

MEASURE CONCENTRATION AND THE WEAK PINSKER PROPERTY

by TIM AUSTIN

ABSTRACT

Let (X, μ) be a standard probability space. An automorphism T of (X, μ) has the weak Pinsker property if for every $\varepsilon > 0$ it has a splitting into a direct product of a Bernoulli shift and an automorphism of entropy less than ε . This property was introduced by Thouvenot, who asked whether it holds for all ergodic automorphisms.

This paper proves that it does. The proof actually gives a more general result. Firstly, it gives a relative version: any factor map from one ergodic automorphism to another can be enlarged by arbitrarily little entropy to become relatively Bernoulli. Secondly, using some facts about relative orbit equivalence, the analogous result holds for all free and ergodic measure-preserving actions of a countable amenable group.

The key to this work is a new result about measure concentration. Suppose now that μ is a probability measure on a finite product space A^n , and endow this space with its Hamming metric. We prove that μ may be represented as a mixture of other measures in which (i) most of the weight in the mixture is on measures that exhibit a strong kind of concentration, and (ii) the number of summands is bounded in terms of the difference between the Shannon entropy of μ and the combined Shannon entropies of its marginals.

CONTENTS

1. Introduction	1
2. Background on kernels, entropy, and mutual information	9
3. Total correlation and dual total correlation	16
4. Background on metrics and measure transportation	24
5. Preliminaries on measure concentration	28
6. A relative of Theorem B	43
7. Completed proof of Theorem B	54
8. Proof of Theorem C	61
9. A reformulation: extremality	71
10. Background on relative entropy in ergodic theory	79
11. Relative finite determination	84
12. Relative extremality	89
13. From extremality to finite determination	91
14. A quantitative step towards Theorem A	96
15. Completed proof of Theorem A	107
16. Some known consequences and open questions	110
Acknowledgements	116
References	116

1. Introduction

The weak Pinsker property. — In this paper, a **measure-preserving system** or **automorphism** consists of (i) a standard probability space (X, μ) and (ii) a measurable and μ -preserving transformation $T : X \rightarrow X$ with a measurable inverse. We write \mathcal{B}_X for the σ -algebra of X when it is needed.

One of the most fundamental invariants of an automorphism is its entropy. Kolmogorov and Sinai first brought the notion of entropy to bear on questions of ergodic theory in [45, 46, 92]. They showed that Shannon's entropy rate for stationary, finite-state

stochastic processes is monotone under factor maps of processes, and deduced several new non-isomorphism results, most famously among Bernoulli shifts.

In [93], Sinai proved a result about entropy and ergodic theory of a different nature. He showed that any ergodic automorphism (X, μ, T) of positive entropy h admits factor maps onto any Bernoulli shift of entropy at most h . This marked a turn in entropy theory towards synthetic results, which deduce the existence of an isomorphism or factor map from some modest assumption about invariants, such as Sinai's inequality between entropies. These synthetic results are generally much more delicate. They produce maps between systems using complicated combinatorics and analysis, and the maps that result are rarely canonical.

The most famous result of this kind is Ornstein's theorem [67, 68], which shows that two Bernoulli shifts of equal (finite or infinite) entropy are isomorphic. Together with Kolmogorov and Sinai's original entropy calculations, this completes the classification of Bernoulli shifts up to isomorphism.

The real heart of Ornstein's work goes far beyond Bernoulli shifts: it provides necessary and sufficient conditions for an arbitrary stationary process to be isomorphic to a Bernoulli shift. A careful introduction to this sophisticated theory can be found in the contemporary monographs [90] and [73]. An intuitive discussion without many proofs can be found in [69], and the historical account in [41] puts Ornstein's work in a broader context. Finally, the survey [107] collects statements and references for much of the more recent work in this area.

Since Ornstein's original work, several other necessary and sufficient conditions for Bernoullicity have also been explored, with far-reaching consequences for the classification of automorphisms. On the one hand, many well-known systems turn out to be isomorphic to Bernoulli shifts: see, for instance, [107, Section 6] and the references given there. On the other hand, many examples can be constructed that are distinct from Bernoulli shifts, in spite of having related properties such as being \mathbf{K} -automorphisms [37, 70, 75].

In particular, the subsequent non-isomorphism results included counterexamples [71, 72] to an early, ambitious conjecture of Pinsker [79]. Let us say that a system (X, μ, T) **splits** into two of its factors (Y, ν, S) and (Y', ν', S') if there is an isomorphism

$$(X, \mu, T) \longrightarrow (Y \times Y', \nu \times \nu', S \times S').$$

Such an isomorphism is called a **splitting**. Pinsker's conjecture asserted that every ergodic automorphism splits into a factor of entropy zero and a \mathbf{K} -automorphism. Even after Ornstein found his counterexamples, the class of systems with this structure attracted considerable interest. Of particular importance is the subclass of systems that split into a zero-entropy factor and a Bernoulli factor.

In general, if

$$\pi : (X, \mu, T) \longrightarrow (Y, \nu, S)$$

is a factor map of automorphisms, then (X, μ, T) is **relatively Bernoulli over** π if there are a Bernoulli system \mathbf{B} and a factor map $\varphi : (X, \mu, T) \rightarrow \mathbf{B}$ such that the combined map $x \mapsto (\pi(x), \varphi(x))$ is a splitting. Since we restrict attention to standard probability spaces, it is equivalent to require that (i) π and φ are independent maps on the probability space (X, μ) and (ii) π and φ together generate \mathcal{B}_X modulo the μ -negligible sets.

The study of systems that split into zero-entropy and Bernoulli factors led Thouvenot to develop a ‘relative’ version of Ornstein’s conditions for Bernoullicity [104, 105]. Then, in [106], he proposed a replacement for the structure in Pinsker’s conjecture. According to Thouvenot, an automorphism has the **weak Pinsker property** if for every $\varepsilon > 0$ it has a splitting into a Bernoulli shift and a system of entropy less than ε . In the Introduction to [106], Thouvenot wrote,

‘The meaning of this definition is that the “structure” of these systems lies in factors of arbitrarily small entropy and that their randomness is essentially driven by a Bernoulli process.’

Although Pinsker’s conjecture and the weak Pinsker property are very close, it is worth being aware of the following important difference. If an ergodic automorphism has a splitting as in Pinsker’s conjecture, then the zero-entropy factor is essentially unique: it must generate the Pinsker σ -algebra of the original automorphism. However, there is nothing canonical about the splittings promised by the weak Pinsker property. This is illustrated very clearly by a result of Kalikow [35]. For a well-known example of a non-Bernoulli K -automorphism, he exhibits two splittings, each into a Bernoulli and a non-Bernoulli factor, such that the σ -algebras generated by the two non-Bernoulli factors have trivial intersection.

In this paper we prove that all ergodic automorphisms have the weak Pinsker property. Our proof actually gives a generalization of this fact, formulated relative to another fixed factor of the automorphism.

Theorem A. — *Let $\pi : (X, \mu, T) \rightarrow (Y, \nu, S)$ be a factor map of ergodic automorphisms. For every $\varepsilon > 0$, this map admits a factorization*

$$\begin{array}{ccc} (X, \mu, T) & \xrightarrow{\pi} & (Y, \nu, S) \\ & \searrow \pi_1 & \nearrow \pi_2 \\ & (\tilde{Y}, \tilde{\nu}, \tilde{S}) & \end{array}$$

in which (X, μ, T) is relatively Bernoulli over π_1 and $(\tilde{Y}, \tilde{\nu}, \tilde{S})$ has relative entropy less than ε over π_2 .

In fact, our method naturally yields an even more general result: the extension of Theorem A to free and ergodic measure-preserving actions of arbitrary countable amenable groups. This is Theorem 15.1 below. In our approach this requires essentially no more work than Theorem A, because even for Theorem A we appeal to some relative orbit equivalence theory which applies equally easily to any free amenable group action.

By a result of Fieldsteel [21], Theorem A also implies the weak Pinsker property for flows; we discuss this further in Subsection 16.3.

Measure concentration. — Within ergodic theory, the proof of Theorem A requires a careful use of Thouvenot’s relative Ornstein theory. Some of the elements we need have apparently not been published before, but they are quite natural extensions of well-known results. We state and prove these results carefully where we use them below, but they should not be thought of as really new.

The main innovations of this paper are results in discrete probability, and do not mention ergodic theory at all. They concern the important phenomenon of measure concentration. This phenomenon is already implicitly used in Ornstein’s original work. Its role in ergodic theory is made clearer by results of Marton and Shields [61], who call it the ‘blowing up property’. (Also, see Section 12 for some older references in Ornstein theory about the related notion of ‘extremality’.) Measure concentration has other important applications across probability theory [97, 99, 114], geometry [27, Chapter 3 $\frac{1}{2}$], and Banach space theory [66].

Let A be a finite alphabet. We study the Cartesian product spaces A^n for large n , endowed with the normalized Hamming metrics d_n (see Subsection 4.1). On the set of all probability measures $\text{Prob}(A^n)$, let \bar{d}_n be the transportation metric arising from d_n (see Subsection 4.2). We focus on a specific form of measure concentration expressed by a transportation inequality. Given $\mu \in \text{Prob}(A^n)$ and also parameters $\kappa > 0$ and $r > 0$, we say μ satisfies the inequality $T(\kappa, r)$ if the following holds for any other $\nu \in \text{Prob}(A^n)$:

$$\bar{d}_n(\nu, \mu) \leq \frac{1}{\kappa} D(\nu \parallel \mu) + r,$$

where D denotes Kullback–Leibler divergence. Such inequalities are introduced more carefully and generally in Section 5.

Inequalities of roughly this kind go back to works of Marton about product measures, starting with [57]. See also [100] for related results in Euclidean spaces. To see how $T(\kappa, r)$ implies concentration in the older sense of Lévy, consider a measurable subset $U \subseteq A^n$ with $\mu(U) \geq 1/2$, and apply $T(\kappa, r)$ to the measure $\nu := \mu(\cdot | U)$. The result is the bound

$$\bar{d}_n(\mu(\cdot | U), \mu) \leq \frac{1}{\kappa} D(\mu(\cdot | U) \parallel \mu) + r \leq \frac{\log 2}{\kappa} + r.$$

If κ is large and r is small, then it follows that $\overline{d}_n(\mu(\cdot | U), \mu)$ is small. This means that there is a coupling λ of $\mu(\cdot | U)$ and μ for which the integral

$$\int d_n(\mathbf{a}, \mathbf{a}') \lambda(d\mathbf{a}, d\mathbf{a}')$$

is small. By Markov's inequality, this λ must mostly be supported on pairs of strings $(\mathbf{a}, \mathbf{a}')$ that are close in d_n . Since the first marginal of λ is supported on U , and the second is the whole of μ , this is possible only if some small Hamming-neighbourhood of U has almost full measure according to μ .

Marton's original paper contains a more complete discussion of the relation between transportation inequalities and more traditional notions of measure concentration on metric spaces. Their basic connection to Ornstein theory is discussed in [61].

Our first new result about concentration asserts that an arbitrary probability measure μ on A^n can be represented as a mixture of other measures, most of which do satisfy a good transportation inequality, and with the number of summands controlled in terms of a quantity that indicates how much μ differs from being a product measure.

Theorem B. — For any $\varepsilon, r > 0$ there exists $c > 0$ such that the following holds. Let $\mu \in \text{Prob}(A^n)$, let $\mu_{\{i\}} \in \text{Prob}(A)$ be its marginals for $i = 1, 2, \dots, n$, and set

$$(1) \quad E := \sum_{i=1}^n H(\mu_{\{i\}}) - H(\mu).$$

Then μ may be written as a mixture

$$(2) \quad \mu = p_1 \mu_1 + \dots + p_m \mu_m$$

so that

- (a) $m \leq ce^{cE}$,
- (b) $p_1 < \varepsilon$, and
- (c) the measure μ_j satisfies $T(rn/1200, r)$ for every $j = 2, 3, \dots, m$.

The constant 1200 appearing here is presumably far from optimal, but I have not tried to improve it.

Beware that the inequality $T(\kappa, r)$ gets stronger as κ increases or r decreases, and so knowing $T(rn/1200, r)$ for one particular value of $r > 0$ does not imply it for any other value of r , larger or smaller.

The constant c provided by Theorem B depends on ε and r , but not on the alphabet A . In fact, if we allow c to depend on A , then Theorem B is easily subsumed by simpler results. To see this we must consider two cases. On the one hand, Marton's concentration result for product measures (see Proposition 5.2 below) gives a small constant

α depending on r such that, if $E < \alpha n$, then μ itself is close in \overline{d}_n to the product of its marginals. Using this, one can find a subset $U \subseteq A^n$ with $\mu(U)$ close to 1 and such that $\mu(\cdot | U)$ is highly concentrated: see Proposition 5.9 below for a precise statement. So in case $E < \alpha n$, we can obtain Theorem B with the small error term plus only one highly concentrated term. On the other hand, if $E \geq \alpha n$, then the partition of A^n into *singletons* has only

$$|A|^n = e^{\log |A| n} \leq e^{\frac{\log |A|}{\alpha} E}$$

parts. This bound has the form of (a) above if we allow c to be $\log |A|/\alpha$, and point masses are certainly highly concentrated. In view of these two arguments, Theorem B is useful only when $|A|$ is large in terms of ε , r and E . That turns out to be the crucial case for our application to ergodic theory.

Once Theorem B is proved, we also prove the following variant.

Theorem C. — *For any $\varepsilon, r > 0$ there exist $c > 0$ and $\kappa > 0$ such that, for any alphabet A , the following holds for all sufficiently large n . Let $\mu \in \text{Prob}(A^n)$, and let E be as in (1). Then there is a partition*

$$A^n = U_1 \cup \dots \cup U_m$$

such that

- (a) $m \leq ce^{cE}$,
- (b) $\mu(U_1) < \varepsilon$, and
- (c) $\mu(U_j) > 0$ and the conditioned measure $\mu(\cdot | U_j)$ satisfies $T(\kappa n, r)$ for every $j = 2, 3, \dots, m$.

Qualitatively, this is stronger than Theorem B: one recovers the qualitative result of Theorem B by letting $p_j = \mu(U_j)$ and $\mu_j = \mu(\cdot | U_j)$ for each j . As in Theorem B, the proof gives a value for κ which is a small constant multiple of r , but this time we have not recorded the precise constant in the statement. Also, the statement of Theorem C requires that n be sufficiently large in terms of ε , r and A , but I do not know whether the dependence on A is really needed.

Theorem C is proved using Theorem B in Section 8. The key idea is a random partitioning procedure which can be used to turn the mixture from Theorem B into a suitable partition.

Theorem C is the key to our proof of Theorem A. In that proof, Theorem C is applied repeatedly for smaller and smaller values of r (and with different choices of A and n each time), and the resulting partitions are then combined into a construction of the new factor π_1 of Theorem A. In fact, for this application it suffices to know a special case of Theorem C which has a shorter proof, because in the ergodic theoretic setting some relevant extra hypotheses can be obtained from the Shannon–McMillan theorem.

However, we choose to prove and then cite the full strength of Theorem C. We do so both for the sake of greater generality and because it reduces the complexity of the proof of Theorem A.

In Theorems B and C, the starting hypothesis is a bound on the difference of entropies in (1). The use of this particular quantity underlies the application to proving Theorem A, because we are able to control this quantity in the ergodic theoretic setting. This control is a simple consequence of the fact that, in a stationary stochastic process, the normalized Shannon entropies of long finite strings converge to the limiting entropy rate of the process. The resulting link from Theorem C to Theorem A is explained in Step 2 of the construction in Subsection 14.1, and specifically by Lemma 14.4. That specific explanation can largely be appreciated before reading the rest of the paper, although the rest of Subsection 14.1 requires more preparation.

The proof of Theorem B is not constructive. It gives no reason to expect that the summands in the resulting mixture are unique, even approximately. The same applies to the partition produced by Theorem C. This issue may reflect the fact that the splittings promised by the weak Pinsker property are not at all canonical. However, I do not know of concrete examples of measures μ on A^n for which one can prove that several decompositions as in Theorem B are possible, none of them in any way ‘better’ than the others. It would be worthwhile to understand this issue more completely.

For our application of these decomposition results to ergodic theory it suffices to let $\varepsilon = r$. However, the statements and proofs of these results are clearer if the roles of these two tolerances are kept separate.

Structure and randomness. — Theorems B and C bear a resemblance to some key structural results in extremal and additive combinatorics. They may be regarded as ‘rough classifications’ of measures on A^n , in the sense proposed by Gowers in [25]. Although Gowers’ principal examples come from those branches of combinatorics, he suggests that rough classifications should be useful across large parts of mathematics: see his Subsection 3.8. Theorems B and C are examples in a new setting.

Let us discuss this further for Theorem B. Theorem B expresses an arbitrary measure on A^n as a mixture of other measures, not too many in number, and most of them highly concentrated. We may fit this description into the following terms:

- A kind of ‘structure’ which has controlled ‘complexity’. A suitable indication of this ‘complexity’ is the number of summands in the mixture (2).
- A kind of ‘randomness’ relative to that structure. In this case measure concentration plays the part of randomness: most of the weight in (2) is on measures μ_i that exhibit strong concentration. It is natural to regard concentration as a manifestation of ‘randomness’ because product measures — the ‘most random’ of all — are very highly concentrated.

To a first approximation, we can identify these notions of ‘structure’ and ‘randomness’ with the terms that Thouvenot used to describe the meaning of the weak Pinsker

property in the sentence from [106] cited above. Put very roughly, the small number of summands in (2) turns into the smallness of the relative entropy $h(\tilde{\nu}, \tilde{S} | \pi_2)$ in Theorem A. (This description ignores several extra constructions and arguments that are needed before the proof of Theorem A is complete, but they are not of the same order of magnitude as Theorem B itself.)

A few years after Gowers’ paper, Tao’s ICM address [102] emphasized that a dichotomy between ‘structure’ and ‘randomness’ can be found in all the different proofs of Szemerédi’s theorem (one of the highlights of additive combinatorics). Tao describes how this dichotomy reappears in related work in ergodic theory and harmonic analysis. The flavour of ergodic theory in Tao’s paper is different from ours: he refers to the theory that originates in Furstenberg’s multiple recurrence theorem [22], which makes no mention of Kolmogorov–Sinai entropy. However, the carrier of ergodic theoretic ‘structure’ is the same in both settings: a special factor (distal in Furstenberg’s work, low-entropy in ours) relative to which the original automorphism exhibits some kind of ‘randomness’ (relative weak mixing, respectively relative Bernoullicity). Thus, independently and after an interval of thirty years, Tao’s identification of this useful dichotomy aligns with Thouvenot’s description of the weak Pinsker property cited above.

The parallel with parts of extremal and additive combinatorics resurfaces inside the proof of Theorem B. In the most substantial step of that proof, the representation of μ as a mixture is obtained by a ‘decrement’ argument in a certain quantity that describes how far μ is from being a product measure. (The ‘decrement’ argument does not quite produce the mixture in (2)—some secondary processing is still required.) This proof is inspired by the ‘energy increment’ proof Szemerédi’s regularity lemma and various similar proofs in additive combinatorics. Indeed, one of the main innovations of the present paper is finding the right quantity in which to seek a suitable decrement. It turns out to be the ‘dual total correlation’, a classical measure of multi-variate mutual information (see Section 3). We revisit the comparison with Szemerédi’s regularity lemma at the end of Subsection 6.1.

Outline of the rest of the paper. — Part I (Sections 2–9) begins with background on Shannon entropy and related quantities, metric probability spaces, and measure concentration. Then the bulk of this part is given to the proofs of Theorems B and C. The final section in this part, Section 9, introduces ‘extremality’, a different notion of concentration which is more convenient for applications to ergodic theory. However, none of Part I itself involves any ergodic theory, and this part may be relevant to other applications in metric geometry or probability. Section 9 is logically independent from the proofs of Theorems B and C.

Part II (Sections 10–13) gives a partial account of Thouvenot’s relative Ornstein theory, tailored to the specific needs of the proof of Theorem A. This part makes no mention of Theorems B or C. Its only reliance on Part I is for some of the general machinery in Section 9.

Finally, Part III (Sections 14–16) completes the proof of Theorem A and its generalization to countable amenable groups, and collects some consequences and remaining open questions. In this part we draw together the threads from Parts I and II. In particular, the relevance of Theorem C to proving Theorem A is explained in Section 14.

Part I

MEASURES ON FINITE CARTESIAN PRODUCTS

2. Background on kernels, entropy, and mutual information

We write $\text{Prob}(X)$ for the set of all probability measures on a standard measurable space X . We mostly use this notation when X is a finite set. In that case, given $\mu \in \text{Prob}(X)$ and $x \in X$, we often write $\mu(x)$ instead of $\mu(\{x\})$.

2.1. Kernels, joint distributions, and fuzzy partitions. — Let Ω and X be two measurable spaces with X standard. A **kernel** from Ω to X is a measurable function $\omega \mapsto \mu_\omega$ from Ω to the set $\text{Prob}(X)$, where the latter is given its usual measurable structure (generated, for instance, by the vague topology resulting from any choice of compact metric on X which generates \mathcal{B}_X). We often denote such a kernel by μ_\bullet . Kernels arise as regular conditional distributions in probability theory; we meet many such cases below.

Suppose in addition that P is a probability measure on Ω . Then P and μ_\bullet together define a probability measure on $\Omega \times X$ with its product σ -algebra, given formally by the integral

$$\int_{\Omega} \delta_\omega \times \mu_\omega P(d\omega).$$

See, for instance, [17, Theorem 10.2.1] for the rigorous construction. We denote this new measure by $P \times \mu_\bullet$. Following some information theorists, we refer to it as the **hookup** of P and μ_\bullet . It is a coupling of P with the averaged measure

$$(3) \quad \mu = \int_{\Omega} \mu_\omega P(d\omega).$$

Such an average of the measures in a kernel is called a **mixture**. Often we start with the measure μ and then consider ways to express it as a mixture. If Ω and μ are non-trivial then there are many ways to do this. Once an expression as in (3) has been chosen, we may also consider a pair of random variables (ζ, ξ) whose joint distribution is $P \times \mu_\bullet$. Then ξ by itself has distribution μ . Loosely motivated by terminology in [20, Section II.5], we call such a pair of random variables a **randomization** of the mixture.

Now consider a standard probability space (X, μ) and a measurable function $\rho : X \rightarrow [0, \infty)$ satisfying

$$0 < \int \rho \, d\mu < \infty.$$

Then we define

$$\mu_{|\rho} := \frac{\rho}{\int \rho \, d\mu} \cdot \mu.$$

Similarly, if $U \subseteq X$ has positive measure, then we often write

$$\mu_{|U} := \mu(\cdot | U).$$

A **fuzzy partition** on X is a finite tuple (ρ_1, \dots, ρ_k) of measurable functions from X to $[0, 1]$ satisfying

$$(4) \quad \rho_1 + \rho_2 + \dots + \rho_k = 1.$$

This is more conventionally called a ‘partition of unity’, but in the present setting I feel ‘fuzzy partition’ is less easily confused with partitions that consist of sets. If P_1, \dots, P_k are a partition of X into measurable subsets, then they give rise to the fuzzy partition $(1_{P_1}, \dots, 1_{P_k})$, which has the special property of consisting of indicator functions.

A fuzzy partition (ρ_1, \dots, ρ_k) gives rise to a decomposition of μ into other measures:

$$\mu = \rho_1 \cdot \mu + \dots + \rho_k \cdot \mu.$$

We may also write this as a mixture

$$(5) \quad p_1 \cdot \mu_1 + \dots + p_k \cdot \mu_k$$

by letting $p_i := \int \rho_i \, d\mu$ and $\mu_j := \mu_{|\rho_j}$, where terms for which $p_i = 0$ are interpreted as zero.

The stochastic vector $p := (p_1, p_2, \dots, p_k)$ may be regarded as a probability distribution on $\{1, 2, \dots, k\}$, and the family $(\mu_j)_{j=1}^k$ may be regarded as a kernel from $\{1, 2, \dots, k\}$ to X . We may therefore construct the hookup $p \times \mu_\bullet$. This hookup has a simple formula in terms of the fuzzy partition:

$$(p \times \mu_\bullet)(\{j\} \times A) = \int_A \rho_j \, d\mu \quad \text{for } 1 \leq j \leq k \text{ and measurable } A \subseteq X.$$

One can also reverse the roles of X and $\{1, 2, \dots, k\}$ here, and regard $p \times \mu_\bullet$ as the hookup of μ to the kernel

$$x \mapsto (\rho_1(x), \dots, \rho_k(x))$$

from X to $\text{Prob}(\{1, 2, \dots, k\})$: the fact that this *is* a kernel is precisely (4).

2.2. Entropy and mutual information. — If ξ and ζ are finite-valued random variables on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$, then $\mathbf{H}(\xi)$, $\mathbf{H}(\xi | \zeta)$ denote Shannon entropy and conditional Shannon entropy, respectively. We also use the notation $\mathbf{H}(\mu)$ for the Shannon entropy of a probability distribution μ on a finite set. We assume familiarity with the basic properties of entropy, particularly the chain rule. A standard exposition is [11, Chapter 2].

More generally, one can define the conditional entropy $\mathbf{H}(\xi | \mathcal{G})$ whenever ξ is a finite-valued random variable and \mathcal{G} is a σ -subalgebra of \mathcal{F} . To write it in terms of unconditional entropy, let $\omega \mapsto \mu_\omega$ be a conditional distribution for ξ given \mathcal{G} . Then we have

$$(6) \quad \mathbf{H}(\xi | \mathcal{G}) = \int \mathbf{H}(\mu_\omega) \mathbf{P}(d\omega).$$

See, for instance, [5, Section 12].

If $F \in \mathcal{F}$ has $\mathbf{P}(F) > 0$, then we sometimes use the notation

$$(7) \quad \mathbf{H}(\xi | F) := \mathbf{H}_{\mathbf{P}(\cdot | F)}(\xi) \quad \text{and} \quad \mathbf{H}(\xi | \zeta; F) := \mathbf{H}_{\mathbf{P}(\cdot | F)}(\xi | \zeta).$$

For a fixed alphabet A , Shannon's entropy function \mathbf{H} is concave on $\text{Prob}(A)$. The next inequality gives a useful way to reverse this concavity, up to an additional error term.

Lemma 2.1. — *Consider a finite mixture*

$$\mu = p_1 \mu_1 + \cdots + p_k \mu_k$$

of probability measures on a finite set A . Then

$$\mathbf{H}(\mu) \leq \sum_{j=1}^k p_j \mathbf{H}(\mu_j) + \mathbf{H}(p_1, \dots, p_k).$$

Proof. — Let (ζ, ξ) be a randomization of the mixture, so this random pair takes values in $\{1, 2, \dots, k\} \times A$. Then ζ and ξ have marginal distributions (p_1, \dots, p_k) and μ , respectively. Now monotonicity of \mathbf{H} and the chain rule give

$$\begin{aligned} \mathbf{H}(\mu) &= \mathbf{H}(\xi) \leq \mathbf{H}(\zeta, \xi) = \mathbf{H}(\xi | \zeta) + \mathbf{H}(\zeta) \\ &= \sum_{j=1}^k p_j \mathbf{H}(\mu_j) + \mathbf{H}(p_1, \dots, p_k). \end{aligned}$$

□

In addition to entropy and conditional entropy, we make essential use of two other related quantities: mutual information, and Kullback–Leibler ('KL') divergence. These are also standard in information theory, but appear less often in ergodic theory. For finite-valued random variables they are introduced alongside entropy in [11, Chapter 2].

If ξ and ζ are finite-valued, then their mutual information is

$$I(\xi ; \zeta) := H(\xi) - H(\xi | \zeta).$$

An exercise in the chain rule gives the alternative formula

$$I(\xi ; \zeta) = H(\xi) + H(\zeta) - H(\xi, \zeta).$$

From this it follows that $I(\xi ; \zeta)$ is symmetric in ξ and ζ . We define a conditional mutual information such as $I(\xi ; \zeta | F)$ or $I(\xi ; \zeta | \alpha)$ by conditioning all the entropies in the formulae above on the set F or random variable α . Mutual information and its conditional version satisfy their own chain rule, similar to that for entropy: see, for instance, [11, Theorem 2.5.2].

KL divergence is a way of comparing two probability measures, say μ and ν , on the same measurable space X . The KL divergence of ν with respect to μ is $+\infty$ unless ν is absolutely continuous with respect to μ . In that case it is given by

$$(8) \quad D(\nu \| \mu) := \int \frac{d\nu}{d\mu} \log \frac{d\nu}{d\mu} d\mu = \int \log \frac{d\nu}{d\mu} d\nu.$$

This may still be $+\infty$, since $\frac{d\nu}{d\mu} \log \frac{d\nu}{d\mu}$ need not lie in $L^1(\mu)$. Since the function $t \mapsto t \log t$ is strictly convex, Jensen's inequality gives that $D(\nu \| \mu) \geq 0$, with equality if and only if $d\nu/d\mu = 1$, hence if and only if $\nu = \mu$.

KL divergence appears in various elementary but valuable formulae for conditional entropy and mutual information. Let the random variables ζ and ξ take values in the finite sets A and B . Let $\mu \in \text{Prob}(A)$, $\nu \in \text{Prob}(B)$, and $\lambda \in \text{Prob}(A \times B)$ be the distribution of ζ , of ξ , and of the pair (ζ, ξ) respectively. Then a standard calculation gives

$$(9) \quad I(\xi ; \zeta) = D(\lambda \| \mu \times \nu).$$

The next calculation is also routine, but we include a proof for completeness.

Lemma 2.2. — *Let (Ω, \mathcal{F}, P) be a probability space and let X be a finite set. Let $\xi : \Omega \rightarrow X$ be measurable and let μ be its distribution. Let $\mathcal{G} \subseteq \mathcal{H}$ be σ -subalgebras of \mathcal{F} , and let μ_\bullet and ν_\bullet be conditional distributions of ξ given \mathcal{G} and \mathcal{H} , respectively. Then*

$$H(\xi | \mathcal{G}) - H(\xi | \mathcal{H}) = \int D(\nu_\omega \| \mu_\omega) P(d\omega).$$

In particular, applying this with \mathcal{G} trivial, we have

$$H(\xi) - H(\xi | \mathcal{H}) = \int D(\nu_\omega \| \mu) P(d\omega).$$

Proof. — Since \mathbf{X} is finite, the definition of KL divergence gives

$$\begin{aligned} \int \mathbf{D}(\nu_\omega \parallel \mu_\omega) \mathbf{P}(d\omega) &= \int \sum_{x \in \mathbf{X}} \nu_\omega(x) \log \frac{\nu_\omega(x)}{\mu_\omega(x)} \mathbf{P}(d\omega) \\ &= \int \sum_{x \in \mathbf{X}} \nu_\omega(x) \log \nu_\omega(x) \mathbf{P}(d\omega) - \int \sum_{x \in \mathbf{X}} \nu_\omega(x) \log \mu_\omega(x) \mathbf{P}(d\omega). \end{aligned}$$

For each $x \in \mathbf{X}$ we have

$$\nu_\omega(x) = \mathbf{P}(\{\xi = x\} \mid \mathcal{H})(\omega) \quad \text{and} \quad \mu_\omega(x) = \mathbf{P}(\{\xi = x\} \mid \mathcal{G})(\omega) \quad \mathbf{P}\text{-a.s.}$$

Therefore, by the tower property of conditional expectation, for each fixed $x \in \mathbf{X}$ the function $\mu_\omega(x)$ is a conditional expectation of $\nu_\omega(x)$ onto \mathcal{G} . Using this in the second integral above, that formula becomes

$$\int \sum_{x \in \mathbf{X}} \nu_\omega(x) \log \nu_\omega(x) \mathbf{P}(d\omega) - \int \sum_{x \in \mathbf{X}} \mu_\omega(x) \log \mu_\omega(x) \mathbf{P}(d\omega).$$

This equals $[-\mathbf{H}(\xi \mid \mathcal{H}) + \mathbf{H}(\xi \mid \mathcal{G})]$, by (6). \square

Now consider a fuzzy partition (ρ_1, \dots, ρ_k) on a finite set \mathbf{X} and the associated mixture (5). Let (ζ, ξ) be a randomization of this mixture. We define

$$(10) \quad \mathbf{I}_\mu(\rho_1, \dots, \rho_k) := \mathbf{I}(\xi; \zeta).$$

This notion of mutual information for a measure and a fuzzy partition plays a central role later in Part I. Using Lemma 2.2 it may be evaluated as follows.

Corollary 2.3. — *With μ and (ρ_1, \dots, ρ_k) as above, we have*

$$\mathbf{I}_\mu(\rho_1, \dots, \rho_k) = \sum_{j=1}^k p_j \cdot \mathbf{D}(\mu_{|\rho_j} \parallel \mu), \quad \text{where } p_j := \int \rho_j d\mu.$$

Proof. — Let (ζ, ξ) be the randomization above, and apply the second formula of Lemma 2.2 with \mathcal{H} the σ -algebra generated by ζ . The result follows because $\mu_{|\rho_j}$ is the conditional distribution of ξ given the event $\{\zeta = j\}$, and because

$$\mathbf{I}_\mu(\rho_1, \dots, \rho_k) = \mathbf{I}(\xi; \zeta) = \mathbf{H}(\xi) - \mathbf{H}(\xi \mid \zeta).$$

\square

It is also useful to have a version of the chain rule for mutual information in terms of fuzzy partitions. To explain this, let us consider two fuzzy partitions $(\rho_i)_{i=1}^k$ and $(\rho'_j)_{j=1}^\ell$ on a finite set \mathbf{X} , and assume there is a partition

$$\{1, 2, \dots, \ell\} = \mathbf{J}_1 \cup \mathbf{J}_2 \cup \dots \cup \mathbf{J}_k$$

into nonempty subsets such that

$$(11) \quad \rho_i = \sum_{j \in J_i} \rho'_j \quad \text{for } i = 1, 2, \dots, k.$$

This is a natural generalization of the notion of refinement for partitions consisting of sets. However, beware that for a general pair of fuzzy partitions there may be several ways to partition $\{1, 2, \dots, \ell\}$ into k subsets so that (11) holds.

From (11), it follows that the tuple of ratios

$$(12) \quad \left(\frac{\rho'_j}{\rho_i} \right)_{j \in J_i}$$

define a new fuzzy partition for each $1 \leq i \leq k$. To be precise, these ratios are defined only on the set $\{\rho_i > 0\}$. Outside that set, each ρ'_j for $j \in J_i$ is also zero, by (11). We extend the new fuzzy partition (12) arbitrarily to the set $\{\rho_i = 0\}$.

Let (ζ, ξ) be a randomization of the mixture obtained from μ and $(\rho'_j)_{j=1}^\ell$. This means that (ζ, ξ) are a random pair in $\{1, 2, \dots, \ell\} \times \mathbf{X}$, and their joint distribution is given by a hookup of the kind constructed at the end of the previous subsection. Define a third random variable $\bar{\zeta}$ taking values in $\{1, 2, \dots, k\}$ by setting

$$\bar{\zeta} = i \quad \text{when } \zeta \in J_i.$$

The pair $(\bar{\zeta}, \xi)$ may be regarded as a ‘coarsening’ of (ζ, ξ) in the first entry. From the relation (11) it follows that the pair $(\bar{\zeta}, \xi)$ are a randomization of the mixture obtained from μ and $(\rho_i)_{i=1}^k$. Similarly, if we condition on the event $\{\bar{\zeta} = i\}$, then (ζ, ξ) become a randomization of the mixture obtained from $\mu|_{\rho_i}$ and the new fuzzy partition (12). Since the set $\{\rho_i = 0\}$ is negligible according to $\mu|_{\rho_i}$, our choice of extension of (12) to this set is unimportant.

Now we can simply write out the chain rule

$$I(\xi; \zeta) = I(\xi; \bar{\zeta}) + I(\xi; \zeta | \bar{\zeta})$$

(see [11, Theorem 2.5.2]) in terms of all these fuzzy partitions. The result is

$$(13) \quad I_\mu(\rho'_1, \dots, \rho'_\ell) = I_\mu(\rho_1, \dots, \rho_k) + \sum_{i=1}^k \left(\int \rho_i d\mu \right) \cdot I_{\mu|_{\rho_i}} \left(\left(\frac{\rho'_j}{\rho_i} \right)_{j \in J_i} \right).$$

An immediate corollary is worth noting in itself: whenever there is a partition J_1, \dots, J_k of $\{1, 2, \dots, \ell\}$ for which (11) holds, we must have

$$(14) \quad I_\mu(\rho'_1, \dots, \rho'_\ell) \geq I_\mu(\rho_1, \dots, \rho_k).$$

Remark. — KL divergence is defined for pairs of measures on general measurable spaces. As a result, we may use (9) to extend the definition of mutual information to general pairs of random variables, not necessarily discrete. The basic properties of mutual information still hold, with proofs that require only small adjustments. Mutual information is already defined and studied this way in the original reference [50]. See [80, Section 2.1] for a comparison with other definitions at this level of generality.

The main theorems of the present paper do not need this extra generality, but KL divergence is central to the background ideas in Section 5 below, where the more general setting is appropriate.

2.3. *A classical variational principle.* — Let (\mathbf{X}, μ) be a probability space and let $f \in L^\infty(\mu)$. Let $\langle \cdot \rangle$ denote integration with respect to μ , and let

$$M_\mu(f) := \langle e^f \rangle \quad \text{and} \quad C_\mu(f) := \log M_\mu(f).$$

Regarding f as a random variable, the functions of $t \in \mathbf{R}$ given by $M_\mu(tf)$ and $C_\mu(tf)$ are its moment generating function and cumulant generating function, respectively.

Definition 2.4. — For any μ and f as above, the associated **Gibbs measure** is the probability measure

$$\mu_{|e^f} = \frac{e^f}{\langle e^f \rangle} \cdot \mu.$$

Note that the normalizing constant in this definition is equal to $M_\mu(f)$.

The importance of Gibbs measures derives from the following classical variational principle. It is essentially Gibbs' own principle from his work on the foundations of statistical physics, phrased in modern, abstract terms. Rigorous mathematical treatments go back at least to Kullback's work [48] in information theory and Sanov's [88] on large deviations. We include the short proof for completeness, since we need the result at one point in Subsection 5.3 below. Closely related results appear in [11, Theorem 12.1.1 and Problem 12.2].

Lemma 2.5. — For μ and f as above, the function of $\nu \in \text{Prob}(\mathbf{X})$ given by the expression

$$D(\nu \parallel \mu) - \int f \, d\nu$$

achieves its unique minimum at $\nu = \mu_{|e^f}$, where it is equal to $-C_\mu(f)$.

Proof. — If ν is not absolutely continuous with respect to μ , then $D(\nu \parallel \mu) = +\infty$, and therefore ν is not a candidate for minimizing the expression in question. So suppose that ν is absolutely continuous with respect to μ .

Since μ and $\mu_{|e^f}$ are mutually absolutely continuous, ν is also absolutely continuous with respect to $\mu_{|e^f}$. Let ρ be the Radon–Nikodym derivative $d\nu/d\mu_{|e^f}$. Then

$$\frac{d\nu}{d\mu} = \rho \cdot \frac{e^f}{\langle e^f \rangle},$$

where $\langle \cdot \rangle$ denotes integration with respect to μ , as before. Substituting this into the definition of $D(\nu \parallel \mu)$, we obtain

$$\begin{aligned} D(\nu \parallel \mu) &= \int \log \left(\rho \frac{e^f}{\langle e^f \rangle} \right) d\nu \\ &= \int \log \rho d\nu + \int f d\nu - \log \langle e^f \rangle \\ &= D(\nu \parallel \mu_{|e^f}) + \int f d\nu - C_\mu(f). \end{aligned}$$

Re-arranging, this becomes

$$D(\nu \parallel \mu) - \int f d\nu = D(\nu \parallel \mu_{|e^f}) - C_\mu(f).$$

The term $D(\nu \parallel \mu_{|e^f})$ is non-negative, and zero if and only if $\nu = \mu_{|e^f}$. □

3. Total correlation and dual total correlation

Let $\xi = (\xi_1, \dots, \xi_n)$ be a tuple of finite-valued random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. If $n = 2$ then the mutual information $I(\xi_1; \xi_2)$ provides a canonical way to quantify the dependence between them. But this has no single generalization to the cases $n > 2$. Rather, a range of options are available, suitable for different purposes. Early studies of these options, and the many relations between them, include [64], [115] and [31]. An idea of the breadth of these options can be obtained from [12] and the references given there.

Two different quantities of this kind play crucial roles in this paper.

3.1. Total correlation. — A simple and natural choice for ‘multi-variate mutual information’ is the difference

$$\sum_{i=1}^n H(\xi_i) - H(\xi).$$

This is always non-negative by the subadditivity of Shannon entropy. It goes back at least to Watanabe’s paper [115], so following him we call it the **total correlation**. It agrees

with mutual information if $n = 2$. It depends only on the joint distribution μ of the tuple ξ , so we may denote it by either $\text{TC}(\xi)$ or $\text{TC}(\mu)$. This leaves the role of the n separate coordinates to the reader's understanding.

If A is a finite set and μ is a probability measure on A^n , then $\text{TC}(\mu)$ is precisely the quantity E that appears in Theorems B and C. It is the key feature of such a finite-dimensional distribution that we are able to control when we apply Theorem C during the proof of Theorem A. This control is exerted in Subsection 14.1.

The total correlation of μ appears in various classical estimates of probability theory and statistics. For instance, it is central to Csiszár's approach to conditional limit theorems for product measures [14], where it is the quantity that must be bounded to prove convergence in information. In these applications it is often written in the alternative form

$$(15) \quad \text{TC}(\mu) = D(\mu \parallel \mu_{\{1\}} \times \cdots \times \mu_{\{n\}}),$$

where $\mu_{\{i\}}$ is the i^{th} marginal of μ on A . This equation for $\text{TC}(\mu)$ generalizes (9). The generalization is easily proved by induction on n , applying (9) at each step.

Since total correlation is given by a difference of entropy values, it is neither convex nor concave as a function of μ . However, we do have the following useful approximate concavity.

Lemma 3.1. — *Consider a finite mixture*

$$\mu = p_1\mu_1 + p_2\mu_2 + \cdots + p_k\mu_k$$

of probability measures on A^n . Then

$$(16) \quad \sum_{j=1}^k p_j \text{TC}(\mu_j) \leq \text{TC}(\mu) + H(p_1, \dots, p_k).$$

Proof. — Let $\xi_i : A^n \rightarrow A$ be the i^{th} coordinate projection for $1 \leq i \leq n$. On the one hand, the concavity of H gives

$$\sum_{j=1}^k p_j H_{\mu_j}(\xi_i) \leq H_{\mu}(\xi_i) \quad \text{for each } i = 1, 2, \dots, n.$$

On the other hand, Lemma 2.1 gives

$$\sum_{j=1}^k p_j H(\mu_j) \geq H(\mu) - H(p_1, \dots, p_k).$$

The desired result is a linear combination of these inequalities. □

Corollary 3.2. — If μ is a probability measure on A^n , $U \subseteq A^n$, and $\mu(U) > 0$, then

$$\mathrm{TC}(\mu|_U) \leq \frac{1}{\mu(U)}(\mathrm{TC}(\mu) + \log 2).$$

Proof. — Simply apply Lemma 3.1 to the decomposition

$$\mu = \mu(U) \cdot \mu|_U + \mu(A^n \setminus U) \cdot \mu|_{A^n \setminus U},$$

then drop the term involving $\mu|_{A^n \setminus U}$ and recall that any two-set partition has entropy at most $\log 2$. \square

A simple example shows that one cannot hope for any useful inequality that reverses (16).

Example 3.3. — For $p \in [0, 1]$, let v_p be the p -biased distribution on $\{0, 1\}$. Now let

$$\mu := \frac{1}{2}(v_p^{\times n} + v_q^{\times n})$$

for some distinct p and q . Then $\mu_{\{i\}} = v_{(p+q)/2}$ for each $1 \leq i \leq n$, and so

$$\mathrm{TC}(\mu) = n\mathrm{H}\left(\frac{p+q}{2}, 1 - \frac{p+q}{2}\right) - \mathrm{H}\left(\frac{v_p^{\times n} + v_q^{\times n}}{2}\right).$$

By Lemma 2.1, we have

$$(17) \quad \mathrm{H}\left(\frac{v_p^{\times n} + v_q^{\times n}}{2}\right) \leq \log 2 + \frac{\mathrm{H}(p, 1-p)n + \mathrm{H}(q, 1-q)n}{2}.$$

Therefore

$$\mathrm{TC}(\mu) \geq \left[\mathrm{H}\left(\frac{p+q}{2}, 1 - \frac{p+q}{2}\right) - \frac{\mathrm{H}(p, 1-p) + \mathrm{H}(q, 1-q)}{2} \right] n - \log 2.$$

This grows linearly with n , since p and q are distinct and H is strictly concave. But μ is an average of just two product measures, each of which has total correlation zero.

