Journal of Logic & Analysis 15:2 (2023) 1–68
ISSN 1759-9008

Covering entropy for types in tracial W*–algebras

DAVID JEKEL

Abstract: We study embeddings of tracial W*–algebras into a ultraproduct of matrix algebras through an amalgamation of free probabilistic and model-theoretic techniques. Jung implicitly and Hayes explicitly defined 1–*bounded entropy* h through the asymptotic covering numbers of Voiculescu's microstate spaces, that is, spaces of matrix tuples $(X_1^{(N)}, X_2^{(N)}, \dots)$ having approximately the same $*$–moments as the generators $(X_1, X_2, \dots)$ of a given tracial W*–algebra. We study the analogous covering entropy for microstate spaces defined through formulas that use suprema and infima, not only $*$–algebra operations and the trace–formulas which arise in the model theory of tracial W*–algebras initiated by Farah, Hart, and Sherman. By relating the new theory with the original 1–bounded entropy, we show that if $\mathcal{M}$ is a separable tracial W*–algebra with $h(\mathcal{N} : \mathcal{M}) \geq 0$, then there exists an embedding of $\mathcal{M}$ into a matrix ultraproduct $\mathcal{Q} = \prod_{n \to \mathcal{U}} M_n(\mathbb{C})$ such that $h(\mathcal{N} : \mathcal{Q})$ is arbitrarily close to $h(\mathcal{N} : \mathcal{M})$. We deduce that if all embeddings of $\mathcal{M}$ into $\mathcal{Q}$ are automorphically equivalent, then $\mathcal{M}$ is strongly 1–bounded and in fact has $h(\mathcal{M}) \leq 0$.

2020 Mathematics Subject Classification 03C66 (primary); 46L51, 46L54, 94A17 (secondary)

Keywords: model theory, von Neumann algebra, free entropy, random matrix, strongly 1–bounded, ultraproduct

1 Introduction

1.1 Overview

The study of W*–algebras or von Neumann algebras is a deep and challenging subject with many connections to fields as diverse as ergodic theory, geometric group theory, random matrix theory, quantum information, and model theory. Our present goal is to bring together two of these facets—the model theory of tracial W*–algebras studied in the work of Farah, Hart, and Sherman [7, 8, 9] and Voiculescu's free entropy theory which, roughly speaking, quantifies the amount of matrix approximations for the generators of W*–algebra (see eg Voiculescu [32, 33], Ge [12], Jung [25] and

Published: May 2023 https://doi.org/10.4115/jla.2023.15.2

Hayes [16]). On the free entropy side, we will work in the framework of Hayes' 1–bounded entropy h [16] which arose out of the work of Jung [25]; for history and motivation, refer to Hayes, Jekel, Nelson, and Sinclair [18, § 2]. The sibling paper of this one [23] develops the analog of Voiculescu's free microstate entropy for the setting of model-theoretic types.

We adapt the framework of 1–bounded entropy from Jung [25] and Hayes [16] to capture data about the generators' model-theoretic *type* and not only their *non-commutative law*. The non-commutative law of a tuple $(X_j)_{j\in\mathbb{N}}$ from $\mathcal{M} = (M, \tau)$ encodes the joint moments $\tau(p(X))$ for non-commutative $*$–polynomials p; laws thus describe tracial von Neumann algebras with chosen generators up to generator-preserving isomorphism. The *type* includes the values of more complicated formulas that involve not only the addition, adjoint, product, and trace operations, but also taking suprema and infima in auxiliary variables over an operator norm ball in $\mathcal{M}$. For instance, the type would include the value of the formula

$$\sup_{Z\in D_1^{\mathcal{M}}} \inf_{Y\in \overline{B}_{\mathcal{M}}(0,1)} [\tau(X_1 Y X_2 Z)^2 + \tau(X_2^2 Y^5)]$$

where $D_1^{\mathcal{M}}$ denotes the closed unit ball with respect to operator norm. The entropy $\mathrm{Ent}^{\mathcal{U}}(\mu)$ of a type μ is defined in terms of the exponential growth rates of the covering numbers of microstate spaces (spaces of matrix tuples with approximately the same type as our chosen generators, as in Voiculescu's work), just like Jung and Hayes' 1–bounded entropy except with types instead of laws. However, we prefer the term "covering entropy" rather than "1–bounded entropy" as a more intrinsic description of the definition. The superscript $\mathcal{U}$ denotes the fact that we take limits with respect to a fixed non-principal ultrafilter $\mathcal{U}$ on $\mathbb{N}$.

Just as in the original definition of the 1–bounded entropy, a key property of the covering entropy $\mathrm{Ent}^{\mathcal{U}}(\mu)$ is that it is invariant under change of coordinates (see § 4.3). More precisely, if $\mathbf{X}$ and $\mathbf{Y}$ are tuples from $\mathcal{M}$ with $\mathrm{W}^*(\mathbf{X}) = \mathrm{W}^*(\mathbf{Y})$, then their types $\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$ and $\mathrm{tp}^{\mathcal{M}}(\mathbf{Y})$ have the same covering entropy (Corollary 4.10). This allows us to define the entropy $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$ of a separable tracial W^*–algebra $\mathcal{N} \subseteq \mathcal{M}$ as the entropy of the type of any generating set. As suggested in Hayes, Jekel, Nelson, and Sinclair [18], we streamline the proof of this invariance property using the result that every tuple $\mathbf{Y}$ from $\mathrm{W}^*(\mathbf{X})$ can be expressed as $\mathbf{f}(\mathbf{X})$ for some quantifier-free definable function (see Jekel [22, § 13]). More generally, we can extend the definition of $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$ to the case where $\mathcal{N} \subseteq \mathcal{M}$ is not separable by setting it to be the supremum of $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N}_0 : \mathcal{M})$ over separable W^*–subalgebras $\mathcal{N}_0 \subseteq \mathcal{N}$ or equivalently, the supremum of $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$ over all tuples $\mathbf{X} \in L^\infty(\mathcal{N})^{\mathbb{N}}$ (see Definition 4.11).

The covering entropy $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$ can be viewed intuitively as a measurement of the amount of tracial W*–embeddings of $\mathcal{N}$ into the matrix ultraproduct $\mathcal{Q} = \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$ that extend to elementary embeddings of $\mathcal{M}$ (compare §4.4). This is the analog of the idea that the 1–bounded entropy $h(\mathcal{N} : \mathcal{M})$ of $\mathcal{N}$ in the presence of $\mathcal{M}$ quantifies the amount of W*–embeddings of $\mathcal{N}$ into $\mathcal{Q}$ that extend to any embedding of $\mathcal{M}$. Thus, our work is motivated in part by the study of embeddings into ultraproducts, which is one theme of recent work on von Neumann algebras; for instance, see Popa [28], Goldbring [14], Ioana and Spaas [20], Atkinson and Kunnawalkam Elayavalli [2], Atkinson, Goldbring, and Kunnawalkam Elayavalli [1], Gao [11].

We make a precise connection between $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$ and 1–bounded entropy as follows. There is a canonical projection π_{qf} from the space of types to the space of non-commutative laws, since a non-commutative law describes the evaluation of quantifier-free formulas (rather than all logical formulas) in a tuple $\mathbf{X}$. Given a non-commutative law (or quantifier-free type) μ, the 1–bounded entropy $h^{\mathcal{U}}(\mu)$ can be expressed through the following variational principle (Corollary 5.4):

$$h^{\mathcal{U}}(\mu) = \sup_{\nu\in\pi_{\mathrm{qf}}^{-1}(\mu)} \mathrm{Ent}^{\mathcal{U}}(\nu) \tag{1–1}$$

Thus, the 1–bounded entropy is the quantifier-free version of the entropy for types.

In a similar way, the 1–bounded entropy of $\mathcal{N}$ in the presence of $\mathcal{M}$ is the version using existential types. Entropy in the presence is described using microstates for a tuple $\mathbf{X}$ in $\mathcal{N}$ such that there exist compatible microstates for a tuple $\mathbf{Y}$ that generates $\mathcal{M}$. In the model-theoretic framework, the existence of such microstates for $\mathcal{M}$ is described through the evaluation of existential formulas in the original generators and their microstates (see §5.4). Similar to the quantifier-free setting, there is a projection $\pi_{\exists}$ from the space of types into the space of existential types, and a similar variational principle expressing the covering entropy of an existential type μ as the supremum of $\mathrm{Ent}^{\mathcal{U}}(\nu)$ over full types $\nu \in \pi_{\exists}^{-1}(\mu)$ (Lemma 5.13).

Altogether these ingredients allow us to prove the following result about ultraproduct embeddings, which is restated and proved in Theorem 5.24:

Theorem 1.1 *Let $c \in \mathbb{R}$. Let $\mathcal{N} \subseteq \mathcal{M}$ be an inclusion of separable tracial* W*–*algebras $h^{\mathcal{U}}(\mathcal{N} : \mathcal{M}) > c$. Then there exists an embedding ι of $\mathcal{M}$ into the matrix ultraproduct $\mathcal{Q} = \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$ such that $\mathrm{Ent}^{\mathcal{U}}(\iota(\mathcal{N}) : \mathcal{Q}) > c$, hence also $h^{\mathcal{U}}(\iota(\mathcal{N}) : \mathcal{Q}) > c$.*

The hypotheses of the theorem hold for instance when $\mathcal{N} = \mathcal{M}$ is a nontrivial free product by Voiculescu [33, Proposition 6.8] and Jung [25, Corollary 3.5] and

Hayes [16, Proposition A.16] since $h^{\mathcal{U}}(\mathcal{M} : \mathcal{M}) = \infty$. They also hold for the von Neumann algebras of groups with non-approximately-inner cocycles by Shlyakhtenko [30, Theorem 3] and Jung [25, Corollary 3.5] and Hayes [16, Proposition A.16].

In particular, since there exists $\mathcal{M}$ with $h^{\mathcal{U}}(\mathcal{M} : \mathcal{M}) = \infty$, the theorem implies that there exist types in $\mathcal{Q}$ with arbitrarily large covering entropy, and therefore, $h^{\mathcal{U}}(\mathcal{Q} : \mathcal{Q}) = \infty$. Similarly, the entropy $\mathrm{Ent}^{\mathcal{U}}(\mathcal{Q} : \mathcal{Q})$ given by Definition 4.11 is infinite.

Corollary 1.2 *Let $\mathcal{U}$ be a free ultrafilter on $\mathbb{N}$ and let $\mathcal{Q} = \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$. Then $h^{\mathcal{U}}(\mathcal{Q} : \mathcal{Q}) = \infty$ and $\mathrm{Ent}^{\mathcal{U}}(\mathcal{Q} : \mathcal{Q}) = \infty$. Hence, $\mathcal{Q}$ is not strongly 1–bounded.*

The following corollary of Theorem 1.1 was communicated to me by Ben Hayes.

Corollary 1.3 *Let $\mathcal{N} \subseteq \mathcal{M}$ be an inclusion of separable tracial W^*–algebras $h(\mathcal{M}) > 0$ such that $\mathcal{N}$ is a II_1 factor (it has trivial center). Then there exists a free ultrafilter $\mathcal{V}$ and an embedding $\iota : \mathcal{M} \to \prod_{n\to\mathcal{V}} M_n(\mathbb{C})$ such that $\mathcal{N}' \cap \prod_{n\to\mathcal{V}} M_n(\mathbb{C}) = \mathbb{C}$.*

Proof The 1–bounded entropy $h(\mathcal{M})$ is the supremum of $h^{\mathcal{V}}(\mathcal{M})$ over free ultrafilters $\mathcal{V}$. Hence, there exists some free ultrafilter $\mathcal{U}$ such that $h^{\mathcal{V}}(\mathcal{M}) > 0$ and by Theorem 1.1 there is an embedding $\iota_0 : \mathcal{M} \to \mathcal{Q} = \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$ with $h^{\mathcal{U}}(\iota_0(\mathcal{M}) : \mathcal{Q}) > 0$.

A general fact about 1–bounded entropy is that if $\mathcal{A} \subseteq \mathcal{B}$ and $\mathcal{A}' \cap \mathcal{B}$ is diffuse, then $h(\mathcal{A} : \mathcal{B}) \leq 0$. Indeed, if $\mathcal{A}' \cap \mathcal{B}$ is diffuse, it contains a diffuse amenable subalgebra $\mathcal{C}$. Let

$$\mathcal{N} = \mathrm{W}^*(u \in \mathcal{Q} \text{ unitary: } u\mathcal{C}u^* \cap \mathcal{C} \text{ is diffuse})$$

(this is known as the step 1 wq-normalizer of $\mathcal{C}$ and was introduced in Gatalan and Popa [10]). Note that $\mathcal{A} \subseteq \mathcal{N}$. Hence, by Hayes [16, Property 1, page 10],

$$h^{\mathcal{U}}(\mathcal{A} : \mathcal{B}) \leq h^{\mathcal{U}}(\mathcal{N} : \mathcal{B}).$$

By Hayes [16, Theorem 2.8 and Proposition 3.2],

$$h^{\mathcal{U}}(\mathcal{N} : \mathcal{B}) = h^{\mathcal{U}}(\mathcal{C} : \mathcal{B}).$$

Then using Hayes [16, Property 1, page 10] again,

$$h^{\mathcal{U}}(\mathcal{C} : \mathcal{B}) \leq h^{\mathcal{U}}(\mathcal{C} : \mathcal{C}) = h^{\mathcal{U}}(\mathcal{C})$$

which is zero since $\mathcal{C}$ is amenable. Hence, $h(\mathcal{A} : \mathcal{B}) \leq 0$.

By the contrapositive, since in our case $h(\iota(\mathcal{M}) : \mathcal{Q}) \geq h^{\mathcal{U}}(\iota_0(\mathcal{M}) : \mathcal{Q}) > 0$, then $\iota_0(\mathcal{M})' \cap \mathcal{Q}$ is not diffuse. Therefore, it contains a minimal projection p. Let $p\mathcal{Q}p$ be

the compression of $\mathcal{Q}$ by p equipped with the trace $\tau_{p\mathcal{Q}p}(x) = \tau_{\mathcal{Q}}(pxp)/\tau(p)$, and let $\iota : \mathcal{M} \to p\mathcal{Q}p$ be the map $\iota(x) = p\iota_0(x)p = \iota_0(x)$. Since p commutes with $\iota_0(\mathcal{M})$, it follows that ι_0 is a $*$–homomorphism, and since $\mathcal{M}$ is a II_1 factor and hence has a unique trace, the map ι must be trace-preserving. Because p was a minimal projection in $\iota_0(\mathcal{M})' \cap \mathcal{Q}$, we know $\iota(\mathcal{M})' \cap p\mathcal{Q}p$ has no nontrivial projections and hence is $\mathbb{C}$.

Finally, note that $p\mathcal{Q}p$ is a matrix ultraproduct $\prod_{n\to\mathcal{V}} M_n(\mathbb{C})$ for some ultrafilter $\mathcal{V}$. Indeed, by stability of projections there exist projections p_n in $M_n(\mathbb{C})$ such that p is the equivalence class of $(p_n)_{n\in\mathbb{N}}$ in $\mathcal{Q}$. Let $k(n)$ be the rank of p_n. One can check that $p\mathcal{Q}p = \prod_{n\to\mathcal{U}} p_n M_n(\mathbb{C}) p_n \cong \prod_{n\to\mathcal{U}} M_{k(n)}(\mathbb{C})$, which is simply a matrix ultraproduct for a different ultrafilter $\mathcal{V}$. □

As shown in Jekel [23, Theorem 1.2], the analogous result holds for free entropy rather than 1–bounded entropy *without* having to change the ultrafilter $\mathcal{U}$ to the ultrafilter $\mathcal{V}$.

1.2 Embeddings into Ultraproducts

Our results relate to recent work and questions about embeddings into ultraproducts. Jung [24] used the study of microstates to show that a separable tracial W*–algebra $\mathcal{A}$ is amenable if and only if all embeddings of $\mathcal{A}$ into $\mathcal{R}^{\mathcal{U}}$ are unitarily conjugate. Atkinson and Kunnawalkam Elayavalli [2] strengthened this result by showing that $\mathcal{A}$ is amenable if and only if all embeddings of $\mathcal{A}$ into $\mathcal{R}^{\mathcal{U}}$ are ucp-conjugate (meaning they are conjugate by an automorphism of $\mathcal{R}^{\mathcal{U}}$ that lifts to a sequence of unital completely positive maps $\mathcal{R} \to \mathcal{R}$). Atkinson, Goldbring, and Kunnawalkam Elayavalli [1] later showed that if a separable II_1 factor $\mathcal{M}$ is Connes-embeddable and all embeddings of $\mathcal{M}$ into $\mathcal{M}^{\mathcal{U}}$ are automorphically conjugate, then $\mathcal{M} \cong \mathcal{R}$.

One can ask similar questions for embeddings into the ultraproduct $\mathcal{Q} = \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$ for some fixed free ultrafilter $\mathcal{U}$. Atkinson and Kunnawalkam Elayavalli [2] showed that if $\mathcal{A}$ is a separable Connes-embeddable tracial W*–algebra and the space of unitary orbits of embeddings $\mathcal{A} \to \mathcal{Q}$ is separable in a certain metric, then $\mathcal{A}$ must be amenable. In particular, if all embeddings $\mathcal{A} \to \mathcal{Q}$ are unitarily conjugate, then $\mathcal{A}$ is amenable. It is an open question whether this result still holds when "unitarily conjugate" is replaced by "automorphically conjugate." However, Theorem 1.1 implies the following result, which was pointed out to me by Srivatsav Kunnawalkam Elayavalli:

Corollary 1.4 *Let $\mathcal{A}$ be a tracial* W*–*algebra. Suppose that any two embeddings $\mathcal{A} \to \mathcal{Q} = \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$ are conjugate by an automorphism of $\mathcal{Q}$. Then $h^{\mathcal{U}}(\mathcal{A}) \leq 0$.*

Proof We proceed by contradiction. Suppose that $h^{\mathcal{U}}(\mathcal{A}) > 0$. By Theorem 1.1, there exists an embedding $\alpha : \mathcal{A} \to \mathcal{Q}$ with $h^{\mathcal{U}}(\alpha(\mathcal{A}) : \mathcal{Q}) > 0$. Moreover, since $\mathcal{A}$ is Connes-embeddable, so is $\mathcal{R}\overline{\otimes}\mathcal{A}$, so there exists some embedding $\beta : \mathcal{R}\overline{\otimes}\mathcal{A} \to \mathcal{Q}$. In particular, $\beta(\mathcal{R}) \subseteq \beta(\mathcal{A})' \cap \mathcal{Q}$. If we assume for contradiction that α and $\beta|_{\mathcal{A}}$ are conjugate by an automorphism, then $\alpha(\mathcal{A})' \cap \mathcal{Q}$ also contains a copy of $\mathcal{R}$, so in particular, $\alpha(\mathcal{A})' \cap \mathcal{Q}$ is diffuse. As pointed out in the proof of Corollary 1.3, this implies that $h(\alpha(\mathcal{A}) : \mathcal{Q}) = 0$, which contradicts our choice of α. □

Intuitively, the corollary says that if the space of embeddings modulo automorphic conjugacy is trivial, then the space of embeddings modulo unitary conjugacy is not too large, since $h^{\mathcal{U}}(\mathcal{A})$ quantifies the "amount" of embeddings $\mathcal{A} \to \mathcal{Q}$ up to unitary conjugacy. The conclusion that $h^{\mathcal{U}}(\mathcal{A}) = 0$ is a weakening of amenability since by Jung's theorem [24] amenability is equivalent to the space of embeddings modulo unitary conjugacy being trivial.

We remark that the free entropy techniques used here to study embeddings into $\mathcal{Q}$ cannot be directly applied to study embeddings into $\mathcal{R}^{\mathcal{U}}$. For instance, Theorem 1.1 does not make sense with $\mathcal{Q}$ replaced by $\mathcal{R}^{\mathcal{U}}$. Indeed, $\mathcal{R}^{\mathcal{U}}$ has property Gamma by Farah Hart and Sherman [9, § 3.2.2], and every tracial W*–algebra with property Gamma has 1–bounded entropy zero (this is a special case of Hayes [16, Corollary 4.6] and it is shown explicitly in Hayes, Jekel, Nelson, and Sinclair [18, § 1.2, Example 4]). Thus, $h(\mathcal{R}^{\mathcal{U}}) = 0$ and therefore, for any subalgebra $\mathcal{M}$ of $\mathcal{R}^{\mathcal{U}}$, we also have $h^{\mathcal{U}}(\mathcal{M} : \mathcal{R}^{\mathcal{U}}) = 0$ by Hayes [16, § 2, Property 1]. Hence, Theorem 1.1 would not hold with $\mathcal{R}^{\mathcal{U}}$ instead of $\mathcal{Q}$. By contrast, many other operator-theoretic and model-theoretic techniques are more easily applied to $\mathcal{R}^{\mathcal{U}}$ than to $\mathcal{Q}$ since $\mathcal{R}^{\mathcal{U}}$ is an ultrapower; see for instance the following works of Goldbring [14, 15, 13].

1.3 Outline

In large part, our goal is to establish communication between the free probabilistic and model theoretic subgroups of operator algebras, and to show that many of the notions in free probability (such as non-commutative laws, microstates spaces in the presence, and relative microstate spaces) arise naturally from the model-theoretic framework. Therefore, we strive to make the exposition largely self-contained and use model-theoretic language throughout.

We start out by explaining the model-theoretic framework for operator algebras in § 2. In particular, we give a more detailed explanation than current literature of the languages

and structures for multiple sorts and multiple domains of quantification for each. Next, in §3, we give a self-contained development of definable predicates and functions of infinite tuples, including the result that every element of a tracial W*–algebra can be realized by applying a quantifier-free definable function to the generators which was observed in the author's thesis [22, 18].

§4 develops the framework of covering entropy for types. We show the invariance of entropy under change of coordinates in §4.3, describe the relationship with ultraproduct embeddings in §4.4, and finally show that adding variables in the (model-theoretic) algebraic closure of given tuple $\mathbf{X}$ does not change its entropy in §4.5.

In §5, we describe the quantifier-free and existential versions of entropy, showing that they agree with the 1–bounded entropy of Hayes. We conclude the proof Theorem 1.1 there.

In §6, we describe a generalization to conditional (or "relative") entropy, which focuses on quantifying the embeddings $\mathcal{N} \to \mathcal{Q}$ which restrict to a fixed embedding $\iota : \mathcal{A} \to \mathcal{Q}$ on a given W*–subalgebra $\mathcal{A}$. The existential version of the conditional covering entropy was studied by Hayes explicitly for $\mathcal{A}$ diffuse abelian and implicitly for $\mathcal{A}$ diffuse amenable [16], in which case it agrees with the unconditional version. However, the conditional covering entropy (for full, quantifier-free, or existential types) makes sense for any diffuse $\mathcal{A}$ with a specified embedding $\alpha : \mathcal{A} \to \mathcal{Q}$ (though, as far as we know, it may depend on the embedding α). Moreover, conditional entropy is natural from the model-theoretic perspective, since it arises from replacing formulas in the original language with formulas that have coefficients from the subalgebra $\mathcal{A}$.

1.4 Acknowledgements

This was work was partially funded by the NSF postdoc grant DMS-2002826. I thank the organizers and hosts of the 2017 workshop on the model theory of operator algebras at University of California, Irvine, and of the 2018 long program Quantitative Linear Algebra at the Institute of Pure and Applied Mathematics at UCLA for conferences that greatly contributed to my knowledge of operator algebras, entropy, and model theory. Special thanks to Ben Hayes for pointing out Corollary 1.3, Srivatsav Kunnawalkam Elayavalli for pointing out Corollary 1.4, and Jennifer Pi for reading the paper in detail and finding many typos. Thanks to Isaac Goldbring, Ben Hayes, and Adrian Ioana for their comments on the paper and advice on references. Finally, thanks to the anonymous referees for numerous emendations to the paper.

2 Continuous model theory for tracial W*–algebras

This section sketches the setup of continuous model theory, or model theory for metric structures of Ben Yaacov, Berenstein, Henson, and Usvyatsov [3, 4] and its application to operator algebras by Farah, Hart, and Sherman [7, 8, 9]. We strive to present a self-contained exposition for two reasons: First, some readers may not be familiar with the model-theoretic terminology. Second, we are following the treatment in [8] which introduces "domains of quantification" to cut down on the number of "sorts," which means that some of the statements need to modified from their original form in [3].

2.1 Background on operator algebras

We start by giving some basic terminology and background on operator algebras. For further detail and history, we suggest consulting the following references: Kadison and Ringrose [26], Dixmier [6], Sakai [29], Takesaki [31], Blackadar [5], and Zhu [34].

C*–algebras:

(1) A (unital) algebra over $\mathbb{C}$ is a unital ring A with a unital inclusion map $\mathbb{C} \to A$.
(2) A (unital) $*$–algebra is an algebra A equipped with a conjugate linear involution $*$ such that $(ab)^* = b^*a^*$.
(3) A unital C*–algebra is a $*$–algebra A equipped with a complete norm $\|\cdot\|$ such that $\|ab\| \leq \|a\|\|b\|$ and $\|a^*a\| = \|a\|^2$ for $a, b \in A$.

A collection of fundamental results in C*–algebra theory establishes that C*–algebras can always be represented as algebras of operators on Hilbert spaces. If H is a Hilbert space, the algebra of bounded operators $B(H)$ is a C*–algebra. Conversely, every unital C*–algebra can be embedded into $B(H)$ by some unital and isometric $*$–homomorphism ρ. By isometric, we mean that $\|\rho(a)\| = \|a\|$, where $\|\rho(a)\|$ is the operator norm on $B(H)$ and $\|a\|$ is the given norm on the C*–algebra A.

W*–algebras: A *von Neumann algebra* is a $*$–subalgebra of $B(H)$ (for some Hilbert space H) that is closed in the *strong operator topology*, the topology of pointwise convergence as functions on H. A W*–*algebra* is a C*–algebra that admits a predual (that is, it is the dual of some Banach space). A deep result of Sakai showed that for a C*–algebra A is a W*–algebra if and only if it is isomorphic to a von Neumann algebra; moreover, the weak-$\star$ topology on a W*–algebra A is uniquely determined by its C*–algebra structure as shown by Sakai [29, Corollary 1.13.3].

Tracial W*–algebras: A *tracial* W*–*algebra* is a W*–algebra M together with a linear map $\tau : M \to \mathbb{C}$ satisfying:

- *positivity*: $\tau(x^*x) \geq 0$ for all $x \in M$
- *unitality*: $\tau(1) = 1$
- *traciality*: $\tau(xy) = \tau(yx)$ for $x, y \in A$
- *faithfulness*: $\tau(x^*x) = 0$ implies $x = 0$ for $x \in A$.
- *weak-$\star$ continuity*: $\tau : M \to \mathbb{C}$ is weak-$\star$ continuous.

We call τ a *faithful normal tracial state*.

The standard representation: Given a tracial W*–algebra (M, τ), we can form a Hilbert space $L^2(M, \tau)$ as the completion of M with respect to the inner product $\langle x, y \rangle_\tau = \tau(x^*y)$; if $x \in M$, then we denote the corresponding element of $L^2(M, \tau)$ by $\widehat{x}$. There is a unique unital $*$–homomorphism $\pi_\tau : M \to B(L^2(M, \tau))$ satisfying $\pi_\tau(x)\widehat{y} = \widehat{xy}$ for $x, y \in M$. Now π_τ is a $*$–homomorphism isometric with respect to the operator norm, and its image is a von Neumann algebra. The construction of $L^2(M, \tau)$ and π_τ is a special case of the *GNS (Gelfand–Naimark–Segal) construction* and π_τ is also known as the *standard representation* of (M, τ). Note the convergence of a net x_i to x in M with respect to the strong operator topology in $B(L^2(M, \tau))$ implies convergence of $\widehat{x_i} = \pi_\tau(x_i)\widehat{1}$ to $\widehat{x} = \pi_\tau(x)\widehat{1}$ in $L^2(M, \tau)$. (It turns out that the converse is true if $(x_i)_{i \in I}$ is bounded in operator norm, but we will not need to use this fact directly.)

$*$–polynomials and generators: Given an index set I, we denote by $\mathbb{C}\langle x_i, x_i^* : i \in I \rangle$ the free unital algebra (or non-commutative polynomial algebra) generated by indeterminates x_i and x_i^* for $i \in I$. We equip $\mathbb{C}\langle x_i, x_i^* : i \in I \rangle$ with the unique $*$–operation sending x_i to x_i^*, thus making it into a $*$–algebra. If A is a unital $*$–algebra and $(a_i)_{i \in I}$ a collection of elements, there is a unique unital $*$–homomorphism $\rho : \mathbb{C}\langle x_i, x_i^* : i \in I \rangle \to A$ mapping x_i to a_i for each $i \in I$. We refer to the elements of $\mathbb{C}\langle x_i, x_i^* : i \in I \rangle$ as *non-commutative $*$–polynomials*, and if $p \in \mathbb{C}\langle x_i, x_i^* : i \in I \rangle$ and $\rho : \mathbb{C}\langle x_i, x_i^* : i \in I \rangle \to A$ is as above, we denote $\rho(p)$ by $p(a_i : i \in I)$. Moreover, the image of ρ is the *$*$–algebra generated by $(a_i)_{i \in I}$*.

If A is a C*–algebra and $(a_i)_{i \in I}$ is a collection of elements of I, then the C*–*algebra generated by* $(a_i)_{i \in I}$ is the norm-closure of the $*$–algebra generated by $(a_i)_{i \in I}$. Similarly, if M is a von Neumann algebra and $(a_i)_{i \in I}$ is a collection of elements of M, then the *von Neumann subalgebra* or W*–*subalgebra generated by* $(a_i)_{i \in I}$ is the strong operator topology closure of the $*$–algebra generated by $(a_i)_{i \in I}$. In particular, we say that $(a_i)_{i \in I}$ *generates* M if the strong operator topology closure is all of M.

2.2 Languages and structures

Next, let us sketch the setup of continuous model theory, or model theory for metric structures from Ben Yaacov, Berenstein, Henson, and Usvyatsov [3, 4]. We will follow the treatment in Farah, Hart, and Sherman [8] which introduces "domains of quantification" to cut down on the number of "sorts" needed.

A *language* $\mathcal{L}$ consists of:

- A set $\mathcal{S}$ whose elements are called *sorts*.
- For each $S \in \mathcal{S}$, a privileged relation symbol d_S (which will represent a metric) and a set $\mathcal{D}_S$ whose elements are called *domains of quantification for* S.
- For each $S \in \mathcal{S}$ and $D, D' \in \mathcal{D}_S$ an assigned constant $C_{D,D'}$.
- A countably infinite set of *variable symbols* for each sort S. We denote the variables by $(x_i)_{i\in\mathbb{N}}$.
- A set of *function symbols*.
- For each function symbol f, an assigned tuple $(S_1, \ldots, S_n)$ of sorts called the *domain*, another sort S called the *codomain*. We call n the *arity of* f.
- For each function symbol f with domain $(S_1, \ldots, S_n)$ and codomain S, and for every $\mathbf{D} = (D_1, \ldots, D_n) \in \mathcal{D}_{S_1} \times \cdots \times \mathcal{D}_{S_n}$, there is an assigned $D_{f,\mathbf{D}} \in \mathcal{D}_S$ (representing a range bound), and assigned moduli of continuity $\omega_{f,\mathbf{D},1}, \ldots, \omega_{f,\mathcal{D},n}$. (Here "modulus of continuity" means a continuous increasing, zero-preserving function $[0,\infty) \to [0,\infty)$).
- A set of *relation symbols*.
- For each relation symbol R, an assigned domain $(S_1, \ldots, S_n)$ as in the case of function symbols.
- For each relation symbol R and for every $\mathbf{D} = (D_1, \ldots, D_n) \in \mathcal{D}_{S_1} \times \cdots \times \mathcal{D}_{S_n}$, an assigned bound $N_{R,\mathbf{D}} \in [0,\infty)$ and assigned moduli of continuity $\omega_{R,\mathbf{D},1}, \ldots, \omega_{R,\mathcal{D},n}$.

Given a language $\mathcal{L}$, an *$\mathcal{L}$–structure* $\mathcal{M}$ assigns an object to each symbol in $\mathcal{L}$, called the *interpretation* of that symbol, in the following manner:

- Each sort $S \in \mathcal{S}$ is assigned a metric space $S^{\mathcal{M}}$, and the symbol d_S is interpreted as the metric $d_S^{\mathcal{M}}$ on $S^{\mathcal{M}}$.
- Each domain of quantification $D \in \mathcal{D}_S$ is assigned a subset $D^{\mathcal{M}} \subseteq S^{\mathcal{M}}$, such that $D^{\mathcal{M}}$ is complete for each D, $S^{\mathcal{M}} = \bigcup_{D\in\mathcal{D}_S} D^{\mathcal{M}}$, and $\sup_{X\in D, Y\in D'} d_S^{\mathcal{M}}(X,Y) \leq C_{D,D'}$.
- Each function symbol f with domain $(S_1, \ldots, S_n)$ and codomain S is interpreted as a function $f^{\mathcal{M}} : S_1^{\mathcal{M}} \times \cdots \times S_n^{\mathcal{M}} \to S^{\mathcal{M}}$. Moreover, for each $\mathbf{D} = (D_1, \ldots, D_n) \in \mathcal{D}_{S_1} \times \cdots \times \mathcal{D}_{S_n}$, the function $f^{\mathcal{M}}$ maps $D_1^{\mathcal{M}} \times \cdots \times D_n^{\mathcal{M}}$ into $D_{f,\mathbf{D}}^{\mathcal{M}}$. Finally,

$f^{\mathcal{M}}$ restricted to $D_1^{\mathcal{M}} \times \cdots \times D_n^{\mathcal{M}}$ is uniformly continuous in the ith variable with modulus of continuity of $\omega_{f,\mathbf{D},i}$.

