Disjoint amalgamation in locally finite AEC[*]

John T. Baldwin[†]　　　Martin Koerwien[‡]　　　Michael C. Laskowski[§]
University of Illinois at Chicago　　　KGRC　　　University of Maryland

June 20, 2015

Abstract

We introduce the concept of a *locally finite* abstract elementary class and develop the theory of excellence (with respect to *disjoint* $(\leq \lambda, k)$-*amalgamation*) for such classes. From this we find a family of complete $L_{\omega_1, \omega}$ sentences ϕ_r that a) homogeneously characterizes $\aleph_r$ (improving results of Hjorth [12] and Laskowski-Shelah [14] and answering a question of [22]), while b) the ϕ_r provide the first examples of a class of models of a complete sentence in $L_{\omega_1, \omega}$ where the spectrum of cardinals in which amalgamation holds is other that none or all.

1　Introduction

Amalgamation[1], finding a model M_2 in a given class $\boldsymbol{K}$ into which each of two extensions M_0, M_1 of a model $M \in \boldsymbol{K}$ can be embedded, has been a theme in model theory in the almost 60 years since the work of Jónsson and Fraïssé. An easy application of compactness shows that amalgamation holds for every triple of models of a complete first order theory. For an $L_{\omega_1, \omega}$-sentence ϕ, the situation is much different; there can be a bound on the cardinality of models of ϕ and whether the amalgamation property holds can depend on the cardinality of the particular models. Shelah generalized the Jónsson context for homogeneous-universal models to that of an abstract elementary class by providing axioms governing the notion of *strong substructure*. He introduced the notion of n-dimensional amalgamation in an infinite cardinal λ and used it to prove that excellence (r-dimensional amalgamation in $\aleph_0$ for every $r < \omega$) implies ϕ has arbitrarily large models and r-dimensional amalgamation in all cardinals. We introduce an analogy to excellence–defining disjoint $(\leq \lambda, k)$-amalgamation for classes of finite structures satisfying a closure of intersections property (Definition 2.1.3.) We strengthen the necessity of amalgamation[2] in all $r < \omega$ by constructing for each r a sentence ϕ_r which satisfies disjoint $(\leq \aleph_0, k)$-amalgamation for $k \leq r$ but which has *no* model in $\aleph_{r+1}$.

In [12], Hjorth found, by an inductive procedure, for each $\alpha < \omega_1$, a countable (finite for finite α) set S_α of complete $L_{\omega_1, \omega}$-sentences such that some $\phi_\alpha \in S_\alpha$ characterizes $\aleph_\alpha$ (ϕ_α has a model of that cardinality but no larger model). It is conjectured [21] that it may be impossible to decide in ZFC which sentence

---

[1] Reference to amalgamation or 2-amalgamation are to this notion; we try to be careful about our variant of what is labeled 'disjoint' (often called strong in the literature) amalgamation.

[2]Hart-Shelah [11] provided an earlier example showing there are ϕ_r categorical up to $\aleph_r$ but then losing categoricity. Those examples have arbitrarily large models and satisfy amalgamation in all cardinals [4].

works. In this note, we show a modification of the Laskowski-Shelah example (see [14, 6]) gives a family of $L_{\omega_1,\omega}$-sentences ϕ_r, which homogeneously (see Definition 2.4.1) characterize $\aleph_r$ for $r < \omega$. Thus for the first time Theorem 3.2.20 establishes in ZFC, the existence of specific sentences ϕ_r characterizing $\aleph_r$.

Our basic objects of study (Section 2) are classes $\boldsymbol{K}_0$ of finite structures with ordinary substructure taken as 'strong'. Given $\boldsymbol{K}_0$, we consider two ancillary classes of structures in the same vocabulary. $(\hat{\boldsymbol{K}}, \leq)$ denotes the structures that are locally (See Definition 2.1.2) in $\boldsymbol{K}_0$; this is what is meant by a *locally finite AEC*. If $(\boldsymbol{K}_0, \leq)$ satisfies amalgamation and there are only countably many isomorphism types, then there is a countable generic model M, which is always *rich* (Definition 2.1.7) and atomic, at least after adding some new relation symbols to describe $L_{\omega_1,\omega}$-definable subsets. The class $\mathbf{R}$ of rich models is the collection of all structures satisfying the Scott sentence ϕ_M of M. Now our principal results go in two directions: building models of $\boldsymbol{K}$ and $\mathbf{R}$ with cardinality up to some $\aleph_r$ and showing there are no larger models.

If $(\boldsymbol{K}_0, \leq)$ satisfies our notion of *disjoint* $(< \aleph_0, r + 1)$ amalgamation then (by a new construction) both $\hat{\boldsymbol{K}}$ and $\mathbf{R}$ have models in $\aleph_r$ and satisfy disjoint $(< \aleph_s, r - s)$ amalgamation for $s \leq r$. We modify [6] to show the existence of homogeneous characterizations (Definition 2.4.1) (arising from [12]); this leads to new examples of joint embedding spectra in [7, 9].

For the other direction (Section 3), if $\hat{\boldsymbol{K}}$ is a locally finite AEC, then for any $M \in \hat{\boldsymbol{K}}$, defining $\mathrm{cl}(A)$ to be the smallest substructure of M containing A for any subset $A \subseteq M$ is a locally finite closure relation. If our AEC $\hat{\boldsymbol{K}}$ forbids an independent subset of size $r + 2$, then a combinatorial argument disallows a model of size $\aleph_{r+1}$. We construct particular examples $\mathbf{R}_r$ for each $r < \omega$ that have such a locally finite closure relation and so homogeneously characterize $\aleph_r$. Automatically they fail *disjoint* $(\leq \aleph_{r-1}, 2)$ amalgamation. Rather technical arguments demonstrate the failure of 'normal' 2-amalgamation in $\aleph_{r-1}$ for $\hat{\boldsymbol{K}}$ and R.

We conclude in Section 4 by putting the results in context and speculating on the number of models in a cardinal characterized by a complete sentence of $L_{\omega_1,\omega}$. These results on characterizing cardinals are intimately connected with spectra of (disjoint) amalgamation. The *finite amalgamation spectrum* of an abstract elementary class $\boldsymbol{K}$ is the set $X_{\boldsymbol{K}}$ of $n < \omega$ such that $\boldsymbol{K}$ has a model in $\aleph_n$, $\aleph_n$ is at least the Löwenheim-Skolem number of $\boldsymbol{K}$, and satisfies amalgamation[3] in $\aleph_n$. We discuss in Section 4 the other known spectra. The paper presents the first spectra of an AEC which is not either an initial or a co-initial interval.

We thank Ioannis Souldatos for conversations leading to clearer formulation of some problems, Alex Kruckman for some useful comments, and Will Boney for a discussion redirecting our focus to excellence.

2   Locally Finite AEC, k-Disjoint Amalgamation, and rich models

We begin by generating two AEC from a class $\boldsymbol{K}_0$ of finite structures that is closed under isomorphisms and substructures. The first, $\hat{\boldsymbol{K}}$ is defined directly from $\boldsymbol{K}_0$. The second is the subclass of rich models $\mathbf{R} \subseteq \hat{\boldsymbol{K}}$ (see Definition 2.1.7) which only exist under additional hypotheses. Of particular interest will be the case where there is a unique countable rich model M which is an atomic model of its first-order theory $Th(M)$. In that case, it follows that every rich model is atomic with respect to $Th(M)$ and we denote the class of rich models as $\mathbf{At} = \mathbf{At}(\boldsymbol{K}_0)$; it can be viewed interchangeably as the class of atomic models of an associated first order theory or the models of a *complete* sentence ϕ in $L_{\omega_1,\omega}$. (See, e.g. Chapter 6 of [3].).

---

[3]For the precise formulations of amalgamation see Definition 2.4.1 and Remark 2.2.5.

2.1 Locally Finite AEC's

We will use the following background fact[4] about isomorphism that holds for all structures. When we write A extends B we mean that for a fixed vocabulary τ, A is a τ-substructure of B.

Fact 2.1.1. *Suppose M is a substructure of N, and that A is a union of structures extending M. Then, there exists a copy N' that is isomorphic to N over M, with $N' \cap A = M$.*

There is nothing mysterious about this remark. To prove it, choose a set M' of cardinality $|M - N|$ and disjoint from A. Let f be a bijection between $M - N$ and M' and the identity on M. Define the relations on M' to be the image of the relations on M under f. However, difficulties arise when one wants to make such a construction inside specified ambient models.

Definition 2.1.2. *Suppose that K_0 is a class of finite structures in a countable vocabulary τ and that K_0 is closed under isomorphisms. We call the class of τ-structures M (including the empty structure) with the property that every finite subset $A \subseteq |M|$ is contained in a finite substructure $N \in K_0$ of M, K_0-locally finite and denote it by $\hat{K}$.*

For reasons explained in Remark 2.2.11, it is necessary to make an additional assumption to transfer amalgamation properties from finite to countable structures. Whenever we speak of the intersection of two structures M and N, we mean that the intersection of their domains M_0 is the universe of τ-substructure of each.

Definition 2.1.3. *An AEC $(K, \prec_K)$ satisfies the non-trivial intersection property if for any $M, N \in K$, (1) the intersection of their domains M_0 is the universe of an element of K and (2) $M_0 \prec_K M$ and $M_0 \prec_K N$.*

This property has a surprisingly strong consequence.

Lemma 2.1.4. *An AEC $(K, \prec_K)$ in a vocabulary τ satisfies the non-trivial intersection property if and only if (a) K is closed under substructure and (b) $\prec_K$ is substructure.*

Proof: Clearly, if (a) and (b) hold, then we have closure under intersection. For the converse, assume $(K, \prec_K)$ is closed under intersection. Let $M \in K$ and let A be an arbitrary τ-substructure of M. By Fact 2.1.1 there is another copy M' of M, whose intersection with M is precisely A. As both $M, M' \in K$, A must be in K as well. The verification of (b) is similar; if A is not strong in M, condition (2) in Definition 2.1.3 fails. $\square_{2.1.4}$

In light of Lemma 2.1.4, we can proceed concretely. *We fix a vocabulary τ and a class K_0 of τ-structures that is closed under substructure.* We associate a locally finite AEC by making one change from the standard ([18] or [3]) definition of abstract elementary class: modify the usual notion of Löwenheim-Skolem number as follows.

