



ADDITIVITY PROPERTIES OF SOFIC ENTROPY AND MEASURES ON MODEL SPACES

TIM AUSTIN

Courant Institute, New York University, New York, NY 10012, USA;
email: tim@cims.nyu.edu

Received 4 December 2015; accepted 12 June 2016

Abstract

Sofic entropy is an invariant for probability-preserving actions of sofic groups. It was introduced a few years ago by Lewis Bowen, and shown to extend the classical Kolmogorov–Sinai entropy from the setting of amenable groups. Some parts of Kolmogorov–Sinai entropy theory generalize to sofic entropy, but in other respects this new invariant behaves less regularly. This paper explores conditions under which sofic entropy is additive for Cartesian products of systems. It is always subadditive, but the reverse inequality can fail. We define a new entropy notion in terms of probability distributions on the spaces of good models of an action. Using this, we prove a general lower bound for the sofic entropy of a Cartesian product in terms of separate quantities for the two factor systems involved. We also prove that this lower bound is optimal in a certain sense, and use it to derive some sufficient conditions for the strict additivity of sofic entropy itself. Various other properties of this new entropy notion are also developed.

2010 Mathematics Subject Classification: 37A35 (primary); 37A50, 60K35, 82B20 (secondary)

1. Introduction

Let G be a discrete sofic group, (X, μ) a standard probability space and $T : G \curvearrowright X$ a measurable action which preserves μ . The triple (X, μ, T) is called a G -**system** or just a **system**.

Fix a sofic approximation $\Sigma = (\sigma_n)_{n \geq 1}$ to the group G . For a system (X, μ, T) which has a finite generating partition, Lewis Bowen defined the ‘sofic entropy relative to Σ ’, denoted by $h_\Sigma(\mu, T)$ [4]. An alternative definition which does not require a finite generating partition, and so generalizes Bowen’s, was given by

© The Author 2016. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

Kerr and Li in [15]. That definition was in terms of operator algebras, but they later gave a more elementary construction of the same invariant: [16, Section 3].

If G is amenable, then sofic entropy agrees with the classical Kolmogorov–Sinai entropy h_{KS} for any choice of sofic approximation (see [7, 16]). If G is not amenable, then sofic entropy can serve as a substitute for Kolmogorov–Sinai entropy for some purposes. Bowen’s original motivation was to show that isomorphic Bernoulli shifts over a sofic group G must have one-dimensional distributions of equal Shannon entropy: this was accomplished for shifts with finite alphabet in [4] and then completed in [14]. (The sufficiency of that condition is also known in many cases, including for any G that contains an infinite amenable subgroup.) However, it is still uncertain how much of classical entropy theory generalizes to sofic groups, or what modifications are necessary.

This paper considers how h_{Σ} behaves under forming Cartesian products of systems. If G is amenable, then Kolmogorov–Sinai entropy is additive under Cartesian products:

$$h_{KS}(\mu \times \nu, T \times S) = h_{KS}(\mu, T) + h_{KS}(\nu, S).$$

Sofic entropy is subadditive for Cartesian products, and indeed for arbitrary joinings; the easy proof of this is given in Subsection 3.3. However, strict additivity can fail: examples showing this will also be given in that subsection.

Suppose that (X, μ, T) is a G -system. The main innovation of this paper is a new quantity, denoted by $h_{\Sigma}^{dq}(\mu, T)$, with the property that

$$h_{\Sigma}(\mu \times \nu, T \times S) \geq h_{\Sigma}^{dq}(\mu, T) + h_{\Sigma}(\nu, S)$$

whenever (Y, ν, S) is another G -system for which $h_{\Sigma}(\nu, S)$ has a certain regularity (it is equal to the ‘lower’ sofic entropy $\underline{h}_{\Sigma}(\nu, S)$, which is recalled below).

The sofic entropy of (X, μ, T) is obtained from a certain family of metric spaces consisting of ‘finitary models’ for (X, μ, T) . The new quantity $h_{\Sigma}^{dq}(\mu, T)$ is also defined in terms of these spaces, but together with another kind of structure: sequences of probability measures μ_n on those model spaces.

Model spaces and sofic entropy. In order to formulate our main results precisely, we first recall the construction of model spaces and the definition of sofic entropy.

Our definition is very close to that in [16, Section 3] but it is adapted slightly better to the purposes of this paper. The following is only a sketch; full details are given in Section 3. The equivalence with the Kerr–Li definition is shown in Subsection 3.2.

First, a **G -process** is a G -system in which $X = \mathcal{X}^G$ for some other standard measurable space \mathcal{X} and S is the right-shift action of G on X :

$$S^g((x_h)_{h \in G}) := (x_{hg})_{h \in G}.$$

This is close to the probabilistic notion of a ‘ G -stationary process’: formally, that would be the collection of coordinate projections $\mathcal{X}^G \rightarrow \mathcal{X}$, regarded as a G -indexed family of \mathcal{X} -valued random variables on the probability space (\mathcal{X}^G, μ) .

When we deal with isomorphism-invariant properties of systems, no generality is lost by confining our attention to G -processes. Indeed, for any system (X, μ, T) , the map

$$\Phi : X \rightarrow X^G : x \mapsto (T^h x)_{h \in G} \quad (1)$$

intertwines T with the right-shift action S on X^G , and converts μ into the shift-invariant measure $\Phi_*\mu$. This map is clearly injective, so it defines a measure-theoretic isomorphism from (X, μ, T) to the G -process $(X^G, \Phi_*\mu, S)$.

If (\mathcal{X}^G, μ, S) is a G -process and $F \subseteq G$, then μ_F denotes the marginal of μ on \mathcal{X}^F . Also, whenever $F \subseteq F' \subseteq G$, we let $\pi_F^{F'}$ denote the coordinate projection $\mathcal{X}^{F'} \rightarrow \mathcal{X}^F$. Thus,

$$\mu_F = (\pi_F^{F'})_*\mu_{F'} = (\pi_F^G)_*\mu.$$

If $x \in \mathcal{X}^{F'}$ then we often write $x|_F$ as a shorthand for $\pi_F^{F'}(x)$.

Next, since \mathcal{X} is standard, its σ -algebra may be generated as the Borel sets for some compact metric d . We refer to such a d as a **compact generating metric** for \mathcal{X} . Although this metric is far from unique, it is a key auxiliary object in the constructions that follow. A **metric G -process** is a quadruple $(\mathcal{X}^G, \mu, S, d)$ in which (\mathcal{X}^G, μ, S) is a G -process and d is such a metric on \mathcal{X} . Sofic entropy is initially defined for metric G -processes, and then one shows that it does not depend on the choice of metric. One can extend this fact to allow more general Polish generating metrics on \mathcal{X} [12], but we do not do so here. Once a metric d has been chosen, it is always implicit that \mathcal{X}^G has the resulting compact metrizable product topology, and similarly for other products of standard spaces for which we have chosen compact generating metrics.

Now let V be a nonempty finite set. The space \mathcal{X}^V carries an associated metric defined by

$$d^{(V)}(\mathbf{x}, \mathbf{x}') = \frac{1}{|V|} \sum_{v \in V} d(x_v, x'_v) \quad \text{for } \mathbf{x} = (x_v)_{v \in V}, \quad \mathbf{x}' = (x'_v)_{v \in V}.$$

We call this the **Hamming average** of d over V . It generalizes the classical normalized Hamming metrics, which arise in this way when (\mathcal{X}, d) is a finite set with the discrete metric.

Now consider again a sofic approximation $\Sigma = (\sigma_n)_{n \geq 1}$ to G , where each σ_n is a map from G to $\text{Sym}(V_n)$ for some finite set V_n (see Section 3 for a complete definition). Given $(\mathcal{X}^G, \mu, \mathcal{S}, d)$ and Σ , we define a sequence of subsets

$$\Omega(\mathcal{O}, \sigma_n) \subseteq \mathcal{X}^{V_n}, \quad n \in \mathbb{N},$$

for each choice of a weak*-neighbourhood \mathcal{O} of μ in $\text{Prob}(\mathcal{X}^G)$. Note that the weak* topology invoked here depends on the topology of \mathcal{X}^G , which in turn depends on the choice of d . The elements of $\Omega(\mathcal{O}, \sigma_n)$ are the ‘ \mathcal{O} -good models of μ over σ_n ’: the elements of \mathcal{X}^{V_n} whose empirical distributions lie in \mathcal{O} , hence ‘close’ to μ . Empirical distributions are defined in Section 3. The ‘quality’ required of these models improves as \mathcal{O} is reduced, and it is clear that

$$\mathcal{O}' \subseteq \mathcal{O} \implies \Omega(\mathcal{O}', \sigma) \subseteq \Omega(\mathcal{O}, \sigma)$$

for any finite V and map $\sigma : G \rightarrow \text{Sym}(V)$.

For any compact metric space (Y, d_Y) , subset $Z \subseteq Y$, and $\delta > 0$, we let $\text{cov}_\delta(Z, d_Y)$ be the minimum cardinality among δ -dense subsets of Z . The sofic entropy $h_\Sigma(\mu)$ is defined to be

$$\sup_{\delta > 0} \inf_{\mathcal{O}} \limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{cov}_\delta(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)}),$$

where \mathcal{O} ranges over all weak*-neighbourhoods of μ . Heuristically, this is approximately the exponential growth rate of the covering numbers

$$\text{cov}_\delta(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)})$$

as $n \rightarrow \infty$, for sufficiently small δ and then for sufficiently small \mathcal{O} depending on δ .

In general, this sequence of covering numbers need not grow at a well-defined exponential rate, so one takes the supremum of those rates over subsequences. It can be important to know when the covering numbers have different asymptotics along other subsequences. To capture this possibility, one also defines the **lower sofic entropy**

$$\underline{h}_\Sigma(\mu) := \sup_{\delta > 0} \inf_{\mathcal{O}} \liminf_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{cov}_\delta(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)}).$$

It can happen that $\underline{h}_\Sigma(\mu) < h_\Sigma(\mu)$: see the end of Subsection 3.1. If $\underline{h}_\Sigma(\mu) = h_\Sigma(\mu)$, then this asserts the following: for every $\varepsilon > 0$ there is a $\delta_0 > 0$ such that for every $\delta \in (0, \delta_0)$ there is a weak*-neighbourhood \mathcal{O}_δ such that for every weak*-neighbourhood $\mathcal{O} \subseteq \mathcal{O}_\delta$ we have

$$\left| \frac{1}{|V_n|} \log \text{cov}_\delta(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)}) - h_\Sigma(\mu) \right| < \varepsilon$$

for all sufficiently large n . More heuristically: if δ is sufficiently small and then \mathcal{O} is sufficiently small depending on δ , the covering numbers do grow at an approximately well-defined exponential rate. Since it suffices to check this for rational δ and for a countable basis of neighbourhoods \mathcal{O} , a simple diagonal argument can always provide a subsequence of $(\sigma_n)_{n \geq 1}$ for which this is the case.

The quantity $h_\Sigma(\mu)$ generally depends on the choice of sofic approximation Σ . But, crucially, it does not depend on the compact metric d that one uses to generate the σ -algebra of \mathcal{X} . In fact, $h_\Sigma(\mu)$ and likewise $\underline{h}_\Sigma(\mu)$ are invariants of the system (\mathcal{X}^G, μ, S) up to measure-theoretic isomorphism, for a fixed choice of the sofic approximation Σ . One may therefore define $h_\Sigma(\mu, T)$ and similarly $\underline{h}_\Sigma(\mu, T)$ for an arbitrary G -system (X, μ, T) , for instance by using the isomorphism (1).

Measures on model spaces. Our new invariant is defined in terms of sequences of measures μ_n on \mathcal{X}^{V_n} which are asymptotically supported on these model spaces and which locally resemble μ at most points of V_n , which we refer to as ‘vertices’.

Section 5 will consider three senses in which a sequence of probability measures μ_n on \mathcal{X}^{V_n} can converge to a measure μ on \mathcal{X}^G : local weak* convergence, quenched convergence, and doubly quenched convergence. All three senses are relative to a particular choice of sofic approximation Σ ; we refer to convergence ‘over Σ ’ if we need to make that choice explicit. They are also relative to a particular choice of compact generating metric d for \mathcal{X} .

Local weak* convergence asserts that, once n is large, the marginals of μ_n around most points of V_n resemble the corresponding marginal of μ in the weak* topology. Here we use that the sofic approximation σ_n gives a way to copy a fixed finite subset of G to a corresponding ‘patch’ around any vertex of V_n , perhaps with errors for a few vertices. This convergence is denoted by $\mu_n \xrightarrow{\text{lw}^*} \mu$. This notion already has an important role in the study of various statistical physics models on random graphs.

Quenched convergence strengthens local weak* convergence by imposing a second condition: that μ_n be mostly supported on individual good models of μ . The term ‘quenched’ is also taken from statistical physics, where it indicates a property that holds among most instances in an ensemble, not just on average. This convergence is denoted by $\mu_n \xrightarrow{q} \mu$. Quenched convergence is strictly stronger than local weak* convergence in general, and the difference between them has a simple characterization in terms of certain random measures constructed from the μ_n (Lemma 5.5). However, using this characterization, it follows that the two notions are equivalent if μ is ergodic (Corollary 5.7).

A closely related notion of convergence for a sequence of measures on model spaces already appears in [6]. Using this notion, that paper gives a new formula

for sofic entropy for certain special examples of probability-preserving systems and sofic approximations.

For any compact metric space (Y, d_Y) , Borel probability measure ν on Y , and $\varepsilon, \delta > 0$, we write

$$\text{cov}_{\varepsilon, \delta}(\nu, d_Y) := \min\{|F| : F \subseteq Y, \nu(B_\delta(F)) > 1 - \varepsilon\},$$

where $B_\delta(F)$ is the δ -neighbourhood of F according to the metric d_Y . Using this quantity in place of the covering numbers of spaces themselves, we define the following analogue of h_Σ :

$$h_\Sigma^q(\mu) := \sup \left\{ \sup_{\delta, \varepsilon > 0} \limsup_{i \rightarrow \infty} \frac{1}{|V_{n_i}|} \log \text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})}) : \right. \\ \left. n_i \uparrow \infty \text{ and } \mu_i \xrightarrow{q} \mu \text{ over } (\sigma_{n_i})_{i \geq 1} \right\}.$$

The outer supremum here is over all subsequences $(\sigma_{n_i})_{i \geq 1}$ of the sofic approximation Σ , and over all sequences of measures μ_i on $\mathcal{X}^{V_{n_i}}$ that quenched-converge to μ over that subsequence. We must allow this supremum over subsequences, because it may be that there is no sequence of measures μ_n such that $\mu_n \xrightarrow{q} \mu$ over the original sofic approximation at all. This will be explained more carefully at the beginning of Section 6.

We call $h_\Sigma^q(\mu)$ the **model-measure sofic entropy of μ rel Σ** . Like sofic entropy, it is an isomorphism-invariant of the G -process (Theorem 6.4), and so in fact it does not depend on the choice of the generating metric d . As a result, its definition can be extended unambiguously to arbitrary G -systems. Since quenched convergence $\mu_n \xrightarrow{q} \mu$ requires that μ_n be mostly supported on good models for μ once n is large, it follows easily that $h_\Sigma \geq h_\Sigma^q$ (see Lemma 6.1). This inequality can be strict.

Like sofic entropy, h_Σ^q is always subadditive under Cartesian products, but may not be strictly additive. However, this defect can be repaired by further restricting the sequences of measures on model spaces that we allow.

If $\mu_n \xrightarrow{\text{lw}^*} \mu$, then it follows easily that $\mu_n \times \mu_n \xrightarrow{\text{lw}^*} \mu \times \mu$. However, the same implication may fail for quenched convergence: even if μ_n is asymptotically mostly supported on good models for μ , the product $\mu_n \times \mu_n$ may not be mostly supported on good models for $\mu \times \mu$. This phenomenon is responsible for cases in which $h_\Sigma^q(\mu^{\times 2}) < 2h_\Sigma^q(\mu)$. However, if we simply *require* the convergence of Cartesian squares

$$\mu_n \times \mu_n \xrightarrow{q} \mu \times \mu,$$

then it turns out that this implies good behaviour for all other Cartesian products with the measures μ_n .

THEOREM A. *Suppose that $\mu_n \xrightarrow{q} \mu$ over Σ . The following are equivalent:*

- (i) $\mu_n \times \mu_n \xrightarrow{q} \mu \times \mu$;
- (ii) *if $(\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$ is another metric G -process and \mathcal{N} is a weak* neighbourhood of $\mu \times \nu$, then there is a weak* neighbourhood \mathcal{O} of ν such that*

$$\inf_{\nu \in \Omega(\mathcal{O}, \sigma_{n_i})} \mu_{n_i} \{ \mathbf{x} \in \mathcal{X}^{V_{n_i}} : (\mathbf{x}, \mathbf{y}) \in \Omega(\mathcal{N}, \sigma_{n_i}) \} \longrightarrow 1 \quad \text{as } i \longrightarrow \infty$$

for any subsequence $n_i \uparrow \infty$ such that $\Omega(\mathcal{O}, \sigma_{n_i}) \neq \emptyset$ for all i (we regard this as vacuously true if there is no such subsequence n_i);

- (iii) *if $(\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$ is another metric G -process, $n_i \uparrow \infty$, and $\nu_i \in \text{Prob}(\mathcal{Y}^{V_{n_i}})$ is a sequence such that $\nu_i \xrightarrow{q} \nu$ over $(\sigma_{n_i})_{i \geq 1}$, then $\mu_{n_i} \times \nu_i \xrightarrow{q} \mu \times \nu$ over $(\sigma_{n_i})_{i \geq 1}$.*

Under any of the above equivalent conditions, we say that $(\mu_n)_{n \geq 1}$ **doubly quenched-converges** to μ , and denote this by $\mu_n \xrightarrow{\text{dq}} \mu$.

Theorem A is analogous to the equivalence among various standard characterizations of weak mixing. Furthermore, since local weak* convergence implies quenched convergence when μ is ergodic, one can deduce that quenched convergence implies doubly quenched convergence when μ is weakly mixing (Lemma 5.15).

Doubly quenched convergence finally leads to the new invariant we need:

$$h_{\Sigma}^{\text{dq}}(\mu) := \sup \left\{ \sup_{\delta, \varepsilon > 0} \limsup_{i \rightarrow \infty} \frac{1}{|V_{n_i}|} \log \text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})}) : \right. \\ \left. n_i \uparrow \infty \text{ and } \mu_i \xrightarrow{\text{dq}} \mu \text{ over } (\sigma_{n_i})_{i \geq 1} \right\}.$$

This is called the **doubly quenched model-measure sofic entropy of μ rel Σ** . One sees easily that $h_{\Sigma}^{\text{dq}} \leq h_{\Sigma}^q$ (Lemma 6.1). The proof that h_{Σ}^q is isomorphism-invariant (and hence independent of the choice of d) gives the same result for h_{Σ}^{dq} , and so the definition of h_{Σ}^{dq} can be extended unambiguously to any G -system.

We can now state our main result for Cartesian products.

THEOREM B. *Suppose that (X, μ, T) and (Y, ν, S) are G -systems such that*

$$h_{\Sigma}(v, S) = \underline{h}_{\Sigma}(v, S). \tag{2}$$

Then

$$h_{\Sigma}(\mu \times \nu, T \times S) \geq h_{\Sigma}^{\text{dq}}(\mu, T) + h_{\Sigma}(v, S). \tag{3}$$

This follows fairly easily from conclusion (ii) of Theorem A.

Of course, by symmetry, one also has the analogous conclusion with the roles of (X, μ, T) and (Y, ν, S) reversed.

One cannot hope for anything like (3) without some assumption such as (2), since in general the relevant entropies $h_{\Sigma}^{\text{dq}}(\mu, T)$ and $h_{\Sigma}(\nu, S)$ could be obtained as limit suprema along disjoint subsequences.

Since $h_{\Sigma}^{\text{dq}} \leq h_{\Sigma}$ and h_{Σ} is always subadditive, the following is an immediate corollary.

COROLLARY B'. *If (X, μ, T) and (Y, ν, S) satisfy*

$$h_{\Sigma}^{\text{dq}}(\mu, T) = h_{\Sigma}(\mu, T) \quad \text{and} \quad \underline{h}_{\Sigma}(\nu, S) = h_{\Sigma}(\nu, S),$$

or vice versa, then

$$h_{\Sigma}(\mu \times \nu, T \times S) = h_{\Sigma}(\mu, T) + h_{\Sigma}(\nu, S).$$

For instance, this condition on (X, μ, T) is satisfied by Bernoulli systems, so we recover the known result [4, Section 8] that forming products with Bernoulli systems always makes the obvious additive contribution to sofic entropy.

It can happen that $h_{\Sigma}^{\text{q}} > h_{\Sigma}^{\text{dq}}$, and there are cases in which one cannot replace $h_{\Sigma}^{\text{dq}}(\mu, T)$ with $h_{\Sigma}^{\text{q}}(\mu, T)$ in Theorem B; see Example 6.3. However, this cannot occur if (X, μ, T) is weakly mixing, simply because quenched convergence itself implies doubly quenched convergence for weakly mixing systems.

Unlike the other notions, doubly quenched model-measure sofic entropy does enjoy a general additivity result for Cartesian products. The only caveat is that we must still assume some analogue of condition (2) in Theorem B. This is conveniently expressed in terms of a ‘lower’ version of doubly quenched model-measure sofic entropy, denoted $\underline{h}_{\Sigma}^{\text{dq}}$.

The additivity result for h_{Σ}^{dq} has an analogue for h_{Σ}^{q} under an additional ergodicity assumption.

THEOREM C. *For any G -systems (X, μ, T) and (Y, ν, S) , it holds that*

$$h_{\Sigma}^{\text{dq}}(\mu \times \nu, T \times S) \leq h_{\Sigma}^{\text{dq}}(\mu, T) + h_{\Sigma}^{\text{dq}}(\nu, S),$$

and similarly with h_{Σ}^{dq} replaced by h_{Σ}^{q} .

If $h_{\Sigma}^{\text{dq}}(\nu, S) = \underline{h}_{\Sigma}^{\text{dq}}(\nu, S)$, then in fact

$$h_{\Sigma}^{\text{dq}}(\mu \times \nu, T \times S) = h_{\Sigma}^{\text{dq}}(\mu, T) + h_{\Sigma}^{\text{dq}}(\nu, S), \quad (4)$$

where we interpret the right-hand side as $-\infty$ if either of its terms is $-\infty$.

If $\mu \times \nu$ is ergodic, then the analogous result holds with h_{Σ}^{dq} and $\underline{h}_{\Sigma}^{\text{dq}}$ replaced by h_{Σ}^{q} and $\underline{h}_{\Sigma}^{\text{q}}$ throughout. In particular, this is the case if both μ and ν are ergodic and one of them is weakly mixing.

Theorems **B** and **C** will be proved in Section 7. The first conclusion of Theorem **C** (subadditivity) actually holds for arbitrary joinings; see Proposition 7.1.

By applying Theorem **C** to copies of a single system (X, μ, T) , we can show that h_{Σ}^{dq} is stable under Cartesian powers:

$$h_{\Sigma}^{\text{dq}}(\mu^{\times k}) = k \cdot h_{\Sigma}^{\text{dq}}(\mu) \quad \forall k \geq 1$$

(see Corollary 7.3). In this case we can do without the assumption that $h_{\Sigma}^{\text{dq}}(\mu, T) = \underline{h}_{\Sigma}^{\text{dq}}(\mu, T)$.

Processes with finite state spaces. In case \mathcal{X} is a finite set and (\mathcal{X}^G, μ, S) is a G -process, we are able to prove another relation between the new invariant h_{Σ}^{dq} and sofic entropies. In view of isomorphism invariance, this will actually apply to any system (X, μ, T) which has a finite generating partition. For an ergodic system, this, in turn, is equivalent to finiteness of the Rokhlin entropy $h^{\text{Rok}}(\mu, T)$, by the results of [23].

For a general system (X, μ, T) , consider the sequence of values

$$\frac{1}{k} h_{\Sigma}(\mu^{\times k}, T^{\times k}), \quad k \geq 1.$$

Since h_{Σ} is subadditive, the sequence $(h_{\Sigma}(\mu^{\times k}, T^{\times k}))_{k \geq 1}$ is subadditive. Therefore the limit

$$h_{\Sigma}^{\text{ps}}(\mu, T) := \lim_{k \rightarrow \infty} \frac{1}{k} h_{\Sigma}(\mu^{\times k}, T^{\times k})$$

exists and satisfies $h_{\Sigma}^{\text{ps}}(\mu, T) \leq h_{\Sigma}(\mu, T)$, by Fekete's lemma. We call it the **power-stabilized sofic entropy rel Σ** . It is clearly an isomorphism-invariant. If $h_{\Sigma}^{\text{ps}}(\mu, T) < h_{\Sigma}(\mu, T)$, then this gap quantifies the failure of additivity of sofic entropy among the Cartesian powers of (X, μ, T) itself.

THEOREM D. *It is always the case that*

$$h_{\Sigma}^{\text{ps}}(\mu, T) \geq h_{\Sigma}^{\text{dq}}(\mu, T),$$

and this is an equality if (X, μ, T) has a finite generating partition.

It is fairly easy to prove that $h_{\Sigma}^{\text{ps}} \geq h_{\Sigma}^{\text{dq}}$, so most of the work goes into proving the reverse inequality. This will be done by showing that individual good models for large Cartesian powers $(\mu^{\times k}, T^{\times k})$ can be converted into measures that doubly quenched-converge to μ .

Theorem D leads to the following sense in which Theorem B is optimal for systems that have finite generating partitions.

COROLLARY D'. *If (X, μ, T) has a finite generating partition and $h_{\Sigma}(\mu^{\times k}, T^{\times k}) = \underline{h}_{\Sigma}(\mu^{\times k}, T^{\times k})$ for all k , then*

$$h_{\Sigma}^{\text{dq}}(\mu, T) = \inf\{h_{\Sigma}(\mu \times \nu, T \times S) - h_{\Sigma}(\nu, S) : (Y, \nu, S) \text{ another } G\text{-system with } h_{\Sigma}(\nu, S) = \underline{h}_{\Sigma}(\nu, S)\}.$$

Thus, if (X, μ, T) has a finite generating partition and $h_{\Sigma}(\mu^{\times k}, T^{\times k}) = \underline{h}_{\Sigma}(\mu^{\times k}, T^{\times k})$ for all k , then no other quantity which depends only on (X, μ, T) can improve on $h_{\Sigma}^{\text{dq}}(\mu, T)$ in Theorem B.

Equality in the case of co-induced systems. For an infinite sofic group G and a Bernoulli process $(\mathcal{X}^G, \nu^{\times G}, S)$, it is fairly easy to show that $h_{\Sigma}(\nu^{\times G})$, $h_{\Sigma}^{\text{q}}(\nu^{\times G})$ and $h_{\Sigma}^{\text{dq}}(\nu^{\times G})$ are all just equal to the Shannon entropy of ν . (In the case of h_{Σ} , this is the calculation that gives the classification of Bernoulli systems in [4, 14].) This supplies some examples for which the three entropy notions coincide.

The final topic of this paper is a generalization of this result to a class of co-inductions. Suppose now that G and H are two sofic groups, and let

$$\Sigma = (\sigma_n : G \longrightarrow \text{Sym}(V_n))_{n \geq 1} \quad \text{and} \quad T = (\tau_n : H \longrightarrow \text{Sym}(W_n))_{n \geq 1}$$

be respective sofic approximations for them. Then $G \times H$ has a **product** sofic approximation $\Sigma \times T = (\sigma_n \times \tau_n)_{n \geq 1}$, where

$$(\sigma_n \times \tau_n)^{(g,h)} := \sigma_n^g \times \tau_n^h \in \text{Sym}(V_n \times W_n).$$

Let (X, μ, T) be a G -system. Then **co-induction** gives a functorial way to construct from it a $(G \times H)$ -system: the new probability space is $(X^H, \mu^{\times H})$, and the action $\text{CInd}_G^{G \times H} T$ is defined by setting

$$(\text{CInd}_G^{G \times H} T)^g := (T^g)^{\times H} \quad \text{for } g \in G$$

and

$$(\text{CInd}_G^{G \times H} T)^h := (x_k)_{k \in H} \quad \text{for } h \in H$$

(so H acts simply by the right shift on X^H). Since these commute, they define an action of $G \times H$. See [13, Subsection II.10.(G)] or [9] for a more general discussion and some previous uses of co-induction in ergodic theory.

If G is the trivial group, then this just produces an H -Bernoulli shift. The following result therefore generalizes our observation about Bernoulli shifts.

THEOREM E. *Let G , H , Σ , T , and (X, μ, T) be as above. Assume that H is infinite. Then*

$$h_{\Sigma \times T}(\mu^{\times H}, \text{CInd}_G^{G \times H} T) = h_{\Sigma \times T}^q(\mu^{\times H}, \text{CInd}_G^{G \times H} T) = h_{\Sigma \times T}^{\text{dq}}(\mu^{\times H}, \text{CInd}_G^{G \times H} T).$$

For a G -system (X, μ, T) , a different condition which is sufficient for $h_{\Sigma}(\mu, T) = h_{\Sigma}^q(\mu, T)$ has appeared previously as [6, Theorem 4.1].

Combining Theorem E with Theorems B and D immediately yields the following.

