Available online at www.sciencedirect.com
SCIENCE DIRECT
ELSEVIER
Annals of Pure and Applied Logic 141 (2006) 108–147
ANNALS OF PURE AND APPLIED LOGIC
www.elsevier.com/locate/apal

Categoricity in abstract elementary classes with no maximal models

Monica VanDieren

Department of Mathematics, University of Michigan, Ann Arbor MI 48109, United States

Received 8 August 2003; received in revised form 13 September 2005; accepted 4 October 2005
Available online 20 December 2005
Communicated by A.J. Wilkie

Abstract

The results in this paper are in a context of abstract elementary classes identified by Shelah and Villaveces in which the amalgamation property is not assumed. The long-term goal is to solve Shelah's Categoricity Conjecture in this context. Here we tackle a problem of Shelah and Villaveces by proving that in their context, the uniqueness of limit models follows from categoricity under the assumption that the subclass of amalgamation bases is closed under unions of bounded, $\prec_{\mathcal{K}}$-increasing chains.

Keywords: Model theory; Abstract elementary classes; Categoricity; Amalgamation; Limit models; Weak diamond

Introduction

The origins of much of pure model theory can be traced back to Łoś' Conjecture [15]. This conjecture was resolved by M. Morley in his Ph.D. thesis in 1962 [17]. Morley then questioned the status of the conjecture for uncountable theories. Building on work of W. Marsh, F. Rowbottom and J.P. Ressayre, S. Shelah proved the statement for uncountable theories in 1970 [19]. Out of Morley and Shelah's proofs the program of *stability theory* or *classification theory* evolved.

While first-order logic has far reaching applications to other fields of mathematics, there are several interesting frameworks which cannot be captured by first-order logic. A classification theory for non-elementary classes will open the door potentially to a multitude of applications of model theory to classical mathematics and provide insight into first-order model theory.

Shelah posed a generalization of Łoś' Conjecture to $L_{\omega_1,\omega}$ as a test question to measure progress in non-first-order model theory. Focus on non-elementary classes began to shift in the late seventies when Shelah, influenced by B. Jónsson's work in universal algebra (see [12,13]), identified the notion of *abstract elementary class (AEC)* to capture many non-first-order logics [24] including $L_{\omega_1,\omega}$ and $L_{\omega_1,\omega}(\mathbf{Q})$. An abstract elementary class is a class of structures of the same similarity type endowed with a morphism satisfying natural properties such as closure under directed limits.

Definition 0.1. $\mathcal{K}$ is an *abstract elementary class (AEC)* iff $\mathcal{K}$ is a class of models for some vocabulary which is denoted by $L(\mathcal{K})$, and the class is equipped with a partial order, $\preceq_{\mathcal{K}}$ satisfying the following:

(1) Closure under isomorphisms.

E-mail address: mvd@umich.edu.

0168-0072/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.apal.2005.10.006

(a) For every $M \in \mathcal{K}$ and every $L(\mathcal{K})$-structure N if $M \cong N$ then $N \in \mathcal{K}$.
(b) Let $N_1, N_2 \in \mathcal{K}$ and $M_1, M_2 \in \mathcal{K}$ such that there exist $f_l : N_l \cong M_l$ (for $l = 1, 2$) satisfying $f_1 \subseteq f_2$ then $N_1 \prec_{\mathcal{K}} N_2$ implies that $M_1 \prec_{\mathcal{K}} M_2$.

(2) $\preceq_{\mathcal{K}}$ refines the submodel relation.
(3) If $M_0, M_1 \preceq_{\mathcal{K}} N$ and M_0 is a submodel of M_1, then $M_0 \preceq_{\mathcal{K}} M_1$.
(4) (Downward Löwenheim–Skolem Axiom) There is a *Löwenheim–Skolem number of* $\mathcal{K}$, denoted $LS(\mathcal{K})$ which is the minimal κ such that for every $N \in \mathcal{K}$ and every $A \subset N$, there exists M with $A \subseteq M \prec_{\mathcal{K}} N$ of cardinality $\kappa + |A|$.
(5) If $\langle M_i \mid i < \delta\rangle$ is a $\prec_{\mathcal{K}}$-increasing and chain of models in $\mathcal{K}$
(a) $\bigcup_{i<\delta} M_i \in \mathcal{K}$,
(b) for every $j < \delta$, $M_j \prec_{\mathcal{K}} \bigcup_{i<\delta} M_i$ and
(c) if $M_i \prec_{\mathcal{K}} N$ for every $i < \delta$, then $\bigcup_{i<\delta} M_i \prec_{\mathcal{K}} N$.

Definition 0.2. For $M, N \in \mathcal{K}$ a monomorphism $f : M \to N$ is called a $\prec_{\mathcal{K}}$*-embedding* or a $\prec_{\mathcal{K}}$*-mapping* iff $f[M] \preceq_{\mathcal{K}} N$.

Notation 0.3. We write $\mathcal{K}_\mu := \{M \in \mathcal{K} \mid \|M\| = \mu\}$.

Remark 0.4. The Hanf number of $\mathcal{K}$ will be formally defined in Definition I.3.8. It is bounded by $\beth_{(2^{2^{LS(\mathcal{K})}})^+}$.

Shelah extended his categoricity conjecture for $L_{\omega_1,\omega}$-theories in the following form in [29], see also [26]:

Conjecture 0.5 (*Shelah's Categoricity Conjecture*). *Let* $\mathcal{K}$ *be an abstract elementary class. If* $\mathcal{K}$ *is categorical in some* $\lambda \geq \text{Hanf}(\mathcal{K})$, *then for every* $\mu \geq \text{Hanf}(\mathcal{K})$, $\mathcal{K}$ *is categorical in* μ.

Definition 0.6. We say $\mathcal{K}$ *is categorical in* λ whenever there exists exactly one model in $\mathcal{K}$ of cardinality λ up to isomorphism.

Despite the existence of over 1000 published pages of partial results towards this conjecture, it remains open. Since the mid-eighties, model theorists have approached Shelah's conjecture from two different directions (see [6] for a short history). Shelah, M. Makkai and O. Kolman attacked the conjecture with set theoretic assumptions [16,14,30]. On the other hand, Shelah also looked at the conjecture under model theoretic assumptions in [28,31,32]. The approach of Shelah and A. Villaveces in [33] involved a balance between set theoretic and model theoretic assumptions. This paper further investigates the context of [33] which we delineate here:

Assumption 0.7. (1) $\mathcal{K}$ is an AEC with no maximal models with respect to the relation $\prec_{\mathcal{K}}$,
(2) $\mathcal{K}$ is categorical in some fixed $\lambda \geq \text{Hanf}(\mathcal{K})$,
(3) GCH holds and
(4) a form of the weak diamond holds, namely $\Phi_{\mu^+}(S^{\mu^+}_{\text{cf}(\mu)})$ holds for every μ with $\mu < \lambda$ (see Definition I.3.2).

The purpose of [33] was to begin investigating the conjecture that the amalgamation property follows from categoricity in a large enough cardinality. All of the other attempts to prove Conjecture 0.5 have made use of the assumption of the amalgamation property which is a sufficient condition to define a reasonable notion of (Galois)-type (see Section 1).

Definition 0.8. Let $\mathcal{K}$ be an abstract elementary class and μ a cardinal $\geq \text{LS}(\mathcal{K})$.

(1) We say that $M \in \mathcal{K}_\mu$ is an *amalgamation base* if for every $N_1, N_2 \in \mathcal{K}_\mu$ and $g_i : M \to N_i$ for $(i = 1, 2)$, there are $\prec_{\mathcal{K}}$-embeddings f_i, $(i = 1, 2)$ and a model N such that the following diagram commutes:

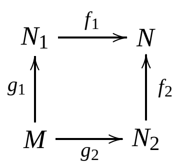

(2) An abstract elementary class $\mathcal{K}$ satisfies the *amalgamation property* iff every $M \in \mathcal{K}$ is an amalgamation base.

(3) We write $\mathcal{K}^{am}$ for the class of amalgamation bases which are in $\mathcal{K}$. We also use $\mathcal{K}^{am}_\mu$ to denote the class of amalgamation bases of cardinality μ.

Remark 0.9. (1) The definition of amalgamation base varies across the literature. Our definition of amalgamation base is weaker than an alternative formulation which does not put any restriction on the cardinality of N_1 and N_2. Under the assumption of the amalgamation property, these definitions are known to be equivalent. However, in this context, where the amalgamation property is not assumed, we cannot guarantee the existence of the stronger form of amalgamation bases.

(2) We get an equivalent definition of amalgamation base, if we additionally require that $g_i \restriction M = id_M$ for $i = 1, 2$, in the definition above. See [7] for details.

It is conjectured that categoricity in a large enough cardinality implies the amalgamation property. However, there are examples of abstract elementary classes which are categorical in ω successive cardinals, but fail to have the amalgamation property in larger cardinalities [11,34]. Shelah constructs an abstract elementary class whose models are bipartite random graphs. Models of cardinality $\aleph_1$ in this class witness the failure of amalgamation. Intriguingly, under the assumption of Martin's Axiom, this class of bipartite graphs is categorical in $\aleph_0$ and $\aleph_1$. On the other hand, if one assumes a version of the weak diamond, Shelah proves that categoricity in $\aleph_0$ and $\aleph_1$ implies amalgamation in $\aleph_1$ ([24] or see [6] for an exposition). There are other natural examples of abstract elementary classes which do not satisfy the amalgamation property but are unstable such as the class of locally finite groups [10].

Limited progress has been made to prove that amalgamation follows from categoricity. Kolman and Shelah manage to prove this for AECs that can be axiomatized by a $L_{\kappa,\omega}$ sentence with κ a measurable cardinal [14]. They first introduce limit models as a substitute for saturated models, and then prove the uniqueness of limit models (see Definition I.2.7).

To better understand the relationship between the amalgamation property, categoricity and the uniqueness of limit models, consider the questions of uniqueness and existence of limit models in classes which satisfy the amalgamation property, but not are not necessarily categorical:

Remark 0.10. Even under the amalgamation property, the uniqueness and existence of limit models do not come for free. The existence requires stability (see [32] or [8]). The question of uniqueness of limit models is tied into (super)stability as well. Even in first-order logic, the uniqueness of limit models fails for un-superstable theories (see [8] or [28] for examples). The uniqueness of limit models has been proven in AECs under the assumption of categoricity ([14,28], and here, Theorem II.9.1). Recently Grossberg, VanDieren and Villaveces identified sufficient conditions (which are consequences of superstability) for the uniqueness of limit models in classes with the amalgamation property [9].

The motivation for this paper is to elaborate on recent work of Shelah and Villaveces in which they strive to prove under weaker assumptions than Kolman and Shelah that the amalgamation property follows from categoricity above the Hanf number. The first step in proving amalgamation is to show the uniqueness of limit models.

The uniqueness of limit models under Assumption 0.7 generalizes Theorem 6.5 of [28] where Shelah assumes the full amalgamation property. The amalgamation property is used in [28] in several forms including the fact that saturated models are model homogeneous and that all reducts of Ehrenfeucht–Mostowski models are amalgamation bases. Shelah then uses the uniqueness of limit models to prove that the union of a chain of μ-saturated models is μ-saturated, provided that the chain is of length $< \mu^+$. This is one of the main steps in proving a downward categoricity transfer theorem for classes with the amalgamation property.

In the Fall of 1999, we identified several problems with Shelah and Villaveces' proof of the uniqueness of limit models from [33]. After two years of correspondence, Shelah and Villaveces conceded that they were not able to resolve these problems. While these issues are undertaken in this paper, to date the proof of the uniqueness of limit models has resisted a complete solution under Assumption 0.7. After presenting a partial solution (Theorem II.9.1) of the uniqueness of limit models and discussing this with Shelah at a Mid-Atlantic Mathematical Logic Seminar in the Fall of 2001, we were not able to remove the extra hypothesis. The extra hypothesis was weakened in [35]. This paper provides a complete proof of an intermediate uniqueness result patching a gap that was found in [35] in the Fall of 2002. The partial solution to the uniqueness of limit models described here is in the context identified in [33] (Assumption 0.7) under the hypothesis:

Hypothesis 1: Every continuous tower inside $\mathfrak{C}$ has an amalgamable extension inside $\mathfrak{C}$ (see Sections 2 and 5 for the definitions).

Remark 0.11. The model $\mathfrak{C}$ in Hypothesis 1 is not the usual monster model. It is a weak substitute for the monster model and is introduced in Section 2. Monster models, as we know them in first-order logic, are model homogeneous. In the absence of the amalgamation property, model homogeneous models may not exist.

In the context of [33], Hypothesis 1 is a consequence of the more natural Hypothesis 2 (see Section 10).

Hypothesis 2: For $\mu < \lambda$, the class of amalgamation bases of cardinality μ (denoted by $\mathcal{K}^{am}_\mu$) is closed under unions of $\prec_\mathcal{K}$-increasing chains of length $< \mu^+$.

It seems reasonable to consider a weakening of Grossberg's Intermediate Categoricity Conjecture which captures Hypothesis 2:

Conjecture 0.12. *Let $\mathcal{K}$ be an AEC. If there exists a $\lambda \geq \text{Hanf}(\mathcal{K})$ such that $\mathcal{K}$ is categorical in λ, then $\mathcal{K}^{am}_\mu$ is closed under unions of length $< \mu^+$ for all μ with $LS(\mathcal{K}) \leq \mu < \lambda$.*

Although Theorem 1.11 of Chapter 4 in [25] addresses a similar problem to Hypothesis 2, this statement may be too ambitious to prove. An alternative hypothesis which also implies Hypothesis 1 is

Hypothesis 3: The union of a $\prec_\mathcal{K}$-increasing chain of length $< \mu^+$ of limit models of cardinality μ is a limit model.

Hypothesis 3 may be more approachable as it is a relative of the first-order consequence of superstability that the union of a $\prec$-increasing chain of $\kappa(T)$-many saturated models is saturated.

Hypothesis 1 has relatives in the literature as well. Indeed, in [24] where the amalgamation property is not assumed, Shelah identifies the link between the existence of maximal elements of $\mathcal{K}^3_{\aleph_0}$ (a specialization of towers of length 1) and $2^{\aleph_1}$ non-isomorphic models in $\aleph_1$.

This paper is divided into three parts outlined below.

Part I. The first part summarizes the necessary definitions and background material. It also includes some new results on μ-splitting.

Section 1 Galois types
Section 2 Limit models
Section 3 Limit models are amalgamation bases
Section 4 μ-splitting
Section 5 Towers

Part II. Here we provide a complete proof of the uniqueness of limit models under Hypothesis 1 and Assumption 0.7.

Section 6 Relatively full towers
Section 7 Continuous $<^c_{\mu,\alpha}$-extensions
Section 8 Refined orderings on towers
Section 9 Uniqueness of limit models

Part III. In this part of the paper we include a partial result in the direction of Hypothesis 1 and discuss reduced towers.

Section 10 $<^c_{\mu,\alpha}$-Extension property for nice towers
Section 11 Reduced towers

Part I. Preliminaries

Throughout this paper, unless otherwise stated, we will make Assumption 0.7 and μ will be a cardinal satisfying $LS(\mathcal{K}) \leq \mu < \lambda$ where λ is the categoricity cardinal.

We introduce the necessary definitions and background from [33]. The reader familiar with [33] may skim through Section 2 where the monster model is introduced and then proceed to Section 4 which includes some new results on μ-splitting.

1. Galois types

In this section we discuss problems that arise when working without the amalgamation property in AECs. The first obstacle is to identify a reasonable notion of type. Because of the category-theoretic definition of abstract elementary classes, the first-order notion of formulas and types cannot be applied. To overcome this barrier, Shelah has suggested identifying types, not with formulas, but with the orbit of an element under the group of automorphisms fixing a given structure. In order to carry out this definition of type, the following binary relation E must be an equivalence relation on triples (a, M, N). In order to avoid confusing this new notion of "type" with the conventional one (i.e. set of formulas) we will follow [6] and [7] and introduce it below under the name of *Galois type*.

Definition I.1.1. For triples $(\bar{a}_l, M_l, N_l)$ where $\bar{a}_l \in N_l$ and $M_l \preceq_{\mathcal{K}} N_l \in \mathcal{K}$ for $l = 1, 2$, we define a binary relation E as follows: $(\bar{a}_1, M_1, N_1)E(\bar{a}_2, M_2, N_2)$ iff $M := M_1 = M_2$ and there exists $N \in \mathcal{K}$ and $\prec_{\mathcal{K}}$-mappings f_1, f_2 such that $f_l : N_l \to N$ and $f_l \restriction M = id_M$ for $l = 1, 2$ and $f_1(\bar{a}_1) = f_2(\bar{a}_2)$:

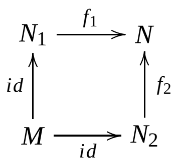

To prove that E is an equivalence relation (more specifically, that E is transitive), we need to restrict ourselves to amalgamation bases.

Remark I.1.2. E is an equivalence relation on the set of triples of the form $(\bar{a}, M, N)$ where $M \preceq_{\mathcal{K}} N$, $\bar{a} \in N$ and $M, N \in \mathcal{K}^{am}_\mu$ for fixed $\mu \geq LS(\mathcal{K})$. To see that E is transitive, consider $(a_1, M, N_1)E(a_2, M, N_2)$ and $(a_2, M, N_2)E(a_3, M, N_3)$ where $M, N_1, N_2, N_3 \in \mathcal{K}^{am}_\mu$. Let $N_{1,2}$ and f_1, f_2 be such that $f_1 : N_1 \to N_{1,2}$; $f_2 : N_2 \to N_{1,2}$ and $f_1 \restriction M = f_2 \restriction M = id_M$ with $f_1(a_1) = f_2(a_2)$. Similarly define g_2, g_3 and $N_{2,3}$ with $g_2(a_2) = g_3(a_3)$. By the Downward Löwenheim–Skolem Axiom, we may assume that $N_{1,2}$ and $N_{2,3}$ have cardinality μ. Consider the following diagram of this situation.

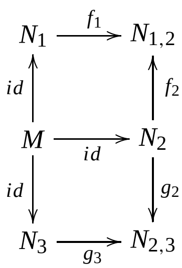

Since N_2 was chosen to be an amalgamation base, we can amalgamate $N_{1,2}$ and $N_{2,3}$ over N_2 with mappings h_1 and h_3 and an amalgam N^* giving us the following diagram:

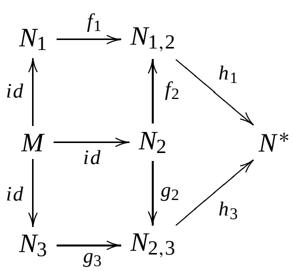

Notice that $h_1(f_1(a_1)) = h_3(g_3(a_3))$. Thus $h_1 \circ f_1$ and $h_3 \circ g_3$ witness that $(a_1, M, N_1)E(a_3, M, N_3)$.

Remark I.1.3 (*Invariance*). If M is an amalgamation base and f is a $\prec_{\mathcal{K}}$-embedding, then $f(M)$ is an amalgamation base.

In AECs with the amalgamation property, we are often limited to speak of types only over models. Here we are further restricted to deal with types only over models which are amalgamation bases.

Definition I.1.4. Let $\mu \geq LS(\mathcal{K})$ be given.

(1) For $M, N \in \mathcal{K}^{am}_\mu$ with $M \preceq_{\mathcal{K}} N$ and $\bar{a} \in {}^{\omega>}|N|$, the *Galois type of $\bar{a}$ in N over M*, written ga-tp$(\bar{a}/M, N)$, is defined to be $(\bar{a}, M, N)/E$.
(2) For $M \in \mathcal{K}^{am}_\mu$,

$$\text{ga-S}^1(M) := \{\text{ga-tp}(a/M, N) \mid M \preceq_{\mathcal{K}} N \in \mathcal{K}^{am}_\mu, a \in N\}.$$

(3) We say $p \in$ ga-S(M) *is realized in M'* whenever $M \prec_{\mathcal{K}} M'$ and there exist $\bar{a} \in M'$ and $N \in \mathcal{K}^{am}_\mu$ such that $p = (\bar{a}, M, N)/E$.
(4) For $M' \in \mathcal{K}^{am}_\mu$ with $M \prec_{\mathcal{K}} M'$ and $q = \text{ga-tp}(\bar{a}/M', N) \in \text{ga-S}(M')$, we define *the restriction of q to M* as $q \restriction M := \text{ga-tp}(\bar{a}/M, N)$.
(5) For $M' \in \mathcal{K}^{am}_\mu$ with $M \prec_{\mathcal{K}} M'$, we say that $q \in$ ga-S(M') *extends* $p \in$ ga-S(M) iff $q \restriction M = p$.
(6) $p \in$ ga-S(M) is said to be *non-algebraic* if no $a \in M$ realizes p.

Notation I.1.5. We will often abbreviate a Galois type, ga-tp$(a/M, N)$ as ga-tp(a/M), when the role of N is not crucial or is clear. This occurs mostly when we are working inside of a fixed structure $\mathfrak{C}$, which we define in Section 2.

Fact I.1.6 (*See [7]*). *When $\mathcal{K} = \text{Mod}(T)$ for T a complete first-order theory, the above definition of* ga-tp$(a/M, N)$ *coincides with the classical first-order definition where c and a have the same type over M iff for every first-order formula $\varphi(x, \bar{b})$ with parameters $\bar{b}$ from M,*

$$N \models \varphi(c, \bar{b}) \text{ iff } N \models \varphi(a, \bar{b}).$$

We will now define Galois stability in an analogous way:

Definition I.1.7. We say that $\mathcal{K}$ is *Galois stable in μ* if for every $M \in \mathcal{K}^{am}_\mu$, $|\text{ga-S}^1(M)| = \mu$.

Fact I.1.8 (*Fact 2.1.3 of [33]*). *If $\mathcal{K}$ is categorical in λ, then for every $\mu < \lambda$, we have that $\mathcal{K}$ is Galois stable in μ.*

By combining results from [33,8] and [2] it is possible to improve this to conclude Galois stability in some cardinals $\geq \lambda$, but it remains open whether or not in AECs categoricity implies Galois stability in all cardinalities above $LS(\mathcal{K})$.

Definition I.1.9. Let $\mu > LS(\mathcal{K})$, M is said to be *μ-saturated* if for every $N \prec_{\mathcal{K}} M$ with $N \in \mathcal{K}^{am}_{<\mu}$ and every Galois type p over N, we have that p is realized in M.

The following fact is proved by showing the equivalence of model homogeneous models and saturated models in classes which satisfy the amalgamation property [31].

Fact I.1.10. *Suppose that $\mathcal{K}$ satisfies the amalgamation property. If M_1 and $M_2 \in \mathcal{K}_\mu$ are μ-saturated and there exists $N \prec_{\mathcal{K}} M_1, M_2$ with $N \in \mathcal{K}_{<\mu}$, then $M_1 \cong M_2$.*

Since we will be working in a context where the amalgamation property is not assumed, we do not have the uniqueness of saturated models at hand. In fact even the existence of saturated models is questionable. The purpose of this paper is to identify a suitable substitute for saturation that is unique up to isomorphism in every cardinality. The candidate is the limit model discussed in the following section. Later we will give an alternative characterization of limit models as the union of a relatively full tower (see Section 6). This characterization plays the role of $\mathbf{F}^a_\kappa$-saturated models from first-order model theory (see Chapter IV of [26]).

2. Limit models

In this section we define limit models and discuss their uniqueness and existence. A local substitute for the monster model is also introduced.

We begin with universal extensions which are central in the definition of limit models. A universal extension captures some properties of saturated models without referring explicitly to types. The notion of universality over countable models was first analyzed by Shelah in Theorem 1.4(3) of [22].

Definition I.2.1. (1) Let κ be a cardinal $\geq LS(\mathcal{K})$. We say that N is *κ-universal over M* iff for every $M' \in \mathcal{K}_\kappa$ with $M \prec_{\mathcal{K}} M'$ there exists a $\prec_{\mathcal{K}}$-embedding $g : M' \to N$ such that $g \restriction M = id_M$:

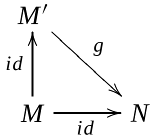

(2) We say N is *universal over M* or N is *a universal extension of M* iff N is $\|M\|$-universal over M.

Notation I.2.2. In diagrams, we will indicate that N is universal over M, by writing $M \overset{id}{\Longrightarrow} N$.

Remark I.2.3. Notice that the definition of N *universal over M* requires all extensions of M of cardinality $\|M\|$ to be embeddable into N. First-order variants of this definition in the literature often involve $\|M\| < \|N\|$. We will be considering the case when $\|M\| = \|N\|$.

Remark I.2.4. Suppose that T is a first-order complete theory that is stable in some regular μ. Then every model M of T of cardinality μ has an elementary extension N of cardinality μ which is universal over M. To see this, define an elementary-increasing and continuous chain of models of T of cardinality μ, $\langle N_i \mid i < \mu \rangle$ such that N_{i+1} realizes all types over N_i. Let $N = \bigcup_{i<\mu} N_i$. By a back-and-forth construction, one can show that N is universal over M.

The existence of universal extensions in AECs follows from categoricity in λ and GCH or categoricity and uses the presentation of the model of cardinality λ as a reduct of an EM-model.

Fact I.2.5 (*Theorem 1.3.1 from [33]*). *Let μ be such that $LS(\mathcal{K}) \leq \mu < \lambda$. Then every element of $\mathcal{K}^{am}_\mu$ has a universal extension in $\mathcal{K}^{am}_\mu$.*

Another existence result that does not use GCH or categoricity can be proved under the assumption of Galois stability and the amalgamation property ([32] or see [8] for a proof).

Notice that the following observation asserts that it is unreasonable to prove a stronger existence statement than Fact I.2.5, without having proved the amalgamation property.

Proposition I.2.6. *If $M \in \mathcal{K}_\mu$ has a universal extension, then M is an amalgamation base.*

We can now define the principal concept of this paper:

Definition I.2.7. For $M', M \in \mathcal{K}_\mu$ and σ a limit ordinal with $\sigma < \mu^+$, we say that M' is a *(μ, σ)-limit over M* iff there exists a $\prec_{\mathcal{K}}$-increasing and continuous sequence of models $\langle M_i \in \mathcal{K}_\mu \mid i < \sigma \rangle$ such that

(1) $M = M_0$,
(2) $M' = \bigcup_{i<\sigma} M_i$
(3) for $i < \sigma$, M_i is an amalgamation base and
(4) M_{i+1} is universal over M_i.

Remark I.2.8. (1) Notice that in Definition I.2.7, for $i < \sigma$ and i a limit ordinal, M_i is a (μ, i)-limit model.
(2) Notice that Condition (3) implies Condition (4) of Definition I.2.7. In our constructions, since the question of whether a particular model is an amalgamation base becomes crucial, we choose to list this as a separate condition.

Definition I.2.9. We say that M' is a *(μ, σ)-limit* iff there is some $M \in \mathcal{K}$ such that M' is a (μ, σ)-limit over M.

While limit models were used is [14] and [28], their use extends to other contexts. There is evidence that the uniqueness of limit models provides a basis for the development of a notion of non-forking and a stability theory for abstract elementary classes. Limit models are used in [8] to develop the notion of non-splitting in tame, Galois-stable AECs. The uniqueness of limit models implies the existence of superlimits in [31]. Additionally, in [32] the uniqueness of limit models appears as an axiom for good frames and the limit models are closely related to brimmed models. In all of these applications, limit models provide a substitute for Galois-saturated models.

By repeated applications of Fact I.2.5, the existence of (μ, ω)-limit models can be proved:

Fact I.2.10 (*Theorem 1.3.1 from [33]*). *Let μ be a cardinal such that $\mu < \lambda$. For every $M \in \mathcal{K}^{am}_{\mu}$, there is a (μ, ω)-limit over M.*

In order to extend this argument further to yield the existence of (μ, σ)-limits for arbitrary limit ordinals $\sigma < \mu^+$, we need to be able to verify that limit models are in fact amalgamation bases. We will examine this in Section 3.

