ^+ -compact object is one such that every family of

subobjects that cover it contains a subfamily of at most κ subobjects that still covers). This is a consequence of the universal property of the κ -classifying topos.

More explicitly: starting with the category of models of a κ -coherent theory, we define the following topological groupoid of models. Fixing a set U of size κ , we consider the set G_0 of all models whose underlying structure has elements from U . Then we consider the set G_1 of isomorphisms between them. We put a topology on G_0 whose basic opens consist of sets of models of the form:

$$(\phi(\mathbf{x}), \mathbf{c}) := \{M \in G_0 : \mathbf{c} \in [[\phi(\mathbf{x})]]^M\}$$

Next, we consider the least topology on G_1 which makes the source and target $s, t : G_1 \rightarrow G_0$ continuous and that contains all sets of the form:

$$(\mathbf{a} \mapsto \mathbf{b}) := \{f \in G_1 : \mathbf{a} \in s(f) \text{ and } f(\mathbf{a}) = \mathbf{b}\}$$

In our case, the topos of equivariant sheaves on our topological groupoid has as objects pairs $(a : A \rightarrow G_0, \alpha)$, where a is a local homeomorphism and $\alpha : G_1 \times_{G_0} A \rightarrow A$ satisfies the conditions:

- $a(\alpha(f, x)) = t(f)$
- $\alpha(1, x) = x$
- $\alpha(g, \alpha(f, x)) = \alpha(gf, x)$

We get in the end:

Theorem 3.2. *The κ -classifying topos of \mathbb{T} is precisely the topos of equivariant sheaves $\mathcal{S}h_{G_1}(G_0)$*

The proof follows essentially the same outline of Theorem 1.4.8 in [AF13], with the only modification that we consider syntactic categories of κ -theories and topological groupoids of models with respect to the κ -logical topology. The only detail which needs a different justification is in showing that the embedding $\mathcal{C}_{\mathbb{T}} \rightarrow \mathcal{S}h_{G_1}(G_0)$ given by sending a formula to its corresponding definable set functor is cover-reflecting when we consider the canonical coverage in $\mathcal{S}h_{G_1}(G_0)$ and the κ -coherent coverage in $\mathcal{C}_{\mathbb{T}}$. But this is exactly given by the completeness theorem of section 3.2, since the models of G_0 are enough.

4 The λ -classifying topos of a κ -theory

In this section fix $\kappa < \lambda$ such that $\kappa^{<\kappa} = \kappa$ and $\lambda^{<\lambda} = \lambda$. Let \mathbb{T} be a κ -coherent theory in $\mathcal{L}_{\kappa^+, \kappa}$, $\mathcal{C}_{\mathbb{T}}$ be its syntactic category and $Mod_{\lambda}(\mathbb{T})$ be the full subcategory of λ -presentable models. Assume that the category of models of \mathbb{T} is λ -accessible (this is the case, e.g., if $\lambda = \kappa^+$ or, more generally, if $\kappa^+ \trianglelefteq \lambda$). Let \mathbb{T}' be the theory in $\mathcal{L}_{\lambda^+, \lambda}$ with the same axioms as those of \mathbb{T} . An important result we will prove here is the following:

Theorem 4.1. *The λ -classifying topos of \mathbb{T}' is equivalent to the presheaf topos $\mathbf{Set}^{Mod_{\lambda}(\mathbb{T})}$. Moreover, the canonical embedding of the syntactic category $\mathcal{C}_{\mathbb{T}'} \hookrightarrow \mathbf{Set}^{Mod_{\lambda}(\mathbb{T})}$ is given by the evaluation functor, which on objects acts by sending (\mathbf{x}, ϕ) to the functor $\{M \mapsto [[\phi]]^M\}$.*

Proof. By hypothesis every model of \mathbb{T}' is a λ -filtered colimit of models in $Mod_\lambda(\mathbb{T})$. Note first that the following diagram:

$$\begin{array}{ccc}
\mathcal{C}_{\mathbb{T}'} & \xrightarrow{ev} & \mathbf{Set}^{Mod_\lambda(\mathbb{T})} \\
M \cong \varinjlim_i M_i \searrow & & \swarrow M' \cong \varinjlim_i ev_{M_i} \\
& \mathbf{Set} &
\end{array}$$

commutes up to invertible 2-cell. Here ev and ev_{M_i} are the evaluation functors, defined on objects as $ev((\mathbf{x}, \phi)) = \{M \mapsto [[\phi]]^M\}$ and $ev_{M_i}(F) = F(M_i)$, respectively, while $\varinjlim M_i$ is the canonical λ -filtered colimit of λ -presentable models associated to the model M . Note also that since λ -filtered colimits commute with λ -small limits, M' will preserve, in addition to all colimits, also λ -small limits.

Let now $\mathbf{Set}[\mathbb{T}']_\lambda$ be the λ -classifying topos of \mathbb{T}' . We shall prove that this latter is equivalent to $\mathbf{Set}^{Mod_\lambda(\mathbb{T})}$ by verifying in this presheaf topos the universal property of $\mathbf{Set}[\mathbb{T}']_\lambda$, i.e., that models of \mathbb{T}' in a λ -topos \mathcal{E} corresponds to λ -geometric morphisms from \mathcal{E} to the presheaf topos. It is enough to prove this universal property in the particular case in which $\mathcal{E} = \mathbf{Set}[\mathbb{T}']_\lambda$.

Given then the λ -classifying topos \mathcal{E} of \mathbb{T}' , by the completeness theorem (see section 3.2) it will have enough λ -points. Hence, there is a conservative λ -geometric morphism with inverse image $E : \mathcal{E} \rightarrow \mathbf{Set}^I$ such that composition with the evaluation at $i \in I$, $ev(i)E$ gives a λ -point of \mathcal{E} . Now each model of \mathbb{T}' in \mathcal{E} , $N : \mathcal{C}_{\mathbb{T}'} \rightarrow \mathcal{E}$ gives rise to models in \mathbf{Set} by considering their images through each $ev(i)E$. These correspond to unique (up to isomorphism) λ -geometric morphisms with inverse image $\mathbf{Set}^{Mod_\lambda(\mathbb{T})} \rightarrow \mathbf{Set}$, which in turn induce a λ -geometric morphism with inverse image $G : \mathbf{Set}^{Mod_\lambda(\mathbb{T})} \rightarrow \mathbf{Set}^I$ and with the property that the composition $G \circ ev : \mathcal{C}_{\mathbb{T}'} \rightarrow \mathbf{Set}^{Mod_\lambda(\mathbb{T})} \rightarrow \mathbf{Set}^I$ is the same (up to isomorphism) as $EN : \mathcal{C}_{\mathbb{T}'} \rightarrow \mathbf{Set}^I$. In other words, considering \mathcal{E} as a subcategory of \mathbf{Set}^I , the image of $G \circ ev$ belongs to \mathcal{E} .

$$\begin{array}{ccc}
\mathcal{C}_{\mathbb{T}'} & \xrightarrow{ev} & \mathbf{Set}^{Mod_\lambda(\mathbb{T})} \\
N \searrow & & \swarrow G \\
& \mathcal{E} & \\
E \downarrow & & \swarrow \\
& \mathbf{Set}^I & \\
ev(i) \downarrow & & \\
& \mathbf{Set} &
\end{array}$$

On the other hand, every object F in $\mathbf{Set}^{Mod_\lambda(\mathbb{T})}$ can be canonically expressed as a colimit of representables, $F \cong \varinjlim_i [M_i, -]$. In turn, and supposing for a moment that each $|M_i| < \lambda$, we have that $M_i : \mathcal{C}_{\mathbb{T}} \rightarrow \mathbf{Set}$ is a λ -small colimit of representables $M_i \cong \varinjlim_j [\phi_{ij}, -]$.

It follows that:

$$F \cong \varinjlim_i [\varinjlim_j [\phi_{ij}, -]_{\mathcal{C}_{\mathbb{T}}}, -]_{Mod_\lambda(\mathbb{T})} \cong \varinjlim_i \varprojlim_j [[\phi_{ij}, -]_{\mathcal{C}_{\mathbb{T}}}, -]_{Mod_\lambda(\mathbb{T})} \cong \varinjlim_i \varprojlim_j ev(\phi_{ij})$$

where the last isomorphism follows from Yoneda lemma. We claim that this even happens if $|M_i| = \lambda$. Indeed, in this case λ is the successor of a singular cardinal μ of cofinality less than θ and M_i has internal size μ , so that by Lemma 2.6 it is a directed colimit of models of cardinality less than μ . Suppose first $\kappa < \mu$. Then we can express:

$$[M_i, -] \cong [\varinjlim_k \varinjlim_j [\phi_{ijk}, -]_{\mathcal{C}_{\mathbb{T}}, -}]_{\text{Mod}_{\lambda}(\mathbb{T})} \cong \varprojlim_k \varprojlim_j \text{ev}(\phi_{ijk})$$

and F has a similar expression, where limits are λ -small. If, on the other hand, $\kappa = \mu$, then we note that the restriction functor $r : \mathbf{Set}^{\mathcal{K}_{\geq \nu, \leq \mu}} \rightarrow \mathbf{Set}^{\mathcal{K}_{\mu}}$ has a right adjoint i such that $ri(F) \cong F$ for each F in $\mathbf{Set}^{\mathcal{K}_{\mu}}$, and expressing now $i(F)$ as a colimit of limits of evaluations, as above, we can now apply r , which preserves colimits and limits, and arrive to the same expression for F .

Now G preserves λ -small limits and colimits, and so we will have:

$$G(F) \cong \varinjlim_i \varprojlim_j G \circ \text{ev}(\phi_{ij})$$

and similarly on arrows. Therefore, G is completely determined (up to isomorphism) by its value on the objects $\text{ev}(\phi_{ij})$. Since the value of G on such objects belongs to \mathcal{E} , and E preserves λ -small limits and colimits, it follows that G itself factors through \mathcal{E} . Moreover, it is the unique (up to isomorphism) inverse image of a λ -geometric morphism corresponding to the given model in \mathcal{E} . This finishes the proof. \square

5 The omitting types theorem for infinite quantifier languages

We will assume in this section that $\kappa^{<\kappa} = \kappa$. By a type we understand a consistent set of formulas in a given tuple of variables. It is complete when the set is maximal. The goal of this section is to prove the following:

Theorem 5.1. (*Omitting types theorem for infinite quantifier languages*) Assume GCH , and let κ be a regular cardinal. Let F be a Boolean κ -fragment of $\mathcal{L}_{\kappa^+, \kappa}$ containing a consistent theory \mathbb{T} whose category of models has directed colimits, and let $\{p_i : i < \kappa\}$ be a set of non-isolated types. Then there is a model of \mathbb{T} that simultaneously omits all the types.

