UNIQUENESS OF LIMIT MODELS IN CLASSES WITH AMALGAMATION

RAMI GROSSBERG, MONICA VANDIEREN, AND ANDRÉS VILLAVECES

Abstract. We prove:

Theorem 0.1 (Main Theorem)**.** *Let $\mathcal{K}$ be an Abstract Elementary Class and $\mu > \mathrm{LS}(\mathcal{K})$. Suppose $\mathcal{K}$ satisfies the disjoint amalgamation property for models of cardinality μ. If $\mathcal{K}$ is μ-Galois-stable, does not have long splitting chains, and satisfies locality of splitting, then any two (μ, σ_ℓ)-limits over M for ($\ell \in \{1, 2\}$) are isomorphic over M.*

This theorem extends results of Shelah from [Sh 394], [Sh 576], [Sh 600], Kolman and Shelah in [KoSh] and Shelah and Villaveces from [ShVi]. Our uniqueness theorem was used by Grossberg and VanDieren to prove a case of Shelah's categoricity conjecture for tame Abstract Elementary Classes in [GrVa2]. Theorem 0.1 stands in for the so far underdeveloped superstability for categorical abstract elementary classes.

1. introduction

In 1977, Shelah, building on the work of Jónsson and Fraïssé, identified a non-elementary context in which a model theoretic analysis could be carried out. Shelah began to study classes of models, together with a partial ordering of the class, which exhibit many of the properties that the models of a first order theory have with respect to the elementary submodel relation. Such classes were named abstract elementary classes. They are broad enough to generalize $L_{\omega_1,\omega}(\mathbf{Q})$. Both classification theory and stability theory may be carried out to some extent within these classes. One strong advantage is that there are no a priori compactness assumptions. We reproduce the definition here.

Definition 1.1. Let $\mathcal{K}$ be a class of structures all in the same similarity type $\mathrm{L}(\mathcal{K})$, and let $\prec_{\mathcal{K}}$ be a partial order on $\mathcal{K}$. The ordered pair $\langle \mathcal{K}, \prec_{\mathcal{K}} \rangle$ is an *abstract elementary class, AEC for short* iff

A0 (Closure under isomorphism)

(a) For every $M \in \mathcal{K}$ and every $\mathrm{L}(\mathcal{K})$-structure N if $M \cong N$ then $N \in \mathcal{K}$.

Date: November 2, 2008.

AMS Subject Classification: Primary: 03C45, 03C52, 03C75. Secondary: 03C05, 03C55 and 03C95.

The third author was partially sponsored for this work by grant No. 1101-05-13605 of the *Instituto Colombiano para el Desarrollo de la Ciencia y la Tecnología*, Colciencias.

(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 $M_1 \prec_{\mathcal{K}} M_2$.

A1 For all $M, N \in \mathcal{K}$ if $M \prec_{\mathcal{K}} N$ then $M \subseteq N$.

A2 Let M, N, M^* be $\mathrm{L}(\mathcal{K})$-structures. If $M \subseteq N$, $M \prec_{\mathcal{K}} M^*$ and $N \prec_{\mathcal{K}} M^*$ then $M \prec_{\mathcal{K}} N$.

A3 (Downward Löwenheim-Skolem) There exists a cardinal $\mathrm{LS}(\mathcal{K}) \geq \aleph_0 + |\,\mathrm{L}(\mathcal{K})|$ such that for every $M \in \mathcal{K}$ and for every $A \subseteq |M|$ there exists $N \in \mathcal{K}$ such that $N \prec_{\mathcal{K}} M$, $|N| \supseteq A$ and $\|N\| \leq |A| + \mathrm{LS}(\mathcal{K})$.

A4 (Tarski-Vaught Chain)

(a) For every regular cardinal μ and every $N \in \mathcal{K}$ if $\{M_i \prec_{\mathcal{K}} N \ : \ i < \mu\} \subseteq \mathcal{K}$ is $\prec_{\mathcal{K}}$-increasing (i.e. $i < j \Longrightarrow M_i \prec_{\mathcal{K}} M_j$) then $\bigcup_{i<\mu} M_i \in \mathcal{K}$ and $\bigcup_{i<\mu} M_i \prec_{\mathcal{K}} N$.

(b) For every regular μ, if $\{M_i \ : \ i < \mu\} \subseteq \mathcal{K}$ is $\prec_{\mathcal{K}}$-increasing then $\bigcup_{i<\mu} M_i \in \mathcal{K}$ and $M_0 \prec_{\mathcal{K}} \bigcup_{i<\mu} M_i$.

For M and $N \in \mathcal{K}$ a monomorphism $f : M \to N$ is called an $\mathcal{K}$-*embedding* iff $f[M] \prec_{\mathcal{K}} N$. Thus, $M \prec_{\mathcal{K}} N$ is equivalent to "id_M is a $\mathcal{K}$-embedding from M into N".

For $M_0 \prec_{\mathcal{K}} M_1$ and $N \in \mathcal{K}$, the formula $f : M_1 \underset{M_0}{\to} N$ stands for f is a $\mathcal{K}$-embedding such that $f \restriction M_0 = \mathrm{id}_{M_0}$.

For a class $\mathcal{K}$ and a cardinal $\mu \geq \mathrm{LS}(\mathcal{K})$ let

$$\mathcal{K}_\mu := \{M \in \mathcal{K} \ : \ \|M\| = \mu\}.$$

In reality, abstract elementary classes were not as approachable as one would hope and much work in non-elementary model theory takes place in contexts which additionally satisfy the amalgamation property:

Definition 1.2. Let $\mu \geq \mathrm{LS}(\mathcal{K})$. We say that $\mathcal{K}$ has the μ-*amalgamation property* (μ-*AP*) iff for any $M_\ell \in \mathcal{K}_\mu$ (for $\ell \in \{0, 1, 2\}$) such that $M_0 \prec_{\mathcal{K}} M_1$ and $M_0 \prec_{\mathcal{K}} M_2$ there are $N \in \mathcal{K}_\mu$ and $\mathcal{K}$-embeddings $f_\ell : M_\ell \to N$ such that $f_\ell \restriction M_0 = \mathrm{id}_{M_0}$ for $\ell = 1, 2$.

A model $M_0 \in \mathcal{K}_\mu$ satisfying the above requirement is called an *amalgamation base*.

We say that $\mathcal{K}$ has the *amalgamation property* (*AP*) iff any triple of models from $\mathcal{K}_{\geq \mathrm{LS}(\mathcal{K})}$ can be amalgamated.

Remark 1.3.
(1) Using the isomorphism axioms we can see that $\mathcal{K}$ has the λ-AP iff for any $M_\ell \in \mathcal{K}_\lambda$ (for $\ell \in \{0, 1, 2\}$) such that $M_0 \prec_{\mathcal{K}} M_\ell$ (for $\ell \in \{1, 2\}$) there are $N \in \mathcal{K}_\lambda$ and $f : M_1 \underset{M_0}{\to} N$ such that $N \succ_{\mathcal{K}} M_2$.
(2) Using the axioms of AECs it is not difficult to prove that if $\mathcal{K}$ has the λ-AP for every $\lambda \geq \mathrm{LS}(\mathcal{K})$ then $\mathcal{K}$ has the AP.

A stronger version of the amalgamation property is

Definition 1.4. Let $\mathcal{K}$ be an abstract class. $\mathcal{K}$ has the λ-*Disjoint Amalgamation Property* iff for every $M_\ell \in \mathcal{K}_\lambda$ (for $\ell = 0, 1, 2$) such that $M_0 \prec_{\mathcal{K}} M_\ell$ (for $\ell = 1, 2$) there are $N \in \mathcal{K}_\lambda$ which is a $\mathcal{K}$-extension of M_2 and a $\mathcal{K}$-embedding $f : M_1 \underset{M_0}{\rightarrow} N$ such that $f[M_1] \cap M_2 = M_0$.

We say that a class has the *disjoint amalgamation property* iff it has the λ-disjoint amalgamation property for every $\lambda \geq \mathrm{LS}(\mathcal{K}) + \aleph_0$. We write DAP for short.

An application of the compactness theorem establishes:

Fact 1.5. *If T is a complete first-order theory then $\langle \mathrm{Mod}(T), \prec\rangle$ has the λ-DAP for all $\lambda \geq |\,\mathrm{L}(T)| + \aleph_0$*

The roots of the following fact can be traced back to Jónsson's 1960 paper [Jo], the present formulation is from [Gr1]:

Fact 1.6. *Let $\langle \mathcal{K}, \prec_{\mathcal{K}}\rangle$ be an AEC and $\lambda \geq \kappa > \mathrm{LS}(\mathcal{K})$ such that $K_{<\lambda}$ has the AP and the JEP. Suppose $M \in \mathcal{K}$.*

If $\lambda^{<\kappa} = \lambda \geq \|M\|$ then there exists $N \succ M$ of cardinality λ which is κ-model-homogeneous.

Thus if an AEC $\mathcal{K}$ has the amalgamation property then like in first-order stability theory we may assume that there is a large model-homogeneous $\mathfrak{C} \in \mathcal{K}$, that acts like a monster model.

We will refer to the model $\mathfrak{C}$ as the *monster model*. All models considered will be of size less than $\|\mathfrak{C}\|$, and we will find realizations of types we construct inside this monster model.

From now on, we assume that the monster model $\mathfrak{C}$ has been fixed.

Notions of types as sets of formulas, even when the class is described in some infinitary logic, do not behave as nicely as in first-order logic. Thus we need a replacement which was introduced by Shelah in [Sh 394], in order to avoid confusion with the classical notion following [Gr2] we call this newer, different notion *Galois-type*.

Since in this paper we deal only with AECs with the AP property, the notion of Galois type has a simpler definition than in the general case.

Definition 1.7 (Galois types)**.** Suppose that $\mathcal{K}$ has the AP.

(1) Given $M \in \mathcal{K}$ consider the action of $\mathrm{Aut}_M(\mathfrak{C})$ on $\mathfrak{C}$, for an element $a \in |\mathfrak{C}|$ let ga-tp(a/M) denote the *Galois type of a over M* which is defined as the orbit of a under $\mathrm{Aut}_M(\mathfrak{C})$.
(2) For $M \in \mathcal{K}$, we let
$$\text{ga-S}(M) = \{\ \text{ga-tp}(a/M)\ :\ a \in |\mathfrak{C}|\}.$$
(3) $\mathcal{K}$ is λ-*Galois-stable* iff
$$N \in \mathcal{K}_\lambda \implies |\,\text{ga-S}(N)| \leq \lambda.$$
(4) Given $p \in \text{ga-S}(M)$ and $N \in \mathcal{K}$ such that $N \succ_{\mathcal{K}} M$, we say that p is *realized* by $a \in N$, iff ga-tp$(a/M) = p$. Just as in the first-order case we will write $a \models p$ when a is a realization of p.

