

CHAPTER 5

***m*-entropy**

We define the m -entropy of T^α to be the infimum of entropies of T^β , over arrangements β in the m -equivalence class of α . Although our earlier work has made it essentially obvious, we will show that m -entropy is upper semi-continuous in m_α and hence this infimum is obtained on a dense G_δ subset of the equivalence class $E_m(\alpha)$. As in earlier work [20, 36], the main goal of this section is to show that for a fixed size m , either the m -entropy of T^α is zero, for all arrangements α , or the m -entropy is the usual entropy for all arrangements α . We will say that each size m is either entropy-free or entropy-preserving.

Let m be a fixed size. Let α be a G -arrangement. Recall that $E_m(\alpha) = \{\beta | \alpha \overset{m}{\sim} \beta\}$ is the m -equivalence class of α .

DEFINITION 5.0.40. *Define the **m-entropy** of T^α to be*

$$h_m(T^\alpha) = \inf\{h(T^\beta); \beta \in E_m(\alpha)\}.$$

We first argue that this m -entropy is attained residually in $E_m(\alpha)$.

The following lemma tells us that if two arrangements are m -close then they are, in fact, close in distribution.

LEMMA 5.0.41. *For every $\varepsilon > 0$ there exists $\delta > 0$ such that if*

$$m(\alpha, \beta) < \delta$$

then, for any partition P ,

$$\|(T^\alpha, P), (T^\beta, P)\|_* < \varepsilon.$$

PROOF. This follows rather easily from lemma 2.2.9 as if $P(T_g^\alpha(x)) \neq P(T_g^\beta(x))$ then in particular $\alpha(x, T_g^\beta(x)) \neq g$. \square

COROLLARY 5.0.42. *Fix any G -arrangement α and partition P . Let $\varepsilon > 0$. There exists $\delta > 0$ such that if $m(\alpha, \beta) < \delta$ (i.e. if $\beta \in B_\delta(\alpha)$), then $h(T^\beta, P) < h(T^\alpha, P) + \varepsilon$.*

PROOF. This is a direct consequence of Lemma 5.0.41 and Theorem 3.0.24. \square

THEOREM 5.0.43. *The set of $\beta \in E_m(\alpha)$ with $h(T^\beta) = h_m(T^\alpha)$ is a dense G_δ in $E_m(\alpha)$. As $E_m(\alpha)$ is Polish, this infimum is achieved residually.*

PROOF. Fix α and P . Let

$$\mathcal{O}_{\varepsilon, P} = \{\beta \in E_m(\alpha); h(T^\beta, P) < h_m(T^\alpha) + \varepsilon\}.$$

If $h(T^\beta) = h_m(T^\alpha)$ then $\beta \in \mathcal{O}_{\varepsilon, P}$ for all $\varepsilon > 0$ and finite partitions P . On the other hand, letting $\{P_i\}$ be a countable family of partitions labelled by finite subsets of \mathbb{N} , dense in the L^1 -metric on partitions. As $h(T^\beta, P)$ is continuous in the L^1 -metric, if $\beta \in \mathcal{O}_{1/k, P_i}$ for all k and i then we will have

$$h(T^\beta, P) \leq h_m(T^\alpha)$$

for all P and hence $h(T^\beta) = h_m(T^\alpha)$.

All we need to see is that the sets $\mathcal{O}_{\varepsilon, P}$ are open. This though is just upper semi-continuity of entropy $h(\tilde{\mu})$ for σ -invariant and ergodic measures on Σ_P^G , our Theorem 3.0.24. □

Our goal now is to show that if ever $h_m(T^\alpha) < h(T^\alpha)$, for some α , then for this size m and for **any** G -arrangement β , we have $h_m(T^\beta) = 0$. Note that this β need not be in the m -equivalence class $E_m(\alpha)$ or even on the orbits of the same free and ergodic action. Although our path to this is somewhat technical we gain a lot of insight into the relation between a size and its entropy along the way.

First we show that if entropy can be lowered by moving to another m -equivalent arrangement, then relative to a partition, it can be lowered by an element of the full group.

LEMMA 5.0.44. *Let $e > 0$. Suppose, for some α ,*

$$h_m(T^\alpha) < h(T^\alpha) - e.$$

Let P be any partition. For any $\varepsilon > 0$, there exists Γ such that $m(\alpha, \phi) < \varepsilon$ and $h(T^{\alpha\phi}, P) < h(T^\alpha) - e$.