COROLLARY E'. *Consider again the setting of Theorem E.*

(1) *If (Y, ν, S) is another $(G \times H)$ -system such that*

$$h_{\Sigma \times T}(\nu, S) = \underline{h}_{\Sigma \times T}(\nu, S),$$

then

$$h_{\Sigma \times T}(\mu^{\times H} \times \nu, \text{CInd}_G^{G \times H} T \times S) = h_{\Sigma \times T}(\mu^{\times H}, \text{CInd}_G^{G \times H} T) + h_{\Sigma \times T}(\nu, S).$$

(2) *If (X, μ, T) has a finite generating partition, then*

$$h_{\Sigma \times T}((\mu^{\times H})^{\times k}, (\text{CInd}_G^{G \times H} T)^{\times k}) = k \cdot h_{\Sigma \times T}(\mu^{\times H}, \text{CInd}_G^{G \times H} T)$$

for all $k \geq 1$.

□

Overview of the paper. The rest of this paper is divided into two parts.

Part I concerns the spaces of good models for a G -process. After collecting some background material in Section 2, model spaces and sofic entropy are defined carefully and studied in Section 3.

The most substantial results of this part are in Section 4, which describes how a factor map between G -processes can be converted into ‘approximately Lipschitz’ maps between their model spaces, endowed with suitable Hamming-like metrics. This is delicate, because an arbitrary measurable factor map must first be approximated by ‘almost continuous’ maps, and then these can be modified to act on the model spaces.

Part II introduces measures on the spaces of models of a G -process. Section 5 defines locally weak*, quenched and doubly quenched convergence, and establishes some basic properties, including Theorem A. Then Section 6 uses such convergent sequences of measures to define the associated entropy notions and prove some preparatory results for them. These then lead to proofs of all the main theorems stated above: Theorems B and C in Section 7; Theorem D in Section 8; and Theorem E in Section 9.

When a factor map between G -processes is converted into maps between the model spaces, those can then be used to convert a convergent sequence of measures for the domain process into a convergent sequence of measures for the target process. This construction is crucial for many of the proofs in Part II, but, as might be expected from Part I, it requires careful control of several different approximations. This is the most technical aspect of the paper.

Section 10 collects some open questions about this new notion of entropy.

Many of the ideas in Part I are just small variations on [4, 14–16]. Certainly none of Section 3 is really original. The principal difference from those works is our explicit development of maps between model spaces corresponding to factor maps between systems. This gives us some very versatile tools for passing between model spaces, as illustrated by their use in the proofs of the main theorems. We develop the somewhat new formalism for sofic entropy in Section 3 in order to make this analysis of maps simpler and more natural, and to have the same effect on our definitions of model-measure sofic entropies in Part II. This is why we include this new formalism, rather than just using the definitions from [16], for example.

Part I

Model spaces and maps

2. Some notation and preliminaries

2.1. Elementary analysis. We use Landau ('big- O ' and 'little- o ') notation without further comment. Among real numbers, we sometimes write ' $a \approx_\varepsilon b$ ' in place of ' $|a - b| < \varepsilon$ '. The notation ' $\varepsilon_n \downarrow 0$ ' means that $(\varepsilon_n)_{n \geq 1}$ is a nonincreasing sequence of strictly positive real numbers which tends to 0.

Now let P be a property that holds for some nondecreasing sequences of nonnegative integers (that is, a subset of the set of nondecreasing members of $\mathbb{N}^{\mathbb{N}}$). We will say that P holds **whenever** $m_1 \leq m_2 \leq \dots$ **grows sufficiently slowly** if there is a fixed nondecreasing sequence $(m_n^\circ)_{n \geq 1}$ with $m_n^\circ \uparrow \infty$ such that, for any

other nondecreasing sequence $(m_n)_{n \geq 1} \in \mathbb{N}^{\mathbb{N}}$, we have

$$[m_n \longrightarrow \infty \text{ and } m_n \leq m_n^\circ \forall n] \implies (m_n)_{n \geq 1} \text{ has P.}$$

The following nomenclature from probabilistic combinatorics will be very convenient. If $(\Omega_n, \mathbb{P}_n)_{n \geq 1}$ is a sequence of probability spaces, and P is a property that holds for some elements of Ω_n for each n , then P holds **with high probability** ('**w.h.p.**') if

$$\mathbb{P}_n\{\omega \in \Omega_n : \text{P holds for } \omega\} \longrightarrow 1 \quad \text{as } n \longrightarrow \infty.$$

If P is written out explicitly in terms of ω , then we may also write that P holds **w.h.p. in ω** . If each Ω_n is finite and \mathbb{P}_n is not specified, then this is to be understood with \mathbb{P}_n equal to the uniform measure on Ω_n .

2.2. Metric spaces and almost-Lipschitz maps. Let (X, d_X) and (Y, d_Y) be metric spaces, let $\varepsilon > 0$, and let $L < \infty$. A map $\varphi : X \longrightarrow Y$ is ε -**almost L -Lipschitz** if it is Borel measurable and satisfies

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

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

The following requires only an immediate check.

LEMMA 2.1. *If $\varphi : X \longrightarrow Y$ is ε -almost L -Lipschitz and $f : Y \longrightarrow \mathbb{R}$ is K -Lipschitz, then $f \circ \varphi$ is $(K\varepsilon)$ -almost (KL) -Lipschitz. \square*

Maps that are η -almost Lipschitz for arbitrarily small η have the following simple characterization.

LEMMA 2.2. *If (X, d_X) and (Y, d_Y) are compact and $f : X \longrightarrow Y$, then f is continuous if and only if it is η -almost Lipschitz for every $\eta > 0$.*

Proof. The reverse implication is simple, so we focus on the forward implication.

Suppose for simplicity that d_Y has diameter at most 1. Given $\eta > 0$, let $\delta > 0$ be so small that if $x, x' \in X$ then

$$d_X(x, x') < \delta \implies d_Y(f(x), f(x')) < \eta.$$

Now define $L := 1/\delta$. For any $x, x' \in X$, we obtain

$$d_Y(f(x), f(x')) \leq \begin{cases} \eta & \text{if } d_X(x, x') < \delta \\ 1 < \eta + L\delta & \text{if } d_X(x, x') \geq \delta \end{cases} \leq \eta + Ld_X(x, x'). \quad \square$$

For a general target space, an ε -almost-Lipschitz map need not be close to a truly Lipschitz map, even if ε is very small. However, this does hold among \mathbb{R} -valued maps.

LEMMA 2.3. *If (X, d) is a metric space, $U \subseteq X$ is nonempty, and $f : X \rightarrow \mathbb{R}$ is a function such that $f|_U$ is ε -almost L -Lipschitz, then there is an L -Lipschitz map $g : X \rightarrow \mathbb{R}$ such that $|f(x) - g(x)| \leq \varepsilon$ for all $x \in U$.*

Proof. The following standard construction gives a suitable approximant:

$$g(x) := \inf_{x' \in U} (f(x') + Ld(x, x')). \quad \square$$

As in the Introduction, if V is a nonempty finite set and (X, d) is a metric space, then $d^{(V)}$ denotes the Hamming-average metric on X^V defined by

$$d^{(V)}(x, x') := \frac{1}{|V|} \sum_{v \in V} d(x_v, x'_v).$$

If $V \subseteq U$ with V nonempty and finite, then $d^{(V)}(x, x')$ is still well defined for pairs $x, x' \in X^U$. It defines a pseudometric on X^U . In particular, if G is a group, then this gives a pseudometric $d^{(F)}$ on X^G for every nonempty finite $F \subseteq G$.

Probability measures on metric spaces will always be defined on their Borel σ -algebras. The set of Borel probability measures on a metric space X is denoted by $\text{Prob}(X)$. If X is compact then $\text{Prob}(X)$ is given the weak* topology.

2.3. Approximating Borel maps by almost-Lipschitz maps. In order to study factor maps between systems, we will need to approximate Borel maps by maps that have a fairly explicit kind of ‘approximate continuity’. For our purposes, the best-adapted approximants seem to be almost-Lipschitz maps. In this subsection we prove the existence of such approximants using Lusin’s theorem.

The following definition is not standard, but will be very convenient. It is a prelude to Definition 4.2, which is a dynamical version.

DEFINITION 2.4. Let (X, d_X) and (Y, d_Y) be compact metric spaces, let $\mu \in \text{Prob}(X)$, let $\varphi : X \rightarrow Y$ be Borel, and let $\eta > 0$. Then an η -almost Lipschitz (or η -AL) **approximation to φ rel μ** is a Borel map $\psi : X \rightarrow Y$ with the following properties:

- (i) the map ψ approximates φ in the sense that

$$\int d_Y(\varphi(x), \psi(x)) \mu(dx) < \eta;$$

- (ii) there is an open subset $U \subseteq X$ such that $\mu(U) > 1 - \eta$, and such that $\psi|_U$ is η -almost Lipschitz from d_X to d_Y .

LEMMA 2.5. *Let (X, d_X) and (Y, d_Y) be compact metric spaces, let $\mu \in \text{Prob}(X)$, and let $\varphi : X \rightarrow Y$ be Borel. Then φ has η -AL approximations rel μ for all $\eta > 0$.*

Proof. We may assume for simplicity that d_Y has diameter at most 1.

Lusin's theorem gives a compact subset $K \subseteq X$ such that $\mu(K) > 1 - \eta$ and $\varphi|_K$ is continuous. Then Lemma 2.2 gives $L < \infty$ such that $\varphi|_K$ is $(\eta/3)$ -almost L -Lipschitz.

Now choose ε so small that $L\varepsilon < \eta/3$, and let $U := B_\varepsilon(K)$. Certainly $\mu(U) > 1 - \eta$. Let $\xi : X \rightarrow K$ be a Borel map such that $\xi|_K = \text{id}_K$ and $d_X(x, \xi(x)) < \varepsilon$ for all $x \in U$, and define $\psi := \varphi \circ \xi$. This gives $\varphi(x) = \psi(x)$ for all $x \in K$, and hence

$$\int d_Y(\varphi(x), \psi(x)) \mu(dx) \leq \mu(X \setminus K) < \eta.$$

Finally, for any $x, x' \in U$, we have $\xi(x), \xi(x') \in K$, and therefore

$$\begin{aligned} d_Y(\psi(x), \psi(x')) &= d_Y(\varphi(\xi(x)), \varphi(\xi(x'))) \leq \eta/3 + Ld_X(\xi(x), \xi(x')) \\ &\leq \eta/3 + 2L\varepsilon + Ld_X(x, x') < \eta + Ld_X(x, x'). \end{aligned}$$

So $\psi|_U$ is η -almost L -Lipschitz. □

REMARK. The key difference between Lusin's theorem itself and Lemma 2.5 is that ψ is almost Lipschitz on an *open* set of large measure. This tweak will be important for some applications of the Portmanteau theorem later. ◁

2.4. Covering and packing numbers. If (X, d) is a metric space and $\delta > 0$, then a subset $F \subseteq X$ is δ -**separated** if any distinct $x, y \in F$ satisfy $d(x, y) \geq \delta$. The δ -**covering** and δ -**packing numbers** of the space are defined by

$$\text{cov}_\delta(X, d) := \min\{|F| : B_\delta(F) = X\}$$

and

$$\text{pack}_\delta(X, d) := \max\{|F| : F \text{ is } \delta\text{-separated in } (X, d)\},$$

where either value may be $+\infty$. More generally, if $Y \subseteq X$ then we abbreviate

$$\text{cov}_\delta(Y, d|_{Y \times Y}) =: \text{cov}_\delta(Y, d) \quad \text{and} \quad \text{pack}_\delta(Y, d|_{Y \times Y}) =: \text{pack}_\delta(Y, d).$$

These definitions lead quickly to the standard inequalities

$$\text{cov}_{\delta/2}(Y, d) \geq \text{pack}_\delta(Y, d) \geq \text{cov}_\delta(Y, d) \quad \forall \delta > 0. \quad (5)$$

Now suppose in addition that $\mu \in \text{Prob}(X)$. For $\varepsilon, \delta > 0$, the (ε, δ) -**covering number of μ according to d** is

$$\text{cov}_{\varepsilon, \delta}(\mu, d) := \min\{\text{cov}_{\delta}(V, d) : V \subseteq X \text{ such that } \mu(V) > 1 - \varepsilon\}.$$

Similarly, the (ε, δ) -**packing number of μ according to d** is

$$\text{pack}_{\varepsilon, \delta}(\mu, d) := \min\{\text{pack}_{\delta}(V, d) : V \subseteq X \text{ such that } \mu(V) > 1 - \varepsilon\}.$$

From these definitions, the inequalities (5) translate immediately into

$$\text{cov}_{\varepsilon, \delta/2}(\mu, d) \geq \text{pack}_{\varepsilon, \delta}(\mu, d) \geq \text{cov}_{\varepsilon, \delta}(\mu, d) \quad \forall \varepsilon, \delta > 0. \quad (6)$$

Now let (X, d_X) and (Y, d_Y) be metric spaces and let μ and ν be Borel probabilities on X and Y , respectively. Assume that X and Y are separable, so that the Borel σ -algebra of $X \times Y$ agrees with the product of their separate Borel σ -algebras. Let d be the **Hamming average** on $X \times Y$ of the metrics d_X and d_Y :

$$d((x, y), (x', y')) := \frac{1}{2}d_X(x, x') + \frac{1}{2}d_Y(y, y').$$

The following is quite close to standard results for covering and packing numbers, but the switch to the setting of probability measures requires a little extra work.

LEMMA 2.6. *The following hold for all $\varepsilon, \delta > 0$:*

(i) *any coupling λ of μ and ν satisfies*

$$\text{cov}_{\varepsilon, \delta}(\lambda, d) \leq \text{cov}_{\varepsilon/2, \delta}(\mu, d_X) \cdot \text{cov}_{\varepsilon/2, \delta}(\nu, d_Y);$$

(ii) *the product measure satisfies*

$$\text{pack}_{\varepsilon, \delta/2}(\mu \times \nu, d) \geq \text{pack}_{\sqrt{\varepsilon}, \delta}(\mu, d_X) \cdot \text{pack}_{\sqrt{\varepsilon}, \delta}(\nu, d_Y).$$

Proof. (i) If $E \subseteq X$ and $F \subseteq Y$ are such that $\mu(B_{\delta}(E)) > 1 - \varepsilon/2$ and $\nu(B_{\delta}(F)) > 1 - \varepsilon/2$ (where these neighbourhoods are taken according to the metrics d_X and d_Y , respectively), then $B_{\delta}(E \times F) \supseteq B_{\delta}(E) \times B_{\delta}(F)$, and so

$$\begin{aligned} \lambda((X \times Y) \setminus B_{\delta}(E \times F)) &\leq \lambda((X \times Y) \setminus (B_{\delta}(E) \times B_{\delta}(F))) \\ &\leq \lambda((X \setminus B_{\delta}(E)) \times Y) + \lambda(X \times (Y \setminus B_{\delta}(F))) \\ &= \mu(X \setminus B_{\delta}(E)) + \nu(Y \setminus B_{\delta}(F)) \\ &< 2(\varepsilon/2) = \varepsilon. \end{aligned}$$

(ii) Let $m := \text{pack}_{\sqrt{\varepsilon}, \delta}(\mu, d_X)$ and $n := \text{pack}_{\sqrt{\varepsilon}, \delta}(\nu, d_Y)$.

Suppose that $V \subseteq X \times Y$ has $(\mu \times \nu)(V) > 1 - \varepsilon$. For each $x \in X$ let

$$V_x := \{y \in Y : (x, y) \in V\}.$$

Fubini's theorem and Chebyshev's inequality imply that the set

$$U := \{x : \nu(V_x) > 1 - \sqrt{\varepsilon}\}$$

has $\mu(U) > 1 - \sqrt{\varepsilon}$. The latter inequality implies that there is some $F \subseteq U$ with cardinality m and which is δ -separated according to d_X . On the other hand, for each $x \in F$, the inequality $\nu(V_x) > 1 - \sqrt{\varepsilon}$ implies that there is some $E_x \subseteq V_x$ with cardinality n and which is δ -separated according to d_Y . Now the set of pairs

$$\{(x, y) : x \in F \text{ and } y \in E_x\}$$

is contained in V , has cardinality nm , and is $(\delta/2)$ -separated according to d . \square

Combining Lemma 2.6 with the inequalities in (6) gives the following.

COROLLARY 2.7. *For all $\varepsilon, \delta > 0$ one has*

$$\text{cov}_{\varepsilon, \delta/4}(\mu \times \nu, d) \geq \text{cov}_{\sqrt{\varepsilon}, \delta}(\mu, d_X) \cdot \text{cov}_{\sqrt{\varepsilon}, \delta}(\nu, d_Y). \quad \square$$

The next lemma gives some control over the covering numbers of the pushforward of a measure under an almost-Lipschitz map. After that, Lemma 2.9 shows that the covering numbers of pushforward measures are somewhat stable under small changes to the maps. These results will be used for the proofs that our new entropy notions are isomorphism-invariant in Subsection 6.1.

LEMMA 2.8. *Let (X, d_X) and (Y, d_Y) be metric spaces, let $\mu \in \text{Prob}(X)$, and let $\varphi : X \rightarrow Y$ be η -almost L -Lipschitz. Then*

$$\text{cov}_{\varepsilon, \eta+L\delta}(\varphi_*\mu, d_Y) \leq \text{cov}_{\varepsilon, \delta}(\mu, d_X).$$

Proof. Choose $F \subseteq X$ such that $\mu(B_\delta(F)) > 1 - \varepsilon$. Since φ is η -almost L -Lipchitz, we have

$$B_{\eta+L\delta}(\varphi(F)) \supseteq \varphi(B_\delta(F)),$$

and hence

$$\varphi^{-1}(B_{\eta+L\delta}(\varphi(F))) \supseteq B_\delta(F).$$

Therefore

$$(\varphi_*\mu)(B_{\eta+L\delta}(\varphi(F))) = \mu(\varphi^{-1}(B_{\eta+L\delta}(\varphi(F)))) \geq \mu(B_\delta(F)) > 1 - \varepsilon. \quad \square$$

LEMMA 2.9. Let (X, μ) be a probability space, let (Y, d) be a separable metric space, and let $\varphi, \psi : X \rightarrow Y$ be measurable functions. Also let $\varepsilon, \varepsilon', \delta > 0$. If

$$\mu\{d(\varphi(\cdot), \psi(\cdot)) > \delta/2\} < \varepsilon'$$

then

$$\text{cov}_{\varepsilon+\varepsilon', \delta}(\varphi_*\mu, d) \leq \text{cov}_{\varepsilon, \delta/2}(\psi_*\mu, d).$$

As before, the separability of Y ensures that the Borel σ -algebra of $Y \times Y$ agrees with the product σ -algebra, and hence that d is measurable with respect to the latter. This is needed for the first displayed inequality above to make sense.

Proof. Let $U := \{d(\varphi(\cdot), \psi(\cdot)) \leq \delta/2\}$, and let $F \subseteq Y$ be such that

$$\mu\{x : \psi(x) \in B_{\delta/2}(F)\} = \psi_*\mu(B_{\delta/2}(F)) > 1 - \varepsilon.$$

Then

$$\begin{aligned} \varphi_*\mu(B_\delta(F)) &\geq \mu\{x : x \in U \text{ and } \varphi(x) \in B_\delta(F)\} \\ &\geq \mu\{x : x \in U \text{ and } \psi(x) \in B_{\delta/2}(F)\} > 1 - \varepsilon - \varepsilon'. \quad \square \end{aligned}$$

3. Empirical distributions, good models, and sofic entropy

3.1. Definitions. Suppose that G is a countable discrete group. Suppose further that V is a finite set and that $\sigma : G \rightarrow \text{Sym}(V)$ is any map. Think of this σ as an ‘attempt’ at a representation of G by permutations of V . Given $g, h \in G$ and $v \in V$, it may not be the case that

$$\sigma^g(\sigma^h(v)) = \sigma^{gh}(v). \quad (7)$$

The ‘quality’ of σ as an attempt at a representation can be quantified by the number of v at which (7) holds, say for some finite list of groups elements g, h of interest.

A **sofic approximation** to G is a sequence of finite sets V_n and maps

$$\sigma_n : G \rightarrow \text{Sym}(V_n), \quad n \geq 1,$$

such that

$$[\sigma_n^g(\sigma_n^h(v)) = \sigma_n^{gh}(v) \quad \text{w.h.p. in } v] \quad \forall g, h \in G \quad (8)$$

and

$$[\sigma_n^g(v) \neq v \quad \text{w.h.p. in } v] \quad \forall g \in G \setminus \{e_G\}, \quad (9)$$

both as $n \rightarrow \infty$. Note the order of the quantifiers: we certainly do *not* ask that

$$[\sigma_n^g(\sigma_n^h(v)) = \sigma_n^{gh}(v) \quad \forall g, h \in G] \quad \text{w.h.p. in } v.$$

The group itself is **sofic** if it has a sofic approximation. This definition essentially follows Weiss [24] and (with a different nomenclature) Gromov [11].

Now consider also a compact metric space (\mathcal{X}, d) , and let V and σ be as above. Elements of \mathcal{X}^V will be denoted by boldface letters, to distinguish them from elements of shift spaces such as \mathcal{X}^G . For $\mathbf{x} = (x_v)_{v \in V} \in \mathcal{X}^V$ and $v \in V$, the **pullback name of \mathbf{x} at v** is defined by

$$\Pi_v^\sigma(\mathbf{x}) := (x_{\sigma^g(v)})_{g \in G} \in \mathcal{X}^G.$$

This defines a map $\Pi_v^\sigma : \mathcal{X}^V \rightarrow \mathcal{X}^G$ for each v . Let S be the right-shift action of G on \mathcal{X}^G . Properties (8) and (9) have the following simple but important consequence.

LEMMA 3.1. *If $(\sigma_n)_{n \geq 1}$ is a sofic approximation to G , $F \subseteq G$ is finite, and $g \in G$, then the following holds w.h.p. in $v \in V_n$:*

$$\Pi_{\sigma_n^g(v)}^{\sigma_n}(\cdot)|_F = (S^g(\Pi_v^{\sigma_n}(\cdot)))|_F.$$

More fully, this conclusion asserts that

$$[\Pi_{\sigma_n^g(v)}^{\sigma_n}(\mathbf{x})|_F = \Pi_v^{\sigma_n}(\mathbf{x})|_{Fg} \quad \forall \mathbf{x} \in \mathcal{X}^{V_n}] \quad \text{w.h.p. in } v,$$

where we identify elements of \mathcal{X}^{Fg} with elements of \mathcal{X}^F in the obvious way.

Proof. It suffices to prove this when F is an arbitrary singleton, say $\{h\}$. Then it holds w.h.p. in v that

$$\sigma_n^h(\sigma_n^g(v)) = \sigma_n^{hg}(v).$$

If v satisfies this, and $\mathbf{x} \in \mathcal{X}^{V_n}$, then

$$(\Pi_{\sigma_n^g(v)}^{\sigma_n}(\mathbf{x}))_h = x_{\sigma_n^h(\sigma_n^g(v))} = x_{\sigma_n^{hg}(v)} = (\Pi_v^{\sigma_n}(\mathbf{x}))_{hg}. \quad \square$$

REMARK. The map σ also gives rise to an ‘adjoint’ map $\rho : G \rightarrow \text{Sym}(\mathcal{X}^V)$:

$$\rho^g((x_v)_{v \in V}) := (x_{\sigma_{g^{-1}}(v)})_{v \in V}.$$

Similarly to the proof above, one can show that, if $(\sigma_n)_{n \geq 1}$ is a sofic approximation, then for most v the pullback-name map $\Pi_v^{\sigma_n}$ approximately intertwines ρ_n^g with the left-shift action \tilde{S} of G on \mathcal{X}^G , defined by

$$\tilde{S}^g((x_h)_{h \in G}) := (x_{g^{-1}h})_{h \in G}.$$

This observation does not seem to be useful unless our measure μ on \mathcal{X}^G is invariant under \tilde{S} as well as S . In that case, each \tilde{S}^g defines an isomorphism from the system (\mathcal{X}^G, μ, S) to itself, and the above relationship to ρ_n^g becomes a special case of a result for general factor maps; see Lemma 4.8 below. \triangleleft

Each $\mathbf{x} \in \mathcal{X}^{V_n}$ has an associated probability measure on \mathcal{X}^G called its **empirical distribution**:

$$P_{\mathbf{x}}^\sigma := \frac{1}{|V|} \sum_{v \in V} \delta_{\Pi_v^\sigma(\mathbf{x})}. \tag{10}$$

Lemma 3.1 will mostly be used through the following consequence, which asserts an approximate invariance for empirical distributions.

LEMMA 3.2. *Let $F \subseteq G$ be finite and $g \in G$. Then*

$$\sup_{\mathbf{x} \in \mathcal{X}^{V_n}} \|(P_{\mathbf{x}}^{\sigma_n})_F - (S_*^g P_{\mathbf{x}}^{\sigma_n})_F\|_{\text{TV}} \longrightarrow 0 \quad \text{as } n \longrightarrow \infty.$$

Proof. Lemma 3.1 gives that

$$(S^g \Pi_v^{\sigma_n}(\cdot))|_F = \Pi_{\sigma_n^g v}^{\sigma_n}(\cdot)|_F \quad \text{w.h.p. in } v,$$

and therefore

$$\begin{aligned} (P_{\mathbf{x}}^{\sigma_n})_F - (S_*^g P_{\mathbf{x}}^{\sigma_n})_F &= \frac{1}{|V_n|} \sum_{v \in V_n} (\delta_{\Pi_v^{\sigma_n}(\mathbf{x})|_F} - \delta_{(S^g \Pi_v^{\sigma_n}(\mathbf{x}))|_F}) \\ &= \frac{1}{|V_n|} \sum_{v \in V_n} (\delta_{\Pi_v^{\sigma_n}(\mathbf{x})|_F} - \delta_{\Pi_{\sigma_n^g v}^{\sigma_n}(\mathbf{x})|_F}) + o(1). \end{aligned}$$

But the sum on this last line vanishes, because σ_n^g is a permutation of V_n . \square

Now suppose that μ is a right-shift-invariant probability measure on \mathcal{X}^G . For any $F \subseteq G$, let μ_F denote the marginal of μ on X^F , as previously.

For G and σ as above, and for any $\mathcal{O} \subseteq \text{Prob}(\mathcal{X}^G)$, let

$$\Omega(\mathcal{O}, \sigma) := \{\mathbf{x} \in \mathcal{X}^V : P_{\mathbf{x}}^\sigma \in \mathcal{O}\}. \tag{11}$$

In case \mathcal{O} is a w^* -neighbourhood of μ , elements of $\Omega(\mathcal{O}, \sigma)$ are called **\mathcal{O} -good models** for μ over σ .

This definition is easiest to visualize in case \mathcal{O} is of the form $\{v : v_F \in \mathcal{O}'\}$ for some finite $F \subseteq G$ and some w^* -neighbourhood \mathcal{O}' of μ_F . Neighbourhoods \mathcal{O}

of this form are a basis of neighbourhoods around μ , so this assumption does not lose much generality. For this \mathcal{O} we can write

$$\Omega(\mathcal{O}, \sigma_n) = \left\{ \mathbf{x} \in \mathcal{X}^{V_n} : \frac{1}{|V_n|} \sum_{v \in V_n} \delta_{\Pi_v^{\sigma_n}(\mathbf{x})|_F} \in \mathcal{O}' \right\}.$$

Given the finite set F , it follows from conditions (8) and (9) that the mapping

$$g \mapsto \sigma_n^g(v)$$

defines a map from F to a corresponding ‘window’ $\sigma_n^F(v)$ around v in V_n , and this map is a bijection w.h.p. in v . Consequently, one may regard the map $\mathbf{x} \mapsto \Pi_v^{\sigma_n}(\mathbf{x})|_F$ as restricting \mathbf{x} to $\mathbf{x}|_{\sigma_n^F(v)}$, and then making a copy of this restriction indexed by F itself. The good models $\Omega(\mathcal{O}, \sigma_n)$ are those \mathbf{x} such that, on average over v , the frequency with which one sees a particular restriction through this window is close to the F -marginal of μ itself. This is the sense in which \mathbf{x} is ‘modelling’ μ .

In many examples \mathcal{X} is a finite alphabet. In that case one can instead work with total-variation neighbourhoods of finite-dimensional marginals of μ . Specifically, if \mathcal{X} is finite then the (F, ε) -good models for μ over σ are the elements of

$$\Omega_\mu(F, \varepsilon, \sigma) := \left\{ \mathbf{x} \in \mathcal{X}^V : \left\| \frac{1}{|V|} \sum_{v \in V} \delta_{\Pi_v^\sigma(\mathbf{x})|_F} - \mu_F \right\|_{\text{TV}} < \varepsilon \right\}.$$

Finally, as in the Introduction, we define the sofic entropy of a metric G -process $(\mathcal{X}^G, \mu, S, d)$ to be

$$h_\Sigma(\mu) := \sup_{\delta > 0} \inf_{\mathcal{O}} \limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{cov}_\delta(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)}),$$

where \mathcal{O} ranges over w^* -neighbourhoods of μ . Similarly, the lower sofic entropy is

$$\underline{h}_\Sigma(\mu) := \sup_{\delta > 0} \inf_{\mathcal{O}} \liminf_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{cov}_\delta(\Omega(\mathcal{O}, \sigma_n), d^{(V_n)}).$$

We do not record d in the notation for these quantities because it turns out that they depend only on the measure-theoretic structure of the process (\mathcal{X}^G, μ, S) , as will be shown in the next subsection.

Examples in which $\underline{h}_\Sigma(\mu) < h_\Sigma(\mu)$ can be obtained from examples in which $h_\Sigma(\mu) \neq h_{\Sigma'}(\mu)$ for two different sofic approximations Σ and Σ' to G . Interleaving Σ and Σ' into a single sofic approximation then gives the former inequality. For instance, [5, Subsection 8.3] includes examples of free-group Markov chains whose f -invariant is finite and negative. Since (i) sofic entropies

can take values only in $\{-\infty\} \cup [0, +\infty]$, and (ii) the f-invariant can be expressed as a kind of average of sofic entropies over randomly chosen sofic approximations [3], it follows that these systems must have some sofic approximations which give nonnegative real values for the sofic entropy, and others which give $-\infty$.

