



PROJECT MUSE®

Sofic homological invariants and the Weak Pinsker Property

Lewis Bowen

American Journal of Mathematics, Volume 144, Number 1, February 2022,
pp. 169-226 (Article)

Published by Johns Hopkins University Press

DOI: <https://doi.org/10.1353/ajm.2022.0003>

AMERICAN JOURNAL
OF MATHEMATICS



Published under the aegis of
JOHNS HOPKINS UNIVERSITY PRESS

JOHNS HOPKINS
UNIVERSITY PRESS

➡ For additional information about this article

<https://muse.jhu.edu/article/844461/summary>

SOFIC HOMOLOGICAL INVARIANTS AND THE WEAK PINSKER PROPERTY

By LEWIS BOWEN

Abstract. A probability-measure-preserving transformation has the Weak Pinsker Property (WPP) if for every $\epsilon > 0$ it is measurably conjugate to the direct product of a transformation with entropy $< \epsilon$ and a Bernoulli shift. In a recent breakthrough, Tim Austin proved that every ergodic transformation satisfies this property. Moreover, the natural analog for amenable group actions is also true. By contrast, this paper provides a counterexample in which the group Γ is a non-abelian free group and the notion of entropy is sofic entropy. The counterexample is a limit of hardcore models on random regular graphs. In order to prove that it does not have the WPP, this paper introduces new measure conjugacy invariants based on the growth of homology of the model spaces of the action. The main result is obtained by showing that any action with the WPP has subexponential homology growth in dimension 0, while the counterexample has exponential homology growth in dimension 0.

Contents.

1. Introduction.
 - 1.1. Homological measure-conjugacy invariants.
 - 1.2. An action with positive zero-dimensional homology growth.
 - 1.3. A brief guide to the paper.
2. Notation and conventions.
3. A review of sofic entropy.
 - 3.1. Sofic approximations.
 - 3.2. Sofic entropy.
 4. Sofic homology.
 - 4.1. Homology theory on the Hamming cube.
 - 4.2. Special subgroups defined by a sofic approximation.
 - 4.3. Main results.
 5. Preliminaries to the proof of Theorem 4.3.
 - 5.1. Almost Lipschitz maps.
 - 5.2. Equivariant maps and their approximations.
 - 5.3. Sofic models.
 6. Proof of Theorem 4.3.
 7. Homology computations.
 - 7.1. Contractible model spaces.
 - 7.2. Edit distance.
 - 7.3. Diffuse sofic approximations.
 - 7.3.1. Diffuse approximations and contractibility.
 - 7.4. Bernoulli shifts.

Manuscript received February 4, 2019; revised February 24, 2020.

Research supported in part by NSF grant DMS-1900386 and a fellowship from the Simons Foundation.

American Journal of Mathematics 144 (2022), 169–226. © 2022 by Johns Hopkins University Press.

- 8. An action without the Weak Pinsker Property.
- 8.1. The configuration model and the permutation model.
- 8.1.1. Expectations.
- 8.2. First moment computations.
- 8.3. The planted model and the uniform model.
- 8.4. Bounding clusters of independent sets.
- 8.5. A variational principle.
- 8.6. Proof of Theorem 1.2.
- 9. Questions.

References.

1. Introduction. This paper is concerned with the general problem of classifying measure-preserving actions of countable groups on probability spaces. To be precise, fix a countable group Γ and let $(X, \mu_X), (Y, \mu_Y)$ be standard probability spaces. Then two actions $\Gamma \curvearrowright (X, \mu_X), \Gamma \curvearrowright (Y, \mu_Y)$ are *measurably conjugate* or *isomorphic* if there exists a measure-preserving isomorphism $\Phi : (X, \mu_X) \rightarrow (Y, \mu_Y)$ that intertwines the actions in the sense that $\Phi(gx) = g\Phi(x)$ for every $g \in \Gamma$ and a.e. $x \in X$.

Anti-classification theorems convincingly show it is not possible to classify all actions up to measure-conjugacy [17, 16]. In spite of this, there are interesting structural results. To explain these, it is necessary to introduce Bernoulli shifts, which are some of the most fundamental actions. Let (K, κ) be a standard probability space and equip the product space $K^\Gamma = \{x : \Gamma \rightarrow K\}$ with the product measure κ^Γ . The group acts on this space by $(gx)(f) = x(g^{-1}f)$. This action is called the *Bernoulli shift over Γ with base (K, κ)* .

A classical example of a general structural result is Sinai's factor Theorem. It states that, when $\Gamma = \mathbb{Z}$, any action with positive entropy factors onto a Bernoulli shift. Moreover, the factor can be chosen so that the relative entropy is zero. These statements have recently been generalized to arbitrary countable groups by Seward [35].

Another example comes from Pinsker. In 1960, Pinsker conjectured that any ergodic measure-preserving transformation $T : X \rightarrow X$ of a standard probability space (X, μ) is measurably conjugate to a direct product $T \approx S \times U$ such that S has zero entropy and U is a K -transformation (which means that every nontrivial factor of U has positive entropy) [31]. This was falsified by Ornstein [30, 29]. The study of such systems led Thouvenot to introduce the *Weak Pinsker Property* (WPP) for measure-preserving transformations: T has the WPP if for every $\epsilon > 0$, T is measurably conjugate to a direct product $S \times U$ such that S has entropy $< \epsilon$ and U is isomorphic to a Bernoulli shift. He asked whether all ergodic transformations have the WPP and proved important structural properties of this class [37].

In recent breakthrough work, Tim Austin has proven that indeed every ergodic transformation has the WPP [3]. Moreover, the analogous statement for measure-preserving actions of amenable groups is also true.

The purpose of this paper is to give an example of an ergodic action of a non-abelian free group without the WPP. In this context there are two main entropy notions: sofic and Rokhlin. Sofic entropy was initiated in [9] and Rokhlin entropy in [33] (see [10, 11] for an introduction and survey). This paper uses sofic entropy although the results also apply with Rokhlin entropy because it upper bounds sofic entropy.

The sofic entropy of a measure-preserving action $\mathbf{a} := \Gamma \curvearrowright (X, \mu)$ depends apriori on a choice of sofic approximation Σ to Γ . So it will be referred to here as Σ -entropy and denoted by $h_\Sigma(\mathbf{a})$. The action has the *Weak Pinsker Property (WPP)* with respect to Σ if for every $\epsilon > 0$, \mathbf{a} is isomorphic to a direct product $\mathbf{b} \times \mathbf{c}$ such that \mathbf{c} has Σ -entropy $< \epsilon$ and \mathbf{b} is isomorphic to a Bernoulli shift. For example, if \mathbf{a} has nonpositive Σ -entropy, then it automatically has the WPP because \mathbf{b} is allowed to be the trivial action (which can be thought of as a Bernoulli shift with trivial base space). The main result of this paper is:

THEOREM 1.1. *Let \mathbb{F}_r denote the free group of rank r . Then there is an r_0 such that for all $r > r_0$, there exists a sofic approximation Σ to \mathbb{F}_r and an ergodic essentially free action $\mathbb{F}_r \curvearrowright (X, \mu)$ that does not have the Weak Pinsker Property with respect to Σ .*

Remark 1. An action \mathbf{a} has the *Weak Pinsker Property with respect to Rokhlin entropy* if for every $\epsilon > 0$, \mathbf{a} is isomorphic to a direct product $\mathbf{b} \times \mathbf{c}$ such that \mathbf{b} has Rokhlin entropy $< \epsilon$ and \mathbf{c} is isomorphic to a Bernoulli shift. Because Rokhlin entropy upper bounds Σ -entropy (for every Σ), this property is apriori stronger than the WPP with respect to Σ . In particular, the action in Theorem 1.1 does not have the WPP with respect to Rokhlin entropy.

1.1. Homological measure-conjugacy invariants. The proof of Theorem 1.1 is in two steps, the first of which is a construction of a family of new measure-conjugacy invariants based on the asymptotic homology of model spaces. Here is a brief sketch in the special case that μ is a shift-invariant measure on \mathcal{X}^Γ where \mathcal{X} is a finite alphabet. In this case, the sofic approximation Σ is a sequence $\Sigma = \{\sigma_n\}_{n \in \mathbb{N}}$ of maps $\sigma_n : \Gamma \rightarrow \text{sym}(V_n)$ where each V_n is a finite set and $\text{sym}(V_n)$ is the symmetric group of V_n . For every open neighborhood \mathcal{O} of μ in the space of probability measures on \mathcal{X}^Γ there is a subset $\Omega(\mathcal{O}, \sigma_n) \subset \mathcal{X}^{V_n}$ consisting of vertex-labelings whose “empirical measure” is in \mathcal{O} . The sets $\Omega(\mathcal{O}, \sigma_n)$ equipped with the normalized Hamming metric are called *model spaces*. The Σ -entropy is the exponential rate of growth of the cardinalities of these model spaces.

Given a bound $\kappa > 0$, each model space $\Omega(\mathcal{O}, \sigma_n)$ is the vertex set of a simplicial complex whose d -simplices consist of subsets $S \subset \Omega(\mathcal{O}, \sigma_n)$ of cardinality $d + 1$ such that the distance between any two elements of S is bounded by κ . Homology is usually defined as cycles mod boundaries. That is also true here with the caveat that the boundaries are defined using parameters $\kappa' \geq \kappa$ and $\mathcal{O}' \supset \mathcal{O}$ in place of κ, \mathcal{O} . So the homology group of the n -th model space depends on four parameters

$\kappa, \kappa', \mathcal{O}, \mathcal{O}'$ in addition to σ_n . The asymptotic behavior of these homology groups provide new invariants. This idea was inspired by Tim Austin's paper [2] which gave an asymptotic notion of connectedness for model spaces. That notion is equivalent to the asymptotic triviality of the 0-dimensional homology groups.

One of the new invariants, denoted $b_{\Sigma,0}(\mathbf{a})$, is the exponential growth rate of the 0-th betti numbers of the model spaces. Intuitively, it estimates the growth rate of the number of “clusters” of good models. If an action \mathbf{a} has the Weak Pinsker Property with respect to Σ then $b_{\Sigma,0}(\mathbf{a}) = 0$. This is because the model spaces for a direct product of the form $\mathbf{b} \times \mathbf{c}$ where \mathbf{b} is Bernoulli contract (in a coarse sense) to model spaces for \mathbf{c} and $b_{\Sigma,0}(\mathbf{c})$ is bounded by the Σ -entropy of \mathbf{c} .

1.2. An action with positive zero-dimensional homology growth. To finish the proof of Theorem 1.1, the next result suffices.

THEOREM 1.2. *There exists r_0 such that if $r > r_0$ then there exists a sofic approximation Σ to \mathbb{F}_r and an invariant measure μ on the shift space $\{0,1\}^{\mathbb{F}_r}$ such that $b_{\Sigma,0}(\mathbb{F}_r \curvearrowright (\{0,1\}^{\mathbb{F}_r}, \mu)) > 0$. In particular, $\Gamma \curvearrowright (\{0,1\}^{\Gamma}, \mu)$ does not have the Weak Pinsker Property.*

The example is based on the geometry of the space of independent subsets of random regular graphs. To be precise, let $G = (V, E)$ be a graph. A subset $W \subset V$ is *independent* if there does not exist an edge between any two vertices of W . The *density* of W is $\#W/\#V$. The maximum density of an independent set is denoted $\alpha(G)$.

Fix an even integer $d \geq 3$ and consider choosing a d -regular graph $\mathbf{G}_{d,n}$ on n vertices uniformly at random (amongst all d -regular graphs on n vertices). The first moment method shows that $\alpha(\mathbf{G}_{d,n})$ is bounded above by $2\log(d)/d + o(\log(d)/d)$ with high probability as $n \rightarrow \infty$ [7]. By a non-constructive argument using Azuma's inequality, Frieze-Łuczak obtained a matching lower bound [18]. More recently, it was shown in [5] that the limit $\lim_{n \rightarrow \infty} \mathbb{E}[\alpha(\mathbf{G}_{d,n})]$ exists and an explicit formula was obtained in [13] by a deep study of the structure of high density independent sets.

There are no known polynomial-time algorithms for constructing independent subsets of $\mathbf{G}_{d,n}$ with density larger than $\log(d)/d$. It is argued in [12] that a reason for this is that there are many independent subsets I with density between $\log(d)/d$ and $2\log(d)/d$ that are maximal in the sense that they are not properly contained in any other independent subsets. Moreover, it is often the case that there does not exist a subset $I' \subset I$ with density larger than $\log(d)/d + \epsilon$ which is contained in an independent subset with density larger than the density of I . So local perturbations cannot be used to increase the density of a given independent subset. To be precise, the paper [12] studies Erdős-Renyi style sparse graphs. However, the same ideas can be adapted to regular graphs.

Another feature established in [12] is that the space of independent sets with a fixed density in between $\log(d)/d$ and $2\log(d)/d$ “shatters” into exponentially

many clusters separated by macroscopic gulfs. A similar phenomenon is used in [19] to show that no “local” algorithm can produce independent subsets of $\mathbf{G}_{d,n}$ with density larger than $(1 + \frac{1}{\sqrt{2}}) \log(d)/d + \epsilon$. This was improved to $\log(d)/d + \epsilon$ in [32]. Shattering is used here to obtain an action of the free group with positive $b_{\Sigma,0}$.

In order to explain how to utilize these results to obtain Theorem 1.1, let $\mathbb{F}_r = \langle a_1, \dots, a_r \rangle$ be the free group of rank $r \geq 2$. Given a homomorphism $\sigma : \mathbb{F}_r \rightarrow \text{sym}(n)$, let $G(\sigma)$ be the multi-graph with vertex set $[n]$ and edges $\{v, \sigma(a_i)v\}$ (over $v \in [n], 1 \leq i \leq r$). The *permutation model* is the random graph $G(\sigma_n)$ where σ_n is a uniformly random homomorphism from \mathbb{F}_r to $\text{sym}(n)$. By [20] the permutation model and the configuration model used in [18, 19] to study $\mathbf{G}_{2r,n}$ are contiguous. This allows results about $\mathbf{G}_{2r,n}$ to be transferred to $G(\sigma_n)$.

A result of Bollobás [6] implies that, with high probability, $\mathbf{G}_{2r,n}$ has few short cycles. Together with the contiguity theorem, this shows the existence of a sofic approximation $\Sigma = \{\sigma_n\}_{n=1}^{\infty}$ to \mathbb{F}_r such that the deterministic graph $G(\sigma_n)$ and the random graph $G(\sigma_n)$ have (with high probability) approximately the same number of independent sets (of some fixed density). Moreover, the space of independent subsets of $G(\sigma_n)$ at a certain fixed density shatters.

An action of the free group is obtained using a non-constructive compactness argument whose proof is related to the proof of the Variational Principle in [26]. The end result is an invariant measure μ on the shift space $\{0, 1\}^{\mathbb{F}_r}$ such that a significant fraction of independent subsets at a certain fixed density of $G(\sigma_n)$ are good models for μ . From this, we conclude $b_{\Sigma,0}(\mu) > 0$.

1.3. A brief guide to the paper.

- §2 explains notational conventions.
- §3 reviews sofic entropy and fixes notation used throughout the paper.
- §4 defines the new homological invariants.
- §5–6 contain the proof that the new invariants are in fact invariant.
- §7 contains proofs that the new invariants trivialize when the group is amenable or the action is Bernoulli. Also in this section is a proof that if \mathbf{a} has the WPP with respect to Σ then $b_{\Sigma,0}(\mathbf{a}) = 0$.
- §8 proves Theorem 1.2.
- §9 is a list of open problems related to the new invariants and the Weak Pinsker Property.

Acknowledgments. The homological invariants introduced in this paper are inspired by [2]. The techniques for proving that they are measure-conjugacy invariants are simplified versions of techniques introduced in [1]. Brandon Seward suggested that it might be possible to use the shattering property to give a counterexample to the WPP. I would like to thank Dylan Airey, Tim Austin and Brandon Seward for many conversations related to this paper. Also thanks to IPAM, the Institute for Pure and Applied Mathematics and the UCLA mathematics dept. A

significant part of this paper was written while I was visiting during the Quantitative Linear Algebra Program in the Spring of 2018. Also thanks to Tim Austin and Chris Shriver for catching errors in previous versions.

2. Notation and conventions. In general, if A, B are sets then A^B denotes the set of all functions $x : B \rightarrow A$. If $x \in A^B$ and $b \in B$ then the notations $x(b)$ and x_b express the same element of A .

All maps and subsets are measurable unless explicitly stated otherwise. As a rule, all measure zero phenomena are ignored.

Given a topological space X , let $\text{Prob}(X)$ denote the space of all Borel probability measures on X endowed with the weak* topology. This is the smallest topology such that for every continuous compactly supported function f on X the map $\mu \mapsto \int f d\mu$ is continuous (for $\mu \in \text{Prob}(X)$). If X is compact then the Banach-Alaoglu Theorem implies $\text{Prob}(X)$ is compact. If $\Gamma \curvearrowright X$ is a continuous action by Γ then let $\text{Prob}_\Gamma(X)$ denote the subspace of Γ -invariant Borel probability measures. This is a closed subspace of $\text{Prob}(X)$.

We write $f(n) = o_n(1)$ to mean $\lim_{n \rightarrow \infty} f(n) = 0$. Similarly, $f(r) = o_r(\log^2(r)/r)$ means $\lim_{r \rightarrow \infty} f(r)(\log^2(r)/r)^{-1} = 0$.

3. A review of sofic entropy. We will use the symbolic approach to sofic entropy with notational conventions similar to Tim Austin's from [1, 2].

3.1. Sofic approximations. Suppose $\sigma : \Gamma \rightarrow \text{sym}(V)$ is a map where V is a finite set and $\text{sym}(V)$ is the group of permutations of V . It is not required that σ is a homomorphism. Let $D \Subset \Gamma$ be finite and $\delta > 0$. Then σ is

- (D, δ) -multiplicative if

$$\#\{v \in V : \sigma_i(gh)v = \sigma_i(g)\sigma_i(h)v \ \forall g, h \in D\} > (1 - \delta)|V|,$$

- (D, δ) -trace preserving if

$$\#\{v \in V : \sigma_i(f)v \neq v \ \forall f \in D \setminus \{1_\Gamma\}\} > (1 - \delta)|V|,$$

- (D, δ) -sofic if it is both (D, δ) -multiplicative and (D, δ) -trace preserving.

A *sofic approximation* to Γ consists of a sequence $\Sigma = \{\sigma_i\}_{i \in \mathbb{N}}$ of maps $\sigma_i : \Gamma \rightarrow \text{sym}(V_i)$ such that for all finite $D \subset \Gamma$, $\delta > 0$ and all but finitely many i , σ_i is (D, δ) -sofic. A group is *sofic* if it admits a sofic approximation.

3.2. Sofic entropy. Throughout, $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ denote compact metric spaces. Given a finite set V , let $d_{\mathcal{X}}^V$ be the normalized ℓ^1 -metric on \mathcal{X}^V defined by

$$d_{\mathcal{X}}^V(x, y) := |V|^{-1} \sum_{v \in V} d_{\mathcal{X}}(x_v, y_v).$$

For any finite set V , map $\sigma : \Gamma \rightarrow \text{sym}(V)$, $x \in \mathcal{X}^V$ and $v \in V$ the *pullback name of x with respect to (σ, v)* is the element $\Pi_v^\sigma(x) \in \mathcal{X}^\Gamma$ defined by

$$\Pi_v^\sigma(x)(g) := x(\sigma(g)^{-1}v).$$

For example, if σ is a homomorphism then $h\Pi_v^\sigma(x) = \Pi_{\sigma(h)v}^\sigma(x)$ so that the map $v \mapsto \Pi_v^\sigma(x)$ is Γ -equivariant.

The *empirical measure of x* is the probability measure P_x^σ on \mathcal{X}^Γ defined by

$$P_x^\sigma := |V|^{-1} \sum_{v \in V} \delta_{\Pi_v^\sigma(x)}.$$

For example, if σ is a homomorphism and $\sigma(\Gamma)$ acts transitively on V then $\{\Pi_v^\sigma(x) : v \in V\}$ is a single Γ -orbit in which case P_x^σ is the uniform measure on a finite Γ -orbit.

Given $\mathcal{O} \subset \text{Prob}(\mathcal{X}^\Gamma)$, an element $x \in \mathcal{X}^V$ is a (\mathcal{O}, σ) -*microstate* if $P_x^\sigma \in \mathcal{O}$. Let $\Omega(\mathcal{O}, \sigma)$ be the set of all (\mathcal{O}, σ) -microstates. The metric space $(\Omega(\mathcal{O}, \sigma), d_{\mathcal{X}}^V)$ is a *model space* for the action $\Gamma \curvearrowright (\mathcal{X}^\Gamma, \mu)$ for any $\mu \in \mathcal{O}$. A major idea introduced in [1, 2] is to derive measure-conjugacy invariants from the asymptotic geometric features of these model spaces.

Recall that a subset Y of a metric space (X, d_X) is ϵ -*covering* if X is the open ϵ -neighborhood of Y . Let $\text{cov}_\epsilon(X, d_X)$ denote the minimum cardinality of an ϵ -covering subset of X .

Let $\Sigma = \{\sigma_i\}_{i \in \mathbb{N}}$ be a sofic approximation to Γ . The Σ -*entropy of $\Gamma \curvearrowright (\mathcal{X}^\Gamma, \mu)$* is defined by

$$h_\Sigma(\mu) := \sup_{\epsilon > 0} \inf_{\mathcal{O} \ni \mu} \limsup_{i \rightarrow \infty} |V_i|^{-1} \log \text{cov}_\epsilon(\Omega(\mathcal{O}, \sigma_i), d_{\mathcal{X}}^{V_i}).$$

See [1] for a proof that this definition is equivalent to previous formulations of sofic entropy given in [9] or [26] for example. For general discussions or when Σ is left implicit, the Σ -entropy is called the *sofic entropy*.

The basic facts about sofic entropy are: it is a measure-conjugacy invariant, it agrees with classical entropy when Γ is amenable, it can depend on the choice of sofic approximation, it can increase under factor maps, the sofic entropy of a Bernoulli shift is the Shannon entropy of the base. See [10] for an introduction.

Remark 2. In the special case in which \mathcal{X} is finite, the definition above reduces to

$$h_\Sigma(\mu) := \inf_{\mathcal{O} \ni \mu} \limsup_{i \rightarrow \infty} |V_i|^{-1} \log \#\Omega(\mathcal{O}, \sigma_i).$$

4. Sofic homology.

4.1. Homology theory on the Hamming cube. Fix a finite set V and compact metric space $(\mathcal{X}, d_{\mathcal{X}})$. For an integer $d \geq 0$, let $C_d(\mathcal{X}^V)$ be the abelian group

generated by all symbols of the form $[x_0, \dots, x_d]$ (with $x_0, \dots, x_d \in \mathcal{X}^V$) subject to the relations:

$$[x_{\pi(0)}, \dots, x_{\pi(d)} v] = (-1)^{\text{sign}(\pi)} [x_0, x_1, \dots, x_d]$$

over all $\pi \in \text{sym}(d+1)$. An element of the form $[x_0, \dots, x_d]$ is an *oriented d-simplex* of \mathcal{X}^V and an element of $C_d(\mathcal{X}^V)$ is called a *d-chain*.

Let $\partial_d : C_d(\mathcal{X}^V) \rightarrow C_{d-1}(\mathcal{X}^V)$ denote the boundary map

$$\partial_d([x_0, \dots, x_d]) = \sum_{i=0}^d (-1)^i [x_0, \dots, \hat{x}_i, \dots, x_d]$$

where \hat{x}_i indicates that x_i is omitted.

There is not much interesting that we can say about the homology of the Hamming cube \mathcal{X}^V . Instead we will focus on special subgroups of $C_d(\mathcal{X}^V)$ defined in terms of sofic approximation data as explained next.

4.2. Special subgroups defined by a sofic approximation. Let $\sigma : \Gamma \rightarrow \text{sym}(V)$ be a map. Given an open subset $\mathcal{O} \subset \text{Prob}(\mathcal{X}^\Gamma)$ and $\kappa > 0$, let $C_d(\mathcal{O}, \kappa, \sigma)$ be the subgroup of $C_d(\mathcal{X}^V)$ generated by all chains of the form $[x_0, \dots, x_d]$ such that each x_i is a (\mathcal{O}, σ) -microstate (that is $P_{x_i}^\sigma \in \mathcal{O}$) and $d_{\mathcal{X}}^V(x_i, x_j) < \kappa$ for all i, j . Let

$$\begin{aligned} Z_d(\mathcal{O}, \kappa, \sigma) &= \ker(\partial_d) \cap C_d(\mathcal{O}, \kappa, \sigma) \\ B_d(\mathcal{O}, \kappa, \sigma) &= \partial_{d+1}(C_{d+1}(\mathcal{O}, \kappa, \sigma)) \end{aligned}$$

be the $(\mathcal{O}, \kappa, \sigma)$ -cycles and boundaries respectively.

The *length* of a *d-chain* $z \in C_d(\mathcal{X}^V)$ is the smallest number of oriented simplices needed to represent z . So if

$$z = \sum_{i=1}^k c_i s_i$$

where $c_i \in \mathbb{Z}$ are coefficients and $s_i = [x_0^i, \dots, x_d^i]$ is an oriented simplex then the length of z is at most k . For $L > 0$, let $Z_d^L(\mathcal{O}, \kappa, \sigma)$ be the subgroup of $Z_d(\mathcal{O}, \kappa, \sigma)$ generated by $(\mathcal{O}, \kappa, \sigma)$ -cycles of length $\leq L$. To be precise, $z \in Z_d^L(\mathcal{O}, \kappa, \sigma)$ if it is possible to write $z = \sum_{i=1}^k c_i z_i$ for some coefficients $c_i \in \mathbb{Z}$ and cycles $z_i \in Z_d(\mathcal{O}, \kappa, \sigma)$ such that each z_i has length $\leq L$.

Given nested open subsets $\mathcal{O}_1 \subset \mathcal{O}_2 \subset \text{Prob}(\mathcal{X}^\Gamma)$, constants $0 < \kappa_1 \leq \kappa_2$ and $L > 0$, define the homology group

$$H_d^L(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma) := \frac{Z_d^L(\mathcal{O}_1, \kappa_1, \sigma)}{Z_d^L(\mathcal{O}_1, \kappa_1, \sigma) \cap B_d(\mathcal{O}_2, \kappa_2, \sigma)}.$$

4.3. Main results. Before stating the main theorem, we mention the following corollary which gives the flavor of the main result without as many quantifiers.

Definition 1. A group H is a *QS-group* of a group G if H is isomorphic to a quotient of a subgroup of G . Let Abel denote the class of abelian groups. A function $F : \text{Abel} \rightarrow \mathbb{R}$ is *monotone* if whenever H is an QS-group of G , $F(H) \leq F(G)$.

The next corollary follows immediately from Theorem 4.3 which is stated below.

COROLLARY 4.1. *Let $F = \{F_i\}_{i \in \mathbb{N}}$ be a sequence of monotone functions $F_i : \text{Abel} \rightarrow \mathbb{R}$. Given an invariant measure $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$, a sofic approximation Σ and $L \in [1, \infty]$, define*

$$F_{d,\Sigma}(\mu) := \sup_{\mathcal{O}_2 \ni \mu} \sup_{\kappa_2 > 0} \inf_{\mu \in \mathcal{O}_1 \subset \mathcal{O}_2} \inf_{\kappa_1 > 0} \sup_{0 < L < \infty} \limsup_{i \rightarrow \infty} F_i(H_d^L(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_i)).$$

If \mathcal{X} is totally disconnected then $F_{d,\Sigma}$ is a measure-conjugacy invariant. In other words, if $(\mathcal{Y}, d_\mathcal{Y})$ is another totally disconnected compact metric space, $\nu \in \text{Prob}_\Gamma(\mathcal{Y}^\Gamma)$ and the actions $\Gamma \curvearrowright (\mathcal{X}^\Gamma, \mu), \Gamma \curvearrowright (\mathcal{Y}^\Gamma, \nu)$ are measurably conjugate then $F_{d,\Sigma}(\mu) = F_{d,\Sigma}(\nu)$.

