UNIVERSITY OF CALIFORNIA
Los Angeles

Entropy of Nonamenable Group Actions

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mathematics

by

Christopher Edgar Shriver

2021

ABSTRACT OF THE DISSERTATION

Entropy of Nonamenable Group Actions

by

Christopher Edgar Shriver
Doctor of Philosophy in Mathematics
University of California, Los Angeles, 2021
Professor Timothy Austin, Chair

Sofic entropy is an isomorphism invariant of measure-preserving actions of sofic groups introduced by Lewis Bowen in [Bow10c]. Its classical analogue was introduced in the 1950s by Kolmogorov and Sinai in order to show that Bernoulli shifts over $\mathbb{Z}$ are nonisomorphic when their base measures have different Shannon entropies. This entropy rate was actively studied over the next few decades and extended to arbitrary amenable groups by [OW87].

On the one hand, amenable groups provide an appropriate setting for entropy theory since they have a way of performing the kind of average used to define an entropy rate. On the other hand, statistical physicists have long been interested in some nonamenable structures, such as the Bethe lattice. The problem of finding an appropriate entropy notion for nonamenable group actions, and in particular the problem of isomorphism of Bernoulli shifts in this setting, remained open until Bowen's work. One way to briefly summarize the idea of sofic entropy is to say that we consider the entropy per site along a sequence of large finite systems which locally approximate the infinite one (called a sofic approximation) rather than large finite subsystems of the infinite one. An interesting problem which arises is what effect the choice of sofic approximation can have on the sofic entropy rate.

This thesis presents related work on several different problems of sofic entropy theory. In Chapter 2 (based on [Shr20b]) we study the f-invariant, a variant of sofic entropy for free-group actions introduced in [Bow10b] which can be defined using a kind of uniform random sofic approximation. We use a relative version of the f-invariant to show that the sofic entropy over a kind of "stochastic block model" random sofic approximation is given by the solution to an entropy-maximization problem. Understanding this optimization problem may shed further light on the dependence of sofic entropy on the sofic approximation.

Chapter 4, based on [Shr20a], uses a new notion of sofic free energy density to study Gibbs measures and Glauber dynamics for nearest-neighbor interacting particle systems on some nonamenable groups. The main results are that, under certain reasonable conditions, every Glauber-invariant, shift-invariant measure is Gibbs and that the Glauber evolution of any shift-invariant measure converges to the set of Gibbs measures. These extend results of [Hol71] for the Ising model on integer lattices.

Chapter 5, based on [Shr21], begins by proving a metastability result for states on finite graphs which are locally similar to the Cayley graph of a finitely-generated group: the Glauber evolution of any state on a finite graph will converge to the unique Gibbs state, but we show that if the initial state is "pseudo-Gibbs" in that it is in some sense consistent with *some* Gibbs measure on the infinite group, then that consistency will tend to persist for a long time. We then return to the entropy-maximization problem raised in Chapter 2. We show that a maximal-entropy joining of two Gibbs measures for nearest-neighbor interactions (not necessarily the same interaction) must be a relative product over the tail σ-algebra, unless every joining has entropy $-\infty$. In particular, if either is tail-trivial then the unique maximal-entropy joining is the product. In the latter case, this provides examples where the sofic entropy over a stochastic block model is equal to the f-invariant. We conclude by using recent results on bisections of random regular graphs [DMS17] to show that, for the free-boundary Ising model, the product self-joining has less than maximal f-invariant for some nontrivial temperature range.

The dissertation of Christopher Edgar Shriver is approved.

Guido Montufar

Georg Menz

Dimitri Shlyakhtenko

Timothy Austin, Committee Chair

University of California, Los Angeles

2021

TABLE OF CONTENTS

1 Introduction **1**

1.1 Equilibrium in statistical mechanics 2

1.1.1 The Ising model on an integer lattice 2

1.1.2 Free energy density beyond integer lattices 4

1.2 Entropy of nonamenable group actions 8

1.2.1 The f-invariant and its relative version 10

1.2.2 Maximal-entropy joinings 12

2 The relative f-invariant **14**

2.1 Introduction, main results 14

2.1.1 Main results 19

2.1.2 Random sofic approximations 21

2.2 Weights 23

2.2.1 Constructing weights and good models 26

2.3 Properties of F and f 27

2.4 Non-vacuity of main theorems 32

2.4.1 Theorem F 32

2.4.2 Theorems C and D 33

2.5 Counting Lemmas 33

2.6 Proof of Theorem F 36

2.6.1 Upper bound 36

2.6.2 Lower bound 38

2.7 Proof of Theorem C 41

2.7.1 Upper bound 42

2.7.2 Lower bound 44

2.8 Proof of Theorem D 48

2.8.1 Lower bound 49

2.8.2 Upper bound 49

2.9 Proof of Proposition 2.1.3 51

2.10 Proof of Lemma 2.2.3 53

2.10.1 The vertex measure 53

2.10.2 The $\mathtt{B}$ half-marginal 55

2.10.3 The edge measure 59

3 Gibbs measures and Glauber dynamics 64

3.1 Interaction 65

3.2 Glauber dynamics 67

3.3 Gibbs measures 68

3.4 Good models for measures on $\mathtt{A}^\Gamma$ 70

3.5 Free energy density 71

3.6 Measuring non-Gibbs-ness 74

3.7 Proofs of statements involving infinitary dynamics 75

3.7.1 Approximate equivariance (Proof of Theorem 3.2.1) 76

4 Free energy, Gibbs measures, and Glauber dynamics for nearest-neighbor interactions on trees 83

4.1 Introduction, main results 83

4.1.1 Related work 84

4.1.2 Precise statements of basic definitions and main theorems 85

4.2 Proof of Theorem B 87

4.3 Connection to property PA 93

4.3.1 Proof of Proposition 3.5.1 95

5 Metastability and maximal-entropy joinings of Gibbs measures on finitely-generated groups 96

5.1 Introduction, main results 96

5.1.1 Overview 101

5.2 Metastability of near-Gibbs-ness 102

5.2.1 Concentration from convexity 102

5.2.2 Controlling lateral motion 103

5.2.3 Proof of Theorem G 109

5.3 Maximal-entropy joinings 110

5.4 Non-optimal Gibbs joinings 118

5.4.1 Proof of Theorem H 122

5.4.2 Concentration of random homomorphisms 124

5.5 Proof of Proposition 5.3.1 126

References 133

LIST OF FIGURES

2.1 Picking the vertex measure 53

2.2 Picking the half-marginal 56

2.3 Picking the edge measure 60

3.1 Graph of F_0 75

ACKNOWLEDGMENTS

Thanks to Tim Austin for the suggestion of the projects in Chapters 2 and 5, and for many helpful conversations and comments on the preprints on which this document is based. Thanks also to Lewis Bowen, Brandon Seward, and Ben Hayes for helpful conversations and literature references.

This material is based upon work supported by the National Science Foundation under Grant No. DMS-1855694.

VITA

2016	B.S. Mathematics (Intensive), Yale University.
2016-2021	Teaching Assistant and Research Assistant, Mathematics Department, UCLA.

CHAPTER 1

Introduction

This chapter serves as an introduction to the main results of the dissertation. Most of them are stated precisely here; all are (re)stated in the following chapters.

In Section 1.1 we discuss notions of equilibrium in interacting particle systems. The situation is well-understood in the euclidean lattice setting, at least under the assumption of shift-invariance, but problems have persisted in the nonamenable setting. The standard approach to the problems we consider here uses a definition of free energy density (or relative entropy density) which fails when applied to trees and other nonamenable graphs. We discuss how, using ideas from sofic entropy theory, we can define a notion of free energy density on trees and use it to prove that a shift-invariant state is Gibbs if and only if it is Glauber-invariant, and that any shift-invariant state converges to the set of Gibbs states when evolved under Glauber dynamics.

In Section 1.2 we discuss sofic entropy and the f-invariant for shift systems. We provide a formula for the relative f-invariant in terms of a random sofic approximation which is a type of stochastic block model. We also show that the sofic entropy over this stochastic block model is the solution to an entropy-maximization problem, and provide a partial solution using the preceding results on Glauber dynamics.

1.1 Equilibrium in statistical mechanics

1.1.1 The Ising model on an integer lattice

The prototypical example of the relevant type of system is the Ising model, an old and well-studied model of magnetism introduced in [Isi25]. In this model we have a rectangular grid $R \subset \mathbb{Z}^r$ of particles, each of which can have 'spin' either $+1$ or -1. Each particle interacts only with its nearest neighbors, and possibly with an external field. The particles should tend to align with the field and, in the ferromagnetic case, with their neighbors.

We distinguish between a specific 'microstate' in $\{\pm 1\}^R$ describing the exact configuration of the system, and a 'state' in $\mathrm{Prob}(\{\pm 1\}^R)$ which describes the statistics of the system.

More generally, we can allow the local state space $\{\pm 1\}$ to be any finite set $\mathtt{A}$, and replace the rectangular grid with a locally finite graph $G = (V, E)$. If V is finite, the interaction is defined using an energy function of the form

$$U(\mathbf{x}) = \sum_{v \in V} h(\mathbf{x}(v)) + \sum_{\{v,w\} \in E} J(\mathbf{x}(v), \mathbf{x}(w)) \qquad \mathbf{x} \in \mathtt{A}^V$$

where $h \colon \mathtt{A} \to \mathbb{R}$ is an arbitrary function representing interaction with an 'external field,' $J \colon \mathtt{A}^2 \to \mathbb{R}$ is a symmetric function representing the 'pairwise interaction', and the second sum is over unordered pairs of adjacent vertices. We can think of the interaction as an ordered pair (h, J). Given a state $\zeta \in \mathrm{Prob}(\mathtt{A}^V)$, we write $\zeta(U)$ for the average energy, that is, the expectation of U.

For such systems there are several notions of 'equilibrium.'

The first equilibrium notion is that of a (variational) *equilibrium state*. This is defined using the *free energy* of a state which, for finite systems, is given by $A(\zeta) = \zeta(U) - \mathrm{H}(\zeta)$. The free energy often includes a dependence on a temperature T (or inverse temperature β), but for our purposes that would unnecessarily complicate many expressions. For infinite systems the energy of a microstate may not be well-defined. In the case $G = \mathbb{Z}^r$ we instead use the *free energy density* defined by

$$\mathrm{a}(\mu) = \limsup_{N\to\infty} \frac{1}{(2N+1)^r} A(\mu^N) \tag{1.1}$$

where μ^N is the marginal on $[-N, N]^r \subset \mathbb{Z}^r$, which we interpret as the state of that subsystem. An equilibrium state is one which minimizes the free energy (density).

For finite systems the free energy is a strictly convex function, so has a unique minimizer. This is called the *Gibbs state*, and it assigns a microstate $\mathbf{x} \in \mathtt{A}^V$ probability proportional to $\exp(-U(\mathbf{x}))$. For infinite systems, the energy is typically infinite so we cannot use the same formula to define an 'infinite-volume Gibbs state.' Instead, we note that the finitary Gibbs state is uniquely determined by specifying the conditional distributions of the spins at individual vertices given the spins elsewhere; this point of view extends naturally to infinite systems. The collection of these conditional distributions is called the 'Gibbs specification,' and the convention is to define infinite-volume Gibbs measures using this specification. A good reference for this framework is [Geo11]; see also Chapter 3 below. In general the Gibbs states form a compact, convex set which we denote $\mathscr{G}$. If we wish to reference a particular interaction $\Phi = (h, J)$ we may write $\mathscr{G}(\Phi)$. The set of shift-invariant Gibbs measures will be denoted $\mathscr{G}^\Gamma(\Phi)$.

A third notion of equilibrium relates to the *Glauber dynamics*, a standard model for how a system evolves over time which may be described as follows: assign each $v \in V$ an alarm clock which rings at random intervals which are distributed as exponential random variables with mean 1 (*i.e.* a Poisson clock). Different clocks and different time intervals are independent. Each time a clock rings, we rerandomize the associated vertex conditioned on the spins elsewhere using the Gibbs specification. In this setting it is natural to say that a system is in equilibrium if it is invariant under the dynamics. A good reference is [Lig05]; see also Chapter 3 below.

The invariance of Gibbs states under Glauber dynamics is fairly straightforward; see for example [Lig05, Theorem IV.2.15]. Other relationships between these three equilibrium notions are less clear for infinite systems.

Holley [Hol71] showed for the Ising model on $\mathbb{Z}^r$ that free energy density is nonincreasing under Glauber dynamics. Furthermore, if the initial state is shift-invariant and non-Gibbs then the free energy density *strictly* decreases and the state weakly converges to $\mathscr{G}$. This implies that any shift-invariant state which is either a variational equilibrium state or Glauber-invariant must also be Gibbs.

His approach appears to rely essentially on the fact that a large finite subsystem can, to a good approximation, be treated as isolated. This is because the subsystem can only interact with its surroundings through its boundary, which is relatively small. This property has long been known to be important in statistical physics. For example, as stated in [LL58]:

> the particles which take part in the interaction of a subsystem with neighboring parts of the system, are mainly those nearest its surface. Their number in comparison with the total number of particles quickly falls with an increase in size of the latter. This fact [...] implies that we can consider [separate subsystems] as independent in a statistical sense.

In other words, the essential fact is really that $\mathbb{Z}^r$ is *amenable*; this is evidenced by the sequence of boxes $\left([-N,N]^r\right)_{N\in\mathbb{N}}$ being a 'Følner sequence,' which is a sequence of finite sets that exhausts the full group and whose boundaries are vanishingly small in proportion to the interiors. When available, such sequences are typically appropriate for averages as in Equation 1.1.

1.1.2 Free energy density beyond integer lattices

The rectangular grid is natural for modeling an arrangement of particles in euclidean space. However, it is also natural to study similar systems with other dependence structures. A particular case of interest is the infinite regular tree, sometimes called the Bethe lattice. Slightly more generally, we could consider a finitely-generated group Γ with generating set $S = \{s_1, \dots, s_r\}$ and identity e.

As mentioned above, Holley's methods do not work when Γ is nonamenable. Lewis Bowen's work on sofic entropy, initiated in [Bow10b, Bow10c], suggests a solution may be found by replacing 'finite subsystems' with 'finite systems which locally look like the infinite system.' In $\mathbb{Z}^r$, finite subsystems satisfy the latter description because a large neighborhood of any vertex far enough away from the boundary looks like $\mathbb{Z}^r$, and most vertices are far from the boundary. For $\Gamma = \mathbb{F}_r$, typical large random $2r$-regular graphs have the desired property (they tend not to have very many small loops).

An immediate problem is that, while we can pass from the state of a system to the state of a subsystem by taking a marginal, there is no obvious analogue of this operation that will produce a state on a separate finite system. More concretely, we must answer a question like: if $G = (V, E)$ is a large $2r$-regular graph and $\mu \in \operatorname{Prob}(\mathtt{A}^{\mathbb{F}_r})$, when can we think of a state $\zeta \in \operatorname{Prob}(\mathtt{A}^V)$ as a finitary version of μ?

Again, we can provide a useful answer to this using the framework of Bowen's sofic entropy theory, with more similarity to recent variants [Alp16, Aus16, Abe18]: the idea is to require that the 'local statistics' of ζ are in some sense consistent with μ. To make this more precise, fix $R \in \mathbb{N}$ and suppose that for most $v \in V$ the radius-R ball $\mathrm{B}^G(v, R)$ in G is isomorphic to the ball $\mathrm{B}^\Gamma(e, R)$ centered at the identity. Then we can identify each of these balls in the finite graph with $\mathrm{B}^\Gamma(e, R)$, so that the marginal of ζ on each of them can be thought of as a measure on $\mathtt{A}^{\mathrm{B}^\Gamma(e,R)}$. The average of these marginals is the *empirical distribution*, denoted $P_\zeta^{G,R} \in \operatorname{Prob}(\mathtt{A}^{\mathrm{B}^\Gamma(e,R)})$.

There is one problem with this description: the choices of isomorphisms $\mathrm{B}^G(v, R) \cong \mathrm{B}^\Gamma(e, R)$ do make a difference in general. We therefore assume more structure of G in order to ensure that there is a canonical isomorphism, and in particular a canonical way to lift marginals of ζ.

The structure can be described as follows: the edges of the Cayley graph of Γ naturally come directed and labeled by the generators. Assume that the edges of G are also directed and labeled, with one s_i-labeled edge coming in and one going out of each vertex for each

i. Now for any $v \in V$ with $\mathrm{B}^G(v,R) \cong \mathrm{B}^\Gamma(e,R)$, there is exactly one isomorphism which respects edge labelings and directions. Using these canonical isomorphisms, we end up with a well-defined empirical distribution.

We will encode such a structure on G by a homomorphism $\sigma \colon \Gamma \to \mathrm{Sym}(V)$, where each $\sigma(s_i)$ is the permutation that sends each $v \in V$ to the vertex on the other end of the s_i-labeled edge coming out of v. Below, we will always refer to elements of $\mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$ rather than directed, edge-labeled graphs with vertex set V, but one should keep this graph-homomorphism correspondence in mind. See Chapter 3 for a more explicit definition of the empirical distribution of a state ζ.

In fact, using such a structure we can pull back labelings of G to labelings of Γ without regard for local similarity; although this local similarity will be very important for other purposes it is convenient to be able to ignore it here. Specifically, given $\mathbf{x} \in \mathtt{A}^V$ and $v \in V$ we define the pullback name $\Pi_v^\sigma \mathbf{x} \in \mathtt{A}^\Gamma$ by

$$\Pi_v^\sigma \mathbf{x}(\gamma) = \mathbf{x}(\sigma^\gamma v).$$

Our version of the "marginal of ζ at v" is then the pushforward $[\mathbf{x} \mapsto \Pi_v^\sigma \mathbf{x}]_* \zeta$, and we define the empirical distribution to be the average over choices of v:

$$P_\zeta^\sigma = \frac{1}{|V|} \sum_{v \in V} [\mathbf{x} \mapsto \Pi_v^\sigma \mathbf{x}]_* \zeta.$$

Instead of truncating at radius R, the "localness" will come from making comparisons in the weak topology on $\mathrm{Prob}(\mathtt{A}^\Gamma)$.

For a weak-open neighborhood $\mathcal{O}$ of μ, we say that ζ is $\mathcal{O}$-consistent with μ (over σ) if $P_\zeta^\sigma \in \mathcal{O}$. Let $\mathbf{\Omega}(\sigma, \mathcal{O})$ denote the set of such ζ; we can think of these as suitable finitary versions of μ with quality parameter $\mathcal{O}$. Note that the finite set V for which $\mathbf{\Omega}(\sigma, \mathcal{O}) \subset \mathrm{Prob}(\mathtt{A}^V)$ is implicitly specified by σ via its codomain.

We can now define our notion of free energy density: let $\Sigma = (\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n)))_{n \in \mathbb{N}}$ be a sequence of homomorphisms such that for each $R \in \mathbb{N}$ the fraction of $v \in V_n$ such that

$\mathrm{B}^{\sigma_n}(v,R) \cong \mathrm{B}^{\Gamma}(e,R)$ approaches 1 as $n \to \infty$. The letter Σ reflects that this is a sofic approximation to Γ. We then define the *free energy density* of $\mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$ relative to Σ by

$$\mathrm{a}_\Sigma(\mu) = \sup_{\mathcal{O} \ni \mu} \limsup_{n\to\infty} \frac{1}{|V_n|} \inf_{\zeta \in \mathbf{\Omega}(\sigma_n, \mathcal{O})} A(\zeta),$$

where the supremum is over weak-open neighborhoods of μ.

Note that $\mathrm{a}_\Sigma(\mu)$ has an implicit dependence on a specific nearest-neighbor interaction, via the energy functions $U \colon \mathtt{A}^{V_n} \to \mathbb{R}$.

Remark. The constraint "$\zeta \in \mathbf{\Omega}(\sigma_n, \mathcal{O})$" will typically prevent the infimum from simply being attained by the true finitary Gibbs state, which has minimal free energy among *all* states on the graph of σ_n, rather than just among those whose empirical distributions are close to μ.

However, in the special case where μ is the local limit of the Gibbs states on σ_n this constraint becomes irrelevant. In particular, we have

$$\mathrm{a}_\Sigma(\mu) = \limsup_{n\to\infty} \frac{1}{|V_n|} \left[-\log Z_n\right]$$

where Z_n is the 'partition function,' i.e. the normalizing factor which appears in the definition of the Gibbs state on σ_n. A strong enough mode of convergence to an explicit limit was established for the Ising model in [DM10] for nonzero external field and in [MMS12] for zero external field. The former paper also shows how to provide an expression for the limiting free energy density. □

Using this notion of free energy density along with ideas from Holley's paper, in Chapter 4 we will prove the following:

Theorem A. *For any choice of* Σ *and any nearest-neighbor interaction, every* $\mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$ *minimizing* a_Σ *is Gibbs for that interaction, unless* a_Σ *is identically* $+\infty$*.*

This will follow from monotonicity of free energy density (Proposition 4.2.2) and from the fact that free energy density is always strictly decreasing under Glauber dynamics as long as the measure is not Gibbs (and $\mathrm{a}_\Sigma \neq +\infty$).

We also show that, under some conditions, Glauber dynamics converges to the set of Gibbs measures:

Theorem B. *Suppose $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$, and let μ_t denote its evolution under Glauber dynamics. If there exist $s \geq 0$ and Σ such that $\mathrm{a}_\Sigma(\mu_s) < +\infty$, then μ_t converges weakly to the set of Gibbs measures as $t \to \infty$.*

Here, shift-invariance means invariance under the natural action of Γ. It is not necessary for μ to be invariant under the full automorphism group of the Cayley graph of Γ.

The requirement that μ_s have finite free energy density with respect to some sofic approximation is always satisfied for many groups, including amenable groups and finite-rank free groups. See Corollary 4.1.1.

A key step in the proof of Theorem B is to show that if a finitary state ζ is consistent with μ then ζ_t is also consistent with μ_t; see Theorem 3.2.1. There is some decay in the quality parameters over time, but this is controlled by the degree of local similarity of the finite system to Γ. This is an analogue of the fact that a large finite subsystem in $\mathbb{Z}^r$ can be treated as isolated from its surroundings.

One application of Theorem B is Theorem E below.

1.2 Entropy of nonamenable group actions

Amenability, mentioned above in the context of free energy density, also plays a central role in entropy theory. If Γ is a countable amenable group, the entropy rate of an element of $\mathrm{Prob}(\mathtt{A}^\Gamma)$ is defined by an averaging process analogous to Equation 1.1. Again, it turns out to be important that the average is taken over a Følner sequence.

The problem of nonamenability stood for several decades until work by Lewis Bowen [Bow10b, Bow10c] extended entropy theory to sofic groups. Most reasonable groups are known to be sofic; there are currently no known non-sofic groups. A sofic group is one which admits a sofic approximation. Sofic approximations were essentially introduced above: a sequence of homomorphisms from Γ to finite symmetric groups such that the graphs of these homomorphisms locally look like the Cayley graph of Γ is a sofic approximation to Γ. More generally the maps should only be required to be asymptotically homomorphisms in some sense, but for simplicity we only consider true homomorphisms here.

Let $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$ be a homomorphism, and let $\mathtt{A}$ be a finite set. Given a microstate $\mathbf{x} \in \mathtt{A}^V$, we can define its empirical distribution by

$$P_{\mathbf{x}}^{\sigma} := P_{\delta_{\mathbf{x}}}^{\sigma} = \frac{1}{|V|} \sum_{v \in V} \delta_{\Pi_v^{\sigma} \mathbf{x}}.$$

We say $\mathbf{x}$ is an $\mathcal{O}$-*good model* for $\mu \in \operatorname{Prob}(\mathtt{A}^{\Gamma})$ if $P_{\mathbf{x}}^{\sigma} \in \mathcal{O}$. Call the set of such microstates $\Omega(\sigma, \mathcal{O})$.

The *sofic entropy* of a measure $\mu \in \operatorname{Prob}(\mathtt{A}^{\Gamma})$ is the exponential growth rate of the number of good models for μ: given a sofic approximation $\Sigma = \big(\sigma_n\big)_{n \in \mathbb{N}}$,

$$\mathrm{h}_{\Sigma}(\mu) := \inf_{\mathcal{O}} \limsup_{n \to \infty} \frac{1}{|V_n|} \log |\Omega(\sigma_n, \mathcal{O})|.$$

A measure with many good models could be seen as being "more random," since more microstates are consistent with it. However, standard intuition about entropy does not always apply for nonamenable groups: in particular, sofic entropy can increase under factor maps (so one system can generate another which is "more random"). Ornstein and Weiss had already observed in [OW87, Appendix C] that any reasonable notion of entropy for free-group actions would have this property.

Sofic entropy may take different values depending on Σ, but in all known occurrences of this phenomenon one of the values is $-\infty$. This is typically considered degenerate, since it means that one of the approximations supports no good models for μ. It is an open problem whether two sofic approximations may assign distinct finite entropy values to the same measure.

1.2.1 The f-invariant and its relative version

A related, but more tractable, problem is to count the expected number of good models over a sequence of random homomorphisms. In [Bow10a], Bowen studied the following case:

Let $\Gamma = \mathbb{F}_r$, and for each n let $V_n = \{1, \dots, n\}$ and let σ_n be a random permutation distributed uniformly on $\mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))$. We consider $\mathbf{\Sigma} = (\sigma_n)_{n \in \mathbb{N}}$ to be a random sofic approximation. The exponential growth rate of the expected number of good models for $\mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$ may then be written

$$\mathrm{h}_{\mathbf{\Sigma}}(\mu) := \inf_{\mathcal{O}} \limsup_{n \to \infty} \frac{1}{|V_n|} \log \mathbb{E} |\Omega(\sigma_n, \mathcal{O})|.$$

Bowen showed that this coincides with the f-invariant of μ, which he had introduced in [Bow10c]. If μ is a homogeneous Markov chain then $f(\mu)$ can be easily calculated.

Now suppose we have two finite sets $\mathtt{A}, \mathtt{B}$, and let $\mu \in \mathrm{Prob}((\mathtt{A} \times \mathtt{B})^\Gamma)$ be shift-invariant. Write the marginals on $\mathtt{A}^\Gamma, \mathtt{B}^\Gamma$ as $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$. The *relative f-invariant* is given by

$$f(\mu \mid \mathtt{B}) = f(\mu) - f(\mu_{\mathtt{B}}),$$

as long as $f(\mu_{\mathtt{B}})$ is finite; see Chapter 2. We may try to interpret this as the growth rate of the average number of good models for μ per good model for $\mu_{\mathtt{B}}$. Theorem C may be seen as a precise version of this statement.

Each good model for μ, as an element of $(\mathtt{A} \times \mathtt{B})^V$, can naturally be split into a pair of good models for $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$. The converse is not quite true: a pair of good models for $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$ combine to form a good model for some joining (i.e. shift-invariant coupling) of the two, but not necessarily μ. Thus it makes sense to ask: given a good model for $\mu_{\mathtt{B}}$, how many extensions does it admit to a good model for μ?

In Chapter 2 we introduce a type of 'stochastic block model' random sofic approximation. This consists of a fixed sequence $(\mathbf{y}_n \in \mathtt{B}^{V_n})_{n \in \mathbb{N}}$ and a sequence of random homomorphisms $\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))$. The distribution of σ_n is uniform on a particular set of homomorphisms chosen to ensure that for any quality parameter $\mathcal{O}$ we have $\mathbf{y}_n \in \Omega(\sigma_n, \mathcal{O})$ for all large n. Thus $\mathbf{y}_n$ is a planted partition of V_n, and $\mu_{\mathtt{B}}$ controls the statistics of edges between the parts.

More specifically, we will define a family of pseudometrics $\{d_k^*\}_{k\in\mathbb{N}}$ which generate the weak* topology on $\mathrm{Prob}(\mathtt{A}^\Gamma)$ and, for $\sigma_0 \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$, $\mathbf{y}_0 \in \mathtt{B}^V$, and $k \in \mathbb{N}$, define

$$\mathtt{SBM}(\sigma_0, \mathbf{y}_0, k) \coloneqq \mathrm{Unif}(\{\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V)) \,:\, d_k^*(P_{\mathbf{y}_0}^{\sigma},\, P_{\mathbf{y}_0}^{\sigma_0}) = 0\}).$$

Theorem C. *Let $\alpha\colon X \to \mathtt{A}$ and $\beta\colon X \to \mathtt{B}$ be finite observables. Let m_n approach infinity as n goes to infinity while satisfying $m_n = o(\log\log n)$. For each n let $\mathbf{y}_n \in \mathtt{B}^n$ and $\sigma_n \in \mathrm{Hom}(G, \mathrm{Sym}(n))$ be such that*

$$d_{m_n}^*(P_{\mathbf{y}_n}^{\sigma_n}, \beta_*^G \mu) = O\big(\tfrac{1}{\log n}\big).$$

Suppose that $f_\mu(T, \beta) > -\infty$. With $\mu_n = \mathtt{SBM}(\sigma_n, \mathbf{y}_n, m_n)$,

$$f_\mu(T, \alpha \mid \beta) = \inf_{\mathcal{O} \ni (\alpha\beta)_*^G \mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x}, \mathbf{y}_n) \in \Omega(\sigma, \mathcal{O})\}|.$$

Using this result, we obtain a formula for the growth rate of the average number of good models a stochastic block model admits for $\mu_{\mathtt{A}}$. If we do not require them to combine with the planted good model $\mathbf{y}_n$ for $\mu_{\mathtt{B}}$ to form a good model for a particular joining μ, then we end up with an optimization over all joinings:

Theorem D. *Let μ_n, α, β be as in the statement of Theorem C. Then*

$$\inf_{\mathcal{O} \ni \alpha_*^G \mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\Omega(\sigma, \mathcal{O})| = \sup_{\lambda \in \mathsf{J}(\alpha_*^G \mu,\, \beta_*^G \mu)} f_\lambda(S, \mathtt{a} \mid \mathtt{b}).$$

The left-hand side here is just the sofic entropy of $\mu_{\mathtt{A}}$ with respect to the random sofic approximation $\mathbf{\Sigma}_{\mathtt{B}} \coloneqq (\mu_n)$. Taking λ to be the product joining $\mu_{\mathtt{A}} \times \mu_{\mathtt{B}}$, we see in particular that

$$\mathrm{h}_{\mathbf{\Sigma}_{\mathtt{B}}}(\mu_{\mathtt{A}}) \geq f(\mu_{\mathtt{A}} \times \mu_{\mathtt{B}}) - f(\mu_{\mathtt{B}}) = f(\mu_{\mathtt{A}}).$$

Thus any stochastic block model supports at least as many good models (on average) as a uniformly random graph.

Now suppose $\mu_{\mathtt{A}}$ and $\mu_{\mathtt{B}}$ are both the free-boundary Ising state at low enough temperature that $f(\mu_{\mathtt{A}}) < 0$ [Bow20a, Section 3.3]. We can take λ to be the diagonal joining to get

$$\mathrm{h}_{\mathbf{\Sigma}_{\mathtt{B}}}(\mu_{\mathtt{A}}) \geq 0.$$

But with $\mathbf{\Sigma}$ defined as at the beginning of this section, we have

$$\mathrm{h}_{\mathbf{\Sigma}_{\mathtt{B}}}(\mu_{\mathtt{A}}) > f(\mu_{\mathtt{A}}) = \mathrm{h}_{\mathbf{\Sigma}}(\mu_{\mathtt{A}}).$$

This is one way to see that there exist random sofic approximations with distinct entropy values, but the difference may be explained by the fact that $\mathbf{\Sigma}$ has a possibility of not supporting any good models for $\mu_{\mathtt{A}}$ at all (otherwise the f-invariant would be nonnegative).

Question: can we find some $\mu_{\mathtt{A}}$ and two stochastic block models $\mathbf{\Sigma}_{\mathtt{B}}, \mathbf{\Sigma}'_{\mathtt{B}}$ which both guarantee good models for $\mu_{\mathtt{A}}$ and with $\mathrm{h}_{\mathbf{\Sigma}_{\mathtt{B}}}(\mu_{\mathtt{A}}) > \mathrm{h}_{\mathbf{\Sigma}'_{\mathtt{B}}}(\mu_{\mathtt{A}})$?

If 'yes', it should be possible to extract two (nonrandom) sofic approximations which give distinct finite entropy values to $\mu_{\mathtt{A}}$.

1.2.2 Maximal-entropy joinings

To answer this question, we will need to understand better the optimization problem which appears in Theorem D. The following theorem is progress in that direction. To state it we need to introduce one last concept: Suppose we have two nearest-neighbor interactions $\Phi^{\mathtt{A}}, \Phi^{\mathtt{B}}$ on $\mathtt{A}^\Gamma$ and $\mathtt{B}^\Gamma$ defined as above using 'external fields' $h_{\mathtt{A}}, h_{\mathtt{B}}$ and 'pair interactions' $J_{\mathtt{A}}, J_{\mathtt{B}}$. We can produce a nearest-neighbor interaction $\Phi^{\mathtt{A}} \oplus \Phi^{\mathtt{B}}$ on $(\mathtt{A} \times \mathtt{B})^\Gamma$ by setting $h_{\mathtt{AB}}((\mathtt{a}, \mathtt{b})) = h_{\mathtt{A}}(\mathtt{a}) + h_{\mathtt{B}}(\mathtt{b})$ and $J_{\mathtt{AB}}((\mathtt{a}_1, \mathtt{b}_1), (\mathtt{a}_2, \mathtt{b}_2)) = J_{\mathtt{A}}(\mathtt{a}_1, \mathtt{a}_2) + J_{\mathtt{B}}(\mathtt{b}_1, \mathtt{b}_2)$. We will call this the *sum interaction*; see also the closely-related definition of a "product specification" in [Geo11, Example 7.18].

Theorem E. *Let λ be a joining of two shift-invariant Gibbs measures $\mu_{\mathtt{A}} \in \mathscr{G}^\Gamma(\Phi^{\mathtt{A}}), \mu_{\mathtt{B}} \in \mathscr{G}^\Gamma(\Phi^{\mathtt{B}})$ for nearest-neighbor interactions $\Phi^{\mathtt{A}}, \Phi^{\mathtt{B}}$. Let Σ be a random sofic approximation to Γ, and assume that there is some joining λ of $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$ with $\mathrm{h}_\Sigma(\lambda) > -\infty$.*

If λ *maximizes* h_Σ *among all joinings of* $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$*, then* $\lambda \in \mathscr{G}^\Gamma(\Phi^{\mathtt{A}} \oplus \Phi^{\mathtt{B}})$.

Note that, for our purposes, deterministic sofic approximations will be considered to be a special case of random sofic approximations.

The proof uses the results above regarding Gibbs states and Glauber dynamics. An additional key result, possibly of independent interest, is that a good model for a Gibbs state μ is likely to stay a good model for μ for a long time when evolved under the associated Glauber dynamics, as long as the finite system has enough local similarity to Γ. This is proven using the main technical result used to prove Theorem B combined with a concentration argument which relies on the fact that the Gibbs states form a face of the simplex of shift-invariant probability measures (Lemma 3.3.1). For a precise statement of this metastability result, see Theorem G.

The extreme (i.e. tail-trivial) Gibbs states for the product interaction are exactly the products of extreme Gibbs states for the factors [Geo11, Equation 7.19]. It follows that

1. The supremum is attained by a joining whose disintegration over the tail σ-algebra consists of product measures, *i.e.* a relatively independent joining over the tail; see Proposition 5.3.1 for the precise interpretation of this.

2. If either $\mu_{\mathtt{A}}$ or $\mu_{\mathtt{B}}$ is extreme then the product joining attains the supremum. In particular, the sofic entropy of $\mu_{\mathtt{A}}$ over a stochastic block model generated by $\mu_{\mathtt{B}}$ is equal to $f(\mu_{\mathtt{A}})$.

CHAPTER 2

The relative f-invariant

The f-invariant is an isomorphism invariant of free-group measure-preserving actions introduced by Lewis Bowen in [Bow10b], where it was used to show that two finite-entropy Bernoulli shifts over a finitely generated free group can be isomorphic only if their base measures have the same Shannon entropy. In [Bow10a] Bowen showed that the f-invariant is a variant of sofic entropy; in particular it is the exponential growth rate of the expected number of good models over a uniform random homomorphism.

In this chapter we present an analogous formula for the relative f-invariant and use it to prove a formula for the exponential growth rate of the expected number of good models over a random sofic approximation which is a type of stochastic block model.

2.1 Introduction, main results

Let Γ denote the the rank-r free group with generating set $\{s_1, \dots, s_r\}$ and identity e, and let (X, μ, T) be a measure-preserving Γ-system, *i.e.* T is a homomorphism from Γ to the automorphism group of the standard probability space (X, μ). We will not need to make explicit use of the σ-algebra on X, so we leave it unnamed.

An *observable* on X is a measurable map with domain X. In this dissertation the codomain will be a finite set endowed with the discrete sigma algebra; in this case we call the map a *finite observable* and the codomain an *alphabet*.

Any observable $\alpha\colon X \to \mathtt{A}$ induces a map $\alpha^\Gamma\colon X \to \mathtt{A}^\Gamma$ by setting

$$(\alpha^\Gamma(x))_g = \alpha(T_g x) \quad \text{for all } g \in \Gamma.$$

We call the $\mathtt{A}$-coloring $\alpha^\Gamma(x)$ of Γ the *itinerary* of x, since it records the observations that will be made over the entire orbit of x under the action of Γ. We also similarly define the map $\alpha^H\colon X \to \mathtt{A}^H$ for any subset H of Γ. We abbreviate $\alpha^n := \alpha^{\mathrm{B}(e,n)}$, where $\mathrm{B}(e,n)$ is the closed ball of radius n centered at the identity in Γ, which is endowed with the word-length metric. If $\beta\colon X \to \mathtt{B}$ is a second finite observable, we denote by $\alpha\beta\colon X \to \mathtt{A} \times \mathtt{B}$ the map $\alpha\beta(x) = (\alpha(x), \beta(x))$.

The (Shannon) entropy of a finite observable $\alpha\colon X \to \mathtt{A}$ is defined by

$$\mathrm{H}_\mu(\alpha) = -\sum_{a \in \mathtt{A}} \alpha_*\mu(a) \log \alpha_*\mu(a),$$

where $\alpha_*\mu \in \mathrm{Prob}(\mathtt{A})$ is the pushforward measure, with the convention $0 \log 0 = 0$. The entropy of α can be interpreted as the expected amount of information revealed by observing α, assuming its distribution $\alpha_*\mu$ is known.

An early application of Shannon's entropy to ergodic theory was its use by Kolmogorov and Sinai to show that there exist nonisomorphic Bernoulli shifts over $\mathbb{Z}$. A Bernoulli shift over $\mathbb{Z}$ is a system of the form $(\mathtt{A}^\mathbb{Z}, \mu^\mathbb{Z}, S)$ for some alphabet $\mathtt{A}$ and $\mu \in \mathrm{Prob}(\mathtt{A})$; S is the shift action of $\mathbb{Z}$. They did this by defining an *entropy rate* for $\mathbb{Z}$-systems, which can be interpreted as the average information per unit time revealed by observing the system. For a Bernoulli shift $(\mathtt{A}^\mathbb{Z}, \mu^\mathbb{Z}, S)$, the entropy rate is simply the "base entropy" $\mathrm{H}_\mu(\alpha)$, where $\alpha\colon \mathtt{A}^n \to \mathtt{A}$ is the "time zero" observable.

Isomorphism invariance of the KS entropy rate is typically proven using the fact that entropy rate is nonincreasing under factor maps (which are surjective homomorphisms of measure-preserving systems). This fact can be interpreted as stating that a system cannot simulate another system that is "more random."

The entropy rate was soon generalized to systems acted on by an arbitrary amenable group (such as $\mathbb{Z}^d$). Extending beyond amenable groups proved more difficult, and in fact it was found to be impossible for such an extension to preserve all desirable properties of the KS entropy rate. In particular, an entropy rate for nonamenable group actions which assigns Bernoulli shifts their base entropy cannot be nonincreasing under factor maps [OW87, Appendix C].

The first invariant to distinguish between Bernoulli shifts over free groups is Lewis Bowen's f-invariant. Following [Bow10a], this can be defined by

$$F_\mu(T,\alpha) = (1-2r)\mathrm{H}_\mu(\alpha) + \sum_{i=1}^{r} \mathrm{H}_\mu(\alpha^{\{e,s_i\}})$$
$$f_\mu(T,\alpha) = \inf_n F_\mu(T,\alpha^n) = \lim_{n\to\infty} F_\mu(T,\alpha^n).$$

The main theorem of [Bow10b] is that $f_\mu(T,\alpha)$ depends on the observable α only through the σ-algebra it generates. In particular, the common value of $f_\mu(T,\alpha)$ among all α which generate the σ-algebra of the measurable space X (assuming such α exist) is a measure-conjugacy invariant of the system (X,μ,T). In the same paper, he showed that the f-invariant of a Bernoulli shift is the Shannon entropy of the base measure; in particular, Bernoulli shifts with different base entropies are nonisomorphic.

In [Bow10a], Bowen gave an alternate formula for the f-invariant, which we now introduce.

For any homomorphism $\sigma\colon \Gamma \to \mathrm{Sym}(n)$ we have a Γ-system $([n], \mathrm{Unif}(n), \sigma)$, and we can consider a labeling $\mathbf{x} \in \mathtt{A}^n$ as an $\mathtt{A}$-valued observable on this system. We denote the law of its itinerary by $P^\sigma_{\mathbf{x}} = \mathbf{x}^G_* \mathrm{Unif}(n)$ and call this the *empirical distribution* of $\mathbf{x}$. We say that $\mathbf{x}$ is a good model for α over σ if it is difficult to distinguish the Γ-systems (X,μ,T) and $([n], \mathrm{Unif}(n), \sigma)$ via their respective observables α and $\mathbf{x}$. To make this precise, we denote

$$\Omega(\sigma,\mathcal{O}) \coloneqq \{\mathbf{x} \in \mathtt{A}^n \,:\, P^\sigma_{\mathbf{x}} \in \mathcal{O}\},$$

which is a set of good models for α over σ if $\mathcal{O}$ is a weak*-open neighborhood of $\alpha^G_*\mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$; the particular set $\mathcal{O}$ quantifies how good the models are. The alphabet $\mathtt{A}$ is given the discrete topology and $\mathtt{A}^\Gamma$ the product topology, so "weak*-close" means marginals on some finite sets are close in total variation norm.

For each $n \in \mathbb{N}$, let $\mu_n = \mathrm{Unif}(\mathrm{Hom}(\Gamma, \mathrm{Sym}(n)))$. Bowen showed in [Bow10a] that the f-invariant is given by

$$f_\mu(T,\alpha) = \inf_{\mathcal{O} \ni \alpha_*^G \mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma \sim \mu_n} |\Omega(\sigma, \mathcal{O})|.$$

To make an analogy with statistical physics, we can think of $\alpha_*^G \mu$ as a macroscopic statistical distribution of the state of a system; then the f-invariant is the exponential growth rate of the expected number of microstates on a sequence of finite random graphs that are consistent with these statistics.

More generally, given any random or deterministic sofic approximation $\Sigma = \{\mu_n\}_{n=1}^\infty$, we can define the sofic entropy relative to Σ by

$$\mathrm{h}_{\Sigma,\mu}(T,\alpha) = \inf_{\mathcal{O} \ni \alpha_*^G \mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma \sim \mu_n} |\Omega(\sigma, \mathcal{O})|.$$

Here each μ_n is a probability measure on the set of functions $\Gamma \to \mathrm{Sym}(n)$ which is supported on functions which are approximately free homomorphisms.

This dissertation is motivated by a desire to better understand the dependence of sofic entropy on the sofic approximation Σ. For any choice of Σ, the sofic entropy agrees with KS entropy if the acting group is amenable [Bow12] and with the Shannon entropy of the base if the system is a Bernoulli shift [Bow10c]. For some systems, the sofic entropy can be finite relative to some sofic approximations and $-\infty$ relative to others. It is unknown whether two deterministic sofic approximations can yield different finite entropy values for the same system.

In this chapter, we express the entropy relative to a type of stochastic block model in terms of the relative f-invariant, which we now introduce.

If α, β are two finite observables, the conditional entropy is

$$\mathrm{H}_\mu(\alpha|\beta) = \mathrm{H}_\mu(\alpha\beta) - \mathrm{H}_\mu(\beta).$$

This can be interpreted as the expected amount of information revealed by observing α if both the value of β and the joint distribution of α and β are known. The relative f-invariant is defined by

$$\begin{aligned} F_\mu(T,\alpha|\beta) &= F_\mu(T,\alpha\beta) - F_\mu(T,\beta) \\ &= (1-2r)\mathrm{H}_\mu(\alpha|\beta) + \sum_{i=1}^{r} \mathrm{H}_\mu(\alpha^{\{e,s_i\}} \mid \beta^{\{e,s_i\}}) \\ f_\mu(T,\alpha|\beta) &= \inf_{k_1\in\mathbb{N}} \sup_{k_2\in\mathbb{N}} F_\mu(T,\alpha^{k_1} \mid \beta^{k_2}). \end{aligned}$$

Both the infimum and supremum can be replaced by limits; this follows from Lemma 2.3.2 below. It follows from Corollary 2.3.5 that we could also directly define

$$f_\mu(T,\alpha|\beta) = f_\mu(T,\alpha\beta) - f_\mu(T,\beta),$$

as long as $f_\mu(T,\beta) > -\infty$.

We now define the relevant type of stochastic block model. If H is a finite subset of Γ, we denote by $d^H(\mu,\nu)$ the total variation distance between the marginals of μ and ν on $\mathtt{A}^H$. Our convention for the total variation distance between measures $\mu,\nu \in \mathrm{Prob}(\mathtt{A})$ is

$$\|\mu-\nu\|_{\mathrm{TV}} = \frac{1}{2}\sum_{a\in\mathtt{A}} |\mu\{a\} - \nu\{a\}|.$$

For each $k \in \mathbb{N}$ we define a pseudometric on $\mathrm{Prob}(\mathtt{A}^\Gamma)$ by

$$d_k^*(\mu,\nu) = \sum_{i\in[r]} d^{\mathrm{B}(e,k)\cup\mathrm{B}(s_i,k)}(\mu,\nu).$$

Note that $\{d_k^*\}_{k\in\mathbb{N}}$ together generate the weak* topology on $\mathrm{Prob}(\mathtt{A}^\Gamma)$. These generalize the function d_σ^* from [Bow10a], which corresponds to the case $k=0$. For $\mathcal{O} = \{\nu \in \mathrm{Prob}(\mathtt{A}^\Gamma) : d_k^*(\alpha_*^G\mu,\nu) < \varepsilon\}$ we write

$$\Omega(\sigma,\mathcal{O}) =: \Omega_k^*(\sigma,\alpha,\varepsilon) \subseteq \mathtt{A}^n.$$

Our stochastic block model is now defined as follows: given $\mathbf{y}_0 \in \mathtt{B}^n$, $\sigma_0 \in \mathrm{Hom}(\Gamma,\mathrm{Sym}(n))$, and $k\in\mathbb{N}$, let

$$\mathtt{SBM}(\sigma_0,\mathbf{y}_0,k) \coloneqq \mathrm{Unif}(\{\sigma \in \mathrm{Hom}(\Gamma,\mathrm{Sym}(n)) \;:\; d_k^*(P_{\mathbf{y}_0}^{\sigma}, P_{\mathbf{y}_0}^{\sigma_0}) = 0\}).$$

The labeling $\mathbf{y}_0$ partitions the elements of $[n]$ into $|\mathtt{B}|$ communities, and we can think of the random homomorphism σ as a random choice of directed edges between and within the communities. Certain statistics of these random edge choices are determined by the reference homomorphism σ_0; note that for $k > 0$ these statistics are more precise than those specified by a standard stochastic block model. In Section 2.2 we define weights, which are the objects used to record the relevant statistics.

2.1.1 Main results

Our main theorems show that the relative f-invariant can be interpreted as the growth rate of the expected number of ways to extend a planted good model for β to a good model for $\alpha\beta$, over a stochastic block model which has statistics determined by β and its planted model.

We first prove that if $\beta_*^G\mu$ is Markov then we can use a stochastic block model which only takes into account "one-step statistics."

Theorem F. *Let $\alpha\colon X \to \mathtt{A}$ and $\beta\colon X \to \mathtt{B}$ be finite observables, and for each n let $\mathbf{y}_n \in \mathtt{B}^n$ and $\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$ be such that*

$$\lim_{n\to\infty} d_0^*(P_{\mathbf{y}_n}^{\sigma_n}, \beta_*^G\mu) = 0.$$

Suppose that $\beta_*^G\mu$ is a Markov measure. With $\mu_n = \mathtt{SBM}(\sigma_n, \mathbf{y}_n, 0)$, we have

$$f_\mu(T, \alpha \mid \beta) = \inf_{\mathcal{O}\ni(\alpha\beta)_*^G\mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x}, \mathbf{y}_n) \in \Omega(\sigma, \mathcal{O})\}|.$$

Proposition 2.1.1. *The assumptions of Theorem F are nonvacuous; that is, for any finite observable $\beta\colon X \to \mathtt{B}$ there exist sequences $\{\mathbf{y}_n \in \mathtt{B}^n\}_{n=1}^\infty$ and $\{\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))\}_{n=1}^\infty$ such that $\lim_{n\to\infty} d_0^*(P_{\mathbf{y}_n}^{\sigma_n}, \beta_*^G\mu) = 0$.*

This follows from the fact that free group actions are "sofic," which is proven for example in [Pau11, DKP13, Pop14]. A more elementary proof is given in Section 2.4 below.

