

SURJUNCTIVITY DOES NOT CHARACTERIZE COSOFICITY OF INVARIANT RANDOM SUBGROUPS

LEWIS BOWEN AND MICHAEL CHAPMAN

ABSTRACT. A group is *surjunctive* if every injective cellular automaton on it is also surjective. Gottschalk famously conjectured that all groups are surjunctive. This remains a central open problem in symbolic dynamics and descriptive set theory. Gromov and Weiss termed the notion of *sofic* groups, and proved that all such groups are surjunctive, providing the largest class of groups which satisfy Gottschalk’s conjecture. It is still open to decide whether all groups are sofic. This became a major open problem in group theory, and is related to other well known problems such as the Aldous–Lyons conjecture in probability theory and to Connes’ embedding problem in the theory of operator algebras.

A complementary natural question to ask is: Does the reverse implication to Gromov and Weiss’ result holds? Namely, are all surjunctive groups sofic? As currently there are no known non-sofic groups, answering this problem in the negative in the category of groups is still out of reach. This paper resolves this problem in the generalized setup of *invariant random subgroups* of free groups (IRSs), where non (co)sofic objects were recently shown to exist by Lubotzky, Vidick and the two authors. Specifically, we prove that there *exists* a surjunctive non (co)sofic IRS, resolving the aforementioned problem in the negative. Our proof uses a complexity theoretic approach, and in particular a recent development due to Manzoor, as well as the theory of Rokhlin entropy developed by Seward and others. As a byproduct of our proof technique, the non (co)sofic IRS we provide satisfies a condition stronger than surjunctivity; it satisfies a version of Seward’s *maximal Rokhlin entropy of Bernoulli Shifts* (RBS) criterion.

1. Introduction

Let Γ be a finitely generated group,¹ Σ a finite set (of colors), and Σ^Γ the Σ -colorings of Γ . A *cellular automaton* on Γ with color palette Σ is a local, element independent recoloring mechanism of Γ ; namely, it is a recoloring map $\Phi: \Sigma^\Gamma \rightarrow \Sigma^\Gamma$, and there are elements $\gamma_1, \dots, \gamma_n$ and a function $\phi: \Sigma^n \rightarrow \Sigma$ such that $\Phi(c)(x) = \phi(c(x\gamma_1), \dots, c(x\gamma_n))$ for every coloring $c: \Gamma \rightarrow \Sigma$ and element $x \in \Gamma$. A group Γ is said to be *surjunctive* if every injective cellular automaton on it is also surjective. A well known open problem due to Gottschalk [Got73] is: Are all finitely generated groups surjunctive? In [Gro99] and [Wei00], Gromov and Weiss defined the notion of a *sofic* group (see Definition 1.4), and showed that all such groups are surjunctive. Subsequently, the following became a major open problem: Are all finitely generated groups sofic? Another natural problem that arises from their work is:

Problem 1.1 ([Gro99, Wei00]). Are all surjunctive groups sofic? Namely, does surjunctivity characterize soficity?

Currently, no non-sofic group is known to exist, so a negative answer to Problem 1.1 is fairly out of reach. But, there is a generalized setup to this one, in which both soficity and surjunctivity are well defined, and also non sofic objects were recently shown to exist [BCLV24, BCV24] — this is the setup of *invariant random subgroups* (IRSs) of free groups.²

¹Often the more generalized setup of discrete countable groups is used. But for the simplicity of various definitions, we prefer to keep the finite generation assumption throughout.

²The study of invariant random subgroups was initiated independently by [Ver10], [AGV14] and [Bow14] (and appeared implicitly in [SZ94]). It has become a fruitful area of research, as several open problems in group theory, number theory and dynamics were resolved by studying IRSs (cf. [Gel18, ABB⁺17, FG23] and the references therein).

An invariant random subgroup of a given group is a probability measure over its subgroups which is conjugate invariant (see Section 1.1 for a detailed definition). Every pair (Γ, S) consisting of a group and a set of generators for it induces a short exact sequence $N \hookrightarrow \mathcal{F} \twoheadrightarrow \Gamma$, where $\mathcal{F} = \mathcal{F}_S$ is the free group with basis S . As $\mathbf{1}_N$, the Dirac measure concentrated at N , is an IRS of $\mathcal{F}$, the mapping $(\Gamma, S) \mapsto \mathbf{1}_N$ embeds the finitely generated groups into IRSs of free groups. This embedding is meaningful in our context, as IRSs can be *cosofic* (see Definition 1.4) and surjunctive (see Definition 1.6), and $\mathbf{1}_N$ is cosofic (respectively, surjunctive³) if and only if $\Gamma \cong \mathcal{F}/N$ is sofic (respectively, surjunctive). Moreover, similar to the result of Gromov and Weiss, one can show (see Corollary 5.4 and Theorem 3.6) that a cosofic IRS is always surjunctive. So, the analogue of Problem 1.1 arises — are all surjunctive IRSs of free groups cosofic? Our main result is a negative answer to this problem:

Theorem 1.2. *There exists a surjunctive non cosofic IRS of a finitely generated free group.*

In the rest of this introduction, we provide more detailed definitions of cosoficity and surjunctivity, as well as describe our proof technique.

1.1. Cosoficity and Surjunctivity of Invariant Random Subgroups. The set $\mathbf{sub}(\mathcal{F})$ consisting of all subgroups of the free group $\mathcal{F}$ is compact when equipped with the product topology. A *random subgroup* is a Borel probability distribution over $\mathbf{sub}(\mathcal{F})$, and we denote the collection of all of them by $\text{Prob}(\mathbf{sub}(\mathcal{F}))$.⁴ The random subgroups can be equipped with the weak* topology, and with this topology, this space is compact. The conjugation action $w.K = wKw^{-1}$ of $\mathcal{F}$ on its subgroups extends naturally to random subgroups.⁵

Definition 1.3 (Invariant Random Subgroups). An *invariant random subgroup* (IRS) of $\mathcal{F}$ is a random subgroup $\pi \in \text{Prob}(\mathbf{sub}(\mathcal{F}))$ such that $w.\pi = \pi$ for every $w \in \mathcal{F}$. The set of all IRSs of $\mathcal{F}$ is denoted by $\text{IRS}(\mathcal{F})$.

For us there will be two main examples of IRSs. First, for every normal subgroup $N \trianglelefteq \mathcal{F}$, the Dirac measure $\mathbf{1}_N$ is an IRS. In addition, whenever $\mathcal{F}$ acts on a finite set Ω , it induces a *finitely described* IRS by choosing a uniformly random stabilizer of this action; namely, letting $\text{Stab}(x) = \{w \in \mathcal{F} \mid w.x = x\}$, the IRS associated with $\mathcal{F} \curvearrowright \Omega$ is

$$(1.1) \quad \pi = \frac{1}{|\Omega|} \sum_{x \in \Omega} \mathbf{1}_{\text{Stab}(x)}.$$

Definition 1.4 (Soficity and Cosoficity). An IRS π is *cosofic* if it is in the weak* closure of the finitely described IRSs, and we denote the set of these by $\text{IRS}_{\text{sof}}(\mathcal{F})$. A finitely generated group Γ is *sofic* if it is isomorphic to $\mathcal{F}/N$ with $\mathcal{F}$ a finitely generated free group, $N \trianglelefteq \mathcal{F}$, and $\mathbf{1}_N$ a cosofic IRS of $\mathcal{F}$.

The Σ -colorings of a group Γ can be equipped with the product topology. In addition, Γ acts naturally on Σ^Γ by shifts; namely for $\gamma \in \Gamma$ and $c: \Gamma \rightarrow \Sigma$, $\gamma.c(x) = c(\gamma^{-1}x)$. With these two observations, a cellular automaton is a recoloring map $\Phi: \Sigma^\Gamma \rightarrow \Sigma^\Gamma$ which is continuous and commutes with the group action, i.e., $\gamma.\Phi(c) = \Phi(\gamma.c)$ for every $\gamma \in \Gamma$ — this property is often called *conjugacy invariance* or Γ -*equivariance*. We now generalize this to IRSs. Given a subgroup $K \leq \mathcal{F}$, a coloring $c: \mathcal{F} \rightarrow \Sigma$ is K -invariant if $k.c = c$ for all $k \in K$ — this data can also be given as a Σ -coloring of the right K -cosets $K \backslash \mathcal{F} = \{Kx \mid x \in \mathcal{F}\}$. Let $\mathbf{sub}(\mathcal{F}, \Sigma)$ be the set

³It would be more natural to call this property *cosurjunctivity*, but we decided to avoid this.

⁴It is often the case that a random subgroup is a random variable with values in subgroups of $\mathcal{F}$ that is distributed according to some law. We do not use this perspective at all along the paper, and a random subgroup (and later random subset) is the distribution law itself.

⁵Because H will be used throughout this paper to denote entropy (of various kinds), we avoid using this notation for subgroups and use K for subgroups instead.

of all Σ -colorings of K -cosets of $\mathcal{F}$ for any subgroup $K \leq \mathcal{F}$; namely, it consists of all pairs (K, c) where K is a subgroup of $\mathcal{F}$ and $c: \mathcal{F} \rightarrow \Sigma$ is K -invariant. As $\mathbf{sub}(\mathcal{F}, \Sigma) \subseteq \mathbf{sub}(\mathcal{F}) \times \Sigma^{\mathcal{F}}$, it can be equipped with the product topology. In addition, the diagonal action $w.(K, c) = (wKw^{-1}, w.c)$ of $\mathcal{F}$ on $\mathbf{sub}(\mathcal{F}) \times \Sigma^{\mathcal{F}}$ preserves $\mathbf{sub}(\mathcal{F}, \Sigma)$ — as, c being K -invariant implies $w.c$ is wKw^{-1} -invariant. For π an IRS of $\mathcal{F}$, let $\text{supp}(\pi) \subseteq \mathbf{sub}(\mathcal{F})$ be its support — i.e., subgroups whose neighborhoods all have positive π -measure. Define the *domain of π* , $\text{Dom}(\pi) \subseteq \mathbf{sub}(\mathcal{F}, \Sigma)$, to consist of all pairs (K, c) where $K \in \text{supp}(\pi)$ — namely, all colorings of right K -cosets of subgroups that π can “actually sample”. Since π is an IRS, $\text{Dom}(\pi)$ is preserved under the action of $\mathcal{F}$. Finally, let $P: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \mathbf{sub}(\mathcal{F})$ be the projection to the first coordinate.

Definition 1.5 (Cellular automaton on an IRS). A function $\Phi: \text{Dom}(\pi) \rightarrow \text{Dom}(\pi)$ is a cellular automaton on π (with color palette Σ) if it is continuous, conjugate invariant (i.e., $w\Phi w^{-1}(K, c) = \Phi(K, c)$ for every $(K, c) \in \text{Dom}(\pi)$ and $w \in \mathcal{F}$) and $P \circ \Phi = P$ (i.e., the Φ -image of (K, c) must be (K, c') , so Φ indeed recolors the right cosets of K for every K in the support of π).

Definition 1.6 (Surjunctive IRS). A $\pi \in \text{IRS}(\mathcal{F})$ is *surjunctive* if every (almost everywhere) injective π -cellular automaton must be surjective.

Remark 1.7. A cellular automaton on an IRS π is a generalization of those defined over finitely generated groups (when taking $\pi = \mathbf{1}_N$ with $\Gamma \cong \mathcal{F}/N$). In addition, $\mathbf{1}_N$ is surjunctive as in Definition 1.6 if and only if $\mathcal{F}/N$ is a surjunctive group.

The reader can now parse Theorem 1.2.

1.2. Proof technique. Describing our proof requires three steps:

- (1) First, we need to recall the argument structure of [BCLV24, BCV24], where non cosofic IRSs were shown to exist.
- (2) Then, we need to describe Manzoor’s recent work [Man25a], in which he used a certain amplification of the above argument structure to deduce the existence of non cosofic IRSs which satisfy Lück’s determinant conjecture [Lüc02].
- (3) Finally, to implement the appropriate amplification a la Manzoor, we use an IRS version of a sufficient condition to surjunctivity due to Seward [Sew19b].

1.2.1. Subgroup Tests and non cosofic IRSs. A *challenge* is a continuous function D from $\mathbf{sub}(\mathcal{F})$ to $\{0, 1\}$. A *subgroup test* $\mathcal{T}$ is a rational probability distribution over finitely many challenges, namely it consists of challenges $D_1, \dots, D_m$ and $\mu \in \text{Prob}([m]) \cap \mathbb{Q}^m$. The *value* of a random subgroup π versus a subgroup test $\mathcal{T}$ is

$$(1.2) \quad \text{val}(\mathcal{T}, \pi) = \mathbb{E}_{K \sim \pi} \mathbb{E}_{i \sim \mu} [D_i(K)].$$

It is immediate that $\text{val}(\mathcal{T}, \cdot)$ is a continuous linear functional on the random subgroups. The sofic value $\text{val}_{\text{sof}}(\mathcal{T})$ of a subgroup test $\mathcal{T}$ is the supremum of $\text{val}(\mathcal{T}, \pi)$ over the finitely described IRSs, and its ergodic value $\text{val}_{\text{erg}}(\mathcal{T})$ is the supremum of $\text{val}(\mathcal{T}, \pi)$ over all IRSs. By the continuity of $\text{val}(\mathcal{T}, \cdot)$, if every IRS is cosofic, then $\text{val}_{\text{sof}}(\mathcal{T}) = \text{val}_{\text{erg}}(\mathcal{T})$. Given $\mathcal{T}$ and a finitely described IRS π , one can calculate $\text{val}(\mathcal{T}, \pi)$. As there are countably many finitely described IRSs, one can list them $\{\pi_i\}_{i=1}^{\infty}$ and evaluate $\text{val}(\mathcal{T}, \pi_i)$ sequentially. The limsup of this sequence is $\text{val}_{\text{sof}}(\mathcal{T})$, and thus:




FIGURE 1.1. Outer and inner algorithmic approximations, together with an incomputability result, imply that the two sets are different.