- Each relation symbol R with domain $(S_1, \ldots, S_n)$ is interpreted as a function $R^{\mathcal{M}} : S_1^{\mathcal{M}} \times \cdots \times S_n^{\mathcal{M}} \to \mathbb{R}$. Moreover, for each $\mathbf{D} = (D_1, \ldots, D_n) \in \mathcal{D}_{S_1} \times \cdots \times \mathcal{D}_{S_n}$, $f^{\mathcal{M}}$ is bounded by $N_{R,\mathbf{D}}$ on $M(D_1) \times \cdots \times M(D_n)$ and uniformly continuous in the ith argument with modulus of continuity of $\omega_{R,\mathbf{D},i}$.

The *language* $\mathcal{L}_{\mathrm{tr}}$ *of tracial* W*–*algebras* can be described as follows. We will also simultaneously describe how a tracial W*–algebra (M, τ) gives rise to an $\mathcal{L}_{\mathrm{tr}}$–structure $\mathcal{M}$, that is, how each symbol will be interpreted.

- A single sort, to be interpreted as the W*–algebra M. If $\mathcal{M} = (M, \tau)$ is a tracial W*–algebra, we denote the interpretation of this sort by $L^\infty(\mathcal{M})$ because of the intuition of tracial W*–algebras as non-commutative measure spaces.
- Domains of quantification $\{D_r\}_{r\in(0,\infty)}$, to be interpreted as the operator norm balls of radius r in M.
- The metric symbol d, to be interpreted as the metric induced by $\|\cdot\|_{2,\tau}$.
- A binary function symbol $+$, to be interpreted as addition.
- A binary function symbol $\cdot$, to be interpreted as multiplication.
- A unary function symbol $*$, to be interpreted as the adjoint operation.
- For each $\lambda \in \mathbb{C}$, a unary function symbol, to be interpreted as multiplication by λ.
- Function symbols of arity 0 (in other words constants) 0 and 1, to be interpreted as additive and multiplicative identity elements.
- Two unary relation symbols $\operatorname{Re}\mathrm{tr}$ and $\operatorname{Im}\mathrm{tr}$, to be interpreted the real and imaginary parts of the trace τ.
- For technical reasons explained by Farah, Hart, and Sherman [8], we also introduce for each d–variable non-commutative polynomial p a symbol $t_p : L^\infty(\mathcal{M})^d$ representing the evaluation of p, along with the appropriate range bounds $N_{t_p,\mathbf{r}}$ given by the supremum of $\|p(X_1, \ldots, X_d)\|$ over all $(X_1, \ldots, X_d)$ in a tracial W*–algebra $\mathcal{M}$.

Each function and relation symbol is assigned range bounds and moduli of continuity that one would expect, eg multiplication is supposed to map $D_r \times D_{r'}$ into $D_{rr'}$ with $\omega_{(D_r,D'_r),1}(t) = r't$ and $\omega_{(D_r,D'_r),2} = rt$.

Although not every $\mathcal{L}_{\mathrm{tr}}$–structure comes from a tracial W*–algebra, one can formulate axioms in the language such that any structure satisfying these axioms comes from a tracial W*–algebra [8, § 3.2]. In order to state this result precisely, we first have to explain formulas and sentences.

2.3 Syntax: Terms, formulas, conditions, and sentences

Terms in a language $\mathcal{L}$ are expressions obtained by iteratively composing the function symbols and variables. For example, if x_1, x_2, ... are variables in a sort S and $f : S \times S \to S$ and $g : S \times S \to S$ are function symbols, then $f(g(x_1, x_2), x_1)$ is a term. Each term has assigned range bounds and moduli of continuity in each variable which are the natural ones computed from those of the individual function symbols making up the composition. Any term f with variables $x_1 \in S_1, \ldots, x_k \in S_k$ and output in S can be interpreted in an $\mathcal{L}$–structure as a function $S_1^{\mathcal{M}} \times \cdots \times S_k^{\mathcal{M}} \to S^{\mathcal{M}}$. For example, in the language $\mathcal{L}_{\mathrm{tr}}$, the terms are expressions obtained from iterating scalar multiplication, addition, multiplication, and the $*$–operation on variables and the unit symbol 1. If (M, τ) is a tracial W*–algebra, then the interpretation of a term in $\mathcal{M}$ is a function represented by a $*$–polynomial.

Basic formulas in a language are obtained by evaluating relation symbols on terms. In other words, if $T_1, \ldots, T_k$ are terms valued in sorts $S_1, \ldots, S_k$, and R is a relation $S_1 \times \cdots \times S_k \to \mathbb{R}$, then $R(T_1, \ldots, T_n)$ is a basic formula. The basic formulas have assigned range bounds and moduli of continuity similar to the function symbols. In an $\mathcal{L}$–structure $\mathcal{M}$, a basic formula ϕ is interpreted as a function $\phi^{\mathcal{M}} : S_1^{\mathcal{M}} \times \cdots \times S_k^{\mathcal{M}} \to \mathbb{R}$. In $\mathcal{L}_{\mathrm{tr}}$, a basic formula can take the form $\operatorname{Re}\operatorname{tr}(f)$ or $\operatorname{Im}\operatorname{tr}(f)$ where f is an expression obtained by iterating the algebraic operations. Thus, when evaluated in a tracial W*–algebra, it corresponds to the real or imaginary part of the trace of a non-commutative $*$–polynomial.

Formulas are obtained from basic formulas by iterating several operations:

- Given a formulas $\phi_1, \ldots, \phi_n$ and $F : \mathbb{R}^n \to \mathbb{R}$ continuous, $F(\phi_1, \ldots, \phi_n)$ is a formula.
- If ϕ is a formula, D is a domain of quantification for some sort S, and x is one of our variables in S, then $\inf_{x \in D} \phi$ and $\sup_{x \in D} \phi$ are formulas.

Each occurrence of a variable in ϕ is either *bound* to a quantifier $\sup_{x \in D}$ or $\inf_{x \in D}$, or else it is *free*. We will often write $\phi(x_1, \ldots, x_n)$ for a formula to indicate that the free variables are $x_1, \ldots, x_n$.

All these formulas also have assigned range bounds and moduli of continuity. The moduli of continuity of $F(\phi_1, \ldots, \phi_n)$ are obtained by composition from the moduli of continuity of F and ϕ_j as in [3, § 2 Appendix and Theorem 3.5]. Next, if $\phi : S_1 \times \cdots \times S_n \to S$ and for $D \in \mathcal{D}_{S_n}$

$$\psi(x_1, \ldots, x_{n-1}) = \sup_{x_n \in D} \phi(x_1, \ldots, x_{n-1}, x_n)$$

then

$$\omega_{\psi,(D_1,\ldots,D_{n-1}),j} = \omega_{\phi,(D_1,\ldots,D_{n-1},D),j}.$$

Each formula has an interpretation in every $\mathcal{L}$–structure $\mathcal{M}$, defined by induction on the complexity of the formula. If $\phi = F(\phi_1, \ldots, \phi_n)$, then $\phi^{\mathcal{M}} = F(\phi_1^{\mathcal{M}}, \ldots, \phi_n^{\mathcal{M}})$. Similarly, if $\psi(x_1, \ldots, x_{n-1}) = \sup_{x_n \in D} \psi(x_1, \ldots, x_n)$, then

$$\psi^{\mathcal{M}}(X_1, X_2, \ldots, X_n) = \sup_{X_n \in D^{\mathcal{M}}} \phi^{\mathcal{M}}(X_1, \ldots, X_{n-1}).$$

Here $X_1, \ldots, X_n$ are elements of the sorts in the $\mathcal{L}$–structure $\mathcal{M}$, rather than formal variables.

Example 2.1 In $\mathcal{L}_{\mathrm{tr}}$, some terms are:

$$x_1 x_2, \quad (x_1 x_2)x_3 + (x_2^* x_3)(x_1 x_3^*)^*$$

A basic formula is:

$$\operatorname{Re}\operatorname{tr}(x_1 x_2 + x_3^*(x_2 x_1)^*)$$

Another formula is:

$$\operatorname{Re}\operatorname{tr}(x_1 x_2) + e^{\operatorname{Im}\operatorname{tr}(x_1^*(x_2 x_3^*)\operatorname{Re}\operatorname{tr}(x_4)}$$

We can also write a formula:

$$\sup_{x_1 \in D_2} [\operatorname{Re}\operatorname{tr}(x_1 x_2) + e^{\operatorname{Im}\operatorname{tr}(x_1^*(x_2 x_3^*)\operatorname{Re}\operatorname{tr}(x_4)}]$$

which will be interpreted as the supremum of the previous formula over x_1 in the ball of radius 2. In this formula, x_1 is bound to the quantifier $\sup_{x_1 \in D_2}$ and the variables x_2 and x_3 are free.

For convenience, we will assume that our formulas do not have two copies of the same variable (ie, if a variable is bound to a quantifier there is no other variable of the same name that is free or bound to a different quantifier). For instance, in the formula

$$\operatorname{Im}\operatorname{tr}(x_1) \sup_{x_1 \in D_1} \operatorname{Re}\operatorname{tr}(x_1 x_2 + x_3 x_1^*)$$

the first occurrence of x_1 is free while the latter two occurrences are bound to the quantifier $\sup_{x_1 \in D_1}$, but we can rewrite this formula equivalently as

$$\operatorname{Im}\operatorname{tr}(x_1) \sup_{y_1 \in D_1} \operatorname{Re}\operatorname{tr}(y_1 x_2 + x_3 y_1^*).$$

We will typically denote the free variables by $(x_i)_{i \in \mathbb{N}}$ and the bound variables by $(y_i)_{i \in \mathbb{N}}$. Lowercase letters will be used for formal variables while uppercase letters will be used for individual operators in operator algebras (or more generally elements of an $\mathcal{L}$–structure).

2.4 Theories, models, and axioms

A *sentence* is a formula with no free variables. If ϕ is a sentence, then the interpretation $\phi^{\mathcal{M}}$ in an $\mathcal{L}$–structure is simply a real number.

A *theory* T in a language $\mathcal{L}$ is a set of sentences. We say that an $\mathcal{L}$–structure $\mathcal{M}$ *models* the theory T, or $\mathcal{M} \models \mathrm{T}$ if $\phi^{\mathcal{M}} = 0$ for all $\phi \in \mathrm{T}$.

If $\mathcal{M}$ is an $\mathcal{L}$–structure, then the *theory of* $\mathcal{M}$, denoted $\mathrm{Th}(\mathcal{M})$ is the set of sentences ϕ such that $\phi^{\mathcal{M}} = 0$. As observed by Farah, Hart, and Sherman [8], the theory of $\mathcal{M}$ also uniquely determines the values $\phi^{\mathcal{M}}$ of all sentences ϕ since $\phi - c$ is a sentence for every constant $c \in \mathbb{R}$.

More generally, if $\mathcal{C}$ is a class of $\mathcal{L}$–structures, then $\mathrm{Th}(\mathcal{C})$ is the set of all sentences ϕ such that $\phi^{\mathcal{M}} = 0$ for all $\mathcal{M}$ in $\mathcal{C}$. The class $\mathcal{C}$ is said to be *axiomatizable* if every $\mathcal{L}$–structure that models $\mathrm{Th}(\mathcal{C})$ is actually in $\mathcal{C}$.

Farah, Hart, and Sherman [8, §3.2] showed that the class of $\mathcal{L}_{\mathrm{tr}}$–structures that represent actual tracial W*–algebras is axiomatizable. The axioms, roughly speaking, encode the fact that $\mathcal{M}$ is a $*$–algebra, the fact that τ is a tracial state, the fact that $\|xy\|_{L^2(\mathcal{M})} \leq r\|y\|_{L^2(\mathcal{M})}$ for $x \in D_r^{\mathcal{M}}$, the relationship between the distance and the trace, the fact that $D_r^{\mathcal{M}}$ is contained in $D_{r'}^{\mathcal{M}}$ for $r < r'$ (that is, $\sup_{x \in D_r} \inf_{y \in D_{r'}} d(x,y) = 0$), and the fact that t_p agrees with the evaluation of the non-commutative polynomial p.

The theory of tracial W*–algebras will be denoted T_{tr}. It is also shown in [8] that II_1 factors (infinite-dimensional tracial W*–algebras with trivial center) are axiomatizable by a theory $\mathrm{T}_{\mathrm{II}_1}$.

2.5 Ultraproducts

An important construction for continuous model theory and for W*–algebras is the ultraproduct. Ultraproducts are a way of constructing a limiting object out of arbitrary sequences (or more generally indexed families) of objects. In order to force limits to exist, one uses a device called an ultrafilter.

Let I be an index set. An *ultrafilter* $\mathcal{U}$ on I is a collection of subsets of I such that

- $\varnothing \notin \mathcal{U}$.
- If $A \subseteq B \subseteq I$ and $A \in \mathcal{U}$, then $B \in \mathcal{U}$.
- If $A, B \in \mathcal{U}$, then $A \cap B \in \mathcal{U}$.
- For each $A \subseteq I$, either $A \in \mathcal{U}$ or $A^c \in \mathcal{U}$.

If $\mathcal{U}$ is an ultrafilter on I, Ω is a topological space, and $f : I \to \Omega$ is a function, then we say that

$$\lim_{i \to \mathcal{U}} f(i) = w$$

if for every neighborhood $O \ni w$ in Ω, the preimage $f^{-1}(O)$ is an element of $\mathcal{U}$. Now if $i \in I$, there is an ultrafilter $\mathcal{U}_i := \{A \subseteq I : i \in A\}$, which is called a *principal ultrafilter*. All other ultrafilters are called *non-principal* or *free ultrafilters*.

The set of all ultrafilters on I can be identified with the Stone-Čech compactification of I, where I is given the discrete topology (see eg Hindman and Strauss [19, § 3]). The principal ultrafilters correspond to the points of the original space I. In particular, this means that if Ω is a compact Hausdorff space and $f : I \to \Omega$ is a function, then $\lim_{i \to \mathcal{U}} f(i)$ exists in Ω.

Now consider a language $\mathcal{L}$. Let I be an index set, $\mathcal{U}$ an ultrafilter on I, and for each $i \in I$, let $\mathcal{M}$ be an $\mathcal{L}$–structure. The ultraproduct $\prod_{i \to \mathcal{U}} \mathcal{M}_i$ is the $\mathcal{L}$–structure $\mathcal{M}$ defined as follows (see [3, § 5]). For each sort S, consider tuples $(X_i)_{i \in I}$ where $X_i \in S^{\mathcal{M}_i}$.

- Let's call $(X_i)_{i \in I}$ *confined* if there exists $D \in \mathcal{D}_S$ such that $X_i \in D^{\mathcal{M}_i}$ for all i.
- Let's call $(X_i)_{i \in I}$ and $(Y_i)_{i \in I}$ *equivalent* if $\lim_{i \to \mathcal{U}} d_S^{\mathcal{M}_i}(X_i, Y_i) = 0$.
- For a confined tuple $(X_i)_{i \in I}$, let $[X_i]_{i \in I}$ denote its equivalence class.

We define $S^{\mathcal{M}}$ to be the set of equivalence classes of confined tuples $(X_i)_{i \in I}$. The metric $d_S^{\mathcal{M}}$ on $S^{\mathcal{M}}$ is then given by

$$d_S^{\mathcal{M}}([X_i]_{i \in I}, [Y_i]_{i \in I}) = \lim_{i \to \mathcal{U}} d_S^{\mathcal{M}_i}(X_i, Y_i).$$

This is independent of the choice of representative for the equivalence classes because of the triangle inequality, and it is finite because if $X_i \in D^{\mathcal{M}_i}$ and $Y_i \in (D')^{\mathcal{M}_i}$ for all i, then $d_S^{\mathcal{M}_i}(X_i, Y_i) \leq C_{D,D'}$. Then $S^{\mathcal{M}}$ is a metric space and $S^{\mathcal{M}} = \bigcup_{D \in \mathcal{D}_S} D^{\mathcal{M}}$, where $D^{\mathcal{M}}$ is the set of classes $[X_i]_{i \in I}$ with $X_i \in D^{\mathcal{M}_i}$ for all i. Moreover, $D^{\mathcal{M}}$ is automatically complete [3, Proposition 5.3].

Each function symbol $f : S_1 \times \cdots \times S_n \to S$ receives its interpretation $f^{\mathcal{M}}$ through

$$f^{\mathcal{M}}([X_{1,i}]_{i \in I}, \ldots, [X_{n,i}]_{i \in I}) = [f^{\mathcal{M}_i}(X_{1,i}, \ldots, X_{n,i})]_{i \in I}$$

which is well-defined because of the uniform continuity of f on each domain of quantification, and similarly, each relation receives its interpretation in $\mathcal{M}$ through

$$R^{\mathcal{M}}([X_{1,i}]_{i \in I}, \ldots, [X_{n,i}]_{i \in I}) = \lim_{i \to \mathcal{U}} R^{\mathcal{M}_i}(X_{1,i}, \ldots, X_{n,i}).$$

One can verify by the same reasoning as [3, § 5] that $\mathcal{M}$ is indeed an $\mathcal{L}$–structure.

One of the reasons ultraproducts are so important is because of the following result, known as (the continuous analog of) Łos's theorem. See [3, Theorem 5.4].

Theorem 2.2 *Let $\mathcal{L}$ be a language, $\mathcal{U}$ an ultrafilter on an index set I, and $\mathcal{M}_i$ an $\mathcal{L}$–structure for each $i \in I$. Let $\mathcal{M} = \prod_{i \to \mathcal{U}} \mathcal{M}_i$. If ϕ is a formula with free variables $x_1 \in S_1, \ldots, x_n \in S_n$, then for any $[X_{1,i}]_{i \in I} \in S_1^{\mathcal{M}}, \ldots, [X_{n,i}]_{i \in I} \in S_n^{\mathcal{M}}$, we have*

$$\phi^{\mathcal{M}}([X_{1,i}]_{i \in I}, \ldots, [X_{n,i}]_{i \in I}) = \lim_{i \to \mathcal{U}} \phi^{\mathcal{M}_i}(X_{1,i}, \ldots, X_{n,i}).$$

Corollary 2.3 *In the situation of the previous theorem, if T is an $\mathcal{L}$–theory, and if $\mathcal{M}_i \models \mathrm{T}$ for all i, then $\mathcal{M} \models \mathrm{T}$.*

In particular, this shows that an ultraproduct of $\mathcal{L}_{\mathrm{tr}}$–structures that are tracial W*–algebras will also be an $\mathcal{L}_{\mathrm{tr}}$–structure that is a tracial W*–algebra. One can verify that the model-theoretic ultraproduct agrees in this case with the ultraproduct of tracial W*–algebras.

3 Definable predicates and functions

This section describes types, definable predicates, and definable functions. The material in §3.1 – §3.4 is largely a mixture of folklore and adaptations of [3]; our main contribution is to write down the results in the setting of *infinite tuples* and *domains of quantification*. In §3.5, we give a characterization of quantifier-free definable functions in $\mathcal{L}_{\mathrm{tr}}$ based on Jekel [22, §13] and Hayes, Jekel, Nelson, and Sinclair [18, §2].

3.1 Types

Definition 3.1 Let $\mathbf{S} = (S_j)_{j \in \mathbb{N}}$ be an $\mathbb{N}$–tuple of sorts in $\mathcal{L}$. Let $\mathcal{F}_{\mathbf{S}}$ be the space of $\mathcal{L}$–formulas with free variables $(x_j)_{j \in \mathbb{N}}$ with x_j from the sort S_j. If $\mathcal{M}$ is an $\mathcal{L}$–structure and $\mathbf{X} \in \prod_{j \in \mathbb{N}} S_j^{\mathcal{M}}$, then the *type* of $\mathbf{X}$ is the map:

$$\mathrm{tp}^{\mathcal{M}}(\mathbf{X}) \colon \mathcal{F}_{\mathbf{S}} \to \mathbb{R}, \quad \phi \mapsto \phi^{\mathcal{M}}(\mathbf{X})$$

Definition 3.2 Let $\mathbf{S} = (S_j)_{j \in \mathbb{N}}$ be an $\mathbb{N}$–tuple of sorts in $\mathcal{L}$, and let T be an $\mathcal{L}$–theory. If $\mathbf{D} \in \prod_{j \in \mathbb{N}} \mathcal{D}_{S_j}$, then we denote by $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ the set of types $\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$ of all $\mathbf{X} \in \prod_{j \in \mathbb{N}} D_j^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$.

Definition 3.3 If $\mathbf{S}$ is an $\mathbb{N}$–tuple of $\mathcal{L}$–sorts, the set $\mathcal{F}_{\mathbb{S}}$ of formulas defines a real vector space. For each $\mathcal{L}$–structure $\mathcal{M}$ and $\mathbf{X} \in \prod_{j \in \mathbb{N}} S_j^{\mathcal{M}}$, the type $\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$ is a (real) linear map $\mathcal{F}_{\mathbf{S}} \to \mathbb{R}$. Thus, for each $\mathcal{L}$–theory T and $\mathbf{D} \in \prod_{j \in \mathbb{N}} \mathcal{D}_{S_j}$, the space $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is a subset of the dual $\mathcal{F}_{\mathbb{S}}^{\dagger}$. We equip $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ with the weak-$\star$ topology (also known as the *logic topology*).

The following observation is well known in continuous model theory; see [3, Corollary 5.12, Proposition 8.6]

Observation 3.4 $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ *is compact in the weak-$\star$ topology.*

Proof Each formula ϕ has a range bound $N_{\phi,\mathbf{D}}$ such that $|\phi^{\mathcal{M}}(\mathbf{X})| \leq N_{\phi,\mathbf{D}}$ for all $\mathcal{L}$–structures $\mathcal{M}$ and all $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j$. Thus, $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is a subset of $\prod_{\phi\in\mathcal{F}_{\mathcal{D}}}[-N_{\phi,\mathbf{D}}, N_{\phi,\mathbf{D}}]$ with the product topology, which is compact by Tychonoff's theorem.

Moreover, $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is a closed subset. While closedness can be expressed in terms of nets, it can also be expressed in terms of ultralimits. A set A is closed if and only if for every I and $f : I \to A$ and ultrafilter $\mathcal{U}$, the limit $\lim_{i\to\mathcal{U}} f(i)$ exists in A. It then follows from Theorem 2.2 that if $[X_i]_{i\in I}$ is an element of an ultraproduct $\mathcal{M}$ of $\mathcal{L}$–structures $\mathcal{M}_i$, then $\mathrm{tp}^{\mathcal{M}}([X_i]_{i\in I}) = \lim_{i\to\mathcal{U}} \mathrm{tp}^{\mathcal{M}_i}(X_i)$. □

Although many times authors choose to work with $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ for each $\mathbf{D}$, we find it convenient to specify a topology on the entire space of types $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$ that extends the topology on each $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$, so that our later results can be stated about $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$ globally. The topology on $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is given by a categorical colimit of the topologies on $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$.

Definition 3.5 For a language $\mathcal{L}$, tuple $\mathbf{S}$ of sorts, and theory T, let $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$ denote the space of $\mathbf{S}$–types for all $\mathcal{M} \models \mathrm{T}$. Note that $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$ is the union of all $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ for all $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$. We say that $\mathcal{O} \subseteq \mathbb{S}_{\mathbf{S}}(\mathrm{T})$ is open if $\mathcal{O} \cap \mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is open for every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$; this defines a topology on $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$, which we will also call the *logic topology*.

Observation 3.6 *For a language $\mathcal{L}$, tuple $\mathbf{S}$ of sorts, theory* T*, and $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$, the inclusion map $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{S}_{\mathbf{S}}(\mathrm{T})$ is a topological embedding.*

Proof Note that $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$ is Hausdorff; indeed, μ and ν are two distinct types, then there exists a formula ϕ with $\mu(\phi) \neq \nu(\phi)$. One can check that the sets

$$U = \{\sigma \in \mathbb{S}_{\mathbf{S}}(\mathrm{T}) : |\sigma(\phi) - \mu(\phi)| < |\sigma(\phi) - \nu(\phi)|\}$$

and

$$V = \{\sigma \in \mathbb{S}_{\mathbf{S}}(\mathrm{T}) : |\sigma(\phi) - \nu(\phi)| < |\sigma(\phi) - \mu(\phi)|\}$$

are open and they separate μ and ν.

Continuity of the inclusion map $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{S}_{\mathbf{S}}(\mathrm{T})$ follows from the definition of open sets in $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$. Then since $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is compact and $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$ is Hausdorff, the map is a topological embedding. □

Observation 3.7 *For a language $\mathcal{L}$, tuple $\mathbf{S}$ of sorts, theory T, and topological space Ω, a function $\psi : \mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \Omega$ is continuous if and only if $\psi|_{\mathbb{S}_{\mathbf{D}}(\mathrm{T})}$ is continuous for every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$.*

Proof This follows from the definition of open sets in $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$. □

3.2 Definable predicates

Next, we describe definable predicates, which are certain limits of formulas. It will turn out that definable predicates correspond precisely to continuous functions $\mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{R}$, and thus they are a natural completion of the space of formulas in the setting of continuous model theory. Our approach to the definition will be semantic rather than syntactic, defining these objects immediately in terms of their interpretations.

Definition 3.8 Let $\mathcal{L}$ be a language and T an $\mathcal{L}$–theory. A *definable predicate relative to* T is a collection of functions $\phi^{\mathcal{M}} : \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}} \to \mathbb{R}$ (for each $\mathcal{M} \models \mathrm{T}$) such that for every collection of domains $\mathbf{D} = (D_j)_{j\in\mathbb{N}}$ and every $\epsilon > 0$, there exists a finite $F \subseteq \mathbb{N}$ and an $\mathcal{L}$–formula $\psi(x_j : j \in F)$ such that whenever $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$, we have

$$|\phi^{\mathcal{M}}(\mathbf{X}) - \psi^{\mathcal{M}}(X_j : j \in F)| < \epsilon.$$

In other words a definable predicate is an object that can be uniformly approximated by a formula on any product of domains of quantification, where the approximation works uniformly for all models of the theory T. This is done relative to T because, for instance, in the study of tracial W^*–algebras we do not care if the definable predicate makes sense to evaluate on arbitrary $\mathcal{L}_{\mathrm{tr}}$–structures, only those which actually come from tracial W^*–algebras.

Note that every formula defines a definable predicate. However, two formulas as defined in the previous section (where the range bounds and moduli of continuity are part of the definition) may reduce to the same definable predicate (especially given the restriction that we work relative to a given theory T).

The next proposition describes definable predicates as continuous functions on the space of types. This is an adaptation of [3, Theorem 9.9] to the setting with domains of quantification.

Proposition 3.9 *Let $\mathcal{L}$ be a language and* T *an $\mathcal{L}$–theory. Let ϕ be a collection of functions $\phi^{\mathcal{M}} : \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}} \to \mathbb{R}$ for each $\mathcal{M} \models$ T. The following are equivalent:*

(1) *ϕ is a definable predicate relative to* T.
(2) *There exists a continuous $\gamma : \mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{R}$ such that $\phi^{\mathcal{M}}(\mathbf{X}) = \gamma(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$ for all $\mathcal{M} \models$ T and $\mathbf{X} \in \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}}$.*

Proof (1) $\implies$ (2). First, suppose that ϕ is a formula. Then by definition of type, $\phi^{\mathcal{M}}(\mathbf{X})$ only depends on the type of $\mathbf{X}$ in $\mathcal{M}$, and hence $\phi^{\mathcal{M}}(\mathbf{X}) = \gamma(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$ for some $\gamma : \mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{R}$. For each $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$, the restriction of γ to a map $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{R}$ is continuous by definition of the weak-$\star$ topology. Hence, by Observation 3.7, γ is a continuous function $\mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{R}$.

Now let ϕ be a general definable predicate. Fix $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$. Then taking $\epsilon = 1/n$ in Definition 3.8, there exists a formula $\phi_{\mathbf{D},n}$ depending on finitely many of the variables x_j, such that

$$|\phi^{\mathcal{M}}(\mathbf{X}) - \phi^{\mathcal{M}}_{\mathbf{D},n}(\mathbf{X})| < \frac{1}{n} \tag{3–1}$$

for all $\mathcal{M} \models$ T and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$. By the previous paragraph, there exists a continuous $\gamma_{\mathbf{D},n} : \mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{R}$ such that $\phi^{\mathcal{M}}_{\mathbf{D},n}(\mathbf{X}) = \gamma_{\mathbf{D},n}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$ for all $\mathcal{M} \models$ T and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$. By (3–1),

$$\sup_{\mu\in\mathbb{S}_{\mathbf{D}}(\mathrm{T})} |\gamma_{\mathbf{D},n}(\mu) - \gamma_{\mathbf{D},m}(\mu)| \leq \frac{1}{n} + \frac{1}{m}$$

which implies that the sequence $\gamma_{\mathbf{D},n}$ converges as $n \to \infty$ to a continuous $\gamma_{\mathbf{D}} : \mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{R}$. Also, by (3–1),

$$\phi^{\mathcal{M}}(\mathbf{X}) = \gamma_{\mathbf{D}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})))$$

for $\mathcal{M} \models$ T and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$. This in turn implies that $\gamma_{\mathbf{D}}$ and $\gamma_{\mathbf{D}'}$ agree on $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \cap \mathbb{S}_{\mathbf{D}'}(\mathrm{T})$ for any $\mathbf{D}$ and $\mathbf{D}' \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$. Thus, for some function $\gamma : \mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{R}$, we have $\gamma_{\mathbf{D}} = \gamma|_{\mathbf{D}}$ for $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$. By Observation 3.7, γ is continuous on $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$.

(2) $\implies$ (1). Assume there exists $\gamma : \mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{R}$ continuous such that $\phi^{\mathcal{M}}(\mathbf{X}) = \gamma(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$ for all $\mathcal{M} \models$ T and $\mathbf{X} \in \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}}$. Fix $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$. Let $\mathcal{A}$ be the set of functions $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{R}$ given by the application of formulas $\phi \in \mathcal{F}_{\mathbf{S}}$. Then $\mathcal{A}$ is a subalgebra of $C(\mathbb{S}_{\mathbf{D}}(\mathrm{T}), \mathbb{R})$ since formulas are closed under sums, products, and scalar multiplication by real numbers. Moreover, $\mathcal{A}$ separates points because by definition two types are the same if they agree on all formulas. Therefore, since $\gamma|_{\mathbb{S}_{\mathbf{D}}(\mathrm{T})}$ is continuous, the Stone-Weierstrass theorem implies that there exists a formula ψ depending on finitely many of the variables x_j such that $|\phi^{\mathcal{M}}(\mathbf{X}) - \psi^{\mathcal{M}}(\mathbf{X})| < \epsilon$ whenever $\mathcal{M} \models$ T and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$. □

Lemma 3.10 *If* $\mathbf{S}$ *is an* $\mathbb{N}$*–tuple of types and* $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$*, then the logic topology on* $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ *agrees with the weak-*$\star$ *topology obtained by viewing* $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ *as a subspace of the dual of the vector space of definable predicates.*

Proof We defined the logic topology as the weak-$\star$ topology generated by the pairing of types with formulas in variables $x_j \in S_j$ for $j \in \mathbb{N}$. Since every formula gives a definable predicate, the weak-$\star$ topology obtained from the pairing with definable predicates is at least as strong as the logic topology. On other hand, for each $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$, every definable predicate can be approximated uniformly by $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$, and hence the pairing with each definable predicate ϕ defines a map $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{R}$ that is continuous with respect to the logic topology, and hence the logic topology is at least as strong as the weak-$\star$ topology obtained from pairing with definable predicates. □

Just like formulas, definable predicates are uniformly continuous on any product of domains of quantification. But to say this properly, we should clarify what "uniform continuity" means for a function of infinitely many variables. If Ω_j is a metric space, then $\prod_{j\in\mathbb{N}} \Omega_j$ with the product topology is metrizable but without a canonical choice of metric. However, we will say that $\phi : \prod_{j\in\mathbb{N}} \Omega_j \to \mathbb{R}$ is uniformly continuous if for every $\epsilon > 0$, there exists a finite $F \subseteq \mathbb{N}$ and $\delta > 0$, such that

$$d_j(x_j, y_j) < \delta \text{ for } j \in F \implies |\phi(x) - \phi(y)| < \epsilon.$$

In other words, uniform continuity is defined with respect to the product *uniform structure* on $\prod_{j\in\mathbb{N}} \Omega_j$. (See for instance James [21] for background on uniform structures.)

Observation 3.11 *If* $\phi = (\phi^{\mathcal{M}})$ *is a definable predicate over* $\mathcal{L}$ *relative to* T*, then* ϕ *satisfies the following uniform continuity property:*

For every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$ *and* $\epsilon > 0$*, there exists a finite* $F \subseteq \mathbb{N}$ *and* $\delta > 0$ *such that, for every* $\mathcal{M} \models \mathrm{T}$ *and* $\mathbf{X}, \mathbf{Y} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$*,*

$$d^{\mathcal{M}}(X_j, Y_j) < \delta \text{ for all } j \in F \implies |\phi^{\mathcal{M}}(\mathbf{X}) - \phi^{\mathcal{M}}(\mathbf{Y})| < \epsilon.$$

Moreover, for every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$*, there exists a constant* C *such that* $|\phi^{\mathcal{M}}| \leq C$ *for all* $\mathcal{M} \models \mathrm{T}$*.*

By construction, this result holds for formulas in finitely many X_j's, and it holds for general definable predicates by the principle that uniform continuity and boundedness are preserved under uniform limits.