If $\beta^G_* \mu$ is not Markov, then the same formula holds with a more precise type of stochastic block model:

Theorem C. *Let $\alpha \colon X \to \mathtt{A}$ and $\beta \colon X \to \mathtt{B}$ be finite observables. Let m_n approach infinity as n goes to infinity while satisfying $m_n = o(\log \log n)$. For each n let $\mathbf{y}_n \in \mathtt{B}^n$ and $\sigma_n \in \mathrm{Hom}(G, \mathrm{Sym}(n))$ be such that*

$$d^*_{m_n}(P^{\sigma_n}_{\mathbf{y}_n}, \beta^G_* \mu) = O\big(\tfrac{1}{\log n}\big).$$

Suppose that $f_\mu(T, \beta) > -\infty$. With $\mu_n = \mathtt{SBM}(\sigma_n, \mathbf{y}_n, m_n)$,

$$f_\mu(T, \alpha \mid \beta) = \inf_{\mathcal{O} \ni (\alpha\beta)^G_* \mu} \limsup_{n \to \infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma \sim \mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x}, \mathbf{y}_n) \in \Omega(\sigma, \mathcal{O})\}|.$$

Proposition 2.1.2. *The assumptions of Theorem C are nonvacuous; that is, for any finite observable $\beta \colon X \to \mathtt{B}$ and any sequence $\{m_n \in \mathbb{N}\}_{n=1}^\infty$ approaching infinity while satisfying $m_n = o(\log \log n)$, there exist sequences $\{\mathbf{y}_n \in \mathtt{B}^n\}_{n=1}^\infty$ and $\{\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))\}_{n=1}^\infty$ such that $\lim_{n\to\infty} d^*_{m_n}(P^{\sigma_n}_{\mathbf{y}_n}, \beta^G_* \mu) = O\left(\frac{1}{\log n}\right)$.*

The expressions appearing on the right-hand sides of Theorems F and C are very comparable to Ben Hayes's definition of "relative sofic entropy in the presence" [Hay16, Definition 2.5]. Some differences are that we consider *expected* numbers of good models over *random* sofic approximations, and that Hayes takes a supremum inside the logarithm over which good model is to be extended, while we fix a sequence $\{\mathbf{y}_n\}$ of planted good models. Hayes also does not restrict to shift systems as we do here.

Using Theorem C we prove the following formula for the growth rate of the expected number of good models over a stochastic block model:

Theorem D. *Let μ_n, α, β be as in the statement of Theorem C. Then*

$$\inf_{\mathcal{O} \ni \alpha^G_* \mu} \limsup_{n \to \infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma \sim \mu_n} |\Omega(\sigma, \mathcal{O})| = \sup_{\lambda \in \mathsf{J}(\alpha^G_* \mu, \beta^G_* \mu)} f_\lambda(S, \mathtt{a} \mid \mathtt{b}).$$

Here $\mathsf{J}(\alpha_*^\Gamma\mu,\, \beta_*^\Gamma\mu)$ is the set of joinings of the Γ-systems $(\mathtt{A}^\Gamma, \alpha_*^G\mu, S)$ and $(\mathtt{B}^\Gamma, \beta_*^G\mu, S)$, *i.e.* shift-invariant probability measures on $(\mathtt{A}\times\mathtt{B})^\Gamma$ whose $\mathtt{A}^\Gamma, \mathtt{B}^\Gamma$ marginals are $\alpha_*^\Gamma\mu$, $\beta_*^\Gamma\mu$, respectively. S denotes the shift action of Γ. We use $\mathtt{a}, \mathtt{b}$ to denote the maps

$$\begin{aligned} \mathtt{a}\colon (\mathtt{A}\times\mathtt{B})^\Gamma &\to \mathtt{A} \\ \big((a_g, b_g)\big)_{g\in\Gamma} &\mapsto a_e \end{aligned} \qquad\qquad \begin{aligned} \mathtt{b}\colon (\mathtt{A}\times\mathtt{B})^\Gamma &\to \mathtt{B} \\ \big((a_g, b_g)\big)_{g\in\Gamma} &\mapsto b_e \end{aligned}$$

which observe the $\mathtt{A}$ (resp. $\mathtt{B}$) label at the identity.

Remark. The supremum is always greater than or equal to $f_\mu(T,\alpha)$, with equality attained by the product joining; this means that the expected number of good models for α over a block model with built-in good models for any β is at least the expected number of good models over a uniformly random homomorphism. It is possible for the supremum to be strictly larger, however. For example, suppose $f_\mu(T,\alpha) < 0$ and $\alpha = \beta$, and let λ be the diagonal joining. Then

$$f_\lambda(S, \mathtt{a} \mid \mathtt{b}) = 0 > f_\mu(T,\alpha).$$

See Chapter 5 for more discussion of this optimization problem.

2.1.2 Random sofic approximations

As noted above, the f-invariant is closely related to another invariant of measure-preserving systems called sofic entropy, which was introduced by Lewis Bowen in [Bow10c].

A homomorphism $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(n))$ is called (D,δ)-sofic for some finite $D\subset\Gamma$ and $\delta > 0$ if

$$|\{j \in [n] \,:\, \sigma(\gamma)j \neq j \;\forall \gamma \in D\setminus\{e\}\}| > (1-\delta)n.$$

A sequence of homomorphisms $\Sigma = \big(\sigma_n \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(n))\big)_{n\in\mathbb{N}}$ is called a sofic approximation if for every (D,δ) the homomorphism σ_n is (D,δ)-sofic for all large enough n.

The sofic entropy relative to Σ is the exponential growth rate of the number of good models over σ_n. Specifically, for any finite observable α on X we have

$$\mathrm{h}_{\Sigma,\mu}(T,\alpha) = \inf_{\mathcal{O}\ni\alpha_*^\Gamma\mu} \limsup_{n\to\infty} \frac{1}{n} \log|\Omega(\sigma_n, \mathcal{O})|.$$

This is an isomorphism invariant of the system (X, μ, T) if α is any *generating* observable, *i.e.* if the σ-algebra of the measurable space X is the coarsest one which is shift-invariant and α-measurable.

By analogy with this expression, we might call the sequences of random homomophisms appearing in expressions above "random sofic approximations." The following proposition provides further justification for this terminology.

Proposition 2.1.3. *If (μ_n) is any of the sequences appearing in Theorems F, C, and D, then for any (D, δ) there exists $\varepsilon > 0$ such that*

$$\mathbb{P}_{\sigma\sim\mu_n}\big(\sigma \text{ is } (D,\delta)\text{-sofic}\big) \geq 1 - n^{-\varepsilon n}$$

for all large enough n.

In particular, if $\sigma_1 \sim \mu_1$, $\sigma_2 \sim \mu_2$ *etc.* are independent then (σ_n) is a sofic approximation with probability 1.

Organization

In Section 2.2 we define weights and discuss some of their useful properties. In Section 2.3 we prove a few basic results about the functions f and F. Some of the results of these two sections are used in Section 2.4 to show that the assumptions of the main theorems are not vacuous. In Section 2.5 we show how the function F is related to the number of homomorphism-labeling pairs $(\sigma, \mathbf{y})$ that realize a given weight, which is the main ingredient of the proofs of Theorems F and C given in the next two sections. In Section 2.8 we show how to deduce Theorem D from Theorem C. Section 2.9 contains a proof of Proposition 2.1.3. The final section contains a proof of Lemma 2.2.3, which asserts that a weight can be approximated by a denominator-n weight with a specified marginal.

2.2 Weights

If $\alpha \colon X \to \mathtt{A}$ is a finite observable, for $a, a' \in \mathtt{A}$ and $i \in [r]$ let

$$W_\alpha(a, a'; i) = \alpha_*^{\{e, s_i\}}\mu(a, a') = \mu\{x \in X \,:\, \alpha(x) = a,\ \alpha(T_{s_i}x) = a'\}$$

and also denote

$$W_\alpha(a) = \alpha_*\mu(a).$$

For $\mathbf{x} \in \mathtt{A}^n$ and $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$ let

$$W_{\sigma,\mathbf{x}}(a, a'; i) = P_{\mathbf{x}}^{\sigma,\{e,s_i\}}(a, a')$$

and $W_{\sigma,\mathbf{x}}(a) = P_{\mathbf{x}}^{\sigma,\{e\}}(a)$.

More abstractly, any $W \in \big(\mathrm{Prob}(\mathtt{A}^2)\big)^r$ is called an $\mathtt{A}$-*weight* if

$$\sum_{a' \in \mathtt{A}} W(a, a'; i) = \sum_{a' \in \mathtt{A}} W(a', a; j)$$

for all $i, j \in [r]$ and $a \in \mathtt{A}$. For each $a \in \mathtt{A}$ we denote this common value $W(a)$. Note that the objects W_α and $W_{\sigma,\mathbf{x}}$ defined above satisfy this condition.

We say that W has denominator n if $n \cdot W(a, a'; i) \in \mathbb{N}$ for all a, a', i.

The measures $W(\cdot, \cdot; i)$ for $i \in [r]$ are called the *edge measures* of W, and $W(\cdot)$ is called the *vertex measure*.

For any alphabet $\mathtt{A}$, we use the metric on $\mathtt{A}$-weights defined by

$$\begin{aligned} d(W_1, W_2) &:= \sum_{i \in [r]} \|W_1(\cdot, \cdot; i) - W_2(\cdot, \cdot; i)\|_{\mathrm{TV}} \\ &= \frac{1}{2} \sum_{i \in [r]} \sum_{a, a' \in \mathtt{A}} |W_1(a, a'; i) - W_2(a, a'; i)|. \end{aligned}$$

We can use weights to count good models up to equivalence under the pseudometrics d_k^* using the following proposition:

Proposition 2.2.1. *If* $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$ *and* $\mathbf{x} \in \mathtt{A}^n$*, then for any observable* $\alpha \colon X \to \mathtt{A}$

$$d(W_{\sigma,\mathbf{x}^k}, W_{\alpha^k}) = d_k^*(P_{\mathbf{x}}^\sigma, \alpha_*^G \mu).$$

Note this implies also that

$$d_k^*(P_{\mathbf{x}}^\sigma, \alpha_*^G \mu) = d_0^*(P_{\mathbf{x}^k}^\sigma, (\alpha^k)_*^G \mu).$$

Proof. By definition of the distance between weights,

$$\begin{aligned} d(W_{\sigma,\mathbf{x}^k}, W_{\alpha^k}) &= \frac{1}{2}\sum_{i\in[r]} \sum_{\mathbf{a},\mathbf{a}'\in \mathtt{A}^{\mathrm{B}(e,k)}} \left| W_{\sigma,\mathbf{x}^k}(\mathbf{a},\mathbf{a}';i) - W_{\alpha^k}(\mathbf{a},\mathbf{a}';i)\right| \\ &= \frac{1}{2}\sum_{i\in[r]} \sum_{\mathbf{a},\mathbf{a}'\in \mathtt{A}^{\mathrm{B}(e,k)}} \left| \frac{1}{n}\left|\left\{ j\in[n] \;:\; \begin{array}{c} (\mathbf{x}^k)_j = \mathbf{a} \\ (\mathbf{x}^k)_{\sigma(s_i)j} = \mathbf{a}' \end{array} \right\}\right| \right. \\ &\qquad \left. - \mu \left\{ x\in X \;:\; \begin{array}{c} \alpha^k(x) = \mathbf{a} \\ \alpha^k(T_{s_i}x) = \mathbf{a}' \end{array} \right\} \right|. \end{aligned}$$

For many 'incompatible' pairs $\mathbf{a},\mathbf{a}'$, both terms will be zero: suppose $g \in \mathrm{B}(e,k)\cap \mathrm{B}(s_i,k)$, so that $gs_i^{-1} \in \mathrm{B}(e,k)$. If the second term in the absolute value is nonzero, then for some $x\in X$ we have $\alpha^k(x) = \mathbf{a}$ and $\alpha^k(T_{s_i}x) = \mathbf{a}'$, and therefore

$$\mathbf{a}'_{gs_i^{-1}} = (\alpha^k(T_{s_i}x))_{gs_i^{-1}} = \alpha(T_{gs_i^{-1}}T_{s_i}x) = \alpha(T_g x) = (\alpha^k(x))_g = \mathbf{a}_g.$$

The same argument shows that $\mathbf{a}'_{gs_i^{-1}} = \mathbf{a}_g$ for all $g\in \mathrm{B}(e,k)\cap\mathrm{B}(s_i,k)$ whenever the first term is nonzero. Therefore we can restrict the sum to pairs $\mathbf{a},\mathbf{a}'$ with $\mathbf{a}'_{gs_i^{-1}} = \mathbf{a}_g$ for all $g\in \mathrm{B}(e,k)\cap\mathrm{B}(s_i,k)$. Equivalently, we can sum over all $\mathbf{A}\in \mathtt{A}^{\mathrm{B}(e,k)\cup\mathrm{B}(s_i,k)}$ to get

$$\begin{aligned} d(W_{\sigma,\mathbf{x}^k}, W_{\alpha^k}) &= \frac{1}{2}\sum_{i\in[r]} \sum_{\mathbf{A}\in \mathtt{A}^{\mathrm{B}(e,k)\cup\mathrm{B}(s_i,k)}} \left| \frac{1}{n}\left|\left\{ j\in[n] \;:\; \left(\mathbf{x}^{\mathrm{B}(e,k)\cup\mathrm{B}(s_i,k)}\right)_j = \mathbf{A}\right\}\right| \right. \\ &\qquad \left. - \mu\left\{x\in X \;:\; \alpha^{\mathrm{B}(e,k)\cup\mathrm{B}(s_i,k)}(x) = \mathbf{A}\right\} \right| \\ &= \sum_{i\in[r]} d^{\mathrm{B}(e,k)\cup\mathrm{B}(s_i,k)}(P_{\mathbf{x}}^\sigma, \alpha_*^\Gamma \mu). \end{aligned}$$

□

It will be useful to consider the pushforward map induced by a map between alphabets: if $\pi\colon \mathtt{A}\to\mathtt{B}$ is a measurable map and W is an $\mathtt{A}$-weight, then πW is the $\mathtt{B}$-weight given by

$$\pi W(b,b';i) = \sum_{a\in\pi^{-1}\{b\}} \sum_{a'\in\pi^{-1}\{b'\}} W(a,a';i).$$

Note that this implies that the vertex measure of W is

$$\pi W(b) = \sum_{a \in \pi^{-1}\{b\}} W(a).$$

For example, let $\pi_{\mathtt{B}} \colon \mathtt{A} \times \mathtt{B} \to \mathtt{B}$ be the projection map. If W is an $\mathtt{A} \times \mathtt{B}$-weight then $\pi_{\mathtt{B}} W$ is given by

$$\pi_{\mathtt{B}} W(b_1) = \sum_{a \in \mathtt{A}} W\big((a, b_1)\big) \qquad \pi_{\mathtt{B}} W(b_1, b_2; i) = \sum_{a_1, a_2 \in \mathtt{A}} W\big((a_1, b_1), (a_2, b_2); i\big).$$

We call this the $\mathtt{B}$-marginal of W.

All weights in the present chapter will be over alphabets of the form $\mathtt{A}^{\mathrm{B}(e,k)} \times \mathtt{B}^{\mathrm{B}(e,k')}$. We use this fact to introduce some simplified notation for projections:

- π_A denotes projection onto the entire $\mathtt{A}$ factor $\mathtt{A}^{\mathrm{B}(e,k)}$; π_B is used similarly.
- For $m < k$ and $m' < k'$, $\pi_{m,m'}$ denotes projection onto $\mathtt{A}^{\mathrm{B}(e,m)} \times \mathtt{B}^{\mathrm{B}(e,m')}$.
- π_m denotes the projection $\mathtt{A}^{\mathrm{B}(e,k)} \to \mathtt{A}^{\mathrm{B}(e,m)}$, except that if $m = 0$ we write π_e.

We define $F(W)$ for an abstract weight W by

$$F(W) = (1 - 2r)\mathrm{H}\big(W(\cdot)\big) + \sum_{i \in [r]} \mathrm{H}\big(W(\cdot, \cdot; i)\big)$$

where H is the Shannon entropy. Note that this is consistent with the above definitions in that, for example,

$$F(W_\alpha) = F_\mu(T, \alpha).$$

We can revisit the definition of our version of the stochastic block model using weights: Let $H \subset \Gamma$ and let W be a denominator-n $\mathtt{B}^{\mathrm{B}(e,k)}$-weight. Suppose there exist $\mathbf{y} \in \mathtt{B}^n$ and $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$ such that $W = W_{\sigma, \mathbf{y}^k}$. Then

$$\mathtt{SBM}(\sigma, \mathbf{y}, k) = \mathrm{Unif}(\{\sigma' \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n)) \,:\, W_{\sigma', \mathbf{y}^k} = W\}),$$

so we can also denote this distribution by $\mathtt{SBM}(\mathbf{y}, W)$. Specifying the distribution by a weight rather than a specific homomorphism will occasionally be more convenient.

2.2.1 Constructing weights and good models

We borrow the first result of this type from [Bow10a]; it allows us to find a denominator-n approximation to a given weight.

Lemma 2.2.2 (Lemma 2.3 of [Bow10a])**.** *There is a constant C such that for any $\mathtt{A}$-weight W there is a denominator-n $\mathtt{A}$-weight within distance $C|\mathtt{A}|^2 r/n$ of W.*

The following lemma allows us not only to construct a denominator-n approximation to a given weight, but also to specify a marginal of this approximation:

Lemma 2.2.3. *Let W be an $\mathtt{A} \times \mathtt{B}$-weight. If $W_{\mathtt{B}}$ is a $\mathtt{B}$-weight of denominator n with $d(W_{\mathtt{B}}, \pi_{\mathtt{B}} W) < \delta$ then there is an $\mathtt{A} \times \mathtt{B}$-weight $W_{\mathtt{AB}}$ with denominator n such that $\pi_{\mathtt{B}} W_{\mathtt{AB}} = W_{\mathtt{B}}$ and $d(W_{\mathtt{AB}}, W) < 265 r(|\mathtt{A} \times \mathtt{B}|^2/n + \delta)$.*

The construction is fairly involved, so is postponed to Section 2.10. The constant 265 is not intended to be optimal.

The definition of a weight $W_{\sigma, \mathbf{x}^k}$ in terms of a homomorphism σ and a labeling $\mathbf{x}$ is straightforward. However, we will also need to know whether a given weight can be realized in this way. The next two results address this inverse problem.

Proposition 2.2.4. *If W is a denominator-n $\mathtt{A}$-weight, then there exist $\mathbf{x} \in \mathtt{A}^n$ and $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(n))$ such that $W = W_{\sigma, \mathbf{x}}$.*

Proof. This is implied by Proposition 2.1 of [Bow10a]. □

Unfortunately, this does not imply that for every denominator-n $\mathtt{A}^{\mathrm{B}(e,k)}$-weight W there is some $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(n))$ and $\mathbf{x} \in \mathtt{A}^n$ such that $W = W_{\sigma, \mathbf{x}^k}$; instead it provides $\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n$ such that $W = W_{\sigma, \mathbf{X}}$.

However, if we already know that W is close to a weight of the form W_{α^k} for some observable α, then the following proposition shows that W is also close to a weight of the form $W_{\sigma, \mathbf{x}^k}$.

Proposition 2.2.5. *Let* $\alpha\colon X \to \mathtt{A}$*,* $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$*, and* $\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n$ *be such that* $d(W_{\sigma,\mathbf{X}}, W_{\alpha^k}) \leq \varepsilon$ *for some* $\varepsilon \geq 0$*. Writing* $\mathbf{x} = \pi_e \mathbf{X} \in \mathtt{A}^n$*, we have*

$$d(W_{\sigma,\mathbf{X}}, W_{\sigma,\mathbf{x}^k}) \leq 2r|\mathrm{B}(e,k)|\varepsilon.$$

An immediate consequence is that $\mathbf{X} \in \Omega_0^*(\sigma, \alpha^k, \varepsilon)$ implies $\pi_e \mathbf{X} \in \Omega_k^*(\sigma, \alpha, c\varepsilon)$ where $c = 1 + 2r|\mathrm{B}(e,k)|$; cf. Claim 2 in the proof of Proposition 3.2 of [Bow10a].

Proof. Claim 4 in the proof of Proposition 3.2 of [Bow10a] implies that

$$|\{j \in [n] \,:\, \mathbf{X}(j) \neq \mathbf{x}^k(j)\}| \leq n|\mathrm{B}(e,k)|\varepsilon.$$

It follows that for any $i \in [r]$

$$\begin{aligned}
&|\{j \in [n] \,:\, \mathbf{X}^{\{e,s_i\}}(j) \neq (\mathbf{x}^k)^{\{e,s_i\}}(j)\}| \\
&\qquad \leq |\{j \in [n] \,:\, \mathbf{X}(j) \neq \mathbf{x}^k(j)\}| + |\{j \in [n] \,:\, \mathbf{X}(\sigma(s_i)j) \neq \mathbf{x}^k(\sigma(s_i)j)\}| \\
&\qquad \leq 2n|\mathrm{B}(e,k)|\varepsilon,
\end{aligned}$$

so

$$\begin{aligned}
d(W_{\sigma,\mathbf{X}}, W_{\sigma,\mathbf{x}^k}) &= \sum_{i\in[r]} \left\| \left(\mathbf{X}^{\{e,s_i\}}\right)_* \mathrm{Unif}(n) - \left((\mathbf{x}^k)^{\{e,s_i\}}\right)_* \mathrm{Unif}(n) \right\|_{\mathrm{TV}} \\
&\leq \sum_{i\in[r]} 2|\mathrm{B}(e,k)|\varepsilon = 2r|\mathrm{B}(e,k)|\varepsilon.
\end{aligned}$$

□

2.3 Properties of F and f

Lemma 2.3.1 (Continuity as weight function)**.** *If* W_1, W_2 *are* $\mathtt{A}$*-weights with* $d(W_1, W_2) \leq \varepsilon \leq 1$ *then*

$$|F(W_1) - F(W_2)| \leq 4r\big(\mathrm{H}(\varepsilon) + \varepsilon \log_2|\mathtt{A}|\big).$$

where $\mathrm{H}(p)$ *denotes the entropy of the probability measure* $(p, 1-p) \in \mathrm{Prob}(\{0,1\})$*.*

Proof. We use Fano's inequality in the following form (Equation (2.139) of [CT06]): suppose X, Y are $\mathtt{A}$-valued random variables defined on the same probability space and let $p_e = \mathbb{P}(X \neq Y)$ be their probability of disagreement. Then

$$\mathrm{H}(X \mid Y) \leq \mathrm{H}(p_e) + p_e \log|\mathtt{A}|.$$

Using the chain rule and nonnegativity of Shannon entropy, we can deduce that

$$|\mathrm{H}(X) - \mathrm{H}(Y)| \leq \mathrm{H}(p_e) + p_e \log|\mathtt{A}|.$$

Let $\mu_1, \mu_2 \in \mathrm{Prob}(\mathtt{A})$ be the respective distributions of X_1, X_2. Because $\|\mu_1 - \mu_2\|_{\mathrm{TV}}$ is the minimum value of $\mathbb{P}(X \neq Y)$ over all possible couplings, if $\|\mu_1 - \mu_2\|_{\mathrm{TV}} < \varepsilon$ then

$$|\mathrm{H}(\mu_1) - \mathrm{H}(\mu_2)| \leq \mathrm{H}(\varepsilon) + \varepsilon \log_2|\mathtt{A}|.$$

The assumed bound $d(W_1, W_2) \leq \varepsilon$ implies that each vertex and edge measure of W_1 is within total variation distance ε of its counterpart in W_2, so

$$\begin{aligned} |F(W_1) - F(W_2)| &\leq |1 - 2r| \cdot \left|\mathrm{H}\big(W_1(\cdot)\big) - \mathrm{H}\big(W_2(\cdot)\big)\right| \\ &\qquad + \sum_{i \in [r]} \left|\mathrm{H}\big(W_1(\cdot, \cdot; i)\big) - \mathrm{H}\big(W_2(\cdot, \cdot; i)\big)\right| \\ &\leq (2r - 1)\left(\mathrm{H}(\varepsilon) + \varepsilon \log_2|\mathtt{A}|\right) \\ &\qquad + r \cdot \left(\mathrm{H}(\varepsilon) + \varepsilon \log_2|\mathtt{A}|^2\right) \\ &\leq 4r\big(\mathrm{H}(\varepsilon) + \varepsilon \log_2|\mathtt{A}|\big). \end{aligned}$$

□

Let $\alpha \colon X \to \mathtt{A}$ and $\beta \colon X \to \mathtt{B}$ be observables. We say that β is a *coarsening* of α if each part of the partition of X induced by β is a union of parts of the partition induced by α (up to null sets). Equivalently, there is some function $g \colon \mathtt{A} \to \mathtt{B}$ such that $\beta = g \circ \alpha$ almost surely. In this situation we can also call α a refinement of β.

A useful property of the Shannon entropy $\mathrm{H}_\mu(\alpha)$ is monotonicity under refinement. The function F does not share this property, but it is monotone under the following particular kind of refinement introduced in [Bow10b]:

We say that β is a *simple splitting* of α if there is some $s \in \{s_1^{\pm 1}, \ldots, s_r^{\pm 1}\}$ and a coarsening $\tilde{\alpha}$ of α such that, up to null sets, the partition induced by β is the coarsest common refinement of the partitions induced by α and $\tilde{\alpha} \circ T_s$.

We say that β is a *splitting* of α if there are observables $\alpha = \beta_0, \beta_1, \ldots, \beta_n = \beta$ such that β_i is a simple splitting of β_{i-1} for $i = 1, 2, \ldots, n$. We will use the following monotonicity properties of the relative version of F:

Lemma 2.3.2 (Monotonicity under splitting)**.**

1. *If α_1 is a splitting of α_2 then $F(\alpha_1|\beta) \leq F(\alpha_2|\beta)$.*
2. *If β_1 is a splitting of β_2 then $F(\alpha|\beta_1) \geq F(\alpha|\beta_2)$.*

Proof. 1. This is essentially Proposition 5.1 of [Bow10b]; conditioning on β makes no difference to the proof.

2. The proof is based on the proof of Part 1, but in place of the chain rule for conditional entropy we use the following bound:

$$\begin{aligned} \mathrm{H}(\alpha \mid \beta_2) &\leq \mathrm{H}(\alpha, \beta_1 \mid \beta_2) && \text{(monotonicity)} \\ &= \mathrm{H}(\beta_1 \mid \beta_2) + \mathrm{H}(\alpha \mid \beta_1, \beta_2) && \text{(chain rule)} \\ &\leq \mathrm{H}(\beta_1 \mid \beta_2) + \mathrm{H}(\alpha \mid \beta_1) && \text{(monotonicity).} \end{aligned}$$

We will also use the following consequence of the previous bound:

$$\begin{aligned} &\mathrm{H}(\alpha^{\{e,s_i\}} \mid \beta_1^{\{e,s_i\}}) - \mathrm{H}(\alpha^{\{e,s_i\}} \mid \beta_2^{\{e,s_i\}}) \\ &\qquad \geq -\mathrm{H}(\beta_1^{\{e,s_i\}} \mid \beta_2^{\{e,s_i\}}) && \text{(previous bound)} \\ &\qquad \geq -\big(\mathrm{H}(\beta_1^{\{s_i\}} \mid \beta_2^{\{e,s_i\}}) + \mathrm{H}(\beta_1 \mid \beta_2^{\{e,s_i\}})\big) && \text{(subadditivity)} \\ &\qquad = -\big(\mathrm{H}(\beta_1 \mid \beta_2^{\{e,s_i^{-1}\}}) + \mathrm{H}(\beta_1 \mid \beta_2^{\{e,s_i\}})\big) && \text{(T-invariance of μ).} \end{aligned}$$

It suffices to check the case where β_1 is a simple splitting of β_2: let $t \in \{s_1^{\pm 1}, \ldots, s_r^{\pm 1}\}$ and let $\tilde{\beta}$ be a coarsening of β_2 such that the partition induced by β_1 is the same as the coarsest common refinement of the partitions induced by β_2 and $\tilde{\beta} \circ T_t$ up to null sets. Then, using the two bounds just derived,

$$
\begin{aligned}
F(\alpha|\beta_1) - F(\alpha|\beta_2) &= (1-2r)\left(\mathrm{H}(\alpha|\beta_1) - \mathrm{H}(\alpha|\beta_2)\right) \\
&\quad + \sum_{i\in[r]} \left(\mathrm{H}(\alpha^{\{e,s_i\}}|\beta_1^{\{e,s_i\}}) - \mathrm{H}(\alpha^{\{e,s_i\}}|\beta_1^{\{e,s_i\}}) \right) \\
&\geq (1-2r)\left(-\mathrm{H}(\beta_1|\beta_2)\right) - \sum_{i\in[r]} \left(\mathrm{H}(\beta_1 \mid \beta_2^{\{e,s_i^{-1}\}}) + \mathrm{H}(\beta_1 \mid \beta_2^{\{e,s_i\}}) \right) \\
&= (2r-1)\mathrm{H}(\beta_1|\beta_2) - \sum_{s\in\{s_1^{\pm 1}\dots s_r^{\pm 1}\}} \mathrm{H}(\beta_1 \mid \beta_2^{\{e,s\}})
\end{aligned}
$$

But

$$
\mathrm{H}(\beta_1 \mid \beta_2^{\{e,t\}}) \leq \mathrm{H}(\beta_1 \mid \beta_2 \tilde{\beta}^{\{t\}}) = 0,
$$

so we can remove the t term from the sum to get

$$
\begin{aligned}
F(\alpha|\beta_1) - F(\alpha|\beta_2) &\geq (2r-1)\mathrm{H}(\beta_1|\beta_2) - \sum_{s\in\{s_1^{\pm 1}\dots s_r^{\pm 1}\}\setminus\{t\}} \mathrm{H}(\beta_1 \mid \beta_2^{\{e,s\}}) \\
&= \sum_{s\in\{s_1^{\pm 1}\dots s_r^{\pm 1}\}\setminus\{t\}} \left(\mathrm{H}(\beta_1|\beta_2) - \mathrm{H}(\beta_1 \mid \beta_2^{\{e,s\}}) \right) \\
&\geq 0.
\end{aligned}
$$

□

One corollary is the following convenient formula:

Corollary 2.3.3. *Let α, β be finite observables such that $\beta_*^G \mu$ is a Markov measure. Then $F_\mu(T, \alpha^{k_1} \mid \beta^{k_2})$ is independent of k_2. In particular,*

$$
f_\mu(T, \alpha \mid \beta) = \inf_k F_\mu(T, \alpha^k \mid \beta).
$$

Proof. By the previous proposition, for any $k \leq k_2$ we have

$$
F_\mu(T, \alpha^{k_1} \mid \beta^k) \leq F_\mu(T, \alpha^{k_1} \mid \beta^{k_2}).
$$

On the other hand, by Theorem 6.1 of [Bow10d] $F_\mu(T, \beta^k) = F_\mu(T, \beta^{k_2})$ so

$$
F_\mu(T, \alpha^{k_1} \mid \beta^k) = F_\mu(T, \alpha^{k_1}\beta^k) - F_\mu(T, \beta^{k_2}).
$$

Applying monotonicity under splitting to the first term on the right gives

$$F_\mu(T, \alpha^{k_1} \mid \beta^k) \geq F_\mu(T, \alpha^{k_1}\beta^{k_2}) - F_\mu(T, \beta^{k_2}) = F_\mu(T, \alpha^{k_1} \mid \beta^{k_2}).$$

This establishes independence of k_2; the formula for f follows. □

Proposition 2.3.4. *Let α, β be finite observables. Then for any $k \in \mathbb{N}$,*

$$F_\mu(T, \alpha^k \mid \beta) \leq \mathrm{H}_\mu\big(\alpha \mid \beta\big).$$

It follows that

$$f_\mu(T, \alpha \mid \beta) \leq \mathrm{H}_\mu\big(\alpha \mid \beta\big).$$

Proof. By Lemma 2.3.2, $F_\mu(T, \alpha^k \mid \beta) \leq F_\mu(T, \alpha \mid \beta)$. Using elementary properties of Shannon entropy, we have

$$\begin{aligned} F_\mu(T, \alpha \mid \beta) &= (1-2r)\mathrm{H}_\mu(\alpha \mid \beta) + \sum_{i \in [r]} \mathrm{H}_\mu\big(\alpha^{\{e,s_i\}} \mid \beta^{\{e,s_i\}}\big) \\ &\leq (1-2r)\mathrm{H}_\mu(\alpha \mid \beta) + \sum_{i \in [r]} \left[\mathrm{H}_\mu\big(\alpha \mid \beta^{\{e,s_i\}}\big) + \mathrm{H}_\mu\big(\alpha^{\{s_i\}} \mid \beta^{\{e,s_i\}}\big)\right] \\ &\leq (1-2r)\mathrm{H}_\mu(\alpha \mid \beta) + \sum_{i \in [r]} \left[\mathrm{H}_\mu\big(\alpha \mid \beta\big) + \mathrm{H}_\mu\big(\alpha^{\{s_i\}} \mid \beta^{\{s_i\}}\big)\right]. \end{aligned}$$

By T-invariance of μ we have

$$\mathrm{H}_\mu\big(\alpha^{\{s_i\}} \mid \beta^{\{s_i\}}\big) = \mathrm{H}_\mu(\alpha \mid \beta),$$

so the first inequality follows.

For any $k_1, k_2 \in \mathbb{N}$ this gives

$$F_\mu(T, \alpha^{k_1} \mid \beta^{k_2}) \leq \mathrm{H}_\mu(\alpha \mid \beta^{k_2}) \leq \mathrm{H}_\mu(\alpha \mid \beta),$$

so the second inequality follows upon taking the supremum over k_2 then the infimum over k_1. □

We can use this bound to give a proof of the chain rule for the relative f-invariant, a version of which first appeared in [Bow10d] (there it is called the Abramov-Rokhlin formula; see also [BG13]):

Corollary 2.3.5 (Chain rule)**.**

$$f_\mu(T, \alpha\beta) = f_\mu(T, \alpha \mid \beta) + f_\mu(T, \beta).$$

Proof. By definition of the relative version of F and the chain rule for conditional entropy, for each k_1, k_2 we have

$$F_\mu(T, \alpha^{k_1}\beta^{k_2}) = F_\mu(T, \alpha^{k_1} \mid \beta^{k_2}) + F_\mu(T, \beta^{k_2}).$$

By Lemma 2.3.2 each term is monotone in k_2, so the limits as $k_2 \to \infty$ exist. By Proposition 2.3.4 all terms are bounded above (recall we only consider finite observables, so in particular all observables have finite entropy), so we can split the limit across the sum on the right to get

$$\lim_{k_2\to\infty} F_\mu(T, \alpha^{k_1}\beta^{k_2}) = \lim_{k_2\to\infty} F_\mu(T, \alpha^{k_1} \mid \beta^{k_2}) + f_\mu(T, \beta).$$

Taking k_1 to infinity gives the result. □

2.4 Non-vacuity of main theorems

2.4.1 Theorem F

Here we prove Proposition 2.1.1, which asserts the nonvacuity of Theorem F. Given $\beta\colon X \to \mathtt{B}$, we need to show that there exist $\mathbf{y}_n \in \mathtt{B}^n$ and $\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$ such that

$$\lim_{n\to\infty} d_0^*(P^{\sigma_n}_{\mathbf{y}_n}, \beta^G_*\mu) = 0.$$

By Lemma 2.2.2, there is a sequence $\{W_n\}_{n=1}^\infty$ of $\mathtt{B}$-weights such that W_n has denominator n for each n and $d(W_n, W_\beta) = o(1)$. By Proposition 2.2.4, for each n we can pick $\mathbf{y}_n, \sigma_n$ such that $W_{\sigma_n,\mathbf{y}_n} = W_n$. Since $d_0^*(P^{\sigma_n}_{\mathbf{y}_n}, \beta^G_*\mu) = d(W_{\sigma_n,\mathbf{y}_n}, W_\beta)$, these suffice.

2.4.2 Theorems C and D

Here we prove Proposition 2.1.2, which asserts the nonvacuity of Theorem C (and by extension Theorem D, since the assumptions are the same).

Let m_n approach infinity as n approaches infinity while satisfying $m_n = o(\log\log n)$ and let $\beta\colon X \to \mathtt{B}$ be a finite observable. We need to show that there exist $\mathbf{y}_n \in \mathtt{B}^n$ and $\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$ such that $d^*_{m_n}(P^{\sigma_n}_{\mathbf{y}_n}, \beta^G_* \mu) = O(\frac{1}{\log n})$.

By Lemma 2.2.2, there is a sequence $\{W_n\}_{n=1}^\infty$ of weights such that W_n is a denominator-n $\mathtt{B}^{\mathrm{B}(e,m_n)}$-weight for each n and $d(W_n, W_{\beta^{m_n}}) = O(\frac{|\mathtt{B}^{\mathrm{B}(e,m_n)}|^2}{n})$. By Proposition 2.2.4, for each n we can pick $\mathbf{Y}_n, \sigma_n$ such that $W_{\sigma_n, \mathbf{Y}_n} = W_n$. Let $\mathbf{y}_n = \pi_e \mathbf{Y}_n$. By Proposition 2.2.5,

$$d^*_{m_n}(P^{\sigma_n}_{\mathbf{y}_n}, \beta^G_* \mu) = d(W_{\sigma_n, \mathbf{y}_n^{m_n}}, W_{\beta^{m_n}}) = O\left(|\mathrm{B}(e,m_n)| \cdot \frac{|\mathtt{B}^{\mathrm{B}(e,m_n)}|^2}{n}\right) = O\left(\frac{1}{\log n}\right).$$

2.5 Counting Lemmas

For a $\mathtt{B}$-weight W, let $Z_n(W)$ denote the number of pairs $(\sigma, \mathbf{y}) \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n)) \times \mathtt{B}^n$ such that $W_{\sigma,\mathbf{y}} = W$.

Proposition 2.5.1. *If W is a $\mathtt{B}$-weight with denominator n then*

$$(3\sqrt{n})^{-r|\mathtt{B}|^2} \leq \frac{Z_n(W)}{e^{F(W)n}(n!)^r n^{(1-r)/2}} \leq (3\sqrt{n})^{r|\mathtt{B}|^2}.$$

Proof. We write

$$Z_n(W) = \sum_\sigma |\{\mathbf{y} \in \mathtt{B}^n \,:\, W_{\sigma,\mathbf{y}} = W\}| = (n!)^r \mathop{\mathbb{E}}_\sigma |\{\mathbf{y} \in \mathtt{B}^n \,:\, W_{\sigma,\mathbf{y}} = W\}|.$$

where $\mathbb{E}_\sigma$ denotes the expectation over a uniform choice of $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$.

Proposition 2.1 of [Bow10a] states that

$$\mathop{\mathbb{E}}_\sigma |\{\mathbf{y} \in \mathtt{B}^n \,:\, W_{\sigma,\mathbf{y}} = W\}| = \frac{n!^{1-r} \prod_{b \in \mathtt{B}} (nW(b))!^{2r-1}}{\prod_{i=1}^r \prod_{b,b' \in \mathtt{B}} (nW(b,b';i))!}.$$

Lemma 2.2 of the same paper gives an estimate of this quantity, but for our purposes we need to be more careful about how the estimate depends on the size of the alphabet.

We use the version of Stirling's approximation

$$k^{k+1/2}e^{-k} \le k! \le 3 \cdot k^{k+1/2}e^{-k},$$

valid for $k \ge 1$. To estimate the products that appear in the expectation, we will need to omit all factors which equal $0! = 1$ since Stirling's approximation is not valid for these. To do this carefully, let

$$\mathtt{B}' = \{b \in \mathtt{B} \,:\, W(b) \ne 0\}$$

and for each $i \in [r]$ let

$$\mathtt{B}'_i = \{(b, b') \in \mathtt{B}^2 \,:\, W(b, b'; i) \ne 0\}.$$

For the numerator of the above expectation we get

$$\begin{aligned} n!^{1-r} \prod_{b \in \mathtt{B}'} (nW(b))!^{2r-1} &\le (3n^{n+1/2}\, e^{-n})^{1-r} \prod_{b \in \mathtt{B}'} \left(3(nW(b))^{nW(b)+1/2} e^{-nW(b)}\right)^{2r-1} \\ &= 3^{1-r+|\mathtt{B}'|(2r-1)}\, n^{rn+1/2-r/2+(2r-1)|\mathtt{B}'|/2} \\ &\qquad \times e^{-rn+(2r-1)[n\sum_{b\in\mathtt{B}'} W(b)\log W(b) + \frac{1}{2}\sum_{b\in\mathtt{B}'} \log W(b)]} \end{aligned}$$

and a lower bound which is identical except missing the first factor. For the denominator, let $S = \sum_{i\in[r]} |\mathtt{B}'_i|$. We get

$$\begin{aligned} \prod_{i=1}^{r} \prod_{(b,b')\in\mathtt{B}'_i} (nW(b,b';i))! &\le \prod_{i=1}^{r} \prod_{(b,b')\in\mathtt{B}'_i} 3(nW(b,b';i))^{nW(b,b';i)+1/2} e^{-nW(b,b';i)} \\ &= 3^S\, n^{nr+S/2} \\ &\qquad \times e^{n\sum_i \sum_{b,b'} W(b,b';i)\log W(b,b';i) + \frac{1}{2}\sum_{i,b,b'} \log W(b,b';i) - nr}, \end{aligned}$$

and again we have a lower bound which is identical except missing the first factor 3^S. Therefore the quotient is bounded above by

$$3^{1-r+|\mathtt{B}'|(2r-1)}\, n^{(1-r)/2+(2r-1)|\mathtt{B}'|/2-S/2}\, e^{-nF(W)+(2r-1)\frac{1}{2}\sum_b \log W(b) - \frac{1}{2}\sum_{i,b,b'} \log W(b,b';i)}$$

and below by

$$3^{-S}\, n^{(1-r)/2+(2r-1)|\mathtt{B}'|/2-S/2}\, e^{-nF(W)+(2r-1)\frac{1}{2}\sum_b \log W(b) - \frac{1}{2}\sum_{i,b,b'} \log W(b,b';i)}.$$

Since W has denominator n, we have

$$0 \geq (2r-1)\frac{1}{2}\sum_{b\in\mathtt{B}'}\log W(b) \geq (2r-1)\frac{1}{2}\sum_{b\in\mathtt{B}'}\log\frac{1}{n} = -\frac{2r-1}{2}|\mathtt{B}'|\log n$$

and

$$0 \leq -\frac{1}{2}\sum_i \sum_{(b,b')\in B_i'}\log W(b,b';i) \leq -\frac{1}{2}\sum_i \sum_{(b,b')\in\mathtt{B}_i'}\log\frac{1}{n} = \frac{S}{2}\log n.$$

Therefore $Z_n(W)$ satisfies

$$3^{-S}n^{((1-r)-S)/2}e^{F(W)n}(n!)^r \leq Z_n(W) \leq 3^{1-r+|\mathtt{B}'|(2r-1)}n^{((1-r)+(2r-1)|\mathtt{B}'|)/2}e^{F(W)n}(n!)^r.$$

Since $S \leq r|\mathtt{B}|^2$ and $|\mathtt{B}'| \leq |\mathtt{B}|$, we conclude that

$$3^{-r|\mathtt{B}|^2}n^{\left((1-r)-r|\mathtt{B}|^2\right)/2}e^{F(W)n}(n!)^r \leq Z_n(W) \leq 3^{1-r+|\mathtt{B}|(2r-1)}n^{((1-r)+(2r-1)|\mathtt{B}|)/2}e^{F(W)n}(n!)^r,$$

and the stated inequality follows. □

The following proposition establishes the connection between the relative version of F and expected numbers of good models over stochastic block models.

Proposition 2.5.2. *Given any denominator-n $(\mathtt{A}\times\mathtt{B}^{\mathrm{B}(e,k)})$-weight $W_{\mathtt{AB}}$, let $W_{\mathtt{B}}$ denote the $\mathtt{B}^{\mathrm{B}(e,k)}$-weight $\pi_{\mathtt{B}}W_{\mathtt{AB}}$. Let $\mathbf{y}\in\mathtt{B}^n$ be a fixed labeling with $p_{\mathbf{y}} = \pi_e W_{\mathtt{B}}(\cdot)$, and let*

$$\mu = \mathtt{SBM}(\mathbf{y}, W_{\mathtt{B}}) = \mathrm{Unif}(\{\sigma\in\mathrm{Hom}(\Gamma,\mathrm{Sym}(n))\,:\,W_{\sigma,\mathbf{y}^k} = W_{\mathtt{B}}\}),$$

assuming $W_{\mathtt{B}}$ is such that the desired support is nonempty. Then

$$\mathcal{E} := \mathop{\mathbb{E}}_{\sigma\sim\mu}\left|\{\mathbf{x}\in\mathtt{A}^n\,:\,W_{\sigma,(\mathbf{x},\mathbf{y}^k)} = W_{\mathtt{AB}}\}\right| = \frac{Z_n(W_{\mathtt{AB}})}{Z_n(W_{\mathtt{B}})}.$$

In particular,

$$\frac{\mathcal{E}}{e^{n(F(W_{\mathtt{AB}})-F(W_{\mathtt{B}}))}} \in \left((9n)^{-r|\mathtt{B}|^2(|\mathtt{A}|^2+1)},\ (9n)^{r|\mathtt{B}|^2(|\mathtt{A}|^2+1)}\right).$$

Lemma 2.5.3. *Let $W_{\mathtt{AB}}$ be a $\mathtt{A}\times\mathtt{B}^{\mathrm{B}(e,k)}$ weight of denominator n. Then*

$$\left|\{(\sigma,\mathbf{x},\mathbf{y})\,:\,W_{\sigma,(\mathbf{x},\mathbf{y}^k)} = W_{\mathtt{AB}}\}\right| \in \left\{0,\ \left|\{(\sigma,\mathbf{x},\mathbf{Y})\,:\,W_{\sigma,(\mathbf{x},\mathbf{Y})} = W_{\mathtt{AB}}\}\right|\right\}.$$

Proof. Suppose $\left|\{(\sigma, \mathbf{x}, \mathbf{y}) \,:\, W_{\sigma,(\mathbf{x},\mathbf{y}^k)} = W_{\mathtt{AB}}\}\right| \neq 0$; we then need to show

$$\left|\{(\sigma, \mathbf{x}, \mathbf{y}) \,:\, W_{\sigma,(\mathbf{x},\mathbf{y}^k)} = W_{\mathtt{AB}}\}\right| = \left|\{(\sigma, \mathbf{x}, \mathbf{Y}) \,:\, W_{\sigma,(\mathbf{x},\mathbf{Y})} = W_{\mathtt{AB}}\}\right|.$$

The inequality $\leq$ is clear, since we have an injection $(\sigma, \mathbf{x}, \mathbf{y}) \mapsto (\sigma, \mathbf{x}, \mathbf{y}^k)$.

The converse inequality holds because $(\sigma, \mathbf{x}, \mathbf{Y}) \mapsto (\sigma, \mathbf{x}, \mathbf{Y}_e)$ in an injection from the set on the right to the set on the left. This follows from the remark at the beginning of the proof of Proposition 2.2.5. □

Proof of Proposition. Let

$$\tilde{\mu} = \mathrm{Unif}(\{(\sigma, \tilde{\mathbf{y}}) \,:\, W_{\sigma,\tilde{\mathbf{y}}^k} = W_{\mathtt{B}}\});$$

then, since $\left|\{\mathbf{x} \in \mathtt{A}^n \,:\, W_{\sigma,(\mathbf{x},\tilde{\mathbf{y}}^k)} = W_{\mathtt{AB}}\}\right|$ is independent of the choice of $\tilde{\mathbf{y}}$ with $p_{\tilde{\mathbf{y}}} = \pi_e W_{\mathtt{B}}(\cdot)$,

$$\begin{aligned}
\mathcal{E} &= \mathop{\mathbb{E}}_{(\sigma,\tilde{\mathbf{y}})\sim\tilde{\mu}} \left|\{\mathbf{x} \in \mathtt{A}^n \,:\, W_{\sigma,(\mathbf{x},\tilde{\mathbf{y}}^k)} = W_{\mathtt{AB}}\}\right| \\
&= \frac{\sum_{\sigma,\tilde{\mathbf{y}}} \left|\{\mathbf{x} \in \mathtt{A}^n \,:\, W_{\sigma,(\mathbf{x},\tilde{\mathbf{y}}^k)} = W_{\mathtt{AB}}\}\right|}{\left|\{(\sigma, \tilde{\mathbf{y}}) \,:\, W_{\sigma,\tilde{\mathbf{y}}^k} = W_{\mathtt{B}}\}\right|} \\
&= \frac{\left|\{(\sigma, \mathbf{x}, \tilde{\mathbf{y}}) \,:\, W_{\sigma,(\mathbf{x},\tilde{\mathbf{y}}^k)} = W_{\mathtt{AB}}\}\right|}{\left|\{(\sigma, \tilde{\mathbf{y}}) \,:\, W_{\sigma,\tilde{\mathbf{y}}^k} = W_{\mathtt{B}}\}\right|} \\
&= \frac{\left|\{(\sigma, \mathbf{x}, \mathbf{Y}) \,:\, W_{\sigma,(\mathbf{x},\mathbf{Y})} = W_{\mathtt{AB}}\}\right|}{\left|\{(\sigma, \mathbf{Y}) \,:\, W_{\sigma,\mathbf{Y}} = W_{\mathtt{B}}\}\right|} && \text{(previous lemma)} \\
&= \frac{Z_n(W_{\mathtt{AB}})}{Z_n(W_{\mathtt{B}})}.
\end{aligned}$$

Note that our assumption that the intended support of μ is nonempty allows us to rule out the “0” case in the application of the lemma.

