

Infinitary Limits of Finite Structures

by

Alex Kruckman

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Thomas Scanlon, Chair

Professor George Bergman

Professor Leo Harrington

Professor Wesley Holliday

Spring 2016

Infinitary Limits of Finite Structures

Copyright 2016
by
Alex Kruckman

Abstract

Infinitary Limits of Finite Structures

by

Alex Kruckman

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Thomas Scanlon, Chair

We study three distinct ways of assigning infinitary limits to classes of finite structures. We are primarily concerned with logically motivated questions about these limit objects and their theories, as well as connections and analogies between them.

In the first part, we consider limits of sequences of finite structures which converge with respect to densities of quantifier-free formulas, generalizing the dense graph limits and similar structural limits studied in combinatorics. A convergent sequence determines, as a limit object, a certain kind of probability measure on the space of countable structures with domain ω , which we call an ergodic structure. After developing the background theory, we take up the case of properly ergodic structures, which do not assign full measure to any single isomorphism class. The main result is a purely logical characterization of those theories in countable fragments of $L_{\omega_1, \omega}$ which arise as the theories of properly ergodic structures.

In the second part, we study categories consisting of finite structures and certain “strong” embeddings between them. We identify a necessary and sufficient condition for the well-definedness of strong embeddings between infinite direct limits from the category. This allows us to develop the natural generalization of classical Fraïssé theory in this context, with a focus on topology and genericity, in the sense of Baire category, in the space of direct limits with domain ω . We elaborate on an analogy between these generic limits and the measure limits in the first part, and we examine model-theoretic properties of generic limits.

In the third part, we take up logical limits, in the sense of ultraproducts and zero-one laws. In contrast to the first two parts, the limit theories here are always pseudofinite. We focus on countably categorical theories, which are essentially generic theories of small Fraïssé classes, and we show that higher amalgamation properties (disjoint n -amalgamation for all n) are sufficient to prove pseudofiniteness by a probabilistic argument. We examine relationships between pseudofiniteness, higher amalgamation, and model theoretic dividing lines, especially simplicity and NSOP₁. Our hope is that in “purely combinatorial” situations, pseudofiniteness should always be explained by randomness, via higher amalgamation; in an attempt to isolate these situations, we define the class of primitive combinatorial theories.

To Emmy

Who has brightened my life every day
of the last five years (and counting!).

Contents

Contents	ii
Acknowledgements	iv
0 Introduction	1
Conventions	7
I Measure	8
1 Quantifier-free convergence	9
1.1 Convergent sequences of finite structures	9
1.2 Measures as limit objects	15
1.3 The Aldous–Hoover–Kallenberg theorem	23
1.4 The theory of an ergodic structure	26
1.5 Trivial definable closure	33
2 Properly ergodic structures	37
2.1 Examples	37
2.2 Pithy Π_2 theories	40
2.3 Morley–Scott analysis of ergodic structures	42
2.4 Rooted models	47
2.5 Constructing properly ergodic structures	50
II Category	58
3 Strong embeddings and direct limits	59
3.1 Strong embedding classes	59
3.2 Direct limits and extendibility	61
3.3 Abstract elementary classes	68
3.4 Amalgamation	71
3.5 Universality and homogeneity	75

4	Generic limits	79
4.1	The space Dir_K and genericity	79
4.2	The generic limit	85
4.3	Genericity in classes without generic limits	90
4.4	Model theory and the generic limit	97
4.5	Strong robustness	101
III Logic		108
5	Pseudofinite countably categorical theories	109
5.1	Zero-one laws and ultraproducts	109
5.2	Small Fraïssé classes	116
5.3	Disjoint n -amalgamation	120
5.4	Two generic theories of equivalence relations	129
5.5	Primitive combinatorial theories	134
Bibliography		144

Acknowledgments

Many people, directly and indirectly, have helped me bring this thesis into the world, and I would like to express my gratitude to them.

First and foremost, my advisor, Tom Scanlon, who has always been extremely generous with his patience, time, and insight. He has taught me a lot of mathematics and a lot about being a mathematician, and I very much enjoyed the years I spent as his student.

George Bergman, Leo Harrington, and Wes Holliday, the other members of my thesis committee, for lending their time and support to this project. Early in my career as a graduate student at Berkeley, I took a course with George on universal algebra and a course with Leo on stability theory, and both had a profound impact on my view of mathematics.

Going back further, Dan Abramovich and Mike Rosen, my undergraduate senior thesis advisors at Brown University, who first encouraged me to learn something about model theory. They would have liked to see a few applications to algebraic geometry or number theory in these pages, but I hope to be able to make it up to them in the future.

Going back furthest of all, my parents, who I cannot thank enough.

My fellow graduate students in logic at Berkeley, especially Reid Dale, Gwyneth Harrison-Shermoen, Sridhar Ramesh, Nick Ramsey, Lawrence Valby, Andy Voellmer, and Michael Wan. They have been great friends, and I have learned an enormous amount from them.

Gabe Conant, Cameron Hill, Darío García, and Caroline Terry, who, for better or for worse, stimulated my interest in infinitary limits of finite structures, especially during the model theory semester at MSRI in Spring 2014.

Nate Ackerman, Cameron Freer, and Rehana Patel, who graciously invited me to collaborate on their project of understanding ergodic structures and taught me much of what I know about the subject. I am also grateful for their permission to include portions of our joint paper *Properly ergodic structures* in this thesis (see Chapter 0 for an indication of which sections are joint work).

Last but far from least, Emmy Brockman, my favorite person, and Finarfin, my favorite cat. Emmy's unflagging confidence and support kept my spirits up through the long hours of writing. Finarfin insisted on sitting on my lap throughout many of those long hours and probably contributed a few typos by walking across the keyboard when I wasn't looking.

Chapter 0

Introduction

This thesis is devoted to three distinct ways of assigning infinitary limit objects to classes of finite structures.

I Measure: Given a finite structure A and a quantifier-free formula $\varphi(\bar{x})$, the density of $\varphi(\bar{x})$ in A is the probability that a tuple \bar{a} sampled uniformly and independently from A satisfies $\varphi(\bar{x})$. A sequence of finite structures $\langle A_n \rangle_{n \in \omega}$ is said to converge if the density of $\varphi(\bar{x})$ in A_n converges for every quantifier-free formula $\varphi(\bar{x})$. Such a convergent sequence determines a probability measure on the space Str_L of L -structures with domain ω ; the measures which arise in this way are called ergodic structures.

For every sentence φ , an ergodic structure μ assigns measure 0 or 1 to the set of models of φ ; hence, μ determines a complete theory. It may happen that μ assigns measure 1 to the isomorphism class of a particular countable structure M . In this case, μ is said to be almost surely isomorphic to M . If this does not happen, μ is said to be properly ergodic.

II Category: A strong embedding class K is a category of finite structures and embeddings between them. If K satisfies an additional hypothesis I call extendibility, then direct limits from K are well behaved; in particular, the class of strong embeddings extends smoothly to K -direct limits. This makes it possible to develop a natural generalization of classical Fraïssé theory in this context. There is also a natural topological space Dir_K of K -direct limits with domain ω , and genericity in the sense of Fraïssé theory has a topological counterpart in Baire category.

If K satisfies the joint embedding property, then for every sentence φ , the set of models of φ is either meager or comeager in Dir_K ; hence, K determines a complete generic theory. Further, if K is countable up to isomorphism and satisfies the weak amalgamation property, then there is a particular countable structure M , called the generic limit of K , such that the isomorphism class of M is comeager in Dir_K . In this case, K is called a generalized Fraïssé class.

III Logic: Given a class K of finite structures and a sequence of probability measures $\langle \mu_n \rangle_{n \in \omega}$ on subsets $\langle \Omega_n \rangle_{n \in \omega}$ of K , K is said to have a zero-one law if the μ_n -measure of the set of models of φ in Ω_n converges to 0 or 1 for every sentence φ . This gives rise to a complete theory, the almost-sure theory of K . Alternatively, given an ultrafilter \mathfrak{U} on a subset Ω of K , we can take the ultraproduct $M = (\prod_{A \in \Omega} A) / \mathfrak{U}$ and look at the theory of M . In either case, we get a complete theory which is K -pseudofinite: every sentence has a finite model in K .

While these limit notions may be useful for studying the combinatorics of finite structures or finite model theory in the classical sense (e.g. [32]), this is not my main interest here. After developing the basic theory of these limit notions (from a somewhat unique perspective), I turn quickly to questions about the limit objects and limit theories themselves, motivated by three main themes.

- (1) The measure/category analogy: The first two limit notions described above seem quite unrelated at first. But when we view the limits as taking place on the spaces Str_L and Dir_K , they can be seen to behave in similar ways. In each case, certain subsets of a space of structures are considered “large”, with respect to a measure or with respect to Baire category, and these large sets give rise to a limit theory (containing a sentence φ if and only if the set of models of φ is large) and sometimes a limit structure M (if the isomorphism class of M is large). In the measure setting, we fix the topological space Str_L and vary the notion of “large sets” (the sets of measure 1) by considering different probability measures μ on Str_L . In the category setting, we fix the notion of “large sets” (the comeager sets) and vary the space Dir_K by considering different strong embedding classes K . Of course, this is but one instances of the analogy between measure and category in general topological spaces, outlined in the book [80], for example. I view it as a useful guiding principle.
- (2) The one model/many models dichotomy: In [2] and [3], Ackerman, Freer, and Patel characterized those structures M (respectively, those theories T) such that there exists an ergodic structure which is almost surely isomorphic to M (respectively, gives measure 1 to the models of T). Here, the properly ergodic case is addressed, and a characterization is given of the theories of properly ergodic structures. As a consequence of this characterization, it is shown that any such theory has continuum-many countable models up to isomorphism. This dichotomy also exists in the category setting, between generalized Fraïssé classes and classes which merely have the joint embedding property. The key property characterizing generalized Fraïssé classes, the weak amalgamation property, was discovered independently by Ivanov [48] and Kechris and Rosendal [56]. I show that this context too, if a strong embedding class has the disjoint amalgamation property, but no generic limit structure, then its (first-order) generic theory has continuum-many countable models up to isomorphism.

- (3) Agreements/disagreements between the limit notions: For certain classes of finite structures, all of our limit notions exist and agree. The canonical example is the class of all finite graphs. It is a Fraïssé class, whose generic limit is the random graph (a.k.a. the Rado graph). There is an ergodic structure μ , corresponding to the Erdős–Rényi random graph process on ω , which is almost surely isomorphic to the random graph, and which is the limit of a natural convergent sequence of finite graphs, the Paley graphs. Finally, the class of finite graphs has a zero-one law with respect to the uniform measures on the sets of labeled graphs of size n , and the almost sure theory is the theory of the random graph. This example is somewhat misleading; its agreeable behavior certainly seems to be the exception rather than the rule. In general, it is very hard to determine whether a given theory, even the generic theory of a Fraïssé class, is pseudofinite: the simplest non-trivial example, the theory of the generic triangle-free graph, is still quite mysterious (see [23]). Here I identify a combinatorial condition, disjoint n -amalgamation for all n , which is sufficient for agreement of our limit notions as in the case of finite graphs. And I suggest at the end of the thesis that in “purely combinatorial” situations, pseudofiniteness of countably categorical theories (and, more specifically, agreement of our limit notions) should always be explained by disjoint n -amalgamation.

Having described the general thrust, I will now give a more detailed outline of the contents of this thesis. Readers may wish to skip the outline and jump right into Part I, but they should not overlook the conventions established at the end of this chapter.

The majority of the material in Chapter 1 is not original to me; see [9], [29], [68], [83] for other presentations (all with very distinct flavors!) of what are essentially the same ideas. What is new is my perspective on the material, emphasizing logic and densities of quantifier-free formulas instead of substructure densities or homomorphism densities. This has several advantages: it allows us to work in a countably infinite language with no additional difficulties, it clarifies the appearance of a probability measure on Str_L as an infinitary limit object, and it leads naturally to a relatively intuitive variant of the Aldous–Hoover–Kallenberg theorem (Section 1.3).

In Section 1.1, I define the notion of convergence for quantifier-free densities and prove some basic properties. In particular, I show that in a finite language, it is equivalent to the notion of convergence for substructure densities studied in combinatorics. In Section 1.2, I explain how a convergent sequence of finite structures induces a probability measure on the space Str_L of structures with domain ω , and I characterize the ergodic structures, the measures which arise in this way. Ergodic structures are ergodic with respect to the “logic action” of S_∞ on Str_L by permutations of ω . It turns out that ergodicity is equivalent to a more elementary property, disjoint-independence, which, unlike ergodicity, is obvious from the limit construction.

Ergodic structures are, in particular, invariant for the logic action. In Section 1.3, I give a variant of the Aldous–Hoover–Kallenberg theorem, adapted to our context, which states that every invariant probability measure on Str_L can be represented as a random process that depends on independent sources of randomness at every finite subset of ω . In the special

case of graph limits, the most commonly studied limit objects are called graphons [68]; they are essentially AHK representations of the corresponding ergodic structures.

Turning to logic in Section 1.4, I show how to associate a complete theory in the infinitary logic $L_{\omega_1, \omega}$ (or any countable fragment) to an ergodic structure. This follows immediately from ergodicity, but I give an original elementary proof directly from the disjoint-independence property. Section 1.5 explains the relevance of trivial definable closure to the theories of ergodic structures and reviews the prior work of Ackerman, Freer, and Patel (in [2] and [3]), where trivial definable closure is the key condition. The work of Ackerman, Freer, and Patel does not emphasize the connection to limits of finite structures, instead studying ergodic structures for their own intrinsic interest. Nevertheless, it is this thread which we pursue in the next chapter.

The topic of Chapter 2 is properly ergodic structures. I provide a number of examples in Section 2.1, which illustrate some of their key features. Section 2.2 contains an exposition of a well-known construction which puts the models of a theory in a countable fragment of $L_{\omega_1, \omega}$ in natural bijection with the models of a first-order theory omitting a countable set of types. This construction, and the induced natural bijection between properly ergodic structures, will be useful in the rest of the chapter.

In Section 2.3, I undertake a Morley–Scott analysis of an ergodic structure μ (based on Morley’s proof [77] that the number of isomorphism classes of countable models of a sentence of $L_{\omega_1, \omega}$ is countable, \aleph_1 , or 2^{\aleph_0}). This gives a notion of Scott rank for ergodic structures and, in the properly ergodic case, allows us to find a countable fragment F of $L_{\omega_1, \omega}$ in which there is a formula $\chi(\bar{x})$ which has positive measure, but which picks out continuum-many F -types, each of which has measure 0. As a corollary of this analysis, for any properly ergodic structure μ , the complete $L_{\omega_1, \omega}$ -theory $\text{Th}(\mu)$ has no models (of any cardinality), but for any countable fragment F , the F -theory $\text{Th}_F(\mu)$ has continuum-many models up to isomorphism. This can be viewed as an analog of Vaught’s Conjecture in this context.

In Section 2.4, I introduce the notion of a rooted model of a theory. A model M is rooted if a collection of non-isolated types (e.g. the continuum-many types of measure 0 coming from the Morley–Scott analysis) has few realizations in M in a precise sense. I use the Aldous–Hoover–Kallenberg theorem to show that a model sampled from a properly ergodic measure is almost surely rooted.

In Section 2.5, I use a single rooted model of a theory T with trivial definable closure to guide the construction, via an inverse limit of finite structures, of a rooted Borel model \mathbb{M} , equipped with an atomless probability measure ν . Then a properly ergodic structure μ is obtained by sampling from (\mathbb{M}, ν) . The inverse limit construction is a refinement of the methods from [2], [3], and [5], which in turn generalized a construction due to Petrov and Vershik [81]. Together, the results of Sections 2.3–2.5 give a characterization of the theories of properly ergodic structures in countable fragments of $L_{\omega_1, \omega}$.

Chapter 2, as well as portions of the preliminary material in Chapter 1, are adapted from the paper *Properly ergodic structures* [6], which I coauthored with Nate Ackerman, Cameron Freer, and Rehana Patel. Since I have chosen to include it in my thesis, I should probably comment on my contribution to this joint work. Section 2.2 and the excerpts from the paper

which appear in Chapter 1 constitute background material; they do not contain original results. The examples in Section 2.1 (except Example 2.1.1, which is otherwise attributed) and the results in Sections 2.3 and 2.4 are solely my work. Section 2.5 is joint work: here, the main ideas and the outline of the construction are due to Nate, Cameron, and Rehana, while I worked out the details and identified the necessity of the rootedness condition in the argument. The actual text was written by me, though my coauthors contributed many valuable suggestions on presentation.

In Chapters 3 and 4, I turn to the category limit. Ever since Hrushovski’s use [45] of a generalized Fraïssé construction to produce a strongly minimal set violating the Trichotomy Conjecture, there has been much interest in “Hrushovski constructions”: generalized Fraïssé constructions for classes of finite structures and strong embeddings, where the class of strong embeddings is defined in terms of a “pre-dimension function”. While an industry developed around generating counterexamples in stability theory using Hrushovski constructions, there has also been some quite general work not aimed at producing specific examples: see, for example, the papers [12], [34], and [82].

My aim here is to work in the greatest generality possible, starting from a totally arbitrary category of finite structures and “strong” embeddings and identifying the minimal assumptions necessary to make the basic theory work. The aim is not applications, but foundations for future variants of the Fraïssé construction. A similar approach was taken in [66], where Kueker and Laskowski did away with pre-dimension functions, but instead required that the class of strong embeddings be appropriately definable by universal conditions. While it is abstract, my approach is not purely categorical, since the objects and arrows in my categories are really finite structures and embeddings. Purely categorical approaches to Fraïssé theory have been studied by others, see [18] or [65], for example.

Strong embedding classes are defined in Section 3.1. The goal, of course, is not to study strong embedding classes, but to study their direct limits, and I explain in Section 3.2 how the strong embeddings between finite structures in a class K naturally induce strong embeddings between K -direct limits. There is a hitch: whether an embedding between infinite K -direct limits is strong may depend on their presentations as K -direct limits; equivalently, strong embeddings may not be closed under isomorphism. I call a strong embedding class extendible if it does not have this defect, and I give a combinatorial condition which is equivalent to extendibility and one which implies it. Extendibility is the bare minimum assumption on a strong embedding class, necessary for essentially everything that comes after.

In 3.3, I show that the class of direct limits of an extendible strong embedding class satisfying the “coherence” condition is an abstract elementary class, and I note some connections with the theory of AECs.

Sections 3.4 and 3.5 adapt the notions of amalgamation and homogeneity from classical Fraïssé theory to our context and prove straightforward generalizations of the usual connections between these notions. The most important adjustment here is the generalization of the usual amalgamation property to the weak amalgamation property and the corresponding notions of weak- K -homogeneity and weak- K -ultrahomogeneity.

Topology gets involved in Chapter 4. In Section 4.1, I introduce the space Dir_K of K -

direct limits with domain ω and compare its topology to that of Str_L . In the case of classical Fraïssé theory, which deals with hereditary classes in finite languages, Dir_K is just a closed subspace of Str_L . I define generic properties of labeled K -direct limits via an infinite game for two players and use the Banach–Mazur game to connect this notion to topological genericity in Dir_K .

In Section 4.2, I prove the theorem (due to Ivanov [48] and Kechris and Rosendal [56] in less general contexts) that the weak amalgamation property (along with the joint embedding property and countability up to isomorphism) is equivalent to the existence of a countable K -direct limit with generic isomorphism class. In Section 4.3, I prove that the joint embedding property is necessary and sufficient for every isomorphism-invariant determined property to be generic or co-generic; hence a class with the joint embedding property has a complete generic theory in $L_{\omega_1, \omega}$ or any countable fragment. I elaborate on the measure/category analogy with Part I and show that if a class has the joint embedding property but lacks the weak amalgamation property, then every generic property is satisfied by continuum-many countable K -direct limits up to isomorphism.

I examine model-theoretic properties of the generic limit in Section 4.4, characterizing countable categoricity and giving sufficient conditions for atomicity, model completeness, and quantifier elimination under an additional topological definability assumption. In Section 4.5, I study a slight generalization (“strong robustness”) of a strengthening (“super-robustness”) due to Hill [42] of a notion (“robustness”) due to Macpherson and Steinhorn [74] which facilitates the study of definability in classes of finite structures. Looking ahead to Chapter 5, a connection is drawn to pseudofiniteness of the generic limit.

Section 5.1 begins with a very general definition of convergence laws and zero-one laws for classes of finite structures. In this generality, I observe that every K -pseudofinite theory T (meaning every sentence in T has a finite model in K) arises both as an ultraproduct of structures in K and as the almost-sure theory of K for some sequence of measures. I also examine several generic theories of small Fraïssé classes in detail, comparing and contrasting our three limit notions, as motivation to focus on pseudofiniteness for countably categorical theories. In Section 5.2, I present the background on small Fraïssé classes, defining the canonical language as a bridge to general countably categorical theories. I also introduce the notions of Fraïssé expansion and filtered Fraïssé class that will be useful in the next sections.

The heart of the chapter is Section 5.3 on disjoint n -amalgamation. After introducing these properties and providing some context about the role of n -amalgamation properties in model theory, I prove that disjoint n -amalgamation for all n is a sufficient condition for pseudofiniteness, and I explain how this theorem generalizes and unifies some previous work. Combining this result with Fraïssé expansions and filtered Fraïssé classes yields a strategy for showing that generic theories of Fraïssé classes are pseudofinite. But I also give a negative result, showing that this method cannot be used to show that the theory of the generic triangle-free graph is pseudofinite.

The method is applied effectively in Section 5.4, however, to show that two generic theories of equivalence relations, T_{feq}^* and T_{CPZ} , are pseudofinite. These results are interesting from the point of view of model-theoretic dividing lines, since to my knowledge they are the

first unsimple countably categorical theories that have been proven to be pseudofinite. Both these theories have the combinatorial property NSOP₁ (demonstrated here for T_{CPZ} for the first time), so they lie just outside the class of simple theories.

Section 5.5 is somewhat more speculative. The idea is that in certain “purely combinatorial” settings (e.g. excluding algebra or stability), pseudofiniteness of countably categorical theories should always be explained by a probabilistic argument facilitated by disjoint n -amalgamation for all n . In an attempt to isolate such a purely combinatorial setting, I define the class of primitive combinatorial theories, which are countably categorical theories with no algebraicity and no interesting definable equivalence relations on realizations of complete types over finite sets. I present a few results and many conjectures about the primitive combinatorial theories.

Portions of Chapter 5 are adapted from my paper *Disjoint n -amalgamation and pseudofinite countably categorical theories* [64], which has been submitted for publication.

Conventions

The following conventions will remain in place throughout the thesis. Starting now, we will refer to ourselves in the plural. L is always a countable relational language. Structures are always L -structures unless otherwise specified. We allow L to contain 0-ary relations (i.e. propositional symbols), and we allow empty structures. Our basic logical language includes a primitive symbol $=$, interpreted as equality, and primitive atomic formulas \top and \perp , interpreted as true and false, respectively. Formulas are formulas of first-order logic unless otherwise specified (we will sometimes deal with infinitary formulas in $L_{\omega_1, \omega}$). Every formula $\varphi(\bar{x})$ comes with a variable context \bar{x} , which includes the variables which appear free in φ , but may be larger. We write $l(\bar{x})$ for the length of the tuple \bar{x} . The variables in \bar{x} are always distinct. A formula $\varphi(\bar{x})$ in the variable context x_1, \dots, x_k is explicitly non-redundant if it implies $x_i \neq x_j$ for all $i \neq j$. Similarly, a type $p(\bar{x})$ is non-redundant if it contains $x_i \neq x_j$ for all $i \neq j$, and a tuple $\bar{a} = (a_1, \dots, a_k)$ is non-redundant if $a_i \neq a_j$ for all $i \neq j$. As usual, A^k denotes the set of k -tuples from A , and we adopt the less common notation $A^{\bar{k}}$ for the set of non-redundant k -tuples from A . We write $\|\bar{a}\|$ for the set of elements in the tuple \bar{a} . We say that a tuple \bar{a} enumerates a set A if \bar{a} is non-redundant and $\|\bar{a}\| = A$. We write $[n]$ for the set $\{0, \dots, n-1\}$.

Part I

Measure

Chapter 1

Quantifier-free convergence

1.1 Convergent sequences of finite structures

We begin with a bit of notation and a basic fact. Given a structure A and a formula $\varphi(\bar{x})$, we define $\varphi(A) = \{\bar{a} \in A^{l(x)} \mid A \models \varphi(\bar{a})\}$. If A and B are structures, we denote by $\binom{B}{A}$ the set of substructures of B which are isomorphic to A , and we denote by $\text{Emb}(A, B)$ the set of embeddings of A into B . $\text{Aut}(A)$ is the group of automorphisms of A .

Lemma 1.1.1. *For all structures A and B , we have $|\text{Emb}(A, B)| = |\text{Aut}(A)| \cdot \binom{B}{A}$. If L is a finite language and A is a finite structure of size k , enumerated by a tuple \bar{a} , there is a quantifier-free formula $\varphi_A(\bar{x})$ such that for all structures C and all k -tuples \bar{c} from C , $C \models \varphi_A(\bar{c})$ if and only if the map $a_i \mapsto c_i$ is an embedding $A \rightarrow C$. Hence $|\varphi_A(C)| = |\text{Emb}(A, C)|$.*

Proof. For each $A' \in \binom{B}{A}$, fix an isomorphism $f_{A'}: A \rightarrow A'$. Then there is a bijection $\text{Aut}(A) \times \binom{B}{A} \cong \text{Emb}(A, B)$ given by $(\sigma, A') \mapsto f_{A'} \circ \sigma$. To see that it is surjective, note that if $g: A \rightarrow B$ is an embedding with range $A' = g(A)$, $(f_{A'}^{-1} \circ g) \in \text{Aut}(A)$, and $(f_{A'}^{-1} \circ g, A') \mapsto g$.

Now set $\text{Diag}(A) = \{\theta(\bar{x}) \mid \theta \text{ is atomic or negated atomic, and } A \models \theta(\bar{a})\}$, the atomic diagram of A . Since L is finite, $\text{Diag}(A)$ is finite, so we can define $\varphi_A(\bar{x})$ to be the conjunction of all the formulas in $\text{Diag}(A)$. \square

Definition 1.1.2. Given a quantifier-free formula $\varphi(\bar{x})$ in n variables and a non-empty finite structure A , the **density** of $\varphi(\bar{x})$ in A , denoted $P(\varphi(\bar{x}); A)$, is the probability that n elements sampled uniformly and independently from A satisfy φ :

$$P(\varphi(\bar{x}); A) = \frac{|\varphi(A)|}{|A|^n}.$$

Lemma 1.1.3. *Quantifier-free densities satisfy the following basic properties.*

- (1) *If $\varphi(\bar{x})$ is equivalent to $\psi(\bar{x})$, then $P(\varphi(\bar{x}); A) = P(\psi(\bar{x}); A)$.*
- (2) *Given a formula $\varphi(\bar{x})$, we may extend its variable context by adding dummy variables \bar{y} , obtaining a new formula $\varphi^*(\bar{x}, \bar{y})$. Then $P(\varphi(\bar{x}); A) = P(\varphi^*(\bar{x}, \bar{y}); A)$.*

(3) Given a formula $\varphi(\bar{x})$, we may replace \bar{x} by a new tuple of variables \bar{y} of the same length, obtaining a new formula $\varphi^*(\bar{y})$. Then $P(\varphi(\bar{x}); A) = P(\varphi^*(\bar{y}); A)$.

(4) $P(\top; A) = 1$, and $P(\perp; A) = 0$.

(5) $P(\varphi(\bar{x}) \vee \psi(\bar{x}); A) = P(\varphi(\bar{x}); A) + P(\psi(\bar{x}); A) - P(\varphi(\bar{x}) \wedge \psi(\bar{x}); A)$.

(6) $P(\varphi(\bar{x}) \wedge \psi(\bar{y}); A) = P(\varphi(\bar{x}); A) \cdot P(\psi(\bar{y}); A)$ if \bar{x} and \bar{y} are disjoint tuples of variables.

Proof. With notation as above, we have

$$(1): P(\varphi(\bar{x}); A) = \frac{|\varphi(A)|}{|A|^{l(\bar{x})}} = \frac{|\psi(A)|}{|A|^{l(\bar{x})}} = P(\psi(\bar{x}); A).$$

$$(2): P(\varphi(\bar{x}); A) = \frac{|\varphi(A)|}{|A|^{l(\bar{x})}} = \frac{|\varphi(A) \times A^{l(\bar{y})}|}{|A|^{l(\bar{x})}|A|^{l(\bar{y})}} = \frac{|\varphi^*(A)|}{|A|^{l(\bar{x})+l(\bar{y})}} = P(\varphi^*(\bar{x}, \bar{y}); A).$$

$$(3): P(\varphi(\bar{x}); A) = \frac{|\varphi(A)|}{|A|^{l(\bar{x})}} = \frac{|\varphi(A)|}{|A|^{l(\bar{y})}} = P(\varphi(\bar{y}); A).$$

$$(4): P(\top; A) = \frac{|A^0|}{|A|^0} = 1, \text{ and } P(\perp; A) = \frac{|\emptyset|}{|A|^0} = 0.$$

$$(5): P(\varphi(\bar{x}) \vee \psi(\bar{x}); A) = \frac{|\varphi(A) \cup \psi(A)|}{|A|^{l(\bar{x})}} = \frac{|\varphi(A)| + |\psi(A)| - |\varphi(A) \cap \psi(A)|}{|A|^{l(\bar{x})}}$$

$$= \frac{|\varphi(A)|}{|A|^{l(\bar{x})}} + \frac{|\psi(A)|}{|A|^{l(\bar{x})}} - \frac{|\varphi(A) \cap \psi(A)|}{|A|^{l(\bar{x})}} = P(\varphi(\bar{x}); A) + P(\psi(\bar{x}); A) - P((\varphi \wedge \psi)(\bar{x}); A).$$

$$(6): P(\varphi(\bar{x}) \wedge \psi(\bar{y}); A) = \frac{|\varphi(A) \times \psi(A)|}{|A|^{l(\bar{x})+l(\bar{y})}} = \frac{|\varphi(A)|}{|A|^{l(\bar{x})}} \cdot \frac{|\psi(A)|}{|A|^{l(\bar{y})}} = P(\varphi(\bar{x}); A) \cdot P(\psi(\bar{y}); A). \quad \square$$

Definition 1.1.4. A sequence $\langle B_n \rangle_{n \in \omega}$ of finite structures **converges (for quantifier-free densities)** if $\lim_{n \rightarrow \infty} P(\varphi(\bar{x}); B_n)$ exists in $[0, 1]$ for every quantifier-free formula $\varphi(\bar{x})$.

Remark 1.1.5. For any non-empty finite structure A , we have $P(x = y; A) = 1/|A|$. Hence if $\langle B_n \rangle_{n \in \omega}$ converges, we have $\lim_{n \rightarrow \infty} P(x = y; B_n) = \lim_{n \rightarrow \infty} 1/|B_n|$, and this limit is either 0 (if $\lim_{n \rightarrow \infty} |B_n| = \infty$) or $1/k$ (if $|B_n|$ is eventually constant, equal to k).

If L is finite, something stronger is true in the latter case: the sequence is eventually constant up to isomorphism. Let A be a structure of size k , and consider the formula $\varphi_A(\bar{x})$ as in Lemma 1.1.1. Then for any other structure B of size k ,

$$P(\varphi_A(\bar{x}); B) = \frac{|\text{Aut}(A)| \binom{B}{A}}{|B|^k} = \begin{cases} \frac{|\text{Aut}(A)|}{k^k} & \text{if } B \cong A \\ 0 & \text{otherwise} \end{cases}.$$

Now since L is finite, there are only finitely many structures of size k up to isomorphism, say $\{A_1, \dots, A_m\}$. So if there is some N such that $|B_n| = k$ for all $n > N$, then one of the formulas $\varphi_{A_i}(\bar{x})$ has a constant positive density infinitely often, and thus by convergence there is $N' \geq N$ such that $P(\varphi_{A_i}(\bar{x}); B_n) > 0$, and hence $B_n \cong A_i$, for all $n > N'$.

On the other hand, in an infinite language, there are convergent sequences of constant size which are not eventually constant up to isomorphism.

Example 1.1.6. Let $L = \{P_n\}_{n \in \omega}$, with each P_n a unary relation symbol. Let $\langle B_n \rangle_{n \in \omega}$ be a sequence of structures of size 1, $B_n = \{*\}$, such that $B_n \models P_i(*)$ if and only if $i \leq n$. This sequence is convergent, since the densities of instances of $P_n(x_i)$ and $x_i = x_j$ converge to 1, while the densities of their negations converge to 0.

Despite the fact that non-trivial behavior can occur on sequences of bounded size in an infinite language, we will be interested primarily in sequences whose sizes grow to infinity.

Lemma 1.1.7. *Let $\langle B_n \rangle_{n \in \omega}$ be a sequence of finite structures such that $\lim_{n \rightarrow \infty} |B_n| = \infty$. Then $\langle B_n \rangle_{n \in \omega}$ converges if and only if $\lim_{n \rightarrow \infty} P(\varphi(\bar{x}); B_n)$ exists for every explicitly non-redundant quantifier-free formula $\varphi(\bar{x})$.*

Proof. First, we claim that for any quantifier-free formula $\psi(\bar{x})$ and any pair of variables x_i and x_j from \bar{x} , $\lim_{n \rightarrow \infty} P(\psi(\bar{x}); B_n) = \lim_{n \rightarrow \infty} P(\psi(\bar{x}) \wedge (x_i \neq x_j); B_n)$. Indeed, for all n ,

$$P(\psi(\bar{x}); B_n) = P(\psi(\bar{x}) \vee (x_i \neq x_j); B_n) - P((x_i \neq x_j); B_n) + P(\psi(\bar{x}) \wedge (x_i \neq x_j); B_n). \quad (\dagger)$$

But by Remark 1.1.5,

$$1 \geq \lim_{n \rightarrow \infty} P(\psi(\bar{x}) \vee (x_i \neq x_j); B_n) \geq \lim_{n \rightarrow \infty} P((x_i \neq x_j); B_n) = 1,$$

so taking the limit on both sides of (\dagger) establishes the claim.

Now assume that $\lim_{n \rightarrow \infty} P(\varphi(\bar{x}); B_n)$ exists for every explicitly non-redundant quantifier-free formula $\varphi(\bar{x})$. Let $\psi(\bar{x})$ be an arbitrary quantifier-free formula in the variables x_1, \dots, x_k , and let $\theta(\bar{x})$ be the formula $\bigwedge_{1 \leq i < j \leq k} x_i \neq x_j$. By repeatedly applying the above claim, we see that

$$\lim_{n \rightarrow \infty} P(\psi(\bar{x}); B_n) = \lim_{n \rightarrow \infty} P(\psi(\bar{x}) \wedge \theta(\bar{x}); B_n),$$

and $\psi(\bar{x}) \wedge \theta(\bar{x})$ is explicitly non-redundant, so the limit exists. \square

The notion of convergence for quantifier-free densities is closely related to other notions of convergence which have been well-studied by combinatorialists, especially in the special case of graphs. We mention here the notion of convergence for substructure densities and point out that, in the case of a finite language, it is equivalent to convergence for quantifier-free densities. There are similar notions of convergence for homomorphism densities and convergence for embedding densities, which are well known to be equivalent to convergence for substructure densities. For a comprehensive reference, see [68].

Definition 1.1.8. Let A and B be finite structures. The **density** of A in B , denoted $p(A; B)$, is the probability that a subset of B of size $|A|$, chosen uniformly at random, is the domain of a substructure isomorphic to A :

$$p(A; B) = \frac{\binom{B}{A}}{\binom{|B|}{|A|}}.$$

If $|A| > |B|$, we set $p(A; B) = 0$. A sequence $\langle B_n \rangle_{n \in \omega}$ of finite structures **converges for substructure densities** if $\lim_{n \rightarrow \infty} p(A; B_n)$ exists in $[0, 1]$ for every finite structure A .

Theorem 1.1.9. *Suppose that L is finite and $\langle B_n \rangle_{n \in \omega}$ is a sequence of finite structures such that $\lim_{n \rightarrow \infty} |B_n| = \infty$. Then $\langle B_n \rangle_{n \in \omega}$ converges for quantifier-free densities if and only if it converges for substructure densities.*

Proof. We first note that for fixed k ,

$$\lim_{n \rightarrow \infty} \frac{\binom{n}{k}}{n^k} = \lim_{n \rightarrow \infty} \frac{n(n-1)\dots(n-k+1)}{n^k k!} = \frac{1}{k!}.$$

Suppose that $\langle B_n \rangle_{n \in \omega}$ converges for quantifier-free densities, and let A be a finite structure of size k . Let $\varphi_A(\bar{x})$ be the formula for A defined by Lemma 1.1.1. Now,

$$\begin{aligned} \lim_{n \rightarrow \infty} p(A; B_n) &= \lim_{n \rightarrow \infty} \frac{\left| \binom{B_n}{A} \right|}{\binom{|B_n|}{|A|}} = \lim_{n \rightarrow \infty} \frac{|\text{Emb}(A, B_n)|}{|\text{Aut}(A)|} \cdot \frac{k!}{|B_n|^k} \\ &= \frac{k!}{|\text{Aut}(A)|} \lim_{n \rightarrow \infty} \frac{|\varphi_A(B_n)|}{|B_n|^k} = \frac{k!}{|\text{Aut}(A)|} \lim_{n \rightarrow \infty} P(\varphi_A(\bar{x}); B_n), \end{aligned}$$

hence the limit exists.

Conversely, suppose that $\langle B_n \rangle_{n \in \omega}$ converges for substructure densities, and let $\psi(\bar{x})$ be a quantifier-free formula. By Lemma 1.1.7, we may assume that $\psi(\bar{x})$ is explicitly non-redundant. Then, since L is finite, $\psi(\bar{x})$ is equivalent to $\bigvee_{i=1}^m \varphi_{A_i}(\bar{x})$, where the A_i are the finitely many structures (up to isomorphism) enumerated by tuples satisfying $\psi(\bar{x})$ and the $\varphi_{A_i}(\bar{x})$ are as in Lemma 1.1.1. Note that these formulas are pairwise inconsistent and that a structure may occur more than once on this list, if it can be enumerated multiple ways by a tuple satisfying $\psi(\bar{x})$. Now we have

$$\begin{aligned} \lim_{n \rightarrow \infty} P(\psi(\bar{x}); B_n) &= \lim_{n \rightarrow \infty} \sum_{i=1}^m P(\varphi_{A_i}(\bar{x}); B_n) \\ &= \lim_{n \rightarrow \infty} \sum_{i=1}^m \frac{|\varphi_{A_i}(B_n)|}{|B_n|^k} \\ &= \sum_{i=1}^m \lim_{n \rightarrow \infty} \frac{|\text{Aut}(A_i)| \left| \binom{B_n}{A_i} \right|}{k! \binom{|B_n|}{k}} \\ &= \sum_{i=1}^m \left(\frac{|\text{Aut}(A_i)|}{k!} \right) \lim_{n \rightarrow \infty} p(A_i, B_n), \end{aligned}$$

hence the limit exists. □

While convergence for quantifier-free densities and convergence for substructure densities are equivalent in a finite language, only the former has a useful generalization to infinite languages, as evidenced by the following example. For this reason, when we speak of convergence below, we will always mean convergence for quantifier-free densities.

Example 1.1.10. Let $\langle B_n \rangle_{n \in \omega}$ be the sequence defined in Example 1.1.6. We saw there that $\langle B_n \rangle_{n \in \omega}$ converges for quantifier-free densities. $\langle B_n \rangle_{n \in \omega}$ also converges for substructure densities, but in a rather trivial way: $\lim_{n \rightarrow \infty} p(A; B_n) = 0$ for all non-empty A . When $|A| > 1$, this is by definition. When $|A| = 1$, $p(A; B_n)$ is positive for at most one n , when $B_n \cong A$. Unlike the limiting quantifier-free densities, the limiting substructure densities fail to capture the asymptotic behavior of $\langle B_n \rangle_{n \in \omega}$ in a meaningful way.

The above example can also be adjusted to produce a sequence which trivially converges for substructure densities as above, but fails to converge for quantifier-free densities. We can code an element $\eta \in 2^\omega$ by a structure $A = \{*\}$ of size 1 by setting $A \models P_n(*)$ if and only if $\eta(n) = 1$. If $\langle \eta_n \rangle_{n \in \omega}$ is a non-convergent sequence in 2^ω in which no element of 2^ω appears infinitely often, then the corresponding sequence $\langle B_n \rangle_{n \in \omega}$ of structures of size 1 has the desired properties.

We can easily obtain similar examples in which the sizes of the structures grow to infinity. Simply replace the structures B_n above by structures B'_n of size n on which each P_i holds of all or none of the elements, according to whether it holds of the unique element in B_n .

Example 1.1.11. The most well-known example of a convergent sequence of finite structures is the Paley graphs. Given a prime number p congruent to 1 mod 4, we define a graph G_p with domain \mathbb{F}_p by setting aRb if and only if $a - b$ is a non-zero square in \mathbb{F}_p (note that -1 is a square in \mathbb{F}_p , so the edge relation is symmetric).

Now letting $\langle p_n \rangle_{n \in \omega}$ be the sequence of primes which are 1 mod 4, $\langle G_{p_n} \rangle_{n \in \omega}$ converges. It is a fact that if $\varphi(\bar{x})$ is a consistent formula describing n edges and m non-edges

$$\left(\bigwedge_{i=1}^n x_{k_i} R x_{k'_i} \right) \wedge \left(\bigwedge_{j=1}^m \neg x_{l_j} R x_{l'_j} \right),$$

then $\lim_{n \rightarrow \infty} P(\varphi(\bar{x}); G_{p_n}) = 2^{-(n+m)}$. Using this, and the fact that the formulas $x = y$, xRx , and $xRy \wedge \neg yRx$ have limiting density 0, it is straightforward to compute the limiting density of any quantifier-free formula.

The first analyses of the Paley graphs used very non-trivial number theory (see, for example, [14]). Chung, Graham, and Wilson in [26] showed that finitely many instances of the limiting densities asserted above suffice to imply the rest, and these finitely many instances, which they listed explicitly, can be checked in an elementary way for the Paley graphs. This insight has led to the burgeoning field of finitely forcible graphons [71].

The Paley graphs can be viewed as a deterministic version of the Erdős–Rényi random graph model, in which two vertices are connected by an edge with probability 1/2, independently over all pairs of vertices. For this reason, the Paley graphs (and all sequences

satisfying the Chung–Graham–Wilson conditions) are called quasi-random. We will revisit this connection in Example 1.4.16 below.

Since we have defined a natural notion of convergence for finite structures, it makes sense to ask what a convergent sequence of finite structures converges *to*. In Section 1.2, we will demonstrate that a certain class of probability measures on the space of countable labeled structures serve as natural limit objects for sequences whose sizes grow to infinity. But first we take a more obvious and elementary approach: since a sequence converges if and only if a limiting density exists for every quantifier-free formula, we take as the limit object the tuple of all these limiting densities.

Let L_{qf} be the (countable) set of all quantifier-free formulas with variables from a fixed countable supply $\{x_n \mid n \in \omega\}$, and consider the space $[0, 1]^{L_{\text{qf}}}$ of L_{qf} -indexed tuples of real numbers in $[0, 1]$. Note that $[0, 1]^{L_{\text{qf}}}$ is compact (Tychonoff's theorem) and metrizable (enumerating the formulas in L_{qf} as $\{\varphi_n \mid n \in \omega\}$, a compatible metric is given by $d((a_{\varphi_n}), (b_{\varphi_n})) = \sum_{n \in \omega} 2^{-n} |a_{\varphi_n} - b_{\varphi_n}|$). In particular, it is a Polish space (see [55, Proposition I.4.2]).

Now to each finite structure B , we can associate the point $P(B) = (P(\varphi(\bar{x}); B))_{\varphi \in L_{\text{qf}}}$ in $[0, 1]^{L_{\text{qf}}}$, and since the product topology on $[0, 1]^{L_{\text{qf}}}$ is the topology of pointwise convergence, the sequence $\langle B_n \rangle_{n \in \omega}$ converges if and only if the sequence $\langle P(B_n) \rangle_{n \in \omega}$ converges in $[0, 1]^{L_{\text{qf}}}$ to the point $(\lim_{n \rightarrow \infty} P(\varphi(\bar{x}); B_n))_{\varphi \in L_{\text{qf}}}$.

From this point of view, letting $F = \{P(B) \mid B \text{ finite}\} \subseteq [0, 1]^{L_{\text{qf}}}$, it makes sense to call the closure \bar{F} , which is again a compact Polish space, the **space of limits of finite structures**. Indeed, since $[0, 1]^{L_{\text{qf}}}$ is a metric space, every point in \bar{F} is the limit of a sequence from F .

As we have seen in the case of the Paley graphs (Example 1.1.11), it can be rather complicated to verify that an explicit sequence of finite structures converges. However, it is an immediate consequence of representation of convergence as convergence in $[0, 1]^{L_{\text{qf}}}$ that convergent sequences are easy to come by.

Proposition 1.1.12. *Every sequence of finite structures has a convergent subsequence.*

Proof. Let $\langle B_n \rangle_{n \in \omega}$ be a sequence of finite structures. Since $[0, 1]^{L_{\text{qf}}}$ is a compact metrizable space, it is sequentially compact, and hence $\langle P(B_n) \rangle_{n \in \omega}$ has a convergent subsequence, and the corresponding subsequence of $\langle B_n \rangle_{n \in \omega}$ also converges. \square

Another way to obtain convergent sequences is to change our definition slightly and consider convergence with respect to an ultrafilter. We briefly recall the basic theory of filter and ultrafilter convergence.

Definition 1.1.13. Let X be a topological space, and let \mathcal{F} be a filter on X . We say that $x \in X$ is an \mathcal{F} -limit point if for every open neighborhood U of x , $U \in \mathcal{F}$.

Proposition 1.1.14. *Let X be a topological space, and let \mathcal{F} be a proper filter on X .*

- (1) *If X is Hausdorff, then there is at most one \mathcal{F} -limit point.*
- (2) *If X is compact and \mathcal{F} is an ultrafilter, then there is at least one \mathcal{F} -limit point.*
- (3) *If \mathcal{F} is supported on S (i.e. $S \in \mathcal{F}$), then every \mathcal{F} -limit point is in the closure of S .*

Proof. Suppose that X is Hausdorff, and $x \neq y$ are both \mathcal{F} -limit points. Then there are disjoint open neighborhoods U_x and U_y of x and y , respectively, and $U_x, U_y \in \mathcal{F}$. But then $U_x \cap U_y = \emptyset \in \mathcal{F}$, contradiction.

Now let \mathcal{F} be an ultrafilter, and suppose for contradiction that X is compact, but there is no \mathcal{F} -limit point. Then for every $x \in X$, there is an open set U_x containing x with $U_x \notin \mathcal{F}$. Note that the U_x cover X , and their complements U_x^c are in \mathcal{F} . By compactness there is a finite subcover $X = \bigcup_{i=1}^n U_{x_i}$. But then $\bigcap_{i=1}^n U_{x_i}^c = \emptyset \in \mathcal{F}$, contradiction.

Finally, let $S \in \mathcal{F}$, and let x be an \mathcal{F} -limit point. Then for every open neighborhood U of x , $U \in \mathcal{F}$, so $U \cap S \in \mathcal{F}$ is non-empty. Hence $x \in \overline{S}$. \square

Given an I -indexed family of structures $\langle A_i \rangle_{i \in I}$ and a filter \mathcal{F} on I , \mathcal{F} pushes forward along P to a filter $P_*\mathcal{F}$ on $[0, 1]^{L_{\text{qf}}}$. Explicitly, for $S \subseteq [0, 1]^{L_{\text{qf}}}$, we have $S \in P_*\mathcal{F}$ if and only if $\{i \in I \mid P(A_i) \in S\} \in \mathcal{F}$. We say that $\langle A_i \rangle_{i \in I}$ \mathcal{F} -converges to $x \in [0, 1]^{L_{\text{qf}}}$ if x is a $P_*\mathcal{F}$ -limit point. By Proposition 1.1.14 (1), the $P_*\mathcal{F}$ -limit point is unique if it exists. Note that this definition agrees with the usual notion of convergence in the case that $I = \omega$ and \mathcal{F} is the cofinite filter.

However, if we take \mathcal{F} to be an ultrafilter on I , then $P_*\mathcal{F}$ is also an ultrafilter on $[0, 1]^{L_{\text{qf}}}$, and by Proposition 1.1.14 (2), every sequence of finite structures \mathcal{F} -converges. Moreover, $P_*\mathcal{F}$ is supported on $F = \{P(B) \mid B \text{ finite}\}$ for every family of finite structures $\langle A_i \rangle_{i \in I}$ and every ultrafilter \mathcal{F} on I . So by Proposition 1.1.14 (3), the limit point of every \mathcal{F} -convergent sequence is contained in \overline{F} . Hence ultrafilter convergence gives rise to the same space \overline{F} of limits of finite structures as ordinary convergence.

1.2 Measures as limit objects

In the previous section, we identified the limit of a convergent sequence of finite structures with a point in the space $[0, 1]^{L_{\text{qf}}}$, namely the sequence of limiting densities of quantifier-free formulas. But it is possible to encode these limiting densities in other, more structured, limit objects. In particular, we will describe how a convergent sequence of finite structures gives rise to a measure on an appropriate space of countably infinite structures, and we characterize the measures which arise in this way.

In the special case of graphs, it is more common to take a “graphon”, a symmetric measurable function $[0, 1]^2 \rightarrow [0, 1]$, as the limit object for a convergent sequence of graphs. Graphons were introduced by Lovasz and Szegedy, see [69] and the book [68]. For the relationship between graphons and exchangeable arrays from probability, which are similar

to our invariant measures, see Diaconis and Janson [29]. Graphons have useful generalizations to hypergraphs and finite relational languages, and Austin [9] gives a good survey. Another option in the context of a general relational language is the notion of a flag algebra in the sense of Razborov [83] (more precisely, the limit object in this context is a positive homomorphism from a flag algebra to \mathbb{R}).

Definition 1.2.1. Str_L is the space of L -structures with domain ω . For any formula $\varphi(\bar{x})$ and tuple \bar{a} from ω , we define $\llbracket \varphi(\bar{a}) \rrbracket = \{M \in \text{Str}_L \mid M \models \varphi(\bar{a})\}$. The topology on Str_L is generated by sets of the form $\llbracket R(\bar{a}) \rrbracket$ and $\llbracket \neg R(\bar{a}) \rrbracket$, where R ranges over the relation symbols in L and \bar{a} ranges over the $\text{ar}(R)$ -tuples from ω .

Note that the sets $\llbracket R(\bar{a}) \rrbracket$ and $\llbracket \neg R(\bar{a}) \rrbracket$ are complements, so both are clopen. Furthermore, a basic open set in Str_L is a finite intersection of these subbasic clopens, and hence has the form

$$\bigcap_{i=1}^n \llbracket \varphi_i(\bar{a}) \rrbracket = \llbracket \bigwedge_{i=1}^n \varphi_i(\bar{a}) \rrbracket,$$

where each $\varphi_i(\bar{x})$ is an instance of R or $\neg R$ on a tuple of variables from \bar{x} .

A point M in Str_L is uniquely determined by whether or not, for each relation symbol R in L of arity $\text{ar}(R)$ and each $\text{ar}(R)$ -tuple \bar{a} from ω , $M \models R(\bar{a})$. It follows that Str_L is homeomorphic to the Cantor space

$$\prod_{R \in L} 2^{(\omega^{\text{ar}(R)})}.$$

In particular, it is a compact Polish space with a basis of clopen sets.

Stone duality [50] tells us that every compact Hausdorff space with a basis of clopen sets is homeomorphic to the Stone space of its Boolean algebra of clopen sets. In this case, it will be useful to us to give a logical interpretation of this Boolean algebra.

Let $\mathcal{B}_{\text{qf}}^\omega$ be the Boolean algebra of quantifier-free formulas in the infinite variable context $\{x_i \mid i \in \omega\}$, modulo logical equivalence. We can view this Boolean algebra as the direct limit $\varinjlim \mathcal{B}_{\text{qf}}^n$ of the Boolean algebras of quantifier-free formulas in the finite variable contexts x_0, \dots, x_{n-1} , modulo logical equivalence, where the connecting embeddings are given by adjoining dummy variables. The Stone space of $\mathcal{B}_{\text{qf}}^\omega$ is S_{qf}^ω , the space of complete consistent quantifier-free types in the variable context $\{x_i \mid i \in \omega\}$, relative to the empty theory. By duality, $S_{\text{qf}}^\omega = \varprojlim S_{\text{qf}}^n$, where S_{qf}^n is the space of complete consistent quantifier-free types in the variables x_0, \dots, x_{n-1} .

Let $S_{\text{qf}}^{\omega, \text{nr}} = \{p \in S_{\text{qf}}^\omega \mid x_i \neq x_j \in p \text{ for all } i \neq j \text{ in } \omega\}$ be the subspace of non-redundant types. The Boolean algebra $\mathcal{B}_{\text{qf}}^{\omega, \text{nr}}$ of clopen sets in $S_{\text{qf}}^{\omega, \text{nr}}$ is isomorphic to $\mathcal{B}_{\text{qf}}^\omega$ modulo the filter \mathcal{F}_{nr} generated by $\{x_i \neq x_j \mid i \neq j \text{ in } \omega\}$.

Now given a structure $M \in \text{Str}_L$, let $\text{Diag}_{\text{qf}}(M) = \{\varphi(x_{a_1}, \dots, x_{a_n}) \mid M \models \varphi(a_1, \dots, a_n)\}$. Note that $\text{Diag}_{\text{qf}}(M)$, the quantifier-free diagram of M , is a non-redundant type in $S_{\text{qf}}^{\omega, \text{nr}}$.

Lemma 1.2.2. *The map Diag_{qf} is a homeomorphism $\text{Str}_L \cong S_{\text{qf}}^\omega$.*

Proof. The map Diag_{qf} is injective, since distinct structures in Str_L differ by some relation on some tuple from ω , which is reflected in their quantifier-free diagrams. It is also surjective, since given any type $p \in S_{\text{qf}}^\omega$, we can form the structure M_p with domain ω , such that $M_p \models R(a_1, \dots, a_n)$ if and only if $R(x_{a_1}, \dots, x_{a_n}) \in p$, and $p = \text{Diag}_{\text{qf}}(M_p)$.

To see that the topologies agree, note that the topology on S_{qf}^ω is generated by clopen sets of the form $U_{\varphi(\bar{x})} = \{p \in S_{\text{qf}}^\omega \mid \varphi(\bar{x}) \in p\}$, for $\varphi(\bar{x})$ quantifier-free. If $\bar{x} = x_{a_1}, \dots, x_{a_n}$, the preimage of this set under Diag_{qf} is $\llbracket \varphi(\bar{a}) \rrbracket$. We would like to show that these are exactly the clopen sets of Str_L .

As noted above, the sets defined by atomic formulas $\llbracket R(\bar{a}) \rrbracket$ are clopen, as are those defined by instances of equality:

$$\llbracket a = b \rrbracket = \begin{cases} \text{Str}_L & \text{if } a = b \\ \emptyset & \text{otherwise.} \end{cases}$$

Now quantifier-free formulas are built up from atomics by finite Boolean combinations, which preserve clopen sets:

$$\begin{aligned} \llbracket \neg \varphi(\bar{a}) \rrbracket &= \text{Str}_L \setminus \llbracket \varphi(\bar{a}) \rrbracket \\ \llbracket \varphi(\bar{a}) \wedge \psi(\bar{a}) \rrbracket &= \llbracket \varphi(\bar{a}) \rrbracket \cap \llbracket \psi(\bar{a}) \rrbracket \\ \llbracket \varphi(\bar{a}) \vee \psi(\bar{a}) \rrbracket &= \llbracket \varphi(\bar{a}) \rrbracket \cup \llbracket \psi(\bar{a}) \rrbracket \end{aligned}$$

Conversely, suppose C is clopen in Str_L . Since C is open, it is a union of basic open sets

$$C = \bigcup_{i \in I} \left[\bigwedge_{j=1}^n \varphi_{ij}(\bar{a}_i) \right],$$

where each $\varphi_{ij}(\bar{x})$ is an instance of R or $\neg R$. But since C is closed, it is compact, and a finite union suffices. So

$$C = \left[\bigvee_{j=1}^m \bigwedge_{i=1}^n \varphi_{ij}(\bar{a}_i) \right],$$

and hence every clopen set in Str_L has the form $\llbracket \varphi(\bar{a}) \rrbracket$ for $\varphi(\bar{x})$ quantifier-free. \square

The space Str_L comes naturally equipped with a group action, which captures the isomorphism relation on structures with domain ω . The measures we are interested in are all invariant for this action.

Definition 1.2.3. The **logic action** is the natural action of S_∞ , the permutation group of ω , on Str_L , defined by pushing forward a structure along a permutation of its domain. Explicitly, if $\sigma: \omega \rightarrow \omega$ is a bijection and $M \in \text{Str}_L$, then

$$\sigma(M) \models R(a_1, \dots, a_n) \quad \text{if and only if} \quad M \models R(\sigma^{-1}(a_1), \dots, \sigma^{-1}(a_n)).$$

Note that $\sigma(M) = N$ if and only if $\sigma: M \rightarrow N$ is an isomorphism, so the orbit of a point M under the logic action is the set of all structures in Str_L which are isomorphic to M .

Definition 1.2.4. Let μ be a Borel probability measure on Str_L .

- μ is **invariant** (for the logic action) if for every Borel set X and every $\sigma \in S_\infty$, we have $\mu(\sigma(X)) = \mu(X)$.
- μ satisfies **disjoint-independence** if $\mu(\llbracket \varphi(\bar{a}) \wedge \psi(\bar{b}) \rrbracket) = \mu(\llbracket \varphi(\bar{a}) \rrbracket)\mu(\llbracket \psi(\bar{b}) \rrbracket)$ whenever $\varphi(\bar{x})$ and $\psi(\bar{x})$ are quantifier-free formulas and \bar{a} and \bar{b} are disjoint tuples.

An **ergodic structure** is an invariant Borel probability measure on Str_L which satisfies disjoint-independence.

The words “ergodic” and “structure” here both require a bit of explanation. First, we use the word “structure” because we would like to think of an invariant measure on Str_L as a random countable structure. Indeed, we will see in the remainder of this chapter and the next that ergodic structures share certain properties of ordinary structures. Second, it follows from Theorem 1.3.2 below that an invariant measure on Str_L satisfies disjoint-independence if and only if it is ergodic in the classical sense, which we now take a moment to recall.

Definition 1.2.5. Let μ be an invariant Borel probability measure on Str_L . A Borel set X is **almost surely invariant** for μ if $\mu(X \Delta \sigma(X)) = 0$ for all $\sigma \in S_\infty$. An invariant measure μ is **ergodic** if for every almost surely invariant Borel set X , either $\mu(X) = 0$ or $\mu(X) = 1$.

Of course, every set of measure 0 or 1 is almost surely invariant. Ergodicity says that these are the only almost surely invariant sets, i.e. the group S_∞ “thoroughly mixes” the space relative to the measure.

Remark 1.2.6. Let S be a space equipped with an action of a group G and a G -invariant measure μ . A Borel set X is **invariant** if $\sigma(X) = X$ for all $\sigma \in G$, and μ is **weakly ergodic** if $\mu(X) = 0$ or 1 for every invariant Borel set X . Many sources call weakly ergodic measures ergodic; this confusion is explained by the fact that the notions are equivalent when G is countable. However, for actions of uncountable groups (such as S_∞), ergodicity as we have defined it above is the more important property. See [53, Appendix A1].

In Example 1.4.17 below, we describe an invariant measure μ on Str_L which is weakly ergodic but not ergodic and an almost surely invariant set (with respect to μ) which does not differ by a set of measure 0 from an invariant Borel set.

We have preferred to give the definition of ergodic structure in terms of the disjoint-independence property here because it is rather more concrete, because it is straightforward to verify in examples, and because it clearly holds of measures obtained as limits of convergent sequences of finite structures (as in the proof of Theorem 1.2.10). We will also use this property directly in several proofs (Theorems 1.2.18 and 1.4.8), while ergodicity is used only indirectly.

