A CLASSIFICATION OF ORBITS ADMITTING A UNIQUE INVARIANT MEASURE

NATHANAEL ACKERMAN, CAMERON FREER,
ALEKSANDRA KWIATKOWSKA, AND REHANA PATEL

ABSTRACT. We consider the space of countable structures with fixed underlying set in a given countable language. We show that the number of ergodic probability measures on this space that are S_∞-invariant and concentrated on a single isomorphism class must be zero, or one, or continuum. Further, such an isomorphism class admits a unique S_∞-invariant probability measure precisely when the structure is highly homogeneous; by a result of Peter J. Cameron, these are the structures that are interdefinable with one of the five reducts of the rational linear order $(\mathbb{Q}, <)$.

1. INTRODUCTION

A countable structure in a countable language can be said to admit a random symmetric construction when there is a probability measure on its isomorphism class (of structures having a fixed underlying set) that is invariant under the logic action of S_∞. Ackerman, Freer, and Patel [AFP16] characterized those structures admitting such invariant measures. In this paper, we further explore this setting

August 26, 2016
Keywords: invariant measure, high homogeneity, unique ergodicity.
2010 *Mathematics Subject Classification:* 03C98, 37L40, 60G09, 20B27.

by determining the possible numbers of such ergodic invariant measures, and by characterizing when there is a unique invariant measure.

A dynamical system is said to be uniquely ergodic when it admits a unique, hence necessarily ergodic, invariant measure. In most classical ergodic-theoretic settings, the dynamical system consists of a measure space along with a single map, or at most a countable semigroup of transformations; unique ergodicity has been of longstanding interest for such systems. In contrast, unique ergodicity for systems consisting of a larger space of transformations (such as the automorphism group of a structure) has been a focus of more recent research, notably that of Glasner and Weiss [GW02], and of Angel, Kechris, and Lyons [AKL14].

When studying continuous dynamical systems, one often considers minimal flows, i.e., continuous actions on compact Hausdorff spaces such that each orbit is dense; [AKL14] examines unique ergodicity in this setting. In the present paper, we are interested in unique ergodicity of actions where the underlying space need not be compact and there is just one orbit: We characterize when the logic action of the group S_∞ on an orbit is uniquely ergodic.

Any transitive S_∞-space is isomorphic to the action of S_∞ on the isomorphism class of a countable structure, restricted to a fixed underlying set. The main result of [AFP16] states that such an isomorphism class admits at least one S_∞-invariant measure precisely when the structure has trivial definable closure. Here we characterize those countable structures whose isomorphism classes admit *exactly one* such measure, and show via a result of Peter J. Cameron that the five reducts of $(\mathbb{Q}, <)$ are essentially the only ones. Furthermore, if the isomorphism class of a countable structure admits more than one S_∞-invariant measure, it must admit continuum-many ergodic such measures.

1.1. **Motivation and main results.** In this paper we consider, for a given countable language L, the collection of countable L-structures having the natural numbers $\mathbb{N}$ as underlying set. This collection can be made into a measurable space, denoted Str_L, in a standard way, as we describe in Section 2.

The group S_∞ of permutations of $\mathbb{N}$ acts naturally on Str_L by permuting the underlying set of elements. This action is known as the *logic action* of S_∞ on Str_L, and has been studied extensively in descriptive set theory. For details, see [BK96, §2.5] or [Gao09, §11.3]. Observe that the S_∞-orbits of Str_L are precisely the isomorphism classes of structures in Str_L.

By an *invariant measure* on Str_L, we will always mean a Borel probability measure on Str_L that is invariant under the logic action of S_∞. We are specifically interested in those invariant measures on Str_L that assign measure 1 to a single orbit, i.e., the isomorphism class in Str_L of some countable L-structure $\mathcal{M}$. In this case we say that the orbit of $\mathcal{M}$ *admits* an invariant measure, or simply that $\mathcal{M}$ admits an invariant measure.

When a countable structure $\mathcal{M}$ admits an invariant measure, this measure can be thought of as providing a symmetric random construction of $\mathcal{M}$. The main result of [AFP16] describes precisely when such a construction is possible: A structure $\mathcal{M} \in \mathrm{Str}_L$ admits an invariant measure if and only if definable closure in $\mathcal{M}$ is trivial, i.e., the pointwise stabilizer in $\mathrm{Aut}(\mathcal{M})$ of any finite tuple fixes no additional elements. But even when there *are* invariant measures concentrated on the orbit of $\mathcal{M}$, it is not obvious how many different ones there are.

If an orbit admits at least two invariant measures, there are trivially always continuum-many such measures, because a convex combination of any two gives an invariant measure on that orbit, and these combinations yield distinct measures. It is therefore useful to count instead the invariant measures that are not decomposable in this way, namely the *ergodic* ones. It is a standard fact that the invariant measures on Str_L form a simplex in which the ergodic invariant measures are precisely the *extreme* points, i.e., those that cannot be written as a nontrivial convex combination of invariant measures. Moreover, every invariant measure is a mixture of these extreme invariant measures. (For more details, see [Kal05, Lemma A1.2 and Theorem A1.3] and [Phe01, Chapters 10 and 12].) Thus when counting invariant measures on an orbit, the interesting quantity to consider is the number of ergodic invariant measures.

Many natural examples admit more than one invariant measure. For instance, consider the Erdős–Rényi [ER59] construction $G(\mathbb{N}, p)$ of the Rado graph, a countably infinite random graph in which edges have independent probability p, where $0 < p < 1$. This yields continuum-many ergodic invariant measures concentrated on the orbit of the Rado graph, as each value of p leads to a different ergodic invariant measure.

On the other hand, some countable structures admit just one invariant measure. One such example is well-known: The rational linear order $(\mathbb{Q}, <)$ admits a unique invariant measure, which can be described as follows. For every finite n-tuple of distinct elements of $\mathbb{Q}$, each of the $n!$-many orderings of the n-tuple must be assigned the same probability, by S_∞-invariance. This collection of finite-dimensional marginal distributions determines a (necessarily unique) invariant measure on Str_L, by the Kolmogorov extension theorem (see, e.g., [Kal02, Theorems 6.14 and 6.16]). This probability measure can be shown to be concentrated on the orbit of $(\mathbb{Q}, <)$ and is sometimes known as the *Glasner–Weiss measure*; for details, see [GW02, Theorems 8.1 and 8.2]. We will discuss this measure and its construction further in Section 3.

In fact, these examples illustrate the only possibilities: Either a countable structure admits no invariant measure, or a unique invariant measure, or continuum-many ergodic invariant measures. Furthermore, a countable structure admits a unique invariant measure precisely when it has the property known as *high homogeneity*. The main result of this paper is the following trichotomy. It refines the dichotomy obtained in [AFP16].

Theorem 1.1. *Let $\mathcal{M}$ be a countable structure in a countable language L. Then exactly one of the following holds:*

(0) *The structure $\mathcal{M}$ has nontrivial definable closure, in which case there is no S_∞-invariant Borel probability measure on* Str_L *that is concentrated on the orbit of $\mathcal{M}$.*

(1) *The structure $\mathcal{M}$ is highly homogeneous, in which case there is a unique S_∞-invariant Borel probability measure on* Str_L *that is concentrated on the orbit of $\mathcal{M}$.*

($2^{\aleph_0}$) *There are continuum-many ergodic S_∞-invariant Borel probability measures on* Str_L *that are concentrated on the orbit of $\mathcal{M}$.*

Moreover, by a result of Peter J. Cameron, the case where $\mathcal{M}$ is highly homogeneous is equivalent to $\mathcal{M}$ being interdefinable with a definable reduct (henceforth *reduct*) of $(\mathbb{Q}, <)$, of which there are five. In particular, this shows the known invariant measures on these five to be canonical.

1.2. **Additional motivation.** The present work has been motivated by further considerations, which we now describe.

Fouché and Nies [Nie13, §15] describe one notion of an algorithmically random presentation of a given countable structure; see also [Fou13], [Fou12], and [FP98]. In the case of the rational linear order $(\mathbb{Q}, <)$, they note that their notion of randomness is in a sense canonical by virtue of the unique ergodicity of the orbit of $(\mathbb{Q}, <)$. Hence one may ask which other orbits of countable structures are uniquely ergodic. Theorem 1.1, along with the result of Cameron, shows that the orbits of $(\mathbb{Q}, <)$ and of its reducts are essentially the only instances.

We also note a connection with "Kolmogorov's example" of a transitive but non-ergodic action of S_∞, described by Vershik in [Ver03]. Many settings in classical ergodic theory permit at most one invariant measure. For example, when a separable locally compact group G acts continuously and transitively on a Polish space X, there is at most one G-invariant probability measure on X [Ver03, Theorem 2]. The action of S_∞, however, allows for continuum-many ergodic invariant measures on the same orbit, as noted above in the case of the Erdős–Rényi constructions; for more details, see [Ver03, §3]. Indeed, many specific orbits with this property are known, but the present work strengthens the sense in which this is typical for S_∞ and uniqueness is rare: There are essentially merely five exceptions to the rule of having either continuum-many ergodic invariant measures or none.