The rest of the result then follows from our estimates on Z_n in Proposition 2.5.1. □

2.6 Proof of Theorem F

2.6.1 Upper bound

Note that we will not rely on the Markov assumption for the upper bound.

For each $k \in \mathbb{N}$,

$$\begin{aligned}
\inf_{\mathcal{O} \ni (\alpha\beta)^G_* \mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} &|\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x}, \mathbf{y}_n) \in \Omega(\sigma, \mathcal{O})\}| \\
&\leq \inf_{\varepsilon} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x}, \mathbf{y}_n) \in \Omega_k^*(\sigma, \alpha\beta, \varepsilon)\}| \\
&= \inf_{\varepsilon} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x}^k, \mathbf{y}_n^k) \in \Omega_0^*(\sigma, (\alpha\beta)^k, \varepsilon)\}| \\
&\leq \inf_{\varepsilon} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, (\mathbf{X}, \mathbf{y}_n^k) \in \Omega_0^*(\sigma, (\alpha\beta)^k, \varepsilon)\}|.
\end{aligned}$$

Write

$$\begin{aligned}
\mathcal{E}_k(n,\varepsilon) &\coloneqq \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, (\mathbf{X}, \mathbf{y}_n^k) \in \Omega_0^*(\sigma, (\alpha\beta)^k, \varepsilon)\}| \\
&= \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, d\big(W_{\sigma,(\mathbf{X},\mathbf{y}_n^k)}, W_{(\alpha\beta)^k}\big) < \varepsilon\}|
\end{aligned}$$

and assume that n is large enough that $m_n \geq k$.

Writing $\mathcal{W}_n(\alpha\beta, k, \varepsilon)$ for the set of all denominator-n weights W with $d(W, W_{(\alpha\beta)^k}) < \varepsilon$,

$$\begin{aligned}
\mathcal{E}_k(n,\varepsilon) &= \mathop{\mathbb{E}}_{\sigma\sim\mu_n} \sum_{W \in \mathcal{W}_n(\alpha\beta,k,\varepsilon)} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, W_{\sigma,(\mathbf{X},\mathbf{y}_n^k)} = W\}| \\
&= \sum_{W \in \mathcal{W}_n(\alpha\beta,k,\varepsilon)} \Big[\mathop{\mathbb{E}}_{\sigma\sim\mu_n} \big[|\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, W_{\sigma,(\mathbf{X},\mathbf{y}_n^k)} = W\}| \big| W_{\sigma,\mathbf{y}_n^k} = \pi_{\mathtt{B}} W \big] \\
&\qquad\qquad \times \mathop{\mathbb{P}}_{\sigma\sim\mu_n} (W_{\sigma,\mathbf{y}_n^k} = \pi_{\mathtt{B}} W) \Big]
\end{aligned}$$

since if $W_{\sigma,\mathbf{y}_n^k} \neq \pi_{\mathtt{B}} W$ then $W_{\sigma,(\mathbf{X},\mathbf{y}_n^k)} \neq W$. But μ_n conditioned on $\{W_{\sigma,\mathbf{y}_n^k} = \pi_{\mathtt{B}} W\}$ is $\mathtt{SBM}(\mathbf{y}_n, \pi_{\mathtt{B}} W)$, so we can bound the expectation above using Proposition 2.5.2, getting

$$\mathcal{E}_k(n,\varepsilon) \leq (9n)^{r|\mathtt{B}^{\mathrm{B}(e,k)}|^2(|\mathtt{A}^{\mathrm{B}(e,k)}|+1)} \sum_{W \in \mathcal{W}_n(\alpha\beta,k,\varepsilon)} e^{n(F(W) - F(\pi_{\mathtt{B}} W))} \mathop{\mathbb{P}}_{\sigma\sim\mu_n} (W_{\sigma,\mathbf{y}_n^k} = \pi_{\mathtt{B}} W).$$

Note $(9n)^{r|\mathtt{B}^{\mathrm{B}(e,k)}|^2(|\mathtt{A}^{\mathrm{B}(e,k)}|+1)} \leq e^{o_{n\to\infty}(n)}$. Fix $\delta > 0$. By continuity of F, for all small enough ε (possibly depending on k) we have

$$\mathcal{E}_k(n,\varepsilon) \leq e^{n(F_\mu(T, \alpha^k | \beta^k) + \delta + o_{n\to\infty}(1))} \sum_{W \in \mathcal{W}_n(\alpha\beta,k,\varepsilon)} \mathop{\mathbb{P}}_{\sigma\sim\mu_n} (W_{\sigma,\mathbf{y}_n^k} = \pi_{\mathtt{B}} W).$$

Bounding each probability by 1, we get

$$\mathcal{E}_k(n,\varepsilon) \leq e^{n(F_\mu(T,\alpha^k|\beta^k)+\delta+o_{n\to\infty}(1))}|\mathcal{W}_n(\alpha\beta,k,\varepsilon)|.$$

But

$$|\mathcal{W}_n(\alpha\beta,k,\varepsilon)| \leq n^{r\left|(\mathtt{A}\times\mathtt{B})^{\mathrm{B}(e,k)}\right|^2} \leq e^{o_{n\to\infty}(n)},$$

so this implies

$$\begin{aligned}\limsup_{n\to\infty}\frac{1}{n}\log\mathcal{E}_k(n,\varepsilon) &\leq F_\mu(T,\alpha^k \mid \beta^k)+\delta\\ &\leq F_\mu(T,\alpha^k \mid \beta^{k_2})+\delta\end{aligned}$$

for any $k_2 \geq k$, by monotonicity under splitting. Taking the limit as $k_2 \to \infty$ followed by the infimum over ε (which takes δ to 0) and k gives

$$\inf_{\varepsilon,k}\limsup_{n\to\infty}\frac{1}{n}\log\mathcal{E}_k(n,\varepsilon) \leq f_\mu(T,\alpha \mid \beta).$$

Since

$$\inf_{\mathcal{O}\ni(\alpha\beta)^G_*\mu}\limsup_{n\to\infty}\frac{1}{n}\log\mathop{\mathbb{E}}_{\sigma\sim\mu_n}|\{\mathbf{x}\in\mathtt{A}^n \,:\, (\mathbf{x},\mathbf{y}_n)\in\Omega(\sigma,\mathcal{O})\}| \leq \inf_{\varepsilon}\limsup_{n\to\infty}\frac{1}{n}\log\mathcal{E}_k(n,\varepsilon)$$

for every k, this completes the upper bound.

2.6.2 Lower bound

Fix $k\in\mathbb{N}$. To estimate

$$\mathcal{E} \coloneqq \mathop{\mathbb{E}}_{\sigma\sim\mu_n}|\{\mathbf{x}\in\mathtt{A}^n \,:\, (\mathbf{x},\mathbf{y}_n)\in\Omega^*_k(\sigma,\alpha\beta,\varepsilon)\}|$$

we bound below using the expected size of

$$\mathcal{X}_k(\sigma,\alpha\beta,\varepsilon \mid \mathbf{y}_n) \coloneqq \{\mathbf{X}\in\left(\mathtt{A}^{\mathrm{B}(e,k)}\right)^n \,:\, (\mathbf{X},\mathbf{y}^k_n)\in\Omega^*_0(\sigma,(\alpha\beta)^k,\varepsilon)\}.$$

This is not a true lower bound but, by Equation 2.1 below, there are constants C,d,c independent of n such that

$$|\mathcal{X}_k(\sigma,\alpha\beta,\varepsilon \mid \mathbf{y}_n)| \leq C\exp\left(nd\varepsilon+n\mathrm{H}(2|\mathrm{B}(e,k)|\varepsilon)\right)\cdot|\{\mathbf{x}\in\mathtt{A}^n \,:\, (\mathbf{x},\mathbf{y}_n)\in\Omega^*_k(\sigma,\alpha\beta,\varepsilon)\}|.$$

The 'error' factor has an exponential growth rate which vanishes as $\varepsilon \to 0$, so will not be a problem.

We now find a lower bound for the expectation of $|\mathcal{X}_k|$. Applying Proposition 2.5.2 as above, we have

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\mathcal{X}_k(\sigma, \alpha\beta, \varepsilon \mid \mathbf{y}_n)| \\
&\quad = \sum_{W \in \mathcal{W}_n(\alpha\beta,k,\varepsilon)} \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, W_{\sigma,(\mathbf{X},\mathbf{y}_n^k)} = W\}| \\
&\quad \geq \sum_{W \in \mathcal{W}_n(\alpha\beta,k,\varepsilon)} \exp\!\big[n(F(W) - F(\pi_{\mathtt{B}}W) - o_n(1))\big] \mathop{\mathbb{P}}_{\sigma\sim\mu_n} \big(\pi_{\mathtt{B}}W = W_{\sigma,\mathbf{y}_n^k}\big).
\end{aligned}$$

For any $\delta > 0$, for small enough $\varepsilon > 0$ (independent of n), by continuity of F this is at least

$$\exp\big[n(F_\mu(\alpha^k \mid \beta^k) - \delta - o_n(1))\big] \sum_{W \in \mathcal{W}_n(\alpha\beta,k,\varepsilon)} \mathop{\mathbb{P}}_{\sigma\sim\mu_n} \big(\pi_{\mathtt{B}}W = W_{\sigma,\mathbf{y}_n^k}\big).$$

We give a lower bound for the sum by first rewriting it as

$$\sum_{W_{\mathtt{B}} \text{ denom.-}n\ \mathtt{B}^{\mathrm{B}(e,k)}\text{-weight}} |\{W \in \mathcal{W}_n(\alpha\beta, k, \varepsilon) \,:\, \pi_{\mathtt{B}}W = W_{\mathtt{B}}\}| \cdot \mathop{\mathbb{P}}_{\sigma\sim\mu_n} (W_{\sigma,\mathbf{y}_n^k} = W_{\mathtt{B}}).$$

Fix $\eta > 0$. By Lemma 2.2.3, for all large enough n the $\mathtt{B}$-weight $W_{\sigma_n,\mathbf{y}_n}$ can be extended to a $\mathtt{B}^{\mathrm{B}(e,k)}$-weight $W_{\mathtt{B}}$ with $d(W_{\mathtt{B}}, W_{\beta^k}) \leq \eta$; to apply the lemma we can think of the extended weight $W_{\mathtt{B}}$ as having alphabet $\mathtt{B}^{\mathrm{B}(e,k)\setminus\{e\}} \times \mathtt{B}$, and recall that we assume $\lim_{n\to\infty} d(W_{\sigma_n,\mathbf{y}_n}, W_\beta) = 0$. Choose $\sigma, \mathbf{Y}$ such that $W_{\mathtt{B}} = W_{\sigma,\mathbf{Y}}$. Since $W_{\mathtt{B}}$ is an extension of $W_{\sigma_n,\mathbf{y}_n}$, we can make this choice in such a way that $\pi_e\mathbf{Y} = \mathbf{y}_n$.

Let $\widetilde{W_{\mathtt{B}}} = W_{\sigma,\mathbf{y}_n^k}$. By Proposition 2.2.5,

$$d(\widetilde{W_{\mathtt{B}}}, W_{\beta^k}) \leq d(\widetilde{W_{\mathtt{B}}}, W_{\mathtt{B}}) + d(W_{\mathtt{B}}, W_{\beta^k}) \leq 2r|\mathrm{B}(e,k)|\eta + \eta.$$

So, as long as η is small enough and n is large enough (depending on ε, k), by Lemma 2.2.3

$$|\{W \in \mathcal{W}_n(\alpha\beta, k, \varepsilon) \,:\, \pi_{\mathtt{B}}W = W_{\mathtt{B}}\}| \geq 1.$$

Now consider the probability appearing in the $\widetilde{W}_{\mathtt{B}}$ term:

$$\mathop{\mathbb{P}}_{\sigma\sim\mu_n}(W_{\sigma,\mathbf{y}_n^k} = \widetilde{W}_{\mathtt{B}}) = \frac{|\{\sigma \,:\, W_{\sigma,\mathbf{y}_n^k} = \widetilde{W}_{\mathtt{B}}\}|}{|\{\sigma \,:\, W_{\sigma,\mathbf{y}_n} = W_{\sigma_n,\mathbf{y}_n}\}|}.$$

By symmetry in choice of $\mathbf{y}$ with the correct letter frequencies, we can write this as

$$\begin{aligned}
\mathop{\mathbb{P}}_{\sigma\sim\mu_n}(W_{\sigma,\mathbf{y}_n^k} = \widetilde{W}_{\mathtt{B}}) &= \frac{\left|\{(\sigma,\mathbf{y}) \,:\, W_{\sigma,\mathbf{y}^k} = \widetilde{W}_{\mathtt{B}}\}\right|}{|\{(\sigma,\mathbf{y}) \,:\, W_{\sigma,\mathbf{y}} = W_{\sigma_n,\mathbf{y}_n}\}|} \\
&= \frac{\left|\{(\sigma,\mathbf{Y}) \,:\, W_{\sigma,\mathbf{Y}} = \widetilde{W}_{\mathtt{B}}\}\right|}{|\{(\sigma,\mathbf{y}) \,:\, W_{\sigma,\mathbf{y}} = W_{\sigma_n,\mathbf{y}_n}\}|} && \text{(Prop. 2.2.5)} \\
&= \frac{Z_n(\widetilde{W}_{\mathtt{B}})}{Z_n(W_{\sigma_n,\mathbf{y}_n})} && \text{(definition of } Z_n\text{)} \\
&\geq \exp\Big(n[F(\widetilde{W}_{\mathtt{B}}) - F(W_{\sigma_n,\mathbf{y}_n})]\Big)\cdot(3\sqrt{n})^{-r(|\mathtt{B}^{\mathrm{B}(e,k)}|^2-|\mathtt{B}|)} && \text{(Prop. 2.5.1)} \\
&= \exp\Big(n[F(\widetilde{W}_{\mathtt{B}}) - F(W_{\sigma_n,\mathbf{y}_n}) - o(1)]\Big)\,.
\end{aligned}$$

By continuity of F, we then get

$$\mathop{\mathbb{P}}_{\sigma\sim\mu_n}(W_{\sigma,\mathbf{y}_n^k} = \widetilde{W}_{\mathtt{B}}) \geq \exp n\big(F_\mu(\beta^k) - F_\mu(\beta) - 2\delta + o(1)\big)$$

for all large enough n and small enough η (again depending on k,ε), with $\delta > 0$ the same as chosen above. Since $\beta_*^G\mu$ is a Markov chain, $F_\mu(\beta^k) = F_\mu(\beta)$.

Putting this all together: for any $k \in \mathbb{N}$, for all $\delta > 0$ we have

$$\mathop{\mathbb{E}}_{\sigma\sim\mu_n}|\mathcal{X}_k(\sigma,\alpha\beta,\varepsilon \mid \mathbf{y}_n)| \geq \exp\left[n(F_\mu(\alpha^k \mid \beta^k) - 3\delta - o(1))\right]$$

for all large enough n and small enough $\varepsilon > 0$.

It follows that for any $k \in \mathbb{N}$

$$\inf_\varepsilon \limsup_{n\to\infty} \frac{1}{n}\log \mathop{\mathbb{E}}_{\sigma\sim\mu_n}|\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x},\mathbf{y}_n) \in \Omega_k^*(\sigma,\alpha\beta,\varepsilon)\}| \geq F_\mu(T,\alpha^k \mid \beta^k).$$

Taking the limit as $k \to \infty$ gives the desired bound, using Corollary 2.3.3 and that the family of pseudometrics $\{d_k^* \,:\, k \in \mathbb{N}\}$ generates the weak* topology.

2.7 Proof of Theorem C

Let $W_n = W_{\sigma_n, \mathbf{y}_n^{m_n}}$, so that

$$\mu_n = \mathtt{SBM}(\mathbf{y}_n, W_n).$$

Note that, by definition of μ_n,

$$\mathbb{P}_{\sigma \sim \mu_n} \left(W_{\sigma, \mathbf{y}_n^{m_n}} = W_n \right) = 1.$$

Lemma 2.7.1. *With W_n as just defined in terms of m_n, σ_n, and $\mathbf{y}_n$, we have*

$$\lim_{n \to \infty} F(W_n) = f_\mu(T, \beta).$$

Proof. The assumption in the theorem statement that $d^*_{m_n}(P^{\sigma_n}_{\mathbf{y}_n}, \beta^G_* \mu) = O\big(\frac{1}{\log n}\big)$ implies the existence of a constant C such that

$$d(W_n, W_{\beta^{m_n}}) \leq \frac{C}{\log n}.$$

By Lemma 2.3.1 we have

$$|F(W_{\sigma, \mathbf{y}^{m_n}}) - F(W_{\beta^{m_n}})| \leq 4r \Big(\mathrm{H}(\tfrac{C}{\log n}) + \tfrac{C}{\log n} |\mathrm{B}(e, m_n)| \log_2 |\mathtt{B}| \Big) = o(1)$$

using that $m_n = o(\log \log n)$. Since m_n approaches infinity as n goes to infinity we have $f_\mu(T, \beta) = \lim_{n \to \infty} F(W_{\beta^{m_n}})$, so the result follows. □

Lemma 2.7.2. *If $m_n = o(\log \log n)$, then for any $k > 0$ and $\varepsilon > 0$ we have $|\mathtt{B}^{\mathrm{B}(e, m_n)}|^k = o(n^\varepsilon)$.*

Proof. This is certainly true if $|\mathtt{B}| = 1$; assume therefore that $|\mathtt{B}| \geq 2$.

Our assumption $m_n = o(\log \log n)$ guarantees that

$$(2r-1)^{m_n} < \frac{r-1}{r} \frac{\varepsilon}{k \log |\mathtt{B}|} \log n$$

for all large enough n. Therefore

$$|\mathrm{B}(e, m_n)| = \frac{r(2r-1)^{m_n} - 1}{r - 1} < \frac{\varepsilon}{k \log |\mathtt{B}|} \log n.$$

This inequality can be rearranged to give

$$|\mathtt{B}^{\mathrm{B}(e,m_n)}|^k < n^\varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, the result follows. □

In the remainder of this section we prove Theorem C by first proving the right-hand side is an upper bound for the left, then proving it is also lower bound.

2.7.1 Upper bound

Just as in the proof of the upper bound in Theorem F, for each $k \in \mathbb{N}$ and $\varepsilon > 0$ we have

$$\inf_{\mathcal{O} \ni (\alpha\beta)^G_* \mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x}, \mathbf{y}_n) \in \Omega(\sigma, \mathcal{O})\}| \leq \limsup_{n\to\infty} \frac{1}{n} \log \mathcal{E}_k(n, \varepsilon),$$

where

$$\begin{aligned}\mathcal{E}_k(n, \varepsilon) &\coloneqq \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, (\mathbf{X}, \mathbf{y}_n^k) \in \Omega_0^*(\sigma, (\alpha\beta)^k, \varepsilon)\}| \\ &= \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, d\big(W_{\sigma,(\mathbf{X},\mathbf{y}_n^k)}, W_{(\alpha\beta)^k}\big) < \varepsilon)\}|.\end{aligned}$$

We assume that n is large enough that $m_n \geq k$.

Since μ_n is $\mathtt{SBM}(\sigma_n, \mathbf{y}_n, m_n)$ rather than $\mathtt{SBM}(\sigma_n, \mathbf{y}_n, k)$, we cannot apply Proposition 2.5.2 directly to this expression. We get around this as follows: Let

$$\mathcal{W}_n(m, m') \coloneqq \left\{ W_{\sigma,(\mathbf{X},\mathbf{y}^{m'})} \,:\, \sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n)),\ \mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,m)})^n,\ \mathbf{y} \in \mathtt{B}^n \right\}.$$

All elements of this set are denominator-n $\mathtt{A}^{\mathrm{B}(e,m)} \times \mathtt{B}^{\mathrm{B}(e,m')}$-weights; we avoid the question of exactly which weights are in this set, but call such weights *attainable*. For $k \leq m$ and $k' \leq m'$ let

$$\mathcal{W}_n(m, m'; \alpha\beta, k, k'; \varepsilon) = \left\{ W \in \mathcal{W}_n(m, m') \,:\, d\big(\pi_{k,k'} W,\ W_{\alpha^k \beta^{k'}}\big) < \varepsilon \right\}$$

denote the set of such weights whose appropriate marginal is within ε of the $(\mathtt{A}^{\mathrm{B}(e,k)} \times \mathtt{B}^{\mathrm{B}(e,k')})$-weight $W_{\alpha^k\beta^{k'}}$. For now we take $m = k = k'$ but we will need more generality below. Then

$$\mathcal{E}_k(n, \varepsilon) = \mathop{\mathbb{E}}_{\sigma\sim\mu_n} \sum_{W \in \mathcal{W}_n(k, m_n; \alpha\beta, k, k; \varepsilon)} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n \,:\, W_{\sigma,(\mathbf{X},\mathbf{y}_n^{m_n})} = W\}|$$

so we can apply Proposition 2.5.2 to get

$$\mathcal{E}_k(n,\varepsilon) \le (9n)^{r|\mathtt{B}^{\mathrm{B}(e,m_n)}|^2(|\mathtt{A}^{\mathrm{B}(e,k)}|+1)} \sum_{W \in \mathcal{W}_n(k,m_n;\alpha\beta,k,k;\varepsilon)} e^{n(F(W)-F(\pi_{\mathtt{B}}W))} \mathbf{1}_{\{\pi_{\mathtt{B}}W = W_n\}}.$$

By Lemma 2.7.2 we have $(9n)^{r|\mathtt{B}^{\mathrm{B}(e,m_n)}|^2(|\mathtt{A}^{\mathrm{B}(e,k)}|+1)} \le e^{o_{n\to\infty}(n)}$. Using this and Lemma 2.7.1 we have

$$\mathcal{E}_k(n,\varepsilon) \le \sum_{W \in \mathcal{W}_n(k,m_n;\alpha\beta,k,k;\varepsilon)} e^{n(F(W)-f(T,\beta)+o_{n\to\infty}(1))} \mathbf{1}_{\{\pi_{\mathtt{B}}W = W_n\}},$$

where the little o is uniform over all terms in the sum. Here we use the assumption that $f_\mu(T,\beta)$ is finite.

By definition of $\mathcal{W}_n(k,m_n)$, for any $W \in \mathcal{W}_n(k,m_n;\alpha\beta,k,k;\varepsilon)$ we can pick some $\sigma \in \mathrm{Hom}(\Gamma,\mathrm{Sym}(n))$, $\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k)})^n$, and $\mathbf{y} \in \mathtt{B}^n$ so that $W = W_{\sigma,(\mathbf{X},\mathbf{y}^{m_n})}$. Then since $\mathbf{X}\mathbf{y}^{m_n}$ is a splitting of $\mathbf{X}\mathbf{y}^k$, by Lemma 2.3.2 we have

$$F(W) = F(\sigma, \mathbf{X}\mathbf{y}^{m_n}) \le F(\sigma, \mathbf{X}\mathbf{y}^k) = F(\pi_{k,k}W).$$

By continuity of F, for all small enough ε (depending on k) we have

$$F(\pi_{k,k}W) \le F(W_{(\alpha\beta)^k}) + \delta = F_\mu(T,(\alpha\beta)^k) + \delta.$$

Along with the above, this implies that

$$\mathcal{E}_k(n,\varepsilon) \le e^{n(F(T,(\alpha\beta)^k)-f(T,\beta)+o_n(1)+\delta)} \sum_{W \in \mathcal{W}_n(k,m_n;\alpha\beta,k,k;\varepsilon)} \mathbf{1}_{\{\pi_{\mathtt{B}}W = W_n\}}.$$

Bounding all terms in the sum by 1, we get

$$\mathcal{E}_k(n,\varepsilon) \le e^{n(F(T,(\alpha\beta)^k)-f_\mu(T,\beta)+o_n(1)+\delta)} \, |\mathcal{W}_n(k,m_n;\alpha\beta,k,k;\varepsilon)|.$$

Using Lemma 2.7.2 we have

$$|\mathcal{W}_n(k,m_n;\alpha\beta,k,k;\varepsilon)| \le |\mathcal{W}_n(k,m_n)| \le n^{r\left|\mathtt{A}^{\mathrm{B}(e,k)}\times\mathtt{B}^{\mathrm{B}(e,m_n)}\right|^2} \le e^{o_{n\to\infty}(n)},$$

so this implies

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathcal{E}_k(n,\varepsilon) \le F_\mu(T,(\alpha\beta)^k) - f_\mu(T,\beta) + \delta.$$

Taking the infimum over ε and k, and using the chain rule for f (Corollary 2.3.5, again using the assumption that $f_\mu(T,\beta)$ is finite), gives

$$\inf_{\varepsilon,k} \limsup_{n\to\infty} \frac{1}{n} \log \mathcal{E}_k(n,\varepsilon) \leq f_\mu(T,\alpha\beta) - f_\mu(T,\beta) = f_\mu(T,\alpha \mid \beta).$$

Since

$$\inf_{\mathcal{O}\ni(\alpha\beta)^G_*\mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x},\mathbf{y}_n) \in \Omega(\sigma,\mathcal{O})\}| \leq \inf_{\varepsilon} \limsup_{n\to\infty} \frac{1}{n} \log \mathcal{E}_k(n,\varepsilon),$$

for every k, this completes the upper bound.

2.7.2 Lower bound

In this section we denote

$$\mathcal{X}_{k_1,k_2}(\sigma,\alpha\beta,\varepsilon \mid \mathbf{y}) \coloneqq \{\mathbf{X} \in \left(\mathtt{A}^{\mathrm{B}(e,k_1)}\right)^n \,:\, (\mathbf{X},\mathbf{y}^{k_2}) \in \Omega_0^*(\sigma,\alpha^{k_1}\beta^{k_2},\varepsilon)\}$$

$$\Omega_k^*(\sigma,\alpha\beta,\varepsilon \mid \mathbf{y}) \coloneqq \{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x},\mathbf{y}) \in \Omega_k^*(\sigma,\alpha\beta,\varepsilon)\}$$

(note the dependence on n is implicitly specified by $\sigma \in \mathrm{Hom}(\Gamma,\mathrm{Sym}(n))$ and $\mathbf{y} \in \mathtt{B}^n$), and with $\Sigma = \{\mu_n\}_{n=1}^\infty$

$$\begin{aligned} \mathrm{h}_{\Sigma,\mu}(T,\alpha \mid \beta \,:\, k,\varepsilon) &\coloneqq \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, (\mathbf{x},\mathbf{y}) \in \Omega_k^*(\sigma,\alpha\beta,\varepsilon)\}| \\ &= \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\Omega_k^*(\sigma,\alpha\beta,\varepsilon \mid \mathbf{y})|. \end{aligned}$$

The following two claims are used to relate the sizes of the sets defined above.

Claim 2.7.3. *Let $k \leq \min(k_1,k_2)$. For any $\sigma, \mathbf{y}$ we have*

$$\pi_e\left[\mathcal{X}_{k_1,k_2}(\sigma,\alpha\beta,\varepsilon \mid \mathbf{y})\right] \subseteq \Omega_k^*(\sigma,\alpha\beta,c\varepsilon \mid \mathbf{y})$$

where $c = 1 + |\mathrm{B}(e,k)|$.

Proof. If $(\mathbf{X},\mathbf{y}^{k_2}) \in \Omega_0^*(\sigma,\alpha^{k_1}\beta^{k_2},\varepsilon)$, then

$$\pi_{k,k}(\mathbf{X},\mathbf{y}^{k_2}) \in \Omega_0^*(\sigma,(\alpha\beta)^k,\varepsilon);$$

this follows from the fact that total variation distance is nonincreasing under pushforwards. Applying Proposition 2.2.5, we get

$$(\pi_e \mathbf{X}, \mathbf{y}) = \pi_e\left(\pi_{k,k}(\mathbf{X}, \mathbf{y}^{k_2})\right) \in \Omega_k^*(\sigma, \alpha\beta, c\varepsilon). \qquad \square$$

Claim 2.7.4. *Fix $\sigma, \mathbf{y}$, and $k \le \min(k_1, k_2)$. As established in the previous claim, we can consider π_e as a map from $\mathcal{X}_{k_1,k_2}(\sigma, \alpha\beta, \varepsilon \mid \mathbf{y})$ to $\Omega_k^*(\sigma, \alpha\beta, c\varepsilon \mid \mathbf{y})$. There are constants C, d independent of n such that π_e is at most $C \exp\left(nd\varepsilon + n\mathrm{H}(2|\mathrm{B}(e,k)|\varepsilon)\right)$-to-one.*

Proof. If $\Omega_k^*(\sigma, \alpha\beta, c\varepsilon \mid \mathbf{y})$ is empty, then the claim is vacuously true. Otherwise, fix $\mathbf{x} \in \Omega_k^*(\sigma, \alpha\beta, c\varepsilon \mid \mathbf{y})$. If $\mathbf{X} \in \pi_e^{-1}\{\mathbf{x}\}$, then $\pi_e(\mathbf{X}, \mathbf{y}^k) = (\mathbf{x}, \mathbf{y})$. By Claim 3 in the proof of Proposition 3.2 of [Bow10a] the number of such pairs $(\mathbf{X}, \mathbf{y}^k)$, and therefore the number of such $\mathbf{X}$, is bounded above by

$$3\sqrt{2}|\mathtt{A} \times \mathtt{B}|^{|\mathrm{B}(e,k)|\left(n|\mathrm{B}(e,k)|\varepsilon - 1\right)} \exp\left(n\mathrm{H}(2|\mathrm{B}(e,k)|\varepsilon)\right)$$

where H is the Shannon entropy. (We give more explicit constants here than in [Bow10a] to make the dependence on n clear). $\square$

Claim 2 implies that

$$|\mathcal{X}_{k_1,k_2}(\sigma, \alpha\beta, \varepsilon \mid \mathbf{y})| \le C \exp\left(nd\varepsilon + n\mathrm{H}(2|\mathrm{B}(e,k)|\varepsilon)\right) \cdot |\Omega_k^*(\sigma, \alpha\beta, c\varepsilon \mid \mathbf{y})|, \tag{2.1}$$

where C, d are independent of n.

We now find a lower bound for the expectation of $|\mathcal{X}|$. Fix $k_1, k_2 \in \mathbb{N}$, and suppose n is large enough that $m_n \geq \max(k_1, k_2)$. Using Proposition 2.5.2 and Lemma 2.7.2, we have

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\mathcal{X}_{k_1,k_2}(\sigma, \alpha\beta, \varepsilon \mid \mathbf{y}_n)| \\
&\quad = \sum_{W \in \mathcal{W}_n(k_1, m_n; \alpha\beta, k_1, k_2; \varepsilon)} \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{X} \in (\mathtt{A}^{\mathrm{B}(e,k_1)})^n \,:\, W_{\sigma,(\mathbf{X},\mathbf{y}_n^{m_n})} = W\}| \\
&\quad \geq \sum_{W \in \mathcal{W}_n(k_1, m_n; \alpha\beta, k_1, k_2; \varepsilon)} \exp\big[n(F(W) - F(\pi_{\mathtt{B}} W) - o_n(1))\big] \mathbf{1}_{\{\pi_{\mathtt{B}} W = W_{\sigma,\mathbf{y}_n^{m_n}}\}} \\
&\quad \geq \inf_{W \in \mathcal{W}_n(k_1, m_n; \alpha\beta, k_1, k_2; \varepsilon)} \exp\big[n(F(W) - F(\pi_{\mathtt{B}} W) - o_n(1))\big] \\
&\qquad \times \sum_{W \in \mathcal{W}_n(k_1, m_n; \alpha\beta, k_1, k_2; \varepsilon)} \mathbf{1}_{\{\pi_{\mathtt{B}} W = W_{\sigma,\mathbf{y}_n^{m_n}}\}}
\end{aligned}$$

We bound the infimum below as follows: Given any $W \in \mathcal{W}_n(k_1, m_n; \alpha\beta, k_1, k_2; \varepsilon)$, we can let $\mathbf{X}, \mathbf{y}, \sigma$ be such that $W = W_{\sigma,(\mathbf{X},\mathbf{y}^{m_n})}$. Then by Lemma 2.3.2 and continuity of F

$$\begin{aligned}
F(W) - F(\pi_{\mathtt{B}} W) &= F(\sigma, \mathbf{X} | \mathbf{y}^{m_n}) \\
&\geq F(\sigma, \mathbf{X} | \mathbf{y}^{k_2}) \\
&= F(\pi_{k_1,k_2} W) - F(\pi_{\mathtt{B}} \pi_{k_1,k_2} W) \\
&\geq F_\mu(T, \alpha^{k_1} | \beta^{k_2}) - \delta
\end{aligned}$$

for any $\delta > 0$ for all small enough ε (with "small enough" dependent only on k_1, k_2). This implies that the infimum is bounded below by

$$\exp\big[n(F_\mu(T, \alpha^{k_1} | \beta^{k_2}) - o_n(1) - \delta)\big].$$

We bound the sum below by first rewriting it as

$$\big|\{W \in \mathcal{W}_n(k_1, m_n; \alpha\beta, k_1, k_2; \varepsilon) \,:\, \pi_{\mathtt{B}} W = W_{\sigma,\mathbf{y}_n^{m_n}}\}\big|.$$

The following claim, then, implies that the sum is bounded below by 1.

Claim 2.7.5. *For all large enough n,*

$$\big\{W \in \mathcal{W}_n(k_1, m_n; \alpha\beta, k_1, k_2; \varepsilon) \,:\, \pi_{\mathtt{B}} W = W_{\sigma,\mathbf{y}_n^{m_n}}\big\} \neq \varnothing.$$

Proof. By Lemma 2.2.3, if

$$n > 680|\mathtt{A}^{\mathrm{B}(e,k_1)} \times \mathtt{B}^{\mathrm{B}(e,m_n)}|^2 r/\varepsilon$$

and $d(W_{\sigma,\mathbf{y}_n^{m_n}}, W_{\beta^{m_n}}) < \frac{\varepsilon}{530r}$ then there is a $(\mathtt{A}^{\mathrm{B}(e,k_1)} \times \mathtt{B}^{\mathrm{B}(e,m_n)})$-weight W with $\pi_{\mathtt{B}} W = W_{\sigma,\mathbf{y}_n^{m_n}}$ and $d(W, W_{\alpha^{k_1}\beta^{m_n}}) < \varepsilon$. By definition of μ_n and Lemma 2.7.2, both conditions are met for all large enough n.

The claim will follow if we show that W is attainable.

With W as chosen above, by Proposition 2.2.4 we can choose $\tilde{\sigma} \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$, $\tilde{\mathbf{X}} \in (\mathtt{A}^{\mathrm{B}(e,k_1)})^n$, and $\tilde{\mathbf{Y}} \in (\mathtt{B}^{\mathrm{B}(e,m_n)})^n$ such that $W = W_{\tilde{\sigma},(\tilde{\mathbf{X}},\tilde{\mathbf{Y}})}$.

Let $\tilde{\mathbf{y}} = \pi_e \tilde{\mathbf{Y}} \in \mathtt{B}^n$. To complete the proof we show that $\tilde{\mathbf{y}}^{m_n} = \tilde{\mathbf{Y}}$, i.e.

$$\tilde{\mathbf{y}}\big(\tilde{\sigma}(g)i\big) = \Big(\tilde{\mathbf{Y}}(i)\Big)_g$$

for all $i \in [n]$ and $g \in \mathrm{B}(e, m_n)$. We prove this by induction on the word length $|g|$.

The base case $|g| = 0$ (i.e. $g = e$) follows immediately from the definition of $\tilde{\mathbf{y}}$.

For the inductive step, write $g = ht$ with $|h| = |g| - 1$ and $t \in \{s_1^{\pm 1}, \ldots, s_r^{\pm 1}\}$. Then, assuming the result holds for h,

$$\tilde{\mathbf{y}}\big(\tilde{\sigma}(g)i\big) = \tilde{\mathbf{y}}\big(\tilde{\sigma}(h)\tilde{\sigma}(t)i\big) = \Big(\tilde{\mathbf{Y}}(\tilde{\sigma}(t)i)\Big)_h .$$

Now since $W_{\tilde{\sigma},\tilde{\mathbf{Y}}} = W_{\sigma_n,\mathbf{y}_n^{m_n}}$, we can pick $j \in [n]$ such that

$$\tilde{\mathbf{Y}}(i) = \mathbf{y}_n^{m_n}(j) \quad \text{and} \quad \tilde{\mathbf{Y}}(\tilde{\sigma}(t)i) = \mathbf{y}_n^{m_n}(\sigma(t)j).$$

This implies

$$\Big(\tilde{\mathbf{Y}}(\tilde{\sigma}(t)i)\Big)_h = \big(\mathbf{y}_n^{m_n}(\sigma(t)j)\big)_h = \mathbf{y}_n(\sigma(g)j) = \big(\mathbf{y}_n^{m_n}(j)\big)_g = \Big(\tilde{\mathbf{Y}}(i)\Big)_g . \qquad \square$$

Hence for all large enough n we have

$$\mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\mathcal{X}_{k_1,k_2}(\sigma, \alpha\beta, \varepsilon \mid \mathbf{y}_n)| \geq \exp\big[n(F_\mu(T, \alpha^{k_1} \mid \beta^{k_2}) - o_n(1) - \delta)\big],$$

and therefore

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\mathcal{X}_{k_1,k_2}(\sigma, \alpha\beta, \varepsilon \mid \mathbf{y}_n)| \geq F_\mu(T, \alpha^{k_1} \mid \beta^{k_2}) - \delta.$$

Combining this lower bound with Equation (2.1) and the definition of $\mathrm{h}_{\Sigma,\mu}(T, \alpha \mid \beta : k, c\varepsilon)$, we get

$$d\varepsilon + \mathrm{H}(2|\mathrm{B}(e,k)|\varepsilon) + \mathrm{h}_{\Sigma,\mu}(T, \alpha \mid \beta : k, c\varepsilon) \geq F_\mu(T, \alpha^{k_1} \mid \beta^{k_2}) - \delta.$$

Taking the inf in ε then letting δ go to zero gives

$$\inf_\varepsilon \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n : (\mathbf{x}, \mathbf{y}_n) \in \Omega_k^*(\sigma, \alpha\beta, \varepsilon)\}| \geq F_\mu(T, \alpha^{k_1} \mid \beta^{k_2})$$

for $k \leq \min(k_1, k_2)$. First take $k_2 \to \infty$, then $k_1 \to \infty$, then take the infimum over k. We get

$$\begin{aligned} f_\mu(T, \alpha \mid \beta) &\leq \inf_{\varepsilon,k} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n : (\mathbf{x}, \mathbf{y}_n) \in \Omega_k^*(\sigma, \alpha\beta, \varepsilon)\}| \\ &= \inf_{\mathcal{O}\ni(\alpha\beta)_*^G\mu} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n : (\mathbf{x}, \mathbf{y}_n) \in \Omega(\sigma, \mathcal{O})\}| \end{aligned}$$

where the last line follows because the collection of pseudometrics $\{d_k^* : k \in \mathbb{N}\}$ generates the weak* topology on $\mathrm{Prob}((\mathtt{A} \times \mathtt{B})^\Gamma)$.

2.8 Proof of Theorem D

By analogy with sofic entropy, we denote $\Sigma := \{\mu_n\}_{n=1}^\infty$ and denote the left-hand side of the formula in the theorem statement as $\mathrm{h}_{\Sigma,\mu}(T, \alpha)$.

Endow $\mathrm{Prob}(\mathtt{A}^\Gamma)$ with the metric

$$d(\lambda, \nu) := \sum_{r=1}^\infty 2^{-r} d^{\mathrm{B}(e,r)}(\lambda, \nu).$$

Note that this induces the weak* topology (where $\mathtt{A}$ is given the discrete topology and $\mathtt{A}^\Gamma$ the product topology).

Writing $\mu_{\mathtt{A}} = \alpha^G_* \mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$, we then have

$$\mathrm{h}_{\Sigma,\mu}(T,\alpha) = \inf_{\varepsilon>0} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, d(P^\sigma_{\mathbf{x}}, \mu_{\mathtt{A}}) < \varepsilon\}|.$$

We will similarly denote $\mu_{\mathtt{B}} = \beta^G_* \mu \in \mathrm{Prob}(\mathtt{B}^\Gamma)$.

2.8.1 Lower bound

Let $\lambda \in \mathrm{Prob}((\mathtt{A}\times\mathtt{B})^\Gamma)$ be any joining of (the shift systems with respective measures) $\mu_{\mathtt{A}}$ and $\mu_{\mathtt{B}}$. Then for any $\mathbf{x} \in \mathtt{A}^n$ and $\mathbf{y} \in \mathtt{B}^n$ we have

$$d(P^\sigma_{\mathbf{x}}, \mu_{\mathtt{A}}) \le d(P^\sigma_{(\mathbf{x},\mathbf{y})}, \lambda),$$

where d is defined on $\mathrm{Prob}((\mathtt{A}\times\mathtt{B})^\Gamma)$ analogously to the definition given on $\mathrm{Prob}(\mathtt{A}^\Gamma)$ above. This inequality holds because total variation distance is nonincreasing under pushforwards. Consequently

$$\mathrm{h}_{\Sigma,\mu}(T,\alpha) \ge \inf_{\varepsilon>0} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, d(P^\sigma_{(\mathbf{x},\mathbf{y}_n)}, \lambda) < \varepsilon\}| = f_\lambda(S, \mathtt{a} \mid \mathtt{b}).$$

Taking the supremum over joinings λ gives the lower bound.

2.8.2 Upper bound

For $\varepsilon > 0$, let

$$\mathsf{J}_\varepsilon \coloneqq \{\lambda \in \mathrm{Prob}^S((\mathtt{A}\times\mathtt{B})^\Gamma) \,:\, d(\mathtt{a}^G_*\lambda, \mu_{\mathtt{A}}) < \varepsilon \text{ and } d(\mathtt{b}^G_*\lambda, \mu_{\mathtt{B}}) < \varepsilon\}$$

be the set of shift-invariant "approximate joinings" of $\mu_{\mathtt{A}}$ and $\mu_{\mathtt{B}}$. Since $\mathrm{Prob}((\mathtt{A}\times\mathtt{B})^\Gamma)$ is compact, for each $\varepsilon > 0$ there exist $\lambda_1, \ldots, \lambda_m \in \mathsf{J}_\varepsilon$ such that

$$\mathsf{J}_\varepsilon \subseteq \bigcup_{i=1}^m \mathrm{B}(\lambda_i, \varepsilon).$$

By definition of μ_n we have $\mathbb{P}_{\sigma\sim\mu_n}(d(P^\sigma_{\mathbf{y}_n}, \mu_{\mathtt{B}}) < \varepsilon) = 1$ for all large enough n. Therefore

$$\begin{aligned}
\mathrm{h}_{\Sigma,\mu}(T,\alpha) &= \inf_\varepsilon \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, P^\sigma_{(\mathbf{x},\mathbf{y}_n)} \in \mathsf{J}_\varepsilon\}| \\
&\leq \inf_\varepsilon \limsup_{n\to\infty} \frac{1}{n} \log \sum_{i=1}^m \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, P^\sigma_{(\mathbf{x},\mathbf{y}_n)} \in \mathrm{B}(\lambda_i, \varepsilon)\}| \\
&= \inf_\varepsilon \max_{1\leq i\leq m} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, P^\sigma_{(\mathbf{x},\mathbf{y}_n)} \in \mathrm{B}(\lambda_i, \varepsilon)\}| \\
&\leq \inf_\varepsilon \sup_{\lambda\in\mathsf{J}_\varepsilon} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, P^\sigma_{(\mathbf{x},\mathbf{y}_n)} \in \mathrm{B}(\lambda, \varepsilon)\}|.
\end{aligned}$$

Note that the entire expression in the inf is decreasing as $\varepsilon \to 0$, so we may replace the inf with a limit. Rather than taking a continuous limit we write

$$\mathrm{h}_{\Sigma,\mu}(T,\alpha) \leq \lim_{m\to\infty} \sup_{\lambda\in\mathsf{J}_{1/m}} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, P^\sigma_{(\mathbf{x},\mathbf{y}_n)} \in \mathrm{B}(\lambda, 1/m)\}|.$$

For each m pick $\lambda_m \in \mathsf{J}_{1/m}$ to get within $1/m$ of the supremum. Then the right-hand side is equal to

$$\lim_{m\to\infty} \limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, P^\sigma_{(\mathbf{x},\mathbf{y}_n)} \in \mathrm{B}(\lambda_m, 1/m)\}|. \tag{$*$}$$

Let λ_{m_j} be a subsequence with weak* limit λ_0. By weak* continuity of pushforwards under projection we have $\lambda_0 \in \mathsf{J}(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$. Now for any $\delta > 0$, for all large enough j we have both $1/m_j < \delta/2$ and $d(\lambda_{m_j}, \lambda_0) < \delta/2$, so by the triangle inequality

$$\mathrm{B}(\lambda_{m_j}, 1/m_j) \subseteq \mathrm{B}(\lambda_0, \delta).$$

It follows that the expression in $(*)$, and hence $h_\Sigma(\alpha)$, is bounded above by

$$\limsup_{n\to\infty} \frac{1}{n} \log \mathop{\mathbb{E}}_{\sigma\sim\mu_n} |\{\mathbf{x} \in \mathtt{A}^n \,:\, P^\sigma_{(\mathbf{x},\mathbf{y}_n)} \in \mathrm{B}(\lambda_0, \delta)\}|.$$

Taking the infimum over δ shows that

$$h_\Sigma(\mu,\alpha) \leq f_{\lambda_0}(S, \mathtt{a} \mid \mathtt{b}) \leq \sup_{\lambda\in\mathsf{J}(\mu_{\mathtt{A}},\mu_{\mathtt{B}})} f_\lambda(S, \mathtt{a} \mid \mathtt{b}).$$

2.9 Proof of Proposition 2.1.3

All sequences of interest are of the form

$$\mu_n = \texttt{SBM}(\sigma_n, \mathbf{y}_n, m_n) = \mathrm{Unif}(\{\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n)) \,:\, W_{\sigma, \mathbf{y}_n^{m_n}} = W_n\})$$

with $\mathbf{y}_n \in \mathtt{B}^n$, $\sigma_n \in \mathrm{Sym}(n)$, $m_n = o(\log\log n)$, and where W_n is the $\mathtt{B}^{\mathrm{B}(e,m_n)}$-weight $W_{\sigma_n, \mathbf{y}_n^{m_n}}$. In the case of Theorem F we simply have $m_n = 0$ for all n.

The theorem will follow from the following:

Lemma 2.9.1. *Let ζ_n denote the uniform measure on* $\mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$*. Then for any finite $D \subset \Gamma$ and $\delta > 0$ there exists $\varepsilon > 0$ such that*

$$\mathop{\mathbb{P}}_{\sigma \sim \zeta_n} (\sigma \textit{ is } (D, \delta)\textit{-sofic}) \geq 1 - n^{-\varepsilon n}$$

for all large enough n. □

This can be proven by making superficial changes to the proof of the similar result in [Bow20b].