For a more detailed discussion of Galois types, extensions and restrictions, equivalent and more general formulations, the reader may consult [Gr2].

The main concept of this paper is Shelah's *limit model* which we will show serves as a substitute for the role of saturation in stability theory. Why do we need substitutes for the role of saturation? In homogeneous abstract elementary classes (see, for example, [GrLe]) where one may study classes of models omitting given sets of types, even the existence of a saturated model presents some problems. Instead, Galois-saturated models which realize as many types as possible are used. However, when stability theory has been ported to contexts more general than first order logic, many situations have appeared when Galois-saturated models do not fulfill the main roles that saturated models play in elementary classes. For example, under the assumption of categoricity (under reasonable stability conditions), uniqueness of Galois-saturated models is not straightforward and is proved by first considering limit models. Thus, looking for notions that may appropriately substitute the role of saturated models is crucial.

We first need to define universal extensions as they are the building blocks of limit models:

Definition 1.8. (1) Let κ be a cardinal $\geq \mathrm{LS}(\mathcal{K})$. We say $M^* \succ_{\mathcal{K}} N$ is *κ-universal over N* iff for every $N' \in \mathcal{K}_\kappa$ with $N \prec_{\mathcal{K}} N'$ there exists a $\mathcal{K}$-embedding $g : N' \underset{N}{\rightarrow} M^*$ such that:

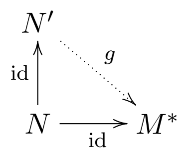

(2) We say M^* is *universal over N* or M^* is *a universal extension of N* iff M^* is $\|N\|$-universal over N.

Theorem 1.9 (Existence)**.** *Let $\mathcal{K}$ be an AEC without maximal models and suppose it is Galois-stable in μ. If $\mathcal{K}$ has the amalgamation property then for every $N \in \mathcal{K}_\mu$ there exists $M^* \gneq_{\mathcal{K}} N$, universal over N of cardinality μ.*

This theorem was stated without proof as Claim 1.16 in [Sh 600], for a proof see [GrVa1] or [Gr1].

In [KoSh] and in [Sh 576] a substitute for saturated was introduced: (μ, α)*-saturated models*. Shelah in [Sh 600] calls this notion *brimmed* and in his later paper with Villaveces [ShVi] the name *limit models* is used. We use the more recent terminology.

Definition 1.10. [Limit models] Let $\mu \geq \mathrm{LS}(\mathcal{K})$ and $\alpha < \mu^+$ a limit ordinal and $N \in \mathcal{K}_\mu$. We say that M is (μ, α)*-limit over N* iff there exists an increasing and continuous chain $\{M_i \mid i < \alpha\} \subseteq \mathcal{K}_\mu$ such that $M_0 = N$, $M = \bigcup_{i<\alpha} M_i$, M_i is a proper $\mathcal{K}$-submodel of M_{i+1} and M_{i+1} is universal over M_i for all $i < \alpha$.

From Theorem 1.9 we get that for $\alpha \leq \mu^+$ there always exists a (μ, α)-limit model provided $\mathcal{K}$ has the AP, has no maximal models and is μ-Galois-stable.

The following theorem partially clarifies the analogy with saturated models:

Theorem 1.11. *Let T be a complete first-order theory and let $\mathcal{K}$ be the elementary class* $\mathrm{Mod}(T)$ *with the usual notion of elementary submodels.*

(1) *Suppose T is superstable. If M is (μ, δ)-limit model for δ a limit ordinal, then M is saturated.*
(2) *Suppose T is stable. If M is (μ, δ)-limit model for δ a limit ordinal with* $\mathrm{cf}\,\delta \geq |T|^+$, *then M is saturated.*

Proof. Exercise to the reader; use an argument similar to the proof of [Sh e, Theorem III 3.11]. $\dashv$

Thus in elementary classes superstability implies that limit models are saturated, in particular are unique. This raises the following natural question for the situation in AECs:

Question 1.12 (Uniqueness problem)**.** *Let $\mathcal{K}$ be an AEC, $\mu \geq \mathrm{LS}(\mathcal{K})$, $\sigma_1, \sigma_2 < \mu^+$, $M \in \mathcal{K}_\mu$ and suppose that N_ℓ (μ, σ_ℓ)-limit models over M. What "reasonable" assumptions on $\mathcal{K}$ will imply that $\exists f : N_1 \cong_M N_2$?*

Using a back and forth argument one can show that when $\mathrm{cf}\,\sigma_1 = \mathrm{cf}\,\sigma_2$ then we get uniqueness without any assumptions on $\mathcal{K}$. More precisely:

Fact 1.13. *Let $\mu \geq \mathrm{LS}(\mathcal{K})$ 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 $g \restriction M = \mathrm{id}_M$. 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 1.14. *Let μ be a cardinal and σ a limit ordinal with $\sigma < \mu^+$. If M is a (μ, σ)-limit model, then M is a $(\mu, \mathrm{cf}(\sigma))$-limit model.*

Therefore, Question 1.12 is non-trivial only for the case where $\mathrm{cf}(\sigma_1) \neq \mathrm{cf}(\sigma_2)$.

Limit models are not necessarily unique even for first order complete stable theories.

Theorem 1.15. *Suppose T is a complete stable theory. Let $\mu \geq 2^{|T|}$ such that $\mu^{|T|} = \mu$. If T is not superstable then there are $M \models T$ and $N \succ M$, $N \neq M$, both of cardinality μ such that N is (μ, ω)-limit over M but not isomorphic to any (μ, κ)-limit model over M for $\mathrm{cf}(\kappa) \geq \kappa(T)$.*

Proof. As T is unsuperstable, by [Sh e, Lemma VII, 3.5 (2)], there are $\lambda := (2^\mu)^+$, $\{\bar{a}_\eta | \eta \in {}^{\omega \geq}\lambda\}$ and $\{\varphi_n(\bar{x}, \bar{y}_n) | n < \omega\}$ such that

$$\forall n < \omega \forall \eta \in {}^{\omega}\lambda \forall \nu \in {}^{n}\lambda \mathfrak{C} \models \varphi_n[\bar{a}_\eta, \bar{a}_\nu] \Leftrightarrow \nu < \eta.$$

By induction on $n < \omega$ define $\{M_n | n < \omega\}$ all of cardinality μ and $\{\eta_n, \nu_n | n < \omega\}$ such that

(1) M_0 is saturated,
(2) M_{n+1} is universal over M_n and saturated of cardinality μ,
(3) $\eta_{n+1} > \eta_n \wedge \nu_{n+1} > \nu_n \wedge \eta_{n+1} \neq \nu_{n+1}$,
(4) $\bar{a}_{\eta_{n+1}}, \bar{a}_{\nu_{n+1}} \in M_{n+1}$
(5) $\text{ga-tp}(\bar{a}_{\eta_{n+1}}/M_n) = \text{ga-tp}(\bar{a}_{\nu_{n+1}}/M_n)$.

This is enough: Let $N = \bigcup_{n<\omega} M_n$. Clearly N is (μ,ω)-limit over M_0. Let $N' \models T$ be (μ,κ)-limit over M_0. By Theorem 1.11(2), N' must be saturated. It is enough to show that N is not saturated. Consider $p := \{\varphi_{n+1}(\bar{x};\bar{a}_{\eta_{n+1}}) \wedge \neg\varphi_{n+1}(\bar{x};\bar{a}_{\nu_{n+1}}) | n < \omega\}$. The set of formulas p is a type since it is realized in $\mathfrak{C}$ by $\bar{a}_\eta$ where $\eta := \bigcup_{n<\omega} \eta_n$. However, if $\bar{a} \in M$ then $\exists n < \omega \bar{a} \in M_n$, by (5) above, so

$$\mathfrak{C} \models \varphi_{n+2}[\bar{a}, \bar{a}_{\eta_{n+1}}] \Leftrightarrow \mathfrak{C} \models \varphi_{n+2}[\bar{a}, \bar{a}_{\nu_{n+1}}],$$

so $\bar{a} \not\models p$.

This is possible: By stability and $\mu^{|T|} = \mu$, T must be μ-stable; using the proof of [Sh e, Th. III 3.12], every model of cardinality μ has a saturated proper elementary extension. Let M_0 be this model and take $\eta_0 = \nu_0 := \langle\rangle$. Given η_n, ν_n, M_n, using Theorem 1.9 let $M^* \succ M$, universal over M_n of cardinality μ. Let $M^{**} \succ M^*$ of cardinality μ containing $\bar{a}_{\eta_n}$ and $\bar{a}_{\nu_n}$. By [Sh e, Th. III 3.12], we can take $M_{n+1} \succ M^{**}$ saturated of cardinality μ. Clearly it is universal over M_n. Consider $F(\alpha) := \text{tp}(\bar{a}_{\eta_n}\hat{\ }\alpha/M_n)$. As λ is regular and $\lambda > |S(M_n)|$, there is $S \subset \lambda$ of cardinality λ such that $\alpha \neq \beta \in S \Rightarrow F(\alpha) = F(\beta)$. Pick $\alpha \neq \beta \in S$ and define $\eta_{n+1} := \eta_n\hat{\ }\alpha$ and $\nu_{n+1} := \nu_n\hat{\ }\beta$. ⊣

A similar proof provides a fine analysis of the difference between limit models and saturated models in the strictly stable case in first order, for uncountable theories:

Theorem 1.16. *Let T be a complete first order theory in an uncountable language. If T is stable in some $\mu > \aleph_0$ but not superstable then there exists a (μ,ω)-limit model that is not isomorphic to any (μ,κ)-limit when* $\text{cf}(\kappa) \geq \kappa(T)$.

Proof. Left to the reader - similar to the previous. ⊣

The main result of this paper is (check Definition 1.4):

Theorem 1.17 (Main Theorem)**.** *Let $\mathcal{K}$ be an AEC without maximal models, and $\mu > \text{LS}(\mathcal{K})$. Suppose $\mathcal{K}$ satisfies the μ-DAP. If $\mathcal{K}$ is μ-Galois-stable, does not have long splitting chains, and satisfies locality of splitting*[1]*, then any two (μ,σ_ℓ)-limits over M for $(\ell \in \{1,2\})$ are isomorphic over M.*

Notice that μ-DAP is occasionally a property we get for free if the class $\mathcal{K}$ has an axiomatization in a logic with sufficient compactness; essentially, Robinson's consistency property is enough. In other occasions DAP is a

[1]See Assumption 2.5 for the precise description of long splitting chains and locality

known corollary of categoricity, even when AP is not assumed. When a class has arbitrarily large models, satisfies the amalgamation property and is categorical, then μ-DAP below the categoricity cardinal holds - amalgamation may be replaced by 'no maximal models' if the disjoint amalgamation needed is only over some distinguished models (see [ShVi] and [Va]).

