The amalgamation spectrum

John T. Baldwin[*]
Department of Mathematics, Statistics and Computer Science
University of Illinois at Chicago
and Alexei Kolesnikov
Towson University
Department of Mathematics
Towson Maryland
and Saharon Shelah[†]
Department of Mathematics
Hebrew University of Jerusalem
Rutgers University

March 15, 2008

Abstract

We study when classes can have the disjoint amalgamation property for a proper initial segment of cardinals.

Theorem A. For every natural number k, there is a class $\boldsymbol{K}_k$ defined by a sentence in $L_{\omega_1,\omega}$ that has no models of cardinality greater than $\beth_{k+1}$, but $\boldsymbol{K}_k$ has the disjoint amalgamation property on models of cardinality less than or equal to $\aleph_{k-3}$ and has models of cardinality $\aleph_{k-1}$.

More strongly, we can have disjoint amalgamation up to $\aleph_\alpha$ for $\alpha < \omega_1$, but have a bound on size of models.

Theorem B. For every countable ordinal α, there is a class $\boldsymbol{K}_\alpha$ defined by a sentence in $L_{\omega_1,\omega}$ that has no models of cardinality greater than $\beth_{\omega_1}$, but $\boldsymbol{K}$ does have the disjoint amalgamation property on models of cardinality less than or equal to $\aleph_\alpha$.

Finally we show that, consistently with ZFC, we can extend the $\aleph_\alpha$ to $\beth_\alpha$ in the second theorem. Similar results hold for arbitrary ordinal α with $|\alpha| = \kappa$ and $L_{\kappa^+,\omega}$.

A sentence ϕ of $L_{\omega_1,\omega}$ is said to *characterize* μ, if ϕ has a model of cardinality μ but no model in any larger cardinalities. There are a number of results [Mor65, Hjo07, Kni77, LS93, Sou] giving various examples that show (the one

due to Hjorth is the most general) any cardinal κ below $\beth_{\omega_1}$ can be characterized by a sentence ϕ_κ of $L_{\omega_1,\omega}$. We consider the effect of requiring that the models of ϕ satisfy the disjoint amalgamation property (up to some point).

For strong properties like categoricity one can show great regularities in the eventual spectrum [She99, She83a, She83b]. These results deduce eventual categoricity from categoricity in a large enough cardinality or from a long enough initial sequence of categoricity cardinals (possibly with additional model-theoretic or set-theoretic hypotheses). These results become even stronger [GV06, Les] as eventual categoricity is deduced from categoricity in $\mathrm{LS}(\boldsymbol{K})^{++}$ or even for countable languages from categoricity in $\mathrm{LS}(\boldsymbol{K})^+$ assuming tameness, the amalgamation property, and arbitrarily large models.

In fact, Shelah's argument for eventual categoricity for categoricity below $\aleph_\omega$ (assuming $2^{\aleph_n} < 2^{\aleph_{n+1}}$ for $n < \omega$) proceeds by showing one has strong amalgamation conditions for finite independent systems of countable models. These conditions involve both existence and uniqueness; but the uniqueness is not needed for constructing arbitrarily large models. In the 1980's, Grossberg [Gro02] raised the issue of studying the amalgamation spectrum.

Definition 0.1 *Let $(\boldsymbol{K}, \prec_{\boldsymbol{K}})$ be an abstract elementary class. The amalgamation spectrum of $\boldsymbol{K}$ is the class of cardinals κ such that if there are $M \prec_{\boldsymbol{K}} N_1$ and $M \prec_{\boldsymbol{K}} N_2$ with all three in $\boldsymbol{K}$ with cardinality $\kappa \geq \mathrm{LS}(\boldsymbol{K})$, then there is a model N_3 necessarily of cardinality κ into which both N_1 and N_2 can be strongly embedded over M.*

In particular, Grossberg [Gro02] asked whether $\beth_{\omega_1}$ is the 'Hanf number for amalgamation' for $L_{\omega_1,\omega}$. That is, does every sentence of $L_{\omega_1,\omega}$ that satisfies the amalgamation property in some cardinal above $\beth_{\omega_1}$ satisfy the amalgamation property eventually. We do not answer that question. Our work is an approximation to saying that this conjecture is optimal for the notion of *disjoint* amalgamation. The aim is to find for each $\alpha < \omega_1$ a sentence ϕ_α that, provably in ZFC, has disjoint amalgamation up to $\beth_\alpha$ but does not have arbitrarily large models. Our examples show this result with the added hypothesis of GCH and in Section 3, we show the result is consistent with many other choices for the function defining cardinal exponentiation.

Our strategy in Section 1 is to allow generalized amalgamation only for systems of less than k models (for appropriate k). In Section 2 we work with $L_{\kappa^+,\omega}$ where $\kappa = \aleph_\delta$ and for each $\alpha < \kappa^+$ construct an example to guarantee amalgamation only up to $\aleph_{\delta+\alpha}$. This section introduces a new idea – the amalgamation of certain ranked systems of models. In each case, the constructed class has a bounded number of models.

Let us note one easy example.

Example 0.2 Let τ contain infinitely many unary predicates P_n and one binary predicate E. Define a first order theory T so that $P_{n+1}(x) \to P_n(x)$, E is an equivalence relation with two classes, which are each represented by exactly one point in $P_n - P_{n+1}$ for each n. Now omit the type of two inequivalent points

that satisfy all the P_i. This gives a sentence of $L_{\omega_1,\omega}$ that is categorical and satisfies amalgamation in all uncountable powers but fails amalgamation in $\aleph_0$.

The first two sections of this paper represent refinements of the same construction and we gradually develop the machinery for stronger results. In the first section we guarantee amalgamation (and existence) up to κ^{+k} (the kth successor of κ) for a sentence in $L_{\kappa^+,\omega}$.

We note several not entirely standard notational conventions. We write $|M|$ to denote the universe of the model M (and $\|M\|$ for its cardinality) where emphasis is needed. We use $\subset$ for *proper* subset and $\subseteq$ when the subset may be the larger set. The symbol $[M]^m$ denotes the set of m-element subsets of M.

1   Amalgamation can first fail at $\kappa^{+(k-2)}$.

For any natural number k and any cardinal κ, we construct a sentence of $L_{\kappa^+,\omega}$ that satisfies disjoint amalgamation up to $\kappa^{+(k-3)}$ (the $(k-3)$-rd successor of κ) but no further; it has no models with cardinality greater than $\beth_k(\kappa)$. When $\kappa = \aleph_0$, we get result stated in the abstract. We construct the class by force to have k-disjoint amalgamation on models of size less than κ. Then we use the relevant aspects of the 'excellence technology' to show this implies 2-amalgamation on $\kappa^{+(k-3)}$.

Definition 1.1 *Fix a natural number k and a cardinal κ.*

1. *Let τ contain n-ary predicates $P_{n;\alpha}$, for each $n \le k$ and $\alpha < \kappa$.*

2. *Let $\boldsymbol{K}_k$ be the class of τ-structures (including the empty structure) such that:*

    (a) *for each $n \le k$, the $P_{n;\alpha}^M$ partition the n-element subsets of M;*

    (b) *there is no sequence of $k+1$ elements of M that are indiscernible for quantifier-free τ-formulas.*

Throughout this paper *indiscernible* means indiscernible for quantifier free formulas. Definition 1.1 implies that the predicates $P_{n;\alpha}^M$ are actually predicates of sets; i.e. they are symmetric and hold only of sequences of distinct points.