It may be an interesting question to further understand the structure of the simplex of invariant measures in the non-uniquely-ergodic case — when, as we show in the case of a single orbit, there are continuum-many ergodic such measures or none. Note, however, that this space will often not be compact, as the actions we consider are on spaces that are usually not compact.

2. Preliminaries

In this paper, L will always be a countable language. We consider the space Str_L of countable L-structures having underlying set $\mathbb{N}$, equipped with the σ-algebra of Borel sets generated by the topology described in Definition 2.1 below. We will often use the notation $\mathbf{x}$ to denote the finite tuple of variables $x_0 \cdots x_{n-1}$, where $n = |\mathbf{x}|$.

Recall that $\mathcal{L}_{\omega_1,\omega}(L)$ denotes the infinitary language based on L consisting of formulas that can have countably infinite conjunctions and disjunctions, but only finitely many quantifiers and free variables; for details, see [Kec95, §16.C].

Definition 2.1. Given a formula $\varphi \in \mathcal{L}_{\omega_1,\omega}(L)$ and $n_0, \ldots, n_{j-1} \in \mathbb{N}$, where j is the number of free variables of φ, define

$$[\![\varphi(n_0, \ldots, n_{j-1})]\!] := \{\mathcal{M} \in \mathrm{Str}_L : \mathcal{M} \models \varphi(n_0, \ldots, n_{j-1})\}.$$

Sets of this form are closed under finite intersection, and form a basis for the topology of Str_L.

Consider S_∞, the permutation group of the natural numbers $\mathbb{N}$. This group acts on Str_L via the **logic action**: For $g \in S_\infty$ and $\mathcal{M}, \mathcal{N} \in \mathrm{Str}_L$, we define $g \cdot \mathcal{M} = \mathcal{N}$ to hold whenever

$$R^{\mathcal{N}}(s_0, \ldots, s_{k-1}) \quad \text{if and only if} \quad R^{\mathcal{M}}\big(g^{-1}(s_0), \ldots, g^{-1}(s_{k-1})\big)$$

for all relation symbols $R \in L$ and $s_0, \ldots, s_{k-1} \in \mathbb{N}$, where k is the arity of R, and similarly with constant and function symbols. Observe that the orbit of a structure under the logic action is its isomorphism class in Str_L; every such orbit is Borel by Scott's isomorphism theorem. For more details on the logic action, see [Kec95, §16.C].

We define an **invariant measure** on Str_L to be a Borel probability measure μ on Str_L that is invariant under the logic action of S_∞ on Str_L, i.e., $\mu(X) = \mu(g \cdot X)$ for every Borel set $X \subseteq \mathrm{Str}_L$ and $g \in S_\infty$. When an invariant measure on Str_L is concentrated on the orbit of some structure in Str_L, we may restrict attention to this orbit, and speak equivalently of an invariant measure on the orbit itself.

In this paper, we are interested in invariant measures that are *ergodic*. Given an action of a group G on a set X, and an element $g \in G$, we write gx to denote the image of $x \in X$ under g, and for $A \subseteq X$ we write $gA := \{gx : x \in A\}$.

Definition 2.2. Consider a Borel action of a Polish group G on a standard Borel space X. A probability measure μ on X is **ergodic** when for every Borel $B \subseteq X$ satisfying $\mu(B \mathbin{\triangle} g^{-1}B) = 0$ for each $g \in G$, either $\mu(B) = 0$ or $\mu(B) = 1$.

In other words, an ergodic measure (with respect to a particular action of G) is one that assigns every almost G-invariant set either zero or full measure. In our setting, X will be one of $\mathbb{R}^\omega$ or Str_L, and we will consider the group action of S_∞ on $\mathbb{R}^\omega$ that permutes coordinates, and the logic action of S_∞ on Str_L.

A structure $\mathcal{M} \in \mathrm{Str}_L$ has *trivial definable closure* when the pointwise stabilizer in $\mathrm{Aut}(\mathcal{M})$ of an arbitrary finite tuple of $\mathcal{M}$ fixes no additional points:

Definition 2.3. Let $\mathcal{M} \in \mathrm{Str}_L$. For a tuple $\mathbf{a} \in \mathcal{M}$, the **definable closure** of $\mathbf{a}$ in $\mathcal{M}$, written $\mathrm{dcl}_{\mathcal{M}}(\mathbf{a})$, is the set of elements of $\mathcal{M}$ that are fixed by every automorphism of $\mathcal{M}$ that fixes $\mathbf{a}$ pointwise. The structure $\mathcal{M}$ has **trivial definable closure** when $\mathrm{dcl}_{\mathcal{M}}(\mathbf{a}) = \mathbf{a}$ for every (finite) tuple $\mathbf{a} \in \mathcal{M}$.

The easier direction of the main theorem of [AFP16] states that any structure admitting an invariant measure must have trivial definable closure.

Theorem 2.4 ([AFP16, Theorem 4.1])**.** *Let $\mathcal{M} \in \mathrm{Str}_L$. If $\mathcal{M}$ does not have trivial definable closure, then $\mathcal{M}$ does not admit an invariant measure.*

This corresponds to case (0) of Theorem 1.1.

2.1. **Canonical structures and interdefinability.** In the proof of our main theorem, we will work in the setting of *canonical languages* and *canonical structures*. We provide a brief description of these notions here; for more details, see [AFP16, §2.5].

Definition 2.5. Let G be a closed subgroup of S_∞, and consider the action of G on $\mathbb{N}$. Define the **canonical language for** G to be the (countable) relational language L_G that consists of, for each $k \in \mathbb{N}$ and G-orbit $E \subseteq \mathbb{N}^k$, a k-ary relation symbol R_E. Then define the **canonical structure for** G to be the structure $\mathcal{C}_G \in \mathrm{Str}_{L_G}$ in which, for each G-orbit E, the interpretation of R_E is the set E.

Definition 2.6. Given a structure $\mathcal{M} \in \mathrm{Str}_L$, define the **canonical language for** $\mathcal{M}$, written $L_{\overline{\mathcal{M}}}$, to be the countable relational language L_G where $G := \mathrm{Aut}(\mathcal{M})$. Similarly, define the **canonical structure for** $\mathcal{M}$, written $\overline{\mathcal{M}}$, to be the countable $L_{\overline{\mathcal{M}}}$-structure $\mathcal{C}_G$.

Structures that are *interdefinable*, in the following sense, will be regarded as interchangeable for purposes of our classification.

Definition 2.7. Let $\mathcal{M}$ and $\mathcal{N}$ be structures in (possibly different) countable languages, both having underlying set $\mathbb{N}$. Then $\mathcal{M}$ and $\mathcal{N}$ are said to be **interdefinable** when they have the same canonical language and same canonical structure.

Note that two structures are interdefinable if and only if there is an $\mathcal{L}_{\omega_1,\omega}$-*interdefinition* between them, in the terminology of [AFP16, Definition 2.11]; see the discussion after [AFP16, Lemma 2.13] for details.

By Definitions 2.6 and 2.7, it is immediate that a structure $\mathcal{M} \in \mathrm{Str}_L$ and its canonical structure $\overline{\mathcal{M}}$ are interdefinable.

Proposition 2.8. *Let $\mathcal{M} \in \mathrm{Str}_L$. There is a Borel bijection, respecting the action of S_∞, between the orbit of $\mathcal{M}$ in Str_L and the orbit of its canonical structure $\overline{\mathcal{M}}$ in $\mathrm{Str}_{L_{\overline{\mathcal{M}}}}$. In particular, this map induces a bijection between the set of ergodic invariant measures on the orbit of $\mathcal{M}$ and the set of ergodic invariant measures on the orbit of $\overline{\mathcal{M}}$.*

Proof. First observe that the orbit of $\mathcal{M}$ and the orbit of $\overline{\mathcal{M}}$ are each Borel spaces that inherit the logic action. The structures $\mathcal{M}$ and $\overline{\mathcal{M}}$ are interdefinable, and so [AFP16, Lemma 2.14] applies. The proof of this lemma provides explicit maps between Str_L and $\mathrm{Str}_{L_{\overline{\mathcal{M}}}}$ which, when restricted respectively to the orbit of $\mathcal{M}$ and of $\overline{\mathcal{M}}$, have the desired property. □

We immediately obtain the following corollary.

Corollary 2.9. *Let $\mathcal{M} \in \mathrm{Str}_L$. Then $\mathcal{M}$ and its canonical structure $\overline{\mathcal{M}}$ admit the same number of ergodic invariant measures.*

The following two results are straightforward.