Approximations to Theorem 1.17 and its relatives were considered by several authors:

In Theorem 6.5 of [Sh 394], Shelah claims uniqueness of limit models of cardinality μ for classes with the amalgamation property under little more than categoricity in some $\lambda > \mu > \mathrm{LS}(\mathcal{K})$ together with existence of arbitrarily large models. The argument in [Sh 394] depends in a crucial way on an analysis of Ehrenfeucht-Mostowski models; however unlike [Sh 394] since we don't assume here categoricity and existence of models above the Hanf number, we are not allowed to use the Ehrenfeucht-Mostowski machinery. Under similar assumptions as Shelah's result mentioned above, more recently, Baldwin in [Ba] (Chapter 11) has used methods based on [Sh 394] to prove that if M_1 and M_2 are (μ, σ_1) and (μ, σ_2) limit models over N then $M_1 \cong M_2$ - Baldwin, however, does not prove that M_1 and M_2 are isomorphic *over* N. Our result is therefore much stronger than that in [Ba]. Shelah's claim in Theorem 6.5 of [Sh 394] (isomorphism over the base) seems too strong for the proof that he suggests. Instead, he proves that (μ, κ)-limit models are Galois saturated.

Question 1.18. *Is it possible (for AEC) to have uniqueness for limit models without it being Galois-saturated? We think that this is possible in singulars like $\beth_\omega$ (Notice that in Theorem 1.15 what is really proved is that if T is strictly stable then there* exists *a λ such that uniqueness for (λ, α)-limit models fails. In the proof we work λ is of the form $(2^\mu)^+$, in particular regular (necessary for the argument).*

Kolman and Shelah in [KoSh] prove the uniqueness of limit models of cardinality μ in λ-categorical AECs that are axiomatized by a $L_{\kappa,\omega}$-sentence where $\lambda > \mu$ and κ is a measurable cardinal. Then Kolman and Shelah use this uniqueness result to prove that amalgamation occurs below the categoricity cardinal in $L_{\kappa,\omega}$-theories with κ measurable. Both the measurability of κ and the categoricity are used integrally in their proof of uniqueness.

Shelah in [Sh 576] (see Claim 7.8) proved a special case of the uniqueness of limit models under the assumption of μ-AP, categoricity in μ and in μ^+ as well as assuming $K_{\mu^{++}} \neq \emptyset$. In that paper Shelah needs to produce *reduced types* and use some of their special properties.

[ShVi] attempted to prove a uniqueness theorem without assuming any form of amalgamation; however, they assumed that $\mathcal{K}$ is categorical in some $\lambda > \mathrm{Hanf}(\mathcal{K}) + \mu$ and that every model in $\mathcal{K}$ has a proper extension. VanDieren in [Va] managed to prove the above uniqueness statement under the assumptions of [ShVi] together with the additional assumption that

$\mathcal{K}^{am} := \{M \in \mathcal{K}_\mu \mid M \text{ is an amalgamation base}\}$ is closed under unions of increasing $\prec_{\mathcal{K}}$ chains.

In [Sh 600] the basic context is that of a *good frame*, which is an axiomatization of the notion of superstability. Its full definition is more than a page long. Shelah's assumptions on the AEC include, among other things, the amalgamation property, the existence of a forking like dependence relation and of a family of types playing a role akin to that of regular types in first order superstable theories – Shelah calls them *bs*-types – and several requirements on the interaction of these types and the dependence relation. One of the axioms of a good frame is the existence of a non-maximal super-limit model. This axiom along with μ stability implies the uniqueness of limit models of cardinality μ. In Claim 4.8 of [Sh 600] he states that in a good frame limit models are unique (i.e. the same conclusion of our Main Theorem). (While we don't claim that we understand Shelah's proof or believe in its correctness, he explicitly uses the interplay between *bs*-types and the forking notion as well as no long forking chains and continuity of forking.) Thus, the main differences are two: first, our assumptions on $\mathcal{K}$ are weaker than what Shelah is using (and our use of various versions of superstability is different from that of [Sh 600], as we do not require the full power of good frames) and second, our methods are quite different from his.

The formal differences between our approaches can be summarized as follows:

(a) Suppose that $\mathcal{K}$ is an AEC satisfying the disjoint amalgamation property and is categorical in λ^+ for some $\lambda > \mathrm{LS}(\mathcal{K})$; we then get uniqueness of limit models and no splitting chains of length ω. This result is used in [GrVa2] to conclude that $\mathcal{K}$ is categorical in all $\mu > \mathrm{LS}(\mathcal{K})^+$. In this case DAP follows from the other assumptions. By way of comparison, in order to get a good frame, Shelah needs results of [Sh 576] (a 99 pages-long paper) and [Sh 705] (220 pages) to conclude that good frames exist from the assumption of categoricity in several consecutive cardinals + several weak-diamonds. On the other hand, all our results are in ZFC.

(b) In the case when $\mathcal{K}$ is the class of models of a complete first order theory T, Shelah's proof in [Sh 600, Claim 4.8] really uses the full power of assuming that T is *super*stable. The proof of uniqueness in this paper just needs, in addition to the stability of T, no splitting chains of length ω. As the main interest of our theorem is for the general case of AEC, rather than just for first order theories, the difference between this paper and [Sh 600, Claim 4.8] is clearer when understood in light of the greater picture as in (a) above.

The reason for these differences is that Shelah's papers [Sh 576], [Sh 600] and [Sh 705] focus on a problem entirely different from [GrVa2]'s. Grossberg and VanDieren's [GrVa2] (as well as [GrVa0]) were written with Shelah's categoricity conjecture in mind where the basic assumption is categoricity in a cardinal above $\mathrm{Hanf}(\mathcal{K})$. Shelah's above mentioned work was motivated by questions asked by Grossberg in fall 1994 aimed to generalize [Sh 87b] to

AECs which are not $\mathrm{PC}_{\aleph_0,\aleph_0}$. As the underlying, driving problems of Shelah and of Grossberg and VanDieren's work are quite different also the methods used to solve them differ, however occasionally the same concepts appear in both.

We are particularly interested in Theorem 1.17 not only for the sake of generalizing Shelah's result from [Sh 576] but due to the fact that the first and second author use this uniqueness theorem along with tools from [Sh 394] in a crucial step to prove:

Theorem 1.19 (Upward categoricity theorem, [GrVa2])**.** *Suppose that $\mathcal{K}$ has arbitrarily large models, is χ-tame and satisfies the amalgamation and joint embedding properties. Let λ be such that $\lambda > \mathrm{LS}(\mathcal{K})$ and $\lambda \geq \chi$. If $\mathcal{K}$ is categorical in λ^+ then $\mathcal{K}$ is categorical in all $\mu \geq \lambda^+$.*

Grossberg and VanDieren's use of the uniqueness of limit models in this theorem hints at a connection between classical definitions of superstability in first order logic and the uniqueness of limit models. This link is explored in further work of VanDieren.

Years after learning Theorem 1.19, Baldwin presented an alternative proof of Theorem 1.19 in [Ba].

We thank John Baldwin for helping to clarify the presentation. We also thank the referee for valuable suggestions and remarks.

2. The Setting

We will prove the uniqueness of limit models in classes which are equipped with a moderately well-behaved dependence relation. Thus we will assume that $\mathcal{K}$ is stable in μ. We will use μ-splitting as the dependence relation, but any dependence relation which is local and has existence, uniqueness and extension properties suffices.

Definition 2.1. A type $p \in \text{ga-S}(M)$ *μ-splits* over $N \in \mathcal{K}_{\leq\mu}$ if and only if N is a $\prec_{\mathcal{K}}$-submodel of M of cardinality μ and there there exist $N_1, N_2 \in \mathcal{K}_\mu$ and a $\mathcal{K}$-mapping h such that $N \prec_{\mathcal{K}} N_l \prec_{\mathcal{K}} M$ for $l = 1,2$ and $h : N_1 \to N_2$ with $h \restriction N = \mathrm{id}_N$ and $p \restriction N_2 \neq h(p \restriction N_1)$.

The existence property for non-μ-splitting types follows from Galois stability in μ:

Fact 2.2 (Claim 3.3 of [Sh 394])**.** *Assume $\mathcal{K}$ is an abstract elementary class and is Galois-stable in μ. For every $M \in \mathcal{K}_{\geq\mu}$ and $p \in \text{ga-S}(M)$, there exists $N \in \mathcal{K}_\mu$ such that p does not μ-split over N.*

The uniqueness and extension property of non-μ-splitting types holds for types over limit models:

Fact 2.3 (Theorem I.4.15 of [Va])**.** *Suppose that $\mathcal{K}$ is an AEC. Let $N, M, M' \in \mathcal{K}_\mu$ be such that M' is universal over M and M is universal over N. If $p \in \text{ga-S}(M)$ does not μ-split over N, then there is a unique $p' \in \text{ga-S}(M')$ such that p' extends p and p' does not μ-split over N.*

A variation of this fact is later used in an induction construction in the proof of Theorem 5.8. We state it explicitly here:

Fact 2.4 (Theorem I.4.10 of [Va]). *Let M, N, M^* be models in $\mathcal{K}_\mu$. Suppose that M is universal over N and that M^* is universal over M. If a type $p = \text{ga-tp}(a/M)$ does not μ-split over N then there exists an automorphism g of $\mathfrak{C}$ fixing N such that $\text{ga-tp}(g(a)/M^*)$ does not μ-split over N and $\text{ga-tp}(g(a)/M) = p$.*

Here are the assumptions of the paper:

Assumption 2.5. $\mathcal{K}$ is an AEC with the μ-DAP and JEP and $\mathcal{K}$ satisfies the following properties:

(1) $\mathcal{K}$ is stable in μ.
(2) All models are submodels of a fixed monster model $\mathfrak{C}$.
(3) μ-splitting in $\mathcal{K}$ satisfies the following locality (also called continuity) and existence properties.
For all infinite α, for every sequence $\langle M_i \mid i < \alpha\rangle$ of limit models of cardinality μ and for every $p \in \text{ga-S}(M_\alpha)$ we have that
  (a) If for every $i < \alpha$, the type $p \restriction M_i$ does not μ-split over M_0, then p does not μ-split over M_0.
  (b) There exists $i < \alpha$ such that p does not μ-split over M_i.