Fact 1.8. *There is an algorithm that approximates $\text{val}_{\text{sof}}(\mathcal{T})$, the sofic value of a subgroup test $\mathcal{T}$, from below.*

On the other hand, in [BCLV24] we provided a downwards sequence of polytopes (with finitely many faces) $\{\mathcal{Q}_r\}_{r=1}^{\infty}$ in $\text{Prob}(\text{sub}(\mathcal{F}))$ whose intersection is $\text{IRS}(\mathcal{F})$.⁶ Using linear programming, one can evaluate $\beta_r = \sup_{\pi \in \mathcal{Q}_r} \text{val}(\mathcal{T}, \pi)$ for each r in finite time. Since $\liminf_{r \rightarrow \infty} \beta_r = \text{val}_{\text{erg}}(\mathcal{T})$, we have:

Fact 1.9. *There is an algorithm that approximates $\text{val}_{\text{erg}}(\mathcal{T})$, the ergodic value of a subgroup test $\mathcal{T}$, from above.*⁷

Combining Facts 1.8 and 1.9, if $\text{val}_{\text{sof}}(\mathcal{T}) = \text{val}_{\text{erg}}(\mathcal{T})$ for every $\mathcal{T}$, then $\text{val}_{\text{sof}}(\mathcal{T})$ can be calculated up to any predetermined accuracy. A theorem in [BCLV24] states that:

Fact 1.10. *The following task is as hard as the Halting Problem: Given a subgroup test $\mathcal{T}$ and a positive rational ε , output a rational v such that $|\text{val}_{\text{sof}}(\mathcal{T}) - v| \leq \varepsilon$.*

Combining all of the above, one deduces that there must be a subgroup test $\mathcal{T}$ satisfying $\text{val}_{\text{sof}}(\mathcal{T}) < \text{val}_{\text{erg}}(\mathcal{T})$, which in turn implies non cosofic IRSs exist. See Figure 1.1 for a visualization of this setup.

1.2.2. Manzoor’s amplification of the complexity theoretic framework. Let $\mathcal{P}$ be some property satisfied by all cosofic IRSs. Let $\text{IRS}_{\mathcal{P}}(\mathcal{F})$ be the set of all IRSs that satisfy this property. Then, this set sits in between the cosofic IRSs and all IRSs — see Figure 1.2. If this set has a natural outer approximation, one deduces immediately from Facts 1.10 and 1.8 that there **exist** non cosofic IRSs **which satisfy** this property. This is how Manzoor [Man25a] found a non cosofic IRS which satisfies Lück’s determinant conjecture [Lüc02]. Hence, it is natural to choose $\text{IRS}_{\text{surj}}(\mathcal{F})$ to be the set of all surjunctive IRSs, and find an outer approximation of this set using polytopes. Unfortunately, we do not know how to do that. But, we can find a **subset** of $\text{IRS}_{\text{surj}}(\mathcal{F})$ that has an outer approximation, which will be sufficient to deduce Theorem 1.2.

⁶These polytopes are the pseudo IRSs with respect to finite subsets. The definition is recalled in Section 5.

⁷Algorithms like this one are commonly called “NPA type hierarchies”, due to [NPA08].




FIGURE 1.2. Outer algorithmic approximation for any set of IRSs which satisfy a certain property $\mathcal{P}$, together with the incomputability result for $\text{val}_{\text{sof}}(\mathcal{T})$ from Fact 1.10, imply that there are non cosofic IRSs with this property $\mathcal{P}$.

Remark 1.11. Manzoor suggests in [Man25a] to study both outer and inner approximations for $\text{IRS}_{\mathcal{P}}(\mathcal{F})$. We believe that approximating $\text{val}_{\text{erg}}(\mathcal{T})$ should be impossible (similar to $\text{val}_{\text{sof}}(\mathcal{T})$ in Fact 1.10; the non-local games analogue of this was recently proved in [Lin25]), and if true, inner approximations should yield IRSs that do not satisfy $\mathcal{P}$. Specifically, we neither know whether $\text{IRS}_{\text{surj}}(\mathcal{F})$ has an outer approximation, nor an inner approximation, but if for some reason it has an inner approximation, together with an inapproximability result for $\text{val}_{\text{erg}}(\mathcal{T})$ one will deduce the existence of a non-surjunctive IRS, which would be very exciting.

Remark 1.12. Our Theorem 1.2 can probably be upgraded to the existence of a surjunctive non cohyperlinear IRS (which is stronger than being non cosofic) using the methods of [Man25b, Man25a]. In that case, the combination of the above strengthening of Theorem 1.2 together with Theorem 1.3 in [Man25a] would provide a non cohyperlinear IRS that satisfies both Gottschalk’s conjecture and Lück’s determinant conjecture simultaneously. This shows another beautiful aspect of Manzoor’s approach, as it has a cumulative nature.

1.2.3. Seward’s Rokhlin entropy criterion for surjunctivity. In [Sew19b], Seward defined what it means for a group Γ to satisfy the *maximal Rokhlin entropy on Bernoulli Shifts criterion* (RBS), and proved that a Γ satisfying RBS must be surjunctive. Here, we suggest an IRS analogue of RBS, and prove that indeed every IRS which satisfies RBS is also surjunctive (Theorem 3.6). Similarly, every cosofic IRS satisfies RBS (Corollary 5.4). In summary, if we let $\text{IRS}_{\text{RBS}}(\mathcal{F})$ be the collection of all IRSs satisfying RBS, then we have

$$(1.3) \quad \text{IRS}_{\text{sof}}(\mathcal{F}) \subseteq \text{IRS}_{\text{RBS}}(\mathcal{F}) \subseteq \text{IRS}_{\text{surj}}(\mathcal{F}).$$

To implement Manzoor’s approach, we need to find an outer approximation to $\text{IRS}_{\text{RBS}}(\mathcal{F})$. Given such an outer approximation, by the previous discussion, a non cosofic IRS which satisfies RBS exists, and in turn it will be surjunctive. In Section 5 we provide the required outer approximation.

Let us demystify the RBS criterion a bit. To that end, we describe first the modern approach to Gromov and Weiss’ result that “sofic groups are surjunctive” (cf. Section 5 of [Bow18]). Following the first author’s work in the measurable setup [Bow10], Kerr and Li [KL11] defined the notion of *topological sofic entropy* of an action $\Gamma \curvearrowright X$ of a sofic group on a compact topological space.⁸ This quantity satisfies the following three facts: (1) As discussed before, for every finite set Σ one has the shift action of Γ on Σ^Γ (which is a compact space), and the sofic entropy of this action is $\log |\Sigma|$. (2) In addition, for every closed strict subset $X \subsetneq \Sigma^\Gamma$ which is Γ -invariant, the action $\Gamma \curvearrowright X$ has sofic entropy which is strictly smaller than $\log |\Sigma|$. (3) Moreover, sofic entropy is a *topological invariant*: Given two actions $\Gamma \curvearrowright X$ and $\Gamma \curvearrowright Y$, a (topological) *factor map* between them is a continuous $\Phi: X \rightarrow Y$ which commutes with the group action. In case a factor map Φ is injective, the sofic entropy of $\Gamma \curvearrowright X$ is smaller or equal than that of $\Gamma \curvearrowright Y$. Hence, given a non surjective cellular automaton $\Phi: \Sigma^\Gamma \rightarrow \Sigma^\Gamma$ (which is a factor map by definition), the action of Γ on $\text{Im}(\Phi) \subseteq \Sigma^\Gamma$ has sofic entropy strictly smaller than $\log |\Sigma|$ by clause (2). On the other hand, if Φ is injective, then by clauses (1) and (3) the sofic entropy of this action is at least $\log |\Sigma|$. Hence, an injective cellular automaton on a sofic group must be surjective.

Seward developed in [Sew19a] a way of associating a new type of entropy, which he termed *Rokhlin entropy*, to a probability measure preserving (p.m.p) action of any finitely generated group Γ (not necessarily a sofic one) on a standard probability space (X, μ) , which generalizes the Kolmogorov–Sinai entropy in the amenable case (similar to sofic entropy). By equipping Σ with a probability measure κ , and Σ^Γ with the product measure κ^Γ , the shift action of Γ on Σ^Γ becomes p.m.p — this action is known as the *Bernoulli shift* of Γ with base (Σ, κ) . The measurable setup has the appropriate analogue of a factor map (with continuity replaced by measurability), and the analogue of the above clause (3) is satisfied for it. The RBS criterion is an analogue of (1), namely, a group Γ satisfies this criterion if the Rokhlin entropy of the Bernoulli shift with base (Σ, κ) is equal to the Shannon entropy of the base $H(\kappa) = -\sum_{a \in \Sigma} \kappa(a) \log(\kappa(a))$, which is its maximal possible value — note that in case κ is the uniform distribution u , this quantity is exactly $\log |\Sigma|$, as it was in the sofic case. Under the assumption of RBS, one can also recover a slightly more intricate variant of (2). Finally, one needs to note that a cellular automaton induces a measurable factor map from $(\Sigma^\Gamma, u^\Gamma)$, and by applying the appropriate version of the above argument (see Corollary 4.1 in [Sew19b]), deduce that a group satisfying RBS must be surjunctive.

1.3. Structure of the paper. Section 2 is devoted to the basics of entropy theory as well as the more intricate theory of Rokhlin entropy. Section 3 defines the maximal Rokhlin entropy of Bernoulli Shifts (RBS) criterion, which was originally defined for groups, in the generalized setup of invariant random subgroups (IRSs). Section 4 provides an equivalent criterion to RBS which is more amenable to analysis. Section 5 provides the desired outer approximation result for IRSs which satisfy RBS. Finally, in Section 6 we collect everything and prove Theorem 1.2.

Acknowledgments. We would like to thank Brandon Seward for helpful discussions, and for suggesting to us to use Fact 3.11 to prove Theorem 3.6. We would also like to thank Aareyan Manzoor for providing us with useful remarks that were incorporated to the final presentation, and in particular the need for a rational approximation as devised in Fact 5.5. LB is partially supported by NSF Grant DMS–2453399. MC would like to thank the Institute for Advanced Study for the fruitful environment, and is supported by NSF Grant DMS–2424441.

⁸Actually, the sofic entropy is defined with respect to a specific sofic approximation of the group Γ . But, as this explanation aims to provide only “the ideas”, we ignore this intricacy.

2. Entropy Theory

We need some basic facts from classic probability theory, ergodic theory and entropy theory. An expert could skip most of Section 2, referring back to it for notation if necessary. We use Section 2 of [Sew16] as a guide, and the reader can refer there, as well as to the Appendix of [Sew19b], for any missing information.

2.1. Basic Notions. Let $(X, \mathcal{B}_X, \mu)$ be a standard probability space. Throughout the paper, we assume X is already equipped with a metric which makes it complete and separable (namely, a Polish space) so that $\mathcal{B}_X$ is the Borel sigma-algebra of X with respect to this metric — the notation $\mathcal{B}_X$ for the Borel sigma-algebra of a topological space will be used throughout. We denote by $x \sim \mu$ a variable x sampled according to μ . Given a Borel measurable map between two standard Borel spaces $\psi: (X, \mathcal{B}_X) \rightarrow (Y, \mathcal{B}_Y)$, Y becomes a standard probability space by equipping it with the pushforward measure $\nu = \psi_*\mu$, defined by $\psi_*\mu(B) = \mu(\psi^{-1}(B))$ for every Borel set $B \in \mathcal{B}_Y$. By the disintegration of measures theorem (cf. Theorem 5.14 in [EW11]), for ν -almost every $y \in Y$, there is a Borel probability measure μ_y^ψ on X which is supported on $\psi^{-1}(y)$, and for every Borel measurable map $f: X \rightarrow \mathbb{R}$ we have $\mathbb{E}_{x \sim \mu} [f(x)] = \mathbb{E}_{y \sim \nu} \mathbb{E}_{x \sim \mu_y^\psi} [f(x)]$. The measures μ_y^ψ are the *conditional distributions* $\mu_y^\psi(\cdot) = \mathbb{P}_{x \sim \mu} [x \in \cdot | \psi(x) = y]$. Note that a measurable map $\psi: (X, \mathcal{B}_X) \rightarrow (Y, \mathcal{B}_Y)$ as before induces on X a sub sigma-algebra $\psi^{-1}(\mathcal{B}_Y)$. Conversely, every sub sigma-algebra $\mathcal{F} \subseteq \mathcal{B}_X$ induces a factor map $\psi = \psi_{\mathcal{F}}: (X, \mathcal{B}_X) \rightarrow (Y, \mathcal{B}_Y)$ which satisfies $\psi^{-1}(\mathcal{B}_Y) = \mathcal{F}$ up to μ -null sets. Hence, conditioning over a sub sigma-algebra makes sense. Let (Σ, κ) be a finite or countable probability space. The *Shannon entropy* of κ is $H(\kappa) = -\sum_{a \in \Sigma} \kappa(a) \log(\kappa(a)) \in [0, \infty]$ — this is the expected amount of information (in bits) required to describe the outcome of a sample according to κ .

2.2. Observables and Partitions. Let $(X, \mathcal{B}_X, \mu)$ be a standard probability space, and Σ a finite or countable set equipped with the discrete topology — Σ will often be called the *color palette*, *alphabet* or *indexing set*. An *observable* is a measurable map $\phi: X \rightarrow \Sigma$; the collection of inverse images $\{\phi^{-1}(a) \mid a \in \Sigma\}$ is often called a (measurable) *partition* of X indexed by Σ . By pushforwarding μ along ϕ , one gets a probability distribution $\phi_*\mu$ over Σ ; we denote the Shannon entropy of this measure by $H_\mu(\phi) = H(\phi_*\mu)$, and call this the Shannon entropy of ϕ with respect to μ — this is the expected amount of information gained from learning $\phi(x)$ for $x \sim \mu$. Given a sub sigma algebra $\mathcal{F} \subseteq \mathcal{B}_X$, and letting $\psi: (X, \mathcal{B}_X) \rightarrow (Y, \mathcal{B}_Y)$ be the associated factor and $\nu = \psi_*\mu$ the pushforward measure, the *conditional Shannon entropy* of $\phi: X \rightarrow \Sigma$ with respect to μ given (or relative to) ψ (or $\mathcal{F}$) is $H_\mu(\phi|\psi) = H_\mu(\phi|\mathcal{F}) = \mathbb{E}_{y \sim \nu} [H_{\mu_y^\psi}(\phi)]$ — this is the amount of additional information we expect to receive when $\phi(x)$ is revealed to us given that $x \sim \mu$ and that $\psi(x)$ was revealed to us first.

