

MODEL THEORY AND METRIC CONVERGENCE I: METASTABILITY AND DOMINATED CONVERGENCE

EDUARDO DUEÑEZ AND JOSÉ N. IOVINO

Dedicated to Ward Henson with gratitude on the occasion of his retirement.

ABSTRACT. We study Tao’s finitary viewpoint of convergence in metric spaces, as captured by the notion of metastability. We adopt the perspective of continuous model theory. We show that, in essence, metastable convergence with a given rate is the only formulation of metric convergence that can be captured by a theory in continuous first-order logic, a result we call the Uniform Metastability Principle. Philosophically, this principle amounts to the following meta-theorem: “*If a classical statement about convergence in metric structures is refined to a statement about metastable convergence with some rate, then the validity of the original statement implies the validity of its metastable version.*” As an instance of this phenomenon, we formulate an abstract version of Tao’s Metastable Dominated Convergence Theorem as a statement about axiomatizable classes of metric structures, and show that it is a direct consequence of the Uniform Metastability Principle.

INTRODUCTION

The concept of convergence in metric spaces is fundamental in analysis. The present article is the first of a series focusing on results, both classical and new, in which the convergence of some sequence(s)—or, more generally, some nets—follows from suitable hypotheses. We shall use the loose nomenclature “convergence theorem” for any result of this kind; the best known such results are the classical Monotone and Dominated Convergence theorems, as well as the ergodic convergence theorems of von Neumann and Birkhoff. For simplicity, we assume that all metric spaces under consideration are complete (an alternative perspective would be the study of theorems about Cauchy sequences and nets in metric spaces not necessarily complete).

Given a convergence theorem, it is natural to ask whether it admits refinements whereby the conclusion states a stricter mode of convergence. When the statement of a convergence theorem involves a collection of sequences, the classical refinement of the property of simultaneous convergence is that of uniform convergence. However, few convergence theorems in analysis admit a natural refinement implying uniform convergence. Furthermore, even in the rare cases when uniform convergence of a family of sequences is implied, the parameters of uniform convergence are rarely universal: Typically, uniform convergence will hold in every structure (consisting of the ambient metric space, sequences therein, plus any other necessary ingredients) satisfying suitable hypotheses, but not uniformly across all such structures. In fact, even if a convergence theorem should refer to the convergence of a single sequence, our focus shall be on the entirety of structures to which the theorem applies, and hence on

Date: July 4, 2018.

2010 Mathematics Subject Classification. Primary: 40A05; Secondary: 03Cxx, 46Bxx, 28-xx.

Key words and phrases. Convergence, metastability, metric structures, finitary.

This research was partially funded by NSF grant DMS-1500615.

all instances of relevant sequences. Thus, even when a “single” sequence is under immediate discussion, we ask whether its mode of convergence admits parameters that are uniform over all instances.

Tao [Tao08] introduced the notion of *metastable convergence* (Definition 1.2), which is an equivalent formulation of the usual Cauchy property (Proposition 1.5). The metastable-convergence viewpoint leads to the notion of *uniformly metastable convergence* (Definition 1.7), which is not only a metastable analogue of the classical property of uniform convergence of a family of sequences, but also a generalization thereof (Remark 1.8). Tao obtains a metastable version of the classical Dominated Convergence Theorem that holds with metastable rates that are universal; this result plays a crucial role in the proof his remarkable result on the convergence of ergodic averages for polynomial abelian group actions [Tao08]. Walsh’s subsequent generalization [Wal12] of Tao’s theorem to polynomial nilpotent group actions relies on a similar convergence theorem.

In this article we prove that metastability with a given rate is the only formulation of metric convergence that can be captured by a theory in continuous first-order logic (Proposition 2.3). This is a precise statement of Tao’s observation that metastable convergence with a prescribed rate is a “finitary” property [Tao]. The conceptual backbone of the manuscript is the Uniform Metastability Principle (Proposition 2.4), which may be formulated as the following meta-theorem: *“If a classical statement about convergence in metric structures is refined to a statement about metastable convergence with some rate, then the validity of the original statement implies the validity of its metastable version.”* As an instance of this phenomenon, we obtain a soft proof of a version of the Metastable Dominated Convergence Theorem. Our proof depends on neither infinitary arguments *à la* Tao [Tao08], nor recursive arguments and constructive analysis *à la* Avigad *et al* [ADR12]. We show that the Uniform Metastability Principle follows directly from the fundamental theorem of model theory of metric structures, namely, the Compactness Theorem. We believe that, in spite of its simplicity, this principle captures a certain philosophical view revealing the scope of applicability of model-theoretic methods to the study of convergence in metric spaces. Although anticipated by Avigad and the second author [AI13], we are not aware of a purely model-theoretical formulation of this principle hitherto.

In order to make the results of this paper accessible to readers with no prior background in logic, particularly to researchers in analysis, the last section (Section 6) is a self-contained tutorial on the basics of model theory of metric structures. The literature contains several equivalent formulations of this theory; the one in current widespread use is the framework of first-order continuous logic developed by Ben Yaacov and Usvyatsov [BYU10, BYBHU08] building upon ideas of Chang-Keisler [CK62, CK66] and Henson [HI02]. We use the language of continuous approximations originally developed by Henson, as we feel that it is simpler and more natural for the applications at hand. The tutorial of Section 6 parallels portions of earlier introductions given by Henson and the second author that emphasize Banach spaces [HI02]; however, the present exposition places greater emphasis on metric structures and on topics of direct interest for the study of convergence theorems, particularly structures that are hybrid in the sense that they include metric sorts alongside discrete sorts. Readers already comfortable with model theory of metric structures from this perspective need to consult Section 6 only for reference.

Section 1, which does not depend on model theory at all, is an introduction to the metastable viewpoint of convergence. As an attempt to shed light into the finitary nature of uniformly metastable convergence, Subsection 1-I discusses the relation between the usual Cauchy criterion and Tao's notion of metastable convergence. In Subsection 1-II we define various notions of metastability, oscillation and uniform metastability for arbitrary sequences or nets in a metric space, and show how these notions relate to classical ones. Proposition 1.5, in particular, states that metastable convergence is equivalent to usual convergence (always in complete spaces).

In Section 2, we begin the discussion of sequences and nets from the viewpoint of model theory. We exhibit a (continuous) first-order axiomatization of the property of metastable convergence with a given uniform rate. Subsequently, we state and prove the Uniform Metastability Principle and a useful corollary thereof (Proposition 2.4 and Corollary 2.5).

In Section 3, we study finite measure structures from a model-theoretic viewpoint. We introduce a certain class of metric structures, which we call Loeb structures due to their close relation to Loeb probability spaces from nonstandard analysis [Loe75].¹ We start by defining the notion of pre-Loeb structure. Roughly speaking, a pre-Loeb structure is a metric structure that satisfies all the first-order properties of probability spaces $(\Omega, \mathcal{A}, \mathbf{P})$, where \mathbf{P} is a probability measure on a Boolean algebra \mathcal{A} of subsets of Ω . Loeb structures are then defined as pre-Loeb structures that satisfy a saturation hypothesis. By basic model theory (Section 6-IX), every metric structure can be extended naturally to a saturated structure. As a by-product of saturation, Loeb structures possess the infinitary properties (i.e., countable additivity) of *bona fide* classical probability spaces (Proposition 3.4).²

In Section 4, we study integration from a model-theoretic viewpoint. We begin by introducing the class of pre-integration structures. True to their name, these are structures satisfying the first-order axioms of the space $\mathcal{L}_\Omega^\infty$ of essentially bounded real functions on, say, a probability space $(\Omega, \mathcal{A}, \mathbf{P})$, endowed with the integration linear functional $I : \mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$ mapping f to $\int_\Omega f(\omega) d\mathbf{P}(\omega)$. (In particular, a pre-integration structure is a pre-Loeb structure.) Of course, most models of those axioms will not correspond to *bona fide*, countably additive probability spaces, nor will I correspond to the operation of integration with respect to a measure in the usual sense. Nevertheless, a Riesz Representation Theorem (Theorem 4.7) for integration structures (that is, for saturated pre-integration structures), asserts that the space $\mathcal{L}_\Omega^\infty$ (whose elements correspond to bounded functions on the induced Loeb probability space) is endowed with a positive functional extending I and given by the usual operation of integration with respect to Loeb measure.

In Section 5, we introduce directed pre-integration structures. In essence, these are pre-integration structures $(\Omega, \mathcal{A}, \mathbf{P}, \mathcal{L}_\Omega^\infty, I)$ further endowed, say, for simplicity, with a bounded sequence $\varphi = (\varphi_n : n \in \mathbb{N})$ of elements of $\mathcal{L}_\Omega^\infty$ (identified with bounded functions on Ω). Directed integration structures are defined as saturated directed pre-integration structures, as expected. By the Riesz Representation Theorem (Theorem 4.7), under the saturation hypothesis, I corresponds to a classical integration operator on $\mathcal{L}_\Omega^\infty$; therefore, the usual proof of the Dominated Convergence Theorem applies verbatim (Proposition 5.3). Obviously, the conclusion of the Dominated Convergence Theorem continues to hold if an additional hypothesis of uniform metastable pointwise convergence is imposed on the family φ ; moreover,

¹This manuscript makes no use of results or notions from nonstandard analysis.

²More precisely, every Loeb structure induces a classical probability space $(\Omega, \mathcal{A}_L, \mathbf{P}_L)$ on the same underlying sample space Ω (Proposition 3.4).

this hypothesis is axiomatizable. It follows from the Uniform Metastability Principle that, under the additional hypothesis of uniform metastable pointwise convergence, the conclusion of the Dominated Convergence Theorem must admit a strengthening to convergence with a certain metastable rate. This yields an immediate proof of Tao’s Metastable Dominated Convergence Theorem (Corollary 5.4).

In the Appendix to this manuscript we state the viewpoint that a property of metric spaces should be considered finitary when it is equivalent to the satisfaction of a collection of axioms in a first-order language for metric structures.

We conclude this introduction by mentioning that Tao has formulated a nonstandard-analysis version of his Metastable Dominated Convergence Theorem in a blog post on *Walsh’s ergodic theorem, metastability, and external Cauchy convergence* [Tao]. Tao’s insightful post served as philosophical motivation for the model-theoretic perspective adopted in the current manuscript.

We gratefully acknowledge support for this research from the National Science Foundation through grant DMS-1500615. We thank the administration and sponsors of the Banff International Research Station for providing an excellent research environment in June 2016 during the meeting of the focused research group on Topological Methods in Model Theory (16frg676). We also thank Xavier Caicedo, Chris Eagle, and Frank Tall for valuable discussions and comments that helped refine some of the ideas and content of this manuscript.

1. A FINITARY FORMULATION OF CONVERGENCE

I. MOTIVATION

Given $\epsilon \geq 0$, we say that a sequence $a = (a_n : n \in \mathbb{N})$ in some metric space X is $[\epsilon]$ -Cauchy if at least one of its tails $a_{\geq N} = (a_n : n \geq N)$ satisfies the inequality $\text{osc}_{\geq N}(a) := \sup_{m,n \geq N} d(a_m, a_n) \leq \epsilon$. The $[\epsilon]$ -Cauchy condition is infinitary in the sense that $\text{osc}_{\geq N}(a)$ depends on the values a_n as n varies over the infinite set $\{N, N+1, \dots\}$. Tao’s criterion for “metastable convergence” [Tao08] imposes small oscillation conditions only on *finite* segments $a_{[N, N']} = (a_n : N \leq n \leq N')$ of a . For every fixed choice of a strictly increasing function $F : \mathbb{N} \rightarrow \mathbb{N}$ we regard the collection $\eta = ([N, F(N)] : N \in \mathbb{N})$ as a “sampling” of \mathbb{N} , one finite segment at a time; abusing nomenclature, we use the name *sampling* to refer to either the collection η or the function F defining it. In Tao’s nomenclature, for $\epsilon > 0$ and sampling F , the sequence a is $[\epsilon, F]$ -metastable if the inequality $\text{osc}_{[N, F(N)]}(a) := \sup_{N \leq m, n \leq F(N)} d(a_m, a_n) \leq \epsilon$ holds for some N . For fixed $\epsilon > 0$ and sampling F , “[ϵ, F]-metastable” is a weaker property than “[ϵ]-Cauchy” inasmuch as the former involves only the values of a on the subsets $a_{[N, F(N)]}$ of the tails $a_{\geq N}$. However, when F varies over all samplings of \mathbb{N} , the conjunction of the corresponding properties of $[\epsilon, F]$ -metastability of a implies that a is $[\epsilon]$ -Cauchy. This leads to Tao’s characterization of convergence (i.e., of the Cauchy property in complete spaces) as $[\epsilon, F]$ -metastability for all $\epsilon > 0$ and all samplings F (Proposition 1.5).

The metastable characterization of convergence is still not quite finitary because the existential statement on $N \in \mathbb{N}$ is infinitary. (We use the term “finitary” as a synonym for “axiomatizable in first-order continuous logic” per the Appendix to this manuscript.) Tao’s concept of *metastability with a given rate* arises by restricting this existential statement to a bounded (finitary) one. After Tao, for fixed $\epsilon > 0$ and sampling F , we call $M \in \mathbb{N}$ an (*upper bound on the*) *rate of $[\epsilon, F]$ -metastable convergence of a* if $\text{osc}_{[N, F(N)]}(a) \leq \epsilon$ holds for

some $N \leq M$. Any collection $M_\bullet = (M_{\epsilon,F}) \subset \mathbb{N}$ of natural numbers, one for each $\epsilon > 0$ and sampling F , is an (*upper bound on the*) *rate of metastability of a* if $M_{\epsilon,F}$ bounds the $[\epsilon, F]$ -rate of metastable convergence of a for all ϵ, F . Evidently, given ϵ, F and $M \in \mathbb{N}$, the property “ M is a rate of $[\epsilon, F]$ -metastability for a ” is finitary.

Given arbitrary ϵ, F and any Cauchy sequence a , it is clear that a admits *some* bound M_\bullet on its rate of metastability—one may take $M_{\epsilon,F}$ to be any N satisfying $\text{osc}_{\geq N}(a) \leq \epsilon$. However, no choice of M_\bullet applies uniformly to *all* Cauchy sequences.

In Subsection II below, we define various notions of metastable convergence for nets in metric spaces. In particular, the notion of metastable convergence with a given rate (Definition 1.7) is the natural finitary notion of convergence of nets extending Tao’s (for sequences).

II. METASTABLE CONVERGENCE OF NETS IN METRIC SPACES

Throughout this section we fix a nonempty directed set (\mathcal{D}, \leq) (that is, \leq is a reflexive, antisymmetric and transitive binary relation on \mathcal{D} such that every pair of elements has an upper bound). We denote by $\mathcal{D}_{\geq i}$ the final segment $\{j \in \mathcal{D} : j \geq i\} \subset \mathcal{D}$.

We recall that a \mathcal{D} -net $a_\bullet = (a_i : i \in \mathcal{D})$ in a topological space X *converges to* $b \in X$ if, for every neighborhood B of b in X , there exists $i \in \mathcal{D}$ such that $a_j \in B$ for all $j \geq i$. If (X, d) is a metric space and $\epsilon \geq 0$, we will say that the \mathcal{D} -net a_\bullet is $[\epsilon]$ -*Cauchy* if there exists $i \in \mathcal{D}$ such that $d(a_j, a_{j'}) \leq \epsilon$ for all $j, j' \geq i$, and that a_\bullet is (ϵ) -*Cauchy* (or ϵ -*Cauchy*) if a_\bullet is $[\epsilon']$ -Cauchy for all $\epsilon' > \epsilon$. The net a_\bullet is Cauchy in the usual sense when it is 0-Cauchy; in this case, a_\bullet converges to b if for all $\epsilon > 0$ there is $i \in \mathcal{D}$ such that $d(a_i, b) \leq \epsilon$ for all $j \geq i$. Every Cauchy net in a complete metric space X converges to some (necessarily unique) element $b \in X$. All metric spaces under consideration shall be complete, so we use the term “convergent” as a synonym of “Cauchy”.

1.1. Definition. A *sampling* of the directed set (\mathcal{D}, \leq) is any collection $\eta = (\eta_i : i \in \mathcal{D})$ of finite subsets of \mathcal{D} (indexed by \mathcal{D} itself), such that η_i is a nonempty finite subset of $\mathcal{D}_{\geq i}$ for each $i \in \mathcal{D}$.³ The collection of all samplings of \mathcal{D} will be denoted $\text{Smpl}(\mathcal{D})$.

1.2. Definition. Fix a metric space (X, d) and directed set (\mathcal{D}, \leq) . For $\eta \in \text{Smpl}(\mathcal{D})$ and $\epsilon > 0$, a \mathcal{D} -net $a_\bullet = (a_i : i \in \mathcal{D})$ in X is:

- *strict ϵ, η -metastable*, or $[\epsilon, \eta]$ -*metastable* (note the square brackets), if there exists $i \in \mathcal{D}$ such that $d(a_j, a_{j'}) \leq \epsilon$ for all $j, j' \in \eta_i$. Any such i is a *witness* of the (strict) $[\epsilon, \eta]$ -metastability of a_\bullet .
- *lax ϵ, η -metastable*, or (ϵ, η) -*metastable* (note the round parentheses), if a_\bullet is $[\epsilon', \eta]$ -metastable for all $\epsilon' > \epsilon$,
- η -*metastable*, if a_\bullet is $(0, \eta)$ -metastable,
- ϵ -*metastable*, if a_\bullet is (ϵ, η') -metastable for all $\eta' \in \text{Smpl}(\mathcal{D})$,
- *metastable* if a_\bullet is 0-metastable.

For fixed ϵ, η , (ϵ, η) -metastability (resp., $[\epsilon, \eta]$ -metastability, ϵ -metastability) implies (ϵ', η) -metastability (resp., $[\epsilon', \eta]$ -metastability, ϵ' -metastability) for all $\epsilon' \geq \epsilon$. It is also clear that strict $[\epsilon, \eta]$ -metastability implies lax (ϵ, η) -metastability, that lax (ϵ', η) -metastability for all $\epsilon' > \epsilon \geq 0$ implies lax (ϵ, η) -metastability, and consequently that ϵ' -metastability for all $\epsilon' > \epsilon \geq 0$ implies ϵ -metastability. We also remark that any net a_\bullet is necessarily C -metastable if C is an upper bound on the distances $d(a_i, a_j)$ ($i, j \in \mathcal{D}$); this is the case, in particular, if $d(a_i, x_0) \leq C/2$ for some fixed $x_0 \in X$ and all $i \in \mathcal{D}$.

³The condition “ η_i is nonempty” is not logically necessary. We impose it for heuristic reasons only.

1.3. **Definition.** Let $a_\bullet = (a_i : i \in \mathcal{D})$ be an arbitrary bounded \mathcal{D} -net on a metric space (X, d) .

For every $\eta \in \text{Smpl}(\mathcal{D})$, the η -oscillation $\text{osc}_\eta(a_\bullet)$ of a_\bullet is

$$\text{osc}_\eta(a_\bullet) = \inf\{\epsilon \geq 0 : a_\bullet \text{ is } [\epsilon, \eta]\text{-metastable}\}.$$

The oscillation $\text{osc}(a_\bullet)$ of a_\bullet is

$$\text{osc}(a_\bullet) = \sup\{\text{osc}_\eta(a_\bullet) : \eta \in \text{Smpl}(\mathcal{D})\}.$$

Note that the notations “osc”, “osc $_\eta$ ” fail to exhibit the dependence of $\text{osc}(a_\bullet)$, $\text{osc}_\eta(a_\bullet)$ on \mathcal{D} . However, when these notations are used, the directed set \mathcal{D} will be fixed, precluding ambiguous interpretations.

1.4. **Proposition.** For any bounded \mathcal{D} -net a_\bullet , $\eta \in \text{Smpl}(\mathcal{D})$ and $\epsilon \geq 0$:

- (1) $\text{osc}_\eta(a_\bullet) \leq \epsilon$ if and only if a_\bullet is (ϵ, η) -metastable,
- (2) $\text{osc}_\eta(a_\bullet) = \inf\{\epsilon \geq 0 : \text{osc}_{\eta_i}(a_\bullet) \leq \epsilon \text{ for some } i \in \mathcal{D}\}$,
- (3) $\text{osc}_\eta(a_\bullet) = \min\{\epsilon \geq 0 : a_\bullet \text{ is } (\epsilon, \eta)\text{-metastable}\}$,
- (4) $\text{osc}_\eta(a_\bullet) = \inf\{\text{osc}_{\eta_i}(a_\bullet) : i \in \mathcal{D}\}$,
- (5) $\text{osc}(a_\bullet) = \min\{\epsilon \geq 0 : a_\bullet \text{ is } \epsilon\text{-metastable}\}$,
- (6) $\text{osc}(a_\bullet) = \sup\{\epsilon \geq 0 : \text{for all } i \in \mathcal{D} \text{ there exist } j, j' \geq i \text{ with } d(a_j, a_{j'}) \geq \epsilon\}$,
- (7) a_\bullet is ϵ -Cauchy if and only if $\text{osc}(a_\bullet) \leq \epsilon$,
- (8) a_\bullet is ϵ -Cauchy if and only if a_\bullet is ϵ -metastable.

It is customary to use (6) above as the definition of $\text{osc}(a_\bullet)$.

Proof. (1) Let $r = \text{osc}_\eta(a_\bullet)$. For any $t > r$ the definition of osc_η implies that a_\bullet is $[s, \eta]$ -metastable for some $s \in (r, t)$, hence also $[t, \eta]$ -metastable; therefore, a_\bullet is (r, η) -metastable, hence *a fortiori* (ϵ, η) -metastable for all $\epsilon \geq r$. Conversely, if a_\bullet is (ϵ, η) -metastable, then it is $[\epsilon', \eta]$ -metastable for all $\epsilon' \geq \epsilon$, so the definition of osc_η implies $r \leq \epsilon$.

(2) Clearly, a_\bullet is $[\epsilon, \eta]$ -metastable iff $\text{osc}_{\eta_i}(a_\bullet) \leq \epsilon$ for some $i \in \mathcal{D}$.

(3) This follows immediately from (1).

(4) Let $r = \text{osc}_\eta(a_\bullet)$ and $s = \inf\{\text{osc}_{\eta_i}(a_\bullet) : i \in \mathcal{D}\}$. By part (2), we have $r \leq s$. Conversely, if $\epsilon \geq \text{osc}_{\eta_i}(a_\bullet)$ for some $i \in \mathcal{D}$, then $\epsilon \geq s$, hence $s \leq r$.

(5) Clearly, the set $\{\epsilon \geq 0 : a_\bullet \text{ is } \epsilon\text{-metastable}\}$ has a least element, say s , since ϵ' -metastability for all $\epsilon' > \epsilon$ is equivalent to ϵ -metastability. Let $r = \text{osc}(a_\bullet)$. If a_\bullet is ϵ -metastable, then it is (ϵ, η) -metastable for all $\eta \in \text{Smpl}(\mathcal{D})$, hence $\text{osc}_\eta(a_\bullet) \leq \epsilon$ by (1), so $r \leq s$ by (3). Conversely, for all $\eta \in \text{Smpl}(\mathcal{D})$, a_\bullet is (r, η) -metastable by (1) and the definition of osc , and thus a_\bullet is r -metastable, so $r \geq s$.

(6) Let $r = \text{osc}(a_\bullet)$ and

$$s = \sup\{\epsilon \geq 0 : \text{for all } i \in \mathcal{D} \text{ there exist } j, j' \geq i \text{ with } d(a_j, a_{j'}) \geq \epsilon\}.$$

If $0 \leq t < s$, then for all $i \in \mathcal{D}$ there exists $\eta_i = \{j, j'\} \subset \mathcal{D}_{\geq i}$ such that $d(a_j, a_{j'}) > t$, hence a_\bullet is not $[t, \eta]$ -metastable for $\eta = (\eta_i : i \in \mathcal{D})$. It follows that $t \leq \text{osc}_\eta(a_\bullet) \leq r$. As this holds for all positive $t < s$, we have $s \leq r$. Conversely, if $0 \leq t < r$, then $\text{osc}_\eta(a_\bullet) > t$ for some $\eta \in \text{Smpl}(\mathcal{D})$. By (1), a_\bullet is not (t, η) -metastable. Hence, there is $t' > t$ such that for all $i \in \mathcal{D}$ there exist $j, j' \in \eta_i \subset \mathcal{D}_{\geq i}$ with $d(a_j, a_{j'}) > t' > t$; hence, $t \leq s$. As this holds for all $t < r$, we have $r \leq s$.

- (7) Let $r = \text{osc}(a_\bullet)$. It follows from (6) that a_\bullet is r -Cauchy, thus also ϵ -Cauchy for all $\epsilon \geq r$. Conversely, if $0 \leq \epsilon < r$, then by (6) there exists $\epsilon' \in (\epsilon, r)$ with the property that for all $i \in \mathcal{D}$ there exist $j, j' \geq i$ with $d(a_j, a_{j'}) \geq \epsilon'$. Then a_\bullet is not ϵ -Cauchy. \square
- (8) This follows immediately from (5) and (7), plus the remarks following Definition 1.2.