However, it is an important open problem whether one can obtain two different finite real values for $h_\Sigma(\mu)$ and $h_{\Sigma'}(\mu)$ using two different sofic approximations.

3.2. Agreement with the Kerr–Li definition. Bowen defined sofic entropy for systems with finite generating partitions in [4]. Kerr and Li give a new definition in [15] in terms of approximate homomorphism between commutative von Neumann algebras, and showed that it gives the same values as Bowen's if there is a finite generating partition. Then, in [16, Section 3], Kerr and Li gave another, more elementary definition of general sofic entropy, and proved its equivalence to their previous definition.

In this subsection we show that our definition gives the same values as the entropy of [16, Definition 3.3]. We refer to that paper for a careful introduction to their definition. If (X, μ, T) is a G -system, let us write $\tilde{h}_\Sigma(\mu, T)$ for the entropy defined there.

It suffices to show that any metric G -process $(\mathcal{X}^G, \mu, S, d)$ satisfies

$$h_\Sigma(\mu) = \tilde{h}_\Sigma(\mu, S).$$

That is, rather than analyse a G -action on an arbitrary probability space (X, μ) , we may restrict our attention to the case of processes, so $X = \mathcal{X}^G$ for another standard measurable space \mathcal{X} which is equipped with a particular compact generating metric d . We may also assume that d has diameter at most 1.

On $X = \mathcal{X}^G$, define the pseudometric

$$\rho(x, y) = d(x_e, y_e) \quad \text{for } x = (x_g)_g, y = (y_g)_g \in \mathcal{X}^G.$$

Let $\Sigma = (\sigma_n: G \rightarrow \text{Sym}(V_n))_{n \geq 1}$ be the sofic approximation. If F is a finite subset of G , L is a finite subset of $C(X)$, and $\varepsilon > 0$, then define $\text{Map}_\mu(\rho, F, L, \varepsilon, \sigma_n)$ to be the set of those $x \in X^{V_n}$ such that:

(i) we have

$$\sqrt{\frac{1}{|V_n|} \sum_{v \in V_n} \rho(x_{\sigma_n^g(v)}, S^g(x_v))^2} < \varepsilon \quad \forall g \in F;$$

(ii) and

$$\frac{1}{V_n} \sum_{v \in V_n} f(x_v) \approx_\varepsilon \int f d\mu \quad \forall f \in L.$$

The pseudometric ρ is clearly dynamically generating (see [16, Section 2] or [17, Section 4]). Therefore [16, Proposition 3.4] gives that

$$\tilde{h}_\Sigma(\mu, S) = \sup_{\delta > 0} \inf_{F, L, \varepsilon} \limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{pack}_\delta(\text{Map}_\mu(\rho, F, L, \varepsilon, \sigma_n), \rho^{2, (V_n)}), \quad (12)$$

where F , L and ε are as above, and $\rho^{2, (V_n)}$ is the ℓ_2 -analogue of the Hamming-average metric on X^{V_n} :

$$\rho^{2, (V_n)}(x, y) := \sqrt{\frac{1}{|V_n|} \sum_{v \in V_n} \rho(x_v, y_v)^2} \quad \text{for } (x_v)_v, (y_v)_v \in X^{V_n}.$$

(I have adjusted some of the notation from [16] to match the present paper.)

PROPOSITION 3.3. *In the setting above we have $h_\Sigma(\mu) = \tilde{h}_\Sigma(\mu, S)$.*

Proof. Step 1. Let

$$\rho^{1, (V_n)}(x, y) := \frac{1}{|V_n|} \sum_{v \in V_n} \rho(x_v, y_v) \quad \text{for } (x_v)_v, (y_v)_v \in X^{V_n}.$$

Since ρ is bounded by 1, we have

$$\rho^{1, (V_n)} \leq \rho^{2, (V_n)} \leq \sqrt{\rho^{1, (V_n)}}.$$

Since the right-hand side of (12) takes a supremum over $\delta > 0$, we may therefore replace $\rho^{2, (V_n)}$ with $\rho^{1, (V_n)}$ in that equation. Also, recalling the inequalities (5), the same reasoning lets us replace pack_δ with cov_δ in (12).

Step 2. For each n , define

$$\Phi_n : X^{V_n} \longrightarrow \mathcal{X}^{V_n} : ((x_{v,g})_{g \in G})_{v \in V_n} \mapsto (x_{v,e})_{v \in V_n}.$$

This is an isometry from the pseudometric $\rho^{1, (V_n)}$ to the true metric $d^{(V_n)}$.

Now suppose that \mathcal{O} is a w^* -neighbourhood of μ . Then there are a finite $L \subseteq C(X)$, say consisting of $[0, 1]$ -valued functions, and $\varepsilon_1 > 0$ such that

$$\mathcal{O} \supseteq \mathcal{O}' := \left\{ \nu \in \text{Prob}(X) : \int f d\nu \approx_{3\varepsilon_1} \int f d\mu \quad \forall f \in L \right\}.$$

Since L is finite, there are a finite set $F \subseteq G$ and an $\varepsilon_2 > 0$ such that, for any $x, y \in X$, we have

$$[d(x_g, y_g) < \varepsilon_2 \quad \forall g \in F] \implies [f(x) \approx_{\varepsilon_1} f(y) \quad \forall f \in L].$$

Given ε_2 , we may now choose $\varepsilon \in (0, \varepsilon_1)$ so small that the following holds. If

$$x = ((x_{v,g})_{g \in G})_{v \in V_n} \in \text{Map}_\mu(\rho, F, L, \varepsilon, \sigma_n),$$

then condition (i) implies that

$$|\{v \in V_n : d(x_{\sigma_n^g(v),e}, x_{v,g}) < \varepsilon_2 \forall g \in F\}| > (1 - \varepsilon_1)|V_n|.$$

It follows that

$$\int f dP_{\Phi_n(x)}^{\sigma_n} = \frac{1}{|V_n|} \sum_{v \in V_n} f((x_{\sigma_n^g(v),e})_{g \in G}) \approx_{2\varepsilon_1} \frac{1}{|V_n|} \sum_{v \in V_n} f((x_{v,g})_{g \in G}),$$

and now condition (ii) gives that this is within ε_1 of $\int f d\mu$. This shows that

$$\Phi_n(\text{Map}_\mu(\rho, F, L, \varepsilon, \sigma_n)) \subseteq \Omega(\mathcal{O}', \sigma_n) \subseteq \Omega(\mathcal{O}, \sigma_n).$$

Since Φ_n is an isometry, it preserves covering numbers. Taking infima over \mathcal{O} (on the right-hand side) or F, L and ε (on the left-hand side), this shows that

$$\tilde{h}_\Sigma(\mu, S) \leq h_\Sigma(\mu).$$

Step 3. The proof of the reverse inequality is very similar. Now we define

$$\Psi_n : \mathcal{X}^{V_n} \longrightarrow X^{V_n} : \mathbf{x} \mapsto (\Pi_v^{\sigma_n}(\mathbf{x}))_{v \in V_n},$$

which is an isometry from $d^{(V_n)}$ to $\rho^{1,(V_n)}$. For any F, L , and $\varepsilon > 0$, there is a w^* -neighbourhood \mathcal{O} of μ such that

$$\Psi_n(\Omega(\mathcal{O}, \sigma_n)) \subseteq \text{Map}_\mu(\rho, F, L, \varepsilon, \sigma_n).$$

This \mathcal{O} can be obtained by reversing the construction in Step 2. The only point worth remarking is that condition (ii) in the definition of $\text{Map}_\mu(\rho, F, L, \varepsilon, \sigma_n)$ is obtained as a consequence of Lemma 3.1.

Taking infima, this leads to

$$h_\Sigma(\mu) \leq \tilde{h}_\Sigma(\mu, S). \quad \square$$

Since Kerr and Li have already proved isomorphism invariance of sofic entropy using their definitions (see [15]), that property follows for ours. This justifies the omission of d from the notation $h_\Sigma(\mu)$. A proof of this invariance using our definition would be similar to the proof that model-measure sofic entropy is isomorphism-invariant, which is given in Subsection 6.1 below.

REMARK. The proof of Proposition 3.3 uses a particular generating pseudometric on $X = \mathcal{X}^G$, obtained from a metric on the single coordinate space \mathcal{X} . However, it turns out that for any abstract G -system (X, μ, T) , any totally bounded and Borel measurable pseudometric ρ on X arises in this way, up to isomorphism. Indeed, letting (\mathcal{X}_0, d_0) be the quotient of X by the relation $\{\rho = 0\}$, and letting (\mathcal{X}, d) be the completion of (\mathcal{X}_0, d_0) , it is easy to check that the quotient map $q : X \rightarrow \mathcal{X}$ is Borel, and now this can be extended to a factor map $q^G : X \rightarrow \mathcal{X}^G$ similarly to (1). If ρ is dynamically generating, then q^G is an isomorphism, and ρ is now obtained by pulling back the identity-coordinate metric d through this isomorphism. \triangleleft

3.3. Subadditivity and failure of additivity.

PROPOSITION 3.4. *Let (\mathcal{X}^G, μ, S) and (\mathcal{Y}^G, ν, S) be two G -processes, and let λ be a joining of them. Then*

$$h_S(\lambda) \leq h_S(\mu) + h_S(\nu).$$

Proof. Let d_X and d_Y be compact generating metrics on \mathcal{X} and \mathcal{Y} respectively, and let d be their Hamming average on $\mathcal{X} \times \mathcal{Y}$. All subsequent topologies are those determined by these metrics.

For any w^* -neighbourhood \mathcal{O}_1 of μ and \mathcal{O}_2 of ν , there is a w^* -neighbourhood \mathcal{N} of λ such that every $\theta \in \mathcal{N}$ has first marginal in \mathcal{O}_1 and second marginal in \mathcal{O}_2 . This implies that

$$\Omega(\mathcal{N}, \sigma_n) \subseteq \Omega(\mathcal{O}_1, \sigma_n) \times \Omega(\mathcal{O}_2, \sigma_n) \quad \forall n \geq 1,$$

and now the inequality

$$\text{cov}_\delta(\Omega(\mathcal{N}, \sigma_n), d^{(V_n)}) \leq \text{cov}_\delta(\Omega(\mathcal{O}_1, \sigma_n), d_X^{(V_n)}) \cdot \text{cov}_\delta(\Omega(\mathcal{O}_2, \sigma_n), d_Y^{(V_n)})$$

completes the proof. \square

In particular,

$$h_S(\mu \times \nu) \leq h_S(\mu) + h_S(\nu).$$

We now describe examples in which this inequality is strict.

EXAMPLE 3.5. Some standard probabilistic estimates are required to justify this example carefully, but we omit these for brevity.

Let $H = \langle a, b \rangle$ be the free group on two generators, let $H' = \langle a', b' \rangle$ be a copy of H , and let $G = H * H'$. Then G is a free group on four generators, and we may regard H as a subgroup of G .

Let T_0 be the trivial H -action on the set $\mathcal{X} = \{0, 1\}$, and endow this set with its discrete metric. Let

$$\mu_0 := \frac{3}{4}\delta_0 + \frac{1}{4}\delta_1 \in \text{Prob}(\mathcal{X}).$$

(It will be clear in what follows that $\frac{3}{4}$ could be replaced with any value in $(1/2, 1)$.)

Now co-induce T_0 to the G -action $\text{CInd}_H^G T_0$ on the space

$$(X, \mu) := (\mathcal{X}^{H \setminus G}, \mu_0^{\times H \setminus G}).$$

This co-induced system is isomorphic to a G -process (\mathcal{X}^G, ν, S) where ν is defined by the following three properties:

- every one-dimensional marginal of ν equals μ_0 ;
- if $Hg = Hg'$ for some $g, g' \in G$, then $x_g = x_{g'}$ for ν -a.e. x ;
- if the cosets Hg_1, \dots, Hg_k are distinct, and x is drawn at random from ν , then the coordinates x_{g_1}, \dots, x_{g_k} are independent.

See [13, Subsection II.10.(G)] or [9] for the definition and basic properties of co-induction.

Now for each n let $U_n := \{1, \dots, 3n\}$, $W_n := \{3n + 1, \dots, 4n\}$ and $V_n := U_n \cup W_n$. Choose four elements of $\text{Sym}(V_n)$ in the following randomized way:

- Let $\sigma_n^{a'}$ and $\sigma_n^{b'}$ be independent, uniformly random elements of $\text{Sym}(V_n)$.
- Let $\tau_{n,0}^a$ and $\tau_{n,0}^b$ be uniformly random elements of $\text{Sym}(U_n)$ and let $\tau_{n,1}^a$ and $\tau_{n,1}^b$ be uniformly random elements of $\text{Sym}(W_n)$, all independent. Let $\sigma_n^a := \tau_{n,0}^a \cup \tau_{n,1}^a$ and $\sigma_n^b := \tau_{n,0}^b \cup \tau_{n,1}^b$. Thus, σ_n^a and σ_n^b are chosen uniformly and independently from among those elements of $\text{Sym}(V_n)$ that preserve the partition $\{U_n, W_n\}$.

For each n , the four permutations $\sigma_n^a, \sigma_n^b, \sigma_n^{a'}$ and $\sigma_n^{b'}$ generate a random homomorphism $\sigma_n : G \rightarrow \text{Sym}(V_n)$. Standard arguments show that the resulting sequence $\Sigma = (\sigma_n)_{n \geq 1}$ is a sofic approximation to G with high probability.

For each n , let $\mathbf{x}_n \in \mathcal{X}^{V_n}$ be the indicator function 1_{W_n} . It is now easily checked that $P_{\mathbf{x}_n}^{\sigma_n} \xrightarrow{\text{weak}^*} \nu$, and so for every w^* -neighbourhood \mathcal{O} of ν we have $\Omega(\mathcal{O}, \sigma_n) \neq \emptyset$ for all sufficiently large n . Therefore, $h_\Sigma(\nu, S) \geq 0$.

The random Schreier graph on V_n generated by the random permutations σ_n^a and σ_n^b is an expander within each of U_n and W_n with high probability; this follows by the usual counting argument (see, for instance, [18, Proposition 1.2.1]). In this case, any other partition of V_n with small edge boundary in this Schreier graph must be very close to the partition $\{U_n, W_n\}$. This implies that any other good

model $\mathbf{y} \in \mathcal{X}^{V_n}$ of ν must be very close to \mathbf{x}_n in normalized Hamming distance as $n \rightarrow \infty$. It follows that in fact $h_\Sigma(\nu, S) = 0$.

(Note that at this point, we are using the fact that the atom sizes of the measure μ_0 correspond to the ratios $|U_n|/|V_n|$ and $|W_n|/|V_n|$. If these were all equal to $1/2$, instead of $3/4$ and $1/4$, then both \mathbf{x}_n and 1_{U_n} would be good models for ν , and all other good models would lie close to one of these two in normalized Hamming distance. In this case, the rest of the argument below can still be completed, but there would be slightly more work to do.)

However, we can now show that $h_\Sigma(\nu \times \nu, S \times S) = -\infty$. If $(\mathbf{y}, \mathbf{y}') \in (\mathcal{X} \times \mathcal{X})^{V_n}$ were a good model for $\nu \times \nu$, then both \mathbf{y} and \mathbf{y}' would be good models for ν , hence close to \mathbf{x}_n in Hamming distance. But this would imply that

$$P_{(\mathbf{y}, \mathbf{y}')}^{\sigma_n} \{(1, 0)\} = \frac{1}{4n} |\{v \in V_n : y_v = 1 \text{ and } y'_v = 0\}|$$

is close to 0, whereas a good model for $\nu \times \nu$ should have this probability close to $\frac{1}{4} \cdot \frac{3}{4}$. So Σ does not provide arbitrarily good models for $\nu \times \nu$ as $n \rightarrow \infty$.

Note that this argument is really only about the H -subaction of S . The only reason to co-induce to G is to make an example which is ergodic overall and free. It would be interesting to know whether one can produce such an example which is totally ergodic by starting with a more subtle choice of H -system.

If one replaces each Σ with $\Sigma' := (\sigma_n^{\times k} : G \rightarrow V_n^{\times k})_{n \geq 1}$ for some fixed $k \geq 1$, then similar reasoning shows that

$$h_\Sigma(\mu^{\times \ell}, S^{\times \ell}) = \begin{cases} 0 & \text{for } \ell \leq k, \\ -\infty & \text{for } \ell > k. \end{cases} \quad \triangleleft$$

EXAMPLE 3.6. Let us speculate about a second example. The details required for its analysis are not available in full, but it would arguably be more natural than Example 3.5.

Let $H = \langle a, b \rangle$ be the free group and let $\mathcal{X} = \{0, 1\}$, as above. We start by constructing some finite quotients of H as a variant of the ‘planted bisection model’. This classical model has a long history in statistical physics and computer science; see [21] for its definition and some references.

Let $\alpha \in (0, 1)$ be a small parameter. Let $V_n = U_n \cup W_n$ be as in Example 3.5, but now construct $\sigma_n^a, \sigma_n^b \in \text{Sym}(V_n)$ as follows.

First let Γ_n be a random 4-regular graph on V_n drawn uniformly from those graphs that have roughly $6n(1 - \alpha)$ edges within U_n , $2n(1 - \alpha)$ edges within W_n , and $8n\alpha$ edges between U_n and W_n .

Using this random graph Γ_n , one can construct the pair of permutations σ_n^a and σ_n^b as follows. First, a simple greedy algorithm finds a disjoint union of

cycles in Γ_n that contains all vertices in V_n . Choose an orientation for each of these cycles, and let those directed edges define the permutation σ_n^a . After removing these edges from Γ_n , the remaining 2-regular graph decomposes into another disjoint union of cycles; orienting those gives the permutation σ_n^b . Let $\sigma_n : G \rightarrow \text{Sym}(V_n)$ be the homomorphism generated by σ_n^a and σ_n^b . Now Γ_n is the Schreier graph of the homomorphism σ_n and generating set $\{a^{\pm 1}, b^{\pm 1}\}$.

As with other simple random-graph models, it should hold that Γ_n looks like a tree in a large neighbourhood around most points of V_n , and this would imply that $\Sigma = (\sigma_n)_{n \geq 1}$ is a sofic approximation to H .

Finally, let $\mathbf{x}_n = 1_{W_n}$ as in Example 3.5, and now use w^* -compactness to choose a subsequence $n_1 < n_2 < \dots$ such that $P_{\mathbf{x}_{n_i}}^{\sigma_{n_i}} \xrightarrow{\text{weak}^*} \mu$ for some $\mu \in \text{Prob}(\mathcal{X}^H)$. Of course, this guarantees that for any w^* -neighbourhood \mathcal{O} of μ we have $\mathbf{x}_{n_i} \in \mathcal{O}(\mathcal{O}, \sigma_{n_i})$ for all sufficiently large i , and so $h_\Sigma(\mu) \geq 0$.

On the other hand, our intuition is that, if α is extremely small, so the graph Γ_n has sufficiently few of its edges crossing from U_n to W_n , then any other partition of V_n into subsets of sizes roughly $3n$ and n and with so few edges between must be very close to $\{U_n, W_n\}$ (up to an error depending on α). Some hope for a proof of this is offered by the recent work [21] on the original planted bisection model, which shows that if the two edge densities in the model are sufficiently well separated, then one can reconstruct the values of those edge densities with high probability if one is given only the output graph Γ_n .

If this prediction is correct, then the same argument as for Example 3.5 will show that $h_\Sigma(\mu \times \mu) = -\infty$. ◁

4. Factor maps and maps between model spaces

Before introducing measures on model spaces, we need to consider how a factor map between systems can be approximated by a sequence of somewhat ‘regular’ maps between their model spaces.

This section is rather technical, but it lays essential foundations for many of the arguments that follow. In particular, it is the basis for the proof that h_Σ^q and h_Σ^{dq} are isomorphism-invariant.

4.1. Approximating factor maps. Suppose that $\varphi : \mathcal{X}^G \rightarrow \mathcal{Y}$ is measurable, so it gives rise to the measurable equivariant map

$$\Phi := (\varphi \circ S^g)_{g \in G} : \mathcal{X}^G \rightarrow \mathcal{Y}^G.$$

More generally, suppose that $E, F \subseteq G$ and that $\varphi : \mathcal{X}^E \rightarrow \mathcal{Y}$. Then one defines

$$\varphi^F : \mathcal{X}^{EF} \rightarrow \mathcal{Y}^F : (x_g)_{g \in EF} \mapsto (\varphi((x_{hg})_{h \in E}))_{g \in F};$$

in this notation, $\Phi = \varphi^G$.

In order to study such equivariant maps, we need the ability to approximate φ by a map which is ‘roughly continuous’. This can be done in two steps. The first is to replace φ with a map depending on only finitely many coordinates.

DEFINITION 4.1. If $\varphi : \mathcal{X}^G \rightarrow \mathcal{Y}$ and $D \subseteq G$ is finite, then φ is **D -local** if it is measurable with respect to π_D . A function is **local** if it is **D -local** for some D .

Similarly, a subset $U \subseteq \mathcal{X}^G$ is **D -local** if it equals $\pi_D^{-1}(V)$ for some $V \subseteq \mathcal{X}^D$.

If a function is described as ‘ D -local’, then it is always implied that D is finite.

We now introduce a choice of metrics on \mathcal{X} and \mathcal{Y} . The next definition is a simple adaptation of Definition 2.4.

DEFINITION 4.2. Let $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ be compact metric spaces, let $\mu \in \text{Prob}^S(\mathcal{X}^G)$, let $\varphi : \mathcal{X}^G \rightarrow \mathcal{Y}$ be a measurable function, and let $\eta > 0$. An **η -almost Lipschitz** (or **η -AL**) **approximation to φ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$** is a measurable map $\psi : \mathcal{X}^G \rightarrow \mathcal{Y}$ with the following properties.

(i) The map ψ approximates φ in the sense that

$$\int d_{\mathcal{Y}}(\varphi(x), \psi(x)) \mu(dx) < \eta. \quad (13)$$

(ii) There is a finite $D \subseteq G$ such that ψ is D -local.

(iii) There is a D -local open subset $U \subseteq \mathcal{X}^G$ such that $\mu(U) > 1 - \eta$ and such that $\psi|_U$ is η -almost Lipschitz from $d_{\mathcal{X}}^{(D)}$ to $d_{\mathcal{Y}}$.

In this definition, since ψ and U are both D local, we may regard $\psi|_U$ as a function on \mathcal{X}^D . Part (iii) of the definition can be understood this way, or by considering almost-Lipschitz functions with respect to the pseudometric $d_{\mathcal{X}}^{(D)}$ on \mathcal{X}^G .

Definition 4.2 really does depend on the measure μ and on the specific metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$. However, we may sometimes drop the qualifier ‘rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ ’ when these data are clear from the context.

Formally, the ψ in this definition is a D -local function on the whole space \mathcal{X}^G . We sometimes commit the abuse of writing $\psi(x|_D)$ in place of $\psi(x)$ when the local nature of the function is important.

LEMMA 4.3. *There exist η -AL approximations to φ for all $\eta > 0$.*

Proof. Let $\eta > 0$. Standard measure theory gives some finite $D \subseteq G$ and a measurable function $\varphi' : \mathcal{X}^D \rightarrow \mathcal{Y}$ such that

$$\int d_{\mathcal{Y}}(\varphi(x), \varphi'(x|_D)) \mu(dx) < \eta.$$

Since η is arbitrary, it now suffices to approximate φ' instead of φ ; or, equivalently, to assume that φ itself is a function on \mathcal{X}^D . Having done so, we apply Lemma 2.5 to this map and the metric spaces $(\mathcal{X}^D, d_{\mathcal{X}}^{(D)})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$. \square

In case \mathcal{X} is a finite set, all finite-dimensional Cartesian powers of \mathcal{X} are finite, and so one can use a simplified form of Definition 4.2. In that case, an η -AL approximation to $\varphi : \mathcal{X}^G \rightarrow \mathcal{Y}$ is simply a local map $\psi : \mathcal{X}^G \rightarrow \mathcal{Y}$ such that $\mu\{\varphi = \psi\} > 1 - \eta$.

It will be helpful to know that Definition 4.2 behaves well in relation to Hamming averages. The next lemma describes this.

LEMMA 4.4. *Suppose that $d_{\mathcal{Y}}$ has diameter at most 1. If ψ is an η -AL approximation to φ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ for some $\eta \in (0, 1)$, then ψ^F is a $(3\sqrt{\eta})$ -AL approximation to $\varphi^F : \mathcal{X}^G \rightarrow \mathcal{Y}^F$ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}}^{(F)})$ for every finite $F \subseteq G$.*

Proof. Let ψ be D -local, and in this proof let us regard ψ as a function on \mathcal{X}^D itself. Let $U \subseteq \mathcal{X}^D$ be an open subset with $\mu_D(U) > 1 - \eta$ and such that $\psi|U$ is η -almost L -Lipschitz from $d_{\mathcal{X}}^{(D)}$ to $d_{\mathcal{Y}}$.

Firstly, the shift invariance of μ and inequality (13) imply that

$$\int d_{\mathcal{Y}}^{(F)}(\varphi^F(x), \psi^F(x)) \mu(dx) = \frac{1}{|F|} \sum_{g \in F} \int d_{\mathcal{Y}}(\varphi(S^g x), \psi(S^g x)) \mu(dx) < \eta. \tag{14}$$

Next, it is clear that ψ^F is (DF) -local. Let

$$U_F := \{x \in \mathcal{X}^{DF} : |\{g \in F : x|_{Dg} \notin U\}| < \sqrt{\eta}|F|\}.$$

This set is open, and

$$\int |\{g \in F : x|_{Dg} \notin U\}| \mu(dx) = \sum_{g \in F} \mu\{x : x|_{Dg} \notin U\} = |F| \cdot \mu_D(\mathcal{X}^D \setminus U) < \eta|F|,$$

so Chebyshev's inequality proves that $\mu_{DF}(U_F) > 1 - \sqrt{\eta}$.

Finally, if $x, x' \in U_F$, then

$$\begin{aligned}
 d_{\mathcal{Y}}^{(F)}(\psi^F(x), \psi^F(x')) &= \frac{1}{|F|} \sum_{g \in F} d_{\mathcal{Y}}(\psi(x|_{D_g}), \psi(x'|_{D_g})) \\
 &\leq \frac{|\{g \in F : x|_{D_g} \notin U \text{ or } x'|_{D_g} \notin U\}|}{|F|} \\
 &\quad + \frac{1}{|F|} \sum_{\substack{g \in F, \\ x|_{D_g}, x'|_{D_g} \in U}} (\eta + Ld_{\mathcal{X}}^{(D)}(x|_{D_g}, x'|_{D_g})) \\
 &\leq 3\sqrt{\eta} + L \frac{1}{|D||F|} \sum_{g \in F, h \in D} d_{\mathcal{X}}(x_{hg}, x'_{hg}).
 \end{aligned}$$

Any point of DF can be represented as a product hg with $h \in D$ and $g \in G$ in no more than $|D|$ ways, and so the last line above is at most

$$3\sqrt{\eta} + L|D|d_{\mathcal{X}}^{(DF)}(x, x').$$

This shows that $\psi^F|_{U_F}$ is $(3\sqrt{\eta})$ -almost $(L|D|)$ -Lipschitz. □

Sometimes it is preferable to use the approximation (14) through its consequence that

$$\mu\{x : d_{\mathcal{Y}}^{(F)}(\varphi^F(x), \psi^F(x)) < \sqrt{\eta}\} > 1 - \sqrt{\eta}, \tag{15}$$

which follows by Chebyshev’s inequality.

Now suppose that $\Phi = \varphi^G : (\mathcal{X}^G, \mu, S) \longrightarrow (\mathcal{Y}^G, \nu, S)$ is a factor map, and that $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ are generating compact metrics on \mathcal{X} and \mathcal{Y} with diameter at most 1. If φ is not continuous for the resulting topologies on \mathcal{X}^G and \mathcal{Y} , then $\Phi_* : \text{Prob}(\mathcal{X}^G) \longrightarrow \text{Prob}(\mathcal{Y}^G)$ cannot be w^* -continuous. However, the next lemma shows that, if η is sufficiently small, then an η -AL approximation to φ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ acts approximately continuously on measures which are w^* -close to μ . This fact will be used several times later.

LEMMA 4.5. *For every w^* -neighbourhood \mathcal{N} of ν there is an $\eta > 0$ with the following property. For any η -AL approximation ψ to φ , there is a w^* -neighbourhood \mathcal{O} of μ such that*

$$(\psi^G)_*(\mathcal{O}) \subseteq \mathcal{N}.$$

REMARK. It is very important that η can be chosen depending only on \mathcal{N} . ◁

Proof. It suffices to prove this for any subbasis of w^* -neighbourhoods of ν , so we may assume that

$$\mathcal{N} = \left\{ \theta \in \text{Prob}(\mathcal{Y}^G) : \int h d\theta \approx_\kappa \int h d\nu \right\},$$

where $h : \mathcal{Y}^G \rightarrow [0, 1]$ is F -local for some finite $F \subseteq G$ and is 1-Lipschitz according to $d_{\mathcal{Y}}^{(F)}$, and $\kappa > 0$.

