

COUNTABLE INFINITARY THEORIES ADMITTING AN INVARIANT MEASURE

NATHANAEL ACKERMAN, CAMERON FREER, AND REHANA PATEL

ABSTRACT. Let L be a countable language. We characterize, in terms of definable closure, those countable theories Σ of $\mathcal{L}_{\omega_1, \omega}(L)$ for which there exists an S_∞ -invariant probability measure on the collection of models of Σ with underlying set \mathbb{N} . Restricting to $\mathcal{L}_{\omega, \omega}(L)$, this answers an open question of Gaifman from 1964, via a translation between S_∞ -invariant measures and Gaifman’s symmetric measure-models with strict equality. It also extends the known characterization in the case where Σ implies a Scott sentence. To establish our result, we introduce machinery for building invariant measures from a directed system of countable structures with measures.

1. Introduction	1
1.1. Main result	2
1.2. Outline of the paper	3
2. Preliminaries	4
2.1. The logic action on the measurable space Str_L	5
2.2. Ergodic structures	5
2.3. Morleyization and Pithy Π_2 theories	7
2.4. Definable closure and duplication of formulas	11
3. Ergodic structures via sampling from Borel structures	14
3.1. Borel L -structures	14
3.2. Weighted homomorphism densities	16
4. Layerings	17
4.1. Basic layerings	17
4.2. Continuous limits	19
4.3. Regular layerings	20
4.4. Approximations to types	21
4.5. Layer Transformation	23
5. Existence of ergodic structures	26
6. Non-existence of ergodic structures	29
7. Classification of ergodic structures	30
Acknowledgements	31
References	31

1. INTRODUCTION

Logic and probability bear many formal resemblances, and there is a long history of model-theoretic approaches to unifying them. A seminal work extending classical model theory to random structures is Gaifman’s 1964 paper [Gai64]. This paper

provides coherence conditions for assigning probabilities to formulas from some first-order language, instantiated from a fixed domain, in a way that respects the logical relationships between formulas. Gaifman calls such an assignment of probabilities a *measure-model*; in the case where the assignment also respects equality, he calls it a *measure-model with strict equality*.

A key case that Gaifman addresses is that of a *symmetric measure-model*, where the probabilities assigned to a formula are invariant under arbitrary finite permutations of the instantiating domain. He shows how to obtain a symmetric measure-model for an arbitrary countable first-order theory, but demonstrates that some of these theories admit only symmetric measure-models without strict equality. This leaves open the following question from [Gai64, §4]:

“The problem of characterizing those theories (i.e. measures having the values 0 and 1) which possess a measure-model with strict equality, satisfying also the symmetry condition, seems to be difficult.”

Gaifman’s paper is concerned with first-order theories of $\mathcal{L}_{\omega,\omega}$, but his question is also natural in the infinitary setting of $\mathcal{L}_{\omega_1,\omega}$ explored by Scott and Krauss in [SK66] and [Kra69]. In the present paper, we answer the more general question for arbitrary countable theories of $\mathcal{L}_{\omega_1,\omega}$ in the case of countable domains, by developing methods that extend our earlier work on invariant measures [AFP16]. In doing so, we answer Gaifman’s original question, for countable domains and languages.

Our setting is the following; for more details, see §§2.1 and 2.2. Let L be a countable language, and write Str_L for the measurable space of L -structures with underlying set \mathbb{N} . The space S_∞ of permutations of \mathbb{N} acts on Str_L by permuting the underlying set. Given a countable collection of sentences $\Sigma \subseteq \mathcal{L}_{\omega_1,\omega}(L)$, we consider when there is an S_∞ -invariant probability measure on Str_L that assigns probability 1 to the class of models of Σ in Str_L .

For L a countable language and X a countable set, the paper [Gai64] shows essentially that every measure-model with strict equality, of formulas of $\mathcal{L}_{\omega_1,\omega}(L)$ instantiated by elements from X , is induced by some probability distribution on the class of L -structures with underlying set X . In the case where the underlying set X is \mathbb{N} , Gaifman’s symmetric measure-models with strict equality correspond to S_∞ -invariant probability measures on Str_L . An $\mathcal{L}_{\omega,\omega}(L)$ -theory Σ has a symmetric measure-model with strict equality when the corresponding S_∞ -invariant probability measure assigns probability 1 to the class of models of Σ in Str_L .

1.1. Main result. Our main theorem, Theorem 1.1, states that for a countable language L and countable theory Σ of $\mathcal{L}_{\omega_1,\omega}(L)$, there is an S_∞ -invariant probability measure concentrated on the class of models of Σ in Str_L precisely when Σ has a completion (in some countable fragment) that has trivial definable closure (for that fragment) — a criterion that is often easy to check in practice.

This theorem is a generalization of the main result of [AFP16], which considered only the case where Σ implies a Scott sentence. Recall that a Scott sentence is a sentence $\sigma \in \mathcal{L}_{\omega_1, \omega}(L)$ that has exactly one countable model up to isomorphism; in particular, $\{\sigma\}$, though not deductively closed, is a complete theory of $\mathcal{L}_{\omega_1, \omega}(L)$ in the sense that it implies φ or implies $\neg\varphi$ for each sentence $\varphi \in \mathcal{L}_{\omega_1, \omega}(L)$. Theorem 1.1 can therefore be viewed as a generalization of [AFP16] to countable theories of $\mathcal{L}_{\omega_1, \omega}(L)$ that are not necessarily complete for $\mathcal{L}_{\omega_1, \omega}(L)$.

For an S_∞ -invariant probability measure μ on Str_L , we say that μ is *ergodic* if every μ -almost S_∞ -invariant subset of Str_L is assigned measure 0 or 1 by μ . Every S_∞ -invariant probability measure can be decomposed as a convex combination of ergodic ones. Further, if there is an S_∞ -invariant probability measure concentrated on a given Borel set $B \subseteq \text{Str}_L$, then there is an ergodic such measure, and so it often suffices to consider only the ergodic S_∞ -invariant probability measures.

Theorem 1.1. *Let L be a countable language, and let $\Sigma \subseteq \mathcal{L}_{\omega_1, \omega}(L)$ be a countable set of sentences. Then the following are equivalent:*

- (1) *There is an S_∞ -invariant probability measure concentrated on the class of models of Σ in Str_L .*
- (2) *There is an ergodic S_∞ -invariant probability measure concentrated on the class of models of Σ in Str_L .*
- (3) *There is a countable fragment $A \subseteq \mathcal{L}_{\omega_1, \omega}(L)$ and a consistent theory $T \subseteq A$ that is complete for A , such that $\Sigma \subseteq T$ and T has trivial A -definable closure.*
- (4) *For all countable fragments $A \subseteq \mathcal{L}_{\omega_1, \omega}(L)$ such that $\Sigma \subseteq A$, there is a consistent theory $T \subseteq A$ that is complete for A , such that $\Sigma \subseteq T$ and T has trivial A -definable closure.*

Note that when Σ is itself complete for some countable fragment $A \subseteq \mathcal{L}_{\omega_1, \omega}(L)$, then Theorem 1.1 implies that there is an S_∞ -invariant probability measure concentrated on the class of models of Σ in Str_L if and only if Σ has trivial A -definable closure.

In the case where Σ is a set of first-order sentences, our main result simplifies to the following.

Theorem 1.2. *Let $\Sigma \subseteq \mathcal{L}_{\omega, \omega}(L)$ be a set of first-order sentences. Then there is an S_∞ -invariant probability measure concentrated on the class of models of Σ in Str_L if and only if there is some complete consistent theory $T \subseteq \mathcal{L}_{\omega, \omega}(L)$ with $\Sigma \subseteq T$ such that T has trivial $\mathcal{L}_{\omega, \omega}(L)$ -definable closure.*

This answers Gaifman's question for countable languages and domains.

1.2. Outline of the paper. We begin, in Section 2, by providing definitions regarding our setting, including the notion of an *ergodic structure*, i.e., an ergodic S_∞ -invariant probability measure on Str_L . We describe a process we call *pithy Π_2 Morleyization* that allows us to work with languages and theories that have

nice properties which allow us to carry out the main construction of our paper. We also provide basic results about trivial definable closure and its relation to duplication of formulas.

In Section 3, we describe a method for building ergodic structures via sampling from a Borel structure equipped with a measure.

Next, in Sections 4 and 5, a countable consistent theory that is complete for a countable fragment and has trivial definable closure for that fragment, we show how to build a Borel structure and measure such that sampling from it yields an *ergodic model* of the theory, i.e., an ergodic structure that concentrates on the collection of models of the theory. We do so by building a special type of directed system of finite structures with measures whose limit is such a Borel structure along with a measure.

In Section 6, we show that if a countable consistent theory that is complete for some countable fragment has non-trivial definable closure for that fragment, then it does not admit an ergodic model.

Finally, in Section 7, we combine the positive and negative results from Sections 5 and 6 to obtain our main result, Theorem 1.1, and its corollary for first-order languages, Theorem 1.2.

2. PRELIMINARIES

In this section, we introduce and develop notions involving invariant measures and certain kinds of theories that will be useful in our constructions.

Throughout this paper, let L be a countable language. We allow relation symbols to be 0-ary. The instantiation of a 0-ary relation in an L -structure is the assignment True or False; such relations allow us to simplify the Morleyization construction in §2.3. For a nested pair of countable languages $L' \subseteq L''$ and an L'' -structure \mathcal{N} , we write $\mathcal{N}|_{L'}$ to denote the reduct of \mathcal{N} to the language L' .

We will formally allow formulas only to contain $\wedge, \bigwedge, \neg, \exists$, and we use \vee, \bigvee, \forall as shorthand in the standard way. This simplification loses no generality, but sometimes allows us a cleaner presentation, as it reduces the number of cases to deal with for inductions on formulas.

Recall that a **fragment** (of $\mathcal{L}_{\omega_1, \omega}(L)$) is a subset $A \subseteq \mathcal{L}_{\omega_1, \omega}(L)$ that is closed under subformulas as well as the logical operations of $\wedge, \vee, \neg, (\exists x)$ and $(\forall x)$.

We say that a set of sentences $T \subseteq \mathcal{L}_{\omega_1, \omega}(L)$ is a **theory** when it is consistent. In general, we do not require theories to be deductively closed or complete (for either a countable fragment or for $\mathcal{L}_{\omega_1, \omega}(L)$).

Let A be a fragment. An **A -theory** T is a theory T such that $T \subseteq A$. An A -theory T is **complete for A** when for every sentence $\sigma \in A$, either $T \models \sigma$ or $T \models \neg\sigma$; in this case we say that T is a **complete A -theory**. Note that we do not require even complete A -theories to be deductively closed; this will be important when we work with A -theories of a restricted syntactic form (namely, pithy Π_2 , as described in §2.3).

For a measure m and singleton set $\{x\}$, we often abbreviate $m(\{x\})$ by the notation $m(x)$, and for a function i we similarly abbreviate the inverse image $i^{-1}(\{x\})$ by $i^{-1}(x)$. Likewise, we write $Y \cup x$ to denote $Y \cup \{x\}$ and $Y \setminus x$ to denote $Y \setminus \{x\}$. We write *qf-type* to mean a quantifier-free type.

2.1. The logic action on the measurable space Str_L . The measurable space Str_L has underlying space the collection of L -structures with underlying set \mathbb{N} , and σ -algebra generated by subbasic open sets of the form

$$\llbracket \varphi(n_1, \dots, n_j) \rrbracket := \{ \mathcal{M} \in \text{Str}_L : \mathcal{M} \models \varphi(n_1, \dots, n_j) \},$$

where φ is an atomic L -formula, j is its number of free variables, and $n_1, \dots, n_j \in \mathbb{N}$. In fact, $\llbracket \varphi(n_1, \dots, n_j) \rrbracket$ is Borel for arbitrary formulas φ of $\mathcal{L}_{\omega_1, \omega}(L)$, by [Kec95, Proposition 16.7]. Given a countable theory T , we write $\llbracket T \rrbracket$ to mean $\llbracket \bigwedge_{\varphi \in T} \varphi \rrbracket$.

Let S_∞ denote the permutation group of \mathbb{N} . The *logic action* of S_∞ on Str_L is the action induced by permutation of the underlying set \mathbb{N} ; for more details, see [Kec95, §16.C]. We say that a Borel probability measure m on Str_L is **invariant** when it is invariant under the action of S_∞ , i.e., $m(X) = m(g \cdot X)$ for every Borel set $X \subseteq \text{Str}_L$ and $g \in S_\infty$; we often simply call such an m an **invariant measure**. A probability measure m is **concentrated** on a Borel set $X \subseteq \text{Str}_L$ when $m(X) = 1$.

The Lopez-Escobar theorem (see, e.g., [Kec95, Theorem 16.8]) states that a set $X \subseteq \text{Str}_L$ is Borel and invariant under the logic action if and only if there is a sentence $\varphi \in \mathcal{L}_{\omega_1, \omega}(L)$ such that $X = \llbracket \varphi \rrbracket$.

We will sometimes speak of an invariant measure concentrated on a sentence φ , by which we mean an invariant measure concentrated on $\llbracket \varphi \rrbracket$, the class of models of φ in Str_L . We will likewise speak of an invariant measure concentrated on a countable theory T , by which we mean an invariant measure concentrated on $\llbracket T \rrbracket$, the class of models of T in Str_L , and in this case say that T **admits** an invariant measure. Given a countable collection Θ of qf-types, we say that an invariant measure **omits** Θ when it is concentrated on the class of structures in Str_L that omit every qf-type in Θ , i.e., on $\llbracket \bigwedge_{p(\bar{x}) \in \Theta} (\forall \bar{x}) \neg p(\bar{x}) \rrbracket$.