Fact 1.2 *For each k, the family $\boldsymbol{K}_k$ is defined by a universal first order theory and the omission of certain family of types. In particular $(\boldsymbol{K}_k, \prec_{\boldsymbol{K}})$ is an abstract elementary class where $\prec_{\boldsymbol{K}}$ denotes the substructure relation.*

Definition 1.3 *A set of $\boldsymbol{K}$-structures $\overline{N} = \langle N_u : u \subset k \rangle$ is a $(< \lambda, k)$-system for $\boldsymbol{K}$ if for $u, v \subset k$:*

1. *each $N_u \in \boldsymbol{K}$ and $\|N_u\| < \lambda$;*

2. *if $u \subset v$ then $N_u \subset N_v$;*

3. $N_u \cap N_v = N_{u \cap v}$.

Because we are constructing classes that are closed under substructure in a relational language we can always amalgamate without introducing new points. That is,

Definition 1.4 *We say that $\boldsymbol{K}$ has* direct $(< \lambda, k)$-amalgamation *if*

1. *$k = 0$ and there is $M \in \boldsymbol{K}$ with $\|M\| = \mu$ for all $\mu < \lambda$.*

2. *$k = 1$ and for all $\mu < \lambda$, each $M \in \boldsymbol{K}$ with $\|M\| = \mu$ has a proper extension.*

3. *$k \geq 2$ and for any $(< \lambda, k)$-system $\overline{N}$ there is a model M with universe $\bigcup_{u \subset k} |N_u|$ such that for every $u \subset k$, N_u is a substructure of M.*

We can replace $< \lambda$ by λ with the obvious modification. Note that when $k = 2$, the direct amalgamation in the previous definition implies what is normally called the *disjoint amalgamation property*. Note that the (λ, k)-amalgamation property holds trivially in λ if there are no models in λ. We construct classes $\boldsymbol{K}$ that have $(\lambda, \leq k)$-amalgamation (note $\leq$) on an initial segment of cardinals but do not have arbitrarily large models.

Definition 1.5   *1. A* special (λ, k)-system $(\overline{N}, \boldsymbol{a})$ *for $\boldsymbol{K}$ is a (λ, k)-system with a* special *sequence of elements $\boldsymbol{a} = \{a_\ell : \ell < k\}$ such that:*

   (a) for each $u \subset k$, $|N_u| = |N_\emptyset| \cup \{a_\ell : \ell \in u\}$ and

   (b) $\|N_\emptyset\| = \lambda$.

2. *We say that $\boldsymbol{K}$ has* special (λ, k)-amalgamation *(or the* special (λ, k)-existence property) *if $k < 2$ and $\boldsymbol{K}$ has direct (λ, k)-amalgamation or $k \geq 2$ and if any special (λ, k)-system can be directly amalgamated.*

In this context (of a universal theory in a relational language), it suffices to study special amalgamations.

Lemma 1.6 *Let $n \geq 2$. If $\boldsymbol{K} \subseteq \boldsymbol{K}_k$ is closed under increasing unions and has special (λ, n)-amalgamation (with respect to substructure as the notion of strong submodel) then it has direct (λ, n)-amalgamation.*

Proof. Fix n and λ. We prove the statement by induction on the number $m \leq n$ of models among $\{N_{\{\ell\}} \mid \ell < n\}$ that are *not* of the form $N_\emptyset \cup \{a_\ell\}$, i.e., are not a one-point extension of the base model. The base case $m = 0$ is the special amalgamation. Suppose now that we are able to amalgamate any (λ, n)-system in which m models $N_{\{\ell\}}$ are arbitrary extensions of $N_\emptyset$, and the rest are one-point extensions.

Let $\overline{N}$ be a (λ, n)-system where the models $\{N_{\{\ell\}} \mid m < \ell < n\}$ are one-point extensions of $N_\emptyset$, and the remaining $m + 1$ are any extensions of cardinality at

most λ. Enumerate the set $|N_{\{m\}}| - |N_\emptyset|$ as $\{b_i \mid i < \lambda\}$. For $u \subset n$, if u contains m, let N_u^i be the substructure of N_u with universe $|N_u| - \{b_j \mid j \geq i\}$ (that is, in each N_u we replace $N_{\{m\}}$ with $N_\emptyset \cup \{b_j \mid j < i\}$).

We build the amalgam N in stages, by induction on $i < \lambda$. Let N^1 be an amalgam of the special (λ, n)-system $\{N_u^1 \mid u \subset n, m \in u\} \cup \{N_u \mid u \subset n, m \notin u\}$. Note that the amalgam is over $N_{\{m\}}^0 = N_\emptyset$.

Having constructed N^i, let N^{i+1} be an amalgam of the special (λ, n)-system

$$\{N_u^{i+1} \mid u \subset n, m \in u\} \cup \{N^i \upharpoonright |N_u| \cup |N_{\{m\}}^i| \mid u \subset n, m \notin u\}.$$

Note that the amalgam is over $N_{\{m\}}^i$. Taking unions at limits, we are done. $\square_{1.6}$

So to establish amalgamation, we need only establish special amalgamation. We first show both the amalgamation and the existence of special systems for models of cardinality less than κ.

Lemma 1.7   1. *The class $\boldsymbol{K}_k$ has special $(< \kappa, s)$-amalgamation for all $s \leq k$.*

2. *For each $s \leq k$, and each $\mu < \kappa$, there is a special (μ, s)-system.*

Proof. We first prove both statements (1) and (2) for $s \leq 1$ by constructing a sequence of models in $\boldsymbol{K}$, $\{M_\beta \mid |M_\beta| = \beta,\ \beta < \kappa\}$, such that $M_\beta \prec M_{\beta+1}$. For $s = 0$, let M_0 be the empty structure. Given M_β, for each tuple $\mathbf{b} \in M_\beta$ and each $t = \lg(\mathbf{b})$, let $P_{t+1,\beta}(\mathbf{b}, \beta)$ hold to construct $M_{\beta+1}$. If β is a limit ordinal, let $M_\beta := \bigcup_{\alpha < \beta} M_\alpha$. By construction, all elements in M_β have different 1-types. (Taking the union $\bigcup_{\beta < \kappa} M_\beta$, we also have shown $\boldsymbol{K}_k$ satisfies $(\kappa, 0)$-amalgamation.)

Now we show statement (1) for $2 \leq s \leq k$. Let $\overline{M} \subset \boldsymbol{K}_k$ be a special $(< \kappa, s)$-system. The crucial point is that since each model in $\overline{M}$ has size less than κ, and the predicates $P_{n,\alpha}$ partition the universe, there are strictly fewer than κ predicates $P_{n,\alpha}$ that are non-empty in the models M_u, $u \subset s$. Let $\gamma = \gamma(\overline{M}) < \kappa$ be such that $P_{s,\gamma}$ is empty in each M_u, $u \subset s$.

Let $|M_s| := \bigcup \overline{M}$. For each $n < s$ and each $\mathbf{b} \in^n |M_s|$, we necessarily have $\mathbf{b} \in |M_u|$ for some u and so the truth of $P_{n,\gamma}(\mathbf{b})$ has been determined. If $n \geq s$, then for $\mathbf{b} \in {}^n|M_s|$, either $\mathbf{b} \in |M_u|$ for some u and so the truth of $P_{s,\alpha}(\mathbf{b})$ has been determined, or $\mathbf{b}$ contains $\boldsymbol{a} := \{a_\ell : \ell < s\}$. In the latter case, we define $P_{s,\gamma}(\mathbf{b})$ for some γ such that $P_{s,\gamma}$ has not been used.