In this case we will show that any $\eta < (\kappa/18)^2$ has the required property. Let ψ be an η -AL approximation to φ , and let $D \subseteq G$ and $U \subseteq \mathcal{X}^G$ be as in Definition 4.2 for this ψ . Let U_F be the (DF) -local open subset of \mathcal{X}^G constructed in the proof of Lemma 4.4, so $\psi^F|_{U_F}$ is $(3\sqrt{\eta})$ -almost Lipschitz according to $d_{\mathcal{X}}^{(DF)}$. Let us abbreviate $3\sqrt{\eta} =: \alpha$.

The composition $h \circ \psi^G$ is (DF) -local. We may regard h as a map $\mathcal{Y}^F \rightarrow \mathbb{R}$ and ψ^F as a map $\mathcal{X}^{DF} \rightarrow \mathcal{Y}^F$. Applying Lemma 2.1, it follows that the restriction $(h \circ \psi^F)|_{U_F}$ is α -almost Lipschitz according to $d_{\mathcal{X}}^{(DF)}$. Therefore, Lemma 2.3 gives a (DF) -local function $f : \mathcal{X}^G \rightarrow [0, 1]$ which is truly Lipschitz according to $d_{\mathcal{X}}^{(DF)}$ and satisfies

$$\|(h \circ \psi^F - f)|_{U_F}\|_\infty \leq \alpha.$$

Let

$$\mathcal{O} := \left\{ \gamma \in \text{Prob}(\mathcal{X}^G) : \gamma(U_F) > 1 - \alpha \text{ and } \int f d\gamma \approx_\alpha \int f d\mu \right\}.$$

Suppose that $\gamma \in \mathcal{O}$. Then

$$\begin{aligned} \left| \int h d(\psi_*^G \gamma) - \int h d\nu \right| &= \left| \int h \circ \psi^G d\gamma - \int h \circ \varphi^G d\mu \right| \\ &\leq \left| \int h \circ \psi^G d\gamma - \int h \circ \psi^G d\mu \right| \\ &\quad + \int |h(\psi^G(x)) - h(\varphi^G(x))| \mu(dx). \end{aligned}$$

We now bound these two terms separately. The first is at most

$$\begin{aligned} &\left| \int f d\gamma - \int f d\mu \right| + \left| \int h \circ \psi^G d\gamma - \int f d\gamma \right| + \left| \int f d\mu - \int h \circ \psi^G d\mu \right| \\ &< \alpha + (\alpha + \gamma(\mathcal{X}^G \setminus U_F)) + (\alpha + \mu(\mathcal{X}^G \setminus U_F)) < 5\alpha, \end{aligned}$$

using the definition of \mathcal{O} . Since h is 1-Lipschitz according to $d_{\mathcal{Y}}^{(F)}$, the second term is at most

$$\int d_{\mathcal{Y}}^{(F)}(\psi^F(x), \varphi^F(x)) \mu(dx) = \int d_{\mathcal{Y}}(\psi(x), \varphi(x)) \mu(dx),$$

and this is at most α according to Definition 2.4. Adding these estimates gives

$$\left| \int h d(\psi_*^G \gamma) - \int h d\nu \right| < 6\alpha < \kappa. \quad \square$$

REMARK. The importance of AL approximations to $\Phi = \varphi^G$ is that we can control their interactions with the w^* -neighbourhoods that appear in the definition of good models. Almost-Lipschitz maps are not the only way to do this, but they are very convenient. On the one hand, we cannot use truly continuous maps in general, since there are choices of \mathcal{X} and \mathcal{Y} for which there are not enough continuous maps $\mathcal{X}^G \rightarrow \mathcal{Y}$. This is why we use maps ψ for which continuity can fail, but only up to an additive error that we control. On the other hand, in the next subsection we will use such approximants ψ to construct a family of maps acting between model spaces, and it will be important to exert some uniform control over the ‘approximate continuity’ of all these other maps. A simple way is to show that they are all η' -almost L' -Lipschitz for some common η' and L' ; see Lemma 4.9 below. \triangleleft

The following lemma and corollary give some simple ways of combining AL approximations.

LEMMA 4.6. *Let $\eta > 0$. Suppose that:*

- $(\mathcal{X}_i^G, \mu_i, S, d_i)$ for $i = 1, 2$ are metric G -process;
- (\mathcal{Y}_i, d'_i) for $i = 1, 2$ are compact metric spaces;
- $\varphi_i : \mathcal{X}_i^G \rightarrow \mathcal{Y}_i$ for $i = 1, 2$ are measurable functions;
- and $\psi_i : \mathcal{X}_i^G \rightarrow \mathcal{Y}_i$ is an η -AL approximation to φ_i rel (μ, d_i, d'_i) for each $i = 1, 2$.

Let d be the Hamming-average metric of d_1 and d_2 on $\mathcal{X}_1 \times \mathcal{X}_2$, and similarly let d' be the Hamming average of d'_1 and d'_2 . Finally, let λ be any joining of μ_1 and μ_2 . Then the map

$$\psi_1 \times \psi_2 : (\mathcal{X}_1 \times \mathcal{X}_2)^G \rightarrow \mathcal{Y}_1 \times \mathcal{Y}_2$$

is a (2η) -AL approximation to $\varphi_1 \times \varphi_2$ rel (λ, d, d') .

Proof. For $i = 1, 2$, let $D_i \subseteq G$ be the finite subsets and $U_i \subseteq \mathcal{X}_i^G$ the open subsets promised by Definition 4.2 for the maps ψ_i . Then (13) and the definition of d give

$$\begin{aligned} & \int d'((\varphi_1(x_1), \varphi_2(x_2)), (\psi_1(x_1), \psi_2(x_2))) \lambda(dx_1, dx_2) \\ &= \frac{1}{2} \left(\int d'_1(\varphi_1(x_1), \psi_1(x_1)) \mu_1(dx_1) + \int d'_2(\varphi_2(x_2), \psi_2(x_2)) \mu_2(dx_2) \right) < \eta. \end{aligned}$$

The map $\psi_1 \times \psi_2$ is $(D_1 \cup D_2)$ -local, and so is the open set $U := U_1 \times U_2$. This set U has

$$\lambda(U) \geq 1 - \mu_1(\mathcal{X}_1^G \setminus U_1) - \mu_2(\mathcal{X}_2^G \setminus U_2) > 1 - 2\eta,$$

and the definition of d implies that $(\psi_1 \times \psi_2)|_U$ is η -almost Lipschitz. □

COROLLARY 4.7. *Let $\eta > 0$. Suppose that:*

- $(\mathcal{X}^G, \mu, S, d)$ is a metric G -process;
- (\mathcal{Y}_i, d'_i) for $i = 1, 2$ are compact metric spaces;
- $\varphi_i : \mathcal{X}^G \rightarrow \mathcal{Y}_i$ for $i = 1, 2$ are measurable functions;
- and $\psi_i : \mathcal{X}^G \rightarrow \mathcal{Y}_i$ is an η -AL approximation to φ_i rel (μ, d, d'_i) for each $i = 1, 2$.

Let d' be the Hamming-average metric of d'_1 and d'_2 on $\mathcal{Y}_1 \times \mathcal{Y}_2$. Then the map

$$(\psi_1, \psi_2) : \mathcal{X}^G \rightarrow \mathcal{Y}_1 \times \mathcal{Y}_2$$

is a (2η) -AL approximation to (φ_1, φ_2) rel (μ, d, d') .

Proof. This is the special case of Lemma 4.6 in which both processes $(\mathcal{X}_i^G, \mu_i, S, d_i)$ are equal to $(\mathcal{X}^G, \mu, S, d)$ and λ is the diagonal joining. □

4.2. Action of approximations on good models. Now suppose that $\psi : \mathcal{X}^G \rightarrow \mathcal{Y}$, that V is a finite set and that $\sigma : G \rightarrow \text{Sym}(V)$. We define a new mapping $\psi^\sigma : \mathcal{X}^V \rightarrow \mathcal{Y}^V$ by

$$\psi^\sigma(\mathbf{x}) := (\psi(\Pi_v^\sigma(\mathbf{x})))_{v \in V}.$$

This is easily visualized if ψ is F -local for some finite $F \subseteq G$, and if v is such that the map $F \rightarrow V : g \mapsto \sigma^g(v)$ is injective. In this case the tuple $\Pi_v^\sigma(\mathbf{x})|_F$ may be regarded as a copy of the restriction $\mathbf{x}|_{\sigma^F(v)}$, ‘pulled back’ so that it is labelled by F itself. Then $\psi^\sigma(\mathbf{x})$ is simply the result of applying ψ to this restriction around each $v \in V$. For a fixed finite F , sofic approximations give that most points $v \in V_n$ satisfy that injectivity requirement once n is large. A general

measurable function $\mathcal{X}^G \rightarrow \mathcal{Y}$ may not be local, but it can be approximated by local functions. As a result, the general definition of ψ^σ still resembles that special case, up to some errors that we have to control from time to time.

The following lemma gives a useful compatibility between ψ^σ and Π_v^σ .

LEMMA 4.8. *Let Σ be a sofic approximation, let $F \subseteq G$ be finite, and suppose that $\psi : \mathcal{X}^G \rightarrow \mathcal{Y}$ is local. Then the following holds w.h.p. in v :*

$$\Pi_v^{\sigma_n}(\psi^{\sigma_n}(\cdot))|_F = \psi^F(\Pi_v^{\sigma_n}(\cdot)).$$

More fully, this conclusion asserts that

$$[\Pi_v^{\sigma_n}(\psi^{\sigma_n}(\mathbf{x}))|_F = \psi^F(\Pi_v^{\sigma_n}(\mathbf{x})) \quad \forall \mathbf{x} \in \mathcal{X}^{V_n}] \quad \text{w.h.p. in } v.$$

Proof. This is similar to the proof of Lemma 3.1. It suffices to prove it when F is an arbitrary singleton, say $\{h\}$. If ψ is D -local, then it holds w.h.p. in v that

$$\sigma_n^g(\sigma_n^h(v)) = \sigma_n^{gh}(v) \quad \forall g \in D.$$

For such v we have

$$\begin{aligned} (\Pi_v^{\sigma_n}(\psi^{\sigma_n}(\mathbf{x})))_h &= (\psi^{\sigma_n}(\mathbf{x}))_{\sigma_n^h(v)} = \psi(\Pi_{\sigma_n^h(v)}^{\sigma_n}(\mathbf{x})) = \psi((x_{\sigma_n^g(\sigma_n^h(v))})_{g \in G}) \\ &= \psi((x_{\sigma_n^{gh}(v)})_{g \in D}) = \psi((\Pi_v^{\sigma_n}(\mathbf{x}))|_{Dh}) = \psi^{\{h\}}(\Pi_v^{\sigma_n}(\mathbf{x})). \quad \square \end{aligned}$$

Now fix compact generating metrics $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$. The next result shows that the maps ψ^σ inherit some regularity from ψ .

LEMMA 4.9. *Suppose that $D \subseteq G$ is finite, that $\psi : \mathcal{X}^G \rightarrow \mathcal{Y}$ is D -local, and that $U \subseteq \mathcal{X}^G$ is a D -local open set such that $\mu(U) > 1 - \eta$ and such that $\psi|_U$ is η -almost L -Lipschitz from $d_{\mathcal{X}}^{(D)}$ to $d_{\mathcal{Y}}$. Then there is a w^* -neighbourhood \mathcal{O} of μ such that*

$$\psi^{\sigma_n} | \Omega(\mathcal{O}, \sigma_n)$$

is 3η -almost $(L|D|)$ -Lipschitz from $d_{\mathcal{X}}^{(V_n)}$ to $d_{\mathcal{Y}}^{(V_n)}$ for all sufficiently large n .

Proof. This is similar to the almost-Lipschitz estimate in the proof of Lemma 4.4. Assume that $d_{\mathcal{Y}}$ has diameter at most 1 for simplicity. Let

$$\mathcal{O} = \{\theta \in \text{Prob}(\mathcal{X}^G) : \theta(U) > 1 - \eta\},$$

which is w^* -open by the Portmanteau theorem. Suppose that $\mathbf{x}, \mathbf{x}' \in \Omega(\mathcal{O}, \sigma_n)$ for some n . Then the definition of \mathcal{O} implies that

$$|\{v \in V : \Pi_v^{\sigma_n}(\mathbf{x}) \notin U\}| = P_{\mathbf{x}}^{\sigma_n}(\mathcal{X}^G \setminus U) < \eta,$$

and similarly for \mathbf{x}' . It follows that

$$\begin{aligned}
 & d_{\mathcal{Y}}^{(V_n)}(\psi^{\sigma_n}(\mathbf{x}), \psi^{\sigma_n}(\mathbf{x}')) \\
 & \leq \frac{1}{|V_n|} \sum_{\substack{v \in V_n, \\ \Pi_v^{\sigma_n}(\mathbf{x}), \Pi_v^{\sigma_n}(\mathbf{x}') \in U}} d_{\mathcal{Y}}(\psi(\Pi_v^{\sigma_n}(\mathbf{x})), \psi(\Pi_v^{\sigma_n}(\mathbf{x}'))) + 2\eta \\
 & \leq \frac{1}{|V_n|} \sum_{\substack{v \in V_n, \\ \Pi_v^{\sigma_n}(\mathbf{x}), \Pi_v^{\sigma_n}(\mathbf{x}') \in U}} Ld_{\mathcal{X}}^{(D)}(\Pi_v^{\sigma_n}(\mathbf{x}), \Pi_v^{\sigma_n}(\mathbf{x}')) + 3\eta \\
 & \leq L|D|d_{\mathcal{X}}^{(V_n)}(\mathbf{x}, \mathbf{x}') + 3\eta. \quad \square
 \end{aligned}$$

Now suppose that $\Phi = \varphi^G : (\mathcal{X}^G, \mu, S) \longrightarrow (\mathcal{Y}^G, \nu, S)$ is a factor map, and that $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ are compact generating metrics on \mathcal{X} and \mathcal{Y} with diameter at most 1. The next proposition shows that AL approximations to φ approximately preserve good models. Related facts can be found within the proofs of [15, Theorem 2.6] and [16, Proposition 3.4], which show that the Kerr–Li definition of sofic entropy is independent of an underlying choice of a ‘dynamically generating’ sequence of bounded functions or of a ‘dynamically generating’ pseudometric.

PROPOSITION 4.10. *For every w^* -neighbourhood \mathcal{N} of ν there is an $\eta > 0$ with the following property. If ψ is an η -AL approximation to φ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$, then there is a w^* -neighbourhood \mathcal{O} of μ such that*

$$\psi^{\sigma_n}(\Omega(\mathcal{O}, \sigma_n)) \subseteq \Omega(\mathcal{N}, \sigma_n)$$

for all sufficiently large n .

Proof. It suffices to prove this for a subbasis of w^* -neighbourhoods \mathcal{N} of ν , so we may assume that

$$\mathcal{N} = \{\theta \in \text{Prob}(\mathcal{Y}^G) : \theta_E \in \mathcal{N}'_1\}$$

for some finite $E \subseteq G$ and w^* -neighbourhood \mathcal{N}'_1 of ν_E .

Now let $\mathcal{N}'_1 \subseteq \mathcal{N}'_1$ be another w^* -neighbourhood of ν_E with the property that any measure which lies sufficiently close to \mathcal{N}'_1 in total variation must lie inside \mathcal{N}'_1 . This is possible because the total-variation norm is stronger than the w^* -topology. Let

$$\mathcal{N}' := \{\theta \in \text{Prob}(\mathcal{Y}^G) : \theta_E \in \mathcal{N}'_1\}.$$

Let $\eta > 0$ be given by Lemma 4.5 so that, for any η -AL approximation ψ to φ , there is a w^* -neighbourhood \mathcal{O} of μ such that

$$(\psi^G)_*(\mathcal{O}) \subseteq \mathcal{N}'.$$

We will show that this η also has the property required for the present proposition, and that we can use the same \mathcal{O} for the function ψ .

Indeed, since ψ is a local function, Lemma 4.8 gives that

$$\Pi_v^{\sigma_n}(\psi^{\sigma_n}(\cdot))|_E = \psi^E(\Pi_v^{\sigma_n}(\cdot)) \quad \text{w.h.p. in } v.$$

Using this, for any $\mathbf{x} \in \mathcal{X}^{V_n}$ the definition of empirical measures gives

$$\begin{aligned} (P_{\psi^{\sigma_n}(\mathbf{x})}^{\sigma_n})_E &= \frac{1}{|V_n|} \sum_{v \in V_n} \delta_{\Pi_v^{\sigma_n}(\psi^{\sigma_n}(\mathbf{x}))|_E} = \frac{1}{|V_n|} \sum_{v \in V_n} \delta_{\psi^E(\Pi_v^{\sigma_n}(\mathbf{x}))} + o(1) \\ &= (\psi^E)_* P_{\mathbf{x}}^{\sigma_n} + o(1) \end{aligned}$$

as $n \rightarrow \infty$, where this approximation is in total variation and is uniform in \mathbf{x} . Therefore, if $\mathbf{x} \in \Omega(\mathcal{O}, \sigma_n)$, then $P_{\mathbf{x}}^{\sigma_n} \in \mathcal{O}$, and so

$$(\psi^G)_* P_{\mathbf{x}}^{\sigma_n} \in \mathcal{N}', \quad \text{that is, } (\psi^E)_* P_{\mathbf{x}}^{\sigma_n} \in \mathcal{N}'_1.$$

Once $(P_{\psi^{\sigma_n}(\mathbf{x})}^{\sigma_n})_E$ lies close enough to $(\psi^E)_* P_{\mathbf{x}}^{\sigma_n}$ in total variation, it lies inside \mathcal{N}'_1 . Since the total-variation estimate above was uniform in \mathbf{x} , this gives that $P_{\psi^{\sigma_n}(\mathbf{x})}^{\sigma_n} \in \mathcal{N}$ for all $\mathbf{x} \in \Omega(\mathcal{O}, \sigma_n)$, for all sufficiently large n . \square

We will sometimes need to compare different approximations to the same factor map.

LEMMA 4.11. *Fix $\eta > 0$, and let ψ and ψ' be two η -AL approximations to φ rel (μ, d_X, d_Y) . Then there is a w^* -neighbourhood \mathcal{O} of μ such that*

$$d_Y^{(V)}(\psi^\sigma(\mathbf{x}), (\psi')^\sigma(\mathbf{x})) < 10\eta \quad \forall \mathbf{x} \in \Omega(\mathcal{O}, \sigma)$$

for any map $\sigma : G \rightarrow \text{Sym}(V)$.

Proof. Recall that we assume d_Y has diameter at most 1.

Let D and U (respectively D' and U') be as in Definition 4.2 for the map ψ (respectively ψ'). By replacing each of D and D' with their union, we may assume they are equal. Having done so, both U and U' are D -local. Consider the function $h : \mathcal{X}^D \rightarrow \mathbb{R}$ defined by

$$h(x) := d_Y(\psi(x), \psi'(x)).$$

This is D -local, and an easy check shows that $h|_{U \cap U'}$ is (2η) -almost Lipschitz according to $d_X^{(D)}$.

Invoking Lemma 2.3, let $f : \mathcal{X}^G \rightarrow [0, 1]$ be a D -local function which is truly Lipschitz according to $d_{\mathcal{X}}^{(D)}$ and has the property that

$$\|(h - f)|_{U \cap U'}\|_{\infty} \leq 2\eta.$$

It follows that

$$\begin{aligned} \int f \, d\mu &\leq \mu(\mathcal{X}^G \setminus (U \cap U')) + 2\eta + \int_{U \cap U'} h \, d\mu \\ &< 4\eta + \int d_{\mathcal{Y}}(\varphi(x), \psi(x)) \, \mu(dx) + \int d_{\mathcal{Y}}(\varphi(x), \psi'(x)) \, \mu(dx) < 6\eta, \end{aligned}$$

using (13), the definition of h and the triangle inequality for $d_{\mathcal{Y}}$.

Now let

$$\mathcal{O} := \left\{ \theta \in \text{Prob}(\mathcal{X}^G) : \theta(U \cap U') > 1 - 2\eta \text{ and } \int f \, d\theta < 6\eta \right\}.$$

This is a w^* -open set which contains μ by construction. The function f was introduced for the sake of defining \mathcal{O} , since h itself may not be strictly continuous and so cannot be used to define a w^* -open set in the same way. For $\mathbf{x} \in \Omega(\mathcal{O}, \sigma)$, it follows that

$$\begin{aligned} d_{\mathcal{Y}}^{(V)}(\psi^{\sigma}(\mathbf{x}), (\psi')^{\sigma}(\mathbf{x})) &= \int h \, dP_{\mathbf{x}}^{\sigma} \\ &\leq \int f \, dP_{\mathbf{x}}^{\sigma} + 2\eta + P_{\mathbf{x}}^{\sigma}(\mathcal{X}^G \setminus (U \cap U')) < 10\eta. \quad \square \end{aligned}$$

4.3. Formulations in terms of sequences. In many of the arguments in Part 2, instead of working with a single AL approximation to a factor map, it will be more convenient to work with a sequence of increasingly good approximations. We therefore make the following relative of Definition 4.2.

DEFINITION 4.12. An **almost Lipschitz** (or **AL**) **approximating sequence** for $\varphi \text{ rel } (\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ is a sequence $(\psi_m)_{m \geq 1}$ such that each ψ_m is an η_m -AL approximation to $\varphi \text{ rel } (\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ for some sequence $\eta_m \downarrow 0$. This situation is denoted by

$$\psi_m \xrightarrow{\text{aL}} \varphi \text{ rel } (\mu, d_{\mathcal{X}}, d_{\mathcal{Y}}).$$

As before, we sometimes drop the qualifier ‘ $\text{rel } (\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ ’ when these data are clear.

We now reformulate a few of the results above in terms of AL approximating sequences. This will assist in their application later in the paper. We will rely on the standard fact that the w^* -topology on $\text{Prob}(\mathcal{X}^G)$ is first countable, since it is metrizable. This will be used again later without further explanation.

We start with the reformulation of Lemma 4.5.

COROLLARY 4.13. *Let μ, ν and φ be as above, let $(\psi_k)_{k \geq 1}$ be an AL approximating sequence to φ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$, and let $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$ and $\mathcal{N}_1 \supseteq \mathcal{N}_2 \supseteq \dots$ be bases for the w^* topologies at μ and ν , respectively. Then whenever the sequence $(k_n)_{n \geq 1}$ grows sufficiently slowly, it holds that whenever the sequence $(m_n)_{n \geq 1}$ grows sufficiently slowly, we have*

$$(\psi_{k_n}^G)_*(\mathcal{O}_n) \subseteq \mathcal{N}_{m_n}$$

for all sufficiently large n .

REMARK. The conclusion here needs to be parsed carefully: both sequences $(k_n)_{n \geq 1}$ and $(m_n)_{n \geq 1}$ must be chosen to grow sufficiently slowly, but the bound on the growth of the second sequence may depend on the choice of the first sequence. ◁

Proof. The desired conclusion is not disrupted if we change \mathcal{N}_1 , so let us assume that $\mathcal{N}_1 = \text{Prob}(\mathcal{Y}^G)$.

Each ψ_k is an η_k -AL approximation to φ for some parameters $\eta_k \downarrow 0$. Therefore, for each m , Lemma 4.5 gives some $K(m)$ such that for every $k \geq K(m)$ there is an $N(k, m)$ for which

$$(\psi_k^G)_*(\mathcal{O}_n) \subseteq \mathcal{N}_m \quad \forall n \geq N(k, m).$$

Since $\mathcal{N}_1 = \text{Prob}(\mathcal{Y}^G)$, we may take $K(1) = 1$ and $N(k, 1) = 1$ for all k .

By replacing each $N(k, m)$ with the value

$$\max\{N(k', m') : 1 \leq k' \leq k, 1 \leq m' \leq m\},$$

we may also assume that

$$N(k_1, m_1) \leq N(k_2, m_2) \quad \text{whenever } k_1 \leq k_2 \text{ and } m_1 \leq m_2.$$

Now for each n define

$$s_n := \max\{s : N(s, s) \leq n\}.$$

This is well defined for all $n \geq 1$ because we have assumed that $N(1, 1) = 1$. Clearly $s_1 \leq s_2 \leq \dots$ and these values tend to ∞ .

Finally, assume that $(k_n)_{n \geq 1}$ is a sequence tending to ∞ for which $k_n \leq s_n$ for all n . Having done so, assume that $(m_n)_{n \geq 1}$ is a sequence tending to ∞ which grows slowly enough that

$$m_n \leq s_n \quad \text{and} \quad K(m_n) \leq k_n \quad \forall n.$$

Then for any n we have $k_n \geq K(m_n)$ and also

$$k_n, m_n \leq s_n \quad \implies \quad N(k_n, m_n) \leq N(s_n, s_n) \leq n.$$

Therefore

$$(\psi_{k_n}^G)_*(\mathcal{O}_n) \subseteq (\psi_{k_n}^G)_*(\mathcal{O}_{N(k_n, m_n)}) \subseteq \mathcal{N}_{m_n}. \quad \square$$

Proposition 4.10 has an analogous reformulation in terms of AL approximating sequences.

COROLLARY 4.14. *Let μ, ν and φ be as above, let $(\psi_k)_{k \geq 1}$ be an AL approximating sequence to φ rel (μ, d_X, d_Y) , and let $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$ and $\mathcal{N}_1 \supseteq \mathcal{N}_2 \supseteq \dots$ be bases for the w^* topologies at μ and ν , respectively. Then whenever the sequence $(k_n)_{n \geq 1}$ grows sufficiently slowly, it holds that whenever the sequence $(m_n)_{n \geq 1}$ grows sufficiently slowly, we have*

$$\psi_{k_n}^{\sigma_n}(\Omega(\mathcal{O}_n, \sigma_n)) \subseteq \Omega(\mathcal{N}_{m_n}, \sigma_n)$$

for all sufficiently large n .

Proof. This follows from Proposition 4.10 just as Corollary 4.13 follows from Lemma 4.5. \square

It will also be convenient to have a version of Lemma 4.8 in terms of sequences.

COROLLARY 4.15. *Let Σ be a sofic approximation, let $F_1 \subseteq F_2 \subseteq \dots$ be finite subsets of G , and let ψ_1, ψ_2, \dots be a sequence of local functions from \mathcal{X}^G to \mathcal{Y} . Then it holds that*

$$\Pi_v^{\sigma_n}(\psi_{k_n}^{\sigma_n}(\cdot))|_{F_{m_n}} = \psi_{k_n}^{F_{m_n}}(\Pi_v^{\sigma_n}(\cdot)) \quad \text{w.h.p. in } v \text{ as } n \longrightarrow \infty$$

for any two sequences $k_1 \leq k_2 \leq \dots$ and $m_1 \leq m_2 \leq \dots$ which both grow sufficiently slowly.

Proof. For each integer $\ell \geq 1$, Lemma 4.8 gives some $N(\ell) \in \mathbb{N}$ such that

$$\begin{aligned} \{|v \in V_n : \Pi_v^{\sigma_n}(\psi_k^{\sigma_n}(\cdot))|_{F_m} = \psi_k^{F_m}(\Pi_v^{\sigma_n}(\cdot)) \forall k, m \in \{1, 2, \dots, \ell\}\} \\ \geq (1 - 2^{-\ell})|V_n| \quad \forall n \geq N(\ell). \end{aligned}$$

Clearly we may assume that $N(1) \leq N(2) \leq \dots$. Now the desired conclusion holds provided

$$k_n, m_n \leq \max(\{1\} \cup \{\ell : N(\ell) \leq n\}) \quad \forall n. \quad \square$$

4.4. A categorial point of view. Section 3 has shown how a metric G -process $(\mathcal{X}^G, \mu, S, d)$ may be converted into the sequences of good model spaces $\Omega(\mathcal{O}, \sigma_n)$ for $n \geq 1$ and w^* -neighbourhoods \mathcal{O} of μ . Then, Section 4 has shown how a factor map

$$\Phi = \varphi^G : (\mathcal{X}^G, \mu, S, d_{\mathcal{X}}) \longrightarrow (\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$$

may be converted into the maps

$$\psi^{\sigma_n} : \mathcal{X}^{V_n} \longrightarrow \mathcal{Y}^{V_n},$$

where ψ is an η -AL approximation to φ for some small η . These maps respect the subsets of good models in the sense of Proposition 4.10.

One can describe all this work as setting up a functor from the category of metric G -processes to another category. The target category here should have objects that are families of sequences of subsets \mathcal{X}^{V_n} , such as our sets $\Omega(\mathcal{O}, \sigma_n)$, or possibly equivalence classes of such families under a kind of ‘asymptotic equivalence’. The morphisms should be (equivalence classes of) families of sequences of maps $\mathcal{X}^{V_n} \longrightarrow \mathcal{Y}^{V_n}$ which respect those sequences of subsets, such as the maps ψ^{σ_n} for the possible choices of η -AL approximation ψ to φ as $\eta \longrightarrow 0$.

Part II

Measures on model spaces and associated entropies

5. Convergence of measures on model spaces

5.1. Local weak* convergence and distributions on measures. Let $\Sigma = (\sigma_n)_{n \geq 1}$ be a sofic approximation to G , and let $(\mathcal{X}^G, \mu, S, d)$ be a metric G -process. It will be clear that all the notions and results of this subsection depend on d , as well as on the process (\mathcal{X}^G, μ, S) .

Suppose also that $\mu_n \in \text{Prob}(\mathcal{X}^{V_n})$ for each $n \geq 1$.

DEFINITION 5.1. The sequence $(\mu_n)_{n \geq 1}$ **locally weak* converges** to μ if for every w^* -neighbourhood \mathcal{O} of μ it is the case that

$$(\Pi_v^{\sigma_n})_* \mu_n \in \mathcal{O} \quad \text{w.h.p. in } v \text{ as } n \longrightarrow \infty.$$

This is denoted by $\mu_n \xrightarrow{\text{lw}^*} \mu$.