Note that this version of the omitting types theorem can be expressed in an entirely semantical way:

Theorem 5.2. Assume GCH . Let \mathbb{T} be a satisfiable theory in a Boolean κ -fragment of $\mathcal{L}_{\kappa^+, \kappa}$ whose category of models has directed colimits, and let p_i , for each $i < \kappa$, be a set of formulas of the fragment. Suppose that whenever ψ is such that $\mathbb{T} \cup \exists \mathbf{x} \psi$ is satisfiable, there is ϕ in p_i such that $\mathbb{T} \cup \exists \mathbf{x}(\psi \wedge \neg \phi)$ is satisfiable. Then the theory:

$$\mathbb{T} \cup \bigwedge_{i < \kappa} \forall \mathbf{x} \bigvee_{\phi \in p_i} \neg \phi(\mathbf{x})$$

is satisfiable.

Proof. Consider the (Boolean) syntactic category of \mathbb{T} in $\mathcal{L}_{\kappa^+, \kappa}$ and the subcategory $\mathcal{C}_{\mathbb{T}}$ given by those formulas in context $[\mathbf{x}, \phi]$ belonging to the κ -fragment F . For each type $p_i = \{\phi_i(\mathbf{x})\}_{i < \kappa}$ consider the family of arrows $U_i = \{[\mathbf{x}, \neg \phi_i] \rightarrow [\mathbf{x}, \top]\}_{i < \kappa}$. Put a κ -Grothendieck topology τ on $\mathcal{C}_{\mathbb{T}}$ generated by:

1. all κ -small jointly epic families of arrows and the κ^+ -small jointly epic families of arrows corresponding to axioms of (the κ -coherent Morleyization of) \mathbb{T}
2. the families U_i for each type p_i

It follows that a κ -flat continuous functor $\mathcal{C}_{\mathbb{T}} \rightarrow \mathbf{Set}$, i.e., a κ -point of the corresponding κ -topos of sheaves, is exactly a model of \mathbb{T} omitting all of the p_i . This topos is clearly κ -separable, so that by the completeness theorem (see section 3.2) it will have enough κ -points. However, we need to verify that it is non-degenerate to guarantee that there will be at least one non-trivial such model (it is easy to see that, if one of the types is isolated, the topos is degenerate, but we will see that this is the only obstruction).

It is enough to verify that the representable functor $[-, 0]$ is a sheaf for any τ -covering family, since then the conservativity of Yoneda embedding will imply that $\mathcal{Sh}(\mathcal{C}_{\mathbb{T}}, \tau)$ is non-degenerate. Now any such τ -covering family is built via pullbacks and transfinite composites from the two types of covers specified above. Clearly, $[-, 0]$ is a sheaf for the first type of covers. The covering U_i becomes, in the κ -classifying topos \mathcal{E} of \mathbb{T} (that is, when sheafifying with respect to only the first type of covers) a family $\{\neg C_i \rightarrow A\}$, not necessarily epimorphic. But since the type p_i was non-isolated, it follows that $\bigwedge_{i < \kappa} C_i = 0$ in \mathcal{E} , or, what is the same, $A = \neg \bigvee_{i < \kappa} \neg C_i$ there. This means that the family $\{\neg C_i \rightarrow A\}$, while not necessarily covering, is covering up to a double negation. It is easy to see that this is also true for the image of the cover through the conservative evaluation functor $\mathcal{E} \rightarrow \mathbf{Set}^{Mod(\mathbb{T}_{\leq \kappa})}$, by using that in $\mathbf{Set}^{Mod(\mathbb{T}_{\leq \kappa})} \cong \mathbf{Set}[\mathbb{T}]_{\kappa^+}$ the sequent $\bigwedge_{i < \kappa} \neg \neg \phi_i \vdash_{\mathbf{x}} \neg \neg \bigwedge_{i < \kappa} \phi_i$ holds (which can be easily verified to be a consequence of the category of models having directed colimits). Indeed, this is equivalent to saying that if the pullback of the cover along a morphism $A \rightarrow (\mathbf{x}, \top)$ is 0, then $A \cong 0$. Since conjunctions of size κ of atomic formulas generate $\mathbf{Set}^{Mod(\mathbb{T}_{\leq \kappa})}$, it is enough to consider those A of that form. Since the statement holds when A is an atomic formula, it must hold still for such a conjunction due to the validity of the sequent $\bigwedge_{i < \kappa} \neg \neg \phi_i \vdash_{\mathbf{x}} \neg \neg \bigwedge_{i < \kappa} \phi_i$ and the construction of pullbacks in the syntactic category $(\mathcal{C}_{\mathbb{T}})_{\kappa^+}$. In summary, our dense covers in \mathcal{E} remain dense in $\mathbf{Set}[\mathbb{T}]_{\kappa^+}$. The same is, of course, true for pullbacks of such families, and in fact for a transfinite composite of such families (for this latter fact we use again the sequent $\bigwedge_{i < \kappa} \neg \neg \phi_i \vdash_{\mathbf{x}} \neg \neg \bigwedge_{i < \kappa} \phi_i$ and the transfinite transitivity rule from [Esp20]). In particular, this means that if the domains of the arrows in a τ -covering family are 0, so is the common codomain. This says precisely that $[-, 0]$ is a sheaf for the τ -covering family, as we wanted. \square

6 Partial isomorphisms and the λ -classifying topos

We will prove now a connection between $\mathcal{L}_{\infty, \lambda}$ -equivalence and λ -classifying toposes which will be useful. It is essentially a consequence of the omitting types theorem we proved before (Theorem 5.1), and is inspired by model-theoretic arguments of Vaught on atomic and prime models and topos-theoretic results from Blass and Ščedrov on Boolean classifying toposes.

Theorem 6.1. *Let \mathbb{T} be a theory in a Boolean λ -fragment of $\mathcal{L}_{\lambda^+, \lambda}$ whose category of models has directed colimits. Then if $\lambda^{<\lambda} = \lambda$, every pair of models of \mathbb{T} of size λ are $\mathcal{L}_{\infty, \lambda}$ -elementarily equivalent if and only if the λ -classifying topos of the theory $\mathbb{T}_{\lambda} := \mathbb{T} \cup \{\text{“there are } \lambda \text{ distinct elements”}\}$ is two-valued and Boolean (alternatively, atomic and connected).*

Proof. (\implies) Suppose any two models of \mathbb{T} of size λ are $\mathcal{L}_{\infty, \lambda}$ -elementarily equivalent and consider the syntactic category \mathcal{C} of the theory $\mathbb{T} \cup \{\text{“there are } \lambda \text{ distinct elements”}\}$, axiomatized

in an appropriate Boolean fragment of $\mathcal{L}_{\lambda^+, \lambda}$ ⁹. This latter theory is clearly complete, since by (downward) Löwenheim-Skolem theorem it follows that all models are $\mathcal{L}_{\lambda^+, \lambda}$ -elementarily equivalent to a model \mathcal{M} of cardinality λ . Therefore, its λ -classifying topos must be two-valued. To see that it is Boolean, we will prove first that \mathcal{C} is atomic, i.e., each Boolean algebra of subobjects of a given object is atomic.

Let $[\mathbf{x}, \psi(\mathbf{x})]$ be non-zero in \mathcal{C} ; then it is satisfiable in a model of cardinality λ by the completeness theorem for $\mathcal{L}_{\lambda^+, \lambda}$, so that there is \mathbf{a} in \mathcal{M} with $\mathcal{M} \models \psi(\mathbf{a})$. Let p be the type $\{\phi(\mathbf{x}) : \mathcal{M} \models \phi(\mathbf{a})\}$. If p was non-isolated, there would be a model omitting it, i.e., there would exist a model \mathcal{N} of:

$$\mathbb{T} \cup \{\text{“there are } \lambda \text{ distinct elements”}\} \cup \{\forall \mathbf{x} \bigvee_{\mathcal{M} \models \phi(\mathbf{a})} \neg \phi(\mathbf{x})\}$$

This is impossible since by hypothesis \mathcal{M} and \mathcal{N} are $\mathcal{L}_{\infty, \lambda}$ -elementarily equivalent. Therefore, p must be isolated by some $\theta(\mathbf{x})$, which must then be an atom in the Boolean algebra of subobjects of $[\mathbf{x}, \top]$. It follows from this that such algebra is atomic. Moreover, the join of all atoms is the top element (as this is the case in any one, and therefore all, models).

Let us now see that the λ -classifying topos must be Boolean. Such a topos is built by considering sheaves on \mathcal{C} when equipped with the κ -topology τ generated by those jointly epic families of cardinality at most λ that corresponds to axioms of (the κ -geometric Morleyization of) the theory. Let \mathcal{C}' be the full subcategory of \mathcal{C} consisting of non-zero objects, and τ' the κ -topology induced by τ . Then the topos $\mathcal{S}h(\mathcal{C}', \tau')$ is still the λ -classifying topos, but now the κ -topology τ' coincides with the coverage ρ consisting of stable nonempty sieves. Indeed, the nontrivial part is showing that a ρ -covering sieve R on an object $[\mathbf{x}, \phi]$ of \mathcal{C}' is also τ' -covering. Since ϕ is a union of at most λ atoms, for each atom there is an arrow in R factoring through it, and since its domain is nonzero, its image must be the whole atom. Choosing one such morphism of R for each atom we get a jointly epic family from τ contained in R . Finally, it follows that the λ -classifying topos is equivalent to the topos of sheaves on \mathcal{C}' for the double negation topology, which is Boolean.

(\Leftarrow) Suppose that the λ -classifying topos of the theory

$$\mathbb{T} \cup \{\text{“there are } \lambda \text{ distinct elements”}\}$$

is two-valued and Boolean. Since it is also λ -separable, it has enough λ -points (see section 3.2), and in particular it must be atomic. Hence, \mathcal{C} is also atomic. Let $p_i = \{\theta(\mathbf{x}_i) : \theta \text{ is an atom in } \mathcal{S}ub([\mathbf{x}_i, \top])\}$, where for each $i < \kappa$, $\mathbf{x}_i = x_0 x_1 \dots$ up to (but excluding) i . Then the family $[\mathbf{x}_i, \theta] \longrightarrow [\mathbf{x}_i, \top]$ is jointly epic. Therefore, any λ -point of the topos corresponds to an atomic model of \mathbb{T} of cardinality at least λ . Since the topos is also two-valued, all such atomic models are $\mathcal{L}_{\lambda^+, \lambda}$ -elementarily equivalent. Hence, a back and forth argument shows that any two such models of cardinality λ must be $\mathcal{L}_{\infty, \lambda}$ -elementarily equivalent.

To complete the proof, notice that atomic toposes are Boolean, while a Boolean topos with enough points must be atomic, and Boolean toposes are two-valued if and only if they are connected. \square

Using Theorem 6.1 we can now get the following:

⁹This can be done through sentences expressing that for each tuple of less than λ distinct elements there is one element different from all those of the tuple.