Lemma 2.10 ([AFP16, Lemma 2.15])**.** *Let $\mathcal{M}$ and $\mathcal{N}$ be interdefinable structures in (possibly different) countable languages, both having underlying set $\mathbb{N}$. Then $\mathcal{M}$ has trivial definable closure if and only if $\mathcal{N}$ does.*

Corollary 2.11. *Let $\mathcal{M} \in \mathrm{Str}_L$. Then $\mathcal{M}$ has trivial definable closure if and only if its canonical structure $\overline{\mathcal{M}}$ has trivial definable closure.*

For any $\mathcal{M} \in \mathrm{Str}_L$, we will see in Lemma 2.17 that $\mathcal{M}$ is highly homogeneous if and only if $\overline{\mathcal{M}}$ is; combining this fact with Corollaries 2.9 and 2.11, when proving Theorem 1.1 it will suffice to consider instead the number of ergodic invariant measures admitted by the canonical structure $\overline{\mathcal{M}}$.

2.2. **Ultrahomogeneous structures.** Countable *ultrahomogeneous* relational structures play an important role throughout this paper, as canonical structures are ultrahomogeneous and canonical languages are relational.

Definition 2.12. A countable structure $\mathcal{M}$ is **ultrahomogeneous** if every isomorphism between finitely generated substructures of $\mathcal{M}$ can be extended to an automorphism of $\mathcal{M}$.

The following fact is folklore (see also the discussion following [AFP16, Proposition 2.17]).

Proposition 2.13. *Let $\mathcal{M} \in \mathrm{Str}_L$. The canonical structure $\overline{\mathcal{M}}$ is ultrahomogeneous.*

Ultrahomogeneous structures can be given particularly convenient $\mathcal{L}_{\omega_1,\omega}(L)$ axiomatizations via *pithy* Π_2 sentences, which can be thought of as "one-point extension axioms".

Definition 2.14 ([AFP16, Definitions 2.3 and 2.4])**.** A sentence in $\mathcal{L}_{\omega_1,\omega}(L)$ is Π_2 when it is of the form $(\forall \mathbf{x})(\exists \mathbf{y})\psi(\mathbf{x},\mathbf{y})$, where the (possibly empty) tuple $\mathbf{xy}$ consists of distinct variables, and $\psi(\mathbf{x},\mathbf{y})$ is quantifier-free. A countable theory T of $\mathcal{L}_{\omega_1,\omega}(L)$ is Π_2 when every sentence in T is Π_2.

A Π_2 sentence $(\forall \mathbf{x})(\exists \mathbf{y})\psi(\mathbf{x},\mathbf{y}) \in \mathcal{L}_{\omega_1,\omega}(L)$, where $\psi(\mathbf{x},\mathbf{y})$ is quantifier-free, is said to be **pithy** when the tuple $\mathbf{y}$ consists of precisely one variable. A countable Π_2 theory T of $\mathcal{L}_{\omega_1,\omega}(L)$ is said to be pithy when every sentence in T is pithy. Note that we allow the degenerate case where $\mathbf{x}$ is the empty tuple and ψ is of the form $(\exists y)\psi(y)$.

The following result follows from essentially the same proof as [AFP16, Proposition 2.17]. It states that a Scott sentence for an ultrahomogeneous relational structure is equivalent to a theory of a particular syntactic form.

Proposition 2.15. *Let L be relational and let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous. There is a countable $\mathcal{L}_{\omega_1,\omega}(L)$-theory, every sentence of which is pithy Π_2, and all of whose countable models are isomorphic to $\mathcal{M}$.*

We will call this theory the **Fraïssé theory of** $\mathcal{M}$.

2.3. **Highly homogeneous structures.** *High homogeneity* is the key notion in our characterization of structures admitting a unique invariant measure.

Definition 2.16 ([Cam90, §2.1])**.** A structure $\mathcal{M} \in \mathrm{Str}_L$ is **highly homogeneous** when, for each $k \in \mathbb{N}$ and for every pair of k-element sets $X, Y \subseteq \mathcal{M}$, there is some $f \in \mathrm{Aut}(\mathcal{M})$ such that $Y = \{f(x) : x \in X\}$.

The following lemma is immediate, and allows us to generalize the notion of high homogeneity to permutation groups.

Lemma 2.17. *A structure $\mathcal{M} \in \mathrm{Str}_L$ is highly homogeneous if and only if its canonical structure is.*

Definition 2.18 ([Cam90, §2.1])**.** A closed permutation group G on $\mathbb{N}$ is called **highly homogeneous** when its canonical structure $\mathcal{C}_G$ is highly homogeneous.

The crucial fact about highly homogeneous structures is the following.

Lemma 2.19. *Let L be relational and let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous. Then* $\mathrm{Aut}(\mathcal{M})$ *is highly homogeneous if and only if for any $k \in \mathbb{N}$, all k-element substructures of $\mathcal{M}$ are isomorphic.*

Proof. Let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous, and let X and Y be arbitrary substructures of $\mathcal{M}$ of size k. If $\mathrm{Aut}(\mathcal{M})$ is highly homogeneous, then X and Y are isomorphic via the restriction of any $f \in \mathrm{Aut}(\mathcal{M})$ such that $Y = \{f(x) : x \in X\}$. Conversely, if there is some isomorphism of structures $g\colon X \to Y$, then by the ultrahomogeneity of $\mathcal{M}$, there is some $f \in \mathrm{Aut}(\mathcal{M})$ extending g to all of $\mathcal{M}$. $\square$

Highly homogeneous structures have been classified explicitly by Cameron [Cam76], and characterized (up to interdefinability) as the five reducts of $(\mathbb{Q}, <)$, as we now describe.

Let $(\mathbb{Q}, <)$ be the set of rational numbers equipped with the usual order. The following three relations are definable within $(\mathbb{Q}, <)$:

(1) The ternary linear *betweenness* relation B, given by

$$B(a, b, c) \iff (a < b < c) \vee (c < b < a).$$

(2) The ternary *circular order* relation K, given by

$$K(a, b, c) \iff (a < b < c) \vee (b < c < a) \vee (c < a < b).$$

(3) The quaternary *separation* relation S, given by

$$\begin{aligned} S(a, b, c, d) \quad \iff \quad & \big(K(a, b, c) \wedge K(b, c, d) \wedge K(c, d, a)\big) \\ & \vee \big(K(d, c, b) \wedge K(c, b, a) \wedge K(b, a, d)\big). \end{aligned}$$

The structure $(\mathbb{Q}, B)$ can be thought of as forgetting the direction of the order, $(\mathbb{Q}, K)$ as gluing the rational line into a circle, and $(\mathbb{Q}, S)$ as forgetting which way is clockwise on this circle.

The following is a consequence of Theorem 6.1 in Cameron [Cam76]; see also (3.10) of [Cam90, §3.4]. For further details, see [Mac11, Theorem 6.2.1].

Theorem 2.20 (Cameron)**.** *Let G be a highly homogeneous structure. Then G is interdefinable with one of the following: the set $\mathbb{Q}$ (in the empty language), $(\mathbb{Q}, <)$, $(\mathbb{Q}, B)$, $(\mathbb{Q}, K)$, or $(\mathbb{Q}, S)$.*

Notice that these five structures all have trivial definable closure; this will imply, in Lemma 3.1, that the orbit of each highly homogeneous structure admits a unique invariant measure.

2.4. **Borel L-structures and ergodic invariant measures.** Aldous, Hoover, and Kallenberg have characterized ergodic invariant measures on Str_L in terms of a certain sampling procedure involving continuum-sized objects; for details, see [Aus08] and [Kal05].

We will obtain ergodic invariant measures via a special case of this procedure, by sampling from a particular kind of continuum-sized structure, called a *Borel L-structure*. For more on the connection between Borel L-structures and the Aldous–Hoover–Kallenberg representation, see [AFP16, §6.1].

Definition 2.21 ([AFP16, Definition 3.1])**.** Let L be relational, and let $\mathcal{P}$ be an L-structure whose underlying set is the set $\mathbb{R}$ of real numbers. We say that $\mathcal{P}$ is a **Borel L-structure** if for all relation symbols $R \in L$, the set

$$\{\mathbf{a} \in \mathcal{P}^j : \mathcal{P} \models R(\mathbf{a})\}$$

is a Borel subset of $\mathbb{R}^j$, where j is the arity of R.

The sampling procedure is given by the following map $\mathcal{F}_{\mathcal{P}}$ that takes each sequence of elements of $\mathcal{P}$ to the corresponding structure with underlying set $\mathbb{N}$.