There are two natural notions of distance between observables. Given two observables ϕ and ψ with the same indexing set Σ , the *total variation* distance between them is $d_{TV}(\phi, \psi) = \frac{1}{2} \sum_{a \in \Sigma} \mu(\phi^{-1}(a) \Delta \psi^{-1}(a))$, where Δ is the symmetric difference of sets — this agrees with the usual total variation distance between $\phi_*\mu$ and $\psi_*\mu$, which measures given $x \sim \mu$, what is the probability that $\phi(x) \neq \psi(x)$. Even without ϕ and ψ having the same indexing set, one can define the *Rokhlin distance* between them by $d_{\text{Rok}}(\phi, \psi) = H_\mu(\phi|\psi) + H_\mu(\psi|\phi)$ — this distance measures how much additional information is gained by revealing one observable after the other. In addition to these distances, there is a natural order induced on the observables by *refinement* and *coarsening* of the partitions: An observable $\phi: X \rightarrow \Sigma$ is a refinement of $\psi: X \rightarrow \Omega$ (or alternatively, ψ is a coarsening of ϕ), denoted by $\phi \geq \psi$, if for every $a \in \Sigma$ there is a $b \in \Omega$ for which $\phi^{-1}(a) \subseteq \psi^{-1}(b)$ up to μ -null sets. The *join* of ϕ and ψ is the observable $\phi \vee \psi: X \rightarrow \Sigma \times \Omega$ given by $\phi \vee \psi(x) = (\phi(x), \psi(x))$. If $\mathcal{F}_1, \mathcal{F}_2$ are two sigma-algebras then their join $\mathcal{F}_1 \vee \mathcal{F}_2$ is the smallest sigma-algebra containing both; this notation is consistent with the one

for observables, as $(\phi \vee \psi)^{-1}(\mathcal{B}_{\Sigma \times \Omega}) = \phi^{-1}(\mathcal{B}_{\Sigma}) \vee \psi^{-1}(\mathcal{B}_{\Omega})$. We often abuse notation and denote $\phi \vee \mathcal{F}$, by which we mean $\phi^{-1}(\mathcal{B}_{\Sigma}) \vee \mathcal{F}$.

The following facts are (mostly) quite immediate

Fact 2.1 (See Lemma 2.1 in [Sew16], [Dow11] and the appendix of [Sew19b]). *Let $(X, \mathcal{B}_X, \mu)$ be a standard probability space (equipped with a metric which makes it into a Polish space, with $\mathcal{B}_X$ the Borel sigma-algebra and μ a Borel measure), $\phi: X \rightarrow \Sigma$ and $\psi: X \rightarrow \Omega$ observables, and $\mathcal{F}, \mathcal{F}', (\mathcal{F}_n)_{n=1}^{\infty}$ sub sigma algebras of $\mathcal{B}_X$. Then:*

1. Zero conditional entropy: $H_{\mu}(\phi|\mathcal{F}) = 0$ if and only if $\phi^{-1}(a) \in \mathcal{F}$ up to a μ -null set for every $a \in \Sigma$.
2. Conditional property: $H_{\mu}(\phi \vee \psi|\mathcal{F}) = H_{\mu}(\psi|\mathcal{F}) + H_{\mu}(\phi|\psi \vee \mathcal{F})$.⁹
3. Finite coarsenings approximate the entropy: $H_{\mu}(\phi|\mathcal{F})$ is the supremum over $H_{\mu}(\tau|\mathcal{F})$ for all finite alphabet observables τ for which $\tau \leq \phi$.
4. More pre-given information implies less entropy: $\mathcal{F}' \subseteq \mathcal{F}$ up to μ -null sets implies that $H_{\mu}(\phi|\mathcal{F}) \leq H_{\mu}(\phi|\mathcal{F}')$.
5. Limit of increasing pre-given information: If for every n we have $\mathcal{F}_n \subseteq \mathcal{F}_{n+1}$, $H(\phi|\mathcal{F}_n) < \infty$ and $\mathcal{F} = \bigvee_n \mathcal{F}_n$, then $H(\phi|\mathcal{F}) = \inf_{n \in \mathbb{N}} H(\phi|\mathcal{F}_n)$.
6. The space of finite entropy observables is complete when equipped with the Rokhlin distance d_{Rok} .¹⁰
7. Distance bounds: If $H_{\mu}(\phi|\mathcal{F}), H_{\mu}(\psi|\mathcal{F}) < \infty$, then $|H_{\mu}(\phi|\mathcal{F}) - H_{\mu}(\psi|\mathcal{F})| \leq d_{\text{Rok}}(\phi, \psi)$. Similarly, if $\sigma: X \rightarrow \Xi$ is another observable, and $H_{\mu}(\sigma|\phi \vee \mathcal{F}), H_{\mu}(\sigma|\psi \vee \mathcal{F}) < \infty$, then $|H_{\mu}(\sigma|\phi \vee \mathcal{F}) - H_{\mu}(\sigma|\psi \vee \mathcal{F})| \leq 2d_{\text{Rok}}(\phi, \psi)$. In addition, for a finite set $A \subseteq \mathcal{F}$ we have $d_{\text{Rok}}(\phi^A, \psi^A) \leq |A| \cdot d_{\text{Rok}}(\phi, \psi)$. Finally, if $\Sigma = \Omega$, then $d_{\text{Rok}}(\phi, \psi) \leq 2d_{TV}(\phi, \psi)$.
8. Every observable has a close by finite alphabet observable: Let $\Sigma = \mathbb{N}$, and for every $n \in \mathbb{N}$ let $\phi_n(x) = \min(n, \phi(x))$. It is clear that $\phi_n \xrightarrow{TV} \phi$, and thus by item 7., $\phi_n \xrightarrow{\text{Rok}} \phi$.
9. Every finite alphabet observable has a close by continuous finite alphabet observable: Let $|\Sigma| < \infty$. Suppose X is a totally disconnected. By the outer regularity of μ , for every $\varepsilon > 0$ there is a continuous observable $\varphi: X \rightarrow \Sigma$ such that $d_{TV}(\varphi, \phi) < \varepsilon$.
10. The continuous finite alphabet observables are dense in all observables: Assume X is totally disconnected. From items 7., 8. and 9. one can deduce that regardless of whether Σ is finite or infinite, for every $\varepsilon > 0$ there is a continuous finite alphabet observable $\varphi: X \rightarrow \Sigma'$ such that $d_{\text{Rok}}(\varphi, \phi) < \varepsilon$.

2.3. Generating observables and Rokhlin entropy. Let Γ be a countable group and fix a probability measure preserving (p.m.p) action $\Gamma \curvearrowright (X, \mu)$ (on a standard probability space) and observable $\phi: X \rightarrow \Sigma$. The group Γ acts naturally on observables, by letting $\phi^{\gamma}: X \rightarrow \Sigma$ be the observable $\phi^{\gamma}(x) = \phi(\gamma^{-1}x)$ for every $\gamma \in \Gamma$. Therefore, a subset $S \subset \Gamma$ defines a map $\phi^S: X \rightarrow \Sigma^S$ by letting $\phi^S = \bigvee_{s \in S} \phi^s$, or in other words $\phi^S(x)(s) = \phi(s^{-1}x)$. This map arises in two special cases: when S is finite, in which case ϕ^S is an observable, and when $S = \Gamma$, in which case it defines a Γ -equivariant map from X to Σ^{Γ} . A sigma algebra $\mathcal{F} \subseteq \mathcal{B}_X$ is said to be Γ -invariant if $B \in \mathcal{F}$ implies $\gamma.B \in \mathcal{F}$ for every $B \in \mathcal{F}$ and $\gamma \in \Gamma$. Similarly, a set B is Γ -invariant if $\gamma.B = B$ for every $\gamma \in \Gamma$, and we denote by $\mathcal{I}_{\Gamma}$ the sigma algebra generated by Γ -invariant sets.

⁹This is the entropic version of $\mathbb{P}[A|B] = \frac{\mathbb{P}[A \cap B]}{\mathbb{P}[B]}$.

¹⁰Here we are slightly abusing terminology. If two observables have Rokhlin distance zero then we consider them to be the same.

Definition 2.2 (Generating observable). An observable ϕ is *generating* with respect to a p.m.p action $\Gamma \curvearrowright (X, \mu)$ if the smallest Γ -invariant sigma-algebra which makes ϕ measurable is $\mathcal{B}_X$ (up to μ -null sets). Equivalently, for $\phi^\Gamma: X \rightarrow \Sigma^\Gamma$ we have that $(\phi^\Gamma)^{-1}(\mathcal{B}_{\Sigma^\Gamma})$ is $\mathcal{B}_X$ up to μ -null sets.

Let $\mathcal{F} \subseteq \mathcal{B}_X$ be a Γ -invariant sub sigma algebra. An observable ϕ is *generating given $\mathcal{F}$* , if $\phi^\Gamma \vee \mathcal{F}$ is $\mathcal{B}_X$ up to μ -null sets

Remark 2.3. There is an “operative” way of thinking about ϕ being generating: If one samples $x \sim \mu$, and reveals to you $\phi^\gamma(x)$ for every $\gamma \in \Gamma$, then you can (almost surely) recover x . Similarly, if $\psi: X \rightarrow Y$ is the factor associated with $\mathcal{F}$, then ϕ generating given $\mathcal{F}$ implies that for $x \sim \mu$, revealing $\psi(x)$ and $\phi^\gamma(x)$ for every $\gamma \in \Gamma$ allows one to (almost surely) recover x .

Definition 2.4 (Rokhlin entropy). Let $\mathcal{F}$ be a Γ -invariant Borel sigma-algebra. The *relative Rokhlin entropy* of this action given $\mathcal{F}$ is

$$H_{\text{Rok}}(\Gamma \curvearrowright (X, \mu) | \mathcal{F}) = \inf H_\mu(\phi | \mathcal{F}),$$

where the infimum is over all observables ϕ which generate given $\mathcal{F}$ (Definition 2.2). In case $\mathcal{F}$ is trivial up to μ -null sets, this quantity is called the *Rokhlin entropy* and we denote it by $H_{\text{Rok}}(\Gamma \curvearrowright (X, \mu))$.

Remark 2.5. Seward usually defines the relative Rokhlin entropy modulo the information given by $\mathcal{I}_\Gamma$, namely as the infimum of $H_\mu(\phi | \mathcal{F} \vee \mathcal{I}_\Gamma)$ for observables ϕ generating given $\mathcal{F} \vee \mathcal{I}_\Gamma$. His definition is the “correct” one in case of general aperiodic actions. But, in our case we mix both the aperiodic and periodic cases, and it would be more natural to use our definition. This will be consistent with the theory developed by Seward and others, since the $\mathcal{F}$ we use includes $\mathcal{I}_\Gamma$ whenever the IRS in question is aperiodic, and in the periodic case we intentionally diverge from the standard definition to be compatible with our version of RBS (Definition 3.5).

Let Σ be a finite set, and κ a probability measure over it. Recall that the *Bernoulli shift* over Γ with base space (Σ, κ) is the probability measure preserving action of Γ on the product space $(\Sigma^\Gamma, \kappa^\Gamma)$ given by $\gamma.c(x) = c(\gamma^{-1}x)$ for $\gamma, x \in \Gamma$ and $c: \Gamma \rightarrow \Sigma$. The *canonical observable* of this action, which we denote by $\phi_\Sigma: \Sigma^\Gamma \rightarrow \Sigma$, is defined by $\phi_\Sigma(c) = c(\mathbf{1}_\Gamma)$. As the pushforward measure $(\phi_\Sigma)_*\kappa^\Gamma$ is κ , we have that $H_{\kappa^\Gamma}(\phi_\Sigma) = H(\kappa)$. In addition, ϕ_Σ is generating (Definition 2.2), as $\phi_\Sigma^\gamma(c) = \phi_\Sigma(\gamma^{-1}.c) = \gamma^{-1}.c(\mathbf{1}_\Gamma) = c(\gamma)$ for every $\gamma \in \Gamma$, which means that we recover c when we get access to all ϕ_Σ^γ (see Remark 2.3). All in all, the Rokhlin entropy (Definition 2.4) of the Bernoulli shift satisfies $H_{\text{Rok}}(\Gamma \curvearrowright (\Sigma^\Gamma, \kappa^\Gamma)) \leq H(\kappa)$. A group Γ satisfies the *maximal Rokhlin entropy of Bernoulli Shifts* (RBS) criterion [Sew19b] if for every finite Σ , and given u the uniform measure over Σ , one has $H_{\text{Rok}}(\Gamma \curvearrowright (\Sigma^\Gamma, u^\Gamma)) = H(u) = \log |\Sigma|$.¹¹