Another useful property is that definable predicates are closed under the same types of operations as formulas. In fact, we can use infinitary rather than finitary operations. Point (1) here is an adaptation of [3, Proposition 9.3].

Lemma 3.12

(1) *If* $F : \mathbb{R}^{\mathbb{N}} \to \mathbb{R}$ *is continuous (where* $\mathbb{R}^{\mathbb{N}}$ *has the product topology) and* $(\phi_j)_{j\in\mathbb{N}}$ *are definable predicates* $\prod_{j\in\mathbb{N}} S_j \to \mathbb{R}$ *in* $\mathcal{L}$ *relative to* T*, then* $F((\phi_j)_{j\in\mathbb{N}})$ *is a definable predicate.*

(2) *If* ϕ *is a definable predicate* $\prod_{j\in\mathbb{N}} S_j \times \prod_{j\in\mathbb{N}} S'_j \to \mathbb{R}$ *in* $\mathcal{L}$ *relative to* T *and* $\mathbf{D}' \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S'_j}$*, then*

$$\psi^{\mathcal{M}}(\mathbf{X}, \mathbf{Y}) := \inf_{\mathbf{Y}\in\prod_{j\in\mathbb{N}}(D'_j)^{\mathcal{M}}} \phi(\mathbf{X}, \mathbf{Y})$$

is also a definable predicate in $\mathcal{L}$ *relative to* T*.*

Proof (1) This follows from 3.9 and the fact that continuity is preserved by composition.

(2) Fix $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$ and $\epsilon > 0$. Then there exist a formula ϕ_0 whose free variables are a finite subset of the x_j's and y_j's, such that $|\phi^{\mathcal{M}} - \phi_0^{\mathcal{M}}| < \epsilon$ on $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}} \times \prod_{j\in\mathbb{N}} (D'_j)^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$. Note that

$$\psi_0^{\mathcal{M}}(\mathbf{X}) = \inf_{\mathbf{Y}\in\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}} \phi_0^{\mathcal{M}}(\mathbf{X}, \mathbf{Y})$$

also defines a formula because the infimum is effectively over only finitely many Y_j's. Also, $|\psi_0^{\mathcal{M}} - \psi^{\mathcal{M}}| \leq \epsilon$ on $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$. Therefore, ψ is a definable predicate. □

We conclude with a brief remark on separability since we will use the separability of $\mathcal{L}_{\mathrm{tr}}$ in the sequel. For a $\mathcal{L}$–theory T, we equip $C(\mathbb{S}_{\mathbf{S}}(\mathrm{T}))$ with the locally convex topology generated by the family of seminorms

$$\phi \mapsto \|\phi|_{\mathbb{S}_{\mathbf{D}}(\mathrm{T})}\|_{C(\mathbb{S}_{\mathbf{D}}(\mathrm{T}))}$$

for $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$. In other words, a net ϕ_i to ϕ converges in this topology if and only if $\|(\phi - \phi_i)|_{\mathbb{S}_{\mathbf{D}}(\mathrm{T})}\|_{C(\mathbb{S}_{\mathbf{D}}(\mathrm{T}))} \to 0$ for all $\mathbf{D}$.

Definition 3.13 A language $\mathcal{L}$ is *separable* if

(1) $\mathcal{L}$ has countably many sorts.

(2) For every $\mathbb{N}$–tuple $\mathbb{S}$ of sorts, the space $C(\mathbb{S}_{\mathbf{S}}(\varnothing))$ is separable, where $\varnothing$ denotes the empty theory.

Observation 3.14 *If $\mathcal{L}$ is a separable language, T is an $\mathcal{L}$–theory, and $\mathbf{D}$ is an $\mathbb{N}$–tuple of domains from an $\mathbb{N}$–tuple of sorts $\mathbf{S}$, then $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is metrizable.*

Proof By separability of $\mathcal{L}$, there is a dense sequence $(\phi_n)_{n\in\mathbb{N}}$ in $C(\mathbb{S}_{\mathbf{S}}(\varnothing))$. Since the restriction maps $C(\mathbb{S}_{\mathbf{S}}(\varnothing)) \to C(\mathbb{S}_{\mathbf{D}}(\varnothing))$ and $C(\mathbb{S}_{\mathbf{D}}(\varnothing)) \to C(\mathbb{S}_{\mathbf{D}}(\mathrm{T}))$ are continuous, (ϕ_n) also defines a dense subset in $C(\mathbb{S}_{\mathbf{D}}(\mathrm{T}))$. For each n, there exists a constant K_n such that $|\mu(\phi_n)| \leq K_n$ for all $\mu \in \mathbb{S}_{\mathbf{D}}(\mathrm{T})$; this holds because ϕ_n can be uniformly approximated on $\mathbf{D}$ by formulas, which are also uniformly bounded. Then we may define a metric on $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ by:

$$d(\mu,\nu) = \sum_{n\in\mathbb{N}} \frac{1}{2^n K_n} |\mu(\phi_n) - \nu(\phi_n)|$$

The verification that this induces the weak-$*$ topology is routine. The types $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ induce linear functionals on $C(\mathbf{S}_{\mathbf{D}}(\mathrm{T}))$, or in other words, $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is contained in the unit ball of the dual of $C(\mathbb{S}_{\mathbf{D}}(\mathrm{T}))$ so convergence on a dense subset of $C(\mathbb{S}_{\mathbf{D}}(\mathrm{T}))$ is equivalent to convergence on all of $C(\mathbb{S}_{\mathbf{D}}(\mathrm{T}))$. □

Observation 3.15 *The language $\mathcal{L}_{\mathrm{tr}}$ of tracial W^*–algebras is separable.*

Proof Consider $\mathcal{L}_{\mathrm{tr}}$ formulas obtained using only scalar multiplication by numbers in $\mathbb{Q}[i]$ rather than $\mathbb{C}$ and using only suprema and infima over D_r for $r \in \mathbb{Q} \cap (0,\infty)$. There are only countably many such formulas, and one can show that these formulas are dense in the space of definable predicates. □

3.3 Definable functions

Although definable functions are often defined only for finite tuples, it is useful for the theory of covering entropy to work with infinite tuples as both the input and the output functions. The following "functional" description of definable functions makes it easy to prove properties relating them with definable predicates and the type space.

Definition 3.16 Let $\mathbf{S}$ and $\mathbf{S}'$ be $\mathbb{N}$–tuples of sorts in the language $\mathcal{L}$. A *definable function* $f : \prod_{j\in\mathbb{N}} S_j \to \prod_{j\in\mathbb{N}} S'_j$ relative to the $\mathcal{L}$–theory T is a collection of maps $\mathbf{f}^{\mathcal{M}} : \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}} \to \prod_{j\in\mathbb{N}} (S'_j)^{\mathcal{M}}$ for $\mathcal{M} \models \mathrm{T}$ satisfying the following conditions:

(1) For each $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$, there exists $\mathbf{D}' \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S'_j}$ such that for every $\mathcal{M} \models \mathrm{T}$, $\mathbf{f}^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ into $\prod_{j\in\mathbb{N}} (D'_j)^{\mathcal{M}}$.

(2) Whenever $\tilde{\mathbf{S}}$ is another tuple of sorts and ϕ is a definable predicate relative to T in the free variables $x'_j \in S'_j$ and $\tilde{x}_j \in \tilde{S}_j$ for $j \in \mathbb{N}$, then $\phi(\mathbf{f}(\mathbf{x}), \tilde{\mathbf{x}})$ is a definable predicate in the variables $\mathbf{x} = (x_j)_{j\in\mathbb{N}}$ and $\tilde{\mathbf{x}} = (\tilde{x}_j)_{j\in\mathbb{N}}$.

The next proposition gives a more down-to-earth characterization of definable functions which can be more easily checked in examples. This is in fact typically used as the definition [3, Definition 9.22].

Proposition 3.17 *Let* $\mathbf{S}$ *and* $\mathbf{S}'$ *be* $\mathbb{N}$*–tuples of types in the language* $\mathcal{L}$ *and let* T *be an* $\mathcal{L}$*–theory. Let* $\mathbf{f} : \prod_{j\in\mathbb{N}} S_j \to \prod_{j\in\mathbb{N}} S'_j$ *be a collection of maps* $\mathbf{f}^{\mathcal{M}} : \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}} \to \prod_{j\in\mathbb{N}} (S'_j)^{\mathcal{M}}$ *for* $\mathcal{M} \models \mathrm{T}$ *satisfying (1) of Definition 3.16. Then* $\mathbf{f}$ *is a definable function if and only if, for each* $k \in \mathbb{N}$*, the map* $\phi_k(\mathbf{x}, y) = d(f_k(\mathbf{x}), y)$ *is a definable predicate on* $\prod_{j\in\mathbb{N}} S_j \times S'_k$.

Proof ($\Longrightarrow$) Let $\tilde{\mathbf{S}} = \mathbf{S}'$, and consider the definable predicate $\phi(\mathbf{x}', \tilde{\mathbf{x}}) = d(x'_k, \tilde{x}_k)$. Taking $\tilde{\mathbf{S}} = \mathbf{S}'$ in Definition 3.16 (2), we see that if $\mathbf{f}$ is a definable function, then $\phi(\mathbf{f}(\mathbf{x}), \tilde{\mathbf{x}}) = d(f_k(\mathbf{x}), \tilde{x}_k)$ is a definable predicate. So substituting y for $\tilde{x}_k$, we have proved the claim.

($\Longleftarrow$) In order to verify (2) of Definition 3.16, let $\tilde{\mathbf{S}}$ be an $\mathbb{N}$–tuple of sorts, and let $\phi(\mathbf{x}', \tilde{\mathbf{x}})$ be a definable predicate on $\prod_{j\in\mathbb{N}} S'_j \times \prod_{j\in\mathbb{N}} \tilde{S}_j$. We need to show that $\psi(\mathbf{x}, \tilde{\mathbf{x}}) = \phi(\mathbf{f}(\mathbf{x}), \tilde{\mathbf{x}})$ is a well-defined definable predicate relative to T. Thus, to check Definition 3.8, fix $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$ and $\tilde{\mathbf{D}} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{\tilde{S}_j}$ and $\epsilon > 0$. Since we assumed that Definition 3.16 (1) holds, there exists $\mathbf{D}'$ such that f maps $\prod_{j\in\mathbb{N}} D_j$ into $\prod_{j\in\mathbb{N}} D'_j$.

By Definition 3.8, there exists a formula η depending on finitely many of the variables x'_j and $\tilde{x}_j$ that approximates ϕ within $\epsilon/2$ on $\prod_{j\in\mathbb{N}} (D'_j)^{\mathcal{M}} \times \prod_{j\in\mathbb{N}} (\tilde{D}_j)^{\mathcal{M}}$. Let F' be the set of indices j such that η depends on x'_j. For $t > 0$, let

$$\psi_t(\mathbf{x}, \tilde{\mathbf{x}}) = \inf_{y_j \in D'_j : j\in F} \left[\eta(y_j : j \in F, \tilde{\mathbf{x}}) + \frac{1}{t} \sum_{j\in F} d(\mathbf{f}_j(\mathbf{x}), y_j) \right]$$

which is a definable predicate by our assumption on $\mathbf{f}$ and by Lemma 3.12 (2).

We want to show that ψ_t is close to ϕ when t is sufficiently small. We automatically have $\psi_t^{\mathcal{M}}(\mathbf{X}, \tilde{\mathbf{X}}) \leq \eta^{\mathcal{M}}(\mathbf{f}(\mathbf{X}), \tilde{\mathbf{X}})$ for $\mathbf{X} \in \prod_{j\in\mathbb{N}} (D'_j)^{\mathcal{M}}$ and $\tilde{\mathbf{X}} \in \prod_{j\in\mathbb{N}} (\tilde{D}_j)^{\mathcal{M}}$ when $\mathcal{M} \models \mathrm{T}$ since $\mathbf{f}_j(\mathbf{X})$ is a value of $\mathbf{Y}$ participating in the infimum. To get a bound in the other direction, first observe that since η is a formula, $|\eta|$ is bounded on $\prod_{j\in\mathbb{N}} D'_j \times \prod_{j\in\mathbb{N}} \tilde{D}_j$ by some constant C. We then observe using the triangle inequality that

$$\eta^{\mathcal{M}}(\mathbf{Y}, \tilde{\mathbf{X}}) + \frac{1}{t} \sum_{j\in F} d^{\mathcal{M}}(Y_j, f_j^{\mathcal{M}}(\mathbf{X})) \geq \eta^{\mathcal{M}}(\mathbf{f}^{\mathcal{M}}(\mathbf{X}), \tilde{\mathbf{X}})$$

unless $\sum_{j\in F} d^{\mathcal{M}}(Y_j, f_j^{\mathcal{M}}(\mathbf{X})) < 2Ct$, and therefore the infimum is witnessed by $\mathbf{Y}$ such that $\sum_{j\in F} d^{\mathcal{M}}(Y_j, f_j^{\mathcal{M}}(\mathbf{X})) < 2Ct$. Furthermore, by the uniform continuity property of the formula η (Observation 3.11), there exists t such that if $\sum_{j\in F} d^{\mathcal{M}}(Y_j, f_j^{\mathcal{M}}(\mathbf{X})) < 2Ct$, then

$$|\eta^{\mathcal{M}}(\mathbf{Y}, \tilde{\mathbf{X}}) - \eta^{\mathcal{M}}(\mathbf{f}(\mathbf{X}), \tilde{\mathbf{X}})| < \frac{\epsilon}{2}.$$

It follows that for this value of t,

$$\eta^{\mathcal{M}}(\mathbf{f}(\mathbf{X}), \tilde{\mathbf{X}}) - \frac{\epsilon}{2} \leq \psi_t^{\mathcal{M}}(\mathbf{X}, \tilde{\mathbf{X}}) \leq \eta^{\mathcal{M}}(\mathbf{f}(\mathbf{X}), \tilde{\mathbf{X}})$$

hence also $$|\psi_t^{\mathcal{M}}(\mathbf{X}, \tilde{\mathbf{X}}) - \psi^{\mathcal{M}}(\mathbf{X}, \tilde{\mathbf{X}})| < \epsilon$$

for $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ and $\tilde{\mathbf{X}} \in \prod_{j\in\mathbb{N}} (\tilde{D}_j)^{\mathcal{M}}$. Since ψ can be approximated in this way by definable predicates, it is a definable predicate, which proves the claim of the proposition. □

Corollary 3.18 *If $f_k : \prod_{j\in\mathbb{N}} S_j \to S'_k$ is a term for each $k \in \mathbb{N}$, then $\mathbf{f} = (f_j)_{j\in\mathbb{N}}$ is a definable function relative to any theory* T.

Proof By definition a term maps a product of domains of quantification into some domain of quantification, which verifies (1) of Definition 3.16. Moreover, for each k, $d(f_k(\mathbf{x}), y_k)$ is a formula, hence a definable predicate, so by the previous proposition $\mathbf{f}$ is a definable function. □

Similar to definable predicates, definable functions are automatically uniformly continuous with respect to d on each product of domains of quantification. This is a straightforward generalization of [3, Proposition 9.23].

Lemma 3.19 *Let $\mathcal{L}$ be a language,* T *an $\mathcal{L}$–theory, $\mathbf{S}$ and $\mathbf{S}'$ $\mathbb{N}$–tuples of sorts, and $\mathbf{f} : \prod_{j\in\mathbb{N}} S_j \to \prod_{j\in\mathbb{N}} S'_j$ a definable function. Then for every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$ and $F \subseteq \mathbb{N}$ finite and $\epsilon > 0$, there exists $E \subseteq \mathbb{N}$ finite and $\delta > 0$ such that*

$$d_{S_i}^{\mathcal{M}}(X_i, Y_i) < \delta \text{ for all } i \in E \implies d_{S'_j}^{\mathcal{M}}(f_j^{\mathcal{M}}(\mathbf{X}), f_j^{\mathcal{M}}(\mathbf{Y})) < \epsilon \text{ for all } j \in F$$

whenever $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X}, \mathbf{Y} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$.

Proof Let $\mathbf{D}' \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$ such that $\mathbf{f}^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ into $\prod_{j\in\mathbb{N}} (D'_j)^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$. Fix $\epsilon > 0$ and $F \subseteq \mathbb{N}$ finite. Then by Lemma 3.12 and Proposition 3.17, the object

$$\phi^{\mathcal{M}}(\mathbf{X}, \mathbf{Y}) = \max_{j\in F} d_{S'_j}^{\mathcal{M}}(f_j^{\mathcal{M}}(\mathbf{X}), Y_j)$$

is a definable predicate relative to T. Hence, by Observation 3.11 there exists $E \subseteq \mathbb{N}$ finite and $\delta > 0$ such that

$$d_{S_i}^{\mathcal{M}}(X_i', X_i) < \delta \text{ for } i \in E \text{ and } d_{S_j'}^{\mathcal{M}}(Y_j, Y_j') < \delta \text{ for } j \in F \implies |\phi^{\mathcal{M}}(\mathbf{X}, \mathbf{Y}) - \phi^{\mathcal{M}}(\mathbf{X}', \mathbf{Y}')| < \epsilon$$

for $\mathbf{X}, \mathbf{X}' \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ and $\mathbf{Y}, \mathbf{Y}' \in \prod_{j\in\mathcal{M}} (D_j')^{\mathcal{M}}$ and $\mathcal{M} \models \mathrm{T}$. Taking $\mathbf{Y} = \mathbf{f}(\mathbf{X})$, we see that

$$|\phi^{\mathcal{M}}(\mathbf{X}, \mathbf{f}^{\mathcal{M}}(\mathbf{X})) - \phi^{\mathcal{M}}(\mathbf{X}', \mathbf{f}^{\mathcal{M}}(\mathbf{X}))| = \max_{j\in F} d_{S_j'}^{\mathcal{M}}(f_j(\mathbf{X}'), f_j(\mathbf{X})) < \epsilon$$

whenever $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X}, \mathbf{X}' \in \prod_{j\in\mathbb{N}} D_j$ and $\max_{i\in E} d_{S_i}^{\mathcal{M}}(X_i, X_i') < \delta$, which is the desired uniform continuity property. □

Next, we describe the relationship between definable functions and types.

Lemma 3.20 *Let* $\mathbf{S}$ *and* $\mathbf{S}'$ *be* $\mathbb{N}$*–tuples of sorts, and let* $\mathbf{f} : \prod_{j\in\mathbb{N}} S_j \to \prod_{j\in\mathbb{N}} S_j'$ *be a definable function relative to* T.

(1) *If* ϕ *is a definable predicate in the variables* $x_j' \in S_j'$ *for* $j \in \mathbb{N}$*, then* $\phi \circ \mathbf{f}$ *is a definable predicate.*

(2) *If* $\mathcal{M} \models \mathrm{T}$ *and* $\mathbf{X} \in \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}}$*, then* $\mathrm{tp}^{\mathcal{M}}(\mathbf{f}^{\mathcal{M}}(\mathbf{X}))$ *is uniquely determined by* $\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$ *and* $\mathbf{f}$.

(3) *Let* $\mathbf{f}_* : \mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{S}_{\mathbf{S}'}(\mathrm{T})$ *be the map such that* $\mathrm{tp}^{\mathcal{M}}(\mathbf{f}^{\mathcal{M}}(\mathbf{X})) = \mathbf{f}_* \,\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$*. Then* $\mathbf{f}_*$ *is continuous with respect to the logic topology.*

Proof (1) Considering another $\mathbb{N}$–tuple $\tilde{\mathbf{S}}$ of sorts, we may view $\phi(\mathbf{x})$ as a definable predicate in $(\mathbf{x}, \tilde{\mathbf{x}})$, and hence by Definition 3.16, $\phi \circ \mathbf{f}$ is a definable predicate.

(2) For every definable predicate ϕ in $\mathbf{x}'$, $\phi \circ \mathbf{f}$ is a definable predicate, and hence $\phi^{\mathcal{M}} \circ \mathbf{f}^{\mathcal{M}}(\mathbf{X})$ is uniquely determined by $\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$ for all $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}}$. Since this is true for every definable predicate ϕ in variables $\mathbf{x}'$, it follows that $\mathrm{tp}^{\mathcal{M}}(\mathbf{f}(\mathbf{X}))$ is uniquely determined by $\mathbf{f}$ and $\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$.

(3) Let $\mathcal{O}$ be an open set in $\mathbb{S}_{\mathbf{S}'}(\mathrm{T})$. By definition of the topology on $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$, in order to show that $(\mathbf{f}_*)^{-1}(\mathcal{O})$ is open, is suffices to show that $(\mathbf{f}_*)^{-1}(\mathcal{O}) \cap \mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is open for every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$.

For any such $\mathbf{D}$, by Definition 3.16 (1), there exists $\mathbf{D}' \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j'}$ such that $\mathbf{f}^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ into $\prod_{j\in\mathbb{N}} (D_j')^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$. This implies that $\mathbf{f}_*$ maps $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ into $\mathbb{S}_{\mathbf{D}'}(\mathrm{T})$. Hence, $(\mathbf{f}_*)^{-1}(\mathcal{O}) \cap \mathbb{S}_{\mathbf{D}}(\mathrm{T}) = (\mathbf{f}_*|_{\mathbb{S}_{\mathbf{D}}(\mathrm{T})})^{-1}(\mathcal{O} \cap \mathbb{S}_{\mathbf{D}'}(\mathrm{T}))$, so to show that

this set is $(\mathbf{f}_*)^{-1}(\mathcal{O}) \cap \mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is open, it suffices to check continuity of $\mathbf{f}_*$ as a map $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{S}_{\mathbf{D}'}(\mathrm{T})$.

By Lemma 3.10, the topology on $\mathbb{S}_{\mathbf{D}'}(\mathrm{T})$ is generated by the pairings of types with every definable predicate ϕ in the variables $\mathbf{x}'$. If ϕ is such a definable predicate, then $\phi \circ \mathbf{f}$ is a definable predicate in $\mathbf{x}$ by (1), and therefore, it is continuous with respect to the logic topology on $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$. Thus, the map $\mathbf{f}_* : \mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{S}_{\mathbf{D}'}(\mathrm{T})$ is continuous as desired. □

Finally, we verify that definable functions are closed under composition.

Observation 3.21 *Let $\mathbf{S}$, $\mathbf{S}'$, and $\mathbf{S}''$ be $\mathbb{N}$–tuples of sorts in the language $\mathcal{L}$. If $\mathbf{f} : \prod_{j\in\mathbb{N}} S_j \to \prod_{j\in\mathbb{N}} S'_j$ and $\mathbf{g} : \prod_{j\in\mathbb{N}} S'_j \to \prod_{j\in\mathbb{N}} S''_j$ are definable functions, then so is $\mathbf{g} \circ \mathbf{f}$.*

Proof If $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$, then since $\mathbf{f}$ is definable, there exists $\mathbf{D}' \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S'_j}$ such that $\mathbf{f}^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ into $\prod_{j\in\mathbb{N}} (D'_j)^{\mathcal{M}}$ for every $\mathcal{M} \models \mathrm{T}$. Similarly, there exists $\mathbf{D}'' \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S''_j}$ such that $\mathbf{g}^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} (D'_j)^{\mathcal{M}}$ into $\prod_{j\in\mathbb{N}} (D''_j)^{\mathcal{M}}$. Hence, $(\mathbf{g}\circ\mathbf{f})^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ into $\prod_{j\in\mathbb{N}} (D''_j)^{\mathcal{M}}$, so that $\mathbf{g} \circ \mathbf{f}$ satisfies (1) of Definition 3.16.

Let $\tilde{\mathbf{S}}$ be another $\mathbb{N}$–tuple of sorts and let ϕ be a definable predicate in the variables $x''_j \in S''_j$ for $j \in \mathbb{N}$ and $\tilde{x}_j$ for $j \in \mathbb{N}$. By the definability of $\mathbf{g}$, $\psi(\mathbf{x}', \tilde{\mathbf{x}}) := \phi(\mathbf{g}(\mathbf{x}'), \tilde{\mathbf{x}})$ is also a definable predicate. Then by definability of $\mathbf{f}$, $\psi(\mathbf{f}(\mathbf{x}), \tilde{\mathbf{x}}) = \phi((\mathbf{g}\circ\mathbf{f})(\mathbf{x}), \tilde{\mathbf{x}})$ is a definable predicate. Therefore, $\mathbf{g} \circ \mathbf{f}$ satisfies (2) of Definition 3.16, so it is a definable function. □

3.4 Quantifier-free types and definable predicates

Quantifier-free formulas, that is, formulas defined without suprema or infima, are the simplest kind of formula and have special significance in our study of tracial W*–algebras.

Definition 3.22 *Quantifier-free formulas* are formulas obtained through the application of relations to terms and iterative application of continuous functions $\mathbb{R}^n \to \mathbb{R}$, that is, formulas obtained without using sup and inf operations. If $\mathbf{S}$ is an $\mathbb{N}$–tuple of sorts, we denote the set of quantifier-free formulas in variables $x_j \in S_j$ for $j \in \mathbb{N}$ by $\mathcal{F}_{\mathrm{qf},\mathbf{S}}$.

Quantifier-free types, the space of quantifier-free types, and quantifier-free definable predicates are defined in the same ways as the analogous objects for types, to wit:

Definition 3.23 Let $\mathbf{S}$ be an $\mathbb{N}$–tuple of sorts in the language $\mathcal{L}$. If $\mathcal{M}$ is an $\mathcal{L}$–structure, then the *quantifier-free type* of $\mathbf{X} \in \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}}$ is the map:

$$\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X})\colon\ \mathcal{F}_{\mathbf{S}} \to \mathbb{R}, \quad \phi \mapsto \phi^{\mathcal{M}}(\mathbf{X})$$

Definition 3.24

- If T is an $\mathcal{L}$–theory and $\mathbf{S}$ is an $\mathbb{N}$–tuple of sorts, then $\mathbb{S}_{\mathrm{qf},\mathbf{S}}(\mathrm{T})$ will denote the set of types $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X})$ for $\mathbf{X} \in \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}}$ and $\mathcal{M} \models \mathrm{T}$.
- If $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathbf{D}_{\mathbf{S}}$, then $\mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$ will denote the set of types $\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$ for $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$.
- We equip $\mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$ with the weak-$\star$ topology as a subset of the dual of $\mathcal{F}_{\mathrm{qf},\mathbf{S}}$.
- We equip $\mathbb{S}_{\mathrm{qf},\mathbf{S}}(\mathrm{T})$ with the topology such that $\mathcal{O}$ is open if and only if $\mathcal{O} \cap \mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$ is open for every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$. We call this the *(quantifier-free) logic topology*.

Definition 3.25 Let T be an $\mathcal{L}$–theory and $\mathbf{S}$ is an $\mathbb{N}$–tuple of sorts. A *quantifier-free definable predicate* is collection of functions $\phi^{\mathcal{M}} : \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}} \to \mathbb{R}$ for $\mathcal{M} \models \mathrm{T}$ such that for every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$ and $\epsilon > 0$, there exists a quantifier-free formula ψ in finitely many of the variables $x_j \in S_j$, such that

$$|\phi^{\mathcal{M}}(\mathbf{X}) - \psi^{\mathcal{M}}(\mathbf{X})| < \epsilon$$

for $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ for $\mathcal{M} \models \mathrm{T}$.

The following can be verified in the same way as for types, when $\mathbf{S}$ is an $\mathbb{N}$–tuple of sorts and T is an $\mathcal{L}$–theory:

- For each $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$, the space $\mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$ is a compact Hausdorff space.
- ϕ is a quantifier-free definable predicate if and only if $\phi^{\mathcal{M}}(\mathbf{X}) = \omega(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$ for some continuous $\omega : \mathbb{S}_{\mathrm{qf},\mathbf{S}}(\mathrm{T}) \to \mathbb{R}$.
- If ϕ_j is a quantifier-free definable predicate for $j \in \mathbb{N}$ and $F : \mathbb{R}^{\mathbb{N}} \to \mathbb{R}$ is continuous, then $F((\phi_j)_{j\in\mathbb{N}})$ is a quantifier-free definable predicate.

Furthermore, the quantifier-free type space and the type space can be related as follows.

Observation 3.26 *Let $\mathbf{S}$ be an $\mathbb{N}$–tuple of sorts in $\mathcal{L}$ and* T *an $\mathcal{L}$–theory. Let $\pi : \mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{S}_{\mathrm{qf},\mathbf{S}}(\mathrm{T})$ be the map that sends a type (as a linear map $\mathcal{F}_{\mathbf{S}} \to \mathbb{R}$) to its restriction to $\mathcal{F}_{\mathrm{qf},\mathbf{S}}$. Then $\pi(\mathbb{S}_{\mathbf{D}}(\mathrm{T})) = \mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$ for each $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$, and π is a topological quotient map.*

Proof First, π is a continuous map $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$ by definition of the weak-$\star$ topology. Then since a set in $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$ is open if and only if its restriction to $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$

is open, and the same holds for the quantifier-free versions, we deduce that π is continuous. It is immediate from the definitions that $\pi(\mathbb{S}_{\mathbf{D}}(\mathrm{T})) = \mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$. Then because $\mathbb{S}_{\mathbf{D}}(\mathrm{T})$ is compact and $\mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$ is Hausdorff, π defines a topological quotient map $\mathbb{S}_{\mathbf{D}}(\mathrm{T}) \to \mathbb{S}_{\mathrm{qf},\mathbf{D}}(\mathrm{T})$. Finally, using the definition of open sets in $\mathbb{S}_{\mathbf{S}}(\mathrm{T})$ and $\mathbb{S}_{\mathrm{qf},\mathbf{S}}(\mathrm{T})$, we deduce that $\mathcal{O} \subseteq \mathbb{S}_{\mathrm{qf},\mathbf{S}}(\mathrm{T})$ is open if and only if $\pi^{-1}(\mathcal{O})$ is open, hence π is a topological quotient map. □

Remark 3.27 A convenient feature of $\mathcal{L}_{\mathrm{tr}}$ is that $\pi^{-1}(\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})) = \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$. Indeed, suppose that $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathbf{X} \in L^\infty(\mathcal{M})^{\mathbb{N}}$ with $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) \in \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$. Then the operator norm of X_j can be recovered from $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X})$ through

$$\|X_j\| = \lim_{k\to\infty} (\mathrm{Re}\,\mathrm{tr}^{\mathcal{M}}((X_j^* X_j)^k))^{1/2k}$$

hence $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}$, so that $\mathrm{tp}^{\mathcal{M}}(\mathbf{X}) \in \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$.

3.5 Quantifier-free definable functions in $\mathcal{L}_{\mathrm{tr}}$

Quantifier-free definable functions are defined analogously to definable functions.

Definition 3.28 T be an $\mathcal{L}$–theory, and let $\mathbf{S}$ and $\mathbf{S}'$ be $\mathbb{N}$–tuples of sorts. A *quantifier-free definable function* $\mathbf{f} : \prod_{j\in\mathbb{N}} S_j \to \prod_{j\in\mathbb{N}} S_j'$ is a collection of functions $\mathbf{f}^{\mathcal{M}} : \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}} \to \prod_{j\in\mathbb{N}} (S_j')^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$ satisfying the following conditions:

(1) For each $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$, there exists $\mathbf{D}' \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j'}$ such that for every $\mathcal{M} \models \mathrm{T}$, $\mathbf{f}^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ into $\prod_{j\in\mathbb{N}} (D_j')^{\mathcal{M}}$.

(2) Whenever $\tilde{\mathbf{S}}$ is another tuple of sorts and ϕ is a quantifier-free definable predicate relative to T in the free variables $x_j' \in S_j'$ and $\tilde{x}_j \in \tilde{S}_j$ for $j \in \mathbb{N}$, then $\phi(\mathbf{f}(\mathbf{x}), \tilde{\mathbf{x}})$ is a quantifier-free definable predicate in the variables $\mathbf{x} = (x_j)_{j\in\mathbb{N}}$ and $\tilde{\mathbf{x}} = (\tilde{x}_j)_{j\in\mathbb{N}}$.

Example 3.29 If f_j is a term in a finite subset of the variables x_j, then $\mathbf{f} = (f_j)_{j\in\mathbb{N}}$ is a quantifier-free definable function relative to any $\mathcal{L}$–theory T. To see this, suppose that ϕ is a quantifier-free definable predicate and $\mathbf{D}$ is a tuple of domains of quantification. Let $\mathbf{f}$ map $\mathbf{D}$ into $\mathbf{D}'$. As a quantifier-free definable predicate, ϕ can, for any given $\epsilon > 0$, be approximated on $\mathbf{D}'$ by a quantifier-free formula ψ with error smaller than ϵ on $\prod_{j\in\mathbb{N}} (D_j')^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$. Then $\psi \circ \mathbf{f}$ is a quantifier-free formula that approximates $\phi \circ \mathbf{f}$ within ϵ on $\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$.

The following facts about quantifier-free definable functions are verified just as in the case of definable functions:

- Quantifier-free definable functions are closed under composition.
- If ϕ is a quantifier-free definable predicate and $\mathbf{f}$ is a quantifier-free definable function, then $\phi \circ \mathbf{f}$ is a quantifier-free definable predicate.
- For each definable function $\mathbf{f} : \prod_{j\in\mathbb{N}} S_j \to \prod_{j\in\mathbb{N}} S'_j$, there is a continuous map $\mathbf{f}_* : \mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{S}_{\mathbf{S}'}(\mathrm{T})$ given by $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{f}(\mathbf{X})) = \mathbf{f}_* \mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X})$.
- If $\mathbf{f}$ is a quantifier-free definable function, then for each $k \in \mathbb{N}$, the object $\phi(\mathbf{x}, y) = d_{S'_k}(f_k(\mathbf{x}), y)$ is a quantifier-free definable predicate.
- Hence, $\mathbf{f}$ is a quantifier-free definable function, then it is a definable function by Proposition 3.17.
- Thus, a quantifier-free definable function is uniformly continuous on every product of domains of quantification.