Definition 2.22 ([AFP16, Definition 3.2])**.** Let L be relational and let $\mathcal{P}$ be a Borel L-structure. The map $\mathcal{F}_{\mathcal{P}}\colon \mathbb{R}^\omega \to \mathrm{Str}_L$ is defined as follows. For $\mathbf{t} = (t_0, t_1, \ldots) \in \mathbb{R}^\omega$, let $\mathcal{F}_{\mathcal{P}}(\mathbf{t})$ be the L-structure with underlying set $\mathbb{N}$ satisfying

$$\mathcal{F}_{\mathcal{P}}(\mathbf{t}) \models R(n_1, \ldots, n_j) \quad \Leftrightarrow \quad \mathcal{P} \models R(t_{n_1}, \ldots, t_{n_j})$$

for all $n_1, \ldots, n_j \in \mathbb{N}$ and for every relation symbol $R \in L$, and for which equality is inherited from $\mathbb{N}$, i.e.,

$$\mathcal{F}_{\mathcal{P}}(\mathbf{t}) \models (m \neq n)$$

if and only if m and n are distinct natural numbers.

The map $\mathcal{F}_{\mathcal{P}}$ is Borel measurable [AFP16, Lemma 3.3]. Furthermore, $\mathcal{F}_{\mathcal{P}}$ is an S_∞-map, i.e., $\sigma\mathcal{F}_{\mathcal{P}}(\mathbf{t}) = \mathcal{F}_{\mathcal{P}}(\sigma\mathbf{t})$ for every $\sigma \in S_\infty$ and $\mathbf{t} \in \mathbb{R}^\omega$.

The pushforward of $\mathcal{F}_{\mathcal{P}}$ gives rise to an ergodic invariant measure, as we will see in Proposition 2.24.

Definition 2.23 ([AFP16, Definition 3.4])**.** Let L be relational, let $\mathcal{P}$ be a Borel L-structure, and let m be a probability measure on $\mathbb{R}$. Define the measure $\mu_{(\mathcal{P},m)}$ on Str_L to be

$$\mu_{(\mathcal{P},m)} := m^\infty \circ {\mathcal{F}_{\mathcal{P}}}^{-1}.$$

Note that $m^\infty\big({\mathcal{F}_{\mathcal{P}}}^{-1}(\mathrm{Str}_L)\big) = 1$, and so $\mu_{(\mathcal{P},m)}$ 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 $\mathbb{R}$.

By [AFP16, Lemma 3.5], $\mu_{(\mathcal{P},m)}$ is an invariant measure on Str_L. In fact, $\mu_{(\mathcal{P},m)}$ is ergodic: Aldous showed that it is ergodic for finite relational languages [Kal05, Lemma 7.35]; we require the following extension of this result to the setting of countable (possibly infinite) relational languages, whose proof we include here for completeness.

Proposition 2.24. *Let L be relational, let $\mathcal{P}$ be a Borel L-structure, and let m be a probability measure on $\mathbb{R}$. Then the measure $\mu_{(\mathcal{P},m)}$ is ergodic.*

Proof. First note that the measure m^∞ on $\mathbb{R}^\omega$ is ergodic by the Hewitt–Savage 0–1 law; for details, see [Kal05, Corollary 1.6] and [Kal02, Theorem 3.15].

Write $\mu := \mu_{(\mathcal{P},m)}$. Let $B \subseteq \mathrm{Str}_L$ be Borel and suppose that $\mu(B \mathbin{\triangle} \sigma^{-1}B) = 0$ for every $\sigma \in S_\infty$. We will show that either $\mu(B) = 0$ or $\mu(B) = 1$.

Let $\mathbf{t} \in \mathbb{R}^\omega$ and $\sigma \in S_\infty$. We have

$$\mathbf{t} \in \sigma^{-1}{\mathcal{F}_{\mathcal{P}}}^{-1}(B) \iff \mathcal{F}_{\mathcal{P}}(\sigma\mathbf{t}) \in B,$$

where σ and σ^{-1} act on $\mathbb{R}^\omega$, and

$$\mathbf{t} \in {\mathcal{F}_{\mathcal{P}}}^{-1}(\sigma^{-1}B) \iff \sigma\mathcal{F}_{\mathcal{P}}(\mathbf{t}) \in B,$$

where σ and σ^{-1} act on Str_L via the logic action.

Now, $\sigma\mathcal{F}_{\mathcal{P}}(\mathbf{t}) = \mathcal{F}_{\mathcal{P}}(\sigma\mathbf{t})$, and so

$$\mathcal{F}_{\mathcal{P}}{}^{-1}(\sigma^{-1}B) = \sigma^{-1}\mathcal{F}_{\mathcal{P}}{}^{-1}(B).$$

Using this fact, we have

$$\begin{aligned}
0 &= \mu(B \mathbin{\triangle} \sigma^{-1}B)\\
&= m^\infty\big(\mathcal{F}_{\mathcal{P}}{}^{-1}(B \mathbin{\triangle} \sigma^{-1}B)\big)\\
&= m^\infty\big(\mathcal{F}_{\mathcal{P}}{}^{-1}(B) \mathbin{\triangle} \mathcal{F}_{\mathcal{P}}{}^{-1}(\sigma^{-1}B)\big)\\
&= m^\infty\big(\mathcal{F}_{\mathcal{P}}{}^{-1}(B) \mathbin{\triangle} \sigma^{-1}\mathcal{F}_{\mathcal{P}}{}^{-1}(B)\big)\\
&= m^\infty(A \mathbin{\triangle} \sigma^{-1}A),
\end{aligned}$$

where $A := \mathcal{F}_{\mathcal{P}}{}^{-1}(B)$.

Because m^∞ is ergodic and $m^\infty(A \mathbin{\triangle} \sigma^{-1}A) = 0$ for every $\sigma \in S_\infty$, either $m^\infty(A) = 0$ or $m^\infty(A) = 1$ must hold. But then as $\mu(B) = m^\infty(A)$, either $\mu(B) = 0$ or $\mu(B) = 1$, as desired. $\square$

Not all ergodic invariant measures are of the form $\mu_{(\mathcal{P},m)}$: For example, it can be shown that the distribution of an Erdős–Rényi graph $G(\mathbb{N}, p)$ for $0 < p < 1$, each of which is concentrated on the orbit of the Rado graph, is not of this form. However, Petrov and Vershik [PV10] have shown that the orbit of the Rado graph admits an invariant measure of the form $\mu_{(\mathcal{P},m)}$ (in our terminology). More generally, the proof of [AFP16, Corollary 6.1] shows that whenever an orbit admits an invariant measure, it admits one of the form $\mu_{(\mathcal{P},m)}$. Note that this class of invariant measures also occurs elsewhere; see Kallenberg's notion of *simple arrays* [Kal99] and, in the case of graphs, the notions of *random-free graphons* [Jan13, §10] or 0–1 *valued graphons* [LS10].

2.5. **Strong witnessing and the existence of invariant measures.** We now consider how to obtain ergodic invariant measures concentrated on a particular orbit. We will do so by obtaining ergodic invariant measures concentrated on the class of models in Str_L of a particular Fraïssé theory T, where this class is the desired orbit.

A measure m on $\mathbb{R}$ is said to be **nondegenerate** when every nonempty open set has positive measure, and **continuous** when it assigns measure zero to every singleton.

Definition 2.25 ([AFP16, Definition 3.8])**.** Let $\mathcal{P}$ be a Borel L-structure and let m be a probability measure on $\mathbb{R}$. Suppose T is a countable pithy Π_2 theory of $\mathcal{L}_{\omega_1,\omega}(L)$. We say that the pair $(\mathcal{P}, m)$ **witnesses** T if for every sentence $(\forall\mathbf{x})(\exists y)\psi(\mathbf{x}, y) \in T$, and for every tuple $\mathbf{a} \in \mathcal{P}$ such that $|\mathbf{a}| = |\mathbf{x}|$, we have either

(i) $\mathcal{P} \models \psi(\mathbf{a}, b)$ for some $b \in \mathbf{a}$, or

(ii) $\quad m\big(\{b \in \mathcal{P} : \mathcal{P} \models \psi(\mathbf{a}, b)\}\big) > 0.$

We say that $\mathcal{P}$ **strongly witnesses** T when, for every nondegenerate continuous probability measure m on $\mathbb{R}$, the pair $(\mathcal{P}, m)$ witnesses T.

Proposition 2.26 ([AFP16, Theorem 3.10]). *Let L be relational, let T be a countable pithy Π_2 theory of $\mathcal{L}_{\omega_1,\omega}(L)$, and let $\mathcal{P}$ be a Borel L-structure. Suppose m is a continuous probability measure on $\mathbb{R}$ such that $(\mathcal{P}, m)$ witnesses T. Then $\mu_{(\mathcal{P},m)}$ is concentrated on the set of structures in Str_L that are models of T.*

The main theorem of [AFP16] states that a countably infinite structure in a countable language admits at least one invariant measure if and only if it has trivial definable closure. The easier direction is stated in Theorem 2.4 above. Proposition 2.27, which we will need in the proof of Theorem 1.1, is the key result used in Theorem 2.28, essentially the harder direction of [AFP16].