2.2. Ergodic structures. In order to determine which countable theories admit an invariant measure, it will suffice to ask which admit an ergodic invariant measure.

Definition 2.1. *We say that a probability measure μ on Str_L is **ergodic** if $\mu(B) \in \{0, 1\}$ for all Borel sets B that satisfy $\mu(B \Delta \tau^{-1}(B)) = 0$ for every $\tau \in S_\infty$.*

In other words, an ergodic probability measure μ is one that does not assign intermediate measure to any μ -almost invariant set.

We will use the following standard result in the proof of Theorem 1.1; we include its proof for completeness.

Lemma 2.2. *If $\Sigma \subseteq \mathcal{L}_{\omega_1, \omega}(L)$ is countable set of sentences that admits an S_∞ -invariant measure, then Σ admits an ergodic S_∞ -invariant measure.*

Proof. Let μ be an invariant measure concentrated on Σ , and suppose that Σ does not admit an ergodic invariant measure. The measure μ can be decomposed into a mixture of ergodic invariant measures on Str_L (see, e.g., [Kal05, Lemma A1.2 and Theorem A1.3]). By hypothesis, none of the measures in the decomposition is concentrated on Σ . By the Lopez-Escobar theorem, $\llbracket \Sigma \rrbracket$ is Borel invariant, and so the measures in the decomposition must therefore all be concentrated on the complement of $\llbracket \Sigma \rrbracket$. But then μ , which is a mixture of the measures in the decomposition, is concentrated on the complement of $\llbracket \Sigma \rrbracket$, a contradiction. \square

Ergodic invariant measures can be thought of as “probabilistic” structures, as we now describe.

Definition 2.3. *An **ergodic L -structure** is an ergodic invariant measure on Str_L . An ergodic L -structure μ is said to **almost surely satisfy** a sentence $\varphi \in \mathcal{L}_{\omega_1, \omega}(L)$ when $\mu(\llbracket \varphi \rrbracket) = 1$.*

Lemma 2.4. *Let μ be an ergodic L -structure, and define*

$$\text{Th}(\mu) := \{\varphi \in \mathcal{L}_{\omega_1, \omega}(L) : \mu(\llbracket \varphi \rrbracket) = 1\}.$$

Then $\text{Th}(\mu)$ is a complete, deductively closed $\mathcal{L}_{\omega_1, \omega}(L)$ -theory.

Proof. Because the measure μ is ergodic, for any sentence $\varphi \in \mathcal{L}_{\omega_1, \omega}(L)$, exactly one of $\varphi \in \text{Th}(\mu)$ or $\neg\varphi \in \text{Th}(\mu)$ holds, by the Lopez-Escobar theorem. Deductive closure follows from σ -additivity, and so $\text{Th}(\mu)$ is also consistent, hence an $\mathcal{L}_{\omega_1, \omega}(L)$ -theory. \square

Definition 2.5. *We say that μ is an **ergodic model** of T when $T \subseteq \text{Th}(\mu)$; in that case we say that $\text{Th}(\mu)$ is the theory of μ .*

An ergodic structure therefore has a complete $\mathcal{L}_{\omega_1, \omega}(L)$ -theory, and it is an ergodic model of this theory. This provides some justification for considering ergodic structures as probabilistic generalizations of classical model-theoretic structures.

Finally, we say that an ergodic L -structure almost surely has a property P when it assigns measure 1 to the collection of elements of Str_L having property P . For example, we say that μ almost surely omits a countable collection of qf-types Θ when $\mu(\llbracket \bigwedge_{p(\bar{x}) \in \Theta} (\forall \bar{x}) \neg p(\bar{x}) \rrbracket) = 1$. We will be especially interested in ergodic models of a given theory that almost surely omit a particular collection of qf-types.

2.3. Morleyization and Pithy Π_2 theories. In our main construction it will be important to work with a first-order theory consisting of sentences of a specific form which can be thought of as “one-point extension axioms”. It will also be important that our measure concentrate on the collection of models of a theory of

one-point extension axioms that omits a countable collection of (non-principal) qf-types, which we obtain by a variant of a standard construction.

The notion of *non-redundant* tuples, formulas, structures, and theories will be important when formulating the notions of *duplication of qf-types* (§2.4) and *layering transformations* (§4.5).

Definition 2.6. *A tuple $a_1 \cdots a_n$ is **non-redundant** if $a_i \neq a_j$ for $i \neq j$. A formula with free variables x_1, \dots, x_n is **non-redundant** if it implies the formula $\bigwedge_{1 \leq i < j \leq n} (x_i \neq x_j)$. A structure in a relational language is **non-redundant** if each relation holds only on non-redundant tuples, and a theory in a relational language is **non-redundant** if all of its models are non-redundant.*

We now introduce the special form of sentences, which we call pithy Π_2 . We then show that for any countable fragment A of $\mathcal{L}_{\omega_1, \omega}(L)$, there is a relational language L_A , a pithy Π_2 first-order L_A -theory Th_A , and a countable collection of qf-types Θ_A of L_A , such that each L -structure has a common definable expansion with some model of Th_A omitting Θ_A .

Definition 2.7. *A sentence $\varphi \in \mathcal{L}_{\omega_1, \omega}(L)$ is said to be **pithy Π_2** if it is of the form $(\forall \bar{x})(\exists y)\psi(\bar{x}, y)$ where $\psi(\bar{x}, y)$ is quantifier-free, and \bar{x} is a finite (possibly empty) sequence of variables. A countable theory T is said to be pithy Π_2 when every sentence in T is pithy Π_2 .*

Throughout the rest of this subsection, let A be a countable fragment of $\mathcal{L}_{\omega_1, \omega}(L)$ and let T be an A -theory.

We now define the relational language L_A . For $m \in \mathbb{N}$ write $[m] := \{1, \dots, m\}$. Let I be the set of triples $(\varphi(\bar{x}), \iota, \kappa)$ for which there exist $n \in \mathbb{N}$ and $r \leq n$ such that:

- $\varphi(\bar{x}) \in A$ is an L -formula with $\bar{x} = x_1 \cdots x_n$ its tuple of distinct free variables,
- $\iota: [n] \rightarrow [r]$ is a surjection, and
- $\kappa: [r] \rightarrow [n]$ is an injection such that $\iota \circ \kappa$ is the identity on $[r]$.

Then define the relational language

$$L_A := \{Q_{\varphi(\bar{x}), \iota, \kappa} : (\varphi(\bar{x}), \iota, \kappa) \in I\},$$

where each $Q_{\varphi(\bar{x}), \iota, \kappa}$ is a relation symbol of arity $|\bar{x}|$.

For each surjection $\iota: [n] \rightarrow [r]$, define

$$\text{eq}_\iota := \bigwedge_{j, \ell \in [n]: \iota(j) = \iota(\ell)} (x_j = x_\ell) \wedge \bigwedge_{j, \ell \in [n]: \iota(j) \neq \iota(\ell)} (x_j \neq x_\ell).$$

Now for each $\eta(\bar{x}) \in A$, define the formula

$$R_{\eta(\bar{x})}(\bar{x}) := \bigvee_{\iota, \kappa: (\eta(\bar{x}), \iota, \kappa) \in I} \{\text{eq}_\iota \wedge Q_{\eta(\bar{x}), \iota, \kappa}(x_{\kappa(1)} \cdots x_{\kappa(r)})\},$$

where r is the arity of $Q_{\eta(\bar{x}),\iota,\kappa}$ and $\bar{x} = x_1 \cdots x_n$ is the tuple of distinct free variables of η .

After two observations, we next define the theory Th_A connecting L_A to L (which will be the pithy Π_2 Morleyization of the empty theory). First, note that we have described multiple relation symbols that our theory will prove equivalent; namely, whenever $(\varphi(\bar{x}), \iota, \kappa), (\varphi(\bar{x}), \iota, \kappa') \in I$ we have

$$\text{Th}_A \models Q_{\varphi(\bar{x}),\iota,\kappa}(\bar{y}) \leftrightarrow Q_{\varphi(\bar{x}),\iota,\kappa'}(\bar{y}),$$

though this does not pose a problem (and simplifies naming of the Q relation symbols).

Second, note that the language L_A is “built from” L in the sense that for each atomic L -formula $\eta(\bar{x})$ there is a first-order quantifier-free L_A -formula $R_\eta(\bar{x})$ that Th_A will prove equivalent to $\eta(\bar{x})$. However, L_A does not literally contain L as a sublanguage (and we will later refer to the language $L_A \cup L$ when needed). Note that when φ is a sentence, then each $Q_{(\varphi(\bar{x}),\iota,\kappa)}$ is a 0-ary relation symbol, and R_φ is a 0-ary formula, whose presence simplifies the technicalities of the pithy Π_2 Morleyization.

Define the pithy Π_2 -theory Th_A to be the collection of sentences

$$\bullet (\forall \bar{y}) [Q_{(\zeta(\bar{u}),\iota,\kappa)}(\bar{y}) \rightarrow \bigwedge_{j,\ell \in [k]: j \neq \ell} (y_j \neq y_\ell)]$$

for all $(\zeta(\bar{u}), \iota, \kappa) \in I$ where $\bar{y} = y_1 \cdots y_k$ is the tuple of distinct free variables in $Q_{(\zeta(\bar{u}),\iota,\kappa)}$, and

- $(\forall \bar{x}) [R_{\neg\psi}(\bar{x}) \leftrightarrow \neg R_\psi(\bar{x})]$,
- $(\forall \bar{x}) [R_{\xi_0 \wedge \xi_1}(\bar{x}) \leftrightarrow (R_{\xi_0}(\bar{w}_0) \wedge R_{\xi_1}(\bar{w}_1))]$,
- $(\forall \bar{x}) [R_{(\exists y)\varphi}(\bar{x}) \leftrightarrow (\exists y) R_\varphi(\bar{x}, y)]$,
- $(\forall \bar{x}) [R_{(\exists y)\psi}(\bar{x}) \leftrightarrow (\exists y) R_\psi(\bar{x})]$, and
- $(\forall \bar{x}) [R_{\bigwedge_{i \in \omega} \eta_i}(\bar{x}) \rightarrow R_{\eta_j}(\bar{z}_j)]$ for all $j \in \omega$,

for all $\xi_0, \xi_1, \psi, \varphi, \bigwedge_{i \in \omega} \eta_i \in A$; where \bar{x} is a tuple containing precisely the free variables of ψ ; where the tuple \bar{w}_i contains precisely the free variables of ξ_i for $i = 0, 1$ and \bar{x} contains precisely the variables occurring in \bar{w}_0 or \bar{w}_1 , where the tuple \bar{z}_i contains precisely the free variables of η_i for each $i \in \omega$ and \bar{x} contains precisely the variables occurring in some \bar{z}_i ; and where the free variables of φ are precisely the variables in $\bar{x}y$, with $y \notin \bar{x}$.

We also define a theory $E_{L,A}$ that will allow us to restrict to the case of relational languages in our main constructions. Let $E_{L,A}$ be the $(L_A \cup L)$ -theory consisting of $(\forall \bar{w}) [R_{\zeta(\bar{w})}(\bar{w}) \leftrightarrow \zeta(\bar{w})]$, where $\zeta(\bar{w})$ is an atomic L -formula whose free variables are precisely those in the tuple \bar{w} .

Lemma 2.8. *Th_A is a non-redundant pithy Π_2 first-order L_A -theory.*

Proof. Th_A is a first-order L_A -theory by definition, and is non-redundant by its first set of axioms (involving every relation symbol $Q_{(\zeta(\bar{u}),\iota,\kappa)}$ in the language). The theory Th_A is also pithy Π_2 by the form of the remaining axioms and the fact

that each formula R_ζ has a quantifier-free definition in terms of the Q relations symbols. \square

We now define a countable collection of qf-types Θ_A in the language L_A . For every formula of the form $\bigwedge_{i \in \omega} \eta_i \in A$, define the qf-type $p_{\bigwedge_{i \in \omega} \eta_i}$ by

$$p_{\bigwedge_{i \in \omega} \eta_i}(\bar{x}) := \{\neg R_{\bigwedge_{i \in \omega} \eta_i}(\bar{x})\} \cup \{R_{\eta_i}(\bar{z}_i) : i \in \omega\},$$

where \bar{x} and \bar{z}_i are tuples of variables as above. Then define the collection

$$\Theta_A := \{p_{\bigwedge_{i \in \omega} \eta_i} : \bigwedge_{i \in \omega} \eta_i \in A\}.$$

Note that if there is some $n \in \omega$ such that

$$\models (\forall \bar{x}) \left(\bigwedge_{i \in \omega} \eta_i(\bar{z}_i) \leftrightarrow \bigwedge_{i \leq n} \eta_i(\bar{z}_i) \right),$$

then $p_{\bigwedge_{i \in \omega} \eta_i}$ is inconsistent with $\text{Th}_A \cup \text{E}_{L,A}$. However, this is not a problem, as then the qf-type is automatically omitted.

Define the L_A -theory Th_A^+ to be

$$\text{Th}_A^+ := \text{Th}_A \cup \{(\forall \bar{x}) \neg p(\bar{x}) : p(\bar{x}) \in \Theta_A\},$$

where the tuple of variables \bar{x} is of the appropriate length for each qf-type $p(\bar{x}) \in \Theta_A$. In fact, a model of Th_A^+ is precisely the *pithy Π_2 Morleyization* of some L -structure (see Definition 2.10 below).