3. The maximal Rokhlin entropy of Bernoulli Shifts criterion for Invariant Random Subgroups

Let $\mathcal{F}$ be a finitely generated free group, and recall from Section 1.1 the space $\mathbf{sub}(\mathcal{F}, \Sigma)$ consisting of all pairs (K, c) , where K is a subgroup of $\mathcal{F}$ and $c: \mathcal{F} \rightarrow \Sigma$ is a K -invariant coloring, together with the action $w.(K, c) = (wKw^{-1}, w.c)$ for $w \in \mathcal{F}$. Let κ be a probability distribution over Σ . For every $K \leq \mathcal{F}$, let $\kappa_K = \kappa^{(\kappa \setminus \mathcal{F})}$ be the product distribution over right K -cosets of $\mathcal{F}$; namely, κ_K samples for every coset Kw a color $c(Kw) \in \Sigma$ according to the distribution κ independently of all other cosets. Given an invariant random subgroup (IRS) π of $\mathcal{F}$ (Definition 1.3), one can define the probability measure κ_π over $\mathbf{sub}(\mathcal{F}, \Sigma)$ to be $\mathbb{E}_{K \sim \pi}[\mathbf{1}_K \times \kappa_K]$ — in words, to sample according to κ_π we need first to sample K according to π , and

¹¹Seward assumes this is true for **every** standard probability base space (Σ, κ) and not only for finite ones equipped with the uniform measure. But, to deduce surjunctivity, both in his case and in the IRS case which is the topic of this paper, this weaker condition suffices.

then sample for every Kw an independent color $c(Kw)$ according to κ , to retrieve $(K, c) \in \mathbf{sub}(\mathcal{F}, \Sigma)$. When equipping $\mathbf{sub}(\mathcal{F}, \Sigma)$ with κ_π , the aforementioned action of $\mathcal{F}$ on $\mathbf{sub}(\mathcal{F}, \Sigma)$ becomes p.m.p, and we call it the *Bernoulli shift* over π with base (Σ, κ) . As before, when $\Gamma \cong \mathcal{F}/N$ for a normal subgroup $N \trianglelefteq \mathcal{F}$, the Bernoulli shift over the IRS $\mathbf{1}_N$ with base (Σ, κ) coincides with the Bernoulli shift over the group Γ with the same base. Note also that, if $P: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \mathbf{sub}(\mathcal{F})$ is the projection on the left coordinate, then the pushforward satisfies $P_*\kappa_\pi = \pi$. We provide now the appropriate IRS analogue to the discussion on RBS from the end of Section 2. To that end we need the following simple fact:

Fact 3.1 (Finite vs. infinite index decomposition of IRSs). *Let π be an IRS of $\mathcal{F}$. There are IRSs π_f and π_∞ such that: • π is a convex combination of them; • $K \sim \pi_f$ is of finite index; • $K \sim \pi_\infty$ is of infinite index.*

We call π_f the finite component of π and π_∞ its infinite component.

Remark 3.2. Let $\mathbf{sub}_f(\mathcal{F})$ be the set of finite index subgroups of $\mathcal{F}$ and $\mathbf{sub}_\infty(\mathcal{F})$ the set of infinite index ones. As the singleton $\{K\}$ for every $K \leq \mathcal{F}$ with $[\mathcal{F} : K] < \infty$ is open, $\mathbf{sub}_f(\mathcal{F})$ is an open set, and thus $\mathbf{sub}_\infty(\mathcal{F})$ is closed. The above condition “a subgroup K sampled according to π_f is of finite index” is equivalent to π_f being a convex combination of finitely described IRSs (1.1), while the other condition “a subgroup K sampled according to π_∞ is of infinite index” is the same as $\text{supp}(\pi_\infty) \subseteq \mathbf{sub}_\infty(\mathcal{F})$.

We denote by $\mathbf{sub}_f(\mathcal{F}, \Sigma)$ all $(H, c) \in \mathbf{sub}(\mathcal{F}, \Sigma)$ such that $H \in \mathbf{sub}_f(\mathcal{F})$, and $\mathbf{sub}_\infty(\mathcal{F}, \Sigma)$ is similarly defined.

Definition 3.3. The *canonical observable* $\phi_\Sigma: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \Sigma$ is defined to be $\phi_\Sigma(K, c) = c(K)$.

Claim 3.4. *Let π be an IRS of $\mathcal{F}$, and (Σ, κ) a finite probability space. Then, the canonical observable ϕ_Σ is generating (Definition 2.2) for the Bernoulli shift over π with base (Σ, κ) **given** P . In addition, $H_{\kappa_\pi}(\phi_\Sigma|P) = H(\kappa)$, which in turn implies $H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), \kappa_\pi)|P) \leq H(\kappa)$.*

Proof. For the generation claim, use the operative framework of Remark 2.3. If $(K, c) \sim \kappa_\pi$, and we get access to $P(K, c) = K$ as well as

$$(3.1) \quad \phi_\Sigma^w(K, c) = \phi_\Sigma(w^{-1} \cdot (K, c)) = \phi_\Sigma(w^{-1}Kw, w^{-1} \cdot c) = w^{-1} \cdot c(w^{-1}Kw) = c(Kw)$$

for every $w \in \mathcal{F}$, then we recover (K, c) completely. Hence ϕ_Σ is generating given P .

For the entropy condition, as $P_*\kappa_\pi = \pi$, and the conditional distributions (Section 2.1) satisfy $(\kappa_\pi)_K^P = \mathbf{1}_K \times \kappa_K$ for every $K \leq \mathcal{F}$, we have that

$$(3.2) \quad H_{\kappa_\pi}(\phi_\Sigma|P) = \mathbb{E}_{K \sim \pi} [H_{\mathbf{1}_K \times \kappa_K}(\phi_\Sigma)] = \mathbb{E}_{K \sim \pi} [H(\underbrace{((\phi_\Sigma)_* \mathbf{1}_K \times \kappa_K)}_{=\kappa})] = H(\kappa).$$

As the relative Rokhlin entropy (Definition 2.4) $H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), \kappa_\pi)|P)$ is the infimum of $H_{\kappa_\pi}(\phi|P)$ over all observables which generate given P , we are done. $\square$

Definition 3.5 (RBS). Let π be an IRS **supported on infinite index subgroups**. We say that π satisfies the Rokhlin entropy of Bernoulli Shifts (RBS) criterion if $H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi)|P) = H(u) = \log |\Sigma|$ for every finite base Σ equipped with the uniform probability distribution u . We will say that a general IRS π (not necessarily supported on infinite index subgroups) satisfies RBS if its infinite component π_∞ (as in Fact 3.1) does. In particular, by definition, if π -a.e. subgroup has finite-index, then π satisfies RBS. We denote the set of IRSs satisfying RBS by $\text{IRS}_{\text{RBS}}(\mathcal{F})$.

By Definition 3.5, every IRS which is supported on finite index subgroups satisfies RBS. But, it is not clear from the definition that $\text{IRS}_{\text{RBS}}(\mathcal{F})$ is a closed subset of $\text{IRS}(\mathcal{F})$. Hence, one needs still to argue why every

cosofic IRS satisfies RBS. In the next section, we provide a condition equivalent to RBS which will allow us both to construct an outer approximation algorithm for $\text{IRS}_{\text{RBS}}(\mathcal{F})$, as well as to show that $\text{IRS}_{\text{RBS}}(\mathcal{F})$ is weak* closed (and thus every cosofic IRS will satisfy RBS). Let us finish this section by proving:

Theorem 3.6. *If an IRS satisfies the RBS criterion (Definition 3.5), then it is surjunctive (Definition 1.6).*

To that end, we need to collect additional definitions and claims. The following fact is straightforward and we leave it without a proof:

Fact 3.7 (Basic surjunctivity properties). *Let π, π_1 and π_2 be IRSs of $\mathcal{F}$. We have:*

- (1) Surjunctivity is preserved under convex combinations: *If π_1 and π_2 are surjunctive (Definition 1.6), then every convex combination of π_1 and π_2 is also surjunctive.*
- (2) Finite index IRSs are surjunctive: *If π is supported on finite index subgroups, then it is surjunctive.*

Let $([0, 1], \lambda)$ be the interval equipped with the Lebesgue measure on it. For a subgroup $K \leq \mathcal{F}$, let $\lambda_K = \lambda^{(K \backslash \mathcal{F})}$ be the distribution that associates to each right K -coset an independent uniform real number between 0 and 1, which we think of as the “time in which this coset appears”. Almost surely, $t: K \backslash \mathcal{F} \rightarrow [0, 1]$ sampled according to λ_K induces a linear order on $K \backslash \mathcal{F}$, and we let $L_{K,t} = \{w \in \mathcal{F} \mid t(K) > t(Kw)\}$ to consist of all “past group elements”. Given an IRS π , let $\lambda_\pi = \mathbb{E}_{K \sim \pi}[\mathbf{1}_K \times \lambda_K]$, namely sample a subgroup K according to π and then sample a uniform real number between 0 and 1 for each K -coset independently.

Every p.m.p action $\mathcal{F} \curvearrowright (X, \mu)$ induces an IRS $\text{Stab}(\mu)$ of $\mathcal{F}$ by letting $K = \text{Stab}(x)$ for $x \sim \mu$. Given an IRS π , we say that $\mathcal{F} \curvearrowright (X, \mu)$ has stabilizer type π , if $\text{Stab}(\mu) = \pi$. Let $\mu_K = \mu_K^{\text{Stab}}$ be the associated conditional distributions (Section 2.1).

Definition 3.8 (Percolation entropy). Let π be an IRS of $\mathcal{F}$, $\mathcal{F} \curvearrowright (X, \mu)$ a p.m.p action with stabilizer type π , $\phi: X \rightarrow \Omega$ an observable (Section 2.2) and $\mathcal{F}$ an $\mathcal{F}$ -invariant sub sigma algebra of $\mathcal{B}_X$. The percolation entropy of ϕ given $\mathcal{F}$ is

$$(3.3) \quad H_{\text{perc}, \mu}(\phi | \mathcal{F}) = \mathbb{E}_{(K,t) \sim \lambda_\pi} [H_{\mu_K}(\phi | \phi^{L_{K,t}} \vee \mathcal{F})].$$

Remark 3.9. Let us unpack the above definition, as it is a bit mysterious. By the definition of conditional distributions, and as π is the stabilizer type of μ , sampling $x \sim \mu$ is the same as first sampling $K \sim \pi$ and then $x \sim \mu_K$. On the right hand side of (3.3), we have one more object sampled, which is a time function $t: K \backslash \mathcal{F} \rightarrow [0, 1]$ sampled according to λ_K . Let ψ be the factor associated to $\mathcal{F}$. Then, the quantity $H_{\mu_K}(\phi | \phi^{L_{K,t}} \vee \mathcal{F})$ measures how much information is gained by learning $\phi(x)$ for $x \sim \mu_K$, given that you have learned the value $\psi(x)$ and the conjugates $\phi^w(x)$ for every past element $w \in L_{K,t}$. Namely, given all evaluations of past conjugates of ϕ (in addition to ψ), how much more information will the evaluation of ϕ provide? The next claim contains the main example we will need, which will also shed some light on this definition.

Claim 3.10. *Let π be an IRS of $\mathcal{F}$ supported on infinite index subgroups, Σ a finite set, $|\Sigma| > 1$, and $P: \text{sub}(\mathcal{F}, \Sigma) \rightarrow \text{sub}(\mathcal{F})$ the projection on the first coordinate. Recall the shift action $w.(K, c) = (wKw^{-1}, w.c)$ of $\mathcal{F}$ on $\text{sub}(\mathcal{F}, \Sigma)$.*

- (1) *Let u be the uniform measure over Σ , $\mathcal{F} \curvearrowright (\text{sub}(\mathcal{F}, \Sigma), u_\pi)$ the associated Bernoulli shift, and $\phi_\Sigma: \text{sub}(\mathcal{F}, \Sigma) \rightarrow \Sigma$ the canonical observable $\phi_\Sigma(K, c) = c(K)$. Then, the stabilizer type of u_π is π , the sigma algebra of $\mathcal{F}$ -invariant sets $\mathcal{I}_{\mathcal{F}}$ is contained in $P^{-1}(\mathcal{B}_{\text{sub}(\mathcal{F})})$ up to u_π -null sets, and $H_{\text{perc}, u_\pi}(\phi_\Sigma | P) = \log |\Sigma|$.*

- (2) Let $\mathcal{F}$ still act on $\mathbf{sub}(\mathcal{F}, \Sigma)$ by shifts, and let μ be a Borel probability measure on $\mathbf{sub}(\mathcal{F}, \Sigma)$ which is preserved under this action. Assume the stabilizer type of μ is $\pi = P_*\mu$ and $H_{\text{perc}, \mu}(\phi_\Sigma | P) = \log |\Sigma|$. Then $\mu = u_\pi$.

Proof. **Claim 1.** The stabilizer type of u_π is π .

Proof of Claim 1. Given a subgroup $K \leq \mathcal{F}$, recall that $u_K = u^{\kappa \setminus \mathcal{F}}$ is the product measure on $\Sigma^{\kappa \setminus \mathcal{F}}$, that $\mathbf{1}_K$ is the Dirac measure concentrated on K , and that $u_\pi = \mathbb{E}_{K \sim \pi}[\mathbf{1}_K \times u_K]$. The stabilizer of $K \leq \mathcal{F}$ with respect to conjugation is the normalizer $\text{Nor}(K)$, and the stabilizer of $c: \mathcal{F} \rightarrow K$ with respect to the shift action is $M(c) = \{w \in \mathcal{F} \mid \forall x \in \mathcal{F}: c(x) = c(w^{-1}x)\}$. Hence, the stabilizer of (K, c) is $\text{Nor}(K) \cap M(c)$. As c is K -invariant, $K \subseteq M(c)$ and thus $K \subseteq \text{Stab}(K, c)$ for every $(K, c) \in \mathbf{sub}(\mathcal{F}, \Sigma)$. By Boole's inequality,

$$\begin{aligned} \mathbb{P}_{(K, c) \sim u_\pi} [K \neq \text{Stab}(K, c)] &= \mathbb{P}_{(K, c) \sim u_\pi} [\exists w \in \mathcal{F}: w \notin K \wedge w \in \text{Stab}(K, c)] \\ &\leq \sum_{w \in \mathcal{F}} \mathbb{P}_{(K, c) \sim u_\pi} [w \notin K \wedge w \in \text{Stab}(K, c)] \\ &= (*). \end{aligned}$$

Let $\text{rad}(\pi) = \bigcap_{K \in \text{supp}(\pi)} K$. Then, for every $w \notin \text{rad}(\pi)$, there is a positive probability that $K \sim \pi$ satisfies $w \notin K$.