While the existence of limit models can be derived from the categoricity and weak diamond assumptions, the uniqueness of limit models is more difficult. Here we recall two easy uniqueness facts which state that limit models of the same length are isomorphic. They are proved using the natural back-and-forth construction of an isomorphism.

Fact I.2.11 (*Fact 1.3.6 from [33]*). *Let $\mu \geq LS(\mathcal{K})$ and $\sigma < \mu^+$. If M_1 and M_2 are (μ, σ)-limits over M, then there exists an isomorphism $g : M_1 \to M_2$ such that $g \restriction M = id_M$. Moreover if M_1 is a (μ, σ)-limit over M_0; N_1 is a (μ, σ)-limit over N_0 and $g : M_0 \cong N_0$, then there exists a $\prec_{\mathcal{K}}$-mapping, $\hat{g}$, extending g such that $\hat{g} : M_1 \cong N_1$.*

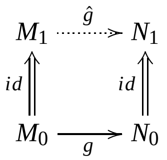

Fact I.2.12 (*Fact 1.3.7 from [33]*). *Let μ be a cardinal and σ a limit ordinal with $\sigma < \mu^+ \leq \lambda$. If M is a (μ, σ)-limit model, then M is a $(\mu, cf(\sigma))$-limit model.*

A more challenging uniqueness question is to prove that two limit models of different lengths ($\sigma_1 \neq \sigma_2$) are isomorphic:

Conjecture I.2.13. *Suppose that $\mathcal{K}$ is categorical in some $\lambda \geq$ Hanf$(\mathcal{K})$ and μ is a cardinal with $LS(\mathcal{K}) \leq \mu < \lambda$. Let σ_1 and σ_2 be limit ordinals $< \mu^+$. Suppose M_1 and M_2 are (μ, σ_1)- and (μ, σ_2)-limits over M, respectively. Then M_1 is isomorphic to M_2 over M.*

The main result of this paper, Theorem II.9.1, is a solution to this conjecture under Assumption 0.7 and Hypothesis 1.

We will need one more notion of limit model, which will later serve as a substitute for a monster model. This is a natural extension of the limit models already defined:

Definition I.2.14. Let μ be a cardinal $< \lambda$, we say that $\check{M}$ is a (μ, μ^+)-*limit over* M iff there exists a $\prec_{\mathcal{K}}$-increasing and continuous chain of models $\langle M_i \in \mathcal{K}^{am}_{\mu} \mid i < \mu^+ \rangle$ such that $M_0 = M$, $\bigcup_{i<\mu^+} M_i = \check{M}$, and for $i < \mu^+$, M_{i+1} is universal over M_i.

Remark I.2.15. While it is known that in our context (μ, θ)-limit models are amalgamation bases when $\theta < \mu^+$, it is open whether or not (μ, μ^+)-limits are amalgamation bases. To avoid confusion between these two concepts of limit models, we will denote (μ, μ^+)-limit models with a $\check{}$ above the model's name (i.e. $\check{M}$). Later we will avoid this confusion by fixing a (μ, μ^+)-limit model and denoting it by $\mathfrak{C}$, since it will substitute the usual notion of a monster model.

The existence of (μ, μ^+)-limit models follows from the fact that (μ, θ)-limit models are amalgamation bases when $\theta < \mu^+$, see Corollary I.3.14. The uniqueness of (μ, μ^+)-limit models (Corollary I.2.20) can be shown using an easy back-and-forth construction as in the proof of Fact I.2.11.

The following theorem indicates that (μ, μ^+)-limits provide some level of homogeneity. First we recall an exercise regarding amalgamation.

Remark I.2.16. Suppose that M_0, M_1 and M_2 can be amalgamated, then by renaming elements, we can choose the amalgam to be a $\prec_{\mathcal{K}}$-extension of M_2.

Theorem I.2.17. *If $\check{M}$ is a (μ, μ^+)-limit, then for every $N \prec_{\mathcal{K}} \check{M}$ with $N \in \mathcal{K}^{am}_\mu$, we have that $\check{M}$ is universal over N. Moreover, $\check{M}$ is a (μ, μ^+)-limit over N.*

Proof. Suppose that $\check{M}$ is a (μ, μ^+)-limit model and $N \prec_{\mathcal{K}} \check{M}$ is such that $N \in \mathcal{K}^{am}_\mu$. Let N' be an extension of N of cardinality μ. Let $\langle M_i \mid i < \mu^+\rangle$ witness that $\check{M}$ is a (μ, μ^+)-limit model. Since N has cardinality μ, there exists $i < \mu^+$, such that $N \prec_{\mathcal{K}} M_i$. Since N is an amalgamation base, we can amalgamate M_i and N' over N with amalgam $M' \in \mathcal{K}_\mu$. By Remark I.2.16, we may assume that $M_i \prec_{\mathcal{K}} M'$.

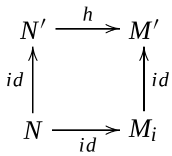

Since M_{i+1} is universal over M_i, there is $g : M' \to M_{i+1}$ such that $g \restriction M_i = id_{M_i}$. Then $g \circ h$ give us the desired mapping from N' into $\check{M}$ over N.

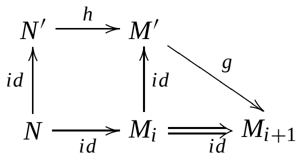
□

Remark I.2.18. If N is not an amalgamation base, then there are no universal models over N.

It is immediate that $\mathfrak{C}$ realizes many types:

Corollary I.2.19. *For every $M \in \mathcal{K}^{am}_\mu$ with $M \prec_{\mathcal{K}} \mathfrak{C}$, we have that $\mathfrak{C}$ is saturated over M.*

Corollary I.2.20. *Suppose $\check{M}_1$ and $\check{M}_2$ are (μ, μ^+)-limits over $M_1, M_2 \in \mathcal{K}^{am}_\mu$, respectively. If there exists an isomorphism $h : M_1 \cong M_2$, then h can be extended to an isomorphism $g : \check{M}_1 \cong \check{M}_2$.*

Since (μ, μ^+)-limit models are unique and are universal over all amalgamation bases of cardinality μ, they are in some sense homogeneous. We will see that if $\check{M}$ is a (μ, μ^+)-limit model and $\text{ga-tp}(a/M, \check{M}) = \text{ga-tp}(b/M, \check{M})$, then there is an automorphism f of $\check{M}$ fixing M such that $f(a) = b$ (Corollary I.2.25). In some ways, (μ, μ^+)-limit models behave like monster models in first-order logic if we restrict ourselves to amalgamation bases and models of cardinality μ. This justifies the following notation.

Notation I.2.21. We fix a cardinal μ with $LS(\mathcal{K}) \leq \mu < \lambda$ and a (μ, μ^+)-limit model and denote it by $\mathfrak{C}$. For $M \prec_{\mathcal{K}} \mathfrak{C}$ we abbreviate

$$\{f \mid f \text{ is an automorphism of } \mathfrak{C} \text{ with } f \restriction M = id_M\}$$

by $\text{Aut}_M(\mathfrak{C})$.

While it is customary to work entirely inside of a fixed monster model $\mathfrak{C}$ in first-order logic, we will sometimes need to consider structures outside of $\mathfrak{C}$ since we do not have the full power of model homogeneity in this context.

We now recall a result from [33] which will be used in our proof of Corollary I.2.25. Although Shelah and Villaveces work without the amalgamation property as an assumption, using weak diamond they prove a weak amalgamation property, which they refer to as *density of amalgamation bases*.

Fact I.2.22 (*Theorem 1.2.4 from [33]*). *Every $M \in \mathcal{K}_{<\lambda}$ has a proper $\mathcal{K}$-extension of the same cardinality which is an amalgamation base.*

We can now improve Fact I.2.5 slightly. This improvement is used throughout this paper.

Lemma I.2.23. *For every μ with $LS(\mathcal{K}) \le \mu < \lambda$, if $M \in \mathcal{K}^{am}_\mu$, $N \in \mathcal{K}$ and $\bar{a} \in {}^{\mu^+>}|N|$ are such that $M \prec_{\mathcal{K}} N$, then there exists $M^{\bar{a}} \in \mathcal{K}^{am}_\mu$ such that $M^{\bar{a}}$ is universal over M and $M \bigcup \bar{a} \subseteq M^{\bar{a}}$.*

Proof. By Axiom 4 of AEC, we can find $M' \prec_{\mathcal{K}} N$ of cardinality μ containing $M \bigcup \bar{a}$. Applying Fact I.2.22, there exists an amalgamation base of cardinality μ, say M'', extending M'. By Fact I.2.5 we can find a universal extension of M'' of cardinality μ, say $M^{\bar{a}}$.

Notice that $M^{\bar{a}}$ is also universal over M. Why? Suppose M^* is an extension of M of cardinality μ. Since M is an amalgamation base we can amalgamate M'' and M^* over M. WLOG we may assume that the amalgam, M^{**}, is an extension of M'' of cardinality μ and a $\prec_{\mathcal{K}}$-mapping $f^* : M^* \to M^{**}$ with $f^* \restriction M = id_M$.

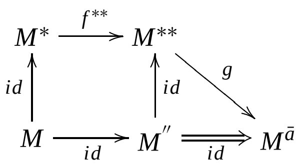

Now, since $M^{\bar{a}}$ is universal over M'', there exists a $\prec_{\mathcal{K}}$-mapping g such that $g : M^{**} \to M^{\bar{a}}$ with $g \restriction M'' = id_{M''}$. Notice that $g \circ f^*$ gives us the desired mapping of M^* into $M^{\bar{a}}$. □

Notice that Lemma I.2.23 is a step closer to proving that $\mathcal{K}^{am}$ satisfies Axiom 4 of the definition of AEC as it gives a weak downward Löwenheim–Skolem property. It is an open question whether or not $\mathcal{K}^{am}$ is an AEC.[1]

An alternative version of Lemma I.2.23 gives us

Lemma I.2.24. *Given amalgamation bases of cardinality μ, M_1 and M_2. If $M_1, M_2 \prec_{\mathcal{K}} \mathfrak{C}$, then there exists an amalgamation base $M' \prec_{\mathcal{K}} \mathfrak{C}$ of cardinality μ that is universal over both M_1 and M_2.*

Proof. Let $\langle M'_i \mid i < \mu^+\rangle$ witness that $\mathfrak{C}$ is a (μ, μ^+)-limit model. Then there exists $i < \mu^+$ such that $M_1, M_2 \prec_{\mathcal{K}} M'_i$. Notice that by choice of the sequence $\langle M'_j \mid j < \mu^+\rangle$, we have that M'_{i+1} is universal over M'_i. Thus, using the assumption that M_1 and M_2 are amalgamation bases, M'_{i+1} is universal over M_1 and M_2. □

The following is a corollary of Theorem I.2.17 and justifies our choice of notation for (μ, μ^+)-limit models.

Corollary I.2.25. *If* ga-tp$(a/M, \mathfrak{C})$ = ga-tp$(b/M, \mathfrak{C})$*, then there is an automorphism f of $\mathfrak{C}$ fixing M such that $f(a) = b$.*

Proof. Suppose that ga-tp$(a/M, \mathfrak{C})$ = ga-tp$(b/M, \mathfrak{C})$. By Theorem I.2.17, $\mathfrak{C}$ is a (μ, μ^+)-limit over M. Let $\langle M_i \in \mathcal{K}^{am}_\mu \mid i < \mu^+\rangle$ witness this. There exists an $i < \mu^+$ such that $a, b \in M_i$. Denote M_i by both M_a and M_b. By definition of types, there is a model N of cardinality μ and $\prec_{\mathcal{K}}$-mappings g, h such that $g(a) = h(b)$ and the following diagram commutes:

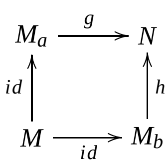

Notice that $\mathfrak{C}$ is universal over M_b. So there is a $\prec_{\mathcal{K}}$-mapping, $f' : N \to \mathfrak{C}$ such that the following diagram commutes:

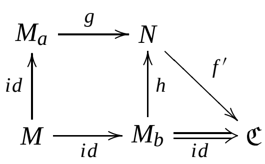

[1] The main difficulty is Axiom 5.

Consider $f' \circ g$. Notice that it is a partial automorphism with domain M_a. By Corollary I.2.20 applied to $f' \circ g(M_a)$ and M_a, the mapping $f' \circ g$ can be extended to an automorphism of $\mathfrak{C}$, call such an extension f. Then, $f \upharpoonright M = id_M$ and $f(a) = f' \circ g(a) = f'(h(b)) = b$, as required. □

3. Limit models are amalgamation bases

While Fact I.2.22 asserts the existence of amalgamation bases, it is useful to identify what other features are sufficient for a model to be an amalgamation base. Makkai and Shelah were able to prove that all existentially closed models are amalgamation bases for $L_{\kappa,\omega}$ theories with κ above a strongly compact cardinal (Corollary 1.6 of [16]). Kolman and Shelah identified a concept called *niceness* which implied amalgamation in categorical $L_{\kappa,\omega}$ theories with κ above a measurable cardinal. (Note: Their notion of niceness is not related to the notion of nice towers appearing in Section 5). They then showed that every model of cardinality $< \lambda$ was nice (see [14]). These results relied heavily on set theoretic assumptions.

In a more general context, Shelah and Villaveces state that every limit model is an amalgamation base (Fact 1.3.10 of [33]), using $\diamondsuit_{\mu^+}(S^{\mu^+}_{\mathrm{cf}(\mu)})$. For completeness, we provide a proof that every (μ,θ)-limit model with $\theta < \mu^+$ is an amalgamation base under a weaker version of diamond $(\Phi_{\mu^+}(S^{\mu^+}_{\mathrm{cf}(\mu)}))$. This is the content of Theorem I.3.13.

Let us first recall the set theoretic and model theoretic machinery necessary for the proof.

Definition I.3.1. Let θ be a regular ordinal $< \mu^+$. We denote

$$S^{\mu^+}_\theta := \{\alpha < \mu^+ \mid \mathrm{cf}(\alpha) = \theta\}.$$

The Φ-principle defined next is known as *Devlin and Shelah's weak diamond* [4].

Definition I.3.2. For μ a cardinal and $S \subseteq \mu^+$ a stationary set, the weak diamond, denoted by $\Phi_{\mu^+}(S)$, is said to hold iff for all $F : {}^{\mu^+>}2 \to 2$ there exists $g : \mu^+ \to 2$ such that for every $f : \mu^+ \to 2$ the set

$$\{\delta \in S \mid F(f \upharpoonright \delta) = g(\delta)\} \text{ is stationary.}$$

We will be using a consequence of $\Phi_{\mu^+}(S)$, called $\Theta_{\mu^+}(S)$ (see [7]).

Definition I.3.3. For μ a cardinal $S \subseteq \mu^+$ a stationary set, $\Theta_{\mu^+}(S)$ is said to hold if and only if for all families of functions

$$\{f_\eta \ : \ \eta \in {}^{\mu^+}2 \text{ where } f_\eta : \mu^+ \to \mu^+\}$$

and for every club $C \subseteq \mu^+$, there exist $\eta \neq \nu \in {}^{\mu^+}2$ and there exists a $\delta \in C \cap S$ such that

(1) $\eta \upharpoonright \delta = \nu \upharpoonright \delta$,
(2) $f_\eta \upharpoonright \delta = f_\nu \upharpoonright \delta$ and
(3) $\eta(\delta) \neq \nu(\delta)$.

The relative strength of these principles is provided below. See [7] for details.

Fact I.3.4. *For S a stationary subset of μ^+, $\diamondsuit_{\mu^+}(S) \Longrightarrow \Phi_{\mu^+}(S) \Longrightarrow \Theta_{\mu^+}(S)$.*

For most regular $\theta < \mu^+$, Fact I.3.4 and the following imply that $\Phi_{\mu^+}(S^{\mu^+}_\theta)$ follows from GCH:

Fact I.3.5 (*[5] for μ Regular and [21] for μ Singular*). *For every $\mu > \aleph_1$, GCH $\Longrightarrow \diamondsuit_{\mu^+}(S)$ where $S = S^{\mu^+}_\theta$ for any regular $\theta \neq \mathrm{cf}(\mu)$.*

Thus, from GCH and $\Phi_{\mu^+}(S^{\mu^+}_{\mathrm{cf}(\mu)})$ we have that $\Phi_{\mu^+}(S^{\mu^+}_\theta)$ holds for every regular $\theta < \mu^+$.

In addition to the weak diamond, we will be using Ehrenfeucht–Mostowski models. Let us recall some facts here.

The following gives a characterization of AECs as *PC*-classes. Fact I.3.7 is often referred to as Shelah's Presentation Theorem.

Definition I.3.6. A class $\mathcal{K}$ of structures is called a *PC-class* if there exists a language L_1, a first-order theory T_1 in the language L_1 and a collection of types without parameters, Γ, such that L_1 is an expansion of $L(\mathcal{K})$ and

$$\mathcal{K} = PC(T_1, \Gamma, L) := \{M \restriction L : M \models T_1 \text{ and } M \text{ omits all types from } \Gamma\}.$$

When $|T_1| + |L_1| + |\Gamma| + \aleph_0 = \chi$, we say that $\mathcal{K}$ is PC_χ. PC-classes are sometimes referred to as *projective classes* or *pseudo-elementary classes*.

Fact I.3.7 (*Lemma 1.8 of [24] or See [7]*). *If $(\mathcal{K}, \prec_{\mathcal{K}})$ is an AEC, then there exists $\chi \leq 2^{LS(\mathcal{K})}$ such that $\mathcal{K}$ is PC_χ.*

The representation of AECs as *PC*-classes allows us to construct Ehrenfeucht–Mostowski models if there are arbitrarily large models in our class.

Definition I.3.8. Given an AEC $\mathcal{K}$, we define the *Hanf number of* $\mathcal{K}$, abbreviated Hanf($\mathcal{K}$), as the minimal κ such that for every $PC_{2^{LS(\mathcal{K})}}$-class, $\mathcal{K}'$, if there exists a model $M \in \mathcal{K}'$ of cardinality κ, then there are arbitrarily large models in $\mathcal{K}'$.

Fact I.3.9 (*Claim 0.6 of [28] or See [7]*). *Assume that $\mathcal{K}$ is an AEC that contains a model of cardinality $\geq \beth_{(2^{2^{LS(\mathcal{K})}})^+}$. Then, there is a Φ, proper for linear orders,*[2] *such that for all linear orders $I \subseteq J$ we have that*

(1) $EM(I, \Phi) \restriction L(\mathcal{K}) \prec_{\mathcal{K}} EM(J, \Phi) \restriction L(\mathcal{K})$ *and*
(2) $\|EM(I, \Phi) \restriction L(\mathcal{K})\| = |I| + LS(\mathcal{K})$.

It is a theorem of C.C. Chang based on a theorem of Morley that Hanf($\mathcal{K}$) $\leq \beth_{(2^{2^{LS(\mathcal{K})}})^+}$ (see Section 4 of Chapter VII of [26]). Morley's proof [18] gives a better upper bound in certain situations: for a class $\mathcal{K}$ that is $PC_{\aleph_0}$, the Hanf number of $\mathcal{K}$ is $\leq \beth_{\omega_1}$.

In our context, since $\mathcal{K}$ has no maximal models, $\mathcal{K}$ has a model of cardinality Hanf($\mathcal{K}$). Then by Fact I.3.9, we can construct Ehrenfeucht–Mostowski models.

We describe an index set which appears often in papers about the categoricity conjecture. This index set appears in several places including [14,28] and [33].

Notation I.3.10. Let $\alpha < \lambda$ be given.

For $X \subseteq \alpha$, we define

$$I_X := \left\{\eta \in {}^{\omega}X : \{n < \omega \mid \eta(n) \neq 0\} \text{ is finite}\right\}.$$

The following fact is proved in several papers e.g. [33].

Fact I.3.11. *If $M \prec_{\mathcal{K}} EM(I_\lambda, \Phi) \restriction L(\mathcal{K})$ is a model of cardinality μ^+ with $\mu^+ < \lambda$, then there exists a $\prec_{\mathcal{K}}$-mapping $f : M \to EM(I_{\mu^+}, \Phi) \restriction L(\mathcal{K})$.*

A variant of this universality property is (implicit in Lemma 3.7 of [14] or see [1]):

Fact I.3.12. *Suppose κ is a regular cardinal. If $M \prec_{\mathcal{K}} EM(I_\kappa, \Phi) \restriction L(\mathcal{K})$ is a model of cardinality $< \kappa$ and $N \prec_{\mathcal{K}} EM(I_\lambda, \Phi) \restriction L(\mathcal{K})$ is an extension of M of cardinality $\|M\|$, then there exists a $\prec_{\mathcal{K}}$-embedding $f : N \to EM(I_\kappa, \Phi) \restriction L(\mathcal{K})$ such that $f \restriction M = id_M$.*

We now prove that limit models are amalgamation bases.

Theorem I.3.13. *Under Assumption 0.7, if M is a (μ, θ)-limit for some θ with $\theta < \mu^+ \leq \lambda$, then M is an amalgamation base.*

Proof. Given μ, suppose that θ is the minimal infinite ordinal $< \mu^+$ such that there exists a model M which is a (μ, θ)-limit and not an amalgamation base. Notice that by Fact I.2.12, we may assume that $\text{cf}(\theta) = \theta$. We assume that all models have as their universe a subset of μ^+.

[2] Also known as a blueprint, see Definition 2.5 of Chapter VII, Section 5 of [26] for a formal definition.

For this proof we will make use of the following notation. We will consider binary sequences ordered by initial segment and denote this ordering by $\lessdot$. For $\eta \in {}^{\alpha}2$ we use $l(\eta)$ as an abbreviation for the length of η, in this case $l(\eta) = \alpha$.

With the intention of eventually applying $\Theta_{\mu^+}(S^{\mu^+}_\theta)$, we will define a tree of structures $\langle M_\eta \in \mathcal{K}_\mu \mid \eta \in {}^{\mu^+>}2\rangle$ such that when $l(\eta)$ has cofinality θ, M_η will be a (μ, θ)-limit model and $M_{\eta\hat{}0}, M_{\eta\hat{}1}$ will witness that M_η is not an amalgamation base. After this tree of structures is defined we will embed each chain of models into a universal model of cardinality μ^+. We will apply $\Theta_{\mu^+}(S^{\mu^+}_\theta)$ to these embeddings. $\Theta_{\mu^+}(S^{\mu^+}_\theta)$ will provide an amalgam for $M_{\eta\hat{}0}$ and $M_{\eta\hat{}1}$ over M_η for some sequence η whose length has cofinality θ, giving us a contradiction.

In order to construct such a tree of models, we will need several conditions to hold throughout the inductive construction:

(1) $M \preceq_{\mathcal{K}} M_{\langle\rangle}$
(2) for $\eta \lessdot \nu \in {}^{\mu^+>}2$, $M_\eta \prec_{\mathcal{K}} M_\nu$
(3) for $l(\eta)$ a limit ordinal with $\mathrm{cf}(l(\eta)) \leq \theta$, $M_\eta = \bigcup_{\alpha<l(\eta)} M_{\eta\restriction\alpha}$
(4) for $\eta \in {}^{\alpha}2$ with $\alpha \in S^{\mu^+}_\theta$,
  (a) M_η is a (μ, θ)-limit model
  (b) $M_{\eta\hat{}\langle 0\rangle}, M_{\eta\hat{}\langle 1\rangle}$ cannot be amalgamated over M_η
  (c) $M_{\eta\hat{}\langle 0\rangle}$ and $M_{\eta\hat{}\langle 1\rangle}$ are amalgamation bases of cardinality μ
(5) for $\eta \in {}^{\alpha}2$ with $\alpha \notin S^{\mu^+}_\theta$,
  (a) M_η is an amalgamation base
  (b) $M_{\eta\hat{}\langle 0\rangle}, M_{\eta\hat{}\langle 1\rangle}$ are universal over M_η and
  (c) $M_{\eta\hat{}\langle 0\rangle}$ and $M_{\eta\hat{}\langle 1\rangle}$ are amalgamation bases of cardinality μ (it may be that $M_{\eta\hat{}\langle 0\rangle} = M_{\eta\hat{}\langle 1\rangle}$ in this case).

This construction is possible:

$\eta = \langle\rangle$: By Fact I.2.22, we can find $M' \in \mathcal{K}^{am}_\mu$ such that $M \prec_{\mathcal{K}} M'$. Define $M_{\langle\rangle} := M'$.

$l(\eta)$ is a limit ordinal: When $\mathrm{cf}(l(\eta)) > \theta$, let $M'_\eta := \bigcup_{\alpha<l(\eta)} M_{\eta\restriction\alpha}$. M'_η is not necessarily an amalgamation base, but for the purposes of this construction, continuity at such limits is not important. Thus by Fact I.2.22 we can find an extension of M'_η, say M_η, of cardinality μ such that M_η is an amalgamation base.

For η with $\mathrm{cf}(l(\eta)) \leq \theta$, we require continuity. Define $M_\eta := \bigcup_{\alpha<l(\eta)} M_{\eta\restriction\alpha}$. We need to verify that if $l(\eta) \notin S^{\mu^+}_\theta$, then M_η is an amalgamation base. In fact, we will show that such a M_η will be a $(\mu, \mathrm{cf}(l(\eta)))$-limit model. Let $\langle \alpha_i \mid i < \mathrm{cf}(l(\eta))\rangle$ be an increasing and continuous sequence of ordinals converging to $l(\eta)$ such that $\mathrm{cf}(\alpha_i) < \theta$ for every $i < \mathrm{cf}(l(\eta))$. Condition (5b) guarantees that for $i < \mathrm{cf}(l(\eta))$, $M_{\eta\restriction\alpha_{i+1}}$ is universal over $M_{\eta\restriction\alpha}$. Additionally, condition (3) ensures us that $\langle M_{\eta\restriction\alpha_i} \mid i < \mathrm{cf}(l(\eta))\rangle$ is continuous. This sequence of models witnesses that M_η is a $(\mu, \mathrm{cf}(l(\eta)))$-limit model. By our minimal choice of θ and our assumption that in this phase of the construction $\mathrm{cf}(l(\eta)) \lneq \theta$, we have that $(\mu, \mathrm{cf}(l(\eta)))$-limit models are amalgamation bases. Thus M_η is an amalgamation base.

$\eta\hat{}\langle i\rangle$ where $l(\eta) \in S^{\mu^+}_\theta$: We first notice that $M_\eta := \bigcup_{\alpha<l(\eta)} M_{\eta\restriction\alpha}$ is a (μ, θ)-limit model. Why? Since $l(\eta) \in S^{\mu^+}_\theta$ and θ is regular, we can find an increasing and continuous sequence of ordinals, $\langle \alpha_i \mid i < \theta\rangle$ converging to $l(\eta)$ such that for each $i < \theta$ we have that $\mathrm{cf}(\alpha_i) < \theta$. Condition (5b) of the construction guarantees that for each $i < \theta$, $M_{\eta\restriction\alpha_{i+1}}$ is universal over $M_{\eta\restriction\alpha_i}$. Thus $\langle M_{\eta\restriction\alpha_i} \mid i < \theta\rangle$ witnesses that M_η is a (μ, θ)-limit model.

Since M_η is a (μ, θ)-limit, we can fix an isomorphism $f : M \cong M_\eta$. By Remark I.1.3, M_η is not an amalgamation base. Thus there exist $M_{\eta\hat{}0}$ and $M_{\eta\hat{}1}$ extensions of M_η which cannot be amalgamated over M_η. WLOG, by the Density of Amalgamation Bases, we can choose $M_{\eta\hat{}\langle 0\rangle}$ and $M_{\eta\hat{}\langle 1\rangle}$ to be elements of $\mathcal{K}^{am}_\mu$.

