Notre Dame Journal of Formal Logic
Volume 60, Number 1, 2019

Disjoint n-Amalgamation and Pseudofinite Countably Categorical Theories

Alex Kruckman

Abstract Disjoint n-amalgamation is a condition on a complete first-order theory specifying that certain locally consistent families of types are also globally consistent. In this article, we show that if a countably categorical theory T admits an expansion with disjoint n-amalgamation for all n, then T is pseudofinite. All theories which admit an expansion with disjoint n-amalgamation for all n are simple, but the method can be extended, using filtrations of Fraïssé classes, to show that certain nonsimple theories are pseudofinite. As case studies, we examine two generic theories of equivalence relations, T^*_{feq} and T_{CPZ}, and show that both are pseudofinite. The theories T^*_{feq} and T_{CPZ} are not simple, but they have NSOP_1. This is established here for T_{CPZ} for the first time.

1 Introduction

The theory T_{RG} of the random graph (also called the *Rado graph*) arises naturally in two distinct ways. First, the random graph is the Fraïssé limit of the class of all finite graphs $\mathcal{G}$: the unique countable ultrahomogeneous graph which embeds a copy of each finite graph. Second, T_{RG} is the almost-sure theory of finite graphs, in the sense of zero-one laws: letting $\mathcal{G}(n)$ be the set of (labeled) graphs of size n and μ_n the uniform measure on $\mathcal{G}(n)$, we have, for every sentence φ,

$$\lim_{n\to\infty} \mu_n\big(\big\{G \in \mathcal{G}(n) \bigm| G \models \varphi\big\}\big) = 1 \iff \varphi \in T_{\mathrm{RG}}.$$

The latter observation shows that T_{RG} is pseudofinite; that is, every sentence in the theory has a finite model. In fact, the probabilistic argument shows that each sentence $\varphi \in T_{\mathrm{RG}}$ has many finite models. For large n, most finite graphs of size n satisfy φ.

Received October 13, 2015; accepted September 21, 2016
First published online January 29, 2019
2010 Mathematics Subject Classification: Primary 03C13; Secondary 03C45
Keywords: amalgamation, pseudofinite theories, countably categorical theories, Fraïssé theory
 10.1215/00294527-2018-0025

The situation is very different for the class $\mathcal{G}_\triangle$ of finite triangle-free graphs. The class $\mathcal{G}_\triangle$ has a Fraïssé limit, the generic triangle-free graph (also called the *Henson graph*; see [18]), as well as a zero-one law for the uniform measures μ_n on $\mathcal{G}_\triangle(n)$, but its almost-sure theory diverges from its generic theory. Indeed, Erdős, Kleitman, and Rothschild [15] showed that almost all large finite triangle-free graphs are bipartite and hence do not contain any cycles of odd length, in contrast to the generic triangle-free graph.

So the probabilistic argument that showed that the theory of the random graph is pseudofinite fails for the generic triangle-free graph. In fact, it is still unknown whether the theory of the generic triangle-free graph is pseudofinite (see Cherlin [10], [11]). This state of affairs suggests the following very general question.

Question 1.1 When does a Fraïssé limit have a pseudofinite theory?

There are, essentially, two ways to show that a theory T is pseudofinite. The first way is to construct finite structures which satisfy arbitrary finite subsets of T. An example in the case of the random graph is the sequence of Paley graphs. For each prime power $q \equiv 1 (\mathrm{mod}\, 4)$, define a graph with domain the finite field $\mathbb{F}_q$, putting an edge between distinct elements a and b just in case $a - b$ is a square in $\mathbb{F}_q$. Then the theories of the Paley graphs converge to T_{RG} (see Blass, Exoo, and Harary [4] for details, and Blass and Rossman [5] for other explicit constructions).

The second way is via a probabilistic argument. Usually, this amounts to specifying a probability measure μ_n on some class $K(n)$ of finite L-structures for all $n \in \omega$, such that

$$\lim_{n\to\infty} \mu_n\big(\{A \in K(n) \mid A \models \varphi\}\big) = 1 \iff \varphi \in T.$$

The first method has the advantage of being more explicit, and the constructions may be of combinatorial interest. But the second method tells us something more, assuming that the measures μ_n are natural enough: not only do the sentences of T have finite models, but *most* large structures in some class satisfy the sentences in T. Of course, the meaning of "natural" is left intentionally vague. For example, the measure μ_n should not concentrate on the nth element of some explicit sequence! We refine our question as follows.

Question 1.2 When does a Fraïssé limit have a pseudofinite theory for a good probabilistic reason? For example, when is it the almost-sure theory for a natural sequence $(K(n), \mu_n)_{n\in\omega}$ of classes of finite structures equipped with probability measures?

An example of a Fraïssé limit which is pseudofinite, but *not* for a good probabilistic reason, is the vector space V of countably infinite dimension over a finite field. The finite models of sentences in $\mathrm{Th}(V)$ are few and far between, existing only in certain finite sizes and unique up to isomorphism in those sizes. This is one of a whole family of examples of a similar character, the smoothly approximable structures, studied by Kantor, Liebeck, and Macpherson in [21] and classified by Cherlin and Hrushovski in [12]. Smoothly approximable structures are essentially algebraic: they are coordinatized by certain geometries coming from vector spaces equipped with bilinear forms.

The main purpose of this article is to advance a claim that "combinatorial" Fraïssé limits (in contrast to the algebraic smoothly approximable structures) which are pseudofinite tend to be pseudofinite for a good probabilistic reason, and moreover that this good probabilistic reason tends to rely on a combinatorial condition, disjoint n-amalgamation, which generalizes the disjoint (or "strong") amalgamation property for Fraïssé classes with trivial algebraic closure (acl).

The starting point is Theorem 3.10, which shows that for a countably categorical theory, disjoint n-amalgamation for all n is a sufficient condition for pseudofiniteness. The hypothesis of disjoint n-amalgamation for all n is very strong; however, almost all examples of "combinatorial" countably categorical theories which are known to be pseudofinite either have disjoint n-amalgamation for all n, or are reducts of theories with disjoint n-amalgamation for all n. Such theories are simple (in the sense of the model-theoretic dividing line; see Theorem 3.14). The only exceptions (that I am aware of at the time of this writing) are built from equivalence relations.

Kim and Pillay [25] made the "rather outrageous conjecture" that every pseudofinite countably categorical theory is simple. The generic theory of a parameterized family of equivalence relations, T^*_{feq}, was suggested by Shelah as a counterexample to this conjecture. However, to my knowledge, no proof that T^*_{feq} is pseudofinite has appeared in the literature.

In this article, I demonstrate pseudofiniteness of T^*_{feq} (see Section 4.2), as well as another generic theory of equivalence relations, T_{CPZ} (see Section 4.3), which was introduced (and shown to not be simple) by Casanovas, Peláez, and Ziegler [8]. In both cases, the argument relies on a method of filtering the relevant Fraïssé class as a union of simpler Fraïssé classes, each of which admits an expansion to a countably categorical theory with disjoint n-amalgamation for all n. This shows that the pseudofiniteness of these examples, too, can be viewed as a consequence of a probabilistic argument involving disjoint n-amalgamation. An interesting feature of this method is that each sentence of the theory is shown to be in the almost-sure theory for a sequence $(K(n), \mu_n)_{n\in\omega}$ of classes of finite structures equipped with probability measures, and hence is pseudofinite for a good probabilistic reason, but different sentences require different sequences.

Countably categorical pseudofinite theories do not have the strict order property. Since I am not aware of a reference for this folklore result, I will give a proof here.

Proposition 1.3 *No countably categorical pseudofinite theory has the strict order property.*

Proof If T has the strict order property, then it interprets a partial order with infinite chains. So it suffices to show that no countably categorical partial order $(P, <)$ with infinite chains is pseudofinite.

By compactness, we can find an infinite increasing chain $\{a_i \mid i \in \omega\}$ with $P \models a_i < a_j$ if and only if $i < j$. In a countably categorical theory, automorphism-invariant properties are definable, so there is a formula $\varphi(x)$, with $\varphi(x) \in \text{tp}(a_i)$ for all i, such that $P \models \varphi(b)$ if and only if there is an infinite increasing chain above b.

Now $P \models \exists x\, \varphi(x) \wedge \forall x\, (\varphi(x) \to \exists y\, (x < y \wedge \varphi(y)))$. But in any partial order, this sentence implies the existence of an infinite increasing chain of elements satisfying $\varphi(x)$, so its conjunction with the partial order axioms has no finite model. □

Džamonja and Shelah [14] introduced the property SOP_1. It is the first in a linearly ordered hierarchy of combinatorial properties called SOP_n (for n-strong order property), which were originally defined by Shelah [29] for $n \geq 3$. A theory has NSOP_n

if it does not have the n-strong order property. As usual in model theory, the named properties are bad: theories with NSOP$_n$ are tamer than theories with SOP$_n$. These properties lie strictly between nonsimplicity and the strict order property (SOP):

$$\text{simple} \implies \text{NSOP}_1 \implies \cdots \implies \text{NSOP}_n \implies \cdots \implies \text{NSOP}.$$

It is worth noting that SOP$_2$ also goes by the name TP$_1$ (the tree property of the first kind; see Kim and Kim [22] for a discussion), and every theory which is known to have SOP$_1$ also has SOP$_3$. So it is possible that NSOP$_1 \stackrel{?}{=}$ NSOP$_2$ = NTP$_1 \stackrel{?}{=}$ NSOP$_3$. The generic triangle-free graph has SOP$_3$ but NSOP$_4$ (see [29]).

Chernikov and Ramsey [13] gave an independence relation criterion for NSOP$_1$ and used it to show that T^*_{feq} has NSOP$_1$. The theory T_{CPZ} was not considered in [13], but the methods there also suffice to show that T_{CPZ} has NSOP$_1$ (see Corollary 4.10). On the other hand, almost nothing is known about pseudofiniteness of countably categorical theories in the region between SOP$_1$ and the strict order property. While acknowledging that we have a paucity of other examples, I think it is reasonable to update the outrageous conjecture of Kim and Pillay in the following way.