Proposition 2.27 ([AFP16, Theorem 3.19 and Lemma 3.20]). *Let L be relational and let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous. If $\mathcal{M}$ has trivial definable closure, then there is a Borel L-structure $\mathcal{P}$ that strongly witnesses the Fraïssé theory of $\mathcal{M}$.*

Theorem 2.28 ([AFP16, Theorem 3.21]). *Let L be relational and let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous. If $\mathcal{M}$ has trivial definable closure, then $\mathcal{M}$ admits an invariant measure.*

Proof. There is a Borel L-structure $\mathcal{P}$ that strongly witnesses the Fraïssé theory of $\mathcal{M}$, by Proposition 2.27. Let m be any nondegenerate continuous probability measure m on $\mathbb{R}$ (e.g., a Gaussian). Then by Proposition 2.26, the invariant measure $\mu_{(\mathcal{P},m)}$ is concentrated on the set of models of the Fraïssé theory of $\mathcal{M}$ in Str_L. In particular, $\mu_{(\mathcal{P},m)}$ is concentrated on the orbit of $\mathcal{M}$. $\square$

Finally, we establish a lemma about measures of the form $\mu_{(\mathcal{P},m)}$. Recall the notation $[\![\varphi]\!]$ from Definition 2.1.

Lemma 2.29. *Let L be relational, let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous, and let T be the Fraïssé theory of $\mathcal{M}$. Suppose that $\mathcal{P}$ is a Borel L-structure that strongly witnesses T. Let m be a nondegenerate continuous probability measure on $\mathbb{R}$. Then for every $n \in \mathbb{N}$ and every $\mathcal{L}_{\omega_1,\omega}(L)$-formula φ having n free variables,*

$$\mu_{(\mathcal{P},m)}\big([\![\varphi(0, \ldots, n-1)]\!]\big) = m^n\big(\{\mathbf{a} \in \mathbb{R}^n : \mathcal{P} \models \varphi(\mathbf{a})\}\big).$$

Proof. By [AFP16, Lemma 3.6], because m is continuous, $\mu_{(\mathcal{P},m)}$ is concentrated on the isomorphism classes in Str_L of countably infinite substructures of $\mathcal{P}$.

Because $\mathcal{M}$ is ultrahomogeneous, for every $\mathcal{L}_{\omega_1,\omega}(L)$-formula $\varphi(\mathbf{x})$ there is some quantifier-free $\psi(\mathbf{x})$ such that

$$\mathcal{M} \models (\forall \mathbf{x})\big(\varphi(\mathbf{x}) \leftrightarrow \psi(\mathbf{x})\big).$$

Because $\mathcal{P}$ strongly witnesses T, by [AFP16, Lemma 3.9] we have that $\mathcal{P} \models T$. Hence

$$\mathcal{P} \models (\forall \mathbf{x})\big(\varphi(\mathbf{x}) \leftrightarrow \psi(\mathbf{x})\big).$$

In particular, if a sequence of reals determines a substructure of $\mathcal{P}$ that is isomorphic to $\mathcal{M}$, then this substructure is in fact $\mathcal{L}_{\omega_1,\omega}(L)$-elementary.

Therefore, as $\mathcal{P}$ strongly witnesses T, by Proposition 2.26 and Definition 2.23 the probability measure m^∞ concentrates on sequences of reals that determine elementary substructures of $\mathcal{P}$. Hence the probability that a structure sampled according to $\mu_{(\mathcal{P},m)}$ satisfies $\varphi(0, \dots, n-1)$ is equal to

$$m^n\big(\{\mathbf{a} \in \mathbb{R}^n : \mathcal{P} \models \varphi(\mathbf{a})\}\big),$$

as desired. □

3. The number of ergodic invariant measures

In this section we prove our main result, Theorem 1.1.

3.1. **Unique invariant measures.** We now show that every ultrahomogeneous highly homogeneous structure admits a unique invariant measure. Recall Cameron's result, Theorem 2.20, that the highly homogeneous structures are (up to interdefinability) precisely the five reducts of the rational linear order $(\mathbb{Q}, <)$.

Lemma 3.1. *Let L be relational and let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous. If $\mathcal{M}$ is highly homogeneous, then there is an invariant measure on the isomorphism class of $\mathcal{M}$ in Str_L.*

Proof. We can check directly that each reduct of $(\mathbb{Q}, <)$ has trivial definable closure. By Theorem 2.20 and the hypothesis that $\mathcal{M}$ is highly homogeneous, $\mathcal{M}$ is interdefinable with one of these five. Hence $\mathcal{M}$ also has trivial definable closure by Lemma 2.10. Therefore by Theorem 2.28, there is an invariant measure on the isomorphism class of $\mathcal{M}$ in Str_L. □

Alternatively, instead of applying Theorem 2.28, there are several more direct ways of constructing an invariant measure on each of the five reducts of $(\mathbb{Q}, <)$. We sketched the construction of the Glasner–Weiss measure on $(\mathbb{Q}, <)$ in §1.1, as the weak limit of the uniform measures on n-element linear orders; each of the other four also arises as the weak limit of uniform measures.

Another way to construct the Glasner–Weiss measure is as the ordering on the set $\mathbb{N}$ of indices induced by an m-i.i.d. sequence of reals, where m is any nondegenerate continuous probability measure on $\mathbb{R}$. The invariant measures on the remaining four reducts may be obtained in a similar way from an i.i.d. sequence on the respective reduct of $\mathbb{R}$.

For example, for the countable dense circular order, the (unique) invariant measure can be obtained as either the weak limit of the uniform measure on circular orders of size n with the (ternary) clockwise-order relation, or from the ternary relation induced on the set $\mathbb{N}$ of indices by the clockwise-ordering of an m-i.i.d. sequence, where m is a nondegenerate continuous probability measure on the unit circle.

Note that the existence of an invariant measure on the orbit of each ultrahomogeneous highly homogeneous structure $\mathcal{M}$ is a consequence of Exercise 5 of [Cam90, §4.10]; this exercise implies that the weak limit of uniform measures on n-element substructures of $\mathcal{M}$ is invariant and concentrated on the orbit of $\mathcal{M}$.

After the following lemma, we will be able to prove that every ultrahomogeneous highly homogeneous structure admits a unique invariant measure. Write S_n to denote the group of permutations of $\{0, 1, \ldots, n-1\}$.

Lemma 3.2. *Let $\mathcal{M} \in \mathrm{Str}_L$. If $\mathcal{M}$ is highly homogeneous, then there is at most one invariant measure on the isomorphism class of $\mathcal{M}$ in Str_L.*

Proof. Let $n \in \mathbb{N}$ and let p be a qf-type of $\mathcal{L}_{\omega_1,\omega}(L)$ in n variables that is realized in $\mathcal{M}$. Because $\mathcal{M}$ is highly homogeneous, for any qf-type q of $\mathcal{L}_{\omega_1,\omega}(L)$ in n variables that is realized in $\mathcal{M}$, there is some $\tau \in S_n$ such that

$$\mathcal{M} \models (\forall x_0 \cdots x_{n-1})\ \big(p(x_0, \ldots, x_{n-1}) \leftrightarrow q(x_{\tau(0)}, \ldots, x_{\tau(n-1)})\big).$$

Suppose μ is an invariant measure on Str_L concentrated on the orbit of $\mathcal{M}$. Then for any $k_0, \ldots, k_{n-1} \in \mathbb{N}$, we have

$$\mu\big(\llbracket p(k_0, \ldots, k_{n-1})\rrbracket\big) = \mu\big(\llbracket q(k_{\tau(0)}, \ldots, k_{\tau(n-1)})\rrbracket\big).$$

By the S_∞-invariance of μ, we have

$$\mu\big(\llbracket q(k_{\tau(0)}, \ldots, k_{\tau(n-1)})\rrbracket\big) = \mu\big(\llbracket q(k_0, \ldots, k_{n-1})\rrbracket\big).$$

Let α_n be the number of distinct qf-types of $\mathcal{L}_{\omega_1,\omega}(L)$ in n-many variables that are realized in $\mathcal{M}$. Note that $\alpha_n \leq n!$ by the high homogeneity of $\mathcal{M}$. Then

$$\mu\big(\llbracket p(k_0, \ldots, k_{n-1})\rrbracket\big) = \frac{1}{\alpha_n}.$$

Sets of the form $\llbracket p(k_0, \ldots, k_{n-1})\rrbracket$ generate the σ-algebra of Borel subsets of the isomorphism class of $\mathcal{M}$ in Str_L, and so μ must be the unique measure determined in this way. $\square$

Putting the previous two results together, we obtain the following.