$\eta\hat{}\langle i\rangle$ where $l(\eta) \notin S^{\mu^+}_\theta$: Since M_η is an amalgamation base, we can choose $M_{\eta\hat{}\langle 0\rangle}$ and $M_{\eta\hat{}\langle 1\rangle}$ to be extensions of M_η such that $M_{\eta\hat{}\langle l\rangle} \in \mathcal{K}^{am}_\mu$ and $M_{\eta\hat{}\langle l\rangle}$ is universal over M_η, for $l = 0, 1$.

This completes the construction. Let C be a club containing $\{\alpha < \mu^+ \mid M_\alpha$ has universe $\alpha\}$.

For every $\eta \in {}^{\mu^+}2$, define $M_\eta := \bigcup_{\alpha<\mu^+} M_{\eta\restriction\alpha}$. Notice that by condition (5b) of the construction, each M_η has cardinality μ^+. By categoricity in λ and Fact I.3.11, we can fix a $\prec_{\mathcal{K}}$-mapping $g_\eta : M_\eta \to EM(I_{\mu^+}, \Phi) \restriction L(\mathcal{K})$ for each $\eta \in {}^{\mu^+}2$. Now apply $\Theta_{\mu^+}(S^{\mu^+}_\theta)$ to find $\eta, \nu \in {}^{\mu^+}2$ and $\alpha \in S^{\mu^+}_\theta \cap C$ such that

· $\rho := \eta \restriction \alpha = \nu \restriction \alpha$,

· $\eta(\alpha) = 0$, $\nu(\alpha) = 1$ and
· $g_\eta \restriction M_\rho = g_\nu \restriction M_\rho$.

Let $N := EM(I_{\mu^+}, \Phi) \restriction L(\mathcal{K})$. Then the following diagram commutes:

$$\begin{array}{ccc} M_{\rho\hat{}\langle 1\rangle} & \xrightarrow{g_\nu \restriction M_{\rho\hat{}\langle 1\rangle}} & N \\ \uparrow{\scriptstyle id} & & \uparrow{\scriptstyle g_\eta \restriction M_{\rho\hat{}\langle 0\rangle}} \\ M_\rho & \xrightarrow[id]{} & M_{\rho\hat{}\langle 0\rangle} \end{array}$$

Notice that $g_\eta \restriction M_{\rho\hat{}\langle 0\rangle}$ and $g_\nu \restriction M_{\rho\hat{}\langle 1\rangle}$ witness that $M_{\rho\hat{}\langle 0\rangle}$ and $M_{\rho\hat{}\langle 1\rangle}$ can be amalgamated over M_ρ. Since $l(\rho) = \alpha \in S^{\mu^+}_\theta$, $M_{\rho\hat{}\langle 0\rangle}$ and $M_{\rho\hat{}\langle 1\rangle}$ were chosen so that they cannot be amalgamated over M_ρ. Thus, we contradict condition (4b) of the construction. □

Now that we have verified that limit models are amalgamation bases, we can use the existence of universal extensions to construct (μ, θ)-limit models for arbitrary $\theta < \mu^+$.

Corollary I.3.14 (*Existence of Limit Models*). *For every cardinal μ and limit ordinal θ with $\theta \leq \mu^+ \leq \lambda$, if M is an amalgamation base of cardinality μ, then there exists a (μ, θ)-limit over M.*

Proof. By repeated applications of Fact I.2.5 (existence of universal extensions) and Theorem I.3.13. □

In addition to the fact that limit models are amalgamation bases, we will use an even stronger amalgamation property of limit models. It is a result of Shelah and Villaveces. The argument provided is a simplification of the original and was suggested by J. Baldwin.

Fact I.3.15 (*Weak Disjoint Amalgamation [33]*). *Given $\lambda > \mu \geq LS(\mathcal{K})$ and $\alpha, \theta_0 < \mu^+$ with θ_0 regular. If M_0 is a (μ, θ_0)-limit and $M_1, M_2 \in \mathcal{K}_\mu$ are $\prec_{\mathcal{K}}$-extensions of M_0, then for every $\bar{b} \in {}^\alpha(M_1 \backslash M_0)$, there exist M_3, a model, and h, a $\prec_{\mathcal{K}}$-embedding, such that*

(1) *$h : M_2 \to M_3$;*
(2) *$h \restriction M_0 = id_{M_0}$ and*
(3) *$h(M_2) \cap \bar{b} = \emptyset$ (equivalently $h(M_2) \cap M_1 = M_0$).*

Proof. Let M_0, M_1 and M_2 be given as in the statement of the claim. First notice that we may assume that M_0, M_1 and M_2 are such that there is a $\delta < \mu^+$ with $M_0 = M_1 \cap (EM(I_\delta, \Phi) \restriction L(\mathcal{K}))$ and $M_1, M_2 \prec_{\mathcal{K}} EM(I_{\mu^+}, \Phi) \restriction L(\mathcal{K})$. Why? Define $\langle N_i \in \mathcal{K}_\mu \mid i < \mu^+ \rangle$ a $\prec_{\mathcal{K}}$-increasing and continuous chain of amalgamation bases such that

(1) $N_0 = M_0$ and
(2) N_{i+1} is universal over N_i.

Let $N_{\mu^+} = \bigcup_{i<\mu^+} N_i$. By categoricity and Fact I.3.11, there exists a $\prec_{\mathcal{K}}$-mapping f such that $f : N_{\mu^+} \to EM(I_{\mu^+}, \Phi) \restriction L(\mathcal{K})$. Consider the club $C = \{\delta < \mu^+ \mid f(N_{\mu^+}) \cap (EM(I_\delta, \Phi) \restriction L(\mathcal{K})) = f(N_\delta)\}$. Let $\delta \in C \cap S^{\mu^+}_{\mathrm{cf}(\theta_0)}$. Notice that $f(N_\delta)$ is a $(\mu, \mathrm{cf}(\theta_0))$-limit model. Since M_0 is also a $(\mu, \mathrm{cf}(\theta_0))$-limit model, there exists $g : M_0 \cong f(N_\delta)$. Since $f(N_{\delta+1})$ is universal over $f(N_\delta)$, we can extend g to g' such that $g' : M_1 \to f(N_{\delta+1})$ with $g'(M_1) \cap EM(I_\delta, \Phi) \restriction L(\mathcal{K}) = g'(M_0)$. Thus we may take M_0, M_1 and M_2 with $M_0 = M_1 \cap EM(I_\delta, \Phi) \restriction L(\mathcal{K})$.

Let δ be such that $M_1 \cap (EM(I_\delta, \Phi) \restriction L(\mathcal{K})) = M_0$ and let $\delta^* < \mu^+$ be such that $M_1, M_2 \prec_{\mathcal{K}} EM(I_{\delta^*}) \restriction L(\mathcal{K})$. Let h be the $\mathcal{K}$ mapping from $EM(I_{\delta^*}) \restriction L(\mathcal{K})$ into $EM(I_{\mu^+}, \Phi) \restriction L(\mathcal{K})$ induced by

$$\alpha \mapsto \delta^* + \alpha$$

for all $\alpha < \delta^*$.

We will show that if $b \in M_1 \backslash M_0$ then $b \notin h(M_2)$. Suppose for the sake of contradiction that $b \in M_1 \backslash M_0$ and $b \in h(M_2)$. Let τ be a Skolem term and let $\bar{\alpha}$, $\bar{\beta}$ be finite sequences such that $\bar{\alpha} \in I_\delta$ and $\bar{\beta} \in I_{\delta^*} \backslash I_\delta$, satisfying $b = \tau(\bar{\alpha}, \bar{\beta})$.

Since $b \in h(M_2)$, there exists a Skolem term σ and finite sequences $\bar{\alpha}' \in I_\delta$ and $\bar{\beta}' \in I_{\mu^+} \backslash I_{\delta^*}$ satisfying $b = \sigma(\bar{\alpha}', \bar{\beta}')$.

Since $\bar{\beta}'$ and $\bar{\beta}$ are disjoint, we can find $\bar{\gamma}'$ and $\bar{\gamma} \in I_\delta$ such that the type of $\bar{\beta}'\hat{\ }\bar{\beta}$ is the same as the type of $\bar{\gamma}'\hat{\ }\bar{\gamma}$ over $\bar{\alpha}'\hat{\ }\bar{\alpha}$ with respect to the lexicographical order of I_{μ^+}. Notice then that the type of $\bar{\beta}'$ and $\bar{\gamma}'$ over $\bar{\gamma}\hat{\ }\bar{\alpha}'\hat{\ }\bar{\alpha}$ are the same with respect to the lexicographical ordering.

Recall

$$EM(I_{\mu^+}, \Phi) \upharpoonright L(\mathcal{K}) \models b = \tau(\bar{\alpha}, \bar{\beta}) = \sigma(\bar{\alpha}', \bar{\beta}').$$

Thus

$$EM(I_{\mu^+}, \Phi) \upharpoonright L(\mathcal{K}) \models \tau(\bar{\alpha}, \bar{\gamma}) = \sigma(\bar{\alpha}', \bar{\gamma}').$$

Since $\bar{\gamma}'$ and $\bar{\beta}'$ look the same over $\bar{\gamma}\hat{\ }\bar{\alpha}'\hat{\ }\bar{\alpha}$, we also have

$$EM(I_{\mu^+}, \Phi) \upharpoonright L(\mathcal{K}) \models \tau(\bar{\alpha}, \bar{\gamma}) = \sigma(\bar{\alpha}', \bar{\beta}').$$

Combining the implications gives us a representation of b with parameters from I_δ. Thus $b \in EM(I_\delta, \Phi) \upharpoonright L(\mathcal{K})$. Since $M_0 = M_1 \cap (EM(I_\delta, \Phi) \upharpoonright L(\mathcal{K}))$, we get that $b \in M_0$ which contradicts our choice of b. □

Let us state an easy corollary of Fact I.3.15 that will simplify future constructions:

Corollary I.3.16. *Suppose μ, M_0, M_1, M_2 and $\bar{b}$ are as in the statement of Fact I.3.15. If $M_1 \prec_{\mathcal{K}} \mathfrak{C}$, then there exists a $\prec_{\mathcal{K}}$-mapping h such that*

(1) $h : M_2 \to \mathfrak{C}$,
(2) $h \upharpoonright M_0 = id_{M_0}$ *and*
(3) $h(M_2) \cap \bar{b} = M_0$ *(equivalently $h(M_2) \cap M_1 = \emptyset$).*

Proof. By Fact I.3.15, there exists a $\prec_{\mathcal{K}}$-mapping g and a model M_3 of cardinality μ such that

- $g : M_2 \to M_3$
- $g \upharpoonright M_0 = id_{M_0}$
- $g(M_2) \cap \bar{b} = M_0$ and
- $M_1 \prec_{\mathcal{K}} M_3$.

Since $\mathfrak{C}$ is universal over M_1, we can fix a $\prec_{\mathcal{K}}$-mapping f such that $f : M_3 \to \mathfrak{C}$ and $f \upharpoonright M_1 = id_{M_1}$. Notice that $h := g \circ f$ is the desired mapping from M_2 into $\mathfrak{C}$. □

4. μ-Splitting

Appearing in [28] is μ-splitting, which is a generalization of the first-order notion of splitting to AECs. Most results concerning μ-splitting are proved under the assumption of categoricity. Just recently Grossberg and VanDieren have made progress without categoricity by considering μ-splitting in Galois-stable, tame AECs (see [8]).

In this section we will develop non-μ-splitting as our dependence relation and prove the extension and uniqueness properties for non-μ-splitting types.

Before defining μ-splitting we need to describe what is meant by the image of a Galois type:

Definition I.4.1. Let M be an amalgamation base and $p \in$ ga-S(M). If h is a $\prec_{\mathcal{K}}$-mapping with domain M we can define $h(p)$ as follows. Since $\mathfrak{C}$ is saturated over M (Corollary I.2.19), we can fix $a \in \mathfrak{C}$ realizing p. By Corollary I.2.20, we can extend h to $\check{h}$ an automorphism of $\mathfrak{C}$. Denote by

$$h(p) := \text{ga-tp}(\check{h}(a)/h(M)).$$

The verification that this definition does not depend on our choices of $\check{h}$ and a is left to the reader.

Definition I.4.2. Let μ be a cardinal with $\mu < \lambda$. For $M \in \mathcal{K}^{am}$ and $p \in$ ga-S(M), we say that *p μ-splits over N* iff $N \prec_{\mathcal{K}} M$ and there exist amalgamation bases $N_1, N_2 \in \mathcal{K}_\mu$ and a $\prec_{\mathcal{K}}$-mapping $h : N_1 \cong N_2$ such that

(1) $N \prec_{\mathcal{K}} N_1, N_2 \prec_{\mathcal{K}} M$,
(2) $h(p \restriction N_1) \neq p \restriction N_2$ and
(3) $h \restriction N = id_N$.

Remark I.4.3. If T is a first-order theory stable in μ and M is saturated, then for all $N \prec M$ of cardinality μ, the first-order type, tp(a/M), does not split (in the first-order sense) over N iff ga-tp(a/M) does not μ-split over N.

Let us state some easy facts concerning μ-splitting.

Remark I.4.4. Let $N \prec_{\mathcal{K}} M \prec_{\mathcal{K}} M'$ be amalgamation bases of cardinality μ such that ga-tp(a/M') does not μ-split over N.

(1) (Monotonicity) Then ga-tp(a/M) does not μ-split over N.
(2) (Invariance) If h is a $\prec_{\mathcal{K}}$-mapping with domain M', $h($ga-tp$(a/M'))$ does not μ-split over $h(N)$.

The following appears in [28] under the assumption of the amalgamation property. The same conclusion holds in this context.

Fact I.4.5 (*Claim 3.3.1 of [28]*). *If $\mathcal{K}$ is μ-Galois stable and $\mathcal{K}$ satisfies the amalgamation property, then for every $M \in \mathcal{K}_{\geq\mu}$ and every $p \in$ ga-S(M), there exists a $N \prec_{\mathcal{K}} M$ of cardinality μ such that $N \in \mathcal{K}$ and p does not μ-split over N.*

Shelah and Villaveces draw connections between categoricity and superstability properties using μ-splitting. Let us recall some first-order consequences of superstability.

Fact I.4.6. *Let T be a countable first-order theory. Suppose T is superstable.*

(1) *If $\langle M_i \mid i \leq \sigma\rangle$ is a $\prec$-increasing and continuous chain of models and σ is a limit ordinal, then for every $p \in S(M_\sigma)$, there exists $i < \sigma$ such that p does not fork over M_i.*
(2) *Let T be a countable first-order theory. Suppose T is superstable. Let $\langle M_i \mid i \leq \sigma\rangle$ be a $\prec$-increasing and continuous chain of models with σ a limit ordinal. If $p \in S(M_\sigma)$ is such that for every $i < \sigma$, $p \restriction M_i$ does not fork over M_0, then p does not fork over M_0.*

These results are consequences of $\kappa(T) = \aleph_0$[3] and the finite character of forking (see Chapter III Section 3 of [26]). It is interesting that Shelah and Villaveces manage to prove analogs of these theorems without having the finite character of μ-splitting or the compactness theorem.

Fact I.4.7 is an analog of Fact I.4.6(1), restated: under the assumption of categoricity there are no long splitting chains. The proof of this fact relies on a combinatorial blackbox principle (see Chapter III of [27].)

Fact I.4.7 (*Theorem 2.2.1 from [33]*). *Under Assumption 0.7, suppose that*

(1) *$\langle M_i \mid i \leq \sigma\rangle$ is $\prec_{\mathcal{K}}$-increasing and continuous,*
(2) *for all $i \leq \sigma$, $M_i \in \mathcal{K}^{am}_\mu$,*
(3) *for all $i < \sigma$, M_{i+1} is universal over M_i and*
(4) *$p \in$ ga-S(M_σ).*

Then there exists an $i < \sigma$ such that p does not μ-split over M_i.

Implicit in Shelah and Villaveces' proof of Fact I.4.7 is a statement similar to Fact I.4.6(2). The proof of Fact I.4.7 is by contradiction. If Fact I.4.7 fails to be true, then there is a counter-example that has one of three properties (cases (a), (b), and (c) of their proof). Each case is separately refuted. Case (a) yields:

Fact I.4.8. *Under Assumption 0.7, suppose that*

(1) *$\langle M_i \mid i \leq \sigma\rangle$ is $\prec_{\mathcal{K}}$-increasing and continuous,*

[3] $\kappa(T)$ is the locality cardinal of non-forking; see Definition 3.1 in Chapter III Section 3 of [26].

(2) *for all* $i \leq \sigma$, $M_i \in \mathcal{K}^{am}_\mu$,
(3) *for all* $i < \sigma$, M_{i+1} *is universal over* M_i,
(4) $p \in$ ga-S(M_σ) *and*
(5) $p \restriction M_i$ *does not* μ*-split over* M_0 *for all* $i < \sigma$.

Then p *does not* μ*-split over* M_0.

The proofs of Fact I.4.7 and Fact I.4.8 use the full power of the categoricity assumption. In particular, Shelah and Villaveces use the fact that every model can be embedded into a reduct of an Ehrenfeucht–Mostowski model. It is open as to whether or not the categoricity assumption can be removed:

Question I.4.9. Can statements similar to Facts I.4.7 and I.4.8 be proved under the assumption of any of the working definitions of Galois superstability?

We now derive the extension and uniqueness properties for non-splitting types (Theorems I.4.10 and I.4.12). These results do not rely on any assumptions on the class. We will use these properties to find extensions of towers, but they are also useful for developing a stability theory for tame abstract elementary classes in [8].

Theorem I.4.10 (*Extension of Non-splitting Types*). *Suppose that* $M \in \mathcal{K}_\mu$ *is universal over* N *and* ga-tp$(a/M, \mathfrak{C})$ *does not* μ*-split over* N, *when* $\mathfrak{C}$ *is a* (μ, μ^+)*-limit containing a* $\bigcup M$.

Let $M' \in \mathcal{K}^{am}_\mu$ *be an extension of* M *with* $M' \prec_{\mathcal{K}} \mathfrak{C}$. *Then there exists a* $\prec_{\mathcal{K}}$*-mapping* $g \in \mathrm{Aut}_M(\mathfrak{C})$ *such that* ga-tp$(a/g(M'))$ *does not* μ*-split over* N. *Equivalently,* $g^{-1} \in \mathrm{Aut}_M(\mathfrak{C})$ *is such that* ga-tp$(g^{-1}(a)/M')$ *does not* μ*-split over* N.

Proof. Since M is universal over N, there exists a $\prec_{\mathcal{K}}$-mapping $h' : M' \to M$ with $h' \restriction N = id_N$. By Corollary I.2.20, we can extend h' to an automorphism h of $\mathfrak{C}$. Notice that by monotonicity, ga-tp$(a/h(M'))$ does not μ-split over N. By invariance,

$$\text{ga-tp}(h^{-1}(a)/M') \text{ does not } \mu\text{-split over } N. \qquad (*)$$

Subclaim I.4.11. ga-tp$(h^{-1}(a)/M) =$ ga-tp(a/M).

Proof. We will use the notion of μ-splitting to prove this subclaim. So let us rename the models in such a way that our application of the definition of μ-splitting will become transparent. Let $N_1 := h^{-1}(M)$ and $N_2 := M$. Let $p :=$ ga-tp$(h^{-1}(a)/h^{-1}(M))$. Consider the mapping $h : N_1 \cong N_2$. By invariance, p does not μ-split over N. Thus, $h(p \restriction N_1) = p \restriction N_2$. Let us calculate this

$$h(p \restriction N_1) = \text{ga-tp}(h(h^{-1}(a))/h(h^{-1}(M))) = \text{ga-tp}(a/M).$$

While,

$$p \restriction N_2 = \text{ga-tp}(h^{-1}(a)/M).$$

Thus ga-tp$(h^{-1}(a)/M) =$ ga-tp(a/M) is as required. □

From the subclaim, we can find a $\prec_{\mathcal{K}}$-mapping $g \in \mathrm{Aut}_M(\mathfrak{C})$ such that $g \circ h^{-1}(a) = a$. Notice that by applying g to $(*)$ we get

$$\text{ga-tp}(a/g(M'), \mathfrak{C}) \text{ does not } \mu\text{-split over } N. \qquad (**)$$

Applying g^{-1} to $(**)$ gives us the *equivalently* clause:

$$\text{ga-tp}(g^{-1}(a)/M', \mathfrak{C}) \text{ does not } \mu\text{-split over } N.$$

Since $g \restriction M = id_M$, we have that

$$\text{ga-tp}(g(a)/M) = \text{ga-tp}(g^{-1}(a)/M) = \text{ga-tp}(a/M). \quad \square$$

Not only do non-splitting extensions exist, but they are unique:

Theorem I.4.12 (*Uniqueness of Non-splitting Extensions*). *Let $N, M, M' \in \mathcal{K}^{am}_\mu$ be such that M' is universal over M and M is universal over N. If $p \in$ ga-S(M) does not μ-split over N, then there is a unique $p' \in$ ga-S(M') such that p' extends p and p' does not μ-split over N.*

Proof. By Theorem I.4.10, there exists $p' \in$ ga-S(M') extending p such that p' does not μ-split over N. Suppose for the sake of contradiction that there exists $q' \neq p' \in$ ga-S(M') extending p such that q' does not μ-split over N. Let a, b be such that $p' =$ ga-tp(a/M') and $q' =$ ga-tp(b/M'). Since M is universal over N, there exists a $\prec_{\mathcal{K}}$-mapping $f : M' \to M$ with $f \restriction N = id_N$. Since p' and q' do not μ-split over N we have

$$\text{ga-tp}(a/f(M')) = \text{ga-tp}(f(a)/f(M')) \text{ and} \tag{$*$}_a$$

$$\text{ga-tp}(b/f(M')) = \text{ga-tp}(f(b)/f(M')). \tag{$*$}_b$$

On the other hand, since $p' \neq q'$, we have that

$$\text{ga-tp}(f(a)/f(M')) \neq \text{ga-tp}(f(b)/f(M')). \tag{$*$}$$

Combining $(*)_a$, $(*)_b$ and $(*)$, we get

$$\text{ga-tp}(a/f(M')) \neq \text{ga-tp}(b/f(M')).$$

Since $f(M') \prec_{\mathcal{K}} M$, this inequality witnesses that

$$\text{ga-tp}(a/M) \neq \text{ga-tp}(b/M),$$

contradicting our choice of p' and q' both extending p. □

Remark I.4.13. Notice that the following follows from the existence and uniqueness of non-splitting extensions: Let $N, M, M' \in \mathcal{K}^{am}_\mu$ with M universal over N and $M \prec_{\mathcal{K}} M'$. If $p \in$ ga-S(M) does not μ-split over N and is non-algebraic, then any $q \in$ ga-S(M') which extends p and does not μ-split over N is also non-algebraic.

The following is a corollary of the existence and uniqueness for non-splitting types. It allows us to extend an increasing chain of non-splitting types. Recall that in an AEC, a type p extending an increasing chain of types $\langle p_i \mid i < \theta\rangle$ does not always exist and may not even be unique when it does exist (see [2]).

Corollary I.4.14. *Suppose that $\langle M_i \in \mathcal{K}^{am}_\mu \mid i < \theta\rangle$ is a $\prec_{\mathcal{K}}$-increasing chain of models and $\langle p_i \in$ ga-S$(M_i) \mid i < \theta\rangle$ is an increasing chain of types such that for every $i < \theta$, p_i does not μ-split over M_0 and M_1 is universal over M_0. If $M = \bigcup_{i<\theta} M_i$ is an amalgamation base, then there exists $p \in$ ga-S(M) such that for each $i < \theta$ $p_i \subset p$. Moreover, p does not μ-split over M_0.*

Proof. Suppose that M is an amalgamation base. By Theorem I.4.10, there is $p \in$ ga-S(M) extending p_1 such that p does not μ-split over M_0. By Theorem I.4.12, we have that for every $i < \theta$, $p_i = p \restriction M_i$. □

5. Towers

While Theorem I.4.10 allows us to find extensions of a non-splitting Galois type in any AEC, Sections 7 and 10 are dedicated to the difficult task of finding non-splitting extensions of α-many types simultaneously under categoricity. The mechanics used to do this include towers.

Shelah introduced chains of towers in [20] and [23] as a tool to build a model of cardinality μ^{++} from models of cardinality μ. Towers are also used in [3] to handle abstract classes which satisfy Axioms 1–4 of AECs, but for which the union axiom, Axiom 5, is not assumed. A particular example of such classes is the class of Banach spaces.

We follow the notation introduced in [33]. In [33] several other towers were defined. The superscript c in the ordering $<^c_{\mu,\alpha}$ and the superscripts $+$ and $*$ in the class ${}^+\mathcal{K}^*_{\mu,\alpha}$ serve as parameters in their paper to distinguish various definitions. In this paper, we will refer to only the towers in ${}^+\mathcal{K}^*_{\mu,\alpha}$ ordered by $<^c_{\mu,\alpha}$.

Definition I.5.1.

$$
{}^{+}\mathcal{K}^{*}_{\mu,\alpha} := \left\{ (\bar{M}, \bar{a}, \bar{N}) \;\middle|\; \begin{array}{l} \bar{M} = \langle M_i \in \mathcal{K}_\mu \mid i < \alpha \rangle \text{ is } \prec_{\mathcal{K}} \text{-increasing;} \\ M_i \text{ is a } (\mu, \theta_i)\text{-limit model for some } \theta_i < \mu^+; \\ a_i \in M_{i+1} \backslash M_i \text{ for } i+1 < \alpha; \\ \bar{N} = \langle N_i \in \mathcal{K}_\mu \mid i+1 < \alpha \rangle \\ N_i \text{ is a } (\mu, \sigma_i)\text{-limit model for some } \sigma_i < \mu^+; \\ \text{for every } i+1 < \alpha,\ N_i \prec_{\mathcal{K}} M_i; \\ M_i \text{ is universal over } N_i \text{ and} \\ \text{ga-tp}(a_i/M_i, M_{i+1}) \text{ does not } \mu\text{-split over } N_i. \end{array} \right\}.
$$

Remark I.5.2. The sequence $\bar{M}$ is not necessarily continuous. The sequence $\bar{N}$ may not be $\prec_{\mathcal{K}}$-increasing or continuous.

Notation I.5.3. We will use the term *continuous tower* to refer to towers of the form $(\bar{M}, \bar{a}, \bar{N})$ with $\bar{M}$ a continuous sequence. If $(\bar{M}, \bar{a}, \bar{N}) \in {}^{+}\mathcal{K}^{*}_{\mu,\alpha}$, we say that $\bigcup_{i<\alpha} M_i$ is the *top of the tower* and that $(\bar{M}, \bar{a}, \bar{N})$ has *length* α.

Notation I.5.4. For θ a limit ordinal $< \mu^+$, we write ${}^{+}\mathcal{K}^{\theta}_{\mu,\alpha}$ for the collection of all towers $(\bar{M}, \bar{a}, \bar{N}) \in {}^{+}\mathcal{K}^{*}_{\mu,\alpha}$ where each M_i is a (μ, θ)-limit model.

Our goal is to simultaneously extend the α non-splitting Galois types, $\{\text{ga-tp}(a_i/M_i, M_{i+1}) \mid i+1 < \alpha\}$. The following ordering on towers captures this.

Definition I.5.5. For $(\bar{M}, \bar{a}, \bar{N})$ and $(\bar{M}', \bar{a}', \bar{N}') \in {}^{+}\mathcal{K}^{*}_{\mu,\alpha}$, we say
$(\bar{M}, \bar{a}, \bar{N}) \leq^{c}_{\mu,\alpha} (\bar{M}', \bar{a}', \bar{N}')$ iff

(1) for $i < \alpha$ either $M'_i = M_i$ or M'_i is universal over M_i,
(2) $\bar{a} = \bar{a}'$ and
(3) $\bar{N} = \bar{N}'$.