Remark 1.2.7. If $\varphi(\bar{x})$ is a formula, \bar{a} is a non-redundant tuple from ω , and $\sigma \in S_\infty$, then

$$\sigma^{-1}(\llbracket \varphi(\bar{a}) \rrbracket) = \{M \in \text{Str}_L \mid \sigma(M) \models \varphi(\bar{a})\} = \{M \in \text{Str}_L \mid M \models \varphi(\sigma^{-1}(\bar{a}))\} = \llbracket \varphi(\sigma^{-1}(\bar{a})) \rrbracket.$$

If μ is invariant, then the value $\mu(\llbracket \varphi(\bar{a}) \rrbracket)$ is independent of the choice of non-redundant tuple \bar{a} , since if $\sigma \in S_\infty$, with $\sigma(\bar{a}) = \bar{b}$, then $\mu(\llbracket \varphi(\bar{a}) \rrbracket) = \mu(\sigma(\llbracket \varphi(\bar{a}) \rrbracket)) = \mu(\llbracket \varphi(\bar{b}) \rrbracket)$. For convenience, we denote this quantity by $\mu(\varphi(\bar{x}))$. Note that under this convention, if $\varphi(\bar{x})$ implies $x_i = x_j$ for some $i \neq j$, then $\mu(\varphi(\bar{x})) = 0$.

We are now ready to explain the sense in which convergent sequences of finite structures converge to ergodic structures.

Definition 1.2.8. Let \mathcal{B} be a Boolean algebra. A finitely additive probability measure on \mathcal{B} is a function $\nu: \mathcal{B} \rightarrow [0, 1]$ such that $\nu(\perp) = 0$, $\nu(\top) = 1$, and $\nu(X \vee Y) = \nu(X) + \nu(Y) - \nu(X \wedge Y)$ for all $X, Y \in \mathcal{B}$.

The following lemma is a straightforward application of the Hahn–Kolmogorov measure extension theorem [92, Theorem 1.7.8 and Exercise 1.7.7].

Lemma 1.2.9. *Let \mathcal{B} be the Boolean algebra of clopen sets in Str_L . Any finitely additive probability measure ν on \mathcal{B} extends to a unique Borel probability measure μ on Str_L .*

Theorem 1.2.10. *Let $\langle B_n \rangle_{n \in \omega}$ be a convergent sequence of finite structures, and assume that $\lim_{n \rightarrow \infty} |B_n| = \infty$. Then there is a unique Borel probability measure μ on Str_L such that $\mu(\llbracket \varphi(\bar{a}) \rrbracket) = \lim_{n \rightarrow \infty} P(\varphi(\bar{x}); B_n)$ for all quantifier-free formulas $\varphi(\bar{x})$ and non-redundant tuples \bar{a} , and μ is an ergodic structure.*

Proof. Recall that $\mathcal{B}_{\text{qf}}^m$ is the Boolean algebra of quantifier-free formulas in the m variables x_0, \dots, x_{m-1} , modulo logical equivalence. We define a measure ν_m on $\mathcal{B}_{\text{qf}}^m$ by $\nu_m(\varphi(\bar{x})) = \lim_{n \rightarrow \infty} P(\varphi(\bar{x}); B_n)$. This is well-defined by Lemma 1.1.3 (1) and finitely additive by Lemma 1.1.3 (4) and (5).

The measures ν_m cohere to a finitely additive probability measure ν_ω on $\mathcal{B}_{\text{qf}}^\omega = \varinjlim \mathcal{B}_{\text{qf}}^m$, where we set $\nu_\omega(\varphi(\bar{x})) = \nu_m(\varphi(\bar{x}))$ for m large enough so that the variables in \bar{x} are among x_0, \dots, x_{m-1} . This is well-defined by Lemma 1.1.3 (2).

Now by Remark 1.1.5, every element of the filter \mathcal{F}_{nr} generated by $\{x_i \neq x_j \mid i \neq j \text{ in } \omega\}$ has measure 1, so ν_ω descends to a finitely additive probability measure on the Boolean algebra $\mathcal{B}_{\text{qf}}^\omega = \mathcal{B}_{\text{qf}}^\omega / \mathcal{F}_{\text{nr}}$, which can be transported to a finitely additive probability measure ν on the Boolean algebra \mathcal{B} of clopen sets in Str_L via the homeomorphism $\text{Str}_L \cong S_{\text{qf}}^\omega$ established in Lemma 1.2.2. By Lemma 1.2.9, ν extends to a Borel probability measure μ on Str_L . In other words, $\mu(\llbracket \varphi(a_1, \dots, a_m) \rrbracket) = \nu(\llbracket \varphi(a_1, \dots, a_m) \rrbracket) = \lim_{n \rightarrow \infty} P(\varphi(x_{a_1}, \dots, x_{a_m}); B_n)$, and μ is unique subject to this condition. It remains to show that μ is an ergodic structure.

For any $\varphi(\bar{a})$ and $\psi(\bar{b})$ such that $\bar{a} = a_1, \dots, a_m$ and $\bar{b} = b_1, \dots, b_l$ are disjoint, let $\bar{x} = x_{a_1}, \dots, x_{a_m}$ and $\bar{x}' = x_{b_1}, \dots, x_{b_l}$. Then, by Lemma 1.1.3 (6), we have

$$\begin{aligned} \mu(\llbracket \varphi(\bar{a}) \wedge \psi(\bar{b}) \rrbracket) &= \lim_{n \rightarrow \infty} P(\varphi(\bar{x}) \wedge \psi(\bar{x}'); B_n) \\ &= \left(\lim_{n \rightarrow \infty} P(\varphi(\bar{x}); B_n) \right) \left(\lim_{n \rightarrow \infty} P(\psi(\bar{x}'); B_n) \right) \\ &= \mu(\llbracket \varphi(\bar{a}) \rrbracket) \mu(\llbracket \psi(\bar{b}) \rrbracket). \end{aligned}$$

So μ satisfies disjoint-independence.

Finally, we show that μ is invariant. Let $\sigma \in S_\infty$. By Remark 1.2.7 and Lemma 1.1.3 (3),

$$\begin{aligned} \mu(\sigma^{-1}(\llbracket \varphi(\bar{a}) \rrbracket)) &= \mu(\llbracket \varphi(\sigma^{-1}(\bar{a})) \rrbracket) \\ &= \lim_{n \rightarrow \infty} P(\varphi(x_{\sigma^{-1}(a_1)}, \dots, x_{\sigma^{-1}(a_m)}); B_n) \\ &= \lim_{n \rightarrow \infty} P(\varphi(x_{a_1}, \dots, x_{a_m}); B_n) \\ &= \mu(\llbracket \varphi(\bar{a}) \rrbracket). \end{aligned}$$

Now μ pushes forward under σ to a Borel probability measure $\sigma_*\mu$, defined by $\sigma_*\mu(X) = \mu(\sigma^{-1}(X))$. But the computation above shows that $\sigma_*\mu(\llbracket \varphi(\bar{a}) \rrbracket) = \mu(\llbracket \varphi(\bar{a}) \rrbracket)$ for any clopen set $\llbracket \varphi(\bar{a}) \rrbracket$ in \mathcal{B} . Hence μ and $\sigma_*\mu$ agree on \mathcal{B} , so by the uniqueness in Lemma 1.2.9, $\mu = \sigma_*\mu$, i.e. $\mu(X) = \mu(\sigma^{-1}(X))$ for all Borel X , and μ is invariant. \square

Definition 1.2.11. Given a convergent sequence $\langle B_n \rangle_{n \in \omega}$ and μ as in Theorem 1.2.10, we say $\langle B_n \rangle_{n \in \omega}$ **converges** to μ .

In fact, the converse to Theorem 1.2.10 is true: every ergodic structure is the limit of a sequence of finite structures. We show this in Theorem 1.2.18 by a probabilistic argument. Intuitively, an ergodic structure μ tells us how to pick a structure with domain ω at random, so it also tells us how to pick a finite structure of size n at random (by taking the induced substructure on $[n]$). If we sample a sequence of finite structures of increasing size from μ in this way, then this sequence converges to μ almost surely.

Versions of this theorem have been proven for the various other types of limit objects discussed at the beginning of this section. See [69, Theorem 2.2] for the graphon context and [83, Theorem 3.3] for the flag algebra context.

We need a few basic facts from probability theory.

Definition 1.2.12. A **probability space** is a set Ω equipped with a σ -algebra of measurable sets and a probability measure μ . A **random variable** X is a measurable function $\Omega \rightarrow \mathbb{R}$. The **expected value** of X is $E(X) = \int_\Omega X d\mu$. The **variance** of X is $\text{Var}(X) = E((X - E(X))^2)$. We write $\text{Pr}(\mathcal{A})$ for the probability that the event \mathcal{A} occurs. That is, $\text{Pr}(\mathcal{A}) = E(1_{\mathcal{A}}) = \mu(\llbracket \mathcal{A} \rrbracket)$, where $1_{\mathcal{A}}$ is the indicator function of \mathcal{A} and $\llbracket \mathcal{A} \rrbracket$ is the subset of Ω on which \mathcal{A} holds. To say an event occurs **almost surely** is to say it has probability 1.

Lemma 1.2.13. *If X is a random variable, then $\text{Var}(X) = E(X^2) - E(X)^2$.*

Proof. Directly from the definition,

$$\begin{aligned} \text{Var}(X) &= E(X^2 - 2E(X)X + E(X)^2) \\ &= E(X^2) - 2E(X)E(X) + E(X)^2 \\ &= E(X^2) - E(X)^2. \end{aligned} \quad \square$$

Lemma 1.2.14 (Chebyshev’s Inequality, [30] (1.3.4)). *If X is a random variable and $b > 0$, then*

$$\Pr(|X - E(X)| \geq b) \leq \frac{\text{Var}(X)}{b^2}.$$

Lemma 1.2.15 (First Borel–Cantelli Lemma, [30] (1.6.1)). *Let $\{\mathcal{A}_n \mid n \in \omega\}$ be a sequence of events such that*

$$\sum_{n \in \omega} \Pr(\mathcal{A}_n) < \infty.$$

Then the probability that infinitely many of the \mathcal{A}_n occur is 0.

We only need the first Borel–Cantelli Lemma now, but the second will be useful later.

Lemma 1.2.16 (Second Borel–Cantelli Lemma, [30] (1.6.6)). *Let $\{\mathcal{A}_n \mid n \in \omega\}$ be a sequence of independent events such that*

$$\sum_{n \in \omega} \Pr(\mathcal{A}_n) = \infty.$$

Then the probability that infinitely many of the \mathcal{A}_n occur is 1.

We denote by $\text{Str}_{L,n}$ the space of L -structures with domain $[n]$. Exactly as in Str_L , the topology is generated by sets of the form $\llbracket R(\bar{a}) \rrbracket$ and $\llbracket \neg R(\bar{a}) \rrbracket$, where R ranges over the relation symbols in L and \bar{a} ranges over the $\text{ar}(R)$ -tuples from $[n]$.

Remark 1.2.17. For every n , there is a continuous map $\upharpoonright_n: \text{Str}_L \rightarrow \text{Str}_{L,n}$ sending a structure M with domain ω to the induced substructure on the subset $[n]$. Any ergodic structure μ pushes forward along this map to a probability measure μ_n on $\text{Str}_{L,n}$, defined by $\mu_n(X) = \mu(\upharpoonright_n^{-1}(X))$.

Theorem 1.2.18. *Let μ be an ergodic structure. Then there is a sequence of finite structures $\langle B_n \rangle_{n \in \omega}$ which converges to μ .*

Proof. Recall that we denote by $[n]^k$ the set of non-redundant k -tuples from $[n]$. The size of $[n]^k$ is given by the “falling factorial”: $|[n]^k| = n^{\underline{k}} = \prod_{i=0}^{k-1} (n - i)$.

Fix an explicitly non-redundant formula $\varphi(\bar{x})$ with k free variables, and for any $n \geq 1$, consider the random variable $P(\varphi(\bar{x}); A)$ on $\text{Str}_{L,n}$, equipped with the measure μ_n defined in Remark 1.2.17. Since $\varphi(\bar{x})$ is explicitly non-redundant, we only need to count non-redundant

tuples in $|\varphi(A)|$, and by the convention in Remark 1.2.7, $\mu_n(\llbracket\varphi(\bar{a})\rrbracket) = \mu(\varphi(\bar{x}))$ whenever \bar{a} is non-redundant. So,

$$\begin{aligned} E(P(\varphi(\bar{x}); A)) &= \int_{A \in \text{Str}_{L,n}} \frac{|\varphi(A)|}{n^k} d\mu_n = \frac{1}{n^k} \sum_{\bar{a} \in [n]^k} \int_{A \in \text{Str}_{L,n}} 1_{\llbracket\varphi(\bar{a})\rrbracket} d\mu_n \\ &= \frac{1}{n^k} \sum_{\bar{a} \in [n]^k} \mu_n(\llbracket\varphi(\bar{a})\rrbracket) = \frac{n^k}{n^k} \mu(\varphi(\bar{x})). \end{aligned}$$

In order to bound the variance of $P(\varphi(\bar{x}); A)$, we bound the expected value of its square. If \bar{y} is a tuple of variables disjoint from \bar{x} , then $P(\varphi(\bar{x}); A)^2 = P(\varphi(\bar{x}) \wedge \varphi(\bar{y}); A)$ by disjoint-independence. Note that the formula $\varphi(\bar{x}) \wedge \varphi(\bar{y})$ is not explicitly non-redundant. For disjoint non-redundant tuples \bar{a} and \bar{b} , $\mu_n(\llbracket\varphi(\bar{a}) \wedge \varphi(\bar{b})\rrbracket) = \mu(\varphi(\bar{x}) \wedge \varphi(\bar{y})) = \mu(\varphi(\bar{x}))^2$ as before, but when \bar{a} and \bar{b} are not disjoint, we bound $\mu_n(\llbracket\varphi(\bar{a}) \wedge \varphi(\bar{b})\rrbracket)$ above by 1.

$$\begin{aligned} E(P(\varphi(\bar{x}); A)^2) &= E(P(\varphi(\bar{x}) \wedge \varphi(\bar{y}); A)) \\ &= \frac{1}{n^{2k}} \sum_{\bar{a}\bar{b} \in ([n]^k)^2} \int_{A \in \text{Str}_{L,n}} 1_{\llbracket\varphi(\bar{a}) \wedge \varphi(\bar{b})\rrbracket} d\mu_n \\ &= \frac{1}{n^{2k}} \left(\sum_{\bar{a}\bar{b} \in [n]^{2k}} \mu_n(\llbracket\varphi(\bar{a}) \wedge \varphi(\bar{b})\rrbracket) + \sum_{\bar{a}\bar{b} \in ([n]^k)^2 \setminus [n]^{2k}} \mu_n(\llbracket\varphi(\bar{a}) \wedge \varphi(\bar{b})\rrbracket) \right) \\ &\leq \frac{1}{n^{2k}} (n^{2k} \mu(\varphi(\bar{x}))^2 + ((n^k)^2 - n^{2k})). \end{aligned}$$

By Lemma 1.2.13,

$$\begin{aligned} \text{Var}(P(\varphi(\bar{x}); A)) &= E(P(\varphi(\bar{x}); A)^2) - E(P(\varphi(\bar{x}); A))^2 \\ &\leq \frac{1}{n^{2k}} (n^{2k} \mu(\varphi(\bar{x}))^2 + ((n^k)^2 - n^{2k})) - \frac{(n^k)^2}{n^{2k}} \mu(\varphi(\bar{x}))^2 \\ &\leq \left(\frac{(n^k)^2 - n^{2k}}{n^{2k}} \right) (1 - \mu(\varphi(\bar{x}))^2) \\ &\leq \frac{(n^k)^2 - n^{2k}}{n^{2k}}. \end{aligned}$$

The numerator of this fraction is a polynomial in n of degree at most $2k - 1$, since the leading terms of $(n^k)^2$ and n^{2k} cancel. Thus there is a constant C_k , depending only on k , such that $\text{Var}(P(\varphi(\bar{x}); A)) \leq \frac{C_k}{n}$.

Consider the product space $\widehat{S} = \prod_{n \in \omega} \text{Str}_{L,n^2}$, and let $\widehat{\mu}$ be the product measure $\prod_{n \in \omega} \mu_{n^2}$ on this space. Intuitively, sampling a point $\langle B_n \rangle_{n \in \omega}$ of \widehat{S} from $\widehat{\mu}$ corresponds to sampling countably many points $\langle M_n \rangle_{n \in \omega}$ of Str_L independently from μ and taking the sequence where B_n is the induced substructure of M_n on $[n^2]$ for each n . We take $|B_n| = n^2$ so that the variances computed above approach 0 rapidly enough.

We will show that the sequence $\langle B_n \rangle_{n \in \omega}$ almost surely converges to μ (i.e. with probability 1 according to $\widehat{\mu}$). In particular, this will establish the existence of some sequence of finite structures which converges to μ . By Lemma 1.1.7, we only need to show convergence for explicitly non-redundant quantifier-free formulas.

Fix an explicitly non-redundant formula $\varphi(\bar{x})$, and fix $\varepsilon > 0$. Let X_n^φ be the random variable $P(\varphi(\bar{x}); B_n)$ on Str_{L, n^2} whose expected value and variance we computed above, and let $\mathcal{A}_n^{\varphi, \varepsilon}$ be the event $|X_n^\varphi - \mu(\varphi(\bar{x}))| \geq \varepsilon$. We would like to show that almost surely, at most finitely many of the events $\mathcal{A}_n^{\varphi, \varepsilon}$ occur. Since $\lim_{n \rightarrow \infty} \frac{(n^2)^k}{(n^2)^k} = 1$, there exists N such that for all $n > N$, $|E(X_n^\varphi) - \mu(\varphi(\bar{x}))| < \varepsilon/2$, and hence by Chebyshev's Inequality,

$$\begin{aligned} \Pr(\mathcal{A}_n^{\varphi, \varepsilon}) &\leq \Pr(|X_n^\varphi - E(X_n^\varphi)| \geq \varepsilon/2) \\ &\leq \frac{\text{Var}(X_n^\varphi)}{(\varepsilon/2)^2} \\ &\leq \frac{4C_k}{\varepsilon^2 n^2}. \end{aligned}$$

Thus we have $\sum_{n \in \omega} \Pr(\mathcal{A}_n^{\varphi, \varepsilon}) < \infty$, so by the first Borel–Cantelli Lemma (Lemma 1.2.15), the probability that infinitely many of the events $\mathcal{A}_n^{\varphi, \varepsilon}$ occur is 0. Since we can take a countable sequence of values of ε decreasing to 0, $\lim_{n \rightarrow \infty} P(\varphi(\bar{x}); B_n) = \mu(\varphi(\bar{x}))$ almost surely by countable additivity. And there are only countably many explicitly non-redundant formulas to check, so by countable additivity again, the sequence $\langle B_n \rangle_{n \in \omega}$ almost surely converges to μ . \square

Remark 1.2.19. In the proof of Theorem 1.2.18, we did not use our assumption that the structures B_n were sampled independently from the measures μ_n . Instead of pushing forward μ individually to the spaces Str_{L, n^2} and then taking the product measure, we can push forward μ forward directly to a measure $\widehat{\mu}$ on \widehat{S} along the natural map $\text{Str}_L \rightarrow \prod_{n \in \omega} \text{Str}_{L, n^2}$, and the proof goes through. Intuitively, this shows that if we sample a countable structure M from our ergodic structure μ and set B_n to be the induced substructure of M on the subset $[n^2]$, then the sequence $\langle B_n \rangle_{n \in \omega}$ almost surely converges to μ .

Also, it was not important that the domains of the structures in our sequence were the sets $[n^2]$. We can replace these sets with any sequence of subsets of ω whose sizes grow fast enough, and the induced substructures will still converge to μ almost surely.

1.3 The Aldous–Hoover–Kallenberg theorem

In this section, we state a version of the Aldous–Hoover–Kallenberg theorem. This theorem, which is a generalization of de Finetti's theorem [30, Theorem 4.6.6] to exchangeable arrays of random variables, was discovered independently by Aldous [8] and Hoover [44], and further developed by Kallenberg [51] and others.

The original settings for the AHK theorem were abstract probabilistic frameworks, which did not involve the space Str_L , the logic action, or spaces of quantifier-free types. So the

version of the theorem presented here is a translation of the original theorem to a form which is convenient for our setting. See [1, Section 2.5] for a discussion of this translation. The survey by Austin [9] also provides details on a slightly different translation of the AHK theorem, and applications to random structures and asymptotics of finite structures.

As above, we denote by S_{qf}^n the Stone space of non-redundant quantifier-free types in the variables x_0, \dots, x_{n-1} . Its basic clopen sets have the form $U_{\varphi(\bar{x})} = \{p \in S_{\text{qf}}^n \mid \varphi(\bar{x}) \in p\}$, where $\varphi(\bar{x})$ is a quantifier-free formula. This space admits an action of the symmetric group S_n (the group of permutations of $[n]$), by $\sigma(p(x_0, \dots, x_{n-1})) = p(x_{\sigma(0)}, \dots, x_{\sigma(n-1)})$ for $\sigma \in S_n$.

Let $(\xi_A)_{A \in \mathcal{P}_{\text{fin}}(\omega)}$ be a collection of independent random variables, each uniformly distributed on $[0, 1]$, indexed by the finite subsets of ω . We think of ξ_A as a source of randomness sitting on the subset A , which we will use to build a random L -structure with domain ω . If $\bar{a} \in \omega^n$, the injective function $i: [n] \rightarrow \omega$ enumerating \bar{a} associates to each $X \in \mathcal{P}([n])$ a subset $i(X) \subseteq \|\bar{a}\|$. We denote by $\widehat{\xi}_{\bar{a}}$ the family of random variables $(\xi_{i(X)})_{X \in \mathcal{P}([n])}$.

An **AHK system** is a collection of measurable functions

$$(f_n: [0, 1]^{\mathcal{P}([n])} \rightarrow S_{\text{qf}}^n)_{n \in \omega}.$$

satisfying the coherence conditions:

- If $\sigma \in S_n$, then $f_n((\xi_{\sigma(X)})_{X \subseteq [n]}) = \sigma(f_n((\xi_X)_{X \subseteq [n]}))$ almost surely.
- If $m < n$, then $f_m((\xi_X)_{X \subseteq [m]}) \subseteq f_n((\xi_Y)_{Y \subseteq [n]})$ almost surely.

That is, f_n takes as input a collection of values in $[0, 1]$, indexed by $\mathcal{P}([n])$, and produces a non-redundant quantifier-free n -type. Using our random variables ξ_A , we have a natural notion of **sampling** from an AHK system to obtain a non-redundant quantifier-free type $r_{\bar{a}} = f_n(\widehat{\xi}_{\bar{a}})$ for every finite tuple \bar{a} from ω . Note that the order in which $\|\bar{a}\|$ is enumerated in the tuple \bar{a} is significant, since f_n is, in general, not symmetric in its arguments. The coherence conditions ensure that the quantifier-free types obtained from the functions f_n cohere appropriately to define a structure with domain ω .

Formally, an AHK system $(f_n)_{n \in \omega}$ gives rise to a finitely additive probability measure ν on the Boolean algebra \mathcal{B} of clopen sets in Str_L , defined by

$$\nu(\llbracket \varphi(\bar{a}) \rrbracket) = \lambda^{\mathcal{P}([n])}(f_n^{-1}(U_{\varphi(\bar{x})})),$$

where $\lambda^{\mathcal{P}([n])}$ is the uniform product measure on $[0, 1]^{\mathcal{P}([n])}$, $\varphi(\bar{x})$ is a quantifier-free formula, and \bar{a} is a non-redundant tuple from ω . This is the probability that $\varphi(\bar{x}) \in r_{\bar{a}}$. The coherence conditions imply that this is well-defined: the first ensures that the order in which we list the variables in $\varphi(\bar{x})$ is irrelevant, and the second ensures that the measure is independent of the variable context \bar{x} .

Since the definition of $\nu(\llbracket \varphi(\bar{a}) \rrbracket)$ does not depend on the choice of non-redundant tuple \bar{a} , ν is manifestly invariant for the logic action, By Lemma 1.2.9, μ^* induces a unique invariant Borel probability measure μ on Str_L . In this case, we say that $(f_n)_{n \in \omega}$ is an **AHK representation** of μ .

Theorem 1.3.1 (Aldous–Hoover–Kallenberg, [53, Theorem 7.22]). *Every invariant probability measure μ on Str_L has an AHK representation.*

The AHK representation produced by Theorem 1.3.1 is not unique, but it is unique up to certain appropriately measure-preserving transformations. See [53, Theorem 7.28] for a precise statement.

The key fact to observe about AHK systems is that if \bar{a} and \bar{b} are tuples from ω whose intersection $\|\bar{a}\| \cap \|\bar{b}\|$ is enumerated by the tuple \bar{c} , then the random quantifier-free types $r_{\bar{a}}$ and $r_{\bar{b}}$ are conditionally independent over $\widehat{\xi}_{\bar{c}}$. And if \bar{a} and \bar{b} are disjoint, then $\widehat{\xi}_{\bar{c}} = \xi_{\emptyset}$.

If μ satisfies disjoint-independence (i.e. if μ is an ergodic structure), then $r_{\bar{a}}$ and $r_{\bar{b}}$ should be fully independent whenever \bar{a} and \bar{b} are disjoint. In other words, the random variable ξ_{\emptyset} should be irrelevant in the AHK representation of μ . It is a less obvious fact that this condition is equivalent to ergodicity of μ . In the probabilistic literature (e.g. [53]) the word “dissociated” is used for the analog of disjoint-independence.

Theorem 1.3.2 ([53, Lemma 7.35]). *Let μ be an invariant probability measure on Str_L . The following are equivalent:*

- (1) μ is ergodic.
- (2) μ satisfies disjoint-independence.
- (3) μ has an AHK representation in which the functions f_n do not depend on the argument indexed by \emptyset .

We’ll now take a moment to explain the connection between AHK representations and graphons.

Example 1.3.3. Let L be the language consisting of a single binary relation symbol R , and suppose μ is an ergodic structure which gives measure 1 to the set of graphs with edge relation R in Str_L . Then much of the information contained in an AHK representation of μ is unnecessary. For $n = 0$ and $n = 1$, there is only one quantifier-free n -type realized in any graph, so the functions f_0 and f_1 are almost surely constant. The quantifier-free type of a non-redundant tuple \bar{a} in a graph is already determined by the quantifier-free 2-types of the pairs from \bar{a} , so the functions f_n for $n \geq 3$ are determined (almost surely) by the function f_2 . Finally, by Theorem 1.3.2, f_2 can be chosen so it does not depend on the input ξ_{\emptyset} .

Hence an AHK representation of μ is essentially determined by a function f_2 in three variables $(\xi_{\{0\}}, \xi_{\{1\}}, \xi_{\{0,1\}})$ with two possible outputs (the non-redundant quantifier-free 2-types determined by xRy and $\neg xRy$). Sampling from this AHK representation amounts to picking a random value $\xi_{\{i\}}$ from $[0, 1]$ uniformly and independently for each $i \in \omega$, then picking another random value $\xi_{\{i,j\}}$ from $[0, 1]$ uniformly and independently for each pair $\{i, j\}$, and setting $\text{qftp}(i, j) = f_2(\xi_{\{i\}}, \xi_{\{j\}}, \xi_{\{i,j\}})$. For fixed values of $\xi_{\{i\}}$ and $\xi_{\{j\}}$, it is irrelevant which sets of values of $\xi_{\{i,j\}}$ produce which of the possible 2-types; what matters is the measure of these sets, i.e. the probabilities out the two outcomes. So we can simplify

further, replacing f_2 with a measurable function $W: [0, 1]^2 \rightarrow [0, 1]$, such that $W(a, b)$ is the probability that $xRy \in f_2(a, b, \xi)$, for a random variable ξ uniformly distributed on $[0, 1]$. Note that this function W is symmetric in its two inputs. A symmetric measurable function $[0, 1]^2 \rightarrow [0, 1]$ is called a graphon [68]. Sampling from the AHK representation of μ agrees with the natural notion of sampling from the graphon: build a graph on ω by picking a value x_i from $[0, 1]$ uniformly and independently for each i , then connect i and j by an edge with probability $W(x_i, x_j)$.

Following this model, it is straightforward to describe analogs of graphons for other classes of structures in languages of bounded arity, by paring away unnecessary details from the AHK representation. For example, in the case of simple directed graphs (i.e. structures in the language consisting of a single binary relation R , which satisfy $\forall x \neg xRx$), there is a unique quantifier-free 0-type and a unique quantifier-free 1-type, and there are four non-redundant quantifier-free 2-types, so an AHK representation can be pared down to a measurable function $W: [0, 1]^2 \rightarrow \{(a, b, c, d) \in [0, 1]^4 \mid a + b + c + d = 1\}$, where the four entries of $W(a, b)$ are the probabilities assigned to the non-redundant quantifier-free 2-types.

1.4 The theory of an ergodic structure

The infinitary logic $L_{\omega_1, \omega}$ is the extension of first-order logic obtained by allowing, as new formula-building operations, the conjunction (\bigwedge) or disjunction (\bigvee) of any countable ($< \omega_1$) family of formulas in the same finite ($< \omega$) variable context. We ensure that all our variables come from a fixed countable supply. For a reference on $L_{\omega_1, \omega}$, see [57].

As before, given a formula $\varphi(\bar{x}) \in L_{\omega_1, \omega}$ and a tuple \bar{a} from ω , we define

$$\llbracket \varphi(\bar{a}) \rrbracket = \{M \in \text{Str}_L \mid M \models \varphi(\bar{a})\}.$$

This is always a Borel set in Str_L . Indeed, negation corresponds to complement, countable conjunction and disjunction correspond to countable intersection and union, and quantifiers over the countable domain also correspond to countable Boolean operations:

$$\begin{aligned} \llbracket \forall x \varphi(\bar{a}, x) \rrbracket &= \bigcap_{b \in \omega} \llbracket \varphi(\bar{a}, b) \rrbracket \\ \llbracket \exists x \varphi(\bar{a}, x) \rrbracket &= \bigcup_{b \in \omega} \llbracket \varphi(\bar{a}, b) \rrbracket. \end{aligned}$$

If φ is a sentence, then the Borel set $\llbracket \varphi \rrbracket$ is invariant for the logic action, since its set of models is closed under isomorphism. A key fact about $L_{\omega_1, \omega}$ is Lopez-Escobar's Theorem, which states that the converse holds.

Theorem 1.4.1 (Lopez-Escobar, [55, Theorem I.16.8]). *If $X \subseteq \text{Str}_L$ is an invariant Borel set, then $X = \llbracket \varphi_X \rrbracket$ for some sentence φ_X of $L_{\omega_1, \omega}$.*

Scott's Theorem, a predecessor to Lopez-Escobar's Theorem, states that $L_{\omega_1, \omega}$ is expressive enough to describe countable structures uniquely up to isomorphism.

Theorem 1.4.2 (Scott, [75, Theorem 2.4.15]). *For any countable structure M , there is a sentence φ_M of $L_{\omega_1, \omega}$, the **Scott sentence of M** , such that for all countable structures N , we have $N \models \varphi_M$ if and only if $N \cong M$.*

Since the set of structures in Str_L which are isomorphic to M is the orbit of M under the logic action, Scott's Theorem can be viewed as a consequence of Lopez-Escobar's theorem together with the general fact [55, Theorem I.15.14] that for any Borel action of a Polish group on a standard Borel space, every orbit is Borel. Alternatively, we can conclude from Scott's Theorem (which has a purely model-theoretic proof) that every isomorphism class of countable structures is Borel in Str_L without using any descriptive set theory.

Let us now return to our measures. If μ is an invariant probability measure on Str_L , we adopt the convention of Remark 1.2.7 for formulas of $L_{\omega_1, \omega}$, writing $\mu(\varphi(\bar{x}))$ for $\mu(\llbracket \varphi(\bar{a}) \rrbracket)$, since this quantity is independent of the choice of non-redundant tuple \bar{a} .

Definition 1.4.3. Let μ be an invariant probability measure on Str_L . If φ is a sentence of $L_{\omega_1, \omega}$, we say μ **almost surely satisfies** φ if $\mu(\varphi) = 1$. We write $\mu \models_{\text{a.s.}} \varphi$, and we define $\text{Th}(\mu) = \{\varphi \in L_{\omega_1, \omega} \mid \mu \models_{\text{a.s.}} \varphi\}$. Similarly, if Σ is a set of sentences of $L_{\omega_1, \omega}$, we write $\mu \models_{\text{a.s.}} \Sigma$ if $\mu \models_{\text{a.s.}} \varphi$ for all $\varphi \in \Sigma$.

Proposition 1.4.4. *If μ is an ergodic structure, then $\text{Th}(\mu)$ is a complete and countably consistent theory of $L_{\omega_1, \omega}$. That is, for every sentence φ , $\varphi \in \text{Th}(\mu)$ or $\neg\varphi \in \text{Th}(\mu)$, and every countable subset $\Sigma \subseteq \text{Th}(\mu)$ has a model.*

Proof. For any sentence φ , the set $\llbracket \varphi \rrbracket$ is an invariant Borel set. In particular, it is almost surely invariant, so by ergodicity (Definition 1.2.5), $\mu(\varphi) = 0$ or 1 , and hence $\varphi \in \text{Th}(\mu)$ or $\neg\varphi \in \text{Th}(\mu)$. Now let Σ be a countable subset of $\text{Th}(\mu)$. Since a countable intersection of measure 1 sets has measure 1, $\mu(\bigwedge_{\varphi \in \Sigma} \varphi) = 1$. In particular, $\llbracket \bigwedge_{\varphi \in \Sigma} \varphi \rrbracket$ is non-empty. \square

Remark 1.4.5. Lopez-Escobar's Theorem (Theorem 1.4.1) tells us that if ν is an invariant probability measure on Str_L , then $\text{Th}(\nu)$ is complete if and only if $\mu(X) = 0$ or 1 for every invariant Borel set X if and only if ν is weakly ergodic (see Remark 1.2.6). But there are weakly ergodic invariant probability measures on Str_L which are not ergodic, i.e. which do not arise as limits of finite structures. See Example 1.4.17 below for a concrete example of this.

We have seen that completeness of $\text{Th}(\mu)$ follows immediately from ergodicity of μ . However, since we view disjoint-independence as a more basic feature of our measures than ergodicity, it is worth noting that we can give a proof of Proposition 1.4.4 directly from disjoint-independence. A similar argument (but in a less general context, and dealing only with the first-order case) appears in a paper of Gaifman's from 1964 [37], where he gives what is essentially a probabilistic and logical analysis of the random L -structure (which example he attributes to Rabin and Scott), predating the work of Glebskii, Kogan, Liogon'kii, and Talanov [38] and Fagin [35].

The idea is to show that disjoint-independence holds not just for quantifier-free formulas, but for all formulas of $\mathcal{L}_{\omega_1, \omega}$, and in particular for sentences. Then the result follows by the standard trick for proving zero-one laws in probability: an event which is independent from itself has measure 0 or 1.

We will need to do an induction on quantifier-rank of formulas. We will only use the notion of quantifier-rank in the proof of Lemma 1.4.7, and we define it in a way that will be most convenient for that proof. For the purposes of Definition 1.4.6 and Lemma 1.4.7, we view \forall and $\exists y$ and shorthands, defined by negation from \bigwedge and $\forall y$.

Definition 1.4.6. Let $\varphi(\bar{x})$ be a formula of $L_{\omega_1, \omega}$. The quantifier rank of $\varphi(\bar{x})$, denoted $\text{QR}(\varphi(\bar{x}))$, is an ordinal, defined by induction on the complexity of $\varphi(\bar{x})$:

- If $\varphi(\bar{x})$ is atomic, $\text{QR}(\varphi(\bar{x})) = 0$.
- If $\varphi(\bar{x}) = \neg\psi(\bar{x})$, $\text{QR}(\varphi(\bar{x})) = \text{QR}(\psi(\bar{x}))$.
- If $\varphi(\bar{x}) = \bigwedge_{i \in \omega} \psi_i(\bar{x})$, $\text{QR}(\varphi(\bar{x})) = \sup_{i \in \omega} (\text{QR}(\psi_i(\bar{x})))$.
- If $\varphi(\bar{x}) = \forall y \psi(\bar{x}, y)$, $\text{QR}(\varphi(\bar{x})) = \text{QR}(\psi(\bar{x}, y)) + 1$.

Note that the quantifier rank of any formula which begins with a quantifier must be a successor ordinal.

Lemma 1.4.7. Let μ be an ergodic structure, let $\varphi(\bar{x})$ and $\psi(\bar{y})$ be formulas of $\mathcal{L}_{\omega_1, \omega}$, and let \bar{a} and \bar{b} be disjoint tuples from ω . Then $\mu(\llbracket \varphi(\bar{a}) \wedge \psi(\bar{b}) \rrbracket) = \mu(\llbracket \varphi(\bar{a}) \rrbracket) \mu(\llbracket \psi(\bar{b}) \rrbracket)$.

Proof. For any ordinal α and any tuple \bar{a} from ω , let $\mathcal{QR}(\alpha, \bar{a}) = \{\llbracket \varphi(\bar{a}) \rrbracket \mid \text{QR}(\varphi(\bar{x})) \leq \alpha\}$. Note that $\mathcal{QR}(\alpha, \bar{a})$ is a σ -algebra (it is closed under complement and countable intersection). We will show by induction that for any pair of ordinals (α, β) and for any disjoint tuples \bar{a} and \bar{b} , the σ -algebras $\mathcal{QR}(\alpha, \bar{a})$ and $\mathcal{QR}(\beta, \bar{b})$ are μ -independent, i.e. the conclusion of the theorem holds for any $\llbracket \varphi(\bar{a}) \rrbracket \in \mathcal{QR}(\alpha, \bar{a})$ and $\llbracket \psi(\bar{b}) \rrbracket \in \mathcal{QR}(\beta, \bar{b})$. We use repeatedly the fact that if \mathcal{S} and \mathcal{T} are families of sets which are closed under finite intersection, and if \mathcal{S} and \mathcal{T} are independent, then the σ -algebras generated by \mathcal{S} and \mathcal{T} are independent [30, Theorem 1.4.2].

For the base case, $\alpha = \beta = 0$, $\mathcal{QR}(\alpha, \bar{a})$ and $\mathcal{QR}(\beta, \bar{b})$ are the σ -algebras generated by the first-order quantifier-free formulas, so the assertion follows immediately from the fact that μ satisfies disjoint-independence.

Now assume without loss of generality that $\alpha > 0$. If α is a limit ordinal, then $\mathcal{QR}(\alpha, \bar{a})$ is the σ -algebra generated by $\bigcup_{\gamma < \alpha} \mathcal{QR}(\gamma, \bar{a})$. Since for all $\gamma < \alpha$, $\mathcal{QR}(\gamma, \bar{a})$ is independent from $\mathcal{QR}(\beta, \bar{b})$, the σ -algebra generated by the union is also independent from $\mathcal{QR}(\beta, \bar{b})$.

If $\alpha = \gamma + 1$, then $\mathcal{QR}(\alpha, \bar{a})$ is generated by $\mathcal{QR}(\gamma, \bar{a}) \cup \{\llbracket \forall z \hat{\varphi}(\bar{a}, z) \rrbracket \mid \text{QR}(\hat{\varphi}(\bar{x}, z)) \leq \gamma\}$ (we chose the universal quantifier here to ensure that this generating family is closed under finite intersection). It suffices to show that for any such $\hat{\varphi}(\bar{x}, z)$ and any $\llbracket \psi(\bar{b}) \rrbracket \in \mathcal{QR}(\beta, \bar{b})$, $\llbracket \forall z \hat{\varphi}(\bar{a}, z) \rrbracket$ and $\llbracket \psi(\bar{b}) \rrbracket$ are independent.

By countable additivity of μ , we have

$$\begin{aligned} \mu(\llbracket \forall z \widehat{\varphi}(\bar{a}, z) \rrbracket) &= \mu\left(\bigcap_{c \in \omega} \llbracket \widehat{\varphi}(\bar{a}, c) \rrbracket\right) \\ &= \inf_{J \subseteq_{\text{fin}} \omega} \mu\left(\bigcap_{c \in J} \llbracket \widehat{\varphi}(\bar{a}, c) \rrbracket\right) \\ &= \inf_{J \subseteq_{\text{fin}} \omega} \mu\left(\llbracket \bigwedge_{c \in J} \widehat{\varphi}(\bar{a}, c) \rrbracket\right). \end{aligned}$$

For each finite set $J \subseteq \omega$, let \bar{c} be a tuple enumerating J . Then there is a permutation $\sigma_J \in S_\infty$ such that $\sigma_J(\bar{a}) = \bar{a}$, and $\sigma_J(\bar{c})$ is disjoint from \bar{b} . By invariance,

$$\mu(\llbracket \forall z \widehat{\varphi}(\bar{a}, z) \rrbracket) = \inf_{J \subseteq_{\text{fin}} \omega} \mu\left(\llbracket \bigwedge_{c \in J} \widehat{\varphi}(\sigma_J(\bar{a}), \sigma_J(c)) \rrbracket\right) = \inf_{J \subseteq_{\text{fin}} \omega} \mu\left(\llbracket \bigwedge_{c \in J} \widehat{\varphi}(\bar{a}, \sigma_J(c)) \rrbracket\right). \quad (\dagger)$$

This means that for every $\varepsilon > 0$, there exists a $J \subseteq_{\text{fin}} \omega$ such that the difference between the measures of $\llbracket \forall z \widehat{\varphi}(\bar{a}, z) \rrbracket$ and $\llbracket \bigwedge_{c \in J} \widehat{\varphi}(\bar{a}, \sigma_J(c)) \rrbracket$ is less than ε . Intersecting with $\llbracket \psi(\bar{b}) \rrbracket$, we also have that the difference between the measures of $\llbracket \forall z \widehat{\varphi}(\bar{a}, z) \rrbracket \cap \llbracket \psi(\bar{b}) \rrbracket$ and $\llbracket \bigwedge_{c \in J} \widehat{\varphi}(\bar{a}, \sigma_J(c)) \rrbracket \cap \llbracket \psi(\bar{b}) \rrbracket$ is less than ε . Hence,

$$\begin{aligned} \mu(\llbracket \forall z \widehat{\varphi}(\bar{a}, z) \wedge \psi(\bar{b}) \rrbracket) &= \inf_{J \subseteq_{\text{fin}} \omega} \mu\left(\llbracket \left(\bigwedge_{c \in J} \widehat{\varphi}(\bar{a}, \sigma_J(c))\right) \wedge \psi(\bar{b}) \rrbracket\right) \\ &= \inf_{J \subseteq_{\text{fin}} \omega} \mu\left(\llbracket \bigwedge_{c \in J} \widehat{\varphi}(\bar{a}, \sigma_J(c)) \rrbracket\right) \mu(\llbracket \psi(\bar{b}) \rrbracket) \\ &= \mu(\llbracket \forall z \widehat{\varphi}(\bar{a}, z) \rrbracket) \mu(\llbracket \psi(\bar{b}) \rrbracket). \end{aligned}$$

The second equality follows by induction, since for all $J \subseteq_{\text{fin}} \omega$, letting \bar{c} enumerate J , $\llbracket \bigwedge_{c \in J} \widehat{\varphi}(\bar{a}, \sigma_J(c)) \rrbracket \in \mathcal{QR}(\gamma, \sigma_J(\bar{a}\bar{c}))$, and $\sigma_J(\bar{a}\bar{c})$ is disjoint from \bar{b} . The last equality is another application of equation (\dagger) above. \square

Theorem 1.4.8. *Let μ be an ergodic structure. For any sentence φ of $\mathcal{L}_{\omega_1, \omega}$, $\mu(\llbracket \varphi \rrbracket) = 0$ or 1 .*

Proof. φ is a sentence, so it has no free variables, and we can apply Lemma 1.4.7:

$$\mu(\llbracket \varphi \rrbracket) = \mu(\llbracket \varphi \wedge \varphi \rrbracket) = \mu(\llbracket \varphi \rrbracket)^2,$$

so $\mu(\llbracket \varphi \rrbracket) = 0$ or 1 . \square

The completeness of $\text{Th}(\mu)$ and the existence of Scott sentences suggests a dichotomy in the behavior of an ergodic structure μ : either some countable structure appears almost surely upon sampling from μ , or every countable structure appears with probability 0.

Definition 1.4.9. If μ is an invariant probability measure on Str_L and M is a countable structure, we say that μ is **almost surely isomorphic to M** if μ assigns measure 1 to the set of structures isomorphic to M ; equivalently, $\mu \models_{\text{a.s.}} \varphi_M$, where φ_M is the Scott sentence of M .

Definition 1.4.10. An ergodic structure μ is **properly ergodic** if it is not almost surely isomorphic to any countable structure M ; equivalently, by completeness of $\text{Th}(\mu)$, μ assigns measure 0 to every orbit of the logic action.

Remark 1.4.11. If μ is properly ergodic, then $\text{Th}(\mu)$ contains $\neg\varphi_M$ for every countable structure M , and thus $\text{Th}(\mu)$ has no countable models. It is still conceivable that $\text{Th}(\mu)$ has uncountable models (Löwenheim–Skolem does not apply to complete theories of $L_{\omega_1, \omega}$), but we will see later (Corollary 2.3.9) that this is not the case; $\text{Th}(\mu)$ has no models of any cardinality. Nevertheless, as noted in Proposition 1.4.4, every countable subset of $\text{Th}(\mu)$ has countable models. This suggests that we should restrict our attention to countable fragments of $L_{\omega_1, \omega}$.

Definition 1.4.12. A **fragment** of $L_{\omega_1, \omega}$ is a set of formulas which contains all atomic formulas and is closed under subformula and the first-order formula-building operations: finite Boolean combinations, quantification, and substitution of free variables (from the countable supply). If F is a fragment of $L_{\omega_1, \omega}$, we define $\text{Th}_F(\mu) = \{\varphi \in F \mid \mu \models_{\text{a.s.}} \varphi\}$.

Remark 1.4.13. Any countable set of formulas generates a countable fragment of $L_{\omega_1, \omega}$, the least fragment containing this set. The minimal fragment, generated by the empty set, is first-order logic FO.

Definition 1.4.14. Let F be a fragment of $L_{\omega_1, \omega}$.

- A set of sentences T is a (complete satisfiable) **F -theory** if T has a model and, for every sentence $\varphi \in F$, either $\varphi \in T$ or $\neg\varphi \in T$. Equivalently, $T = \{\psi \in F \mid M \models \psi\}$ for some structure M .
- A set of formulas $p(\bar{x})$ is an **F -type** if there is a structure M and a tuple \bar{a} from M such that $p(\bar{x}) = \{\psi(\bar{x}) \in F \mid M \models \psi(\bar{a})\}$. We say that \bar{a} **realizes** p in M .
- An F -type p is **consistent** with an F -theory T if it is realized in some model of T . We write $S_F^n(T)$ for the space of F -types in the variables x_0, \dots, x_{n-1} which are consistent with T , and $S_F^n(T)$ for the subspace of non-redundant types.

Remark 1.4.15. The Löwenheim–Skolem Theorem *does* hold for countable fragments of $L_{\omega_1, \omega}$ (see [57, Theorem 1.5.4]). Thus, if F is countable, every F -theory has a countable model and every F -type which is consistent with T is realized in a countable model of T .

We'll finish this section with several examples.

Example 1.4.16. For any real number p in the interval $(0, 1)$, we build a graph with domain ω using the Erdős–Rényi random graph model with edge probability p : we connect i and j by an edge with probability p , independently for each pair $\{i, j\}$ from ω . This random construction describes a probability measure μ_p on Str_L . Formally, if $\varphi(\bar{x})$ is a consistent quantifier-free formula describing n edges and m non-edges $(\bigwedge_{i=1}^n x_{k_i} R x_{k'_i}) \wedge (\bigwedge_{j=1}^m \neg x_{l_j} R x_{l'_j})$ and \bar{a} is any non-redundant tuple from ω , then $\mu(\llbracket \varphi(\bar{a}) \rrbracket) = p^{n+m}$. We also set $\mu(\llbracket aRa \rrbracket) = 0$ and $\mu(\llbracket aRb \wedge \neg bRa \rrbracket) = 0$ (since the edge relation of a graph is anti-reflexive and symmetric), and all this data is sufficient to determine the measure of any instance of a quantifier-free formula, and hence it induces a probability measure on Str_L by Lemma 1.2.9.

Furthermore, μ_p is an ergodic structure. Invariance follows from the fact that all edges are assigned with equal probability, and disjoint-independence from the fact that all edges are assigned independently. Note that $\mu_{1/2}$ is the limit of the sequence of Paley graphs defined in Example 1.1.11.

For all $n, m \in \omega$, we define the **extension axiom** $\psi_{n,m}$:

$$\forall x_1, \dots, x_n \forall y_1, \dots, y_m \left(\bigwedge_{i=1}^n \bigwedge_{j=1}^m x_i \neq y_j \rightarrow \left(\exists z \bigwedge_{i=1}^n x_i R z \wedge \bigwedge_{j=1}^m \neg y_j R z \right) \right).$$

Together, the extension axioms assert that any finite set of points A can be extended by a new point b in such a way that b has an arbitrary non-redundant quantifier-free type over A (consistent with anti-reflexivity and symmetry). Note that if $n = m = 0$, $\psi_{n,m}$ simply asserts that there is at least one element ($\top \rightarrow \exists z \top$).

For all p , $\mu_p(\psi_{n,m}) = 1$. Indeed, for all non-redundant $\bar{a}\bar{b}$ and all $c \notin \llbracket \bar{a}\bar{b} \rrbracket$,

$$\mu_p \left(\left[\bigwedge_{i=1}^n a_i R c \wedge \bigwedge_{j=1}^m \neg b_j R c \right] \right) = \beta$$

for some fixed $\beta > 0$, and these events are independent for distinct values of c , so by the second Borel–Cantelli Lemma (Lemma 1.2.16), almost surely infinitely many occur. In particular,

$$\mu_p \left(\left[\exists z \bigwedge_{i=1}^n a_i R z \wedge \bigwedge_{j=1}^m \neg b_j R z \right] \right) = 1.$$

Then $\llbracket \psi_{n,m} \rrbracket$ is a countable intersection of measure 1 sets, so $\mu_p(\psi_{n,m}) = 1$.

Now by a standard application of the back-and-forth method (see, for example, [75, Theorem 2.4.2]) the extension axioms $\{\psi_{n,m} \mid n, m \in \omega\}$ together with the theory of graphs (a single sentence asserting that R is symmetric and anti-reflexive, which is also satisfied almost surely by μ_p) axiomatize a complete countably categorical first-order theory T_{RG} , which is the model companion of the theory of graphs. Since μ_p almost surely satisfies T_{RG} , μ_p is almost surely isomorphic to the unique countable model of T_{RG} , known as the **Rado graph**, or the **random graph**.

Example 1.4.17. Next, consider the following two step random construction. First, pick a parameter p from $[0, 1]$ uniformly at random, then build a graph with domain ω using the Erdős–Rényi construction with edge probability p . This describes an invariant probability measure μ on Str_L , a “mixture” of the measures μ_p from Example 1.4.16. For any Borel set $X \subseteq \text{Str}_L$, $\mu(X) = \int_0^1 \mu_p(X) dp$.

This μ is also almost surely isomorphic to the Rado graph, since almost every choice of parameter p (excluding 0 and 1) generates the Rado graph almost surely. So $\text{Th}(\mu)$ is complete (it is the full theory of the Rado graph in $L_{\omega_1, \omega}$), and μ is weakly ergodic (see Remark 1.4.5). But μ is not an ergodic structure.

The easiest way to see this is to note that μ does not satisfy disjoint-independence. Indeed, $\mu(xRy) = \int_0^1 p dp = 1/2$, while $\mu((xRy) \wedge (x'Ry')) = \int_0^1 p^2 dp = 1/3 \neq (1/2)^2$. The initial random choice of p in the construction plays the role of the random variable ξ_\emptyset in an AHK representation of μ (Theorem 1.3.2).

It is also interesting to see that μ is not ergodic by exhibiting an almost surely invariant Borel set $X_{1/2}$ such that $\mu(X_{1/2})$ is not equal to 0 or 1. The point is that while we almost surely get the Rado graph (up to isomorphism) for all values of p , we can nevertheless recover p almost surely from a labeled copy of the Rado graph built with edge density p .

Let $\langle A_n \rangle_{n \in \omega}$ be some sequence of subsets of ω such that $|A_n| = n^2$ for all n . For any $M \in \text{Str}_L$, let B_n be the induced substructure of M with domain A_n . Define

$$X_{1/2} = \{M \in \text{Str}_L \mid \lim_{n \rightarrow \infty} P(xRy; B_n) < 1/2\}.$$

Observe first that $X_{1/2}$ is Borel. Indeed, it is the union over $N \in \omega$ of the intersection over $n > N$ of the event that $P(xRy; B_n) < 1/2$, and this event can be expressed as the union over the finitely many diagrams of B_n with edge density less than $1/2$.

Now by Theorem 1.2.18 and Remark 1.2.19, $\mu_p(X_{1/2}) = 0$ if $p < 1/2$, and $\mu_p(X_{1/2}) = 1$ if $p \geq 1/2$. So $\mu(X_{1/2}) = \int_0^1 \mu_p(X_{1/2}) dp = 1/2$.

Finally, we claim that $X_{1/2}$ is almost surely invariant. Fixing $\sigma \in S_\infty$, we have

$$\sigma(X_{1/2}) = \{M \in \text{Str}_L \mid \lim_{n \rightarrow \infty} P(xRy; \sigma(B_n)) < 1/2\},$$

where $\sigma(B_n)$ is the induced substructure of M on $\sigma(A_n)$. By Remark 1.2.19, the measure $\mu_p(X_{1/2})$ does not depend on the choice of the sets A_n , so we again have $\mu_p(\sigma(X_{1/2})) = 0$ if $p < 1/2$, and $\mu_p(\sigma(X_{1/2})) = 1$ if $p \geq 1/2$. In either case, $\mu_p(X_{1/2} \triangle \sigma(X_{1/2})) = 0$, and hence $\mu(X_{1/2} \triangle \sigma(X_{1/2})) = \int_0^1 \mu_p(X_{1/2} \triangle \sigma(X_{1/2})) dp = 0$.

Finally, we give an example of a properly ergodic structure, the **kaleidoscope random graph**. The kaleidoscope random graph comes from a class of examples introduced in [5]. Many more examples of properly ergodic structures are discussed in Section 2.1.

Example 1.4.18. Let L be the language $\{R_n \mid n \in \omega\}$, where each R_n is a binary relation symbol. The interpretation of each R_n will be anti-reflexive and symmetric.

We build an L -structure with domain ω , setting iR_nj with probability $1/2$, independently for each pair $\{i, j\}$ from ω and each $n \in \omega$. Just as in Example 1.4.16, this construction describes an ergodic structure μ . But in contrast to Example 1.4.16, μ is properly ergodic.

There are continuum-many non-redundant quantifier-free 2-types, $\{p_A(x, y) \mid A \subseteq \omega\}$, given by $p_A(x, y) = \{xR_ny \mid n \in A\} \cup \{\neg xR_ny \mid n \notin A\}$. Each type p_A is realized with probability 0, so any particular countable L -structure, which must realize some countable collection of these types, occurs with probability 0 up to isomorphism.

In fact, for every set $A \subseteq \omega$, the theory $\text{Th}(\mu)$ contains the sentence

$$\neg \exists x, y \bigwedge_{n \in A} xR_ny \wedge \bigwedge_{n \notin A} \neg xR_ny.$$

Since all possible quantifier-free 2-types are ruled out by $\text{Th}(\mu)$, this theory has no models of any cardinality. Note, however, that any countable fragment F of $L_{\omega_1, \omega}$ only contains countably many of the sentences above, so $\text{Th}_F(\mu)$ will only rule out countably many of the quantifier-free 2-types.

Restricting to the first-order fragment, the theory $\text{Th}_{\text{FO}}(\mu)$ has several nice properties. It is the model companion of the universal theory asserting that each R_n is anti-reflexive and symmetric. It can be axiomatized by extension axioms, analogous to those in Example 1.4.16, asserting, for each finite sublanguage $L^* \subseteq L$, that any finite set of points A can be extended by a new element b with an arbitrary non-redundant quantifier-free type over A in the language L^* . The reduct of $\text{Th}_{\text{FO}}(\mu)$ to any of these finite sublanguages is countably categorical, but $\text{Th}_{\text{FO}}(\mu)$ has continuum-many countable models (clearly, since there are continuum-many quantifier-free 2-types consistent with $\text{Th}_{\text{FO}}(\mu)$).

This example can be modified to produce the **kaleidoscope random n -hypergraph** for any n . We call the case $n = 1$ the **kaleidoscope random predicate**.

1.5 Trivial definable closure

Ackerman, Freer, and Patel showed [3, 2] that the characteristic property satisfied by the theory (in a countable fragment of $L_{\omega_1, \omega}$) of an ergodic structure is trivial definable closure.

Definition 1.5.1. Let F be a fragment of $L_{\omega_1, \omega}$. An F -theory T has **trivial definable closure (dcl)** if there is no formula $\varphi(\bar{x}, y)$ in F such that T contains

$$\exists \bar{x} \exists !y \left(\left(\bigwedge_{i=1}^n y \neq x_i \right) \wedge \varphi(\bar{x}, y) \right).$$

Here $\exists !y$ is the standard abbreviation for “there exists a unique y ”.

Remark 1.5.2. Recall that our F -theories are always complete. If $M \models T$, then T has trivial dcl if and only if the closure operator dcl_F is trivial on M : $\text{dcl}_F(A) = A$ for all $A \subseteq M$, where $\text{dcl}_F(A)$ is the set of all $b \in M$ such that b is the unique element of M satisfying some formula in F with parameters from A .

The easy direction of the Ackerman–Freer–Patel characterization is that the theory of an ergodic structure always has trivial dcl. This comes down to the fact that if a measure is invariant for the action of some group G , then no positive-measure set can have infinitely many almost surely disjoint images under the action of G . The following theorem began its life in [2, Theorem 4.1], which only covered the case when μ is almost surely isomorphic to a countable structure; the full form appears in [3].

Theorem 1.5.3. *Let μ be an ergodic structure and F a fragment of $L_{\omega_1, \omega}$. Then $\text{Th}_F(\mu)$ has trivial dcl.*

Proof. Suppose there is a formula $\varphi(\bar{x}, y)$ in F such that

$$\mu(\exists \bar{x} \exists! y (\bigwedge_{i=1}^n y \neq x_i) \wedge \varphi(\bar{x}, y)) = 1.$$

Let $\psi(\bar{x}, y)$ be the formula $(\bigwedge_{i=1}^n y \neq x_i) \wedge \varphi(\bar{x}, y)$.

By countable additivity of μ , there is a tuple \bar{a} from ω such that

$$\mu(\llbracket \exists! y \psi(\bar{a}, y) \rrbracket) > 0.$$

Let $\theta(\bar{a})$ be the formula $\forall z_1 \forall z_2 (\psi(\bar{a}, z_1) \wedge \psi(\bar{a}, z_2) \rightarrow (z_1 = z_2))$, so that $\exists! y \psi(\bar{a}, y)$ is equivalent to

$$\exists y \psi(\bar{a}, y) \wedge \theta(\bar{a}).$$

Since this formula has positive measure, countable additivity again implies that there is some $b \in \omega \setminus \|\bar{a}\|$ such that

$$\beta = \mu(\llbracket \psi(\bar{a}, b) \wedge \theta(\bar{a}) \rrbracket) > 0.$$

By invariance, for any $c \in \omega \setminus \|\bar{a}\|$, we also have

$$\mu(\llbracket \psi(\bar{a}, c) \wedge \theta(\bar{a}) \rrbracket) = \beta.$$

But θ ensures that $\psi(\bar{a}, b) \wedge \theta(\bar{a})$ and $\psi(\bar{a}, c) \wedge \theta(\bar{a})$ are inconsistent when $b \neq c$, so, computing the measure of the disjoint union,

$$\mu\left(\bigcup_{b \in \omega \setminus \|\bar{a}\|} \llbracket \psi(\bar{a}, b) \wedge \theta(\bar{a}) \rrbracket\right) = \sum_{b \in \omega \setminus \|\bar{a}\|} \beta = \infty,$$

which is a contradiction. □

The main result of [2] is a characterization of those countable structures M such that there exists an ergodic structure μ which is almost surely isomorphic to M . That characterization was given in terms of trivial “group-theoretic” dcl (the group is $\text{Aut}(M)$).

