



HAL
open science

Continuous first order logic and local stability

Itai Ben Yaacov, Alexander Usvyatsov

► **To cite this version:**

Itai Ben Yaacov, Alexander Usvyatsov. Continuous first order logic and local stability. Transactions of the American Mathematical Society, 2010, 362 (10), pp.5213-5259. 10.1090/S0002-9947-10-04837-3 . hal-00220454

HAL Id: hal-00220454

<https://hal.science/hal-00220454>

Submitted on 28 Jan 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CONTINUOUS FIRST ORDER LOGIC AND LOCAL STABILITY

ITAI BEN YAACOV AND ALEXANDER USVYATSOV

ABSTRACT. We develop continuous first order logic, a variant of the logic described in [CK66]. We show that this logic has the same power of expression as the framework of open Hausdorff cats, and as such extends Henson's logic for Banach space structures. We conclude with the development of local stability, for which this logic is particularly well-suited.

INTRODUCTION

A common trend in modern model theory is to generalise model-theoretic notions and tools to frameworks that go beyond that of first order logic and elementary classes and properties. In doing this, there is usually a trade-off: the more general the framework, the weaker the available tools, and one finds oneself many times trying to play this trade-off, looking for the most general framework in which a specific argument can be carried through. The authors admit having committed this sin not once.

The present paper is somewhat different, though: we do present what seems to be a new framework, or more precisely, a new logic, but in fact we prove that it is completely equivalent to one that has been previously defined elsewhere, namely that of (metric) open Hausdorff cats (see [Ben05]).

Another logic dealing with metric structures is Henson's logic of positive bounded formulae and approximate satisfaction (see for example [HI02]). Even though Henson's logic was formulated for unbounded Banach space structures while ours deal with bounded metric structures, it is fair to say that the two logics are equivalent. First of all, Henson's approach makes perfect sense in the bounded setting in which case the two logics are indeed equivalent. Banach space structures can (in most cases) be reduced for model-theoretic purposes to their closed unit ball (see Example 4.5). Moreover, there exists an *unbounded* variant of continuous logic which is equivalent with to (a somewhat extended) Henson's logic for arbitrary (bounded or unbounded) metric structures. It can be reduced back to continuous logic as studied here (i.e., bounded) via the addition of a

Date: 28 January 2008.

2000 Mathematics Subject Classification. 03C90,03C45.

Key words and phrases. continuous logic, metric structures, stability, local stability.

Research of the first author supported by NSF grant DMS-0500172.

The authors would like to thank C. Ward Henson for stimulating discussions, and Sylvia Carlisle and Eric Owiesny for a careful reading of the manuscript.

single point at infinity. This beyond the scope of the present paper and is discussed in detail in [Benc].

Finally, this logic is *almost* a special case of the continuous first order logic that Chang and Keisler studied in [CK66]. We do differ with their definitions on several crucial points, where we find they were too general, or not general enough. Our logic is a special case in that instead of allowing any compact Hausdorff space X as a set of truth values, we find that letting X be the unit interval $[0, 1]$ alone is enough. Indeed, as every compact Hausdorff space embeds into a power of the interval, there is no loss of generality. Similarly, as unit interval admits a natural complete linear ordering, we may eliminate the plethora of quantifiers present in [CK66], and the arbitrary choices involved, in favour of two canonical quantifiers, \sup and \inf , which are simply the manifestations in this setting of the classical quantifiers \forall and \exists . On the other hand, extending Chang and Keisler, we allow the “equality symbol” to take any truth value in $[0, 1]$. Thus, from an *equality* symbol it becomes a *distance* symbol, allowing us to interpret metric structures in the modified logic.

However, continuous first order logic has significant advantages over previous formalisms for metric structures. To begin with, it is an immediate generalisation of classical first order logic, more natural and less technically involved than previous formalisms. More importantly, it allows us to beat the above-mentioned trade-off. Of course, if two logics have the same power of expression, and only differ in presentation, then an argument can be carried in one if and only if it can be carried in the other; but it may still happen that notions which arise naturally from one of the presentations are more useful, and render clear and obvious what was obscure with the other one. This indeed seems to be the case with continuous first order logic, which further supports our contention that it is indeed the “true and correct” generalisation of classical first order logic to the context of metric structures arising in analysis.

An example for this, which was part of the original motivation towards these ideas, is a question by C. Ward Henson, which can be roughly stated as “how does one generalise local (i.e., formula-by-formula) stability theory to the logic of positive bounded formulae?”. The short answer, as far as we can see, is “one doesn’t.” The long answer is that positive bounded formulae may not be the correct analogues of first order formulae for these purposes, whereas continuous first order formulae are.

In Section 1 we define the syntax of continuous first order logic: signatures, connectives, quantifiers, formulae and conditions.

In Section 2 we define the semantics: pre-structures, structures, the special role of the metric and truth values.

In Section 3 we discuss types and definable predicates. The family of definable predicates is the completion, in some natural sense, of the family of continuous formulae.

In Section 4 we discuss continuous first order theories and some basic properties such as quantifier elimination. We also compare continuous first order theories with previous notions such as open Hausdorff cats.

In Section 5 we discuss imaginaries as canonical parameters to formulae and definable predicates.

In Section 6 we define φ -types, i.e., types which only depend on values of instances of a formula φ .

In Section 7 we develop local stability, answering Henson's question.

In Section 8 we show how to deduce the standard global theory of independence from the local one in a stable theory.

We also have two appendices:

Appendix A contains a remark concerning an alternative (and useful) presentation of continuity moduli.

Appendix B deals with the case of a formula which is stable in a single model of a theory.

1. CONTINUOUS FIRST ORDER FORMULAE

In classical (first order) logic there are two possible truth values: “true”, sometimes denoted by \top or T , and “false”, denoted by \perp or F . Often enough one associates the classical truth values with numerical values, and the most common choice is probably to assign T the value 1 and F the value 0. This assignment is not sacred, however, and for our purposes the opposite assignment, i.e., $T = 0$ and $F = 1$, fits more elegantly.

The basic idea of this paper is to repeat the development of first order logic with one tiny difference: we replace the *finite* set of truth values $\{0, 1\}$ with the *compact* set $[0, 1]$. Everything else should follow naturally from this modification. We will refer to the classical framework also as *discrete logic*, whereas the one we develop here will be referred to as *continuous logic*.

As in classical logic, a *continuous signature* \mathcal{L} is a set of function symbols and predicate symbols, each of which having an associated arity $n < \omega$. In an actual continuous structure, the function symbols will be interpreted as functions from the structure into itself, and the predicate symbols as functions to the set of truth values, i.e., the interval $[0, 1]$.

For the definition of pure syntax we may restrict ourselves to non-metric signatures, which are the analogues of classical signatures without equality.

Definition 1.1. A *non-metric continuous signature* consists of a set of function symbols and predicate symbols, and for each function symbol f or predicate symbol P , its arity $n_f < \omega$ or $n_P < \omega$.

We may also consider multi-sorted signatures, in which case the arity of each symbol specify how many arguments are in each sort, such that the total is finite, and each function symbol has a target sort.