Now we must show that M_s does not have a sequence of $k+1$ indiscernibles. But this is straightforward. Note that any candidate sequence must contain $\boldsymbol{a} = \{a_\ell : \ell < s\}$ or else the sequence is contained in some M_u and so cannot be indiscernible since $M_u \in \boldsymbol{K}_k$. Thus, we have $P_{s,\gamma(\overline{M})}^{M_s}(\boldsymbol{a})$. But any other s-element sequence from M_s is contained in some M_u (by the definition of a special system) and so satisfies $P_{s,\alpha'}^{M_u}$ where $P_{s,\alpha'}^{M_u} \neq \emptyset$ so $\alpha(\overline{M}) \neq \alpha'$ and no indiscernible sequence can contain $\boldsymbol{a}$.

Finally, we show statement (2) for $2 \leq s \leq k$. This result follows by induction using (1) from: if special $(< \kappa, \leq s)$-amalgamation holds there is a special

$(< \kappa, s + 1)$-system. And this is easy to see. For any $\mu < \kappa$, since $(\mu, 1)$-amalgamation holds, there are in fact $s + 1$ one point extensions of a given model N of size μ. Now applying $(< \kappa, \leq s)$-amalgamation, we construct the N_u for $u \subset s$. $\qquad \square_{1.7}$

The argument in the last paragraph of the above proof can be used to show the following corollary.

Corollary 1.8 *For $s \geq 1$, if $\boldsymbol{K}$ has $(< \lambda, \leq s)$-amalgamation, then for all $n \leq s + 1$ there are special $(< \lambda, n)$-systems.*

We have disjoint amalgamation below κ; we want to show this property holds in larger cardinals. We turn to the techniques of [She83a, She83b], expounded in Part IV of [Bal]. The essential point is to trade $(s + 1)$-amalgamation in λ for s-amalgamation in λ^+.

For this induction step we need the notion of a filtration of a (λ, n)-system. In the current situation, we could give a somewhat simpler description using special systems. But little would be gained and we prefer to conform with the more general formulation of e.g. [She83b, Bal]. We add countably many points at each step; some may be from the leaves and some from the heart. Recall that we write $||M||$ for the cardinality of the universe of a structure M.

Definition 1.9 *Suppose $\mathcal{S} = \langle M_s : s \in \mathcal{P}^-(n) \rangle$ is a (λ, n)-system. A filtration of $\mathcal{S}$ is a system $\mathcal{S}^\alpha = \langle M_s^\alpha : s \in \mathcal{P}^-(n), \alpha < \lambda \rangle$ such that:*

1. *each $||M_s^\alpha|| = \alpha^* = |\alpha| + \aleph_0$;*

2. *for each s in $\mathcal{P}^-(n)$, $\{M_s^\alpha : \alpha < \lambda\}$ is a filtration of M_s;*

3. *for each α, $\mathcal{S}^\alpha$ is an (α^*, n)-system.*

Claim 1.10 *For any $s < \omega$, if $\boldsymbol{K}_k$ has the $(\lambda, \leq s+1)$ amalgamation property, then it has the $(\lambda^+, \leq s)$-amalgamation property.*

Proof. We first note that the case $s = 0$ is easy. We get a model in λ^+ by taking the union of increasing chain of smaller models by $(\lambda, 1)$-amalgamation.

Now let $s \geq 1$ and let $\mathcal{S} = \overline{M}$ be a (λ^+, s)-system. Choose a filtration $\mathcal{S}^\alpha$ of $\mathcal{S}$. We can further choose N^α, Q^α for $\alpha < \lambda$ such that:

1. $||N^\alpha|| = |\alpha| + \lambda$;

2. $N^\alpha \cap M = M_\emptyset^\alpha$;

3. $Q^\alpha \in \boldsymbol{K}_k$ has universe $N^\alpha \cup M_s^{\alpha+1}$.

We construct N^α by (λ, s)-amalgamation and then Q^α by $(\lambda, s+1)$-amalgamation.
$\qquad \square_{1.10}$

We can show inductively that (λ^+, s)-system exist. Filter a $(\lambda^+, s - 1)$-system as in the last paragraph of the proof of Lemma 1.8; choose an extension of N_0 and then iteratively apply $(\lambda^+, s + 1)$-amalgamation.

Theorem 1.11   *1. K_k has no models of cardinality greater than $\beth_k(\kappa)$.*

2. *K_k has the disjoint amalgamation property on models of cardinality less than or equal to $\kappa^{+(k-3)}$.*

3. *K_k has models of cardinality $\kappa^{+(k-1)}$.*

4. *there is a cardinal μ strictly less than $\beth_k(\kappa)$ where disjoint amalgamation fails.*

Proof. Statement 1 is immediate from the Erdoš–Rado theorem. If we partition the $(k+1)$-element subsets of a set of cardinality $(\beth_k(\kappa))^+$ into κ sets determined by the $P_{n;\alpha}$, then there is a set of cardinality κ^+ ($k+1$ would be enough!) homogeneous for the partition.

Lemmas 1.6 and 1.7 give us direct $(< \kappa, k)$ amalgamation. Statement 2 then follows from Claim 1.10. Thus, we have K_k satisfies direct amalgamation property for $(\kappa^{+m}, k-1-m)$-systems for $2 \leq m \leq k-3$.

Finally, for any λ, direct amalgamation of $(\lambda, \leq 2)$-systems implies every model of cardinality λ^+ extends to a model with cardinality λ^{++}; see Lemma 1.8. So we have statement 3. But this extension property implies that if the disjoint amalgamation property holds for each $\mu < \beth_k(\kappa)$, then there is a model of cardinality $(\beth_k(\kappa))^+$, contradicting part 1. Thus statement 4 holds. $\square_{1.11}$

Note that the arguments in both [She83a, She83b] and [Bal] are carried out in the context of the models of a complete sentence of $L_{\omega_1,\omega}$ (atomic models of a first order theory). The sentence here is far from complete. We use essentially that our class is a universal theory with omitting types; there is no way to guarantee that all consistent types are realized. This does not affect the amalgamation arguments in the proof of Theorem 1.11.

The example is an abstract elementary class with $\prec_K$ taken as τ-substructure. Note that we specifically constructed the class to have a bound on the number of models. There is no hope that assuming the class has arbitrarily large models yields a theorem which implies disjoint amalgamation on many cardinals. As, we can take the disjoint union of the example here with 'a completely well-behaved' class: the result has arbitrarily large models but the same disjoint amalgamation spectrum as the example here.

2   Disjoint Amalgamation can hold at $\aleph_\alpha$ but fail

We improve the result of Section 1 to show that for every $\alpha < \omega_1$ there is a sentence of $L_{\omega_1,\omega}$ that has disjoint amalgamation up to $\aleph_\alpha$ and no model of cardinality greater than $\beth_{\omega_1}$. We can get the analogous result for $L_{\kappa^+,\omega}$ beginning at $\kappa = \aleph_\delta$ with no real change in the proof so we do the more general case. Thus for each κ and α with $\kappa \leq \alpha < \kappa^+$, we are defining below a class $\mathcal{M}_{\kappa,\alpha}$ of finite models and a class of models $K_{\kappa,\alpha}$. Since the dependence on κ and α is completely uniform, we will just write $\mathcal{M}$ and K.