Definition 1.5.4. A countable structure M has **trivial group-theoretic dcl** if for any finite subset $A \subseteq M$ and element $b \in M \setminus A$, there is an automorphism $\sigma \in \text{Aut}(M)$ such that $\sigma(a) = a$ for all $a \in A$, but $\sigma(b) \neq b$.

Theorem 1.5.5 (AFP [2]). *Let M be a countable structure. There exists an ergodic structure which is almost surely isomorphic to M if and only if M has trivial group-theoretic dcl.*

Remark 1.5.6. In [2], Theorem 1.5.5 is stated for invariant probability measures on Str_L instead of ergodic structures. However, the method used in [2] and [3] of obtaining a measure via i.i.d. sampling from a Borel structure, which we use again in Section 2.5, always produces an ergodic structure. This was mentioned in passing in [2]. A proof is given in Lemma 2.5.2 below.

Theorem 1.5.7 (AFP [3]). *Let F be a countable fragment of $L_{\omega_1, \omega}$, and let T be an F -theory. Then $T = \text{Th}_F(\mu)$ for some ergodic structure μ if and only if T has trivial dcl.*

Remark 1.5.8. Theorem 1.5.5 follows from Theorem 1.5.7. Indeed, let M be a countable structure, let φ_M be its Scott sentence, and let F_M be the countable fragment of $L_{\omega_1, \omega}$ generated by φ_M . Then (as a consequence of the proof of Scott's Theorem) $\text{Th}_{F_M}(M)$ has trivial dcl if and only if M has trivial group-theoretic dcl, and (since $\text{Th}_{F_M}(M)$ is axiomatized by φ_M) $\text{Th}_{F_M}(M) = \text{Th}_{F_M}(\mu)$ if and only if $\mu \models \varphi_M$.

The next example illustrates the important distinction between trivial dcl and trivial group-theoretic dcl in our context.

Example 1.5.9. Let T be the first-order theory of the kaleidoscope random predicate (see Example 1.4.18 above) in the language $\{P_n \mid n \in \omega\}$. The theory T says that for every $m \in \omega$ and every subset $A \subseteq [m]$, there exists an element x such that $P_n(x)$ holds if and only if $n \in A$.

Now let T' be T together with the infinitary sentence

$$\forall x \forall y (x \neq y) \rightarrow \bigvee_{n \in \omega} \neg(P_n(x) \leftrightarrow P_n(y)).$$

The kaleidoscope random predicate almost surely satisfies T' . Each of the continuum-many quantifier-free 1-types is realized with probability 0, and since the quantifier-free 1-types of distinct elements of ω are independent, almost surely no quantifier-free 1-type is realized more than once.

In a model M of T' , no two elements have the same quantifier-free 1-type. Hence $\text{Aut}(M)$ is the trivial group, and M has non-trivial group-theoretic dcl. But the countable fragment F of $L_{\omega_1, \omega}$ generated by T' does not contain the conjunctions of the form $\bigwedge_{n \in A} P_n(x) \wedge \bigwedge_{n \notin A} \neg P_n(x)$ for $A \subseteq \omega$ needed to pin down elements uniquely. In fact, any completion of T' to an F -theory has trivial dcl, by Theorem 1.5.3.

Theorem 1.5.7 gives a satisfying characterization of the theories of ergodic structures, and the ‘‘AFP construction’’ from [2] and [3] gives a very flexible method for building ergodic structures. However, the version of the construction in [2] always produces an ergodic structure which is almost surely isomorphic to a countable structure, and the version in [3]

is not precise enough to control whether the resulting ergodic structure is properly ergodic. In the paper [5], the construction was modified to produce properly ergodic structures, but only satisfying a very restrictive class of “approximately countably categorical” theories.

In the next chapter, we seek to answer two natural questions. How can we construct properly ergodic structures? And which F -theories arise as the theories of properly ergodic structures?

Chapter 2

Properly ergodic structures

2.1 Examples

We begin with a host of examples of properly ergodic structures. These illustrate some of the key features discussed in the following sections. They will be described informally as probabilistic constructions of countable structures. Just as in Examples 1.4.16 and 1.4.18 above, these constructions specify probability measures on Str_L , and it should be straightforward to check invariance and disjoint-independence.

In the descriptions below, all random choices are made independently, unless otherwise specified. When we say that we pick a random element $A \in 2^\omega$, we always refer to the natural measure on 2^ω , the infinite product of the Bernoulli(1/2) measure on $2 = \{0, 1\}$. We identify such an $A \in 2^\omega$ with both a subset of ω and an infinite binary sequence.

Example 2.1.1 (Random geometric graphs). Bonato and Janssen [15] introduced a new geometric random graph model: given a countable dense set V of points in some metric space (S, d) , and a real number $p \in (0, 1)$, construct a graph on V by connecting x and y by an edge with probability p , for each pair $\{x, y\}$ from V such that $d(x, y) \leq 1$. If, in addition, we obtain the set V by i.i.d. sampling from some distribution on S with a strictly positive density function, then this random construction describes an ergodic structure, called the **random geometric graph on S with edge probability p** (we suppress the distribution on S , since the particular choice of distribution turns out to be irrelevant).

Bonato and Janssen showed that if we take S to be ℓ_∞^n for some n , then the random geometric graph on S with edge probability p is almost surely isomorphic to a single countable graph, denoted GR_n . Later, Balister, Bollobás, Gunderson, Leader, and Walters [13] showed that the spaces ℓ_∞^n are the unique normed spaces with this property: if S is a normed space that is not isometric to ℓ_∞^n for any n , then the random geometric graph on S with edge probability p is properly ergodic.

We will see that the presence of a formula $\chi(\bar{x})$ of positive measure, such that every type containing χ has probability 0 of being realized, is a characteristic feature of properly ergodic

structures. In the kaleidoscope random graph, Example 1.4.18, $x \neq y$ is such a formula, since every non-redundant quantifier-free 2-type is realized with probability 0. In contrast to the kaleidoscope random graph, Example 2.1.2 shows that these types of probability 0 may almost surely have infinitely many realizations if they are realized at all.

On the other hand, in Example 2.1.3, we see that requiring each of a family of continuum-many 1-types to be realized infinitely many times (if at all) can rule out proper ergodicity. This shows that the consistency of continuum-many types in a theory with trivial dcl is not sufficient for the existence of a properly ergodic model of the theory. These phenomena motivate the definition of rootedness in Section 2.4.

Example 2.1.2 (The max graph). As in Example 1.4.18, let $L = \{R_n \mid n \in \omega\}$, where each R_n is a binary relation with an anti-reflexive and symmetric interpretation. We build a random L -structure with domain ω . For each $i \in \omega$, choose a random element $A_i \in 2^\omega$. Now for each pair $\{i, j\}$, let $A_{ij} = \max(A_i, A_j)$, where we give 2^ω its lexicographic order. We set iR_nj if and only if $n \in A_{ij}$.

We have continuum-many non-redundant quantifier-free 2-types $\{p_A \mid A \in 2^\omega\}$, where $xR_ny \in p_A$ if and only if $n \in A$, and each is realized with probability 0, since if (i, j) realizes p_A , we must have $A_i = A$ or $A_j = A$.

As long as A_i is not the constant 0 sequence (which appears with probability 0), there is a positive probability, conditioned on the choice of A_i , that if $j \neq i$, then $A_j \leq A_i$, and hence $\text{qftp}(i, j) = p_{A_i}$. Moreover, all these events are conditionally independent. Thus, almost surely, any quantifier-free 2-type that is realized is realized infinitely many times. However, since the probability that $A_i = A_j$ when $i \neq j$ is 0, almost surely all realizations of p_{A_i} have a common intersection, namely the vertex i .

Example 2.1.3 (A non-example). Let $L = \{E\} \cup \{P_n \mid n \in \omega\}$, and let T be the first-order model companion of the universal theory asserting that E is an equivalence relation and the P_n are unary predicates respecting E (if xEy , then $P_n(x)$ if and only if $P_n(y)$). This is similar to the first-order theory of the kaleidoscope random predicate, but with each element replaced by an infinite E -class.

There is no properly ergodic structure that satisfies T almost surely. Indeed, suppose $\mu \models_{\text{a.s.}} T$. Then for every quantifier-free 1-type p , there is some probability $\mu(p)$ that p is the quantifier-free type of the element $i \in \omega$, and, by invariance, $\mu(p)$ does not depend on the choice of i . We denote by $S_{\text{qf}}^1(\mu)$ the countable set of quantifier-free 1-types with positive measure. If $\sum_{p \in S_{\text{qf}}^1(\mu)} \mu(p) = 1$, then almost surely only the types in $S_{\text{qf}}^1(\mu)$ are realized, since $\mu \models_{\text{a.s.}} \forall x \bigwedge_{p \in S_{\text{qf}}^1(\mu)} \bigwedge_{\varphi \in p} \varphi(x)$. Further, μ determines, for each $p \in S_{\text{qf}}^1(\mu)$, the number of E -classes on which p is realized (among $\{1, 2, \dots, \aleph_0\}$), since each of the countably many choices is expressible by a sentence of $L_{\omega_1, \omega}$. The data of which quantifier-free 1-types are realized, and how many E -classes realize each, determines a unique L -structure up to isomorphism, so μ is not properly ergodic.

On the other hand, if $\sum_{p \in S_{\text{qf}}^1(\mu)} \mu(p) < 1$, then almost surely some types that are not in $S_{\text{qf}}^1(\mu)$ are realized. Any such type p is realized with probability 0, and, by disjoint-

independence, the quantifier-free 1-types of distinct elements of ω are independent. So, almost surely, each of the types of probability 0 is realized at most once. This contradicts the fact that any realized type must be realized on an entire infinite E -class.

Example 2.1.4 (The necessity of infinitary logic). Let $L = \{P\} \cup \{R_j^i \mid i, j \in \omega\}$, where P is a unary predicate and the R_j^i are binary relations, and let T be the first-order model companion of the following universal theory:

- $\forall x \forall y R_j^i(x, y) \rightarrow (P(x) \wedge \neg P(y))$ for all i and j .
- $\forall x \forall y \neg R_0^i(x, y) \wedge R_0^{i'}(x, y)$ for all $i \neq i'$.
- $\forall x \forall y R_{j+1}^i(x, y) \rightarrow R_j^i(x, y)$ for all i and j .

Thus, a model of T is a bipartite graph in which each edge from x to y is labeled by some $i \in \omega$ (in the superscript) and the set of all $j < k$ for some $k \in \omega_+ \cup \{\infty\}$ (in the subscript). Here $\omega_+ = \omega \setminus \{0\}$.

Now T is a complete theory with quantifier elimination and with only countably many types over \emptyset . Hence, by countable additivity, if μ is an ergodic structure that satisfies T almost surely, then there is no positive-measure first-order formula $\chi(\bar{x})$ such that every type containing χ has measure 0. Nevertheless, we will describe a properly ergodic structure that almost surely satisfies T .

First, for each $x \in \omega$, set $P(x)$ with probability $1/2$, and pick $A_x \in 2^\omega$ at random. Next, for each pair $x \neq y$, if $P(x)$ and $\neg P(y)$, then we choose which of the R_j^i to assign to (x, y) . Choose $i \in \omega$ at random, setting $i = n$ with probability $2^{-(n+1)}$. If $i \in A_x$, choose $k \in \omega_+ \cup \{\infty\}$ at random, setting $k = \infty$ with probability $1/2$ and $k = n$ with probability $2^{-(n+1)}$ for $n \in \omega_+$. If $i \notin A_x$, choose $k \in \omega_+$ at random, setting $k = n$ with probability 2^{-n} . Finally, set $R_j^i(x, y)$ for all $j < k$.

In the resulting random structure, we can almost surely recover A_x from every $x \in P$, since if $i \in A_x$, then almost surely there is some y such that $R_j^i(x, y)$ for all $j \in \omega$ (that is, the choice $k = \infty$ was made for the pair (x, y)), whereas this outcome is impossible if $i \notin A_x$. Thus the structure encodes a countable set of elements of 2^ω , each of which occurs with probability 0.

The information encoding A_x is part of the 1-type of x in any countable fragment of $L_{\omega_1, \omega}$ containing the infinitary formulas $\{\exists y \bigwedge_{j \in \omega} R_j^i(x, y) \mid i \in \omega\}$, but it is not expressible in first-order logic.

With the exception of Example 2.1.1, the preceding examples have all used infinite languages, as this is the most convenient setting in which to split the measure over continuum-many types. We conclude with an elementary example in the language with a single binary relation.

Example 2.1.5 (An example in a finite language). Let $L = \{R\}$, where R is a binary relation. In our probabilistic construction, we will enforce the following almost surely:

- Let $O = \{x \mid R(x, x)\}$, and $P = \{x \mid \neg R(x, x)\}$. Then O and P are both infinite sets.
- If $R(x, y)$, then either x and y are both in O , or x is in P and y is in O .
- R is a preorder on O . Denote by xEy the induced equivalence relation $R(x, y) \wedge R(y, x)$. Then E has infinitely many infinite classes, and R linearly orders the E -classes with order type ω .
- Given $x \in P$ and $y, z \in O$, if $R(x, y)$ and yEz , then $R(x, z)$. So R relates each element of P to some subset of the E -classes.

Thus we can interpret the kaleidoscope random predicate on P , where the n^{th} predicate P_n holds of x if and only if x is R -related to the n^{th} class in the linear order on O .

Now it is straightforward to describe the probabilistic construction. For each $i \in \omega$, let $R(i, i)$ hold with probability $1/2$. This determines whether i is in O or P . If $i \in O$, we choose which E -class to put i in, under the order induced by R , selecting the n^{th} class with probability $2^{-(n+1)}$. These choices determine all the R -relations between elements of O . On the other hand, if $i \in P$, we pick $A_i \in 2^\omega$ at random and relate i to each element of the n^{th} class in O if and only if $n \in A_i$.

This describes an ergodic structure μ , since the quantifier-free types of disjoint tuples are independent. We obtain the properties described in the bullet points above almost surely, and since ω is rigid, any isomorphism between structures satisfying these properties must preserve the order on the E -classes. For any subset of the E -classes, the probability is 0 that there is an element of P which is related to exactly those E -classes, and so μ is properly ergodic.

2.2 Pithy Π_2 theories

It is a well-known fact, originally observed by Chang, that if T is a theory in a countable fragment F of $L_{\omega_1, \omega}$, then the models of T are exactly the reducts to L of the models of a first-order theory T' in a larger language $L' \supseteq L$ that omit a countable set of types Q .

The idea is to “Morleyize”: introduce a new relation symbol R_φ for every formula $\varphi(\bar{x})$ in F and encode the intended interpretations of the R_φ in the theory T' . The role of the countable set of types Q is to achieve this for infinitary conjunctions and disjunctions, which cannot be accounted for in first-order logic. There are two features of this construction that will be useful for us: First, it reduces F -types to quantifier-free types. Second, T' can be axiomatized by pithy Π_2 sentences, also called “one-point extension axioms”.

Definition 2.2.1. A first-order sentence is **pithy Π_2** if it is universal (Π_1) or if it has the form $\forall \bar{x} \exists y \varphi(\bar{x}, y)$, where $\varphi(\bar{x}, y)$ is quantifier-free, \bar{x} is a tuple of variables (possibly empty), and y is a single variable. A **pithy Π_2 theory** is a set of pithy Π_2 sentences.

Note that, for us, all pithy Π_2 theories are first-order.

Theorem 2.2.2. *Let F be a countable fragment of $L_{\omega_1, \omega}$ and T an F -theory. Then there is a language $L' \supseteq L$, a pithy Π_2 L' -theory T' , and a countable set of partial quantifier-free L' -types Q with the following properties.*

- (a) *There is a bijection between formulas $\varphi(\bar{x})$ in F and atomic L' -formulas $R_\varphi(\bar{x})$ which are not in L , such that if $M \models T'$ omits all the types in Q , then $M \models \forall \bar{x} \varphi(\bar{x}) \leftrightarrow R_\varphi(\bar{x})$.*
- (b) *The reduct to L is a bijection between the class of models of T' omitting all the types in Q and the class of models of T .*

Proof. Let $L' = L \cup \{R_\varphi \mid \varphi(\bar{x}) \in F\}$, where the arity of the relation symbol R_φ is the number of free variables in $\varphi(\bar{x})$. Recall that we allow 0-ary relation symbols (i.e., propositional symbols). Thus, we include a 0-ary relation R_ψ symbol for every sentence $\psi \in F$.

Let T_{def} be the theory consisting of the following axioms, for each formula $\varphi(\bar{x}) \in F$:

- (1) $\forall \bar{x} (R_\varphi(\bar{x}) \leftrightarrow \varphi(\bar{x}))$, if $\varphi(\bar{x})$ is atomic.
- (2) $\forall \bar{x} (R_\varphi(\bar{x}) \leftrightarrow \neg R_\psi(\bar{x}))$, if $\varphi(\bar{x})$ is of the form $\neg\psi(\bar{x})$.
- (3) $\forall \bar{x} (R_\varphi(\bar{x}) \leftrightarrow R_\psi(\bar{x}) \wedge R_\theta(\bar{x}))$, if $\varphi(\bar{x})$ is of the form $\psi(\bar{x}) \wedge \theta(\bar{x})$.
- (4) $\forall \bar{x} (R_\varphi(\bar{x}) \leftrightarrow R_\psi(\bar{x}) \vee R_\theta(\bar{x}))$, if $\varphi(\bar{x})$ is of the form $\psi(\bar{x}) \vee \theta(\bar{x})$.
- (5) $\forall \bar{x} (R_\varphi(\bar{x}) \rightarrow R_{\psi_i}(\bar{x}))$ for all i , if $\varphi(\bar{x})$ is of the form $\bigwedge_{i \in I} \psi_i(\bar{x})$.
- (6) $\forall \bar{x} (R_{\psi_i}(\bar{x}) \rightarrow R_\varphi(\bar{x}))$ for all i , if $\varphi(\bar{x})$ is of the form $\bigvee_{i \in I} \psi_i(\bar{x})$.
- (7) $\forall \bar{x} (R_\varphi(\bar{x}) \leftrightarrow \forall y R_\psi(\bar{x}, y))$, if $\varphi(\bar{x})$ is of the form $\forall y \psi(\bar{x}, y)$.
- (8) $\forall \bar{x} (R_\varphi(\bar{x}) \leftrightarrow \exists y R_\psi(\bar{x}, y))$, if $\varphi(\bar{x})$ is of the form $\exists y \psi(\bar{x}, y)$.

Note that all the axioms of T_{def} are first-order and universal except for those of type (7) and (8), which are first-order and pithy Π_2 when put in prenex normal form.

The implications in the axioms of type (5) and (6) cannot be made into bi-implications, since arbitrary countable infinite conjunctions and disjunctions are not expressible in first-order logic. To ensure that the corresponding R_φ have their intended interpretation, we let Q consist of the partial quantifier-free types:

- (i) $q_\varphi(\bar{x}) = \{R_{\psi_i}(\bar{x}) \mid i \in I\} \cup \{\neg R_\varphi(\bar{x})\}$, for all $\varphi(\bar{x})$ of the form $\bigwedge_{i \in I} \psi_i(\bar{x})$
- (ii) $q_\varphi(\bar{x}) = \{\neg R_{\psi_i}(\bar{x}) \mid i \in I\} \cup \{R_\varphi(\bar{x})\}$, for all $\varphi(\bar{x})$ of the form $\bigvee_{i \in I} \psi_i(\bar{x})$.

It is now straightforward to show by induction on the complexity of formulas that if a model $M \models T_{\text{def}}$ omits every type in Q , then for all $\varphi(\bar{x})$ in F and all \bar{a} from M , $M \models \varphi(\bar{a})$ if and only if $M \models R_\varphi(\bar{a})$. This establishes (a) and implies that every L -structure N admits a unique expansion to an L' -structure N' which satisfies T_{def} and omits every type in Q .

Set $T' = T_{\text{def}} \cup \{R_\psi \mid \psi \in T\}$. If M is a model of T' which omits every type in Q , then the reduct $M \upharpoonright L$ is a model of T . And if $N \models T$, then $N' \models T'$, where N' is the canonical expansion of N . This establishes (b). \square

Corollary 2.2.3. *There is a bijection between the ergodic L -structures (measures on Str_L) which almost surely satisfy T and the ergodic L' -structures (measures on $\text{Str}_{L'}$) which almost surely satisfy T' and omit all the types in Q . This bijection sends ergodic structures to ergodic structures and properly ergodic structures to properly ergodic structures.*

Proof. The reduct \upharpoonright_L is a continuous map $\text{Str}_{L'} \rightarrow \text{Str}_L$, since the preimages of clopen sets in Str_L are also clopen sets in $\text{Str}_{L'}$. By Theorem 2.2.2, \upharpoonright_L is a bijection between the subspace X' of $\text{Str}_{L'}$ consisting of models of T' which omit all the types in Q and the subspace X of Str_L consisting of models of T . Upon restricting to these subspaces, the inverse of \upharpoonright_L is a Borel map, since the image of clopen set in X' (described by a first-order quantifier-free formula) is a Borel set in X (described by a formula of $L_{\omega_1, \omega}$). Hence \upharpoonright_L is a Borel isomorphism between these subspaces, and it induces a bijection between the set of probability measures on $\text{Str}_{L'}$ concentrating on X' and the sets of probability measures on Str_L concentrating on X . Moreover, \upharpoonright_L preserves the logic action, so the induced bijection on measures preserves invariance, ergodicity, and proper ergodicity. \square

2.3 Morley–Scott analysis of ergodic structures

Throughout this section, let μ be an ergodic structure. We perform a “Morley–Scott analysis” of μ , based on Morley’s proof [77] that the number of isomorphism classes of countable models of a sentence of $L_{\omega_1, \omega}$ is countable, \aleph_1 , or 2^{\aleph_0} . This allows us to assign the analog of a Scott rank and Scott fragment to μ , and studying the theory of μ in its Scott fragment clarifies the properly ergodic dichotomy. We also obtain an analog of Vaught’s Conjecture for ergodic structures.

If F is a countable fragment of $L_{\omega_1, \omega}$ and p is an F -type, then we denote by $\theta_p(\bar{x})$ the conjunction of all the formulas in p , $\bigwedge_{\varphi \in p} \varphi(\bar{x})$. This is a formula of $L_{\omega_1, \omega}$ (although not a formula of F in general), so it is assigned a measure by μ . We write $\mu(p)$ as shorthand for $\mu(\theta_p(\bar{x}))$. This is the probability, according to μ , that any given non-redundant tuple from ω satisfies p . In particular, if $\mu(p) > 0$, then p is non-redundant.

Definition 2.3.1. We denote by $S_F^n(\mu)$ the set $\{p \mid \mu(p) > 0\}$ of **positive-measure F -types** in the variables x_0, \dots, x_{n-1} . In the case $n = 0$, $S_F^0(\mu)$ has one element, namely $\text{Th}_F(\mu)$.

Lemma 2.3.2. *For all $n \in \omega$, $|S_F^n(\mu)| \leq \aleph_0$.*

Proof. Fix a non-redundant tuple \bar{a} from ω . The sets $\{\llbracket \theta_p(\bar{a}) \rrbracket \mid p \in S_F^n(\mu)\}$ are disjoint sets of positive measure in Str_L . By additivity of μ , $T_m = \{p \in S_F^n(\mu) \mid \mu(p) \geq 1/m\}$ is finite (of size at most m) for all $m \in \omega$, so $S_F^n(\mu) = \bigcup_{m \in \omega} T_m$ is countable. \square

We build a sequence $\langle F_\alpha \rangle_{\alpha \in \omega_1}$ of countable fragments of $L_{\omega_1, \omega}$ of length ω_1 , depending on the ergodic structure μ :

$$\begin{aligned} F_0 &= \text{FO, the first-order fragment.} \\ F_{\alpha+1} &= \text{the fragment generated by } F_\alpha \cup \left\{ \theta_p(\bar{x}) \mid p \in \bigcup_{n \in \omega} S_{F_\alpha}^n(\mu) \right\}. \\ F_\gamma &= \bigcup_{\alpha < \gamma} F_\alpha, \text{ if } \gamma \text{ is a limit ordinal.} \end{aligned}$$

Definition 2.3.3. We say that $p \in S_{F_\alpha}^n(\mu)$ **splits at** $\beta > \alpha$ if $\mu(q) < \mu(p)$ for all types $q \in S_{F_\beta}^n(\mu)$ such that $p \subseteq q$. We say that p **splits later** if there exists β such that p splits at β . We say that μ **has stabilized** at γ if for all $n \in \omega$, no type in $S_{F_\gamma}^n(\mu)$ splits later.

Lemma 2.3.4. *Let $\alpha < \beta < \gamma$.*

- (1) *If a type $p \in S_{F_\alpha}^n$ splits at β , then p also splits at γ .*
- (2) *Suppose $p \in S_{F_\beta}^n$ splits at γ . Then $p' = p \cap F_\alpha$ also splits at γ .*
- (3) *If μ has stabilized at α , then μ has stabilized at β .*

Proof. (1): Pick $q \in S_{F_\gamma}^n(\mu)$ with $p \subseteq q$, and let $q' = q \cap F_\beta$. Then $\mu(q) \leq \mu(q') < \mu(p)$, since p splits at β .

(2): First, $\mu(p) \leq \mu(p')$, so $p' \in S_{F_\alpha}^n(\mu)$. Pick $q \in S_{F_\gamma}^n(\mu)$ such that $p' \subseteq q$. If $p \subseteq q$, then $\mu(q) < \mu(p) \leq \mu(p')$, since p splits at γ . And if $p \not\subseteq q$, then $\mu(q) \leq \mu(p') - \mu(p) < \mu(p')$, since $\mu(p) > 0$. In either case, $\mu(q) < \mu(p')$, so p' splits at γ .

(3): If μ has not stabilized at β , then some F_β -type p splits later. By (2), $p' = p \cap F_\alpha$ also splits later, so μ has not stabilized at α . \square

Lemma 2.3.5. *There is some countable ordinal γ such that μ has stabilized at γ .*

Proof. Fix $n \in \omega$. For $\alpha \in \omega_1$, let

$$\begin{aligned} \text{Sp}(\alpha) &= \{p \in S_{F_\alpha}^n(\mu) \mid p \text{ splits later}\}, \\ r_\alpha &= \sup\{\mu(p) \mid p \in \text{Sp}(\alpha)\}. \end{aligned}$$

Note that $\text{Sp}(\alpha)$ is countable, since $S_{F_\alpha}^n(\mu)$ is. If $\text{Sp}(\alpha)$ is non-empty, then $r_\alpha > 0$, and in fact the supremum is achieved by finitely many types, since $\sum_{p \in \text{Sp}(\alpha)} \mu(p) \leq 1$.

By Lemma 2.3.4 (2), the measure of any type in $\text{Sp}(\beta)$ is bounded above by the measure of a type in $\text{Sp}(\alpha)$, namely its restriction to F_α . So we have $r_\beta \leq r_\alpha$ whenever $\alpha < \beta$.

Now assume for contradiction that $\text{Sp}(\alpha)$ is non-empty for all α . We build a cofinal sequence $\langle \alpha_\delta \rangle_{\delta \in \omega_1}$ in ω_1 , such that $\langle r_{\alpha_\delta} \rangle_{\delta \in \omega_1}$ is strictly decreasing.

At each successor stage, we are given $\alpha = \alpha_\delta$, and we seek $\beta = \alpha_{\delta+1}$ with $r_\beta < r_\alpha$. Since $\text{Sp}(\alpha)$ is non-empty, there are finitely many types p_1, \dots, p_n of maximal measure r_α . For each i , pick $\beta_i > \alpha$ such that p_i splits at $\beta_i > \alpha$, and let $\beta = \max(\beta_1, \dots, \beta_n)$. By Lemma 2.3.4 (1), each p_i splits at β . Let q be a type in $\text{Sp}(\beta)$ with $\mu(q) = r_\beta$, and let $q' = q \cap F_\alpha$. By Lemma 2.3.4 (2), $q' \in \text{Sp}(\alpha)$. If q' is one of the p_i , then $\mu(q) < \mu(p_i) = r_\alpha$, since p_i splits at β . If not, then $\mu(q) \leq \mu(q') < r_\alpha$. In either case, $r_\beta = \mu(q) < r_\alpha$.

If λ is a countable limit ordinal, let $\alpha_\lambda = \sup_{\delta < \lambda} \alpha_\delta$. This is an element of ω_1 , since ω_1 is regular. And for all $\delta < \lambda$, since $\alpha_{(\delta+1)} < \alpha_\lambda$, $r_{\alpha_\lambda} \leq r_{\alpha_{\delta+1}} < r_{\alpha_\delta}$.

Of course, there is no strictly decreasing sequence of real numbers of length ω_1 , since \mathbb{R} contains a countable dense set. Hence there is some $\gamma_n \in \omega_1$ such that $\text{Sp}(\gamma_n)$ is empty, i.e., no type in $S_{F_{\gamma_n}}^n(\mu)$ splits later. Let $\gamma = \sup_{n \in \omega} \{\gamma_n\} \in \omega_1$. Then μ has stabilized at γ . \square

We can think of the minimal ordinal γ such that μ has stabilized at γ as an analog of the Scott rank for the ergodic structure μ . Since no F_γ -type splits later, every positive-measure $F_{\gamma+1}$ -type q is isolated by the $F_{\gamma+1}$ -formula θ_p for its restriction $p = q \cap F_\gamma$. Lemma 2.3.6 says that if every tuple satisfies one of these positive-measure types almost surely, then μ almost surely satisfies a Scott sentence.

Lemma 2.3.6. *Suppose that μ has stabilized at γ , and that for all $n \in \omega$,*

$$\sum_{p \in S_{F_\gamma}^n(\mu)} \mu(p) = 1.$$

Then μ is almost surely isomorphic to a countable structure.

Proof. For each type $r(\bar{x}) \in S_{F_\gamma}^n(\mu)$ (we include the case $n = 0$), let E_r be the set of types $q(\bar{x}, y) \in S_{F_\gamma}^{n+1}(\mu)$ with $r \subseteq q$. Fix a type $p(\bar{x}) \in S_{F_\gamma}^n(\mu)$, and let φ_p be the sentence

$$\forall \bar{x} \left(\theta_p(\bar{x}) \rightarrow \forall (y \notin \bar{x}) \bigvee_{q \in E_p} \theta_q(\bar{x}, y) \right),$$

and let ψ_p be the sentence

$$\forall \bar{x} \left(\theta_p(\bar{x}) \rightarrow \bigwedge_{q \in E_p} \exists (y \notin \bar{x}) \theta_q(\bar{x}, y) \right).$$

Here $\forall (y \notin \bar{x}) \rho(\bar{x}, y)$ and $\exists (y \notin \bar{x}) \rho(\bar{x}, y)$ are shorthand for $\forall y ((\bigwedge_{i=1}^n y \neq x_i) \rightarrow \rho(\bar{x}, y))$ and $\exists y ((\bigwedge_{i=1}^n y \neq x_i) \wedge \rho(\bar{x}, y))$, respectively. We would like to show that μ satisfies φ_p and ψ_p almost surely.

By assumption, and since every $q \in S_{F_\gamma}^{n+1}(\mu)$ is in E_r for a unique $r \in S_{F_\gamma}^n(\mu)$,

$$1 = \sum_{q \in S_{F_\gamma}^{n+1}(\mu)} \mu(q) = \sum_{r \in S_{F_\gamma}^n(\mu)} \sum_{q \in E_r} \mu(q).$$

Then for all $r \in S_{F_\gamma}^n(\mu)$, we must have

$$\mu(r) = \sum_{q \in E_r} \mu(q).$$

In particular, this is true for $r = p$, so for any tuple \bar{a} and any b not in \bar{a} , $\left[\bigvee_{q \in E_p} \theta_q(\bar{a}, b) \right]$ has full measure in $\llbracket \theta_p(\bar{a}) \rrbracket$ (this is true even when \bar{a} is redundant, since in that case $\llbracket \theta_p(\bar{a}) \rrbracket$ has measure 0). A countable intersection (over $b \in \omega \setminus \|\bar{a}\|$) of subsets of $\llbracket \theta_p(\bar{a}) \rrbracket$ with full measure still has full measure, so

$$\mu \left(\left[\left[\theta_p(\bar{a}) \rightarrow \forall (y \notin \bar{a}) \bigvee_{q \in E_p} \theta_q(\bar{a}, y) \right] \right] \right) = 1.$$

Taking another countable intersection over all tuples \bar{a} , we have $\mu \models_{\text{a.s.}} \varphi_p$.

We turn now to ψ_p . Since μ stabilizes at γ , there is a (necessarily unique) extension of p to a type $p^* \in S_{F_{\gamma+1}}^n(\mu)$ with $\mu(p^*) = \mu(p)$. Let $q(\bar{x}, y)$ be any type in E_p , and let $v_q(\bar{x}) \in F_{\gamma+1}$ be the formula $(\exists y \notin \bar{x}) \theta_q(\bar{x}, y)$. Note that $\theta_q(\bar{x}, y)$ implies $v_q(\bar{x})$ and $v_q(\bar{x})$ implies $\theta_p(\bar{x})$. So $\mu(v_q(\bar{x})) \geq \mu(q) > 0$. Then we must have $v_q(\bar{x}) \in p^*$, otherwise $\mu(p^*) \leq \mu(p) - \mu(v_q(\bar{x}))$. Finally, we conclude that for any tuple \bar{a} , $\llbracket v_q(\bar{a}) \rrbracket$ has full measure in $\llbracket \theta_p(\bar{a}) \rrbracket$, since $\mu(p) = \mu(p^*) \leq \mu(v_q(\bar{x})) \leq \mu(p)$.

As before, a countable intersection (of the sets $\llbracket v_q(\bar{a}) \rrbracket$ for $q \in E_p$, this time) of subsets with full measure has full measure, so

$$\mu \left(\left[\left[\theta_p(\bar{a}) \rightarrow \bigwedge_{q \in E_p} \exists (y \notin \bar{a}) \theta_q(\bar{a}, y) \right] \right] \right) = 1.$$

Taking another countable intersection over all tuples \bar{a} , we have $\mu \models_{\text{a.s.}} \psi_p$.

Let $T = \text{Th}_{F_\gamma}(\mu) \cup \left\{ \varphi_p, \psi_p \mid p \in \bigcup_{n \in \omega} S_{F_\gamma}^n(\mu) \right\}$ and note that T is countable. Since μ almost surely satisfies T , it suffices to show that any two countable models of T are isomorphic. This is a straightforward back-and-forth argument, using φ_p and ψ_p to extend a partial F_γ -elementary isomorphism defined on a realization of p by one step: φ_p tells us that each one-point extension in one model realizes one of the types in E_p , and ψ_p tells us that every type in E_p is realized in a one-point extension in the other model. To start, the empty tuples in the two models of T satisfy the same type in $S_{F_\gamma}^0$, namely $\text{Th}_{F_\gamma}(\mu)$. \square

Theorem 2.3.7. *Let μ be an ergodic structure. Then μ is properly ergodic if and only if for every countable fragment F of $L_{\omega_1, \omega}$, there is a countable fragment $F' \supseteq F$ and a formula $\chi(\bar{x})$ in F' such that $\mu(\chi(\bar{x})) > 0$, but $\mu(p) = 0$ for every F' -type $p(\bar{x})$ containing $\chi(\bar{x})$.*

Proof. Suppose μ is properly ergodic. By Lemma 2.3.5, μ stabilizes at some γ , and by Lemma 2.3.6, there is some n such that $\sum_{p \in S_{F_\gamma}^n(\mu)} \mu(p) < 1$. Let $\chi(\bar{x})$ be the formula $\bigwedge_{p \in S_{F_\gamma}^n(\mu)} \neg \theta_p(\bar{x})$. Then $\mu(\chi(\bar{x})) > 0$.

Let F' be the countable fragment generated by $F \cup F_\gamma \cup \{\chi(\bar{x})\}$, and suppose that $p(\bar{x})$ is an F' -type containing $\chi(\bar{x})$. Let $q = p \cap F_\gamma$. Since q is an F_γ -type that is consistent with $\chi(\bar{x})$, $q \notin S_{F_\gamma}^n(\mu)$, so $\mu(p) \leq \mu(q) = 0$.

Conversely, suppose we have such a fragment F' and a formula $\chi(\bar{x})$. Since $\mu(\chi(\bar{x})) > 0$, $\mu \models_{\text{a.s.}} \exists \bar{x} \chi(\bar{x})$. Let M be a countable structure. If M contains no tuple satisfying $\chi(\bar{x})$, then μ assigns measure 0 to the isomorphism class of M . On the other hand, if $M \models \chi(\bar{a})$, then since μ assigns measure 0 to the set of structures realizing $\text{tp}_{F'}(\bar{a})$, μ also assigns measure 0 to the isomorphism class of M . So μ is properly ergodic. \square

By countable additivity, if a sentence φ of $L_{\omega_1, \omega}$ has only countably many countable models up to isomorphism, then any ergodic structure μ that almost surely satisfies φ is almost surely isomorphic to one of its models. That is, no ergodic model of φ is properly ergodic. We show now that the same is true if φ is a counterexample to Vaught's conjecture, i.e., a sentence with uncountably many, but fewer than continuum-many, countable models.

Corollary 2.3.8 (“Vaught's Conjecture for ergodic structures”). *Let φ be a sentence of $L_{\omega_1, \omega}$. If there is a properly ergodic structure μ such that $\mu \models_{\text{a.s.}} \varphi$, then φ has continuum-many countable models up to isomorphism.*

Proof. This is a consequence of Theorem 2.3.7 and an observation due to Morley [77]: for any countable fragment F of $L_{\omega_1, \omega}$ containing φ and any $n \in \omega$, the set $S_F^n(\varphi)$ of F -types consistent with φ is an analytic subset of 2^F . Since analytic sets have the Perfect Set Property, if $|S_F^n(\varphi)| > \aleph_0$, then $|S_F^n(\varphi)| = 2^{\aleph_0}$. And since a countable structure realizes only countably many n -types, if $|S_F^n(\varphi)| = 2^{\aleph_0}$, then φ has continuum-many countable models up to isomorphism.

Now let μ be the given properly ergodic structure, let F be a countable fragment containing φ , and let F' and $\chi(\bar{x})$ be as in Theorem 2.3.7, let $n = l(\bar{x})$ be the length of the tuple \bar{x} , and suppose for a contradiction that $|S_{F'}^n(\varphi)| \leq \aleph_0$. Let $U_\chi = \{p \in S_{F'}^n(\varphi) \mid \chi \in p\}$. Then U_χ is countable, and, by our choice of $\chi(\bar{x})$, we have $\mu(p) = 0$ for all $p \in U_\chi$. Since $\mu(\llbracket \varphi \rrbracket) = 1$, for any non-redundant tuple \bar{a} , we have

$$0 < \mu(\llbracket \chi(\bar{a}) \rrbracket) = \mu(\llbracket (\varphi \wedge \chi)(\bar{a}) \rrbracket) = \mu\left(\bigcup_{p \in U_\chi} \llbracket \theta_p(\bar{a}) \rrbracket\right) = \sum_{p \in U_\chi} \mu(p),$$

which is a contradiction, by countable additivity of μ . \square

Kechris has observed (in private communication) that Corollary 2.3.8 also follows from a result in descriptive set theory [55, Exercise I.17.14]: an analog for measure of a result of Kuratowski about category [67].

Recall that $\text{Th}(\mu)$ is the complete $L_{\omega_1, \omega}$ -theory of μ . As noted in Remark 1.4.11, μ is properly ergodic if and only if $\text{Th}(\mu)$ has no countable models. In fact, if μ is properly ergodic, then $\text{Th}(\mu)$ has no models at all. This is stronger, since the Löwenheim–Skolem theorem fails for complete theories of $L_{\omega_1, \omega}$.

Corollary 2.3.9. *If μ is properly ergodic, then $\text{Th}(\mu)$ has no models (of any cardinality). However, for any countable fragment F of $L_{\omega_1, \omega}$, $\text{Th}_F(\mu)$ has continuum-many countable models up to isomorphism.*

Proof. Starting with any countable fragment F , let F' and $\chi(\bar{x})$ be as in Theorem 2.3.7. Then $\mu(\chi(\bar{x})) > 0$, so $\exists \bar{x} \chi(\bar{x}) \in \text{Th}(\mu)$. Now if $\text{Th}(\mu)$ has a model M , then there is some tuple \bar{a} from M satisfying χ . Let p be the F' -type of \bar{a} . Since p contains $\chi(\bar{x})$, we have $\mu(p) = 0$, and so $\neg \exists \bar{x} \theta_p(\bar{x}) \in \text{Th}(\mu)$, contradiction.

The rest follows immediately from Corollary 2.3.8, taking $\varphi = \bigwedge_{\psi \in \text{Th}_F(\mu)} \psi$. □

2.4 Rooted models

The Morley–Scott analysis in Section 2.3 shows that proper ergodicity of μ can always be explained by a positive-measure formula $\chi(\bar{x})$ such that any type containing $\chi(\bar{x})$ has measure 0. In a countable structure sampled from μ , each of these types of measure 0 will be realized “rarely”. Sometimes “rarely” means “at most once”, as in Examples 1.4.18 and 1.5.9. But in Example 2.1.2 we saw that a type p of measure 0 can be realized by infinitely many tuples, all of which share a common element $i \in \omega$. This is explained by the fact that p has positive measure after conditioning on a random choice “living at” some finite set containing i . In this section, we will use the AHK representation (see Section 1.3) to show that this behavior, which we call rootedness, is typical.

Throughout this section, let F be a countable fragment of $L_{\omega_1, \omega}$, and let T be an F -theory. Recall that $S_F^n(T)$ is the space of F -types in the variables x_0, \dots, x_{n-1} which are consistent with T , and $S_F^n(T)$ is the subspace of non-redundant types. Similarly, $S_{\text{qf}}^n(T)$ and $S_{\text{qf}}^n(T)$ denote the spaces first-order quantifier-free and non-redundant first-order quantifier-free types in the variables x_0, \dots, x_{n-1} which are consistent with T .

Definition 2.4.1. Let $p \in S_F^n(T)$ be a type realized in $M \models T$. An element $a \in M$ is called a **root** of p in M if a is an element of every tuple realizing p in M . We use the same terminology for quantifier-free types in $S_{\text{qf}}^n(T)$.

Remark 2.4.2. If a type p has a unique realization in M , then p has a root in M (take any element of the unique tuple realizing p). When $n = 1$, the converse is true: a realized type $p(x) \in S_F^1(T)$ (or $S_{\text{qf}}^1(T)$) has a root in M if and only if it has a unique realization in M .

Definition 2.4.3. Let $\chi(\bar{x})$ be a formula in F such that $\chi(\bar{x}) \wedge (\bigwedge_{i \neq j} x_i \neq x_j)$ is consistent with T . Then a model $M \models T$ is **χ -rooted** if every type $p(\bar{x}) \in S_F^n(T)$ which contains $\chi(\bar{x})$ and is realized in M has a root in M .

Remark 2.4.4. We note that the property of χ -rootedness is expressible by a sentence of $L_{\omega_1, \omega}$, which asserts that for every non-redundant tuple \bar{a} satisfying $\chi(\bar{x})$, there is some

element a_i of the tuple such that every other tuple \bar{b} with the same F -type as \bar{a} contains a_i :

$$\forall \bar{x} \left(\chi(\bar{x}) \wedge \bigwedge_{i \neq j} x_i \neq x_j \right) \rightarrow \left(\bigvee_{i=1}^n \forall \bar{y} \left(\left(\bigwedge_{\varphi \in F} \varphi(\bar{x}) \leftrightarrow \varphi(\bar{y}) \right) \rightarrow \left(\bigvee_{j=1}^n x_i = y_j \right) \right) \right),$$

and hence the set of χ -rooted models of T is a Borel set in Str_L .

Let M be a χ -rooted model of T . Suppose that $p(\bar{x}) \in S_F^n(T)$ contains $\chi(\bar{x})$ and is realized in M , and let a be a root of p in M . If $M \models p(a, \bar{b})$, then a is the unique element of M satisfying $p(x, \bar{b})$, since if $c \neq a$, then a is not in $c\bar{b}$, so $c\bar{b}$ does not realize p . This implies that M has non-trivial group-theoretic definable closure, since every automorphism of M fixing \bar{b} also fixes a . Note that T may still have trivial definable closure, since p is an F -type and, in general, is not equivalent to a formula in F .

We can conclude, however, that if an F -theory T with trivial definable closure has a χ -rooted model, then no non-redundant type that contains $\chi(\bar{x})$ is isolated. Thus isolated types are not dense in $S_F^n(T)$. By the standard model-theoretic arguments, which go through in countable fragments of $L_{\omega_1, \omega}$ (see [57]), this implies that T does not have a prime model with respect to F -elementary embeddings, and that there are continuum-many types in $S_F^n(T)$ containing $\chi(\bar{x})$.

The next theorem says that this situation is typical of properly ergodic structures μ : the theory $\text{Th}_F(\mu)$ has many χ -rooted models for some $\chi(\bar{x})$ and (by Theorem 1.5.3) trivial dcl. Of course, given any non-redundant F -type p containing $\chi(\bar{x})$, we can try to bring χ -rootedness into direct conflict with trivial dcl by moving to a larger countable fragment F' which contains the formula $\theta_p(\bar{x}) = \bigwedge_{\varphi \in p} \varphi(\bar{x})$ isolating p . But, as we will see, every non-redundant type containing the formula $\chi(\bar{x})$ has measure 0, so $\text{Th}_{F'}(\mu)$ contains the sentence $\forall \bar{x} \neg \theta_p(\bar{x})$, ruling out troublesome realizations of p .

It is also worth noting that the countable fragment F' can only isolate and rule out countably many of the continuum-many types of measure 0 containing $\chi(\bar{x})$. In the example of the kaleidoscope random graph (Example 1.4.18), we could extend from the first-order fragment to a countable fragment F of $L_{\omega_1, \omega}$ containing some of the conjunctions $\bigwedge_{n \in A} xR_ny \wedge \bigwedge_{n \notin A} \neg xR_ny$, for $A \subseteq \omega$. Then the theory $\text{Th}_F(\mu)$ is essentially the same as $\text{Th}_{\text{FO}}(\mu)$, but with countably many of the continuum-many quantifier-free 2-types forbidden.

Theorem 2.4.5. *Let μ be a properly ergodic structure, F a countable fragment of $L_{\omega_1, \omega}$, and $\chi(\bar{x})$ a formula in F such that $\mu(\chi(\bar{x})) > 0$. Suppose that $\mu(p) = 0$ for every F -type $p(\bar{x})$ containing $\chi(\bar{x})$. Then μ assigns measure 1 to the set of χ -rooted models of $\text{Th}_F(\mu)$.*

Proof. Let L' , T' , and Q be the language, pithy Π_2 theory, and countable set of partial quantifier-free types obtained from Theorem 2.2.2. By Corollary 2.2.3, μ corresponds to a properly ergodic L' -structure μ' , which gives measure 1 to the set of models of T' which omit all the types in Q .

For such models, each formula $\varphi(\bar{x})$ in F is equivalent to the atomic L' -formula $R_\varphi(\bar{x})$, so we have $\mu'(R_\chi(\bar{x})) > 0$, and for every quantifier-free type $q(\bar{x})$ containing $R_\chi(\bar{x})$, $\mu'(q) = 0$. It

suffices to show that for an L' -structure M sampled from μ' , almost surely every quantifier-free type $p(\bar{x}) \in S_{\text{qf}}^n(T')$ that contains $R_\chi(\bar{x})$ and is realized in M has a root in M .

By Theorem 1.3.1, μ' has an AHK representation $(f_m)_{m \in \omega}$. We adopt the notation of Section 1.3 for the random variables $(\xi_A)_{A \in \mathcal{P}_{\text{fin}}(\omega)}$. Since μ' is ergodic, by Theorem 1.3.2, we can pick the functions f_m so they do not depend on the argument indexed by \emptyset .

Let $p \in S_{\text{qf}}^n(T')$. Given a subset $X \subseteq [n]$, we can separate the inputs to f_n , $(x_A)_{A \in \mathcal{P}([n])}$, into those contained in X and those not contained in X : $((x_A)_{A \subseteq X}, (x_A)_{A \not\subseteq X})$. Then if we fix values $(r_A)_{A \subseteq X}$ from $[0, 1]$, we say p is **likely given $(r_A)_{A \subseteq X}$ in position** if there is a positive probability that $f_n((r_A)_{A \subseteq X}, (\zeta_A)_{A \not\subseteq X}) = p$, when $(\zeta_A)_{A \not\subseteq X}$ are uniform i.i.d. random variables with values in $[0, 1]$.

Now for any $0 \leq m \leq n$, and any values $(r_A)_{A \in \mathcal{P}([m])}$, we say p is **likely given $(r_A)_{A \in \mathcal{P}([m])}$ in position**. We make some observations about this definition:

- In terms of the random variables $(\xi_A)_{A \in \mathcal{P}_{\text{fin}}(\omega)}$ that we use to sample from the AHK representation of μ' , for a tuple \bar{b} of length at most n , p is likely given $\widehat{\xi}_{\bar{b}}$ (see Section 1.3 for this notation) if, conditioned on the values of $\widehat{\xi}_{\bar{b}}$, there is a positive probability that p is realized on some tuple containing all the elements of \bar{b} .
- Note that this does not depend on the order in which the set $B = \|\bar{b}\|$ is enumerated as a tuple. Abusing notation, we say p is likely given $\widehat{\xi}_B$ if p is likely given $\widehat{\xi}_{\bar{b}}$ for some tuple \bar{b} enumerating B .
- If $|B| = n$, then all the inputs of f_n are fixed, and p is likely given $\widehat{\xi}_B$ if and only if p is realized on B (meaning that some enumeration of B as a tuple realizes p).
- If $B = \emptyset$, then, since f_n does not depend on the argument indexed by \emptyset , none of the relevant inputs of f_n are fixed, and p is likely given $\widehat{\xi}_B$ if and only if $\mu'(p) > 0$.

Claim 1: Almost surely, for every $B \in \mathcal{P}_{\text{fin}}(\omega)$ with $|B| \leq n$, and for every $A \in \mathcal{P}_{\text{fin}}(\omega)$, if p is likely given $\widehat{\xi}_B$, then there is an extension $B \subseteq C$ with $|C| = n$ and $C \cap A = B \cap A$ such that p is realized on C .

This is clear if $|B| = n$, taking $C = B$. If not, let $(C_i)_{i \in \omega}$ be extensions of B of size n , such that $C_i \cap C_j = B$ when $i \neq j$ and such that $C_i \cap A = B \cap A$ for all i . Let \mathcal{A}_i be the event that p is realized on C_i . The \mathcal{A}_i are conditionally independent over $\widehat{\xi}_B$, and each has the same positive probability, so almost surely infinitely many occur.

Claim 2: Almost surely, for every B in $\mathcal{P}_{\text{fin}}(\omega)$ with $|B| \leq n$, and for every pair of extensions $B \subseteq C$ and $B \subseteq D$ with $C \cap D = B$ and $|C| = |D| = n$, if p is realized on both C and D , then p is likely given $\widehat{\xi}_B$.

Again, this is clear if $|B| = n$, since then $C = D = B$. If not, then since the quantifier-free types realized on C and D are conditionally independent over $\widehat{\xi}_B$, the probability that the

same probability 0 event (the realization of any type p which is not likely given $\widehat{\xi}_B$) happens on both C and D is 0.

Claim 3: Almost surely, for all A and B in $\mathcal{P}_{\text{fin}}(\omega)$ with $|A| \leq n$ and $|B| \leq n$, if p is likely given $\widehat{\xi}_A$ and likely given $\widehat{\xi}_B$, then p is likely given $\widehat{\xi}_{A \cap B}$.

This follows from the last two claims. Suppose p is likely given $\widehat{\xi}_A$ and likely given $\widehat{\xi}_B$. By Claim 1, almost surely, there is an extension $A \subseteq A'$ with $|A'| = n$ such that $A' \cap B = A \cap B$ and p is realized on A' . By Claim 1 again, almost surely, there is an extension $B \subseteq B'$ with $|B'| = n$ such that $B' \cap A' = B \cap A'$ and p is realized on B' . But then $A' \cap B' = A \cap B$, so by Claim 2, almost surely p is likely given $\widehat{\xi}_{A \cap B}$.

Let M be the random structure obtained by sampling from the AHK representation. For any quantifier-free type p containing $R_\chi(\bar{x})$ which is realized in M , p is likely given $\widehat{\xi}_{\bar{a}}$ for all \bar{a} realizing p , and by Claim 3, almost surely the sets $\{\|\bar{a}\| \mid p \text{ is likely given } \widehat{\xi}_{\bar{a}}\}$ are closed under intersection. But since $\mu'(p) = 0$, p is not likely given ξ_\emptyset . Hence the intersection of all the realizations of p is almost surely non-empty, i.e., p has a root in M . \square

2.5 Constructing properly ergodic structures

In this section, we will use a single χ -rooted model M of an F -theory T to construct a properly ergodic structure which almost surely satisfies T . The strategy is to build a Borel structure \mathbb{M} as an inverse limit of finite structures. We use M as a guide to ensure that \mathbb{M} is also χ -rooted. A coherent system of probability measures on the finite structures gives rise to a probability measure ν on \mathbb{M} . Then our ergodic structure μ is obtained by i.i.d. sampling of countably many points from \mathbb{M} according to ν and taking the induced substructure.

Definition 2.5.1. A **Borel structure** \mathbb{M} is an L -structure whose domain is a standard Borel space such that for every relation symbol R of arity $\text{ar}(R)$ in L , the subset $R \subseteq M^{\text{ar}(R)}$ is Borel. A **measured structure** is a Borel structure \mathbb{M} equipped with an atomless probability measure ν .

Given a measured structure (\mathbb{M}, ν) , there is a canonical measure $\mu_{\mathbb{M}, \nu}$ on Str_L , obtained by sampling a countable ν -i.i.d. sequence (of almost surely distinct points) from \mathbb{M} and taking the induced substructure. Somewhat more formally, $\mu_{\mathbb{M}, \nu}$ is the distribution of a random structure in Str_L whose atomic diagram on ω is given by that of the random substructure of \mathbb{M} with underlying set $(v_i)_{i \in \omega}$, where $(v_i)_{i \in \omega}$ is a ν -i.i.d. sequence of elements in \mathbb{M} .

We can also view the measure $\mu_{\mathbb{M}, \nu}$ as the result of sampling from an AHK system. Choose a measure-preserving Borel isomorphism h from $[0, 1]$ equipped with the uniform measure to the domain of \mathbb{M} equipped with ν , and define functions $f_n: [0, 1]^{\mathcal{P}_{\text{fin}}([n])} \rightarrow S_{\text{qf}}^n(L)$ by

$$f_n((\xi_A)_{A \subseteq [n]}) = \text{qftp}(h(\xi_{\{0\}}), \dots, h(\xi_{\{n-1\}})).$$

Informally, these functions ignore the random variables ξ_A when $|A| \neq 1$ and view the $(\xi_{\{a\}})_{a \in \omega}$ as independent random variables with distribution ν taking their values in \mathbb{M} . To be totally formal, we should also choose a dummy output for each f_n on the measure 0 set where $\xi_{\{a\}} = \xi_{\{b\}}$ for some $a \neq b$.

Now $(f_n)_{n \in \omega}$ is an AHK system, so it induces an invariant measure on Str_L . This measure is clearly the same as $\mu_{\mathbb{M}, \nu}$ described above via sampling of a random substructure. Since the sampling is done independently, $\mu_{\mathbb{M}, \nu}$ satisfies disjoint-independence. Equivalently, since the f_n do not depend on the argument indexed by \emptyset , the measure $\mu_{\mathbb{M}, \nu}$ is ergodic (Theorem 1.3.2). In any case, we have established the following lemma.

Lemma 2.5.2. *Given a measured structure (\mathbb{M}, ν) , the measure $\mu_{\mathbb{M}, \nu}$ on Str_L is an ergodic structure.*

In fact, the ergodic structure $\mu_{\mathbb{M}, \nu}$ is “random-free”. This terminology comes from the world of graphons: a graphon is said to be **random-free** when it is $\{0, 1\}$ -valued almost everywhere [49, Section 10]. This can be thought of as “having randomness” only at the level of vertices (and not at higher levels, i.e. at the level of edges). See also 0–1 valued graphons in [70] and the simple arrays of [52]. A graphon is random-free if and only if the corresponding ergodic structure is random-free in the following sense.

Definition 2.5.3. An AHK system $(f_n)_{n \in \omega}$ is **random-free** if each function f_n depends only on the singleton variables $\xi_{\{a\}}$ for $a \in \omega$. An ergodic structure μ is **random-free** if it has a random-free AHK representation.

We would like to transfer properties of \mathbb{M} to almost-sure properties of $\mu_{\mathbb{M}, \nu}$. It is not true in general that $\mu_{\mathbb{M}, \nu} \models_{\text{a.s.}} \text{Th}_F(\mathbb{M})$. But the following property will allow us to transfer satisfaction in \mathbb{M} to satisfaction in $\mu_{\mathbb{M}, \nu}$ for pithy Π_2 sentences. This will be sufficient, using the trick of Section 2.2.

Definition 2.5.4. Let T be a pithy Π_2 theory, and let (\mathbb{M}, ν) be a measured structure. We say that (\mathbb{M}, ν) **satisfies T with strong witnesses** if the following hold for all $\varphi \in T$:

- If φ is universal, then $\mathbb{M} \models \varphi$.
- If $\varphi = \forall \bar{x} \exists y \rho(\bar{x}, y)$, then for every tuple \bar{a} from \mathbb{M} , the set $\rho(\bar{a}, \mathbb{M}) = \{b \in \mathbb{M} \mid \rho(\bar{a}, b)\}$ either contains an element of the tuple \bar{a} or has positive ν -measure.

Note that if (\mathbb{M}, ν) satisfies T with strong witnesses, then $\mathbb{M} \models T$.

Lemma 2.5.5. *Let (\mathbb{M}, ν) be a measured structure, let $\mu = \mu_{\mathbb{M}, \nu}$.*

- (i) *Let Q be a countable set of partial quantifier-free types. If \mathbb{M} omits all the types in Q , then almost surely μ omits all the types in Q .*
- (ii) *Let T be a pithy Π_2 theory. If (\mathbb{M}, ν) satisfies T with strong witnesses, then μ almost surely satisfies T .*

(iii) Further, if there is a quantifier-free formula $\chi(\bar{x})$ such that \mathbb{M} is χ -rooted with respect to quantifier-free types, then μ is properly ergodic.

Proof. (i): If no tuple from \mathbb{M} realizes a quantifier-free type $q \in Q$, then no tuple from any countable substructure sampled from \mathbb{M} realizes q .

(ii): Every universal sentence $\forall \bar{x} \psi(\bar{x})$ in T is almost surely satisfied by μ , since every tuple \bar{v} from \mathbb{M} satisfies the quantifier-free formula $\psi(\bar{x})$.

So it suffices to consider sentences of the form $\forall \bar{x} \exists y \rho(\bar{x}, y)$. Fix a tuple \bar{a} from ω . Corresponding to \bar{a} , we have the tuple of random elements $\bar{v} = (v_{a_1}, \dots, v_{a_n})$ sampled from \mathbb{M} . By Fubini's theorem, it suffices to show that for each possible value of this random tuple, there is almost surely some $b \in \omega$ such that $\mathbb{M} \models \rho(\bar{v}, v_b)$.

By strong witnesses, $\rho(\bar{v}, \mathbb{M})$ either contains an element v_{a_i} of the tuple \bar{v} or has positive measure. In the first case, v_{a_i} serves as our witness. In the second case, since there are infinitely many other independent random elements $(v_b)_{b \in \omega \setminus \|\bar{a}\|}$, almost surely infinitely many of them land in the positive-measure set $\rho(\bar{v}, \mathbb{M})$.