Corollary 6.2. *In a κ -separable topos, whose category of κ -points has directed colimits, two models of size κ are $\mathcal{L}_{\infty, \kappa}$ -elementarily equivalent if and only if the topos is two-valued and Boolean (alternatively, atomic and connected). In particular, categoricity in κ implies that any morphism between models of size at least κ is $\mathcal{L}_{\infty, \kappa}$ -elementary.*

Proof. As explained in [Esp20], the topos κ -classifies a κ -coherent theory \mathbb{T} . If we let \mathbb{T}_B be the theory obtained from \mathbb{T} by adding all instances of excluded middle over $\mathcal{L}_{\kappa^+, \kappa}$, then we have a stable surjection $s : \mathbf{Set}[\mathbb{T}_B]_{\kappa} \twoheadrightarrow \mathbf{Set}[\mathbb{T}]_{\kappa}$ in the category of Grothendieck toposes (indeed, given the pullback \mathcal{F} of s along $\mathcal{G} \twoheadrightarrow \mathbf{Set}[\mathbb{T}]$, if \mathcal{G} classifies the geometric theory \mathbb{S} , then \mathcal{F} will classify the theory \mathbb{S}_B obtained by adding instances of excluded middle over $\mathcal{L}_{\infty, \omega}$, which is a conservative extension). Since two-valued Boolean toposes are atoms in the lattice of subtoposes of a given topos, it is enough to show that $\mathbf{Set}[\mathbb{T}]_{\kappa}$ has no proper non-degenerate subtoposes if any two models of size κ are isomorphic. Suppose this latter condition holds; consider a subtopos $i : \mathcal{E} \hookrightarrow \mathbf{Set}[\mathbb{T}]$, and pull it back along s . We get a geometric morphism $t : \mathcal{T} \twoheadrightarrow \mathbf{Set}[\mathbb{T}_B]_{\kappa}$ and a surjection $s' : \mathcal{T} \twoheadrightarrow \mathcal{E}$. Now $\mathbf{Set}[\mathbb{T}]_{\kappa}$ has the same models of size κ as $\mathbf{Set}[\mathbb{T}_B]_{\kappa}$ has; if we prove that the sequent $\bigwedge_{i < \kappa} \neg\neg\phi_i \vdash_{\mathbf{x}} \neg\neg\bigwedge_{i < \kappa} \phi_i$ holds in $(\mathcal{C}_{\mathbb{T}_B})_{\kappa^+}$, we could run the same proof of Theorem 6.1 (since we can apply the omitting types theorem) to show that $\mathbf{Set}[\mathbb{T}_B]_{\kappa}$ is Boolean and two-valued, equivalent in turn to having no proper non-degenerate subtoposes. Therefore, by considering the surjection-embedding factorization of t , this is equivalent to either t being a surjection or \mathcal{T} being degenerate. In the first case, it follows that the composite st is a surjection, and since $st \simeq is'$, that i must be a surjection, in which case \mathcal{E} is equivalent to $\mathbf{Set}[\mathbb{T}]_{\kappa}$. In the second case, since s' is a surjection, it follows that \mathcal{E} must be degenerate.

It remains to prove that $\bigwedge_{i < \kappa} \neg\neg\phi_i \vdash_{\mathbf{x}} \neg\neg\bigwedge_{i < \kappa} \phi_i$ holds in $(\mathcal{C}_{\mathbb{T}_B})_{\kappa^+}$, for which it is enough to show that the sheafification functor from $\mathbf{Set}[\mathbb{T}_B]_{\kappa^+}$ to its double negation subtopos is a κ^+ -geometric morphism. Now, the pullback, in the 2-category of κ^+ -toposes and κ^+ -geometric morphisms, of the double negation subtopos \mathcal{S} of $\mathbf{Set}[\mathbb{T}]_{\kappa^+}$ along $s' : \mathbf{Set}[\mathbb{T}_B]_{\kappa^+} \twoheadrightarrow \mathbf{Set}[\mathbb{T}]_{\kappa^+}$ is equivalent to \mathcal{S} .¹⁰ Therefore, the double negation subtopos has an embedding into $\mathbf{Set}[\mathbb{T}_B]_{\kappa^+}$. Moreover, the embedding is dense, since this latter is two-valued. It follows that $\mathbf{Set}[\mathbb{T}_B]_{\kappa^+}$ and $\mathbf{Set}[\mathbb{T}]_{\kappa^+}$ have equivalent double negation subtoposes, which completes the proof. \square

Recall from [Ros97] that an object M of an accessible category is said to be λ -closed if all morphisms $M \twoheadrightarrow N$ are λ -pure. In model-theoretic terms, this means that all morphisms $M \twoheadrightarrow N$ reflect the truth of λ -coherent existential formulas with (less than λ) parameters from M whenever the category is axiomatizable in λ -coherent logic. One consequence of Corollary 6.2 that will be of use is the following:

Lemma 6.3. *Assume $\lambda^{<\lambda} = \lambda$. If \mathcal{K} is categoricity in $\lambda > \kappa$, the subcategory of $\mathcal{K}_{\geq \lambda}$ where we just consider the monomorphisms is closed under directed colimits.*

Proof. Let \mathbb{T} be the $\mathcal{L}_{\kappa, \theta}$ theory axiomatizing \mathcal{K} . The λ -classifying topos of \mathbb{T}_{λ} is Boolean by Corollary 6.2, so that the evaluation functor $\mathcal{C}_{\mathbb{T}_{\lambda}} \twoheadrightarrow \mathbf{Set}^{\mathcal{K}_{\geq \lambda}}$ is Boolean and the sequent $\top \vdash_{\mathbf{x}} \phi \vee \neg\phi$ is forced in every model above the categoricity cardinal. In particular, all morphisms in the restriction of \mathcal{K}_{λ} to the monomorphisms will be furthermore ω -pure, and since directed colimits of ω -pure morphisms are ω -pure, such restriction is closed under directed colimits. \square

Continuing with the same notation, we have now:

Lemma 6.4. *Assume GCH. If λ is a categoricity cardinal, then the model of size λ is λ -closed.*

¹⁰If the site for \mathcal{S} is $\mathcal{C}_{\mathcal{S}}$, then the site of the pullback is the free κ -Boolean extension of $\mathcal{C}_{\mathcal{S}}$, which coincides with $\mathcal{C}_{\mathcal{S}}$ if this latter is already κ -Boolean.

Proof. Assume first that λ is regular. Then the model M of size λ is λ -atomic. Given a tuple \mathbf{c} and an embedding $M \rightarrow N$ with $N \models \exists \mathbf{x} \phi(\mathbf{x}, \mathbf{c})$, let ϕ_0 be the complete formula that isolates the type of \mathbf{c} in M . Then, since $\phi_0(\mathbf{y}) \wedge \exists \mathbf{x} \phi(\mathbf{x}, \mathbf{y})$ is consistent with \mathbb{T}_λ , we must have $\phi_0(\mathbf{y}) \vdash_{\mathbf{y}} \exists \mathbf{x} \phi(\mathbf{x}, \mathbf{y})$ within \mathbb{T}_λ , and in particular $M \models \exists \mathbf{x} \phi(\mathbf{x}, \mathbf{c})$.

In the case when λ is limit, we consider first the topos $\mathbf{Set}[\mathbb{T}_\lambda^B]_\lambda$, where \mathbb{T}_λ^B is obtained from \mathbb{T}_λ by adding all instances of excluded middle for $< \lambda$ -coherent formulas. Apply now the same proof as in Theorem 5.1 to the λ -type $\{\phi_i(\mathbf{x})\}_{i < \lambda}$ realized by \mathbf{c} in M to show that, if it is not isolated, there is a Grothendieck coverage τ on $(\mathcal{C}_{\mathbb{T}_\lambda^B})_\lambda := \lim\{(\mathcal{C}_{\mathbb{T}_\kappa^B})_\kappa : \kappa < \lambda\}$ which contains covers corresponding to the omission of the type. Now consider the natural functor $((\mathcal{C}_{\mathbb{T}_\lambda^B})_\lambda, \tau) \rightarrow ((\mathcal{C}_{\mathbb{T}_\lambda^B})_{\lambda^+}, \tau)$ inducing a morphism of the corresponding toposes. The codomain topos has enough λ^+ -points, and the sequent $\bigwedge_{i < \kappa} \neg\neg \phi_i \vdash_{\mathbf{x}} \neg\neg \bigwedge_{i < \kappa} \phi_i$ holds there (as can be seen by noticing that $\mathbf{Set}[\mathbb{T}_\lambda^B]_{\lambda^+}$ and $\mathbf{Set}[\mathbb{T}_\lambda]_{\lambda^+}$ have equivalent double negation subtoposes, proved analogously to the last paragraph of the proof of Corollary 6.2), so the type in question must be omitted by some model of size at most λ^+ . But this model must contain the model of size λ as a submodel which realizes the type, which is absurd. This proves that the type is isolated by a complete Boolean formula and hence that $\mathbf{Set}[\mathbb{T}_\lambda^B]_\lambda$ is two-valued and Boolean. With the same argument as in Corollary 6.2, we can conclude that also $\mathbf{Set}[\mathbb{T}_\lambda]_\lambda$ is two-valued and Boolean, whence atomic (as the point of size λ must be a surjection) and then we can proceed as in the regular case. \square

Remark 6.5. *It follows with the same proof idea of Lemma 6.4 that if \mathbb{T} is λ -categorical then $\mathbf{Set}[\mathbb{T}_\lambda]_\lambda$ is two-valued and Boolean, even if λ is singular (in which case by $\mathbf{Set}[\mathbb{T}_\lambda]_\lambda$ we understand the limit of the toposes $\mathbf{Set}[\mathbb{T}_\nu]_\nu$ for regular $\nu < \lambda$).*

We get now:

Lemma 6.6. *Assume GCH and amalgamation. Then categoricity cardinals are closed in the class of all cardinals.*

Proof. Let λ be a limit of categoricity cardinals. Any model of size λ is a directed colimit of models of size ν for ν a categoricity cardinal. By Lemma 6.4, each of the models in the colimit is ν -closed. Since the Löwenheim-Skolem theorem (namely, that every subset $A \subseteq M$ of a model of size κ^+ with $|A| = \kappa$ contains a submodel $N \supseteq A$ of size κ) still holds in terms of internal size, and under SCH, for those cardinals κ which have cofinality less than θ , it follows that the model of size λ is λ -closed. Since amalgamation holds, it is also λ -saturated, and there is only one such model of size λ up to isomorphism (cf. [Ros97]). \square

7 Classifying toposes for saturated models

Recall from [Ros97] the categorical definition of κ -saturated model as a model M such that for every $p : N \rightarrow N'$ with N, N' κ -presentable, each $q : N \rightarrow M$ extends to some $q' : N' \rightarrow M$:

$$\begin{array}{ccc} & N' & \\ & \uparrow p & \\ N & \xrightarrow{q} & M \end{array} \quad \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} q' \\ \text{---} \\ \text{---} \end{array}$$

This notion coincides with the usual model-theoretic notion in finitary theories, and can be used in the infinitary case to simplify considerably the arguments of [SM90].