Notation 2.1 *Fix for this section a cardinal $\kappa = \aleph_\delta$ and an ordinal α with $\kappa \leq \alpha < \kappa^+$. Let $\tau = \tau_{\kappa,\alpha}$ contain unary predicates $P_{1;\gamma,\alpha+2}$ with $\gamma \leq \kappa$ and n-ary relation symbols $P_{n;\gamma,\beta}$ for $2 \leq n < \omega$, $\gamma < \kappa$, and $\beta \leq \alpha + 1$. As in Section 1, the τ-structures in our class are such that $[M]^n$ is partitioned by the $P_{n;\gamma,\beta}$.*

The role of the parameter β in $P_{n;\gamma,\beta}$ is to provide a "rank" for finite sets of indiscernibles (every singleton has the maximal rank by our definition). The role of γ is to make sure that there are κ predicates of the each rank.

Definition 2.2 *We define the set* ID *as the set of all quantifier-free types of finite indiscernible sequences of decreasing rank. Thus each member of* ID *is a collection p of formulas $P_{1;\zeta_1,\alpha+2}(x_1), P_{2;\zeta_2,\alpha_2}(x_1, x_2) \ldots P_{k,\zeta_k,\alpha_k}(x_1, \ldots x_k)$, where $\alpha + 2 > \alpha_2 > \cdots > \alpha_k$ and each ζ_i is less than κ.*

Definition 2.3 *The* rank *of a finite indiscernible sequence $\boldsymbol{a}$ is the ordinal γ such that $P_{n,\zeta,\gamma}(\boldsymbol{a})$ for $n = \lg(\boldsymbol{a})$ and some $\zeta < \kappa$.*

Note that since for every finite indiscernible sequence $\boldsymbol{a}$ the rank of increasing segments of $\boldsymbol{a}$ is decreasing, then for some $\lambda < \beth_{(2^\kappa)^+}$ the class $\boldsymbol{K}$ does not have models of size greater than or equal to λ. Otherwise, we would be able to find a model containing an infinite indiscernible sequence by a standard argument. But then we would have a decreasing sequence of ordinals.

We have guaranteed that all 1-tuples have the maximal rank $\alpha + 2$ (since the third index on a unary predicate is $\alpha + 2$).

Notation 2.4   *1. Let $(\overline{N}, \boldsymbol{a})$ be a special system of models with $\boldsymbol{a} = \langle a_0, \ldots, a_{k-1} \rangle$. Let $a_u = \langle a_i \mid i \in u \rangle$. We say that sequence $\boldsymbol{a}$ is* formally indiscernible *if for every $u, v \subset k$ if $|u| = |v|$, then $\mathrm{tp}_{qf}(a_u, N_u) = \mathrm{tp}_{qf}(a_v, N_v)$.*

    2. By the rank *of a formally indiscernible sequence $\boldsymbol{a}$ we mean the rank of the indiscernible sequence a_u for one (any) $u \subset k$ with $|u| = k - 1$.*

Definition 2.5 *Define a collection of finite structures $\mathcal{M} = \mathcal{M}_{\kappa,\alpha}$ by induction on the size of the structure. For $n \geq 1$, let $\mathcal{M}_0$ contain the empty structure and let $\mathcal{M}_1$ be the set of all one-element τ-structures.*

Let $\mathcal{M}_n$ be the set of all n-element τ-structures M such that:

    1. for $k \leq n$ the $P_{k,\gamma,\beta}$ partition $[|M|]^k$ as γ, β vary;

    2. every $(n-1)$-element substructure of M is in $\mathcal{M}_{n-1}$;

    3. if $|M| = \{a_0, \ldots, a_{n-1}\}$ and the sequence $\{a_0, \ldots, a_{n-1}\}$ is formally indiscernible of rank β, then $M \models P_{n;\zeta,\gamma}(a_0, \ldots, a_{n-1})$ for some $\gamma < \beta$ and $\zeta < \kappa$;

Finally, $\mathcal{M} = \bigcup_n \mathcal{M}_n$.

Note that clauses 1) and 4) imply that for any $M \in \mathcal{M}$, the n-element subsets of M are partitioned by the $P_{n;\gamma,\beta}$ (some predicates may be empty). The class described in this section is the class of all τ-structures M such that the every finite substructure of M is a member of $\mathcal{M} = \mathcal{M}_{\kappa,\alpha}$.

Definition 2.6 *Let $\boldsymbol{K}_{\kappa,\alpha} = \boldsymbol{K} := \{M \mid every\ finite\ substructure\ of\ M\ is\ in\ \mathcal{M}\}$.*

The terms and notation for special (λ, k)-systems have been defined in Section 1; but our demand here on amalgamation of certain systems is weaker because we only require systems of sufficiently large rank to be amalgamated.

Definition 2.7 *We say that $\boldsymbol{K}$ has the $(\aleph_\zeta, k, \beta)$-amalgamation property if for every special $(\aleph_\zeta, k)$-system $(\overline{N}, \boldsymbol{a})$:*

1. *if $\boldsymbol{a}$ is formally indiscernible of rank strictly greater than β, then there is an amalgam N of $(\overline{N}, \boldsymbol{a})$ in $\boldsymbol{K}$.*

2. *if $\boldsymbol{a}$ is not formally indiscernible, then there is an amalgam N of the system $(\overline{N}, \boldsymbol{a})$ in $\boldsymbol{K}$.*

Note two important but obvious consequences of this definition. The second holds because our requirement that all singletons have rank $\alpha + 2$ means that there is no restriction on which special two systems can be amalgamated. As in Section 1, it suffices to prove special amalgamation.

Remark 2.8   *1. The $(\aleph_\zeta, k, \beta)$-amalgamation property immediately yields the $(\aleph_\zeta, k, \gamma)$-amalgamation property for any γ with $\beta \leq \gamma \leq \alpha + 1$.*

2. *$(\aleph_\zeta, 2, \alpha + 1)$-amalgamation property immediately yields disjoint amalgamation of models with cardinality $\aleph_\zeta$.*

The rest of the proof has two steps. We first show that we have a 'pseudo-amalgamation' on the class of models of cardinality less than $\kappa = \aleph_\delta$, namely $(< \aleph_\delta, k, 0)$-amalgamation for all k. Then we show this 'pseudo-amalgamation' extends up to $\aleph_{\delta+\alpha}$; we prove by induction that $(< \aleph_{\delta+\beta}, k, \beta)$ amalgamation holds for all k and all $\beta \leq \alpha$. This is not k-amalgamation on all systems of size $\aleph_{\delta+\beta}$; only on those systems of sufficiently large rank.

Lemma 2.9 *The class $\boldsymbol{K}$ has the $(< \kappa, k, 0)$-amalgamation property for all $k < \omega$.*

Proof. We can construct models to satisfy the $(< \kappa, s, 0)$-amalgamation property for $s \leq 1$ as in the proof of Lemma 1.7, using that there are always new predicates to define the extension.

Let $(\overline{N}, \boldsymbol{a})$ be a special $(< \kappa, k)$-system. If the sequence $\boldsymbol{a} := \{a_0, \ldots, a_{k-1}\}$ is not formally indiscernible, then we can define a τ-structure N^* on $|N| \cup \boldsymbol{a}$ in an arbitrary way. The resulting model will be a member of $\boldsymbol{K}$ because we are not adding any indiscernible sequences.