Remark 2.6. In the context of an AEC with the full amalgamation property and JEP, categoricity in a cardinal $\lambda > \mu$ implies all parts of Assumption 2.5. For a proof of Assumption 2.5.1 from categoricity, see Claim 1.7 of [Sh 394] or [Ba]. The observation that assumption 2.5(3a) follows from categoricity is a consequence of Observation 6.2 and Main Lemma 9.4 of [Sh 394]. Lemma 6.3 of [Sh 394] is the statement that assumption 2.5(3b) follows from categoricity when the cofinality of the categoricity cardinal is larger than μ.

Theorem 2.7 (Long existence follows from stability in FO). *Suppose that T is first order complete. If T is stable then Assumption 2.5(3b) holds for α such that $\text{cf}(\alpha) \geq |T|^+$.*

Proof. Let $\langle M_i | i \leq \alpha\}$ be an increasing sequence of saturated models $p \in S(M_\alpha)$ such that $\forall i < \alpha$, p μ-splits over M_i. Let $\varphi_i(\bar{x}, \bar{y})$ be a formula witnessing the splitting of $p \restriction M_{i+1}$ over M_i, and let $\bar{a}_i, \bar{b}_i \in M_{i+1}$. As $\text{cf}(\alpha) \geq |T|^+$, $\exists S \subset \alpha$ infinite such that $i, j \in S \Rightarrow \varphi_i = \varphi_j$.

Wlog, suppose that $\langle M_n | n \leq \omega\rangle$ is an increasing sequence of saturated models, and $p \in S_\varphi(M_\omega)$ is such that $\bar{a}_i, \bar{b}_i \in M_{i+1}$ witness that $p \restriction M_{i+1}$ splits over M_i. Then $p(x_1, \bar{y}_1, \bar{z}_1, x_2, \bar{y}_2, \bar{z}_2)$ and $\{\bar{d}_i | i < \omega\}$ witness that p has the order property, where $\bar{d}_i = \bar{a}_i \hat{\ } \bar{b}_i \hat{\ } c_i$ and

$$c_i \in M_{i+2} \models p \restriction \{\bar{a}_k, \bar{b}_k | k \leq i\} \cup \{d_k | k < i\}.$$

See [Gr1, Lemma IV, 2.12]. $\dashv$

Remark 2.8. Assumption 2.5 holds in contexts without the assumption of categoricity. The Disjoint Amalgamation Property (DAP) comes for free in first order classes of the form $(Mod(T), \prec)$ for complete T. DAP also holds in homogeneous classes (see [Sh 3] or [Po]), in excellent classes (see [Sh 87b]) and is an axiom in the definition of finitary classes (see [HyKe]). It also holds for cats consisting of existentially closed models of positive Robinson theories ([Za]). In each of these contexts dependence relations satisfying Assumption 2.5 have been developed. Finally, the locality and existence of non-μ-splitting extensions are akin to consequences of superstability in first order logic.

3. Strong Types

Under the assumption of μ-stability, we can define *strong types* as in [ShVi]. These strong types will allow us to achieve a better control of extensions of towers of models than what we obtain using just Galois types.

Definition 3.1 (Definition 3.2.1 of [ShVi])**.** For M a (μ, θ)-limit model (see definition 1.10),

(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 $M' \in \mathcal{K}_\mu$ 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.

(3) Two strong types $(p_1, N_1) \in \mathfrak{St}(M_1)$ and $(p_2, N_2) \in \mathfrak{St}(M_2)$ are *parallel* iff for every M' of cardinality μ extending M_1 and M_2 there exists $q \in$ ga-S(M') such that q extends both p_1 and p_2 and q does not μ-split over N_1 and N_2. We use the notation $(p_1, N_1) \| (p_2, N_2)$ for (p_1, N_1) is parallel to (p_2, N_2).

Remark 3.2. Under the assumption of the existence of universal extensions, it is equivalent to say two strong types $(p_1, N_1) \in \mathfrak{St}(M_1)$ and $(p_2, N_2) \in \mathfrak{St}(M_2)$ are *parallel* iff for some M' of cardinality μ universal over some common extension of M_1 and M_2 there exists $q \in$ ga-S(M') such that q extends both p_1 and p_2 and q does not μ-split over N_1 and N_2.

Lemma 3.3 (Monotonicity of parallel types)**.** *Suppose $M_0, M_1 \in \mathcal{K}_\mu$ and $M_0 \prec_{\mathcal{K}} M_1$ and $(p, N) \in \mathfrak{St}(M_1)$. If M_0 is universal over N, then we have $(p \restriction M_0, N) \| (p, N)$.*

Proof. Straightforward using the uniqueness of non-μ-splitting extensions. ⊣

Notation 3.4. Let $M, M' \in \mathcal{K}_\mu$, Suppose that M is a $\prec_{\mathcal{K}}$-submodel of M'. For $(p, N) \in \mathfrak{St}(M')$, if M is universal over N, we define the restriction $(p, N) \upharpoonright M \in \mathfrak{St}(M)$ to be $(p \upharpoonright M, N)$.

If we write $(p, N) \upharpoonright M$, we mean that p does not μ-split over N and M is universal over N.

Notice that $\sim$ is an equivalence relation on $\mathfrak{St}(M)$ (see [Va]). Stability in μ implies that there are few strong types over any model of cardinality μ:

Fact 3.5 (Claim 3.2.2 (3) of [ShVi])**.** *If $\mathcal{K}$ is Galois-stable in μ, then for any $M \in \mathcal{K}$ of cardinality μ, $|\mathfrak{St}(M)/\sim| \leq \mu$.*

4. Towers

To each (μ, θ)-limit model M we can naturally associate a continuous chain $\bar{M} = \langle M_i \in \mathcal{K}_\mu \mid i < \theta\rangle$ witnessing that M is a (μ, θ)-limit model (that is, $\bigcup_{i<\theta} M_i = M$ and M_{i+1} is universal over M_i). Further, by Facts 1.13 and 1.14 we can require that this chain satisfy additional requirements such as M_{i+1} is a limit model over M_i . In this section we will be adding more information to this chain of models, resulting in a tower (see Definition 4.1). But first, we will describe how the towers will be used to prove the main theorem of this paper.

To prove the uniqueness of limit models we will construct a model which is simultaneously a (μ, θ_1)-limit model over some fixed model M and a (μ, θ_2)-limit model over M. Notice that, by Fact 1.13, it is enough to construct a model M^* that is simultaneously a (μ, ω)-limit model and a (μ, θ)-limit model for arbitrary θ. By Fact 1.14 we may assume that θ is a limit ordinal $< \mu^+$ such that $\theta = \mu \cdot \theta$.

So, we actually construct an array of models with $\omega + 1$ rows and the number of columns of this array will have the same cofinality as θ. See the big picture of the construction on page 22. We intend to carry out the construction **down** and **to the right** in that picture. In this array, the bottom right hand corner (M^*) will be a (μ, ω)-limit model witnessed by a chain of models as described in the first paragraph of this section. This chain will appear in the last column of the array. We will see that M^* is a (μ, θ)-limit model by examining the last (the ωth) row of the array. This last row will be an $\prec_{\mathcal{K}}$-increasing sequence of models, $\bar{M}^*$ whose length will have the same cofinality as θ. However we will not be able to guarantee that M^*_{i+1} is universal over M^*_i in this last row. Thus we need another method to conclude that M^* is a (μ, θ)-limit model. This involves attaching more information to our sequence $\bar{M}^*$. We call this accessorized sequence of models a tower (see Definition 4.1 below). Each row in our construction of the array of models will be such a tower.

Under the assumption of Galois-stability, given any sequence $\langle a_i \mid i < \theta\rangle$ of elements with $a_i \in M_{i+1} \backslash M_i$, we can identify $N_i \prec_{\mathcal{K}} M_i$ such that ga-tp(a_i/M_i) does not μ-split over N_i. Furthermore, by Assumption 2.5, we may choose this N_i such that M_i is a limit model over N_i. We abbreviate this situation by a tower $(\bar{M}, \bar{a}, \bar{N})$:

Definition 4.1 (Towers)**.** Given a well ordering $(I, <)$ of cardinality $< \mu^+$, we will denote the successor of i in the ordering I by $i+1$ when it is clear. Then, we define a *tower* to be a triple $(\bar{M}, \bar{a}, \bar{N})$ where $\bar{M} = \langle M_i \mid i < \theta \rangle$ is a $\prec_{\mathcal{K}}$-increasing sequence of limit models of cardinality μ; $\bar{a} = \langle a_i \mid i < \theta \rangle$ and $\bar{N} = \langle N_i \mid i < \theta \rangle$ satisfy $a_i \in M_{i+1} \backslash M_i$; ga-tp$(a_i/M_i)$ does not μ-split over N_i; M_i is a (μ, σ)-limit model over N_i.

Notation 4.2. We denote by $\mathcal{K}^*_{\mu,I}$ the set of towers of the form $(\bar{M}, \bar{a}, \bar{N})$ where the sequences $\bar{M}$, $\bar{a}$ and $\bar{N}$ are indexed by I. Occasionally, I will be an ordinal θ with the usual ordering, and we write $\mathcal{K}^*_{\mu,\theta}$ for this set of towers. At times, we will be considering towers based on different well orderings I and I' simultaneously. In these contexts if $i \in I \bigcap I'$, the notation $i+1$ is not necessarily well-defined so we will use the notation $\mathrm{succ}_I(i)$ for the successor of i in the ordering I. Finally when I is a sub-order of I' for any $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,I'}$ we write $(\bar{M}, \bar{a}, \bar{N}) \restriction I$ for the tower in $\mathcal{K}^*_{\mu,I}$ given by the subsequences $\langle M_i \mid i \in I \rangle$, $\langle N_i \mid i \in I \rangle$ and $\langle a_i \mid i \in I \rangle$.

In addition to having control over the last row of the array, we also need to be able to guarantee that the last column of the tower witnesses that M^* is a (μ, ω)-limit model. This will be done by prescribing the following ordering on rows of the array:

Definition 4.3. For towers $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}_{\mu,I}$ and $(\bar{M}', \bar{a}', \bar{N}') \in \mathcal{K}_{\mu,I'}$ with $I \subseteq I'$, we write $(\bar{M}, \bar{a}, \bar{N}) < (\bar{M}', \bar{a}', \bar{N}')$ if and only if for every $i \in I$, $a_i = a'_i$, $N_i = N'_i$ and M'_i is a proper universal extension of M_i.