1.5. Proposition (Metastable characterization of the Cauchy property). *A net in a metric space X is Cauchy if and only if it is metastable.*

Proof. This is the particular case $\epsilon = 0$ of (8) in Proposition 1.4. \square

1.6. Remark. Propositions 1.4 and 1.5 remain true if we consider only samplings η such that η_i consists of no more than two elements of $\mathcal{D}_{\geq i}$ for all $i \in \mathcal{D}$. This more restrictive definition of sampling could be used in all further developments without any essential changes.

1.7. Definition. Fix a directed set \mathcal{D} . The collection of all finite subsets of \mathcal{D} will be denoted $\mathcal{P}_{\text{fin}}(\mathcal{D})$. Let $a_\bullet = (a_i : i \in \mathcal{D})$ be a \mathcal{D} -net in a metric space X .

Given $\epsilon > 0$ and $\eta \in \text{Smpl}(\mathcal{D})$, a set $E \in \mathcal{P}_{\text{fin}}(\mathcal{D})$ is called a *(bound on the) rate of $[\epsilon, \eta]$ -metastability of a_\bullet* if there exists a witness $i \in E$ of the $[\epsilon, \eta]$ -metastability of a_\bullet . (No sequence a_\bullet has this property if E is empty.)

For $\epsilon > 0$, a collection $E_\bullet = (E_\eta : \eta \in \text{Smpl}(\mathcal{D})) \subset \mathcal{P}_{\text{fin}}(\mathcal{D})$ is called a *(bound on the) rate of $[\epsilon]$ -metastability of a_\bullet* if E_η is a rate of $[\epsilon, \eta]$ -metastability of a_\bullet for all $\eta \in \text{Smpl}(\mathcal{D})$.

For $r \geq 0$ and $\eta \in \text{Smpl}(\mathcal{D})$, a collection $E_\bullet = (E_\epsilon : \epsilon > r)$ in $\mathcal{P}_{\text{fin}}(\mathcal{D})$ is called a *(bound on the) rate of (r, η) -metastability of a_\bullet* if E_ϵ is a rate of $[\epsilon, \eta]$ -metastability of a_\bullet for all $\epsilon > r$.

For $r \geq 0$, a collection $E_\bullet = (E_{\epsilon, \eta} : \eta \in \text{Smpl}(\mathcal{D}), \epsilon > r)$ in $\mathcal{P}_{\text{fin}}(\mathcal{D})$ is called a *(bound on the) rate of r -metastability of a_\bullet* , if $E_{\epsilon, \eta}$ is a rate of (r, η) -metastability of a_\bullet for all $\eta \in \text{Smpl}(\mathcal{D})$, where $E_{\epsilon, \eta} = (E_{\epsilon, \eta} : \epsilon > r)$. When $r = 0$, we say simply that E_\bullet is a (bound on the) rate of metastability of a_\bullet .

If \mathcal{C} is any collection of \mathcal{D} -nets in X , we say that:

- E is a uniform (bound on the) rate of $[\epsilon, \eta]$ -metastability for \mathcal{C} , or \mathcal{C} is E -uniformly $[\epsilon, \eta]$ -metastable, if a_\bullet is E -uniformly $[\epsilon, \eta]$ -metastable for all $a_\bullet \in \mathcal{C}$;
- E_\bullet is a uniform (bound on the) rate of $[\epsilon]$ -metastability for \mathcal{C} , or \mathcal{C} is E_\bullet -uniformly $[\epsilon]$ -metastable, if a_\bullet is E_\bullet -uniformly $[\epsilon]$ -metastable for all $a_\bullet \in \mathcal{C}$;
- E_\bullet is a uniform (bound on the) rate of r -metastability for \mathcal{C} , or \mathcal{C} is E_\bullet -uniformly r -metastable, if a_\bullet is E_\bullet -uniformly r -metastable for all $a_\bullet \in \mathcal{C}$. (When $r = 0$, we usually omit it.)

1.8. Remark. We show that the concept of uniform metastability generalizes that of uniform convergence. Let \mathcal{C} be any collection of \mathcal{D} -nets. Given $M \in \mathcal{D}$ and $\epsilon > 0$, the collection \mathcal{C} is M -uniformly $[\epsilon]$ -Cauchy if $\text{osc}_{\geq M}(a_\bullet) := \sup_{j, j' \geq M} d(a_j, a_{j'}) \leq \epsilon$ for each $\epsilon > 0$ and $a_\bullet \in \mathcal{C}$. Given $M_\bullet = (M_\epsilon : \epsilon > 0) \subset \mathcal{D}$, the collection \mathcal{C} is M_\bullet -uniformly Cauchy if every $a_\bullet \in \mathcal{C}$ is M_ϵ -uniformly $[\epsilon]$ -Cauchy for each $\epsilon > 0$. Corresponding to M_\bullet there is a rate of metastability $E_\bullet = (E_{\epsilon, \eta})$ defined by $E_{\epsilon, \eta} = \{M_\epsilon\} \in \mathcal{P}_{\text{fin}}(\mathcal{D})$. Under this identification, the collection \mathcal{C} is M_\bullet -uniformly Cauchy if and only if it is E_\bullet -metastable.

Clearly, metastability rates E_\bullet obtained from a collection M_\bullet as above are very special. The following example, due to Avigad *et al.* [ADR12], exhibits a family of uniformly metastable sequences that are not uniformly convergent. Every monotonically increasing sequence in $[0, 1]$ is convergent. Let \mathcal{C} be the collection of all such sequences. Clearly, there is no rate M_\bullet such that all sequences $a_\bullet \in \mathcal{C}$ are M_\bullet -uniformly Cauchy; in fact, for $\epsilon \in (0, 1)$, the sequence a_\bullet with $a_m = 0$ for $m \leq M_\epsilon$ and $a_m = 1$ for $m > M_\epsilon$ satisfies $\text{osc}_{\geq E_\epsilon}(a_\bullet) = 1 > \epsilon$.

On the other hand, for any $\epsilon > 0$ and a sampling of \mathbb{N} given as a strictly increasing function $F : \mathbb{N} \rightarrow \mathbb{N}$ (as per Section 1.1), let $E_{\epsilon, F} = \{m \in \mathbb{N} : m \leq F^{(k)}(0)\}$ where $k = \lceil \epsilon^{-1} \rceil$ is the smallest integer no smaller than ϵ^{-1} and $F^{(k)}(0)$ is the k -fold iterate of F applied to 0. Since $k\epsilon \geq 1$, at least one of the k differences $a_{F^{(j+1)}(0)} - a_{F^{(j)}(0)}$ ($j = 0, 1, \dots, k-1$) must not exceed ϵ whenever $a_{\bullet} \in \mathcal{C}$, hence \mathcal{C} is $E_{\epsilon, F}$ -uniformly $[\epsilon]$ -metastable. The collection $E_{\bullet} = (E_{\epsilon, F})$ is a uniform metastability rate for \mathcal{C} .

The concept of uniform metastability is crucial to our applications. As discussed in Remark 1.8 above, uniform metastability is a proper generalization of the classical notion of uniform convergence. Moreover, uniform metastability with a given rate is axiomatizable in the logic of metric structures (Proposition 2.3 below). This allows for powerful applications of model theory, particularly of compactness (e.g., the Uniform Metastability Principle, Proposition 2.4). We believe that many convergence results in analysis follow from hypotheses captured by the semantics of first-order logic for metric structures; consequently, such results ought to admit refinements to convergence with a metastability rate—and moreover the rate ought to be universal, i.e., independent of the structure to which the theorem is applied. Tao’s Metastable Dominated Convergence Theorem, as stated and proved in Section 5 below, is but one example of this philosophy.

2. CONVERGENCE OF NETS IN METRIC STRUCTURES AND UNIFORM METASTABILITY

Throughout Sections 2 to 5, we will assume that the reader is familiar with the material presented in Section 6. In particular, we assume familiarity with the notions of Henson metric structure, signature, positive bounded formula, approximate satisfaction, the Compactness Theorem 6.31, and saturated structures. We will deal with multi-sorted structures that contain discrete sorts alongside nondiscrete ones. Recall that \mathbb{R} , equipped with its ordered field structure and a constant for each rational number, occurs tacitly as a sort of every metric structure, and that discrete predicates in a structure are seen as $\{0, 1\}$ -valued functions. If (M, d, a) is a pointed metric space and $C \geq 0$, the set $\{x \in M \mid d(x, a) \leq C\}$ will be denoted $M^{[C]}$.

Throughout the paper, L will be a many-sorted signature with sorts $(\mathbb{S}_i : i < \alpha)$ (for some ordinal $\alpha > 1$), and \mathbb{S}_0 will be the special sort designated for \mathbb{R} . For notational convenience, we will identify \mathbb{S}_0 with \mathbb{R} .

Hereafter, (\mathcal{D}, \leq) will denote a directed set with least element j_0 . We will regard (\mathcal{D}, \leq) as a discrete metric structure with sorts \mathcal{D} and \mathbb{R} , and the point j_0 will be regarded as the anchor for the sort \mathcal{D} . We will refer to (\mathcal{D}, \leq, j_0) as a *pointed directed set*.

2.1. Definition (Directed structure). Fix a pointed directed set (\mathcal{D}, \leq, j_0) . Let L be a many-sorted signature with sorts $(\mathbb{S}_i : i < \alpha)$ (for some ordinal $\alpha > 1$), where $\mathbb{S}_0 = \mathbb{R}$. The sort \mathbb{S}_1 will be called the *directed sort*; it will be denoted by \mathbb{D} henceforth. Let L include a symbol $[\cdot \leq \cdot]$ for a function $\mathbb{D} \times \mathbb{D} \rightarrow \{0, 1\}$. Let L also include distinct constant symbols c_j , one for each $j \in \mathcal{D}$. In addition to the function symbols for the sort metrics and the operations on \mathbb{R} plus constants for rational numbers, L may include other function and constant symbols, as well as any other sorts than those already mentioned. A \mathcal{D} -directed structure is any metric L -structure \mathcal{M} such that

- the sort $\mathbb{D}^{\mathcal{M}}$ is discrete,

- the interpretation of $\llbracket \cdot \leq \cdot \rrbracket$ induces an order on $\mathbb{D}^{\mathcal{M}}$, denoted \leq (by a slight abuse of notation), such that $(\mathbb{D}^{\mathcal{M}}, \leq, \mathbf{c}_{j_0}^{\mathcal{M}})$ is an anchored directed set and, for all $i, j \in \mathbb{D}^{\mathcal{M}}$, $i \leq j$ holds precisely when $\llbracket i \leq j \rrbracket = 1$,
- the map $\mathcal{D} \rightarrow \mathbb{D}^{\mathcal{M}}$ defined by $j \mapsto \mathbf{c}_j^{\mathcal{M}}$ is an order-preserving injection;

It should be clear that the class of \mathcal{D} -directed structures is axiomatizable in the semantics of approximate satisfaction of positive bounded L -formulas. We note that the embedding $\mathcal{D} \hookrightarrow \mathbb{D}^{\mathcal{M}}$ is usually not surjective.

In order to discuss nets in \mathcal{D} -directed structures, we need to extend the language L with a function symbol $\mathbf{s} : \mathbb{D} \rightarrow \mathbb{S}_l$ interpreted as a function from the directed sort \mathbb{D} into some other sort \mathbb{S}_l of L . In this context, T shall denote any fixed uniform $L[\mathbf{s}]$ -theory extending the theory of \mathcal{D} -directed L -structures. (Extensions of a Henson language are discussed in Section 6.X.)

2.2. Definition. Fix a \mathcal{D} -directed L -structure \mathcal{M} , a new function symbol $\mathbf{s} : \mathbb{D} \rightarrow \mathbb{S}_l$, and a uniform $L[\mathbf{s}]$ -theory T . An (*external*) *net* in \mathcal{M} is any \mathcal{D} -net $\mathbf{s}_\bullet = (s_i : i \in \mathcal{D})$ taking values in some sort of L . An *internal net (modulo T)* is a function $s : \mathbb{D}^{\mathcal{M}} \rightarrow \mathbb{S}_l^{\mathcal{M}}$ such that (\mathcal{M}, s) is a model of T .

Note that an internal net s yields an external \mathcal{D} -net $\mathbf{s}_\bullet = (s_i : i \in \mathcal{D})$, letting $s_i = s(i)$ for $i \in \mathcal{D}$.

For any rational $\epsilon > 0$, $\eta \in \text{Smpl}(\mathcal{D})$ and internal net $s = \mathbf{s}^{\mathcal{M}}$, the \mathcal{D} -net \mathbf{s}_\bullet is strict $[\epsilon, \eta]$ -metastable if and only if there exists $i \in \mathcal{D}$ such that

$$\mathcal{M} \approx \xi_i^\eta(\epsilon) \quad (\text{equivalently, } \mathcal{M} \models \xi_i^\eta(\epsilon)),$$

where ξ_i^η is the positive bounded L -formula

$$\xi_i^\eta(\mathbf{t}) : \bigwedge_{j, j' \in \eta_i} (d_l(\mathbf{s}(j), \mathbf{s}(j')) \leq \mathbf{t}), \quad (1)$$

as follows from the semantics of approximate satisfaction and part (3) of Proposition 1.4.⁴ We emphasize that ξ_i^η is a *bona fide* L -formula since it is a finite conjunction of atomic formulas. Nevertheless, the property “ \mathbf{s}_\bullet is $[\epsilon, \eta]$ -metastable” is not L -axiomatizable since the asserted existence of the witness $i \in \mathcal{D}$ amounts to an infinite disjunction of formulas when \mathcal{D} is itself infinite.

For fixed $\eta \in \text{Smpl}(\mathcal{D})$ and $\epsilon > 0$, call \mathbf{s}_\bullet an $[\epsilon, \eta]$ -*unstable* \mathcal{D} -net if, for all $i \in \mathcal{D}$, there exist $j, j' \in \eta_i$ with $d_l(s_j, s_{j'}) \geq \epsilon$. For rational $\epsilon \geq 0$, the assertion “ \mathbf{s}_\bullet is $[\epsilon, \eta]$ -unstable” is equivalent to

$$\mathcal{M} \approx \mathbb{W} \xi_i^\eta(\epsilon) \quad \text{for all } i \in \mathcal{D}, \quad (\text{alternatively, } \mathcal{M} \models \mathbb{W} \xi_i^\eta(\epsilon))$$

where

$$\mathbb{W} \xi_i^\eta(\mathbf{t}) : \bigvee_{j, j' \in \eta_i} (d_l(\mathbf{s}(j), \mathbf{s}(j')) \geq \mathbf{t}). \quad (2)$$

Note that $[\epsilon, \eta]$ -metastability is consistent with $[\epsilon', \eta]$ -instability precisely when $\epsilon' \leq \epsilon$. In contrast to the property of strict metastability, the property of $[\epsilon, \eta]$ -instability for given

⁴For notational convenience, in (1) we write \mathbf{s}_k for $\mathbf{s}(\mathbf{c}_k)$, where \mathbf{c}_k is the constant denoting the element $k \in \mathcal{D}$. Similar simplifications will be usually made without comment whenever the intended strict syntax is otherwise clear.

$\eta \in \mathcal{D}$ and $\epsilon \geq 0$ is axiomatized by the collection

$$\{\mathbb{W}\xi_i^\eta(\epsilon') : \text{rational } \epsilon' < \epsilon \text{ and } i \in \mathcal{D}\}$$

of L -sentences, in the semantics of approximate (or discrete) satisfaction.

Recall that if \mathcal{M} is a structure, the complete theory $\text{Th}(\mathcal{M})$ of \mathcal{M} is the set of all sentences satisfied by \mathcal{M} . It should be clear from the preceding observations that if s_\bullet is a \mathcal{D} -net underlying an internal net s , then classical properties of s_\bullet , including: “ s_\bullet is $[\epsilon, \eta]$ -metastable”, “ s_\bullet is (ϵ, η) -metastable”, “ s_\bullet is η -metastable”, “ $r \leq \text{osc}(s_\bullet) \leq s$ ”, etc., depend only on the complete theory T of \mathcal{M} , though some of these properties are not axiomatizable.

2.3. Proposition (Axiomatizability of uniform metastability). *Fix a pointed directed set (\mathcal{D}, \leq, j_0) , a language L for \mathcal{D} -directed structures, a new function symbol $\mathbf{s} : \mathbb{D} \rightarrow \mathbb{S}_\iota$ and a uniform $L[\mathbf{s}]$ -theory T . For fixed reals $\epsilon > r \geq 0$ and sampling $\eta \in \text{Smpl}(\mathcal{D})$, the following properties of the \mathcal{D} -net s_\bullet induced by an internal net $s = \mathbf{s}^\mathcal{M}$ are $L[\mathbf{s}]$ -axiomatizable (modulo T):*

- “ E is a rate of $[\epsilon, \eta]$ -metastability for s_\bullet ”, for any fixed $E \in \mathcal{P}_{\text{fin}}(\mathcal{D})$;
- “ E_\bullet is a rate of (r, η) -metastability for s_\bullet ”, for any family $E_\bullet = (E_\epsilon : \epsilon > r)$ in $\mathcal{P}_{\text{fin}}(\mathcal{D})$;
- “ E_\bullet is a rate of r -metastability for s_\bullet ”, for any family $E_\bullet = (E_{\epsilon, \eta} : \eta \in \text{Smpl}(\mathcal{D}), \epsilon > r)$ in $\mathcal{P}_{\text{fin}}(\mathcal{D})$.

Hence, each of the preceding properties characterizes an axiomatizable subclass of models of T . In fact, in each case the collection of axioms is independent of T and L : It depends only on the underlying directed set \mathcal{D} .

Proof. For $E \in \mathcal{P}_{\text{fin}}(\mathcal{D})$ and $\eta \in \mathcal{D}$, let

$$\xi_E^\eta = \bigvee_{i \in E} \xi_i^\eta \quad \text{with } \xi_i^\eta \text{ given by (1).}$$

The first stated property is axiomatized by the formulas $\xi_E^\eta(\epsilon')$ for rational $\epsilon' > \epsilon$, the second by all formulas $\xi_{E_\epsilon}^\eta(\epsilon')$ for rational $\epsilon' > \epsilon$, and the third by all formulas $\xi_{E_{\epsilon, \eta}}^\eta(\epsilon')$ for $\epsilon > r$, rational $\epsilon' > \epsilon$, and all $\eta \in \text{Smpl}(\mathcal{D})$. \square

2.4. Proposition (Uniform Metastability Principle). *Fix a pointed directed set (\mathcal{D}, \leq, j_0) , a language L for \mathcal{D} -directed structures, a new function symbol $\mathbf{s} : \mathbb{D} \rightarrow \mathbb{S}_\iota$ and a uniform $L[\mathbf{s}]$ -theory T . Let $\eta \in \text{Smpl}(\mathcal{D})$ and $r \geq 0$. If s_\bullet is (r, η) -metastable whenever $(\mathcal{M}, s) \approx T$, then there exists a collection $E_\bullet = (E_\epsilon : \epsilon > r)$ in $\mathcal{P}_{\text{fin}}(\mathcal{D})$ such that E_\bullet is a rate of (r, η) -metastability for s_\bullet , uniformly over all models (\mathcal{M}, s) of T .*

Loosely speaking, if all internal sequences are metastable, then they are uniformly metastable. (Of course the notion of internal sequence is in reference to a fixed uniform theory T .) On the other hand, the uniform rate E_\bullet is certainly dependent on T . The uniform theory T implies an upper bound $C \geq 0$ such that $s(j) \in (\mathbb{S}_\iota)^{[C]}$ for all $j \in \mathbb{D}^\mathcal{M}$ whenever $(\mathcal{M}, s) \approx T$.

Proof. It is enough to show that a bound E_ϵ on the rate of $[\epsilon, \eta]$ -metastability exists for each rational $\epsilon > r$, uniformly for all s_\bullet arising from $\mathbf{s}^\mathcal{M}$ for arbitrary $\mathcal{M} \approx T$. Assume no such E_ϵ exists for some $\epsilon > r$. With the notation of Proposition 2.3, given $E \in \mathcal{P}_{\text{fin}}(\mathcal{D})$ there exists $(\mathcal{M}, s) \approx T$ and some rational $\epsilon' > \epsilon$ such that $(\mathcal{M}, s) \not\approx \xi_E^\eta(\epsilon')$. Consequently, $(\mathcal{M}, s) \approx \mathbb{W}\xi_E^\eta(\epsilon) : \bigwedge_{i \in E} \mathbb{W}\xi_i^\eta(\epsilon)$, with $\mathbb{W}\xi_i^\eta$ as in (2). Thus, the collection

$$X_\epsilon^\eta = \{\mathbb{W}\xi_i^\eta(\epsilon) : i \in \mathcal{D}\}$$

is finitely jointly satisfiable with T . By the Compactness Theorem 6.31, there exists a model (\mathcal{M}, s) of $T \cup X_\epsilon^\eta$. On the one hand, s_\bullet is $[\epsilon, \eta]$ -unstable because $\mathcal{M} \approx X_\epsilon^\eta$; on the other hand, s_\bullet is (r, η) -metastable by hypothesis, since $(\mathcal{M}, s) \approx T$. This is a contradiction since $r < \epsilon$. \square

The Uniform Metastability Principle may be formulated as the following dichotomy: For a fixed uniform $L[\mathbf{s}]$ -theory T , real $r \geq 0$ and $\eta \in \text{Smpl}(\mathcal{D})$,

- Either: There exists a rate $E_\bullet = (E_\epsilon : \epsilon > r)$ of (r, η) -metastability for s_\bullet valid uniformly for all models (\mathcal{M}, s) of T ;
- Or else: There exists a model (\mathcal{M}, s) of T such that s_\bullet fails to be (r, η) -metastable.

Below, we state a form of the Uniform Metastability Principle that applies when many internal nets are realized in the same structure. Fix a pointed directed set (\mathcal{D}, \leq, j_0) and a signature L for \mathcal{D} -directed structures. Extend L with function symbols $\mathbf{s} : \mathbb{D} \rightarrow \mathbb{S}_i$ $\sigma : \mathbb{A} \times \mathbb{D} \rightarrow \mathbb{S}_i$ where \mathbb{A} is of the form $\mathbb{S}_{i_1} \times \cdots \times \mathbb{S}_{i_n}$ (a Cartesian product of sorts of L) and \mathbb{S}_i is any sort of L . We will treat \mathbb{A} itself as a sort, writing $\mathbb{A}^{\mathcal{M}}$ for $\mathbb{S}_{i_1}^{\mathcal{M}} \times \cdots \times \mathbb{S}_{i_n}^{\mathcal{M}}$ in any L -structure \mathcal{M} . If (\mathcal{M}, σ, s) is an $L[\sigma, \mathbf{s}]$ -structure, we will say that s *admits σ -parameters of size $\leq C$* if there exists $\bar{a} \in (\mathbb{A}^{\mathcal{M}})^{[C]}$ such that $s(j) = \sigma(\bar{a}, j)$ for all $j \in \mathbb{D}^{\mathcal{M}}$. The slightly more general notion that s *approximately admits σ -parameters of size $\lesssim C$* means that, for all $\epsilon > 0$, there exists $\bar{a} \in (\mathbb{A}^{\mathcal{M}})^{[C+\epsilon]}$ such that $d(s(j), \sigma(\bar{a}, j)) \leq \epsilon$ for all $j \in \mathbb{D}^{\mathcal{M}}$. If C is rational, the latter property of $L[\sigma, \mathbf{s}]$ -structures (\mathcal{M}, σ, s) is captured by a single axiom

$$v_C : (\exists_C \bar{a})(\forall j)(d(s(j), \sigma(\bar{a}, j)) \leq 0). \quad (3)$$

In general, the property is axiomatized by the scheme $\{v_D : \text{rational } D > C\}$.