Conjecture 1.4 *Every pseudofinite countably categorical theory has NSOP$_1$.*

In Section 2, I review the relevant background on Fraïssé theory. I introduce disjoint n-amalgamation in Section 3.1 and prove Theorem 3.10 in Section 3.2. Section 3.3 contains some context about the role of n-amalgamation properties in model theory, as well as an explanation of how Theorem 3.10 generalizes and unifies previous work. In Section 4, I introduce the notion of a filtered Fraïssé class and give the applications to generic theories of equivalence relations (T^*_{feq} and T_{CPZ}), along with a negative result, Proposition 4.4, showing that this method cannot be used to show that the generic triangle-free graph is pseudofinite.

2 Preliminaries

In this section, I give a brief review of Fraïssé theory. The "canonical language" described in Definition 2.8 provides the bridge to general countably categorical theories. (For proofs, see Cameron [7, Sections 2.6–2.8] or Hodges [19, Section 7.1].)

Let L be a relational language (not necessarily finite), and let K be a class of finite L-structures which is closed under isomorphism.

(a) K has the *hereditary property* if it is closed under substructure.
(b) K has the *joint embedding property* if for all $A, B \in K$, there exist $C \in K$ and embeddings $A \hookrightarrow C$ and $B \hookrightarrow C$.
(c) K has the *amalgamation property* (or 2-*amalgamation*) if for all $A, B, C \in K$ and embeddings $f\colon A \hookrightarrow B$ and $g\colon A \hookrightarrow C$, there exist $D \in K$ and embeddings $f'\colon B \hookrightarrow D$ and $g'\colon C \hookrightarrow D$ such that $f' \circ f = g' \circ g$.
(d) K has the *disjoint amalgamation property* (or *disjoint* 2-*amalgamation*) if, in the definition of the amalgamation property, the images of B and C in D can additionally be taken to be disjoint over the image of A in D: $(f' \circ f)[A] = (g' \circ g)[A] = f'[B] \cap g'[C]$.
(e) K is a *weak Fraïssé class* if it is countable up to isomorphism and has the hereditary property, the joint embedding property, and the amalgamation property.
(f) K is a *Fraïssé class* if it is a weak Fraïssé class and additionally K contains only finitely many structures of size n up to isomorphism for all $n \in \omega$.

Remark 2.1 What I call a weak Fraïssé class here is often simply called a Fraïssé class. However, as we are only interested in Fraïssé classes with countably categorical generic theory, it is convenient to include the finiteness condition in the definition. Note that in a finite relational language, the notions coincide. In many sources the disjoint amalgamation property is called the *strong amalgamation property*.

Definition 2.2 Let M be a countable L-structure.

(1) The *age* of M is the class of all finite structures which embed in M.
(2) M is *ultrahomogeneous* if every isomorphism between finite substructures of M extends to an automorphism of M.
(3) M has *trivial acl* if $\operatorname{acl}(A) = A$ for all $A \subseteq M$.

Theorem/Definition 2.3 The class K is a weak Fraïssé class if and only if there is a countable ultrahomogeneous structure M_K with age K. In this case, M_K is unique up to isomorphism and is called the *Fraïssé limit* of K. We call $T_K = \operatorname{Th}(M_K)$ the *generic theory* of K. We have that K is a (strong) Fraïssé class if and only if T_K is countably categorical. In this case, T_K has quantifier elimination. A Fraïssé class K has the disjoint amalgamation property if and only if M_K has trivial acl.

Given a Fraïssé class K and $n \in \omega$, I will write $K(n)$ for the (finite) set of structures in K with domain $[n] = \{1, \dots, n\}$. Note that I include the empty structure in the case $n = 0$. The set $K(n)$ contains $(n!/|\operatorname{Aut}(A)|)$-many isomorphic copies of every structure A in K. It will be convenient to identify these structures with their quantifier-free n-types: for $A \in K(n)$,

$$\operatorname{qftp}(A) = \big\{\varphi(x_1, \dots, x_n) \;\big|\; \varphi \text{ is quantifier-free, and } A \models \varphi(1, \dots, n)\big\}.$$

Since T_K has quantifier elimination, we can further identify the structures in $K(n)$ with the set of first-order n-types over the empty set relative to T_K which are nonredundant, in the sense that they contain the formulas $\{x_i \neq x_j \mid i \neq j\}$.

Now each n-type relative to T_K is isolated by a quantifier-free formula. In other words, each structure $A \in K(n)$ is distinguished from the others by a single quantifier-free formula $\theta_A(x_1, \dots, x_n)$. In the case that L is finite, we may take θ_A to be the conjunction of the atomic diagram of A. If L is infinite, then a large enough part of the atomic diagram suffices.

Theorem 2.4 *In this notation, the generic theory T_K can be explicitly axiomatized as follows.*

(1) *The universal theory of K: This amounts to the sentences, for $n \in \omega$,*

$$\forall x_1, \dots, x_n \Big(\Big(\bigwedge_{i \neq j} x_i \neq x_j\Big) \to \Big(\bigvee_{A \in K(n)} \theta_A(\overline{x})\Big)\Big),$$

together with, if L is infinite, the information about how θ_A determines the other quantifier-free formulas. That is, for each n, $A \in K(n)$, and quantifier-free formula $\varphi(\overline{x}) \in \operatorname{qftp}(A)$,

$$\forall \overline{x} \left(\theta_A(\overline{x}) \to \varphi(\overline{x})\right).$$

(2) *One-point extension axioms: For all $A \in K(n)$ and $B \in K(n+1)$, we say that (A, B) is a* one-point extension *if A is the induced substructure of B with domain $[n]$. Given a one-point extension (A, B), we have the axiom*

$$\forall \overline{x}\, \exists y\, \big(\theta_A(\overline{x}) \to \theta_B(\overline{x}, y)\big).$$

Definition 2.5 The theory $T_{K,n}$ is the (incomplete) theory axiomatized by:

(1) the sentences in the universal theory of K in at most n universal quantifiers;
(2) all one-point extension axioms for K (with no restriction on the sizes of A and B).

A model of $T_{K,n}$ satisfies all the one-point extension axioms over substructures satisfying one of the formulas θ_A for $A \in K$, but its age need only agree with K up to substructures of size at most n. We will see in Theorem 3.10 below that basic disjoint amalgamation up to level n implies pseudofiniteness of $T_{K,n}$.

It will be useful to consider expansions of T_K at the level of the Fraïssé class K.

Definition 2.6 Let K and K' be Fraïssé classes in languages L and L', respectively, such that $L \subseteq L'$. We say that K' is a *Fraïssé expansion* of K if

(1) $K = \{A \restriction L \mid A \in K'\}$,
(2) for all one-point extensions (A, B) in K, and every expansion of A to a structure A' in K', there is an expansion of B to a structure B' in K' such that (A', B') is a one-point extension in K'.

Theorem 2.7 *The class K' is a Fraïssé expansion of K if and only if the Fraïssé limit $M_{K'}$ of K' is an expansion of the Fraïssé limit M_K of K.*

Proof Suppose that $M_{K'} \restriction L = M_K$. Then $K = \mathrm{Age}(M_K) = \{A \restriction L \mid A \in \mathrm{Age}(M_{K'})\}$, and $\mathrm{Age}(M_{K'}) = K'$. Given a one-point extension (A, B) and an expansion A' of A, we can find a substructure of $M_{K'}$ isomorphic to A'. In the reduct, this substructure is isomorphic to A, and, since the one-point extension axiom for (A, B) is true of M_K, it extends to a copy of B. We can take B' to be the L'-structure on this subset of $M_{K'}$.

Conversely, to show that $M_{K'}$ is an expansion of M_K, by countable categoricity it suffices to show that $M_{K'} \restriction L$ satisfies the theory T_K. It clearly satisfies the universal part, since $\mathrm{Age}(M_{K'} \restriction L) = \{A \restriction L \mid A \in K'\} = K$. For the extension axioms, suppose that (A, B) is a one-point extension, and we have a copy of A in $M_{K'} \restriction L$. Let A' be the L'-structure on this subset of $M_{K'}$. Since K' is a Fraïssé expansion of K, we can find an expansion B' of B in K' such that (A', B') is a one-point extension, and, since the one-point extension axiom for (A', B') is true of $M_{K'}$, our copy of A' extends to a copy of B'. Hence, in the reduct, our copy of A extends to a copy of B. $\square$

Definition 2.8 Let T be any countably categorical L-theory, and let M be its unique countable model. The *canonical language* for T is the language L' with one n-ary relation symbol R_p for each n-type $p(\overline{x})$ realized in M.

We make M into an L'-structure M' in the natural way by setting $M' \models R_p(\overline{a})$ if and only if $\overline{a}$ realizes $p(\overline{x})$ in M. Let $T' = \mathrm{Th}_{L'}(M')$. Then T and T' are interdefinable theories, M' is ultrahomogeneous, and hence is the Fraïssé limit of its age K, and K has the disjoint amalgamation property if and only if M has trivial acl. Note that for each $A \in K(n)$, we may take the isolating formula θ_A to be one of the basic n-ary relation symbols R_p.

3 Disjoint n-Amalgamation

3.1 Definitions To fix notation, $[n] = \{1, \dots, n\}$, $\mathcal{P}([n])$ is the powerset of $[n]$, and $\mathcal{P}^-([n])$ is the set of all proper subsets of $[n]$. A family $\mathcal{F} \subseteq \mathcal{P}([n])$ of subsets of $[n]$ is *downward closed* if $S' \in \mathcal{F}$ whenever $S' \subseteq S$ and $S \in \mathcal{F}$.