Remark 4.4. The ordering $<$ on towers is identical to the ordering $<^c_\mu$ defined in [ShVi]. The superscript was used by Shelah and Villaveces to distinguish this ordering from others. We only use one ordering on towers, so we omit the superscripts and subscripts here.

Once we have established an ordering on towers, we can define the union of an increasing sequence of towers. Suppose that $\langle (\bar{M}, \bar{a}, \bar{N})^\gamma \in \mathcal{K}^*_{\mu,I_\gamma} \mid \gamma < \beta \rangle$ is an increasing sequence of towers such that the index set I_γ of $(\bar{M}, \bar{a}, \bar{N})^\gamma$ is a subordering of the index set $I_{\gamma'}$ for $(\bar{M}, \bar{a}, \bar{N})^{\gamma'}$ whenever $\gamma < \gamma'$. Let $I_\beta := \bigcup_{\gamma<\beta} I_\gamma$. Then denote by $(\bar{M}, \bar{a}, \bar{N})^\beta \in \mathcal{K}^*_{\mu,I_\beta}$ the union of the sequence of towers where

$$\bar{a}^\beta_i = a_i^{\min\{\gamma \mid i \in \gamma\}},$$

$$\bar{N}^\beta_i = N_i^{\min\{\gamma \mid i \in \gamma\}}$$

and

$$\bar{M}^\beta = \langle M^\beta_i \mid i \in \bigcup_{\gamma<\beta} I_\gamma \rangle$$

with

$$M_i^\beta = \bigcup_{\gamma<\beta} \bigcup_{i\in I_\gamma} M_i^\gamma.$$

Notice that Assumption 2.5.3a is needed to guarantee that $(\bar{M}, \bar{a}, \bar{N})^\beta$ is in fact a tower.

Notice that we do not assume an individual tower to be continuous. Nor do we assume that inside of a tower M_{i+1} is universal over M_i. If one considers the approach of defining an array of models row by row, then generally (even in the first order case) even if all rows are continuous and satisfy the universality property mentioned in this paragraph, it is not true that the the union of these rows will be a tower in which every model is universal over its predecesors.

For a tower $(\bar{M}, \bar{a}, \bar{N})$, it was shown in [ShVi], that even if M_{i+1} is not universal over M_i, one can conclude that $\bigcup_{i<\theta} M_i$ is a (μ,θ)-limit model provided that all types over each of the M_i are realized by a sufficient number of a_js in the tower. Unfortunately constructing such a tower meeting these along with all of our other requirements is beyond reach. However, in [Va], VanDieren showed that slightly less was needed (see Definition 4.5). Here we provide a proof since we can avoid many complications that arise in [Va] because we have at our disposal the amalgamation property.

Definition 4.5 (Relatively Full Towers)**.** Suppose that I is a well-ordered set such that there exists a cofinal sequence $\langle i_\alpha \mid \alpha < \theta\rangle$ of I of order type θ such that there are $\mu\cdot\omega$ many elements between i_α and $i_{\alpha+1}$.
Let $(\bar{M}, \bar{a}, \bar{N})$ be a tower indexed by I such that each M_i is a (μ,σ)-limit model. For each i, let $\langle M_i^\gamma \mid \gamma < \sigma\rangle$ witness that M_i is a (μ,σ)-limit model. The tower $(\bar{M}, \bar{a}, \bar{N})$ is *full relative to* $(M_i^\gamma)_{\gamma<\sigma, i\in I}$ iff for every $\gamma<\sigma$ and every $(p, M_i^\gamma) \in \mathfrak{St}(M_i)$ with $i_\alpha \leq i < i_{\alpha+1}$, there exists $j\in I$ with $i\leq j < i_{\alpha+1}$ such that $(\text{ga-tp}(a_j/M_j), N_j)$ and (p, M_i^γ) are parallel.

Fact 4.6. *Let θ be a limit ordinal $< \mu^+$ satisfying $\theta = \mu\cdot\theta$. Suppose that I is a well-ordered set as in Definition 4.5.*

Let $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,I}$ be a tower made up of (μ,σ)-limit models. If $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}_{\mu,I}$ is full relative to $(M_i^\gamma)_{i\in I,\gamma<\sigma}$, then $M := \bigcup_{i\in I} M_i$ is a (μ,θ)-limit model.

Proof. Without loss of generality we may assume that $\bar{M}$ is continuous. Let M' be a (μ,θ)-limit model over M_{i_0} witnessed by $\langle M'_\alpha \mid \alpha<\theta\rangle$. By μ-Disjoint Amalgamation, we may assume that $M'\cap M = M_{i_0}$. Since $\theta=\mu\cdot\theta$, we may also arrange things so that the universe of M'_α is $\mu\cdot\alpha$ and $\alpha\in M'_{\alpha+1}$.

We will construct an isomorphism between M and M' by induction on $\alpha<\theta$. Define an increasing and continuous sequence of $\prec_{\mathcal{K}}$-mappings $\langle h_\alpha \mid \alpha<\theta\rangle$ such that

(1) $h_\alpha : M_{i_\alpha+j} \to M'_{\alpha+1}$ for some $j<\mu\cdot\omega$
(2) $h_0 = \text{id}_{M_{i_0}}$ and
(3) $\alpha \in \text{rg}(h_{\alpha+1})$.

For $\alpha = 0$ take $h_0 = \mathrm{id}_{M_{i_0}}$. For α a limit ordinal let $h_\alpha = \bigcup_{\beta<\alpha} h_\beta$. Since $\bar{M}$ is continuous, the induction hypothesis gives us that h_α is a $\prec_{\mathcal{K}}$-mapping from M_{i_α} into M'_α allowing us to satisfy condition (1) of the construction.

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

Now consider the case when $\alpha \notin \mathrm{rg}(h_\alpha)$. Since $\langle M^\gamma_{i_\alpha+j} \mid \gamma < \sigma\rangle$ witnesses that $M_{i_\alpha+j}$ is a (μ,σ)-limit model, by Assumption 2.5, there exists $\gamma < \sigma$ such that ga-tp$(\alpha/M_{i_\alpha+j})$ does not μ-split over $M^\gamma_{i_\alpha+j}$. By our choice of $\bar{M}'$ disjoint from $\bar{M}$ outside of M_{i_0}, we know that $\alpha \notin M_{i_\alpha+j}$. Thus ga-tp$(\alpha/M_{i_\alpha+j})$ is non-algebraic. By relative fullness of $(\bar{M},\bar{a},\bar{N})$, there exists j' with $j \leq j' < i_{\alpha+1}$ such that (ga-tp$(\alpha/M_{i_\alpha+j'}), M^\gamma_{i_\alpha+j})$ is parallel to (ga-tp$(a_{i_{\alpha+1}+j'}/M_{i_{\alpha+1}+j'}), N_{i_{\alpha+1}+j'})$. In particular we have that

$$(*) \quad \text{ga-tp}(a_{i_{\alpha+1}+j'}/M_{i_\alpha+j}) = \text{ga-tp}(\alpha/M_{i_\alpha+j}).$$

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

$$(**) \quad \text{ga-tp}(h'(a_{i_{\alpha+1}+j'})/h_\alpha(M_{i_\alpha+j})) = \text{ga-tp}(\alpha/h_\alpha(M_{i_\alpha+j})).$$

Since $M'_{\alpha+2}$ is universal over $h_\alpha(M_{i_\alpha})$, we may extend h_α to a $\mathcal{K}$-mapping $h_{\alpha+1} : M_{i_{\alpha+1}+j'} \to M'_{\alpha+2}$ such that $h_{\alpha+1}(a_{i_{\alpha+1}+j'}) = \alpha$.

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

5. Uniqueness of Limit Models

We now begin the construction of our array of models and M^*. The goal will be to build an array of models with $\omega + 1$ rows so that the bottom row of the array is a relatively full tower. To get a relatively full tower at the end of the construction, we will be adding elements to the index set of towers row by row so that the $(\omega + 1)^{\mathrm{st}}$ row will have a tower indexed by I as in Definition 4.5 and will have realizations of all required parallel types. At stage n of our construction the tower that we build is indexed by I_n described here:

Notation 5.1. The index sets I_n will be defined inductively so that $\langle I_n \mid n < \omega+1\rangle$ is an increasing and continuous chain of well-ordered sets. We fix I_0 to be an index set of order type $\theta+1$ and will denote it by $\langle i_\alpha \mid \alpha \leq \theta\rangle$. We will refer to the members of I_0 by name in many stages of the construction. These indices serve as anchors for the members of the remaining index sets in the array. Next we demand that for each $n < \omega$, $\{j \in I_n \mid i_\alpha < j < i_{\alpha+1}\}$ has order type $\mu \cdot n$ such that each I_n has a supremum i_θ. An example of such $\langle I_n \mid n \leq \omega\rangle$ is $I_n = \theta \times (\mu \cdot n) \bigcup\{i_\theta\}$ ordered lexicographically, where i_θ is an element $\geq$ each $i \in \bigcup_{n<\omega} I_n$.

To prove the main theorem of the paper, we need to prove that for a fixed $M \in \mathcal{K}$ of cardinality μ any (μ, θ) and (μ, ω) limit models over M are isomorphic over M. Fix $M \in \mathcal{K}_\mu$ and θ such that $\mu \cdot \theta = \theta$. Without loss of generality, M is a limit model. We define by induction on $n \leq \omega$ a $<$-increasing and continuous sequence of towers $(\bar{M}, \bar{a}, \bar{N})$ such that

(1) $(\bar{M}, \bar{a}, \bar{N})^0$ is a tower with $M_0^0 = M$.
(2) $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu, I_n}$
(3) For every $(p, N) \in \mathfrak{St}(M_i^n)$ with $i_\alpha \leq i < i_{\alpha+1}$ there is $j \in I_{n+1}$ with $i < j < i_{\alpha+1}$ so that $(\text{ga-tp}(a_j/M_j^{n+1}), N_j^{n+1})$ and (p, N) are parallel.
(4) $M^{n+1}_{i_{\alpha+1}}$ is a (μ, μ)-limit model over $\bigcup_{j < i_{\alpha+1}} M_j^{n+1}$.