Definition 2.1.5. *A class $(K, \leq)$ of τ-structures and the relation $\leq$ as ordinary substructure is a locally finite abstract elementary class if it satisfies the normal axioms for an AEC except the usual Löwenheim-Skolem condition is replaced by: If $M \in K$ and A is a finite subset of M, then there is a finite $N \in K$ with $A \subset N \leq M$.*

---

[4]This notion appears in some philosophy papers with the evocative name: push-through construction [10].

One example of a locally finite AEC with the non-trivial intersection property follows. Given a class $(\mathbf{K}_0, \leq)$ of finite τ-structures, closed under isomorphism and substructure, we let $\hat{\mathbf{K}}$ denote the class of $\mathbf{K}_0$-locally finite τ-structures. Then $(\hat{\mathbf{K}}, \leq)$ is a locally finite AEC. Somewhat surprisingly, these are the only examples.

Proposition 2.1.6. *Let $(\mathbf{K}, \leq)$ be any locally finite AEC in a vocabulary τ with the non-trivial intersection property and let $\mathbf{K}_0$ denote the class of finite structures in $\mathbf{K}$. Then:*

1. *Both $\mathbf{K}$ and $\mathbf{K}_0$ are closed under substructures; and*

2. *$\mathbf{K}$ is equal to the class of $\mathbf{K}_0$-locally finite τ-structures.*

Proof: The first clause follows immediately from Lemma 2.1.4. For the second, note that given any $M \in \mathbf{K}$ and finite subset $A \subseteq M$, the Löwenheim-Skolem condition on $\mathbf{K}$ yields a finite substructure $B \leq M$ with $A \subseteq B$. As $\mathbf{K}$ is closed under substructures, B is in $\mathbf{K}$, and hence in $\mathbf{K}_0$. Thus, M is $\mathbf{K}_0$-locally finite.

Conversely, we prove by induction on cardinals λ that every $\mathbf{K}_0$-locally finite structure of size λ is in $\mathbf{K}$. This is immediate when λ is finite, so fix an infinite λ and assume that $\mathbf{K}$ contains every $\mathbf{K}_0$-locally finite structure of size less than λ. Choose any $\mathbf{K}_0$-locally finite structure M of size λ. As M is locally finite, we can find a continuous chain $\langle N_\alpha : \alpha < \lambda \rangle$ of substructures of M, each of size less than λ, whose union is M. It is easily verified that each N_α is $\mathbf{K}_0$-locally finite, hence each $N_\alpha \in \mathbf{K}$ by our inductive hypothesis. As $\mathbf{K}$ is closed under unions of chains, it follows that $M \in \mathbf{K}$. $\quad \square_{2.1.6}$

The following notions are only used when the class of finite structures $(\mathbf{K}_0, \leq)$ has the joint embedding property (JEP).

Definition 2.1.7. *Let $(\mathbf{K}_0, \leq)$ denote a class of finite τ-structures and let $(\hat{\mathbf{K}}, \leq)$ denote the associated locally finite abstract elementary class.*

1. *A model $M \in \hat{\mathbf{K}}$ is finitely $\mathbf{K}_0$-homogeneous or* rich *if for all finite $A \leq B \in \mathbf{K}_0$, every embedding $f : A \to M$ extends to an embedding $g : B \to M$. We denote the class of rich models in $\hat{\mathbf{K}}$ as $\mathbf{R}$.*

2. *The model $M \in \hat{\mathbf{K}}$ is* generic *if M is rich and M is an increasing union of a chain of finite substructures, each of which are in $\mathbf{K}_0$.*

It is easily checked that if rich models exist for a class $(\mathbf{K}_0, \leq)$ of finite structures with JEP, then $(\mathbf{R}, \leq)$ is an AEC with Löwenheim-Skolem number equal to the number of isomorphism classes of $\mathbf{K}_0$ (provided this number is infinite). As well, any two rich models are $L_{\infty,\omega}$-equivalent. Also, a rich model M is generic if and only if M is countable.

We will be interested in cases where a generic model M exists, and that M is an atomic model of its first-order theory. Curiously, this second condition has nothing to do with the structure embeddings on the class $\mathbf{K}_0$, but rather with our choice of vocabulary. The following condition is needed when, for some values of n, $\mathbf{K}_0$ has infinitely many isomorphism types of structures of size n.

Definition 2.1.8. *A class $\mathbf{K}_0$ of finite structures in a countable vocabulary is* separable *if, for each $A \in \mathbf{K}_0$ and enumeration $\mathbf{a}$ of A, there is a quantifier-free first order formula $\phi_{\mathbf{a}}(\mathbf{x})$ such that:*

- *$A \models \phi_{\mathbf{a}}(\mathbf{a})$; and*

- *for all $B \in \mathbf{K}_0$ and all tuples $\mathbf{b}$ from B, $B \models \phi_A(\mathbf{b})$ if and only if $\mathbf{b}$ enumerates a substructure B' of B and the map $\mathbf{a} \mapsto \mathbf{b}$ is an isomorphism.*

4

In practice, we will apply the observation that if for each $A \in \boldsymbol{K}_0$ and enumeration $\boldsymbol{a}$ of A, there is a quantifier-free formula $\phi'_{\boldsymbol{a}}(\mathbf{x})$ such that there are only finitely many $B \in \boldsymbol{K}_0$ with cardinality $|A|$ that under some enumeration $\mathbf{b}$ satisfy $\phi'_{\boldsymbol{a}}(\mathbf{b})$, then $\boldsymbol{K}_0$ is separable.

Lemma 2.1.9. *Suppose τ is countable and $\boldsymbol{K}_0$ is a class of finite τ-structures that is closed under substructure, satisfies amalgamation, and JEP, then a $\boldsymbol{K}_0$-generic (and so rich) model M exists. Moreover, if $\boldsymbol{K}_0$ is separable, M is an atomic model of $Th(M)$. Further, $\mathbf{R} = \mathbf{At}$, i.e., every rich model N is an atomic model of $Th(M)$.*

Proof: Since the class $\boldsymbol{K}_0$ of finite structures is separable it has countably many isomorphism types, and thus a $\boldsymbol{K}_0$-generic M exists by the usual Fraïssé construction. To show that M is an atomic model of $Th(M)$, it suffices to show that any finite tuple $\boldsymbol{a}$ from M can be extended to a larger finite tuple $\mathbf{b}$ whose type is isolated by a complete formula. Coupled with the fact that M is $\boldsymbol{K}_0$-locally finite, we need only show that for any finite substructure $A \leq M$, any enumeration $\boldsymbol{a}$ of A realizes an isolated type. Since every isomorphism of finite substructures of M extends to an automorphism of M, the formula $\phi_{\boldsymbol{a}}(\mathbf{x})$ isolates $\mathrm{tp}(\boldsymbol{a})$ in M.

The final sentence follows since any two rich models are $L_{\infty,\omega}$-equivalent. $\square_{2.1.9}$

Using Definition 2.1.8 and Lemma 2.1.9 as a guide, we can see what we need to expand our vocabulary to ensure that a generic model will become atomic with respect to its theory.

Lemma 2.1.10. *Let $(\boldsymbol{K}_0, \leq)$ be any class of finite τ-structures, closed under substructure, for which a generic model M exists. Then there is a vocabulary $\tau^* \supseteq \tau$ and a related class $(\boldsymbol{K}_0^*, \leq)$ of finite τ^*-structures satisfying:*

- *Every $A \in \boldsymbol{K}_0$ has a canonical expansion to an $A^* \in \boldsymbol{K}_0^*$;*

- *The class $(\hat{\boldsymbol{K}}^*, \leq)$ consisting of all $\boldsymbol{K}_0^*$-locally finite τ^* structures is a locally finite AEC. Moreover, every $N \in \hat{\boldsymbol{K}}$ has a canonical expansion to an $N^* \in \hat{\boldsymbol{K}}^*$;*

- *An element $N \in \hat{\boldsymbol{K}}$ is $\boldsymbol{K}_0$-rich if and only if its canonical expansion N^* is $\boldsymbol{K}_0^*$-rich. In particular, the canonical expansion M^* of the $\boldsymbol{K}_0$-generic is $\boldsymbol{K}_0^*$-generic;*

- *M^* is an atomic model of $Th(M^*)$.*

Proof. For each n, for every isomorphism type $A \in \boldsymbol{K}_0$ of cardinality n, and for every enumeration $\boldsymbol{a}$ of A, add a new n-ary predicate $R_{\boldsymbol{a}}(\mathbf{x})$ to τ^*. The canonical expansion B^* of any $B \in \boldsymbol{K}_0$ is formed by positing that $R_{\boldsymbol{a}}(\mathbf{b})$ holds of some $\mathbf{b} \in (B^*)^n$ if and only if the bijection $\boldsymbol{a} \to \mathbf{b}$ is a τ-isomorphism. Let $\boldsymbol{K}_0^*$ be the class of all B^* for $B \in \boldsymbol{K}_0$ and let $\hat{\boldsymbol{K}}^*$ be the class of all $\boldsymbol{K}_0^*$-locally finite τ^*-structures. That M^* is an atomic model of $Th(M^*)$ follows from Lemma 2.1.9. $\square_{2.1.10}$

2.2 k-configurations

Within the context of Assumption 2.2.2, we develop a simpler analog of Shelah's notion of excellence. Excellence was first formulated [16, 17] in an ω-stable context that takes place entirely in the context of atomic models. There are two complementary features: n-existence implies there are arbitrarily large atomic models; n-uniqueness gives more control of the models and the analog of Morley's theorem. We have separated these functions. $(< \aleph_0, n)$-disjoint amalgamation plays the role of n-existence. But there is no uniqueness. Shelah develops there a substantial apparatus to define 'independence' and excellence concerns

'independent systems'. He develops an abstract version of these notions for 'universal classes' in [19]. Closer to our context here is the study of $(< \lambda, k)$ systems in [8]. The applications here require much less machinery that either of these, because we are able to exploit disjoint amalgamation and our classes are closed under substructure.

Notation 2.2.1. *For a given vocabulary τ, a τ-structure A is* minimal *if it has no proper substructure. If τ has no constant symbols, we allow A to be the empty τ-structure.*

Assumption 2.2.2. *Throughout this subsection we have a fixed vocabulary τ with a fixed, minimal τ-structure A. We consider classes $(\boldsymbol{K}, \leq)$ of τ-structures, where $\leq$ denotes 'substructure' and every $M \in \boldsymbol{K}$ is locally finite and has A as a substructure. We additionally assume that $\boldsymbol{K}$ is closed under substructures, isomorphisms fixing A pointwise, and unions of continuous chains of arbitrary ordinal length.*

We establish some notation that is useful for comparing finite and infinite structures.