The next result follows by setting $F_i(G) := |V_i|^{-1} \log \dim_{\mathbb{Q}}(G \otimes_{\mathbb{Z}} \mathbb{Q})$ in the previous corollary.

COROLLARY 4.2. *Given $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$, define the d -th betti number of μ with respect to Σ by*

$$\begin{aligned} b_{d,\Sigma}(\mu) := & \sup_{\mathcal{O}_2 \ni \mu} \sup_{\kappa_2 > 0} \inf_{\mu \in \mathcal{O}_1 \subset \mathcal{O}_2} \inf_{\kappa_1 > 0} \sup_{0 < L < \infty} \limsup_{i \rightarrow \infty} |V_i|^{-1} \\ & \times \log \dim_{\mathbb{Q}}(H_d^L(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_i) \otimes_{\mathbb{Z}} \mathbb{Q}). \end{aligned}$$

If \mathcal{X} is totally disconnected then $b_{d,\Sigma}(\mu)$ is a measure-conjugacy invariant.

The main definition is:

Definition 2. Let $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$, $\nu \in \text{Prob}_\Gamma(\mathcal{Y}^\Gamma)$, $L, d \geq 0$. Then the d -dimensional sofic homology of ν is less than or equal to the d -dimensional sofic homology of μ if for every open neighborhood $\mathcal{O}_{2,\nu} \ni \nu$, every $\kappa_{2,\nu} > 0$ there exist an open neighborhood $\mathcal{O}_{2,\mu} \ni \mu$ and $\kappa_{2,\mu} > 0$ such that for every open $\mathcal{O}_{1,\mu}$ with $\mu \in \mathcal{O}_{1,\mu} \subset \mathcal{O}_{2,\mu}$ and every $\kappa_{1,\mu}$ with $0 < \kappa_{1,\mu} \leq \kappa_{2,\mu}$ there exist an open neighborhood $\mathcal{O}_{1,\nu}$ with $\nu \in \mathcal{O}_{1,\nu} \subset \mathcal{O}_{2,\nu}$ and $\kappa_{1,\nu}$ with $0 < \kappa_{1,\nu} \leq \kappa_{2,\nu}$ such that for every $0 < L < \infty$ and all but finitely many n , $H_d^L(\mathcal{O}_{1,\nu}, \mathcal{O}_{2,\nu}, \kappa_{1,\nu}, \kappa_{2,\nu}, \sigma_n)$ is a QS-group of $H_d^L(\mathcal{O}_{1,\mu}, \mathcal{O}_{2,\mu}, \kappa_{1,\mu}, \kappa_{2,\mu}, \sigma_n)$. The d -dimensional sofic homology theories of μ and ν are equivalent if the d -dimensional sofic homology of μ is less than or equal to the d -dimensional sofic homology of ν and vice versa.

The main theorem is:

THEOREM 4.3. *The homology groups defined above yield a measure-conjugacy invariant as follows. Suppose \mathcal{X}, \mathcal{Y} are totally disconnected compact metric spaces, $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$, $\nu \in \text{Prob}_\Gamma(\mathcal{Y}^\Gamma)$ and $\Gamma \curvearrowright (\mathcal{X}^\Gamma, \mu)$ is measurably conjugate to $\Gamma \curvearrowright (\mathcal{Y}^\Gamma, \nu)$. Then μ and ν have equivalent d -dimensional sofic homology theories with respect to every approximation Σ and for every dimension d .*

Remark 3. All of the definitions could be changed by setting $L = \infty$ throughout. The analog of Theorem 4.3 still holds under this change with essentially the same proof. However, we do not know how to compute this homology except in degenerate cases.

5. Preliminaries to the proof of Theorem 4.3.

5.1. Almost Lipschitz maps.

Definition 3. Let (X, d_X) and (Y, d_Y) be metric spaces, let $\epsilon > 0$, and let $L < \infty$. A map $\phi : X \rightarrow Y$ is ϵ -almost L -Lipschitz if

$$d_Y(\phi(x), \phi(x')) \leq \epsilon + Ld_X(x, x') \quad \forall x, x' \in X.$$

A map is ϵ -almost Lipschitz if it is so for some L .

LEMMA 5.1. *A uniformly continuous map from a bounded metric space to another bounded metric space is η -almost Lipschitz for every $\eta > 0$.*

Proof. Let $\phi : \mathcal{X} \rightarrow \mathcal{Y}$ be a uniformly continuous map from a bounded space $(\mathcal{X}, d_{\mathcal{X}})$ to a bounded metric space $(\mathcal{Y}, d_{\mathcal{Y}})$ and let $\eta > 0$. Let $\epsilon > 0$ be small enough so that if $d_{\mathcal{X}}(x, y) < \epsilon$ then $d_{\mathcal{Y}}(\phi x, \phi y) < \eta$.

Now let $x, y \in X$ be arbitrary. If $d_{\mathcal{X}}(x, y) \geq \epsilon$ then

$$d_{\mathcal{Y}}(\phi x, \phi y) \leq \text{diam}(\mathcal{Y}, d_{\mathcal{Y}}) \leq \eta + \frac{\text{diam}(\mathcal{Y}, d_{\mathcal{Y}})}{\epsilon} d_{\mathcal{X}}(x, y).$$

So ϕ is η -almost $\frac{\text{diam}(\mathcal{Y}, d_{\mathcal{Y}})}{\epsilon}$ -Lipschitz. □

5.2. Equivariant maps and their approximations.

Notation 1. If $x \in \mathcal{X}^\Gamma$ and $g \in \Gamma$, then $S^g x = gx \in \mathcal{X}^\Gamma$ is defined by $S^g x(f) = x(g^{-1} f)$. We will also write $S^g x$ if $x \in \mathcal{Y}^\Gamma$. So S^g is the shift by g .

Definition 4. A map $\Phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}^\Gamma$ is *equivariant* if $\Phi(gx) = g\Phi(x)$ for a.e. $x \in \mathcal{X}^\Gamma$ and every $g \in \Gamma$. Given a map $\psi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ we define an equivariant map $\psi^\Gamma : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}^\Gamma$ by $\psi^\Gamma(x)(h) = \psi(S^{h^{-1}} x)$. For example, if $\Phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}^\Gamma$ is equivariant and $\phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ is defined by $\phi(x) = \Phi(x)(1_\Gamma)$ then $\Phi = \phi^\Gamma$.

Definition 5. For a subset $D \subset \Gamma$, let $\text{Res}^D : \mathcal{X}^\Gamma \rightarrow \mathcal{X}^D$ denote the restriction map. If $\phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ and $D \subseteq \Gamma$ is finite, then ϕ is D -local if it is measurable with respect to Res^D . A function is local if it is D -local for some D .

Definition 6. As above, we let $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ be bounded Polish spaces. Also let $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$, $\phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ be a measurable function, and $\eta > 0$. An η -uniformly continuous (or η -UC) approximation to ϕ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ is a measurable map $\tilde{\phi} : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ with the following properties.

- (i) The map $\tilde{\phi}$ approximates ϕ in the sense that

$$(1) \quad \int d_{\mathcal{Y}}(\tilde{\phi}(x), \phi(x)) d\mu(x) < \eta.$$

- (ii) There is a finite $D \subseteq \Gamma$ such that $\tilde{\phi}$ is D -local.

(iii) Regarded as a map from \mathcal{X}^D to \mathcal{Y} , $\tilde{\phi}$ is uniformly continuous with respect to $d_{\mathcal{X}}^D$ and $d_{\mathcal{Y}}$ (where $d_{\mathcal{X}}^D$ is the normalized ℓ^1 -metric on \mathcal{X}^D as defined in the beginning of §3.2).

LEMMA 5.2. *Suppose that $\mathcal{X}, \mathcal{Y}, \mu, \phi$ are as in Definition 6, \mathcal{X} is totally disconnected and both \mathcal{X}, \mathcal{Y} are compact. Then there exist η -UC approximations to ϕ for all $\eta > 0$.*

Proof. After rescaling if necessary, we may assume that the diameter of \mathcal{Y} is bounded by 1. Because \mathcal{Y} is compact, there exists a finite open cover $\mathcal{O} = \{O_1, \dots, O_n\}$ of \mathcal{Y} by sets of diameter $< \eta/3$.

A subset $X \subset \mathcal{X}^\Gamma$ is D -local if its characteristic function $1_X : \mathcal{X}^\Gamma \rightarrow \mathbb{R}$ is D -local. Because \mathcal{X} is totally disconnected, for every $1 \leq i \leq n$, there exist a finite subset $D_i \subset \Gamma$ and a D_i -local clopen subset $\tilde{C}_i \subset \mathcal{X}^\Gamma$ such that

$$\mu(\tilde{C}_i \Delta \phi^{-1}(O_i)) < \frac{\eta}{3n}.$$

For $1 \leq i \leq n$, let

$$C_i := \tilde{C}_i \setminus \bigcup_{j=i+1}^n \tilde{C}_j.$$

Also let $C_0 = \mathcal{X}^\Gamma \setminus \bigcup_{i=1}^n C_i$. Then $\{C_i\}_{i=0}^n$ is a clopen partition of \mathcal{X}^Γ . Setting $D = \cup_i D_i$, we see that C_i is D -local for every $0 \leq i \leq n$.

Choose a point $p_i \in O_i$ for all $1 \leq i \leq n$ and also let $p_0 \in \mathcal{Y}$ be an arbitrary point. Define $\tilde{\phi} : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ by $\tilde{\phi}(x) = p_i$ if $x \in C_i$. By construction, $\tilde{\phi}$ is D -local. It is uniformly continuous because it is continuous and \mathcal{X}^Γ is compact. To finish the proof, it suffices to estimate the error in the approximation to ϕ :

$$\int d_{\mathcal{Y}}(\phi(x), \tilde{\phi}(x)) d\mu(x) = \sum_{i=0}^n \int_{C_i} d_{\mathcal{Y}}(\phi(x), \tilde{\phi}(x)) d\mu(x).$$

Since $C_0 = \mathcal{X}^\Gamma \setminus \cup_{i=1}^n \tilde{C}_i$,

$$\mu(C_0) \leq \sum_{i=1}^n \mu(\tilde{C}_i \Delta \phi^{-1}(O_i)) \leq \eta/3.$$

Since the diameter of $(\mathcal{Y}, d_{\mathcal{Y}})$ is bounded by 1, $\int_{C_0} d_{\mathcal{Y}}(\phi(x), \tilde{\phi}(x)) d\mu(x) \leq \eta/3$.

For any $1 \leq i \leq n$, $C_i \subset \tilde{C}_i$. Therefore, $C_i \subset \phi^{-1}(O_i) \cup (\tilde{C}_i \Delta \phi^{-1}(O_i))$. If $x \in C_i \cap \phi^{-1}(O_i)$ then $d_{\mathcal{Y}}(\phi x, \tilde{\phi} x) < \eta/3$ since O_i has diameter $< \eta/3$. So

$$\begin{aligned} \int_{C_i} d_{\mathcal{Y}}(\phi(x), \tilde{\phi}(x)) d\mu(x) &= \int_{C_i \cap \phi^{-1}(O_i)} d_{\mathcal{Y}}(\phi(x), \tilde{\phi}(x)) d\mu(x) \\ &\quad + \int_{C_i \setminus \phi^{-1}(O_i)} d_{\mathcal{Y}}(\phi(x), \tilde{\phi}(x)) d\mu(x) \\ &\leq \mu(C_i)\eta/23 + \mu(\tilde{C}_i \Delta \phi^{-1}(O_i)) \leq \mu(C_i)\eta/3 + \frac{\eta}{3n}. \end{aligned}$$

Since $\sum_{i=1}^n \mu(C_i)\eta/3 + \frac{\eta}{3n} \leq 2\eta/3$,

$$\int d_{\mathcal{Y}}(\phi(x), \tilde{\phi}(x)) d\mu(x) \leq \eta.$$

Since η is arbitrary, this implies the lemma. \square

Definition 7. Let $F \subset \Gamma$ be finite and $\phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$. Then $\phi^F : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}^F$ is defined by $\phi^F = \text{Res}^F \circ \phi^\Gamma$. So for any $f \in F$,

$$\phi^F(x)(f) = \phi^\Gamma(x)(f) = \phi(S^{f^{-1}}x).$$

LEMMA 5.3. Suppose that $\mathcal{X}, \mathcal{Y}, \mu, \phi$ are as in Definition 6 and $(\mathcal{Y}, d_{\mathcal{Y}})$ has diameter at most 1. If $\tilde{\phi}$ is an η -UC approximation to ϕ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ for some $\eta \in (0, 1)$, then $\tilde{\phi}^F$ is an η -UC approximation to $\phi^F : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}^F$ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}}^F)$ for every finite $F \subseteq \Gamma$.

Proof. This lemma is similar to [1, Lemma 4.4] but it is easier since we work with UC maps.

Firstly, the shift-invariance of μ and inequality (1) imply that

$$(2) \quad \int d_{\mathcal{Y}}(\phi^F(x), \tilde{\phi}^F(x)) d\mu(x) = \frac{1}{|F|} \sum_{g \in F^{-1}} \int d_{\mathcal{Y}}(\phi(gx), \tilde{\phi}(gx)) d\mu(x) < \eta.$$

Let $\tilde{\phi}$ be D -local for some finite $D \subset \Gamma$. Then $\tilde{\phi}^F$ is FD -local since for any $f \in F$, $\tilde{\phi}^F(x)_f = \tilde{\phi}(S^{f^{-1}}x)$ depends only on the restriction of $S^{f^{-1}}x$ to D . However, for $d \in D$, $S^{f^{-1}}x(d) = x(fd)$. So $\tilde{\phi}^F(x)$ depends only on the restriction of x to FD .

Lastly, we claim $\tilde{\phi}^F$ is uniformly continuous as a map from \mathcal{X}^{FD} to \mathcal{Y}^F . To see this, let $\epsilon > 0$. Since $\tilde{\phi}$ is uniformly continuous as a map from \mathcal{X}^D to \mathcal{Y} , there is a

$\delta > 0$ such that if $x, y \in \mathcal{X}^D$ satisfy $d_{\mathcal{X}}^D(x, y) < \sqrt{\delta}|D|$ then $d_{\mathcal{Y}}(\tilde{\phi}(x), \tilde{\phi}(y)) < \epsilon/2$. By choosing δ smaller if necessary we may assume $\sqrt{\delta} < \epsilon/2$.

For every $g \in FD$ the number of pairs $(f, d) \in F \times D$ such that $fd = g$ is at most $|D|$. Therefore,

$$\begin{aligned} d_{\mathcal{X}}^{FD}(x, y) &= |FD|^{-1} \sum_{g \in FD} d_{\mathcal{X}}(x_g, y_g) \geq |FD|^{-1}|D|^{-1} \sum_{f \in F} \sum_{d \in D} d_{\mathcal{X}}(x_{fd}, y_{fd}) \\ &= |FD|^{-1} \sum_{f \in F} d_{\mathcal{X}}^D(S^{f^{-1}}x, S^{f^{-1}}y). \end{aligned}$$

Suppose $x, y \in \mathcal{X}^{FD}$ satisfy $d_{\mathcal{X}}^{FD}(x, y) < \delta$. By the previous inequality,

$$|F|^{-1} \sum_{f \in F} d_{\mathcal{X}}^D(S^{f^{-1}}x, S^{f^{-1}}y) \leq \frac{|FD|}{|F|} d_{\mathcal{X}}^{FD}(x, y) < \frac{|FD|}{|F|} \delta \leq |D| \delta.$$

By Markov's inequality, there exists a subset $F' \subset F$ such that $|F'| \geq (1 - \sqrt{\delta})|F|$ and $d_{\mathcal{X}}^D(S^{g^{-1}}x, S^{g^{-1}}y) < \sqrt{\delta}|D|$ for all $g \in F'$. By choice of δ , if $f \in F'$ then $d_{\mathcal{Y}}(\tilde{\phi}(S^{f^{-1}}x), \tilde{\phi}(S^{f^{-1}}y)) \leq \epsilon/2$. Because the diameter of $(\mathcal{Y}, d_{\mathcal{Y}})$ is at most 1,

$$\begin{aligned} d_{\mathcal{Y}}^F(\tilde{\phi}^F x, \tilde{\phi}^F y) &= |F|^{-1} \sum_{g \in F^{-1}} d_{\mathcal{Y}}(\tilde{\phi}(S^g x), \tilde{\phi}(S^g y)) \leq \frac{(\epsilon/2)|F'| + |F \setminus F'|}{|F|} \\ &\leq \epsilon/2 + \sqrt{\delta} \leq \epsilon. \end{aligned}$$

This shows $\tilde{\phi}^F$ is uniformly continuous as a map from $(\mathcal{X}^{FD}, d_{\mathcal{X}}^{FD})$ to $(\mathcal{Y}^F, d_{\mathcal{Y}}^F)$. \square

LEMMA 5.4. *Suppose that $\mathcal{X}, \mathcal{Y}, \mu, \phi$ are as in Definition 6 and $(\mathcal{Y}, d_{\mathcal{Y}})$ has diameter at most 1. Let $\nu = \phi_* \mu \in \text{Prob}_{\Gamma}(\mathcal{Y}^{\Gamma})$. Suppose $(\mathcal{Z}, d_{\mathcal{Z}})$ is also a bounded Polish space with diameter 1 and $\psi : \mathcal{Y}^{\Gamma} \rightarrow \mathcal{Z}$ is measurable. Let $\tilde{\psi}$ be an η_{ψ} -UC approximation to ψ and $\tilde{\phi}$ an η_{ϕ} -UC approximation to ϕ . Then $\tilde{\psi} \circ \tilde{\phi}^{\Gamma}$ is an η -UC-approximation to $\psi \circ \phi^{\Gamma}$ where $\eta = \eta(\tilde{\phi}, \tilde{\psi})$ tends to $2\sqrt{\eta_{\psi}} + \eta_{\phi}$ as η_{ϕ} tends to zero with $\tilde{\psi}$ fixed.*

Proof. By definition there exists a finite subset $D_{\psi} \subset \Gamma$ such that $\tilde{\psi}$ is D_{ψ} -local and $\tilde{\psi}$ regarded as a map from $\mathcal{Y}^{D_{\psi}}$ to \mathcal{Z} is uniformly continuous. Moreover,

$$\int d_{\mathcal{Z}}(\psi y, \tilde{\psi} y) d\nu(y) < \eta_{\psi}.$$

As mentioned in Lemma 5.1, because $\tilde{\psi}$ is uniformly continuous, it is η_{ψ} -almost L_{ψ} -Lipschitz for some constant L_{ψ} .

Suppose $\tilde{\phi}$ is an η_{ϕ} -UC approximation to ϕ . By Lemma 5.3, $\tilde{\phi}^{D_{\psi}}$ is a η_{ϕ} -UC approximation to $\phi^{D_{\psi}}$. So there exists a finite subset $D_{\phi} \subset \Gamma$ such that $\tilde{\phi}^{D_{\psi}}$ is

D_ϕ -local, $\tilde{\phi}^{D_\psi}$ regarded as a map from \mathcal{X}^{D_ϕ} to \mathcal{Y}^{D_ψ} is uniformly continuous and

$$\int d_{\mathcal{Y}}^{D_\psi}(\phi^{D_\psi}x, \tilde{\phi}^{D_\psi}x) d\mu(x) < \eta_\phi.$$

It is immediate that $\tilde{\psi} \circ \tilde{\phi}^\Gamma$ is $D_\phi D_\psi$ -local and when regarded as a map from $\mathcal{X}^{D_\phi D_\psi}$ to \mathcal{Z} , it is uniformly continuous. Let

$$\begin{aligned} G_\psi &= \{y \in \mathcal{Y}^\Gamma : d_{\mathcal{Z}}(\tilde{\psi}y, \psi y) < \sqrt{\eta_\psi}\} \\ G_\phi &= \{x \in \mathcal{X}^\Gamma : d_{\mathcal{Y}}^{D_\psi}(\tilde{\phi}^{D_\psi}x, \phi^{D_\psi}x) < \sqrt{\eta_\phi}\}. \end{aligned}$$

Suppose $x \in G_\phi \cap (\phi^\Gamma)^{-1}(G_\psi)$.

$$\begin{aligned} d_{\mathcal{Z}}(\tilde{\psi}\tilde{\phi}^\Gamma x, \psi\phi^\Gamma x) &\leq d_{\mathcal{Z}}(\tilde{\psi}\tilde{\phi}^\Gamma x, \tilde{\psi}\phi^\Gamma x) + d_{\mathcal{Z}}(\tilde{\psi}\phi^\Gamma x, \psi\phi^\Gamma x) \\ &\leq (\eta_\psi + L_\psi\sqrt{\eta_\phi}) + \sqrt{\eta_\psi}. \end{aligned}$$

The first term above occurs because $\tilde{\psi}$ is η_ψ -almost L_ψ -Lipschitz as a map from \mathcal{Y}^{D_ψ} to \mathcal{Z} and $d_{\mathcal{Y}}^{D_\psi}(\tilde{\phi}^{D_\psi}x, \phi^{D_\psi}x) < \sqrt{\eta_\phi}$. The second term occurs because $\phi^\Gamma(x) \in G_\psi$.

It follows that

$$\begin{aligned} \int d_{\mathcal{Z}}(\tilde{\psi}\tilde{\phi}^\Gamma x, \psi\phi^\Gamma x) d\mu(x) &\leq (1 - \mu(G_\phi \cap (\phi^\Gamma)^{-1}(G_\psi))) \text{diam}(\mathcal{Z}) \\ &\quad + \eta_\psi + L_\psi\sqrt{\eta_\phi} + \sqrt{\eta_\psi}. \end{aligned}$$

By Markov's inequality, $\mu(G_\phi) > 1 - \sqrt{\eta_\phi}$ and $\nu(G_\psi) > 1 - \sqrt{\eta_\psi}$. Because $\phi_*\mu = \nu$ it follows that

$$1 - \mu(G_\phi \cap (\phi^\Gamma)^{-1}(G_\psi)) < \sqrt{\eta_\phi} + \sqrt{\eta_\psi}.$$

Since $\text{diam}(\mathcal{Z}) = 1$,

$$\int d_{\mathcal{Z}}(\tilde{\psi}\tilde{\phi}^\Gamma x, \psi\phi^\Gamma x) d\mu(x) \leq (L_\psi + 1)\sqrt{\eta_\phi} + 2\sqrt{\eta_\psi} + \eta_\psi. \quad \square$$

5.3. Sofic models. Recall that the pullback name of $x \in \mathcal{X}^V$ with respect to $\sigma : \Gamma \rightarrow \text{sym}(V)$ and $v \in V$ is

$$\Pi_v^\sigma(x)(g) := x(\sigma(g)^{-1}v).$$

Given a map $\sigma : \Gamma \rightarrow \text{sym}(V)$, where V is a finite set and a map $\phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$, define $\phi^\sigma : \mathcal{X}^V \rightarrow \mathcal{Y}^V$ by

$$\phi^\sigma(x)_v = \phi(\Pi_v^\sigma(x)).$$

LEMMA 5.5. Suppose $\phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ is D -local for some finite set $D \subset \Gamma$ and regarded as a map from $\mathcal{X}^D \rightarrow \mathcal{Y}$ is η -almost L -Lipschitz. Then ϕ^σ is η -almost L -Lipschitz regarded as map from $(\mathcal{X}^V, d_\mathcal{X}^V)$ to $(\mathcal{Y}^V, d_\mathcal{Y}^V)$.

Proof. Let $x, y \in \mathcal{X}^V$. Then

$$\begin{aligned} d_\mathcal{Y}^V(\phi^\sigma x, \phi^\sigma y) &= |V|^{-1} \sum_{v \in V} d_\mathcal{Y}((\phi^\sigma x)_v, (\phi^\sigma y)_v) \\ &\leq |V|^{-1} \sum_{v \in V} \eta + L d_\mathcal{X}^D(\text{Res}^D \Pi_v^\sigma(x), \text{Res}^D \Pi_v^\sigma(y)) \\ &= \eta + L|V|^{-1} \sum_{v \in V} |D|^{-1} \sum_{g \in D} d_\mathcal{X}(\Pi_v^\sigma(x)_g, \Pi_v^\sigma(y)_g) \\ &= \eta + L|V|^{-1} \sum_{v \in V} |D|^{-1} \sum_{g \in D} d_\mathcal{X}(x(\sigma(g^{-1})v), y(\sigma(g^{-1})v)) \\ &= \eta + L d_\mathcal{X}^V(x, y) \end{aligned}$$

where the last equality holds because for each $v \in V$ the number of pairs $(g, w) \in D \times V$ such that $\sigma(g^{-1})w = v$ equals $|D|$. Because x, y are arbitrary, this implies the lemma. \square

LEMMA 5.6. Suppose $\phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ is D_ϕ -local for some finite set $D_\phi \subset \Gamma$ and $\psi : \mathcal{Y}^\Gamma \rightarrow \mathcal{Z}$ is D_ψ -local for some finite set D_ψ . Then for all $x \in \mathcal{X}^V$,

$$\begin{aligned} &\{v \in V : (\psi\phi^\Gamma)^\sigma(x)_v \neq \psi^\sigma\phi^\sigma(x)_v\} \\ &\subset \{v \in V : \exists h \in D_\phi, g \in D_\psi, \sigma(h^{-1})\sigma(g^{-1})v \neq \sigma(h^{-1}g^{-1})v\}. \end{aligned}$$

In particular, if $1_\Gamma \in D_\phi \cap D_\psi$ and σ is a $(D_\phi^{-1}D_\psi^{-1}, \delta)$ -sofic approximation to Γ then

$$\#\{v \in V : (\psi\phi^\Gamma)^\sigma(x)_v \neq \psi^\sigma\phi^\sigma(x)_v\} \leq \delta|V|.$$

Proof. Fix $v \in V$. Suppose $\sigma(h^{-1})\sigma(g^{-1})v = \sigma(h^{-1}g^{-1})v$ for all $h \in D_\phi$ and $g \in D_\psi$. It suffices to show $(\psi\phi^\Gamma)^\sigma(x)_v = \psi^\sigma\phi^\sigma(x)_v$. Observe that

$$\begin{aligned} (\psi\phi^\Gamma)^\sigma(x)_v &= \psi(\phi^\Gamma(\Pi_v^\sigma x)) \\ \psi^\sigma\phi^\sigma(x)_v &= \psi(\Pi_v^\sigma(\phi^\sigma(x))). \end{aligned}$$

Since ψ is D_ψ -local, it suffices to show that for every $g \in D_\psi$,

$$\phi^\Gamma(\Pi_v^\sigma x)_g = \Pi_v^\sigma(\phi^\sigma(x))_g.$$

Observe that

$$\begin{aligned}\phi^\Gamma(\Pi_v^\sigma x)_g &= \phi(S^{g^{-1}}\Pi_v^\sigma(x)) \\ \Pi_v^\sigma(\phi^\sigma(x))_g &= \phi^\sigma(x)(\sigma(g^{-1})v) = \phi(\Pi_{\sigma(g^{-1})v}^\sigma(x)).\end{aligned}$$

Since ϕ is D_ϕ -local it suffices to show that for every $h \in D_\phi$,

$$S^{g^{-1}}\Pi_v^\sigma(x)_h = \Pi_{\sigma(g^{-1})v}^\sigma(x)_h.$$

The left-hand side simplifies as follows:

$$S^{g^{-1}}\Pi_v^\sigma(x)_h = \Pi_v^\sigma(x)_{gh} = x(\sigma((gh)^{-1})v) = x(\sigma(h^{-1}g^{-1})v).$$

The right-hand side simplifies to

$$\Pi_{\sigma(g^{-1})v}^\sigma(x)_h = x(\sigma(h^{-1})\sigma(g^{-1})v).$$

Therefore if $\sigma(h^{-1}g^{-1})v = \sigma(h^{-1})\sigma(g^{-1})v$ for every $h \in D_\phi$ and $g \in D_\psi$ then $(\psi\phi^\Gamma)^\sigma(x)_v = \psi^\sigma\phi^\sigma(x)_v$. \square

Definition 8. The *total variation distance* between two measures μ and ν on the same σ -algebra \mathcal{F} is

$$d_{\text{TV}}(\mu, \nu) = \sup_{A \in \mathcal{F}} |\mu(A) - \nu(A)|.$$

Roughly speaking, the next lemma shows that closeness in total variation distance of restricted measures implies closeness in the weak* topology.