Given a continuous signature \mathcal{L} , we define \mathcal{L} -terms and atomic \mathcal{L} -formulae as usual. However, since the truth values of predicates are going to be in $[0, 1]$, rather than in $\{0, 1\}$, we need to adapt our connectives and quantifiers accordingly.

Let us start with connectives. In the discrete setting we use a somewhat fixed set of unary and binary Boolean connectives, from which we can construct any n -ary Boolean expression. In other words, any mapping from $\{0, 1\}^n \rightarrow \{0, 1\}$ can be written using these connectives (otherwise, we would have introduced additional ones). By analogy, an n -ary continuous connective should be a *continuous* mapping from $[0, 1]^n \rightarrow [0, 1]$, and we would like to have a set of connectives with which we can construct every continuous mapping $[0, 1]^n \rightarrow [0, 1]$, for every n . However, this may be problematic, as continuum many connectives would give rise to uncountably many formulae even in a countable signature. To avoid this anomaly we will content ourselves with a set of connectives which merely allows to construct arbitrarily good approximations of every continuous mapping $[0, 1]^n \rightarrow [0, 1]$.

Common connectives we use, by arity:

- Constants in $[0, 1]$.
- $\neg x = 1 - x$, and $\frac{x}{2}$.
- $x \wedge y = \min\{x, y\}$, $x \vee y = \max\{x, y\}$, $x \dot{-} y = (x - y) \vee 0$, $x \dot{+} y = (x + y) \wedge 1$, $|x - y|$.

We can express the non-constant connectives above in terms of the connectives \neg and $\dot{-}$:

$$\begin{aligned} x \wedge y &= x \dot{-} (x \dot{-} y) \\ x \vee y &= \neg(\neg x \wedge \neg y) \\ x \dot{+} y &= \neg(\neg x \dot{-} y) \\ |x - y| &= (x \dot{-} y) \vee (y \dot{-} x) = (x \dot{-} y) \dot{+} (y \dot{-} x) \end{aligned}$$

The expression $x \dot{-} ny$ is a shorthand for $((x \dot{-} y) \dot{-} y) \dots \dot{-} y$, n times. We would also like to point out to the reader that the expression $x \dot{-} y$ is the analogue of the Boolean expression $y \rightarrow x$. For example, the continuous Modus Ponens says that if both y and $x \dot{-} y$ are true, i.e., equal to zero, then so is x .

Definition 1.2. (i) A *system of continuous connectives* is a sequence $\mathcal{F} = \{F_n : n < \omega\}$ where each F_n is a collection of continuous functions from $[0, 1]^n$ to $[0, 1]$.

(ii) We say that a system of continuous connectives \mathcal{F} is *closed* if it satisfies:

- (a) For all $m < n < \omega$, the projection on the m th coordinate $\pi_{n,m} : [0, 1]^n \rightarrow [0, 1]$ belongs to F_n .
- (b) Let $f \in F_n$, and $g_0, \dots, g_{n-1} \in F_m$. Then the composition $f \circ (g_0, \dots, g_{n-1}) \in F_m$.

If \mathcal{F} is any system of continuous connectives, then $\bar{\mathcal{F}}$ is the closed system it generates.

(iii) We say that a closed system of continuous connectives \mathcal{F} is *full* if for every $0 < n < \omega$, the set F_n is dense in the set of all continuous functions $\{f : [0, 1]^n \rightarrow [0, 1]\}$ in the compact-open (i.e., uniform convergence) topology. An arbitrary system of continuous connectives \mathcal{F} is full if $\bar{\mathcal{F}}$ is.

(We exclude $n = 0$ in order to allow full systems of connectives without truth constants, i.e., in which F_0 is empty.)

Fact 1.3 (Stone-Weierstrass Theorem, lattice version). *Let X be a compact Hausdorff space containing at least two points, $I \subseteq \mathbb{R}$ an interval, and equip $\mathfrak{A} = C(X, I)$ with the uniform convergence topology. Let $\mathfrak{B} \subseteq \mathfrak{A}$ be a sub-lattice, such that for every distinct $x, y \in X$, $a, b \in I$, and $\varepsilon > 0$, there is $f \in \mathfrak{B}$ such that $|f(x) - a|, |f(y) - b| < \varepsilon$. Then \mathfrak{B} is dense in \mathfrak{A} .*

Proof. The proof of this or very similar results should appear in almost any analysis textbook. We will nonetheless include the proof for completeness.

Let $f \in \mathfrak{A}$ and $\varepsilon > 0$ be given. For each pair of points $x, y \in X$ we can by hypothesis find $g_{x,y} \in \mathfrak{B}$ for which $|g_{x,y}(x) - f(x)|, |g_{x,y}(y) - f(y)| < \varepsilon$. (In case $x = y$ we take $g_{x,x} = g_{x,z}$ for any $z \neq x$.) The set $V_{x,y} = \{z \in X : f(z) - \varepsilon < g_{x,y}(z)\}$ is an open neighbourhood of y .

Let us fix x . The family $\{V_{x,y} : y \in X\}$ is an open covering of X , and admits a finite sub-covering $\{V_{x,y_i} : i < n\}$. Let $g_x = \bigvee_{i < n} g_{x,y_i} \in \mathfrak{B}$. Then $f(z) - \varepsilon < g_x(z)$ for all $z \in X$ and $|g_x(x) - f(x)| < \varepsilon$. Thus $U_x = \{z \in X : g_x(z) < f(z) + \varepsilon\}$ is an open neighbourhood of x .

Now let x vary. The family $\{U_x : x \in X\}$ is an open covering of X admitting a finite sub-covering $\{U_{x_j} : j < m\}$. Let $g = \bigwedge_{j < m} g_{x_j}$. Then $f(z) - \varepsilon < g(z) < f(z) + \varepsilon$ for all $z \in X$, i.e., $\|g(z) - f(z)\| < \varepsilon$ as desired. ■_{1.3}

It will be more convenient to use the following consequence, which is analogous to the Stone-Weierstrass characterisation of dense algebras of functions:

Corollary 1.4. *Let X be a compact Hausdorff space and let $\mathfrak{A} = C(X, [0, 1])$. Assume that $\mathfrak{B} \subseteq \mathfrak{A}$ is closed under \neg and \div , separates points in X (i.e., for every two distinct $x, y \in X$ there is $f \in \mathfrak{B}$ such that $f(x) \neq f(y)$), and satisfies either of the following two additional properties:*

- (i) *The set $C = \{c \in [0, 1] : \text{the constant } c \text{ is in } \mathfrak{B}\}$ is dense in $[0, 1]$.*
- (ii) *\mathfrak{B} is closed under $x \mapsto \frac{x}{2}$.*

Then \mathfrak{B} is dense in \mathfrak{A} .

Proof. Since \mathfrak{B} is closed under \neg and \div it is also closed under \vee and \wedge , so it is a sub-lattice of \mathfrak{A} .

Since \mathfrak{B} separates points it is in particular non-empty, so we have $0 = f \div f \in \mathfrak{B}$ for any $f \in \mathfrak{B}$, whereby $1 = \neg 0 \in \mathfrak{B}$. In case \mathfrak{B} is closed under $\frac{x}{2}$ we conclude that $1/2^n \in \mathfrak{B}$ for all n , and since \mathfrak{B} is also closed under \div , \mathfrak{B} contains all the dyadic constants in $[0, 1]$ which are dense in $[0, 1]$. We may therefore assume that \mathfrak{B} contains a dense set of constants.