Let T be a theory, and let A be a set of parameters in a model of T. We say that a type $p(\overline{x})$ over A in the variables $\{x_i \mid i \in I\}$ is *nonredundant* if it contains the formulas $\{x_i \neq x_j \mid i \neq j \in I\}$ and $\{x_i \neq a \mid i \in I, a \in A\}$. Given a downward closed family of subsets $\mathcal{F} \subseteq \mathcal{P}([n])$, and variables $\overline{x}_1, \dots, \overline{x}_n$, a *coherent $\mathcal{F}$-family of types over A* is a set $\{p_S \mid S \in \mathcal{F}\}$ such that each p_S is a nonredundant type over A in the variables $\overline{x}_S = \{\overline{x}_i \mid i \in S\}$, and $p_{S'} \subseteq p_S$ when $S' \subseteq S$. Here each $\overline{x}_i$ is a tuple of variables, possibly empty or infinite, but such that $\overline{x}_i$ is disjoint from $\overline{x}_j$ when $i \neq j$.

For $n \geq 2$, a *disjoint n-amalgamation problem* is a coherent $\mathcal{P}^-([n])$-family of types over a set A. A *basic* disjoint n-amalgamation problem is a disjoint n-amalgamation problem over the empty set in the singleton variables $x_1, \dots, x_n$.

A *solution* to a (basic) disjoint n-amalgamation problem is an extension of the coherent $\mathcal{P}^-([n])$-family of types to a coherent $\mathcal{P}([n])$-family of types; that is, a nonredundant type $p_{[n]}$ such that $p_S \subseteq p_{[n]}$ for all S. We say T has *(basic) disjoint n-amalgamation* if every (basic) n-amalgamation problem has a solution.

If we replace $\mathcal{P}^-([n])$ by another downward closed family of subsets $\mathcal{F}$ in the definitions above, then we call the amalgamation problem *partial*.

First, we offer some remarks on the definitions.

Remark 3.1 In any coherent $\mathcal{F}$-family of types over A, the type $p_\varnothing$ is a 0-type in the empty tuple of variables, which simply specifies the elementary diagram of the parameters A.

Remark 3.2 To specify a disjoint n-amalgamation problem, it would be sufficient to give the types p_S for all S with $|S| = n - 1$ and check that they agree on intersections, in the sense that $p_S \restriction \overline{x}_{S \cap S'} = p'_S \restriction \overline{x}_{S \cap S'}$ for all S and S'. However, it is sometimes notationally convenient to keep the intermediate stages around.

Remark 3.3 A Fraïssé class K has the disjoint amalgamation property if and only if T_K has disjoint 2-amalgamation. Indeed, given $A, B, C \in K$ and embeddings $f\colon A \hookrightarrow B$ and $g\colon A \hookrightarrow C$, we take A to be the base set of parameters, so $p_\varnothing = \mathrm{qftp}(A)$, and we set $p_{\{1\}}(\overline{x}_1) = \mathrm{qftp}((B \setminus A)/A)$ and $p_{\{2\}}(\overline{x}_2) = \mathrm{qftp}((C \setminus A)/A)$, identifying A with its images in B and C under f and g. By quantifier elimination, these quantifier-free types determine complete types relative to T_K. A solution to this disjoint 2-amalgamation problem is the same as a structure D in K into which B and C embed disjointly over the image of A.

Remark 3.4 Given a Fraïssé class K, recall that we have identified $K(n)$, the structures in K with domain $[n]$, with the set of nonredundant quantifier-free n-types relative to T_K. A basic disjoint n-amalgamation problem relative to T_K is a coherent $\mathcal{P}^-([n])$-family of quantifier-free types $P = \{p_S \mid S \in \mathcal{P}^-([n])\}$ in the variables $x_1, \dots, x_n$, where each type p_S corresponds to a structure A_S in K of size $|S|$. We write $K(n, P) = \{p_{[n]}(x_1, \dots, x_n) \in K(n) \mid p_S \subseteq p_{[n]} \text{ for all } S \in \mathcal{P}^-([n])\}$ for the set of solutions to the amalgamation problem P, each of which corresponds to a structure $A_{[n]}$ in K of size n which contains all the A_S as substructures. To say that T_K has basic disjoint n-amalgamation is to say that $K(n, P)$ is nonempty for all P.

It will be useful to observe that disjoint amalgamation gives solutions to partial amalgamation problems as well.

Lemma 3.5 *Suppose that T has (basic) disjoint k-amalgamation for all $2 \leq k \leq n$. Then every partial (basic) disjoint n-amalgamation problem has a solution.*

Proof I will consider the general case. The same proof works in the basic case.

We are given a partial disjoint n-amalgamation problem over A in variables $\overline{x}_1, \ldots, \overline{x}_n$; that is, a coherent $\mathcal{F}$-family of types $\{p_S \mid S \in \mathcal{F}\}$, with $\mathcal{F} \subseteq \mathcal{P}^-([n])$ downward closed.

We build a solution to the partial disjoint n-amalgamation problem from the bottom up. By induction on $1 \leq k \leq n$, I claim that we can extend this family to a coherent $\mathcal{F}_k$-family of types, where $\mathcal{F}_k = \mathcal{F} \cup \{S \subseteq [n] \mid |S| \leq k\}$. When $k = n$, we have a coherent $\mathcal{P}([n])$-family of types, as desired.

When $k = 1$, if there is any i such that $i \notin S$ for all $S \in \mathcal{F}$, then the original $\mathcal{F}$-family of types says nothing about the variables $\overline{x}_i$. We add $\{i\}$ into $\mathcal{F}_1$ and choose any nonredundant type $p_{\{i\}}$ over A in the variables $\overline{x}_i$. If $\varnothing \notin \mathcal{F}$ (which only happens if $\mathcal{F}$ is empty), then we also add it into $\mathcal{F}_1$, along with the unique 0-type $p_\varnothing$ containing the elementary diagram of A.

Given a coherent $\mathcal{F}_{k-1}$-family of types by induction, with $2 \leq k \leq n$, we wish to extend to a coherent $\mathcal{F}_k$-family of types. If there is any set $S \subseteq [n]$ with $|S| = k$ such that $S \notin \mathcal{F}_{k-1}$, then all proper subsets of S are in $\mathcal{F}_{k-1}$. Hence we have types $\{p_R \mid R \in \mathcal{P}^-(S)\}$ which form a coherent $\mathcal{P}^-(S)$-family. Using k-amalgamation, we can find a nonredundant type p_S in the variables $\overline{x}_S$ extending the types p_R. Doing this for all such S gives a coherent $\mathcal{F}_k$-family of types, as desired. $\square$

Disjoint n-amalgamation is more general and seems more natural, but it is basic disjoint n-amalgamation which is relevant in the proof of Theorem 3.10. We are largely interested in theories with disjoint n-amalgamation for all n, and in this case the two notions agree.

Proposition 3.6 *A theory T has disjoint n-amalgamation for all n if and only if T has basic disjoint n-amalgamation for all n.*

Proof One direction is clear, since basic disjoint n-amalgamation is a special case of disjoint n-amalgamation.

In the other direction, note first that there is a solution to the disjoint n-amalgamation problem $\{p_S \mid S \in \mathcal{P}^-([n])\}$ if and only if the partial type

$$\{x \neq x' \mid x, x' \text{ distinct}\} \cup \bigcup_{S \in \mathcal{P}^-([n])} p_S(\overline{x}_S)$$

is consistent (actually, we could omit the formulas asserting nonredundancy when $n > 2$). Hence, by compactness, we can reduce to the case that A is finite and each tuple of variables $\overline{x}_i$ is finite.

Let $N = |A| + \sum_{i=1}^n |\overline{x}_i|$, where $|\overline{x}_i|$ is the length of the tuple $\overline{x}_i$. Introduce variables $y_1, \ldots, y_N$, where $y_1, \ldots, y_{|A|}$ enumerate A and the remaining variables relabel the x variables. Now each type p_S over A determines a type in some subset of the y variables, by replacing the parameters from A and the x variables by the appropriate y variables. Closing downward under restriction to smaller sets of variables, we obtain a partial basic disjoint N-amalgamation problem over the empty set in the singleton variables $y_1, \ldots, y_N$. By Lemma 3.5 and basic disjoint N-amalgamation, this partial amalgamation problem has a solution, a type $p_{[N]}(y_1, \dots, y_N)$ over the empty set. Once again replacing the y variables with the original parameters from A and x variables, we obtain a type $p_{[n]}$ over A which is a solution to the original n-amalgamation problem. □

Example 3.7 The class $\mathcal{G}_\triangle$ of triangle-free graphs has disjoint 2-amalgamation: if A embeds in B and C, then we can amalgamate B and C "freely" over A by not adding any new edge relations between B and C. But it does not have disjoint 3-amalgamation: the nonredundant 2-types determined by $x_1 R x_2$, $x_2 R x_3$, and $x_1 R x_3$ cannot be amalgamated.

Generalizing, let K_n^k be the class of n-free k-hypergraphs: the language consists of a single k-ary relation $R(x_1, \dots, x_k)$, and the structures in K_n^k are hypergraphs (so R is symmetric and anti-reflexive) such that for every n-tuple $\overline{a}$ of distinct elements, there is some subtuple $\overline{b}$ of length k such that $\neg R(\overline{b})$ holds. Note that $\mathcal{G}_\triangle$ is K_3^2.

For $n > k$, K_n^k satisfies basic disjoint m-amalgamation for $m < n$, but fails basic disjoint n-amalgamation, since the first forbidden configuration has size n. However, K_n^k already fails disjoint $(k+1)$-amalgamation. Over a base set A consisting of a complete hypergraph on $(n-k-1)$ vertices, the k-type over A which describes, together with A, a complete hypergraph on $(n-1)$ vertices is consistent, but $(k+1)$ copies of it cannot be amalgamated.

Example 3.8 There are countably categorical theories which do not have disjoint n-amalgamation for all n, but which admit countably categorical expansions with disjoint n-amalgamation for all n.