Proposition 3.3. *Let L be relational and let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous. If $\mathcal{M}$ is highly homogeneous, then there is a unique invariant measure on the isomorphism class of $\mathcal{M}$ in Str_L.*

Proof. By Lemma 3.1, there is an invariant measure on the isomorphism class of $\mathcal{M}$ in Str_L. On the other hand, by Lemma 3.2, this is the only invariant measure on the isomorphism class of $\mathcal{M}$ in Str_L. $\square$

3.2. **Continuum-many ergodic invariant measures.** We now show that when a countable ultrahomogeneous structure in a relational language admits an invariant measure but is not highly homogeneous, there are continuum-many ergodic invariant measures on its orbit. We do this by constructing a continuum-sized class of *reweighted* measures $m^{\mathcal{W}}$ that give rise to distinct measures $\mu_{(\mathcal{P},m^{\mathcal{W}})}$ on the orbit of the structure, for some appropriate $\mathcal{P}$. This will allow us to complete the proof of our main result, Theorem 1.1. We start with some definitions.

Definition 3.4. A **partition** of $\mathbb{R}$ is a collection of subsets of $\mathbb{R}$ that are non-overlapping and whose union is $\mathbb{R}$. By **half-open interval**, we mean a non-empty, left-closed, right-open interval of $\mathbb{R}$, including the cases $\mathbb{R}$, $(-\infty, c)$, and $[c, \infty)$ for $c \in \mathbb{R}$. A **weight** $\mathcal{W}$ consists of a partition of $\mathbb{R}$ into a finite set $\mathcal{I}_{\mathcal{W}}$ of finite unions of half-open intervals, along with a map $u_{\mathcal{W}} \colon \mathcal{I}_{\mathcal{W}} \to \mathbb{R}^+$ that assigns a positive real number to each element of $\mathcal{I}_{\mathcal{W}}$ and satisfies

$$\sum_{I \in \mathcal{I}_{\mathcal{W}}} u_{\mathcal{W}}(I) = 1.$$

Given a measure m on $\mathbb{R}$, the **reweighting** $m^{\mathcal{W}}$ of m by a weight $\mathcal{W}$ is the measure on $\mathbb{R}$ defined by

$$m^{\mathcal{W}}(B) = \sum_{I \in \mathcal{I}_{\mathcal{W}}} u_{\mathcal{W}}(I)\, \frac{m(B \cap I)}{m(I)}$$

for all Borel sets $B \subseteq \mathbb{R}$.

The following is immediate from the definition of a weight.

Lemma 3.5. *Let m be a nondegenerate continuous probability measure on $\mathbb{R}$, and let $\mathcal{W}$ be a weight. Then $m^{\mathcal{W}}$, the reweighting of m by $\mathcal{W}$, is also a nondegenerate continuous probability measure on $\mathbb{R}$.*

We then obtain the following corollary.

Corollary 3.6. *Let L be relational, let $\mathcal{P}$ be a Borel L-structure that strongly witnesses a pithy Π_2 theory T, and let m be a nondegenerate continuous probability measure on $\mathbb{R}$. Let $\mathcal{W}$ be a weight. Then $\mu_{(\mathcal{P},m^{\mathcal{W}})}$ is concentrated on the set of structures in Str_L that are models of T, just as $\mu_{(\mathcal{P},m)}$ is.*

Proof. Let L, $\mathcal{P}$, T, m, and $\mathcal{W}$ be as stated. By Lemma 3.5, $m^{\mathcal{W}}$ is also a nondegenerate continuous probability measure. Therefore, because $\mathcal{P}$ strongly witnesses T, both $(\mathcal{P}, m)$ and $(\mathcal{P}, m^{\mathcal{W}})$ witness T. Hence both $\mu_{(\mathcal{P},m)}$ and $\mu_{(\mathcal{P},m^{\mathcal{W}})}$ are concentrated on the set of models of T in Str_L by Proposition 2.26. ☐

We now show that when $\mathcal{M}$ is not highly homogeneous but admits an invariant measure, reweighting can be used to obtain continuum-many ergodic invariant measures on the isomorphism class of $\mathcal{M}$ in Str_L. Specifically, suppose L is relational, $\mathcal{M} \in \mathrm{Str}_L$ is ultrahomogeneous, and T is the Fraïssé theory of $\mathcal{M}$. Then as $\mathcal{W}$ ranges over weights, we will see that there are continuum-many measures $\mu_{(\mathcal{P},m^{\mathcal{W}})}$, where $\mathcal{P}$ is a Borel L-structure that strongly witnesses T and m is a nondegenerate continuous probability measure on $\mathbb{R}$.

We start with two technical results. Recall that S_n is the group of permutations of $\{0, 1, \ldots, n-1\}$.

Lemma 3.7. *Fix $n, \ell \in \mathbb{N}$. Suppose $\{a_s\}_{s\in\{0,1,\ldots,\ell\}^n}$ is a collection of non-negative reals with the following properties:*

(a) *For each $\sigma \in S_n$ and $s, t \in \{0, 1, \ldots, \ell\}^n$, if $s \circ \sigma = t$ then $a_s = a_t$.*
(b) *For some $s, t \in \{0, 1, \ldots, \ell\}^n$, we have $a_s \neq a_t$.*

Then as the variables $\lambda_0, \ldots, \lambda_\ell$ range over positive reals such that

$$\lambda_0 + \cdots + \lambda_\ell = 1,$$

the polynomial

$$\sum_{s\in\{0,1,\ldots,\ell\}^n} a_s \lambda_{s(0)} \cdots \lambda_{s(n-1)} \tag{$\spadesuit$}$$

assumes continuum-many values.

Proof. In ($\spadesuit$) substitute $1 - \sum_{i=0}^{\ell-1} \lambda_i$ for λ_ℓ to obtain a polynomial P in ℓ-many variables $\lambda_0, \ldots, \lambda_{\ell-1}$. We will show that P is a non-constant polynomial, and therefore assumes continuum-many values as $\lambda_0, \ldots, \lambda_{\ell-1}$ range over positive reals such that

$$\lambda_0 + \cdots + \lambda_{\ell-1} < 1.$$

Suppose towards a contradiction that P is a constant polynomial. Let $a^* := a_u$, where $u \in \{0, 1, \ldots, \ell\}^n$ is the constant function taking the value ℓ. Consider, for $0 \leq k \leq n$, the following claim ($\diamondsuit_k$).

($\diamondsuit_k$) For every j such that $0 \leq j \leq k$, whenever $s \in \{0, 1, \ldots, \ell\}^n$ is such that exactly j-many of $s(0), \ldots, s(n-1)$ are different from ℓ, then $a_s = a^*$.

The statement ($\diamondsuit_n$) implies that for every $s \in \{0, 1, \ldots, \ell\}^n$, we have $a_s = a^*$, thereby contradicting (b). Hence it suffices to prove ($\diamondsuit_n$), which we now do by induction on k.

The statement ($\diamondsuit_0$) is clear. Now let k be such that $1 \leq k \leq n$, and suppose that ($\diamondsuit_{k-1}$) holds. We will show that ($\diamondsuit_k$) holds. Let $s \in \{0, 1, \ldots, \ell\}^n$ be such that exactly k-many of $s(0), \ldots, s(n-1)$ are different from ℓ; we must prove that $a_s = a^*$.

Since ($\diamondsuit_{k-1}$) holds, by (a) we may assume without loss of generality that none of $s(0), s(1), \ldots, s(k-1)$ equals ℓ and that

$$s(k) = s(k+1) = \cdots = s(n-1) = \ell.$$

For $0 \leq r \leq \ell - 1$ let k_r denote the number of times that r appears in the sequence $s(0), \ldots, s(k-1)$. In particular,

$$\lambda_{s(0)} \lambda_{s(1)} \cdots \lambda_{s(k-1)} = \lambda_0^{k_0} \lambda_1^{k_1} \cdots \lambda_{\ell-1}^{k_{\ell-1}},$$

and $k = k_0 + k_1 + \cdots + k_{\ell-1}$.