2.5. Corollary (Uniform metastability of parametrized sequences). *Fix a pointed directed set (\mathcal{D}, \leq, j_0) and a signature L for \mathcal{D} -directed structures. Extend L with function symbols $\mathbf{s} : \mathbb{D} \rightarrow \mathbb{S}_i$ and $\sigma : \mathbb{A} \times \mathbb{D} \rightarrow \mathbb{S}_i$ where \mathbb{A} is of the form $\mathbb{S}_{i_1} \times \cdots \times \mathbb{S}_{i_n}$ (a Cartesian product of sorts of L) and \mathbb{S}_i is any sort of L . Fix a rational $C \geq 0$ and a uniform $L[\sigma, \mathbf{s}]$ -theory T such that s approximately admits σ -parameters of size $\lesssim C$ for every model $(\mathcal{M}, \sigma, s) \approx T$ (i.e., $T \models v_D$ for all rational $D > C$ with v_D defined in (3) above). For some $r \geq 0$ and $\eta \in \text{Smpl}(\mathcal{D})$, assume that the external \mathcal{D} -net s_\bullet is (r, η) -metastable whenever $(\mathcal{M}, \sigma, s) \approx T$. Then there exists a collection $E_\bullet = (E_\epsilon : \epsilon > r)$ in $\mathcal{P}_{\text{fin}}(\mathcal{D})$ such that s_\bullet is E_\bullet -uniformly (r, η) -metastable whenever $(\mathcal{M}, \sigma, s) \approx T$. Moreover, E_\bullet depends only on T .*

As a simple application of saturation, the notions “ s admits σ -parameters of size $\leq C$ ” and “ s approximately admits σ -parameters of size $\lesssim C$ ” are seen to be equivalent in any ω^+ -saturated $L[\sigma, \mathbf{s}]$ -structure (\mathcal{M}, σ, s) . Since every model of T admits an ω^+ -saturated elementary extension, the hypotheses of Corollary 2.5 may be relaxed to state that the external \mathcal{D} -nets s_\bullet are (r, η) -metastable whenever s admits (exact) parameters of size $\leq C$ via σ (without changing the conclusion).

Proof. By the Uniform Metastability Principle 2.4, if the collection E_\bullet did not exist, there would be a model (\mathcal{M}, σ, s) of T such that s_\bullet is not (r, η) -metastable, contradicting the hypotheses. Hence, the metastability rate E_\bullet must exist, and depends only on T . (Note that E_\bullet implicitly depends on C .) \square

3. LOEB STRUCTURES

In this section, L will be a many-sorted signature with sorts $(\mathbb{S}_i : i < \alpha)$ (for some ordinal α), where $\mathbb{S}_0 = \mathbb{R}$ (as convened), $\mathbb{S}_1 = \Omega$, and $\mathbb{S}_2 = \mathcal{A}$. The sort Ω will be interpreted as a set with anchor point ω_0 , and \mathcal{A} will be interpreted as a Boolean algebra with its least element as anchor point.

Recall that, by default, L comes equipped with the following symbols:

- a binary function symbol $d_i : \mathbb{S}_i \times \mathbb{S}_i \rightarrow \mathbb{R}$ ($i < \alpha$) for the metric of \mathbb{S}_i (we may dispense with d_0 since $d_0(x, y) = |y - x|$ for $x, y \in \mathbb{R}$, and we shall denote d_1 by d_Ω , d_2 by $d_{\mathcal{A}}$);
- binary symbols $+\mathbb{R}$, $-\mathbb{R}$, $\cdot\mathbb{R}$ for the operations of addition, subtraction, multiplication of \mathbb{R} and a monadic function symbol $|\cdot|_{\mathbb{R}}$ for the absolute value of \mathbb{R} ;
- binary symbols $\wedge_{\mathbb{R}}$ (“minimum”) and $\vee_{\mathbb{R}}$ (“maximum”) for the lattice operations of \mathbb{R} ;
- an \mathbb{R} -valued constant c_r for each rational number $r \in \mathbb{R}$.

We will require that, in addition to these symbols, the signature L include the following:

- a symbol for the *measure*:

$$\mu : \mathcal{A} \rightarrow \mathbb{R};$$

- function symbols

$$\cup : \mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A}, \quad \cap : \mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A}, \quad {}^c : \mathcal{A} \rightarrow \mathcal{A}$$

for the union (join), intersection (meet) and complementation operations of the interpretation of \mathcal{A} ;

- a symbol $[\cdot \in \cdot] : \Omega \times \mathcal{A} \rightarrow \{0, 1\}$;
- \mathcal{A} -valued constants \emptyset, Ω for the zero (null element) and unity (universal element) of the interpretation of \mathcal{A} ;
- an Ω -valued constant c_{ω_0} for the distinguished element ω_0 of the interpretation of Ω .

The signature L may include other function and constant symbols, as well as many other sorts than those already mentioned.

In order to simplify the notation, we normally omit the subscripts in the operations $+\mathbb{R}$, $-\mathbb{R}$, $\cdot\mathbb{R}$, $\wedge_{\mathbb{R}}$, $\vee_{\mathbb{R}}$, and $|\cdot|_{\mathbb{R}}$ of \mathbb{R} . Also, for a rational number r , we denote the constant c_r simply as r .

If \mathcal{M} is a fixed L -structure, we will denote the sorts of \mathcal{M} corresponding to \mathcal{A} and Ω also as \mathcal{A} and Ω . Moreover, if the context allows it, we will identify the symbols of L with their interpretation in \mathcal{M} ; thus, for instance, we denote $\cup^{\mathcal{M}}$ and $\Omega^{\mathcal{M}}$ simply as \cup and Ω (respectively).

3.1. Definition (Loeb structure). A signature L as described above will henceforth be called a *Loeb signature*. A *pre-Loeb finite measure structure* is an L -structure \mathcal{M} such that:

- the metrics d_Ω on Ω and $d_{\mathcal{A}}$ on \mathcal{A} are discrete;
- $[\cdot \in \cdot]$ is a $\{0, 1\}$ -valued function identifying $(\mathcal{A}, \emptyset, \Omega, \cup, \cap, \cdot^c)$ with an algebra of subsets of Ω , i.e., for all $A, B \in \mathcal{A}$ and $\omega \in \Omega$:
 - $d_{\mathcal{A}}(A, B) = \sup_{x \in \Omega} |[\![x \in A]\!] - [\![x \in B]\!]|$;
 - $[\![\omega \in \emptyset]\!] = 0$;
 - $[\![\omega \in \Omega]\!] = 1$;
 - $[\![\omega \in A \cup B]\!] = [\![\omega \in A]\!] \vee [\![\omega \in B]\!]$;
 - $[\![\omega \in A \cap B]\!] = [\![\omega \in A]\!] \wedge [\![\omega \in B]\!]$;

- $\llbracket \omega \in A \rrbracket + \llbracket \omega \in A^c \rrbracket = 1$;
- the interpretation μ of $\boldsymbol{\mu}$ is a finitely additive measure on \mathcal{A} . For some $C \geq 0$:
 - $\mu(\emptyset) = 0$ and $0 \leq \mu(A) \leq \mu(\Omega) =: \|\mu\| \leq C$ for all $A \in \mathcal{A}$;
 - $\mu(A \cup B) + \mu(A \cap B) = \mu(A) + \mu(B)$ for all $A, B \in \mathcal{A}$.

In a *pre-Loeb signed measure structure* \mathcal{M} , the last axioms are replaced by:

- For some $C \geq 0$, $\mu = \boldsymbol{\mu}^{\mathcal{M}}$ is a signed measure on \mathcal{A} with total variation at most C :
 - $\mu(\emptyset) = 0$ and $\mu(A \cup B) + \mu(A \cap B) = \mu(A) + \mu(B)$ for all $A, B \in \mathcal{A}$;
 - $\|\mu\| := \sup\{|\mu(A)| + |\mu(B)| - |\mu(A \cap B)| : A, B \in \mathcal{A}\} \leq C$.

A *pre-Loeb probability structure* is a pre-Loeb finite measure structure \mathcal{M} with $\|\mu\| = 1$. A *Loeb probability (finite measure, signed measure) structure* is any λ^+ -saturated pre-Loeb probability (finite measure, signed measure) structure, where $\lambda = \text{card}(L)$.

Clearly, a positive pre-Loeb structure is a signed pre-Loeb structure. It can be easily verified that the equality $\|\mu\| = \sup\{|\mu(A)| + |\mu(B)| - |\mu(A \cap B)| : A, B \in \mathcal{A}\}$ holds in any pre-Loeb measure structure (probability, finite or positive). On the other hand, in signed measure structures only the inequality $|\mu(\Omega)| \leq \|\mu\|$ holds in general.

Recall (Definition 6.22) that if L' is a signature, a class \mathcal{C} of metric structures is said to be L' -axiomatizable if it consists of the models of a (positive bounded) L' -theory.

3.2. Proposition. *Let L be a signature for Loeb structures. For every fixed $C \geq 0$, the class of all pre-Loeb probability (finite measure, signed measure) L -structures \mathcal{M} such that $\|\mu\| \leq C$ is axiomatizable.*

(The assumption $C = 1$ in the case of probability structures is tacit.)

Proof. Given $C \geq 0$, we have to show that the clauses of Definition 3.1 are equivalent to the (approximate) satisfaction of positive bounded L -sentences. Clearly, it suffices to do so under the assumption that C is rational (if C is irrational, simply take the union of all axiom schemes for rational $D > C$). This is a routine exercise, so we give only one example. The condition $\|\mu\| \leq C$ amounts to the approximate satisfaction of the sentence⁵

$$(\forall A, B)(|\boldsymbol{\mu}(A)| + |\boldsymbol{\mu}(B)| - |\boldsymbol{\mu}(A \cap B)| \leq C).$$

The reader is invited to write down formulas axiomatizing the remaining clauses. □

3.3. Definition. For any pre-Loeb structure \mathcal{M} , define

$$\begin{aligned} \mathcal{A}^{\mathcal{M}} &\rightarrow \mathcal{P}(\Omega^{\mathcal{M}}) \\ A &\mapsto [A]^{\mathcal{M}} = \{\omega \in \Omega^{\mathcal{M}} : \llbracket \omega \in A \rrbracket = 1\}. \end{aligned}$$

The collection

$$[\mathcal{A}]^{\mathcal{M}} = \{[A]^{\mathcal{M}} : A \in \mathcal{A}\}$$

of subsets of $\Omega^{\mathcal{M}}$ is the *induced (external) algebra in $\Omega^{\mathcal{M}}$* . It follows from the definition of pre-Loeb structure that $[\mathcal{A}]^{\mathcal{M}}$ is an algebra of subsets of $\Omega^{\mathcal{M}}$ (the map $A \mapsto [A]^{\mathcal{M}}$ is an isomorphism between the Boolean algebras \mathcal{A} and $[\mathcal{A}]^{\mathcal{M}}$).

The *(external) measure induced by μ on \mathcal{M}* is the real-valued function

$$\begin{aligned} [\mu]^{\mathcal{M}} : [\mathcal{A}]^{\mathcal{M}} &\rightarrow \mathbb{R} \\ [A]^{\mathcal{M}} &\mapsto \mu(A). \end{aligned}$$

⁵We remark that the value $D = \|\mu\|$ is itself uniquely characterized by the formula $(\forall A, B)(|\boldsymbol{\mu}(A)| + |\boldsymbol{\mu}(B)| - |\boldsymbol{\mu}(A \cap B)| \leq D) \wedge (\exists A, B)(|\boldsymbol{\mu}(A)| + |\boldsymbol{\mu}(B)| - |\boldsymbol{\mu}(A \cap B)| \geq D)$.

$[\mu]^{\mathcal{M}}$ is well defined because $A \mapsto [A]^{\mathcal{M}}$ is injective.

When \mathcal{M} is fixed, we usually write $[\mu]$, $[\mathcal{A}]$, $[A]$ for $[\mu]^{\mathcal{M}}$, $[\mathcal{A}]^{\mathcal{M}}$, $[A]^{\mathcal{M}}$. We may also write $[\Omega]$ instead of $\Omega^{\mathcal{M}}$, for emphasis.

3.4. Proposition. *Fix a Loeb structure \mathcal{M} and let \mathcal{A}_L be the σ -algebra of subsets of Ω generated by $[\mathcal{A}]$.*

- (1) *The function $[\mu]$ extends to a unique countably additive positive [signed] measure μ_L on \mathcal{A}_L with total variation $\text{var}(\mu_L) = \|\mu\|$.*
- (2) *For all $S \in \mathcal{A}_L$ and $\epsilon > 0$ there exist $A, B \in \mathcal{A}$ such that $[A] \subset S \subset [B]$ and $\|\mu \upharpoonright (B \setminus A)\| \leq \epsilon$, where*

$$\begin{aligned} \mu \upharpoonright X : \mathcal{A} &\rightarrow \mathbb{R} \\ Y &\mapsto \mu(Y \cap X) \end{aligned}$$

for any $X \in \mathcal{A}$.

- (3) *For every $S \in \mathcal{A}_L$ there exists $A \in \mathcal{A}$ such that $|\mu_L|(S \Delta [A]) = 0$, where $X \Delta Y = (X \setminus Y) \cup (Y \setminus X)$ is the symmetric set difference and $|\mu_L|$ is the (nonnegative) absolute measure of μ_L .*

Proof. We prove the statements for μ a positive finite measure, leaving the case of a signed measure to the reader.

- (1) Clearly, $[\mu]$ is a finitely additive nonnegative measure on $([\Omega], [\mathcal{A}])$ with finite total variation $\text{var}([\mu]) = \mu(\Omega) = \|\mu\|$. Suppose that $([A_i] : i \in \mathbb{N})$ is a descending chain in $[\mathcal{A}]$ (for some descending chain (A_i) in \mathcal{A}). Assume that, for every $i \in \mathbb{N}$, $[A_i]$ is nonempty. Choose $\omega_i \in [A_i]$. By saturation, since (A_i) is descending, we have $\llbracket \omega \in A_i \rrbracket = 1$ for some ω and all $i \in \mathbb{N}$. Certainly, $\omega \in \bigcap_{i \in \mathbb{N}} [A_i]$. Therefore, if $\bigcap_{i \in \mathbb{N}} [A_i] = \emptyset$, we must have $A_j = \emptyset$ for some $j \in \mathbb{N}$, hence $\inf_{i \in \mathbb{N}} [\mu]([A_i]) = [\mu]([A_j]) = \mu(A_j) = \mu(\emptyset) = 0$. By taking relative complements, the preceding argument implies that $[\mu]$ is countably additive on $[\mathcal{A}]$, hence a premeasure thereon (in fact, if a set $B \in [\mathcal{A}]$ is a countable union of sets in $[\mathcal{A}]$, then B is necessarily a finite union of such sets). By the Carathéodory Extension Theorem, $[\mu]$ admits an extension to a countably-additive measure on $\sigma[\mathcal{A}] = \mathcal{A}_L$ with total variation $\text{var}(\mu_L) = \text{var}([\mu]) = \|\mu\|$. The extension is unique because $[\mu]$ has finite total variation.
- (2) This assertion also follows from Carathéodory's theorem.
- (3) Let $S \in \mathcal{A}_L$. By part (2), for each $n \in \mathbb{N}$ we may choose $A_n, B_n \in \mathcal{A}$ with $[A_n] \subset S \subset [B_n]$ and $\mu(B_n \setminus A_n) \leq 1/(n+1)$. Clearly, $\inf_n \mu(B_n \setminus A_n) = 0$. Without loss of generality, (A_n) is increasing and (B_n) decreasing. By saturation, there exists $A \in \mathcal{A}$ such that $A_n \subset A \subset B_n$ for all $n \in \mathbb{N}$. Since (A_n) is increasing and (B_n) decreasing, $U = \bigcup_{n \in \mathbb{N}} [A_n] \subset [A] \subset \bigcap_{n \in \mathbb{N}} [B_n] = V$, and also $U \subset S \subset V$. We have $\mu_L(S \Delta [A]) \leq \mu_L(V \setminus U) = \inf_n \mu(B_n \setminus A_n) = 0$, so A is as required. \square

3.5. Definition. We call \mathcal{A}_L the *Loeb algebra* of the Loeb structure \mathcal{M} , and μ_L the *Loeb measure* on $[\Omega]$ (i.e., on $\Omega^{\mathcal{M}}$).

3.6. Remarks.

- (1) The proof of Proposition 3.4 given above is an adaptation of the classical construction of Loeb measures [Loe75] (see also the articles by Cutland [Cut83, Cut00], and Ross [Ros97], on which we base our approach). Our context differs from the classical one in the sense that we do not need to use nonstandard universes or hyperreals, and our measures need not be probability measures.

- (2) Our definition of Loeb algebra differs from the classical definition, according to which the Loeb algebra is the completion $\overline{\mathcal{A}_L}$ of \mathcal{A}_L relative to μ_L and every subset U of an μ_L -null set is declared to be $\overline{\mathcal{A}_L}$ -measurable and null. While the classical definition has the advantage that the converse statements to (2) and (3) of Proposition 3.4 hold, we prefer to avoid completing \mathcal{A}_L so its definition is independent of μ .

4. INTEGRATION STRUCTURES

In this section, L will denote a signature for Loeb structures (with sorts \mathbb{R} , Ω , \mathcal{A}) such that, in addition to all the constant and function symbols required for Loeb structures, L includes a sort $\mathcal{L}_\Omega^\infty$ (whose interpretation will be a Banach lattice-algebra) and the following function symbols:

- a monadic symbol $\|\cdot\|_{\mathcal{L}_\Omega^\infty} : \mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$ (to be interpreted as the Banach norm of sort $\mathcal{L}_\Omega^\infty$);
- function symbols $\sup : \mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$ and $\inf : \mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$;
- binary symbols

$$\begin{aligned} +_{\mathcal{L}_\Omega^\infty} &: \mathcal{L}_\Omega^\infty \times \mathcal{L}_\Omega^\infty \rightarrow \mathcal{L}_\Omega^\infty, \\ -_{\mathcal{L}_\Omega^\infty} &: \mathcal{L}_\Omega^\infty \times \mathcal{L}_\Omega^\infty \rightarrow \mathcal{L}_\Omega^\infty, \\ \cdot_{\mathcal{L}_\Omega^\infty} &: \mathcal{L}_\Omega^\infty \times \mathcal{L}_\Omega^\infty \rightarrow \mathcal{L}_\Omega^\infty, \end{aligned}$$

and $\cdot_{\mathcal{L}_\Omega^\infty} : \mathbb{R} \times \mathcal{L}_\Omega^\infty \rightarrow \mathcal{L}_\Omega^\infty$ (to be interpreted as the algebra operations of $\mathcal{L}_\Omega^\infty$ and the scalar multiplication of the interpretation of $\mathcal{L}_\Omega^\infty$, respectively);

- binary symbols

$$\wedge_{\mathcal{L}_\Omega^\infty} : \mathcal{L}_\Omega^\infty \times \mathcal{L}_\Omega^\infty \rightarrow \mathcal{L}_\Omega^\infty, \quad \vee_{\mathcal{L}_\Omega^\infty} : \mathcal{L}_\Omega^\infty \times \mathcal{L}_\Omega^\infty \rightarrow \mathcal{L}_\Omega^\infty$$

(to be interpreted as the lattice operations of the interpretation of $\mathcal{L}_\Omega^\infty$);

- monadic symbols $\sup_\Omega : \mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$ and $\inf_\Omega : \mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$;
- a binary symbol $\text{ev}_\Omega : \mathcal{L}_\Omega^\infty \times \Omega \rightarrow \mathbb{R}$;
- a monadic symbol $\chi : \mathcal{A} \rightarrow \mathcal{L}_\Omega^\infty$;
- $I : \mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$;
- $\mathcal{L}_\Omega^\infty$ -valued constant symbols $0_{\mathcal{L}_\Omega^\infty}$ and $1_{\mathcal{L}_\Omega^\infty}$.

A signature including the sorts and symbols above will be called a *signature for integration structures*. Such a signature may include many other sorts, function symbols, and constant symbols.

4.1. Notation. If L is a signature for integration structures, \mathcal{M} is an L -structure, $f \in (\mathcal{L}_\Omega^\infty)^\mathcal{M}$, $\omega \in \Omega^\mathcal{M}$ and $A \in \mathcal{A}^\mathcal{M}$, we write $\text{ev}_\Omega^\mathcal{M}(f, \omega)$ for $f(\omega)$, χ_A for $\chi(A)$, and If for $I(f)$.

If \mathcal{M} is an L -structure and \mathcal{M} is fixed by the context, in order to simplify the notation, we will denote the sort corresponding to $\mathcal{L}_\Omega^\infty$ in \mathcal{M} as $\mathcal{L}_\Omega^\infty$, rather than $(\mathcal{L}_\Omega^\infty)^\mathcal{M}$. Carrying this simplification one step further, we will remove the \mathcal{M} -superscript from the interpretations of the function symbols in \mathcal{M} ; thus, for instance, we write $\|\cdot\|_{\mathcal{L}_\Omega^\infty}$ and $+_{\mathcal{L}_\Omega^\infty}$ instead of $(\|\cdot\|_{\mathcal{L}_\Omega^\infty})^\mathcal{M}$ and $(+_{\mathcal{L}_\Omega^\infty})^\mathcal{M}$, respectively.

4.2. Definition (Integration structure). Let L be a signature for integration structures. A *positive pre-integration structure* is an L -structure \mathcal{M} such that:

- The interpretation of $(\mathcal{L}_\Omega^\infty, +_{\mathcal{L}_\Omega^\infty}, -_{\mathcal{L}_\Omega^\infty}, \cdot_{\mathcal{L}_\Omega^\infty}, \cdot_{\mathcal{L}_\Omega^\infty}, \wedge_{\mathcal{L}_\Omega^\infty}, \vee_{\mathcal{L}_\Omega^\infty}, \|\cdot\|_{\mathcal{L}_\Omega^\infty}, 0_{\mathcal{L}_\Omega^\infty})$ in \mathcal{M} is a Banach algebra and Banach lattice with anchor $0_{\mathcal{L}_\Omega^\infty}$ (the *zero function*);

- for all $f \in \mathcal{L}_\Omega^\infty$ and $\omega \in \Omega$, we have:
 - the metric on $\mathcal{L}_\Omega^\infty$ is that induced by the norm $\|\cdot\|_{\mathcal{L}_\Omega^\infty}$;
 - $\|f\|_{\mathcal{L}_\Omega^\infty} = \sup_{\omega' \in \Omega} |f(\omega')|$;
 - $\text{ev}_\Omega(\cdot, \omega)$ is an algebra and lattice homomorphism;
 - $\sup f = \sup_{\omega' \in \Omega} f(\omega')$ and $\inf f = \inf_{\omega' \in \Omega} f(\omega')$;
 - $0_{\mathcal{L}_\Omega^\infty}(\omega) = 0$ and $1_{\mathcal{L}_\Omega^\infty}(\omega) = 1$;
- each $f \in \mathcal{L}_\Omega^\infty$ is *approximately \mathcal{A} -measurable* ($\mathcal{A} \approx \text{measurable}$) in the following sense:
 - For any reals $u < v$ there exists $A \in \mathcal{A}$ such that $f(\omega) \leq v$ if $\llbracket \omega \in A \rrbracket = 1$, and $f(\omega) \geq u$ if $\llbracket \omega \in A \rrbracket = 0$;
- $\chi_A(\omega) = \llbracket \omega \in A \rrbracket$ for all $A \in \mathcal{A}$ and $\omega \in \Omega$;
- for some $C \geq 0$:
 - \mathcal{M} is a pre-Loeb finite measure structure satisfying $\mu(A) = I\chi_A$ for all $A \in \mathcal{A}$, and $\|\mu\| := \mu(\Omega) \leq C$ (we also define $\|I\| = \|\mu\|$);
 - the *integration operation* I is a $\|\mu\|$ -Lipschitz linear functional $\mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$;
 - $\|\mu\| \inf f \leq If \leq \|\mu\| \sup f$ for all $f \in \mathcal{L}_\Omega^\infty$ (in particular, I is a positive functional: $If \geq 0$ if $f \geq 0$).

In a *signed pre-integration structure*, the last axiom becomes:

- For some $C \geq 0$:
 - \mathcal{M} is a pre-Loeb signed measure structure satisfying $\|\mu\| \leq C$ and $\mu(A) = I\chi_A$ for all $A \in \mathcal{A}$ (we define $\|I\| := \|\mu\|$);
 - I is a $\|\mu\|$ -Lipschitz linear functional $\mathcal{L}_\Omega^\infty \rightarrow \mathbb{R}$;
 - $|If| \leq \|\mu\| \|f\|_{\mathcal{L}_\Omega^\infty}$ for all $f \in \mathcal{L}_\Omega^\infty$.

A *probability pre-integration structure* is a positive integration structure with $\|\mu\| = 1$. A *probability* (resp., *positive*, *signed*) *integration structure* is any λ^+ -saturated probability (resp., positive, signed) pre-integration structure, where $\lambda = \text{card}(L)$.

For $f \in \mathcal{L}_\Omega^\infty$, we let $f_+ = (f \vee 0)$, $f_- = (-f)_+$ (so $f = f_+ + f_-$ with $f_+, f_- \geq 0$), and $|f| = f_+ + f_-$ (so $|f|(\omega) = |f(\omega)|$ for all $\omega \in \Omega$). For syntactic convenience, we treat the order of the lattice $\mathcal{L}_\Omega^\infty$ as part of the language, so we write $f \leq g$ to mean $\|(f \wedge g) - f\|_{\mathcal{L}_\Omega^\infty} \leq 0$.