As a simple example, consider the theory of a single equivalence relation with k infinite classes. Transitivity is a failure of disjoint 3-amalgamation: the nonredundant 2-types determined by $x_1 E x_2$, $x_2 E x_3$, and $\neg x_1 E x_3$ cannot be amalgamated. But if we expand the language by adding k new unary relations $C_1, \dots, C_k$ in such a way that each class is named by one of the C_i, the resulting theory has disjoint n-amalgamation for all n.

For a more interesting example, the random graph (which is easily seen to have disjoint n-amalgamation for all n) in its canonical language has a reduct to an ultrahomogeneous 3-hypergraph, where the relation $R(a, b, c)$ holds if and only if there are an *odd* number of the three possible edges between a, b, and c. This structure turns out to be ultrahomogeneous in the language $\{R\}$, and its age is the class of all finite 3-hypergraphs with the property that on any four vertices a, b, c, and d, there are an *even* number of the four possible 3-edges. Hence this class fails to have disjoint 4-amalgamation. (For more information on this example, see Macpherson [27], where it is called the "homogeneous two-graph." More examples of this kind can be found in the literature on reducts of homogeneous structures; see, e.g., Thomas [30].)

3.2 Pseudofiniteness

Definition 3.9 A theory T is *pseudofinite* if for every sentence φ such that $T \models \varphi$, φ has a finite model.

Theorem 3.10 below is stated in a fine-grained way: amalgamation just up to level n gives pseudofiniteness of the theory $T_{K,n}$ (see Definition 2.5). The proof involves a probabilistic construction of a structure of size N for each N *from the bottom up*. This is the same idea as in the proof of Lemma 3.5, but there we could fix an arbitrary

k-type extending a given coherent family of l-types for $l < k$. Here we introduce randomness by choosing an extension uniformly at random.

The probabilistic calculation is essentially the same as the one used in the classical proofs of the zero-one laws for graphs and general L-structures (see [19, Lemma 7.4.6]). The key point is that the amalgamation properties allow us to make all choices as independently as possible: the quantifier-free types assigned to subsets A and B of $[N]$ are independent when conditioned on the quantifier-free type assigned to $A \cap B$. It is this independence which makes the calculation go through.

Formally, we construct a probability measure on the space $L[N]$ of L-structures with domain $[N]$. Given a formula $\varphi(\overline{x})$ and a tuple $\overline{a}$ from $[N]$, we write $[\varphi(\overline{a})] = \{M \in L[N] \mid M \models \varphi(\overline{a})\}$. The space $L[N]$ is topologized by taking the instances of the atomic and negated atomic formulas $[(\neg)R(\overline{a})]$ as subbasic open sets. Of course, if L is finite, then $L[N]$ is a finite discrete space.

Theorem 3.10 *Let K be a Fraïssé class whose generic theory T_K has basic disjoint k-amalgamation for all $2 \leq k \leq n$. Then every sentence in $T_{K,n}$ has a finite model. If T_K has basic disjoint k-amalgamation for all k, then every sentence in T_K has a finite model in K.*

Proof I will define a probability measure μ_N on $L[N]$ for each $N \in \omega$ by describing a probabilistic construction of a structure $M_N \in L[N]$. Recall the notation above and in Remark 3.4.

We assign quantifier-free k-types to each subset of size k from $[N]$ by induction. When $k = 0$, there is no choice: by hereditarity and the joint embedding property, there is a unique empty structure in $K(0)$. When $k = 1$, for each $i \in [N]$, choose the quantifier-free 1-type of $\{i\}$ uniformly at random from $K(1)$. Now proceed inductively: having assigned quantifier-free l-types to all subsets of size l with $l < k$, we wish to assign quantifier-free k-types. For each k-tuple $i_1, \ldots, i_k$ of distinct elements from $[N]$, let $P = \{p_S \mid S \in \mathcal{P}^-([k])\}$ be the collection of quantifier-free types assigned to all proper subtuples, that is, $p_S(\overline{x}_S) = \mathrm{qftp}(\{i_j \mid j \in S\})$. If T_K has basic disjoint k-amalgamation, $K(k, P)$ is nonempty and finite, and we may choose the quantifier-free k-type of $i_1, \ldots, i_k$ uniformly at random from $K(k, P)$.

Now if T_K has basic disjoint k-amalgamation for all k, we can continue this construction all the way up to $k = N$, so that the resulting structure M_N is in $K(N)$. Call this the *unbounded* case. On the other hand, if T_K has basic disjoint k-amalgamation only for $k \leq n$, then we stop at $k = n$. To complete the construction, we assign any remaining relations completely freely at random. That is, for each relation R (of arity $r > n$) and r-tuple $i_1, \ldots, i_r$ containing at least $n + 1$ distinct elements, we set $R(i_1, \ldots, i_r)$ with probability $1/2$. The result is an L-structure M_N which may not be in K, but the induced structures of size at most n are guaranteed to be in K. Call this the *bounded* case.

I claim that if φ is one of the axioms of $T_{K,n}$ (in the bounded case) or T_K (in the unbounded case), then $\lim_{N\to\infty} \mu_N([\varphi]) = 1$. Each universal axiom φ has the form $\forall x_1, \ldots, x_k\, \psi(\overline{x})$ (with $k \leq n$ in the bounded case), where ψ is quantifier-free and true on all k-tuples from structures in K. Since all substructures of our random structure of size at most k are in K, φ is always satisfied by M_N, and so $\mu_N([\varphi]) = 1$ for all N.

Now suppose that φ is the one-point extension axiom $\forall \overline{x}\, \exists y\, (\theta_A(\overline{x}) \to \theta_B(\overline{x}, y))$. Let $\overline{a}$ be a tuple of $|A|$-many distinct elements from $[N]$, and let b be any other element. Conditioning on the event that $M_n \models \theta_A(\overline{a})$, there is a positive probability ε that $M_N \models \theta_B(\overline{a}, b)$.

Indeed, in the unbounded case, or when $|A| < n$ in the bounded case, θ_B specifies the quantifier-free $|B|$-type of the tuple $\overline{a}b$ among those allowed by K. There is a positive probability $(1/|K(1)|)$ that the correct 1-type is assigned to b, and, given that the correct l-type has been assigned to all subtuples of $\overline{a}b$ involving b of length $l < k$, there is a positive probability $(1/|K(k, P)|$ for the appropriate basic disjoint k-amalgamation problem $P)$ that the correct k-type is assigned to a given subtuple of length k. Then ε is the product of all these probabilities for $1 \leq k \leq |B|$. When $|A| \geq n$ in the unbounded case, the above reasoning applies for the subtuples of $\overline{a}b$ of length at most n. On longer tuples, since θ_B only mentions finitely many relations, and the truth values of these relations are assigned freely at random, there is some additional positive probability that these will be decided in a way satisfying θ_B (at least $1/2^m$, where m is the minimum number of additional instances of relations which need to be decided positively or negatively to ensure satisfaction of θ_B).

Moreover, for distinct elements b and b', the events that $\overline{a}b$ and $\overline{a}b'$ satisfy θ_B are conditionally independent, since the quantifier-free types of tuples involving elements from $\overline{a}$ and b but not b' are decided independently from those of tuples involving elements from $\overline{a}$ and b' but not b, conditioned on the quantifier-free type assigned to $\overline{a}$.

Now we compute the probability that φ is *not* satisfied by M_N. Conditioned on the event that $M_N \models \theta_A(\overline{a})$, the probability that $M_N \not\models \exists y\, \theta_B(\overline{a}, y)$ is $(1-\varepsilon)^{N-|A|}$, since there are $N - |A|$ choices for the element b, each with independent probability $(1 - \varepsilon)$ of failing to satisfy θ_B. Removing the conditioning, the probability that $M_N \not\models \exists y\, (\theta_A(\overline{a}) \rightarrow \theta_B(\overline{a}, y))$ for any given $\overline{a}$ is at most $(1 - \varepsilon)^{N-|A|}$, since the formula is vacuously satisfied when $\overline{a}$ does not satisfy θ_A. Finally, there are $N^{|A|}$ possible tuples $\overline{a}$, so the probability that $M_N \not\models \forall \overline{x}\, \exists y\, (\theta_A(\overline{x}) \rightarrow \theta_B(\overline{x}, y))$ is at most $N^{|A|}(1 - \varepsilon)^{N-|A|}$. Since $|A|$ is constant, the exponential decay dominates the polynomial growth, and $\lim_{N\to\infty} \mu_N([\neg\varphi]) = 0$, so $\lim_{N\to\infty} \mu_N([\varphi]) = 1$.

To conclude, any sentence $\psi \in T_{K,n}$ is a logical consequence of finitely many of the axioms $\varphi_1, \ldots, \varphi_m$ considered above. We need only pick N large enough so that $\mu_N([\varphi_i]) > 1 - 1/m$ for all i. Then $\mu_N([\bigwedge_{i=1}^m \varphi_i]) > 0$, so the conjunction $\bigwedge_{i=1}^m \varphi_i$, and hence also ψ, has a model of size N. In the unbounded case, our construction ensures that this model is in K. □

Corollary 3.11 *Any countably categorical theory T with disjoint n-amalgamation for all $n \geq 2$ is pseudofinite.*

Proof Let T' be the equivalent of T in the canonical language. Then it suffices to show that T' is pseudofinite, since pseudofiniteness is preserved under interdefinability. But T' is the generic theory for a Fraïssé class with basic disjoint n-amalgamation for all n, so by Theorem 3.10, it is pseudofinite. □

Remark 3.12 Since pseudofiniteness is preserved under reduct, the examples described in Example 3.8 are pseudofinite.

3.3 Relationship to other notions The notion of n-amalgamation has been studied in other model-theoretic contexts, usually in the form of *independent n-amalgamation*. Given some notion of independence, $\mathop{\smile\!\!\!\!\!\!\mid}$, the main example being nonforking independence in a simple theory, an independent n-amalgamation problem is given by a coherent $\mathcal{P}^-([n])$-family of types over A, with the nonredundancy condition replaced by the condition that any realization $\{\overline{a}_i \mid i \in S\}$ of $p_S(\overline{x}_S)$ is an independent set over A with respect to $\mathop{\smile\!\!\!\!\!\!|}\;$.