Suppose now that $\boldsymbol{a}$ is formally indiscernible of an arbitrary rank $\eta > 0$. We now interpret predicates on $\bigcup_{u \subset k} |N_u|$ to define a model $N \in \boldsymbol{K}$. Define $P_{k,\varsigma,0}(\boldsymbol{a})$ for some unused predicate $P_{k,\varsigma,0}$. It is easy to define the rest of the new predicates (say $P_{k+m,\varsigma_1,\epsilon}$) on sequences $\boldsymbol{ab}$ where $\mathbf{b} \in N$ and $\lg(\mathbf{b}) = m$) to guarantee that for each n, the n-tuples are partitioned. There is no indiscernible sequence in N except $\boldsymbol{a}$ that is not in some N_u. As, any such sequence would have to contain $\boldsymbol{a}$ but $\boldsymbol{a}$ has rank 0. Thus, $N \in \boldsymbol{K}$. $\qquad \square_{2.9}$

Lemma 2.10 *If $\beta + k - 1 \le \alpha + 2$ and $\boldsymbol{K}$ satisfies $(< \lambda, 1, \alpha+1)$-amalgamation, then there is a (λ, k)-system $(\overline{N}, \boldsymbol{a})$ such that the sequence $\boldsymbol{a}$ is formally indiscernible of rank $\beta + 1$.*

Proof. Choose $\boldsymbol{a}$ as an indiscernible sequence of length k with the ranks of s-element subsequences chosen as $\beta + k - s$. In particular, the rank of the entire sequence is β. By Lemma 2.9 there are extensions of $\boldsymbol{a}$ to models in $\boldsymbol{K}$ of every cardinality below κ and then by $(< \lambda, 1, \alpha + 1)$-amalgamation to every power below λ (recall that all singletons have the rank $\alpha + 2$). Let N be union of this chain and let $N_\emptyset$ denote $N - \boldsymbol{a}$. Choose the N_u as $N_\emptyset \cup \{a_i : i \in u\}$. $\qquad \square_{2.10}$

Lemma 2.11 *For all $k < \omega$, if the class $\boldsymbol{K}$ has the $(< \aleph_{\delta+\beta}, k, \beta)$-amalgamation property, then $\boldsymbol{K}$ has the $(\aleph_{\delta+\beta}, k-1, \beta+1)$-amalgamation property.*

Moreover, any special $(\aleph_{\delta+\beta}, k - 1)$-system $(\overline{N}, \boldsymbol{a})$, where $\boldsymbol{a}$ is formally indiscernible of rank strictly greater than $\beta + 1$, can be amalgamated so that the rank of $\boldsymbol{a}$ in the amalgam is $\beta + 1$.

Proof. Let $(\overline{N}, \boldsymbol{a})$ be a special $(\aleph_{\delta+\beta}, k - 1)$-system. If the sequence $\boldsymbol{a} := \{a_0, \ldots, a_{k-2}\}$ is not formally indiscernible, then we can define a τ-structure on $N \cup \boldsymbol{a}$ in an arbitrary way. The resulting model will be a member of $\boldsymbol{K}$ because we are not adding any indiscernible sequences.

Suppose now that $\boldsymbol{a}$ is formally indiscernible of rank strictly greater than $\beta + 1$. Without loss of generality, we can assume that $\bigcup_{u \subset k} |N_u| = \aleph_\beta$ and that the special sequence $\boldsymbol{a}$ is the first $k - 1$ elements of $\aleph_\beta$. We need to define a τ-structure, on $\aleph_\beta$ that extends the N_u.

For some $\eta < \kappa$, assign a predicate $P_{k-1;\eta,\beta+1}$ to hold of $a_0, \ldots, a_{k-2}$. Thus, when we finish the construction, we will satisfied the "moreover" clause of the lemma. Call this structure N^{k-1}. Proceed by induction on $k-1 \le i < \aleph_\beta$ to define the amalgam N^{i+1} with domain $i + 1$. For this, let $\boldsymbol{a}' = \langle a_0, \ldots, a_{k-2}, i \rangle$. And let $(\overline{N}^{i+1}, \boldsymbol{a}')$ be the special $(|i|, k)$-system

$$(\langle \hat{N}_u : u \subset \{a_0, \ldots, a_{k-2}, i\}, |u| = k - 1 \rangle, \boldsymbol{a}')$$

where $\hat{N}_u = N_{u-\{i\}} \cap (i+1)$ if $i \in u$ and $\hat{N}_u = N^i$ if $i \notin u$.

If the sequence $\boldsymbol{a}'$ is not formally indiscernible, then we can amalgamate without a problem. Otherwise, the rank of any $(k-1)$-tuple from $\boldsymbol{a}'$ is $\beta + 1$ (by our choice of $P_{k-1;\eta,\beta+1}$). $(\overline{N}^{i+1}, \boldsymbol{a}')$ is a special $(|i|, k)$-system, and $|i| = \aleph_\epsilon$ with $\epsilon < \delta + \beta$. The amalgam then exists by the $(< \aleph_{\delta+\beta}, k, \beta)$-amalgamation property. The union $\bigcup_{i<\aleph_\beta} N^i$ is the desired amalgam. $\qquad \square_{2.11}$

10

The first sentence of the next theorem is an easy induction using Lemmas 2.9 and 2.11. The second is immediate from Definition 2.7 and Remark 2.8.2.

Theorem 2.12 *Recall* $\kappa = \aleph_\delta$. *For all* $\beta \leq \alpha + 1$, *the class* **K** *has the* $(\aleph_{\delta+\beta}, k, \beta + 1)$-*amalgamation property for all* $k < \omega$. *In particular, the class* **K** *has the* $(\aleph_{\delta+\alpha}, k, \alpha + 1)$-*amalgamation property for all* $k < \omega$ *and so* **K** *has disjoint amalgamation over models of size* $\leq \aleph_{\delta+\alpha}$. *But* **K** *has no model of power* $\beth_{(2^\kappa)^+}$.

The third sentence of the theorems follows from the standard bounds e.g [She78a, Bal] for finding indiscernible sequences. It would be interesting to find a smaller upper bound on the size of the models.

3   Disjoint amalgamation can hold up to $\beth_\alpha$

Under GCH we have shown in Section 2 that for each $\alpha < \omega_1$ there is a sentence of $L_{\omega_1,\omega}$ that has disjoint amalgamation up to $\aleph_\alpha = \beth_\alpha$ but does not have arbitrarily large models. In this section we replace the use of GCH by the use of a generalized Martin's axiom to obtain the result with much less stringent requirements on cardinal exponentiation. Thus, we begin the section by describing and developing some consequence of a generalized Martin's Axiom. In Theorem 3.10, we return to the model theoretic goal. As usual, we can and do carry out the argument for the more general situation of $L_{\kappa^+,\omega}$.

Definition 3.1 ($\mathrm{Ax}_0(\lambda)$) *Let* λ *be a regular cardinal. Let* $\mathbb{P}$ *be a poset such that*

1. *Any decreasing sequence of length less than* λ *has the greatest lower bound.*

2. *(Strong form of* λ^+-*cc) Given any* $\langle p_i : i < \lambda^+ \rangle$ *there is a regressive* h *on* λ^+ *and a club* C *such that for* $\alpha, \beta \in C \cap cof(\lambda)$ *if* $h(\alpha) = h(\beta)$, *then* p_α *and* p_β *are compatible.*

The axiom $\mathrm{Ax}_0(\lambda)$ *states that for any* $\mathbb{P}$ *satisfying (1) and (2), for any family of fewer than* 2^λ *dense sets in* $\mathbb{P}$ *there is a filter meeting them all.*

The following is a standard fact.