(iii): By (ii), $\mu \models_{\text{a.s.}} T$, and since $\chi(\bar{x}) \wedge (\bigwedge_{i \neq j} x_i \neq x_j)$ is consistent with T , $\mu(\chi(\bar{x})) > 0$. Let p be any type containing $\chi(\bar{x})$, and let q be its restriction to the first-order quantifier-free formulas. To show that $\mu(p) = 0$, it suffices to show that $\mu(q) = 0$.

Now since \mathbb{M} is χ -rooted with respect to quantifier-free types, q has a root a in \mathbb{M} . The probability that a tuple sampled from \mathbb{M} satisfies q is bounded above by the probability that the tuple contains a . This is 0, since ν is atomless. Hence, by Theorem 2.3.7, μ is properly ergodic. \square

Thus, after applying the pithy Π_2 transformation from Section 2.2 to an F -theory T , we have reduced the problem of constructing a properly ergodic structure almost surely satisfying T to that of constructing a measured structure with the properties in Lemma 2.5.5.

Theorem 2.5.6. *Let F be a countable fragment of $L_{\omega_1, \omega}$, T a complete F -theory with trivial dcl, $\chi(\bar{x})$ a formula in F , and M a χ -rooted model of T . Then there is a properly ergodic structure μ such that $\mu \models_{\text{a.s.}} T$.*

Proof. We begin by applying Theorem 2.2.2 to obtain a language $L' \supseteq L$, a pithy Π_2 theory T' , and a countable set of partial quantifier-free types Q . Let M' be the natural expansion of M to an L' -structure. Then M' is R_χ -rooted, where $R_\chi(\bar{x})$ is the atomic L' -formula corresponding to the L -formula $\chi(\bar{x})$. By Corollary 2.2.3, it suffices to construct a properly ergodic L' -structure which almost surely satisfies T' and omits the types in Q .

Part 1: The inverse system

We construct a sequence $(A_k)_{k \in \omega}$ of finite L' -structures, each of which is identified with a substructure of M' . Given a structure A , we define the structure A^* to have underlying set

$A \cup \{*\}$, where no new relations hold involving $*$. For each k , we equip the underlying set of A_k^* with a discrete probability measure ν_k that assigns positive measure to every element, and we fix a finite sublanguage L_k of L' .

Finally, we define surjective connecting maps $f_k: A_{k+1}^* \rightarrow A_k^*$ such that $f_k(*) = *$ for all k , which preserve the measures and certain quantifier-free types:

- (1) $\nu_{k+1}(f_k^{-1}(X)) = \nu_k(X)$ for all $X \subseteq A_k^*$.
- (2) If \bar{a} is a non-redundant tuple from A_{k+1} such that $f_k(\bar{a})$ is a non-redundant tuple from A_k , then $\text{qftp}_{L_k}(\bar{a}) = \text{qftp}_{L_k}(f_k(\bar{a}))$. Note that we make no requirement if f_k is not injective on \bar{a} or if any element of \bar{a} is mapped to $*$.

We call a pithy Π_2 sentence in T' an extension axiom if it is not universal. We enumerate the extension axioms in T' as $\langle \varphi_k \rangle_{k \in \omega}$ and the types in Q as $\langle q_i \rangle_{k \in \omega}$ with redundancies, so that each axiom and each type appears infinitely often in its list. We also enumerate the symbols in the language L' as $\langle R_k \rangle_{k \in \omega}$, without redundancies.

At stage 0, we set $A_0 = \emptyset$, with its induced structure as the empty substructure of M' . Then $A_0^* = \{*\}$, and we set $\nu_0(\{*\}) = 1$ and $L_0 = \emptyset$.

At stage $k + 1$, we are given A_k , ν_k , and L_k . We define A_{k+1} , ν_{k+1} , L_{k+1} , and the connecting map f_k in four steps.

Step 1: Splitting the elements of A_k .

Enumerate the elements of A_k as $\{a_1, \dots, a_m\}$. We build intermediate substructures of M' , $B_i = \{a_1, \dots, a_m, a'_1, \dots, a'_i\}$ for $0 \leq i \leq m$, where each new element a'_j is a ‘‘copy’’ of a_j to be defined. We start with $B_0 = A_k$.

Given B_i , let $\varphi_{B_i}(x_1, \dots, x_m, x'_1, \dots, x'_i)$ be the conjunction of all atomic and negated atomic L_k -formulas holding on B_i , so that φ encodes the quantifier-free L_k -type of B_i . Now there is an L -formula ψ_{B_i} in F such that ψ_{B_i} has the same realizations as φ_{B_i} in M' . Since $T = \text{Th}_F(M)$ has trivial dcl, we can find another realization $a'_{i+1} \neq a_{i+1}$ of $\psi_{B_i}(a_1, \dots, x_{i+1}, \dots, a_m, a'_1, \dots, a'_i)$ in $M' \setminus B_i$. Set $B_{i+1} = B_i \cup \{a'_{i+1}\}$. We have

$$\text{qftp}_{L_k}(a_1, \dots, a_{i+1}, \dots, a_m, a'_1, \dots, a'_i) = \text{qftp}_{L_k}(a_1, \dots, a'_{i+1}, \dots, a_m, a'_1, \dots, a'_i). \quad (\dagger)$$

At the end of Step 1, we have a structure $B_m = \{a_1, \dots, a_m, a'_1, \dots, a'_m\}$.

Step 2: Splitting $$.*

The current extension axiom φ_k in our enumeration has the form $\forall \bar{x} \exists y \rho(\bar{x}, y)$, where \bar{x} is a tuple of length j and ρ is quantifier-free. Suppose there is a tuple \bar{a} from B_m such that $B_m \models \neg \exists y \rho(\bar{a}, y)$. Then, since $M' \models \exists y \rho(\bar{a}, y)$, we can choose some witness $c_{\bar{a}}$ to the existential quantifier in $M' \setminus B_m$. Let $W = \{c_{\bar{a}} \mid \bar{a} \in B_m^j \text{ and } B_m \models \neg \exists y \rho(\bar{a}, y)\}$ be the (finite) set of chosen witnesses. Note that if \bar{x} is the empty tuple of variables, then W is either empty or consists of a single witness, depending on whether $B_m \models \exists y \rho(y)$.

Let $A_{k+1} = B_m \cup W$ if W is non-empty, and otherwise let $A_{k+1} = B_m \cup \{c\}$, where c is any new element in $M' \setminus B_m$.

Step 3: Defining f_k and ν_{k+1} .

Recall that f_k is to be a map from A_{k+1}^* to A_k^* . We set $f_k(a_i) = f_k(a'_i) = a_i$ and $f_k(c) = f_k(*) = *$ for $c \in A_{k+1} \setminus B_m$.

We define ν_{k+1} by splitting the measure of an element of A_k^* evenly among its preimages under f_k . So $\nu_{k+1}(a_i) = \nu_{k+1}(a'_i) = \frac{1}{2}\nu_k(a_i)$, and $\nu_{k+1}(c) = \nu_{k+1}(*) = \frac{1}{N}\nu_k(*)$, where $N = |A_{k+1}^* \setminus B_m| \geq 2$. Note that every element of A_{k+1}^* has positive measure, by induction.

Step 4: Defining L_{k+1} .

We expand the current language L_k to L_{k+1} by adding finitely many new symbols from L' . Let R_k and q_k be the current symbol in L' and type in Q , respectively, in our enumerations.

- (a) Add R_k to L_{k+1} if it is not already included.
- (b) Since A_{k+1} is a substructure of M' , no tuple from A_{k+1} realizes q_k . That is, for every tuple \bar{a} from A_{k+1} , there is some quantifier-free formula $\varphi_{\bar{a}}(\bar{x}) \in q_k$ such that $M' \models \neg\varphi_{\bar{a}}(\bar{a})$. Add the finitely many relation symbols appearing in $\varphi_{\bar{a}}$ to L_{k+1} .
- (c) Let n be the number of free variables in $\chi(\bar{x})$. For every pair of non-redundant n -tuples \bar{a} and \bar{b} from A_{k+1} that realize distinct quantifier-free L' -types in M' , there is some relation symbol $R_{\bar{a},\bar{b}}$ that separates their types. Add $R_{\bar{a},\bar{b}}$ to L_{k+1} .

This completes stage $k + 1$ of the construction. Let us check that conditions (1) and (2) above are satisfied by the connecting map f_k .

- (1): Since ν_k and ν_{k+1} are discrete measures on finite spaces, it suffices to check that $\nu_k(a) = \sum_{b \in f_k^{-1}(\{a\})} \nu_{k+1}(b)$ for every singleton $a \in A_k^*$. This follows immediately from our definitions of f_k and ν_{k+1} .
- (2): Let \bar{b} be a non-redundant tuple from A_{k+1} . The assumption that $f_k(\bar{b})$ is a non-redundant tuple from A_k means that every element of \bar{b} is in B_m (since the other elements are mapped to $*$) and that a_i and a'_i are not both in \bar{b} for any i . For any function $\gamma: [m] \rightarrow [2]$, let \bar{a}^γ be the m -tuple which contains a_i if $\gamma(i) = 0$ and a'_i if $\gamma(i) = 1$. Then, expanding \bar{b} to an m -tuple of the form \bar{a}^γ , it suffices to show that $\text{qftp}_{L_k}(\bar{a}^\gamma) = \text{qftp}_{L_k}(f_k(\bar{a}^\gamma)) = \text{qftp}_{L_k}(\bar{a})$. This follows by several applications of instances of the equality (†) above.

Part 2: The measured structure

Let \mathbb{X} be the inverse limit of the system of finite sets A_k^* and surjective connecting maps f_k . For each k , let π_k be the projection map $\mathbb{X} \rightarrow A_k^*$. Then \mathbb{X} is a profinite set, so it has a natural topological structure as a Stone space, in which the basic clopen sets are exactly the preimages under the maps π_k of subsets of the sets A_k^* . Note that \mathbb{X} is separable, so it is a standard Borel space.

Let ν^* be the finitely additive measure on the Boolean algebra \mathcal{B}^* of clopen subsets of \mathbb{X} defined by $\nu^*(\pi_k^{-1}(X)) = \nu_k(X)$. This is well-defined by condition (1). By the Hahn–Kolmogorov Measure Extension Theorem [92, Theorem 1.7.8], ν^* extends to a Borel probability measure ν on \mathbb{X} .

Now each element a of A_k^* has at least 2 preimages in A_{k+1}^* , each of which have measure at most $\frac{1}{2}\nu_k(a)$. Hence, by induction, the measure of each element of A_k^* is at most 2^{-k} . So for all $x \in \mathbb{X}$, x is contained in a basic clopen set $X_k = \pi_k^{-1}(\{\pi_k(x)\})$ with $\nu(X_k) \leq 2^{-k}$ for all k . This implies that $\nu(\{x\}) = 0$ and ν is non-atomic.

Note that there is a unique element $*$ of \mathbb{X} with the property that $\pi_k(*) = *$ for all k . We define a Borel L' -structure \mathbb{M} with domain $\mathbb{X} \setminus \{*\}$ (which is also a standard Borel space). Since we have only removed a measure 0 set from \mathbb{X} , ν restricts to a probability measure on \mathbb{M} , which we also call ν .

We define the structure on \mathbb{M} by giving the quantifier-free type of every non-redundant tuple from \mathbb{M} . By Step 4 (a), $\bigcup_{k=0}^{\infty} L_k = L'$. Given a non-redundant tuple \bar{a} from \mathbb{M} and a quantifier-free formula $\varphi(\bar{x})$, we choose k large enough so that L_k contains all of the relation symbols appearing in $\varphi(\bar{x})$ and so that $\pi_k(\bar{a})$ is a non-redundant tuple from A_k . We set $\mathbb{M} \models \varphi(\bar{a})$ if and only if $A_k \models \varphi(\pi_k(\bar{a}))$. This is well-defined by condition (2).

According to this definition, to determine whether a quantifier-free formula $\varphi(\bar{x})$ holds of a redundant tuple \bar{a} , we can remove the redundancies from \bar{a} and replace the corresponding variables in \bar{x} . For example, if $a_i = a_j$, we can remove a_j and replace instances of x_j with x_i . This is equivalent to choosing k large enough so that distinct elements of \bar{a} are mapped by π_k to distinct elements of A_k and checking whether $A_k \models \varphi(\pi_k(\bar{a}))$.

The interpretation of a relation symbol R is a Borel subset of $\mathbb{M}^{\text{ar}(R)}$. Indeed, fixing k , the set of tuples \bar{a} such that distinct elements of \bar{a} are mapped by π_k to distinct elements of A_k and $\pi_k(\bar{a})$ satisfies R is closed (the finite union of certain boxes intersected with certain diagonals), and the interpretation of R is the countable union (over k) of these sets. Hence \mathbb{M} is a Borel structure.

We now verify the conditions of Lemma 2.5.5 for the measured structure (\mathbb{M}, ν) , the pithy Π_2 theory T' , the quantifier-free types Q , and the quantifier-free formula $R_\chi(\bar{x})$.

(i): \mathbb{M} omits all the types in Q .

Let $q(\bar{x})$ be a type in Q , and let \bar{a} be a tuple from \mathbb{M} . Let k be large enough so that π_k maps distinct elements of \bar{a} to distinct elements of A_k . Since q appears infinitely many times in our enumeration of Q , there is some $l > k$ such that $q = q_l$. Then π_{l+1} maps distinct elements of \bar{a} to distinct elements of A_{l+1} . Let $\bar{b} = \pi_{l+1}(\bar{a})$. In Step 4 (b) of stage $l + 1$ of the construction, we ensured that L_{l+1} includes the relation symbols appearing in a quantifier-free formula $\varphi_{\bar{b}}(\bar{x}) \in q_k$ such that $A_{l+1} \models \neg\varphi_{\bar{b}}(\bar{b})$. Then also $\mathbb{M} \models \neg\varphi_{\bar{b}}(\bar{a})$, and hence \bar{a} does not realize q .

(ii): (\mathbb{M}, ν) satisfies T' with strong witnesses.

Let φ be an axiom of T' . Then φ has the form $\forall \bar{x} \psi(\bar{x})$, where ψ is quantifier-free or ψ has a single existential quantifier. Let \bar{a} be a tuple from \mathbb{M} . Let k be large enough

so that all the symbols in φ are in L_k and π_k maps distinct elements of \bar{a} to distinct elements of A_k .

If ψ is quantifier-free, then $\mathbb{M} \models \psi(\bar{a})$ if and only if $A_k \models \psi(\pi_k(\bar{a}))$. The latter holds, since A_k is a substructure of M' , and $M' \models \varphi$.

If ψ is existential, then it has the form $\exists y \rho(\bar{x}, y)$, and since φ appears infinitely many times in our enumeration of the extension axioms in T' , there is some $l > k$ such that $\varphi = \varphi_l$. Then π_l maps distinct elements of \bar{a} to distinct elements of A_l . Let $\bar{b} = \pi_{l+1}(\bar{a})$. In Step 2 of stage $l + 1$ of the construction, we ensured that there was some witness $c_{\bar{b}}$ such that $A_{l+1} \models \rho(\bar{b}, c_{\bar{b}})$. If $c_{\bar{b}}$ is not an element of the tuple \bar{b} , then for any $c \in \mathbb{M}$ such that $\pi_{l+1}(c) = c_{\bar{b}}$, we have $\mathbb{M} \models \rho(\bar{a}, c)$. Since $\nu(\pi_{l+1}^{-1}(\{c\})) = \nu_{l+1}(c) > 0$, the set of witnesses for $\exists y \rho(\bar{a}, y)$ has positive ν -measure. On the other hand, if $c_{\bar{b}}$ is an element of the tuple \bar{b} , say b_i , then $\mathbb{M} \models \rho(\bar{a}, a_i)$.

(iii): \mathbb{M} is R_χ -rooted with respect to quantifier-free types.

To show that \mathbb{M} is R_χ -rooted, we need to show that every non-redundant quantifier-free n -type containing R_χ that is realized in \mathbb{M} has a root in \mathbb{M} . Suppose not. Then there is a quantifier-free type $p(\bar{x})$ and a family of tuples $(\bar{a}^i)_{i \in I}$ from \mathbb{M} such that each \bar{a}^i realizes p , but there is no element a which is in every \bar{a}^i . Note that if such a family exists, then we can find one containing only finitely many tuples: picking some \bar{a} in the family, for each element a_j in \bar{a} there is another tuple in the family which does not contain a_j , so $(n + 1)$ tuples suffice.

Let $(\bar{a}^1, \dots, \bar{a}^m)$ be our finite family of tuples. Let k be large enough so that $R_\chi \in L_k$ and π_k maps all the elements of $(\bar{a}^1, \dots, \bar{a}^m)$ to distinct elements of A_k . For all i , let $\bar{b}^i = \pi_k(\bar{a}^i)$. Then all of the tuples \bar{b}^i realize the same quantifier-free L_k -type $p' = p \upharpoonright L_k$ in A_k , and p' contains R_χ . By Step 4 (c) of stage k of our construction, the tuples \bar{b}^i must actually realize the same quantifier-free L' type $q \supseteq p'$ in M' (which may be distinct from p). But there is no element which appears in all of these tuples, contradicting the fact that M' is R_χ -rooted.

Let $\mu = \mu_{\mathbb{M}, \nu}$. By Lemma 2.5.2 and Lemma 2.5.5, μ is a properly ergodic structure that almost surely satisfies T' and omits the types in Q . \square

Theorem 2.5.6, along with the results of the previous sections, gives a characterization of those theories which admit properly ergodic models.

Theorem 2.5.7. *Suppose Σ is a set of sentences in some countable fragment F of $L_{\omega_1, \omega}$. The following are equivalent:*

- (1) *There is a properly ergodic structure μ such that $\mu \models_{\text{a.s.}} \Sigma$.*
- (2) *There is a countable fragment $F' \supseteq F$ of $L_{\omega_1, \omega}$, a complete F' -theory $T \supseteq \Sigma$ with trivial dcl, a formula $\chi(\bar{x})$ in F' , and a model $M \models T$ which is χ -rooted.*

Proof. (1) \rightarrow (2): Theorem 2.3.7 gives us a countable fragment $F' \supseteq F$, and a formula $\chi(\bar{x})$ in F' such that $\mu(\chi) > 0$, but for every F' -type $p(\bar{x})$ containing χ , $\mu(p) = 0$. Let $T = \text{Th}_F(\mu)$. Then $\Sigma \subseteq T$, and T has trivial dcl by Theorem 1.5.3. Now by Theorem 2.4.5, the set of χ -rooted models of T has measure 1. In particular, it is non-empty.

(2) \rightarrow (1): By Theorem 2.5.6, there is a properly ergodic structure μ such that $\mu \models_{\text{a.s.}} T$, and $\Sigma \subseteq T$. \square

Remark 2.5.8. The conditions in Theorem 2.5.7 (2) can sometimes be satisfied with $F' = F$. In fact, for many of the examples in Section 2.1, we could take F' to be first-order logic. However, Example 2.1.4 shows that, in general, the move to a larger fragment of $L_{\omega_1, \omega}$ is necessary.

The following corollaries, which may be of interest independently of Theorem 2.5.7, follow immediately from its proof in the case that μ is properly ergodic and from the analogous construction in [2] in the case that μ is almost surely isomorphic to a countable structure.

Corollary 2.5.9. *If μ is an ergodic structure, then for any countable fragment F of $L_{\omega_1, \omega}$, the theory $\text{Th}_F(\mu)$ has a Borel model (of cardinality 2^{\aleph_0}).*

Corollary 2.5.10. *For every countable fragment F of $L_{\omega_1, \omega}$, every ergodic structure μ is F -elementarily equivalent to a random-free ergodic structure μ' . That is, there exists a random-free ergodic structure μ' such that $\text{Th}_F(\mu) = \text{Th}_F(\mu')$.*

Part II
Category

Chapter 3

Strong embeddings and direct limits

3.1 Strong embedding classes

Definition 3.1.1. A **strong embedding class** $K = (S_K, E_K)$ is a class S_K of finite L -structures, closed under isomorphism, together with a class E_K of embeddings, called **strong embeddings** or **K -embeddings**, between members of S_K , such that:

- (1) E_K is closed under composition.
- (2) Every isomorphism between members of S_K is in E_K .

Let \mathcal{C}_L be the category whose objects are L -structures and whose arrows are embeddings. Then Definition 3.1.1 says that K is a subcategory of \mathcal{C}_L which is full with respect to isomorphisms, such that every object in K is finite.

Remark 3.1.2. An immediate consequence of Definition 3.1.1 is that the class E_K of strong embeddings is closed under isomorphism. That is, if the following diagram in \mathcal{C}_L is commutative, then f is strong if and only if g is strong.

$$\begin{array}{ccc} B & \xleftrightarrow{\sim} & B' \\ f \uparrow & & \uparrow g \\ A & \xleftrightarrow{\sim} & A' \end{array}$$

Notation 3.1.3. If A is a substructure of B and the inclusion $A \rightarrow B$ is a strong embedding, we write $A \preceq_K B$ and say that A is a strong substructure of B . We drop the K from the notation if it is clear from context.

This definition generalizes the setting of classical Fraïssé theory (see, e.g. [17, Sections 2.6–8] or [43, Section 7.1]), which deals with hereditary classes.

Definition 3.1.4. A strong embedding class K is a **hereditary class** if S_K is closed under substructure and E_K consists of all embeddings between structures in S_K . Equivalently, if $B \in S_K$ and $f: A \rightarrow B$ is an embedding, then $A \in S_K$ and $f \in E_K$.

Our convention in this thesis is to work exclusively with relational languages. But since a strong embedding class K need not be hereditary, we can simulate an n -ary function symbol f using an $(n + 1)$ -ary relation symbol R_f by including in K only those structures in which R_f is interpreted as the graph of a function. A drawback to this approach is that our framework only handles *finite* structures, while classical Fraïssé theory in languages with function symbols can be extended easily to classes of *finitely generated* structures.

Definition 3.1.5. A strong embedding class K is a **chain class** if there is a chain

$$A_0 \preceq A_1 \preceq A_2 \preceq \dots$$

such that every structure in K is isomorphic to A_i for some i . Note that we do *not* require that every strong embedding is isomorphic to one of the inclusions $A_i \rightarrow A_j$ for $i < j$.

In some sense, hereditary classes and chain classes are opposite extremes. Chain classes contain at most one structure of size n up to isomorphism for all n , while heredity classes are very rich, containing many non-isomorphic structures of each size (except in trivial cases). Of course, “many” may still mean finitely many, as in the following definition.

Definition 3.1.6. Let K be a strong embedding class, or just a class of finite structures (without specified embeddings). Define $K_n = \{A \text{ in } K \mid |A| = n\}$. K is **small** if K_n is finite up to isomorphism for all $n \in \omega$.

Of course, if the language L is finite, there are only finitely many L -structures of size n up to isomorphism, so every class of finite structures is small. But even if the language is infinite, smallness of K gives us an analog of Lemma 1.1.1.

Lemma 3.1.7. *Suppose K is a small class of finite structures. Then for every L -structure A of size n (not necessarily in K), enumerated by a tuple \bar{a} , there is an explicitly non-redundant quantifier-free formula $\varphi_A(\bar{x})$ such that $A \models \varphi(\bar{a})$, and if B is in K_n , enumerated by a tuple \bar{b} , then $B \models \varphi_A(\bar{b})$ if and only if the map $a_i \mapsto b_i$ is an isomorphism $A \rightarrow B$ (so if A is not in K , $B \models \neg\varphi_A(\bar{b})$ for all B in K_n).*

Proof. We may assume A has domain $[n]$ and \bar{a} is the tuple $(0, \dots, n-1)$. Let $\{B_1, \dots, B_k\}$ be the set of structures in K with domain $[n]$ which are not equal to A (a structure may appear multiple times on this list up to isomorphism, but the list is finite, since K is small). Now for all $1 \leq i \leq k$, there is some quantifier-free formula $\varphi_i(\bar{x})$ such that $A \models \varphi_i(0, \dots, n-1)$, but $B_i \models \neg\varphi_i(0, \dots, n-1)$. Let $\varphi_A(\bar{x})$ be $\left(\bigwedge_{i=1}^k \varphi_i(\bar{x})\right) \wedge \left(\bigwedge_{i \neq j} x_i \neq x_j\right)$.

For all B in K_n , enumerated by \bar{b} , if the map $i \mapsto b_i$ is an isomorphism $A \rightarrow B$, then since $A \models \varphi_A(0, \dots, n-1)$, $B \models \varphi_A(\bar{b})$. Conversely, if the map is not an isomorphism, then the map $\sigma: b_i \mapsto i$ is not an isomorphism $B \rightarrow A$, so σ is an isomorphism $B \rightarrow B_i$ for some i . Then $B \models \neg\varphi_i(\bar{b})$, so $B \models \neg\varphi_A(\bar{b})$. \square

Sometimes we will have occasion to pare down a strong embedding class to a subclass.


Definition 3.1.8. Let $K = (S_K, E_K)$ and $K' = (S_{K'}, E_{K'})$ be strong embedding classes.

- K' is a **subclass** of K if $S_{K'} \subseteq S_K$ and $E_{K'} \subseteq E_K$. That is, K' is a subcategory of K .
- K' is a **full subclass** of K if additionally when $A, B \in S_{K'}$ and $f: A \rightarrow B$ is in E_K , then f is in $E_{K'}$. That is, K' is a full subcategory of K .
- K' is a **cofinal subclass** of K if K' is a full subclass of K and for all A in K there exists B in K' and a strong embedding $A \rightarrow B$ in E_K .

We will consider the following examples repeatedly in later sections.

Example 3.1.9. Let $K_0 = (S_{K_0}, E_{K_0})$, where S_{K_0} is the class of all finite graphs and E_{K_0} is the class of all embeddings between finite graphs. Then K_0 is a strong embedding class.

Let $K_1 = (S_{K_1}, E_{K_1})$, where S_{K_1} is the class of all finite acyclic graphs of degree at most two and E_{K_1} is the class of all embeddings between members of S_{K_1} . Graphs in S_{K_1} are disjoint unions of finite chains. For example:



Then K_1 is a full subclass of K_0 , but K_1 is not cofinal in K_0 .

Let $K_2 = (S_{K_2}, E_{K_2})$, where S_{K_2} is the class of all finite *connected* acyclic graphs of degree at most two, and E_{K_2} is again the class of all embeddings between members of S_{K_2} . K_2 is a cofinal subclass of K_1 , since we can always add new vertices joining the connected components of a graph in K_1 into a single chain.

Let $K_3 = (S_{K_3}, E_{K_3})$, where $S_{K_3} = S_{K_1}$, but E_{K_3} is the class of all embeddings $A \rightarrow B$ such that the image of a connected component of A is a connected component of B . K_3 is a subclass of K_1 which is not full.

K_0 and K_1 are hereditary classes, but K_2 is not: graphs in K_2 may have subgraphs which are not connected. However, K_2 is a chain class (this terminology is unrelated to the fact that we describe the graphs in K_2 as chains!). Indeed, there is exactly one structure C_n in K_2 of size n up to isomorphism, and C_n embeds in C_{n+1} for all n .

Having established the terminology for strong embedding classes, we move on to study their direct limits.

3.2 Direct limits and extendibility

The category \mathcal{C}_L has all direct limits (colimits along directed systems). To avoid any confusion, we recall the definitions here.

A **directed set** (I, \leq) is a poset in which every finite subset of I has an upper bound. This includes the empty subset, so I is nonempty. We view (I, \leq) as a category with at most one arrow between any two objects. A **directed system** in \mathcal{C}_L is a functor F from a directed set (I, \leq) to \mathcal{C}_L . If $A_i = F(i)$ for $i \in I$, we denote the directed system by (A_i) ,

suppressing the index set I and the embeddings in the notation. If $i \leq j$, we denote the image of the unique arrow from i to j in (I, \leq) by $e_{ij}: A_i \rightarrow A_j$.

The **direct limit** $\varinjlim(A_i)$ of the directed system (A_i) has underlying set $(\bigsqcup_{i \in I} A_i) / \sim$, where, for $a \in A_i$ and $b \in A_j$, $a \sim b$ if and only if $e_{ik}(a) = e_{jk}(b)$ in A_k for some k with $i \leq k$ and $j \leq k$. This is an equivalence relation (transitivity holds using directedness), and the choice of k does not matter, since all maps in \mathcal{C}_L are injective. If all the connecting maps e_{ij} are inclusions (so $A_i \subseteq A_j$ when $i \leq j$ in I), the equivalence relation \sim is just equality, and we take the underlying set of $\varinjlim(A_i)$ to be the union of the underlying sets of the A_i .

If R is a relation symbol, $\varinjlim(A_i) \models R(a_1, \dots, a_n)$ if and only if $A_k \models R(b_1, \dots, b_n)$ for some k such that the \sim equivalence classes a_1, \dots, a_n have representatives b_1, \dots, b_n in A_k . Such a k exists by directedness, and the choice of k does not matter, since all maps in \mathcal{C}_L are embeddings.

Notation 3.2.1. If $M = \varinjlim(A_i)$, we write π_i^M for the map $A_i \rightarrow M$ sending each element to its equivalence class. We can check whether elements of A_i are equal in M or whether relations hold between elements of A_i in M by looking at A_i itself, so π_i^M is an embedding. For notational convenience, we often identify each A_i with its image in M under π_i^M , so that all of the embeddings π_i^M and e_{ij} are inclusions, and $M = \bigcup_{i \in I} A_i$.

The direct limit satisfies a universal property. Let $M = \varinjlim(A_i)$. Given embeddings $f_i: A_i \rightarrow N$ for all $i \in I$, such that $f_j \circ e_{ij} = f_i$ whenever $i \leq j$ in I , there is a unique embedding $f: M \rightarrow N$ such that $f \circ \pi_i^M = f_i$ for all $i \in I$.

In defining the structure on $\varinjlim(A_i)$, we used the fact that a finite subset of $\varinjlim(A_i)$ is contained in some A_k . A slight strengthening of this will be constantly useful, so we state it here for the record.

Lemma 3.2.2. *If $M = \varinjlim(A_i)$, then for any A_j and any finite $B \subseteq M$, there is a k such that $j \leq k$ (so A_j embeds in A_k in the directed system), and $B \subseteq A_k$.*

Proof. For each $b \in B$, pick A_{i_b} in the directed system so that b is an element of A_{i_b} . Then by directedness, there is a k such that $j \leq k$ and $i_b \leq k$ for all b . \square

We are interested in the direct limits of finite structures in K along strong embeddings.

Definition 3.2.3. We say a structure M is a **K -direct limit** if it is isomorphic to the direct limit in \mathcal{C}_L of some directed system in the subcategory K (i.e. all objects are in K and all embeddings are strong).

If M is a K -direct limit, then M is isomorphic to $\varinjlim(A_i)$ for some directed system (A_i) in K . If we fix such an isomorphism, which we call a **presentation** of M as a K -direct limit, we adopt Notation 3.2.1, suppressing the isomorphism and identifying the A_i with substructures of M .

Remark 3.2.4. Every structure in K is a K -direct limit (along a directed system with just one object), and every finite K -direct limit is in K (by Lemma 3.2.2, if $M = \varinjlim(A_i)$ is finite, then there is some k such that the image in A_k in M is all of M , and hence $M \cong A_k$).

We now extend our class of strong embeddings to K -direct limits in the natural way, taking the K -embeddings and the embeddings π_i^M to be strong and closing under composition (this tells us when an embedding from a finite structure in K to a K -direct limit is strong) and the universal property of the direct limit (an embedding out of a K -direct limit is strong if it is induced by strong embeddings out of the structures in its directed system). This process is distilled in the concrete definition below. In categorical language, we take the category K_∞ of K -direct limits and strong embeddings to be equivalent to the category $\text{ind-}K$ of formal direct limits in K (See [50, Section VI.1]).

Definition 3.2.5. Let M and N be K -direct limits, given together with direct limit presentations, $M \cong \varinjlim(A_i)$ and $N \cong \varinjlim(B_j)$, and let $f: M \rightarrow N$ be an embedding. Then f is a **strong embedding** if for all A_i , $f \circ \pi_i^M$ factors as $\pi_j^N \circ g$ for some B_j and some strong embedding $g: A_i \rightarrow B_j$.

$$\begin{array}{ccc} M & \xrightarrow{f} & N \\ \pi_i^M \uparrow & & \uparrow \pi_j^N \\ A_i & \xrightarrow{g} & B_j \end{array}$$

Note that this definition depends on the presentations of M and N as K -direct limits.

Proposition 3.2.6. *The class of strong embeddings between K -direct limits is closed under composition.*

Proof. Let $M \cong \varinjlim(A_i)$, $M' \cong \varinjlim(B_j)$, and $M'' \cong \varinjlim(C_k)$ be K -direct limits, given together with direct limit presentations, and let $f: M \rightarrow M'$ and $f': M' \rightarrow M''$ be strong embeddings.

$$\begin{array}{ccccc} M & \xrightarrow{f} & M' & \xrightarrow{f'} & M'' \\ \pi_i^M \uparrow & & \uparrow \pi_j^{M'} & & \uparrow \pi_k^{M''} \\ A_i & \xrightarrow{g} & B_j & \xrightarrow{g'} & C_k \end{array}$$

For any A_i , we have $f' \circ f \circ \pi_i^M = f' \circ \pi_j^{M'} \circ g$ for some B_j and strong $g: A_i \rightarrow B_j$, since f is strong, and $f' \circ \pi_j^{M'} \circ g = \pi_k^{M''} \circ g' \circ g$ for some C_k and strong $g': B_j \rightarrow C_k$, since f' is strong. So $(f' \circ f) \circ \pi_i^M = \pi_k^{M''} \circ (g' \circ g)$, and the embedding $(g' \circ g): A_i \rightarrow C_k$ is strong since E_K is closed under composition. \square

However, the class of strong embeddings between K -direct limits as defined above may not contain all isomorphisms. We rule out this undesirable situation with a definition.

Definition 3.2.7. A strong embedding class K is **extendible** if whenever M and N are K -direct limits with direct limit presentations, $M \cong \varinjlim(A_i)$ and $N \cong \varinjlim(B_j)$, if $f: M \rightarrow N$ is an isomorphism, then f is a strong embedding.

Proposition 3.2.8. *Let K be an extendible strong embedding class. Let M and N be K -direct limits, and let $f: M \rightarrow N$ be an embedding. Whether f is strong does not depend on the presentations of M and N as K -direct limits.*

Proof. Given presentations $M \cong \varinjlim(A_i)$ and $N \cong \varinjlim(B_j)$ of M and N as K -direct limits, to say $f: M \rightarrow N$ is strong is to say that the corresponding embedding $g: \varinjlim(A_i) \rightarrow \varinjlim(B_j)$ (obtained by the identifications of M and N with the direct limits) is strong.

Similarly, if we pick two other presentations $M \cong \varinjlim(A'_i)$ and $N \cong \varinjlim(B'_j)$ of M and N as K -direct limits, then f is strong with respect to these presentations if and only if $g': \varinjlim(A'_i) \rightarrow \varinjlim(B'_j)$ is. Then we have the following commutative diagram.

$$\begin{array}{ccccc} \varinjlim(B_j) & \longleftrightarrow & N & \longleftrightarrow & \varinjlim(B'_j) \\ \uparrow g & & \uparrow f & & \uparrow g' \\ \varinjlim(A_i) & \longleftrightarrow & M & \longleftrightarrow & \varinjlim(A'_i) \end{array}$$

Since strong embeddings are closed under composition (Proposition 3.2.6) and isomorphisms are strong, g is strong if and only if g' is strong. \square

Notation 3.2.9. Let K be an extendible class. We write $K_\infty = (S_{K_\infty}, E_{K_\infty})$, where S_{K_∞} is the class of K -direct limits and E_{K_∞} is the class of strong embeddings between them (this is well-defined by Proposition 3.2.8). As in the finite case, if M is a substructure of N and the inclusion $M \rightarrow N$ is a strong embedding, we write $M \preceq_K N$, and we drop the K if it is clear from context.

Remark 3.2.10. As a special case of Definition 3.2.5, if A is in K , $N \cong \varinjlim(B_j)$ is a K -direct limit, and $A \subseteq N$, then $A \preceq N$ if and only if $A \preceq B_j$ for some j .

It is convenient to have a more concrete criterion for extendibility, which refers only to K , not to its direct limits.

Proposition 3.2.11. *A strong embedding class K is extendible if and only if whenever a “ladder diagram” of the following form commutes in \mathcal{C}_L , where the f_i and g_i are strong embeddings, but the α_i and β_i are arbitrary embeddings,*

$$\begin{array}{ccccccc} A_0 & \xrightarrow{f_0} & A_1 & \xrightarrow{f_1} & A_2 & \xrightarrow{f_2} & \dots \\ & \searrow \alpha_0 & \nearrow \beta_0 & \searrow \alpha_1 & \nearrow \beta_1 & \searrow \alpha_2 & \nearrow \beta_2 \\ & & B_0 & \xrightarrow{g_0} & B_1 & \xrightarrow{g_1} & B_2 & \xrightarrow{\quad} & \dots \end{array}$$

then, for some n , the unique map $h: A_0 \rightarrow B_n$ coming from the diagram is a strong embedding.


Proof. Suppose K is extendible. Given a ladder diagram, let $M = \varinjlim(A_i)$ and $N = \varinjlim(B_i)$. Then the embeddings $\pi_i^N \circ \alpha_i: A_i \rightarrow N$ induce an embedding $\alpha: M \rightarrow N$ (since they cohere with the connecting maps: $(\pi_{i+1}^N \circ \alpha_{i+1}) \circ f_i = \pi_{i+1}^N \circ g_i \circ \alpha_i = \pi_i^N \circ \alpha_i$ for all i), and similarly the embeddings $\pi_{i+1}^M \circ \beta_i: B_i \rightarrow M$ induce an embedding $\beta: N \rightarrow M$. Moreover, α and β are inverses: given an element $a \in M$ such that $a = \pi_i^M(a')$ with $a' \in A_i$, we have $\beta(\alpha(a)) = \beta(\pi_i^N(\alpha_i(a'))) = \pi_{i+1}^M(\beta_i(\alpha_i(a'))) = \pi_{i+1}^M(f_i(a')) = \pi_i^M(a') = a$, and similarly $\alpha(\beta(b)) = b$ for $b \in N$. So α is an isomorphism, and hence a strong embedding. Then the map $\alpha \circ \pi_0^M: A_0 \rightarrow N$ must factor as $\pi_n^N \circ h'$ for some n via a strong embedding $h': A_0 \rightarrow B_n$. Now for any $a \in A_0$, $\pi_n^N(h'(a)) = \alpha(\pi_0^M(a)) = \pi_0^N(\alpha_0(a)) = \pi_n^N(h'(a))$, and π_n^N is injective, so $h' = h$.

Conversely, suppose the condition on ladder diagrams is satisfied, and let $\varphi: M \rightarrow N$ be an isomorphism between K -direct limits, with presentations $M \cong \varinjlim(A_i)$ and $N \cong \varinjlim(B_j)$. Identify the A_i and B_j with substructures of M and N , respectively. We would like to show that φ is strong, so let A_{i_0} be among the A_i . Now the image of A_{i_0} under φ is a finite subset of N , so by Proposition 3.2.2, we can find B_{j_0} such that $\varphi(A_{i_0}) \subseteq B_{j_0}$. Let $\alpha_0 = (\varphi \upharpoonright A_{i_0}): A_{i_0} \rightarrow B_{j_0}$. This is an embedding, which is not necessarily strong.

Similarly, the image of B_{j_0} under φ^{-1} is a finite subset of M , so by Proposition 3.2.2, we can find A_{i_1} such that $A_{i_0} \preceq A_{i_1}$ and $\varphi^{-1}(B_{j_0}) \subseteq A_{i_1}$. Let f_0 be the strong inclusion $A_{i_0} \rightarrow A_{i_1}$, and let $\beta_0 = \varphi^{-1} \upharpoonright B_{j_0}: B_{j_0} \rightarrow A_{i_1}$. Continuing in this way, we build a ladder diagram, and we conclude that for some n , A_{i_0} embeds strongly in B_{i_n} via the map coming from the diagram, i.e. the map $\varphi \circ \pi_{i_0}^M$ factors as $\pi_{i_n}^N \circ h_n$, where $h: A_{i_0} \rightarrow B_{i_n}$ is strong. Hence the condition of Definition 3.2.5 for φ to be a strong embedding is satisfied. \square

This condition on the class K is still a bit unwieldy. We now give a finitary condition on K , smoothness, which is simple to verify and implies extendibility. We also define coherence, a weaker, but natural, condition.

Definition 3.2.12. Let K be a strong embedding class.



- K is **smooth** if whenever A, B, C are in K and $f: A \rightarrow B$ and $g: B \rightarrow C$ are embeddings such that $g \circ f$ is strong, then f is also strong.
- K is **coherent** if whenever A, B, C , are in K and $f: A \rightarrow B$ and $g: B \rightarrow C$ are embeddings such that g is strong and $g \circ f$ is strong, then f is also strong.

Similarly, if K is extendible, we say K_∞ is smooth or coherent if the same conditions hold for all $A, B, C \in K_\infty$.

Remark 3.2.13. Every hereditary class is smooth, coherent, and extendible, since all embeddings are strong.

Proposition 3.2.14. *Let K be an extendible strong embedding class. If K is smooth, then so is K_∞ . If K is coherent, then so is K_∞ .*

Proof. Suppose we have embeddings of K -direct limits $f: M_1 \rightarrow M_2$ and $g: M_2 \rightarrow M_3$. Further, we assume that $g \circ f$ is strong. For convenience, we identify M_1 and M_2 with their images in M_3 , so $M_1 \subseteq M_2$, $M_2 \subseteq M_3$, and $M_1 \preceq_K M_3$. Let $M_1 \cong \varinjlim(A_i)$, $M_2 \cong \varinjlim(B_j)$, $M_3 \cong \varinjlim(C_k)$.

For smoothness, to show that $M_1 \preceq_K M_2$, it suffices to show that for any A_i , there is a B_j such that $A_i \preceq_K B_j$. Choose any B_j such that $A_i \subseteq B_j$. Since $M_1 \preceq_K M_3$, there is some C_k such that $A_i \preceq_K C_k$, and we may choose C_k so that $B_j \subseteq C_k$ by Lemma 3.2.2. Now applying smoothness in K , we have $A_i \preceq_K B_j$.

For coherence, we further assume that $M_2 \preceq_K M_3$. Again, we need to show that for any A_i , there is a B_j such that $A_i \preceq_K B_j$, and we choose any B_j such that $A_i \subseteq B_j$. Now there is some C_k such that $A_i \preceq_K C_k$ and some C'_k such that $B_j \preceq_K C'_k$. By directedness of the (C_k) system, there is some C''_k such that $A_i \preceq_K C_k \preceq_K C''_k$ and $B_j \preceq_K C'_k \preceq_K C''_k$. Now applying coherence in K , we have $A_i \preceq_K B_j$. \square

Remark 3.2.15. We borrow the term smooth from Kueker and Laskowski [66] who use it for a slightly stronger condition. In their definition, a class K is smooth if for all $A \in K$, there is a universal first-order theory T_A with parameters from A such that if B is in K and A is a substructure of B , then $A \preceq B$ if and only if $B \models T_A$. A smooth class in the sense of Kueker and Laskowski is clearly smooth in our sense, since the truth of universal sentences is preserved under substructure.

The sort of definability condition on the class of strong embeddings assumed by Kueker and Laskowski is sometimes useful (see our similar use of generic semi-definability, Definition 4.4.4, in Section 4.4), but one of our goals here is to do as much as possible without such assumptions. Observe, however, that if the language L is finite, then any smooth class in our sense is also smooth in the sense of Kueker and Laskowski. For every B in K such that A is a substructure of B but $A \not\preceq B$, let \bar{b} enumerate $B \setminus A$, let $\varphi_B(\bar{a}, \bar{b})$ be the conjunction of the diagram of B (as in Lemma 1.1.1), and include the sentence $\forall \bar{y} \neg \varphi_B(\bar{a}, \bar{y})$ in T_A . Suppose C is a structure in K and A is a substructure of C . If $A \not\preceq C$, then $C \models \exists \bar{y} \varphi_C(\bar{a}, \bar{y})$, so $C \not\models T_A$. On the other hand, if $A \preceq C$, then for any tuple \bar{c} from C such that $\bar{a}\bar{c}$ enumerates a structure in K , $A \preceq \llbracket \bar{a}\bar{c} \rrbracket \subseteq C$ (by smoothness), so $C \models \neg \varphi_B(\bar{a}, \bar{c})$ for all B in K such that A is a substructure of B but $A \not\preceq B$. And if \bar{c} is a tuple from C such that $\bar{a}\bar{c}$ is redundant or $\llbracket \bar{a}\bar{c} \rrbracket$ is not in K , then $C \models \neg \varphi_B(\bar{a}, \bar{c})$ for all B in K . Hence $C \models T_A$.


Smooth implies coherent and extendible, but the latter conditions are independent.

Proposition 3.2.16. *Every smooth class is coherent and extendible.*

Proof. It is clear from the definition that smoothness implies coherence. For extendibility, we use the characterization of Proposition 3.2.11. Given a ladder diagram, smoothness implies that since the map f_0 is strong, then already the map $\alpha_0: A_0 \rightarrow B_0$ is strong. \square

Example 3.2.17 (A class which is extendible but not coherent). Let L be the empty language, and let K be the class of all finite sets. Say an embedding $A \rightarrow B$ is strong when $|B| \neq |A| + 1$.

K is extendible: Given a partial ladder diagram, with f_0 and g_0 strong embeddings and $\alpha_0, \beta_0, \alpha_1$ arbitrary embeddings,



we have $|A_0| \leq |B_0| \leq |A_1| \leq |B_1|$. So either $|A_1| = |A_0|$, in which case also $|B_0| = |A_0|$, and α_0 is a strong embedding, or $|A_1| \geq |A_0| + 2$, in which case also $|B_1| \geq |A_0| + 2$, and $\alpha_1 \circ f_0: A_0 \rightarrow B_1$ is a strong embedding.

K is not coherent, and hence not smooth: Let A be any finite set, let $f: A \rightarrow B$ be an embedding of A into a set B of cardinality $|A| + 1$, and let $g: B \rightarrow C$ be an embedding of B into a structure C of cardinality $|A| + 3$. Then both g and $g \circ f$ are strong embeddings, but f is not.

Remark 3.2.18. Taking Proposition 3.2.16 and Example 3.2.17 together, the moral is that there are two easy ways to obtain extendibility. Smoothness says that strong substructure is closed downward ($A \preceq C$ and $A \subseteq B \subseteq C$ implies $A \preceq B$), and this easily implies extendibility. But if (proper) strong substructure is closed upward instead ($A \prec B$ and $B \subseteq C$ implies $A \prec C$), this also implies that K is extendible.

Example 3.2.19 (A class which is coherent but not extendible). Again, let L be the empty language, and let K be the class of all finite sets. Say an embedding $A \rightarrow B$ is strong when $|A|$ and $|B|$ are both even or $|A|$ and $|B|$ are both odd.

K is coherent: If $f: A \rightarrow B$ and $g: B \rightarrow C$ are embeddings such that g is strong and $g \circ f$ are strong, then $|B|$ and $|C|$ have the same parity and $|A|$ and $|C|$ have the same parity, so $|A|$ and $|B|$ have the same parity. Hence f is strong.

K is not extendible, and hence not smooth: Set $A_n = [2n]$ and $B_n = [2n + 1]$, and let $\alpha_n: A_n \rightarrow B_n$ and $\beta_n: B_n \rightarrow A_{n+1}$ be the inclusions. This describes a ladder diagram in which all the composite embeddings $A_n \rightarrow A_{n+1}$ and $B_n \rightarrow B_{n+1}$ are strong, but no embedding $A_0 \rightarrow B_n$ is strong.

Remark 3.2.20. Example 3.2.19 illustrates the somewhat unintuitive behavior of non-extendible classes. In this example, a countable set which is presented as a direct limit of even-sized finite sets is “different” than a countable set which is presented as a direct limit of odd-sized finite sets: no bijection between them is strong. In an extendible class, on the other hand, the presentation of a structure as a direct limit carries no extra information that is not captured in the structure’s isomorphism type as an L -structure. In fact, for an extendible class K , every K -direct limit has a canonical maximal presentation, and hence

any two presentations of the same structure as a K -direct limit fit together in a natural way. This is the content of the next proposition, which is the key fact about extendible classes.

Proposition 3.2.21. *Let K be an extendible class, and let M be a K -direct limit. The set of all finite strong substructures of M , together with the inclusions between them which are strong embeddings, is a directed system D_M in K , the direct limit of which is M .*

If M and N are K -direct limits, an embedding $f: M \rightarrow N$ is strong if and only if f identifies D_M with a subsystem of D_N .

Proof. To show that D_M is a directed system, we just need to show that for any $A \preceq_K M$ and $B \preceq_K M$, there is some $C \preceq_K M$ with $A \preceq_K C$ and $B \preceq_K C$. Pick any presentation of M as a K -direct limit, $M \cong \varinjlim (C_i)$. Then there are some i and j such that $A \preceq_K C_i$ and $B \preceq_K C_j$. By directedness of the (C_i) system, there is some C_k with $A \preceq_K C_i \preceq_K C_k$ and $B \preceq_K C_j \preceq_K C_k$.

The inclusions of the structures in D_M into M induce an embedding $\varinjlim D_M \rightarrow M$, which is also surjective, since for each $m \in M$, $m \in C_i$ for some $C_i \in D_M$. So $\overline{M} \cong \varinjlim D_M$.

For the second claim, consider an embedding $f: M \rightarrow N$ with $M \cong \varinjlim D_M$ and $N \cong \varinjlim D_N$. If $f: M \rightarrow N$ is strong, then for all A in the system D_M , $f(A) \preceq_K N$, so $f(A)$ appears in the system D_N . And f carries inclusions to inclusions, so f identifies D_M with a subsystem of D_N . Conversely, if for all A in the system D_M , $f(A)$ is in the system D_N , then $f \circ \pi^M$ factors as $\pi^N \circ f_A$, where f_A is the (strong) isomorphism $A \rightarrow f(A)$, so f is strong. \square

3.3 Abstract elementary classes

We will now take a moment to observe that extendibility and coherence are sufficient to ensure that the class of K -direct limits, equipped with the strong substructure relation, is an abstract elementary class. For more on AECs, see [10].

Definition 3.3.1 (Shelah, [85]). An **abstract elementary class (AEC)** in the language L is a class C of L -structures together with a relation \preceq such that:

- (1) \preceq is a partial order on C .
- (2) If $M \preceq N$, then M is a substructure of N .
- (3) C and \preceq are closed under isomorphism: if $M \in C$, $f: M \rightarrow M'$ is an isomorphism, and $N \preceq M$, then $M' \in C$ and $\sigma(N) \preceq M'$.
- (4) If $M_1, M_2, M_3 \in C$ such that $M_1 \subseteq M_2$, $M_2 \preceq M_3$, and $M_1 \preceq M_3$, then $M_1 \preceq M_2$.
- (5) If $\{M_\alpha \mid \alpha < \gamma\}$ satisfies $M_\alpha \preceq M_\beta$ for all $\alpha \leq \beta < \gamma$, then

- a) $\bigcup_{\alpha < \gamma} M_\alpha \in C$, and

- b) if $M_\alpha \preceq N$ for all $\alpha < \gamma$, then $\bigcup_{\alpha < \gamma} M_\alpha \preceq N$.
- (6) There is an infinite cardinal $\text{LS}(C)$, called the **Löwenheim–Skolem number** of C , such that if $M \in C$ and $A \subseteq M$, then there exists $N \preceq M$ such that $A \subseteq N$ and $|N| \leq |A| + \text{LS}(C)$.

Proposition 3.3.2. *Let K be an extendible and coherent strong embedding class. Then $(S_{K_\infty}, \preceq_K)$ is an AEC with Löwenheim–Skolem number \aleph_0 . In particular, this is true when K is smooth. If K is extendible but not coherent, only condition (4) fails.*

Proof. (1): This holds because K_∞ is a subcategory of \mathcal{C}_L (identity maps are strong, and strong embeddings are closed under composition).

(2): By definition.

(3): The class of K -direct limits is closed under isomorphism, and so is the class of strong embeddings between them, as shown in Proposition 3.2.8.

(4): This is just coherence for K_∞ . We have shown in Proposition 3.2.14 that if K is coherent, so is K_∞ .

(5): For (a), we must show that the direct limit of a chain of K -direct limits is again a K -direct limit. For each M_α , consider the maximal presentation of $M_\alpha \cong \varinjlim D_\alpha$ as the K -direct limit, as in Proposition 3.2.21. By the same proposition, D_α is a subsystem of D_β when $\alpha \leq \beta$, since $M_\alpha \preceq_K M_\beta$. Taking the union along the chain of directed systems $\{D_\alpha \mid \alpha < \gamma\}$, we obtain a directed system D_γ , the direct limit along which is $\bigcup_{\alpha < \gamma} M_\alpha$.

For (b), observe that if D_N is the maximal presentation of N as a K -direct limit, then all of the directed systems D_α are subsystems of D_N . Then also D_γ is a subsystem of D_N , so $\bigcup_{\alpha < \gamma} M_\alpha \preceq_K N$ by Proposition 3.2.21.

(6): Take $\text{LS}(C) = \aleph_0$. Let $M \in K_\infty$ and $A \subseteq M$. If A is empty, we can take any finite strong substructure of M . Otherwise, for each element $a \in A$, pick some finite $B_{\{a\}} \preceq_K M$ such that $a \in B_{\{a\}}$. Now we close the family $\{B_{\{a\}} \mid a \in A\}$ to a directed system of finite strong substructures of M , indexed by the directed set of non-empty finite subsets of A , by picking, for each non-empty finite $X \subseteq A$, a finite structure $B_X \preceq_K M$ such that $B_Y \preceq_K B_X$ whenever $Y \subseteq X$. This can be done by induction on the size of X , using the fact (Proposition 3.2.21 again) that the strong substructures of M form a directed system. The direct limit $N = \varinjlim (B_X)$ embeds as a strong substructure of M . If A is finite, then N is just B_A , which is finite, and $|N| < \aleph_0 = |A| + \aleph_0$. If A is infinite, then there are $|A|$ finite subsets of A , each contributing finitely many elements to the direct limit, so $|N| \leq |A| = |A| + \aleph_0$. \square

As the proof of (6) shows, the Löwenheim–Skolem axiom holds in a strong form for the AEC $(S_{K_\infty}, \preceq_K)$. This condition is called local finiteness.

Definition 3.3.3 (Baldwin, Koerwein, and Laskowski [11]). An AEC (C, \preceq) is **locally finite** if for all $M \in C$ and finite $A \subseteq M$, there is a finite $N \in C$ with $A \subseteq N \preceq M$. Equivalently, every structure in C is the direct limit of its finite strong substructures.

The term “finitary AEC” has been used by Hyttinen and Kesälä (see [47], for example), for an altogether different notion. To avoid confusion, and as a demonstration of how tameness properties of AECs may interact with properties of the class K , we give their definition here and explore which clauses must hold and which may fail in our context.

Definition 3.3.4 (Hyttinen and Kesälä [47]). An AEC is **finitary** if:

- (1) It has Löwenheim–Skolem number \aleph_0 ,
- (2) It has arbitrarily large models,
- (3) It has the amalgamation property,
- (4) It has the joint embedding property, and
- (5) It has finite character.

We have already noted that an AEC of the form $(S_{K_\infty}, \preceq_K)$ satisfies condition (1). We will now define finite character and show that condition (5) is also satisfied. But Examples 3.3.7 and 3.3.8 below will show that conditions (2), (3), and (4) may fail. For the definitions of the amalgamation property and the joint embedding property, see Definition 3.4.1 below.

Definition 3.3.5. An AEC has **finite character** if $M_1 \preceq M_2$ if and only if $M_1 \subseteq M_2$ and for all tuples \bar{a} from M_1 , there exist N and strong embeddings $f: M_1 \rightarrow N$ and $g: M_2 \rightarrow N$ such that $f(\bar{a}) = g(\bar{a})$.

The intuition is that the existence of the strong embeddings f and g witness that $\text{tp}_{M_1}(\bar{a}) = \text{tp}_N(\bar{a}) = \text{tp}_{M_2}(\bar{a})$, for some abstract notion of “type” suitable to AECs, and thus the condition says that to check that $M_1 \preceq M_2$, it suffices to check that $\text{tp}_{M_1}(\bar{a}) = \text{tp}_{M_2}(\bar{a})$ for any finite tuple \bar{a} .

Proposition 3.3.6. *If K is an extendible and coherent strong embedding class, then the AEC $(S_{K_\infty}, \preceq_K)$ has finite character.*

Proof. Let $M_1 \subseteq M_2$ be K -direct limits satisfying the hypothesis on finite tuples from M_1 . To show that $M_1 \preceq_K M_2$, we need to show that for all finite $A \preceq_K M_1$, $A \preceq_K M_2$. Let \bar{a} enumerate A . Then we have strong embeddings $f: M_1 \rightarrow N$ and $g: M_2 \rightarrow N$ such that $f(\bar{a}) = g(\bar{a})$. Note that since f is strong, $f(A) \preceq_K N$. Then we have $g(A) \subseteq g(M_2)$, $g(A) = f(A) \preceq_K N$, and $g(M_2) \preceq_K N$. By coherence, $g(A) \preceq_K g(M_2)$, and by isomorphism invariance, $A \preceq_K M_2$. \square

Example 3.3.7 (A hereditary class K such that K_∞ has no uncountable models). Consider the class K_2 from Example 3.1.9 (finite connected acyclic graphs of degree at most two). Any structure in $(K_2)_\infty$ is again a connected acyclic graph of degree at most two. But there are only two such infinite graphs, and both are countable: an infinite chain of vertices with one endpoint, A_1 , and an infinite chain of vertices with no endpoints, A_0 . The latter is maximal, in the sense that any embedding from A_0 to a structure in $(K_2)_\infty$ is an isomorphism.


Example 3.3.8 (A hereditary class which fails to have the joint embedding property and the amalgamation property). Let $L = \{P\}$, where P is a unary predicate, and let K be the class of all finite sets such that every element is in P or no element is in P , equipped with all embeddings between structures in K . If A and B are nonempty structures such that all elements of A satisfy P but no elements of B satisfy P , then there is no C in K such that A and B both embed in C , and the joint embedding property fails. Further, the unique empty structure in K embeds in both A and B , so the amalgamation property fails. Since K is contained in K_∞ , K_∞ also fails to satisfy these properties.

3.4 Amalgamation


We now introduce the joint embedding property, the amalgamation property, and an important variant, the weak amalgamation property.

Definition 3.4.1. Let K be a strong embedding class.

- K has the **joint embedding property** if for all A and B in K there exist strong embeddings $f: A \rightarrow C$ and $g: B \rightarrow C$ for some C in K .



- A strong embedding $f: A \rightarrow B$ is an **amalgamation embedding** if for all strong embeddings $g_1: B \rightarrow C_1$ and $g_2: B \rightarrow C_2$, there exist strong embeddings $h_1: C_1 \rightarrow D$ and $h_2: C_2 \rightarrow D$ for some D in K , such that $h_1 \circ g_1 \circ f = h_2 \circ g_2 \circ f$ (as in the diagram on the left).



If the identity map id_A is an amalgamation embedding, we say that A is an **amalgamation base**. Equivalently, for all strong embeddings $g_1: A \rightarrow C_1$ and $g_2: A \rightarrow C_2$, there exist strong embeddings $h_1: C_1 \rightarrow D$ and $h_2: C_2 \rightarrow D$ for some D in K , such that $h_1 \circ g_2 = h_2 \circ g_1$ (as in the diagram on the right).

- K has the **weak amalgamation property** if for all A in K there exists an amalgamation embedding $f: A \rightarrow B$ for some B in K .
- K has the **amalgamation property** if every A is an amalgamation base (i.e. we can always take $B = A$ and $f = \text{id}_A$ in the definition of the weak amalgamation property).

Remark 3.4.2. It is tempting to view the joint embedding property as essentially a special case of the amalgamation property and argue that we may as well include the empty structure in K and adopt the convention that all embeddings out of the empty structure are strong. This argument of convenience fails when we recall that our logical formalism allows 0-ary relation symbols, so there may be multiple non-isomorphic empty structures.