In the case $n = 3$, independent 3-amalgamation over models is often called the *independence theorem*. It is a well-known theorem of Kim and Pillay [24] that the independence theorem, along with a few other natural properties, characterizes forking in simple theories.

Theorem 3.13 ([24, Theorem 4.2]) *Let T be a complete theory, and let $\mathop{\smile\!\!\!\!\!\!|}$ be a ternary relation, written $a \mathop{\smile\!\!\!\!\!\!|}_A B$, where a is a finite tuple and A and B are sets. As usual, all tuples and sets come from some highly saturated model of T. Suppose that $\mathop{\smile\!\!\!\!\!\!|}$ satisfies the following properties.*

(1) *(Invariance) If $a \mathop{\smile\!\!\!\!\!\!|}_A B$ and $\mathrm{tp}(a'A'B') = \mathrm{tp}(aAB)$, then $a' \mathop{\smile\!\!\!\!\!\!|}_{A'} B'$.*
(2) *(Local character) For all a, B, there is $A \subseteq B$ such that $|A| \leq |T|$ and $a \mathop{\smile\!\!\!\!\!\!|}_A B$.*
(3) *(Finite character) We have that $a \mathop{\smile\!\!\!\!\!\!|}_A B$ if and only if for every finite tuple b from B, $a \mathop{\smile\!\!\!\!\!\!|}_A Ab$.*
(4) *(Extension) For all a, A, and B, there is a' such that $\mathrm{tp}(a'/A) = \mathrm{tp}(a/A)$ and $a' \mathop{\smile\!\!\!\!\!\!|}_A B$.*
(5) *(Symmetry) If $a \mathop{\smile\!\!\!\!\!\!|}_A Ab$, then $b \mathop{\smile\!\!\!\!\!\!|}_A Aa$.*
(6) *(Transitivity) If $A \subseteq B \subseteq C$, then $a \mathop{\smile\!\!\!\!\!\!|}_A B$ and $a \mathop{\smile\!\!\!\!\!\!|}_B C$ if and only if $a \mathop{\smile\!\!\!\!\!\!|}_A C$.*
(7) *(Independence theorem) Let $M \models T$ be a model, let a and a' be tuples such that $\mathrm{tp}(a/M) = \mathrm{tp}(a'/M)$, and let A and B be sets. If $A \mathop{\smile\!\!\!\!\!\!|}_M B$, $a \mathop{\smile\!\!\!\!\!\!|}_M A$, and $a' \mathop{\smile\!\!\!\!\!\!|}_M B$, then there exists a'' such that $\mathrm{tp}(Aa''/M) = \mathrm{tp}(Aa/M)$, $\mathrm{tp}(Ba''/M) = \mathrm{tp}(Ba'/M)$, and $a'' \mathop{\smile\!\!\!\!\!\!|}_M AB$.*

Then T is simple, and $\mathop{\smile\!\!\!\!\!\!|}$ is nonforking ($\mathop{\smile\!\!\!\!\!\!|} = \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{f}}$).

Disjoint n-amalgamation is a strong form of independent amalgamation, where the relevant independence relation is the disjointness relation $\mathop{\smile\!\!\!\!\!\!|}^{=}$, defined by $A \mathop{\smile\!\!\!\!\!\!|}^{=}_C B$ if and only if $A \cap B \subseteq C$. We say that a theory has *trivial forking* if $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{f}} = \mathop{\smile\!\!\!\!\!\!|}^{=}$.

Theorem 3.14 *A countably categorical theory T with disjoint 2-amalgamation (i.e., trivial acl) and disjoint 3-amalgamation is simple with trivial forking.*

Proof We can use Theorem 3.13 to show that $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{f}} = \mathop{\smile\!\!\!\!\!\!|}^{=}$. Most of the conditions are straightforward to check, so I will only remark on a few of them. For local character, we can take $A = a \cap B$, so A is finite and $a \mathop{\smile\!\!\!\!\!\!|}^{=}_A B$. For extension, we find a' by realizing the type $\mathrm{tp}(a/A) \cup \{a_i \neq b \mid a_i$ from a such that $a_i \notin A$, and $b \in B\}$. This is consistent by trivial acl and compactness. Finally, for the independence theorem, we apply disjoint 3-amalgamation to amalgamate the three 2-types $p_{\{12\}} = \mathrm{tp}(aA/M)$, $p_{\{13\}} = \mathrm{tp}(a'B/M)$, and $p_{\{23\}} = \mathrm{tp}(AB/M)$ (first removing any redundant elements of M from a, a', A, and B). □

Remark 3.15 A consequence of Theorem 3.14 is that the class of triangle-free graphs $\mathcal{G}_\triangle$ does not admit a Fraïssé expansion to a class with disjoint n-amalgamation for all n, since its generic theory is not simple (see [29]).

Remark 3.16 Motivated by the fact that many examples of simple theories (such as T_{RG} and ACFA; see Chatzidakis and Hrushovski [9]) satisfy independent

n-amalgamation for $n \geq 3$, Kolesnikov [26] and Kim, Kolesnikov, and Tsuboi [23] developed a hierarchy of notions of n-simplicity for $1 \leq n \leq \omega$, where 1-simplicity coincides with simplicity. If a countably categorical theory T has disjoint k-amalgamation for all $2 \leq k \leq n$, then it is $(n-2)$-simple with trivial forking, and if it has disjoint n-amalgamation for all n, then it is ω-simple.

Several other appearances of n-amalgamation properties in model theory are worth mentioning. In the context of abstract elementary classes (AECs), independent n-amalgamation of models goes by the name "excellence" (see, e.g., Baldwin [2]). Disjoint n-amalgamation for classes of finite structures has also been studied by Baldwin, Koerwien, and Laskowski [3] with applications to AECs. And in the context of stable theories, Goodrick, Kim, and Kolesnikov [17] have uncovered a connection between existence and uniqueness of independent n-amalgamation and definable polygroupoids, generalizing earlier work of Hrushovski [20] on independent 3-amalgamation and groupoids.

The observation that disjoint n-amalgamation is sufficient for pseudofiniteness generalizes and unifies a number of earlier observations. I will note a few here.

(i) Oberschelp [28] identified an unusual syntactic condition which is sufficient for the almost-sure theory of a class of finite structures under the uniform measures to agree with its generic theory. A universal sentence is called *parametric* if it is of the form $\forall x_1, \ldots, x_n\, ((\bigwedge_{i \neq j} x_i \neq x_j) \to \varphi(\overline{x}))$, where φ is a Boolean combination of atomic formulas $R(y_1, \ldots, y_m)$ such that each variable x_i appears among the y_j. For example, reflexivity $\forall x\, R(x,x)$ and symmetry $\forall x, y\, (x \neq y \to (R(x,y) \leftrightarrow R(y,x)))$ are parametric conditions, while transitivity $\forall x, y, z\, ((R(x,y) \wedge R(y,z)) \to R(x,z))$ is not a parametric condition, since each atomic formula appearing only involves two of the three quantified variables. A *parametric class* is the class of finite models of a set of parametric axioms.

Any parametric class has disjoint n-amalgamation for all n. It is easiest to see this by checking basic disjoint n-amalgamation: the restrictions imposed by a parametric theory on the relations involving nonredundant n-tuples and m-tuples are totally independent when $n \neq m$.

(ii) In their work on the random simplicial complex, Brooke-Taylor and Testa [6] introduced the notion of a *local Fraïssé class* and showed that the generic theory of a local Fraïssé class is pseudofinite, by methods similar to those in the proof of Theorem 3.10. A universal sentence is called *local* if it is of the form $\forall x_1, \ldots, x_n\, (R(x_1, \ldots, x_n) \to \psi(\overline{x}))$, where R is a relation in the language and ψ is quantifier-free. A *local class* is the class of finite models of a set of local axioms.

Again, any local class has n-amalgamation for all n. A local theory only imposes restrictions on tuples which satisfy some relation. So disjoint n-amalgamation problems can be solved "freely" by simply not adding any further relations.

(iii) Ahlman [1] has investigated countably categorical theories in a binary relational language (one with no relation symbols of arity greater than 2) which are simple with SU-rank 1 and trivial pregeometry. In the case when $\mathrm{acl}^{\mathrm{eq}}(\varnothing) = \varnothing$, this agrees with what I call a simple theory with trivial forking ($\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{f}} = \mathop{\smile\!\!\!\!\!\!|}^{=}$) above.

Ahlman shows that in such a theory T there is a $\varnothing$-definable equivalence relation ξ with finitely many infinite classes such that T can be axiomatized by certain (ξ, Δ)-*extension properties* describing the possible relationships between elements in different classes. Further, he shows that these theories are pseudofinite. The definition of (ξ, Δ)-extension property is somewhat technical, so I will not give it here. But this condition implies that T has an expansion (obtained by naming the finitely many classes of ξ) with n-amalgamation for all n. The fact that the language is binary ensures that describing the possible relationships between pairs of elements suffices.

4 Generic Theories of Equivalence Relations

4.1 Filtered Fraïssé classes We will extend the disjoint n-amalgamation argument for pseudofiniteness to certain nonsimple theories, using the notion of a filtered Fraïssé class.

Definition 4.1 A Fraïssé class K is *filtered* by a chain $K_0 \subseteq K_1 \subseteq K_2 \subseteq \cdots$ if each K_n is a Fraïssé class, and $\bigcup_{n\in\omega} K_n = K$.