That is, once n is large, the local marginals of μ_n resemble those of μ around most vertices of V_n . This kind of convergence depends on d through the resulting w^* topology on $\text{Prob}(\mathcal{X}^G)$.

Bernoulli shifts give the obvious examples.

LEMMA 5.2. *If $(\mathcal{X}^G, v^{\times G}, S)$ is a Bernoulli process over G and d is any choice of compact generating metric on \mathcal{X} , then $v^{\times V_n} \xrightarrow{\text{lw}^*} v^{\times G}$. \square*

Definition 5.1 and some relatives have an established role in statistical physics. They appear naturally in analyses of the asymptotic behaviour of some classical statistical physics models, such as the Ising model, over sequences of sparse graphs. In case those underlying graphs converge to a limiting infinite graph in a suitable sense, one can ask whether the Gibbs measures constructed over them converge to a Gibbs measure over that infinite graph. See, for instance, [20, Subsection 2.2], where local weak* convergence is called ‘[weak] local convergence in probability’. Such convergence for sequences of measures is in much the same spirit as Benjamini–Schramm convergence for sparse random graphs themselves [2].

The next definition connects local weak* convergence and model spaces.

DEFINITION 5.3. The sequence $(\mu_n)_{n \geq 1}$ **quenched-converges** to μ if

(i) $\mu_n \xrightarrow{\text{lw}^*} \mu$, and

(ii) $\mu_n(\Omega(\mathcal{O}, \sigma_n)) \rightarrow 1$ as $n \rightarrow \infty$ for any w^* -neighbourhood \mathcal{O} of μ .

This is denoted by $\mu_n \xrightarrow{q} \mu$.

One can strengthen Lemma 5.2 to show that $v^{\times V_n} \xrightarrow{q} v^{\times G}$. This can be proved directly using the Law of Large Numbers, but we will deduce it after developing some more general theory; see Corollary 5.8 below.

It is sometimes more convenient to replace Definition 5.1 or 5.3 with the following variants. The proofs are immediate, and are omitted.

LEMMA 5.4. *Let $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$ be a fixed basis of w^* -neighbourhoods of μ . Then $\mu_n \xrightarrow{\text{lw}^*} \mu$ if and only if it holds that*

$$(\Pi_v^{\sigma_n})_* \mu_n \in \mathcal{O}_{k_n} \quad \text{w.h.p. in } v \text{ as } n \rightarrow \infty \tag{16}$$

whenever the sequence $k_1 \leq k_2 \leq \dots$ grows sufficiently slowly. Similarly, $\mu_n \xrightarrow{q} \mu$ if and only if we have both (16) and

$$\mu_n(\Omega(\mathcal{O}_{k_n}, \sigma_n)) \longrightarrow 1 \quad \text{as } n \longrightarrow \infty$$

whenever the sequence $k_1 \leq k_2 \leq \dots$ grows sufficiently slowly. \square

In general, it may happen that $\mu_n \xrightarrow{\text{lw}^*} \mu$ as in Definition 5.1, but condition (ii) of Definition 5.3 is not satisfied. In this case, consider the empirical-distribution maps

$$P^{\sigma_n} : \mathcal{X}^{V_n} \longrightarrow \text{Prob}(\mathcal{X}^G) : \mathbf{x} \mapsto P_{\mathbf{x}}^{\sigma_n}.$$

Pushing forward through these maps gives a sequence of distributions on measures

$$P_*^{\sigma_n} \mu_n \in \text{Prob}(\text{Prob}(\mathcal{X}^G)).$$

Since the weak* topology on $\text{Prob}(\mathcal{X}^G)$ is compact and metrizable, this space of distributions on measures carries a weak* topology of its own, which is also compact and metrizable. In the sequel it should always be clear which of these weak* topologies is being referred to.

The distributions on measures $P_*^{\sigma_n} \mu_n$ give the following useful characterization of the difference between Definitions 5.1 and 5.3.

LEMMA 5.5. *If $\mu_n \xrightarrow{\text{lw}^*} \mu$, then $\mu_n \xrightarrow{q} \mu$ if and only if*

$$P_*^{\sigma_n} \mu_n \xrightarrow{\text{weak}^*} \delta_{\mu}.$$

Proof. For any w^* -neighbourhood \mathcal{O} of μ , we have

$$\mu_n(\Omega(\mathcal{O}, \sigma_n)) = \mu_n\{\mathbf{x} \in \mathcal{X}^{V_n} : P_{\mathbf{x}}^{\sigma_n} \in \mathcal{O}\} = (P_*^{\sigma_n} \mu_n)(\mathcal{O}).$$

By the Portmanteau theorem, the w^* -convergence of $P_*^{\sigma_n} \mu_n$ to δ_{μ} is equivalent to the convergence $(P_*^{\sigma_n} \mu_n)(\mathcal{O}) \longrightarrow 1$ for every such \mathcal{O} . \square

Given only that $\mu_n \xrightarrow{\text{lw}^*} \mu$, the distributions on measures $P_*^{\sigma_n} \mu_n$ converge to a decomposition of μ into other invariant measures. The next lemma describes this.

LEMMA 5.6. *If $\theta \in \text{Prob}(\text{Prob}(\mathcal{X}^G))$ is a subsequential w^* -limit of the sequence of distributions on measures $(P_*^{\sigma_n} \mu_n)_{n \geq 1}$, then $\theta(\text{Prob}^S(\mathcal{X}^G)) = 1$, and the barycentre of θ is equal to μ , meaning that*

$$\int v \theta(dv) = \mu.$$

Proof. By passing to a subsequence we may simply assume that

$$P_*^{\sigma_n} \mu_n \xrightarrow{\text{weak}^*} \theta. \tag{17}$$

Support. For each local function $f \in C(\mathcal{X}^G)$ and each $g \in G$, Lemma 3.2 gives that

$$\sup_{\mathbf{x} \in \mathcal{X}^{V_n}} \left| \int f dP_{\mathbf{x}}^{\sigma_n} - \int f \circ S^g dP_{\mathbf{x}}^{\sigma_n} \right| \longrightarrow 0.$$

Since local functions are uniformly dense in $C(\mathcal{X}^G)$, this implies that

$$(P_*^{\sigma_n} \mu_n) \left\{ \nu \in \text{Prob}(\mathcal{X}^G) : \int f d\nu \approx_\varepsilon \int f \circ S^g d\nu \right\} \longrightarrow 1$$

for all $f \in C(X^G)$ and $\varepsilon > 0$. By the Portmanteau theorem, it follows that θ is supported on $\text{Prob}^S(\mathcal{X}^G)$.

Barycentre. For any $f \in C(\mathcal{X}^G)$, the weak* convergence in (17) gives

$$\begin{aligned} \iint f(x) \nu(dx) \theta(d\nu) &= \lim_{n \rightarrow \infty} \iint f(x) \nu(dx) (P_*^{\sigma_n} \mu_n)(d\nu) \\ &= \lim_{n \rightarrow \infty} \iint f(x) P_{\mathbf{x}}^{\sigma_n}(dx) \mu_n(d\mathbf{x}) \\ &= \lim_{n \rightarrow \infty} \frac{1}{|V_n|} \sum_{v \in V_n} \int f(\Pi_v^{\sigma_n}(\mathbf{x})) \mu_n(d\mathbf{x}) \\ &= \lim_{n \rightarrow \infty} \frac{1}{|V_n|} \sum_{v \in V_n} \int f d((\Pi_v^{\sigma_n})_* \mu_n), \end{aligned}$$

and this converges to $\int f d\mu$ because $\mu_n \xrightarrow{\text{lw}^*} \mu$. □

The two previous lemmas have the following immediate consequence.

COROLLARY 5.7. *If $(\mathcal{X}^G, \mu, S, d)$ is ergodic then*

$$\mu_n \xrightarrow{\text{lw}^*} \mu \implies \mu_n \xrightarrow{q} \mu.$$

Proof. Since μ is ergodic, it has no nontrivial representation as a barycentre of other invariant measures. In view of Lemma 5.6, it follows that the only possible subsequential limit of $(P_*^{\sigma_n} \mu_n)_{n \geq 1}$ is the Dirac mass at μ itself. Now Lemma 5.5 completes the proof. □

In light of Lemma 5.2, we can immediately deduce the following.

COROLLARY 5.8. *If $(\mathcal{X}^G, \nu^{\times G}, S)$ is a Bernoulli process over G and d is any choice of compact generating metric on \mathcal{X} , then $\nu^{\times V_n} \xrightarrow{q} \nu^{\times G}$. \square*

This result is already implicit in the proofs of [4, Theorem 8] and [14, Lemma 2.2], which calculate the sofic entropies of Bernoulli processes.

REMARK. Another notion of convergence, also introduced in [20], is ‘local-on-average weak* convergence’. It requires only that

$$\frac{1}{|V_n|} \sum_{v \in V_n} (\Pi_v^{\sigma_n})_* \mu_n \xrightarrow{\text{weak}^*} \mu.$$

This is clearly weaker than local weak* convergence. Indeed, any sequence of measures which satisfy only condition (ii) in Definition 5.3 must have this property, and a variation of the proof of Corollary 5.7 shows that these are equivalent if μ is ergodic. We do not use this kind of convergence in this paper. \triangleleft

5.2. Doubly quenched convergence. For any $\mu_n \in \text{Prob}(\mathcal{X}^{V_n})$ we have

$$(\Pi_v^{\sigma_n})_*(\mu_n \times \mu_n) = (\Pi_v^{\sigma_n})_* \mu_n \times (\Pi_v^{\sigma_n})_* \mu_n \quad \forall v \in V_n.$$

Therefore

$$\mu_n \xrightarrow{\text{lw}^*} \mu \implies \mu_n \times \mu_n \xrightarrow{\text{lw}^*} \mu \times \mu.$$

However, the analogous implication can fail for quenched convergence.

EXAMPLE 5.9. Let G be a residually finite group and let $G > G_1 > G_2 > \dots$ be finite-index normal subgroups whose intersection is $\{e\}$. Let X be the compact inverse limit of the tower of finite groups

$$\dots \rightarrow G/G_2 \rightarrow G/G_1.$$

This X is a compact group. Let m be its Haar measure, and let T be the action of G on X by left rotations. Let d be a left-invariant compact group metric on X . The map (1) gives a metric G -process (X^G, μ, S, d) isomorphic to the Kronecker system (X, m, T) .

For each n , let $\mu_{(n)}$ be the pushforward of the measure μ under the coordinate-wise factor map $X^G \rightarrow (G/G_n)^G$. Then $\mu_{(n)}$ is supported on the G_n -periodic elements of $(G/G_n)^G$.

Let $V_n := G/G_n$ and let $\sigma_n : G \rightarrow \text{Sym}(V_n)$ be the left-rotation action of G for each n . These together give a sofic approximation Σ to G .

For each n we now construct a measure $\mu_n \in \text{Prob}(X^{V_n})$ as follows. Let $S_n \subseteq G$ be a cross-section of G_n in G . The identity mapping $V_n \rightarrow G/G_n$ may be regarded as an element $\mathbf{z}_n \in (G/G_n)^{V_n}$. Let $\mu_n^\circ \in \text{Prob}((G/G_n)^{V_n})$ be the law of a random rotate of \mathbf{z}_n : that is,

$$\mu_n^\circ := \frac{1}{|S_n|} \sum_{g \in S_n} \delta_{\mathbf{z}_n \circ \sigma_n^g}.$$

Finally, let μ_n be any lift of μ_n° to a measure on X^{V_n} .

For each n , the measure μ_n° is the Haar measure on the G -orbit of \mathbf{z}_n . That orbit is a free and transitive (G/G_n) -space. Therefore, the G -action on the finitely supported measure μ_n° is isomorphic to the left-rotation action on G/G_n with Haar measure. Composing with the map (1), this isomorphism converts the elements $\mathbf{z}_n \circ \sigma_n^g \in (G/G_n)^{V_n}$ into the G_n -periodic elements of $(G/G_n)^G$. Therefore, the points in the support of μ_n° have empirical distribution that actually equals $\mu_{(n)}$, and now the local marginals of μ_n° are also all equal to $\mu_{(n)}$. It follows that $\mu_n \xrightarrow{q} \mu$ as $n \rightarrow \infty$.

However, we also have

$$\mu_n^\circ \times \mu_n^\circ = \frac{1}{|S_n|^2} \sum_{g,h \in S_n} \delta_{\mathbf{z}_n \circ \sigma_n^g} \times \delta_{\mathbf{z}_n \circ \sigma_n^h} = \frac{1}{|S_n|^2} \sum_{g,h \in S_n} \delta_{(\mathbf{z}_n \circ \sigma_n^g, \mathbf{z}_n \circ \sigma_n^h) \circ (\text{id}_{V_n} \times \sigma_n^h)}.$$

From this we can calculate that the distribution

$$P_*^{\sigma_n}(\mu_n^\circ \times \mu_n^\circ) \in \text{Prob}(\text{Prob}((G/G_n \times G/G_n)^G))$$

is the law of the random measure

$$\int \delta_{(x, S^h x)} \mu_{(n)}(dx),$$

where h is a uniform random element of S_n . This shows that $P_*^{\sigma_n}(\mu_n \times \mu_n)$ does not converge weakly* to $\delta_{\mu \times \mu}$, but rather to the disintegration of $\mu \times \mu$ into the ergodic components supported on the cosets of the diagonal subgroup in $X \times X$. Therefore, by Lemma 5.5, $\mu_n \times \mu_n$ does not quenched-converge to $\mu \times \mu$.

This example is particularly striking if G has Kazhdan’s property (T). In that case, the Schreier graphs of the quotients σ_n are expanders (see [18, Section 3.3] or [1, Section 6.1]). Therefore, the sofic approximation Σ cannot support arbitrarily good models for any nonergodic G -system; in particular, $h_\Sigma(\mu \times \mu) = -\infty$. In this case there can be no sequence of measures $\nu_n \in \text{Prob}(X^{V_n})$ such that $\nu_n \times \nu_n \xrightarrow{q} \mu \times \mu$. ◁

To rule out examples like these, we make the following definition.

DEFINITION 5.10. The sequence $(\mu_n)_{n \geq 1}$ **doubly-quenched-converges** to μ if

$$\mu_n \times \mu_n \xrightarrow{q} \mu \times \mu.$$

This is denoted by $\mu_n \xrightarrow{\text{dq}} \mu$.

Once again, the obvious positive examples are Bernoulli processes. Indeed, if $(\mathcal{X}^G, \nu^{\times G}, S)$ is a Bernoulli process, then $\nu^{\times G} \times \nu^{\times G}$ may be identified with $(\nu \times \nu)^{\times G}$ and $\nu^{\times V_n} \times \nu^{\times V_n}$ may be identified with $(\nu \times \nu)^{\times V_n}$. Therefore, applying Corollary 5.8 directly to $(\nu \times \nu)^{\times G}$ gives the following.

LEMMA 5.11. *If $(\mathcal{X}^G, \nu^{\times G}, S)$ is a Bernoulli process over G and d is any choice of compact generating metric on \mathcal{X} , then $\nu^{\times V_n} \xrightarrow{\text{dq}} \nu^{\times G}$. \square*

The ‘quenched’ condition in Definition 5.1 asserts that, once n is large, μ_n is mostly supported on individual models \mathbf{x} whose empirical distribution is close to μ . This is a kind of ‘equidistribution’ of the points that support μ , and is an analogue of ergodicity in the setting of a locally weak* convergent sequence of probability measures. With this in mind, doubly-quenched convergence is the analogue of weak mixing.

The main result of this subsection is Theorem A, which gives two other conditions that are equivalent to $\mu_n \xrightarrow{\text{dq}} \mu$. It shows that doubly-quenched convergence is preserved by other Cartesian products as well. This is analogous to some of the classical equivalent conditions for weak mixing of a probability-preserving transformation; see, for instance, [22, Theorem 2.6.1].

The proof of Theorem A requires little more than a few applications of the Cauchy–Bunyakovsky–Schwarz inequality, via the following easy consequence.

LEMMA 5.12. *Let (X, μ) be a probability space, let H be a real Hilbert space with inner product $\langle \cdot, \cdot \rangle$, let $a : X \rightarrow H$ be strongly measurable, and let $b \in H$ have norm at most 1. Then*

$$\int |\langle a(x), b \rangle| \mu(dx) \leq \sqrt{\iint |\langle a(x), a(x') \rangle| \mu(dx) \mu(dx')}. \quad \square$$

Proof of Theorem A. (i) \implies (ii). It suffices to prove (ii) for a subbasis of w^* -neighbourhoods of $\mu \times \nu$, so suppose that

$$\mathcal{N} = \left\{ \lambda \in \text{Prob}(\mathcal{X}^G \times \mathcal{Y}^G) : \int f \otimes h d\lambda \approx_\varepsilon \int f d\mu \int h d\nu \right\}$$

for some $f \in C(\mathcal{X}^G)$ and $h \in C(\mathcal{Y}^G)$ with $\|f\|_\infty, \|h\|_\infty \leq 1$ and some $\varepsilon > 0$.

Let

$$\mathcal{O} := \left\{ \theta \in \text{Prob}(\mathcal{Y}^G) : \int h d\theta \approx_{\varepsilon/2} \int h d\nu \right\}.$$

We will show that this \mathcal{O} suffices. Suppose that $n_i \uparrow \infty$ and $\Omega(\mathcal{O}, \sigma_{n_i}) \neq \emptyset$ for all i . By relabelling the subsequence, we may assume that $n_i = i$ for all i , and so ignore the indexing by i . Now conclusion (ii) will follow if we prove that

$$\mu_n \left\{ \mathbf{x} \in \mathcal{X}^{V_n} : \frac{1}{|V_n|} \sum_{v \in V_n} f(\Pi_v^{\sigma_n} \mathbf{x}) h(\Pi_v^{\sigma_n} \mathbf{y}_n) \approx_{\varepsilon} \int f d\mu \int h d\nu \right\} \longrightarrow 1$$

for any sequence of models $\mathbf{y}_n \in \Omega(\mathcal{O}, \sigma_n)$.

Since $\|f\|_{\infty} \leq 1$, we have

$$\begin{aligned} & \left| \frac{1}{|V_n|} \sum_{v \in V_n} f(\Pi_v^{\sigma_n} \mathbf{x}) h(\Pi_v^{\sigma_n} \mathbf{y}_n) - \int f d\mu \int h d\nu \right| \\ & \leq \left| \frac{1}{|V_n|} \sum_{v \in V_n} \left(f(\Pi_v^{\sigma_n} \mathbf{x}) - \int f d\mu \right) h(\Pi_v^{\sigma_n} \mathbf{y}_n) \right| \\ & \quad + \left| \frac{1}{|V_n|} \sum_{v \in V_n} h(\Pi_v^{\sigma_n} \mathbf{y}_n) - \int h d\nu \right|. \end{aligned}$$

The last term here is less than $\varepsilon/2$ by the definition of \mathcal{O} . It therefore remains to prove that

$$\mu_n \left\{ \mathbf{x} \in \mathcal{X}^{V_n} : \left| \frac{1}{|V_n|} \sum_{v \in V_n} f(\Pi_v^{\sigma_n} \mathbf{x}) h(\Pi_v^{\sigma_n} \mathbf{y}_n) \right| < \varepsilon/2 \right\} \longrightarrow 1$$

under the extra assumption that $\int f d\mu = 0$. By Chebyshev's Inequality, this will follow if we prove that

$$\int \left| \frac{1}{|V_n|} \sum_{v \in V_n} f(\Pi_v^{\sigma_n} \mathbf{x}) h(\Pi_v^{\sigma_n} \mathbf{y}_n) \right| \mu_n(d\mathbf{x}) \longrightarrow 0$$

under that extra assumption.

To do this, consider the Hilbert spaces $H_n := \ell^2(V_n)$ with the inner products

$$\langle a, b \rangle_n := \frac{1}{|V_n|} \sum_{v \in V_n} a_v b_v.$$

Let

$$b := (h(\Pi_v^{\sigma_n} \mathbf{y}_n))_{v \in V_n}$$

and define

$$a : \mathcal{X}^{V_n} \longrightarrow H_n : \mathbf{x} \mapsto (f(\Pi_v^{\sigma_n} \mathbf{x}))_{v \in V_n}.$$

Now Lemma 5.12 gives

$$\begin{aligned} \int \left| \frac{1}{|V_n|} \sum_{v \in V_n} f(\Pi_v^{\sigma_n} \mathbf{x}) h(\Pi_v^{\sigma_n} \mathbf{y}_n) \right| \mu_n(d\mathbf{x}) &= \int |\langle a(\mathbf{x}), b \rangle_n| \mu_n(d\mathbf{x}) \\ &\leq \sqrt{\iint |\langle a(\mathbf{x}), a(\mathbf{x}') \rangle_n| \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}')} \end{aligned}$$

However,

$$\iint |\langle a(\mathbf{x}), a(\mathbf{x}') \rangle_n| \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') = \iint \left| \int f \otimes f dP_{(\mathbf{x}, \mathbf{x}')}^{\sigma_n} \right| \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}'),$$

and this converges to

$$\int f \otimes f d(\mu \times \mu) = \left(\int f d\mu \right)^2 = 0$$

as $n \rightarrow \infty$, by assumption (i).

(ii) \implies (iii). We have seen that $\mu_{n_i} \times \nu_i \xrightarrow{\text{lw}^*} \mu \times \nu$; the only issue is to show that $\mu_{n_i} \times \nu_i$ is asymptotically supported on good models for $\mu \times \nu$. This now follows from conclusion (ii) and Fubini’s theorem, since ν_i is asymptotically supported on good models for ν .

(iii) \implies (i). Clearly (i) is a special case of (iii). □

Conclusion (ii) of Theorem A has a corollary whose conclusion does not involve measures on model spaces. It asserts that, asymptotically as $n \rightarrow \infty$, every sufficiently good model for ν in \mathcal{Y}^{V_n} can be lifted to a good model for $\mu \times \nu$ in $(\mathcal{X} \times \mathcal{Y})^{V_n}$.

COROLLARY 5.13. *Let $(\mathcal{X}^G, \mu, S, d_{\mathcal{X}})$ and $(\mathcal{Y}^G, \mu, S, d_{\mathcal{Y}})$ be metric G -processes, and suppose that $\mu_n \xrightarrow{\text{dq}} \mu$ over Σ . Let $\pi : (\mathcal{X} \times \mathcal{Y})^G \rightarrow \mathcal{Y}$ be the projection onto the \mathcal{Y} -component of the identity coordinate, so that $\pi^{\sigma_n} : (\mathcal{X} \times \mathcal{Y})^{V_n} \rightarrow \mathcal{Y}^{V_n}$ is the coordinate projection for each n . Then for every w^* -neighbourhood \mathcal{N} of $\mu \times \nu$ there is a w^* -neighbourhood \mathcal{O} of ν such that*

$$\pi^{\sigma_n}(\Omega(\mathcal{N}, \sigma_n)) \supseteq \Omega(\mathcal{O}, \sigma_n)$$

for all sufficiently large n . □

The next theorem gives another equivalent characterization of the convergence $\mu_n \xrightarrow{\text{dq}} \mu$. We separate it from Theorem A because it is not involved in the rest of this paper.

THEOREM 5.14. *Assume that $\mu_n \xrightarrow{q} \mu$. Then $\mu_n \xrightarrow{\text{dq}} \mu$ if and only if*

$$(\Pi_v^{\sigma_n}, \Pi_{v'}^{\sigma_n})_* \mu_n \xrightarrow{\text{weak}^*} \mu \times \mu \quad \text{w.h.p. in } (v, v') \in V_n \times V_n \text{ as } n \rightarrow \infty.$$

Proof. (\implies). Let $f_1, f_2 \in C(\mathcal{X}^G)$, and suppose that $\int f_1 d\mu = 0$. Then

$$\begin{aligned} & \frac{1}{|V_n|^2} \sum_{v, v' \in V_n} \left(\int f_1(\Pi_v^{\sigma_n} \mathbf{x}) f_2(\Pi_{v'}^{\sigma_n} \mathbf{x}) \mu_n(d\mathbf{x}) \right)^2 \\ &= \frac{1}{|V_n|^2} \sum_{v, v' \in V_n} \iint f_1(\Pi_v^{\sigma_n} \mathbf{x}) f_1(\Pi_v^{\sigma_n} \mathbf{x}') f_2(\Pi_{v'}^{\sigma_n} \mathbf{x}) f_2(\Pi_{v'}^{\sigma_n} \mathbf{x}') \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') \\ &= \frac{1}{|V_n|} \sum_{v \in V_n} \iint f_1(\Pi_v^{\sigma_n} \mathbf{x}) f_1(\Pi_v^{\sigma_n} \mathbf{x}') \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') \\ &\quad \times \frac{1}{|V_n|} \sum_{v' \in V_n} \iint f_2(\Pi_{v'}^{\sigma_n} \mathbf{x}) f_2(\Pi_{v'}^{\sigma_n} \mathbf{x}') \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') \\ &= \iint \left(\int f_1 \otimes f_1 dP_{(\mathbf{x}, \mathbf{x}')}^{\sigma_n} \right) \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') \\ &\quad \times \iint \left(\int f_2 \otimes f_2 dP_{(\mathbf{x}, \mathbf{x}')}^{\sigma_n} \right) \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}'). \end{aligned}$$

Doubly-quenched convergence implies that the first integral in this product tends to

$$\int f_1 \otimes f_1 d(\mu \times \mu) = \left(\int f_1 d\mu \right)^2 = 0.$$

Therefore, Chebyshev’s Inequality gives that

$$\int f_1(\Pi_v^{\sigma_n} \mathbf{x}) f_2(\Pi_{v'}^{\sigma_n} \mathbf{x}) \mu_n(d\mathbf{x}) \rightarrow 0 \quad \text{w.h.p. in } (v, v').$$

Finally, adjusting by constants as in the proof of (i) \implies (ii) in Theorem A, it follows that

$$\int f_1(\Pi_v^{\sigma_n} \mathbf{x}) f_2(\Pi_{v'}^{\sigma_n} \mathbf{x}) \mu_n(d\mathbf{x}) \rightarrow \int f_1 \otimes f_2 d(\mu \times \mu) \quad \text{w.h.p. in } (v, v')$$

for arbitrary $f_1, f_2 \in C(\mathcal{X}^G)$.

(\Leftarrow). Let $f_1 := f_2 := f \in C(\mathcal{X}^G)$ have mean zero according to μ . Reversing the chain of equalities in the previous step, we see that the assumed weak* convergence implies that

$$\iint \left(\int f \otimes f dP_{(x,x')}^{\sigma_n} \right) \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') \longrightarrow 0 \quad \text{as } n \longrightarrow \infty.$$

Also, it is clear that

$$\begin{aligned} & \iint \left(\int f_1 \otimes f_2 dP_{(x,x')}^{\sigma_n} \right) \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') \\ &= \iint \left(\int f_2 \otimes f_1 dP_{(x,x')}^{\sigma_n} \right) \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') \end{aligned}$$

for any other $f_1, f_2 \in C(\mathcal{X}^G)$, and so the Polarization Identity gives that these integrals also tend to 0 if $\int f_1 d\mu = \int f_2 d\mu = 0$.

Finally, this convergence generalizes to

$$\iint \left(\int f_1 \otimes f_2 dP_{(x,x')}^{\sigma_n} \right) \mu_n(d\mathbf{x}) \mu_n(d\mathbf{x}') \longrightarrow \int f_1 d\mu \int f_2 d\mu$$

for arbitrary $f_1, f_2 \in C(\mathcal{X}^G)$, because the assumption that $\mu_n \xrightarrow{q} \mu$ handles the case when either f_1 or f_2 is constant. □

Theorem 5.14 continues the analogy between doubly-quenched convergence and weak mixing: it corresponds to the classical fact that an ergodic probability-preserving transformation (X, μ, T) is weakly mixing if and only if

$$\frac{1}{n} \sum_{m=1}^n \left| \int f(x)g(T^m x) \mu(dx) - \int f d\mu \int g d\mu \right| \longrightarrow 0 \quad \forall f, g \in L^2(\mu)$$

as $n \longrightarrow \infty$ (see again [22, Theorem 2.6.1]).

Doubly-quenched convergence is not only analogous to weak mixing, but also logically related to it.

LEMMA 5.15. *If (\mathcal{X}^G, μ, S) is weakly mixing, then*

$$\mu_n \xrightarrow{q} \mu \implies \mu_n \xrightarrow{dq} \mu.$$

Proof. This follows from Corollary 5.7, since if μ is weakly mixing then $\mu \times \mu$ is ergodic. □

REMARK. Another property which seems related to doubly-quenched convergence is ‘replica symmetry’. In the study of spin glasses and other disordered systems in statistical physics, the term ‘replica symmetry’ (or its negation, ‘replica symmetry breaking’) is used for a variety of phenomena that are expected to occur or fail together in most models of interest.

Often such a model consists of a special sequence of probability measures μ_n on $\{\pm 1\}^{V_n}$ for some finite sets V_n , for instance given by a particular Hamiltonian. In spin glasses, the measures (and the Hamiltonians) are usually random themselves. For such a sequence of measures, one popular meaning of ‘replica symmetry’ is that the sequence of ‘overlaps’

$$R(\sigma, \sigma') := \frac{1}{n} \sum_{i=1}^n \sigma_i \sigma'_i, \quad (\sigma, \sigma') \in \{\pm 1\}^n \times \{\pm 1\}^n,$$

regarded as a sequence of random variables for the probabilities $\mu_n \times \mu_n$, should concentrate as $n \rightarrow \infty$.