LEMMA 5.7. *For any $\mu \in \text{Prob}(\mathcal{X}^\Gamma)$ and any weak* open set $O \subset \text{Prob}(\mathcal{X}^\Gamma)$ with $\mu \in O$, there exists a finite set $E \subset \Gamma$ and $\delta > 0$ such that*

$$\{\nu \in \text{Prob}(\mathcal{X}^\Gamma) : d_{\text{TV}}(\text{Res}_*^E \mu, \text{Res}_*^E \nu) < \delta\} \subset O.$$

Proof. By definition of the weak* topology, there are continuous functions f_1, \dots, f_k on \mathcal{X}^Γ and $\epsilon > 0$ such that

$$\left\{ \nu \in \text{Prob}(\mathcal{X}^\Gamma) : \left| \int f_i d\nu - \int f_i d\mu \right| < \epsilon \ \forall 1 \leq i \leq k \right\} \subset O.$$

Because each \mathcal{X}^Γ is the inverse limit of the compact spaces \mathcal{X}^D over finite $D \subset \Gamma$, there exist a finite $E \subset \Gamma$ and E -local continuous functions f'_1, \dots, f'_k on \mathcal{X}^Γ such that $|f_i - f'_i| < \epsilon/3$ for all i . By the triangle inequality,

$$\left\{ \nu \in \text{Prob}(\mathcal{X}^\Gamma) : \left| \int f'_i d\nu - \int f'_i d\mu \right| \leq \epsilon/3 \ \forall 1 \leq i \leq k \right\} \subset O.$$

Now suppose $\nu \in \text{Prob}(\mathcal{X}^\Gamma)$ and $d_{\text{TV}}(\text{Res}_*^E \mu, \text{Res}_*^E \nu) < \frac{\epsilon}{6M}$ where $M = \max_{1 \leq i \leq k} \|f'_i\|_{\text{sup}}$. By abuse of notation, we may consider each f'_i to be a continuous function on \mathcal{X}^E . So

$$\begin{aligned} \left| \int f'_i d\nu - \int f'_i d\mu \right| &= \left| \int f'_i d\text{Res}_*^E \nu - \int f'_i d\text{Res}_*^E \mu \right| \\ &\leq 2 \|f'_i\|_{\text{sup}} d_{\text{TV}}(\text{Res}_*^E \mu, \text{Res}_*^E \nu) \leq \epsilon/3. \end{aligned}$$

By the previous inclusion, this implies $\nu \in O$ and completes the lemma (with $\delta = \frac{\epsilon}{6M}$). \square

LEMMA 5.8. *Suppose $\phi: \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ is D -local for some finite set $D \subset \Gamma$. Also let $E \subset \Gamma$ be finite and $\delta > 0$. Then there exist a finite set $F \subset \Gamma$ and $\epsilon > 0$ (depending only on D, E, δ) such that if σ is an (F, ϵ) -sofic approximation to Γ , then for all $x \in \mathcal{X}^V$,*

$$\#\{v \in V: (\phi^\Gamma)(\Pi_v^\sigma(x))(g) \neq \Pi_v^\sigma(\phi^\sigma(x))(g) \ \forall g \in E\} \leq \delta|V|.$$

In particular, the total variation distance between the restricted empirical measures $\text{Res}_^E \phi_*^F P_x^\sigma$ and $\text{Res}_*^E P_{\phi^\sigma(x)}^\sigma$ is bounded by δ .*

Proof. Let $F \subset \Gamma$ and $\epsilon > 0$ be such that if σ is (F, ϵ) -sofic then

$$\#\{v \in V: \sigma(gh)^{-1}v = \sigma(h)^{-1}\sigma(g)^{-1}v \ \forall h \in D, \forall g \in E\} \geq (1 - \delta)\#V.$$

Suppose σ is (F, ϵ) -sofic and fix $v \in V$. By Definition 4,

$$(\phi^\Gamma)(\Pi_v^\sigma(x))(g) = \phi(S^{g^{-1}}\Pi_v^\sigma(x)).$$

By the definitions of pullback and ϕ^σ ,

$$\Pi_v^\sigma(\phi^\sigma(x))(g) = \phi^\sigma(x)(\sigma(g)^{-1}v) = \phi(\Pi_{\sigma(g)^{-1}v}^\sigma(x)).$$

Because ϕ is D -local, if

$$(S^{g^{-1}}\Pi_v^\sigma(x))(h) = \Pi_{\sigma(g)^{-1}v}^\sigma(x)(h)$$

for all $h \in D$ then $(\phi^\Gamma)(\Pi_v^\sigma(x))(g) = \Pi_v^\sigma(\phi^\sigma(x))(g)$.

We compute

$$(S^{g^{-1}}\Pi_v^\sigma(x))(h) = \Pi_v^\sigma(x)(gh) = x(\sigma(gh)^{-1}v)$$

and

$$\Pi_{\sigma(g)^{-1}v}^\sigma(x)(h) = x(\sigma(h)^{-1}\sigma(g)^{-1}v).$$

Because σ is (F, ϵ) -sofic, there is a $(1 - \delta)$ -fraction of vertices v such that $(S^{g^{-1}}\Pi_v^\sigma(x))(h) = \Pi_{\sigma(g)^{-1}v}^\sigma(x)(h)$ for all $g \in E, h \in D$. Again, since ϕ is D -local,

this condition implies $(\phi^\Gamma)(\Pi_v^\sigma(x))(g) = \Pi_v^\sigma(\phi^\sigma(x))(g)$ for all $g \in E$. This proves the first claim.

The last claim is implied by the first. In fact, $\text{Res}_*^\Gamma \phi_*^\Gamma P_x^\sigma$ is the law of $\text{Res}^E(\phi^\Gamma(\Pi_v^\sigma(x)))$ where $v \in V$ is chosen uniformly at random, while $\text{Res}_*^E P_{\phi^\sigma(x)}^\sigma$ is the law of $\text{Res}^E(\Pi_v^\sigma(\phi^\sigma(x)))$ where $v \in V$ is chosen uniformly at random. Since $\text{Res}^E(\phi^\Gamma(\Pi_v^\sigma(x))) = \text{Res}^E(\Pi_v^\sigma(\phi^\sigma(x)))$ for all but a δ -fraction of vertices v , the total variation distance between the restricted empirical measures $\text{Res}_*^\Gamma \phi_*^\Gamma P_x^\sigma$ and $\text{Res}_*^E P_{\phi^\sigma(x)}^\sigma$ is bounded by δ . \square

6. Proof of Theorem 4.3. We need one more lemma before the proof of the main theorem. Let $(\mathcal{Y}, d_{\mathcal{Y}})$ be a compact metric space. Given a map $\lambda : \mathcal{Y}^V \rightarrow \mathcal{Y}^V$ let $\lambda_* : C_d(\mathcal{Y}^V) \rightarrow C_d(\mathcal{Y}^V)$ be the corresponding homomorphism of chain groups. Note λ_* commutes with all boundary maps $\partial_d : C_d(\mathcal{Y}^V) \rightarrow C_{d-1}(\mathcal{Y}^V)$.

LEMMA 6.1. *Let $\lambda : \mathcal{Y}^V \rightarrow \mathcal{Y}^V$ be given and suppose there is a constant $\kappa' > 0$ such that $d_{\mathcal{Y}}^V(y, \lambda(y)) < \kappa'$ for all $y \in \mathcal{Y}^V$. If $z \in Z_d(\mathcal{O}_1, \kappa, \sigma)$ and $\lambda(\Omega(\mathcal{O}_1, \sigma)) \subset \Omega(\mathcal{O}_2, \sigma)$ (for some $\mathcal{O}_1, \mathcal{O}_2, \kappa, \sigma$) then*

$$z - \lambda_* z \in B_d(\mathcal{O}_1 \cup \mathcal{O}_2, \kappa + 2\kappa', \sigma).$$

Proof. It will be convenient to work with ordered simplices rather than ordinary simplices. An *ordered k -simplex* is an ordered $(k+1)$ -tuple (x_0, \dots, x_k) with $x_i \in \mathcal{Y}^V$. Given an ordered k -simplex (x_0, \dots, x_k) define $P(x_0, \dots, x_k) \in C_{k+1}(\mathcal{Y}^V)$ by

$$P(x_0, \dots, x_k) := \sum_{i=0}^k (-1)^i [x_0, \dots, x_i, \lambda(x_i), \dots, \lambda(x_k)].$$

Given an oriented simplex $[x_0, \dots, x_k] \in C_k(\mathcal{Y}^V)$ define $P([x_0, \dots, x_k]) \in C_{k+1}(\mathcal{Y}^V)$ by

$$P([x_0, \dots, x_k]) := |\text{sym}(k+1)|^{-1} \sum_{\pi \in \text{sym}(k+1)} \text{sign}(\pi) P(x_{\pi(0)}, \dots, x_{\pi(k)}).$$

We extend P linearly so that it is well defined as a homomorphism from $C_k(\mathcal{Y}^V)$ to $C_{k+1}(\mathcal{Y}^V)$ (for every k too).

Claim. $\partial P = \lambda_* - I + P\partial$ where I denotes the identity map.

Proof. It suffices to show that for any x_0, \dots, x_k ,

$$\partial P([x_0, \dots, x_k]) = [\lambda(x_0), \lambda(x_1), \dots, \lambda(x_k)] - [x_0, \dots, x_k] - P(\partial[x_0, \dots, x_k]).$$

The proof is by direct inspection of the coefficients. Details are in the proof of [21, Theorem 2.10]. \square

If $z \in Z_d(\mathcal{O}_1, \kappa, \sigma)$ then $Pz \in C_d(\mathcal{O}_1 \cup \mathcal{O}_2, \kappa + 2\kappa', \sigma)$ since $\lambda(\Omega(\mathcal{O}_1, \sigma)) \subset \Omega(\mathcal{O}_2, \sigma)$ and $d_{\mathcal{Y}}^V(y, \lambda(y)) < \kappa'$ for all y . The claim implies

$$\partial P(z) = \lambda_* z - z - P\partial z = \lambda_* z - z.$$

Therefore $\lambda_* z - z \in B_d(\mathcal{O}_1 \cup \mathcal{O}_2, \kappa + 2\kappa', \sigma)$. \square

Proof of Theorem 4.3. Suppose $\Phi : (\mathcal{X}^\Gamma, \mu) \rightarrow (\mathcal{Y}^\Gamma, \nu)$ is a measure-conjugacy. We may assume without loss of generality that $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ have diameter 1. Let $\mathcal{O}_{2,\nu}$ and $\kappa_{2,\nu}$ be given.

We need to choose $\mathcal{O}_{2,\mu}$ and $\kappa_{2,\mu}$. Before doing this, define $\phi : \mathcal{X}^\Gamma \rightarrow \mathcal{Y}$ by $\phi(x) := \Phi(x)_e$ (where $e \in \Gamma$ is the identity element). Note $\phi^\Gamma = \Phi$. We choose a UC-approximation $\tilde{\phi}$ to ϕ as follows. Choose $0 < \eta_\phi < 1$ small enough so that $3\sqrt{\eta_\phi} < \frac{\kappa_{2,\nu}}{8}$. By Lemma 5.2 there exists an η_ϕ -UC-approximation $\tilde{\phi}$ to ϕ . Because $\tilde{\phi}$ is continuous, there exists an open neighborhood $\mathcal{O}_{2,\mu}$ of μ such that the closure of $\tilde{\phi}_*^\Gamma \mathcal{O}_{2,\mu}$ is contained in $\mathcal{O}_{2,\nu}$. By definition of UC-approximation, there is a finite set $D_\phi \subset \Gamma$ such that $\tilde{\phi}$ is D_ϕ -local. By Lemma 5.1, $\tilde{\phi}$ is η_ϕ -almost L_ϕ -Lipschitz for some L_ϕ (when regarded as a map from \mathcal{X}^{D_ϕ} to \mathcal{Y}). Now choose $\kappa_{2,\mu} > 0$ so that $\eta_\phi + L_\phi \kappa_{2,\mu} < \frac{\kappa_{2,\nu}}{4}$.

Next we let $\mathcal{O}_{1,\mu}$ be an arbitrary open set with $\mu \in \mathcal{O}_{1,\mu} \subset \mathcal{O}_{2,\mu}$ and let $\kappa_{1,\mu}$ be an arbitrary constant with $0 < \kappa_{1,\mu} \leq \kappa_{2,\mu}$.

We need to choose $\mathcal{O}_{1,\nu}$ and $\kappa_{1,\nu}$. Before doing this, define $\psi : \mathcal{Y}^\Gamma \rightarrow \mathcal{X}$ by $\psi(y) := \Phi^{-1}(y)_e$. Note $\psi^\Gamma = \Phi^{-1}$. We choose a UC-approximation $\tilde{\psi}$ to ψ as follows. Choose $0 < \eta_\psi < 1$ so that $3\sqrt{\eta_\psi} < \frac{\kappa_{1,\mu}}{2}$. By Lemma 5.4 (and using $3\sqrt{\eta_\phi} < \frac{\kappa_{2,\nu}}{8}$) we can choose η_ψ smaller if necessary so that for any η_ψ -UC approximation $\tilde{\psi}$ to ψ , the composition $\tilde{\phi}\tilde{\psi}^\Gamma$ is a $\frac{\kappa_{2,\nu}}{8}$ -UC approximation to $\phi\psi^\Gamma$ which is the identity-coordinate projection. Fix such a UC-approximation $\tilde{\psi}$. Because $\tilde{\psi}$ and therefore $\tilde{\psi}^\Gamma$ are continuous, there is an open neighborhood $\mathcal{O}_{1,\nu}$ of ν such that $\mathcal{O}_{1,\nu} \subset \mathcal{O}_{2,\nu}$ and the closure of $\tilde{\psi}_*^\Gamma \mathcal{O}_{1,\nu}$ is contained in $\mathcal{O}_{1,\mu}$. By choosing $\mathcal{O}_{1,\nu}$ smaller if necessary we may assume that

$$\int d(\tilde{\phi}\tilde{\psi}^\Gamma z, z_e) d\nu'(z) < \kappa_{2,\nu}/8$$

for every $\nu' \in \mathcal{O}_{1,\nu}$. This is because the inequality holds if $\nu' = \nu$ (since $\tilde{\phi}\tilde{\psi}^\Gamma$ is a $\kappa_{2,\nu}/8$ -UC approximation to the identity coordinate-projection) and the map $z \mapsto d(\tilde{\phi}\tilde{\psi}^\Gamma z, z_e)$ is continuous.

By definition of UC-approximation, there is a finite set $D_\psi \subset \Gamma$ such that $\tilde{\psi}$ is D_ψ -local. By Lemma 5.1, $\tilde{\psi}$ is η_ψ -almost L_ψ -Lipschitz for some L_ψ (when regarded as a map from \mathcal{Y}^{D_ψ} to \mathcal{X}). Finally, choose $\kappa_{1,\nu} > 0$ so that $\eta_\psi + L_\psi \kappa_{1,\nu} < \kappa_{1,\mu}$ and $\kappa_{1,\nu} < \frac{\kappa_{2,\nu}}{2}$.

To simplify notation, let

$$\begin{aligned} Z(n, \nu) &:= Z_d^L(\mathcal{O}_{1,\nu}, \kappa_{1,\nu}, \sigma_n), & Z(n, \mu) &:= Z_d^L(\mathcal{O}_{1,\mu}, \kappa_{1,\mu}, \sigma_n) \\ B(n, \nu) &:= B_d(\mathcal{O}_{2,\nu}, \kappa_{2,\nu}, \sigma_n), & B(n, \mu) &:= B_d(\mathcal{O}_{2,\mu}, \kappa_{2,\mu}, \sigma_n) \\ H(n, \nu) &:= H_d^L(\mathcal{O}_{1,\nu}, \mathcal{O}_{2,\nu}, \kappa_{1,\nu}, \kappa_{2,\nu}, \sigma_n), \\ H(n, \mu) &:= H_d^L(\mathcal{O}_{1,\mu}, \mathcal{O}_{2,\mu}, \kappa_{1,\mu}, \kappa_{2,\mu}, \sigma_n). \end{aligned}$$

Also let

$$q_{n,\nu} : Z(n, \nu) \longrightarrow H(n, \nu), \quad q_{n,\mu} : Z(n, \mu) \longrightarrow H(n, \mu)$$

be the quotient maps.

Next we must verify that for all but finitely many n , $H(n, \nu)$ is a QS-group of $H(n, \mu)$. Our strategy is as follows. First we show that $\tilde{\psi}_*^{\sigma_n}(Z(n, \nu)) \subset Z(n, \mu)$. Let $S = q_{n,\mu} \circ \tilde{\psi}_*^{\sigma_n}(Z(n, \nu)) \leq H(n, \mu)$. Next we show that $\tilde{\phi}^{\sigma_n}$ induces a map, denoted by $\tilde{\phi}^H$, from S back to $H(n, \nu)$ that is surjective. This shows that $H(n, \nu)$ is a quotient of the subgroup $S \leq H(n, \mu)$ and thereby completes the proof.

It is convenient to first show that $\tilde{\psi}^{\sigma_n}$ and $\tilde{\phi}^{\sigma_n}$ behave well with respect to the Hamming metrics and empirical distributions.

Claim 1. For all but finitely many $n \in \mathbb{N}$ the following holds. For any $x, x' \in \Omega(\mathcal{O}_{2,\mu}, \sigma_n)$ and $y, y' \in \Omega(\mathcal{O}_{1,\nu}, \sigma_n)$,

- (1) $\tilde{\phi}^{\sigma_n}(x) \in \Omega(\mathcal{O}_{2,\nu}, \sigma_n)$,
- (2) if $d_{\mathcal{X}}^{V_n}(x, x') < \kappa_{2,\mu}$ then $d_{\mathcal{Y}}^{V_n}(\tilde{\phi}^{\sigma_n}(x), \tilde{\phi}^{\sigma_n}(x')) < \frac{\kappa_{2,\nu}}{4}$.
- (3) $\tilde{\psi}^{\sigma_n}(y) \in \Omega(\mathcal{O}_{1,\mu}, \sigma_n)$,
- (4) If $d_{\mathcal{Y}}^{V_n}(y, y') < \kappa_{1,\nu}$ then $d_{\mathcal{X}}^{V_n}(\tilde{\psi}^{\sigma_n}(y), \tilde{\psi}^{\sigma_n}(y')) < \kappa_{1,\mu}$.

Proof of Claim 1. We chose $\mathcal{O}_{2,\mu}$ so that the *closure* of $\tilde{\phi}_*^\Gamma \mathcal{O}_{2,\mu}$ is contained in $\mathcal{O}_{2,\nu}$. By compactness and Lemma 5.7, there exists a finite set $E \subset \Gamma$ and $\delta > 0$ such that for any $\alpha \in \tilde{\phi}_*^\Gamma \mathcal{O}_{2,\mu}$, if $\beta \in \text{Prob}(\mathcal{Y}^\Gamma)$ is such that the restricted measures $\text{Res}_*^E \alpha, \text{Res}_*^E \beta$ have total variation distance $< \delta$ then $\beta \in \mathcal{O}_{2,\nu}$.

Lemma 5.8 implies that if n is sufficiently large (independent of x) then the total variation distance of the restricted measures $\text{Res}_*^E P_{\tilde{\phi}^{\sigma_n}(x)}^{\sigma_n}$ and $\text{Res}_*^E \tilde{\phi}_*^\Gamma P_x^{\sigma_n} \in \mathcal{O}_{2,\nu}$ is bounded by $\delta > 0$. Since $P_x^{\sigma_n} \in \mathcal{O}_{2,\mu}$, $\tilde{\phi}_*^\Gamma P_x^{\sigma_n} \in \tilde{\phi}_*^\Gamma \mathcal{O}_{2,\mu}$. So this implies $P_{\tilde{\phi}^{\sigma_n}(x)}^{\sigma_n} \in \mathcal{O}_{2,\nu}$. This proves (1).

Since $d_{\mathcal{X}}^{V_n}(x, x') < \kappa_{2,\mu}$ and $\tilde{\phi}$ is η_ϕ -almost L_ϕ -Lipschitz, Lemma 5.5 implies

$$d_{\mathcal{Y}}^{V_n}(\tilde{\phi}^{\sigma_n}(x), \tilde{\phi}^{\sigma_n}(x')) \leq L_\phi d_{\mathcal{X}}^{V_n}(x, x') + \eta_\phi < L_\phi \kappa_{2,\mu} + \eta_\phi < \frac{\kappa_{2,\nu}}{4}.$$

This proves (2). The other statements are proven similarly. \square

Claim 2. For all but finitely many $n \in \mathbb{N}$, for every $z \in Z(n, \nu)$,

$$\tilde{\phi}_*^{\sigma_n} \tilde{\psi}_*^{\sigma_n}(z) - z \in B(n, \nu).$$

Proof of Claim 2. We will verify the hypotheses of Lemma 6.1 with $\lambda = \tilde{\phi}^{\sigma_n} \tilde{\psi}^{\sigma_n}$. So let $y \in \Omega(\mathcal{O}_{1,\nu}, \sigma_n)$.

By Claim 1(3), $\tilde{\psi}^{\sigma_n}(y) \in \Omega(\mathcal{O}_{1,\mu}, \sigma_n)$. Since $\mathcal{O}_{1,\mu} \subset \mathcal{O}_{2,\mu}$, this implies $\tilde{\psi}^{\sigma_n}(y) \in \Omega(\mathcal{O}_{2,\mu}, \sigma_n)$. So Claim 1(1) implies $\tilde{\phi}^{\sigma_n} \tilde{\psi}^{\sigma_n}(y) \in \Omega(\mathcal{O}_{2,\nu}, \sigma_n)$.

By Lemma 5.6 for all but finitely many n ,

$$\#\{v \in V_n : (\tilde{\phi} \tilde{\psi}^\Gamma)^{\sigma_n}(y)_v \neq \tilde{\phi}^{\sigma_n} \tilde{\psi}^{\sigma_n}(y)_v\} \leq (\kappa_{2,\nu}/8) |V_n|.$$

Because \mathcal{Y} has diameter 1,

$$d_{\mathcal{Y}}^{V_n}(\tilde{\phi}^{\sigma_n} \tilde{\psi}^{\sigma_n}(y), y) \leq (\kappa_{2,\nu}/8) + d_{\mathcal{Y}}^{V_n}((\tilde{\phi} \tilde{\psi}^\Gamma)^{\sigma_n}(y), y).$$

Note

$$d_{\mathcal{Y}}^{V_n}((\tilde{\phi} \tilde{\psi}^\Gamma)^{\sigma_n}(y), y) = \int d_{\mathcal{Y}}((\tilde{\phi} \tilde{\psi}^\Gamma)(z), z_e) dP_y^{\sigma_n}(z).$$

Since $P_y^{\sigma_n} \in \mathcal{O}_{1,\nu}$, the choice of $\mathcal{O}_{1,\nu}$ implies

$$\int d_{\mathcal{Y}}((\tilde{\phi} \tilde{\psi}^\Gamma)(z), z_e) dP_y^{\sigma_n}(z) < \kappa_{2,\nu}/8.$$

Thus,

$$d_{\mathcal{Y}}^{V_n}(\tilde{\phi}^{\sigma_n} \tilde{\psi}^{\sigma_n}(y), y) \leq \kappa_{2,\nu}/4$$

for every $y \in \Omega(\mathcal{O}_{1,\nu}, \sigma_n)$. Claim 2 now follows from Lemma 6.1 with $\kappa' = \kappa_{2,\nu}/4$ and $\kappa = \kappa_{1,\nu} < \kappa_{2,\nu}/2$. \square

Claim 3. For all but finitely many $n \in \mathbb{N}$,

$$(\tilde{\psi}_*^{\sigma_n})^{-1}(B(n, \mu)) \cap Z(n, \nu) \subset B(n, \nu).$$

Proof of Claim 3. Let $z \in Z(n, \nu)$ and suppose $\tilde{\psi}_*^{\sigma_n}(z) \in B(n, \mu)$. By Claim 1(1,2), $\tilde{\phi}_*^{\sigma_n} \tilde{\psi}_*^{\sigma_n}(z) \in B(n, \nu)$. By Claim 2, $z - \tilde{\phi}_*^{\sigma_n} \tilde{\psi}_*^{\sigma_n}(z) \in B(n, \nu)$. Therefore, $z \in B(n, \nu)$ as required. \square

Since $\tilde{\psi}^{\sigma_n}$ commutes with the boundary map ∂_d , Claim 1 (3,4) implies that $\tilde{\psi}_*^{\sigma_n}(Z(n, \nu)) \subset Z(n, \mu)$. Therefore $S := q_{n,\mu} \circ \tilde{\psi}_*^{\sigma_n}(Z(n, \nu))$ is a well-defined subgroup of $H(n, \mu)$. Define $\tilde{\phi}^H : S \rightarrow H(n, \nu)$ as follows. Given $w \in S$, let $z \in Z(n, \nu)$ be such that $q_{n,\mu} \circ \tilde{\psi}_*^{\sigma_n}(z) = w$. Then define $\tilde{\phi}^H(w) := z + B(n, \nu) \cap Z(n, \nu)$.

To see that $\tilde{\phi}^H$ is well defined, suppose that $z' \in Z(n, \nu)$ also satisfies $q_{n, \mu} \circ \tilde{\psi}_*^{\sigma_n}(z') = w$. Then

$$\tilde{\psi}_*^{\sigma_n}(z) - \tilde{\psi}_*^{\sigma_n}(z') = \tilde{\psi}_*^{\sigma_n}(z - z') \in B(n, \mu).$$

By Claim 3, this implies $z - z' \in B(n, \nu)$. This implies $\tilde{\phi}^H$ is well defined.

To check that $\tilde{\phi}^H$ is surjective, let $z \in Z(n, \nu)$. If $w = q_{n, \mu} \tilde{\psi}_*^{\sigma_n}(z) \in S$ then $\tilde{\phi}^H(w) = z + B(n, \mu) \cap Z(n, \nu)$. This shows $\tilde{\phi}^H$ is surjective. So $H(n, \nu)$ is a QS-group of $H(n, \mu)$ as required. \square

7. Homology computations.

7.1. Contractible model spaces. Let $(\mathcal{X}, d_{\mathcal{X}})$ be compact totally disconnected metric space, $\mu \in \text{Prob}_{\Gamma}(\mathcal{X}^{\Gamma})$, and Σ be a sofic approximation to Γ .