4.3. Remark. The usual definition of measurability suggests postulating that for every $f \in \mathcal{L}_\Omega^\infty$ and every interval $J \subset \mathbb{R}$ there shall exist $A \in \mathcal{A}$ such that $\omega \in A$ if and only if $f(\omega) \in J$. However, exact measurability in this sense is not axiomatizable by positive bounded L -sentences. On the other hand, approximate measurability is axiomatizable as shown below in Proposition 4.6. Perhaps surprisingly, this postulate fails even in (saturated) probability structures. By working in saturated probability spaces and externally enlarging $[\mathcal{A}]$ to a σ -algebra $\sigma[\mathcal{A}]$ of subsets of Ω , the $\mathcal{A} \approx \text{measurability}$ axiom implies the exact $\sigma[\mathcal{A}]$ -measurability of all functions $[f] : \Omega \rightarrow \mathbb{R}$ for $f \in \mathcal{L}_\Omega^\infty$. See Proposition 4.6 below.

4.4. Proposition. *Let L be a language for pre-measure structures. For every $C \geq 0$, the class of all (probability, positive, or signed) pre-measure L -structures \mathcal{M} with $\|\mu^{\mathcal{M}}\| \leq C$ is axiomatizable in the logic of approximate satisfaction of positive bounded formulas.*

Proof. It is a routine exercise to verify that the axioms in the definition of integration structure can be written as sets of positive bounded L -sentences. As an example, the $\mathcal{A} \approx \text{measurability}$ condition amounts to the axiom schema

$$(\forall Df)(\exists A)(\forall \omega)(r\chi_{A^c} - D\chi_A \leq f(\omega) \leq r\chi_A + D\chi_{A^c})$$

for all rational $D > 0$ and r , as easily seen from the semantics of approximate satisfaction. The remaining axioms are handled similarly. \square

4.5. Notation. Given any function $G : \mathbb{R} \rightarrow \mathbb{R}$, an element $f \in \mathcal{L}_\Omega^\infty$, and a real number $t \in \mathbb{R}$, we let

$$\{G(f) \leq t\} := \{\omega \in \Omega : G(f(\omega)) \leq t\},$$

with similar definitions for $\{G(f) < t\}$, $\{G(f) \geq t\}$, and $\{G(f) > t\}$.

4.6. Proposition. *Let \mathcal{M} be an integration L -structure. For every $f \in \mathcal{L}_\Omega^\infty$, the function $[f] = [f]^M : \Omega \rightarrow \mathbb{R}$ defined by $\omega \mapsto f(\omega)$ is \mathcal{A}_L -measurable.*

Proof. Since \mathcal{A}_L is a σ -algebra, it suffices to show that $\{f \leq t\} \in \mathcal{A}_L$ for any fixed real number t .

Let $(u_n : n \in \mathbb{N})$ be a strictly decreasing sequence of rational numbers such that $\inf_n u_n = t$. By \mathcal{A} -measurability and saturation, for each $n \in \mathbb{N}$ there exists $A_n \in \mathcal{A}$ such that $f(\omega) \leq u_n$ if $\omega \in [A_n]$, and $f(\omega) \geq u_n$ if $\omega \notin [A_n]$. Clearly, $\{f \leq t\} = \bigcap_{n \in \mathbb{N}} [A_n] \in \mathcal{A}_L$. \square

4.7. Theorem (Riesz Representation Theorem for integration structures). *Let \mathcal{M} be an integration L -structure (positive or signed). For every $f \in \mathcal{L}_\Omega^\infty$, $[f]$ is Loeb-integrable and*

$$If = \int [f] d\mu_L.$$

Proof. We assume that I is positive, leaving the signed case to the reader. Let $C = \|I\| = \|\mu\|$ and $D = \|f\|_{\mathcal{L}_\Omega^\infty}$. $[f]$ is Loeb integrable because it is \mathcal{A}_L -measurable (Proposition 4.6) and bounded (by D). The Loeb measure μ_L is also positive with total variation $\text{var}(\mu_L) = \mu_L(\Omega) = \mu(\Omega) = \|\mu\| = C$. Let us write $\int F$ for $\int F d\mu_L$. If $A \in \mathcal{A}$, we have $I\chi_A = \mu(A) = \mu_L([A]) = \int [\chi_A]$. Fix $\epsilon > 0$ and let $(J_i : i < k)$ be any finite collection of disjoint intervals, each having length at most ϵ , such that $\bigcup_{i < k} J_i \supset [-D, D]$. For $i < k$, let $S_i = \{f \in J_i\} \subset \Omega$. By measurability of $[f]$ (Proposition 4.6), the collection $(S_i : i < k)$ is an \mathcal{A}_L -measurable disjoint cover of Ω , hence $\sum_{i < k} \mu_L(S_i) = \mu_L(\Omega) = C$. Choose rational numbers r_i ($i < k$) such that $J_i \subset [r_i, r_i + 2\epsilon]$. The \mathcal{A}_L -simple function $F_\epsilon = \sum_{i < k} r_i \chi_{S_i} : \Omega \rightarrow \mathbb{R}$ satisfies $F_\epsilon \leq [f] \leq F_\epsilon + 2\epsilon$, hence $\int F_\epsilon \leq \int [f] \leq \int F_\epsilon + 2C\epsilon$. By internal approximability (Proposition 3.4(3)), there exist $A_i \in \mathcal{A}$ ($i < k$) such that $\mu_L([A_i] \Delta S_i) = 0$; thus $\sum_{i < k} \mu(A_i) = \sum_{i < k} \mu_L(S_i) = C$. Without loss of generality we may assume that the sets A_i are pairwise disjoint. Let $B = (\bigcup_{i < k} A_i)^c$, so $\mu(B) = 0$. Let $f_\epsilon = f\chi_B + \sum_{i < k} r_i \chi_{A_i}$. Certainly, $f_\epsilon \leq f \leq f_\epsilon + 2\epsilon$, so $If_\epsilon \leq If \leq If_\epsilon + 2C\epsilon$. Since $\mu(B) = 0$, we have $If_\epsilon = \sum_{i < k} r_i \mu(A_i) = \sum_{i < k} r_i \mu_L(S_i) = \int F_\epsilon$. Thus, If and $\int [f]$ both lie in $[If_\epsilon, If_\epsilon + 2C\epsilon]$, so $|If - \int [f]| \leq 2C\epsilon$. Since C is fixed and this holds for all $\epsilon > 0$, we conclude that $If = \int [f]$. \square

5. DIRECTED INTEGRATION STRUCTURES

5.1. Definition (Directed integration structure). Fix a pointed directed set (\mathcal{D}, \leq, j_0) . Let L be both a signature for integration structures (with sorts $\mathbb{R}, \Omega, \mathcal{A}, \mathcal{L}_\Omega^\infty$) and also for \mathcal{D} -directed structures (with directed sort \mathbb{D}). In addition, assume that L has a function symbol

$$\varphi : \mathbb{D} \rightarrow \mathcal{L}_\Omega^\infty.$$

Such L will be called a *signature for directed integration structures* (and it may include any other sorts and symbols than those named).

A \mathcal{D} -directed pre-integration structure is an L -structure \mathcal{M} that is both a \mathcal{D} -directed L -structure and a pre-integration L -structure (positive or signed). A \mathcal{D} -directed integration structure is an λ^+ -saturated \mathcal{D} -directed pre-integration structure, where $\lambda = \text{card}(L)$.

Given any \mathcal{D} -directed pre-integration structure \mathcal{M} , let $\varphi = \varphi^{\mathcal{M}}$. Note that the definition of L -structure implies that φ is uniformly bounded on the discrete sort $\mathbb{D}^{\mathcal{M}}$. We define

$$\|\varphi\| := \sup_{j \in \mathbb{D}} \|\varphi(j)\|_{\mathcal{L}_{\Omega}^{\infty}}.$$

For $j \in \mathbb{D}$, we will denote $\varphi(j)$ by φ_j . Each $\omega \in \Omega$ defines an external \mathcal{D} -net $\varphi_{\bullet}(\omega) := (\varphi_j(\omega) : j \in \mathcal{D})$ in $\mathcal{L}_{\Omega}^{\infty}$. We also have a real-valued \mathcal{D} -net $I\varphi_{\bullet} := (I(\varphi_j) : j \in \mathcal{D})$.

5.2. Proposition. *Let L be a language for directed integration structures and $C \geq 0$. The class of all (positive or signed) \mathcal{D} -directed pre-integration L -structures \mathcal{M} such that $\|\varphi\| \leq C$ and $\|\mu\| \leq C$ is axiomatizable in the logic of approximate satisfaction.*

Proof. This follows from Propositions 3.2 and 4.4, upon remarking that the condition $\|\varphi\| \leq C$ is equivalent to the approximate satisfaction of the axioms⁶

$$(\forall j)(\|\varphi_j\|_{\mathcal{L}_{\Omega}^{\infty}} \leq r) \quad \text{for all rational } r > C. \quad \square$$

5.3. Proposition (Dominated Convergence Theorem). *Assume that the directed set \mathcal{D} is countable. Let \mathcal{M} be a (saturated) \mathcal{D} -directed integration L -structure, for a suitable language L . Then*

$$\text{osc}(I\varphi_{\bullet}) \leq \|I\| \sup_{\omega \in \Omega} \text{osc}(\varphi_{\bullet}(\omega)) \quad \text{for all } \omega \in \Omega.$$

In particular, $I\varphi_{\bullet}$ is convergent if $\varphi_{\bullet}(\omega)$ is convergent for each $\omega \in \Omega$.

Proof. We prove the inequality $\text{osc}(I\varphi_{\bullet}) \leq \sup_{\omega \in \Omega} \text{osc}(\varphi_{\bullet}(\omega))$ in the case when \mathcal{M} is a probability integration structure (with $\|I\| = 1$), leaving the general case to the reader. Let $r = \sup_{\omega \in \Omega} \text{osc}(\varphi_{\bullet}(\omega))$.

For $\epsilon \in \Omega$ and $i, j, j' \in \mathcal{D}$, let

$$\Omega_{\epsilon}^{j,j'} = \{\omega \in \Omega : |\varphi_{j'}(\omega) - \varphi_j(\omega)| \leq r + \epsilon\}$$

and

$$\Omega_{\epsilon}^i = \bigcap_{j,j' \in \mathcal{D}_{\geq i}} \Omega_{\epsilon}^{j,j'} = \{\omega \in \Omega : |\varphi_{j'}(\omega) - \varphi_j(\omega)| \leq r + \epsilon \text{ for all } j, j' \in \mathcal{D}_{\geq i}\}.$$

The functions $\varphi_j(\cdot)$ and $\varphi_{j'}(\cdot)$ are \mathcal{A}_L -measurable, by Proposition 4.6. It is clear that each set $\Omega_{\epsilon}^{j,j'}$ is also \mathcal{A}_L -measurable, and so is the (countable) intersection Ω_{ϵ}^i . Clearly, $\Omega_{\epsilon}^i \subset \Omega_{\epsilon}^{i'}$ for $i \leq i'$. Since $\omega \in \Omega_{\epsilon}^i$ implies $\text{osc}_{\eta_i}(\varphi_{\bullet}(\omega)) \leq r + \epsilon$, while $\text{osc}(\varphi_{\bullet}(\omega)) \leq r$ for all ω by hypothesis, part (2) of Proposition 1.4 gives:

$$\bigcup_{i \in \mathcal{D}} \Omega_{\epsilon}^i = \Omega.$$

⁶The value $D = \|\varphi\|$ is characterized by the approximate satisfaction of the formula

$$(\forall j)(\|\varphi_j\|_{\mathcal{L}_{\Omega}^{\infty}} \leq D) \wedge (\exists j)(\|\varphi_j\|_{\mathcal{L}_{\Omega}^{\infty}} \geq D).$$

Since \mathcal{D} is countable and μ_L is a probability measure, we have $\sup_{i \in \mathcal{D}} \mu_L(\Omega_\epsilon^i) = \mu_L(\Omega) = 1$, hence $\mu(\Omega_\epsilon^{i_0}) \geq 1 - \epsilon/(\|\varphi\| + 1)$ for some $i_0 \in \mathcal{D}$. For $j, j' \geq i_0$, we have:

$$\begin{aligned} |I(\varphi_{j'}) - I(\varphi_j)| &\leq I(|\varphi_{j'} - \varphi_j|) = \int |\varphi_{j'} - \varphi_j| \chi_{\Omega_\epsilon^{i_0}} + \int |\varphi_{j'} - \varphi_j| \chi_{(\Omega_\epsilon^{i_0})^c} \\ &\leq (r + \epsilon)\mu_L(\Omega_\epsilon^{i_0}) + 2\|\varphi\|\mu_L((\Omega_\epsilon^{i_0})^c) \leq r + 3\epsilon. \end{aligned}$$

Therefore, $\text{osc}(I\varphi_\bullet) \leq \text{osc}_{\eta_i}(I\varphi_\bullet) \leq r + 3\epsilon$ (again, by part (2) of Proposition 1.4). Since ϵ is an arbitrary positive number, $\text{osc}(I\varphi_\bullet) \leq r = \sup_{\omega \in \Omega} \text{osc}(\varphi_\bullet(\omega))$. \square

Note that the proof above uses only the standard theory of integration.

5.4. Corollary (Metastable Dominated Convergence Theorem). *Let (\mathcal{D}, \leq, j_0) be a directed set with \mathcal{D} countable. Let \mathcal{M} be a (not necessarily saturated) \mathcal{D} -directed pre-integration L -structure, for a suitable signature L . Fix a real number $s \geq 0$. Let T be any uniform theory including the axiom $\|I\| \leq s$ (that is, the axioms $\|I\| \leq u$ for all rationals $u > s$) and extending the theory of \mathcal{D} -directed pre-integration L -structures. Given $r \geq 0$ and any collection $E_\bullet^r = (E_{\epsilon, \eta} : \eta \in \text{Smpl}(\mathcal{D}), \epsilon > r)$ in $\mathcal{P}_{\text{fin}}(\mathcal{D})$, there exists another collection $\widetilde{E}_\bullet^{rs} = (\widetilde{E}_{\epsilon, \eta} : \eta \in \text{Smpl}(\mathcal{D}), \epsilon > rs)$ such that every model \mathcal{M} of T satisfies the following property:*

If every external \mathcal{D} -net in the collection $\mathcal{C} = (\varphi_\bullet(\omega) : \varphi \in (\mathcal{L}_\Omega^\infty)^{[1]}, \omega \in \Omega)$ is E_\bullet^r -uniformly r -metastable, then $\widetilde{E}_\bullet^{rs}$ is a rate of rs -metastability for the collection $(I\varphi_\bullet : \varphi \in (\mathcal{L}_\Omega^\infty)^{[1]})$.

In fact, one such $\widetilde{E}_\bullet^{rs}$ may be found depending only on r, s , and E_\bullet^r .

Proof. Let $r, s \geq 0$ and E_\bullet^r be given. Restrict the signature L so it only names the sorts and symbols strictly required for a signature of \mathcal{D} -directed pre-integration structures. Let $\widetilde{T} = T[\boldsymbol{\sigma}, \mathbf{a}]$ where $\mathbf{a} : \mathbb{D} \rightarrow \mathbb{R}$ and $\boldsymbol{\sigma} : \mathcal{L}_\Omega^\infty \times \Omega \times \mathbb{D} \rightarrow \mathbb{R}$ are new function symbols, and let T be the theory of pre-integration L -structures augmented with the following \widetilde{L} -axioms:

- $\|I\| \leq C$, for all rational $C > s$;
- $(\forall_C \varphi)(\forall j)(\forall \omega)(|\boldsymbol{\sigma}(\varphi, \omega, j) - \varphi_j(\omega)| \leq 0)$, for each rational $C \geq 0$;
- $(\exists_1 \varphi)(\exists \omega)(\forall j)(|\mathbf{a}(j) - \boldsymbol{\sigma}(\varphi, \omega, j)| \leq 0)$.

Clearly, T is a uniform \widetilde{L} -theory, and $\boldsymbol{\sigma}$ is $(T \upharpoonright L[\boldsymbol{\sigma}])$ -definable in L ; in fact, every model \mathcal{M} of $T \upharpoonright L$ admits a unique expansion to a model $(\mathcal{M}, \boldsymbol{\sigma})$ of $T \upharpoonright L[\boldsymbol{\sigma}]$, and any such model admits some expansion $(\mathcal{M}, \boldsymbol{\sigma}, a)$ to a model of T via any $a : \mathbb{D}^{\mathcal{M}} \rightarrow \mathbb{R}$ that approximately admits σ -parameters of size $\lesssim 1$. Using Proposition 2.3, extend T to a (necessarily uniform) \widetilde{L} -theory \widetilde{T} such that a_\bullet is E_\bullet^r -uniformly r -metastable for all models $(\mathcal{M}, \boldsymbol{\sigma}, a)$ of \widetilde{T} . Henceforth all L' -structures for a language $L' \subset \widetilde{L}$ are assumed to be models of $\widetilde{T} \upharpoonright L'$. If $(\mathcal{M}, \boldsymbol{\sigma}) \approx \widetilde{T} \upharpoonright L[\boldsymbol{\sigma}]$, every external sequence $\varphi_\bullet(\omega)$ with $\|\varphi\| \leq 1$ admits σ -parameters (φ, ω) of size ≤ 1 , and every such $\varphi_\bullet(\omega)$ is of the form a_\bullet for some expansion $(\mathcal{M}, \boldsymbol{\sigma}, a) \approx \widetilde{T}$ of $(\mathcal{M}, \boldsymbol{\sigma})$. By Proposition 1.4, $\text{osc}(a_\bullet) \leq r$, hence $\text{osc}(\varphi_\bullet(\omega)) \leq r$ whenever $\|\varphi\| \leq 1$. Since $\|I\| \leq s$, we have $\text{osc}(I\varphi_\bullet) \leq rs$ by Proposition 5.3. Let $\mathbf{b} : \mathbb{D} \rightarrow \mathbb{R}$ and $\boldsymbol{\tau} : \mathcal{L}_\Omega^\infty \times \mathbb{D} \rightarrow \mathbb{R}$ be new function symbols. Expand \widetilde{T} to an $\widetilde{L}[\boldsymbol{\tau}, \mathbf{b}]$ -theory \widetilde{T}' by adding the axioms

- $(\forall_C \varphi)(\forall j)(|\boldsymbol{\tau}(\varphi, j) - I\varphi_j| \leq 0)$, for each rational $C \geq 0$;
- $(\exists_1 \varphi)(\forall j)(|\mathbf{b}(j) - \boldsymbol{\tau}(\varphi, j)| \leq 0)$.

Clearly, \widetilde{T}' is uniform, $\boldsymbol{\tau}$ is $(\widetilde{T}' \upharpoonright \widetilde{L}[\boldsymbol{\tau}])$ -definable in \widetilde{L} , and an expansion $(\mathcal{M}, \boldsymbol{\tau}, b)$ of a model $(\mathcal{M}, \boldsymbol{\tau})$ of $\widetilde{T}' \upharpoonright \widetilde{L}[\boldsymbol{\tau}]$ to a model of \widetilde{T}' is via $b : \mathbb{D}^{\mathcal{M}} \rightarrow \mathbb{R}$ that approximately admits τ -parameters of size $\lesssim 1$. Henceforth, all L' -structures for any signature $L' \subset \widetilde{L}[\boldsymbol{\tau}, \mathbf{b}]$ are assumed to

be models of $\widetilde{T}' \upharpoonright L'$. Since \widetilde{T} implies that $\text{osc}(I\varphi_\bullet) \leq rs$ whenever $\|\varphi\| \leq 1$, \widetilde{T}' implies⁷ that $\text{osc}(b_\bullet) \leq rs$. By uniform metastability of parametrized sequences (Corollary 2.5), there exists a collection $\widetilde{E}_\bullet^{rs}$ depending only on \widetilde{T}' such that all sequences \widetilde{b}_\bullet , and hence all sequences $I\varphi_\bullet$ for $\|\varphi\| \leq 1$, are $\widetilde{E}_\bullet^{rs}$ -uniformly rs -metastable. Moreover, $\widetilde{E}_\bullet^{rs}$ depends only on \widetilde{T}' , hence only on r, s and the given rate E_\bullet^r . \square

6. BACKGROUND ON METRIC MODEL THEORY

This section describes the general framework for the classes of structures that are the focus of the paper. We refer to these structures as *metric structures*.

HENSON METRIC STRUCTURES

Recall that a *pointed metric space* is a triple (M, d, a) , where (M, d) is a metric space and a is a distinguished element of M called the *anchor* of M . If (M, d, a) is a pointed metric space, the closed ball of radius r around the anchor point a will be denoted $B_M[r]$, or simply $B[r]$ if the ambient space M is clear from the context; the corresponding open ball will be denoted $M^{(r)}$ or $B(r)$. If $(M_1, d_1, a_1), \dots, (M_n, d_n, a_n)$ are pointed metric spaces, we regard the product $\prod_{i=1}^n (M_i, d_i, a_i)$ tacitly as a pointed metric space by taking (a_1, \dots, a_n) as its anchor and using the supremum metric.

6.1. Definition. A *metric* (or *Henson*) *structure* \mathcal{M} (often just called a structure in this manuscript) consists of the following items:

- A family $(M^{(s)} \mid s \in \mathbf{S})$ of pointed metric spaces,
- A collection of functions of the form

$$F : M^{(s_1)} \times \dots \times M^{(s_n)} \rightarrow M^{(s_0)},$$

each of which is locally uniformly continuous, i.e., uniformly continuous on each bounded subset of its domain.

The spaces $M^{(s)}$ are called the *sorts* of \mathcal{M} . We say that \mathcal{M} is *based on* the collection $(M^{(s)} \mid s \in \mathbf{S})$ of its sorts.

We do require that every metric structure contain the set \mathbb{R} of real numbers, equipped with the usual distance and 0 as an anchor point, as a distinguished sort. We also require that the given metric on each sort of \mathcal{M} be included in the list of functions of \mathcal{M} , and that the anchor of each sort be included as a (constant) function.

If \mathcal{M} is based on $(M^{(s)} \mid s \in \mathbf{S})$ an element of $M^{(s)}$ will be called an *element* of \mathcal{M} of sort s . The *cardinality* of \mathcal{M} , denoted $\text{card}(\mathcal{M})$, is defined as $\sum_{s \in \mathbf{S}} \text{card}(M^{(s)})$.

Some of the sorts M of a structure may be discrete metric spaces, with the respective metric $d : M \times M \rightarrow \{0, 1\}$ taking the value 1 at every pair of distinct points. If all the sorts of \mathcal{M} are discrete, we will say that \mathcal{M} is a discrete structure. Similarly, if the sorts of \mathcal{M} are bounded, we will say that \mathcal{M} is a bounded metric structure.

Some of the functions of a structure \mathcal{M} may have arity 0. Such functions correspond to distinguished elements of \mathcal{M} . We will call these elements the *constants* of the structure. If

⁷*A priori*, \widetilde{T}' only implies that $\text{osc}(b_\bullet) \leq rs$ when b_\bullet admits τ -parameters of size ≤ 1 , so $b_\bullet : j \mapsto \tau(\varphi, j) = I\varphi_j$ for some $\|\varphi\| \leq 1$. If this is the case, however, then any b_\bullet that approximately admits τ -parameters of size $\lesssim 1$ must still satisfy $\text{osc}(b_\bullet) \leq rs$. Alternatively, the rest of the proof may proceed applying Corollary 2.5 in the stronger version that only assumes metastability when b_\bullet admits exact parameters of size ≤ 1 .

F is a $\{0, 1\}$ -valued function of \mathcal{M} , we will identify F with a subset of its domain, namely, $F^{-1}(1)$. Such a function will be called a *relation*, or a *predicate*, of \mathcal{M} .

We will require that the special sort \mathbb{R} should come equipped with the field operations of \mathbb{R} , the order relation and the lattice operations ($\max(x, y)$ and $\min(x, y)$), plus a constant for each rational number.

If a structure \mathcal{M} is based on $(M^{(s)} \mid s \in \mathbf{S})$ and $(F^i \mid i \in I)$ is a list of the functions of \mathcal{M} , we write

$$\mathcal{M} = (M^{(s)}, F_i \mid s \in \mathbf{S}, i \in I).$$

For notational simplicity, the real sort \mathbb{R} , the metrics on the sorts of \mathcal{M} , and their respective anchors need not be listed explicitly in this notation. We will only list them when needed for emphasis.

The structures that we will be dealing with are “hybrid” in the sense that some of their sorts are discrete, while others are genuine metric spaces. Typically, the nondiscrete structures will be Banach algebras or Banach lattices; in these the natural anchor point is 0. The discrete sorts that we will encounter include partial orders and purely algebraic structures; in structures of this type, the particular choice of anchor point is often inconsequential.

I. HENSON SIGNATURES AND STRUCTURE ISOMORPHISMS

We will need a formal way to index the sorts and functions of any given structure \mathcal{M} . This is accomplished through the concept of *signature* of a metric structure.