We say $(\bar{M}, \bar{a}, \bar{N}) <^{c}_{\mu,\alpha} (\bar{M}', \bar{a}', \bar{N}')$ iff $(\bar{M}, \bar{a}, \bar{N}) \leq^{c}_{\mu,\alpha} (\bar{M}', \bar{a}', \bar{N}')$ and $M'_i \neq M_i$ for every $i < \alpha$.

Remark I.5.6. Notice that in Definition I.5.5, if $(\bar{M}, \bar{a}, \bar{N}) <^{c}_{\mu,\alpha} (\bar{M}', \bar{a}, \bar{N})$, then for every $i < \alpha$, ga-tp$(a_i/M'_i, M'_{i+1})$ does not μ-split over N_i.

Notation I.5.7. We will often be looking at extensions of an initial segment of a tower. We introduce the following notation for this. Suppose $(\bar{M}, \bar{a}, \bar{N}) \in {}^{+}\mathcal{K}^{*}_{\mu,\alpha}$. Let $\beta < \alpha$. We write $\bar{M} \restriction \beta$ for the sequence $\langle M_i \mid i < \beta \rangle$. Similarly, $\bar{a} \restriction \beta = \langle a_i \mid i+1 < \beta \rangle$ and $\bar{N} \restriction \beta = \langle N_i \mid i+1 < \beta \rangle$. Then $(\bar{M}, \bar{a}, \bar{N}) \restriction \beta$ will represent the tower $(\bar{M} \restriction \beta, \bar{a} \restriction \beta, \bar{N} \restriction \beta) \in {}^{+}\mathcal{K}^{*}_{\mu,\beta}$. If $(\bar{M}', \bar{a}', \bar{N}')$ is a $<^{c}_{\mu,\beta}$-extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \beta$, we refer to it as a *partial extension* of $(\bar{M}, \bar{a}, \bar{N})$.

The requirement that M'_i is universal over M_i in the definition of $<^{c}_{\mu,\alpha}$ allows us to conclude that the models in the union of a $<^{c}_{\mu,\alpha}$-increasing chain of towers are limit models. In particular, the union of a $<^{c}_{\mu,\alpha}$-increasing chain (of length $< \mu^+$) of towers is a tower.

Definition I.5.8. We say that $\mathcal{K}$ satisfies the $<^{c}_{\mu,\alpha}$-extension property iff every tower in ${}^{+}\mathcal{K}^{*}_{\mu,\alpha}$ has a $<^{c}_{\mu,\alpha}$-extension.

The $<^{c}_{\mu,\alpha}$-extension property serves as a weak substitute for the extension property of non-forking in first-order model theory, but is much stronger than the extension property for non-splitting. Notice that for towers with $\alpha = 1$, Theorem I.4.10 and the existence of universal extensions (Fact I.2.5) give the $<^{c}_{\mu,1}$ extension property. Actually it is possible to derive the $<^{c}_{\mu,n}$-extension property for all $n \leq \omega$ with no more than the existence of universal extensions and the extension property for non-splitting types.

It is open whether or not every $\mathcal{K}$ satisfying Assumption 0.7 has the $<^{c}_{\mu,\alpha}$-extension property for $\alpha > \omega$. The difficulties concern discontinuous towers. Notice that if $(\bar{M}, \bar{a}, \bar{N})$ is not continuous, then for some limit ordinal $i < \alpha$, we may have that $\bigcup_{j<i} M_j$ is not an amalgamation base. Suppose that we have constructed a partial extension of $(\bar{M}, \bar{a}, \bar{N})$ up to i. It may be the case that this extension and M_i may not be amalgamated over $\bigcup_{j<i} M_j$. This

would rule out much hope for using the partial extension as a base for a longer extension of the entire tower $(\bar{M}, \bar{a}, \bar{N})$. With this in mind, it is natural to restrict ourselves to continuous towers. However, it is not easy to prove that every continuous tower has a continuous extension. In fact, we can only prove this under an extra assumption, Hypothesis 1 (see Section 7).

In addition to the continuous towers, we have identified two subclasses of ${}^{+}\mathcal{K}^*_{\mu,\alpha}$, amalgamable and nice towers, for which a $<^c_{\mu,\alpha}$-extension property can be proven.

Definition I.5.9. We say that $(\bar{M}, \bar{a}, \bar{N}) \in {}^{+}\mathcal{K}^*_{\mu,\alpha}$ is *nice* iff whenever $i < \alpha$ is a limit ordinal, $\bigcup_{j<i} M_j$ is an amalgamation base.

Remark I.5.10. Since every M_i is a (μ, θ_i)-limit for some limit ordinal $\theta_i < \mu^+$, by Theorem I.3.13, we have that every M_i is also an amalgamation base. So *nice* only is a requirement for limit ordinals i where $\bar{M}$ is not continuous at i. Thus, if $(\bar{M}, \bar{a}, \bar{N})$ is a continuous tower, then $(\bar{M}, \bar{a}, \bar{N})$ is nice.

Notice that the definition of nice does not require that the top of the tower $\left(\bigcup_{i<\alpha} M_i\right)$ be an amalgamation base. For these towers we introduce the terminology:

Definition I.5.11. We say that $(\bar{M}, \bar{a}, \bar{N}) \in {}^{+}\mathcal{K}^*_{\mu,\alpha}$ is *amalgamable* iff it is nice and $\bigcup_{i<\alpha} M_i$ is an amalgamation base.

We use the word amalgamable to refer to such towers, because any two $<^c_{\mu,\alpha}$-extensions of an amalgamable tower $(\bar{M}, \bar{a}, \bar{N})$ can be amalgamated over $\bigcup_{i<\alpha} M_i$.

Notice that the classes of amalgamable and nice towers both avoid the problematic towers described above. Namely, if $(\bar{M}, \bar{a}, \bar{N})$ is discontinuous at i, we require that $\bigcup_{j<i} M_j$ is an amalgamation base. We can show that every nice tower has an amalgamable extension (Theorem III.10.1). In particular, every continuous tower has an amalgamable extension. However, this amalgamable extension may not be continuous. Furthermore, if we instead restrict ourselves to amalgamable towers, we will run into the difficulty that the union of a $<^c_{\mu,\alpha}$-increasing chain of amalgamable towers need not be amalgamable (or even nice). But, with a little help from Hypothesis 1, we are able to carry through the strategy of restricting ourselves to continuous towers. By carefully stacking the amalgamable extensions in Section 7, we construct continuous extensions of continuous towers.

Notation I.5.12. Ultimately, we will be constructing a $<^c_{\mu,\alpha}$-extension, $(\bar{M}', \bar{a}', \bar{N}')$ of a tower $(\bar{M}, \bar{a}, \bar{N})$, but we will allow the extension to live on a larger index set, $(\bar{M}', \bar{a}', \bar{N}') \in {}^{+}\mathcal{K}^*_{\mu,\alpha'}$ for some $\alpha' > \alpha$. We will also like to arrange the construction so that α is not identified with an initial segment of α', but as some other scattered, increasing subsequence of α'. Therefore, we will prefer to consider the relative structure of these index sets in addition to their order types. We make the following convention that will be justified in later constructions. When α and δ are ordinals, $\alpha \times \delta$ with the lexicographical ordering ($<_{lex}$), is well ordered. Recall that $\mathrm{otp}(\alpha \times \delta, <_{lex}) = \delta \cdot \alpha$ where $\cdot$ is ordinal multiplication. For easier notation in future constructions, we will identify $\alpha \times \delta$ with the interval of ordinals $[0, \delta \cdot \alpha)$ and ${}^{+}\mathcal{K}^*_{\mu,\alpha\times\delta}$ will refer to the collection of towers ${}^{+}\mathcal{K}^*_{\mu,\delta\cdot\alpha}$. The notation will be more convenient when we compare towers in ${}^{+}\mathcal{K}^*_{\mu,\alpha\times\delta}$ with those in ${}^{+}\mathcal{K}^*_{\mu,\alpha'\times\delta'}$ for $\alpha' \geq \alpha$ and $\delta' \geq \delta$.

We will make use of the following proposition concerning $<^c_{\mu,\alpha}$ throughout the paper.

Proposition I.5.13. *If $(\bar{M}', \bar{a}, \bar{N})$ is a $<^c_{\mu,\alpha}$-extension of $(\bar{M}, \bar{a}, \bar{N})$, then for every $i \leq j < \alpha$, we have that M'_j is universal over M_i.*

Proof. By definition of $<^c_{\mu,\alpha}$, we have that M'_i is universal over M_i. Since $\bar{M}'$ is increasing, $M'_i \preceq_{\mathcal{K}} M'_j$. So M'_j is universal over M_i as well. □

Part II. Uniqueness of limit models

We will use towers to prove the uniqueness of limit models by producing a model which is simultaneously a (μ, θ_1)-limit model and a (μ, θ_2)-limit model. The construction of such a model is sufficient to prove the uniqueness of limit models by Fact I.2.11 and involves building an increasing and continuous chain of towers.

The idea is to build a two-dimensional array (with the cofinality of the height $= \theta_1$ and the cofinality of the width $= \theta_2$) of models such that the bottom corner of the array (M^*) is a (μ, θ_1)-limit model witnessed by the last column and a (μ, θ_2)-limit model witnessed by the last row of the array. The actual construction involves increasing the length of the towers as we go from one row to the next.

The construction of this array is done by identifying each row of the array with a tower and then building a $<^c_{\mu,\alpha}$-increasing and continuous chain of towers (where α will vary throughout our construction).

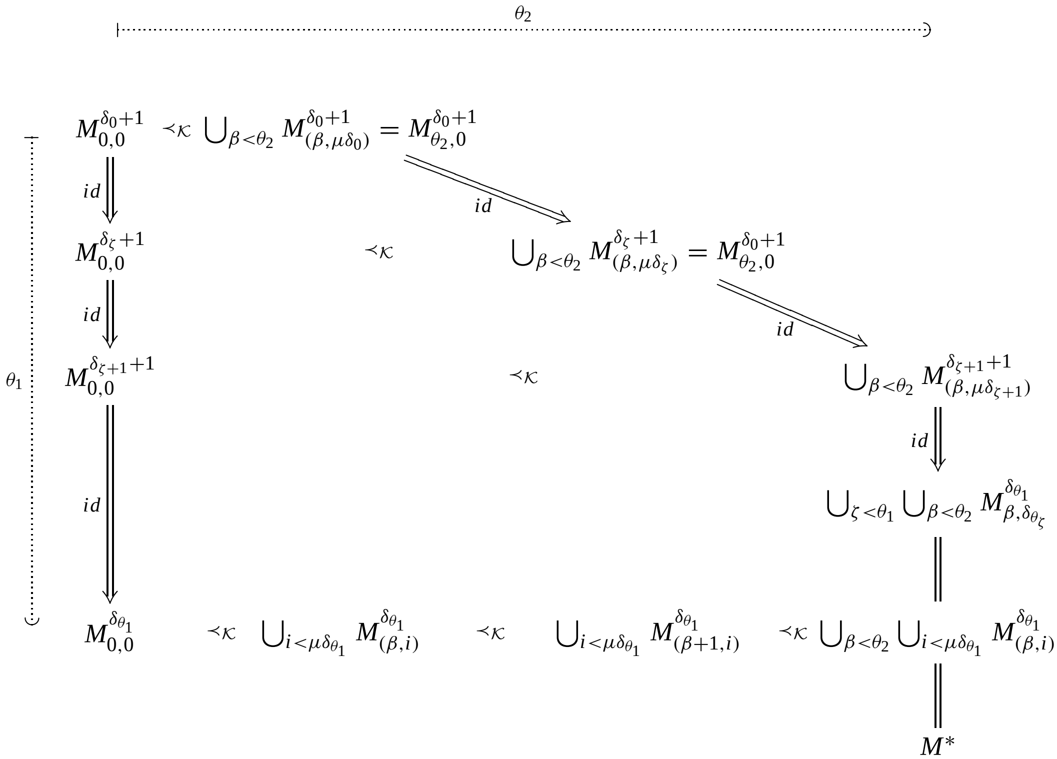

In order to witness that M^* is a (μ, θ_1)-limit model, we will need for our towers to be *increasing* in such a way that the models in the $\delta + 1$st tower are universal over the models in the δth tower. This is possible if we can prove that every continuous tower has a continuous $<^c_{\mu,\alpha}$-extension. This is the subject of Section 7 and related material appears in Section 10.

While M^* is built up by a chain of cofinality θ_2, it may not be a (μ, θ_2)-limit model. In order to conclude that M^* is a (μ, θ_2)-limit model, we show in Section 6, that the top of a continuous, relatively full tower of length θ_2 is a (μ, θ_2)-limit model. We will construct the relatively full tower by requiring that at every stage of our construction of the array, we realize all the strong types over the previous tower in a systematic way. Section 8 provides the technical machinery to carry this through. The actual construction of M^* is carried out in Section 9.

6. Relatively full towers

We begin this section by recalling the definition of *strong types* from [33].

Definition II.6.1 (*Definition 3.2.1 of [33]*). For M a (μ, θ)-limit model,

(1) Let

$$\mathfrak{St}(M) := \left\{ (p, N) \;\middle|\; \begin{array}{l} N \prec_{\mathcal{K}} M; \\ N \text{ is a } (\mu,\theta) - \text{limit model}; \\ M \text{ is universal over } N; \\ p \in \text{ga-S}(M) \text{ is non-algebraic}; \\ \text{and } p \text{ does not } \mu - \text{split over } N. \end{array} \right\}.$$

(2) For types $(p_l, N_l) \in \mathfrak{St}(M)$ $(l = 1, 2)$, we say $(p_1, N_1) \sim (p_2, N_2)$ iff for every $M' \in \mathcal{K}^{am}_\mu$ extending M there is a $q \in$ ga-S(M') extending both p_1 and p_2 such that q does not μ-split over N_1 and q does not μ-split over N_2.

Notation II.6.2. Suppose $M \prec_{\mathcal{K}} M'$ are amalgamation bases of cardinality μ. For $(p, N) \in \mathfrak{St}(M')$, if M is universal over N, we define the restriction $(p, N) \restriction M \in \mathfrak{St}(M')$ to be $(p \restriction M, N)$.

We write $(p, N) \restriction M$ only when p does not μ-split over N and M is universal over N.

Notice that $\sim$ is an equivalence relation on $\mathfrak{St}(M)$. To see that $\sim$ is a transitive relation on $\mathfrak{St}(M)$, suppose that $(p_1, N_1) \sim (p_2, N_2)$ and $(p_2, N_2) \sim (p_3, N_3)$. Let $M' \in \mathcal{K}^{am}_\mu$ be an extension of M and fix $q_{ij} \in$ ga-S(M') extending both p_i and p_j and q_{ij} does not μ-split over both N_i and N_j (for $\langle i, j\rangle = \langle 1, 2\rangle, \langle 2, 3\rangle$). Since p_2 has a unique non-splitting extension to M' (Theorem I.4.12), we know that $q_{12} = q_{23}$. Then q_{12} witnesses that $(p_1, N_1) \sim (p_3, N_3)$ since it is an extension of both p_1 and p_3 and does not μ-split over both N_1 and N_3.

The following lemma is used to provide a bound on the number of strong types.

Lemma II.6.3. *Given $M \in \mathcal{K}^{am}_\mu$, and $(p, N), (p', N') \in \mathfrak{St}(M)$. Let $M' \in \mathcal{K}^{am}_\mu$ be a universal extension of M. To show that $(p, N) \sim (p', N')$ it suffices to find $q \in$ ga-S(M') such that q extends both p and p' and such that q does not μ-split over N and N'.*

Proof. Suppose $q \in$ ga-S(M') extends both p and p' and does not μ-split over N and N'. Let $M^* \in \mathcal{K}^{am}_\mu$ be an extension of M. By universality of M', there exists $f : M^* \to M'$ such that $f \restriction M = id_M$. Consider $f^{-1}(q)$. It extends p and p' and does not μ-split over N and N' by invariance. Thus $(p, N) \sim (p', N')$. □

The following appears as a Fact 3.2.2(3) in [33]. We provide a proof here for completeness.

Fact II.6.4. *For $M \in \mathcal{K}^{am}_\mu$, $|\mathfrak{St}(M)/\sim| \le \mu$.*

Proof of Fact II.6.4. Suppose for the sake of contradiction that $|\mathfrak{St}(M)/\sim| > \mu$.

Let $\{(p_i, N_i) \in \mathfrak{St}(M) \mid i < \mu^+\}$ be pairwise non-equivalent. By Galois stability (Fact I.1.8) and the pigeon-hole principle, there exist $p \in$ ga-S(M) and $I \subset \mu^+$ of cardinality μ^+ such that $i \in I$ implies $p_i = p$. Set $p :=$ ga-tp(a/M) with $a \in \mathfrak{C}$.

Fix $M' \in \mathcal{K}^{am}_\mu$ a universal extension of M inside $\mathfrak{C}$. We will show that there are $\ge \mu^+$ types over M'. This will provide us with a contradiction since $\mathcal{K}$ is Galois stable in μ (Fact I.1.8).

For each $i \in I$, by the extension property of non-splitting (Theorem I.4.10), there exists $f_i \in \mathrm{Aut}_M(\mathfrak{C})$ such that

- ga-tp$(f_i(a)/M')$ does not μ-split over N_i and
- ga-tp$(f_i(a)/M')$ extends ga-tp(a/M).

Claim II.6.5. *For $i \ne j \in I$, we have that the types,* ga-tp$(f_i(a)/M')$ *and* ga-tp$(f_j(a)/M')$, *are not equal.*

Proof of Claim II.6.5. Otherwise ga-tp$(f_i(a)/M')$ does not μ-split over N_i and does not μ-split over N_j. By Lemma II.6.3, this implies that $(p, N_i) \sim (p, N_j)$ contradicting our choice of non-$\sim$-equivalent strong types. □

This completes the proof as $\{$ga-tp$(f_i(a)/M') \mid i \in I\}$ is a set of μ^+ distinct types over M', contradicting μ-Galois stability. □

We can now consider towers which are saturated with respect to strong types (from $\mathfrak{St}(M)$). These towers are called relatively full.

Definition II.6.6. Let α, δ and θ be limit ordinals $< \mu^+$. Suppose $\langle \bar{M}_{\beta,i} \mid (\beta, i) \in \alpha \times \delta\rangle$ is such that each $\bar{M}_{\beta,i}$ is a sequence of limit models, $\langle M^\gamma_{\beta,i} \mid \gamma < \theta\rangle$, with $M^{\gamma+1}_{\beta,i}$ universal over $M^\gamma_{\beta,i}$ for all $(\beta, i) \in \alpha \times \delta$.

A tower $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^\theta_{\mu,\alpha\times\delta}$ is said to be *full relative to* $\langle \bar{M}^\gamma \mid \gamma < \theta\rangle$ iff for all $(\beta, i) \in \alpha \times \delta$

(1) $\bar{M}_{\beta,i}$ witnesses that $M_{\beta,i}$ is a (μ, θ)-limit model and
(2) for all $(p, N^*) \in \mathfrak{St}(M_{\beta,i})$ with $N^* = M^\gamma_{\beta,i}$ for some $\gamma < \theta$, there is a $j < \delta$ such that $($ga-tp$(a_{\beta+1,j}/M_{\beta+1,j}), N_{\beta+1,j}) \restriction M_{\beta,i} \sim (p, N^*)$.

$$
\begin{array}{cccc}
M^0_{0,0} & \prec_{\mathcal{K}} \cdots \prec_{\mathcal{K}} & M^0_{\beta,i} & \prec_{\mathcal{K}} \cdots \\
\Big\Downarrow id & & \Big\Downarrow id & \\
M^{\gamma}_{0,0} & \prec_{\mathcal{K}} \cdots \prec_{\mathcal{K}} & M^{\gamma}_{\beta,i} & \prec_{\mathcal{K}} \cdots \\
\Big\Downarrow id & & \Big\Downarrow id & \\
M^{\gamma+1}_{0,0} & \prec_{\mathcal{K}} \cdots \prec_{\mathcal{K}} & M^{\gamma+1}_{\beta,i} & \prec_{\mathcal{K}} \cdots \\
\Big\Downarrow id & & \Big\Downarrow id & \\
M_{0,0} = \bigcup_{\gamma<\theta} M^{\gamma}_{0,0} & \prec_{\mathcal{K}} \cdots \prec_{\mathcal{K}} & M_{\beta,i} = \bigcup_{\gamma<\theta} M^{\gamma}_{\beta,i} & \prec_{\mathcal{K}} \cdots
\end{array}
$$

Notation II.6.7. We say that $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^{\theta}_{\mu,\alpha\times\delta}$ is *relatively full* iff there exists $\langle \bar{M}_{\beta,i} \mid (\beta, i) \in \alpha \times \delta\rangle$ as in Definition II.6.6 such that $(\bar{M}, \bar{a}, \bar{N})$ is full relative to $\langle \bar{M}_{\beta,i} \mid (\beta, i) \in \alpha \times \delta\rangle$.

Remark II.6.8. A strengthening of Definition II.6.6 appears in [33] under the name full towers (see Definition 3.2.3 of their paper). Consider the statement:

$$\forall M \in \mathcal{K}^{am}_{\mu} \text{ and } \forall (p, N), (p', N') \in \mathfrak{St}(M),\ (p, N) \sim (p', N') \text{ iff } p = p'. \tag{$*$}$$

Notice that for $M \in \mathcal{K}^{am}_{\mu}$, if $(p, N) \sim (p', N') \in \mathfrak{St}(M)$, then necessarily $p = p'$. To see this, take $M' \in \mathcal{K}^{am}_{\mu}$ some extension of M and $q \in$ ga-S(M') such that q extends both p and p' and does not μ-split over N and N'. Then $q \restriction M = p$ and $q \restriction M = p'$. So p and p' must be equal.

However we do not know that $(*)$ holds in our context. Shelah has implicitly shown, with much work, that it does hold in categorical AECs which satisfy the amalgamation property [28]. It is a consequence of transitivity of non-splitting.

Property $(*)$ implies that relatively full towers are full. We use relatively full towers since the construction of full towers by an increasing chain of towers in this context has been seen to be problematic.

The following proposition is immediate from the definition of relative fullness.

Proposition II.6.9. *Let α and δ be limit ordinals $< \mu^+$. If $(\bar{M}, \bar{a}, \bar{N}) \in {}^{+}\mathcal{K}^{\theta}_{\mu,\alpha\times\delta}$ is full relative to $\langle \bar{M}_{\beta,i} \mid (\beta, i) \in \alpha \times \delta\rangle$, then for every limit ordinal $\beta < \alpha$, we have that the restriction $(\bar{M}, \bar{a}, \bar{N}) \restriction \beta \times \delta$ is full relative to $\langle \bar{M}_{\beta',i'} \mid (\beta', i') \in \beta \times \delta\rangle$.*

The following theorem is proved in [33] for full towers (Theorem 3.2.4 of their paper). Our strengthening provides us with an alternative characterization of limit models as the top of a relatively full tower.

Theorem II.6.10. *Let α be an ordinal $< \mu^+$ such that $\alpha = \mu \cdot \alpha$. Suppose $\delta < \mu^+$. If $(\bar{M}, \bar{a}, \bar{N}) \in {}^{+}\mathcal{K}^{\theta}_{\mu,\alpha\times\delta}$ is full relative to $\langle \bar{M}_{\beta,i} \mid (\beta, i) \in \alpha \times \delta\rangle$ and $\bar{M}$ is continuous, then $M := \bigcup_{i<\alpha\cdot\delta} M_i$ is a $(\mu, \mathrm{cf}(\alpha))$-limit model over M_0.*

Proof. Let $M' \prec_{\mathcal{K}} \mathfrak{C}$ be a (μ, α)-limit over $M_{0,0}$ witnessed by $\langle M'_i \mid i < \alpha\rangle$. By Weak Disjoint Amalgamation and renaming elements, we can arrange that $\bigcup_{i<\alpha} M'_i \cap \bigcup_{i<\alpha\cdot\delta} M_i = M_{0,0}$ and that for each $i < \alpha$ we can identify the universe of M'_i with $\mu(1 + i)$. Notice that since $\alpha = \mu \cdot \alpha$, we have that $i \in M'_{i+1}$ for every $i < \alpha$. We will construct an isomorphism from M into M'.

Now we define by induction on $i < \alpha$ a increasing and continuous sequence of $\prec_{\mathcal{K}}$-mappings $\langle h_i \mid i < \alpha\rangle$ such that

(1) $h_i : M_{i,j} \to M'_{i+1}$ for some $j < \delta$
(2) $h_0 = id_{M_{0,0}}$ and
(3) $i \in \mathrm{rg}(h_{i+1})$.

For $i = 0$ take $h_0 = id_{M_{0,0}}$. For i a limit ordinal let $\check{h}_i = \bigcup_{i'<i} h_{i'}$. Since $\bar{M}$ is continuous, we know that $\bigcup_{\substack{i'<i\\ j<\delta}} M_{i',j}$ is an amalgamation base. Thus the induction hypothesis gives us that h_i is a $\prec_{\mathcal{K}}$-mapping from $M_{i,0} = \bigcup_{\substack{i'<i\\ j<\delta}} M_{i',j}$ into M'_i allowing us to satisfy condition (1) of the construction.

Suppose that h_i has been defined. Let $j < \delta$ be such that $h_i : M_{i,j} \to M'_{i+1}$. There are two cases: either $i \in \mathrm{rg}(h_i)$ or $i \notin \mathrm{rg}(h_i)$. First suppose that $i \in \mathrm{rg}(h_i)$. Since M'_{i+2} is universal over M'_{i+1}, it is also universal over $h_i(M_{i,j})$. This allows us to extend h_i to $h_{i+1} : M_{i+1,0} \to M'_{i+2}$.