PROOF. As the full-group acts minimally on $E_m(\alpha)$ and preserves entropy, those $\beta \in E_m(\alpha)$ with

$$h(T^\beta) < h(T^\alpha) - e$$

are a dense set. Thus those β with

$$h(T^\beta, P) < h(T^\alpha) - e$$

are a dense and open set. On the other hand, those $\alpha\phi$ with

$$m(\alpha, \phi) < \varepsilon$$

are dense in $B_\varepsilon(\alpha)$, the ball of radius ε in m about α . Intersecting these two sets yields the result. \square

Next we show that if entropy can be lowered, then it can be lowered on an i.i.d. process. We are following the argument of [20] and [36] here. Knowing that the m -entropy of a Bernoulli action is less than its entropy will allow us to conclude this for all ergodic systems of positive entropy.

LEMMA 5.0.45. *Let α be any arrangement. Suppose there exists $e > 0$ such that $h_m(T^\alpha) < h(T^\alpha) - e$. Let (U^β, P) be any i.i.d. process with $h(U^\beta, P) = h(T^\alpha)$. For any $\varepsilon > 0$, there exists ϕ in the full-group of U^β such that $m(\beta, \phi) < \varepsilon$ and $h(U^{\beta\phi}, P) < h(U^\beta, P) - e$.*

PROOF. Fix $\alpha \in \mathcal{A}$ with $h(T^\alpha) > 0$. Suppose $e > 0$ such that

$$h_m(T^\alpha) < h(T^\alpha) - e.$$

Let (U^β, P) be any i.i.d. process with $h(U^\beta, P) = h(T^\alpha)$. Let $\varepsilon > 0$.

Let $\eta > 0$, so that $h_m(T^\alpha) < h(T^\alpha) - e - \eta$. At this point, as in both [20] and [36] we seem forced to introduce a circularity in that we want to apply Sinai's Theorem to say we can obtain a full-entropy i.i.d. factor of T^α . This is of course a known theorem, so no circularity is really obtained. We also note that even within the confines of our work here no circularity is obtained as for the size m of conjugacy, m -entropy is of course entropy and the Sinai's Theorem (our Theorem 7.3.3) for this strongest of equivalences will hold without further ado concerning m -entropy. Apply the Ornstein-Weiss, Sinai Theorem [28] to see that there exists a partition P' such that $h(T^\alpha, P') = h(T^\alpha)$ and (T^α, P') has the same i.i.d. distribution as (U^β, P) , that is to say

$$\|(T^\alpha, P'), (U^\beta, P)\|_* = 0.$$

By Lemma 5.0.44 for any $\varepsilon > 0$ there exists a full-group element ϕ' with $m(\alpha, \phi') < \varepsilon$ such that $m(\alpha, \phi') < \frac{\varepsilon}{2}$ and

$$h(T^{\alpha\phi'}, P') < h(T^\alpha) - e - \eta.$$

If, in fact, this ϕ' were measurable with respect to $\bigvee_{g \in G} T_{g^{-1}}^\alpha(P')$ then ϕ' would have a "version" on (U^β, P) , which would complete the proof. In general, of course, this need not hold. What we need of course is to "copy" ϕ' into this process.

Without loss of generality, we may assume that ϕ' is bounded. Select $\delta_1 > 0$, $\delta_1 = \delta_1(\alpha, \phi', \varepsilon, \eta)$, from Axiom 3 and from Theorem 3.0.24.

Thus if

$$\|(T^\alpha, g_{(\alpha, \phi')} \vee P'), (T^\alpha, g_{(\alpha, \phi)} \vee P')\|_* < \delta_1.$$

then both

$$m(\alpha, \phi) < m(\alpha, \phi') + \frac{\varepsilon}{2}$$

and

$$h(T^{\alpha\phi}, P') < h(T^{\alpha\phi'}) + \eta.$$

Apply Corollary 4.0.34 with “ Q ” = P' and “ ϕ ” = ϕ' to select Γ , measurable with respect to $\mathcal{H} = \bigvee_{g \in G} T_{g^{-1}}^{\alpha}(P')$ such that

$$\|(T^{\alpha, g(\alpha, \phi')} \vee P'), (T^{\alpha, g(\alpha, \phi)} \vee P')\|_* < \delta_1.$$

Hence

$$h(T^{\alpha\phi}, P') < h(T^{\alpha}) - e,$$

which completes the result. \square

Next, we see that if the entropy of a process can be perturbed down by arbitrarily small (in m) elements of the full-group, then the m -entropy of the action is in fact less than the standard entropy. This provides a converse to Lemma 5.0.45

LEMMA 5.0.46. *Let $e > 0$. Suppose, for some partition P and arrangement α , there exists $\{\phi_i\} \subseteq \Gamma$ with $m(\alpha, \phi_i) \rightarrow 0$ and*

$$h(T^{\alpha\phi_i}, P) < h(T^{\alpha}, P) - e$$

for all i . Then $h_m(T^{\alpha}) \leq h(T^{\alpha}) - e$.