Let $x, y \in X$ be distinct, $a, b \in [0, 1]$ and $\varepsilon > 0$. Let us first treat the case where $X = [0, 1]$ and $\text{id}_{[0,1]} \in \mathfrak{B}$. Then $x, y \in [0, 1]$, and without loss of generality we may

assume that $x < y$. Assume first that $a \geq b$. Let $m \in \mathbb{N}$ be such that $\frac{a}{y-x} < m$, and let $f_0(t) = (a \dot{-} m(t \dot{-} x)) \vee b$. Then $f_0(x) = a \vee b = a$, and $f_0(y) = 0 \vee b = b$. Replacing the constants $x, a, b \in [0, 1]$ with close enough approximations from C we obtain $f \in \mathfrak{B}$ such that $|f(x) - a|, |f(y) - b| < \varepsilon$. If $a < b$, find $f \in \mathfrak{B}$ such that $|f(x) - \neg a|, |f(y) - \neg b| < \varepsilon$, and then $\neg f \in \mathfrak{B}$ is as required.

We now return to the general case where X can be any compact Hausdorff space. Since $x \neq y$ there is $g \in \mathfrak{B}$ such that $g(x) \neq g(y)$, and without loss of generality we may assume that $g(x) < g(y)$. As in the previous paragraph construct $f: [0, 1] \rightarrow [0, 1]$ such that $|f(g(x)) - a|, |f(g(y)) - b| < \varepsilon$, and observe that $f \circ g \in \mathfrak{B}$.

We have shown that \mathfrak{B} satisfies the hypotheses of Fact 1.3 and is therefore dense in \mathfrak{A} . ■_{1.4}

Corollary 1.5. *Let $C \subseteq [0, 1]$ be dense, $1 \in C$. Then the following system is full:*

- (i) $F_0 = C$ (i.e., a truth constant for each $c \in C$).
- (ii) $F_2 = \{\dot{-}\}$.
- (iii) $F_n = \emptyset$ otherwise.

Corollary 1.6. *The following system is full:*

- (i) $F_1 = \{\neg, \frac{x}{2}\}$.
- (ii) $F_2 = \{\dot{-}\}$.
- (iii) $F_n = \emptyset$ otherwise.

The system appearing in Corollary 1.5 can be viewed as the continuous analogue of the full system of Boolean connectives $\{T, F, \rightarrow\}$ (T and F being truth constants), while that of Corollary 1.6 is reminiscent of $\{\neg, \rightarrow\}$. We will usually use the latter (i.e., $\{\neg, \frac{x}{2}, \dot{-}\}$), which has the advantage of being finite. Note however that for this we need to introduce an additional unary connective $\frac{x}{2}$ which has no counterpart in classical discrete logic.

Remark 1.7. Unlike the discrete case, the family $\{\neg, \vee, \wedge\}$ is not full, and this cannot be remedied by the addition of truth constants. Indeed, it can be verified by induction that every function $f: [0, 1]^n \rightarrow [0, 1]$ constructed from these connectives is 1-Lipschitz in every argument.

This takes care of connectives: any full system would do. We will usually prefer to work with countable systems of connectives, so that countable signatures give countable languages. When making general statements (e.g., the axioms for pseudo-metrics and uniform continuity we give below) it is advisable to use a minimal system of connectives, and we will usually use the one from Corollary 1.6 consisting of $\{\neg, \frac{x}{2}, \dot{-}\}$. On the other hand, when spelling out actual theories, it may be convenient (and legitimate) to admit additional continuous functions from $[0, 1]^n$ to $[0, 1]$ as connectives.

As for quantifiers, the situation is much simpler: we contend that the transition from $\{T, F\}$ to $[0, 1]$ imposes a single pair of quantifiers, or rather, imposes a re-interpretation of the classical quantifiers \forall and \exists (on this point we differ quite significantly from [CK66]). In order to see this, let us look for a construction of the discrete quantifiers \forall and \exists .

Let M be a set, and $\mathcal{R}_M(n)$ be the set of all n -ary relations on M ; we may view each $R \in \mathcal{R}_M(n)$ as a property of n free variables x_0, \dots, x_{n-1} . For every $R \in \mathcal{R}_M(n)$, let $j(R) \in \mathcal{R}_M(n+1)$ be defined as the same relation, with an additional dummy variable x_n . Similarly, for every $R \in \mathcal{R}_M(n+1)$ we have relations $\exists x_n R$ and $\forall x_n R$ in $\mathcal{R}_M(n)$. Then for every $R \in \mathcal{R}_M(n)$ and $Q \in \mathcal{R}_M(n+1)$ the following two properties hold (here “ \models ” means “implies”):

$$\begin{aligned} Q \models j(R) &\iff \exists x_n Q \models R \\ j(R) \models Q &\iff R \models \forall x_n Q. \end{aligned}$$

These properties actually determine the relations $\exists x_n Q$ and $\forall x_n Q$, and can therefore be used as the definition of the semantics of the quantifiers.

Replacing $\{T, F\}$ with $[0, 1]$, let $\mathcal{C}_M(n)$ be the set of all functions from M^n to $[0, 1]$. We define $j: \mathcal{C}_M(n) \rightarrow \mathcal{C}_M(n+1)$ as above, and $\inf_{x_n}, \sup_{x_n}: \mathcal{C}_M(n+1) \rightarrow \mathcal{C}_M(n)$ in the obvious manner. Since we identify T with 0 and F with 1, the relation \models should be replaced with \geq , and we observe that for every $f \in \mathcal{C}_M(n)$ and $g \in \mathcal{C}_M(n+1)$:

$$\begin{aligned} g \geq j(f) &\iff \inf_{x_n} g \geq f \\ j(f) \geq g &\iff f \geq \sup_{x_n} g. \end{aligned}$$

Therefore, as in discrete logic, we will have two quantifiers, whose semantics are defined by the properties above. We will use the symbols \inf and \sup , respectively, to denote the quantifiers, as these best describe their semantics. Make no mistake, though: these are not “new” quantifiers that we have “chosen” for continuous logic, but rather the only possible re-interpretation of the discrete quantifiers \exists and \forall in continuous logic. (Remark 2.11 below will relate our quantifiers to Henson’s sense of approximate satisfaction of quantifiers, further justifying our choice of quantifiers.)

Once we have connectives and quantifiers, we define the set of continuous first order formulae in the usual manner.

Definition 1.8. A *condition* is an expression of the form $\varphi = 0$ where φ is a formula.

A condition is *sentential* if φ is a sentence.

A condition is *universal* if it is of the form $\sup_{\bar{x}} \varphi = 0$ where φ is quantifier-free

If r is a dyadic number then $\varphi \leq r$ is an abbreviation for the condition $\varphi \div r = 0$, and similarly $\varphi \geq r$ for $r \div \varphi = 0$. (Thus $\sup_{\bar{x}} \varphi \leq r$ and $\inf_{\bar{x}} \varphi \geq r$ abbreviate universal conditions.) With some abuse of notation we may use $\varphi \leq r$ for an arbitrary $r \in [0, 1]$ as an abbreviation for the *set* of conditions $\{\varphi \leq r': r' \geq r \text{ dyadic}\}$. We define $\varphi \geq r$ and $\varphi = r$ as abbreviations for sets of conditions similarly.