A simple example of a hereditary class with the amalgamation property but without the joint embedding property is the class of all finite structures and all embeddings in the language with a single 0-ary relation symbol.

We will see in Section 4.3 that the joint embedding property corresponds to the existence of a *complete* generic theory (relative to which, in particular, each 0-ary relation symbols must be designated to be true or false).

Definition 3.4.3. A **generalized Fraïssé class** is an extendible strong embedding class with the joint embedding property and the weak amalgamation property which contains only countably many structures up to isomorphism.

This definition generalizes the classical notion of a Fraïssé class.

Definition 3.4.4. A **Fraïssé class** is a hereditary class with the joint embedding property and the amalgamation property which contains only countably many structures up to isomorphism.

Proposition 3.4.5. *The class of amalgamation embeddings is stable under composition with strong embeddings. That is, if $f: A \rightarrow B$ is an amalgamation embedding, and $f^*: A' \rightarrow A$ and $f_*: B \rightarrow B'$ are strong embeddings, then $f_* \circ f$ and $f \circ f^*$ are amalgamation embeddings.*

Consequently, the class of amalgamation embeddings is closed under isomorphism, and K has the amalgamation property if and only if every strong embedding is an amalgamation embedding.

Proof. It suffices to show that $f_* \circ f \circ f^*: A' \rightarrow B'$ is an amalgamation embedding, since we can take f^* or f_* to be identity maps. Given strong embeddings $g_1: B' \rightarrow C_1$ and $g_2: B' \rightarrow C_2$, $(g_1 \circ f_*)$ and $(g_2 \circ f_*)$ are strong embeddings out of B , so there exist $h_1: C_1 \rightarrow D$ and $h_2: C_2 \rightarrow D$ for some D in K , such that $h_1 \circ (g_1 \circ f_*) \circ f = h_2 \circ (g_2 \circ f_*) \circ f$. But then $h_1 \circ g_1 \circ (f_* \circ f \circ f^*) = h_2 \circ g_2 \circ (f_* \circ f \circ f^*)$, as desired.

Since isomorphisms are strong embeddings, it follows that the class of amalgamation embeddings is closed under isomorphism, just as in Remark 3.1.2.


For the last claim, if K has the amalgamation property, then id_A is an amalgamation embedding for all A . Then for any $f: A \rightarrow B$, $f = f \circ \text{id}_A$ is an amalgamation embedding. \square

The easiest way to verify that a class has the weak amalgamation property is to find a cofinal subclass with the amalgamation property, as in the following proposition.

Proposition 3.4.6. *If a strong embedding class K contains a cofinal subclass K' with the amalgamation property, then K has the weak amalgamation property.*

Proof. Let A be in K , and let $f: A \rightarrow B$ be any K -embedding with B in K' . Such an f exists since K' is cofinal in K . We claim that f is an amalgamation embedding. Given K -embeddings $g_1: B \rightarrow C_1$ and $g_2: B \rightarrow C_2$, by cofinality we can find K -embeddings $g'_1: C_1 \rightarrow C'_1$ and $g'_2: C_2 \rightarrow C'_2$, with C'_1 and C'_2 in K' . Since K' is a full subclass of K , $g'_1 \circ g_1$ and $g'_2 \circ g_2$ are K' -embeddings. Now B is an amalgamation base in K' , so we can find D in K' and embeddings $h_1: C'_1 \rightarrow D$ and $h_2: C'_2 \rightarrow D$ such that $h_1 \circ (g'_1 \circ g_1) = h_2 \circ (g'_2 \circ g_2)$. But then $(h_1 \circ g'_1) \circ g_1 \circ f = (h_2 \circ g'_2) \circ g_2 \circ f$, so D amalgamates C_1 and C_2 over A as well. \square


The proof of the proposition shows that if K is a cofinal subclass of K and B is an amalgamation base in K' , then B is an amalgamation base in K . In this situation, we are always able to make the following diagram commute:



This is a stronger condition (amalgamation over B) than what is required in the definition of amalgamation embedding (amalgamation over A). Intuitively, if $f: A \rightarrow B$ is an amalgamation embedding, then the way A embeds into B includes enough information about A to ensure amalgamation over A . But if we only have the weak amalgamation property, we may not yet have enough information about B (i.e. B may not be an amalgamation base). The following example shows that this situation is possible: there is a generalized Fraïssé class which does not contain a cofinal class with the amalgamation property. In fact, it may contain no amalgamation bases at all.

Example 3.4.7. Let $L = \{R, \mathbf{1}, \mathbf{2}, \mathbf{3}, \mathbf{4}, \mathbf{5}\}$, where R is a binary relation and the other symbols are unary relations. Let K be the class of finite non-empty connected acyclic

graphs (with edge relation R), such that each vertex is colored by exactly one of the unary relations, and which omit the following five subgraphs:



For the class of strong embeddings, we take all embeddings between structures in K . Then K is a smooth class, in particular it is extendible.

We say a vertex labeled \mathbf{i} is **determined** if it has a neighbor labeled $\mathbf{i+1}$ or $\mathbf{i+2}$ (here addition is interpreted cyclically mod 5). If a vertex in a structure A is determined, say by having a neighbor labeled $\mathbf{i+1}$, it cannot be connected to a new neighbor labeled $\mathbf{i+2}$ in any extension of A .

If a structure A contains a vertex v labeled \mathbf{i} which is undetermined, then A is not an amalgamation base. Indeed, we can embed A into structures B_1 and B_2 by adding a new vertex connected only to v and labeled $\mathbf{i+1}$ in B_1 and similarly labeled $\mathbf{i+2}$ in B_2 , and these two embeddings cannot be amalgamated over A .

We claim that no structure in K is an amalgamation base. Suppose for contradiction that A is an amalgamation base. Then every vertex in A is determined. Choose an arbitrary vertex v_0 (A is non-empty). Then v_0 has a neighbor which determines it, call this neighbor v_1 . Continue in this way, defining $\{v_i \mid i \in \omega\}$ such that v_{i+1} determines v_i . Note that since v_{i+1} determines v_i , v_i cannot determine v_{i+1} (this is why we needed five predicates). Since also A contains no cycles, $\{v_i \mid i \in \omega\}$ is an infinite path through A , contradicting finiteness.

However, K has the weak amalgamation property. Indeed, given A , let B be a structure obtained from A by adjoining a new vertex connected only to v and determining v for each undetermined $v \in A$. Then any two extensions C_1, C_2 of B can be amalgamated “freely” over A , by not identifying any elements of $C_1 \setminus A$ and $C_2 \setminus A$, and adding no edge relations between these sets. Each element of $B \setminus A$ has two copies in the amalgam, one in C_1 and one in C_2 .

To finish the verification that K is a generalized Fraïssé class, note that K contains only countably many structures up to isomorphism (it is a class of finite structures in a finite relational language), and K has the joint embedding property: we can embed A and B into the disjoint union $A \sqcup B \sqcup \{*\}$, ensuring that the result is connected and acyclic by connecting $*$ arbitrarily to one vertex of A and one vertex of B , and labeling $*$ by any legal \mathbf{i} .


The weak amalgamation property was identified independently by Ivanov [48] (who calls it the “almost amalgamation property”) and Kechris and Rosendal [56]. Of course, these papers both have slightly different settings, and they do not work in the generality of strong embedding classes. The significance of this property is that it is a necessary and sufficient condition for the existence of a generic countable K -direct limit, analogous to the Fraïssé limit of an ordinary Fraïssé class. This is Theorem 4.2.2 below. The weakening of the amalgamation property corresponds to a weakening of the usual ultrahomogeneity property of Fraïssé limits, which we will describe in the next section.

3.5 Universality and homogeneity

Throughout this section, let K be an extendible strong embedding class.


Definition 3.5.1. Let M be a K -direct limit.

- The K -age of M is the class of all structures in K which embed strongly in M .
- M is K -universal if every A in K embeds strongly in M (i.e. the K -age of M is S_K).
- M is **weakly- K -homogeneous** if for all $A \preceq M$, there exists B such that $A \preceq B \preceq M$, and for all strong embeddings $g: B \rightarrow C$ such that C is in the K -age of M , there exists a strong embedding $h: C \rightarrow M$ such that, naming the inclusions $i: A \rightarrow B$ and $j: B \rightarrow M$, we have $h \circ g \circ i = j \circ i$ (as in the diagram on the left). We say B **witnesses weak- K -homogeneity** for A .



M is K -homogeneous if we can always take $B = A$ (as in the diagram on the right).

- M is **weakly- K -ultrahomogeneous** if for all $A \preceq M$, there exists B such that $A \preceq B \preceq M$, and for all strong embeddings $g: B \rightarrow M$, there is an automorphism $\sigma \in \text{Aut}(M)$ such that, naming the inclusions $i: A \rightarrow B$ and $j: B \rightarrow M$, we have $\sigma \circ j \circ i = g \circ i$ (as in the diagram on the left). We say B **witnesses weak- K -ultrahomogeneity** for A .



M is K -ultrahomogeneous if we can always take $B = A$ (as in the diagram on the right).

Theorem 3.5.2. *If M and N are countably infinite K -direct limits which are weakly- K -homogeneous and which have the same K -age, then for any $A \preceq B \preceq M$ such that B witnesses weak- K -homogeneity for A , and any strong embedding $g: B \rightarrow N$, there is an isomorphism $\sigma: M \cong N$ such that, naming the inclusions $i: A \rightarrow B$ and $j: B \rightarrow M$, we have $\sigma \circ j \circ i = g \circ i$.*

Proof. We go back and forth, building the isomorphism σ as a union of a chain of partial isomorphisms $\{\sigma_k \mid k \in \omega\}$, such that σ_k has domain $A_k \preceq M$ and range $A'_k \preceq N$. Along the way, we also define a sequence of partial isomorphisms $\{\tau_k \mid k \in \omega\}$, such that τ_k has domain $B_k \preceq M$ and range $B'_k \preceq N$, with $A_k \preceq B_k$ and $A'_k \preceq B'_k$, and $\sigma_k \subseteq \tau_k$. We will *not* require that $\tau_k \subseteq \tau_{k+1}$ for all k . We will, however, ensure that if k is even, then B_k witnesses weak- K -homogeneity for A_k (in M), and if k is odd, then B'_k witnesses weak- K -homogeneity for A'_k (in N).

To begin, let $A_0 = A$, $B_0 = B$, $A'_0 = g(A)$, $B'_0 = g(B)$, $\sigma_0 = g \upharpoonright A$, and $\tau_0 = g$. By assumption, B_0 witnesses weak- K -homogeneity for A_0 . And already σ_0 carries A to $g(A)$, so this ensures that we will have $\sigma \circ j \circ i = g \circ i$ at the end of the day.

Enumerate $M = \{m_k \mid k \in \omega\}$ and $N = \{n_k \mid k \in \omega\}$. At stage $2k + 1$, we extend the given partial isomorphism, σ_{2k} , to a partial isomorphism σ_{2k+1} which includes n_k in its range. We are given $\sigma_{2k}: A_{2k} \rightarrow A'_{2k}$ and $\tau_{2k}: B_{2k} \rightarrow B'_{2k}$. Choose $A'_{2k+1} \preceq N$ such that $B'_{2k} \preceq A'_{2k+1}$ and $n_k \in A'_{2k+1}$, and choose B'_{2k+1} witnessing weak- K -homogeneity for A'_{2k+1} (so $A'_{2k+1} \preceq B'_{2k+1} \preceq N$). Let $h: B_{2k} \rightarrow B'_{2k+1}$ be the composition of τ_{2k} with the inclusion $B'_{2k} \rightarrow B'_{2k+1}$.

Since $2k$ is even, B_{2k} witnesses weak- K -homogeneity for A_{2k} , and B'_{2k+1} is in the age of M , since M and N have the same K -age, so there exists a strong embedding $l: B'_{2k+1} \rightarrow M$ over A_{2k} . Let $A_{2k+1} = l(A'_{2k+1})$, $\sigma_{2k+1} = (l \upharpoonright A'_{2k+1})^{-1}$, $B_{2k+1} = l(B'_{2k+1})$, and $\tau_{2k+1} = l^{-1}$. Then $\sigma_{2k} \subseteq \sigma_{2k+1} \subseteq \tau_{2k+1}$ and B'_{2k+1} witnesses weak- K -homogeneity for A'_{2k+1} .

The even stage $2k + 2$ is similar: the roles of M and N are reversed, and we ensure that σ_{2k+1} includes m_k in its domain. \square

Corollary 3.5.3. *Let M be a countably infinite K -direct limit. M is weakly- K -homogeneous if and only if M is weakly- K -ultrahomogeneous, and for all $A \preceq B \preceq M$, B witnesses weak- K -homogeneity for A if and only if B witnesses weak- K -ultrahomogeneity for A . In particular, M is K -homogeneous if and only if M is K -ultrahomogeneous.*

Proof. Suppose M is weakly- K -homogeneous. Then for any $A \preceq M$, any B witnessing weak- K -homogeneity for A also witnesses weak- K -ultrahomogeneity for A . Indeed, for any strong embedding $g: B \rightarrow M$, taking $N = M$ in Theorem 3.5.2 (M and N clearly have the same K -age) gives an isomorphism $\sigma: M \cong N$, which is the desired automorphism of M .

Conversely, suppose M is weakly- K -ultrahomogeneous. Then for any $A \preceq M$, any B witnessing weak- K -ultrahomogeneity for A also witnesses weak- K -homogeneity for A . Indeed, for any strong embedding $g: B \rightarrow C$, with C in the K -age of M , there is some strong embedding $l: C \rightarrow M$. Then $l \circ g$ is a strong embedding $B \rightarrow M$, so by weak- K -

ultrahomogeneity, there is an automorphism σ of M such that $\sigma^{-1} \circ (l \circ g) \circ i = j \circ i$. Let $h = \sigma^{-1} \circ l$. This is a strong embedding $C \rightarrow M$, and $h \circ g \circ i = j \circ i$, as desired.

Finally, we have shown that every A in the K -age of M witnesses weak- K -homogeneity for itself if and only if every A in the K -age of M witnesses weak- K -ultrahomogeneity for itself. In other words, M is K -homogeneous if and only if M is K -ultrahomogeneous. \square

Corollary 3.5.4. *If there exists a countably infinite K -direct limit which is K -universal and weakly- K -homogeneous, it is unique up to isomorphism.*

Proof. Suppose M and N are countably infinite K -universal weakly- K -homogeneous K -direct limits. M and N have the same K -age, namely K . Let $A \preceq_K M$ be any finite strong substructure, and let B witness weak- K -homogeneity for A . Since N is K -universal, there is a strong embedding $B \rightarrow N$, so by Theorem 3.5.2, M and N are isomorphic. \square

We now connect universality and (weak) homogeneity to the joint embedding and (weak) amalgamation properties.

Theorem 3.5.5. *Let K be an extendible strong embedding class. The following are equivalent:*

- (1) K is countable up to isomorphism, and K has the joint embedding property.
- (2) K has a cofinal chain subclass.
- (3) There is a countable (possibly finite) K -direct limit which is K -universal.

Proof. (1) \rightarrow (2): Enumerate the isomorphism classes in K as $\{A_i \mid i \in \omega\}$. If K is finite up to isomorphism, simply enumerate the classes with repetitions. We build a chain $C_0 \preceq C_1 \preceq C_2 \preceq \dots$ by induction. Let $C_0 = A_0$. Given C_n , use the joint embedding property to find some C_{n+1} in K such that $C_n \preceq C_{n+1}$ and $f_{n+1}: A_{n+1} \rightarrow C_{n+1}$ is a strong embedding. Let K' be the full subclass of K with structures $S_{K'} = \{A \in K \mid A \cong C_i \text{ for some } i \in \omega\}$.

Clearly K' is a chain class. We show that K' is cofinal in K . Given B in K , B is isomorphic to A_i for some $i \in \omega$. Then B embeds strongly in C_i , composing the isomorphism $B \cong A_i$ with the embedding $f_i: A_i \rightarrow C_i$.

(2) \rightarrow (3): Let K' be the cofinal chain subclass, with $C_0 \preceq_{K'} C_1 \preceq_{K'} C_2 \preceq_{K'} \dots$ the witnessing chain. The inclusions in the chain are K -embeddings, since K' is a subclass of K . Let $M = \varinjlim(C_i)$ be the direct limit of the chain. M is a countable K -direct limit, and given $A \in K$, A embeds strongly in some C_i , by cofinality of K' , so M is K -universal.

(3) \rightarrow (1): Let M be the countable K -universal K -direct limit. Since every structure in K is isomorphic to a strong substructure of M , and M has only countably many finite subsets, K is countable up to isomorphism. For the joint embedding property, let A and B be in K . We may identify A and B with strong substructures of M . By Proposition 3.2.21, we may find $C \preceq M$ such that $A \preceq C$ and $B \preceq C$, as desired. \square

Theorem 3.5.6. *Let K be an extendible strong embedding class. If there is a countable (possibly finite) K -direct limit M which is K -universal and weakly K -homogeneous, then K is a generalized Fraïssé class. If, additionally, M is K -homogeneous, then K has the amalgamation property.*

Proof. K is an extendible strong embedding class by assumption, and by Theorem 3.5.5, the existence of a countable universal K -direct limit implies that K is countable up to isomorphism and has the joint embedding property. It remains to verify the weak amalgamation property.

Let A be in K . Identifying A with a strong substructure of M , we can find $A \preceq B \preceq M$ such that B witnesses weak- K -homogeneity for A . We will show that the inclusion $f: A \rightarrow B$ is an amalgamation embedding. Indeed, let $g_1: B \rightarrow C_1$ and $g_2: B \rightarrow C_2$ be strong embeddings. By weak- K -homogeneity, we can embed C_1 and C_2 into M over A , by $h_1: C_1 \rightarrow M$ and $h_2: C_2 \rightarrow M$. By Proposition 3.2.21, there is $D \preceq_K M$ such that $h_1(C_1) \preceq_K D$ and $h_2(C_2) \preceq_K D$. Then D amalgamates C_1 and C_2 over A .

If M is K -homogeneous, then we could take $B = A$, and the same argument shows that K has the amalgamation property. \square

The converse is true, but we will delay giving the proof (Theorem 4.2.2) until after we have discussed generic constructions in Section 4.1.

Chapter 4

Generic limits

4.1 The space Dir_K and genericity

Throughout this section, let K be an extendible strong embedding class. We will strengthen this standing assumption below in Convention 4.1.9. We describe a space Dir_K of countably infinite K -direct limits, in analogy with the space Str_L of countably infinite L -structures studied in Part I.

Definition 4.1.1. A **labeled structure** is a (finite or infinite) structure whose domain is a subset of ω .

If we have a directed system (A_i) of labeled structures, and all the connecting maps are inclusions, then the underlying set of $M = \varinjlim(A_i)$ is the union of the underlying sets of the A_i so M is a labeled structure. Conversely, if M is a labeled K -direct limit, then M is the direct limit of its strong substructures in K , as in Proposition 3.2.21, i.e. a direct limit of finite labeled structures.

Definition 4.1.2. Dir_K is the space of K -direct limits with domain ω . The topology on Dir_K is generated by open sets $U_A = \{M \in \text{Dir}_K \mid A \preceq M\}$, for all finite labeled structures A in K .

Note that the domain of an infinite labeled K -direct limit M may be a proper subset of ω . Then M is not a point of Dir_K . The points of Dir_K are, in particular, L -structures with domain ω , so $\text{Dir}_K \subseteq \text{Str}_L$. The topology on Dir_K is typically strictly finer than the topology on Str_L (see Section 1.2), though they agree in the setting of classical Fraïssé theory (Proposition 4.1.4 (5)).

Definition 4.1.3. K is **semi-definable** if the topology on Dir_K agrees with the subspace topology inherited from Str_L .

The notion of a semi-definable class generalizes the notion of a separable class introduced by Baldwin, Koerwein, and Laskowski in [11]; the name is justified in Proposition 4.1.4 (4).

It is a much weaker condition than other common definability conditions, such as smoothness in the sense of Kueker and Laskowski (Remark 3.2.15), but it is strong enough to have some model-theoretic consequences about K -direct limits, as in Theorem 4.4.6 below.

We now establish some basic properties of the topology on Dir_K .

Proposition 4.1.4. *Let K be an extendible strong embedding class, and let \mathcal{T} be the topology on Dir_K .*

- (1) $\mathcal{U} = \{U_A \mid A \text{ is a finite labeled structure in } K\}$ is a basis for \mathcal{T} . That is, every finite intersection of sets in \mathcal{U} is equal to a union of sets in \mathcal{U} .
- (2) \mathcal{T} is Hausdorff and finer than the subspace topology inherited from Str_L .
- (3) If K is coherent, then the sets in \mathcal{U} are clopen.
- (4) K is semi-definable if and only if for all A in K and every infinite K -direct limit M with $A \preceq M$, there is a first-order quantifier-free formula $\psi_A(\bar{x}, \bar{y})$ such that $M \models \exists \bar{y} \psi(\bar{a}, \bar{y})$, where \bar{a} is a tuple enumerating A , and whenever N is an infinite K -direct limit, if $N \models \exists \bar{y} \psi_A(\bar{b}, \bar{y})$, then the map $a_i \mapsto b_i$ is a strong embedding $A \rightarrow N$.
- (5) If K is a small hereditary class, then K is semi-definable, and Dir_K is a closed subspace of Str_L .

Proof. (1) Let $V = \bigcap_{i=1}^m U_{A_i}$, where the A_i are finite labeled structures in K . It suffices to show that if $M \in V$, then there is some set $U_B \in \mathcal{U}$ such that $M \in U_B \subseteq V$. By Proposition 3.2.21, there is some finite $B \preceq M$ such that $A_i \preceq B \preceq M$ for all i . Then $M \in U_B$, and $U_B \subseteq V$, since if $N \in U_B$, we have $A_i \preceq B \preceq N$ for all i .

(2) We show that \mathcal{T} is finer than the subspace topology. If $[\varphi(\bar{a})]$ be a basic open set in Str_L (so $\varphi(\bar{x})$ is an atomic or negated atomic formula), we must show that $[\varphi(\bar{a})] \cap \text{Dir}_K$ is open in Dir_K . Let $M \in [\varphi(\bar{a})] \cap \text{Dir}_K$, and find a finite $B \preceq M$ such that B contains \bar{a} . For any $N \in U_B$, $N \models \varphi(\bar{a})$, so U_B is an open neighborhood of M contained in $[\varphi(\bar{a})] \cap \text{Dir}_K$. It follows that \mathcal{T} is Hausdorff, since a subspace of a Hausdorff space is Hausdorff, and any topology which is finer than a Hausdorff topology is Hausdorff.

(3) Suppose K is coherent. Let A be a finite labeled structure in K with underlying set $X \subseteq \omega$. We must show that $V = \text{Dir}_K \setminus U_A$ is open. Let $M \in V$, and choose a finite $B \preceq M$ with $X \subseteq B$. Then $A \not\preceq B$, since otherwise we have $A \preceq B \preceq M$, and $M \in U_A$. For any other $N \in U_B$, $N \notin U_A$, since otherwise we have $A \preceq N$ and $B \preceq N$, so $A \preceq B$ by coherence. Hence U_B is an open neighborhood of M contained in V .

(4) We have already shown in (2) that \mathcal{T} is finer than the subspace topology. The converse holds if and only if every basic open set U_A in Dir_K (corresponding to a finite labeled structure A in K) is open in the subspace topology. And this is true if and only if every point $M \in U_A$ has a basic open neighborhood in the subspace topology contained in U_A .

Such a neighborhood is given by the intersection of a basic clopen set in Str_L with Dir_K , and such a basic clopen set has the form $\llbracket \psi_A(\bar{a}, \bar{c}) \rrbracket$ for $\psi_A(\bar{x}, \bar{y})$ first-order quantifier-free. Here, without loss of generality, we may ensure that \bar{a} is a tuple enumerating A and that ψ_A is explicitly non-redundant.

So assume we have this topological condition. Let A' be in K , and let M' be an infinite K -direct limit with $A' \preceq M'$. Let M'' be a countably infinite K -direct limit with $A' \preceq M'' \preceq M'$ (by Proposition 3.3.2 (6)). By isomorphism invariance, we may assume that the domain of M'' is ω , so $M'' \in \text{Dir}_K$ and A' is a finite labeled structure in K . Let $\llbracket \psi_{A'}(\bar{a}, \bar{c}) \rrbracket$ be our open neighborhood of M'' in the subspace topology contained in $U_{A'}$. Then $M'' \models \exists \bar{y} \psi_{A'}(\bar{a}, \bar{y})$, and since M'' is a substructure of M' and $\psi_{A'}$ is quantifier-free, also $M' \models \exists \bar{y} \psi_{A'}(\bar{a}, \bar{y})$. Now let N be any infinite K -direct limit such that $N \models \exists \bar{y} \psi_{A'}(\bar{b}, \bar{y})$, so $N \models \psi_{A'}(\bar{b}, \bar{d})$ for some tuple \bar{d} . Let $N' \preceq N$ be a countably infinite K -direct limit containing \bar{b} and \bar{d} . Again, by isomorphism invariance, we may assume that the domain of N' is ω and that $\bar{b} = \bar{a}$ and $\bar{d} = \bar{c}$. Then $N' \in \llbracket \psi_{A'}(\bar{a}, \bar{c}) \rrbracket \subseteq U_{A'}$, so $A' \preceq N' \preceq N$.

Conversely, suppose we have the condition in (4). Let A be a finite labeled structure in K enumerated by the tuple \bar{a} , and let $M \in U_A$. By assumption there is a formula $\psi(\bar{x}, \bar{y})$ such that $M \models \exists \bar{y} \psi(\bar{a}, \bar{y})$, and whenever N is an infinite K -direct limit, if $N \models \exists \bar{y} \psi(\bar{b}, \bar{y})$, then the map $a_i \mapsto b_i$ is a strong embedding $A \rightarrow N$. Let \bar{c} be such that $M \models \psi(\bar{a}, \bar{c})$. Then if $N \in \llbracket \psi(\bar{a}, \bar{c}) \rrbracket$, $N \models \exists \bar{y} \psi(\bar{a}, \bar{y})$, so $A \preceq N$, and $N \in U_A$. Hence $\llbracket \psi(\bar{a}, \bar{c}) \rrbracket$ is a neighborhood of M in the subspace topology contained in U_A .

- (5) Hereditary classes are, of course, extendible (Remark 3.2.13). We show that the condition in (4) holds, and we can even take \bar{y} to be the empty tuple of variables. Let \bar{a} be a tuple enumerating the domain of A , and let $\phi_A(\bar{x})$ be the quantifier-free formula provided by Lemma 3.1.7, so $A \models \varphi_A(\bar{a})$. If $A \preceq M$, then also $M \models \psi_A(\bar{a})$, and if $N \models \psi_A(\bar{b})$ then $B = \llbracket \bar{b} \rrbracket$ is in K (K is hereditary), so the map $a_i \mapsto b_i$ is an isomorphism $A \rightarrow B$, and hence an embedding $A \rightarrow N$, and all embeddings are strong. In other words, $U_A = \llbracket \psi_A(\bar{a}) \rrbracket \cap \text{Dir}_K$.

If K is hereditary, then a structure $M \in \text{Str}_L$ is a K -direct limit if and only if every finite substructure of M is in K if and only if M is a model of the universal theory $T_K = \{\forall \bar{x} \neg \varphi_A(\bar{x}) \mid A \notin S_K\}$, where $\varphi_A(\bar{x})$ is formula from Lemma 3.1.7. Indeed, if M is a K -direct limit, then for every tuple \bar{b} from M , $B = \llbracket \bar{b} \rrbracket$ is in K , so $B \models \neg \varphi_A(\bar{b})$ for all A not in K , and M agrees. And if M is not a K -direct limit, then there is some tuple \bar{b} from M such that $B = \llbracket \bar{b} \rrbracket$ is not in K , and hence $B \models \exists \bar{x} \varphi_B(\bar{x})$, and M agrees.

Each axiom of T_K picks out a closed subset of Str_L (the intersection over all \bar{a} of the clopen sets $\llbracket \neg \varphi_A(\bar{a}) \rrbracket$), so Dir_K is closed in Str_L . \square

To avoid trivialities when studying infinite K -direct limits, we make some further restrictions on the class K .

Definition 4.1.5. A strong embedding class K is **pruned** if, for every A in K , there is a proper strong embedding $f: A \rightarrow B$ (i.e. f is not an isomorphism).

Proposition 4.1.6. *Let K be an extendible strong embedding class. The following are equivalent:*

- (1) K is pruned.
- (2) Every A in K embeds strongly in an infinite K -direct limit.
- (3) Every basic open set U_A in Dir_K is non-empty.

Proof. (1) \rightarrow (2): Let $A_0 = A$. For all i , given A_i , let $f_i: A_i \rightarrow A_{i+1}$ be a proper strong embedding, so $|A_{i+1}| > |A_i|$. Then $M = \varinjlim(A_i)$ is an infinite K -direct limit in which A embeds strongly.

(2) \rightarrow (3): Let A be a finite labeled structure in K . Then there is a strong embedding $A \rightarrow M$ with M infinite. Since K_∞ has Löwenheim–Skolem number \aleph_0 (Proposition 3.3.2 (6)), we may assume M is countably infinite, and replacing M with an isomorphic structure over A , we may assume that the domain of M is ω and $A \preceq M$. Then $M \in U_A$.

(3) \rightarrow (1): Let B be in K with $|B| = n$. Then B is isomorphic to a finite labeled structure A with domain $[n]$. Let $M \in U_A$, and find a finite $C \preceq M$ such that $A \preceq C$ and $n \in C$ (so A is a proper strong substructure of C). Then composing the isomorphism $B \rightarrow A$ with the inclusion $A \rightarrow C$ gives a proper strong embedding $B \rightarrow C$. \square

Remark 4.1.7. Every extendible strong embedding class K contains a canonical full pruned subclass K_p , consisting of those structures A in K such that A embeds strongly in an infinite K -direct limit, together with all K -embeddings between them. Note that every infinite K -direct limit $M \cong \varinjlim(A_i)$ is also an infinite K_p -direct limit, since the directed system (A_i) in \mathcal{C}_K is also a directed system in \mathcal{C}_{K_p} .

Remark 4.1.8. If K has the joint embedding property and contains arbitrarily large finite structures, then K is pruned. Indeed, for any A in K , we can pick B in K with $|B| > |A|$ and strong embeddings $f: A \rightarrow C$ and $g: B \rightarrow C$. Then $|C| \geq |B| > |A|$, so f is proper.

Convention 4.1.9. For the remainder of this chapter, we assume that K is a non-empty pruned extendible strong embedding class.

In particular, there are infinite K -direct limits, so Dir_K is non-empty. We will study genericity in the sense of Baire category in Dir_K (See [55, Section I.8]). Recall that a set is meager (Baire called meager sets “first category”) if it is a countable union of nowhere dense sets. The meager sets form a σ -ideal: they are closed under subset and countable union. A set is comeager if its complement is meager; equivalently, if it contains a countable intersection of dense open sets. We view meager sets as very small, in analogy with the null sets for a probability measure, and we the comeager sets as very large, or “generic”, in analogy with the measure 1 sets. The fact that the meager sets form a σ -ideal plays the role

of countable additivity for a measure. As a warm-up, we show directly that Dir_K is a Baire space: every comeager set is dense.

Theorem 4.1.10. *Dir_K is a Baire space.*

Proof. If X is comeager, then $X \supseteq \bigcap_{i \in \omega} D_i$, where $\langle D_i \rangle_{i \in \omega}$ is a countable family of dense open sets. To show that X is dense, let U_A be any basic open set. We build a sequence $A = A_0 \preceq A_1 \preceq A_2 \preceq \dots$ of finite labeled structures by induction, such that $\varinjlim(A_i) = M$ and $M \in U_A \cap \bigcap_{i \in \omega} D_i \subseteq U_A \cap X$. To start, set $A_0 = A$. Given A_i , consider the set D_i in our family. Since D_i is dense and U_{A_i} is non-empty (K is pruned), there is some $N_i \in U_{A_i} \cap D_i$. Since D_i is open, there is a basic open set U_{B_i} such that $N_i \in U_{B_i} \subseteq D_i$. Pick a strong substructure $A_{i+1} \preceq N_i$ such that $A_i \preceq A_{i+1}$, $B_i \preceq A_{i+1}$, and $i \in A_{i+1}$. Then the domain of $M = \varinjlim(A_i)$ is all of ω , $A \preceq M$, so $M \in U_A$, and $B_i \preceq A_{i+1} \preceq M$, so $M \in U_{B_i} \subseteq D_i$. \square

Let P be a property of labeled K -direct limits. We prefer to be flexible about what counts as a property, usually just giving a description in words. Formally, a property can be identified with the set of labeled K -direct limits satisfying it, and, conversely, for any set X of labeled K -direct limits, we can consider the property of being in X . Note that a property may be satisfied by finite as well as infinite labeled K -direct limits, and properties need not be invariant under isomorphisms of (unlabeled) L -structures. For example, for $n \in \omega$, “ n is in the domain” is a property of labeled K -direct limits.

Given a property P , we define a game $G_K(P)$ for two players. We follow standard terminology for infinite games (see, for example, [55, Definition I.8.10]). Players I and II take turns playing finite labeled structures in K in a chain $A_0 \preceq A_1 \preceq A_2 \preceq \dots$, so Player I chooses A_i when i is even, and Player II chooses A_i when i is odd. Given a run of the game $(A_i)_{i \in \omega}$, we let $M = \varinjlim(A_i)$. Player II wins if M satisfies P , while Player I wins if M does not satisfy P .

Definition 4.1.11. The property P is **generic** if Player II has a winning strategy in the game $G_K(P)$.

Notation 4.1.12. If P is a property of labeled K -direct limits, define $\llbracket P \rrbracket = \{M \in \text{Dir}_K \mid M \text{ satisfies } P\}$.

Given our standing assumption that K is pruned, the game $G_K(P)$ is really just a convenient rephrasing of the Banach–Mazur game $G^{**}(A, X)$ with target set $A = \llbracket P \rrbracket$ on the space $X = \text{Dir}_K$ (see [55, Section I.8.H]). In this game, Players I and II take turns playing non-empty open sets in a descending chain $V_0 \supseteq V_1 \supseteq V_2 \supseteq \dots$, Player II wins if $\bigcap_{i \in \omega} V_i \subseteq A$, and Player I wins otherwise. It is a basic fact that Player II has a winning strategy in $G^{**}(A, X)$ if and only if A is comeager in the space X [55, Theorem I.8.33].

Theorem 4.1.13. *For any property P , the game $G_K(P)$ is equivalent to the Banach–Mazur game $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$. That is, a player has a winning strategy in $G_K(P)$ if and only if the same player has a winning strategy in $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$. Consequently, P is generic if and only if $\llbracket P \rrbracket$ is comeager in Dir_K .*

Proof. We call the players Player Y and Player Z (so Y and Z are I and II, in some order).

Suppose Player Z has a winning strategy in $G_K(P)$. Players Y and Z play a game of $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$ with run $(V_i)_{i \in \omega}$, and Player Z simultaneously simulates a game of $G_K(P)$ with run $(A_i)_{i \in \omega}$. Player Z ensures that $U_{A_i} \subseteq V_i$ for all i , that $U_{A_k} = V_k$ if turn k is their turn, and that $[k] \subseteq A_k$ if turn k is their opponent's turn.

On their turn, Player Z follows their winning strategy in the simulated $G_K(P)$, playing a finite labeled structure A_k in K , and also plays $V_k = U_{A_k}$ in $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$. Note that V_k is non-empty, since K is pruned. If it is turn 0, this a legal move. If not, then since $A_{k-1} \preceq A_k$, we have $V_k = U_{A_k} \subseteq U_{A_{k-1}} \subseteq V_{k-1}$, so this is a legal move.

On the opponent's turn, Player Y plays some non-empty open set V_k . If it is turn 0, Player Z finds a basic open set $U_{A_0} \subseteq V_0$ and records A_0 as Player Y's move in the simulated $G_K(P)$. If it is not turn 0, then $V_k \subseteq V_{k-1} = U_{A_{k-1}}$. Player Z picks a point $M \in V_k$ and a basic open neighborhood $M \in U_{B_k} \subseteq V_k$. Then since $A_{k-1} \preceq M$ and $B_k \preceq M$, there is some $A_k \preceq M$ such that $A_{k-1} \preceq A_k$, $B_k \preceq A_k$, and $[k] \subseteq A_k$. Note that $U_{A_k} \subseteq U_{B_k} \subseteq V_k$. Player Z records A_k as Player Y's move in the simulated $G_K(P)$.

Now since $[k] \subseteq A_k$ and $U_{A_{k+1}} = V_{k+1}$ for cofinally many k , the domain of $M = \varinjlim(A_i)$ is all of ω , so $M \in \text{Dir}_K$, and $\bigcap_{i \in \omega} V_i = \bigcap_{i \in \omega} U_{A_i} = \{M\}$. Since Player Z followed a winning strategy in $G_K(P)$, M satisfies P if Z is II ($\{M\} \subseteq \llbracket P \rrbracket$), and M does not satisfy P if Z is I ($\{M\} \not\subseteq \llbracket P \rrbracket$). So this is a winning strategy for Player Z in $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$.

Conversely, suppose Player Z has a winning strategy in $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$. Players Y and Z play a game of $G_K(P)$ with run $(A_i)_{i \in \omega}$, and Player Z simultaneously simulates a game of $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$ with run $(V_i)_{i \in \omega}$. Player Z ensures that $U_{A_i} \subseteq V_i$ for all i , that $U_{A_k} = V_k$ if turn k is their opponent's turn, and that $[k] \subseteq A_k$ if turn k is their turn.

On their turn, Player Z follows their winning strategy in the simulated $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$, playing a non-empty open set V_k . If it is turn 0, they find some basic open set $U_{A_0} \subseteq V_0$ and play A_0 in $G_K(P)$. If it is not turn 0, then $V_k \subseteq V_{k-1} = U_{A_{k-1}}$. Player Z picks a point $M \in V_k$ and a basic open neighborhood $M \in U_{B_k} \subseteq V_k$. Then $A_{k-1} \preceq M$, and $B_k \preceq M$, there is some $A_k \preceq M$ such that $A_{k-1} \preceq A_k$, $B_k \preceq A_k$, and $[k] \subseteq A_k$. Note that $U_{A_k} \subseteq U_{B_k} \subseteq V_k$. Player Z plays A_k in $G_K(P)$.

On the opponent's turn, Player Y plays some finite labeled structure A_k in K . Player Z records $V_k = U_{A_k}$ as Player Y's move in the simulated $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$. Note that V_k is non-empty, since K is pruned. If it is turn 0, this is automatically a legal move. If not, then since $A_{k-1} \preceq A_k$, we have $V_k = U_{A_k} \subseteq U_{A_{k-1}} \subseteq V_{k-1}$, so this is a legal move.

Now since $[k] \subseteq A_k$ for cofinally many k , the domain of $M = \varinjlim(A_i)$ is all of ω , and $\bigcap_{i \in \omega} V_i = \bigcap_{i \in \omega} U_{A_i} = \{M\}$. Since Player Z followed a winning strategy in $G^{**}(\llbracket P \rrbracket, \text{Dir}_K)$, $\{M\} \subseteq \llbracket P \rrbracket$ if Z is II (M satisfies P), and $\{M\} \not\subseteq \llbracket P \rrbracket$ if Z is I (M does not satisfy P). So this is a winning strategy for Player Z in $G_K(P)$.

The equivalence of genericity and comeagerness follows from the standard characterization of comeager sets via the Banach–Mazur game [55, Theorem I.8.33]. \square

Remark 4.1.14. Theorem 4.1.10 is actually a consequence of Theorem 4.1.13. Let X be a comeager set in Dir_K , and let P be the property $\llbracket P \rrbracket = X$. Then for any finite labeled

structure A in K , Player II can win the game $G_K(P)$ even if Player I's first move is A . That is, there is a structure $M \in U_A \cap X$. So X is dense.

Remark 4.1.15. Comeager sets are closed under countable intersection, by definition. So a consequence of Theorem 4.1.13 is that a countable conjunction of generic properties is generic. Game-theoretically, this corresponds to the fact that Player II can “interleave” countably many winning strategies in order to force the direct limit to satisfy countably many generic properties. But the topological formulation in terms of comeager sets allows us to avoid describing such interleaved strategies explicitly.

Remark 4.1.16. In reducing genericity of properties of labeled K -direct limits (whose domains may be proper subsets of ω , or even finite) to comeagerness in Dir_K , we are really using the fact that K is pruned, and hence that the property of having domain ω is generic. Indeed, if P is this property, then $\llbracket P \rrbracket$ is all of Dir_K . But no requirement that the direct limit $\varinjlim(A_i)$ has domain ω is built into the game $G_K(P)$, so we don't have to worry about this when describing strategies explicitly. This is taken care of for us in the proof of Theorem 4.1.13.

4.2 The generic limit

Convention 4.1.9 remains in effect in this section (and throughout this chapter). Our new language of genericity allows us to formulate a strong converse to Theorem 3.5.6. It was the goal of characterizing condition (4) in the theorem below that led Ivanov [48] and Kechris and Rosendal [56] to identify the weak amalgamation property.

Notation 4.2.1. For any K -direct limit M , we write Iso_M for the property “isomorphic to M ”.

Theorem 4.2.2. *The following are equivalent:*

- (1) K is a generalized Fraïssé class.
- (2) K -universality and weak- K -homogeneity are generic properties.
- (3) There exists $M \in \text{Dir}_K$ such that M is K -universal and weakly- K -homogeneous.
- (4) There exists $M \in \text{Dir}_K$ such that Iso_M is a generic property.

Proof. (1) \rightarrow (2): Let K be a generalized Fraïssé class, and let P be the property “universal and weakly- K -homogeneous”. We will show that P is implied by a countable conjunction of properties, and we will show that each of these properties is generic by describing a winning strategy for Player II in the relevant game. This suffices by Remark 4.1.15.

- Let AE be the property of a labeled K -direct limit M that for every strong substructure $A \preceq M$, there is some B such that $A \preceq B \preceq M$ and the inclusion $A \rightarrow B$ is an amalgamation embedding.

We describe a winning strategy for Player II in $G_K(\text{AE})$. Suppose Player I has just played A_k (so k is even), Player II picks an amalgamation embedding $A_k \rightarrow B$ in K (K has the weak amalgamation property). Since amalgamation embeddings are closed under isomorphism, we may assume that B is a finite labeled structure and the embedding is an inclusion. Player II plays $A_{k+1} = B$.

Given a run (A_i) with this strategy, let $M = \varinjlim(A_i)$, and let $A \preceq M$ be any strong substructure. Then there is some k such that $A \preceq A_k$, and we may assume that k is even. The inclusion $A_k \rightarrow A_{k+1}$ is an amalgamation embedding, and since amalgamation embeddings are stable under composition (Proposition 3.2.6), the inclusion $A \rightarrow A_{k+1}$ is also an amalgamation embedding. So M satisfies AE.

- For all B in K , let Emb_B be the property of a labeled K -direct limit M that there exists a strong embedding $B \rightarrow M$. Note that if $B \cong D$, then M satisfies Emb_B if and only if it satisfies Emb_D .

We describe a winning strategy for Player II in $G_K(\text{Emb}_B)$. Suppose that on turn 0, Player I plays A_0 . Pick strong embeddings $f: A_0 \rightarrow C$ and $g: B \rightarrow C$ for some C in K (K has the joint embedding property). By composing these embeddings with an isomorphism, we may assume that C is a finite labeled structure and $f: A_0 \rightarrow C$ is an inclusion. Player II plays $A_1 = C$ and plays arbitrarily thereafter.

Given a run (A_i) with this strategy, let $M = \varinjlim(A_i)$. B embeds strongly in A_1 , and hence in M .

- Let $A \preceq B \preceq C$ be finite labeled structures such that the inclusion $i: A \rightarrow B$ is an amalgamation embedding. We also name the inclusion $j: B \rightarrow C$. Let $\text{WH}_{A,B,C}$ be the property of a labeled K -direct limit M that $A \preceq B \preceq M$ and that there is a strong embedding $f: C \rightarrow M$ such that $f \circ j \circ i$ is equal to the inclusion of A in M .

We describe a winning strategy for Player II in $G_K(\text{WH}_{A,B,C})$. Suppose Player I has just played A_k (so k is even). If $B \not\preceq A_k$, Player II plays arbitrarily. On the other hand, if $B \preceq A_k$ with inclusion j' , there is some structure D and strong embeddings $h_1: A_k \rightarrow D$ and $h_2: C \rightarrow D$ such that $h_1 \circ j' \circ i = h_2 \circ j \circ i$. We may assume that D is a finite labeled structure and $h_1: A_k \rightarrow D$ is an inclusion. Player II plays $A_{k+1} = D$ and plays arbitrarily thereafter.

Given a run (A_i) with this strategy, let $M = \varinjlim(A_i)$. Suppose that $A \preceq B \preceq M$. Then there is some k such that $B \preceq A_k$, and we may assume that k is even. Then we have a strong embedding $h_2: C \rightarrow A_{k+1}$, and hence a strong embedding $f: C \rightarrow M$, such that $f \circ j \circ i$ is equal to the inclusion of A in M .

Since K is countable up to isomorphism, we can enumerate the isomorphism types of finite structures in K as $\{B_i \mid i \in \omega\}$, and there are only countably many chains of finite labeled structures $A \preceq B \preceq C$ in K . So we have a countable list of generic properties: AE, Emb_{B_i} for all $i \in \omega$, and all the properties $\text{WH}_{A,B,C}$. It remains to show that if M satisfies all of these properties, then it is K -universal and weakly- K -homogeneous.

K -universality is taken care of by the properties $\{\text{Emb}_{B_i} \mid i \in \omega\}$. For weak- K -homogeneity, suppose $A \preceq M$. By AE, there exists $A \preceq B \preceq M$ such that the inclusion $i: A \rightarrow B$ is an amalgamation embedding. We claim that B witnesses weak- K -homogeneity for A . Suppose we have a strong embedding $j: B \rightarrow C$ with C in K . We may assume that C is a finite labeled structure and that j is an inclusion. Then by $\text{WH}_{A,B,C}$, there is a strong embedding $f: C \rightarrow M$ such that $f \circ j \circ i$ is equal to the inclusion of A in M .

(2) \rightarrow (3): Since K is pruned and non-empty, Dir_K is non-empty, and any comeager set in Dir_K is non-empty. So there is a countable K -direct limit M satisfying the generic properties of K -universality and weak- K -homogeneity.

(3) \rightarrow (1): This is Theorem 3.5.6.

(2) \rightarrow (4): Let P be the property “ K -universal and weakly- K -homogeneous”. Pick some $M \in \llbracket P \rrbracket$. By Corollary 3.5.4, any other $N \in \llbracket P \rrbracket$ is isomorphic to M , so $\llbracket P \rrbracket = \llbracket \text{Iso}_M \rrbracket$, and Iso_M is also generic by Theorem 4.1.13.

(4) \rightarrow (1): Let $M \in \text{Dir}_K$ be such that Iso_M is generic, so Player II has a winning strategy in the game $G_K(\text{Iso}_M)$. We must show that K is countable up to isomorphism and has the joint embedding and weak amalgamation properties.


First, we claim that M is K -universal. Then the existence of M implies, by Theorem 3.5.5, that K is countable up to isomorphism and has the joint embedding property. If M is not K -universal, there is some A in K such that A does not embed strongly in M . In $G_K(\text{Iso}_M)$, Player I may play $A_0 = A$ and play arbitrarily thereafter. Then A embeds strongly in $\varinjlim(A_i)$, so this is a winning strategy for Player I, and a contradiction.

Now assume that K does not have the weak amalgamation property. Then there is a structure A in K such that no strong embedding $A \rightarrow B$ is an amalgamation embedding. We may assume that A is a labeled structure. Let Players I and II play two copies of the game $G_K(\text{Iso}_M)$ at once, so on their turn, a player makes a move in both games. At the end of the joint game, we have two chains $(A_i)_{i \in \omega}$ and $(A'_i)_{i \in \omega}$ and two labeled K -direct limits $N = \varinjlim(A_i)$ and $N' = \varinjlim(A'_i)$. We will show that Player I has a strategy in the joint game which ensures $N \not\cong N'$.

Player I begins by playing $A_0 = A'_0 = A$. On each subsequent turn for Player I, Player II has just played two finite labeled structures A_k and A'_k . For each strong embedding $f: A \rightarrow A'_k$ (of which there are only finitely many, and at least one), Player I adds f to the end of a queue.

Player I then pulls an embedding $f': A \rightarrow A'_l$ from the front of the queue. Note that $l \leq k$. Composing f' with the inclusion $A'_l \rightarrow A'_k$, we have a strong embedding $f_2: A \rightarrow A'_k$.

We also have the strong inclusion $f_1: A \rightarrow A_k$. Consider the following diagram:



If there do not exist B in K and strong embeddings g_1 and g_2 making the above diagram commute, then Player I plays arbitrarily. But if some B , g_1 , and g_2 exist, then we have a single embedding $f: A \rightarrow B$ coming from the diagram. By our assumption, f is not an amalgamation embedding. Hence there exist strong embeddings $h_1: B \rightarrow C_1$ and $h_2: B \rightarrow C_2$ in K which cannot be amalgamated over A . We may assume that C_1 and C_2 are labeled structures and that the embeddings $h_1 \circ g_1: A_k \rightarrow C_1$ and $h_2 \circ g_2: A'_k \rightarrow C_2$ are inclusions. Player I plays $A_{k+1} = C_1$ and $A'_{k+1} = C_2$.

Now consider a play of the joint game in which Player II follows a winning strategy for $G_K(\text{Iso}_M)$ in both games and Player I follows the strategy described above. Two labeled K -direct limits N and N' are produced, and Player II's strategy ensures that both N and N' are isomorphic to M , so there is an isomorphism $\sigma: N \cong N'$.

Thanks to Player I's first turn, $A \preceq N$, and hence $\sigma(A) \preceq N'$. Then $\sigma(A) \preceq A'_l$ for some l , and we may assume l is odd. Player I added the strong embedding $f' = (\sigma \upharpoonright A): A \rightarrow A'_l$ to the queue on turn l , and at some turn $k \geq l$, Player I pulled f' from the front of the queue.

On this turn, Player I considered the diagram above, where f_1 is the inclusion $A \rightarrow A_k$ and f_2 is the composition of f' with the inclusion $A'_l \rightarrow A'_k$, so f_2 agrees with σ on A . Since $A_k \preceq N$, $\sigma(A_k) \preceq N'$, and there is a finite $B \preceq N'$ such that $\sigma(A_k) \preceq B$ and $A'_k \preceq B$. Taking $g_1 = (\sigma \upharpoonright A_k): A_k \rightarrow B$ and g_2 the inclusion $A'_k \rightarrow B$ makes the diagram commute (both maps $A \rightarrow B$ agree with σ on A). Then Player I played $A_{k+1} = C_1$ and $A'_{k+1} = C_2$ which cannot be amalgamated over A . But we can use σ to amalgamate A_{k+1} and A'_{k+1} over A exactly as we used it to amalgamate A_k and A'_k over A . This is a contradiction. \square

Definition 4.2.3. Let K be a generalized Fraïssé class. We call a K -direct limit M the **generic limit** of K if Iso_M is generic. If K is a Fraïssé class, then M is called the **Fraïssé limit** of K .

Remark 4.2.4. Since comeager sets have non-empty intersection in Dir_K , the generic limit of K is unique up to isomorphism, and it is also the unique K -universal and weakly- K -homogeneous K -direct limit up to isomorphism.

Remark 4.2.5. The proofs of Theorem 3.5.6 and Theorem 4.2.2 show that if M is the generic limit K , and $A \preceq B \preceq M$, then B witnesses weak- K -homogeneity for A if and only if the inclusion $A \rightarrow B$ is an amalgamation embedding. In particular, K has the amalgamation property if and only if M is K -homogeneous.

Theorem 4.2.6. *If K is a generalized Fraïssé class and K' is a cofinal subclass of K , then K' is also a generalized Fraïssé class, and the generic limits of K and K' are isomorphic.*

Proof. Let M be the generic limit of K . By Theorem 4.2.2 and Remark 4.2.4, it suffices to show that M is a K' -direct limit and that M is K' -universal and weakly- K' -homogeneous.

First, we claim that if $A \preceq_K M$, then there is some C in K' such that $A \preceq_K C \preceq_K M$. Indeed, there is some $A \preceq_K B \preceq_K M$ such that B witnesses weak- K -homogeneity for A , and since K' is cofinal in K , there exists C' in K' and a K' -embedding $B \rightarrow C'$. K' embeddings are also K embeddings, so C' K -embeds in M over A , and we can take C to be the image of C' under this embedding.

Now let (A_i) be the directed system of all K -strong substructures of M together with all K -strong inclusions between them, and let (B_j) be the subsystem consisting of the structures and inclusions in (A_i) which are also in K' . Then (B_j) is a directed system in K' . Indeed, for any $B \preceq_K M$ and $B' \preceq_K M$ such that B and B' are in K' , there is some $A \preceq_K M$ such that $B \preceq_K A$ and $B' \preceq_K A$. By the claim, there is some C in K' such that $A \preceq_K C \preceq_K M$, and hence, since K' is full, $B \preceq_{K'} C$ and $B' \preceq_{K'} C$.

For any $a \in M$, there is some $A \preceq_K M$ such that $a \in A$, and by the claim there is some $A \preceq_K C \preceq_K M$ with C in K' , so $a \in C$. Then (B_j) covers M , and $M = \varinjlim (B_j)$. Further, if A is in K' , then $A \preceq_K M$ if and only if $A \preceq_{K'} M$.

K' -universality follows immediately from K -universality, since K' is a full subclass of K .

For weak- K' -homogeneity, let $A \preceq_{K'} M$. Then there is some $A \preceq_K B \preceq_K M$ such that B witnesses weak- K -homogeneity for A . By the claim, there is some $B \preceq_K C \preceq_K M$ such that C is in K' . Then C witnesses weak- K -homogeneity for A (one can see this directly or use Remark 4.2.5 and the fact that amalgamation embeddings are closed under composition, Proposition 3.4.5). And C also witnesses weak- K' -homogeneity for A , since K' is a subclass of K and every K -embedding $D \rightarrow M$ for D in K' is also a K' -embedding. \square

Example 4.2.7. We revisit the classes from Example 3.1.9, which are all (pruned) generalized Fraïssé classes, and discuss their generic limits.

K_0 is the class of all finite graphs, with all embeddings. It is hereditary and has the amalgamation property, so it is even a Fraïssé class. Its Fraïssé limit is the random graph. A probabilistic construction of the random graph was discussed in Example 1.4.16.

K_1 is the class of all finite acyclic graphs of degree at most two, with all embeddings, and K_2 is the class of all connected graphs in K_1 , with all embeddings.

K_1 is hereditary, but it does not have the amalgamation property: Consider the graph A consisting of two disconnected vertices, v and w . We can embed A into graphs B_2 and B_3 in which v and w are connected by paths of length 2 and 3, respectively. But B_2 and B_3 cannot be amalgamated over A , since any way of doing so would introduce a cycle.

K_2 is not hereditary, but it does have the amalgamation property: A structure A in K_2 is just a chain of some length n . Picking one side of A to call the left, an embedding of A into B in K_2 just extends the left side of the chain by l vertices and the right side of the chain by r vertices. Given two such embeddings $A \rightarrow B_1$ and $A \rightarrow B_2$, we can amalgamate them

by extending A on the left by the maximum of l_1 and l_2 and on the right by the maximum of r_1 and r_2 .

Since K_2 is cofinal in K_1 , K_1 has the weak amalgamation property by Proposition 3.4.6. Any embedding of A in K_1 into a connected graph in K_2 is an amalgamation embedding. And by Theorem 4.2.6, K_1 and K_2 have the same generic limit, the infinite chain without endpoints:



This generic limit is K_2 -ultrahomogeneous, since any two embeddings of a finite connected chain are conjugate by an automorphism (slide, and possibly flip). But it is only weakly- K_1 -ultrahomogeneous, since the distances between the connected components of a graph in K_1 may be different for different embeddings. If A is a substructure of the generic limit in K_1 , then weak- K_1 -ultrahomogeneity is witnessed for A by the connected substructure B in K_2 obtained by including all the vertices lying between connected components of A .

Finally, K_3 is the class with the same structures as K_2 , but with strong embeddings $A \rightarrow B$ such that the image of every connected component of A is a connected component of B . We can always add more connected components, of any finite size, but connected components cannot grow once they have been added. Hence the generic limit of K_3 is a disjoint union of infinitely many chains of each finite length.

In the next section, we will study genericity in Dir_K for classes which are not generalized Fraïssé classes and explain some connections with Part I. But we will return to the generic limits of generalized Fraïssé classes in Sections 4.4 and 4.5.

4.3 Genericity in classes without generic limits

In Section 1.4, we found that we could attach a complete theory of $L_{\omega_1, \omega}$ to an ergodic structure, even if the ergodic structure is not almost surely isomorphic to a countable structure. Indeed, if μ is an ergodic structure, then every Borel set which is invariant for the logic action has measure 1 or measure 0. In this section, we seek to do the same for a strong embedding class, hoping that every reasonable isomorphism-invariant property is generic or cogenic, i.e. comeager or meager in Dir_K . It turns out that the joint embedding property is exactly what we need to make this work.

Definition 4.3.1. The **logic action** is the natural action of S_∞ on Dir_K . It is the restriction of the logic action of S_∞ on Str_L (Definition 1.2.3.)

The logic action is well-defined since the class of K -direct limits is closed under isomorphism, and it respects the topology on Dir_K , since the class of strong embeddings is closed under isomorphism. As in Section 1.2 the orbit of a structure M in Dir_K is the set of structures in Dir_K which are isomorphic to M , and a subset of Dir_K is invariant for the logic action if and only if it is closed under isomorphism in Dir_K .

Definition 4.3.2. Let P be a property of labeled K -direct limits.

- P is **open** if $\llbracket P \rrbracket$ is an open set in Dir_K , and P is **Borel** if $\llbracket P \rrbracket$ is a Borel set in Dir_K .
- P is **invariant** if $\sigma(\llbracket P \rrbracket) = \llbracket P \rrbracket$ for all $\sigma \in S_\infty$.
- P is **generically invariant** if $\llbracket P \rrbracket \triangle \sigma(\llbracket P \rrbracket)$ is meager in Dir_K for all $\sigma \in S_\infty$.
- P is **determined** if the game $G(P)$ is determined, i.e. one of the players has a winning strategy.
- P is **strongly determined** if either P or $\neg P$ is generic, i.e. Player II has a winning strategy in one of the games $G(P)$ and $G(\neg P)$, and $\llbracket P \rrbracket$ is either meager or comeager.

Recall that a set A in a topological space X has the Baire property if there is an open set $U \subseteq X$ such that $U \triangle A$ is meager. The family of sets with the Baire property is a σ -algebra which contains the Borel sets (in fact, it is the σ -algebra generated by the open sets and the meager sets) [55, Proposition I.8.22].

In the analogy between measure and category, the sets with the Baire property correspond to the measurable sets (the σ -algebra generated by the open sets and the null sets for a given measure μ). This is reflected in the following proposition (one half of [55, Exercise I.8.35]), which says that the “size” of a set with the Baire property is witnessed by a winning strategy for one of the players in the Banach–Mazur game.

Proposition 4.3.3. *If A has the Baire property, then the Banach–Mazur game $G^{**}(A, X)$ is determined. Consequently, if $\llbracket P \rrbracket$ has the Baire property, then P is determined.*

Proof. Let $B = X \setminus A$. Then B also has the Baire property, so there is an open set U such that $U \triangle B$ is meager. If U is empty, then B is meager, so A is comeager, and Player II has a winning strategy in $G^{**}(A, X)$. On the other hand, if U is non-empty, then $B \cap U$ is comeager in U , so Player II has a winning strategy in $G^{**}(B \cap U, U)$. Player I can “steal” this strategy to obtain a winning strategy in $G^{**}(A, X)$.