Hence,

$$(*) = \sum_{w \notin \text{rad}(\pi)} \mathbb{P}_{(K, c) \sim u_\pi} [w \notin K] \cdot \mathbb{P}[w \in \text{Stab}(K, c) \mid w \notin K].$$

Hence, if we can show that $\mathbb{P}[w \in \text{Stab}(K, c) \mid w \notin K] = 0$ for every infinite index subgroup K , then we are done. Because K has infinite index in $\mathcal{F}$, there is an infinite $I \subset \mathcal{F}$ such that $\Theta = \{K\gamma, Kw^{-1}\gamma \mid \gamma \in I\}$ are all different K -cosets. As c was sampled from u_K , the color of each coset in Θ is independent of every other coset. In particular, $\mathbb{P}[c(\gamma) = c(w^{-1}\gamma)] = \frac{1}{|\Sigma|} < 1$. Therefore, with probability 1 there is a $\gamma \in I$ for which $c(\gamma) \neq c(w^{-1}\gamma) = w.c(\gamma)$, and $w \notin \text{Stab}(K, c)$. $\square$

Claim 2. The sigma algebra of $\mathcal{F}$ -invariant sets $\mathcal{I}_{\mathcal{F}}$ is contained in $P^{-1}(\mathcal{B}_{\mathbf{sub}(\mathcal{F})})$ up to u_π -null sets.

Proof of Claim 2. Recall that $\mathbf{sub}_\infty(\mathcal{F})$ (Remark 3.2) is the closed set of infinite index subgroups of $\mathcal{F}$. There is a sequence of continuous maps $\phi_n: \mathbf{sub}_\infty(\mathcal{F}) \rightarrow \mathcal{F}$ such that for every $K \in \mathbf{sub}_\infty(\mathcal{F})$ and for any finite set $W \subseteq \mathcal{F}$, there exists n_0 such that $\phi_n(K)^{-1} \notin \bigcup_{w, w' \in W} Kw w'^{-1}$ for all $n \geq n_0$ — for example, ϕ_n can lexicographically order the sphere of radius n in the Schreier graph $\text{Sch}(\mathcal{F}, K, S)$ and output the shortest representative of the first vertex in this order.¹²

For $f \in L^1(\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi)$ and $K \in \text{supp}(\pi) \subseteq \mathbf{sub}_\infty(\mathcal{F})$, let $\tilde{f}(K) = \mathbb{E}_{c \sim u_K} [f(K, c)]$ and let $f \circ \phi_n(K, c) = f(\phi_n(K), (K, c))$. We claim the following mixing property: For every $f_1, f_2 \in L^1(\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi)$, the covariance of $f_1 \circ \phi_n$ and f_2 given P tends to zero, namely

$$(3.4) \quad \lim_{n \rightarrow \infty} \mathbb{E}_{(K, c) \sim u_\pi} [f_1 \circ \phi_n(K, c) \cdot f_2(K, c) - \widetilde{f_1 \circ \phi_n}(K) \cdot \tilde{f}_2(K)] = 0.$$

In order to prove this statement, it suffices to prove it in the special case in which both f_1, f_2 are continuous, because continuous functions are dense in $L^1(\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi)$. Assume f_1, f_2 are continuous. Then there exists a

¹²These functions are continuous as the Schreier graph up to radius n can be constructed by only knowing the intersection of K with a ball of radius $2n + 1$.

finite set $W \subset \mathcal{F}$ such that $f_1(K, c)$ and $f_2(K, c)$ depend only on $K \cap W$ and the restriction of c to W . For a large enough n we have $\phi_n(K)^{-1} \notin \bigcup_{w, w' \in W} K w w'^{-1}$, which implies in particular that the cosets $\{K w\}_{w \in W}$ and $\{K \phi_n(K)^{-1} w\}_{w \in W}$ are disjoint. As u_K colors each coset independently, we deduce that $c|_W$ and $\phi_n(K).c|_W = c|_{\phi_n(K)^{-1}.W}$ are independent when $c \sim u_K$. Thus, $\mathbb{E}_{(K, c) \sim u_\pi} [f_1 \circ \phi_n(K, c) \cdot f_2(K, c) - \widetilde{f_1 \circ \phi_n(K)} \cdot \widetilde{f_2(K)}] = 0$, which proves (3.4).

Now let $E \subseteq \text{sub}(\mathcal{F}, \Sigma)$ be a $\mathcal{F}$ -invariant measurable subset. Let $f_1 = \mathbf{1}_E$ and $f_2 = 1 - f_1$. As E is $\mathcal{F}$ -invariant, $f_1 = f_1 \circ \phi_n$ for every n , and we have $(f_1 \circ \phi_n) \cdot f_2 = 0$ almost everywhere. So (3.4) implies that for π -almost every K , $\widetilde{f_1}(K) = 0$ or 1 — in other words, for π -almost every K , either for u_K -almost all $c: K \backslash \mathcal{F} \rightarrow \Sigma$ we have $(K, c) \in E$, or for almost none of them. This implies that E is in $P^{-1}(\mathcal{B}_{\text{sub}(\mathcal{F})})$ up to u_π -null sets, as required. $\square$

Claim 3. $H_{\text{perc}, u_\pi}(\phi_\Sigma | P) = \log |\Sigma|$.

Proof of Claim 3. By definition,

$$H_{\text{perc}, u_\pi}(\phi_\Sigma | P) = \mathbb{E}_{(K, t) \sim \lambda_\pi} [H_{\mathbf{1}_K \times u_K}(\phi_\Sigma | \phi_\Sigma^{L_{K, t}} \vee P)].$$

When sampling $(K, c) \sim \mathbf{1}_K \times u_K$, $P(K, c) = K$ gives us no new information, and $H_{\mathbf{1}_K \times u_K}(\phi_\Sigma | \phi_\Sigma^{L_{K, t}} \vee P) = H_{\mathbf{1}_K \times u_K}(\phi_\Sigma | \phi_\Sigma^{L_{K, t}})$. In addition, as $\text{Id}_{\mathcal{F}} \notin L_{K, t}$, ϕ_Σ is independent of $\phi_\Sigma^{L_{K, t}}$ and

$$H_{\mathbf{1}_K \times u_K}(\phi_\Sigma | \phi_\Sigma^{L_{K, t}}) = H_{\mathbf{1}_K \times u_K}(\phi_\Sigma) = H(\phi_{\Sigma*}(\mathbf{1}_K \times u_K)) = H(u) = \log |\Sigma|. \quad \square$$

Claim 4. Assume $H_{\text{perc}, \mu}(\phi_\Sigma | P) = \log |\Sigma|$. Then $\mu = u_\pi$.

Proof of Claim 4. As μ both has stabilizer type π and $P_*\mu = \pi$, we have

$$H_{\text{perc}, \mu}(\phi_\Sigma | P) = \mathbb{E}_{(K, t) \sim \lambda_\pi} [H_{\mu_K}(\phi_\Sigma | \phi_\Sigma^{L_{K, t}} \vee P)],$$

where $\mu = \mathbb{E}_{K \sim \pi} [\mathbf{1}_K \times \mu_K]$. Because ϕ_Σ takes values in Σ , $H_{\mu_K}(\phi_\Sigma | \phi_\Sigma^{L_{K, t}} \vee P) \leq \log |\Sigma|$, and equality holds if and only if ϕ_Σ is independent of $\phi_\Sigma^{L_{K, t}} \vee P$ (for a.e. (K, t)) and $\phi_{\Sigma*}\mu_K$ is the uniform measure on Σ . But this implies that if $c \sim \mu_K$ is a random sample, then $c(K)$ is independent of $c(Kw)$ for all non-identity cosets Kw (since each coset Kw appears in $L_{K, t}$ with positive probability). Moreover, the variables $c(Kw)$ are uniformly distributed on Σ . Thus μ_K is u_K and $\mu = \pi$. $\square$

The reason we have defined percolation entropy is because of the following result:

Fact 3.11 (Corollary 7.3 in [Sew16]). *Let π be an IRS of $\mathcal{F}$ supported on infinite index subgroups, $\mathcal{F} \curvearrowright (X, \mu)$ a p.m.p action with stabilizer type π , $\phi: X \rightarrow \Omega$ a **generating** observable (Section 2.2 and Definition 2.2) and $\mathcal{F}$ an $\mathcal{F}$ -invariant sub sigma algebra of $\mathcal{B}_X$ which contains the sub algebra $\mathcal{I}_{\mathcal{F}}$ of $\mathcal{F}$ -invariant sets. Then*

$$H_{\text{Rok}}(\mathcal{F} \curvearrowright (X, \mu) | \mathcal{F}) \leq H_{\text{perc}, \mu}(\phi | \mathcal{F}).$$

Proof of Theorem 3.6. By Fact 3.7, it is enough to prove the theorem assuming π is supported on infinite index subgroups. Let $\Phi: \text{Dom}(\pi) \rightarrow \text{Dom}(\pi)$ be an injective cellular automaton. Since u_π is supported on $\text{Dom}(\pi)$, we can push it forward along Φ and get $\mu = \Phi_* u_\pi$. As Φ is $\mathcal{F}$ -equivariant, we have

$$\forall (K, c) \in \text{sub}(\mathcal{F}, \Sigma), w \in \text{Stab}(\Phi(K, c)) : \quad \Phi(K, c) = w.\Phi(K, c) = \Phi(w.(K, c)),$$

but it is assumed to be injective and thus $(K, c) = w.(K, c)$ and $w \in \text{Stab}(K, c)$. Hence, $\text{Stab}(\Phi(K, c)) = \text{Stab}(K, c)$ and μ has the same stabilizer type as u_π , which is π by Claim 3.10. Now, the shift $\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), \mu)$ is a p.m.p action which is isomorphic to the original Bernoulli shift (as Φ is injective). Since π satisfies RBS (Definition 3.5), we have

$$H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), \mu) | P) = H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi) | P) = \log |\Sigma|.$$

By Claim 3.10, $\mathcal{I}_{\mathcal{F}} \subseteq P^{-1}(\mathcal{B}_{\mathbf{sub}(\mathcal{F})})$, and thus by Fact 3.11, we get that $H_{\text{perc}, \mu}(\phi_\Sigma | P) \geq \log |\Sigma|$. As

$$H_{\mu_K}(\phi_\Sigma | \phi_\Sigma^{L_{K,t}} \vee P) \leq H_{\mu_K}(\phi_\Sigma) = H((\phi_\Sigma)_* \mu_K) \leq \log |\Sigma|$$

(where the first inequality is due item 4. in Fact 2.1 and the last is because $(\phi_\Sigma)_* \mu_K$ is some probability distribution over Σ), we deduce that $H_{\text{perc}, \mu}(\phi_\Sigma | P)$ is also upper bounded by $\log |\Sigma|$. Hence, $H_{\text{perc}, \mu}(\phi_\Sigma | P) = \log |\Sigma|$, and by (2) in Claim 3.10 we have $\mu = u_\pi$. Since Φ is continuous, its image agrees with $\text{supp}(\mu) = \text{supp}(u_\pi) = \text{Dom}(\pi)$ and it is thus surjective. $\square$

4. A condition equivalent to RBS

In this section we prove a condition equivalent to RBS. In some sense, this result is the key for providing an outer approximation (Theorem 5.3) and thus for proving Theorem 1.2. This is because it translates the RBS condition, which is about measurable (which is a non-finitary condition) generating (a global condition) observables, to a condition for general continuous observables, allowing local analysis.

Theorem 4.1. *An IRS $\mathcal{F}$ satisfies RBS (Definition 3.5) if and only if for every pair of finite sets Σ, Ω and every continuous observable $\psi: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \Omega$ we have*

$$(4.1) \quad \log |\Sigma| = H(u) \leq H_{u_\pi}(\psi | P) + H_{u_\pi}(\phi_\Sigma | \psi^\mathcal{F} \vee P),$$

where u is the uniform measure over Σ and ϕ_Σ is the canonical observable (Definition 3.3).

Proof. Let us first prove that satisfying condition (4.1) for every continuous observable implies the RBS criterion. Assume π does not satisfy RBS, and let $\pi = \alpha_\infty \pi_\infty + \alpha_f \pi_f$ be its decomposition to infinite and finite components as in Fact 3.1. Then, $\alpha_\infty > 0$, and there is a (measurable) observable $\psi: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \Omega$ which is generating given P , namely $\psi^\mathcal{F} \vee P$ is the Borel sigma algebra of $\mathbf{sub}(\mathcal{F}, \Sigma)$ up to π_∞ -null sets, and also $H_{u_{\pi_\infty}}(\psi | P) < \log |\Sigma|$. Let us define a new observable $\psi': \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \Omega \sqcup \Sigma$:

$$\psi'(K, c) = \begin{cases} \phi_\Sigma(K, c) & [\mathcal{F} : K] < \infty, \\ \psi(K, c) & [\mathcal{F} : K] = \infty, \end{cases}$$

where ϕ_Σ is the canonical observable (Definition 3.3). We claim that: $\bullet$ ψ' is measurable, $\bullet$ it is generating given P up to π -null sets, $\bullet$ $H_{u_\pi}(\psi' | P) < \log |\Sigma|$. To that end, recall the notations $\mathbf{sub}_f(\mathcal{F})$ and $\mathbf{sub}_\infty(\mathcal{F})$ (Remark 3.2) for the finite and infinite index subgroups of $\mathcal{F}$ respectively. The inverse image of $a \in \Sigma$ has the form $\phi_\Sigma^{-1}(a) \cap \mathbf{sub}_f(\mathcal{F}, \Sigma)$ which is the intersection of two open subsets of $\mathbf{sub}(\mathcal{F}, \Sigma)$ and is therefore open. The inverse image of $b \in \Omega$ is $\psi^{-1}(b) \cap \mathbf{sub}_\infty(\mathcal{F}, \Sigma)$ which is the intersection of a Borel subset and a closed subset and is therefore Borel. All in all, ψ' is measurable. For generation, let $(K, c) \sim u_\pi$, and assume we are given the information $\psi'^w(K, c)$ for every $w \in \mathcal{F}$ and $P(K, c) = K$. In case K is of finite index, then $\psi'^w(K, c) = c(Kw)$ (3.1) which determines c completely. In case K is of infinite index, then it is in the support of π_∞ and by our assumption that ψ generates with respect to π_∞ , as $\psi'^w(K, c) = \psi^w(K, c)$ in this case, c is completely determined.