Notation 1.9. Given a formula φ we will use $\forall \bar{x} \varphi = 0$ as an alternative notation for $\sup_{\bar{x}} \varphi = 0$. While this may be viewed as a mere notational convention, the semantic contents of $\forall \bar{x} (\varphi = 0)$ is indeed equivalent to that of $(\sup_{\bar{x}} \varphi) = 0$ (notice how the

parentheses move, though). Similarly, we may write $\forall \bar{x} \varphi \leq r$ for $\sup_{\bar{x}} \varphi \leq r$ and $\forall \bar{x} \varphi \geq r$ for $\inf_{\bar{x}} \varphi \geq r$.

2. CONTINUOUS STRUCTURES

In classical logic one usually has a distinguished binary predicate symbol $=$, and the logic requires that this symbol always be interpreted as actual equality. The definition we gave for a non-metric continuous signature is the analogue of a discrete signature *without equality*. The analogue of a discrete signature with equality is somewhat trickier, since the symbol taking equality's place need no longer be discrete. Discrete equality always satisfies the equivalence relation axioms:

$$(ER) \quad \begin{aligned} & \forall x \, x = x \\ & \forall xy \, x = y \rightarrow y = x \\ & \forall xyz \, x = y \rightarrow (y = z \rightarrow x = z) \end{aligned}$$

Still within the discrete framework, let us replace the symbol $=$ with the symbol d . Recalling that $T = 0$, $F = 1$ we obtain the discrete metric:

$$d(a, b) = \begin{cases} 0 & a = b \\ 1 & a \neq b \end{cases}$$

Let us now translate ER to continuous logic, recalling that \div is the analogue of implication:

$$(PM) \quad \begin{aligned} & \sup_x d(x, x) = 0 \\ & \sup_{xy} d(x, y) \div d(y, x) = 0 \\ & \sup_{xyz} (d(x, z) \div d(y, z)) \div d(x, y) = 0 \end{aligned}$$

Following Notation 1.9, we can rewrite PM equivalently as the axioms of a pseudo-metric, justifying the use of the symbol d :

$$(PM') \quad \begin{aligned} & \forall x \, d(x, x) = 0 \\ & \forall xy \, d(x, y) = d(y, x) \\ & \forall xyz \, d(x, y) \leq d(x, z) + d(z, y) \end{aligned}$$

By the very definition of equality it is also a congruence relation for all the other symbols, which can be axiomatised as:

$$(CR) \quad \begin{aligned} & \forall \bar{x} \bar{y} zw \, (z = w \rightarrow f(\bar{x}, z, \bar{y}) = f(\bar{x}, w, \bar{y})) \\ & \forall \bar{x} \bar{y} zw \, (z = w \rightarrow (P(\bar{x}, z, \bar{y}) \rightarrow P(\bar{x}, w, \bar{y}))) \end{aligned}$$

Translating CR to continuous logic as we did with ER above would yield axioms saying that every symbol is 1-Lipschitz with respect to d in each of the variables. While we could leave it like this, there is no harm in allowing other moduli of uniform continuity.

Definition 2.1. (i) A *continuity modulus* is a function $\delta: (0, \infty) \rightarrow (0, \infty)$ (for our purposes a domain of $(0, 1]$ would suffice).

(ii) Let $(X_1, d_1), (X_2, d_2)$ be metric spaces. We say that a mapping $f: X_1 \rightarrow X_2$ is *uniformly continuous with respect to a continuity modulus δ* (or just that f *respects δ*) if for all $\varepsilon > 0$ and for all $x, y \in X_1$: $d_1(x, y) < \delta(\varepsilon) \implies d_2(f(x), f(y)) \leq \varepsilon$.

(For a different approach to the definition of continuity moduli see Appendix A.)

Thus, for each n -ary symbol s and each $i < n$ we will fix a continuity modulus $\delta_{s,i}$, and the congruence relation property will be replaced with the requirement that as a function of its i th argument, s should respect $\delta_{s,i}$. As above this can be written in pure continuous logic or be translated to a more readable form:

$$\begin{aligned}
 (\text{UC}_{\mathcal{L}}) \quad & \sup_{x_{<i}, y_{<n-i-1}, z, w} (\delta_{f,i}(\varepsilon) \dot{-} d(z, w)) \wedge (d(f(\bar{x}, z, \bar{y}), f(\bar{x}, w, \bar{y})) \dot{-} \varepsilon) = 0 \\
 & \sup_{x_{<i}, y_{<n-i-1}, z, w} (\delta_{P,i}(\varepsilon) \dot{-} d(z, w)) \wedge ((P(\bar{x}, z, \bar{y}) \dot{-} P(\bar{x}, w, \bar{y})) \dot{-} \varepsilon) = 0 \\
 (\text{UC}'_{\mathcal{L}}) \quad & \forall x_{<i}, y_{<n-i-1}, z, w \left(d(z, w) < \delta_{f,i}(\varepsilon) \rightarrow d(f(\bar{x}, z, \bar{y}), f(\bar{x}, w, \bar{y})) \leq \varepsilon \right) \\
 & \forall x_{<i}, y_{<n-i-1}, z, w \left(d(z, w) < \delta_{P,i}(\varepsilon) \rightarrow (P(\bar{x}, z, \bar{y}) \dot{-} P(\bar{x}, w, \bar{y})) \leq \varepsilon \right)
 \end{aligned}$$

Here $x_{<i}$ denotes the tuple x_0, \dots, x_{i-1} , and similarly for $y_{<n-i-1}$, etc.

Remark 2.2. The axiom scheme $\text{UC}_{\mathcal{L}}$ can be reformulated to mention only constants in some dense set $C \subseteq [0, 1]$ (say rational or dyadic numbers): simply, for every $\varepsilon > 0$, and every $r, q \in C$ such that $r > \varepsilon$ and $q < \delta_{s,i}(\varepsilon)$ (where s is either f or P):

$$\begin{aligned}
 (\text{UC}''_{\mathcal{L}}) \quad & \sup_{\bar{x}, \bar{y}, z, w} (q \dot{-} d(z, w)) \wedge (d(f(\bar{x}, z, \bar{y}), f(\bar{x}, w, \bar{y})) \dot{-} r) = 0 \\
 & \sup_{\bar{x}, \bar{y}, z, w} (q \dot{-} d(z, w)) \wedge ((P(\bar{x}, z, \bar{y}) \dot{-} P(\bar{x}, w, \bar{y})) \dot{-} r) = 0
 \end{aligned}$$

This leads to the following definition:

Definition 2.3. A *(metric) continuous signature* is a non-metric continuous signature along with the following additional data:

- (i) One binary predicate symbol, denoted d , is specified as the distinguished *distance symbol*.
- (ii) For each n -ary symbol s , and for each $i < n$ a continuity modulus $\delta_{s,i}$, called the *uniform continuity modulus* of s with respect to the i th argument.