Claim 3.2 *The condition (2) in the above definition is equivalent to the following: Given any* $\langle p_i : i < \lambda^+ \rangle$ *there is* $h : \lambda^+ \to A$ *such that*

1. $A = \bigcup_{i < \lambda^+} A_i$;

2. $|A_i| \leq \lambda$;

3. *the sequence of* A_i *is increasing and is continuous at points of cofinality* λ

4. $h(j) \in A_j$ *for all* j;

11

and there is a club C such that for $\alpha, \beta \in C \cap cof(\lambda)$ if $h(\alpha) = h(\beta)$, then p_α and p_β are compatible.

Definition 3.3 *Suppose GCH holds in the ground model, and pick a cardinal κ such that $\kappa^{<\kappa} = \kappa$.*

Let $\{\lambda_i \mid i \leq \alpha\}$ be an increasing continuous sequence of cardinals such that $\lambda_0 = \kappa$, if λ_i is singular, then $\lambda_{i+1} = \lambda_i^+$, and if λ_i is regular, then λ_{i+1} is a regular cardinal greater than λ_i.

Then we say the sequence λ_i is a permissible partial exponentiation function.

Note that the cardinals λ_i, $i \leq \alpha$, will become $\beth_i(\kappa)$ in the forcing extension.

The following was established in [She80], Theorem 4.12. The argument for that theorem calls on [She78b].

Fact 3.4 *Suppose GCH holds in the ground model, and pick a cardinal κ such that $\kappa^{<\kappa} = \kappa$.*

If $\{\lambda_i \mid i \leq \alpha\}$ is a permissible partial exponential function then there is a cardinal-preserving forcing extension V^{P_α} of the ground model such that

1. *for all $i < \alpha$, we have $2^{\lambda_i} = \lambda_{i+1}$;*

2. *for each regular cardinal λ_i, the generalized Martin's axiom $\mathrm{Ax}_0(\lambda_i)$ holds.*

Remark 3.5 *The paper [She78b] demands the following additional condition on partial orders in the statement of the generalized Martin's axiom:*

3. *If $p, q \in \mathbb{P}$ are compatible, then there is the greatest lower bound of p and q in $\mathbb{P}$.*

The paper [She80] states the consistency result without requiring the condition (3), but the proof refers the reader to [She78b]. The proposition below establishes that the condition (3) is not necessary.

Proposition 3.6 *Suppose that $\mathrm{Ax}_0(\lambda)$ holds for posets satisfying the conditions (1)–(3). Let $\mathbb{Q}$ be a poset satisfying only the conditions (1) and (2). Then for any family of fewer than 2^λ dense sets in $\mathbb{Q}$ there is a filter meeting them all, i.e., $\mathrm{Ax}_0(\lambda)$ holds for all posets satisfying (1) and (2).*

Proof. Take $\mathbb{Q}$ satisfying conditions (1) and (2). Let

$$\mathbb{P} := \{X \subset \mathbb{Q}, \mid |X| < \lambda, \text{ elements of } X \text{ have a common lower bound in } \mathbb{Q}\},$$

the partial order is $X \leq Y$ (X is stronger than Y) if $X = Y$ or $X \supset Y$ and X contains a common lower bound for Y in $\mathbb{Q}$.

It is easy to check that $\mathbb{P}$ satisfies (1)–(3) and that, given a dense subset D of $\mathbb{Q}$, the set $C := \{X \in \mathbb{P} \mid X \cap D \neq \emptyset\}$ is dense in $\mathbb{P}$. Now it is easy to get a filter in $\mathbb{Q}$ meeting a collection of fewer than 2^λ dense subsets. $\qquad\square$

From (1) in Fact 3.4, we have that $\lambda_i = \beth_i(\kappa)$. We will need a stronger property of cardinal exponentiation implied by Ax_0: it turns out that, for every regular λ_i, we have $2^\mu = 2^{\lambda_i}$ for all $\lambda_i \leq \mu < 2^{\lambda_i}$. We prove this property after establishing the following two claims. In Corollary 3.9 we conclude in particular that $\lambda_i^{<\lambda_i} = \lambda_i$ whenever λ_i is regular.

Claim 3.7 *Let $\lambda^{<\lambda} = \lambda$, let $\mathrm{Ax}_0(\lambda)$ hold, and let $\lambda \le \mu < 2^\lambda$. For any family $\mathcal{X} = \{X_i \mid i < \mu\}$ of almost disjoint subsets of λ, where $|X_i| = \lambda$, there is a set $X^* \subset \lambda$, $|X^*| = \lambda$, which is almost disjoint from each X_i, $i < \mu$, i.e., $|X^* \cap X_i| < \lambda$ for all $i < \mu$.*

Proof. Let

$$P := \{(f, \mathcal{F}) \mid f \subset \lambda, |f| < \lambda, \mathcal{F} \subset \mathcal{X}, |\mathcal{F}| < \lambda\}.$$

The order is $(f_1, \mathcal{F}_1) \le (f_2, \mathcal{F}_2)$ if $f_1 \supseteq f_2$, $\mathcal{F}_1 \supseteq \mathcal{F}_2$ and $(f_1 - f_2) \cap \bigcup \mathcal{F}_2 = \emptyset$. (The intuition is that the first component is an approximation to the subset we are building and the second component lists the subsets to which we promise not to add any more elements.)

From regularity of λ, it follows that any decreasing chain of fewer than λ conditions in $\mathbb{P}$ has the greatest lower bound. We claim that P has the strong λ^+-chain condition. This is easy given that $\lambda^{<\lambda} = \lambda$. Indeed, let $\{(f_i, \mathcal{F}_i) \mid i < \lambda^+\}$ be a sequence in P, take $A_i := \lambda^{<\lambda}$ for all $i < \lambda^+$ and define $h(i) := f_i \in A_i$. Then whenever $h(i) = h(j)$ we have $f_i = f_j$, and so $(f_i, \mathcal{F}_i)$ and $(f_j, \mathcal{F}_j)$ are compatible: the element $(f_i, \mathcal{F}_i \cup \mathcal{F}_j)$ is below both of them.

Now we define the dense sets D_i, $i < \lambda$, and E_i, $i < \mu$:

$$D_\alpha := \{(f, \mathcal{F}) \mid f \cap [\alpha, \lambda) \ne \emptyset\}, \quad E_i := \{(f, \mathcal{F}) \mid X_i \in \mathcal{F}\}.$$

It is easy to check that the subsets are dense. Given a condition $(f, \mathcal{F})$, there is a set $X \in \mathcal{X} - \mathcal{F}$, so in particular $|X| = \lambda$ and X is almost disjoint from every set in $\mathcal{F}$. Thus, $(\bigcup \mathcal{F}) \cap X$ is bounded in λ. So there is a point $\delta \in X \cap [\alpha, \lambda)$. Now $(f \cup \{\delta\}, \mathcal{F}) \in D_\alpha$ is a condition that extends $(f, \mathcal{F})$. Density of E_i is immediate.