Clearly this holds in case the sets V_n are associated to some sofic approximation of a group G and $\mu_n \times \mu_n \xrightarrow{q} \mu \times \mu$ for some shift-invariant measure μ on $\{\pm 1\}^G$. But doubly-quenched convergence could be stronger in general. We should be careful about bringing the term ‘replica symmetry’ into ergodic theory, since it does have several meanings for the physicists and it is not yet clear under what conditions they coincide. However, it would be very interesting to know whether such ideas or models can shed further light on doubly-quenched convergence.

An introduction to replica symmetry and replica symmetry breaking from a physical point of view can be found in [19, Chs 8, 12 and 19]. \triangleleft

5.3. Behaviour under factor maps. This subsection considers how local weak* and quenched convergence behave under applying AL approximants to factor maps.

PROPOSITION 5.16. *Let $\Phi = \varphi^G : (\mathcal{X}^G, \mu, S, d_{\mathcal{X}}) \rightarrow (\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$ be a factor map of metric G -processes. Suppose that $\mu_n \xrightarrow{\text{lw}^*} \mu$, and let $\psi_k \xrightarrow{\text{aL}} \varphi$ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$. Then*

$$(\psi_{k_n}^{\sigma_n})_* \mu_n \xrightarrow{\text{lw}^*} \nu$$

whenever the sequence $k_1 \leq k_2 \leq \dots$ grows sufficiently slowly. The same holds if both instances of ‘ $\xrightarrow{\text{lw}^}$ ’ are replaced with ‘ \xrightarrow{q} ’ or if both are replaced with ‘ $\xrightarrow{\text{dq}}$ ’.*

Proof. First consider the case of quenched convergence. Suppose $\mu_n \xrightarrow{q} \mu$.

Let $\mathcal{N}_1 \supseteq \mathcal{N}_2 \supseteq \dots$ be a basis of w^* -neighbourhoods at ν . By shrinking each \mathcal{N}_m if necessary, we may assume that each of them has the form

$$\mathcal{N}'_m = \{\theta \in \text{Prob}(\mathcal{Y}^G) : \theta_{F_m} \in \mathcal{N}'_m\}$$

for some finite set $F_m \subseteq G$ and w^* -neighbourhood \mathcal{N}'_m of ν_{F_m} .

Using Lemma 5.4, we can choose a basis $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$ of w^* -neighbourhoods at μ for which

$$(\Pi_v^{\sigma_n})_* \mu_n \in \mathcal{O}_n \quad \text{w.h.p. in } \nu \text{ as } n \rightarrow \infty \tag{18}$$

and

$$\mu_n(\Omega(\mathcal{O}_n, \sigma_n)) \rightarrow 1 \quad \text{as } n \rightarrow \infty. \tag{19}$$

By the combination of Corollaries 4.13 and 4.14, if we now choose $(k_n)_{n \geq 1}$ growing sufficiently slowly, then any choice of $(m_n)_{n \geq 1}$ which grows sufficiently slowly (depending on $(k_n)_{n \geq 1}$) gives

$$(\psi_{k_n}^G)_*(\mathcal{O}_n) \subseteq \mathcal{N}_{m_n} \quad \text{and} \quad \psi_{k_n}^{\sigma_n}(\Omega(\mathcal{O}_n, \sigma_n)) \subseteq \Omega(\mathcal{N}_{m_n}, \sigma_n) \tag{20}$$

for all sufficiently large n . In addition, by Corollary 4.15, if we choose $(k_n)_{n \geq 1}$ and $(m_n)_{n \geq 1}$ both growing slowly enough, then

$$((\Pi_v^{\sigma_n})_*(\psi_{k_n}^{\sigma_n})_* \mu_n)_{F_{m_n}} = (\psi_{k_n}^{F_{m_n}})_*(\Pi_v^{\sigma_n})_* \mu_n \quad \text{w.h.p. in } \nu \text{ as } n \rightarrow \infty.$$

Having chosen such sequences $(k_n)_{n \geq 1}$ and $(m_n)_{n \geq 1}$, we obtain from (18) and (20) that the following all hold w.h.p. in ν as $n \rightarrow \infty$:

$$\begin{aligned} ((\Pi_v^{\sigma_n})_*(\psi_{k_n}^{\sigma_n})_* \mu_n)_{F_{m_n}} &= ((\psi_{k_n}^G)_*(\Pi_v^{\sigma_n})_* \mu_n)_{F_{m_n}} \in ((\psi_{k_n}^G)_*(\mathcal{O}_n))_{F_{m_n}} \\ &\subseteq \{\theta_{F_{m_n}} : \theta \in \mathcal{N}'_{F_{m_n}}\} = \mathcal{N}'_{F_{m_n}}, \end{aligned}$$

and hence $(\psi_{k_n}^{\sigma_n})_* \mu_n \xrightarrow{\text{lw}^*} \nu$. Similarly, from (19) and (20) we obtain that

$$\begin{aligned} ((\psi_{k_n}^{\sigma_n})_* \mu_n)(\Omega(\mathcal{N}_{m_n}, \sigma_n)) &\geq ((\psi_{k_n}^{\sigma_n})_* \mu_n)(\psi_{k_n}^{\sigma_n}(\Omega(\mathcal{O}_n, \sigma_n))) \\ &\geq \mu_n(\Omega(\mathcal{O}_n, \sigma_n)) \rightarrow 1 \end{aligned}$$

as $n \rightarrow \infty$, so in fact $(\psi_{k_n}^{\sigma_n})_* \mu_n \xrightarrow{q} \nu$, by Lemma 5.4.

For the case of local weak* convergence, we argue in the same way, except ignoring the lower bounds on $((\psi_{k_n}^{\sigma_n})_* \mu_n)(\Omega(\mathcal{N}_{m_n}, \sigma_n))$ and omitting the appeal to Corollary 4.14.

Finally, Lemma 4.6 gives that $\psi_k \times \psi_k \xrightarrow{\text{al}} \varphi \times \varphi \text{ rel } \mu \times \mu$. Therefore, if $\mu_n \xrightarrow{\text{dq}} \mu$, then we may apply the argument for the quenched case to $\mu_n \times \mu_n$. \square

If $\varphi : \mathcal{X}^G \rightarrow \mathcal{Y}$ is itself local and continuous, then by Lemma 2.2 we may regard φ as an η -AL approximation to itself for every $\eta > 0$. Therefore, Proposition 5.16 has the following special case.

COROLLARY 5.17. *If $\Phi = \varphi^G$ is as in Proposition 5.16 with φ local and continuous, and $\mu_n \xrightarrow{\text{lw}^*} \mu$, then*

$$\varphi_*^{\sigma_n} \mu_n \xrightarrow{\text{lw}^*} \nu,$$

and similarly for quenched and doubly-quenched convergence.

In particular, suppose that λ is a joining of the metric G -processes $(\mathcal{X}^G, \mu, S, d_{\mathcal{X}})$ and $(\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$, and that $\lambda_n \in \text{Prob}((\mathcal{X} \times \mathcal{Y})^{V_n})$ satisfies

$$\lambda_n \xrightarrow{\text{lw}^*} \lambda.$$

If μ_n and ν_n are the marginals of λ_n on \mathcal{X}^{V_n} and \mathcal{Y}^{V_n} respectively, then

$$\mu_n \xrightarrow{\text{lw}^*} \mu \quad \text{and} \quad \nu_n \xrightarrow{\text{lw}^*} \nu.$$

The same conclusions hold if locally weak* convergence is replaced with quenched or doubly-quenched convergence throughout. \square

The mere existence of a convergent sequence of measures on model spaces is a feature of a process that can be useful. Proposition 5.16 has the following important consequence for this feature.

COROLLARY 5.18. *For a given metric G -process $(\mathcal{X}^G, \mu, S, d_{\mathcal{X}})$ and sofic approximation Σ , the property that there exists a sequence $\mu_n \in \text{Prob}(\mathcal{X}^{V_n})$ which locally weak* (respectively, quenched or doubly-quenched) converges to μ is preserved by all factor maps, including all isomorphisms. In particular, it is a property of the process (\mathcal{X}^G, μ, S) , not depending on the choice of the metric $d_{\mathcal{X}}$.*

Therefore, the definitions of these properties may be extended unambiguously to abstract G -systems. \square

Beware of the following distinction: this corollary tells us that the existence of a sequence $\mu_n \xrightarrow{\text{lw}^*} \mu$ is independent of $d_{\mathcal{X}}$, but whether a given sequence μ_n satisfies this certainly does depend on $d_{\mathcal{X}}$.

6. Model-measure sofic entropies

Let (\mathcal{X}^G, μ, S) be a G -process and d a compact generating metric on \mathcal{X} . As in the Introduction, we define the **quenched model-measure sofic entropy** to be

$$h_{\Sigma}^q(\mu) := \sup \left\{ \sup_{\delta, \varepsilon > 0} \limsup_{i \rightarrow \infty} \frac{1}{|V_{n_i}|} \log \text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})}) : \right. \\ \left. n_i \uparrow \infty \text{ and } \mu_i \xrightarrow{q} \mu \text{ over } (\sigma_{n_i})_{i \geq 1} \right\},$$

and the **doubly-quenched model-measure sofic entropy** to be

$$h_{\Sigma}^{\text{dq}}(\mu) := \sup \left\{ \sup_{\delta, \varepsilon > 0} \limsup_{i \rightarrow \infty} \frac{1}{|V_{n_i}|} \log \text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})}) : \right. \\ \left. n_i \uparrow \infty \text{ and } \mu_i \xrightarrow{\text{dq}} \mu \text{ over } (\sigma_{n_i})_{i \geq 1} \right\}.$$

As for sofic entropy, it turns out that these do not depend on the choice of compact generating metric, and so it is omitted from the notation.

In general, it could happen that there are

$$n_1 < m_1 < n_2 < m_2 < \dots,$$

a w^* -neighbourhood \mathcal{O} of μ , and a sequence of measure $\mu_i \in \text{Prob}(\mathcal{X}^{V_{n_i}})$ such that

$$\mu_{n_i} \xrightarrow{q} \mu \quad \text{over } (\sigma_{n_i})_{i \geq 1},$$

but on the other hand

$$\Omega(\mathcal{O}, \sigma_{m_j}) = \emptyset \quad \forall j \geq 1. \tag{21}$$

The first of these conditions implies that $\Omega(\mathcal{O}, \sigma_{n_i}) \neq \emptyset$ for all sufficiently large i , and hence $h_{\Sigma}(\mu) \geq 0$. Thus, the presence of any sofic subapproximation along which one can find good models gives a lower bound on the sofic entropy. We wish to define the model-measure sofic entropies so that they have the analogous property.

However, if (21) holds then there is no way to insert $(\mu_i)_{i \geq 1}$ into a sequence of measures which quenched-converges to μ over the whole of the original sofic approximation. This is why we must explicitly allow a supremum over arbitrary subsequences $(n_i)_{i \geq 1}$ in the definitions of h_{Σ}^q and h_{Σ}^{dq} , in addition to taking a limit supremum along those subapproximations.

Just as for sofic entropy, it is sometimes important to know whether a restriction to sofic subapproximations is really necessary in computing h_{Σ}^{dq} . This can be

expressed by comparing with the **lower doubly-quenched model-measure sofic entropy**:

$$h_{\Sigma}^{\text{dq}}(\mu) := \sup \left\{ \sup_{\delta, \varepsilon > 0} \liminf_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{cov}_{\varepsilon, \delta}(\mu_n, d^{(V_n)}) : \mu_n \xrightarrow{\text{dq}} \mu \text{ over } \Sigma \right\}.$$

This time we do insist that $\mu_n \xrightarrow{\text{dq}} \mu$ over the whole of the original sofic approximation Σ , and then we take a limit infimum as $n \rightarrow \infty$, rather than a limit supremum.

The simplest relationship between sofic entropy and its model-measure variants is as follows.

LEMMA 6.1. *For any G -system $(\mathcal{X}^G, \mu, S, d)$ we have*

$$h_{\Sigma}(\mu) \geq h_{\Sigma}^{\text{q}}(\mu) \geq h_{\Sigma}^{\text{dq}}(\mu).$$

Proof. Firstly, suppose that $n_i \uparrow \infty$ and that $\mu_i \xrightarrow{\text{q}} \mu$ over $(\sigma_{n_i})_{i \geq 1}$. For any $\varepsilon, \delta > 0$ and any w^* -neighbourhood \mathcal{O} of μ , condition (ii) in Definition 5.3 gives that

$$\mu_i(\Omega(\mathcal{O}, \sigma_{n_i})) > 1 - \varepsilon \quad \text{and hence } \text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})}) \leq \text{cov}_{\delta}(\Omega(\mathcal{O}, \sigma_{n_i}), d^{(V_{n_i})})$$

for all sufficiently large i . This shows that $h_{\Sigma}^{\text{q}}(\mu) \leq h_{\Sigma}(\mu)$.

Second, since doubly-quenched convergence implies quenched convergence, the supremum defining $h_{\Sigma}^{\text{dq}}(\mu)$ is over a subset of that defining $h_{\Sigma}^{\text{q}}(\mu)$, so $h_{\Sigma}^{\text{dq}}(\mu) \leq h_{\Sigma}^{\text{q}}(\mu)$. \square

EXAMPLE 6.2. Let $\mathcal{X} = \{0, 1\}$, let G be the free group on four generators, and let (\mathcal{X}^G, μ, S) be the G -process constructed in Example 3.5. We saw in that example that $h_{\Sigma}(\mu) = 0$, but also that, once n is large, any sufficiently good model for μ in \mathcal{X}^{V_n} must be very close to the particular model $\mathbf{x}_n = 1_{w_n}$. Therefore, if the measures $\mu_n \in \text{Prob}(\mathcal{X}^{V_n})$ are asymptotically supported on good models of μ , then they must be mostly supported on smaller and smaller Hamming balls around \mathbf{x}_n . Once n is large this has the following consequence: for most vertices $v \in V_n$, the marginal of μ_n around v is close to either the Dirac mass at 0 or the Dirac mass at 1. This violates the definition of local weak* convergence, and so there are no subsequence $n_1 < n_2 < \dots$ and measures satisfying $\mu_i \xrightarrow{\text{q}} \mu$ over $(\sigma_{n_i})_{i \geq 1}$. Thus $h_{\Sigma}^{\text{q}}(\mu) = -\infty$. \triangleleft

Of course, Lemma 5.15 gives immediately that $h_{\Sigma}^{\text{q}}(\mu) = h_{\Sigma}^{\text{dq}}(\mu)$ if (\mathcal{X}^G, μ, S) is weakly mixing. Kronecker systems can be examples in which $h_{\Sigma}^{\text{q}}(\mu) > h_{\Sigma}^{\text{dq}}(\mu)$.

EXAMPLE 6.3. If G is a group with Kazhdan's property (T) in Example 5.9, then we produced a sequence $\mu_n \xrightarrow{q} \mu$ over the given sofic approximation Σ , but also showed that there can be no sequence $\nu_n \xrightarrow{\text{dq}} \mu$. The latter argument still holds over any sofic subapproximation. Therefore, $h_\Sigma^q(\mu) \geq 0$ (indeed, simple estimates show that it equals 0) but $h_\Sigma^{\text{dq}}(\mu) = -\infty$ in that example.

We also saw that $h_\Sigma(\mu \times \mu) = -\infty$ for this system. Therefore, this is an example in which one cannot replace h_Σ^{dq} with h_Σ^q in the conclusion of Theorem B. \triangleleft

6.1. Invariance under isomorphism.

THEOREM 6.4. *For a fixed sofic approximation Σ , the quantities $h_\Sigma^q(\mu)$ and $h_\Sigma^{\text{dq}}(\mu)$ are isomorphism invariants of the metric G -process $(\mathcal{X}^G, \mu, S, d)$. In particular, they do not depend on the choice of d .*

Proposition 5.16 is the key to this result, together with the following.

PROPOSITION 6.5. *Suppose that $\Phi = \varphi^G : (\mathcal{X}^G, \mu, S) \rightarrow (\mathcal{Y}^G, \nu, S)$ is a factor map, that d_X and d_Y are compact generating metrics of diameter at most 1, that $\mu_n \xrightarrow{q} \mu$, and that $\psi_m \xrightarrow{\text{al}} \varphi$ rel (μ, d_X, d_Y) . For any $\varepsilon, \delta \in (0, 1)$ there is a $\delta' > 0$ for which the following holds. Provided $m_1 \leq m_2 \leq \dots$ grows sufficiently slowly, we have*

$$\text{cov}_{\varepsilon, \delta}((\psi_{m_n}^{\sigma_n})_* \mu_n, d_Y^{(V_n)}) \leq \text{cov}_{\varepsilon/4, \delta'}(\mu_n, d_X^{(V_n)})$$

for all sufficiently large n .

Proof. Suppose that ψ_m is an η_m -AL approximation to φ for each m , where $\eta_m \downarrow 0$. We may assume that $(\eta_m)_{m \geq 1}$ is nonincreasing. Then there are finite sets $D_m \subseteq G$, D_m -local open sets $U_m \subseteq \mathcal{X}^G$, and constants $L_m < \infty$ such that each $\psi_m : \mathcal{X}^G \rightarrow \mathcal{Y}$ is η_m -almost L_m -Lipschitz from $d_X^{(D_m)}$ to d_Y .

For each m , Lemma 4.9 gives w^* -neighbourhoods \mathcal{N}_m of μ such that

$$\psi_m^{\sigma_n} | \Omega(\mathcal{N}_m, \sigma_n)$$

is $(3\eta_m)$ -almost $(L_m | D_m|)$ -Lipschitz for all sufficiently large n . Choose m' large enough that $3\eta_{m'} < \delta/1000$. Now choose δ' small enough that

$$3\eta_{m'} + L_{m'} | D_{m'} | \delta' < \delta/2. \quad (22)$$

Once n is sufficiently large, we have

$$\mu_n(\Omega(\mathcal{N}_{m'}, \sigma_n)) > 1 - \varepsilon/4 > 3/4.$$

This implies that

$$\text{cov}_{3\varepsilon/4, \delta/2}((\psi_{m'}^{\sigma_n})_* \mu_n, d_{\mathcal{Y}}^{(V_n)}) \leq \text{cov}_{\varepsilon/2, \delta/2}((\psi_{m'}^{\sigma_n})_* (\mu_{n|\Omega(\mathcal{N}_{m'}, \sigma_n)}), d_{\mathcal{Y}}^{(V_n)}), \tag{23}$$

where $\mu_{n|\Omega(\mathcal{N}_{m'}, \sigma_n)}$ is the measure μ_n conditioned on the subset $\Omega(\mathcal{N}_{m'}, \sigma_n)$. It also implies that

$$\text{cov}_{\varepsilon/2, \delta'}(\mu_{n|\Omega(\mathcal{N}_{m'}, \sigma_n)}, d_{\mathcal{X}}^{(V_n)}) \leq \text{cov}_{\varepsilon/4, \delta'}(\mu_n, d_{\mathcal{X}}^{(V_n)}), \tag{24}$$

because if $\mu_n(F) < \varepsilon/4$ then

$$\mu_{n|\Omega(\mathcal{N}_{m'}, \sigma_n)}(F) \leq \frac{\mu_n(F)}{\mu_n(\Omega(\mathcal{N}_{m'}, \sigma_n))} < \frac{\varepsilon/4}{1 - \varepsilon/4} < \varepsilon/2.$$

For sufficiently large n , the fact that $\psi_{m'}^{\sigma_n}|\Omega(\mathcal{N}_{m'}, \sigma_n)$ is $(3\eta_{m'})$ -almost $(L_{m'}|D_{m'}|)$ -Lipschitz may be combined with Lemma 2.8 to conclude that

$$\text{cov}_{\varepsilon/2, 3\eta_{m'} + L_{m'}|D_{m'}| \delta'}((\psi_{m'}^{\sigma_n})_* \nu, d_{\mathcal{Y}}^{(V_n)}) \leq \text{cov}_{\varepsilon/2, \delta'}(\nu, d_{\mathcal{X}}^{(V_n)})$$

for any Borel probability measure ν supported on $\Omega(\mathcal{N}_{m'}, \sigma_n)$. Applying this with $\nu = \mu_{n|\Omega(\mathcal{N}_{m'}, \sigma_n)}$ and using (22) and (24), we obtain

$$\text{cov}_{\varepsilon/2, \delta/2}((\psi_{m'}^{\sigma_n})_* (\mu_{n|\Omega(\mathcal{N}_{m'}, \sigma_n)}), d_{\mathcal{Y}}^{(V_n)}) \leq \text{cov}_{\varepsilon/4, \delta'}(\mu_n, d_{\mathcal{X}}^{(V_n)}) \tag{25}$$

for all sufficiently large n .

Finally, since η_m decreases as $m \rightarrow \infty$, Lemma 4.11 gives w^* -neighbourhoods \mathcal{O}_m of μ such that

$$d_{\mathcal{Y}}^{(V_n)}(\psi_{m'}^{\sigma_n}(\mathbf{x}), \psi_{m'}^{\sigma_n}(\mathbf{x})) < 10\eta_{m'} < \delta/2 \quad \forall \mathbf{x} \in \mathcal{O}_m \quad \forall m \geq m'.$$

By forming running intersections, we may assume that $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$. Since $\mu_n(\Omega(\mathcal{O}_m, \sigma_n))$ tends to 1 for every m , it follows that, if $m_1 \leq m_2 \leq \dots$ grows sufficiently slowly, then

$$\mu_n\{\mathbf{x} : d_{\mathcal{Y}}^{(V_n)}(\psi_{m_1}^{\sigma_n}(\mathbf{x}), \psi_{m_2}^{\sigma_n}(\mathbf{x})) > \delta/2\} \rightarrow 0 \quad \text{as } n \rightarrow \infty.$$

Therefore, for such a slowly growing sequence, Lemma 2.9 gives

$$\text{cov}_{\varepsilon, \delta}((\psi_{m_n}^{\sigma_n})_* \mu_n, d_{\mathcal{Y}}^{(V_n)}) \leq \text{cov}_{3\varepsilon/4, \delta/2}((\psi_{m'}^{\sigma_n})_* \mu_n, d_{\mathcal{Y}}^{(V_n)})$$

for all sufficiently large n . This completes the proof in combination with inequalities (23) and (25). □

Proof of Theorem 6.4. Let $\Phi = \varphi^G : (\mathcal{X}^G, \mu, S) \longrightarrow (\mathcal{Y}^G, \nu, S)$ be an isomorphism, choose compact generating metrics $d_{\mathcal{X}}$ for \mathcal{X} and $d_{\mathcal{Y}}$ for \mathcal{Y} of diameter at most 1, and suppose that $\Phi^{-1} = \tilde{\varphi}^G$. Let d be the Hamming average of $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$, and let $d_{\mathcal{X}}^{(2)}$ be the Hamming average of two copies of $d_{\mathcal{X}}$.

We will show that $h_{\Sigma}^q(\mu) \leq h_{\Sigma}^q(\nu)$; the reverse must also hold by symmetry. An exactly analogous argument gives the proof for h_{Σ}^{dq} .

If the left-hand side is $-\infty$ then there is nothing to prove. So let $n_i \uparrow \infty$ be a subsequence, suppose that $\mu_{n_i} \xrightarrow{q} \mu$ over $(\sigma_{n_i})_{i \geq 1}$, and let $\varepsilon > 0$ and $\delta > 0$. We will produce measures $\nu_{n_i} \xrightarrow{q} \nu$ over $(\sigma_{n_i})_{i \geq 1}$ and a $\delta' > 0$ such that

$$\text{cov}_{\varepsilon, \delta}(\mu_{n_i}, d_{\mathcal{X}}^{(V_{n_i})}) \leq \text{cov}_{\varepsilon/8, \delta'}(\nu_{n_i}, d_{\mathcal{Y}}^{(V_{n_i})})$$

for all sufficiently large i . To this end, it suffices to consider only the sofic subapproximation $(\sigma_{n_i})_{i \geq 1}$, so we may relabel this subapproximation and simply assume that it equals $(\sigma_n)_{n \geq 1}$.

Now let $\psi_k \xrightarrow{\text{al}} \varphi$ rel $(\mu, d_{\mathcal{X}}, d_{\mathcal{Y}})$ and $\tilde{\psi}_m \xrightarrow{\text{al}} \tilde{\varphi}$ rel $(\nu, d_{\mathcal{Y}}, d_{\mathcal{X}})$. In addition, let $\xi : \mathcal{X}^G \longrightarrow \mathcal{X}$ be the projection onto the e -indexed coordinate, so $\xi^G = \text{id}_{\mathcal{X}^G}$. Finally, let

$$\lambda := \int_{\mathcal{X}^G} \delta_{(\alpha, \Phi(x))} \mu(dx) \quad \text{and} \quad \hat{\lambda} := \int_{\mathcal{X}^G} \delta_{(x, x)} \mu(dx).$$

These are the graphical joining of μ and ν corresponding to the factor map Φ , and the diagonal joining of μ with itself, respectively.

Since ξ may be regarded as a constant AL approximating sequence to itself, Lemma 4.6 and Corollary 4.7 give that

$$(\xi, \psi_k) \xrightarrow{\text{al}} (\xi, \varphi) \quad \text{rel } (\mu, d_{\mathcal{X}}, d) \tag{26}$$

and

$$\xi \times \tilde{\psi}_m \xrightarrow{\text{al}} \xi \times \tilde{\varphi} \quad \text{rel } (\lambda, d, d_{\mathcal{X}}^{(2)}). \tag{27}$$

By Proposition 5.16 and (26), if we choose $k_1 \leq k_2 \leq \dots$ growing sufficiently slowly, then

$$\lambda_n := ((\xi, \psi_{k_n})^{\sigma_n})_* \mu_n \xrightarrow{q} \lambda. \tag{28}$$

Fix such a sequence $(k_n)_{n \geq 1}$, and define $\nu_n := (\psi_{k_n}^{\sigma_n})_* \mu_n$, so this is the marginal of λ_n on the space \mathcal{Y}^{ν_n} . By Corollary 5.17 we also have $\nu_n \xrightarrow{q} \nu$.

We now apply Proposition 5.16 with (27) and (28), and also Proposition 6.5. According to those propositions, if we choose $m_1 \leq m_2 \leq \dots$ growing sufficiently slowly, then we have

$$((\xi \times \tilde{\psi}_{m_n})^{\sigma_n})_* \lambda_n \xrightarrow{q} (\xi^G \times \tilde{\varphi})_* \lambda = \hat{\lambda}, \tag{29}$$

and also there is a $\delta' > 0$ such that

$$\text{cov}_{\varepsilon/2, \delta/2}((\tilde{\psi}_{m_n}^{\sigma_n})_* \nu_n, d_{\mathcal{X}}^{(V_n)}) \leq \text{cov}_{\varepsilon/8, \delta'}(\nu_n, d_{\mathcal{Y}}^{(V_n)}) \quad (30)$$

for all sufficiently large n .

We finish the proof by comparing μ_n with its image measure

$$(\tilde{\psi}_{m_n}^{\sigma_n})_* \nu_n = (\tilde{\psi}_{m_n}^{\sigma_n})_*(\psi_{k_n}^{\sigma_n})_* \mu_n.$$

This is where we need the joining $\hat{\lambda}$. On $\mathcal{X}^G \times \mathcal{X}^G$, let F be the continuous function

$$F((x_g)_{g \in G}, (x'_g)_{g \in G}) := d_{\mathcal{X}}(x_e, x'_e).$$

Then by (29) and a simple calculation we have

$$\begin{aligned} \int d_{\mathcal{X}}^{(V_n)}(\mathbf{x}, \tilde{\psi}_{m_n}^{\sigma_n}(\psi_{k_n}^{\sigma_n}(\mathbf{x}))) \mu_n(d\mathbf{x}) &= \iint F dP_{(\mathbf{x}, \tilde{\psi}_{m_n}^{\sigma_n}(\psi_{k_n}^{\sigma_n}(\mathbf{x})))}^{\sigma_n} \mu_n(d\mathbf{x}) \\ &= \iint F dP_{(\mathbf{x}, \tilde{\psi}_{m_n}^{\sigma_n}(\mathbf{y}))}^{\sigma_n} \lambda_n(d\mathbf{x}, d\mathbf{y}) \\ &\longrightarrow \int F d\hat{\lambda} = 0. \end{aligned}$$

Therefore, Lemma 2.9 gives

$$\text{cov}_{\varepsilon, \delta}(\mu_n, d_{\mathcal{X}}^{(V_n)}) \leq \text{cov}_{\varepsilon/2, \delta/2}((\tilde{\psi}_{m_n}^{\sigma_n})_* \nu_n, d_{\mathcal{X}}^{(V_n)})$$

for all sufficiently large n . Combining with (30) completes the proof. \square

7. Entropy of Cartesian products

7.1. Proof of Theorem B.

Proof of Theorem B. Let $(\mathcal{X}^G, \mu, S, d_{\mathcal{X}})$ and $(\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$ be metric G -processes, and suppose that $\eta > 0$. Let d be the Hamming average of $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ on $\mathcal{X} \times \mathcal{Y}$. Let $h_1 := h_{\Sigma}^{\text{dq}}(\mu)$ and $h_2 = h_{\Sigma}(\nu)$. We assume the latter is equal to $\underline{h}_{\Sigma}(\nu)$, so there is a $\delta > 0$ such that for any w^* -neighbourhood \mathcal{O} of ν we have

$$\text{cov}_{\delta}(\Omega(\mathcal{O}, \sigma_n), d_{\mathcal{Y}}^{(V_n)}) \geq e^{(h_2 - \eta)|V_n|} \quad (31)$$

for all sufficiently large n .