Hence, by Remark 2.3, ψ' is generating. Finally, as ψ' is generating, we have $H_{u_\pi}(\phi_\Sigma|\psi'^{\mathcal{F}} \vee P) = 0$ by item 1. of Fact 2.1. Moreover, by the definition of conditional entropy (Section 2.2), we have

$$(4.2) \quad H_{u_\pi}(\psi'|P) = \underbrace{\alpha_f H_{u_{\pi_f}}(\phi_\Sigma|P)}_{=(3.2)\log|\Sigma|} + \underbrace{\alpha_\infty H_{u_{\pi_\infty}}(\psi|P)}_{<\log|\Sigma|} < \log|\Sigma|.$$

It may seem that we are done with this direction, as ψ' does not satisfy (4.1), but ψ' is only measurable and not continuous. Let us amend this. Since $H_{u_\pi}(\phi_\Sigma|\psi'^{\mathcal{F}} \vee P) = 0$, by item 5. in Fact 2.1, for every $\varepsilon > 0$ there is some finite subset $A \subseteq \mathcal{F}$ for which $H_{u_\pi}(\phi_\Sigma|\psi'^A \vee P) < \varepsilon$. By item 7. in Fact 2.1, if ψ'' is an observable which satisfies $d_{\text{Rok}}(\psi'', \psi') \leq \delta$, then $d_{\text{Rok}}(\psi'', \psi'^A) \leq |A|\delta$ and

$$\begin{aligned} |H_{u_\pi}(\psi''|P) - H_{u_\pi}(\psi'|P)| &\leq d_{\text{Rok}}(\psi'', \psi') \leq \delta, \\ |H_{u_\pi}(\phi_\Sigma|\psi''^A \vee P) - H_{u_\pi}(\phi_\Sigma|\psi'^A \vee P)| &\leq 2d_{\text{Rok}}(\psi'', \psi'^A) \leq 2|A|\delta. \end{aligned}$$

Hence,

$$\begin{aligned} H_{u_\pi}(\psi''|P) + H_{u_\pi}(\phi_\Sigma|\psi''^A \vee P) &\leq H_{u_\pi}(\psi'|P) + \delta + H_{u_\pi}(\phi_\Sigma|\psi'^A \vee P) + 2|A|\delta \\ &\leq H_{u_\pi}(\psi'|P) + (2|A| + 1)\delta + \varepsilon. \end{aligned}$$

As $H_{u_\pi}(\psi'|P) < \log|\Sigma|$ (4.2), we can choose ε and then δ so that $H_{u_\pi}(\psi''|P) + H_{u_\pi}(\phi_\Sigma|\psi''^A \vee P) < \log|\Sigma|$. By item 10. in Fact 2.1, we can choose ψ'' to be continuous and are thus done with this direction.

Since

$$\begin{aligned} H_{u_\pi}(\psi|P) + H_{u_\pi}(\phi_\Sigma|\psi^{\mathcal{F}} \vee P) &= \alpha_\infty(H_{u_{\pi_\infty}}(\psi|P) + H_{u_{\pi_\infty}}(\phi_\Sigma|\psi^{\mathcal{F}} \vee P)) \\ &\quad + \alpha_f(H_{u_{\pi_f}}(\psi|P) + H_{u_{\pi_f}}(\phi_\Sigma|\psi^{\mathcal{F}} \vee P)), \end{aligned}$$

where $\pi = \alpha_\infty\pi_\infty + \alpha_f\pi_f$ is the decomposition as in Fact 3.1, we can prove that RBS implies (4.1) separately for the finite index and infinite index cases.

Assume π samples only finite index subgroups (so it satisfies RBS by Definition 3.5). Then, for every observable $\psi: \text{sub}(\mathcal{F}, \Sigma) \rightarrow \Omega$ and finite index subgroup $K \leq \mathcal{F}$ we have

$$\begin{aligned} (4.3) \quad \log(\Sigma) &= \frac{[\mathcal{F} : K] \cdot \log(\Sigma)}{[\mathcal{F} : K]} \\ &= \frac{H_{u_K}(\phi_\Sigma^{\mathcal{F}})}{[\mathcal{F} : K]} \\ &= \frac{H_{u_K}(\phi_\Sigma^{\mathcal{F}} \vee \psi^{\mathcal{F}})}{[\mathcal{F} : K]} \\ &= \frac{H_{u_K}(\psi^{\mathcal{F}}) + H_{u_K}(\phi_\Sigma^{\mathcal{F}}|\psi^{\mathcal{F}})}{[\mathcal{F} : K]} \\ &\leq H_{u_K}(\psi) + H_{u_K}(\phi_\Sigma|\psi^{\mathcal{F}}) \end{aligned}$$

where the second equality is since $\phi_{\Sigma^*}^{\mathcal{F}} u_K = u^{[\mathcal{F}:K]}$, the third equality is since $\phi_\Sigma^{\mathcal{F}}$ was already providing all information, the fourth equality is from the conditional property in item 2. of Fact 2.1, and the last inequality is a combination of the conditional property and item 4. in Fact 2.1. As $H_{u_\pi}(\psi|P) + H_{u_\pi}(\phi_\Sigma|\psi^{\mathcal{F}} \vee P) = \mathbb{E}_{K \sim \pi}[H_{u_K}(\psi) + H_{u_K}(\phi_\Sigma|\psi^{\mathcal{F}})]$, and π was assumed to sample only finite index subgroups, this case is clear from (4.3).

Assume now that π is supported on infinite index subgroups and satisfies RBS, and that $\psi: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \Omega$ is an observable. As ϕ_Σ is generating relative to P , it is also generating given P and $\psi^\mathcal{F}$. By the definition of Rokhlin entropy (Definition 2.4), we thus have

$$(4.4) \quad H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi) | \psi^\mathcal{F} \vee P) \leq H_{u_\pi}(\phi_\Sigma | \psi^\mathcal{F} \vee P).$$

Since π is supported on infinite index subgroups, and the p.m.p action $(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi)$ has stabilizer type π (Claim 3.10), it is in particular aperiodic. Moreover, the algebra $(\psi^\mathcal{F})^{-1}(\mathcal{B}_{\Omega^\mathcal{F}}) \vee P^{-1}(\mathcal{B}_{\mathbf{sub}(\mathcal{F})})$ is $\mathcal{F}$ -invariant. Corollary 1.4 in [AS21] states that in such a case, for every $\varepsilon > 0$ there is an observable $\alpha: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \Xi$ which generates this action given $\psi^\mathcal{F} \vee P$, namely $(\psi^\mathcal{F})^{-1}(\mathcal{B}_{\Omega^\mathcal{F}}) \vee P^{-1}(\mathcal{B}_{\mathbf{sub}(\mathcal{F})}) \vee (\alpha^\mathcal{F})^{-1}(\mathcal{B}_{\Xi^\mathcal{F}})$ is $\mathcal{B}_{\mathbf{sub}(\mathcal{F}, \Sigma)}$ up to u_π -null sets, and

$$(4.5) \quad H_{u_\pi}(\alpha) \leq H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi) | \psi^\mathcal{F} \vee P) + \varepsilon.$$

Note that this is somewhat surprising, and is indeed a non trivial fact. From Definition 2.4 it is clear that there is such an α for which $H(\alpha | \psi^\mathcal{F} \vee P)$ is bounded from above by the Rokhlin entropy plus ε , but the entropy of α itself without access to $\psi^\mathcal{F} \vee P$ is potentially much higher. Nevertheless, Alpeev and Seward showed that there is such an observable satisfying (4.5). Note that α generating given $\psi^\mathcal{F} \vee P$ is the same as $\psi \vee \alpha$ generating given P . Therefore,

$$\begin{aligned} H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi) | P) &\leq H_{u_\pi}(\psi \vee \alpha | P) \\ &= H_{u_\pi}(\psi | P) + \underbrace{H_{u_\pi}(\alpha | \psi \vee P)}_{\leq H_{u_\pi}(\alpha)} \\ &\stackrel{(4.5)}{\leq} H_{u_\pi}(\psi | P) + H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi) | \psi^\mathcal{F} \vee P) + \varepsilon \\ &\stackrel{(4.4)}{\leq} H_{u_\pi}(\psi | P) + H_{u_\pi}(\phi_\Sigma | \psi^\mathcal{F} \vee P) + \varepsilon. \end{aligned}$$

Since this is true for every $\varepsilon > 0$, and as we assumed π satisfies RBS — namely that $H_{\text{Rok}}(\mathcal{F} \curvearrowright (\mathbf{sub}(\mathcal{F}, \Sigma), u_\pi) | P) = \log |\Sigma|$ — we deduce (4.1) is satisfied in this case. $\square$

5. Outer approximation

Let us recall some notions from [BCLV24]. Let $\mathcal{F} = \mathcal{F}(S)$ be the free group with finite generating set S , and let $B \subseteq \mathcal{F}$ be a finite subset. For $Y, N \subseteq B$, let $C_B(Y, N) = \{A \subseteq B \mid Y \subseteq A \wedge N \cap A = \emptyset\} \subseteq \{0, 1\}^B$. A subset $A \subseteq B$ is a *pseudo subgroup* if there is a subgroup $K \leq \mathcal{F}$ such that $B \cap K = A$, and we denote the set of them by $P\mathbf{sub}(B)$. A random subset of B , say π , is a *pseudo IRS* if it is supported on pseudo subgroups and for every two subsets $Y, N \subseteq B$ which satisfy $sYs^{-1}, sNs^{-1} \subseteq B$ for every $s \in S \sqcup S^{-1}$, we have

$$\forall s \in S \sqcup S^{-1} : \pi(C_B(Y, N)) = \pi(C_B(sYs^{-1}, sNs^{-1})).$$

Denote by $\text{PIRS}(B) \subseteq \text{Prob}(\{0, 1\}^B)$ the pseudo IRSs of B . Recall from [BCLV24] the notation $\pi \cap B$ for the pushforward of π along the restriction map $R_B(K) = K \cap B$, and let $\widetilde{\text{PIRS}}(B)$ consist of all random subgroups π such that $\pi \cap B \in \text{PIRS}(B)$. In Section 2.2 of [BCLV24], it was shown that $\bigcap_{B \subseteq \mathcal{F}, |B| < \infty} \widetilde{\text{PIRS}}(B) = \text{IRS}(\mathcal{F})$.

Remark 5.1. Recalling a footnote from the Introduction, we do not think of a random subset (or a random subgroup) as a random variable with values in subsets, but as the distribution itself — this is a slight abuse of terminology, but is consistent throughout this paper (and [BCLV24]), so we keep with it here.

Here we define a natural subset of $\text{PIRS}(B)$ which captures a local form of (4.1), which was shown to be equivalent to RBS (Theorem 4.1). Let $W, W' \subseteq B$ be subsets; we think of these sets as “windows”. A (W, W') -continuous observable on B is a map $\psi: \{0, 1\}^B \times \Sigma^B \rightarrow \Omega$ such that ψ “only depends on the W, W' local view”; namely, if $(K_1, c_1), (K_2, c_2) \in \text{sub}(B, \Sigma)$ and $(K_1 \cap W, c_1|_{W'}) = (K_2 \cap W, c_2|_{W'})$, then $\psi(K_1, c_1) = \psi(K_2, c_2)$. If $\gamma \in \mathcal{F}$, let $\gamma.c: B \rightarrow \Sigma$ be

$$\forall x \in B : \quad \gamma.c(x) = \begin{cases} c(\gamma^{-1}x) & \gamma^{-1}x \in B, \\ \sigma & \gamma^{-1}x \notin B, \end{cases}$$

where σ is some fixed element in Σ . If $\gamma \in \mathcal{F}$ is such that $\gamma W \gamma^{-1}, \gamma W' \subseteq B$, we can define a $(\gamma W \gamma^{-1}, \gamma W')$ -continuous observable $\psi^\gamma: \{0, 1\}^B \times \Sigma^B \rightarrow \Omega$ by letting $\psi^\gamma(K, c) = \psi(\gamma^{-1}K \cap B, \gamma^{-1}.c)$; indeed, if $K_1 \cap \gamma W \gamma^{-1} = K_2 \cap \gamma W \gamma^{-1}$ and $c_1|_{\gamma W'} = c_2|_{\gamma W'}$, then $\gamma^{-1}K_1 \cap W = \gamma^{-1}K_2 \cap W$ and $\gamma^{-1}.c_1(w) = c_1(\gamma w) = c_2(\gamma w) = \gamma^{-1}.c_2(w)$ for every $w \in W'$, namely $\gamma^{-1}.c_1|_{W'} = \gamma^{-1}.c_2|_{W'}$, and thus

$$\begin{aligned} \psi^\gamma(K_1, c_1) &= \psi(\gamma^{-1}K_1 \cap B, \gamma^{-1}.c_1) \\ &=_{\psi \text{ is } (W, W') \text{ continuous}} \psi(\gamma^{-1}K_2 \cap B, \gamma^{-1}.c_2) \\ &= \psi^\gamma(K_2, c_2). \end{aligned}$$

If $A \subseteq \mathcal{F}$ is such that $\bigcup_{\gamma \in A} \gamma W \gamma^{-1}, \bigcup_{\gamma \in A} \gamma W' \subseteq B$, then we can define $\psi^A: \{0, 1\}^B \times \Sigma^B \rightarrow \Omega^A$ by $\psi^A = \bigvee_{\gamma \in A} \psi^\gamma$, which is $(\bigcup_{\gamma \in A} \gamma W \gamma^{-1}, \bigcup_{\gamma \in A} \gamma W')$ -continuous.