Notation 2.2.3. *It is convenient to let $\aleph_{-1}$ be a synonym for 'finite'. For A any set, we write $\aleph(A) = \aleph_{-1}$ if and only if A is finite. For infinite sets A, $\aleph(A)$ denotes the usual cardinality $|A|$. Also, the successor of $\aleph_{-1}$ is $\aleph_0$, i.e., $(\aleph_{-1})^+ = \aleph_0$.*

The basic objects of study are k-configurations from a class $\boldsymbol{K}$ satisfying Assumption 2.2.2. Unlike Shelah's development of k-systems of models indexed by the set $\mathcal{P}(k)^-$ (which can be thought of as being $2^k - 1$ vertices of a k-dimensional cube) with the requirement that $u \subset v$ implies $N_u \subset N_v$ (Definition 1.3 of [8]), we consider here just the k 'maximal vertices' and make no restrictions on the intersections. Since the only requirement on the cardinalities of the M_i is that one be λ, our notion of amalgamation is inherently cumulative; disjoint (λ, k)-amalgamation is not defined.

Definition 2.2.4. *For $k \geq 1$, a k-configuration is a sequence $\overline{M} = \langle M_i : i < k \rangle$ of models (not isomorphism types) from $\boldsymbol{K}$. We say $\overline{M}$ has* power λ *if the cardinality of $\bigcup_{i<k} M_i$ is λ. An* extension *of $\overline{M}$ is any $N \in \boldsymbol{K}$ such that every M_i is a substructure of N.*

Remark 2.2.5. Whether or not a given k-configuration $\overline{M}$ has an extension depends on more than the sequence of isomorphism types of the constituent M_i's, as the pattern of intersections is relevant as well. For example, a 2-configuration $\langle M_0, M_1 \rangle$ with neither contained in the other has an extension if and only if the triple of structures $\langle M_0 \cap M_1, M_0, M_1 \rangle$ has an extension amalgamating them disjointly. Thus our next definition of $(< \lambda, k)$-disjoint amalgamation is very *different* from that in [8, 13] for $k > 2$; they agree for $k = 2$.

Definition 2.2.6. *Fix a cardinal $\lambda = \aleph_\alpha$ for $\alpha \geq -1$. We define the notion of a class $(\boldsymbol{K}, \leq)$ having $(\leq \lambda, k)$-disjoint amalgamation in two steps:*

1. *$(\boldsymbol{K}, \leq)$ has $(\leq \lambda, 0)$-disjoint amalgamation if there is $N \in \boldsymbol{K}$ of power λ;*

2. *For $k \geq 1$, $(\boldsymbol{K}, \leq)$ has $(\leq \lambda, k)$-disjoint amalgamation if it has $(\lambda, 0)$-disjoint amalgamation and every k-configuration $\overline{M}$ of power λ has an extension $N \in \boldsymbol{K}$ such that every M_i is a proper substructure of N.*

For $\lambda \geq \aleph_0$, we define $(< \lambda, k)$-disjoint amalgamation analogously.

Note that $(\leq \lambda, k+1)$-disjoint amalgamation immediately implies $(\leq \lambda, k)$-disjoint amalgamation, as we are allowed to repeat an M_i.

Remark 2.2.7.
- *By employing Fact 2.1.1, we see that if X is any pre-determined set, then if a k-configuration $\overline{M}$ has an extension, then it also has an extension N such that $N \setminus \bigcup \overline{M}$ is disjoint from X.*

- *Thus, $(\leq \lambda, 1)$-disjoint amalgamation asserts that $\mathbf{K}$ has a model of size λ and that every model of size λ has a proper extension.*

Remark 2.2.7 yields the following simplifying lemma.

Lemma 2.2.8. *Assuming, $(\leq \lambda, 1)$-disjoint amalgamation, for $k \geq 2$, in order to inductively establish $(\leq \lambda, k)$-disjoint amalgamation, it suffices to prove that every k-configuration of power λ has an extension.*

Proof. Once we have some extension N, using $(\leq \lambda, 1)$-disjoint amalgamation, we get a proper extension N' of N. $\quad \square_{2.2.8}$

We need two definitions to prove the next proposition.

Definition 2.2.9. *Fix a k-configuration $\overline{M} = \langle M_i : i < k \rangle$.*

1. *A subconfiguration of $\overline{M}$ is a k-configuration $\overline{C} = \langle C_i : i < k \rangle$ such that C_i is a substructure of M_i for each $i < k$.*

2. *A filtration of $\overline{M}$ is a sequence $\langle \overline{C}_\alpha : \alpha < \lambda \rangle$ of subconfigurations of $\overline{M}$ such that*

   (a) *For every $\alpha < \lambda$, $\overline{C}_\alpha = \langle C_i^\alpha : i < k \rangle$ is a subconfiguration of $\overline{M}$ of power less than λ; and*
   (b) *for every $i < k$, the sequence $\langle C_i^\alpha : \alpha < \lambda \rangle$ is a continuous chain of submodels of M_i whose union is M_i.*

We most definitely do not require that $C_i^{\alpha+1}$ properly extend C_i^α! Indeed, if λ is regular and some M_i has power less than λ, then the sequence $\langle C_i^\alpha : \alpha < \lambda \rangle$ will necessarily be constant on a tail of α's.

Proposition 2.2.10. *Suppose $(\mathbf{K}, \leq)$ satisfies Assumption 2.2.2. For all cardinals $\lambda \geq \aleph_0$ and for all $k \in \omega$, if $\mathbf{K}$ has $(< \lambda, k+1)$-disjoint amalgamation, then it also has $(\leq \lambda, k)$-disjoint amalgamation.*

Proof. Fix $\lambda \geq \aleph_0$ and $k \in \omega$. Assume that $\mathbf{K}$ has $(< \lambda, k+1)$-disjoint amalgamation. If $k = 0$, then we construct some $N \in \mathbf{K}$ of power λ as the union of a continuous, increasing chain of models $\langle C_\alpha : \alpha < \lambda \rangle$, where each C_α has power less than λ. So assume $k \geq 1$. From our comments above, it suffices to show that every k-configuration $\overline{M} = \langle M_i : i < k \rangle$ of power λ has an extension.

Claim 1. A filtration of $\overline{M}$ exists.

Proof. As $\mathbf{K}$ is locally finite, the minimal τ-structure A from Assumption 2.2.2 is necessarily finite. So begin the filtration by putting $\overline{C}_0 := \langle C_i^0 : i < k \rangle$, where $C_i^0 = A$ for each $i < k$. By bookkeeping, it suffices to show that every subconfiguration $\overline{C}$ of $\overline{M}$ of power less than λ, and every $a \in \bigcup \overline{M}$, there is a subconfiguration $\overline{C}'$ of $\overline{M}$, $\aleph(\bigcup \overline{C}') = \aleph(\bigcup \overline{C})$, such that C_i is a substructure of C_i' for each $i < k$ and $a \in \bigcup \overline{C}'$.

To see that we can accomplish this, fix $\overline{C}$ and $a \in \bigcup \overline{M}$ as above. Take

$$Y = \{a\} \cup \bigcup \{C_i : i < k\}$$

and, for each $i < k$ let C_i' be the smallest substructure of M_i containing $Y \cap M_i$. Note that since $(\mathbf{K}, \leq)$ is locally finite, $\aleph(C_i') = \aleph(C_i)$ for each $i < k$.

Having proved Claim 1, Fix a filtration $\langle \overline{C}_\alpha : \alpha < \lambda \rangle$ of $\overline{M}$, and let $X = \bigcup \overline{M}$. We recursively construct a continuous chain $\langle D_\alpha : \alpha < \lambda \rangle$ of elements of $\mathbf{K}$ such that

- $\aleph(D_\alpha) = \aleph(\bigcup \overline{C}_\alpha)$; and

- Each D_α is an extension of $\overline{C}_\alpha$ that is disjoint from X over $\overline{C}_\alpha$.

But this is easy. For $\alpha = 0$, use $(< \lambda, k)$-disjoint amalgamation on $\overline{C}_0$ to choose D_0. For $\alpha < \lambda$ a non-zero limit, there is nothing to check (given that $\boldsymbol{K}$ is closed under unions of chains). Finally, suppose $\alpha < \lambda$ and D_α has been constructed. Take $D_{\alpha+1}$ to be an extension of the $(k+1)$-configuration $\overline{C}_{\alpha+1}{}^\frown D_\alpha$. $\square_{2.2.10}$

Remark 2.2.11. Recall that in [17, 3], one obtains a simultaneous uniform filtration of each model in the system being approximated. For infinite successor cardinals, the filtration is obtained by a use of club sets. In approximating countable models by finite ones we don't have such a tool. The technique here was developed to overcome this difficulty. It is, in fact, notably simpler but works only under strong hypotheses, such as Assumption 2.2.2.

Proposition 2.2.10 does not hold for rich models. But the following lemma allows us to construct them in Corollary 2.3.1.

Proposition 2.2.12. *Suppose $(\boldsymbol{K}, \leq)$ satisfies Assumption 2.2.2. Fix a cardinal $\lambda \geq \aleph_0$ and assume that $\boldsymbol{K}$ has $(< \lambda, 2)$-disjoint amalgamation. Then, for any $M, B \in \boldsymbol{K}$ with $\|M\| \leq \lambda$ and $\|B\| < \lambda$, if $M \cap B \in \boldsymbol{K}$, then the 2-configuration $\langle M, B \rangle$ has an extension N of power λ.*

Proof. Let $E = M \cap B$. Choose a filtration $\langle C_\alpha : \alpha < \lambda \rangle$ of M with $C_0 = E$. That is, $\langle C_\alpha : \alpha < \lambda \rangle$ is a continuous chain of substructures of M, each of power less than λ, whose union is M. Then, as in the proof of Proposition 2.2.10, use $(< \lambda, 2)$-disjoint amalgamation and Remark 2.2.7 to construct a continuous chain $\langle D_\alpha : \alpha < \lambda \rangle$, where $D_0 = B$ and, for each $\alpha < \lambda$, $D_{\alpha+1}$ is an extension of $C_{\alpha+1}{}^\frown D_\alpha$ disjoint from M of power $\aleph(D_{\alpha+1}) = \aleph(C_{\alpha+1} \cup D_\alpha)$. Then $N = \bigcup \{D_\alpha : \alpha < \lambda\}$ is an extension of $\langle M, B \rangle$ of power λ. $\square_{2.2.12}$

2.3 Rich and Atomic Models

We use the results from the previous subsection to show the existence of rich and atomic models in various contexts. Here, we need to bound the number of isomorphism types of finite models. Let $\boldsymbol{K}_0$ be a class of finite τ-structures, each of which extends a given minimal A, that is closed under substructures and isomorphisms fixing A pointwise. Then $(\hat{\boldsymbol{K}}, \leq)$, the associated locally finite AEC consisting of all $\boldsymbol{K}_0$-locally finite τ-structures, satisfies Assumption 2.2.2. We let $\mathbf{R}$ denote the subclass of rich models. Recall that $\mathbf{R} = \mathbf{At}$ whenever the class $\boldsymbol{K}_0$ is separable (Definition 2.1.8). We begin with the following corollary to Proposition 2.2.12.