Definition 9. (Contractible model spaces) We say μ has *contractible model spaces* with respect to Σ if for every open neighborhood \mathcal{O}_2 of μ in $\text{Prob}(\mathcal{X}^{\Gamma})$, and every $\delta > 0$ there exists an open neighborhood $\mathcal{O}_1 \subset \text{Prob}(\mathcal{X}^{\Gamma})$ with $\mu \in \mathcal{O}_1 \subset \mathcal{O}_2$ such that for every $0 < K < \infty$ and all but finitely many n , if $x_1, \dots, x_K \in \Omega(\mathcal{O}_1, \sigma_n)$ then there exist $x_i^{(j)}$ (for $j \geq 0$) such that

- (1) $x_i^{(j)} \in \Omega(\mathcal{O}_2, \sigma_n)$ for all i, j ,
- (2) $x_i^{(0)} = x_i$ for all i ,
- (3) $d_{\mathcal{X}}^{V_n}(x_i^{(j+1)}, x_k^{(j+1)}) \leq d_{\mathcal{X}}^{V_n}(x_i^{(j)}, x_k^{(j)})$ for all i, j, k ,
- (4) $d_{\mathcal{X}}^{V_n}(x_i^{(j)}, x_i^{(j+1)}) < \delta$ for all i, j
- (5) there exists M such that $x_1^{(M)} = \dots = x_K^{(M)}$. This M may depend on x_1, \dots, x_K . In particular, M may depend on n .

Definition 10. (Vanishing homology) The measure μ is said to have *vanishing reduced homology in dimension d with respect to Σ* if \forall open neighborhoods $\mathcal{O}_2 \ni \mu$, $\forall \kappa_2 > 0$ there exist an open neighborhood \mathcal{O}_1 with $\mu \in \mathcal{O}_1 \subset \mathcal{O}_2$ and $\kappa_1 > 0$ such that for every $L \in \mathbb{N}$, either $d > 0$ and $H_d^L(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n) = 0$ for all but finitely many n or $d = 0$ and $H_d^L(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n) \cong \mathbb{Z}$ for all but finitely many n . By Theorem 4.3, this notion is a measure-conjugacy invariant. Hence it can also be applied to measure-preserving systems of the form $\Gamma \curvearrowright (X, \mu)$ in which (X, μ) is a standard probability space without any additional structure.

PROPOSITION 7.1. *If μ has contractible model spaces with respect to Σ then μ has vanishing reduced homology in every dimension with respect to Σ .*

Remark 4. The proof of this Proposition is the only place in this paper where the finiteness of the parameter L in the definition of the homology groups $H_d^L(\cdot)$ is used directly.

Proof. Let \mathcal{O}_2 be an open neighborhood of μ and $\kappa_2 > 0$. Choose $0 < \delta < \kappa_2/3$ and let $\kappa_1 := \delta$. Let \mathcal{O}_1 be as in the definition of contractibility. Let $d \in \mathbb{N}$ be a dimension and $L > 0$. Let $K \geq (d+1)L$. Let

$$z_n = \sum_{i=1}^k c_i s_i \in Z_d^L(\mathcal{O}_1, \kappa_1, \sigma_n)$$

be a cycle of length $k \leq L$. If $d > 0$ then it suffices to show that $z_n \in B_d(\mathcal{O}_2, \kappa_2, \sigma_n)$.

Let $x_0, \dots, x_m \in \Omega(\mathcal{O}_1, \sigma_n)$ be an enumeration of the vertices contained in the oriented simplices s_1, \dots, s_k . Note $m \leq (d+1)L \leq K$ is bounded independently of n .

Let $x_i^{(j)}$ be as in the definition of contractible. The map $\psi^{(j)} : \{x_0, \dots, x_m\} \rightarrow \mathcal{X}^{V_n}$ defined by $\psi^{(j)}(x_i) = x_i^{(j)}$ is well defined. We can extend it to a map on all of \mathcal{X}^{V_n} by setting $\psi^{(j)}(x) = x$ for all $x \notin \{x_0, \dots, x_m\}$. In particular, $\psi_*^{(j)}(s_i)$ is well defined for every simplex s_i and therefore $\psi_*^{(j)}(z_n)$ is also well defined as an element of $C_d(\mathcal{X}^{V_n})$.

We claim that $\psi_*^{(j)}(z_n) \in Z_d(\mathcal{O}_2, \kappa_1, \sigma_n)$ for all j . For $j = 0$ this is true since $\psi_*^{(0)}(z_n) = z_n$. Assuming it is true for some $j \geq 0$, observe that by Lemma 6.1,

$$\psi_*^{(j+1)}(z_n) - \psi_*^{(j)}(z_n) \in B_d(\mathcal{O}_2, \kappa_2, \sigma_n)$$

(this uses $\kappa_1 < \kappa_2/3$). Moreover $\psi_*^{(j+1)}(s_i) \in C_d(\mathcal{O}_2, \kappa_1, \sigma_n)$ by property (3) in the definition of contractibility. Thus $\psi_*^{(j+1)}(z_n) \in Z_d(\mathcal{O}_2, \kappa_1, \sigma_n)$. This completes the induction.

Moreover, we showed that $\psi_*^{(j)}(z_n) - z_n \in B_d(\mathcal{O}_2, \kappa_2, \sigma_n)$ for all j . If $d > 0$ then $\psi_*^{(M)}(z_n)$ is trivial (by property (4)) and so $z_n \in B_d(\mathcal{O}_2, \kappa_2, \sigma_n)$. This completes the proof in the case $d > 0$.

If $d = 0$ then every element $C_0(\mathcal{O}_1, \kappa_1, \sigma_n)$ is a cycle (so $Z_0^L(\mathcal{O}_1, \kappa_1, \sigma_n)$ is the free abelian group generated by $\Omega(\mathcal{O}_1, \sigma_n)$). We have shown for any $x_1, x_2 \in \Omega(\mathcal{O}_1, \sigma_n)$ there is a δ -path in $\Omega(\mathcal{O}_2, \sigma_n)$ connecting them (namely $x_1^{(0)}, \dots, x_1^{(M)}, x_2^{(M)}, \dots, x_2^{(0)}$). Therefore $x_1 - x_2 \in B_0(\mathcal{O}_2, \kappa_1, \sigma_n)$ which implies $H_d^L(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n) \cong \mathbb{Z}$ for all but finitely many n . \square

7.2. Edit distance. The goal of this section is to show that the homological invariants defined above depend on the sofic approximation only up to edit-distance zero as defined next.

For each finite set $S \subset \Gamma$ and finite set V there is a pseudo-metric d_S on the set of all maps $\sigma : \Gamma \rightarrow \text{sym}(V)$ defined by

$$d_S(\sigma_1, \sigma_2) := |V|^{-1} \# \{v \in V : \exists s \in S \text{ such that } \sigma_1(s)v \neq \sigma_2(s)v\}.$$

Sofic approximations $\Sigma = \{\sigma_n : \Gamma \rightarrow V_n\}$ and $\Sigma' = \{\sigma'_n : \Gamma \rightarrow V_n\}$ with the same target sets $\{V_n\}_n$ are said to be at *edit-distance zero* if for every finite $S \subset \Gamma$,

$$\limsup_{n \rightarrow \infty} d_S(\sigma_n, \sigma'_n) = 0.$$

PROPOSITION 7.2. *Suppose Σ, Σ' are at edit-distance zero and $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$. If μ has contractible model spaces with respect to Σ then μ also has contractible model spaces with respect to Σ' .*

For the proof it will be helpful to have the next definition.

Definition 11. An open subset $\mathcal{O} \subset \text{Prob}(\mathcal{X}^\Gamma)$ is *D-local* (where $D \subset \Gamma$ is finite) if there is an open set $\tilde{\mathcal{O}} \subset \text{Prob}(\mathcal{X}^D)$ such that \mathcal{O} is the inverse image of $\tilde{\mathcal{O}}$ under the projection map $\text{Prob}(\mathcal{X}^\Gamma) \rightarrow \text{Prob}(\mathcal{X}^D)$. An open subset $\mathcal{O} \subset \text{Prob}(\mathcal{X}^\Gamma)$ is *local* if it is *D-local* for some D .

The next lemma can be used to show that many sofic invariants (such as entropy) depend on Σ only up to edit-distance zero. We will use it to prove Proposition 7.2.

LEMMA 7.3. *Suppose Σ, Σ' have edit-distance zero and $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$. Then for any open neighborhood $\mathcal{O} \ni \mu$ there exists an open neighborhood \mathcal{O}' with $\mu \in \mathcal{O}' \subset \mathcal{O}$ such that $\Omega(\mathcal{O}', \sigma'_n) \subset \Omega(\mathcal{O}, \sigma_n)$ for all but finitely many n .*

Proof. Let an open set $\mathcal{O} \ni \mu$ be given. Let \mathcal{O}' be an open set containing μ such that the weak* closure of \mathcal{O}' is contained in \mathcal{O} . By compactness and Lemma 5.7, there exists a finite set $E \subset \Gamma$ and $\epsilon > 0$ such that if $\nu \in \mathcal{O}'$ and $\nu' \in \text{Prob}(\mathcal{X}^E)$ is such that $d_{\text{TV}}(\text{Res}_*^E \nu, \text{Res}_*^E \nu') < \epsilon$ then $\nu' \in \mathcal{O}$.

Because Σ, Σ' have edit-distance zero, $d_E(\sigma_n, \sigma'_n) < \epsilon$ for all but finitely n . This condition implies $d_{\text{TV}}(\text{Res}_*^E(P_x^{\sigma_n}), \text{Res}_*^E(P_x^{\sigma'_n})) < \epsilon$ for all $x \in \mathcal{X}^{V_n}$. In particular, if $x \in \Omega(\mathcal{O}', \sigma'_n)$ then $x \in \Omega(\mathcal{O}, \sigma_n)$. Thus $\Omega(\mathcal{O}', \sigma'_n) \subset \Omega(\mathcal{O}, \sigma_n)$. \square

Proof of Proposition 7.2. Let an open set $\mathcal{O}'_2 \ni \mu$ and $\delta > 0$ be given. By Lemma 7.3 there exists an open set \mathcal{O}_2 such that $\mu \in \mathcal{O}_2 \subset \mathcal{O}'_2$ and $\Omega(\mathcal{O}_2, \sigma_n) \subset \Omega(\mathcal{O}'_2, \sigma'_n)$ for all but finitely many n .

Let $\mathcal{O}_1 \subset \mathcal{O}_2$ be an open neighborhood of μ satisfying the definition of contractible model spaces with respect to Σ . So for every $0 < K < \infty$ and all but finitely many n , if $x_i \in \Omega(\mathcal{O}_1, \sigma_n)$ ($1 \leq i \leq K$) there exist $x_i^{(j)}$ (for $j \geq 0$) such that

- (1) $x_i^{(j)} \in \Omega(\mathcal{O}_2, \sigma_n)$ for all i, j ,
- (2) $x_i^{(0)} = x_i$ for all i ,
- (3) $d_{\mathcal{X}}^{V_n}(x_i^{(j+1)}, x_k^{(j+1)}) \leq d_{\mathcal{X}}^{V_n}(x_i^{(j)}, x_k^{(j)})$ for all i, j, k ,
- (4) $d_{\mathcal{X}}^{V_n}(x_i^{(j)}, x_i^{(j+1)}) < \delta$ for all i, j ,
- (5) there exists M such that $x_1^{(M)} = \dots = x_K^{(M)}$.

By Lemma 7.3 there exists an open set \mathcal{O}'_1 such that $\mu \in \mathcal{O}'_1 \subset \mathcal{O}_1$ and $\Omega(\mathcal{O}'_1, \sigma'_n) \subset \Omega(\mathcal{O}_1, \sigma_n)$ for all but finitely many n .

Now let $x_1, \dots, x_K \in \Omega(\mathcal{O}'_1, \sigma'_n)$. Since $\Omega(\mathcal{O}'_1, \sigma'_n) \subset \Omega(\mathcal{O}_1, \sigma_n)$, there exist $x_i^{(j)}$ for $j \geq 0$ satisfying the above conditions (if n is sufficiently large). In particular,

$$x_i^{(j)} \in \Omega(\mathcal{O}_2, \sigma_n) \subset \Omega(\mathcal{O}'_2, \sigma'_n)$$

for n sufficiently large (independent of x_1, \dots, x_K). This proves μ has contractible model spaces with respect to Σ' . \square

7.3. Diffuse sofic approximations. One of the main goals of this section is to prove that if Γ is amenable then every $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$ has contractible model spaces. In fact, more is true, one only needs that the sofic approximation Σ is diffuse. This condition, explained below, holds automatically if Γ is amenable. Moreover, even if Γ is non-amenable then diffuse sofic approximations can be constructed out of arbitrary sofic approximations.

The *disjoint union* of maps $\sigma_i : \Gamma \rightarrow \text{sym}(V_i)$ (for $i = 1, 2$) is the map $\sigma_1 \sqcup \sigma_2 : \Gamma \rightarrow \text{sym}(V_1 \sqcup V_2)$ defined by $\sigma_1 \sqcup \sigma_2(g)v = \sigma_i(g)v$ if $v \in V_i$.

Definition 12. A sofic approximation $\Sigma = \{\sigma_n : \Gamma \rightarrow \text{sym}(V_n)\}_{n \in \mathbb{N}}$ is *diffuse* if there exists a sofic approximation $\Sigma' = \{\sigma'_n : \Gamma \rightarrow \text{sym}(V_n)\}_{n \in \mathbb{N}}$ such that

- Σ and Σ' have edit-distance zero,
- for every n , σ'_n can be expressed as a disjoint union $\sigma'_n := \sigma'_{n,1} \sqcup \dots \sqcup \sigma'_{n,m_n}$ such that if $\sigma'_{n,i} : \Gamma \rightarrow \text{sym}(V_{n,i})$ then

$$\max_{1 \leq i \leq m_n} |V_{n,i}| = o(|V_n|).$$

Example 1. Let $\Sigma = \{\sigma_n : \Gamma \rightarrow \text{sym}(V_n)\}_{n \in \mathbb{N}}$ by any sofic approximation to any group Γ and let $\{W_n\}_n$ be a sequence of finite sets with $|W_n| \rightarrow \infty$ as $n \rightarrow \infty$. Define $\sigma'_n : \Gamma \rightarrow \text{sym}(V_n \times W_n)$ by

$$\sigma'_n(g)(v, w) = (\sigma_n(g)v, w).$$

In other words, σ'_n is the direct product of σ_n with the trivial homomorphism $\Gamma \rightarrow \text{sym}(W_n)$. Then $\Sigma' = \{\sigma'_n\}$ is diffuse. In fact σ'_n is the disjoint union of W_n copies of σ_n .

LEMMA 7.4. *If Γ is amenable then every sofic approximation Σ to Γ is diffuse.*

Proof. The special case in which Γ is finitely generated is a direct consequence of [15, Proposition 2.8]. The general case follows from the finitely generated case by a diagonalization argument. \square

7.3.1. Diffuse approximations and contractibility.

THEOREM 7.5. *If Σ is diffuse and $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$ then μ has contractible model spaces with respect to Σ . In particular, if Γ is amenable then μ has contractible model spaces with respect to every sofic approximation.*

Because Σ is diffuse, after replacing it with another sofic approximation at edit-distance zero, we may assume that each set V_n comes equipped with a partition $V_n = \sqcup_i V_{n,i}$ so that the image of $\sigma_n : \Gamma \rightarrow \text{sym}(V_n)$ lies in the direct product $\prod_i \text{sym}(V_{n,i})$. The idea behind the proof is to modify $x_1, \dots, x_K \in \mathcal{X}^{V_n}$ on each of the $V_{n,i}$'s to form paths $(x_i^{(j)})_{i,j}$ that merge together. However, we must be careful so that each of the $x_i^{(j)}$'s has empirical distribution close to μ . To accomplish this, we coarsen the given partition so that each x_i restricted to any part of the good partition has empirical distribution close to μ . This explains why the next two lemmas are needed. (No attempt has been made to optimize the constants below).

LEMMA 7.6. *Let (Ω, \mathbb{P}) be a standard probability space and $f : \Omega \rightarrow [0, 1]$ be a measurable random variable. Let $\delta \in (0, 1/4)$, $0 < \epsilon < \delta/400$ and suppose that $\mathbb{P}(\{\omega\}) < \epsilon$ for every $\omega \in \Omega$. Then there exists a finite measurable partition \mathcal{P} of Ω such that*

- (1) $\mathbb{P}(P) \leq 100\epsilon/\delta$ for every $P \in \mathcal{P}$,
- (2) $\|\mathbb{E}[f|\mathcal{P}] - \mathbb{E}[f]\|_{L^\infty(\Omega, \mathbb{P})} \leq \delta$

where \mathbb{E} denotes expectation with respect to \mathbb{P} .

Proof. After passing to an image of the measure space (Ω, \mathbb{P}) if necessary, we may assume without loss of generality that Ω is a finite set. So let $\Omega = \{\omega_1, \dots, \omega_n\}$ be ordered so that $f(\omega_1) \leq f(\omega_2) \leq \dots \leq f(\omega_n)$. After making a small perturbation if necessary, we may also assume that $f(\omega) \neq \mathbb{E}[f]$ for any $\omega \in \Omega$.

Define a piecewise constant function $F : [0, 1] \rightarrow [0, 1]$ by

$$F(x) = f(\omega_{h(x)})$$

where $h(x)$ is the smallest number such that $x \leq \mathbb{P}(\{\omega_1, \dots, \omega_{h(x)}\})$. Note that $\int F dx = \mathbb{E}[f]$ and F is monotone increasing. We will first solve the problem with (Ω, \mathbb{P}) and f replaced by Lebesgue measure on the unit interval and F .

Claim 1. There exists a natural number m such that $20\epsilon/\delta \leq 1/m \leq 21\epsilon/\delta$.

Proof. Let $m \in \mathbb{N}$ be such that $\frac{1}{m+1} < 20\epsilon/\delta \leq \frac{1}{m}$. Because $\epsilon < \delta/400$, $20\epsilon/\delta < 1/20$ and $m \geq 20$.

It suffices to show $1/m \leq 21\epsilon/\delta$. Equivalently, it suffices to show $21m\epsilon/\delta \geq 1$. The condition $\frac{1}{m+1} < 20\epsilon/\delta$ implies $20(m+1)\epsilon/\delta > 1$. Because $m \geq 20$, we have $21m \geq 20(m+1)$. So $21m\epsilon/\delta \geq 20(m+1)\epsilon/\delta > 1$. \square

Claim 2. There is a finite partition \mathcal{Q} of $[0, 1]$ such that every $Q \in \mathcal{Q}$ satisfies

- (1) Q is a union of at most two disjoint intervals;
- (2) $20\epsilon/\delta \leq |Q| \leq 21\epsilon/\delta$ where $|\cdot|$ denotes Lebesgue measure;
- (3) $|Q|^{-1} \int_Q F(x) dx = \mathbb{E}[f]$.

Proof. Because $f(\omega) \neq \mathbb{E}[f]$ for any $\omega \in \Omega$, $F(t) \neq \int_0^1 F(x) dx$ for any t . Since F is monotone increasing, there exists a unique $x_0 \in [0, 1]$ such that for all a and b with $0 \leq a < x_0 < b \leq 1$,

$$F(a) < \int_0^1 F dx < F(b).$$

Because F is monotone, for every $0 \leq a < x_0$, there exists a unique $g(a)$ with $x_0 < g(a) \leq 1$ such that

$$\frac{1}{g(a)-a} \int_a^{g(a)} F dx = \int_0^1 F dx = \mathbb{E}[f].$$

Note g is an orientation-reversing homeomorphism of $[0, x_0]$ onto $(x_0, 1]$.

Using Claim 1, choose points $0 = a_1 < a_2 < \dots < a_m < x_0$ such that for each $1 \leq i < m$,

$$20\epsilon/\delta \leq 1/m = (a_{i+1} - a_i) + g(a_i) - g(a_{i+1}) \leq 21\epsilon/\delta$$

and

$$20\epsilon/\delta \leq 1/m = g(a_m) - a_m \leq 21\epsilon/\delta.$$

Let \mathcal{Q} be the partition containing $[a_m, g(a_m)]$ and $[a_i, a_{i+1}] \cup (g(a_{i+1}), g(a_i)]$ for $i < m$. \square

Claim 2 solves the problem for F . To solve it for the original function f , we will obtain a partition \mathcal{P} that approximates \mathcal{Q} .

Recall that for $x \in [0, 1]$, $h(x)$ is the smallest index such that $x \leq \mathbb{P}(\{\omega_1, \dots, \omega_{h(x)}\})$. Define $h_\Omega : [0, 1] \rightarrow \Omega$ by $h_\Omega(x) = \omega_{h(x)} \in \Omega$. For $X \subset [0, 1]$, let

$$\begin{aligned} X^+ &= h_\Omega^{-1}(h_\Omega(X)), \\ X^- &= \{x \in X : h_\Omega^{-1}(h_\Omega(x)) \subset X\}. \end{aligned}$$

Let $\tilde{\mathcal{P}}$ be an h_Ω -measurable partition of $[0, 1]$ such that for every $P \in \tilde{\mathcal{P}}$ there exists $Q \in \mathcal{Q}$ such that $Q^- \subset P \subset Q^+$. Because Q is a disjoint union of at most two intervals, $h_\Omega(Q^+) \setminus h_\Omega(Q^-)$ contains at most 4 elements of Ω . Since each element of Ω has measure $< \epsilon$ and h_Ω is measure-preserving, $|Q \Delta P| \leq 4\epsilon$. So

$$|P| \leq |Q| + 4\epsilon \leq 21\epsilon/\delta + 4\epsilon \leq 100\epsilon/\delta.$$

Also

$$\left| |P|^{-1} \int_P F dx - \mathbb{E}[f] \right| = \left| |P|^{-1} \int_P F(x) dx - |Q|^{-1} \int_Q F(x) dx \right|.$$

We may decompose the first integral as $\int_P = \int_{P \cap Q} + \int_{P \setminus Q}$ and the second one similarly. From this, we see that the above is bounded by

$$\begin{aligned} & \left| |P|^{-1} - |Q|^{-1} \right| \int_{P \cap Q} F(x) dx + |P|^{-1} \int_{P \setminus Q} F(x) dx + |Q|^{-1} \int_{Q \setminus P} F(x) dx \\ & \leq \left| |P|^{-1} - |Q|^{-1} \right| |P \cap Q| + |P|^{-1} |P \setminus Q| + |Q|^{-1} |Q \setminus P| \\ & \leq 4\epsilon \left(\frac{|Q \cap P|}{|P||Q|} + |P|^{-1} + |Q|^{-1} \right). \end{aligned}$$

We use the bounds $|P \cap Q| \leq 21\epsilon/\delta$ and $|Q| \geq 20\epsilon/\delta$, $|P| \geq 20\epsilon/\delta - 4\epsilon \geq 19\epsilon/\delta$ to obtain

$$\left| |P|^{-1} \int_P F dx - \mathbb{E}[f] \right| < \delta.$$

Since P is arbitrary, this shows $\|\mathbb{E}[F|\tilde{\mathcal{P}}] - \mathbb{E}[F]\|_\infty \leq \delta$. Since $\tilde{\mathcal{P}}$ is h_Ω -measurable, it induces a partition $\mathcal{P} = h_\Omega(\tilde{\mathcal{P}})$ on Ω . Because h_Ω is measure-preserving, the required properties of \mathcal{P} follow from the corresponding properties of $\tilde{\mathcal{P}}$. \square

LEMMA 7.7. *Let (Ω, \mathbb{P}) be a standard probability space and $f_1, \dots, f_m : \Omega \rightarrow [0, 1]$ be random variables. Let $0 < \delta < 1/4$, $0 < \epsilon$, and suppose that $\mathbb{P}(\{\omega\}) < \epsilon$ for every $\omega \in \Omega$. Assume $\epsilon < (\delta/400)^m$. Then there exists a measurable partition \mathcal{P} of Ω such that*

- (1) $\mathbb{P}(P) \leq \epsilon(100/\delta)^m$ for every $P \in \mathcal{P}$,
- (2) $\|\mathbb{E}[f_i|\mathcal{P}] - \mathbb{E}[f_i]\|_{L^\infty(\Omega, \mathbb{P})} \leq \delta$ for every $1 \leq i \leq m$.

Proof. We prove this by induction on m . The previous lemma establishes the base case $m = 1$.

For the inductive step, assume $m \geq 2$ and there is a measurable partition \mathcal{Q} of Ω such that

- (1) $\mathbb{P}(Q) \leq \epsilon(100/\delta)^{m-1}$ for every $Q \in \mathcal{Q}$,
- (2) $\|\mathbb{E}[f_i|\mathcal{Q}] - \mathbb{E}[f_i]\|_{L^\infty(\Omega, \mathbb{P})} \leq \delta$ for every $1 \leq i \leq m-1$.

Apply the previous lemma with $(\mathcal{Q}, \mathbb{P})$ in place of (Ω, \mathbb{P}) , $\mathbb{E}[f_m|\mathcal{Q}]$ in place of f and $\epsilon(100/\delta)^{m-1}$ in place of ϵ to obtain a partition \mathcal{P} of Ω that coarsens \mathcal{Q} and satisfies

- (1) $\mathbb{P}(P) \leq \epsilon(100/\delta)^m$ for every $P \in \mathcal{P}$,
- (2) $\|\mathbb{E}[\mathbb{E}[f_m|\mathcal{Q}]|\mathcal{P}] - \mathbb{E}[f_m]\|_{L^\infty(\Omega, \mathbb{P})} = \|\mathbb{E}[f_m|\mathcal{P}] - \mathbb{E}[f_m]\|_{L^\infty(\Omega, \mathbb{P})} \leq \delta$.

Since \mathcal{P} coarsens \mathcal{Q} , $\|\mathbb{E}[f_i|\mathcal{P}] - \mathbb{E}[f_i]\|_{L^\infty(\Omega, \mathbb{P})} \leq \delta$ holds for every $1 \leq i \leq m-1$ too. \square

Proof of Theorem 7.5. By Proposition 7.2, we may assume without loss of generality that for every n , σ_n can be expressed as a disjoint union $\sigma_n := \sigma_{n,1} \sqcup \dots \sqcup \sigma_{n,m_n}$ such that if $\sigma_{n,i} : \Gamma \rightarrow \text{sym}(V_{n,i})$ then

$$\max_{1 \leq i \leq m_n} |V_{n,i}| = o(|V_n|).$$

Let $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$, \mathcal{O}_2 be an open neighborhood of μ in $\text{Prob}(\mathcal{X}^\Gamma)$ and $\delta > 0$ be given. By choosing \mathcal{O}_2 smaller if necessary, we may assume it is convex. Let $\mathcal{O}_1 \subset \mathcal{O}_2$ be an open neighborhood of μ such that the closure of \mathcal{O}_1 is contained in \mathcal{O}_2 . It follows that there exist $\delta_2 > 0$ and continuous functions $\tilde{f}_1, \dots, \tilde{f}_m : \mathcal{X}^\Gamma \rightarrow [0, 1]$ such that if $\nu \in \mathcal{O}_1$ and ν' is such that $|\nu(\tilde{f}_i) - \nu'(\tilde{f}_i)| < \delta_2$ for all $1 \leq i \leq m$ then $\nu' \in \mathcal{O}_2$.