Let β be the coefficient of $\lambda_0^{k_0} \lambda_1^{k_1} \cdots \lambda_{\ell-1}^{k_{\ell-1}}$ in P. For $t_0, \ldots, t_{\ell-1} \in \mathbb{N}$ such that $t_0 + \cdots + t_{\ell-1} \leq n$, let $\Gamma(t_0, t_1, \ldots, t_{\ell-1}) \in \{0, 1, \ldots, \ell\}^n$ be the non-decreasing sequence of length n consisting of t_0-many 0's, t_1-many 1's, ..., $t_{\ell-1}$-many $t_{\ell-1}$'s, and $(n - \sum_{i=0}^{\ell-1} t_i)$-many ℓ's. Define $C := \frac{n!}{k_0! k_1! \cdots k_{\ell-1}! (n-k)!}$. Then

$$\beta = C \sum_{t_i \leq k_i} a_{\Gamma(t_0, t_1, \ldots, t_{\ell-1})} \binom{k_0}{t_0} \cdots \binom{k_{\ell-1}}{t_{\ell-1}} (-1)^{k - \sum_{i=0}^{\ell-1} t_i}. \tag{$\heartsuit$}$$

Note that $a_s = a_{\Gamma(k_0, k_1, \ldots, k_{\ell-1})}$. By $(\diamondsuit_{k-1})$, we also have $a^* = a_{\Gamma(t_0, t_1, \ldots, t_{\ell-1})}$ if $\sum_{i=0}^{\ell-1} t_i < k$, in particular whenever each $t_i \leq k_i$ for $0 \leq i \leq \ell - 1$ and $(t_0, t_1, \ldots, t_{\ell-1}) \neq (k_0, k_1, \ldots, k_{\ell-1})$. In other words, all subexpressions $a_{\Gamma(t_0, t_1, \ldots, t_{\ell-1})}$ appearing in ($\heartsuit$) other than (possibly) $a_{\Gamma(k_0, k_1, \ldots, k_{\ell-1})}$ are equal to a^*.

By the multinomial and binomial theorems,

$$\begin{aligned} \sum_{t_i \leq k_i} \binom{k_0}{t_0} \cdots \binom{k_{\ell-1}}{t_{\ell-1}} (-1)^{k - \sum_{i=0}^{\ell-1} t_i} &= \sum_{t \leq k} \sum_{\substack{t_i \leq k_i \\ \sum_i t_i = t}} \binom{k_0}{t_0} \cdots \binom{k_{\ell-1}}{t_{\ell-1}} (-1)^{k - \sum_{i=0}^{\ell-1} t_i} \\ &= \sum_{t \leq k} \binom{k}{t} (-1)^{k-t} = (1-1)^k = 0. \end{aligned}$$

Therefore

$$\beta = 0 + (a_s - a^*) \binom{k_0}{k_0} \cdots \binom{k_{\ell-1}}{k_{\ell-1}} (-1)^{k-k} = a_s - a^*.$$

But by the assumption that P is a constant polynomial, $\beta = 0$, and so $a_s = a^*$, as desired. □

Using this lemma, we can prove the following.

Proposition 3.8. *Let m be a nondegenerate continuous probability measure, and let n be a positive integer. Suppose $A \subseteq \mathbb{R}^n$ is an S_n-invariant Borel set such that $0 < m^n(A) < 1$. Then the family of reals $\{(m^{\mathcal{W}})^n(A) \colon \mathcal{W}$ is a weight$\}$ has cardinality equal to the continuum.*

Proof. Because $m^n(A) > 0$, we may define $\widetilde{m}$, the conditional distribution of m^n given A, by

$$\widetilde{m}(B) := \frac{m^n(A \cap B)}{m^n(A)}$$

for every Borel set $B \subseteq \mathbb{R}^n$.

Because $m^n(A) < 1$, we have $m^n(\mathbb{R}^n - A) = 1 - m^n(A) > 0$. Furthermore $\widetilde{m}(\mathbb{R}^n - A) = 0$, and so $m^n \neq \widetilde{m}$. Therefore there are some half-open intervals $X_0, X_1, \dots, X_{n-1}$ such that

$$m^n\big(\prod_{i=0}^{n-1} X_i\big) \neq \widetilde{m}\big(\prod_{i=0}^{n-1} X_i\big).$$

Because m is nondegenerate, $m(X_i) > 0$ for each $i \leq n-1$.

Define the partition $\mathcal{J}$ of $\mathbb{R}$ to be the family of non-empty sets of the form $X_0^{e_0} \cap \ldots \cap X_{n-1}^{e_{n-1}}$ for some $e_0, \dots, e_{n-1} \in \{+, -\}$, where $X_j^+ := X_j$ and $X_j^- := \mathbb{R} - X_j$. Let $\ell := |\mathcal{J}| - 1$, and let $Y_0, \dots, Y_\ell$ be some enumeration of $\mathcal{J}$.

For $s \in \{0, 1, \dots, \ell\}^n$, set

$$a_s := \frac{m^n\big(A \cap \prod_{i=0}^{n-1} Y_{s(i)}\big)}{m^n\big(\prod_{i=0}^{n-1} Y_{s(i)}\big)}.$$

Note that there exists some $v \in \{0, 1, \dots, \ell\}^n$ such that

$$m^n\big(\prod_{i=0}^{n-1} Y_{v(i)}\big) \neq \widetilde{m}\big(\prod_{i=0}^{n-1} Y_{v(i)}\big),$$

i.e.,

$$m^n\big(\prod_{i=0}^{n-1} Y_{v(i)}\big) \neq \frac{m^n\big(A \cap \prod_{i=0}^{n-1} Y_{v(i)}\big)}{m^n(A)},$$

and hence

$$m^n(A) \neq \frac{m^n\big(A \cap \prod_{i=0}^{n-1} Y_{v(i)}\big)}{m^n\big(\prod_{i=0}^{n-1} Y_{v(i)}\big)} = a_v.$$

Observe that if for $s \in \{0, 1, \dots, \ell\}^n$, the values a_s are all equal, then this value is $m^n(A)$. But we have just shown that $a_v \neq m^n(A)$, and so $a_s \neq a_t$ for some s, t. Further, since A is S_n-invariant, from the definition of a_s we have that for every $\sigma \in S_n$, and every s and t, if $s \circ \sigma = t$ then $a_s = a_t$.

Hence the assumptions of Lemma 3.7 are satisfied, and so the expression

$$\sum_{s \in \{0,1,\dots,\ell\}^n} a_s \lambda_{s(0)} \cdots \lambda_{s(n-1)}$$

takes continuum-many values as $\lambda_0, \dots, \lambda_\ell$ range over positive reals satisfying $\lambda_0 + \cdots + \lambda_\ell = 1$. Each such $\lambda_0, \dots, \lambda_\ell$ together with the partition $\mathcal{J}$ yields a weight $\mathcal{W}$ via $u_{\mathcal{W}}(Y_i) := \lambda_i$ for $i \leq \ell$. Then the corresponding reweightings $m^{\mathcal{W}}$ satisfy

$$\begin{aligned} (m^{\mathcal{W}})^n(A) &= \sum_{s \in \{0,1,\dots,\ell\}^n} \frac{m^n\big(A \cap \prod_{i=0}^{n-1} Y_{s(i)}\big)}{m^n\big(\prod_{i=0}^{n-1} Y_{s(i)}\big)} \lambda_{s(0)} \cdots \lambda_{s(n-1)} \\ &= \sum_{s \in \{0,1,\dots,\ell\}^n} a_s \lambda_{s(0)} \cdots \lambda_{s(n-1)}. \end{aligned}$$

We conclude that the family $\{(m^{\mathcal{W}})^n(A) \colon \mathcal{W} \text{ is a weight}\}$ has cardinality equal to the continuum. $\square$

Now we may prove our main result about weights.

Proposition 3.9. *Let L be relational, let $\mathcal{M} \in \mathrm{Str}_L$ be ultrahomogeneous, and let T be the Fraïssé theory of $\mathcal{M}$. Suppose that $\mathcal{M}$ is not highly homogeneous. Further suppose that $\mathcal{P}$ is a Borel L-structure that strongly witnesses T, and that m is a nondegenerate continuous probability measure on $\mathbb{R}$. Then there are continuum-many measures $\mu_{(\mathcal{P},m^{\mathcal{W}})}$, as $\mathcal{W}$ ranges over weights.*

Proof. By Proposition 2.26, $\mu_{(\mathcal{P},m)}$ is concentrated on the orbit of $\mathcal{M}$. Because $\mathcal{M}$ is not highly homogeneous, by Lemma 2.19 there are non-isomorphic n-element substructures A_0, A_1 of $\mathcal{M}$ for some $n \in \mathbb{N}$. Fix an enumeration of each of A_0, A_1, and for $i = 0, 1$ let φ_i be a quantifier-free $\mathcal{L}_{\omega_1,\omega}(L)$-formula in n-many free variables that is satisfied by A_i and not by A_{1-i} (in their respective enumerations). Note that

$$0 < \mu_{(\mathcal{P},m)}\Big(\llbracket \bigvee_{\sigma\in S_n} \varphi_i(\sigma(0),\dots,\sigma(n-1))\rrbracket\Big)$$

for $i = 0, 1$, as φ_i is realized in $\mathcal{M}$. Furthermore, as

$$\mu_{(\mathcal{P},m)}\Big(\llbracket \bigvee_{\sigma\in S_n} \varphi_0(\sigma(0),\dots,\sigma(n-1))\rrbracket\Big) + \mu_{(\mathcal{P},m)}\Big(\llbracket \bigvee_{\sigma\in S_n} \varphi_1(\sigma(0),\dots,\sigma(n-1))\rrbracket\Big) \leq 1,$$

we have

$$\mu_{(\mathcal{P},m)}\Big(\llbracket \bigvee_{\sigma\in S_n} \varphi_i(\sigma(0),\dots,\sigma(n-1))\rrbracket\Big) < 1$$

for $i = 0, 1$.