We are now going to prove the following:

Theorem 7.1. *Assume GCH and amalgamation. Then the κ^+ -classifying topos $\mathbf{Set}^{[\mathbb{T}_{\kappa^+}^{\text{sat}}]_{\kappa^+}}$ of the theory of κ^+ -saturated models is precisely $\mathcal{Sh}(\mathcal{K}_{\kappa}^{\text{op}}, \tau_D)$, where τ_D is the dense topology (i.e., it is the double negation subtopos of $\mathbf{Set}^{\mathcal{K}_{\kappa}}$).*

Proof. Consider the following diagram:

$$\begin{array}{ccc}
 (\mathcal{C}_{\mathbb{T}_{\kappa}})_{\kappa^+} & \xrightarrow{ev} & \mathbf{Set}^{\mathcal{K}_{\kappa}} \\
 & & \searrow M \cong \varinjlim_i ev_{N_i} \\
 & \downarrow f^* & \mathbf{Set} \\
 & \mathcal{Sh}(\mathcal{K}_{\kappa}^{\text{op}}, \tau_D) & \nearrow \text{dashed}
 \end{array}$$

The existence of directed colimits readily implies that the dense topology in $(\mathcal{C}_{\mathbb{T}_{\kappa}})_{\kappa^+}$ is a κ^+ -topology (transfinite composites up to κ of dense covers are dense). Whence, the double negation subtopos is a κ^+ -topos and the sheafification functor preserves κ^+ -small colimits. Since there is amalgamation, we have that a model M is κ^+ -saturated if and only if for every $p : N \rightarrow N'$ in \mathcal{K}_{κ} , each $q : N \rightarrow M$ extends to some $q' : N' \rightarrow M$ (see the diagram above). This is the same as saying that $M : \mathbf{Set}^{\mathcal{K}_{\kappa}} \rightarrow \mathbf{Set}$ maps $p^* : [N', -] \rightarrow [N, -]$ to an epimorphism, since, writing $M \cong \varinjlim_i ev_{N_i}$ as a κ^+ -filtered colimit of evaluations, and given that any N in \mathcal{K}_{κ} is κ^+ -presentable, we have:

$$\varinjlim_i ev_{N_i}([N, -]) = \varinjlim_i [N, N_i] \cong [N, \varinjlim_i N_i] \cong [N, M]$$

It follows that M is κ^+ -saturated if and only if it factors through $f^* : \mathbf{Set}^{\mathcal{K}_{\kappa}} \rightarrow \mathcal{Sh}(\mathcal{K}_{\kappa}^{\text{op}}, \tau_D)$, since the amalgamation property in \mathcal{K}_{κ} is precisely the right Ore condition on $\mathcal{K}_{\kappa}^{\text{op}}$, and so it implies that the dense topology coincides with the atomic topology (every nonempty sieve covers). This finishes the proof. \square

Remark 7.2. *The description of the $< \lambda$ -classifying topos for $< \lambda$ -saturated models can be given, with a similar proof as that of Theorem 7.1, and with the same hypothesis, as the double negation subtopos of the topos of presheaves on models of size less than λ , which will therefore be equivalent to the pseudolimit of the κ -classifying toposes for κ -saturated models.*

8 Categoricity and amalgamation

Before proving Grossberg conjecture, we will now show that some of the properties that hold in AEC' s also hold in the context of accessible categories with directed colimits. We start with the following:

Lemma 8.1. *Assume that $(\kappa^+)^{\kappa} = \kappa^+$. Then \mathcal{K}_{κ} has the amalgamation property if and only if $\mathbf{Set}^{\mathcal{K}_{\kappa}}$ forces the sequent $\top \vdash_{\mathbf{x}} \neg \exists \mathbf{y} \phi(\mathbf{x}, \mathbf{y}) \vee \neg \neg \exists \mathbf{y} \phi(\mathbf{x}, \mathbf{y})$ for every κ^+ -coherent formula $\phi(\mathbf{x}, \mathbf{y})$.*

Proof. If \mathcal{K}_{κ} has the amalgamation property then $\mathbf{Set}^{[\mathbb{T}_{\kappa}]_{\kappa^+}} \cong \mathbf{Set}^{\mathcal{K}_{\kappa}}$ is a De Morgan topos (see [Joh79]), so that in particular it forces the relevant sequent. Conversely, suppose that such a sequent is forced and consider two morphisms $M \rightarrow N$ and $M \rightarrow S$. If they cannot be amalgamated, using the Löwenheim-Skolem property, it follows by the completeness theorem (section 3.2) that the $\mathcal{L}_{\kappa^{++}, \kappa^+}$ -theory axiomatized by $\mathbb{T}_{\kappa} \cup \text{Diag}^+(N) \cup \text{Diag}^+(S)$ is inconsistent (where we use the same constants for elements coming from M in the definition of the positive

diagrams). Therefore, there is a κ^+ -coherent formula ϕ with constant symbols \mathbf{c} only from M such that $\mathbb{T}_\kappa \cup \text{Diag}^+(N)$ proves $\neg \exists \mathbf{y} \phi(\mathbf{c}, \mathbf{y})$ (if $\kappa \notin S$ we slightly modify the diagrams). Hence, in $\mathbf{Set}^{\mathcal{K}_\kappa}$, $N \Vdash \neg \exists \mathbf{y} \phi(\mathbf{c}, \mathbf{y})$. Since we have $M \Vdash \neg \exists \mathbf{y} \phi(\mathbf{c}, \mathbf{y}) \vee \neg \neg \exists \mathbf{y} \phi(\mathbf{c}, \mathbf{y})$, we must have $M \Vdash \neg \exists \mathbf{y} \phi(\mathbf{c}, \mathbf{y})$. This however contradicts the fact that $S \Vdash \exists \mathbf{y} \phi(\mathbf{c}, \mathbf{y})$. \square

We now prove that the model-theoretic hypothesis of amalgamation for accessible categories can be deduced from categoricity in a high enough cardinal, assuming that there is a proper class of strongly compact cardinals. We have:

Theorem 8.2. *Assume GCH. Let κ be a strongly compact cardinal and let \mathcal{K} be equivalent to the category of models of some $\mathcal{L}_{\kappa, \kappa}$ theory \mathbb{T} . If \mathcal{K} is categorical at $\lambda \geq \kappa$, then $\mathcal{K}_{\geq \kappa}$ has the amalgamation property.*

Proof. Since λ is a categoricity cardinal, it follows by Lemma 6.4 that N_λ , the model of size λ , is λ -closed, so in particular it is κ -closed. Thus, we have that for any κ -coherent $\sigma(\mathbf{x})$, in the topos $\mathbf{Set}^{\mathcal{K}_{\geq \kappa, \leq \lambda}}$ the model N_λ forces the sentence $\forall \mathbf{x} (\sigma(\mathbf{x}) \vee \neg \sigma(\mathbf{x}))$, so that $\mathbf{Set}^{\mathcal{K}_{\geq \kappa, \leq \lambda}}$ forces its double negation. By the intuitionistic completeness theorem of [Esp19], there is a conservative Heyting embedding $ev : (\mathcal{C}_{\mathbb{T}_\kappa})_\kappa \longrightarrow \mathbf{Set}^{\mathcal{K}_{\geq \kappa, \leq \lambda}}$. Since $(\mathcal{C}_{\mathbb{T}_\kappa})_\kappa$ is two-valued (by categoricity in λ), $(\mathcal{C}_{\mathbb{T}_\kappa})_\kappa$ satisfies $\forall \mathbf{x} (\sigma(\mathbf{x}) \vee \neg \sigma(\mathbf{x}))$. This implies that all morphisms in $\mathcal{K}_{\geq \kappa}$ are κ -pure; amalgamation in $\mathcal{K}_{\geq \kappa}$ follows now as in the finitary case. \square

Note that *SCH* holds above a strongly compact cardinal, so that it is also enough to assume the latter for our results. In the rest of this section we will work to remove the large cardinal hypothesis.

Studying eventual categoricity in an accessible category with directed colimits which is categorical in some cardinal is the same as studying it in the μ -AEC obtained by restricting the morphisms to monomorphisms, since by Lemma 6.3 this restriction also has directed colimits. The advantage is that we can now consider the functorial expansion (see [BGL⁺16] for the definition) of the μ -AEC in question defined by Vasey, who called them substructure functorial expansion, in which we add a κ^+ -small arity predicate P whose interpretation in a model M consists of the image of the underlying structure of a model N of size κ embedded in M through a morphism in the μ -AEC. This particular expansion, which gives rise to an isomorphic μ -AEC, has the property that morphisms coincide with substructure embeddings. Moreover, its models of size at least κ can be axiomatized as follows, extending further the language with the symbol \subseteq :

$$\begin{aligned} \top \vdash_{\mathbf{x}} \exists \mathbf{y} \left(\bigvee_{M_0 \in S} \psi_{M_0}(\mathbf{y}) \wedge \mathbf{x} \subseteq \mathbf{y} \wedge P(\mathbf{y}) \right) \\ \top \vdash_{\mathbf{xy}} (\mathbf{x} \subseteq \mathbf{y} \wedge P(\mathbf{x}) \wedge P(\mathbf{y})) \longrightarrow \bigvee_{(M_0, M_1) \in T} \psi_{(M_0, M_1)}(\mathbf{x}, \mathbf{y}) \\ \top \vdash_{\mathbf{xy}} (P(\mathbf{y}) \wedge \psi_{(M_0, M_1)}(\mathbf{x}, \mathbf{y})) \longrightarrow P(\mathbf{x}) \\ \top \vdash_{\mathbf{xy}} \mathbf{x} \subseteq \mathbf{y} \leftrightarrow \bigwedge_{i \in I} \bigvee_{j \in J} x_i = y_j \end{aligned}$$

Here S is a skeleton of the subcategory of models of size κ , T is the set of pairs (M_0, M_1) with a morphism in the μ -AEC and $M_0, M_1 \in S$, while $\psi_{M_0}, \psi_{M_0, M_1}$ are conjunctions of atomic and negated atomic formulas of the extended language such that $\psi_{M_0}(\mathbf{z})$ holds if and only if \mathbf{z} is

isomorphic to M_0 , and $\psi_{M_0, M_1}(\mathbf{z}, \mathbf{w})$ holds if and only if (\mathbf{z}, \mathbf{w}) is isomorphic to (M_0, M_1) . With these definitions, it should be clear that the previous sequents really axiomatize the μ -AEC in question: the first sequent is just an axiomatic reformulation of the downward Löwenheim-Skolem theorem, while the second and third just expresses that the morphisms are determined solely by the morphisms between models of size κ and that P holds of every strong embedding. The last sequent is separately added for convenience in what follows.