Let $x_0, \dots, x_K \in \Omega(\mathcal{O}_1, \sigma_n)$. Let $\Omega = \{1, \dots, m_n\}$ and \mathbb{P} be the probability measure on Ω given by $\mathbb{P}(k) = |V_{n,k}|/|V_n|$.

Let $f_{i,j} : V_n \rightarrow [0, 1]$ be given by

$$f_{i,j}(v) = \tilde{f}_i(\Pi_v^{\sigma_n} x_j).$$

Define $\tilde{f}_{i,j} : \Omega \rightarrow [0, 1]$ by

$$\tilde{f}_{i,j}(k) = |V_{n,k}|^{-1} \sum_{v \in V_{n,k}} f_{i,j}(v).$$

For sufficiently large n , apply Lemma 7.7 to (Ω, \mathbb{P}) and the functions $(\tilde{f}_{i,j})$ to obtain a partition \mathcal{P}_n of Ω . Let \mathcal{Q}_n be the partition of V_n defined by pulling back the partition \mathcal{P}_n under the map $V_n \rightarrow \Omega$ defined by $v \mapsto i$ if $v \in V_{n,i}$. Then

- (1) \mathcal{Q}_n coarsens the partition $V_n = \sqcup V_{n,i}$.
- (2) For each i, j , $\|\mathbb{E}[f_{i,j} | \mathcal{Q}_n] - \mathbb{E}[f_{i,j}] \|_{L^\infty(\Omega, \mathbb{P})} \leq \delta_2$. By the choice of $f_{i,j}$ and δ_2 , this implies that for every $Q \in \mathcal{Q}_n$, the empirical measure of $x_j \upharpoonright Q$ with respect to $\sigma_n \upharpoonright Q$ lies in \mathcal{O}_2 (where $\sigma_n \upharpoonright Q$ is the map $\Gamma \rightarrow \text{sym}(Q)$ obtained by restriction). In symbols, $P_{x_j \upharpoonright Q}^{\sigma_n \upharpoonright Q} \in \mathcal{O}_2$.

(3) $\max_{Q \in \mathcal{Q}_n} |Q| \leq \epsilon_n |V_n|$ for some constants $\epsilon_n > 0$ with $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$.

For $0 \leq j$, define $\Psi_j : \mathcal{X}^{V_n} \rightarrow \mathcal{X}^{V_n}$ by $\Psi_j(x)(v) = x_0(v)$ if $v \in Q_i$ for some $i \leq j$ and $\Psi_j(x)(v) = x(v)$ otherwise. Set $x_i^{(j)} = \Psi_j(x_i)$.

Since $P_{x_i \upharpoonright Q}^{\sigma_n \upharpoonright Q} \in \mathcal{O}_2$ for every $Q \in \mathcal{Q}_n$ and \mathcal{O}_2 is convex, $x_i^{(j)} \in \Omega(\mathcal{O}_2, \sigma_n)$ for all i, j .

If n is large enough then $\epsilon_n \text{diam}(\mathcal{X}) < \delta$ and therefore $d_{\mathcal{X}}^{V_n}(x_i^{(j)}, x_i^{(j+1)}) < \delta$ for all i, j . The map Ψ_j is distance contracting, Ψ_0 is the identity and Ψ_{m_n} maps all of $\{x_0, \dots, x_K\}$ to x_0 . This verifies all of the conditions in the definition of contractible model spaces. \square

7.4. Bernoulli shifts.

THEOREM 7.8. *Let Γ be a countably infinite group, Σ a sofic approximation, $(\mathcal{X}, d_{\mathcal{X}}), (\mathcal{Y}, d_{\mathcal{Y}})$ totally disconnected compact metric spaces and $\beta \in \text{Prob}(\mathcal{X})$ a probability measure on \mathcal{X} . Then the Bernoulli shift $\Gamma \curvearrowright (\mathcal{X}, \beta)^\Gamma$ has contractible model spaces with respect to Σ . Moreover, for any $\nu \in \text{Prob}_\Gamma(\mathcal{Y}^\Gamma)$, the 0-dimensional sofic homology theories of ν and $\beta^\Gamma \times \nu$ are equivalent.*

In [2], Tim Austin proved a similar result. The proof here follows the same strategy.

Remark 5. Unfortunately, it is not clear whether the d -dimensional sofic homology theories of ν and $\beta^\Gamma \times \nu$ are equivalent for $d > 0$. To explain why, suppose $z, z' \in Z_d^L(\mathcal{O}_1, \kappa_1, \sigma)$ are each representable as a weighted sum of at most L simplices and $z - z' \in B_d(\mathcal{O}_2, \kappa_2, \sigma)$. Then there is a $(d+1)$ -chain w such that $z - z' = \partial_d w$. But the proof in the 0-dimensional case uses that there exists a sequence $z = z_1, \dots, z_k = z'$ of cycles interpolating between z and z' such that there is a uniform bound on the ℓ^1 -norm $\|z_i - z_{i+1}\|_1$ for each i and each z_i is representable as a weighted sum of at most L' cycles for some L' that does not depend on z, z' . There does not appear to be any good reason why this property should hold if $d > 0$.

We will use special neighborhoods of $\beta^\Gamma \times \nu$ defined as follows. For any finite $D \subset \Gamma$, let $\mathcal{B}_{\mathcal{X}^D}$ be the smallest Borel sub-sigma-algebra on $\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma$ such that the projection $\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma \rightarrow \mathcal{X}^D \times \mathcal{Y}^D$ is $\mathcal{B}_{\mathcal{X}^D}$ -measurable.

A subset $\mathcal{F} \subset C(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma)$ of continuous functions is *hereditary* if there is some finite $D \subset \Gamma$ such that every $f \in \mathcal{F}$ is D -local and the conditional expectation $\mathbb{E}_{\beta^\Gamma \times \nu}[f | \mathcal{B}_{\mathcal{X}^C}] \in \mathcal{F}$ for every $C \subset D$. Moreover, we require that f is 1-Lipschitz in the \mathcal{X}^D -variable as a function from $\mathcal{X}^D \times \mathcal{Y}^D$ to \mathbb{R} . To be precise this means that

$$|f(x_1, y) - f(x_2, y)| \leq d_{\mathcal{X}}^D(x_1, x_2)$$

for every $x_1, x_2 \in \mathcal{X}^D$ and $y \in \mathcal{Y}^D$ (where we have abused notation by identifying f with its projection to $\mathcal{X}^D \times \mathcal{Y}^D$).

A neighborhood $\mathcal{O} \subset \text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma)$ is *hereditary* if there is a finite hereditary subset $\mathcal{F} \subset C(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma)$ and $\delta > 0$ such that

$$\mathcal{O} = \{\mu' \in \text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma) : |\mu'(f) - \beta^\Gamma \times \nu(f)| < \delta \ \forall f \in \mathcal{F}\}.$$

The next proposition shows that the measure $\beta^\Gamma \times \nu$ satisfies a property that is a kind of relative version of having contractible model spaces.

PROPOSITION 7.9. *Let $\mathcal{O}_1 \subset \mathcal{O}_2 \subset \text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma)$ be open neighborhoods of $\beta^\Gamma \times \nu$. Suppose that the closure of \mathcal{O}_1 is contained in \mathcal{O}_2 and that \mathcal{O}_1 is hereditary.*

Then for any $\kappa > 0$ and $K > 0$ there exists N such that if $n > N$ then for any $(x_i, y_i) \in \Omega(\mathcal{O}_1, \sigma_n)$ ($1 \leq i \leq K$) there exist $x_i^{(j)}$ (for $j \geq 0$) such that for all i, j, k ,

- (1) $(x_i^{(j)}, y_i) \in \Omega(\mathcal{O}_2, \sigma_n)$,
- (2) $x_i^{(0)} = x_i$,
- (3) $d_{\mathcal{X}}^{V_n}(x_i^{(j)}, x_i^{(j+1)}) < \kappa$,
- (4) $d_{\mathcal{X}}^{V_n}(x_i^{(j+1)}, x_k^{(j+1)}) \leq d_{\mathcal{X}}^{V_n}(x_i^{(j)}, x_k^{(j)})$ and
- (5) there exists M such that $x_1^{(M)} = \dots = x_K^{(M)}$. In fact we may choose the $x_i^{(j)}$'s so that $M = \lceil \text{diam}(\mathcal{X})/\kappa \rceil + 1$.

LEMMA 7.10. *Let $\mathcal{O} \subset \text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma)$ be an open neighborhood of $\beta^\Gamma \times \nu$. Then there exists an open neighborhood $\mathcal{O}' \subset \text{Prob}(\mathcal{Y}^\Gamma)$ of ν such that for all but finitely many n , if $y \in \Omega(\mathcal{O}', \sigma_n)$ then there exists $x \in \mathcal{X}^{V_n}$ such that $(x, y) \in \Omega(\mathcal{O}, \sigma_n)$.*

Proof. This is equivalent to saying that $\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma \rightarrow \mathcal{X}^\Gamma$ is model-surjective (in the language of [2]). This is proven implicitly in [9, Theorem 8.1] and [2, Theorem 6.8]. \square

Proof of Theorem 7.8 given Proposition 7.9. The fact that Bernoulli shifts have contractible model spaces is implied by the special case of Proposition 7.9 in which ν is the Dirac mass on a fixed point.

Before proving the second statement, note that the projection map $\text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma) \rightarrow \text{Prob}(\mathcal{Y}^\Gamma)$ is open (this is implied by [14, Theorem 2.5]). So if $\mathcal{O} \subset \text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma)$ is open then its image, which we denote by $\text{Proj}_{\mathcal{Y}^\Gamma}(\mathcal{O})$, is an open subset of $\text{Prob}(\mathcal{Y}^\Gamma)$.

First we will show that the 0-dimensional sofic homology of ν is greater than or equal to the 0-dimensional sofic homology of $\beta^\Gamma \times \nu$. In order to define the homology of $\beta^\Gamma \times \nu$, we identify $\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma$ with $(\mathcal{X} \times \mathcal{Y})^\Gamma$ and use the metric $d_{\mathcal{X} \times \mathcal{Y}}((x, y), (x', y')) = d_{\mathcal{X}}(x, x') + d_{\mathcal{Y}}(y, y')$.

Let \mathcal{O}_2 be an arbitrary open neighborhood of $\beta^\Gamma \times \nu$ in $\text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma)$. Also let $\kappa_2 > 0$. By Lemma 7.10 there exists an open neighborhood $\mathcal{O}'_2 \subset \text{Proj}_{\mathcal{Y}^\Gamma}(\mathcal{O}_2)$ of ν such that for all but finitely many n , if $y \in \Omega(\mathcal{O}'_2, \sigma_n)$ then there exists $x \in \mathcal{X}^{V_n}$ such that $(x, y) \in \Omega(\mathcal{O}_2, \sigma_n)$. Let $\kappa'_2 = \kappa_2$. Let $\mathcal{O}'_1 \subset \mathcal{O}'_2$ be an arbitrary open neighborhood of ν . Also let $0 < \kappa'_1$ be arbitrary. Choose an open neighborhood \mathcal{O}_1 of $\beta^\Gamma \times \nu$ so that its closure is contained in \mathcal{O}_2 and $\text{Proj}_{\mathcal{Y}^\Gamma}(\mathcal{O}_1)$ is contained in \mathcal{O}'_1 . Because hereditary neighborhoods form a basis, we may also choose \mathcal{O}_1 to be hereditary. Let $\kappa_1 = \kappa'_1$.

Let $\pi : (\mathcal{X} \times \mathcal{Y})^{V_n} \rightarrow \mathcal{Y}^{V_n}$ be the projection map. Let S be the subgroup of $H_0(\mathcal{O}'_1, \mathcal{O}'_2, \kappa'_1, \kappa'_2, \sigma_n)$ generated by the set of all 0-chains of the form $\pi_*([x, y]) = [y]$ for $(x, y) \in \Omega(\mathcal{O}_1, \sigma_n)$. We claim that the map

$$\pi_*([x, y]) \longmapsto [x, y]$$

from S to $H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n)$ is well defined. It suffices to show that if (x_1, y_1) and (x_2, y_2) are both in $\Omega(\mathcal{O}_1, \sigma_n)$ and $y_1 - y_2 \in B_0(\mathcal{O}'_2, \kappa_2, \sigma_n)$ then there is a κ_2 -path from (x_1, y_1) to (x_2, y_2) in $\Omega(\mathcal{O}_2, \sigma_n)$.

Because $y_1 - y_2 \in B_0(\mathcal{O}'_2, \kappa_2, \sigma_n)$ there exists a κ_2 -path $w_1, \dots, w_m \in \Omega(\mathcal{O}'_2, \sigma_n)$ from $w_1 = y_1$ to $w_m = y_2$. By Lemma 7.10, there exist $u_1, \dots, u_m \in \mathcal{X}^{V_n}$ such that $(u_i, w_i) \in \Omega(\mathcal{O}_2, \sigma_n)$ for all i . We may assume $u_1 = x_1$, $u_m = x_2$.

Since $(x_1, y_1) - (x_2, y_2) = \sum_{i=1}^{m-1} (u_i, w_i) - (u_{i+1}, w_{i+1})$, it suffices to show that $(u_i, w_i) - (u_{i+1}, w_{i+1}) \in B_0(\mathcal{O}_2, \kappa_2, \sigma_n)$ for all $1 \leq i < m$.

So fix i with $1 \leq i < m$. By Proposition 7.9 (with $K = 2$, $\kappa = \kappa_1$ and all large n) for $j \geq 0$ there exist elements $u_i^{(j)}, u_{i+1}^{(j)}$ such that for all $k \in \{i, i+1\}$ and $j \geq 0$,

- (1) $(u_k^{(j)}, w_k) \in \Omega(\mathcal{O}_2, \sigma_n)$,
- (2) $u_k^{(0)} = u_k$,
- (3) $d_{\mathcal{X}^{V_n}}(u_k^{(j)}, u_k^{(j+1)}) < \kappa_1$,
- (4) there exists M_i such that $u_i^{(M_i)} = u_{i+1}^{(M_i)}$.

It follows that

$$(u_i, w_i), (u_i^{(1)}, w_i), \dots, (u_i^{(M_i)}, w_i), (u_{i+1}^{(M_i)}, w_{i+1}), \dots, (u_{i+1}^{(1)}, w_{i+1}), (u_{i+1}, w_{i+1})$$

is a κ_2 -path from (u_i, w_i) to (u_{i+1}, w_{i+1}) in $\Omega(\mathcal{O}_2, \sigma_n)$ (with respect to the metric $d_{\mathcal{X} \times \mathcal{Y}}^{V_n}$). Thus $(u_i, w_i) - (u_{i+1}, w_{i+1}) \in B_0(\mathcal{O}_2, \kappa_2, \sigma_n)$ as required.

So the map from S to $H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n)$ is well defined. It is also surjective by construction. So $H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n)$ is a QS-group of $H_0(\mathcal{O}'_1, \mathcal{O}'_2, \kappa'_1, \kappa'_2, \sigma_n)$ and therefore, the 0-dimensional sofic homology of $\beta^\Gamma \times \nu$ is bounded by the 0-dimensional sofic homology of ν .

To finish the proof, we will show that the 0-dimensional sofic homology of $\beta^\Gamma \times \nu$ is greater than or equal to the 0-dimensional sofic homology of ν . So let \mathcal{O}'_2 be an arbitrary open neighborhood of ν in $\text{Prob}(\mathcal{Y}^\Gamma)$. Also let $\kappa'_2 > 0$. Let $\mathcal{O}_2 = \text{Prob}(\mathcal{X}^\Gamma) \times \mathcal{O}'_2$ and $\kappa_2 = \kappa'_2$. Let $\mathcal{O}_1 \subset \mathcal{O}_2$ be an arbitrary open neighborhood of $\beta^\Gamma \times \nu$. Let $0 < \kappa_1 \leq \kappa_2$. By Lemma 7.10, there exists an open neighborhood \mathcal{O}'_1 of ν such that $\mathcal{O}'_1 \subset \mathcal{O}'_2$ and for all but finitely many n , if $y \in \Omega(\mathcal{O}'_1, \sigma_n)$ then there exists $x \in \mathcal{X}^{V_n}$ such that $(x, y) \in \Omega(\mathcal{O}_1, \sigma_n)$. Set $\kappa'_1 = \kappa_1$.

Let S be the subgroup of $H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n)$ generated by all 0-cycles of the form $[x, y]$ with $y \in \Omega(\mathcal{O}'_1, \sigma_n)$ and $(x, y) \in \Omega(\mathcal{O}_1, \sigma_n)$. The map $[x, y] \rightarrow [y]$ from S to $H_0(\mathcal{O}'_1, \mathcal{O}'_2, \kappa'_1, \kappa'_2, \sigma_n)$ is well defined because if $[x_1, y_1] - [x_2, y_2] \in B_0(\mathcal{O}_2, \kappa_2, \sigma_n)$ then there is a κ_2 -path from $[x_1, y_1]$ to $[x_2, y_2]$ in $\Omega(\mathcal{O}_2, \sigma_n)$. The projection of this path to \mathcal{Y}^{V_n} is a κ'_2 -path from $[y_1]$ to $[y_2]$ in $\Omega(\mathcal{O}'_2, \sigma_n)$. So $[y_1]$ and $[y_2]$ represent the same element of $H_0(\mathcal{O}'_1, \mathcal{O}'_2, \kappa'_1, \kappa'_2, \sigma_n)$.

We claim the map $[x, y] \rightarrow [y]$ from S to $H_0(\mathcal{O}'_1, \mathcal{O}'_2, \kappa'_1, \kappa'_2, \sigma_n)$ is surjective. Suppose $[y] \in H_0(\mathcal{O}'_1, \mathcal{O}'_2, \kappa'_1, \kappa'_2, \sigma_n)$. Then $y \in \Omega(\mathcal{O}'_1, \sigma_n)$. By choice of \mathcal{O}_1 , this implies the existence of $x \in \mathcal{X}^{V_n}$ with $(x, y) \in \Omega(\mathcal{O}_1, \sigma_n)$. Thus $[x, y] \in S$. Since $[y]$ is arbitrary this implies the claimed surjectivity. So $H_0(\mathcal{O}'_1, \mathcal{O}'_2, \kappa'_1, \kappa'_2, \sigma_n)$ is a QS-group of $H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n)$ as required. \square

We will use a well-known concentration inequality on Hamming cubes to prove Proposition 7.9. First we need some notation. Let λ_s denote the probability measure on $\{0, 1\}$ given by $\lambda_s(1) = s$, $\lambda_s(0) = 1 - s$. Let $\mathbb{P}_s = \lambda_s^{V_n} \times \beta^{V_n}$ be a probability measure on $\{0, 1\}^{V_n} \times \mathcal{X}^{V_n}$. Let \mathbb{E}_s denote expectation with respect to \mathbb{P}_s .

PROPOSITION 7.11. *There exists a constant $C > 0$ (depending only on the diameter of \mathcal{X}) such that for any $\epsilon > 0$, any $s \in [0, 1]$ and any 1-Lipschitz function $F : \{0, 1\}^{V_n} \times \mathcal{X}^{V_n} \rightarrow \mathbb{R}$,*

$$\mathbb{P}_s \left\{ |F - \mathbb{E}[F]| \geq \epsilon \right\} \leq 2e^{-C\epsilon^2|V_n|}.$$

Here we are using the normalized Hamming metric $d_{\{0,1\} \times \mathcal{X}}^{V_n}$ given by

$$\begin{aligned} d_{\{0,1\} \times \mathcal{X}}^{V_n}((\chi_1, x_1), (\chi_2, x_2)) &= d_{\{0,1\}}^{V_n}(\chi_1, \chi_2) + d_{\mathcal{X}}^{V_n}(x_1, x_2) \\ &= |V_n|^{-1} \# \{v \in V_n : \chi_1(v) \neq \chi_2(v)\} \\ &\quad + |V_n|^{-1} \sum_{v \in V_n} d_{\mathcal{X}}(x_1(v), x_2(v)). \end{aligned}$$

For the proof see [27, Corollary 1.17].

Proposition 7.9 is proven by letting $x_i^{(j)}$ (for $j = 0, 1, \dots$) be the result of a coupled random walk on \mathcal{X}^{V_n} . To define this coupled random walk, for $x \in \mathcal{X}^{V_n}$ define

$$\xi^x : \{0, 1\}^{V_n} \times \mathcal{X}^{V_n} \longrightarrow \mathcal{X}^{V_n}$$

by

$$\xi^x(\chi, z)_v := \begin{cases} z_v & \text{if } \chi(v) = 1 \\ x_v & \text{if } \chi(v) = 0. \end{cases}$$

We will think of ξ^x as a random variable taking values in \mathcal{X}^{V_n} . More precisely, we choose a random subset of V_n with each vertex being chosen with probability s (independently). Then we randomize the value of x at each chosen vertex. This produces the new random element ξ^x .

LEMMA 7.12. *Fix notation as in Proposition 7.9. Then for any $s \in [0, 1]$,*

$$\liminf_{n \rightarrow \infty} \left\{ \mathbb{P}_s((\xi^x, y) \in \Omega(\mathcal{O}_2, \sigma_n)) : (x, y) \in \Omega(\mathcal{O}_1, \sigma_n) \right\} = 1.$$

Proof. Because \mathcal{O}_1 is hereditary, there is a finite hereditary subset $\mathcal{F} \subset C(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma)$ and $\delta > 0$ such that

$$\mathcal{O}_1 = \left\{ \mu' \in \text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma) : |\mu'(f) - \beta^\Gamma \times \nu(f)| < \delta \ \forall f \in \mathcal{F} \right\}.$$

Because \mathcal{O}_2 contains the closure of \mathcal{O}_1 there is a $\delta' > \delta$ such that

$$\mathcal{O}_2 \supset \{\mu' \in \text{Prob}(\mathcal{X}^\Gamma \times \mathcal{Y}^\Gamma) : |\mu'(f) - \beta^\Gamma \times \nu(f)| < \delta' \ \forall f \in \mathcal{F}\}.$$

Fix $(x, y) \in \Omega(\mathcal{O}_1, \sigma_n)$ and $f \in \mathcal{F}$. It suffices to obtain a lower bound on

$$\mathbb{P}_s \left(|P_{(\xi^x, y)}^{\sigma_n}(f) - \beta^\Gamma \times \nu(f)| < \delta' \right)$$

that tends to 1 as $n \rightarrow \infty$ but does not depend on (x, y) .

Let $D \subset \Gamma$ be a finite subset such that f is D -local. Let $W_n \subset V_n$ be the set of all vertices $v \in V_n$ such that the map

$$g \in D \longmapsto \sigma_n(g)^{-1}v$$

is injective. For any $v \in W_n$,

$$\mathbb{E}_s [f(\Pi_v^{\sigma_n}(\xi^x, y))] = \sum_{S \subset D} s^{|S|} (1-s)^{|D \setminus S|} \mathbb{E}_{\beta^\Gamma \times \nu} [f | \mathcal{B}_{\mathcal{X}^{D \setminus S}}] (\Pi_v^{\sigma_n}(x, y)).$$

To see this, let $\mathbf{S} : \{0, 1\}^{V_n} \rightarrow 2^D$ be the random subset of D defined by

$$\mathbf{S}(\chi) = \{g \in D : \chi(\sigma_n(g)^{-1}v) = 1\}.$$

Then for any $g \in D$, $g \in \mathbf{S}$ with probability s and these events are jointly independent over $g \in D$. Moreover, conditioned on $\mathbf{S}(\chi) = S$, the expected value of $f(\Pi_v^{\sigma_n}(\xi^x, y))$ is $\mathbb{E}_{\beta^\Gamma \times \nu} [f | \mathcal{B}_{\mathcal{X}^{D \setminus S}}] (\Pi_v^{\sigma_n}(x, y))$.

Because Σ is a sofic approximation, $\lim_{n \rightarrow \infty} |W_n|/|V_n| = 1$. So

$$\begin{aligned} \mathbb{E}_s \left[P_{(\xi^x, y)}^{\sigma_n}(f) \right] &= |V_n|^{-1} \sum_{v \in V_n} \mathbb{E}_s [f(\Pi_v^{\sigma_n}(\xi^x, y))] \\ &= |V_n|^{-1} \sum_{v \in V_n} \sum_{S \subset D} s^{|S|} (1-s)^{|D \setminus S|} \mathbb{E}_{\beta^\Gamma \times \nu} [f | \mathcal{B}_{\mathcal{X}^{D \setminus S}}] (\Pi_v^{\sigma_n}(x, y)) + O(|V_n|^{-1}) \\ &= \sum_{S \subset D} s^{|S|} (1-s)^{|D \setminus S|} |V_n|^{-1} \sum_{v \in V_n} \mathbb{E}_{\beta^\Gamma \times \nu} [f | \mathcal{B}_{\mathcal{X}^{D \setminus S}}] (\Pi_v^{\sigma_n}(x, y)) + O(|V_n|^{-1}) \\ &= \sum_{S \subset D} s^{|S|} (1-s)^{|D \setminus S|} P_{(x, y)}^{\sigma_n} (\mathbb{E}_{\beta^\Gamma \times \nu} [f | \mathcal{B}_{\mathcal{X}^{D \setminus S}}]) + O(|V_n|^{-1}). \end{aligned}$$

Because \mathcal{F} is hereditary, $\mathbb{E}_{\beta^\Gamma \times \nu} [f | \mathcal{B}_{\mathcal{X}^{D \setminus S}}] \in \mathcal{F}$. So

$$\left| P_{(x, y)}^{\sigma_n} (\mathbb{E}_{\beta^\Gamma \times \nu} [f | \mathcal{B}_{\mathcal{X}^{D \setminus S}}]) - \beta^\Gamma \times \nu(f) \right| < \delta$$

for all $S \subset D$. Thus

$$\left| \mathbb{E}_s \left[P_{(\xi^x, y)}^{\sigma_n}(f) \right] - \beta^\Gamma \times \nu(f) \right| < \delta + O(|V_n|^{-1}).$$

Let $F : \{0,1\}^{V_n} \times \mathcal{X}^{V_n} \rightarrow \mathbb{R}$ be the function

$$F(\chi, z) = P_{(\xi^x(\chi, z), y)}^{\sigma_n}(f).$$

To finish the proof, it suffices (by Proposition 7.11 and the previous inequality) to prove that F is C -Lipschitz for some constant $C > 0$ (that does not depend on n , x or y but may depend on other parameters). In fact,

$$\begin{aligned} & |F(\chi, z) - F(\chi', z')| \\ &= |V_n|^{-1} \left| \sum_{v \in V_n} f(\Pi_v^{\sigma_n}(\xi^x(\chi, z), y)) - f(\Pi_v^{\sigma_n}(\xi^x(\chi', z'), y)) \right| \\ &\leq |V_n|^{-1} \sum_{v \in V_n} |f(\Pi_v^{\sigma_n}(\xi^x(\chi, z), y)) - f(\Pi_v^{\sigma_n}(\xi^x(\chi', z'), y))| \\ &\leq |V_n|^{-1} \sum_{v \in V_n} d_{\mathcal{X} \times \mathcal{Y}}^D(\Pi_v^{\sigma_n}(\xi^x(\chi, z), y), \Pi_v^{\sigma_n}(\xi^x(\chi', z'), y)). \end{aligned}$$

The last inequality above occurs because f is 1-Lipschitz as a function from $\mathcal{X}^D \times \mathcal{Y}^D$ to \mathbb{R} .