A filter G meeting all the dense sets D_α and E_i gives the desired subset of λ: we let $X^* := \bigcup_{(f, \mathcal{F}) \in G} f$. For each of the sets $X_i \in \mathcal{X}$, there is a condition $(f, \mathcal{F}) \in G$ such that $\mathcal{F}$ contains X_i. By the definition of the order, the intersection $X^* \cap X_i$ is equal to $f \cap X_i$. Since the filter meets every D_α, the set X^* is unbounded in λ. $\qquad\square$

Claim 3.8 *Let $\lambda^{<\lambda} = \lambda$, suppose that $\mathrm{Ax}_0(\lambda)$ holds, and let $\lambda \le \mu < 2^\lambda$. For any μ with $\lambda \le \mu < 2^\lambda$ we have $2^\mu = 2^\lambda$.*

Proof. Let $\{X_i \mid i < \mu\}$ be a family of almost disjoint subsets of λ, each of cardinality λ. For each $S \subset \mu$ we construct a subset $A_S \subset \lambda$ such that A_S has an unbounded intersection with X_i for every $i \in S$, and is almost disjoint from X_j for every $j \notin S$. The resulting map $A_S \mapsto S$ gives a surjection from (a part of) $\mathcal{P}(\lambda)$ to $\mathcal{P}(\mu)$.

The argument is the same as the one in the previous claim: let

$$P := \{(f, \mathcal{F}) \mid f \subset \lambda, |f| < \lambda, \mathcal{F} \subset \{X_i \mid i \in S\}, |\mathcal{F}| < \lambda\}.$$

The order is $(f_1, \mathcal{F}_1) \le (f_2, \mathcal{F}_2)$ if $f_1 \supseteq f_2$, $\mathcal{F}_1 \supseteq \mathcal{F}_2$ and $(f_1 - f_2) \cap \bigcup \mathcal{F}_2 = \emptyset$. The conditions on the partial order are easily verified, just as in the previous claim.

13

The relevant dense sets in this case are $D_{i,\alpha}$ for $\alpha < \lambda$ and $i \in S$, and E_i, for $i \in \mu - S$ where

$$D_{i,\alpha} := \{(f, \mathcal{F}) \mid f \cap X_i \cap [\alpha, \lambda) \neq \emptyset\}, \ i \in S, \ \alpha < \lambda$$

and

$$E_j := \{(f, \mathcal{F}) \mid X_j \in \mathcal{F}\}, \ j \in \mu - S.$$

The dense sets $D_{i,\alpha}$ make sure that we meet the set X_i, $i \in S$, unboundedly many times, and the sets E_j ensure that every set X_j, $j \notin S$, has a bounded intersection with the set X_S. $\qquad\square$

Corollary 3.9 *Let κ be a cardinal such that $\kappa^{<\kappa} = \kappa$, let $\alpha < \kappa^+$, and let $\{\lambda_i \mid i \leq \alpha\}$ be a permissible partial exponential function.*

Let V be a set-theoretic universe such that

1. *for all $i < \alpha$, we have $2^{\lambda_i} = \lambda_{i+1}$;*

2. *for each of the regular cardinals λ_i, the axiom $\mathrm{Ax}_0(\lambda_i)$ holds.*

Then for each of the regular cardinals λ_i, $i < \alpha$, we have $\lambda_i^{<\lambda_i} = \lambda_i$.

Proof. We prove this by induction on $i < \alpha$. For $i = 0$, the statement is given by the assumption on κ. If i is a limit ordinal, then λ_i is singular, since $\alpha < \kappa^+$. If $i = j + 1$, then there are two cases. If i is a limit ordinal, then $\lambda_j^{<\lambda_j} = (2^{\lambda_i})_i^\lambda = 2^{\lambda_i} = \lambda_j$. If i is a successor, then $\mathrm{Ax}_0(\lambda_i)$ holds, and by Claim 3.8 $2^\mu = 2^{\lambda_i}$ for all μ such that $\lambda_i \leq \mu < 2^{\lambda_i}$. But then $\lambda_j^{<\lambda_j} = \lambda_j$. $\qquad\square$

Let α be a limit ordinal, and κ a cardinal such that $\kappa < \alpha < \kappa^+$. Let $\boldsymbol{K} = \boldsymbol{K}_{\kappa,\alpha}$ be the class built in Section 2.

The main result of this section is the following:

Theorem 3.10 *For any limit ordinal α, where $\kappa \leq \alpha < \kappa^+$, it is relatively consistent with ZFC+ the exponential function below $\beth_\alpha$ is any permissible function that the class $\boldsymbol{K} = \boldsymbol{K}_{\kappa,\alpha}$ has the disjoint amalgamation property in every cardinal up to $\beth_\alpha(\kappa)$ but $\boldsymbol{K}$ does not have arbitrarily large models.*

For the remainder of the section, we work in V^{P_α} from Fact 3.4. The proof of the theorem will use the following steps.

By Lemma 2.9 in Section 2, for every $\mu < \kappa$ and every $k < \omega$ there are (μ, k)-systems of models in $\boldsymbol{K}$, and the class $\boldsymbol{K}$ has $(< \kappa, k, 0)$-amalgamation property for all $k < \omega$.

Next we restate Lemma 2.11 in the form appropriate for this section replacing $\aleph_\alpha$ by λ (the proof is identical). This allows us in the succeeding Lemma to move amalgamation below λ not just below λ^+ but below 2^λ (using our additional set theoretic hypotheses).

Lemma 3.11 *For all $k < \omega$, if the class $\boldsymbol{K}$ has the $(< \lambda, k, \beta)$-amalgamation property, then $\boldsymbol{K}$ has the $(\lambda, k-1, \beta+1)$-amalgamation property.*

Moreover, any special $(\lambda, k-1)$-system $(\overline{N}, \boldsymbol{a})$, where $\boldsymbol{a}$ is formally indiscernible of rank strictly greater than $\beta + 1$, can be amalgamated so that the rank of $\boldsymbol{a}$ in the amalgam is $\beta + 1$.

The new ingredient is the following lemma.

Lemma 3.12 *Suppose $\lambda \geq \kappa$ is a regular cardinal, and $\mathrm{Ax}_0(\lambda)$ holds. For all $k < \omega$, if $\boldsymbol{K}$ has $(< \lambda, \leq k+2, \beta)$-amalgamation property, then for all μ, $\lambda \leq \mu < 2^\lambda$, the class $\boldsymbol{K}$ has $(\mu, k, \beta+2)$-amalgamation property.*

Proof. Let $\lambda \leq \mu < 2^\lambda$. Fix a special system $\overline{N} = \{N_u \mid u \subset k\}$, where $\bigcup N_u = \mu$ and $a_\ell = \ell$ for $\ell < k$. If $\{a_\ell \mid \ell < k\}$ are not formally indiscernible, we can amalgamate easily, so suppose $\{a_\ell \mid \ell < k\}$ are formally indiscernible with rank of a_u strictly greater than $\beta + 2$ for all $u \subset k$. Our goal is to make $\bigcup |N_u|$ into a model.

Define a partially ordered set $\mathbb{P}$.

Definition 3.13 *Let $\mathbb{P}$ be a set of models $M \in \boldsymbol{K}$ such that*

1. *M has the universe $B \subset \mu$, $|B| < \lambda$;*

2. *$\{a_l \mid l < k\} \subseteq B$;*

3. *if $u \subset k$, then $N_u \upharpoonright (B \cap |N_u|)$ is a substructure of M;*

4. *the rank of $\boldsymbol{a}$ in M is at least $\beta + 2$;*

5. *the rank of every $(k+1)$-element indiscernible sequence extending $\boldsymbol{a}$ in M is at least $\beta + 1$.*

The partial order is the reverse $\boldsymbol{K}$-submodel relation.