Corollary 2.3.1. *Suppose $(\boldsymbol{K}, \leq)$ satisfies Assumption 2.2.2. Fix $\lambda \geq \aleph_0$. If $\boldsymbol{K}$ has $(< \lambda, 2)$-disjoint amalgamation and has at most λ isomorphism types of finite structures, then*

1. *every $M \in \boldsymbol{K}$ of power λ can be extended to a rich model $N \in \boldsymbol{K}$, which is also of power λ.*

2. *and consequently there is a rich model in λ^+.*

Proof. This follows immediately from Proposition 2.2.12 and bookkeeping. Specifically, given $M \in \boldsymbol{K}$ of power λ, use Proposition 2.2.12 repeatedly to construct a continuous chain $\langle M_i : i < \lambda \rangle$ of elements of $\boldsymbol{K}$, each of size λ. At a given stage $i < \lambda$, focus on a specific finite substructure $A \subseteq M_i$ and a particular

8

finite extension $B \in \boldsymbol{K}$ of A. By Fact 2.1.1, by replacing B by a conjugate copy over A, we may assume $B \cap M_i = A$. Then apply Proposition 2.2.12 to get an extension M_{i+1} of $\langle M_i, B \rangle$ of power λ. As there are only λ constraints, we can organize this construction so that $N = \bigcup \{M_i : i < \lambda\}$ is rich. For 2), iterating this procedure λ^+ times we get a rich model in λ^+. $\quad \square_{2.3.1}$

In particular if, $\boldsymbol{K}_0$ has $(< \aleph_0, 2)$-disjoint amalgamation, then both $\hat{\boldsymbol{K}}$ and $\mathbf{R}$ have models in $\aleph_1$. As an aside, note that *disjoint* amalgamation is essential here. If we take the class of finite linear orders under end-extension, the amalgamation property holds; but, the generic, $(\omega, <)$, is the only model of its Scott sentence.

Theorem 2.3.2. *Suppose $1 \leq r < \omega$ and $\boldsymbol{K}_0$ has the $(< \aleph_0, r+1)$-disjoint amalgamation property. Then for every $0 \leq s \leq r$, $(\hat{\boldsymbol{K}}, \leq)$ has the $(\leq \aleph_s, r-s)$-disjoint amalgamation property. In particular, $\hat{\boldsymbol{K}}$ has models of power $\aleph_r$. Moreover, if there are only countably many isomorphism types in $\boldsymbol{K}_0$, then rich models of power $\aleph_r$ exist and the class $\mathbf{R}$ also has $(\leq \aleph_s, r-s)$-disjoint amalgamation.*

Proof. Fix $1 \leq r < \omega$ and $0 \leq s \leq r$. The first conclusion follows immediately by iterating Proposition 2.2.10 $(s+1)$ times. The existence of a model in $\hat{\boldsymbol{K}}$ of size $\aleph_r$ follows by taking $s = r$ and recalling the definition of $(\aleph_r, 0)$-disjoint amalgamation.

To establish the 'Moreover' statement requires splitting into cases. If $s < r$, then from above, $\hat{\boldsymbol{K}}$ has $(< \aleph_s, r-s+1)$-disjoint amalgamation. As $r-s+1 \geq 2$, it follows from Corollary 2.3.1 that every model $M \in \hat{\boldsymbol{K}}$ of size $\aleph_s$ can be extended to a rich model of size $\aleph_s$. So, given any $r-s$-configuration $\overline{M}$ of rich models of size at most $\aleph_s$, use $(\leq \aleph_s, r-s)$-disjoint amalgamation of $\hat{\boldsymbol{K}}$ to find an extension $M^* \in \hat{\boldsymbol{K}}$ of size $\aleph_s$, and then use Corollary 2.3.1 to extend M^* to a rich model N that is also of size $\aleph_s$. Finally, if $s = r$, then as $r - s = 0$, we need only construct a rich model M of size $\aleph_r$. Apply Corollary 2.3.1.2. $\quad \square_{2.3.2}$

We close this section with a result that is typical of excellent classes.

Theorem 2.3.3. *If $\boldsymbol{K}_0$ has $(< \aleph_0, k)$-disjoint amalgamation for each $k \in \omega$, then both $\hat{\boldsymbol{K}}$ and the subclass $\mathbf{R}$ of rich models have arbitrarily large models.*

Proof. By induction on α, use Proposition 2.2.10 to show that $\boldsymbol{K}$ has $(< \aleph_\alpha, k)$-disjoint amalgamation for each $k \in \omega$. From this, it follows immediately that $\hat{\boldsymbol{K}}$ has models in every infinite cardinality. If we let κ denote the number of isomorphism types of finite structures in $\boldsymbol{K}$, then it follows from Corollary 2.3.1 that every model $M \in \hat{\boldsymbol{K}}$ of power at least κ has a rich extension N of the same cardinality. $\quad \square_{2.3.3}$

Remark 2.3.4. *In both Theorem 2.3.2 and Theorem 2.3.3, if the subclass $\boldsymbol{K}_0$ of $\hat{\boldsymbol{K}}$ consisting of all finite structures is closed under substructure, is separable and satisfies JEP, then, simply because $\mathbf{R} = \mathbf{At}$ by Lemma 2.1.9, we get the same spectrum of $(< \lambda, k)$-disjoint amalgamation for $\mathbf{At}$ as we had for $\mathbf{R}$.*

2.4   Homogeneous characterizability

Hjorth's characterization of each $\aleph_\alpha$ for $\alpha < \omega_1$ depended heavily on a notion, *homogeneous characterizability* introduced by Baumgartner and Malitz and resurrected in [22].

Definition 2.4.1. *Let τ be a vocabulary with a distinguished unary predicate V. A complete $L_{\omega_1, \omega}$-sentence ϕ_κ homogeneously characterizes a cardinal κ if:*

1. In the unique countable model $M \models \phi_\kappa$, $V(M)$ is an infinite set of absolute indiscernibles *(i.e., every permutation of $V(M)$ extends to an automorphism of M); and*

2. ϕ_κ has a model M of size κ with $|V(M)| = \kappa$ but no larger model.

One way of generating complete sentences that homogeneously characterize a cardinal is to start with an appropriate class of finite τ-structures $\mathbf{K}_0^h$ that has a generic model M and take ϕ to be the Scott sentence of M. In [6], a general method for producing such classes of finite structures was introduced.

Construction 2.4.2. *Suppose $\mathbf{K}_0$ is a class of finite structures in a vocabulary τ that is closed under isomorphism and substructure, has countably many isomorphism types, satisfies JEP, and at least $(< \aleph_0, 2)$- disjoint amalgamation. Given such a class, extend the vocabulary to $\tau_h = \tau \cup \{U, V, p\}$, where U, V are unary predicates and p is a unary function symbol, none of which appear in τ. We define an associated class $\mathbf{K}_0^h$ to be all finite τ_h-structures A^h such that*

- *The vocabulary of τ is defined on $U(A^h)$;*

- *The reduct of $U(A^h)$ to τ is an element of $\mathbf{K}_0$; and*

- *$p : U(A^h) \to V(A^h)$ is a projection[5].*

In [6] it was shown that this induced class satisfies JEP and $(< \aleph_0, 2)$-disjoint amalgamation; hence has a generic model M^h; and that $V(M^h)$ is an infinite set of absolute indiscernibles. As in [6], one constructs further examples by 'merging' two theories on the set of absolute indiscernibles; applications requiring the ability to fix the cardinality of the set of absolute indiscernibles appear in [7, 9]. We say N is an $(\aleph_r, \aleph_k)$- model if $\|N\| = \aleph_r$ and $|V(N)| = \aleph_k$. In order to prove the set of amalgamations attains power $\aleph_r$ in a *rich* model we need a further variant of amalgamation.

Definition 2.4.3. *Let $(\mathbf{K}, \leq)$ be a class of structures satisfying Assumption 2.2.2. Given a cardinal λ and $k \in \omega$, we say that $\mathbf{K}$ has* frugal $(\leq \lambda, k)$-disjoint amalgamation *if it has $(\leq \lambda, k)$-disjoint amalgamation and, when $k \geq 2$, every k-configuration $\langle M_i : i < k \rangle$ has an extension $N \in \mathbf{K}$ with universe $\bigcup_{i<k} M_i$.*

The following result is the main goal of this section.

Theorem 2.4.4. *Suppose $r \geq 1$ and a class of finite structures $(\mathbf{K}_0, \leq)$ is separable, satisfies JEP, and has frugal $(< \aleph_0, r + 1)$-disjoint amalgamation. The associated class of models $\mathbf{K}_0^h$ has a countable generic model M that has a set V of absolute indiscernibles. For every $0 \leq s \leq r$, there is a* maximal $(\aleph_r, \aleph_s)$- *model N of the Scott sentence of M.*

The key idea is that disjoint amalgamation allows us to push up the cardinality while in $\mathbf{K}^h$ frugal amalgamation allows us to do so while fixing the set V of indiscernibles.

Definition 2.4.5. *Suppose $(\mathbf{K}, \leq)$ satisfies Assumption 2.2.2 and τ has a distinguished predicate V. $\mathbf{K}$ has V-conservative amalgamation in λ if for any $M \in \mathbf{K}$ of cardinality λ and finite $A \subset B \in \mathbf{K}$, with $A \subset M$ there is an $N \in \mathbf{K}$ containing M and an embedding g of B into N such that $N = M \cup [\operatorname{im} g]$ and $g{\restriction}V(B) \subseteq V(M)$.*

The frugality is crucial in the next lemma. The indiscernability of the elements of V allows us to map $V(B) - M$ into $V(M)$. The frugality ensures that the amalgamation does not push other elements into V.