For fixed $v \in V_n$,

$$\begin{aligned} & d_{\mathcal{X} \times \mathcal{Y}}^D(\Pi_v^{\sigma_n}(\xi^x(\chi, z), y), \Pi_v^{\sigma_n}(\xi^x(\chi', z'), y)) \\ &\leq |D|^{-1} \sum_{g \in D} d_{\mathcal{X}}(z(\sigma_n(g)^{-1}v), z'(\sigma_n(g)^{-1}v)) + \text{diam}(\mathcal{X}) \mathbf{1}_{\chi(\sigma_n(g)^{-1}v) \neq \chi'(\sigma_n(g)^{-1}v)}. \end{aligned}$$

Summing over all v , we obtain

$$\begin{aligned} |F(\chi, z) - F(\chi', z')| &\leq |V_n|^{-1} \sum_{v \in V_n} d_{\mathcal{X}}(z(v), z'(v)) + \text{diam}(\mathcal{X}) \mathbf{1}_{\chi_v \neq \chi'_v} \\ &= d_{\mathcal{X}}^{V_n}(z, z') + \text{diam}(\mathcal{X}) d_{\{0,1\}}^{V_n}(\chi, \chi'). \end{aligned}$$

So F is $\max(1, \text{diam}(\mathcal{X}))$ -Lipschitz. \square

Proof of Proposition 7.9. Fix notation as in the statement of Proposition 7.9. Let Leb denote Lebesgue measure on the unit interval $[0, 1]$. Let $\mathbb{P} = \text{Leb}^{V_n} \times \beta^{V_n}$ be the product measure on $[0, 1]^{V_n} \times \mathcal{X}^{V_n}$. Let \mathbb{E} denote expectation with respect to \mathbb{P} .

For $x \in \mathcal{X}^{V_n}$ and $s \in [0, 1]$, define

$$\zeta(x, s | \cdot, \cdot) : [0, 1]^{V_n} \times \mathcal{X}^{V_n} \longrightarrow \mathcal{X}^{V_n}$$

by

$$\zeta(x, s | \tau, z)_v := \begin{cases} z_v & \text{if } \tau_v \leq s \\ x_v & \text{if } \tau_v > s \end{cases}$$

The distribution of $\zeta(x, s | \cdot, \cdot)$ (with respect to $\text{Leb}^{V_n} \times \beta^{V_n}$) is the same as the distribution of ξ^x (with respect to \mathbb{P}_s). So Lemma 7.12 applies.

Fix a natural number $k > \text{diam}(\mathcal{X})/\kappa$. Let (τ, z) be a random variable with distribution $\text{Leb}^{V_n} \times \beta^{V_n}$. Then with high probability (whp) as $n \rightarrow \infty$ the following events occur:

- (1) For every $0 \leq j \leq k$,

$$\#\{v \in V_n : \tau_v \in (j/k, j/k + 1/k]\} < \kappa|V_n|/\text{diam}(\mathcal{X}).$$

- (2) $\zeta(x_i, j/k | \cdot, \cdot) \in \Omega(\mathcal{O}_2, \sigma_n)$ for all $1 \leq i \leq K$ and $0 \leq j \leq k$.

- (3) $\tau_v \neq 0$ for all $v \in V_n$.

This first condition holds whp by the law of large numbers, the second by Lemma 7.12, and the last occurs with probability 1.

So there is some $(\tau, z) \in [0, 1]^{V_n} \times \mathcal{X}^{V_n}$ such that all of the above conditions hold. Set $x_i^{(j)} = \zeta(x_i, j/k | \tau, z)$. This first 3 conclusions of Proposition 7.9 are immediate. The fourth occurs by definition of $\zeta(x_i, j/k | \tau, z)$. The fifth occurs with $M = k$. \square

LEMMA 7.13. $b_{d, \Sigma}(\nu) \leq (d+1)h_\Sigma(\nu)$ for all d, Σ, ν .

Proof. Let $\mathcal{O}_1 \subset \mathcal{O}_2 \subset \text{Prob}(\mathcal{Y}^\Gamma)$ be open neighborhoods of ν and $0 < \kappa_1 \leq \kappa_2/3$. Let $S_n \subset \Omega(\mathcal{O}_1, \sigma_n)$ be a subset whose κ_1 -neighborhood contains $\Omega(\mathcal{O}_1, \sigma_n)$ and

$$|S_n| = \text{cov}_{\kappa_1}(\Omega(\mathcal{O}_1, \sigma_n), d_{\mathcal{Y}}^{V_n}).$$

So there exists a map $\Psi : \Omega(\mathcal{O}_1, \sigma_n) \rightarrow S_n$ such that $d_{\mathcal{Y}}^{V_n}(\Psi(y), y) < \kappa_1$ for all y . If $y \in \mathcal{Y}^{V_n} \setminus \Omega(\mathcal{O}_1, \sigma_n)$, then define $\Psi(y) = y$ so that now we can consider Ψ as a map from \mathcal{Y}^{V_n} to itself. By Lemma 6.1, $\Psi_*(z) - z \in B_d(\mathcal{O}_2, \kappa_2, \sigma_n)$ for any $z \in Z_d^L(\mathcal{O}_1, \kappa_1, \sigma_n)$. This uses $\mathcal{O}_1 \subset \mathcal{O}_2$ and $3\kappa_1 \leq \kappa_2$. It follows that

$$\dim_{\mathbb{Q}} (H_d^L(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n) \otimes_{\mathbb{Z}} \mathbb{Q}) \leq \#S_n^{d+1} = \text{cov}_{\kappa_1}(\Omega(\mathcal{O}_1, \sigma_n), d_{\mathcal{Y}}^{V_n})^{d+1}.$$

The lemma now follows from the definitions of $b_{d, \Sigma}(\nu)$ and $h_\Sigma(\nu)$. \square

COROLLARY 7.14. *If $\Gamma \curvearrowright (X, \mu)$ has the Weak Pinsker Property then $b_{0, \Sigma}(\mu) = 0$.*

Proof. Let $\epsilon > 0$. Then $\Gamma \curvearrowright (X, \mu)$ is isomorphic to the direct product of a Bernoulli shift and an action with entropy $< \epsilon$. By Theorem 7.8, the 0-dimensional

sofic homology of $\Gamma \curvearrowright (X, \mu)$ is equivalent to the 0-dimensional sofic homology of an action with entropy $< \epsilon$. By Lemma 7.13, this shows $b_{0, \Sigma}(\mu) < \epsilon$. \square

8. An action without the Weak Pinsker Property. This section proves Theorem 1.2. Here is an outline:

- §8.1 shows that two models of random graphs, the configuration model and the permutation model, are closely related, allowing the transfer of results about one to the other. This is useful here because results in the literature are generally proven for the configuration model, but it is the permutation model that gives actions of the free group.
- §8.2 is about first moment computations of the numbers of independent sets and pairs of independent sets of a random regular graph.
- §8.3 explains two models of random pairs (G, I) where G is a regular graph and $I \subset V$ is an independent set. These models are called the planted and the uniform model. Usually it is easy to estimate probabilities with respect to the planted model but not with respect to the uniform model. However, sofic entropy is more closely related to the uniform model. Fortunately, there is an inequality relating the two models.
- §8.4 has the main technical result bounding the size of clusters of independent sets in a random regular graph.
- §8.5 provides a general result for obtaining an invariant measure μ on \mathcal{X}^Γ whose model spaces $\Omega(\mathcal{O}, \sigma_n)$ have large intersections with fixed subsets $\mathcal{W}_n \subset \mathcal{X}^{V_n}$. This is applied later with \mathcal{W}_n equal to the set of “good” independent subsets to obtain the invariant measure in Theorem 1.2.
- §8.6 finishes the proof of Theorem 1.2.

8.1. The configuration model and the permutation model. The proof of Theorem 8.11 is made simpler by borrowing results about independent sets on the configuration model of random regular graphs and transferring them to the permutation model. This section explains the two models and two key theorems linking them together.

Definition 13 (The configuration model). Let $d \geq 3, n \geq 1$ be integers such that dn is even. Let π be a uniformly random perfect matching on $[n] \times [d]$. Let $G_{\text{conf}}(\pi)$ be the random multi-graph with vertex set $[n]$ such that the number of edges from i to j equals the number of edges between $\{i\} \times [d]$ and $\{j\} \times [d]$ in the matching π . This is called the *configuration model* [7, 24]. It gives a random d -regular multi-graph on n vertices. Let $\mathbb{P}_{d,n}^{\text{conf}}$ denote the law of $G_{\text{conf}}(\pi)$ and let $\mathbb{E}_{d,n}^{\text{conf}}$ be its expectation operator. So $\mathbb{P}_{d,n}^{\text{conf}}$ is a probability measure on $\text{Graphs}(d, n)$, which is the set of all d -regular multi-graphs on $[n]$.

Definition 14 (The permutation model). Let $\Gamma = \mathbb{F}_r = \langle a_1, \dots, a_r \rangle$ be the rank r free group. Let $\mathbb{P}_{r,n}^{\text{perm}}$ be the uniform probability measure on the set

$\text{Hom}(\mathbb{F}_r, \text{sym}(n))$ of homomorphisms from \mathbb{F}_r to $\text{sym}(n)$. Also let $\mathbb{E}_{r,n}^{\text{perm}}$ be its expectation operator. For $\sigma \in \text{Hom}(\mathbb{F}_r, \text{sym}(n))$, let $G(\sigma)$ be the multi-graph with vertex set $[n]$ and edges $\{v, \sigma(a_i)v\}$ (over $v \in [n]$, $1 \leq i \leq r$). If σ is random with law $\mathbb{P}_{r,n}^{\text{perm}}$ then $G(\sigma)$ is called the *permutation model*. The law of $G(\sigma)$ is a probability measure on $\text{Graphs}(2r, n)$, which by abuse of notation, we will also denote by $\mathbb{P}_{r,n}^{\text{perm}}$.

One of the main results of [20] is:

THEOREM 8.1. *Let $A_n \subset \text{Graphs}(d, n)$ be any sequence of subsets. Suppose $d \geq 4$ is even and let $2r = d$. Then $\lim_n \mathbb{P}_{d,n}^{\text{conf}}(A_n) = 1$ if and only if $\lim_n \mathbb{P}_{r,n}^{\text{perm}}(A_n) = 1$. Equivalently, the permutation and configuration models model are contiguous.*

COROLLARY 8.2. *Let σ_n be a random homomorphism from \mathbb{F}_r to $\text{sym}(n)$ with law $\mathbb{P}_{r,n}^{\text{perm}}$. Then for every finite $D \subset \mathbb{F}_r$ and $\delta > 0$,*

$$\lim_{n \rightarrow \infty} \mathbb{P}_{r,n}^{\text{perm}}(\sigma_n \text{ is } (D, \delta)\text{-sofic}) = 1.$$

Proof. It suffices to show that for any nontrivial $w \in \mathbb{F}_r$ the number of $v \in \{1, \dots, n\}$ such that $\sigma_n(w)v = v$ is $o(n)$ with high probability as $n \rightarrow \infty$. To phrase this a different way, it suffices to prove that for any $L > 0$ the number of simple closed cycles of length $\leq L$ in $G(\sigma_n)$ is $o(n)$ with high probability as $n \rightarrow \infty$. This statement is proven in [6] for the configuration model. Since the two models are contiguous, it also holds for the permutation model. \square

8.1.1. Expectations. The next result shows that the first moment method applied to counting vertex-labelings of either the configuration or permutation model results in the same calculation up to subexponential factors. To explain further we need some notation.

Definition 15. (Admissible pairs) Let \mathcal{X} be a finite set. A pair of vectors $(\pi^{\text{vert}}, \pi^{\text{edge}}) \in \text{Prob}(\mathcal{X}) \times \text{Prob}(\mathcal{X} \times \mathcal{X}) \subset \mathbb{R}^{\mathcal{X}} \times \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is *admissible* if both π^{vert} and π^{edge} are probability vectors, π^{edge} is symmetric in the sense that $\pi^{\text{edge}}(p, q) = \pi^{\text{edge}}(q, p)$ for all p, q and both marginals of π^{edge} equal π^{vert} . The latter condition means for every $p \in \mathcal{X}$,

$$\pi^{\text{vert}}(p) = \sum_{q \in \mathcal{X}} \pi^{\text{edge}}(p, q) = \sum_{q \in \mathcal{X}} \pi^{\text{edge}}(q, p).$$

Let $\mathcal{A}(\mathcal{X}) \subset \mathbb{R}^{\mathcal{X}} \times \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ be the compact space of all admissible pairs.

Example 2. Suppose $G = (V, E)$ is a finite d -regular multi-graph and $x : V \rightarrow \mathcal{X}$ a map. Choose a vertex v and a directed edge e independently and uniformly at random. Then the distributions of $x(v)$ and $x(e)$ form a pair of admissible vectors. Let $(\pi_x^{\text{vert}}, \pi_x^{\text{edge}}) \in \mathcal{A}(\mathcal{X})$ denote this pair of distributions.

Definition 16 ($\Omega(\pi^{\text{vert}}, \pi^{\text{edge}}; G)$). Given a finite graph $G = (V, E)$ and an admissible pair of vectors $(\pi^{\text{vert}}, \pi^{\text{edge}}) \in \mathcal{A}(\mathcal{X})$, let $\Omega(\pi^{\text{vert}}, \pi^{\text{edge}}; G)$ be the set of all vertex-labelings $x : V \rightarrow \mathcal{X}$ such that

$$\pi^{\text{vert}} = \pi_x^{\text{vert}}, \quad \pi^{\text{edge}} = \pi_x^{\text{edge}}.$$

Write $\#\Omega(\pi^{\text{vert}}, \pi^{\text{edge}})$ for the random variable

$$\mathbf{G} \longmapsto \#\Omega(\pi^{\text{vert}}, \pi^{\text{edge}}; \mathbf{G})$$

where \mathbf{G} is either a random sample of the configuration or the permutation model, depending on context.

Definition 17. For $x > 0$, let $\eta(x) = -x \log(x)$. Extend this by continuity so that $\eta(0) = 0$. Given a vector $\vec{p} = (p_1, \dots, p_k)$ of nonnegative real numbers, let

$$H(\vec{p}) = H(p_1, \dots, p_k) = \sum_{i=1}^k \eta(p_i)$$

be the *Shannon entropy* of \vec{p} . For example, it is well known that if \vec{p} is a probability vector then the associated multinomial coefficients satisfy

$$\lim_{n \rightarrow \infty} n^{-1} \log \binom{n}{m_{n,1}, \dots, m_{n,k}} = H(p_1, \dots, p_k)$$

where $m_{n,1}, \dots, m_{n,k}$ are any choice of nonnegative integers satisfying $m_{n,1} + \dots + m_{n,k} = n$ and $\lim_{n \rightarrow \infty} m_{n,i}/n = p_i$ for all i .

LEMMA 8.3. *Let $(\pi^{\text{vert}}, \pi^{\text{edge}}) \in \mathcal{A}(\mathcal{X})$ be an admissible pair of vectors. If $d \geq 2$, $n \geq 1$ are integers, dn is even, π^{vert} takes values in $\frac{1}{n}\mathbb{Z}$ and π^{edge} takes values in $\frac{2}{dn}\mathbb{Z}$ then*

$$n^{-1} \log \mathbb{E}_{d,n}^{\text{conf}} [\#\Omega(\pi^{\text{vert}}, \pi^{\text{edge}})] = (d/2)H(\pi^{\text{edge}}) - (d-1)H(\pi^{\text{vert}}) + o_n(1)$$

and for any integer $r \geq 1$,

$$n^{-1} \log \mathbb{E}_{r,n}^{\text{perm}} [\#\Omega(\pi^{\text{vert}}, \pi^{\text{edge}})] = rH(\pi^{\text{edge}}) - (2r-1)H(\pi^{\text{vert}}) + o_n(1).$$

Proof. The case of the permutation model is handled in [8]. It seems likely that the configuration model result has been known for some time because it is essentially the “first moment method” which has been a standard tool in this area of probabilistic combinatorics for decades. However, I have been unable to find a suitable reference.

The proof will reduce to [4, Lemma 4.1] once we verify that $\mathbb{E}_{d,n}^{\text{conf}} [\#\Omega(\pi^{\text{vert}}, \pi^{\text{edge}})]$ is positive.

Because π^{vert} takes values in $\frac{1}{n}\mathbb{Z}$, there exists a map $x : [n] \rightarrow \mathcal{X}$ such that $\pi_x^{\text{vert}} = \pi^{\text{vert}}$. Let $\tilde{x} : [n] \times [d] \rightarrow \mathcal{X}$ be the lift defined by $\tilde{x}(v, i) = x(v)$. Because

π^{edge} takes values in $\frac{2}{dn}\mathbb{Z}$, there exists some perfect matching μ on $[n] \times [d]$ such that the distribution on pairs of vertex-labels from x assigned to the endpoints of a uniformly random oriented edge of μ is π^{edge} . Here are some details to justify this claim. First $\pi^{\text{vert}} \in (\frac{2}{dn}\mathbb{Z})^{\mathcal{X}}$ because π^{edge} takes values in $\frac{2}{dn}\mathbb{Z}$ and π^{vert} is the first marginal of π^{edge} . So there exists a partition

$$[n] \times [d] = P \sqcup Q$$

such that

$$|P \cap \tilde{x}^{-1}(p)| = |Q \cap \tilde{x}^{-1}(p)| = \pi^{\text{vert}}(p)nd/2$$

for all $p \in \mathcal{X}$. For every $p \in \mathcal{X}$, choose partitions

$$P \cap \tilde{x}^{-1}(p) = \sqcup_{q \in \mathcal{X}} P_{p,q}, \quad Q \cap \tilde{x}^{-1}(p) = \sqcup_{q \in \mathcal{X}} Q_{q,p}$$

such that

$$\#P_{p,q} = \#Q_{q,p} = \pi^{\text{edge}}(p, q)nd/2.$$

Finally let μ be any perfect matching of $[n] \times [d]$ that restricts to a perfect matching from $P_{p,q}$ to $Q_{q,p}$ for all $p, q \in \mathcal{X}$. This matching satisfies the claim.

By [4, Lemma 4.1], the probability that a uniformly random perfect matching of $[n] \times [d]$ has edge distribution π^{edge} is

$$e^{(dn/2)H(\pi^{\text{edge}})} e^{-dnH(\pi^{\text{vert}})}$$

(up to a multiplicative factor that is subexponential in n). To see this, replace n in [4, Lemma 4.1] with dn , μ with π^{vert} and ν with π^{edge} . Since the number of ways to choose x is $\exp(nH(\pi^{\text{vert}}))$ (up to a multiplicative factor that is subexponential in n), this implies the result. \square

THEOREM 8.4. *Fix an integer $d \geq 2$. For $n \in \mathbb{N}$ with dn even, suppose $\mathcal{K}_n \subset \mathcal{A}(\mathcal{X})$ satisfies the following: if $(\pi^{\text{vert}}, \pi^{\text{edge}}) \in \mathcal{K}_n$ then π^{vert} takes values in $\frac{1}{n}\mathbb{Z}$ and π^{edge} takes values in $\frac{2}{dn}\mathbb{Z}$. Suppose that \mathcal{K}_n converges to a closed subspace \mathcal{M} in the Hausdorff topology (on the space of all closed subsets of $\mathcal{A}(\mathcal{X})$) as $n \rightarrow \infty$. Then*

$$\begin{aligned} \lim_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{d,n}^{\text{conf}} [\# \{x \in \mathcal{X}^n : (\pi_x^{\text{vert}}, \pi_x^{\text{edge}}) \in \mathcal{K}_n\}] \\ = \max \{(d/2)H(\pi^{\text{edge}}) - (d-1)H(\pi^{\text{vert}}) : (\pi^{\text{vert}}, \pi^{\text{edge}}) \in \mathcal{M}\}. \end{aligned}$$

Moreover, if $d = 2r$ for some integer $r \geq 1$ then the quantity above also equals

$$\lim_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r,n}^{\text{perm}} [\# \{x \in \mathcal{X}^n : (\pi_x^{\text{vert}}, \pi_x^{\text{edge}}) \in \mathcal{K}_n\}].$$

Proof. This follows from the previous Lemma since the function

$$(\pi^{\text{vert}}, \pi^{\text{edge}}) \longmapsto (d/2)H(\pi^{\text{edge}}) - (d-1)H(\pi^{\text{vert}})$$

is uniformly continuous and there are at most a polynomial (in n) number of admissible pairs $(\pi^{\text{vert}}, \pi^{\text{edge}}) \in \mathcal{K}_n$ such that π^{edge} takes values in $\frac{2}{dn}\mathbb{Z}$. In fact this number is bounded by $(dn+1)^{|\mathcal{X}|^2}$. \square

8.2. First moment computations. Let $\Gamma = \mathbb{F}_r = \langle a_1, \dots, a_r \rangle$ denote the free group of rank $r \geq 2$. For $s, s', t \in [0, 1]$ and $\sigma \in \text{Hom}(\mathbb{F}_r, \text{sym}(n))$, let $\mathcal{I}_{s,s',t}(\sigma)$ be the set of all of pairs (W, W') of independent subsets of $G(\sigma)$ satisfying

$$(3) \quad \#W = \lfloor sn \rfloor, \quad \#W' = \lfloor s'n \rfloor, \quad \#(W \cap W') = \lfloor tn \rfloor.$$

In the arguments to follow, $\#\mathcal{I}_{s,s',t}$ is regarded as a random variable with respect to the permutation model. Let $f(r, s, s', t)$ denote the upper exponential growth rate of the expected value of $\#\mathcal{I}_{s,s',t}$.

$$f(r, s, s', t) := \limsup_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r,n}^{\text{perm}} [\#\mathcal{I}_{s,s',t}].$$

THEOREM 8.5. *The function f is uniformly continuous. Moreover, the limit exists so that*

$$f(r, s, s', t) := \lim_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r,n}^{\text{perm}} [\#\mathcal{I}_{s,s',t}].$$

Proof. Let $\mathcal{X} = \{(0,0), (0,1), (1,0), (1,1)\}$. Let \mathcal{K}_n be the set of all admissible pairs $(\pi^{\text{vert}}, \pi^{\text{edge}})$ such that π^{vert} takes values in $\frac{1}{n}\mathbb{Z}$, π^{edge} takes values in $\frac{2}{dn}\mathbb{Z}$ ($d = 2r$) and these linear equations are satisfied:

$$\begin{aligned} \pi^{\text{vert}}(1,0) + \pi^{\text{vert}}(1,1) &= \lfloor sn \rfloor / n \\ \pi^{\text{vert}}(0,1) + \pi^{\text{vert}}(1,1) &= \lfloor s'n \rfloor / n \\ \pi^{\text{vert}}(1,1) &= \lfloor tn \rfloor / n \\ \pi^{\text{edge}}((i_1, j_1), (i_2, j_2)) &= 0 \text{ if either } (i_1 = i_2 = 1) \text{ or } (j_1 = j_2 = 1). \end{aligned}$$

A pair (W, W') of independent subsets of $G(\sigma)$ satisfies (3) if and only if $(\pi_x^{\text{vert}}, \pi_x^{\text{edge}}) \in \mathcal{K}_n$ where $x = (1_W, 1_{W'})$.

Let $\mathcal{M}(r, s, s', t)$ be the set of all admissible pairs $(\pi^{\text{vert}}, \pi^{\text{edge}})$ satisfying these linear equations:

$$\begin{aligned} \pi^{\text{vert}}(1,0) + \pi^{\text{vert}}(1,1) &= s \\ \pi^{\text{vert}}(0,1) + \pi^{\text{vert}}(1,1) &= s' \\ \pi^{\text{vert}}(1,1) &= t \\ \pi^{\text{edge}}((i_1, j_1), (i_2, j_2)) &= 0 \text{ if either } (i_1 = i_2 = 1) \text{ or } (j_1 = j_2 = 1). \end{aligned}$$

Because \mathcal{K}_n converges to $\mathcal{M}(r, s, s', t)$ in the Hausdorff topology, Theorem 8.4 implies

$$(4) \quad f(r, s, s', t) = \max_{(\pi^{\text{vert}}, \pi^{\text{edge}}) \in \mathcal{M}(r, s, s', t)} -(2r-1)H(\pi^{\text{vert}}) + rH(\pi^{\text{edge}})$$

$$(5) \quad = \lim_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r, n}^{\text{perm}} [\#\mathcal{I}_{s, s', t}].$$

Continuity of f now follows from continuity of

$$(\pi^{\text{vert}}, \pi^{\text{edge}}) \mapsto -(2r-1)H(\pi^{\text{vert}}) + rH(\pi^{\text{edge}})$$

and

$$(r, s, s', t) \mapsto \mathcal{M}(r, s, s', t)$$

where the latter is with respect to the Hausdorff topology on closed subsets of the space of admissible vector pairs. Uniform continuity follows from continuity and compactness. \square

We are most interested in the special case in which s' is close to s . In this case we will use the continuity result to reduce the study of this function to the special case $s = s'$. To simplify notation, let $f(r, s, t) := f(r, s, s, t)$.

For $\sigma \in \text{Hom}(\mathbb{F}_r, \text{sym}(n))$, let $\mathcal{I}_s(\sigma)$ be the collection of independent sets W of the graph $G(\sigma)$ such that $\#W = \lfloor sn \rfloor$. Also let

$$f(r, s) := \limsup_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r, n}^{\text{perm}} [\#\mathcal{I}_s]$$

be the upper exponential growth rate of the expected number of independent sets of cardinality close to sn .

THEOREM 8.6. *Fix $0 \leq \bar{t} \leq \bar{s} \leq 1$. If $t := t(r) := \bar{t} \frac{\log(2r)}{r}$ and $s := s(r) := \bar{s} \frac{\log(2r)}{r}$ then*

$$\begin{aligned} f(r, s) &= \eta(s) - rs^2 + O(\log(r)/r) \\ f(r, s, t) &= \eta(t) + 2\eta(s-t) + r[t^2 - 2s^2] + O(\log(r)/r) \end{aligned}$$

where the error term implicit in the big $O(\cdot)$ notation does not depend on \bar{s}, \bar{t} .

Proof. We keep notation as in the previous proof. Observe that $f(r, s, s) = f(r, s, s, s)$ counts pairs of identical independent sets. So $f(r, s, s) = f(r, s)$. So it suffices to obtain the estimate for $f(r, s, t)$. The calculation is similar to one in Gamarnik-Sudan [19].