Claim 3.14 *The partially ordered set $\mathbb{P}$ satisfies the conditions of the generalized Martin's axiom $\mathrm{Ax}_0(\lambda)$.*

Proof of Claim. The cardinality of $\mathbb{P}$ is $\mu^{<\lambda} < 2^\lambda$. The poset $\mathbb{P}$ is closed under descending chains of length less than λ.

It remains to check the strong chain condition. Let $\{M_i \mid i < \lambda^+\}$ be a sequence of elements in $\mathbb{P}$. For $i < \lambda^+$, let A_i be the set of $\boldsymbol{K}$-models of size strictly less than λ with the universe contained in $\bigcup_{j<i} |M_j|$. There are $\lambda^{<\lambda} = \lambda$ choices for the universe of such a $\boldsymbol{K}$-model; and for each universe, there are $2^{<\lambda}$ choices for each predicate; and κ predicates, so $\max\{\kappa, 2^{<\lambda}\} \leq \lambda$ ways to define a τ-structure.

Thus, the cardinality of each set A_i is λ; continuity at points of cofinality λ follows from regularity of λ.

Let $A := \bigcup_{i < \lambda^+} A_i$. Define $h : \lambda^+ \to A$ by letting

$$h(i) := M_i \upharpoonright (|M_i| \cap \bigcup_{j < i} |M_j|).$$

If $h(\delta_1) = h(\delta_2)$, then $M_1 := M_{\delta_1}$ and $M_2 := M_{\delta_2}$ agree on their intersection. We will show that these two elements of $\mathbb{P}$ are compatible, i.e., there is a **K**-model $N \in \mathbb{P}$ extending both M_1 and M_2. To show that such a model exists, we need to use $(k+2)$-amalgamation.

Without loss of generality, we may assume that $M_1 = M_0 \cup \{b_1\}$ and $M_2 = M_0 \cup \{b_2\}$, where $b_1 \neq b_2$. To construct the model N, we amalgamate the special system:

$$\{N_u \upharpoonright (|N_u| \cap (|M_1| \cup |M_2|)) \mid u \subset k, \ |u| = k-1\} \cup \{M_1, M_2\}$$

with the root M_0 and the special sequence $\mathbf{b} = \{a_0, \dots, a_{k-1}, b_1, b_2\}$.

If the sequence $\mathbf{b}$ is not formally indiscernible, we amalgamate without a problem; otherwise the rank of the subsequence $\{a_0, \dots, a_{k-1}, b_1\}$ has to be at least $\beta + 1$ by (5) in the definition of $\mathbb{P}$. So the amalgam N exists by the special $(< \lambda, k+2, \beta)$-amalgamation. $\qquad \square_{3.14}$

To finish the proof of Lemma 3.12, we need to show that for each $i < \mu$ the set $D_i := \{M \in \mathbb{P} \mid i \in |M|\}$ is dense in $\mathbb{P}$.

This is easy by $(< \lambda, k+1, \beta+1)$-amalgamation. If $M \in \mathbb{P}$ and $i \notin |M|$, we amalgamate the special system

$$\{N_u \upharpoonright (|N_u| \cap (|M| \cup \{i\})) \mid u \subset k, \ |u| = k-1\} \cup \{M\}$$

with the root $N_\emptyset \cap M$ and the special sequence $\mathbf{b} = \{a_0, \dots, a_{k-1}, i\}$. If $\mathbf{b}$ is not formally indiscernible, we can amalgamate without a problem, otherwise the rank of $\mathbf{b}$ has to be strictly greater than $\beta + 1$. By Lemma 2.11, in particular by the "moreover" clause, we get the amalgam M' of the system in which the rank of $\mathbf{b}$ is $\beta + 1$. Thus, $M' \in \mathbb{P}$.

By $\mathrm{Ax}_0(\lambda)$, there is a filter G on $\mathbb{P}$ meeting all the dense sets D_i, $i < \mu$. The desired model $M^* := \bigcup G$ is the needed amalgam. $\qquad \square_{3.12}$

Proof of Theorem 3.10. Using Lemmas 2.9, 3.11, and 3.12, we have the following by induction on $\gamma \leq \alpha$:

1. if $\gamma = 0$ or γ is a limit ordinal, then **K** has $(< \beth_\gamma(\kappa), k, \gamma)$-amalgamation for all $k < \omega$;

2. if $\gamma = \beta + 1$ where β is a limit ordinal, then **K** has $(< \beth_\gamma(\kappa), k, \gamma+1)$-amalgamation for all $k < \omega$;

3. if $\gamma = \beta + 1$ where β is a successor ordinal, then **K** has $(< \beth_\gamma(\kappa), k, \gamma+2)$-amalgamation for all $k < \omega$.

In particular, we have $(\beth_\gamma(\kappa), 2, \alpha + 1)$-amalgamation for all $\gamma \leq \alpha$. By Remark 2.8.2, it means that we have the disjoint amalgamation over the models of size up to $\beth_\alpha(\kappa)$. $\qquad \square_{3.10}$

This completes our arguments showing that 'disjoint amalgamation' can amalgamation can fail late. Among the intriguing questions is deleting the word disjoint.

References

[Bal]   John T. Baldwin. Categoricity. www.math.uic.edu/~ jbaldwin.

[Gro02]   Rami Grossberg. Classification theory for non-elementary classes. In Yi Zhang, editor, *Logic and Algebra*, pages 165–204. AMS, 2002. Contemporary Mathematics 302.

[GV06]   R. Grossberg and M. VanDieren. Shelah's categoricity conjecture from a successor for tame abstract elementary classes. *The Journal of Symbolic Logic*, 71:553–568, 2006.

[Hjo07]   Greg Hjorth. Knight's model, its automorphism group, and characterizing the uncountable cardinals. *Notre Dame Journal of Formal Logic*, 2007.

[Kni77]   J.F. Knight. A complete $l_{\omega_1,\omega}$-sentence characterizing $\aleph_1$. *Journal of Symbolic Logic*, 42:151–161, 1977.

[Les]   Olivier Lessmann. Upward categoricity from a successor cardinal for an abstract elementary class with amalgamation. preprint.

[LS93]   Michael C. Laskowski and Saharon Shelah. On the existence of atomic models. *J. Symbolic Logic*, 58:1189–1194, 1993.

[Mor65]   M. Morley. Omitting classes of elements. In Addison, Henkin, and Tarski, editors, *The Theory of Models*, pages 265–273. North-Holland, Amsterdam, 1965.

[She78a] S. Shelah. *Classification Theory and the Number of Nonisomorphic Models*. North-Holland, 1978.

[She78b] Saharon Shelah. A weak generalization of MA to higher cardinals. *Israel J Math*, 30:297–306, 1978.

[She80]   Saharon Shelah. Simple unstable theories. *Annals of Mathematical Logic*, 19:177–203, 1980.

[She83a] S. Shelah. Classification theory for nonelementary classes. I. the number of uncountable models of $\psi \in L_{\omega_1\omega}$ part A. *Israel Journal of Mathematics*, 46:3:212–240, 1983. paper 87a.

[She83b] S. 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;3:241–271, 1983. paper 87b.

[She99] S. Shelah. Categoricity for abstract classes with amalgamation. *Annals of Pure and Applied Logic*, 98:261–294, 1999. paper 394. Consult Shelah for post-publication revisions.

[Sou] Ioannis Souldatos. Notes on cardinals that characterizable by a complete (scott) sentence. preprint.