An easy induction on L -formulas shows that in any model \mathcal{K} of $\text{Th}_A^+ \cup \text{E}_{L,A}$, for all formulas $\psi \in A$, we have

$$\mathcal{K} \models (\forall \bar{x}) (\psi(\bar{x}) \leftrightarrow R_\psi(\bar{x})),$$

where \bar{x} contains precisely the free variables of ψ , i.e., \mathcal{K} omits Θ_A . In fact, we have the following.

Lemma 2.9. *For every L -structure \mathcal{M} there is a unique expansion of \mathcal{M} to an $(L_A \cup L)$ -structure \mathcal{M}_A such that $\mathcal{M}_A \models \text{Th}_A^+ \cup \text{E}_{L,A}$. In addition, for every L_A -structure \mathcal{N} that satisfies Th_A^+ , there is a unique expansion of \mathcal{N} to an $(L_A \cup L)$ -structure $\mathcal{N}^E \models \text{E}_{L,A}$. Further, $(\mathcal{M}_A|_{L_A})^E = \mathcal{M}_A$, and $(\mathcal{N}^E|_L)_A = \mathcal{N}^E$.*

Proof. Let \mathcal{M}_A be the $(L_A \cup L)$ -structure that satisfies $\text{Th}_A^+ \cup \text{E}_{L,A}$ and whose reduct to L is \mathcal{M} . Such a structure clearly exists and is the unique $(L_A \cup L)$ -structure with these properties.

Likewise, let \mathcal{N}^E be the unique $(L_A \cup L)$ -structure that satisfies $\text{Th}_A^+ \cup \text{E}_{L,A}$ and whose reduct to L_A is \mathcal{N} . \square

Definition 2.10. *Let \mathcal{M} be an L -structure. The **pithy Π_2 Morleyization** of \mathcal{M} for A is the L_A -structure $\mathcal{M}_A|_{L_A}$.*

Observe that any pithy Π_2 Morleyization is in a relational language (by the construction of L_A) and is non-redundant (by Lemma 2.8).

Define the countable pithy Π_2 first-order L_A -theory $T_A := \text{Th}_A \cup \{R_\sigma : \sigma \in T\}$, which we will use in Section 5. Further define the $\mathcal{L}_{\omega_1, \omega}(L_A)$ -theory $T_A^+ := \text{Th}_A^+ \cup \{R_\sigma : \sigma \in T\}$, and let A^+ be the smallest L_A -fragment containing Th_A^+ . Observe that A^+ is countable and T_A^+ is an A^+ -theory; we will later see in Corollary 2.16 that T_A^+ is a complete A^+ -theory.

Note that an L_A -structure satisfies T_A^+ if and only if it satisfies T_A and omits Θ_A . The theory T_A^+ has the following key property.

Lemma 2.11. *Let \mathcal{M} be an L -structure. For every L -sentence $\sigma \in A$, we have $\mathcal{M} \models \sigma$ if and only if $\mathcal{M}_A \models R_\sigma$. In particular, \mathcal{M}_A is a definable expansion of \mathcal{M} for A , and $\mathcal{M} \models T$ if and only if $\mathcal{M}_A \models T_A^+$.*

Proof. We have $\mathcal{M}_A \models \text{Th}_A^+ \cup \text{E}_{L,A}$ by Lemma 2.9. Hence for atomic formulas σ , we have $\mathcal{M} \models \sigma$ if and only if $\mathcal{M}_A \models R_\sigma$. This equivalence holds for all σ via an easy induction on formulas, using the way that Th_A^+ was constructed. \square

For every L -type p over A that is consistent with T , define the qf L_A -type $p_A(\bar{x}) := \{R_\psi(\bar{x}) : \psi \in p\}$. The following statement that types lift to qf-types is immediate from Lemma 2.11.

Corollary 2.12. *Let p be a L -type over A that is consistent with T . Then p_A is consistent with Th_A^+ and such that $\mathcal{M} \models p(\bar{a})$ iff $\mathcal{M}_A \models p_A(\bar{a})$ for every model \mathcal{M} of T and tuple $\bar{a} \in \mathcal{M}$ whose length is the number of free variables of p_A .*

We now establish a tight connection between models of T and of T_A^+ .

Corollary 2.13. *The map $\mathcal{M} \mapsto \mathcal{M}_A|_{L_A}$ restricts to a Borel bijection between $\llbracket T \rrbracket$ and $\llbracket T_A^+ \rrbracket$ that commutes with the logic action and preserves convex combinations of probability measures.*

Proof. By Lemma 2.9 and Corollary 2.12, the map $\mathcal{M} \mapsto \mathcal{M}_A|_{L_A}$ is a bijection between Str_L and the set of models of Th_A^+ within Str_{L_A} . This bijection is clearly Borel; note that it commutes with the logic action and preserves convex combinations of probability measures. Further, by Lemma 2.11, the image of $\llbracket T \rrbracket$ is $\llbracket T_A^+ \rrbracket$. Finally, note that every model of T_A^+ is the pithy Π_2 Morleyization of some model of T . \square

This yields the following corollary for ergodic structures, since the ergodic structures are precisely those invariant measures that cannot be decomposed as a non-trivial convex combination of others (again see [Kal05, Lemma A1.2 and Theorem A1.3]).

Corollary 2.14. *T has an ergodic model if and only if T_A^+ does.*

Proof. By Corollary 2.13, there is a bijection between the set of ergodic L -structures that almost surely satisfy T and the ergodic L_A -structures that almost

surely satisfy T_A^+ . In particular, one such set of ergodic structures is non-empty if and only if the other is non-empty. \square

This corollary reduces the problem of determining whether there is an ergodic model of an A -theory to that of determining whether there is an ergodic model of a corresponding first-order theory that omits a certain countable set of qf-types.

The following result will be useful when considering trivial A -definable closure in §2.4.

Corollary 2.15. *For any sentence $\sigma \in A$, we have*

$$T \models \sigma \quad \text{if and only if} \quad T_A^+ \models R_\sigma.$$

Proof. By the completeness theorem for sentences of $\mathcal{L}_{\omega_1, \omega}(L)$, we have $T \models \sigma$ if and only if $\llbracket T \cup \{\neg\sigma\} \rrbracket$ is empty. By Corollary 2.14, $\llbracket T \cup \{\neg\sigma\} \rrbracket$ is empty if and only if $\llbracket T_A^+ \cup \{R_{\neg\sigma}\} \rrbracket$ is empty.

Again by completeness, $T_A^+ \models R_\sigma$ if and only if $\llbracket T_A^+ \cup \{\neg R_\sigma\} \rrbracket$ is empty. But

$$\llbracket T_A^+ \cup \{\neg R_\sigma\} \rrbracket = \llbracket (T_A^+ \cup \{R_{\neg\sigma}\}) \rrbracket,$$

and so we are done. \square

The following corollary will be used in the proof of our main result, Theorem 1.1.

Corollary 2.16. *For any complete A -theory T , the theory T_A^+ is a complete A^+ -theory.*

Proof. By an easy induction on formulas, for every formula $\varphi \in A^+$, there is a relation $R \in L_A$, of arity the number of free variables of φ , such that $\text{Th}_A^+ \models (\forall \bar{x})(R(\bar{x}) \leftrightarrow \varphi(\bar{x}))$. Therefore the result follows by Corollary 2.15. \square

2.4. Definable closure and duplication of formulas. The main result of [AFP16] is that the only obstacle to the existence of an ergodic structure concentrated on the isomorphism class of a given structure \mathcal{M} is the presence of a tuple whose definable closure in \mathcal{M} includes additional elements. The main result of the present paper shows that the only obstruction to the existence of an ergodic structure concentrated on the collection of models of a complete A -theory $T \subseteq \mathcal{L}_{\omega_1, \omega}(L)$ is the presence of a formula that T proves *uniformly* witnesses the non-triviality of definable closure in the models of T .

In this subsection we introduce the notion of *trivial A -definable closure* over a fragment A as well as a useful equivalent concept, that of having *duplication of formulas* in A .

Definition 2.17. *Let A be a fragment and let \mathcal{M} be an L -structure. Let $T^{\mathcal{M}}$ be the complete A -theory of \mathcal{M} . The **A -definable closure** of a set $X \subseteq \mathcal{M}$ is defined to be the set $\text{dcl}_A(X)$ of all elements $b \in \mathcal{M}$ for which there is a formula $\varphi(\bar{x}, y) \in A$ and a tuple \bar{a} of length $|\bar{x}|$ all of whose elements are in X , such that*

$$\mathcal{M} \models \varphi(\bar{a}, b) \quad \text{and} \quad T^{\mathcal{M}} \models (\forall \bar{x})(\exists^{\leq 1} y) \varphi(\bar{x}, y),$$

where $(\exists^=1y)$ means “there exists a unique y ”. The A -**algebraic closure** of X , written $\text{acl}_A(X)$ is defined similarly except with $(\exists^=1y)$ replaced by $(\exists^{<\omega}y)$, i.e., “there exist finitely many y ”.

Recall that the Scott sentence $\sigma_{\mathcal{M}} \in \mathcal{L}_{\omega_1, \omega}(L)$ of a countable L -structure \mathcal{M} determines \mathcal{M} up to isomorphism among countable structures. When A contains the Scott sentence of \mathcal{M} , then $\text{dcl}_A(X)$ is the same as the usual “group-theoretic” notion of definable closure in \mathcal{M} , i.e., the set $\text{dcl}_{\mathcal{M}}^G(X)$ of those $b \in \mathcal{M}$ for which

$$\{g(b) : g \in \text{Aut}(\mathcal{M}) \text{ s.t. } (\forall c \in X) g(c) = c\} = \{b\}.$$

The structure \mathcal{M} is said to have trivial group-theoretic definable closure when $\text{dcl}_{\mathcal{M}}^G(X) = X$ for all finite $X \subseteq \mathcal{M}$. Indeed, this is the notion of definable closure used in [AFP16], which only considers the case of Scott sentences.

Definition 2.18. *Let A be a fragment and let \mathcal{M} be an L -structure. The structure \mathcal{M} has **trivial A -definable closure** when $\text{dcl}_A(X) = X$ for all $X \subseteq \mathcal{M}$. Note that \mathcal{M} already has trivial A -definable closure if $\text{dcl}_A(X) = X$ for all finite $X \subseteq \mathcal{M}$. Also note that the analogously defined notion of trivial A -algebraic closure is identical to trivial A -definable closure.*

A complete A -theory T has **trivial A -definable closure** when all models $\mathcal{M} \models T$ have trivial A -definable closure.

One can also perform a formula-by-formula analysis of trivial A -definable closure.

Definition 2.19. *Let A be a countable fragment of $\mathcal{L}_{\omega_1, \omega}(L)$. A complete A -theory T has trivial $\varphi(\bar{x}, y)$ -definable closure for a formula $\varphi(\bar{x}, y)$ of $\mathcal{L}_{\omega_1, \omega}(L)$ if $T \models \neg(\exists \bar{x})(\exists^=1y) \varphi(\bar{x}, y)$.*

The following lemma is immediate.

Lemma 2.20. *Let A be a countable fragment of $\mathcal{L}_{\omega_1, \omega}(L)$. A complete A -theory T has trivial A -definable closure if and only if T has trivial ψ -definable closure for all non-redundant formulas $\psi \in A$.*

It was shown in [AFP16] that \mathcal{M} has trivial group-theoretic definable closure if and only if there is an ergodic structure concentrated on its orbit. When we move to A -theories it is not the case that having an ergodic model of the theory ensures that all (classical) models of the theory have trivial group-theoretic definable closure. In fact, [AFNP16] showed that there are many first-order theories on which some invariant measure is concentrated, but for which almost every (classical) structure sampled according to the measure has non-trivial group-theoretic definable closure. For further examples of this phenomenon, see [AFKrP17].

Lemma 2.21. *Let A be a countable fragment of $\mathcal{L}_{\omega_1, \omega}(L)$, and let T be a complete A -theory. Then T has trivial A -definable closure if and only if T_A^+ has trivial A^+ -definable closure.*

Proof. For every sentence $\sigma \in A^+$, there is a relation symbol $R \in L_A$ such that

$$T_A^+ \models \sigma \leftrightarrow R.$$

However, by Corollary 2.15, for every formula $\varphi \in A$, the theory T has trivial φ -definable closure if and only if T_A^+ has trivial R_φ -definable closure.

In particular, T has trivial φ -definable closure for all formulas $\varphi \in A$ if and only if T_A^+ has trivial ψ -definable closure for all formulas $\psi \in A^+$. Therefore, by Lemma 2.20, T has trivial A -definable closure if and only if T_A^+ has trivial A^+ -definable closure. \square

Rather than working with the notion of trivial A -definable closure directly, it will sometimes be convenient to use the equivalent property of having *duplication of formulas* in A , which is closely related to a notion that first appeared in [AFP16, §3.3]. Recall the notion of a non-redundant formula (Definition 2.6).