Now consider the case when $i \notin \mathrm{rg}(h_i)$. We illustrate the construction for this case:

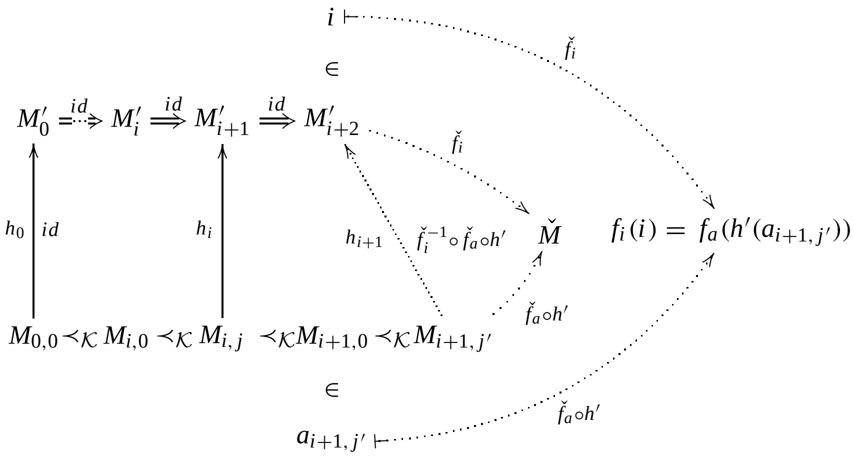

Since $\langle M^{\gamma}_{i,j} \mid \gamma < \theta \rangle$ witness that $M_{i,j}$ is a (μ, θ)-limit model, by Fact I.4.7, there exists $\gamma < \theta$ such that ga-tp$(i/M_{i,j})$ does not μ-split over $M^{\gamma}_{i,j}$. By our choice of $\bar{M}'$ disjoint from $\bar{M}$ outside of M_0, we know that $i \notin M_{i,j}$. Thus ga-tp$(i/M_{i,j})$ is non-algebraic and by relative fullness of $(\bar{M}, \bar{a}, \bar{N})$, there exists $j' < \delta$ such that

$$(\text{ga-tp}(i/M_{i,j}), M^{\gamma}_{i,j}) \sim (\text{ga-tp}(a_{i+1,j'}/M_{i+1,j'}), N_{i+1,j'}) \restriction M_{i,j}.$$

In particular we have that

$$\text{ga-tp}(a_{i+1,j'}/M_{i,j}) = \text{ga-tp}(i/M_{i,j}). \tag{$*$}$$

We can extend h_i to an automorphism h' of $\mathfrak{C}$. An application of h' to $(*)$ gives us

$$\text{ga-tp}(h'(a_{i+1,j'})/h_i(M_{i,j})) = \text{ga-tp}(i/h_i(M_{i,j})). \tag{$**$}$$

By $(**)$, there exist $M^* \in \mathcal{K}^{am}_{\mu}$ a $\mathcal{K}$-substructure of $\mathfrak{C}$ containing $M_{i,j}$ and $\prec_{\mathcal{K}}$-mappings $f_a : h'(M_{i+1,j'+1}) \to M^*$ and $f_i : M'_{i+2} \to M^*$ such that $f_a(h'(a_{i+1,j'})) = f_i(i)$ and $f_a \restriction h_i(M_{i,j}) = f_i \restriction h_i(M_{i,j}) = id_{h_i(M_{i,j})}$. Since M'_{i+2} is universal over M'_{i+1}, it is also universal over $h_i(M_{i,j})$. So we may assume that $M^* = M'_{i+2}$. Since $\mathfrak{C}$ is a (μ, μ^+)-limit model, we can extend f_a and f_i to automorphisms of $\mathfrak{C}$, say $\check{f}_a$ and $\check{f}_i$. Let $h_{i+1} : M_{i+1,j'+1} \to M'_{i+2}$ be defined as $\check{f}_i^{-1} \circ \check{f}_a \circ h'$. Notice that $h_{i+1}(a_{i+1,j'}) = i$. This completes the construction.

Let $h := \bigcup_{i<\alpha} h_i$. Clearly $h : M \to M'$. To see that h is an isomorphism, notice that condition (3) of the construction forces h to be surjective. □

Remark II.6.11. Theorem II.6.10 can be improved by replacing the assumption of continuity of $(\bar{M}, \bar{a}, \bar{N})$ with niceness. The same proof works with a minor adjustment at the limit stage. We lift the requirement that $\langle h_i \mid i < \alpha \rangle$ is continuous and use the fact that M'_{i+1} is universal over M'_i to carry out the construction at limits.

7. Existence of continuous $<^c_{\mu,\alpha}$-extensions

Our proof of the uniqueness of limit models will involve a $<^c_{\mu,\alpha}$-increasing chain of continuous towers such that the index sets of the towers grow throughout the chain. The purpose of this section and of Section 8 is to develop the machinery that will allow us to construct such a chain of continuous towers while refining the index sets along the way. While we will only use the fact that every continuous tower has a continuous extension, we prove the stronger statement to fuel the induction of Theorem II.7.1.

The claim that every continuous tower has a continuous extension still alludes a full solution. Hypothesis 1 is sufficient to derive the extension property. It is an open problem if this hypothesis can be removed.

Hypothesis 1: Every continuous tower of length α inside $\mathfrak{C}$ has an amalgamable $<^c_{\mu,\alpha}$-extension inside $\mathfrak{C}$.

Theorem II.7.1 (*Existence of Continuous Extensions*). *Let $(\bar{M}, \bar{a}, \bar{N})$ be a nice tower of length α in $\mathfrak{C}$. Under Hypothesis 1, there exists a continuous, amalgamable tower $(\bar{M}^*, \bar{a}, \bar{N})$ inside $\mathfrak{C}$ such that $(\bar{M}, \bar{a}, \bar{N}) <^c_{\mu,\alpha} (\bar{M}^*, \bar{a}, \bar{N})$.*

Furthermore, if $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\beta}$ is a continuous partial extension of $(\bar{M}, \bar{a}, \bar{N})$, then there exist a $\prec_{\mathcal{K}}$-mapping f and a continuous tower $(\bar{M}^*, \bar{a}, \bar{N})$ extending $(\bar{M}, \bar{a}, \bar{N})$ so that $f(M'_i) \preceq_{\mathcal{K}} M^*_i$ for all $i < \beta$.

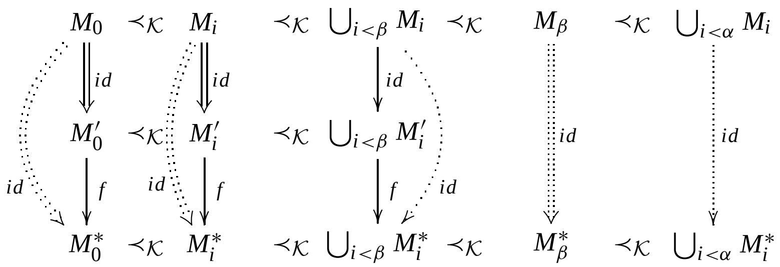

The proof of Theorem II.7.1 is by induction on α. Notice that for $\alpha \leq \omega$, there is little to do since all towers of length $\leq \omega$ are vacuously continuous. If α is the successor of a successor, then the induction hypothesis and the extension property for non-μ-splitting types (Theorem I.4.10) produce a continuous extension. We take care of the case that α is a limit ordinal by taking direct limits of partial continuous extensions. The difficult case is when α is the successor of a limit ordinal. This case employs Hypothesis 1. We will build an increasing chain of continuous towers throwing in a particular element at each stage so that in the end we will have added enough (μ-many, predetermined) elements to have a universal extension over $\bigcup_{i<\delta} M_i$. The following proposition allows us to add in the new elements in this stage of the inductive proof of Theorem II.7.1 (when $\alpha = \delta + 1$ and δ is a limit ordinal).

Proposition II.7.2. *Suppose that Theorem II.7.1 holds for all amalgamable towers of length δ for some limit ordinal $\delta < \mu^+$. Let $(\bar{M}, \bar{a}, \bar{N})$ be an amalgamable tower of length δ inside $\mathfrak{C}$. For every $b \in \mathfrak{C}$, there exists a continuous, amalgamable tower $(\bar{M}^*, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\delta}$ inside $\mathfrak{C}$ such that $b \in \bigcup_{i<\delta} M^*_i$ and $(\bar{M}, \bar{a}, \bar{N}) <^c_{\mu,\delta} (\bar{M}^*, \bar{a}, \bar{N})$.*

Furthermore, if $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\beta}$ is a continuous partial extension of $(\bar{M}, \bar{a}, \bar{N})$, we can choose $(\bar{M}^*, \bar{a}, \bar{N})$ such that there exists a $\prec_{\mathcal{K}}$-mapping f with $f(M'_i) \preceq_{\mathcal{K}} M^*_i$ for all $i < \beta$.

Proof. We begin by defining by induction on $\zeta < \delta$ a $<^c_{\mu,\delta}$-increasing and continuous sequence of towers, $\langle(\bar{M}, \bar{a}, \bar{N})^\zeta \in {}^+\mathcal{K}^*_{\mu,\delta} \mid \zeta \leq \delta\rangle$ such that

(1) $(\bar{M}, \bar{a}, \bar{N}) \leq^c_{\mu,\delta} (\bar{M}, \bar{a}, \bar{N})^0$,
(2) $(\bar{M}, \bar{a}, \bar{N})^\zeta$ is continuous and
(3) if we are given $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\beta}$ a continuous partial extension of $(\bar{M}, \bar{a}, \bar{N})$, then there is a $\prec_{\mathcal{K}}$-mapping f with $f(M'_i) \preceq_{\mathcal{K}} M^0_i$ for all $i < \beta$.

This produces a δ-by-$(\delta + 1)$-array of models which we will diagonalize.

Why is this construction possible? Since $(\bar{M}, \bar{a}, \bar{N})$ is amalgamable, by the hypothesis of the proposition, $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$ has a continuous extension $(\bar{M}^0, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\delta}$. Furthermore, if we are given $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\beta}$ as above, then by condition (2) of Theorem II.7.1, we may find f such that $f(M'_i) \preceq_{\mathcal{K}} M^0_i$ for all $i < \beta$. At successor stages we can find continuous extensions by the hypothesis of the proposition and the fact that continuous towers are nice. When ζ is a limit ordinal, we take unions. The unions will be continuous, since the union of an increasing chain of continuous towers is continuous.

Since $\bigcup_{i<\delta} M_i$ is an amalgamation base, we can find an isomorphic copy of this chain of towers inside $\mathfrak{C}$. WLOG, for $\zeta < \delta$, $M^\delta_\zeta \prec_{\mathcal{K}} \mathfrak{C}$.

Consider the diagonal sequence $\langle M^\zeta_\zeta \mid \zeta < \delta\rangle$. Notice that this is a $\prec_{\mathcal{K}}$-increasing sequence of amalgamation bases. For $\zeta < \delta$, we have $M^{\zeta+1}_{\zeta+1}$ is universal over M^ζ_ζ. Why? From the definition of $<^c$, $M^{\zeta+1}_\zeta$ is universal over M^ζ_ζ. Since $M^{\zeta+1}_\zeta \prec_{\mathcal{K}} M^{\zeta+1}_{\zeta+1}$, we have that $M^{\zeta+1}_{\zeta+1}$ is also universal over M^ζ_ζ (see Proposition I.5.13).

By construction, each $\bar{M}^\zeta$ is continuous. Thus the sequence $\langle M^\zeta_\zeta \mid \zeta < \delta\rangle$ is continuous. Then $\langle M^\zeta_\zeta \mid \zeta < \delta\rangle$ witnesses that $\bigcup_{\zeta<\delta} M^\zeta_\zeta$ is a (μ, δ)-limit model. Let M^b be a limit model inside $\mathfrak{C}$ that is universal over $\bigcup_{\zeta<\delta} M^\zeta_\zeta$ and contains b.

Because $\bigcup_{\zeta<\delta} M^\zeta_\zeta$ is a limit model, we can apply Fact I.4.7 to ga-tp$\left(b/\bigcup_{\zeta<\delta} M^\zeta_\zeta, M^\delta_\delta\right)$. Let $\xi < \delta$ be such that

$$\text{ga-tp}\left(b \Big/ \bigcup_{\zeta<\delta} M^\zeta_\zeta, M^b\right) \text{ does not } \mu\text{-split over } M^\xi_\xi. \qquad (*)$$

Notice that $(\langle M^i_i \mid i < \xi\rangle, \bar{a}, \bar{N}) \restriction \xi$ is a $<^c_{\mu,\xi}$-extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \xi$.

We will find a $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a}, \bar{N})$ by defining a $\prec_{\mathcal{K}}$-increasing chain of models $\langle N^*_i \mid i < \alpha\rangle$ and an increasing chain of $\prec_{\mathcal{K}}$-mappings $\langle h_i \mid i < \alpha\rangle$ with the intention that the pre-image of N^*_i under an extension of $\bigcup_{i<\alpha} h_i$ will form a sequence $\bar{M}^*$ such that $(\bar{M}, \bar{a}, \bar{N}) <^c_{\mu,\delta} (\bar{M}^*, \bar{a}, \bar{N})$, $b \in M^*_{\xi+1}$ and $M^*_i = M^i_i$ for all $i < \xi$. We choose by induction on $i < \delta$ a $\prec_{\mathcal{K}}$-increasing and continuous chain of limit models $\langle N^*_i \in \mathcal{K}_\mu \mid i < \delta\rangle$ and an increasing and continuous sequence of $\prec_{\mathcal{K}}$-mappings $\langle h_i \mid i < \delta\rangle$ satisfying

(1) N^*_{i+1} is a limit model and is universal over N^*_i
(2) $h_i : M^i_i \to N^*_i$
(3) $h_i(M^i_i) \prec_{\mathcal{K}} M^{i+1}_i$
(4) ga-tp$(h_{i+1}(a_i/N^*_i)$ does not μ-split over $h_i(N_i)$
(5) $M^b \prec_{\mathcal{K}} N^*_{\xi+1}$ and
(6) for $i \leq \xi$, $N^*_i = M^i_i$ with $h_i = id_{M^i_i}$.

We depict the construction below. The inverse image of the sequence of N^*'s will form the required continuous $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a}, \bar{N})$.

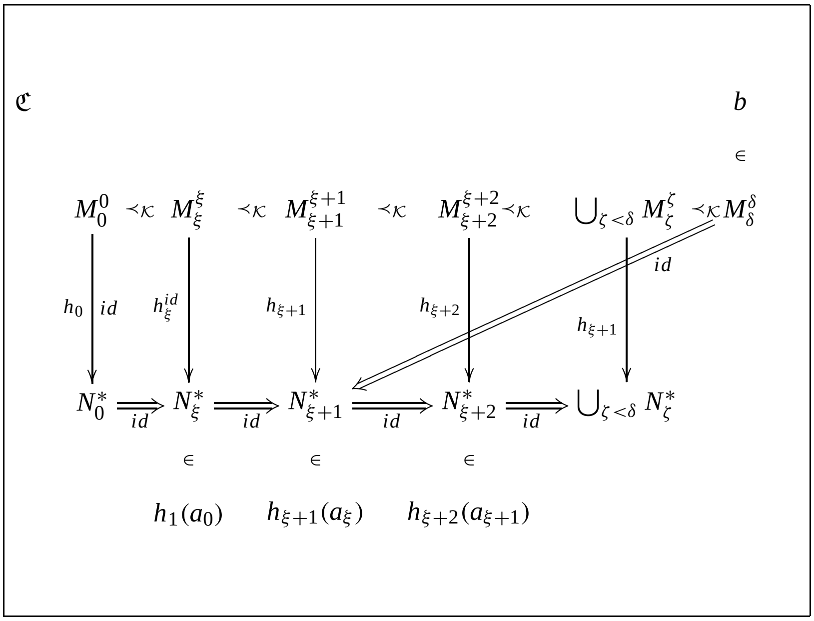

The requirements determine the definition of N_i^* for $i \le \xi$. We proceed with the rest of the construction by induction on i. If i is a limit ordinal $\ge \xi$, let $N_i^* = \bigcup_{j<i} N_j^*$ and $h_i = \bigcup_{j<i} h_j$.

Suppose that we have defined h_i and N_i^* satisfying the conditions of the construction. We now describe how to define N_{i+1}^*. First, we extend h_i to $\bar{h}_i \in \mathrm{Aut}(\mathfrak{C})$. We can assume that $\bar{h}_i(a_i) \in M_{i+1}^{i+2}$. This is possible since M_{i+1}^{i+2} is universal over $h_i(M_i^i)$ by construction.

Since ga-tp(a_i/M_i^i) does not μ-split over N_i, by invariance we have that ga-tp$(\bar{h}_i(a_i)/h_i(M_i^i))$ does not μ-split over $h_i(N_i)$. We now adjust the proof of the existence property for non-splitting extensions.

Claim II.7.3. *We can find $g \in \mathrm{Aut}(\mathfrak{C})$ such that* ga-tp$(g(\bar{h}_i(a_i))/N_i^*)$ *does not μ-split over $h_i(N_i)$ and* $g(\bar{h}_i(M_{i+1}^{i+1})) \prec_{\mathcal{K}} M_{i+1}^{i+2}$.

Proof of Claim II.7.3. First we find a $\prec_{\mathcal{K}}$-mapping f such that $f : N_i^* \to h_i(M_i^i)$ such that $f \restriction h_i(N_i) = id_{h_i(N_i)}$ which is possible since $h_i(M_i^i)$ is universal over $h_i(N_i)$. Notice that ga-tp$(f^{-1}(\bar{h}_i(a_i))/N_i^*)$ does not μ-split over $h_i(N_i)$ and

$$\text{ga-tp}(f^{-1}(\bar{h}_i(a_i))/h_i(M_i^i)) = \text{ga-tp}(\bar{h}_i(a_i)/h_i(M_i^i)) \tag{+}$$

by a non-splitting argument as in the proof of Theorem I.4.12.

Let N^+ be a limit model of cardinality μ containing $f^{-1}(\bar{h}_i(a_i))$ with $f^{-1}(\bar{h}_i(M_{i+1}^{i+1})) \prec_{\mathcal{K}} N^+$. Now using the equality of types (+) and the fact that M_{i+1}^{i+2} is universal over $h_i(M_i^i)$ with $\bar{h}_i(a_i) \in M_{i+1}^{i+2}$, we can find a $\prec_{\mathcal{K}}$-mapping $f^+ : N^+ \to M_{i+1}^{i+2}$ such that $f^+ \restriction h_i(M_i^i) = id_{h_i(M_i^i)}$ and $f^+(f^{-1}(\bar{h}_i(a_i))) = \bar{h}_i(a_i)$. Now set $g := f^+ \circ f^{-1} : \bar{h}(M_{i+1}^{i+1}) \to M_{i+1}^{i+2}$. □

Fix such a g as in the claim and set $h_{i+1} := g \circ \bar{h}_i \restriction M_{i+1}^{i+1}$. Let N_{i+1}^* be a $\prec_{\mathcal{K}}$ extension of N_i^*, M^b and $h_{i+1}(M_{i+1}^{i+1})$ of cardinality μ inside $\mathfrak{C}$. Choose N_{i+1}^* to additionally be a limit model and universal over N_i^*.

This completes the construction.

We now argue that the construction of these sequences is enough to find a $<^c_{\mu,\delta}$-extension, $(\bar{M}^*, \bar{a}, \bar{N})$, of $(\bar{M}, \bar{a}, \bar{N})$ such that $b \in M_\zeta^*$ for some $\zeta < \delta$.

Let $h_\delta := \bigcup_{i<\delta} h_i$. We will be defining for $i < \delta$, M_i^* to be pre-image of N_i^* under some extension of h_δ. The following claim allows us to choose the pre-image so that M_ζ^* contains b for some $\zeta < \delta$.

Claim II.7.4. *There exists $h \in \mathrm{Aut}(\mathfrak{C})$ extending $\bigcup_{i<\delta} h_i$ such that $h(b) = b$.*

Proof of Claim II.7.4. Let $h_\delta := \bigcup_{i<\delta} h_i$. Consider the increasing and continuous sequence $\langle h_\delta(M_i^i) \mid i < \delta\rangle$. By invariance, $h_\delta(M_{i+1}^{i+1})$ is universal over $h_\delta(M_i^i)$ and each $h_\delta(M_i^i)$ is a limit model.

Furthermore, from our choice of ξ, we know that ga-tp(b/M_i^δ) does not μ-split over M_ξ^ξ. Since $h_i(M_i^i) \prec_{\mathcal{K}} M_i^{i+1} \prec_{\mathcal{K}} \bigcup_{j<\delta} M_j^\delta$, monotonicity of non-splitting allows us to conclude that

$$\text{ga-tp}(b/h_\delta(M_i^i)) \text{ does not } \mu\text{-split over } M_\xi^\xi.$$

This allows us to apply Fact I.4.8, to ga-tp $\left(b/\bigcup_{i<\delta} h_\delta(M_i^i)\right)$ yielding

$$\text{ga-tp}\left(b \middle/ \bigcup_{i<\delta} h_\delta(M_i^i)\right) \text{ does not } \mu\text{-split over } M_\xi^\xi. \tag{**}$$

Notice that $\bigcup_{i<\delta} M_i^i$ is a limit model witnessed by $\langle M_j^j \mid j < \delta\rangle$. So we can apply Corollary I.2.20 and extend $\bigcup_{i<\delta} h_i$ to an automorphism h^* of $\mathfrak{C}$. We will first show that

$$\text{ga-tp}\left(b \middle/ h^*\left(\bigcup_{i<\delta} M_i^i\right), \mathfrak{C}\right) = \text{ga-tp}\left(h^*(b) \middle/ h^*\left(\bigcup_{i<\delta} M_i^i\right), \mathfrak{C}\right). \tag{***}$$

By invariance and our choice of ξ in $(*)$,

$$\text{ga-tp}\left(h^*(b) \Big/ h^*\left(\bigcup_{i<\delta} M_i^i\right), \mathfrak{C}\right) \text{ does not } \mu\text{-split over } M_\xi^\xi.$$

We will use non-splitting to derive $(***)$. To make the application of non-splitting more transparent, let $N^1 := \bigcup_{i<\delta} M_i^i$, $N^2 := h^*\left(\bigcup_{i<\delta} M_i^i\right)$ and $p := \text{ga-tp}(b/N^2)$. By $(**)$, we have that $p \restriction N^2 = h^*(p \restriction N^1)$. In other words,

$$\text{ga-tp}\left(b \Big/ h^*\left(\bigcup_{i<\delta} M_i^i\right), \mathfrak{C}\right) = \text{ga-tp}\left(h^*(b) \Big/ h^*\left(\bigcup_{i<\delta} M_i^i\right), \mathfrak{C}\right),$$

as desired.

From $(***)$ and Corollary I.2.25, we can find an automorphism f of $\mathfrak{C}$ such that $f(h^*(b)) = b$ and $f \restriction h^*\left(\bigcup_{i<\delta} M_i^i\right) = id_{h^*(\bigcup_{i<\delta} M_i^i)}$. Notice that $h := f \circ h^*$ satisfies the conditions of the claim. □

Now that we have an automorphism h fixing b and $\bigcup_{i<\delta} M_i$, we can define for each $i < \delta$, $M_i^* := h^{-1}(N_i^*)$.

Claim II.7.5. *$(\bar{M}^*, \bar{a}, \bar{N})$ is a $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a}, \bar{N})$ such that $b \in M^*_{\xi+1}$.*

Proof of Claim II.7.5. By construction $b \in M_\delta^\delta \subseteq N^*_{\xi+1}$. Since $h(b) = b$, this implies $b \in M^*_{\xi+1}$. To verify that we have a $\leq^c_{\mu,\delta}$-extension we need to show for $i < \delta$:

i. M_i^* is universal over M_i
ii. $a_i \in M^*_{i+1} \backslash M_i$ for $i + 1 < \delta$ and
iii. ga-tp(a_i/M_i^*) does not μ-split over N_i whenever $i, i + 1 \leq \delta$.

Item i follows from the fact that M_i^i is universal over M_i and $M_i^i \prec_{\mathcal{K}} M_i^*$. Item iii follows from invariance and our construction of the N_i^*'s. Finally, recalling that a non-splitting extension of a non-algebraic type is also non-algebraic (Remark I.4.13) we see that Item iii implies $a_i \notin M_i^*$. By our choice of $h_{i+1}(a_i) \in M^{i+2}_{i+1} \prec_{\mathcal{K}} N^*_{i+1}$, we have that $a_i \in M^*_{i+1}$. Thus Item ii is satisfied as well. □ □

Before beginning the proof of Theorem II.7.1, recall that we will be building a directed system of partial extensions to take care of the induction step when α is a limit ordinal. Let us establish a few facts about directed systems here. Using the axioms of AEC and Shelah's Presentation Theorem, one can show that Axiom 5 of the definition of AEC has an alternative formulation (see [24] or Chapter 13 of [7]):

Definition II.7.6. A partially ordered set $(I, \leq)$ is *directed* iff for every $a, b \in I$, there exists $c \in I$ such that $a \leq c$ and $b \leq c$.

Fact II.7.7 (*P.M. Cohn 1965*). *Let $(I, \leq)$ be a directed set. If $\langle M_t \mid t \in I\rangle$ and $\{h_{t,r} \mid t \leq r \in I\}$ are such that*

(1) *for $t \in I$, $M_t \in \mathcal{K}$*
(2) *for $t \leq r \in I$, $h_{t,r} : M_t \to M_r$ is a $\prec_{\mathcal{K}}$-embedding and*
(3) *for $t_1 \leq t_2 \leq t_3 \in I$, $h_{t_1,t_3} = h_{t_2,t_3} \circ h_{t_1,t_2}$ and $h_{t,t} = id_{M_t}$,*

then, whenever $s = \lim_{t\in I} t$, there exist $M_s \in \mathcal{K}$ and $\prec_{\mathcal{K}}$-mappings $\{h_{t,s} \mid t \in I\}$ such that

$$h_{t,s} : M_t \to M_s,\ M_s = \bigcup_{t<s} h_{t,s}(M_t) \text{ and}$$
$$\text{for } t_1 \leq t_2 \leq s,\ h_{t_1,s} = h_{t_2,s} \circ h_{t_1,t_2} \text{ and } h_{s,s} = id_{M_s}.$$

Definition II.7.8. (1) $(\langle M_t \mid t \in I\rangle, \{h_{t,s} \mid t \leq s \in I\})$ from Fact II.7.7 is called a *directed system*.
(2) We say that M_s together with $\langle h_{t,s} \mid t \leq s\rangle$ satisfying the conclusion of Fact II.7.7 is *a direct limit of* $(\langle M_t \mid t < s\rangle, \{h_{t,r} \mid t \leq r < s\})$.

Later we will generalize these systems by producing directed systems of towers instead of models.

Now we use Proposition II.7.2 to prove Theorem II.7.1.

Proof of Theorem II.7.1. We prove that every amalgamable tower has a continuous extension by induction on α.
$\boldsymbol{\alpha = 0}$**:** By Theorem I.3.13 and Corollary I.3.14, we can find a (μ, ω)-limit over M_0. Fix such a model and call it M'_0.
$\boldsymbol{\alpha = \delta + 1}$ **and** $\boldsymbol{\delta}$ **is a limit ordinal:** The strategy is to start out with a continuous extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$ (which we call $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$.) If we are lucky, the top of $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ will be universal over M_δ. Since this cannot be guaranteed, we will repeatedly add new elements into extensions of $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ until the top of one of these extensions is universal over M_δ.

By the induction hypothesis, we can find $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta) \in {}^{+}\mathcal{K}^*_{\mu,\delta}$ such that

- $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ is a $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$ and
- and if $(\bar{M}', \bar{a} \restriction \beta, \bar{N} \restriction \beta)$ is a continuous $<^c_{\mu,\beta}$-extension of $(\bar{M}, \bar{a}, \bar{N})$, then we can choose $\bar{M}^{**}$ such that there exists a $\prec_{\mathcal{K}}$-mapping f with $f(M'_i) \prec_{\mathcal{K}} M^{**}_i$ for all $i < \beta$.

Notice that since $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ is continuous, we can apply the induction hypothesis δ-many times to find an $<^c_{\mu,\delta}$-increasing chain of continuous towers of length δ. In addition to being continuous, the top of this chain will be an amalgamable extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$. Why? The top of this tower will be a (μ, δ)-limit model witnessed by the diagonal. Thus WLOG we may assume that $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ is amalgamable and continuous.

We construct a continuous $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a}, \bar{N})$ by the induction hypothesis and repeated applications of Proposition II.7.2.

Let M'_δ be a limit model and universal over M_δ inside $\mathfrak{C}$. Enumerate M'_δ as $\{b_\zeta \mid \zeta < \delta\mu\}$. We will add these elements into extensions of $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ by defining by induction on $\zeta \leq \delta\mu$ a $<^c_{\mu,\delta}$-increasing and continuous chain of towers $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\zeta \in {}^{+}\mathcal{K}^*_{\mu,\delta}$ such that

(1) $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\zeta$ is a $<^c_{\mu,\delta}$-extension of $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$
(2) $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\zeta$ is continuous and
(3) $b_\zeta \in \bigcup_{i<\delta} M^{\zeta+1}_i \prec_{\mathcal{K}} \mathfrak{C}$.