If we assume that we have categoricity at κ , we can now get an axiomatization of an isomorphic μ -AEC which can be entirely rewritten through sequents in the $(2^\kappa)^+$ - Reg_{\neg} fragment¹¹. This is an intuitionistic fragment of first-order logic which contains no disjunctions, obtained from the $(2^\kappa)^+$ -regular fragment by adding \perp , together with the axioms $\perp \vdash_{\mathbf{x}} \phi$ and the axioms for \neg that make it into a negation operator. Indeed, in the first sequent above the disjunction reduces to a single disjunct since we have categoricity at κ , while the last three sequents above have the general form of universal sentences $\forall \mathbf{z} \bigvee_{i \in I} \bigwedge_{j \in J} \psi_{ij}$, and each such sentence is equivalent to the set of sequents $\{\exists \mathbf{z} \bigwedge_{i \in I} \neg \psi_{i, f(i)} \vdash \perp\}_{f \in J^I}$.

The $(2^\kappa)^+$ - Reg_{\neg} fragment contains the $(2^\kappa)^+$ - Reg_{\perp} subfragment, not containing the symbol \neg . The syntactic category \mathcal{D} of any $(2^\kappa)^+$ - Reg_{\neg} theory can be studied through the category $\mathcal{K}_{\geq (2^\kappa)^+}^r$ of its $(2^\kappa)^+$ - Reg_{\perp} models (models of the $(2^\kappa)^+$ - Reg_{\perp} internal theory of \mathcal{D} , also known as the $(2^\kappa)^+$ - Reg_{\perp} Morleyization of the $(2^\kappa)^+$ - Reg_{\neg} theory). These latter are in particular $(2^\kappa)^+$ -regular models for the extended signature in which there is an extra propositional symbol \perp and one predicate symbol S for each negated atomic formula $\neg R$ and where the axioms of the theory contain all axioms obtained from formally replacing $\neg R$ by S in each $(2^\kappa)^+$ - Reg_{\neg} axiom and, additionally, all those axioms of the form $\perp \vdash_{\mathbf{x}} \phi$ and $R \wedge S \vdash_{\mathbf{x}} \perp$. Adding the axioms $\top \vdash_{\mathbf{x}} R \vee S$ gives us a $(2^\kappa)^+$ -coherent theory which axiomatizes the μ -AEC.

If $(\mathcal{C})_{\lambda+}^r$ is the syntactic category of the λ^+ - Reg_{\perp} theory with the same axioms as the $(2^\kappa)^+$ - Reg_{\perp} theory (whose syntactic category is \mathcal{C}) axiomatizing the μ -AEC, then its λ^+ -classifying topos $\mathcal{S}h((\mathcal{C})_{\lambda+}^r, \tau)$ (where τ is the λ^+ - Reg_{\perp} coverage) will be precisely equivalent to the presheaf topos $\mathbf{Set}^{\mathcal{K}_{\geq (2^\kappa)^+, \leq \lambda}^r}$, as can be seen as a special case of Theorem 4.1 when λ is big enough. In particular, the embedding $(\mathcal{C})_{\lambda+}^r \longrightarrow \mathbf{Set}^{\mathcal{K}_{\geq (2^\kappa)^+, \leq \lambda}^r}$ will preserve \neg since it can be identified with Yoneda embedding, which preserves any right adjoint to pullback functors that might exist, see [BJ98]).

Using the compactness of $(2^\kappa)^+$ - Reg_{\perp} logic, it is also easy to verify that the canonical functor $F : \mathcal{C} \longrightarrow (\mathcal{C})_{\lambda+}^r$ also preserves \neg . For if given a λ^+ -regular sentence $\exists \mathbf{x} \bigwedge_{i < \lambda} \phi_i$ we have $\exists \mathbf{y} \bigwedge_{i < \lambda} \phi_i \wedge R \vdash_{\mathbf{x}} \perp$ in λ^+ - Reg_{\perp} logic, there must be a $(2^\kappa)^+$ -regular sentence $\exists_{i \in T} y_i \bigwedge_{i \in T} \phi_i$, for some subset $T \subset \lambda$ of size at most 2^κ , such that $\exists_{i \in T} y_i \bigwedge_{i \in T} \phi_i \wedge R \vdash_{\mathbf{x}} \perp$ in $(2^\kappa)^+$ - Reg_{\perp} logic, from which our result follows.

It follows, in fact, that the evaluation functor $ev : \mathcal{C} \longrightarrow \mathbf{Set}^{\mathcal{K}_{\geq (2^\kappa)^+, \leq \lambda}^r}$, the composite of Yoneda embedding with F , preserves \neg ,¹² which in particular means that the interpretation of S in the presheaf topos will be precisely that of $\neg R$.

Note that, if we add to the $(2^\kappa)^+$ - Reg_{\neg} axiomatization above all instances of excluded middle, we get an axiomatization of (an isomorphic copy of) the μ -AEC. This fact will be essential for the proof of the following:

¹¹I would like to thank Will Boney and Nate Ackerman for catching a glitch in the original version of the argument in the rest of the section, as well as for their careful reading and interest.

¹²It is also possible to give a direct proof of this fact, using the compactness of $(2^\kappa)^+$ - Reg_{\perp} logic, with the same arguments as in the proof of Joyal's theorem, according to which $ev : \mathcal{C} \longrightarrow \mathbf{Set}^{Mod(\mathcal{C})}$ preserves universal quantification when $Mod(\mathcal{C})$ is the category of coherent models of the Heyting category \mathcal{C} . This is worked out in the author PhD thesis for the more general disjunction-free fragment.

Theorem 8.3. (*Grossberg conjecture for ν -AEC's*) Assume SCH (only if $\nu > \omega$), and let ρ be the Hanf number for categoricity. Then categoricity in $\kappa \geq \rho$ and in $\mu > 2^\kappa$ implies that $\mathcal{K}_{\geq \mu}$ has amalgamation.

Proof. Choose a regular $\lambda \geq \mu$, and suppose first that $(\lambda^+)^{\lambda} = \lambda^+$. Let \mathbb{T}^r be the κ^+ - Reg_{\perp} Morleyization of the κ^+ - Reg_{\neg} axiomatization of \mathbb{T}_{κ^+} and $\mathcal{K}_{\geq \kappa^+, \leq \lambda}^r$ its category of κ^+ - Reg_{\perp} models. Then, by the observation above, the subtopos $\mathbf{Set}^{\mathcal{K}_{\geq \kappa^+, \leq \lambda}} \hookrightarrow \mathbf{Set}^{\mathcal{K}_{\geq \kappa^+, \leq \lambda}^r}$ is dense, as it λ^+ -classifies the quotient theory obtained by adding axioms of the form $\top \vdash_x R \vee S$, which are instances of excluded middle for atomic formulas R , since $[[S]] = [[\neg R]]$ in $\mathbf{Set}^{\mathcal{K}_{\geq \kappa^+, \leq \lambda}^r}$. Also, the sheafification functor is given by the transpose i^* of the inclusion functor $i : \mathcal{K}_{\geq \kappa^+, \leq \lambda} \rightarrow \mathcal{K}_{\geq \kappa^+, \leq \lambda}^r$, as can be verified syntactically, and this will be a λ^+ -geometric morphism since i^* preserves all limits and colimits. Note now that the sheafification functor $j : \mathbf{Set}^{\mathcal{K}_{\geq \kappa^+, \leq \lambda}^r} \rightarrow \mathcal{Sh}(\mathcal{K}_{\geq \kappa^+, \leq \lambda}^{op}, \tau_D)$ must send $[N', -] \rightarrow [N, -]$ to an epimorphism whenever $N \rightarrow N'$ in $\mathcal{K}_{\geq \kappa^+, \leq \lambda}^r$ is κ^+ -pure. Indeed, in that case it can be amalgamated with any morphism $N \rightarrow N''$ in $\mathcal{K}_{\geq \kappa^+, \leq \lambda}^r$ using κ^+ -compactness of κ^+ - Reg_{\perp} logic, making $N \rightarrow N'$ a cover for the dense topology. Given then a morphism $N \rightarrow N'$ in $\mathcal{K}_{\geq \mu, \leq \lambda}$, it must be μ - Reg_{\neg} -elementary by categoricity in μ , so that it is in particular a κ^+ -pure embedding in $\mathcal{K}_{\geq \kappa^+, \leq \lambda}^r$. It follows that the sheafification functor $j' : \mathbf{Set}^{\mathcal{K}_{\geq \kappa^+, \leq \lambda}} \rightarrow \mathcal{Sh}(\mathcal{K}_{\geq \kappa^+, \leq \lambda}^{op}, \tau_D)$ must send $[N', -] \rightarrow [N, -]$ to an epimorphism, so that $N \rightarrow N'$ generates a cover for the dense topology in $\mathcal{K}_{\geq \kappa^+, \leq \lambda}$, i.e., amalgamation at $|N|$ holds. We postpone for the last section the arguments that will allow us to eliminate the hypothesis $(\lambda^+)^{\lambda} = \lambda^+$ and other uses of GCH in the case of AEC's and replace it with SCH for general ν -AEC's. \square


In the following section we will make use of Theorem 8.3 and assume that eventual amalgamation holds whenever we have categoricity in a high enough cardinal.

9 Eventual categoricity

We can finally prove the following:

Theorem 9.1. Assume GCH and amalgamation. Suppose the accessible category \mathcal{K} has directed colimits and is categorical in some cardinal $\kappa > \chi$ (where χ is the Hanf number for categoricity). Then it is also categorical at any $\delta > \kappa$.

Proof. Assume first $\delta = \kappa^+$ for a cardinal κ of cofinality at least θ ; let us see that $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^+}$ has a unique κ^+ -point of size κ^+ , up to isomorphism. Using Lemma 6.3, we will assume that all morphisms are monomorphisms, since the categoricity spectrum does not change with such restriction. Consider the diagram:



for a model N_λ of size $\lambda > \kappa^+$ that is a categoricity cardinal. Note that $Sh(K_\kappa^{op}, \tau_D)$ is a κ^+ -topos since τ_D is a κ^+ -topology (which follows in turn from the fact that K_κ has κ^+ -directed colimits). Here $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^+}$ is a dense subtopos of $\mathbf{Set}[\mathbb{T}_\kappa]_{\kappa^+} \cong \mathbf{Set}^{K_\kappa}$ since any model in K_κ embeds in a model of size κ^+ and thus every non-zero κ^+ -coherent sentence in \mathbf{Set}^{K_κ} must be non-zero in $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^+}$. It follows that the double negation subtopos factors through $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^+}$. Moreover, it is in fact a subtopos of $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^+}$ and the sheafification functor (which is the same as the embedding into \mathbf{Set}^{K_κ} followed by the sheafification of this latter) must preserve κ^+ -small limits.

Now N_λ is λ -closed, whence it is κ^+ -closed, and therefore κ^+ -saturated, inducing the bottom morphism above. We will prove that any sequent $\vartheta \vdash_{\mathbf{x}} \psi$ valid in $Sh(K_\kappa^{op}, \tau_D)$ is already valid in $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^+}$. It would then follow that $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^+} \cong Sh(K_\kappa^{op}, \tau_D)$ and every model of size κ^+ is κ^+ -saturated. Categoricity in κ^+ follows (see [Ros97]).