Lemma 2.4.6. *If $\mathbf{K}$ has frugal $(\leq \lambda)$-disjoint amalgamation, then $\mathbf{K}^h$ has V-conservative amalgamation.*

---

[5]We do not insist in the finite models that p is onto but it will be onto in the generic model.

Proof. Let $M \in \boldsymbol{K}$ have cardinality λ and finite $A \subset B$ be in $\boldsymbol{K}$, with $A \subset M$. Let f be an arbitrary 1-1 map from $V(B) - M$ into $V(M) - V(A)$. Let A' be the structure imposed on $A \cup [\text{im } f]$ by M and let B' be the structure with universe $A \cup [\text{im } f] \cup [B - (M \cup V(B))]$, and the natural structure making $B \approx B'$ over the isomorphism of A' and A. Now N, the frugal amalgamation of M and B' over A', is the required V-conservative amalgam of M and B, as witnessed by $g = f \cup [\text{id} \upharpoonright (B - (V(B) - M))]$. $\square_{2.4.6}$

Proof of Theorem 2.4.4: For $s \leq r$ we can construct an $(\aleph_s, \aleph_s)$-model M of $\boldsymbol{K}^h$ by Lemma 2.3.2. To extend it to a rich $(\aleph_r, \aleph_s)$-model, do the same construction as in Lemma 2.3.1 but use the V-conservative 2-amalgamation guaranteed by Lemma 2.4.6. $\square_{2.4.4}$

3   Characterizing $\aleph_r$ and the amalgamation spectra

In [14] Laskowski-Shelah constructed by a Fraïssé construction, which is easily seen to satisfy disjoint amalgamation in the finite, a complete sentence ϕ_{LS} in $L_{\omega_1,\omega}$ such that every model in $\aleph_1$ is maximal and so characterizes $\aleph_1$. The idea was to exhibit a countable family of binary functions that give rise to a locally finite closure relation, and then to note that any locally finite closure relation on a set of size $\aleph_2$ must have an independent subset of size three. They also asserted that they had a similar construction to characterize all cardinals up to $\aleph_\omega$, but there was a mistake in the proof.[6] We remedy that now.

3.1   Constructing Atomic Models

Now we apply the general methods of Section 2 to a construct a particular family of examples. For a fixed $r \geq 1$, let τ_r be the (countable) vocabulary consisting of countably many $(r + 1)$-ary functions f_n and countably many $(r + 1)$-ary relations R_n. We consider the class $\boldsymbol{K}_0^r$ of finite τ_r-structures (including the empty structure) that satisfy the following three sentences of $L_{\omega_1,\omega}$:

- The relations $\{R_n : n \in \omega\}$ partition the $(r + 1)$-tuples;

- For every $(r + 1)$-tuple $\boldsymbol{a} = (a_0, \ldots, a_r)$, if $R_n(\boldsymbol{a})$ holds, then $f_m(\boldsymbol{a}) = a_0$ for every $m \geq n$;

- There is no independent subset of size $r + 2$.

The third condition refers to the closure relation on a τ_r-structure M defined by iteratively applying the functions $\{f_n\}$, i.e., for every subset $A \subseteq M$, $\text{cl}(A)$ is the smallest substructure of M containing A. A set B is *independent* if, for every $b \in B$, $b \notin \text{cl}(B \setminus \{b\})$.

As we allow the empty structure, which is obviously minimal and is a substructure of every element of $\boldsymbol{K}_0^r$, the class $(\hat{\boldsymbol{K}}^r, \leq)$ consisting of all $\boldsymbol{K}_0^r$-locally finite τ_r-structures satisfies Assumption 2.2.2. The disjointness of the predicates R_n and the triviality condition on f_m for $m \geq n$ implies that a model in $\boldsymbol{K}_0^r$ of cardinality n has only finitely many possible isomorphism types once the R_n have been specified. By the observation after Definition 2.1.8, $\boldsymbol{K}_0^r$ is separable.

We begin by studying the (frugal disjoint amalgamation spectrum of the class $\boldsymbol{K}_0^r$ of finite structures. Then we use the general properties established in Section 2 to construct models of $\hat{\boldsymbol{K}}^r$ in certain cardinals. We then recall the combinatorial Fact 3.2.1 and deduce from it the negative results: non-existence and failure of amalgamation in certain cardinals.

---

[6]Specifically, the closure relation defined in Lemma 0.7 of [14] is not locally finite.

Theorem 3.1.1. *For each $r \geq 1$, $\hat{\boldsymbol{K}}^r$ has frugal $(< \aleph_0, r+1)$-disjoint amalgamation. Further, $\hat{\boldsymbol{K}}^r$ does not have $(< \aleph_0, r+2)$-disjoint amalgamation.*

Proof. Fix an $(r+1)$-configuration $\overline{M} = \langle M_i : i \leq r \rangle$ of finite structures from $\boldsymbol{K}_0^r$. Let $D = \bigcup \overline{M}$ and let $\langle d_i : i < t \rangle$ enumerate the set D, where $t = |D|$. Call an $(r+1)$-tuple $\mathbf{b} \in D^{r+1}$ *unspecified* if $\mathbf{b} \notin M_j^{r+1}$ for any $j \leq r$. If our amalgam N is to have universe D and have each M_j be a substructure, we need only define the functions f_n and the relations R_n on unspecified tuples.

For every unspecified $\mathbf{b} = (b_0, \ldots, b_r)$, put $R_t^N(\mathbf{b})$ and define

$$f_i^N(\mathbf{b}) = \begin{cases} d_i & \text{if } i < t \\ b_0 & \text{otherwise} \end{cases}$$

To see that $N \in \boldsymbol{K}_0^r$, we show that there are no independent subsets of size $r+2$. Choose any $B \subseteq N$ of size $r+2$. If $B \subseteq M_j$ for some j, then B is not independent since $N_j \in \boldsymbol{K}_0^r$. However, if B is not contained in any M_j, then there is an unspecified $\mathbf{b} \in B^{r+1}$. In this case, B is contained in the closure of $\mathbf{b}$, so again B is not independent.

For the second sentence, start with any family $\{C_i : i < r+2\}$ of $(r+2)$ elements of $\boldsymbol{K}_0^r$, each of whose universes are non-empty, but pairwise disjoint. For each $i < r+2$, choose an element $c_i \in C_i$. Note that $c_i \notin C_j$ for any $j \neq i$.

Now, for each $i < r+2$, let $N_i \in \boldsymbol{K}_0^r$ be a frugal amalgam of the $(r+1)$-configuration $\langle C_j : j \neq i \rangle$. Note that for each i, $c_i \in N_j$ if and only if $j \neq i$. We argue that the $(r+2)$-configuration $\langle N_i : i < r+2 \rangle$ has no disjoint amalgam in $\boldsymbol{K}_0^r$. To see this, let M be any τ_r structure having each N_i as a substructure. We argue that the set $C^* = \{c_i : i < r+2\}$ is independent in M. Indeed, fix any $i < r+2$. From above, the $(r+1)$-element set $C^* \setminus \{i\}$ is a subset of N_i, while $c_i \notin N_i$. As N_i is a substructure of M, it follows that $c_i \notin \mathrm{cl}_M(C^* \setminus \{c_i\})$. Thus, C^* is independent in M. $\square_{3.1.1}$

Corollary 3.1.2. *For each $r \geq 1$, the class $\boldsymbol{K}_0^r$ has a Fraïssé limit M^r. It is generic, and is an element of $\hat{\boldsymbol{K}}^r$. It is the unique countable atomic model of its first order theory. Moreover, a τ_r-structure N is in $\mathbf{At}^r$ if and only if (1) N is $\boldsymbol{K}_0^r$-locally finite; and (2) N is $\boldsymbol{K}_0^r$-rich.*

As notation, for each $r \geq 1$, let T^r denote the first-order theory of the $\boldsymbol{K}_0^r$-generic M^r, and let ϕ_r denote the Scott sentence of M^r.

Now we apply Theorem 2.3.2 to the example at hand.

Theorem 3.1.3. *For each r with $1 \leq r < \omega$ and $0 \leq s \leq r$, $\hat{\boldsymbol{K}}^r$ has frugal $(\aleph_s, r-s)$-disjoint amalgamation in $\aleph_s$. In particular, there are atomic models of T^r of cardinality $\aleph_s$ for $s \leq r$.*

3.2 Bounding the Cardinality and Blocking Amalgamation

We use a slight modification of Lemma 2.3 of [6] to obtain our negative results. The proof of the following Fact is included as a convenience for the reader.

Fact 3.2.1. *For every $k \in \omega$, if cl is a locally finite closure relation on a set X of size $\aleph_k$, then there is an independent subset of size $k+1$.*

Proof. By induction on k. When $k = 0$, take any singleton not included in $\mathrm{cl}(\emptyset)$. Assuming the Fact for k, given any locally finite closure relation cl on a set X of size $\aleph_{k+1}$, fix a cl-closed subset $Y \subseteq X$ of size $\aleph_k$ and choose any $a \in X \setminus Y$. Define a locally finite closure relation cl_a on Y by $\mathrm{cl}_a(Z) = \mathrm{cl}(Z \cup \{a\}) \cap Y$. It is easily checked that if $B \subseteq Y$ is cl_a-independent, then $B \cup \{a\}$ is cl-independent. $\square_{3.2.1}$

Lemma 3.2.2. *Suppose each model of $\hat{K}$ admits a locally finite closure relation* cl *such that there is no* cl-*independent subset of size $r + 2$. Then $\hat{K}$ has only maximal models in $\aleph_r$ and so 2-amalgamation is trivially true in $\aleph_r$.*

Proof. Fix $M \in \hat{K}$ of size $\aleph_r$ and assume by way of contradiction that it had a proper extension $N \in \hat{K}$. As in the proof of Fact 3.2.1, fix $a \in N \setminus M$ and define a closure relation cl_a on M by $\text{cl}_a(Z) = \text{cl}(Z \cup \{a\}) \cap M$. As cl_a is locally finite, it follows from Fact 3.2.1 that there is a cl_a-independent subset $B \subseteq M$ of size $r + 1$. But then, $B \cup \{a\}$ would be a cl-independent subset of N of size $r + 2$, contradicting $N \in \hat{K}$. $\square_{3.2.2}$

In this section, we say that a class $(K, \leq)$ has $(\lambda, 2)$-(disjoint) amalgamation if any triple of distinct models of cardinality λ can be (disjointly) amalgamated. This usage is distinct from the earlier use in this paper which because of its cumulative nature gave no meaning to amalgamation in a fixed cardinal. Thus, trivially, if all models in λ are maximal $(\lambda, 2)$-(disjoint) amalgamation holds. We write 2-amalgamation for the usual notion which allows identifications. Thus difficult parts in this section are to show that amalgamation allowing identifications is impossible.