Given M, we can find a tower $(\bar{M}, \bar{a}, \bar{N})^0 \in \mathcal{K}^*_{\mu, I_0}$ with $M_0^0 = M$ because of the existence of universal extensions and because of Assumption 2.5.3b. The last pages (Page 22 onward) of this section provide a picture of this construction of an array of models, explanations for carrying out the final stage of the construction and a proof that this is sufficient to prove the main theorem. We spend most of the remainder of this section verifying that it is possible to carry out the induction step of the construction. This is a particular case of Theorem II.7.1 of [Va]. But since our context is somewhat easier, we do not encounter so many obstacles as in [Va] and we provide a different, more direct proof here:

Theorem 5.2 (Dense $<$-extension property)**.** *Given $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu, I_n}$ there exists $(\bar{M}', \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu, I_{n+1}}$ such that $(\bar{M}', \bar{a}, \bar{N}) < (\bar{M}, \bar{a}, \bar{N})$ and for each $(p, N) \in \mathfrak{St}(M_i)$ with $i_\alpha \leq i < i_{\alpha+1}$, there exists $j \in I_{n+1}$ with $i < j < i_{\alpha+1}$ such that $(\text{ga-tp}(a_j/M_j'), N_j)$ and (p, N) are parallel. Here, the M_i's are defined for $i \in I_n$ and the M_j' are defined for $j \in I_{n+1}$.*

Before we prove Theorem 5.2, we prove a slightly weaker extension property, one in which we can find an extension of the tower $(\bar{M}, \bar{a}, \bar{N})$ of the same index set:

Lemma 5.3 ($<$-extension property)**.** *Given $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu, I}$, there exists a $<$-extension $(\bar{M}', \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu, I}$ of $(\bar{M}, \bar{a}, \bar{N})$ such that for each limit i, M_i' is a (μ, μ)-limit model over $\bigcup_{j<i} M_j'$.*

Proof. Given $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu, I}$ we will define a $<$-extension $(\bar{M}', \bar{a}, \bar{N})$ by induction on $i \in I$. Notice that a straightforward induction proof is not sufficient here for if we have defined $\langle M_j \mid j \leq i \rangle$ as a tower extending $(\bar{M}, \bar{a}, \bar{N})$ restricted to $\langle j \mid j \leq i \rangle$ and are at the stage of defining M'_{i+1}, we may be faced with an impossible task: During our construction we may have inadvertently placed inside M_i' witnesses for the splitting of the type of a_{i+1} over N_{i+1}; this would prevent us from extending M_i' to M'_{i+1} so that $\text{ga-tp}(a_{i+1}/M'_{i+1})$ does not μ-split over N_{i+1}. Therefore, we will instead define approximations (M_i^+) for M_i' by induction on $i \in I$ and at each stage i of the induction we will make adjustments of the previously defined approximation M_j^+ for $j < i$. This leads us into defining M_i^+ and a directed system of $\prec_{\mathcal{K}}$-embeddings $\langle f_{j,i} \mid j < i \in I\rangle$ such that for $i \in I$, $M_i \prec_{\mathcal{K}} M_i^+$ for $j \leq i$, $f_{j,i} : M_j^+ \to M_i^+$ and $f_{j,i} \restriction M_j = \mathrm{id}_{M_j}$. We further require that M_{i+1}^+ is a limit model over $f_{i,i+1}(M_i^+)$ and ga-tp$(a_i/f_{i,i+1}(M_i^+))$ does not μ-split over N_i. When i is a limit, we choose M_i^+ to be a (μ,μ)-limit model over $\bigcup_{j<i} f_{j,i}(M_j^+)$

This construction is done by induction on $i \in I$ using the existence of non-μ-splitting extensions. Suppose that $\langle M_k^+ \mid k \leq i\rangle$ and $\langle f_{k,l} \mid k \leq l \leq i\rangle$ have been defined. We explain how to define M_{i+1}^+ and $f_{i,i+1}$. The rest of the definitions required for the $i+1^{\text{st}}$ stage are dictated by the requirement that we are forming a directed system. Let M_{i+1}^* be an limit model over both M_i^+ and M_{i+1}. Since ga-tp(a_{i+1}/M_{i+1}) does not μ-split over N_{i+1}, by Fact 2.4 there exists $f \in \mathrm{Aut}(\mathfrak{C}/M_{i+1})$ so that ga-tp$(a_{i+1}/f(M_{i+1}^*)$ does not μ-split over N_{i+1}. Take $M_{i+1}^+ := f(M_{i+1}^*)$ and $f_{i,i+1} := f \restriction M_i^+$.

(Note: The use of the directed system is crucial in getting the non-splitting. At each stage of his suggestion it is possible that the M_i' has some splitting information of a future a_j that would prevent the possibility of extending M_i' to M_j' with the required non-splitting.)

At limit stages we take direct limits so that $f_{j,i} \restriction M_j = \mathrm{id}_{M_j}$. This is possible by Subclaims II.7.10 and II.7.11 of [Va] or see Claim 2.17 of [GrVa2]. Take an extension of the direct limit that is both universal over M_i and is a (μ,μ)-limit over $\bigcup_{j<i} f_{j,i}(M_j)$ and call this M_i^+. We forfeit continuity of the tower at this point, but it will be recovered later using reduced towers.

Let $f_{j,\sup\{I\}}$ and $M'_{\sup\{I\}}$ be the direct limit of this system such that $f_{j,\sup\{I\}} \restriction M_j = \mathrm{id}_{M_j}$. We can now define $M_j' := f_{j,\sup\{I\}}(M_j^+)$ for each $j \in I$. The details of the verification that $(\bar{M}', \bar{a}, \bar{N})$ is as required are left to the reader, and can be found in [Va].

$\dashv$

We can now use the extension property for towers of the same index set from Lemma 5.3 to prove the dense extension property which allows us to grow the index set as we add elements to the models in the extension.

Proof of Theorem 5.2. Given $(\bar{M}, \bar{a}, \bar{N})^n \in \mathcal{K}^*_{\mu,I_n}$, let $(\bar{M}', \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,I_n}$ be an extension of $(\bar{M}, \bar{a}, \bar{N})$ as in Lemma 5.3 so that each $M'_{i_{\alpha+1}}$ is a (μ,μ)-limit model over $\bigcup_{j<i_{\alpha+1}} M_j'$.

For each i_α, let $\langle M_l' \mid l \in I_{n+1},\ i_\alpha + \mu\cdot n < l < i_{\alpha+1}\rangle$ witness that $M'_{i_{\alpha+1}}$ is a (μ,μ)-limit model over $\bigcup_{j<i_{\alpha+1}} M_j'$. Without loss of generality we may assume that each of these M_l' is a limit model over its predecessor.

Fix $\{(p,N)^l_{i_\alpha} \mid i_\alpha + \mu\cdot n < l < i_{\alpha+1}\}$ an enumeration of $\bigcup\{\mathfrak{S}\mathrm{t}(M_i) : i \in I_n, i_\alpha \leq i < i_{\alpha+1}\}$. By our choice of I_{n+1} and stability in μ, such an enumeration is possible. Since $M'_{\mathrm{succ}_{I_{n+1}}(l)}$ is universal over M_l', there exists a realization in $M'_{\mathrm{succ}_{I_{n+1}}}(l)$ of the non-μ-splitting extension of $p^l_{i_\alpha}$ to M'_l. Let a_l be this realization and take $N_l := N^l_{i_\alpha}$.

Notice that $(\langle M'_j \mid j \in I_{n+1}\rangle, \langle a_j \mid j \in I_{n+1}\rangle, \langle N_j \mid j \in I_{n+1}\rangle)$ provide the desired extension of $(\bar{M}, \bar{a}, \bar{N})$ in $\mathcal{K}^*_{\mu, I_{n+1}}$. ⊣

We are almost ready to carry out the complete construction. However, notice that Theorem 5.2 does not provide us with a continuous extension. Therefore the bottom (i.e. the $\omega + 1^{st}$) row of our array may not be continuous which would prevent us from applying Fact 4.6 to conclude that M^* is a (μ, θ)-limit model. So we will further require that the towers that occur in the rows of our array are all continuous. This can be guaranteed by restricting ourselves to reduced towers as in [ShVi] and [Va].

Definition 5.4. A tower $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,I}$ is said to be *reduced* provided that for every $(\bar{M}', \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,I}$ with $(\bar{M}, \bar{a}, \bar{N}) \leq (\bar{M}', \bar{a}, \bar{N})$ we have that for every $i \in I$,

$$(*)_i \quad M'_i \cap \bigcup_{j \in I} M_j = M_i.$$

If we take a $<$-increasing chain of reduced towers, the union will be reduced. The following fact appears as Theorem 3.1.14 of [ShVi]. We provide the proof for completeness.

Fact 5.5. *Let $\langle (\bar{M}, \bar{a}, \bar{N})^\gamma \in \mathcal{K}^*_{\mu, I_\gamma} \mid \gamma < \beta \rangle$ be a $<$-increasing and continuous sequence of reduced towers such that the sequence is continuous in the sense that for a limit $\gamma < \beta$, the tower $(\bar{M}, \bar{a}, \bar{N})^\gamma$ is the union of the towers $(\bar{M}, \bar{a}, \bar{N})^\zeta$ for $\zeta < \gamma$. Then the union of the sequence of towers $\langle (\bar{M}, \bar{a}, \bar{N})^\gamma \in \mathcal{K}^*_{\mu, I_\gamma} \mid \gamma < \beta \rangle$ is itself a reduced tower.*

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

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

Fact 5.6 (Density of reduced towers)**.** *There exists a reduced $<$-extension of every tower in $\mathcal{K}^*_{\mu,I}$.*

Proof. Assume for the sake of contradiction that no $<$-extension of $(\bar{M}, \bar{a}, \bar{N})$ is reduced. This allows us to construct a $\leq$-increasing and continuous sequence of towers $\langle (\bar{M}, \bar{a}, \bar{N})^\zeta \in \mathcal{K}^*_{\mu,I} \mid \zeta < \mu^+ \rangle$ such that $(\bar{M}, \bar{a}, \bar{N})^{\zeta+1}$ witnesses that $(\bar{M}, \bar{a}, \bar{N})^\zeta$ is not reduced. The construction is done inductively in the obvious way.