By shrinking δ further if necessary, and choosing $\varepsilon > 0$ sufficiently small, we may also assume that there are a sequence $n_i \uparrow \infty$ and a sequence of measures $\mu_i \in \text{Prob}(\mathcal{X}^{V_{n_i}})$ such that $\mu_{n_i} \xrightarrow{\text{dq}} \mu$ over $(\sigma_{n_i})_{i \geq 1}$ and

$$\text{cov}_{\varepsilon, \delta}(\mu_i, d_{\mathcal{X}}^{(V_{n_i})}) \geq e^{(h_1 - \eta)|V_{n_i}|}$$

for all sufficiently large i .

Let \mathcal{N} be any w^* -neighbourhood of $\mu \times \nu$. We will prove that

$$\limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{pack}_{\delta/2}(\Omega(\mathcal{N}, \sigma_n), d^{(V_n)}) \geq h_1 + h_2 - 2\eta.$$

Since $\eta > 0$ is arbitrary, and recalling the inequalities (5), this will complete the proof.

By conclusion (ii) of Theorem A, there is a w^* -neighbourhood \mathcal{O} of ν such that

$$\inf_{\mathbf{y} \in \Omega(\mathcal{O}, \sigma_{n_i})} \mu_i\{\mathbf{x} : (\mathbf{x}, \mathbf{y}) \in \Omega(\mathcal{N}, \sigma_{n_i})\} > 1 - \varepsilon \quad (32)$$

for all sufficiently large i (note that (31) guarantees that $\Omega(\mathcal{O}, \sigma_{n_i})$ is nonempty for all sufficiently large i).

Having chosen \mathcal{O} , the inequalities (5) and the lower bound (31) let us choose subsets $F_n \subseteq \Omega(\mathcal{O}, \sigma_n)$ which are δ -separated according to the metrics $d_{\mathbf{y}}^{(V_n)}$ and such that

$$|F_n| \geq e^{(h_2 - \eta)|V_n|}$$

for all sufficiently large n .

For each $\mathbf{y} \in F_n$, let

$$G_{n,\mathbf{y}} := \{\mathbf{x} \in \mathcal{X}^{V_n} : (\mathbf{x}, \mathbf{y}) \in \Omega(\mathcal{N}, \sigma_n)\},$$

so (32) implies that once i is sufficiently large we have $\mu_i(G_{n_i,\mathbf{y}}) > 1 - \varepsilon$ for all $\mathbf{y} \in F_{n_i}$. This requires that

$$\text{cov}_{\delta}(G_{n_i,\mathbf{y}}, d_{\mathcal{X}}^{(V_{n_i})}) \geq \text{cov}_{\varepsilon,\delta}(\mu_i, d_{\mathcal{X}}^{(V_{n_i})}) \quad \forall \mathbf{y} \in F_{n_i}$$

once i is sufficiently large.

Using the inequalities (5) again, we may therefore find further subsets $H_{n,\mathbf{y}} \subseteq G_{n,\mathbf{y}}$ for each $\mathbf{y} \in F_n$ which are δ -separated according to the metrics $d_{\mathcal{X}}^{(V_n)}$ and such that

$$|H_{n_i,\mathbf{y}}| \geq e^{(h_1 - \eta)|V_{n_i}|} \quad \forall \mathbf{y} \in F_{n_i}$$

for all sufficiently large i .

Finally, defining

$$K_n := \{(\mathbf{x}, \mathbf{y}) : \mathbf{y} \in F_n, \mathbf{x} \in H_{n,\mathbf{y}}\}$$

for each n , it follows that these sets are $(\delta/2)$ -separated according to $d^{(V_n)}$; that they are contained in $\Omega(\mathcal{N}, \sigma_n)$; and that

$$\limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log |K_n| \geq \limsup_{i \rightarrow \infty} \frac{1}{|V_{n_i}|} \log |K_{n_i}| \geq h_1 + h_2 - 2\eta. \quad \square$$

REMARK. The above proof really shows that any G -systems (X, μ, T) and (Y, ν, S) satisfy

$$h_{\Sigma}(\mu \times \nu, T \times S) \geq h_{\Sigma}^{\text{dq}}(\mu, T) + \underline{h}_{\Sigma}(\nu, S).$$

This conclusion may also be deduced formally from the statement of Theorem B, by first restricting to a sofic subapproximation which nearly realizes the value $h_{\Sigma}^{\text{dq}}(\mu, T)$, and then restricting to a further subapproximation along which the sofic entropy and lower sofic entropy of (Y, ν, S) agree. \triangleleft

7.2. Proof of Theorem C. The first assertion of Theorem C, that model-measure sofic entropy is subadditive, holds for arbitrary joinings, similarly to sofic entropy itself (see Proposition 3.4).

PROPOSITION 7.1. *Let (X, μ, T) and (Y, ν, S) be G -systems and let λ be a joining of them. Then*

$$h_{\Sigma}^{\text{q}}(\lambda, T \times S) \leq h_{\Sigma}^{\text{q}}(\mu, T) + h_{\Sigma}^{\text{q}}(\nu, S),$$

and similarly if h_{Σ}^{q} is replaced with h_{Σ}^{dq} throughout.

Proof. It suffices to consider two metric G -processes $(\mathcal{X}^G, \mu, S, d_{\mathcal{X}})$ and $(\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$. Let d be the Hamming average of $d_{\mathcal{X}}$ and $d_{\mathcal{Y}}$ on $\mathcal{X} \times \mathcal{Y}$.

Suppose $(\sigma_n)_{i \geq 1}$ is a sofic subapproximation, that $\varepsilon, \delta > 0$, and that $\lambda_i \xrightarrow{\text{q}} \lambda$ over $(\sigma_n)_{i \geq 1}$. By relabelling the subapproximation, we may simply assume that it equals $(\sigma_n)_{n \geq 1}$, and write $\lambda_n \xrightarrow{\text{q}} \lambda$.

Let μ_n and ν_n be the marginals of λ_n on \mathcal{X}^{V_n} and \mathcal{Y}^{V_n} , so Corollary 5.17 gives that $\mu_n \xrightarrow{\text{q}} \mu$ and $\nu_n \xrightarrow{\text{q}} \nu$. Part (i) of Lemma 2.6 gives

$$\text{cov}_{\varepsilon, \delta}(\lambda_n, d^{(V_n)}) \leq \text{cov}_{\varepsilon/2, \delta}(\mu_n, d_{\mathcal{X}}^{(V_n)}) \cdot \text{cov}_{\varepsilon/2, \delta}(\nu_n, d_{\mathcal{Y}}^{(V_n)}).$$

Taking logarithms and normalizing by $|V_n|$, this completes the proof, since ε, δ and λ_n were arbitrary.

The argument is the same if quenched convergence is replaced with doubly-quenched convergence. \square

The reverse inequality required for Theorem C relies on the following lemma, which is another manifestation of the difference between quenched and doubly-quenched convergence.

LEMMA 7.2. If $\mu_n \in \text{Prob}(\mathcal{X}^{V_n})$ and $\nu_n \in \text{Prob}(\mathcal{Y}^{V_n})$ satisfy

$$\text{both } \mu_n \xrightarrow{\text{dq}} \mu \quad \text{and} \quad \nu_n \xrightarrow{\text{dq}} \nu \quad \text{over } \Sigma,$$

then

$$\mu_n \times \nu_n \xrightarrow{\text{dq}} \mu \times \nu \quad \text{over } \Sigma.$$

Proof. Repeatedly applying the implication (i) \implies (iii) in Theorem A gives

$$\begin{aligned} \nu_n \times \nu_n \xrightarrow{q} \nu \times \nu &\implies \mu_n \times \nu_n \times \nu_n \xrightarrow{q} \mu \times \nu \times \nu \\ &\implies \mu_n \times \mu_n \times \nu_n \times \nu_n \xrightarrow{q} \mu \times \mu \times \nu \times \nu, \end{aligned}$$

and hence $\mu_n \times \nu_n \xrightarrow{\text{dq}} \mu \times \nu$. \square

Proof of Theorem C. Let $(\mathcal{X}^G, \mu, S, d_{\mathcal{X}})$, $(\mathcal{Y}^G, \nu, S, d_{\mathcal{Y}})$ and d be as in the proof of Proposition 7.1. That proposition has given subadditivity, so it remains to prove the reverse inequality.

We need only consider the case in which $h_{\Sigma}^{\text{dq}}(\mu), h_{\Sigma}^{\text{dq}}(\nu) \geq 0$, for otherwise subadditivity gives also that $h_{\Sigma}^{\text{dq}}(\mu \times \nu) = -\infty$.

Since we assume that $h_{\Sigma}^{\text{dq}}(\nu) = \underline{h}_{\Sigma}^{\text{dq}}(\nu)$, for any $\eta > 0$ there are $\varepsilon, \delta > 0$ and a sequence $\nu_n \in \text{Prob}(\mathcal{Y}^{V_n})$ such that $\nu_n \xrightarrow{\text{dq}} \nu$ over Σ and

$$\text{cov}_{\sqrt{\varepsilon}, \delta}(\nu_n, d_{\mathcal{Y}}^{(V_n)}) \geq e^{(h_{\Sigma}^{\text{dq}}(\nu) - \eta)|V_n|}$$

for all sufficiently large n . Of course, this remains true if we pass to any sofic subapproximation.

After shrinking ε and δ if necessary, now let $(\sigma_{n_i})_{i \geq 1}$ be a sofic subapproximation and let $\mu_i \in \text{Prob}(\mathcal{X}^{(V_{n_i})})$ be a sequence such that $\mu_i \xrightarrow{\text{dq}} \mu$ over $(\sigma_{n_i})_{i \geq 1}$ and

$$\text{cov}_{\sqrt{\varepsilon}, \delta}(\mu_i, d_{\mathcal{Y}}^{(V_{n_i})}) \geq e^{(h_{\Sigma}^{\text{dq}}(\mu) - \eta)|V_{n_i}|}$$

for all sufficiently large i .

Lemma 7.2 gives that $\mu_i \times \nu_{n_i} \xrightarrow{\text{dq}} \mu \times \nu$ over $(\sigma_{n_i})_{i \geq 1}$, and Corollary 2.7 gives

$$\text{cov}_{\varepsilon, \delta/4}(\mu_i \times \nu_{n_i}, d^{(V_{n_i})}) \geq \text{cov}_{\sqrt{\varepsilon}, \delta}(\mu_i, d_{\mathcal{X}}^{(V_{n_i})}) \cdot \text{cov}_{\sqrt{\varepsilon}, \delta}(\nu_{n_i}, d_{\mathcal{Y}}^{(V_{n_i})}).$$

Taking logarithms and normalizing by $|V_{n_i}|$, this completes the proof.

If $\mu \times \nu$ is ergodic then we may argue in just the same way using only quenched convergence, because in that case

$$\mu_i \xrightarrow{q} \mu \quad \text{and} \quad \nu_{n_i} \xrightarrow{q} \nu \quad \text{imply} \quad \mu_i \times \nu_{n_i} \xrightarrow{q} \mu \times \nu. \quad \square$$

Example 6.3 shows that the assumption that $\mu \times \nu$ is ergodic cannot be dropped. The proof of Theorem C gives the following important special case.

COROLLARY 7.3. *For any G -system (X, μ, T) we have*

$$h_{\Sigma}^{\text{dq}}(\mu^{\times k}, T^{\times k}) = k \cdot h_{\Sigma}^{\text{dq}}(\mu, T) \quad \forall k \geq 1.$$

Here we do not need the assumption that $h_{\Sigma}^{\text{dq}} = \underline{h}_{\Sigma}^{\text{dq}}$. This is because we are now combining (X, μ, T) with itself, so there is no risk that two different sofic subapproximations are needed to obtain the entropies of the ingredient systems.

Proof. If $h_{\Sigma}^{\text{dq}}(\mu, T) = -\infty$ then the result is trivial, so suppose otherwise.

Let $(\mathcal{X}^G, \mu, S, d)$ be a metric G -process. For any $\eta > 0$, there are $\varepsilon, \delta > 0$, a sofic subapproximation $\Sigma' = (\sigma_{n_i})_{i \geq 1}$, and a sequence of measures $\mu_{n_i} \xrightarrow{\text{dq}} \mu$ over Σ' such that

$$\text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})}) \geq e^{(h_{\Sigma'}^{\text{dq}}(\mu) - \eta)|V_{n_i}|}$$

for all sufficiently large i . It follows that

$$h_{\Sigma}^{\text{dq}}(\mu) \geq h_{\Sigma'}^{\text{dq}}(\mu) \geq \underline{h}_{\Sigma'}^{\text{dq}}(\mu) \geq h_{\Sigma'}^{\text{dq}}(\mu) - \eta.$$

Since η was arbitrary, a simple diagonal argument now gives a sofic subapproximation Σ'' such that in fact

$$h_{\Sigma''}^{\text{dq}}(\mu) = \underline{h}_{\Sigma''}^{\text{dq}}(\mu) \geq h_{\Sigma}^{\text{dq}}(\mu) - \eta.$$

A k -fold application of Theorem C with this sofic subapproximation gives that

$$h_{\Sigma}^{\text{dq}}(\mu^{\times k}) \geq h_{\Sigma''}^{\text{dq}}(\mu^{\times k}) = k \cdot h_{\Sigma''}^{\text{dq}}(\mu) \geq k \cdot (h_{\Sigma}^{\text{dq}}(\mu) - \eta).$$

Since η was arbitrary, this completes the proof. \square

8. Processes with finite state spaces

8.1. Alternative formulae for the entropies. Now let (\mathcal{X}^G, μ, S) be a G -process with finite state space \mathcal{X} . Let d be the discrete metric on \mathcal{X} (all distances are zero or one). For this process there are alternative, simpler formulae for the sofic entropy and model-measure sofic entropy. These will be essential in the proof of Theorem D.

PROPOSITION 8.1. *For a G -process (\mathcal{X}^G, μ, S) with $|\mathcal{X}| < \infty$, we have*

$$h_{\Sigma}(\mu) = \inf_{\mathcal{O}} \limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log |\Omega(\mathcal{O}, \sigma_n)|,$$

where \mathcal{O} ranges over w^* -neighbourhoods of μ , and

$$h_{\Sigma}^{\text{dq}}(\mu) := \sup \left\{ \sup_{\varepsilon > 0} \limsup_{i \rightarrow \infty} \frac{1}{|V_{n_i}|} \log \text{cov}_{\varepsilon}(\mu_i) : \right. \\ \left. n_i \uparrow \infty \text{ and } \mu_i \xrightarrow{\text{dq}} \mu \text{ over } (\sigma_{n_i})_{i \geq 1} \right\},$$

where

$$\text{cov}_{\varepsilon}(\mu_i) = \min\{|F| : F \subseteq \mathcal{X}^{V_{n_i}} \text{ with } \mu_i(F) > 1 - \varepsilon\}.$$

This proposition is a consequence of the following bound on the volumes of Hamming balls, which is implied by standard estimates in Information Theory: see, for instance, [10, Section 5.4].

LEMMA 8.2. *If (\mathcal{X}, d) is a nonempty finite set with its discrete metric, then for every $\eta > 0$ there is a $\delta > 0$ such that the following holds. If V is a nonempty finite set and $\mathbf{x} \in \mathcal{X}^V$, then*

$$|B_{\delta}(\mathbf{x})| \leq e^{\eta|V|},$$

where $B_{\delta}(\mathbf{x})$ is the δ -ball around \mathbf{x} for the metric $d^{(V)}$. □

Proof of Proposition 8.1. The formula for $h_{\Sigma}(\mu)$ is easily seen to be equivalent to Bowen’s original definition of sofic entropy for systems with finite generating partitions [4]. Its agreement with $h_{\Sigma}(\mu)$ is therefore contained in [14].

The reasoning for model-measure sofic entropy is very similar. Clearly we always have

$$\text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})}) \leq \text{cov}_{\varepsilon}(\mu_i),$$

and so $h_{\Sigma}^{\text{dq}}(\mu)$ is bounded from above by the right-hand side of the desired formula. On the other hand, for any $\eta > 0$, Lemma 8.2 gives a $\delta > 0$ such that all δ -balls in the metric $d^{(V_{n_i})}$ have size at most $e^{\eta|V_{n_i}|}$. This implies that

$$|B_{\delta}(F)| \leq e^{\eta|V_{n_i}|} |F| \quad \forall F \subseteq \mathcal{X}^{V_{n_i}},$$

and hence that

$$\frac{1}{|V_{n_i}|} \log \text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})}) \geq \frac{1}{|V_{n_i}|} \log \text{cov}_{\varepsilon}(\mu_i) - \eta \quad \forall i \geq 1.$$

Since η can be made arbitrarily small, $h_{\Sigma}^{\text{dq}}(\mu)$ is also bounded from below by the right-hand side of the desired formula. □

8.2. Proof of Theorem D. One half of Theorem D is true for general systems.

LEMMA 8.3. Any G -system (X, μ, T) satisfies

$$h_{\Sigma}^{dq}(\mu, T) \leq h_{\Sigma}^{ps}(\mu, T).$$

Proof. Combining Lemma 6.1 and Corollary 7.3 gives

$$h_{\Sigma}^{dq}(\mu, T) = \frac{1}{k} h_{\Sigma}^{dq}(\mu^{\times k}, T^{\times k}) \leq \frac{1}{k} h_{\Sigma}(\mu^{\times k}, T^{\times k}) \quad \forall k \geq 1,$$

and hence $h_{\Sigma}^{dq}(\mu) \leq h_{\Sigma}^{ps}(\mu)$. □

We prove the reverse half of Theorem D using the assumption that $|\mathcal{X}|$ is finite. In this case we always endow \mathcal{X} with the discrete metric d . As remarked in the Introduction, this gives the result for any G -system that has a finite generating partition, including any ergodic G -system for which $h^{\text{Rok}}(\mu, T) < \infty$ (see [23]).

The proof relies on producing model measures for μ out of good models for $\mu^{\times k}$ for large values of k . This is done using the following proposition.

PROPOSITION 8.4. Let V be any finite set and $\sigma : G \rightarrow \text{Sym}(V)$ any map.

For any $\varepsilon > 0$ and w^* -neighbourhood \mathcal{O} of $\mu \times \mu$, the following holds for any sufficiently large $k \geq 1$: there is a w^* -neighbourhood \mathcal{N} of $\mu^{\times k}$ such that, if

$$\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in \Omega(\mathcal{N}, \sigma),$$

then the measure

$$\varrho := \frac{1}{k} \sum_{i=1}^k \delta_{\mathbf{x}_i} \in \text{Prob}(\mathcal{X}^{V_n})$$

satisfies

$$|\{v : (\Pi_v^{\sigma})_*(\varrho \times \varrho) \in \mathcal{O}\}| > (1 - \varepsilon)|V| \quad \text{and} \quad (\varrho \times \varrho)(\Omega(\mathcal{O}, \sigma)) > 1 - \varepsilon.$$

We consider a neighbourhood \mathcal{O} of $\mu \times \mu$, rather than μ , to ensure that the resulting model measures not only quenched-converge but doubly-quenched-converge.

Proof. It suffices to prove this for a subbasic family of neighbourhoods \mathcal{O} , so we may assume that

$$\mathcal{O} = \left\{ \theta \in \text{Prob}(\mathcal{X}^G \times \mathcal{X}^G) : \int f \otimes h d\theta \approx_{2\kappa} \int f d\mu \int h d\mu \right\}$$

for some $h, f \in C(\mathcal{X}^G)$ with $\|f\|_{\infty}, \|h\|_{\infty} \leq 1$ and some $\kappa > 0$.

Part 1. For each k , define two continuous functions $(\mathcal{X}^G)^k \rightarrow \mathbb{R}$ by

$$F_k(x_1, \dots, x_k) = \frac{1}{k} \sum_{i=1}^k f(x_i) \quad \text{and} \quad H_k(x_1, \dots, x_k) = \frac{1}{k} \sum_{i=1}^k h(x_i),$$

and let

$$U_k := \left\{ (x_1, \dots, x_k) \in (\mathcal{X}^G)^k : F_k(x_1, \dots, x_k) \approx_\kappa \int f d\mu \right. \\ \left. \text{and } H_k(x_1, \dots, x_k) \approx_\kappa \int h d\mu \right\}.$$

This U_k is open, and the Law of Large Numbers gives that $\mu^{\times k}(U_k) \rightarrow 1$ as $k \rightarrow \infty$. Letting

$$\mathcal{N}_{1,k} := \{ \nu \in \text{Prob}((\mathcal{X}^G)^k) : \nu(U_k) > 1 - \varepsilon \},$$

it follows that this is a w^* -neighbourhood of $\mu^{\times k}$ for all sufficiently large k .

If $(\mathbf{x}_1, \dots, \mathbf{x}_k) \in \Omega(\mathcal{N}_{1,k}, \sigma)$, then this asserts that

$$P_{(\mathbf{x}_1, \dots, \mathbf{x}_k)}^\sigma(U_k) > 1 - \varepsilon,$$

and hence that the set

$$V_k^{\text{good}} := \left\{ v \in V : F_k(\Pi_v^\sigma(\mathbf{x}_1), \dots, \Pi_v^\sigma(\mathbf{x}_k)) \approx_\kappa \int f d\mu \right. \\ \left. \text{and } H_k(\Pi_v^\sigma(\mathbf{x}_1), \dots, \Pi_v^\sigma(\mathbf{x}_k)) \approx_\kappa \int h d\mu \right\}$$

has $|V_k^{\text{good}}| > (1 - \varepsilon)|V|$.

Now let ϱ be as in the statement of the proposition, and observe that

$$\int f \otimes h d((\Pi_v^\sigma)_*(\varrho \times \varrho)) = \int (f \circ \Pi_v^\sigma) d\varrho \cdot \int (h \circ \Pi_v^\sigma) d\varrho.$$

The first of these right-hand factors is equal to

$$\int (f \circ \Pi_v^\sigma) d\varrho = \frac{1}{k} \sum_{i=1}^k f(\Pi_v^\sigma(\mathbf{x}_i)) = F_k(\Pi_v^\sigma(\mathbf{x}_1), \dots, \Pi_v^\sigma(\mathbf{x}_k)),$$

and similarly the second is equal to $H_k(\Pi_v^\sigma(\mathbf{x}_1), \dots, \Pi_v^\sigma(\mathbf{x}_k))$. Therefore any $v \in V_k^{\text{good}}$ satisfies

$$\int f \otimes h d((\Pi_v^\sigma)_*(\varrho \times \varrho)) = F_k(\Pi_v^\sigma(\mathbf{x}_1), \dots, \Pi_v^\sigma(\mathbf{x}_k)) \cdot H_k(\Pi_v^\sigma(\mathbf{x}_1), \dots, \Pi_v^\sigma(\mathbf{x}_k)) \approx_{2\kappa} \int f d\mu \int h d\mu,$$

where the last estimate by 2κ uses the fact that both of the factors here lie in $[-1, 1]$.

Hence

$$|\{v : (\Pi_v^\sigma)_*(\varrho \times \varrho) \in \mathcal{O}\}| \geq |V_k^{\text{good}}| > (1 - \varepsilon)|V|$$

once k is sufficiently large.

Part 2. The second part is simpler. The set

$$\mathcal{N}_{2,k} := \left\{ v \in \text{Prob}((\mathcal{X}^G)^k) : \int f(x_i)h(x_j) v(dx_1, \dots, dx_k) \approx_{2\kappa} \int f d\mu \int h d\mu \right. \\ \left. \text{whenever } i, j \in \{1, \dots, k\} \text{ are distinct} \right\}$$

is another a w^* -neighbourhood of $\mu^{\times k}$ for every k . If $(\mathbf{x}_1, \dots, \mathbf{x}_k) \in \Omega(\mathcal{N}_{2,k}, \sigma)$ and $i \neq j$, then

$$\int f \otimes h dP_{(\mathbf{x}_i, \mathbf{x}_j)}^\sigma = \int f(x_i)h(x_j) P_{(\mathbf{x}_1, \dots, \mathbf{x}_k)}^\sigma(dx_1, \dots, dx_k) \approx_{2\kappa} \int f d\mu \int h d\mu :$$

that is, $P_{(\mathbf{x}_i, \mathbf{x}_j)}^\sigma \in \mathcal{O}$. Therefore

$$(\rho \times \rho)(\Omega(\mathcal{O}, \sigma)) = \frac{|\{(i, j) \in \{1, \dots, k\}^2 : P_{(\mathbf{x}_i, \mathbf{x}_j)}^\sigma \in \mathcal{O}\}|}{k^2} \geq \frac{k(k-1)}{k^2},$$

and this is greater than $1 - \varepsilon$ once k is large enough.

Completion. Choose k large enough to satisfy both parts above, and set

$$\mathcal{N} := \mathcal{N}_{1,k} \cap \mathcal{N}_{2,k}. \quad \square$$

Proof of Theorem D. One inequality is already given by Lemma 8.3, so we focus on the other.

Let $\varepsilon > 0$. Let $\mathcal{O}_1 \supseteq \mathcal{O}_2 \supseteq \dots$ be a basis of w^* -neighbourhoods of μ . By Proposition 8.4, there are integers $1 \leq k_1 \leq k_2 \leq \dots$ tending to ∞ and w^* -neighbourhoods \mathcal{N}_j of $\mu^{\times k_j}$ for every j such that the following holds. If

$$\vec{\mathbf{x}} := (\mathbf{x}_1, \dots, \mathbf{x}_{k_j}) \in \Omega(\mathcal{N}_j, \sigma_n) \quad \text{for some } j \text{ and } n,$$

then the measure

$$\varrho_{\vec{x}}^{n,j} := \frac{1}{k_j} \sum_{i=1}^{k_j} \delta_{x_i}$$

satisfies

$$\begin{aligned} & |\{v \in V_n : (\Pi_v^{\sigma_n})_*(\varrho_{\vec{x}}^{n,j} \times \varrho_{\vec{x}}^{n,j}) \in \mathcal{O}_j\}| \\ & > (1 - 2^{-j})|V_n| \quad \text{and} \quad (\varrho_{\vec{x}}^{n,j} \times \varrho_{\vec{x}}^{n,j})(\Omega(\mathcal{O}_j, \sigma_n)) > 1 - 2^{-j}. \end{aligned} \tag{33}$$

Next, by the definition of h_{Σ}^{ps} , we may also choose a subsequence $n_1 < n_2 < \dots$ such that

$$|\Omega(\mathcal{N}_j, \sigma_{n_j})| \geq \exp(k_j(h_{\Sigma}^{\text{ps}}(\mu) - \varepsilon)|V_{n_j}|) \geq 1 \tag{34}$$

for all j . Now set

$$\mu_j := \frac{1}{|\Omega(\mathcal{N}_j, \sigma_{n_j})|} \sum_{\vec{x} \in \Omega(\mathcal{N}_j, \sigma_{n_j})} \varrho_{\vec{x}}^{n_j,j} \quad \text{for } j \geq 1.$$

Since the sets \mathcal{O}_j form a basis of neighbourhoods around $\mu \times \mu$, the bounds (33) imply that for any sequence of single k_j -tuples $\vec{x}_j \in \Omega(\mathcal{N}_j, \sigma_{n_j})$, we have

$$\rho_{\vec{x}_j}^{n_j,j} \xrightarrow{\text{dq}} \mu \quad \text{along } (\sigma_{n_j})_{j \geq 1}.$$

We can quickly strengthen this conclusion as follows: if $\vec{y}_j \in \Omega(\mathcal{N}_j, \sigma_{n_j})$ is any other sequence of k_j -tuples, then the implication (i) \implies (iii) of Theorem A gives that

$$\rho_{\vec{x}_j}^{n_j,j} \times \rho_{\vec{y}_j}^{n_j,j} \xrightarrow{\text{q}} \mu \times \mu.$$

By simply averaging this last assertion, it follows that

$$\mu_j \times \mu_j = \frac{1}{|\Omega(\mathcal{N}_j, \sigma_{n_j})|^2} \sum_{\vec{x}, \vec{y} \in \Omega(\mathcal{N}_j, \sigma_{n_j})} \varrho_{\vec{x}}^{n_j,j} \times \varrho_{\vec{y}}^{n_j,j} \xrightarrow{\text{q}} \mu \times \mu,$$

and hence $\mu_j \xrightarrow{\text{dq}} \mu$.

Finally, let

$$H(2\varepsilon', 1 - 2\varepsilon') := -2\varepsilon' \log(2\varepsilon') - (1 - 2\varepsilon') \log(1 - 2\varepsilon') \quad \text{for } \varepsilon' \in (0, 1/2),$$

and choose ε' so small that

$$2\varepsilon' \log |\mathcal{X}| + H(2\varepsilon', 1 - 2\varepsilon') < \varepsilon.$$

Consider the covering numbers $\text{cov}_{\varepsilon'}(\mu_j)$. For each j , let $F_j \subseteq \mathcal{X}^{V_{n_j}}$ be a minimum-size subset for which $\mu_j(F_j) > 1 - \varepsilon'$. By the definition of μ_j and Chebyshev's Inequality, this implies that at least half of the tuples $(\mathbf{x}_1, \dots, \mathbf{x}_{k_j}) \in \Omega(\mathcal{N}_{n_j}, \sigma_{n_j})$ satisfy

$$Q_{(\mathbf{x}_1, \dots, \mathbf{x}_{k_j})}^{n_j, j}(F_j) = \frac{|\{i \in \{1, \dots, k_j\} : \mathbf{x}_i \in F_j\}|}{k_j} > 1 - 2\varepsilon'.$$

On the other hand, a simple estimate using volumes of Hamming balls (see [10, Section 5.4]) gives

$$\left| \left\{ (\mathbf{x}_1, \dots, \mathbf{x}_{k_j}) \in (\mathcal{X}^{V_{n_j}})^{k_j} : \frac{|\{i \in \{1, \dots, k_j\} : \mathbf{x}_i \in F_j\}|}{k_j} > 1 - 2\varepsilon' \right\} \right| \leq |F_j|^{k_j} \cdot |\mathcal{X}|^{2\varepsilon' \cdot k_j \cdot |V_{n_j}|} \cdot 2^{H(2\varepsilon', 1-2\varepsilon')k_j},$$

where the last factor estimates the number of ways of choosing at most $2\varepsilon'k_j$ coordinates $i \in \{1, \dots, k_j\}$ at which to allow $\mathbf{x}_i \notin F_j$.