Definition 2.22. *Let T be a theory and $\varphi(\bar{x}, y) \in \mathcal{L}_{\omega_1, \omega}(L)$ be a non-redundant formula consistent with T . The theory T **duplicates** φ when*

$$T \models (\exists \bar{x}, y, z) (\varphi(\bar{x}, y) \wedge \varphi(\bar{x}, z) \wedge (y \neq z)).$$

*The theory T has **duplication of quantifier-free formulas** when it duplicates every non-redundant quantifier-free first-order formula $\varphi(\bar{x}, y) \in \mathcal{L}_{\omega, \omega}(L)$ consistent with T .*

*Let A be a fragment and suppose that T is an A -theory. The theory T has **duplication of formulas in A** when T duplicates every non-redundant formula in A consistent with T .*

We then have the following.

Lemma 2.23. *Let A be a fragment, and suppose that T has duplication of formulas in A . For any non-redundant formula $\varphi(x_1, \dots, x_k) \in A$ such that $(\exists x_1, \dots, x_k) \varphi(x_1, \dots, x_k)$ is consistent with T and any $\ell \leq k$, there is a non-redundant formula $\psi(x_1^0, x_1^1, \dots, x_\ell^0, x_\ell^1)$ such that*

- $(\exists x_1^0, x_1^1, \dots, x_\ell^0, x_\ell^1, x_{\ell+1}, \dots, x_k) \psi(x_1^0, x_1^1, \dots, x_\ell^0, x_\ell^1, x_{\ell+1}, \dots, x_k)$ is consistent with T .
- For all maps $\alpha: \{1, \dots, \ell\} \rightarrow \{0, 1\}$,

$$T \models (\forall x_1^0, x_1^1, \dots, x_\ell^0, x_\ell^1, x_{\ell+1}, \dots, x_k) (\psi(x_1^0, x_1^1, \dots, x_\ell^0, x_\ell^1, x_{\ell+1}, \dots, x_k) \rightarrow \varphi(x_1^{\alpha(1)}, \dots, x_\ell^{\alpha(\ell)}, x_{\ell+1}, \dots, x_k)).$$

Proof. The case where $\ell = 1$ is immediate from the definition of duplication of formulas. For $\ell \geq 1$, the result follows from an easy induction. \square

In our main construction, where we assume trivial definable closure, we will actually use the (potentially weaker) notion of duplication of formulas.

Lemma 2.24. *Let A be a fragment and let T be a complete A -theory. Suppose that T has trivial A -definable closure. Then T has duplication of formulas in A .*

Proof. Let $\varphi(\bar{x}, y)$ be an A -formula consistent with T . Then $T \models (\exists \bar{x}, y) \varphi(\bar{x}, y)$. But T has trivial A -definable closure, and so by Lemma 2.20 we have $T \models \neg(\exists \bar{x})(\exists^{\neq 1} y) \varphi(\bar{x}, y)$. Therefore $T \models (\exists \bar{x}, y, z) (\varphi(\bar{x}, y) \wedge \varphi(\bar{x}, z) \wedge (y \neq z))$. \square

Note that the other direction of Lemma 2.24 does not necessarily hold.

Lemma 2.25. *Let T be a complete A -theory, and suppose that T has non-trivial A -definable closure. Then T has non-trivial $\varphi(\bar{x}, y)$ -definable closure for some non-redundant $\varphi(\bar{x}, y) \in A$, i.e.,*

$$T \models (\exists \bar{x})(\exists^{\neq 1} y) \varphi(\bar{x}, y).$$

Proof. Because T is a complete A -theory, if T has non-trivial $\bigvee_{i \leq n} \psi_i$ -definable closure for some $\bigvee_{i \leq n} \psi_i \in A$, then T has non-trivial ψ_i -definable closure for some $i \leq n$. But every formula in A is equivalent to a disjunct of finitely many non-redundant formulas in A , and so the result follows. \square

3. ERGODIC STRUCTURES VIA SAMPLING FROM BOREL STRUCTURES

In this section we describe a general framework for constructing ergodic structures by sampling from continuum-sized structures, which dates back to work of Aldous [Ald81] and Hoover [Hoo79], and which has been used more recently in [PV10], [AFP16], [AFNP16], [AFKwP17], and [AFKrP17].

We begin by defining Borel L -structures, and then describe how to obtain an ergodic structure via a sampling procedure from a Borel L -structure. We then describe weighted homomorphism densities, which will allow us to show convergence of finite sampled substructures to the desired measure on countably infinite structures.

3.1. Borel L -structures. When considering whether a given theory has an ergodic model, without loss of generality we may restrict to the case of relational languages (by Corollary 2.14, as T_A^+ is in a relational language).

Throughout this subsection, L will denote a countable relational language.

Definition 3.1. *Let \mathcal{P} be an L -structure. We say that \mathcal{P} is a **Borel L -structure** if there is a Borel σ -algebra on its underlying set $\mathcal{B}_{\mathcal{P}}$ such that for all relation symbols $R \in L$, the set $\{\bar{a} \in \mathcal{P}^j : R^{\mathcal{P}}(\bar{a})\}$ is a Borel subset of $\mathcal{B}_{\mathcal{P}}^j$, where j is the arity of R .*

We next describe a map taking an element of \mathcal{N}^ω to an L -structure with underlying set \mathbb{N} . The application of this map to an appropriate random sequence of elements of a Borel L -structure will induce a random L -structure with underlying set \mathbb{N} .

Definition 3.2. *Let \mathcal{N} be an L -structure (of arbitrary cardinality). Define the function $\mathcal{F}_{\mathcal{N}}: \mathcal{N}^\omega \rightarrow \text{Str}_L$ as follows. For $\mathcal{A} = (a_i)_{i \in \omega} \in \mathcal{N}^\omega$, let $\mathcal{F}_{\mathcal{N}}(\mathcal{A})$ be the*

L -structure with underlying set \mathbb{N} satisfying

$$\mathcal{F}_{\mathcal{N}}(\mathcal{A}) \models R(n_1, \dots, n_j) \iff \mathcal{N} \models R(a_{n_1}, \dots, a_{n_j})$$

for every relation symbol $R \in L$ and for all $n_1, \dots, n_j \in \mathbb{N}$, where j is the arity of R .

Note that we consider equality as a logical symbol, not a relation symbol (so that equality is inherited from the underlying set).

Observe that when \mathcal{P} is a Borel L -structure, $\mathcal{F}_{\mathcal{P}}$ is a Borel measurable function (see [AFP16, Lemma 3.3]).

Definition 3.3. Let \mathcal{P} be a Borel L -structure, and let m be a probability measure on $\mathcal{B}_{\mathcal{P}}$. The measure $\mu_{(\mathcal{P}, m)}$ on Str_L is defined to be $m^\infty \circ \mathcal{F}_{\mathcal{P}}^{-1}$, where m^∞ is the product measure on $\mathcal{B}_{\mathcal{P}}^\omega$.

Note that m^∞ is invariant under arbitrary reordering of the indices. We will obtain an invariant measure on Str_L by taking the distribution of the *random* structure with underlying set \mathbb{N} corresponding to an m -i.i.d. sequence of elements of \mathcal{P} .

Note that $\mu_{(\mathcal{P}, m)}$ is a probability measure, namely the distribution of a random element in Str_L induced via $\mathcal{F}_{\mathcal{P}}$ by an m -i.i.d. sequence on $\mathcal{B}_{\mathcal{P}}$. The invariance of m^∞ under the action of S_∞ on $\mathcal{B}_{\mathcal{P}}^\omega$ yields the invariance of $\mu_{(\mathcal{P}, m)}$ under the logic action. In particular, $\mu_{(\mathcal{P}, m)}$ is a S_∞ -invariant measure.

Lemma 3.4 ([AFP16, Lemma 3.5]). Let \mathcal{P} be a Borel L -structure, and let m be a probability measure on $\mathcal{B}_{\mathcal{P}}$. Then $\mu_{(\mathcal{P}, m)}$ is S_∞ -invariant.

Recall that a measure is said to be **continuous** (or *atomless*) if it assigns measure zero to every singleton. The following lemma describes several key properties of the invariant measures obtained by sampling using continuous measures.

Lemma 3.5. Let \mathcal{P} be a Borel L -structure, and let m be a continuous probability measure on \mathcal{B} . Then $\mu_{(\mathcal{P}, m)}$ is an ergodic structure that is concentrated on the union of isomorphism classes of countably infinite substructures of \mathcal{P} .

Proof. Let $\mathcal{A} = (a_i)_{i \in \omega}$ be an m -i.i.d. sequence of elements of \mathcal{B} . Note that the induced countable structure $\mathcal{F}_{\mathcal{P}}(\mathcal{A})$ is now a *random L -structure*, i.e., an Str_L -valued random variable, whose distribution is $\mu_{(\mathcal{P}, m)}$. Because m is continuous, and since for any $k \neq \ell$ the random variables a_k and a_ℓ are independent, the sequence \mathcal{A} has no repeated entries, almost surely. Hence $\mathcal{F}_{\mathcal{P}}(\mathcal{A})$ is almost surely isomorphic to a countably infinite (induced) substructure of \mathcal{P} .

The measure $\mu_{(\mathcal{P}, m)}$ is ergodic, as shown in [AFKwP17, Proposition 2.24]. \square

The fact that $\mu_{(\mathcal{P}, m)}$ is concentrated on substructures of \mathcal{P} (up to isomorphism) will be extended in Lemma 4.9, where we show how to make it almost surely satisfy a given theory.

3.2. Weighted homomorphism densities. Weighted homomorphism densities will allow us in Lemma 4.4 to show that the weak limit of finite sampled substructures of a Borel L -structure achieves the desired invariant measure on countably infinite structures.

Definition 3.6. *Suppose $L' \subseteq L$, and let \mathcal{M} be an L' -structure and \mathcal{N} an L -structure. Define a map f from the underlying set of \mathcal{M} to the underlying set of \mathcal{N} to be a **homomorphism** if it is a homomorphism from \mathcal{M} to $\mathcal{N}|_{L'}$ (the L' -reduct of \mathcal{N}), and a **full homomorphism** if it also preserves all non-relations of L' . Write $\text{Hom}(\mathcal{M}, \mathcal{N})$ and $\text{Full}(\mathcal{M}, \mathcal{N})$ to denote these respective classes of maps.*

Definition 3.7. *Suppose $L' \subseteq L$, with L' finite, and let \mathcal{M} be a finite L' -structure and \mathcal{N} an arbitrary L -structure. Let m be a probability measure on the underlying set of \mathcal{N} . Define the **weighted full homomorphism density** $t_{\text{full}}(\mathcal{M}, (\mathcal{N}, m))$ to be the probability that assigning the elements of \mathcal{M} to elements of \mathcal{N} in an m -i.i.d. way yields a full homomorphism, i.e., a map in which relations and non-relations are preserved. Specifically,*

$$t_{\text{full}}(\mathcal{M}, (\mathcal{N}, m)) := \int_{\text{Full}(\mathcal{M}, \mathcal{N})} dm^{|\mathcal{M}|},$$

which reduces to

$$t_{\text{full}}(\mathcal{M}, (\mathcal{N}, m)) = \sum_{f \in \text{Full}(\mathcal{M}, \mathcal{N})} \prod_{a \in \mathcal{M}} m(f(a))$$

when \mathcal{N} is countable.

This notion extends the case where \mathcal{M} is a simple graph and \mathcal{N} is an edge-weighted graph, as in, e.g., [Lov12, §5.2.1]. For more details on full homomorphism densities, see [AFNP16, §3.1].

Proposition 3.8. *Let \mathcal{N} be a finite L -structure and m a probability measure on the underlying set of \mathcal{N} . Then $\mu_{(\mathcal{N}, m)}$ is completely determined by the sequence of numbers*

$$\{t_{\text{full}}(\mathcal{M}, (\mathcal{N}, m))\}_{\mathcal{M}},$$

where \mathcal{M} ranges over finite L' -structures, where $L' \subseteq L$ and L' is finite.

Proof. For any finite $L' \subseteq L$ and any L' structure \mathcal{M} with underlying set $\{0, \dots, |\mathcal{M}| - 1\}$, let $p_{\mathcal{M}}$ be the qf-type of $\{0, \dots, |\mathcal{M}| - 1\}$ in \mathcal{M} . For any finite $L' \subseteq L$ and quantifier-free L' -formula φ with $|\mathcal{M}|$ -many free variables, let $\nu_{L'}(\llbracket \varphi(n_0, \dots, n_{|\mathcal{M}|-1}) \rrbracket)$ be the quantity

$$\sum \{t_{\text{full}}(\mathcal{M}, (\mathcal{N}, m)) : \mathcal{M} \text{ is an } L'\text{-structure and } \models p_{\mathcal{M}} \rightarrow \varphi\}.$$

Let L' and L'' be finite languages such that $L' \subseteq L'' \subseteq L$. Let ψ be a quantifier-free L' -formula; in particular, ψ is also an L'' -formula. Note that

$\nu_{L''}(\llbracket \psi(n_0, \dots, n_{|\mathcal{M}|-1}) \rrbracket)$ is the probability that an m -i.i.d. sequence of elements of \mathcal{N} of length $|\mathcal{M}|$ satisfies ψ . Hence $\nu_{L''}(\llbracket \psi(n_0, \dots, n_{|\mathcal{M}|-1}) \rrbracket)$ does not depend on the choice of L'' extending L' , and so we may define $\nu(\llbracket \psi(n_0, \dots, n_{|\mathcal{M}|-1}) \rrbracket)$ to be $\nu_{L'}(\llbracket \psi(n_0, \dots, n_{|\mathcal{M}|-1}) \rrbracket)$.