In case κ has cofinality less than θ , since models in K_κ could have cardinality κ^+ , the topos $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^+}$ must be replaced by a topos whose κ^+ -points are precisely those models of internal size at least κ^+ . We can take instead the topos of equivariant sheaves on the topological groupoid of models of internal size κ^+ (cf. section 3.5), which will be precisely the κ^+ -classifying topos of their theory $\mathbb{T}_{\kappa^+}^{int}$. Indeed, if that was not the case, the topos would have some κ^+ point of \mathbf{Set}^{K_κ} which is not of internal size κ^+ , and the only possibility is that it has the model M of internal size κ as the extra κ^+ -point, in which case it would coincide with \mathbf{Set}^{K_κ} , which would then have a conservative embedding into a topos with a jointly conservative set of models of internal size κ^+ . Let us see that this implies that M has a κ^+ -pure embedding into a model of internal size κ^+ (which is impossible, of course). Let ϕ be the conjunction of the diagram of the subset \mathbf{c} of cardinality κ in M which is a concrete directed colimit of models of smaller size, plus the negation of each existential sentence with parameters from \mathbf{c} which does not hold in M . If no model of size κ^+ satisfied ϕ , the sequent $\phi \vdash_{\mathbf{x}} \perp$ would be provable in $\mathbb{T}_{\kappa^+}^{int}$, whence by conservativity it would also be provable in \mathbf{Set}^{K_κ} , which is not possible since $M \models \phi(\mathbf{c})$. In summary, the topos of equivariant sheaves on the topological groupoid of models of internal size κ^+ is precisely the κ^+ -classifying topos of their theory $\mathbb{T}_{\kappa^+}^{int}$. Moreover, $\mathbb{T}_{\kappa^+}^{int}$ is a quotient of \mathbb{T}_{κ^+} , and thus this sheaf topos is a subtopos of \mathbf{Set}^{K_κ} and it is easy to check that the embedding is dense. Therefore, we can continue with the proof above replacing \mathbb{T}_{κ^+} with $\mathbb{T}_{\kappa^+}^{int}$.

Now consider the κ^+ -coherent theory whose axioms are those of \mathbb{T}_{κ^+} . By the completeness theorem of section 3.2, this theory admits a jointly conservative set of κ^+ -coherent prime theories in a new language where we add κ^+ many new constants, each of which is obtained from \mathbb{T}_{κ^+} by adding κ^+ -coherent sentences as axioms. We build now a poset of such prime theories in extended languages intended to form a Kripke model, as explained in section 3.4. We make however a modification to the construction: for each prime theory Γ over a set \mathbf{c} of κ^+ many

constants we consider the theory $\Gamma' := \Gamma \cup \{d = c_i \vdash \perp\}_{i < \kappa^+}$ where d is a new constant, and for this latter we compute the conservative prime theories over the extended language and define them as the theories at the successor nodes.

At the level of the syntactic categories, $\mathcal{C}_{\Gamma'}$ is the (pseudo-)colimit of $\mathcal{C}_{\Gamma_j} := \mathcal{C}_{\Gamma} / (\{\}, \bigwedge_{i < j} d \neq c_i)$ in the 2-category of κ^+ -coherent theories and κ^+ -coherent morphisms. Since this colimit is in fact a κ^+ -Heyting category, it will be equivalent to the (pseudo-)colimit computed in the 2-category of κ^+ -Heyting theories and κ^+ -Heyting morphisms.

At limit ordinals, we consider the union of the theories corresponding to each branch; this might not be prime, but if consistent, we extend it to prime theories to form the nodes of the next level of the poset. For the resulting frame to be a Kripke model, we need to prove that Γ' is conservative over Γ . For this, in turn, it is enough to show that any κ^+ -coherent model of \mathbb{T}_{κ^+} has a proper κ^+ -pure extension. We claim that, moreover, we can always find a proper κ^+ -Boolean extension of any model M of internal size κ^+ . To prove it, suppose otherwise. Then the topos $\mathbf{Set}[\mathbb{T}_{\kappa^+}^B]_{\kappa^{++}} / [M, -] \cong \mathbf{Set}^{\mathcal{K}_{\kappa^+}^B} / [M, -] \cong \mathbf{Set}^{M/\mathcal{K}_{\kappa^+}^B}$, (where $\mathcal{K}_{\kappa^+}^B$ consists of the models in \mathcal{K}_{κ^+} and all its κ^+ -Boolean homomorphisms, and where $\mathbb{T}_{\kappa^+}^B$ is \mathbb{T}_{κ^+} plus all those instances of excluded middle for κ^+ -coherent formulas), would be two-valued and Boolean. On the other hand, we have a stable surjection $\mathbf{Set}^{M/\mathcal{K}_{\kappa^+}^B} \twoheadrightarrow \mathbf{Set}^{M/\mathcal{K}_{\kappa^+}}$; this can be seen by considering first the stable surjection $\mathbf{Set}^{\mathcal{K}_{\kappa^+}^B} \cong \mathbf{Set}[\mathbb{T}_{\kappa^+}^B]_{\kappa^{++}} \twoheadrightarrow \mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^{++}} \cong \mathbf{Set}^{\mathcal{K}_{\kappa^+}}$ (which, at the level of the theories that they κ^{++} -classify, just adds instances of excluded middle for κ^+ -coherent formulas). Then we consider the pullback functor to the slice $\mathbf{Set}^{\mathcal{K}_{\kappa^+}} \twoheadrightarrow \mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^{++}} / [M, -] \cong \mathbf{Set}^{M/\mathcal{K}_{\kappa^+}}$, which is a geometric morphism along whose direct image we take the following (pseudo-)pullback:

$$\begin{array}{ccc} \mathbf{Set}^{\mathcal{K}_{\kappa^+}^B} & \xrightarrow{\quad} & \mathbf{Set}^{\mathcal{K}_{\kappa^+}} \\ \uparrow & & \uparrow \\ \mathbf{Set}[\mathbb{T}_{\kappa^+}^B]_{\kappa^{++}} / [M, -] & \xrightarrow{\quad} & \mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\kappa^{++}} / [M, -] \end{array}$$

Then the (pseudo-)pullback is precisely $\mathbf{Set}[\mathbb{T}_{\kappa^+}^B]_{\kappa^{++}} / [M, -] \cong \mathbf{Set}^{M/\mathcal{K}_{\kappa^+}^B}$, as can be verified using the universal property of the slice. This proves that indeed $\mathbf{Set}^{M/\mathcal{K}_{\kappa^+}^B} \twoheadrightarrow \mathbf{Set}^{M/\mathcal{K}_{\kappa^+}}$ is a stable surjection. Since being two-valued and Boolean is equivalent to having no proper non-degenerate subtoposes, by an argument analogous to that of the proof of Corollary 6.2, we would conclude that, since $\mathbf{Set}^{M/\mathcal{K}_{\kappa^+}^B}$ is Boolean (whence atomic and a coproduct of two-valued Boolean toposes), so would $\mathbf{Set}^{M/\mathcal{K}_{\kappa^+}}$ be, which is absurd since M is not maximal (which is in turn a consequence of categoricity and amalgamation). This finishes the proof of our claim.

Next, we claim that in the poset of theories there is a branch of height κ^{++} consisting of consistent theories whose union, in the logic $\mathcal{L}_{\kappa^{+++}, \kappa^{++}}$, is still consistent. To see this, note first that the corresponding poset of prime theories can be presented as a Kripke model of the internal κ^+ -Heyting theory¹³ of the category $(\mathcal{C}_{\mathbb{T}_{\kappa^+}})_{\kappa^+}$ (where the underlying set of each node consists of the constants in the prime theory, cf. section 3.4). This follows from the property that a prime theory proves (from axioms in the internal κ^+ -Heyting theory) any κ^+ -Heyting formula evaluated in constants \mathbf{c} if and only if that formula is forced at \mathbf{c} by the corresponding node of the tree. In turn, this property follows by a straightforward induction on the complexity of that formula

¹³When defining the fragment, we make sure to include all κ^+ -coherent formulas which arise in the κ^+ -classifying topos as instances of universal quantification, so that the syntactic category is κ^+ -Heyting. Then we can define its internal κ^+ -Heyting theory as the set of all κ^+ -Heyting sequents, in the same signature, which are true in the category.

using that the successor theories are jointly conservative over a given node in the extended language, as in section 3.4. Now each prime theory is determined by \mathbb{T}_{κ^+} and the conjunction of the rest of its axioms $\phi_f(\mathbf{c})$ (we can assume these axioms are atomic sentences using the primeness of the theory). We claim that $\phi_f \vdash_{\mathbf{y}_f} \neg \neg \bigvee_{g \in \gamma^{\beta+1}, g|_\beta = f} \exists \mathbf{x}_g \phi_g$ in $(\mathcal{C}_{\mathbb{T}_{\kappa^+}})_{\kappa^{++}}$. To see this, note that for any model of $\phi_f(\mathbf{c})$ and any embedding into some other model, this latter is not maximal, and thus has an embedding into a model of $\phi_f(\mathbf{c}) \wedge \bigwedge_i d \neq c_i$; then, a straightforward Löwenheim-Skolem argument (cf. section 3.3) shows that such model satisfying $\phi_f(\mathbf{c}) \wedge \bigwedge_i d \neq c_i$ for some \mathbf{c} must satisfy $\phi_g(\mathbf{d})$ for some g and \mathbf{d} , i.e., it satisfies $\bigvee_{g \in \gamma^{\beta+1}, g|_\beta = f} \exists \mathbf{x}_g \phi_g$. Whence, in $\mathbf{Set}^{\mathcal{K}_{\kappa^+}}$ we see that any model forcing $\phi_f(\mathbf{c})$ must force $\neg \neg \bigvee_{g \in \gamma^{\beta+1}, g|_\beta = f} \exists \mathbf{x}_g \phi_g$. This shows that indeed the premise $\phi_f \vdash_{\mathbf{y}_f} \neg \neg \bigvee_{g \in \gamma^{\beta+1}, g|_\beta = f} \exists \mathbf{x}_g \phi_g$ (that of the transfinite transitivity rule up to double negation) holds in $(\mathcal{C}_{\mathbb{T}_{\kappa^+}})_{\kappa^{++}}$. Since the premises in the transfinite transitivity rule for limit ordinals, up to double negation, also hold, trivially, the conclusion of the rule tells us that the theories of the minimal nodes of any bar are, up to a double negation, jointly conservative over the theory at the root (we are using that the sequent $\bigwedge_{i < \kappa^+} \neg \neg \phi_i \vdash_{\mathbf{x}} \neg \neg \bigwedge_{i < \kappa^+} \phi_i$ holds in $(\mathcal{C}_{\mathbb{T}_{\kappa^+}})_{\kappa^{++}}$, a consequence of \mathcal{K}_{κ^+} having directed colimits). This proves that a branch of height κ^{++} has to exist, since otherwise there would be a bar composed of inconsistent theories, while ϕ_\emptyset is consistent.