PROOF. To begin, as the assumed inequality above is strict, it still holds for some $e' > e$. By Theorem 3.0.24 conditional entropy is upper semi-continuous with respect to the distribution pseudometric. Thus for any finite partition Q , we have that

$$\limsup_{i \rightarrow \infty} h(T^{\alpha\phi_i}, P \vee Q) \leq \limsup_{i \rightarrow \infty} (h(T^{\alpha\phi_i}, P) + h(T^{\alpha}, Q|P)).$$

Thus, by hypothesis, for any finite partition Q ,

$$\limsup_{i \rightarrow \infty} h(T^{\alpha\phi_i}, P \vee Q) < h(T^{\alpha}, P \vee Q) - e',$$

a strict inequality. In particular this tells us that for any finite partition Q there is a partition $R = P \vee Q$ refining Q and elements of the full group ϕ_i with $m(\alpha, \phi_i) \rightarrow 0$ and

$$\limsup_{i \rightarrow \infty} h(T^{\alpha\phi_i}, R) < h(T^{\alpha}, R) - e'.$$

Let P_j be a countable list of finite partitions dense in L^1 in the space of all partitions taking values in \mathbb{N} . We will define a sequence

of sets. Let \mathcal{O}_j consist of those $\beta \in E_m(\alpha)$ for which there is some partition R refining P_j and

$$h(T^\beta, R) < h(T^\alpha, R) - e'.$$

Upper semi-continuity of entropy tells us this is an open set. Using $\beta = \alpha\phi_i$ for i large enough we see that \mathcal{O}_j has α as an accumulation point. Consider those ϕ in the full-group that have finite order. By Theorem 4.0.27 such $\alpha\phi$ with ϕ of finite order are m -dense in $E_m(\alpha)$. We show that any such $\alpha\phi$ is an accumulation point of \mathcal{O}_j showing it to be dense. Fix such a ϕ and notice that any partition now is refined by a finite ϕ -invariant partition (by invariant here we mean one whose elements are permuted by ϕ .) As $(T^{\alpha\phi}, R\phi)$ is identical in distribution to (T^α, R) for such a ϕ -invariant partition we must have

$$h(T^{\alpha\phi}, R) = h(T^\alpha, R).$$

As $T^{\alpha\phi}$ is conjugate to T^α for any such R there must be a further refinement R' , which we can once more assume is ϕ -invariant, and full-group elements ψ_i with $m(\alpha\phi, \psi_i) \rightarrow 0$ and

$$\limsup_{i \rightarrow \infty} h(T^{\alpha\phi\psi_i}, R') < h(T^{\alpha\phi}, R') - e' = h(T^\alpha, R') - e'.$$

This implies that once i is large enough $\alpha\phi\psi_i \in \mathcal{O}_j$ and $\alpha\phi$ is an accumulation point of \mathcal{O}_j and these open sets are dense.

We know $E_m(\alpha)$ is Polish in the m -topology and hence these sets have nontrivial intersection in $E_m(\alpha)$. For any β in this intersection of course, and for all partitions P , there will be a refinement R of P with

$$h(T^\beta, R) < h(T^\alpha, R) - e'.$$

It follows easily that

$$h(T^\beta) \leq h(T^\alpha) - e' < h(T^\alpha) - e.$$

□

LEMMA 5.0.47. *Let α be any arrangement. Suppose there exists $e > 0$ such that $h_m(T^\alpha) < h(T^\alpha) - e$. Let β be any arrangement on any free and ergodic orbit relation for which $h(T^\beta) \geq h(T^\alpha)$. Then $h_m(T^\beta) < h(T^\beta) - e$.*

PROOF. Apply the Ornstein-Weiss Sinai Theorem [28] to obtain a partition P such that (T^β, P) is i.i.d. with $h(T^\beta, P) = h(T^\alpha)$. By

lemma 5.0.45, there exists a sequence $\{\phi_i\} \subseteq \Gamma$ such that $m(\beta, \phi_i) \rightarrow 0$ and, for every i ,

$$h(T^{\beta\phi_i}, P) < h(T^\beta, P) - e.$$

By lemma 5.0.46, $h_m(T^\beta) < h(T^\beta) - e$. □

The next result gives some insight into the relation between a size m , rearrangements and entropy. This is precisely where we use the copying lemma Theorem 4.0.38 through Corollary 4.0.39.