Theorem 4.2 *Let K be a Fraïssé class filtered by $\{K_n \mid n \in \omega\}$. Then $\varphi \in T_K$ if and only if $\varphi \in T_{K_n}$ for all sufficiently large n.*

Proof It suffices to check for each of the axioms of T_K given in Theorem 2.4. Since each K_n is a subclass of K, every universal sentence in T_K is also in T_{K_n}. Let (A, B) be a one-point extension with corresponding axiom φ. For large enough n, the structures A and B are in K_n, so (A, B) is also a one-point extension in K_n, and $\varphi \in T_{K_n}$. $\square$

Pseudofiniteness is preserved in filtered Fraïssé classes.

Corollary 4.3 *If a Fraïssé class K is filtered by $\{K_n \mid n \in \omega\}$ and each generic theory T_{K_n} is pseudofinite, then the generic theory T_K is pseudofinite.*

Proof Each sentence φ in T_K is also in T_{K_n} for sufficiently large n, and hence φ has a finite model. $\square$

As a consequence, if K is filtered by $\{K_n \mid n \in \omega\}$, and each K_n admits a Fraïssé expansion with disjoint n-amalgamation for all n, then K is pseudofinite. This argument is used in the next two sections to establish pseudofiniteness of the theories T^*_{feq} and T_{CPZ}.

It is worth noting that this method cannot be used to show that the theory of the generic triangle-free graph is pseudofinite. Let G_1, G_2, and G_3 be the graphs on three vertices with a single edge, two edges, and three edges, respectively. For any filtration $\{K_n \mid n \in \omega\}$ of the Fraïssé class $\mathcal{G}_\triangle$ of triangle-free graphs, some K_n must include the graphs G_1 and G_2 but not G_3. But Proposition 4.4 shows that such a class does not admit a Fraïssé expansion with disjoint n-amalgamation for all n.

Proposition 4.4 *Let K be a Fraïssé class consisting of graphs (in the language with a single edge relation E), and suppose that K contains the graphs G_1 and G_2 but not G_3. Then no Fraïssé expansion of K has disjoint* 2*-amalgamation and disjoint* 3*-amalgamation.*

Proof Suppose for contradiction that K has a Fraïssé expansion K' in the language L' with disjoint 2-amalgamation and disjoint 3-amalgamation. Let $p(x)$ be any quantifier-free 1-type in K'. Then by disjoint 2-amalgamation we can find some quantifier-free 2-type $q(x, y)$ in K' such that $q(x, y) \models p(x) \wedge p(y) \wedge x \neq y$. Now, letting $p_\varnothing$ be the unique quantifier-free 0-type in K', the family of types $\{p_\varnothing, p(x), p(y), p(z), q(x, y), q(y, z), q(x, z)\}$ is a basic disjoint 3-amalgamation problem for K'. Then we must have $q(x, y) \models \neg xRy$, for otherwise the reduct to L of any solution to the 3-amalgamation problem would be a copy of G_3 in K.

Let H be the graph on two vertices, v_1 and v_2, with no edge. Note that H is in K. Labeling the vertices of G_1 by v_1, v_2, v_3, so that the unique edge is v_2Rv_3, G_1 is a one-point extension of H. Now H admits an expansion to a structure in K' (described by $q(v_1, v_2)$) in which both vertices v_1 and v_2 have quantifier-free type p, so since K' is a Fraïssé expansion of K, G_1 admits a compatible expansion to a structure in K'; call it G_1'. Let $p'(y) = \mathrm{qftp}_{G_1'}(v_3)$, and let $q_i(x, y) = \mathrm{qftp}_{G_1'}(v_i, v_3)$ for $i = 1, 2$. Note that we have $q_i(x, y) \models p(x) \wedge p'(y)$ for $i = 1, 2$, but $q_1(x, y) \models \neg xRy$, while $q_2(x, y) \models xRy$. That is, the pair of 1-types $p(x)$ and $p'(y)$ are consistent with both xRy and $\neg xRy$. We will use this situation to build a triangle.

Again, labeling the vertices of G_2 by v_1, v_2, v_3, so that $\neg v_1Rv_2$, G_2 is a one-point extension of H. Since H admits an expansion to a structure H' in K' so $H' \models q_1(v_1, v_2)$, G_2 admits a compatible expansion to a structure G_2' in K'. Let $p'' = \mathrm{qftp}_{G_2'}(v_3)$, $r_1(x, z) = \mathrm{qftp}_{G_2'}(v_1, v_3)$, and $r_2(y, z) = \mathrm{qftp}_{G_2'}(v_2, v_3)$. Note that $r_1(x, z) \models p(x) \wedge p''(z) \wedge xRz$ and $r_2(y, z) \models p'(x) \wedge p''(z) \wedge yRz$.

Now the family of types $\{p_\varnothing, p(x), p'(y), p''(z), q_2(x, y), r_1(x, z), r_2(y, z)\}$ is a basic disjoint 3-amalgamation problem for K'. But the reduct to L of any solution is a copy of G_3 in K. $\square$

4.2 The theory T^*_{feq} Let L be the language with two sorts, O and P (for "objects" and "parameters"), and a ternary relation $E_x(y, z)$, where x is a variable of sort P and y and z are variables of sort O. Then K_{feq} is the class of finite L-structures with the property that for all a of sort P, $E_a(y, z)$ is an equivalence relation on O.

The class K_{feq} is a Fraïssé class. We define T^*_{feq} to be the generic theory of K_{feq}. Our aim is to show that it is pseudofinite. Before giving the details of the proof, I will describe the simple idea. Filter the class K_{feq} by the subclasses K_n in which each equivalence relation in the parameterized family has at most n classes. Expand these classes by parameterized predicates naming each class. The resulting class has n-amalgamation for all n, and hence has pseudofinite generic theory.

Theorem 4.5 *The theory T^*_{feq} is pseudofinite.*

Proof For $n \geq 1$, let K_n be the subclass of K_{feq} consisting of those structures with the property that for all a of sort P, the equivalence relation E_a has at most n classes. Let us check that K_n is a Fraïssé class.

It clearly has the hereditary property. For the disjoint amalgamation property, suppose that we have embeddings $f\colon A \hookrightarrow B$ and $g\colon A \hookrightarrow C$ of structures in K_n. We specify a structure D with domain $A \cup (B \setminus f[A]) \cup (C \setminus g[A])$ into which B and C embed in the obvious way over A. That is, for each parameter a in $P(D)$, we must specify an equivalence relation on $O(D)$. If a is in $P(A)$, then it already defines equivalence relations on B and C. First, number the E_a-classes in A by $1, \dots, l$. Then, if there are further unnumbered E_a-classes in B and C, number them by $l+1, \dots, m_B$ and $l+1, \dots, m_C$, respectively. Note that $m_B, m_C \leq n$. Now define E_a in $O(D)$ to have $\max(m_B, m_C)$ classes by merging the classes assigned the same number in the obvious way. The situation is even simpler if a is not in $P(A)$. Say without loss of generality that it is in $P(B)$. Then we can extend E_a to $O(C)$ by adding all elements of $O(C \setminus g[A])$ to a single existing E_a-class. The joint embedding property follows from the amalgamation property by taking A to be the empty structure.

For any structure A in K_{feq}, if $|O(A)| = N$, then for all $a \in P(A)$, the equivalence relation E_a has at most N classes, so $A \in K_N$. Hence $K_{\text{feq}} = \bigcup_{n=1}^{\infty} K_n$. So K_{feq} is a filtered Fraïssé class, and by Corollary 4.3, it suffices to show that each T_{K_n} is pseudofinite.

Let L'_n be the expanded language which includes, in addition to the relation E, n binary relation symbols $C_1(x,y), \dots, C_n(x,y)$, where x is a variable of sort P and y is a variable of sort O. Let K'_n be the class of finite L'_n-structures which are expansions of structures in K_n such that for all a of sort P, each of the E_a-classes is picked by one of the formulas $C_i(a,y)$.

We need to check that K'_n is a Fraïssé expansion of K_n. Certainly we have $K_n = \{A \upharpoonright L \mid A \in K'_n\}$, since every structure in K_n can be expanded to one in K'_n by labeling the classes for each equivalence relation arbitrarily. Suppose now that (A, B) is a one-point extension in K_n and A' is an expansion of A to a structure in K'_n. If the new element $b \in B$ is in $P(B)$, then it defines a new equivalence relation E_b on $O(A) = O(B)$, and we can expand B to B' in K'_n by labeling the E_b-classes arbitrarily. On the other hand, suppose that b is in $O(B)$. Then for each parameter a, either b is an existing E_a-class labeled by $C_i(a,y)$, in which case we set $C_i(a,b)$, or b is in a new E_a-class, in which case we set $C_j(a,b)$ for some unused C_j.

Finally, note that $T_{K'_n}$ has disjoint 2-amalgamation, since it is a Fraïssé class with the disjoint amalgamation property. I claim that it also has disjoint n-amalgamation for all $n \geq 3$. Indeed, the behavior of the ternary relation $E_x(y,z)$ is entirely determined by the behavior of the binary relations $C_i(x,y)$, and an L'_n-structure $(P(A), O(A))$ is in K'_n if and only if for every a in $P(A)$ and b in $O(a)$, $C_i(a,b)$ holds for exactly one i. So any inconsistency is already ruled out at the level of the 2-types. Since in a coherent $\mathcal{P}^-([n])$-family of types for $n \geq 3$, every pair of variables is contained in one of the types, we conclude that there are no inconsistencies, and every disjoint n-amalgamation problem has a solution.

So $T_{K'_n}$ has disjoint n-amalgamation for all n, and hence it and its reduct T_{K_n} are pseudofinite by Theorem 3.10. □

A natural question is whether T^*_{feq} is, in fact, the almost-sure theory for the class K_{feq} for the uniform measures. It is not, as the following proposition shows. Of course, since we have described K_{feq} in a two-sorted language, there is some ambiguity as to what we mean by the uniform measures. For maximum generality, let us fix two increasing functions $f, g \colon \omega \to \omega$. For $n \in \omega$, let $K_{\text{feq}}(f(n), g(n))$ be the structures in K_{feq} with object sort of size $f(n)$ and parameter sort of size $g(n)$, and let $\mu_{f(n),g(n)}$ be the uniform measure on $K_{\text{feq}}(f(n), g(n))$.