The following diagram depicts the construction:

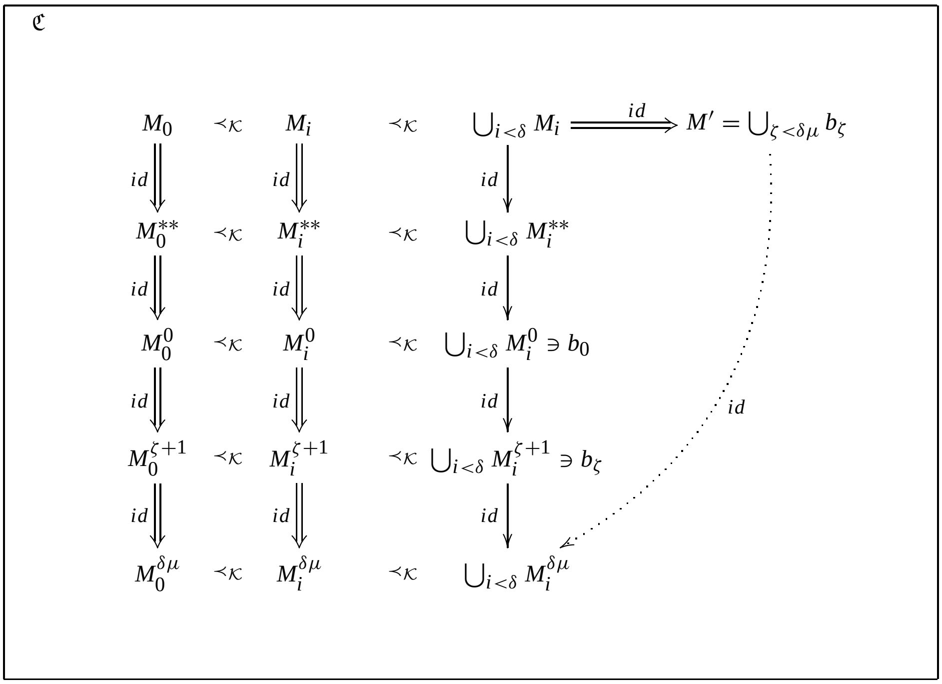

The construction is possible by the induction hypothesis and Proposition II.7.2:

$\zeta = 0$: Since $\bigcup_{i<\delta} M_i^{**}$ is an amalgamation base, we can apply Proposition II.7.2 and find a $<^c_{\mu,\delta}$-extension $(\bar{M}^0, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ in $\mathfrak{C}$ such that $b_0 \in \bigcup_{i<\delta} M_i^0$.
$\zeta + 1$: Suppose that $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\zeta$ has been defined. It is a continuous tower of length δ. If $\bigcup_{i<\delta} M_i^\zeta$ is an amalgamation base, by the induction hypothesis we can apply Proposition II.7.2 to find a $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\zeta$, say $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^{\zeta+1}$ inside $\mathfrak{C}$ such that $b_\zeta \in \bigcup_{i<\delta} M_i^{\zeta+1}$.

Suppose on the other hand, that $\bigcup_{i<\delta} M_i^\zeta$ is not an amalgamation base. This may occur when ζ is a limit ordinal of a different cofinality than the cofinality of δ. By Hypothesis 1, there is an amalgamable extension of $(\bar{M}, \bar{a}, \bar{N})^\zeta$ inside $\mathfrak{C}$. Apply Proposition II.7.2 to the amalgamable extension and b_ζ. The proposition will produce a $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\zeta$, say $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^{\zeta+1}$ inside $\mathfrak{C}$ such that $b_\zeta \in \bigcup_{i<\delta} M_i^{\zeta+1}$.
ζ a limit ordinal: If ζ is a limit ordinal we can set $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\zeta := \bigcup_{\xi<\zeta} (\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\xi$. It is a continuous tower since all the $(\bar{M}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)^\xi$'s are continuous. This completes the construction.

Now consider the tower $(\bar{M}^*, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\delta+1}$ defined by $M_i^* := M_i^{\delta\mu}$ for all $i < \delta$ and $M_\delta^* := \bigcup_{i<\delta} M_i^{\delta\mu}$. Since M_δ^* contains M'_δ, it is universal over M_δ. Thus $(\bar{M}^*, \bar{a}, \bar{N})$ is a $<^c_{\mu,\delta+1}$-extension of $(\bar{M}, \bar{a}, \bar{N})$. Since $(\bar{M}^{\delta\mu}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ is continuous, we have that $(\bar{M}^*, \bar{a}, \bar{N})$ is also continuous. Notice that $(\bar{M}^*, \bar{a}, \bar{N})$ is amalgamable as well. By construction for every $i < \delta$, M_δ^* is a limit model. For the case $i = \delta$, we see that M_δ^* is a (μ, δ)-limit model witnessed by the diagonal $\langle M_i^{i\mu} \mid i < \delta \rangle$.
$\alpha = \delta + 1$ and δ is a successor ordinal: By the induction hypothesis we can find a continuous, amalgamable extension $(\bar{M}^{**}, \bar{a} \restriction \delta, \bar{N} \restriction \delta)$ of $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$ and if we are given $(\bar{M}', \bar{a} \restriction \beta, \bar{N} \restriction \beta)$ as in part (2) of the statement of the theorem, we may assume that there is a $\prec_\mathcal{K}$-mapping f^* such that $f^*(M'_i) \prec_\mathcal{K} M_i^{**}$ for all $i < \beta$.

Since $M_{\delta-1}^{**}$ and M_δ are both $\mathcal{K}$-substructures of $\mathfrak{C}$, we can apply the Downward-Löwenheim Axiom for AECs to find M_δ^{**} (a first approximation to M_δ^*) a model of cardinality μ extending both $M_{\delta-1}^{**}$ and M_δ. WLOG by Theorem I.2.17 and Lemma I.2.24 we may assume that M_δ^{**} is a limit model of cardinality μ and M_δ^{**} is universal over both $M_{\delta-1}^{**}$ and M_δ. By Theorem I.4.10, we can find a $\prec_\mathcal{K}$-mapping $h : M_\delta^{**} \to \mathfrak{C}$ such that $h \restriction M_\delta = id_{M_\delta}$ and ga-tp$(a_\delta / h(M_\delta^{**}))$ does not μ-split over N_δ. Set $M_i^* := h(M_i^{**})$ for all $i \leq \delta$. Notice that by invariance $(\bar{M}^*, \bar{a}, \bar{N}) \restriction \delta$ is a $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a}, \bar{N})$. To conclude that $(\bar{M}^*, \bar{a}, \bar{N})$ is the required $<^c_{\mu,\alpha}$-extension of $(\bar{M}, \bar{a}, \bar{N})$ with $f = h \circ f^*$ if appropriate, it remains to check that

Subclaim II.7.9. $a_\delta \notin M_\delta^*$.

Proof of Subclaim II.7.9. Suppose that $a_\delta \in M_\delta^*$. Since M_δ is universal over N_δ, there exists a $\prec_\mathcal{K}$-mapping, $g : M_\delta^* \to M_\delta$ such that $g \restriction N_\delta = id_{N_\delta}$. Since ga-tp$(a_\delta / M_\delta^*)$ does not μ-split over N_δ, we have that

$$\text{ga-tp}(a_\delta / g(M_\delta^*)) = \text{ga-tp}(g(a_\delta) / g(M_\delta^*)). \quad (*)$$

Notice that because $g(a_\delta) \in g(M_\delta^*)$, $(*)$ implies that $a_\delta = g(a_\delta)$. Thus $a_\delta \in g(M_\delta^*) \prec_\mathcal{K} M_\delta$. This contradicts the definition of towers: $a_\delta \notin M_\delta$.

α is a limit ordinal $> \omega$: We will construct a directed system of partial extensions of $(\bar{M}, \bar{a}, \bar{N})$, $\langle (\bar{M}, \bar{a}, \bar{N})^\zeta \mid \zeta < \alpha \rangle$ and $\langle f_{\xi,\zeta} \mid \xi \leq \zeta < \alpha \rangle$ satisfying the following conditions:

(1) $(\bar{M}, \bar{a}, \bar{N}) \restriction \zeta <^c_{\mu,\zeta} (\bar{M}, \bar{a}, \bar{N})^\zeta$
(2) $(\bar{M}, \bar{a}, \bar{N})^\zeta$ is continuous
(3) $(\bar{M}, \bar{a}, \bar{N})^\zeta$ lies in $\mathfrak{C}$
(4) $f_{\xi,\zeta} \restriction M_i^\xi : M_i^\xi \to M_i^\zeta$ for $i < \xi \leq \zeta$
(5) for all $\xi < \zeta$, M_ξ^ζ is universal over $f_{\xi,\zeta}\left(\bigcup_{i<\xi} M_i^\xi\right)$ and
(6) $f_{\xi,\zeta} \restriction M_\xi = id_{M_\xi}$ for all $\xi < \zeta < \alpha$.

The construction is possible by the induction hypothesis and Proposition II.7.2. We provide the details here.
$\zeta = 0$: Set $\bar{M}^0$ equal to the empty sequence and $f_{0,0}$ equal to the empty mapping.
$\zeta = \xi + 1$: Suppose that $(\bar{M}, \bar{a}, \bar{N})^\xi$ and $\langle f_{\gamma,\gamma'} \mid \gamma \leq \gamma' \leq \xi \rangle$ have been defined accordingly. Then by the induction hypothesis applied to $(\bar{M}, \bar{a}, \bar{N}) \restriction \zeta$ and the partial extension $(\bar{M}, \bar{a}, \bar{N})^\xi$, we can find a $\prec_\mathcal{K}$-mapping f and a continuous extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \zeta$. By applying the induction hypothesis again to this continuous extension, we

can find $(\bar{M}, \bar{a}, \bar{N})^\zeta \in {}^+\mathcal{K}^*_{\mu,\zeta}$ inside $\mathfrak{C}$ such that for all $i < \xi$, $f(M^\xi_i) \prec_\mathcal{K} M^\zeta_i$, $f \restriction M_i = id_{M_i}$ and M^ζ is universal over $f\left(\bigcup_{i<\xi} M^\xi_i\right)$. Notice that by setting $f_{\gamma,\xi+1} = f \circ f_{\gamma,\xi}$ and $f_{\zeta,\zeta} = id_{\bigcup_{\xi<\zeta} M^\zeta_\xi}$ we have completed the successor stage of the construction.

ζ a limit ordinal: By the induction hypothesis we have constructed a directed system $\langle \bigcup_{i<\gamma} M^\gamma_i \mid \gamma < \zeta\rangle$ with $\langle f_{\gamma,\xi} \mid \gamma \le \xi < \zeta\rangle$. By Fact II.7.7 we can find a direct limit to this system, $M^{**}_\zeta \in \mathcal{K}$ and $\prec_\mathcal{K}$-mappings $\langle f^{**}_{\gamma,\zeta} \mid \gamma \le \zeta\rangle$. First notice that

Subclaim II.7.10. *$\langle f^{**}_{\gamma,\zeta} \restriction M_\gamma \mid \gamma < \zeta\rangle$ is increasing.*

Proof. Let $\gamma < \xi < \zeta$ be given. By construction

$$f_{\gamma,\xi} \restriction M_\gamma = id_{M_\gamma}.$$

An application of $f^{**}_{\xi,\zeta}$ yields

$$f^{**}_{\xi,\zeta} \circ f_{\gamma,\xi} \restriction M_\gamma = f^{**}_{\gamma,\zeta} \restriction M_\gamma.$$

Since $f^{**}_{\gamma,\zeta}$ and $f^{**}_{\xi,\zeta}$ come from a direct limit of the system which includes the mapping $f_{\gamma,\xi}$, we have

$$f^{**}_{\gamma,\zeta} \restriction M_\gamma = f^{**}_{\xi,\zeta} \circ f_{\gamma,\zeta} \restriction M_\gamma.$$

Combining the equalities yields

$$f^{**}_{\gamma,\zeta} \restriction M_\gamma = f^{**}_{\xi,\zeta} \restriction M_\gamma.$$

This completes the proof of Subclaim II.7.10. □

By the subclaim, we have that $f := \bigcup_{\gamma<\zeta} f^{**}_{\gamma,\zeta} \restriction M_\gamma$ is a $\prec_\mathcal{K}$-mapping from $\bigcup_{\gamma<\zeta} M_\gamma$ onto $\bigcup_{\gamma<\zeta} f^{**}_{\gamma,\zeta}(M_\gamma)$. Since $\mathfrak{C}$ is a (μ,μ^+)-limit model and since $\bigcup_{\gamma<\zeta} M_\gamma$ is an amalgamation base (as $(\bar{M}, \bar{a}, \bar{N})$ is nice) we can assume that f is a partial automorphism of $\mathfrak{C}$ and extend it to $F \in \mathrm{Aut}(\mathfrak{C})$ by Corollary I.2.20.

Now consider the direct limit defined by $M^\zeta_\zeta := F^{-1}(M^{**}_\zeta)$ with $\langle f^*_{\xi,\zeta} := F^{-1} \circ f^{**}_{\xi,\zeta} \mid \xi < \zeta\rangle$ and $f^*_{\zeta,\zeta} = id_{M^*_\zeta}$. Let $M^\zeta_i := f_{\xi,\zeta}(M^\xi_i)$ for all $i < \xi$. This is well-defined since $f_{\xi,\zeta}$ is part of the direct limit of a directed system. Notice that $f^*_{\xi,\zeta} \restriction M_\xi = F^{-1} \circ f^{**}_{\xi,\zeta} \restriction M_\xi = id_{M_\xi}$ for $\xi < \zeta$.

Subclaim II.7.11. $(\bar{M}, \bar{a}, \bar{N}) \restriction \zeta <^c_{\mu,\zeta} (\bar{M}, \bar{a}, \bar{N})^\zeta$.

Proof of Subclaim II.7.11. We need to verify that for all $\xi < \zeta$,

(1) $M^\zeta_\xi \prec_\mathcal{K} M^\zeta_{\xi+1}$,
(2) $a_\xi \in M^\zeta_{\xi+1} \backslash M^\zeta_\xi$ and
(3) ga-tp(a_ξ / M^ζ_ξ) does not μ-split over N_ξ.

To see that $\bar{M}^\zeta$ is increasing, by the induction hypothesis,

$$f_{\xi,\xi+1}\left(\bigcup_{i<\xi} M^\xi_i\right) \prec_\mathcal{K} M^{\xi+1}_\xi.$$

Applying $f_{\xi+1,\zeta}$ to both sides of this equation gives us for every $j < \xi$,

$$M^\zeta_j \prec_\mathcal{K} f_{\xi,\zeta}\left(\bigcup_{i<\xi} M^\xi_i\right) = f_{\xi+1,\zeta}\left(f_{\xi,\xi+1}\left(\bigcup_{i<\xi} M^\xi_i\right)\right) \prec_\mathcal{K} f_{\xi+1,\zeta}(M^{\xi+1}_\xi) = M^\zeta_\xi.$$

By the induction hypothesis for all $\xi < \zeta$, $a_\xi \notin M^{\xi+2}_\xi$ and ga-tp$(a_\xi / M^{\xi+2}_\xi)$ does not μ-split over N_ξ. Since $f_{\xi+2,\zeta} \restriction M_{\xi+1} = id_{M_{\xi+1}}$, invariance gives us $f_{\xi+2,\zeta}(a_\xi) = a_\xi \notin f_{\xi+2,\xi}(M^{\xi+2}_\xi) = M^\zeta_\xi$ and ga-tp(a_ξ / M^ζ_ξ) does not μ-split over N_ξ. □

Notice that $(\bar{M}, \bar{a}, \bar{N})^\zeta$ is continuous since it is formed from the direct limit of a continuous system. To see that $(\bar{M}, \bar{a}, \bar{N})^\zeta$ is amalgamable, notice that condition (5) of the construction guarantees that $\bigcup_{\xi<\zeta} M^\zeta_\xi$ is a (μ, ζ)-limit witnessed by $\langle f_{\xi,\zeta}\left(\bigcup_{i<\xi} M^\xi_i\right) \mid \xi < \zeta\rangle$. This completes the construction.

Why is the construction sufficient to produce $(\bar{M}', \bar{a}, \bar{N})$ as required? We have constructed a directed system $\langle\bigcup_{i<\gamma} M^\gamma_i \mid \gamma \leq \xi < \alpha\rangle$ with $\langle f_{\gamma,\xi} \mid \gamma \leq \xi < \alpha\rangle$. By Fact II.7.7 and Subclaim II.7.10 we can find a direct limit to this system, M^*_α and $\prec_{\mathcal{K}}$-mappings $\langle f_{\gamma,\alpha} \mid \gamma \leq \alpha\rangle$ such that $f_{\gamma,\alpha} \restriction M_i = id_{M_i}$ for all $i < \alpha$. If $(\bar{M}, \bar{a}, \bar{N})$ is amalgamable, then M^*_α can be chosen to lie in $\mathfrak{C}$. Define for all $\zeta < \alpha$, $M^*_\zeta := f_{\zeta+1,\alpha}(M^{\zeta+1}_\zeta)$. Notice that as in Subclaim II.7.11, $(\bar{M}, \bar{a}, \bar{N}) <^c_{\mu,\alpha} (\bar{M}^*, \bar{a}, \bar{N})$. And, as in the limit stage of the construction, we see that $(\bar{M}^*, \bar{a}, \bar{N})$ is continuous and amalgamable.

The second part of the statement of the theorem is obtained by modifying our construction by setting $(\bar{M}, \bar{a}, \bar{N})^\beta = (\bar{M}', \bar{a}, \bar{N})$ and proceeding with the construction from $\beta + 1$. □

8. Refined orderings on towers

In this section we further develop the machinery of towers which will be used to construct a relatively full tower in Section 9.

Definition II.8.1. For ordinals $\alpha, \alpha', \delta, \delta' < \mu^+$ with $\alpha \leq \alpha'$ and $\delta \leq \delta'$. We say that $(\bar{M}', \bar{a}', \bar{N}') \in {}^+\mathcal{K}^*_{\mu,\alpha'\times\delta'}$ is a $<^c$-extension of $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha\times\delta}$ iff

- for every $\beta < \alpha$ and every $i < \delta$, $M'_{\beta,i}$ is universal over $M_{\beta,i}$
- for every $\beta < \alpha$ and $i + 1 < \delta$, $a_{\beta,i} = a'_{\beta,i}$ and $N_{\beta,i} = N'_{\beta,i}$.

The following theorem is used to construct relatively full towers by adding realizations of strong types between $M_{\beta,i}$ and $M_{\beta+1,0}$ in an $<^c$-extension of the tower $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha\times\delta}$.

Theorem II.8.2. *Under Hypothesis 1, given α an ordinal $< \mu^+$ and a nice tower, $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha\times\mu\alpha}$, we can find an amalgamable, continuous extension $(\bar{M}', \bar{a}', \bar{N}') \in {}^+\mathcal{K}^*_{\mu,\alpha+1\times\mu(\alpha+1)}$ of $(\bar{M}, \bar{a}, \bar{N})$ such that for a fixed enumeration, $\{(p, N)^\zeta_l \mid l < \mu\}$, of $\bigcup_{i<\mu\alpha} \mathfrak{St}(M_{\zeta,i})$ for each $\zeta < \alpha$, we have that*

$$(p, N)^\zeta_l \sim (\text{ga-tp}(a_{\zeta+1,l+1}/M'_{\zeta+1,l+1}), N_{\zeta+1,l+1}) \restriction \text{dom}(p^\zeta_l). \qquad (*)$$

Proof. We begin by constructing $(\bar{M}', \bar{a}, \bar{N})$, a continuous, amalgamable $<^c_{\mu,\alpha\times\mu\alpha}$-extension of $(\bar{M}, \bar{a}, \bar{N})$, such that for $\zeta+1 < \alpha$, $M'_{\zeta+1,0}$ is a (μ, μ)-limit over $\bigcup_{i<\mu\alpha} M'_{\zeta,i}$. The construction of $(\bar{M}', \bar{a}, \bar{N})$ is done by defining a directed system of amalgamable, continuous partial extensions of $(\bar{M}, \bar{a}, \bar{N})$ using Theorem II.7.1. Specifically, Theorem II.7.1 allows us to define by induction on ζ, a directed system $\langle(\bar{M}, \bar{a}, \bar{N})^\zeta \mid 1 \leq \zeta \leq \alpha\rangle$ and $\langle f_{\xi,\zeta} \mid 1 \leq \xi \leq \zeta \leq \alpha\rangle$ satisfying the following conditions:

(1) $(\bar{M}, \bar{a}, \bar{N}) \restriction (\zeta \times \mu\alpha) <^c_{\mu,\zeta\times\mu\alpha} (\bar{M}, \bar{a}, \bar{N})^\zeta$
(2) $(\bar{M}, \bar{a}, \bar{N})^\zeta$ is continuous and amalgamable
(3) $(\bar{M}, \bar{a}, \bar{N})^\zeta$ lies in $\mathfrak{C}$ for $\zeta < \alpha$
(4) $M^{\zeta+1}_{\zeta+1,0}$ is a (μ, μ)-limit over $\bigcup_{i<\mu\alpha} M^\zeta_{\zeta,i}$
(5) for all $\xi < \zeta$, M^ζ_ξ is universal over $f_{\xi,\zeta}\left(\bigcup_{i<\xi} M^\xi_i\right)$
(6) $f_{\xi,\zeta} \restriction M^\xi_i : M^\xi_i \to M^\zeta_i$ for $i < \xi \leq \zeta$ and
(7) $f_{\xi,\zeta} \restriction M_\xi = id_{M_\xi}$ for all $\xi < \zeta < \alpha$.

The details of the direct limit construction are similar to the direct limit construction in the limit case of Theorem II.7.1.

The construction is sufficient: Let $(\bar{M}', \bar{a}, \bar{N}) := (\bar{M}, \bar{a}, \bar{N})^\alpha$. For each $\zeta + 1 < \alpha$, fix a sequence $\langle M^*_{\zeta,i} \mid i < \mu\rangle$ witnessing that $M'_{\zeta+1,0}$ is a (μ, μ)-limit over $\bigcup_{i<\mu\alpha} M'_{\zeta,i}$. Define $M'_{\zeta,\mu\alpha+i} := M^*_{\zeta,i}$ for each $i < \mu$ and $\zeta + 1 < \alpha$.

$$\begin{array}{ccccccccc}
M_{0,0} \prec_{\mathcal{K}} & M_{0,i} & \prec_{\mathcal{K}} \bigcup_{i<\mu\alpha} M_{0,i} & & & \prec_{\mathcal{K}} & & & M_{1,0} \\
\Downarrow id & \Downarrow id & \Downarrow id & & & & & & \Downarrow id \\
M'_{0,0} \prec_{\mathcal{K}} & M'_{0,i} & \prec_{\mathcal{K}} \bigcup_{i<\mu\alpha} M'_{0,i} & = & M^*_{0,0} \overset{id}{\dashrightarrow} M^*_{0,j} & \overset{id}{\dashrightarrow} & M^*_{0,j+1} \overset{id}{\dashrightarrow} & \bigcup_{j<\mu} M^*_{0,j} & = M'_{1,0}
\end{array}$$

For each $\zeta + 1 < \alpha$ and each $l < \mu$, by the Theorem I.4.10, we can find $q \in \text{ga-S}(M'_{\zeta+1,\mu\alpha+l})$ extending p^ζ_l such that q does not μ-split over N^ζ_l. Since $M'_{\zeta+1,\mu\alpha+l+1}$ is universal over $M'_{\zeta+1,\mu\alpha+l}$, there is $a \in M'_{\zeta+1,\mu\alpha+l+1}$ realizing q. Set $a_{\zeta+1,\mu\alpha+l} = a$ and $N_{\zeta+1,\mu\alpha+l} = N^\zeta_l$. This gives us a definition of $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha\times\mu(\alpha+1)}$. To extend this tower to a tower with index set $(\alpha + 1) \times \mu(\alpha + 1)$, we use the fact that $(\bar{M}', \bar{a}, \bar{N})$ is amalgamable to fix M^* a $(\mu, \mu(\alpha + 1))$-limit model over $\bigcup_{i<\mu\alpha,\zeta<\alpha} M'_{\alpha,i}$. Let $\langle M'_{\alpha,i} \mid i < \mu(\alpha + 1)\rangle$ witness this. WLOG we may assume that $M'_{\alpha,i+1}$ is a (μ, ω)-limit over $M'_{\alpha,i}$ for each $i < \mu(\alpha + 1)$. For each $i < \mu(\alpha + 1)$, fix $a_{\alpha,i} \in M'_{\alpha,i+1}\backslash M'_{\alpha,i}$. By Fact I.4.7 and our choice of $M'_{\alpha,i}$ as a limit model, there is an $N \prec_{\mathcal{K}} M'_{\alpha,i}$ such that $M'_{\alpha,i}$ is universal over N and $\text{ga-tp}(a_{\alpha,i}/M'_{\alpha,i})$ does not μ-split over N. Set $N_{\alpha,i} = N$. Notice that $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,(\alpha+1)\times\mu(\alpha+1)}$ is as required. □

9. Uniqueness of limit models

Recall the running assumptions:

(1) $\mathcal{K}$ is an abstract elementary class,
(2) $\mathcal{K}$ has no maximal models,
(3) $\mathcal{K}$ is categorical in some $\lambda > LS(\mathcal{K})$,
(4) GCH and $\Phi_{\mu^+}(S^{\mu^+}_{\text{cf}(\mu)})$ holds for every cardinal $\mu < \lambda$.

Under these assumptions and Hypothesis 1, we can prove the uniqueness of limit models using the results from Sections 6 and 8.

Theorem II.9.1 (*Uniqueness of Limit Models*). *Let μ be a cardinal θ_1, θ_2 limit ordinals such that $\theta_1, \theta_2 < \mu^+ \leq \lambda$. Under Hypothesis 1, if M_1 and M_2 are (μ, θ_1) and (μ, θ_2) limit models over M, respectively, then there exists an isomorphism $f : M_1 \cong M_2$ such that $f \restriction M = id_M$.*

Proof. Let $M \in \mathcal{K}^{am}_\mu$ be given. By Fact I.2.11, it is enough to show that there exists a θ_2 such that for every θ_1 a limit ordinal $< \mu^+$, we have that a (μ, θ_1)-limit model over M is isomorphic to a (μ, θ_2)-limit model over M. Take θ_2 such that $\theta_2 = \mu\theta_2$. Fix θ_1 a limit ordinal $< \mu^+$. By Fact I.2.12, we may assume that θ_1 is regular. Using Fact I.2.11 again, it is enough to construct a model M^* which is simultaneously a (μ, θ_1)-limit model over M and a (μ, θ_2)-limit model over M.

The idea is to build a (scattered) array of models such that at some point in the array, we will find a model which is a (μ, θ_1)-limit model witnessed by its height in the array and is a (μ, θ_2)-limit model witnessed by its horizontal position in the array, relative fullness and continuity. We will define a chain of length μ^+ of continuous towers while increasing the index set of the towers in order to realize strong types as we proceed with the goal of producing many relatively full rows.