6.2. Definition. Let \mathcal{M} be a structure based on $(M^{(s)} \mid s \in \mathbf{S})$. A *Henson signature* L for \mathcal{M} consists of the following items:

- A sort index set \mathbf{S} ,
- A special element $s_{\mathbb{R}} \in \mathbf{S}$ such that $M^{(s_{\mathbb{R}})} = \mathbb{R}$,
- For each function $F : M^{(s_1)} \times \cdots \times M^{(s_n)} \rightarrow M^{(s_0)}$, a triple of the form

$$((s_1, \dots, s_n), f, s_0),$$

where f is a purely syntactic symbol called a *function symbol* for F . We write $F = f^{\mathcal{M}}$ and call F the *interpretation* of f in \mathcal{M} . We call $s_1 \times \cdots \times s_n$ and s_0 the *domain* and *range* of f , respectively. We express this by writing (purely formally)

$$f : s_1 \times \cdots \times s_n \rightarrow s_0.$$

The number n is called the *arity* of the function symbol f . If $n = 0$ and the constant value of $f^{\mathcal{M}}$ in $M^{(s_0)}$ is c , we call f a *constant symbol* for c .

We express the fact that L is a signature for \mathcal{M} by saying that \mathcal{M} is an *L -structure*. A structure \mathcal{M} for some Henson signature L will also be called a *Henson structure*. The *cardinality* of a signature L , denoted $\text{card}(L)$, is defined as

$$\text{card}(\mathbf{S}) + \text{card}(\{f \mid f \text{ is a function symbol of } L\}) + \aleph_0,$$

where \mathbf{S} is the sort index set of L .

6.3. Definition. If L and L' are signatures, we say that L is a *subsignature* of L' (or that L' is an *extension* of L), and write $L \subseteq L'$, if the following conditions hold:

- The sort index set of L is a subset of the sort index set of L' ,
- Every triple of the form $((s_1, \dots, s_n), f, s_0)$ that is in L is also in L' .

If L, L' are signatures, we say that L' is an *extension by constants* of L if L and L' have the same sort index set and every function symbol of L' that is not in L is a constant symbol. If the set of such constant symbols is C , we denote L' as $L[C]$.

6.4. Definition. Let L be a signature and let \mathcal{M} and \mathcal{N} be L -structures based on $(M^{(s)} \mid s \in \mathbf{S})$ and $(N^{(s)} \mid s \in \mathbf{S})$, respectively.

- (1) \mathcal{M} is a *substructure* of \mathcal{N} if $M^{(s)} \subseteq N^{(s)}$ and, for each function symbol f , the interpretation $f^{\mathcal{N}}$ of f in \mathcal{N} is an extension of $f^{\mathcal{M}}$.
- (2) \mathcal{M} and \mathcal{N} are *isomorphic* if there exists a family $\mathcal{J} = (\mathcal{J}^{(s)} \mid s \in \mathbf{S})$ of maps (called an *isomorphism* from \mathcal{M} into \mathcal{N}) such that for each $s \in \mathbf{S}$, $\mathcal{J}^{(s)} : M^{(s)} \rightarrow N^{(s)}$ is a bijection that commutes with the interpretations of the function symbols of L , in the sense that if $f : s_1 \times \cdots \times s_n \rightarrow s_0$, then $\mathcal{J}^{(s_0)}(f^{\mathcal{M}}(a_1), \dots, f^{\mathcal{M}}(a_n)) = f^{\mathcal{N}}(\mathcal{J}^{(s_1)}(a_1), \dots, \mathcal{J}^{(s_n)}(a_n))$. If a is an element of $M^{(s)}$ and the sort index s need not be made specific, we may write $\mathcal{J}(a)$ instead of $\mathcal{J}^{(s)}(a)$.
- (3) An *automorphism* of \mathcal{M} is an isomorphism between \mathcal{M} and \mathcal{M} .

II. UNIFORM CLASSES AND ULTRAPRODUCTS OF METRIC STRUCTURES

Recall that a *filter* on a nonempty set Λ is a collection \mathcal{F} of subsets of Λ such that (i) $\Lambda \in \mathcal{F}$ and $\emptyset \notin \mathcal{F}$, (ii) $A \cap B \in \mathcal{F}$ if $A, B \in \mathcal{F}$, and (iii) $A \in \mathcal{F}$ if $B \in \mathcal{F}$ and $A \supset B$. An *ultrafilter* on Λ is a maximal filter \mathcal{U} on Λ ; equivalently, \mathcal{U} is a filter such that (iv) $A \in \mathcal{U}$ or $\Lambda \setminus A \in \mathcal{U}$ for all $A \subset \Lambda$. If Λ is an index set and \mathcal{F} is a filter on Λ , we will say that a subset of Λ is \mathcal{F} -*large* if it is in \mathcal{F} . An ultrafilter \mathcal{U} on Λ is *principal* if there exists $\lambda_0 \in \Lambda$ such that $A \in \mathcal{U}$ iff $A \ni \lambda_0$ for all $A \subset \Lambda$; otherwise, \mathcal{U} is *nonprincipal*. If X is a topological space, $(x_\lambda)_{\lambda \in \Lambda}$ is a family of elements of X , and \mathcal{F} is a filter on Λ , we will say that $(x_\lambda)_{\lambda \in \Lambda}$ *converges* to an element $y \in X$ with respect to \mathcal{F} if for every neighborhood U of y , the set $\{\lambda \in \Lambda \mid x_\lambda \in U\}$ is \mathcal{F} -large. If X is compact Hausdorff, then for every family $x_\bullet = (x_\lambda)_{\lambda \in \Lambda}$ and every ultrafilter \mathcal{U} on Λ there exists a unique $y \in X$ such that $(x_\lambda)_{\lambda \in \Lambda}$ converges to y with respect to \mathcal{U} ; this element y is called the \mathcal{U} -*limit* of $(x_\lambda)_{\lambda \in \Lambda}$ and is denoted $\mathcal{U}\text{lim } x_\bullet$ or $\mathcal{U}\text{lim}_\lambda x_\lambda$.

Let $(X_\lambda, d_\lambda)_{\lambda \in \Lambda}$ be a family of metric spaces and let \mathcal{U} be an ultrafilter on Λ . The \mathcal{U} -*ultraproduct* of $(X_\lambda, d_\lambda)_{\lambda \in \Lambda}$ is the metric space defined in the following manner. Let $\ell^\infty(X_\lambda, d_\lambda)_\Lambda := \ell^\infty(X_\lambda, d_\lambda \mid \lambda \in \Lambda)$ be the set of all elements of $\prod_{\lambda \in \Lambda} X_\lambda$ that are bounded (when regarded as families indexed by Λ in the natural way). For $x = (x_\lambda)_{\lambda \in \Lambda}$, $y = (y_\lambda)_{\lambda \in \Lambda}$ in $\ell^\infty(X_\lambda, d_\lambda)_\Lambda$, and an ultrafilter \mathcal{U} on Λ , define

$$d(x, y) = \mathcal{U}\text{lim}_\lambda d_\lambda(x_\lambda, y_\lambda).$$

Since elements of $\ell^\infty(X_\lambda, d_\lambda)_\Lambda$ are bounded families, it is clear that d is well defined. It is also easy to verify that d is a pseudometric on $\ell^\infty(X_\lambda, d_\lambda)_\Lambda$. Now we can turn d into a metric in the usual way, namely by identifying any two elements $x, y \in \ell^\infty(X_\lambda, d_\lambda)_\Lambda$ such that $d(x, y) = 0$. For $x \in \ell^\infty(X_\lambda, d_\lambda)_\Lambda$, we let $(x)_\mathcal{U}$ denote the equivalence class of x under this identification, and for any two equivalence classes $(x)_\mathcal{U}, (y)_\mathcal{U}$, we define $d((x)_\mathcal{U}, (y)_\mathcal{U})$ as $d(x, y)$. The resulting metric space is called the \mathcal{U} -*ultraproduct* of the family $(X_\lambda, d_\lambda)_{\lambda \in \Lambda}$. It will be denoted $(\prod_{\lambda \in \Lambda} X_\lambda)_\mathcal{U}$.

If the spaces (X_λ, d_λ) are identical to the same space (X, d) for $\lambda \in \Lambda$, the \mathcal{U} -ultraproduct $(\prod_{\lambda \in \Lambda} X_\lambda)_\mathcal{U}$ is called the \mathcal{U} -*ultrapower* of (X, d) , denoted $(X)_\mathcal{U}$. Note that the map from X into $(X)_\mathcal{U}$ that assigns to each $x \in X$ the equivalence class of the constant family $(x \mid \lambda \in \Lambda)$ is an isometric embedding. When the ultrafilter \mathcal{U} is principal or the space (X, d) is compact

this map is surjective, though it is not so in general. The verification of these statements is left to the reader.

In the definition of ultraproduct, we lifted the metrics from the family $(X_\lambda, d_\lambda)_{\lambda \in \Lambda}$ to $(\prod_{\lambda \in \Lambda} X_\lambda)_\mathcal{U}$ by taking \mathcal{U} -limits. Doing the same for more general functions requires additional hypotheses. Let us introduce the concept of *uniform family of functions*.

6.5. Definition. Suppose that (X, d, a) and (Y, ρ, b) are pointed pseudometric spaces, B is a subset of X , and $F : X \rightarrow Y$ is uniformly continuous and bounded on B .

(1) A *bound for F on B* is a number $\Omega \geq 0$ such that

$$x \in B \quad \Rightarrow \quad F(x) \in B_Y(\Omega).$$

(2) A *modulus of uniform continuity for F on B* is a function $\Delta : (0, \infty) \rightarrow [0, \infty)$ such that, for all $x, y \in B$ and $\epsilon > 0$,

$$d(x, y) < \Delta(\epsilon) \quad \Rightarrow \quad \rho(F(x), F(y)) \leq \epsilon.$$

6.6. Definition. Let L be a signature and let \mathcal{C} be a class of L -structures. We will say that \mathcal{C} is a *uniform class* if the following two conditions hold for every function symbol $f : s_1 \times \cdots \times s_n \rightarrow s_0$ of L and every $r > 0$:

(1) (Local equiboundedness condition for \mathcal{C} .) There exists $\Omega = \Omega_{f,r} \in [0, \infty)$ such that, for every structure \mathcal{M} of \mathcal{C} , the number Ω is a bound for $f^{\mathcal{M}}$ on $B_{M_\lambda^{(s_1)}}(r) \times \cdots \times B_{M_\lambda^{(s_n)}}(r)$.

(2) (Local equicontinuity condition for \mathcal{C} .) There exists $\Delta = \Delta_{f,r} : (0, \infty) \rightarrow [0, \infty)$ such that for every structure \mathcal{M} of \mathcal{C} , the function Δ is a modulus of uniform continuity for $f^{\mathcal{M}}$ on $B_{M_\lambda^{(s_1)}}(r) \times \cdots \times B_{M_\lambda^{(s_n)}}(r)$.

Any collection $(\Omega_{r,f}, \Delta_{r,f} \mid r > 0)$ will be called a family of *moduli of local uniform continuity* for f . A collection $\mathbb{U} = (\Omega_{r,f}, \Delta_{r,f} \mid r > 0, f \in \mathbf{F})$, with f ranging over the collection \mathbf{F} of function symbols of L , will be called a *modulus of uniformity* for L -structures.

6.7. Remark. Clearly, any single L -structure \mathcal{M} admits some modulus of uniformity \mathbb{U} ; however, no single such \mathbb{U} is a modulus of uniformity for every L -structure. This is quite analogous to the fact that every Cauchy sequence is metastable with some uniform rate E_\bullet , but no single such rate of uniform metastability applies to all Cauchy sequences (refer to the discussion in Section 1-I).

Let \mathcal{C} be a uniform class of L -structures. Let $(\mathcal{M}_\lambda)_{\lambda \in \Lambda}$ be family of structures in \mathcal{C} such that \mathcal{M}_λ is based on $(M_\lambda^{(s)} \mid s \in \mathbf{S})$ for each $\lambda \in \Lambda$. If $f : s_1 \times \cdots \times s_n \rightarrow s_0$ is a function symbol of L , then for any ultrafilter \mathcal{U} on Λ we define a function

$$\left(\prod_{\lambda \in \Lambda} f^{\mathcal{M}_\lambda} \right)_\mathcal{U} : \left(\prod_{\lambda \in \Lambda} M_\lambda^{(s_1)} \right)_\mathcal{U} \times \cdots \times \left(\prod_{\lambda \in \Lambda} M_\lambda^{(s_n)} \right)_\mathcal{U} \rightarrow \left(\prod_{\lambda \in \Lambda} M_\lambda^{(s_0)} \right)_\mathcal{U}$$

naturally as follows: If $(x_\lambda^i)_{\lambda \in \Lambda} \in \ell^\infty(M_\lambda^{(s_i)}, d_\lambda^{(s_i)})_\Lambda$ for $i = 1, \dots, n$, we let

$$\left(\prod_{\lambda \in \Lambda} f^{\mathcal{M}_\lambda} \right)_\mathcal{U} \left(((x_\lambda^1)_{\lambda \in \Lambda})_\mathcal{U}, \dots, ((x_\lambda^n)_{\lambda \in \Lambda})_\mathcal{U} \right) = ((f(x_\lambda^1, \dots, x_\lambda^n))_{\lambda \in \Lambda})_\mathcal{U}. \quad (4)$$

The uniformity of \mathcal{C} implies that if $(x_\lambda^i)_{\lambda \in \Lambda} \in B_r(\ell^\infty(M_\lambda^{(s_i)}, d_\lambda^{(s_i)})_\Lambda)$, then $(f(x_\lambda^1, \dots, x_\lambda^n))_{\lambda \in \Lambda} \in B_\Omega(\ell^\infty(M_\lambda^{(s_0)}, d_\lambda^{(s_0)})_\Lambda)$ for some $\Omega > 0$, hence the right-hand side of (4) is an element of

$(\prod_{\lambda \in \Lambda} X_\lambda)_u$. Thus, if Ω is a uniform bound for $f^{\mathcal{M}_\lambda}$ on $B_{M_\lambda^{(s_1)}}(r_1) \times \cdots \times B_{M_\lambda^{(s_n)}}(r_n)$ for all $\lambda \in \Lambda$, then Ω is also a bound for $(\prod_{\lambda \in \Lambda} f^{\mathcal{M}_\lambda})_u$ on $B_{(\prod_{\lambda \in \Lambda} M_\lambda^{(s_1)})_u}(r_1) \times \cdots \times B_{(\prod_{\lambda \in \Lambda} M_\lambda^{(s_n)})_u}(r_n)$. It is also trivial to verify that, if Δ is a uniform continuity modulus for $f^{\mathcal{M}_\lambda}$ on $B_{M_\lambda^{(s_1)}}(r_1) \times \cdots \times B_{M_\lambda^{(s_n)}}(r_n)$ for all $\lambda \in \Lambda$, then Δ is also a modulus of uniform continuity for $(\prod_{\lambda \in \Lambda} f^{\mathcal{M}_\lambda})_u$ on $B_{(\prod_{\lambda \in \Lambda} M_\lambda^{(s_1)})_u}(r_1) \times \cdots \times B_{(\prod_{\lambda \in \Lambda} M_\lambda^{(s_n)})_u}(r_n)$. Thus, equation (4) defines the interpretation of functions in $(\prod_{\lambda \in \Lambda} X_\lambda)_u$ well. The following proposition summarizes the preceding discussion.

6.8. Proposition. *Let \mathcal{C} be a uniform class of L -structures and let $(\mathcal{M}_\lambda)_{\lambda \in \Lambda}$ be a family of structures in \mathcal{C} such that for each $\lambda \in \Lambda$ the structure \mathcal{M}_λ is based on $(M_\lambda^{(s)} \mid s \in \mathbf{S})$. If \mathcal{U} is an ultrafilter on Λ , we obtain an L -structure $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_u$ based on $((\prod_{\lambda \in \Lambda} M_\lambda^{(s)})_u \mid s \in \mathbf{S})$ by interpreting any function symbol f of L in $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_u$ as $(\prod_{\lambda \in \Lambda} f^{\mathcal{M}_\lambda})_u$.*

Furthermore, any modulus of uniformity for \mathcal{C} is also a modulus of uniformity for $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_u$.

6.9. Definition. The structure $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_u$ in Proposition 6.8 is called the \mathcal{U} -ultraproduct of the family $(\mathcal{M}_\lambda)_{\lambda \in \Lambda}$.

Note that the hypothesis that a class \mathcal{C} is uniform asserts that the collection $(f^{\mathcal{M}} \mid \mathcal{M} \in \mathcal{C})$ of interpretations of a given function symbol f is an equicontinuous and equibounded family, precisely as in the statement of the Arzelà-Ascoli theorem. This classical result is thus subsumed under the fact that any ultraproduct of a uniform family of L -structures is itself an L -structure. Just as the hypotheses of equicontinuity and equiboundedness are both necessary for the conclusion of the Arzelà-Ascoli theorem to hold, it should be clear that ultraproduct structures are defined, in general, only for subfamilies of some *uniform* class of structures—otherwise, the right-hand side of equation (4) may fail to define an element lying at finite distance from the anchor $((a_\lambda)_{\lambda \in \Lambda})_u$, or else the function so defined may fail to be continuous.

III. HENSON LANGUAGES AND SEMANTICS: FORMULAS AND SATISFACTION

We now focus our attention on the precise connection between metric structures and their ultrapowers and, more generally, between families of metric structures and their ultraproducts. These connections are intimately connected to notions from model theory, a branch of mathematical logic.

In the current literature, there are two formally different but equivalent logical frameworks to study metric structures from a model-theoretic perspective. One of these frameworks is that of *continuous model theory* [BYU10, BYBHU08], which uses real-valued logic, and the other is the logic of approximate truth, introduced in the 1970s by C. W. Henson [Hen75, Hen76] and developed further by Henson and the second author [HI02, Iov14]. We have adopted the latter because, despite its less widespread use in the current literature, it has strong syntactic advantages, as it allows dealing with unbounded metric structures such as Banach spaces in a natural fashion, without having to replace the metric by an equivalent bounded metric.⁸

Let L be a fixed signature. In analogy with languages of traditional (discrete) first-order logic, we construct a language, called a *Henson language*, which is suitable for discussing properties of metric structures. The language consists of syntactic expressions called *positive*

⁸For a proof of the equivalence among various formulations, see [Iov01, Iov09, CI14].

bounded formulas of L , or L -formulas. These are strings or symbols built from a basic alphabet that includes the following symbols:

- The function symbols of the signature L ,
- For each sort index $s \in \mathbf{S}$ of L , a countable collection of symbols called *variables of sort s* , or variables *bound to the sort s* .
- Logical connectives \vee and \wedge , and for each positive rational number r , quantifiers \forall_r and \exists_r .
- Parentheses and the comma symbol.

First we define the concept of L -term. Intuitively, a term is a string of symbols that may be interpreted by elements of L -structures. Since elements of structures occur inside sorts, each term must have a sort associated with it. Thus we define the concept of *s -valued term*:

6.10. Definition. An *s -valued L -term* is any finite string of symbols that can be obtained by finitely many applications of the following rules of formation:

- (1) Every variable of sort s is an s -valued term,
- (2) If f is a function symbol with $f : s_1 \times \cdots \times s_n \rightarrow s$ and t_1, \dots, t_n are such that t_i is an s_i -valued for $i = 1, \dots, n$, then $f(t_1, \dots, t_n)$ is an s -valued term.

If t is a term and x_1, \dots, x_n is a list of variables that contains all the variables occurring in t , we write t as $t(x_1, \dots, x_n)$.

A *real-valued term* is an $s_{\mathbb{R}}$ -valued term. A *term* is string that is an s -valued term for some $s \in \mathbf{S}$.

6.11. Definition. Let \mathcal{M} be an L -structure based on $(M^{(s)} \mid s \in \mathbf{S})$ and let $t(x_1, \dots, x_n)$ be an L -term, where x_i is a variable of sort s_i , for $i = 1, \dots, n$. If a_1, \dots, a_n are elements of \mathcal{M} such that a_i is of sort s_i , for $i = 1, \dots, n$, the *evaluation* of t in \mathcal{M} at a_1, \dots, a_n , denoted $t^{\mathcal{M}}[a_1, \dots, a_n]$, is defined by induction on the length of t as follows:

- (1) If t is x_i , then $t^{\mathcal{M}}[a_1, \dots, a_n]$ is a_i ,
- (2) If t is $f(t_1, \dots, t_n)$, where f is a function symbol and t, \dots, t_n are terms of lower length, then $t^{\mathcal{M}}[a_1, \dots, a_n]$ is

$$f^{\mathcal{M}}(t_1^{\mathcal{M}}[a_1, \dots, a_n], \dots, t_n^{\mathcal{M}}[a_1, \dots, a_n]).$$

As an addendum to (2), by a slight abuse of notation, if f is a nullary $M^{(s)}$ -valued function symbol (i.e., a constant symbol) we usually interpret f as be the (unique) element $a \in M^{(s)}$ in the range of the function $f^{\mathcal{M}}$.

6.12. Notation. Recall from the definition of signature that every signature L must include a special sort index $s_{\mathbb{R}}$ and constant symbol for each rational number. Informally we will identify each rational number with its constant symbol in L . More generally, since L includes function symbols for the addition and multiplication in \mathbb{R} , for every polynomial $p(x_1, \dots, x_n) \in \mathbb{Q}[x_1, \dots, x_n]$ there exists a real-valued L -term $t(x_1, \dots, x_n)$ such that $t^{\mathcal{M}}[a_1, \dots, a_n] = p(a_1, \dots, a_n)$ for any L -structure \mathcal{M} and $a_1, \dots, a_n \in \mathbb{R}$. We will identify t and p . Thus, if t_1, \dots, t_n are L terms and $p(x_1, \dots, x_n) \in \mathbb{Q}[x_1, \dots, x_n]$, we may refer to the L -term $p(t_1, \dots, t_n)$.

6.13. Definition. A *positive bounded L -formula* (or simply an *L -formula*) is any finite string of symbols that can be obtained by finitely many applications of the following rules of formation:

- (1) If t is a real-valued L -term and r is a rational number, then the expressions

$$t \leq r \quad \text{and} \quad t \geq r$$

are L -formulas. These are the *atomic* L -formulas.

- (2) If φ and ψ are positive bounded L -formulas, then the expressions

$$(\varphi \wedge \psi) \quad \text{and} \quad (\varphi \vee \psi)$$

are positive bounded L -formulas. These are the *conjunction* and *disjunction*, respectively, of φ and ψ .

- (3) If φ is positive bounded L -formula, r is a positive rational, and x is a variable, then the expressions

$$\exists_r x \varphi \quad \text{and} \quad \forall_r x \varphi$$

are positive bounded L -formulas.

6.14. Notation. Whenever possible, we shall omit parentheses according to the usual syntactic simplification rules. If t is a real-valued term and r_1, r_2 are rational numbers, we will write $r_1 \leq t \leq r_2$ as an abbreviation of the conjunction $(r_1 \leq t \wedge t \leq r_2)$. Similarly, we regard $t = r$ as an abbreviation of the conjunction $(t \leq r \wedge t \geq r)$. If t_1 and t_2 are real-valued terms, we regard $t_1 \leq t_2$ as an abbreviation of $0 \leq t_2 - t_1$ and, if t_1, t_2 are s -valued terms, we regard the expression $t_1 = t_2$ as an abbreviation of $d(t_1, t_2) \leq 0$, where d is the function symbol designating the metric of the sort indexed by s . If $\varphi_1, \dots, \varphi_n$ are formulas, we may write $\bigwedge_{i=1}^n \varphi_i$ and $\bigvee_{i=1}^n \varphi_i$ as abbreviations of $\varphi_1 \wedge \dots \wedge \varphi_n$ and $\varphi_1 \vee \dots \vee \varphi_n$, respectively. If t is an s -valued term and d, a are the designated function symbol and constant symbol, respectively, for the metric and the anchor of this sort, we shall regard the expression $t \in B_r$ as an abbreviation of the formula $d(t, a) \leq r$.

6.15. Definition. A *subformula* of a formula φ is a substring of φ that is itself a formula. If φ is a formula and x is a variable, we say that x occurs *free* in φ if there is at least one occurrence of x in φ that is not within any subformula of the form $\forall_r x \varphi$ or $\exists_r x \varphi$. If x_1, \dots, x_n are variables, we write φ as $\varphi(x_1, \dots, x_n)$ if all the free variables of φ are among x_1, \dots, x_n . A positive bounded L -sentence is a positive bounded formula without any free variables.

The definition below introduces the most basic concept of model theory, namely, the satisfaction relation \models between structures and formulas. Intuitively, if \mathcal{M} is an L structure, $\varphi(x_1, \dots, x_n)$ is an L - formula and a_1, \dots, a_n are elements of \mathcal{M} ,

$$\mathcal{M} \models \varphi[a_1, \dots, a_n]$$

means that φ is true \mathcal{M} if x_1, \dots, x_n are interpreted as a_1, \dots, a_n , respectively. Evidently, for this to be meaningful, the variable x_i must be of the same sort as the element a_i , for $i = 1, \dots, n$.