THEOREM 5.0.48. *Suppose T^α is an ergodic G -action, P is a finite partition, $e > 0$. Suppose $\{\phi_i\} \subseteq FG(\mathcal{O})$ with*

- (i) $m(\alpha, \phi_i) \xrightarrow{i} 0$, and
- (ii) $h(T^\alpha, P|P \circ \phi_i^{-1}) > e$.

Then there exists $\{\psi_i\} \subseteq FG(\mathcal{O})$ such that

- (iii) $m(\alpha, \psi_i) \xrightarrow{i} 0$, and
- (iv) $h(T^\alpha, P \circ \psi_i^{-1}) = h(T^{\alpha\psi_i}, P) < h(T^\alpha, P) - e$.

PROOF. Given $\varepsilon > 0$, we will construct ψ with

$$m(\alpha, \psi) < \varepsilon \text{ and } h(T^{\alpha\psi}, P) < h(T^\alpha, P) - e.$$

Choose ϕ_i with $m(\alpha, \phi_i) < \varepsilon$. Axiom 3 tells us that there is a $\delta > 0$ so that if

$$\|(\alpha, \phi_i), (\alpha, \psi)\|_* < \delta$$

then we will still have $m(\alpha, \psi) < \varepsilon$. Corollary 4.0.39 tells us we can construct a ψ with

$$h(T^{\alpha\psi}, P) = h(T^\alpha, P \circ \psi^{-1}) < h(T^\alpha, P) - e$$

and with

$$\|(\alpha, \phi_i), (\alpha, \psi)\|_* < \delta$$

giving the conclusion. □

The preceding result can be read as saying that if one can find full-group elements of arbitrarily small m -size that move a partition in a way visible to entropy, then one can actually lower entropy and as we are about to see, push entropy to zero. This result will be central to the rest of our proof of the dichotomy of entropy preserving and entropy free sizes.

LEMMA 5.0.49. *Let $\epsilon > 0$. Let T_1 and T_2 be two ergodic G -actions such that $h(T_1 \times T_2) > 0$ and $h_m(T_1 \times T_2) < h(T_1 \times T_2) - \epsilon$. Then either $h_m(T_1) < h(T_1) - \frac{\epsilon}{2}$ or $h_m(T_2) < h(T_2) - \frac{\epsilon}{2}$.*

PROOF. Let α_1 and α_2 be two G -arrangements with $T_1 = T^{\alpha_1}$ and $T_2 = T^{\alpha_2}$. Associated with $T^{\alpha_1} \times T^{\alpha_2}$ is a $G \times G$ -arrangement α such that $T^\alpha = T^{\alpha_1} \times T^{\alpha_2}$.

Thus, according to the hypotheses, we have a $G \times G$ -arrangement α such that $h(T^\alpha) > 0$ and $h_m(T^\alpha) < h(T^\alpha) - \epsilon$.

Apply lemma 5.0.45 to see that for any i.i.d. process (U^β, P) , if $h(U^\beta, P) = h(T^\alpha)$ then for any $\epsilon > 0$, there exists ϕ in the full-group of U^β with $m(\beta, \phi) < \epsilon$ and $h(U^{\beta\phi}, P) < h(U^\beta, P) - \epsilon$.

For $i = 1, 2$, let (U^{β_i}, P_i) be i.i.d. processes with $h(U^{\beta_i}, P_i) = h(T^{\alpha_i})$. Let $P = P_1 \times P_2$, β be so that $U^\beta = U^{\beta_1} \times U^{\beta_2}$. Then for the ϕ found above $m(\beta, \phi) < \epsilon$ and $h(U^{\beta\phi}, P_1 \times P_2) < h(U^\beta, P_1 \times P_2) - \epsilon$.

Note that, in general, $U^{\beta\phi}$ is not a direct product. The above calculation together with basic entropy facts, imply that (writing 2 to mean the trivial algebra) one of the following must hold; either

1. $h(U^\beta, P_1 \times 2 | (P_1 \times 2) \circ \phi^{-1}) > \frac{\epsilon}{2}$, or
2. $h(U^\beta, 2 \times P_2 | (2 \times P_2) \circ \phi^{-1}) > \frac{\epsilon}{2}$.