Player I begins by playing $V_0 = U$. Thereafter, Player I simulates a game of $G^{**}(B \cap U, U)$ with run $(V'_i)_{i \in \omega}$, following the winning strategy for Player II in this game, and ensuring that $V'_i = V_{i+1}$ for all i . Then $\bigcap_{i \in \omega} V_i = \bigcap_{i \in \omega} V'_i \subseteq B \cap U \subseteq X \setminus A$, so this is a winning strategy for Player I.

The last statement follows directly from Theorem 4.1.13. □

Remark 4.3.4. Recall from Section 1.4 that if $\varphi(\bar{x})$ is a formula of $L_{\omega_1, \omega}$ and \bar{a} is a tuple from ω , then $\llbracket \varphi(\bar{a}) \rrbracket \subseteq \text{Str}_L$ is a Borel set. Since the topology on Dir_K is finer than the subspace topology inherited from Str_L , $\llbracket \varphi(\bar{a}) \rrbracket \cap \text{Dir}_K$ is also a Borel set. In particular, $\varphi(\bar{a})$ is a Borel property, and hence a determined property.

It will be convenient to define a variant of our game $G(P)$. Given a property P and a finite labeled structure B , the game $G(P, B)$ is just like the game $G(P)$, except that on turn

0, Player I is required to play some A_0 such that $B \preceq A_0$. Here are some easy observations about this game:

- (1) If Player II has a winning strategy in $G(P)$, then Player II has a winning strategy in $G(P, B)$ for any B . Player I's opening move is only more constrained.
- (2) If Player II has a winning strategy in $G(P, B)$ for some B , then Player I has a winning strategy in $G(\neg P)$. This is a strategy stealing argument just as in the proof of Proposition 4.3.3: Player I opens with $A_0 = B$ and then follows the strategy with the roles of the players reversed.
- (3) If Player I has a winning strategy in $G(P)$ with opening play $A_0 = B$, then Player II has a winning strategy in $G(\neg P, B)$. Player II steals the winning strategy, imagining that they are player I and their opening play was B .
- (4) Player II has a winning strategy in the game $G(P, A)$ if and only if $\llbracket P \rrbracket \cap U_A$ is comeager in U_A . The proof is exactly the same as in Theorem 4.1.13.

Remark 4.3.5. To say that P is determined is to say that if Player II cannot force the direct limit to satisfy P , then Player I can force the direct limit to satisfy $\neg P$. To say that P is strongly determined is to say that if Player II cannot force the direct limit to satisfy P , then Player II can force the direct limit to satisfy $\neg P$. This is stronger, since by observations (1) and (2) above, if Player II has a winning strategy in $G(\neg P)$, then Player I has a winning strategy in $G(P)$.

As in the proof of Theorem 4.2.2, for any A in K , we denote by Emb_A the property of labeled K -direct limits M that A embeds strongly in M .

Theorem 4.3.6. *The following are equivalent:*

- (1) K has the joint embedding property.
- (2) Every generically invariant determined property P is strongly determined.
- (3) Every invariant determined property P is strongly determined.
- (4) For every A in K , Emb_A is generic.

Proof. (1)→(2): Suppose K has the joint embedding property. Let P be a generically invariant determined property, and assume that P is not generic. Since P is determined, Player I has a winning strategy in the game $G(P)$. We will show that Player II has a winning strategy in the game $G(\neg P)$, establishing that $\neg P$ is generic.

Suppose that the first move in Player I's winning strategy for $G(P)$ is the finite labeled structure B . Then Player II has a winning strategy in the game $G(\neg P, B)$, and $\llbracket \neg P \rrbracket \cap U_B$ is comeager in U_B .

Suppose Player I opens with A_0 in our game $G(\neg P)$. Since K has the joint embedding property, there is some C in K admitting strong embeddings $f: A_0 \rightarrow C$ and $g: B \rightarrow C$, and we may assume that C is a finite labeled structure and f is an inclusion. Player II fixes some $\sigma \in S_\infty$ such that σ agrees with g on its restriction to the domain of B . Now $\sigma(U_B) = \{M \in \text{Dir}_K \mid B \preceq \sigma^{-1}(M)\} = \{M \in \text{Dir}_K \mid g(B) \preceq M\}$, and $g(B) \preceq C$, so $U_C \subseteq \sigma(U_B)$. Since $\sigma(\llbracket \neg P \rrbracket) \cap \sigma(U_B)$ is comeager in $\sigma(U_B)$, also $\sigma(\llbracket \neg P \rrbracket) \cap U_C$ is comeager in U_C .

Let P' be the property of $M \in \text{Dir}_K$ that $\sigma^{-1}(M)$ satisfies P , so $\llbracket \neg P' \rrbracket = \sigma(\llbracket \neg P \rrbracket)$. Let Q be the property of $M \in \text{Dir}_K$ that M satisfies P if and only if $\sigma^{-1}(M)$ satisfies P . The complement of $\llbracket Q \rrbracket$ is $\llbracket P \rrbracket \Delta \sigma(\llbracket P \rrbracket)$, which is meager, since P is generically invariant. So $\llbracket Q \rrbracket$ is comeager in Dir_K , and $\llbracket Q \rrbracket \cap \llbracket \neg P' \rrbracket \cap U_C$ is comeager in U_C . Hence Player II has a winning strategy in the game $G(Q \wedge \neg P', C)$. Player II plays $A_1 = C$ and follows the winning strategy thereafter.

The resulting K -direct limit $M = \varinjlim(A_i)$ satisfies $\neg P'$, so $\sigma^{-1}(M)$ satisfies $\neg P$. But M also satisfies Q , so M satisfies $\neg P$, as desired.

(2) \rightarrow (3): Invariant properties are generically invariant.

(3) \rightarrow (4): Emb_A is invariant, since strong embeddings are isomorphism-invariant. And letting S be the set of labeled finite structures isomorphic to A , $\llbracket \text{Emb}_A \rrbracket = \bigcup_{A' \in S} U_{A'}$. So Emb_A is an open property, and hence it is determined, by Proposition 4.3.3.

So Emb_A is strongly determined. But Player I has a winning strategy in $G(\neg \text{Emb}_A)$ (play a labeled copy of A on turn 0 and play arbitrarily thereafter), so $\neg \text{Emb}_A$ is not generic, and hence Emb_A is generic.

(4) \rightarrow (1): If A and B are in K , then the conjunction of Emb_A and Emb_B is generic, so then there is some labeled K -direct limit M admitting strong embeddings $f: A \rightarrow M$ and $g: B \rightarrow M$. Any finite $C \preceq M$ with $f(A) \preceq C$ and $g(B) \preceq C$ witnesses the joint embedding property for A and B . \square

So we have an analogy between genericity in the spaces Dir_K and invariant measures on the space Str_L . Classes with the joint embedding property are like ergodic structures: every invariant Borel set is either very large or very small. If such a class is a generalized Fraïssé class, it is like a measure which is almost surely isomorphic to a countable structure: it concentrates on a single orbit of the logic action. And those classes with the joint embedding property which fail to be generalized Fraïssé classes are like properly ergodic structures (see Corollary 4.3.9 and Proposition 4.3.13 below).

The equivalence between (2) and (3) in Theorem 4.3.6 demonstrates that in the context of strong embedding classes, there is no analog of the distinction between ergodic and weakly ergodic measures (see Remarks 1.2.6 and 1.4.5). The generically invariant Borel sets are the analogs of almost surely invariant Borel sets, but the theorem shows that if all of the invariant open properties $\{\text{Emb}_B \mid B \text{ in } K\}$ are strongly determined, this already implies that K has the joint embedding property and all generically invariant determined properties are strongly determined.

Definition 4.3.7. The **generic theory** of K is $T^{\text{gen}}(K) = \{\varphi \in L_{\omega_1, \omega} \mid \llbracket \varphi \rrbracket \text{ is generic}\}$. If F is a fragment of $L_{\omega_1, \omega}$, define $T_F^{\text{gen}}(K) = \{\varphi \in F \mid \llbracket \varphi \rrbracket \text{ is generic}\}$.

The next corollary is the analog of Proposition 1.4.4.

Corollary 4.3.8. *If K has the joint embedding property, then $T^{\text{gen}}(K)$ is a complete and countably consistent theory. That is, for every sentence φ , $\varphi \in T^{\text{gen}}(K)$ or $\neg\varphi \in T^{\text{gen}}(K)$, and every countable subset $\Sigma \subseteq \text{Th}(\mu)$ has a model. In particular, for any countable fragment F , some countable K -direct limit is a model of $T_F^{\text{gen}}(K)$.*

Proof. Every sentence φ of $L_{\omega_1, \omega}$ is an invariant determined property (Remark 4.3.4), so by Theorem 4.3.6, $\varphi \in T^{\text{gen}}(K)$ or $\neg\varphi \in T^{\text{gen}}(K)$. And if Σ is a countable subset, then since a countable intersection of comeager sets is comeager, $\llbracket \bigwedge_{\varphi \in \Sigma} \varphi \rrbracket$ is comeager. In particular, it is non-empty in Dir_K . \square

Corollary 4.3.9. *Suppose that K has the joint embedding property, but K is not a generalized Fraïssé class. Then there is no labeled K -direct limit which satisfies every invariant generic property, and no countable K -direct limit is a model of $T^{\text{gen}}(K)$.*

Proof. Let M be any countable K -direct limit. If $T^{\text{gen}}(K)$ contains the Scott sentence φ_M , then the property Iso_M is generic, and M is the generic limit of K . By Theorem 4.2.2, this contradicts our assumption that K is not a generalized Fraïssé class.

Then since $T^{\text{gen}}(K)$ is complete, it contains $\neg\varphi_M$, and M fails to satisfy the generic property $\neg\text{Iso}_M$. \square

There are two reasons why a class K with the joint embedding property could fail to be a generalized Fraïssé class: K could be uncountable up to isomorphism, or it could fail the weak amalgamation property. In the first case, since Emb_A is an invariant generic property for all A in K , it is clear that no countable K -direct limit can satisfy every invariant generic property.

Example 4.3.10. Let L be the language $\{R_n \mid n \in \omega\}$, where each R_n is a binary relation symbol. Let K be the class of all finite L -structures in which each R_n is anti-reflexive and symmetric, together with all embeddings between structures in K .

Then K is a (pruned) hereditary class, and it satisfies the joint embedding property and the amalgamation property (for example, given embeddings $f: A \rightarrow B$ and $g: A \rightarrow C$, take the union of B and C , identifying the substructures $f(A)$ and $g(A)$, and do not add any new relations). But it is not a generalized Fraïssé class, since it is not countable up to isomorphism: already there are continuum-many structures in K of size 2.

This is a Baire category version of the measure construction of the kaleidoscope random graph, our most basic example of a properly ergodic structure, as described in Example 1.4.18. The generic theory first-order theory $T_{\text{FO}}^{\text{gen}}(K)$ agrees with the the first-order theory $\text{Th}_{\text{FO}}(\mu)$ of the kaleidoscope random graph. It has continuum-many countable models, but its reduct to any finite sublanguage is countably categorical (and the reducts of the structures in K to any finite sublanguage form a Fraïssé class).

Next we give an example of a class K which is countable up to isomorphism and has the joint embedding property, but which does not have the weak amalgamation property.

Example 4.3.11. Let $L = \{R, P\}$, let S_K be the class of all finite acyclic graphs with edge relation R of degree at most 2 (see the class K_1 in Examples 3.1.9, and 4.2.7), such that P is a unary relation picking out an arbitrary subset, and let E_K be the class of all embeddings between structures in S_K .

It is easy to see that K is hereditary and countable up to isomorphism. It also has the joint embedding property: given A and B in K , we can embed them both in C simply by connecting one end of A to one end of B and setting P or $\neg P$ arbitrarily on the new elements.

But K does *not* have the weak amalgamation property. Indeed, given any non-empty A in K and any embedding $A \rightarrow B$ in K , choose a connected component D of B containing a connected component of A . Let C_1 be the extension of B obtained by adding two new vertices labeled P , one connected to each endpoint of D , and let C_2 be the extension of B obtained by adding two new vertices labeled $\neg P$, one connected to each endpoint of D . The embeddings $B \rightarrow C_1$ and $B \rightarrow C_2$ cannot be amalgamated over A .

The generic first-order theory $T_{\text{FO}}^{\text{gen}}(K)$ is simply the complete theory of a chain, infinite in both directions, which embeds every finite P -labeled chain. And while connectedness is not expressible by a first-order sentence, it is a generic invariant Borel property (expressible by a sentence of $L_{\omega_1, \omega}$), as shown for K_1 in Example 4.2.7.

Remark 4.3.12. At this juncture, it is worth noting a connection with forcing. If K is a pruned, non-empty, extendible class, then forcing with the poset of finite labeled structures, ordered by strong inclusions, gives rise to a new K -direct limit M_G in the forcing extension $V[G]$ which has all generic Borel properties present in V . We can reinterpret the theorems of this section and the last: K has the joint embedding property if and only if the set of invariant generic Borel properties in V which are satisfied by M_G does not depend on the generic G (e.g. M_G satisfies $T^{\text{gen}}(K)$ computed in V , which is complete for $L_{\omega_1, \omega}$ in V). And K is a generalized Fraïssé class if and only if M_G is isomorphic to a structure which is already in V (the generic limit of K).

In Example 4.3.11, given a K -direct limit M which is infinite in both directions, choosing a starting point v and a direction to read gives an infinite binary sequence $r \in 2^\omega$, with $r(n) = 1$ if and only if the n^{th} vertex after v in the chain satisfies P . If $M \in V[G]$ is obtained by forcing with the poset of finite labeled structures, then r is a Cohen real. Indeed, Cohen forcing adds a new real which is in every comeager Borel set in 2^ω which is in V , and one can check that for a given comeager set $X \subseteq 2^\omega$, the property “any real obtained from M as described above lands in X ” is generic.

We can also consider a measure version of this discussion, analogous to random real forcing. Given an ergodic structure μ on Str_L , forcing with the measure algebra on Str_L (the Boolean algebra of Borel sets modulo the ideal of measure 0 sets) gives rise to a new L -structure M_G with domain ω in $V[G]$ which is in every measure 1 Borel set in V . So M_G satisfies $\text{Th}(\mu)$ computed in V , which is complete for $L_{\omega_1, \omega}$ in V . M is not isomorphic to any structure in V if and only if μ is properly ergodic.

Now if K has the joint embedding property but is not a generalized Fraïssé class, no property which is satisfied by only countably many countable K -direct limits up to isomorphism is generic. Indeed, for any countable set $\{M_i \mid i \in \omega\}$ of K -direct limits, with Scott sentences φ_{M_i} , $\bigcup_{i \in \omega} \llbracket \varphi_{M_i} \rrbracket$ is a meager set. However, we can ask whether the analog of Corollary 2.3.8 (which we called an analog of Vaught’s conjecture for ergodic structures) holds: is every generic property satisfied by continuum-many countable K -direct limits up to isomorphism? The answer is yes, assuming that K does not have the weak amalgamation property.

Proposition 4.3.13. *Suppose that K has the joint embedding property but does not have the weak amalgamation property. Then if P is a generic property, P is satisfied by continuum-many countable K -direct limits up to isomorphism.*

Proof. Since P is generic, Player II has a winning strategy in the game $G(P)$. We use this strategy and the “isomorphism squashing” strategy of Player I from the (4)→(1) direction of Theorem 4.2.2 to build an infinite binary tree $\langle A_\eta \rangle_{\eta \in 2^{<\omega}}$ of finite labeled structures. We ensure that $A_\eta \preceq A_{\eta 0}$ and $A_\eta \preceq A_{\eta 1}$ for all $\eta \in 2^{<\omega}$, so that we can build K -direct limits $M_\lambda = \varinjlim (A_{\lambda \upharpoonright n})$ for all $\lambda \in 2^\omega$. And we will further ensure that M_λ satisfies P and $M_\lambda \not\cong M_{\lambda'}$ for all $\lambda \neq \lambda'$ in 2^ω .

Since K does not have the weak amalgamation property, there is a structure A in K such that no strong embedding $A \rightarrow B$ is an amalgamation embedding. We may assume that A is a labeled structure. At stage 0, Player I plays $A_\emptyset = A$ (here $\langle \rangle$ is the empty sequence, the unique element of 2^0).

At stage n , for n odd, it is Player II’s turn. For each $\eta \in 2^n$, we have a partial run of the game $G(P)$, given by $A_{\eta \upharpoonright 0} \preceq A_{\eta \upharpoonright 1} \preceq \dots \preceq A_\eta$. Player II chooses structures $A_{\eta 0} = A_{\eta 1}$, according to the winning strategy for $G(P)$.

At stage $n > 0$, for n even, it is Player I’s turn. For each $\eta \in 2^n$ and each strong embedding $f: A \rightarrow A_\eta$, Player I adds f to the end of a queue.

Player I then pulls an embedding $f': A \rightarrow A_\nu$ from the front of the queue. Note that $\nu \in 2^l$ for some $l \leq n$. Enumerate the pairs (ζ, ξ) such that $\xi \in 2^n$ extends ν but $\zeta \in 2^n$ does not extend ν as (ζ_i, ξ_i) for $i < k$. Handling each pair in turn, Player I extends each A_η , $\eta \in 2^n$, by a sequence of intermediate extensions $A_\eta = B_\eta^0 \preceq B_\eta^1 \preceq \dots \preceq B_\eta^k$.

For each pair $(\zeta, \xi) = (\zeta_i, \xi_i)$, Player I looks at the current intermediate extensions B_ζ^i and B_ξ^i . Composing f' with the inclusion $A_\nu \rightarrow A_\xi \rightarrow B_\xi^i$, we have a strong embedding $f: A \rightarrow B_\xi^i$. Exactly as in Theorem 4.2.2, with B_ζ^i playing the role of A_k and B_ξ^i playing the role of A'_k , Player I finds extensions B_ζ^{i+1} and B_ξ^{i+1} which cannot be amalgamated over A (i.e. over the inclusion $A \rightarrow B_\zeta^{i+1}$ and the strong embedding $f: A \rightarrow B_\xi^i$). For all η not equal to ζ or ξ , Player I simply sets $B_\eta^{i+1} = B_\eta^i$.

After handling each pair, Player I sets $A_{\eta 0} = A_{\eta 1} = B_\eta^k$ for all $\eta \in 2^n$.

Now for all $\lambda \in 2^\omega$, $(A_{\lambda \upharpoonright n})_{n \in \omega}$ is a run of the game $G(P)$ in which Player II played according to a winning strategy. So $M_\lambda = \varinjlim (A_{\lambda \upharpoonright n})$ satisfies P . Suppose for contradiction that there is an isomorphism $\sigma: M_\lambda \rightarrow M_{\lambda'}$ for some $\lambda \neq \lambda'$. Then σ restricts to a strong

embedding $f': A \rightarrow A_{\lambda' \upharpoonright l}$ for some l , and we may pick l even and large enough so that $(\lambda \upharpoonright l) \neq (\lambda' \upharpoonright l)$.

The embedding f' was added to the queue by Player I at stage l and pulled from the queue at some later stage n , and the pair (ζ, ξ) with $\zeta = (\lambda \upharpoonright n)$ and $\xi = (\lambda' \upharpoonright n)$ was considered at this stage. Just as in Theorem 4.2.2, the fact that Player I ensured that the extensions $A_\zeta \preceq A_{\lambda \upharpoonright (n+1)}$ and $A_\xi \preceq A_{\lambda' \upharpoonright (n+1)}$ cannot be amalgamated over A contradicts the assumption that σ is an isomorphism. \square

4.4 Model theory and the generic limit

In the previous sections, we have used the word “category” in two unrelated ways, referring to *categories* of finite structures and to Baire *category*. We now add a third sense of the word to the mix, also unrelated to the others, characterizing when the first-order theory $\text{Th}(M) = T_{\text{FO}}^{\text{gen}}(K)$ of the generic limit M is countably *categorical*.

We then consider other model-theoretic properties of M and $\text{Th}(M)$: atomicity, model completeness, and quantifier elimination.

The key condition in the characterization of countable categoricity is the existence of a weak Löwenheim–Skolem function. A version of this notion was introduced (in the context of smooth classes with the amalgamation property) by Kueker and Laskowski [66], who observed the connection with countable categoricity, but we believe the name is due to Hill [42]. As far as we are aware, all other sources which prove versions of Theorem 4.4.3 make use of unnecessary definability assumptions on the class of strong embeddings.

Definition 4.4.1. Let K be a generalized Fraïssé class, and let M be its generic limit. A function $l: \omega \rightarrow \omega$ is a **weak Löwenheim–Skolem function** for K if for all finite subsets $A \subseteq M$, there exist $B \preceq C \preceq M$ such that $A \subseteq B$, the inclusion $B \rightarrow C$ is an amalgamation embedding, and $|C| \leq l(|A|)$.

We can give an equivalent finitary condition which doesn’t mention the generic limit.

Proposition 4.4.2. *The function $l: \omega \rightarrow \omega$ is a weak Löwenheim–Skolem function for K if and only if for all D in K and $A \subseteq D$, there exists a strong embedding $f: D \rightarrow D'$, with D' in K , and strong substructures $B \preceq C \preceq D'$, such that $f(A) \subseteq B$, the inclusion $B \rightarrow C$ is an amalgamation embedding, and $|C| \leq l(|A|)$.*

Proof. Suppose that l is a weak Löwenheim–Skolem function for K . Let D be in K and $A \subseteq D$. We identify D with a strong substructure of the generic limit M , so also $A \subseteq M$. Now there exist $B \preceq C \preceq M$ such that $A \subseteq B$, the inclusion $B \rightarrow C$ is an amalgamation embedding, and $|C| \leq l(|A|)$. By Proposition 3.2.21, we may pick $D' \preceq M$ such that $D \preceq D'$ and $C \preceq D'$, and hence the condition in the statement holds, with $f: D \rightarrow D'$ the strong inclusion.

Conversely, suppose the condition in the statement holds, and let A be a finite subset of the generic limit M . Pick $E \preceq D \preceq M$ such that $A \subseteq E$ and D witnesses weak- K -homogeneity for E , and let $i: D \rightarrow E$ be the inclusion. Now there is a strong embedding $f: D \rightarrow D'$ and $B \preceq C \preceq D'$ in K , such that $f(A) \subseteq B$, the inclusion $B \rightarrow C$ is an amalgamation embedding, and $|C| \leq l(|A|)$. By weak- K -homogeneity, there is a strong embedding $g: D' \rightarrow M$ such that $g \circ f \circ i$ is the inclusion of E in M . Let $B' = g(B)$ and $C' = g(C)$. Then since $A \subseteq E$ and $f(A) \subseteq B$, $A = g(f(A)) \subseteq g(B) = B'$. Moreover, the inclusion $B' \rightarrow C'$ is an amalgamation embedding, and $|C'| \leq l(|A|)$, so l is a weak Löwenheim–Skolem function for K . \square

We use the Ryll–Nardzewski characterization of countably categorical theories as those with only finitely many n -types over the empty set for all n [75, Theorem 4.4.1].

Theorem 4.4.3. *Let K be a generalized Fraïssé class. Then $T_{\text{FO}}^{\text{gen}}(K)$ is countably categorical if and only if*

- (1) K is small, and
- (2) K has a weak Löwenheim–Skolem function.

Proof. Let M be the generic limit of K , and suppose $T_{\text{FO}}^{\text{gen}}(K) = \text{Th}(M)$ is countably categorical. For (1), if K is not small, then there is some n such that K_n is infinite up to isomorphism. But since M is K -universal, there are already infinitely many quantifier-free n -types realized in M , contradicting Ryll–Nardzewski.

For (2), we define $l(n)$ as follows. Let S_n be the (finite) set of complete n -types realized in M . For each $p(\bar{x}) \in S_n$, choose a tuple $\bar{a}_p \in M$ realizing p , and let k_p be minimal such that there exists $B \preceq C \preceq M$ such that \bar{a}_p is contained in B , the inclusion $B \rightarrow C$ is an amalgamation embedding, and $|C| = k_p$. Let $l(n) = \max_{p \in S_n} k_p$.

Now if $A \subseteq M$ with $|A| = n$, enumerating A as \bar{a} , we have $M \models p(\bar{a})$ for some $p(\bar{x}) \in S_n$. Since countably categorical structures are strongly homogeneous, there is an automorphism σ of M such that $\sigma(\bar{a}) = \bar{a}_p$, the realization of p chosen above. Now there exists $B \preceq C \preceq M$ such that \bar{a}_p is contained in B , the inclusion $B \rightarrow C$ is an amalgamation embedding, and $|C| = k_p \leq l(n)$. Letting $C' = \sigma^{-1}(C)$ and $B' = \sigma^{-1}(B)$, we have $A \subseteq B' \preceq C' \preceq M$, the inclusion $B' \rightarrow C'$ is an amalgamation embedding, and $|C'| \leq l(n)$. So l is a weak Löwenheim–Skolem function for K .

For the converse, we will show assuming (1) and (2) that only finitely many n -types are realized in M for each n . The suffices by Ryll–Nardzewski. For each n -tuple $\bar{a} \in M$, let $A = \|\bar{a}\|$ and $n' = |A| \leq n$. Choose some $B_{\bar{a}} \preceq C_{\bar{a}} \preceq M$ such that $A \subseteq B_{\bar{a}}$, the inclusion $B_{\bar{a}} \rightarrow C_{\bar{a}}$ is an amalgamation embedding, and $|C_{\bar{a}}| \leq l(n')$. Given two such tuples \bar{a} and \bar{a}' , if $f: C_{\bar{a}} \rightarrow C_{\bar{a}'}$ is an isomorphism with $f(\bar{a}) = \bar{a}'$, then by weak- K -ultrahomogeneity and Remark 4.2.5, there is an automorphism σ of M extending $f \upharpoonright B_{\bar{a}}$. So $\sigma(\bar{a}) = \bar{a}'$, and hence \bar{a} and \bar{a}' have the same type in M .

Hence the type of \bar{a} is determined by the isomorphism type of $C_{\bar{a}}$ and the way \bar{a} sits inside $C_{\bar{a}}$. There are only finitely many choices for the size of $C_{\bar{a}}$, only finitely many isomorphism

types of each size, and only finitely many n -tuples from each isomorphism type, so there are only finitely many n -types realized in M . \square

Semi-definability (Definition 4.1.3) is connected to the model-theoretic properties of atomicity and model completeness. Since we are only interested in the generic limit M , we only need to consider semi-definability restricted to the set $[[\text{Iso}_M]]$ of labeled K -direct limits isomorphic to M .

Definition 4.4.4. Let K be a generalized Fraïssé class with generic limit M . Recall that $[[\text{Iso}_M]]$ is the subset of Dir_K consisting of those structures isomorphic to M . K is **generically semi-definable** if the subspace topology on $[[\text{Iso}_M]]$ inherited from Dir_K is equal to the subspace topology on $[[\text{Iso}_M]]$ inherited from Str_L .

Remark 4.4.5. K is generically semi-definable if and only if for every $A \preceq M$, enumerated by a tuple \bar{a} , there is a first-order quantifier-free formula $\psi_A(\bar{x}, \bar{y})$ such that $M \models \exists \bar{y} \psi_A(\bar{a}, \bar{y})$, and if $M \models \exists \bar{y} \psi_A(\bar{b}, \bar{y})$, then the function $a_i \mapsto b_i$ is a strong embedding $A \rightarrow M$. The argument is exactly as in Proposition 4.1.4 (4), except that we are only interested in the single structure M up to isomorphism.

Theorem 4.4.6. *Let K be a generalized Fraïssé class with generic limit M . Then K is generically semi-definable if and only if every type realized in M is isolated by an existential formula. As a consequence, if K is generically semi-definable, then M is atomic, and the converse holds if $T_{\text{FO}}^{\text{gen}}(K)$ is model complete.*

Proof. Suppose K is generically semi-definable, let \bar{c} be a tuple from M , and pick finite strong substructures $A \preceq B \preceq M$ such that A contains \bar{c} and B witnesses weak- K -ultrahomogeneity for A . Let \bar{a} enumerate the elements of A not in \bar{c} , and let \bar{b} enumerate the elements of B not in A . Let $\exists \bar{z} \psi_B(\bar{c}, \bar{a}, \bar{b}, \bar{z})$ be the formula given by generic semi-definability for B . We claim that $\exists \bar{x} \exists \bar{y} \exists \bar{z} \psi_B(\bar{w}, \bar{x}, \bar{y}, \bar{z})$ isolates the type of \bar{c} . Indeed, suppose that $M \models \exists \bar{x} \exists \bar{y} \exists \bar{z} \psi_B(\bar{c}', \bar{x}, \bar{y}, \bar{z})$ for some other tuple \bar{c}' . Then $M \models \exists \bar{z} \psi_B(\bar{c}', \bar{a}', \bar{b}', \bar{z})$ for some tuples \bar{a}', \bar{b}' , so the map $c_i \mapsto c'_i, a_j \mapsto a'_j, b_k \mapsto b'_k$ is a strong embedding $B \rightarrow M$. Since B witnesses weak- K -ultrahomogeneity for A , there is an automorphism σ of M such that $\sigma(\bar{c}\bar{a}) = \bar{c}'\bar{a}'$. In particular, \bar{c} and \bar{c}' have the same type in M .

Conversely, suppose every type realized in M is isolated by an existential formula. For any $A \preceq M$, enumerated by a tuple \bar{a} , let $\exists \bar{y} \psi(\bar{x}, \bar{y})$ isolate the type of \bar{a} . We show that $\psi(\bar{x}, \bar{y})$ witnesses generic semi-definability for $A \preceq M$. Certainly $M \models \exists \bar{y} \psi(\bar{a}, \bar{y})$, and if $M \models \exists \bar{y} \psi(\bar{b}, \bar{y})$ for some other tuple \bar{b} , then \bar{a} and \bar{b} have the same type in M . Since countable atomic models are strongly homogeneous [75, Lemma 4.2.14], there is an automorphism σ of M with $\sigma(\bar{a}) = \bar{b}$. By isomorphism invariance, since $A \preceq M$, $B = \sigma(A) \preceq M$, so the restriction of σ to A is a strong embedding.

We have shown that if K is generically semi-definable, then M is atomic. Conversely, if $T_{\text{FO}}^{\text{gen}}(K)$ is model complete, then it implies that every formula is equivalent to an existential formula. So if M is atomic, then every type realized in M is isolated by an existential formula, and hence K is generically semi-definable. \square

Corollary 4.4.7. *Let K be a generalized Fraïssé class, and suppose that $T_{\text{FO}}^{\text{gen}}(K)$ is countably categorical. Then K is generically semi-definable if and only if $T_{\text{FO}}^{\text{gen}}(K)$ is model complete.*

Proof. Suppose that $T_{\text{FO}}^{\text{gen}}(K)$ is model complete. Then, since the unique countable model of a countably categorical theory is atomic, K is generically semi-definable by Theorem 4.4.6.

Conversely, suppose K is generically semi-definable. Then by Theorem 4.4.6, every type realized in M is isolated by an existential formula. Since $T_{\text{FO}}^{\text{gen}}(K)$ is countably categorical, every formula $\varphi(\bar{x})$ with n free variables is equivalent to the disjunction of the finitely many complete n -types containing $\varphi(\bar{x})$. The finite disjunction of the existential formulas isolating these types is again an existential formula which is equivalent to $\varphi(\bar{x})$, so $T_{\text{FO}}^{\text{gen}}(K)$ is model complete. \square

Quantifier-elimination is a natural strengthening of model completeness, and it holds for classical Fraïssé limits (of small Fraïssé classes). But we should not expect it to hold for generalized Fraïssé classes which are not Fraïssé classes. Indeed, if K is not hereditary, then we may need to look outside of a finite tuple \bar{a} to a strong substructure $A \preceq M$ containing \bar{a} to determine the type of \bar{a} in M . And similarly, if K has the weak amalgamation property but not the amalgamation property, then we may need to look outside of a finite $A \preceq M$ to a larger $A \preceq B \preceq M$ which witnesses K -ultrahomogeneity for A to determine the type of A in M .

So assuming that K is a Fraïssé class, we state the classical theorem for the record (with a proof that fits smoothly into the setting we have developed here).

Corollary 4.4.8. *Suppose K is a Fraïssé class. Then $T_{\text{FO}}^{\text{gen}}(K)$ is countably categorical if and only if K is small. In this case, $T_{\text{FO}}^{\text{gen}}(K)$ has quantifier elimination.*

Proof. Using the fact that K is hereditary and has the amalgamation property, we show that $l(n) = n$ is a weak Löwenheim–Skolem function for K . Indeed, for any finite $A \subseteq M$, we have $A \preceq M$ (K is closed under substructure and all embeddings are strong), and id_A is an amalgamation embedding (K has the amalgamation property). So we may take $B = C = A$ in the definition of weak Löwenheim–Skolem function, and $|C| = |A| = l(|A|)$.

Hence condition (2) in Theorem 4.4.3 is trivially satisfied, and $T_{\text{FO}}^{\text{gen}}(K)$ is countably categorical if and only if K is small.

For quantifier elimination, note that K is a small hereditary class. Then, as shown in Proposition 4.1.4 (5), K is semi-definable (hence generically semi-definable), and the witnessing formulas can be taken to be quantifier-free instead of existential. Following the proof of Theorem 4.4.6, but with $\|c\| = A = B$, we find that the type of any tuple \bar{c} in M is isolated by a quantifier-free formula. And using countable categoricity as in the proof of Corollary 4.4.7, every formula is equivalent to finite disjunction of isolating quantifier-free formulas. \square

We close this section by returning once more to the classes discussed in Examples 3.1.9. See Example 4.2.7 for descriptions of their generic limits.

Example 4.4.9. K_0 is a small Fraïssé class, and its generic limit, the random graph, is countably categorical with quantifier elimination.

K_1 and K_2 are small, but they fail to have weak Löwenheim–Skolem functions, so the theory of their generic limit M , the doubly infinite chain, is not countably categorical. Indeed, there is no acceptable value for $l(2)$: if we take two points v and w at distance k in M , any substructure in K_2 containing both of them must have size at least k , and any substructure in K_1 witnessing weak- K_1 -homogeneity for the substructure v, w must have size at least k . However, K_1 is hereditary, hence semi-definable, and the doubly infinite chain is atomic, with the type of a tuple \bar{a} isolated by the existential formula describing the distances between the connected components of the substructure on $\|\bar{a}\|$.

K_3 also fails to have a weak Löwenheim–Skolem function. There is no acceptable value of $l(1)$: given a point v in the generic limit M , every strong substructure of M containing v must contain the entire connected component of v , which could be arbitrarily large. M is atomic, since the type of a tuple \bar{a} is isolated by the formula describing positions of the elements of \bar{a} in their connected components, and the sizes of these components. But these formulas are not equivalent to existential formulas; $T_{\text{FO}}^{\text{gen}}(K_3)$ is not model complete, and K is not generically semi-definable.

4.5 Strong robustness

The notion of a robust class of finite structures was introduced by Macpherson and Steinhorn [74] and developed by Marshall in his PhD thesis [76]. Hill [42] introduced a stronger condition, super-robustness, inspired by the definition due to Macpherson and Steinhorn of a robust class of chain complexity 0.

Much of the theory of robust classes could be adapted to the context of strong embedding classes, but we will not pursue this here. Instead, we follow Hill, translating his definition of super-robustness to our context and introducing a new notion, which we call strong robustness. Strong robustness is equivalent to super-robustness for chain classes, but it is a weaker notion in general. Nevertheless, we find that the introduction of strong robustness clarifies some issues in [42], allowing us to get slightly stronger theorems with rather simpler proofs. In particular, we demonstrate a tight connection between strong robustness and pseudofiniteness of the generic theory for generalized Fraïssé classes. We will return to pseudofiniteness (in the more restrictive context of small Fraïssé classes) in Part III.

Convention 4.1.9 remains in effect in this section: K is a non-empty pruned extendible strong embedding class.

Definition 4.5.1. Given A in K , the **cone above A in K** is the class of all B in K such that A embeds strongly in B .

Definition 4.5.2. K is **strongly robust** if for every formula $\varphi(\bar{x})$ there is some A_φ in K such that for all strong embeddings $f: A \rightarrow B$ and $g: B \rightarrow C$ in K and all \bar{b} from B , $B \models \varphi(\bar{b})$ if and only if $C \models \varphi(g(\bar{b}))$.

In other words, strong embeddings are elementary with respect to $\varphi(\bar{x})$ in the cone above A_φ in K .

Definition 4.5.3. Given A in K , the **rank** of A , $\text{rk}(A)$, is the largest n such that there is a chain of strong proper strong substructures $A_0 \prec A_1 \prec \dots \prec A_n = A$.

Remark 4.5.4. A few easy observations about rank:

- If A embeds strongly in B , then $\text{rk}(A) \leq \text{rk}(B)$.
- K contains structures of arbitrarily high rank, since it is pruned and non-empty.
- Since $A \prec B$ implies $|A| < |B|$, every finite structure in K has a rank, and $\text{rk}(A) \leq |A|$.

Definition 4.5.5 (Hill [42]). K is **super-robust** if for every formula $\varphi(\bar{x})$ there exists $e_\varphi \in \omega$ such that for all B in K with $\text{rk}(B) \geq e_\varphi$, all \bar{b} from B , and all strong embeddings $f: B \rightarrow C$ in K , $B \models \varphi(\bar{b})$ if and only if $C \models \varphi(f(\bar{b}))$.

In other words, strong embeddings are elementary with respect to $\varphi(\bar{x})$ above rank e_φ .

Strong robustness and super-robustness are closely related. We now show that the latter implies the former, they are equivalent in the case of chain classes, and they are equivalent up to cofinality, assuming the joint embedding property.

Proposition 4.5.6. *If K is super-robust, then it is strongly robust.*

Proof. Given a formula $\varphi(\bar{x})$, let e_φ be the rank provided by super-robustness. Let A_φ be any structure with $\text{rk}(A_\varphi) \geq e_\varphi$. Then for all strong embeddings $f: A_\varphi \rightarrow B$ and $g: B \rightarrow C$ in K and all $\bar{b} \in B$, we have $\text{rk}(B) \geq \text{rk}(A_\varphi) \geq e_\varphi$, so by super-robustness $B \models \varphi(\bar{b})$ if and only if $C \models \varphi(f(\bar{b}))$. \square

Proposition 4.5.7. *If K is a strongly robust chain class, then K is super-robust.*

Proof. If K is a chain class, then the structures in K are linearly ordered under embeddability, and this order agrees with the linear order on their ranks. That is, B is in the cone above A in K if and only if $\text{rk}(B) \geq \text{rk}(A)$. So if A_φ witnesses strong robustness for $\varphi(\bar{x})$, then $e_\varphi = \text{rk}(A_\varphi)$ witnesses super-robustness for $\varphi(\bar{x})$. \square

Proposition 4.5.8. *If K is strongly robust and K' is a cofinal subclass of K , then K' is strongly robust.*

Proof. For any formula $\varphi(\bar{x})$, let A_φ witness strong robustness for $\varphi(\bar{x})$ in K , and pick any A'_φ in K' into which A_φ embeds strongly. Since the cone above A'_φ in K is a subclass of the cone above A_φ in K , A'_φ witnesses strong robustness for $\varphi(\bar{x})$ in K' . \square

Theorem 4.5.9. *Suppose K is countable up to isomorphism and has the joint embedding property. Then K has a cofinal strongly robust subclass if and only if K has a cofinal super-robust subclass.*

Proof. Any cofinal super-robust subclass is strongly robust by Proposition 4.5.6. Conversely, let K' be a cofinal strongly robust subclass of K . Then K' is also countable up to isomorphism and has the joint embedding property: if A and B are in K' , then A and B K -embed into some C in K , and C K -embeds into some D in K' , so A and B K' -embed into D .

So by Theorem 3.5.5, K' has a cofinal chain class K'' . By Proposition 4.5.8, K'' is strongly robust, and by Proposition 4.5.7, K'' is super-robust. Finally, K'' is a cofinal subclass of K , since if A is in K , then A embeds strongly in some B in K' , which embeds strongly in some C in K'' . \square

Now that we have sorted out the relationship between strong robustness and super-robustness (which just came down to the relationship between cones and ranks), we will address some easy consequences of these properties for K -direct limits. The inspiration for the following theorems comes entirely from [42]. Theorem 4.5.10 is the analog of Lemma 1.11 in [42], and the proof is essentially the same.

Theorem 4.5.10. *Let M be a K -universal K -direct limit. If K is strongly robust, then for every formula $\varphi(\bar{x})$ there is some finite $A_\varphi^* \preceq M$ such that for all finite B with $A_\varphi^* \preceq B \preceq M$ and all \bar{b} from B , $B \models \varphi(\bar{b})$ if and only if $M \models \varphi(\bar{b})$. The converse holds if K is a generalized Fraïssé class and M is its generic limit.*

Proof. Suppose K is strongly robust. We define A_φ^* by induction on the complexity of $\varphi(\bar{x})$:

- If $\varphi(\bar{x})$ is atomic, pick any finite $A_\varphi^* \preceq M$. For all finite B with $B \preceq M$ and all \bar{b} from B , we have $B \models \varphi(\bar{b})$ if and only if $M \models \varphi(\bar{b})$
- If $\varphi(\bar{x})$ is $\neg\psi(\bar{x})$, pick $A_\varphi^* = A_\psi^*$. For all finite B with $A_\varphi^* \preceq B \preceq M$ and all \bar{b} from B , $B \models \psi(\bar{b})$ if and only if $M \models \psi(\bar{b})$, so $B \models \varphi(\bar{b})$ if and only if $M \models \varphi(\bar{b})$.
- If $\varphi(\bar{x})$ is $\psi_1(\bar{x}) \wedge \psi_2(\bar{x})$, pick any finite $A_\varphi^* \preceq M$ such that $A_{\psi_1}^* \preceq A_\varphi^*$ and $A_{\psi_2}^* \preceq A_\varphi^*$. For all finite B with $A_\varphi^* \preceq B \preceq M$ and all \bar{b} from B , $B \models \psi_1(\bar{b})$ if and only if $M \models \psi_1(\bar{b})$, and $B \models \psi_2(\bar{b})$ if and only if $M \models \psi_2(\bar{b})$, so $B \models \varphi(\bar{b})$ if and only if $M \models \varphi(\bar{b})$.
- If $\varphi(\bar{x})$ is $\exists y \psi(\bar{x}, y)$, let A_φ witness strong robustness for $\varphi(\bar{x})$ in K . Since M is K -universal, there is a strong embedding $f: A_\varphi \rightarrow M$. Identify A_φ with its image $f(A_\varphi) \preceq M$. We are also given $A_\psi^* \preceq M$ by induction. Pick any finite $A_\varphi^* \preceq M$ such that $A_\varphi \preceq A_\varphi^*$ and $A_\psi^* \preceq A_\varphi^*$. Consider any finite B with $A_\varphi^* \preceq B \preceq M$ and any \bar{b} from B . If $B \models \varphi(\bar{b})$, then there is some $c \in B$ such that $B \models \psi(\bar{b}, c)$. Since $A_\psi^* \preceq B$, $M \models \psi(\bar{b}, c)$, so $M \models \varphi(\bar{b})$. Conversely, if $M \models \varphi(\bar{b})$, then there is some $c \in M$ such that $M \models \psi(\bar{b}, c)$, and there is some finite $C \preceq M$ such that $c \in C$ and $B \preceq C$. Since $A_\psi^* \preceq C$, we have $C \models \psi(\bar{b}, c)$, so $C \models \varphi(\bar{b})$. Then since $A_\varphi \preceq B \preceq C$, also $B \models \varphi(\bar{b})$.

For the converse, suppose K is a generalized Fraïssé class, M is its generic limit, and for each formula $\varphi(\bar{x})$ there is some A_φ^* as above. Pick any finite A_φ such that $A_\varphi^* \preceq A_\varphi \preceq M$

and A_φ witnesses weak- K -homogeneity for A_φ^* . Then A_φ witnesses strong robustness for $\varphi(\bar{x})$. Indeed, for all strong embeddings $f: A_\varphi \rightarrow B$ and $g: B \rightarrow C$ in K , by weak- K -homogeneity, we can embed C in M over A_φ^* , so that $A_\varphi^* \preceq B' \preceq C' \preceq M$, where B' and C' are isomorphic copies of B and C . Then we have $B' \models \varphi(\bar{b}')$ if and only if $M \models \varphi(\bar{b}')$ if and only if $C' \models \varphi(\bar{b})$, and hence, for any $\bar{b} \in B$, $B \models \varphi(\bar{b})$ if and only if $C \models \varphi(\bar{b})$. \square

Definition 4.5.11. Let M be a K -direct limit.

- We denote by $K(M)$ the set of all finite strong substructures of M . This is like the K -age of M , except that it is not closed under isomorphism.
- We define $K_+(M) = \{A \in K(M) \mid A \text{ is non-empty}\}$.
- Given $A \in K(M)$, the **cone above A in M** is $K_A(M) = \{B \in K(M) \mid A \preceq B\}$.
- The **cone filter**, denoted \mathfrak{C} , is the filter on $K(M)$ generated by the cones. Explicitly,

$$\mathfrak{C} = \{X \subseteq K(M) \mid K_A(M) \subseteq X \text{ for some } A \in K(M)\}.$$

Remark 4.5.12. If $A, B \in K(M)$, there is some finite $C \preceq M$ such that $A \preceq C$ and $B \preceq C$, so $K_C(M) \subseteq K_A(M) \cap K_B(M)$. Hence $\{X \subseteq K(M) \mid K_A(M) \subseteq X \text{ for some } A \in K(M)\}$ is closed under intersection, and \mathfrak{C} is a proper filter.

If $A \preceq M$ is finite and non-empty, then $K_A(M) \subseteq K_+(M)$, so $K_+(M) \in \mathfrak{C}$. Hence \mathfrak{C} restricts to a proper filter $\mathfrak{C}_+ = \{X \cap K_+(M) \mid X \in \mathfrak{C}\}$ on $K_+(M)$.

We consider ultraproducts of the finite structures in $K_+(M)$, modulo ultrafilters extending the cone filter.

Proposition 4.5.13. *If M is a K -direct limit and \mathfrak{U} is any ultrafilter on $K_+(M)$ extending \mathfrak{C}_+ , then there is a canonical embedding $j: M \rightarrow (\prod_{A \in K_+(M)} A)/\mathfrak{U}$.*

Proof. For each $A \in K_+(M)$, let $j_A: M \rightarrow A$ be any function such that $j_A(a) = a$ for all $a \in A$. Let $j: M \rightarrow (\prod_{A \in K_+(M)} A)/\mathfrak{U}$ be the product of the maps j_A , composed with the quotient by the ultrafilter.

If $(j'_A)_{A \in K_+(M)}$ is any other family of maps such that $j'_A(a) = a$ for all $a \in A$, then the induced map j' is equal to j . Indeed, for all $a \in M$, let $A \preceq M$ be any finite strong substructure with $a \in A$. Then $\{B \in K_+(M) \mid j_B(a) = j'_B(a)\} \supseteq K_A(M) \in \mathfrak{C}_+$, so $j(a) = j'(a)$ in $(\prod_{A \in K_+(M)} A)/\mathfrak{U}$.

It remains to check that j is an embedding. Let $\varphi(\bar{x})$ be an atomic formula, let \bar{a} be a tuple from M , and let $A \preceq M$ be any finite strong substructure containing \bar{a} . Then $M \models \varphi(\bar{a})$ if and only if $B \models \varphi(\bar{a})$ for all $B \in K_A(M)$. So $\{B \in K_+(M) \mid \varphi(j_B(\bar{a}))\} \supseteq K_A(M) \in \mathfrak{C}_+$, and by Los's Theorem, $(\prod_{A \in K_+(M)} A)/\mathfrak{U} \models \varphi(j(\bar{a}))$. \square

The embedding j is not elementary in general, since if $\varphi(\bar{x})$ is a formula with quantifiers, M may disagree with its finite strong substructures containing \bar{a} about the truth of $\varphi(\bar{a})$. But j is an elementary embedding when the conclusion of Theorem 4.5.10 holds, and in particular when K is strongly robust and M is K -universal.

Proposition 4.5.14. *Let M be a K -direct limit, and suppose that for every formula $\varphi(\bar{x})$ there is some finite $A_\varphi^* \preceq M$ such that for all B in K with $A_\varphi^* \preceq B \preceq M$ and all \bar{b} from B , $B \models \varphi(\bar{b})$ if and only if $M \models \varphi(\bar{b})$. Then for any ultrafilter \mathfrak{U} on $K_+(M)$ extending \mathfrak{C}_+ , the canonical embedding $j: M \rightarrow (\prod_{A \in K_+(M)} A)/\mathfrak{U}$ is an elementary embedding.*

Proof. Let $\bar{a} \in M$, and suppose $M \models \varphi(\bar{a})$. Pick any finite $B \preceq M$ such that B contains \bar{a} and $A_\varphi^* \preceq B$. Then for all $C \in K_B(M)$, we have $A_\varphi^* \preceq C$, so $C \models \varphi(\bar{a})$. Hence $\{C \in K_+(M) \mid C \models \varphi(j_C(\bar{a}))\} \supseteq K_B(M) \in \mathfrak{C}_+$, so $(\prod_{A \in K_+(M)} A)/\mathfrak{U} \models \varphi(j(\bar{a}))$. \square

The converse of Proposition 4.5.14 does not hold in general. If j is an elementary embedding for any \mathfrak{U} extending \mathfrak{C}_+ , then for any formula $\varphi(\bar{b})$ holding in M , $\varphi(\bar{b})$ holds on the cone above some $A \preceq M$. But the cone in question may depend on the tuple \bar{b} , not merely on the formula $\varphi(\bar{x})$ (i.e. we only get robustness, not strong robustness).

However, under the assumption that M is the generic limit of K and $\text{Th}(M)$ is model complete, the property that certain sentences (not formulas) in $\text{Th}(M)$ hold on a cone in M will be enough to ensure that K is strongly robust.

Theorem 4.5.15. *Let T be a model complete theory. Suppose further that for every sentence $\theta \in T$ there is some A_θ^* in K such that θ is true on the cone above A_θ^* in K , i.e. for every strong embedding $f: A_\theta^* \rightarrow B$ in K , $B \models \theta$. Then K is strongly robust.*

Proof. Let $\varphi(\bar{x})$ be any formula. Since T is model complete, there is an existential formula $\psi_\exists(\bar{x})$ and a universal formula $\psi_\forall(\bar{x})$ such that the sentence $\forall \bar{x} (\varphi(\bar{x}) \leftrightarrow \psi_\exists(\bar{x}) \leftrightarrow \psi_\forall(\bar{x}))$ is in T . Call this sentence θ_φ . We let $A_\varphi = A_{\theta_\varphi}^*$.

Then for all strong embeddings $f: A_\varphi \rightarrow B$ and $g: B \rightarrow C$ in K , $B \models \theta_\varphi$ and $C \models \theta_\varphi$. So for all \bar{b} from B , if $B \models \varphi(\bar{b})$, then $B \models \psi_\exists(\bar{b})$, so $C \models \psi_\exists(g(\bar{b}))$, as $\psi_\exists(\bar{x})$ is existential, so $C \models \varphi(g(\bar{b}))$. Similarly, if $C \models \varphi(g(\bar{b}))$, then $C \models \psi_\forall(g(\bar{b}))$, so $B \models \psi_\forall(\bar{b})$, as $\psi_\forall(\bar{x})$ is universal, so $B \models \varphi(\bar{b})$. \square

The next theorem summarizes our observations.

Theorem 4.5.16. *Suppose K is a generalized Fraïssé class with generic limit M such that $T_{\text{FO}}^{\text{gen}}(K) = \text{Th}(M)$ is model complete. The following are equivalent:*

- (1) K is strongly robust.
- (2) The canonical embedding $j: M \rightarrow (\prod_{A \in K_+(M)} A)/\mathfrak{U}$ is an elementary embedding for any ultrafilter \mathfrak{U} on $K_+(M)$ extending \mathfrak{C}_+ .

- (3) M is elementarily equivalent to $(\prod_{A \in K_+(M)} A)/\mathfrak{U}$ for any ultrafilter \mathfrak{U} on $K_+(M)$ extending \mathfrak{C}_+ .
- (4) For every sentence $\theta \in T_{\text{FO}}^{\text{gen}}(K)$ there is some A_θ^* in K such that θ is true on the cone above A_θ^* in K .

Proof. (1) \rightarrow (2) is Theorem 4.5.10 and Proposition 4.5.14, using the fact that M is K -universal.

(2) \rightarrow (3): Elementary embeddings witness elementary equivalence.

(3) \rightarrow (4): Suppose (4) fails. Then there is some $\theta \in T_{\text{FO}}^{\text{gen}}(K)$ such that for every A in K , there is some strong embedding $f: A \rightarrow B$ in K such that $B \models \neg\theta$.

We claim that $X_{-\theta} = \{C \in K_+(M) \mid C \models \neg\theta\}$ has non-empty intersection with every set in \mathfrak{C}_+ . It suffices to show that $X_{-\theta}$ has non-empty intersection with $K_A(M) \cap K_+(M)$ for all $A \in K(M)$. Indeed, for all $A \in K(M)$, there is some non-empty B in K such that $A \preceq B \preceq M$ and B witnesses weak- K -homogeneity for A . By our assumption, B embeds strongly in some C in K such that $C \models \neg\theta$. By weak- K -homogeneity, C embeds strongly in M over A . Identifying C with its image under this embedding, $C \in X_{-\theta} \cap (K_A(M) \cap K_+(M))$.

Hence $\mathfrak{C}_+ \cup \{X_{-\theta}\}$ extends to an ultrafilter \mathfrak{U} on $K_+(M)$, and $(\prod_{A \in K_+(M)} A)/\mathfrak{U} \models \neg\theta$, which contradicts (3).

(4) \rightarrow (1) is Theorem 4.5.15, using the fact that T is model complete. \square

Strong robustness is a very strong condition, but many more strong embedding classes (such as the pseudofinite Fraïssé classes discussed in Part III) have cofinal strongly robust subclasses. We can characterize this situation by applying Theorem 4.5.16 to cofinal subclasses; the new condition (5) is the most interesting.

Corollary 4.5.17. *Suppose K is a generalized Fraïssé class with generic limit M such that $\text{Th}(M)$ is model complete. The following are equivalent:*

- (0) K has a cofinal super-robust subclass.
- (1) K has a cofinal strongly robust subclass.
- (2) K has a cofinal subclass K' such that the canonical embedding $j: M \rightarrow \prod_{A \in K'_+(M)} A/\mathfrak{U}$ is an elementary embedding for any ultrafilter \mathfrak{U} on $K'_+(M)$ extending \mathfrak{C}_+ .
- (3) K has a cofinal subclass K' such that M is elementarily equivalent to $(\prod_{A \in K'_+(M)} A)/\mathfrak{U}$ for any ultrafilter \mathfrak{U} on $K'_+(M)$ extending \mathfrak{C}_+ .
- (4) K has a cofinal subclass K' such that for every sentence $\theta \in T_{\text{FO}}^{\text{gen}}(K)$ there is some A_θ^* in K' such that θ is true on the cone above A_θ^* in K' .
- (5) For all $\theta \in T_{\text{FO}}^{\text{gen}}(K)$, the models of θ are cofinal in K . That is, for every A in K , there is some strong embedding $f: A \rightarrow B$ in K such that $B \models \theta$.

Proof. (0) \leftrightarrow (1) is Theorem 4.5.9, using the fact that K is countable up to isomorphism and has the joint embedding property.

(1) \leftrightarrow (2) \leftrightarrow (3) \leftrightarrow (4) is by Theorem 4.5.16, using the fact that if K' is a cofinal subclass of K , then K' is also a generalized Fraïssé class with generic limit M (Theorem 4.2.6).

(4) \rightarrow (5): Suppose K has a cofinal subclass K' , such that for any $\theta \in \text{Th}(M)$, there is some A_θ^* in K' such that θ is true on the cone above A_θ^* in K' . For any A in K , by the joint embedding property there is some B in K such that A_θ^* and A K -embed in B . And since K' is cofinal in K , B K -embeds in some C in K' . But since A_θ^* K' -embeds in C , $C \models \theta$.

(5) \rightarrow (4): We build a chain class K' , cofinal in K , witnessing (4). Enumerate $\text{Th}(M)$ as $\{\theta_i \mid i \in \omega\}$, and enumerate the isomorphism classes in K as $\{B_i \mid i \in \omega\}$. To start, by assumption, B_0 embeds strongly in some A_0 in K such that $A_0 \models \theta_0$.

Given A_{n-1} , let $\psi_n = \bigwedge_{j=0}^n \theta_j$. Then $\psi_n \in \text{Th}(M)$. By the joint embedding property, A_{n-1} and B_n both embed strongly in some C_n in K , and by assumption C_n embeds strongly in some A_n in K such that $A_n \models \psi_n$. We may assume that the composite strong embedding $A_{n-1} \rightarrow A_n$ is an inclusion.

We have constructed a chain $A_0 \preceq A_1 \preceq A_2 \preceq \dots$, and we let K' be the full subclass of K consisting of those structures in K which are isomorphic to some A_i . K' is cofinal in K , since for any B in K , B is isomorphic to some B_n , and B_n embeds strongly in A_n . And for any $\theta \in \text{Th}(M)$, $\theta = \theta_n$ for some n , and θ is true in the cone above A_n in K' , since for all $m \geq n$, $A_m \models \psi_m$ and ψ_m implies θ_n . \square

This characterization takes on a particularly simple form for small Fraïssé classes.

Corollary 4.5.18. *Let K be a small Fraïssé class. Then K has a cofinal strongly robust subclass if and only if every sentence in $T_{\text{FO}}^{\text{gen}}(K)$ has a finite model in K .*

Proof. By Corollary 4.4.8, $T_{\text{FO}}^{\text{gen}}(K)$ has quantifier-elimination and hence is model complete.

Suppose K has a cofinal strongly robust subclass. By Corollary 4.5.17, for all θ in $T_{\text{FO}}^{\text{gen}}(K)$, the models of θ are cofinal in $T_{\text{FO}}^{\text{gen}}(K)$. In particular, θ has a finite model.

Conversely, suppose that every sentence in $T_{\text{FO}}^{\text{gen}}(K)$ has a finite model. By Lemma 3.1.7, for every A in K_n , enumerated by a tuple \bar{a} , there is an explicitly non-redundant quantifier-free formula $\varphi_A(\bar{x})$ such that if B is in K_n , then $B \models \varphi(\bar{b})$ if and only if the map $a_i \mapsto b_i$ is an isomorphism $A \rightarrow B$. And $\exists \bar{x} \varphi_A(\bar{x})$ is in $T_{\text{FO}}^{\text{gen}}(K)$, since the generic limit M of K is K -universal, and letting \bar{a}' enumerate a substructure of M isomorphic to A , $M \models \varphi_A(\bar{a}')$.

Let $\theta \in T_{\text{FO}}^{\text{gen}}(K)$. Then $\theta \wedge (\exists \bar{x} \varphi_A(\bar{x}))$ is in $T_{\text{FO}}^{\text{gen}}(K)$. So this sentence has a finite model $B \models \theta$, and there is a tuple \bar{b} from B such that $B \models \varphi_A(\bar{b})$. Letting $B' = \|\bar{b}\|$, B' is in K_n , since K is hereditary and $\varphi_A(\bar{x})$ is explicitly non-redundant, and $B' \models \varphi_A(\bar{b})$, so B' is isomorphic to A , and A embeds in B by the map $a_i \mapsto b_i$. This embedding is strong since K is hereditary. Hence the models of θ are cofinal in K , and by Corollary 4.5.17, K has a cofinal strongly robust subclass. \square

Part III

Logic

Chapter 5

Pseudofinite countably categorical theories

5.1 Zero-one laws and ultraproducts

We now define our third and final notion of limit for a class of finite structures. We are no longer interested in classes of embeddings between our structures, but rather the first-order sentences they satisfy, obtaining limit theories via zero-one laws or ultraproducts. Zero-one laws were first studied by Glebskii, Kogan, Liogon'kii, and Talanov [38] and Fagin [35], and the subject is now one of the core areas of finite model theory. It is possible to study zero-one laws for other logics (see the survey [27], for example), but we restrict ourselves to the first-order case here; all formulas, sentences, and theories are first-order. In the first-order context, the almost-sure theory of a class with a zero-one law agrees with the theory of an ultraproduct of structures from the class.