Part 1) of the following lemma follows from Lemma 3.2.2. The second part depends on the special properties of the current example that we exploit in the remainder of this section.

Lemma 3.2.3. *For every $r \geq 1$, $\hat{K}^r$*

1. *has only maximal models in $\aleph_r$ and so 2-amalgamation is trivially true in $\aleph_r$;*

2. *fails 2-amalgamation in $\aleph_{r-1}$.*

Proof. Part 1) is Lemma 3.2.2. For the second part, we first modify this argument to show that an amalgamation problem (M_0, M_1, M_2) for elements from $\hat{K}^r$ of size $\aleph_{r-1}$ with M_0 a substructure of both M_1 and M_2 is solvable if and only if M_i is embeddable into M_{3-i} over M_0 for either $i = 1$ or 2. To see the non-trivial direction, suppose a triple (M_0, M_1, M_2) are chosen as above and there is $M_3 \in \hat{K}^r$ and embeddings $f : M_1 \to M_3$ and $g : M_2 \to M_3$, each over M_0, with elements $f(a_1) \notin g(M_2)$ and $g(a_2) \notin f(M_1)$. We obtain a contradiction by considering the closure relation on M_0 defined by

$$\text{cl}^*(Z) = \text{cl}(Z \cup \{f(a_1), g(a_2)\}) \cap M_0$$

with the latter closure relation cl computed in M_3. As cl^* is locally finite, it follows from Fact 3.2.1 that there is a cl^*-independent subset $B \subseteq M_0$ of size r. We obtain a contradiction to $M_3 \in \hat{K}^r$ by showing that the set $B \cup \{f(a_1), g(a_2)\}$ is cl-independent. To see this, first note that since $B \cup \{f(a_1)\} \subseteq f(M_1)$, $g(a_2) \notin \text{cl}(B \cup \{f(a_1)\})$. Similarly, $f(a_1) \notin \text{cl}(B \cup \{g(a_2)\})$. But, for any $b \in B$, the cl^*-independence of B implies that $b \notin \text{cl}((B \setminus \{b\} \cup \{f(a_1), g(a_2)\})$, completing the argument characterizing which amalgamation problems are solvable.

To complete part 2) we must establish:

(*) There exists a triple (N_0, N_1, N_2) consisting of elements of $\hat{K}^r$, each of size $\aleph_{r-2}$ such that for $i = 1, 2$, each of N_{3-i} contains an element whose type over N_0 is not realized in the other.

In fact (*) holds in every cardinal below $\aleph_{r-1}$; the exact cardinality $\aleph_{r-1}$ was used to apply Fact 3.2.1 in the reduction to (*). Choose $M \in \hat{K}^r$ of cardinality $\aleph_{r-1}$ and two finite structures $M_1, M_2 \in \hat{K}^r$ such that $M_1 \cap M = M_2 \cap M = N_0$ but M_1 and M_2 are not isomorphic over N_0. Now disjointly amalgamate each M_i with M over N_0 to obtain $\hat{M}_i$. The resulting triple $(N_0, \hat{M}_1, \hat{M}_2)$ is as required since the disjoint amalgamation guarantees that the type of N_i/N_0 is not realized in $\hat{M}_{1-i}$. $\square_{3.2.3}$

13

Part 1) of Lemma 3.2.3 extends immediately to structures in $\mathbf{At}^r$. We now give the considerably more involved argument that $\mathbf{At}^r$ fails 2-amalgamation in $\aleph_{r-1}$. The difficulty is that we must establish (*) for $\mathbf{At}^r$ rather than $\hat{\boldsymbol{K}}^r$. For most of this discussion, we work with the class $\hat{\boldsymbol{K}}^r$, passing to $\mathbf{At}^r$ only at the end. We require the following new notion.

Definition 3.2.4. *Fix any infinite $U \subseteq \omega$. For any $s \leq r - 2$, and any s-system $\langle N_u : u \in \mathcal{P}^-(s)\rangle$ from $\hat{\boldsymbol{K}}^r$, a U-amalgam is any frugal, disjoint amalgam M with the additional property that for every unspecified $(r+1)$-tuple $\mathbf{b}$, if $M \models R_t(\mathbf{b})$, then $t \in U$. (As in the proof of Theorem 3.1.1,* unspecified *means here that $\mathbf{b}$ is not contained in any element of the system.)*

Lemma 3.2.5. *For any infinite $U \subseteq \omega$, Theorems 3.1.1 and 3.1.3 go through word for word, using 'U-amalgamation' in place of 'frugal, disjoint amalgamation'. In particular, by Theorem 3.1.3, any triple (N_0, N_1, N_2) of elements of $\hat{\boldsymbol{K}}^r$ with $N_1 \cap N_2 = N_0$, each of size at most $\aleph_{r-2}$ has a U-amalgamation.*

Definition 3.2.6. *Suppose $N \in \hat{\boldsymbol{K}}^r$, $B \subseteq N$, and $a \in N \setminus B$. A relevant $(r+1)$-tuple is an element of $(B \cup \{a\})^{r+1}$ with exactly one occurrence of a. We define the* species *of a over B in N.*

$$\mathrm{sp}_N(a/B) = \{t \in \omega : N \models R_t(\mathbf{c}) \text{ for some relevant } \mathbf{c}\}$$

The subscript N is necessary since while the species of a/B is a property of sequences from $B \cup \{a\}$, the predicates are specified on those elements only by the model N. The next few lemmas describe basic constructions with U-amalgamation. The definition of U-amalgamation allows us to 'extend the base'.

Lemma 3.2.7. *Fix any infinite subset $U \subseteq \omega$. For any $N_0 \subseteq N_1$ and $N_0 \subseteq N_0'$, all from $\hat{\boldsymbol{K}}^r$, with $N_1 \cap N_0' = N_0$, every U-amalgam $N_1' \in \hat{\boldsymbol{K}}^r$ satisfies for every $a \in N_1 \setminus N_0$:*

$$\mathrm{sp}_{N_1'}(a/N_0') \setminus \mathrm{sp}_{N_1}(a/N_0) \subseteq U.$$

Definition 3.2.8. *Let $M \subseteq M_1 \subseteq M_2$ each be in $\mathbf{R}$ and $V \subseteq \omega$. We write $M_1 \prec_{M,V} M_2$ if for each $a \in M_2 - M_1$, $\mathrm{sp}_{M_2}(a/M) \cap V$ is finite.*

Lemma 3.2.9. *Fix infinite disjoint $U, V \subseteq \omega$. Suppose $N_1 \in \hat{\boldsymbol{K}}^r$ is of size at most $\aleph_{r-2}$, $M \subseteq M' \in \boldsymbol{K}^r$ (hence finite) with $M' \cap N_1 = M$. Then any U-amalgam N_2 of M' and N_1 over M satisfies: $N_1 \prec_{M,V} N_2$.*

Proof. If $a \in N_1$, there are no new relevant tuples to be assigned values. Any $a \in N_2 \setminus N_1$ is in $M' \setminus M$. The only sequences that are not assigned a value from U are from $M \cup \{a\} \subseteq M'$, and there are only finitely many of them. $\square_{3.2.9}$

Lemma 3.2.10. *Fix an infinite-coinfinite $U \subset \omega$. Suppose $M_0 \subseteq M_1$ are in $\hat{K}^r$ and have cardinality at most $\aleph_{r-2}$. Then then there is an $M^* \in \mathbf{At}^r$ such that for any $V \subset \omega$ that is disjoint from U, $M_1 \prec_{M_0,V} M^*$*

Proof. We obtain M^* by modifying the induction of the second paragraph of the proof of Theorem 3.1.3 to build a sequence of models M_α with union M^*. At each stage α, U-amalgamate a finite $A \in K^r$ with M_α over $A \cap M_\alpha$. Then for any $a \in M^*$, it first appears in some M_α and the species of a/M_0 in M_β is defined only for $\beta \geq \alpha$. Further, for any $\beta > \alpha$, $\mathrm{sp}_{M^*}(a/M_0) = \mathrm{sp}_{M_\beta}(a/M_0) = \mathrm{sp}_{M_\alpha}(a/M_0)$ and $\mathrm{sp}_{M_\alpha}(a/M) \cap V$ is finite. It is easy to organize the construction so $M^* \in \mathbf{At}^r$. $\square_{3.2.10}$

We now define the notion of obstruction (to amalgamation). Of course, we can overcome that obstruction in cardinals at most $\aleph_{r-2}$. But we will construct an obstruction of cardinality $\aleph_{r-1}$ which by Lemma 3.2.3 cannot be overcome.

14

Definition 3.2.11. Fix pairwise disjoint, infinite $U, V, W \subseteq \omega$. An *obstruction* is a triple (N_0, N_1, N_2) consisting of elements of $\hat{\boldsymbol{K}}^r$, such that $N_1 \cap N_2 = N_0$ and

- There is exactly one $a \in N_1 \setminus N_0$ such that $\mathrm{sp}_{N_1}(a/N_0)$ has infinite intersection with V and $N_0 \prec_{N_0, V} N_2$. (That is, for any $b \in N_2 \setminus N_0$, $\mathrm{sp}_{N_1}(b/N_0) \cap V$ is finite.)

- There is exactly one $b \in N_2 \setminus N_1$ that $\mathrm{sp}_{N_2}(b/N_0)$ has infinite intersection with W and $N_0 \prec_{N_0, W} N_1$. (That is, and for any $a \in N_1 \setminus N_2$, $\mathrm{sp}_{N_1}(a/N_0) \cap W$ is finite.)

An obstruction is said to have cardinality κ if each N_i has cardinality κ. An *extension* of an obstruction (N_0, N_1, N_2) is an obstruction (N_0', N_1', N_2') satisfying $N_i \subseteq N_i'$ and $N_i' \cap (N_1 \cup N_2) = N_i$ for each i. An extension is *proper* if all three of the structures increase. An *atomic obstruction* is an obstruction (M_0, M_1, M_2) in which every M_i is in $\mathbf{At}^r$.