3.2. Dual total correlation. — In this subsection we start to use the notation

$$[n] := \{1, 2, \dots, n\}.$$

It reappears often later in the paper.

Our second multi-variate generalization of mutual information is the quantity

$$(18) \quad \mathrm{H}(\xi) - \sum_{i=1}^n \mathrm{H}(\xi_i | \xi_{[n] \setminus i}),$$

where

$$\xi_{[n]\setminus i} := (\xi_1, \dots, \xi_{i-1}, \xi_{i+1}, \dots, \xi_n).$$

When $n = 2$, this again agrees with the mutual information $I(\xi_1 ; \xi_2)$. This generalization seems to originate in Han's paper [31], where it is called the **dual total correlation**. Like total correlation, it depends only on the joint distribution μ of ξ . We denote it by either $\text{DTC}(\xi)$ or $\text{DTC}(\mu)$.

By applying the chain rule to the first term in (18) we obtain

$$\begin{aligned} \text{DTC}(\xi) &= \sum_{i=1}^n H(\xi_i | \xi_1, \dots, \xi_{i-1}) - \sum_{i=1}^n H(\xi_i | \xi_{[n]\setminus i}) \\ &= \sum_{i=1}^n [H(\xi_i | \xi_1, \dots, \xi_{i-1}) - H(\xi_i | \xi_{[n]\setminus i})]. \end{aligned}$$

All these terms are non-negative, since Shannon entropy cannot increase under extra conditioning, so $\text{DTC} \geq 0$.

The ergodic theoretic setting produces measures with control on their TC. The relevance of DTC to our paper is less obvious. However, we find below that $\text{DTC}(\mu)$ is more easily related than $\text{TC}(\mu)$ to concentration properties of μ . As a result, we first prove an analog of Theorem B with $\text{DTC}(\mu)$ in place of $\text{TC}(\mu)$: Theorem 7.1 below. We then convert this into Theorem B itself by showing that a bound on $\text{TC}(\mu)$ implies a related bound on $\text{DTC}(\mu')$, where μ' is a projection of μ onto a slightly lower-dimensional product of copies of A : see Lemma 3.7 below. The first of these two stages is the more subtle. It is based on a 'decrement' argument for dual total correlation, in which we show that if μ fails to be concentrated then it may be written as a mixture of two other measures whose dual total correlations are strictly smaller on average. This special property of DTC is the key to the whole proof, and it seems to have no analog for TC. This is why we need the two separate stages described above.

Now suppose that ζ is another finite-valued random variable on the same probability space as the tuple ξ . Then the **conditional dual total correlation** of the tuple ξ given ζ is obtained by conditioning all the Shannon entropies that appear in (18):

$$\text{DTC}(\xi | \zeta) := H(\xi | \zeta) - \sum_{i=1}^n H(\xi_i | \xi_{[n]\setminus i}, \zeta).$$

Lemma 3.4. — *In the setting above, we have*

$$(19) \quad \text{DTC}(\xi) - \text{DTC}(\xi | \zeta) = I(\xi ; \zeta) - \sum_{i=1}^n I(\xi_i ; \zeta | \xi_{[n]\setminus i}).$$

Proof. — This follows from these standard identities:

$$H(\xi) - H(\xi | \zeta) = I(\xi; \zeta)$$

and

$$H(\xi_i | \xi_{[n]\setminus i}) - H(\xi_i | \xi_{[n]\setminus i}, \zeta) = I(\xi_i; \zeta | \xi_{[n]\setminus i}). \quad \square$$

We sometimes refer to the right-hand side of (19) as the **DTC decrement** associated to ξ and ζ . This quantity can actually take either sign, but it is positive in the cases that we need later.

Remark. — For each i , another appeal to the chain rule gives

$$H(\xi) = H(\xi_{[n]\setminus i}) + H(\xi_i | \xi_{[n]\setminus i}).$$

Using this to substitute for $H(\xi_i | \xi_{[n]\setminus i})$ in $\text{DTC}(\xi)$ and simplifying, we obtain

$$\text{DTC}(\xi) = \sum_{i=1}^n H(\xi_{[n]\setminus i}) - (n-1)H(\xi).$$

In this form, the non-negativity of $\text{DTC}(\xi)$ is a special case of some classic inequalities proved by Han in [32] (see also [11, Section 17.6]).

Han's inequalities were part of an abstract study of non-negativity among various multivariate generalizations of mutual information. Shearer proved an even more general inequality at about the same time with a view towards combinatorial applications, which then appeared in [9]. For any subset $S = \{i_1, \dots, i_\ell\} \subseteq [n]$, let us write $\xi_S = (\xi_{i_1}, \dots, \xi_{i_\ell})$. If \mathcal{S} is a family of subsets of $[n]$ with the property that every element of $[n]$ lies in at least k members of \mathcal{S} , then Shearer's inequality asserts that

$$(20) \quad \sum_{S \in \mathcal{S}} H(\xi_S) \geq kH(\xi).$$

A simple proof is given in the survey [81], together with several applications to combinatorics.

When \mathcal{S} consists of all subsets of $[n]$ of size $n-1$, the inequality (20) is just the non-negativity of DTC. However, in principle one could use any family \mathcal{S} and integer k as above to define a kind of mutual information among the random variables ξ_1, \dots, ξ_n : simply use the gap between the two sides in (20).

I do not see any advantage to these other quantities over DTC for the purposes of this paper. But in settings where the joint distribution of (ξ_1, \dots, ξ_n) has some known special structure, it could be worth exploring how that structure is detected by the right choice of such a generalized mutual information, and what consequences it has for other questions about those random variables.

Remark. — The sum of conditional entropies that we subtract in (18) has its own place in the information theory literature. Verdú and Weissman [110] call it the ‘erasure entropy’ of (ξ_1, \dots, ξ_n) , and use it to define the limiting erasure entropy rate of a stationary ergodic source. They derive operational interpretations in connection with (i) Gibbs sampling from a Markov random field and (ii) recovery of a signal passed through a binary erasure channel in the limit of small erasure probability. Their paper also makes use of the sum of conditional mutual informations that appears in the second right-hand term in (19), which they call ‘erasure mutual information’. For this quantity they give an operational interpretation in terms of the decrease in channel capacity due to sporadic erasures of the channel outputs.

Concerning erasure entropy, see also Proposition 16.15 below.

3.3. *The relationship between total correlation and dual total correlation.* — Although $\text{TC}(\mu)$ and $\text{DTC}(\mu)$ both quantify some kind of correlation among the coordinates under μ , they can behave quite differently.

Example 3.5. — Let ν_p be the p -biased distribution on $\{0, 1\}$, let $p \neq q$, and let $\mu = (\nu_p^{\times n} + \nu_q^{\times n})/2$, as in Example 3.3. In that example $\text{TC}(\mu)$ is at least $c(p, q)n - \log 2$, where

$$c(p, q) = \text{H}\left(\frac{p+q}{2}, 1 - \frac{p+q}{2}\right) - \frac{\text{H}(p, 1-p) + \text{H}(q, 1-q)}{2} > 0.$$

On the other hand, once n is large, the two measures $\nu_p^{\times(n-1)}$ and $\nu_q^{\times(n-1)}$ are very nearly disjoint. More precisely, there is a partition $(\mathbf{U}, \mathbf{U}^c)$ of $\{0, 1\}^{n-1}$ such that $\nu_q^{\times(n-1)}(\mathbf{U})$ and $\nu_p^{\times(n-1)}(\mathbf{U}^c)$ are both exponentially small in n . (Finding the best such partition is an application of the classical Neyman–Pearson lemma in hypothesis testing [11, Theorem 11.7.1].) It follows that $\mu\{\xi_{[n-1]} \in \mathbf{U}\}$ and $\mu\{\xi_{[n-1]} \in \mathbf{U}^c\}$ are both exponentially close to $1/2$. Using Bayes’ theorem and this partition, we obtain that the conditional distribution of ξ_i given the event $\{\xi_{[n]\setminus i} = \mathbf{z}\}$ is very close to ν_p if $\mathbf{z} \in \mathbf{U}$, and very close to ν_q if $\mathbf{z} \in \mathbf{U}^c$. In both of these approximations the total variation of the error is exponentially small in n . Therefore, re-using the estimate (17),

$$\begin{aligned} \text{DTC}(\mu) &= \text{H}_\mu(\xi) - \sum_{i=1}^n \text{H}_\mu(\xi_i | \xi_{[n]\setminus i}) \\ &\leq \log 2 + \frac{\text{H}(p, 1-p)n + \text{H}(q, 1-q)n}{2} \\ &\quad - n \cdot \mu\{\xi_{[n-1]} \in \mathbf{U}\} \cdot \text{H}(p, 1-p) \\ &\quad - n \cdot \mu\{\xi_{[n-1]} \in \mathbf{U}^c\} \cdot \text{H}(q, 1-q) + \text{O}(ne^{-c'n}) \\ &\leq \log 2 + \text{O}(ne^{-c'n}) \end{aligned}$$

for some $c' > 0$.

Thus, in this example, $\text{TC}(\mu)$ is larger than $\text{DTC}(\mu)$ by roughly a multiplicative factor of n .

Example 3.6. — Let $A := \mathbf{Z}/q\mathbf{Z}$ for some q , and let μ be the uniform distribution on the subgroup

$$U := \{(a_1, \dots, a_n) \in A^n : a_1 + \dots + a_n = 0\}.$$

Let $\xi = (\xi_1, \dots, \xi_n)$ be the n -tuple of A -valued coordinate projections. Then

$$H(\mu) = H_\mu(\xi) = (n-1) \log q$$

and

$$H_\mu(\xi_i) = \log q \quad \text{for each } i,$$

so

$$\text{TC}(\mu) = \log q.$$

On the other hand, given an element of U , any $n-1$ of its coordinates determine the last coordinate, and so

$$H(\xi_i | \xi_{[n]\setminus i}) = 0 \quad \text{for each } i.$$

Therefore

$$\text{DTC}(\mu) = (n-1) \log q.$$

In this example, $\text{DTC}(\mu)$ is larger than $\text{TC}(\mu)$ almost by a factor of n . In fact, this example also achieves the largest possible DTC with an alphabet of size q , since one always has

$$\begin{aligned} \text{DTC}(\mu) &= H_\mu(\xi) - \sum_{i=1}^n H_\mu(\xi_i | \xi_{[n]\setminus i}) \\ &= H_\mu(\xi_{[n-1]}) + H_\mu(\xi_n | \xi_{[n-1]}) - \sum_{i=1}^n H_\mu(\xi_i | \xi_{[n]\setminus i}) \\ &= H_\mu(\xi_{[n-1]}) - \sum_{i=1}^{n-1} H_\mu(\xi_i | \xi_{[n]\setminus i}) \\ &\leq H_\mu(\xi_{[n-1]}) \\ &\leq (n-1) \log q. \end{aligned}$$

As described in the previous subsection, our route to Theorem B goes through the analogous theorem with DTC in place of TC: Theorem 7.1 below. After proving Theorem 7.1, we need some estimate relating DTC and TC in order to deduce Theorem B. Example 3.6 shows that this estimate cannot simply bound DTC by a multiple of TC, since the factor of n which appears in that example is too large to lead to the correct dependence in Theorem B, part (a).

The key here is that a much better bound on DTC by TC is available if we allow ourselves to discard a few coordinates in A^n .

Lemma 3.7 (Discarding coordinates to control DTC using TC). — For any $\mu \in \text{Prob}(A^n)$ and $r > 0$ there exists $S \subseteq [n]$ with $|S| \geq (1 - r)n$ and such that

$$\text{DTC}(\mu_S) \leq r^{-1} \text{TC}(\mu).$$

Let $\xi = (\xi_1, \dots, \xi_n)$ have joint distribution μ . Beware that, on the left-hand side of the inequality above, $\text{DTC}(\mu_S)$ refers to the dual total correlation among the smaller family of random variables $(\xi_i)_{i \in S}$. In the proof below, we use DTC and TC for measures on product spaces of several different dimensions.

Proof. — Let us write $\alpha := r^{-1} \text{TC}(\mu)/n$ for brevity.

By the definition of DTC and the subadditivity of H , we have

$$(21) \quad \text{DTC}(\mu) \leq \sum_{i=1}^n [H(\xi_i) - H(\xi_i | \xi_{[n] \setminus i})].$$

If the measure μ happens to satisfy $H(\xi_i | \xi_{[n] \setminus i}) \geq H(\xi_i) - \alpha$ for every i , then (21) immediately gives $\text{DTC}(\mu) \leq \alpha n$.

So suppose there is an $i_1 \in [n]$ for which $H(\xi_{i_1} | \xi_{[n] \setminus i_1}) < H(\xi_{i_1}) - \alpha$. Then that failure and the chain rule give

$$\begin{aligned} \text{TC}(\mu_{[n] \setminus i_1}) &= \left(\sum_{i \in [n] \setminus i_1} H(\xi_i) \right) - H(\xi_{[n] \setminus i_1}) \\ &= \left(\sum_{i \in [n] \setminus i_1} H(\xi_i) \right) - [H(\xi) - H(\xi_{i_1} | \xi_{[n] \setminus i_1})] \\ &= \left(\sum_{i=1}^n H(\xi_i) \right) - H(\xi) - [H(\xi_{i_1}) - H(\xi_{i_1} | \xi_{[n] \setminus i_1})] \\ &< \text{TC}(\mu) - \alpha \end{aligned}$$

for the projection $\mu_{[n] \setminus i_1}$ of μ to $A^{[n] \setminus i_1}$.

We can now repeat the argument above. If every $i \in [n] \setminus i_1$ satisfies

$$H(\xi_i | \xi_{[n] \setminus \{i_1, i\}}) \geq H(\xi_i) - \alpha,$$

then the $(n - 1)$ -dimensional analog of (21) gives $\text{DTC}(\mu_{[n] \setminus i_1}) \leq \alpha(n - 1) \leq \alpha n$, and we select $S := [n] \setminus i_1$. Otherwise, there is an $i_2 \in [n] \setminus i_1$ such that

$$\mathbf{H}(\xi_{i_2}) - \mathbf{H}(\xi_{i_2} \mid \xi_{[n] \setminus \{i_1, i_2\}}) > \alpha,$$

and then removing i_2 leads to

$$\text{TC}(\mu_{[n] \setminus \{i_1, i_2\}}) < \text{TC}(\mu_{[n] \setminus i_1}) - \alpha < \text{TC}(\mu) - 2\alpha.$$

Since total correlations are non-negative, we cannot repeat this argument more than k times, where k is the largest integer for which

$$\text{TC}(\mu) - k\alpha > 0.$$

From the definition of α , this amounts to $k < rn$. Thus, after removing at most this many elements from $[n]$, we are left with a subset S of $[n]$ having both the required properties. \square

Lemma 3.7 is nicely illustrated by Example 3.6. In that example, the projection of μ to any $n - 1$ coordinates is a product measure, so the DTC of any such projection collapses to zero.

In Subsection 7.2 we use Lemma 3.7 (together with Lemma 5.10 below) to finish the deduction of Theorem B from Theorem 7.1.

4. Background on metrics and measure transportation

4.1. Normalized Hamming metrics. — For any finite set A and $n \in \mathbf{N}$, the **normalized Hamming metric** on A^n is defined by

$$(22) \quad d_n(\mathbf{a}, \mathbf{a}') := \frac{|\{i \in [n] : a_i \neq a'_i\}|}{n},$$

where $\mathbf{a} = (a_1, \dots, a_n)$, $\mathbf{a}' = (a'_1, \dots, a'_n) \in A^n$. All the metric spaces that appear in Parts II and III have the form (A^n, d_n) or some slight modification of it. We never use Hamming metrics that are not normalized.

More generally, if (K, d) is any metric space, then K^n may be endowed with the **normalized Hamming average** of copies of d , defined by

$$(23) \quad d_n(\mathbf{x}, \mathbf{x}') := \frac{1}{n} \sum_{i=1}^n d(x_i, x'_i) \quad \forall \mathbf{x}, \mathbf{x}' \in K^n.$$

We refer to this generalization occasionally in the sequel, but our main results concern the special case (22).

Remark. — In the alternative formula (15) for total correlation, the right-hand side makes sense for a measure μ on \mathbf{K}^n for any measurable space \mathbf{K} . This suggests an extension of Theorem B to measures μ on (\mathbf{K}^n, d_n) , where (\mathbf{K}, d) is a compact metric space and d_n is as in (23). In fact, this generalization is easily obtained from Theorem B itself by partitioning \mathbf{K} into finitely many subsets of diameter at most δ , using this partition to quantize the measure μ , and then letting $\delta \rightarrow 0$.

On the other hand, various steps in the proof of Theorem B are simpler in the discrete case, particularly once we switch to arguing about DTC instead of TC. We therefore leave this generalization of Theorem B aside.

4.2. Transportation metrics. — Let (\mathbf{K}, d) be a compact metric space. Let $\text{Prob}(\mathbf{K})$ be the space of Borel probability measures on \mathbf{K} , and endow it with the transportation metric

$$\bar{d}(\mu, \nu) := \inf \left\{ \int d(x, y) \lambda(dx, dy) : \lambda \text{ a coupling of } \mu \text{ and } \nu \right\}.$$

Let $\text{Lip}_1(\mathbf{K})$ be the space of all 1-Lipschitz functions $\mathbf{K} \rightarrow \mathbf{R}$. The transportation metric is indeed a metric on $\text{Prob}(\mathbf{K})$, and it generates the vague topology. See, for instance, [17, Section 11.8].

The transportation metric has a long history, and is variously associated with the names Monge, Kantorovich, Rubinstein, Wasserstein and others. Within ergodic theory, the special case when (\mathbf{K}, d) is a Hamming metric space (A^n, d_n) is the finitary version of Ornstein’s ‘d-bar’ metric, which Ornstein introduced independently in the course of his study of isomorphisms between Bernoulli shifts. Here I have adopted a notation based on that connection to ergodic theory. For a Hamming metric space (A^n, d_n) , the associated transportation metric on $\text{Prob}(A^n)$ is \bar{d}_n .

The transportation metric is closely related to many branches of geometry and to the field of optimal transportation: see, for instance, [27, Section 3 $\frac{1}{2}$.10], [23], and [54] for more on these connections. Within ergodic theory, it has also been explored beyond Ornstein’s original uses for it, particularly by Vershik and his co-workers [111–113].

The heart of the theory is the following dual characterization of this metric.

Theorem 4.1 Monge–Kantorovich–Rubinstein duality. — Any $\mu, \nu \in \text{Prob}(\mathbf{K})$ satisfy

$$\bar{d}(\mu, \nu) = \sup_{f \in \text{Lip}_1(\mathbf{K})} \left[\int f d\nu - \int f d\mu \right].$$

The precise statement of this result seems to originate in the papers [38–40]; see [17, Theorem 11.8.2] for a standard modern treatment.

One immediate consequence of Theorem 4.1 is that $\bar{d}(\nu, \mu)$ is always bounded by $\frac{1}{2}\|\nu - \mu\|$ times the diameter of (\mathbf{K}, d) . The following corollary improves this slightly, giving a useful estimate for certain purposes later.

Corollary 4.2. — Let $\mu, \nu \in \text{Prob}(\mathbf{K})$, and let $\mathbf{B}_1, \dots, \mathbf{B}_m$ be a partition of \mathbf{K} into Borel subsets. Then

$$\bar{d}(\nu, \mu) \leq \sum_{i=1}^m \mu(\mathbf{B}_i) \cdot \text{diam}(\mathbf{B}_i) + \frac{1}{2} \left(\sum_{i=1}^m |\nu(\mathbf{B}_i) - \mu(\mathbf{B}_i)| \right) \cdot \text{diam}(\mathbf{K}).$$

Proof. — Let $f \in \text{Lip}_1(\mathbf{K})$. For each $i = 1, 2, \dots, m$ we have

$$\begin{aligned} \int_{\mathbf{B}_i} f \, d\nu - \int_{\mathbf{B}_i} f \, d\mu &\leq \nu(\mathbf{B}_i) \cdot \sup(f|_{\mathbf{B}_i}) - \mu(\mathbf{B}_i) \cdot \inf(f|_{\mathbf{B}_i}) \\ &= \mu(\mathbf{B}_i) \cdot (\sup(f|_{\mathbf{B}_i}) - \inf(f|_{\mathbf{B}_i})) \\ &\quad + (\nu(\mathbf{B}_i) - \mu(\mathbf{B}_i)) \cdot \sup(f|_{\mathbf{B}_i}) \\ &\leq \mu(\mathbf{B}_i) \cdot \text{diam}(\mathbf{B}_i) + |\nu(\mathbf{B}_i) - \mu(\mathbf{B}_i)| \cdot \|f\|, \end{aligned}$$

where $\|\cdot\|$ is the supremum norm. Summing over i , this gives

$$\int f \, d\nu - \int f \, d\mu \leq \sum_{i=1}^m \mu(\mathbf{B}_i) \cdot \text{diam}(\mathbf{B}_i) + \sum_{i=1}^m |\nu(\mathbf{B}_i) - \mu(\mathbf{B}_i)| \cdot \|f\|.$$

Since $\nu(\mathbf{K}) = \mu(\mathbf{K}) = 1$, the left-hand side here is unchanged if we add any constant value to f . We may therefore shift f by a constant so that $\|f\| \leq \frac{1}{2} \text{diam}(\mathbf{K})$. We complete the proof by substituting this bound on $\|f\|$ above and then taking the supremum over all f . \square

To recover the bound by $\frac{1}{2} \|\nu - \mu\| \text{diam}(\mathbf{K})$ from this corollary, observe that if we take $\mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_m$ to be a partition into sets of very small diameter, then we can make the first right-hand term in Corollary 4.2 as small as we please, and the second term is still always bounded by $\frac{1}{2} \|\nu - \mu\| \text{diam}(\mathbf{K})$.

Here is another important estimate for our later purposes. It gives an upper bound on the transportation distance from a measure on a product space to a product of two other measures. It is well known, but we include a proof for completeness.

Lemma 4.3. — Let $(\mathbf{K}, d_{\mathbf{K}})$ and $(\mathbf{L}, d_{\mathbf{L}})$ be compact metric spaces, let $0 < \alpha < 1$, and let d be the following metric on $\mathbf{K} \times \mathbf{L}$:

$$(24) \quad d((x, y), (x', y')) := \alpha d_{\mathbf{K}}(x, x') + (1 - \alpha) d_{\mathbf{L}}(y, y').$$

Let $\mu \in \text{Prob}(\mathbf{K})$, $\nu \in \text{Prob}(\mathbf{L})$ and $\lambda \in \text{Prob}(\mathbf{K} \times \mathbf{L})$. Finally, let $\lambda_{\mathbf{K}}$ be the marginal of λ on \mathbf{K} , and let $x \mapsto \lambda_{\mathbf{L},x}$ be a conditional distribution for the \mathbf{L} -coordinate given the \mathbf{K} -coordinate under λ .

Then

$$\bar{d}(\lambda, \mu \times \nu) \leq \alpha \bar{d}_{\mathbf{K}}(\lambda_{\mathbf{K}}, \mu) + (1 - \alpha) \int_{\mathbf{K}} \bar{d}_{\mathbf{L}}(\lambda_{\mathbf{L},x}, \nu) \lambda_{\mathbf{K}}(dx).$$

Proof. — This can be proved directly from the definition or using Theorem 4.1. The latter approach is slightly shorter.

Let $f : \mathbf{K} \times \mathbf{L} \rightarrow \mathbf{R}$ be 1-Lipschitz for the metric d . Define the function $\bar{f} : \mathbf{K} \rightarrow \mathbf{R}$ by

$$\bar{f}(x) := \int_{\mathbf{L}} f(x, y) \nu(dy).$$

From (24) it follows that the function $f(x, \cdot)$ is $(1 - \alpha)$ -Lipschitz on the metric space $(\mathbf{L}, d_{\mathbf{L}})$ for each fixed $x \in \mathbf{K}$. Therefore

$$\begin{aligned} \int_{\mathbf{L}} [f(x, y) - \bar{f}(x)] \lambda_{\mathbf{L}, x}(dy) &= \int_{\mathbf{L}} f(x, y) \lambda_{\mathbf{L}, x}(dy) - \int_{\mathbf{L}} f(x, y) \nu(dy) \\ &\leq (1 - \alpha) \bar{d}_{\mathbf{L}}(\lambda_{\mathbf{L}, x}, \nu). \end{aligned}$$

Also, the function \bar{f} satisfies

$$|\bar{f}(x) - \bar{f}(x')| \leq \int_{\mathbf{L}} |f(x, y) - f(x', y)| \nu(dy) \leq \alpha d_{\mathbf{K}}(x, x') \quad \forall x, x' \in \mathbf{K},$$

so \bar{f} is α -Lipschitz on $(\mathbf{K}, d_{\mathbf{K}})$. Using the estimates above and Fubini's theorem, we now obtain

$$\begin{aligned} &\int_{\mathbf{K} \times \mathbf{L}} f d\lambda - \int_{\mathbf{K} \times \mathbf{L}} f d(\mu \times \nu) \\ &= \int_{\mathbf{K} \times \mathbf{L}} f d\lambda - \int_{\mathbf{K}} \bar{f} d\lambda_{\mathbf{K}} + \int_{\mathbf{K}} \bar{f} d\lambda_{\mathbf{K}} - \int_{\mathbf{K}} \bar{f} d\mu \\ &= \int_{\mathbf{K}} \left[\int_{\mathbf{L}} [f(x, y) - \bar{f}(x)] \lambda_{\mathbf{L}, x}(dy) \right] \lambda_{\mathbf{K}}(dx) + \int_{\mathbf{K}} \bar{f} d\lambda_{\mathbf{K}} - \int_{\mathbf{K}} \bar{f} d\mu \\ &\leq (1 - \alpha) \int_{\mathbf{K}} \bar{d}_{\mathbf{L}}(\lambda_{\mathbf{L}, x}, \nu) \lambda_{\mathbf{K}}(dx) + \alpha \bar{d}_{\mathbf{K}}(\lambda_{\mathbf{K}}, \mu). \end{aligned}$$

Taking the supremum over f , Theorem 4.1 completes the proof. \square

4.3. *A first connection to total correlation.* — The simple lemma below is needed only for one example later in the paper, but it also begins to connect this section with the previous ones.

Lemma 4.4. — *Let $\nu \in \text{Prob}(A)$, let $\mu \in \text{Prob}(A^n)$, and let $\delta := \bar{d}_n(\mu, \nu^{\times n})$. Then*

$$\text{TC}(\mu) \leq 2(\text{H}(\delta, 1 - \delta) + \delta \log |A|)n.$$

Proof. — Regard the coordinates in $A^n \times A^n$ as a pair of copies of each of the coordinates in A^n . For each $i = 1, 2, \dots, n$, let $\xi_i : A^n \times A^n \rightarrow A$ (respectively, ζ_i) be the projection to the first (respectively, second) copy of the i^{th} coordinate in A^n . Let $\lambda \in \text{Prob}(A^n \times A^n)$ be a coupling of μ and $\nu^{\times n}$ such that

$$\int d_n d\lambda = \frac{1}{n} \sum_{i=1}^n \lambda\{\xi_i \neq \zeta_i\} = \delta.$$

Let $\delta_i := \lambda\{\xi_i \neq \zeta_i\}$ for each i . The tuple ξ has distribution μ under λ , and the tuple ζ has distribution $\nu^{\times n}$.

Now the chain rule, monotonicity under conditioning, and several applications of Fano's inequality [11, Section 2.10] give

$$H_\lambda(\xi_i) \leq H_\lambda(\zeta_i) + H(\delta_i, 1 - \delta_i) + \delta_i \log |A|$$

for each i , and also

$$H_\lambda(\zeta) \leq H_\lambda(\xi) + \sum_{i=1}^n H_\lambda(\zeta_i | \xi_i) \leq H_\lambda(\xi) + \sum_{i=1}^n (H(\delta_i, 1 - \delta_i) + \delta_i \log |A|).$$

Therefore

$$\begin{aligned} \text{TC}(\mu) = \text{TC}(\xi) &= \sum_{i=1}^n H_\lambda(\xi_i) - H_\lambda(\xi) \\ &\leq \text{TC}(\zeta) + 2 \sum_{i=1}^n (H(\delta_i, 1 - \delta_i) + \delta_i \log |A|). \end{aligned}$$

Since $\text{TC}(\zeta) = 0$ and H is concave, this is at most

$$2(H(\delta, 1 - \delta) + \delta \log |A|)n. \quad \square$$

5. Preliminaries on measure concentration

Many of the first important applications of measure concentration were to the local theory of Banach spaces. Classic accounts with an emphasis on these applications can be found in [65, 66]: in particular, [66, Chapter 6] is an early introduction to the general metric-space framework and gives several examples of it. Independently, a concentration inequality due to Margulis found an early application to information theory in [2]. More recently, Gromov has described a very broad class of concentration phenomena for general metric spaces in [28, Chapter 3 $\frac{1}{2}$]. Within probability theory, measure concentration is now a large subject in its own right, with many aspects and many applications. Ledoux'

monograph [53] is dedicated to these. Some of the most central results for product measures can be found in Talagrand’s classic work [98], which contributed some very refined estimates in that setting. Concentration now appears as a branch of probability theory in a few advanced textbooks: Vershynin’s [114] is a recent example.

Although we need only a few of the important ideas here, they must be altered slightly from their usual presentation, in ways that are explained more carefully below. I therefore give complete proofs for all but the most classical results in this section.

In this paper, our principal notion of measure concentration takes the form of certain transportation inequalities that we call ‘T-inequalities’: see Definition 5.1 below. These tie together the metric and measure theoretic ideas of the preceding sections. Subsections 5.1 and 5.2 discuss some history and basic examples. Then Subsection 5.3 introduces an equivalent formulation of T-inequalities in terms of bounds on the exponential moments of 1-Lipschitz functions. This equivalent formulation gives an easy way to establish some basic stability properties of T-inequalities.

Although Theorems B and C are focused on T-inequalities, at a key point in the proof of Theorem B we must switch to concentration inequalities of a different kind. We call them L-inequalities because of their relation to logarithmic Sobolev inequalities. They are formulated in Definition 5.11, and the rest of Subsection 5.4 explains how they are related to T-inequalities.

The rest of this section concerns triples (\mathbf{K}, d, μ) consisting of a compact metric space (\mathbf{K}, d) and a measure $\mu \in \text{Prob}(\mathbf{K})$. We refer to such a triple as a **metric probability space**. We shall not consider ‘metric measure spaces’ that are not compact or have mass other than 1; indeed, the main theorems later in the paper all concern metrics on finite sets.

5.1. Transportation inequalities. — Here is the definition from the Introduction, placed in a more general setting:

Definition 5.1. — Let $\kappa, r > 0$. The metric probability space (\mathbf{K}, d, μ) satisfies the **T-inequality with parameters κ and r** , or $\mathbf{T}(\kappa, r)$, if any other $\nu \in \text{Prob}(\mathbf{K})$ satisfies

$$\bar{d}(\nu, \mu) \leq \frac{1}{\kappa} D(\nu \parallel \mu) + r.$$

In this definition, ‘T’ stands for ‘transportation’, but beware that other authors use ‘T’ for a variety of different inequalities involving transportation metrics. Indeed, Definition 5.1 may be seen as a linearization of a more standard family of transportation inequalities: those of the form

$$(25) \quad \bar{d}(\nu, \mu) \leq \sqrt{\frac{1}{\kappa} D(\nu \parallel \mu)}.$$

Here we refer to (25) as a **square-root transportation inequality**. The square root is a natural and important feature in many settings where such inequalities have been proved, such as product measures and Hamming metrics (discussed further below).

The papers [58] and [6] include good accounts of how square-root transportation inequalities relate to other notions of concentration. Those accounts are easily adapted to the T-inequalities of Definition 5.1. A survey of recent developments on a wide variety of transportation inequalities can be found in [26].

If (\mathbf{K}, d, μ) satisfies (25), then it also satisfies a whole family of T-inequalities as in Definition 5.1. This is because the inequality of arithmetic and geometric means gives

$$(26) \quad \sqrt{\frac{1}{\kappa} \mathbf{D}(v \parallel \mu)} \leq \frac{1}{4\kappa r} \mathbf{D}(v \parallel \mu) + r$$

for any $r > 0$. However, in our Theorems B and C it is important that we fix $r > 0$ in advance and then seek an inequality of the form $\mathbf{T}(\kappa n, r)$ for some κ . This fixed value of r may be viewed as a ‘distance cutoff’: the T-inequality we obtain is informative only about distances greater than or equal to r . Example 5.3 in the next subsection shows that this limitation is necessary.

If (\mathbf{K}, d, μ) has diameter at most 1 then it clearly satisfies $\mathbf{T}(\kappa, 1)$ for every $\kappa > 0$. Such a space also satisfies T-inequalities with arbitrarily small values of r . To see this, first observe from the diameter bound that

$$\bar{d}(v, \mu) \leq \frac{1}{2} \|v - \mu\| \quad \text{for any } v, \mu \in \text{Prob}(\mathbf{K}).$$

Next, Pinsker’s classical inequality [80, p. 15] (with the optimal constant established later in [13], [49] and [42, Theorem 6.11]) provides

$$(27) \quad \|\mu - v\| \leq \sqrt{2\mathbf{D}(v \parallel \mu)}.$$

Combining these inequalities with (26), we obtain that (\mathbf{K}, d, μ) satisfies $\mathbf{T}(8r, r)$ for every $r > 0$.

If an inequality $\mathbf{T}(\kappa, r)$ improves on this trivial estimate, then its strength lies in the trade-off between r and κ . Product spaces with Hamming metrics and product measures provide classical examples of stronger T-inequalities. Let (\mathbf{K}, d) be a metric space, and suppose for simplicity that its diameter is at most 1. Let d_n be the metric on \mathbf{K}^n given by (23). Starting with the simple estimate $\mathbf{T}(8r, r)$ for any probability measure on (\mathbf{K}, d) , one shows by induction on n that any product measure on (\mathbf{K}^n, d_n) satisfies a T-inequality with constants that improve as n increases. The inductive argument amounts to combining Lemma 4.3 with the chain rule for KL divergence. This elegant approach to measure concentration for product spaces is the main contribution of Marton [57] (with improvements and generalizations in [58]). The result is as follows.

Proposition 5.2 (See [57, Lemma 1] and [58, eqn (1.4)]). — If (\mathbf{K}, d) has diameter at most 1, if d_n is given by (23), and if μ is a product measure on \mathbf{K}^n , then

$$(28) \quad \bar{d}_n(v, \mu) \leq \sqrt{\frac{1}{2n} D(v \parallel \mu)} \quad \forall v \in \text{Prob}(\mathbf{K}^n).$$

In particular, combining with (26), the space (\mathbf{K}^n, d_n, μ) satisfies $T(8rn, r)$ for all $r > 0$.

Marton's paper [58] also describes how these inequalities may be seen as sharper and cleaner versions of estimates that Ornstein already uses in his work on isomorphisms between Bernoulli shifts.

In view of the equality (15), Proposition 5.2 implies a kind of reversal of Lemma 4.4.

We use Proposition 5.2 directly at one point in the proof of Theorem C below. It is also a strong motivation for the specific form of the T-inequalities obtained in Theorems B and C.

Square-root transportation inequalities have been proved for various other probability measures on metric spaces, such as Gaussian measures in Euclidean spaces [100]. Marton and others have also extended her work to various classes of weakly dependent stochastic processes: see, for instance, [16, 59, 60, 87].

5.2. Two families of examples. — The next examples show that in Theorem B we must fix the distance cutoff r in advance.

Example 5.3. — Fix a sequence of values $1 > \delta_1 > \delta_2 > \dots$ such that

$$(29) \quad \delta_n \longrightarrow 0 \quad \text{but} \quad \sqrt{n}\delta_n \longrightarrow \infty.$$

For each n , let $C_n \subseteq \{0, 1\}^n$ be a code of maximal cardinality such that its minimum distance is at least δ_n in the normalized Hamming metric. See, for instance, [116, Chapter 4] for these notions from coding theory. By the Gilbert–Varshamov bound [116, Section 4.2] and simple estimates on Hamming-ball volumes, we have

$$(30) \quad |C_n| \geq 2^{(1-o(1))n}.$$

Letting μ_n be the uniform distribution on C_n , it follows that

$$(31) \quad \text{TC}(\mu_n) \leq n \log 2 - \log |C_n| = o(1).$$

Now let ν be any measure which is absolutely continuous with respect to μ_n . Then ν is supported on C_n . Fix $\kappa > 0$. Assume that ν satisfies the square-root transportation inequality with constant κn : that is, the analog of (25) holds with ν as the reference measure instead of μ and with κn in place of κ . If n is sufficiently large, then we can deduce from this that ν must put most of its mass on a single point. To do this, consider any subset $U \subseteq C_n$ such that $\nu(U) \leq 1/2$, and set $a := \nu(U)$. Any coupling of $\nu|_U$ and ν must

transport mass at least a from U to U^c . Since any two elements of C_n are at least distance δ_n apart, this implies

$$(32) \quad \bar{d}_n(\nu, \nu|_U) \geq a\delta_n.$$

On the other hand, since $\nu(U) \geq a$, a simple calculation gives

$$D(\nu|_U \| \nu) = \log \frac{1}{\nu(U)} \leq \log \frac{1}{a}.$$

Therefore our assumed square-root transportation inequality gives

$$\bar{d}_n(\nu, \nu|_U) \leq \sqrt{\frac{\log(1/a)}{\kappa}} \cdot \frac{1}{\sqrt{n}}.$$

By the second part of (29), this is compatible with (32) only if a is $o(1)$ as $n \rightarrow \infty$. It follows that no choice of partition (U, U^c) can separate ν into two substantial pieces once n is large, and hence that ν must be mostly supported on a single point.

Now suppose we have a mixture which represents μ_n and with most of the weight on terms that satisfy a good square-root transportation inequality. Then the argument above shows that those terms must be close to delta masses. This is possible only if the number of terms is roughly $|C_n|$. By (30) and (31), this number is much larger than $\exp(O(\text{TC}(\mu_n)))$.

This example shows that one cannot hope to prove Theorem B with a square-root transportation inequality in place of $T(\kappa n, r)$. Simply by choosing δ_n to decay sufficiently slowly, many similar possibilities can also be ruled out: for instance, replacing the square root with an even smaller power cannot repair the problem.

Of course, this example does not contradict Theorem B itself. For any *fixed* $r > 0$, the lower bound in (32) is less than r once n is large enough, and for those n the terms in the mixture provided by Theorem B need not consist of delta masses. Indeed, I suspect that the classical ‘random’ construction of good codes C_n (see again [116, Chapter 4]) yields measures μ_n that already satisfy some inequality of the form $T(\kappa n, r)$, provided n is large enough in terms of r . It could be interesting to try to prove this.

In view of Proposition 5.2, it is worth emphasizing that many measures other than product measures also exhibit concentration. The next examples are very simple, but already hint at the diversity of such measures.

Example 5.4. — Let $n = 2\ell$ be even, let A be a finite set, and let $B := A \times A$. Then we have an obvious bijection

$$\begin{aligned} F: B^\ell &\longrightarrow A^n : ((a_{1,1}, a_{1,2}), (a_{2,1}, a_{2,2}), \dots, (a_{\ell,1}, a_{\ell,2})) \\ &\longmapsto (a_{1,1}, a_{1,2}, a_{2,1}, a_{2,2}, \dots, a_{\ell,1}, a_{\ell,2}). \end{aligned}$$

Let $\nu \in \text{Prob}(\mathbb{B})$ and let $\mu := F_*\nu^{\times \ell}$.

If d_ℓ and d_n denote the normalized Hamming metrics on \mathbb{B}^ℓ and \mathbb{A}^n respectively, then F is 1-Lipschitz from d_ℓ to d_n . This is because each coordinate in \mathbb{B}^ℓ gives two output coordinates under F , but the normalization constant in d_n is twice that in d_ℓ . It follows that μ inherits any T-inequality that is satisfied by $\nu^{\times \ell}$, for instance by using Proposition 5.6 below. Since Proposition 5.2 gives that $\nu^{\times \ell}$ satisfies $T(8r\ell, r)$, which we may write as $T(4rn, r)$, the same is true of μ .

However, if ν is not a product measure on $\mathbb{A} \times \mathbb{A}$, then μ is not truly a product measure on the product space \mathbb{A}^n . In fact, μ can be far away from any true product measure on \mathbb{A}^n . In simple cases we can deduce this from Lemma 4.4. For instance, suppose that ν is the uniform distribution on the diagonal $\{(a, a) : a \in \mathbb{A}\}$. Then

$$H(\mu) = H(\nu^{\times \ell}) = \ell \log |\mathbb{A}| = \frac{n}{2} \log |\mathbb{A}|.$$

But on each coordinate in \mathbb{A}^n , μ projects to the uniform distribution. Therefore

$$\sum_{i=1}^n H(\mu_{\{i\}}) = n \log |\mathbb{A}|.$$

Putting these calculations together, we obtain

$$\text{TC}(\mu) = n \log |\mathbb{A}| / 2.$$

In view of Lemma 4.4, this implies a uniform lower bound on $\overline{d}_n(\mu, \mu_0^{\times n})$ for any $\mu_0 \in \text{Prob}(\mathbb{A})$.

The last example also satisfies $H_\mu(\xi_i | \xi_{[n] \setminus i}) = 0$ for every i , since each coordinate in $\mathbb{A} \times \mathbb{A}$ determines the other under the diagonal distribution ν . Therefore we also have

$$\text{DTC}(\mu) = H(\mu) - 0 = n \log |\mathbb{A}| / 2.$$

Other choices of ν lead to even more striking properties. Let us sketch some of these without giving complete details. Assume $|\mathbb{A}|$ is very large, and let $\mathbb{C} \subseteq \mathbb{A} \times \mathbb{A}$ have size $|\mathbb{A}|^2/2$ and be ε -uniform for some very small ε , in the sense of quasirandomness for bipartite graphs (see, for instance, [7, Section IV.5]). For any fixed $\varepsilon > 0$, if $|\mathbb{A}|$ is large enough, then a random choice of \mathbb{C} is ε -uniform with high probability. Let ν be the uniform distribution on \mathbb{C} , and let $\mu := F_*\nu^{\times \ell}$.

An easy calculation gives that $\text{TC}(\mu)$ is at most $(n/2) \log 2$. In fact this is roughly an equality, up to a multiple of n which is small depending on the uniformity of \mathbb{C} , and $\text{DTC}(\mu)$ is also roughly $(n/2) \log 2$ up to a similar error.

We need the following consequence of uniformity: if M is any large constant, and if \mathbb{C} is sufficiently uniform in terms of M , then there is a small constant $c > 0$ depending only on M such that

$$(33) \quad (\mu_1 \times \mu_2)((\mathbb{A} \times \mathbb{A}) \setminus \mathbb{C}) \geq c$$

whenever $H(\mu_1)$ and $H(\mu_2)$ are both at least $\log |A| - M$. To see this, let γ be the uniform distribution on A , and recall that with this choice we have $D(\mu_i \parallel \gamma) = \log |A| - H(\mu_i)$. Therefore the assumed lower bound on $H(\mu_i)$ implies the upper bound

$$D(\mu_i \parallel \gamma) = \sum_{a \in A} \gamma(a) \frac{\mu_i(a)}{\gamma(a)} \log \frac{\mu_i(a)}{\gamma(a)} \leq M \quad \text{for } i = 1, 2.$$

Using the strict convexity of the function $t \mapsto t \log t$, one obtains from this bound a positive constant $c_1 \geq 1$, depending only on M , such that the set

$$D_i := \{a : \gamma(a)/c_1 \leq \mu_i(a) \leq c_1 \gamma(a)\}$$

satisfies $\mu_i(D_i) \geq 1/c_1$ for each $i = 1, 2$. In view of the definition of D_i , this set must also satisfy

$$(34) \quad \gamma(D_i) = \frac{|D_i|}{|A|} \geq \frac{1}{c_1} \mu_i(D_i) \geq \frac{1}{c_1^2}.$$

Because of these sets, we have

$$(35) \quad (\mu_1 \times \mu_2)((A \times A) \setminus C) \geq (\mu_1 \times \mu_2)((D_1 \times D_2) \setminus C) \geq \frac{|(D_1 \times D_2) \setminus C|}{c_1^2 |A|^2}.$$

If C is sufficiently uniform in terms of c_1 , then it and its complement must both intersect the product set $D_1 \times D_2$ in roughly half its elements. The uniformity can be applied here because (34) includes a fixed lower bound on the ratios $|D_i|/|A|$. It then follows that the fraction at the end of (35) is at least

$$\frac{|D_1 \times D_2|}{4c_1^2 |A|^2} \geq \frac{1}{4c_1^4},$$

by another appeal to (34). This proves (33) with $c := 1/4c_1^4$.