We now use the letter K for a class of finite structures, not for a strong embedding class as in Part II. In all examples of interest, K is closed under isomorphism, but we do not need to assume this explicitly. Later, K will be a small Fraïssé class, but then K is hereditary, and it is understood to be equipped with all embeddings between structures in K .

Definition 5.1.1. Let K be a class of finite structures. For each $n \in \omega$, let Ω_n be a **sample space** from K : a non-empty set of structures in K , equipped with a σ -algebra \mathcal{A} , according to which $\{A \in \Omega_n \mid A \models \varphi\} \in \mathcal{A}$ for every sentence φ . If a K -sample space Ω_n is finite or countably infinite, we will not specify the σ -algebra \mathcal{A} , instead taking the discrete algebra generated by the singletons. For each $n \in \omega$, let μ_n be a probability measure on Ω_n .

The class K has a **convergence law** (for the measures μ_n on the sample spaces Ω_n) if $\lim_{n \rightarrow \infty} \mu_n(\{A \in \Omega_n \mid A \models \varphi\})$ exists for every sentence φ . And K has a **zero-one law** (for the measures μ_n on the sample spaces Ω_n) if $\lim_{n \rightarrow \infty} \mu_n(\{A \in \Omega_n \mid A \models \varphi\}) = 0$ or 1 for every sentence φ . The **almost-sure theory** of K is the theory $T^{\text{a.s.}}(K)$ consisting of all sentences φ such that $\lim_{n \rightarrow \infty} \mu_n(\{A \in \Omega_n \mid A \models \varphi\}) = 1$.

This notion of convergence is reminiscent of the notion of convergence for quantifier-free densities defined in Section 1.1. But there are two main differences. First, we consider only the satisfaction of sentences (which have no free variables) by finite structures, instead of the satisfaction of formulas by tuples from finite structures. Second, our sentences have quantifiers, which range over the finite structures. Intuitively, we are measuring the convergence of global properties, rather than local behavior.

When K has a zero-one law, we obtain a complete theory $T^{\text{a.s.}}(K)$. This theory has a rather different character than the complete theory $\text{Th}_{\text{FO}}(\mu)$, associated to a convergent sequence of finite structures (or rather, to its limit, the ergodic structure μ) in Section 1.4, and the complete theory $T_{\text{FO}}^{\text{gen}}(K)$, associated to a strong embedding class satisfying relevant properties in Section 4.3. Namely, the criterion for one of the latter two theories to contain a sentence φ involves evaluating the quantifiers in φ on infinite structures (in Str_L and Dir_K , respectively), while for $T^{\text{a.s.}}(K)$, the quantifiers are evaluated on the finite structures themselves. We elaborate on this contrast in Example 5.1.9 below.

A third difference between convergence in the sense of Definition 5.1.1 and convergence for quantifier-free densities is that we sample finite structures from the spaces Ω_n at random according to the measures μ_n , instead of taking a fixed sequence of finite structures. But this is not a major difference: we could have also defined quantifier-free convergence relative to a sequence of measures μ_n on a sequence of sample spaces Ω_n in Section 1.1, by defining the density of $\varphi(\bar{x})$ in Ω_n to be the probability that $A \models \varphi(\bar{a})$, if we sample a structure A from μ_n and then sample a tuple \bar{a} uniformly at random from A . The main reason we did not work in this additional level of generality in Section 1.1 is that we would have lost the key property of disjoint-independence which led us to ergodic structures.

We could view both Definition 5.1.1 and Definition 1.1.4 as special cases of a more general definition: fixing a set of first-order formulas Δ , which may or may not have free variables and which may or may not have quantifiers, we say that a sequence of finite structures (or a sequence of probability measures on sample spaces of finite structures) converges if the probability that a tuple sampled from the finite structure B_n (or a random B_n , itself sampled from the sample space Ω_n according to the measure μ_n) satisfies $\varphi(\bar{x})$ converges for all $\varphi(\bar{x}) \in \Delta$. This idea has been pursued by Nešetřil and Ossona de Mendez [78], who have introduced a flexible class of infinitary limit objects, which they call “modelings”, for these general notions of convergence.

Definition 5.1.2. let T be a complete first-order theory, and let K be a class of finite structures. T is **K -pseudofinite** if every sentence φ in T has a model in K . If K is the class of all finite L -structures, we say T is **pseudofinite**.

It is well known that K -pseudofinite theories are exactly the theories of ultraproducts of structures in K . We now observe that a theory T is K -pseudofinite if and only if there is some sequence of measures on sample spaces from K such that T is the almost-sure theory of K with respect to these measures.

Theorem 5.1.3. *Let T be a complete theory, and let K be a class of finite structures. The following are equivalent:*

- (1) T is K -pseudofinite.
- (2) Some ultraproduct of structures in K is a model of T .
- (3) There is a sequence of measures μ_n on sample spaces Ω_n from K such that K has a zero-one law with respect to the μ_n , and $T = T^{\text{a.s.}}(K)$.

Proof. (1) \rightarrow (3): Enumerate T as $\{\varphi_n \mid n \in \omega\}$. For each $n \in \omega$, the sentence $\bigwedge_{i=1}^n \varphi_n$ has a finite model A_n . Let $\Omega_n = \{A_n\}$, and let μ_n be the unique probability measure on this singleton set. Then for any $\varphi_n \in T$, $\mu_m(\{A \in \Omega_m \mid A \models \varphi_n\}) = 1$ for all $m \geq n$, so K has a zero-one law (since T is complete) and $T = T^{\text{a.s.}}(K)$.

(3) \rightarrow (2): Let $\Omega = \bigcup_{n \in \omega} \Omega_n$, and let $\mathcal{F} = \{X \subseteq \Omega \mid \lim_{n \rightarrow \infty} \mu_n(X \cap \Omega_n) = 1\}$. Given X_1, \dots, X_k in \mathcal{F} , let n be large enough so that $\mu_n(X_i \cap \Omega_n) > 1 - 1/k$ for all i . Then $\mu\left(\left(\bigcap_{i=1}^k X_i\right) \cap \Omega_n\right) > 0$, and $\bigcap_{i=1}^k X_i$ is non-empty. So \mathcal{F} has the finite intersection property. Let \mathcal{U} be any ultrafilter on Ω extending \mathcal{F} . Since $\{A \in \Omega \mid A \models \varphi\} \in \mathcal{U}$ for all $\varphi \in T$, $(\prod_{A \in \Omega} A)/\mathcal{U} \models T$.

(2) \rightarrow (1): Suppose Ω is a set of structures in K and \mathcal{U} is an ultrafilter on Ω , such that $(\prod_{A \in \Omega} A)/\mathcal{U} \models T$. Then for all $\varphi \in T$, $\{A \in \Omega \mid A \models \varphi\}$ is in \mathcal{U} . In particular, it is non-empty, so φ has a model in K . \square

Of course, the measures constructed in the theorem are rather contrived. We've really just found a sequence of structures which satisfy larger and larger finite subsets of T . This is the fault of our very flexible notion of logical convergence, which allowed any measures on any sample spaces. For small classes K , the most well-studied sequences of measures are the uniform measures on the structures in K with domain $[n]$.

Notation 5.1.4. Let K be a class of finite structures. For $n \in \omega$, we denote by $K(n)$ the set of structures in K with domain $[n]$. We include the case $n = 0$: $K(0)$ consists of the empty structures in K .

Recall that a class K of finite structures is **small** if $K_n = \{A \text{ in } K \mid |A| = n\}$ is finite up to isomorphism (Definition 3.1.6). Note that Lemma 3.1.7 holds for small classes of finite structures, and if K is small, then $K(n)$ is finite for all n .

Definition 5.1.5. Let K be a small class of finite structures such that $K(n)$ is non-empty for all $n \in \omega$. K has a **uniform zero-one law** if it has a zero-one law for the uniform measures μ_n on the finite sample spaces $\Omega_n = K(n)$, defined by

$$\mu_n(X) = \frac{|X|}{|K(n)|}$$

for all $X \subseteq K(n)$.

Strictly speaking, this is usually called a uniform *labeled* zero-one law, since we are counting structures with domain labeled by $[n]$. If K has a zero-one law for the measures μ_n , when μ_n is the uniform measure on the set of isomorphism classes in K_n , then K is said to have a uniform *unlabeled* zero-one law. In many natural classes of finite structures, labeled and unlabeled zero-one laws are equivalent, because almost all finite structures in the class are rigid (see [40]).

The uniform measures on $K(n)$ are certainly natural, but in Sections 5.3 and 5.4, we will construct certain sequences of measures on $K(n)$ which are not uniform, but are nonetheless also natural in some sense. The following definition captures this naturality and makes a connection with Part I. We will return to this connection at the end of Section 5.5.

Recall that $\text{Str}_{L,n}$ is the space of L -structures with domain $[n]$. As described in Remark 1.2.17, an ergodic structure ν pushes forward to a measure ν_n on $\text{Str}_{L,n}$ via the restriction map $\upharpoonright_n: \text{Str}_L \rightarrow \text{Str}_{L,n}$. If Ω_n is a subspace of $\text{Str}_{L,n}$ and μ_n is a measure on Ω_n , then μ_n also pushes forward to a measure $i_*\mu_n$ on $\text{Str}_{L,n}$ via the inclusion map $i: \Omega_n \rightarrow \text{Str}_{L,n}$.

Definition 5.1.6. Let K be a class of finite structures. For all $n \in \omega$, let Ω_n be a sample space from K , such that Ω_n is a subspace of $\text{Str}_{L,n}$, and let μ_n be a probability measure on Ω_n . The sequence of measures $\langle \mu_n \rangle_{n \in \omega}$ is **coherent** if there is an ergodic structure ν such that the measures $i_*\mu_n$ and ν_n on $\text{Str}_{L,n}$ are equal for all n .

Remark 5.1.7. By Lemma 1.2.9, an ergodic structure ν is uniquely determined by the values $\nu(\llbracket \varphi(\bar{a}) \rrbracket)$ for all quantifier-free formulas $\varphi(\bar{x})$ and all tuples \bar{a} from ω . And since, for any N is greater than all the elements of \bar{a} , $\nu(\llbracket \varphi(\bar{a}) \rrbracket) = \nu_N(\llbracket \varphi(\bar{a}) \rrbracket)$, ν is uniquely determined by the measures $\langle \nu_n \rangle_{n \in \omega}$. Hence, if is an ergodic structure ν witnessing coherence of $\langle \mu_n \rangle_{n \in \omega}$, then ν is unique.

Exactly as in the arguments in Theorem 1.2.10 and Section 1.3, the measures $\langle \mu_n \rangle_{n \in \omega}$ cohere to an ergodic structure if and only if the following conditions hold:

- Agreement: For all quantifier-free formulas $\varphi(\bar{x})$, all tuples \bar{a} , and all $n \leq m$, both greater than all the elements of \bar{a} , we have $i_*\mu_n(\llbracket \varphi(\bar{a}) \rrbracket) = i_*\mu_m(\llbracket \varphi(\bar{a}) \rrbracket)$.
- Invariance: For all quantifier-free formulas $\varphi(\bar{x})$, all non-redundant tuples \bar{a} and \bar{b} , and all n , $i_*\mu_n(\llbracket \varphi(\bar{a}) \rrbracket) = i_*\mu_n(\llbracket \varphi(\bar{b}) \rrbracket)$.
- Disjoint-independence: For all quantifier-free formulas $\varphi(\bar{x})$ and $\psi(\bar{y})$, all disjoint tuples \bar{a} and \bar{b} , and all n , $i_*\mu_n(\llbracket \varphi(\bar{a}) \wedge \psi(\bar{b}) \rrbracket) = (i_*\mu_n(\llbracket \varphi(\bar{a}) \rrbracket))(i_*\mu_n(\llbracket \psi(\bar{b}) \rrbracket))$.

The next three examples motivate our concerns for the remainder of the chapter.

Example 5.1.8. Let \mathcal{G} be the class of all finite graphs. We saw in Example 4.2.7 that \mathcal{G} is a small Fraïssé class. Its Fraïssé limit is the random graph, whose countably categorical first-order theory we denote T_{RG} . We saw in Example 1.4.16 that there is a sequence of finite graphs (the Paley graphs) which converge to an ergodic structure μ which is almost

surely isomorphic to the random graph. This ergodic structure μ describes the Erdős–Rényi random graph process on ω with edge probability $1/2$.

For all n , μ pushes forward to a measure μ_n on $\text{Str}_{L,n}$, which describes the Erdős–Rényi random graph process on $[n]$ with edge probability $1/2$. The measures μ_n agree with the uniform measures on $\mathcal{G}(n)$ (since exactly half of the graphs in $\text{Str}_{L,n}$ have the edge aRb , exactly half of those graphs have some other edge cRd , etc.). So the uniform measures are coherent.

Moreover, \mathcal{G} has a uniform zero-one law with almost-sure theory $T^{\text{a.s.}}(\mathcal{G}) = T_{\text{RG}}$. For a simple proof, which we will generalize in Section 5.3, see [75, Theorem 2.4.4].

The class of finite graphs exemplifies the ideal behavior; in some sense all three kinds of limits we have discussed exist and agree. Essentially the same thing happens for the class of all finite L -structures when L is finite, which is also a small Fraïssé class, and for the kaleidoscope random graph class from Example 4.3.10, which is not. But as the next two examples show, things aren't always so perfect, even for small Fraïssé classes.

Example 5.1.9. Let \mathcal{L} be the class of all finite linear orders. \mathcal{L} is a small Fraïssé class. Its Fraïssé limit is the unique countable dense linear order without endpoints (isomorphic to (\mathbb{Q}, \leq)), whose countably categorical first-order theory is DLO.

Note that \mathcal{L} contains exactly one structure of size n up to isomorphism for all $n \in \omega$, namely the finite order L_n . The sequence $\langle L_n \rangle_{n \geq 1}$ is convergent (for quantifier-free densities). The easiest way to see this is to show that the sequence is convergent for substructure densities and use Theorem 1.1.9. Indeed, for all $m \leq n$, we have $p(L_n; L_m) = 1$, since every substructure of L_m of size n is isomorphic to L_n .

Let μ be the ergodic structure which is the limit of $\langle L_n \rangle_{n \geq 1}$. In this example, too, the measures μ_n on $\text{Str}_{L,n}$ agree with the uniform measures on $\mathcal{L}(n)$, so the sequence of uniform measures is coherent. Indeed, each structure in $\mathcal{L}(n)$ is the unique structure satisfying the formula $x_{\sigma(0)} < x_{\sigma(1)} < \cdots < x_{\sigma(n-1)}$ for some permutation σ of $[n]$, and each of these formulas must receive the same measure by invariance, so we have

$$\mu(\llbracket x_{\sigma(0)} < x_{\sigma(1)} < \cdots < x_{\sigma(n-1)} \rrbracket) = 1/n! = 1/|\mathcal{L}(n)|.$$

Now μ is almost surely isomorphic to (\mathbb{Q}, \leq) , so $\text{Th}_{\text{FO}}(\mu) = \text{DLO}$. One way to see this is to observe that we can realize the measure μ by an explicit process: sample countably many points independently from the uniform measure on $([0, 1], \leq)$ and take the induced substructure. The result of this process is almost surely a dense linear order without endpoints.

It is an interesting fact [4] that there is a unique ergodic structure which is almost surely isomorphic to (\mathbb{Q}, \leq) , and this structure and its reducts are the only countable structures with this property (up to interdefinability).

\mathcal{L} also has a uniform zero-one law, but its almost-sure theory is *not* DLO. The sentence asserting that there are greatest and least elements, that every non-greatest element has a successor, and that every non-least element has a predecessor, is true in every finite linear order. This sentence together with sentences asserting that there are at least n elements

(each of which has limiting probability 1) axiomatizes a complete theory, the theory of infinite discrete linear orders with endpoints.

As noted above, the difference between this almost-sure theory of \mathcal{L} and its limit theory for both measure and category, DLO, comes down to where the quantifiers are evaluated. DLO is not pseudofinite: for example, the sentence $\forall x \forall y ((x < y) \rightarrow \exists z (x < z < y))$ asserting density does not hold in any finite linear order of size at least 2. But this sentence is in $T_{\text{FO}}^{\text{gen}}(\mathcal{L})$ and $\text{Th}_{\text{FO}}(\mu)$. In the first case, this is because for any two points $a < a'$ in a finite linear order A , we can embed A into a larger finite linear order in which there is a point between a and a' . In the second case, this is because if we sample many points uniformly and independently from a very large linear order, with high probability, one of the points sampled lies strictly between the first two points sampled, and this probability converges to 1 as the number of points sampled and the size of the linear order grows.

In the case of linear orders, we find differences between the limit theories associated to the class, but at least the situation is completely understood. The next example shows there are apparently very simple classes for which we understand very little.

Example 5.1.10. Let \mathcal{G}_Δ be the class of all finite triangle-free graphs. \mathcal{G}_Δ is a small Fraïssé class. Its Fraïssé limit is the **generic triangle-free graph** M_Δ , also called the Henson graph, since it was first studied in [41].

\mathcal{G}_Δ has a uniform zero-one law. But Erdős, Kleitman, and Rothschild showed in [33] that, asymptotically, almost all triangle-free graphs are bipartite. That is, $T^{\text{a.s.}}(\mathcal{G}_\Delta)$ contains sentences asserting that there are no cycles of any odd length. This is in contrast to the generic theory $T_\Delta = \text{Th}(M_\Delta)$, since M_Δ is \mathcal{G}_Δ -universal, and hence contains cycles of all odd lengths ≥ 5 .

The uniform measures μ_n on $\mathcal{G}_\Delta(n)$ are not coherent. For example, $\mu_2(\llbracket x_0 R x_1 \rrbracket) = 1/2$, but $\mu_3(\llbracket x_0 R x_1 \rrbracket) = 3/7$. Of the 7 triangle-free graphs with domain [3], only 3 of them have an edge between 0 and 1.

Now the generic triangle-free graph has trivial group-theoretic definable closure, so, by Theorem 1.5.5, there is some ergodic structure ν which is almost surely isomorphic to M_Δ (in fact, such an ergodic structure was first constructed by Petrov and Vershik [81], and this construction was the inspiration for [2]). This ergodic structure ν gives rise to a coherent sequence of measures ν_n on $\text{Str}_{L,n}$ (which must be distinct from the uniform measures), but it is difficult to extract explicit descriptions of the ν_n from the construction of ν , and it is unknown whether \mathcal{G}_Δ has a convergence law for this sequence of measures, much less whether the almost-sure theory is equal to T_Δ .

So the natural approach to showing that T_Δ is pseudofinite, namely realizing it as the almost-sure theory for some sequence of measures, fails for the uniform measures and is beyond our reach for the coherent measures arising from an ergodic structure concentrating on M_Δ . In fact, the question of whether T_Δ is pseudofinite is a major open problem (see [22] and [23]).

Since the question of pseudofiniteness is already hard enough for countably categorical theories (as evidenced by Example 5.1.10), we will restrict our attention to these theories. And in Section 5.2, we will explain (via the canonical language, Definition 5.2.3) that studying countably categorical theories is essentially the same as studying the generic theories of small Fraïssé classes.

We have already done a bit of work in this direction: we connected K -pseudofiniteness of the generic theory of a small Fraïssé class K to the existence of a strongly robust subclass in Corollary 4.5.18.

A suggestive observation is that each of the motivating examples above have different statuses relative to model theoretic (Shelahian) dividing lines: T_{RG} is a simple theory [84], and in fact most known pseudofinite countably categorical theories are simple. On the other hand, DLO has the strict order property [86].

The argument that DLO is not pseudofinite can be generalized to any countably categorical theory with the strict order property. It seems that this fact is folklore, but being unaware of a reference, we will give a quick proof here.

Proposition 5.1.11. *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. Since interpretations preserve countable categoricity and pseudofiniteness, it suffices to show that no countably categorical partial order (P, \leq) 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 $P \models (\exists x \varphi(x)) \wedge (\forall x (\varphi(x) \rightarrow \exists y (x < y \wedge \varphi(y))))$. But this sentence cannot hold in any finite structure, since it implies the existence of an infinite increasing chain of elements satisfying $\varphi(x)$. \square

In [31], Džamonja and Shelah 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 for $n \geq 3$ in [88]. 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 simplicity and the Strict Order Property (SOP):

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

It is worth noting that SOP_2 also goes by the name TP_1 (the Tree Property of the first kind, see [58] for a discussion), and every theory which is known to have SOP_1 also has SOP_3 . So it is possible that $\text{NSOP}_1 \stackrel{?}{=} \text{NSOP}_2 = \text{NTP}_1 \stackrel{?}{=} \text{NSOP}_3$. The generic triangle-free graph has SOP_3 but NSOP_4 [88].

In Section 5.4, we will study two examples of countably categorical unsimple NSOP₁ theories and show that both are pseudofinite. To our knowledge, these are the first unsimple countably categorical theories which have been proven to be pseudofinite (though Shelah earlier asserted in an unpublished note [89] that one of them, T_{feq}^* , is pseudofinite). Moreover, as far as we know, for *every* countably categorical theory T which has SOP₁ and NSOP (e.g. the generic triangle-free graph), it is an open problem whether T is pseudofinite.

5.2 Small Fraïssé classes

Recall that a small Fraïssé class K is a hereditary strong embedding class with the joint embedding property and the amalgamation property, such that K_n is finite for all $n \in \omega$. As in Chapter 4, we assume that every Fraïssé class K is pruned and non-empty; equivalently, by Remark 4.1.8 and hereditariness, K_n is non-empty for all n .

We are now in the setting of classical Fraïssé theory. By Theorem 4.2.2 and Remark 4.2.5, every Fraïssé class K has a Fraïssé limit M which is K -universal and K -homogeneous. By Corollary 4.4.8, $T_{\text{FO}}^{\text{gen}}(K) = \text{Th}(M)$ is countably categorical and has quantifier elimination.

Notation 5.2.1. If K is a small Fraïssé class, we write M_K for the Fraïssé limit of K and T_K for the generic first-order theory of K .

The canonical language gives us a bridge between Fraïssé limits of small Fraïssé classes and general countably categorical theories.

Definition 5.2.2. The **age** of a structure M , $\text{Age}(M)$, is the class of all finite structures which are isomorphic to substructures of M .

Definition 5.2.3. 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 $p(\bar{x})$ realized in M . We make M into an L' -structure M' in the natural way by setting $M' \models R_p(\bar{a})$ if and only if \bar{a} realizes $p(\bar{x})$ in M , and we set $T' = \text{Th}_{L'}(M')$.

Proposition 5.2.4. *Let T , T' , M , and M' be as in Definition 5.2.3. Let $K = \text{Age}(M')$. Then T and T' are interdefinable theories, T' has quantifier elimination, K is a small Fraïssé class, and M' is its Fraïssé limit.*

Proof. Since T is countably categorical, every n -type $p(\bar{x})$ consistent with T is realized in M and isolated by a formula $\theta_p(\bar{x})$. We can define the atomic L' -formula $R_p(\bar{x})$ by the L -formula $\theta_p(\bar{x})$. Conversely, if $R(\bar{x})$ is a relation symbol in L , there are finitely many n -types $p_1(\bar{x}), \dots, p_k(\bar{x})$ which contain $R(\bar{x})$, so $R(\bar{x})$ is equivalent to $\bigvee_{i=1}^k \theta_{p_i}(\bar{x})$, and we can define the atomic L -formula $R(\bar{x})$ by the L' -formula $\bigvee_{i=1}^k R_{p_i}(\bar{x})$.

Note that tuples \bar{a} and \bar{b} have the same type in M if and only if they have the same quantifier-free type in M' if and only if the map $a_i \mapsto b_i$ is an isomorphism $A \rightarrow B$, where $A = \|\bar{a}\|$ and $B = \|\bar{b}\|$.

Now M' is a K -direct limit (where we equip K with the class of all embeddings between structures in K), since $K = \text{Age}(M')$. And since there are only finitely many n -types realized in M for all $n \in \omega$, K is small.

We will show that M' is K -universal and K -ultrahomogeneous. This will establish that K is a Fraïssé class and M' is its Fraïssé limit. Since K is small, this also implies that $T' = \text{Th}(M')$ has quantifier elimination. K -universality is clear, since $K = \text{Age}(M')$. For K -ultrahomogeneity, suppose we have a substructure A of M' , enumerated by \bar{a} , and an embedding $f: A \rightarrow M'$. Then $B = \|f(\bar{a})\|$ is isomorphic to A , so $\text{tp}_L(\bar{a}) = \text{tp}_L(f(\bar{a}))$. By homogeneity of M , there is an automorphism σ of M extending f , and σ is also an automorphism of M' by interdefinability. \square

The topic of Section 5.3 is a family of properties called disjoint n -amalgamation. We begin here with the first non-trivial case, $n = 2$, phrased in the language of Fraïssé limits. It is equivalent to trivial acl, another important property for us.

Definition 5.2.5. Let K be a Fraïssé class. K has the **disjoint amalgamation property** (or **disjoint 2-amalgamation**) if for all A, B, C in K and embeddings $f: A \rightarrow B$ and $g: A \rightarrow C$, there exists D in K and embeddings $f': B \rightarrow D$ and $g': C \rightarrow D$ such that $f' \circ f = g' \circ g$, and the images of B and C in D are disjoint over the image of A in D : $(f' \circ f)(A) = (g' \circ g)(A) = f'(B) \cap g'(C)$.

In many sources (e.g. [17] and [43]), the disjoint amalgamation property is called the strong amalgamation property.

Definition 5.2.6. Let M be a countable structure and A a finite subset of M . $\text{Aut}(M/A)$ is the group of automorphisms of M which fix A pointwise. The **group-theoretic algebraic closure** of A in M , $\text{acl}(A)$, is the set of all $b \in M$ such that b has only finitely many images under the action of $\text{Aut}(M/A)$. M has **trivial acl** if $\text{acl}(A) = A$ for all finite $A \subseteq M$.

Theorem 5.2.7 ([17, (2.15)]). *A small Fraïssé class K has the disjoint amalgamation property if and only if M_K has trivial acl.*

Remark 5.2.8. In Part I, we defined trivial (group-theoretic) dcl (Definition 1.5.4) for a countable structure M . It is clear that if M has trivial acl, then M has trivial dcl, since a point b which is fixed by all automorphisms in $\text{Aut}(M/A)$ has finitely many images (just one) under these automorphisms. But the converse is true as well: if $b \notin A$ has finitely many images $\{b, b_1, \dots, b_n\}$ under $\text{Aut}(M/A)$, then letting $B = A \cup \{b_1, \dots, b_n\}$, b is fixed by all automorphisms in $\text{Aut}(M/B)$. So the notions are equivalent; we referred to trivial dcl in Part I to be consistent with [2], and we refer to trivial acl in this chapter to be consistent with the literature on homogeneous structures.

In the case that $T = \text{Th}(M)$ is countably categorical, it is an easy consequence of the Ryll–Nardzewski theorem and homogeneity of M that for every finite set A from M , the group-theoretic acl of A in M agrees with the usual model theoretic acl of A , i.e. the set of all b which satisfy some formula $\varphi(x, \bar{a})$ with parameters from A such that $\varphi(M, \bar{a})$ is finite

(see [43, Corollary 7.3.4]). Then M has trivial group-theoretic acl if and only if T has trivial model-theoretic acl in the sense that $\text{acl}(A) = A$ for any set $A \subseteq N \models T$; indeed, a failure of trivial model-theoretic acl in any model N is witnessed by a finite set A (the parameters in the algebraic formula), and hence is witnessed over a finite set A' in M realizing the type of A .

We now give an explicit pithy Π_2 (see Section 2.2) axiomatization of T_K for any small Fraïssé class K . For each structure A in $\text{Str}_{L,n}$, Lemma 3.1.7 provides a quantifier-free formula $\varphi_A(\bar{x})$, such that $A \models \varphi_A(0, \dots, n-1)$, and if $B = \|\bar{b}\|$ is in K , then $B \models \varphi_A(\bar{b})$ if and only if $i \mapsto b_i$ is an isomorphism $A \rightarrow B$. The formula $\varphi_A(\bar{x})$ is the conjunction of enough of the quantifier-free diagram of A to distinguish it from the other structures in $K(n)$.

Definition 5.2.9. (A, B) is a **one-point extension** in K if $A \in K(n)$ and $B \in K(n+1)$ for some n , and A is the substructure of B with domain $[n]$. We include the case $n = 0$, when A is an empty structure.

Theorem 5.2.10. *Let K be a small Fraïssé class. The generic theory T_K can be axiomatized by universal axioms and one-point extension axioms:*

- For all $n \in \omega$ and all A in $\text{Str}_{L,n}$ but not in K , $\forall x_1, \dots, x_n \neg \varphi_A(\bar{x})$.
- For all one-point extensions (A, B) in K , $\forall \bar{x} \exists y \varphi_A(\bar{x}) \rightarrow \varphi_B(\bar{x}, y)$.

Proof. It suffices to show that if M is a countable model of T_K , then $\text{Age}(M) = K$ and M is K -homogeneous.

Suppose A is a finite substructure of M , enumerated by \bar{a} . Then $M \models \varphi_A(\bar{a})$. If A were not in K , this would contradict the axiom $\forall \bar{x} \neg \varphi_A(\bar{x})$. Hence $\text{Age}(M) \subseteq K$. Thus, for B in K and any tuple \bar{b} from M , if $M \models \varphi_B(\bar{b})$, then $\|\bar{b}\|$ is isomorphic to B .

Since K is hereditary and has the joint embedding property, it contains exactly one empty structure E , which embeds in all other structures in K . E is also the substructure of M induced on the empty set, since $\text{Age}(M) \subseteq K$. So if we show that M is K -homogeneous, this will also show that $\text{Age}(M) = K$ (any A in K embeds in M over the unique embedding of its empty substructure E in M).

Suppose A and B are in K and $f: A \rightarrow M$ and $g: A \rightarrow B$ are embeddings. We may assume that $A \in K(n)$, $B \in K(m)$ for some $m \geq n$, and g is the inclusion. For $0 \leq i \leq m-n$, let B_i be the substructure of B with domain $[n+i]$, so $B_0 = A$. If B_i embeds in M , with image enumerated by \bar{b} , then $M \models \varphi_{B_i}(\bar{b})$, so by the one-point extension axiom for $(B_i, B_{(i+1)})$, there is some c with $M \models \varphi_{B_{(i+1)}}(\bar{b}, c)$. Then $\|\bar{b}c\|$ is an isomorphic copy of $B_{(i+1)}$, embedded over the image of B_i . By induction, we can embed each B_i in M over A , ending with $B_{(m-n)} = B$. \square

Example 5.2.11. In the case of \mathcal{G}_Δ , the triangle-free graphs, the universal part of T_Δ can be axiomatized by a single universal sentence with only three quantifiers, asserting that the

graph relation R is symmetric and anti-reflexive, and that there are no triangles. We pointed out in Example 5.1.10 that it is unknown whether T_Δ is pseudofinite, and the difficulty lies in determining whether the one-point extension axioms have finite models. In fact, while examples are known of finite triangle-free graphs satisfying the one-point extension axioms (A, B) for $|A| \leq 3$, it is already an open problem whether this can be extended to $|A| \leq 4$. See [23] for a detailed discussion of this problem.

Definition 5.2.12. $T_{K,n}$ is the (incomplete) theory axiomatized by

- The sentences in the universal theory of K with at most n universal quantifiers.
- 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 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 conclude this section with two useful observations about relationships between small Fraïssé classes.

Definition 5.2.13. Let K and K' be small Fraïssé classes in the languages L and L' , respectively, such that $L \subseteq L'$. We say that K' is a **Fraïssé expansion** of K if

- (1) $K = \{A \upharpoonright L \mid A \in K'\}$.
- (2) For every one-point extension (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 5.2.14. *Let K and K' be small Fraïssé classes in the languages L and L' , respectively, such that $L \subseteq L'$. 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 $M_{K'} \upharpoonright L = M_K$. Then $K = \text{Age}(M_K) = \{A \upharpoonright L \mid A \in K'\}$, since $\text{Age}(M_{K'}) = K'$. 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 M_K , 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 $M_{K'} \upharpoonright L$ satisfies the theory T_K . It clearly satisfies the universal part, since $\text{Age}(M_{K'} \upharpoonright L) = \{A \upharpoonright L \mid A \in K'\} = K$. For the extension axioms, suppose (A, B) is a one-point extension, and we have a copy of A in $M_{K'} \upharpoonright L$. 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 5.2.15. A Fraïssé class K is **filtered** by a chain $K_0 \subseteq K_1 \subseteq K_2 \subseteq \dots$ if each K_n is a Fraïssé class (all in the same language L), and $\bigcup_{n \in \omega} K_n = K$.

Theorem 5.2.16. *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 5.2.10. Since each K_n is a subclass of K , every universal axiom of T_K is also in T_{K_n} . Let (A, B) be a one-point extension. For large enough n , the structures A and B are both in K_n , so (A, B) is also a one-point extension in K_n , and the one-point extension axiom for (A, B) is in T_{K_n} . \square

5.3 Disjoint n -amalgamation

Let T be a complete theory and A a set of parameters in a model of T . We extend the notion of non-redundancy to types with parameters: a type $p(\bar{x})$ over A is **non-redundant** if it contains the formulas $x_i \neq x_j$ for all $i \neq j$ and $x_i \neq a$ for all i and all $a \in A$.

Definition 5.3.1. A family $\mathcal{F} \subseteq \mathcal{P}([n])$ of subsets of $[n]$ is **downwards closed** if $S' \in \mathcal{F}$ whenever $S' \subseteq S$ and $S \in \mathcal{F}$. Given a downwards closed family of subsets $\mathcal{F} \subseteq \mathcal{P}([n])$, and variables $\bar{x}_0, \dots, \bar{x}_{n-1}$, 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 non-redundant type over A in the variables $\bar{x}_S = \{\bar{x}_i \mid i \in S\}$, and $p_{S'} \subseteq p_S$ when $S' \subseteq S$. Here each \bar{x}_i is a tuple of variables, possibly empty or infinite.

We denote by $\mathcal{P}^-([n])$ the set of all proper subsets of $[n]$. For $n \geq 2$, a **disjoint n -amalgamation problem** is a coherent $\mathcal{P}^-([n])$ -family of types over some set A . A **basic disjoint n -amalgamation problem** is a disjoint n -amalgamation problem over the empty set in the singleton variables x_0, \dots, x_{n-1} . If we replace $\mathcal{P}^-([n])$ by another downwards closed family of subsets \mathcal{F} , we call the amalgamation problem **partial**.

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 solution is determined by a non-redundant 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.

Remark 5.3.2. In any coherent \mathcal{F} -family of types over A , the type p_\emptyset is a 0-type in the empty tuple of variables. It simply specifies the elementary diagram of the parameters A .

Remark 5.3.3. To specify a (basic) disjoint n -amalgamation problem, it would be sufficient to give the types p_S for all S such that $|S| = n - 1$ and check that they agree on intersections, in the sense that $p_S \upharpoonright \bar{x}_{S \cap S'} = p_{S'} \upharpoonright \bar{x}_{S \cap S'}$ for all S and S' . However, it is sometimes notationally convenient to keep the intermediate stages around.

Disjoint amalgamation up to level n can be used to find solutions to partial disjoint n -amalgamation problems as well.

Lemma 5.3.4. *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. We 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 $\bar{x}_0, \dots, \bar{x}_{n-1}$, i.e. a coherent \mathcal{F} -family of types $\{p_S \mid S \in \mathcal{F}\}$, with $\mathcal{F} \subseteq \mathcal{P}^-([n])$ downwards closed. By induction on $1 \leq k \leq n$, we 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 $k = 1$, for any $i \in [n]$ such that $\{i\} \notin \mathcal{F}$, $i \notin S$ for all $S \in \mathcal{F}$, so the original \mathcal{F} -family of types says nothing about the variables \bar{x}_i . We add $\{i\}$ into \mathcal{F}_1 and choose any non-redundant type $p_{\{i\}}$ over A in the variables \bar{x}_i . If $\emptyset \notin \mathcal{F}$ (which only happens if \mathcal{F} is empty) we also add \emptyset into \mathcal{F}_1 , along with the unique 0-type p_\emptyset consisting of 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. For any set $S \subseteq [n]$ with $|S| = k$ such that $S \notin \mathcal{F}_{k-1}$, 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 disjoint k -amalgamation, we can find a non-redundant type p_S in the variables \bar{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 5.3.14. We are largely interested in theories with disjoint n -amalgamation for all n , and in this case the two notions agree.

Proposition 5.3.5. *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 n -amalgamation problem $\{p_S \mid S \in \mathcal{P}^-([n])\}$ if and only if the partial type

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

is consistent (actually, we could omit the formulas asserting non-redundancy when $n > 2$, since they are already contained in the types p_S with $|S| = 2$). Hence, by compactness, we can reduce to the case that A is finite and each tuple of variables \bar{x}_i is finite.

Let $N = |A| + \sum_{i=1}^n l(\bar{x}_i)$. Introduce variables y_0, \dots, y_{N-1} , where $y_0, \dots, y_{|A|-1}$ enumerate A and the remaining variables relabel the x variables. Now each type p_S in our family 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 subsets of variables, we obtain a partial disjoint N -amalgamation problem over the empty set in the

singleton variables y_0, \dots, y_{N-1} . By Lemma 5.3.4 and basic disjoint N -amalgamation, this partial amalgamation problem has a solution, a type $p_{[N]}(y_0, \dots, y_{N-1})$ 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. \square

To make the bridge between these definitions and Fraïssé theory, it will be convenient to identify the structures in a small Fraïssé class K with their quantifier-free types: for $A \in K(n)$, $\text{qftp}(A) = \{\varphi(x_0, \dots, x_{n-1}) \mid \varphi \text{ is quantifier-free, and } A \models \varphi(0, \dots, n-1)\}$. Since T_K has quantifier elimination, and since a non-redundant quantifier-free n -type is consistent with T_K if and only if it is $\text{qftp}(A)$ for some $A \in K(n)$, we can further identify $K(n)$ with the set of complete non-redundant n -types over the empty set consistent with T_K .

Remark 5.3.6. A small 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 $f: A \rightarrow B$ and $g: A \rightarrow C$, we take A to be the base set of parameters, so $p_\emptyset = \text{qftp}(A)$, and we set $p_{\{0\}}(\bar{x}_0) = \text{qftp}((B \setminus A)/A)$ and $p_{\{1\}}(\bar{x}_1) = \text{qftp}((C \setminus A)/A)$, identifying A with its images in B and C under f and g . These quantifier-free types determine complete types relative to T_K , and a solution to this disjoint 2-amalgamation problem determines a structure D in K into which B and C embed disjointly over the image of A .

The converse is a similar translation from disjoint 2-amalgamation of types $p_{\{0\}}(\bar{x}_0)$ and $p_{\{1\}}(\bar{x}_1)$ over A to disjoint amalgamation of structures in K , at least when the parameter set A and the variable contexts \bar{x}_0 and \bar{x}_1 are finite. The general infinitary case follows by compactness, as in Proposition 5.3.5.

Notation 5.3.7. Given a small Fraïssé class 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_0, \dots, x_{n-1} , where each type p_S corresponds to a structure A_S in K of size $|S|$. We write $K(n, P) = \{p_{[n]}(x_0, \dots, x_{n-1}) \in K(n) \mid p_S \subseteq p_{[n]}\}$ 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 non-empty for all P .

Example 5.3.8. The class \mathcal{G}_Δ of triangle-free graphs has disjoint 2-amalgamation: if A embeds in B and C , 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 non-redundant 2-types determined by x_1Rx_2 , x_2Rx_3 , and x_1Rx_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 non-redundant n -tuple \bar{a} , there is some subtuple \bar{b} of length k such that $\neg R(\bar{b})$ holds. Note that \mathcal{G}_Δ is K_3^2 .

For $n > k$, K_n^k satisfies basic disjoint m -amalgamation for $m < n$, but fails basic disjoint n -amalgamation. 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 5.3.9. 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 E with k infinite classes. Transitivity is a failure of disjoint 3-amalgamation: the non-redundant 2-types determined by x_1Ex_2 , x_2Ex_3 , and $\neg x_1Ex_3$ 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 in its canonical language (Definition 5.2.3) has a reduct to a 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 . The age K of this hypergraph M 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. Further, M is K -homogeneous, so K is a small Fraïssé class. It is easy to see that T_K fails to have disjoint 4-amalgamation, though it is a reduct of T_{RG} , which has disjoint n -amalgamation for all n . For more information on this example, see [73], where it is called the homogeneous two-graph. More examples of this kind can be found in the literature on reducts of homogeneous structures, e.g. [93].

The general notion of n -amalgamation has been studied in other model-theoretic contexts, usually in the form of *independent* n -amalgamation. An **independence relation** is a ternary relation, often written $a \downarrow_A B$, where a is a finite tuple and A and B are sets. All tuples and sets come from some highly saturated model of T , and we follow the standard notational convention in model theory of concatenating the names of sets and tuples when we mean to take unions. For example, we write Aa for the set $A \cup \|a\|$.

Given some independence relation \downarrow , the main example being non-forking independence \downarrow^f in a simple theory, an independent n -amalgamation problem is given by a coherent $\mathcal{P}^-([n])$ -family of types over A , with the non-redundancy condition replaced by the condition that for all $S \in \mathcal{P}^-([n])$, any realization $\{\bar{a}_i \mid i \in S\}$ of $p_S(\bar{x}_S)$ is an independent set over A with respect to \downarrow , i.e. $\bar{a}_i \downarrow_A A\bar{a}_1 \dots \bar{a}_{i-1}\bar{a}_{i+1} \dots \bar{a}_n$ for all i . See [19] for the definition of non-forking independence and background on independence relations and simple theories.

In the case $n = 3$, independent 3-amalgamation over models is often called the independence theorem. Kim and Pillay showed that the independence theorem, along with a few other natural properties, characterizes non-forking independence in simple theories.

Theorem 5.3.10 ([60], Theorem 4.2). *Let T be a complete theory and \downarrow an independence relation. Suppose that \downarrow satisfies the following properties:*

- (Invariance) *If $a \downarrow_A B$ and $\text{tp}(a'A'B') = \text{tp}(aAB)$, then $a' \downarrow_{A'} B'$.*
- (Local character) *For all a, B , there is $A \subseteq B$ such that $|A| \leq |T|$ and $a \downarrow_A B$.*

- (Finite character) $a \downarrow_A B$ if and only if for every finite tuple b from B , $a \downarrow_A Ab$.
- (Extension) For all a , A , and B , there is a' such that $\text{tp}(a'/A) = \text{tp}(a/A)$ and $a' \downarrow_A B$.
- (Symmetry) If $a \downarrow_A Ab$, then $b \downarrow_A Aa$.
- (Transitivity) If $A \subseteq B \subseteq C$, then $a \downarrow_A B$ and $a \downarrow_B C$ if and only if $a \downarrow_A C$.
- (Independence theorem) Let $M \models T$ be a model, A and B sets, and a and a' tuples such that $\text{tp}(a/M) = \text{tp}(a'/M)$. If $A \downarrow_M B$, $a \downarrow_M A$, and $a' \downarrow_M B$, then there exists a'' such that $\text{tp}(Aa''/M) = \text{tp}(Aa/M)$, $\text{tp}(Ba''/M) = \text{tp}(Ba'/M)$, and $a'' \downarrow_M AB$.

Then T is simple, and \downarrow is non-forking independence.

Disjoint n -amalgamation is a strong form of independent n -amalgamation, where the relevant independence relation is the disjointness relation \downarrow^{\neq} , defined by $A \downarrow^{\neq}_C B$ if and only if $A \cap B \subseteq C$. We say a theory with trivial acl has **trivial forking** if $\downarrow^{\neq} = \downarrow$.

Theorem 5.3.11. *A countably categorical theory T with disjoint 2-amalgamation (equivalently, trivial acl) and disjoint 3-amalgamation is simple with trivial forking.*

Proof. The equivalence of disjoint 2-amalgamation and trivial acl for countably categorical theories follows from the same equivalence for theories of Fraïssé limits (Theorem 5.2.7 and Remark 5.3.6), using the canonical language.

We can use Theorem 5.3.10 to show that $\downarrow^{\neq} = \downarrow$. Most of the conditions are straightforward to check, so I'll only remark on a few of them. For local character, we can take $A = a \cap B$, so A is finite and $a \downarrow^{\neq}_A B$. For extension, we find a' by realizing the type $\text{tp}(a/A) \cup \{a_i \neq b \mid a_i \text{ from } a \text{ and } b \in B\}$. This is consistent by trivial acl and compactness. Finally, for the independence theorem, we apply disjoint 3-amalgamation to amalgamate the three 2-types $p_{\{12\}} = \text{tp}(aA/M)$, $p_{\{13\}} = \text{tp}(a'B/M)$, $p_{\{23\}} = \text{tp}(AB/M)$ (first removing any redundant elements of M from a , a' , A , and B). \square

Remark 5.3.12. Motivated by the fact that many examples of simple theories (such as T_{RG} and ACFA [21]) satisfy independent n -amalgamation for $n \geq 3$, Kolesnikov [63] and Kim, Kolesnikov, and Tsuboi [59] developed notions of n -simplicity for $1 \leq n \leq \omega$, where 1-simplicity coincides with simplicity. Generalizing Theorem 5.3.11, for any $n \geq 3$, if a countably categorical theory T has disjoint k -amalgamation for all $2 \leq k \leq n$, then it is $(n - 2)$ -simple, 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, independent n -amalgamation of models goes by the name “excellence” (see [10], for example). Disjoint n -amalgamation for classes of finite structures has also been studied by Baldwin, Koerwien, and Laskowski with applications to AECs [11]. And in the context of stable theories, Goodrick, Kim, and Kolesnikov have uncovered a connection between existence and uniqueness of independent

n -amalgamation and definable polygroupoids [39], generalizing earlier work of Hrushovski on independent 3-amalgamation and groupoids [46].

We will now prove our theorem connecting disjoint n -amalgamation to pseudofiniteness. The theorem is stated in a fine-grained way: amalgamation just up to level n gives pseudofiniteness of the theory $T_{K,n}$ (see Definition 5.2.12). 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 5.3.4, 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.

Formally, we construct a probability measure μ_N on $\text{Str}_{L,N}$ for all $N \in \omega$. In the case that T_K has basic disjoint n -amalgamation for all n , the μ_N are actually measures on $K(N)$, and K has a zero-one law with respect to these measures, with almost-sure theory T_K .

The probabilistic calculation is a straightforward generalization of the one used in the classical proofs of the zero-one laws for graphs and L -structures (e.g. [43, Lemma 7.4.6]). The key point is that disjoint amalgamation allows 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 $A \cap B$.

Remark 5.3.13. This may remind the reader of the independence condition in the AHK representations of Section 1.3. Indeed, the sequence of measures $\langle \mu_N \rangle_{N \in \omega}$ cohere (in the sense of Definition 5.1.6) to an ergodic structure μ which is almost-surely isomorphic to the Fraïssé limit M_K ; see Theorem 5.5.19 below. And we can view the probabilistic construction in the proof of Theorem 5.3.14 as describing an explicit AHK representation of μ .

Theorem 5.3.14. *Let K be a small 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 T_K is K -pseudofinite.*

Proof. We will define a probability measure μ_N on $\text{Str}_{L,N}$ for each $N \in \omega$ by describing a probabilistic construction of a structure M_N with domain $[N]$. Recall Notation 5.3.7 and our identification of structures in $K(k)$ with complete k -types consistent with T_K .

We assign k -types to each subset of size k from $[N]$ by induction. When $k = 0$, there is no choice: by heredity and the joint embedding property, there is a unique 0-type in $K(0)$. When $k = 1$, for each $i \in [N]$, choose the 1-type of $\{i\}$ uniformly at random from $K(1)$. Having assigned l -types to all subsets of size l with $l < k$, we wish to assign k -types. For each k -tuple i_1, \dots, i_k of distinct elements from $[N]$, let $P = \{p_S \mid S \in \mathcal{P}^-([k])\}$ be the collection of types assigned to all proper subtuples: $p_S(\bar{x}_S) = \text{qftp}(\{i_j \mid j \in S\})$. If T_K has basic disjoint k -amalgamation, $K(k, P)$ is non-empty and finite, and we may choose the k -type of i_1, \dots, 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, \dots, i_r from $[N]$ containing at least $n + 1$ distinct elements, we set $R(i_1, \dots, i_r)$ with probability $1/2$. The resulting L -structure M_N 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.

We claim that if θ is any of the axioms of $T_{K,n}$ (in the bounded case) or T_K (in the unbounded case), then $\lim_{N \rightarrow \infty} \mu_N(\llbracket \theta \rrbracket) = 1$.

Each universal axiom θ has the form $\forall x_1, \dots, x_k \psi(\bar{x})$ ($k \leq n$ in the bounded case), where $\psi(\bar{x})$ is quantifier-free and true on all k -tuples from structures in K . Since all substructures of M_N of size at most k are in K , θ is always satisfied by M_N , so $\mu_N(\llbracket \theta \rrbracket) = 1$ for all N .

Now suppose that θ is the one-point extension axiom $\forall \bar{x} \exists y \varphi_A(\bar{x}) \rightarrow \varphi_B(\bar{x}, y)$. Let \bar{a} be a non-redundant $|A|$ -tuple from $[N]$ and b any other element from $[N]$. Conditioning on the event that $M_N \models \varphi_A(\bar{a})$, there is a positive probability ε that $M_N \models \varphi_B(\bar{a}, b)$.

Indeed, in the unbounded case, or when $|A| < n$ in the bounded case, φ_B specifies the $|B|$ -type of the tuple $\bar{a}b$ 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 $\bar{a}b$ 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 $1 \leq k \leq |B|$. In the bounded case, when $|A| \geq n$, the above reasoning applies for the subtuples of $\bar{a}b$ of length at most n . On longer tuples, since φ_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 φ_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 φ_B).

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

Now we compute the probability that θ is *not* satisfied by M_N . Conditioned on the event that $M_N \models \varphi_A(\bar{a})$, the probability that $M_N \not\models \exists y \varphi_B(\bar{a}, y)$ 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 φ_B . Removing the conditioning, the probability that $M_N \not\models \exists y \varphi_A(\bar{a}) \rightarrow \varphi_B(\bar{a}, y)$ for any given \bar{a} is at most $(1 - \varepsilon)^{N - |A|}$, since the formula is vacuously satisfied when \bar{a} does not satisfy φ_A . Finally, there are $N^{|A|}$ possible tuples \bar{a} , so the probability that $M_N \not\models \forall \bar{x} \exists y \varphi_A(\bar{x}) \rightarrow \varphi_B(\bar{x}, y)$ is at most $N^{|A|} (1 - \varepsilon)^{N - |A|}$. Since $|A|$ is constant, the exponential decay dominates the polynomial growth, and $\lim_{N \rightarrow \infty} \mu_N(\llbracket \neg \theta \rrbracket) = 0$, so $\lim_{N \rightarrow \infty} \mu_N(\llbracket \theta \rrbracket) = 1$.

To conclude, any sentence $\psi \in T_{K,n}$ is a logical consequence of finitely many of the axioms $\theta_1, \dots, \theta_m$ considered above. We need only pick N large enough so that $\mu_N(\llbracket \theta_i \rrbracket) > 1 - \frac{1}{m}$ for all i . Then $\mu_N(\llbracket \bigwedge_{i=1}^m \theta_i \rrbracket) > 0$, so the conjunction $\bigwedge_{i=1}^m \theta_i$, and hence also ψ , has a model of size N . In the unbounded case, our construction ensures that this model is in K . \square

Corollary 5.3.15. *Any countably categorical theory T with disjoint n -amalgamation for all $n \geq 2$ is pseudofinite.*

Proof. Let T' be the version 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 of a small Fraïssé class and has basic disjoint n amalgamation for all n , so by Theorem 5.3.14, it is pseudofinite. \square

The observation that disjoint n -amalgamation is sufficient for pseudofiniteness generalizes and unifies a number of earlier observations. Here are a few examples.

- (1) Oberschelp [79] identified an unusual syntactic condition which is sufficient for a small class of finite structures K to have a zero-one law with respect to the uniform measures. Under this condition, the uniform measures are coherent, K is a Fraïssé class, and the almost-sure theory of K agrees with the generic theory of K . A universal sentence is called **parametric** if it is of the form $\forall x_1, \dots, x_n (\bigwedge_{i \neq j} x_i \neq x_j \rightarrow \varphi(\bar{x}))$ where φ is a Boolean combination of atomic formulas $R(y_1, \dots, y_m)$ such that each variable x_i appears among the y_j . For example, reflexivity $\forall x R(x, x)$ and symmetry $\forall x, y R(x, y) \rightarrow R(y, x)$ are parametric conditions, while transitivity $\forall x, y, z (R(x, y) \wedge R(y, z)) \rightarrow R(x, z)$ is not a parametric condition, since each atomic formula appearing in the Boolean combination involves only 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 non-redundant n -tuples and m -tuples are totally independent when $n \neq m$.

- (2) In their work on the random simplicial complex, Brooke-Taylor & Testa [16] introduced the notion of a local Fraïssé class and showed that the generic theory of a small local Fraïssé class is pseudofinite, by methods similar to those in the proof of Theorem 5.3.14. A universal sentence is called **local** if it is of the form $\forall x_1, \dots, x_n (R(x_1, \dots, x_n) \rightarrow \psi(\bar{x}))$, where R is a relation in the language and ψ is quantifier-free. A **local Fraïssé class** is the class of finite models of a set of local axioms.

Again, any local Fraïssé class has disjoint 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.

- (3) In [7], Ahlman investigated countably categorical theories in a finite 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 that $\text{acl}^{\text{eq}}(\emptyset) = \emptyset$, this agrees with what we call a simple theory with trivial forking above.

Ahlman shows that in such a theory T there is a \emptyset -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 we 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.

Corollary 5.3.16. *If a countably categorical theory T admits an expansion with disjoint n -amalgamation for all n , then T is simple and pseudofinite with trivial acl.*

Proof. This follows immediately from Theorem 5.3.11, Corollary 5.3.15, and Remark 5.2.8, since simplicity, pseudofiniteness, and trivial acl are preserved under reduct. \square

Corollary 5.3.16 applies, for example, to the theory of the homogeneous two-graph described in Example 5.3.9.

Corollary 5.3.17. *If a small Fraïssé class K is filtered by $\{K_n \mid n \in \omega\}$ and each generic theory T_{K_n} is (K_n) -pseudofinite, then the generic theory T_K is (K) -pseudofinite.*

Proof. Each sentence φ in T_K is also in T_{K_n} for some n , and hence φ has a finite model. If T_{K_n} is K_n -pseudofinite, then this model is in K_n , and hence is also in K . \square

Corollaries 5.3.16 and 5.3.17 give us a strategy for showing that the generic theory T_K of a small Fraïssé class K is pseudofinite: filter K by $\{K_n \mid n \in \omega\}$ and show that each K_n admits a Fraïssé expansion with disjoint n -amalgamation for all n . We will use this strategy in Section 5.4 to establish pseudofiniteness of two countably categorical unsimple theories.

It is worth noting that this strategy 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 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}_Δ of triangle-free graphs, some K_n must include the graphs G_1 and G_2 but not G_3 . But Proposition 5.3.18 shows that such a class does not admit a Fraïssé expansion with disjoint n -amalgamation for all n .

Proposition 5.3.18. *Let K be a Fraïssé class consisting of graphs (in the language with a single edge relation R), 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_\emptyset be the unique quantifier-free 0-type in K' , the family of types $\{p_\emptyset, p(x), p(y), p(z), q(x, y), q(y, z), q(x, z)\}$ is a basic disjoint 3-amalgamation problem for $T_{K'}$. Then we must have $\neg xRy \in q(x, y)$, for otherwise the

reduct to L of any solution to the 3-amalgamation problem would be a copy of G_3 in K , contradicting our assumption.

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 . Since K' is a Fraïssé expansion of K , G_1 admits a compatible expansion to a structure in K' , call it G'_1 . Let $p'(y) = \text{qftp}_{G'_1}(v_3)$, and let $q_i(x, y) = \text{qftp}_{G'_1}(v_i, v_3)$ for $i = 1, 2$. Note that we have $q_i(x, y) \models p(x) \wedge p'(y)$ for $i = 1, 2$, but $\neg xRy \in q_1(x, y)$, while $xRy \in q_2(x, y)$. That is, the pair of quantifier-free 1-types $p(x)$ and $p'(y)$ are consistent with both xRy and $\neg xRy$. We will use this situation to build a triangle.

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 that $H' \models q_1(v_1, v_2)$, G_2 admits a compatible expansion to a structure G'_2 in K' . Let $p'' = \text{qftp}_{G'_2}(v_3)$, $r_1(x, z) = \text{qftp}_{G'_2}(v_1, v_3)$, and $r_2(y, z) = \text{qftp}_{G'_2}(v_2, v_3)$. Note that $r_1(x, z) \models p(x) \wedge p''(z) \wedge xRz$ and $r_2(y, z) \models p'(y) \wedge p''(z) \wedge yRz$.

Now the family of types $\{p_\emptyset, p(x), p'(y), p''(z), q_2(x, y), r_1(x, z), r_2(y, z)\}$ 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 , which is a contradiction. \square

5.4 Two generic theories of equivalence relations

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 .

K_{feq} is a Fraïssé class, and T_{feq}^* is its generic theory. T_{feq}^* was introduced and shown to be unsimple by Shelah [87]. In [90], Shelah and Usvyatsov attempted to show that it is NSOP₁, but their proof was incomplete due to some confusions around tree indiscernibles. The situation was clarified by Chernikov and Ramsey, who provided a correct proof of NSOP₁ in [25]. Shelah also claimed in an unpublished note [89] that T_{feq}^* is pseudofinite, but the proof given there was incorrect. To our knowledge, no proof that T_{feq}^* is pseudofinite has appeared in the literature until now.

The idea is simple, using the strategy described at the end of Section 5.3: 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. Then the resulting class has n -amalgamation for all n , and hence has pseudofinite generic theory.

Theorem 5.4.1. 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. We check that K_n is a Fraïssé class.

It is clearly a hereditary class. For the disjoint amalgamation property, suppose we have embeddings $f: A \rightarrow B$ and $g: A \rightarrow C$ 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 . For each parameter a in $P(D)$, we must specify an equivalence relation on $O(D)$. If a is in $P(A)$, it already defines equivalence relations on B and C . First, number the E_a -equivalence 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 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 $K_{\text{feq}} = \bigcup_{n=1}^{\infty} K_n$, and K_{feq} is a filtered Fraïssé class. By Corollary 5.3.17, it suffices to show that each T_{K_n} is pseudofinite.

Let L'_n be the 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 defined by one of the formulas $C_i(a, y)$.

We check that K'_n is a Fraïssé expansion of K_n . Certainly $K_n = \{A \upharpoonright L \mid A \in K'_n\}$, 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 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 K'_n is a Fraïssé class with the disjoint amalgamation property. We 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 5.3.14. \square

A natural question is whether T_{feq}^* is, in fact, the almost-sure theory for a uniform zero-one law on K_{feq} . It is not, as the following proposition shows. Of course, since we have described K_{feq} as 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 $f, g: \omega \rightarrow \omega$. 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 5.4.2. *There is a sentence φ in T_{feq}^* such that*

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

Proof. An example of such a sentence φ is

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

which expresses that any two equivalence classes for distinct equivalence relations intersect. φ is in T_{feq}^* , since for any A in K_{feq} with parameters $a \neq a'$ and objects b, b' (possibly $b = b'$), we can embed A in a structure B in K_{feq} with an object 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.

We will sketch the asymptotics: the measure $\mu_{f(n), g(n)}$ on $K_{\text{feq}}(f(n), g(n))$ 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 $\frac{n}{\log(n)}(1 + o(1))$ [36, Proposition VIII.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 non-trivially for all distinct a and b converges to 0. \square

Proposition 5.4.2 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 5.4.3. Does the class K_{feq} have a 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 ? If not, does it have a convergence law?

The theory T_{CPZ}

Let L be the language with a relation symbol $E_n(\bar{x}; \bar{y})$ of arity $2n$ for all $n \geq 1$. Then K_{CPZ} is the class of finite L -structures A such that for all n , E_n is an equivalence relation on A^n , and there is a single E_n -equivalence class consisting of all the redundant n -tuples.