Let $K \subseteq B$ be a pseudo subgroup of B , and let Σ be a finite set. A coloring $c: B \rightarrow \Sigma$ is K -invariant if for every $x, y \in B$ for which $xy^{-1} \in K$ we have $c(x) = c(y)$. Let $P\text{sub}(B, \Sigma) \subseteq \{0, 1\}^B \times \Sigma^B$ consist of all pairs (K, c) where K is a pseudo subgroup of B and $c: B \rightarrow \Sigma$ is K -invariant. For a fixed $K \in P\text{sub}(B)$, let ν_K be the uniform measure over all K -invariant colorings $c: B \rightarrow \Sigma$ (this is the local analogue of u_K). If we abuse notation and denote by ${}_K \backslash^B$ the smallest equivalence classes on B induced by the relations $x \sim_K y$ if $xy^{-1} \in K$, then a K -invariant coloring is a map $c: {}_K \backslash^B \rightarrow \Sigma$. Given π a pseudo IRS on B , let ν_π be the probability distribution over $P\text{sub}(B, \Sigma)$ which first samples $K \sim \pi$ and then a K -invariant coloring c uniformly at random (this is the local analogue of u_π for an IRS π); namely, $\nu_\pi = \mathbb{E}_{K \sim \pi}[\mathbf{1}_K \times \nu_K]$. As usual, $P: P\text{sub}(B, \Sigma) \rightarrow P\text{sub}(B)$ is the projection to the first coordinate.

Now, because there is no natural group action in this local setup, we cannot discuss generating observables. But, the whole point of Theorem 4.1 was to translate the RBS condition to the continuous non-generating condition (4.1). The two components of (4.1) are $H_{u_\pi}(\psi|P)$ and $H_{u_\pi}(\phi_\Sigma|\psi^\mathcal{F} \vee P) = \inf_{A \subseteq \mathcal{F}, |A| < \infty} (H_{u_\pi}(\phi_\Sigma|\psi^A \vee P))$, and they have a local view version. If $\text{Id}_\mathcal{F} \in B$, let $\phi_\Sigma: \{0, 1\}^B \times \Sigma^B \rightarrow \Sigma$ be $\phi_\Sigma(K, c) = c(\text{Id}_\mathcal{F})$.

Definition 5.2. Let $\text{Id}_\mathcal{F} \in B$ be a finite subset of $\mathcal{F}$. A pseudo IRS π of B is *locally satisfying the RBS condition* if for every three subsets $W, W', A \subseteq B$, finite sets Σ, Ω and function $\psi: \{0, 1\}^B \times \Sigma^B \rightarrow \Omega$ satisfying

- (1) for all $\gamma \in A$ we have $\gamma W \gamma^{-1}, \gamma W' \subseteq B$;
- (2) $W'W'^{-1} \subseteq W$;
- (3) $|\Sigma|, |\Omega| \leq |B|$;
- (4) ψ is (W, W') -continuous;

we have

$$(5.1) \quad H_{\nu_\pi}(\psi|P) + H_{\nu_\pi}(\phi_\Sigma|\psi^A \vee P) \geq \log |\Sigma|.$$

We denote the set of these by $PRBS(B)$, and the set of random subgroups π such that $\pi \cap B$ is in $PRBS(B)$ by $\widetilde{PRBS}(B)$.

Theorem 5.3. *For every $\text{Id}_{\mathcal{F}} \in B \subseteq \mathcal{F}$ where $|B| \leq \mathcal{F}$, the set $PRBS(B)$ of pseudo IRSs which locally satisfy RBS is a compact convex polytope in $\mathbb{R}^{\{0,1\}^B}$ whose defining half spaces are computable from B , and*

$$\bigcap_{\substack{\text{Id}_{\mathcal{F}} \in B \subseteq \mathcal{F} \\ |B| < \infty}} \widetilde{PRBS}(B) = \text{IRS}_{\text{RBS}}(\mathcal{F}),$$

where $\text{IRS}_{\text{RBS}}(\mathcal{F})$ is the set of IRSs satisfying the RBS condition (Definition 3.5). In particular, $\text{IRS}_{\text{RBS}}(\mathcal{F})$ is weak* closed.

Before proving Theorem 5.3 let us spell an immediate corollary of it:

Corollary 5.4. *Every cosofic IRS (Definition 1.4) satisfies the RBS criterion (Definition 3.5).*

Proof. By Definition 3.5, every finitely described IRS of $\mathcal{F}$ is in $\text{IRS}_{\text{RBS}}(\mathcal{F})$. By Theorem 5.3, $\text{IRS}_{\text{RBS}}(\mathcal{F})$ is weak* closed. Therefore, the weak* closure of $\text{IRS}_{\text{fd}}(\mathcal{F})$, which is the set of cosofic IRSs, is contained in $\text{IRS}_{\text{RBS}}(\mathcal{F})$. $\square$

Proof of Theorem 5.3. In [BCLV24], we showed that $PIRS(B)$ is a compact convex polytope with halfspaces computable from B . The set $PRBS(B)$ is a subset of $PIRS(B)$, and is defined by satisfying extra conditions of the form (5.1); these are linear inequalities on $\pi: \{0,1\}^B \rightarrow \mathbb{R}$, since:¹³

$$(5.2) \quad \begin{aligned} H_{\nu_{\pi}}(\psi|P) &= \mathbb{E}_{K \sim \pi} [H_{\mathbf{1}_K \times \nu_K}(\psi)] = \sum_{K \subseteq B} H_{\mathbf{1}_K \times \nu_K}(\psi) \cdot \pi(K), \\ H_{\nu_{\pi}}(\phi_{\Sigma}|\psi^A \vee P) &= \mathbb{E}_{K \sim \pi} [H_{\mathbf{1}_K \times \nu_K}(\phi_{\Sigma}|\psi^A)] = \sum_{K \subseteq B} H_{\mathbf{1}_K \times \nu_K}(\phi_{\Sigma}|\psi^A) \cdot \pi(K). \end{aligned}$$

Moreover, recalling the notation μ_y^{ψ} for conditional distributions from Section 2.1, for $K \subseteq B$ we have

$$\begin{aligned} \forall a \in \Omega : \quad \psi_* \mathbf{1}_K \times \nu_K(a) &= \mathbb{P}_{c \sim \nu_K} [\psi(K, c) = a] \\ &= \frac{|\{c: K \setminus B \rightarrow \Sigma \mid \psi(K, c) = a\}|}{|\Sigma|^{|K \setminus B|}}, \\ \forall a \in \Omega, \alpha \in \Omega^A : \quad \phi_{\Sigma*}(\mathbf{1}_K \times \nu_K)_{\alpha}^{\psi^A}(a) &= \mathbb{P}_{c \sim \nu_K} [\phi_{\Sigma}(K, c) = a \mid \psi^A(c) = \alpha] \\ &= \frac{|\{c: K \setminus B \rightarrow \Sigma \mid \phi_{\Sigma}(c) = a \wedge \psi^A(c) = \alpha\}|}{|\{c: K \setminus B \rightarrow \Sigma \mid \psi^A(c) = \alpha\}|}. \end{aligned}$$

These distributions are thus computable given ψ , A and K , and it is straightforward to calculate their appropriate entropies. Thus, the coefficients in (5.2) are computable from ψ , A and K , which demonstrates that the inequality

¹³The following is actually a more general phenomenon: The map from π to ν_{π} is convex-linear in the sense that if, for $i = 1, 2$, $\pi^{(i)}$ are given and $t \in [0, 1]$, then

$$\nu_{t\pi_1 + (1-t)\pi_2} = t\nu_{\pi_1} + (1-t)\nu_{\pi_2}.$$

Furthermore, Shannon entropy is convex linear as a function of the measure. That is

$$H_{t\mu_1 + (1-t)\mu_2}(\alpha|\mathcal{F}) = tH_{\mu_1}(\alpha|\mathcal{F}) + (1-t)H_{\mu_2}(\alpha|\mathcal{F})$$

in general (for any measures μ_1, μ_2 on the same space and any fixed $\alpha, \mathcal{F}$). Since the composition of convex linear functions is convex linear, it follows that the conditions of (5.1) are indeed linear inequalities.

(5.1) is computable from them. Since there are only finitely many W, W', A, Σ, Ω and ψ for every B ,¹⁴ the sets $\text{PRBS}(B)$ are defined by some finite set of extra computable linear inequalities, and are therefore compact convex computable polytopes.

Assume the IRS π satisfies the RBS criterion. Then, $\pi \cap B$ is a pseudo IRS by [BCLV24]. Let W, W', A, Σ, Ω and ψ be as in the theorem assumptions. Let $\tilde{\psi}: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \Omega$ be the following observable:

$$\tilde{\psi}(K, c) = \psi(K \cap B, c|_B).$$

The observable $\tilde{\psi}$ is continuous by definition, as it depends only on the restrictions of K and c to B . Since $W'W'^{-1} \subseteq W$, the distributions $\nu_{\pi \cap B} \cap (W, W')$ and $u_\pi \cap (W, W')$ are identical — for $K \leq \mathcal{F}$, the local cosets $\langle K \cap B \rangle \setminus^B$ are a priori finer than the intersection $(K \setminus^B)$, since there may be $x, y \in B$ such that $x \not\sim_{\langle K \cap B \rangle} y$, as there is no sequence of $z_1, \dots, z_t \in B$ for which $xz_1^{-1}, z_1z_2^{-1}, \dots, z_t y^{-1} \in K \cap B$, but $xy^{-1} \in K$. So, a $K \cap B$ -invariant coloring is a looser condition than being the restriction to B of a K -invariant coloring. Nevertheless, this is not the case for the restrictions of these colorings to W' , as we forced $W \subseteq B$ to include all words of the form xy^{-1} from W' . In particular, the distributions u_π and $\nu_{\pi \cap B}$ agree when restricted to the (W, W') local view. Furthermore, as $\gamma W \gamma^{-1}, \gamma W' \subseteq B$ for every $\gamma \in A$, we have that $\tilde{\psi}^\gamma(K, c) = \psi^\gamma(K \cap B, c|_B)$, and thus $\tilde{\psi}^A(K, c) = \psi^A(K \cap B, c|_B)$. All in all, we get that $H_{\nu_{\pi \cap B}}(\psi|P) = H_{u_\pi}(\tilde{\psi}|P)$ and $H_{\nu_{\pi \cap B}}(\phi_\Sigma|\psi^A \vee P) = H_{u_\pi}(\phi_\Sigma|\tilde{\psi}^A \vee P)$. Since π satisfies RBS, by Theorem 4.1 we have

$$\begin{aligned} \log |\Sigma| &\leq H_{u_\pi}(\tilde{\psi}|P) + H_{u_\pi}(\phi_\Sigma|\tilde{\psi}^\mathcal{F} \vee P) \\ &\leq H_{u_\pi}(\tilde{\psi}|P) + H_{u_\pi}(\phi_\Sigma|\tilde{\psi}^A \vee P) \\ &\leq H_{\nu_{\pi \cap B}}(\psi|P) + H_{\nu_{\pi \cap B}}(\phi_\Sigma|\psi^A \vee P). \end{aligned}$$

As this analysis was independent of the specific W, W', A, Σ, Ω and ψ we started with, $\pi \cap B$ satisfies all the defining relations of $\text{PRBS}(B)$ for every B , and is thus in $\bigcap_{\substack{\text{Id}_{\mathcal{F}} \in B \subseteq \mathcal{F} \\ |B| < \infty}} \widetilde{\text{PRBS}}(B)$.

Assume the IRS π does not satisfy RBS. Then, by Theorem 4.1, there is some set Σ and a continuous observable $\psi: \mathbf{sub}(\mathcal{F}, \Sigma) \rightarrow \Omega$ such that

$$H_{u_\pi}(\psi|P) + H_{u_\pi}(\phi_\Sigma|\psi^\mathcal{F} \vee P) < \log |\Sigma|.$$

By item 5. in Fact 2.1, $H_{u_\pi}(\phi_\Sigma|\psi^\mathcal{F} \vee P) = \inf_{A \subseteq \mathcal{F}, |A| < \infty} H_{u_\pi}(\phi_\Sigma|\psi^A \vee P)$, and there is thus some large enough (yet finite) A for which

$$H_{u_\pi}(\psi|P) + H_{u_\pi}(\phi_\Sigma|\psi^A \vee P) < \log |\Sigma|.$$

Since ψ is continuous, there are finite sets $W, W' \subseteq \mathcal{F}$ such that $\psi(K_1, c_1) = \psi(K_2, c_2)$ whenever $(K_1 \cap W, c_1|_{W'}) = (K_2 \cap W, c_2|_{W'})$. Without loss of generality, we can assume W contains $W'W'^{-1}$. Let B be a large enough finite subset of $\mathcal{F}$ such that $W, W', \gamma W \gamma^{-1}, \gamma W' \subseteq B$ for every $\gamma \in A$, and $|\Sigma|, |\Omega| \leq |B|$. Let $\hat{\psi}: \{0, 1\}^B \times \Sigma^B \rightarrow \Omega$ be

$$\hat{\psi}(K, c) = \begin{cases} \psi(\tilde{K}, \tilde{c}) & \exists (\tilde{K}, \tilde{c}) \in \mathbf{sub}(\mathcal{F}, \Sigma): K \cap W = \tilde{K} \cap W \wedge c|_{W'} = \tilde{c}|_{W'}, \\ \omega & \text{otherwise,} \end{cases}$$

¹⁴We are a bit sloppy here. There are infinitely many sets Ω, Σ such that $|\Omega|, |\Sigma| \leq |B|$. But, we can assume the sets we use are always of the form $\{1, \dots, k\}$ for $k \leq |B|$, and they will induce the same sets $\text{PRBS}(B)$, as the restrictions depend only on the size of Σ and Ω and not the sets themselves.

where ω is some fixed element of Ω . In the notations of the first direction of this proof, we chose $\widehat{\psi}$ so that $\widetilde{(\widehat{\psi})} = \psi$. As so, by our choice of B and the analysis in the proof of the first direction, we have $H_{\nu_{\pi \cap B}}(\widehat{\psi}|P) = H_{u_\pi}(\psi|P)$ and $H_{\nu_{\pi \cap B}}(\phi_\Sigma|\widehat{\psi}^A \vee P) = H_{u_\pi}(\phi_\Sigma|\psi^A \vee P)$ and thus

$$H_{\nu_{\pi \cap B}}(\widehat{\psi}|P) + H_{\nu_{\pi \cap B}}(\phi_\Sigma|\widehat{\psi}^A \vee P) < \log |\Sigma|.$$

Hence, $\pi \cap B \notin \text{PRBS}(B)$ and we are done. $\square$

To prove Theorem 1.2 in the next section, we would need to run linear programs on the polytopes $\text{PRBS}(B)$. As opposed to [BCLV24], where all equations (and inequalities) defining the pseudo IRSs $\text{PIRS}(B)$ were integral, the sets $\text{PRBS}(B)$ have irrational coefficients in their defining inequalities. This means that the algorithm calculating the optimum over these sets would first need to approximate them, and that leads to errors. There are standard continuity results for linear programming (cf. [Wet85]), which can still guarantee the algorithmic outer approximation we need. But, we decided to be more direct and define an alternative outer approximation to $\text{PRBS}(B)$ which uses only rational coefficients in its defining inequalities.