6.16. Definition. Let \mathcal{M} be an L -structure based on $(M^{(s)} \mid s \in \mathbf{S})$ and let $\varphi(x_1, \dots, x_n)$ be an L -formula, where s_i is a variable of sort s_i , for $i = 1, \dots, n$. If a_1, \dots, a_n are elements of \mathcal{M} such that a_i is of sort s_i , for $i = 1, \dots, n$, the (*discrete*) *satisfaction relation* $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ is defined inductively as follows:

- (1) If $\varphi(x_1, \dots, x_n)$ is $t \leq r$, where $t = t(x_1, \dots, x_n)$ is a real-valued term and r is rational, then $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ if and only if $t^{\mathcal{M}}[a_1, \dots, a_n] \leq r$.

- (2) If $\varphi(x_1, \dots, x_n)$ is $t \geq r$, where $t = t(x_1, \dots, x_n)$ is a real-valued term and r is rational, then $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ if and only if $t^{\mathcal{M}}[a_1, \dots, a_n] \geq r$.
- (3) If $\varphi(x_1, \dots, x_n)$ is $(\psi_1 \wedge \psi_2)$, where ψ_1 and ψ_2 are L -formulas, then $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ if and only if

$$\mathcal{M} \models \psi_1[a_1, \dots, a_n] \quad \text{and} \quad \mathcal{M} \models \psi_2[a_1, \dots, a_n].$$

- (4) If $\varphi(x_1, \dots, x_n)$ is $(\psi_1 \vee \psi_2)$, where ψ_1 and ψ_2 are L -formulas, then $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ if and only if

$$\mathcal{M} \models \psi_1[a_1, \dots, a_n] \quad \text{or} \quad \mathcal{M} \models \psi_2[a_1, \dots, a_n].$$

- (5) If $\varphi(x_1, \dots, x_n)$ is $\exists_r x \psi(x, x_1, \dots, x_n)$, where r is a positive rational, x is a variable of sort s , and $\psi(x, x_1, \dots, x_n)$ is an L -formula, then $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ if and only if

$$\mathcal{M} \models \psi[a, a_1, \dots, a_n] \quad \text{for some } a \in B_{M(s)}[r].$$

- (6) If $\varphi(x_1, \dots, x_n)$ is $\forall_r x \psi(x, x_1, \dots, x_n)$, where r is a positive rational, x is a variable of sort s , and $\psi(x, x_1, \dots, x_n)$ is an L -formula, then $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ if and only if

$$\mathcal{M} \models \psi[a, a_1, \dots, a_n] \quad \text{for every } a \in B_{M(s)}(r).$$

If $\mathcal{M} \models \varphi[a_1, \dots, a_n]$, we say that \mathcal{M} *satisfies* φ at a_1, \dots, a_n .

Note that universal (resp., existential) quantification is interpreted in open (resp., closed) balls.

6.17. Definition. If Φ is a set of formulas, we denote it by $\Phi(x_1, \dots, x_n)$ if all the free variables of all the formulas in Φ are among x_1, \dots, x_n . If \mathcal{M} is a structure and a_1, \dots, a_n are elements of \mathcal{M} , we write $\mathcal{M} \models \Phi[a_1, \dots, a_n]$ if $\mathcal{M} \models \varphi[a_1, \dots, a_n]$ for every $\varphi \in \Phi$.

IV. APPROXIMATIONS AND APPROXIMATE SATISFACTION

We begin this subsection by defining a strict partial ordering of positive bounded L -formulas, namely the relation “ ψ is an approximation of φ ”, denoted $\varphi \Rightarrow \psi$ (or $\psi \Leftarrow \varphi$). Roughly speaking, this means that ψ arises when every estimate occurring in φ is relaxed. The formal definition of the approximation relation is by induction on the complexity of φ , as given by the following table.

<u>If φ is:</u>	<u>The approximations of φ are:</u>
$t \leq r$	$t \leq r'$ where $r' > r$
$t \geq r$	$t \geq r'$ where $r' < r$
$(\xi \wedge \psi)$	$(\xi' \wedge \psi')$ where $\xi \Rightarrow \xi'$ and $\psi \Rightarrow \psi'$
$(\xi \vee \psi)$	$(\xi' \vee \psi')$ where $\xi \Rightarrow \xi'$ and $\psi \Rightarrow \psi'$
$\exists_r x \psi$	$\exists_{r'} x \psi'$ where $\psi \Rightarrow \psi'$ and $r' > r$
$\forall_r x \psi$	$\forall_{r'} x \psi'$ where $\psi \Rightarrow \psi'$ and $r' < r$

6.18. Definition. Let \mathcal{M} be an L -structure based on $(M^{(s)} \mid s \in \mathbf{S})$ and let $\varphi(x_1, \dots, x_n)$ be an L -formula, where x_i is a variable of sort s_i , for $i = 1, \dots, n$. If a_1, \dots, a_n are elements of \mathcal{M} such that a_i is of sort s_i , for $i = 1, \dots, n$, we say that \mathcal{M} *approximately satisfies* φ at a_1, \dots, a_n , and write

$$\mathcal{M} \approx \varphi[a_1, \dots, a_n],$$

if

$$\mathcal{M} \models \varphi'[a_1, \dots, a_n], \quad \text{for every } \varphi' \Leftarrow \varphi.$$

If $\Phi(x_1, \dots, x_n)$ is a set of formulas, we say that \mathcal{M} *approximately satisfies* Φ at a_1, \dots, a_n , and write $\mathcal{M} \approx \Phi[a_1, \dots, a_n]$, if $\mathcal{M} \approx \varphi[a_1, \dots, a_n]$ for every $\varphi \in \Phi$.

Clearly, approximate satisfaction is a weaker notion of truth than discrete satisfaction. For nondiscrete metric space structures, the approximate satisfaction relation \approx is the “correct” notion of truth, in the sense that for these structures it is not the notion of discrete satisfaction, but rather that of approximate satisfaction, that yields a well-behaved model theory. (For discrete structures, the two relations are clearly equivalent.)

The negation connective (“ \neg ”) is not allowed in positive bounded formulas. However, for every positive bounded formula φ there is a positive bounded formula $\mathbb{W}\varphi$, called the *weak negation* of φ , that plays a role analogous to that played by the negation of φ .

6.19. Definition. The unary pseudo-connective \mathbb{W} of *weak negation* of L -formulas is defined recursively as follows.

<u>If φ is:</u>	<u>$\mathbb{W}\varphi$ is:</u>
$t \leq r$	$t \geq r$
$t \geq r$	$t \leq r$
$(\xi \wedge \psi)$	$(\mathbb{W}\xi \vee \mathbb{W}\psi)$
$(\xi \vee \psi)$	$(\mathbb{W}\xi \wedge \mathbb{W}\psi)$
$\forall_r x \psi$	$\exists_r x \mathbb{W}\psi$
$\exists_r x \psi$	$\forall_r x \mathbb{W}\psi$.

6.20. Remarks.

- (1) If φ, φ' are positive bounded formulas, then $\varphi \Rightarrow \varphi'$ if and only if $\mathbb{W}\varphi' \Rightarrow \mathbb{W}\varphi$.
- (2) Although languages for metric structures do not include a connective interpreted as the implication “if φ then ψ ”, a formula of the form “ $\mathbb{W}\varphi \vee \psi$ ” may be regarded as a weak conditional.
- (3) If \mathcal{M} is an L -structure and $\varphi(x_1, \dots, x_n)$ is a positive bounded L -formula such that $\mathcal{M} \not\models \varphi[a_1, \dots, a_n]$, then $\mathcal{M} \models \mathbb{W}\varphi[a_1, \dots, a_n]$. If φ' is an approximation of φ such that $\mathcal{M} \models \mathbb{W}\varphi'[a_1, \dots, a_n]$, then $\mathcal{M} \not\models \varphi[a_1, \dots, a_n]$.

6.21. Proposition. *Let \mathcal{M} be an L -structure based on $(M^{(s)} \mid s \in \mathbf{S})$, let $\varphi(x_1, \dots, x_n)$ be an L -formula, where x_i is a variable of sort s_i , for $i = 1, \dots, n$, and let a_1, \dots, a_n be elements of \mathcal{M} such that a_i is of sort s_i , for $i = 1, \dots, n$. Then, $\mathcal{M} \not\approx \varphi[a_1, \dots, a_n]$ if and only there exists a formula $\varphi' \Leftarrow \varphi$ such that $\mathcal{M} \approx \mathbb{W}\varphi'[a_1, \dots, a_n]$.*

Proof. In order to simplify the nomenclature, let us suppress the lists x_1, \dots, x_n and a_1, \dots, a_n from the notation.

If $\mathcal{M} \not\approx \varphi$, there exists $\varphi' \Leftarrow \varphi$ such that $\mathcal{M} \not\models \varphi'$. We have $\mathcal{M} \models \mathbb{W}\varphi'$ and hence $\mathcal{M} \approx \mathbb{W}\varphi'$. Conversely, assume that there exists $\varphi' \Leftarrow \varphi$ such that $\mathcal{M} \approx \mathbb{W}\varphi'$. Take sentences ψ, ψ' such that $\varphi \Rightarrow \psi \Rightarrow \psi' \Rightarrow \varphi'$. Then $\mathcal{M} \models \mathbb{W}\psi'$ (by Remark 6.20-(3)) and hence $\mathcal{M} \not\models \psi$, so $\mathcal{M} \not\approx \varphi$. \square

V. THEORIES, ELEMENTARY EQUIVALENCE, AND ELEMENTARY SUBSTRUCTURES

In this section we include some basic definitions from model theory.

6.22. Definition. Let L be a signature.

- (1) An L -theory (or simply a *theory*) is a set of L -sentences.
- (2) If T is a theory and $\mathcal{M} \approx T$, we say that \mathcal{M} is a *model* of T . A theory is *satisfiable* (or *consistent*) if it has a model. The class of all models of a theory T is denoted $\text{Mod}(T)$. If \mathcal{C} is any class of L -structures, the class $\mathcal{C} \cap \text{Mod}(T)$ of models of T in \mathcal{C} is denoted $\text{Mod}_{\mathcal{C}}(T)$.
- (3) An *axiomatizable class* \mathcal{C} (or *elementary class*) is one that consists of all the models of a fixed theory T . We say that \mathcal{C} is *L -axiomatizable*, or \mathcal{C} is *axiomatizable by T* , when the language or the theory need to be specified.
- (4) An L -theory T is *uniform* if the class of all models of T is uniform. (See Proposition 6.24.)
- (5) The *complete L -theory of a structure* \mathcal{M} , denoted $\text{Th}(\mathcal{M})$, is the set of all L -sentences φ such that $\mathcal{M} \approx \varphi$. A *complete L -theory* is the complete L -theory of any L -structure \mathcal{M} .
- (6) The *complete L -theory of class \mathcal{C} of L -structures* is $\text{Th}(\mathcal{C}) = \bigcap_{\mathcal{M} \in \mathcal{C}} \text{Th}(\mathcal{M})$.
- (7) Two L -structures \mathcal{M}, \mathcal{N} are *elementarily equivalent*, written $\mathcal{M} \equiv \mathcal{N}$, if they have the same complete theory, i.e., if

$$\mathcal{M} \approx \varphi \iff \mathcal{N} \approx \varphi, \quad \text{for every } L\text{-sentence } \varphi.$$

- (8) If \mathcal{M} and \mathcal{N} are L -structures and \mathcal{M} is a substructure of \mathcal{N} , we say that \mathcal{M} is an *elementary substructure* of \mathcal{N} , and we write $\mathcal{M} \prec \mathcal{N}$, if whenever a_1, \dots, a_n are elements of \mathcal{M} and $\varphi(x_1, \dots, x_n)$ is an L -formula such that a_i is of the same sort as x_i , for $i = 1, \dots, n$, we have

$$\mathcal{M} \approx \varphi[a_1, \dots, a_n] \iff \mathcal{N} \approx \varphi[a_1, \dots, a_n].$$

- 6.23. Remarks.**
- (1) When \mathcal{M} is an L -structure, the interpretation of each function symbol is locally bounded and locally uniformly continuous by definition, hence $\text{Th}(\mathcal{M})$ is necessarily a uniform theory.
 - (2) Any satisfiable theory T admits some extension to a uniform theory $T' = \text{Th}(\mathcal{M})$, where \mathcal{M} is any model of T . Neither the extension nor a modulus of uniformity thereof is uniquely determined by T in general.
 - (3) If T' extends a \mathbb{U} -uniform theory T , then T' is also \mathbb{U} -uniform.

6.24. Proposition. *Let L be a signature. Given a modulus of uniformity \mathbb{U} , the class of L -structures \mathcal{M} such that \mathbb{U} is a modulus of uniformity for \mathcal{M} is axiomatizable.*

Proof. Let $\mathbb{U} = (\Omega_{r,f}, \Delta_{r,f} \mid r > 0, f \in \mathbf{F})$ be a modulus of uniformity. For a function symbol $f : s_1 \times \dots \times s_n \rightarrow s_0$ and rational numbers $u, v, w > 0$, consider the L -sentences

$$\begin{aligned} \chi_{f,u,v} &: \forall_u x_1 \dots \forall_u x_n (d(f(x_1, \dots, x_n), a_0) \leq v), \\ \xi_{f,u,v,w} &: \forall_u x_1 \forall_u y_1 \dots \forall_u x_n \forall_u y_n (\bar{d}(\bar{x}, \bar{y}) \geq v \vee d(f(\bar{x}), f(\bar{y})) \leq w), \end{aligned}$$

where \bar{x}, \bar{y} denote the n -tuples x_1, \dots, x_n and y_1, \dots, y_n , respectively, \bar{d} denotes the supremum distance on $\mathcal{M}^{(s_1)} \times \dots \times \mathcal{M}^{(s_n)}$, and a_0 the constant symbol for the anchor of the sort indexed

by s_0 . It should be clear that for any signature L , the class of L -structures \mathcal{M} such that \mathbb{U} is a modulus of uniformity for \mathcal{M} is axiomatized by the union of the following sets of sentences:

$$\begin{aligned} & \{ \chi_{f,u,v} \mid f \in \mathbf{F}, u, v \in \mathbb{Q}_+, u < r, v > \Omega_{r,f} \text{ for some } r \in (0, \infty) \}, \\ & \{ \xi_{f,u,v,w} \mid f \in \mathbf{F}, u, v \in \mathbb{Q}_+, u < r, v < t, w > \Delta_{r,f}(t) \text{ for some } r, t \in (0, \infty) \}. \quad \square \end{aligned}$$

It is easy to construct examples showing that the meta-property “*the class \mathcal{C} is uniform*” (without specifying a modulus of uniformity) is not axiomatizable in general.

The following proposition gives simple conditions to verify \equiv and \prec .

6.25. Proposition. *Let \mathcal{M} and \mathcal{N} be L -structures.*

(1) $\mathcal{M} \equiv \mathcal{N}$ if and only if for every L -sentence φ ,

$$\mathcal{M} \vDash \varphi \quad \Rightarrow \quad \mathcal{N} \vDash \varphi.$$

(2) (Tarski-Vaught test for \prec). *A substructure \mathcal{M} of the structure \mathcal{N} is an elementary substructure if and only if the following condition holds: If $\varphi(x_1, \dots, x_n, y)$ is an L -formula, a_1, \dots, a_n are elements of \mathcal{M} with a_i of the same sort as x_i for $i = 1, \dots, n$ such that $\mathcal{N} \vDash \exists_r y \varphi[a_1, \dots, a_n]$ for some $r > 0$, then $\mathcal{M} \vDash \exists_r y \varphi[a_1, \dots, a_n]$.*

Proof. For part (1), the direct implication follows by definition of elementary equivalence, while the converse follows from Proposition 6.21. For (2), the direct implication is trivial, and the converse follows by induction on the complexity of φ . \square

The following is an immediate consequence of Proposition 6.25-(2):

6.26. Proposition (Downward Löwenheim-Skolem Theorem). *Let \mathcal{M} be a structure and let A be a set of elements of \mathcal{M} . Then there exists a substructure \mathcal{M}_0 of \mathcal{M} such that*

- $\mathcal{M}_0 \prec \mathcal{M}$,
- Every element of A is an element of \mathcal{M}_0 ,
- $\text{card}(\mathcal{M}_0) \leq \text{card}(A) + \text{card}(L)$.

6.27. Definition. An *elementary chain* is a family $(\mathcal{M}_i)_{i \in I}$ of structures, indexed by some linearly ordered set $(I, <)$, such that $\mathcal{M}_i \prec \mathcal{M}_j$ if $i < j$.

Another useful consequence of Proposition 6.25 is the elementary chain property:

6.28. Proposition (The Elementary Chain Property). *If $(\mathcal{M}_i)_{i \in I}$ is an elementary chain, then $\bigcup_{i \in I} \mathcal{M}_i$ is an elementary extension of \mathcal{M}_j for every $j \in I$.*

Proof. By Proposition 6.25. \square

VI. ŁOŚ’ THEOREM

The following fundamental theorem, proved by J. Łoś in the 1950s [Łoś55], intuitively states that a formula φ is satisfied by an ultraproduct of a family $(\mathcal{M}_\lambda)_{\lambda \in \Lambda}$ of structures if and only if every approximation of φ is satisfied by almost all of the structures \mathcal{M}_λ . Łoś proved the theorem for discrete structures (i.e., traditional first-order logic), where approximations are not needed; however, essentially the same argument holds for arbitrary metric structures.

6.29. Theorem (Łoś’ Theorem for Metric Structures). *Let L be a signature and let $(\mathcal{M}_\lambda)_{\lambda \in \Lambda}$ be a family of L -structures in a uniform class such that for each $\lambda \in \Lambda$ the structure \mathcal{M}_λ is based on $(M_\lambda^{(s)} \mid s \in \mathbf{S})$. Let $(a_{1,\lambda})_{\lambda \in \Lambda}, \dots, (a_{n,\lambda})_{\lambda \in \Lambda}$ be such that $(a_{i,\lambda})_{\lambda \in \Lambda} \in \ell^\infty(M_\lambda^{(s_i)})_\Lambda$*

for $i = 1, \dots, n$ and let $\varphi(x_1, \dots, x_n)$ be an L -formula such that x_i is of sort s_i . Then, for any ultrafilter \mathcal{U} on Λ ,

$$\left(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda \right)_{\mathcal{U}} \approx \varphi[((a_{1,\lambda})_{\lambda \in \Lambda})_{\mathcal{U}}, \dots, ((a_{n,\lambda})_{\lambda \in \Lambda})_{\mathcal{U}}]$$

if and only if for every approximation $\varphi' \approx \varphi$, the set

$$\{ \lambda \in \Lambda \mid \mathcal{M}_\lambda \models \varphi'[a_{1,\lambda}, \dots, a_{n,\lambda}] \}$$

is \mathcal{U} -large.

Proof. Using the definition of \mathcal{U} -ultraproduct of structures and the interpretation of function symbols therein (Proposition 6.8 and Definition 6.9), Łoś' Theorem follows by induction on the complexity of φ . \square

An important corollary of Łoś' theorem is the special case when all the structures \mathcal{M}_λ equal the same structure \mathcal{M} . In this case, the \mathcal{U} -ultraproduct $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_{\mathcal{U}}$ is the \mathcal{U} -ultrapower of \mathcal{M} . Hence we have the following:

6.30. Corollary. *Every metric structure is an elementary substructure of its ultrapowers.*

VII. COMPACTNESS

The Compactness Theorem is arguably the most distinctive theorem of first-order logic.

For a set φ of formulas, let φ_{\approx} denote the set of all approximations of formulas in φ .

6.31. Theorem (Compactness Theorem). *Let \mathcal{C} be a uniform class of structures and let T be an L -theory. If every finite subset of T_{\approx} has a model in \mathcal{C} in the semantics of discrete satisfaction, then T has a model in the semantics of approximate satisfaction. This model may be taken to be an ultraproduct of structures in \mathcal{C} that admits the same modulus of uniformity as the class \mathcal{C} .*

Proof. Let Λ be the set of finite subsets of T_{\approx} , and for each λ in Λ , let \mathcal{M}_λ be a model in \mathcal{C} of all the sentences in λ in the semantics of discrete satisfaction. For every finite subset φ of T_{\approx} , let $\Lambda_{\supseteq \varphi}$ be the set of all $\lambda \in \Lambda$ such that $\lambda \supseteq \varphi$. Then $\mathcal{M}_\lambda \models \varphi$ for every $\lambda \in \Lambda_{\supseteq \varphi}$. Note that the collection of subsets of Λ of the form $\Lambda_{\supseteq \varphi}$ is closed under finite intersections since $\Lambda_{\supseteq \varphi} \cap \Lambda_{\supseteq \psi} = \Lambda_{\supseteq \varphi \cup \psi}$. Let \mathcal{U} be an ultrafilter on Λ extending this collection. Then, by Łoś' Theorem (Theorem 6.29), we have $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_{\mathcal{U}} \approx T$. Furthermore, $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_{\mathcal{U}}$ admits the same modulus of uniformity as the family $(\mathcal{M}_\lambda)_{\lambda \in \Lambda}$, by Proposition 6.8. \square

The following corollary amounts to a restatement of the Compactness Theorem that does not explicitly mention approximation of formulas.

6.32. Corollary. *If Υ is a uniform theory and T is any collection of sentences such that every finite subset of T is approximately satisfied by a model of Υ , then $\Upsilon \cup T$ has a model.*

Proof. Every finite subset φ of $(\Upsilon \cup T)_{\approx}$ must be a subset of $(\Upsilon \cup T')_{\approx}$ for some finite subset $T' \subseteq T$. By hypothesis, $\Upsilon \cup T'$ has a model, which is thus also a model of φ in the semantics of discrete satisfaction. By the Compactness Theorem (Theorem 6.31), $\Upsilon \cup T$ has a model. \square

6.33. Remarks. Fix a signature L and a class \mathcal{C} of L -structures. For any set Φ of L -formulas, let $\text{Mod}_{\mathcal{C}}(\Phi) = \mathcal{C} \cap \text{Mod}(\Phi)$ be the subclass of \mathcal{C} consisting of models of Φ . We topologize \mathcal{C} by taking as a basis of closed sets the collection of all classes of the form $\text{Mod}_{\mathcal{C}}(\varphi)$ as φ

varies over L -formulas; thus, closed sets are of the form $\text{Mod}_{\mathcal{C}}(\Phi)$. We refer to this as the *logic topology on structures*.

Corollary 6.32 is equivalent to the following statement: If $\mathcal{C} = \text{Mod}(\Upsilon)$ is the class of models of a uniform L -theory Υ , and T is any collection of L -sentences such that the family

$$(\text{Mod}_{\mathcal{C}}(\Phi) \mid \Phi \text{ is a finite subset of } T)$$

has the finite intersection property, then

$$\bigcap_{\substack{\Phi \subset T \\ \Phi \text{ finite}}} \text{Mod}_{\mathcal{C}}(\Phi) \neq \emptyset.$$

Thus, an axiomatizable uniform class \mathcal{C} is compact in the logic topology. Moreover, if $(\mathcal{M}_\lambda)_{\lambda \in \Lambda}$ is any family in \mathcal{C} and \mathcal{U} is an ultrafilter on Λ , then a \mathcal{U} -limit of $(\mathcal{M}_\lambda)_{\lambda \in \Lambda}$ in this topology is given by the ultraproduct $(\prod_{\lambda \in \Lambda} \mathcal{M}_\lambda)_{\mathcal{U}}$. This explains why the Compactness Theorem is so named, as well as the naturality of the ultraproduct construction. More generally, if \mathcal{C} is any uniform class endowed with the above topology, then \mathcal{C} is a relatively compact (dense) subset of the compact class $\overline{\mathcal{C}} = \text{Mod}(\text{Th}(\mathcal{C}))$. We emphasize that uniformity is a necessary condition for precompactness.

To conclude these remarks, we note that the logic topology on any uniform class \mathcal{C} of L -structures is regular (although usually not Hausdorff); this follows from the fact that \mathcal{C} is a subspace of the compact (hence regular) space $\overline{\mathcal{C}} = \text{Mod}(\text{Th}(\mathcal{C}))$.

Now we present three useful applications of the Compactness Theorem. The first one (Corollary 6.34) gives a finitary condition for a theory to be of the form $\text{Th}(\mathcal{M})$ for some structure \mathcal{M} . The second one (Corollary 6.37), states that any two models of a complete theory T can be jointly elementarily embedded in a single model of T .

Recall (Definition 6.22) that the complete theory of a structure \mathcal{M} is denoted $\text{Th}(\mathcal{M})$.