Define Θ to be the collection of $\llbracket \varphi(n_0, \dots, n_{r-1}) \rrbracket$ where $n_0, \dots, n_{r-1}, r \in \mathbb{N}$ and φ is a quantifier-free L -formula with r -many free variables. Note that Θ is a ring and ν is a pre-measure on Θ that agrees with $\mu_{(\mathcal{N}, m)}$. By the Carathéodory extension theorem, the unique extension of ν to a measure must be $\mu_{(\mathcal{N}, m)}$. \square

For similar results in the special case of graphs, see [Lov12, Chapter 11]. For related results and discussion, see also [Kru16, §1.1–1.2].

4. LAYERINGS

Now that we have a method for producing an ergodic structure from a Borel L -structure with an associated measure, we need a way to construct Borel L -structures such that the resulting ergodic structures almost surely satisfy the desired theory. Our general method is to build a Borel L -structure as the limit of a directed system of finite structures, as we describe in this section. This construction is similar in spirit to the inverse limit construction of [AFNP16].

Throughout this section, we restrict L to be a countable *relational* language.

4.1. Basic layerings. Our methods rely on a notion called a *layering*, which we introduce here.

Definition 4.1. *Suppose $\mathcal{P}_0, \mathcal{P}_1$ are Borel L -structures on underlying Borel spaces $\mathcal{B}_0, \mathcal{B}_1$ respectively. We say that a map $f: \mathcal{P}_0 \rightarrow \mathcal{P}_1$ is a **Borel homomorphism** when f is a homomorphism of L -structures and is a Borel function.*

We will be interested in a specific kind of Borel L -structure which arises as the limit of a sequence of Borel homomorphisms.

By *countable Borel L -structure*, we mean a countable (finite or infinite) L -structure with the discrete σ -algebra (i.e., that generated by all singletons). We define the notion of an L -layering in terms of a directed system of countable Borel L -structures, but in the L -layerings that we build in Section 5, these Borel L -structures will always be finite.

Definition 4.2. *We define an **L -layering** (\mathcal{N}, m, i) to be a directed system indexed by $(\omega, <)$ satisfying the following, for each $n \in \omega$.*

- *The structure \mathcal{N}_n is a countable Borel L -structure, whose underlying Borel set we denote by N_n .*
- *The measure m_n is a probability measure on N_n whose only measure zero set is \emptyset .*
- *For each $k \leq n$, the map $i_{n,k}: \mathcal{N}_n \rightarrow \mathcal{N}_k$ is a measure-preserving surjective homomorphism. In particular we have $m_n(i_{n,k}^{-1}(a)) = m_k(a)$ for every $a \in N_k$.*

For $n \in \omega$, we refer to (\mathcal{N}_n, m_n) as **level** n of (\mathcal{N}, m, i) .

We now describe the limit of a layering.

Definition 4.3. Suppose $\widehat{\mathcal{N}} = (\mathcal{N}, m, i)$ is an L -layering. We define the **limit** of $\widehat{\mathcal{N}}$ to be the pair $(\mathcal{N}_\omega, m_\omega)$ satisfying the following.

- The L -structure \mathcal{N}_ω has underlying set

$$\mathcal{N}_\omega := \left\{ \{a_k\}_{k \in \omega} : (\forall n \in \omega) \left(a_n \in \mathcal{N}_n \wedge (\forall k \leq n) (i_{n,k}(a_n) = a_k) \right) \right\}.$$

- For $n \in \omega$, the map $i_{\omega,n}: \mathcal{N}_\omega \rightarrow \mathcal{N}_n$ sends the tuple $\{a_k\}_{k \in \omega}$ to a_n .
- The Borel structure on \mathcal{N}_ω is that generated by sets of the form $i_{\omega,n}^{-1}(b)$, where $b \in \mathcal{N}_n$ and $n \in \omega$.
- For any relation symbol $R \in L$,

$$\mathcal{N}_\omega \models \neg R(\{a_k^1\}_{k \in \omega}, \dots, \{a_k^j\}_{k \in \omega})$$

if and only if

$$\mathcal{N}_n \models \neg R(a_n^1, \dots, a_n^j)$$

for some $n \in \omega$, where j is the arity of R .

- The measure m_ω is the unique probability measure satisfying

$$m_\omega(\{\{a_k\}_{k \in \omega} : a_j = b\}) = m_j(b)$$

for all $j \in \omega$ and $b \in \mathcal{N}_j$.

It is easy to check that this is well-defined; in particular, by the Carathéodory extension theorem the measure is uniquely determined by specifying the measures of the sets of the form $i_{\omega,n}^{-1}(b)$, where $b \in \mathcal{N}_n$ and $n \in \omega$.

Observe that \mathcal{N}_ω is the limit in the category of Borel L -structures and Borel homomorphisms of the directed system $\widehat{\mathcal{N}}$. The key property of $(\mathcal{N}_\omega, m_\omega)$ is that the ergodic structure $\mu_{(\mathcal{N}_\omega, m_\omega)}$ obtained by sampling from $(\mathcal{N}_\omega, m_\omega)$ as in Lemma 3.4 is the weak limit of the sequence of ergodic structures obtained by sampling from each of the (\mathcal{N}_n, m_n) for $n \in \omega$. We now make this precise. Let \xrightarrow{w} denote convergence in the weak (a.k.a. weak-* or vague) topology.

Lemma 4.4. *The sequence of measures $\{\mu_{(\mathcal{N}_n, m_n)}\}_{n \in \omega}$ converges weakly to $\mu_{(\mathcal{N}_\omega, m_\omega)}$, i.e., we have $\mu_{(\mathcal{N}_n, m_n)} \xrightarrow{w} \mu_{(\mathcal{N}_\omega, m_\omega)}$.*

Proof. Fix a finite sublanguage $L' \subseteq L$ and a finite L' -structure \mathcal{M} . For each $k \in \omega$, consider the weighted full homomorphism density $t_{\text{full}}(\mathcal{M}, (\mathcal{N}_k, m_k))$ of \mathcal{M} in (\mathcal{N}_k, m_k) . Observe that because the map $i_{k+1,k}: \mathcal{N}_{k+1} \rightarrow \mathcal{N}_k$ is a homomorphism, the probability that the random map $\mathcal{M} \rightarrow \mathcal{N}_k$ (that independently assigns each element of \mathcal{M} an element of \mathcal{N}_k according to m_k) is a homomorphism is at least the probability that the corresponding random map $\mathcal{M} \rightarrow \mathcal{N}_{k+1}$ is a homomorphism. Therefore the sequence of positive reals $\{t_{\text{full}}(\mathcal{M}, (\mathcal{N}_k, m_k))\}_{k \in \omega}$ is non-increasing. In particular, the limit $\lim_{k \rightarrow \infty} t_{\text{full}}(\mathcal{M}, (\mathcal{N}_k, m_k))$ exists and is

equal to $t_{\text{full}}(\mathcal{M}, (\mathcal{N}_\omega, m_\omega))$, because $(\mathcal{N}_\omega, m_\omega)$ is the limit of the directed system formed by $\{(\mathcal{N}_k, m_k)\}_{k \in \mathbb{N}}$ and the maps $\{i_{k+1,k}\}_{k \in \mathbb{N}}$ in the category of probability spaces. Hence for any finite $L' \subseteq L$ and any quantifier-free L' -formula φ , and any $n_0, \dots, n_{r-1} \in \mathbb{N}$, where φ has r -many free variables, we have

$$\lim_{k \rightarrow \infty} \mu_{(\mathcal{N}_k, m_k)}(\llbracket \varphi(n_0, \dots, n_{r-1}) \rrbracket) = \mu_{(\mathcal{N}_\omega, m_\omega)}(\llbracket \varphi(n_0, \dots, n_{r-1}) \rrbracket)$$

by Proposition 3.8. Note that every clopen set is of the form $\llbracket \varphi(n_0, \dots, n_{r-1}) \rrbracket$. Therefore by the Portmanteau theorem, we have $\mu_{(\mathcal{N}_n, m_n)} \xrightarrow{w} \mu_{(\mathcal{N}_\omega, m_\omega)}$. \square

4.2. Continuous limits. We will mainly be interested in layerings whose limit measure is a continuous measure on the limit structure, i.e., where each singleton of the limit structure is a nullset of the limit measure. In this case we say that the layering is **continuous**. We now characterize such layerings.

Lemma 4.5. *Suppose $\widehat{\mathcal{N}} = (\mathcal{N}, m, i)$ is an L -layering with limit $(\mathcal{N}_\omega, m_\omega)$. Then the following are equivalent.*

- (a) *The probability measure m_ω is continuous.*
- (b) *For all $k \in \omega$ and $a \in \mathcal{N}_k$, there are $n > k$ and $X, Y \subseteq \mathcal{N}_n$ such that*
 - $X, Y \subseteq i_{n,k}^{-1}(a)$,
 - $X \cap Y = \emptyset$,
 - $m_n(X) > 1/3 \cdot m_k(a)$, and
 - $m_n(Y) > 1/3 \cdot m_k(a)$.

Proof. Suppose m_ω is continuous. Towards a contradiction, assume that $k \in \omega$ and $a \in \mathcal{N}_k$ are such that for all $n > k$ there are no $X, Y \subseteq \mathcal{N}_n$ satisfying the conditions of (b).

If for some $n > k$ every element $b \in \mathcal{N}_n$ satisfied $m_n(b) < 2/3 \cdot m_n(a)$, then we could construct X and Y satisfying the conditions of (b) by adding elements in decreasing mass to the set X until its mass exceeded $1/3 \cdot m_k(a)$ and allowing $Y = \mathcal{N}_n \setminus X$.

Hence for every $n > k$ there is some $b_n \in \mathcal{N}_n$ such that $i_{n,k}(b_n) = a$ and $m_n(b_n) > 2/3 \cdot m_k(a)$; note that is exactly one such b_n . Therefore we have $i_{n',n}(b_n) = b_{n'}$ for $n' > n > k$. In particular, if we let $b_j = i_{k,j}(a)$ for $j \leq k$, then $m_\omega(\{b_j\}_{j \in \omega}) > 2/3 \cdot m_k(a)$, contradicting the continuity of m_ω .

Now suppose (b) holds, and let $\{a_j\}_{j \in \omega} \in \mathcal{N}_\omega$. Then for every $k \in \omega$ there is some $n > k$ such that $m_n(a_n) \leq 2/3 \cdot m_k(a_k)$. Hence $m_\omega(\{a_j\}_{j \in \omega}) = \inf_j m_j(a_j)$, which is at most $(\frac{2}{3})^{-q} \cdot m_0(a_0)$ for all $q \in \omega$, and so $m_\omega(\{a_j\}_{j \in \omega}) = 0$. \square

4.3. Regular layerings. We will need two other specific conditions on a layering to allow us sufficient control over the limit structure and measure of the layering.

First, we want to be able to define our layerings without specifying where all the mass is assigned at any given level; this will allow us at later levels to *add new elements*. We will accomplish this via a distinguished “virtual” element whose mass at each level consists of that mass not assigned to other “real” elements.

Second, at each level we want to declare that certain first-order quantifier-free formulas hold, and we want such formulas to continue to hold at later levels. Unfortunately, in arbitrary layerings different formulas may hold at different levels, as we required only that the maps between levels were homomorphisms, and not necessarily embeddings. We will address this by assigning each level a language such that the maps between a level and all higher ones preserve the qf-type (in the language of the level) of every tuple that is non-redundant and whose image is also non-redundant. In other words, for such tuples, formulas remain *set*, once set at a lower level. Because we want to allow ourselves to split elements as we move up the levels of the layering, we do not require the qf-type of redundant tuples (or those whose image is redundant) to be set.

We now make both of these conditions precise.

Definition 4.6. *A regular L -layering is a layering $\widehat{\mathcal{N}} = (\mathcal{N}, m, i)$ along with the following additional data.*

- *For each $n \in \omega$ there is an element $*_n \in \mathcal{N}_n$, called the **sink** of \mathcal{N}_n , such that*
 - (a) *all relations in L hold of any tuple in \mathcal{N}_n containing $*_n$,*
 - (b) *for all $k \leq n$ we have $i_{n,k}(*_n) = *_k$, and*
 - (c) *for all $\varepsilon > 0$ there is an $n \in \omega$ such that $m_n(*_n) < \varepsilon$.*
- *There is a sequence $\{L_n\}_{n \in \omega}$ of countable relational languages such that for all $n \in \omega$,*
 - (d) $L_n \subseteq L_{n+1}$,
 - (e) $\bigcup_{j \in \omega} L_j = L$,
 - (f) *for all $k \leq n$ and all non-redundant tuples $\bar{a} \in \mathcal{N}_n$, if $i_{n,k}(\bar{a})$ is a non-redundant tuple in \mathcal{N}_k , then the qf L_k -type of \bar{a} is the same as that of $i_{n,k}(\bar{a})$, and*
 - (g) *all tuples $\bar{a} \in \mathcal{N}_n$ and all $R \in L$ of arity $|\bar{a}|$, if $\mathcal{N}_n \models \neg R(\bar{a})$ then $R \in L_n$ and \bar{a} is non-redundant.*

We can think of a level as containing information about tuples in the resulting limit structure. Condition (a) asserts that the sink $*_n$ is a *neutral element* in that it doesn't force any structure at higher levels. Condition (b) is straightforward, while (c) ensures that $m_\omega(*_\omega) = 0$, where $*_\omega := \{*_j\}_{j \in \omega}$. Conditions (d) and (e) say that L is an increasing union of sublanguages L_n , and (f) ensures that if a decision is made about a non-redundant tuple, that decision propagates to higher levels. Condition (g) states that \mathcal{N}_n is *neutral with respect to $L \setminus L_n$* and makes no decisions about redundant tuples.