'Obstruction' has a hidden parameter: 'with respect to disjoint V, W'. In order to work with the notion we need a third set U disjoint from each of them. **We will keep fixed disjoint, infinite subsets U, V, W of ω in the following construction.** Now we will create obstructions.

Lemma 3.2.12. *For any $N_0 \in \hat{\boldsymbol{K}}^r$ of size $\aleph_0$, there is a 1-point extension $N_1 \in \hat{\boldsymbol{K}}^r$ with universe $N_0 \cup \{a\}$ such that $\mathrm{sp}_{N_1}(a/N_0)$ is an infinite subset of V.*

Proof. Fix any nested sequence $\langle M_k : k \in \omega \rangle$ of elements of $\boldsymbol{K}^r$ (hence finite) such that $N_0 = \bigcup_k M_k$. Inductively assume that we have constructed an element $M_{k,a}$ of $\boldsymbol{K}$ with universe $M_k \cup \{a\}$ satisfying $\mathrm{sp}_{M_{k,a}}(a/M_k) \subseteq V$ and $|\mathrm{sp}_{M_{k,a}}(a/M_k)| \geq k$. Fix any element $t^* \in V$ that is larger than both $\max\{\mathrm{sp}_{M_{k,a}}(a/M_k)\}$ and $||M_{k+1}||$. Then construct an extension with universe $M_{k+1} \cup \{a\}$ using the procedure in the proof of Theorem 3.1.1, saying that $R_{t^*}(\mathbf{b})$ holds for every unspecified $(r+1)$-tuple. $\square_{3.2.12}$

Lemma 3.2.13. *For every $N \in \hat{\boldsymbol{K}}^r$ of size $\aleph_0$, there is a 2-point extension N^* with universe $N \cup \{a, b\}$ such that*

- $\mathrm{sp}_{N^*}(a/N)$ *is an infinite subset of V; and*

- $\mathrm{sp}_{N^*}(b/N)$ *is an infinite subset of W.*

Thus there is a countable obstruction in $\hat{K}^r$.

Proof. Apply Lemma 3.2.12 with V and W to create two one point extensions and then use U-amalgamation to construct N^*. The required obstruction is (N, Na, Nb). $\square_{3.2.13}$

Lemma 3.2.14. *Suppose $r \geq 2$. Any obstruction (N_0, N_1, N_2) of cardinality $\leq \aleph_{r-2}$ has a proper atomic extension.*

Proof. Let (N_0, N_1, N_2) be an obstruction with U-amalgam N^*. Choose by Theorem 3.1.3 a proper extension $N_4 \in \mathbf{At}^r$ of N_0 such that $N_4 \cap N^* = N_0$. Let N_7 be a U-amalgam of N_4 and N^*; it contains U-amalgams N_5 of N_4 and N_1 and N_6 of N_4 and N_2. By Lemma 3.2.7, (N_4, N_5, N_6) is an obstruction. Apply Lemma 3.2.10 twice to construct $N_5', N_6' \in \mathbf{At}^r$ such that $N_5 \prec_{N_4, V} N_5'$ and $N_6 \prec_{N_4, W} N_6'$. Then, (N_4, N_5', N_6') is as required. $\square_{3.2.14}$

Lemma 3.2.15. *For $r \geq 1$ there is an atomic obstruction (M_0, M_1, M_2) of cardinality $\aleph_{r-1}$.*

Proof. Build by Lemma 3.2.14 an increasing $\aleph_{r-1}$ chain of atomic obstructions. To get started, use Lemma 3.2.13. $\square_{3.2.15}$

Proposition 3.2.16. *For any $r \geq 1$, $\mathbf{At}^r$ fails amalgamation in $\aleph_{r-1}$.*

Proof. By Lemma 3.2.15 when $r \geq 2$, we have an atomic obstruction (M_0, M_1, M_2) of cardinality $\aleph_{r-1}$. For such an obstruction, M_i does not embed into M_{3-i} over M_0 for $i = 1, 2$, so by the argument justifying the second sentence of 3.2.3, this triple cannot be amalgamated into any element of $\hat{\boldsymbol{K}}^r$, much less an element of $\mathbf{At}^r$. $\square_{3.2.16}$

We close this section by counting the number of models of $\mathbf{At}^r$ in the uncountable cardinals up to $\aleph_r$. The proofs will be by induction on $1 \leq s \leq r$. The inductive steps are routine, but the base case, proving that there are $2^{\aleph_1}$ non-isomorphic models in $\mathbf{At}^r$ of size $\aleph_1$ requires the main theorem of [1] whose statement requires two ancillary definitions.

Definition 3.2.17. *Suppose T is a complete theory in a countable language and let $\mathbf{At}$ denote its class of atomic models. Fix $M \in \mathbf{At}$ and $\boldsymbol{a}, b$ from M. Then:*

- *A complete formula $\phi(x, \boldsymbol{a})$ is* pseudo-algebraic *in M if $\boldsymbol{a}$ is from M, and $\phi(N, \boldsymbol{a}) = \phi(M, \boldsymbol{a})$ for every countable $N \in \mathbf{At}$ with $N \succeq M$.*

- *We write $b \in \mathrm{pcl}(\boldsymbol{a}, M)$ if $\mathrm{tp}(b/\boldsymbol{a})$ contains a pseudo-algebraic formula in M.*

- *A complete formula $\phi(x, \boldsymbol{a})$ is* pseudo-minimal *if it is not pseudo-algebraic, but for every $\boldsymbol{a}^* \supseteq \boldsymbol{a}$ and c from M and for every $b \in \phi(M, \boldsymbol{a})$, if $c \in \mathrm{pcl}(\boldsymbol{a}^* b, M)$ but $c \notin \mathrm{pcl}(\boldsymbol{a}^*, M)$, then $b \in \mathrm{pcl}(\boldsymbol{a}^* c, M)$.*

- *The class $\mathbf{At}$ has* density of pseudo-minimal types *if for some/every $M \in \mathbf{At}$, for every non-pseudo-algebraic formula $\phi(x, \boldsymbol{a})$, there is $\boldsymbol{a}^* \supseteq \boldsymbol{a}$ from M and a pseudo-minimal formula $\psi(x, \boldsymbol{a}^*)$ such that $\psi(x, \boldsymbol{a}^*) \vdash \phi(x, \boldsymbol{a})$.*

Theorem 3.2.18 ([1]). *Let T be any complete first-order theory in a countable language with an atomic model that is not minimal. If the pseudo-minimal types are not dense, then there are $2^{\aleph_1}$ pairwise non-isomorphic atomic models of T, each of size $\aleph_1$.*

It is obvious that if there is no pseudominimal formula then density of pseudo-minimal types fails, and hence there are many non-isomorphic atomic models of size $\aleph_1$. Lemma 3.2.19 shows this for each of the examples $\mathbf{At^r}$. Indeed, in these examples the notion of pseudo-algebraicity can be greatly simplified.

Lemma 3.2.19. *We need two facts about $\mathbf{At^r}$.*

1. *For any $M \in \mathbf{At}^r$ and for any a, B from M, $a \in \mathrm{pcl}(B)$ iff $a \in \mathrm{acl}(B)$ iff $a \in \mathrm{dcl}(B)$.*

2. *There are no pseudo-minimal formulas for $\mathbf{At^r}$.*

Proof. 1) Note first that $a \in \mathrm{dcl}(B)$ iff $a = \tau(\mathbf{b})$ where τ is a composition of our functions f_n and $\mathbf{b}$ is from B. Now all the right to left implications are obvious.

For the other direction, first note we may assume B is finite. Thus, the smallest substructure of M containing B is finite. By replacing B by the substructure generated by B, we may assume that B itself is a finite substructure of M. So, if $a \notin Bs$, then by countably many disjoint amalgamations we can find a model $N \in \mathbf{At^r}$ contain containing B, but not a. This N witnesses that a $\notin \mathrm{pcl}(B)$.

2) Fix a non-pseudo algebraic formula $\psi(x, \mathbf{b})$; by adding parameters we may assume that $\mathbf{b}$ enumerates a finite substructure B of M and, as it is not pseudo-algebraic, by (1), there is an $a \in (M - B)$ witnessing

$\psi(x, \mathbf{b})$. To see that the formula is not pseudo-minimal, we need only find two elements $c, d \in M$ and find a finite structure such that $d \in \mathrm{dcl}(Bca)$, but $a \notin \mathrm{dcl}(Bdc)$. This is easy to accomplish in $\boldsymbol{K}_0$ and since M is a model of the theory of the generic, we can choose the elements in M. $\quad\square_{3.2.19}$

Now we sum up the properties of the example. Collectively, these results give a correct proof of Theorem 0.6 of [14] and improve [12] by giving an explicit sentence characterizing each $\aleph_r$ for $r < \omega$.

Theorem 3.2.20. *For every $r \geq 1$, the class $\mathbf{At^r}$ satisfies:*

1. *there is a model of size $\aleph_r$, but no larger models;*

2. *every model of size $\aleph_r$ is maximal, and so 2-amalgamation is trivially true in $\aleph_r$;*

3. *disjoint 2-amalgamation holds up to $\aleph_{r-2}$;*

4. *2-ap fails in $\aleph_{r-1}$.*

5. *Each of the classes $\hat{\boldsymbol{K}}^r$ and $\mathbf{At}^r$ have $2^{\aleph_s}$ models in $\aleph_s$ for $1 \leq s \leq r$. In addition, $\hat{\boldsymbol{K}}^r$ has $2^{\aleph_0}$ models in $\aleph_0$.*

Proof. Parts 1-4 are immediate from Theorem 3.1.3, Lemma 3.2.3, and Proposition 3.2.16. We sketch the proof of part 5. First, by relativizing the basic construction of Definition 3.2.5 to permit indices on the R_n only for n in a fixed infinite $W \subset \omega$ gives a situation that differs from the original only by notation. The generic model has tuples satisfying R_n if and only if $n \in W$. So any family of distinct sets W_i for $i < \omega$ gives continuum many countable models of $\hat{\boldsymbol{K}}^r$.