6.34. Corollary. *The following conditions are equivalent for a theory T :*

- (1) *The theory T is complete, i.e., there exists a structure \mathcal{M} such that $T = \text{Th}(\mathcal{M})$.*
- (2) (a) *There exists a uniform theory Υ such that every finite subset of $\Upsilon \cup T$ is satisfiable, and*
 (b) *For every L -sentence φ , if $\varphi \notin T$, then there exists $\varphi' \approx \varphi$ such that $\mathfrak{M} \models \varphi' \in T$.*

Proof. We only have to prove (2) \Rightarrow (1). By part (a) of (2) and Corollary 6.32, T has a model \mathcal{M} whose theory $\text{Th}(\mathcal{M})$ extends T . Part (b) of (2) gives $\text{Th}(\mathcal{M}) \subseteq T$, proving (1). \square

6.35. Notation. If \mathcal{M} is an L -structure and $A = (a_i)_{i \in I}$ is an indexed family of elements of \mathcal{M} , we denote by $(\mathcal{M}, a_i \mid i \in I)$ the expansion of \mathcal{M} that results from adding a distinct constant for a_i , for each $i \in I$. In particular, if A is a set of elements of \mathcal{M} , we denote by

$$(\mathcal{M}, a \mid a \in A)$$

the expansion of \mathcal{M} that results from adding a constant for each $a \in A$. If, in this case, C is a set of constant symbols not already in L that includes one constant symbol designating each constant $a \in A$, we informally refer to the signature $L[C]$ (see Definition 6.3) as $L[A]$. Thus, we informally call the preceding expansion of \mathcal{M} an “ $L[A]$ -structure”. If A consists of all the elements of \mathcal{M} , we write $(\mathcal{M}, a \mid a \in A)$ as $(\mathcal{M}, a \mid a \in \mathcal{M})$ and denote the expanded signature by $L[\mathcal{M}]$.

6.36. Definition. Let \mathcal{M} and \mathcal{N} be L -structures. An *elementary embedding* of \mathcal{M} into \mathcal{N} is a map e that assigns to each element a of \mathcal{M} an element $e(a)$ of \mathcal{N} such that, whenever a_1, \dots, a_n are elements of \mathcal{M} and $\varphi(x_1, \dots, x_n)$ is an L -formula such that a_i is of the same sort as x_i , for $i = 1, \dots, n$, we have

$$\mathcal{M} \models \varphi[a_1, \dots, a_n] \iff \mathcal{N} \models \varphi[e(a_1), \dots, e(a_n)].$$

Note that e is an elementary embedding of \mathcal{M} into \mathcal{N} if and only if the $L[A]$ -structures $(\mathcal{M}, a \mid a \in A)$ and $(\mathcal{N}, e(a) \mid a \in A)$ are elementarily equivalent.

6.37. Corollary. *The following conditions are equivalent for two L -structures $\mathcal{M}, \mathcal{M}'$.*

- (1) $\mathcal{M} \equiv \mathcal{M}'$.
- (2) *There exists a structure \mathcal{N} such that $\mathcal{M} \prec \mathcal{N}$ and there is an elementary embedding of \mathcal{M}' into \mathcal{N} . Moreover, \mathcal{N} can be taken to be an ultrapower of \mathcal{M} .*

Proof. It suffices to prove the direct implication, since the inverse is clear. Assume $\mathcal{M} \equiv \mathcal{M}'$. Let T be the complete $L[\mathcal{M}]$ -theory of \mathcal{M} and T' the complete $L[\mathcal{M}']$ -theory of \mathcal{M}' . Both T and T' are uniform theories, by Remark 6.23. First we show that every finite subset of $T' \cup T$ has a model. Since T' is closed under conjunctions, it suffices to show that $T \cup \{\psi\}$ has a model whenever $\psi \in T'$. Any given formula in T' is of the form $\varphi'(c_{b_1}, \dots, c_{b_n})$, where φ' is an approximation of a formula φ in T and c_{b_i} is a constant for an element b_i of \mathcal{M}' , for $i = 1, \dots, n$. Since $\mathcal{M}' \models \varphi[b_1, \dots, b_n]$ by assumption, then there exists $r > 0$ such that $\mathcal{M}' \models \exists_r x_1 \dots \exists_r x_n \varphi(x_1, \dots, x_n)$. Now, since $\mathcal{M} \equiv \mathcal{M}'$ and $\varphi' \approx \varphi$, the semantics of approximate satisfaction ensure the existence of elements a_1, \dots, a_n of \mathcal{M} such that $\mathcal{M} \models \varphi'[a_1, \dots, a_n]$; hence, $(\mathcal{M}, a \mid a \in \mathcal{M})$ admits an expansion to a model $\tilde{\mathcal{M}}$ of $T \cup \{\psi\}$ simply by letting $b_1^{\tilde{\mathcal{M}}} := a_1, \dots, b_n^{\tilde{\mathcal{M}}} := a_n$. By the Compactness Theorem 6.31, $T \cup T'$ has a model. Let

$$(\mathcal{N}, \tilde{a}, \tilde{b} \mid a \in \mathcal{M}, b \in \mathcal{M}')$$

be a model of this theory. The maps $a \mapsto \tilde{a}$ and $b \mapsto \tilde{b}$ are elementary embeddings of \mathcal{M} and \mathcal{M}' , respectively, into \mathcal{N} . Without loss of generality, we may assume both that \mathcal{N} is an ultrapower of \mathcal{M} , and $\tilde{a} = a$ for all elements a of \mathcal{M} . This proves (2). \square

Essentially the same argument used to prove Corollary 6.37 proves the following result:

6.38. Corollary. *Given any family $(\mathcal{M}_i)_{i \in I}$ of elementarily equivalent L -structures there exists an L -structure \mathcal{N} such that, for each $i \in I$, there is an elementary embedding of \mathcal{M}_i into \mathcal{N} .*

6.39. Proposition. *If \mathcal{M}, \mathcal{N} are elementarily equivalent structures, then there exist structures $\hat{\mathcal{M}}, \hat{\mathcal{N}}$ such that $\mathcal{M} \prec \hat{\mathcal{M}}, \mathcal{N} \prec \hat{\mathcal{N}}$, and $\hat{\mathcal{M}}$ is isomorphic to $\hat{\mathcal{N}}$.*

Furthermore, if $(a_i)_{i \in I}, (b_i)_{i \in I}$ are elements of \mathcal{M} and \mathcal{N} such that

$$(\mathcal{M}, a_i \mid i \in I) \equiv (\mathcal{N}, b_i \mid i \in I),$$

then there exist elementary extensions $\hat{\mathcal{M}} \succ \mathcal{M}$ and $\hat{\mathcal{N}} \succ \mathcal{N}$ and an isomorphism \mathcal{J} from $\hat{\mathcal{M}}$ into $\hat{\mathcal{N}}$ such that $\mathcal{J}(a_i) = b_i$ for all $i \in I$.

Proof. Inductively, for every ordinal $n < \omega$, use Corollary 6.37 to construct structures $\mathcal{M}_n, \mathcal{N}_n$, and maps e_n, f_n , such that

- (i) $\mathcal{M}_0 = \mathcal{M}, \mathcal{N}_0 = \mathcal{N}$,

- (ii) e_n is an elementary embedding of \mathcal{M}_n into \mathcal{N}_{n+1} and f_n is an elementary embedding of \mathcal{N}_{n+1} into \mathcal{M}_{n+1} ,
- (iii) $f_{n+1}(e_n(a)) = a$ for every element a of the universe of \mathcal{M}_n .

Let $\hat{\mathcal{M}} = \bigcup_{n < \omega} \mathcal{M}_n$, $\hat{\mathcal{N}} = \bigcup_{n < \omega} \mathcal{N}_n$, and $e = \bigcup_{n < \omega} e_n$. The e is an isomorphism from $\hat{\mathcal{M}}$ into $\hat{\mathcal{N}}$.

The second part of the statement is given by the preceding construction, since e_0 can be chosen to map a_i to b_i for each $i \in I$. Alternatively, let $C = (c_i \mid i \in I)$ be a collection of new constants and apply the result just proved to the $L[C]$ -structures $\tilde{\mathcal{M}} = (\mathcal{M}, a_i \mid i \in I)$ and $\tilde{\mathcal{N}} = (\mathcal{N}, b_i \mid i \in I)$. \square

VIII. TYPES

We begin this subsection defining the notion of *finite satisfiability* of a set of formulas.

6.40. Definition. If $\varphi(x_1, \dots, x_n)$ is a set of L -formulas and \mathcal{M} is an L -structure, we say that φ is *finitely satisfiable* in \mathcal{M} if there exists r such that for every finite $\varphi_0 \subseteq \varphi$,

$$\mathcal{M} \approx \exists_r x_1 \dots \exists_r x_n \bigwedge_{\varphi \in \varphi_0} \varphi(x_1, \dots, x_n). \quad (*)$$

Note that if $(*)$ holds, then for every finite $\varphi_0 \subseteq \varphi_{\approx}$ there exist c_1, \dots, c_n in \mathcal{M} such that $\mathcal{M} \models \bigwedge_{\varphi \in \varphi_0} \varphi[c_1, \dots, c_n]$ (discrete satisfaction). However, the tuple (c_1, \dots, c_n) depends on φ_0 .

We now introduce one of the central concepts of model theory: that of *type*.

6.41. Definition. If \mathcal{M} is an L -structure, A is a set of elements of \mathcal{M} , and (c_1, \dots, c_n) is a tuple of elements of \mathcal{M} , the *type of (c_1, \dots, c_n) over A* , denoted $\text{tp}_A(c_1, \dots, c_n)$, is the set of all $L[A]$ -formulas $\varphi(x_1, \dots, x_n)$ such that $(\mathcal{M}, a \mid a \in A) \approx \varphi[c_1, \dots, c_n]$. We denote $\text{tp}_{\emptyset}(c_1, \dots, c_n)$ by $\text{tp}_L(c_1, \dots, c_n)$.

If T is a complete L -theory and $t = t(x_1, \dots, x_n)$ is a set of $L[C]$ -formulas, where C is a set of constant symbols not in L , we say that t is an *n -type of T* if there exists an $L[C]$ -model \mathcal{M} of T and elements c_1, \dots, c_n in \mathcal{M} such that $t = \text{tp}_A(c_1, \dots, c_n)$ where A is the subset of \mathcal{M} consisting of elements interpreting the constants in C . In this case, we say t is *realized* in \mathcal{M} , and that the n -tuple (c_1, \dots, c_n) *realizes t* in \mathcal{M} .

6.42. Remarks.

- (1) If T is a complete theory and t is a type for T , then $T \subseteq t$. In fact, if \mathcal{M} is a model of T and t is a type over a set A of elements of \mathcal{M} , then the set of sentences in t is precisely the complete $L[A]$ -theory $\text{Th}(\mathcal{M}, a \mid a \in A)$, which extends $T = \text{Th}(\mathcal{M})$.
- (2) The notation $\text{tp}_A(c_1, \dots, c_n)$ is imprecise in the sense that it does not make reference to the structure \mathcal{M} where the elements c_1, \dots, c_n and the set A “live”. However, since T is a complete L -theory by assumption, precise knowledge of \mathcal{M} is to a large extent unnecessary. In fact, if we are given a family $(t_i)_{i \in I}$ of types of T such that, for $i \in I$, the type t_i is over a set of parameters A_i realized in an L -structure \mathcal{M}_i , then by Corollary 6.38 there exists a single structure \mathcal{N} in which all the structures \mathcal{M}_i are elementary embedded. In fact, as we shall see in Proposition 6.46 below, given any cardinal κ we can fix a “big” model \mathcal{N} of T that is an elementary extension of every L -structure \mathcal{M} with cardinality $\text{card}(\mathcal{M}) < \kappa$, and such that \mathcal{N} realizes every type over any of its subsets A with $\text{card}(A) < \kappa$. Furthermore, the model

\mathcal{N} can be taken with the following additional homogeneity property: If c_1, \dots, c_n and c'_1, \dots, c'_n are elements of \mathcal{N} and A is a set of elements of \mathcal{N} with $\text{card}(A) < \kappa$, then $\text{tp}_A(c_1, \dots, c_n) = \text{tp}_A(c'_1, \dots, c'_n)$ if and only if there is an automorphism of \mathcal{N} carrying c_1, \dots, c_n to c'_1, \dots, c'_n and fixing A pointwise. This allows viewing types as orbits on the big model under the action of its group of automorphisms, and enables a Galois-theoretic viewpoint of complete theories. It also explains the use of the word “type”. We will discuss this in more detail in part XI.

- (3) If T is a complete theory, \mathcal{M} is an arbitrary model of T , A is a set of elements of \mathcal{M} , and t is a type of T over A , there is no guarantee that t is realized in \mathcal{M} . However:
- (a) The equivalence between (1) and (2) of Proposition 6.43 below shows that \mathcal{M} has an elementary extension where t is realized. In particular, every elementary extension of \mathcal{M} has a further elementary extension where t is realized.
 - (b) t is finitely satisfiable in every model of $\text{Th}(\mathcal{M}, a \mid a \in A)$.

6.43. Proposition. *Let T be a complete theory and let $t(x_1, \dots, x_n)$ be a set of $L[C]$ -formulas, where C is a set of constant symbols not in L . The following conditions are equivalent.*

- (1) *There exists a model \mathcal{M} of T , a set A of elements of \mathcal{M} , and elements c_1, \dots, c_n of \mathcal{M} such that $t(x_1, \dots, x_n) = \text{tp}_A(c_1, \dots, c_n)$.*
- (2) *For every model \mathcal{N} of T there exists an elementary extension \mathcal{N}' of \mathcal{N} , a set A of elements of \mathcal{N}' , and elements c_1, \dots, c_n of \mathcal{N}' such that $t(x_1, \dots, x_n) = \text{tp}_A(c_1, \dots, c_n)$.*
- (3) (a) *There exists a positive rational r such that, for every finite subset $\varphi \subset t$, the formula $\exists_r x_1 \dots \exists_r x_n \bigwedge_{\varphi \in \varphi} \varphi(x_1, \dots, x_n)$ is in t .*
 (b) *For every $L[C]$ -formula $\varphi(x_1, \dots, x_n)$, if $\varphi \notin t$, then there exists $\varphi' \not\equiv \varphi$ such that $\varphi' \in t$.*

Proof. The equivalence between (1) and (3) is given by Corollary 6.34 (by replacing the variables x_1, \dots, x_n with constant symbols not already in $L[A]$). The equivalence between (3) and (2) is given by Corollary 6.37. \square

6.44. Remark. Part (3) of Proposition 6.43 gives a purely syntactic criterion for a given set of formulas to be a type.

IX. SATURATED AND HOMOGENEOUS MODELS

Strictly for reasons of notational simplicity, we will focus our attention on one-sorted structures throughout this subsection. When a structure \mathcal{M} has only one sort, this sort is called the *universe* of \mathcal{M} . The *cardinality* of a one-sorted structure is defined as the cardinality of its universe.

The observations in Remark 6.43 lead naturally to the concept of *saturated structure*:

6.45. Definition. Let \mathcal{M} is a structure with universe M and let κ be an infinite cardinal with $\kappa \leq \text{card}(M)$. We say that \mathcal{M} is κ -*saturated* if whenever t is a type of $\text{Th}(\mathcal{M})$ over a subset of M of cardinality strictly less than κ , there is a realization of t in \mathcal{M} .

In this context, we may abuse notation and use ω as a synonym for \aleph_0 , the cardinality of (infinite) countable sets. Thus, an ω^+ -saturated structure is one that realizes types over any countable subset of the universe. The informal terminology “countable saturation” shall mean ω^+ -saturation, and not ω -saturation.

Notice that an L -structure \mathcal{M} is κ -saturated if and only if, whenever A is a subset of the universe of \mathcal{M} of cardinality less than κ and $\varphi(x_1, \dots, x_n)$ is a set of $L[A]$ -formulas that is finitely satisfiable in \mathcal{M} , there exist c_1, \dots, c_n in \mathcal{M} such that $\varphi(x_1, \dots, x_n) \subseteq \text{tp}_A(c_1, \dots, c_n)$.

6.46. Proposition. *If \mathcal{M} is κ -saturated and \mathcal{N} is a structure of cardinality less than κ such that $\mathcal{N} \equiv \mathcal{M}$, then \mathcal{N} can be elementarily embedded in \mathcal{M} .*

Proof. Assume that \mathcal{M} is κ -saturated, $\mathcal{N} \equiv \mathcal{M}$, and $\text{card}(\mathcal{N}) < \kappa$. It follows easily by induction that, if α is an ordinal satisfying $\alpha < \kappa$ and $(a_i)_{i < \alpha}$ is a list of elements of \mathcal{N} , then there exists $(a'_i)_{i < \alpha}$ in \mathcal{M} such that $(\mathcal{N}, a_i \mid i < \alpha) \equiv (\mathcal{M}, a'_i \mid i < \alpha)$. Thus, in the case when $(a_i)_{i < \alpha}$ lists all the elements of \mathcal{N} , the map $a_i \mapsto a'_i$ is an elementary embedding of \mathcal{N} into \mathcal{M} . \square

6.47. Proposition. *If \mathcal{M} is an \aleph_1 -saturated L -structure, then \approx and \models coincide on \mathcal{M} , i.e., for an L -formula $\varphi(x_1, \dots, x_n)$ and elements a_1, \dots, a_n of suitable sorts, we have*

$$\mathcal{M} \approx \varphi[a_1, \dots, a_n] \quad \Leftrightarrow \quad \mathcal{M} \models \varphi[a_1, \dots, a_n].$$

Proof. By induction on the complexity of φ . \square

Recall that, if κ is a cardinal, then κ^+ denotes the smallest cardinal larger than κ .

6.48. Proposition. *If κ is an infinite cardinal, then every structure has a κ^+ -saturated elementary extension.*

Proof. Fix an infinite cardinal κ and an L -structure \mathcal{M} . Applying Remark 6.42-(2), we construct inductively, for every ordinal $i < \kappa^+$, a structure \mathcal{M}_i , such that

- (i) $\mathcal{M}_0 = \mathcal{M}$,
- (ii) $\mathcal{M}_i \prec \mathcal{M}_{i+1}$ and every type over a subset of the universe of \mathcal{M}_i of cardinality less than κ is realized in \mathcal{M}_{i+1} ,
- (iii) If $j < \kappa^+$ is a limit ordinal, then $\mathcal{M}_j = \bigcup_{i < j} \mathcal{M}_i$.

It follows from the Elementary Chain Property (Proposition 6.28) that $(\mathcal{M}_i)_{i < \kappa^+}$ is an elementary chain, and that $\bigcup_{i < \kappa^+} \mathcal{M}_i$ is an elementary extension of \mathcal{M} , which is clearly κ^+ -saturated. \square

Suppose that \mathcal{M} is κ -saturated and let α be an ordinal with $\alpha < \kappa$. It follows directly from the κ -saturation of \mathcal{M} that if $(a_i)_{i < \alpha}$, $(a'_i)_{i < \alpha}$ are families in \mathcal{M} such that $(\mathcal{M}, a_i \mid i < \alpha) \equiv (\mathcal{M}, a'_i \mid i < \alpha)$, then for every element b of \mathcal{M} there exists an element b' such that

$$(\mathcal{M}, b, a_i \mid i < \alpha) \equiv (\mathcal{M}, b', a'_i \mid i < \alpha).$$

A structure \mathcal{M} that has this extension property for all pairs of families $(a_i)_{i < \alpha}$, $(a'_i)_{i < \alpha}$ with $\alpha < \kappa$ is said to be κ -homogeneous. Note that if \mathcal{M} is κ -homogeneous with $\kappa = \text{card}(\mathcal{M})$, then the κ -homogeneity can be used iteratively to extend the map $a_i \mapsto a_i$ ($i < \alpha$) to an automorphism of \mathcal{M} . This suggests the following definition.

6.49. Definition. Let κ be an infinite cardinal. A structure \mathcal{M} is *strongly κ -homogeneous* if whenever α is an ordinal with $\alpha < \kappa$ and $(a_i)_{i < \alpha}$, $(a'_i)_{i < \alpha}$ are families in the universe of \mathcal{M} such that $(\mathcal{M}, a_i \mid i < \alpha) \equiv (\mathcal{M}, a'_i \mid i < \alpha)$, there exists an automorphism J of \mathcal{M} such that $J(a_i) = a'_i$ for all $i < \alpha$.

The following theorem shows that, for arbitrarily large κ , every structure has elementary ultrapowers that are κ -saturated and κ -homogeneous.

6.50. Theorem. *For every infinite cardinal κ there exists an ultrafilter \mathcal{U} with the following property: Whenever \mathcal{M} is a metric structure of cardinality at most 2^κ , the \mathcal{U} -ultrapower of \mathcal{M} is both κ^+ -saturated and κ^+ -homogeneous.*

We omit the proof of Theorem 6.50. The construction of the ultrafilter is due to S. Shelah [She71], and it builds on ideas of H. J. Keisler and K. Kunen. Shelah’s epochal proof is for traditional first-order (i.e., discrete structures), but his argument was adapted by C. W. Henson and the second author for structures based on Banach spaces [HI02, Corollary 12.3]. The proof for Banach structures applies to general metric structures without significant changes.

X. EXTENDING THE LANGUAGE

Given a signature L , a sort index set \mathbf{S} , indices $s_0, s_1, \dots, s_n \in \mathbf{S}$, and a new function symbol $f : s_1 \times \dots \times s_1 \rightarrow s_0$, we will denote by $L[f]$ the signature that results from adding f to L . If \mathcal{M} is an L structure based on $(M^{(s)} \mid s \in \mathbf{S})$ and F is a function from $M^{(s_1)} \times \dots \times M^{(s_n)}$ into $M^{(s_0)}$, then will denote by (\mathcal{M}, F) the expansion of \mathcal{M} to $L[f]$ that results from defining $f^{\mathcal{M}}$ as F .

Let L be a signature and consider the signature $L[f]$, where f is a function symbol not in L . Fix also an $L[f]$ -theory T (which could be empty). In standard mathematical practice there are cases where, in models of T , the new symbol can be dispensed with because it is already definable through L -formulas. This leads to the notion of *definability*, which is the main concern of this subsection. There are several ways to formalize the concept of f being T -definable; for instance, one could say that

If $(\mathcal{M}, F) \approx T$, where \mathcal{M} is an L -structure, then F is determined uniquely by T , i.e., if $(\mathcal{M}, F), (\mathcal{M}, F') \approx T$, then $F = F'$.

or, alternatively,

In all models of T , formulas involving f can be approximated by L -formulas.

Below, we prove that these two conditions are equivalent (see Theorem 6.53); but first we must formalize what we mean by “approximated” in the preceding statement.

Recall that if (M, d, a) is a pointed metric space, the open ball of radius r around the anchor point a is denoted or $B_M(r)$, or $B(r)$ if M is given by the context. If \mathcal{M} is a metric structure and a is in $B_{M^{(s)}}(r)$, where $M^{(s)}$ is one of the sorts of \mathcal{M} , we may informally say that a is an element of $B_{\mathcal{M}}(r)$.

6.51. Definition. Let L be a signature, let f be a real-valued n -ary function symbol, and let T be a uniform $L[f]$ -theory. We will say that f is *explicitly defined by T in L* if the following condition holds for every $r \in \mathbb{R}$ and every pair of nonempty intervals $I \subset J \subseteq \mathbb{R}$ with I closed and J open: There exists an L -formula $\varphi_{r,I,J} = \varphi(x_1, \dots, x_n)$ such that, whenever $(\mathcal{M}, F) \approx T$ and a_1, \dots, a_n are elements of $B_{\mathcal{M}}(r)$, we have

$$\begin{aligned} F(a_1, \dots, a_n) \in I &\Rightarrow \mathcal{M} \approx \varphi[a_1, \dots, a_n], \\ \mathcal{M} \approx \varphi[a_1, \dots, a_n] &\Rightarrow F(a_1, \dots, a_n) \in J. \end{aligned}$$

The collection $(\varphi_{r,I,J} \mid r \in \mathbb{R}, \emptyset \neq I \subset J)$ is a *definition scheme* for F (modulo T).

6.52. Remark. A definition scheme $\Sigma = (\varphi_{r,I,J} : r \in \mathbb{R}, \emptyset \neq I \subset J)$ (I closed, J open) characterizes F uniquely in any structure $(\mathcal{M}, F) \approx T$. Namely, if $a_1, \dots, a_n \in B_{\mathcal{M}}(r)$, then $F(a_1, \dots, a_n)$ is the unique real number t with the following property: $\mathcal{M} \approx \varphi_{r,I,J}[a_1, \dots, a_n]$ whenever $J \supset I \ni t$. Certainly, $F(a_1, \dots, a_n)$ is one such number t (since Σ is a definition scheme for F). Conversely, if t has the stated property, we have $F(a_1, \dots, a_n) \in J$ whenever $J \supset I \ni t$, hence $t = F(a_1, \dots, a_n)$.

The theorem below is the most fundamental result about first-order definability. The topological proof we give here is not commonly known. We thank Xavier Caicedo for pointing out an error in an earlier draft of this manuscript and suggesting a correction.