4.4. Approximations to types. We now introduce a condition on a continuous regular layering (\mathcal{N}, m, i) which will ensure that $\mu_{(\mathcal{N}_\omega, m_\omega)}$ omits a certain kind of type, which can be appropriately approximated in sublanguages. We will also see that satisfying a pithy Π_2 sentence and omitting a qf-type can be expressed as omitting a type of this kind.

Definition 4.7. Let L be a countable relational language and let $\{L_j\}_{j \in \omega}$ be a nondecreasing sequence of languages whose union is L . Let \hat{p} be a sequence $\{p_j\}_{j \in \omega}$ of types such that each p_j is a partial L_j -type in the same variables, and such that $\models p_k \rightarrow p_j$ for $j < k$. Then \hat{p} is an **approximable** $\{L_j\}_{j \in \omega}$ -**type** if for all $j < k$, for all L_k -structures \mathcal{M}_k and L_j -structures \mathcal{M}_j , and for all surjective embeddings $i_{k,j}$ from \mathcal{M}_k to \mathcal{M}_j , if

- $\bar{a} \in \mathcal{M}_k$ is a tuple satisfying p_k , and
- $i_{k,j}$ is injective on the underlying set of \bar{a} ,

then the tuple $i_{k,j}(\bar{a})$ satisfies p_j .

Elements of the limit of a regular layering correspond to sequences with one element from each level of the layering, where the elements of the higher levels of the sequence project down to the elements at the lower levels of the sequence, via the maps of the layering. As such, a tuple of elements in the limit corresponds to a sequence of tuples from each level, on which the maps of the layering are injective past some point.

The idea behind the notion of an approximable type is to capture information about elements in the limit of a regular layering by considering properties of the corresponding sequence. However, because the approximable type is expressed in a way that is independent of the layering, we need it to cohere with any possible layering it could be asked to describe. This is the motivation behind the coherence condition with respect to arbitrary surjective embeddings.

Definition 4.8. Let $\hat{\mathcal{N}} = (\mathcal{N}, m, i)$ be a continuous regular L -layering, and suppose that $\{p_j\}_{j \in \omega}$ is an approximable $\{L_j\}_{j \in \omega}$ -type. We say that $\hat{\mathcal{N}}$ **asymptotically omits** $\{p_j\}_{j \in \omega}$ if for all $k \in \omega$ there is an $n_k > k$ such that

$$\mathcal{N}_{n_k} \models (\forall \bar{x} \in i_{n_k, k}^{-1}(\mathcal{N}_k \setminus *_{n_k})) \neg p_k(\bar{x}).$$

The following key technical lemma will allow us to ensure that a limit measure satisfies or omits formulas of several particular forms.

Lemma 4.9. Let $\hat{\mathcal{N}} = (\mathcal{N}, m, i)$ be a continuous regular L -layering, and suppose $\{p_j\}_{j \in \omega}$ is an approximable $\{L_j\}_{j \in \omega}$ -type. If $\hat{\mathcal{N}}$ asymptotically omits $\{p_j\}_{j \in \omega}$ then

$$\mu_{(\mathcal{N}_\omega, m_\omega)} \left(\left[\left(\forall \bar{x} \right) \bigvee_{j \in \omega} \neg p_j(\bar{x}) \right] \right) = 1.$$

Proof. Let \bar{x} be the variables of $\{p_j\}_{j \in \omega}$ and let $k = |\bar{x}|$. For $j \in \omega$, define $\alpha_j := m_j(*_j) + (1 - \beta_j)$, where β_j denotes the probability that sampling k elements independently from m_j yields a non-redundant tuple. Note that α_j is an upper bound on the probability that the sink $*_j$ is selected at least once or that some element is selected at least twice. Because $\hat{\mathcal{N}}$ is regular, $\lim_j m_j(*_j) = 0$. Because $\hat{\mathcal{N}}$ is continuous, $\lim_j \beta_j = 1$. Hence $\lim_j \alpha_j = 0$.

Note that if $\ell \geq j$ and \bar{a} is a non-redundant tuple of $\mathcal{N}_\ell \setminus *_\ell$ that satisfies p_j , then if $i_{\ell,j}(\bar{a})$ is also non-redundant it satisfies p_j . Hence, by the asymptotic omission of $\{p_j\}_{j \in \omega}$, if $g \geq n_j$ or $g = \omega$, then the probability that a sample from the k -fold product measure m_g^k satisfies p_j is at most α_j .

Therefore the probability that a sample from the k -fold product measure m_ω^k satisfies $(\exists \bar{x}) \bigwedge_{j \in \omega} p_j(\bar{x})$ is at most $\inf_j \alpha_j = 0$. Hence a sample from m_ω^k almost surely satisfies $(\forall \bar{x}) \bigvee_{j \in \omega} \neg p_j(\bar{x})$, and so

$$\mu_{(\mathcal{N}_\omega, m_\omega)} \left(\left[\left[(\forall \bar{x}) \bigvee_{j \in \omega} \neg p_j(\bar{x}) \right] \right] \right) = 1,$$

as desired. \square

We now consider two special classes of approximable types that can be asymptotically omitted. Asymptotically omitting an approximable type of the first class in a continuous regular layering (\mathcal{N}, m, i) ensures that the limit measure $\mu_{(\mathcal{N}_\omega, m_\omega)}$ satisfies a given pithy Π_2 -theory.

Definition 4.10. Let $\widehat{\mathcal{N}} = (\mathcal{N}, m, i)$ be a continuous regular L -layering. Suppose that $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y) \in \mathcal{L}_{\omega_1, \omega}(L)$ is a pithy Π_2 sentence such that there exists some $\ell \in \omega$ for which $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y) \in \mathcal{L}_{\omega_1, \omega}(L_\ell)$. Consider the approximable type $\{p_j\}_{j \in \omega}$ where $p_j(\bar{x}) = (\bar{x} = \bar{x})$ for $j < \ell$ and $p_j(\bar{x}) = (\forall y)\neg\varphi(\bar{x}, y)$ for $j \geq \ell$. We say that $\widehat{\mathcal{N}}$ **asymptotically satisfies** $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$ when $\widehat{\mathcal{N}}$ asymptotically omits $\{p_j\}_{j \in \omega}$. (Note that this does not depend on the choice of ℓ .)

We say that $\widehat{\mathcal{N}}$ asymptotically satisfies a pithy Π_2 -theory T of $\mathcal{L}_{\omega_1, \omega}(L)$ when it asymptotically satisfies every sentence of T .

The following corollary is immediate from Lemma 4.9.

Corollary 4.11. Let $\widehat{\mathcal{N}} = (\mathcal{N}, m, i)$ be a continuous regular L -layering, and suppose T is a pithy Π_2 -theory of $\mathcal{L}_{\omega_1, \omega}(L)$. If $\widehat{\mathcal{N}}$ asymptotically satisfies T then $\mu_{(\mathcal{N}_\omega, m_\omega)}(\llbracket T \rrbracket) = 1$, i.e., $\mu_{(\mathcal{N}_\omega, m_\omega)}$ is concentrated on T .

Asymptotically omitting an approximable type of the second class results in a limit measure $\mu_{(\mathcal{N}_\omega, m_\omega)}$ that omits a given qf L -type.

Definition 4.12. Let $\widehat{\mathcal{N}} = (\mathcal{N}, m, i)$ be a continuous regular L -layering, and let p be a qf L -type. We say that $\widehat{\mathcal{N}}$ **asymptotically omits** p when $\widehat{\mathcal{N}}$ asymptotically omits the approximable type $\left\{ \bigwedge (p \cap \mathcal{L}_{\omega, \omega}(L_j)) \right\}_{j \in \omega}$.

Again the following corollary is immediate from Lemma 4.9.

Corollary 4.13. Let $\widehat{\mathcal{N}} = (\mathcal{N}, m, i)$ be a continuous regular L -layering, and let $p(\bar{x})$ be a qf L -type. If $\widehat{\mathcal{N}}$ asymptotically omits p then $\mu_{(\mathcal{N}_\omega, m_\omega)}(\llbracket (\forall \bar{x})\neg p(\bar{x}) \rrbracket) = 1$, i.e., $\mu_{(\mathcal{N}_\omega, m_\omega)}$ omits the singleton set of types $\{p\}$.

4.5. Layer Transformation. We have reduced the problem of constructing an ergodic structure almost surely satisfying a theory to that of constructing a regular continuous layering which asymptotically omits a countable collection of types and asymptotically satisfies a pithy Π_2 -theory. Now we give a general method for constructing such layerings.

We do this by isolating a relationship between successive levels of a layering, such that when it occurs infinitely often along pairs of consecutive levels of the layering, the limit measure has the desired properties.

An *approximation* consists of the data needed to be a level of a regular L -layering.

Definition 4.14. *Let \mathcal{M} be a non-redundant L -structure. An **approximation** to \mathcal{M} is a triple $(L_+, \mathcal{N}_+, m_+)$ such that*

- $L_+ \subseteq L$;
- the underlying set N_+ of \mathcal{N}_+ satisfies $N_+ = M_+ \cup *_+$, where $M_+ \subseteq M$, the underlying set of \mathcal{M} ;
- $\mathcal{M} \models R(\bar{a})$ if and only if $\mathcal{N}_+ \models R(\bar{a})$ for any tuple $\bar{a} \in M_+$ and any relation symbol $R \in L_+$ of arity $|\bar{a}|$;
- all relations in L hold of any tuple in \mathcal{N}_+ containing $*_n$;
- $\mathcal{N}_+ \models R(\bar{a})$ for all tuples $\bar{a} \in \mathcal{N}_+$ and all $R \in L \setminus L_+$ of arity $|\bar{a}|$; and
- m_+ is a measure on N_+ .

A layer transformation is something which takes an approximation to \mathcal{M} , i.e., something which could be a level of a regular layering, and returns a new approximation along with a map to the old one.

Definition 4.15. *Let \mathcal{M} be a non-redundant L -structure. A **layer transformation** is a function f which takes an approximation $(L_+, \mathcal{N}_+, m_+)$ of \mathcal{M} to a pair $(\iota_+, (L_\dagger, \mathcal{N}_\dagger, m_\dagger))$ such that*

- $(L_\dagger, \mathcal{N}_\dagger, m_\dagger)$ is an approximation to \mathcal{M} ;
- \mathcal{N}_+ has sink $*_+$ and \mathcal{N}_\dagger has sink $*_\dagger$;
- $L_+ \subseteq L_\dagger$;
- $\iota_+: \mathcal{N}_\dagger \rightarrow \mathcal{N}_+$ is a surjective homomorphism and $i_\dagger(*_\dagger) = *_+$;
- for all $\bar{a} \in \mathcal{N}_\dagger$, if $i_\dagger(\bar{a})$ is a non-redundant tuple in \mathcal{N}_+ , then the qf L_+ -type of \bar{a} is the same as that of $i_\dagger(\bar{a})$; and
- for all $a \in \mathcal{N}_+$, we have $m_\dagger(i_\dagger^{-1}(\{a\})) = m_+(\{a\})$.

We will show that properties of layer transformation transfer to properties of limit measures of continuous regular layerings which have the layer transformations as consecutive pairs of layers.

Definition 4.16. *We say that a layer transformation **forces** a property P if whenever (\mathcal{N}, m, i) is a L -layering such that for infinitely many $n \in \mathbb{N}$, the pair $(i_{n+1,n}, (L_{n+1}, \mathcal{N}_{n+1}, m_{n+1}))$ is the image of $(L_n, \mathcal{N}_n, m_n)$ under the transformation, then (\mathcal{N}, m, i) has property P .*

We say that a layer transformation **continuously forces** a property P if whenever (\mathcal{N}, m, i) is a continuous regular L -layering such that for infinitely many $n \in \mathbb{N}$, the pair $(i_{n+1, n}, (L_{n+1}, \mathcal{N}_{n+1}, m_{n+1}))$ is the image of $(L_n, \mathcal{N}_n, m_n)$ under the transformation, then (\mathcal{N}, m, i) has property P .

Let \mathcal{M} be a non-redundant L -structure. We now introduce properties of layer transformations in terms of \mathcal{M} which, when interleaved in a layering, will force the layering to be continuous, regular, and concentrated on a specified theory.

We first consider two properties which force a layering to be continuous and regular.

The *scrapwork transformation* simply adds a new element to the structure and gives it half the mass of the sink of the previous layer.

Definition 4.17. We define a **scrapwork transformation** to be any layer transformation which takes an approximation $(L_+, \mathcal{N}_+, m_+)$ and returns a pair $(\iota_+, (L_+, \mathcal{N}_+, m_+))$ such that

- ι_+ is the identity on $N_+ \setminus *_+$ and takes the element a to $*_+$;
- $L_+ = L_+$;
- $N_+ = N_+ \cup \{a\}$ for some $a \in \mathcal{M}$; and
- $m_+(\{b\}) = m_+(\{b\})$ for any $b \in N_+ \setminus *_+$ and $m_+(a) = m_+(*_+) = \frac{1}{2}m_+(*_+)$.

The following is then immediate.

Lemma 4.18. Any scrapwork transformation forces $m(*_\omega) = 0$ and hence forces a layering to be regular.