By repeating the argument with the consistent κ^{++} -coherent theory corresponding to a node of height κ^{++} (but now in the logic $\mathcal{L}_{\kappa^{+++}, \kappa^{++}}$), we can build a new Kripke model up to κ^{+++} , and so on. At limit cardinals $\rho = \sup_n \rho_n$ (with ρ_n successor cardinals), note that $\mathbf{Set}[\mathbb{T}_{\kappa^+}]_\rho := \lim \mathbf{Set}[\mathbb{T}_{\kappa^+}]_{\rho_n} \cong \lim \mathbf{Set}^{\mathcal{K}_{\geq \kappa^+, \leq \rho_n - 1}}$ and a similar computation to the one above shows that consistent theories of cardinality ρ do exist. We perform for them a similar construction, since we can adapt the completeness theorem (as mentioned at the end of section 3.2) for those theories in $\mathcal{L}_{\rho^+, \rho}$ when the axioms and the conclusions are in ρ -fragments of $\mathcal{L}_{\rho^+, \rho}$ (which are, by definition, unions of all ρ_n -fragments); cf. the last paragraph of section 3.2. The corresponding conservative extensions of each model of size ρ follows now from the fact that any model of size ρ has a proper $< \rho$ -pure extension, as can be proven similarly to the arguments above.

Eventually, we reach a consistent theory of cardinality λ , and we can construct its corresponding Kripke model. Now the theory at the root proves that there are λ distinct elements. Since the model of \mathbb{T}_{κ^+} of size λ satisfies $\forall \mathbf{x}(\vartheta \longrightarrow \psi)$, this sentence will be forced in the Kripke model, and so will be proved by the corresponding prime theory at the root node. This shows that $\neg \forall \mathbf{x}(\vartheta \longrightarrow \psi)$ cannot be 1 in $(\mathcal{C}_{\mathbb{T}_{\kappa^+}})_{\kappa^+}$, as otherwise all successive consistent theories of each cardinality over which we constructed the Kripke models would prove $\neg \forall \mathbf{x}(\vartheta \longrightarrow \psi)$, in particular the prime theory at the root node of the latest Kripke model. Therefore, as $(\mathcal{C}_{\mathbb{T}_{\kappa^+}})_{\kappa^+}$ is two-valued, $\forall \mathbf{x}(\vartheta \longrightarrow \psi)$ has to be 1 there, as we wanted to show.

As a last point, note that the same proof above which allowed us to conclude categoricity in κ^+ from categoricity in κ also let us prove categoricity in κ^{++} and so on, by an inductive procedure. Categoricity in a δ which is a limit cardinal is easily handled knowing that \mathcal{K} will be categorical at all γ_i for a cofinal sequence of successors $\kappa < \gamma_i < \delta$. Indeed, since \mathcal{K} is γ_i -categorical, the model of size γ_i is γ_i -saturated, which allows us to successively find a set of compatible isomorphisms between submodels of any two models of size δ (using directed colimits at limit steps), proving that they are indeed isomorphic (see [Ros97]). \square

Remark 9.2. *The Kripke models built during the proof of Theorem 9.1 allow to find the parallel with Proposition 4.27 in [SM90], established there with the use of stability theory. More specifically, if we have categoricity in κ but not in κ^+ (i.e., there is a κ -saturated model which is not κ^+ -saturated), we have that some formula $\neg \forall \mathbf{x}(\vartheta \longrightarrow \psi)$ is forced at the root of every Kripke model whose root node has cardinality μ ; whence we can conclude that there is a non-saturated*

model in each cardinality $\mu > \kappa^+$.

We will prove in the next section how to replace *GCH* with *SCH* in the general case of accessible categories with directed colimits, and to remove it completely in the case of AEC's. Our main result, in view of Theorem 9.1, is the following:

Theorem 9.3. (*Shelah's eventual categoricity conjecture for accessible categories with directed colimits*). Assume *SCH*, and let \mathcal{K} be an accessible category with directed colimits. Then there exists a cardinal μ_0 such that if \mathcal{K} is categorical in some $\lambda \geq \mu_0$, it is categorical in all $\lambda' \geq \mu_0$.

Proof. It is enough to take μ_0 the maximum of the Hanf numbers for categoricity and non-categoricity. \square

Corollary 9.4. (*Morley's categoricity theorem for infinitary theories*) Assume *SCH*, and let \mathbb{T} be a $\mathcal{L}_{\kappa,\theta}$ -theory whose category of models has directed colimits. Then there exists a cardinal μ_0 such that if \mathbb{T} is categorical in some $\lambda \geq \mu_0$ in S , it is categorical in all $\lambda' \geq \mu_0$ in S .

Example 9.5. In the case in which we consider the cardinality of the underlying model instead of the internal size, the exceptions of Corollary 9.4 are necessary. Indeed, generalizing a result of [LRV19], the category of μ -Hilbert spaces is defined as follows. Consider a μ -field \mathbf{R} , that is a field of hyperreals containing all ordinals up to μ . The construction of such a μ -field proceeds with the following steps:

- Take the initial segment of the ordinals up to μ . The natural (Hessenberg) sum and product is defined setting $a + b$ (resp. $a \cdot b$) as the maximum order type of a linear order extending the partial order given by the disjoint union (resp. the direct product). They are associative, commutative and the product distributes over the sum. At each following step, the sum and product operations can be defined similarly to the construction of the real numbers.
- Build the corresponding ring of μ -integers as pairs of ordinals (a, b) .
- Build the field of fractions of that ring.
- Take the μ -completion of that field considering all μ -Cauchy μ -sequences of fractions.

A μ -Hilbert space is then a Hilbert space over the μ -field \mathbf{R} . The category of μ -Hilbert spaces and isometries is then axiomatizable in $\mathcal{L}_{\mu^+, \mu^+}$ (e.g., generalizing the axiomatization of Hilbert spaces described in [LRV19]). Given an orthonormal base of size λ , each element of the μ -Hilbert space has at most μ nonzero coordinates. As a result, the cardinality is of the form λ^μ . Assuming *GCH* (or merely *SCH*), we have:

$$\lambda^\mu = \begin{cases} \lambda & \text{if } cf(\lambda) > \mu \text{ and } 2^\mu < \lambda \\ \lambda^+ & \text{if } cf(\lambda) \leq \mu \text{ and } 2^\mu < \lambda \end{cases}$$

It follows from this that there are no models at cardinals of cofinality less than μ , while eventually there is exactly one μ -Hilbert space of cardinality λ whenever λ is regular but not a successor of a cardinal of cofinality less than μ , and there are two μ -Hilbert spaces (of internal sizes λ and λ^+) if λ is such a successor. (On the other hand, it is categorical in every λ with respect to internal size).

In the context of AEC's, we get:

Corollary 9.6. (*Shelah's eventual categoricity conjecture for AEC's*). *Let \mathcal{K} be an AEC. Then there exists a cardinal μ_0 such that if \mathcal{K} is categorical in some $\lambda > \mu_0$, it is categorical in all $\lambda' > \mu_0$.*

Proof. It is enough to note that internal and external sizes coincide, so we can use the result of Theorem 9.3. \square

10 Removing GCH

So far we have proven our results assuming that *GCH* holds, which in almost every case was needed to guarantee that every regular cardinal κ satisfies $\kappa^{<\kappa} = \kappa$. Our method of proof is such that in several cases we can use forcing to make this cardinal equality true by collapsing $\kappa^{<\kappa}$ to κ without affecting the main properties of the models of cardinality less than κ . As a consequence, in many of the results on AEC's, *GCH* can be removed, while for μ -AEC's it can be downgraded to *SCH*. This general strategy is illustrated in the following:

Theorem 10.1. *In all the previous applications, Theorem 4.1 holds without any cardinal arithmetic assumptions.*

Proof. The assumptions we have used were $\kappa^{<\kappa} = \kappa$ and $\lambda^{<\lambda} = \lambda$. The first one was there only to guarantee that the cardinality of formulas of the appropriate κ -fragment of $\mathcal{L}_{\kappa^+, \kappa}$ was less than λ . Since \mathcal{K} is axiomatizable in $\mathcal{L}_{\kappa, \theta}(\Sigma)$ (for a signature Σ with only binary relations), it is enough to require $\kappa^{<\theta} < \lambda$, but as we are using this theorem above the Hanf number for model existence, this is always the case.

To handle the second assumption, let $\lambda = \gamma^+$. Consider the forcing extension $V[G]$ in which we collapse $\lambda^{<\lambda}$ to λ . Since this forcing is $< \lambda$ -distributive, models of size less than λ , and their embeddings remain unchanged (we assume they are properly coded by ordinals). Moreover, subsets definable by λ -geometric formulas of $\mathcal{L}_{\infty, \lambda}$ in $V[G]$ remain definable in V . Indeed, any λ -geometric formula $\phi(\mathbf{x})$ is of the form $\bigvee_{j < 2^\gamma} \exists_{i < \gamma} x_i \bigwedge_{i < \gamma} \psi_i^j$, where the ψ_i^j are atomic formulas with free variables amongst $\{x_k : k < i\} \cup \mathbf{x}$ and $\mathbf{x} = \bigcup_{i < \gamma} \mathbf{x}_i^j$ is a union (not necessarily disjoint) of the sets of free variables of each ψ_i^j with $|\mathbf{x}_i^j| = 2$. The formula defines a subset $S \subseteq \bigcup_{i < \gamma} |M|^{\mathbf{x}_i^j}$ in a model M if and only if there is a function $f : S \rightarrow V$ whose range is precisely the set of subformulas $\{\exists_{i < \gamma} x_i \bigwedge_{i < \gamma} \psi_i^j : j < 2^\gamma\}$ such that each $s \in S$ satisfies $f(s)$. Since the model is coded by ordinals (i.e., the underlying set, the functions and relations are all coded by ordinals less than γ), so can S be coded by ordinals less than γ , and thus f is in V and the subset is definable in the ground model.

We can now apply the theorem and deduce that the λ -classifying topos of \mathbb{T}' is equivalent to the presheaf topos $\mathbf{Set}^{Mod_\lambda(\mathbb{T})}$, which reduces to say that the following two conditions are satisfied:

1. The evaluations ev_ϕ in λ -geometric formulas form a generator of $\mathbf{Set}_{V[G]}^{Mod_\lambda(\mathbb{T})}$
2. The evaluation functor $ev : \mathcal{C}_{\mathbb{T}'} \rightarrow \mathbf{Set}_{V[G]}^{Mod_\lambda(\mathbb{T})}$ is full on subobjects (here $\mathcal{C}_{\mathbb{T}'}$ is the full syntactic category, not just the one restricted to the λ -fragment)

Let us see that both these statements can be expressed as assertions that only involve the existence of functions $f_i : \gamma \rightarrow V$; then, by $< \lambda$ -distributivity of the forcing, it follows that the functions f_i are actually in V , and therefore the statements hold already in the ground model.