6.53. Theorem (Beth-Svenonius Definability Theorem). *Let L be a signature, let f be a real-valued function symbol, and let T be a uniform $L[f]$ -theory. The following conditions are equivalent.*

- (1) f is explicitly defined by T in L .
- (2) If \mathcal{M} is an L -structure, then

$$(\mathcal{M}, F), (\mathcal{M}, F') \models T \quad \Rightarrow \quad F = F'.$$

Before proving Theorem 6.53, let us make the following observation about general topological spaces. If Z is a topological space, two points $x, y \in Z$ are said to be *topologically indistinguishable*, denoted $x \equiv y$, if every neighborhood of x contains y and every neighborhood of y contains x . Now, we observe that if X, Y are regular topological spaces with X compact and $g : X \rightarrow Y$ is a continuous bijection, then g is a homeomorphism if and only if

$$g(x) \equiv g(y) \quad \Rightarrow \quad x \equiv y. \quad (5)$$

Indeed, any such g maps indistinguishable points to indistinguishable points (by continuity), so g induces a continuous map $\bar{g} : \bar{X} \rightarrow \bar{Y}$ between the spaces $\bar{X} = X/\equiv$ and $\bar{Y} = Y/\equiv$. Since topologically distinguishable points of a normal space have disjoint neighborhoods, \bar{X} and \bar{Y} are Hausdorff spaces, with \bar{X} compact; moreover, \bar{g} is a bijection, by (5), and hence a homeomorphism. Clearly, g is a homeomorphism also.

Proof. (1) \Rightarrow (2): This is an immediate consequence of Remark 6.52.

As a preliminary step to showing (2) \Rightarrow (1), we prove the following:

Claim. Assume that (2) holds. If L' is a signature extending L , and \mathcal{M}, \mathcal{N} are L' -structures admitting $L'[f]$ -expansions $(\mathcal{M}, F), (\mathcal{N}, G)$ that are models of T , then

$$\mathcal{M} \equiv_{L'} \mathcal{N} \quad \Rightarrow \quad (\mathcal{M}, F) \equiv_{L'[f]} (\mathcal{N}, G).$$

To prove the claim, let L' extend L and let \mathcal{M}, \mathcal{N} be elementarily equivalent L' -structures admitting expansions $(\mathcal{M}, F), (\mathcal{N}, G)$ that are models of T . Using Theorem 6.50, fix an ultrafilter \mathcal{U} such that the \mathcal{U} -ultrapowers $(\mathcal{M})_{\mathcal{U}}, (\mathcal{N})_{\mathcal{U}}$ are isomorphic L' -structures. By Łoś's Theorem 6.29, the \mathcal{U} -ultrapower $(\mathcal{M}, F)_{\mathcal{U}} = ((\mathcal{M})_{\mathcal{U}}, (F)_{\mathcal{U}})$ is an $L'[f]$ -structure elementarily equivalent to (\mathcal{M}, F) ; similarly, $(\mathcal{N}, G)_{\mathcal{U}} = ((\mathcal{N})_{\mathcal{U}}, (G)_{\mathcal{U}}) \equiv_{L'[f]} (\mathcal{N}, G)$. Let \mathcal{J} be an L' -isomorphism from $(\mathcal{M})_{\mathcal{U}}$ into $(\mathcal{N})_{\mathcal{U}}$. *A fortiori*, \mathcal{J} is an L -isomorphism between (the L -reducts of) $(\mathcal{M})_{\mathcal{U}}$ and $(\mathcal{N})_{\mathcal{U}}$. By (2), we have $\mathcal{J}((F)_{\mathcal{U}}) = (G)_{\mathcal{U}}$; thus, $(\mathcal{M}, F)_{\mathcal{U}}$ and $(\mathcal{N}, G)_{\mathcal{U}}$ are isomorphic $L'[f]$ -structures, so $(\mathcal{M}, F) \equiv_{L'[f]} (\mathcal{N}, G)$. This proves the claim.

(2) \Rightarrow (1): Assume that (2) holds. The signature of $L[f]$ specifies that $f : s_1 \times \cdots \times s_n \rightarrow s_{\mathbb{R}}$ for some sorts s_1, \dots, s_n . Fix a new tuple $\bar{c} = c_1, \dots, c_n$ of constant symbols such that c_i is of sort s_i for $i = 1, \dots, n$ and let $L' = L[\bar{c}]$. For any fixed rational $r > 0$, let T'_r be the $L'[f]$ -theory obtained by adding to T the sentences $d(c_i, a_{s_i}) \leq r$, for $i = 1, \dots, n$, where a_{s_i} is (the symbol for) the anchor of sort s_i . Clearly, T'_r is a uniform $L'[f]$ -theory. Let \mathcal{C} be the class of all models of T'_r , i.e., of $L'[f]$ -structures of the form $(\mathcal{M}, b_1, \dots, b_n, F)$, where (\mathcal{M}, F) is a model of T , and b_i is an element of sort s_i of \mathcal{M} that satisfies $d(b_i, a_{s_i}) \leq r$ for $i = 1, \dots, n$. Let \mathcal{D} be the class of all L' -structures \mathcal{N} that are L' -reducts of some $L'[f]$ -structure $(\mathcal{N}, F) \in \mathcal{C}$ and let $g : \mathcal{C} \rightarrow \mathcal{D}$ be the map $(\mathcal{N}, F) \mapsto \mathcal{N}$. Note that g is surjective by definition.

We regard the classes \mathcal{C} and \mathcal{D} as topological spaces as follows. A basis for the closed classes of \mathcal{C} is given by all subclasses of the form $\text{Mod}_{\mathcal{C}}(\varphi)$ where φ is an $L'[f]$ -sentence (the union of finitely many such basic closed subclasses is of the same form, since the logic is closed under disjunction); in other words, the closed subclasses of \mathcal{C} are those of the form $\text{Mod}_{\mathcal{C}}(\Phi)$, where Φ is an $L'[f]$ -theory. Similarly, the closed subclasses of \mathcal{D} are those of the form $\text{Mod}_{\mathcal{D}}(\Phi)$, where Φ is an L' -theory.⁹ Both \mathcal{C} and \mathcal{D} are regular, with \mathcal{C} compact, by Remark 6.33, and the surjection g is clearly continuous.

Let $\mathcal{N} \in \mathcal{D}$, so \mathcal{N} is the L' -reduct of some model $(\mathcal{N}, F) \models T'_r$. Let $\mathcal{M} = \mathcal{N} \upharpoonright L$ be the L -reduct of \mathcal{N} . Clearly, (\mathcal{M}, F) is a model of $T'_r \upharpoonright L[f] = T$; moreover, by (2), (\mathcal{M}, F) is the unique expansion of \mathcal{M} to a model of T . Since T'_r extends T , (\mathcal{N}, F) is necessarily the unique preimage of \mathcal{N} under g , so g is injective; thus, g is a continuous bijection. By the claim and the topological observation immediately following the statement of the theorem, we conclude that g is a homeomorphism.

Fix intervals $\emptyset \neq I = [p, q] \subset J = (u, v)$ with p, q, u, v rational, and let

$$K = \text{Mod}_{\mathcal{C}}(p \leq f(\bar{c}) \leq q), \quad K' = \text{Mod}_{\mathcal{C}}(f(\bar{c}) \leq u \vee f(\bar{c}) \geq v).$$

Since K and K' are closed and disjoint, so are the subsets $g(K)$ and $g(K')$ of the homeomorphic image $\mathcal{D} = g(\mathcal{C})$, which is compact since \mathcal{C} is. Since any compact regular space is normal, there exists a closed neighborhood Q of $g(K)$ disjoint from $g(K')$. By compactness of $g(K')$, Q may be taken of the form $\text{Mod}_{\mathcal{D}}(\varphi(\bar{c}))$ for some L -formula $\varphi_{I,J,r} = \varphi(\bar{x})$. Since the interpretation of \bar{c} is arbitrary in $B_{\mathcal{M}}(r)$, the scheme $(\varphi_{I,J,r} : r \in \mathbb{R}, \emptyset \neq I \subset J)$ defines f explicitly by T . \square

6.54. Definition. If L is a signature, f is a real-valued function symbol, and T is a uniform $L[f]$ -theory, we will say that f is T -definable in L if it satisfies the equivalent conditions of Theorem 6.53. When L is given by the context, we may simply say that f is T -definable; furthermore, if L and T are given by the context, we may say that f is “definable”.

If T is a uniform theory, Theorem 6.53 allows us to see the real-valued functions that are T -definable as those that are left fixed by automorphisms of sufficiently saturated models of T . This observation yields Corollary 6.55 below.

For notational convenience, if f is a function symbol for a given function $F : \mathbb{R}^n \rightarrow \mathbb{R}$, we will liberally identify f with its interpretation F .

6.55. Corollary. *Let T be a uniform theory.*

- (1) *A composition of functions that are T -definable is T -definable.*
- (2) *Every continuous function $F : \mathbb{R}^n \rightarrow \mathbb{R}$ is T -definable.*
- (3) *If $F : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ is T -definable by T , so are the functions $G_r : \mathbb{R}^n \rightarrow \mathbb{R}$ and $H_r : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by*

$$G_r(x_1, \dots, x_n) = \sup_{y \in B(r)} F(x_1, \dots, x_n, y),$$

$$H_r(x_1, \dots, x_n) = \inf_{y \in B(r)} F(x_1, \dots, x_n, y).$$

Proof. Clauses (1) and (3) follow from the definitions. To prove (2), note that, since every signature comes equipped with the ordered field and lattice structure for \mathbb{R} plus a constant for each rational, all polynomial functions with rational coefficients are definable in any theory. Therefore, (2) follows from the Stone-Weierstrass theorem. \square

⁹Later, in Subsection XI, we shall refer to this as the *logic topology*.

6.56. **Definition.** If L is a signature, f is an n -ary function symbol (not necessarily real-valued), and T is a uniform $L[f]$ -theory, we will say that f is T -definable in L if the real-valued function $d(f(x_1, \dots, x_n), y)$ is T -definable in L in the sense of Definition 6.54. As in Definition 6.54, if L or T are given by the context, we may omit them from the nomenclature.

6.57. **Remark.** It is not difficult to verify that, if f is a real-valued function symbol, there is no conflict between the notions of definability for f given by Definitions 6.54 and 6.56. (This is because if f is n -ary and F is an interpretation of f , then for any real number r , $d(F(a_1, \dots, a_n), b) = r$ if and only if $b = \pm F(a_1, \dots, a_n)$.)

XI. SPACES OF TYPES AND THE MONSTER MODEL

In this section, T will denote a fixed complete L -theory and we will denote by \mathcal{T} the set of types that are realized in a model of T . Note that, if $(t_i)_{i \in I}$ is a family of types in \mathcal{T} and t_i is realized in \mathcal{M}_i then, by Corollary 6.38, there exists a model \mathcal{N} of T such that $\mathcal{M}_i \prec \mathcal{N}$ for every $i \in I$. Thus, each t_i is realized in \mathcal{N} .

By Remark 6.42-(2), if $t = t(x_1, \dots, x_n)$ is a type in \mathcal{T} , then there exists a positive rational r such that

$$\exists_r x_1 \dots \exists_r x_n \bigwedge_{\varphi \in \varphi} \varphi(x_1, \dots, x_n) \in t, \quad \text{for each finite } \varphi \subseteq t.$$

For each choice of n and r , we will denote by $\mathcal{T}_n^{(r)}$ be the set of all $t(x_1, \dots, x_n) \in \mathcal{T}$ that satisfy this condition.

If $\varphi(x_1, \dots, x_n)$ is an L -formula, let

$$[\varphi] = \{t \in \mathcal{T} \mid \varphi \in t\}.$$

The collection of sets of this form is closed under finite unions and intersections. We define a topology on \mathcal{T} by letting these be the basic closed sets. We will refer to this topology as the *logic topology*. We shall always regard \mathcal{T} as a topological space via the logic topology.

6.58. **Remark.** The Compactness Theorem 6.31 says exactly that $\mathcal{T}_n^{(r)}$ is compact for each n and r .

The complement of a set of the form $[\varphi]$ is not necessarily of the same form; however, if $t \in \mathcal{T}$, by Proposition 6.43-(3b), we have $t \notin [\varphi]$ if and only if there exists $\psi \in t$ and $\psi' > \psi$ such that $t \in [\psi'] \subseteq [\varphi]^c$. Thus, for $t \in \mathcal{T}$, the sets of the form $[\psi']$ where ψ' is an approximation of a formula $\psi \in t$ form a local neighborhood base around t .

Let $\varphi(x_1, \dots, x_n)$ be the formula $\exists_r x_1 \dots \exists_r x_n \bigwedge_{1 \leq i \leq n} x_i = x_i$. Then, $[\varphi] = \mathcal{T}_n^{(r)}$. Every type contains a formula φ of this kind, for some n and some r . Therefore, the space \mathcal{T} is locally compact.

Let A be a set of parameters, i.e., A is a set of elements of some fixed model \mathcal{M} of T . If $(t_i)_{i \in I}$ is a family in $\mathcal{T}_n^{(r)}$ such that t_i is a type over A for each $i \in I$ and \mathcal{U} is an ultrafilter of I , then, by the compactness of $\mathcal{T}_n^{(r)}$, the limit $\mathcal{U} \lim_i t_i$ is a type over A . Conversely, if $t(\bar{x})$ is a type over A , since t is finitely satisfiable in \mathcal{M} (see Definition 6.40), there exists a family $(\bar{a}_i)_{i \in I}$ in \mathcal{M} and an ultrafilter \mathcal{U} on I such that $t = \mathcal{U} \lim_i \text{tp}_A(\bar{a}_i)$. Thus, types over model \mathcal{M} can be viewed as ultrafilter limits of types realized within \mathcal{M} .

Let κ be a cardinal that is both larger than the cardinality of every model of T needed in our proofs and larger than the cardinality of any indexed family of types mentioned. Let now \mathfrak{C} be a κ -saturated, κ -homogeneous model of T (see Theorem 6.50). Then we can assume that:

- (1) Every model of T occurs as an elementary submodel of \mathfrak{C} (see Proposition 6.46).
- (2) Every type (over a set of parameters) is a model of T is realized in \mathfrak{C} .
- (3) If (a_1, \dots, a_n) and (b_1, \dots, b_n) are n -tuples of elements of \mathfrak{C} such that $\text{tp}_A(a_1, \dots, a_n) = \text{tp}_A(b_1, \dots, b_n)$, then there exists an automorphism f of \mathfrak{C} such that $f(a_i) = b_i$ for $i = 1, \dots, n$.
- (4) A real-valued function f is definable in T over a set of parameters A if and only if f is fixed by every automorphism of \mathfrak{C} that fixes A pointwise fixes the graph of f . (See the remarks preceding Corollary 6.55.)

We shall henceforth refer to \mathfrak{C} as a *big model* or a *monster model* for T . Monster models are a time-saving device, and what makes them convenient are the properties listed above. For each complete theory considered, we will always work within a fixed monster model. The particular choice of monster model will be irrelevant for our discussions since any two such models can be embedded elementarily in a common one. Thus, we may informally refer to “the” monster model of T .

There is another topology on types that plays an important role in the model theory of metric spaces. Let $t(x_1, \dots, x_n)$ and $t'(x_1, \dots, x_n)$ are types consistent with T (i.e., t and t' are realized in the monster model) over the same set of parameters (which can be thought of as a subset of the monster model). We define $d(t, t')$ as the infimum of all distances $d(\bar{a}, \bar{a}')$ where \bar{a} realizes t and \bar{a}' realizes t' . Note that this infimum is realized by some pair \bar{a}, \bar{a}' in the monster model. It can be readily verified, using the saturation of the monster model, that d is a metric and that $(\mathcal{T}_n^{(r)}, d)$ is compact and complete.

APPENDIX: ON THE NOTION OF “FINITARY” PROPERTIES OF METRIC STRUCTURES

We conclude this manuscript with some general remarks on the meaning of the informal term “finitary” in our context. Describing a certain mathematical property as finitary presupposes, in our view, the existence of a formal language \mathcal{L} in which the property can be formulated (here we use the term “language” in an abstract sense not restricted to first-order languages or Henson’s language of positive bounded formulas). There is an inherent tradeoff between the expressive power of logical languages and the strength of their model-theoretic properties. On one hand, if the language \mathcal{L} is rich enough (for example, if \mathcal{L} admits infinite conjunctions and disjunctions), it may capture complex properties with, say, a single formula; however, a powerful model-theoretic property like the Compactness Theorem (Theorem 6.31) is bound to fail for such \mathcal{L} . On the other hand, compactness of first-order-like logics (including Henson’s logic) is, in essence, a reflection of the limited expressive power of the language. In fact, for metric structures, there is no logic strictly more expressive than Henson logic satisfying both the Compactness Theorem and the elementary chain property (Proposition 6.28) [Iov01].

A strong feature of the notion of approximate satisfaction of positive bounded formulas is that it inherently captures “asymptotically approximable” properties of L -structures: By definition, the approximate satisfaction of an L -formula φ amounts to the discrete satisfaction of the full set φ_+ of formulas ψ approximating φ . Now, if Φ is a set of positive bounded L -formulas such that $\Phi \approx \varphi$ (i.e., every model of Φ satisfies φ approximately), then every approximation of φ admits a finite-length proof from Φ plus the axioms for the real numbers and metric spaces.¹⁰ but φ itself may not admit a single such proof. A property P captured

¹⁰This was first observed by C. Ward Henson in the 1980s. He gave lectures on this material, but his lecture notes were not formally published. The idea of finite provability via approximations can be traced

by the approximate satisfaction of a single formula φ should by all rights be called finitary, although P may not admit a finite-length proof. Still, it is enough for each rational $\epsilon > 0$ to prove φ_ϵ , the formula obtained from φ by ϵ -relaxing every inequality and quantifier bound in φ . This approach, though *sensu stricto* infinitary, is much the one used when a proof in analysis starts with the incantation: “Let $\epsilon > 0$ be given.” (Of course, it is always desirable that “the same” proof works for all $\epsilon > 0$ in the sense that ϵ only appears in formulas used in the proof as a parameter, so that a uniform scheme proves all φ_ϵ .)

More generally, any property P equivalent to the simultaneous satisfaction of a collection Ψ of positive bounded formulas is finitary in the following sense: If P holds for all structures in a class \mathcal{C} axiomatized by a theory T , then P admits, in principle, a proof scheme (say, a syntactic proof for each $\psi \in \Psi_+$). Equivalently, the *failure* of P in a structure \mathcal{M} is equivalent to the discrete satisfaction $\mathcal{M} \models^{\text{w}} \psi$ of the weak negation of a single formula $\psi \in \Psi_+$. Thus, a finitary property is witnessed by the absence of counterexamples having finite-length proofs.¹¹

The property that a sequence (a_n) in a metric L -structure be convergent cannot, in general, be captured by the simultaneous satisfaction of any collection (whether finite or infinite) of positive bounded L -formulas. However, the property “ (a_n) is convergent” is equivalent to the disjunction of the properties “ E_\bullet is a rate of metastability for (a_n) ” over all possible metastability rates E_\bullet . For a specific E_\bullet , the latter property is equivalent to the conjunction of the (infinitely many) properties “ $E_{\epsilon,\eta}$ is a rate of $[\epsilon, \eta]$ -metastability for (a_n) ” for all $\epsilon > 0$ and samplings η . Thus, from the perspective discussed above, metastable convergence with rate E_\bullet is a finitary property of metric structures, while convergence with no specified rate is not.

REFERENCES

- [ADR12] Jeremy Avigad, Edward T. Dean, and Jason Rute. A metastable dominated convergence theorem. *J. Log. Anal.*, 4:Paper 3, 19, 2012.
- [AI13] Jeremy Avigad and José Iovino. Ultraproducts and metastability. *New York J. Math.*, 19:713–727, 2013.
- [BYBHU08] Itai Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov. Model theory for metric structures. In *Model theory with applications to algebra and analysis. Vol. 2*, volume 350 of *London Math. Soc. Lecture Note Ser.*, pages 315–427. Cambridge Univ. Press, Cambridge, 2008.
- [BYU10] Itai Ben Yaacov and Alexander Usvyatsov. Continuous first order logic and local stability. *Trans. Amer. Math. Soc.*, 362(10):5213–5259, 2010.
- [CI14] Xavier Caicedo and José N. Iovino. Omitting uncountable types and the strength of $[0, 1]$ -valued logics. *Ann. Pure Appl. Logic*, 165(6):1169–1200, 2014.
- [CK62] Chen-chung Chang and H. Jerome Keisler. Model theories with truth values in a uniform space. *Bull. Amer. Math. Soc.*, 68:107–109, 1962.
- [CK66] Chen-chung Chang and H. Jerome Keisler. *Continuous model theory*. Annals of Mathematics Studies, No. 58. Princeton Univ. Press, Princeton, N.J., 1966.

back to Mostowski, who proved [Mos62] that for any rational $r < 1$, the set of sentences of Łukasiewicz logic that take truth value $> r$ in all structures is recursively enumerable. Simultaneously, E. Specker’s student B. Scarpellini had proved in his 1961 dissertation that the valid sentences of Łukasiewicz logic are not axiomatizable.

¹¹As pointed out by Avigad *et al.* [ADR12], from the perspective of constructive analysis, this is an instance of Kreisel’s no-counterexample interpretation [Kre51, Kre52], which is in turn a particular case of Gödel’s Dialectica interpretation [G58].

- [Cut83] Nigel J. Cutland. Nonstandard measure theory and its applications. *Bull. London Math. Soc.*, 15(6):529–589, 1983.
- [Cut00] Nigel J. Cutland. *Loeb measures in practice: recent advances*, volume 1751 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 2000.
- [G58] Kurt Gödel. über eine bisher noch nicht benützte Erweiterung des finiten Standpunktes. *Dialectica*, 12:280–287, 1958.
- [Hen75] C. Ward Henson. When do two Banach spaces have isometrically isomorphic nonstandard hulls? *Israel J. Math.*, 22(1):57–67, 1975.
- [Hen76] C. Ward Henson. Nonstandard hulls of Banach spaces. *Israel J. Math.*, 25(1-2):108–144, 1976.
- [HI02] C. Ward Henson and José Iovino. Ultraproducts in analysis. In *Analysis and logic (Mons, 1997)*, volume 262 of *London Math. Soc. Lecture Note Ser.*, pages 1–110. Cambridge Univ. Press, Cambridge, 2002.
- [Iov01] José Iovino. On the maximality of logics with approximations. *J. Symbolic Logic*, 66(4):1909–1918, 2001.
- [Iov09] José Iovino. Analytic structures and model theoretic compactness. In *The many sides of logic*, volume 21 of *Stud. Log. (Lond.)*, pages 187–199. Coll. Publ., London, 2009.
- [Iov14] José Iovino. *Applications of model theory to functional analysis*. Dover Publications, Inc., Mineola, NY, 2014.
- [Kre51] G. Kreisel. On the interpretation of non-finitist proofs. I. *J. Symbolic Logic*, 16:241–267, 1951.
- [Kre52] G. Kreisel. On the interpretation of non-finitist proofs. II. Interpretation of number theory. Applications. *J. Symbolic Logic*, 17:43–58, 1952.
- [Loe75] Peter A. Loeb. Conversion from nonstandard to standard measure spaces and applications in probability theory. *Trans. Amer. Math. Soc.*, 211:113–122, 1975.
- [Loś55] Jerzy Loś. Quelques remarques, théorèmes et problèmes sur les classes définissables d’algèbres. In *Mathematical interpretation of formal systems*, pages 98–113. North-Holland Publishing Co., Amsterdam, 1955.
- [Mos62] A. Mostowski. Axiomatizability of some many valued predicate calculi. *Fund. Math.*, 50:165–190, 1961/1962.
- [Ros97] David A. Ross. Loeb measure and probability. In *Nonstandard analysis (Edinburgh, 1996)*, volume 493 of *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, pages 91–120. Kluwer Acad. Publ., Dordrecht, 1997.
- [She71] Saharon Shelah. Every two elementarily equivalent models have isomorphic ultrapowers. *Israel J. Math.*, 10:224–233, 1971.
- [Tao] Terence Tao. Walsh’s ergodic theorem, metastability, and external Cauchy convergence. <http://terrytao.wordpress.com>.
- [Tao08] Terence Tao. Norm convergence of multiple ergodic averages for commuting transformations. *Ergodic Theory Dynam. Systems*, 28(2):657–688, 2008.
- [Wal12] Miguel N. Walsh. Norm convergence of nilpotent ergodic averages. *Ann. of Math. (2)*, 175(3):1667–1688, 2012.

DEPARTMENT OF MATHEMATICS, THE UNIVERSITY OF TEXAS AT SAN ANTONIO, ONE UTSA CIRCLE,
SAN ANTONIO, TX 78249-0664, U.S.A.

E-mail address: eduardo.duenez@utsa.edu

DEPARTMENT OF MATHEMATICS, THE UNIVERSITY OF TEXAS AT SAN ANTONIO, ONE UTSA CIRCLE,
SAN ANTONIO, TX 78249-0664, U.S.A.

E-mail address: jose.iovino@utsa.edu