Proof. First note that $m_+(*_+) = \frac{1}{2}m_+(*_+)$, and so each application of the scrapwork transformation reduces then measure of the sink by one half. In particular, if there are infinitely many such applications, then $m_\omega(*_\omega) = 0$. \square

A *splitting transformation* splits each element into two new elements that each stand in the same relationship to all other elements as the original element did.

Definition 4.19. Suppose that for every quantifier-free formula ψ , the $\mathcal{L}_{\omega_1, \omega}(L)$ -theory of \mathcal{M} has trivial ψ -definable closure. Fix an enumeration b_1, \dots, b_k of $N_+ \setminus *_+$ and let q be the quantifier free L_+ -type of b_1, \dots, b_k . By Lemma 2.23 there are tuples $a_1^0 \cdots a_k^0$ and $a_1^1 \cdots a_k^1$ containing $2k$ -many distinct elements of \mathcal{M} such that $\mathcal{M} \models q(a_1^{\alpha(1)}, \dots, a_k^{\alpha(k)})$ for any function $\alpha: \{1, \dots, k\} \rightarrow \{0, 1\}$.

We define a **splitting transformation** to be any layer transformation which takes an approximation $(L_+, \mathcal{N}_+, m_+)$ and returns a pair $(\iota_+, (L_+, \mathcal{N}_+, m_+))$ such that

- $i_+(a_\ell^j) = b_\ell$;
- $L_+ = L_+$;
- $N_+ = \{a_\ell^j : j \in \{0, 1\} \text{ and } 1 \leq \ell \leq k\} \cup *_+$; and
- $m_+(a_\ell^j) = m_+(b_\ell)/2$ for $j \in \{0, 1\}$ and $1 \leq \ell \leq k$ and let $m_+(*_+) = m_+(*_+)$

for some choice of elements a_j^i and b_j as above.

Lemma 4.20. *Any splitting transformation forces the measure m_ω to be continuous outside of the sink $*_\omega$.*

Proof. Let $X_n = \sup\{m_n(\{a\}) : a \in \mathcal{N}_n \setminus *_{n}\}$. If the relationship between levels n and $n + 1$ is a splitting transformation, then $X_{n+1} = \frac{1}{2}X_n$. Therefore $\lim_{n \rightarrow \infty} X_n = 0$, and so the splitting transformation forces m_ω to be continuous on $\mathcal{N}_\omega \setminus *_\omega$. \square

Sampling from any structure with respect to a measure having a point mass will yield a random structure that almost surely has an indiscernible set. As such, if we hope to find a measure concentrated on structures or theories without an indiscernible set, then we need to be able to sample continuous measures — which is why we need our theory to have trivial ψ -definable closure for some ψ .

In particular, if a layering has cofinally many splitting and scrapwork transformations, then it must be a continuous regular transformation.

We now describe a layering transformation that will ensure that the limit measure satisfies a given pithy Π_2 sentence.

Definition 4.21. *Let $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$ be a pithy Π_2 sentence, and suppose $\mathcal{M} \models (\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$. Define a $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$ -**satisfaction transformation** to be any layer transformation which takes an approximation $(L_+, \mathcal{N}_+, m_+)$ and returns $(\iota_\dagger, (L_\dagger, \mathcal{N}_\dagger, m_\dagger))$ such that*

- ι_\dagger is the identity on $N_+ \setminus *_+$ and takes $N_\dagger \setminus N_+$ to $*_\dagger$;
- $L_\dagger = L_+$;
- \mathcal{N}_\dagger is such that for every $\bar{a} \in \mathcal{N}_+ \setminus *_+$ there is some $b \in \mathcal{N}_\dagger$ such that $\mathcal{M} \models \varphi(\bar{a}, b)$; and
- $m_\dagger(b) = m_+(b)$ for any $b \in N_+ \setminus *_+$, and $m_\dagger(a) > 0$ otherwise.

Note that we can always find such an \mathcal{N}_\dagger as $\mathcal{M} \models (\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$.

Lemma 4.22. *Any $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$ -satisfaction transformation continuously forces the layering to asymptotically satisfy $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$.*

Proof. Suppose (\mathcal{N}, m, i) is a continuous regular layering and that $\bar{a} \in \mathcal{N}_\omega$. Let n be such that $i_{\omega, n}(a_i) \neq i_{\omega, n}(a_j)$ for distinct $a_i, a_j \in \bar{a}$. Let $n' > n$ be a stage at which a $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$ -satisfaction transformation occurs. There is therefore some $b \in \mathcal{N}_{n'} \setminus *_{n'}$ such that $\mathcal{N}_{n'} \models \varphi(i_{\omega, n'}(a_0), \dots, i_{\omega, n'}(a_{k-1}), b)$.

Let $b_\omega \in \mathcal{N}_\omega$ be such that $i_{\omega, n'}(b_\omega) = b$. Then $\mathcal{N}_\omega \models \varphi(\bar{a}, b_\omega)$ and hence $\mathcal{N}_\omega \models (\exists y)\varphi(\bar{a}, y)$. But as \bar{a} was arbitrary, we have $\mathcal{N}_\omega \models (\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$. \square

We now describe the final layering transformation, which will ensure that the limit measure almost surely omits a given type.

Definition 4.23. *Let $(\forall \bar{x})(\exists y)\varphi(\bar{x}, y)$ be a pithy Π_2 sentence, and suppose $p(\bar{x})$ is a qf -type omitted in \mathcal{M} . Define a **omitting- $p(\bar{x})$ transformation** to be any*

layer transformation which takes an approximation $(L_+, \mathcal{N}_+, m_+)$ and returns $(\iota_\dagger, (L_\dagger, \mathcal{N}_\dagger, m_\dagger))$ such that

- $\iota_\dagger = \text{id}$;
- L_\dagger is a language such that for all $\bar{a} \in \mathcal{N}_\dagger$ of the same arity as p , there is a qf-free L_\dagger -formula $\eta(\bar{x})$ such that $\neg\eta(\bar{x}) \in p(\bar{x})$ and $\mathcal{M} \models \eta(\bar{a})$;
- $N_\dagger = N_+$
- $m_\dagger = m_+$.

Note that we can always find such a language, since \mathcal{M} omits $p(\bar{x})$.

Lemma 4.24. *Any omitting- $p(\bar{x})$ transformation continuously forces the layering to asymptotically omit $p(\bar{x})$.*

Proof. Suppose (\mathcal{N}, m, i) is a continuous regular layering and that $\bar{a} \in \mathcal{N}_\omega$. Let n be such that $i_{\omega,n}(a_i) \neq i_{\omega,n}(a_j)$ for distinct $a_i, a_j \in \bar{a}$. Let $n' > n$ be a stage at which the omitting- $p(\bar{x})$ transformation occurs. Hence there is some quantifier-free formula $\eta(\bar{x})$ such that $\neg\eta(\bar{x}) \in p(\bar{x})$ and $\mathcal{N}_{n'+1} \models \eta(i_{\omega,n'}(a_0), \dots, i_{\omega,n'}(a_{k-1}))$. But then $\mathcal{N}_\omega \models \eta(\bar{a})$, and so $\mathcal{N}_\omega \models \neg p(\bar{a})$. However, because \bar{a} was arbitrary, this means that \mathcal{N}_ω omits $p(\bar{x})$. \square

5. EXISTENCE OF ERGODIC STRUCTURES

Having developed appropriate layer transformations, we may now show that for a relational language L , if a complete A -theory T has trivial A -definable closure, then there is an ergodic model of T .

We begin by using duplication of quantifier-free formulas to construct an appropriate layering.

Proposition 5.1. *Let L be a countable relational language, and $\mathcal{M} \in \text{Str}_L$ be a non-redundant L -structure. Further let T be a countable pithy Π_2 theory, and Θ be a countable set of qf-types such that $\mathcal{M} \models T$ and \mathcal{M} omits Θ . Suppose that T has duplication of quantifier-free formulas. Then there is a continuous regular layering (\mathcal{N}, m, i) such that $\mu_{(\mathcal{N}_\omega, m_\omega)}$ is an ergodic model of T that omits Θ .*

Proof. Let $\{Q_j\}_{j \in \omega}$ be an enumeration of the relation symbols of L , and let $\{p_j\}_{j \in \omega}$ be an enumeration of Θ in which each qf-type is enumerated infinitely often. Let $\{(\forall \bar{x})(\exists y)\psi_j(\bar{x}, y)\}_{j \in \omega}$ be an enumeration of T in which each sentence is enumerated infinitely often and such that for every $j \in \omega$, each relation symbol of ψ_j is among $\{Q_k : k < j\}$.

We construct a continuous regular layering (\mathcal{N}, m, i) in stages, building one level at every stage. At each stage n , we will identify a finite L -substructure \mathcal{M}_n of \mathcal{M} (necessarily non-redundant). Let M_n denote the underlying set of \mathcal{M}_n , and choose the sink $*_n$ to not be in the underlying set M of \mathcal{M} . The structure \mathcal{N}_n will be the unique L -structure with underlying set $M_n \cup *_n$ such that

- the sink $*_n$ is a neutral element;

- for any relation symbol R in $L \setminus L_n$ and any tuple \bar{a} of length k , where k is the arity of R , we have $\mathcal{N}_n \models R(\bar{a})$; and
- $\mathcal{M}_n|_{L_n}$ equals the substructure of $\mathcal{N}_n|_{L_n}$ having underlying set M_n .

Stage -1 :

Define $L_{-1} := \emptyset$. Let \mathcal{M}_{-1} be the unique structure on the empty set (which necessarily omits Θ), and let $m_0(*_{-1}) = 1$.

Stage $4n$: Scratch work.

Define $L_{4n} := L_{4n-1}$. Let $a \in M \setminus M_{4n-1}$ and define \mathcal{M}_{4n} to be \mathcal{M} restricted to $M_{4n} \cup \{a\}$.

Define $i_{4n,4n-1}$ to be the identity on M_{4n-1} and such that $i_{4n,4n-1}(a) = *_{4n-1}$. Finally, let m_{4n} be the probability measure that agrees with m_{4n-1} on all subsets of M_{4n-1} and assigns $m_{4n}(a) = m_{4n}(*_{4n}) = \frac{1}{2}m_{4n-1}(*_{4n-1})$.

Stage $4n + 1$: Duplication.

Define $L_{4n+1} := L_{4n}$. Let $k = |M_{4n}|$ and fix an enumeration b_1, \dots, b_k of M_{4n} . Let q be the quantifier-free L_{4n+1} -type of b_1, \dots, b_k . By hypothesis, the theory T has duplication of quantifier-free formulas. Hence by Lemma 2.23 there are tuples $a_1^0 \cdots a_k^0$ and $a_1^1 \cdots a_k^1$ containing $2k$ -many distinct elements of M such that $\mathcal{M} \models q(a_1^{\alpha(1)}, \dots, a_k^{\alpha(k)})$ for any function $\alpha : \{1, \dots, k\} \rightarrow \{0, 1\}$. Define \mathcal{M}_{4n+1} to be the L_A -substructure of \mathcal{M} with underlying set

$$\{a_\ell^j : j \in \{0, 1\} \text{ and } 1 \leq \ell \leq k\}.$$

Finally, let $i_{4n+1,4n}(a_\ell^j) = b_\ell$ and $m_{4n+1}(a_\ell^j) = \frac{1}{2}m_{4n}(b_\ell)$ for $j \in \{0, 1\}$ and $1 \leq \ell \leq k$.

Stage $4n + 2$: Witnesses to Π_2 -statements.

Define $L_{4n+2} := L_{4n+1}$. Let $B \subseteq M$ be a finite set be such that for every $j \leq n$,

$$\mathcal{M} \models (\forall \bar{x} \in M_{4n+1}) (\exists y \in B) \psi_j(\bar{x}, y);$$

such a B exists, as $\mathcal{M} \models T$ and $(\forall \bar{x})(\exists y)\psi_j(\bar{x}, y) \in T$. Define \mathcal{M}_{4n+2} to be \mathcal{M} restricted to have underlying set $M_{4n+1} \cup B$.

Let $i_{4n+2,4n+1}$ be the identity on M_{4n+1} and such that $i_{4n+2,4n+1}(b) = *_{4n+1}$ for all $b \in B \setminus M_{4n+1}$. Let m_{4n+2} be the probability measure that agrees with m_{4n+1} on all subsets of M_{4n+1} and evenly assigns positive mass to $*_{4n+2}$ and to every element of $B \setminus M_{4n+1}$.

Stage $4n + 3$: Omitting qf-types.

Define $\mathcal{M}_{4n+3} := \mathcal{M}_{4n+2}$. Let $L_{4n+3} \supseteq L_{4n+2}$ be a finite language containing $\{Q_j : j \leq n\}$ and such that \mathcal{M}_{4n+3} omits every qf-type in $\{p_j \cap \mathcal{L}_{\omega,\omega}(L_{4n+3}) : j \leq n\}$; this is possible as each qf-type in Θ is quantifier-free and \mathcal{M} omits every qf-type

in Θ . Let $i_{4n+3,4n+2} := \text{id}$ and $m_{4n+3} := m_{4n+2}$.

Finally, as required, for all n , and $k < n$, define $i_{n,k} := i_{k+1,k} \circ i_{k+2,k+1} \circ \dots \circ i_{n,n-1}$. This completes the construction of (N, m, i) .