Having arranged (33), we now consider the function on A^n defined by

$$f(\mathbf{a}) := \frac{1}{n} \left| \{i \in \{1, 2, \dots, \ell\} : (a_{2i-1}, a_{2i}) \notin C\} \right|.$$

It is 1-Lipschitz for the metric d_n . Using (33), Markov's inequality, and some simple double counting, one finds another small but absolute constant c' such that the following holds: for any true product measure $\mu_1 \times \dots \times \mu_n$ on A^n , if

$$H(\mu_1 \times \dots \times \mu_n) = H(\mu_1) + \dots + H(\mu_n) \geq n \log |A| - nM/100,$$

then also

$$\int f d(\mu_1 \times \dots \times \mu_n) \geq c'.$$

Therefore, by Theorem 4.1 applied with the 1-Lipschitz function f , we have

$$(36) \quad \overline{d}_n(\mu_1 \times \cdots \times \mu_n, \mu') \geq c'$$

for any measure μ' which is supported on $F(C^\ell) = \{f = 0\}$.

On the other hand, $\mu := F_*\nu^{\times\ell}$ is supported on $F(C^\ell)$. If we write it as a mixture of any other measures, they must all be supported on $F(C^\ell)$ too. So now suppose we have such a mixture

$$\mu = p_1\mu_1 + \cdots + p_m\mu_m$$

with $m = \exp(O(\text{TC}(\mu))) = \exp(O(n))$. Since

$$H(\mu) = \ell \log |C| = n \log |A| - (n/2) \log 2,$$

and since any measure on A^n has entropy at most $n \log |A|$, Lemma 2.1 implies that most of the mass in the mixture is on terms μ_j for which $H(\mu_j) > n \log |A| - O(n)$, where ‘ $O(n)$ ’ is independent of $|A|$. In particular, if we previously chose M large enough, then this lower bound on entropy is greater than $n \log |A| - nM/100$. Combining this with (36), we deduce that most of the mass in this mixture is on measures μ_j that are at least distance c' from any product measure.

This gives an example μ that is highly concentrated, but cannot be written as a mixture of measures that are close in \overline{d}_n to product measures and where the number of terms is $\exp(O(\text{TC}(\mu)))$. This shows that the ‘good’ summands in Theorem B cannot be limited to measures that are close in \overline{d}_n to products—other examples of highly concentrated measures are sometimes necessary.

The ideas in Example 5.4 can be generalized in many ways. For instance, many other subsets and measures can be constructed by imposing more complicated pairwise constraints on the coordinates in A^n , or by imposing constraints on triples or even larger sets of coordinates. I expect that many such constructions exhibit measure concentration, although certainly many others do not.

However, all the cases in which I know how to prove measure concentration share a certain structure: like the measures $F_*\nu^{\times\ell}$ in Example 5.4, they are $O(1)$ -Lipschitz images of product measures on other product spaces. In that example, we use the map F to ‘hide’ the product structure of $\nu^{\times\ell}$, with the effect that $F_*\nu^{\times\ell}$ is far away from any product measure on A^n . But one might feel that this still has the essence of a product measure, and it suggests the following general question.

Question 5.5. — *Is it true that for any $r, \kappa, \varepsilon > 0$ there exists an $L > 0$ for which the following holds for all $n \in \mathbf{N}$?*

If (A^n, d_n, μ) satisfies $T(\kappa n, r)$, then there exist an integer $\ell \in [n/L, Ln]$, a finite set B , a distribution ν on B , and an L -Lipschitz map $F : B^\ell \rightarrow A^n$ such that

$$\overline{d}_n(\mu, F_*\nu^{\times \ell}) < \varepsilon.$$

The methods of the present paper do not seem to shed any light on this question. An answer in either direction would be very interesting. If the answer is Yes, then we have found a rather strong characterization of concentrated measures on (A^n, d_n) . If it is No, then finding a counterexample seems to require finding a new property of product measures which (i) survives under pushforward by Lipschitz maps but (ii) is not implied by concentration alone.

If Question 5.5 has a positive answer, then it could be regarded as an analog of a result about stationary processes: the equivalence between measure concentration of the finite-dimensional distributions and being a factor of a Bernoulli shift (see [91, Subsection IV.3.c]). However, Question 5.5 assumes much less structure than that result. This, in turn, could have consequences for the ergodic theory of non-amenable acting groups, for which factors of Bernoulli systems are far from being understood.

5.3. Exponential moment bounds and some consequences. — Using Monge–Kantorovich–Rubinstein duality, a T-inequality is equivalent to an exponential moment bound for 1-Lipschitz functions. This fact is essentially due to Bobkov and Götze [6], although they work with square-root transportation inequalities. The statement and proof in our linearized setting are easily obtained by modifying their argument.

Proposition 5.6 (Bobkov–Götze equivalence). — For a metric probability space (K, d, μ) , the following are equivalent:

- (a) The space (K, d, μ) satisfies $T(\kappa, r)$;
- (b) Any 1-Lipschitz function f satisfies

$$M_\mu(\kappa f) \leq e^{\kappa \langle f \rangle + \kappa r}.$$

Proof. — Theorem 4.1 implies that condition (a) is equivalent to

$$\kappa \int f \, d\nu - \kappa \int f \, d\mu \leq D(\nu \parallel \mu) + \kappa r$$

for all $\nu \in \text{Prob}(K)$ and 1-Lipschitz f . Re-arranging, this is equivalent to

$$(37) \quad -\kappa \int f \, d\mu \leq \kappa r + \left(D(\nu \parallel \mu) - \kappa \int f \, d\nu \right)$$

for all such ν and f .

For a given f , Lemma 2.5 asserts that the difference

$$D(\nu \parallel \mu) - \kappa \int f \, d\nu$$

is minimized by the Gibbs measure $\nu = \mu_{|\kappa f}$, and for this choice of ν that difference is equal to $-C_\mu(\kappa f)$. Therefore (37) is equivalent to the inequality

$$C_\mu(\kappa f) \leq \kappa \int f \, d\mu + \kappa r.$$

Exponentiating, we arrive at condition (b). \square

Bobkov and Götze's original result [6, Theorem 1.3] asserts that the square-root transportation inequality (25) is equivalent to the inequality

$$M_\mu(tf) \leq e^{t(f) + t^2/4\kappa}$$

for all 1-Lipschitz $f : \mathbf{K} \rightarrow \mathbf{R}$ and all $t \in \mathbf{R}$. This inequality is also easily linearized to arrive at part (b) of Proposition 5.6.

A first application of Bobkov and Götze's equivalence is a second proof of Proposition 5.2. Let (\mathbf{K}^n, d_n, μ) be as in that proposition. By the equivalence, it suffices to prove that any 1-Lipschitz function $f : \mathbf{K}^n \rightarrow \mathbf{R}$ and $t \in \mathbf{R}$ satisfy

$$(38) \quad \int e^{tf} \, d\mu \leq e^{t(f) + t^2/8n}.$$

Because μ is a product measure, this can be proved by the method of bounded martingale differences, using the filtration $(\mathcal{F}_i)_{i=1}^n$ of \mathbf{K}^n where \mathcal{F}_i consists of the Borel sets that depend on only the first i coordinates. One can also phrase this proof as an induction on n , using a conditional version of (38) with $n = 1$ at each step. The result when $n = 1$ is a classical lemma of Hoeffding [33], and its extension to the setting of 1-Lipschitz functions on product spaces is due to McDiarmid: see [63, Lemma 5.8] and the arguments that follow it. By a special case of the equivalence in Proposition 5.6, that lemma of Hoeffding is equivalent to Pinsker's inequality (27) with the optimal constant.

This use of the inequality (38) fits into a more general formalism of bounding the 'Laplace functional' of a metric probability space: see [53, Section 1.6].

The rest of this subsection gives three stability results for T-inequalities under different kinds of perturbation to a measure. These lemmas can all be proved directly from Definition 5.1, but the proofs become simpler and quicker using condition (b) from Proposition 5.6.

Lemma 5.7. — *Suppose that (\mathbf{K}, d, μ) satisfies $T(\kappa, r)$, and let $\nu \in \text{Prob}(\mathbf{K})$ be absolutely continuous with respect to μ and satisfy*

$$\frac{d\nu}{d\mu} \leq M \quad \mu\text{-a.s.}$$

for some $M \in [1, \infty)$. Then (\mathbf{K}, d, ν) satisfies

$$T\left(\kappa, 2\frac{\log M}{\kappa} + 2r\right).$$

Proof. — We verify the relevant version of condition (b) from Proposition 5.6. Let $\langle \cdot \rangle_\mu$ and $\langle \cdot \rangle_\nu$ denote integration with respect to μ and ν , respectively. Let $f \in \text{Lip}_1(\mathbf{K})$. Then

$$\int e^{\kappa f} d\nu = \int e^{\kappa f} \frac{d\nu}{d\mu} d\mu \leq M \int e^{\kappa f} d\mu = e^{\kappa(\log M/\kappa)} \int e^{\kappa f} d\mu.$$

By assumption, this is at most $\exp(\kappa \langle f \rangle_\mu + \kappa r + \kappa(\log M/\kappa))$. On the other hand, we have

$$D(\nu \parallel \mu) = \int \log \frac{d\nu}{d\mu} d\nu \leq \log M,$$

and so the original definition of $T(\kappa, r)$ gives

$$\langle f \rangle_\mu \leq \langle f \rangle_\nu + \bar{d}(\nu, \mu) \leq \langle f \rangle_\nu + \frac{1}{\kappa} \log M + r.$$

Combining these inequalities completes the proof. \square

Lemma 5.8. — Suppose that (\mathbf{K}, d, μ) satisfies $T(\kappa, r)$, let $\nu \in \text{Prob}(\mathbf{K})$, and assume there exists a coupling λ of μ and ν which is supported on the set

$$\{(x, y) \in \mathbf{K} \times \mathbf{K} : d(x, y) \leq \delta\}.$$

Then ν satisfies $T(\kappa, r + 2\delta)$.

Proof. — We verify the relevant version of condition (b) from Proposition 5.6. Let $\langle \cdot \rangle_\mu$ and $\langle \cdot \rangle_\nu$ denote integration with respect to μ and ν , respectively. Then

$$\int e^{\kappa f} d\nu = \int_{\mathbf{K} \times \mathbf{K}} e^{\kappa f(y)} \lambda(dx, dy).$$

Since $d(x, y) \leq \delta$ for λ -almost every (x, y) , and f is 1-Lipschitz, the last integral is bounded above by

$$\int_{\mathbf{K} \times \mathbf{K}} e^{\kappa f(x) + \kappa \delta} \lambda(dx, dy) = e^{\kappa \delta} \int e^{\kappa f} d\mu.$$

By assumption, this is at most $\exp(\kappa \langle f \rangle_\mu + \kappa r + \kappa \delta)$. Finally, a repeat of the same reasoning about λ and f gives $\langle f \rangle_\mu \leq \langle f \rangle_\nu + \delta$. Combining these inequalities completes the proof. \square

Proposition 5.9. — Suppose that (\mathbf{K}, d, μ) satisfies $T(\kappa, r)$, let $\nu \in \text{Prob}(\mathbf{K})$, and assume that $\bar{d}(\nu, \mu) \leq \delta^2$ for some $\delta \in [0, 1/8)$. Then there exists a Borel set $U \subseteq \mathbf{K}$ satisfying $\nu(U) \geq 1 - 4\delta$ and such that $\nu|_U$ satisfies

$$T\left(\kappa, \frac{8\delta + 2\log 4}{\kappa} + 4r + 4\delta\right).$$

Proof. — Let λ be a coupling of μ and ν such that $\int d \, d\lambda \leq \delta^2$, and let

$$W := \{(x, y) : d(x, y) \leq \delta\}.$$

Then Markov's inequality gives $\lambda(W) \geq 1 - \delta$. Let $\lambda_1 := \lambda|_W$, and let μ_1 and ν_1 be the first and second marginals of λ_1 , respectively.

For these new measures, we have

$$\frac{d\lambda_1}{d\lambda} = \frac{1}{\lambda(W)} \cdot 1_W \leq (1 - \delta)^{-1} \leq 1 + 2\delta,$$

and hence also $d\mu_1/d\mu \leq 1 + 2\delta$. Therefore, by Lemma 5.7, the measure μ_1 still satisfies $T(\kappa, r_1)$, where

$$r_1 := 2 \frac{\log(1 + 2\delta)}{\kappa} + 2r \leq 4\delta/\kappa + 2r.$$

Then, since λ_1 is supported on W , Lemma 5.8 promises that ν_1 still satisfies $T(\kappa, r_2)$, where $r_2 = r_1 + 2\delta$.

Let $\rho := d\nu_1/d\nu$, and let $f := 1 + 2\delta - \rho$. Since $\rho \leq 1 + 2\delta$, f is non-negative. On the other hand,

$$\int f \, d\nu = 1 + 2\delta - 1 = 2\delta.$$

Therefore another appeal to Markov's inequality gives that the set

$$U := \{\rho \geq 1/2 + 2\delta\} = \{f \leq 1/2\}$$

satisfies $\nu(U) \geq 1 - 4\delta$.

Finally, observe that

$$\nu|_U = \frac{1}{\nu(U)} \cdot 1_U \cdot \nu \leq \frac{1}{1 - 4\delta} \cdot (2\rho) \cdot \nu = \frac{2}{1 - 4\delta} \cdot \nu_1 \leq 4\nu_1.$$

By a second appeal to Lemma 5.7, this implies that $\nu|_U$ satisfies $T(\kappa, r_3)$, where

$$r_3 = 2r_2 + 2 \frac{\log 4}{\kappa} = 2(r_1 + 2\delta) + 2 \frac{\log 4}{\kappa} \leq \frac{8\delta + 2\log 4}{\kappa} + 4r + 4\delta.$$

□

Remarks.

1. Different estimates can be obtained by choosing different thresholds for ρ when defining the set U in the proof above. The choice we have made here is simple, but not canonical. This version suits our later purposes quite well, because in those cases κ is very large while r is only fairly small, so the expression with κ in the denominator is also very small.
2. The qualitative conclusion of Proposition 5.9 cannot be substantially improved: in particular, we cannot conclude that ν itself satisfies $T(\kappa, r)$ for some controlled value of r . To see this, suppose that K has diameter 1, and that ν equals $(1 - \delta^2)\mu + \delta^2\mu'$, where μ satisfies $T(\kappa, \delta)$ for some huge κ , but μ' is another measure whose support is at distance 1 from the support of μ . Then $\bar{d}(\nu, \mu) = \delta^2$, but that small multiple of μ' which appears in ν is hard to transport to the rest of ν . This prevents ν from satisfying $T(\kappa', r')$ with $r' \ll 1$ and $\kappa' \gg 1/\delta^2$, no matter how large κ is.

We finish this subsection with another stability property. This one applies to measures on Hamming spaces. It shows that T -inequalities survive when we lift those measures to spaces of slightly higher dimension.

Lemma 5.10 (Stability under lifting). — *Let μ be a probability measure on a Hamming metric space (A^n, d_n) . Let S be a nonempty subset of $[n]$, let μ_S be the projection of μ to A^S , and define $a \in [0, 1)$ by $|S| = (1 - a)n$. If μ_S satisfies $T(\kappa, r)$, then μ itself satisfies*

$$T\left(\frac{\kappa}{1 - a}, (1 - a)r + a\right).$$

Proof. — Let d_S be the normalized Hamming metric on A^S . Suppose that $\nu \in \text{Prob}(A^n)$ has $D(\nu \parallel \mu) < \infty$, and let ν_S be the projection of ν to A^S . Then $D(\nu_S \parallel \mu_S) \leq D(\nu \parallel \mu)$, because $d\nu_S/d\mu_S$ is the conditional μ -expectation of $d\nu/d\mu$ onto the σ -algebra generated by the coordinates in S , and so we may apply the conditional Jensen inequality in the definition (8). Therefore our assumption on μ_S gives

$$\bar{d}_S(\nu_S, \mu_S) \leq \frac{1}{\kappa} D(\nu \parallel \mu) + r.$$

Let λ_S be a coupling of ν_S and μ_S satisfying $\int d_S d\lambda_S = \bar{d}_S(\nu_S, \mu_S)$, and let λ be any lift of λ_S to a coupling of ν and μ . Then

$$\begin{aligned} \int d_n d\lambda &= \frac{|S|}{n} \int d_S d\lambda_S + \frac{n - |S|}{n} \int d_{[n] \setminus S}(\mathbf{x}_{[n] \setminus S}, \mathbf{y}_{[n] \setminus S}) \lambda(d\mathbf{x}, d\mathbf{y}) \\ &\leq (1 - a) \left(\frac{1}{\kappa} D(\nu_S \parallel \mu_S) + r \right) + a \\ &\leq (1 - a) \left(\frac{1}{\kappa} D(\nu \parallel \mu) + r \right) + a. \end{aligned} \quad \square$$

5.4. *A related functional inequality.*

Definition 5.11. — Fix (\mathbf{K}, d) as before, and let $\alpha > 0$ and $0 \leq \kappa_0 \leq \kappa \leq \infty$. Then $\mu \in \text{Prob}(\mathbf{K})$ satisfies the **L-inequality with parameters** κ_0, κ **and** α , or **L** $([\kappa_0, \kappa], \alpha)$, if

$$(39) \quad D(\mu|_{e^f} \parallel \mu) \leq \alpha t^2$$

for any 1-Lipschitz function f on \mathbf{K} and any real value $t \in [\kappa_0, \kappa]$.

This definition is closely related to another important branch of measure concentration theory: logarithmic Sobolev inequalities. Put roughly, given a metric probability space (\mathbf{K}, d, μ) and a suitable class of functions on \mathbf{K} , a logarithmic Sobolev inequality asserts that

$$(40) \quad D(\mu|_{e^f} \parallel \mu) \leq \alpha \int |\nabla f|^2 d\mu|_{e^f}$$

for all f in that class, where $|\nabla f|(x)$ is some notion of the ‘local gradient’ of f at the point $x \in \mathbf{K}$. If (\mathbf{K}, d) is a Riemannian manifold, then $|\nabla f|$ is often the norm of the true gradient function ∇f on \mathbf{K} . For various other metric spaces (\mathbf{K}, d) , including many discrete spaces, alternative notions are available.

Logarithmic Sobolev inequalities form a large area of study in functional analysis in their own right. A good introduction to these inequalities and their relationship to measure concentration can be found at the beginning of [6], which also gives many further references. Several more recent developments are covered in [26, Section 8].

Product spaces exhibit certain logarithmic Sobolev inequalities that improve with dimension, and these offer a third alternative route to Proposition 5.2: see [52, Section 4]. This approach, more than Marton’s original one, played an important role in the discovery of the proof of Theorem B. This is discussed further in Subsection 6.3.

In general, if f is Lipschitz, then any sensible choice for $|\nabla f|$ should be bounded by the Lipschitz constant of f . Thus, although we do not introduce any quantity that plays the role of $|\nabla f|$ in this paper, we can regard (39) as a stunted logarithmic Sobolev inequality. This is the reason for choosing the letter L.

L-inequalities imply T-inequalities. This link is one of the key tools in our proof of Theorem B below. To explain it, we begin with the following elementary estimate, which is also needed again later by itself.

Lemma 5.12. — Let (\mathbf{X}, μ) be a probability space, let $a < b$ be real numbers, and let $f : \mathbf{X} \rightarrow \mathbf{R}$ be a measurable function satisfying $a \leq f \leq b$ almost surely. Then

$$D(\mu|_{e^f} \parallel \mu) \leq (b - a)^2.$$

Proof. — Both sides of the desired inequality are unchanged if we add a constant to f , so we may assume that $a = 0$. Let $\langle \cdot \rangle$ denote integration with respect to μ . Since we are now assuming that $f \geq 0$ almost surely, we must have $\langle e^f \rangle \geq 1$, and therefore

$$\begin{aligned} \int f e^f d\mu &= \int f(e^f - \langle e^f \rangle) d\mu + \langle f \rangle \langle e^f \rangle \\ &\leq \int f(e^f - 1) d\mu + \langle f \rangle \langle e^f \rangle. \end{aligned}$$

Since $e^f - 1 \leq f e^f$ (for instance, by the convexity of \exp), the above implies that

$$\int f e^f d\mu \leq \int f^2 e^f d\mu + \langle f \rangle \langle e^f \rangle \leq b^2 \langle e^f \rangle + \langle f \rangle \langle e^f \rangle.$$

Dividing by $\langle e^f \rangle$ and re-arranging, this inequality completes the proof. \square

Lemma 5.12 and its proof are well known, but I do not know their origins, and have included them for completeness. The proof actually gives the stronger inequality

$$D(\mu_{|e^f} \parallel \mu) \leq \int (f - \min f)^2 d\mu_{|e^f}.$$

This version may be regarded as a logarithmic Sobolev inequality of the kind in (40). For the present paper we need only the simpler statement in Lemma 5.12.

Now we show how to turn an L-inequality into a T-inequality. Intuitively, we find that an L-inequality is a ‘differential version’ of a T-inequality. This is made precise by the following proof, which follows the classic Herbst argument from the study of logarithmic Sobolev inequalities, with some slight adjustments to the present setting. See [3, 51] for early expositions of Herbst’s original (unpublished) argument.

Proposition 5.13. — *If (\mathbf{K}, d) has diameter at most 1 and $0 < r < \kappa < \infty$, then $L([r, \kappa], r/\kappa)$ implies $T(\kappa, 2r)$.*

Proof. — Let f be 1-Lipschitz, and define $\varphi : [0, \kappa] \rightarrow \mathbf{R}$ by

$$\varphi(t) := \begin{cases} \langle f \rangle & \text{if } t = 0, \\ \frac{1}{t} C_\mu(f) & \text{if } 0 < t \leq \kappa. \end{cases}$$

A simple argument shows that φ is continuous at 0, and a little calculus gives

$$\varphi'(t) = \frac{1}{t^2} \left(t \int f d\mu_{|e^f} - C_\mu(f) \right) = \frac{1}{t^2} D(\mu_{|e^f} \parallel \mu) \quad \text{for } t > 0,$$

where the second equality comes from Lemma 2.5. Since (\mathbb{K}, d) has diameter at most 1 and f is 1-Lipschitz, Lemma 5.12 gives $\varphi'(t) \leq 1$ for any t . On the other hand, $L([r, \kappa], r/\kappa)$ gives $\varphi'(t) \leq r/\kappa$ whenever $r \leq t \leq \kappa$. Combining these bounds, we obtain

$$\begin{aligned} C_\mu(\kappa f) &= \kappa \varphi(\kappa) = \kappa \varphi(0) + \kappa \int_0^\kappa \varphi'(t) dt \\ &= \kappa \varphi(0) + \kappa \int_0^r \varphi'(t) dt + \kappa \int_r^\kappa \varphi'(t) dt \\ &\leq \kappa \varphi(0) + 2\kappa r = \kappa \langle f \rangle + 2\kappa r. \end{aligned}$$

This completes the proof via part (b) of Proposition 5.6. □

5.5. *The role of the inequalities in the rest of this paper.* — T-inequalities already appear in the formulation of Theorem B. However, the proof of that theorem involves L-inequalities as well. Put very roughly, in part of the proof of that theorem, we assume that a given measure μ does not already satisfy $T(rn/1200, r)$, deduce from Proposition 5.13 that a related L-inequality also fails, and then use the function f and parameter t from that failure to write μ as a mixture of two ‘better’ measures. This part of the proof of Theorem B is the subject of the next section. I do not know an approach that works directly with the violation of the T-inequality, and avoids the L-inequality.

6. A relative of Theorem B

The proof of Theorem B follows two other decomposition results, which come progressively closer to the conclusion we really want. The biggest part of the proof goes towards the first of those auxiliary decompositions.

In the rest of this section, we always use $\langle \cdot \rangle$ to denote expectation with respect to a measure μ on A^n . Here is the first auxiliary decomposition:

Theorem 6.1. — *Let $\varepsilon, r > 0$. For any $\mu \in \text{Prob}(A^n)$ there exists a fuzzy partition (ρ_1, \dots, ρ_k) such that*

- (a) $I_\mu(\rho_1, \dots, \rho_k) \leq 2 \cdot \text{DTC}(\mu)$ (recall (10) for the notation on the left here),
- (b) $\langle \rho_1 \rangle < \varepsilon$, and
- (c) $\langle \rho_j \rangle > 0$ and the measure $\mu_{|\rho_j}$ satisfies $T(rn/200, r)$ for every $j = 2, 3, \dots, k$.

This differs from Theorem B in two respects. Firstly, it gives control only over the mutual information $I_\mu(\rho_1, \dots, \rho_k)$, not the actual number of terms k . Secondly, that control is in terms of $\text{DTC}(\mu)$, not $\text{TC}(\mu)$ as in the statement of Theorem B. The use of DTC here is crucial. Constructing the fuzzy partition in Theorem 6.1 requires careful

estimates on the behaviour of DTC, and no suitable estimates seem to be available for TC.

In the next section we improve Theorem 6.1 to Theorem 7.1, which does control the number of functions in the fuzzy partition, but still using $\text{DTC}(\mu)$. Finally, we complete the proof of Theorem B using Lemma 3.7, which finds a slightly lower-dimensional projection of μ whose DTC is controlled in terms of $\text{TC}(\mu)$. I do not know a more direct way to use $\text{TC}(\mu)$ for the proof of Theorem B.

6.1. The DTC-decrement argument. — The key idea for the proof of Theorem 6.1 is this: if μ itself does not satisfy $T(rn/200, r)$, then we can turn that failure of concentration into a representation of μ as a mixture whose terms have substantially smaller DTC on average. In this mixture, if much of the mass is on pieces that still fail $T(rn/200, r)$, then those can be decomposed again. This process cannot continue indefinitely, because the average DTC of the terms in these mixtures must remain non-negative. Thus, after a finite number of steps, we reach a mixture whose terms mostly do satisfy $T(rn/200, r)$. Moreover, our specific estimate on the decrement in average DTC at each step turns into a bound on the mutual information of this mixture: up to a constant, it is at most the total reduction we achieved in the average DTC, which in turn is at most the initial value $\text{DTC}(\mu)$.

In Subsection 6.3, we try to give an intuitive reason for expecting that such an argument can be carried out using DTC. I do not know of any reason to hope for a similar argument using TC. The reader may wish to consult Subsection 6.3 before finishing the rest of this section.

Here is the proposition that provides this decrement in average DTC:

Proposition 6.2. — *If $\mu \in \text{Prob}(A^n)$ does not satisfy $T(rn/200, r)$, then there is a fuzzy partition (ρ_1, ρ_2) such that*

$$I_\mu(\rho_1, \rho_2) \geq r^2 n^{-1} e^{-n}$$

and

$$(41) \quad \langle \rho_1 \rangle \cdot \text{DTC}(\mu|_{\rho_1}) + \langle \rho_2 \rangle \cdot \text{DTC}(\mu|_{\rho_2}) \leq \text{DTC}(\mu) - \frac{1}{2} I_\mu(\rho_1, \rho_2).$$

It is worth noting an immediate corollary: if a measure μ has $\text{DTC}(\mu) < r^2 n^{-1} e^{-n}/2$, then μ must already satisfy $T(rn/200, r)$, simply because the left-hand side of (41) cannot be negative. This corollary can be seen as a robust version of concentration for product measures. However, since the value $r^2 n^{-1} e^{-n}/2$ is so small as n grows, I doubt that this version has any advantages over more standard results. For instance, the combination of Propositions 5.2 and 5.9 shows that if $\text{TC}(\mu)$ is $o(n)$ then we can trim off a small piece of μ and be left with strong measure concentration. By another of Han's inequalities from [32], $\text{TC}(\mu)$ is always at most $(n-1) \cdot \text{DTC}(\mu)$, so we obtain the same

conclusion if $\text{DTC}(\mu)$ is $o(1)$. Some recent stronger results are cited in Subsection 7.3 below. Qualitatively, this is the best one can do: the mixture of two product measures in Example 3.5 has DTC roughly $\log 2$ and does not exhibit any non-trivial concentration. The particular inequality $\text{DTC}(\mu) < r^2 n^{-1} e^{-n}/2$ does not play any role in the rest of this paper.

Most of this section is spent proving Proposition 6.2. Before doing so, let us explain how it implies Theorem 6.1. This implication starts with the following corollary.

Corollary 6.3. — *Let $\mu \in \text{Prob}(A^n)$, and let $(\rho_i)_{i=1}^k$ be a fuzzy partition. Let S be the set of all $i \in [k]$ for which $\mu_{|\rho_i}$ does not satisfy $T(rn/200, r)$. Then there is another fuzzy partition $(\rho'_j)_{j=1}^\ell$ with the following properties:*

(a) *(growth in mutual information)*

$$I_\mu(\rho'_1, \dots, \rho'_\ell) \geq I_\mu(\rho_1, \dots, \rho_k) + r^2 n^{-1} e^{-n} \sum_{i \in S} \langle \rho_i \rangle;$$

(b) *(proportional decrease in average DTC)*

$$\begin{aligned} & \sum_{j=1}^{\ell} \langle \rho'_j \rangle \cdot \text{DTC}(\mu_{|\rho'_j}) \\ & \leq \sum_{i=1}^k \langle \rho_i \rangle \cdot \text{DTC}(\mu_{|\rho_i}) - \frac{1}{2} [I_\mu(\rho'_1, \dots, \rho'_\ell) - I_\mu(\rho_1, \dots, \rho_k)]. \end{aligned}$$

Proof. — For each $i \in S$, let $(\rho_{i,1}, \rho_{i,2})$ be the fuzzy partition obtained by applying Proposition 6.2 to the measure $\mu_{|\rho_i}$. Now let $(\rho'_j)_{j=1}^\ell$ be an enumeration of the following functions:

$$\text{all } \rho_i \text{ for } i \in [k] \setminus S \text{ and all } \rho_i \cdot \rho_{i,1} \text{ and } \rho_i \cdot \rho_{i,2} \text{ for } i \in S.$$

These functions satisfy

$$\sum_{i \in [k] \setminus S} \rho_i + \sum_{i \in S} \rho_i \cdot \rho_{i,1} + \sum_{i \in S} \rho_i \cdot \rho_{i,2} = \sum_{i \in [k] \setminus S} \rho_i + \sum_{i \in S} \rho_i = 1,$$

so $(\rho'_j)_{j=1}^\ell$ is a new fuzzy partition.

Let us show that this new fuzzy partition has properties (a) and (b). Both of these follow from the version of the chain rule for mutual information and fuzzy partitions presented in equation (13). In the present setting, that rule gives

$$(42) \quad I_\mu(\rho'_1, \dots, \rho'_\ell) - I_\mu(\rho_1, \dots, \rho_k) = \sum_{i \in S} \langle \rho_i \rangle \cdot I_{\mu_{|\rho_i}}(\rho_{i,1}, \rho_{i,2}).$$

By the first conclusion of Proposition 6.2, this is at least

$$r^2 n^{-1} e^{-n} \sum_{i \in S} \langle \rho_i \rangle.$$

This proves property (a).

On the other hand, the second conclusion of Proposition 6.2 gives

$$\begin{aligned} & \left(\int \rho_{i,1} d\mu_{|\rho_i} \right) \cdot \text{DTC}(\mu_{|\rho_i \rho_{i,1}}) + \left(\int \rho_{i,2} d\mu_{|\rho_i} \right) \cdot \text{DTC}(\mu_{|\rho_i \rho_{i,2}}) \\ & \leq \text{DTC}(\mu_{|\rho_i}) - \frac{1}{2} \mathbf{I}_{\mu_{|\rho_i}}(\rho_{i,1}, \rho_{i,2}). \end{aligned}$$

We multiply this inequality by $\langle \rho_i \rangle$, observe that

$$\langle \rho_i \rangle \cdot \left(\int \rho_{i,a} d\mu_{|\rho_i} \right) = \langle \rho_i \rho_{i,a} \rangle \quad \text{for } a \in \{1, 2\},$$

and add the results over $i \in S$. This gives

$$\begin{aligned} & \sum_{j=1}^{\ell} \langle \rho'_j \rangle \cdot \text{DTC}(\mu_{|\rho'_j}) \\ & = \sum_{i \in [k] \setminus S} \langle \rho_i \rangle \cdot \text{DTC}(\mu_{|\rho_i}) \\ & \quad + \sum_{i \in S} \left(\langle \rho_i \rho_{i,1} \rangle \cdot \text{DTC}(\mu_{|\rho_i \rho_{i,1}}) + \langle \rho_i \rho_{i,2} \rangle \cdot \text{DTC}(\mu_{|\rho_i \rho_{i,2}}) \right) \\ & \leq \sum_{i \in [k] \setminus S} \langle \rho_i \rangle \cdot \text{DTC}(\mu_{|\rho_i}) + \sum_{i \in S} \langle \rho_i \rangle \cdot \text{DTC}(\mu_{|\rho_i}) \\ & \quad - \frac{1}{2} \sum_{i \in S} \langle \rho_i \rangle \cdot \mathbf{I}_{\mu_{|\rho_i}}(\rho_{i,1}, \rho_{i,2}). \end{aligned}$$

Substituting from (42), the right-hand side here is equal to

$$\sum_{i=1}^k \langle \rho_i \rangle \cdot \text{DTC}(\mu_{|\rho_i}) - \frac{1}{2} [\mathbf{I}_{\mu}(\rho'_1, \dots, \rho'_\ell) - \mathbf{I}_{\mu}(\rho_1, \dots, \rho_k)].$$

This proves property (b). □

Proof of Theorem 6.1 from Corollary 6.3. — We produce a finite sequence of fuzzy partitions $(\rho_{t,j})_{j=1}^{k_t}$ for $t = 0, 1, 2, \dots, t_0$ by the following recursion.

To start the recursion, let $k_0 = 1$ and let $\rho_{0,1}$ be the constant function 1. If μ already satisfies $T(rn/200, r)$, then we let $t_0 = 0$ and stop the recursion here.

Now suppose that μ does not satisfy $T(rn/200, r)$ and that we have already constructed the fuzzy partition $(\rho_{t,j})_{j=1}^{k_t}$ for some $t \geq 0$. Let S be the set of all $j \in [k_t]$ for which $\mu|_{\rho_{t,j}}$ does not satisfy $T(rn/200, r)$. If

$$(43) \quad \sum_{j \in S} \langle \rho_{t,j} \rangle < \varepsilon,$$

then set $t_0 := t$ and stop the recursion. Otherwise, let $(\rho_{t+1,j})_{j=1}^{k_{t+1}}$ be a new fuzzy partition produced from $(\rho_{t,j})_{j=1}^{k_t}$ by Corollary 6.3.

If this recursion does not stop at stage t , then the negation of (43) must hold at that stage. Combining this fact with both conclusions of Corollary 6.3, it follows that

$$\sum_{j=1}^{k_{t+1}} \langle \rho_{t+1,j} \rangle \cdot \text{DTC}(\mu|_{\rho_{t+1,j}}) \leq \sum_{j=1}^{k_t} \langle \rho_{t,j} \rangle \cdot \text{DTC}(\mu|_{\rho_{t,j}}) - \frac{1}{2} \varepsilon r^2 n^{-1} e^{-n}.$$

Since the average DTC on the left cannot be negative, and the average DTC decreases here by at least the fixed amount $\varepsilon r^2 n^{-1} e^{-n}/2$, the recursion must stop at some finite stage.

For each $t = 0, 1, \dots, t_0 - 1$, conclusion (b) of Corollary 6.3 gives that

$$\begin{aligned} & I_\mu(\rho_{t+1,1}, \dots, \rho_{t+1,k_{t+1}}) - I_\mu(\rho_{t,1}, \dots, \rho_{t,k_t}) \\ & \leq 2 \left[\sum_{j=1}^{k_t} \langle \rho_{t,j} \rangle \cdot \text{DTC}(\mu|_{\rho_{t,j}}) - \sum_{j=1}^{k_{t+1}} \langle \rho_{t+1,j} \rangle \cdot \text{DTC}(\mu|_{\rho_{t+1,j}}) \right]. \end{aligned}$$

Summing these inequalities over those values of t , we obtain

$$(44) \quad I_\mu(\rho_{t_0,1}, \dots, \rho_{t_0,k_{t_0}}) \leq 2 \cdot \text{DTC}(\mu) - 2 \sum_{j=1}^{k_0} \langle \rho_{t_0,j} \rangle \cdot \text{DTC}(\mu|_{\rho_{t_0,j}}) \leq 2 \cdot \text{DTC}(\mu).$$

To finish, define (ρ_1, \dots, ρ_k) by letting ρ_1 be the sum of all those functions $\rho_{t_0,j}$ with $1 \leq j \leq k_{t_0}$ for which $\mu|_{\rho_{t_0,j}}$ does not satisfy $T(rn/200, r)$, and letting ρ_2, \dots, ρ_k be an enumeration of the remaining entries in $(\rho_{t_0,1}, \dots, \rho_{t_0,k_{t_0}})$. This new fuzzy partition (ρ_1, \dots, ρ_k) satisfies all three of the required conclusions. Conclusion (a) follows from (44) and an application of inequality (14). Conclusions (b) and (c) are written into the conditions for stopping the recursion. \square

Remark. — The above proof gives a bound on the number of entries k in the resulting fuzzy partition, but the bound is too poor to be useful beyond guaranteeing that the recursion terminates. This is because the quantity $r^2 n^{-1} e^{-n}$ appearing in Proposition 6.2 is so small.

Remark. — The recursion above terminates at the first stage t_0 when most terms in the mixture exhibit sufficient concentration. Then conclusion (a) is proved using the fact that the average DTC must remain non-negative. However, there is no guarantee that most of the resulting measures $\mu_{|\rho_j}$ have DTC close to zero. It may be that we have obtained all the measure concentration we want, but many of the values $\text{DTC}(\mu_{|\rho_j})$ are still large. Indeed, if μ is one of the measures constructed using an ε -uniform set in Example 5.4, then μ already satisfies a better T-inequality than those obtained in Theorem 6.1, and it cannot be decomposed into nontrivial summands using Proposition 6.2. However, as discussed in that example, its TC and DTC are both of order n , and it cannot be decomposed efficiently into approximate product measures.

Remark. — ‘Increment’ and ‘decrement’ arguments have other famous applications, particularly in extremal combinatorics. Perhaps the best known is the proof of Szemerédi’s regularity lemma [95, 96], which uses an argument that is now often called the ‘energy increment’. See [7, Section IV.5] for a modern textbook treatment of Szemerédi’s lemma, and see [101] and [103, Chapters 10 and 11] for a broader discussion of increment arguments in additive combinatorics.

However, having made this connection, we should also stress the following difference. In Szemerédi’s regularity lemma, the vertex set of a large finite graph is partitioned into a controlled number of cells so that most *pairs* among those cells have a property called ‘quasirandomness’. This pairwise requirement on the cells leads to the extremely large tower-type bounds that necessarily appear in Szemerédi’s lemma: see [24]. By contrast, Theorem 6.1 produces summands that mostly have a desired property—the T-inequality—individually, but are not required to interact with each other in any particular way. This makes for less dramatic bounds: in particular, for the simple relationship in part (a) of that theorem.

Another precedent for our ‘decrement’ argument is the work of Linial, Samorodnitsky and Wigderson [56] giving a strongly polynomial time algorithm for permanent estimation. Their algorithm relies on obtaining a substantial decrement in the permanent of a nonnegative matrix under a procedure called matrix scaling.

6.2. *Proof of the DTC decrement.* — The rest of this section is spent proving Proposition 6.2.

Given μ and a fuzzy partition (ρ_1, ρ_2) , let (ζ, ξ) be a randomization of the resulting mixture

$$\mu = \langle \rho_1 \rangle \cdot \mu_{|\rho_1} + \langle \rho_2 \rangle \cdot \mu_{|\rho_2}.$$

Thus, (ζ, ξ) takes values in $\{1, 2\} \times \mathbf{X}$. For this pair of random variables, the left-hand side of (19) is

$$\text{DTC}(\mu) - \langle \rho_1 \rangle \cdot \text{DTC}(\mu_{|\rho_1}) - \langle \rho_2 \rangle \cdot \text{DTC}(\mu_{|\rho_2}).$$

The proof of Proposition 6.2 rests on a careful analysis of the right-hand side of (19) for this pair (ζ, ξ) , which reads

$$(45) \quad \mathbf{I}(\xi; \zeta) - \sum_{i=1}^n \mathbf{I}(\xi_i; \zeta \mid \xi_{[n] \setminus i}).$$

We next re-write this expression in terms of μ , ρ_1 and ρ_2 .

By Corollary 2.3, the first term in (45) is equal to

$$(46) \quad \langle \rho_1 \rangle \cdot \mathbf{D}(\mu_{|\rho_1} \parallel \mu) + \langle \rho_2 \rangle \cdot \mathbf{D}(\mu_{|\rho_2} \parallel \mu).$$

The remaining terms in (45) can be expressed similarly. To this end, for each i , let $(\theta_{i,\mathbf{z}} : \mathbf{z} \in A^{[n] \setminus i})$ be a conditional distribution for ξ given $\xi_{[n] \setminus i}$ according to μ . Thus, the only remaining randomness under $\theta_{i,\mathbf{z}}$ is in the coordinate ξ_i . This conditional distribution represents μ as the following mixture:

$$(47) \quad \mu = \int \theta_{i,\mathbf{z}} \mu_{[n] \setminus i}(\mathrm{d}\mathbf{z}).$$

For each i and $\mathbf{z} \in A^{[n] \setminus i}$, let $\langle \cdot \rangle_{i,\mathbf{z}}$ denote integration with respect to $\theta_{i,\mathbf{z}}$. If we condition on the event $\{\xi_{[n] \setminus i} = \mathbf{z}\}$, then the pair (ζ, ξ) becomes a randomization of the mixture

$$\theta_{i,\mathbf{z}} = \langle \rho_1 \rangle_{i,\mathbf{z}} \cdot (\theta_{i,\mathbf{z}})_{|\rho_1} + \langle \rho_2 \rangle_{i,\mathbf{z}} \cdot (\theta_{i,\mathbf{z}})_{|\rho_2}.$$

This is because the conditional probability of the event $\{(\zeta, \xi) = (j, \mathbf{x})\}$ given the event $\{\xi_{[n] \setminus i} = \mathbf{z}\}$ equals

$$\frac{\rho_j(\mathbf{x}) \cdot \mu(\mathbf{x})}{\mu_{[n] \setminus i}(\mathbf{z})} = \rho_j(\mathbf{x}) \cdot \theta_{i,\mathbf{z}}(\mathbf{x})$$

whenever $j \in \{1, 2\}$, $\mathbf{x} \in A^n$ and $\mathbf{z} = \mathbf{x}_{[n] \setminus i}$. Now another appeal to Corollary 2.3 gives

$$(48) \quad \begin{aligned} \mathbf{I}(\xi_i; \zeta \mid \xi_{[n] \setminus i} = \mathbf{z}) &= \mathbf{I}(\xi; \zeta \mid \xi_{[n] \setminus i} = \mathbf{z}) \\ &= \langle \rho_1 \rangle_{i,\mathbf{z}} \cdot \mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_1} \parallel \theta_{i,\mathbf{z}}) + \langle \rho_2 \rangle_{i,\mathbf{z}} \cdot \mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_2} \parallel \theta_{i,\mathbf{z}}) \end{aligned}$$

for $\mu_{[n] \setminus i}$ -almost every \mathbf{z} . Averaging over \mathbf{z} , this becomes

$$(49) \quad \begin{aligned} \mathbf{I}(\xi_i; \zeta \mid \xi_{[n] \setminus i}) &= \int \langle \rho_1 \rangle_{i,\mathbf{z}} \cdot \mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_1} \parallel \theta_{i,\mathbf{z}}) \mu_{[n] \setminus i}(\mathrm{d}\mathbf{z}) \\ &\quad + \int \langle \rho_2 \rangle_{i,\mathbf{z}} \cdot \mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_2} \parallel \theta_{i,\mathbf{z}}) \mu_{[n] \setminus i}(\mathrm{d}\mathbf{z}). \end{aligned}$$

For the proof of Proposition 6.2, we use these calculations in the following special case: suppose that $f : A^n \rightarrow [0, 1]$ is 1-Lipschitz and that $0 \leq t \leq n/200$, and let

$$(50) \quad \rho_1 := \frac{1}{2} e^{-tf} \quad \text{and} \quad \rho_2 := 1 - \frac{1}{2} e^{-tf}.$$

Note that

$$0 < \rho_1 \leq \frac{1}{2} \quad \text{and} \quad \frac{1}{2} \leq \rho_2 < 1.$$

These bounds simplify the proof below, and are the reason for the factor of $\frac{1}{2}$ in the definition of ρ_1 .