If we work with a multi-sorted signature then each sort S has its own distinguished distance symbol d_S .

Definition 2.4. Let \mathcal{L} be a continuous signature. A (*continuous*) \mathcal{L} -*pre-structure* is a set M equipped, for every n -ary function symbol $f \in \mathcal{L}$, with a mapping $f^M: M^n \rightarrow M$, and for every n -ary relation symbol $P \in \mathcal{L}$, with a mapping $P^M: M^n \rightarrow [0, 1]$, such that the pseudo-metric and uniform continuity axioms PM, UC $_{\mathcal{L}}$ (or equivalently PM', UC' $_{\mathcal{L}}$) hold.

An \mathcal{L} -*structure* is a pre-structure M in which d^M is a complete metric (i.e., $d^M(a, b) = 0 \implies a = b$ and every Cauchy sequence converges).

The requirement that d^M be a metric corresponds to the requirement that $=^M$ be equality. Completeness, on the other hand, has no analogue in discrete structures, since every discrete metric is trivially complete; still, it turns out to be the right thing to require.

As in classical logic, by writing a term τ as $\tau(\bar{x})$ we mean that all variables occurring in τ appear in \bar{x} . Similarly, for a formula φ the notation $\varphi(\bar{x})$ means that the tuple \bar{x} contains all free variables of φ .

Definition 2.5. Let $\tau(x_{<n})$ be a term, M a \mathcal{L} -pre-structure. The *interpretation* of $\tau(\bar{x})$ in M is a function $\tau^M: M^n \rightarrow M$ defined inductively as follows:

- If $\tau = x_i$ then $\tau^M(\bar{a}) = a_i$.
- If $\tau = f(\sigma_0, \dots, \sigma_{m-1})$ then $\tau^M(\bar{a}) = f^M(\sigma_0^M(\bar{a}), \dots, \sigma_{m-1}^M(\bar{a}))$.

Definition 2.6. Let $\varphi(x_{<n})$ be a term, M a \mathcal{L} -pre-structure. The *interpretation* of $\varphi(\bar{x})$ in M is a function $\varphi^M: M^n \rightarrow [0, 1]$ defined inductively as follows:

- If $\varphi = P(\tau_0, \dots, \tau_{m-1})$ is atomic then $\varphi^M(\bar{a}) = P^M(\tau_0^M(\bar{a}), \dots, \tau_{m-1}^M(\bar{a}))$.
- If $\varphi = \lambda(\psi_0, \dots, \psi_{m-1})$ where λ is a continuous connective then $\varphi^M(\bar{a}) = \lambda(\psi_0^M(\bar{a}), \dots, \psi_{m-1}^M(\bar{a}))$.
- If $\varphi = \inf_y \psi(y, \bar{x})$ $\varphi^M(\bar{a}) = \inf_{b \in M} \psi^M(b, \bar{a})$, and similarly for sup.

Proposition 2.7. *Let M be an \mathcal{L} -pre-structure, $\tau(x_{<n})$ a term, $\varphi(x_{<n})$ a formula. Then the mappings $\tau^M: M^n \rightarrow M$ and $\varphi^M: M^n \rightarrow [0, 1]$ are uniformly continuous in each of their arguments. Moreover, τ^M and φ^M respect uniform continuity moduli which depend only on τ and φ but not on M .*

Proof. In the case of terms, this is just an inductive argument using the fact that a composition of uniformly continuous mappings is uniformly continuous. In the case of formulae one needs two more facts. First, all connectives are uniformly continuous as continuous mappings from a compact space. Second, if $\varphi(\bar{x}) = \inf_y \psi(y, \bar{x})$ then any uniform continuity modulus $\psi(y, \bar{x})$ respects with respect to x_i is also respected by φ .

As the uniform continuity proof above does not depend on M in any way, uniform continuity moduli for terms and formulae can be extracted from the inductive argument. ■_{2.7}

Definition 2.8. (i) Let $s(\bar{x})$ be a condition $\varphi(\bar{x}) = 0$, M an \mathcal{L} -(pre-)structure and $\bar{a} \in M$. We say that s is satisfied by \bar{a} in M , in symbols $M \models s(\bar{a})$, or even $\bar{a} \models s$ (in case the ambient structure M is clear from the context) if $\varphi^M(\bar{a}) = 0$.

- (ii) A set of conditions $\Sigma(\bar{x})$ is *satisfied* by a tuple $\bar{a} \in M$ (again denoted $M \models \Sigma(\bar{a})$ or $\bar{a} \models \Sigma$) if all conditions in Σ are satisfied by \bar{a} in M . This makes sense just as well in case Σ involves infinitely many free variables.
- (iii) Earlier on we defined $\varphi(\bar{x}) \leq r$, $\varphi(\bar{x}) \geq r$, etc., as abbreviations for conditions or sets thereof. Notice that $\varphi^M(\bar{a}) \leq r$ if and only if the corresponding (set of) condition(s) holds for \bar{a} , and similarly for the other abbreviations. We may therefore extend the definition above without ambiguity to satisfaction of such abbreviations.
- (iv) If $\Sigma(\bar{x})$ is a set of conditions and $s(\bar{x})$ a condition, we say that s is a *logical consequence* of Σ , or that Σ *implies* s , in symbols $\Sigma \models s$, if for every M and $\bar{a} \in M$: $M \models \Sigma(\bar{a}) \implies M \models s(\bar{a})$.
- (v) A set of conditions Σ is *satisfiable* if there is a structure M and $\bar{a} \in M$ such that $M \models \Sigma(\bar{a})$. It is *finitely satisfiable* if every finite $\Sigma_0 \subseteq \Sigma$ is satisfiable. We may further say that Σ is *approximately finitely satisfiable* if for every finite subset $\Sigma_0 \subseteq \Sigma$, which we may assume to be of the form $\{\varphi_i = 0 : i < n\}$, and for every $\varepsilon > 0$, the set of conditions $\{\varphi_i \leq \varepsilon : i < n\}$ is satisfiable.

Definition 2.9. A *morphism* of \mathcal{L} -pre-structures is a mapping of the underlying sets which preserves the interpretations of the symbols. It is *elementary* if it preserves the truth values of formulae as well.

Proposition 2.10. *Let M be an \mathcal{L} -pre-structure. Let $M_0 = M/\{d^M(x, y) = 0\}$, and let d_0 denote the metric induced by d^M on M_0 . Let (\hat{M}_0, \hat{d}_0) be the completion of the metric space (M_0, d_0) (which is for all intents and purposes unique).*

Then there exists a unique way to define an \mathcal{L} -structure \hat{M} on the set \hat{M}_0 such that $d^{\hat{M}} = \hat{d}_0$ and the natural mapping $M \rightarrow \hat{M}$ is a morphism. We call \hat{M} the \mathcal{L} -structure associated to M .

Moreover:

- (i) *If N is any other \mathcal{L} -structure, then any morphism $M \rightarrow N$ factors uniquely through \hat{M} .*
- (ii) *The mapping $M \rightarrow \hat{M}$ is elementary.*

Another way of saying this is that the functor $M \mapsto \hat{M}$ is the left adjoint of the forgetful functor from the category of \mathcal{L} -structures to that of \mathcal{L} -pre-structures, and that it sends elementary morphisms to elementary morphisms.