To prove Proposition 2.1.3, it now suffices to show that for any $\varepsilon > 0$

$$\mathop{\mathbb{P}}_{\sigma \sim \zeta_n} (W_{\sigma, \mathbf{y}_n^{m_n}} = W_n) \geq n^{-\varepsilon n}$$

for all large enough n. To do this, first note that the left-hand side here depends only on the vector $p_{\mathbf{y}_n} \in \mathrm{Prob}(\mathtt{B})$ of letter frequencies. Therefore

$$\begin{aligned}\mathop{\mathbb{P}}_{\sigma \sim \zeta_n} (\exists \mathbf{y} \in \mathtt{B}^n \text{ s.t. } W_{\sigma, \mathbf{y}^{m_n}} = W_n) &\leq \sum_{\mathbf{y} \,:\, p_{\mathbf{y}} = p_{\mathbf{y}_n}} \mathop{\mathbb{P}}_{\sigma \sim \zeta_n} (W_{\sigma, \mathbf{y}^{m_n}} = W_n) \\ &= \exp\{n \mathrm{H}(p_{\mathbf{y}_n}) + o(n)\} \mathop{\mathbb{P}}_{\sigma \sim \zeta_n} (W_{\sigma, \mathbf{y}_n^{m_n}} = W_n).\end{aligned}$$

But by Proposition 2.2.5, if $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$ and $\mathbf{Y} \in (\mathtt{B}^{\mathrm{B}(e,m_n)})^n$ are such that $W_{\sigma, \mathbf{Y}} = W_n = W_{\sigma_n, \mathbf{y}_n^{m_n}}$, then the projection $\mathbf{Y}_e \in \mathtt{B}^n$ satisfies $(\mathbf{Y}_e)^{m_n} = \mathbf{Y}$. Therefore for each σ

$$\big|\{\mathbf{Y} \in (\mathtt{B}^{\mathrm{B}(e,m_n)})^n \,:\, W_{\sigma, \mathbf{Y}} = W_n\}\big| = |\{\mathbf{y} \in \mathtt{B}^n \,:\, W_{\sigma, \mathbf{y}^{m_n}} = W_n\}|.$$

Hence

$$\mathop{\mathbb{E}}_{\sigma\sim\zeta_n}\big|\{\mathbf{Y}\in(\mathtt{B}^{\mathrm{B}(e,m_n)})^n \,:\, W_{\sigma,\mathbf{Y}}=W_n\}\big| = \mathop{\mathbb{E}}_{\sigma\sim\zeta_n}|\{\mathbf{y}\in\mathtt{B}^n \,:\, W_{\sigma,\mathbf{y}^{m_n}}=W_n\}|$$
$$\leq |\mathtt{B}|^n \mathop{\mathbb{P}}_{\sigma\sim\zeta_n}(\exists\mathbf{y}\in\mathtt{B}^n \text{ s.t. } W_{\sigma,\mathbf{y}^{m_n}}=W_n).$$

Combining these last few statements, we see that

$$\mathop{\mathbb{P}}_{\sigma\sim\zeta_n}(W_{\sigma,\mathbf{y}_n^{m_n}}=W_n)\geq \exp\{-2n\log|\mathtt{B}|+o(n)\}\mathop{\mathbb{E}}_{\sigma\sim\zeta_n}\big|\{\mathbf{Y}\in(\mathtt{B}^{\mathrm{B}(e,m_n)})^n \,:\, W_{\sigma,\mathbf{Y}}=W_n\}\big|.$$

We can ignore the first factor here since it only decays exponentially fast. By Proposition 2.5.1,

$$\mathop{\mathbb{E}}_{\sigma\sim\zeta_n}\big|\{\mathbf{Y}\in(\mathtt{B}^{\mathrm{B}(e,m_n)})^n \,:\, W_{\sigma,\mathbf{Y}}=W_n\}\big| = \frac{Z_n(W_n)}{(n!)^r}\geq (3\sqrt{n})^{-r|\mathtt{B}^{\mathrm{B}(e,m_n)}|^2}e^{F(W_n)n}n^{(1-r)/2}.$$

The third factor is clearly not a problem and can also be ignored. For the first factor,

$$\frac{1}{n\log n}\log(3\sqrt{n})^{-r|\mathtt{B}^{\mathrm{B}(e,m_n)}|^2} = -r\frac{|\mathtt{B}^{\mathrm{B}(e,m_n)}|^2}{n}\frac{\log 3\sqrt{n}}{\log n}\to 0 \text{ as } n\to\infty$$

using Lemma 2.7.2. For the second factor, first note that by definition of $F(W_n)$ we have

$$\begin{aligned} F(W_n) &= (1-2r)\mathrm{H}\big(W_n(\cdot)\big)+\sum_{i\in[r]}\mathrm{H}\big(W_n(\cdot,\cdot;i)\big) \\ &\geq -2r\mathrm{H}\big(W_n(\cdot)\big) \\ &\geq -2r\log\big|\mathtt{B}^{\mathrm{B}(e,m_n)}\big|. \end{aligned}$$

So

$$\frac{1}{n\log n}\log e^{F(W_n)n} = \frac{F(W_n)}{\log n}\geq -2r\frac{\log\big|\mathtt{B}^{\mathrm{B}(e,m_n)}\big|}{\log n}\to 0 \text{ as } n\to\infty,$$

again using Lemma 2.7.2. This implies that for every $\varepsilon>0$ we have

$$(3\sqrt{n})^{-r|\mathtt{B}^{\mathrm{B}(e,m_n)}|^2}e^{F(W_n)n}\geq n^{-\varepsilon n}$$

for all large enough n, which implies the result.

	a_0	a_1	$\cdots$
b_0	$\to$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$
b_1	$\to$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$
$\vdots$	$\to$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$

Figure 2.1: Picking entries of the vertex measure $W_{\mathtt{AB}}(\cdot)$. First choose entries of the form $W_{\mathtt{AB}}((a,b))$ for $a \neq a_0$ by rounding down $W((a,b))$, then fill in the first column in a way that guarantees the correct $\mathtt{B}$-marginal.

2.10 Proof of Lemma 2.2.3

We show how to construct a denominator-n weight $W_{\mathtt{AB}}$ that has a given $\mathtt{B}$-marginal $W_{\mathtt{B}}$ and is close to a given $(\mathtt{A}\times\mathtt{B})$-weight W whose $\mathtt{B}$-marginal $\pi_{\mathtt{B}}W$ is close to $W_{\mathtt{B}}$. As in the theorem statement, we assume

$$d(\pi_{\mathtt{B}}W, W_{\mathtt{B}}) < \delta.$$

To minimize the appearance of factors of $\frac{1}{2}$, in this section we work with the ℓ^1 distance on weights, which is twice the distance defined above. Therefore the previous assumption becomes

$$d_1(\pi_{\mathtt{B}}W, W_{\mathtt{B}}) = \sum_{i\in[r]}\sum_{b,b'\in\mathtt{B}} |\pi_{\mathtt{B}}W(b,b';i) - W_{\mathtt{B}}(b,b';i)| < 2\delta.$$

We fix distinguished elements $a_0 \in \mathtt{A}$ and $b_0 \in \mathtt{B}$ which will be referred to throughout this section.

2.10.1 The vertex measure

We first define the weight's vertex measure by

$$\begin{aligned} W_{\mathtt{AB}}((a,b)) &= \tfrac{1}{n}\lfloor n\cdot W((a,b))\rfloor && a\in\mathtt{A}\setminus\{a_0\},\ b\in\mathtt{B} \\ W_{\mathtt{AB}}((a_0,b)) &= W_{\mathtt{B}}(b) - \sum_{a\neq a_0} W_{\mathtt{AB}}((a,b)) && b\in\mathtt{B}. \end{aligned}$$

See Figure 2.1.

Note that $|W_{\mathtt{AB}}((a,b)) - W((a,b))| \le 1/n$ for $a \neq a_0$ and

$$|W_{\mathtt{AB}}((a_0,b)) - W((a_0,b))| \le |W_{\mathtt{B}}(b) - \pi_{\mathtt{B}} W(b)| + |\mathtt{A}|/n.$$

Therefore the ℓ^1 distance between the vertex measures is

$$\begin{aligned}\sum_{a,b} |W_{\mathtt{AB}}((a,b)) - W((a,b))| &\le |\mathtt{A}||\mathtt{B}|/n + \sum_{b\in\mathtt{B}} \big(|W_{\mathtt{B}}(b) - \pi_{\mathtt{B}} W(b)| + |\mathtt{A}|/n\big) \\ &\le 2\delta + 2|\mathtt{A}||\mathtt{B}|/n.\end{aligned}$$

2.10.1.1 Nonnegativity

The terms defined by rounding down W using the floor function are guaranteed to be nonnegative, but the others are not. In the following we show how to repair any negativity.

Let $-R/n$ denote the sum of all negative terms in the vertex measure. Since W contains only nonnegative terms we have

$$\mathbf{1}_{\{W_{\mathtt{AB}}((a,b))<0\}} \cdot |W_{\mathtt{AB}}((a,b))| \le |W_{\mathtt{AB}}((a,b)) - W((a,b))| \quad \text{for all } a,b.$$

Therefore

$$R/n \le \sum_{b\in\mathtt{B}} |W_{\mathtt{AB}}((a_0,b)) - W((a_0,b))| \le 2\delta + |\mathtt{A}||\mathtt{B}|/n.$$

Suppose there is some $b \in \mathtt{B}$ such that $W_{\mathtt{AB}}((a_0,b)) < 0$. Since $W_{\mathtt{AB}}$ has denominator n, we must have $W_{\mathtt{AB}}((a_0,b)) \le -1/n$. By construction, we have

$$\sum_{a\in A} W_{\mathtt{AB}}((a,b)) = W_{\mathtt{B}}(b) \ge 0,$$

so there exists some $a^+ \in \mathtt{A}$ with $W_{\mathtt{AB}}((a^+,b)) \ge 1/n$. Increase $W_{\mathtt{AB}}((a_0,b))$ by $1/n$ and decrease $W_{\mathtt{AB}}((a^+,b))$ by $1/n$.

The number of times we must repeat this step before all terms are nonnegative is exactly R, and each step moves the measure by ℓ^1 distance $2/n$; therefore the final edited vertex measure is distance at most $2R/n$ from the original $W_{\mathtt{AB}}$. If we now let $W_{\mathtt{AB}}$ denote the new, nonnegative vertex measure, by the above bound on R/n we get

$$\sum_{a,b} |W_{\mathtt{AB}}((a,b)) - W((a,b))| \le 6\delta + 4|\mathtt{A}||\mathtt{B}|/n.$$

2.10.2 The B half-marginal

For the purposes of this construction we use the B "half-marginal," which we denote

$$W(b,(a',b');i) := \sum_{a\in\mathtt{A}} W((a,b),(a',b');i).$$

This is an element of Prob $\big((\mathtt{B}\times(\mathtt{A}\times\mathtt{B}))^r\big)$.

Before constructing the edge measure of $W_{\mathtt{AB}}$, in this section we first construct what will be its half-marginal.

For each $i\in[r]$, $b,b'\in\mathtt{B}$, and $a'\in\mathtt{A}$ we define

$$W_{\mathtt{AB}}(b,(a',b');i) = \tfrac{1}{n}\lfloor n\cdot W(b,(a',b');i)\rfloor \qquad \text{for } a'\ne a_0,\ b\ne b_0, \tag{2.2}$$

$$W_{\mathtt{AB}}(b,(a_0,b');i) = W_{\mathtt{B}}(b,b';i) - \sum_{a'\ne a_0} W_{\mathtt{AB}}(b,(a',b');i) \qquad \text{for } b\ne b_0, \tag{2.3}$$

$$W_{\mathtt{AB}}(b_0,(a',b');i) = W_{\mathtt{AB}}((a',b')) - \sum_{b\ne b_0} W_{\mathtt{AB}}(b,(a',b');i). \tag{2.4}$$

See Figure 2.2 for a representation of which terms are defined by each equation.

The definition of the terms in (2.4) ensures that

$$\sum_{b\in\mathtt{B}} W_{\mathtt{AB}}(b,(a',b');i) = W_{\mathtt{AB}}((a',b')) \quad \text{for all } a',b',i.$$

This will ensure that $W_{\mathtt{AB}}$ has the correct vertex measure. Note also that by line (2.3)

$$\sum_{a'\in\mathtt{A}} W_{\mathtt{AB}}(b,(a',b');i) = W_{\mathtt{B}}(b,b';i) \quad \text{for all } b\in\mathtt{B} \text{ and } b'\in\mathtt{B}\setminus\{b_0\}.$$

Using this and definition (2.4) we also get

$$\sum_{a'\in\mathtt{A}} W_{\mathtt{AB}}(b_0,(a',b');i) = W_{\mathtt{B}}(b_0,b';i).$$

	(a_0,b_0)	(a_1,b_0)	(a_2,b_0)	(a_0,b_1)	(a_1,b_1)	(a_2,b_1)	(a_0,b_2)	(a_1,b_2)	(a_2,b_2)
b_0	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$
b_1	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$
b_2	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$

Figure 2.2: A diagram of how the half-marginal $W_{\mathtt{AB}}(\cdot,(\cdot,\cdot);i)$ is chosen if $\mathtt{A}=\{a_0,a_1,a_2\}$ and $\mathtt{B}=\{b_0,b_1,b_2\}$. First obtain the entries marked $\lfloor\cdot\rfloor$ by rounding down W. Then choose the entries marked $\rightarrow$ according to Equation 2.3 which ensures that the $\mathtt{B}$-marginal is $W_{\mathtt{B}}$. Then choose the entries marked $\downarrow$ according to Equation 2.4 which ensures that the vertex weight is the one we chose above.

This will ensure that the $\mathtt{B}$-marginal of $W_{\mathtt{AB}}$ is $W_{\mathtt{B}}$.

We show now that the half-marginal $W_{\mathtt{AB}}(\cdot,(\cdot,\cdot);i)$ is ℓ^1-close to $W(\cdot,(\cdot,\cdot);i)$ by considering separately the contributions to the ℓ^1 distance from terms defined using Equations (2.2), (2.3), and (2.4).

(2.2) terms: Each of the terms of $W_{\mathtt{AB}}$ defined using the floor in Equation (2.2) is distance at most $1/n$ from the corresponding term of W; therefore the total contribution of these terms to the ℓ^1 distance is

$$\sum_{\substack{b\in\mathtt{B}\setminus\{b_0\}\\ a'\in\mathtt{A}\setminus\{a_0\},b'\in\mathtt{B}\\ i\in[r]}} |W_{\mathtt{AB}}(b,(a',b');i)-W(b,(a',b');i)| \le |\mathtt{A}||\mathtt{B}|^2 r/n.$$

(2.3) terms: By the triangle inequality,

$$\begin{aligned}
&|W_{\mathtt{AB}}(b,(a_0,b');i)-W(b,(a_0,b');i)|\\
&\quad=\left|\left(W_{\mathtt{B}}(b,b';i)-\sum_{a'\neq a_0}W_{\mathtt{AB}}(b,(a',b');i)\right)-\right.\\
&\qquad\qquad\left.\left(\pi_{\mathtt{B}}W(b,b';i)-\sum_{a'\neq a_0}W(b,(a',b');i)\right)\right|\\
&\quad\le|W_{\mathtt{B}}(b,b';i)-\pi_{\mathtt{B}}W(b,b';i)|+\sum_{a'\neq a_0}|W_{\mathtt{AB}}(b,(a',b');i)-W(b,(a',b');i)|.
\end{aligned}$$

The total contribution of such terms is therefore

$$\sum_{\substack{b \in \mathtt{B} \setminus \{b_0\},\ b' \in \mathtt{B} \\ i \in [r]}} |W_{\mathtt{AB}}(b,(a_0,b');i) - W(b,(a_0,b');i)|$$

$$\leq \overbrace{\sum_{\substack{b \in \mathtt{B} \setminus \{b_0\},\ b' \in \mathtt{B} \\ i \in [r]}} |W_{\mathtt{B}}(b,b';i) - (\pi_{\mathtt{B}})_* W(b,b';i)|}^{\leq d_1(W_{\mathtt{B}}, \pi_{\mathtt{B}} W)}$$

$$+ \overbrace{\sum_{\substack{b \in \mathtt{B} \setminus \{b_0\} \\ a' \in \mathtt{A} \setminus \{a_0\},\ b' \in \mathtt{B} \\ i \in [r]}} |W_{\mathtt{AB}}(b,(a',b');i) - W(b,(a',b');i)|}^{=\text{contribution from (2.2) terms}}$$

$$\leq 2\delta + |\mathtt{A}||\mathtt{B}|^2 r/n.$$

(2.4) terms: Again applying the triangle inequality,

$$|W_{\mathtt{AB}}(b_0,(a,b');i) - W(b_0,(a,b');i)|$$
$$\leq |W_{\mathtt{AB}}((a,b')) - W((a,b'))| + \sum_{b \neq b_0} |W_{\mathtt{AB}}(b,(a,b');i) - W(b,(a,b');i)|.$$

Summing over all $a \in \mathtt{A}$, $b' \in \mathtt{B}$ and $i \in [r]$, we see that the total contribution of such terms is bounded by

$$\sum_{\substack{a\in\mathtt{A},b'\in\mathtt{B}\\ i\in[r]}}\left[|W_{\mathtt{AB}}((a,b'))-W((a,b'))|+\sum_{b\neq b_0}|W_{\mathtt{AB}}(b,(a,b');i)-W(b,(a,b');i)|\right]$$

$$= \sum_{i\in[r]}\overbrace{\sum_{\substack{a\in\mathtt{A}\\ b\in\mathtt{B}}}|W_{\mathtt{AB}}((a,b))-W((a,b))|}^{\text{vertex measure}}$$

$$+\overbrace{\sum_{\substack{b\in\mathtt{B}\setminus\{b_0\}\\ a'\in\mathtt{A}\setminus\{a_0\},\ b'\in\mathtt{B}\\ i\in[r]}}|W_{\mathtt{AB}}(b,(a',b');i)-W(b,(a',b');i)|}^{\text{(2.2) terms}}$$

$$+\overbrace{\sum_{\substack{b\in\mathtt{B}\setminus\{b_0\},\ b'\in\mathtt{B}\\ i\in[r]}}|W_{\mathtt{AB}}(b,(a_0,b');i)-W(b,(a_0,b');i)|}^{\text{(2.3) terms}}$$

$$\leq r\cdot[6\delta+4|\mathtt{A}||\mathtt{B}|/n]+\left[|\mathtt{A}||\mathtt{B}|^2r/n\right]+\left[2\delta+|\mathtt{A}||\mathtt{B}|^2r/n\right]$$

$$\leq 8r\delta+6|\mathtt{A}||\mathtt{B}|^2r/n.$$

Adding up the contributions of the three types of terms, we see that the ℓ^1 distance between the half-marginals of W and $W_{\mathtt{AB}}$ is bounded by

$$10r\delta+8|\mathtt{A}||\mathtt{B}|^2r/n.$$

2.10.2.1 Nonnegativity

Again, the preceding construction does not guarantee that all terms are nonnegative. In the following we describe how to correct negativity.

Let $-R/n$ be the sum of all negative terms of the half-marginal. As above, we get

$$R/n\leq 10r\delta+7|\mathtt{A}||\mathtt{B}|^2r/n.$$

Suppose there is some $b_-\in B$, $(a'_-,b'_-)\in\mathtt{A}\times\mathtt{B}$, and $i\in[r]$ such that $W_{\mathtt{AB}}(b_-,(a'_-,b'_-);i)<$

0. Then $W_{\mathtt{AB}}(b_-, (a'_-, b'_-); i) \leq -1/n$. Since

$$\sum_{a' \in \mathtt{A}} W_{\mathtt{AB}}(b_-, (a', b'_-); i) = W_{\mathtt{B}}(b_-, b'_-; i) \geq 0$$

and

$$\sum_{b \in \mathtt{B}} W_{\mathtt{AB}}(b, (a'_-, b'_-); i) = W_{\mathtt{AB}}((a'_-, b'_-)) \geq 0$$

there exist $a'_+ \in \mathtt{A}$ and $b_+ \in \mathtt{B}$ such that

$$W_{\mathtt{AB}}(b_-, (a'_+, b'_-); i) \geq 1/n \quad \text{and} \quad W_{\mathtt{AB}}(b_+, (a'_-, b'_-); i) \geq 1/n.$$

Decrease both of these terms by $1/n$, and increase $W_{\mathtt{AB}}(b_-, (a'_-, b'_-); i)$ and $W_{\mathtt{AB}}(b_+, (a'_+, b'_-); i)$ by $1/n$. This moves the half-marginal by ℓ^1 distance $4/n$.

$$\sum_{a' \in \mathtt{A}} W_{\mathtt{AB}}(b, (a', b'); i) = W_{\mathtt{B}}(b, b'; i) \quad \text{and} \quad \sum_{b \in \mathtt{B}} W_{\mathtt{AB}}(b, (a', b'); i) = W_{\mathtt{AB}}((a', b')).$$

This step must be done at most R times to eliminate all negative entries, so the final half-marginal satisfies

$$\begin{aligned}\sum_{i \in [r]} \sum_{b \in \mathtt{B}} \sum_{(a', b') \in \mathtt{A} \times \mathtt{B}} |W_{\mathtt{AB}}(b, (a', b'); i) - W(b, (a', b'); i)| &\leq (10r\delta + 8|\mathtt{A}||\mathtt{B}|^2 r/n) + R \cdot 4/n \\ &\leq 50r\delta + 36|\mathtt{A}||\mathtt{B}|^2 r/n.\end{aligned}$$

2.10.3 The edge measure

Finally, we define the edge measure of $W_{\mathtt{AB}}$ by

$$W_{\mathtt{AB}}((a, b), (a', b'); i) = \tfrac{1}{n}\lfloor n \cdot W((a, b), (a', b'); i) \rfloor \qquad \text{for } a \neq a_0 \text{ and } (a', b') \neq (a_0, b_0), \tag{2.5}$$

$$W_{\mathtt{AB}}((a_0, b), (a', b'); i) = W_{\mathtt{AB}}(b, (a', b'); i) - \sum_{a \neq a_0} W_{\mathtt{AB}}((a, b), (a', b'); i) \qquad \text{for } (a', b') \neq (a_0, b_0), \tag{2.6}$$

$$W_{\mathtt{AB}}((a, b), (a_0, b_0); i) = W_{\mathtt{AB}}((a, b)) - \sum_{(a', b') \neq (a_0, b_0)} W_{\mathtt{AB}}((a, b), (a', b'); i). \tag{2.7}$$

	(a_0,b_0)	(a_1,b_0)	(a_2,b_0)	(a_0,b_1)	(a_1,b_1)	(a_2,b_1)	(a_0,b_2)	(a_1,b_2)	(a_2,b_2)
(a_0,b_0)	$\rightarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$
(a_1,b_0)	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$
(a_2,b_0)	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$
(a_0,b_1)	$\rightarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$
(a_1,b_1)	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$
(a_2,b_1)	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$
(a_0,b_2)	$\rightarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$	$\downarrow$
(a_1,b_2)	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$
(a_2,b_2)	$\rightarrow$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$	$\lfloor\cdot\rfloor$

Figure 2.3: A diagram of how the edge measure $W_{\mathtt{AB}}((\cdot,\cdot),(\cdot,\cdot);i)$ is chosen if $\mathtt{A} = \{a_0,a_1,a_2\}$ and $\mathtt{B} = \{b_0,b_1,b_2\}$. First obtain the entries marked $\lfloor\cdot\rfloor$ by rounding down entries of W. Then choose entries marked $\downarrow$ according to Equation 2.6, which ensures that the $\mathtt{B}$ half-marginal is the one chosen above. Then choose entries marked $\rightarrow$ according to Equation 2.7, which ensures that the vertex measure is the one chosen above.

See Figure 2.3.

It follows from this definition that $W_{\mathtt{AB}}$ is a (signed) weight with $\mathtt{B}$-marginal $W_{\mathtt{B}}$.

We now check that $W_{\mathtt{AB}}$ is ℓ^1-close to W. We consider separately the contribution to the ℓ^1 distance of terms defined in equations (2.5), (2.6), and (2.7):

(2.5) terms: Each term of $W_{\mathtt{AB}}$ defined using the floor function in equation (2.5) is distance at most $1/n$ from the corresponding W term. The total contribution of these terms to the ℓ^1 distance is therefore at most $|\mathtt{A}|^2|\mathtt{B}|^2r/n$.

(2.6) terms: Applying the triangle inequality to terms defined in equation (2.6),

$$\begin{aligned}
&|W_{\mathtt{AB}}((a_0,b),(a',b');i) - W((a_0,b),(a',b');i)| \\
&\quad \leq |W_{\mathtt{AB}}(b,(a',b');i) - W(b,(a',b');i)| \\
&\qquad + \sum_{a\neq a_0} |W_{\mathtt{AB}}((a,b),(a',b');i) - W((a,b),(a',b');i)| \\
&\quad \leq |W_{\mathtt{AB}}(b,(a',b');i) - W(b,(a',b');i)| + |\mathtt{A}|/n.
\end{aligned}$$

By the ℓ^1 bound on the distance between the half-marginals, the total contribution of all such terms is therefore

$$\begin{aligned}
&\sum_{i\in[r]}\sum_{b}\sum_{(a',b')\neq(a_0,b_0)} (|W_{\mathtt{AB}}(b,(a',b');i) - W(b,(a',b');i)| + |\mathtt{A}|/n) \\
&\quad \leq [50r\delta + 36|\mathtt{A}|^2|\mathtt{B}|^2r/n] + |\mathtt{A}|^2|\mathtt{B}|^2r/n \\
&\quad = 50r\delta + 37|\mathtt{A}|^2|\mathtt{B}|^2r/n
\end{aligned}$$

(2.7) terms: Applying the triangle inequality to terms defined in equation (2.7):

$$\begin{aligned}
&|W_{\mathtt{AB}}((a,b),(a_0,b_0);i) - W_{\mathtt{AB}}((a,b),(a_0,b_0);i)| \\
&\qquad \leq |W_{\mathtt{AB}}((a,b)) - W((a,b))| \\
&\qquad\quad + \sum_{(a',b')\neq(a_0,b_0)} |W_{\mathtt{AB}}((a,b),(a',b');i) - W((a,b),(a',b');i)|.
\end{aligned}$$

Therefore the total contribution of all such terms is

$$\sum_{i\in[r]}\sum_{a,b}|W_{\mathtt{AB}}((a,b),(a_0,b_0);i) - W_{\mathtt{AB}}((a,b),(a_0,b_0);i)|$$
$$= \sum_{i\in[r]}\sum_{a,b}\Bigg[|W_{\mathtt{AB}}((a,b)) - W((a,b))|$$
$$+ \sum_{(a',b')\neq(a_0,b_0)}|W_{\mathtt{AB}}((a,b),(a',b');i) - W((a,b),(a',b');i)|\Bigg]$$
$$= \overbrace{\sum_{i\in[r]}\sum_{a,b}|W_{\mathtt{AB}}((a,b)) - W((a,b))|}^{\text{vertex measure}}$$
$$+ \overbrace{\sum_{i\in[r]}\sum_{a\neq a_0}\sum_{b}\sum_{(a',b')\neq(a_0,b_0)}|W_{\mathtt{AB}}((a,b),(a',b');i) - W((a,b),(a',b');i)|}^{\text{(2.5) terms}}$$
$$+ \overbrace{\sum_{i\in[r]}\sum_{b}\sum_{(a',b')\neq(a_0,b_0)}|W_{\mathtt{AB}}((a_0,b),(a',b');i) - W((a_0,b),(a',b');i)|}^{\text{(2.6) terms}}\Bigg]$$
$$\leq r\cdot[6\delta + 3|\mathtt{A}||\mathtt{B}|/n] + \left[|\mathtt{A}|^2|\mathtt{B}|^2 r/n\right] + \left[50r\delta + 37|\mathtt{A}|^2|\mathtt{B}|^2 r/n\right]$$
$$\leq 56r\delta + 41|\mathtt{A}|^2|\mathtt{B}|^2 r/n.$$

Summing up the contributions from terms of all three types, we get that

$$d_1(W_{\mathtt{AB}}, W) \leq 106r\delta + 79|\mathtt{A}|^2|\mathtt{B}|^2 r/n.$$

2.10.3.1 Nonnegativity

We can modify a solution with negative entries to get a nonnegative one similarly to above. Let $-R/n$ be the sum of all negative entries; then

$$R/n \leq 106r\delta + 78|\mathtt{A}|^2|\mathtt{B}|^2 r/n.$$

Suppose there is some entry

$$W_{\mathtt{AB}}((a_-,b_-),(a'_-,b'_-);i) \leq -1/n.$$

We want to increment this term by $1/n$ without affecting the vertex measure or the $\mathtt{B}$ marginal. Since

$$\sum_{(a',b')\in\mathtt{A}\times\mathtt{B}} W_{\mathtt{AB}}((a_-,b_-),(a',b');i) = W_{\mathtt{AB}}((a_-,b_-)) \geq 0$$

there exists some $(a'_+,b'_+) \in \mathtt{A}\times\mathtt{B}$ such that $W_{\mathtt{AB}}((a_-,b_-),(a'_+,b'_+);i) \geq 1/n$; similarly since

$$\sum_{a\in A} W_{\mathtt{AB}}((a,b_-),(a',b'_-);i) = W_{\mathtt{AB}}(b_-,(a'_-,b'_-);i) \geq 0$$

there exists some a_+ such that $W_{\mathtt{AB}}((a_+,b_-),(a'_-,b'_-);i) \geq 1/n$. Increase

$$W_{\mathtt{AB}}((a_-,b_-),(a'_-,b'_-);i) \quad \text{and} \quad W_{\mathtt{AB}}((a_+,b_-),(a'_+,b'_+);i)$$

by $1/n$, and decrease

$$W_{\mathtt{AB}}((a_-,b_-),(a'_+,b'_+);i) \quad \text{and} \quad W_{\mathtt{AB}}((a_+,b_-),(a'_-,b'_-);i)$$

by $1/n$. This moves the weight by ℓ^1 distance $4/n$.

Since R is the maximum number of times we need to do this before there are no more negative entries, the final weight satisfies

$$d_1(W_{\mathtt{AB}},W) \leq 106r\delta + 79|\mathtt{A}|^2|\mathtt{B}|^2 r/n + 4R/n \leq 530r\delta + 391|\mathtt{A}|^2|\mathtt{B}|^2 r/n.$$

To simplify, we write

$$d_1(W_{\mathtt{AB}},W) \leq 530r(\delta + |\mathtt{A}\times\mathtt{B}|^2/n),$$

or

$$d(W_{\mathtt{AB}},W) \leq 265r(\delta + |\mathtt{A}\times\mathtt{B}|^2/n).$$

CHAPTER 3

Gibbs measures and Glauber dynamics

This chapter contains definitions and fundamental results which are common to Chapters 4 and 5.

As above, Γ will denote a group with a fixed set of r generators $\{s_1, \dots, s_r\}$. We will also use the symbol Γ to denote the group's left Cayley graph, which has vertex set Γ and an i-labeled directed edge $(\gamma, s_i\gamma)$ for every $i \in [r]$ and $\gamma \in \Gamma$.

For some finite alphabet $\mathtt{A}$, we define the shift action of Γ on $\mathtt{A}^\Gamma$ by

$$(\beta \mathbf{y})(\gamma) = \mathbf{y}(\gamma\beta)$$

for $\beta, \gamma \in \Gamma$. We can think of this as moving the center of the labeling to β^{-1}. We say that a measure $\mu \in \operatorname{Prob}(\mathtt{A}^\Gamma)$ is shift-invariant if $\beta_*\mu = \mu$ for any $\beta \in \Gamma$, where β_* denotes the pushforward. We denote the set of shift-invariant probability measures by $\operatorname{Prob}^\Gamma(\mathtt{A}^\Gamma)$.

If V is a finite set, we can consider the set $\operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$ of homomorphisms from Γ to the group of permutations of V. Given $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$, we write the permutation which is the image of $\gamma \in \Gamma$ by either σ^γ or $\sigma(\gamma)$. We can associate to σ a directed graph with vertex set V and an i-labeled edge $(v, \sigma^{s_i}(v))$ for each $i \in [r]$ and $v \in V$.

The graph of any σ can be thought of as a finite system which locally looks like Γ, just as a large rectangular grid locally looks like the integer lattice $\mathbb{Z}^r$. The labeling of the edges gives a canonical way to lift elements of $\mathtt{A}^V$ to elements of $\mathtt{A}^\Gamma$; see below.

Either Γ or the graph of some σ can be endowed with a natural graph distance: the distance between a pair of vertices is defined to be the minimal number of edges in a path between them, ignoring edge directions. Let $\mathrm{B}^\sigma(v,R)$ denote the closed radius-R ball centered at $v \in V$, and similarly define $\mathrm{B}^\Gamma(\gamma, R)$ for $\gamma \in \Gamma$.

Give $\mathtt{A}^\Gamma$ the metric

$$d(\mathbf{x},\mathbf{y}) = \sum_{\gamma\in\Gamma} (3r)^{-|\gamma|} \mathbf{1}_{\{\mathbf{x}(\gamma)\neq \mathbf{y}(\gamma)\}};$$

the factor 3 is chosen to ensure convergence. Note that $\operatorname{diam} \mathtt{A}^\Gamma \leq 3$. This metric induces the product topology (with $\mathtt{A}$ having the discrete topology).

Let $\bar{d}$ denote the corresponding transportation metric on $\operatorname{Prob}(\mathtt{A}^\Gamma)$ (the set of Borel probability measures); specifically, with $\operatorname{Lip}_1(\mathtt{A}^\Gamma)$ denoting the set of 1-Lipschitz real-valued functions, we define

$$\bar{d}(\mu,\nu) = \sup\left\{ |\mu f - \nu f| \,:\, f \in \operatorname{Lip}_1(\mathtt{A}^\Gamma)\right\}.$$

Here μf denotes the integral of f with respect to μ. Note that d generates the product topology on $\mathtt{A}^\Gamma$ (which is compact), and $\bar{d}$ generates the weak topology induced by the pairing with continuous functions (which is also compact).

For any set V and any $\mathbf{x} \in \mathtt{A}^V$, $v \in V$, $\mathtt{a} \in \mathtt{A}$ we let $\mathbf{x}^{v\to\mathtt{a}} \in \mathtt{A}^V$ be given by

$$\mathbf{x}^{v\to\mathtt{a}}(w) = \begin{cases} \mathbf{x}(w), & w \neq v \\ \mathtt{a}, & w = v. \end{cases}$$

Recall that an element of $\mathtt{A}^V$ is referred to as a *microstate* and an element of $\operatorname{Prob}(\mathtt{A}^V)$ as a *state*.

3.1 Interaction

Let V be an at most countable set and fix $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$. We will apply this in two cases: when V is finite, and when $V = \Gamma$ and σ is the action of Γ on itself by right multiplication. We will distinguish between these cases by giving notation superscripts of σ or Γ respectively (e.g. Ω^σ versus Ω^Γ).

A nearest-neighbor interaction with alphabet $\mathtt{A}$ is a pair $\Phi = (J, h)$ where $J \colon \mathtt{A}^2 \to \mathbb{R}$ is symmetric and $h \colon \mathtt{A} \to \mathbb{R}$. With $S = \{s_1, \ldots, s_r, s_1^{-1}, \ldots, s_r^{-1}\}$, for $v \in V$ let $\Phi_v \colon \mathtt{A}^V \to \mathbb{R}$ be given by

$$\Phi_v(\mathbf{x}) = h(\mathbf{x}(v)) + \sum_{s \in S} J(\mathbf{x}(v),\, \mathbf{x}(\sigma^s v)).$$

If V is finite then we can define the energy $U \colon \mathtt{A}^V \to \mathbb{R}$ by

$$U(\mathbf{x}) = \sum_{v \in V} h(\mathbf{x}(v)) + \sum_{v \in V} \sum_{i \in [r]} J(\mathbf{x}(v),\, \mathbf{x}(\sigma^{s_i} v)).$$

This can also be written

$$U(\mathbf{x}) = \sum_{v \in V} U_v(\mathbf{x})$$

where

$$U_v(\mathbf{x}) = h(\mathbf{x}(v)) + \frac{1}{2} \sum_{s \in S} J(\mathbf{x}(v),\, \mathbf{x}(\sigma^s v)).$$

Note that U_v can be thought of as "energy per vertex at v." In contrast, Φ_v might be described as "energy due to interactions involving v."

Note that if we define

$$u^{\max} = \max_{\mathtt{a} \in \mathtt{A}} \Big(h(\mathtt{a}) + r \max_{\mathtt{b} \in \mathtt{A}} J(\mathtt{a}, \mathtt{b}) \Big)$$

and

$$u^{\min} = \min_{\mathtt{a} \in \mathtt{A}} \Big(h(\mathtt{a}) + r \min_{\mathtt{b} \in \mathtt{A}} J(\mathtt{a}, \mathtt{b}) \Big)$$

then for any V, σ and any $\mathbf{x} \in \mathtt{A}^V$ we have

$$u^{\min} \leq \frac{1}{|V|} U(\mathbf{x}) \leq u^{\max}.$$

An Ising model with no external field has $\mathtt{A} = \{-1, 1\}$, $J(\mathbf{x}) = \beta \mathtt{ab}$, and $h \equiv 0$ for some $\beta \geq 0$ (the inverse temperature). The Bernoulli shift with base measure $p \in \operatorname{Prob}(\mathtt{A})$ also fits into this framework by taking $J \equiv 0$ and $h(\mathtt{a}) = -\log p(\{\mathtt{a}\})$.

The present framework does not include systems with hard constraints, like the 0-temperature Ising model or the hardcore model.

3.2 Glauber dynamics

For $\mathtt{a} \in \mathtt{A}$ let

$$c_v(\mathbf{x}, \mathtt{a}) = Z_v(\mathbf{x})^{-1} \exp\{-\Phi_v(\mathbf{x}^{v\to\mathtt{a}})\},$$

where $Z_v(\mathbf{x})$ is the normalizing factor which makes $c_v(\mathbf{x}, \cdot)$ a probability measure on $\mathtt{A}$. We can think of $c_v(\mathbf{x}, \cdot)$ as the transition rates for the spin at v conditioned on the current state of the system being $\mathbf{x}$. Note that this only depends on the coordinates of $\mathbf{x}$ at vertices adjacent to v.

The Glauber dynamics is the continuous-time Markov process with state space $\mathtt{A}^V$ and generator Ω given by

$$\Omega f(\mathbf{x}) = \sum_{v \in V} \sum_{\mathtt{a} \in \mathtt{A}} c_v(\mathbf{x}, \mathtt{a})[f(\mathbf{x}^{v\to\mathtt{a}}) - f(\mathbf{x})].$$

If V is finite then this gives a well-defined linear operator on $C(\mathtt{A}^V)$. Otherwise we need to first define Ω on a 'core' of 'smooth' functions for which the sum converges, then take the closure of Ω; see [Lig05] for details. The generator induces a Markov semigroup denoted $\{S(t) : t \geq 0\}$.

Given $\mathbf{x} \in \mathtt{A}^V$, random or deterministic, we let $\mathbf{x}_t$ denote the $\mathtt{A}^V$-valued random variable which is the evolution of $\mathbf{x}$ to time t.

For any continuous function $f \colon \mathtt{A}^V \to \mathbb{R}$ we interpret $S(t)f(\mathbf{x})$ as the expected value of $f(\mathbf{x}_t)$.

The semigroup also acts on probability measures, but on the right: $\mu S(t)$ is interpreted as the evolution of $\mu \in \mathrm{Prob}(\mathtt{A}^V)$ to time t. We will also often write $\mu_t \coloneqq \mu S(t)$; the relevant semigroup will typically be clear from context. The right action convention is appropriate because $[\mu S(t)]f = \mu[S(t)f]$, where μf denotes the integral of f.

Further details of the construction of the dynamics will only be needed for proofs of the following two results. The relevant details are contained in Section 3.7.

There is an approximate equivariance between the Glauber semigroups and the empirical distribution:

Theorem 3.2.1. *There is a constant* $M > 0$ *such that for any* $\mathbf{x} \in \mathtt{A}^V$, $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$, *and* $t \geq 0$

$$\bar{d}\left(S^\sigma(t) P^\sigma_{\mathbf{x}},\ P^\sigma_{\mathbf{x}} S^\Gamma(t)\right) \leq \Delta^\sigma \cdot t e^{Mt}.$$

A proof is given in Section 3.7.1.

It may be helpful to clarify that the first term on the left, $S^\sigma(t) P^\sigma_{\mathbf{x}}$, is the evolution to time t of the function $P^\sigma_\bullet \colon \mathtt{A}^V \to \mathrm{Prob}(\mathtt{A}^\Gamma)$ evaluated at $\mathbf{x} \in \mathtt{A}^V$. The second term is the evolution of the empirical distribution $P^\sigma_{\mathbf{x}}$. So this theorem says that the expected empirical distribution after running the finitary dynamics for time t is close to the result of evolving the original empirical distribution for time t, as long as σ locally looks like Γ.

We also use the following Lipschitz bound on the Markov semigroup:

Lemma 3.2.2. *If* $\mu, \nu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$ *then*

$$\bar{d}\big(\mu S^\Gamma(t),\, \nu S^\Gamma(t)\big) \leq \exp(Mt)\, \bar{d}(\mu, \nu).$$

3.3 Gibbs measures

If V is finite, the Gibbs measure $\xi_V \in \mathrm{Prob}(\mathtt{A}^V)$ is defined by

$$\xi_V\{\mathbf{x}\} = Z_V^{-1} \exp\{-U(\mathbf{x})\}$$

where Z_V is the normalizing constant. Note that

$$\xi_V(\mathbf{y}(v) = \mathtt{a} \mid \mathbf{y}(w) = \mathbf{x}(w)\ \forall w \neq v) = \frac{\exp\{-U(\mathbf{x}^{v \to \mathtt{a}})\}}{\sum_{\mathtt{b} \in \mathtt{A}} \exp\{-U(\mathbf{x}^{v \to \mathtt{b}})\}} = c_v(\mathbf{x}, \mathtt{a}),$$

since

$$\frac{\exp\{-U(\mathbf{x}^{v \to \mathtt{a}})\}}{\exp\{-U(\mathbf{x}^{v \to \mathtt{b}})\}} = \frac{\exp\{-\Phi_v(\mathbf{x}^{v \to \mathtt{a}})\}}{\exp\{-\Phi_v(\mathbf{x}^{v \to \mathtt{b}})\}}.$$

On the infinite graph Γ we must use a different approach, since the sum defining the total energy will not converge. We use a natural generalization of [Lig05, Definition IV.1.5]; see also [Geo11] for a much more general treatment of infinite-volume Gibbs measures.

Let $\mathscr{T}_\gamma$ denote the σ-algebra on $\mathtt{A}^\Gamma$ generated by the coordinate maps corresponding to all vertices except for $\gamma \in \Gamma$. We call $\mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$ a Gibbs measure if for each $\gamma \in \Gamma$ and $\mathtt{a} \in \mathtt{A}$, the function $\mathbf{y} \mapsto c_\gamma(\mathbf{y}, \mathtt{a})$ is a version of the conditional expectation $\mu(\{\mathbf{x} : \mathbf{x}(\gamma) = \mathtt{a}\} \mid \mathscr{T}_v)(\mathbf{y})$. This means that for every integrable $f \colon \mathtt{A}^\Gamma \to \mathbb{R}$ and $\gamma \in \Gamma$ we have

$$\int \sum_{\mathtt{a} \in \mathtt{A}} c_\gamma(\mathbf{x}, \mathtt{a}) f(\mathbf{x}^{\gamma \to \mathtt{a}}) \mu(d\mathbf{x}) = \int f(\mathbf{x})\, \mu(d\mathbf{x}).$$

We may also describe this relation by saying that μ is invariant under re-randomizing the spin at γ using the kernel c_γ.

Note that [Geo11] requires all finite-dimensional conditional expectations to be specified by the potential in a particular way, not just the single-site ones; see Proposition 5.5.2 below for a proof that the definitions are equivalent in this setting.

If μ is Gibbs then for any 'smooth' f

$$\mu \Omega^\Gamma f = \int \left(\sum_{\gamma, \mathtt{a}} c_\gamma(\mathbf{x}, \mathtt{a}) [f(\mathbf{x}^{\gamma \to \mathtt{a}}) - f(\mathbf{x})] \right) \mu(d\mathbf{x}) = 0.$$

It follows that $\mu \Omega^\Gamma = 0$, which means Gibbs measures are Glauber-invariant.

We will denote the set of all Gibbs measures for the interaction Φ by $\mathscr{G}(\Phi)$, or just $\mathscr{G}$ if the specific Φ is clear from context or irrelevant. The shift-invariant Gibbs measures will be denoted by $\mathscr{G}^\Gamma(\Phi)$ or $\mathscr{G}^\Gamma$.

The fact that $\mathscr{G}^\Gamma$ is a face of the simplex $\mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ will be important:

Lemma 3.3.1. *Let* $\theta \in \mathrm{Prob}(\mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma))$ *and suppose* $\int \mu\, \theta(d\mu) \in \mathscr{G}^\Gamma$. *Then* $\theta(\mathscr{G}^\Gamma) = 1$.

This is stated in the case $\Gamma = \mathbb{Z}^r$ in Georgii's book [Geo11, Theorem 14.15(c)]. The proof works just as well in our generality, and goes as follows: It suffices to show that if

$\mu, \nu \in \operatorname{Prob}(\mathtt{A}^\Gamma)$ are shift-invariant, $\mu \in \mathscr{G}^\Gamma$, and ν is absolutely continuous to μ then ν is also Gibbs. Under these assumptions, since $\nu \ll \mu$ we can write $\nu = f\mu$ for some measurable f. But since ν, μ are shift-invariant, f must be μ-a.s. equal to a shift-invariant function. Since μ is shift-invariant, the σ-algebra of shift-invariant measurable subsets of $\mathtt{A}^\Gamma$ is contained in the tail σ-algebra up to μ-null sets. Therefore f is μ-a.s. equal to a tail-measurable function. From this we can conclude that ν is Gibbs.

3.4 Good models for measures on $\mathtt{A}^\Gamma$

Let V be a finite set and let $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$. A labeling $\mathbf{x} \in \mathtt{A}^V$ is said to be a *good model* for $\mu \in \operatorname{Prob}(\mathtt{A}^\Gamma)$ over σ if the empirical distribution $P_{\mathbf{x}}^\sigma$ is close to μ in the weak topology. More precisely, we can say $\mathbf{x}$ is $\mathcal{O}$-good if $P_{\mathbf{x}}^\sigma \in \mathcal{O}$ for some weak-open neighborhood $\mathcal{O} \ni \mu$. The set of such $\mathbf{x}$ is denoted $\Omega(\sigma, \mathcal{O})$. An interpretation of this relationship is that average local quantities of the finite system are consistent with μ.

We define the empirical distribution of a state $\zeta \in \operatorname{Prob}(\mathtt{A}^V)$ by

$$P_\zeta^\sigma \coloneqq \zeta P_{\mathbf{x}}^\sigma = \int P_{\mathbf{x}}^\sigma \, \zeta(d\mathbf{x}) \in \operatorname{Prob}^\Gamma(\mathtt{A}^\Gamma)$$

and say that ζ is $\mathcal{O}$-consistent with μ (for some neighborhood $\mathcal{O} \ni \mu$) over σ if $P_\zeta^\sigma \in \mathcal{O}$. We can still interpret this in terms of averages of local quantities: now the average also involves a random microstate $\mathbf{x}$ with law ζ. We denote the set of such states by $\mathbf{\Omega}(\sigma, \mathcal{O})$. This way of lifting a finitary state is used in [Alp16]; it also essentially appears in the notion of "local convergence on average" introduced in [MMS12, Definition 2.3].

This consistency is stable under Glauber dynamics in the following sense:

Proposition 3.4.1. *Suppose $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$, $\zeta \in \operatorname{Prob}(\mathtt{A}^V)$, and $\mu \in \operatorname{Prob}(\mathtt{A}^\Gamma)$. Let ζ_t, μ_t denote their evolutions under Glauber dynamics on σ, Γ respectively. Then for any $t \geq 0$*

$$\bar{d}\big(P_{\zeta_t}^\sigma, \mu_t\big) \leq \big[\Delta^\sigma t + \bar{d}\big(P_\zeta^\sigma,\, \mu\big)\big] \exp(Mt),$$

where $M > 0$ depends only on the interaction and Γ.

We give the proof here, since it uses only results already stated:

Proof. For any $f \in \mathrm{Lip}_1(\mathtt{A}^\Gamma)$, the triangle inequality gives

$$\begin{aligned}\left|P^\sigma_{\zeta_t} f - \mu_t f\right| &\le \left|P^\sigma_{\zeta_t} f - P^\sigma_\zeta S^\Gamma(t) f\right| + \left|P^\sigma_\zeta S^\Gamma(t) f - \mu_t f\right| \\ &= \left|\zeta\big[S^\sigma(t) P^\sigma_{\mathbf{x}} f - P^\sigma_{\mathbf{x}} S^\Gamma(t) f\big]\right| + \left|P^\sigma_\zeta S^\Gamma(t) f - \mu S^\Gamma(t) f\right|.\end{aligned}$$

Using that ζ is a probability measure and the definition of $\bar{d}$, this implies the bound

$$\left|P^\sigma_{\zeta_t} f - \mu_t f\right| \le \max_{\mathbf{x} \in \mathtt{A}^V} \bar{d}\big(S^\sigma(t) P^\sigma_{\mathbf{x}},\, P^\sigma_{\mathbf{x}} S^\Gamma(t)\big) + \bar{d}\big(P^\sigma_\zeta S^\Gamma(t),\, \mu S^\Gamma(t)\big).$$

The first term may be controlled with Theorem 3.2.1 and the second with Lemma 3.2.2 to get

$$\left|P^\sigma_{\zeta_t} f - \mu_t f\right| \le \big[\Delta^\sigma t + \bar{d}\big(P^\sigma_\zeta,\, \mu\big)\big] \exp(Mt).$$

The result then follows by taking the supremum over $f \in \mathrm{Lip}_1(\mathtt{A}^\Gamma)$. □

3.5 Free energy density

Given $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$, the *free energy* of $\zeta \in \mathrm{Prob}(\mathtt{A}^V)$ is given by

$$A(\zeta) = \zeta(U) - \mathrm{H}(\zeta),$$

where $\zeta(U) = \int U(\mathbf{x})\, \zeta(d\mathbf{x})$ is the average energy and $\mathrm{H}(\zeta)$ is the Shannon entropy.