For this choice of fuzzy partition, we need an upper bound for the right-hand side in (49). It has two terms, which we estimate separately. The key to both estimates is the following geometric feature of the present setting: the measure $\theta_{i,\mathbf{z}}$ is supported on the set

$$S_{i,\mathbf{z}} := \{\mathbf{x} \in \mathbb{A}^n : \mathbf{x}_{[n]\setminus i} = \mathbf{z}\},$$

which has diameter $1/n$ in the normalized Hamming metric.

Lemma 6.4. — For the choice of ρ_1 and ρ_2 in (50), we have

$$\int \langle \rho_1 \rangle_{i,\mathbf{z}} \cdot \mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_1} \parallel \theta_{i,\mathbf{z}}) \mu_{[n]\setminus i}(\mathrm{d}\mathbf{z}) \leq \frac{t^2}{n^2} \langle \rho_1 \rangle.$$

Proof. — The function tf is t -Lipschitz, so for any i and \mathbf{z} we have

$$\max \{ \rho_1(\mathbf{y}) : \mathbf{y} \in S_{i,\mathbf{z}} \} \leq e^{t/n} \min \{ \rho_1(\mathbf{y}) : \mathbf{y} \in S_{i,\mathbf{z}} \}.$$

Therefore Lemma 5.12 gives

$$\mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_1} \parallel \theta_{i,\mathbf{z}}) \leq t^2/n^2.$$

Substituting into the desired integral, we obtain

$$\int \langle \rho_1 \rangle_{i,\mathbf{z}} \cdot \mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_1} \parallel \theta_{i,\mathbf{z}}) \mu_{[n]\setminus i}(\mathrm{d}\mathbf{z}) \leq \frac{t^2}{n^2} \int \langle \rho_1 \rangle_{i,\mathbf{z}} \mu_{[n]\setminus i}(\mathrm{d}\mathbf{z}),$$

and this right-hand integral is equal to $\langle \rho_1 \rangle$ by (47). □

Lemma 6.5. — For the choice of ρ_1 and ρ_2 in (50), we have

$$\int \langle \rho_2 \rangle_{i,\mathbf{z}} \cdot \mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_2} \parallel \theta_{i,\mathbf{z}}) \mu_{[n]\setminus i}(\mathrm{d}\mathbf{z}) \leq 32 \frac{t^2}{n^2} \langle \rho_1 \rangle.$$

Beware that the average on the right here is $\langle \rho_1 \rangle$, not $\langle \rho_2 \rangle$. Also, the factor of 32 is chosen to be simple, not optimal.

Proof. — We certainly have $\langle \rho_2 \rangle_{i,\mathbf{z}} \leq 1$ for all i and \mathbf{z} , so it suffices to show that

$$\int \mathbf{D}((\theta_{i,\mathbf{z}})_{|\rho_2} \parallel \theta_{i,\mathbf{z}}) \mu_{[n]\setminus i}(\mathrm{d}\mathbf{z}) \leq 32 \frac{t^2}{n^2} \langle \rho_1 \rangle.$$

This is the work of the rest of the proof. It turns out to be slightly easier to work with an integral over all of A^n , so let us re-write the last integral as

$$\int D((\theta_{i,\mathbf{x}_{[n]\setminus i}})_{|\rho_2} \parallel \theta_{i,\mathbf{x}_{[n]\setminus i}}) \mu(d\mathbf{x}).$$

Now suppose that $\mathbf{x} \in A^n$ and that $\mathbf{y}, \mathbf{y}' \in S_{i,\mathbf{x}_{[n]\setminus i}}$. Then \mathbf{x}, \mathbf{y} and \mathbf{y}' agree in all but possibly the i^{th} coordinate. Since f is 1-Lipschitz, it follows that

$$\rho_2(\mathbf{y}') - \rho_2(\mathbf{y}) = \frac{1}{2}(e^{-f(\mathbf{y})} - e^{-f(\mathbf{y}')}) \leq \frac{1}{2}e^{-f(\mathbf{x})}(e^{t/n} - e^{-t/n}).$$

Since $0 < t/n \leq 1/200$, we have $e^{t/n} - e^{-t/n} < 4t/n$. Also, we have arranged that ρ_2 is always at least $1/2$. Therefore the last estimate implies that

$$\rho_2(\mathbf{y}') \leq \rho_2(\mathbf{y}) + \frac{2t}{n}e^{-f(\mathbf{x})} \leq \rho_2(\mathbf{y}) \left(1 + \frac{4t}{n}e^{-f(\mathbf{x})}\right) \leq \rho_2(\mathbf{y}) \exp\left(\frac{4t}{n}e^{-f(\mathbf{x})}\right)$$

for all $\mathbf{y}, \mathbf{y}' \in S_{i,\mathbf{x}_{[n]\setminus i}}$. Since $\theta_{i,\mathbf{x}_{[n]\setminus i}}$ is supported on $S_{i,\mathbf{x}_{[n]\setminus i}}$, we may combine this estimate with Lemma 5.12 to obtain

$$D((\theta_{i,\mathbf{x}_{[n]\setminus i}})_{|\rho_2} \parallel \theta_{i,\mathbf{x}_{[n]\setminus i}}) \leq \frac{16t^2}{n^2}e^{-2f(\mathbf{x})} \leq \frac{16t^2}{n^2}e^{-f(\mathbf{x})} = 32\frac{t^2}{n^2}\rho_1(\mathbf{x}).$$

For the second inequality here, notice that we have bounded $e^{-2f(\mathbf{x})}$ by $e^{-f(\mathbf{x})}$. This is extremely crude, but it suffices for the present argument.

Substituting into the desired integral, we obtain

$$\int D((\theta_{i,\mathbf{x}_{[n]\setminus i}})_{|\rho_2} \parallel \theta_{i,\mathbf{x}_{[n]\setminus i}}) \mu(d\mathbf{x}) \leq 32\frac{t^2}{n^2} \int \rho_1(\mathbf{x}) \mu(d\mathbf{x}). \quad \square$$

Combining the preceding lemmas with (45), (46) and (49), we have shown the following.

Corollary 6.6. — For the choice of (ρ_1, ρ_2) in (50) we have

$$\begin{aligned} & \text{DTC}(\mu) - \langle \rho_1 \rangle \cdot \text{DTC}(\mu_{|\rho_1}) - \langle \rho_2 \rangle \cdot \text{DTC}(\mu_{|\rho_2}) \\ & \geq \text{I}_\mu(\rho_1, \rho_2) - n \left(\frac{t^2}{n^2} \langle \rho_1 \rangle + 32 \frac{t^2}{n^2} \langle \rho_1 \rangle \right) \\ & = \langle \rho_1 \rangle \cdot D(\mu_{|\rho_1} \parallel \mu) + \langle \rho_2 \rangle \cdot D(\mu_{|\rho_2} \parallel \mu) - 33 \frac{t^2}{n} \langle \rho_1 \rangle. \end{aligned}$$

Proof of Proposition 6.2. — Since d_n has diameter at most 1, any measure on (A^n, d_n) satisfies $\text{T}(\kappa, 1)$ for all $\kappa > 0$. We may therefore assume that $r < 1$. If μ does not satisfy $\text{T}(rn/200, r)$, then by Proposition 5.13 it also does not satisfy $\text{L}([r/2, rn/200], 100/n)$.

This means there are a 1-Lipschitz function $f : A^n \rightarrow \mathbf{R}$ and a value $t \in [r/2, rn/200] \subseteq [0, n/200]$ such that

$$D(\mu_{|e^f} \parallel \mu) > 100 \frac{t^2}{n}.$$

Replacing f with $-f$ and then adding a constant if necessary, we may assume that f takes values in $[0, 1]$ and satisfies

$$D(\mu_{|e^{-f}} \parallel \mu) = D(\mu_{|(e^{-f}/2)} \parallel \mu) > 100 \frac{t^2}{n}.$$

Now construct ρ_1 and ρ_2 from this function f as in (50). Combining the above lower bound on $D(\mu_{|(e^{-f}/2)} \parallel \mu) = D(\mu_{|\rho_1} \parallel \mu)$ with Corollary 6.6, we obtain

$$\begin{aligned} & \text{DTC}(\mu) - \langle \rho_1 \rangle \cdot \text{DTC}(\mu_{|\rho_1}) - \langle \rho_2 \rangle \cdot \text{DTC}(\mu_{|\rho_2}) \\ & \geq \langle \rho_1 \rangle \cdot D(\mu_{|\rho_1} \parallel \mu) + \langle \rho_2 \rangle \cdot D(\mu_{|\rho_2} \parallel \mu) - 33 \frac{t^2}{n} \langle \rho_1 \rangle \\ & > \langle \rho_1 \rangle \cdot D(\mu_{|\rho_1} \parallel \mu) + \langle \rho_2 \rangle \cdot D(\mu_{|\rho_2} \parallel \mu) - \frac{1}{2} \langle \rho_1 \rangle \cdot D(\mu_{|\rho_1} \parallel \mu) \\ & \geq \frac{1}{2} \left(\langle \rho_1 \rangle \cdot D(\mu_{|\rho_1} \parallel \mu) + \langle \rho_2 \rangle \cdot D(\mu_{|\rho_2} \parallel \mu) \right) \\ & = \frac{1}{2} \mathbf{I}_\mu(\rho_1, \rho_2). \end{aligned}$$

The last line here follows from another appeal to Corollary 2.3.

Finally, since $f \leq 1$ and $r/2 \leq t \leq rn/200 \leq n$, we also obtain

$$\begin{aligned} \mathbf{I}_\mu(\rho_1, \rho_2) & \geq \langle \rho_1 \rangle \cdot D(\mu_{|\rho_1} \parallel \mu) \geq \int \frac{1}{2} e^{-f} d\mu \cdot \left(100 \frac{r^2}{4n} \right) \\ & \geq \frac{r^2 e^{-t}}{n} \geq r^2 n^{-1} e^{-n}. \end{aligned} \quad \square$$

Remark. — The proof of Proposition 6.2 actually exploits the failure of μ to satisfy the stronger L-inequality $L([r/2, rn/200], 100/n)$, rather than the T-inequality in the statement of that proposition. Our proof therefore gives the conclusion of Theorem 6.1 with that L-inequality in place of the T-inequality for the measures $\mu_{|\rho_j}$, $2 \leq j \leq k$. However, later steps in the proof of Theorem B use some properties that we know only for T-inequalities, such as those from Subsection 5.3. We do not refer to L-inequalities again after the present section.

6.3. *The use of DTC in this section.* — Both TC and DTC are notions of multi-variate mutual information for a measure μ on A^n . In searching for a ‘decrement’ proof of Theorem B, it is natural that we try such a quantity. This is because we regard product measures, which exhibit very strong concentration, as an extreme case, and they are precisely the measures for which any good notion of multi-variate mutual information should be zero.

It is more subtle to describe why DTC, rather than TC or any other kind of multi-variate mutual information, is the right quantity for the decrement. A valuable hint in this direction comes from inspecting the terms that define DTC and comparing them with older proofs of logarithmic Sobolev inequalities on Hamming cubes. Let us discuss this with reference to the L-inequalities of Subsection 5.4, which are really a special case of logarithmic Sobolev inequalities.

Let μ be the uniform distribution on $\{0, 1\}^n$, and let $\xi_i : \{0, 1\}^n \rightarrow \{0, 1\}$ be the i^{th} coordinate projection for $1 \leq i \leq n$. Following [52, Section 4] (where the argument is credited to Bobkov), one obtains an L-inequality for the space $(\{0, 1\}^n, d_n)$ and measure μ from the following pair of estimates:

- (a) For any other $\nu \in \text{Prob}(\{0, 1\}^n)$, we have

$$D(\nu \parallel \mu) \leq \sum_{i=1}^n D(\nu \parallel \mu \mid \xi_{[n] \setminus i}),$$

where $D(\nu \parallel \mu \mid \xi_{[n] \setminus i})$ denotes a conditional KL divergence (see [11, Section 2.5]; the summands on the right in [52, Proposition 4.1] have this form).

- (b) If $\nu = \mu_{|e^{-f}}$ for some 1-Lipschitz function f on $\{0, 1\}^n$, then

$$D(\nu \parallel \mu \mid \xi_{[n] \setminus i}) \leq \frac{t^2}{n^2} \quad \text{for each } i.$$

Combining (a) and (b), it follows that $D(\mu_{|e^{-f}} \parallel \mu) \leq t^2/n$ for any 1-Lipschitz function f .

Now consider instead an arbitrary measure μ on $\{0, 1\}^n$, and suppose that $\rho_1 = e^{-f}$ and $\rho_2 = 1 - \rho_1$ for a 1-Lipschitz function f and some t (ignoring the technical factor of $1/2$ in (50)). Of the two estimates above, (b) still holds simply by Lemma 5.12, but (a) is often false unless μ is a product measure.

Observe that the quantity $D(\mu_{|e^{-f}} \parallel \mu)$ appears in the first term of (46), the formula for $I(\xi; \zeta)$ in the previous subsection. Likewise, a simple re-write of the quantity $D(\mu_{|e^{-f}} \parallel \mu \mid \xi_{[n] \setminus i} = \mathbf{z})$ shows that it is equal to $D((\theta_{i, \mathbf{z}})_{|\rho_1} \parallel \theta_{i, \mathbf{z}})$, again using the notation of the previous subsection. This quantity appears in the first term of (48), the formula for $I(\xi_i; \zeta \mid \xi_{[n] \setminus i} = \mathbf{z})$.

So for a general measure μ the difference

$$(51) \quad D(\mu_{|e^{-f}} \parallel \mu) - \sum_{i=1}^n D(\mu_{|e^{-f}} \parallel \mu \mid \xi_{[n] \setminus i})$$

resembles part of the difference

$$(52) \quad \mathbf{I}(\xi; \zeta) - \sum_{i=1}^n \mathbf{I}(\xi_i; \zeta \mid \xi_{[n] \setminus i}) = \text{DTC}(\mu) - \langle \rho_1 \rangle \cdot \text{DTC}(\mu_{|\rho_1}) - \langle \rho_2 \rangle \cdot \text{DTC}(\mu_{|\rho_2}).$$

The difference in (51) does not quite appear in (52) because the coefficients of the various KL divergences here do not match. But the resemblance is enough to suggest the following. If a suitable L-inequality fails, so one can find a function f and parameter t for which $\mathbf{D}(\mu_{|e^{-tf}} \parallel \mu)$ is large, then the difference in (51) must also be large, and then one might hope to show that the DTC-decrement in (52) is also large. The work of Subsection 6.2 consists of the tweaks and technicalities that are needed to turn this hope into rigorous estimates.

Ledoux uses the logarithmic Sobolev inequalities in [52, Section 4] to give a new proof of an exponential moment bound for Lipschitz functions on product spaces. This bound is in turn equivalent to Marton’s transportation inequalities (Proposition 5.2) via the Bobkov–Götze equivalence (see the discussion after Proposition 5.6). Marton’s original proof of Proposition 5.2 is quite different, and arguably more elementary. She uses an induction on the dimension n , and no functional inequalities such as logarithmic Sobolev inequalities appear. As remarked following Proposition 5.6, yet another proof can be given using the Bobkov–Götze equivalence and the method of bounded martingale differences. This last proof does involve bounding an exponential moment, but it still seems more elementary than the logarithmic Sobolev approach.

Question 6.7. — *Is there a proof of Theorem 6.1, or more generally of Theorem B, which uses an induction on n and either (i) some variants of Marton’s ideas, or (ii) the Bobkov–Götze equivalence and some more elementary way of controlling exponential moments of Lipschitz functions?*

I expect that such an alternative proof would offer valuable additional insight into the phenomena behind Theorem B.

7. Completed proof of Theorem B

We still have to turn Theorem 6.1 into Theorem B. This is the work of the present section, which has two stages. The first, in Subsection 7.1, continues to work solely with DTC. The second, in Subsection 7.2, takes us from DTC back to TC.

7.1. Another auxiliary decomposition.

Theorem 7.1. — *For any $\varepsilon, r > 0$ there exists $c > 0$ such that the following holds. Any $\mu \in \text{Prob}(A^n)$ can be written as a mixture*

$$\mu = p_1 \mu_1 + \cdots + p_m \mu_m$$

so that

- (a) $m \leq c \exp(c \cdot \text{DTC}(\mu))$,
- (b) $p_1 < \varepsilon$, and
- (c) the measure μ_j satisfies $T(rn/600, r)$ for every $j = 2, 3, \dots, m$.

Letting $\rho_j := d(p_j \mu_j) / d\mu$, we always have

$$I_\mu(\rho_1, \dots, \rho_m) \leq H(p_1, \dots, p_m) \leq \log m.$$

Therefore Theorem 7.1 is a strengthening of Theorem 6.1. We deduce Theorem 7.1 from Theorem 6.1 by using a simple sampling argument to ‘coarsen’ the representation of μ given by Theorem 6.1 and allowing a slight degradation in the error estimates.

Proposition 7.2 (Sampling from a low-information mixture). — Let (\mathbf{X}, μ) be a standard probability space, and let μ be written as a mixture $\int \mu_\bullet \, d\mathbf{P}$ using some other probability space (Ω, \mathbf{P}) and a kernel μ_\bullet from Ω to \mathbf{X} . Let $\varepsilon \in (0, 1/2)$, let $\Omega_1 \subseteq \Omega$ be measurable with $\mathbf{P}(\Omega_1) > 1 - \varepsilon/2$, and assume that

$$I := \int D(\mu_\omega \parallel \mu) \mathbf{P}(d\omega) < \infty.$$

Finally, let $m := \lceil 16\varepsilon^{-2} e^{16(I+1)/\varepsilon} \rceil$. Then there are elements $\omega_1, \dots, \omega_m \in \Omega_1$, not necessarily distinct, such that

$$\left\| \frac{1}{m} \sum_{j=1}^m \mu_{\omega_j} - \mu \right\| < 3\varepsilon.$$

This is proved using the probabilistic method, together with a simple truncation argument. This combination is reminiscent of classical proofs of the weak law of large numbers for random variables without bounded second moments [19, Section X.2].

Proof. — *Step 1: setup and truncation.* Let $\mathbf{Q} := \mathbf{P} \times \mu_\bullet$, the hookup introduced in Subsection 2.2, and let F be the Radon–Nikodym derivative $d\mathbf{Q}/d(\mathbf{P} \times \mu)$. Then we have

$$\mu_\omega = F(\omega, \cdot) \cdot \mu \quad \text{for P-a.e. } \omega$$

and

$$I = \int F \log F \, d(\mathbf{P} \times \mu).$$

The function $t \log t$ on $[0, \infty)$ has a global minimum at $t = e^{-1}$ and its value there is $-e^{-1}$. Therefore

$$(53) \quad \int F \log^+ F \, d(\mathbf{P} \times \mu) \leq I + e^{-1} < I + 1.$$

Now define

$$F' := \min\{F, e^{8(I+1)/\varepsilon}\} \quad \text{and} \quad \mu'_\omega := F'(\omega, \cdot) \cdot \mu \quad \text{for each } \omega.$$

Each μ'_ω is a positive measure bounded by μ_ω , so Markov's inequality and (53) give

$$\begin{aligned} \int \|\mu_\omega - \mu'_\omega\| \mathbf{P}(d\omega) &= \int (F - F') \, d(\mathbf{P} \times \mu) \\ &\leq \int_{\{F > e^{8(I+1)/\varepsilon}\}} F \, d(\mathbf{P} \times \mu) \\ &\leq \frac{\varepsilon}{8(I+1)} \int F \log^+ F \, d(\mathbf{P} \times \mu) \\ &< \varepsilon/8. \end{aligned}$$

Define

$$\mu' := \int \mu'_\omega \mathbf{P}(d\omega) \quad \text{and} \quad f'(x) := \int F'(\omega, x) \mathbf{P}(d\omega).$$

Then μ' is a positive measure bounded by μ which satisfies $\|\mu - \mu'\| \leq \varepsilon/8$, and f' is a version of the Radon–Nikodym derivative $d\mu'/d\mu$.

Step 2: the probabilistic method. Let Y_1, \dots, Y_m be i.i.d. random elements of Ω , each with distribution \mathbf{P} , and let \mathbf{P} denote the underlying probability measure for those random elements. Let \mathbf{E} and Var denote expectation and variance with respect to \mathbf{P} .

For each $x \in \mathbf{X}$, consider the random variable

$$\mathbf{X}_x := \frac{1}{m} \sum_{j=1}^m F'(Y_j, x).$$

This is an average of i.i.d. random variables. Each of them satisfies

$$\mathbf{E}F'(Y_j, x) = \int F'(\omega, x) \mathbf{P}(d\omega) = f'(x).$$

Also, each is bounded by $e^{8(I+1)/\varepsilon}$, and hence satisfies $\text{Var}(F'(Y_j, x)) \leq e^{16(I+1)/\varepsilon}$.

Now consider the empirical average of the measures μ'_{Y_j} : it is given by

$$\frac{1}{m} \sum_{j=1}^m \mu'_{Y_j} = \frac{1}{m} \sum_{j=1}^m (F'(Y_j, \cdot) \cdot \mu) = \left(\frac{1}{m} \sum_{j=1}^m F'(Y_j, \cdot) \right) \cdot \mu.$$

Therefore, by Fubini's theorem,

$$\begin{aligned}
\mathbf{E} \left\| \frac{1}{m} \sum_{j=1}^m \mu'_{Y_j} - \mu' \right\| &= \mathbf{E} \int \left| \frac{1}{m} \sum_{j=1}^m F'(Y_j, \cdot) - f' \right| d\mu \\
&= \int \mathbf{E} |X_x - \mathbf{E}X_x| \mu(dx) \\
&\leq \int \sqrt{\text{Var}(X_x)} \mu(dx) \\
&\leq m^{-1/2} e^{8(I+1)/\varepsilon}.
\end{aligned}$$

Combining the bounds above, we obtain

$$\begin{aligned}
&\mathbf{E} \left\| \frac{1}{m} \sum_{j=1}^m \mu_{Y_j} - \mu \right\| \\
&\leq \frac{1}{m} \sum_{j=1}^m \mathbf{E} \|\mu_{Y_j} - \mu'_{Y_j}\| + \mathbf{E} \left\| \frac{1}{m} \sum_{j=1}^m \mu'_{Y_j} - \mu' \right\| + \|\mu' - \mu\| \\
&= \int \|\mu_\omega - \mu'_\omega\| \mathbf{P}(d\omega) + \mathbf{E} \left\| \frac{1}{m} \sum_{j=1}^m \mu'_{Y_j} - \mu' \right\| + \|\mu' - \mu\| \\
&< \varepsilon/4 + m^{-1/2} e^{8(I+1)/\varepsilon} \leq \varepsilon/2.
\end{aligned}$$

Therefore Markov's inequality gives

$$\mathbf{P} \left\{ \left\| \frac{1}{m} \sum_{j=1}^m \mu_{Y_j} - \mu \right\| < \varepsilon \right\} > \frac{1}{2}.$$

On the other hand, we have

$$\mathbf{E} |\{j \in [m] : Y_j \in \Omega \setminus \Omega_1\}| = m\mathbf{P}(\Omega \setminus \Omega_1) < m\varepsilon/2,$$

so another appeal to Markov's inequality gives

$$\mathbf{P} \{ |\{j \in [m] : Y_j \in \Omega_1\}| > (1 - \varepsilon)m \} > \frac{1}{2}.$$

Combining this probability lower bound with the previous one, it follows that the intersection of these events has positive probability. Therefore some possible values $\omega_1, \dots, \omega_m$ of Y_1, \dots, Y_m satisfy both

$$\left\| \mu - \frac{1}{m} \sum_{j=1}^m \mu_{\omega_j} \right\| < \varepsilon$$

and

$$|\{j \in [m] : \omega_j \in \Omega_1\}| > (1 - \varepsilon)m.$$

To complete the proof, we simply discard any values ω_j that lie in $\Omega \setminus \Omega_1$ and replace them with arbitrary members of Ω_1 . This incurs an additional error of at most 2ε in the total-variation approximation to μ . \square

Proof of Theorem 7.1. — By shrinking ε if necessary, we may assume that

$$(54) \quad \varepsilon < \frac{1}{6} \quad \text{and also} \quad 400 \log((1 - 3\varepsilon/2)^{-1}) < r^2.$$

Now apply Theorem 6.1 with $\varepsilon^2/2$ in place of ε and with the present value of r . Let (ρ_1, \dots, ρ_k) be the resulting fuzzy partition. Let $\Omega := [k]$, and let \mathbf{P} be the probability distribution on this set defined by $\mathbf{P}(j) := p_j := \langle \rho_j \rangle$. The formula from Corollary 2.3 and conclusion (a) of Theorem 6.1 give

$$\mathbf{I} := \mathbf{I}_\mu(\rho_1, \dots, \rho_k) = \sum_{j \in \Omega} p_j \cdot \mathbf{D}(\mu_{|\rho_j} \| \mu) \leq 2 \cdot \mathbf{DTC}(\mu).$$

Also, let $\Omega_1 = \{2, 3, \dots, k\}$, so conclusion (b) of Theorem 6.1 gives $\mathbf{P}(\Omega_1) > 1 - \varepsilon^2/2$.

We now apply Proposition 7.2 to μ and its representation as a mixture of the measures $\mu_{|\rho_j}$, which we abbreviate to μ_j . We apply that proposition with ε^2 in place of ε . It provides a constant c which depends on ε (and hence also on r , because of (54)), an integer $m \leq ce^d$, and elements $i_1, \dots, i_m \in \Omega_1$ such that

$$(55) \quad \|\mu - \mu'\| < 3\varepsilon^2 \quad \text{where} \quad \mu' := \frac{1}{m} \sum_{j=1}^m \mu_{i_j}.$$

Using the Jordan decomposition of $\mu - \mu'$, we may now write

$$\mu = \gamma + \nu \quad \text{and} \quad \mu' = \gamma + \nu'$$

for some measures γ , ν and ν' such that ν and ν' are mutually singular. These measures satisfy

$$\|\nu\| = \|\nu'\| = \frac{1}{2} \|\mu - \mu'\| < 3\varepsilon^2/2 \quad \text{and hence} \quad \|\gamma\| > 1 - 3\varepsilon^2/2.$$

Let f be the Radon–Nikodym derivative $d\gamma/d\mu'$. Then $0 \leq f \leq 1$, and we have

$$\frac{1}{m} \sum_{j=1}^m \int f d\mu_{i_j} = \int f d\mu' = \|\gamma\| > 1 - 3\varepsilon^2/2.$$

Therefore, by Markov's inequality, the set \mathbf{J} of all $j \in [m]$ which satisfy

$$(56) \quad \int f \, d\mu_j > 1 - 3\varepsilon/2$$

has cardinality at least $(1 - \varepsilon)m$.

If $j \in \mathbf{J}$, then (56) implies that

$$\frac{d(\mu_j)_f}{d\mu_j} = \frac{f}{\int f \, d\mu_j} < (1 - 3\varepsilon/2)^{-1}.$$

For these j , Lemma 5.7 gives that the measure $(\mu_j)_f$ still satisfies

$$\mathsf{T}(rn/200, 400 \log((1 - 3\varepsilon/2)^{-1})/rn + 2r).$$

By the second upper bound in (54), this implies $\mathsf{T}(rn/200, 3r)$.

So now we can write

$$(57) \quad \begin{aligned} \mu &= \nu + \gamma \\ &= \nu + f \cdot \mu' \\ &= \nu + \frac{1}{m} \sum_{j \in [m] \setminus \mathbf{J}} f \cdot \mu_j + \frac{1}{m} \sum_{j \in \mathbf{J}} f \cdot \mu_j. \end{aligned}$$

In this last sum, the first few terms satisfy

$$\left\| \nu + \frac{1}{m} \sum_{j \in [m] \setminus \mathbf{J}} f \cdot \mu_j \right\| < 3\varepsilon/2 + \frac{m - |\mathbf{J}|}{m} < 4\varepsilon.$$

On the other hand, the remainder

$$\frac{1}{m} \sum_{j \in \mathbf{J}} f \cdot \mu_j$$

is a non-negative linear combination of probability measures that all satisfy $\mathsf{T}(rn/200, 3r)$. So now let us combine the first few terms in (57) into a single 'bad' term. This gives a mixture of at most m measures which has the desired properties, except that the bad term has total variation bounded by 4ε rather than ε , and the parameter r has been replaced by $3r$ throughout. Since $\varepsilon > 0$ and $r > 0$ are both arbitrary, this completes the proof. \square

Our route to Theorem 7.1 is quite indirect. This is because the proof of Theorem 6.1 relates the decrement in the DTC to a change in the mutual information of the fuzzy partition $(\rho_j)_j$, not a change in the entropy of the probability vector $(\langle \rho_j \rangle)_j$. As a result, Theorem 6.1 might give a representation of μ as a mixture with far too many summands, and we must then go back and find a more efficient representation by sampling as in Proposition 7.2.

Question 7.3. — Can one give a more direct proof of Theorem 7.1 by improving some of the estimates in the proof of Theorem 6.1?

7.2. *Completed proof of Theorem B.* — Let μ be as in the statement of Theorem B. To deduce Theorem B from Theorem 7.1, the last step is to combine that theorem with Lemmas 3.7 and 5.10.

Proof of Theorem B. — As remarked previously, (A^n, d_n, μ) must satisfy $T(\kappa, 1)$ for all $\kappa > 0$, so we may assume that $r < 1$.

In the statement of Theorem B we have $E = \text{TC}(\mu)$. Let $S \subseteq [n]$ be the subset provided by Lemma 3.7, so $|S| \geq (1 - r)n$ and $\text{DTC}(\mu_S) \leq E/r$. Let c_1 be the constant given by Theorem 7.1 for the current values of ε and r . Applying Theorem 7.1 to μ_S , we can write it as a mixture

$$\mu_S = \rho_1 \cdot \mu_S + \cdots + \rho_m \cdot \mu_S$$

for some fuzzy partition $(\rho_j)_{j=1}^m$ on A^S such that

- (a) $m \leq c_1 \exp(c_1 \cdot \text{DTC}(\mu_S)) \leq ce^{cE}$ where $c := \max\{c_1/r, 1\}$,
- (b) $\int \rho_1 d\mu_S < \varepsilon$, and
- (c) the measure $(\mu_S)_{|\rho_j}$ is defined and satisfies $T(r|S|/600, r)$ for every $j = 2, 3, \dots, m$.

Let $\rho'_j(\mathbf{x}) = \rho_j(\mathbf{x}_S)$ for every $\mathbf{x} \in A^n$. Then the resulting mixture

$$\mu = \rho'_1 \cdot \mu + \cdots + \rho'_m \cdot \mu$$

still satisfies the analogues of properties (a) and (b) above. Finally, property (c) above combines with Lemma 5.10 to give that $\mu_{|\rho'_j}$ satisfies

$$T\left(\frac{n}{|S|} \cdot \frac{r|S|}{600}, \frac{|S|}{n}r + r\right)$$

for every $j = 2, 3, \dots, m$. This inequality implies $T(rn/600, 2r)$, and they coincide in case $S = [n]$. Since $r > 0$ is arbitrary, this completes the proof. \square

7.3. *Aside: another possible connection.* — Theorem 7.1 is worth comparing with recent results of Ellis, Friedgut, Kindler and Yehudayoff in [18]. In our terminology, they prove that if $\text{DTC}(\mu)$ is very small, then μ must itself be close to a product measure in a rather strong sense (certainly strong enough to imply a good concentration inequality). However, they prove this only when $\text{DTC}(\mu)$ is bounded by a fixed and sufficiently small tolerance ε , independently of the dimension n . By contrast, in Part III we need to apply Theorem B when $\text{DTC}(\mu)$ is of order n . This is far outside the domain covered by the results of [18].

It would be interesting to understand whether our present work and the proofs in [18] have some underlying structure in common.

8. Proof of Theorem C

To prove Theorem C, we carve out the desired partition of A^n one set at a time.

Proposition 8.1. — For any $r > 0$ there exist $c, \kappa > 0$ such that, for any alphabet A , the following holds for all sufficiently large n . If $\mu \in \text{Prob}(A^n)$, then there is a subset $V \subseteq A^n$ such that

$$\mu(V) \geq \exp(-c \cdot \text{TC}(\mu))$$

and such that $\mu|_V$ satisfies $T(\kappa n, r)$.

Proof of Theorem C from Proposition 8.1. — Consider $\varepsilon, r > 0$ as in the statement of Theorem C. Clearly we may assume that $\varepsilon < 1$. For this value of r and for the alphabet A , let n be large enough to apply Proposition 8.1. Let c_1 and κ be the new constants given by that proposition.

We now construct a finite disjoint sequence of subsets V_1, V_2, \dots, V_{m-1} of A^n by the following recursion.

To start, let V_1 be a subset such that

$$\mu(V_1) \geq \exp(-c_1 \cdot \text{TC}(\mu))$$

and such that $\mu|_{V_1}$ satisfies $T(\kappa n, r)$, as provided by Proposition 8.1.

Now suppose we have already constructed V_1, \dots, V_ℓ for some $\ell \geq 1$, and let $W := A^n \setminus (V_1 \cup \dots \cup V_\ell)$. If $\mu(W) < \varepsilon$, then stop the recursion and set $m := \ell + 1$. Otherwise, let $\mu' := \mu|_W$, and apply Proposition 8.1 again to this new measure. By Corollary 3.2 we have

$$\text{TC}(\mu') \leq \frac{1}{\mu(W)} (\text{TC}(\mu) + \log 2) \leq \varepsilon^{-1} (\text{TC}(\mu) + \log 2),$$

so Proposition 8.1 gives a new subset $V_{\ell+1}$ of A^n such that

$$\begin{aligned} \mu(V_{\ell+1}) &\geq \mu(W) \cdot \mu'(V_{\ell+1}) \geq \varepsilon \exp(-c_1 \cdot \text{TC}(\mu')) \\ &\geq \varepsilon 2^{-c_1/\varepsilon} \exp(-c_1 \cdot \text{TC}(\mu)/\varepsilon) \end{aligned}$$

and such that $\mu'|_{V_{\ell+1}}$ satisfies $T(\kappa n, r)$. Since μ' is supported on W , we may intersect $V_{\ell+1}$ with W without disrupting either of these conclusions, and so assume that $V_{\ell+1} \subseteq W$. Having done so, we have $\mu'|_{V_{\ell+1}} = \mu|_{V_{\ell+1}}$. This continues the recursion.

The sets V_j are pairwise disjoint, and they all have measure at least

$$\varepsilon 2^{-c_1/\varepsilon} \exp(-c_1 \cdot \text{TC}(\mu)/\varepsilon).$$

Therefore the recursion must terminate at some finite value ℓ which satisfies

$$\begin{aligned} m = \ell + 1 &\leq \varepsilon^{-1} 2^{c_1/\varepsilon} \exp(c_1 \cdot \text{TC}(\mu)/\varepsilon) + 1 \\ &\leq (\varepsilon^{-1} 2^{c_1/\varepsilon} + 1) \exp(c_1 \cdot \text{TC}(\mu)/\varepsilon). \end{aligned}$$

Once it has terminated, let $U_j := V_{j-1}$ for $j = 2, 3, \dots, m$, and let U_1 be the complement of all these sets.

If we set $c := \max\{(2^{-c_1/\varepsilon}/\varepsilon + 1), c_1/\varepsilon\}$, then this partition U_1, \dots, U_m has all three of the desired properties. \square

It remains to prove Proposition 8.1. Fix $r > 0$ for the rest of the section, and consider $\mu \in \text{Prob}(A^n)$. Let $\kappa := r/1200$, and let c_B be the constant provided by Theorem B with the input parameters r and $\varepsilon := 1/2$. Clearly we may assume that $c_B \geq 1$ without disrupting the conclusions of Theorem B. We prove Proposition 8.1 with this value for κ but with $9r$ in place of r and with a new constant c derived from r and c_B . Since $r > 0$ was arbitrary this still completes the proof. Since any measure on A^n satisfies $T(\kappa n, 9r)$ if $r \geq 1/9$, we may also assume that $r < 1/9$.

To lighten notation, let us define

$$(58) \quad \delta := \min\{r^2/42, 1/18\}.$$

We make frequent references to this small auxiliary quantity below. Its definition results in a correct dependence on r for certain estimates near the end of the proof.

At a few points in this section it is necessary that n be sufficiently large in terms of r and $|A|$. These points are explained as they arise.

There are three different cases in the proof of Proposition 8.1, in which the desired set V exists for different reasons. Two of those cases are very simple. All of the real work goes into the third case, including an appeal to Theorem B. The analysis of that third case could be applied to any measure μ on A^n , but I do not see how to control the necessary estimates unless we assume the negations of the first two cases — this is why we separate the proof into cases at all.

The first and simplest case is when μ has a single atom of measure at least

$$(59) \quad \exp\left(-\frac{161c_B}{\delta^2} \cdot \text{TC}(\mu)\right).$$

Then we just take V to be that singleton. So in the rest of our analysis we may assume that all atoms of μ weigh less than this. Of course, the constant that appears in front of $\text{TC}(\mu)$ in (59) is chosen to meet our needs when we come to apply this assumption later.

The next case is when $\text{TC}(\mu)$ is sufficiently small. This is dispatched by the following lemma.

Lemma 8.2. — *If $\text{TC}(\mu) \leq r^4 n$, and if n is sufficiently large in terms of r , then there is a subset V of A^n satisfying $\mu(V) \geq 1/2$ such that $\mu|_V$ satisfies $T(8rn, 9r)$.*

Proof. — Let $\mu_{\{i\}}$ be the i^{th} marginal of μ for $i = 1, 2, \dots, n$. Recall from (15) that

$$\text{TC}(\mu) = \text{D}(\mu \parallel \mu_{\{1\}} \times \cdots \times \mu_{\{n\}}).$$

Therefore Marton's inequality from the first part of Proposition 5.2 gives

$$\bar{d}_n(\mu, \mu_{\{1\}} \times \cdots \times \mu_{\{n\}}) \leq \sqrt{\frac{\text{TC}(\mu)}{2n}} \leq r^2,$$

where in the end we ignore a factor of $\sqrt{1/2}$. Since our assumptions imply that $r < 1/8$, we may now apply the second part of Proposition 5.2 and then Proposition 5.9. These give a subset V of A^n satisfying $\mu(V) \geq 1 - 4r \geq 1/2$ and such that $\mu|_V$ satisfies

$$\text{T}\left(8rn, \frac{8r + 2 \log 4}{8rn} + 4r + 4r\right).$$

If n is sufficiently large, this implies $\text{T}(8rn, 9r)$. □

So now we may assume that μ has no atoms of measure at least (59), and also that $\text{TC}(\mu)$ is greater than $r^4 n$. We keep these assumptions in place until the proof of Proposition 8.1 is completed near the end of this section.

This leaves the case in which we must use Theorem B. This argument is more complicated. It begins by deriving some useful structure for the measure μ from the two assumptions above. This is done in two steps which we formulate as Lemma 8.3 and Proposition 8.4. Lemma 8.3 can be seen as a quantitative form of the asymptotic equipartition property for a big piece of the measure μ .

Lemma 8.3. — *Consider a measure μ satisfying the assumptions above, and let $M := \lceil n \log |A| \rceil$. If n is sufficiently large, then there are a subset $P \subseteq A^n$ and a constant $h \geq 160c_B \text{TC}(\mu)/\delta^2$ such that $\mu(P) \geq 1/4M$,*

$$\text{TC}(\mu|_P) \leq 4\text{TC}(\mu),$$

and

$$e^{-h} \geq \mu|_P(\mathbf{x}) > e^{-h-1} \quad \forall \mathbf{x} \in P.$$

Proof. — Let $E := \text{TC}(\mu)$ and let $h_0 := 161c_B E/\delta^2$ (the exponent from (59)). Let

$$P_j := \{\mathbf{x} : e^{-h_0-j} \geq \mu(\mathbf{x}) > e^{-h_0-j-1}\} \quad \text{for } j = 0, 1, \dots, M-1$$

and let

$$P_M := \{\mathbf{x} : e^{-h_0-M} \geq \mu(\mathbf{x})\}.$$

Since, by assumption, all atoms of μ weigh less than e^{-h_0} , the sets P_0, P_1, \dots, P_M constitute a partition of A^n . Also, from its definition, the last of these sets must satisfy

$$(60) \quad \mu(P_M) \leq e^{-h_0-M} |P_M| \leq e^{-h_0} |A|^{-n} |A|^n = e^{-161c_B E/\delta^2}.$$

Since we are also assuming that $E > r^4 n$, this bound is less than $1/4$ provided n is large enough in terms of r .

Using this partition, we may write

$$\mu = \sum_{j=0}^M \mu(P_j) \cdot \mu|_{P_j}.$$

By Lemma 3.1, this leads to

$$(61) \quad \text{TC}(\mu) + H(\mu(P_0), \dots, \mu(P_M)) \geq \sum_{j=0}^M \mu(P_j) \cdot \text{TC}(\mu|_{P_j}).$$

The left-hand side here is at most

$$\text{TC}(\mu) + \log(M+1) \leq \text{TC}(\mu) + \log \log |A| + \log n + 2.$$

Since we are assuming that $\text{TC}(\mu) > r^4 n$, this upper bound is less than $2\text{TC}(\mu)$ provided n is sufficiently large in terms of r and $|A|$.

Therefore, provided n is sufficiently large, the right-hand side of (61) is at most $2\text{TC}(\mu)$. By Markov's inequality, it follows that more than half of μ is supported on cells P_j which satisfy

$$(62) \quad \text{TC}(\mu|_{P_j}) \leq 4\text{TC}(\mu).$$

By (60), it follows that more than a quarter of μ is supported on such cells P_j with $j \leq M-1$.

Now choose $j \leq M-1$ for which (62) holds and such that $\mu(P_j)$ is maximal, and let $P := P_j$. This maximal value must be at least $1/4M$. By the definition of P , each $\mathbf{x} \in P$ satisfies

$$e^{a-h_0-j} \geq \mu|_P(\mathbf{x}) > e^{a-h_0-j-1}, \quad \text{where } e^a := 1/\mu(P).$$

Finally, let $h := h_0 + j - a$, and observe that

$$h \geq h_0 - a \geq 161c_B E/\delta^2 - \log(4\lceil n \log |A| \rceil).$$

Since we are assuming that $E > r^4 n$, this is at least $160c_B E/\delta^2$ provided n is sufficiently large in terms of r and $|A|$. \square

Proposition 8.4. — Consider a measure μ as in the preceding lemma. Let \mathbf{P} be the set and h the constant obtained there. Let $m := \lceil (1 - \delta)n \rceil$. Then there are subsets

$$\mathbf{S} \subseteq [n], \quad \mathbf{R} \subseteq A^n \quad \text{and} \quad \mathbf{Q} \subseteq A^{\mathbf{S}}$$

such that

- (a) $|\mathbf{S}| = m$,
- (b) $\mu(\mathbf{R}) \geq \delta/32M$,
- (c) $\mathbf{x}_{\mathbf{S}} \in \mathbf{Q}$ for all $\mathbf{x} \in \mathbf{R}$,
- (d) every $\mathbf{z} \in \mathbf{Q}$ has the property that

$$\max_{\mathbf{x} \in \mathbf{R}: \mathbf{x}_{\mathbf{S}} = \mathbf{z}} \mu_{|\mathbf{R}}(\mathbf{x}) \leq e^{-\delta h/4} \mu_{|\mathbf{R}}\{\mathbf{x} : \mathbf{x}_{\mathbf{S}} = \mathbf{z}\},$$

and

- (e) $\text{TC}(\mu_{|\mathbf{R}}) \leq 33\text{TC}(\mu)/\delta$.

Proof. — Let $\mu' := \mu_{|\mathbf{P}}$, and let $\xi_i : A^n \rightarrow A$ be the coordinate projection for $i = 1, \dots, n$.

By re-ordering the coordinates of A^n if necessary, let us assume that

$$(63) \quad H_{\mu'}(\xi_1) \leq \dots \leq H_{\mu'}(\xi_n).$$

Having done so, let $\mathbf{S} := \{1, 2, \dots, m\}$. Let $\mu'_{[\mathbf{m}]}$ be the projection of μ' to $A^{\mathbf{m}}$.