Without loss of generality, suppose (1) holds. By Theorem 5.0.48, given any $\bar{\epsilon} > 0$, we could have selected ϵ in such a way that there exists ψ in the full-group of U^β with $m(\beta, \psi) < \bar{\epsilon}$ and

$$h(U^\beta, (P_1 \times 2) \circ \psi^{-1}) = h(U^{\beta\psi}, P_1 \times 2) < h(U^\beta, P_1 \times 2) - \frac{\epsilon}{2}.$$

Apply Corollary 4.0.34, with “ Q ” = $P_1 \times 2$ and

$$\mathcal{H} = \vee_{g \in G} U_{g^{-1}}^\beta(P_1 \times 2) = \vee_{g \in G} U_{1^{-1}g^{-1}}^{\beta_1}(P_1).$$

Thus for any $\delta > 0$ there exists $\tilde{\psi}$ in the full-group of $U_1^{\beta_1}$

$$\|(U^\beta, g_{(\beta, \psi)} \vee (P_1 \times 2)), (U_1^{\beta_1}, g_{(\beta, \tilde{\psi})} \vee (P_1))\|_* < \delta.$$

By Axiom 3, we may select δ so small that also $m(\beta_1, \tilde{\psi}) < \epsilon$. Furthermore, by the upper semi-continuity of entropy, for any η we may select δ so small that $h(U_1^{\beta_1 \tilde{\psi}}, P_1) < h(U^{\beta\psi}, P_1 \times 2) + \eta$. Thus, in fact, the δ may be selected so small that we still have

$$h(U_1^{\beta_1 \tilde{\psi}}, P_1) < h(U^\beta, P_1 \times 2) - \frac{\epsilon}{2} = h(U_1^{\beta_1}, P_1) - \frac{\epsilon}{2}.$$

In particular, for any $\varepsilon > 0$, we have produced an element $\tilde{\psi}$ in the full-group of $U_1^{\beta_1}$ such that $m(\beta_1, \tilde{\psi}) < \varepsilon$ and $h(U_1^{\beta_1 \tilde{\psi}}, P_1) < h(U_1^{\beta_1}, P_1) - \frac{\varepsilon}{2}$.

Now, lemma 5.0.46 implies that $h_m(T_1) < h(T_1) - \frac{\varepsilon}{2}$ as T_1 must have a factor conjugate to the Bernoulli shift $U_1^{\beta_1}$. \square

We are now ready to complete the proof that there are two kinds of sizes - those which are entropy preserving, in which case the m -entropy is simply entropy, and those which are entropy free, in which case the m -entropy is always zero.

THEOREM 5.0.50. *Let $e > 0$. Suppose there exists some G -arrangement α such that $h(T^\alpha) > 0$ and $h_m(T^\alpha) < h(T^\alpha) - e$. Then for any free and ergodic G -action S^β , $h_m(S^\beta) = 0$.*

PROOF. Suppose for some arrangement β , $h_m(T^\beta) = B > 0$. By theorem 5.0.43, we may select $\gamma \in E_m(\beta)$ such that $h(S^\gamma) = h_m(S^\beta) = B > 0$. In particular, $h(S^\gamma) = h_m(S^\gamma) = B > 0$.

Let (U, P) be any Bernoulli process with $h(U) = h(U, P) = B$. Select k so large (but finite!) that the k -fold product $(U \times \cdots \times U, P \times \cdots \times P)$ has entropy $h(U \times \cdots \times U, P \times \cdots \times P) \geq h(T^\alpha)$. By lemma 5.0.47, $h_m(U \times \cdots \times U) < h(U \times \cdots \times U) - e$. Applying lemma 5.0.49 inductively, $k - 1$ times, we see that $h_m(U) < h(U) - \frac{e}{2^{k-1}}$. (In other words, we can lower the entropy on any Bernoulli process.) Since $h(S^\gamma) = h(U)$, lemma 5.0.47 now implies that $h_m(S^\gamma) < h(S^\gamma) - \frac{e}{2^{k-1}}$. But this is a contradiction, since $h_m(S^\gamma) = h(S^\gamma)$. Hence, we must have that $h_m(S^\beta) = 0$ universally. \square

This completes our technical work on m -entropy.

DEFINITION 5.0.51. *Let m be a given size.*

1. *If $h_m(T^\alpha) = h(T^\alpha)$, for all arrangements α , then m is called an **entropy-preserving size**.*
2. *If $h_m(T^\alpha) = 0$ for all arrangements α , then m is called an **entropy-free size**.*

Theorem 5.0.50 has shown us that there are only these two possibilities, a size is either entropy-preserving or it is entropy-free.