Define by induction on $0 < \alpha < \mu^+$ the $<^c$-increasing sequence of towers, $\langle(\bar{M}, \bar{a}, \bar{N})^\alpha \in {}^+\mathcal{K}^*_{\mu,\alpha\times\mu\alpha} \mid \alpha < \mu^+\rangle$, such that

(1) $M \prec_{\mathcal{K}} M^\alpha_{0,0}$,
(2) $(\bar{M}, \bar{a}, \bar{N})^\alpha$ is continuous and amalgamable,
(3) $(\bar{M}, \bar{a}, \bar{N})^\alpha := \bigcup_{\beta<\alpha}(\bar{M}, \bar{a}, \bar{N})^\beta$ for α a limit ordinal and
(4) In successor stages in new intervals of length μ, put in representatives of all $\mathfrak{St}$-types from the previous stages. More formally, if $(p, N) \in \mathfrak{St}(M^\alpha_{\beta,i})$ for $i < \mu\alpha$ and $\beta < \alpha$, there exists $j \in [\mu\alpha, \mu(\alpha + 1)]$ such that

$$(p, N) \sim (\text{ga-tp}(a_{\beta+1,j}/M^{\alpha+1}_{\beta+1,j}), N_j) \restriction M^\alpha_{\beta,i}.$$

This construction is possible:
$\alpha = 1$: We can choose $\bar{M}^* = \langle M^*_i \mid i < \mu\rangle$ to be an $\prec_{\mathcal{K}}$ increasing continuous sequence of limit models of cardinality μ with $M^*_0 = M$ and M^*_{i+1} universal over M^*_i. For each $i < \mu$, fix $a^1_{0,i} \in M^*_{i+1}\backslash M^*_i$. Now consider ga-tp$(a^1_{0,i}/M^*_i)$. Since M^*_i is a limit model, we can apply Fact I.4.7 to fix $N^1_{0,i} \in \mathcal{K}^{am}_\mu$ such that ga-tp$(a^1_{0,i}/M^*_i)$ does not μ-split over $N^1_{0,i}$ and M^*_i is universal over $N^1_{0,i}$. Let $\bar{a}^1 := \langle a^1_{0,i} \mid i < \mu\rangle$ and $\bar{N}^1 = \langle N^1_{0,i} \mid i < \mu\rangle$.
α a limit ordinal: Take $(\bar{M}, \bar{a}, \bar{N})^\alpha := \bigcup_{\beta<\alpha}(\bar{M}, \bar{a}, \bar{N})^\beta$. Clearly $(\bar{M}, \bar{a}, \bar{N})^\alpha$ is continuous. To see that $(\bar{M}, \bar{a}, \bar{N})^\alpha$ is also amalgamable, we notice that $\bigcup_{\beta,i\in\alpha\times\mu\alpha} M^\alpha_{(\beta,i)}$ is a (μ, α)-limit model witnessed by $\langle\bigcup_{i<\mu\beta} M^\beta_{\beta,i} \mid \beta < \alpha\rangle$.
$\alpha = \beta + 1$: Suppose that $(\bar{M}, \bar{a}, \bar{N})^\beta$ has been defined. By Fact II.6.4, for every $\gamma < \beta$, we can enumerate $\bigcup_{k<\mu\beta} \mathfrak{St}(M^\beta_{\gamma,k})$ as $\{(p, N)^\gamma_l \mid l < \mu\}$. By Theorem II.8.2, we can find a continuous, amalgamable extension $(\bar{M}, \bar{a}, \bar{N})^{\beta+1} \in {}^+\mathcal{K}^*_{\mu,\beta+1\times\mu(\beta+1)}$ of $(\bar{M}, \bar{a}, \bar{N})^\beta$ such that for every $l < \mu$ and $\gamma < \beta$,

$$(p, N)^\gamma_l \sim (\text{ga-tp}(a_{\gamma+1,\mu\beta+l}/M^{\beta+1}_{\gamma+1,\mu\beta+l}), N_{\gamma+1,\mu\beta+l}) \restriction \text{dom}(p^\gamma_l).$$

This completes the construction.

We now want to identify all the rows of the array which are relatively full.

Claim II.9.2. *For δ a limit ordinal $< \mu^+$, we have that $(\bar{M}, \bar{a}, \bar{N})^\delta$ is full relative to $\langle \bar{M}^\delta_{\beta,i} \mid (\beta, i) \in \delta \times \mu\delta\rangle$ where $\bar{M}^\delta_{\beta,i} := \langle M^\gamma_{\beta,i} \mid \gamma < \delta\rangle$.*

Proof. Let $(p, N) \in \mathfrak{St}(M^\delta_{\beta,i})$ be given such that $N = M^\gamma_{\beta,i}$ for some $\gamma < \delta$, $\beta < \delta$ and $i < \mu\delta$. Since our construction is increasing and continuous, there exists $\delta' < \delta$ such that $(\beta, i) \in \delta' \times \mu\delta'$ and $\gamma < \delta'$. Notice then that $M^{\delta'}_{\beta,i}$ is universal over N. Furthermore, $p \restriction M^{\delta'}_{\beta,i}$ does not μ-split over N. Thus $(p, N) \restriction M^{\delta'}_{\beta,i} \in \mathfrak{St}(M^{\delta'}_{\beta,i})$. By condition (4) of the construction, there exists $j < \mu(\delta' + 1)$, such that

$$(p, N) \restriction M^{\delta'}_{\beta,i} \sim (\text{ga-tp}(a_{\beta+1,j}/M^{\beta+1}_{\beta+1,j}), N_{\beta+1,j}) \restriction M^{\delta'}_{\beta,i}.$$

Since $M^{\beta+1}_{\beta+1,j} \prec_{\mathcal{K}} M^\delta_{\beta+1,j}$ and ga-tp$(a_{\beta+1,j}/M^\delta_{\beta+1,j})$ does not μ-split over $N_{\beta+1,j}$, we can replace $M^{\beta+1}_{\beta+1,j}$ with $M^\delta_{\beta+1,j}$:

$$(p, N) \restriction M^{\delta'}_{\beta,i} \sim (\text{ga-tp}(a_{\beta+1,j}/M^\delta_{\beta+1,j}), N_{\beta+1,j}) \restriction M^{\delta'}_{\beta,i}.$$

Let M' be a universal extension of $M^\delta_{\beta+1,j}$. By definition of $\sim$, there exists $q \in$ ga-S(M') such that q extends $p \restriction M^{\delta'}_{\beta,i} =$ ga-tp$(a_{\beta+1,j}/M^{\delta'}_{\beta,i})$ and q does not μ-split over N and $N_{\beta+1,j}$. By the uniqueness of non-splitting extensions (Theorem I.4.12), since p does not μ split over N, we have that $q \restriction M^\delta_{\beta,i} = p$. Also, since ga-tp$(a_{\beta+1,j}/M^\delta_{\beta+1,j})$ does not μ-split over $N_{\beta+1,j}$, Theorem I.4.12 gives us $q \restriction M^\delta_{\beta+1,j} =$ ga-tp$(a_{\beta+1,j}/M^\delta_{\beta+1,j})$. By definition of $\sim$ and Lemma II.6.3, q also witnesses that

$$(\text{ga-tp}(a_{\beta+1,j}/M^\delta_{\beta+1,j}), N_{\beta+1,j}) \restriction M^\delta_{\beta,i} \sim (p, N).$$

Since (p, N) was chosen arbitrarily, we have verified that $(\bar{M}, \bar{a}, \bar{N})^\delta$ satisfies the definition of relative fullness. □

Take $\langle\delta_\zeta < \mu^+ \mid \zeta \leq \theta_1\rangle$ to be an increasing and continuous sequence of limit ordinals $> \theta_2$. We will consider the restrictions (in the sense of Notation II.9.3) of $(\bar{M}, \bar{a}, \bar{N})^{\delta_\zeta}$ to $\theta_2 \times \mu\delta_\zeta$:

Notation II.9.3. For θ and δ ordinals $< \mu^+$ and a sequence $\bar{M}$ indexed by a superset of $\theta \times \mu\delta$, we will abbreviate $\langle M_{\beta,i} \mid \beta < \theta \text{ and } i < \mu\delta\}$ by $\bar{M} \restriction^{\theta\times\mu\delta}$.

Define

$$M^* := \bigcup_{\zeta<\theta_1} \bigcup_{i\in\theta_2\times\mu\delta_\zeta} M^{\delta_\zeta}_i = \bigcup_{i\in\theta_2\times\mu\delta_{\theta_1}} M^{\delta_{\theta_1}}_i.$$

We will now verify that M^* is a (μ, θ_1)-limit over M and a (μ, θ_2)-limit over M.

Notice that $\langle \bigcup_{i\in\theta_2\times\mu\delta_\zeta} M_i^{\delta_\zeta} \mid \zeta < \theta_1\rangle$ witnesses that M^* is a (μ,θ_1) limit. Since $M \prec_{\mathcal{K}} M_{0,0}^{\delta_0}$, M^* is a (μ,θ_1)-limit over M.

By Claim II.9.2 and the fact that the restriction of a relatively full tower is relatively full (Proposition II.6.9), we have that

$$(\bar{M},\bar{a},\bar{N})^{\delta_{\theta_1}} \restriction^{\theta_2\times\mu\delta_\zeta} \text{ is full relative to } \langle \bar{M}_{\beta,i}^{\delta_{\theta_1}} \mid (\beta,i)\in\theta_2\times\mu\delta_{\theta_1}\rangle,$$

where $\bar{M}_{\beta,i}^{\delta_{\theta_1}} := \langle M_{\beta,i}^{\gamma} \mid \gamma < \delta_{\theta_1}\rangle$. Furthermore, we see that $(\bar{M},\bar{a},\bar{N})^{\delta_{\theta_1}} \restriction^{\theta_2\times\mu\delta_{\theta_1}}$ is continuous. Since $\theta_2 = \mu\cdot\theta_2$, we can apply Theorem II.6.10 to conclude that M^* is a (μ,θ_2)-limit model over M. □

Remark II.9.4. The above proof implicitly shows the decomposition of a relatively full tower into a resolution of θ' many towers for every limit $\theta' < \mu^+$.

Part III. Conclusion

We provide a partial proof of Hypothesis 1. We also discuss reduced towers, which appear in the [33] and may be useful as a tool to prove the amalgamation property for categorical AECs with no maximal models. We will continue to make Assumption 0.7.

10. $<^c_{\mu,\alpha}$-Extension property for nice towers

In [33], Shelah and Villaveces claim that every tower in ${}^+\mathcal{K}^*_{\mu,\alpha}$ has a proper $<^c_{\mu,\alpha}$ extension. This proof does not converge. Here we prove a weaker extension property. Namely, we show that every *nice* tower in ${}^+\mathcal{K}^*_{\mu,\alpha}$ has a proper $<^c_{\mu,\alpha}$-extension (Corollary III.10.6). This is a proof of an approximation to the statement of Hypothesis 1 which states that every continuous tower has an amalgamable extension inside $\mathfrak{C}$.

Theorem III.10.1. *Let μ be a cardinal and α,γ ordinals such that $\gamma < \alpha < \mu^+ \le \lambda$. If $(\bar{M},\bar{a},\bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ is nice and $(\bar{M}'',\bar{a},\bar{N}) \restriction \gamma$ is an amalgamable partial extension of $(\bar{M},\bar{a},\bar{N})$, then there exists an amalgamable $(\bar{M}^*,\bar{a},\bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ and a $\prec_{\mathcal{K}}$-mapping f such that*

(1) $(\bar{M},\bar{a},\bar{N}) <^c_{\mu,\alpha} (\bar{M}^*,\bar{a},\bar{N})$
(2) $f(M_i'') = M_i^*$ *for all* $i < \gamma$ *and*
(3) $f \restriction M_i = id_{M_i}$ *for all* $i < \gamma$.

Furthermore if $\bigcup_{i<\alpha} M_i \prec_{\mathcal{K}} \mathfrak{C}$ and $\bar{b} \in {}^{\le\mu}\mathfrak{C}$ is such that $\bar{b}\cap\bigcup_{i<\alpha} M_i = \emptyset$, then we can find $(\bar{M}^*,\bar{a},\bar{N})$ as above with $\bar{b}\cap\bigcup_{i<\alpha} M_i^* = \emptyset$.

Remark III.10.2. If $(\bar{M},\bar{a},\bar{N})$ is amalgamable and $\bigcup_{i<\alpha} M_i \prec_{\mathcal{K}} \mathfrak{C}$, then we can find an extension $(\bar{M}',\bar{a},\bar{N})$ such that $\bigcup_{i<\alpha} M_i' \prec_{\mathcal{K}} \mathfrak{C}$.

Theorem III.10.1 is stronger than the $<^c_{\mu,\alpha}$-extension property since it allows us to avoid μ-many elements ($\bar{b}$). This is possible due to Weak Disjoint Amalgamation, Fact I.3.15.

Proof of Theorem III.10.1. Let an amalgamable $(\bar{M},\bar{a},\bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ be given.

As in the proofs of Theorems II.7.1 and II.8.2, we will define by induction on $i < \alpha$ a direct system of models $\langle M_i' \mid i < \alpha\rangle$ and $\prec_{\mathcal{K}}$-mappings, $\langle f_{j,i} \mid j < i < \alpha\rangle$ such that for $i \le \alpha$:

(1) $(\langle f_{j,i}(M_j') \mid j \le i\rangle, \bar{a} \restriction i+1, \bar{N} \restriction i+1)$ is a $<^c_{\mu,i+1}$-extension of $(\bar{M},\bar{a},\bar{N}) \restriction (i+1)$,
(2) $(\langle M_j' \mid j < i\rangle, \langle f_{j,i} \mid j \le i\rangle)$ forms a directed system,
(3) M_i' is universal over M_i,
(4) M_{i+1}' is universal over $f_{i,i+1}(M_i')$,
(5) $f_{j,i} \restriction M_j = id_{M_j}$.

Notice that the M_i''s will not necessarily form an extension of the tower $(\bar{M}, \bar{a}, \bar{N})$. Rather, for each $i < \alpha$, we find some image of $\langle M_j' \mid j < i \rangle$ which will extend the initial segment of length i of $(\bar{M}, \bar{a}, \bar{N})$ (see condition (1) of the construction).

The construction is possible:

$i = 0$: Since M_0 is an amalgamation base, we can find $M_0'' \in \mathcal{K}_\mu^*$ (a first approximation of the desired M_0') such that M_0'' is universal over M_0. By Theorem I.4.10, we may assume that ga-tp(a_0/M_0'') does not μ-split over N_0 and $M_0'' \prec_{\mathcal{K}} \mathfrak{C}$. Since $a_0 \notin M_0$ and ga-tp(a_0/M_0) does not μ-split over N_0, we know that $a_0 \notin M_0''$. But, we might have that for some $l > 0$, $a_l \in M_0''$ or $\bar{b} \cap M_0'' \neq \emptyset$. We use Weak Disjoint Amalgamation to avoid $\{a_l \mid 0 < l < \alpha\}$ and $\bar{b}$. By the Downward Löwenheim–Skolem Axiom for AECs (Axiom 4) we can choose $M^2 \in \mathcal{K}_\mu$ such that M_0'', $M_1 \prec_{\mathcal{K}} M^2 \prec_{\mathcal{K}} \mathfrak{C}$.

By Corollary I.3.16 (applied to M_1, M_α, M^2 and $\langle a_l \mid 0 < l < \alpha \rangle \cup \bar{b}$), we can find a $\prec_{\mathcal{K}}$-mapping h such that

- $h : M^2 \to \mathfrak{C}$
- $h \restriction M_1 = id_{M_1}$
- $h(M^2) \cap (\{a_l \mid 0 < l < \alpha\} \cup \bar{b}) = \emptyset$.

Define $M_0' := h(M_0'')$. Notice that $a_0 \notin M_0'$ because $a_0 \notin M_0''$ and $h(a_0) = a_0$. Clearly $M_0' \cap (\{a_l \mid 0 \leq l < \alpha\} \cup \bar{b}) = \emptyset$, since $M_0'' \prec_{\mathcal{K}} M^2$ and $h(M^2) \cap \{a_l \mid 0 < l < \alpha\} = \emptyset$. We need only verify that ga-tp(a_0/M_0') does not μ-split over N_0. By invariance, ga-tp(a_0/M_0'') does not μ-split over N_0 implies that ga-tp$(h(a_0)/h(M_0''))$ does not μ-split over N_0. But recall $h(a_0) = a_0$ and $h(M_0'') = M_0'$. Thus ga-tp(a_0/M_0') does not μ-split over N_0.

Set $f_{0,0} := id_{M_0'}$.

Below is a diagram of the successor stage of the construction.

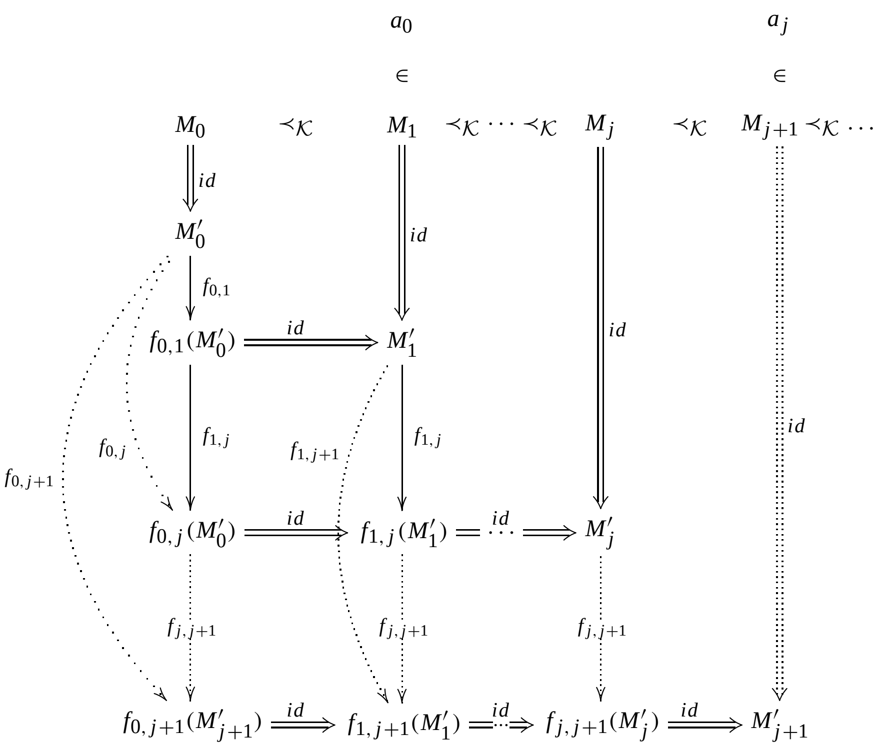

$i = j + 1$: Suppose that we have completed the construction for all $k \leq j$. Since M_j' and M_{j+1} are both $\mathcal{K}$-substructures of $\mathfrak{C}$, we can apply the Downward-Löwenheim Axiom for AECs to find M_{j+1}''' (a first approximation to M_{j+1}') a model of cardinality μ extending both M_j' and M_{j+1}. WLOG by Theorem I.2.17 and Lemma I.2.24 we may assume that M_{j+1}''' is a limit model of cardinality μ and M_{j+1}''' is universal over M_{j+1} and M_j'. By Theorem I.4.10, we can find a $\prec_{\mathcal{K}}$-mapping $f : M'''_{j+1} \to \mathfrak{C}$ such that $f \restriction M_{j+1} = id_{M_{j+1}}$ and ga-tp$(a_{j+1}/f(M'''_{j+1}))$ does not μ-split over N_{j+1}. Set $M''_{j+1} := f(M'''_{j+1})$.

Subclaim III.10.3. $a_{j+1} \notin M''_{j+1}$.

Proof of Subclaim III.10.3. Suppose that $a_{j+1} \in M''_{j+1}$. Since M_{j+1} is universal over N_{j+1}, there exists a $\prec_{\mathcal{K}}$-mapping, $g : M''_{j+1} \to M_{j+1}$ such that $g \restriction N_{j+1} = id_{N_{j+1}}$. Since ga-tp$(a_{j+1}/M''_{j+1})$ does not μ-split over N_{j+1}, we have that

$$\text{ga-tp}(a_{j+1}/g(M''_{j+1}) = \text{ga-tp}(g(a_{j+1})/g(M''_{j+1})).$$

Notice that because $g(a_{j+1}) \in g(M''_{j+1})$, we have that $a_{j+1} = g(a_{j+1})$. Thus $a_{j+1} \in g(M''_{j+1}) \prec_{\mathcal{K}} M_{j+1}$. This contradicts the definition of towers: $a_{j+1} \notin M_{j+1}$. □

M''_{j+1} may serve us well if it does not contain any a_l for $j+1 \le l < \alpha$ or any part of $\bar{b}$, but this is not guaranteed. So we need to make an adjustment. Let M^2 be a model of cardinality μ such that $M_{j+2}, M''_{j+1} \prec_{\mathcal{K}} M^2 \prec_{\mathcal{K}} \mathfrak{C}$. Notice that $\mathfrak{C}$ is universal over M_{j+2}. Thus we can apply Corollary I.3.16 to M_{j+2}, M_α, M^2 and $\langle a_l \mid j+2 \le l < \alpha\rangle \cup \bar{b}$. This yields a $\prec_{\mathcal{K}}$-mapping h such that

- $h : M^2 \to \mathfrak{C}$
- $h \restriction M_{j+2} = id_{M_{j+2}}$ and
- $h(M^2) \cap (\{a_l \mid j+2 \le l < \alpha\} \cup \bar{b}) = \emptyset$.

Set $M'_{j+1} := h(M''_{j+1})$. Notice that by invariance, ga-tp(a_{j+1}/M''_{j+1}) does not μ-split over N_{j+1} implies that ga-tp$(h(a_{j+1})/h(M''_{j+1}))$ does not μ-split over $h(N_{j+1})$. Recalling that $h \restriction M_{j+2} = id_{M_{j+2}}$ we have that ga-tp(a_{j+1}/M''_{j+1}) does not μ-split over N_{j+1}. We need to verify that $a_{j+1} \notin M'_{j+1}$. This holds because $a_{j+1} \notin M''_{j+1}$ and $h(a_{j+1}) = a_{j+1}$.

Set $f_{j+1,j+1} = id_{M_{j+1}}$ and $f_{j,j+1} := h \circ f \restriction M'_j$. To guarantee that we have a directed system, for $k < j$, define $f_{k,j+1} := f_{j,j+1} \circ f_{k,j}$.

i is a limit ordinal: Suppose that $(\langle M'_j \mid j < i\rangle, \langle f_{k,j} \mid k \le j < i\rangle)$ have been defined. Since it is a directed system, we can take direct limits.

Subclaim III.10.4. *We can choose a direct limit* $(M^*_i, \langle f^*_{j,i} \mid j \le i\rangle)$ *of* $(\langle M'_j \mid j < i\rangle, \langle f_{k,j} \mid k \le j < i\rangle)$ *such that*

(1) $M^*_i \prec_{\mathcal{K}} \mathfrak{C}$
(2) $f^*_{j,i} \restriction M_j = id_{M_j}$ *for every* $j < i$.

Proof of Subclaim III.10.4. This follows from Subclaim II.7.10 and the assumption that $(\bar{M}, \bar{a}, \bar{N})$ is nice. □

By Condition (4) of the construction, notice that M^*_i is a (μ, i)-limit model witnessed by $\langle f^*_{j,i}(M'_j) \mid j < i\rangle$. Hence M^*_i is an amalgamation base. Since M^*_i and M_i both live inside of $\mathfrak{C}$, we can find $M'''_i \in \mathcal{K}^*_\mu$ which is universal over M_i and universal over M^*_i.

By Theorem I.4.10 we can find a $\prec_{\mathcal{K}}$-mapping $f : M'''_i \to \mathfrak{C}$ such that $f \restriction M_i = id_{M_i}$ and ga-tp$(a_i/f(M'''_i))$ does not μ-split over N_i. Set $M''_i := f(M'''_i)$. By a similar argument to Subclaim III.10.3, we can see that $a_i \notin M''_i$.

M''_i may contain some a_l when $i \le l < \alpha$ or part of $\bar{b}$. We need to make an adjustment using Weak Disjoint Amalgamation. Let M^2 be a model of cardinality μ such that $M''_i, M_{i+1} \prec_{\mathcal{K}} M^2 \prec_{\mathcal{K}} \mathfrak{C}$. By Corollary I.3.16 applied to M_i, M_α, M^2 and $\langle a_l \mid i < l < \alpha\rangle \cup \bar{b}$ we can find $h : M''_i \to \mathfrak{C}$ such that $h \restriction M_{i+1} = id_{M_{i+1}}$ and $h(M^2) \cap (\{a_l \mid i < l < \alpha\} \cup \bar{b}) = \emptyset$.

Set $M'_i := h(M''_i)$. We need to verify that $a_i \notin M'_i$ and ga-tp(a_i/M'_i) does not μ-split over N_i. Since $a_i \notin M''_i$ and $h(a_i) = a_i$, we have that $a_i \notin h(M''_i) = M'_i$. By invariance of non-splitting, ga-tp(a_i/M''_i) not μ-splitting over N_i implies that ga-tp$(h(a_i)/h(M''_i))$ does not μ-split over $h(N_i)$. Recalling our definition of h and M'_i, this yields ga-tp(a_i/M'_i) does not μ-split over N_i.

As in the proof of Theorem II.7.1, we see that $(\langle (f_{j,i}(M_j) \mid j \le i\rangle, \bar{a} \restriction i, \bar{N} \restriction i)$ is a $<^c_{\mu,i}$-extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction i$.

Set $f_{i,i} := id_{M_i,i}$, and for $j < i$, $f_{j,i} := h \circ f \circ f^*_{j,i}$. This completes the construction.

The construction is enough: We have constructed a directed system $\langle \bigcup_{i<\gamma} M'_i \mid i < \alpha \rangle$ with $\langle f_{i,j} \mid i \le j < \alpha \rangle$. By Fact II.7.7 and Subclaim II.7.10 we can find a direct limit to this system, M^*_α and $\prec_{\mathcal{K}}$-mappings $\langle f_{i,\alpha} \mid i \le \alpha \rangle$ such that $f_{i,\alpha} \restriction M_i = id_{M_i}$ for all $i < \alpha$ and M^*_α avoids $\bar{b}$. If $(\bar{M}, \bar{a}, \bar{N})$ is amalgamable, then M^*_α can be chosen to lie in $\mathfrak{C}$. Define for all $j < \alpha$, $M^*_j := f_{j+1,\alpha}(M'_j)$. Notice that as in Subclaim II.7.11, $(\bar{M}, \bar{a}, \bar{N}) <^c_{\mu,\alpha} (\bar{M}^*, \bar{a}, \bar{N})$. And, as in the limit stage of the construction, we see that $(\bar{M}^*, \bar{a}, \bar{N})$ is continuous and amalgamable. □

Remark III.10.5. Notice that in Theorem III.10.1 if the partial extension $(\bar{M}', \bar{a}, \bar{N})$ is continuous, then we can choose $\bar{M}''$ such that it is continuous below γ, that is for every $i < \gamma$ with i a limit ordinal, $M''_i = \bigcup_{j<i} M''_j$.

Corollary III.10.6 (*The $<^c_{\mu,\alpha}$-Extension Property for Nice Towers*). *If $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ is nice, then there exists an amalgamable $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ such that $(\bar{M}, \bar{a}, \bar{N}) <^c_{\mu,\alpha} (\bar{M}', \bar{a}, \bar{N})$.*

Proof. Take $\gamma = 0$ in Theorem III.10.1 □

Remark III.10.7. Notice that Hypothesis 3 implies that every tower is amalgamable. Thus Hypothesis 3 together with Corollary III.10.6 imply the $<^c_{\mu,\alpha}$-extension property for all towers.

11. Reduced towers

Shelah and Villaveces introduce the notion of reduced towers in order to show the density of continuous towers. While there are difficulties with Shelah and Villaveces' approach, we discuss reduced towers because they have characteristics similar to strongly minimal types in first-order model theory. Additionally, they generalize reduced triples used in [31] to develop a notion of non-forking.

Definition III.11.1. A tower $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ is said to be *reduced* provided that for every $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ with $(\bar{M}, \bar{a}, \bar{N}) \le^c_{\mu,\alpha} (\bar{M}', \bar{a}, \bar{N})$ we have that for every $i < \alpha$,

$$M'_i \cap \bigcup_{j<\alpha} M_j = M_i. \qquad (*)_i$$

If we take a $<^c$-increasing chain of reduced towers, the union will be reduced. The following proposition appears in [33] (Theorem 3.1.14 of [33]) for reduced towers. We provide the proof for completeness.