Given a sequence $\Sigma = (\sigma_n)_{n \in \mathbb{N}}$ with $\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))$ such that $|V_n| \to \infty$ and $\Delta^{\sigma_n} \to 0$, we define the *free energy density* of $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ relative to Σ by

$$\mathrm{a}_\Sigma(\mu) = \lim_{\mathcal{O} \downarrow \mu} \limsup_{n \to \infty} \inf_{\zeta \in \mathbf{\Omega}(\sigma_n, \mathcal{O})} \frac{1}{|V_n|} A(\zeta).$$

We follow the convention that the infimum of the empty set is $+\infty$.

The outermost limit is over the net of weak-open neighborhoods of μ, partially ordered by inclusion. Note that for each n the expression $\inf_{\zeta \in \mathbf{\Omega}(\sigma_n, \mathcal{O})} \frac{1}{|V_n|} A(\zeta)$ is nondecreasing as $\mathcal{O} \downarrow \mu$, so the limit exists and is equal to the supremum over $\mathcal{O} \ni \mu$.

It is straightforward to check from the definitions that $\zeta(U)/|V| = P^{\sigma}_{\zeta}(U_e)$. Consequently we have

$$\mathrm{a}_\Sigma(\mu) = \mu(U_e) - \mathrm{h}^{\mathrm{mod}}_\Sigma(\mu),$$

where $\mathrm{h}^{\mathrm{mod}}$ is the 'modified sofic entropy' in [Alp16], except with a lim inf instead of lim sup. Since this connection will not be used below, we omit the proof. It may also be interesting to investigate other sofic free energy densities with $\mathrm{h}^{\mathrm{mod}}$ replaced by a different type of sofic entropy.

Since the map $\mu \mapsto \mathrm{a}_\Sigma(\mu)$ is defined in terms of a supremum over neighborhoods of μ, it is lower semi-continuous. Consequently, it attains its minimum on $\mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$.

Note also that as long as $\mathbf{\Omega}(\sigma_n, \mathcal{O})$ is nonempty, for every ζ in this set we have

$$u^{\min} - \log|\mathtt{A}| \leq \frac{1}{|V_n|}\big[\zeta(U) - \mathrm{H}(\zeta)\big] \leq u^{\max}.$$

In particular,

$$\mathrm{a}_\Sigma(\mu) \in \big[u^{\min} - \log|\mathtt{A}|,\ u^{\max}\big] \cup \{+\infty\}.$$

The case $\mathrm{a}_\Sigma(\mu) = +\infty$ can actually occur, for example if μ is ergodic and the sofic entropy relative to Σ is $-\infty$. This is because, in the ergodic case, ζ being consistent with μ is the same as being mostly supported on labelings which are good models for μ, but if the sofic entropy is $-\infty$ then there are no good models.

Note, however, that the function $\mu \mapsto \mathrm{a}_\Sigma(\mu)$ is not identically $+\infty$ for any choice of Σ, since the point mass at a constant labeling in $\mathtt{A}^\Gamma$ always has good models.

If μ were not shift-invariant then the expression defining $\mathrm{a}_\Sigma(\mu)$ would still make sense, but would take the value $+\infty$ for any Σ. This is because empirical distributions are always shift-invariant and $\mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ is closed so, no matter what Σ we choose, any small enough neighborhood of μ contains no empirical distributions. In fact, we will see below that for some Γ (for example $\Gamma = \mathbb{F}_2 \times \mathbb{F}_2$) there even exist shift-invariant measures which cannot be approximated by empirical distributions over any Σ. In these cases the obstruction is that empirical distributions always have finite support.

Following [Bow03], a shift-invariant measure in $\mathrm{Prob}^{\Gamma}(\mathtt{A}^{\Gamma})$ is called *periodic* if it has finite support, and a group Γ is said to have *property* PA if the set of periodic measures is dense in $\mathrm{Prob}^{\Gamma}(\mathtt{A}^{\Gamma})$ for every finite alphabet $\mathtt{A}$; in other words, if every shift-invariant probability measure on $\mathtt{A}^{\Gamma}$ has Periodic Approximations.

In Section 4.3 we prove the following:

Proposition 3.5.1. *A group Γ has property* PA *if and only if for every finite alphabet $\mathtt{A}$ and every $\mu \in \mathrm{Prob}^{\Gamma}(\mathtt{A}^{\Gamma})$ there exists a sequence $\Sigma = \big(\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))\big)_{n\in\mathbb{N}}$ such that $\Delta^{\sigma_n} \to 0$ and $\mathrm{a}_{\Sigma}(\mu) < +\infty$.*

Property PA was proved to hold for free groups by Bowen in [Bow03, Theorem 3.4]. Kechris later studied another equivalent property in [Kec12] which he called "MD"; see also the survey [BK20] for more recent information on which groups are known to have this property. In particular:

- Amenable groups have property PA.
- If two nontrivial groups are both either finite or property-PA, then their free product has property PA.
- The recent negative solution to Connes' embedding conjecture [JNV20] implies that the direct product $\mathbb{F}_2 \times \mathbb{F}_2$ does *not* have property PA.

3.5.0.1 Calculation

In some cases it can be possible to calculate the free energy density. Fix a sequence Σ with $\Delta^{\sigma_n} \to 0$, and for each n let $\xi_n \in \mathrm{Prob}(\mathtt{A}^{V_n})$ denote the unique finitary Gibbs measure. Suppose that $P^{\sigma_n}_{\xi_n} \xrightarrow{wk} \mu$. Then for any $\mathcal{O} \ni \mu$ we have $\xi_n \in \mathbf{\Omega}(\sigma_n, \mathcal{O})$ for all large enough n. But ξ_n, by virtue of being the Gibbs measure, has minimal free energy among all probability measures on $\mathtt{A}^{V_n}$, so

$$\inf_{\zeta \in \mathbf{\Omega}(\sigma_n, \mathcal{O})} A(\zeta) = A(\xi_n) = -\log Z_n$$

where Z_n is the normalizing constant appearing in the definition of ξ_n. It follows that

$$\mathrm{a}_\Sigma(\mu) = -\liminf_{n\to\infty} \frac{1}{|V_n|} \log Z_n.$$

The main theorem of [MMS12] implies that, for the Ising model with no external field on a regular tree, the weak limit of $P_{\xi_n}^{\sigma_n}$ is $\frac{1}{2}(\mu_+ + \mu_-)$ for any Σ, where μ_+, μ_- are the Gibbs measures with $+,-$ boundary conditions respectively. In this case $\frac{1}{|V_n|}\log Z_n$ actually converges, and the limit can be written down explicitly; see [DM10].

3.6 Measuring non-Gibbs-ness

As in [Hol71], we make use of the function

$$F_0(s) = \begin{cases} s - s\log s - 1, & s > 0 \\ -1, & s = 0 \end{cases}$$

which appears in an expression for the time derivative of free energy [Proposition 4.2.1]. This function is concave, nonpositive, and equal to 0 if and only if $s = 1$. A graph is included in Figure 3.1. If $\mu_R \in \operatorname{Prob}(\mathtt{A}^{\mathrm{B}(e,R)})$ has full support, we define

$$\Delta^R_{\mathtt{a}}(\mu_R) = \sum_{\mathbf{y}\in\mathtt{A}^{\mathrm{B}(e,R)}} \mu_R\{\mathbf{y}\} \cdot F_0\left(\frac{\exp\{-\Phi_e(\mathbf{y})\}}{\exp\{-\Phi_e(\mathbf{y}^{e\to\mathtt{a}})\}} \frac{\mu_R\{\mathbf{y}^{e\to\mathtt{a}}\}}{\mu_R\{\mathbf{y}\}} \right).$$

This measures the average failure of μ_R to be consistent with the Gibbs specification.

Lemma 3.6.1. *Suppose $\mu \in \operatorname{Prob}(\mathtt{A}^\Gamma)$ is a translation-invariant measure such that for every $R \geq 1$ the marginal $\mu_R \in \operatorname{Prob}(\mathtt{A}^{\mathrm{B}(e,R)})$ has full support and $\Delta^R_{\mathtt{a}}(\mu_R) = 0$ for every $\mathtt{a} \in \mathtt{A}$. Then μ is Gibbs.*

Proof. Fix $R \geq 1$, and let $\mathscr{S}_{\mathrm{B}(e,R)\setminus\{e\}}$ denote the σ-algebra generated by sites in $\mathrm{B}(e,R)\setminus\{e\}$. Then by definition of conditional expectation, for any $\mathtt{a} \in \mathtt{A}$

$$\mu\big(\{\mathbf{x} \in \mathtt{A}^\Gamma \,:\, \mathbf{x}(e) = \mathtt{a}\} \mid \mathscr{S}_{\mathrm{B}(e,R)\setminus\{e\}}\big)(\mathbf{y}) = \frac{\mu_R\{\mathbf{y}^{e\to\mathtt{a}}\}}{\sum_{\mathtt{b}\in\mathtt{A}} \mu_R\{\mathbf{y}^{e\to\mathtt{b}}\}},$$

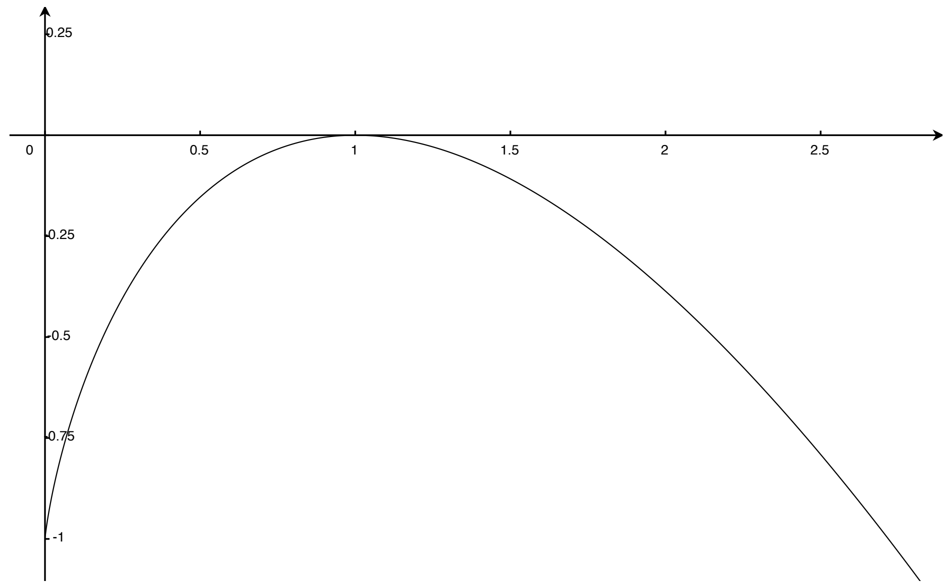

Figure 3.1: Graph of F_0

where on the right-hand side we use the shorthand $\mu_R\{\mathbf{y}\} = \mu_R\{\mathbf{y}\restriction_{\mathrm{B}(e,R)}\}$ for $\mathbf{y} \in \mathtt{A}^\Gamma$. Our assumption that μ_R has full support and $\Delta^R_{\mathtt{a}}(\mu_R) = 0$ implies that

$$\frac{\mu_R\{\mathbf{y}^{e\to\mathtt{a}}\}}{\sum_{\mathtt{b}\in\mathtt{A}} \mu_R\{\mathbf{y}^{e\to\mathtt{b}}\}} = \left(\sum_{\mathtt{b}\in\mathtt{A}} \frac{\mu_R\{\mathbf{y}^{e\to\mathtt{b}}\}}{\mu_R\{\mathbf{y}^{e\to\mathtt{a}}\}}\right)^{-1} = \left(\sum_{\mathtt{b}\in\mathtt{A}} \frac{\exp\{-\Phi_e(\mathbf{y}^{e\to\mathtt{b}})\}}{\exp\{-\Phi_e(\mathbf{y}^{e\to\mathtt{a}})\}}\right)^{-1} = c_e(\mathbf{y}, \mathtt{a})$$

Taking R to infinity, by martingale convergence we get

$$\mu(\cdot \mid \mathscr{T}_e)(\mathbf{y}) = c_e(\mathbf{y}, \cdot).$$

By translation invariance, it follows that μ is Gibbs. $\square$

3.7 Proofs of statements involving infinitary dynamics

We first give some additional setup regarding the Glauber dynamics on $\mathtt{A}^\Gamma$. First, recall that on an infinite graph we must first define the Markov generator on a 'core' of 'smooth' functions. Let $C(\mathtt{A}^\Gamma)$ denote the space of continuous real-valued functions on $\mathtt{A}^\Gamma$, with the supremum norm $\|\cdot\|_\infty$. The smooth functions are defined as follows: given $f \in C(\mathtt{A}^\Gamma)$ and

$v \in \Gamma$, let

$$\Delta_f(v) = \sup\{|f(\eta_1) - f(\eta_2)| \,:\, \eta_1(u) = \eta_2(u)\ \forall u \neq v\},$$

$$|||f||| = \sum_{v\in\Gamma} \Delta_f(v),$$

and

$$D(\mathtt{A}^\Gamma) = \{f \in C(\mathtt{A}^\Gamma) \,:\, |||f||| < \infty\}.$$

Every function which depends on only finitely many coordinates is in $D(\mathtt{A}^\Gamma)$, so $D(\mathtt{A}^\Gamma)$ is dense in $C(\mathtt{A}^\Gamma)$. For every $f \in D(\mathtt{A}^\Gamma)$ the series defining Ωf converges absolutely, and $\Omega f \in C(\mathtt{A}^\Gamma)$.

Note that the condition $|||f||| < \infty$ does not imply that f is continuous; in fact for every tail-measurable f we have $|||f||| = 0$.

Continuing to follow mostly the notation from Liggett's book, let

$$c_u(v) = \sup\left\{\|c_u(\eta_1,\cdot) - c_u(\eta_2,\cdot)\|_{TV} \,:\, \eta_1(\gamma) = \eta_2(\gamma)\ \forall \gamma \neq v\right\}.$$

Then

$$\Theta\beta(u) = \sum_{v\in\Gamma} \beta(v) c_u(v)$$

defines a bounded linear operator on $\ell^1(\Gamma)$.

The closure $\bar{\Omega}^\Gamma$ is a Markov generator, so its domain is a dense subset of the continuous functions $C(\mathtt{A}^\Gamma)$ and the range of $I - \lambda\bar{\Omega}^\Gamma$ is all of $C(\mathtt{A}^\Gamma)$ for all $\lambda \geq 0$. We also have $\|f\| \leq \|(I - \lambda\bar{\Omega}^\Gamma)f\|$ for all $\lambda \geq 0$ [Lig05, comment after Definition 2.1]. In particular $I - \lambda\bar{\Omega}^\Gamma$ is injective. An important consequence is that we have a contraction $(I - \lambda\bar{\Omega}^\Gamma)^{-1} \colon C(\mathtt{A}^\Gamma) \to C(\mathtt{A}^\Gamma)$.

3.7.1 Approximate equivariance (Proof of Theorem 3.2.1)

For $\beta \in \mathbb{R}^\Gamma$, let

$$\|\beta\| = \sup_{\gamma\in\Gamma} |\beta(\gamma)|(3r)^{|\gamma|}.$$

Lemma 3.7.1. *For any continuous* $g\colon \mathtt{A}^\Gamma \to \mathbb{R}$*,*

$$\frac{1}{3}|||g||| \le \|\Delta_g\| = |g|_{\mathrm{Lip}}.$$

In particular, every Lipschitz function is in $D(\mathtt{A}^\Gamma)$.

Proof. For the inequality:

$$|||g||| = \sum_\gamma \Delta_g(\gamma) \le \left(\sup_\gamma \Delta_g(\gamma)(3r)^{|\gamma|}\right)\sum_\gamma (3r)^{-|\gamma|} \le 3\sup_\gamma \Delta_g(\gamma)(3r)^{|\gamma|}.$$

Now similarly, for any $\mathbf{x}, \mathbf{y} \in \mathtt{A}^\Gamma$

$$\begin{aligned}|g(\mathbf{x}) - g(\mathbf{y})| &\le \sum_{\gamma\,:\,\mathbf{x}(\gamma)\neq\mathbf{y}(\gamma)} \Delta_g(\gamma) && \text{(using continuity)}\\ &\le \sup_\gamma \Delta_g(\gamma)(3r)^{|\gamma|} \sum_{\gamma\,:\,\mathbf{x}(\gamma)\neq\mathbf{y}(\gamma)} (3r)^{-|\gamma|}\\ &= \sup_\gamma \Delta_g(\gamma)(3r)^{|\gamma|}\cdot d(\mathbf{x},\mathbf{y}),\end{aligned}$$

so

$$|g|_{\mathrm{Lip}} \le \sup_\gamma \Delta_g(\gamma)(3r)^{|\gamma|} = \|\Delta_g\|.$$

The converse inequality follows from the fact that for any $\gamma \in \Gamma$

$$\Delta_g(\gamma) \le (3r)^{-|\gamma|}|g|_{\mathrm{Lip}}. \qquad \square$$

Lemma 3.7.2. *With respect to the norm* $\|\cdot\|$ *on* $\mathbb{R}^\Gamma$*,* Θ *has operator norm at most*

$$M \coloneqq \sup_\gamma \sum_{h\in\Gamma} c_h(\gamma)(3r)^{d(h,\gamma)} < \infty.$$

Proof. For any $\gamma \in \Gamma$,

$$\begin{aligned}[\Theta\beta](\gamma)\cdot(3r)^{|\gamma|} &\le \sum_{h\in\Gamma}|\beta(h)|c_h(\gamma)(3r)^{|\gamma|}\\ &\le \sum_{h\in\Gamma}|\beta(h)|c_h(\gamma)(3r)^{|h|+d(h,\gamma)}\\ &\le \|\beta\|\sum_{h\in\Gamma}c_h(\gamma)(3r)^{d(h,\gamma)},\end{aligned}$$

so, taking the supremum over γ, we see that $\|\Theta\beta\| \le \|\beta\| M$ and hence the operator norm is bounded by M.

Finiteness of M follows from the fact that always $c_h(\gamma) \le 2$, and $c_h(\gamma) = 0$ if h, γ are not adjacent. So for any γ

$$\sum_{h\in\Gamma} c_h(\gamma)(3r)^{d(h,\gamma)} \le 2 \cdot 2r \cdot (3r)^1 = 12r^2. \qquad \square$$

We can now give a proof of Lemma 3.2.2:

Proof of Lemma 3.2.2. By [Lig05, Theorem 3.9(c)],

$$\Delta_{S^\Gamma(t)f} \le \exp(t\Theta)\Delta_f \quad \forall f \in D(\mathtt{A}^\Gamma).$$

Taking $\|\cdot\|$ norms of both sides gives, by Lemma 3.7.1,

$$|S^\Gamma(t)f|_{\text{Lip}} \le \exp(Mt)|f|_{\text{Lip}}.$$

The result follows from this and the definition of $\bar{d}$. $\square$

Proposition 3.7.3. *For all small enough $\lambda > 0$, for all $g \in D(\mathtt{A}^\Gamma)$ we have*

$$|(I - \lambda\bar{\Omega})^{-k}g|_{\text{Lip}} \le [1 - \lambda M]^{-k}|g|_{\text{Lip}}.$$

Proof. Recall from [Lig05, proof of Theorem 3.9] that for all small enough $\lambda > 0$

$$\Delta_{(I-\lambda\bar{\Omega})^{-k}g} \le [(1+\lambda\varepsilon)I - \lambda\Theta]^{-k}\Delta_g$$

for any $g \in D(\mathtt{A}^\Gamma)$. If we apply the $\|\cdot\|$ norm to both sides we get, by Lemma 3.7.1,

$$|(I - \lambda\bar{\Omega})^{-k}g|_{\text{Lip}} \le [1 - \lambda(M - \varepsilon)]^{-k}|g|_{\text{Lip}}.$$

The stated bound follows after dropping ε, which is positive. $\square$

Define

$$|||f|||_R = \sum_{|\gamma|\ge R} \Delta_f(\gamma).$$

If $f \in D(\mathtt{A}^\Gamma)$ then $\lim_{R\to\infty} \| f \|_R = 0$; if f is Lipschitz then for any $R \geq 0$

$$\| f \|_R \leq 3|f|_{\mathrm{Lip}}(2/3)^R.$$

The following result establishes an approximate equivariance of $P^\sigma_{\mathbf{x}}$ with the generator:

Proposition 3.7.4. *Let V be a finite set and $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$. For any $f \in D(\mathtt{A}^\Gamma)$, $R \in \mathbb{N}$, and $\mathbf{x} \in \mathtt{A}^V$,*

$$\left|\Omega^\sigma P^\sigma_{\mathbf{x}} f - P^\sigma_{\mathbf{x}} \Omega^\Gamma f\right| \leq 3\| f \|_R + 2\delta^\sigma_R \| f \|.$$

Proof. From the definitions of Ω^σ and $P^\sigma_{\mathbf{x}}$,

$$\begin{aligned}\Omega^\sigma P^\sigma_{\mathbf{x}} f &= \sum_{v\in V}\sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a})\left[P^\sigma_{\mathbf{x}^{v\to\mathtt{a}}} f - P^\sigma_{\mathbf{x}} f\right] \\ &= \frac{1}{|V|}\sum_{v,w\in V}\sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a})\left[f(\Pi^\sigma_w \mathbf{x}^{v\to\mathtt{a}}) - f(\Pi^\sigma_w \mathbf{x})\right].\end{aligned}$$

We can compare this to if the sum is restricted to pairs v, w which are nearby in the graph σ:

$$\begin{aligned}&\left|\Omega^\sigma P^\sigma_{\mathbf{x}} f - \frac{1}{|V|}\sum_{w\in V}\sum_{v\in \mathrm{B}^\sigma(w,R-1)}\sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a})\left[f(\Pi^\sigma_w \mathbf{x}^{v\to\mathtt{a}}) - f(\Pi^\sigma_w \mathbf{x})\right]\right| \\ &\qquad \leq \frac{1}{|V|}\sum_{w\in V}\sum_{v\notin \mathrm{B}^\sigma(w,R-1)}\sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a})\left|f(\Pi^\sigma_w \mathbf{x}^{v\to\mathtt{a}}) - f(\Pi^\sigma_w \mathbf{x})\right|\end{aligned}$$

Now since the labelings $\mathbf{x}^{v\to\mathtt{a}}$ and $\mathbf{x}$ differ only at v, their lifted labelings $\Pi^\sigma_w \mathbf{x}^{v\to\mathtt{a}}$ and $\Pi^\sigma_w \mathbf{x}$ differ only at preimages of v under the map $\gamma \mapsto \sigma^\gamma w$. Let $\Pi^\sigma_w\{v\} \subset \Gamma$ denote the set of these preimages. Then the above is bounded by

$$\frac{1}{|V|}\sum_{w\in V}\sum_{v\notin \mathrm{B}^\sigma(w,R-1)}\sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a}) \sum_{\gamma\in\Pi^\sigma_w\{v\}} \Delta_f(\gamma) = \frac{1}{|V|}\sum_{w\in V}\sum_{v\notin \mathrm{B}^\sigma(w,R-1)}\sum_{\gamma\in\Pi^\sigma_w\{v\}} \Delta_f(\gamma).$$

Since for each w the sets in the collection $\{\Pi^\sigma_w\{v\} \,:\, v \notin \mathrm{B}^\sigma(w, R-1)\}$ are disjoint and contained in the complement of $\mathrm{B}^\Gamma(e, R-1)$, we can bound this by

$$\frac{1}{|V|}\sum_{w\in V}\sum_{\gamma\notin \mathrm{B}^\Gamma(e,R-1)} \Delta_f(\gamma) = \| f \|_R.$$

Now suppose $w \in V$ is such that $\mathrm{B}^\sigma(w,R) \cong \mathrm{B}^\Gamma(e,R)$: then for each $v \in \mathrm{B}^\sigma(w,R-1)$ the intersection $\mathrm{B}^\Gamma(e,R-1) \cap \Pi^\sigma_w\{v\}$ consists of a single point, which we call γ^v. We then have $c_v(\mathbf{x},\mathtt{a}) = c_{\gamma^v}(\Pi^\sigma_w \mathbf{x},\mathtt{a})$. From this we can get

$$\Bigg| \sum_{v\in \mathrm{B}^\sigma(w,R-1)} \sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a})\big[f(\Pi^\sigma_w \mathbf{x}^{v\to\mathtt{a}}) - f(\Pi^\sigma_w\mathbf{x})\big] \\ - \sum_{\gamma\in \mathrm{B}^\Gamma(e,R-1)} \sum_{\mathtt{a}\in\mathtt{A}} c_\gamma(\Pi^\sigma_w\mathbf{x},\mathtt{a})\big[f\big((\Pi^\sigma_w\mathbf{x})^{\gamma\to\mathtt{a}}\big) - f\big(\Pi^\sigma_w\mathbf{x}\big)\big]\Bigg| \qquad (*)$$

$$= \Bigg| \sum_{v\in \mathrm{B}^\sigma(w,R-1)} \sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a})[f(\Pi^\sigma_w \mathbf{x}^{v\to\mathtt{a}}) - f\big((\Pi^\sigma_w\mathbf{x})^{\gamma^v\to\mathtt{a}}\big)]\Bigg|$$

$$\leq \sum_{v\in \mathrm{B}^\sigma(w,R-1)} \sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a})\big|f(\Pi^\sigma_w \mathbf{x}^{v\to\mathtt{a}}) - f\big((\Pi^\sigma_w x)^{\gamma^v\to\mathtt{a}}\big)\big|$$

Now our construction also guarantees that the labelings $\Pi^\sigma_w \mathbf{x}^{v\to\mathtt{a}}$ and $(\Pi^\sigma_w\mathbf{x})^{\gamma^v\to\mathtt{a}}$ differ only at sites in $\Pi^\sigma_w\{v\}$ other than γ^v, all of which lie outside $\mathrm{B}^\Gamma(e,R-1)$. Therefore we can continue

$$\begin{aligned}(*) &\leq \sum_{v\in \mathrm{B}^\sigma(w,R-1)} \sum_{\mathtt{a}\in\mathtt{A}} c_v(\mathbf{x},\mathtt{a}) \sum_{\substack{\gamma\in\Pi^\sigma_w\{v\}\\ |\gamma|\geq R}} \Delta_f(\gamma)\\ &\leq \sum_{|\gamma|\geq R} \Delta_f(\gamma)\\ &= |||f|||_R.\end{aligned}$$

In the second-to-last line, we have again used that $\Pi^\sigma_w\{v_1\}$ and $\Pi^\sigma_w\{v_2\}$ are disjoint if $v_1 \neq v_2$.

For other 'bad' w where $\mathrm{B}^\sigma(w,R) \ncong \mathrm{B}^\Gamma(e,R)$, approximating the sum over v by the sum over γ in this way may be inaccurate, but the fraction of $w \in V$ which are bad is only δ^σ_R. For these w we note that the magnitudes of both sums in $(*)$ can be bounded by $\sum_{\gamma\in\Gamma}\Delta_f(\gamma) = |||f|||$.

So far we have shown that

$$\Bigg|\Omega^\sigma P^\sigma_{\mathbf{x}} f - \frac{1}{|V|}\sum_{w\in V}\sum_{|\gamma|<R}\sum_{\mathtt{a}\in\mathtt{A}} c_\gamma(\Pi^\sigma_w\mathbf{x},\mathtt{a})\big[f\big((\Pi^\sigma_w\mathbf{x})^{\gamma\to\mathtt{a}}\big) - f\big(\Pi^\sigma_w\mathbf{x}\big)\big]\Bigg| \leq 2|||f|||_R + 2\delta^\sigma_R|||f|||.$$

To finish, we compare the second term on the left-hand side to $P^{\sigma}_{\mathbf{x}}\Omega^{\Gamma} f$ using an approach similar to above:

$$\begin{aligned}
&\Bigg| \frac{1}{|V|} \sum_{w \in V} \sum_{|\gamma| < R} \sum_{\mathtt{a} \in \mathtt{A}} c_\gamma(\Pi^\sigma_k \mathbf{x}, \mathtt{a}) \big[f\big((\Pi^\sigma_k \mathbf{x})^{\gamma \to \mathtt{a}} \big) - f\big(\Pi^\sigma_k \mathbf{x} \big) \big] \\
&\qquad - \frac{1}{|V|} \sum_{w \in V} \sum_{\gamma \in \Gamma} \sum_{\mathtt{a} \in \mathtt{A}} c_\gamma(\Pi^\sigma_w \mathbf{x}, \mathtt{a}) \big[f\big((\Pi^\sigma_w \mathbf{x})^{\gamma \to \mathtt{a}} \big) - f\big(\Pi^\sigma_w \mathbf{x} \big) \big] \Bigg| \\
&\le \frac{1}{|V|} \sum_{w \in V} \sum_{|\gamma| \ge R} \sum_{\mathtt{a} \in \mathtt{A}} c_\gamma(\Pi^\sigma_w \mathbf{x}, \mathtt{a}) \big| f\big((\Pi^\sigma_w \mathbf{x})^{\gamma \to \mathtt{a}} \big) - f\big(\Pi^\sigma_w \mathbf{x} \big) \big| \\
&\le \frac{1}{|V|} \sum_{w \in V} \sum_{|\gamma| \ge R} \Delta_f(\gamma) \\
&= \vert\!\vert\!\vert f \vert\!\vert\!\vert_R. \qquad \square
\end{aligned}$$

Lemma 3.7.5. *For all small enough $\lambda > 0$, for any $m \in \mathbb{N}$ and $g \in \mathrm{Lip}(\mathtt{A}^\Gamma)$ we have*

$$\|(I - \lambda\Omega^\sigma)^{-m} P^\sigma_{\mathbf{x}} g - P^\sigma_{\mathbf{x}} (I - \lambda\bar{\Omega}^\Gamma)^{-m} g\|_{\ell^\infty(\mathtt{A}^V)} \le \lambda \Delta^\sigma |g|_{\mathrm{Lip}} \sum_{k=1}^{m} (1 - \lambda M)^{-k}.$$

Proof. We use induction on m, starting with the base case $m = 1$. Throughout, we assume λ is small enough for Proposition 3.7.3 to apply.

Given $g \in \mathrm{Lip}(\mathtt{A}^\Gamma)$, let $f = (I - \lambda\bar{\Omega}^\Gamma)^{-1} g$. Then for any $R \in \mathbb{N}$

$$\begin{aligned}
\|P^\sigma_{\mathbf{x}} g - (I - \lambda\Omega^\sigma)[P^\sigma_{\mathbf{x}} f]\|_{\ell^\infty(\mathtt{A}^V)} &= \|P^\sigma_{\mathbf{x}}[f - \lambda\bar{\Omega}^\Gamma f] - (I - \lambda\Omega^\sigma)[P^\sigma_{\mathbf{x}} f]\|_{\ell^\infty(\mathtt{A}^V)} \\
&= \lambda \|\Omega^\sigma P^\sigma_{\mathbf{x}} f - P^\sigma_{\mathbf{x}} \bar{\Omega}^\Gamma f\|_{\ell^\infty(\mathtt{A}^V)} \\
&\le \lambda (3 \vert\!\vert\!\vert f \vert\!\vert\!\vert_R + 2\delta^\sigma_R \vert\!\vert\!\vert f \vert\!\vert\!\vert) && \text{(Prop. 3.7.4)} \\
&\le \lambda (9 \cdot (2/3)^R + 6\delta^\sigma_R) |f|_{\mathrm{Lip}}.
\end{aligned}$$

Taking the infimum over R gives

$$\begin{aligned}
\|P^\sigma_{\mathbf{x}} g - (I - \lambda\Omega^\sigma)[P^\sigma_{\mathbf{x}} f]\|_{\ell^\infty(\mathtt{A}^V)} &\le \lambda \Delta^\sigma |f|_{\mathrm{Lip}} \\
&\le \lambda \Delta^\sigma (1 - \lambda M)^{-1} |g|_{\mathrm{Lip}}. && \text{(Prop. 3.7.3)}
\end{aligned}$$

Since $(I - \lambda\Omega^\sigma)^{-1}$ is a contraction on $\ell^\infty(\mathtt{A}^V)$,

$$\begin{aligned}
\|(I - \lambda\Omega^\sigma)^{-1} P^\sigma_{\mathbf{x}} g - P^\sigma_{\mathbf{x}} (I - \lambda\bar{\Omega}^\Gamma)^{-1} g\|_{\ell^\infty} &= \|(I - \lambda\Omega^\sigma)^{-1} \left[P^\sigma_{\mathbf{x}} g - (I - \lambda\Omega^\sigma)[P^\sigma_{\mathbf{x}} f] \right]\|_{\ell^\infty} \\
&\le \lambda \Delta^\sigma (1 - \lambda M)^{-1} |g|_{\mathrm{Lip}}.
\end{aligned}$$

This proves the base case.

Now assuming the m case and the base case, we prove the $m+1$ case:

$$\begin{aligned}
&\|(I-\lambda\Omega^\sigma)^{-(m+1)}P^\sigma_{\mathbf{x}}g - P^\sigma_{\mathbf{x}}(I-\lambda\bar\Omega^\Gamma)^{-(m+1)}g\|_{\ell^\infty} \\
&\qquad = \Big\|(I-\lambda\Omega^\sigma)^{-1}\left[(I-\lambda\Omega^\sigma)^{-m}P^\sigma_{\mathbf{x}}g - P^\sigma_{\mathbf{x}}(I-\lambda\bar\Omega^\Gamma)^{-m}g\right] \\
&\qquad\quad + \left[(I-\lambda\Omega^\sigma)^{-1}P^\sigma_{\mathbf{x}}(I-\lambda\bar\Omega^\Gamma)^{-m}g - P^\sigma_{\mathbf{x}}(I-\lambda\bar\Omega^\Gamma)^{-1}(I-\lambda\bar\Omega^\Gamma)^{-m}g\right]\Big\|_{\ell^\infty} \\
&\qquad \le \|(I-\lambda\Omega^\sigma)^{-m}P^\sigma_{\mathbf{x}}g - P^\sigma_{\mathbf{x}}(I-\lambda\bar\Omega^\Gamma)^{-m}g\|_\infty && \text{(contraction)} \\
&\qquad\quad + \lambda\Delta^\sigma(1-\lambda M)^{-1}|(I-\lambda\bar\Omega^\Gamma)^{-m}g|_{\mathrm{Lip}} && \text{(base case)} \\
&\qquad \le \lambda\Delta^\sigma|g|_{\mathrm{Lip}}\sum_{k=1}^{m+1}(1-\lambda M)^{-k}. && \text{(inductive hyp., Prop. 3.7.3)}
\end{aligned}$$

This completes the induction. □

Proof of Theorem 3.2.1

Given $g \in \mathrm{Lip}_1(\mathtt{A}^\Gamma)$, for all large enough m we can apply the previous lemma with $\lambda = t/m$, which gives

$$\|(I-\tfrac{t}{m}\Omega^\sigma)^{-m}P^\sigma_{\mathbf{x}}g - P^\sigma_{\mathbf{x}}(I-\tfrac{t}{m}\bar\Omega^\Gamma)^{-m}g\|_{\ell^\infty(\mathtt{A}^V)} \le \tfrac{t}{m}\Delta^\sigma|g|_{\mathrm{Lip}}\sum_{k=1}^{m}(1-\tfrac{t}{m}M)^{-k}.$$

Let $m\to\infty$. The left-hand side converges (by Hille-Yosida; [Lig05, Theorem 2.9(b)]) to $\|S^\sigma(t)P^\sigma_{\mathbf{x}}g - P^\sigma_{\mathbf{x}}S^\Gamma(t)g\|_\infty$ while the lim sup of the right-hand side is bounded by $\Delta^\sigma te^{Mt}|g|_{\mathrm{Lip}}$, since

$$\limsup_{m\to\infty}\frac{1}{m}\sum_{k=1}^{m}(1-\tfrac{t}{m}M)^{-k} \le \limsup_{m\to\infty}\frac{1}{m}\sum_{k=1}^{m}(1-\tfrac{t}{k}M)^{-k} = e^{Mt}.$$

Since $g \in \mathrm{Lip}(\mathtt{A}^\Gamma)$ was arbitrary, the inequality of transportation distance follows.

CHAPTER 4

Free energy, Gibbs measures, and Glauber dynamics for nearest-neighbor interactions on trees

We extend results of R. Holley beyond the integer lattice to a large class of groups which includes free groups. In particular we show that a shift-invariant measure is Gibbs if and only if it is Glauber invariant. Moreover, any shift-invariant measure converges weakly to the set of Gibbs measures when evolved under Glauber dynamics. These results are proven using the notion of free energy density relative to a sofic approximation by homomorphisms, introduced in Chapter 3. We also show that any shift-invariant measure which minimizes free energy density is Gibbs.

4.1 Introduction, main results

In the present chapter we focus on extending results of Holley in [Hol71] to the nonamenable setting. He studied a natural notion of free energy density for systems with sites indexed by $\mathbb{Z}^r$, and used it to relate Gibbs measures and Glauber dynamics. His approach does not seem to work for nonamenable groups due to non-negligibility of the boundary of large finite subsystems.

More specifically, his approach can be viewed as dependent on the following approximate equivariance:

$$\begin{array}{ccc} \text{state on } \mathbb{Z}^r & \xrightarrow{\text{Glauber dynamics}} & \text{state on } \mathbb{Z}^r \\ \Big\downarrow{\scriptstyle\text{marginal}} & \approx\circlearrowright & \Big\downarrow{\scriptstyle\text{marginal}} \\ \text{state on large rectangle} & \xrightarrow{\text{Glauber dynamics}} & \text{state on large rectangle} \end{array}$$

We can imagine that this works because, in a large rectangle, most vertices are far from the boundary. Therefore the dynamics in the rectangle can, to a good approximation, be treated as isolated from the exterior of the rectangle. The composition $\begin{smallmatrix}\cdot & \to & \cdot\\ & & \downarrow\\ \cdot & & \cdot\end{smallmatrix}$ corresponds to running the infinitary dynamics with influence from outside the rectangle while the composition $\begin{smallmatrix}\cdot & & \cdot\\ \downarrow & & \\ \cdot & \to & \cdot\end{smallmatrix}$ corresponds to first isolating the rectangle then running the finitary dynamics.

There is no stated result in Holley's paper which directly corresponds to this phenomenon, but it can be compared to Theorem 3.2.1 above.

To work with nonamenable groups we will instead use an "extrinsic" approach to free energy density which is inspired by recent work on the entropy theory of nonamenable group actions, initiated by Lewis Bowen [Bow10c] to solve similar problems which appear in that area.

4.1.1 Related work

In one respect, Holley [Hol71] worked in slightly more generality than we do here: he considered finite-range interactions, not just nearest-neighbor interactions. Higuchi and Shida [HS75] extended his results to spin systems on $\mathbb{Z}^r$ which may have infinite-range interactions, but the strength of the interactions is assumed to decay sufficiently quickly.

The method of the present thesis may be compatible with such generalizations, but for the sake of simplicity we choose not to pursue them here.

More recently, Jahnel and Külske [JK19] have extended the free energy density approach to non-reversible dynamics on integer-lattice systems.

Caputo and Martinelli [CM06] have shown that if we evolve the product of plus-biased

Bernoulli measures by Ising Glauber dynamics on an infinite tree, then it converges weakly to the "plus boundary conditions" Gibbs measure.

There has been some other work on notions of free energy density for Ising models on nonamenable groups, but these notions do not appear to have the properties we want for our present purposes. Dembo and Montanari [DM10] consider, as we do below, a sequence of finite graphs that locally converge to an infinite tree. Their work differs from ours in that they study the limiting free energy density of the (unique) Gibbs measures on these finite graphs, while we study the free energy density of finitary measures which are locally consistent with a chosen infinitary measure (which is not necessarily Gibbs).

4.1.2 Precise statements of basic definitions and main theorems

The correspondence between finite and infinite systems is established using *empirical distributions*, which we recall here. Let $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$ and $\mathbf{x} \in \mathtt{A}^V$. For any $v \in V$ there is a natural way to lift $\mathbf{x}$ to a labeling $\Pi_v^\sigma \mathbf{x} \in \mathtt{A}^\Gamma$, starting by lifting $\mathbf{x}_v$ to the root e. More precisely,

$$\big(\Pi_v^\sigma \mathbf{x}\big)(\gamma) = \mathbf{x}\big(\sigma^\gamma(v)\big).$$

The empirical distribution of $\mathbf{x}$ is defined by

$$P_{\mathbf{x}}^\sigma = (v \mapsto \Pi_v^\sigma \mathbf{x})_* \,\mathrm{Unif}(V) = \frac{1}{|V|}\sum_{v \in V} \delta_{\Pi_v^\sigma \mathbf{x}} \in \mathrm{Prob}\big(\mathtt{A}^\Gamma\big).$$

This captures the 'local statistics' of $\mathbf{x}$. This notation was used in the approach to sofic entropy in [Aus16].

To state our results we use the following way of measuring local similarity of a finite graph σ to Γ: for $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$, we say $\mathrm{B}^\sigma(v,R) \cong \mathrm{B}^\Gamma(e,R)$ if there is an isomorphism of the induced subgraphs $\mathrm{B}^\sigma(v,R), \mathrm{B}^\Gamma(e,R)$ which respects both edge labels and directions.

Define

$$\delta_R^\sigma = \frac{1}{|V|}|\{v \in V \ : \ \mathrm{B}^\sigma(v,R) \ncong \mathrm{B}^\Gamma(e,R)\}|$$

$$\Delta^\sigma = \inf_R \big(9 \cdot (2/3)^R + 6\delta_R^\sigma\big).$$

The constants which appear here are connected to the choice of metric d in Chapter 3 above. If Δ^σ is small, then σ looks like Γ to a large radius near most vertices. This is a way of saying that the action of σ is approximately free. Note also that a sequence σ_n Benjamini-Schramm converges to the infinite tree Γ (or equivalently is a sofic approximation to the group Γ) if and only if $\Delta^{\sigma_n} \to 0$.

As mentioned above, our central tool is a notion of "free energy density" of a measure $\mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$ with respect to a nearest-neighbor potential. This free energy density is defined relative to a choice of $\Sigma = (\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))_{n\in\mathbb{N}}$ with $\Delta^{\sigma_n} \to 0$. It may be $+\infty$, but if it is finite then it is nonincreasing as μ evolves according to Glauber dynamics (Proposition 4.2.2). Moreover if μ is not Gibbs then it is strictly decreasing; Proposition 4.2.3 gives a stronger version of this claim. For every choice of Σ there exist measures with finite free energy density, so this implies the following:

Theorem A. *For any choice of Σ and any nearest-neighbor interaction, every $\mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$ minimizing a_Σ is Gibbs for that interaction, unless a_Σ is identically $+\infty$.*

The converse is false, since a Gibbs measure may have free energy density $+\infty$ with respect to some Σ. It is unclear whether a Gibbs measure may have finite but non-minimal free energy density.

Theorem B. *Suppose $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$, and let μ_t denote its evolution under Glauber dynamics. If there exist $s \geq 0$ and Σ such that $\mathrm{a}_\Sigma(\mu_s) < +\infty$, then μ_t converges weakly to the set of Gibbs measures as $t \to \infty$.*

It is possible to avoid the degenerate case of infinite free energy density by an appropriate choice of Σ when Γ has a property called "property PA"; see Section 3.5 for a definition and the relevant result (Proposition 3.5.1). Hence we have the following:

Corollary 4.1.1. *If Γ has property* PA*, then a shift-invariant measure is Gibbs if and only if it is Glauber-invariant.*

4.2 Proof of Theorem B

Proposition 4.2.1. *Let $\zeta_0 \in \mathrm{Prob}(\mathtt{A}^V)$, and let ζ_t denote its evolution under Glauber dynamics. Then for all $t > 0$, ζ_t has full support and*

$$\frac{d}{dt}A(\zeta_t) = \sum_{\mathbf{x},v,\mathtt{a}} F_0\left(\frac{\exp\{-\Phi_v(\mathbf{x})\}}{\exp\{-\Phi_v(\mathbf{x}^{v\to\mathtt{a}})\}}\frac{\zeta_t\{\mathbf{x}^{v\to\mathtt{a}}\}}{\zeta_t\{\mathbf{x}\}}\right)\zeta_t\{\mathbf{x}\}c_v(\mathbf{x},\mathtt{a}).$$

Our proof of this proposition is based on the proof of the analogous result in Holley's paper [Hol71], with some minor changes.

For $\mathbf{x} \in \mathtt{A}^V$, write

$$P(\mathbf{x}) = \exp\{-U(\mathbf{x})\}.$$

This is just the Gibbs measure on V, except without the normalizing factor. It is easy to see that

$$A(\zeta) = \sum_{\mathbf{x}} \zeta\{\mathbf{x}\}\log\frac{\zeta\{\mathbf{x}\}}{P(\mathbf{x})}.$$

Proof of proposition. A calculation using the Markov generator shows that

$$\frac{d}{dt}\zeta_t\{\mathbf{x}\} = \sum_{v,\mathtt{a}}\left[\zeta_t\{\mathbf{x}^{v\to\mathtt{a}}\}c_v(\mathbf{x}^{v\to\mathtt{a}},\mathbf{x}(v)) - \zeta_t\{\mathbf{x}\}c_v(\mathbf{x},\mathtt{a})\right].$$

In particular, if $\mathbf{x}$ is such that $\zeta_t\{\mathbf{x}\} = 0$ but there exist $v,\mathtt{a}$ such that $\zeta_t\{\mathbf{x}^{v\to\mathtt{a}}\} > 0$, then $\frac{d}{dt}\zeta_t\{\mathbf{x}\} > 0$. Unless $t = 0$ this would imply the existence of times where ζ_s gives negative mass to $\{x\}$. Therefore ζ_t has full support for all $t > 0$.