From (63) and the subadditivity of entropy, we obtain

$$H_{\mu'}(\xi_{[\mathbf{m}]}) \leq \sum_{i=1}^m H_{\mu'}(\xi_i) \leq \frac{m}{n} \sum_{i=1}^n H_{\mu'}(\xi_i) = \frac{m}{n} (H(\mu') + \text{TC}(\mu')).$$

Provided n is sufficiently large in terms of r , this is at most

$$(1 - 3\delta/4)H(\mu') + (1 - 3\delta/4)\text{TC}(\mu') \leq (1 - 3\delta/4)H(\mu') + \text{TC}(\mu').$$

The first conclusion of Lemma 8.3 gives that $\text{TC}(\mu') \leq 4\text{TC}(\mu) = 4E$. On the other hand, the second conclusion gives that

$$(64) \quad h + 1 > H(\mu') \geq h,$$

and h is greater than $16E/\delta$ because of our earlier assumption that $c_{\mathbf{B}} \geq 1$. Combining these inequalities, we deduce that

$$(65) \quad \begin{aligned} H_{\mu'}(\xi_{[\mathbf{m}]}) &\leq (1 - 3\delta/4)H(\mu') + 4E < (1 - 3\delta/4)H(\mu') + \frac{\delta}{4}H(\mu') \\ &\leq (1 - \delta/2)H(\mu') < (1 - \delta/2)(h + 1). \end{aligned}$$

Let

$$\begin{aligned} \mathbf{Q} &:= \{\mathbf{z} \in A^m : \mu'_{[m]}(\mathbf{z}) \geq e^{-(1-\delta/4)h}\} \\ &= \{\mathbf{z} \in A^m : -\log \mu'_{[m]}(\mathbf{z}) \leq (1-\delta/4)h\}. \end{aligned}$$

Since

$$H_{\mu'}(\xi_{[m]}) = \sum_{\mathbf{z} \in A^m} \mu'_{[m]}(\mathbf{z}) \cdot (-\log \mu'_{[m]}(\mathbf{z})),$$

we deduce from (65) and Markov's inequality that

$$\mu'_{[m]}(\mathbf{Q}) \geq 1 - \frac{(1-\delta/2)(h+1)}{(1-\delta/4)h} = \frac{\delta/4}{(1-\delta/4)} - \frac{(1-\delta/2)}{(1-\delta/4)h}.$$

Since our assumptions on μ give

$$h \geq 160c_B E / \delta^2 \geq 160c_B r^4 n / \delta^2,$$

it follows that $\mu'_{[m]}(\mathbf{Q}) \geq \delta/8$ provided n is large enough in terms of r .

Let $\mathbf{R} := \{\mathbf{x} \in \mathbf{P} : \mathbf{x}_{[m]} \in \mathbf{Q}\}$. Since $\mathbf{R} \subseteq \mathbf{P}$, we have $\mu'_{|\mathbf{R}} = \mu_{|\mathbf{R}}$.

This completes the construction of \mathbf{S} , \mathbf{R} and \mathbf{Q} . Among the desired properties, (a) and (c) are clear from the construction. Property (b) holds because

$$\mu(\mathbf{R}) = \mu_{|\mathbf{P}}(\mathbf{R}) \cdot \mu(\mathbf{P}) = \mu'_{[m]}(\mathbf{Q}) \cdot \mu(\mathbf{P}) \geq \frac{\delta}{8} \cdot \frac{1}{4M}.$$

To prove property (d), observe that every $\mathbf{x} \in \mathbf{R} \subseteq \mathbf{P}$ satisfies

$$\mu_{|\mathbf{R}}(\mathbf{x}) = \frac{\mu'(\mathbf{x})}{\mu'(\mathbf{R})} = \frac{\mu'(\mathbf{x})}{\mu'_{[m]}(\mathbf{Q})} \leq \frac{e^{-h}}{\mu'_{[m]}(\mathbf{Q})},$$

whereas every $\mathbf{z} \in \mathbf{Q}$ satisfies

$$\mu_{|\mathbf{R}}\{\mathbf{x} : \mathbf{x}_{[m]} = \mathbf{z}\} = \frac{\mu'\{\mathbf{x} : \mathbf{x}_{[m]} = \mathbf{z}\}}{\mu'(\mathbf{R})} = \frac{\mu'_{[m]}(\mathbf{z})}{\mu'_{[m]}(\mathbf{Q})} \geq \frac{e^{-(1-\delta/4)h}}{\mu'_{[m]}(\mathbf{Q})}.$$

Finally, towards property (e), Corollary 3.2 and Lemma 8.3 give

$$\text{TC}(\mu_{|\mathbf{R}}) \leq \frac{1}{\mu'(\mathbf{R})} (\text{TC}(\mu') + \log 2) \leq \frac{8}{\delta} (4\text{TC}(\mu) + \log 2).$$

Since we assume $\text{TC}(\mu) > r^4 n$, this last upper bound is less than $33\text{TC}(\mu)/\delta$ provided n is sufficiently large in terms of r . \square

Next we explain how the sets \mathbf{S} , \mathbf{R} and \mathbf{Q} obtained by Proposition 8.4 can be turned into a set \mathbf{V} as required by Proposition 8.1.

Let $\nu := \mu_{\mathbb{R}}$, and let $\langle \cdot \rangle_{\nu}$ denote integration with respect to ν . We apply Theorem B to ν with input parameters r and $\varepsilon := 1/2$. The number of terms in the resulting mixture is at most $c_{\text{B}} \exp(c_{\text{B}} \cdot \text{TC}(\nu))$, and by conclusion (e) of Proposition 8.4 this is at most $c_{\text{B}} \exp(33c_{\text{B}}E/\delta)$. At least half the weight in the mixture must belong to terms that satisfy $T(\kappa n, r)$. Therefore there is at least one term, which we may write in the form $\rho \cdot \nu$, which satisfies $T(\kappa n, r)$ and also

$$(66) \quad \langle \rho \rangle_{\nu} \geq \frac{1}{2c_{\text{B}}} e^{-33c_{\text{B}}E/\delta}.$$

For each $\mathbf{z} \in \mathcal{Q}$, let

$$\mathbf{B}_{\mathbf{z}} := \{\mathbf{x} \in \mathbb{R} : \mathbf{x}_{\text{S}} = \mathbf{z}\},$$

and now let

$$\mathcal{Q}_{\downarrow} := \left\{ \mathbf{z} \in \mathcal{Q} : \int_{\mathbf{B}_{\mathbf{z}}} \rho \, d\nu \geq \delta \cdot \nu(\mathbf{B}_{\mathbf{z}}) \cdot \langle \rho \rangle_{\nu} \right\}.$$

The next step in the construction of \mathbf{V} is to choose a subset $\mathbf{U} \subseteq \mathbb{R}$ at random so that the events $\{\mathbf{U} \ni \mathbf{x}\}$ have independent probabilities $\rho(\mathbf{x})$ for $\mathbf{x} \in \mathbb{R}$. In the final step, \mathbf{V} is obtained from \mathbf{U} by trimming it slightly. Let \mathbf{E} and Var denote expectation and variance with respect to the random choice of \mathbf{U} .

Lemma 8.5. — *Provided n is sufficiently large in terms of r , we have*

$$\mathbf{E} \left| \nu(\mathbf{U} \cap \mathbf{B}_{\mathbf{z}}) - \int_{\mathbf{B}_{\mathbf{z}}} \rho \, d\nu \right| \leq \delta \int_{\mathbf{B}_{\mathbf{z}}} \rho \, d\nu$$

for every $\mathbf{z} \in \mathcal{Q}_{\downarrow}$.

Proof. — In this proof we need the inequality

$$(67) \quad e^{-\delta h/4} \leq \frac{\delta^3}{2c_{\text{B}}} \exp(-33c_{\text{B}}E/\delta),$$

where h is the quantity from Lemma 8.3. Since $h \geq 160c_{\text{B}}E/\delta^2$, this holds provided

$$\exp(-40c_{\text{B}}E/\delta) \leq \frac{\delta^3}{2c_{\text{B}}} \exp(-33c_{\text{B}}E/\delta).$$

Given our current assumption that $E > r^4 n$, this last requirement holds for all n that are sufficiently large in terms of r .

So now assume that (67) holds. Then for each $\mathbf{z} \in \mathcal{Q}_{\downarrow}$ we may combine (66), (67), conclusion (d) from Proposition 8.4, and the definition of \mathcal{Q}_{\downarrow} to obtain

$$(68) \quad \max_{\mathbf{x} \in \mathbf{B}_{\mathbf{z}}} \nu(\mathbf{x}) \leq \frac{\delta^3}{2c_{\text{B}}} e^{-33c_{\text{B}}E/\delta} \nu(\mathbf{B}_{\mathbf{z}}) \leq \delta^3 \cdot \nu(\mathbf{B}_{\mathbf{z}}) \cdot \langle \rho \rangle_{\nu} \leq \delta^2 \int_{\mathbf{B}_{\mathbf{z}}} \rho \, d\nu.$$

Now consider the random variable

$$\mathbf{X} := \nu(\mathbf{U} \cap \mathbf{B}_z) - \int_{\mathbf{B}_z} \rho \, d\nu = \sum_{\mathbf{x} \in \mathbf{B}_z} \nu(\mathbf{x}) \cdot (1_{\{\mathbf{U} \ni \mathbf{x}\}} - \rho(\mathbf{x})).$$

From the distribution of the random set \mathbf{U} , it follows that $\mathbf{E}\mathbf{X} = 0$ and

$$\text{Var}(\mathbf{X}) = \sum_{\mathbf{x} \in \mathbf{B}_z} (\nu(\mathbf{x}))^2 \cdot \text{Var}(1_{\{\mathbf{U} \ni \mathbf{x}\}}) \leq \sum_{\mathbf{x} \in \mathbf{B}_z} (\nu(\mathbf{x}))^2 \cdot \rho(\mathbf{x}).$$

By (68), this is at most

$$\delta^2 \int_{\mathbf{B}_z} \rho \, d\nu \cdot \sum_{\mathbf{x} \in \mathbf{B}_z} \nu(\mathbf{x}) \rho(\mathbf{x}) = \delta^2 \left(\int_{\mathbf{B}_z} \rho \, d\nu \right)^2.$$

Therefore, by the monotonicity of Lebesgue norms,

$$\mathbf{E}|\mathbf{X}| \leq \sqrt{\mathbf{E}(\mathbf{X}^2)} = \sqrt{\text{Var}(\mathbf{X})} \leq \delta \int_{\mathbf{B}_z} \rho \, d\nu. \quad \square$$

Corollary 8.6. — *Provided n is sufficiently large in terms of r , we have*

$$\mathbf{E}|\nu(\mathbf{U}) - \langle \rho \rangle_\nu| \leq \sum_{z \in \mathbf{Q}} \mathbf{E} \left| \nu(\mathbf{U} \cap \mathbf{B}_z) - \int_{\mathbf{B}_z} \rho \, d\nu \right| \leq 3\delta \langle \rho \rangle_\nu.$$

Proof. — The left-hand inequality follows from the triangle inequality and the fact that $(\mathbf{B}_z : z \in \mathbf{Q})$ is a partition of \mathbf{R} . So consider the right-hand inequality. By separating \mathbf{Q} into \mathbf{Q}_1 and $\mathbf{Q} \setminus \mathbf{Q}_1$, the sum over \mathbf{Q} is at most

$$\begin{aligned} (69) \quad & \sum_{z \in \mathbf{Q}_1} \mathbf{E} \left| \nu(\mathbf{U} \cap \mathbf{B}_z) - \int_{\mathbf{B}_z} \rho \, d\nu \right| + \sum_{z \in \mathbf{Q} \setminus \mathbf{Q}_1} \mathbf{E} \nu(\mathbf{U} \cap \mathbf{B}_z) + \sum_{z \in \mathbf{Q} \setminus \mathbf{Q}_1} \int_{\mathbf{B}_z} \rho \, d\nu \\ & = \sum_{z \in \mathbf{Q}_1} \mathbf{E} \left| \nu(\mathbf{U} \cap \mathbf{B}_z) - \int_{\mathbf{B}_z} \rho \, d\nu \right| + 2 \sum_{z \in \mathbf{Q} \setminus \mathbf{Q}_1} \int_{\mathbf{B}_z} \rho \, d\nu. \end{aligned}$$

Using the previous lemma, the first sum in (69) is at most

$$\delta \sum_{z \in \mathbf{Q}_1} \int_{\mathbf{B}_z} \rho \, d\nu \leq \delta \langle \rho \rangle_\nu.$$

By the definition of \mathbf{Q}_1 , the second sum in (69) is at most

$$2\delta \sum_{z \in \mathbf{Q} \setminus \mathbf{Q}_1} \nu(\mathbf{B}_z) \cdot \langle \rho \rangle_\nu \leq 2\delta \langle \rho \rangle_\nu. \quad \square$$

Lemma 8.7. — *Provided n is sufficiently large in terms of r , we have*

$$\mathbf{E}[\nu(\mathbf{U}) \cdot \bar{d}_n(\nu_{|\mathbf{U}}, \nu_{|\rho})] \leq 7\delta \langle \rho \rangle_\nu.$$

Proof. — In the event that $\nu(\mathbf{U}) > 0$, we can bound the distance $\bar{d}_n(\nu_{|\mathbf{U}}, \nu_{|\rho})$ using Corollary 4.2 with the partition $(\mathbf{B}_z : z \in \mathbf{Q})$ of \mathbf{R} . Each \mathbf{B}_z has diameter at most δ according to d_n , and the diameter of the whole of \mathbf{A}^n is 1, so that corollary gives

$$\bar{d}_n(\nu_{|\mathbf{U}}, \nu_{|\rho}) \leq \delta \sum_{z \in \mathbf{Q}} \nu_{|\rho}(\mathbf{B}_z) + \frac{1}{2} \sum_{z \in \mathbf{Q}} |\nu_{|\mathbf{U}}(\mathbf{B}_z) - \nu_{|\rho}(\mathbf{B}_z)|.$$

Since the values $\nu_{|\rho}(\mathbf{B}_z)$ for $z \in \mathbf{Q}$ must sum to 1, the first term on the right is just δ . Let us also remove the factor $1/2$ from the second right-hand term for simplicity.

Multiplying by $\nu(\mathbf{U})$, we obtain

$$\begin{aligned} \nu(\mathbf{U}) \cdot \bar{d}_n(\nu_{|\mathbf{U}}, \nu_{|\rho}) &\leq \delta \nu(\mathbf{U}) + \nu(\mathbf{U}) \sum_{z \in \mathbf{Q}} |\nu_{|\mathbf{U}}(\mathbf{B}_z) - \nu_{|\rho}(\mathbf{B}_z)| \\ &= \delta \nu(\mathbf{U}) + \sum_{z \in \mathbf{Q}} |\nu(\mathbf{U} \cap \mathbf{B}_z) - \nu(\mathbf{U}) \nu_{|\rho}(\mathbf{B}_z)| \\ &\leq \delta \nu(\mathbf{U}) + \sum_{z \in \mathbf{Q}} \left| \nu(\mathbf{U} \cap \mathbf{B}_z) - \int_{\mathbf{B}_z} \rho \, d\nu \right| \\ &\quad + \sum_{z \in \mathbf{Q}} |\nu(\mathbf{U}) - \langle \rho \rangle_\nu| \cdot \nu_{|\rho}(\mathbf{B}_z), \end{aligned}$$

using that $\int_{\mathbf{B}_z} \rho \, d\nu = \langle \rho \rangle_\nu \cdot \nu_{|\rho}(\mathbf{B}_z)$ for the last step. These inequalities also hold in case $\nu(\mathbf{U}) = 0$.

Taking expectations, we obtain

$$\begin{aligned} \mathbf{E}[\nu(\mathbf{U}) \cdot \bar{d}_n(\nu_{|\mathbf{U}}, \nu_{|\rho})] &\leq \delta \mathbf{E} \nu(\mathbf{U}) + \sum_{z \in \mathbf{Q}} \mathbf{E} \left| \nu(\mathbf{U} \cap \mathbf{B}_z) - \int_{\mathbf{B}_z} \rho \, d\nu \right| \\ &\quad + \mathbf{E} |\nu(\mathbf{U}) - \langle \rho \rangle_\nu| \cdot \sum_{z \in \mathbf{Q}} \nu_{|\rho}(\mathbf{B}_z) \\ &= \delta \langle \rho \rangle_\nu + \sum_{z \in \mathbf{Q}} \mathbf{E} \left| \nu(\mathbf{U} \cap \mathbf{B}_z) - \int_{\mathbf{B}_z} \rho \, d\nu \right| + \mathbf{E} |\nu(\mathbf{U}) - \langle \rho \rangle_\nu|. \end{aligned}$$

Corollary 8.6 bounds both the second and the third terms here: each is at most $3\delta \langle \rho \rangle_\nu$. Adding these estimates completes the proof. \square

Proof of Proposition 8.1. — We may assume that $r < 1/8$ without loss of generality.

As explained previously, there are three cases to consider. In case μ has an atom \mathbf{x} of measure at least (59), we can let $V := \{\mathbf{x}\}$. In case $\text{TC}(\mu) \leq r^4 n$, we use the set V provided by Lemma 8.2.

So now suppose that neither of those cases holds. This assumption provides the hypotheses for Proposition 8.4. Let R be the set provided by that proposition, and let $\nu := \mu|_R$. Choose ρ and construct a random subset U of R as above. By Corollary 8.6, Lemma 8.7, and two applications of Markov's inequality, the random set U satisfies each of the inequalities

$$|\nu(U) - \langle \rho \rangle_\nu| \leq 9\delta \langle \rho \rangle_\nu$$

and

$$\nu(U) \cdot \overline{d}_n(\nu|_U, \nu|_\rho) \leq 21\delta \langle \rho \rangle_\nu$$

with probability at least $2/3$.

Therefore there exists a subset $U \subseteq R$ which satisfies both of these inequalities. For that set U , and recalling that we chose $\delta \leq 1/18$ (see (58)), the first inequality gives $\nu(U) \geq \langle \rho \rangle_\nu / 2$. Using this, the second inequality implies that

$$\overline{d}_n(\nu|_U, \nu|_\rho) \leq 42\delta.$$

Since we also chose $\delta \leq r^2/42$ (see (58) again), and since $\nu|_\rho$ satisfies $T(\kappa n, r)$, we may now repeat the proof of Lemma 8.2 (based on Proposition 5.9). It gives another subset V of A^n with $\nu|_U(V) \geq 1 - 4r \geq 1/2$ and such that the measure $(\nu|_U)|_V$ satisfies $T(\kappa n, 9r)$, provided n is sufficiently large in terms of κ and r . Since $\nu|_U$ is supported on U , we may replace V with $U \cap V$ if necessary and then assume that $V \subseteq U$. Having done so, we have $(\nu|_U)|_V = (\mu|_U)|_V = \mu|_V$.

Finally we must prove the lower bound on $\mu(V)$ for a suitable constant c . Recalling (66), it follows from the estimates above that

$$\begin{aligned} \mu(V) &\geq \mu(U) \cdot \mu(V|U) = \mu(R) \cdot \nu(U) \cdot \nu(V|U) \\ &\geq \frac{\delta}{32M} \cdot \frac{\langle \rho \rangle_\nu}{2} \cdot \frac{1}{2} \geq \frac{\delta}{256c_B M} \cdot e^{-33c_B E/\delta}. \end{aligned}$$

Recalling the definition of M in Lemma 8.3, this last expression is equal to

$$\frac{\delta}{256c_B \lceil n \log |A| \rceil} \cdot e^{-33c_B E/\delta}.$$

Since we have restricted attention to the last of our three cases, we know that $E > r^4 n$. Therefore, for any fixed $c > 33c_B/\delta$, the last expression must be greater than e^{-cE} for all sufficiently large n , since the negative exponential is stronger than the prefactor as $n \rightarrow \infty$. \square

Remark. — In Part III we apply Theorem C during the proof of Theorem A. In that application, the measure μ is the conditional distribution of n consecutive letters in a stationary ergodic process, given the corresponding letters in another process coupled to the first. In this case, the Shannon–McMillan theorem enables us to obtain conclusion (d) of Proposition 8.4 after just trimming the measure μ slightly, without using Lemma 8.3. This makes for a slightly easier proof of Theorem C in this case.

Remark. — In the proof of Theorem C, we first obtain the new measure $\nu = \mu|_{\mathbb{R}}$. Then we apply Theorem B to it, consider just one of the summands $\rho \cdot \nu$ in the resulting mixture, and use a random construction to produce a set U such that $\overline{d}_n(\nu|_U, \nu|_{\rho})$ is small.

An alternative is to consider the whole mixture

$$(70) \quad \nu = \rho_1 \cdot \nu_1 + \cdots + \rho_m \cdot \nu_m$$

given by Theorem B, and then use a random partition U_1, \dots, U_m so that (i) each $\mathbf{x} \in A^n$ chooses independently which cell to belong to and (ii) each event $\{U_j \ni \mathbf{x}\}$ has probability $\rho_j(\mathbf{x})$. Using similar estimates to those above, one can show that, in expectation, most of the weight in the mixture is on terms for which $\overline{d}_n(\nu|_{U_j}, \nu|_{\rho_j})$ is small.

However, in this alternative approach, we must still begin by passing from μ to $\nu = \mu|_{\mathbb{R}}$. It does not give directly a partition as required by Theorem C, since we have $\nu|_{U_j} = \mu|_{\mathbb{R} \cap U_j}$, and the sets $\mathbb{R} \cap U_j$ do not cover the whole of A^n . For this reason, it seems easier to use just one well-chosen term from (70), as we have above.

9. A reformulation: extremality

The work of Parts II and III is much easier if we base it on a slightly different notion of measure concentration than that studied so far. This alternative notion comes from an idea of Thouvenot in ergodic theory, described in [107, Definition 6.3]. It is implied by a suitable T-inequality, so we can still bring Theorem C to bear when we need it later, but this new notion is more robust and therefore easier to carry from one setting to another.

9.1. Extremal measures. — Here is our new notion of measure concentration.

Definition 9.1. — Let (\mathbf{K}, d, μ) be a metric probability space. It is (κ, r) -**extremal** if any representation of μ as a mixture $\int \mu_{\bullet} d\mathbf{P}$ satisfies

$$\int \overline{d}(\mu_{\omega}, \mu) \mathbf{P}(d\omega) \leq \frac{1}{\kappa} \int \mathbf{D}(\mu_{\omega} \| \mu) \mathbf{P}(d\omega) + r.$$

In the first place, notice that if r is at least the diameter of (\mathbf{K}, d) then (\mathbf{K}, d, μ) is (κ, r) -extremal for every $\kappa > 0$.

If (\mathbf{K}, d, μ) satisfies $\Gamma(\kappa, r)$ then it is also (κ, r) -extremal: $\Gamma(\kappa, r)$ gives an inequality for each $\bar{d}(\mu_\omega, \mu)$, and we can simply integrate this with respect to $\mathbb{P}(d\omega)$. However, extremality is generally slightly weaker than a Γ -inequality. This becomes clearer from various stability properties of extremality. In particular, the next lemma shows that extremality survives with roughly the same constants under sufficiently small perturbations in \bar{d} , whereas Γ -inequalities can degrade substantially unless we also condition the perturbed measure on a subset (Proposition 5.9). This deficiency of Γ -inequalities may be seen in the examples sketched in the second remark after the proof of Proposition 5.9.

Lemma 9.2 (*Stability under perturbation in transportation*). — *Let $\mu, \mu' \in \text{Prob}(\mathbf{K})$, assume that μ is (κ, r) -extremal, and let $\delta := \bar{d}(\mu, \mu')$. Then μ' is $(\kappa, r + 2\delta)$ -extremal.*

Proof. — By the vague compactness of $\text{Prob}(\mathbf{K})$, there is a coupling λ of μ' and μ such that

$$\int d(x', x) \lambda(dx', dx) = \delta.$$

Let us disintegrate λ over the first coordinate, thus:

$$(71) \quad \lambda = \int_{\mathbf{K}} \delta_{x'} \times \lambda_{x'} \mu'(dx').$$

Now suppose that μ' is equal to the mixture $\int \mu'_\bullet d\mathbb{P}$. Combining this with (71), we obtain

$$\mu = \int_{\mathbf{K}} \lambda_{x'} \mu'(dx') = \int_{\Omega} \left[\int_{\mathbf{K}} \lambda_{x'} \mu'_\omega(dx') \right] \mathbb{P}(d\omega) = \int_{\Omega} \mu_\omega \mathbb{P}(d\omega),$$

where

$$(72) \quad \mu_\omega := \int_{\mathbf{K}} \lambda_{x'} \mu'_\omega(dx').$$

Thus, we have turned our representation of μ' as a mixture into a corresponding representation of μ . Therefore the assumed extremality gives

$$\int \bar{d}(\mu_\omega, \mu) \mathbb{P}(d\omega) \leq \frac{1}{\kappa} \int D(\mu_\omega \| \mu) \mathbb{P}(d\omega) + r.$$

The triangle inequality for \bar{d} gives

$$\begin{aligned} & \int \bar{d}(\mu'_\omega, \mu') \mathbb{P}(d\omega) \\ & \leq \int \bar{d}(\mu'_\omega, \mu_\omega) \mathbb{P}(d\omega) + \int \bar{d}(\mu_\omega, \mu) \mathbb{P}(d\omega) + \bar{d}(\mu, \mu'). \end{aligned}$$

The last right-hand term here is δ . We have just found a bound for the middle term. Finally, by the definition (72), the first term is at most

$$\begin{aligned} \int \left[\iint d(x', x) \lambda_{x'}(dx) \mu'_\omega(dx') \right] \mathbb{P}(d\omega) &= \iint d(x', x) \lambda_{x'}(dx) \mu'(dx') \\ &= \int d \, d\lambda = \delta. \end{aligned}$$

Combining these estimates gives

$$\int \bar{d}(\mu'_\omega, \mu') \mathbb{P}(d\omega) \leq \frac{1}{\kappa} \int \mathbb{D}(\mu_\omega \parallel \mu) \mathbb{P}(d\omega) + r + 2\delta.$$

Finally, for the remaining integral on the right-hand side, we apply the fact that KL divergence is non-increasing when one forms the compounds of two probability measures with the same kernel. See, for instance, [11, Section 4.4, part 1] for this standard consequence of the chain rule. This gives

$$\begin{aligned} \int \mathbb{D}(\mu_\omega \parallel \mu) \mathbb{P}(d\omega) &= \int \mathbb{D}\left(\int_{\mathbb{K}} \lambda_{x'} \mu'_\omega(dx') \parallel \int_{\mathbb{K}} \lambda_{x'} \mu'(dx') \right) \mathbb{P}(d\omega) \\ &\leq \int \mathbb{D}(\mu'_\omega \parallel \mu') \mathbb{P}(d\omega). \quad \square \end{aligned}$$

Next we show how extremality is inherited by product measures.

Lemma 9.3 (*Products of extremal measures*). — *Let $(\mathbb{K}, d_{\mathbb{K}})$ and $(\mathbb{L}, d_{\mathbb{L}})$ be compact metric spaces, let $0 < \alpha < 1$, and let d be the following metric on $\mathbb{K} \times \mathbb{L}$:*

$$(73) \quad d((x, y), (x', y')) := \alpha d_{\mathbb{K}}(x, x') + (1 - \alpha) d_{\mathbb{L}}(y, y').$$

Let $\mu \in \text{Prob}(\mathbb{K})$ and $\nu \in \text{Prob}(\mathbb{L})$, and assume that they are $(\alpha\kappa, r_{\mathbb{K}})$ -extremal and $((1 - \alpha)\kappa, r_{\mathbb{L}})$ -extremal, respectively. Let $r := \alpha r_{\mathbb{K}} + (1 - \alpha) r_{\mathbb{L}}$. Then $\mu \times \nu$ is (κ, r) -extremal on the metric space $(\mathbb{K} \times \mathbb{L}, d)$.

Proof. — Consider a representation

$$\mu \times \nu = \int_{\Omega} \lambda_\omega \mathbb{P}(d\omega).$$

For each ω , let $\lambda_{\mathbb{K}, \omega}$ be the marginal of λ_ω on \mathbb{K} , and consider the disintegration of λ_ω over the first coordinate:

$$\lambda_\omega = \int_{\mathbb{K}} (\delta_x \times \lambda_{\mathbb{L}, (\omega, x)}) \lambda_{\mathbb{K}, \omega}(dx).$$

We apply Lemma 4.3 to each ω separately and then integrate. This gives

$$(74) \quad \begin{aligned} & \int \bar{d}(\lambda_\omega, \mu \times \nu) \mathbb{P}(d\omega) \\ & \leq \alpha \int \bar{d}_K(\lambda_{K,\omega}, \mu) \mathbb{P}(d\omega) \\ & \quad + (1 - \alpha) \iint \bar{d}_L(\lambda_{L,(\omega,x)}, \nu) \lambda_{K,\omega}(dx) \mathbb{P}(d\omega). \end{aligned}$$

The measure μ is equal to the mixture $\int \lambda_{K,\omega} d\mathbb{P}$. Since μ is $(\alpha\kappa, r_K)$ -extremal, it follows that the first right-hand term in (74) is at most

$$\frac{1}{\kappa} \int \mathbb{D}(\lambda_{K,\omega} \| \mu) \mathbb{P}(d\omega) + \alpha r_K.$$

On the other hand, let $\Omega' := \Omega \times K$, and let $\mathbb{Q} := \mathbb{P} \times \lambda_{K,\bullet}$. Then ν is equal to the mixture

$$\iint \lambda_{L,(\omega,x)} \lambda_{K,\omega}(dx) \mathbb{P}(d\omega) = \int \lambda_{L,(\omega,x)} \mathbb{Q}(d\omega, dx).$$

Applying the fact that ν is $((1 - \alpha)\kappa, r_L)$ -extremal to this mixture, it follows that the second right-hand term in (74) is at most

$$\frac{1}{\kappa} \iint \mathbb{D}(\lambda_{L,(\omega,x)} \| \nu) \lambda_{K,\omega}(dx) \mathbb{P}(d\omega) + (1 - \alpha)r_L.$$

Adding these two estimates, we obtain

$$\begin{aligned} & \int \bar{d}(\lambda_\omega, \mu \times \nu) \mathbb{P}(d\omega) \\ & \leq \frac{1}{\kappa} \int \mathbb{D}(\lambda_{K,\omega} \| \mu) \mathbb{P}(d\omega) + \frac{1}{\kappa} \iint \mathbb{D}(\lambda_{L,(\omega,x)} \| \nu) \lambda_{K,\omega}(dx) \mathbb{P}(d\omega) + r. \end{aligned}$$

By the chain rule for KL divergence, this is equal to

$$\frac{1}{\kappa} \int \mathbb{D}(\lambda_\omega \| \mu \times \nu) \mathbb{P}(d\omega) + r. \quad \square$$

9.2. Extremal kernels. — In Parts II and III, we use extremality not just for one measure on a compact metric space (K, d) , but for a $\text{Prob}(K)$ -valued kernel μ_\bullet on another probability space (Y, ν) . In those applications, the important property is that the individual measures μ_y are highly extremal for most $y \in Y$ according to ν . An exceptional set of small probability is allowed.

With this in mind, let us extend Definition 9.1 to kernels.

Definition 9.4. — Let (Y, ν) be a probability space, let (K, d) be a compact metric space, and let μ_\bullet be a kernel from Y to K . The pair (ν, μ_\bullet) is (κ, r) -**extremal** if there is a measurable function

$$Y \longrightarrow (0, \infty) : y \mapsto r_y$$

such that μ_y is (κ, r_y) -extremal for every $y \in Y$ and such that

$$(75) \quad \int r_y \nu(dy) \leq r.$$

Thus, we control the extremality of the individual measures μ_y by the average of their parameters in (75), rather than by explicitly introducing a ‘bad set’ of small probability. This turns out to give much simpler estimates in many of the proofs below, starting with Lemma 9.7. Of course, Definition 9.4 is easily related to such a ‘bad set’ via Markov’s inequality:

Lemma 9.5. — If (ν, μ_\bullet) is (κ, r) -extremal, then there is a subset $Y_1 \subseteq Y$ such that $\nu(Y_1) \geq 1 - \sqrt{r}$ and such that μ_y is (κ, \sqrt{r}) -extremal for every $y \in Y_1$.

On the other hand, if (K, d) has diameter r_0 , $Y_1 \subseteq Y$ is measurable, and μ_y is (κ, r) -extremal for all $y \in Y_1$, then (ν, μ_\bullet) is extremal with parameters

$$(\kappa, r + r_0 \nu(Y \setminus Y_1)).$$

Example 9.6. — Let $p_1 \mu_1 + \dots + p_m \mu_m$ be a mixture given by Theorem B. Regard $p = (p_1, \dots, p_m)$ as a measure on $[m]$ and $(\mu_j)_{j=1}^m$ as a kernel from $[m]$ to A^n . Then the pair (p, μ_\bullet) is $(rn/1200, r + \varepsilon)$ -extremal.

Definition 9.4 clearly depends on the kernel μ_\bullet only up to agreement ν -almost everywhere. Therefore it may be regarded unambiguously as a property of the measure $\nu \times \mu_\bullet$ on the product space $Y \times K$. In the sequel we sometimes write (κ, r) -extremality as a property of a measure on $Y \times K$ rather than of a pair such as (ν, μ_\bullet) .

The next result is an easy extension of Lemma 9.2 to this setting.

Lemma 9.7 (*Stability under perturbation in transportation*). — Let (Y, ν) be a probability space, let (K, d) be a compact metric space, and let μ_\bullet and μ'_\bullet be two kernels from Y to K . Assume that (ν, μ_\bullet) is (κ, r) -extremal, and let

$$\delta := \int \bar{d}(\mu_y, \mu'_y) \nu(dy).$$

Then (ν, μ'_\bullet) is $(\kappa, r + 2\delta)$ -extremal.

Proof. — Let $y \mapsto r_y$ be the function promised by Definition 9.4 for (ν, μ_\bullet) . By applying Lemma 9.2 for each y separately, it follows that the function

$$y \mapsto r_y + 2\bar{d}(\mu_y, \mu'_y)$$

serves the same purpose for (ν, μ'_\bullet) . This function has integral at most $r + 2\delta$. \square

Lemma 9.3 also has an easy generalization to kernels.

Lemma 9.8 (Pointwise products of extremal measures). — Let $(\mathbf{K}, d_{\mathbf{K}})$ and $(\mathbf{L}, d_{\mathbf{L}})$ be compact metric spaces, let $0 < \alpha < 1$, and let d be the metric (73) on $\mathbf{K} \times \mathbf{L}$.

Let μ_\bullet and μ'_\bullet be kernels from (Y, ν) to \mathbf{K} and \mathbf{L} , respectively, and assume that (ν, μ_\bullet) and (ν, μ'_\bullet) are $(\alpha\kappa, r_{\mathbf{K}})$ -extremal and $((1 - \alpha)\kappa, r_{\mathbf{L}})$ -extremal, respectively. Let $r := \alpha r_{\mathbf{K}} + (1 - \alpha)r_{\mathbf{L}}$.

Then the measure ν and pointwise-product kernel $y \mapsto \mu_y \times \mu'_y$ are (κ, r) -extremal on the metric space $(\mathbf{K} \times \mathbf{L}, d)$.

Proof. — Let $y \mapsto r_{\mathbf{K},y}$ and $y \mapsto r_{\mathbf{L},y}$ be functions as promised by Definition 9.4 for μ_\bullet and μ'_\bullet , respectively. For each $y \in Y$, apply Lemma 9.3 to conclude that $\mu_y \times \mu'_y$ is $(\kappa, \alpha r_{\mathbf{K},y} + (1 - \alpha)r_{\mathbf{L},y})$ -extremal. Finally, observe that

$$\int (\alpha r_{\mathbf{K},y} + (1 - \alpha)r_{\mathbf{L},y}) \nu(dy) \leq \alpha r_{\mathbf{K}} + (1 - \alpha)r_{\mathbf{L}}. \quad \square$$

We often use extremality when μ_\bullet is a conditional distribution of a \mathbf{K} -valued random variable F given \mathcal{G} , where F is defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and \mathcal{G} is a σ -subalgebra of \mathcal{F} . Then the following nomenclature is useful.

Definition 9.9. — In the setting above, F is (κ, r) -**extremal over** \mathcal{G} if $\mathbb{P} \times \mu_\bullet$ is (κ, r) -extremal, where μ_\bullet is a conditional distribution of F given \mathcal{G} .

In this setting, the important consequence of extremality is the following.

Lemma 9.10. — Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, let $\mathcal{G} \subseteq \mathcal{H}$ be σ -subalgebras of \mathcal{F} that are countably generated modulo \mathbb{P} , let (\mathbf{K}, d) be a finite metric space, and let $F : \Omega \rightarrow \mathbf{K}$ be measurable and (κ, r) -extremal over \mathcal{G} . Let μ_\bullet and ν_\bullet be conditional distributions for F given \mathcal{G} and \mathcal{H} , respectively. Then

$$\int \bar{d}(\nu_\omega, \mu_\omega) \mathbb{P}(d\omega) \leq \frac{1}{\kappa} [\mathbb{H}(F | \mathcal{G}) - \mathbb{H}(F | \mathcal{H})] + r.$$

One can drop the assumption of countable generation, but we omit the details.

Proof. — Let $\omega \mapsto r_\omega$ be the function promised by Definition 9.4 for $(\mathbb{P}, \mu_\bullet)$.

Since \mathcal{H} is countably generated modulo \mathbb{P} , there is a measurable map $\psi : \Omega \rightarrow Y$ to a standard measurable space such that ψ generates \mathcal{H} modulo \mathbb{P} . Up to agreement \mathbb{P} -a.e., we may write ν_ω as $\nu'_{\psi(\omega)}$ for some kernel ν'_\bullet from Y to \mathbf{K} .

In addition, let $\mathbf{Q}_\bullet : \Omega \rightarrow \text{Prob}(Y)$ be a conditional distribution for ψ given \mathcal{G} . Such a \mathbf{Q}_\bullet exists because Y is standard. Since $\mathcal{H} \supseteq \mathcal{G}$, the tower property of conditional expectation gives

$$(76) \quad \mu_\omega = \int_Y v'_y \mathbf{Q}_\omega(dy) \quad \text{for P-a.e. } \omega.$$

Therefore, for P-a.e. ω , we may apply the (κ, r_ω) -extremality of μ_ω to obtain

$$\int \bar{d}(v'_y, \mu_\omega) \mathbf{Q}_\omega(dy) \leq \frac{1}{\kappa} \int D(v'_y \parallel \mu_\omega) \mathbf{Q}_\omega(dy) + r_\omega.$$

Integrating with respect to $P(d\omega)$, this gives

$$\begin{aligned} \int \bar{d}(v_\omega, \mu_\omega) P(d\omega) &= \iint \bar{d}(v'_y, \mu_\omega) \mathbf{Q}_\omega(dy) P(d\omega) \\ &\leq \frac{1}{\kappa} \iint D(v'_y \parallel \mu_\omega) \mathbf{Q}_\omega(dy) P(d\omega) + \int r_\omega P(d\omega) \\ &\leq \frac{1}{\kappa} \int D(v_\omega \parallel \mu_\omega) P(d\omega) + r. \end{aligned}$$

On the other hand, using again that $\mathcal{H} \supseteq \mathcal{G}$, Lemma 2.2 gives

$$\int D(v_\omega \parallel \mu_\omega) P(d\omega) = H(F \mid \mathcal{G}) - H(F \mid \mathcal{H}).$$

□

Combining Lemmas 9.10 and 9.7 gives the following useful consequence.

Corollary 9.11 (Inheritance of extremality). — *In the setting of Lemma 9.10, let*

$$a := H(F \mid \mathcal{G}) - H(F \mid \mathcal{H}).$$

Then F is $(\kappa, 2a/\kappa + 3r)$ -extremal over \mathcal{H} .

Finally, here is an analog of Lemma 5.10 for extremal kernels.

Lemma 9.12 (Stability under lifting). — *Let d_n be the normalized Hamming metric on A^n for some finite set A , and let $\nu \times \mu_\bullet$ be a probability measure on $Y \times A^n$. Let $S \subseteq [n]$ satisfy $|S| \geq (1 - a)n$, and let $\mu_{S,y}$ be the projection of μ_y to A^S for each $y \in Y$. If $\nu \times \mu_{S,\bullet}$ is (κ, r) -extremal, then $\nu \times \mu_\bullet$ is $(n\kappa/|S|, r + a)$ -extremal.*

Proof. — First consider the case when Y is a single point. Effectively this means we have a single measure μ on A^n for which μ_S is (κ, r) -extremal. Suppose μ is represented as the mixture $\int \mu_\omega P(d\omega)$. Projecting to A^S , this becomes $\mu_S = \int \mu_{S,\omega} P(d\omega)$,

and we have $D(\mu_{S,\omega} \parallel \mu_S) \leq D(\mu_\omega \parallel \mu)$ for each ω as in the proof of Lemma 5.10. So our assumption on μ_S gives

$$\begin{aligned} \int \overline{d}_S(\mu_{S,\omega}, \mu_S) P(d\omega) &\leq \frac{1}{\kappa} \int D(\mu_{S,\omega} \parallel \mu_S) P(d\omega) + r \\ &\leq \frac{1}{\kappa} \int D(\mu_\omega \parallel \mu) P(d\omega) + r. \end{aligned}$$

Now, just as in the proof of Lemma 5.10, we may lift a family of couplings which realize the distances $\overline{d}_S(\mu_{S,\omega}, \mu_S)$, and thereby turn the above inequalities into

$$\int \overline{d}_n(\mu_\omega, \mu) P(d\omega) \leq \frac{|S|}{n} \left(\frac{1}{\kappa} \int D(\mu_\omega \parallel \mu) P(d\omega) + r \right) + a.$$

Casually taking r outside the parentheses, this shows $(n\kappa/|S|, r+a)$ -extremality.

In the case of general Y and a kernel μ_\bullet , simply apply the special case above to each μ_y separately. \square

Part II RELATIVE BERNOULLICITY

To prove Theorem A, we construct a factor of (X, μ, T) using a sequence of applications of Theorem C, and then show that the measure concentration promised by that theorem implies relative Bernoullicity. The second of these steps rests on Thouvenot's relative version of Ornstein theory. In this part of the paper we recall the main results of that theory, and provide some small additions which assist in its application later.

At the heart of Ornstein's original work on Bernoullicity are two necessary and sufficient conditions for an ergodic finite-state process to be isomorphic to a Bernoulli shift: the finitely determined property, and the very weak Bernoulli property. More recently, other necessary and sufficient conditions have been added to this list, such as extremality [107, Definition 6.3] or a version of measure concentration [61] (historically referred to as the 'almost blowing-up property' in ergodic theory).

This theory can be generalized to characterize relative Bernoullicity of an ergodic measure-preserving system over a distinguished factor map. This was realized by Thouvenot, who proved the first key extensions of Ornstein's results in [104, 105] (see also [44] for another treatment). Thouvenot's more recent survey [107] describes the broader context. In the first instance, the theory is based on a relative version of the finitely determined property. Thouvenot's theorem that relative finite determination implies relative Bernoullicity is recalled as Theorem 11.5 below.

In general, a proof that a given system is relatively Bernoulli over some factor map has two stages: (i) a proof of relative finite determination, and then (ii) an appeal

to Theorem 11.5. To prove Theorem A, we construct the new factor map π_1 that it promises, and then show relative Bernoullicity over π_1 following those two stages.

Stage (ii) simply amounts to citing Theorem 11.5 from Thouvenot’s work. That theorem is the most substantial fact that we bring into play from outside this paper.

There is more variety among instances of stage (i) in the literature: that is, proofs of relative finite determination itself. For the construction in our proof of Theorem A, we obtain it from a relative version of extremality. Relative extremality is a much more concrete property than finite determination, and is built on the ideas of Section 9.

Regarding these two stages, let us make a curious observation. The assumption of ergodicity in Theorem A is needed only as a necessary hypothesis for Theorem 11.5. We do not use it anywhere in stage (i) of our proof of Theorem A. In particular, because we have chosen to prove Theorem C in general, rather than just for the examples coming from ergodic theory, we make no direct use of the Shannon–McMillan theorem in this paper (see the first remark at the end of Section 8). Ergodicity does play a crucial role within Thouvenot’s proof of Theorem 11.5 via the ergodic and Shannon–McMillan theorems.

This part of the paper has the following outline. Section 10 recalls some standard theory of Kolmogorov–Sinai entropy relative to a factor. Sections 11 and 12 introduce relative finite determination and relative extremality, respectively. Then Section 13 proves the crucial implication from the latter to the former.