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

$$i(b) := \min \left\{ i \in I \mid b \in \bigcup_{\zeta < \mu^+} \bigcup_{j \leq i} M^\zeta_j \right\} \text{ and}$$

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

$\zeta(\cdot)$ can be viewed as a function from μ^+ to μ^+. Since $|I| = \mu$ and each M^{ζ}_i has cardinality μ, there exists a club $E = \{\delta < \mu^+ \mid \forall b \in \bigcup_{i\in I} M^{\delta}_i,\ \zeta(b) < \delta\}$. 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 $i \in I$ and $b \in M^{\delta+1}_i \cap \bigcup_{j\in I} M^{\delta}_j$ 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 $i(b) \leq i$, this implies 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. ⊣

By revising the proof of Lemma 5.3, we can conclude:

Lemma 5.7. *Suppose that $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,I}$ is reduced. If I_0 is an initial segment of I, then $(\bar{M}, \bar{a}, \bar{N}) \restriction I_0$ is reduced.*

Proof. Suppose that $(\bar{M}, \bar{a}, \bar{N}) \restriction I_0$ is not reduced. Let $(\bar{M}', \bar{a} \restriction I_0, \bar{N} \restriction I_0)$ and $\delta < j \in I_0$ with $b \in (M'_{\delta} \bigcap M_j)\backslash M_{\delta}$ witness this. We can apply the inductive step of Lemma 5.3 to find $(\bar{M}'', \bar{a}, \bar{N})$ an extension of $(\bar{M}, \bar{a}, \bar{N})$ such that there is a $\prec_{\mathcal{K}}$-mapping f from the models of $\bar{M}'$ into the models of $\bar{M}''$ with $f \restriction M_j = \mathrm{id}_{M_j}$. Notice that $(\bar{M}'', \bar{a}, \bar{N})$ and b, δ, j will witness that $(\bar{M}, \bar{a}, \bar{N})$ is not reduced. ⊣

Theorem 5.8 (Reduced towers are continuous)**.** *If $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,I}$ is reduced, then it is continuous, namely for each limit i in I, $M_i = \bigcup_{j<i} M_j$.*

Proof of Theorem 5.8. Suppose the claim fails for μ. Let δ be the minimal limit ordinal such that there exists an I and $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,I}$ a reduced tower discontinuous at the δ $^{\text{th}}$ element of I. We can apply Lemma 5.7 to assume without loss of generality that $I = \delta + 1$.

Fix $(\bar{M}, \bar{a}, \bar{N}) \in \mathcal{K}^*_{\mu,\delta+1}$ reduced and discontinuous at δ with $b \in M_{\delta}\backslash \bigcup_{i<\delta} M_i$.

Claim 5.9. *There exists a $<$-extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$, say $(\bar{M}', \bar{a} \restriction \delta, \bar{N} \restriction \delta) \in \mathcal{K}^*_{\mu,\delta}$ containing b.*

Notice that Claim 5.9 yields Theorem 5.8: 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.

The proof of Claim 5.9 requires several steps:

Proof of Claim 5.9. We use the minimality of δ and the density of reduced towers (Fact 5.6) to build a $<$-increasing and continuous sequence of reduced (and continuous) towers $\langle(\bar{M}, \bar{a}, \bar{N})^{\zeta} \in \mathcal{K}^*_{\mu,\delta} \mid \zeta < \delta\rangle$ such that $(\bar{M}, \bar{a}, \bar{N})^0 := (\bar{M}, \bar{a}, \bar{N}) \restriction \delta$. This gives us a δ by δ array of models. If b appears in this array, we are done. So let us suppose that ga-tp$(b/\bigcup_{i<\delta} M^i_i)$ is non-algebraic. Since $\bigcup_{i<\delta} M^i_i$ is a (μ, δ)-limit model (witnessed by the diagonal of this array), we conclude by Assumption 2.5.3b that there exists $\xi < \delta$ such that ga-tp$(b/\bigcup_{i<\delta} M_i^i)$ does not μ-split over M_ξ^ξ.

We will find a $<$-extension of $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta$ by defining an $\prec_{\mathcal{K}}$-increasing chain of models $\langle N_i' \mid i < \alpha \rangle$ and an increasing chain of $\prec_{\mathcal{K}}$-mappings $\langle h_i \mid i < \delta \rangle$. The pre-image of N_i' under an extension of $\bigcup_{i<\delta} h_i$ will form a sequence $\bar{M}'$ such that $(\bar{M}, \bar{a}, \bar{N}) \restriction \delta < (\bar{M}', \bar{a} \restriction \delta, \bar{N} \restriction \delta)$, $b \in M_{\xi+1}'$ 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) $h_i(M_i^i) \prec_{\mathcal{K}} M_i^{i+1}$
(4) ga-tp$(h_{i+1}(a_i)/N_i')$ does not μ-split over $h_i(N_i)$
(5) $b \in N_{\xi+1}'$ and
(6) for $i \leq \xi$, $N_i' = M_i^i$ with $h_i = \mathrm{id}_{M_i^i}$.

The construction: The requirements determine immediately the definition of N_i' for $i \leq \xi$. We proceed with the rest of the construction by induction on i for $\xi < i < \delta$. To guarantee continuity of the N' and h's, if i is a limit ordinal $\geq \xi$, let $N_i' = \bigcup_{j<i} N_j'$ and $h_i = \bigcup_{j<i} h_j$.

At readers request, we have provided the details of the construction for the critical steps $i = \xi + 1$ and $i = \xi + 2$, although they are not different for future successor stages. To carry out the construction for successors beyond $\xi + 2$, follow the steps for stage $\xi + 2$ described below.

Defining $h_{\xi+1}$ and $N_{\xi+1}'$: Suppose that N_ξ' and h_ξ have been defined according to the construction, specifically satisfying condition (6). By construction, ga-tp(a_ξ/M_ξ^ξ) does not μ-split over N_ξ. Since $N_\xi' = M_\xi^\xi$ and $h_\xi = id_{M_\xi}$, we get condition (4) by invariance so long as $h_{\xi+1}$ extends $h_\xi (= \mathrm{id}_{M_\xi})$. So it suffices to choose $N_{\xi+1}'$ to be some limit model over $M_{\xi+1}^{\xi+2}$ containing b. Because $M_{\xi+1}^{\xi+2}$ is universal over M_ξ^ξ we can fix a mapping $h_{\xi+1} : M_{\xi+1}^{\xi+1} \to M_{\xi+1}^{\xi+2}$ which is the identity on M_ξ^ξ.

Defining $h_{\xi+2}$ and $N_{\xi+2}'$: For this next stage of the construction we need to take a little more care to obtain condition (4). We will begin with finding the mapping $h_{\xi+2}$. First, fix $\bar{h}_{\xi+1}$ an automorphism of $\mathfrak{C}$ extending $h_{\xi+1}$. Because we started out with a δ by δ array built with towers, we have that ga-tp$(a_{\xi+1}/M_{\xi+1}^{\xi+1})$ does not μ split over $N_{\xi+1}$. Thus invariance gives us $p :=$ ga-tp$(\bar{h}_{\xi+1}(a_{\xi+1})/\bar{h}_{\xi+1}(M_{\xi+1}^{\xi+1}))$ does not μ split over $\bar{h}_{\xi+1}(N_{\xi+1})$. By definition of $h_{\xi+1}$ and condition (2), the model $N_{\xi+1}'$ is an extension of the domain of p $(= h_{\xi+1}(M_{\xi+1}^{\xi+1}))$. Thus we can apply Lemma 2.4 to obtain $g \in \mathrm{Aut}(\mathfrak{C}/\bar{h}_{\xi+1}(M_{\xi+1}^{\xi+1}))$ so that $g(\bar{h}_{\xi+1}(a_{\xi+1}))$ realizes the non-μ-splitting extension of p to $N_{\xi+1}'$ (first get an automorphism fixing $\bar{h}_{\xi+1}(N_{\xi+1})$ by

Lemma 2.4, and then use the fact that the image thus obtained realizes an extension of p). Because $M^{\xi+3}_{\xi+2}$ is universal over $M^{\xi+2}_{\xi+1}$ which contains $\bar{h}_{\xi+1}(M^{\xi+1}_{\xi+1})$, we can assume that the image of $(g\circ\bar{h}_{\xi+1})\restriction M^{\xi+2}_{\xi+2}$ is in $M^{\xi+3}_{\xi+2}$. Define $h_{\xi+2}$ to be $(g\circ\bar{h}_{\xi+1})\restriction M^{\xi+2}_{\xi+2}$. All that is required now is to choose $N'_{\xi+2}$ so that it is a limit model over $N'_{\xi+1}$ containing $M^{\xi+3}_{\xi+2}$.

We now argue that the construction of the sequences $\langle N'_i \mid i<\delta\rangle$ and $\langle h_i \mid i<\delta\rangle$ is enough to find a $<$-extension, $(\bar{M}',\bar{a},\bar{N})\in\mathcal{K}^*_{\mu,\delta}$, of $(\bar{M},\bar{a},\bar{N})\restriction\delta$ 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 5.10. *There exists $h\in\mathrm{Aut}(\mathfrak{C})$ extending h_δ such that $h(b)=b$.*

Proof of Claim 5.10. From our choice of ξ, we know that ga-tp$(b/\bigcup_{i<\delta} M^i_i)$ does not μ-split over M^ξ_ξ. Extend h_δ to an automorphism h^* of $\mathfrak{C}$. We will show:

Subclaim 5.11. ga-tp$(b/h^*(\bigcup_{i<\delta} M^i_i),\mathfrak{C})=$ ga-tp$(h^*(b)/h^*(\bigcup_{i<\delta} M^i_i),\mathfrak{C})$.

Proof of Subclaim 5.11. We will use non-splitting to derive the subclaim. Following notation used in the definition of splitting, let $N^1=\bigcup_{i<\delta} M^i_i$, $N^2=h^*(\bigcup_{i<\delta} M^i_i)$ and $p=$ ga-tp(b/N^1). Since $h_i(M^i_i)\prec_{\mathcal{K}} M^{i+1}_i\prec_{\mathcal{K}}\bigcup_{j<\delta} M^j_j$, we have $N^2\prec_{\mathcal{K}} N^1$ Since p does not μ-split over M^ξ_ξ, we have that $p\restriction N^2=h^*(p\restriction N^1)$. In other words,

$$\text{ga-tp}(b/h^*(\bigcup_{i<\delta} M^i_i),\mathfrak{C})=\text{ga-tp}(h^*(b)/h^*(\bigcup_{i<\delta} M^i_i),\mathfrak{C}),$$

as desired. ⊣

From Subclaim 5.11, we can find an automorphism f of $\mathfrak{C}$ such that $f(h^*(b))=b$ and $f\restriction h^*(\bigcup_{i<\delta} M^i_i)=\mathrm{id}_{h^*(\bigcup_{i<\delta} M^i_i)}$. Notice that $h:=f\circ h^*$ satisfies the conditions of Claim 5.10: $h(b)=f(h^*(b))=b$ and $h\supset h_\delta$ as $f\restriction h^*(\bigcup_{i<\delta} M^i_i)=\mathrm{id}_{h^*(\bigcup_{i<\delta} M^i_i)}$.