Therefore

$$\begin{aligned}\frac{d}{dt}A(\zeta_t) &= \sum_{\mathbf{x}\in\mathtt{A}^V} \frac{d}{dt}\left[\zeta_t\{\mathbf{x}\}\log\frac{\zeta_t\{\mathbf{x}\}}{P(\mathbf{x})}\right] \\ &= \sum_{\mathbf{x}\in\mathtt{A}^V} \frac{d}{dt}\left[\zeta_t\{\mathbf{x}\}\right]\log\frac{\zeta_t\{\mathbf{x}\}}{P(\mathbf{x})} \\ &= \sum_{\mathbf{x}\in\mathtt{A}^V}\sum_{v,\mathtt{a}} \left[\zeta_t\{\mathbf{x}^{v\to\mathtt{a}}\}c_v(\mathbf{x}^{v\to\mathtt{a}},\mathbf{x}(v)) - \zeta_t\{\mathbf{x}\}c_v(\mathbf{x},\mathtt{a})\right]\log\frac{\zeta_t\{\mathbf{x}\}}{P(\mathbf{x})}.\end{aligned}$$

For $\mathbf{x},\mathbf{y}\in\mathtt{A}^V$, define

$$\mathfrak{A}(\mathbf{x},\mathbf{y}) = \begin{cases} -\sum\limits_{\substack{\mathtt{a},v \\ \mathtt{a}\neq\mathbf{x}(v)}} c_v(\mathbf{y},\mathtt{a}), & \mathbf{x}=\mathbf{y} \\ c_v(\mathbf{y},\mathtt{a}), & \mathbf{x}\neq\mathbf{y},\ \mathbf{x}=\mathbf{y}^{v\to\mathtt{a}} \text{ for some } v\in V, \mathtt{a}\in\mathtt{A} \\ 0, & \text{else.}\end{cases}$$

This has the following useful properties:

$$\sum_{\mathbf{x}}\mathfrak{A}(\mathbf{x},\mathbf{y}) = 0 \quad \forall\mathbf{y}. \tag{4.1}$$

Proof. For any $\mathbf{y}$,

$$\sum_{\mathbf{x}}\mathfrak{A}(\mathbf{x},\mathbf{y}) = \mathfrak{A}(\mathbf{y},\mathbf{y}) + \sum_{v\in V}\sum_{\mathtt{a}\neq\mathbf{y}(v)}\mathfrak{A}(\mathbf{y}^{v\to\mathtt{a}},\mathbf{y}) + 0 = 0. \qquad \triangleleft$$

$$\sum_{\mathbf{y}}\mathfrak{A}(\mathbf{x},\mathbf{y})P(\mathbf{y}) = 0 \quad \forall\mathbf{x}. \tag{4.2}$$

Proof. For any $\mathbf{x}$,

$$\begin{aligned}\sum_{\mathbf{y}}\mathfrak{A}(\mathbf{x},\mathbf{y})P(\mathbf{y}) &= \mathfrak{A}(\mathbf{x},\mathbf{x})P(\mathbf{x}) + \sum_{v\in V}\sum_{\mathtt{a}\neq\mathbf{x}(v)}\mathfrak{A}(\mathbf{x},\mathbf{x}^{v\to\mathtt{a}})P(\mathbf{x}^{v\to\mathtt{a}}) + 0 = 0 \\ &= \sum_{\substack{\mathtt{a},v \\ \mathtt{a}\neq\mathbf{x}(v)}}\left[-c_v(\mathbf{x},\mathtt{a})\exp\{-U(\mathbf{x})\} + c_v(\mathbf{x}^{v\to\mathtt{a}},\mathbf{x}(v))\exp\{-U(\mathbf{x}^{v\to\mathtt{a}})\}\right] \\ &= 0.\end{aligned}$$

In fact every term of this last sum is zero because

$$\frac{c_v(\mathbf{x},\mathtt{a})\exp\{-U(\mathbf{x})\}}{c_v(\mathbf{x}^{v\to\mathtt{a}},\mathbf{x}(v))\exp\{-U(\mathbf{x}^{v\to\mathtt{a}})\}} = \frac{\exp\{-\Phi_v(\mathbf{x}^{v\to\mathtt{a}})\}\exp\{-U(\mathbf{x})\}}{\exp\{-\Phi_v(\mathbf{x})\}\exp\{-U(\mathbf{x}^{v\to\mathtt{a}})\}} = 1. \qquad \triangleleft$$

Using these two properties of $\mathfrak{A}$, and the fact that ζ_t has full support, we see that

$$\begin{aligned}\frac{d}{dt}A(\zeta_t) &= \sum_{\mathbf{x},\mathbf{y}} \mathfrak{A}(\mathbf{x},\mathbf{y})\zeta_t\{\mathbf{y}\}\log\frac{\zeta_t\{\mathbf{x}\}}{P(\mathbf{x})} \\ &= \sum_{\mathbf{x},v,\mathtt{a}} F_0\left(\frac{P(\mathbf{x})}{P(\mathbf{x}^{v\to\mathtt{a}})}\frac{\zeta_t\{\mathbf{x}^{v\to\mathtt{a}}\}}{\zeta_t\{\mathbf{x}\}}\right)\zeta_t\{\mathbf{x}\}c_v(\mathbf{x},\mathtt{a}).\end{aligned}$$

The desired formula follows from the fact that $\frac{P(\mathbf{x})}{P(\mathbf{x}^{v\to\mathtt{a}})} = \frac{\exp\{-\Phi_v(\mathbf{x})\}}{\exp\{-\Phi_v(\mathbf{x}^{v\to\mathtt{a}})\}}$, also used above. $\square$

We first use the previous result to show that free energy density is nonincreasing.

Proposition 4.2.2. *Suppose* $\mu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$*, and let* $\Sigma = \big(\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))\big)_{n\in\mathbb{N}}$ *such that* $\Delta^{\sigma_n} \to 0$*.*

Then $\mathrm{a}_\Sigma(\mu_0) \geq \mathrm{a}_\Sigma(\mu_t)$ *for all* $t \geq 0$*.*

Proof. If $\mathrm{a}_\Sigma(\mu_0) = +\infty$ then the result is trivial, so suppose $\mathrm{a}_\Sigma(\mu_0) < +\infty$. This means that for any $\mathcal{O} \ni \mu_0$ we have $\mathbf{\Omega}(\sigma_n, \mathcal{O}) \neq \varnothing$ for all large enough n.

Let $\mathcal{U}_t$ be an arbitrary weak-open neighborhood of μ_t. By Proposition 3.4.1 there exists $\mathcal{U}_0 \ni \mu_0$ such that, for all large enough n, we have $\zeta_t \in \mathbf{\Omega}(\sigma_n, \mathcal{U}_t)$ whenever $\zeta_0 \in \mathbf{\Omega}(\sigma_n, \mathcal{U}_0)$.

Suppose n is large enough that $\mathbf{\Omega}(\sigma_n, \mathcal{U}_0) \neq \varnothing$. Since $F_0 \leq 0$, the previous proposition implies that for any $\zeta_0 \in \mathrm{Prob}(\mathtt{A}^{V_n})$ and any $t > 0$

$$\frac{d}{dt}A(\zeta_t) \leq 0.$$

Therefore for any $t \geq 0$

$$A(\zeta_0) \geq A(\zeta_t),$$

and hence

$$\inf_{\zeta\in\mathbf{\Omega}(\sigma_n,\mathcal{U}_0)} A(\zeta) \geq \inf_{\zeta\in\mathbf{\Omega}(\sigma_n,\mathcal{U}_t)} A(\zeta).$$

Now by definition of a_Σ we have

$$\begin{aligned}\mathrm{a}_\Sigma(\mu_0) &= \sup_{\mathcal{O}\ni\mu_0} \limsup_{n\to\infty} \frac{1}{|V_n|} \inf_{\zeta\in\mathbf{\Omega}(\sigma_n,\mathcal{O})} A(\zeta)\\ &\geq \limsup_{n\to\infty} \frac{1}{|V_n|} \inf_{\zeta\in\mathbf{\Omega}(\sigma_n,\mathcal{U}_0)} A(\zeta)\\ &\geq \limsup_{n\to\infty} \frac{1}{|V_n|} \inf_{\zeta\in\mathbf{\Omega}(\sigma_n,\mathcal{U}_t)} A(\zeta).\end{aligned}$$

Taking the supremum over $\mathcal{U}_t$ gives the result. □

By a more careful analysis we can get the following proposition, the second part of which may be interpreted as semicontinuity of the time derivative of $\mathrm{a}_\Sigma(\mu_t)$ as a function of the measure:

Proposition 4.2.3. *Suppose $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ is not Gibbs. Then there exist $c, T > 0$ and an open neighborhood $\mathcal{O} \ni \mu$ such that if $\mu_0 \in \mathcal{O}$ then*

1. *There exists $\delta > 0$ such that for any $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$ with $\Delta^\sigma < \delta$ and any $\zeta_0 \in \mathbf{\Omega}(\sigma, \mathcal{O})$ we have $A(\zeta_0) \geq A(\zeta_t) + ct|V|$ for all $t \in [0, T]$.*
2. *If $\Sigma = \big(\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))\big)_{n\in\mathbb{N}}$ is such that $\Delta^{\sigma_n} \to 0$ then $\mathrm{a}_\Sigma(\mu_0) \geq \mathrm{a}_\Sigma(\mu_t) + ct$ for all $t \in [0, T]$.*

Here we take the convention that $+\infty + ct = +\infty$.

Proof. Since μ is not Gibbs, there exists R such that either μ_R does not have full support or $\Delta^R_{\mathtt{a}}(\mu_R) < 0$ for some $\mathtt{a} \in \mathtt{A}$. We will come back to these two cases in a moment, but for now let R be fixed so that one of them occurs. We may assume $R \geq 1$.

Fix $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$. Let

$$s = \min\left\{ \frac{\exp\{-\Phi_e(\mathbf{x})\}}{\sum_{\mathtt{b}\in\mathtt{A}} \exp\{-\Phi_e(\mathbf{x}^{e\to\mathtt{b}})\}} \;:\; \mathbf{x} \in \mathtt{A}^{\mathrm{B}(e,1)} \right\} > 0,$$

and call $v \in V$ *good* if $\mathrm{B}^\sigma(v, R) \cong \mathrm{B}^\Gamma(e, R)$, and let V' be the set of such v. Then

$$\frac{d}{dt} A(\zeta_t) \leq s \sum_{\substack{\mathbf{x},v,\mathtt{a} \\ v \text{ good}}} F_0\left(\frac{\exp\{-\Phi_v(\mathbf{x})\}}{\exp\{-\Phi_v(\mathbf{x}^{v\to\mathtt{a}})\}} \frac{\zeta_t\{\mathbf{x}^{v\to\mathtt{a}}\}}{\zeta_t\{\mathbf{x}\}} \right) \zeta_t\{\mathbf{x}\}.$$

Let $\widetilde{P^{\sigma,R}_\zeta} \in \mathrm{Prob}(\mathtt{A}^{\mathrm{B}^\Gamma(e,R)})$ be given by

$$\widetilde{P^{\sigma,R}_\zeta}\{\mathbf{y}\} = \frac{1}{|V'|} \sum_{\substack{\mathbf{x},v \\ v \text{ good} \\ \mathbf{x}\restriction_{\mathrm{B}(v,R)}=\mathbf{y}}} \zeta\{\mathbf{x}\}$$

Note that $\widetilde{P^{\sigma,R}_\zeta}$ is close to the $\mathrm{B}^\Gamma(e,R)$-marginal of P^σ_ζ in total variation distance if most vertices are good. Then, applying Jensen's inequality,

$$\begin{aligned}\frac{d}{dt}A(\zeta_t) &\leq s\sum_{\mathtt{a}\in\mathtt{A}} \sum_{\mathbf{y}\in\mathtt{A}^{\mathrm{B}^\Gamma(e,R)}} \sum_{\substack{\mathbf{x},v \\ v \text{ good} \\ \mathbf{x}\restriction_{\mathrm{B}(v,R)}=\mathbf{y}}} F_0\left(\frac{\exp\{-\Phi_v(\mathbf{x})\}}{\exp\{-\Phi_v(\mathbf{x}^{v\to\mathtt{a}})\}}\frac{\zeta_t\{\mathbf{x}^{v\to\mathtt{a}}\}}{\zeta_t\{\mathbf{x}\}}\right)\zeta_t\{\mathbf{x}\} \\ &\leq s|V'|\sum_{\mathtt{a}\in\mathtt{A}} \sum_{\mathbf{y}\in\mathtt{A}^{\mathrm{B}^\Gamma(e,R)}} \widetilde{P^{\sigma,R}_{\zeta_t}}\{\mathbf{y}\}\cdot F_0\left(\frac{\exp\{-\Phi_e(\mathbf{y})\}}{\exp\{-\Phi_e(\mathbf{y}^{v\to\mathtt{a}})\}}\frac{\widetilde{P^{\sigma,R}_{\zeta_t}}\{\mathbf{y}^{e\to\mathtt{a}}\}}{\widetilde{P^{\sigma,R}_{\zeta_t}}\{\mathbf{y}\}}\right).\end{aligned}$$

We now consider the two cases mentioned above. First, if μ_R does not have full support, there exist $\mathbf{y} \in \mathtt{A}^{\mathrm{B}(e,R)}$, $v \in \mathrm{B}(e,R)$ and $\mathtt{a} \in \mathtt{A}$ such that $\mu_R\{\mathbf{y}\} \neq 0$ but $\mu_R\{\mathbf{y}^{v\to\mathtt{a}}\} = 0$. Using translation-invariance of μ, we may assume $v = e$. Then

$$\mu_R\{\mathbf{y}\}\cdot F_0\left(\frac{\exp\{-\Phi_e(\mathbf{y})\}}{\exp\{-\Phi_e(\mathbf{y}^{e\to\mathtt{a}})\}}\frac{\mu_R\{\mathbf{y}^{e\to\mathtt{a}}\}}{\mu_R\{\mathbf{y}\}}\right) = -\mu_R\{\mathbf{y}\} < 0.$$

By continuity of F_0, there exists $\varepsilon > 0$ such that

$$|a - \mu_R\{\mathbf{y}\}| < \varepsilon \text{ and } 0 \leq b < \varepsilon \;\Rightarrow\; a\cdot F_0\left(\frac{\exp\{-\Phi_e(\mathbf{y})\}}{\exp\{-\Phi_e(\mathbf{y}^{e\to\mathtt{a}})\}}\frac{b}{a}\right) < -\frac{\mu_R\{\mathbf{y}\}}{2}.$$

In particular, if $\|\widetilde{P^{\sigma,R}_{\zeta_t}} - \mu_R\|_{\mathrm{TV}} < \varepsilon$ then

$$\frac{d}{dt}A(\zeta_t) < -s|V'|\frac{\mu_R\{\mathbf{y}\}}{2}.$$

In this case we will take $c = s\,\mu_R\{\mathbf{y}\}/4$.

Now consider the other case, in which $\Delta^R_{\mathtt{a}}(\mu_R) < 0$ for some $\mathtt{a}$. By definition of $\Delta^R_{\mathtt{a}}(\mu_R)$, we can pick $\mathbf{y}$ such that

$$C := \mu_R\{\mathbf{y}\}\cdot F_0\left(\frac{\exp\{-\Phi_e(\mathbf{y})\}}{\exp\{-\Phi_e(\mathbf{y}^{e\to\mathtt{a}})\}}\frac{\mu_R\{\mathbf{y}^{e\to\mathtt{a}}\}}{\mu_R\{\mathbf{y}\}}\right) < 0.$$

and proceed in the same way, picking $\varepsilon > 0$ such that

$$|a - \mu_R\{\mathbf{y}\}| < \varepsilon \text{ and } |b - \mu_R\{\mathbf{y}^{e\to\mathtt{a}}\}| < \varepsilon \;\Rightarrow\; a \cdot F_0\left(\frac{\exp\{-\Phi_e(\mathbf{y})\}}{\exp\{-\Phi_e(\mathbf{y}^{e\to\mathtt{a}})\}}\frac{b}{a}\right) < -\frac{C}{2}.$$

In particular, if $\|\widetilde{P^{\sigma,R}_{\zeta_t}} - \mu_R\|_{\mathrm{TV}} < \varepsilon$ then

$$\frac{d}{dt}A(\zeta_t) < -s|V'|\frac{C}{2}.$$

In this case we will take $c = s\,C/4$.

In either case, we now have chosen $\varepsilon > 0$ such that if $\|\widetilde{P^{\sigma,R}_{\zeta_t}} - \mu_R\|_{\mathrm{TV}} < \varepsilon$ then

$$\frac{d}{dt}A(\zeta_t) < -2|V'|c.$$

Let $\mathcal{O}_1 = \{\nu \in \operatorname{Prob}(\mathtt{A}^\Gamma) \,:\, \|\nu_R - \mu_R\|_{\mathrm{TV}} < \varepsilon\}$. By continuity of $(\mu_0, t) \mapsto \mu_t$, we can pick $\mathcal{O}, T$ such that if $\mu_0 \in \mathcal{O}$ then $\mu_t \in \mathcal{O}_1$ for all $t \in [0, T]$.

Fix $\mu_0 \in \mathcal{O}$. By Proposition 3.4.1, for any $\eta > 0$ there exists $\mathcal{U} \ni \mu_0$ and $\delta > 0$ such that if $\zeta_0 \in \mathbf{\Omega}(\sigma_n, \mathcal{U})$ and $\Delta^\sigma < \delta$ then $\bar{d}(P^\sigma_{\zeta_t}, \mu_t) < \eta$ for all $t \in [0, T]$. If we pick η small enough, this implies $\zeta_t \in \mathbf{\Omega}(\sigma, \mathcal{O}_1)$ for all $t \in [0, T]$. Hence if $\zeta_0 \in \mathbf{\Omega}(\sigma_n, \mathcal{U})$ and $\Delta^\sigma < \delta$ then for any $t \in [0, T]$

$$A(\zeta_t) - A(\zeta_0) \leq -2|V'|ct.$$

Now by definition we have

$$|V'| = (1 - \delta^\sigma_R)|V|.$$

If δ is small enough then $\Delta^\sigma < \delta$ implies $\delta^\sigma_R < 1/2$, so that $2|V'| > |V|$. This completes the proof of the first part.

For the second part, for all large enough n we have $\Delta^{\sigma_n} < \delta$ and therefore

$$\begin{aligned} A(\zeta_0) &\geq A(\zeta_t) + |V_n|ct \\ &\geq \inf_{\zeta \in \mathbf{\Omega}(\sigma_n, \mathrm{B}(\mu_t, \eta))} A(\zeta) + |V_n|ct \end{aligned}$$

for all $t \in [0, T]$. Then, since the first limit in the definition of a_Σ is a supremum

$$\begin{aligned} \mathrm{a}_\Sigma(\mu_0) &\geq \limsup_{n\to\infty} \inf_{\zeta_0 \in \mathbf{\Omega}(\sigma_n, \mathcal{U})} \frac{1}{|V_n|} A(\zeta_0) \\ &\geq \limsup_{n\to\infty} \inf_{\zeta \in \mathbf{\Omega}(\sigma_n, \mathrm{B}(\mu_t, \eta))} \frac{1}{|V_n|} A(\zeta) + ct. \end{aligned}$$

Taking η to zero gives the result. □

Proof of Theorem. Suppose for the sake of contradiction that μ_t does not converge to the set of Gibbs measures; then we can pick some $\nu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ which is a limit point of $\{\mu_t : t \geq 0\}$ but not a Gibbs measure.

By Proposition 4.2.3, we can pick $\delta, T > 0$ and an open neighborhood $\mathcal{O} \ni \nu$ such that for every t with $\mu_t \in \mathcal{O}$ we have $\mathrm{a}_\Sigma(\mu_t) \geq \mathrm{a}_\Sigma(\mu_{t+T}) + \delta T$.

On the other hand, under the assumption that $\mathrm{a}_\Sigma(\mu_s) < +\infty$, the set $\{\mathrm{a}_\Sigma(\mu_t) : t \geq s\}$ is bounded: an upper bound is $\mathrm{a}_\Sigma(\mu_s)$ by Proposition 4.2.2, and a lower bound is $u^{\min} - \log|\mathtt{A}|$. This is a contradiction, so the theorem follows. □

4.3 Connection to property PA

We first prove the following:

Proposition 4.3.1. *A group Γ has property* PA *if and only if for any $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ there exists a sequence $\Sigma = \big(\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))\big)_{n\in\mathbb{N}}$ and a sequence $(\mathbf{x}_n \in \mathtt{A}^{V_n})_{n\in\mathbb{N}}$ with $P^{\sigma_n}_{\mathbf{x}_n} \xrightarrow{wk} \mu$ and $\Delta^{\sigma_n} \to 0$.*

Some ideas for this proof were shared with me by Lewis Bowen.

Proof. The 'if' direction is clear, since each $P^{\sigma_n}_{\mathbf{x}_n}$ is periodic.

For the other direction, suppose Γ has property PA and let $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$. By definition of property PA, we can pick a sequence of periodic measures $(\mu_n)_{n\in\mathbb{N}}$ converging to μ.

Fix $n \in \mathbb{N}$. The support of μ_n consists of finitely many orbits under Γ; let $\{\mathbf{y}_1, \ldots, \mathbf{y}_k\} \subset \mathtt{A}^\Gamma$ be a set which contains exactly one element of each orbit, and denote the (finite) orbits by $\Gamma\mathbf{y}_i$. Then we can write

$$\mu_n = \sum_{i=1}^k a_i \operatorname{Unif}(\Gamma\mathbf{y}_i).$$

Pick natural numbers $m_1, \ldots, m_k$, and let V_n be the disjoint union of m_i copies of $\Gamma\mathbf{y}_i$ for each i. Let $\sigma_n \in \operatorname{Hom}(G, \operatorname{Sym}(V_n))$ act separately on each copy of each orbit. Let $\mathbf{x}_n \in \mathtt{A}^{V_n}$ be given by

$$\mathbf{x}_n(v) = v(e).$$

Then

$$P_{\mathbf{x}_n}^{\sigma_n} = \sum_{i=1}^k \frac{m_i}{\sum_{j=1}^k m_j} \operatorname{Unif}(\Gamma\mathbf{y}_i),$$

so if $m_1, \ldots, m_k$ are chosen appropriately then we can ensure $P_{\mathbf{x}_n}^{\sigma_n} \to \mu$.

Now we need to show that we can ensure $\Delta^{\sigma_n} \to 0$. Let $\nu \in \operatorname{Prob}(\{0,1\}^\Gamma)$ be the product measure with uniform base. Then then above argument implies the existence of sequences $\big(\sigma_n \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V_n))\big)_{n\in\mathbb{N}}$ and $\mathbf{z}_n \in (\mathtt{A} \times \{0,1\})^{V_n}$ with $P_{\mathbf{z}_n}^{\sigma_n} \to \mu \times \nu$. If we write $\mathbf{z}_n = (\mathbf{x}_n, \mathbf{y}_n)$ with $\mathbf{x}_n \in \mathtt{A}^{V_n}$ and $\mathbf{y}_n \in \{0,1\}^{V_n}$, then $P_{\mathbf{x}_n}^{\sigma_n} \to \mu$ and $P_{\mathbf{y}_n}^{\sigma_n} \to \nu$. We will show that the latter fact implies $\Delta^{\sigma_n} \to 0$. Suppose $v \in V_n$, $\gamma \in \Gamma$ are such that $\sigma_n^\gamma v = v$. Then for any $\beta \in \Gamma$,

$$\big(\Pi_v^{\sigma_n}\mathbf{y}_n\big)(\beta\gamma) = \mathbf{y}_n\big(\sigma_n^{\beta\gamma}v\big) = \mathbf{y}_n\big(\sigma_n^{\beta}v\big) = \big(\Pi_v^{\sigma_n}\mathbf{y}_n\big)(\beta).$$

In particular, for any finite set $D \subset \Gamma$ we have

$$P_{\mathbf{y}_n}^{\sigma_n}\{\mathbf{w} \in \{0,1\}^\Gamma \,:\, \mathbf{w}(\beta\gamma) = \mathbf{w}(\beta)\ \forall \beta \in D\} \geq \frac{1}{|V_n|}|\{v \in V_n \,:\, \sigma_n^\gamma v = v\}|.$$

But by assumption, as $n \to \infty$ the left-hand side converges to

$$\begin{aligned}\nu\{\mathbf{w} \in \{0,1\}^\Gamma \,:\, \mathbf{w}(\beta\gamma) = \mathbf{w}(\beta)\ \forall \beta \in D\} &\leq \nu\{\mathbf{w} \,:\, \mathbf{w}(\beta\gamma) = \mathbf{w}(\beta)\ \forall \beta \in D \text{ s.t. } \beta\gamma \notin D\} \\ &= 2^{-|D\gamma \setminus D|},\end{aligned}$$

and hence

$$\limsup_{n\to\infty} \frac{1}{|V_n|}|\{v \in V_n \,:\, \sigma_n^\gamma v = v\}| \leq 2^{-|D\gamma \setminus D|}.$$

As long as $\gamma \neq e$, the set D can be chosen to make $|D\gamma \setminus D|$ arbitrarily large, so that

$$\lim_{n\to\infty} \frac{1}{|V_n|}|\{v \in V_n \,:\, \sigma_n^\gamma v = v\}| = 0.$$

For any $R \in \mathbb{N}$, it can be checked that if $\sigma_n^\gamma v \neq v$ for all $\gamma \neq e$ such that $|\gamma| \leq 2R+1$ then the map

$$\mathrm{B}^\Gamma(e,R) \to \mathrm{B}^\sigma(v,R)$$
$$\gamma \mapsto \sigma^\gamma v$$

is an isomorphism of the (labeled, directed) induced subgraphs. Therefore

$$\delta_R^{\sigma_n} \leq \sum_{\gamma \in \mathrm{B}^\Gamma(e,2R+1)\setminus\{e\}} \frac{1}{|V_n|}|\{v \in V_n \,:\, \sigma_n^\gamma v = v\}| \to 0,$$

which, since R is arbitrary, implies $\Delta^{\sigma_n} \to 0$. □

4.3.1 Proof of Proposition 3.5.1

Note $\mathrm{a}_\Sigma(\mu) = +\infty$ if and only if there exists an open neighborhood $\mathcal{O} \ni \mu$ such that $\mathbf{\Omega}(\sigma_n, \mathcal{O})$ is empty for infinitely many n. Therefore if $\mathrm{a}_\Sigma(\mu) < +\infty$ there exists a sequence $\zeta_n \in \mathrm{Prob}(\mathtt{A}^{V_n})$ with $P_{\zeta_n}^{\sigma_n} \to \mu$. Since each $P_{\zeta_n}^{\sigma_n}$ is periodic, this shows that if for every μ there exists Σ with $\mathrm{a}_\Sigma(\mu) < +\infty$ then Γ has property PA.

Conversely, if Γ has property PA and $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ is given, then by the above proposition we can pick Σ and $(\mathbf{x}_n)_{n\in\mathbb{N}}$ with $\Delta^{\sigma_n} \to 0$ and $P_{\mathbf{x}_n}^{\sigma_n} \to \mu$. But then for any open $\mathcal{O} \ni \mu$ we have $\delta_{\mathbf{x}_n} \in \mathbf{\Omega}(\sigma_n, \mathcal{O})$ for all large enough n, so $\mathrm{a}_\Sigma(\mu) < +\infty$.

CHAPTER 5

Metastability and maximal-entropy joinings of Gibbs measures on finitely-generated groups

We prove a metastability result for finitary microstates which are good models for a Gibbs measure for a nearest-neighbor interaction on a finitely-generated group. This is used to show that any maximal-entropy joining of two such Gibbs states is a relative product over the tail σ-algebra, except in degenerate cases.

We also use results on extremal cuts of random graphs to further investigate optimal self-joinings of the Ising model on a free group.

5.1 Introduction, main results

As above, let Γ be a countably infinite group with r generators $s_1, \dots, s_r$, and let $\mathtt{A}$ be a finite set. We will also use Γ to denote the left Cayley graph of the group, which has vertex set Γ and an s_i-labeled directed edge $(\gamma, s_i\gamma)$ for each $i \in [r] = \{1, 2, \dots, r\}$.

The group Γ acts on itself by right multiplication; note that this action consists of isomorphisms of the Cayley graph which preserve edge labels and directions. We also let Γ act on the set of labelings $\mathtt{A}^\Gamma$: given $\mathbf{x} \in \mathtt{A}^\Gamma$ and $\beta \in \Gamma$, the shifted labeling $\beta\mathbf{x}$ is given by

$$\big(\beta\mathbf{x}\big)(\gamma) = \mathbf{x}(\gamma\beta).$$

This also induces an action on $\operatorname{Prob}(\mathtt{A}^\Gamma)$ by pushforwards. A probability measure invariant under this action will be called shift-invariant; the set of such measures will be denoted $\operatorname{Prob}^\Gamma(\mathtt{A}^\Gamma)$.

We will think of a measure $\mu \in \operatorname{Prob}^\Gamma(\mathtt{A}^\Gamma)$ as specifying local statistics of finite systems according to the following paradigm:

Given a finite set V and a homomorphism $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$, we can construct a multigraph with an s_i-labeled directed edge $(v, \sigma^{s_i} v)$ for each $v \in V$ and $i \in [r]$; this will be called the graph of σ.

If $\mathbf{x} \in \mathtt{A}^V$ is any labeling of V by elements of $\mathtt{A}$, we can pull back $\mathbf{x}$ to a labeling of Γ. This is called a pullback name of $\mathbf{x}$, and is denoted

$$\Pi_v^\sigma \mathbf{x} \coloneqq \big(\mathbf{x}(\sigma^\gamma v)\big)_{\gamma \in \Gamma} \in \mathtt{A}^\Gamma.$$

The *empirical distribution* of $\mathbf{x}$ over σ is the distribution of these pullback names if the basepoint v is chosen uniformly at random:

$$P_{\mathbf{x}}^\sigma \coloneqq \big(v \mapsto \Pi_v^\sigma \mathbf{x}\big)_* \operatorname{Unif}(V) = \frac{1}{|V|} \sum_{v \in V} \delta_{\Pi_v^\sigma \mathbf{x}} \in \operatorname{Prob}^\Gamma(\mathtt{A}^\Gamma).$$

The shift-invariance of every empirical distribution is the reason we assumed μ above was shift-invariant.

By analogy with statistical physics we will call $\mathbf{x}$ a microstate, and we will call it a good model for μ if its empirical distribution (over some given σ) is close to μ. More specifically, if $\mathcal{O}$ is some weak-open neighborhood of μ then we say $\mathbf{x}$ is an $\mathcal{O}$-microstate if $P_{\mathbf{x}}^\sigma \in \mathcal{O}$. We call the set of such $\mathbf{x}$

$$\Omega(\sigma, \mathcal{O}) = \{\mathbf{x} \in \mathtt{A}^V \,:\, P_{\mathbf{x}}^\sigma \in \mathcal{O}\}.$$

This is equivalent to Lewis Bowen's framework of "approximating partitions" introduced in [Bow10c] to define sofic entropy. We will discuss entropy below.

This notion of "good model" is most meaningful when the graph of σ has a high degree of local similarity to Γ. We will measure this in the following way: given $R \in \mathbb{N}$, define

$$\delta_R^\sigma = \frac{1}{V} |\{v \in V \,:\, \mathrm{B}^\sigma(v, R) \ncong \mathrm{B}^\Gamma(e, R)\}|.$$

Here the isomorphism is between the subgraphs induced by the radius-R balls centered at v in the graph of σ and those centered at the identity in the Cayley graph of Γ. Recall that we consider edges of the graph of σ and of the Cayley graph to be directed and labeled by the generators of Γ; we require isomorphisms to respect this structure.

We then make the slightly more ad hoc definition

$$\Delta^\sigma = \inf_R \left(9 \cdot (2/3)^R + 6\delta_R^\sigma\right).$$

The particular constants appearing here come from our choice of metric on $\mathtt{A}^\Gamma$ (see Chapter 3) and from the proof of Theorem 3.2.1. If Δ^σ is small, then the graph of σ looks like Γ to a large radius near most vertices. Note that the notation Δ^σ does not need to explicitly specify which Γ the graph of σ is being compared to, since the relevant Γ is always the domain of σ.

Let $\Sigma = (\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n))_{n\in\mathbb{N}}$ be a sequence of homomorphisms, with V_n finite sets. Recall that Σ is a *sofic approximation* to Γ if $\lim_{n\to\infty} \Delta^{\sigma_n} = 0$. The sofic entropy of $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ relative to Σ is defined by

$$\mathrm{h}_\Sigma(\mu) = \inf_{\mathcal{O}\ni\mu} \limsup_{n\to\infty} \frac{1}{|V_n|} \log|\Omega(\sigma_n, \mathcal{O})|,$$

where the infimum is over weak-open neighborhoods of μ. Informally, we would expect $|\Omega(\sigma_n, \mathcal{O})|$ to grow exponentially with $|V_n|$, with a higher exponential growth rate indicating fewer constraints imposed by μ on its good models (so μ is "more random"). In general, though, sofic entropy may behave in counterintuitive ways. While it is an isomorphism invariant, an example of Ornstein and Weiss [OW87] shows that it may increase under factor maps when Γ is not amenable.

The assumption $\Delta^{\sigma_n} \to 0$ is interpreted here as a kind of Benjamini-Schramm convergence, but we can also view it as requiring the actions $\Gamma \curvearrowright^{\sigma_n} V_n$ to be "asymptotically free." More generally we could only require that they be "asymptotically actions" (see for example [Bow20a]) but for simplicity we only consider true homomorphisms here.

Here, we restrict attention to measures μ which are Gibbs for some nearest-neighbor interaction; relevant definitions are given in Chapter 3. For a nearest-neighbor interaction Φ, we denote the set of Gibbs measures by $\mathscr{G}(\Phi) \subset \operatorname{Prob}(\mathtt{A}^\Gamma)$. The set of shift-invariant Gibbs measures is denoted $\mathscr{G}^\Gamma(\Phi)$. An interaction also comes with an associated "Glauber dynamics" which is a natural and useful model for the random evolution of a system over time. We will use subscripts to denote evolution under Glauber dynamics; for example $\mathbf{x}_s$ is the (random) evolution of the microstate $\mathbf{x}_0$.

Our first main result, Theorem G, establishes the metastability of Gibbs microstates under Glauber dynamics:

Theorem G. *Let $\mu \in \mathscr{G}^\Gamma(\Phi)$ for some nearest-neighbor interaction Φ. Denote its evolution under the Glauber dynamics for Φ as $\{\mu_t \,:\, t \geq 0\}$.*

Given any neighborhood $\mathcal{U}_1$ of μ and $t, \varepsilon > 0$, there exists a neighborhood $\mathcal{U}_0$ of μ and $\delta > 0$ such that, for any finite set V and any homomorphism $\sigma \colon \Gamma \to \operatorname{Sym}(V)$, if $\mathbf{x}_0 \in \Omega(\sigma, \mathcal{U}_0)$ and $\Delta^\sigma < \delta$ then $\mathbf{x}_s \in \Omega(\sigma, \mathcal{U}_1)$ for all $s \in [0,t]$ with probability at least $1 - \varepsilon$.

We call this "metastability" because, if we let the Glauber dynamics run forever, the law of $\mathbf{x}$ will converge to the (unique) Gibbs measure on $\mathtt{A}^V$. In particular, we will eventually lose control of its empirical distribution. Theorem G only says that for any *fixed* time t, it can be arranged for the empirical distribution to stay close to μ for time t with probability as close to 1 as desired. The only requirements are that Δ^σ be small enough and that $P^\sigma_{\mathbf{x}}$ start close enough to μ.

Recall that Theorem 3.2.1 gives a type of equivariance between the Glauber dynamics on Γ and on graphs of homomorphisms σ with small Δ^σ: it implies that if $\mathbf{x}$ is a good model for μ (not necessarily Gibbs) then the *expected* empirical distribution of the evolved microstate $\mathbf{x}_t$ stays close to the evolved measure μ_t. The rate at which it drifts away is controlled by Δ^σ. But if μ is Gibbs then it is Glauber-invariant, so in fact the expected empirical distribution stays close to μ.

It turns out to be somewhat difficult to conclude that the empirical distribution actually stays close to μ with high probability. We do this in two steps: first we use the fact that the Gibbs measures form a face of the convex set $\mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$, combined with the mentioned equivariance result, to show that the empirical distribution of $\mathbf{x}$ stays approximately Gibbs for the desired amount of time with high probability. We then use this approximate Gibbsness to show that the empirical distribution tends to move slowly, so typically stays close μ.

Using Theorem G we establish Theorem E, which says that any maximal-entropy joining of two Gibbs measures (possibly for different interactions) must itself be a Gibbs measure for a natural "sum interaction," except in degenerate cases:

Theorem E. *Let λ be a joining of two shift-invariant Gibbs measures $\mu_{\mathtt{A}} \in \mathscr{G}^\Gamma(\Phi^{\mathtt{A}}), \mu_{\mathtt{B}} \in \mathscr{G}^\Gamma(\Phi^{\mathtt{B}})$ for nearest-neighbor interactions $\Phi^{\mathtt{A}}, \Phi^{\mathtt{B}}$. Let Σ be a random sofic approximation to Γ, and assume that there is some joining λ of $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$ with $\mathrm{h}_\Sigma(\lambda) > -\infty$.*

If λ maximizes h_Σ among all joinings of $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$, then $\lambda \in \mathscr{G}^\Gamma(\Phi^{\mathtt{A}} \oplus \Phi^{\mathtt{B}})$.

Here, a random sofic approximation is a sequence of random homomorphisms such that for any $\delta > 0$ the probability of the event $\{\Delta^{\sigma_n} < \delta\}$ approaches 1 superexponentially fast; see Section 5.3. The f-invariant, introduced in [Bow10b], can be written as the sofic entropy relative to a random sofic approximation to a free group [Bow10a].

By Lemma 5.3.1, we can equivalently say that a maximal-entropy joining of two Gibbs measures must be a relative product over the tail σ-algebra.

We also mention two brief corollaries: Corollary 5.3.4 shows that if $\mu_{\mathtt{A}}$ is a shift-invariant extreme point of $\mathscr{G}(\Phi^{\mathtt{A}})$ and $\mu_{\mathtt{B}}$ is any element of $\mathscr{G}^\Gamma(\Phi^{\mathtt{B}})$, then in fact their product joining is the only joining which is Gibbs for the sum interaction. In particular, for any Σ the product joining is the joining with maximal h_Σ. By "extreme" here we mean extreme in the convex set of all Gibbs measures, which is equivalent to triviality on the tail σ-algebra. In particular, ergodicity is not a sufficient condition. As discussed below, one example illustrating this is the Ising model on a free group, which is ergodic for all $\varepsilon > 0$.

Corollary 5.3.5 shows that, except in degenerate cases, Gibbs measures have nonzero sofic entropy over any deterministic sofic approximation.

Our final main result is Theorem H, which asserts that, for free-boundary Ising models at low temperatures, the self-joining with maximal f-invariant is neither the product nor the diagonal joining. Non-maximality of the diagonal joining actually follows in much greater generality from Theorem E, since the diagonal joining is Gibbs only in degenerate cases. The product joining is always Gibbs for the sum interaction. But for temperatures low enough that the f-invariant is negative, the product joining cannot be maximal because it has smaller f-invariant than the diagonal.

Theorem H actually extends non-maximality of the product to slightly higher temperatures. To do this, we show that if the product joining of μ has optimal f-invariant, then a typical random homomorphism supports good models for μ. We can rule out this possibility for free-boundary Ising models at low temperatures using [DMS17].

It remains open whether non-maximality of the product holds all the way up to the reconstruction threshold, at and above which the product joining is maximal by Corollary 5.3.4. A similar type of result in the recent paper [CLM20] suggests that it may.

5.1.1 Overview

In Section 5.2 we prove Theorem G, our main metastability result. In Section 5.3 we give an application of this theorem, characterizing which joinings of two Gibbs states have maximal sofic entropy over a random sofic approximation. Finally, in Section 5.4 we show that, below a certain (nontrivial) temperature, the product self-joining of a free-boundary Ising state does not have maximal f-invariant.

5.2 Metastability of near-Gibbs-ness

If we apply Proposition 3.4.1 with $\zeta = \delta_{\mathbf{x}}$ and $\mu \in \mathscr{G}^\Gamma$ we get

$$\bar{d}\big(S^\sigma(t)P^\sigma_{\mathbf{x}},\ \mu\big) \leq \big(\bar{d}\big(P^\sigma_{\mathbf{x}},\ \mu\big) + \Delta^\sigma t\big)e^{Mt}. \tag{5.1}$$

In particular, if $\mathbf{x}$ is a good model over σ for a Gibbs measure μ, then the *expected* empirical distribution of $\mathbf{x}_t$ stays close to μ for a long time. The first main theorem of the present chapter, Theorem G, states that, in fact, the empirical distribution itself stays close to μ for a long time with high probability.

The remainder of this section is devoted to the proof of this theorem. First we use Lemma 3.3.1 to show that Equation 5.1 implies $P^\sigma_{\mathbf{x}_t}$ must stay close to $\mathscr{G}^\Gamma$ for a long time with high probability. We then control the 'lateral motion,' showing that as long as $P^\sigma_{\mathbf{x}_t}$ stays close to $\mathscr{G}^\Gamma$ it does not tend to move much at all.

5.2.1 Concentration from convexity

Let I denote the weak*-continuous map

$$\begin{aligned} I\colon\ \mathrm{Prob}(\mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)) &\to \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma) \\ \xi &\mapsto \int \nu\, \xi(d\nu). \end{aligned}$$

Lemma 3.3.1 stated that $\theta(\mathscr{G}^\Gamma) = 1$ whenever $I(\theta) \in \mathscr{G}^\Gamma$. The following result is an approximate version of this: if $I(\theta)$ is close to $\mathscr{G}^\Gamma$, then most of the mass of θ must be close to $\mathscr{G}^\Gamma$.

Proposition 5.2.1. *Given any weak* neighborhood $\mathcal{W}$ of $\mathscr{G}^\Gamma$ and $\varepsilon > 0$, there exists a weak* neighborhood $\mathcal{U}$ of $\mathscr{G}^\Gamma$ such that if $I(\xi) \in \mathcal{U}$ then $\xi(\mathcal{W}) > 1 - \varepsilon$.*

Proof. By the portmanteau theorem, the set $\mathcal{E} = \{\xi\ :\ \xi(\mathcal{W}) > 1 - \varepsilon\}$ is weak*-open, and it clearly contains the set $\mathrm{Prob}(\mathscr{G}^\Gamma)$ of probability measures supported on $\mathscr{G}^\Gamma$. We complete the proof by contradiction: suppose that for each neighborhood $\mathcal{U}$ of $\mathscr{G}^\Gamma$ the intersection $(I^{-1}\mathcal{U}) \cap \mathcal{E}^c$ is nonempty.

For each $n \in \mathbb{N}$, let $\mathcal{U}_n$ be the set of measures within $\bar{d}$-distance $1/n$ of $\mathscr{G}^\Gamma$. By assumption, we can pick a sequence $\xi_n \in (I^{-1}\mathcal{U}_n) \cap \mathcal{E}^c$. Now $\operatorname{Prob}(\operatorname{Prob}^\Gamma(\mathtt{A}^\Gamma))$ is compact, so ξ_n has some convergent subsequence ξ_{n_j}. Note the limit of this sequence must still be in the closed set $\mathcal{E}^c$. By definition of the sets $\mathcal{U}_n$ and continuity of I, the limit must also be in $I^{-1}(\mathscr{G}^\Gamma)$.

But since $\mathscr{G}^\Gamma$ is a face of $\operatorname{Prob}^\Gamma(\mathtt{A}^\Gamma)$ (Lemma 3.3.1), in fact $I^{-1}(\mathscr{G}^\Gamma) = \operatorname{Prob}(\mathscr{G}^\Gamma) \subseteq \mathcal{E}$. This is a contradiction, so there must exist some neighborhood $\mathcal{U}$ of $\mathscr{G}^\Gamma$ with $I^{-1}\mathcal{U} \subseteq \mathcal{E}$. □

Proposition 5.2.2. *Let $\mathcal{W}$ be a weak* neighborhood of $\mathscr{G}^\Gamma$. Let $\varepsilon, t > 0$. Then there exists a weak* neighborhood $\mathcal{U}$ of $\mathscr{G}^\Gamma$ and $\delta > 0$ such that if $P^\sigma_{\mathbf{x}_0} \in \mathcal{U}$ and $\Delta^\sigma < \delta$, then $P^\sigma_{\mathbf{x}_t} \in \mathcal{W}$ with probability at least $1 - \varepsilon$.*

Proof. The previous proposition guarantees the existence of a neighborhood $\mathcal{V}$ of $\mathscr{G}^\Gamma$ such that if $S^\sigma(t)P^\sigma_{\mathbf{x}_0} \in \mathcal{V}$ then $P^\sigma_{\mathbf{x}_t} \in \mathcal{W}$ with probability at least $1 - \varepsilon$.

Since $\mathscr{G}^\Gamma$ is compact, we can pick $\eta > 0$ such that $\bigcup_{\mu \in \mathscr{G}^\Gamma} \mathrm{B}^{\bar{d}}(\mu, 2\eta) \subset \mathcal{V}$. Let $\mathcal{U} = \bigcup_{\mu \in \mathscr{G}^\Gamma} \mathrm{B}^{\bar{d}}(\mu, \eta e^{-Mt})$ and let $\delta = \eta e^{-Mt}/t$. Then by (5.1) whenever $P^\sigma_{\mathbf{x}_0} \in \mathcal{U}$ and $\Delta^\sigma < \delta$ we have $S^\sigma(t)P^\sigma_{\mathbf{x}} \in \mathcal{V}$. □

5.2.2 Controlling lateral motion

Having shown that Glauber dynamics tends to stay within the set of good models for near-Gibbs measures, we now show that it tends to move slowly within this region.

Given $\mathbf{x}_0 \in \mathtt{A}^V$, $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$, $g \colon \mathtt{A}^\Gamma \to \mathbb{R}$, and $\tau > 0$, we define a martingale $(M^{g,\tau}_k)_{k=0}^\infty$ by

$$M^{g,\tau}_k = P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_0} g - \sum_{s=0}^{k-1} L_\tau P^\sigma_{\mathbf{x}_{s\tau}} g$$

where

$$L_\tau \coloneqq S^\sigma(\tau) - I.$$

We first show that the terms in the sum stay small as long as $P^{\sigma}_{\mathbf{x}_{s\tau}}$ stays close to $\mathscr{G}^{\Gamma}$ (which we know is likely to happen as long as $P^{\sigma}_{\mathbf{x}_0}$ is close enough to $\mathscr{G}^{\Gamma}$), then we show that the martingale itself likely stays small by bounding the variance. This will imply that $P^{\sigma}_{\mathbf{x}_t} g$ tends to stay near its initial value.

5.2.2.1 Bounding deviation from martingale

It is straightforward from the definitions to show that $\mu \in \mathscr{G}^{\Gamma}$ then $\mu\Omega^{\Gamma} = 0$. We now show that if μ is near $\mathscr{G}^{\Gamma}$ then $\mu\Omega^{\Gamma}$ is near 0.

Let $\|f\|_{\mathrm{BL}} = \max\{|f|_{\mathrm{Lip}}, \|f\|_{\infty}\}$ denote the bounded Lipschitz norm of a real-valued function on $\mathtt{A}^{\Gamma}$. Under this norm, the set $\{f \,:\, \|f\|_{\mathrm{BL}} < \infty\}$ is a Banach space which we call BL. Every $\mu \in \mathrm{Prob}(\mathtt{A}^{\Gamma})$ induces a continuous linear functional I_{μ} on BL defined by

$$I_{\mu} f = \int f \, d\mu.$$

If we endow the continuous dual BL* with the standard dual (operator) norm, it is easy to see that

$$\bar{d}(\mu, \nu) = \|I_{\mu} - I_{\nu}\|_{\mathrm{BL}^*}.$$

Since $\Omega^{\Gamma} g$ is a continuous function whenever $g \in \mathrm{BL}$, for any $\mu \in \mathrm{Prob}(\mathtt{A}^{\Gamma})$ we can define $\mu\Omega^{\Gamma} \in \mathrm{BL}^*$ by

$$\big(\mu\Omega^{\Gamma}\big) g \coloneqq \int \Omega^{\Gamma} g \, d\mu \quad \forall g \in \mathrm{BL}\,.$$

Lemma 5.2.3. *The map*

$$\begin{aligned} \mathrm{Prob}(\mathtt{A}^{\Gamma}) &\to \mathrm{BL}^* \\ \mu &\mapsto \mu\Omega^{\Gamma} \end{aligned}$$

is continuous.

Proof. We first show that the family $F = \{\Omega^{\Gamma} f \,:\, \|f\|_{\mathrm{BL}} \leq 1\}$ is uniformly bounded and equicontinuous. Uniform boundedness is fairly straightforward. We now establish equicontinuity: Suppose $\mathbf{x}, \mathbf{y} \in \mathtt{A}^{\Gamma}$ are such that $d(\mathbf{x}, \mathbf{y}) < (3r)^{-k}$; then $\mathbf{x}$ and $\mathbf{y}$ agree on $\mathrm{B}(e, k)$ so for all $\gamma \in \mathrm{B}(e, k-1)$ and $\mathtt{a} \in \mathtt{A}$ we have

$$c_\gamma(\mathbf{x}, \mathtt{a}) = c_\gamma(\mathbf{y}, \mathtt{a}).$$

So for such γ, if $\|f\|_{\mathrm{BL}} \le 1$ we have

$$\begin{aligned}
&\left|c_\gamma(\mathbf{x}, \mathtt{a})\left[f(\mathbf{x}^{\gamma\to\mathtt{a}}) - f(\mathbf{x})\right] - c_\gamma(\mathbf{y}, \mathtt{a})\left[f(\mathbf{y}^{\gamma\to\mathtt{a}}) - f(\mathbf{y})\right]\right| \\
&\qquad \le c_\gamma(\mathbf{x}, \mathtt{a})\big[|f(\mathbf{x}) - f(\mathbf{y})| + |f(\mathbf{x}^{\gamma\to\mathtt{a}}) - f(\mathbf{y}^{\gamma\to\mathtt{a}})|\big] \\
&\qquad \le c_\gamma(\mathbf{x}, \mathtt{a})\big[2 \cdot (2r)^{-k}\big].
\end{aligned}$$

For $\gamma \notin \mathrm{B}(e, k-1)$ we have $d(\mathbf{x}^{\gamma\to\mathtt{a}}, \mathbf{x}) = (3r)^{-|\gamma|}$, so

$$\left|c_\gamma(\mathbf{x}, \mathtt{a})\left[f(\mathbf{x}^{\gamma\to\mathtt{a}}) - f(\mathbf{x})\right] - c_\gamma(\mathbf{y}, \mathtt{a})\left[f(\mathbf{y}^{\gamma\to\mathtt{a}}) - f(\mathbf{y})\right]\right| \le \left|c_\gamma(\mathbf{x}, \mathtt{a}) - c_\gamma(\mathbf{y}, \mathtt{a})\right|(3r)^{-|\gamma|}.$$

Hence

$$\begin{aligned}
|\Omega^\Gamma f(\mathbf{x}) - \Omega^\Gamma f(\mathbf{y})| &\le \sum_{\gamma \in \mathrm{B}(e,k-1)} \sum_{\mathtt{a} \in \mathtt{A}} 2 c_\gamma(\mathbf{x}, \mathtt{a}) \cdot (3r)^{-k} \\
&\qquad + \sum_{\gamma \notin \mathrm{B}(e,k-1)} \sum_{\mathtt{a} \in \mathtt{A}} (c_\gamma(\mathbf{x}, \mathtt{a}) + c_\gamma(\mathbf{y}, \mathtt{a}))(3r)^{-|\gamma|} \\
&\le 2|\mathrm{B}(e, k-1)|(3r)^{-k} + 2\sum_{s=k}^{\infty} r^s (3r)^{-s} \\
&= 2|\mathrm{B}(e, k-1)|(3r)^{-k} + 3 \cdot 3^{-k} \\
&= o_{k\to\infty}(1).
\end{aligned}$$

Since this bound is uniform over $f \in F$, the family F is equicontinuous.