The proof of ($\Longleftarrow$) in Proposition 3.17 uses quantifiers (specifically infima) and thus does not directly adapt to the quantifier-free setting. This is why we argued that terms are quantifier-free functions directly in Example 3.29 rather than as in Corollary 3.18.

In the special case of $\mathcal{L}_{\mathrm{tr}}$ and T_{tr}, we have the following characterizations of quantifier-free definable functions. Recall that $\mathcal{L}_{\mathrm{tr}}$ has one type S and the domains of quantification are given by D_r for each $r > 0$. Given $\mathbf{r} = (r_j)_{j\in\mathbb{N}} \in (0,\infty)^{\mathbb{N}}$, we write $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ for the set of quantifier-free types of $\mathbb{N}$–tuples in $\prod_{j\in\mathbb{N}} D_{r_j}$ in $\mathcal{L}_{\mathrm{tr}}$ relative to T_{tr}. A variant of this theorem was proved in the author's Ph.D. thesis [22, Proposition 13.6.4].

Theorem 3.30 *Let* $\mathbf{f}$ *be a collection of functions* $\mathbf{f}^{\mathcal{M}} : L^\infty(\mathcal{M})^{\mathbb{N}} \to L^\infty(\mathcal{M})^{\mathbb{N}}$ *for each* $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$. *Suppose that for every* $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, *there exists* $\mathbf{r}' \in (0,\infty)^{\mathbb{N}}$ *such that* $\mathbf{f}^{\mathcal{M}}$ *maps* $\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r_j}$ *into* $\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r'_j}$; *assume that for each* $\mathbf{r}$ *a corresponding* $\mathbf{r}'$ *has been chosen, which we will refer to below. Then the following are equivalent.*

(1) $\mathbf{f}$ *is a quantifier-free definable function in* $\mathcal{L}_{\mathrm{tr}}$ *relative to* T_{tr}.

(2) *For each* $k \in \mathbb{N}$, *the object* $\phi(\mathbf{x}, y) = d_{S'_k}(f_k(\mathbf{x}), y)$ *is a quantifier-free definable predicate.*

(3) *For each* $k \in \mathbb{N}$, $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, *and quantifier-free type* $\mu \in \mathbb{S}_{\mathrm{qf},\mathbf{D}_{\mathbf{r}}}(\mathrm{T}_{\mathrm{tr}})$ *and* $\epsilon > 0$, *there exists a term* g *and an open neighborhood* $\mathcal{O}$ *of* μ *in* $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ *such that, for all* $\mathcal{M} \models \mathrm{T}$ *and* $\mathbf{X} \in \prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r_j}$,

$$g^{\mathcal{M}}(\mathbf{X}) \in D_{r'_k}$$

and
$$\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}) \in \mathcal{O} \implies d^{\mathcal{M}}(f^{\mathcal{M}}_k(\mathbf{X}), g^{\mathcal{M}}(\mathbf{X})) < \epsilon.$$

(4) *For each* $k \in \mathbb{N}$, $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, *and* $\epsilon > 0$, *there exist* $m \in \mathbb{N}$, *quantifier-free formulas* $\phi_1, \ldots, \phi_m$, *and terms* $g_1, \ldots, g_m$ *such that*

$$\sum_{j=1}^{m} \phi^{\mathcal{M}}_j(\mathbf{X}) g^{\mathcal{M}}_j(\mathbf{X}) \in D_{r'_k} \text{ and } d^{\mathcal{M}}\left(f^{\mathcal{M}}_k(\mathbf{X}), \sum_{j=1}^{m} \phi^{\mathcal{M}}_j(\mathbf{X}) g^{\mathcal{M}}_j(\mathbf{X})\right) < \epsilon$$

whenever $\mathcal{M} \models \mathrm{T}$ *and* $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}$.

Proof (1) $\implies$ (2) follows as in Proposition 3.17.

(2) $\implies$ (3). Fix $\mathbf{r}$ and μ and ϵ. Let $\mathbf{X}$ be an $\mathbb{N}$–tuple from a tracial W*–algebra $\mathcal{M}$ which has quantifier-free type μ. Note that if ϕ is a quantifier-free formula and $\mathcal{N}$ is a W*–subalgebra of $\mathcal{M}$ containing $\mathbf{X}$, then $\phi^{\mathcal{N}}(\mathbf{X}) = \phi^{\mathcal{M}}(\mathbf{X})$. Hence, the quantifier-free type of $\mathbf{X}$ in $\mathcal{N}$ is the same as the quantifier-free type in $\mathcal{M}$. In particular, we can replace $\mathcal{M}$ with $\mathcal{N} = \mathrm{W}^*(\mathbf{X})$ and thus assume without loss of generality that $\mathcal{M} = \mathrm{W}^*(\mathbf{X})$.

Recall that $\mathcal{M}$ is faithfully represented on the GNS space $L^2(\mathcal{M})$ (see § 2.1, the standard representation). By the Kaplansky density theorem [27] (also explained in Kadison and Ringrose [26, Theorem 5.3.5]), the ball of radius r_k' in the C*–algebra $\mathrm{C}^*(\mathbf{X})$ generated by $\mathbf{X}$ is dense in the ball of radius r_k' in $\mathcal{M} = \mathrm{W}^*(\mathbf{X})$ with respect to the strong operator topology. Since approximation in the strong operator topology implies approximation in the 2–norm associated to the trace, it follows that there exists $Z \in \mathrm{C}^*(\mathbf{X})$ such that $\|Z\| \leq r_k'$ and $\|Z - Y\|_2 = d^{\mathcal{M}}(Z, Y) < \epsilon/2$.

Next, we must obtain a term g bounded by r_k' such that $d^{\mathcal{M}}(g^{\mathcal{M}}(\mathbf{X}), Y) < \epsilon$. Because we want g to be bounded by r_k' on $\prod_{j\in\mathbb{N}} D_{r_j}$ for all $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$, we view the $*$–polynomials in infinitely many indeterminates as part of a universal C*–algebra. For $*$–polynomials p in infinitely many variables $x_j : j \in \mathbb{N}$, let:

$$\|p\|_u = \sup\left\{\|p(\mathbf{X})\| : \mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}, \mathcal{M} \models \mathrm{T}_{\mathrm{tr}}\right\}$$

This defines a C*–norm on $\mathbb{C}\langle x_j, x_j^* : j \in \mathbb{N}\rangle$. Let A be the completion of $\mathbb{C}\langle x_j, x_j^* : j \in \mathbb{N}\rangle$ into a C*–algebra. If $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}$, then $\|p(\mathbf{X})\| \leq \|p\|_u$ by definition, so there is a $*$–homomorphism $\pi : A \to \mathrm{C}^*(\mathbf{X})$ mapping $x_j \in A$ to $X_j \in \mathcal{M}$ for each $j \in \mathbb{N}$. By Blackadar [5, II.5.1.5], there exists $z \in A$ such that $\pi(z) = Z$ and $\|z\|_A \leq r_k'$.

Now by definition $\mathbb{C}\langle x_j, x_j^* : j \in \mathbb{N}\rangle$ is dense in A. It follows that every element of the r_k'–ball of A can be approximated by non-commutative $*$–polynomials in the r_k'–ball. In particular, there exists some $g \in \mathbb{C}\langle x_j, x_j^* : j \in \mathbb{N}\rangle$ with $\|g - z\|_A < \epsilon/2$, and we can also arrange that $\|g\|_A \leq r_k'$.

Then g is a term such that $\|g^{\mathcal{N}}(\mathbf{Y})\| \leq r_k'$ for all $\mathbf{Y} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{N}}$ for all $\mathcal{N} \models \mathrm{T}_{\mathrm{tr}}$ and such that $d^{\mathcal{M}}(f_k^{\mathcal{M}}(\mathbf{X}), g^{\mathcal{M}}(\mathbf{X})) < \epsilon$ for our particular choice of $\mathcal{M} = \mathrm{W}^*(\mathbf{X})$ with $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) = \mu$. Now observe that $\psi(\mathbf{x})$ is a quantifier-free definable predicate since the term g is a quantifier-free definable function. Let $\mathcal{O} = \{\nu \in \mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) : \nu(\psi) < \epsilon\}$.

Then $\mathcal{O} \ni \mu$ by our choice of g and $\mathcal{O}$ is open by continuity of ψ. Moreover, by definition, if $\mathcal{N} \models \mathrm{T}$ and $\mathbf{Y} \in \prod_{j\in\mathbb{N}} D_{r_j}$ and $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{N}}(\mathbf{Y}) \in \mathcal{O}$, then $d^{\mathcal{N}}(f_k^{\mathcal{N}}(\mathbf{Y}), g^{\mathcal{N}}(\mathbf{Y})) < \epsilon$.

(3) $\implies$ (4). Fix $k \in \mathbb{N}$, $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, and $\delta > 0$. For each $\mu \in \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$, there exists an open $\mathcal{O}_\mu \subseteq \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ and a term g_μ such that for $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}$, $g_\mu^{\mathcal{M}}(\mathbf{X}) \in D_{r_k'}$ and

$$\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) \in \mathcal{O} \implies d^{\mathcal{M}}(g_\mu^{\mathcal{M}}(\mathbf{X}), f_k^{\mathcal{M}}(\mathbf{X})) < \delta.$$

The sets $\mathcal{O}_\mu$ form an open cover of the compact set $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$, and hence there exists $m \in \mathbb{N}$ and $\mu_1, \ldots, \mu_m$ such that $\mathcal{O}_{\mu_1}, \ldots, \mathcal{O}_{\mu_m}$ cover $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$. Let $\mathcal{O}_j = \mathcal{O}_{\mu_j}$ and $g_j = g_{\mu_j}$.

Since $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ is a compact Hausdorff space, there exists a continuous partition of unity $\psi_1, \ldots, \psi_m$ subordinated to the cover $\mathcal{O}_1, \ldots, \mathcal{O}_m$. In other words, there exist $\gamma_1, \ldots, \gamma_m \in C(\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}))$ such that $\gamma_j \geq 0$, $\gamma_j|_{\mathcal{O}_j^c} = 0$, and $\sum_{j=1}^m \gamma_j = 1$. Therefore, for $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}$,

$$\begin{aligned} d^{\mathcal{M}}\left(\sum_{j=1}^m \gamma_j(\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}))g_j^{\mathcal{M}}(\mathbf{X}), f_k^{\mathcal{M}}(\mathbf{X})\right) &= \left\|\sum_{j=1}^m \gamma_j(\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}))(g_j^{\mathcal{M}}(\mathbf{X}) - f_k^{\mathcal{M}}(\mathbf{X}))\right\|_{L^2(\mathcal{M})} \\ &\leq \sum_{j=1}^m \gamma_j(\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}))\left\|g_j^{\mathcal{M}}(\mathbf{X}) - f_k^{\mathcal{M}}(\mathbf{X})\right\|_{L^2(\mathcal{M})} \\ &\leq \sum_{j=1}^m \gamma_j(\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}))\delta \\ &= \delta \end{aligned}$$

and $$\left\|\sum_{j=1}^m \gamma_j(\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}))g_j^{\mathcal{M}}(\mathbf{X})\right\|_{L^\infty(\mathcal{M})} \leq \sum_{j=1}^m \gamma_j(\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}))\|g_j^{\mathcal{M}}(\mathbf{X})\|_{L^\infty(\mathcal{M})} \leq r_k'.$$

Because quantifier-free formulas comprise a dense subset of $C(\mathbb{S}_{\mathrm{qf}}(\mathrm{T}))$ by the Stone–Weierstrass Theorem, there exist quantifier-free formulas $\psi_1, \ldots, \psi_m$ such that

$$|\psi_j^{\mathcal{M}}(\mathbf{X}) - \gamma_j(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))| \leq \frac{\delta}{m} \text{ for } \mathcal{M} \models \mathrm{T}_{\mathrm{tr}} \text{ and } \mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}.$$

It follows that

$$\left\|\sum_{j=1}^{m}\psi_j^{\mathcal{M}}(\mathbf{X})g_j^{\mathcal{M}}(\mathbf{X})\right\|_{L^\infty(\mathcal{M})}\leq r_k'\sum_{j=1}^{m}|\psi_j^{\mathcal{M}}(\mathbf{X})|\leq r_k'(1+\delta)$$

and
$$\left\|\sum_{j=1}^{m}\gamma_j(\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}))g_j^{\mathcal{M}}(\mathbf{X})-\sum_{j=1}^{m}\psi_j^{\mathcal{M}}(\mathbf{X})g_j^{\mathcal{M}}(\mathbf{X})\right\|_{L^2(\mathcal{M})}\leq r_k'\delta.$$

Therefore, let $\phi_j=(1+\delta)^{-1}\psi_j$. Then

$$\left\|\sum_{j=1}^{m}\phi_j^{\mathcal{M}}(\mathbf{X})g_j^{\mathcal{M}}(\mathbf{X})\right\|_{L^\infty(\mathcal{M})}\leq(1+\delta)^{-1}r_k'\sum_{j=1}^{m}|\psi_j^{\mathcal{M}}(\mathbf{X})|\leq r_k'$$

and
$$\left\|\sum_{j=1}^{m}\phi_j^{\mathcal{M}}(\mathbf{X})g_j^{\mathcal{M}}(\mathbf{X})-\sum_{j=1}^{m}\psi_j^{\mathcal{M}}(\mathbf{X})g_j^{\mathcal{M}}(\mathbf{X})\right\|_{L^2(\mathcal{M})}\leq r_k'(1-(1+\delta)^{-1})=\frac{\delta r_k'}{1+\delta}.$$

Hence, by the triangle inequality:

$$\left\|\sum_{j=1}^{m}\phi_j^{\mathcal{M}}(\mathbf{X})g_j^{\mathcal{M}}(\mathbf{X})-f^{\mathcal{M}}(\mathbf{X})\right\|_{L^2(\mathcal{M})}\leq\delta+\delta r_k'+\frac{\delta r_k'}{1+\delta}$$

By choosing δ sufficiently small, we can guarantee that the right-hand side is smaller than a given $\epsilon>0$, so the quantifier-free formulas $\phi_1,\ldots,\phi_m$ have the desired properties for (4).

(4) $\implies$ (2). Fix k and we will show that $\phi(\mathbf{x},y)=d(f_k(\mathbf{x}),y)$ is a quantifier-free definable predicate. To this end, fix $\mathbf{r}\in(0,\infty)^{\mathbb{N}}$, $r'>0$, and $\epsilon>0$, and we will approximate ϕ by a quantifier-free formula on $\prod_{j\in\mathbb{N}}D_j^{\mathcal{M}}\times D_{r'}$ within ϵ for $\mathcal{M}\models\mathrm{T}_{\mathrm{tr}}$. Let $m\in\mathbb{N}$ and $\phi_1,\ldots,\phi_m$ and $g_1,\ldots,g_m$ be as in (4) for our given $\mathbf{r}$ and ϵ, and let

$$h^{\mathcal{M}}(\mathbf{X})=\sum_{j=1}^{m}\phi_j^{\mathcal{M}}(\mathbf{X})g_j^{\mathcal{M}}(\mathbf{X}).$$

Then for $\mathcal{M}\models\mathrm{T}_{\mathrm{tr}}$ and $\mathbf{X}\in\prod_{j\in\mathbb{N}}D_{r_j}^{\mathcal{M}}$ and $Y\in D_{r'}$, we have

$$|d^{\mathcal{M}}(f_k^{\mathcal{M}}(\mathbf{X}),Y)-d^{\mathcal{M}}(h^{\mathcal{M}}(\mathbf{X}),Y)|\leq d^{\mathcal{M}}(f_k^{\mathcal{M}}(\mathbf{X}),h^{\mathcal{M}}(\mathbf{X}))<\epsilon.$$

But note that

$$d^{\mathcal{M}}(h^{\mathcal{M}}(\mathbf{X}), Y) = \left\| \sum_{j=1} \phi_j^{\mathcal{M}}(\mathbf{X}) g_j^{\mathcal{M}}(\mathbf{X}) - Y \right\|_{L^2(\mathcal{M})}$$
$$= \Bigg(\sum_{j,k=1}^{m} \phi_j^{\mathcal{M}}(\mathbf{X}) \phi_k^{\mathcal{M}}(\mathbf{X}) \operatorname{Re} \operatorname{tr}^{\mathcal{M}}(g_j^{\mathcal{M}}(\mathbf{X})^* g_k^{\mathcal{M}}(\mathbf{X}))$$
$$- 2 \sum_{j=1}^{m} \phi_j^{\mathcal{M}}(\mathbf{X}) \operatorname{Re} \operatorname{tr}^{\mathcal{M}}(g_j^{\mathcal{M}} Y) + \operatorname{Re} \operatorname{tr}^{\mathcal{M}}(Y^* Y) \Bigg)^{1/2}$$

which is a quantifier-free formula by inspection.

(3) $\implies$ (1). Let ϕ be a quantifier-free definable predicate, and we will show that $\phi \circ \mathbf{f}$ defines a continuous function on $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ for each $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$ for each $\mathbf{r}$, and hence $\phi \circ \mathbf{f}$ is a quantifier-free definable predicate. To this end, it suffices to show that for each $\mathbf{r} \in (0,\infty)$, each $\mu \in \mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$, and each $\epsilon > 0$, there exists an neighborhood $\mathcal{O}$ of μ in $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ such that $|\phi^{\mathcal{M}} \circ \mathbf{f}^{\mathcal{M}}(\mathbf{X}) - \mu(\phi \circ \mathbf{f})| < \epsilon$ whenever $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) \in \mathcal{O}$.

Fix $\mathbf{r}$, μ, and ϵ. Let $\mathbf{r}' \in (0,\infty)^{\mathbb{N}}$ be such that $\mathbf{f}^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}$ into $\prod_{j\in\mathbb{N}} D_{r_j'}^{\mathcal{M}}$ for $\mathcal{M} \models \mathrm{T}$. By the uniform continuity property of definable predicates, there exists $F \subseteq \mathbb{N}$ finite and $\delta > 0$ such that:

$$\mathbf{Y}, \mathbf{Y}' \in \prod_{j\in\mathbb{N}} D_{r_j'}^{\mathcal{M}} \text{ and } \max_{k\in F} d^{\mathcal{M}}(Y_k, Y_k') < \delta \implies |\phi^{\mathcal{M}}(\mathbf{Y}) - \phi^{\mathcal{M}}(\mathbf{Y}')| < \frac{\epsilon}{3}$$

By (3), for each $k \in F$, choose a term g_k and open $\mathcal{O}_k \subseteq \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ such that, for all $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}$, $g_k^{\mathcal{M}}(\mathbf{X}) \in D_{r_k'}$ and $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) \in \mathcal{O}_k \implies d^{\mathcal{M}}(f_k^{\mathcal{M}}(\mathbf{X}), g^{\mathcal{M}}(\mathbf{X})) < \delta$.

For $k \notin F$, let $g_k = 0$. Then, by our choice of δ and $g_k^{\mathcal{M}}$, for all $\mathcal{M} \models \mathrm{T}$ and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}$:

$$\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) \in \bigcap_{k\in F} \mathcal{O}_k \implies |\phi^{\mathcal{M}}(\mathbf{f}^{\mathcal{M}}(\mathbf{X})) - \phi^{\mathcal{M}}(\mathbf{g}^{\mathcal{M}}(\mathbf{X}))| < \frac{\epsilon}{3}$$

Moreover, $\mathbf{g} = (g_k)_{k\in\mathbb{N}}$ is an $\mathbb{N}$–tuple of terms, hence $\mathbf{g}$ is a quantifier-free definable function. This implies that $\phi \circ \mathbf{g}$ is a quantifier-free definable predicate. This implies that $\mathcal{O}' := \{\nu \in \mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) : |\nu(\phi \circ \mathbf{g}) - \mu(\phi \circ \mathbf{g})| < \epsilon/3\}$ is open in $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$. Let:

$$\mathcal{O} := \mathcal{O}' \cap \bigcap_{k=1}^{m} \mathcal{O}_k$$

Then $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) \in \mathcal{O}$ implies that

$$\begin{aligned}|\phi^{\mathcal{M}} \circ \mathbf{f}^{\mathcal{M}}(\mathbf{X}) - \mu(\phi \circ \mathbf{f})| &\leq |\phi^{\mathcal{M}} \circ \mathbf{f}^{\mathcal{M}}(\mathbf{X}) - \phi^{\mathcal{M}} \circ \mathbf{g}^{\mathcal{M}}(\mathbf{X})| + |\phi^{\mathcal{M}} \circ \mathbf{g}^{\mathcal{M}}(\mathbf{X}) - \mu(\phi \circ \mathbf{g})| \\ &\quad + |\mu(\phi \circ \mathbf{g}) - \mu(\phi \circ \mathbf{f})| \\ &< \epsilon\end{aligned}$$

as desired. □

Example 3.31 Suppose $\rho_j \in C(\mathbb{R})$, and let $f_j^{\mathcal{M}}(\mathbf{X}) = \rho_j(\frac{1}{2}(X_j + X_j^*))$ for $\mathcal{M} \models \mathrm{T}$, where $\rho_j(\frac{1}{2}(X_j + X_j^*))$ is defined through functional calculus. By approximating ρ_j uniformly on $[-r_j, r_j]$ by a polynomial g_j and applying the spectral theorem, we can verify Theorem 3.30 (3) and hence conclude that $\mathbf{f}$ is a quantifier-free definable function relative to T_{tr}. Similarly, if $f_j(\mathbf{X}) = X_j \rho_j(X_j^* X_j)$, then $\mathbf{f}$ is a quantifier-free definable function, an observation that we will use in the proof of Proposition 3.32. In this way, continuous functional calculus fits into the larger model-theoretic framework of quantifier-free definable functions.

Building on Theorem 3.30 and Example 3.31, we can show that every element of the W*–algebra can be realized as a quantifier-free definable function applied to the generators. This fact will be use later on to show that covering entropy remains invariant under change of generators for a tracial W*–algebra. This is a version of [22, Proposition 13.6.6] and [18, Proposition 2.4], and the idea behind the proof is a "forced limit" construction (see Ben Yaacov and Usvyatsov [4, §3.2] or Ben Yaacov, Berenstein, Henson, and Usvyatsov [3, §9, definable predicates]) applied to quantifier-free definable functions rather than quantifier-free definable predicates.

Proposition 3.32 *If $\mathcal{M} = (M, \tau)$ is a tracial W*–algebra and $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}$ generates M and $\mathbf{Y} \in \prod_{j\in\mathbb{N}} D_{r_j'}^{\mathcal{M}}$, then there exists a quantifier-free definable function $\mathbf{f}$ in $\mathcal{L}_{\mathrm{tr}}$ relative to T_{tr} such that $\mathbf{Y} = \mathbf{f}(\mathbf{X})$. In fact, $\mathbf{f}$ can be chosen so that $f_k^{\mathcal{M}}$ maps $\prod_{j\in\mathbb{N}} L^\infty(\mathcal{M})$ into $\prod_{j\in\mathbb{N}} D_{r_j'}^{\mathcal{M}}$ for all $\mathcal{M} \models \mathrm{T}$.*

Proof Arguing as in (2) $\implies$ (3) of Theorem 3.30, for each $k \in \mathbb{N}$ and $m \in \mathbb{N}$ there exists a non-commutative polynomial $g_{k,m}$ such that $\|g_{k,m}^{\mathcal{N}}(\mathbf{X}')\|_{L^\infty(\mathcal{N})} \leq r_k'$ for $\mathcal{N} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathbf{X}' \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{N}}$, and:

$$d^{\mathcal{M}}(g_{k,m}^{\mathcal{M}}(\mathbf{X}), Y_k) < \frac{1}{2^{m+1}}$$

Hence also:

$$d^{\mathcal{M}}(g^{\mathcal{M}}_{k,m}(\mathbf{X}), g^{\mathcal{M}}_{k,m+1}(\mathbf{X})) < \frac{3}{2^{m+2}} < \frac{1}{2^m}$$

Let $F_m : \mathbb{R} \to \mathbb{R}$ be the continuous function such that $F_m(t) = 0$ for $t \leq 3/2^{m+2}$ and $F_m(t) = 1$ for $t \geq 1/2^m$ and F_m is affine on $[3/2^{m+2}, 1/2^m]$. Then

$$\phi_{k,m}(\mathbf{x}) = F_k(d(g_{k,m}(\mathbf{x}), g_{k,m+1}(\mathbf{x})))$$

is a quantifier-free formula. Moreover, by construction, $\phi^{\mathcal{M}}_{k,m}(\mathbf{X}) = 1$ for our given $\mathcal{M}$ and $\mathbf{X}$, while at the same time $\phi^{\mathcal{N}}_{k,m}(\mathbf{X}')$ is zero whenever $\|g^{\mathcal{N}}_{k,m}(\mathbf{X}') - g^{\mathcal{N}}_{k,m+1}(\mathbf{X}'))\|_{L^2(\mathcal{N})} > 1/2^m$ for any $\mathcal{N} \models \mathrm{T_{tr}}$ and $\mathbf{X}' \in \prod_{j\in\mathbb{N}} D^{\mathcal{N}}_{r_j}$. Let

$$\psi_{k,m} = \phi_{k,1}\phi_{k,2}\ldots\phi_{k,m}.$$

Then $\psi_{k,m}$ satisfies the same properties that we just showed for $\phi_{k,m}$ with the additional property that $\psi_{k,m+1} \leq \psi_{k,m}$.

For $\mathcal{N} \models \mathrm{T_{tr}}$ and $\mathbf{X}' \in L^\infty(\mathcal{N})^{\mathbb{N}}$, define:

$$\begin{aligned} f^{\mathcal{N}}_{k,m}(\mathbf{X}') &:= g^{\mathcal{N}}_{k,1}(\mathbf{X}') + \sum_{j=1}^{m-1} \psi^{\mathcal{N}}_{k,j}(\mathbf{X}')(g^{\mathcal{N}}_{k,j+1}(\mathbf{X}') - g^{\mathcal{N}}_{k,j}(\mathbf{X}')) \\ &= (1-\psi_{k,1})^{\mathcal{N}}(\mathbf{X}')g^{\mathcal{N}}_{k,1}(\mathbf{X}') + \sum_{j=2}^{m-1}(\psi_{k,j-1} - \psi_{k,j})^{\mathcal{N}}(\mathbf{X}')g^{\mathcal{N}}_{k,j}(\mathbf{X}') \\ &\quad + \psi^{\mathcal{N}}_{k,m-1}(\mathbf{X}')g^{\mathcal{N}}_{k,m}(\mathbf{X}') \end{aligned}$$

Then $\mathbf{f}_{\cdot,m} = (f_{k,m})_{k\in\mathbb{N}}$ is a quantifier-free definable function by Theorem 3.30 since it is equal to a finite sum of quantifier-free formulas multiplied by terms. Observe that for $\mathcal{N} \models \mathrm{T_{tr}}$ and $\mathbf{X}' \in \prod_{j\in\mathbb{N}} D^{\mathcal{N}}_{r_j}$,

$$\|f^{\mathcal{N}}_{k,m}(\mathbf{X}')\|_{L^\infty(\mathcal{N})} \leq (1-\psi_{k,1})^{\mathcal{N}}(\mathbf{X}')r'_k + \sum_{j=2}^{m-1}(\psi_{k,j-1} - \psi_{k,j})^{\mathcal{N}}(\mathbf{X}')r'_k + \psi^{\mathcal{N}}_{k,m-1}(\mathbf{X}')r'_k = r'_k$$

relying on the fact that $\psi_{k,j} \leq \psi_{k,j-1}$. Furthermore, for $\mathbf{X}' \in \prod_{j\in\mathbb{N}} D^{\mathcal{N}}_{r_j}$, we have:

$$\|f^{\mathcal{N}}_{k,m}(\mathbf{X}') - f^{\mathcal{N}}_{k,m+1}(\mathbf{X}')\|_{L^2(\mathcal{N})} = \psi^{\mathcal{N}}_{k,m}(\mathbf{X}')\|g^{\mathcal{N}}_{k,m}(\mathbf{X}') - g^{\mathcal{N}}_{k,m+1}(\mathbf{X}')\|_{L^2(\mathcal{N})} \leq \frac{1}{2^m}$$

This implies that for $\mathbf{X}' \in \prod_{j\in\mathbb{N}} D^{\mathcal{N}}_{r_j}$, the sequence $f^{\mathcal{N}}_{k,m}(\mathbf{X}')$ converges as $m \to \infty$ to some $f^{\mathcal{N}}_k(\mathbf{X}')$ with:

$$\|f^{\mathcal{N}}_k(\mathbf{X}') - f^{\mathcal{N}}_{k,m}(\mathbf{X}')\|_{L^2(\mathcal{N})} \leq \frac{1}{2^{m-1}}$$

Furthermore, for our given $\mathcal{M}$ and $\mathbf{X}$, we have

$$f_k^{\mathcal{M}}(\mathbf{X}) = g_{k,1}^{\mathcal{M}}(\mathbf{X}) + \sum_{j=1}^{\infty}(g_{k+1,j}^{\mathcal{M}}(\mathbf{X}) - g_{k,j}^{\mathcal{M}}(\mathbf{X})) = Y_k$$

because we assumed that $g_{k,j}^{\mathcal{M}}(\mathbf{X}) \to Y_k$ as $j \to \infty$.

Now this f_k is only well-defined a priori on $\prod_{j\in\mathbb{N}} D_{r_j}$ for our fixed choice of $\mathbf{r}$. In order to extend it to a global function, we use a cut-off trick based on Example 3.31. Let $\rho_j : \mathbb{R} \to \mathbb{R}$ be the function:

$$\rho_j(t) = \begin{cases} 1, & t \leq r_j^2, \\ r_j t^{-1/2}, & t \geq r_j^2 \end{cases}$$

Let $\mathbf{h}$ be given by $h_j^{\mathcal{N}}(\mathbf{X}') = X_j'\rho_j((X_j')^*X_j)$. Then $\mathbf{h}$ is a quantifier-free definable function relative to $\mathrm{T_{tr}}$ by Example 3.31. Moreover, if $\mathcal{N} \models \mathrm{T_{tr}}$ and $\mathbf{X}' \in$, we have

$$\begin{aligned}\|h_j^{\mathcal{N}}(\mathbf{X}')\|_{L^\infty(\mathcal{N})}^2 &= \|h_j^{\mathcal{N}}(\mathbf{X}')^* h_j^{\mathcal{N}}(\mathbf{X}')\|_{L^\infty(\mathcal{N})} \\ &= \|\rho_j((X_j')^*X_j'))(X_j')^*X_j'\rho_j((X_j')^*X_j')\|_{L^\infty(\mathcal{N})} \leq r_j^2\end{aligned}$$

since $\rho_j(t)^2 t \leq r_j^2$. Therefore, $\mathbf{h}^{\mathcal{N}}$ maps $L^\infty(\mathcal{N})^{\mathbb{N}}$ into $\prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{N}}$ for all $\mathcal{N} \models \mathrm{T_{tr}}$. Also, $\mathbf{h}^{\mathcal{N}}(\mathbf{X}') = \mathbf{X}'$ for $\mathbf{X}' \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{N}}$.

Now $\mathbf{f}_{,m} \circ \mathbf{h}$ is a quantifier-free definable function since it is a composition of quantifier-free definable functions. Because $f_{k,m}$ converges to f_k uniformly on $\prod_{j\in\mathbb{N}} D_{r_j}$ as $m \to \infty$, we see that $f_{k,m} \circ \mathbf{h}$ converges uniformly to $f_k \circ \mathbf{h}$ globally as $m \to \infty$. This implies that $\mathbf{f} \circ \mathbf{h}$ is a quantifier-free definable function because quantifier-free functions are closed under limits that are uniform on each product of domains (for instance, using Theorem 3.30 (3) or (4)). Moreover, $\mathbf{f}^{\mathcal{M}}(\mathbf{h}^{\mathcal{M}}(\mathbf{X})) = \mathbf{f}^{\mathcal{M}}(\mathbf{X}) = \mathbf{Y}$ by construction. Finally, $(\mathbf{f} \circ \mathbf{h})^{\mathcal{N}}$ maps into $\prod_{j\in\mathbb{N}} D_{r_j'}^{\mathcal{N}}$ for all $\mathcal{N} \models \mathrm{T_{tr}}$ since $(f_{k,m} \circ \mathbf{h})^{\mathcal{N}}$ maps $\prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{N}}$ into $D_{r_k'}^{\mathcal{N}}$. □

Remark 3.33 We can also deduce from the proof that every continuous function γ on $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T_{tr}})$ extends to a continuous function on $\mathbb{S}_{\mathrm{qf}}(\mathrm{T_{tr}})$, namely $\gamma \circ \mathbf{h}_*$ where $\mathbf{h}$ is as in the proof. In other words, every quantifier-free definable predicate on $\prod_{j\in\mathbb{N}} D_{r_j}$ relative to $\mathrm{T_{tr}}$ extends to a global quantifier-free definable predicate. The same can be said for definable predicates, dropping the word "quantifier-free" in this argument.