For every n the map

$$(L_{4n-1}, N_{4n-1}, m_{4n-1}) \mapsto (i_{4n,4n-1}, (L_{4n}, N_{4n}, m_{4n}))$$

is a scrapwork transformation, and so by Lemma 4.18, (\mathcal{N}, m, i) is a regular layering.

By Lemma 3.4, $\mu_{(\mathcal{N}_\omega, m_\omega)}$ is an invariant probability measure, and by Lemma 3.5 it is ergodic.

For every n , the map

$$(L_{4n}, N_{4n}, m_{4n}) \mapsto (i_{4n+1,4n}, (L_{4n+1}, N_{4n+1}, m_{4n+1}))$$

is a splitting transformation, and so by Lemma 4.20, $(\mathcal{N}_\omega, m_\omega)$ is a continuous layering, since $m_\omega(*_\omega) = 0$.

For every n , the map

$$(L_{4n+1}, N_{4n+1}, m_{4n+1}) \mapsto (i_{4n+2,4n+1}, (L_{4n+2}, N_{4n+2}, m_{4n+2}))$$

is a $(\forall \bar{x})(\exists y)\varphi_n(\bar{x}, y)$ -satisfaction transformation, and so by Lemma 4.22, (\mathcal{N}, m, i) asymptotically satisfies each sentence in T . Hence by Corollary 4.11, $\mu_{(\mathcal{N}_\omega, m_\omega)}$ is concentrated on T .

For every n , the map

$$(L_{4n+2}, N_{4n+2}, m_{4n+2}) \mapsto (i_{4n+3,4n+2}, (L_{4n+3}, N_{4n+3}, m_{4n+3}))$$

is a omitting- $p_n(\bar{x})$ transformation, and so by Lemma 4.24, (\mathcal{N}, m, i) asymptotically omits each qf-type in Θ . Hence by Corollary 4.13, $\mu_{(\mathcal{N}_\omega, m_\omega)}$ omits Θ . \square

Using this layering, we may now build an ergodic model of the theory.

Proposition 5.2. *Let L be a countable language. Suppose A is a countable fragment of $\mathcal{L}_{\omega_1, \omega}(L)$ and T is a complete A -theory. If T has trivial A -definable closure then there is an ergodic model of T .*

Proof. Recall from §2.3 the relational language L_A , the non-redundant pithy Π_2 first-order L_A -theory Th_A , and the countable collection of qf-types Θ_A of L_A . Further recall the pithy Π_2 non-redundant L_A -theory $T_A = \text{Th}_A \cup \{R_\varphi : \varphi \in T\}$.

Suppose T has trivial A -definable closure. Then T_A has trivial A^+ -definable closure as well by Lemma 2.21. By Corollary 2.14 there is an ergodic model of T if and only if there is an ergodic model of T_A^+ . Hence our problem is reduced to that of constructing an ergodic model of T_A^+ , i.e., an ergodic L_A -structure almost surely satisfying T_A and omitting Θ_A .

The A -theory T is countable and consistent; hence let \mathcal{K} be a model of T . By Lemma 2.11, the $(L_A \cup L)$ -structure \mathcal{K}_A is a model of T_A that omits Θ_A . Define

the L_A -structure $\mathcal{M} := \mathcal{K}_A|_{L_A}$, i.e., the pithy Π_2 Morleyization of \mathcal{K} , which is also a model of T_A that omits Θ_A . Further, \mathcal{M} is non-redundant (because Th_A is non-redundant by Lemma 2.8), as is every finite substructure of \mathcal{M} .

The theory T_A has trivial A^+ -definable closure, and so by Lemma 2.24 it also has duplication of formulas in A^+ . In particular, T_A has duplication of quantifier-free formulas.

By Proposition 5.1, there is a continuous regular layering (\mathcal{N}, m, i) such that $\mu_{(\mathcal{N}_\omega, m_\omega)}$ is an ergodic model of T_A that omits Θ_A , as desired. \square

6. NON-EXISTENCE OF ERGODIC STRUCTURES

We now show that if a complete A -theory has non-trivial A -definable closure then it has no ergodic model. As in Section 5, we work with countable *relational* languages L , but in the proof of Theorem 1.1 in Section 7, we will see that this holds for arbitrary countable languages.

Proposition 6.1. *Let L be a countable relational language, and let A be a countable fragment of $\mathcal{L}_{\omega_1, \omega}(L)$. Suppose T is a complete A -theory that has non-trivial A -definable closure. Then there does not exist an ergodic model of T .*

Proof. By Lemmas 2.20 and 2.25, there is some non-redundant formula $\varphi(\bar{x}, y)$ such that $T \models (\exists \bar{x})(\exists^{=1} y) \varphi(\bar{x}, y)$.

Suppose, towards a contradiction, that there is an ergodic model μ of T . Define $P := \llbracket (\exists^{=1} y) \varphi(0, 1, \dots, |\bar{x}| - 1, y) \rrbracket$. First note that if $\mu(P) = 0$, then by invariance

$$\mu(\llbracket (\exists^{=1} y) \varphi(n_0, n_1, \dots, n_{|\bar{x}|-1}, y) \rrbracket) = 0$$

for any distinct elements $n_0, n_1, \dots, n_{|\bar{x}|-1} \in \omega$. Hence $\mu(\llbracket (\exists \bar{x})(\exists^{=1} y) \varphi(\bar{x}, y) \rrbracket) = 0$, contradicting the fact that $T \models (\exists \bar{x})(\exists^{=1} y) \varphi(\bar{x}, y)$. Therefore $\mu(P) > 0$.

Define the probability measure μ^* to be the conditional distribution of μ given that P holds, i.e.,

$$\mu^*(X) = \frac{\mu(P \cap X)}{\mu(P)}$$

for all Borel sets X . Note that μ^* is invariant under any permutation of ω that fixes $0, 1, \dots, |\bar{x}| - 1$.

By the definition of P , we have

$$\llbracket \varphi(0, 1, \dots, |\bar{x}| - 1, i) \rrbracket \cap \llbracket \varphi(0, 1, \dots, |\bar{x}| - 1, j) \rrbracket \cap P = \emptyset$$

for i, j satisfying $|\bar{x}| - 1 < i < j$. Because φ is non-redundant, we also have

$$\bigcup_{i \leq |\bar{x}|-1} \llbracket \varphi(0, 1, \dots, |\bar{x}| - 1, i) \rrbracket = \emptyset.$$

Hence

$$1 = \mu^*(P) = \sum_{i \geq |\bar{x}|} \mu^*(\llbracket \varphi(0, 1, \dots, |\bar{x}| - 1, i) \rrbracket).$$

But if $\alpha = \mu^*(\llbracket \varphi(0, 1, \dots, |\bar{x}| - 1, |\bar{x}|) \rrbracket)$ then by invariance of μ^* we also have $\alpha = \mu^*(\llbracket \varphi(0, 1, \dots, |\bar{x}| - 1, i) \rrbracket)$ for any $i \geq |\bar{x}|$. Hence $1 = \sum_{i \geq |\bar{x}|} \alpha$, which is a contradiction.

Therefore there is no ergodic model of T . \square

7. CLASSIFICATION OF ERGODIC STRUCTURES

Putting the results of Sections 5 and 6 together, we obtain our main theorem.

Proof of Theorem 1.1. Let L be a countable language (not necessarily relational), and let $\Sigma \subseteq \mathcal{L}_{\omega_1, \omega}(L)$ be countable.

First observe that (2) immediately implies (1). On the other hand, (1) implies (2) by Lemma 2.2. Further, (4) implies (3) by letting A be any countable fragment such that $\Sigma \subseteq A$ (for example, the fragment generated by Σ). To conclude the proof, we will show that (3) implies (2) and that (2) implies (4).

Let A and T be as in (3). Consider the countable relational language L_A , the countable fragment A^+ , and the A^+ -theory T_A^+ . Recall, by Corollary 2.16, that T_A^+ is a complete A^+ -theory. By Lemma 2.21, T_A^+ has trivial A^+ -definable closure, as T has trivial A -definable closure. Hence by Proposition 5.2, there is an ergodic model of T_A^+ . By Corollary 2.14, there is also an ergodic model of T , and hence of Σ , and so (2) holds.

Now suppose that (2) holds, and let μ be an ergodic model of Σ . Let A be an arbitrary countable fragment such that $\Sigma \subseteq A$. By Lemma 2.4, $\text{Th}(\mu)$ is a complete $\mathcal{L}_{\omega_1, \omega}(L)$ -theory. Hence $\text{Th}(\mu) \cap A$ is a complete A -theory that admits an invariant measure, and is such that $\Sigma \subseteq \text{Th}(\mu) \cap A$. By Proposition 6.1, the theory $\text{Th}(\mu) \cap A$ must have trivial A -definable closure, and so (4) holds. \square

We obtain Theorem 1.2 as a corollary by specializing to the fragment $\mathcal{L}_{\omega, \omega}(L)$ of first-order L -formulas.

Proof of Theorem 1.2. Suppose Σ admits an invariant measure; then (1) of Theorem 1.1 holds. Hence taking $A = \mathcal{L}_{\omega, \omega}(L)$ in (4), there is a complete first-order theory $T \subseteq \mathcal{L}_{\omega, \omega}(L)$ with $\Sigma \subseteq T$ such that T has trivial $\mathcal{L}_{\omega, \omega}(L)$ -definable closure.

Conversely, suppose $T \subseteq \mathcal{L}_{\omega, \omega}(L)$ is a complete first-order theory such that $\Sigma \subseteq T$ and T has trivial $\mathcal{L}_{\omega, \omega}(L)$ -definable closure. Then (3) of Theorem 1.1 holds. Therefore by (1) there is an invariant measure concentrated on T and hence on Σ . \square

ACKNOWLEDGEMENTS

The authors would like to thank Alex Kruckman for helpful discussions and detailed comments on a draft, and Haim Gaifman for illuminating conversations about the genesis of his work relating logic and probability.

This research was facilitated by participation in the Trimester Program on Universality and Homogeneity of the Hausdorff Research Institute for Mathematics at the University of Bonn (September–December 2013), the LMS–EPSRC Durham Symposium on Permutation Groups and Transformation Semigroups at Durham University (July 2015), the workshop on Logic and Random Graphs at the Lorentz Center (August–September 2015), and the workshop on Homogeneous Structures at the Banff International Research Station for Mathematical Innovation and Discovery (November 2015).

Work on this publication by C. F. was made possible through the support of ARO grant W911NF-13-1-0212 and grants from the John Templeton Foundation and Google. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the U.S. Government or the John Templeton Foundation.

REFERENCES

- [AFKRP17] N. Ackerman, C. Freer, A. Kruckman, and R. Patel, *Properly ergodic structures*, Preprint (2017).
- [AFKwP17] N. Ackerman, C. Freer, A. Kwiatkowska, and R. Patel, *A classification of orbits admitting a unique invariant measure*, *Annals of Pure and Applied Logic* **168** (2017), no. 1, 19–36.
- [AFNP16] N. Ackerman, C. Freer, J. Nešetřil, and R. Patel, *Invariant measures via inverse limits of finite structures*, *Eur. J. Combin.* **52** (2016), 248–289.
- [AFP16] N. Ackerman, C. Freer, and R. Patel, *Invariant measures concentrated on countable structures*, *Forum Math. Sigma* **4** (2016), no. e17, 59 pp.
- [Ald81] D. J. Aldous, *Representations for partially exchangeable arrays of random variables*, *J. Multivariate Anal.* **11** (1981), no. 4, 581–598.
- [Gai64] H. Gaifman, *Concerning measures in first order calculi*, *Israel J. Math.* **2** (1964), 1–18.
- [Hoo79] D. N. Hoover, *Relations on probability spaces and arrays of random variables*, Preprint, Institute for Advanced Study, Princeton, NJ, 1979.
- [Kal05] O. Kallenberg, *Probabilistic symmetries and invariance principles*, *Probability and its Applications*, Springer-Verlag, New York, 2005.
- [Kec95] A. S. Kechris, *Classical descriptive set theory*, *Graduate Texts in Mathematics*, vol. 156, Springer-Verlag, New York, 1995.
- [Kra69] P. H. Krauss, *Representation of symmetric probability models*, *J. Symbolic Logic* **34** (1969), 183–193.
- [Kru16] A. Kruckman, *Infinitary limits of finite structures*, Ph.D. thesis, University of California, Berkeley, 2016.
- [Lov12] L. Lovász, *Large networks and graph limits*, *American Mathematical Society Colloquium Publications*, vol. 60, American Mathematical Society, Providence, RI, 2012.

- [PV10] F. Petrov and A. Vershik, *Uncountable graphs and invariant measures on the set of universal countable graphs*, *Random Struct. Algor.* **37** (2010), no. 3, 389–406.
- [SK66] D. Scott and P. Krauss, *Assigning probabilities to logical formulas*, *Aspects of inductive logic* (ed. Jaakko Hintikka and Patrick Suppes), *Studies in Logic and the Foundations of Mathematics*, North-Holland Publishing Co., Amsterdam, 1966, pp. 219–264.

DEPARTMENT OF MATHEMATICS, HARVARD UNIVERSITY, ONE OXFORD STREET,
CAMBRIDGE, MA 02138, USA

E-mail address: nate@math.harvard.edu

REMINE, 2722 MERRILEE DRIVE SUITE 300, FAIRFAX, VA 22031, USA

E-mail address: cameron@remine.com

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, WHEATON COLLEGE,
NORTON, MA, 02766, USA

E-mail address: patel_rehana@wheatoncollege.edu