Proposition 4.6 *There is a sentence φ in T^*_{feq} such that*

$$\lim_{n\to\infty} \mu_{f(n),g(n)}\big(\{A \in K_{\mathrm{feq}}(f(n), g(n)) \mid A \models \varphi\}\big) = 0.$$

Proof An example of such a sentence φ is

$$\forall(x : P)\,\forall(x' : P)\,\forall(y : O)\,\forall(y' : O)\,\exists(z : O)$$
$$\big(y \neq y' \to \big(E_x(y, z) \wedge E_{x'}(y', z)\big)\big),$$

which expresses that any two equivalence classes for distinct equivalence relations intersect. We have that φ is in T^*_{feq}, since for any structure in K_{feq} with parameters $a \neq a'$ and objects b, b' (possibly $b = b'$), we can add a new object element c which is E_a-equivalent to b and $E_{a'}$-equivalent to b', so φ is implied by the relevant one-point extension axioms.

I will sketch the asymptotics. The measure $\mu_{f(n),g(n)}$ amounts to picking $g(n)$ equivalence relations on a set of size $f(n)$ uniformly and independently. The expected number of equivalence classes in an equivalence relation on a set of size n, chosen uniformly, grows asymptotically as $\frac{n}{\log(n)}(1 + o(1))$ (see Flajolet and Sedgewick [16, Proposition 8.8]). Thus, most of the $g(n)$ equivalence relations have equivalence classes which are much smaller (with average size approximately $\log(n)$) than the number of classes, and the probability that every E_a-class is large enough to intersect every E_b-class nontrivially for all distinct a and b converges to zero. □

Proposition 4.6 shows that T^*_{feq} is not the almost-sure theory of K_{feq} for the measures $\mu_{f(n),g(n)}$, but it would be interesting to know whether such an almost-sure theory exists.

Question 4.7 Does the class K_{feq} have a first-order zero-one law for the measures $\mu_{f(n),g(n)}$? If so, does the almost-sure theory depend on the relative growth rates of f and g?

4.3 The theory T_{CPZ} Let L be the language with a symbol $E_n(\overline{x}; \overline{y})$ of arity $2n$ for all $n \geq 1$. Then K_{CPZ} is the class of finite L-structures with the property that E_n is an equivalence relation on n-tuples for all n, and there is a single E_n-class consisting of all n-tuples which do *not* consist of n distinct elements.

The class K_{CPZ} is a Fraïssé class. We define T_{CPZ} to be the generic theory of K_{CPZ}. In [8], Casanovas, Peláez, and Ziegler introduced the theory T_{CPZ} and showed that it has NSOP_2 and is not simple. For completeness, I will show how to combine the "independence lemma" from [8] with the 3-amalgamation criterion due to Chernikov and Ramsey [13] to show that, in fact, T_{CPZ} has NSOP_1.

We write $\mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{u}}$ for coheir independence: given a model M and tuples a and b, $a \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{u}}_M b$ if and only if $\mathrm{tp}(a/Mb)$ is finitely satisfiable in M; that is, for every formula $\varphi(x, m, b) \in \mathrm{tp}(a/Mb)$, there exists $m' \in M$ such that $\models \varphi(m', m, b)$.

Theorem 4.8 ([13, Theorem 5.7]) *T has $NSOP_1$ if and only if for every $M \models T$ and $b_0c_0 \equiv_M b_1c_1$ such that $c_1 \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{u}}_M c_0$, $c_0 \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{u}}_M b_0$, and $c_1 \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{u}}_M b_1$, there exists b such that $bc_0 \equiv_M b_0c_0 \equiv_M b_1c_1 \equiv_M bc_1$.*

For our purposes, the reader can take the independent 3-amalgamation condition in Theorem 4.8 as the definition of NSOP_1. (For the original definition and further discussion of this property, see [13] or [14].)

Lemma 4.9 ([8, Lemma 4.2]) *Let a, b, c, d', d'' be tuples, and let F be a finite set from a model $M \models T_{\mathrm{CPZ}}$. Assume that a and c have only elements of F in common ($a \mathop{\smile\!\!\!\!\!\mid}^{=}_{F} c$). If $d'a \equiv_F d'b \equiv_F d''b \equiv_F d''c$, then there exists d such that $da \equiv_F d'a \equiv_F d''c \equiv_F dc$.*

Corollary 4.10 *The theory T_{CPZ} has NSOP$_1$.*

Proof Suppose that we are given $M \models T_{\mathrm{CPZ}}$ and $d'a \equiv_M d''c$ such that $c \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{u}}_{M} a$, $a \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{u}}_{M} d'$, and $c \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{u}}_{M} d''$. Let $p(x, y) = \mathrm{tp}(d'a/M) = \mathrm{tp}(d''c/M)$. To verify the condition in Theorem 4.8, we need to show that $p(x, a) \cup p(x, c)$ is consistent.

Suppose it is inconsistent. Then there is some finite subset $F \subseteq M$ such that letting $q(x, y) = \mathrm{tp}(d'a/F) = \mathrm{tp}(d''c/F)$, $q(x, a) \cup q(x, c)$ is inconsistent. Since $c \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{u}}_{M} a$, we certainly have $c \mathop{\smile\!\!\!\!\!\mid}^{=}_{M} a$. By increasing F, we may assume that $c \mathop{\smile\!\!\!\!\!\mid}^{=}_{F} a$. By countable categoricity, q is isolated by a single formula $\theta(x, y)$ over F, and $\theta(d', y) \in \mathrm{tp}(a/Md')$, so by finite satisfiability there exists b in M satisfying $q(d', b)$. Since $d' \equiv_M d''$, we also have $\models q(d'', b)$.

Now the assumptions of Lemma 4.9 are satisfied, and we can find d satisfying $q(d, a)$ and $q(d, c)$, which contradicts inconsistency. □

Now we turn to pseudofiniteness of T_{CPZ}. The strategy is the same as in Section 4.2: filter the Fraïssé class K_{CPZ} by bounding the number of equivalence classes, and expand to a class with disjoint n-amalgamation for all n by naming the classes.

Theorem 4.11 *The theory T_{CPZ} is pseudofinite.*

Proof For $n \geq 1$, let K_n be the subclass of K_{CPZ} consisting of those structures with the property that for all k, the equivalence relation E_k has at most n classes, in addition to the class of redundant tuples.

The class K_n has the hereditary property, and the joint embedding property follows from the amalgamation property by taking A to be the empty structure. For the disjoint amalgamation property, we wish to amalgamate embeddings $f\colon A \hookrightarrow B$ and $g\colon A \hookrightarrow C$ of structures in K_n. We specify a structure D with domain $A \cup (B \setminus f[A]) \cup (C \setminus g[A])$ into which B and C embed in the obvious way over A. Since the relations E_k are independent, we can do this separately for each. Make sure to put all redundant k-tuples into the E_k-class reserved for them, number the E_k-classes which intersect A nontrivially, then go on to number the classes which just appear in B and C, and merge those classes which are assigned the same number, exactly as in Theorem 4.5.

For any structure A in K_{CPZ}, if $|A| = N$, then the number of n-tuples consisting of distinct elements from A reaches its maximum of $N!$ when $n = N$. When $n > N$, every n-tuple from A contains repeated elements. So the number of E_n-classes is bounded above by $N! + 1$ for all n, and $A \in K_{N!+1}$. Hence $K_{\mathrm{feq}} = \bigcup_{n=1}^{\infty} K_n$. So K_{CPZ} is a filtered Fraïssé class, and by Corollary 4.3, it suffices to show that each T_{K_n} is pseudofinite.

Let L'_n be the expanded language which includes, in addition to the relations E_k, $(n+1)$ k-ary relation symbols $C^0_k(\overline{x}), \ldots, C^n_k(\overline{x})$ for each k. Let K'_n be the class of finite L'_n-structures which are expansions of structures in K_n such that for all k, each

E_k-class is picked out by one of the C_k^i, with the class of redundant tuples picked out by C_k^0.

We have $K_n = \{A \upharpoonright L \mid A \in K_n'\}$, since every structure in K_n can be expanded to one in K_n' by labeling the classes for each equivalence relation. Suppose now that (A, B) is a one-point extension in K_n, and A' is an expansion of A to a structure in K_n'. If any k-tuple involving the new element b is part of a class which exists in A, then we label it by the appropriate C_k^i. If adding the new element adds new E_k-classes, then we simply label these classes by unused C_k^j (by the bound n on the number of classes, there will always be enough of the C_k^j). So K_n' is a Fraïssé expansion of K_n.

It remains to show that $T_{K_n'}$ has disjoint n-amalgamation for all n. Suppose we have a coherent $\mathscr{P}^-([n])$-family of types. As noted before, the relations E_k are independent, so we can handle them each separately. And the behavior of E_k is entirely determined by the behavior of the relations C_k^i, so it suffices to set these. But the only restriction here is that every k-tuple should satisfy exactly one C_k^i, and it should be C_k^0 if and only if the tuple contains repeated elements. So to solve our amalgamation problem, we simply assign relations from the C_k^i arbitrarily to those nonredundant k-tuples which are not already determined by the types in the family.

Hence $T_{K_n'}$ has disjoint n-amalgamation for all n, so it and its reduct T_{K_n} are pseudofinite. □

Proposition 4.12 *There is a sentence φ in T_{CPZ} such that*

$$\lim_{n\to\infty} \mu_n\big(\{A \in K_{\mathrm{CPZ}}(n) \mid A \models \varphi\}\big) = 0.$$

Proof An example of such a sentence φ is $\forall x\, \forall y\, \forall y'\, \exists z\, (E_1(x, z) \wedge E_2(y, y'; x, z))$. This sentence says that for all x, the function ρ_x mapping an element z in the E_1-class of x to the E_2-class of xz is surjective onto the E_2-classes. The sentence φ is in T_{CPZ}, since for any A in K_{feq} and elements a, b, and b' in A, we can embed A in a structure B in K_{feq} with an object c such that c is E_1-equivalent to a and (a, c) is E_2-equivalent to (b, b'). If $b = b'$, then we must take $a = c$; otherwise, we can add a new element satisfying this condition. So φ is implied by the relevant one-point extension axioms.