Proposition 3.32 also leads to a proof of the following fact, which is well-known among W*–algebraists:

Lemma 3.34 *Let* $\mathbf{X}$ *be an* $\mathbb{N}$*–tuple in a tracial* W*–algebra* $\mathcal{M}$ *and* $\mathbf{Y}$ *an* $\mathbb{N}$*–tuple in a tracial* W*–algebra* $\mathcal{N}$*. Let* $\mathrm{W}^*(\mathbf{X})$ *and* $\mathrm{W}^*(\mathbf{Y})$ *be the* W*–subalgebras generated by* $\mathbf{X}$ *and* $\mathbf{Y}$ *with the traces obtained from restricting the traces on* $\mathcal{M}$ *and* $\mathcal{N}$ *respectively. Then the following are equivalent:*

(1) $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) = \mathrm{tp}_{\mathrm{qf}}^{\mathcal{N}}(\mathbf{Y})$.

(2) *There exists a trace-preserving* $*$*–isomorphism* $\sigma : \mathrm{W}^*(\mathbf{X}) \to \mathrm{W}^*(\mathbf{Y})$ *such that* $\sigma(\mathbf{X}) = \mathbf{Y}$.

Proof (2) $\implies$ (1). If such a $*$–isomorphism σ exists, then for every $p \in \mathbb{C}\langle x_i, x_i^* : i \in \mathbb{N}\rangle$, we have $\tau_{\mathcal{M}}(p(\mathbf{X})) = \tau_{\mathcal{N}}(\sigma(p(\mathbf{X}))) = \tau_{\mathcal{N}}(p(\mathbf{Y}))$. Hence, every atomic formula evaluates to the same thing on $\mathbf{X}$ and on $\mathbf{Y}$. Since general quantifier-free formulas are obtained by applying continuous connectives to atomic formulas, it follows by induction on complexity that $\phi^{\mathcal{M}}(\mathbf{X}) = \phi^{\mathcal{N}}(\mathbf{Y})$ for any quantifier-free formula $\mathbf{Y}$, and hence $\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) = \mathrm{tp}_{\mathrm{qf}}^{\mathcal{N}}(\mathbf{Y})$.

(1) $\implies$ (2). Let $\mathcal{F}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ be the set of quantifier-free definable functions in $\mathcal{L}_{\mathrm{tr}}$ with respect to T_{tr}. Since quantifier-free functions are closed under composition, $\mathcal{F}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ is a $*$–algebra. Moreover, the evaluation maps $\alpha : \mathcal{F}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}}) \to \mathcal{M}, f \mapsto f^{\mathcal{M}}(\mathbf{X})$ and $\beta : \mathcal{F}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}}) \to \mathcal{N}, f \mapsto f^{\mathcal{N}}(\mathbf{Y})$ are $*$–homomorphisms, and by the previous proposition the images of α and β are $\mathrm{W}^*(\mathbf{X})$ and $\mathrm{W}^*(\mathbf{Y})$ respectively. Since $\operatorname{Re}\operatorname{tr}(f)$ and $\operatorname{Im}\operatorname{tr}(f)$ are quantifier-free definable predicates, $\tau_{\mathcal{M}} \circ \alpha(f) = \operatorname{tr}(f)^{\mathcal{M}}(\mathbf{X}) = \operatorname{tr}(f)^{\mathcal{N}}(\mathbf{Y}) = \tau_{\mathcal{N}} \circ \beta(f)$ for $f \in \mathcal{F}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$, hence also $\|\alpha(f)\|_{L^2(\mathcal{M})} = \|\beta(f)\|_{L^2(\mathcal{N})}$. This implies that $\ker\alpha = \ker\beta$. Therefore, we obtain a $*$–isomorphism $\mathrm{W}^*(\mathbf{X}) \cong \mathcal{F}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})/\ker\alpha = \mathcal{F}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})/\ker\beta \cong \mathrm{W}^*(\mathbf{Y})$, which is trace-preserving since $\tau_{\mathcal{M}} \circ \alpha = \tau_{\mathcal{N}} \circ \beta$. □

4 Entropy for types

We define a version of Hayes' 1–bounded entropy for types rather than only quantifier-free types. Later, in §5, we will see that Hayes' 1–bounded entropy of $\mathcal{N}$ in the presence of $\mathcal{M}$ (denoted $h^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$) can be realized as a special case of entropy for a closed subset of the type space.

4.1 Definition of covering entropy

If $\mathcal{K}$ is a subset of the type space $\mathbb{S}(\mathrm{T}_{\mathrm{tr}})$ and $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, we define

$$\Gamma_{\mathbf{r}}^{(n)}(\mathcal{K}) = \left\{ \mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{M_n(\mathbb{C})} : \mathrm{tp}^{M_n(\mathbb{C})}(\mathbf{X}) \in \mathcal{K} \right\}.$$

We view this as a microstate space as in Voiculescu's free entropy theory. We will then define the entropy of $\mathcal{O}$ through covering numbers of $\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O})$ up to unitary conjugation.

Definition 4.1 (Orbital covering numbers) Given $\Omega \subseteq M_n(\mathbb{C})^{\mathbb{N}}$ and a finite $F \subseteq \mathbb{N}$ and $\epsilon > 0$, we define $N_{F,\epsilon}^{\mathrm{orb}}(\Omega)$ to be the set of $\mathbf{Y} \in M_n(\mathbb{C})^{\mathbb{N}}$ such that there exists a unitary U in $M_n(\mathbb{C})$ and $\mathbf{X} \in \Omega$ such that $\|Y_i - UX_iU^*\|_2 < \epsilon$ for all $i \in F$. If $\Omega \subseteq N_{F,\epsilon}^{\mathrm{orb}}(\Omega')$, we say that Ω' *orbitally* (F,ϵ)*–covers* Ω. We denote by $K_{F,\epsilon}^{\mathrm{orb}}(\Omega)$ the minimum cardinality of a set Ω' that orbitally (F,ϵ)–covers Ω.

Definition 4.2 Fix a non-principal ultrafilter $\mathcal{U}$ on $\mathbb{N}$. For a subset $\mathcal{K}$ of the $\mathbb{S}(\mathrm{T}_{\mathrm{tr}})$ and $F \subseteq I$ finite and $\epsilon > 0$, we define

$$\mathrm{Ent}_{\mathbf{r},F,\epsilon}^{\mathcal{U}}(\mathcal{K}) = \inf_{\text{open } \mathcal{O} \supseteq \mathcal{K}} \lim_{n \to \mathcal{U}} \frac{1}{n^2} \log K_{F,\epsilon}^{\mathrm{orb}}(\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O})).$$

Observation 4.3 (Monotonicity) *Let $\mathcal{K}' \subseteq \mathcal{K} \subseteq \mathbb{S}(\mathrm{T}_{\mathrm{tr}})$, let $F' \subseteq F \subseteq \mathbb{N}$ finite, let $0 < \epsilon \leq \epsilon'$, and let $\mathbf{r}, \mathbf{r}' \in (0,\infty)^{\mathbb{N}}$ with $r_j' \leq r_j$. Then*

$$\mathrm{Ent}_{\mathbf{r}',F',\epsilon'}^{\mathcal{U}}(\mathcal{K}') \leq \mathrm{Ent}_{\mathbf{r},F,\epsilon}^{\mathcal{U}}(\mathcal{K}).$$

In particular, if $\mathcal{O} \subseteq \mathbb{S}(\mathrm{T}_{\mathrm{tr}})$ is open, then

$$\mathrm{Ent}_{\mathbf{r},F,\epsilon}^{\mathcal{U}}(\mathcal{O}) = \lim_{n \to \mathcal{U}} \frac{1}{n^2} \log K_{F,\epsilon}^{\mathrm{orb}}(\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O})).$$

Definition 4.4 (Entropy for types) For $\mathcal{K} \subseteq \mathbb{S}(\mathrm{T})$, define

$$\mathrm{Ent}_{\mathbf{r}}^{\mathcal{U}}(\mathcal{K}) := \sup_{\substack{\text{finite } F \subseteq \mathbb{N} \\ \epsilon > 0}} \mathrm{Ent}_{\mathbf{r},F,\epsilon}^{\mathcal{U}}(\mathcal{K})$$

and

$$\mathrm{Ent}^{\mathcal{U}}(\mathcal{K}) := \sup_{\mathbf{r} \in (0,\infty)^{\mathbb{N}}} \mathrm{Ent}_{\mathbf{r}}^{\mathcal{U}}(\mathcal{K}).$$

Moreover, if $\mu \in \mathbb{S}(\mathrm{T}_{\mathrm{tr}})$, we define $\mathrm{Ent}^{\mathcal{U}}(\mu) = \mathrm{Ent}^{\mathcal{U}}(\{\mu\})$.

4.2 Variational principle

In this section, we show that the covering entropy defines an upper semi-continuous function on the type space, and then deduce a variational principle for the entropy of a closed set, in the spirit of various results in the theory of entropy and large deviations.

Lemma 4.5 (Upper semi-continuity) *For each $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, $F \subseteq \mathbb{N}$ finite, and $\epsilon > 0$, the function $\mu \mapsto \mathrm{Ent}_{\mathbf{r},F,\epsilon}^{\mathcal{U}}(\mu)$ is upper semi-continuous on $\mathbb{S}(\mathrm{T}_{\mathrm{tr}})$.*

Proof For each open $\mathcal{O} \subseteq \mathbb{S}(\mathrm{T}_{\mathrm{tr}})$, let

$$f_{\mathcal{O}}(\mu) = \begin{cases} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}), & \mu \in \mathcal{O} \\ \infty, & \text{otherwise.} \end{cases}$$

Since $\mathcal{O}$ is open, $f_{\mathcal{O}}$ is upper semi-continuous. Moreover, $\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mu)$ is the infimum of $f_{\mathcal{O}}(\mu)$ over all open $\mathcal{O} \subseteq \mathbb{S}(\mathrm{T}_{\mathrm{tr}})$, and the infimum of a family of upper semi-continuous functions is upper semi-continuous. □

Proposition 4.6 (Variational principle) *Let $\mathcal{K} \subseteq \mathbb{S}(\mathrm{T}_{\mathrm{tr}})$ and let $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, $F \subseteq \mathbb{N}$ finite, and $\epsilon > 0$. Then*

$$\sup_{\mu\in\mathcal{K}} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mu) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{K}) \leq \sup_{\mu\in\mathrm{cl}(\mathcal{K})} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mu). \tag{4–1}$$

Hence,

$$\sup_{\mu\in\mathcal{K}} \mathrm{Ent}^{\mathcal{U}}(\mu) \leq \mathrm{Ent}^{\mathcal{U}}(\mathcal{K}) \leq \sup_{\mu\in\mathrm{cl}(\mathcal{K})} \mathrm{Ent}^{\mathcal{U}}(\mu). \tag{4–2}$$

Proof If $\mu \in \mathcal{K}$, then by monotonicity (Observation 4.3), $\mathrm{Ent}^{\mathcal{U}}_{F,\epsilon}(\{\mu\}) \leq \mathrm{Ent}^{\mathcal{U}}_{F,\epsilon}(\mathcal{K})$. Taking the supremum over $\mu \in \mathcal{K}$, we obtain the first inequality of (4–1).

For the second inequality of (4–1), let $C = \sup_{\mu\in\mathrm{cl}(\mathcal{K})} \mathrm{Ent}^{\mathcal{U}}_{F,\epsilon}(\mu)$. If $C = \infty$, there is nothing to prove. Otherwise, let $C' > C$. For each $\mu \in \mathrm{cl}(\mathcal{K}) \cap \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$, there exists some open neighborhood $\mathcal{O}_\mu$ of μ in $\mathbb{S}(\mathrm{T}_{\mathrm{tr}})$ such that $\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_\mu) < C'$. Since $\{\mathcal{O}_\mu\}_{\mu\in\mathrm{cl}(\mathcal{K})\cap\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})}$ is an open cover of the compact set $\mathrm{cl}(\mathcal{K}) \cap \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$, there exist μ_1, …, $\mu_k \in \mathrm{cl}(\mathcal{K}) \cap \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ such that

$$\mathcal{K} \cap \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) \subseteq \bigcup_{j=1}^{k} \mathcal{O}_{\mu_j}.$$

Let $\mathcal{O} = \bigcup_{j=1}^{k} \mathcal{O}_{\mu_j}$. Then

$$K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O})) \leq \sum_{j=1}^{k} K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}_{\mu_j})) \leq k \max_j K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}_{\mu_j})).$$

Thus,

$$\frac{1}{n^2} \log K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O})) \leq \frac{1}{n^2} \log k + \max_j \frac{1}{n^2} \log K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}_{\mu_j})).$$

Taking the limit as $n \to \mathcal{U}$,

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathrm{cl}(\mathcal{K})) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}) \leq \max_j \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_{\mu_j}) \leq C'.$$

Since $C' > C$ was arbitrary,

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{K}) \leq \mathrm{Ent}_{\mathbf{r},F,\epsilon}(\mathrm{cl}(\mathcal{K})) \leq C = \sup_{\mu \in \mathrm{cl}(\mathcal{K})} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mu)$$

completing the proof of (4–1). Taking the supremum over F and ϵ and $\mathbf{r}$ in (4–1), we obtain (4–2). □

4.3 Invariance under change of coordinates

Next, we prove certain invariance properties of the covering entropy. First, $\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mu)$ is independent of $\mathbf{r}$ provided that $\mu \in \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$. Second, if $\mathrm{W}^*(\mathbf{X}) = \mathrm{W}^*(\mathbf{Y})$ inside $\mathcal{M}$, then $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{Y}))$, which allows us to define $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N}, \mathcal{M})$ for a W^*–subalgebra $\mathcal{N}$ inside $\mathcal{M}$. Both of these properties are deduced from the following lemma about push-forward under definable functions. This is closely related to Hayes [16, Lemma A.8 and Theorem A.9].

Proposition 4.7 (Monotonicity under push-forward) *Let* $\mathbf{f}$ *be a definable function relative to* T_{tr}, *let* $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, *and let* $\mathbf{r}' \in (0,\infty)^{\mathbb{N}}$ *be such that* $\mathbf{f}$ *maps* $\prod_{j\in\mathbb{N}} D_{r_j}$ *into* $\prod_{j\in\mathbb{N}} D_{r'_j}$. *Let* $\mathcal{K} \subseteq \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ *be closed. Then*

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}'}(\mathbf{f}_*(\mathcal{K})) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mathcal{K}).$$

Remark 4.8 The analogous monotonicity property does not hold for the original 1–bounded entropy h of a quantifier-free type, but it does hold for 1–bounded entropy in the presence. The monotonicity property holds for the full type and for the existential type of $\mathbf{X}$ because those types already encode information about how $\mathbf{X}$ interacts with the ambient algebra. For more information, see Remark 5.16.

Proof Let $F' \subseteq \mathbb{N}$ finite and $\epsilon' \in (0,1)$ be given. Because $\mathbf{f}$ is a definable function, it is uniformly continuous by Lemma 3.19, hence there exists a finite $F \subseteq I$ and $\epsilon > 0$ such that for every $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathbf{X}, \mathbf{Y} \in \prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r_j}$,

$$\text{(4–3)} \qquad \|X_j - Y_j\|_2 < \epsilon \text{ for all } j \in F \implies \|f_{j'}(\mathbf{X}) - f_{j'}(\mathbf{Y})\|_2 < \epsilon'/3 \text{ for all } j' \in F'.$$

Let $\mathcal{O}$ be a neighborhood of $\mathcal{K}$ in $\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$. By Urysohn's lemma, there exists a continuous function $\psi : \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) \to [0,1]$ such that $\phi = 0$ on $\mathcal{K}$ and $\phi = 1$ on $\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) \setminus \mathcal{O}$. As in Proposition 3.9, there exists a formula η such that $|\eta^{\mathcal{M}} - \phi^{\mathcal{M}}| < \epsilon'/3$ on $\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r_j}$. Next, define $\psi^{\mathcal{M}} : \prod_{j\in\mathbb{N}} S^{\mathcal{M}}_{r'_j} \to \mathbb{R}$ by

$$\psi^{\mathcal{M}}(\mathbf{Y}) = \inf_{\mathbf{X} \in \prod_{j\in\mathbb{N}} D^{\mathcal{M}}_j} \left(\eta^{\mathcal{M}}(\mathbf{X}) + \max_{j' \in F'} d^{\mathcal{M}}(f^{\mathcal{M}}_{j'}(\mathbf{X}), Y_{j'}) \right)$$

which is a definable predicate relative to T_{tr} by Lemma 3.12.

Viewing ψ as a continuous function on $\mathbb{S}_{\mathbf{r}'}(\mathrm{T}_{\mathrm{tr}})$, let $\mathcal{O}' = \psi^{-1}((-\infty, 2\epsilon'/3))$. Note that $\mathbf{f}_*(\mathcal{K}) \subseteq \mathcal{O}'$ since if $\mathbf{Y} = \mathbf{f}^{\mathcal{M}}(\mathbf{X})$, then we can take this value of $\mathbf{X}$ in the infimum defining ψ and obtain that $\psi^{\mathcal{M}}(\mathbf{Y}) \leq \epsilon'/3$. Meanwhile, if $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathbf{Y} \in \prod_{j\in\mathbb{N}} D_{r'_j}^{\mathcal{M}}$ with $\mathrm{tp}^{\mathcal{M}}(\mathbf{Y}) \in \mathcal{O}'$, then there exists $\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{M}}$ with

$$\eta^{\mathcal{M}}(\mathbf{X}) + \max_{j'\in F'} d^{\mathcal{M}}(f_{j'}^{\mathcal{M}}(\mathbf{X}), \mathbf{Y}_{j'}) < \frac{2\epsilon'}{3}$$

which implies that $\mathrm{tp}^{\mathcal{M}}(\mathbf{X}) \in \mathcal{O}$ and $\max_{i'\in F'} \|f_{i'}^{\mathcal{M}}(\mathbf{Y}) - \mathbf{X}_{i'}\| < 2\epsilon'/3$. Applying this with $\mathcal{M} = M_n(\mathbb{C})$, we obtain $\Gamma^{(n)}(\mathcal{O}') \subseteq N_{2\epsilon'/3}(\mathbf{f}^{M_n(\mathbb{C})}(\Gamma^{(n)}(\mathcal{O})))$. If Ω is an (F,ϵ)–cover of $\Gamma^{(n)}(\mathcal{O})$, then by (4–3) and the fact that $\mathbf{f}$ is equivariant with respect to conjugation of an $\mathbb{N}$–tuple by a fixed unitary, $\mathbf{f}_*(\Omega)$ is an orbital $(F',\epsilon'/3)$–cover of $\mathbf{f}_*(\Gamma^{(n)}(\mathcal{O}))$, and therefore also an orbital (F',ϵ')–cover of $\Gamma^{(n)}(\mathcal{O}')$. It follows that

$$K^{\mathrm{orb}}_{\mathbf{r}',F',\epsilon'}(\Gamma^{(n)}(\mathcal{O}')) \leq K^{\mathrm{orb}}_{\mathbf{r},F,\epsilon}(\Gamma^{(n)}(\mathcal{O})).$$

Hence,

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}',F',\epsilon'}(\mathbf{f}_*(\mathcal{K})) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}',F',\epsilon'}(\mathcal{O}') \leq \mathrm{Ent}_{\mathbf{r},F,\epsilon}(\mathcal{O}).$$

Since $\mathcal{O}$ was an arbitrary neighborhood of $\mathcal{K}$, we obtain

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}',F',\epsilon'}(\mathbf{f}_*(\mathcal{K})) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{K}) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mathcal{K}).$$

Since F' and ϵ' were arbitrary, we conclude that $\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}'}(\mathbf{f}_*(\mathcal{K})) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mathcal{K})$, as desired. □

Corollary 4.9 *If $\mathcal{K}$ is a closed subset of $\mathbf{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$, then* $\mathrm{Ent}^{\mathcal{U}}(\mathcal{K}) = \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mathcal{K})$.

Proof By definition, $\mathrm{Ent}^{\mathcal{U}}(\mathcal{K}) \geq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mathcal{K})$. On the other hand, fix some $\mathbf{r}' \in (0,\infty)^{\mathbb{N}}$ and let $\mathbf{r}'' = \max(\mathbf{r}', \mathbf{r})$. By Observation 4.3, $\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}'}(\mathcal{K}) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}''}(\mathcal{K})$. Now applying Proposition 4.7 to the identity map, since id maps $\prod_{j\in\mathbb{N}} D_{r_j}$ into $\prod_{j\in\mathbb{N}} D_{r''_j}$, it follows that $\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}''}(\mathcal{K}) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mathcal{K})$. Since $\mathbf{r}'$ was arbitrary, $\mathrm{Ent}^{\mathcal{U}}(\mathcal{K}) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mathcal{K})$. □

Corollary 4.10 *Let $\mathcal{M} = (M,\tau)$ be a tracial* W*–*algebra and* $\mathbf{X}, \mathbf{Y} \in M^{\mathbb{N}}$. *If* $\mathbf{Y} \in \mathrm{W}^*(\mathbf{X})^{\mathbb{N}}$, *then*

$$\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{Y})) \leq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

In particular, if $\mathrm{W}^*(\mathbf{X}) = \mathrm{W}^*(\mathbf{Y})$, *then* $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{Y}))$.

Proof By Proposition 3.32, there exists a quantifier-free definable function $\mathbf{f}$ relative to T_{tr} such that $\mathbf{Y} = \mathbf{f}^{\mathcal{M}}(\mathbf{X})$. Now $\mathrm{tp}^{\mathcal{M}}(\mathbf{Y}) = \mathbf{f}_* \mathrm{tp}^{\mathcal{M}}(\mathbf{X})$. Hence, applying Proposition 4.7 to $\mathcal{K} = \{\mathrm{tp}^{\mathcal{M}}(\mathbf{X})\}$ (for an appropriate choice of $\mathbf{r}$), we obtain $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{Y})) \leq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$. The second claim follows by symmetry. □

With this invariance result in hand, it seems natural to define the covering entropy for a separable W*–subalgebra of $\mathcal{M}$ as the entropy of any $\mathbb{N}$–tuple of generators. However, the following definition works even in the non-separable case.

Definition 4.11 If $\mathcal{M} = (M, \tau)$ is a tracial W*–algebra and $\mathcal{N}$ is a W*–subalgebra, we define

$$\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M}) = \sup_{\mathbf{X} \in L^\infty(\mathcal{N})^{\mathbb{N}}} \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

Observation 4.12 *Let $\mathcal{M}$ be a tracial* W*–*algebra, and let $\mathcal{N}$ be a* W*–*subalgebra. If $\mathbf{X} \in L^\infty(\mathcal{N})^{\mathbb{N}}_{\mathrm{sa}}$ generates $\mathcal{N}$, then for any $\mathbf{Y} \in \mathcal{N}^{\mathbb{N}}$, we have $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{Y})) \leq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$ by Corollary 4.10, and therefore,*

$$\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M}) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

Moreover, if $\mathcal{P}$ is a W*–*subalgebra of $\mathcal{N}$, then $\mathrm{Ent}^{\mathcal{U}}(\mathcal{P} : \mathcal{M}) \leq \mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$.*

Remark 4.13 Furthermore, it is evident from Definition 4.11 that $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$ only depends on the set of types in $\mathcal{M}$ that are realized in $L^\infty(\mathcal{N})^{\mathbb{N}}$. Hence, if two embeddings $\mathcal{N} \to \mathcal{M}_1$ and $\mathcal{N} \to \mathcal{M}_2$ are elementarily equivalent—meaning that for every definable predicate ϕ and $\mathbf{X} \in L^\infty(\mathcal{N})^{\mathbb{N}}$, we have $\phi^{\mathcal{M}_1}(\mathbf{X}) = \phi^{\mathcal{M}_2}(\mathbf{X})$—then $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M}_1) = \mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M}_2)$.

4.4 Entropy and ultraproduct embeddings

Lemma 4.14 (Ultraproduct realization of types) *Let $\mathcal{Q} = \prod_{n \to \mathcal{U}} M_n(\mathbb{C})$. Let $\mu \in \mathbb{S}(\mathrm{T}_{\mathrm{tr}})$. Then $\mathrm{Ent}^{\mathcal{U}}(\mu)$ is either nonnegative or it is $-\infty$. Moreover, $\mathrm{Ent}^{\mathcal{U}}(\mu) \geq 0$ if and only if there exists $\mathbf{X} \in L^\infty(\mathcal{Q})^{\mathbb{N}}$ such that $\mathrm{tp}^{\mathcal{Q}}(\mathbf{X}) = \mu$.*

Proof Note that $\log K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}))$ is either ≥ 0 or it is $-\infty$. Therefore, $\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mu)$ is either nonnegative or it is $-\infty$. It remains to show the second claim of the lemma.

($\implies$) In light of the foregoing argument, if $\mathrm{Ent}^{\mathcal{U}}(\mu) \geq 0$, then $\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mu) \geq 0$ for some $\mathbf{r}$. By Observations 3.14 and 3.15, $\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ is metrizable, hence there is a sequence $(\mathcal{O}_k)_{k \in \mathbb{N}}$ of neighborhoods of μ in $\mathbb{S}(\mathrm{T})$ such that $\overline{\mathcal{O}}_{k+1} \subseteq \mathcal{O}_k$ and $\bigcap_{k \in \mathbb{N}} \mathcal{O}_k = \{\mu\}$. For $k \in \mathbb{N}$, let $E_k = \{n \in \mathbb{N} : \Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}_k) \neq \varnothing\}$. Now choose $\mathbf{X}^{(n)} \in M_N(\mathbb{C})^{\mathbb{N}}$ as follows. For each $n \notin E_1$, set $\mathbf{X}^{(n)} = 0$. For each $n \in E_k \setminus E_{k+1}$, let $\mathbf{X}^{(n)}$ be an element of $\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}_k)$. If $n \in \bigcap_{k \in \mathbb{N}} E_k$, that means that $\Gamma^{(n)}(\{\mu\}) \neq \varnothing$, so in this case we may choose $\mathbf{X}^{(n)} \in M_n(\mathbb{C})^{\mathbb{N}}$ with $\mathrm{tp}^{M_n(\mathbb{C})}(\mathbf{X}^{(n)}) = \mu$.

Since $\mathcal{U}$ is an ultrafilter, either $E_k \in \mathcal{U}$ or $E_k^c \in \mathcal{U}$. If we had $E_k^c \in \mathcal{U}$, then $\lim_{n\to\mathcal{U}}(1/n^2)\log K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}(\mathcal{O}_k))$ would be $-\infty$ since the set would be empty for $n \in E_k^c$. Hence, $E_k \in \mathcal{U}$. For $n \in E_k$, we have $\mathrm{tp}^{M_n(\mathbb{C})}(\mathbf{X}^{(n)}) \in \mathcal{O}_k$. Therefore, $\lim_{n\to\mathcal{U}} \mathrm{tp}^{M_n(\mathbb{C})}(\mathbf{X}^{(n)}) \in \overline{\mathcal{O}}_k$. Since this holds for all k, $\lim_{n\to\mathcal{U}} \mathrm{tp}^{M_n(\mathbb{C})}(\mathbf{X}^{(n)}) = \mu$. Let $\mathbf{X} = [\mathbf{X}^{(n)}]_{n\in\mathbb{N}} \in L^\infty(\mathcal{Q})^{\mathbb{N}}$. Then

$$\mathrm{tp}^{\mathcal{Q}}(\mathbf{X}) = \lim_{n\to\mathcal{U}} \mathrm{tp}^{M_n(\mathbb{C})}(\mathbf{X}^{(n)}) = \mu.$$

($\Longleftarrow$) Suppose that $\mathbf{X}$ is an element of the ultraproduct with type μ. Let $r_j = \|X_j\|_\infty$. Express $\mathbf{X}$ as $[\mathbf{X}^{(n)}]_{n\in\mathbb{N}}$ for some $\mathbf{X}^{(n)} \in M_n(\mathbb{C})^{\mathbb{N}}$ with $\|X_j^{(n)}\| \le r_j$. Since the type of $\mathbf{X}^{(n)}$ converges to the type of $\mathbf{X}$, for every neighborhood $\mathcal{O}$ of μ, there exists $E \in \mathcal{U}$ such that $\mathrm{tp}^{M_n(\mathbb{C})}(\mathbf{X}^{(n)}) \in \mathcal{O}$ for all $n \in E$, and in particular, $\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}) \neq \varnothing$ for $n \in E$. This implies that $\mathrm{Ent}^{\mathcal{U}}_{F,\epsilon}(\mathcal{O}) \ge 0$ for every F and ϵ. Hence, $\mathrm{Ent}^{\mathcal{U}}(\mu) \ge 0$. □

Recall that an embedding $\mathcal{M} \to \mathcal{Q}$ of tracial W*–algebras is said to be *elementary* if for every definable predicate ϕ and $\mathbf{X} \in L^\infty(\mathcal{M})^{\mathbb{N}}$, we have $\phi^{\mathcal{Q}}(\mathbf{X}) = \phi^{\mathcal{M}}(\mathbf{X})$. This in particular implies that $\mathcal{M}$ and $\mathcal{Q}$ are elementarily equivalent, that is, they have the same theory.

Corollary 4.15 *Suppose that $\mathcal{M}$ is a separable tracial* W*–*algebra and $\mathcal{N} \subseteq \mathcal{M}$ is a* W*–*subalgebra. If* $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M}) \ge 0$, *then there exists an elementary embedding $\iota : \mathcal{M} \to \mathcal{Q}$.*

Remark 4.16 Since the embedding $\iota : \mathcal{M} \to \mathcal{Q}$ is elementary, in particular the embeddings $\mathcal{N} \to \mathcal{M}$ and $\mathcal{N} \to \mathcal{Q}$ are elementarily equivalent, and hence $\mathrm{Ent}^{\mathcal{U}}(\iota(\mathcal{N}) : \mathcal{Q}) = \mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$.

Proof of Corollary 4.15 By Observation 4.12, $\mathrm{Ent}^{\mathcal{U}}(\mathcal{M} : \mathcal{M}) \ge \mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M}) \ge 0$. Let $\mathbf{X} \in L^\infty(\mathcal{M})^{\mathbb{N}}$ generate $\mathcal{M}$. Then by the previous lemma, there exists $\mathbf{X}' \in \mathcal{Q}$ with the same type of $\mathbf{X}$. In particular, since $\mathbf{X}$ and $\mathbf{X}'$ have the same quantifier-free type, Lemma 3.34 shows that there is an embedding $\iota : \mathcal{M} \to \mathcal{Q}$ with $\iota(\mathbf{X}) = \mathbf{X}'$. To show that ι is elementary, suppose that $\mathbf{Y} \in L^\infty(\mathcal{M})^{\mathbb{N}}$ and ϕ is a definable predicate. By Proposition 3.32, there exists a quantifier-free definable function $\mathbf{f}$ such that $\mathbf{Y} = \mathbf{f}^{\mathcal{M}}(\mathbf{X})$. Since $\mathbf{f}$ is quantifier-free, $d^{\mathcal{Q}}(\iota(Y_j), f_j^{\mathcal{Q}}(\mathbf{X}')) = d^{\mathcal{M}}(Y_j, f_j^{\mathcal{M}}(\mathbf{X})) = 0$, hence $\iota(\mathbf{Y}) = \mathbf{f}^{\mathcal{Q}}(\iota(\mathbf{X}))$. Therefore, $\phi^{\mathcal{Q}}(\iota(\mathbf{Y})) = (\phi\circ\mathbf{f})^{\mathcal{Q}}(\mathbf{X}') = (\phi\circ\mathbf{f})^{\mathcal{M}}(\mathbf{X}) = \phi^{\mathcal{M}}(\mathbf{Y})$, where the middle equality follows because $\mathrm{tp}^{\mathcal{Q}}(\mathbf{X}') = \mathrm{tp}^{\mathcal{M}}(\mathbf{X})$, and therefore the embedding is elementary. □

4.5 Entropy and Algebraicity

In this section, we show that $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M}) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{acl}(\mathcal{N}) : \mathcal{M})$, where $\mathrm{acl}(\mathcal{N})$ is the *algebraic closure* of continuous model theory. At present, very little is known about algebraic closures for tracial W*–algebras. Nonetheless, it is natural to study how the model-theoretic 1–bounded entropy behaves under this model-theoretic operation, analogously to how Hayes studied the behavior of 1–bounded entropy under various W*–algebraic operations. (See Hayes [16, § 2] and Hayes, Jekel, Nelson, and Sinclair [18, § 2.3].)

First, we explain the definition of algebraic closure.

Definition 4.17 (Algebraicity) Let $\mathcal{M}$ be a structure in some language $\mathcal{L}$, and let $\mathcal{N}$ be a substructure. Let S be a sort in $\mathcal{L}$.

- A map $\phi : S^{\mathcal{M}} \to \mathbb{R}$ is a *definable predicate in $\mathcal{M}$ over $\mathcal{N}$* if for every $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$ and $\epsilon > 0$, there exists a formula ψ in variables x_j from S_j for $j \in \mathbb{N}$ and y_j from S'_j for $j \in \mathbb{N}$, and there exists $\mathbf{Y} \in \prod_{j\in\mathbb{N}} (S'_j)^{\mathcal{N}}$ such that
$$|\phi(\mathbf{X}) - \psi^{\mathcal{M}}(\mathbf{X}, \mathbf{Y})| < \epsilon \text{ for all } \mathbf{X} \in \prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}.$$
- If $A \subseteq S^{\mathcal{M}}$, we say that A is *definable in $\mathcal{M}$ over $\mathcal{N}$* if the map $S^{\mathcal{M}} \to \mathbb{R} : X \mapsto d^{\mathcal{M}}(X, A)$ is definable in $\mathcal{M}$ over $\mathcal{N}$.
- If $a \in S^{\mathcal{M}}$, we say that a is *algebraic over $\mathcal{N}$* if there exists a compact set $A \subseteq S^{\mathcal{M}}$ such that $a \in A$ and A is definable in $\mathcal{M}$ over $\mathcal{N}$.