The ideas in these sections have all been known in this branch of ergodic theory for many years, at least as folklore. For the sake of completeness, I include full proofs of all but the most basic results that we subsequently use, with the exception of Theorem 11.5. Nevertheless, this part of the paper covers only those pieces of relative Ornstein theory that are needed in Part III, and omits many other topics. For example, a knowledgeable reader will see a version of relative very weak Bernoullicity at work in the proof of Proposition 13.3 below, and will recognize a relative version of ‘almost block independence’ in its conclusion: see the remark following that proof. But we do not introduce these properties formally, and leave aside the implications between them. Relative very weak Bernoullicity already appears in [104], and is known to be equivalent to relative finite determination by results of Thouvenot [104], Rahe [82] and Kieffer [43].

10. Background on relative entropy in ergodic theory

This section contains the most classical facts that we need. We recount some of them carefully for ease of reference later and to establish notation, but we omit most of the proofs.

10.1. Factors, observables, and relative entropy. — For any set A we write T_A for the leftward coordinate shift on $A^{\mathbb{Z}}$. In the sequel the letters A and B always denote finite alphabets.

Most traditional accounts of entropy in ergodic theory are explained in terms of finite measurable partitions. Here we adopt the slightly different language of observables, which is closer to the spirit of coding and information theory (see, for instance, [5, Chapter 5]).

Let (\mathbf{X}, μ, T) be a measure-preserving system. An A -valued **observable** on \mathbf{X} is simply a measurable map $\mathbf{X} \rightarrow A$. An A -valued **process** is a bi-infinite sequence

$$\xi = (\dots, \xi_{-1}, \xi_0, \xi_1, \dots)$$

of observables which together define a factor map from (\mathbf{X}, μ, T) to $(A^{\mathbf{Z}}, \xi_*\mu, T_A)$. We often adopt the point of view that a process is a special kind of factor map. By equivariance, it follows that these maps must satisfy

$$(77) \quad \xi_n(x) = \xi_0(T^n x) \quad \mu\text{-a.s.}$$

Thus, up to agreement modulo μ , ξ is completely specified by the time-zero coordinate ξ_0 . On the other hand, any A -valued observable ξ_0 gives rise to a unique A -valued process: simply use (77) as the definition of all the other ξ_n s.

Given a sequence ξ as above and any subset $F \subseteq \mathbf{Z}$, we let

$$\xi_F := (\xi_n)_{n \in F}.$$

We use this notation mostly when F is a discrete interval. For any integers a and b , we write

$$[a; b] = (a - 1; b] = [a; b + 1) = (a - 1; b + 1) := \{a, a + 1, \dots, b\},$$

interpreting this as \emptyset if $b < a$. We extend this notation to allow $a = -\infty$ or $b = \infty$ in the obvious way.

Given a system (\mathbf{X}, μ, T) and a process ξ , the **finite-dimensional distributions** of ξ are the distributions of the observables ξ_F corresponding to finite subsets $F \subseteq \mathbf{Z}$. In case $\mathbf{X} = A^{\mathbf{Z}}$, $T = T_A$, and ξ is the canonical process, we denote the distribution of ξ_F by μ_F .

The following conditional version of this notation is less standard, but also useful later in the paper.

Definition 10.1 (Block kernels). — Given a system (\mathbf{X}, μ, T) , processes $\xi : \mathbf{X} \rightarrow A^{\mathbf{Z}}$ and $\pi : \mathbf{X} \rightarrow B^{\mathbf{Z}}$, and $n \in \mathbf{N}$, the (ξ, π, n) -**block kernel** is the conditional distribution

$$\mu(\xi_{[0;n]} = \mathbf{a} \mid \pi_{[0;n]} = \mathbf{b}) \quad (\mathbf{a} \in A^n, \mathbf{b} \in B^n),$$

regarded as a kernel from B^n to A^n . We define it arbitrarily for any \mathbf{b} for which $\mu(\pi_{[0;n]} = \mathbf{b}) = 0$.

In case $\mathbf{X} = B^{\mathbf{Z}} \times A^{\mathbf{Z}}$, $T = T_{B \times A}$, and ξ and π are the canonical A - and B -valued processes on \mathbf{X} , we may refer simply to the n -**block kernel**, and denote it by $\mu_{[0;n]}^{\text{block}}(\mathbf{a} \mid \mathbf{b})$.

Now let $\xi : X \rightarrow A^{\mathbf{Z}}$ be a process and let $\pi : (X, \mu, T) \rightarrow (Y, \nu, S)$ be a general factor map. The **relative Kolmogorov–Sinai** (**KS**) **entropy of** (X, μ, T, ξ) **over** π is the quantity

$$h(\xi, \mu, T | \pi) = \lim_{n \rightarrow \infty} \frac{1}{n} H_{\mu}(\xi_{[0;n]} | \pi).$$

The limit exists by the usual subadditivity argument. In case π is trivial it simplifies to

$$h(\xi, \mu, T) = h(\xi_* \mu, T_A),$$

the KS entropy of the factor system $(A^{\mathbf{Z}}, \xi_* \mu, T_A)$, by the Kolmogorov–Sinai theorem.

Henceforth we frequently omit the subscript μ from H_{μ} , provided this measure is clear from the context.

In general, if π and φ are measurable maps on (X, μ) , then we say that φ is **π -measurable** if φ is measurable with respect to the σ -algebra generated by π up to μ -negligible sets.

The next lemma describes a standard alternative approach to relative KS entropy in case the factor map π is also a process.

Lemma 10.2. — *If $\pi : X \rightarrow B^{\mathbf{Z}}$ is a process, then*

$$H(\xi_{[0;n]} | \pi_{[0;n]}) \geq H(\xi_{[0;n]} | \pi)$$

for every n , both of these sequences are subadditive in n , and both have the form

$$h(\xi, \mu, T | \pi) \cdot n + o(n) \quad \text{as } n \rightarrow \infty.$$

If π is an arbitrary factor map, then

$$h(\xi, \mu, T | \pi) = \inf \{ h(\xi, \mu, T | \pi') : \pi' \text{ is a } \pi\text{-measurable process} \}.$$

If φ is another general factor map of (X, μ, T) , then we define

$$h(\varphi, \mu, T | \pi) := \sup \{ h(\xi, \mu, T | \pi) : \xi \text{ is a } \varphi\text{-measurable process} \}.$$

In particular,

$$\begin{aligned} h(\mu, T | \pi) &:= h(\text{id}_X, \mu, T | \pi) \\ &= \sup \{ h(\xi, \mu, T | \pi) : \xi \text{ a process on } (X, \mu, T) \}. \end{aligned}$$

Once again, an important special case arises when $X = B^{\mathbf{Z}} \times A^{\mathbf{Z}}$, $T = T_{B \times A}$, and π is the canonical B -valued process on X . In this case we frequently abbreviate

$$h(\mu, T | \pi) =: h(\mu | \pi).$$

If $\pi : X \rightarrow Y$ and $\varphi : X \rightarrow Z$ are maps between measurable spaces, then we write $\pi \vee \varphi$ for the map

$$X \rightarrow Y \times Z : x \mapsto (\pi(x), \varphi(x)).$$

In the case of observables, this notation deliberately suggests the common refinement of a pair of partitions. We extend it to larger collections of maps or observables in the obvious way.

Relative KS entropy is a dynamical cousin of conditional Shannon entropy, and it has many analogous properties. The next result is an analog of the chain rule for Shannon entropy. It follows at once from that rule, Lemma 10.2, and the definitions above.

Lemma 10.3 (Chain rule). — For any factor maps π , φ and ψ on (X, μ, T) , we have

$$(78) \quad h(\varphi \vee \psi, \mu, T | \pi) = h(\varphi, \mu, T | \pi) + h(\psi, \mu, T | \varphi \vee \pi).$$

Corollary 10.4 (Chain inequality). — For any factor maps π , φ and ψ , we have

$$h(\psi, \mu, T | \varphi \vee \pi) \geq h(\psi, \mu, T | \pi) - h(\varphi, \mu, T | \pi).$$

Proof. — This follows by re-arranging (78) and using the monotonicity

$$h(\varphi \vee \psi, \mu, T | \pi) \geq h(\psi, \mu, T | \pi). \quad \square$$

The next lemma generalizes the well-known equality between the KS entropy of a process and the Shannon entropy of its present conditioned on its past (see, for instance, [5, Equation (12.5)]). The method of proof is the same. That method also appears in the proof of a further generalization in Lemma 10.6 below.

Lemma 10.5. — For any process ξ and factor map π , we have

$$h(\xi, \mu, T | \pi) = H(\xi_0 | \pi \vee \xi_{(-\infty; 0)}).$$

10.2. The use of a periodic set. — At a few points later we need a slightly less standard variation on the calculations above. To formulate it, we say that a measurable subset $F \subseteq X$ is **m -periodic under T modulo μ** for some $m \in \mathbf{N}$ if (i) the sets $F, T^{-1}F, \dots, T^{-(m-1)}F$ form a partition of X modulo μ -negligible sets and (ii) $T^{-m}F = F$ modulo μ .

Lemma 10.6. — If π is a factor map of (X, μ, T) , ξ is a process, and F is π -measurable and m -periodic under T modulo μ , then

$$H(\xi_{[0; m]} | \pi; F) \geq H(\xi_{[0; m]} | \pi \vee \xi_{(-\infty; 0)}; F) = h(\xi, \mu, T | \pi) \cdot m.$$

(Recall (7) for the notation $H(\cdot | \cdot ; F)$.)

Proof. — The inequality on the left follows from the monotonicity of H under conditioning, so we focus on the equality on the right. In case $m = 1$, this is precisely the statement of Lemma 10.5. To handle the general case we simply mimic the usual proof of that special case.

First, for any $k \in \mathbf{N}$, the chain rule for Shannon entropy gives

$$(79) \quad H(\xi_{[0;km]} | \pi ; F) = \sum_{i=0}^{k-1} H(\xi_{[im;(i+1)m]} | \pi \vee \xi_{[0;im]} ; F).$$

For each i in this sum we have $T^{-im}F = F$ modulo μ , because F is m -periodic. Since the σ -algebra generated by π is globally T -invariant, the conditional distribution of $\xi_{[im;(i+1)m]}$ given $\pi \vee \xi_{[0;im]}$ is the same as the conditional distribution of $\xi_{[0;m]}$ given $\pi \vee \xi_{[-im;0]}$, up to shifting by T . Therefore the right-hand side of (79) is equal to

$$\sum_{i=0}^{k-1} H(\xi_{[0;m]} | \pi \vee \xi_{[-im;0]} ; F).$$

By the usual appeal to the martingale convergence theorem (see, for instance, [5, Theorem 12.1]), this sum is equal to

$$(80) \quad k \cdot H(\xi_{[0;m]} | \pi \vee \xi_{(-\infty;0]} ; F) + o(k) \quad \text{as } k \longrightarrow \infty.$$

On the other hand, consider some $t \in \{1, \dots, m-1\}$. Then, reasoning again from the T -invariance of the σ -algebra generated by π , we have

$$\begin{aligned} H(\xi_{[0;km]} | \pi ; T^{-t}F) &= H(\xi_{[t;t+km]} | \pi ; F) \\ &= H(\xi_{[0;km]} | \pi ; F) + O(t \log |A|) = H(\xi_{[0;km]} | \pi ; F) + O(m \log |A|), \end{aligned}$$

since each of the observables $\xi_{[t;t+km]}$ and $\xi_{[0;km]}$ determines all but t coordinates of the other. Therefore

$$\begin{aligned} H(\xi_{[0;km]} | \pi) &= \frac{1}{m} \sum_{t=0}^{m-1} H(\xi_{[0;km]} | \pi ; T^{-t}F) \\ &= H(\xi_{[0;km]} | \pi ; F) + O(m \log |A|), \end{aligned}$$

where the first equality holds because the partition $F, T^{-1}F, \dots, T^{-(m-1)}F$ is π -measurable and all these sets have measure $1/m$. Combining this with (79) and (80), dividing by k , and then letting $k \longrightarrow \infty$, we obtain

$$\begin{aligned} \mathbf{H}(\xi_{[0;m]} | \pi \vee \xi_{(-\infty;0)}; \mathbf{F}) &= \lim_{k \rightarrow \infty} \frac{1}{k} \mathbf{H}(\xi_{[0;km]} | \pi; \mathbf{F}) \\ &= \lim_{k \rightarrow \infty} \frac{1}{k} \mathbf{H}(\xi_{[0;km]} | \pi) = h(\xi, \mu, \mathbf{T} | \pi). \end{aligned} \quad \square$$

Remark. — Lemma 10.6 is really a relative-entropy version of Abramov’s formula for the entropy of an induced transformation [1], in the special case where we induce on a periodic set.

11. Relative finite determination

11.1. Extensions, joinings, and relative $\bar{\mathbf{d}}$ -distance. — Before we define relative finite determination, we need some more notation in the setting of finite-alphabet processes. Let \mathbf{A} and \mathbf{B} be finite sets, and let β be the canonical process from $\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}$ to $\mathbf{B}^{\mathbf{Z}}$. Let ν be a shift-invariant probability measure on $\mathbf{B}^{\mathbf{Z}}$, and let $\text{Ext}(\nu, \mathbf{A})$ be the set of all shift-invariant probability measures λ on $\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}$ which satisfy $\beta_* \lambda = \nu$.

Given $\lambda, \theta \in \text{Ext}(\nu, \mathbf{A})$, let

$$\text{Join}(\lambda, \theta | \beta)$$

be the set of all invariant probability measures γ on $\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}$ which satisfy

$$\gamma\{(b, a, a') : (b, a) \in \mathbf{U}\} = \lambda(\mathbf{U}) \quad \text{and} \quad \gamma\{(b, a, a') : (b, a') \in \mathbf{U}\} = \theta(\mathbf{U})$$

for all measurable $\mathbf{U} \subseteq \mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}$. Members of $\text{Join}(\lambda, \theta | \beta)$ are called **relative joinings of λ and θ over β** .

The set $\text{Join}(\lambda, \theta | \beta)$ always contains the relative product of λ and θ over the map β , so this set is never empty. If λ and θ are both ergodic, then $\text{Join}(\lambda, \theta | \beta)$ also contains all ergodic components of that relative product, so in particular it has some ergodic members.

Definition 11.1. — In the setting above, the **relative $\bar{\mathbf{d}}$ -distance between λ and θ over β** is the quantity

$$\bar{\mathbf{d}}(\lambda, \theta | \beta) = \inf \left\{ \gamma\{(b, a, a') : a_0 \neq a'_0\} : \gamma \in \text{Join}(\lambda, \theta | \beta) \right\}.$$

A standard vague-compactness argument in the space $\text{Join}(\lambda, \theta | \beta)$ shows that this infimum is achieved by some γ .

In order to estimate $\bar{\mathbf{d}}(\lambda, \theta | \beta)$, it is helpful to relate it to transportation distances and the finite-dimensional distributions $\lambda_{[0;n]}$ and $\theta_{[0;n]}$. We may do this in terms of the associated block kernels from Definition 10.1.

Lemma 11.2. — For any $\lambda, \theta \in \text{Ext}(v, A)$ we have

$$\bar{d}(\lambda, \theta | \beta) = \lim_{n \rightarrow \infty} \int \bar{d}_n \left(\lambda_{[0;n]}^{\text{block}}(\cdot | b_{[0;n]}), \theta_{[0;n]}^{\text{block}}(\cdot | b_{[0;n]}) \right) v(db),$$

Proof. — This is a relative version of a well-known formula for Ornstein's original \bar{d} -metric: see, for instance, [91, Theorem I.9.7]. The existence of the right-hand limit is part of the conclusion to be proved.

We see easily that $\bar{d}(\lambda, \theta | \beta)$ is at least the limit supremum of the right-hand side: if γ is any element of $\text{Join}(\lambda, \theta | \beta)$ which achieves the value of $\bar{d}(\lambda, \theta | \beta)$, then its finite-dimensional distributions $\gamma_{[0;n]}$ give an upper bound for that limit supremum. So now let δ be the limit infimum of the right-hand side, and let us show that $\bar{d}(\lambda, \theta | \beta) \leq \delta$. Suppose δ is the true limit along the subsequence indexed by $n_1 < n_2 < \dots$.

For each $n \in \mathbf{N}$ and $\mathbf{b} \in B^n$, select an optimal coupling of $\lambda_{[0;n]}^{\text{block}}(\cdot | \mathbf{b})$ and $\theta_{[0;n]}^{\text{block}}(\cdot | \mathbf{b})$. These couplings together define a kernel from B^n to $A^n \times A^n$ for each n . Hooking these kernels up to the measures $v_{[0;n]}$, we obtain a sequence of probability measures γ_n on $B^n \times A^n \times A^n$ such that γ_n couples $\lambda_{[0;n]}$ and $\theta_{[0;n]}$ over a common copy of B^n and such that

$$\begin{aligned} (81) \quad & \int_{B^{n_k} \times A^{n_k} \times A^{n_k}} d_{n_k}(\mathbf{a}, \mathbf{a}') \gamma_{n_k}(d\mathbf{b}, d\mathbf{a}, d\mathbf{a}') \\ &= \frac{1}{n_k} \sum_{i=1}^{n_k} \gamma_{n_k} \{(\mathbf{b}, \mathbf{a}, \mathbf{a}') : a_i \neq a'_i\} \longrightarrow \delta \quad \text{as } k \longrightarrow \infty. \end{aligned}$$

For each n , let γ'_n be any probability measure on $B^{\mathbf{Z}} \times A^{\mathbf{Z}} \times A^{\mathbf{Z}}$ whose projection to $B^{[0;n]} \times A^{[0;n]} \times A^{[0;n]}$ agrees with γ_n , and let γ''_n be the following averaged measure on $B^{\mathbf{Z}} \times A^{\mathbf{Z}} \times A^{\mathbf{Z}}$:

$$\gamma''_n := \frac{1}{n} \sum_{t=0}^{n-1} (\mathbb{T}_{B \times A \times A}^t)_* \gamma'_n.$$

Because of this averaging, the bound (81) implies that

$$(82) \quad \gamma''_{n_k} \{(b, a, a') : a_0 \neq a'_0\} = \frac{1}{n_k} \sum_{i=0}^{n_k-1} \gamma'_{n_k} \{(b, a, a') : a_i \neq a'_i\} \longrightarrow \delta \quad \text{as } k \longrightarrow \infty.$$

Finally, let γ'' be any subsequential vague limit of the sequence $(\gamma''_{n_k})_{k=1}^{\infty}$. Then γ'' is invariant, because the averaged measures γ''_n satisfy

$$\|\gamma''_n - (\mathbb{T}_{B \times A \times A})_* \gamma''_n\| = O(1/n).$$

Taking the limit in (82), we have

$$\gamma'' \{(b, a, a') : a_0 \neq a'_0\} = \delta.$$

It remains to show that γ'' is a coupling of λ and θ over β . We show that the first projection of γ'' to $\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}$ is λ ; the argument for the second projection is the same.

Fix $\ell \in \mathbf{N}$ and two strings $\mathbf{b} \in \mathbf{B}^\ell$ and $\mathbf{a} \in \mathbf{A}^\ell$, and let

$$\mathbf{Y} := \{(b, a) \in \mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}} : b_{[0;\ell]} = \mathbf{b} \text{ and } a_{[0;\ell]} = \mathbf{a}\}.$$

For each n , the definition of γ_n'' gives

$$(83) \quad \gamma_n''(\mathbf{Y} \times \mathbf{A}^{\mathbf{Z}}) = \frac{1}{n} \sum_{t=0}^{n-1} \gamma_n'(\mathbb{T}_{\mathbf{B} \times \mathbf{A} \times \mathbf{A}}^{-t}(\mathbf{Y} \times \mathbf{A}^{\mathbf{Z}})).$$

For $0 \leq t \leq n - \ell$, the set

$$\mathbb{T}_{\mathbf{B} \times \mathbf{A} \times \mathbf{A}}^{-t}(\mathbf{Y} \times \mathbf{A}^{\mathbf{Z}}) = \{(b, a, a') : b_{[t;t+\ell]} = \mathbf{b} \text{ and } a_{[t;t+\ell]} = \mathbf{a}\}$$

depends only on coordinates with indices in $[0; n)$, and so we have

$$\begin{aligned} \gamma_n'(\mathbb{T}_{\mathbf{B} \times \mathbf{A} \times \mathbf{A}}^{-t}(\mathbf{Y} \times \mathbf{A}^{\mathbf{Z}})) &= \lambda_{[0;n]} \{(b_0, \dots, b_{n-1}, a_0, \dots, a_{n-1}) : \\ &\quad (b_t, \dots, b_{t+\ell-1}) = \mathbf{b} \text{ and } (a_t, \dots, a_{t+\ell-1}) = \mathbf{a}\} \\ &= \lambda(\mathbf{Y}). \end{aligned}$$

Inserting this into (83), we obtain

$$\frac{1}{n} \sum_{t=0}^{n-\ell} \lambda(\mathbf{Y}) + \mathcal{O}(\ell/n) = \lambda(\mathbf{Y}) + \mathcal{O}(\ell/n).$$

Taking $n = n_k$ here and letting $k \rightarrow \infty$, we conclude that $\gamma''(\mathbf{Y} \times \mathbf{A}^{\mathbf{Z}}) = \lambda(\mathbf{Y})$ for any such choice of ℓ , \mathbf{b} and \mathbf{a} . Therefore the first projection of γ'' to $\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}$ is equal to λ . \square

It follows from Lemma 11.2 that the function $\bar{d}(\cdot, \cdot | \beta)$ is indeed a metric on $\text{Ext}(v, \mathbf{A})$, but we do not use this particular fact later in the paper. See [91, Section I.9] for the analogous proof in the non-relative setting.

11.2. Relative finite determination. — The relative version of the finitely determined property has been studied in various papers, particularly [104], [82] and [44]. Here we essentially follow the ideas of [104], except that we also define a quantitative version of this property, in which various small parameters are fixed rather than being subject to a universal or existential quantifier. We do this because several arguments later in the paper require that we keep track of these parameters explicitly.

We define relative finite determination in two steps: for shift-invariant measures, and then for processes. Suppose first that $\nu \in \text{Prob}(\mathbf{B}^{\mathbf{Z}})$ is shift-invariant, and let β be the canonical \mathbf{B} -valued process on $\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}$.

Definition 11.3. — Let $\delta > 0$. A measure $\lambda \in \text{Ext}(\nu, \mathbb{A})$ is **relatively δ -finitely determined** (**relatively δ -FD**) **over** β if there exist $\varepsilon > 0$ and $n \in \mathbf{N}$ for which the following holds: if another measure $\theta \in \text{Ext}(\nu, \mathbb{A})$ satisfies

- (i) $h(\theta | \beta) > h(\lambda | \beta) - \varepsilon$ and
- (ii) $\|\theta_{[0;n]} - \lambda_{[0;n]}\| < \varepsilon$,

then $\bar{d}(\lambda, \theta | \beta) < \delta$.

The measure λ is **relatively FD over** β if it is relatively δ -FD over β for every $\delta > 0$.

Now consider a general system (X, μ, T) and a pair of processes $\pi : X \rightarrow B^{\mathbf{Z}}$ and $\xi : X \rightarrow A^{\mathbf{Z}}$.

Definition 11.4. — The process ξ is **relatively δ -FD** (resp. **relatively FD**) **over** π if the joint distribution

$$\lambda = (\pi \vee \xi)_* \mu$$

is relatively δ -FD (resp. relatively FD) over β .

The part of this definition which quantifies over all δ agrees with the definition in [44, 104].

The following result is the heart of Thouvenot's original work [104].

Theorem 11.5. — If the system $(B^{\mathbf{Z}} \times A^{\mathbf{Z}}, \lambda, T_{B \times A})$ is ergodic and λ is relatively FD over β , then $(B^{\mathbf{Z}} \times A^{\mathbf{Z}}, \lambda, T_{B \times A})$ is relatively Bernoulli over β .

We use this theorem as a 'black box' in Part III. It is the most significant result that we cite from outside this paper.

Lemma 11.6. — Let (X, μ, T) be a system, let

$$\xi : X \rightarrow A^{\mathbf{Z}}, \quad \pi : X \rightarrow B^{\mathbf{Z}} \quad \text{and} \quad \pi' : X \rightarrow (B')^{\mathbf{Z}}$$

be three processes, and assume that π and π' generate the same factor of (X, μ, T) modulo μ -negligible sets. If ξ is relatively δ -FD over π , then it is also relatively δ -FD over π' .

Proof. — Since π and π' generate the same σ -subalgebra of \mathcal{B}_X modulo negligible sets, and since our measurable spaces are all standard, there is a commutative diagram of factor maps

$$\begin{array}{ccc} & (X, \mu, T) & \\ \pi \swarrow & & \searrow \pi' \\ (B^{\mathbf{Z}}, \nu, T_B) & \xleftarrow{\varphi} & ((B')^{\mathbf{Z}}, \nu', T_{B'}) \end{array}$$

where φ is an isomorphism of systems. (See, for instance, the subsection on isomorphism vs. conjugacy in [5, Section 5].)

Let

$$\lambda := (\pi \vee \xi)_* \mu \quad \text{and} \quad \lambda' := (\pi' \vee \xi)_* \mu.$$

Let β be the coordinate projection $\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}} \longrightarrow \mathbf{B}^{\mathbf{Z}}$, and define β' analogously. Let α be the canonical process on $\mathbf{A}^{\mathbf{Z}}$, so that α_F is the coordinate projection $\mathbf{A}^{\mathbf{Z}} \longrightarrow A^F$ for any $F \subseteq \mathbf{Z}$.

For the given value of δ and for the processes ξ and π , let $\varepsilon > 0$ and $n \in \mathbf{N}$ be the values promised by Definition 11.4.

Let $\varepsilon' := \varepsilon/3$. Since $\varphi_{[0;n]} : (\mathbf{B}')^{\mathbf{Z}} \longrightarrow \mathbf{B}^{[0;n]}$ is measurable, there exist $m \in \mathbf{N}$, a map $\Phi : (\mathbf{B}')^{\mathbf{Z}} \longrightarrow \mathbf{B}^{[0;n]}$ that depends only on coordinates in $[-m; n+m)$, and a measurable subset $U \subseteq (\mathbf{B}')^{\mathbf{Z}}$ with $\nu'(U) > 1 - \varepsilon/6$ such that

$$(84) \quad \varphi_{[0;n]}|U = \Phi|U.$$

Let $n' := n + 2m$, and let $U^c := (\mathbf{B}')^{\mathbf{Z}} \setminus U$.

Now suppose that $\theta' \in \text{Ext}(\nu', \mathbf{A})$ satisfies

- (i') $h(\theta' | \beta') > h(\lambda' | \beta') - \varepsilon'$ and
- (ii') $\|\theta'_{[0;n']} - \lambda'_{[0;n']}\| < \varepsilon'$, hence also

$$(85) \quad \|\theta'_{[-m;n+m]} - \lambda'_{[-m;n+m]}\| < \varepsilon'$$

by shift-invariance.

Let $\theta := (\varphi \times \alpha)_* \theta' \in \text{Prob}(\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}})$. We shall deduce from (i') and (ii') that θ and λ satisfy the analogous conditions (i) and (ii) in Definition 11.3.

Towards (i), the chain rule and the Kolmogorov–Sinai theorem give that

$$h(\lambda | \beta) = h(\lambda) - h(\nu) = h(\lambda') - h(\nu') = h(\lambda' | \beta'),$$

and similarly $h(\theta | \beta) = h(\theta' | \beta')$. Therefore, by (i'),

$$h(\theta | \beta) > h(\lambda | \beta) - \varepsilon' > h(\lambda | \beta) - \varepsilon.$$

To verify (ii), consider the estimates

$$(86) \quad \begin{aligned} \|\theta_{[0;n]} - \lambda_{[0;n]}\| &= \|(\varphi_{[0;n]} \times \alpha_{[0;n]})_*(\theta' - \lambda')\| \\ &\leq \|(\varphi_{[0;n]} \times \alpha_{[0;n]})_*(1_{U^c \times \mathbf{A}^{\mathbf{Z}}} \cdot \theta')\| \\ &\quad + \|(\varphi_{[0;n]} \times \alpha_{[0;n]})_*(1_{U \times \mathbf{A}^{\mathbf{Z}}} \cdot (\theta' - \lambda'))\| \\ &\quad + \|(\varphi_{[0;n]} \times \alpha_{[0;n]})_*(1_{U^c \times \mathbf{A}^{\mathbf{Z}}} \cdot \lambda')\|. \end{aligned}$$

Since $\theta'(U \times A^{\mathbf{Z}}) = \lambda'(U \times A^{\mathbf{Z}}) = \nu'(U)$, the first and last terms here are both less than $\varepsilon/6$. On the other hand, by (84), the middle term is equal to

$$\|(\Phi \times \alpha_{[0;n]})_*(1_{U \times A^{\mathbf{Z}}} \cdot (\theta' - \lambda'))\|,$$

which is at most

$$(87) \quad \|(\Phi \times \alpha_{[0;n]})_*(\theta' - \lambda')\| + 2 \cdot (\varepsilon/6)$$

by the same reasoning that gave (86). Since Φ depends only on coordinates in $[-m; n+m)$, the first term of (87) is bounded by the left-hand side of (85), hence is at most $\varepsilon/3$. Adding up these estimates, we obtain condition (ii):

$$\|\theta_{[0;n]} - \lambda_{[0;n]}\| < \varepsilon.$$

Having shown conditions (i) and (ii) for θ and λ , our choice of ε and n gives $\bar{d}(\lambda, \theta | \beta) < \delta$. Let γ be a triple joining that witnesses this inequality, and let

$$\gamma' := (\varphi^{-1} \times \alpha \times \alpha)_* \gamma.$$

This new triple joining witnesses that $\bar{d}(\lambda', \theta' | \beta') < \delta$. □

12. Relative extremality

Extremality is one of the characterizations of Bernoullicity for a finite-state ergodic process. It was introduced by Thouvenot in the mid 1970s, and began to circulate in unpublished notes by Feldman and in Antoine Lamotte's PhD thesis, also unpublished. The earliest published reference I know is [84], where Rudolph uses this characterization of Bernoullicity without proof. Ornstein and Weiss introduce extremality and prove its equivalence to Bernoullicity in the general setting of amenable groups in [77, Section III.4], and use it for their extension of Ornstein theory to that setting. Extremality is also defined and compared to other characterizations of Bernoullicity in [107, Definition 6.3], and a complete account of its place in the theory for a single automorphism appears in [36, Chapter 5].

In this paper we need a relative version of extremality for one process over another. We define this as a natural sequel to Definition 9.9. The resulting definition is superficially a little different from Thouvenot's, but equivalent in all important respects.

Consider again a general system (X, μ, T) and a pair of processes $\xi : X \rightarrow A^{\mathbf{Z}}$ and $\pi : X \rightarrow B^{\mathbf{Z}}$.

Definition 12.1. — Let $r, \kappa > 0$. The process ξ is **relatively** (κ, r) -**extremal over** π if the map $\xi_{[0;n]}$ is $(\kappa n, r)$ -extremal over $\pi_{[0;n]}$ on the probability space (X, μ) for all sufficiently large n .

This definition was the motivation for our work in Section 9. Unpacking Definition 9.9, we can write Definition 12.1 more explicitly as follows. Let $\nu := \pi_*\mu$, let $\lambda := (\pi \vee \xi)_*\mu$, and let $\lambda_{[0;n]}^{\text{block}}$ be the (ξ, π, n) -block kernel from Definition 10.1. The process ξ is (κ, r) -extremal over π if, for every sufficiently large n , there is a real-valued function $\mathbf{b} \mapsto r_{\mathbf{b}}$ on \mathbf{B}^n such that

- the conditional measure $\lambda_{[0;n]}^{\text{block}}(\cdot | \mathbf{b})$ is a $(\kappa n, r_{\mathbf{b}})$ -extremal measure on the metric space (A^n, d_n) for each $\mathbf{b} \in \mathbf{B}^n$, and
- we have

$$\int r_{\mathbf{b}} \nu_{[0;n]}(d\mathbf{b}) \leq r.$$

Extremality is the notion that forms the key link between measure concentration as studied in Part I and relative Bernoullicity. As such, it is the backbone of this whole paper.

The next lemma is a dynamical version of Corollary 9.11. It is used during the proof of Theorem A in Subsection 15.1. During that proof, we need to construct a new observable with respect to which another is fairly extremal, and then enlarge that new observable a little further without losing too much extremality.

Lemma 12.2. — *Let ξ and π be processes on $(\mathbf{X}, \mu, \mathbf{T})$ and suppose that ξ is relatively (κ, r) -extremal over π . Let π' be another process such that π'_0 refines π_0 , and assume that*

$$(88) \quad h(\xi, \mu, \mathbf{T} | \pi) - h(\xi, \mu, \mathbf{T} | \pi') \leq \kappa r.$$

Then ξ is relatively $(\kappa, 6r)$ -extremal over π' .

Proof. — Let n be large enough that

- $\xi_{[0;n]}$ is $(\kappa n, r)$ -extremal over $\pi_{[0;n]}$, and
- $H(\xi_{[0;n]} | \pi_{[0;n]}) \leq (h(\xi, \mu, \mathbf{T} | \pi) + \kappa r/2) \cdot n$.

Then Lemma 10.2 and the assumed bound (88) give

$$\begin{aligned} & H(\xi_{[0;n]} | \pi_{[0;n]}) - H(\xi_{[0;n]} | \pi'_{[0;n]}) \\ & \leq (h(\xi, \mu, \mathbf{T} | \pi) - h(\xi, \mu, \mathbf{T} | \pi') + \kappa r/2) \cdot n \leq 3\kappa r n/2. \end{aligned}$$

Now apply Corollary 9.11 to $\xi_{[0;n]}$, $\pi_{[0;n]}$ and $\pi'_{[0;n]}$. It gives that $\xi_{[0;n]}$ is extremal over $\pi'_{[0;n]}$ with parameters

$$\left(\kappa n, \frac{2 \cdot (3\kappa r n/2)}{\kappa n} + 3r \right) = (\kappa n, 6r). \quad \square$$

13. From extremality to finite determination

Consider again a system (X, μ, T) and processes $\xi : X \rightarrow A^{\mathbb{Z}}$ and $\pi : X \rightarrow B^{\mathbb{Z}}$. The main result of this section is that extremality implies finite determination, with some explicit dependence between the quantitative versions of these properties.

Proposition 13.1. — *Let $r > 0$. If ξ is relatively (κ, r) -extremal over π for some $\kappa > 0$, then ξ is relatively $(7r)$ -FD over π .*

The constant 7 that appears here is convenient, but certainly not optimal.

13.1. Comparison with a product of block kernels. — Let λ be a measure on $B^{\mathbb{Z}} \times A^{\mathbb{Z}}$, and let ν be its marginal on $B^{\mathbb{Z}}$. Let α and β be the canonical A - and B -valued processes on $B^{\mathbb{Z}} \times A^{\mathbb{Z}}$, respectively. In this subsection we write H for H_λ .

The key to Proposition 13.1 is a comparison between the block kernel $\lambda_{[0;kn]}^{\text{block}}$ and the product of k copies of $\lambda_{[0;n]}^{\text{block}}$. This is given in Proposition 13.3. This comparison is also used for another purpose later in the paper: see the proof of Lemma 14.6. For that later application, we need a little extra generality: we can assume that λ is invariant only under the n -fold shift $T_{B \times A}^n$, not necessarily under the single shift $T_{B \times A}$. We can still define the block kernels $\lambda_{[0;kn]}^{\text{block}}$ as in Definition 10.1. We return to the setting of true shift-invariance in Subsection 13.2, where we apply Proposition 13.3 with $\lambda := (\pi \vee \xi)_* \mu$ to prove Proposition 13.1.

So now assume that λ is $T_{B \times A}^n$ -invariant, but not necessarily $T_{B \times A}$ -invariant.

Lemma 13.2. — *We have*

$$H(\alpha_{[0;n]} | \beta \vee \alpha_{(-\infty;0)}) = h(\alpha_{[0;n]}, \lambda, T_{B \times A}^n | \beta).$$

Proof. — This just requires the right point of view. The triple $(B^{\mathbb{Z}} \times A^{\mathbb{Z}}, \lambda, T_{B \times A}^n)$ is a measure-preserving system. In this system, we may regard the map $\alpha_{[0;n]}$ as a single observable taking values in A^n . Under the transformation $T_{B \times A}^n$, this observable generates the process

$$(\dots, \alpha_{[-2n;n]}, \alpha_{[n;0]}, \alpha_{[0;n]}, \alpha_{[n;2n]}, \dots).$$

This is just a copy of α , divided into consecutive blocks of n letters which we now regard as single letters in the enlarged alphabet A^n . We may identify β with a process generated by $\beta_{[0;n]}$ in the same way.

Now apply Lemma 10.5 to this system and pair of processes. □

Proposition 13.3. — *In the situation above, assume that*

- (i) (*extremality*) $\alpha_{[0;n]}$ is relatively $(\kappa n, r)$ -extremal over $\beta_{[0;n]}$ according to λ , and

(ii) (conditional Shannon entropy is close to its Kolmogorov–Sinai limit)

$$H(\alpha_{[0;n]} | \beta_{[0;n]}) < h(\alpha_{[0;n]}, \lambda, T_{\mathbf{B} \times \mathbf{A}}^n | \beta) + \kappa rn.$$

Then

$$\int \overline{d}_{kn} \left(\lambda_{[0;kn]}^{\text{block}}(\cdot | b_{[0;kn]}), \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[0;n]}) \times \cdots \times \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[(k-1)n;kn]}) \right) \nu(db) < 2r$$

for every $k \in \mathbf{N}$. Here \overline{d}_{kn} is the transportation metric associated to $(\mathbf{A}^{kn}, d_{kn})$.

If λ is actually shift-invariant, then hypothesis (ii) of Proposition 13.3 may be rewritten as

$$H(\alpha_{[0;n]} | \beta_{[0;n]}) < (h(\lambda, T_{\mathbf{B} \times \mathbf{A}} | \beta) + \kappa r) \cdot n.$$

Therefore, in this special case, hypothesis (ii) holds for all sufficiently large n by Lemma 10.2.

Proof. — We prove this by induction on k , but first we need to formulate a slightly more general inductive hypotheses. Suppose that $k \in \mathbf{N}$ and that \mathbf{S} is any subset of \mathbf{Z} . Then we write

$$\mathbf{B}^{\mathbf{S}} \longrightarrow \text{Prob}(\mathbf{A}^{kn}) : \mathbf{b} \mapsto \lambda_{[0;kn]|\mathbf{S}}^{\text{block}}(\cdot | \mathbf{b})$$

for a version of the conditional distribution of $\alpha_{[0;kn]}$ given $\beta_{\mathbf{S}}$ under λ . In this notation we have $\lambda_{[0;kn]}^{\text{block}} = \lambda_{[0;kn]|\{0;kn\}}^{\text{block}}$. We show by induction on k that

$$(89) \quad \int \overline{d}_{kn} \left(\lambda_{[0;kn]|\mathbf{S}}^{\text{block}}(\cdot | b_{\mathbf{S}}), \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[0;n]}) \times \cdots \times \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[(k-1)n;kn]}) \right) \nu(db) < 2r$$

whenever $k \in \mathbf{N}$ and $\mathbf{S} \supseteq [0; kn]$. The case $\mathbf{S} = [0; kn]$ gives Proposition 13.3.

The proof of (89) when $k = 1$ is a slightly degenerate version of the proof when $k > 1$, so we explain them together.

Thus, assume we already know (89) with $k - 1$ in place of k and whenever $\mathbf{S} \supseteq [0; (k - 1)n]$. If $k = 1$, then regard this inductive assumption as vacuous. To extend the desired conclusion to k we use Lemma 4.3. Assume now that $\mathbf{S} \supseteq [0; kn]$.

For each $\mathbf{b} \in \mathbf{B}^{\mathbf{S}}$, we have a probability measure $\lambda_{[0;kn]|\mathbf{S}}^{\text{block}}(\cdot | \mathbf{b})$ on

$$\mathbf{A}^{kn} = \underbrace{\mathbf{A}^n \times \cdots \times \mathbf{A}^n}_k.$$

The projection of this measure onto the first $k - 1$ copies of \mathbf{A}^n is $\lambda_{[0;(k-1)n]|\mathbf{S}}^{\text{block}}(\cdot | \mathbf{b})$. This is where we need the additional flexibility that comes from choosing \mathbf{S} separately: the

projection of $\lambda_{[0;kn]}^{\text{block}}(\cdot | \mathbf{b})$ is $\lambda_{[0;(k-1)n] | [0;kn]}^{\text{block}}(\cdot | \mathbf{b})$, and in general this can be different from $\lambda_{[0;(k-1)n]}^{\text{block}}(\cdot | \mathbf{b}_{[0;(k-1)n]})$.

We now disintegrate the measure $\lambda_{[0;kn] | \mathcal{S}}^{\text{block}}(\cdot | \mathbf{b})$ further: for each $\mathbf{b} \in \mathcal{B}^{\mathcal{S}}$, let

$$A^{(k-1)n} \longrightarrow \text{Prob}(A^n) : \mathbf{a} \mapsto \lambda'(\cdot | \mathbf{b}, \mathbf{a})$$

be a conditional distribution under $\lambda_{[0;kn] | \mathcal{S}}^{\text{block}}(\cdot | \mathbf{b})$ of the last n coordinates in A^{kn} , given that the first $(k-1)n$ coordinates agree with \mathbf{a} . If $k=1$, then simply ignore \mathbf{a} and set $\lambda'(\cdot | \mathbf{b}) := \lambda_{[0;n] | \mathcal{S}}^{\text{block}}(\cdot | \mathbf{b})$.

In terms of this further disintegration, we may apply Lemma 4.3 to obtain

$$\begin{aligned} & \int \overline{d}_{kn} \left(\lambda_{[0;kn] | \mathcal{S}}^{\text{block}}(\cdot | b_{\mathcal{S}}), \bigotimes_{j=0}^{k-1} \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}) \right) \nu(db) \\ & \leq \frac{k-1}{k} \int \overline{d}_{(k-1)n} \left(\lambda_{[0;(k-1)n] | \mathcal{S}}^{\text{block}}(\cdot | b_{\mathcal{S}}), \bigotimes_{j=0}^{k-2} \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}) \right) \nu(db) \\ & \quad + \frac{1}{k} \iint \overline{d}_n(\lambda'(\cdot | b_{\mathcal{S}}, \mathbf{a}), \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[(k-1)n;kn]})) \lambda_{[0;(k-1)n] | \mathcal{S}}^{\text{block}}(d\mathbf{a} | b_{\mathcal{S}}) \nu(db). \end{aligned}$$

If $k=1$, then simply ignore the first term on the right of this inequality.

By the inductive hypothesis, the first right-hand term here is less than $2(k-1)r/k$. To finish the proof, we show that

$$\iint \overline{d}_n(\lambda'(\cdot | b_{\mathcal{S}}, \mathbf{a}), \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[(k-1)n;kn]})) \lambda_{[0;(k-1)n] | \mathcal{S}}^{\text{block}}(d\mathbf{a} | b_{\mathcal{S}}) \nu(db) < 2r.$$

The two kernels appearing inside the transportation distance here,

$$\lambda'(\cdot | b_{\mathcal{S}}, \mathbf{a}) \quad \text{and} \quad \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[(k-1)n;kn]}),$$

are conditional distributions for $\alpha_{[(k-1)n;kn]}$ under λ given the σ -algebras

$$\mathcal{H} \text{ generated by } \beta_{\mathcal{S}} \vee \alpha_{[0;(k-1)n]} \quad \text{and} \quad \mathcal{G} \text{ generated by } \beta_{[(k-1)n;kn]},$$

respectively.

Clearly $\mathcal{H} \supseteq \mathcal{G}$. Therefore, by Lemma 9.10, the desired estimate follows from the $(\kappa n, r)$ -extremality of $\alpha_{[0;n]}$ over $\beta_{[0;n]}$, provided we show that

$$\mathbf{H}(\alpha_{[(k-1)n;kn]} | \mathcal{H}) > \mathbf{H}(\alpha_{[(k-1)n;kn]} | \mathcal{G}) - \kappa r n.$$

This, in turn, holds because of assumption (ii) in the statement of the current proposition:

$$\begin{aligned}
\mathbf{H}(\alpha_{[(k-1)n;kn]} \mid \mathcal{H}) &= \mathbf{H}(\alpha_{[(k-1)n;kn]} \mid \beta_S \vee \alpha_{[0;(k-1)n]}) \\
&\geq \mathbf{H}(\alpha_{[(k-1)n;kn]} \mid \beta \vee \alpha_{(-\infty;(k-1)n]}) \\
&= \mathbf{h}(\alpha_{[0;n]}, \lambda, \mathbf{T}_{\mathbf{B} \times \mathbf{A}}^n \mid \beta) \\
&> \mathbf{H}(\alpha_{[(k-1)n;kn]} \mid \mathcal{G}) - \kappa rn,
\end{aligned}$$

where the third line is obtained from Lemma 13.2 and invariance under $\mathbf{T}_{\mathbf{B} \times \mathbf{A}}^n$, and the fourth is our appeal to assumption (ii).