The measure μ_n amounts to picking an equivalence relation on the k-tuples of distinct elements from a set of size n for each k uniformly and independently. Since our sentence only involves E_1 and E_2, we just need to consider the equivalence relations on 1-tuples (there are n of them) and the nonredundant 2-tuples (of which there are $n^2 - n$). Citing again the fact that the expected number of equivalence classes in a random equivalence relation grows asymptotically as $\frac{n}{\log(n)}(1+o(1))$ (see [16, Proposition 8.8]), we see that with high probability there are more E_2-classes $(\frac{n^2}{\log(n^2-n)}(1+o(1)))$ than the size of the average E_1-class $(\log(n))$, in which case the function ρ_x is not surjective for all x, and the probability that φ is satisfied converges to zero. □

In this case, too, it would be interesting to know whether there is a zero-one law for the uniform measures.

References

[1] Ahlman, O., "Simple structures axiomatized by almost sure theories," *Annals of Pure and Applied Logic*, vol. 167 (2016), pp. 435–56. Zbl 06552723. MR 3465369. DOI 10.1016/j.apal.2016.02.001. 151

[2] Baldwin, J. T., *Categoricity*, vol. 50 of *University Lecture Series*, American Mathematical Society, Providence, 2009. Zbl 1183.03002. MR 2532039. DOI 10.1090/ulect/050. 151

[3] Baldwin, J. T., M. Koerwien, and M. C. Laskowski, "Disjoint amalgamation in locally finite AEC," *Journal of Symbolic Logic*, vol. 82 (2017), pp. 98–119. Zbl 06725055. MR 3631278. DOI 10.1017/jsl.2016.25. 151

[4] Blass, A., G. Exoo, and F. Harary, "Paley graphs satisfy all first-order adjacency axioms," *Journal of Graph Theory*, vol. 5 (1981), pp. 435–39. Zbl 0472.05058. MR 0635707. DOI 10.1002/jgt.3190050414. 140

[5] Blass, A., and B. Rossman, "Explicit graphs with extension properties," *Bulletin of the European Association for Theoretical Computer Science EATCS*, vol. 86 (2005), pp. 166–75. Zbl 1169.05373. MR 2163472. 140

[6] Brooke-Taylor, A., and D. Testa, "The infinite random simplicial complex," preprint, arXiv:1308.5517v1 [math.LO]. 151

[7] Cameron, P. J., *Oligomorphic Permutation Groups*, vol. 152 of *London Mathematical Society Lecture Note Series*, Cambridge University Press, Cambridge, 1990. Zbl 0813.20002. MR 1066691. DOI 10.1017/CBO9780511549809. 142

[8] Casanovas, E., R. Peláez, and M. Ziegler, "On many-sorted ω-categorical theories," *Fundamenta Mathematicae*, vol. 214 (2011), pp. 285–94. Zbl 1258.03035. MR 2845727. DOI 10.4064/fm214-3-5. 141, 155, 156

[9] Chatzidakis, Z., and E. Hrushovski, "Model theory of difference fields," *Transactions of the American Mathematical Society*, vol. 351 (1999), pp. 2997–3071. Zbl 0922.03054. MR 1652269. DOI 10.1090/S0002-9947-99-02498-8. 150

[10] Cherlin, G., "Combinatorial problems connected with finite homogeneity," pp. 3–30 in *Proceedings of the International Conference on Algebra, Part 3 (Novosibirsk, 1989)*, edited by L. A. Bokut, A. I. Kostrikin, and L. Ershov, vol. 131 of *Contemporary Mathematics*, American Mathematical Society, Providence, 1992. Zbl 0762.05078. MR 1175870. 140

[11] Cherlin, G., "Two problems on homogeneous structures, revisited," pp. 319–415 in *Model Theoretic Methods in Finite Combinatorics*, edited by M. Grohe and J. A. Makowsky, vol. 558 of *Contemporary Mathematics*, American Mathematical Society, Providence, 2011. Zbl 1261.03115. MR 3418640. DOI 10.1090/conm/558/11055. 140

[12] Cherlin, G., and E. Hrushovski, *Finite Structures with Few Types*, vol. 152 of *Annals of Mathematics Studies*, Princeton University Press, Princeton, 2003. Zbl 1024.03001. MR 1961194. 140

[13] Chernikov, A., and N. Ramsey, "On model-theoretic tree properties," *Journal of Mathematical Logic*, vol. 16 (2016). Zbl 06673303. MR 3580894. DOI 10.1142/S0219061316500094. 142, 155

[14] Džamonja, M., and S. Shelah, "On $\triangleleft^*$-maximality," *Annals of Pure and Applied Logic*, vol. 125 (2004), pp. 119–58. Zbl 1040.03029. MR 2033421. DOI 10.1016/j.apal.2003.11.001. 141, 155

[15] Erdős, P., D. J. Kleitman, and B. L. Rothschild, "Asymptotic enumeration of K_n-free graphs," pp. 19–27 in *Colloquio Internazionale sulle Teorie Combinatorie, Tomo II (Rome, 1973)*, vol. 17 of *Atti dei Convegni Lincei*, Accademia Nazionale dei Lincei, Rome, 1976. Zbl 0358.05027. MR 0463020. 140

[16] Flajolet, P., and R. Sedgewick, *Analytic Combinatorics*, Cambridge University Press, Cambridge, 2009. Zbl 1165.05001. MR 2483235. DOI 10.1017/CBO9780511801655. 155, 157

[17] Goodrick, J., B. Kim, and A. Kolesnikov, "Type-amalgamation properties and polygroupoids in stable theories," *Journal of Mathematical Logic*, vol. 15 (2015). Zbl 1372.03062. MR 3367778. DOI 10.1142/S021906131550004X. 151

[18] Henson, C. W., "A family of countable homogeneous graphs," *Pacific Journal of Mathematics*, vol. 38 (1971), pp. 69–83. Zbl 0219.05064. MR 0304242. 140

[19] Hodges, W., *Model Theory*, vol. 42 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, Cambridge, 1993. Zbl 0789.03031. MR 1221741. DOI 10.1017/CBO9780511551574. 142, 148

[20] Hrushovski, E., "Groupoids, imaginaries and internal covers," *Turkish Journal of Mathematics*, vol. 36 (2012), pp. 173–98. Zbl 1260.03065. MR 2912035. 151

[21] Kantor, W. M., M. W. Liebeck, and H. D. Macpherson, "$\aleph_0$-categorical structures smoothly approximated by finite substructures," *Proceedings of the London Mathematical Society (3)*, vol. 59 (1989), pp. 439–63. Zbl 0649.03018. MR 1014866. DOI 10.1112/plms/s3-59.3.439. 140

[22] Kim, B., and H.-J. Kim, "Notions around tree property 1," *Annals of Pure and Applied Logic*, vol. 162 (2011), pp. 698–709. Zbl 1241.03039. MR 2794257. DOI 10.1016/j.apal.2011.02.001. 142

[23] Kim, B., A. S. Kolesnikov, and A. Tsuboi, "Generalized amalgamation and n-simplicity," *Annals of Pure and Applied Logic*, vol. 155 (2008), pp. 97–114. Zbl 1153.03008. MR 2455560. DOI 10.1016/j.apal.2008.03.005. 151

[24] Kim, B., and A. Pillay, "Simple theories," *Annals of Pure and Applied Logic*, vol. 88 (1997), pp. 149–64. Zbl 0897.03036. MR 1600895. DOI 10.1016/S0168-0072(97)00019-5. 150

[25] Kim, B., and A. Pillay, "From stability to simplicity," *Bulletin of Symbolic Logic*, vol. 4 (1998), pp. 17–36. Zbl 0897.03034. MR 1609191. DOI 10.2307/421004. 141

[26] Kolesnikov, A. S., "n-simple theories," *Annals of Pure and Applied Logic*, vol. 131 (2005), pp. 227–61. Zbl 1067.03047. MR 2097228. DOI 10.1016/j.apal.2004.06.001. 151

[27] Macpherson, D., "A survey of homogeneous structures," *Discrete Mathematics*, vol. 311 (2011), pp. 1599–1634. Zbl 1238.03032. MR 2800979. DOI 10.1016/j.disc.2011.01.024. 147

[28] Oberschelp, W., "Monotonicity for structure numbers in theories without identity," pp. 297–308 in *Combinatoire et représentation du groupe symétrique (Strasbourg, 1976)*, vol. 579 of *Lecture Notes in Mathematics*, Springer, Berlin, 1977. Zbl 0365.02038. MR 0465846. 151

[29] Shelah, S., "Toward classifying unstable theories," *Annals of Pure and Applied Logic*, vol. 80 (1996), pp. 229–55. Zbl 0874.03043. MR 1402297. DOI 10.1016/0168-0072(95)00066-6. 141, 142, 150

[30] Thomas, S., "Reducts of random hypergraphs," *Annals of Pure and Applied Logic*, vol. 80 (1996), pp. 165–93. Zbl 0865.03025. MR 1402977. DOI 10.1016/0168-0072(95)00061-5. 147

Acknowledgments

This work is part of my Ph.D. dissertation "Infinitary limits of finite structures," completed at UC Berkeley under the supervision of Tom Scanlon, who I would like to thank for his support and for many helpful discussions. Nick Ramsey also had a great influence on this article: he told me about the theory T_{CPZ}, suggested Conjecture 1.4, and, after some effort, convinced me that SOP_1 is a natural dividing line.

Department of Mathematics
Indiana University
831 East 3rd St.
Bloomington, Indiana 47405
USA
akruckma@indiana.edu
http://pages.iu.edu/~akruckma