By (4),

$$f(r, s, t) = \max_{(\pi^{\text{vert}}, \pi^{\text{edge}}) \in \mathcal{M}(r, s, s, t)} -(2r-1)H(\pi^{\text{vert}}) + rH(\pi^{\text{edge}}).$$

Let $\pi = (\pi^{\text{vert}}, \pi^{\text{edge}}) \in \mathcal{M}(r, s, s, t)$. We claim:

$$\begin{aligned}\pi^{\text{vert}}(0, 0) &= 1 - 2s + t & \pi^{\text{vert}}(0, 1) &= s - t \\ \pi^{\text{vert}}(1, 0) &= s - t & \pi^{\text{vert}}(1, 1) &= t.\end{aligned}$$

By definition, $\pi^{\text{vert}}(1, 1) = t$. Since $\pi^{\text{vert}}(1, 0) + \pi^{\text{vert}}(1, 1) = s$ this determines the value of $\pi^{\text{vert}}(1, 0)$. The other two values hold by symmetry and the fact that π^{vert} is a probability measure. So

$$H(\pi^{\text{vert}}) = \eta(t) + 2\eta(s - t) + \eta(1 - 2s + t).$$

Because Shannon entropy $H(\cdot)$ is strictly concave and $\mathcal{M}(r, s, s, t)$ is convex, it follows that there is a unique $\pi_0 \in \mathcal{M}(r, s, s, t)$ such that

$$f(r, s, t) = -(2r - 1)H(\pi_0^{\text{vert}}) + rH(\pi_0^{\text{edge}}).$$

The space $\mathcal{X} \times \mathcal{X}$ is invariant under these two symmetries:

$$\begin{aligned}((a, b), (c, d)) &\longmapsto ((c, d), (a, b)), \\ ((a, b), (c, d)) &\longmapsto ((b, a), (d, c)).\end{aligned}$$

These generate a group of order 4. This induces a group of symmetries on $\text{Prob}(\mathcal{X} \times \mathcal{X})$. Moreover, if a measure $\nu \in \text{Prob}(\mathcal{X} \times \mathcal{X})$ has both marginals equal to π^{vert} , then each of its images under this group also has marginals equal to π^{vert} (for any $\pi \in \mathcal{M}(r, s, s, t)$). Moreover, the Shannon entropy is preserved under the action of this group. Because $\pi_0^{\text{edge}} \in \text{Prob}(\mathcal{X} \times \mathcal{X})$ is unique, it is necessarily invariant under the action of this group. So if $x_0 = \pi_0^{\text{edge}}((0, 0), (0, 1))$ then

$$x_0 = \pi_0^{\text{edge}}((0, 0), (1, 0)) = \pi_0^{\text{edge}}((0, 1), (0, 0)) = \pi_0^{\text{edge}}((1, 0), (0, 0)).$$

Since both marginals of π^{edge} equal π^{vert} , this implies

$$\begin{aligned}\pi_0^{\text{edge}}((0, 0), (0, 0)) &= 1 - 2s - 2x_0, \\ \pi_0^{\text{edge}}((0, 0), (1, 1)) &= \pi_0^{\text{edge}}((1, 1), (0, 0)) = t \\ \pi_0^{\text{edge}}((1, 0), (0, 1)) &= \pi_0^{\text{edge}}((0, 1), (1, 0)) = s - t - x_0.\end{aligned}$$

So $H(\pi_0^{\text{edge}}) = G(x_0)$ where

$$G(x) = 2\eta(t) + 4\eta(x) + 2\eta(s - t - x) + \eta(1 - 2s - 2x).$$

Use $\frac{d}{dx}\eta(x) = -1 - \log(x)$, to obtain:

$$\begin{aligned}\frac{\partial G(x)}{\partial x} &= \frac{\partial}{\partial x}(4\eta(x) + 2\eta(s-t-x) + \eta(1-2s-2x)) \\ &= 2\log\left(\frac{(s-t-x)(1-2s-2x)}{x^2}\right).\end{aligned}$$

Since π_0 is maximizing, $\frac{\partial G(x)}{\partial x}(x_0) = 0$. So $(s-t-x_0)(1-2s-2x_0) = x_0^2$.

To simplify notation, we let $x = x_0$ from now on. By collecting terms with a factor of x , we obtain

$$\begin{aligned}H(\pi_0^{\text{edge}}) &= 2\eta(t) + 4\eta(x) + 2\eta(s-t-x) + \eta(1-2s-2x) \\ &= 2\eta(t) - 4x\log(x) - 2(s-t-x)\log(s-t-x) \\ &\quad - (1-2s-2x)\log(1-2s-2x) \\ &= 2\eta(t) + 2x\log\left(\frac{(s-t-x)(1-2s-2x)}{x^2}\right) \\ &\quad - 2(s-t)\log(s-t-x) - (1-2s)\log(1-2s-2x) \\ &= 2\eta(t) - 2(s-t)\log(s-t-x) - (1-2s)\log(1-2s-2x).\end{aligned}$$

Use the above and collect terms with a factor of r to obtain

$$\begin{aligned}f(r, s, t) &= -(2r-1)[\eta(t) + 2\eta(s-t) + \eta(1-2s+t)] \\ &\quad + r[2\eta(t) - 2(s-t)\log(s-t-x) - (1-2s)\log(1-2s-2x)] \\ &= \eta(t) + 2\eta(s-t) - (2r-1)\eta(1-2s+t) \\ &\quad + r[-4\eta(s-t) - 2(s-t)\log(s-t-x) - (1-2s)\log(1-2s-2x)]\end{aligned}$$

Next, we estimate x . Since $(s-t-x)(1-2s-2x) = x^2$,

$$x^2 - (1-2t)x + (s-t)(1-2s) = 0.$$

Let $\Delta = s-t$. Since $x \leq s-t$,

$$\begin{aligned}2x &= 1-2t - \sqrt{(1-2t)^2 - 4(s-t)(1-2s)} \\ &= 1-2s + 2\Delta - \sqrt{(1-2s+2\Delta)^2 - 4\Delta(1-2s)} \\ &= 1-2s + 2\Delta - \sqrt{(1-2s)^2 + 4\Delta^2}.\end{aligned}$$

By Taylor series expansion, if $C > 0$ is a constant then $\sqrt{C^2 + \epsilon} = C + \frac{\epsilon}{2C} + O(\epsilon^2)$. Use this with $C = 1-2s$, $\epsilon = 4\Delta^2$ to obtain

$$(6) \quad 2x = 2\Delta - \frac{2\Delta^2}{(1-2s)} + O(\Delta^4).$$

Thus

$$\begin{aligned}\log(s-t-x) &= \log\left(\frac{(s-t)^2}{1-2s} + O(\Delta^4)\right) \\ &= 2\log(s-t) + 2s + O((s-t)^2) \\ &= 2\log(s-t) + 2s + O(s^2).\end{aligned}$$

Thus

$$\begin{aligned}-4\eta(s-t) - 2(s-t)\log(s-t-x) &= 2(s-t)(2\log(s-t) - \log(s-t-x)) \\ &= -4(s-t)s + O(s^3).\end{aligned}$$

Plugging this back in to the equation for f gives

$$\begin{aligned}f(r,s,t) &= \eta(t) + 2\eta(s-t) - (2r-1)\eta(1-2s+t) \\ &\quad + r[-4(s-t)s - (1-2s)\log(1-2s-2x)] + O(rs^3).\end{aligned}$$

We use the Taylor series estimates

$$\begin{aligned}\eta(1-2s+t) &= 2s-t - (2s-t)^2/2 + O(s^3) \\ \log(1-2s-2x) &= -(2s+2x) - (2s+2x)^2/2 + O(s^3) \\ &= -2[s+x+(s+x)^2] + O(s^3) \\ &= -2(1+s+x)(s+x) + O(s^3)\end{aligned}$$

to obtain

$$\begin{aligned}f(r,s,t) &= \eta(t) + 2\eta(s-t) - (2r-1)[2s-t - (2s-t)^2/2] \\ &\quad - 4rs(s-t) + 2r(1-2s)(1+s+x)(s+x) + O(rs^3).\end{aligned}$$

If r is large enough then $\log^2(r) > r$ which implies $rs^3 \leq s$. So we can replace the $O(rs^3)$ with $O(s)$. This allows us to replace the coefficient $(2r-1)$ on the bracketed term with $2r$. We can also simplify the last term by

$$2r(1-2s)(1+s+x)(s+x) = 2r(1-s+x)(s+x) + O(s)$$

to obtain

$$\begin{aligned}f(r,s,t) &= \eta(t) + 2\eta(s-t) - 4rs(s-t) - 2r[2s-t - (2s-t)^2/2] \\ &\quad + 2r(1-s+x)(s+x) + O(s).\end{aligned}$$

By factoring out a $2r$ we obtain

$$\begin{aligned}f(r,s,t) &= \eta(t) + 2\eta(s-t) \\ &\quad + 2r[-2s(s-t) - 2s+t + (2s-t)^2/2 + (1-s+x)(s+x)] + O(s).\end{aligned}$$

After multiplying out and collecting terms, the quantity in brackets simplifies to

$$[-s - s^2 + t + t^2/2 + x + x^2].$$

By (6), this simplifies to $[-s^2 + t^2/2 + O(s^3)]$. So

$$f(r, s, t) = \eta(t) + 2\eta(s-t) + r[t^2 - 2s^2] + O(s).$$

Since $s = O(\log(r)/r)$, this implies the theorem. \square

COROLLARY 8.7. *Keep notation as in Theorem 8.6. In addition, assume $2/3 < \bar{s} < 1$. Then*

$$\begin{aligned} & \max \{f(r, s, t) - f(r, s) : s/2 \leq t \leq s\} \\ &= f(r, s) - \bar{s}(1 - \bar{s}) \frac{\log^2(r)}{r} + O\left(\frac{\log(r) \log \log(r)}{r}\right). \end{aligned}$$

Proof. By Theorem 8.6,

$$f(r, s, t) - f(r, s) = f(r, s) + \eta(t) + 2\eta(s-t) - 2\eta(s) + rt^2 + O\left(\frac{\log(r)}{r}\right).$$

Since $s = \bar{s} \frac{\log(2r)}{r}$,

$$\eta(s) = -s \log(s) = \bar{s} \frac{\log^2(r)}{r} + O\left(\frac{\log(r) \log \log(r)}{r}\right).$$

Similar estimates hold for $\eta(t)$ and $\eta(s-t)$ and the constant implicit in the $O(\cdot)$ notation is uniform over \bar{s}, \bar{t} . Therefore,

$$\eta(t) + 2\eta(s-t) - 2\eta(s) = -\bar{t} \frac{\log^2(r)}{r} + O\left(\frac{\log(r) \log \log(r)}{r}\right).$$

Since $rt^2 = \bar{t}^2 \frac{\log^2(r)}{r}$, this implies

$$f(r, s, t) - f(r, s) = f(r, s) - (\bar{t} - \bar{t}^2) \frac{\log^2(r)}{r} + O\left(\frac{\log(r) \log \log(r)}{r}\right).$$

The minimum value of $(\bar{t} - \bar{t}^2)$ for $\bar{t} \in [\bar{s}/2, \bar{s}]$ is attained when $\bar{t} = \bar{s}$. This is because $x \mapsto x - x^2$ is concave, so the minimum is achieved at either $\bar{s}/2$ or \bar{s} . But since $\bar{s} \in (2/3, 1)$, $(\bar{s} - \bar{s}^2) < (\bar{s}/2 - \bar{s}^2/4)$. Substituting \bar{s} for \bar{t} finishes the corollary. \square

8.3. The planted model and the uniform model. There are two related models of random independent sets on random regular graphs that we need to consider to prove Theorem 8.11. To explain, recall that for $\sigma \in \text{Hom}(\mathbb{F}_r, \text{sym}(n))$, $\mathcal{I}_s(\sigma)$ is the collection of independent sets W of the graph $G(\sigma)$ such that $\#W = \lfloor s\#V \rfloor$. The *uniform model* is the probability measure $\mathbb{P}_{r,s,n}^{\text{unif}}$ on $\text{Hom}(\mathbb{F}_r, \text{sym}(n)) \times 2^n$ defined by:

$$\mathbb{P}_{r,s,n}^{\text{unif}}(\sigma, W) = \frac{1}{\#\text{Hom}(\mathbb{F}_r, \text{sym}(n)) \times \#\mathcal{I}_s(\sigma)}$$

if $W \in \mathcal{I}_s(\sigma)$ and $\mathbb{P}_{r,s,n}^{\text{unif}}(\sigma, W) = 0$ otherwise. A random sample (σ, \mathbf{W}) with law $\mathbb{P}_{r,s,n}^{\text{unif}}$ can be obtained by first choosing $\sigma \in \text{Hom}(\mathbb{F}_r, \text{sym}(n))$ uniformly at random and then choosing $\mathbf{W} \in \mathcal{I}_s(\sigma)$ uniformly at random.

The *planted model* is the probability measure $\mathbb{P}_{r,s,n}^{\text{plant}}$ on $\text{Hom}(\mathbb{F}_r, \text{sym}(n)) \times 2^n$ that is uniformly distributed on pairs (σ, W) such that $W \in \mathcal{I}_s(\sigma)$. Thus

$$\mathbb{P}_{r,s,n}^{\text{plant}}(\sigma, W) = \frac{1}{\#\text{Hom}(\mathbb{F}_r, \text{sym}(n)) \times \mathbb{E}_{r,n}^{\text{perm}}[\#\mathcal{I}_s(\sigma)]}$$

if $W \in \mathcal{I}_s(\sigma)$ and 0 otherwise. Let $\mathbb{E}_{r,s,n}^{\text{unif}}$ and $\mathbb{E}_{r,s,n}^{\text{plant}}$ be the corresponding expectation operators.

It is relatively easy to compute probabilities with respect to the planted model instead of the uniform model. However, to prove Theorem 8.11 we need to work with the uniform model. The next result forms a bridge between the two models.

THEOREM 8.8. *Fix $\bar{s} \in (0, 1)$ and set $s = \bar{s} \frac{\log(2r)}{r}$. Let $R \geq r$ be an integer satisfying*

$$\liminf_n \mathbb{P}_{R,n}^{\text{perm}}[\#\mathcal{I}_s > 0] = 1.$$

Let

$$X_{R,r,s,n} = \left\{ \sigma \in \text{Hom}(\mathbb{F}_r, \text{sym}(n)) : \#\mathcal{I}_s(\sigma) \geq \frac{\mathbb{E}_{r,n}^{\text{perm}}[\#\mathcal{I}_s]}{2\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s]} \right\}.$$

Then

$$\liminf_{n \rightarrow \infty} \mathbb{P}_{r,n}^{\text{perm}}(X_{R,r,s,n}) = 1.$$

Moreover, if $A \subset X_{R,r,s,n} \times 2^n$ then

$$\mathbb{P}_{r,s,n}^{\text{unif}}(A) \leq \mathbb{P}_{r,s,n}^{\text{plant}}(A) \times 2\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s].$$

Remark 6. The proof below is modeled after an analogous result for sparse Erdős-Renyi graphs obtained in [12].

Proof of Theorem 8.8. Let $\mathbb{F}_R = \langle a_1, \dots, a_R \rangle$ and $\mathbb{F}_r = \langle a_1, \dots, a_r \rangle$. Thus \mathbb{F}_r is a subgroup of \mathbb{F}_R . Let σ be a sample of $\mathbb{P}_{R,n}^{\text{perm}}$ (so σ is a uniformly random homomorphism from \mathbb{F}_R to $\text{sym}(n)$). Let $\sigma \upharpoonright \mathbb{F}_r$ denote the restriction of σ to \mathbb{F}_r .

Any independent set of $G(\sigma)$ is automatically an independent set of $G(\sigma \upharpoonright \mathbb{F}_r)$ since the latter is a subgraph of the former. If $W \subset [n]$ is an independent set for $G(\sigma \upharpoonright \mathbb{F}_r)$ then the probability that W is also an independent set for $G(\sigma)$ depends only on $|W|, (R-r)$ and n . It is easy to derive an exact expression for this probability, but we do not need it.

Let Y_n be the set of all $\sigma \in \text{Hom}(\mathbb{F}_R, \text{sym}(n))$ such that the restriction $\sigma \upharpoonright \mathbb{F}_r \notin X_{R,r,s,n}$. Since $\mathbb{P}_{R,n}^{\text{perm}}(Y_n) = 1 - \mathbb{P}_{r,n}^{\text{perm}}(X_{R,r,s,n})$, it suffices to prove $\mathbb{P}_{R,n}^{\text{perm}}(Y_n) \rightarrow 0$ as $n \rightarrow \infty$. Since the probability that an independent subset W of $G(\sigma \upharpoonright \mathbb{F}_r)$ is an independent subset of $G(\sigma)$ does not depend on $\sigma \upharpoonright \mathbb{F}_r$,

$$\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s|Y_n] = \mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s(\sigma \upharpoonright \mathbb{F}_r)|Y_n] \frac{\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s]}{\mathbb{E}_{r,n}^{\text{perm}}[\#\mathcal{I}_s]}.$$

By definition of Y_n and $X_{R,r,s,n}$,

$$\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s(\sigma \upharpoonright \mathbb{F}_r)|Y_n] < \frac{\mathbb{E}_{r,n}^{\text{perm}}[\#\mathcal{I}_s]}{2\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s]}.$$

Combine this with the equality above to obtain

$$(7) \quad \mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s|Y_n] < \frac{1}{2}.$$

By Markov's inequality,

$$\mathbb{P}_{R,n}^{\text{perm}}(\#\mathcal{I}_s \geq 2\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s|Y_n]|Y_n) \leq \frac{1}{2}.$$

Multiply both sides of the inequality above by -1 and add 1 to obtain

$$1/2 \leq \mathbb{P}_{R,n}^{\text{perm}}(\#\mathcal{I}_s < 2\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s|Y_n]|Y_n) \leq \mathbb{P}_{R,n}^{\text{perm}}(\#\mathcal{I}_s < 1|Y_n).$$

The second inequality above follows from (7). Since $\mathbb{P}_{R,n}^{\text{perm}}(\#\mathcal{I}_s < 1|Y_n) \leq \frac{\mathbb{P}_{R,n}^{\text{perm}}(\#\mathcal{I}_s < 1)}{\mathbb{P}_{R,n}^{\text{perm}}(Y_n)}$, multiply denominators in the inequality $1/2 \leq \frac{\mathbb{P}_{R,n}^{\text{perm}}(\#\mathcal{I}_s < 1)}{\mathbb{P}_{R,n}^{\text{perm}}(Y_n)}$ to obtain

$$\mathbb{P}_{R,n}^{\text{perm}}(Y_n) \leq 2\mathbb{P}_{R,n}^{\text{perm}}(\#\mathcal{I}_s < 1).$$

However, $\mathbb{P}_{R,n}^{\text{perm}}(\#\mathcal{I}_s < 1)$ tends to zero as $n \rightarrow \infty$ by assumption. This shows $\mathbb{P}_{R,n}^{\text{perm}}(Y_n) \rightarrow 0$ and therefore $\mathbb{P}_{r,n}^{\text{perm}}(X_{R,r,s,n}) \rightarrow 1$ as $n \rightarrow \infty$.

To verify the last statement, let $(\sigma, W) \in X_{R,r,s,n} \times 2^n$ be such that $W \in \mathcal{I}_s(\sigma)$. Then

$$\begin{aligned}\mathbb{P}_{r,s,n}^{\text{unif}}(\sigma, W) &= \frac{1}{\#\mathcal{I}_s(\sigma)\#\text{Hom}(\mathbb{F}_r, \text{sym}(n))} \leq \frac{2\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s]}{\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s]\#\text{Hom}(\mathbb{F}_r, \text{sym}(n))} \\ &= \mathbb{P}_{r,s,n}^{\text{plant}}(\sigma, W) \times 2\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s].\end{aligned}$$

These inequalities are justified in turn by the definition of $\mathbb{P}_{r,s,n}^{\text{unif}}$, the assumption that $\sigma \in X_{R,r,s,n}$ and the definition of $\mathbb{P}_{r,s,n}^{\text{plant}}$. \square

Define $\text{MAXIND} : \text{Hom}(\mathbb{F}_r, \text{sym}(n)) \rightarrow \mathbb{R}$ by

$$\text{MAXIND}(\sigma) = n^{-1} \max \{ \#W : W \text{ is an independent subset of } G(\sigma) \}.$$

PROPOSITION 8.9. *For all r sufficiently large, there is a constant $\alpha(r)$ such that*

$$\inf_{\delta > 0} \lim_{n \rightarrow \infty} \mathbb{P}_{r,n}^{\text{perm}} \left(\text{MAXIND} \in (\alpha(r) - \delta, \alpha(r) + \delta) \right) = 1.$$

Moreover,

$$\alpha(r) = \frac{\log(r)}{r} + O(\log \log(r)/r).$$

Remark 7. The limit was proven to exist in [5]. The asymptotic statement follows from earlier results of Frieze-Luczak [18] in the case of the configuration model. Since the two models are contiguous [20] this implies the proposition. An exact formula for $\alpha(r)$ (for sufficiently large r) was recently obtained by Ding-Sly-Sun [13].

COROLLARY 8.10. *Fix $\bar{s} \in (0, 1)$ and set $s = \bar{s} \frac{\log(2r)}{r}$. Let $\sigma_n \sim \mathbb{P}_{r,n}^{\text{perm}}$. Then*

$$n^{-1} \log \#\mathcal{I}_s(\sigma_n) \geq f(r, s) - o_n(1) - O\left(\frac{\log(r) \log \log(r)}{r}\right)$$

with probability tending to 1 as $n \rightarrow \infty$. To be precise, this means that there exists $\delta(r) = O\left(\frac{\log(r) \log \log(r)}{r}\right)$ such that for every $\epsilon > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P}_{r,n}^{\text{perm}} \left(n^{-1} \log \#\mathcal{I}_s(\sigma_n) \geq f(r, s) - \epsilon - \delta(r) \right) = 1.$$

Proof. Let R be the largest integer such that $\alpha(R) > s$. By Proposition 8.9,

$$s < \alpha(R) = \frac{\log(R)}{R} + O(\log \log(R)/R).$$

Observe that the function $x \mapsto \frac{\log(x)}{x}$ is monotone increasing for all $x \geq e$. So if $x \geq e$ then

$$\begin{aligned} 0 \leq \frac{\log(x)}{x} - \frac{\log(x+1)}{x+1} &= \frac{\log(x)}{x} - \frac{\log(x) + \log(1+1/x)}{x+1} \\ &= \frac{\log(x)}{x(x+1)} - \frac{\log(1+1/x)}{x+1} \leq \frac{\log(x)}{x^2}. \end{aligned}$$

Since $\frac{\log(x)}{x^2} \leq \frac{\log \log(x)}{x}$, this implies

$$\alpha(R) - \alpha(R+1) = O(\log \log(R)/R)$$

and therefore we can improve the previous inequality to an equality: $s = \frac{\log(R)}{R} + O(\log \log(R)/R)$.

By Theorem 8.6,

$$f(R, s) = \eta(s) - s^2 R + O(\log(R)/R).$$

Since $s = \frac{\log(R)}{R} + O(\log \log(R)/R)$,

$$\begin{aligned} \eta(s) &= \frac{\log^2(R)}{R} + O\left(\frac{\log(R) \log \log(R)}{R}\right) \\ s^2 R &= \frac{\log^2(R)}{R} + O\left(\frac{\log(R) \log \log(R)}{R}\right). \end{aligned}$$

So $f(R, s) = O\left(\frac{\log(R) \log \log(R)}{R}\right)$. Since $x \mapsto \frac{\log(x) \log \log(x)}{x}$ is monotone decreasing (for all large enough x), this implies

$$(8) \quad f(R, s) = O\left(\frac{\log(r) \log \log(r)}{r}\right).$$

Theorem 8.8 implies that if $\sigma_n \sim \mathbb{P}_{r,n}^{\text{perm}}$ then with probability tending to 1 as $n \rightarrow \infty$,

$$\begin{aligned} n^{-1} \log \#\mathcal{I}_s(\sigma_n) &\geq n^{-1} \log \mathbb{E}_{r,n}^{\text{perm}}[\#\mathcal{I}_s] - n^{-1} \log(2\mathbb{E}_{R,n}^{\text{perm}}[\#\mathcal{I}_s]) \\ &= f(r, s) - f(R, s) - o_n(1) \\ &= f(r, s) - o_n(1) - O\left(\frac{\log(r) \log \log(r)}{r}\right). \end{aligned} \quad \square$$

8.4. Bounding clusters of independent sets. For $(\sigma, W) \in \text{Hom}(\mathbb{F}_r, \text{sym}(n)) \times 2^n$, let $\text{Cl}_{s,\epsilon}(\sigma, W) \subset 2^n$ be the collection of independent subsets W' of the graph $G(\sigma)$ such that $|\#W'/n - s| \leq \epsilon$ and $|W \cap W'| \geq (s/2)n$. Informally, $\text{Cl}_{s,\epsilon}(\sigma, W)$ represents the cluster containing W in $\bigcup_{s' \in (s-\epsilon, s+\epsilon)} \mathcal{I}_{s'}(\sigma)$.

THEOREM 8.11. *Let $\bar{s} \in (\frac{2+\sqrt{2}}{4}, 1)$ and $s := \bar{s} \frac{\log(2r)}{r}$. Then there exists $\gamma = \gamma(r, \bar{s}) > 0$, $\epsilon = \epsilon(r, \bar{s}) > 0$ and $0 < b_1 < b_2 < s$, $b_2 > s/2$, such that the following holds. Let (σ_n, \mathbf{W}_n) be random with law $\mathbb{P}_{r,s,n}^{\text{unif}}$. Then with probability tending to 1 as $n \rightarrow \infty$,*

- (1) *there does not exist an independent set W' of $G(\sigma_n)$ such that $\#W'/n \in (s - \epsilon, s + \epsilon)$ and $b_1 n \leq |\mathbf{W}_n \cap W'| \leq b_2 n$,*
- (2) *$n^{-1} \log \# \text{Cl}_{s,\epsilon}(\sigma_n, \mathbf{W}_n) \leq f(r, s) - \gamma$.*

Proof. The first claim was proven by Gamarnik-Sudan [19, Theorem 2.6] with the configuration model in place of the permutation model. This uses the hypothesis $\bar{s} > \frac{2+\sqrt{2}}{4}$. Since the two models are contiguous by [20], this implies the first claim. The proof of the second claim given below is modeled after the proof of an analogous result for sparse Erdős-Renyi graphs in [12].

For $\sigma \in \text{Hom}(\mathbb{F}_r, \text{sym}(n))$, let $\mathcal{I}_{s;\epsilon}(\sigma)$ be the set of all pairs (W, W') satisfying:

- (1) both W, W' are independent subsets of G_σ ,
- (2) $\#W = \lfloor sn \rfloor$, $|\#W'/n - s| \leq \epsilon$,
- (3) $|W \cap W'| \geq (s/2)n$.