Then, by Lemma 2.29, we have

$$0 < m^n\Big(\big\{(a_0,\dots,a_{n-1}) \in \mathbb{R}^n : \mathcal{P} \models \bigvee_{\sigma\in S_n} \varphi_0(a_{\sigma(0)},\dots,a_{\sigma(n-1)})\big\}\Big) < 1.$$

Hence by Proposition 3.8, as $\mathcal{W}$ ranges over weights,

$$(m^{\mathcal{W}})^n\Big(\big\{(a_0,\dots,a_{n-1}) \in \mathbb{R}^n : \mathcal{P} \models \bigvee_{\sigma\in S_n} \varphi_0(a_{\sigma(0)},\dots,a_{\sigma(n-1)})\big\}\Big)$$

takes on continuum-many values. Again by Lemma 2.29,

$$\mu_{(\mathcal{P},m^{\mathcal{W}})}\Big(\llbracket \bigvee_{\sigma\in S_n} \varphi_i(\sigma(0),\dots,\sigma(n-1))\rrbracket\Big)$$

takes on continuum-many values as $\mathcal{W}$ ranges over weights; in particular, the $\mu_{(\mathcal{P},m^{\mathcal{W}})}$ constitute continuum-many different measures. □

We are now able to complete the proof of our main theorem.

Proof of Theorem 1.1. Given a countable structure $\mathcal{N}$ in a countable language, by Corollaries 2.9 and 2.11 and Lemma 2.17, its canonical structure $\overline{\mathcal{N}}$ admits the same number of ergodic invariant measures as $\mathcal{N}$, is highly homogeneous if and only if $\mathcal{N}$ is, and has trivial definable closure if and only if $\mathcal{N}$ does. Hence it suffices to prove the theorem in the case where $\mathcal{M} \in \mathrm{Str}_L$ is the canonical structure of some countable structure in a countable language; in particular, where $\mathcal{M}$ is ultrahomogeneous (by Proposition 2.13) and L is relational (by Definition 2.6).

By Theorem 2.4, if $\mathcal{M}$ has nontrivial definable closure then its orbit does not admit an invariant measure, as claimed in (0).

By Proposition 3.3, if $\mathcal{M}$ is highly homogeneous then its orbit admits a unique invariant measure, as claimed in (1).

Clearly, the orbit of $\mathcal{M}$ admitting 0, 1, or continuum-many invariant measures are mutually exclusive possibilities. Hence it remains to show that if $\mathcal{M}$ is not highly homogeneous and its orbit admits an invariant measure, then this orbit admits continuum-many ergodic invariant measures.

Again by Theorem 2.4, because the orbit of $\mathcal{M}$ admits an invariant measure, $\mathcal{M}$ must have trivial definable closure. Since $\mathcal{M}$ is ultrahomogeneous and L is relational, by Proposition 2.27 there is a Borel L-structure $\mathcal{P}$ that strongly witnesses the Fraïssé theory of $\mathcal{M}$.

Let m be a nondegenerate continuous probability measure on $\mathbb{R}$. By Proposition 3.9, as $\mathcal{W}$ ranges over weights, there are continuum-many different measures $\mu_{(\mathcal{P},m^{\mathcal{W}})}$. By Corollary 3.6, each is an invariant probability measure concentrated on the orbit of $\mathcal{M}$, and by Proposition 2.24, each is ergodic. Finally, there are at most continuum-many Borel measures on Str_L. $\square$

Acknowledgments

The authors would like to thank Willem Fouché, Yonatan Gutman, Alexander Kechris, André Nies, Arno Pauly, Jan Reimann, and Carol Wood for helpful conversations.

This research was facilitated by the Dagstuhl Seminar on Computability, Complexity, and Randomness (January 2012), the conference on Graphs and Analysis at the Institute for Advanced Study (June 2012), the Buenos Aires Semester in Computability, Complexity, and Randomness (March 2013), the Arbeitsgemeinschaft on Limits of Structures at the Mathematisches Forschungsinstitut Oberwolfach (March–April 2013), the Trimester Program on Universality and Homogeneity of the Hausdorff Research Institute for Mathematics at the University of Bonn (September–December 2013), and the workshop on Analysis, Randomness, and Applications at the University of South Africa (February 2014).

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 John Templeton Foundation.

References

[AFP16] N. Ackerman, C. Freer, and R. Patel, *Invariant measures concentrated on countable structures*, Forum Math. Sigma **4** (2016), no. e17, 59 pp.

[AKL14] O. Angel, A. S. Kechris, and R. Lyons, *Random orderings and unique ergodicity of automorphism groups*, J. Eur. Math. Soc. **16** (2014), no. 10, 2059–2095.

[Aus08] T. Austin, *On exchangeable random variables and the statistics of large graphs and hypergraphs*, Probab. Surv. **5** (2008), 80–145.

[BK96] H. Becker and A. S. Kechris, *The descriptive set theory of Polish group actions*, London Math. Soc. Lecture Note Series, vol. 232, Cambridge Univ. Press, Cambridge, 1996.

[Cam76] P. J. Cameron, *Transitivity of permutation groups on unordered sets*, Math. Z. **148** (1976), no. 2, 127–139.

[Cam90] ———, *Oligomorphic permutation groups*, London Math. Soc. Lecture Note Ser., vol. 152, Cambridge Univ. Press, Cambridge, 1990.

[ER59] P. Erdős and A. Rényi, *On random graphs. I*, Publ. Math. Debrecen **6** (1959), 290–297.

[Fou12] W. L. Fouché, *Algorithmic randomness and Ramsey properties of countable homogeneous structures*, Logic, Language, Information and Computation, Lecture Notes in Comput. Sci., vol. 7456, Springer, Heidelberg, 2012, pp. 246–256.

[Fou13] ———, *Martin-Löf randomness, invariant measures and countable homogeneous structures*, Theory Comput. Syst. **52** (2013), no. 1, 65–79.

[FP98] W. L. Fouché and P. H. Potgieter, *Kolmogorov complexity and symmetric relational structures*, J. Symbolic Logic **63** (1998), no. 3, 1083–1094.

[Gao09] S. Gao, *Invariant descriptive set theory*, Pure and Applied Mathematics, vol. 293, CRC Press, Boca Raton, FL, 2009.

[GW02] E. Glasner and B. Weiss, *Minimal actions of the group* $\mathbb{S}(\mathbb{Z})$ *of permutations of the integers*, Geom. Funct. Anal. **12** (2002), no. 5, 964–988.

[Jan13] S. Janson, *Graphons, cut norm and distance, couplings and rearrangements*, New York Journal of Mathematics. NYJM Monographs, vol. 4, Univ. at Albany, State Univ. of New York, Albany, NY, 2013.

[Kal99] O. Kallenberg, *Multivariate sampling and the estimation problem for exchangeable arrays*, J. Theoret. Probab. **12** (1999), no. 3, 859–883.

[Kal02] ———, *Foundations of modern probability*, 2nd ed., Probability and its Applications, Springer-Verlag, New York, 2002.

[Kal05] ———, *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.

[LS10] L. Lovász and B. Szegedy, *Regularity partitions and the topology of graphons*, An irregular mind: Szemerédi is 70, Bolyai Soc. Math. Stud., vol. 21, János Bolyai Math. Soc., Budapest, 2010, pp. 415–446.

[Mac11] D. Macpherson, *A survey of homogeneous structures*, Discrete Math. **311** (2011), no. 15, 1599–1634.

[Nie13] A. Nies (editor), *Logic Blog 2012*, ArXiv e-print 1302.3686 (2013).

[Phe01] R. R. Phelps, *Lectures on Choquet's theorem*, 2nd ed., Lecture Notes in Math., vol. 1757, Springer-Verlag, Berlin, 2001.

[PV10] F. Petrov and A. Vershik, *Uncountable graphs and invariant measures on the set of universal countable graphs*, Random Structures Algorithms **37** (2010), no. 3, 389–406.

[Ver03] A. M. Vershik, *The Kolmogorov example (an overview of actions of infinite-dimensional groups with an invariant probability measure)*, Teor. Veroyatnost. i Primenen. **48** (2003), no. 2, 386–391.

Department of Mathematics
Harvard University
Cambridge, MA 02138
USA
E-mail address: nate@math.harvard.edu

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139
USA
E-mail address: freer@mit.edu

Mathematical Institute
University of Bonn
Endenicher Allee 60
53115 Bonn
Germany
E-mail address: akwiatk@math.uni-bonn.de

Franklin W. Olin College of Engineering
Needham, MA 02492
USA
E-mail address: rehana.patel@olin.edu