Suppose $(\mu_n)_{n=1}^\infty$ is a sequence of probability measures with weak* limit ν. For any $\varepsilon > 0$, by Arzelà-Ascoli we can pick a finite collection $\Omega^\Gamma f_1, \ldots, \Omega^\Gamma f_k \in F$ which is uniformly ε-dense in F. Hence

$$\begin{aligned}
\limsup_{n\to\infty} \|\mu_n \Omega^\Gamma - \nu \Omega^\Gamma\|_{\mathrm{BL}^*} &= \limsup_{n\to\infty} \sup \left\{ \left| \int \Omega^\Gamma f \, d\mu - \int \Omega^\Gamma f \, d\nu \right| \,:\, \|f\|_{BL} \le 1 \right\} \\
&\le \limsup_{n\to\infty} \left[\max \left\{ \left| \int \Omega^\Gamma f_i \, d\mu_n - \int \Omega^\Gamma f_i \, d\nu \right| \,:\, 1 \le i \le k \right\} + 2\varepsilon \right] \\
&= 2\varepsilon.
\end{aligned}$$

Since ε is arbitrary, this shows that $\mu_n \Omega^\Gamma$ converges to $\nu \Omega^\Gamma$. $\square$

Proposition 5.2.4. *For any* $\mathbf{x} \in \mathtt{A}^V$, $\tau > 0$, *and* g *with* $|g|_{\mathrm{Lip}} \leq 1$

$$|L_\tau(P^\sigma_{\mathbf{x}} g)| \leq \tau \left[\Delta^\sigma e^{M\tau} + \left\| \frac{S^\Gamma(\tau) g - g}{\tau} - \Omega^\Gamma g \right\|_\infty + \|P^\sigma_{\mathbf{x}} \Omega^\Gamma\|_{\mathrm{BL}^*} \right].$$

Proof. For any g,

$$\begin{aligned} |L_\tau(P^\sigma_{\mathbf{x}} g)| &= |S^\sigma(\tau) P^\sigma_{\mathbf{x}} g - P^\sigma_{\mathbf{x}} g| \\ &\leq |S^\sigma(\tau) P^\sigma_{\mathbf{x}} g - P^\sigma_{\mathbf{x}} S^\Gamma(\tau) g| + |P^\sigma_{\mathbf{x}} S^\Gamma(\tau) g - P^\sigma_{\mathbf{x}} g|. \end{aligned}$$

By Theorem 3.2.1, if $|g|_{\mathrm{Lip}} \leq 1$ then the first term here is bounded by $\Delta^\sigma \tau e^{M\tau}$. For the second term, we have

$$\begin{aligned} |P^\sigma_{\mathbf{x}} S^\Gamma(\tau) g - P^\sigma_{\mathbf{x}} g| &\leq \tau \left[P^\sigma_{\mathbf{x}} \left| \frac{S^\Gamma(\tau) g - g}{\tau} - \Omega^\Gamma g \right| + \left| P^\sigma_{\mathbf{x}} \Omega^\Gamma g \right| \right] \\ &\leq \tau \left[\left\| \frac{S^\Gamma(\tau) g - g}{\tau} - \Omega^\Gamma g \right\|_\infty + \|P^\sigma_{\mathbf{x}} \Omega^\Gamma\|_{\mathrm{BL}^*} \right]. \end{aligned}$$

□

5.2.2.2 Martingale concentration

Proposition 5.2.5. *Fix* $t > 0$ *and* g *with* $|g|_{\mathrm{Lip}} \leq 1$. *Then for any* $m \in \mathbb{Z}$ *we have*

$$\mathbb{E}\left[(M^{g,t/m}_m)^2 \right] \leq 9t \left[\frac{1}{|V|} + \frac{t}{m} \right].$$

Proof. Let $\tau = t/m$. Let $\xi_1, \xi_2, \ldots$ denote the martingale increments given by

$$\xi_k = M^{g,\tau}_k - M^{g,\tau}_{k-1} = P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g - L_\tau P^\sigma_{\mathbf{x}_{(k-1)\tau}} g.$$

Let K_k be the number of times a spin changes in the Glauber dynamics starting at $\mathbf{x}_0$ during the time interval $[(k-1)\tau, k\tau)$. We will use that K_k is Poisson with mean $\tau|V|$.

We need the following two lemmas:

Lemma 5.2.6. *If* $\mathbf{x}, \mathbf{x}' \in \mathtt{A}^V$ *differ at exactly one site* $w \in V$, *then*

$$|P^\sigma_{\mathbf{x}} g - P^\sigma_{\mathbf{x}'} g| \leq \frac{3}{|V|}.$$

Proof. Recall that we are assuming $|g|_{\mathrm{Lip}} \le 1$. Using the definitions of empirical distribution and the distance on $\mathtt{A}^V$,

$$\begin{aligned} |P^\sigma_{\mathbf{x}} g - P^\sigma_{\mathbf{x}'} g| \le \frac{3}{|V|} &\le \frac{1}{|V|} \sum_{v\in V} |g(\Pi^\sigma_v \mathbf{x}) - g(\Pi^\sigma_v \mathbf{x}')| \\ &\le \frac{1}{|V|} \sum_{v\in V} d(\Pi^\sigma_v \mathbf{x}, \Pi^\sigma_v \mathbf{x}) \\ &= \frac{1}{|V|} \sum_{v\in V} \sum_{\gamma\in\Gamma} (3r)^{-|\gamma|} \mathbf{1}_{\{\mathbf{x}(\sigma^\gamma v) \ne \mathbf{x}'(\sigma^\gamma v)\}}. \end{aligned}$$

By assumption, $\mathbf{x}(\sigma^\gamma v) \ne \mathbf{x}'(\sigma^\gamma v)$ if and only if $\sigma^\gamma v = w$. Using this fact and changing the order of summation gives

$$|P^\sigma_{\mathbf{x}} g - P^\sigma_{\mathbf{x}'} g| \le \frac{3}{|V|} \le \frac{1}{|V|} \sum_{\gamma\in\Gamma} (3r)^{-|\gamma|} \sum_{v\in V} \mathbf{1}_{\{\sigma^\gamma v = w\}}.$$

But σ^γ is a permutation, so $\sum_{v\in V} \mathbf{1}_{\{\sigma^\gamma v = w\}} = 1$. The result now follows from the bound $\sum_{\gamma\in\Gamma} (3r)^{-|\gamma|} \le 3$. □

Lemma 5.2.7. *For any $k \in \mathbb{N}$,*

$$\mathbb{E}[\xi_k^2] \le 9|V|^{-2}\, \mathbb{E}[K_k^2].$$

Proof. For each k let $\mathcal{F}_k$ be the σ-algebra generated by $(\mathbf{x}_0, \mathbf{x}_\tau, \dots, \mathbf{x}_{k\tau})$.

We first expand out ξ_k^2 using its definition, and then simplify the resulting expression using that $L_\tau P^\sigma_{\mathbf{x}_{(k-1)\tau}} g$ is $\mathcal{F}_{k-1}$-measurable:

$$\begin{aligned} \mathbb{E}[\xi_k^2 \mid \mathcal{F}_{k-1}] &= \mathbb{E}\left[\left(P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g - L_\tau P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)^2 \mid \mathcal{F}_{k-1}\right] \\ &= \mathbb{E}\left[\left(P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)^2 \mid \mathcal{F}_{k-1}\right] + \mathbb{E}\left[\left(L_\tau P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)^2 \mid \mathcal{F}_{k-1}\right] \\ &\quad + \mathbb{E}\left[2\left(P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)\left(-L_\tau P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right) \mid \mathcal{F}_{k-1}\right] \\ &= \mathbb{E}\left[\left(P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)^2 \mid \mathcal{F}_{k-1}\right] + \left(L_\tau P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)^2 \\ &\quad - 2 L_\tau P^\sigma_{\mathbf{x}_{(k-1)\tau}} g \cdot \mathbb{E}\left[P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g \mid \mathcal{F}_{k-1}\right] \\ &= \mathbb{E}\left[\left(P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)^2 \mid \mathcal{F}_{k-1}\right] - \left(L_\tau P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)^2. \end{aligned}$$

Dropping the second term, we're left with

$$\mathbb{E}[\xi_k^2 \mid \mathcal{F}_{k-1}] \le \mathbb{E}\left[\left(P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g\right)^2 \mid \mathcal{F}_{k-1}\right].$$

By the previous lemma, each of the K_k spin flips moves $P^\sigma_{\mathbf{x}} g$ by at most $3/|V|$, so we have

$$|P^\sigma_{\mathbf{x}_{k\tau}} g - P^\sigma_{\mathbf{x}_{(k-1)\tau}} g| \le \tfrac{3}{|V|} K_k.$$

Putting this into the previously obtained bound and taking expectations gives the claimed result. □

Using Lemma 5.2.7 and that $K_k \sim \mathrm{Pois}(\tau|V|)$, we see that

$$\mathbb{E}[\xi_k^2] \le 9|V|^{-2}\,\mathbb{E}[K_k^2] = \tfrac{9}{|V|^2}[\tau|V| + (\tau|V|)^2] = 9\tau[\tfrac{1}{|V|} + \tau].$$

Therefore, since the martingale increments are uncorrelated,

$$\mathbb{E}\left[(M_m^{g,t/m})^2\right] = \sum_{k=1}^m m\,\mathbb{E}[\xi_k^2] \le 9t[\tfrac{1}{|V|} + \tfrac{t}{m}]. \qquad \square$$

By similar methods we can prove the following lemma, which controls the empirical distribution at times between multiples of t/m:

Lemma 5.2.8. *For any $t, \kappa > 0$, g with $\|g\|_{BL} \le 1$, $m \in \mathbb{N}$, and $0 \le j \le m-1$,*

$$\mathbb{P}\left(\exists s \in [jt/m, (j+1)t/m] \text{ with } |P^\sigma_{\mathbf{x}_s} g - P^\sigma_{\mathbf{x}_{jt/m}} g| > \kappa\right) \le 9\frac{t}{m}\left[\frac{1}{|V|} + \frac{t}{m}\right].$$

Proof. By Lemma 5.2.6 the probability is bounded above by

$$\mathbb{P}(K_j > |V|\kappa/3).$$

The result follows from applying Chebyshev's inequality, using that K_j has law $\mathrm{Pois}(t|V|/m)$. □

5.2.3 Proof of Theorem G

Fix $\kappa > 0$ such that $\mathrm{B}(\mu, 9\kappa) \subset \mathcal{U}_1$. Let F be a finite κ-dense (in uniform norm) subset of $\{f : \|f\|_{BL} \leq 1\}$; we showed above that this set is compact in the uniform norm. Then for any $\mu, \nu \in \mathrm{Prob}(\mathtt{A}^\Gamma)$,

$$\bar{d}(\mu, \nu) \leq \sup\{|\mu g - \nu g| : g \in F\} + 2\kappa.$$

For any given $t, \varepsilon > 0$, by Proposition 5.2.5 and Doob's maximal inequality we can pick $M \in \mathbb{N}$ such that for any $\mathbf{x}_0 \in \mathtt{A}^V$

$$\mathbb{P}\left(\max_{g \in F} \max_{0 \leq j \leq m} |M_j^{g, t/m}| \leq \kappa \right) \geq 1 - \varepsilon \tag{*}$$

whenever $|V|, m \geq M$ (recall that the martingale has an implicit dependence on a choice of initial microstate $\mathbf{x}_0 \in \mathtt{A}^V$). By Lemma 5.2.8 we can make M larger if necessary to also ensure that for each $0 \leq j \leq m$ we have

$$\mathbb{P}\big(\max_{g \in F} \max_{s \in [0,t]} |P^\sigma_{\mathbf{x}_s} g - P^\sigma_{\mathbf{x}_{\lfloor sm/t \rfloor t/m}} g| \leq \kappa \big) \geq 1 - \varepsilon. \tag{**}$$

Assume that m is also large enough that

$$\left\| \frac{S^\Gamma(t/m) g - g}{t/m} - \Omega^\Gamma g \right\|_\infty \leq \kappa/t$$

for every $g \in F$ and that $e^{Mt/m} \leq 2$. Assume also that $\Delta^\sigma \leq \kappa/t$ and let $\mathcal{W} = \{\nu : \|\nu\Omega^\Gamma\|_{\mathrm{BL}^*} < \kappa/t\}$; this is an open neighborhood of $\mathscr{G}^\Gamma$ by continuity of the map $\nu \mapsto \nu\Omega^\Gamma$ (Lemma 5.2.3). Then, by Proposition 5.2.4, $P^\sigma_{\mathbf{x}} \in \mathcal{W}$ implies

$$|L_{t/m}(P^\sigma_{\mathbf{x}} g)| \leq \tfrac{4\kappa}{m}.$$

We have also shown (Proposition 5.2.2) that there exist a weak neighborhood $\mathcal{U}$ of $\mathscr{G}^\Gamma$ and $\delta > 0$ such that if $P^\sigma_{\mathbf{x}_0} \in \mathcal{U}$ and $\Delta^\sigma < \delta$, then for each $s \leq t$ we have $\mathbb{P}(P^\sigma_{\mathbf{x}_s} \in \mathcal{W}) \geq 1 - \varepsilon/m$. Therefore under these assumptions

$$\mathbb{P}(P^\sigma_{\mathbf{x}_{kt/m}} \in \mathcal{W} \ \forall 0 \leq k \leq m) \geq 1 - \varepsilon. \tag{***}$$

Suppose that $\mathbf{x}_0 \in \Omega(\sigma, \mathcal{U})$. Then the the probability that the events appearing in (*), (**), and (***) all occur is at least $1 - 3\varepsilon$. Assume they do all occur. Given $g \in F$ and $s \in [0, t]$, pick $j = \lfloor sm/t \rfloor$. Then $0 \leq j \leq m$ so we have

$$\begin{aligned}
|P^\sigma_{\mathbf{x}_s} g - \mu g| &\leq |P^\sigma_{\mathbf{x}_{jt/m}} g - \mu g| + \kappa \\
&\leq |M^{g,t/m}_j| + |P^\sigma_{\mathbf{x}_0} g - \mu g| + \sum_{k=0}^{j-1} |L_{t/m}(P^\sigma_{\mathbf{x}_{kt/m}} g)| + \kappa \\
&\leq \kappa + \bar{d}(P^\sigma_{\mathbf{x}_0}, \mu) + j \cdot \tfrac{4\kappa}{m} + \kappa.
\end{aligned}$$

So if also $\bar{d}(P^\sigma_{\mathbf{x}_0}, \mu) \leq \kappa$ then for any $s \in [0, t]$ we have

$$\sup\{|P_{\mathbf{x}_s} g - \mu g| \,:\, g \in F\} \leq 7\kappa$$

so

$$\bar{d}(P^\sigma_{\mathbf{x}_s}, \mu) \leq 9\kappa$$

and hence $P^\sigma_{\mathbf{x}_s} \in \mathcal{U}_1$.

In summary: let $\mathcal{U}_0 = \mathcal{U} \cap \mathrm{B}(\mu, \kappa)$. If $P^\sigma_{\mathbf{x}_0} \in \mathcal{U}_0$, $|V| \geq M$, and $\Delta^\sigma < \delta$, then with probability at least $1 - 3\varepsilon$ we have $P^\sigma_{\mathbf{x}_s} \in \mathcal{U}_1$ for all $s \in [0, t]$. Since Δ^σ can only be small if $|V|$ is large, we can remove the explicit requirement of a lower bound on $|V|$ by making δ smaller if necessary.

5.3 Maximal-entropy joinings

Recall that we call a sequence of random homomorphisms $\Sigma = (\sigma_n \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V_n)))_{n \in \mathbb{N}}$ a *random sofic approximation* to Γ if for any $\delta > 0$ there exists $c > 0$ such that

$$\mathbb{P}(\Delta^{\sigma_n} > \delta) < n^{-cn}.$$

Examples include deterministic sofic approximations by homomorphisms, uniformly random homomorphisms, and stochastic block models (see Proposition 2.1.3). The assumption that the maps be true homomorphisms has been adopted for simplicity and with a particular application in mind, but is probably not necessary; see [ABL19] for a more general definition.

Recall also that the exponential growth rate for the expected number of good models for $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ is the sofic entropy

$$\mathrm{h}_\Sigma(\mu) = \inf_{\mathcal{O} \ni \mu} \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega(\sigma_n, \mathcal{O})|.$$

If every term of Σ is deterministic then this is the standard sofic entropy. If Γ is a free group and each term of Σ is uniform then this is the f-invariant [Bow10a].

Given two measures $\mu_\mathtt{A} \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ and $\mu_\mathtt{B} \in \mathrm{Prob}^\Gamma(\mathtt{B}^\Gamma)$, which joinings of the two maximize h_Σ for a fixed Σ? This question arose in Chapter 2 and may be of more general interest.

The following theorem provides some information in the case where both systems are Gibbs measures for nearest-neighbor interactions. To state it we need one definition, which is essentially equivalent to [Geo11, Example 7.18]: given two nearest-neighbor interactions $\Phi^\mathtt{A} = (J^\mathtt{A}, h^\mathtt{A}), \Phi^\mathtt{B} = (J^\mathtt{B}, h^\mathtt{B})$ with respective finite alphabets $\mathtt{A}, \mathtt{B}$, define their sum $\Phi^\mathtt{A} \oplus \Phi^\mathtt{B}$ to be the pair $(J^\mathtt{A} \oplus J^\mathtt{B}, h^\mathtt{A} \oplus h^\mathtt{B})$, where

$$[J^\mathtt{A} \oplus J^\mathtt{B}]\big((\mathtt{a}_1, \mathtt{b}_1), (\mathtt{a}_2, \mathtt{b}_2)\big) = J^\mathtt{A}(\mathtt{a}_1, \mathtt{a}_2) + J^\mathtt{B}(\mathtt{b}_1, \mathtt{b}_2)$$

$$[h^\mathtt{A} \oplus h^\mathtt{B}](\mathtt{a}, \mathtt{b}) = h^\mathtt{A}(\mathtt{a}) + h^\mathtt{B}(\mathtt{b}).$$

This is a nearest-neighbor interaction with alphabet $\mathtt{A} \times \mathtt{B}$.

We will use $c_v^\mathtt{A}(\mathbf{x}, \cdot) \in \mathrm{Prob}(\mathtt{A})$ to refer to the transition rates for the Glauber dynamics of $\Phi^\mathtt{A}$, and $U^\mathtt{A} \colon \mathtt{A}^V \to \mathbb{R}$ to refer to the energy, and similarly for $c^\mathtt{B}, U^\mathtt{B}$. Without superscripts, c, U will refer to $\Phi^\mathtt{A} \oplus \Phi^\mathtt{B}$. Note that if $\mathbf{x} \in (\mathtt{A} \times \mathtt{B})^V$ then $c_v(\mathbf{x}, (\mathtt{a}, \mathtt{b})) = c_v^\mathtt{A}(\mathbf{x}_\mathtt{A}, \mathtt{a}) c_v^\mathtt{B}(\mathbf{x}_\mathtt{B}, \mathtt{b})$ and $U(\mathbf{x}) = U^\mathtt{A}(\mathbf{x}_\mathtt{A}) + U^\mathtt{B}(\mathbf{x}_\mathtt{B})$. In particular, the Glauber dynamics for $\Phi^\mathtt{A} \oplus \Phi^\mathtt{B}$ is a coupling of the Glauber dynamics of the summands.

If $\mu_\mathtt{A} \in \mathscr{G}^\Gamma(\Phi^\mathtt{A})$ and $\mu_\mathtt{B} \in \mathscr{G}^\Gamma(\Phi^\mathtt{B})$, then $\mu_\mathtt{A} \times \mu_\mathtt{B} \in \mathscr{G}^\Gamma(\Phi^\mathtt{A} \oplus \Phi^\mathtt{B})$; in particular, there always exist joinings which are Gibbs for the sum interaction.

Theorem E. *Let* λ *be a joining of two shift-invariant Gibbs measures* $\mu_{\mathtt{A}} \in \mathscr{G}^\Gamma(\Phi^{\mathtt{A}}), \mu_{\mathtt{B}} \in \mathscr{G}^\Gamma(\Phi^{\mathtt{B}})$ *for nearest-neighbor interactions* $\Phi^{\mathtt{A}}, \Phi^{\mathtt{B}}$*. Let* Σ *be a random sofic approximation to* Γ*, and assume that there is some joining* λ *of* $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$ *with* $\mathrm{h}_\Sigma(\lambda) > -\infty$*.*

If λ *maximizes* h_Σ *among all joinings of* $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$*, then* $\lambda \in \mathscr{G}^\Gamma(\Phi^{\mathtt{A}} \oplus \Phi^{\mathtt{B}})$*.*

In particular, since h_Σ is upper semicontinuous, there is a Gibbs joining which has maximal h_Σ among all joinings.

The following proposition shows that we can interpret the conclusion of Theorem E as saying that λ is a relatively independent joining over the tail.

Proposition 5.3.1. *Suppose* $\lambda \in \mathscr{G}(\Phi^{\mathtt{A}} \oplus \Phi^{\mathtt{B}})$ *has* $\mathtt{A},\mathtt{B}$*-marginals* $\mu_{\mathtt{A}} \in \mathscr{G}(\Phi^{\mathtt{A}})$*,* $\mu_{\mathtt{B}} \in \mathscr{G}(\Phi^{\mathtt{B}})$ *respectively. Then for* λ*-a.e.* $(\mathbf{x},\mathbf{y}) \in (\mathtt{A}\times\mathtt{B})^\Gamma$*, for every measurable* $E_{\mathtt{A}} \subset \mathtt{A}^\Gamma, E_{\mathtt{B}} \subset \mathtt{B}^\Gamma$ *we have*

$$\begin{aligned}\lambda\big(E_{\mathtt{A}} \times E_{\mathtt{B}} \big| \mathscr{T}^{\mathtt{AB}}\big)(\mathbf{x},\mathbf{y}) &= \lambda\big(E_{\mathtt{A}} \times E_{\mathtt{B}} \big| \mathscr{T}^{\mathtt{A}} \otimes \mathscr{T}^{\mathtt{B}}\big)(\mathbf{x},\mathbf{y}) \\ &= \mu_{\mathtt{A}}\big(E_{\mathtt{A}} \big| \mathscr{T}^{\mathtt{A}}\big)(\mathbf{x}) \cdot \mu_{\mathtt{B}}\big(E_{\mathtt{B}} \big| \mathscr{T}^{\mathtt{B}}\big)(\mathbf{y}),\end{aligned}$$

where $\mathscr{T}^{\mathtt{AB}}, \mathscr{T}^{\mathtt{A}}, \mathscr{T}^{\mathtt{B}}$ *respectively denote the tail* σ*-algebras on* $(\mathtt{A}\times\mathtt{B})^\Gamma, \mathtt{A}^\Gamma, \mathtt{B}^\Gamma$*.*

This proposition is proven in Section 5.5.

Recall Theorem B, which states that the Glauber evolution of any shift-invariant measure converges to the set of Gibbs measures, as long as the group Γ has property PA.

One could prove our Theorem E using Theorem B roughly as follows: Let $\Phi = \Phi^{\mathtt{A}} \oplus \Phi^{\mathtt{B}}$. Starting with an arbitrary $\lambda \in \mathsf{J}(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$, if we evolve under Glauber dynamics for Φ then eventually λ_t will become as close as we like to $\mathscr{G}^\Gamma(\Phi)$, while staying in $\mathsf{J}(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$ (since the marginals are invariant). If we evolve a collection of good models for λ for the same amount of time, our metastability result (Theorem G) implies that they mostly stay good models for approximate joinings. It can also be shown that the evolved collection is almost as large as the initial one, and that most of the evolved states are good models for Gibbs states. From this we could conclude that there is a Gibbs state with at least as many good models as λ.

However, it turns out to be easier to directly use Proposition 4.2.3, which was the main technical result used to prove Theorem B.

Here is a brief summary of the proof of Theorem E:

Suppose λ is a joining of $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$ which is not Gibbs. Fix n and $\mathcal{O} \ni \lambda$. Let p_0 be the uniform distribution on $\Omega(\sigma_n, \mathcal{O}) \subset (\mathtt{A} \times \mathtt{B})^{V_n}$, and let p_t denote its evolution under the Glauber dynamics for Φ.

Since the marginals of λ are Gibbs, and hence invariant under the Glauber dynamics, the average energy $p_t(U)$ is approximately constant over time. But we know that the free energy $A(p_t)$ is strictly decreasing since λ is not Gibbs. This means that the Shannon entropy of p_t must be strictly increasing (up to a small error). But the Shannon entropy of p_0 is $\log|\Omega(\sigma_n, \mathcal{O})|$, and p_t is mostly supported on good models for approximate joinings of $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$ (with the quality of the approximation getting better as $\Delta^{\sigma_n} \to 0$).

The evolved measure p_t having strictly larger entropy means that its support, which is mostly good models for approximate joinings of $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$, must be strictly larger than the set of good models for the particular joining λ. This will imply that $\mathrm{h}_\Sigma(\lambda)$ is not maximal.

Note that we do not know whether p_t stays mostly supported on good models for λ_t; we just know that its expected empirical distribution is near λ_t. So we cannot simply say that $\mathrm{h}_\Sigma(\lambda_t)$ is increasing.

The connection between entropy and the size of support is made using the following variant of Fano's inequality, standard versions of which can be found in [CT06].

Lemma 5.3.2. *Let $\mathcal{F}$ be a finite set and let $p \in \mathrm{Prob}(\mathcal{F})$. If $E \subset \mathcal{F}$ satisfies $p(E) \geq 1 - \varepsilon$ for some $\varepsilon > 0$ then*

$$\log|E| \geq \mathrm{H}(p) - [\log 2 + \varepsilon \log|\mathcal{F}|].$$

Proof. Using the definition of Shannon entropy and splitting terms according to E and its complement,

$$\begin{aligned}
\mathrm{H}(p) &= -\sum_{x\in E} p\{x\}\log p\{x\} - \sum_{x\notin E} p\{x\}\log p\{x\} \\
&= -\left[p(E)\sum_{x\in E}\frac{p\{x\}}{p(E)}\log\frac{p\{x\}}{p(E)} + p(E)\log p(E)\right] \\
&\quad -\left[(1-p(E))\sum_{x\notin E}\frac{p\{x\}}{1-p(E)}\log\frac{p\{x\}}{1-p(E)} + (1-p(E))\log(1-p(E))\right].
\end{aligned}$$

Let $p_E \in \operatorname{Prob}(E)$ denote the renormalized restriction of p to E, and similarly define $p_{E^c} \in \operatorname{Prob}(E^c)$. Then the above can be written

$$\begin{aligned}
\mathrm{H}(p) &= p(E)\mathrm{H}(p_E) + (1-p(E))\mathrm{H}(p_{E^c}) + \mathrm{H}(p(E), 1-p(E)) \\
&\leq \log|E| + \varepsilon\log|\mathcal{F}| + \log 2.
\end{aligned}$$

Rearranging gives the claimed inequality. □

The following proposition shows that the number of good models for any non-Gibbs joining is strictly smaller (by an exponential factor) than the number of good models for approximate joinings.

Proposition 5.3.3. *Suppose $\lambda \in \mathsf{J}(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$ is not in $\mathscr{G}(\Phi)$. There exist constants $C_1, C_2 > 0$ such that for any $\varepsilon, \eta > 0$ there exist $\delta > 0$ and $\mathcal{O} \ni \lambda$ such that if $\sigma \in \operatorname{Hom}(\Gamma, \operatorname{Sym}(V))$ satisfies $\Delta^\sigma < \delta$ then*

$$|\Omega(\sigma, \mathsf{J}^\eta(\mu_{\mathtt{A}}, \mu_{\mathtt{B}}))| \geq |\Omega(\sigma, \mathcal{O})| \cdot \tfrac{1}{2}\exp\bigl[|V|(C_1 - \varepsilon C_2)\bigr].$$

Proof. Note that if $\Omega(\sigma, \mathcal{O})$ is empty then the inequality is trivially satisfied, so we will assume below that this is not the case.

First, using that λ is not Gibbs, pick $\delta, c, T > 0$ and $\mathcal{O} \ni \lambda$ as appear in Proposition 4.2.3. Fix $t \in (0, T]$ arbitrarily.

Note that for convenience we may assume $\varepsilon < \eta$. By Theorem G, by making $\delta, \mathcal{O}$ smaller if necessary we can ensure that if $\mathbf{x}_0 \in \Omega(\sigma, \mathcal{O})$ and $\Delta^\sigma < \delta$ then $P^\sigma_{\mathbf{x}_t} \in \mathsf{J}^\varepsilon(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$ with probability at least $1-\varepsilon$. Consequently, if we let $p_0 = \operatorname{Unif}(\Omega(\sigma, \mathcal{O}))$ then $p_t(\Omega(\sigma, \mathsf{J}^\varepsilon(\mu_{\mathtt{A}}, \mu_{\mathtt{B}}))) > 1-\varepsilon$;

note that since $\operatorname{diam} \mathtt{A}^\Gamma, \operatorname{diam} \mathtt{B}^\Gamma \leq 3$ this implies $P^\sigma_{p_t} \in \mathsf{J}^{4\varepsilon}(\mu_\mathtt{A}, \mu_\mathtt{B})$. Also, for convenience we may shrink $\mathcal{O}$ if necessary to ensure $\mathcal{O} \subset \mathsf{J}^{\varepsilon}(\mu_\mathtt{A}, \mu_\mathtt{B})$.

We now show that the entropy of p_t is increasing, up to a small error. By choice of $\delta, c, t, \mathcal{O}$ we have

$$A(p_t) \leq A(p_0) - c|V|t,$$

or equivalently

$$\mathrm{H}(p_t) \geq \mathrm{H}(p_0) + c|V|t - [p_0(U) - p_t(U)].$$

Since the empirical distributions of p_0, p_t have approximately the same marginals, the difference in average energy is small. Specifically, since $P^\sigma_{p_0} \in \mathsf{J}^{\varepsilon}(\mu_\mathtt{A}, \mu_\mathtt{B})$ and $P^\sigma_{p_t} \in \mathsf{J}^{4\varepsilon}(\mu_\mathtt{A}, \mu_\mathtt{B})$

$$\begin{aligned}
|p_t(U) - p_0(U)| &= |V| \cdot |P^\sigma_{p_t} U_e - P^\sigma_{p_0} U_e| \\
&\leq |V| \cdot \big(|P^\sigma_{p_t} U_e - [\mu_\mathtt{A} U_e^\mathtt{A} + \mu_\mathtt{B} U_e^\mathtt{B}]| + |[\mu_\mathtt{A} U_e^\mathtt{A} + \mu_\mathtt{B} U_e^\mathtt{B}] - P^\sigma_{p_0} U_e| \big) \\
&\leq |V| \cdot \big(|\pi_\mathtt{A} P^\sigma_{p_t} U_e^\mathtt{A} - \mu_\mathtt{A} U_e^\mathtt{A}| + |\pi_\mathtt{B} P^\sigma_{p_t} U_e^\mathtt{B} - \mu_\mathtt{B} U_e^\mathtt{B}| \\
&\quad + |\pi_\mathtt{A} P^\sigma_{p_0} U_e^\mathtt{A} - \mu_\mathtt{A} U_e^\mathtt{A}| + |\pi_\mathtt{B} P^\sigma_{p_0} U_e^\mathtt{B} - \mu_\mathtt{A} U_e^\mathtt{B}| \big) \\
&\leq 5\varepsilon |V| (|U_e^\mathtt{A}|_{\mathrm{Lip}} + |U_e^\mathtt{B}|_{\mathrm{Lip}})
\end{aligned}$$

so

$$\mathrm{H}(p_t) \geq \mathrm{H}(p_0) + c|V|t - 5\varepsilon |V| (|U_e^\mathtt{A}|_{\mathrm{Lip}} + |U_e^\mathtt{B}|_{\mathrm{Lip}}).$$

By Lemma 5.3.2,

$$\begin{aligned}
&\log|\Omega(\sigma, \mathsf{J}^{\varepsilon}(\mu_\mathtt{A}, \mu_\mathtt{B}))| \\
&\quad \geq \mathrm{H}(p_t) - [\log 2 + \varepsilon |V| \log|\mathtt{A} \times \mathtt{B}|] \\
&\quad \geq \mathrm{H}(p_0) + c|V|t - [5\varepsilon |V| (|U_e^\mathtt{A}|_{\mathrm{Lip}} + |U_e^\mathtt{B}|_{\mathrm{Lip}})] - [\log 2 + \varepsilon |V| \log|\mathtt{A} \times \mathtt{B}|] \\
&\quad = \log|\Omega(\sigma, \mathcal{O})| + |V| \big(ct - \varepsilon [5 |U_e^\mathtt{A}|_{\mathrm{Lip}} + 5 |U_e^\mathtt{B}|_{\mathrm{Lip}} + \log|\mathtt{A} \times \mathtt{B}|] \big) - \log 2.
\end{aligned}$$

Since $\Omega(\sigma, \mathsf{J}^{\varepsilon}(\mu_\mathtt{A}, \mu_\mathtt{B})) \subset \Omega(\sigma, \mathsf{J}^{\eta}(\mu_\mathtt{A}, \mu_\mathtt{B}))$, exponentiating both sides gives the claimed inequality with $C_1 = ct$ and $C_2 = 5|U_e^\mathtt{A}|_{\mathrm{Lip}} + 5|U_e^\mathtt{B}|_{\mathrm{Lip}} + \log|\mathtt{A} \times \mathtt{B}|$. □

Proof of Theorem E. Suppose λ is not Gibbs, and pick $\varepsilon, \eta > 0$. By Proposition 5.3.3 we can pick $\delta > 0$ and $\mathcal{O} \ni \lambda$ such that if $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$ satisfies $\Delta^\sigma < \delta$ then

$$|\Omega(\sigma, \mathsf{J}^\eta(\mu_{\mathtt{A}}, \mu_{\mathtt{B}}))| \geq |\Omega(\sigma, \mathcal{O})| \cdot \tfrac{1}{2} \exp\big[|V|(C_1 - \varepsilon C_2)\big].$$

Since the probability that $\Delta^{\sigma_n} < \delta$ approaches 1 superexponentially fast in n, this implies

$$\begin{aligned}
\limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega(\sigma_n, \mathsf{J}^\eta(\mu_{\mathtt{A}}, \mu_{\mathtt{B}}))| &\geq \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega(\sigma_n, \mathcal{O})| + (C_1 - \varepsilon C_2) \\
&\geq \inf_{\mathcal{O}\ni\lambda} \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega(\sigma_n, \mathcal{O})| + C_1 - \varepsilon C_2 \\
&= \mathrm{h}_\Sigma(\lambda) + C_1 - \varepsilon C_2
\end{aligned}$$

Since $\varepsilon > 0$ was arbitrary,

$$\limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega(\sigma_n, \mathsf{J}^\eta(\mu_{\mathtt{A}}, \mu_{\mathtt{B}}))| \geq \mathrm{h}_\Sigma(\lambda) + C_1.$$

Taking the infimum over $\eta > 0$ gives

$$\mathrm{h}_\Sigma(\lambda) + C_1 \leq \inf_{\eta>0} \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega(\sigma_n, \mathsf{J}^\eta(\mu_{\mathtt{A}}, \mu_{\mathtt{B}}))|.$$

The remainder of the proof is analogous to the proof of Theorem D.

By compactness, we can let $\mathcal{F} \subset \mathsf{J}^\eta(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$ be a finite set with $\mathsf{J}^\eta(\mu_{\mathtt{A}}, \mu_{\mathtt{B}}) \subset \bigcup_{\theta\in\mathcal{F}} \mathrm{B}^{\bar{d}}(\theta, \eta)$. Then

$$\begin{aligned}
\mathrm{h}_\Sigma(\lambda) + C_1 &\leq \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}\left|\Omega\left(\sigma_n, \bigcup_{\theta\in\mathcal{F}} \mathrm{B}(\theta, \eta)\right)\right| \\
&= \max_{\theta\in\mathcal{F}} \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega\left(\sigma_n, \mathrm{B}(\theta, \eta)\right)| \\
&\leq \sup_{\theta\in\mathsf{J}^\eta} \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega\left(\sigma_n, \mathrm{B}(\theta, \eta)\right)|.
\end{aligned}$$

Now for each $m \in \mathbb{N}$ take $\eta = 1/m$, and let $\theta_m \in \mathsf{J}^\eta(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$ get within $1/m$ of the supremum in the last line of the previous display. By compactness, we can pass to a weakly-convergent subsequence θ_{m_k} with limit θ_∞, which must lie in $\mathsf{J}(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$. Given $\mathcal{O} \ni \theta_\infty$, for m large enough we have $\mathrm{B}(\theta_m, 1/m) \subset \mathcal{O}$. Therefore

$$\begin{aligned}\mathrm{h}_\Sigma(\lambda) + C_1 &\leq \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega\left(\sigma_n, \mathrm{B}(\theta_m, 1/m)\right)| + \tfrac{1}{m} \\ &\leq \limsup_{n\to\infty} \frac{1}{|V_n|} \log \mathbb{E}|\Omega\left(\sigma_n, \mathcal{O}\right)| + \tfrac{1}{m}.\end{aligned}$$

Taking m to infinity then the infimum over $\mathcal{O}$ gives

$$\mathrm{h}_\Sigma(\lambda) + C_1 \leq \mathrm{h}_\Sigma(\theta_\infty).$$

Since $C_1 > 0$ and $\theta_\infty \in \mathsf{J}(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})$, this means that $\mathrm{h}_\Sigma(\lambda)$ is not maximal, unless every joining has $\mathrm{h}_\Sigma = -\infty$. □

In some cases, this allows us to say exactly which measure maximizes h_Σ:

Corollary 5.3.4. *Suppose* $\mu_{\mathtt{A}} \in \left[\mathrm{ex}\,\mathscr{G}(\Phi^{\mathtt{A}})\right] \cap \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ *and* $\mu_{\mathtt{B}} \in \mathscr{G}^\Gamma(\Phi^{\mathtt{B}})$. *Then*

$$\mathscr{G}(\Phi^{\mathtt{A}} \oplus \Phi^{\mathtt{B}}) \cap \mathsf{J}(\mu_{\mathtt{A}}, \mu_{\mathtt{B}}) = \{\mu_{\mathtt{A}} \times \mu_{\mathtt{B}}\}.$$

In particular,

$$\sup_{\lambda \in \mathsf{J}(\mu_{\mathtt{A}}, \mu_{\mathtt{B}})} \mathrm{h}_\Sigma(\lambda) = \mathrm{h}_\Sigma(\mu_{\mathtt{A}} \times \mu_{\mathtt{B}}).$$

Note that we require $\mu_{\mathtt{A}}$ to be an extreme point of the set of all Gibbs measures, not just the shift-invariant ones.

Proof. By [Geo11, Equation (7.19)], we have

$$\mathrm{ex}\,\mathscr{G}(\Phi) = \{\mu \times \nu \,:\, \mu \in \mathrm{ex}\,\mathscr{G}(\Phi^{\mathtt{A}}),\ \nu \in \mathrm{ex}\,\mathscr{G}(\Phi^{\mathtt{B}})\}.$$

Let λ be a joining of $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$ which is in $\mathscr{G}(\Phi)$, and write its extreme decomposition in $\mathscr{G}(\Phi)$ as

$$\lambda = \int \mu \times \nu \, \xi(d\mu, d\nu).$$

Then taking the marginal on $\mathtt{A}^\Gamma$ gives

$$\mu_{\mathtt{A}} = \int \mu \, \xi(d\mu, d\nu),$$

so extremality of $\mu_{\mathtt{A}}$ implies that ξ gives full mass to the set $\{(\mu_{\mathtt{A}}, \nu) : \nu \in \mathrm{ex}\,\mathscr{G}(\Phi^{\mathtt{B}})\}$. Therefore

$$\lambda = \mu_{\mathtt{A}} \times \int \nu\, \xi(d\mu, d\nu) = \mu_{\mathtt{A}} \times \mu_{\mathtt{B}}. \qquad \square$$

For example, at and above the reconstruction threshold, the free-boundary Ising Gibbs measure μ^{FB} is extreme [BRZ95, Iof96]. Therefore given any other fixed Gibbs measure (possibly for another nearest-neighbor potential and temperature), the product joining with μ^{FB} has maximal h_Σ.

We also note the following corollary:

Corollary 5.3.5. *If* $\mu \in \mathrm{Prob}^\Gamma(\mathtt{A}^\Gamma)$ *is a Gibbs measure and* $|\mathtt{A}| > 1$*, then for every deterministic sofic approximation* Σ *we have* $\mathrm{h}_\Sigma(\mu) \neq 0$*.*

Since for deterministic sofic approximations we always have $\mathrm{h}_\Sigma(\mu) \in \{-\infty\} \cup [0, +\infty)$ we could also write the conclusion as "either $\mathrm{h}_\Sigma(\mu) = -\infty$ or $\mathrm{h}_\Sigma(\mu) > 0$." Informally, we could then say that any deterministic sofic approximation either supports no good models for μ at all, or else the number of good models has a strictly positive (upper) exponential growth rate.

Proof. Suppose $\mathrm{h}_\Sigma(\mu) \neq -\infty$. Since the diagonal self-joining $\mu \bigtriangleup \mu$ is not Gibbs, Theorem E implies the existence of some other self-joining λ with $\mathrm{h}_\Sigma(\lambda) > \mathrm{h}_\Sigma(\mu \bigtriangleup \mu)$. But then

$$\mathrm{h}_\Sigma(\mu) = \mathrm{h}_\Sigma(\mu \bigtriangleup \mu) < \mathrm{h}_\Sigma(\lambda) \leq 2\mathrm{h}_\Sigma(\mu),$$

where the last inequality depends on Σ being deterministic. $\square$

5.4 Non-optimal Gibbs joinings

One might wonder whether the converse of Theorem E is true: does every joining of two Gibbs measures which is Gibbs for their sum interaction maximize entropy? In particular, Theorem D leads us to wonder when the product joining is non-maximal, since in this case the sofic entropy over the relevant stochastic block model is different from the f-invariant.

In this section we restrict to a particular random sofic approximation: Assume that Γ is the rank-r free group, and let $\sigma_n \in \operatorname{Hom}(\Gamma, \operatorname{Sym}([n]))$ be uniformly random. The paper [Bow10a] shows that h_Σ is the f-invariant introduced in [Bow10b]; see also the survey [Bow20a] for more information on the f-invariant.

A particularly useful property, not shared by all variants of sofic entropy, is additivity: $f(\mu \times \nu) = f(\mu) + f(\nu)$.

We also restrict to a particular class of Gibbs measures: the (free boundary conditions) Ising measure with transition probability $\varepsilon \in (0, 1/2]$ is the Γ-indexed, $\{-1, +1\}$-valued stationary Markov chain with uniform single-vertex marginals and transition matrix

$$\begin{pmatrix} 1-\varepsilon & \varepsilon \\ \varepsilon & 1-\varepsilon \end{pmatrix}.$$

We denote the distribution by $\mathsf{is}_\varepsilon \in \operatorname{Prob}^\Gamma(\{\pm 1\}^\Gamma)$. For each ε, the measure is_ε is Gibbs for the nearest-neighbor interaction with $h \equiv 0$ and $J(\mathsf{a}, \mathsf{b}) = -\beta \mathsf{ab}$, where the "inverse temperature" β is determined by the relation

$$\frac{\varepsilon}{1-\varepsilon} = \exp(-2\beta).$$

If ε is small then β is large, so we think of this as "low temperature." We can also think of is_ε as a model for broadcasting information, where we start with a uniformly random bit at the identity and transmit it across edges with error probability ε.

Since is_ε is a Markov chain, its f-invariant can be easily calculated. It is given by

$$f(\mathsf{is}_\varepsilon) = \log 2 + r(\mathrm{H}(\varepsilon) - \log 2) \tag{5.2}$$

where $\mathrm{H}(\varepsilon) = -[\varepsilon \log \varepsilon + (1-\varepsilon)\log(1-\varepsilon)]$ [Bow20a, Section 3.3]. In particular, $f(\mathsf{is}_\varepsilon) < 0$ for small enough ε. It is also not too difficult to show that if $\mathsf{is}_\varepsilon \bigtriangleup \mathsf{is}_\varepsilon$ is the diagonal self-joining then

$$f(\mathsf{is}_\varepsilon \triangle \mathsf{is}_\varepsilon) = f(\mathsf{is}_\varepsilon).$$

Therefore if $f(\mathsf{is}_\varepsilon) < 0$ then the product joining is not optimal, since $2f(\mathsf{is}_\varepsilon) < f(\mathsf{is}_\varepsilon)$. We can extend this to the case $f(\mathsf{is}_\varepsilon) = 0$, since Theorem E implies that the diagonal joining is non-optimal.

This already answers the question posed at the beginning of this section in the negative: the product joining is always Gibbs for the sum interaction, but is not maximal for small enough ε. In the rest of this section we extend further the range of ε where this is true.

To do this, we will use a result of [DMS17] to argue that, for some ε below the reconstruction threshold but above where the f-invariant is 0, the optimal Ising self-joining is not the product or the diagonal joining.

We first introduce some relevant terminology. For a finite graph $G = (V, E)$, a bisection is a partition $V = V_1 \sqcup V_2$ where $|V_1| = |V_2|$ if $|V|$ is even, or the sizes differ by 1 if $|V|$ is odd. The cut size of a bipartition $V = V_1 \sqcup V_2$ is the number of edges whose endpoints lie in different parts. The smallest cut size of any bisection of G is denoted $\mathsf{mcut}(G)$. For the graph of $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$, we will simply write $\mathsf{mcut}(\sigma)$.

The relevant result we will use is the following:

Theorem 5.4.1 (modification of [DMS17, Theorem 1.5])**.** *Let $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$ be chosen uniformly at random. Then as $|V| \to \infty$,*

$$\frac{\mathsf{mcut}(\sigma)}{|V| r} \xrightarrow{\mathbb{P}} \varepsilon_c = \frac{1}{2} - \mathsf{P}_* \frac{1}{\sqrt{2r}} + o_{r\to\infty}(r^{-1/2}),$$

where $\mathsf{P}_* \approx 0.7632$.

Here "$\xrightarrow{\mathbb{P}}$" denotes convergence in probability. Note that the existence of some related limits was established earlier in [BGT13], but the particular form of the asymptotic ε_c (found in [DMS17]) is useful here due to its comparability to ε_f (defined below).

The constant P_* is the limiting ground state energy density of the Sherrington-Kirkpatrick model; we will not need its precise definition here.

Proof of Theorem. Let $G^{\mathrm{reg}}(V,d)$ denote a d-regular graph with vertex set V, chosen uniformly at random (undefined unless $|V|d$ is even). Theorem 1.5 of [DMS17] states that

$$\frac{\mathsf{mcut}(G^{\mathrm{reg}}(V,2r))}{|V|r} \xrightarrow{\mathbb{P}} \varepsilon_c.$$

They actually prove the stronger result that this holds when $G^{\mathrm{reg}}(V,d)$ is a random multigraph chosen according to the configuration model. By the main theorems of [GJK02], the same holds with $G^{\mathrm{reg}}(V,d)$ replaced by a uniformly random $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$. □

The main result of this section is the following:

Theorem H. *If $\varepsilon < \varepsilon_c$ then the product self-joining of is_ε has non-maximal f-invariant.*

Let $\varepsilon_f < 1/2$ be the smaller solution to $f(\mathsf{is}_\varepsilon) = 0$. If $\varepsilon \leq \varepsilon_f$ then $f(\mathsf{is}_\varepsilon) \leq 0$; we have remarked above that this implies non-optimality of the product joining. A Taylor expansion of H yields from Equation 5.2

$$\varepsilon_f = \frac{1}{2} - \sqrt{\log 2}\frac{1}{\sqrt{2r}} + o_{r\to\infty}(r^{-1/2}).$$

Since $\sqrt{\log 2} \approx 0.8326 > 0.7632 \approx \mathsf{P}_*$,

$$\varepsilon_f < \varepsilon_c \quad \text{for all large } r.$$

Therefore this theorem does, in fact, extend the range of non-maximality of the product (for large enough r).