The definition of the planted model implies:

$$(9) \quad \mathbb{E}_{r,s,n}^{\text{plant}} [\# \text{Cl}_{s,\epsilon}] = \frac{\mathbb{E}_{r,n}^{\text{perm}} [\# \mathcal{I}_{s;\epsilon}]}{\mathbb{E}_{r,n}^{\text{perm}} [\# \mathcal{I}_s]}.$$

By Theorem 8.4 and (9),

$$\begin{aligned} \lim_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r,s,n}^{\text{plant}} [\# \text{Cl}_{s,\epsilon}] \\ = \max \{ f(r, s, s', t) - f(r, s) : s' \in [s - \epsilon, s + \epsilon], s/2 \leq t \leq s \}. \end{aligned}$$

By Theorem 8.5, f is uniformly continuous. So there exists $\epsilon = \epsilon(r, \bar{s}) > 0$ such that

$$\max_{s' \in [s - \epsilon, s + \epsilon], s/2 \leq t \leq s} f(r, s, s', t) - \max_{s/2 \leq t \leq s} f(r, s, t) \leq \frac{\log(r) \log \log(r)}{r}.$$

Assume from now on that $|s - s'| < \epsilon$. Then

$$\begin{aligned} \lim_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r,s,n}^{\text{plant}} [\# \text{Cl}_{s,\epsilon}] &= \max \{ f(r, s, t) - f(r, s) : s/2 \leq t \leq s \} \\ &\quad + O \left(\frac{\log(r) \log \log(r)}{r} \right). \end{aligned}$$

By Corollary 8.7,

$$(10) \quad \begin{aligned} \lim_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r,s,n}^{\text{plant}} [\# \text{Cl}_{s,\epsilon}] &= f(r, s) - \bar{s}(1 - \bar{s}) \frac{\log^2(r)}{r} \\ &\quad + O \left(\frac{\log(r) \log \log(r)}{r} \right). \end{aligned}$$

Next these estimates are transferred from the planted model to the uniform model. To simplify notation, let $X_n = X_{R,r,s,n}$ be as in Theorem 8.8 where R is the largest integer such that $\alpha(R) \geq s$. By Theorem 8.8,

$$\begin{aligned} \mathbb{E}_{r,s,n}^{\text{unif}}[\#\text{Cl}_{s,\epsilon}|X_n] &= \mathbb{P}_{r,n}^{\text{perm}}(X_n)^{-1} \mathbb{E}_{r,s,n}^{\text{unif}}[\#\text{Cl}_{s,\epsilon} 1_{X_n}] \\ &\leq \mathbb{P}_{r,n}^{\text{perm}}(X_n)^{-1} \mathbb{E}_{r,s,n}^{\text{plant}}[\#\text{Cl}_{s,\epsilon} 1_{X_n}] \times 2 \mathbb{E}_{R,n}^{\text{perm}}(\#\mathcal{I}_s) \\ &\leq \mathbb{P}_{r,n}^{\text{perm}}(X_n)^{-1} \mathbb{E}_{r,s,n}^{\text{plant}}[\#\text{Cl}_{s,\epsilon}] \times 2 \mathbb{E}_{R,n}^{\text{perm}}(\#\mathcal{I}_s). \end{aligned}$$

Recall that $f(R,s) = \lim_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{R,n}^{\text{perm}}(\#\mathcal{I}_s)$. Combined with (10), the fact that $\mathbb{P}_{r,n}^{\text{perm}}(X_n) \rightarrow 1$ as $n \rightarrow \infty$ and equation (8), this implies

$$\begin{aligned} \limsup_{n \rightarrow \infty} n^{-1} \log \mathbb{E}_{r,s,n}^{\text{unif}}[\#\text{Cl}_{s,\epsilon}|X_n] &\leq f(r,s) - \bar{s}(1-\bar{s}) \frac{\log^2(r)}{r} \\ (11) \quad &\quad + O\left(\frac{\log(r) \log \log(r)}{r}\right). \end{aligned}$$

By Markov's inequality, this implies the theorem with any choice of $\gamma = \gamma(r, \bar{s})$ satisfying

$$0 < \gamma(r, \bar{s}) < \bar{s}(1-\bar{s}) \frac{\log^2(r)}{r} + O\left(\frac{\log(r) \log \log(r)}{r}\right). \quad \square$$

8.5. A variational principle. The next result will be used in the proof of Theorem 1.2 to obtain an invariant measure that is, in some sense, a subsequential limit of the uniform models. Its proof uses the same ideas that are behind Kerr-Li's proof of the Variational Principle for sofic entropy [26].

PROPOSITION 8.12. *Let $\Sigma = \{\sigma_n\}_{n=1}^\infty$ be a sofic approximation to Γ , \mathcal{X} a finite set and for each $n \in \mathbb{N}$, let $\mathcal{W}_n \subset \mathcal{X}^{V_n}$ be given. Then there exists an invariant measure $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$ satisfying*

$$h_\Sigma(\mu) \geq \inf_{\mathcal{O} \ni \mu} \limsup_{n \rightarrow \infty} |V_n|^{-1} \log \#(\Omega(\mathcal{O}, \sigma_n) \cap \mathcal{W}_n) = \limsup_{n \rightarrow \infty} |V_n|^{-1} \log \#\mathcal{W}_n.$$

Proof. The inequality is trivial, so it suffices to prove the equality.

The space $\text{Prob}(\mathcal{X}^\Gamma)$ is compact and metrizable in the weak* topology. So fix a metric on $\text{Prob}(\mathcal{X}^\Gamma)$ with diameter ≤ 1 . Choose a sequence $\mathcal{U}^{(j)}$ ($j = 1, 2, \dots$) of finite open covers of $\text{Prob}_\Gamma(\mathcal{X}^\Gamma) \subset \text{Prob}(\mathcal{X}^\Gamma)$ by open balls of radius 2^{-j} . For convenience, set $\mathcal{U}^{(0)} = \{\text{Prob}(\mathcal{X}^\Gamma)\}$.

We will inductively construct a sequence $\{\mathcal{O}^{(j)}\}_{j=0}^\infty$ of open sets $\mathcal{O}^{(j)} \in \mathcal{U}^{(j)}$ and positive constants $\{m^{(j)}\}_{j=1}^\infty$ satisfying: for every $J \in \mathbb{N}$,

- (1) $\mathcal{O}^{(j)} \cap \mathcal{O}^{(j+1)} \neq \emptyset$ for all $0 \leq j < J$,
- (2) for every $0 \leq j \leq J$, there exist infinitely many n such that

$$\#(\mathcal{W}_n \cap \Omega(\mathcal{O}^{(j)}, \sigma_n)) \geq \#\mathcal{W}_n / m^{(j)}.$$

For the base case, set $\mathcal{O}^{(0)} = \text{Prob}(\mathcal{X}^\Gamma)$ and $m^{(0)} = 1$. For induction, suppose there is some $J \geq 1$, open subsets $\mathcal{O}^{(1)}, \dots, \mathcal{O}^{(J)}$ and constants $\{m^{(j)}\}_{j=1}^J$ satisfying the criteria above.

For each n , let $\mathcal{O}_n^{(J+1)} \in \mathcal{U}^{(J+1)}$ be an open set such that $\mathcal{O}_n^{(J+1)} \cap \mathcal{O}^{(J)} \neq \emptyset$ and

$$\#(\mathcal{W}_n \cap \Omega(\mathcal{O}_n^{(J+1)}, \sigma_n)) \geq \#(\mathcal{W}_n \cap \Omega(\mathcal{O}', \sigma_n))$$

for all $\mathcal{O}' \in \mathcal{U}^{(J+1)}$ such that $\mathcal{O}' \cap \mathcal{O}^{(J)} \neq \emptyset$. Let $\mathcal{O}^{(J+1)} \in \mathcal{U}^{(J+1)}$ be an open subset such that there exists infinitely many n satisfying:

$$\begin{aligned} \mathcal{O}^{(J+1)} &= \mathcal{O}_n^{(J+1)} \\ \#(\mathcal{W}_n \cap \Omega(\mathcal{O}^{(J)}, \sigma_n)) &\geq \#\mathcal{W}_n / m^{(J)}. \end{aligned}$$

Let

$$m^{(J+1)} := m^{(J)} |\mathcal{U}^{(J+1)}|.$$

Since the number of open sets $\mathcal{O}' \in \mathcal{U}^{(J+1)}$ that intersect $\mathcal{O}^{(J)}$ nontrivially is at most $|\mathcal{U}^{(J+1)}|$ and

$$\Omega(\mathcal{O}^{(J)}, \sigma_n) \subset \bigcup \left\{ \Omega(\mathcal{O}', \sigma_n) : \mathcal{O}' \cap \mathcal{O}^{(J)} \neq \emptyset \right\},$$

it follows that for infinitely many n ,

$$\begin{aligned} \#(\mathcal{W}_n \cap \Omega(\mathcal{O}^{(J+1)}, \sigma_n)) &\geq |\mathcal{U}^{(J+1)}|^{-1} \# \left(\mathcal{W}_n \cap \Omega(\mathcal{O}^{(J)}, \sigma_n) \right) \\ &\geq \frac{1}{m^{(J)} |\mathcal{U}^{(J+1)}|} \#\mathcal{W}_n. \end{aligned}$$

This proves the inductive step and the claim.

Since each $\mathcal{U}^{(j)}$ is a covering by balls of radius 2^{-j} , if $\mu_j \in \mathcal{O}^{(j)}$ is arbitrary, then $\{\mu_j\}_j$ is a Cauchy sequence. Let $\mu = \lim_{j \rightarrow \infty} \mu_j$. If $\mathcal{O} \subset \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$ is any open subset containing μ then \mathcal{O} contains $\mathcal{O}^{(j)}$ for some j . Thus

$$\begin{aligned} \limsup_{n \rightarrow \infty} |V_n|^{-1} \log \#\mathcal{W}_n &\geq \inf_{\mathcal{O} \ni \mu} \limsup_{n \rightarrow \infty} |V_n|^{-1} \log \#(\Omega(\mathcal{O}, \sigma_n) \cap \mathcal{W}_n) \\ &\geq \limsup_{n \rightarrow \infty} |V_n|^{-1} \log \#(\Omega(\mathcal{O}^{(j)}, \sigma_n) \cap \mathcal{W}_n) \\ &\geq \limsup_{n \rightarrow \infty} |V_n|^{-1} \log \frac{\#\mathcal{W}_n}{m^{(j)}} = \limsup_{n \rightarrow \infty} |V_n|^{-1} \log \#\mathcal{W}_n. \end{aligned}$$

This proves the equality. \square

8.6. Proof of Theorem 1.2.

Proof of Theorem 1.2. Choose constants $\bar{s}, s, r, b_1, b_2, \epsilon, \gamma$ satisfying Theorem 8.11. By choosing $\epsilon > 0$ smaller if necessary it may be assumed that $\epsilon/5 < b_2 - b_1$ and $(\epsilon/10)\log(2) + H(\epsilon/10, 1 - \epsilon/10) < \gamma/6$ where $H(a, b) = -a\log(a) - b\log(b)$ for any $a, b > 0$.

Given a homomorphism $\sigma : \mathbb{F}_r \rightarrow \text{sym}(n)$, let $\mathcal{W}(\sigma) \subset \mathcal{I}_s(\sigma)$ be the collection of independent subsets W satisfying

(1) there does not exist an independent set W' of $G(\sigma_n)$ such that $\#W'/n \in (s - \epsilon, s + \epsilon)$ and $b_1n \leq |W \cap W'| \leq b_2n$,

(2)

$$(12) \quad n^{-1} \log(\#\text{Cl}_{s, \epsilon}(\sigma_n, W)) \leq f(r, s) - \gamma/2.$$

An independent set W of $G(\sigma_n)$ is identified with its indicator function $1_W \in \{0, 1\}^n$. So by abuse of notation one may consider $\mathcal{W}(\sigma)$ to be a subset of $\{0, 1\}^n$.

By Corollaries 8.2, 8.10 and Theorem 8.11 there exists a sofic approximation $\Sigma = \{\sigma_n\}_{n \in \mathbb{N}}$ with $\sigma_n : \mathbb{F}_r \rightarrow \text{sym}(n)$ such that if $\mathcal{W}_n := \mathcal{W}(\sigma_n)$ then

$$\limsup_{n \rightarrow \infty} n^{-1} \log(\#\mathcal{W}_n) \geq f(r, s).$$

By Proposition 8.12, there exists an invariant measure $\mu \in \{0, 1\}^{\mathbb{F}_r}$ such that

$$h_\Sigma(\mu) \geq \inf_{\mathcal{O} \ni \mu} \limsup_{n \rightarrow \infty} n^{-1} \log \#(\Omega(\mathcal{O}, \sigma_n) \cap \mathcal{W}_n) = \limsup_{n \rightarrow \infty} n^{-1} \log \#\mathcal{W}_n \geq f(r, s).$$

For $g \in \mathbb{F}_r$, let $\pi_g : \{0, 1\}^{\mathbb{F}_r} \rightarrow \{0, 1\}$ be the coordinate projection. Let \mathcal{O}_2 be the set of measures $\nu \in \text{Prob}(\{0, 1\}^{\mathbb{F}_r})$ satisfying

(1) $\nu(\pi_e = 1) \in (s - \epsilon/2, s + \epsilon/2)$

(2) $\nu(\text{either } (\pi_e, \pi_{a_i}) = (1, 1) \text{ or } (\pi_e, \pi_{a_i^{-1}}) = (1, 1) \text{ for some } i) < \epsilon/10$.

Let $0 < \kappa_2$ be a constant with $b_2 - \kappa_2 - \epsilon/5 > b_1$.

It suffices to show that if $0 < \kappa_1 \leq \kappa_2$ is any constant and $\mathcal{O}_1 \subset \text{Prob}(\{0, 1\}^{\mathbb{F}_r})$ is any open neighborhood of μ with $\mathcal{O}_1 \subset \mathcal{O}_2$ then

$$\limsup_n n^{-1} \log \dim_{\mathbb{Q}}(H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n) \otimes_{\mathbb{Z}} \mathbb{Q}) \geq \gamma/3.$$

In the notation above, the superscript L is omitted because all 0-cycles are finite sums of length one 0-cycles. So the parameter L is irrelevant to studying 0-dimensional homology.

For $x \in \Omega(\mathcal{O}_1, \sigma_n) \cap \mathcal{W}_n$, let $C(x)$ be the set of all $y \in \Omega(\mathcal{O}_1, \sigma_n)$ such that

$$x - y \in B_0(\mathcal{O}_2, \kappa_2, \sigma_n).$$

In other words, $y \in C(x)$ if and only if there exists a path $x = x_0, x_1, \dots, x_k = y$ such that $x_i \in \Omega(\mathcal{O}_2, \sigma_n)$ for all i and $d(x_i, x_{i+1}) < \kappa_2$. Observe that $\Omega(\mathcal{O}_1, \sigma_n) \cap \mathcal{W}_n$ is the disjoint union sets of the form $C(x) \cap \mathcal{W}_n$.

We will estimate the cardinality $\#C(x)$ by showing that $C(x)$ is contained in the $\epsilon/10$ -neighborhood of $\text{Cl}_{s,\epsilon}(\sigma_n, W_1)$. So suppose $y \in C(x)$ and let $x_0, x_1, \dots, x_k = y$ be a path as above. Because $x_i \in \Omega(\mathcal{O}_2, \sigma_n)$, its empirical measure is in \mathcal{O}_2 , which implies that the “bad” set

$$\begin{aligned} & B(x_i) \\ & := \{v \in [n] : x(v) = 1 \text{ and either } x(\sigma(a_i)v) = 1 \text{ or } x(s(a_i^{-1})v) = 1 \text{ for some } i\} \end{aligned}$$

has cardinality at most $(\epsilon/10)n$. Let $W_i = x_i^{-1}(1) \setminus B(x_i)$. Then W_i is an independent subset of $G(\sigma_n)$ with density $|W_i|/n \in (s - \epsilon, s + \epsilon)$. Observe that $x^{-1}(1) = W_1$ since $x \in \mathcal{W}_n$ is an independent set.

Let d_{Hamm} denote the normalized Hamming metric on subsets of $[n]$. So

$$d_{\text{Hamm}}(W, W') = |W \Delta W'|/n.$$

Then

$$d_{\text{Hamm}}(W_i, W_{i+1}) \leq d(x_i, x_{i+1}) + |B(x_i)|/n + |B(x_{i+1})|/n < \kappa_2 + \epsilon/5.$$

for all i .

Let $\text{Cl}(x)$ be the collection of all independent sets W of $G(\sigma_n)$ such that $|W|/n \in (s - \epsilon, s + \epsilon)$ and $|W \cap W_1|/n \geq b_2$. I claim that $W_k \in \text{Cl}(x)$. If not, then there exists a smallest number $j \geq 2$ such that $W_j \notin \text{Cl}(x)$. Since $W_{j-1} \in \text{Cl}(x)$, $|W_{j-1} \cap W_1|/n \geq b_2$. Thus

$$|W_j \cap W_1|/n \geq |W_{j-1} \cap W_1|/n - d_{\text{Hamm}}(W_{j-1}, W_j) > b_2 - \kappa_2 - \epsilon/5 > b_1.$$

Since $x \in \mathcal{W}_n$, the definition of \mathcal{W}_n implies $|W_j \cap W_1|/n \geq b_2$ which implies $W_j \in \text{Cl}(x)$. This contradiction proves the claim.

Since $W_k \in \text{Cl}(x) \subset \text{Cl}_{s,\epsilon}(\sigma_n, W_1)$ and $|y^{-1}(1) \Delta W_k| \leq (\epsilon/10)n$, it follows that $C(x)$ is contained in the $\epsilon/10$ -neighborhood of $\text{Cl}_{s,\epsilon}(\sigma_n, W_1)$. Thus

$$\#C(x) \leq 2^{(\epsilon/10)n} \binom{n}{\lceil \epsilon n/10 \rceil} \# \text{Cl}_{s,\epsilon}(\sigma_n, W_1).$$

Combine the previous inequality with (12) to obtain

$$\begin{aligned} n^{-1} \log \#C(x) & \leq (\epsilon/10) \log(2) + H(\epsilon/10, 1 - \epsilon/10) + f(r, s) - \gamma/2 + o_n(1) \\ & < f(r, s) - \gamma/3 + o_n(1). \end{aligned}$$

Since $\Omega(\mathcal{O}_1, \sigma_n) \cap \mathcal{W}_n$ is a disjoint union of sets of the form $C(x) \cap \mathcal{W}_n$ and

$$\limsup_{n \rightarrow \infty} n^{-1} \log \#(\Omega(\mathcal{O}_1, \sigma_n) \cap \mathcal{W}_n) \geq f(r, s),$$

it follows that the number of different subsets of the form $C(x)$ for $x \in \Omega(\mathcal{O}_1, \sigma_n) \cap \mathcal{W}_n$ is at least $e^{(\gamma/3)n}$ up to subexponential factors. However each subset of the form

$C(x)$ contributes a dimension to the homology group $H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n) \otimes_{\mathbb{Z}} \mathbb{Q}$. Thus

$$\limsup_{n \rightarrow \infty} n^{-1} \log \dim_{\mathbb{Q}} (H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n) \otimes_{\mathbb{Z}} \mathbb{Q}) \geq \gamma/3 > 0. \quad \square$$

9. Questions.

(1) Does there exist an action whose sofic homology does not vanish in dimension 1 or in some higher dimension? Does there exist such an example which is a Markov chain over a free group?

(2) How does the d -dimensional sofic homology change under standard operations or perturbations of group actions, such as taking a direct product, passing to a subgroup of the acting group, coinducing from an action of a subgroup, ergodic decomposition, direct limits, inverse limits, taking a weak* limit of invariant measures or a d -bar limit?

(3) Is $b_{\Sigma,0}(\mu)$ a continuous or semi-continuous function of $\mu \in \text{Prob}(\mathcal{X}^\Gamma)$ if \mathcal{X} is finite? Given a positive number $0 < t < \log |\mathcal{X}|$ does there exist an invariant measure $\mu \in \text{Prob}_\Gamma(\mathcal{X}^\Gamma)$ with $b_{\Sigma,0}(\mu) = t$?

(4) It is well known that sofic entropy can increase under a factor map. To correct for this, several authors have defined the sofic entropy of a factor relative to the source. This notion was inspired by Kerr's approach to sofic entropy in [25] and has been variously called outer sofic entropy, extension entropy or entropy in the presence [28, 22, 23, 34, 36]. It seems likely that there should be an analogous definition of outer sofic homology. If so, this might be useful for defining relative sofic homology.

(5) Let \mathbb{K} be a compact abelian group and Γ a countable group. Then \mathbb{K} is identified with the subgroup of constants of \mathbb{K}^Γ . This subgroup is Γ -invariant and therefore $\mathbb{K}^\Gamma/\mathbb{K}$ is a compact abelian group on which Γ acts by automorphisms. The action $\Gamma \curvearrowright \mathbb{K}^\Gamma/\mathbb{K}$ is called a *Popa factor*. In [2] Tim Austin proved that if Γ has property (T), is residually finite and $\mathbb{K} = \mathbb{R}/\mathbb{Z}$ then there is a sofic approximation Σ relative to which the Popa factor does not have connected model spaces. This means that its 0-dimensional sofic homology is not trivial. In spite of this, I conjecture that $b_{\Sigma,0}(\Gamma \curvearrowright \mathbb{K}^\Gamma/\mathbb{K}) = 0$ for any sofic approximation Σ . In fact, my guess is that the sofic homology groups satisfy a bound of the form:

$$\dim_{\mathbb{Q}} (H_0(\mathcal{O}_1, \mathcal{O}_2, \kappa_1, \kappa_2, \sigma_n) \otimes_{\mathbb{Z}} \mathbb{Q}) \leq B$$

where B does not depend on n (but is allowed to depend on everything else).

(6) I conjecture that any strongly ergodic distal action $\Gamma \curvearrowright (X, \mu)$ has the property that its 0-dimensional sofic homology is not trivial. These actions have zero Σ -entropy and therefore have 0-th exponential Betti number equal to zero. My guess is that, like the Popa factors, such actions satisfy a constant bound on the growth of their homology.

(7) Fix a sofic approximation $\Sigma = \{\sigma_i : \Gamma \rightarrow \text{sym}(V_i)\}$. Let $F_i : \text{Abel} \rightarrow \mathbb{R}$ be the function $F_i(G) = \frac{\log \dim_{\mathbb{Q}}(G \otimes_{\mathbb{Z}} \mathbb{Q})}{\log \#V_i}$. The invariant $F_{\Sigma, d}(\mu)$ defined in Corollary 4.1 is the polynomial growth rate of the d -dimensional sofic homology. Given $t > 0$ and a dimension $d \geq 0$ does there exist an invariant measure μ with $F_{\Sigma, d}(\mu) = t$?

(8) The present paper shows that free groups of sufficiently large rank admit actions without the Weak Pinsker Property. Does the same result hold for all non-amenable sofic groups?

(9) If the 0-dimensional sofic homology of an ergodic action vanishes, then does the action have the Weak Pinsker Property?

(10) Does there exist an ergodic action with positive entropy that has no non-trivial direct Bernoulli factors? I conjecture that the frozen model introduced in [13] has this property.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TEXAS AT AUSTIN, AUSTIN, TX 78712

E-mail: lpbowen@math.utexas.edu

REFERENCES

- [1] T. Austin, Additivity properties of sofic entropy and measures on model spaces, *Forum Math. Sigma* **4** (2016), Paper No. e25, 79.
- [2] ———, The geometry of model spaces for probability-preserving actions of sofic groups, *Anal. Geom. Metr. Spaces* **4** (2016), no. 1, 160–186.
- [3] ———, Measure concentration and the weak Pinsker property, *Publ. Math. Inst. Hautes Études Sci.* **128** (2018), 1–119.
- [4] Á. Backhausz and B. Szegedy, On large-girth regular graphs and random processes on trees, *Random Structures Algorithms* **53** (2018), no. 3, 389–416.
- [5] M. Bayati, D. Gamarnik, and P. Tetali, Combinatorial approach to the interpolation method and scaling limits in sparse random graphs, *Ann. Probab.* **41** (2013), no. 6, 4080–4115.
- [6] B. Bollobás, A probabilistic proof of an asymptotic formula for the number of labelled regular graphs, *European J. Combin.* **1** (1980), no. 4, 311–316.
- [7] ———, *Random Graphs*, second ed., *Cambridge Stud. Adv. Math.*, vol. 73, Cambridge University Press, Cambridge, 2001.
- [8] L. Bowen, The ergodic theory of free group actions: entropy and the f -invariant, *Groups Geom. Dyn.* **4** (2010), no. 3, 419–432.
- [9] ———, Measure conjugacy invariants for actions of countable sofic groups, *J. Amer. Math. Soc.* **23** (2010), no. 1, 217–245.
- [10] ———, A brief introduction of sofic entropy theory, *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. III. Invited Lectures*, World Sci. Publ., Hackensack, NJ, 2018, pp. 1847–1866.
- [11] ———, Examples in the entropy theory of countable group actions, *Ergodic Theory Dynam. Systems* **40** (2020), no. 10, 2593–2680.
- [12] A. Coja-Oghlan and C. Efthymiou, On independent sets in random graphs, *Random Structures Algorithms* **47** (2015), no. 3, 436–486.
- [13] Jian Ding, Allan Sly, and Nike Sun, Maximum independent sets on random regular graphs, *Acta Math.* **217** (2016), no. 2, 263–340.

- [14] L. Q. Eifler, Open mapping theorems for probability measures on metric spaces, *Pacific J. Math.* **66** (1976), no. 1, 89–97.
- [15] G. Elek and E. Szabó, Sofic representations of amenable groups, *Proc. Amer. Math. Soc.* **139** (2011), no. 12, 4285–4291.
- [16] M. Foreman, D. J. Rudolph, and B. Weiss, The conjugacy problem in ergodic theory, *Ann. of Math.* (2) **173** (2011), no. 3, 1529–1586.
- [17] M. Foreman and B. Weiss, An anti-classification theorem for ergodic measure preserving transformations, *J. Eur. Math. Soc. (JEMS)* **6** (2004), no. 3, 277–292.
- [18] A. M. Frieze and T. Łuczak, On the independence and chromatic numbers of random regular graphs, *J. Combin. Theory Ser. B* **54** (1992), no. 1, 123–132.
- [19] D. Gamarnik and M. Sudan, Limits of local algorithms over sparse random graphs [extended abstract], *ITCS'14—Proceedings of the 2014 Conference on Innovations in Theoretical Computer Science*, ACM, New York, 2014, pp. 369–375.
- [20] C. Greenhill, S. Janson, J. H. Kim, and N. C. Wormald, Permutation pseudographs and contiguity, *Combin. Probab. Comput.* **11** (2002), no. 3, 273–298.
- [21] A. Hatcher, *Algebraic Topology*, Cambridge University Press, Cambridge, 2002.
- [22] B. Hayes, Mixing and spectral gap relative to Pinsker factors for sofic groups, *Proceedings of the 2014 Maui and 2015 Qinhuangdao Conferences in Honour of Vaughan F. R. Jones' 60th Birthday*, *Proc. Centre Math. Appl. Austral. Nat. Univ.*, vol. 46, Austral. Nat. Univ., Canberra, 2017, pp. 193–221.
- [23] ———, Relative entropy and the Pinsker product formula for sofic groups, preprint, <https://arxiv.org/abs/1605.01747>.
- [24] S. Janson, T. Łuczak, and A. Rucinski, *Random Graphs*, Wiley-Interscience Ser. Discrete Math. Optim., Wiley-Interscience, New York, 2000.
- [25] D. Kerr, Sofic measure entropy via finite partitions, *Groups Geom. Dyn.* **7** (2013), no. 3, 617–632.
- [26] D. Kerr and H. Li, Entropy and the variational principle for actions of sofic groups, *Invent. Math.* **186** (2011), no. 3, 501–558.
- [27] M. Ledoux, *The Concentration of Measure Phenomenon*, Math. Surveys Monogr., vol. 89, American Mathematical Society, Providence, RI, 2001.
- [28] H. Li and B. Liang, Sofic mean length, *Adv. Math.* **353** (2019), 802–858.
- [29] D. S. Ornstein, A K automorphism with no square root and Pinsker's conjecture, *Adv. Math.* **10** (1973), 89–102.
- [30] ———, A mixing transformation for which Pinsker's conjecture fails, *Adv. Math.* **10** (1973), 103–123.
- [31] M. S. Pinsker, Dynamical systems with completely positive or zero entropy, *Soviet Math. Dokl.* **1** (1960), 937–938.
- [32] M. Rahman and B. Virág, Local algorithms for independent sets are half-optimal, *Ann. Probab.* **45** (2017), no. 3, 1543–1577.
- [33] B. Seward, Krieger's finite generator theorem for actions of countable groups I, *Invent. Math.* **215** (2019), no. 1, 265–310.
- [34] ———, Krieger's finite generator theorem for actions of countable groups II, *J. Mod. Dyn.* **15** (2019), 1–39.
- [35] ———, Positive entropy actions of countable groups factor onto Bernoulli shifts, *J. Amer. Math. Soc.* **33** (2020), no. 1, 57–101.
- [36] ———, Weak containment and Rokhlin entropy, preprint, <https://arxiv.org/abs/1602.06680>.
- [37] J.-P. Thouvenot, On the stability of the weak Pinsker property, *Israel J. Math.* **27** (1977), no. 2, 150–162.