K_{CPZ} is a Fraïssé class, and T_{CPZ} is its generic theory. In [20], Casanovas, Peláez, and Ziegler introduced the theory T_{CPZ} and showed that it is NSOP₂ and not simple. We will show how to combine the “independence lemma” from [20] with the 3-amalgamation criterion for NSOP₁ due to Chernikov and Ramsey [25] to show that T_{CPZ} is NSOP₁.

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

Theorem 5.4.4 ([25], Theorem 5.7). *T is NSOP₁ if and only if for every $M \models T$ and $b_0c_0 \equiv_M b_1c_1$ such that $c_1 \perp_M^u c_0$, $c_0 \perp_M^u b_0$, and $c_1 \perp_M^u b_1$, 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 simply take the independent 3-amalgamation condition in Theorem 5.4.4 as the definition of NSOP₁. For the original definition and further discussion, see [25] and [31].

Lemma 5.4.5 ([20], Lemma 4.2). *Let a, b, c, d', d'' be tuples and F a finite set from a model $M \models T_{\text{CPZ}}$. Assume that a and c have only elements of F in common, i.e. $a \perp_F^u c$. 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 5.4.6. *T_{CPZ} is NSOP₁.*

Proof. Suppose we are given $M \models T_{\text{CPZ}}$ and $d'a \equiv_M d''c$ such that $c \perp_M^u a$, $a \perp_M^u d'$, and $c \perp_M^u d''$. Let $p(x, y) = \text{tp}(d'a/M) = \text{tp}(d''c/M)$. To verify the condition in Theorem 5.4.4, 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 $q(x, y) = \text{tp}(d'a/F) = \text{tp}(d''c/F)$, $q(x, a) \cup q(x, c)$ is inconsistent. Since $c \perp_M^u a$, we certainly have $c \perp_F^u a$. By increasing F , we may assume that $c \perp_F^u a$. By countable categoricity, q is isolated by a single formula $\theta(x, y)$ over F , and $\theta(d', y) \in \text{tp}(a/Md')$, so by finite satisfiability there exists b in M satisfying $q(d', b)$. Since $d' \equiv_M d''$, we also have $q(d'', b)$.

Now the assumptions of Lemma 5.4.5 are satisfied, and we can find d satisfying $q(d, a)$ and $q(d, c)$, which contradicts inconsistency. \square

Now we turn to pseudofiniteness of T_{CPZ} . The strategy is the same as in Theorem 5.4.1: 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 5.4.7. *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.

K_n is clearly a hereditary class, 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 $f: A \rightarrow B$ and $g: A \rightarrow C$ 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 non-trivially, 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 5.4.1.

For any structure A in K_{CPZ} , if $|A| = N$, then the number of non-redundant n -tuples from A reaches its maximum of $N!$ when $n = N$. When $n > N$, there are no non-redundant n -tuples from A . So the number of E_n -equivalence classes is bounded above by $N! + 1$ for all n , and $A \in K_{N!+1}$. Hence $K_{\text{feq}} = \bigcup_{n=1}^{\infty} K_n$, and K_{CPZ} is a filtered Fraïssé class. By Corollary 5.3.17, it suffices to show that each T_{K_n} is pseudofinite.

Let L'_n be the language which includes, in addition to the relations E_k , $(n+1)$ k -ary relation symbols $C_k^0(\bar{x}), \dots, C_k^n(\bar{x})$ 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 defined by one of the $C_k^i(\bar{x})$, with the class of redundant tuples defined by $C_k^0(\bar{x})$.

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 any k -tuple involving the new element b is part of a class which exists in A , we label it by the appropriate C_k^i . If adding the new element adds new E_k -classes, 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 $\mathcal{P}^-([n])$ -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 is redundant. To solve our amalgamation problem, we simply assign relations from the C_k^i arbitrarily to those non-redundant k -tuples which are not already determined by the types in the family.

So $T_{K'_n}$ has disjoint n -amalgamation for all n , and hence it and its reduct T_{K_n} are pseudofinite by Theorem 5.3.14. \square

Proposition 5.4.8. *There is a sentence φ in T_{CPZ} such that*

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

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. 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 $b = b'$, 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 on $K_{\text{CPZ}}(n)$ 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 consider the equivalence relations on elements (of which there are n) and the non-redundant 2-tuples (of which there are $n^2 - n$). Using again the fact that the expected number of equivalence classes in a random equivalence relation grows asymptotically as $\frac{n}{\log(n)}(1 + o(1))$ [36, Proposition VIII.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 0. \square

The analog of Question 5.4.3 is interesting in this case too.

Question 5.4.9. Does K_{CPZ} have a uniform zero-one law? If not, does it have a uniform convergence law?

5.5 Primitive combinatorial theories

Upon first hearing the result of Corollary 5.3.15, one might naively ask whether the converse is true. Does every countably categorical pseudofinite theory have disjoint n -amalgamation for all n ? The answer is no, but we believe that this direction is nevertheless worth exploring.

The first class of counterexamples to the naive question comes from the fact that pseudofiniteness is preserved under reduct, but disjoint n -amalgamation is not. The homogeneous two-graph described in Example 5.3.9 is pseudofinite but fails to have disjoint 4-amalgamation. So we should instead ask whether every countably categorical pseudofinite theory admits an expansion with disjoint n -amalgamation for all n .

Based on the known examples, there is an intuition that pseudofiniteness for countably categorical theories always arises from “tame combinatorics”/randomness, or from “tame algebra”/stability; see [73, remarks on p.20] for an example of this sentiment. On the algebraic side, which we have neglected due to our interest in disjoint n -amalgamation in relational languages, the smoothly approximable structures studied in [24] and [54] provide many interesting examples. But these algebraic examples present another class of obvious counterexamples to the naive question: disjoint 2-amalgamation is equivalent to trivial acl, and there are many countably categorical pseudofinite theories with non-trivial acl.

So if we wish to connect pseudofiniteness to disjoint n -amalgamation, we will have to concentrate on the combinatorial paradigm and assume trivial acl. It turns out, however, that trivial acl alone is not a very strong hypothesis; as the following remark shows, any structure at all can be replaced by an essentially equivalent one with trivial acl. This is accomplished by hiding the algebraicity in a quotient by a definable equivalence relation.

Remark 5.5.1. Any (countably categorical) structure M in a relational language can be “blown up” to a (countably categorical) structure \widehat{M} with trivial acl which interprets M . Simply add a new equivalence relation E to the language and replace each element of M by an infinite E -class, so that $M \cong \widehat{M}/E$. Then given any finite set $A \subseteq \widehat{M}$ and any $b \notin A$, there is an automorphism of \widehat{M} fixing A and moving b , by permuting its equivalence class.

In fact, every example of a countably categorical theory with trivial acl which we currently know to be pseudofinite is either a reduct of a countably categorical theory with disjoint n -amalgamation for all n , or is built from equivalence relations (e.g. the blow-up of a pseudofinite countably categorical theory as in Remark 5.5.1, or the theories T_{feq}^* and T_{CPZ} in Section 5.4). In an attempt to avoid the hidden algebraicity in blow-ups and to rule out counterexamples like T_{feq}^* and T_{CPZ} , we will outlaw all interesting equivalence relations.

Definition 5.5.2. A **primitive combinatorial theory** is a countably categorical theory with trivial acl, such that for any finite set A and any complete n -type $p(\bar{x})$ over A , every A -definable equivalence relation on realizations of p is \emptyset -definable in the empty language.

Note that there are only two equivalence relations on singletons that are \emptyset -definable in the empty language: the equality relation $x = y$ and the trivial relation \top which holds of all pairs. The word primitive is often used in the context of homogeneous structures (see [73]) to mean that these are the only two \emptyset -definable equivalence relations on singletons. In particular, a primitive homogeneous structure has a unique 1-type over \emptyset . Our usage of the word primitive is different: we also rule out interesting equivalence relations on n -tuples which are definable with parameters, but we restrict our attention to the realizations of a complete type over the parameters, allowing equivalence relations across distinct types.

We now show that we only need to check primitivity on 1-types.

Proposition 5.5.3. *Let T be a complete theory with trivial acl, and suppose that for every (finite) set A and every complete 1-type $p(x)$ over A , every A -(invariant/definable) equivalence relation on the realizations of p is either equality or trivial. If q is a complete n -type over a (finite) set A and E is an A -(invariant/definable) equivalence relation on the realizations of q , then there is a subset $I \subseteq \{1, \dots, n\}$ of relevant coordinates and a group Σ of permutations of I such that $(\bar{x})E(\bar{y})$ if and only if $\bigvee_{\sigma \in \Sigma} \bigwedge_{i \in I} x_i = y_{\sigma(i)}$. In particular, E is \emptyset -definable in the empty language.*

Proof. By induction on n . The base case, $n = 1$, is true by assumption: for the equality relation, take $I = \{1\}$, and for the trivial relation, take $I = \emptyset$. In either case, the group Σ is trivial.

Suppose $n > 1$, let q be a complete n -type over A , and let E be an equivalence relation on the realizations of q which is A -(invariant/definable). For any $i \in \{1, \dots, n\}$, let q^i be the restriction of q to the variables $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n$ (omitting x_i), and for any i and any $\bar{a} \models q^i$, let $p_{\bar{a}}^i(x_i)$ be the 1-type $q(a_1, \dots, x_i, \dots, a_n)$ over $A\bar{a}$. If $b \models p_{\bar{a}}^i$, we write $\bar{a}b$ for the tuple $a_1, \dots, b, \dots, a_n$. Consider the equivalence relation $E_{\bar{a}}^i$ on the realizations of $p_{\bar{a}}^i$, given by $(b)E_{\bar{a}}^i(b')$ if and only if $(\bar{a}b)E(\bar{a}b')$. $E_{\bar{a}}^i$ is $A\bar{a}$ -invariant, and if E is A -definable, then $E_{\bar{a}}^i$ is $A\bar{a}$ -definable. By assumption, for each $\bar{a} \models q^i$, $E_{\bar{a}}^i$ is equality or trivial, and by A -invariance of E , which case we are in does not depend on the choice of \bar{a} satisfying q^i .

Case 1: For some i and all $\bar{a} \models q^i$, $E_{\bar{a}}^i$ is trivial. This includes the case when $p_{\bar{a}}^i$ is an algebraic type; by trivial acl, $p_{\bar{a}}^i$ is algebraic if and only if it contains the formula $x_i = c$ for

$c \in A$ or $x_i = a_j$ for some $j \neq i$, hence it has exactly one realization, and the trivial relation is the only equivalence relation on realizations of $p_{\bar{a}}^i$.

Consider the relation E^i on the realizations of q^i , defined by $(\bar{a})E^i(\bar{a}')$ if and only if for all $b \models p_{\bar{a}}^i$ and $b' \models p_{\bar{a}'}^i$, we have $(\bar{a}b)E(\bar{a}'b')$. We claim that also $(\bar{a})E^i(\bar{a}')$ if and only if there exist $b \models p_{\bar{a}}^i$ and $b' \models p_{\bar{a}'}^i$, such that $(\bar{a}b)E(\bar{a}'b')$. Indeed, if there exist such b and b' , then for any other $d \models p_{\bar{a}}^i$ and $d' \models p_{\bar{a}'}^i$, we have $(b)E_{\bar{a}}^i(d)$ and $(b')E_{\bar{a}'}^i(d')$ by triviality of $E_{\bar{a}}^i$ and $E_{\bar{a}'}^i$, so $(\bar{a}d)E(\bar{a}b)E(\bar{a}'b')E(\bar{a}'d')$.

Now E^i is an A -invariant equivalence relation, and we claim that if E is A -definable, then so is E^i . We have seen that $(\bar{a})E^i(\bar{a}')$ if and only if there exist realizations $b \models p_{\bar{a}}^i$ and $b' \models p_{\bar{a}'}^i$ such that $(\bar{a}b)E(\bar{a}'b')$. Since type-definable sets are closed under existential quantification, E^i is type-definable over A . But also, $\neg(\bar{a})E^i(\bar{a}')$ if and only if there exist realizations $b \models p_{\bar{a}}^i$ and $b' \models p_{\bar{a}'}^i$ such that $\neg(\bar{a}b)E(\bar{a}'b')$. Hence the complement of E^i is also type-definable over A , and, by compactness, E^i is A -definable.

By induction, E^i is definable by a formula of the specified form. And we have shown that for any $\bar{a}b$ and $\bar{a}'b'$ satisfying q , $(\bar{a}b)E(\bar{a}'b')$ if and only if $(\bar{a})E^i(\bar{a}')$. Thus E is definable by the same formula as E^i , up to variable substitution, with i inserted as an additional irrelevant coordinate.

Case 2: For all i and all $\bar{a} \models q^i$, $p_{\bar{a}}^i$ is non-algebraic, and $E_{\bar{a}}^i$ is the equality relation. We claim that if $(\bar{d})E(\bar{d}')$, then the tuple \bar{d}' is a permutation of the tuple \bar{d} . Since we are in Case 2, the tuples \bar{d} and \bar{d}' are non-redundant, so it suffices to show that every element of \bar{d} is an element of \bar{d}' .

Write \bar{d} as $\bar{a}b$, where b is the i^{th} coordinate of \bar{d} . Then b is the unique realization of its type over $A\bar{a}\bar{d}'$. Indeed, if $b \equiv_{A\bar{a}\bar{d}'} b'$, then since $(\bar{a}b)E(\bar{d}')$ and E is A -invariant, we also have $(\bar{a}b')E(\bar{d}')$. But then $(\bar{a}b)E(\bar{a}b')$, so $(b)E_{\bar{a}}^i(b')$, and hence $b = b'$. By trivial acl, b must be an element of $A\bar{a}\bar{d}'$. Since we are in Case 2, q does not contain the formulas $x_i = c$ for $c \in A$ or $x_i = x_j$ for $j \neq i$, so b is not an element of A or the tuple \bar{a} . Thus it appears in \bar{d}' .

Let $\Sigma = \{\sigma \in S_n \mid \sigma(\bar{d}) \models q \text{ and } (\bar{d})E(\sigma(\bar{d}))\}$. By A -invariance of E , Σ does not depend on the choice of \bar{d} . Hence $\bar{x}E\bar{y}$ if and only if $\bigvee_{\sigma \in \Sigma} \bigwedge_{i=1}^n x_i = y_{\sigma(i)}$. Reflexivity, symmetry, and transitivity of E imply that Σ contains the identity, is closed under inverse, and is closed under composition, respectively, so Σ is a subgroup of the symmetric group S_n . \square

Remark 5.5.4. Since the theory $T_{=}$ of an infinite set in the empty language satisfies the hypotheses of Proposition 5.5.3, we have also classified the definable equivalence relations on realizations of complete types in this theory.

Anticipating that it could be useful in more general contexts (and it has been already [28]), we have stated Proposition 5.5.3 in a flexible way, not assuming countable categoricity and allowing the reader to choose between infinite or finite sets of parameters and invariant or definable relations. But in our setting, it immediately implies the reduction of primitivity to 1-types.

Corollary 5.5.5. *If T is countably categorical with trivial acl, then T is a primitive combinatorial theory if and only if for every finite set A and every complete 1-type p over A , every A -definable equivalence relation on realizations of p is either equality or trivial.*

Example 5.5.6. Let T_Δ be the theory of the generic triangle-free graph. We have seen that T_Δ is countably categorical with trivial acl and quantifier elimination (since G_Δ is a small Fraïssé class with disjoint 2-amalgamation). We claim that it is a primitive combinatorial theory.

Let A be a finite subset of a model of T_Δ , and let $p(x)$ be a complete 1-type over A . We may assume that $p(x)$ is non-algebraic, otherwise it has at most one realization. Then, by quantifier elimination, there are at most two 2-types $q(x, y)$ over A consistent with $p(x) \cup p(y) \cup \{x \neq y\}$, determined by xRy and $\neg xRy$, respectively. xEy must be defined by one of these relations.

If $p(x)$ contains the formula xRc for any $c \in A$, then xRy is inconsistent with $p(x) \cup p(y) \cup \{x \neq y\}$, so E is equality (if it is defined by xRy) or trivial (if it is defined by $\neg xRy$). On the other hand, if $p(x)$ contains $\neg xRc$ for all $c \in A$, then the formulas $xRy \wedge yRz \wedge \neg xRz$ and $\neg xRy \wedge \neg yRz \wedge xRz$ are both consistent with $p(x) \cup p(y) \cup p(z) \cup \{x \neq y, y \neq z, x \neq z\}$, so neither relation is transitive on realizations of p .

The following lemma, which generalizes of a theorem of Macpherson [72, Proposition 1.3], shows that primitive combinatorial theories exhibit no non-trivial stable behavior. This suggests that we may view primitive combinatorial theories as “purely combinatorial purely unstable” theories, and supports our motivation for defining this class: to isolate the combinatorial paradigm for pseudofiniteness. We also believe that the study of primitive combinatorial theories could be interesting independently of investigations of pseudofiniteness, as a testing ground for combinatorial behaviors in unstable model theory.

Lemma 5.5.7. *Let T be a primitive combinatorial theory, and let $p(x)$ be a complete 1-type over \emptyset . If $\varphi(x, \bar{y})$ is a stable formula, then any instance $\varphi(x, \bar{b})$ of φ defines a finite or cofinite subset of the realizations of p .*

Proof. Let M be the unique countable model of T , and let θ_p be a formula isolating p . Suppose for contradiction that $\varphi(x, \bar{b})$ defines an infinite and coinfinite subset of $\theta_p(M)$.

An **instance** of φ is a formula of the form $\varphi(x, \bar{b}')$ for some tuple \bar{b}' from M . A φ -**formula** is a Boolean combination of instances of φ . A φ -formula $\psi(x, \bar{c})$ is φ -**minimal** in p if $\theta_p(x) \wedge \psi(x, \bar{c})$ is non-algebraic and, for any other φ -formula $\psi'(x, \bar{c}')$, one of $\theta_p(x) \wedge \psi(x, \bar{c}) \wedge \psi'(x, \bar{c}')$ or $\theta_p(x) \wedge \psi(x, \bar{c}) \wedge \neg \psi'(x, \bar{c}')$ is algebraic. Now there is a φ -formula $\psi(x, \bar{c})$ such that $\psi(x, \bar{c}) \rightarrow \varphi(x, \bar{b})$ and $\psi(x, \bar{c})$ is φ -minimal in p . If not, we could build a complete binary tree of φ -formulas, contradicting stability of $\varphi(x, \bar{y})$.

Let $q(\bar{z}) = \text{tp}(\bar{c}/\emptyset)$. For any conjugate \bar{c}' of \bar{c} (i.e. for any realization of q), we let $X_{\bar{c}'} = \theta_p(M) \cap \psi(M, \bar{c}')$. For any other conjugate \bar{c}'' , since $\psi(x, \bar{c}')$ is also φ -minimal, either $X_{\bar{c}'} \cap X_{\bar{c}''}$ or $X_{\bar{c}'} \setminus X_{\bar{c}''}$ is finite. By trivial acl, this finite set (whichever it is) is contained in the set of parameters $\|\bar{c}'\bar{c}''\|$.

The idea is that the sets $X_{\bar{c}}$ partition $\theta_p(M)$ into the classes of a definable equivalence relation, up to the finite noise coming from the parameters.

Let $\theta_q(\bar{z})$ be a formula isolating q . Let E be the relation on $\theta_p(M)$ defined by aEa' if and only if there is some \bar{c}' realizing q such that $a, a' \in X_{\bar{c}'}$ and $a, a' \notin \|\bar{c}'\|$:

$$xEa' \leftrightarrow \exists \bar{z} \left(\theta_q(\bar{z}) \wedge \left(\bigwedge_{i=1}^k (x \neq z_i) \wedge (x' \neq z_i) \right) \wedge \psi(x, \bar{z}) \wedge \psi(x', \bar{z}) \right).$$

We claim that E is an equivalence relation on $\theta_p(M)$ which is not equality or trivial.

Reflexivity: Let $a \in \theta_p(M)$. Since $X_{\bar{c}}$ is infinite, there is some $a' \in X_{\bar{c}}$ such that $a' \notin \|\bar{c}\|$. Then a is a conjugate of a' , so there is some conjugate \bar{c}' of \bar{c} such that $a \in X_{\bar{c}'}$ and $a \notin \|\bar{c}'\|$. This \bar{c}' witnesses aEa .

Symmetry: Clear from the definition.

Transitivity: Let a, a' , and a'' be distinct elements of $\theta_p(M)$, and suppose that \bar{c}' witnesses aEa' and \bar{c}'' witnesses $a'Ea''$. We may assume that $a \notin \|\bar{c}''\|$. Indeed, $\text{tp}(a/a'a'')$ is not algebraic, so it has some realization a^* which is not in $\|\bar{c}''\|$. Then a^* is a conjugate of a over $a'a''$, so \bar{c}'' has a conjugate over $a'a''$ which also witnesses $a'Ea''$ and does not contain a .

Now $a' \in X_{\bar{c}'} \cap X_{\bar{c}''}$ and $a' \notin \|\bar{c}'\bar{c}''\|$, so, by trivial acl, $X_{\bar{c}'} \cap X_{\bar{c}''}$ is infinite. Hence $X_{\bar{c}'} \setminus X_{\bar{c}''}$ is finite. But $a \notin \|\bar{c}'\bar{c}''\|$ and $a \in X_{\bar{c}'}$, so, by trivial acl again, we must have $a \in X_{\bar{c}''}$. Then \bar{c}'' witnesses aEa'' .

Not equality: There are infinitely many elements in $X_{\bar{c}}$ which are not in $\|\bar{c}\|$, so \bar{c} witnesses that they are all equivalent.

Not trivial: Let $a \in X_{\bar{c}}$ with $a \notin \|\bar{c}\|$, and let $a' \in \theta_p(M) \setminus X_{\bar{c}}$ with $a' \notin \|\bar{c}\|$. The latter is possible, since $\varphi(M, \bar{b})$ has infinite complement in $\theta_p(M)$. Suppose for contradiction that there is some \bar{c}' witnessing aEa' . Then $a \in X_{\bar{c}'} \cap X_{\bar{c}}$ and $a' \in X_{\bar{c}'} \setminus X_{\bar{c}}$ and neither a nor a' are in $\|\bar{c}'\|$, so both sets are infinite, contradiction. \square

Let L_n^P be the language $\{P_1, \dots, P_n\}$, where each P_i is a unary relation, and let T_n^P be the theory asserting that the P_i pick out infinite disjoint sets and every element is in some P_i . Then T_n^P is a stable primitive combinatorial theory.

Theorem 5.5.8. *Any stable primitive combinatorial theory is interdefinable with T_n^P for some n .*

Proof. Let T be a primitive combinatorial theory in the language L , and suppose that there are n complete 1-types over \emptyset , isolated by the formulas $\theta_1(x), \dots, \theta_n(x)$. These formulas pick out infinite sets by trivial acl. Then T interprets T_n^P with domain defined by $\top(x)$ and $P_i(x)$ defined by $\theta_i(x)$.

Conversely, we claim that every L -formula $\varphi(\bar{x})$ is equivalent modulo T to a Boolean combination of the θ_i and instances of equality. From this it follows immediately that T_n^P interprets T with domain defined by $\top(x)$. We prove this by induction on the length l of the tuple \bar{x} . When $l = 0$, every sentence is equivalent modulo T to \top or \perp (since T is complete). When $l = 1$, every formula $\varphi(x)$ is equivalent to a disjunction of some of the $\theta_i(x)$.

Suppose we have established the claim at length l . Consider a formula $\varphi(x, \bar{y})$, where \bar{y} has length l . Let $\psi_1(\bar{y}), \dots, \psi_m(\bar{y})$ be formulas isolating the finitely many l -types consistent with T . Fixing j with $1 \leq j \leq m$, let \bar{b} be a tuple satisfying $\psi_j(\bar{y})$. For $1 \leq i \leq n$, by Lemma 5.5.7, $\varphi(x, \bar{b})$ defines a finite or cofinite subset of the realizations of $\theta_i(x)$. By trivial acl, this set or its complement (whichever is finite) is contained in $\|\bar{b}\|$, so it is defined by $\left(\bigvee_{k \in S_{ij}} x = b_k\right)$ or $\neg \left(\bigvee_{k \in S_{ij}} x = b_k\right)$, for some subset S_{ij} of $[l]$. This definition does not depend on the choice of \bar{b} satisfying $\psi_j(\bar{y})$.

Then $\psi(x, \bar{y})$ is equivalent to $\bigwedge_{i=1}^n \bigwedge_{j=1}^m \left((\theta_i(x) \wedge \psi_j(\bar{y})) \rightarrow (\neg)_{ij} \left(\bigvee_{k \in S_{ij}} x = b_k \right) \right)$, where $(\neg)_{ij}$ indicates that the negation may or may not appear, depending on the values of i and j . We are done since, by induction, the formulas $\psi_j(\bar{y})$ are equivalent to Boolean combinations of the θ_i and instances of equality. \square

Remark 5.5.9. Note that for Lemma 5.5.7 and Theorem 5.5.8, we only used the hypothesis that there are no interesting A -definable equivalence relations on complete 1-types over A in the case $A = \emptyset$.

We now turn from stability to simplicity. Recall that tuples a and b have the same **strong type** over A , written $\text{stp}(a/A) = \text{stp}(b/A)$, if aEb for every bounded (equivalently finite) A -definable equivalence relation. And tuples a and b have the same **Lascar strong type** over A , written $\text{Lstp}(a/A) = \text{Lstp}(b/A)$, if aEb for every bounded A -invariant equivalence relation. See [19, Chapter 9] for more background on these notions. Since we are working with countably categorical theories, invariant equivalence relations over finite sets are definable, and Lascar strong types over finite sets are just strong types. And since we have ruled out interesting definable equivalence relations, strong types over finite sets in primitive combinatorial theories are just types.

Proposition 5.5.10. *Let T be a primitive combinatorial theory, let A be any finite set, and let a and b be tuples. Then $\text{Lstp}(a/A) = \text{Lstp}(b/A)$ if and only if $\text{tp}(a/A) = \text{tp}(b/A)$.*

Proof. In any theory, if $\text{Lstp}(a/A) = \text{Lstp}(b/A)$, then $\text{tp}(a/A) = \text{tp}(b/A)$ (the relation of having the same type over A is a bounded A -invariant equivalence relation). So suppose $\text{tp}(a/A) = \text{tp}(b/A)$, and call this type p . We may assume that p is not algebraic, since otherwise it has a unique realization. Let E be a bounded A -invariant equivalence relation; we might as well restrict E to the realizations of p . In a countably categorical theory, a relation which is invariant over a finite set A is definable over A , so E is A -definable, and hence E is \emptyset -definable in the empty language.

But given the description in Proposition 5.5.3 of all equivalence relations which are \emptyset -definable in the empty language, we see that the only one which is bounded is the trivial relation. Indeed, since p is not algebraic, we can find unboundedly many realizations $\langle c_\alpha \rangle_{\alpha < \kappa}$ of p which are disjoint over A . If $\alpha \neq \beta$, the only coordinates of c_α and c_β which agree are those which are in A , and these coordinates agree for all pairs of realizations of p . Hence the c_α are only related by the trivial relation on realizations of p , so E is trivial, and aEb . \square

In Theorem 5.3.10 above, we stated the independence theorem over models. Simple theories also satisfy a form of independent 3-amalgamation over arbitrary sets, given an additional condition involving Lascar strong types.

Theorem 5.5.11 ([19, Corollary 10.10]). *Let T be a simple theory. Suppose we have sets A , B , and C such that $B \downarrow_A^f C$ and tuples a and a' such that $a \downarrow_A^f B$, $a' \downarrow_A^f C$, and $\text{Lstp}(a/A) = \text{Lstp}(a'/A)$. Then there exists a'' such that $a'' \downarrow_A^f BC$, $\text{tp}(a''/AB) = \text{tp}(a/AB)$, and $\text{tp}(a''/AC) = \text{tp}(a'/AC)$.*

Theorem 5.5.12. *For a primitive combinatorial theory T , the following are equivalent:*

- (1) T has disjoint 3-amalgamation.
- (2) T is simple with trivial forking.

Proof. One direction is Theorem 5.3.11, since a primitive combinatorial theory also has disjoint 2-amalgamation. For the converse, suppose we have a coherent $\mathcal{P}^-([3])$ -family of types over A . Let bc realize $p_{\{1,2\}}(x_1, x_2)$. We would like to show that $p_{\{0,1\}}(x_0, b) \cup p_{\{0,2\}}(x_0, c)$ is consistent. By compactness, we may assume that A is finite, b and c are finite tuples, and x_0 is a finite tuple of variables. Let a realize $p_{\{0,1\}}(x_0, b)$, and let a' realize $p_{\{0,2\}}(x_0, c)$, so both a and a' realize $p_{\{0\}}(x_0)$. The types $p_{\{i,j\}}$ are non-redundant, so, by trivial forking, $b \downarrow_A^f c$, $a \downarrow_A^f b$, and $a' \downarrow_A^f c$. By Proposition 5.5.10, $\text{Lstp}(a/A) = \text{Lstp}(a'/A)$. So by Theorem 5.5.11, there is some a'' such that $\text{tp}(a''/Ab) = \text{tp}(a/Ab)$ and $\text{tp}(a''/Ac) = \text{tp}(a'/Ac)$, i.e. a'' realizes $p_{\{0,1\}}(x_0, b) \cup p_{\{0,2\}}(x_0, c)$. \square

We do not know of any simple primitive combinatorial theories which fail to satisfy the equivalent conditions of Theorem 5.5.12, and we suspect that there are none.

Conjecture 5.5.13. Every simple primitive combinatorial theory has trivial forking and disjoint 3-amalgamation.

Conjecture 5.5.13 is related to the stable forking conjecture, which states roughly that every instance of forking in a simple theory is witnessed by a stable formula [62]. Since primitive combinatorial theories have no non-trivial stable behavior, simple primitive combinatorial theories should have trivial forking.

By Theorem 5.5.12, if a primitive combinatorial theory is not simple, then it exhibits a failure of disjoint 3-amalgamation. The basic examples of failures of disjoint 3-amalgamation in small Fraïssé classes are transitivity of equivalence relations (which we've ruled out), graphs and directed graphs omitting triangles (the generic theories of these classes have SOP_3), and transitivity in partial orders (the generic theories of these classes have the strict order property, which is even stronger than SOP_3). Evans and Wong have shown [34] that a certain class of Hrushovski constructions cannot produce generic theories which are un-simple but have NSOP_3 . The pattern exhibited by these examples motivated the following conjecture.

Conjecture 5.5.14. If a primitive combinatorial theory is not simple, then it has SOP_3 .

The NIP theories are a class of theories which contain the stable theories but also include theories with the strict order property (in fact, a theory is stable if and only if it has NIP and NSOP [86]). Distal theories were introduced by Simon [91] with the goal of characterizing the “purely unstable” NIP theories (e.g. DLO, the other o-minimal theories, and $\text{Th}(\mathbb{Q}_p)$). Following our intuition that primitive combinatorial theories have no non-trivial stable behavior, we might expect that unstable NIP primitive combinatorial theories should be distal. Of course, if we put a copy of DLO next to a pure set, the result will be unstable, but not distal. We need to ensure that no part of the structure is stable.

Definition 5.5.15. A primitive combinatorial theory is **purely unstable** if for every complete 1-type $p(x)$ over \emptyset , the induced structure on the realizations of p (which is also a primitive combinatorial theory) is unstable.

By Theorem 5.5.8, this is equivalent to asking that the induced structure on p is not interdefinable with a pure set.

Conjecture 5.5.16. Every NIP purely unstable primitive combinatorial theory is distal.

Conjecture 5.5.17. Returning now to our original motivation, we collect here our conjectures on pseudofinite countably categorical theories

- (1) Every pseudofinite primitive combinatorial theory is a reduct of a primitive combinatorial theory with disjoint n -amalgamation for all n .
- (2) Every pseudofinite primitive combinatorial theory is simple.
- (3) If the theory T_K of a small Fraïssé class K with the disjoint amalgamation property is pseudofinite, then K can be filtered by $\{K_n \mid n \in \omega\}$, where each K_n has a Fraïssé expansion with disjoint n -amalgamation for all n .
- (4) Every pseudofinite countably categorical theory has NSOP_1 .

(1) is the idea that motivated our definition of primitive combinatorial theories. Again, it says that in “purely combinatorial purely unstable” settings, pseudofiniteness should always be explained by probabilistic argument given in Theorem 5.3.14. Perhaps a more convincing motivation for the conjecture is the lack of known counterexamples. (2) is significantly weaker, asking only for independent 3-amalgamation (equivalently, disjoint 3-amalgamation, if Conjecture 5.5.13 is true), instead of disjoint n -amalgamation for all n .

(3) and (4) are parallel to (1) and (2), but dropping the primitivity assumption. Essentially, they say that the worst that can happen outside of the realm of primitive combinatorial theories is that we get behavior like that of T_{feq}^* and T_{CPZ} . Early in the development of simple theories Kim and Pillay [61] made the “rather outrageous conjecture” that every pseudofinite countably categorical theory is simple. Of course, T_{feq}^* and T_{CPZ} are counterexamples, but

they still have NSOP_1 . The “new outrageous conjecture” (4), which asks for NSOP_1 instead of simplicity, was suggested to us by Ramsey.

Remark 5.5.18. Any of the parts of Conjecture 5.5.17 would prove that the theory T_Δ of the generic triangle-free graph is *not* pseudofinite. Indeed, T_Δ is a primitive combinatorial theory (Example 5.5.6), but it is not simple, so (2) implies it is not pseudofinite, as does (1) (using Theorem 5.3.11). T_Δ is countably categorical, but it has SOP_1 (even SOP_3), so (4) suffices as well. And T_Δ admits no filtration by Fraïssé classes as in (3), by Proposition 5.3.18, so we can draw the same conclusion from (3).

It is fitting to end this thesis with a question that ties together the three types of infinitary limits that we have studied in Parts I, II, and III. If a small Fraïssé class K has disjoint n -amalgamation for all n , then (1) the sequence of measures $\langle \mu_n \rangle_{n \in \omega}$ on the spaces $\text{Str}_{L,n}$ produced in the proof of Theorem 5.3.14 cohere to an ergodic structure μ , (2) μ is almost surely isomorphic to the Fraïssé limit M_K , and (3) K has a zero-one law with respect to the μ_n , with almost-sure theory T_K . We noted this in Remark 5.3.13, and we now give a proof.

Theorem 5.5.19. *Let K be a small Fraïssé class with disjoint n -amalgamation for all n . Let $\langle \mu_n \rangle_{n \in \omega}$ be the sequence of measures on the spaces $\text{Str}_{L,n}$ constructed in the proof of Theorem 5.3.14. Then the μ_n cohere (in the sense of Definition 5.1.6) to an ergodic structure μ which is almost surely isomorphic to M_K .*

Proof. It is easy to check the conditions for coherence given in Remark 5.1.7. Agreement and invariance follow from the fact that, in the definition of μ_n , the random procedure for determining the quantifier-free type of a non-redundant tuple \bar{a} does not depend on n or \bar{a} (just the length of \bar{a}). Disjoint-independence holds because all decisions about quantifier-free types of disjoint tuples are made independently.

To show that μ is almost surely isomorphic to M_K , we just need to check that $\mu(\llbracket \psi \rrbracket) = 1$ for every axiom ψ of T_K , since T_K is countably categorical. This is very similar to the proof that $\lim_{n \rightarrow \infty} \mu_n(\llbracket \psi \rrbracket) = 1$ for all $\psi \in T_K$ given in Theorem 5.3.14.

If ψ is a universal axiom of T_K , then ψ has the form $\forall \bar{x} \neg \varphi_A(\bar{x})$, for some A in $\text{Str}_{L,n}$ but not in K . Since a countable intersection of measure 1 sets has measure 1, it suffices to show that $\mu(\llbracket \neg \varphi_A(\bar{a}) \rrbracket) = 1$ for all tuples \bar{a} from ω . Choosing N greater than the largest element of \bar{a} , $\mu(\llbracket \neg \varphi_A(\bar{a}) \rrbracket) = \mu_N(\llbracket \neg \varphi_A(\bar{a}) \rrbracket) = 1$. This is clear from the definition of μ_N and was demonstrated in the proof of Theorem 5.3.14.

If ψ is a one-point extension axiom of T_K , then ψ has the form $\forall \bar{x} \exists y \varphi_A(\bar{x}) \rightarrow \varphi_B(\bar{x}, y)$. Again, it suffices to show that $\mu(\llbracket \exists y \varphi_A(\bar{a}) \rightarrow \varphi_B(\bar{a}, y) \rrbracket) = 1$ for all tuples \bar{a} from ω . The formula is vacuously satisfied by any redundant tuple, so fix a non-redundant tuple \bar{a} . Let b and b' be elements not in $\llbracket \bar{a} \rrbracket$, and let N be larger than any of these natural numbers. We saw in the proof of Theorem 5.3.14 that according to μ_N , if we condition on the event $\llbracket \varphi_A(\bar{a}) \rrbracket$, then the events $\llbracket \varphi_B(\bar{a}, b) \rrbracket$ and $\llbracket \varphi_B(\bar{a}, b') \rrbracket$ are conditionally independent and have the same positive probability, which does not depend on N . Moving from $\text{Str}_{L,N}$ to Str_L , our measure μ agrees with μ_N , so the events $\{\llbracket \varphi_B(\bar{a}, b) \rrbracket \mid b \in \omega \setminus \llbracket \bar{a} \rrbracket\}$ each have positive

probability when conditioned on $\llbracket \varphi_A(\bar{a}) \rrbracket$ and are conditionally independent. Almost surely (conditioned on $\llbracket \varphi_A(\bar{a}) \rrbracket$) infinitely many of them occur, and (removing the conditioning) $\mu(\llbracket \exists y \varphi_A(\bar{a}) \rightarrow \varphi_B(\bar{a}, y) \rrbracket) = 1$. \square

This is a perfect situation: all three kinds of limits exist and agree.

Definition 5.5.20. Let K be a small Fraïssé class, and let μ be an ergodic structure. The pair (K, μ) is **perfect** if μ is almost surely isomorphic to the Fraïssé limit M_K of K , and μ induces a coherent sequence μ_n of measures on $\text{Str}_{L,n}$ such that K has a zero-one law with respect to the μ_n and the almost-sure theory $T^{\text{a.s.}}(K)$ is equal to the generic theory T_K .

Question 5.5.21. Which small Fraïssé classes K are part of a perfect pair (K, μ) ? Which ergodic structures μ pair perfectly with a given small Fraïssé class K ? What is the relationship with disjoint n -amalgamation?

If our readers are disappointed that we have not solved this problem, we hope that the anticlimax will motivate them to work on it (as well as the many conjectures in this section). In our view, the mark of a good story is that it raises more questions than it answers.

Bibliography

- [1] Nathanael Ackerman. “Representations of $\text{Aut}(M)$ -invariant measures: part 1”. Sept. 2015. [arXiv:1509.06170\[math.LO\]](https://arxiv.org/abs/1509.06170). URL: <http://arxiv.org/abs/1509.06170>.
- [2] Nathanael Ackerman, Cameron Freer, and Rehana Patel. “Invariant measures concentrated on countable structures”. June 2012. [arXiv:1206.4011\[math.LO\]](https://arxiv.org/abs/1206.4011). URL: <http://arxiv.org/abs/1206.4011>.
- [3] Nathanael Ackerman, Cameron Freer, and Rehana Patel. “A classification of sentences admitting an invariant measure”. In preparation. 2016.
- [4] Nathanael Ackerman, Cameron Freer, Aleksandra Kwiatkowska, and Rehana Patel. “A classification of orbits admitting a unique invariant measure”. Dec. 2014. [arXiv:1412.2735\[math.LO\]](https://arxiv.org/abs/1412.2735). URL: <http://arxiv.org/abs/1412.2735>.
- [5] Nathanael Ackerman, Cameron Freer, Jaroslav Nešetřil, and Rehana Patel. “Invariant measures via inverse limits of finite structures”. *European J. Combin.* 52.B (2016), pp. 248–289. URL: [http://arxiv.org/abs/1310.8147](https://arxiv.org/abs/1310.8147).
- [6] Nathanael Ackerman, Cameron Freer, Alex Kruckman, and Rehana Patel. “Properly ergodic structures”. In preparation. 2016.
- [7] Ove Ahlman. “Simple structures axiomatized by almost sure theories”. Aug. 2015. [arXiv:1508.06039\[math.LO\]](https://arxiv.org/abs/1508.06039). URL: <http://arxiv.org/abs/1508.06039>.
- [8] David J. Aldous. “Representations for partially exchangeable arrays of random variables”. *J. Multivariate Anal.* 11.4 (1981), pp. 581–598. URL: [http://dx.doi.org/10.1016/0047-259X\(81\)90099-3](http://dx.doi.org/10.1016/0047-259X(81)90099-3).
- [9] Tim Austin. “On exchangeable random variables and the statistics of large graphs and hypergraphs”. *Probab. Surv.* 5 (2008), pp. 80–145. URL: <http://arxiv.org/abs/0801.1698>.
- [10] John T. Baldwin. *Categoricity*. Vol. 50. University Lecture Series. American Mathematical Society, Providence, RI, 2009.
- [11] John T. Baldwin, Martin Koerwien, and Michael C. Laskowski. “Disjoint amalgamation in locally finite AEC”. Preprint. June 2015. URL: <http://homepages.math.uic.edu/~jrbaldwin/pub/cut618.pdf>.

- [12] John T. Baldwin and Niandong Shi. “Stable generic structures”. *Ann. Pure Appl. Logic* 79.1 (1996), pp. 1–35. URL: [http://dx.doi.org/10.1016/0168-0072\(95\)00027-5](http://dx.doi.org/10.1016/0168-0072(95)00027-5).
- [13] Paul Balister, Béla Bollobás, Karen Gunderson, Imre Leader, and Mark Walters. “Random geometric graphs and isometries of normed spaces”. Apr. 2015. [arXiv:1504.05324](https://arxiv.org/abs/1504.05324)[math.FA]. URL: <http://arxiv.org/abs/1504.05324>.
- [14] Béla Bollobás and Andrew Thomason. “Graphs which contain all small graphs”. *European J. Combin.* 2.1 (1981), pp. 13–15. URL: [http://dx.doi.org/10.1016/S0195-6698\(81\)80015-7](http://dx.doi.org/10.1016/S0195-6698(81)80015-7).
- [15] Anthony Bonato and Jeannette Janssen. “Infinite random geometric graphs”. *Ann. Comb.* 15.4 (2011), pp. 597–617. URL: <http://arxiv.org/abs/0908.2590>.
- [16] Andrew Brooke-Taylor and Damiano Testa. “The infinite random simplicial complex”. Aug. 2013. [arXiv:1308.5517](https://arxiv.org/abs/1308.5517)[math.LO]. URL: <http://arxiv.org/abs/1308.5517>.
- [17] Peter J. Cameron. *Oligomorphic permutation groups*. Vol. 152. London Mathematical Society Lecture Note Series. Cambridge University Press, Cambridge, 1990.
- [18] Olivia Caramello. “Fraïssé’s construction from a topos-theoretic perspective”. *Log. Univers.* 8.2 (2014), pp. 261–281. URL: <http://arxiv.org/abs/0805.2778>.
- [19] Enrique Casanovas. *Simple Theories and Hyperimaginaries*. Vol. 39. Lecture Notes in Logic. Cambridge University Press, Cambridge, 2011.
- [20] Enrique Casanovas, Rodrigo Peláez, and Martin Ziegler. “On many-sorted ω -categorical theories”. *Fund. Math.* 214.3 (2011), pp. 285–294. URL: <http://arxiv.org/abs/1012.4370>.
- [21] Zoé Chatzidakis and Ehud Hrushovski. “Model theory of difference fields”. *Trans. Amer. Math. Soc.* 351.8 (1999), pp. 2997–3071. URL: <http://dx.doi.org/10.1090/S0002-9947-99-02498-8>.
- [22] Gregory Cherlin. “Combinatorial problems connected with finite homogeneity”. In: *Proceedings of the International Conference on Algebra, Part 3 (Novosibirsk, 1989)*. Vol. 131. Contemp. Math. Amer. Math. Soc., Providence, RI, 1992, pp. 3–30.
- [23] Gregory Cherlin. “Two problems on homogeneous structures, revisited”. In: *Model theoretic methods in finite combinatorics*. Vol. 558. Contemp. Math. Amer. Math. Soc., Providence, RI, 2011, pp. 319–415. URL: <http://www.rci.rutgers.edu/~cherlin/Publication/census.pdf>.
- [24] Gregory Cherlin and Ehud Hrushovski. *Finite structures with few types*. Vol. 152. Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 2003.
- [25] Artem Chernikov and Nicholas Ramsey. “On model-theoretic tree properties”. May 2015. [arXiv:1505.00454](https://arxiv.org/abs/1505.00454)[math.LO]. URL: <http://arxiv.org/abs/1505.00454>.
- [26] F. R. K. Chung, R. L. Graham, and R. M. Wilson. “Quasi-random graphs”. *Combinatorica* 9.4 (1989), pp. 345–362. URL: <http://dx.doi.org/10.1007/BF02125347>.

- [27] Kevin J. Compton. “0-1 laws in logic and combinatorics”. In: *Algorithms and order (Ottawa, ON, 1987)*. Kluwer Acad. Publ., Dordrecht, 1989, pp. 353–383. URL: http://dx.doi.org/10.1007/978-94-009-2639-4_10.
- [28] Gabriel Conant. “Neostability in countable homogeneous metric spaces”. Apr. 2015. arXiv:1504.02427[math.LO]. URL: <http://arxiv.org/abs/1504.02427>.
- [29] Persi Diaconis and Svante Janson. “Graph limits and exchangeable random graphs”. *Rend. Mat. Appl. (7)* 28.1 (2008), pp. 33–61. URL: <http://arxiv.org/abs/0712.2749>.
- [30] Rick Durrett. *Probability: theory and examples*. Fourth edition. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2010.
- [31] Mirna Džamonja and Saharon Shelah. “On \triangleleft^* -maximality”. *Ann. Pure Appl. Logic* 125.1-3 (2004), pp. 119–158. URL: <http://arxiv.org/abs/math/0009087>.
- [32] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2006.
- [33] P. Erdős, D. J. Kleitman, and B. L. Rothschild. “Asymptotic enumeration of K_n -free graphs”. In: *Colloquio Internazionale sulle Teorie Combinatorie (Rome, 1973), Tomo II*. Atti dei Convegni Lincei No. 17. Accad. Naz. Lincei, Rome, 1976, pp. 19–27. URL: https://www.renyi.hu/~p_erdos/1976-03.pdf.
- [34] David M. Evans and Mark Wing Ho Wong. “Some remarks on generic structures”. *J. Symbolic Logic* 74.4 (2009), pp. 1143–1154. URL: <http://dx.doi.org/10.2178/jsl/1254748684>.
- [35] Ronald Fagin. “Probabilities on finite models”. *J. Symbolic Logic* 41.1 (1976), pp. 50–58. URL: <http://dx.doi.org/10.2307/2272945>.
- [36] Philippe Flajolet and Robert Sedgewick. *Analytic combinatorics*. Cambridge University Press, Cambridge, 2009.
- [37] Haim Gaifman. “Concerning measures in first order calculi”. *Israel J. Math.* 2 (1964), pp. 1–18. URL: <http://dx.doi.org/10.1007/BF02759729>.
- [38] Yu.V. Glebskii, D.I. Kogan, M.I. Liogon’kii, and V.A. Talanov. “Range and degree of realizability of formulas in the restricted predicate calculus”. English translation. *Cybernetics* 5.2 (1969), pp. 142–154. URL: <http://dx.doi.org/10.1007/BF01071084>.
- [39] John Goodrick, Byunghan Kim, and Alexei Kolesnikov. “Type-amalgamation properties and polygroupoids in stable theories”. *J. Math. Log.* 15.1 (2015). URL: <http://arxiv.org/abs/1404.1525>.
- [40] Yuri Gurevich. “Zero-one laws”. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS* 46 (1992), pp. 90–106. URL: <http://research.microsoft.com/en-us/um/people/gurevich/Opera/95.pdf>.

- [41] C. Ward Henson. “A family of countable homogeneous graphs”. *Pacific J. Math.* 38 (1971), pp. 69–83. URL: <http://dx.doi.org/10.2140/pjm.1971.38.69>.
- [42] Cameron Donnay Hill. “On semi/tame classes of finite structures”. Preprint. 2016. URL: <http://cdhill.web.wesleyan.edu/TameClasses.pdf>.
- [43] Wilfrid Hodges. *Model theory*. Vol. 42. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, 1993.
- [44] Douglas N. Hoover. “Relations on probability spaces and arrays of random variables”. Unpublished. 1979. URL: <https://www.stat.berkeley.edu/~aldous/Research/hoover.pdf>.
- [45] Ehud Hrushovski. “A new strongly minimal set”. *Ann. Pure Appl. Logic* 62.2 (1993). Stability in model theory, III (Trento, 1991), pp. 147–166. URL: [http://dx.doi.org/10.1016/0168-0072\(93\)90171-9](http://dx.doi.org/10.1016/0168-0072(93)90171-9).
- [46] Ehud Hrushovski. “Groupoids, imaginaries and internal covers”. *Turkish J. Math.* 36.2 (2012), pp. 173–198. URL: <http://arxiv.org/abs/math/0603413>.
- [47] T. Hyttinen and M. Kesälä. “Independence in finitary abstract elementary classes”. *Ann. Pure Appl. Logic* 143.1-3 (2006), pp. 103–138. URL: <http://dx.doi.org/10.1016/j.apal.2006.01.009>.
- [48] A. A. Ivanov. “Generic expansions of ω -categorical structures and semantics of generalized quantifiers”. *J. Symbolic Logic* 64.2 (1999), pp. 775–789. URL: <http://dx.doi.org/10.2307/2586500>.
- [49] Svante Janson. *Graphons, cut norm and distance, couplings and rearrangements*. Vol. 4. NYJM Monographs. New York Journal of Mathematics, Albany, NY, 2013. URL: <http://arxiv.org/abs/1009.2376>.
- [50] Peter T. Johnstone. *Stone spaces*. Vol. 3. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge, 1982.
- [51] Olav Kallenberg. “Symmetries on random arrays and set-indexed processes”. *J. Theoret. Probab.* 5.4 (1992), pp. 727–765. URL: <http://dx.doi.org/10.1007/BF01058727>.
- [52] Olav Kallenberg. “Multivariate sampling and the estimation problem for exchangeable arrays”. *J. Theoret. Probab.* 12.3 (1999), pp. 859–883. URL: <http://dx.doi.org/10.1023/A:1021692202530>.
- [53] Olav Kallenberg. *Probabilistic symmetries and invariance principles*. Probability and its Applications (New York). Springer, New York, 2005.
- [54] W. M. Kantor, Martin W. Liebeck, and Dugald Macpherson. “ \aleph_0 -categorical structures smoothly approximated by finite substructures”. *Proc. London Math. Soc. (3)* 59.3 (1989), pp. 439–463. URL: <http://dx.doi.org/10.1112/plms/s3-59.3.439>.
- [55] Alexander S. Kechris. *Classical descriptive set theory*. Vol. 156. Graduate Texts in Mathematics. Springer-Verlag, New York, 1995.

- [56] Alexander S. Kechris and Christian Rosendal. “Turbulence, amalgamation, and generic automorphisms of homogeneous structures”. *Proc. Lond. Math. Soc.* 94.2 (2007), pp. 302–350. URL: <http://arxiv.org/abs/math/0409567>.
- [57] H. Jerome Keisler and Julia F. Knight. “Barwise: infinitary logic and admissible sets”. *Bull. Symbolic Logic* 10.1 (2004), pp. 4–36. URL: <http://dx.doi.org/10.2178/bsl/1080330272>.
- [58] Byunghan Kim and Hyeung-Joon Kim. “Notions around tree property 1”. *Ann. Pure Appl. Logic* 162.9 (2011), pp. 698–709. URL: <http://dx.doi.org/10.2178/bsl/1080330272>.
- [59] Byunghan Kim, Alexei S. Kolesnikov, and Akito Tsuboi. “Generalized amalgamation and n -simplicity”. *Ann. Pure Appl. Logic* 155.2 (2008), pp. 97–114. URL: <http://dx.doi.org/10.1016/j.apal.2008.03.005>.
- [60] Byunghan Kim and Anand Pillay. “Simple theories”. *Ann. Pure Appl. Logic* 88.2-3 (1997). Joint AILA-KGS Model Theory Meeting (Florence, 1995), pp. 149–164. URL: [http://dx.doi.org/10.1016/S0168-0072\(97\)00019-5](http://dx.doi.org/10.1016/S0168-0072(97)00019-5).
- [61] Byunghan Kim and Anand Pillay. “From stability to simplicity”. *Bull. Symbolic Logic* 4.1 (1998), pp. 17–36. URL: <http://dx.doi.org/10.2307/421004>.
- [62] Byunghan Kim and Anand Pillay. “Around stable forking”. *Fund. Math.* 170.1-2 (2001), pp. 107–118. URL: <http://dx.doi.org/10.4064/fm170-1-6>.
- [63] Alexei S. Kolesnikov. “ n -simple theories”. *Ann. Pure Appl. Logic* 131.1-3 (2005), pp. 227–261. URL: <http://dx.doi.org/10.1016/j.apal.2004.06.001>.
- [64] Alex Kruckman. “Disjoint n -amalgamation and pseudofinite countably categorical theories”. Submitted to *Notre Dame J. Form. Log.* Oct. 2015. arXiv:1510.03539[math.LO]. URL: <http://arxiv.org/abs/1510.03539>.
- [65] Wiesław Kubiś. “Fraïssé sequences: category-theoretic approach to universal homogeneous structures”. *Ann. Pure Appl. Logic* 165.11 (2014), pp. 1755–1811. URL: <http://arxiv.org/abs/0711.1683>.
- [66] David W. Kueker and Michael C. Laskowski. “On generic structures”. *Notre Dame J. Formal Logic* 33.2 (1992), pp. 175–183. URL: <http://dx.doi.org/10.1305/ndjfl/1093636094>.
- [67] K. Kuratovskii. “A certain theorem on ideals, and its applications to the Baire property of Polish spaces”. English translation. *Russian Math. Surveys* 31.5 (1976), pp. 124–127. URL: <http://dx.doi.org/10.1070/RM1976v031n05ABEH004192>.
- [68] László Lovász. *Large networks and graph limits*. Vol. 60. American Mathematical Society Colloquium Publications. American Mathematical Society, Providence, RI, 2012.
- [69] László Lovász and Balázs Szegedy. “Limits of dense graph sequences”. *J. Combin. Theory Ser. B* 96.6 (2006), pp. 933–957. URL: <http://arxiv.org/abs/math/0408173>.

- [70] László Lovász and Balázs Szegedy. “Regularity partitions and the topology of graphons”. In: *An irregular mind*. Vol. 21. Bolyai Soc. Math. Stud. János Bolyai Math. Soc., Budapest, 2010, pp. 415–446. URL: <http://arxiv.org/abs/1002.4377>.
- [71] László Lovász and Balázs Szegedy. “Finitely forcible graphons”. *J. Combin. Theory Ser. B* 101.5 (2011), pp. 269–301. URL: <http://arxiv.org/abs/0901.0929>.
- [72] Dugald Macpherson. “Finite axiomatizability and theories with trivial algebraic closure”. *Notre Dame J. Formal Logic* 32.2 (1991), pp. 188–192. URL: <http://dx.doi.org/10.1305/ndjfl/1093635744>.
- [73] Dugald Macpherson. “A survey of homogeneous structures”. *Discrete Math.* 311.15 (2011), pp. 1599–1634. URL: <http://dx.doi.org/10.1016/j.disc.2011.01.024>.
- [74] Dugald Macpherson and Charles Steinhorn. “Definability in classes of finite structures”. In: *Finite and algorithmic model theory*. Vol. 379. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge, 2011, pp. 140–176. URL: <http://dx.doi.org/10.1017/CB09780511974960.005>.
- [75] David Marker. *Model theory: an introduction*. Vol. 217. Graduate Texts in Mathematics. Springer-Verlag, New York, 2002.
- [76] Richard Marshall. “Robust classes of finite structures”. PhD thesis. University of Leeds, Dec. 2007.
- [77] Michael Morley. “The number of countable models”. *J. Symbolic Logic* 35 (1970), pp. 14–18. URL: <http://dx.doi.org/10.2307/2271150>.
- [78] Jaroslav Nešetřil and Patrice Ossona de Mendez. “A unified approach to structural limits, and limits of graphs with bounded tree-depth”. Mar. 2013. arXiv:1303.6471[math.CO]. URL: <http://arxiv.org/abs/1303.6471>.
- [79] Walter Oberschelp. “Monotonicity for structure numbers in theories without identity”. In: *Combinatoire et représentation du groupe symétrique (Actes Table Ronde CNRS, Univ. Louis-Pasteur Strasbourg, Strasbourg, 1976)*. Vol. 579. Lecture Notes in Math. Springer, Berlin, 1977, pp. 297–308. URL: <http://dx.doi.org/10.1007/BFb0090025>.
- [80] John C. Oxtoby. *Measure and category*. Second edition. Vol. 2. Graduate Texts in Mathematics. Springer-Verlag, New York-Berlin, 1980.
- [81] Fedor Petrov and Anatoly Vershik. “Uncountable graphs and invariant measures on the set of universal countable graphs”. *Random Structures Algorithms* 37.3 (2010), pp. 389–406. URL: <http://arxiv.org/abs/0804.3386>.
- [82] Massoud Pourmahdian. “Simple generic structures”. *Ann. Pure Appl. Logic* 121.2-3 (2003), pp. 227–260. URL: [http://dx.doi.org/10.1016/S0168-0072\(02\)00114-8](http://dx.doi.org/10.1016/S0168-0072(02)00114-8).
- [83] Alexander A. Razborov. “Flag algebras”. *J. Symbolic Logic* 72.4 (2007), pp. 1239–1282. URL: <http://dx.doi.org/10.2178/jsl/1203350785>.
- [84] Saharon Shelah. “Simple unstable theories”. *Ann. Math. Logic* 19.3 (1980), pp. 177–203. URL: [http://dx.doi.org/10.1016/0003-4843\(80\)90009-1](http://dx.doi.org/10.1016/0003-4843(80)90009-1).

- [85] Saharon Shelah. “Classification of nonelementary classes. II. Abstract elementary classes”. In: *Classification theory (Chicago, IL, 1985)*. Vol. 1292. Lecture Notes in Math. Springer, Berlin, 1987, pp. 419–497. URL: <http://dx.doi.org/10.1007/BFb0082243>.
- [86] Saharon Shelah. *Classification theory and the number of nonisomorphic models*. Second edition. Vol. 92. Studies in Logic and the Foundations of Mathematics. North-Holland Publishing Co., Amsterdam, 1990.
- [87] Saharon Shelah. “The universality spectrum: consistency for more classes”. In: *Combinatorics, Paul Erdős is eighty, Vol. 1*. Bolyai Soc. Math. Stud. János Bolyai Math. Soc., Budapest, 1993, pp. 403–420. URL: <http://arxiv.org/abs/math/9412229>.
- [88] Saharon Shelah. “Toward classifying unstable theories”. *Ann. Pure Appl. Logic* 80.3 (1996), pp. 229–255. URL: <http://arxiv.org/abs/math/9508205>.
- [89] Saharon Shelah. “Pseudofinite, \aleph_0 -categorical but not simple”. Unpublished note. 2001.
- [90] Saharon Shelah and Alexander Usvyatsov. “More on SOP_1 and SOP_2 ”. *Ann. Pure Appl. Logic* 155.1 (2008), pp. 16–31. URL: <http://arxiv.org/abs/math/0404178>.
- [91] Pierre Simon. “Distal and non-distal NIP theories”. *Ann. Pure Appl. Logic* 164.3 (2013), pp. 294–318. URL: <http://arxiv.org/abs/1103.2239>.
- [92] Terence Tao. *An introduction to measure theory*. Vol. 126. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2011.
- [93] Simon Thomas. “Reducts of random hypergraphs”. *Ann. Pure Appl. Logic* 80.2 (1996), pp. 165–193. URL: [http://dx.doi.org/10.1016/0168-0072\(95\)00061-5](http://dx.doi.org/10.1016/0168-0072(95)00061-5).