Remark 4.18 It will be convenient in our arguments that for tracial W*–algebras $\mathcal{M}$ and $\mathcal{N}$, if a function $\phi : L^\infty(\mathcal{M}) \to \mathbb{R}$ is definable in $\mathcal{M}$ over $\mathcal{N}$, then there exists a definable predicate θ and $\mathbf{Y} \in L^\infty(\mathcal{N})^{\mathbb{N}}$ such that $\phi(X) = \theta(X, \mathbf{Y})$. This follows by a forced-limit argument similar to Proposition 3.32: Since ϕ is definable in $\mathcal{M}$ over $\mathcal{N}$, then for each $k \in \mathbb{N}$, there exists a formula θ_k and $\mathbf{Y}_k \in L^\infty(\mathcal{N})^{\mathbb{N}}$ such that
$$|\phi(X) - \theta_k(X, \mathbf{Y}_k)| < \frac{1}{2^k} \text{ for } X \in D_k^{\mathcal{M}}.$$
Let $\mathbf{Y}$ by an $\mathbb{N}$–tuple obtained by joining together the $\mathbf{Y}_k$'s into a single tuple, so that θ_k can be viewed as a definable predicate in $(X, \mathbf{Y})$. Similar to the proof of Proposition 3.32, there exists a definable predicate ψ_k such that $\psi_k^{\mathcal{M}}(X, \mathbf{Y}) = 1$ and $\psi_k \cdot (\theta_{k+1} - \theta_k) < 2/2^k$ on D_k. Then
$$\theta := \theta_1 + \sum_{k=1}^{\infty} \psi_k(\theta_{k+1} - \theta_k)$$
converges uniformly on every domain D_r and satisfies $\phi(X) = \theta^{\mathcal{M}}(X, \mathbf{Y})$.

Definition 4.19 (Algebraic closure) Let $\mathcal{M}$ be an $\mathcal{L}$–structure and $\mathcal{N}$ an $\mathcal{L}$–substructure. We define $S^{\mathrm{acl}(\mathcal{N})}$ to be the set of $a \in S^{\mathcal{M}}$ that are algebraic in $\mathcal{M}$ over $\mathcal{N}$. We let $\mathrm{acl}(\mathcal{N}) = (S^{\mathrm{acl}(\mathcal{N})})_{S\in\mathcal{S}}$. (Although we omit $\mathcal{M}$ from the notation, the algebraic closure a priori depends on the ambient structure $\mathcal{M}$.)

For the properties of algebraic closure, see Ben Yaacov, Berenstein, Henson, and Usvyatsov [3, § 10]. In particular, one can show that if $\mathcal{N}_1 \subseteq \mathrm{acl}(\mathcal{N}_2)$, then $\mathrm{acl}(\mathcal{N}_1) \subseteq \mathrm{acl}(\mathcal{N}_2)$ ("what is algebraic over the algebraic closure of $\mathcal{N}_2$ is algebraic over $\mathcal{N}_2$"). Moreover, one can verify directly from Definition 4.17 that f is a term and $Y_1, \ldots, Y_k \in \mathcal{N}$, then $f^{\mathcal{M}}(Y_1, \ldots, Y_k) \in \mathrm{acl}(\mathcal{N})$. By combining these properties, it follows that $\mathrm{acl}(\mathcal{N})$ is an $\mathcal{L}$–substructure of $\mathcal{M}$.

Thus, in particular, if $\mathcal{N} \subseteq \mathcal{M}$ are tracial W*–algebras, then the algebraic closure $\mathrm{acl}(\mathcal{N})$ of $\mathcal{N}$ in $\mathcal{M}$ is a tracial W*–subalgebra of $\mathcal{M}$ as well. We will show that $\mathrm{Ent}^{\mathcal{U}}(\mathrm{acl}(\mathcal{N}) : \mathcal{M}) = \mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$. We first consider the case of adjoining to an $\mathbb{N}$–tuple $\mathbf{X}$ a single element Y that is algebraic over $\mathrm{W}^*(\mathbf{X})$, and this case takes the bulk of the work.

Theorem 4.20 *Let* $\mathbf{X}$ *be an* $\mathbb{N}$*–tuple in* $\mathcal{M} = (M, \tau)$*. Let* $Y \in M$ *be algebraic over* $\mathrm{W}^*(\mathbf{X})$*. Then*

$$\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(Y, \mathbf{X})) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

The inequality $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})) \leq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(Y, \mathbf{X}))$ follows from Proposition 4.7, so we only need to prove the opposite inequality.

The idea of the argument is that Y comes from a definable compact set A. We can cover A by some finite number k of ϵ–balls. Transferring this to the microstate approximations would tell us that for each matrix approximation $\mathbf{X}'$ for $\mathbf{X}$, the possible matrix approximations for Y can be covered by k many ϵ–balls. So the covering number for the microstate space of $(Y, \mathbf{X})$ would be at most k times that of $\mathbf{X}$; the factor of k is negligible in the large-n limit because we will take the logarithm and divide by n^2.

Proof $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})) \leq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(Y, \mathbf{X}))$ holds by Proposition 4.7.

By algebraicity of Y and Remark 4.18, there exists a compact $A \subseteq \mathcal{M}$, a definable predicate ϕ relative to T_{tr}, and $\mathbf{X}' \in \mathrm{W}^*(\mathbf{X})^{\mathbb{N}}$ such that $Y \in A$ and $d^{\mathcal{M}}(Z, A) = \phi^{\mathcal{M}}(Z, \mathbf{X}')$. Since $\mathbf{X}' = \mathbf{f}(\mathbf{X})$ for some quantifier-free definable function $\mathbf{f}$, we have

$$d^{\mathcal{M}}(Z, A) = \phi^{\mathcal{M}}(Z, \mathbf{f}(\mathbf{X})) = \psi^{\mathcal{M}}(Z, \mathbf{X})$$

where ψ is the definable predicate given by composing ψ with $\mathbf{f}$ in the coordinates $2, 3, \ldots$

Fix $\mathbf{r} = (r_j)_{j\in\mathbb{N}}$ such that $X_j \in D_{r_j}$ and fix r such that $Z \in D_r$. We want to show that:

$$\sup_{(F',\epsilon')} \inf_{\mathcal{O}\ni \mathrm{tp}^{\mathcal{M}}(Y,\mathbf{X})} \lim_{n\to\mathcal{U}} \frac{1}{n^2} \log K^{\mathrm{orb}}_{F',\epsilon'}(\Gamma^{(n)}_{r,\mathbf{r}}(\mathcal{O}')) \leq \sup_{(F,\epsilon)} \inf_{\mathcal{O}\ni \mathrm{tp}^{\mathcal{M}}(Y,\mathbf{X})} \lim_{n\to\mathcal{U}} \frac{1}{n^2} \log K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}))$$

Here we regard $\mathbb{N}$ as starting at 1, and we index the tuple $(Y, \mathbf{X})$ by $\{0\} \sqcup \mathbb{N}$, where the 0 index corresponds to Y. Fix $F' \subseteq \{0\} \sqcup \mathbb{N}$ finite and $\epsilon' > 0$. Since enlarging F' would only increase the quantity inside the $\sup_{(F',\epsilon')}$, assume without loss of generality that F' contains the index 0 corresponding to the variable Y, hence $F' = \{0\} \sqcup F_1$ for some $F_1 \subseteq \mathbb{N}$.

By compactness of A, there exists $k \in \mathbb{N}$ and there exist $Y_1, \ldots, Y_k \in A$ such that the $\epsilon'/4$–balls centered at $Y_1, \ldots, Y_k$ cover A. This implies that every point within a distance of $\epsilon'/4$ from A is within a distance of $\epsilon'/2$ from one of the points $Y_1, \ldots, Y_k$, and therefore

$$\sup_{Y\in D_r^{\mathcal{M}}} \min\Big(\epsilon'/4 - \psi^{\mathcal{M}}(Y,\mathbf{X}), \min(d^{\mathcal{M}}(Y,Y_1), \ldots, d^{\mathcal{M}}(Y,Y_k)) - \epsilon'/2\Big) \leq 0.$$

Choose $t_j \in (0,\infty)$ such that $Y_j \in D_{t_j}$. Let α and β be the definable predicates

$$\alpha(y_1,\ldots,y_k,\mathbf{x}) = \sup_{y\in D_r^{\mathcal{M}}} \min(\epsilon'/4 - \psi(y,\mathbf{x}), \min(d(y,y_1),\ldots,d(y,y_k)) - \epsilon'/2)$$

and

$$\beta(\mathbf{x}) = \inf_{y_1\in D_{t_1}} \cdots \inf_{y_k\in D_{t_k}} \alpha(y_1,\ldots,y_k,\mathbf{x})$$

so that $\alpha^{\mathcal{M}}(Y_1,\ldots,Y_k,\mathbf{X}) \leq 0$ and $\beta^{\mathcal{M}}(\mathbf{X}) \leq 0$.

By uniform continuity (Observation 3.11), there exists $F_2 \subseteq \mathbb{N}$ finite and $\delta > 0$ such that for all $\mathcal{N} \models \mathrm{T}_{\mathrm{tr}}$, all $Y_1' \in D_{t_1}, \ldots, Y_k' \in D_{t_k}$, and all $\mathbf{X}', \mathbf{X}'' \in \prod_{j\in\mathbb{N}} D_{r_j}^{\mathcal{N}}$, we have:

$$\max_{j\in F_2} d^{\mathcal{N}}(X_j', X_j'') < \delta \implies |\alpha^{\mathcal{N}}(Y_1',\ldots,Y_k',\mathbf{X}') - \alpha^{\mathcal{N}}(Y_1',\ldots,Y_k',\mathbf{X}'')| < \frac{\epsilon'}{16}$$

Fix a neighborhood $\mathcal{O}$ of $\mathrm{tp}^{\mathcal{M}}(\mathbf{X})$, and let

$$\mathcal{O}' = \{\mathrm{tp}^{\mathcal{N}}(Y',\mathbf{X}') : \mathcal{N} \models \mathrm{T}_{\mathrm{tr}}, \psi^{\mathcal{N}}(Y',\mathbf{X}') < \epsilon'/8, \beta^{\mathcal{N}}(\mathbf{X}) < \epsilon'/16, \mathrm{tp}^{\mathcal{N}}(\mathbf{X}') \in \mathcal{O}\}$$

which is a neighborhood of $\mathrm{tp}^{\mathcal{M}}(Y,\mathbf{X})$. Let $\epsilon = \min(\delta,\epsilon')$ and $F = F_1 \cup F_2$. We claim that:

$$K^{\mathrm{orb}}_{F',\epsilon'}(\Gamma_{r,\mathbf{r}}(\mathcal{O}')) \leq k\, K^{\mathrm{orb}}_{F,\epsilon}(\Gamma_{\mathbf{r}}(\mathcal{O}))$$

There exists an orbital (F,ϵ)–cover Ω of $\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O})$ with $\Omega \subseteq \Gamma_{\mathbf{r}}^{(n)}(\mathcal{O})$ and $|\Omega| \leq K_{F,\epsilon/2}(\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O}))$. Indeed, we can let Ω_0 be any orbital $(F,\epsilon/2)$–cover of $\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O})$ not necessarily contained in $\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O})$ and let Ω contain one point $\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O}) \cap N_{F,\epsilon/2}(\mathbf{Y})$ for each $\mathbf{Y} \in \Omega'$ where the intersection is nonempty.

For each $\mathbf{X}' \in \Gamma_{\mathbf{r}}^{(n)}(\mathcal{O})$, we have

$$\beta^{M_n(\mathbb{C})}(\mathbf{X}') < \frac{\epsilon'}{16}$$

and therefore, there exist $Y_1' \in D_{r_1}^{M_n(\mathbb{C})}, \ldots, Y_k' \in D_{r_k}^{M_n(\mathbb{C})}$ such that

$$\alpha^{M_n(\mathbb{C})}(Y_1', \ldots, Y_k', \mathbf{X}') < \frac{\epsilon'}{16}.$$

Choose for each $\mathbf{X}' \in \Omega$ a corresponding $Y_1'(\mathbf{X}'), \ldots, Y_k'(\mathbf{X}')$, and let

$$\Omega' = \{(Y_1'(\mathbf{X}'), \mathbf{X}'), \ldots, (Y_k'(\mathbf{X}'), \mathbf{X}') : \mathbf{X}' \in \Omega\}.$$

We claim that Ω' is an orbital (F',ϵ')–cover of $\Gamma_{r,\mathbf{r}}^{(n)}(\mathcal{O}')$. Let $(Y'', \mathbf{X}'') \in \Gamma_{r,\mathbf{r}}^{(n)}(\mathcal{O}')$. Then $\mathbf{X}'' \in \Gamma_{\mathbf{r}}(\mathcal{O})$. Therefore, there exists a unitary U and $\mathbf{X}' \in \Omega$ such that $U\mathbf{X}''U^* \in N_{F,\epsilon}(\mathbf{X}')$. Let $Y_1' = Y_1'(\mathbf{X}')$, $\ldots$, $Y_k' = Y_k'(\mathbf{X}')$, and note that because $d^{M_n(\mathbb{C})}(X_j', X_j'') < \delta$ for $j \in F_2$, we have

$$|\alpha^{M_n(\mathbb{C})}(Y_1', \ldots, Y_k', U\mathbf{X}''U^*) - \alpha^{M_n(\mathbb{C})}(Y_1', \ldots, Y_k', \mathbf{X}')| < \frac{\epsilon'}{16}$$

hence

$$\alpha^{M_n(\mathbb{C})}(Y_1', \ldots, Y_k', U\mathbf{X}''U^*) < \frac{\epsilon'}{16} + \frac{\epsilon'}{16} = \frac{\epsilon'}{8}.$$

By definition of α, this means that

$$\sup_{Y' \in D_r^{M_n}(\mathbb{C})} \min\Big(\epsilon'/4 - \psi^{M_n(\mathbb{C})}(Y', U\mathbf{X}''U^*), \min(d^{M_n(\mathbb{C})}(Y', Y_1'), \ldots, d^{M_n(\mathbb{C})}(Y', Y_k')) - \epsilon'/2\Big) < \frac{\epsilon'}{8}.$$

Now because $(Y'', \mathbf{X}'') \in \Gamma_{r,\mathbf{r}}^{(n)}(\mathcal{O}')$, we have

$$\psi^{M_n(\mathbb{C})}(UY''U^*, U\mathbf{X}''U^*) = \psi^{M_n(\mathbb{C})}(Y'', \mathbf{X}'') < \epsilon'/8.$$

It follows that

$$\frac{\epsilon'}{4} - \psi^{M_n(\mathbb{C})}(UY''U^*, U\mathbf{X}''U^*) > \frac{\epsilon'}{8}$$

and therefore

$$\min(d^{M_n(\mathbb{C})}(UY''U^*, Y_1'), \ldots, d^{M_n(\mathbb{C})}(UY''U^*, Y_k')) - \frac{\epsilon'}{2} < \frac{\epsilon'}{8}$$

hence $d(UY''U^*, Y_i) < \epsilon'/2 + \epsilon'/8 < \epsilon'$ for some $i \in \{1, \ldots, k\}$. Therefore, overall

$$U\mathbf{X}''U^* \in N_{F,\epsilon}(\mathbf{X}') \subseteq N_{F_1,\epsilon'}(\mathbf{X}') \text{ and } UY''U^* \in N_{\epsilon'}(Y_i')$$

and thus $(Y'', \mathbf{X}'') \in N_{F',\epsilon'}(Y_i', \mathbf{X}')$, which shows that $\Gamma^{(n)}_{r,\mathbf{r}'}(\mathcal{O}') \subseteq N^{\mathrm{orb}}_{F',\epsilon'}(\Omega')$.

We conclude that

$$K_{F',\epsilon'}(\Gamma^{(n)}_{r,\mathbf{r}}(\mathcal{O}')) \leq k\, K_{F,\epsilon/2}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O})).$$

Hence, applying $\lim_{n\to\mathcal{U}}(1/n^2)\log$ to both sides,

$$\mathrm{Ent}^{\mathcal{U}}_{(r,\mathbf{r}'),F',\epsilon'}(\mathcal{O}') \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}).$$

Because for every $\mathcal{O}$ there exists such an $\mathcal{O}'$, we obtain

$$\mathrm{Ent}^{\mathcal{U}}_{(r,\mathbf{r}),F',\epsilon'}(\mathrm{tp}^{\mathcal{M}}(Y, \mathbf{X})) \leq \mathrm{Ent}_{\mathbf{r},F,\epsilon/2}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

Then because for every (F', ϵ') there exists such an (F, ϵ), we get

$$\mathrm{Ent}^{\mathcal{U}}_{r,\mathbf{r}}(\mathrm{tp}^{\mathcal{M}}(Y, \mathbf{X})) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

Taking the supremum over $\mathbf{r}$ and r completes the proof. □

Corollary 4.21 *Let $\mathcal{M}$ be a tracial W^*–algebra and $\mathcal{N}$ a tracial W^*–subalgebra. Then*

$$\mathrm{Ent}^{\mathcal{U}}(\mathrm{acl}(\mathcal{N}) : \mathcal{M}) = \mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M}).$$

Proof The inequality $\mathrm{Ent}^{\mathcal{U}}(\mathrm{acl}(\mathcal{N}) : \mathcal{M}) \geq \mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$ holds by Observation 4.12.

On the other hand, suppose that $\mathbf{Y}$ is an $\mathbb{N}$–tuple in $\mathrm{acl}(\mathcal{N})$. Using Remark 4.18, each Y_k is algebraic over some separable W^*–subalgebra of $\mathcal{N}$. Let $\mathcal{N}_0 \subseteq \mathcal{N}$ be the join of all these subalgebras, so that $\mathcal{N}_0$ is separable and $\mathbf{Y}$ is algebraic over $\mathcal{N}_0$. Let $\mathbf{X} \in L^\infty(\mathcal{N}_0)^{\mathbb{N}}$ generate $\mathcal{N}_0$. Since Y_1 is algebraic over $\mathcal{N}_0$, we have

$$\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, Y_1)) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

Similarly, since Y_2 is algebraic over $\mathrm{W}^*(\mathbf{X}, Y_1)$, we have

$$\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, Y_1, Y_2)) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

Continuing inductively, for $k \in \mathbb{N}$,

$$\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, Y_1, \ldots, Y_k)) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

Now to analyze $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, \mathbf{Y}))$, suppose $\mathbf{r} \in (0, \infty)^{\mathcal{N} \sqcup \mathbb{N}}$ and $\epsilon > 0$ and $F \subseteq \mathbb{N} \sqcup \mathbb{N}$ is finite. Then $F \subseteq \mathbb{N} \sqcup \{1, \ldots, k\}$ for some $k \in \mathbb{N}$. For every neighborhood $\mathcal{O}$ of $\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, Y_1, \ldots, Y_k)$, there is a corresponding neighborhood $\mathcal{O}'$ of $\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, \mathbf{Y})$ given as

the preimage of $\mathcal{O}$ under map restricting the type of an $\mathbb{N} \sqcup \mathbb{N}$–tuple to the type of the $\mathbb{N} \sqcup \{1, \ldots, k\}$–subtuple. Since $F \subseteq \mathbb{N} \sqcup \{1, \ldots, k\}$, then

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}') = \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}',F,\epsilon}(\mathcal{O})$$

where $\mathbf{r}'$ is the restriction of $\mathbf{r}$ to $\mathbb{N} \sqcup \{1, \ldots, k\}$. This implies that:

$$\begin{aligned}\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, \mathbf{Y})) &\leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r}',F,\epsilon}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, Y_1, \ldots, Y_k)) \\ &\leq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, Y_1, \ldots, Y_k)) = \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))\end{aligned}$$

Since $\mathbf{r}$, F, and ϵ were arbitrary, $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, \mathbf{Y})) \leq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$. Also,

$$\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{Y})) \leq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}, \mathbf{Y}))$$

by Corollary 4.10. Since $\mathbf{Y}$ was an arbitrary $\mathbb{N}$–tuple in $\mathrm{acl}(\mathcal{N})$, we obtain $\mathrm{Ent}^{\mathcal{U}}(\mathrm{acl}(\mathcal{N}) : \mathcal{M}) \leq \mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M})$ as desired. □

5 Entropy for quantifier-free and existential types

In this section, we explain how Hayes' 1–bounded entropy (or covering entropy for non-commutative laws) relates to the entropy for types in this paper. Specifically, the 1–bounded entropy for laws corresponds is the version for quantifier-free types and the 1–bounded entropy of $\mathcal{N}$ in the presence of a larger W*–algebra $\mathcal{M}$ is the version for existential types.

5.1 Entropy for quantifier-free types

We begin with the quantifier-free version, essentially the same as orbital version of $h(\mathcal{M})$ in Hayes [16, Appendix A].

Definition 5.1 (Entropy for quantifier-free types) For $\mathcal{K} \subseteq \mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ and $\mathbf{r} \in (0, \infty)^{\mathbb{N}}$, we define:

$$\Gamma^{(n)}_{\mathbf{r}}(\mathcal{K}) = \left\{ \mathbf{X} \in \prod_{j \in \mathbb{N}} D_{r_j}^{M_n(\mathbb{C})} : \mathrm{tp}^{M_n(\mathbb{C})}_{\mathrm{qf}}(\mathbf{X}) \in \mathcal{K} \right\}$$

Then we define, for $F \subseteq \mathbb{N}$ finite and $\epsilon > 0$,

$$\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf},\mathbf{r},F,\epsilon}(\mathcal{K}) = \inf_{\mathcal{O} \supseteq \mathcal{K} \text{ open in } \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})} \lim_{n \to \mathcal{U}} \frac{1}{n^2} \log K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}))$$

and we set $$\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathcal{K}) = \sup_{\mathbf{r},F,\epsilon} \mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf},\mathbf{r},F,\epsilon}(\mathcal{K}).$$

For $\mu \in \mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$, let $\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mu) = \mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\{\mu\})$.

Some earlier works such as Hayes [16] phrased the definition of $\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mu)$ in terms of particular open sets $\mathcal{O}_k$ (for instance, those defined by looking at moments of order up to k being within some distance $1/k$ of the moments of μ), following the same conventions as Voiculescu [32] originally used. This does not change the definition because of the following lemma.

Lemma 5.2 *Let $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$. Let $\mathcal{K} \subseteq \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$. Let $(\mathcal{O}_k)_{k\in\mathbb{N}}$ be a sequence of open subsets of $\mathbb{S}(\mathrm{T}_{\mathrm{tr}})$ such that $\overline{\mathcal{O}}_{k+1} \subseteq \mathcal{O}_k$ and $\bigcap_{k=1}^{\infty} \mathcal{O}_k = \mathcal{K}$. Then*

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{K}) = \lim_{k\to\infty} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_k) = \inf_{k\in\mathbb{N}} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_k).$$

The same holds with $\mathbb{S}(\mathrm{T}_{\mathrm{tr}})$ and $\mathrm{Ent}^{\mathcal{U}}$ replaced by their quantifier-free versions.

Proof By Observation 4.3,

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{K}) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_{k+1}) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_k)$$

so that

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{K}) \leq \inf_{k\in\mathbb{N}} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_k) = \lim_{k\to\infty} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_k).$$

For the inequality in the other direction, fix $\mathcal{O} \supseteq \mathcal{K}$ open. Then $\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) \setminus \mathcal{O}$ is closed and disjoint from k. Moreover, it is contained in $\mathcal{K}^c = \bigcup_{k\in\mathbb{N}} \mathcal{O}_k^c = \bigcup_{k\in\mathbb{N}} \overline{\mathcal{O}}_k^c$. By compactness, there is a finite subcollection of $\overline{\mathcal{O}}_k^c$'s that covers $\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) \setminus \mathcal{O}$. The $\mathcal{O}_k$'s are nested, so there exists some k such that $\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) \setminus \mathcal{O} \subseteq \overline{\mathcal{O}_k}^c$, hence $\mathcal{O}_k \cap \mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) \subseteq \mathcal{O}$. Therefore,

$$\inf_{k\in\mathbb{N}} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_k) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}).$$

Since $\mathcal{O}$ was arbitrary,

$$\inf_{k\in\mathbb{N}} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{O}_k) \leq \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{K}).$$

The argument for the quantifier-free case is identical. □

This lemma also allows us to relate the entropy $\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}$ for quantifier-free types directly to the entropy for types $\mathrm{Ent}^{\mathcal{U}}$.

Lemma 5.3 *Let $\pi_{\mathrm{qf}} : \mathbb{S}(\mathrm{T}_{\mathrm{tr}}) \to \mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ be the canonical restriction map. Let $\mathcal{K} \subseteq \mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ be closed. Then*

$$\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathcal{K}) = \mathrm{Ent}^{\mathcal{U}}(\pi_{\mathrm{qf}}^{-1}(\mathcal{K})).$$

Proof Fix $\mathbf{r} \in (0,\infty)^{\mathcal{N}}$, and let $\mathcal{K}_{\mathbf{r}} = \mathcal{K} \cap \mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$. Since $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ is metrizable, there exists a sequence of open sets $\mathcal{O}_k$ in $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ such that $\overline{\mathcal{O}}_{k+1} \subseteq \mathcal{O}_k$ and $\bigcap_{k\in\mathbb{N}} \mathcal{O}_k = \mathcal{K}_{\mathbf{r}}$ (and these can be extended to open sets in $\mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ since the inclusion of $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ is a topological embedding). Now $\pi_{\mathrm{qf}}^{-1}(\mathcal{O}_k)$ is open in $\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ and $\overline{\pi_{\mathrm{qf}}^{-1}(\mathcal{O}_{k+1})} \subseteq \pi_{\mathrm{qf}}^{-1}(\overline{\mathcal{O}}_{k+1}) \subseteq \pi_{\mathrm{qf}}^{-1}(\mathcal{O}_k)$ and $\bigcap_{k\in\mathbb{N}} \pi_{\mathrm{qf}}^{-1}(\mathcal{O}_k) = \pi_{\mathrm{qf}}^{-1}(\mathcal{K}_{\mathbf{r}})$. Note that $\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O}_k) = \Gamma_{\mathbf{r}}^{(n)}(\pi_{\mathrm{qf}}^{-1}(\mathcal{O}_k))$. Thus, using the previous lemma:

$$\begin{aligned}
\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf},\mathbf{r},F,\epsilon}(\mathcal{K}) &= \mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf},\mathbf{r},F,\epsilon}(\mathcal{K}_{\mathbf{r}}) \\
&= \inf_{k\in\mathbb{N}} \mathrm{Ent}_{\mathrm{qf},\mathbf{r},F,\epsilon}(\mathcal{O}_k) \\
&= \inf_{k\in\mathbb{N}} \mathrm{Ent}_{\mathbf{r},F,\epsilon}(\pi_{\mathrm{qf}}^{-1}(\mathcal{O}_k)) \\
&= \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\pi_{\mathrm{qf}}^{-1}(\mathcal{K}_{\mathbf{r}})) \\
&= \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\pi_{\mathrm{qf}}^{-1}(\mathcal{K}))
\end{aligned}$$

Taking the supremum over $\mathbf{r}$, F, and ϵ completes the argument. □

In particular, by combining this with the variational principle (Proposition 4.6), we obtain the following corollary.

Corollary 5.4 *Let $\pi_{\mathrm{qf}} : \mathbb{S}(\mathrm{T}_{\mathrm{tr}}) \to \mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ be the restriction map. If $\mu \in \mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$, then*

$$\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mu) = \sup_{\nu\in\pi_{\mathrm{qf}}^{-1}(\mu)} \mathrm{Ent}^{\mathcal{U}}(\nu).$$

This also implies that the quantifier-free entropy of $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X})$ only depends on $\mathrm{W}^*(\mathbf{X})$, which is an important property of 1–bounded entropy previously established by Hayes in [16, Theorem A.9].

Corollary 5.5 *Let $\mathcal{M} = (M,\tau)$ be a tracial* W*–algebra. Let $\mathbf{X}$, $\mathbf{Y} \in M^{\mathbb{N}}$. If $\mathrm{W}^*(\mathbf{X}) = \mathrm{W}^*(\mathbf{Y})$, then*

$$\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X})) = \mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{Y})).$$

Proof By Proposition 3.32, there exist quantifier-free definable functions $\mathbf{f}$ and $\mathbf{g}$ such that $\mathbf{f}^{\mathcal{M}}(\mathbf{X}) = \mathbf{Y}$ and $\mathbf{g}^{\mathcal{M}}(\mathbf{Y}) = \mathbf{X}$. If $\mu \in \pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}))$, then $\mathbf{f}_*\mu \in \pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{Y}))$ since $\pi_{\mathrm{qf}} \circ \mathbf{f}_* = \mathbf{f}_* \circ \pi_{\mathrm{qf}}$. Similarly, $\mathbf{g}_*$ maps $\pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{Y}))$ into $\pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}))$.

Since $d(\mathbf{f}\circ\mathbf{g}(\mathbf{x}),\mathbf{x})$ is a quantifier-free definable predicate, its value only depends on the quantifier-free type of the input, and thus $\mathbf{g}^{\mathcal{N}} \circ \mathbf{f}^{\mathcal{N}}(\mathbf{Z}) = \mathbf{Z}$ whenever $\mathcal{N} \models \mathrm{T}_{\mathrm{tr}}$ and

$\mathrm{tp}^{\mathcal{N}}_{\mathrm{qf}}(\mathbf{Z}) = \mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X})$. In particular, if $\mu \in \pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}))$, then $\mathbf{g}_*\mathbf{f}_*\mu = \mu$. The same holds for $\mathbf{g} \circ \mathbf{f}$. Hence, $\mathbf{f}$ and $\mathbf{g}$ define mutually inverse maps between $\pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}))$ and $\pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{Y}))$. Note also that $\pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}))$ is contained in $\mathbb{S}_{\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ for some $\mathbf{r}$ by Remark 3.27. Therefore, by Proposition 4.7,

$$\mathrm{Ent}^{\mathcal{U}}(\pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{U}}_{\mathrm{qf}}(\mathbf{X}))) = \mathrm{Ent}^{\mathcal{U}}(\pi_{\mathrm{qf}}^{-1}(\mathrm{tp}^{\mathcal{U}}_{\mathrm{qf}}(\mathbf{Y})))$$

which implies the claimed result by Lemma 5.3. □

Furthermore, since the quantifier-free type does not depend on the ambient tracial W*–algebra $\mathcal{M}$, it follows that if $\mathbf{X}$ and $\mathbf{Y}$ in *different* tracial W*–algebras generate isomorphic tracial W*–algebras, then their quantifier-free types have the same entropy. Hence, it is consistent to define for a separable tracial W*–algebra $\mathcal{M}$,

$$\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathcal{M}) = \mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}))$$

where $\mathbf{X}$ is an $\mathbb{N}$–tuple of generators for $\mathcal{M}$ (for the definition of $\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathcal{M})$ in the case of non-separable $\mathcal{M}$, see Remark 5.21 below). However, Remark 5.16 shows that there is no quantifier-free analog of monotonicity under pushforward (Proposition 4.7).

5.2 Existential types

Now we turn our attention to existential types.

Definition 5.6 An *existential formula* in a language $\mathcal{L}$ is a formula of the form

$$\phi(\mathbf{x}) = \inf_{y_1 \in D_1, \ldots, y_k \in D_k} \psi(\mathbf{x}, y_1, \ldots, y_k)$$

where ψ is a quantifier-free formula and $D_1, \ldots, D_k$ are domains of quantification in the appropriate sorts. Similarly, we say that ϕ is an *existential definable predicate relative to* T if

$$\phi^{\mathcal{M}}(\mathbf{X}) = \inf_{\mathbf{Y} \in \prod_{j \in \mathbb{N}} D_j^{\mathcal{M}}} \psi^{\mathcal{M}}(\mathbf{X}, \mathbf{Y})$$

for $\mathcal{M} \models \mathrm{T}$, where ψ is a quantifier-free definable predicate.

Observation 5.7 *Any existential definable predicate can be approximated uniformly on each product of domains of quantification by an existential formula.*

Definition 5.8 Let $\mathcal{M}$ be an $\mathcal{L}$–structure, $\mathbf{S}$ an $\mathbb{N}$–tuple of sorts, and $\mathbf{X} \in \prod_{j\in\mathbb{N}} S_j^{\mathcal{M}}$. Let $\mathcal{F}_{\exists,\mathbf{S}}$ denote the space of existential formulas. The *existential type* $\mathrm{tp}_{\exists}^{\mathcal{M}}(\mathbf{X})$ is the map

$$\mathrm{tp}_{\exists}^{\mathcal{M}}(\mathbf{X}) : \mathcal{F}_{\exists,\mathbf{S}} \to \mathbb{R}, \phi \mapsto \phi^{\mathcal{M}}(\mathbf{X}).$$

If T is an $\mathcal{L}$–theory, we denote the set of existential types that arise in models of T by $\mathbb{S}_{\exists,\mathbf{S}}(\mathrm{T})$.

The topology for existential types, however, is not simply the weak-$\star$ topology on $\mathbb{S}_{\exists,\mathbf{D}}(\mathrm{T})$ for each tuple of domains. Rather, we define neighborhoods of a type $\mu = \mathrm{tp}^{\mathcal{M}}(\mathbf{X})$ using sets of the form $\{\nu : \nu(\phi) < \mu(\phi) + \epsilon\}$. The idea is that if $\phi^{\mathcal{M}}(\mathbf{X}) = \inf_{\mathbf{Y}\in\prod_{j\in\mathbb{N}} D_j^{\mathcal{M}}} \psi^{\mathcal{M}}(\mathbf{X},\mathbf{Y})$ for some quantifier-free definable predicate ϕ, then $\mu(\phi) \leq c$ means that there exists $\mathbf{Y}$ such that $\psi^{\mathcal{M}}(\mathbf{X},\mathbf{Y}) < c + \delta$ for any $\delta > 0$. Thus, a neighborhood corresponds to types ν where there exists $\mathbf{Y}$ that gets within ϵ of the infimum achieved by μ.