Proof. Straightforward using standard facts about metrics, pseudo-metrics and completions. ■_{2.10}

We say that two formulae are *equivalent*, denoted $\varphi \equiv \psi$ if they define the same functions on every \mathcal{L} -structure (equivalently: on every \mathcal{L} -pre-structure). For example, let $\varphi[t/x]$ denote the free substitution of t for x in φ . Then if y does not appear in φ , then $\sup_x \varphi \equiv \sup_y \varphi[y/x]$ (this is bound substitution of y for x in $\sup_x \varphi$). Similarly, provided that x is not free in φ we have $\varphi \wedge \sup_x \psi \equiv \sup_x \varphi \wedge \psi$, $\varphi \dot{\wedge} \sup_x \psi \equiv \sup_x (\varphi \dot{\wedge} \psi)$, etc.

Using these and similar observations, it is easy to verify that all formulae written using the full system of connectives $\{\neg, \frac{x}{2}, \div\}$ have equivalent prenex forms. In other words, for every such formula φ there is an equivalent formula of the form $\psi = \sup_x \inf_y \sup_z \dots \varphi$, where φ is quantifier-free. The same would hold with any other system of connectives which are monotone in each of their arguments.

Remark 2.11. We can extend Notation 1.9 to all conditions in prenex form, and thereby to all conditions. Consider a condition $\psi \leq r$ (a condition of the basic form $\psi = 0$ is equivalent to $\psi \leq 0$). Write ψ in prenex form, so the condition becomes:

$$\sup_x \left(\inf_y \left(\sup_z \dots \varphi(x, y, z, \dots) \right) \right) \leq r$$

A reader familiar with Henson's logic [HI02] will not find it difficult to verify, by induction on the number of quantifiers, that this is equivalent to the *approximate satisfaction* of:

$$\forall x \left(\exists y \left(\forall z \dots \varphi(x, y, z, \dots) \leq r \right) \right)$$

(Notice how the parentheses move, though.)

We view this as additional evidence to the analogy between the continuous quantifiers \inf and \sup and the Boolean quantifiers \exists and \forall .

Remark 2.12. Unlike the situation in Henson's logic, there are no bounds on the quantifiers as everything in our logic is already assumed to be bounded. For a fuller statement of equivalence between satisfaction in continuous logic and approximate satisfaction in positive bounded logic, see the section on unbounded structure in [Benc]. In particular we show there that under appropriate modifications necessitated by the fact that Henson's logic considers unbounded structures, it has the same power of expression as continuous logic.

Definition 2.13. Let M be an \mathcal{L} -structure. A *formula with parameters in M* is something of the form $\varphi(\bar{x}, \bar{b})$, where $\varphi(\bar{x}, \bar{y})$ is a formula in the tuples of variables \bar{x} and \bar{y} , and $\bar{b} \in M$. Such a formula can also be viewed as an $\mathcal{L}(M)$ -formula, where $\mathcal{L}(M)$ is obtained from \mathcal{L} by adding constant symbols for the elements of M , in which case it may be denoted by $\varphi(\bar{x})$ (i.e., the parameters may be "hidden").

Definition 2.14 (Ultraproducts). Let $\{M_i : i \in I\}$ be \mathcal{L} -structures (or even pre-structures), and \mathcal{U} an ultrafilter on I . Let $N_0 = \prod_i M_i$, and interpret the function and predicate symbols on it as follows:

$$\begin{aligned} f^{N_0}((a_i), (b_i), \dots) &= (f^{M_i}(a_i, b_i, \dots)) \\ P^{N_0}((a_i), (b_i), \dots) &= \lim_{\mathcal{U}} P^{M_i}(a_i, b_i, \dots) \end{aligned}$$

Recall that a sequence in a compact set has a unique limit modulo an ultrafilter: for any open set $U \subseteq [0, 1]$, we have:

$$P^{N_0}((a_i), (b_i), \dots) \in U \iff \{i \in I : P^{M_i}(a_i, b_i, \dots) \in U\} \in \mathcal{U}$$

It is immediate to verify that N_0 satisfies PM and $\text{UC}_{\mathcal{L}}$, so N_0 is an \mathcal{L} -pre-structure. Finally, define $\prod_i M_i/\mathcal{U} = N = \tilde{N}_0$, and call it the *ultraproduct* of $\{M_i : i \in I\}$ modulo \mathcal{U} .

Theorem 2.15 (Łoś's Theorem for continuous logic). *Let $N = \prod_i M_i/\mathcal{U}$ as above. For every tuple $(a_i) \in \prod M_i$ let $[a_i]$ be its image in N . Then for every formula $\varphi(\bar{x})$ we have:*

$$\varphi^N([a_i], [b_i], \dots) = \lim_{\mathcal{U}} \varphi^{M_i}(a_i, b_i, \dots)$$

Proof. By induction on the complexity of φ . See also [CK66, Chapter V]. ■_{2.15}

Corollary 2.16 (Compactness Theorem for continuous first order logic). *Let Σ be a family of conditions (possibly with free variables). Then Σ is satisfiable in an \mathcal{L} -structure if and only if it is finitely so, and furthermore if and only if it is approximately finitely so (see Definition 2.8).*

Proof. The proof is essentially the same as in discrete logic. Replacing free variables with new constant symbols we may assume all conditions are sentential. Enumerate $\Sigma = \{\varphi_i = 0 : i < \lambda\}$, and let $I = \{(w, \varepsilon) : w \subseteq \lambda \text{ is finite and } \varepsilon > 0\}$. For every $(w, \varepsilon) \in I$ choose $M_{w, \varepsilon}$ in which the conditions $\varphi_i \leq \varepsilon$ hold for $i \in w$. For $(w, \varepsilon) \in I$, let $J_{w, \varepsilon} = \{(w', \varepsilon') \in I : w' \supseteq w, \varepsilon' \leq \varepsilon\}$. Then the collection $\mathcal{U}_0 = \{J_{w, \varepsilon} : (w, \varepsilon) \in I\}$ generates a proper filter on I which may be extended to an ultrafilter \mathcal{U} . Let $M = \prod M_{w, \varepsilon}/\mathcal{U}$. By Łoś's Theorem we have $M \models \varphi_i \leq \varepsilon$ for every $i < \lambda$ and $\varepsilon > 0$, so in fact $M \models \varphi_i = 0$. Thus $M \models \Sigma$. ■_{2.16}

Fact 2.17 (Tarski-Vaught Test). *Let M be a structure, $A \subseteq M$ a closed subset. Then the following are equivalent:*

- (i) *The set A is (the domain of) an elementary substructure of M : $A \preceq M$.*
- (ii) *For every formula $\varphi(y, \bar{x})$ and every $\bar{a} \in A$:*

$$\inf\{\varphi(b, \bar{a})^M : b \in M\} = \inf\{\varphi(b, \bar{a})^M : b \in A\}.$$

Proof. One direction is by definition. For the other, we first verify that A is a substructure of M , i.e., closed under the function symbols. Indeed, in order to show that $\bar{a} \in A \implies f(\bar{a}) \in A$ we use the assumption for the formula $d(y, f(\bar{a}))$ and the fact that A is complete. We then proceed to show that $\varphi(\bar{a})^A = \varphi(\bar{a})^M$ for all $\bar{a} \in A$ and formula φ by induction on φ , as in the first order case. ■_{2.17}

When measuring the size of a structure we will use its density character (as a metric space), denoted $\|M\|$, rather than its cardinality.