Next, from Theorem 3.2.18 and Lemma 3.2.19, we see that $\mathbf{At^r}$ and *a fortiori* $\hat{\boldsymbol{K}}^r$ has $2^{\aleph_1}$ models in $\aleph_1$. Now the techniques for violating amalgamation in Section 3 and induction on s using properties of U-amalgamation show $\hat{\boldsymbol{K}}^r$ has $2^{\aleph_s}$ models in $\aleph_s$ for $1 \leq s \leq r$. Finally, when $s > 1$ a similar argument allows one to choose non-isomorphic $M, N \in \hat{\boldsymbol{K}}^r$ to non-isomorphic models in $\mathbf{At^r}$. The last two arguments rely on the notion of a V-component: A V-component E of $a \in N$ is a maximal uncountable subset of N such that for any tuple $\mathbf{e} \in E$ if every permutation $\mathbf{f}$ of $\mathbf{e}a$ satisfies $R_t(\mathbf{f})$ then $t \in V$. $\quad\square_{3.2.20}$

4   Context and Conclusions

Spectrum functions are investigated along three axes: the spectrum might be of existence, amalgamation, joint embedding, maximal models etc.; the class might be defined as an AEC, a (complete) sentence of $L_{\omega_1,\omega}$, etc.; the result may be in ZFC or not. We place our work in the context of continuing work on these issues.

With respect to the existence spectrum for complete sentences of $L_{\omega_1,\omega}$, we extended a generalized Fraïssé method introduced by [12, 14] and combined it with our notion of disjoint $(\leq \lambda, k)$ amalgamation (inspired by Shelah's notion of excellence). Hjorth's proof requires an inductive choice between sentences ϕ_α and ψ_α at each α, which depend on the sentence that characterizes $\aleph_\alpha$. Either one of them homogenously characterizes $\aleph_\alpha$ or the other characterizes $\aleph_{\alpha+1}$. But for $\alpha > 1$, it is unknown which sentence does which. By Theorem 2.4.4 we have specified a sentence to give a homogenous characterization of each $\aleph_r$.

The *finite amalgamation spectrum* of an abstract elementary class $\boldsymbol{K}$ is the set $X_{\boldsymbol{K}}$ of $n < \omega$ (for[7] $\aleph_n \geq$

---

[7]The need for this restriction was pointed out to us by David Kueker who noticed that variants on the well-order examples allow exotic spectra if one requires amalgamation over models smaller than $\mathrm{LS}(\boldsymbol{K})$.

$LS(\mathbf{K})$), and $\mathbf{K}$ satisfies amalgamation[8] in $\aleph_n$. There are many examples[9] where the finite amalgamation spectrum of a complete sentence of $L_{\omega_1,\omega}$ is either $\emptyset$ or ω.

As detailed in Theorem 3.2.20 for each $1 \leq r < \omega$, we gave the first example of such a sentence with a non-trivial spectrum: amalagation holds up to $\aleph_{r-2}$, but fails in $\aleph_{r-1}$. It holds (trivially) in $\aleph_r$ (since all models are maximal); there is no model in $\aleph_{r+1}$.

As one would expect, there are more possibilities if we drop completeness or drop the restriction to sentences of $L_{\omega_1,\omega}$. The previous best result for an incomplete $L_{\omega_1,\omega}$-sentence had disjoint amalgamation as defined in [8] up to $\aleph_{k-3}$, and no model in $\beth_k$. Kolesnikov and Lambie-Hanson [13] study a family of AEC's called coloring classes. Both of these papers construct classes that fail amalgamation at higher cardinals but the connection between the cardinalities where amalgamation fails and of the maximal models is much less tight than in the current paper. The examples of Kolesnikov and Lambie-Hanson are distinctive as amalgamation is equivalent to disjoint amalgamation: some results depend on a generalized Martin axiom. The construction of non-trivial spectra of disjoint embedding [2] and of maximal models for complete sentences [9] rely on the current paper.

There is only a bit more known if one allows arbitrary AEC. Well-orderings of order type at most $\aleph_r$ under end extension have amalgamation in $\{\aleph_0, \aleph_1, \dots, \aleph_r\}$. But these classes are not $L_{\omega_1,\omega}$-axiomatizable. An incomplete sentence with finite amalgamation spectrum $\omega - \{0\}$ is given in [8].

Baldwin and Boney [5] have shown that the Hanf number for amalgamation is no more than the first strongly compact cardinal. The immense gap between the results here show how open the amalgamation spectra is. There are three evident areas: a) try to move the techniques here beyond $\aleph_\omega$; b) tighten the bounds in [8, 13]; c) going beyond $\beth_{\omega_1}$ in ZFC would require totally new ideas.

We noted above that if an AEC has disjoint $(\leq \aleph_s, 2)$-amalgamation it has a model in $\aleph_{s+2}$. Thus, on general grounds we knew $\hat{\mathbf{K}}^r$ fails disjoint $(\leq \lambda, k)$-amalgamation in $\aleph_{r-1}$. But to show ordinary 2- amalgamation failed we had to use our particular combinatorics in Lemma 3.2.3.2. We don't have a 'soft' argument that 'ordinary' amalgamation must fail in $\aleph_{r-1}$. But there is a connection between the amalgamation and existence spectra.

A rough picture of Shelah's vision of the spectrum function for AEC is that model classes are wide or tall. We could summarize that in a hyper-strong Shelah-style conjecture: If a (complete) sentence of $L_{\omega_1,\omega}$ characterizes κ^+ then it has $2^{(\kappa^+)}$ models in κ^+. This conjecture is closely connected to the status of amalgamation in κ.

Lemma 4.0.1. *If $\mathbf{K}$ has only maximal models in κ^+ and has amalgamation in κ then it has at most 2^κ models in κ^+.*

Proof. It is well-known (Lemma 2.7 of [20]) that if an AEC $\mathbf{K}$ has the amalgamation property in κ and all models in κ^+ are maximal, pairs of models in κ^+ can be amalgamated over a submodel of size κ. Thus, there is a 1-1 map from models of cardinality κ^+ to models of cardinality κ: Map M of cardinality κ^+ to a submodel M' of cardinality κ. If M and N map to the same model, they have a common extension. But both are maximal, so they must be isomorphic and we have the Lemma. $\square_{4.0.1}$

Consideration of this conjecture for our examples motivated Part 5 of Theorem 3.2.20, which with Lemma 4.0.1 gives a second proof of Proposition 3.2.16. We close with two questions.

Question 4.0.2. *1. Is there a (complete) sentence of $L_{\omega_1,\omega}$ which characterizes $\kappa > \aleph_0$ and has fewer than 2^κ models of cardinality κ?*

---

[8] We say amalgamation holds in κ in the trivial special case when all models in κ are maximal. We say amalgamation fails in κ if there are no models to amalgamate.

[9] Kueker [15] gave the first example of a complete sentence failing amalgamation in $\aleph_0$.

2. *Is there any AEC, in particular defined by a complete sentence in $L_{\omega_1,\omega}$, whose finite non-trivial amalgamation spectrum is not an interval?*

References

[1] J. Baldwin, C. Laskowski, and S. Shelah. Constructing many uncountable atomic models in $\aleph_1$. submitted, 2015.

[2] John Baldwin, Martin Koerwien, and Ioannis Souldatos. The joint embedding property and maximal models. preprint.

[3] John T. Baldwin. *Categoricity*. Number 51 in University Lecture Notes. American Mathematical Society, Providence, USA, 2009. www.math.uic.edu/~jbaldwin.

[4] John T. Baldwin and Alexei Kolesnikov. Categoricity, amalgamation, and tameness. *Israel Journal of Mathematics*, 170, 2009. also at `www.math.uic.edu/\~\jbaldwin`.

[5] J.T. Baldwin and William Boney. The Hanf number for amalgamation and joint embedding in aec's. 2014.

[6] J.T. Baldwin, Sy Friedman, M. Koerwien, and C. Laskowski. Three red herrings around Vaught's conjecture. to appear: Transactions of the American Math Society, 2013.

[7] J.T. Baldwin, M. Koerwien, and I. Souldatos. The joint embedding property and maximal models. preprint, 2014.

[8] J.T. Baldwin, A. Kolesnikov, and S. Shelah. The amalgamation spectrum. *Journal of Symbolic Logic*, 74:914–928, 2009.

[9] J.T. Baldwin and I. Souldatos. Complete $L_{\omega_1,\omega}$-sentences with maximal models in multiple cardinalities. preprint, 2015.

[10] T. Button and S. Walsh. Ideas and results in model theory: Reference, realism, structure and categoricity. 2015 manuscript:`http://faculty.sites.uci.edu/seanwalsh/files/2015/01/button-walsh-arXiv-submit.1150992.pdf`, 2015.

[11] Bradd Hart and Saharon Shelah. Categoricity over P for first order T or categoricity for $\phi \in l_{\omega_1\omega}$ can stop at $\aleph_k$ while holding for $\aleph_0, \cdots, \aleph_{k-1}$. *Israel Journal of Mathematics*, 70:219–235, 1990.

[12] Greg Hjorth. Knight's model, its automorphism group, and characterizing the uncountable cardinals. *Journal of Mathematical Logic*, pages 113–144, 2002.

[13] Alexei Kolesnikov and Christopher Lambie-Hanson. Hanf numbers for amalgamation of coloring classes. preprint, 2014.

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

[15] J. Malitz. The Hanf number for complete $L_{\omega_1,\omega}$ sentences. In J. Barwise, editor, *The syntax and semantics of infinitary languages*, LNM 72, pages 166–181. Springer-Verlag, 1968.

[16] 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.

[17] S. Shelah. Classification theory for nonelementary classes. II. 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.

[18] S. Shelah. *Classification Theory for Abstract Elementary Classes*. Studies in Logic. College Publications www.collegepublications.co.uk, 2009. Binds together papers 88r, 600, 705, 734 with introduction E53.

[19] S. Shelah. *Classification Theory for Abstract Elementary Classes: II*. Studies in Logic. College Publications <www.collegepublications.co.uk>, 2010. Binds together papers 300 A-G, E46, 838.

[20] Saharon Shelah. Classification of nonelementary classes II, abstract elementary classes. In J.T. Baldwin, editor, *Classification theory (Chicago, IL, 1985)*, pages 419–497. Springer, Berlin, 1987. paper 88: Proceedings of the USA–Israel Conference on Classification Theory, Chicago, December 1985; volume 1292 of *Lecture Notes in Mathematics*.

[21] Ioannis Souldatos. Characterizing the powerset by a complete (Scott) sentence. *Fundamenta Mathematica*, 222:131–154, 2013.

[22] Ioannis Souldatos. Notes on cardinals that characterizable by a complete (Scott) sentence. *Notre Dame Journal of Formal Logic*, 55:533–551, 2013.