Therefore

$$\begin{aligned} \frac{1}{2} |\Omega(\mathcal{N}_{n_j}, \sigma_{n_j})| &\leq |F_j|^{k_j} \cdot |\mathcal{X}|^{2\varepsilon' \cdot k_j \cdot |V_{n_j}|} \cdot 2^{H(2\varepsilon', 1-2\varepsilon')k_j} \\ &= (\text{cov}_{\varepsilon'}(\mu_j))^{k_j} \cdot |\mathcal{X}|^{2\varepsilon' \cdot k_j \cdot |V_{n_j}|} \cdot 2^{H(2\varepsilon', 1-2\varepsilon')k_j}. \end{aligned}$$

Combining this with (34) and rearranging, we obtain

$$\frac{1}{|V_{n_j}|} \log \text{cov}_{\varepsilon'}(\mu_j) \geq h_{\Sigma}^{\text{ps}}(\mu) - \varepsilon - 2\varepsilon' \log |\mathcal{X}| - \frac{H(2\varepsilon', 1 - 2\varepsilon') \cdot \log 2}{|V_{n_j}|} - o\left(\frac{1}{|k_j|}\right),$$

and this lower bound is greater than $h_{\Sigma}^{\text{ps}}(\mu) - 2\varepsilon$ for all sufficiently large j . Since ε was arbitrary, this shows that $h_{\Sigma}^{\text{dq}}(\mu) \geq h_{\Sigma}^{\text{ps}}(\mu)$. □

Proof of Corollary D'. Theorem B has already proved that

$$h_{\Sigma}^{\text{dq}}(\mu, T) \leq h_{\Sigma}(\mu \times \nu, T \times S) - h_{\Sigma}(\nu, S)$$

for any other G -system (Y, ν, S) satisfying $h_{\Sigma}(\nu, S) = \underline{h}_{\Sigma}(\nu, S)$.

On the other hand, if (X, μ, T) has a finite generating partition, then Theorem D shows that

$$\frac{1}{k} h_{\Sigma}(\mu^{\times k}, T^{\times k}) \longrightarrow h_{\Sigma}^{\text{dq}}(\mu, T).$$

In particular, for any $\varepsilon > 0$, there must be infinitely many k for which

$$h_{\Sigma}(\mu^{\times(k+1)}, T^{\times(k+1)}) \leq h_{\Sigma}(\mu^{\times k}, T^{\times k}) + h_{\Sigma}^{\text{dq}}(\mu, T) + \varepsilon.$$

Letting $(Y, \nu, S) := (X^k, \mu^{\times k}, T^{\times k})$ and rearranging, this becomes

$$h_{\Sigma}(\mu \times \nu, T \times S) - h_{\Sigma}(\nu, S) \leq h_{\Sigma}^{\text{dq}}(\mu, T) + \varepsilon.$$

Since we also have $h_{\Sigma}(\nu, S) = \underline{h}_{\Sigma}(\nu, S)$ for this system (Y, ν, S) by assumption, these examples complete the proof. \square

Another corollary seems worth including at this point. For any G -system (X, μ, T) , the definition of h_{Σ}^{ps} gives that

$$h_{\Sigma}^{\text{ps}}(\mu^{\times k}, T^{\times k}) = \lim_{n \rightarrow \infty} \frac{1}{n} h_{\Sigma}(\mu^{\times kn}, T^{\times kn}) = k \cdot h_{\Sigma}^{\text{ps}}(\mu, T) \quad \forall k \geq 1.$$

If (Y, ν, S) is another system with the property that

$$h_{\Sigma}(\nu^{\times k}, S^{\times k}) = \underline{h}_{\Sigma}(\nu^{\times k}, S^{\times k}) \quad \forall k \geq 1, \quad (35)$$

then we may take k th Cartesian powers of both systems in Theorem B to obtain

$$\frac{1}{k} h_{\Sigma}((\mu \times \nu)^{\times k}, (T \times S)^{\times k}) \geq h_{\Sigma}^{\text{ps}}(\mu, T) + \frac{1}{k} h_{\Sigma}(\nu^{\times k}, S^{\times k}) \quad \forall k \geq 1.$$

Letting $k \rightarrow \infty$, we conclude that h_{Σ}^{ps} is strictly additive.

COROLLARY 8.5. *If (X, μ, T) has a finite generating partition and (35) is satisfied then*

$$h_{\Sigma}^{\text{ps}}(\mu \times \nu, T \times S) = h_{\Sigma}^{\text{ps}}(\mu, T) + h_{\Sigma}^{\text{ps}}(\nu, S). \quad \square$$

8.3. Some remarks on systems without finite generating partitions. I do not know whether $h_{\Sigma}^{\text{ps}} = h_{\Sigma}^{\text{dq}}$ for arbitrary systems. If this is so, then Corollaries D' and 8.5 can be extended to them. Also, Theorem B could be rewritten with h_{Σ}^{dq} replaced by h_{Σ}^{ps} , so that model measures are not needed to give a meaningful lower bound on the sofic entropy of a Cartesian product.

The difficulty in the general case seems to be the following. Written out in full, the power-stabilized entropy for a metric G -process $(\mathcal{X}^G, \mu, S, d)$ is

$$h_{\Sigma}^{\text{ps}}(\mu) := \limsup_{k \rightarrow \infty} \frac{1}{k} \sup_{\delta > 0} \inf_{\text{Int}(\mathcal{O}) \ni \mu^{\times k}} \limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{cov}_{\delta}(\Omega(\mathcal{O}, \sigma_n), (d^{(k)})^{(V_n)}).$$

If this is nonnegative, then for every $k \geq 1$ and $\varepsilon > 0$ there is a $\delta_k > 0$ such that

$$\inf_{\text{Int}(\mathcal{O}) \ni \mu^{\times k}} \limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{cov}_{\delta_k}(\Omega(\mathcal{O}, \sigma_n), (d^{(k)})^{(V_n)}) \geq k(h_{\Sigma}^{\text{ps}}(\mu) - \varepsilon).$$

However, as far as I know, it could happen that we must choose smaller and smaller values of δ_k as $k \rightarrow \infty$. On the other hand, in order to make contact with h_{Σ}^{dq} , we must find *fixed* $\varepsilon, \delta > 0$ and measures satisfying $\mu_i \xrightarrow{\text{dq}} \mu$ along $(\sigma_{n_i})_{i \geq 1}$ such that $\text{cov}_{\varepsilon, \delta}(\mu_i, d^{(V_{n_i})})$ grows fast enough as $i \rightarrow \infty$. If the choice of δ_k tends to 0 as $k \rightarrow \infty$, and then we construct the measures μ_j as in the proof of Theorem D, we do not obtain control over $\text{cov}_{\varepsilon, \delta}(\mu_j, d^{(V_{n_j})})$ for any fixed $\delta > 0$.

To get around this problem, one could simply redefine h_{Σ}^{ps} so that the supremum over δ appears on the outside: let us set

$$\tilde{h}_{\Sigma}^{\text{ps}}(\mu) := \sup_{\delta > 0} \left[\limsup_{k \rightarrow \infty} \frac{1}{k} \inf_{\text{Int}(\mathcal{O}) \ni \mu^{\times k}} \limsup_{n \rightarrow \infty} \frac{1}{|V_n|} \log \text{cov}_{\delta}(\Omega(\mathcal{O}, \sigma_n), (d^{(k)})^{(V_n)}) \right].$$

Using this quantity, the construction used to prove Theorem D does generalize quite easily, leading to the inequality

$$\tilde{h}_{\Sigma}^{\text{ps}}(\mu) \leq h_{\Sigma}^{\text{dq}}(\mu).$$

However, now the argument of Lemma 8.3 seems to run into difficulty, and I cannot show that

$$\tilde{h}_{\Sigma}^{\text{ps}}(\mu) \geq h_{\Sigma}^{\text{dq}}(\mu).$$

Thus, Theorem D will hold for arbitrary systems if the supremum over δ may be exchanged with the limit supremum over k in the formula for $h_{\Sigma}^{\text{ps}}(\mu)$. If \mathcal{X} is finite and d is the discrete metric, then Proposition 8.1 lets one switch to counting individual models, so that δ disappears altogether from $h_{\Sigma}^{\text{ps}}(\mu)$ and $h_{\Sigma}^{\text{dq}}(\mu)$. This is why the proofs above could be completed when \mathcal{X} is finite.

9. Co-induced systems

Consider the setting of Theorem E. In view of Lemma 6.1, that theorem will follow if we show that

$$h_{\Sigma \times T}(\mu^{\times H}, \text{CInd}_G^{G \times H} T) \leq h_{\Sigma \times T}^{\text{dq}}(\mu^{\times H}, \text{CInd}_G^{G \times H} T). \tag{36}$$

As usual, we can assume that we start with a metric G -process $(\mathcal{X}^G, \mu, S, d)$. After co-induction this simply becomes $(\mathcal{X}^{G \times H}, \mu^{\times H}, \tilde{S}, d)$, where \tilde{S} is the right-shift action of $G \times H$.

To prove Theorem E, we also need to use the left action \tilde{T} of H on $\mathcal{X}^{G \times H}$:

$$\tilde{T}^h((x_{g,k})_{(g,k) \in G \times H}) = (x_{g,h^{-1}k})_{(g,k) \in G \times H}.$$

The product measure $\mu^{\times H}$ is invariant under this action, as well as under \tilde{S} : this special feature of the measure is crucial for the proof. Since H is infinite, the H -system $(\mathcal{X}^{G \times H}, \mu^{\times H}, \tilde{T})$ is weakly mixing.

The action \tilde{T} commutes with \tilde{S} , and so each transformation \tilde{T}^h is an automorphism of $(X^H, \mu^{\times H}, \tilde{S})$. Therefore, the results of Section 4 apply to each of these transformations. Each \tilde{T}^h is already defined coordinate-wise by the $\{(e_G, h^{-1})\}$ -local map

$$\mathcal{X}^{G \times H} \longrightarrow \mathcal{X} : (x_{g,k})_{g,k} \mapsto x_{e_G, h^{-1}},$$

which is 1-Lipschitz from $d^{((e_G, h^{-1}))}$ to d . Therefore, there is no need to introduce other AL approximations to these maps.

For each σ_n, τ_n and $h \in H$, the map

$$\text{id}_{V_n} \times \tau_n^h : V_n \times W_n \longrightarrow V_n \times W_n$$

has an ‘approximate adjoint’ defined by setting

$$\rho_n^h : \mathcal{X}^{V_n \times W_n} \longrightarrow \mathcal{X}^{V_n \times W_n} : (x_{v,w})_{v \in V_n, w \in W_n} \mapsto (x_{v, \sigma_n^{h^{-1}}(w)})_{v \in V_n, w \in W_n}.$$

Such maps were already discussed in the remark following Lemma 3.1: as explained there, they become useful only now that our measure $\mu^{\times H}$ is also left-shift-invariant. In terms of these, a special case of Lemma 4.8 translates as follows.

LEMMA 9.1. *If $F \subseteq G \times H$ and $E \subseteq H$ are finite, then the following holds w.h.p in $(v, w) \in V_n \times W_n$:*

$$\Pi_{(v,w)}^{\sigma_n \times \tau_n}(\rho_n^h(\cdot))|_F = (\tilde{T}^h(\Pi_{(v,w)}^{\sigma_n \times \tau_n}(\cdot)))|_F \quad \forall h \in E. \quad \square$$

LEMMA 9.2. *Let E_1, E_2, \dots be any sequence of finite subsets of H satisfying $|E_m| \longrightarrow \infty$. Also, suppose that $\theta_n \in \text{Prob}(\mathcal{X}^{V_n \times W_n})$ is any sequence of measures satisfying*

$$\theta_n(\Omega(\mathcal{O}, \sigma_n \times \tau_n)) \longrightarrow 1 \quad \text{as } n \longrightarrow \infty \quad (37)$$

for any w^* -neighbourhood \mathcal{O} of $\mu^{\times H}$.

Provided the sequence $m_1 \leq m_2 \leq \dots$ grows sufficiently slowly, the sequence of measures

$$\mu_n := \frac{1}{|E_{m_n}|} \sum_{h \in E_{m_n}} (\rho_n^h)_* \theta_n$$

doubly-quenched-converges to $\mu^{\times H}$ over $\Sigma \times \mathbb{T}$.

Proof. It suffices to show that $\mu_n \xrightarrow{\text{lw}^*} \mu^{\times H}$; since the co-induced system is weakly mixing, this implies doubly-quenched convergence.

Also, it suffices to consider a subbasic w^* -neighbourhood of $\mu^{\times H}$, so let

$$\mathcal{O} := \left\{ \nu \in \text{Prob}(\mathcal{X}^{G \times H}) : \int f d\nu \approx_\kappa \int f d\mu^{\times H} \right\}$$

for some local function $f \in C(\mathcal{X}^{G \times H})$ and $\kappa > 0$.

Now for each m let

$$U_m = \left\{ x \in \mathcal{X}^{G \times H} : \frac{1}{|E_m|} \sum_{h \in E_m} f(\tilde{T}^h x) \approx_\kappa \int f d\mu \right\}.$$

Since $|E_m| \rightarrow \infty$, the Law of Large Numbers gives that $\mu^{\times H}(U_m) \rightarrow 1$ as $m \rightarrow \infty$ (note that this works even if there is no ergodic theorem over the sets E_m for general H -systems). Choose real values $\alpha_m < \mu^{\times H}(U_m)$ which still tend to 1, and for each m let

$$\mathcal{N}_m := \{ \nu \in \text{Prob}(\mathcal{X}^{G \times H}) : \nu(U_m) > \alpha_m \}.$$

Each \mathcal{N}_m is a w^* -neighbourhood of $\mu^{\times H}$, and so (37) implies that

$$\theta_n(\Omega(\mathcal{N}_{m_n}, \sigma_n \times \tau_n)) \rightarrow 1 \quad \text{as } n \rightarrow \infty$$

provided $m_1 \leq m_2 \leq \dots$ grows sufficiently slowly. In terms of empirical distributions, this implies that

$$\begin{aligned} & \int \frac{|\{(v, w) \in V_n \times W_n : \Pi_{(v,w)}^{\sigma_n \times \tau_n}(\mathbf{x}) \in U_{m_n}\}|}{|V_n \times W_n|} \theta_n(d\mathbf{x}) \\ &= \frac{1}{|V_n \times W_n|} \sum_{(v,w) \in V_n \times W_n} \theta_n\{\mathbf{x} : \Pi_{(v,w)}^{\sigma_n \times \tau_n}(\mathbf{x}) \in U_{m_n}\} \rightarrow 1 \end{aligned} \quad (38)$$

as $n \rightarrow \infty$.

Finally, we have

$$\int f d((\Pi_{(v,w)}^{\sigma_n \times \tau_n})_* \mu_n) = \int_{\mathcal{X}^{V_n \times W_n}} \left(\frac{1}{|E_{m_n}|} \sum_{h \in E_{m_n}} f(\Pi_{(v,w)}^{\sigma_n \times \tau_n}(\rho_n^h(\mathbf{x}))) \right) \theta_n(d\mathbf{x}).$$

Since f is a local function, another appeal to Lemma 9.1 gives that, w.h.p. in (v, w) , this is equal to

$$\int_{\mathcal{X}^{V_n \times W_n}} \left(\frac{1}{|E_{m_n}|} \sum_{h \in E_{m_n}} f(\tilde{T}^h(\Pi_{(v,w)}^{\sigma_n \times \tau_n}(\mathbf{x}))) \right) \theta_n(d\mathbf{x}).$$

Recalling the definition of U_{m_n} and the inequality (38), this, in turn, lies within κ of $\int f d\mu$ w.h.p. in (v, w) . That is,

$$(\Pi_{(v,w)}^{\sigma_n \times \tau_n})_* \mu_n \in \mathcal{O} \quad \text{w.h.p. in } (v, w),$$

as required. □

Proof of Theorem E. We need only prove the inequality (36) for the co-induced process. If $h_{\Sigma \times T}(\mu^{\times H}) = -\infty$ then the result is trivial, so suppose otherwise, let $h_1 < h_{\Sigma \times T}(\mu^{\times H})$ be arbitrary, and let h_2 lie strictly between these two values.

Let d be a compact generating metric on \mathcal{X} , and let E_1, E_2, \dots be finite subsets of H with $|E_m| \rightarrow \infty$.

From the definition of $h_{\Sigma \times T}(\mu^{\times H})$, it follows that there are $\delta > 0$ and a sequence of subsets $A_n \subseteq \mathcal{X}^{V_n \times W_n}$ such that:

- (i) each A_n is δ -separated according to $d^{(V_n \times W_n)}$;
- (ii) for every w^* -neighbourhood \mathcal{O} of $\mu^{\times H}$ we have

$$A_n \subseteq \Omega(\mathcal{O}, \sigma_n \times \tau_n)$$

for all sufficiently large n ; and

- (iii) $|A_n| \geq \exp(h_2 |V_n| |W_n|)$ for infinitely many n .

By passing to a subsequence $n_1 < n_2 < \dots$, we may now assume that (iii) holds for all sufficiently large n , and in particular that $A_n \neq \emptyset$ for every n .

Let θ_n be the uniform measure on A_n for each n . Then condition (ii) above shows that these satisfy the hypotheses of Lemma 9.2, and so

$$\mu_n := \frac{1}{|E_{m_n}|} \sum_{h \in E_{m_n}} (\rho_n^h)_* \theta_n \xrightarrow{\text{dq}} \mu^{\times H}$$

for some $m_1 \leq m_2 \leq \dots$ tending slowly to ∞ .

Finally, for each n let $B_n \subseteq \mathcal{X}^{V_n \times W_n}$ be a subset of minimal cardinality such that

$$\mu_n(B_{\delta/2}(B_n)) > 1/2.$$

From the definition of μ_n , this requires that

$$\theta_n((\rho_n^h)^{-1}(B_n)) = \frac{|A_n \cap (\rho_n^h)^{-1}(B_n)|}{|A_n|} > \frac{1}{2} \quad \text{for some } h \in F_{m_n}.$$

Since each ρ_n^h is an isometry of the metric $d^{(V_n \times W_n)}$, this and property (i) require that

$$|B_n| \geq |A_n|/2.$$

Therefore, since property (iii) now holds for all sufficiently large n , we have

$$\text{cov}_{1/2, \delta/2}(\mu_n, d^{(V_n \times W_n)}) = |B_n| > \exp(h_1 |V_n| |W_n|)$$

for all sufficiently large n , and thus

$$h_{\Sigma \times T}^{\text{dq}}(\mu^{\times H}) \geq h_1.$$

Since $h_1 < h_{\Sigma \times T}(\mu^{\times H})$ was arbitrary, this completes the proof. \square

In case G is trivial, Theorem E just asserts that all our sofic entropy notions coincide for Bernoulli H -systems. However, in that case one could give a much simpler proof: if $(\mathcal{X}^H, \nu^{\times H}, S)$ is a Bernoulli H -process and d a compact generating metric on \mathcal{X} , then Lemma 5.11 gives that $\nu^{\times W_n} \xrightarrow{\text{dq}} \nu^{\times H}$, and this sequence of measures achieves the full sofic entropy of the process, which just equals the Shannon entropy of ν .

It is worth comparing Theorem E with [6, Theorem 4.1], which gives other sufficient conditions for a G -system (X, μ, T) to satisfy $h_{\Sigma}(\mu, T) = h_{\Sigma}^{\text{q}}(\mu, T)$. Bowen's assumptions are that G is residually finite and that Σ consists of quotients by finite-index normal subgroups of G . This looks quite different from Theorem E: in the first place, Theorem E assumes that the system is of a special kind, whereas Bowen's condition is mostly about the sofic approximation. However, the proofs do have a similar flavour. A key point in the proof of Theorem E is that the left-shift action of H on X^H commutes with the co-induced $(G \times H)$ -action and is weakly mixing. Bowen's proof also requires that there be a 'sufficiently large' group commuting with a given action: in his case, that the G -actions on the sofic approximations V_n commute with some transitive actions from the other side. In both cases, the proof uses this large centralizer for some auxiliary averaging, which converts single good models into measures. It would be interesting to find some way of unifying these two sufficient conditions. It would also be worth knowing whether Bowen's condition can be generalized in such a way that every sofic group has some sofic approximation which satisfies it.

In light of the role played by the left-shift H -action in the proof of Theorem E, I suspect it might have a far-reaching generalization as follows. Let (X, μ, T) be a G -system. Let $\text{Aut}(X, \mu)$ denote the group of all measure-preserving automorphisms of the probability space (X, μ) , up to agreement μ -almost everywhere. This is a Polish group in its coarse topology. The G -action T defines a homomorphism $G \rightarrow \text{Aut}(X, \mu)$, and we define the **centralizer** of T to be the subgroup of elements of $\text{Aut}(X, \mu)$ which commute with the image of that homomorphism.

CONJECTURE 9.3. *If the centralizer of T is ergodic, then $h_{\Sigma}(\mu, T) = h_{\Sigma}^q(\mu, T)$. If the centralizer is weakly mixing, then $h_{\Sigma}(\mu, T) = h_{\Sigma}^{\text{dq}}(\mu, T)$.*

For example, the centralizer of the co-induced $(G \times H)$ -system $(X^H, \mu^{\times H}, \text{CInd}_G^{G \times H} T)$ includes the left-shift action of H , which is Bernoulli and therefore weakly mixing.

On the other hand, I do not believe that the conditions in Conjecture 9.3 are necessary. On the contrary, if G is amenable, then one should always have

$$h_{\Sigma}(\mu, T) = h_{\Sigma}^q(\mu, T) = h_{\Sigma}^{\text{dq}}(\mu, T) = h_{\text{KS}}(\mu, T)$$

for any sofic approximation Σ . The equality of the first and last values here is shown in [7], and I think similar methods are able to prove the others.

10. Some directions for further study

Just as for sofic entropy itself, the following basic fact is not known for the model-measure sofic entropies. It suggests a major gap in our present understanding.

QUESTION 10.1. Are there a sofic group G , a G -system (X, μ, T) , and two different sofic approximations Σ and Σ' to G such that

$$0 \leq h_{\Sigma}^q(\mu, T) < h_{\Sigma'}^q(\mu, T),$$

or similarly with h^q replaced by h^{dq} ? What if G is a free group?

There are cases in which some sofic approximations give a nonnegative value, while others give $-\infty$, just as there are for sofic entropy itself.

The following related questions are also open.

QUESTION 10.2. Are there a sofic group G , a G -system (X, μ, T) , and a sofic approximation Σ such that at least two of the quantities

$$h_{\Sigma}(\mu, T), \quad h_{\Sigma}^q(\mu, T) \quad \text{and} \quad h_{\Sigma}^{\text{dq}}(\mu, T)$$

are nonnegative, but are not equal? What if G is a free group?

QUESTION 10.3. Are there a sofic group G , a G -system (X, μ, T) , and a sofic approximation Σ such that the sequence

$$\frac{1}{k} h_{\Sigma}(\mu^{\times k}, T^{\times k}), \quad k \geq 1,$$

contains at least two distinct nonnegative values? How about for h_{Σ}^q ? Are there examples in which

$$0 \leq h_{\Sigma}^{\text{ps}}(\mu, T) < h_{\Sigma}(\mu, T)?$$

What if G is a free group?

Another possibility that might be worth pursuing is that certain choices of sofic approximation give some simplification of the entropy theories.

QUESTION 10.4. Let G be a sofic group. Is there a sofic approximation Σ to G such that

$$h_{\Sigma}(\mu, T) = h_{\Sigma}^{\text{dq}}(\mu, T)$$

for all G -systems (X, μ, T) ?

In this case, I think an obvious candidate is to start with an arbitrary sofic approximation $\Sigma_0 = (\sigma_n : G \rightarrow \text{Sym}(V_n))_{n \geq 1}$, and then let Σ be the sequence

$$\sigma_n^{\times m_n} : G \rightarrow \text{Sym}(V_n^{m_n})$$

for some slowly growing sequence $m_1 \leq m_2 \leq \dots$. It might be that for this Σ , some variation of the averaging argument used to prove Theorem E would answer the above question positively. I have not pursued this idea very far.

In case G is a free group, Bowen introduced another entropy-like invariant called the ‘ f -invariant’ in [8], and denoted it by $f(\mu, T)$. In [3], he then showed that $f(\mu, T)$ may be expressed as a kind of average of sofic entropies over random sofic approximations. As a result, the f -invariant may have better behaviour than the sofic entropy along any given sofic approximation. It would be interesting to study its additivity properties using the method of the present paper.

Acknowledgements

I am grateful for several discussions with Miklós Abért, Lewis Bowen, David Fisher, Alex Lubotzky and Brandon Seward. After the first versions of this paper appeared, I learned that Miklós Abért and Benjamin Weiss have been studying some quite similar definitions using measures on model spaces. I also thank the referee for correcting several small errors. This research was supported partly by a fellowship from the Clay Mathematics Institute and partly by the Simons Collaboration on Algorithms and Geometry.

References

- [1] B. Bekka, P. de la Harpe and A. Valette, *Kazhdan's Property (T)*, New Mathematical Monographs, 11 (Cambridge University Press, Cambridge, 2008).
- [2] I. Benjamini and O. Schramm, 'Recurrence of distributional limits of finite planar graphs', *Electron. J. Probab.* **6**(23) (2001), 13 pp (electronic).
- [3] L. Bowen, 'The ergodic theory of free group actions: entropy and the f -invariant', *Groups Geom. Dyn.* **4**(3) (2010), 419–432.
- [4] L. Bowen, 'Measure conjugacy invariants for actions of countable sofic groups', *J. Amer. Math. Soc.* **23**(1) (2010), 217–245.
- [5] L. Bowen, 'Non-abelian free group actions: Markov processes, the Abramov–Rohlin formula and Yuzvinskii's formula', *Ergodic Theory Dynam. Systems* **30**(6) (2010), 1629–1663.
- [6] L. Bowen, 'Entropy for expansive algebraic actions of residually finite groups', *Ergodic Theory Dynam. Systems* **31**(3) (2011), 703–718.
- [7] L. Bowen, 'Sofic entropy and amenable groups', *Ergodic Theory Dynam. Systems* **32**(2) (2012), 427–466.
- [8] L. P. Bowen, 'A measure-conjugacy invariant for free group actions', *Ann. of Math. (2)* **171**(2) (2010), 1387–1400.
- [9] A. H. Dooley and G. Zhang, 'Co-induction in dynamical systems', *Ergodic Theory Dynam. Systems* **32**(3) (2012), 919–940.
- [10] C. M. Goldie and R. G. E. Pinch, *Communication Theory*, London Mathematical Society Student Texts, 20 (Cambridge University Press, Cambridge, 1991).
- [11] M. Gromov, 'Endomorphisms of symbolic algebraic varieties', *J. Eur. Math. Soc. (JEMS)* **1**(2) (1999), 109–197.
- [12] B. Hayes, 'Polish models and sofic entropy', *J. Inst. Math. Jussieu*. Available on CJO 2016 doi:[10.1017/S1474748015000468](https://doi.org/10.1017/S1474748015000468).
- [13] A. S. KeCHRIS, *Global Aspects of Ergodic Group Actions*, Mathematical Surveys and Monographs, 160 (American Mathematical Society, Providence, RI, 2010).
- [14] D. Kerr and H. Li, 'Bernoulli actions and infinite entropy', *Groups Geom. Dyn.* **5**(3) (2011), 663–672.
- [15] D. Kerr and H. Li, 'Entropy and the variational principle for actions of sofic groups', *Invent. Math.* **186**(3) (2011), 501–558.
- [16] D. Kerr and H. Li, 'Soficity, amenability, and dynamical entropy', *Amer. J. Math.* **135**(3) (2013), 721–761.
- [17] H. Li, 'Compact group automorphisms, addition formulas and Fuglede–Kadison determinants', *Ann. of Math. (2)* **176**(1) (2012), 303–347.
- [18] A. Lubotzky, *Discrete Groups, Expanding Graphs and Invariant Measures*, Modern Birkhäuser Classics (Birkhäuser Verlag, Basel, 2010), with an appendix by Jonathan D. Rogawski, reprint of the 1994 edition.
- [19] M. Mézard and A. Montanari, *Information, Physics, and Computation*, Oxford Graduate Texts (Oxford University Press, Oxford, 2009).
- [20] A. Montanari, E. Mossel and A. Sly, 'The weak limit of Ising models on locally tree-like graphs', *Probab. Theory Related Fields* **152**(1–2) (2012), 31–51.
- [21] E. Mossel, J. Neeman and A. Sly, 'Reconstruction and estimation in the planted partition model', *Probab. Theory Related Fields* **162**(3–4) (2015), 431–461.
- [22] K. E. Petersen, *Ergodic Theory*, (Cambridge University Press, Cambridge, 1983).
- [23] B. Seward, Krieger's finite generator theorem for actions of countable groups. I. Preprint, [arXiv:1405.3604](https://arxiv.org/abs/1405.3604).
- [24] B. Weiss, 'Sofic groups and dynamical systems', *Sankhyā Ser. A* **62**(3) (2000), 350–359. Ergodic theory and harmonic analysis (Mumbai, 1999).