We leave the following results as an exercise to the reader:

Fact 2.18 (Upward Löwenheim-Skolem). *Let M be a non-compact structure (as a metric space). Then for every cardinal κ there is an elementary extension $N \succeq M$ such that $\|N\| \geq \kappa$.*

Fact 2.19 (Downward Löwenheim-Skolem). *Let M be a structure, $A \subseteq M$ a subset. Then there exists an elementary substructure $N \preceq M$ such that $A \subseteq N$ and $\|N\| \leq |A| + |\mathcal{L}|$.*

Fact 2.20 (Elementary chain). *Let α be an ordinal and $(M_i : i < \alpha)$ an increasing chain of structures such that $i < j < \alpha \implies M_i \preceq M_j$. Let $M = \bigcup_i M_i$. Then $M_i \preceq M$ for all i .*

3. TYPES AND DEFINABLE PREDICATES

We fix a continuous signature \mathcal{L} , as well as a full system of connectives (which might as well be $\{\neg, \frac{x}{2}, \div\}$).

3.1. Spaces of complete types. Recall that a condition (or abbreviation thereof) is something of the form $\varphi = 0$, $\varphi \leq r$ or $\varphi \geq r$ where φ is a formula and $r \in [0, 1]$ is dyadic.

- Definition 3.1.**
- (i) Let M be a structure and $\bar{a} \in M^n$. We define the *type of \bar{a} in M* , denoted $\text{tp}^M(\bar{a})$ (or just $\text{tp}(\bar{a})$ when there is no ambiguity about the ambient structure), as the set of all conditions in $\bar{x} = x_{<n}$ satisfied in M by \bar{a} .
 - (ii) Observe that equality of types $\text{tp}^M(\bar{a}) = \text{tp}^N(\bar{b})$ is equivalent to the elementary equivalence $(M, \bar{a}) \equiv (N, \bar{b})$ (where the tuples are named by new constant symbols). If the ambient structures are clear for the context we may shorten this to $\bar{a} \equiv \bar{b}$.
 - (iii) A *complete n -type*, or just an *n -type*, is a maximal satisfiable set of conditions in the free variables $x_{<n}$. The space of all n -types is denoted S_n .
 - (iv) If p is an n -type and $\bar{a} \in M^n$ is such that $M \models p(\bar{a})$, we say that \bar{a} *realises* p (in M).

We start with a few easy observations whose proof we leave to the reader.

Lemma 3.2. (i) *If \bar{a} is an n -tuple (in some structure) then $\text{tp}(\bar{a})$ is a (complete) n -type.*

Conversely, every n -type can be obtained as the type of an n -tuple.

- (ii) *Let $p \in S_n$ and $\varphi(x_{<n})$ be a formula. Then for every realisation $\bar{a} \models p$, the value of $\varphi(\bar{a})$ depends only on p . It will be denoted φ^p (the value of φ according to p).*

Conversely, the mapping $\varphi \mapsto \varphi^p$, where φ varies over all formulae in the variables $x_{<n}$, determines p .

Every formula $\varphi(x_{<n})$ defines a mapping $p \mapsto \varphi^p: S_n \rightarrow [0, 1]$. With some abuse of notation this function will also be denoted by φ . This is legitimate, since it is clear that two formulae $\varphi(\bar{x})$ and $\psi(\bar{x})$ are equivalent if and only if the functions $\varphi, \psi: S_n \rightarrow [0, 1]$ are equal.

We equip S_n with the minimal topology in which all the functions of this form are continuous, which is sometimes called the *logic topology*.

Lemma 3.3. *With the topology given above, S_n is a compact and Hausdorff space.*

Proof. Let $\mathcal{L}(n)$ be the family of all formulae in the variables $x_{<n}$, and consider the mapping $S_n \rightarrow [0, 1]^{\mathcal{L}(n)}$ given by $p \mapsto (\varphi^p: \varphi \in \mathcal{L}(n))$. As observed in Lemma 3.2 this mapping is injective, and the topology on S_n is the topology induced on its image from the product topology on $[0, 1]^{\mathcal{L}(n)}$, which is Hausdorff. It follows from the Compactness Theorem (Corollary 2.16) that the image is closed in $[0, 1]^{\mathcal{L}(n)}$, so it is compact. ■_{3.3}

Since we are only interested in formulae up to equivalence, it is legitimate to identify a formula with such a mapping. We know that in discrete first order logic the equivalence classes of formulae are in bijection with clopen sets of types, i.e., with continuous mappings from the type space to $\{T, F\}$. Here we cannot claim as much, and it only holds up to uniform approximations:

Proposition 3.4. *A function $f: S_n \rightarrow [0, 1]$ is continuous if and only if it can be uniformly approximated by formulae.*

Proof. For right to left, we know that every formula defines a continuous mapping on the type space, and a uniform limit of continuous mappings is continuous.

Left to right is a consequence of Corollary 1.4 and the fact that formulae separate types. ■_{3.4}

Given a formula $\varphi(x_{<n})$ and $r \in [0, 1]$, we define $[\varphi < r]^{S_n} = \{p \in S_n: \varphi^p < r\}$. We may omit S_n from the notation when it may not cause ambiguity. We could define $[\varphi > r]$ similarly, but as it is equal to $[\neg\varphi < 1 - r]$ this would not introduce any new sets. All sets of this form are clearly open in S_n . Similarly, we define sets of the form $[\varphi \leq r]$, $[\varphi \geq r]$, which are closed. (Since $\varphi \leq r$ is a condition we can also characterise $[\varphi \leq r]$ as the set $\{p: \text{“}\varphi \leq r\text{”} \in p\}$).

Lemma 3.5. *The family of sets of the form $[\varphi < r]$ forms a basis of open sets for the topology on S_n . Equivalently, the family of sets of the form $[\varphi \leq r]$ forms a basis of closed sets.*

Moreover, if U is a neighbourhood of p , we can always find a formula $\varphi(x_{<n})$ such that $p \in [\varphi = 0] \subseteq [\varphi < 1/2] \subseteq U$.

Proof. We prove the moreover part, which clearly implies the rest. Assume that $p \in U \subseteq S_n$ and U is open. By Urysohn’s Lemma there is a continuous function $f: S_n \rightarrow [0, 1]$ such that $f(p) = 0$ and $f|_{U^c} = 1$. We can then find a formula φ_0 such that $|f - \varphi_0| \leq 1/4$. Then the formula $\varphi = \varphi_0 \div 1/4$ would do. ■_{3.5}

3.2. Definable predicates. The discussion above, which is semantic in nature, should convince the reader that uniform limits of formulae are interesting objects, which we would like to call *definable predicates*. But, as formulae are first defined syntactically and only later interpreted as truth value mappings from structures or from type spaces, it will be more convenient later on to first define definable predicates syntactically. Since uniform convergence of the truth values is a semantic notion it cannot be brought into consideration on the syntactic level, so we use instead a trick we call *forced convergence*.