The connection between the Ising model and mcut is that if a graph G admits a good model for is_ε, then $\mathsf{mcut}(G)$ must not be much bigger than $\varepsilon|V|r$: since the single-vertex marginal of is_ε is uniform, this good model must approximately bisect V, and since the transition probability is ε, the cut size of the corresponding partition must be approximately $\varepsilon|V|r$ (since $|V|r$ is the total number of edges). More precisely:

Lemma 5.4.2. *For every $\delta > 0$ there exists a neighborhood $\mathcal{O} \ni \mathsf{is}_\varepsilon$ such that for every $\sigma \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(V))$*

$$\Omega(\sigma, \mathcal{O}) \neq \varnothing \quad \Rightarrow \quad \mathsf{mcut}(\sigma) < |V|r(\varepsilon + \delta).$$

Proof. Let $\rho = \frac{\delta}{2r+1}$, and let $\mathcal{O}$ be the set of $\nu \in \mathrm{Prob}^\Gamma(\{\pm1\}^\Gamma)$ whose marginal on $\mathrm{B}^\Gamma(e,1)$ is within total variation distance ρ of the same marginal of is_ε.

Suppose we have $\mathbf{x} \in \Omega(\sigma, \mathcal{O})$. Then

$$\left| \tfrac{1}{|V|} |\{v \in V \,:\, \mathbf{x}(v) = +1\}| - \frac{1}{2} \right| < \rho$$

so we can pick $\mathbf{y} \in \{\pm1\}^V$ with

$$\left| |\{v \in V \,:\, \mathbf{y}(v) = +1\}| - \frac{|V|}{2} \right| \leq 1 \quad \text{and} \quad \frac{1}{|V|} |\{v \in V \,:\, \mathbf{y}(v) \neq \mathbf{x}(v)\}| < \rho.$$

Now $\mathbf{y}$ induces a bisection of V, and

$$\begin{aligned} \left| \frac{1}{|V|r} \sum_{v \in V} \sum_{i \in r} \mathbf{1}_{\{\mathbf{y}(v) \neq \mathbf{y}(\sigma^i v)\}} - \varepsilon \right| &\leq \left| \frac{1}{|V|r} \sum_{v \in V} \sum_{i \in r} \mathbf{1}_{\{\mathbf{x}(v) \neq \mathbf{x}(\sigma^i v)\}} - \varepsilon \right| \\ &\quad + \frac{1}{|V|r} \sum_{v \in V} \sum_{i \in r} \left| \mathbf{1}_{\{\mathbf{y}(v) \neq \mathbf{y}(\sigma^i v)\}} - \mathbf{1}_{\{\mathbf{x}(v) \neq \mathbf{x}(\sigma^i v)\}} \right|. \end{aligned}$$

The first term is at most ρ by definition of $\mathcal{O}$: to see this, write

$$\frac{1}{|V|r} \sum_{v \in V} \sum_{i \in r} \mathbf{1}_{\{\mathbf{x}(v) \neq \mathbf{x}(\sigma^i v)\}} = \int \frac{1}{r} \sum_{i \in [r]} \mathbf{1}_{\{\mathbf{z}(e) \neq \mathbf{z}(s_i)\}} \, P^\sigma_{\mathbf{x}}(d\mathbf{z}).$$

To bound the second term, write

$$\begin{aligned} &\sum_{v \in V} \sum_{i \in r} \left| \mathbf{1}_{\{\mathbf{y}(v) \neq \mathbf{y}(\sigma^i v)\}} - \mathbf{1}_{\{\mathbf{x}(v) \neq \mathbf{x}(\sigma^i v)\}} \right| \\ &\qquad \leq \sum_{v \in V} \sum_{i \in r} \left[\mathbf{1}_{\{\mathbf{y}(v) \neq \mathbf{x}(v)\}} + \mathbf{1}_{\{\mathbf{y}(\sigma^i v) \neq \mathbf{x}(\sigma^i v)\}} \right] \\ &\qquad = 2r \sum_{v \in V} \mathbf{1}_{\{\mathbf{y}(v) \neq \mathbf{x}(v)\}} \leq 2r|V|\rho. \end{aligned}$$

Therefore the cut size of the bisection induced by $\mathbf{y}$ is at most

$$|V|r\varepsilon + |V|r\rho + 2r|V|\rho = |V|r(\varepsilon + \delta).$$

□

5.4.1 Proof of Theorem H

Non-optimality of the product joining for $\varepsilon < \varepsilon_c$ follows from the next two lemmas.

Lemma 5.4.3. *Suppose that* $\mathsf{is}_\varepsilon \times \mathsf{is}_\varepsilon$ *has maximal* f *among all self-joinings of* is_ε*. Then for any* $\delta > 0$

$$\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\tfrac{\mathsf{mcut}(\sigma_n)}{rn} < \varepsilon + \delta\right) \geq 0.$$

Proof. A standard argument shows that

$$\inf_{\mathcal{O}\ni \mathsf{is}_\varepsilon} \limsup_{n\to\infty} \frac{1}{n} \log \mathbb{E}\left[|\Omega(\sigma_n, \mathcal{O})|^2\right] = \sup_{\lambda \in \mathsf{J}(\mathsf{is}_\varepsilon, \mathsf{is}_\varepsilon)} f(\lambda) = 2f(\mathsf{is}_\varepsilon),$$

where the second equality uses our assumption that the product joining is optimal. Therefore for any η, for all small enough $\mathcal{O}$ we have

$$\mathbb{E}\left[|\Omega(\sigma_n, \mathcal{O})|^2\right] < \exp\left[n(2f(\mathsf{is}_\varepsilon) + \eta)\right]$$

for all large enough n. Similarly, since $f(\mathsf{is}_\varepsilon) = \inf_{\mathcal{O}\ni \mathsf{is}_\varepsilon} \limsup_{n\to\infty} \frac{1}{n} \log \mathbb{E}|\Omega(\sigma_n, \mathcal{O})|$, for any $\mathcal{O} \ni \mathsf{is}_\varepsilon$ we have

$$\mathbb{E}|\Omega(\sigma_n, \mathcal{O})| > \exp\left[n(f(\mathsf{is}_\varepsilon) - \eta)\right]$$

for infinitely many n.

By Lemma 5.4.2, for all small enough $\mathcal{O} \ni \mathsf{is}_\varepsilon$ we have

$$\mathbb{P}(\tfrac{\mathsf{mcut}(\sigma_n)}{rn} < \varepsilon + \delta) \leq \mathbb{P}(\Omega(\sigma_n, \mathcal{O}) \neq \varnothing).$$

Using the Paley-Zygmund inequality,

$$\begin{aligned} \mathbb{P}(\Omega(\sigma_n, \mathcal{O}) \neq \varnothing) &\geq \mathbb{P}\left(|\Omega(\sigma_n, \mathcal{O})| > \tfrac{1}{2}\, \mathbb{E}|\Omega(\sigma_n, \mathcal{O})|\right) \\ &\geq (1 - \tfrac{1}{2})^2 \frac{\left[\mathbb{E}\, |\Omega(\sigma_n, \mathcal{O})|\right]^2}{\mathbb{E}\left[|\Omega(\sigma_n, \mathcal{O})|^2\right]} \\ &> \tfrac{1}{4} \exp\left[-2\eta n\right] \end{aligned}$$

for infinitely many n. Hence

$$\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}\left(\tfrac{\mathsf{mcut}(\sigma_n)}{rn} < \varepsilon + \delta\right) > -2\eta$$

and, since $\eta > 0$ is arbitrary, the result follows. ∎

Lemma 5.4.4. *If $\varepsilon < \varepsilon_c$ then for all small enough $\delta > 0$*

$$\liminf_{n\to\infty} \frac{1}{n} \log \mathbb{P}(\tfrac{\mathsf{mcut}(\sigma_n)}{rn} < \varepsilon + \delta) < 0.$$

Proof. Theorem 5.4.1 implies that $\lim_{n\to\infty} \mathbb{E}\, \frac{\mathsf{mcut}(\sigma_n)}{rn} = \varepsilon_c$, so if δ is small enough that $\varepsilon+\delta < \varepsilon_c$ then for any $0 < t < \varepsilon_c - (\varepsilon+\delta)$

$$\begin{aligned}
\mathbb{P}(\tfrac{\mathsf{mcut}(\sigma_n)}{rn} < \varepsilon + \delta) &= \mathbb{P}(\tfrac{\mathsf{mcut}(\sigma_n)}{rn} - \mathbb{E}\,\tfrac{\mathsf{mcut}(\sigma_n)}{rn} < \varepsilon + \delta - \mathbb{E}\,\tfrac{\mathsf{mcut}(\sigma_n)}{rn}) \\
&\le \mathbb{P}(\tfrac{\mathsf{mcut}(\sigma_n)}{rn} - \mathbb{E}\,\tfrac{\mathsf{mcut}(\sigma_n)}{rn} < -t) \qquad \text{(for all large } n\text{)} \\
&\le \mathbb{P}(\left|\tfrac{\mathsf{mcut}(\sigma_n)}{rn} - \mathbb{E}\,\tfrac{\mathsf{mcut}(\sigma_n)}{rn}\right| \ge t).
\end{aligned}$$

By a standard "switching" argument (Lemma 5.4.5), we have

$$\mathbb{P}(|\tfrac{\mathsf{mcut}(\sigma_n)}{rn} - \mathbb{E}\,\tfrac{\mathsf{mcut}(\sigma_n)}{rn}| \ge t) \le 2\exp(-t^2 nr/8) \quad \forall t > 0.$$

The result follows. □

5.4.2 Concentration of random homomorphisms

Here we develop an analogue of [Wor99, Theorem 2.19], which proves exponential concentration for functions which are not changed much under "switching." Similar concentration techniques also appear in the survey [McD98].

Given $\tau_1, \tau_2 \in \mathrm{Sym}(n)$, we write $\tau_1 \sim \tau_2$ if

$$|\{j \in [n] \,:\, \tau_1(j) \ne \tau_2(j)\}| = 2.$$

Note that 2 is the smallest positive number of disagreements between two permutations. If $\tau_1 \sim \tau_2$ and $i, j \in [n]$ are the points where they disagree, then it must be that $\tau_1(i) = \tau_2(j)$ and $\tau_2(i) = \tau_1(j)$. For this reason we say they differ by a *switching*.

We extend this to homomorphisms $\sigma_1, \sigma_2 \colon \mathbb{F}_r \to \mathrm{Sym}(n)$ by saying $\sigma_1 \sim \sigma_2$ whenever there is exactly one $i_0 \in [r]$ with $\sigma_1^{i_0} \sim \sigma_2^{i_0}$ and for all $i \ne i_0$ we have $\sigma_1^i = \sigma_2^i$.

If $\sigma_1 \sim \sigma_2$ then $|\mathsf{mcut}(\sigma_1) - \mathsf{mcut}(\sigma_2)| \le 2$. The following lemma establishes concentration for functions with this property.

Lemma 5.4.5. *Suppose g is a real-valued function on* $\mathrm{Hom}(\mathbb{F}_r, \mathrm{Sym}(n))$ *such that* $|g(\sigma_1) - g(\sigma_2)| \leq c$ *whenever* $\sigma_1 \sim \sigma_2$. *Then if σ is chosen uniformly at random*

$$\mathbb{P}\left(|g(\sigma) - \mathbb{E}\, g(\sigma)| > t\right) \leq 2\exp\left(\frac{-t^2}{2nrc^2}\right).$$

Proof. We choose σ by picking $\sigma^i(j)$ in lexicographic order on $(i,j) \in [r] \times [n]$ uniformly from all allowable choices. Let

$$\{\varnothing, \mathrm{Hom}(\mathbb{F}_r, \mathrm{Sym}(n))\} = \mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots \subseteq \mathcal{F}_{nr} = \mathcal{P}(\mathrm{Hom}(\mathbb{F}_r, \mathrm{Sym}(n)))$$

be the filtration induced by these choices. If we show that

$$|\mathbb{E}[g(\sigma) \mid \mathcal{F}_k] - \mathbb{E}[g(\sigma) \mid \mathcal{F}_{k-1}]| \leq c \quad \text{for all } k,$$

then the result will follow from Azuma-Hoeffding.

Fix $k = i_0 r + j_0 \in [nr]$, so that $\mathcal{F}_k$ records the choice of $\sigma^{i_0}(j_0)$ and all previous choices. It is helpful to think of, for $\sigma_0 \in \mathrm{Hom}(\Gamma, \mathrm{Sym}(n))$,

$$\mathbb{E}[g(\sigma) \mid \mathcal{F}_k](\sigma_0) = \mathbb{E}[g(\sigma) \mid \sigma^i(j) = \sigma_0^i(j)\ \forall (i,j) \leq (i_0, j_0)]$$
$$\mathbb{E}[g(\sigma) \mid \mathcal{F}_{k-1}](\sigma_0) = \mathbb{E}[g(\sigma) \mid \sigma^i(j) = \sigma_0^i(j)\ \forall (i,j) < (i_0, j_0)].$$

We need to show that the difference between these two quantities is bounded by c for each fixed σ_0.

Let $A \subset [n]$ be the set of allowed values for $\sigma^i(j)$ given the event $U := \{\sigma^i(j) = \sigma_0^i(j)\ \forall (i,j) < (i_0, j_0)\}$. For each $a \in A$ let $U_a = U \cap \{\sigma^i(j) = a\}$. Note that each U_a has the same probability, namely $\frac{1}{|A|}\mathbb{P}(U)$. For convenience write $a_0 = \sigma_0^i(j)$. Then we can rewrite the above quantities as

$$\mathbb{E}[g(\sigma) \mid U_{a_0}] \quad \text{and} \quad \mathbb{E}[g(\sigma) \mid U].$$

Then

$$\begin{aligned} |\mathbb{E}[g(\sigma) \mid U_{a_0}] - \mathbb{E}[g(\sigma) \mid U]| &= \left|\mathbb{E}[g(\sigma) \mid U_{a_0}] - \frac{1}{|A|}\sum_{a \in A} \mathbb{E}[g(\sigma) \mid U_a]\right| \\ &\leq \frac{1}{|A|}\sum_{a \in A}\Big|\mathbb{E}[g(\sigma) \mid U_{a_0}] - \mathbb{E}[g(\sigma) \mid U_a]\Big| \end{aligned}$$

For $\sigma \in U$ and $a \in A$, let $S_a\sigma$ denote the unique switching of σ with $(S_a\sigma)^i(j) = a$ (or take $S_a\sigma = \sigma$ if $\sigma \in U_a$ already). Note that $\sigma \in U$ implies $S_a\sigma \in U_a$. Moreover, if $\sigma \sim \mathrm{Unif}(U_{a_0})$ then $S_a\sigma \sim \mathrm{Unif}(U_a)$ (since S_a is a bijection). Therefore

$$\begin{aligned}
&\Big|\mathbb{E}[g(\sigma) \mid U_{a_0}] - \mathbb{E}[g(\sigma) \mid U_a]\Big| \\
&\qquad\qquad \leq \Big|\mathbb{E}[g(S_a\sigma) \mid U_{a_0}] - \mathbb{E}[g(\sigma) \mid U_{a_0}]\Big| + \Big|\mathbb{E}[g(S_a\sigma) \mid U_{a_0}] - \mathbb{E}[g(\sigma) \mid U_a]\Big| \\
&\qquad\qquad \leq c + 0,
\end{aligned}$$

so the result follows. □

5.5 Proof of Proposition 5.3.1

First fix $E_{\mathtt{A}} \subset \mathtt{A}^\Gamma$ and $E_{\mathtt{B}} \subset \mathtt{B}^\Gamma$, and let $(\Lambda_n)_{n\in\mathbb{N}}$ be an increasing sequence of finite subsets of Γ with $\bigcup_{n\in\mathbb{N}} \Lambda_n = \Gamma$. It suffices to assume that $E_{\mathtt{A}}, E_{\mathtt{B}}$ are cylinder sets, *i.e.* each depends on only finitely many coordinates.

Extending the convention used above when defining Gibbs measures (and following [Geo11]), for a finite set $\Lambda \subset \Gamma$ we let $\mathscr{T}^{\mathtt{AB}}_\Lambda$ be the sub-σ-algebra of the Borel σ-algebra on $(\mathtt{A} \times \mathtt{B})^\Gamma$ generated by the coordinate maps for $v \in \Gamma \setminus \Lambda$. We similarly define $\mathscr{T}^{\mathtt{A}}_\Lambda, \mathscr{T}^{\mathtt{B}}_\Lambda$.

Then the tail σ-algebra on $(\mathtt{A} \times \mathtt{B})^\Gamma$ is $\mathscr{T}^{\mathtt{AB}} = \bigcap_{n\in\mathbb{N}} \mathscr{T}^{\mathtt{AB}}_{\Lambda_n}$, so by backwards martingale convergence we have

$$\lambda\big(E_{\mathtt{A}} \times E_{\mathtt{B}} \big| \mathscr{T}^{\mathtt{AB}}\big)(\mathbf{x}, \mathbf{y}) = \lim_{n\to\infty} \lambda\big(E_{\mathtt{A}} \times E_{\mathtt{B}} \big| \mathscr{T}^{\mathtt{AB}}_{\Lambda_n}\big)(\mathbf{x}, \mathbf{y})$$

for λ-a.e. $(\mathbf{x}, \mathbf{y})$.

Lemma 5.5.1. *If $E_{\mathtt{A}}, E_{\mathtt{B}} \subset \mathtt{A}^\Gamma, \mathtt{B}^\Gamma$ are cylinder sets and n is large enough that both depend only on coordinates in Λ_n, then*

$$\lambda\big(E_{\mathtt{A}} \times E_{\mathtt{B}} \big| \mathscr{T}^{\mathtt{AB}}_{\Lambda_n}\big)(\mathbf{x}, \mathbf{y}) = \mu_{\mathtt{A}}\big(E_{\mathtt{A}} \big| \mathscr{T}^{\mathtt{A}}_{\Lambda_n}\big)(\mathbf{x}) \cdot \mu_{\mathtt{B}}\big(E_{\mathtt{B}} \big| \mathscr{T}^{\mathtt{B}}_{\Lambda_n}\big)(\mathbf{y})$$

for λ-a.e. $(\mathbf{x}, \mathbf{y}) \in (\mathtt{A} \times \mathtt{B})^\Gamma$.

The assumption on n may not be necessary, but it is convenient and sufficient for our purposes.

Proof. If $m > n$ then we write $\mathscr{F}^{\mathtt{AB}}_{m\setminus n}$ for the σ-algebra generated by coordinates in $\Lambda_m \setminus \Lambda_n$. For $\mathbf{w} \in \mathtt{A}^{\Lambda_n}$ we will write $\mathbf{w} \in E_{\mathtt{A}}$ to mean that any extension of $\mathbf{w}$ to an element of $\mathtt{A}^{\Gamma}$ would be in $E_{\mathtt{A}}$; this makes sense because we assume that $E_{\mathtt{A}}$ is only determined by coordinates in Λ_n. For $(\mathbf{x}, \mathbf{y}) \in (\mathtt{A} \times \mathtt{B})^{\Gamma}$ we write $\lambda_m\{(\mathbf{x}, \mathbf{y})\}$ for the λ-measure of all labelings which agree with $(\mathbf{x}, \mathbf{y})$ on Λ_m.

The first step is to show that

$$\lambda\big(E_{\mathtt{A}} \times E_{\mathtt{B}}\big|\mathscr{F}^{\mathtt{AB}}_{m\setminus n}\big)(\mathbf{x}, \mathbf{y}) = \frac{\displaystyle\sum_{\substack{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\, \mathbf{z}\in\mathtt{B}^{\Lambda_n}\\ \mathbf{w}\in E_{\mathtt{A}},\, \mathbf{z}\in E_{\mathtt{B}}}} \lambda_m\big\{\big(\mathbf{x}^{\Lambda_n\to\mathbf{w}}, \mathbf{y}^{\Lambda_n\to\mathbf{z}}\big)\big\}}{\displaystyle\sum_{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\mathbf{z}\in\mathtt{B}^{\Lambda_n}} \lambda_m\big\{\big(\mathbf{x}^{\Lambda_n\to\mathbf{w}}, \mathbf{y}^{\Lambda_n\to\mathbf{z}}\big)\big\}}.$$

It is clear that this is $\mathscr{F}^{\mathtt{AB}}_{m\setminus n}$-measurable. To show that it is a version of the conditional expectation, let $F \in \mathscr{F}^{\mathtt{AB}}_{m\setminus n}$. Pick $(\mathbf{u}_1, \mathbf{v}_1), \ldots, (\mathbf{u}_k, \mathbf{v}_k) \in (\mathtt{A} \times \mathtt{B})^{\Lambda_m\setminus\Lambda_n}$ such that a labeling is in F if and only if its restriction to $\Lambda_m \setminus \Lambda_n$ is one of those k options. Then

$$\begin{aligned}
&\int_F \left[\frac{\displaystyle\sum_{\substack{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\, \mathbf{z}\in\mathtt{B}^{\Lambda_n}\\ \mathbf{w}\in E_{\mathtt{A}},\, \mathbf{z}\in E_{\mathtt{B}}}} \lambda_m\big\{\big(\mathbf{x}^{\Lambda_n\to\mathbf{w}}, \mathbf{y}^{\Lambda_n\to\mathbf{z}}\big)\big\}}{\displaystyle\sum_{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\mathbf{z}\in\mathtt{B}^{\Lambda_n}} \lambda_m\big\{\big(\mathbf{x}^{\Lambda_n\to\mathbf{w}}, \mathbf{y}^{\Lambda_n\to\mathbf{z}}\big)\big\}}\right] d\lambda(\mathbf{x}, \mathbf{y}) \\
&\quad= \sum_{i=1}^{k} \sum_{\mathbf{w}'\in\mathtt{A}^{\Lambda_n},\mathbf{z}'\in\mathtt{B}^{\Lambda_n}} \lambda_m\big\{\big(\mathbf{u}_i\mathbf{w}', \mathbf{v}_i\mathbf{z}'\big)\big\} \frac{\displaystyle\sum_{\substack{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\, \mathbf{z}\in\mathtt{B}^{\Lambda_n}\\ \mathbf{w}\in E_{\mathtt{A}},\, \mathbf{z}\in E_{\mathtt{B}}}} \lambda_m\big\{\big(\mathbf{u}_i\mathbf{w}, \mathbf{v}_i\mathbf{z}\big)\big\}}{\displaystyle\sum_{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\mathbf{z}\in\mathtt{B}^{\Lambda_n}} \lambda_m\big\{\big(\mathbf{u}_i\mathbf{w}, \mathbf{v}_i\mathbf{z}\big)\big\}} \\
&\quad= \sum_{i=1}^{k} \sum_{\substack{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\, \mathbf{z}\in\mathtt{B}^{\Lambda_n}\\ \mathbf{w}\in E_{\mathtt{A}},\, \mathbf{z}\in E_{\mathtt{B}}}} \lambda_m\big\{\big(\mathbf{u}_i\mathbf{w}, \mathbf{v}_i\mathbf{z}\big)\big\} \\
&\quad= \sum_{i=1}^{k} \sum_{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\, \mathbf{z}\in\mathtt{B}^{\Lambda_n}} \mathbf{1}_{\{\mathbf{w}\in E_{\mathtt{A}},\, \mathbf{z}\in E_{\mathtt{B}}\}} \lambda_m\big\{\big(\mathbf{u}_i\mathbf{w}, \mathbf{v}_i\mathbf{z}\big)\big\} \\
&\quad= \int_F \mathbf{1}_{\{\mathbf{x}\in E_{\mathtt{A}},\, \mathbf{y}\in E_{\mathtt{B}}\}} d\lambda(\mathbf{x}, \mathbf{y}).
\end{aligned}$$

Now let $v \in \Lambda_n$. Let $\lambda_{m\setminus v}$ be interpreted like λ_m but looking at agreement on $\Lambda_m \setminus \{v\}$. Then by the (single-site) Gibbs property and product structure

$$\frac{\lambda_m\big\{\big(\mathbf{x},\mathbf{y}\big)\big\}}{\sum_{\substack{\mathtt{a}\in\mathtt{A}\\ \mathtt{b}\in\mathtt{B}}}\lambda_m\big\{\big(\mathbf{x}^{v\to\mathtt{a}},\mathbf{y}^{v\to\mathtt{b}}\big)\big\}} = \frac{\lambda_{m\setminus v}\big\{\big(\mathbf{x},\mathbf{y}\big)\big\}\, c_v(\mathbf{x}\mathbf{y},\mathbf{x}(v)\mathbf{y}(v))}{\sum_{\substack{\mathtt{a}\in\mathtt{A}\\ \mathtt{b}\in\mathtt{B}}}\lambda_{m\setminus v}\big\{\big(\mathbf{x},\mathbf{y}\big)\big\}c_v(\mathbf{x}\mathbf{y},\mathtt{ab})}$$
$$= c_v^{\mathtt{A}}(\mathbf{x},\mathbf{x}(v))c_v^{\mathtt{B}}(\mathbf{y},\mathbf{x}(v)).$$

Similarly

$$\frac{\mu_{\mathtt{A},m}\{\mathbf{x}\}}{\sum_{\mathtt{a}\in\mathtt{A}}\mu_{\mathtt{A},m}\{\mathbf{x}^{v\to\mathtt{a}}\}} = c_v^{\mathtt{A}}(\mathbf{x},\mathbf{x}(v))$$

and similarly for $\mu_{\mathtt{B}}$. Given $\big(\mathbf{x},\mathbf{y}\big)$ and $\big(\mathbf{x}',\mathbf{y}'\big)$ which agree on $\Lambda_m \setminus \Lambda_n$, let

$$\big(\mathbf{x},\mathbf{y}\big) = \big(\mathbf{x}_0,\mathbf{y}_0\big),\ldots,\big(\mathbf{x}_l,\mathbf{y}_l\big) = \big(\mathbf{x}',\mathbf{y}'\big)$$

be such that for each $i = 0,\ldots,l-1$ the labelings $\big(\mathbf{x}_i,\mathbf{y}_i\big)$, $\big(\mathbf{x}_{i+1},\mathbf{y}_{i+1}\big)$ differ at exactly one site $v_i \in \Lambda_n$, with the former agreeing with $(\mathbf{x},\mathbf{y})$ at that site and the latter with $(\mathbf{x}',\mathbf{y}')$. Then

$$\begin{aligned}
\frac{\lambda_m\big\{\big(\mathbf{x},\mathbf{y}\big)\big\}}{\lambda_m\big\{\big(\mathbf{x}',\mathbf{y}'\big)\big\}} &= \prod_{i=0}^{l-1}\frac{\lambda_m\big\{\big(\mathbf{x}_i,\mathbf{y}_i\big)\big\}}{\lambda_m\big\{\big(\mathbf{x}_{i+1},\mathbf{y}_{i+1}\big)\big\}}\\
&= \prod_{i=0}^{l-1}\frac{c_{v_i}^{\mathtt{A}}(\mathbf{x}_i,\mathbf{x}_i(v_i))c_{v_i}^{\mathtt{B}}(\mathbf{y}_i,\mathbf{y}_i(v_i))}{c_{v_i}^{\mathtt{A}}(\mathbf{x}_{i+1},\mathbf{x}_{i+1}(v_i))c_{v_i}^{\mathtt{B}}(\mathbf{y}_{i+1},\mathbf{y}_{i+1}(v_i))}\\
&= \prod_{i=0}^{l-1}\frac{c_{v_i}^{\mathtt{A}}(\mathbf{x}_i,\mathbf{x}_i(v_i))}{c_{v_i}^{\mathtt{A}}(\mathbf{x}_{i+1},\mathbf{x}_{i+1}(v_i))}\prod_{i=0}^{l-1}\frac{c_{v_i}^{\mathtt{B}}(\mathbf{y}_i,\mathbf{y}_i(v_i))}{c_{v_i}^{\mathtt{B}}(\mathbf{y}_{i+1},\mathbf{y}_{i+1}(v_i))}\\
&= \frac{\mu_{\mathtt{A},m}\{\mathbf{x}\}}{\mu_{\mathtt{A},m}\{\mathbf{x}'\}}\frac{\mu_{\mathtt{B},m}\{\mathbf{y}\}}{\mu_{\mathtt{B},m}\{\mathbf{y}'\}}.
\end{aligned}$$

So if we fix $\big(\mathbf{x}',\mathbf{y}'\big)$ then for any $\big(\mathbf{x},\mathbf{y}\big)$ agreeing with it on $\Lambda_m \setminus \Lambda_n$ we have

$$\lambda_m\big\{\big(\mathbf{x},\mathbf{y}\big)\big\} = C_m\mu_{\mathtt{A},m}\{\mathbf{x}\}\mu_{\mathtt{B},m}\{\mathbf{y}\}$$

where

$$C_m = \frac{\lambda_m\big\{\big(\mathbf{x}',\mathbf{y}'\big)\big\}}{\mu_{\mathtt{A},m}\{\mathbf{x}'\}\mu_{\mathtt{B},m}\{\mathbf{y}'\}}.$$

Returning to the first result of the present proof, this gives

$$\lambda\big(E_{\mathtt{A}} \times E_{\mathtt{B}} \big| \mathscr{F}^{\mathtt{AB}}_{m\setminus n}\big)(\mathbf{x},\mathbf{y}) = \frac{\displaystyle\sum_{\substack{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\,\mathbf{z}\in\mathtt{B}^{\Lambda_n}\\ \mathbf{w}\in E_{\mathtt{A}},\,\mathbf{z}\in E_{\mathtt{B}}}} \lambda_m\big\{\big(\mathbf{x}^{\Lambda_n\to\mathbf{w}},\mathbf{y}^{\Lambda_n\to\mathbf{z}}\big)\big\}}{\displaystyle\sum_{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\mathbf{z}\in\mathtt{B}^{\Lambda_n}} \lambda_m\big\{\big(\mathbf{x}^{\Lambda_n\to\mathbf{w}},\mathbf{y}^{\Lambda_n\to\mathbf{z}}\big)\big\}}$$

$$= \frac{\displaystyle\sum_{\substack{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\,\mathbf{z}\in\mathtt{B}^{\Lambda_n}\\ \mathbf{w}\in E_{\mathtt{A}},\,\mathbf{z}\in E_{\mathtt{B}}}} C_m\mu_{\mathtt{A},m}\big\{\mathbf{x}^{\Lambda_n\to\mathbf{w}}\big\}\mu_{\mathtt{B},m}\big\{\mathbf{y}^{\Lambda_n\to\mathbf{z}}\big\}}{\displaystyle\sum_{\mathbf{w}\in\mathtt{A}^{\Lambda_n},\mathbf{z}\in\mathtt{B}^{\Lambda_n}} C_m\mu_{\mathtt{A},m}\big\{\mathbf{x}^{\Lambda_n\to\mathbf{w}}\big\}\mu_{\mathtt{B},m}\big\{\mathbf{y}^{\Lambda_n\to\mathbf{z}}\big\}}$$

$$= \frac{\displaystyle\sum_{\substack{\mathbf{w}\in\mathtt{A}^{\Lambda_n}\\ \mathbf{w}\in E_{\mathtt{A}}}} \mu_{\mathtt{A},m}\big\{\mathbf{x}^{\Lambda_n\to\mathbf{w}}\big\}\sum_{\substack{\mathbf{z}\in\mathtt{B}^{\Lambda_n}\\ \mathbf{z}\in E_{\mathtt{B}}}} \mu_{\mathtt{B},m}\big\{\mathbf{y}^{\Lambda_n\to\mathbf{z}}\big\}}{\displaystyle\sum_{\mathbf{w}\in\mathtt{A}^{\Lambda_n}} \mu_{\mathtt{A},m}\big\{\mathbf{x}^{\Lambda_n\to\mathbf{w}}\big\}\sum_{\mathbf{z}\in\mathtt{B}^{\Lambda_n}} \mu_{\mathtt{B},m}\big\{\mathbf{y}^{\Lambda_n\to\mathbf{z}}\big\}}$$

$$= \mu_{\mathtt{A}}\big(E_{\mathtt{A}}\big|\mathscr{F}^{\mathtt{A}}_{m\setminus n}\big)(\mathbf{x}) \cdot \mu_{\mathtt{B}}\big(E_{\mathtt{B}}\big|\mathscr{F}^{\mathtt{B}}_{m\setminus n}\big)(\mathbf{y}),$$

where the last line follows from the same argument used above on the conditional expectation of λ.

Taking m to infinity gives the result, by upwards martingale convergence; on the right-hand side we have convergence for $\mu_{\mathtt{A}}$-a.e. $\mathbf{x}$ and $\mu_{\mathtt{B}}$-a.e. $\mathbf{y}$, which is λ-a.e. $(\mathbf{x},\mathbf{y})$ because λ has marginals $\mu_{\mathtt{A}}, \mu_{\mathtt{B}}$. □

Throughout the thesis we have used a 'single-site' definition of Gibbs states which is a slight generalization of [Lig05, Definition IV.1.5]. It is comforting to know, and sometimes useful above, that in our setting this is equivalent to [Geo11, Definition 2.9] which requires all finite-dimensional conditional distributions to be specified by the potential. The following proposition establishes this using some work done in the previous proof.

Proposition 5.5.2. *If $\mu_{\mathtt{A}} \in \mathscr{G}(\Phi)$ as defined in Chapter 3, then all finite-dimensional conditional expectations are specified by Φ as in the definition of Gibbs states in [Geo11].*

Proof. We begin by recalling from the proof above that if m is large enough that Λ_m contains

Λ, then for any $\mathbf{x}, \mathbf{x}'$ which agree on Λ

$$\frac{\mu_{\mathtt{A},m}\{\mathbf{x}\}}{\mu_{\mathtt{A},m}\{\mathbf{x}'\}} = \prod_{i=0}^{l-1} \frac{c_{v_i}^{\mathtt{A}}(\mathbf{x}_i, \mathbf{x}_i(v_i))}{c_{v_i}^{\mathtt{A}}(\mathbf{x}_{i+1}, \mathbf{x}_{i+1}(v_i))}$$

where $\mathbf{x}_0, \ldots \mathbf{x}_l \in \mathtt{A}^\Gamma$ and $v_0, \ldots, v_l \in \Lambda$ are chosen as above, except now we add the assumption that $\Lambda = \{v_0, \ldots, v_l\}$. The normalizing factors from the $c_{v_i}^{\mathtt{A}}$ terms in the the same factor of the product cancel (since they are independent of the label at v_i), so using the definitions from Chapter 3 we can write this as

$$\begin{aligned}\frac{\mu_{\mathtt{A},m}\{\mathbf{x}\}}{\mu_{\mathtt{A},m}\{\mathbf{x}'\}} &= \prod_{i=0}^{l-1} \frac{\exp\big(-\Phi_{v_i}(\mathbf{x}_i)\big)}{\exp\big(-\Phi_{v_i}(\mathbf{x}_{i+1})\big)} \\ &= \prod_{i=0}^{l-1} \frac{\exp\big(-h(\mathbf{x}_i(v_i)) - \sum_{s\in S} J(\mathbf{x}_i(v_i), \mathbf{x}_i(\sigma^s v_i)\big)}{\exp\big(-h(\mathbf{x}_{i+1}(v_i)) - \sum_{s\in S} J(\mathbf{x}_{i+1}(v_i), \mathbf{x}_{i+1}(\sigma^s v_i)\big)}.\end{aligned}$$

Using the definition of the $\mathbf{x}_i$'s and v_i's, and separating the h and J terms, we can write this as

$$\prod_{i=0}^{l-1} \frac{\exp\big(-h(\mathbf{x}(v_i))\big)}{\exp\big(-h(\mathbf{x}'(v_i))\big)} \prod_{i=0}^{l-1} \frac{\exp\big(-\sum_{s\in S} J(\mathbf{x}(v_i), \mathbf{x}_i(\sigma^s v_i)\big)}{\exp\big(-\sum_{s\in S} J(\mathbf{x}'(v_i), \mathbf{x}_{i+1}(\sigma^s v_i)\big)}.$$

For the first piece, since $\mathbf{x}, \mathbf{x}'$ only differ at $v_0, \ldots, v_l$ we have

$$\prod_{i=0}^{l-1} \frac{\exp\big(-h(\mathbf{x}(v_i))\big)}{\exp\big(-h(\mathbf{x}'(v_i))\big)} = \frac{\exp\big(-\sum_{v\in\Lambda} h(\mathbf{x}(v))\big)}{\exp\big(-\sum_{v\in\Lambda} h(\mathbf{x}'(v))\big)}.$$

To deal with the second piece, for the numerator we have

$$\begin{aligned}&\sum_{i=0}^{l-1} \sum_{s\in S} J(\mathbf{x}(v_i), \mathbf{x}_i(\sigma^s v_i)\big) \\ &\quad = \sum_{i=0}^{l-1} \Bigg[\sum_{\substack{s\in S \\ \sigma^s v_i = v_j \text{ for some } j<i}} J(\mathbf{x}(v_i), \mathbf{x}(\sigma^s v_i)\big) + \sum_{\substack{s\in S \\ \sigma^s v_i = v_j \text{ for some } j>i}} J(\mathbf{x}(v_i), \mathbf{x}'(\sigma^s v_i)) \\ &\qquad\qquad + \sum_{\substack{s\in S \\ \sigma^s v_i \neq v_j \ \forall j}} J(\mathbf{x}(v_i), \mathbf{x}(\sigma^s v_i) \Bigg]\end{aligned}$$

and for the denominator we have

$$\sum_{i=0}^{l-1}\sum_{s\in S} J(\mathbf{x}'(v_i), \mathbf{x}_{i+1}(\sigma^s v_i)\big)$$

$$= \sum_{i=0}^{l-1}\left[\sum_{\substack{s\in S\\ \sigma^s v_i = v_j \text{ for some } j<i+1}} J(\mathbf{x}'(v_i), \mathbf{x}(\sigma^s v_i)\big) + \sum_{\substack{s\in S\\ \sigma^s v_i = v_j \text{ for some } j\geq i+1}} J(\mathbf{x}'(v_i), \mathbf{x}'(\sigma^s v_i)) \right.$$

$$\left. + \sum_{\substack{s\in S\\ \sigma^s v_i \neq v_j\ \forall j}} J(\mathbf{x}'(v_i), \mathbf{x}'(\sigma^s v_i)\right].$$

Assuming $e \notin S$, the conditions on the first two sums here are equivalent to those on the first two sums in the numerator.

Now consider the second sum from the exponent in the numerator:

$$\begin{aligned}\sum_{i=0}^{l-1}\sum_{\substack{s\in S\\ \sigma^s v_i = v_j \text{ for some } j>i}} J(\mathbf{x}(v_i), \mathbf{x}'(\sigma^s v_i) &= \sum_{i=0}^{l-1}\sum_{s\in S}\sum_{j=0}^{l-1} \mathbf{1}_{\{j>i\}}\mathbf{1}_{\{\sigma^s v_i = v_j\}} J(\mathbf{x}(v_i), \mathbf{x}'(v_j)) \\ &= \sum_{j=0}^{l-1}\sum_{i=0}^{l-1}\sum_{s\in S} \mathbf{1}_{\{i<j\}}\mathbf{1}_{\{v_i = \sigma^{s^{-1}} v_j\}} J(\mathbf{x}(v_i), \mathbf{x}'(v_j)) \\ &= \sum_{j=0}^{l-1}\sum_{\substack{s\in S\\ \sigma^{s^{-1}} v_j = v_i \text{ for some } i<j}} J(\mathbf{x}(\sigma^{s^{-1}} v_j), \mathbf{x}'(v_j)).\end{aligned}$$

Since $S = \{s_1, \ldots, s_r, s_1^{-1}, \ldots, s_r^{-1}\}$ is closed under taking inverses and J is symmetric, this is equal to the first sum from the exponent in the denominator. Therefore these terms cancel, and we have

$$\frac{\mu_{\mathtt{A},m}\{\mathbf{x}\}}{\mu_{\mathtt{A},m}\{\mathbf{x}'\}} = \frac{\exp\big(-\Phi_\Lambda(\mathbf{x})\big)}{\exp\big(-\Phi_\Lambda(\mathbf{x}')\big)}$$

where

$$\Phi_\Lambda(\mathbf{x}) = \sum_{v\in\Lambda} h(\mathbf{x}(v)) + \sum_{i=0}^{l-1}\left[\sum_{\substack{s\in S\\ \sigma^s v_i = v_j \text{ for some } j<i}} J(\mathbf{x}(v_i), \mathbf{x}(\sigma^s v_i)\big) + \sum_{\substack{s\in S\\ \sigma^s v_i \neq v_j\ \forall j}} J(\mathbf{x}(v_i), \mathbf{x}(\sigma^s v_i)\right].$$

Note that the first sum in the square brackets counts pairwise interactions between adjacent sites in Λ (counting each pair once) while the second sum counts interactions with the outer boundary.

By an argument similar to the one at the end of the previous proof, we can use this to show that $\mu_{\mathtt{A}}(\cdot|\mathscr{T}_{\Lambda}^{\mathtt{A}})(\mathbf{x})$ is specified by Φ_{Λ} in the desired way. $\square$

REFERENCES

[Abe18] Miklos Abert. “A Measured Approach to Sofic Entropy.” Conference Talk, IPAM, May 2018.

[ABL19] Dylan Airey, Lewis Bowen, and Frank Lin. “A Topological Dynamical System with Two Different Positive Sofic Entropies.” *arXiv:1911.08272 [math, stat]*, November 2019.

[Alp16] A. Alpeev. “The Entropy of Gibbs Measures on Sofic Groups.” *Journal of Mathematical Sciences*, **215**(6):649–658, June 2016.

[Aus16] Tim Austin. “Additivity Properties of Sofic Entropy and Measures on Model Spaces.” *Forum of Mathematics, Sigma*, **4**, 2016.

[BG13] Lewis Bowen and Yonatan Gutman. “Nonabelian Free Group Actions: Markov Processes, the Abramov–Rohlin Formula and Yuzvinskii’s Formula – CORRIGENDUM.” *Ergodic Theory and Dynamical Systems*, **33**(2):643–645, April 2013.

[BGT13] Mohsen Bayati, David Gamarnik, and Prasad Tetali. “Combinatorial Approach to the Interpolation Method and Scaling Limits in Sparse Random Graphs.” *Ann. Probab.*, **41**(6):4080–4115, November 2013.

[BK20] Peter J. Burton and Alexander S. Kechris. “Weak Containment of Measure-Preserving Group Actions.” *Ergodic Theory and Dynamical Systems*, **40**(10):2681–2733, October 2020.

[Bow03] Lewis Bowen. “Periodicity and Circle Packings of the Hyperbolic Plane.” *Geometriae Dedicata*, **102**(1):213–236, December 2003.

[Bow10a] Lewis Bowen. “The Ergodic Theory of Free Group Actions: Entropy and the f-Invariant.” *Groups, Geometry, and Dynamics*, pp. 419–432, 2010.

[Bow10b] Lewis Bowen. “A Measure-Conjugacy Invariant for Free Group Actions.” *Annals of Mathematics*, **171**(2):1387–1400, March 2010.

[Bow10c] Lewis Bowen. “Measure Conjugacy Invariants for Actions of Countable Sofic Groups.” *Journal of the American Mathematical Society*, **23**(1):217–217, January 2010.

[Bow10d] Lewis Bowen. “Non-Abelian Free Group Actions: Markov Processes, the Abramov–Rohlin Formula and Yuzvinskii’s Formula.” *Ergodic Theory and Dynamical Systems*, **30**(6):1629–1663, December 2010.

[Bow12] Lewis Bowen. “Sofic Entropy and Amenable Groups.” *Ergodic Theory and Dynamical Systems*, **32**(2):427–466, April 2012.

[Bow20a] Lewis Bowen. "Examples in the Entropy Theory of Countable Group Actions." *Ergodic Theory and Dynamical Systems*, **40**(10):2593–2680, October 2020.

[Bow20b] Lewis Bowen. "Sofic Homological Invariants and the Weak Pinsker Property." *arXiv:1807.08191 [math]*, February 2020.

[BRZ95] P. M. Bleher, J. Ruiz, and V. A. Zagrebnov. "On the Purity of the Limiting Gibbs State for the Ising Model on the Bethe Lattice." *Journal of Statistical Physics*, **79**(1-2):473–482, April 1995.

[CLM20] Amin Coja-Oghlan, Philipp Loick, Balázs F. Mezei, and Gregory B. Sorkin. "The Ising Antiferromagnet and Max Cut on Random Regular Graphs." *arXiv:2009.10483 [cs, math]*, November 2020.

[CM06] Pietro Caputo and Fabio Martinelli. "Phase Ordering after a Deep Quench: The Stochastic Ising and Hard Core Gas Models on a Tree." *Probability Theory and Related Fields*, **136**(1):37–80, September 2006.

[CT06] T. M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, NJ, second edition, 2006.

[DKP13] Ken Dykema, David Kerr, and Mikaël Pichot. "Sofic Dimension for Discrete Measured Groupoids." *Trans. Amer. Math. Soc.*, **366**(2):707–748, September 2013.

[DM10] Amir Dembo and Andrea Montanari. "Ising Models on Locally Tree-like Graphs." *The Annals of Applied Probability*, **20**(2):565–592, April 2010.

[DMS17] Amir Dembo, Andrea Montanari, and Subhabrata Sen. "Extremal Cuts of Sparse Random Graphs." *Ann. Probab.*, **45**(2):1190–1217, March 2017.

[Geo11] Hans-Otto Georgii. *Gibbs Measures and Phase Transitions*. Number 9 in De Gruyter Studies in Mathematics. de Gruyter, Berlin, second edition, 2011.

[GJK02] Catherine Greenhill, Svante Janson, Jeong Han Kim, and Nicholas C. Wormald. "Permutation Pseudographs and Contiguity." *Combinatorics, Probability and Computing*, **11**(3):273–298, May 2002.

[Hay16] Ben Hayes. "Relative Entropy and the Pinsker Product Formula for Sofic Groups." *Groups, Geometry, and Dynamics*, May 2016.

[Hol71] Richard Holley. "Free Energy in a Markovian Model of a Lattice Spin System." *Communications in Mathematical Physics*, **23**(2):87–99, June 1971.

[HS75] Yasunari Higuchi and Tokuzo Shiga. "Some Results on Markov Processes of Infinite Lattice Spin Systems." *Journal of Mathematics of Kyoto University*, **15**(1):211–229, 1975.

[Iof96] Dmitry Ioffe. "On the Extremality of the Disordered State for the Ising Model on the Bethe Lattice." *Letters in Mathematical Physics*, **37**(2):137–143, June 1996.

[Isi25] Ernst Ising. "Beitrag zur Theorie des Ferromagnetismus." *Z. Physik*, **31**(1):253–258, February 1925.

[JK19] Benedikt Jahnel and Christof Külske. "Attractor Properties for Irreversible and Reversible Interacting Particle Systems." *Communications in Mathematical Physics*, **366**(1):139–172, February 2019.

[JNV20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. "MIP*=RE." *arXiv:2001.04383 [quant-ph]*, January 2020.

[Kec12] Alexander S. Kechris. "Weak Containment in the Space of Actions of a Free Group." *Israel Journal of Mathematics*, **189**(1):461–507, June 2012.

[Lig05] Thomas M. Liggett. *Interacting Particle Systems.* Classics in Mathematics. Springer, Berlin, 2005.

[LL58] L.D. Landau and E.M. Lifshitz. *Statistical Physics*, volume 5 of *Course of Theoretical Physics.* Pergamon Press, London, 1958.

[McD98] Colin McDiarmid. "Concentration." In Michel Habib, Colin McDiarmid, Jorge Ramirez-Alfonsin, and Bruce Reed, editors, *Probabilistic Methods for Algorithmic Discrete Mathematics*, pp. 195–248. Springer, Berlin, Heidelberg, 1998.

[MMS12] Andrea Montanari, Elchanan Mossel, and Allan Sly. "The Weak Limit of Ising Models on Locally Tree-like Graphs." *Probability Theory and Related Fields*, **152**(1-2):31–51, February 2012.

[OW87] Donald S. Ornstein and Benjamin Weiss. "Entropy and Isomorphism Theorems for Actions of Amenable Groups." *Journal d'Analyse Mathématique*, **48**(1):1–141, December 1987.

[Pau11] Liviu Păunescu. "On Sofic Actions and Equivalence Relations." *Journal of Functional Analysis*, **261**(9):2461–2485, November 2011.

[Pop14] Sorin Popa. "Independence Properties in Subalgebras of Ultraproduct II 1 Factors." *Journal of Functional Analysis*, **266**(9):5818–5846, May 2014.

[Shr20a] Christopher Shriver. "Free Energy, Gibbs Measures, and Glauber Dynamics for Nearest-Neighbor Interactions on Trees." *arXiv:2011.00653 [math]*, November 2020.

[Shr20b] Christopher Shriver. "The relative f-invariant and non-uniform random sofic approximations." *arXiv:2003.00663 [math]*, March 2020.

[Shr21] Christopher Shriver. "Metastability and Maximal-Entropy Joinings of Gibbs Measures on Finitely-Generated Groups." *arXiv:2103.04467 [math-ph]*, March 2021.

[Wor99] N. C. Wormald. "Models of Random Regular Graphs." In J. D. Lamb and D. A. Preece, editors, *Surveys in Combinatorics, 1999*, pp. 239–298. Cambridge University Press, first edition, July 1999.