Given a real number a and a positive integer n , let $\rho(a, n) = \frac{\lceil 10^n \cdot a \rceil}{10^n}$ — this is the n -digits upper decimal expansion of a . Similarly, for an inequality $\alpha: \sum a_i X_i + b \geq 0$ with real coefficients a_i, b and variables X_i , we let $\rho(\alpha, n): \sum \rho(a_i, n) X_i + \rho(b, n) \geq 0$. Note that a non-negative assignment $\forall i, X_i = v_i \geq 0$ which satisfies α must also satisfy $\rho(\alpha, n)$ for every n , since $\rho(b, n) \geq b$ and $\rho(a_i, n) X_i \geq a_i X_i$. Let $\text{PRBS}(B, n)$ be $\text{PIRS}(B)$ with the inequalities $\rho((5.1), n)$ ¹⁵ instead of (5.1) for every relevant tuple $W, W', A, \Sigma, \Omega, \psi$. For $B \subseteq C \subseteq \mathcal{F}$, and $A \subseteq C$, recall the notation $R_{B \subseteq C}(A) = A \cap B \subseteq B$ for the restriction map from C to B . The following is quite immediate, and we leave it without a proof:

Fact 5.5. *Let $\text{Id}_{\mathcal{F}} \in B \subseteq C \subseteq \mathcal{F}$ and $n \in \mathbb{N}$ such that $|B|, |C| < \infty$. Then:*

- $\text{PRBS}(B, n+1) \subseteq \text{PRBS}(B, n)$;
- $R_{B \subseteq C} \text{PRBS}(C) = \{R_{B \subseteq C} \pi \mid \pi \in \text{PRBS}(C)\} \subseteq \text{PRBS}(B)$;
- $R_{B \subseteq C} \text{PRBS}(C, n) = \{R_{B \subseteq C} \pi \mid \pi \in \text{PRBS}(C, n)\} \subseteq \text{PRBS}(B, n)$;
- $\bigcap_{n=1}^{\infty} \text{PRBS}(B, n) = \text{PRBS}(B)$.

By combining these observations with Theorem 5.3, and by denoting $\widetilde{\text{PRBS}}(B, n)$ for all random subgroups $\pi \in \text{Prob}(\text{sub}(\mathcal{F}))$ satisfying $R_{B*} \pi = \pi \cap B \in \text{PRBS}(B, n)$, we have

$$\bigcap_{\substack{\text{Id}_{\mathcal{F}} \in B \subseteq \mathcal{F} \\ |B| < \infty}} \widetilde{\text{PRBS}}(B, |B|) = \text{IRS}_{\text{RBS}}(\mathcal{F}).$$

6. Proof of Theorem 1.2

Recall the notion of a subgroup test $\mathcal{T}$ (Section 1.2.1) consisting of challenges $\{D_i\}_{i=1}^m$ and a rational probability distribution $\mu \in \text{Prob}[m]$, and specifically its value versus an IRS defined in equation (1.2). As challenges are continuous functions $D_i: \text{sub}(\mathcal{F}) \rightarrow \{0, 1\}$, there are finite subsets $\Delta_i \subseteq \mathcal{F}$ and functions $d_i: P\text{sub}(\Delta_i) \rightarrow \{0, 1\}$ such that $D_i(K) = d_i(K \cap \Delta_i)$. Let $\Delta = \bigcup_{i=1}^m \Delta_i$ be the *mutual continuity domain* of $\mathcal{T}$. For every finite subset $B \subseteq \mathcal{F}$ such that $\Delta \subseteq B$, one can define a value of a random pseudo subgroup π of B by letting

$$(6.1) \quad \text{val}(\mathcal{T}, \pi) = \mathbb{E}_{i \sim \mu} \mathbb{E}_{K \sim \pi} [d_i(K \cap \Delta_i)].$$

¹⁵As the constant $\log |\Sigma|$ is on the right hand side of (5.1), one first needs to move it to the left and then apply $\rho(\cdot, n)$.

It is straightforward to check that $\text{val}(\mathcal{T}, \pi)$ as in (1.2) agrees with $\text{val}(\mathcal{T}, \pi \cap B)$ as in (6.1), and so it makes sense to use the same notation. Let $\text{val}_{\text{RBS}}(\mathcal{T})$ be the supremum of $\text{val}(\mathcal{T}, \pi)$ when running over all IRSs π which satisfy the RBS criterion.

Theorem 6.1. *There is an algorithm which takes as input a subgroup test $\mathcal{T}$ and outputs a non-increasing sequence of rationals $\{\rho_r\}_{r=r_0}^\infty$ such that $\lim_{r \rightarrow \infty} \rho_r = \text{val}_{\text{RBS}}(\mathcal{T})$.*

Proof. Let Δ be the mutual continuity domain of $\mathcal{T}$. Now, for every finite set $B \subseteq \mathcal{F}$ which contains Δ and $\text{Id}_{\mathcal{F}}$, $\text{val}(\mathcal{T}, \cdot)$ induces a linear functional from $\mathbb{R}^{\{0,1\}^B}$ to $\mathbb{R}$. Maximizing this functional over a compact computable rational polytope of $\mathbb{R}^{\{0,1\}^B}$ can be done in finite time (this is exactly the task of linear programming). Hence, finding $\rho_B = \max\{\text{val}(\mathcal{T}, \pi) \mid \pi \in \text{PRBS}(B, |B|)\}$ can be calculated in finite time. As Δ is a finite set of words in the free group, there is an integer r_0 which is the length of the longest word in Δ . So, $\Delta \subseteq B_r(\text{Id}_{\mathcal{F}})$ for every $r \geq r_0$. Let $\rho_r = \rho_{B_r}$ for every $r \geq r_0$. By Fact 5.5 and the fact (1.2) and (6.1) agree, the limit of the sequence ρ_r is indeed $\text{val}_{\text{RBS}}(\mathcal{T})$. $\square$

By combining Theorem 6.1 with Fact 1.8 and Fact 1.10, we deduce Theorem 1.2. $\square$

REFERENCES

- [ABB⁺17] M. Abert, N. Bergeron, I. Biringer, T. Gelander, N. Nikolov, J. Raimbault, and I. Samet, *On the growth of L^2 -invariants for sequences of lattices in Lie groups*, Ann. of Math. (2) **185** (2017), no. 3, 711–790. MR3664810 [↑1](#)
- [AGV14] M. Abért, Y. Glasner, and B. Virág, *Kesten’s theorem for invariant random subgroups*, Duke Mathematical Journal **163** (2014), no. 3, 465–488. [↑1](#)
- [AS21] A. Alpeev and B. Seward, *Krieger’s finite generator theorem for actions of countable groups III*, Ergodic Theory Dynam. Systems **41** (2021), no. 10, 2881–2917. MR4308159 [↑16](#)
- [BCLV24] L. Bowen, M. Chapman, A. Lubotzky, and T. Vidick, *The Aldous–Lyons Conjecture I: Subgroup Tests*, preprint (2024). [↑1](#), [3](#), [4](#), [16](#), [18](#), [19](#), [20](#)
- [BCV24] L. Bowen, M. Chapman, and T. Vidick, *The Aldous–Lyons Conjecture II: Undecidability*, preprint (2024). [↑1](#), [3](#)
- [Bow10] L. Bowen, *Measure conjugacy invariants for actions of countable sofic groups*, Journal of the American Mathematical Society **23** (2010), no. 1, 217–245. [↑6](#)
- [Bow14] L. Bowen, *Random walks on random coset spaces with applications to Furstenberg entropy*, Invent. Math. **196** (2014), no. 2, 485–510. MR3193754 [↑1](#)
- [Bow18] L. P. Bowen, *A brief introduction to sofic entropy theory*, Proceedings of the International Congress of Mathematicians (ICM 2018), 2018, pp. 1847–1866. [↑6](#)
- [Dow11] T. Downarowicz, *Entropy in dynamical systems*, Vol. 18, Cambridge University Press, 2011. [↑8](#)
- [EW11] M. Einsiedler and T. Ward, *Ergodic theory: with a view towards number theory*, Graduate Texts in Mathematics, Springer, London, 2011. [↑7](#)

- [FG23] M. Fraczyk and T. Gelander, *Infinite volume and infinite injectivity radius*, Annals of Mathematics **197** (2023), no. 1, 389–421. [↑1](#)
- [Gel18] T. Gelander, *A view on invariant random subgroups and lattices*, Proceedings of the international congress of mathematicians: Rio de janeiro 2018. World Scientific, 2018, pp. 1321–1344. [↑1](#)
- [Got73] W. Gottschalk, *Some general dynamical notions*, Recent Advances in Topological Dynamics: Proceedings of the Conference on Topological Dynamics, held at Yale University, June 19–23, 1972, in honor of Professor Gustav Arnold Hedlund on the occasion of his retirement, 1973, pp. 120–125. [↑1](#)
- [Gro99] M. Gromov, *Endomorphisms of symbolic algebraic varieties*, J. Eur. Math. Soc. (JEMS) **1** (1999), no. 2, 109–197. MR1694588 [↑1](#)
- [KL11] D. Kerr and H. Li, *Entropy and the variational principle for actions of sofic groups*, Invent. Math. **186** (2011), no. 3, 501–558. MR2854085 [↑6](#)
- [Lin25] J. Lin, *Mipco= core*, arXiv preprint arXiv:2510.07162 (2025). [↑5](#)
- [Lüc02] W. Lück, *L2-invariants: theory and applications to geometry and K-theory*, Vol. 44, Springer, 2002. [↑3](#), [4](#)
- [Man25a] A. Manzoor, *Invariant Random Subgroups, Soficity, and Lück’s determinant conjecture*, arXiv preprint arXiv:2508.15154 (2025). [↑3](#), [4](#), [5](#)
- [Man25b] A. Manzoor, *There is an equivalence relation whose von neumann algebra is not connes embeddable*, arXiv preprint arXiv:2502.06697 (2025). [↑5](#)
- [NPA08] M. Navascués, S. Pironio, and A. Acín, *A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations*, New Journal of Physics **10** (2008), no. 7, 073013. [↑4](#)
- [Sew16] B. Seward, *Weak containment and Rokhlin entropy*, arXiv preprint arXiv:1602.06680 (2016). [↑7](#), [8](#), [13](#)
- [Sew19a] B. Seward, *Krieger’s finite generator theorem for actions of countable groups I*, Inventiones mathematicae **215** (2019), no. 1, 265–310. [↑6](#)
- [Sew19b] B. Seward, *Krieger’s finite generator theorem for actions of countable groups II*, J. Mod. Dyn **15** (2019), 1–39. [↑3](#), [5](#), [6](#), [7](#), [8](#), [9](#)
- [SZ94] G. Stuck and R. J. Zimmer, *Stabilizers for ergodic actions of higher rank semisimple groups*, Annals of Mathematics-Second Series **139** (1994), no. 3, 723–748. [↑1](#)
- [Ver10] A. M. Vershik, *Nonfree actions of countable groups and their characters*, Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI) **378** (2010), no. Teoriya Predstavleniĭ, Dinamicheskie Sistemy, Kombinatornye Metody. XVIII, 5–16, 228. MR2749291 [↑1](#)
- [Wei00] B. Weiss, *Sofic groups and dynamical systems*, Sankhya Ser. A **62** (2000), no. 3, 350–359. Ergodic theory and harmonic analysis (Mumbai, 1999). MR1803462 (2001j:37022) [↑1](#)
- [Wet85] R. J-B Wets, *On the continuity of the value of a linear program and of related polyhedral-valued multifunctions*, Mathematical Programming Essays in Honor of George B. Dantzig Part I, 1985, pp. 14–29. [↑20](#)

LEWIS BOWEN

DEPARTMENT OF MATHEMATICS

1 UNIVERSITY STATION C1200

UNIVERSITY OF TEXAS AT AUSTIN

AUSTIN, TX, 78712 USA.

Email address: `lpbowen@math.utexas.edu`

MICHAEL CHAPMAN

SCHOOL OF MATHEMATICS

INSTITUTE FOR ADVANCED STUDY,

1 EINSTEIN DRIVE, PRINCETON, NJ 08540, USA.

Email address: `mchapman@ias.edu`