⊣

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 5.12. *$(\bar{M}',\bar{a},\bar{N})$ is a $<$-extension of $(\bar{M},\bar{a},\bar{N})$ such that $b\in M'_{\xi+1}$.*

Proof of Claim 5.12. By construction $b\in N'_{\xi+1}$. Since $h(b)=b$, this implies $b\in M'_{\xi+1}$. To verify that we have a $\leq$-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, 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. ⊣

⊣

⊣

Corollary 5.13. *In Theorem 5.2, we can choose* $(\bar{M}, \bar{a}, \bar{N})^{\alpha+1}$ *to be reduced, and hence continuous.*

The construction:

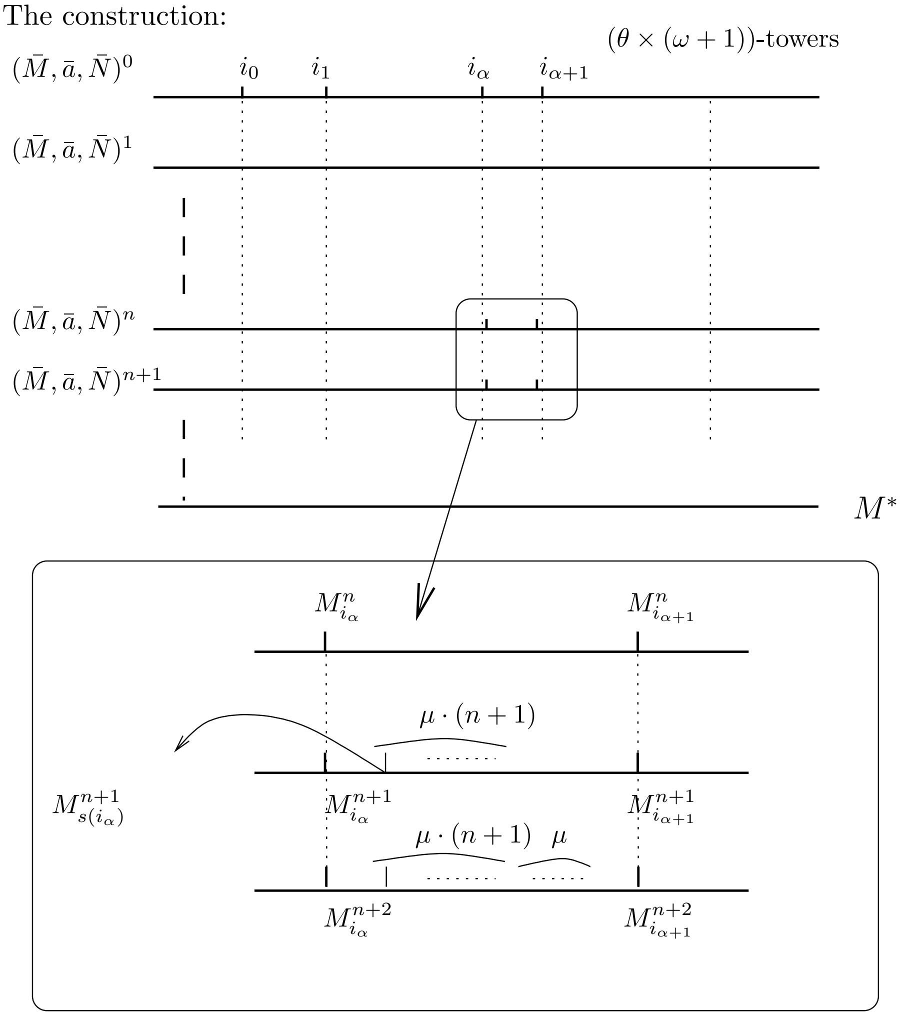

Corollary 5.13 tells us that the construction of our array of models as an increasing sequence of towers is possible in successor cases. In the limit case,

let $I_\omega = \bigcup_{m<\omega} I_m$, and simply define, $(\bar{M}, \bar{a}, \bar{N})^\omega \in \mathcal{K}^*_{\mu,\mathbf{I}_\omega}$ to be the union of the towers $(\bar{M}, \bar{a}, \bar{N})^n$.

To see that the construction suffices for what we need, first notice that the last column of the array, $\langle M^n_{i_\theta} \mid n < \omega\rangle$, witnesses that M^* is a (μ, ω)-limit model. In light of Fact 4.6 we need only verify that the last row of the array is a relatively full tower of cofinality θ.

Claim 5.14. $(\bar{M}, \bar{a}, \bar{N})^\omega$ *is full relative to* $(M^n_i)_{n<\omega, i\in I_\omega}$.

Proof. Given i with $i_\alpha \leq i < i_{\alpha+1}$, let (p, M^n_i) be some strong type in $\mathfrak{St}(M^\omega_i)$. Notice that by monotonicity of non-splitting $(p \restriction M^{n+1}_i, M^n_i) \in \mathfrak{St}(M^{n+1}_i)$. By construction there is a $j \in I_{n+1}$ with $i < j < i_{\alpha+1}$ such that $(\text{ga-tp}(a_j/M^{n+2}_j), N^{n+2}_j)$ is parallel to $p \restriction M^{n+1}_i$. We will show that $(\text{ga-tp}(a_j/M^\omega_j), N^\omega_j)$ is parallel to (p, N).

First notice that $\text{ga-tp}(a_j/M^\omega_j)$ does not μ-split over $N^\omega_j = N^{n+2}_j$ because $(\bar{M}, \bar{a}, \bar{N})^\omega$ is a tower. Since $(\text{ga-tp}(a_j/M^{n+2}_j), N^{n+2}_j)$ is parallel to $(p \restriction M^{n+1}_i, M^n_i)$ there is $q \in \text{ga-S}(M^\omega_j)$ such that q extends both $p \restriction M^{n+1}_i$ and $\text{ga-tp}(a_j/M^{n+2}_j)$. By two separate applications of the uniqueness of non-μ-splitting extensions we know that $q \restriction M^\omega_i = p$ and $q = \text{ga-tp}(a_j/M^\omega_j)$. To see that (q, N^ω_j) is parallel to (p, M^n_i), let M' be an extension of M^ω_j of cardinality μ. Since $(p \restriction M^{n+1}_i, M^n_i)$ and $(q \restriction M^{n+2}_j, N^{n+2}_j)$ are parallel, there is $q' \in \text{ga-S}(M')$ extending both $p \restriction M^{n+1}_i$ and $q \restriction M^{n+2}_j$ and not μ-splitting over both M^n_i and N^{n+2}_j. By the uniqueness of non-μ-splitting extensions, we have that q' is also an extension of q and p. Thus q' witnesses that (q, N^ω_j) and (p, M^n_i) are parallel. $\dashv$

This completes the proof of Theorem 1.17.

References

[Ba] J. Baldwin. **Categoricity in Abstract Elementary Classes**, 246 pages, in preparation. Check http://www.math.uic.edu/ jbaldwin/pub/AEClec.pdf

[Gr1] Rami Grossberg, **Classification theory for Abstract Elementary Classes**, A book in preparation.

[Gr2] Rami Grossberg, Classification theory for non-elementary classes, *Contemporary Mathematics*, **302**, (2002) AMS, pp. 165–204.

[GrLe] Rami Grossberg and Olivier Lessmann. Shelah's Stability Spectrum and Homogeneity Spectrum, *Archive for mathematical Logic*, **41**, (2002) 1, 1-31.

[GrVa0] Rami Grossberg and Monica VanDieren. Shelah's Categoricity Conjecture from a successor for Tame Abstract Elementary Classes, *Journal of Symbolic Logic.* **71**, (2006) 2, 553–568.

[GrVa1] Rami Grossberg and Monica VanDieren, Galois-stability for Tame Abstract Elementary Classes, *Journal of Mathematical Logic.* **6**, No. 1 (2006) 25–49.

[GrVa2] Rami Grossberg and Monica VanDieren, Categoricity from one successor cardinal in Tame Abstract Elementary Classes. (21 pages) *Journal of Mathematical Logic*, accepted.

[HyKe] Tapani Hyttinen and Meeri Kesälä. Independence in Finitary AECs.

[Jo] Bjarni Jónsson, Homogeneous universal systems, *Math. Scand.*, **8**, 1960, 137–142.

[KoSh] Oren Kolman and Saharon Shelah, Categoricity of Theories in $L_{\kappa,\omega}$ when κ is a measurable cardinal. Part I. *Fundamentae Mathematicae*, **151**, 209-240, 1996.

[Po] Poizat, Poizat's book.

[Sh e] Saharon Shelah. Classification Theory. Revised edition. North Holland, 1990.

[Sh 3] Saharon Shelah. Shelah paper 3.

[Sh 87b] Saharon Shelah. Classification theory for nonelementary classes. I. The number of uncountable models of $\psi \in L_{\omega_1,\omega}$. Part B, *Israel Journal of Mathematics*, **46**:241–273, 1983.

[Sh 394] Saharon Shelah, Categoricity of abstract classes with amalgamation, *Annals of Pure and Applied Logic*, 1999.

[Sh 576] Saharon Shelah. Categoricity of an abstract elementary class in two successive cardinals. *Israel Journal of Mathematics*,**126**, (2001), 29–128.

[Sh 600] Saharon Shelah. Categoricity in abstract elementary classes: going up inductive step. Preprint. arXiv:math.LO/0011215 (November 2000 version), 82 pages.

[Sh 705] Saharon Shelah. Toward classification theory of good λ frames and abstract elementary classes.

[ShVi] Saharon Shelah and Andrés Villaveces. Categoricity in abstract elementary classes with no maximal models. *Annals of Pure and Applied Logic*, **97**, 1999, no.1–3, 1–25.

[Va] Monica VanDieren. Categoricity in Abstract Elementary Classes with No Maximal Models, *Annals of Pure and Applied Logic*, **141**, 2006, pages 108–147.

[Za] Pedro Zambrano. Cats, docilidad, y la propiedad de amalgamación disyunta. Submitted.

E-mail address, Rami Grossberg: `rami@cmu.edu`

Department of Mathematics, Carnegie Mellon University, Pittsburgh PA 15213

E-mail address, Monica VanDieren: `vandieren@rmu.edu`

Department of Mathematics, Robert Morris University, Moon Township PA 15108

E-mail address, Andrés Villaveces: `avillavecesn@unal.edu.co`

Departamento de Matemáticas, Universidad Nacional de Colombia, Bogotá, Colombia