Theorem III.11.2. *If $\langle (\bar{M}, \bar{a}, \bar{N})^\gamma \in {}^+\mathcal{K}^*_{\mu,\alpha} \mid \gamma < \beta \rangle$ is a $<^c_{\mu,\alpha}$-increasing and continuous sequence of reduced towers, then the union of this sequence of towers is a reduced tower.*

Proof. Denote by $(\bar{M}, \bar{a}, \bar{N})^\beta$ the union of the sequence of towers. That is $\bar{a}^\beta = \bar{a}^0$, $\bar{N}^\beta = \bar{N}^0$ and $\bar{M}^\beta = \langle M^\beta_i \mid i < \alpha \rangle$ where $M^\beta_i = \bigcup_{\gamma<\beta} M^\gamma_i$.

Suppose that $(\bar{M}, \bar{a}, \bar{N})^\beta$ is not reduced. Let $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ witness this. Then there exists an $i < \alpha$ and an element b such that $b \in \left(M'_i \cap \bigcup_{j<\alpha} M^\beta_j\right) \setminus M^\beta_i$. There exists $\gamma < \beta$ such that $b \in \bigcup_{j<\alpha} M^\gamma_j \setminus M^\gamma_i$. Notice that $(\bar{M}', \bar{a}, \bar{N})$ witnesses that $(\bar{M}, \bar{a}, \bar{N})^\gamma$ is not reduced. □

The following appears in [33] (Theorem 3.1.13).

Fact III.11.3 (*Density of Reduced Towers*). *There exists a reduced $<^c_{\mu,\alpha}$-extension of every nice tower in ${}^+\mathcal{K}^*_{\mu,\alpha}$.*

Proof. Suppose for the sake of contradiction that no $<^c_{\mu,\alpha}$-extension of the tower $(\bar{M}, \bar{a}, \bar{N})$ is reduced. This allows us to construct a $\le^c_{\mu,\alpha}$-increasing and continuous sequence of towers $\langle (\bar{M}, \bar{a}, \bar{N})^\zeta \in {}^+\mathcal{K}^*_{\mu,\alpha} \mid \zeta < \mu^+ \rangle$ such that $(\bar{M}, \bar{a}, \bar{N})^{\zeta+1}$ witnesses that $(\bar{M}, \bar{a}, \bar{N})^\zeta$ is not reduced.

The construction: Since $(\bar{M}, \bar{a}, \bar{N})$ is nice, we can apply Corollary III.10.6 to find $(\bar{M}, \bar{a}, \bar{N})^0$ a $<^c_{\mu,\alpha}$ extension of $(\bar{M}, \bar{a}, \bar{N})$.

Suppose that $(\bar{M}, \bar{a}, \bar{N})^\zeta$ has been defined. Since it is a $<^c_{\mu,\alpha}$-extension of $(\bar{M}, \bar{a}, \bar{N})$, we know it is not reduced. Let $(\bar{M}, \bar{a}, \bar{N})^{\zeta+1} \in {}^+\mathcal{K}^*_{\mu,\alpha}$ be a $\leq^c_{\mu,\alpha}$-extension of $(\bar{M}, \bar{a}, \bar{N})^\zeta$, witnessing this.

For ζ a limit ordinal, let $(\bar{M}, \bar{a}, \bar{N})^\zeta = \bigcup_{\gamma<\zeta} (\bar{M}, \bar{a}, \bar{N})^\gamma$. This completes the construction.

For each $b \in \bigcup_{\zeta<\mu^+, i<\alpha} M^\zeta_i$ define

$$i(b) := \min \left\{ i < \alpha \;\middle|\; b \in \bigcup_{\zeta<\mu^+} \bigcup_{j<i} M^\zeta_j \right\} \text{ and}$$

$$\zeta(b) := \min \left\{ \zeta < \mu^+ \mid b \in M^\zeta_{i(b)} \right\}.$$

$\zeta(\cdot)$ can be viewed as a function from μ^+ to μ^+. Thus there exists a club $E = \left\{\delta < \mu^+ \mid \forall b \in \bigcup_{i<\alpha} M^\delta_i,\ \zeta(b) < \delta\right\}$. Actually, all we need is for E to be non-empty.

Fix $\delta \in E$. By construction $(\bar{M}, \bar{a}, \bar{N})^{\delta+1}$ witnesses the fact that $(\bar{M}, \bar{a}, \bar{N})^\delta$ is not reduced. So we may fix $i < \alpha$ and $b \in M^{\delta+1}_i \cap \bigcup_{j<\alpha} M^\delta_j$ such that $b \notin M^\delta_i$. Since $b \in M^{\delta+1}_i$, we have that $i(b) \leq i$. Since $\delta \in E$, we know that there exists $\zeta < \delta$ such that $b \in M^\zeta_{i(b)}$. Because $\zeta < \delta$ and $i(b) < i$, we have that $b \in M^\delta_i$ as well. This contradicts our choice of i and b witnessing the failure of $(\bar{M}, \bar{a}, \bar{N})^\delta$ to be reduced. □

A variation of the following theorem was claimed in [33] for reduced towers. Unfortunately, their proof does not converge. Under Hypothesis 3, we resolve their problems here.

Theorem III.11.4 (*Reduced Towers are Continuous*). *Under Hypothesis* 3, *if* $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ *is reduced, then it is continuous.*

The keys to resolving problems of [33] are the extra conditions in the main construction and the following lemma which is a consequence of Theorem III.10.1 and the definition of reduced tower.

Lemma III.11.5. *Suppose that* $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha}$ *is reduced and nice, then for every* $\beta < \alpha$, $(\bar{M}, \bar{a}, \bar{N}) \restriction \beta$ *is reduced.*

Notice that without the full $<^c_{\mu,\alpha}$-extension property, it is conceivable to have a discontinuous reduced tower with non-reduced restrictions.

Proof of Theorem III.11.4. Suppose the claim fails for μ and δ is the minimal limit ordinal for which it fails. More precisely, δ is the minimal element of

$$S = \left\{ \delta < \mu^+ \;\middle|\; \begin{array}{l} \delta \text{ is a limit ordinal such that there exists} \\ \text{an } \alpha < \mu^+ \text{ and} \\ \text{a nice, reduced tower } (\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\alpha} \\ \text{with } M_\delta \supsetneq_{\mathcal{K}} \bigcup_{i<\delta} M_i \end{array} \right\}.$$

Let α witness that $\delta \in S$. Hypothesis 3 implies that every tower is amalgamable. Thus we can apply Lemma III.11.5, to assume that $\alpha = \delta+1$. Fix $(\bar{M}, \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\delta+1}$ witnessing that $\delta \in S$. Let $b \in M_\delta \setminus \bigcup_{i<\delta} M_i$ be given. By Fact III.11.3, Hypothesis 3 and the minimality of δ, every nice tower of length δ has a continuous extension. Combining this with the fact that $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$ is amalgamable, we can apply Proposition II.7.2 to $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$ and b to find a $<^c_{\mu,\delta}$-extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$, say $(\bar{M}', \bar{a} \restriction \delta, \bar{N} \restriction \delta) \in {}^+\mathcal{K}^*_{\mu,\delta}$, in $\mathfrak{C}$ containing b. Let $M'_\delta \prec_{\mathcal{K}} \mathfrak{C}$ be a limit model universal over M_δ containing $\bigcup_{i<\delta} M'_i$. Notice that $(\bar{M}', \bar{a}, \bar{N}) \in {}^+\mathcal{K}^*_{\mu,\delta+1}$ is an extension of $(\bar{M}, \bar{a}, \bar{N})$ witnessing that $(\bar{M}, \bar{a}, \bar{N})$ is not reduced. □

Positive solutions to the following questions would allow us to adjust the previous proof to conclude that every nice tower has a continuous extension without any extra hypothesis.

Question III.11.6. Is it possible to remove Hypothesis 3 in the proof of Theorem III.11.4? Alternatively, can one show the density of nice, reduced towers?

The next step towards Shelah's Categoricity Conjecture is to show that the uniqueness of limit models implies the amalgamation property in this context.

Acknowledgements

We thank R. Grossberg for his kind and generous guidance on Chapter II of the Ph.D. thesis [35] on which this paper is based. We are indebted to J. Baldwin for his unending dedication to improving the clarity and correctness of this paper and for chairing the thesis examination committee for [35]. A. Villaveces' comments and discussions on this paper and on [35] were invaluable. Additionally, we express our gratitude to the referee whose patience and comments contributed to a much needed revision of the original work.

References

[1] J. Baldwin, Ehrenfeucht–Mostowski Models in Abstract Elementary Classes, Preprint available at http://www2.math.uic.edu/~jbaldwin/model.html.
[2] J. Baldwin, D. Kueker, M. VanDieren, Upward stability transfer (8 pages), Notre Dame J. Formal Logic (in press). Preprint available at: http://www.math.lsa.umich.edu/~mvd/home.html.
[3] J. Baldwin, S. Shelah, Abstract classes with few models have 'Homogeneous-Universal' models, JSL 60 (1995).
[4] K.J. Devlin, S. Shelah, A weak version of $\diamond$ which follows from $2^{\aleph_0} < 2^{\aleph_1}$, Israel J. Math. 29 (1978) 239–247.
[5] J. Gregory, Higher souslin trees and the generalized continuum hypothesis, JSL 41 (1976) 663–671.
[6] R. Grossberg, Classification theory for non-elementary classes, in: Yi Zhang (Ed.), Logic and Algebra, in: Contemporary Mathematics, vol. 302, AMS, 2002, pp. 165–204.
[7] R. Grossberg, A Course in Model Theory, Available at http://www.math.cmu.edu/~rami.
[8] R. Grossberg, M. VanDieren, Stability in Tame Abstract Elementary Classes (23 pages), J. Symbolic Logic (in press). Preprint available at: www.math.lsa.umich.edu/~mvd/home.html.
[9] R. Grossberg, M. VanDieren, A. Villaveces, Limit models in classes with amalgamation (18 pages). Preprint available at: www.math.lsa.umich.edu/~mvd/home.html (submitted for publication).
[10] P. Hall, Some constructions for locally finite groups, J. London Math Soc. 34 (1959) 305–319.
[11] B. Hart, S. Shelah, Categoricity over P for first order T or categoricity for $\varphi \in \mathrm{L}_{\omega_1\omega}$ can stop at $\aleph_k$ while holding for $\aleph_0, \cdots, \aleph_{k-1}$, Israel J. Math. 70 (1990) 219–235.
[12] B. Jónsson, Universal relational systems, Math. Scand. 4 (1956) 193–208.
[13] B. Jónsson, Homogeneous universal systems, Math. Scand. 8 (1960) 137–142.
[14] O. Kolman, S. Shelah, Categoricity of theories in $L_{\kappa,\omega}$ when κ is a measurable cardinal. Part I, Fund. Math. 151 (1996) 209–240.
[15] J. Łoś, On the categoricity in power of elementary deductive systems and related problems, Colloq. Math. 3 (1954) 58–62.
[16] M. Makkai, S. Shelah, Categoricity of theories $L_{\kappa\omega}$ with κ a compact cardinal, Ann. Pure Appl. Logic 47 (1990) 41–97.
[17] M. Morley, Categoricity in power, Trans. Amer. Math. Soc. 114 (1965) 514–538.
[18] M. Morley, Omitting classes of elements, in: J.A. Addison, L.H. Henkin, A.T. Tarski (Eds.), The Theory of Models, North-Holland Publ Co, 1965, pp. 265–273.
[19] S. Shelah, Categoricity of uncountable theories, in: L.A. Henkin et al. (Eds.) Proceedings of the Tarski Symposium, Providence, R.I, 1974, pp. 187–203.
[20] S. Shelah, Categoricity in $\aleph_1$ of sentences in $L_{\omega_1,\omega}(\mathbf{Q})$, Israel J. Math. (1975).
[21] S. Shelah, On successors of singular cardinals, in: M.B. Boffa, D.r. van Dalen, K.M. McAloon (Eds.), Logic Colloquium 78, in: Studies in Logic and the Foundation of Mathematics col 97, North-Holland, New York, 1979, pp. 357–380.
[22] S. Shelah, Classification theory for nonelementary classes, I. The number of uncountable models of $\psi \in L_{\omega_1,\omega}$, Part A, Israel J. Math. 46 (1983) 212–240.
[23] S. Shelah, Classification theory for nonelementary classes, I. The number of uncountable models of $\psi \in L_{\omega_1,\omega}$, Part B, Israel J. Math. 46 (1983) 241–273.
[24] S. Shelah, Classification of nonelementary classes, II. Abstract Elementary Classes, in: J.T. Baldwin (Ed.), Classification Theory (Chicago IL, 1985) (Proceedings of the USA-Israel Conference on Classification Theory, Chicago, December 1985), in: Lecture Notes in Mathematics, vol. 1292, Springer, Berlin, 1987, pp. 419–497.
[25] S. Shelah, Universal Classes, in: Classification Theory of Lecture Notes in Mathematics, vol. 1292, Springer, Berlin, 1987, pp. 264–418.
[26] S. Shelah, Classification Theory and the Number of Non-isomorphic Models, 2nd edition, North-Holland, Amsterdam, 1990.
[27] S. Shelah, Cardinal Arithmetic, in: Oxford Logic Guides, vol. 29, Oxford University Press, 1994.
[28] S. Shelah, Categoricity of abstract classes with amalgamation, Ann. Pure Appl. Logic 98 (1999) 261–294.
[29] S. Shelah, On what I do not understand (and have something to say), model theory, Math Japonica 51 (2000) 329–377.
[30] S. Shelah, Categoricity of Theories in $L_{\kappa^*,\omega}$ when κ^* is a measurable cardinal, Part II, Fund. Math. 170 (1–2) (2001) 165–196. Dedicated to the memory of Jerzy Łoś.
[31] S. Shelah, Categoricity of an abstract elementary class in two successive cardinals, Israel J. Math 126 (2001) 29–128.
[32] S. Shelah, Categoricity in abstract elementary classes: going up inductive step, 92 pages (preprint).
[33] S. Shelah, A. Villaveces, Categoricity in abstract elementary classes with no maximal models, Ann. Pure Appl. Logic 97 (1–3) (1999) 1–25.
[34] S. Shelah, A. Villaveces, Categoricity may fail late (submitted for publication).
[35] M. VanDieren, Categoricity and stability in abstract elementary classes, Ph.D. Thesis, May 9, 2002.

Update

Annals of Pure and Applied Logic
Volume 164, Issue 2, February 2013, Page 131–133

DOI: https://doi.org/10.1016/j.apal.2012.09.003

Annals of Pure and Applied Logic 164 (2013) 131–133

Contents lists available at SciVerse ScienceDirect

Annals of Pure and Applied Logic

www.elsevier.com/locate/apal

Erratum

Erratum to "Categoricity in abstract elementary classes with no maximal models" [Ann. Pure Appl. Logic 141 (2006) 108–147]

Monica M. VanDieren [1]

Robert Morris University, 6001 University Blvd, Moon Township, PA 15108, United States

ARTICLE INFO

Article history:
Received 1 June 2012
Received in revised form 10 July 2012
Available online 27 September 2012
Communicated by A.J. Wilkie

ABSTRACT

In the paper "Categoricity in abstract elementary classes with no maximal models", we address gaps in Saharon Shelah and Andrés Villaveces' (1999) proof in [4] of the uniqueness of limit models of cardinality μ in λ-categorical abstract elementary classes with no maximal models, where λ is some cardinal larger than μ. Both [4] (Shelah and Villaveces, 1999) and [5] (VanDieren, 2006) employ set theoretic assumptions, namely GCH and $\Phi_{\mu^+}(S^{\mu^+}_{\mathrm{cf}(\mu)})$.

Recently, Tapani Hyttinen pointed out a problem in an early draft of [3] (Grossberg et al., 2011) to Villaveces. This problem stems from the proof in Shelah and Villaveces' (1999) [4] that reduced towers are continuous. Residues of this problem also infect the proof of Proposition II.7.2 in VanDieren (2006) [5]. We respond to the issues in Shelah and Villaveces (1999) [4] and VanDieren (2006) [5] with alternative proofs under the strengthened assumption that the abstract elementary class is categorical in μ^+.

1. Introduction

A tower is the main construct in [4] and [5] that is used to prove the uniqueness of limit models. A tower is a sequence of length α of amalgamation bases (specifically limit models), denoted by $\bar{M} = \langle M_i \in \mathcal{K}^*_\mu \mid i < \alpha \rangle$, along with a sequence of designated elements $\bar{a} = \langle a_{i+1} \in M_{i+1} \setminus M_i \mid i < \alpha \rangle$ and a sequence of designated submodels $\bar{N} = \langle N_{i+1} \mid i < \alpha \rangle$ for which $M_i \prec_{\mathcal{K}} M_{i+1}$, ga-tp$(a_i/M_i)$ does not μ-split over N_i, and M_i is universal over N_i (see Definition I.5.1 of [5]). Notice that the sequence $\bar{M}$ is not required to be continuous. In fact, many times we will not have continuous towers. For instance, discontinuous towers arise in the proof that an amalgamable tower $(\bar{M}, \bar{a}, \bar{N})$ can be extended to a tower $(\bar{M}', \bar{a}', \bar{N}')$ so that $\bar{a} = \bar{a}'$, $\bar{N} = \bar{N}'$, and the models M'_i are universal extensions of M_i (see Theorem III.10.1 of [5]).

There are a couple of reasons why continuous towers are utilized in the proof the uniqueness of limit models. Because in [4] and [5] we do not have the full amalgamation property, the continuity of towers allows us to avoid models that are not amalgamation bases when we take a union of an increasing chain of towers. Even in an environment which admits full amalgamation, the structure of the proof of the uniqueness of limit models requires a construction of an array of models in which the last row and last column of this array need to be continuous (see the first figure in Part II of [5]). Finding continuous extensions of towers is intrinsic in the construction of the array of models in the proof of the uniqueness of limit models in [3–5].

In [5] we explore two approaches to produce continuous towers. One method is to consider reduced towers and verify that they are continuous (Theorem III.11.4 from [5]) and dense. The other approach is to explicitly construct continuous extensions. This was attempted in Theorem II.7.1 of [5]. In both of these approaches, a gap appears which we fix here assuming

DOI of original article: http://dx.doi.org/10.1016/j.apal.2005.10.006.
E-mail address: vandieren@rmu.edu.
[1] The author was partially sponsored for this work by grant DMS 0801313 of the *National Science Foundation*.

0168-0072/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.apal.2012.09.003

categoricity in μ^+. In Section 2 we provide a proof that reduced towers are continuous in a μ^+-categorical abstract elementary class with no maximal models, and in Section 3 we provide an alternative proof that continuous extensions of towers exist in μ^+-categorical classes. As a result, the following theorem replaces the main result of [4] and Theorem II.9.1 of [5].

Theorem 1. *Assume that $\mathcal{K}$ is a μ^+-categorical abstract elementary class with no maximal models, for some $\mu \geqslant \mathrm{LS}(\mathcal{K})$. Further assume that GCH and $\Phi_{\mu^+}(S^{\mu^+}_{\mathrm{cf}(\mu)})$ hold. Let θ_1 and θ_2 be limit ordinals $< \mu^+$. Under Hypothesis 1,[2] if M_1 and M_2 are (μ, θ_1)- and (μ, θ_2)-limit models over M, respectively, then there exists an isomorphism $f : M_1 \cong M_2$ such that $f \restriction M$ is the identity mapping.*

For the remainder of this document, we assume that $\mathcal{K}$ is an abstract elementary class with no maximal models that is categorical in μ^+, for some $\mu \geqslant \mathrm{LS}(\mathcal{K})$. While this assumption is stronger than assuming categoricity in some λ larger than μ, in the broader context of categoricity results for abstract elementary classes, it is routine to work in a class that is categorical in a successor cardinal. Furthermore, this assumption is sufficient for the application of the uniqueness of limit models in the upward categoricity transfer theorems in [1] and [2].

In this paper, we will also use facts from [5] that follow from GCH and $\Phi_{\mu^+}(S^{\mu^+}_{\mathrm{cf}(\mu)})$: specifically, limit models are amalgamation bases and every amalgamation base of cardinality μ has a universal extension of the same cardinality. We refer the reader to [5] for definitions and notation.

2. Reduced towers are continuous

One method of generating continuous extensions of towers is to restrict all towers to be reduced towers and show that these towers are continuous. The assertion that reduced towers are continuous is made in Theorem 3.1.15 of [4].

There are two problems with the proof of Theorem 3.1.15 in [4]. The first involves inadvertently constructing models that are not amalgamation bases. This problem is fixed in [5] by the introduction of nice towers. The other issue was not known or addressed in [5]. It was first identified by Hyttinen when the problem was reproduced in an early draft of [3]. The difficulty occurs at the induction step of the construction. This step of the construction is isolated as Proposition II.7.2 in [5]. The proposition states that given a tower $(\bar{M}, \bar{a}, \bar{N})$ and an element b outside of the tower, one can find an extension $(\bar{M}', \bar{a}, \bar{N})$ of $(\bar{M}, \bar{a}, \bar{N})$ which contains b. There are several conditions that must be satisfied simultaneously in the construction. It is not clear how all of these conditions can simultaneously hold under the given assumptions.

Below, we prove a variation of Theorem 3.1.15 of [4] that is sufficient to carry out the uniqueness of limit models proof in [5]. It replaces Theorem III.11.4 from [5]. This proof can also be adapted to fix Proposition II.7.2 in [5] under the assumption of categoricity in μ^+.

Theorem 2. *Under the running assumptions of this paper, most notably the assumption that $\mathcal{K}$ is categorical in μ^+, if $(\bar{M}, \bar{a}, \bar{N})$ is a nice, reduced tower constructed of models of cardinality μ, then $(\bar{M}, \bar{a}, \bar{N})$ is continuous.*

Proof. Suppose the theorem fails. Let $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,\alpha}$ be a counter-example of minimal length, α. Notice that by Lemma III.11.5 of [5], we can conclude that $\alpha = \delta + 1$ for some limit ordinal δ and that the failure of continuity must occur at δ. Let $b \in M_\delta \setminus \bigcup_{i<\delta} M_i$ witness the discontinuity of the tower.

By the minimality of α and the density of reduced towers (Theorem III.11.2 of [5]) we can construct a $<^c_{\mu,\delta}$-increasing and continuous chain of reduced, continuous towers $\langle (\bar{M}, \bar{a}, \bar{N})^i \in \mathcal{K}^*_{\mu,\delta} \mid i < \mu^+ \rangle$ with $(\bar{M}, \bar{a}, \bar{N})^0 := (\bar{M}, \bar{a}, \bar{N}) \restriction \delta$. Let $\check{M} := \bigcup_{i<\mu^+, \beta<\delta} M^i_\beta$. Because $\check{M}$ is a model of cardinality μ^+, our categoricity assumption tells us that it must be Galois-saturated. Let $\check{b} \in \check{M}$ realize $\text{ga-tp}(b/\bigcup_{\beta<\delta} M_\beta)$. Fix i so that $\check{b} \in \bigcup_{\beta<\delta} M^i_\beta$. By the equality of the types of b and $\check{b}$ over $\bigcup_{\beta<\delta} M_\beta$, we can fix a $\mathcal{K}$-mapping f so that $f(\check{b}) = b$ and $f \restriction \bigcup_{\beta<\delta} M_\beta$ is the identity. Now consider the tower $(\bar{M}', \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,\alpha}$ defined by setting $M'_\beta := f(M^i_\beta)$ for $\beta < \delta$ and choosing M'_δ to be a limit model which extends $\bigcup_{\beta<\delta} M'_\beta$ and is universal over M_δ. Notice that $(\bar{M}', \bar{a}, \bar{N})$ and b witness that $(\bar{M}, \bar{a}, \bar{N})$ is not reduced. □

3. Continuous extensions exist

Proposition II.7.2 in [5] is used to prove Theorem II.7.1 which asserts that every nice tower has a continuous extension. Avoiding the problem in the proof of Proposition II.7.2, we provide an alternative and simpler proof of the existence of continuous extensions of continuous towers under the assumption of categoricity.

Theorem 3. *Assuming categoricity in μ^+ and given a continuous tower $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,\delta}$, there exists a continuous tower $(\bar{M}', \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,\delta}$ such that $(\bar{M}, \bar{a}, \bar{N}) <^c_{\mu,\delta} (\bar{M}', \bar{a}, \bar{N})$.*

---

[2] Hypothesis 1 of [5] is the statement: every continuous tower has an amalgamable extension inside $\mathfrak{C}$. A more natural statement that implies Hypothesis 1 is that the class of amalgamation bases of cardinality μ is closed under unions of $\prec_{\mathcal{K}}$-increasing chains of length $< \mu^+$. In particular the amalgamation property implies Hypothesis 1.

Proof. We proceed to prove by induction on δ that each continuous tower of length δ has a continuous extension. By the definition of the ordering on towers, $<^c_{\mu,\delta}$, the only difficult stage of the induction is when $\delta = \alpha + 1$ and α is a limit ordinal.

Fix $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,\alpha+1}$ a continuous tower. By our induction hypothesis, the subtower $(\bar{M}, \bar{a}, \bar{N}) \restriction \alpha$ has continuous extensions of length α. Build a $<^c_{\mu,\alpha}$-increasing and continuous chain of continuous towers, $\langle (\bar{M}, \bar{a}, \bar{N})^i \in \mathcal{K}^*_{\mu,\alpha} \mid i < \mu^+ \rangle$, so that $(\bar{M}, \bar{a}, \bar{N})^0 := (\bar{M}, \bar{a}, \bar{N}) \restriction \alpha$. We will show that one of the towers in this chain can be lengthened to a continuous tower that extends $(\bar{M}, \bar{a}, \bar{N})$.

First, extend each one of the towers of length α in this chain $\langle (\bar{M}, \bar{a}, \bar{N})^i \in \mathcal{K}^*_{\mu,\alpha} \mid i < \mu^+ \rangle$ to a tower of length $\alpha + 1$ by defining the last model in the extended tower to be $M^i_\alpha := \bigcup_{\beta<\alpha} M^i_\beta$ for $i < \mu^+$. Consider the top model in this sequence of towers, $\check{M} := \bigcup_{i<\mu^+} M^i_\alpha$. It is a model of cardinality μ^+; so by our categoricity assumption it is Galois-saturated and universal over every model of cardinality μ. Moreover, $\check{M}$ is a (μ, μ^+)-limit model and universal over M_α. In particular, there exists an $i < \mu^+$ such that M^i_α is universal over M_α. Fix such an i and define the tower $(\bar{M}', \bar{a}, \bar{N})$ of length $\alpha + 1$ by setting $M'_\beta := M^i_\beta$ for $\beta \leqslant \alpha$. By our selection of i and by $(\bar{M}, \bar{a}, \bar{N})^i$ being selected as a partial extension of $(\bar{M}, \bar{a}, \bar{N})$, the tower $(\bar{M}', \bar{a}, \bar{N})$ is a continuous extension of $(\bar{M}, \bar{a}, \bar{N})$, as required. □

References

[1] Rami Grossberg, Monica VanDieren, Shelah's categoricity conjecture from a successor for tame abstract elementary classes, J. Symbolic Logic 71 (2) (2006) 553–568.
[2] Rami Grossberg, Monica VanDieren, Categoricity from one successor cardinal in tame abstract elementary classes, J. Math. Log. 6 (2) (2006) 181–201.
[3] Rami Grossberg, Monica VanDieren, Andrés Villaveces, Uniqueness of limit models in abstract elementary classes, preprint, 2011.
[4] Saharon Shelah, Andrés Villaveces, Toward categoricity for classes with no maximal models, Ann. Pure Appl. Logic 97 (1–3) (1999) 1–25.
[5] Monica VanDieren, Categoricity in abstract elementary classes with no maximal models, Ann. Pure Appl. Logic 141 (2006) 108–147.