Definition 5.9 Let T be an $\mathcal{L}$–theory, $\mathbf{S}$ an $\mathbb{N}$–tuple of sorts, and $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$. We say that $\mathcal{O} \subseteq \mathbb{S}_{\exists,\mathbf{D}}(\mathrm{T})$ is open if for every $\mu \in \mathcal{O}$, there exist existential formulas ϕ_1, $\ldots$, ϕ_k and $\epsilon_1, \ldots, \epsilon_k > 0$ such that

$$\{\nu \in \mathbb{S}_{\exists,\mathbf{D}}(\mathrm{T}) : \nu(\phi_j) < \mu(\phi_j) + \epsilon_j \text{ for } j = 1,\ldots,k\} \subseteq \mathcal{O}.$$

Moreover, we say that $\mathcal{O} \subseteq \mathbb{S}_{\exists,\mathbf{S}}(\mathrm{T})$ is open if $\mathcal{O} \cap \mathbb{S}_{\exists,\mathbf{D}}(\mathrm{T})$ is open in $\mathbb{S}_{\exists,\mathbf{D}}(\mathrm{T})$ for all $\mathbf{D} \in \prod_{j\in\mathbb{N}} \mathcal{D}_{S_j}$.

Observation 5.10

- *Any set of the form* $\{\nu : \nu(\phi_1) < c_1, \ldots, \nu(\phi_k) < c_k\}$*, where* $\phi_1, \ldots, \phi_k$ *are existential definable predicates, is open in* $\mathbb{S}_{\exists,\mathbf{S}}(\mathrm{T})$.
- *The same holds if* ϕ_j *is an existential definable predicate rather than existential formula, since it can be uniformly approximated by existential formulas on each product of domains of quantification, hence existential definable predicates may be used in Definition 5.9 without changing the definition.*
- *The inclusion* $\mathbb{S}_{\exists,\mathbf{D}}(\mathrm{T}) \to \mathbb{S}_{\exists,\mathbf{S}}(\mathrm{T})$ *is a topological embedding since each of the basic open sets in* $\mathbb{S}_{\exists,\mathbf{D}}(\mathrm{T})$ *given by* $\nu(\phi_j) < \mu(\phi_j) + \epsilon_j$ *for* $j = 1, \ldots, k$ *extends to an open set in* $\mathbb{S}_{\exists,\mathbf{S}}(\mathrm{T})$.
- *The restriction map* $\mathbb{S}_{\mathbf{S}}(\mathrm{T}) \to \mathbb{S}_{\exists,\mathbf{S}}(\mathrm{T})$ *is continuous.*

Remark 5.11 Like the Zariski topology on the space of ideals in a commutative ring, the topology on $\mathbb{S}_{\exists,\mathbf{S}}(\mathrm{T})$ is often non-Hausdorff. For instance, the closure of a point is given by

$$\overline{\{\mu\}} = \{\nu : \nu(\phi) \geq \mu(\phi) \text{ for all } \phi \in \mathcal{F}_{\exists,\mathbf{S}}\}.$$

Meanwhile, the intersection of all neighborhoods of μ is given by

$$\mathcal{K}_\mu = \{\nu : \nu(\phi) \leq \mu(\phi) \text{ for all } \phi \in \mathcal{F}_{\exists,\mathbf{S}}\}. \tag{5–1}$$

We say that ν *extends* μ if $\nu(\phi) \leq \mu(\phi)$ for all existential formulas ϕ, which is equivalent to saying that for $\phi \in \mathcal{F}_{\exists,\mathbf{S}}$, we have $\mu(\phi) = 0$ implies that $\nu(\phi) = 0$ (since $\max(\phi - c, 0)$ is an existential formula if ϕ is). Then $\{\mu\} = \mathcal{K}_\mu$ if and only if it does have any proper extension, or it is *maximal*. These closed points correspond to existential types from existentially closed models (see Goldbring [15, §6.2]), and such maximal existential types in $\mathbb{S}_{\exists,\mathbf{D}}(\mathrm{T})$ form a compact Hausdorff space. However, our present goal is to work with general tracial W*–algebras, not only those that are existentially closed.

5.3 Entropy for existential types

Here we define the entropy for existential types which corresponds to Hayes' entropy of $\mathcal{N}$ in the presence of $\mathcal{M}$. We explain our definition in this subsection, and in the next one we relate it with Hayes' definition.

Definition 5.12 For $\mathcal{K} \subseteq \mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}})$, let

$$\Gamma^{(n)}_{\mathbf{r}}(\mathcal{K}) = \{\mathbf{X} \in \prod_{j\in\mathbb{N}} D_{r_j}^{M_n(\mathbb{C})} : \mathrm{tp}_\exists^{M_n(\mathbb{C})}(\mathbf{X}) \in \mathcal{K}\}$$

and define for $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, $F \subseteq \mathbb{N}$, finite, and $\epsilon > 0$,

$$\mathrm{Ent}^{\mathcal{U}}_{\exists,\mathbf{r},F,\epsilon}(\mathcal{K}) = \inf_{\mathcal{O}\supseteq\mathcal{K} \text{ open}} \lim_{n\to\mathcal{U}} \frac{1}{n^2} \log K^{\mathrm{orb}}_{F,\epsilon}(\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O})).$$

Then let

$$\mathrm{Ent}^{\mathcal{U}}_\exists(\mathcal{K}) = \sup_{\mathbf{r},F,\epsilon} \mathrm{Ent}^{\mathcal{U}}_{\exists,\mathbf{r},F,\epsilon}(\mathcal{K}).$$

Because of the non-Hausdorff nature of $\mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}})$, we will be content to focus on the existential entropy for an individual existential type rather than for a closed set of existential types.

Lemma 5.13 *Let $\mu \in \mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}})$, and let $\mathcal{K}_\mu$ be given by (5–1). Let $\pi : \mathbb{S}(\mathrm{T}_{\mathrm{tr}}) \to \mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}})$ be the canonical restriction map. Then*

$$\mathrm{Ent}^{\mathcal{U}}_\exists(\mu) = \mathrm{Ent}^{\mathcal{U}}\left(\pi^{-1}\left(\mathcal{K}_\mu\right)\right) = \sup_{\nu\in\pi^{-1}(\mathcal{K}_\mu)} \mathrm{Ent}^{\mathcal{U}}(\nu).$$

Proof Fix $\mathbf{r} \in (0,\infty)^{\mathbb{N}}$, $F \subseteq \mathbb{N}$ finite, and $\epsilon > 0$. If $\mathcal{O}$ is a neighborhood of μ in $\mathbb{S}_{\exists}(\mathrm{T}_{\mathrm{tr}})$, then it contains $\mathcal{K}_\mu$, and hence $\pi^{-1}(\mathcal{O})$ is a neighborhood of $\pi^{-1}(\mathcal{K}_\mu)$ in $\mathbb{S}(\mathrm{T}_{\mathrm{tr}})$. Moreover, $\Gamma^{(n)}_{\mathbf{r}}(\mathcal{O}) = \Gamma^{(n)}_{\mathbf{r}}(\pi^{-1}(\mathcal{O}))$, hence

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}\left(\pi^{-1}\left(\mathcal{K}_\mu\right)\right) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists,\mathbf{r},F,\epsilon}(\mu).$$

It remains to show the reverse inequality. Since the space of definable predicates on $\prod_{j\in\mathbb{N}} D_{r_j}$ relative to T_{tr} is separable with respect to the uniform metric, so is the space of existential definable predicates. Let $(\phi_j)_{j\in\mathbb{N}}$ be a sequence of existential definable predicates that are dense in this space. Let

$$\mathcal{O}_k = \left\{\nu \in \mathbb{S}_{\exists,\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) : \nu(\phi_j) < \mu(\phi_j) + \frac{1}{k} \text{ for } j \leq k\right\}.$$

Note that

$$\bigcap_{k\in\mathbb{N}} \mathcal{O}_k = \{\nu \in \mathbb{S}_{\exists,\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) : \nu(\phi_k) \leq \mu(\phi_k) \text{ or } k \in \mathbb{N}\} = \mathcal{K}_\mu.$$

Moreover,

$$\overline{\pi^{-1}(\mathcal{O}_{k+1})} \subseteq \left\{\nu \in \mathbb{S}(\mathrm{T}_{\mathrm{tr}}) : \nu(\phi_j) \leq \mu(\phi_j) + \frac{1}{k+1} \text{ for } j \leq k+1\right\} \subseteq \pi^{-1}(\mathcal{O}_k).$$

Therefore, by Lemma 5.2 applied to $\pi^{-1}(\mathcal{O}_k)$, we have

$$\mathrm{Ent}^{\mathcal{U}}_{\exists,\mathbf{r},F,\epsilon}(\mu) \leq \inf_{k\in\mathbb{N}} \mathrm{Ent}^{\mathcal{U}}_{\exists,\mathbf{r},F,\epsilon}(\mathcal{O}_k) = \inf_{k\in\mathbb{N}} \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\pi^{-1}(\mathcal{O}_k)) = \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\pi^{-1}(\mathcal{K}_\mu))$$

where the last equality follows from the density of $\{\phi_k : k \in \mathbb{N}\}$. Thus, $\mathrm{Ent}^{\mathcal{U}}_{\exists,\mathbf{r},F,\epsilon}(\mu) = \mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\pi^{-1}(\mathcal{K}_\mu))$. Taking the supremum over $\mathbf{r}$, F, and ϵ yields the first asserted equality $\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mu) = \mathrm{Ent}^{\mathcal{U}}(\pi^{-1}(\mathcal{K}_\mu))$. The second equality follows from the applying the variational principle (Proposition 4.6) to the closed set $\pi^{-1}(\mathcal{K}_\mu)$. □

Like the entropy for full types, the entropy for existential types satisfies a certain monotonicity under pushforwards. First, to clarify the meaning of pushforward, note that if $\mathbf{f}$ is a quantifier-free definable function and ϕ is an existential definable predicate, say

$$\phi^{\mathcal{M}}(\mathbf{X}) = \inf_{\mathbf{Y}\in\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r'_j}} \psi^{\mathcal{M}}(\mathbf{X},\mathbf{Y}) \text{ for } \mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$$

where ψ is a quantifier-free definable predicate; then

$$(\phi\circ\mathbf{f})^{\mathcal{M}}(\mathbf{X}) = \inf_{\mathbf{Y}\in\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r'_j}} \psi^{\mathcal{M}}(\mathbf{f}(\mathbf{X}),\mathbf{Y})$$

is also an existential definable predicate. Hence, there is a well-defined pushforward map $\mathbf{f}_* : \mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}}) \to \mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}})$ given by $\mathbf{f}_*\mu(\phi) = \mu(\phi \circ \mathbf{f})$. Furthermore, $\mathbf{f}_*$ is continuous with respect to the topology on $\mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}})$ for the same reason that $\phi \circ \mathbf{f}$ is an existential definable predicate whenever ϕ is an existential definable predicate and $\mathbf{f}$ is a quantifier-free definable function.

The following lemma can be proved directly in a similar way to Proposition 4.7, as was essentially done by Hayes in [16]; compare also the proof of Proposition 5.23 below. However, as one of our main goals is to illuminate the model-theoretic nature of the existential entropy, we will give an argument to deduce this from Proposition 4.7.

Lemma 5.14 *Let $\mu \in \mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}})$ and let $\mathbf{f}$ be a quantifier-free definable function relative to T_{tr}. Then*

$$\mathrm{Ent}^{\mathcal{U}}_\exists(\mathbf{f}_*\mu) \leq \mathrm{Ent}^{\mathcal{U}}_\exists(\mu).$$

Proof Let $\pi : \mathbb{S}(\mathrm{T}_{\mathrm{tr}}) \to \mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}})$ be the restriction map. Let

$$\mathcal{K} = \mathcal{K}_\mu = \{\nu \in \mathbb{S}_\exists(\mathrm{T}_{\mathrm{tr}}) : \nu(\phi) \leq \mu(\phi) \text{ for existential } \phi\}$$

and similarly, let $\mathcal{K}' = \mathcal{K}_{\mathbf{f}_*\mu}$. By Lemma 5.13:

$$\mathrm{Ent}^{\mathcal{U}}_\exists(\mu) = \mathrm{Ent}^{\mathcal{U}}_\exists(\mathcal{K}) = \mathrm{Ent}^{\mathcal{U}}(\pi^{-1}(\mathcal{K}))$$
$$\mathrm{Ent}^{\mathcal{U}}_\exists(\mathbf{f}_*\mu) = \mathrm{Ent}^{\mathcal{U}}_\exists(\mathcal{K}') = \mathrm{Ent}^{\mathcal{U}}(\pi^{-1}(\mathcal{K}'))$$

Meanwhile, by Proposition 4.7, Corollary 4.9, and Remark 3.27,

$$\mathrm{Ent}^{\mathcal{U}}(\mathbf{f}_*\pi^{-1}(\mathcal{K})) \leq \mathrm{Ent}^{\mathcal{U}}(\pi^{-1}(\mathcal{K})).$$

Therefore, it suffices to show that $\pi^{-1}(\mathcal{K}') = \mathbf{f}_*(\pi^{-1}(\mathcal{K}))$.

By continuity of the pushforward on the space of existential types, it follows that $\mathbf{f}_*(\mathcal{K}) \subseteq \mathcal{K}'$, and hence

$$\mathbf{f}_*(\pi^{-1}(\mathcal{K})) \subseteq \pi^{-1}(\mathbf{f}_*(\mathcal{K})) \subseteq \pi^{-1}(\mathcal{K}').$$

To prove the reverse inclusion, fix $\nu \in \pi^{-1}(\mathcal{K}')$. Fix $\mathbf{r}$ such that $\mathcal{K} \subseteq \mathbb{S}_{\exists,\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ and $\mathbf{r}'$ such that $\mathbf{f}$ maps $\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r_j}$ into $\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r'_j}$ for $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$. For $F \subseteq \mathbb{N}$ finite and $\phi_1, \ldots, \phi_k$ existential definable predicates, consider the definable predicate:

$$\psi^{\mathcal{M}}(\mathbf{Y}) = \inf_{\mathbf{X}\in\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r'_j}} \left[\sum_{j\in F} d^{\mathcal{M}}(f_j^{\mathcal{M}}(\mathbf{X}), Y_j) + \sum_{j=1}^{k} \max(0, \phi_j^{\mathcal{M}}(\mathbf{X}) - \mu(\phi_j))\right]$$

Then ψ is an existential definable predicate: indeed, if

$$\phi_i^{\mathcal{M}}(\mathbf{X}) = \inf_{\mathbf{Z}\in\prod_{j\in D_{r_{i,j}}}} \eta_i^{\mathcal{M}}(\mathbf{X}, \mathbf{Z})$$

where η is quantifier-free then:

$$\psi^{\mathcal{M}}(\mathbf{Y}) = \inf_{\mathbf{X}\in\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r'_j}} \inf_{\substack{\mathbf{Z}_i\in\prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r_{i,j}} \\ \text{for } i=1,\dots,k}} \left[\sum_{j\in F} d^{\mathcal{M}}(f_j^{\mathcal{M}}(\mathbf{X}), Y_j) + \sum_{i=1}^{k} \max(0, \eta_i^{\mathcal{M}}(\mathbf{X}, \mathbf{Z}_i) - \mu(\phi_i)) \right]$$

Since $\nu \in \pi^{-1}(\mathcal{K}')$, it follows that

$$nu(\psi) \leq \mathbf{f}_*\mu(\psi) = 0;$$

this last equality holds because if $\mathrm{tp}^{\mathcal{N}}_{\exists}(\mathbf{X}') = \mu$ and $\mathbf{Y}' = \mathbf{f}_*(\mathbf{X}')$, then $\psi^{\mathcal{N}}(\mathbf{Y}') = 0$ since $\mathbf{X}'$ participates in the infimum defining $\psi^{\mathcal{N}}(\mathbf{Y}')$.

Unwinding the definition of $\nu(\psi) \leq 0$, we have shown that for every $\epsilon > 0$ and $F \subseteq \mathbb{N}$ finite and $\phi_1, \dots, \phi_k$ existential definable predicates, there exist $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathbf{Y} \in L^\infty(\mathcal{M})^{\mathbb{N}}$ and $\mathbf{X} \in L^\infty(\mathcal{M})^{\mathbb{N}}$ with $\mathrm{tp}^{\mathcal{M}}(\mathbf{Y}) = \nu$ and

$$\sum_{j\in F} d^{\mathcal{M}}(f_j^{\mathcal{M}}(\mathbf{X}), Y_j) + \sum_{j=1}^{k} \max(0, \phi_j^{\mathcal{M}}(\mathbf{X}) - \mu(\phi_j)) < \epsilon.$$

Using an ultraproduct argument (or equivalently using the compactness theorem in continuous model theory, [3, Theorem 5.8], [4, Corollary 2.16]), there exists some $\mathcal{M}$, $\mathbf{X}$ and $\mathbf{Y}$ such that $\mathrm{tp}^{\mathcal{M}}(\mathbf{Y}) = \nu$, $d^{\mathcal{M}}(Y_j, f_j^{\mathcal{M}}(\mathbf{X})) = 0$, and $\phi^{\mathcal{M}}(\mathbf{X}) \leq \mu(\phi)$ for all existential definable predicates ϕ.

This implies that $\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X}) \in \mathcal{K}_\mu = \mathcal{K}$, hence $\mathrm{tp}^{\mathcal{M}}(\mathbf{X}) \in \pi^{-1}(\mathcal{K})$. Therefore, $\nu = \mathrm{tp}^{\mathcal{M}}(\mathbf{Y}) = \mathbf{f}_* \mathrm{tp}^{\mathcal{M}}(\mathbf{X}) \in \mathbf{f}_*(\pi^{-1}(\mathcal{K}))$ as desired. □

The next corollary follows from Lemma 5.14 and Proposition 3.32.

Corollary 5.15 *If* $\mathcal{M} \models \mathrm{T}_{\mathrm{tr}}$ *and* $\mathbf{X}, \mathbf{Y} \in L^\infty(\mathcal{M})^{\mathbb{N}}$ *and* $\mathrm{W}^*(\mathbf{Y}) \subseteq \mathrm{W}^*(\mathbf{X})$, *then*

$$\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}}(\mathbf{Y})) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X})).$$

In particular, if $\mathrm{W}^*(\mathbf{X}) = \mathrm{W}^*(\mathbf{Y})$, *then* $\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}}(\mathbf{Y})) = \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}}(\mathbf{X}))$.

Remark 5.16 The monotonicity property fails for the quantifier-free entropy. For instance, let $\mathcal{M}$ be the von Neumann algebra of the free group $\mathbb{F}_2$ and $\mathcal{R}$ the hyperfinite II_1 factor. Then $\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathcal{M}) = \infty$ but $\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathcal{M}\overline{\otimes}\mathcal{R}) = 0$ (by the same reasoning as in Corollary 1.4). The proof of Lemma 5.19 breaks down because if $\pi : \mathbb{S}(\mathrm{T}_{\mathrm{tr}}) \to \mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ is the restriction map, then $\pi^{-1}(\mathbf{f}_*\mu) \neq \mathbf{f}_*(\pi^{-1}(\mu))$ in general. Given a $\mathbf{Y}$ with $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{Y}) = \mathbf{f}_*\mu$, in order to show the existence of some $\mathbf{X}$ with $\mathbf{f}^{\mathcal{M}}(\mathbf{X}) \approx \mathbf{Y}$, we would have to use an existential formula in $\mathbf{Y}$.

Now come to the definition of existential entropy for $\mathcal{N} \subseteq \mathcal{M}$, which we will show in § 5.4 is equivalent to Hayes' $h(\mathcal{N} : \mathcal{M})$.

Definition 5.17 Let $\mathcal{M}$ be a tracial W*–algebra and $\mathcal{N} \subseteq \mathcal{M}$ a W*–subalgebra. Then define

$$\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{N} : \mathcal{M}) := \sup_{\mathbf{X} \in L^\infty(\mathcal{N})^{\mathbb{N}}} \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})).$$

The following is immediate from Corollary 5.15.

Corollary 5.18 *Let $\mathcal{M}$ be a tracial W*–algebra and $\mathcal{N} \subseteq \mathcal{M}$ a W*–subalgebra. If $\mathbf{X} \in L^\infty(\mathcal{N})^{\mathbb{N}}$ generates $\mathcal{N}$, then $\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{N} : \mathcal{M}) = \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X}))$.*

Lemma 5.19 *Let $\mathcal{M}_1 \subseteq \mathcal{M}_2 \subseteq \mathcal{M}_3$ be a tracial W*–algebras. Then*

$$\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{M}_1 : \mathcal{M}_3) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{M}_2 : \mathcal{M}_3) \text{ and } \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{M}_1 : \mathcal{M}_3) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{M}_1 : \mathcal{M}_2).$$

Proof The first inequality is immediate from Definition 5.17. For the second inequality, note that for every $\mathbf{X} \in L^\infty(\mathcal{M}_1)^{\mathbb{N}}$ and every existential formula ϕ, we have $\phi^{\mathcal{M}_3}(\mathbf{X}) \leq \phi^{\mathcal{M}_2}(\mathbf{X})$ since the first is the infimum over a larger set than the second. In other words, $\mathrm{tp}^{\mathcal{M}_3}_{\exists}(\mathbf{X})$ is an extension of $\mathrm{tp}^{\mathcal{M}_2}_{\exists}(\mathbf{X})$, and hence every neighborhood of $\mathrm{tp}^{\mathcal{M}_2}_{\exists}(\mathbf{X})$ is also a neighborhood of $\mathrm{tp}^{\mathcal{M}_3}_{\exists}(\mathbf{X})$. This implies that $\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}_3}_{\exists}(\mathbf{X})) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}_2}_{\exists}(\mathbf{X}))$. Since this holds for all $\mathbf{X} \in L^\infty(\mathcal{M})^{\mathbb{N}}$, we obtain $\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{M}_1 : \mathcal{M}_3) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{M}_1 : \mathcal{M}_2)$. □

Next, we show that the quantifier-free entropy can be expressed in terms of the existential entropy.

Lemma 5.20 *Let $\mathcal{M}$ be a separable tracial W*–algebra. Then*

$$\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathcal{M}) = \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{M} : \mathcal{M}).$$

Proof Suppose $\mathbf{X} \in L^\infty(\mathcal{M})^{\mathbb{N}}$ generates $\mathcal{M}$. Fix $\mathbf{r}$ such that $\mathbf{X} \in \prod_{j \in \mathbb{N}} D^{\mathcal{M}}_{r_j}$. Let $\pi : \mathbb{S}_{\exists}(\mathrm{T}_{\mathrm{tr}}) \to \mathbb{S}_{\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$ be the restriction map. It suffices to show that $\mathrm{Ent}_{\exists,\mathbf{r},F,\epsilon}(\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})) = \mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf},\mathbf{r},F,\epsilon}(\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}))$ for all $\mathbf{r}$, F, and ϵ, which in turn will follow if we prove that every neighborhood $\mathcal{O}'$ of $\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})$ in $\mathbb{S}_{\exists,\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ contains $\pi^{-1}(\mathcal{O})$ for some neighborhood $\mathcal{O}$ of $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X})$ in $\mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}})$ and vice versa.

Let $\mathcal{O}$ be a neighborhood of $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X})$. By the definition / properties of the weak-$\star$ topology, there exist some quantifier-free definable predicates $\phi_1, \ldots, \phi_k$ and intervals (a_k, b_k) such that

$$\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}) \in \{\nu \in \mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T}_{\mathrm{tr}}) : \nu(\phi_i) \in (a_i, b_i) \text{ for } i = 1, \ldots, k\} \subseteq \mathcal{O}.$$

Then ϕ_i and $-\phi_i$ are both existential definable predicates, hence

$$\pi^{-1}(\mathcal{O}) \supseteq \mathcal{O}' := \{\nu \in \mathbb{S}_{\exists,\mathbf{r}}(\mathrm{T_{tr}}) : \nu(\phi_i) < b_i, \nu(-\phi_k) < -a_i \text{ for } i = 1,\dots,k\}.$$

Conversely, let $\mathcal{O}'$ be a neighborhood of $\mathrm{tp}_{\exists}^{\mathcal{M}}(\mathbf{X})$ in $\mathbb{S}_{\exists,\mathbf{r}}(\mathbf{X})$. Then there exists existential definable predicates $\phi_1,\dots,\phi_k$ and $c_1,\dots,c_k \in \mathbb{R}$ such that

$$\mathrm{tp}_{\exists}^{\mathcal{M}}(\mathbf{X}) \in \{\nu \in \mathbb{S}_{\exists,\mathbf{r}}(\mathrm{T_{tr}}) : \nu(\phi_i) < c_i \text{ for } i = 1,\dots,k\} \subseteq \mathcal{O}'.$$

Suppose that

$$\phi_i^{\mathcal{N}}(\mathbf{X}') = \inf_{\mathbf{Y}' \in \prod_{j\in\mathbb{N}} D_{r_{i,j}}^{\mathcal{N}}} \psi_i^{\mathcal{N}}(\mathbf{X}', \mathbf{Y}')$$

for all $\mathcal{N} \models \mathrm{T_{tr}}$ and $\mathbf{X}' \in L^\infty(\mathcal{N})^{\mathbb{N}}$, where ψ_i is quantifier-free. Because $\phi_i^{\mathcal{M}}(\mathbf{X}) < c_i$, there exists $\mathbf{Y}_i \in \prod_{j\in\mathbb{N}} D_{r_{i,j}}^{\mathcal{M}}$ with $\psi_i^{\mathcal{M}}(\mathbf{X}, \mathbf{Y}_i) < c_i$. By Proposition 3.32, there exists a quantifier-free definable function $\mathbf{f}_i$ such that $\mathbf{Y}_i = \mathbf{f}_i^{\mathcal{M}}(\mathbf{X})$ and $\mathbf{f}_i^{\mathcal{N}}$ maps into $\prod_{j\in\mathbb{N}} D_{r_{i,j}}^{\mathcal{N}}$ for all $\mathcal{N} \models \mathrm{T_{tr}}$. Let

$$\eta_i^{\mathcal{N}}(\mathbf{X}') = \psi_i^{\mathcal{N}}(\mathbf{X}', \mathbf{f}^{\mathcal{N}}(\mathbf{X}')) \geq \phi_i^{\mathcal{N}}(\mathbf{X}').$$

Then η_i is quantifier-free. Thus,

$$\mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{X}) \in \mathcal{O} := \{\nu \in \mathbb{S}_{\mathrm{qf},\mathbf{r}}(\mathrm{T_{tr}}) : \nu(\eta_i) < c_i\}$$

and $$\pi^{-1}(\mathcal{O}) \subseteq \{\nu \in \mathbb{S}_{\exists,\mathbf{r}}(\mathrm{T_{tr}}) : \nu(\phi_i) < c_i \text{ for } i = 1,\dots,k\} \subseteq \mathcal{O}'$$

as desired. □

Remark 5.21 Therefore, it is natural to define $\mathrm{Ent}_{\mathrm{qf}}^{\mathcal{U}}(\mathcal{M})$ for general (not necessarily separable $\mathcal{M}$) by $\mathrm{Ent}_{\mathrm{qf}}^{\mathcal{U}}(\mathcal{M}) := \mathrm{Ent}_{\exists}^{\mathcal{U}}(\mathcal{M} : \mathcal{M})$.

5.4 Existential entropy and entropy in the presence

Let us finally explain why the existential entropy defined here agrees with (the ultrafilter version of) Hayes' 1–bounded entropy of $\mathcal{N}$ in the presence of $\mathcal{M}$ in [16]. The definition is given in terms of Voiculescu's microstate spaces for some $\mathbf{X}$ in the presence of $\mathbf{Y}$ from [33].

Definition 5.22 (Hayes [16]) Let $\mathcal{M}$ be a tracial W*–algebra. Let I and J be arbitrary index sets and let $\mathbf{X} \in L^\infty(\mathcal{M})^I$ and $\mathbf{Y} \in L^\infty(\mathcal{M})^J$. Let $\mathbf{r} \in (0,\infty)^I$ and $\mathbf{r}' \in (0,\infty)^J$ such that $\|X_j\| \leq r_j$ and $\|Y_j\| \leq r_j'$. Let $\mathbb{S}_{\mathbf{r},\mathbf{r}',\mathrm{qf}}(\mathrm{T_{tr}})$ be the set of quantifier-free

types of tuples from $\prod_{i\in I} D_{r_i} \times \prod_{j\in J} D_{r'_j}$ equipped with the weak-$\star$ topology. Let $p : M_n(\mathbb{C})^{I\sqcup J} \to M_n(\mathbb{C})^I$ be the canonical coordinate projection. Then we define

$$h^{\mathcal{U}}_{\mathbf{r},\mathbf{r}'}(\mathbf{X} : \mathbf{Y}) := \sup_{\epsilon>0} \sup_{F\subseteq I \text{ finite}} \inf_{\mathcal{O}\ni \mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X},\mathbf{Y})} \lim_{n\to\mathcal{U}} \frac{1}{n^2} \log K^{\mathrm{orb}}_{F,\epsilon}(p[\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'}(\mathcal{O})])$$

where $\mathcal{O}$ ranges over all neighborhoods of $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X},\mathbf{Y})$ in $\mathbb{S}_{\mathbf{r},\mathbf{r}',\mathrm{qf}}(\mathrm{T}_{\mathrm{tr}})$.

Here we use arbitrary index sets I and J rather than $\mathbb{N}$ because we do not assume that $\mathcal{M}$ is separable. This is a technical issue we will have to consider when proving that our definition using $\mathbb{N}$–tuples agrees with Hayes.' Apart from that, the idea of the proof is that a matrix tuple $\mathbf{X}'$ is in the projection $p[\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'}(\mathcal{O})]$ if and only if *there exists* some $\mathbf{Y}'$ such that $\mathrm{tp}^{M_n(\mathbb{C})}_{\mathrm{qf}}(\mathbf{X}',\mathbf{Y}') \in \mathcal{O}$. If $\mathbf{X}',\mathbf{Y}'$ being in $\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'}(\mathcal{O})$ can be detected by a quantifier-free formula being less than some c (using Urysohn's lemma), then $\mathbf{X}'$ being in $p[\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'}(\mathcal{O})]$ can be detected by an existential formula.

Proposition 5.23 *In the setup of Definition 5.22, we have* $h^{\mathcal{U}}_{\mathbf{r},\mathbf{r}'}(\mathbf{X} : \mathbf{Y}) = \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{W}^*(\mathbf{X}) : \mathrm{W}^*(\mathbf{X},\mathbf{Y}))$.

We remark at the start of the proof that all the facts we proved about definable predicates and functions work for arbitrary index sets, so long as they do not invoke metrizability of the type space. We also leave some details to the reader for the sake of space.

Proof We may assume without loss of generality that $\mathcal{M} = \mathrm{W}^*(\mathbf{X},\mathbf{Y})$ since restricting to a smaller W*–algebra does not change the quantifier-free type of $(\mathbf{X},\mathbf{Y})$.