Condition 1 boils down to saying that for each functor $F : Mod_\lambda(\mathbb{T}) \rightarrow \mathbf{Set}_{V[G]}$ and proper subfunctor $S \rightarrow F$ there is a λ -geometric formula ϕ and maps $\eta_M : ev_\phi(M) \rightarrow F(M)$ satisfying

naturality requirements such that η does not factor through S . In particular, this is true for all functors $F : \text{Mod}_\lambda(\mathbb{T}) \rightarrow \mathbf{Set}_V$. But in that case, since the underlying sets of the models, their functions and relations are all coded by ordinals less than γ , when $F(M) \subset V$ it follows from $< \lambda$ -distributivity that η_M is in V as well. Thus, the evaluations in λ -geometric formulas generate $\text{Set}_V^{\text{Mod}_\lambda(\mathbb{T})}$ and condition 1 is true in the ground model.

Condition 2 is easier to handle, since it is equivalent to saying that any distinguished subset of each λ -presentable model (which must have size less than λ) that is preserved by model homomorphisms is definable by a λ -geometric formula, which is a statement involving definability of fixed subsets of (the underlying sets of) these models, and since those subsets are unchanged by the forcing extension, the statement remains true in the ground model. \square

In analogous ways, several uses of *GCH* that have been made so far can be eliminated:

- The use of *GCH* in Lemma 8.1 comes from the equality $(\kappa^+)^{\kappa} = \kappa^+$, whose only purpose was to ensure that the completeness theorem for $\mathcal{L}_{\kappa^{++}, \kappa^+}$ can be applied, at the same time that the $\text{Diag}^+(N)$ and $\text{Diag}^+(S)$ contain at most $\kappa^{<\theta} \leq \kappa^+$ atomic formulas. To remove the need for *GCH* in this case, we proceed with the same strategy as in Theorem 10.1. That is, if the equality does not hold, we consider the forcing extension $V[G]$ in which we collapse $(\kappa^+)^{\kappa}$ to κ^+ . This forcing is $< \kappa^+$ -distributive and thus it does not change the category \mathcal{K}_κ . Whence, since we know that \mathcal{K}_κ satisfies amalgamation in $V[G]$, it already satisfies it in V . An analogous argument allows to conclude in Theorem 8.2 amalgamation at κ even when $\lambda^{<\lambda} > \lambda$, and to eliminate the uses of *GCH* in Theorem 8.3. The rest of the uses of *GCH* in Theorem 9.1, with one exception that we treat separately below, can be dealt with by using a forcing argument collapsing $(\lambda^+)^{\lambda}$ to λ^+ to derive the downward categoricity transfer.
- Likewise, the use of *GCH* in Theorem 7.1 can be avoided since it was only there to justify the applicability of Theorem 4.1, where it is not needed by Theorem 10.1. Also, in Lemma 6.4, *GCH* is used for the equality $(\lambda^+)^{\lambda} = \lambda^+$, but note that we can conclude λ -closedness without *GCH*, by similar forcing arguments (so that Lemma 6.6 does not need *GCH* either). On the other hand, the use of the equality $\kappa^{<\kappa} = \kappa$ in Lemma 6.3 was to guarantee the applicability of Corollary 6.2, where it was crucial to prove the Booleanness of the topos $\mathbf{Set}[\mathbb{T}_\kappa]_\kappa$, and to make sure that cardinality coincides with internal size. But this latter condition holds in cardinals in S , and the proof of that corollary does not need that equality though to prove that the topos is two-valued, so using this fact we can prove that the morphisms are ω -pure by using the following argument: we know that the topos $\mathbf{Set}^{\mathcal{K}_\lambda}$ is two-valued, and since each model of size λ embeds into a ω -closed model of size λ , $\mathbf{Set}^{\mathcal{K}_\lambda}$ forces $\neg \neg \forall \mathbf{x}(\phi \vee \neg \phi)$ and thus $\forall \mathbf{x}(\phi \vee \neg \phi)$, from which our claim follows.

Finally, consider the use of *GCH* from the proof of Theorem 9.1 for the Kripke completeness theorem. To remove it, define $f(\alpha) = \alpha^{<\alpha}$ and $C(\alpha) = \sup\{f^n(\alpha) : n \in \mathbb{N}\}$. Note that $C(\alpha) \times C(\alpha)$ has a well-ordering which is the union of the canonical well-orderings of each $f^n(\alpha) \times f^n(\alpha)$ for $n \geq 1$; this can be used in the completeness proof of section 3.2 to find for each α -coherent theory a set of prime $C(\alpha)$ -coherent theories in a language extended with $C(\alpha)$ many constants and which are jointly conservative over the original theory in the logic $\mathcal{L}_{(C(\alpha))^+, C(\alpha)}$. It suffices then to modify the proof as follows: we define each slice $\mathcal{C}' \cong \mathcal{C}_{\mathbb{T} \cup \phi(\mathbf{c})}$ as before, but now in the definition of the set of jointly epic families of morphisms over each object of the syntactic category, we make sure to put first the set of α many covers $\{C_i \rightarrow A\}_{i < \alpha}$ generated, as an α -Grothendieck topology, from axioms of the theory, and when considering the

set of sections $s_B : B \longrightarrow A'$ which are used in the transfinite construction of section 3.2, we make sure to put first those arrows, if any, represented by α -coherent formulas. We then use a well-ordering $\alpha \times \alpha \longrightarrow \alpha$ to get as the first α elements of each family $\mathcal{F}(B)$ the pullbacks of the first α coverings along the first α sections.

In each of the term models obtained from \mathcal{C}' by the transfinite construction of section 3.2 we choose a submodel of size α containing \mathbf{c} (this is possible due to the Löwenheim-Skolem property). Each of these submodels M determines, by the Löwenheim-Skolem theorem of section 3.3, a branch b and a further submodel M_b which belongs to the original set of term models, and has cardinality α . Let B_0 be the set of all such branches, and let T_0 be the subtree formed by the branches in B_0 . Since any model of \mathcal{C}' has, as a submodel, some term model of some branch in B_0 , the morphisms from \mathcal{C}' to the slices over elements of any bar in T_0 must be jointly conservative. We claim that, moreover, the set of sections $s_B : B \longrightarrow A'$ which were used in the transfinite construction have cardinality at most α if B is a node in T_0 . The reason for this latter claim is that we have restricted ourselves to accessible categories where all morphisms are monomorphisms, so that the formula $\forall xy(x = y \vee x \neq y)$ is forced by any object B of \mathcal{C}' ; thus, two sections $s_B : B \longrightarrow A'$ which are different will represent different arrows in the pseudocolimit.

Let now start with the syntactic category of the theory in α -coherent logic, and take its slice \mathcal{C}'' . Perform the transfinite construction as before but only up to α , using only the α covers $\{C_i \longrightarrow A\}_{i < \alpha}$ arising from axioms of the theory and the α sections $s_B : B \longrightarrow A'$ represented by α -coherent formulas. It follows that the set of covering families $\mathcal{F}(B)$ over each B in T_0 contains, as the first α elements, only those pullbacks, along the α sections $s_B : B \longrightarrow A'$ represented by α -coherent formulas, of the α covers $\{C_i \longrightarrow A\}_{i < \alpha}$ arising from axioms of the theory. Then, the term models produced by this construction at level α in each branch of T_0 will have cardinality at most α ; likewise, their α -coherent theories will also have α many axioms, allowing for the construction to proceed. In this way we can get a jointly conservative set of α -coherent prime theories over a language extended with α many constants. This finishes the proof that *GCH* can be removed in the particular case of AEC's, and downgraded to *SCH* in general accessible categories with directed colimits.

Remark 10.2. *The argument of the previous paragraphs can be alternatively used to show how to avoid GCH while using the completeness theorem of section 3.2 for all the partial results on μ -AEC's, replacing it with SCH and, in the particular case of AEC's, eliminating it completely. As can be seen from the argument, the key is the availability of the downward Löwenheim-Skolem property, which allows us to circumvent cardinal arithmetic assumptions. These can be considered, in a sense, to be built in into the Löwenheim-Skolem property.*

References

- [AF13] Steve Awodey and Henrik Forssell, *First-order logical duality*, Annals of Pure and Applied Logic **164** (2013), no. 3, 319–348.
- [BGL⁺16] Will Boney, Rami Grossberg, Michael Lieberman, Jiří Rosický, and Sebastien Vasey, *μ -abstract elementary classes and other generalizations*, Journal of Pure and Applied Algebra **220** (2016), no. 9, 3048–3066.
- [BJ98] Carsten Butz and Peter Johnstone, *Classifying toposes for first-order theories*, Annals of Pure and Applied Logic **91** (1998), no. 1, 33–58.
- [BR12] Tibor Beke and Jiří Rosický, *Abstract elementary classes and accessible categories*, Annals of Pure and Applied Logic **163** (2012), no. 12, 2008–2017.

- [Esp19] Christian Espíndola, *Infinitary first-order categorical logic*, Annals of Pure and Applied Logic **170** (2019), no. 2, 137–162.
- [Esp20] ———, *Infinitary generalizations of Deligne’s completeness theorem*, The Journal of Symbolic Logic **85** (2020), no. 3, 1147–1162.
- [Jec03] Thomas Jech, *Set theory. The Third Millenium Edition, revised and expanded*, Springer Monographs in Mathematics. Springer-Verlag, Berlin (2003).
- [Joh79] Peter Johnstone, *Conditions related to De Morgan’s law*, Applications of sheaves, Springer, 1979, pp. 479–491.
- [Kar64] Carol Karp, *Languages with expressions of infinite length*, North-Holland Publishing Company, 1964.
- [Kel89] Gregory Maxwell Kelly, *Elementary observations on 2-categorical limits*, Bulletin of the Australian Mathematical Society **39** (1989), no. 2, 301–317.
- [LRV19] Michael Lieberman, Jiří Rosický, and Sebastien Vasey, *Internal sizes in μ -abstract elementary classes*, Journal of Pure and Applied Algebra **223** (2019), no. 10, 4560–4582.
- [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.
- [Moe88] Ieke Moerdijk, *The classifying topos of a continuous groupoid. i*, Transactions of the American Mathematical Society **310** (1988), no. 2, 629–668.
- [Ros97] Jiří Rosický, *Accessible categories, saturation and categoricity*, The Journal of Symbolic Logic **62** (1997), no. 3, 891–901.
- [She09] Saharon Shelah, *Classification theory for abstract elementary classes*, Studies in Logic: Mathematical logic and foundations, vol. 18, College Publications, 2009.
- [SM90] Saharon Shelah and Michael Makkai, *Categoricity of theories in $\mathcal{L}_{\kappa,\omega}$, with κ a compact cardinal*, Annals of Pure and Applied Logic **47** (1990), no. 1, 41–97.
- [SS18] Shelah Saharon and Vasey Sebastien, *Categoricity and multidimensional diagrams*, <https://arxiv.org/abs/1805.06291> (2018).
- [Vas18] Sebastien Vasey, *The categoricity spectrum of large abstract elementary classes*, <https://arxiv.org/abs/1805.04068> (2018).