This continues the induction, and hence completes the proof. \square

Remark. — Proposition 13.3 may be seen as a relative version of one of the standard implications of Ornstein theory: that extremality implies almost block independence. We do not formulate a precise relative version of almost block independence in this paper, but this could easily be done along the same lines as the non-relative version: see, for instance, [91, Section IV.1] or [36, Definitions 473 and 476] (the latter reference calls this the ‘independent concatenations’ property). Then the conclusion of Proposition 13.3 should imply that λ is $(2r)$ -almost block independent. I expect this leads to another characterization of relative Bernoullicity for this λ , but we do not pursue this idea here.

In the next subsection we use Proposition 13.3 to prove Proposition 13.1. At another point later we also need a different consequence of Proposition 13.3:

Corollary 13.4 (Stretching property of extremality). — Assume hypotheses (i) and (ii) from the statement of Proposition 13.3. Then the observable $\alpha_{[0;kn]}$ is relatively $(\kappa kn, 5r)$ -extremal over $\beta_{[0;kn]}$ for all $k \in \mathbf{N}$.

Proof. — Let us write a general element of \mathbf{B}^{kn} as $(\mathbf{b}_1, \dots, \mathbf{b}_k)$, where $\mathbf{b}_j \in \mathbf{B}^n$ for each $j = 1, 2, \dots, k$. A k -fold application of Lemma 9.8 gives that the measure $\nu_{[0;kn]}$ and the pointwise-product kernel

$$(\mathbf{b}_1, \dots, \mathbf{b}_k) \mapsto \lambda_{[0;n]}^{\text{block}}(\cdot \mid \mathbf{b}_1) \times \dots \times \lambda_{[0;n]}^{\text{block}}(\cdot \mid \mathbf{b}_k)$$

are $(\kappa kn, r)$ -extremal. Now combine this fact with Proposition 13.3 and Lemma 9.7. \square

13.2. From extremality to finite determination. — We return to the system $(\mathbf{X}, \mu, \mathbf{T})$ and processes ξ and π that appear in Proposition 13.1. Let $\nu := \pi_* \mu$ and $\lambda := (\pi \vee \xi)_* \mu$. Let α and β be the canonical \mathbf{A} - and \mathbf{B} -valued processes on $\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}$, respectively.

Proof of Proposition 13.1. — Choose $n \in \mathbf{N}$ so large that both of the following hold:

- (a) $\alpha_{[0;n]}$ is $(\kappa n, r)$ -extremal over $\beta_{[0;n]}$, and
- (b) $\mathbf{H}_\lambda(\alpha_{[0;n]} \mid \beta_{[0;n]}) < (\mathbf{h}(\alpha, \lambda, \mathbf{T}_{\mathbf{B} \times \mathbf{A}} \mid \beta) + \kappa r) \cdot n = (\mathbf{h}(\lambda \mid \beta) + \kappa r) \cdot n$.

The former holds for all sufficiently large n by assumption. The latter holds for all sufficiently large n by Lemma 10.2.

By Proposition 13.3, conditions (a) and (b) imply that

$$(90) \quad \int \overline{d}_{kn} \left(\lambda_{[0;kn]}^{\text{block}}(\cdot | b_{[0;kn]}), \bigotimes_{j=0}^{k-1} \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}) \right) \nu(db) < 2r$$

for all $k \in \mathbf{N}$.

Now suppose that $\theta \in \text{Ext}(\nu, \mathbf{A})$ satisfies the conditions

- (i) $h(\theta | \beta) > h(\lambda | \beta) - \varepsilon$, and
- (ii) $\|\theta_{[0;n]} - \lambda_{[0;n]}\| < \varepsilon$

for some $\varepsilon > 0$.

Since $\lambda_{[0;n]}$ and $\theta_{[0;n]}$ have the same marginal on \mathbf{B}^n (namely, $\nu_{[0;n]}$), assumption (ii) implies that the following also holds:

- (ii') (conditional distributions close in \overline{d}_n)

$$\int \overline{d}_n(\lambda_{[0;n]}^{\text{block}}(\cdot | b_{[0;n]}), \theta_{[0;n]}^{\text{block}}(\cdot | b_{[0;n]})) \nu(db) < \varepsilon.$$

If ε is sufficiently small in terms of r , then we can apply Lemma 9.7 to conclude that

$$(91) \quad \alpha_{[0;n]} \text{ is } (\kappa n, 2r)\text{-extremal over } \beta_{[0;n]} \text{ according to } \theta,$$

in addition to the extremality according to λ .

On the other hand, the quantity

$$\mathbf{H}_\theta(\alpha_{[0;n]} | \beta_{[0;n]}) = \mathbf{H}_\theta(\alpha_{[0;n]} \vee \beta_{[0;n]}) - \mathbf{H}_\theta(\beta_{[0;n]})$$

is a continuous function of the joint distribution $\theta_{[0;n]}$. Therefore, provided we chose ε sufficiently small, assumption (ii) implies that

$$\mathbf{H}_\theta(\alpha_{[0;n]} | \beta_{[0;n]}) < \mathbf{H}_\lambda(\alpha_{[0;n]} | \beta_{[0;n]}) + \kappa r n / 2.$$

Provided also that we chose $\varepsilon < \kappa r / 2$, condition (b) and assumption (i) give that this right-hand side is less than

$$(h(\lambda | \beta) + \kappa r + \kappa r / 2) \cdot n < (h(\theta | \beta) + 2\kappa r) \cdot n.$$

Combining these inequalities, we obtain

$$(92) \quad \mathbf{H}_\theta(\alpha_{[0;n]} | \beta_{[0;n]}) < (h(\theta | \beta) + 2\kappa r) \cdot n.$$

Conclusions (91) and (92) now provide the hypotheses for another application of Proposition 13.3, this time to the system $(\mathbf{B}^{\mathbf{Z}} \times \mathbf{A}^{\mathbf{Z}}, \theta, \mathbf{T}_{\mathbf{B} \times \mathbf{A}})$ and with r replaced by $2r$. We conclude that

$$(93) \quad \int \overline{d}_{kn} \left(\theta_{[0;kn]}^{\text{block}}(\cdot | b_{[0;kn]}), \bigotimes_{j=0}^{k-1} \theta_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}) \right) \nu(db) < 4r$$

for all $k \in \mathbf{N}$.

Finally, for any $k \in \mathbf{N}$ and $b \in \mathbf{B}^{\mathbf{Z}}$, a $(k-1)$ -fold application of Lemma 4.3 gives

$$\begin{aligned} & \overline{d}_{kn} \left(\bigotimes_{j=0}^{k-1} \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}), \bigotimes_{j=0}^{k-1} \theta_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}) \right) \\ & \leq \frac{1}{k} \sum_{j=0}^{k-1} \overline{d}_n \left(\lambda_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}), \theta_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}) \right). \end{aligned}$$

Here we have used only the special case of Lemma 4.3 that compares two product measures. Integrating this inequality with respect to $\nu(db)$, and recalling condition (ii'), it follows that

$$(94) \quad \int \overline{d}_{kn} \left(\bigotimes_{j=0}^{k-1} \lambda_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}), \bigotimes_{j=0}^{k-1} \theta_{[0;n]}^{\text{block}}(\cdot | b_{[jn;(j+1)n]}) \right) \nu(db) < \varepsilon,$$

which we now assume is less than r . Combined with (90), (93), and the triangle inequality for \overline{d}_{kn} , we obtain that

$$\int \overline{d}_{kn} \left(\lambda_{[0;kn]}^{\text{block}}(\cdot | b_{[0;kn]}), \theta_{[0;kn]}^{\text{block}}(\cdot | b_{[0;kn]}) \right) \nu(db) < 7r.$$

Letting $k \rightarrow \infty$, Lemma 11.2 completes the proof. \square

Part III THE WEAK PINSKER PROPERTY

14. A quantitative step towards Theorem A

In this section, we consider a factor map $\pi : (X, \mu, \mathbf{T}) \rightarrow (Y, \nu, \mathbf{S})$ of ergodic systems such that (Y, ν, \mathbf{S}) satisfies the following special conditions:

The entropy $h(\nu, \mathbf{S})$ is finite, and \mathbf{S} has an N -periodic set modulo ν for every $N \in \mathbf{N}$.

We refer to this as **assumption (A)**. There certainly exist atomless and ergodic automorphisms that satisfy these conditions. For instance, let

$$Y := \prod_{N \geq 1} (\mathbf{Z}/N\mathbf{Z}),$$

let S be the rotation of this compact group by the element $(1, 1, \dots)$, and let ν be any S -invariant ergodic measure.

Since $h(\nu, S)$ is finite, Krieger's generator theorem [47] tells us that (Y, ν, S) is isomorphic to a shift-system with a finite state space. Since the conclusion of Theorem A depends on (Y, ν, S) only up to isomorphism, we are free to replace (Y, ν, S) with that shift-system: that is, we may assume that $Y = \mathbf{B}^{\mathbf{Z}}$ and $S = T_{\mathbf{B}}$ for some finite set \mathbf{B} . Accordingly, π is now a \mathbf{B} -valued process on (X, μ, T) . We retain these assumptions until the proof of Theorem A under assumption (A) is completed in Subsection 15.1.

Now let $\xi : X \rightarrow \mathbf{A}^{\mathbf{Z}}$ be another process in addition to π .

Proposition 14.1. — *Suppose assumption (A) holds, and let $r > 0$ and $\varepsilon > 0$. Then there exist a new process φ on (X, μ, T) and a value $\kappa > 0$ such that*

- (a) $h(\varphi, \mu, T) < \varepsilon$, and
- (b) ξ is relatively (κ, r) -extremal over $\pi \vee \varphi$.

In the light of Proposition 13.1 and Theorem 11.5, this should be seen as a quantitative step towards relative Bernoullicity for ξ over $\pi \vee \varphi$. In Subsection 15.1, we apply Proposition 14.1 repeatedly with smaller and smaller values of r to prove Theorem A under assumption (A). Then it only remains to use some orbit equivalence theory to deduce Theorem A in full.

Throughout this section we abbreviate H_{μ} to H and $h(\xi, \mu, T | \pi)$ to $h(\xi | \pi)$, and similarly. When the measure is omitted from the notation, the correct choice is always μ .

Remark. — It is possible to prove Theorem A in a single step, without ever treating the special case of assumption (A), by replacing periodic sets with sufficiently good Rokhlin sets (see, for instance, [30, p71]). However, the use of Rokhlin's lemma introduces another error tolerance, which then complicates several of the estimates in the proof. Since we use orbit equivalences to extend Theorem A to general amenable groups anyway, there seems to be little value in trying to avoid assumption (A). Of course, Rokhlin's lemma is still implicitly at work within the results that we cite from orbit equivalence theory.

14.1. *The construction of a new observable.* — Before proving Proposition 14.1, we formulate and prove a proposition that gives a new observable φ with some related but more technical properties.

Proposition 14.2. — *Suppose assumption (A) holds, and let $r > 0$ and $\varepsilon > 0$. Then there exist a new process φ on $(\mathbf{X}, \mu, \mathbf{T})$, a positive integer \mathbf{N} , a value $\kappa > 0$, and a φ_0 -measurable and \mathbf{N} -periodic set \mathbf{F} modulo μ such that:*

- (a) $h(\varphi) < \varepsilon$;
- (b) $\xi_{[0;\mathbf{N}]}$ is relatively $(\kappa\mathbf{N}, r)$ -extremal over $(\pi \vee \varphi)_{[0;\mathbf{N}]}$ according to $\mu(\cdot | \mathbf{F})$;
- (c) $\mathbf{H}(\xi_{[0;\mathbf{N}]} | (\pi \vee \varphi)_{[0;\mathbf{N}]}; \mathbf{F}) < (h(\xi | \pi \vee \varphi) + r\kappa) \cdot \mathbf{N}$.

The proof of Proposition 14.2 occupies this subsection. It is the longest and trickiest part of the proof of Theorem A. It includes the key point at which we apply Theorem C.

A warning is in order at this juncture. In the statement of Theorem C we denote the alphabet by \mathbf{A} . However, in this subsection our application of Theorem C does not use the alphabet \mathbf{A} of the process ξ , but rather some large Cartesian power \mathbf{A}^ℓ . The integer ℓ is chosen in Step 1 of the proof below, based on the behaviour of the process ξ (see choice (P2)).

The data \mathbf{N} and \mathbf{F} given by Proposition 14.2 do not appear in Proposition 14.1, but they play an essential auxiliary role in the construction of the new process φ and the proof that it has all the desired properties. This is why we formulate and prove Proposition 14.2 separately.

Conclusions (b) and (c) of Proposition 14.2 match hypotheses (i) and (ii) of Proposition 13.3, but for the conditioned measure $\mu(\cdot | \mathbf{F})$ rather than μ . In the next subsection we show how true extremality of ξ over $\pi \vee \varphi$ can be deduced from Proposition 14.2 using Corollary 13.4.

We break the construction of \mathbf{F} and φ into several steps, and complete the proof of Proposition 14.2 at the end of this subsection.

Step 1: choice of parameters. — By reducing ε if necessary, we may assume it is less than $\log 2$. We now choose several more auxiliary parameters as follows.

- (P1) (Constants provided by Theorem C.) Let c_C and κ_C be the constants provided by Theorem C when both the input error tolerances in that theorem are set equal to r . Now choose some positive $\delta < \min\{\varepsilon/c_C, r\}$.
- (P2) (First time-scale.) Choose $\ell \in \mathbf{N}$ so large that

$$\mathbf{H}(\xi_{[0;\ell]} | \pi_{[0;\ell]}) < (h(\xi | \pi) + \delta^2/2) \cdot \ell.$$

This choice of ℓ is possible by Lemma 10.2. Set $\kappa := \kappa_C/\ell$.

- (P3) (Second time-scale.) Lastly, choose $n \in \mathbf{N}$ so large that all of the following hold:
 - (P3.1) (entropy estimate)

$$\mathbf{H}(\xi_{[0;n\ell]} | \pi_{[0;n\ell]}) < (h(\xi | \pi) + \kappa r/2) \cdot n \cdot \ell;$$

(P3.2) (requirement for Theorem C) Theorem C applies to the n -fold power of the alphabet A^ℓ when both of the input tolerances equal r ;

(P3.3) (another auxiliary estimate needed below)

$$c_C \exp(c_C \delta n \ell) < \exp(\varepsilon n \ell).$$

The fact that (P3.1) holds for all sufficiently large n is a consequence of Lemma 10.2. The fact that (P3.2) can be satisfied is precisely Theorem C.

Let $N := n\ell$.

Each of these choices is needed for a particular step in the argument below, and is explained when we reach that step. The relation between ℓ and n is very important: it allows us to regard certain conditional distributions on $A^{n\ell}$ as measures on the n -fold product of the alphabet A^ℓ for which we have control over the total correlation. This control is established in Lemma 14.4 below.

Step 2: introducing a periodic set. — Having chosen N , it is time to invoke assumption (A). It provides a π -measurable set $F_0 \subseteq X$ which is N -periodic under T modulo μ . Modulo negligible sets, we can now define a π -measurable factor map $\tau : X \rightarrow \mathbf{Z}/N\mathbf{Z}$ to a finite rotation by requiring that

$$\tau(x) = j \pmod{N} \quad \text{when} \quad T^{-j}x \in F_0.$$

Let $F_j := \{\tau = j \pmod{N}\}$ for each $j \in \mathbf{Z}$, so this equals $T^j F_0$ modulo μ , and let $\bar{F}_j := \{\tau = j \pmod{\ell}\}$, which agrees with

$$F_j \cup T^{-\ell}F_j \cup \dots \cup T^{-(n-1)\ell}F_j$$

modulo μ . Each F_j is N -periodic under T modulo μ , each \bar{F}_j is ℓ -periodic under T modulo μ , and $\bar{F}_{j+\ell} = \bar{F}_j$ modulo μ for every $j \in \mathbf{Z}$.

Lemma 14.3. — *There exists $j \in \{0, 1, \dots, N-1\}$ satisfying both*

$$(a) \quad H(\xi_{[0;N]} | \pi_{[0;N]}; F_j) < (h(\xi | \pi) + \kappa r) \cdot N$$

and

$$(b) \quad \frac{1}{n} \sum_{i=0}^{n-1} H(\xi_{[0;\ell]} | \pi_{[0;\ell]}; T^{-i\ell}F_j) < (h(\xi | \pi) + \delta^2) \cdot \ell.$$

Proof. — Up to a negligible set, the sets F_j for $0 \leq j \leq N-1$ agree with the partition of X generated by τ . Therefore we can compute $H(\xi_{[0;N]} | \tau \vee \pi_{[0;N]})$ by first conditioning onto each of the sets F_j , and then refining that partition by $\pi_{[0;N]}$. This gives

$$\frac{1}{N} \sum_{j=0}^{N-1} H(\xi_{[0;N]} | \pi_{[0;N]}; F_j) = H(\xi_{[0;N]} | \tau \vee \pi_{[0;N]})$$

$$\leq \mathbf{H}(\xi_{[0;N]} | \pi_{[0;N]}) < (\mathbf{h}(\xi | \pi) + \kappa r/2) \cdot N,$$

by assumption (P3.1). On the other hand, each of the sets F_j is π -measurable and N -periodic under T , so the terms in the left-hand average here all satisfy

$$\mathbf{H}(\xi_{[0;N]} | \pi_{[0;N]}; F_j) \geq \mathbf{H}(\xi_{[0;N]} | \pi; F_j) \geq \mathbf{h}(\xi | \pi) \cdot N,$$

using Lemma 10.6 for the second inequality. Therefore, by Markov's inequality, inequality (a) is satisfied for more than half of the values $j \in \{0, 1, \dots, N-1\}$.

Turning to inequality (b), we start by observing the following. Since $\bar{F}_{j+\ell} = \bar{F}_j$ modulo μ and N is a multiple of ℓ , we have

$$\frac{1}{N} \sum_{j=0}^{N-1} \mathbf{H}(\xi_{[0;\ell]} | \pi_{[0;\ell]}; \bar{F}_j) = \frac{1}{\ell} \sum_{j=0}^{\ell-1} \mathbf{H}(\xi_{[0;\ell]} | \pi_{[0;\ell]}; \bar{F}_j).$$

This lets us re-apply the argument for inequality (a). This time we use the ℓ -periodic sets \bar{F}_j in place of F_j , the maps $\xi_{[0;\ell]}$ and $\pi_{[0;\ell]}$, and assumption (P2) in place of (P3.1). We deduce that more than half of the values $j \in \{0, 1, \dots, N-1\}$ satisfy

$$(95) \quad \mathbf{H}(\xi_{[0;\ell]} | \pi_{[0;\ell]}; \bar{F}_j) < (\mathbf{h}(\xi | \pi) + \delta^2) \cdot \ell.$$

On the other hand, for each j , the sets $T^{-i\ell}F_j$ for $0 \leq i \leq n-1$ are a partition of \bar{F}_j into n further subsets of equal measure (again up to a negligible set), and this partition of \bar{F}_j is generated by the restriction $\tau|_{\bar{F}_j}$. From this, the chain rule, and monotonicity under conditioning, we deduce that

$$\begin{aligned} \frac{1}{n} \sum_{i=0}^{n-1} \mathbf{H}(\xi_{[0;\ell]} | \pi_{[0;\ell]}; T^{-i\ell}F_j) &= \mathbf{H}(\xi_{[0;\ell]} | \pi_{[0;\ell]} \vee \tau; \bar{F}_j) \\ &\leq \mathbf{H}(\xi_{[0;\ell]} | \pi_{[0;\ell]}; \bar{F}_j). \end{aligned}$$

Therefore any value of j which satisfies (95) also satisfies inequality (b).

Combining the conclusions above, there is at least one j for which both (a) and (b) hold. \square

With j as given by Lemma 14.3, fix $F := F_j$.

Step 3: some properties of the conditional marginals. — Set $\tilde{A} := A^\ell$, so that we may identify $A^{n\ell} = A^N$ with \tilde{A}^n . Now define a new kernel by

$$\tilde{\lambda}(\mathbf{a} | \mathbf{b}) := \mu(\xi_{[0;N]} = \mathbf{a} | \{\pi_{[0;N]} = \mathbf{b}\} \cap F) \quad (\mathbf{a} \in A^N = \tilde{A}^n, \mathbf{b} \in B^N).$$

This is an instance of an N -block kernel, as in Definition 10.1, except for the extra conditioning on F . At this point it is important that we regard this as a kernel from B^N to \tilde{A}^n , rather than to A^N . In addition, define the probability measure $\tilde{\nu}$ on B^N by $\tilde{\nu}(\mathbf{b}) := \mu(\pi_{[0;N]} = \mathbf{b} | F)$.

Lemma 14.4. — *Let W be the set of all $\mathbf{b} \in \mathbb{B}^N$ which satisfy*

$$\tilde{\nu}(\mathbf{b}) > 0 \quad \text{and} \quad \text{TC}(\tilde{\lambda}(\cdot | \mathbf{b})) \leq \delta N.$$

Then $\tilde{\nu}(W) = \mu(\pi_{[0;N]} \in W | \mathbb{F}) > 1 - \delta$.

Proof. — Consider a string $\mathbf{b} \in \mathbb{B}^N$ for which $\tilde{\nu}(\mathbf{b}) > 0$. Then $\tilde{\lambda}(\cdot | \mathbf{b})$ is the conditional distribution of $\xi_{[0;N]}$ given the event $\{\pi_{[0;N]} \stackrel{\sim}{=} \mathbf{b}\} \cap \mathbb{F}$, where we regard $\xi_{[0;N]}$ as taking values in $\tilde{\mathbb{A}}^n$. Accordingly, the n marginals of $\tilde{\lambda}(\cdot | \mathbf{b})$ on the space $\tilde{\mathbb{A}}$ are the conditional distributions of the $\tilde{\mathbb{A}}$ -valued observables

$$\xi_{[0;\ell]}, \xi_{[\ell;2\ell]}, \dots, \xi_{[(n-1)\ell;n\ell]}$$

given the event $\{\pi_{[0;N]} = \mathbf{b}\} \cap \mathbb{F}$. Therefore, by the definition of total correlation,

$$\begin{aligned} \text{TC}(\tilde{\lambda}(\cdot | \mathbf{b})) &= \sum_{i=1}^n \mathbb{H}(\xi_{[(i-1)\ell;i\ell]} | \{\pi_{[0;N]} = \mathbf{b}\} \cap \mathbb{F}) \\ &\quad - \mathbb{H}(\xi_{[0;N]} | \{\pi_{[0;N]} = \mathbf{b}\} \cap \mathbb{F}) \\ &= \sum_{i=1}^n \mathbb{H}_{\mu(\cdot | \mathbb{F})}(\xi_{[(i-1)\ell;i\ell]} | \{\pi_{[0;N]} = \mathbf{b}\}) \\ &\quad - \mathbb{H}_{\mu(\cdot | \mathbb{F})}(\xi_{[0;N]} | \{\pi_{[0;N]} = \mathbf{b}\}). \end{aligned}$$

Integrating this equality with respect to $\tilde{\nu}(d\mathbf{b})$, we obtain:

$$(96) \quad \int \text{TC}(\tilde{\lambda}(\cdot | \mathbf{b})) \tilde{\nu}(d\mathbf{b}) = \sum_{i=1}^n \mathbb{H}(\xi_{[(i-1)\ell;i\ell]} | \pi_{[0;N]}; \mathbb{F}) - \mathbb{H}(\xi_{[0;N]} | \pi_{[0;N]}; \mathbb{F}).$$

By the monotonicity of Shannon entropy under conditioning, and then by the T-invariance of μ , this is at most

$$(97) \quad \begin{aligned} &\sum_{i=1}^n \mathbb{H}(\xi_{[(i-1)\ell;i\ell]} | \pi_{[(i-1)\ell;i\ell]}; \mathbb{F}) - \mathbb{H}(\xi_{[0;N]} | \pi_{[0;N]}; \mathbb{F}) \\ &= \sum_{i=0}^{n-1} \mathbb{H}(\xi_{[0;\ell]} | \pi_{[0;\ell]}; T^{-i\ell}\mathbb{F}) - \mathbb{H}(\xi_{[0;N]} | \pi_{[0;N]}; \mathbb{F}). \end{aligned}$$

By inequality (b) from Lemma 14.3, the sum of positive terms here is less than

$$(\mathfrak{h}(\xi | \pi) + \delta^2) \cdot \ell \cdot n = (\mathfrak{h}(\xi | \pi) + \delta^2) \cdot N.$$

On the other hand, $H(\xi_{[0;N]} | \pi_{[0;N]}; F)$ is at least $H(\xi_{[0;N]} | \pi; F)$ by the monotonicity of Shannon entropy under conditioning, and the latter is at least $h(\xi | \pi) \cdot N$ by Lemma 10.6. Therefore the whole of (97) is strictly less than

$$N \cdot (h(\xi | \pi) + \delta^2) - N \cdot h(\xi | \pi) = \delta^2 N,$$

and hence so is the integral on the left-hand side of (96). So $\tilde{\nu}(W) > 1 - \delta$ by Markov's inequality. \square

Step 4: application of Theorem C. — For each $\mathbf{b} \in W$, we now apply Theorem C to the measure $\tilde{\lambda}(\cdot | \mathbf{b})$ with the choice of parameters in (P1) above. For each $\mathbf{b} \in W$, that theorem provides a partition

$$(98) \quad \tilde{A}^n = U_{\mathbf{b},1} \cup \dots \cup U_{\mathbf{b},m_{\mathbf{b}}}$$

with the following properties:

- (a) The number of cells satisfies $m_{\mathbf{b}} \leq c_C \exp(c_C \delta N)$, which is less than $e^{\varepsilon N}$ by our choice in (P3.3).
- (b) The first cell satisfies $\tilde{\lambda}(U_{\mathbf{b},1} | \mathbf{b}) < r$.
- (c) The twice-conditioned measure $(\tilde{\lambda}(\cdot | \mathbf{b}))_{|U_{\mathbf{b},j}}$ is well-defined and satisfies $T(\kappa_C n, r)$ for every $j = 2, 3, \dots, m_{\mathbf{b}}$, where the underlying metric is the normalized Hamming metric d_n on the n -fold product \tilde{A}^n . This is greater than or equal to the normalized Hamming metric d_N when we identify \tilde{A}^n with the N -fold product A^N , so $(\tilde{\lambda}(\cdot | \mathbf{b}))_{|U_{\mathbf{b},j}}$ satisfies $T(\kappa_C n, r)$ for that choice of metric also. Recalling that $\kappa = \kappa_C / \ell$, let us write this last inequality as $T(\kappa N, r)$.

By including a few copies of the empty set, let us suppose that $m_{\mathbf{b}}$ equals a fixed value m for each $\mathbf{b} \in W$, with the modification that conclusion (c) above holds only when $U_{\mathbf{b},j} \neq \emptyset$. Also, for $\mathbf{b} \in B^N \setminus W$, choose an arbitrary partition of \tilde{A}^n into m subsets, so that the notation in (98) still makes sense for those \mathbf{b} .

We now switch back from \tilde{A}^n to A^N in our discussion.

In the rest of the construction, we need the right-hand side of (98) to be indexed by strings over some alphabet, rather than by the integers $1, \dots, m$. Since we have chosen $\varepsilon < \log 2$, we have $m \leq 2^N$. So let $S \subseteq \{0, 1\}^N$ be any subset of cardinality m , and re-write the partitions on the right-hand side of (98) so that the parts are indexed by the strings in S :

$$(99) \quad A^N = \bigcup_{\mathbf{c} \in S} U_{\mathbf{b},\mathbf{c}}.$$

Do this in such a way that some fixed element $\mathbf{c}_1 \in S$ has $U_{\mathbf{b},\mathbf{c}_1}$ equal to the ‘bad’ cell $U_{\mathbf{b},1}$ from property (b) for each \mathbf{b} .

Define $\Psi : B^N \times A^N \longrightarrow S$ by

$$\Psi(\mathbf{b}, \mathbf{a}) = \mathbf{c} \quad \text{whenever } \mathbf{a} \in U_{\mathbf{b}, \mathbf{c}},$$

and define $\Phi : F \longrightarrow S$ by

$$\Phi(x) := \Psi(\pi_{[0;N]}(x), \xi_{[0;N]}(x)).$$

Lemma 14.5. — *The map $\xi_{[0;N]}$ is $(\kappa N, 3r)$ -extremal over $\pi_{[0;N]} \vee \Phi$ according to $\mu(\cdot | F)$.*

Proof. — The point is that the measures

$$\lambda_{\mathbf{b}, \mathbf{c}} := (\tilde{\lambda}(\cdot | \mathbf{b}))|_{U_{\mathbf{b}, \mathbf{c}}}$$

constitute a conditional distribution of $\xi_{[0;N]}$ given $\pi_{[0;N]} \vee \Phi$ according to $\mu(\cdot | F)$. To see this, let $\mathbf{a} \in A^N$, $\mathbf{b} \in B^N$ and $\mathbf{c} \in S$, and suppose that the conditional probability

$$\begin{aligned} & \frac{\mu(\xi_{[0;N]} \vee \pi_{[0;N]} \vee \Phi = (\mathbf{a}, \mathbf{b}, \mathbf{c}) | F)}{\mu(\pi_{[0;N]} \vee \Phi = (\mathbf{b}, \mathbf{c}) | F)} \\ &= \frac{\mu(\{\xi_{[0;N]} \vee \pi_{[0;N]} \vee \Phi = (\mathbf{a}, \mathbf{b}, \mathbf{c})\} \cap F)}{\mu(\{\pi_{[0;N]} \vee \Phi = (\mathbf{b}, \mathbf{c})\} \cap F)} \end{aligned}$$

has nonzero denominator. Then this fraction satisfies the following:

- It is zero if $\mathbf{c} \neq \Psi(\mathbf{b}, \mathbf{a})$, or equivalently if $\mathbf{a} \notin U_{\mathbf{b}, \mathbf{c}}$. In this case $\lambda_{\mathbf{b}, \mathbf{c}}(\mathbf{a})$ is also zero.
- If $\mathbf{c} = \Psi(\mathbf{b}, \mathbf{a})$, then the fraction equals

$$\begin{aligned} & \frac{\mu(\{\pi_{[0;N]} = \mathbf{b}\} \cap \{\xi_{[0;N]} = \mathbf{a}\} \cap F)}{\mu(\{\pi_{[0;N]} = \mathbf{b}\} \cap \{\Psi(\mathbf{b}, \xi_{[0;N]}) = \mathbf{c}\} \cap F)} \\ &= \frac{\mu(\xi_{[0;N]} = \mathbf{a} | \{\pi_{[0;N]} = \mathbf{b}\} \cap F)}{\mu(\xi_{[0;N]} \in U_{\mathbf{b}, \mathbf{c}} | \{\pi_{[0;N]} = \mathbf{b}\} \cap F)} = \frac{\tilde{\lambda}(\mathbf{a} | \mathbf{b})}{\tilde{\lambda}(U_{\mathbf{b}, \mathbf{c}} | \mathbf{b})} = \lambda_{\mathbf{b}, \mathbf{c}}(\mathbf{a}). \end{aligned}$$

Having identified this conditional distribution, we may use the following consequence of our appeal to Theorem C:

$$\begin{aligned} & \mu(\pi_{[0;N]} \vee \Phi \in \{(\mathbf{b}, \mathbf{c}) : \lambda_{\mathbf{b}, \mathbf{c}} \text{ satisfies } T(\kappa N, r)\} | F) \\ & \geq 1 - \mu(\pi_{[0;N]} \notin W \text{ or } \xi_{[0;N]} \in U_{\pi_{[0;N]}, \mathbf{c}_1} | F) \\ & \geq 1 - \mu(\pi_{[0;N]} \notin W | F) - \int_W \tilde{\lambda}(U_{\mathbf{b}, \mathbf{c}_1} | \mathbf{b}) \tilde{\nu}(d\mathbf{b}). \end{aligned}$$

By Lemma 14.4 and conclusion (b) obtained from Theorem C, this is at least $1 - \delta - r$, which is at least $1 - 2r$. Now Lemma 9.5 completes the proof. \square

Step 5: finishing the construction. — It remains to construct the new process φ . It is obtained from τ together with another new observable $\psi_0 : \mathbf{X} \rightarrow \{0, 1\}$. The idea behind ψ_0 is that names should be written as length- N blocks that are given by Φ , where the blocks start at the return times to F . In notation, we wish to define ψ_0 so that

$$(100) \quad \psi_{[0;N]}|F = \Phi.$$

This is possible because F is N -periodic — indeed, this is precisely why we introduce an N -periodic set into the construction. The formal definition that achieves (100) is as follows. Given $x \in \mathbf{X}$, choose $t \in [0; N)$ so that $T^{-t}x \in F$, and now let $\psi_0(x)$ be the $(t+1)$ th letter of the string

$$\Phi(T^{-t}x).$$

(We need the $(t+1)$ th letter because we index $\{0, 1\}^N$ using $\{1, \dots, N\}$, not $[0; N)$.)

Let $\psi : \mathbf{X} \rightarrow \{0, 1\}^{\mathbf{Z}}$ be the process generated by ψ_0 . Now let $\varphi := \tau \vee \psi$, and observe that $\pi \vee \varphi$ and $\pi \vee \psi$ generate the same factor of (\mathbf{X}, μ, T) .

Proof of Proposition 14.2. — First, (100) and Lemma 10.6 give that

$$h(\varphi) = h(\psi) \leq \frac{1}{N}H(\psi_{[0;N]} | F) = \frac{1}{N}H(\Phi | F) \leq \frac{\log |S|}{N} < \varepsilon.$$

This proves (a).

Next, the definition (100) has the consequence that the conditional distribution of $\xi_{[0;N]}$ given $\pi_{[0;N]} \vee \varphi_{[0;N]}$ according to $\mu(\cdot | F)$ is the same as its conditional distribution given $\pi_{[0;N]} \vee \Phi$ according to $\mu(\cdot | F)$. In particular, τ has no effect on this conditional distribution, because $\tau|F$ is constant. Therefore $\xi_{[0;N]}$ is relatively $(\kappa N, 3r)$ -extremal over $(\pi \vee \varphi)_{[0;N]}$ according to $\mu(\cdot | F)$, by Lemma 14.5.

Now we turn to (c). On the event F , we have from (100) that the map $\psi_{[0;N]}$ agrees with Φ , which is a function of $\pi_{[0;N]} \vee \xi_{[0;N]}$. Therefore the chain rule for Shannon entropy gives

$$\begin{aligned} H(\xi_{[0;N]} | \varphi_{[0;N]} \vee \pi_{[0;N]}; F) &= H(\xi_{[0;N]} | \psi_{[0;N]} \vee \pi_{[0;N]}; F) \\ &= H(\xi_{[0;N]} | \Phi \vee \pi_{[0;N]}; F) \\ &= H(\xi_{[0;N]} | \pi_{[0;N]}; F) - H(\Phi | \pi_{[0;N]}; F) \\ &= H(\xi_{[0;N]} | \pi_{[0;N]}; F) - H(\psi_{[0;N]} | \pi_{[0;N]}; F). \end{aligned}$$

By inequality (a) from Lemma 14.3, this is less than

$$(h(\xi | \pi) + \kappa r) \cdot N - H(\psi_{[0;N]} | \pi_{[0;N]}; F).$$

On the other hand, Lemma 10.6 gives

$$H(\psi_{[0;N]} | \pi_{[0;N]}; F) \geq h(\psi | \pi) \cdot N,$$

so it follows that

$$H(\xi_{[0;N]} | \varphi_{[0;N]} \vee \pi_{[0;N]}; F) < (h(\xi | \pi) - h(\psi | \pi) + \kappa r) \cdot N.$$

Finally, $\pi \vee \psi$ and $\pi \vee \varphi$ generate the same σ -algebra modulo μ and are both measurable with respect to $\pi \vee \xi$, so Lemma 10.3 can be applied and re-arranged to give

$$h(\xi | \pi) - h(\psi | \pi) = h(\xi \vee \varphi | \pi) - h(\varphi | \pi) = h(\xi | \pi \vee \varphi).$$

Substituting this into the previous inequality, we have shown that

$$H(\xi_{[0;N]} | (\pi \vee \varphi)_{[0;N]}; F) < (h(\xi | \pi \vee \varphi) + \kappa r) \cdot N.$$

Since $r < 3r$, this completes the proof of all three desired properties with $3r$ in place of r . This suffices because $r > 0$ was arbitrary. \square

14.2. Completion of the quantitative step. — We now complete the proof of Proposition 14.1. Let $r' := r/17$, and apply Proposition 14.2 with r' in place of r . Let φ , $N \in \mathbf{N}$, $\kappa > 0$ and $F \subseteq X$ be the data given by that proposition.

We show that this φ satisfies the conclusions of Proposition 14.1. Essentially, this requires that we replace conclusion (b) of Proposition 14.2 with extremality that holds without reference to the periodic set F and over arbitrarily long time-scales.

This replacement is enabled by the next two lemmas. Let $\tilde{\pi} := \pi \vee \varphi$, and suppose it takes values in C^Z . The first lemma gives extremality over longer time-scales, but still according to $\mu(\cdot | F)$.

Lemma 14.6. — *For any $k \in \mathbf{N}$, the observable $\xi_{[0;kN]}$ is relatively $(\kappa kN, 5r')$ -extremal over $\tilde{\pi}_{[0;kN]}$ according to $\mu(\cdot | F)$.*

Proof. — This is our application of the ‘stretching’ result for extremality given by Corollary 13.4. Let $\mu' := \mu(\cdot | F)$. Since F is N -periodic, μ' is invariant under T^N , and so we may apply Corollary 13.4 to the joint distribution of ξ and $\tilde{\pi}$ under μ' . Conclusions (b) and (c) of Proposition 14.2 give precisely the hypotheses (i) and (ii) needed by that corollary. The result follows. \square

Using Lemma 14.6, we now prove a similar result for each of the shifted sets $T^{-i}F$, and for all sufficiently large time-scales rather than just multiples of N .

Lemma 14.7. — *Let $i \in \{0, 1, \dots, N-1\}$. For all sufficiently large t , the observable $\xi_{[0;t]}$ is $(\kappa t, r)$ -extremal over $\tilde{\pi}_{[0;t]}$ according to $\mu(\cdot | T^{-i}F)$.*

Proof. — For this we must combine Lemma 14.6, the invariance of μ , and some of the stability properties of extremality from Section 9.

Step 1. Assume $t \geq N$, and let $k \geq 0$ be the largest integer satisfying $i + kN \leq t$. Let $j := i + kN$. Then $[i; j)$ is the largest subinterval of $[0; t)$ that starts at i and has length a multiple of N .

For any $(\mathbf{c}, \mathbf{a}) \in (\mathbf{C} \times \mathbf{A})^{kN}$, the T -invariance of μ gives

$$\mu((\tilde{\pi} \vee \xi)_{[i;j)} = (\mathbf{c}, \mathbf{a}) \mid T^{-i}\mathbf{F}) = \mu((\tilde{\pi} \vee \xi)_{[0;kN)} = (\mathbf{c}, \mathbf{a}) \mid \mathbf{F}).$$

Therefore, by Lemma 14.6, $\xi_{[i;j)}$ is $(\kappa kN, 5r')$ -extremal over $\tilde{\pi}_{[i;j)}$ according to $\mu(\cdot \mid T^{-i}\mathbf{F})$.

Step 2. The observable $\tilde{\pi}_{[0;t)}$ generates a finer partition of $\mathbf{C}^{\mathbf{Z}}$ than the observable $\tilde{\pi}_{[i;j)}$. On the other hand,

$$\begin{aligned} \mathbf{H}(\xi_{[i;j)} \mid \tilde{\pi}_{[i;j)}) - \mathbf{H}(\xi_{[i;j)} \mid \tilde{\pi}_{[0;t)}) &\leq \mathbf{H}(\tilde{\pi}_{[0;t)} \mid \tilde{\pi}_{[i;j)}) \leq \mathbf{H}(\tilde{\pi}_{[0;t)} \vee \tilde{\pi}_{[i;t)}) \\ &\leq 2\mathbf{NH}(\tilde{\pi}_0), \end{aligned}$$

where the last estimate holds by stationarity and because $i \leq N$ and $t - j \leq N$. Therefore we may combine Step 1 with Corollary 9.11 to conclude that $\xi_{[i;j)}$ is extremal over $\tilde{\pi}_{[0;t)}$ according to $\mu(\cdot \mid T^{-i}\mathbf{F})$ with parameters

$$\left(\kappa kN, \frac{4\mathbf{NH}(\tilde{\pi}_0)}{\kappa kN} + 15r' \right).$$

If t is sufficiently large, and hence k is sufficiently large, then these parameters are at least as strong as $(\kappa kN, 16r')$.

Step 3. Finally, assume t is so large that $kN \geq (1 - r')t$. Since $\xi_{[i;j)}$ is an image of $\xi_{[0;t)}$ under coordinate projection, we may apply Lemma 9.12. Starting from the conclusion of Step 2, that lemma gives that $\xi_{[0;t)}$ is $(\kappa t, 17r')$ -extremal over $\tilde{\pi}_{[0;t)}$ according to $\mu(\cdot \mid T^{-i}\mathbf{F})$. Since $r = 17r'$, this completes the proof. \square

Proof of Proposition 14.1. — Let $\theta := (\tilde{\pi} \vee \xi)_* \mu$ and $\gamma := \tilde{\pi}_* \mu$. For $t \in \mathbf{N}$, consider the t -block kernel of θ :

$$\theta_{[0;t)}^{\text{block}}(\mathbf{a} \mid \mathbf{c}) := \mu(\xi_{[0;t)} = \mathbf{a} \mid \tilde{\pi}_{[0;t)} = \mathbf{c}) \quad (\mathbf{a} \in \mathbf{A}^t, \mathbf{c} \in \mathbf{C}^t)$$

(see Definition 10.1). (It does not matter how we interpret this in case $\mu(\tilde{\pi}_{[0;t)} = \mathbf{c})$ is zero.)

Each of the sets $T^{-i}\mathbf{F}$ is measurable with respect to $\tilde{\pi}_0$, hence certainly with respect to $\tilde{\pi}_{[0;t)}$. So there is a partition

$$\mathbf{C}^t = \mathbf{G}_1 \cup \dots \cup \mathbf{G}_N$$

such that $T^{-i}\mathbf{F} = \{\tilde{\pi}_{[0;t)} \in \mathbf{G}_i\}$ modulo μ for each $i = 1, 2, \dots, N$. By Lemma 14.7 the measure $\gamma_{[0;t)}(\cdot \mid \mathbf{G}_i)$ and kernel $\theta_{[0;t)}^{\text{block}}(\cdot \mid \cdot)$ are $(\kappa t, r)$ -extremal for each i . Since the \mathbf{G}_i s are a partition of \mathbf{C}^t , it follows directly from Definition 9.4 that $\gamma_{[0;t)}$ itself and $\theta_{[0;t)}^{\text{block}}(\cdot \mid \cdot)$ are $(\kappa t, r)$ -extremal. \square

15. Completed proof of Theorem A

15.1. *Proof under assumption (A).* — Let $\pi : (\mathbf{X}, \mu, \mathbf{T}) \rightarrow (\mathbf{Y}, \nu, \mathbf{S})$ be the factor map in the statement of Theorem A. We must prove the relative weak Pinsker property over this factor. We first do this under the additional assumption (A) from the previous section. Note that we do *not* assume that $(\mathbf{X}, \mu, \mathbf{T})$ itself has finite entropy.

Proof of Theorem A under assumption (A). — As in the previous section, under assumption (A) we may assume that π is a finite-valued process, not just an arbitrary factor map.

Let $\xi^{(k)} : \mathbf{X} \rightarrow \mathbf{A}_k^{\mathbf{Z}}$, $k \geq 1$, be a sequence of finite-valued processes which together generate the whole of $\mathcal{B}_{\mathbf{X}}$. By replacing each $\xi^{(k)}$ with the common refinement $\xi^{(1)} \vee \dots \vee \xi^{(k)}$ if necessary, we may assume that $\xi^{(k)}$ is $\xi^{(k+1)}$ -measurable for every k . For the sake of notation, let $\xi^{(0)}$ be a trivial process.

(In case $h(\mu, \mathbf{T}) < \infty$, Krieger's generator theorem provides a single generating process, but this makes the rest of the proof only slightly simpler.)