First, let us show that $h^{\mathcal{U}}_{\mathbf{r},\mathbf{r}'}(\mathbf{X} : \mathbf{Y}) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{W}^*(\mathbf{X}) : \mathcal{M})$. Let $F \subseteq I$ finite and $\epsilon > 0$. First, to deal with changing index sets from I to $\mathbb{N}$, let $\alpha : F \to \mathbb{N}$ be an injective function and let $\mathbf{f}$ be the quantifier-free definable function that sends an I–tuple $\mathbf{X}'$ to the $\mathbb{N}$–tuple obtained by putting X'_j into the $\alpha(j)$th entry for $j \in F$ and fills the other entries with zeros. Let $\mathbf{Z} = \mathbf{f}^{\mathcal{M}}(\mathbf{X})$, fix some $\mathrm{T} \in (0,\infty)^{\mathbb{N}}$ with $\mathbf{Z} \in \prod_{j\in\mathbb{N}} D^{\mathcal{N}}_{t_j}$, and let $\mathcal{O}$ be a neighborhood of $\mu = \mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{Z})$ in $\mathbb{S}_{\exists,\mathrm{T}}(\mathrm{T}_{\mathrm{tr}})$. Then there exist existential definable predicates $\phi_1,\ldots,\phi_k$ and $\epsilon_1,\ldots,\epsilon_k > 0$ such that

$$\{\nu \in \mathbb{S}_{\exists,\mathrm{T}}(\mathrm{T}_{\mathrm{tr}}) : \nu(\phi_j) \leq \mu(\phi_j) + \epsilon_j \text{ for } j = 1,\ldots,k\} \subseteq \mathcal{O}.$$

There exist quantifier-free definable predicates $\psi_1,\ldots,\psi_k$ such that

$$\phi^{\mathcal{N}}_j(\mathbf{Z}') = \inf_{\mathbf{W}'\in\prod_{i\in\mathbb{N}} D^{\mathcal{N}}_{t_{i,j}}} \psi^{\mathcal{N}}_j(\mathbf{Z}',\mathbf{W}')$$

for all $\mathcal{N}$ and $\mathbf{Z}' \in L^\infty(\mathcal{N})^{\mathbb{N}}$. Moreover, for our particular $\mathcal{M}$ and $\mathbf{Z}$, there exists $\mathbf{W}_j \in \prod_{i\in\mathbb{N}} D^{\mathcal{M}}_{t_{i,j}}$ such that

$$\psi_j^{\mathcal{M}}(\mathbf{Z}, \mathbf{W}_j) < \mu(\phi_j) + \epsilon_j.$$

By Proposition 3.32, $\mathbf{W}_j = \mathbf{g}_j(\mathbf{X}, \mathbf{Y})$ for some quantifier-free definable function $\mathbf{g}_j$. Let $\psi_j(\mathbf{f}, \mathbf{g}_j)$ denote the quantifier-free definable predicate defined for $I \sqcup J$–tuples by applying $\mathbf{f}$ to the I–tuple and $\mathbf{g}_j$ to the $I \sqcup J$–tuple and then applying ψ_j. Then

$$\mathcal{O}' := \bigcap_{j=1}^{k} \{\sigma \in \mathbb{S}_{\mathrm{qf},\mathbf{r},\mathbf{r}'}(\mathrm{T}_{\mathrm{tr}}) : \sigma(\psi_j(\mathbf{f}, \mathbf{g}_j)) < \mu(\phi_j) + \epsilon_j\}$$

is a neighborhood of $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}, \mathbf{Y})$ in $\mathbb{S}_{\mathrm{qf},\mathbf{r},\mathbf{r}'}(\mathrm{T}_{\mathrm{tr}})$ such that

$$p[\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'}(\mathcal{O}')] \subseteq (\mathbf{f}^{M_n(\mathbb{C})})^{-1}[\Gamma^{(n)}_{\mathrm{T}}(\mathcal{O})].$$

Therefore,

$$K^{\mathrm{orb}}_{F,\epsilon}(p[\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'}(\mathcal{O}')]) \leq K^{\mathrm{orb}}_{\alpha(F),\epsilon}(\Gamma^{(n)}_{\mathrm{T}}(\mathcal{O})).$$

Because for every such $\mathcal{O}$, there exists such an $\mathcal{O}'$, we obtain that

$$\inf_{\mathcal{O}' \ni \mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X},\mathbf{Y})} \lim_{n\to\mathcal{U}} \frac{1}{n^2} \log K^{\mathrm{orb}}_{F,\epsilon}(p[\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'}(\mathcal{O}')]) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists,\alpha(F),\epsilon}(\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{f}(\mathbf{X}))) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{W}^*(\mathbf{X}) : \mathcal{M}).$$

Since F and ϵ were arbitrary, we are done with the first inequality.

To prove the second inequality, we must show that for all $\mathbf{Z} \in \mathrm{W}^*(\mathbf{X})^{\mathbb{N}}$, we have $\mathrm{Ent}_{\exists}(\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{Z})) \leq h_{\mathbf{r},\mathbf{r}'}(\mathbf{X} : \mathbf{Y})$. Fix $\mathbf{Z}$, let $\mathrm{T} \in (0,\infty)^{\mathbb{N}}$ with $\|Z_j\| \leq t_j$, and write $\mathbf{Z} = \mathbf{f}^{\mathcal{M}}(\mathbf{X})$ for some quantifier-free definable function $\mathbf{f}$ depending on countably many coordinates of $\mathbf{X}$. Let $\mathcal{O}'$ be a neighborhood of $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}, \mathbf{Y})$. Note that $\mathcal{O}'$ contains a neighborhood of μ that depends only on finitely many coordinates of $\mathbf{X}$ and $\mathbf{Y}$. By Urysohn's lemma and Remark 3.33, there exists a quantifier-free definable predicate ψ with values in $[0,1]$ (depending on only finitely many coordinates) such that $\psi^{\mathcal{M}}(\mathbf{X}, \mathbf{Y}) = 0$ and $\{\sigma \in \mathbb{S}_{\mathrm{qf},\mathbf{r},\mathbf{r}'}(\mathrm{T}_{\mathrm{tr}}) : \sigma(\psi) < 1\} \subseteq \mathcal{O}'$.

Fix $F \subseteq \mathbb{N}$ finite and $\epsilon \in (0,2)$, and consider the existential formula:

$$\phi^{\mathcal{N}}(\mathbf{Z}') = \inf_{\mathbf{X}' \in \prod_{i\in I} D^{\mathcal{N}}_{r_i}} \inf_{\mathbf{Y}' \in \prod_{j\in J} D^{\mathcal{N}}_{r'_j}} \left(\sum_{k\in F} d^{\mathcal{N}}(\mathbf{f}^{\mathcal{N}}_k(\mathbf{X}'), \mathbf{Z}'_k) + \psi^{\mathcal{N}}(\mathbf{X}', \mathbf{Y}') \right)$$

Because $\mathbf{f}$ and ψ only depend on countably many coordinates, the infima can be expressed using only countably many variables, so this expression is a valid existential definable predicate. Moreover, note that for $\mathbf{Z}' \in \prod_{k\in\mathbb{N}} D^{M_n(\mathbb{C})}_{t_k}$, we have:

$$\phi^{M_n(\mathbb{C})}(\mathbf{Z}') < \frac{\epsilon}{2} \implies \mathbf{Z}' \in N_{\epsilon/2}(\mathbf{f}^{M_n(\mathbb{C})} \circ p(\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'})(\mathcal{O}'))$$

Let $\mathcal{O} = \{\mu \in \mathbb{S}_{\exists,\mathrm{T}}(\mathrm{T}_{\mathrm{tr}}) : \mu(\phi) < \epsilon/2\}$. By applying the uniform continuity of $\mathbf{f}$ (Lemma 3.19) for the given $\epsilon/2$ and F in the target space, we obtain a corresponding $F' \subseteq I$ and $\delta > 0$ such that

$$K^{\mathrm{orb}}_{\epsilon,F}(\Gamma^{(n)}_{\mathrm{T}}(\mathcal{O})) \leq K^{\mathrm{orb}}_{\epsilon/2,F}(\mathbf{f}^{M_n(\mathbb{C})} \circ p(\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'})(\mathcal{O}')) \leq K^{\mathrm{orb}}_{\delta,F'}(p(\Gamma^{(n)}_{\mathbf{r},\mathbf{r}'}(\mathcal{O}'))).$$

Applying the definitions of the appropriate limits, suprema, and infima shows that $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{Z})) \leq h_{\mathbf{r},\mathbf{r}'}(\mathbf{X} : \mathbf{Y})$. □

5.5 Applications to ultraproduct embeddings

Theorem 5.24 *Let $\mathcal{N} \subseteq \mathcal{M}$ be separable tracial W^*–algebras, and let $\mathcal{Q} = \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$. Suppose that $\mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{N} : \mathcal{M}) \geq 0$. Then for every $c < \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{N} : \mathcal{M})$, there exists an embedding $\iota : \mathcal{M} \to \mathcal{Q}$ such that*

$$\mathrm{Ent}^{\mathcal{U}}_{\exists}(\iota(\mathcal{N}) : \mathcal{Q}) \geq \mathrm{Ent}^{\mathcal{U}}(\iota(\mathcal{N}) : \mathcal{Q}) > c.$$

Proof Let $\mathbf{X}$ be an $\mathbb{N}$–tuple of generators for $\mathcal{N}$. Let $\pi : \mathbb{S}(\mathrm{T}_{\mathrm{tr}}) \to \mathbb{S}_{\exists}(\mathrm{T}_{\mathrm{tr}})$ be the restriction map. By Lemma 5.13,

$$c < \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})) = \sup_{\mu \in \pi^{-1}(\mathcal{K}_{\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})})} \mathrm{Ent}^{\mathcal{U}}(\mu)$$

so there exists a type μ such that $\pi(\mu) \in \mathcal{K}_{\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})}$ and $\mathrm{Ent}^{\mathcal{U}}(\mu) > c$. By Lemma 4.14, there exists $\mathbf{X}' \in \mathcal{Q}$ with $\mathrm{tp}^{\mathcal{Q}}(\mathbf{X}') = \mu$. As in Corollary 4.15, there exists an embedding $\iota : \mathcal{N} \to \mathcal{Q}$ with $\iota(\mathbf{X}) = \mathbf{X}'$. Observe that

$$\begin{aligned}\mathrm{Ent}^{\mathcal{U}}_{\exists}(\iota(\mathcal{N}) : \mathcal{Q}) = \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathrm{tp}^{\mathcal{Q}}_{\exists}(\mathbf{X}')) &\geq \mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{Q}}(\mathbf{X}')) \\ &= \mathrm{Ent}^{\mathcal{U}}(\iota(\mathcal{N}) : \mathcal{Q}) = \mathrm{Ent}^{\mathcal{U}}(\mu) > c\end{aligned}$$

where we apply in order Corollary 5.18, Lemma 5.13, Observation 4.12, and the choice of μ and $\mathbf{X}'$.

It only remains to show that ι extends to an embedding of $\mathcal{M}$. Let $\mathbf{Y} \in L^\infty(\mathcal{M})^{\mathbb{N}}$ be a set of generators. Let $\mathbf{r}$ and $\mathbf{r}'$ be such that $\mathbf{X} \in \prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r_j}$ and $\mathbf{Y} \in \prod_{j\in\mathbb{N}} D^{\mathcal{M}}_{r'_j}$. Since the quantifier-free type space $\mathbb{S}_{\mathrm{qf},\mathbf{r},\mathbf{r}'}(\mathrm{T}_{\mathrm{tr}})$ for $\mathbb{N} \sqcup \mathbb{N}$–tuples is metrizable, there exists a nonnegative continuous function on $\mathbb{S}_{\mathrm{qf},\mathbf{r},\mathbf{r}'}(\mathrm{T}_{\mathrm{tr}})$ that equals zero at and only at the point $\mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X},\mathbf{Y})$. By Remark 3.33, this continuous function extends to a global quantifier-free definable predicate ϕ. Let ψ be the existential predicate given by

$$\psi^{\mathcal{N}}(\mathbf{Z}) = \inf_{\mathbf{W} \in \prod_{j\in\mathbb{N}} D^{\mathcal{N}}_{r'_j}} \phi^{\mathcal{N}}(\mathbf{Z},\mathbf{W})$$

for $\mathcal{N} \models \mathrm{T}_{\mathrm{tr}}$ and $\mathbf{Z} \in L^\infty(\mathcal{N})^{\mathbb{N}}$. Thus, $\psi^{\mathcal{M}}(\mathbf{X}) \leq \phi^{\mathcal{M}}(\mathbf{X}, \mathbf{Y}) = 0$.

Because ψ is existential and $\mathrm{tp}^{\mathcal{Q}}_{\exists}(\mathbf{X}') \in \mathcal{K}_{\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})}$, we have $\psi^{\mathcal{Q}}(\mathbf{X}') \leq \psi^{\mathcal{M}}(\mathbf{X}) = 0$. We may write $\mathbf{X}' = [\mathbf{X}^{(n)}]_{n\in\mathbb{N}}$ where $\mathbf{X}^{(n)} \in \prod_{j\in\mathbb{N}} D_{r_j}^{M_n(\mathbb{C})}$ (this follows for instance from the construction of $\mathbf{X}'$ through Lemma 4.14). Then $\lim_{n\to\mathcal{U}} \psi^{M_n(\mathbb{C})}(\mathbf{X}^{(n)}) = \psi^{\mathcal{Q}}(\mathbf{X}') = 0$, hence there exists $\mathbf{Y}^{(n)} \in \prod_{j\in\mathbb{N}} D_{r'_j}^{M_n(\mathbb{C})}$ such that $\lim_{n\to\mathcal{U}} \phi^{M_n(\mathbb{C})}(\mathbf{X}^{(n)}, \mathbf{Y}^{(n)}) = 0$.

Let $\mathbf{Y}' = [\mathbf{Y}^{(n)}]_{n\in\mathbb{N}} \in \prod_{j\in\mathbb{N}} D_{r'_j}^{\mathcal{Q}}$. Then $\phi^{\mathcal{Q}}(\mathbf{X}', \mathbf{Y}') = 0$, and therefore, $\mathrm{tp}^{\mathcal{Q}}_{\mathrm{qf}}(\mathbf{X}', \mathbf{Y}') = \mathrm{tp}^{\mathcal{M}}_{\mathrm{qf}}(\mathbf{X}, \mathbf{Y})$. Hence, by Lemma 3.34, there exists an embedding $\iota' : \mathcal{M} \to \mathcal{Q}$ with $\iota'(\mathbf{X}, \mathbf{Y}) = (\mathbf{X}', \mathbf{Y}')$. This is the desired extension of ι. □

Remark 5.25 Note that $\mathrm{Ent}^{\mathcal{U}}_{\exists}(\iota(\mathcal{N}) : \mathcal{Q}) \leq \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{N} : \mathcal{M})$ for any such embedding ι. Thus, the point of the theorem is that *some* ι can be chosen to make this inequality close to an equality. It is not obvious that there is an existential type in $\mathcal{Q}$ extending the $\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})$ with close to the same amount of entropy of μ. The key ingredient is the variational principle (Proposition 4.6) applied through Lemma 5.13, which gives us not only an existential type $\mathrm{tp}^{\mathcal{Q}}_{\exists}(\mathbf{X}')$ extending $\mathrm{tp}^{\mathcal{M}}_{\exists}(\mathbf{X})$ with large entropy, but even the full type $\mathrm{tp}^{\mathcal{Q}}(\mathbf{X}')$ with large entropy.

In particular, the theorem shows that if $\mathrm{Ent}^{\mathcal{U}}_{\mathrm{qf}}(\mathcal{M}) > 0$, then there exists an embedding of $\mathcal{M}$ into $\mathcal{Q}$ with $\mathrm{Ent}^{\mathcal{U}}(\mathcal{M} : \mathcal{Q}) > 0$, and hence also $h(\mathcal{M} : \mathcal{Q}) = \mathrm{Ent}^{\mathcal{U}}_{\exists}(\mathcal{M} : \mathcal{Q}) > 0$.

6 Remarks on conditional entropy

In this section, we sketch how the previous results could be adapted to the setting of entropy relative to a W*–subalgebra. However, we will not give the arguments in detail because we will not be giving any new applications of the conditional version of entropy. Our goal is mainly to complete our translation between the different flavors of microstate spaces in free entropy theory and the different flavors of types in the conditional setting.

Hayes' original definition of 1–bounded entropy used microstate spaces relative to a fixed microstate sequence for some self-adjoint element with diffuse spectrum. He then showed that this was equivalent to the 1–bounded entropy defined through unitary orbits (the definition that we have used so far in this paper). As remarked by Hayes, Jekel, and Kunnawalkam Elayavalli [17, §4.1], the same reasoning shows that orbital 1–bounded entropy is equivalent to 1–bounded entropy relative to fixed microstates for any diffuse amenable W*–subalgebra $\mathcal{P}$ of $\mathcal{M}$. In fact, one can formulate the definition

of 1–bounded entropy of $\mathcal{M}$ relative to any W*–subalgebra $\mathcal{A}$ with a fixed choice of microstates $\mathbf{Y}^{(n)}$ with $\lim_{n\to\mathcal{U}} \mathrm{tp}_{\mathrm{qf}}^{M_n(\mathbb{C})}(\mathbf{Y}^{(n)}) = \mathrm{tp}_{\mathrm{qf}}^{\mathcal{M}}(\mathbf{Y})$. Unlike the case where $\mathcal{A}$ is amenable, the 1–bounded entropy relative to $\mathcal{A}$ may, as far as we know, depend on the choice of microstates for $\mathcal{A}$, and in general it will not coincide with the orbital 1–bounded entropy. Nonetheless, relative 1–bounded entropy for general $\mathcal{A}$ has a natural motivation in terms of ultraproduct embeddings: Fixing $\mathcal{A} \subseteq \mathcal{N} \subseteq \mathcal{M}$ and an embedding $\iota : \mathcal{A} \to \mathcal{Q} := \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$, a relative 1–bounded entropy $h(\mathcal{N} : \mathcal{M}|\mathcal{A}, \iota)$ would quantify the amount of embeddings of $\iota' : \mathcal{N} \to \mathcal{Q}$ that extend ι and which admit some extension $\iota'' : \mathcal{M} \to \mathcal{Q}$.

Just as we have interpreted the entropy in the presence as corresponding to existential types in the model-theoretic framework, relative entropy naturally corresponds to types over $\mathcal{A}$. Types over $\mathcal{A}$ represent types in a language $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$ described as follows. Let $\mathcal{A}$ be a separable tracial W*–algebra. Let $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$ be the language obtained by adding to $\mathcal{L}_{\mathrm{tr}}$ a constant symbol $\alpha(a) \in D_{\|a\|}$ for each $a \in \mathcal{A}$.

Let $\mathrm{T}_{\mathrm{tr},\mathcal{A}}$ be the $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$ theory obtained from T_{tr} by adding the (infinite family of) axioms

- $\alpha(a+b) = \alpha(a) + \alpha(b)$ for each $a, b \in \mathcal{A}$.
- $\alpha(\lambda a) = \lambda\alpha(a)$ for $a \in \mathcal{A}$ and $\lambda \in \mathbb{C}$.
- $\alpha(ab) = \alpha(a)\alpha(b)$ for $a, b \in \mathcal{A}$.
- $\alpha(a^*) = \alpha(a)^*$ for $a \in \mathcal{A}$.
- $\alpha(1) = 1$.
- $\mathrm{Re}\,\mathrm{tr}\,\alpha(a) = \tau_{\mathcal{A}}(a)$ where $\tau_{\mathcal{A}}$ is the given trace on the tracial W*–algebra $\mathcal{A}$.

We leave it as an exercise to the reader to verify that every model of $\mathrm{T}_{\mathrm{tr},\mathcal{A}}$ is given by a tracial W*–algebra $\mathcal{M}$ together with an embedding (unital, trace-preserving $*$–homomorphism) $\alpha : \mathcal{A} \to \mathcal{M}$, and conversely every such embedding defines a model of $\mathrm{T}_{\mathrm{tr},\mathcal{A}}$. Given a tracial W*–algebra $\mathcal{M}$ and an inclusion $\alpha : \mathcal{A} \to \mathcal{M}$, the $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$–type of a tuple $\mathbf{X}$ is also known as the *type of* $\mathbf{X}$ *over* $\mathcal{A}$ and denoted $\mathrm{tp}^{\mathcal{M}}(\mathbf{X}/\mathcal{A})$.

Next, we want to define versions of entropy for quantifier-free, full, and existential types over $\mathcal{A}$, using covering numbers for microstate spaces corresponding to neighborhoods of the type over $\mathcal{A}$. Unfortunately, we cannot use neighborhoods in the space of $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$–types $\mathbb{S}_{\mathcal{A}}(\mathrm{T}_{\mathrm{tr},\mathcal{A}})$ because the matrix algebra $M_n(\mathbb{C})$ could never be a model of $\mathrm{T}_{\mathrm{tr},\mathcal{A}}$ since it cannot contain a copy of $\mathcal{A}$ unless $\mathcal{A}$ is finite-dimensional. In other words, the issue is that we must work with approximate embeddings $\alpha_n : \mathcal{A} \to M_n(\mathbb{C})$ rather than literal embeddings, since the latter may not exist. Thus, we will look at $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$ structures that satisfy T_{tr} but not necessarily $\mathrm{T}_{\mathrm{tr},\mathcal{A}}$, which are tracial von Neumann algebras together with a function $\alpha : \mathcal{A} \to \mathcal{M}$ that is not necessarily is a $*$–homomorphism or even linear but does satisfy $\|\alpha(a)\| \leq \|a\|$ for $a \in \mathcal{A}$. We

will denote by $\mathbb{S}_{\mathcal{A}}(T_{\mathrm{tr}})$ the set of $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$–types that arise from models of T_{tr}, so that $\mathbb{S}_{\mathcal{A}}(\mathrm{T}_{\mathrm{tr}}) \supseteq S_{\mathcal{A}}(\mathrm{T}_{\mathrm{tr},\mathcal{A}})$.

Given a sequence of functions $\alpha_n : \mathcal{A} \to M_n(\mathbb{C})$ and $\mathcal{O} \subseteq S_{\mathcal{A}}(T_{\mathrm{tr}})$, we define the microstate space

$$\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O} \mid \alpha_n) = \{\mathbf{X} \in M_n(\mathbb{C})^{\mathbb{N}} : \mathrm{tp}^{M_n(\mathbb{C}),\alpha_n}(\mathbf{X}) \in \mathcal{O}\}$$

where $\mathrm{tp}^{M_n(\mathbb{C}),\alpha_n}(\mathbf{X})$ is the $\mathcal{L}_{\mathcal{A}}$ type of $\mathbf{X}$ in the $\mathcal{L}_{\mathcal{A}}$ structure given by $M_n(\mathbb{C})$ and $\alpha_n : \mathcal{A} \to M_n(\mathbb{C})$. We are interested only in the case when $(\alpha_n)_{n\in\mathbb{N}}$ defines a trace-preserving $*$–homomorphism $\alpha : \mathcal{A} \to \mathcal{Q} = \prod_{n\to\mathcal{U}} M_n(\mathbb{C})$. Then for a closed set $\mathcal{K} \subseteq \mathbb{S}_{\mathcal{A}}(\mathrm{T}_{\mathrm{tr},\mathcal{A}}) \subseteq \mathbb{S}_{\mathcal{A}}(\mathrm{T}_{\mathrm{tr}})$, we define

$$\mathrm{Ent}^{\mathcal{U}}_{\mathbf{r},F,\epsilon}(\mathcal{K} \mid \alpha) = \inf_{\mathcal{O}\supseteq\mathcal{K} \text{ open}} \lim_{n\to\mathcal{U}} \frac{1}{n^2} \log K_{F,\epsilon}(\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O} \mid \alpha)$$

where the infimum is over all open neighborhoods of $\mathcal{K}$ in $S_{\mathcal{A}}(\mathrm{T}_{\mathrm{tr}})$, and then let $\mathrm{Ent}^{\mathcal{U}}(\mathcal{K} \mid \alpha)$ be the supremum over $\mathbf{r}$, F, and ϵ.

As the notation above suggests, it turns out that this quantity only depends on the embedding $\alpha : \mathcal{A} \to \mathcal{Q}$, not on the particular lift $(\alpha_n)_{n\in\mathbb{N}}$. To see this, suppose β_n is another such lift, so that for every $a \in \mathcal{A}$ we have $d^{M_n(\mathbb{C})}(\alpha_n(a), \beta_n(a)) \to 0$ as $n \to \mathcal{U}$. Using Urysohn's lemma, taking a smaller neighborhood if necessary, we can assume the neighborhood $\mathcal{O}$ is given by $\phi < \delta$ for some nonnegative formula $\phi(x_1, x_2, \dots)$ in $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$. Then ϕ can be equivalently viewed as an $\mathcal{L}_{\mathrm{tr}}$ formula in the variables x_j together with additional variables corresponding to the elements of $\mathcal{A}$. By uniform continuity of the formulas, $|\phi(\mathbf{X}, \alpha_n(a))_{a\in\mathcal{A}}) - \phi(\mathbf{X}, (\beta_n(a))_{a\in\mathcal{A}})| < \delta/2$ for n in a small enough neighborhood of $\mathcal{U}$. Thus, if the neighborhood $\mathcal{O}'$ is given by $\phi < \delta/2$, we get $\Gamma_{\mathbf{r}}^{(n)}(\mathcal{O}' \mid \beta_n) \subseteq \Gamma_{\mathbf{r}}^{(n)}(\mathcal{O} \mid \alpha_n)$. The argument is finished by taking the appropriate infima over $\mathcal{O}$ and limits.[1]

We remark that the approximate embedding $\alpha_n : \mathcal{A} \to M_n(\mathbb{C})$ can be thought of as a choice of microstates for *every* element of $\mathcal{A}$. But, as in Hayes original description of relative 1–bounded entropy, we could instead fix a generating set $\mathbf{A}$ for $\mathcal{A}$, fix microstates $\mathbf{A}^{(n)}$ for that *generating set*, and define microstate spaces of matrix tuples $\mathbf{X}$ such that the $\mathcal{L}_{\mathrm{tr}}$–type of $(\mathcal{A}^{(n)}, \mathbf{X})$ is in a certain neighborhood $\mathcal{O}$ of the set $\mathcal{K}$. It is

[1]For the analog of this argument in the existential case, we would work only with the case when $\mathcal{K} = \mathcal{K}_\mu$ for a single existential $\mathcal{L}_{\mathrm{tr},\mathcal{A}}$–type. The only issue adapting the above argument to the existential case is in finding, for a given a neighborhood $\mathcal{O}$ of $\mathcal{K}_\mu$, a sub-neighborhood of the form $\phi^{-1}((-\infty, \epsilon))$ for an existential formula ϕ. Since the space of existential types is not Hausdorff, we cannot apply Urysohn's lemma, but rather must work with the existential formulas directly to construct such a neighborhood.

a technical exercise to show that these definitions are equivalent, the key point being that every element of a can be expressed as a quantifier-free definable function of the generating tuple $\mathbf{A}$.

Most of the properties we showed for $\mathrm{Ent}^{\mathcal{U}}$ adapt to the relative version with the same method of proof. For instance, it satisfies the analog of the variational principle (Proposition 4.6) and monotonicity under pushforward (4.7). Thus, given $\mathcal{A} \subseteq \mathcal{N} \subseteq \mathcal{M}$ and an embedding $\alpha : \mathcal{A} \to \mathcal{Q}$, we can define $\mathrm{Ent}^{\mathcal{U}}(\mathcal{N} : \mathcal{M} \mid \alpha)$ as the supremum of $\mathrm{Ent}^{\mathcal{U}}(\mathrm{tp}^{\mathcal{M},\alpha}(\mathbf{X}))$ for tuples $\mathbf{X}$ from $\mathcal{N}$. Analogously to Lemma 4.14, if $\mathrm{Ent}^{\mathcal{U}}(\mathcal{M} \mid \alpha) \geq 0$, then there is an embedding of $\mathcal{N}$ into $\mathcal{Q}$ that restricts to α on $\mathcal{A}$ and extends to an elementary embedding of $\mathcal{M}$. The quantifier-free and existential version of conditional entropy are defined in a similar way, and the relationship between them works the same way as it does for the unconditional version.

References

[1] **S Atkinson**, **I Goldbring**, **S Kunnawalkam Elayavalli**, *Factorial relative commutants and the generalized Jung property for II1 factors*, Advances in Mathematics 396 (2022) 108107; https://doi.org/10.1016/j.aim.2021.108107

[2] **S Atkinson**, **S Kunnawalkam Elayavalli**, *On Ultraproduct Embeddings and Amenability for Tracial von Neumann Algebras*, International Mathematics Research Notices 2021 (2020) 2882–2918; https://doi.org/10.1093/imrn/rnaa257

[3] **I Ben Yaacov**, **A Berenstein**, **C W Henson**, **A Usvyatsov**, *Model theory for metric structures*, from: "Model Theory with Applications to Algebra and Analysis, Vol. II", (Z C et al, editor), London Mathematical Society Lecture Notes Series 350, Cambridge University Press (2008) 315–427; https://doi.org/10.1016/j.aim.2021.108107

[4] **I Ben Yaacov**, **A Usvyatsov**, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society 362 (2010) 5213–5259; https://doi.org/10.1090/S0002-9947-10-04837-3

[5] **B Blackadar**, *Operator Algebras: Theory of C^*-algebras and von Neumann algebras*, volume 122 of *Encyclopaedia of Mathematical Sciences*, Springer-Verlag, Berlin, Heidelberg (2006); https://doi.org/10.1007/3-540-28517-2

[6] **J Dixmier**, *Les C^*-algèbres et leurs représentations*, volume 29 of *Cahiers Scientifiques*, 2 edition, Gauthier-Villars, Paris (1969)Reprinted by Editions Jacques Gabay, Paris, 1996. Translated as C^*-algebras, North-Holland, Amsterdam, 1977. First Edition 1964

[7] **I Farah**, **B Hart**, **D Sherman**, *Model theory of operator algebras I: stability*, Bulletin of the London Mathematical Society 45 (2013) 825–838; https://doi.org/10.1112/blms/bdt014

[8] **I Farah**, **B Hart**, **D Sherman**, *Model theory of operator algebras II: model theory*, Israel Journal of Mathematics 201 (2014) 477–505; https://doi.org/10.1007/s11856-014-1046-7

[9] **I Farah**, **B Hart**, **D Sherman**, *Model theory of operator algebras III: elementary equivalence and* II_1 *factors*, Bulletin of the London Mathematical Society 46 (2014) 609–628; https://doi.org/10.1112/blms/bdu012

[10] **A Galatan**, **S Popa**, *Smooth bimodules and cohomology of* II_1 *factors*, J. Inst. Math. Jussieu (2015) 1–33; https://doi.org/10.1017/S1474748015000122

[11] **W Gao**, *Relative Embeddability of von Neumann Algebras and Amalgamated Free Products* (2020) With an appendix by Weichen Gao and Marius Junge; arXiv:arXiv:2012.07940

[12] **L Ge**, *Applications of free entropy to finite von Neumann algebras. II*, Ann. of Math. (2) 147 (1998) 143–157; https://doi.org/10.2307/120985

[13] **I Goldbring**, *Non-embeddable* II_1 *factors resembling the hyperfinite* II_1 *factor*. Preprint; arXiv:2101.10467

[14] **I Goldbring**, *On Popa's factorial commutant embedding problem*, Proc. Amer. Math. Soc. 148 (2020) 5007–5012; https://doi.org/https://doi.org/10.1090/proc/15141

[15] **I Goldbring**, *Enforceable operator algebras*, Journal of the Institute of Mathematics of Jussieu 20 (2021) 31–63; https://doi.org/10.1017/S1474748019000112

[16] **B Hayes**, *1-Bounded entropy and regularity problems in von Neumann algebras*, International Mathematics Research Notices 1 (2018) 57–137; https://doi.org/10.1093/imrn/rnw237

[17] **B Hayes**, **D Jekel**, **S Kunnawalkam Elayavalli**, *Property (T) and strong* 1*-boundedness for von Neumann algebras*. Preprint; arXiv:2107.03278

[18] **B Hayes**, **D Jekel**, **B Nelson**, **T Sinclair**, *A random matrix approach to absorption theorems for free products*, International Mathematics Research Notices 2021 (2021) 1919–1979; https://doi.org/10.1093/imrn/rnaa191

[19] **N Hindman**, **D Strauss**, *Algebra in the Stone-Čech Compactification: Theory and Applications*, De Gruyter, Berlin, Boston (2011); https://doi.org/doi:10.1515/9783110258356

[20] **A Ioana**, **P Spaas**, II_1 *factors with exotic central sequence algebras*, Journal of the Institute of Mathematics of Jussieu 20 (2021) 1671–1696; https://doi.org/10.1017/S1474748019000653

[21] **I M James**, *Introduction to Uniform Spaces*, volume 144 of *London Mathematical Society Lecture Note Series*, Cambridge University Press (1990); https://doi.org/10.1017/CBO9780511721519

[22] **D Jekel**, *Evolution equations in non-commutative probability*, PhD thesis, University of California, Los Angeles (2020)

[23] **D Jekel**, *Free probability and model theory of tracial* W**-algebras*. Preprint; to appear in *Model Theory of Operator Algebras*, ed. I. Goldbring, DeGruyter (2022); arXiv:2208.13867,

[24] **K Jung**, *Amenability, tubularity, and embeddings into* $\mathcal{R}^\omega$, Math. Ann. 338 (2007) 241–248; https://doi.org/10.1007/s00208-006-0074-y

[25] **K Jung**, *Strongly* 1*-bounded von Neumann algebras*, Geom. Funct. Anal. 17 (2007) 1180–1200; https://doi.org/10.1007/s00039-007-0624-9

[26] **R V Kadison**, **J R Ringrose**, *Fundamentals of the Theory of Operator Algebras I*, volume 15 of *Graduate Studies in Mathematics*, American Mathematical Society, Providence (1983)

[27] **I Kaplansky**, *Modules over operator algebras*, American Journal of Mathematics 75 (1953) 839–858; https://doi.org/10.2307/2372552

[28] **S Popa**, *Asymptotic Orthogonalization of Subalgebras in* II_1 *1 Factors*, Publ. Res. Inst. Math. Sci. 55 (2019) 795–809; https://doi.org/10.4171/PRIMS/55-4-5

[29] **S Sakai**, C**-algebras and* W**-algebras*, volume 60 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*, Springer-Verlag, Berlin Heidelberg (1971)

[30] **D Shlyakhtenko**, *Lower estimates on microstates free entropy dimension*, Analysis & PDE 2 (2009) 119–146; https://doi.org/10.2140/apde.2009.2.119

[31] **M Takesaki**, *Theory of Operator Algebras I*, volume 124 of *Encyclopaedia of Mathematical Sciences*, Springer-Verlag, Berlin Heidelberg (2002)

[32] **D-V Voiculescu**, *The analogues of entropy and of Fisher's information in free probability, II*, Inventiones Mathematicae 118 (1994) 411–440; https://doi.org/10.1007/BF01231539

[33] **D-V Voiculescu**, *The analogues of entropy and of Fisher's information measure in free probability, III: Absence of Cartan subalgebras*, Geometric and Functional Analysis 6 (1996) 172–199; https://doi.org/10.1007/BF02246772

[34] **K Zhu**, *An Introduction to Operator Algebras*, Studies in Advanced Mathematics, CRC Press, Ann Arbor (1993)

Department of Mathematics, University of California, San Diego

djekel@ucsd.edu

https://sites.google.com/ucsd.edu/djekel/home?authuser=0

Received: 19 July 2022 Revised: 11 April 2023