Construction. Fix a sequence $(k_i)_{i=1}^{\infty} \in \mathbf{N}^{\mathbf{N}}$ in which every natural number appears infinitely often. We use a recursion to construct sequences of finite-valued processes $\varphi^{(i)}$ and real values $\kappa_i > 0$ such that the following hold for every i :

(a) we have

$$h(\varphi^{(i)}, \mu, \mathbf{T}) < 2^{-i} \min \{ \varepsilon, \kappa_1, \dots, \kappa_{i-1} \}.$$

(b) $\xi^{(k_i)}$ is conditionally $(\kappa_i, 2^{-i})$ -extremal over the process

$$\psi^{(i)} := \pi \vee \xi^{(k_{i-1})} \vee \varphi^{(1)} \vee \dots \vee \varphi^{(i)}.$$

To start the recursion, apply Proposition 14.1 to the processes $\pi \vee \xi^{(k_1-1)}$ and $\xi^{(k_1)}$. If $i \geq 2$ and we have already found $\varphi^{(j)}$ for all $j < i$, then apply Proposition 14.1 to the processes

$$\pi \vee \xi^{(k_{i-1})} \vee \varphi^{(1)} \vee \dots \vee \varphi^{(i-1)} \quad \text{and} \quad \xi^{(k_i)}.$$

Completion of the proof. Consider the factor of $(\mathbf{X}, \mu, \mathbf{T})$ generated by π and all the new processes $\varphi^{(i)}$. By property (a), it satisfies

$$\begin{aligned} h\left(\pi \vee \bigvee_{i \geq 1} \varphi^{(i)}, \mu, \mathbf{T} \mid \pi\right) &\leq \lim_{\ell \rightarrow \infty} h(\varphi^{(1)} \vee \dots \vee \varphi^{(\ell)}, \mu, \mathbf{T}) \\ &\leq \lim_{\ell \rightarrow \infty} \sum_{j=1}^{\ell} h(\varphi^{(j)}, \mu, \mathbf{T}) < \varepsilon. \end{aligned}$$

Therefore, by Krieger's generator theorem, there is a finite-valued process $\tilde{\pi}$ that generates the same factor of $(\mathbf{X}, \mu, \mathbf{T})$ as $\pi \vee \bigvee_{i \geq 1} \varphi^{(i)}$, modulo μ -negligible sets.

More generally, for any $i \in \mathbf{N}$, a similar estimate using property (a) gives

$$(101) \quad h(\tilde{\pi}, \mu, T | \pi \vee \varphi^{(1)} \vee \dots \vee \varphi^{(i)}) \leq \sum_{j=i+1}^{\infty} h(\varphi^{(j)}, \mu, T) < \sum_{j=i+1}^{\infty} 2^{-j} \kappa_i = 2^{-i} \kappa_i.$$

Now consider one of the processes $\xi^{(k)}$. Since k appears in the sequence $(k_i)_i$ infinitely often, there are arbitrarily large $i \in \mathbf{N}$ for which $k_i = k$. We consider such a choice of i , and argue as follows:

- According to property (b) above, $\xi^{(k)} = \xi^{(k_i)}$ is conditionally $(\kappa_i, 2^{-i})$ -extremal over $\psi^{(i)}$.
- Using estimate (101) and Lemma 12.2, it follows that $\xi^{(k)}$ is also $(\kappa_i, 6 \cdot 2^{-i})$ -extremal over $\tilde{\pi} \vee \psi^{(i)}$.
- Next, Proposition 13.1 gives that $\xi^{(k)}$ is relatively $(42 \cdot 2^{-i})$ -FD over $\tilde{\pi} \vee \psi^{(i)}$.
- Finally, Lemma 11.6 gives the same conclusion for $\xi^{(k)}$ over $\tilde{\pi} \vee \xi^{(k-1)}$, since the two processes $\tilde{\pi} \vee \psi^{(i)}$ and $\tilde{\pi} \vee \xi^{(k-1)}$ generate the same factor of (\mathbf{X}, μ, T) modulo μ .

Thus, we have shown that $\xi^{(k)}$ is relatively $(42 \cdot 2^{-i})$ -FD over $\tilde{\pi} \vee \xi^{(k-1)}$ for arbitrarily large values of i . Now Thouvenot's theorem (Theorem 11.5) gives that it is also relatively Bernoulli.

For each k , this fact promises an i.i.d. process $\zeta^{(k)} : \mathbf{X} \rightarrow \mathbf{C}_k^{\mathbf{Z}}$ such that (i) the map $\zeta^{(k)}$ is independent from the map

$$\tilde{\pi} \vee \xi^{(1)} \vee \dots \vee \xi^{(k-1)},$$

and (ii) the maps

$$\tilde{\pi} \vee \xi^{(1)} \vee \dots \vee \xi^{(k-1)} \vee \zeta^{(k)} \quad \text{and} \quad \tilde{\pi} \vee \xi^{(1)} \vee \dots \vee \xi^{(k-1)} \vee \xi^{(k)}$$

generate the same factor of (\mathbf{X}, μ, T) modulo μ . It follows that the factor map

$$\bigvee_{k \geq 1} \zeta^{(k)} : \mathbf{X} \rightarrow \prod_{k \geq 1} \mathbf{C}_k^{\mathbf{Z}}$$

has image a Bernoulli shift (possibly of infinite entropy), is independent from the factor map $\tilde{\pi}$, and is such that

$$\tilde{\pi} \vee \bigvee_{k \geq 1} \zeta^{(k)} \quad \text{and} \quad \tilde{\pi} \vee \bigvee_{k \geq 1} \xi^{(k)}$$

generate the same factor of (\mathbf{X}, μ, T) modulo μ . This factor is the whole of $\mathcal{B}_{\mathbf{X}}$ by our choice of the $\xi^{(k)}$ s, so this completes the proof. \square

Remark. — An important part of the above proof is our appeal to Lemma 12.2. Given the extremality of $\xi^{(k_i)}$ over $\psi^{(i)}$ and the entropy bound (101), that lemma promises that $\xi^{(k_i)}$ is still extremal, with slightly worse parameters, over the limiting factor map $\tilde{\pi} \vee \psi^{(i)}$.

It is natural to ask whether a more qualitative fact could be used in place of this argument. Specifically, suppose that ξ is a process and $\pi^{(i)}$ is a sequence of factor maps such that $\pi^{(i)}$ is $\pi^{(i+1)}$ -measurable for every i . If ξ is relatively Bernoulli over each $\pi^{(i)}$, is it also relatively Bernoulli over $\bigvee_{i \geq 1} \pi^{(i)}$? Unfortunately, Jean-Paul Thouvenot has indicated to me an example in which the answer is No. It seems that some more quantitative argument, like our appeal to Lemma 12.2, really is needed.

15.2. Removing assumption (A), and general amenable groups. — We now complete the proof of Theorem A in full.

In fact, for the work that is left to do, it makes no difference if we prove a result for actions of a general countable amenable group G . So suppose now that $T = (T^g)_{g \in G}$ is a free and ergodic measure-preserving action of G on a standard probability space (X, μ) , and let

$$\pi : (X, \mu, T) \longrightarrow (Y, \nu, S)$$

be a factor map of measure-preserving G -actions.

Theorem 15.1. — *For every $\varepsilon > 0$, there is another factor map $\tilde{\pi}$ of the G -system (X, μ, T) such that π is $\tilde{\pi}$ -measurable, (X, μ, T) is relatively Bernoulli over $\tilde{\pi}$, and*

$$h(\tilde{\pi}, \mu, T | \pi) < \varepsilon.$$

The deduction of Theorem 15.1 from the special case of the preceding subsection uses some machinery from orbit equivalence theory. The results we need can mostly be found in Danilenko and Park's development [15] of the 'orbital approach' to Ornstein theory for amenable group actions.

Proof of Theorem 15.1, and hence of Theorem A. — First, we may enlarge π with as little additional entropy as we please in order to assume that the G -system (Y, ν, S) is free (see [15, Theorem 5.4]).

Having done so, Connes, Feldman and Weiss' classic generalization of Dye's theorem [10] provides a single automorphism S_1 of (Y, ν) that has the same orbits as S . Moreover, we may choose S_1 to be isomorphic to any atomless and ergodic automorphism we like. In particular, we can insist that S_1 satisfy the conditions in assumption (A).

Since (Y, ν, S) is free, the factor map π restricts to a bijection on almost every individual orbit of T . Using these bijections we obtain a unique lift of S_1 to an automorphism T_1 of (X, μ) with the same orbits as T .

Now let

$$\tilde{\pi} : (\mathbf{X}, \mu, T_1) \longrightarrow (\tilde{\mathbf{Y}}, \tilde{\nu}, \tilde{S}_1)$$

be a new factor map provided by the special case of Theorem A that we have already proved. Then $\tilde{\pi}$ is also a factor map for the original G -action T , because the σ -algebra generated by $\tilde{\pi}$ contains the σ -algebra generated by π , and the orbit-equivalence cocycle is π -measurable. By Rudolph and Weiss' relative entropy theorem [86], we have

$$h(\tilde{\pi}, \mu, T | \pi) = h(\tilde{\pi}, \mu, T_1 | \pi) < \varepsilon.$$

On the other hand, the automorphism T_1 is relatively Bernoulli over $\tilde{\pi}$, and the orbit-equivalence cocycle is also $\tilde{\pi}$ -measurable. To finish, [15, Proposition 3.8] shows that this property of relative Bernoullicity really depends on only (i) the orbit equivalence relation on $(\tilde{\mathbf{Y}}, \tilde{\nu})$ inherited from (\mathbf{X}, μ, T_1) , and (ii) the associated cocycle of automorphisms of (\mathbf{X}, μ) . Therefore the automorphism T_1 is relatively Bernoulli over $\tilde{\pi}$ if and only if this holds for the G -action T . \square

16. Some known consequences and open questions

Since Thouvenot proposed the weak Pinsker property, it has been shown to have many consequences in ergodic theory. Here we recall some of these, now stated as unconditional results about ergodic automorphisms. We also present a couple of new results and collect a few open questions. Jean-Paul Thouvenot contributed a great deal to the contents of this section.

16.1. Results about joinings, factors, and isomorphisms. —

Proposition 16.1 (From [108]). — *Let \mathbf{X} and \mathbf{X}' be ergodic automorphisms, let \mathbf{B} be a Bernoulli shift with finite entropy, and suppose that $\mathbf{X} \times \mathbf{B}$ is isomorphic to $\mathbf{X}' \times \mathbf{B}$. Then \mathbf{X} and \mathbf{X}' are themselves isomorphic.*

(The statement of Proposition 16.1 in [108] is restricted to finite-entropy automorphisms, but one sees easily from the proof that the assumption of finite entropy for \mathbf{X} and \mathbf{X}' can be removed once we have Theorem A.)

The following question, also posed by Thouvenot, remains open.

Question 16.2. — *Does Proposition 16.1 still hold if both occurrences of ‘isomorphic’ are replaced with ‘Kakutani equivalent’?*

It is not clear whether the weak Pinsker property offers any insight here. The question is open even when \mathbf{X} and \mathbf{X}' both have entropy zero. The following related question may be even more basic, and is also open.

Question 16.3. — *Does every ergodic automorphism have an induced transformation that splits into a system of entropy zero and a Bernoulli system?*

Curiously, the weak Pinsker property also gives a result about canceling more general K -automorphisms from direct products in the opposite direction to Proposition 16.1.

Proposition 16.4. — *Let \mathbf{X} and \mathbf{X}' be ergodic systems of equal entropy h , possibly infinite, and let $\varepsilon > 0$. Then there is a system \mathbf{Z} of entropy less than ε such that*

$$\mathbf{X} \times \mathbf{Z} \quad \text{is isomorphic to} \quad \mathbf{X}' \times \mathbf{Z}.$$

If the two original systems are K -automorphisms, then \mathbf{Z} may be chosen to be K .

Proof. — Let $\mathbf{X}_0 := \mathbf{X}$ and $\mathbf{X}'_0 := \mathbf{X}'$. A recursive application of the weak Pinsker property gives sequences of systems \mathbf{X}_i and \mathbf{B}_i , $i \geq 1$, such that \mathbf{X}_i has entropy less than $2^{-i-1}\varepsilon$, each \mathbf{B}_i is Bernoulli, and such that

$$\mathbf{X}_i \quad \text{is isomorphic to} \quad \mathbf{X}_{i+1} \times \mathbf{B}_{i+1} \quad \forall i \geq 0.$$

If \mathbf{X}_0 has the K -property, then so do all its factors \mathbf{X}_i . Apply the same argument to \mathbf{X}'_0 to obtain \mathbf{X}'_i and \mathbf{B}'_i for $i \geq 1$. If \mathbf{X}'_0 has the K -property, then so does every \mathbf{X}'_i .

Now let

$$\mathbf{Z} := \mathbf{X}_1 \times \mathbf{X}_2 \times \cdots \times \mathbf{X}'_1 \times \mathbf{X}'_2 \times \cdots .$$

This has entropy less than

$$\sum_{i=1}^{\infty} 2^{-i-1}\varepsilon + \sum_{i=1}^{\infty} 2^{-i-1}\varepsilon = \varepsilon.$$

If \mathbf{X}_0 and \mathbf{X}'_0 have the K -property, then \mathbf{Z} is a direct product of K -automorphisms, so it also has the K -property.

To finish the proof, consider the direct product

$$\mathbf{X}_0 \times \mathbf{Z} = \mathbf{X}_0 \times \mathbf{X}_1 \times \mathbf{X}_2 \times \cdots \times \mathbf{X}'_1 \times \mathbf{X}'_2 \times \cdots .$$

In the product on the right, we know that each coordinate factor \mathbf{X}_i may be replaced with $\mathbf{X}_{i+1} \times \mathbf{B}_{i+1}$ up to isomorphism. Therefore this product is isomorphic to

$$(\mathbf{X}_1 \times \mathbf{B}_1) \times (\mathbf{X}_2 \times \mathbf{B}_2) \times \cdots \times \mathbf{X}'_1 \times \mathbf{X}'_2 \times \cdots .$$

Re-ordering the factors, this is isomorphic to

$$\mathbf{X}_1 \times \mathbf{X}_2 \times \cdots \times \mathbf{X}'_1 \times \mathbf{X}'_2 \times \cdots \times \mathbf{B}_1 \times \mathbf{B}_2 \times \cdots ,$$

which is equal to the direct product of \mathbf{Z} with a Bernoulli shift. Similarly, $\mathbf{X}'_0 \times \mathbf{Z}$ is isomorphic to the direct product of \mathbf{Z} and a Bernoulli shift. Since $\mathbf{X}_0 \times \mathbf{Z}$ and $\mathbf{X}'_0 \times \mathbf{Z}$ have the same entropy, it follows that they are isomorphic to each other. \square

Proposition 16.5 (From [94]). — *If \mathbf{X} is ergodic and has positive entropy, then it has two Bernoulli factors that together generate the whole σ -algebra modulo μ ; equivalently, \mathbf{X} is isomorphic to a joining of two Bernoulli systems. For any $\varepsilon > 0$, we may choose one of those Bernoulli factors to have entropy less than ε .*

The next result answers a question of Benjamin Weiss which has circulated informally for some time. I was introduced to the question by Yonatan Gutman. The neat proof below was shown to me by Jean-Paul Thouvenot. Thouvenot discusses the question further as problem (1) in [109].

Proposition 16.6. — *If \mathbf{X} and \mathbf{X}' are K -automorphisms of equal entropy h , possibly infinite, then they have a common extension that is also a K -automorphism of entropy h .*

Proof. — The desired conclusion depends on our two systems only up to isomorphism. Therefore, by the weak Pinsker property, we may assume that

$$\mathbf{X} = \mathbf{X}_1 \times \mathbf{B}$$

and

$$\mathbf{X}' = \mathbf{X}'_1 \times \mathbf{B}'$$

where \mathbf{B} and \mathbf{B}' are Bernoulli and where the entropies $a := h(\mathbf{X}_1)$ and $b := h(\mathbf{X}'_1)$ satisfy $a + b < h$. The new systems \mathbf{X}_1 and \mathbf{X}'_1 inherit the K -property from the original systems.

Let \mathbf{B}'' be a Bernoulli shift of entropy $h - a - b$, interpreting this as ∞ if $h = \infty$. Now consider the system

$$\tilde{\mathbf{X}} := \mathbf{X}_1 \times \mathbf{X}'_1 \times \mathbf{B}''.$$

This is a product of K -automorphisms, hence still has the K -property. Its entropy is

$$a + b + (h - a - b) = h.$$

Finally, it admits both of our original systems as factors. Indeed, by Sinai's theorem, there is a factor map from \mathbf{X}'_1 to another Bernoulli shift \mathbf{B}'_1 of entropy b . Applying this map to the middle coordinate of $\tilde{\mathbf{X}}$ gives a factor map from $\tilde{\mathbf{X}}$ to

$$\mathbf{X}_1 \times \mathbf{B}'_1 \times \mathbf{B}'',$$

which is isomorphic to $\mathbf{X}_1 \times \mathbf{B}$ and hence to \mathbf{X} . The argument for \mathbf{X}' is analogous. \square

Given an automorphism (X, μ, T) , another object of interest in ergodic theory is its **centralizer**: the group of all other μ -preserving transformations of X that commute with T . Informally, it is known that the centralizer of a Bernoulli system is very large. A precise result in this direction is due to Rudolph, who proved in [85] that if T is

Bernoulli and S is another automorphism that commutes with the whole centralizer of T , then S must be a power of T .

If (X, μ, T) is any ergodic automorphism with positive entropy, then Theorem A provides a non-trivial splitting

$$(102) \quad (X, \mu, T) \xrightarrow{\text{iso}} (Y, \nu, S) \times \mathbf{B}.$$

Under this splitting, we obtain many elements of the centralizer of T of the form $S^m \times U$, where $m \in \mathbf{Z}$ and U is any centralizer element of \mathbf{B} . In the first place, this answers a simple question posed in [85]: if (X, μ, T) has positive entropy, then its centralizer contains many elements that are not just powers of T . I understand from Jean-Paul Thouvenot that by combining these centralizer elements with unpublished work of Ferenczi and the ideas in [94, Remark 2], one can also extend Rudolph's main result from [85] to any ergodic automorphism of positive entropy. We do not explore this fully here, but leave the reader with the next natural question:

Question 16.7. — *Is there a K -automorphism (X, μ, T) such that every element of the centralizer arises as $S^m \times U$ for some splitting as in (102)?*

An old heuristic of Ornstein asserts that the classification of factors within a fixed Bernoulli shift should mirror the complexity of the classification of general K -automorphisms: see the discussion in [74] and the examples in [34]. Can the methods of the present paper shed additional light on the lattice of all factors of a fixed Bernoulli shift?

Finally, here is a question that seems to lie strictly beyond the weak Pinsker property. Theorem A promises that any ergodic system with positive entropy may be split into a non-trivial direct product in which one factor is Bernoulli.

Question 16.8. — *Is there a non-Bernoulli K -automorphism with the property that any non-trivial splitting of it consists of one Bernoulli and one non-Bernoulli factor?*

16.2. Results about generating observables and processes. — Recall that a process $\alpha : (X, \mu, T) \longrightarrow (A^{\mathbf{Z}}, \alpha_*\mu, T_A)$ is uniquely determined by its time-zero observable α_0 . That observable is called **generating** if α is an isomorphism of systems.

Proposition 16.9 (From [108]). — *Any ergodic system of finite entropy has a generating observable whose two-sided tail equals the Pinsker factor.*

Proposition 16.10 (From [78]). — *Any ergodic system of finite entropy has a generating observable whose one-sided super-tail equals the Pinsker factor (see [78] for the relevant definitions).*

Proposition 16.11 (From [111]). — *For any ergodic finite-state process, Vershik’s secondary entropy (defined in [111]) grows slower than any doubly exponential function (although it can grow faster than any single exponential [62]).*

The next result is a consequence of Proposition 16.6. It was brought to my attention by Yonatan Gutman, who established the connection together with Michael Hochman.

Corollary 16.12 (From [29]). — *The class of finite-entropy, non-Bernoulli K -processes has no finitely observable invariants other than entropy (see [29] for the relevant definitions).*

If (X, μ, T) is a system and $\alpha : X \rightarrow A^{\mathbb{Z}}$ is a finite-valued generating process, then the conditional entropy $H_{\mu}(\alpha_0 | \alpha_{(-\infty; 0)})$ equals the Kolmogorov–Sinai entropy $h(\mu, T)$. However, if the process is real-valued (so $A = \mathbf{R}$), then in general one knows only that

$$H_{\mu}(\alpha_0 | \alpha_{(-\infty; 0)}) \leq h(\mu, T).$$

In connection with these quantities, the next result answers an early question of Rokhlin and Sinai [83, paragraph 12.7].

Proposition 16.13. — *If (X, μ, T) is a finite-entropy K -automorphism, then for any $c \in (0, h(\mu, T)]$ it has a real-valued generating process $\alpha : X \rightarrow \mathbf{R}^{\mathbb{Z}}$ such that*

$$H_{\mu}(\alpha_0 | \alpha_{(-\infty; 0)}) = c.$$

Proof. — If $c = h(\mu, T)$ then we simply take a finite-valued generating observable, as provided by Krieger’s generator theorem. So let us assume that $c < h(\mu, T)$.

In case (X, μ, T) is Bernoulli, the result is proved by Lindenstrauss, Peres and Schlag in [55] using an elegant construction of α in terms of Bernoulli convolutions.

For the general case, by Theorem A we may split (X, μ, T) into a factor of entropy less than $h(\mu, T) - c$ and a Bernoulli factor. Take any finite-valued generating observable of the low-entropy factor, and apply the result from [55] to the Bernoulli factor. The resulting product observable gives conditional entropy equal to c . \square

Another natural question about generating processes asks for a refinement of Theorem A. It was suggested to me by Gábor Pete.

Question 16.14. — *Let $(A^{\mathbb{Z}}, \mu, T_A)$ be an ergodic finite-state process and let $\varepsilon > 0$. Must the process be finitarily isomorphic to a direct product of the form*

$$\underbrace{(C^{\mathbb{Z}}, \nu, T_C)}_{\text{entropy} < \varepsilon} \times \underbrace{(B^{\mathbb{Z}}, \rho^{\times \mathbb{Z}}, T_B)}_{\text{i.i.d.}}?$$

If this does not always hold, are there natural sufficient conditions?

What if we ask instead for an isomorphism that is unilateral in one direction or the other?

16.3. *More general groups and actions.* — Many of the corollaries listed above should extend to actions of general countable amenable groups. Indeed, the orbital methods of [15] yield some such generalizations quite easily: for instance, Proposition 16.5 above should imply its extension to actions of amenable groups using the same arguments as for [15, Theorem 6.1].

Here is another consequence of the weak Pinsker property for actions of an arbitrary countable amenable group G :

Proposition 16.15. — *Let G be a countable amenable group. For any $\varepsilon > 0$, any free and ergodic G -system (X, μ, T) of finite entropy h has a generating observable α with the property that*

$$(103) \quad H_\mu\left(\alpha \mid \bigvee_{g \in G \setminus \{e\}} \alpha \circ T^g\right) > h - \varepsilon.$$

To prove this, simply combine the two time-zero observables in a direct product of a low-entropy system and a Bernoulli shift. Proposition 16.15 resulted from discussions with Brandon Seward about ‘percolative entropy’, a quantity of current interest in the ergodic theory of non-amenable groups. In case $G = \mathbf{Z}$, the left-hand side of (103) is the same as Verdú and Weissman’s ‘erasure entropy rate’ from [110]. In that setting, Proposition 16.15 asserts that any stationary ergodic source with a finite alphabet admits generating observables whose erasure entropy rate is arbitrarily close to the Kolmogorov–Sinai entropy. This can be seen as an opposite to Ornstein and Weiss’ celebrated result that any such source has a generating observable which is bilaterally deterministic [76].

Fieldsteel showed in [21, Theorem 3] that an ergodic flow $(T^t)_{t \in \mathbf{R}}$ (that is, a bi-measurable and measure-preserving action of \mathbf{R}) has the weak Pinsker property if and only if the map T^t is ergodic and has the weak Pinsker property for at least one value of t . Combined with Theorem A, this immediately gives the following.

Proposition 16.16. — *Every ergodic flow has the weak Pinsker property.*

Do Theorem 15.1 and Proposition 16.16 have a common generalization to actions of a larger class of locally compact, second countable groups? It would be natural to start with those that have a ‘good entropy theory’ [77]. Some relevant techniques may be available in [4]. In another direction, is there a sensible generalization to non-free actions?

Beyond the class of amenable groups one cannot hope for too much. For sofic groups, one can ask after a version of the weak Pinsker property with sofic entropy [8] or Rokhlin entropy [89] in place of KS entropy, but Lewis Bowen has found natural examples of actions of non-amenable free groups for which those properties fail.

Acknowledgements

My sincere thanks go to Jean-Paul Thouvenot for sharing his enthusiasm for the weak Pinsker property with me, and for patient and thoughtful answers to many questions over the last few years. I have also been guided throughout this project by some very wise advice from Dan Rudolph, whom I had the great fortune to meet before the end of his life. More recently, I had helpful conversations with Brandon Seward, Mike Saks, Ofer Zeitouni and Lewis Bowen.

Lewis Bowen also read an early version of this paper carefully and suggested several important changes and corrections. Other valuable suggestions came from Yonatan Gutman, Jean-Paul Thouvenot and an anonymous referee.

Much of the work on this project was done during a graduate research fellowship from the Clay Mathematics Institute.

REFERENCES

1. L. M. ABRAMOV, The entropy of a derived automorphism, *Dokl. Akad. Nauk SSSR*, **128** (1959), 647–650.
2. R. AHLWEDE, P. GÁCS and J. KÖRNER, Bounds on conditional probabilities with applications in multi-user communication, *ℳ. Wahrscheinlichkeitstheor. Verw. Geb.*, **34** (1976), 157–177.
3. S. AIDA, T. MASUDA and I. SHIGEKAWA, Logarithmic Sobolev inequalities and exponential integrability, *J. Funct. Anal.*, **126** (1994), 83–101.
4. N. AVNI, Entropy theory for cross-sections, *Geom. Funct. Anal.*, **19** (2010), 1515–1538.
5. P. BILLINGSLEY, *Ergodic Theory and Information*, Wiley, New York, 1965.
6. S. G. BOBKOV and F. GÖTZE, Exponential integrability and transportation cost related to logarithmic Sobolev inequalities, *J. Funct. Anal.*, **163** (1999), 1–28.
7. B. BOLLOBÁS, *Modern Graph Theory*, Springer, Berlin, 1998.
8. L. BOWEN, Measure conjugacy invariants for actions of countable sofic groups, *J. Am. Math. Soc.*, **23** (2010), 217–245.
9. F. R. K. CHUNG, R. L. GRAHAM, P. FRANKL and J. B. SHEARER, Some intersection theorems for ordered sets and graphs, *J. Comb. Theory, Ser. A*, **43** (1986), 23–37.
10. A. CONNES, J. FELDMAN and B. WEISS, An amenable equivalence relation is generated by a single transformation, *Ergod. Theory Dyn. Syst.*, **1** (1982), 431–450, 1981.
11. T. M. COVER and J. A. THOMAS, *Elements of Information Theory*, 2nd ed., Wiley–Interscience, Hoboken, 2006.
12. G. CROOKS, On measures of entropy and information. Technical note, available online at threeplusone.com/on_information.pdf.
13. I. CSISZÁR, Information-type measures of difference of probability distributions and indirect observations, *Studia Sci. Math. Hung.*, **2** (1967), 299–318.
14. I. CSISZÁR, Sanov property, generalized I-projection and a conditional limit theorem, *Ann. Probab.*, **12** (1984), 768–793.
15. A. I. DANILENKO and K. K. PARK, Generators and Bernoullian factors for amenable actions and cocycles on their orbits, *Ergod. Theory Dyn. Syst.*, **22** (2002), 1715–1745.
16. A. DEMBO, Information inequalities and concentration of measure, *Ann. Probab.*, **25** (1997), 927–939.
17. R. M. DUDLEY, *Real Analysis and Probability*, Cambridge Studies in Advanced Mathematics, vol. 74, Cambridge Univ. Press, Cambridge, 2002, revised reprint of the 1989 original.
18. D. ELLIS, E. FRIEDGUT, G. KINDLER and A. YEHUDAYOFF, Geometric stability via information theory, *Discrete Anal.*, **10** (2016), 28.
19. W. FELLER, *An Introduction to Probability Theory and Its Applications, vol. I*, 3rd ed., Wiley, New York, 1968.
20. W. FELLER, *An Introduction to Probability Theory and Its Applications, vol. II*, 2nd ed., Wiley, New York, 1971.
21. A. FIELDSTEEL, Stability of the weak Pinsker property for flows, *Ergod. Theory Dyn. Syst.*, **4** (1984), 381–390.

22. H. FURSTENBERG, Ergodic behaviour of diagonal measures and a theorem of Szemerédi on arithmetic progressions, *J. Anal. Math.*, **31** (1977), 204–256.
23. W. GANGBO and R. J. McCANN, The geometry of optimal transportation, *Acta Math.*, **177** (1996), 113–161.
24. W. T. GOWERS, Lower bounds of tower type for Szemerédi’s uniformity lemma, *Geom. Funct. Anal.*, **7** (1997), 322–337.
25. W. T. GOWERS, Rough structure and classification, in *Geom. Funct. Anal., Special Volume, Part I (GAFA 2000, Tel Aviv, 1999)*, pp. 79–117, 2000.
26. N. GOZLAN and C. LÉONARD, Transport inequalities. A survey, *Markov Process. Relat. Fields*, **16** (2010), 635–736.
27. M. GROMOV, Endomorphisms of symbolic algebraic varieties, *J. Eur. Math. Soc.*, **1** (1999), 109–197.
28. M. GROMOV, *Metric Structures for Riemannian and Non-Riemannian Spaces*, 2nd ed., Birkhäuser, Boston, 2001.
29. Y. GUTMAN and M. HOCHMAN, On processes which cannot be distinguished by finite observation, *Isr. J. Math.*, **164** (2008), 265–284.
30. P. R. HALMOS, *Lectures on Ergodic Theory*, Chelsea, New York, 1960.
31. T. S. HAN, Linear dependence structure of the entropy space, *Inf. Control*, **29** (1975), 337–368.
32. T. S. HAN, Nonnegative entropy measures of multivariate symmetric correlations, *Inf. Control*, **36** (1978), 133–156.
33. W. HOEFFDING, Probability inequalities for sums of bounded random variables, *J. Am. Stat. Assoc.*, **58** (1963), 13–30.
34. C. HOFFMAN, The behavior of Bernoulli shifts relative to their factors, *Ergod. Theory Dyn. Syst.*, **19** (1999), 1255–1280.
35. S. KALIKOW, Non-intersecting splitting σ -algebras in a non-Bernoulli transformation, *Ergod. Theory Dyn. Syst.*, **32** (2012), 691–705.
36. S. KALIKOW and R. McCUTCHEON, *An Outline of Ergodic Theory*, Cambridge Studies in Advanced Mathematics, vol. 122, Cambridge Univ. Press, Cambridge, 2010.
37. S. A. KALIKOW, T, T^{-1} transformation is not loosely Bernoulli, *Ann. Math. (2)*, **115** (1982), 393–409.
38. L. V. KANTOROVICH, On a problem of Monge, *Žap. Nauč. Semin. POMI*, **312** (2004), 15–16.
39. L. V. KANTOROVICH and G. v. RUBINSHTEIN, On a functional space and certain extremum problems, *Dokl. Akad. Nauk SSSR*, **115** (1957), 1058–1061.
40. L. V. KANTOROVITCH, On the translocation of masses, *C. R. (Dokl.) Acad. Sci. URSS*, **37** (1942), 199–201.
41. A. KATOK, Fifty years of entropy in dynamics: 1958–2007, *J. Mod. Dyn.*, **1** (2007), 545–596.
42. J. H. B. KEMPERMAN, On the optimum rate of transmitting information, in *Proc. Int. Sympos. on Probability and Information Theory (McMaster Univ., Hamilton, ON, 1968)*, pp. 126–169, Springer, Berlin, 1969.
43. J. C. KIEFFER, A direct proof that VWB processes are closed in the \bar{d} -metric, *Isr. J. Math.*, **41** (1982), 154–160.
44. J. C. KIEFFER, A simple development of the Thouvenot relative isomorphism theory, *Ann. Probab.*, **12** (1984), 204–211.
45. A. N. KOLMOGOROV, A new metric invariant of transient dynamical systems and automorphisms in Lebesgue spaces, *Dokl. Akad. Nauk SSSR*, **119** (1958), 861–864.
46. A. N. KOLMOGOROV, Entropy per unit time as a metric invariant of automorphisms, *Dokl. Akad. Nauk SSSR*, **124** (1959), 754–755.
47. W. KRIEGER, On entropy and generators of measure-preserving transformations, *Trans. Am. Math. Soc.*, **149** (1970), 453–464.
48. S. KULLBACK, Certain inequalities in information theory and the Cramér–Rao inequality, *Ann. Math. Stat.*, **25** (1954), 745–751.
49. S. KULLBACK, A lower bound for discrimination information in terms of variation, *IEEE Trans. Inf. Theory*, **T-13** (1967), 126.
50. S. KULLBACK and R. A. LEIBLER, On information and sufficiency, *Ann. Math. Stat.*, **22** (1951), 79–86.
51. M. LEDOUX, Remarks on logarithmic Sobolev constants, exponential integrability and bounds on the diameter, *J. Math. Kyoto Univ.*, **35** (1995), 211–220.
52. M. LEDOUX, On Talagrand’s deviation inequalities for product measures, *ESAIM Probab. Stat.*, **1** (1995), 63–87.
53. M. LEDOUX, *The Concentration of Measure Phenomenon*, Mathematical Surveys and Monographs, vol. 89, Am. Math. Soc., Providence, 2001.
54. V. L. LEVIN, General Monge–Kantorovich problem and its applications in measure theory and mathematical economics, in *Functional Analysis, Optimization, and Mathematical Economics*, pp. 141–176, Oxford Univ. Press, New York, 1990.
55. E. LINDENSTRAUSS, Y. PERES and W. SCHLAG, Bernoulli convolutions and an intermediate value theorem for entropies of K -partitions, *J. Anal. Math.*, **87** (2002), 337–367, dedicated to the memory of Thomas H. Wolff.
56. N. LINIAL, A. SAMORODNITSKY and A. WIGDERSON, A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents, *Combinatorica*, **20** (2000), 545–568.

57. K. MARTON, A simple proof of the blowing-up lemma, *IEEE Trans. Inf. Theory*, **32** (1986), 445–446.
58. K. MARTON, Bounding \bar{d} -distance by informational divergence: a method to prove measure concentration, *Ann. Probab.*, **24** (1996), 857–866.
59. K. MARTON, A measure concentration inequality for contracting Markov chains, *Geom. Funct. Anal.*, **6** (1996), 556–571.
60. K. MARTON, Measure concentration for a class of random processes, *Probab. Theory Relat. Fields*, **110** (1998), 427–439.
61. K. MARTON and P. C. SHIELDS, The positive-divergence and blowing-up properties, *Isr. J. Math.*, **86** (1994), 331–348.
62. K. MARTON and P. C. SHIELDS, How many future measures can there be? *Ergod. Theory Dyn. Syst.*, **22** (2002), 257–280.
63. C. McDIARMID, On the method of bounded differences, in *Surveys in Combinatorics (Norwich, 1989)*, London Math. Soc. Lecture Note Ser., vol. 141, pp. 148–188, Cambridge Univ. Press, Cambridge, 1989.
64. W. J. MCGILL, Multivariate information transmission, *Trans. IRE Profess. Group Inf. Theory*, **4** (1954), 93–111.
65. V. D. MILMAN, The heritage of P. Lévy in geometrical functional analysis, *Astérisque*, **157–158** (1988), 273–301.
66. V. D. MILMAN and G. SCHECHTMAN, *Asymptotic Theory of Finite-Dimensional Normed Spaces*, Lecture Notes in Mathematics, vol. 1200, Springer, Berlin, 1986, with an appendix by M. Gromov.
67. D. ORNSTEIN, Bernoulli shifts with the same entropy are isomorphic, *Adv. Math.*, **4** (1970), 337–352, 1970.
68. D. ORNSTEIN, Two Bernoulli shifts with infinite entropy are isomorphic, *Adv. Math.*, **5** (1970), 339–348, 1970.
69. D. ORNSTEIN, Newton’s laws and coin tossing, *Not. Am. Math. Soc.*, **60** (2013), 450–459.
70. D. S. ORNSTEIN, An example of a Kolmogorov automorphism that is not a Bernoulli shift, *Adv. Math.*, **10** (1973), 49–62.
71. D. S. ORNSTEIN, A K automorphism with no square root and Pinsker’s conjecture, *Adv. Math.*, **10** (1973), 89–102.
72. D. S. ORNSTEIN, A mixing transformation for which Pinsker’s conjecture fails, *Adv. Math.*, **10** (1973), 103–123.
73. D. S. ORNSTEIN, *Ergodic Theory, Randomness, and Dynamical Systems*, Yale Mathematical Monographs, vol. 5, Yale Univ. Press, New Haven, 1974.
74. D. S. ORNSTEIN, Factors of Bernoulli shifts, *Isr. J. Math.*, **21** (1975), 145–153, 1974.
75. D. S. ORNSTEIN and P. C. SHIELDS, An uncountable family of K-automorphisms, *Adv. Math.*, **10** (1973), 63–88.
76. D. S. ORNSTEIN and B. WEISS, Every transformation is bilaterally deterministic, *Isr. J. Math.*, **21** (1975), 154–158, 1974.
77. D. S. ORNSTEIN and B. WEISS, Entropy and isomorphism theorems for actions of amenable groups, *J. Anal. Math.*, **48** (1987), 1–141.
78. K. PETERSEN and J.-P. THOUVENOT, Tail fields generated by symbol counts in measure-preserving systems, *Colloq. Math.*, **101** (2004), 9–23.
79. M. S. PINSKER, Dynamical systems with completely positive or zero entropy, *Sov. Math. Dokl.*, **1** (1960), 937–938.
80. M. S. PINSKER, *Information and Information Stability of Random Variables and Processes*, Holden-Day, San Francisco, 1964.
81. J. RADHAKRISHNAN, Entropy and counting, in J. Mishra (ed.) *Computational Mathematics, Modelling and Applications, IIT Kharagpur, Golden Jubilee Volume*, pp. 146–168, Narosa, New York, 2003, available online via www.tcs.tifr.res.in/~jaikumar/Papers.
82. M. RAHE, Relatively finitely determined implies relatively very weak bernoulli, *Can. J. Math.*, **30** (1978), 531–548.
83. V. A. ROHLIN, Lectures on the entropy theory of transformations with invariant measure, *Usp. Mat. Nauk*, **22** (1967), 3–56.
84. D. J. RUDOLPH, If a two-point extension of a Bernoulli shift has an ergodic square, then it is Bernoulli, *Isr. J. Math.*, **30** (1978), 159–180.
85. D. J. RUDOLPH, The second centralizer of a Bernoulli shift is just its powers, *Isr. J. Math.*, **29** (1978), 167–178.
86. D. J. RUDOLPH and B. WEISS, Entropy and mixing for amenable group actions, *Ann. Math. (2)*, **151** (2000), 1119–1150.
87. P.-M. SAMSON, Concentration of measure inequalities for Markov chains and Φ -mixing processes, *Ann. Probab.*, **28** (2000), 416–461.
88. I. N. SANOV, On the probability of large deviations of random variables, in *Select. Transl. Math. Statist. and Probability, Vol. 1*, pp. 213–244, Inst. Math. Statist. and Amer. Math. Soc., Providence, R.I., 1961.
89. B. SEWARD, Krieger’s finite generator theorem for actions of countable groups I. [arXiv:1405.3604](https://arxiv.org/abs/1405.3604).
90. P. SHIELDS, *The Theory of Bernoulli Shifts*, Chicago Lectures in Mathematics, Univ. of Chicago Press, Chicago, 1973.
91. P. SHIELDS, *The Ergodic Theory of Discrete Sample Paths*, Graduate Studies in Mathematics, vol. 13, Am. Math. Soc., Providence, 1996.
92. J. SINAI, On the concept of entropy for a dynamic system, *Dokl. Akad. Nauk SSSR*, **124** (1959), 768–771.
93. J. G. SINAI, On a weak isomorphism of transformations with invariant measure, *Mat. Sb. (N.S.)*, **63** (1964), 23–42.

94. M. SMORODINSKY and J.-P. THOUVENOT, Bernoulli factors that span a transformation, *Isr. J. Math.*, **32** (1979), 39–43.
95. E. SZEMERÉDI, On sets of integers containing no k elements in arithmetic progression, *Acta Arith.*, **27** (1975), 199–245.
96. E. SZEMERÉDI, Regular partitions of graphs, in *Problèmes combinatoires et théorie des graphes*, (Univ. Orsay, Orsay, 1976). Colloq. Internat. CNRS, vol. 260, pp. 399–401, CNRS, Paris, 1978.
97. M. TALAGRAND, On Russo’s approximate zero-one law, *Ann. Probab.*, **22** (1994), 1576–1587.
98. M. TALAGRAND, Concentration of measure and isoperimetric inequalities in product spaces, *Inst. Hautes Études Sci. Publ. Math.*, **81** (1995), 73–205.
99. M. TALAGRAND, A new look at independence, *Ann. Probab.*, **24** (1996), 1–34.
100. M. TALAGRAND, Transportation cost for Gaussian and other product measures, *Geom. Funct. Anal.*, **6** (1996), 587–600.
101. T. TAO, Szemerédi’s regularity lemma revisited, *Contrib. Discrete Math.*, **1** (2006), 8–28.
102. T. TAO, The dichotomy between structure and randomness, arithmetic progressions, and the primes, in *International Congress of Mathematicians, vol. I*, pp. 581–608, Eur. Math. Soc., Zürich, 2007.
103. T. TAO and V. VU, *Additive Combinatorics*, Cambridge Univ. Press, Cambridge, 2006.
104. J.-P. THOUVENOT, Quelques propriétés des systèmes dynamiques qui se décomposent en un produit de deux systèmes dont l’un est un schéma de Bernoulli, *Isr. J. Math.*, **21** (1975), 177–207.
105. J.-P. THOUVENOT, Une classe de systèmes pour lesquels la conjecture de Pinsker est vraie, *Isr. J. Math.*, **21** (1975), 208–214, 1974.
106. J.-P. THOUVENOT, On the stability of the weak Pinsker property, *Isr. J. Math.*, **27** (1977), 150–162.
107. J.-P. THOUVENOT, Entropy, isomorphism and equivalence in ergodic theory, in *Handbook of Dynamical Systems, vol. 1A*, pp. 205–238, North-Holland, Amsterdam, 2002.
108. J.-P. THOUVENOT, Two facts concerning the transformations which satisfy the weak Pinsker property, *Ergod. Theory Dyn. Syst.*, **28** (2008), 689–695.
109. J.-P. THOUVENOT, Relative spectral theory and measure-theoretic entropy of Gaussian extensions, *Fundam. Math.*, **206** (2009), 287–298.
110. S. VERDÚ and T. WEISSMAN, The information lost in erasures, *IEEE Trans. Inf. Theory*, **54** (2008), 5030–5058.
111. A. M. VERSHIK, Dynamic theory of growth in groups: entropy, boundaries, examples, *Usp. Mat. Nauk*, **55** (2000), 59–128.
112. A. M. VERSHIK, The Kantorovich metric: the initial history and little-known applications, *Žap. Nauč. Semin. POMI*, **312** (2004), 69–85.
113. A. M. VERSHIK, Dynamics of metrics in measure spaces and their asymptotic invariants, *Markov Process. Relat. Fields*, **16** (2010), 169–184.
114. R. VERSHYNIN, High-Dimensional Probability: An Introduction with Applications in Data Science, Cambridge Univ. Press (in press), draft available online at <https://www.math.uci.edu/~rvershyn/papers/HDP-book/HDP-book.html>.
115. S. WATANABE, Information theoretical analysis of multivariate correlation, *IBM J. Res. Dev.*, **4** (1960), 66–82.
116. D. WELSH, *Codes and Cryptography*, Oxford Science Publications, Clarendon/Oxford Univ. Press, New York, 1988.

T. A.
 Department of Mathematics,
 University of California, Los Angeles,
 Los Angeles, CA 90095-1555, USA
tim@math.ucla.edu

Manuscrit reçu le 16 octobre 2017
Manuscrit accepté le 3 février 2018
publié en ligne le 15 février 2018.