This will be particularly beneficial later on when we need to consider sequences of formulae which sometimes (i.e., in some structures, or with some parameters) converge, and sometimes do not.

Forced conversion is first of all an operation on sequences in $[0, 1]$, always yielding a number in $[0, 1]$. The forced limit coincides with the limit if the sequence converges fast enough. More precisely, if the sequence converges fast enough, we take its limit; otherwise, we find the closest sequence which does converge fast enough, and take *its* limit. Formally:

Definition 3.6. Let $(a_n : n < \omega)$ be a sequence in $[0, 1]$. We define a modified sequence $(a_{\mathcal{F}\text{lim},n} : n < \omega)$ by induction:

$$a_{\mathcal{F}\text{lim},0} = a_0$$

$$a_{\mathcal{F}\text{lim},n+1} = \begin{cases} a_{\mathcal{F}\text{lim},n} + 2^{-n-1} & a_{\mathcal{F}\text{lim},n} + 2^{-n-1} \leq a_{n+1} \\ a_{n+1} & a_{\mathcal{F}\text{lim},n} - 2^{-n-1} \leq a_{n+1} \leq a_{\mathcal{F}\text{lim},n} + 2^{-n-1} \\ a_{\mathcal{F}\text{lim},n} - 2^{-n-1} & a_{\mathcal{F}\text{lim},n} - 2^{-n-1} \geq a_{n+1} \end{cases}$$

The sequence $(a_{\mathcal{F}\text{lim},n} : n < \omega)$ is always a Cauchy sequence, satisfying $n \leq m < \omega \implies |a_{\mathcal{F}\text{lim},n} - a_{\mathcal{F}\text{lim},m}| \leq 2^{-n}$.

We define the *forced limit* of the original sequence $(a_n : n < \omega)$ as:

$$\mathcal{F}\text{lim} a_i \stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} a_{\mathcal{F}\text{lim},n}.$$

Lemma 3.7. *The function $\mathcal{F}\text{lim}: [0, 1]^\omega \rightarrow [0, 1]$ is continuous, and if $(a_n : n < \omega)$ is a sequence such that $|a_n - a_{n+1}| \leq 2^{-n}$ for all n then $\mathcal{F}\text{lim} a_n = \lim a_n$.*

In addition, if $a_n \rightarrow b \in [0, 1]$ fast enough so that $|a_n - b| \leq 2^{-n}$ for all n , then $\mathcal{F}\text{lim} a_n = b$.

Proof. Continuity follows from the fact that if (a_n) and (b_n) are two sequences, and for some m we have $|a_n - b_n| < 2^{-m}$ for all $n \leq m$, then we can show by induction that $|a_{\mathcal{F}\text{lim},n} - b_{\mathcal{F}\text{lim},n}| < 2^{-m}$ for all $n \leq m$, whereby $|\mathcal{F}\text{lim} a_n - \mathcal{F}\text{lim} b_n| < 3 \cdot 2^{-m}$.

For the second condition, one again shows by induction that $|a_{\mathcal{F}\text{lim},n} - b| < 2^{-n}$ whence the conclusion. ■_{3.7}

We may therefore think of $\mathcal{F}\text{lim}$ as an *infinitary* continuous connective.

Definition 3.8. A *definable predicate* is a forced limit of a sequence of formulae, i.e., an (infinite) expression of the form $\mathcal{F}\text{lim}_{n \rightarrow \infty} \varphi_n$.

We say that a variable x is free in $\mathcal{F}\text{lim} \varphi_n$ if it is free in any of the φ_n .

Note that a definable predicate may have infinitely (yet countably) many free variables. In practise, we will mostly consider forced limits of formulae with a fixed (finite) tuple of free variables, but possibly with parameters, so the limit might involve infinitely many parameters. We may write such a definable predicate as $\psi(\bar{x}, B) = \mathcal{F}\text{lim} \varphi_n(\bar{x}, \bar{b}_n)$, and say that $\psi(\bar{x}, B)$ is obtained from $\psi(\bar{x}, Y)$ ($= \mathcal{F}\text{lim} \varphi_n(\bar{x}, \bar{y}_n)$) by substituting the infinite

tuple of parameters $B = \bigcup \bar{b}_n$ in place of the parameter variables $Y = \bigcup \bar{y}_n$. (Another way to think about this is to view parameter variables as constant symbols for which we do not have yet any interpretation in mind: then indeed all the free variables are in \bar{x} .) Later on in Section 5 we will construct canonical parameters for such instances of $\psi(\bar{x}, Y)$. We shall see then that dealing with infinitely many parameters (or parameter variables) is not too difficult as long as the tuple \bar{x} of actually free variable is finite. So from now on, we will only consider definable predicates in finitely many variables.

The semantic interpretation of definable predicates is as expected:

Definition 3.9. Let $\psi(\bar{x}) = \mathcal{F}\text{lim } \varphi_n(\bar{x})$ be a definable predicate in $\bar{x} = x_{< m}$. Then for every structure M and $\bar{a} \in M^m$, we define $\psi^M(\bar{a}) = \mathcal{F}\text{lim } \varphi_n^M(\bar{a})$. Similarly, for $p \in S_m$ we define $\psi^p = \mathcal{F}\text{lim } \varphi_n^p$.

As with formulae, if $\bar{a} \in M$ realises $p \in S_m$, then $\psi^M(\bar{a}) = \psi^p$. Therefore, two definable predicates $\psi(\bar{x})$ and $\chi(\bar{x})$ are equivalent (i.e., $\psi^M = \chi^M$ for all M) if and only if they are equal as functions of types.

In the same manner we tend to identify two equivalent formulae, we will tend to identify two equivalent definable predicates. Thus, rather than viewing definable as infinite syntactic objects we will rather view them as *semantic* objects, i.e., as continuous functions on the corresponding type space. If one insists on making a terminological distinction then the syntactic notion of a forced limit of formulae could be called a *limit formula*.

Proposition 3.10. *The continuous functions $S_m \rightarrow [0, 1]$ are precisely those given by definable predicates in m free variables.*

Proof. Let $\psi(x_{< m}) = \mathcal{F}\text{lim } \varphi_n$ be a definable predicate. Then the mapping $p \mapsto (\varphi_n^p : n < \omega)$ is continuous, and composed with the continuous mapping $\mathcal{F}\text{lim} : [0, 1]^\omega \rightarrow [0, 1]$ we get $\psi : S_m \rightarrow [0, 1]$, which is therefore continuous.

Conversely, let $f : S_m \rightarrow [0, 1]$ be continuous. Then by Proposition 3.4 f can be uniformly approximated by formulae: for every n we can find $\varphi_n(x_{< m})$ such that $|f - \varphi_n| \leq 2^{-n}$. Then $f = \lim \varphi_n = \mathcal{F}\text{lim } \varphi_n$. ■_{3.10}

Viewed as continuous functions of types, it is clear that the family of definable predicates is closed under all continuous combinations, finitary or not (so in particular under the infinitary continuous connective $\mathcal{F}\text{lim}$), as well as under uniform limits. Finally, let $\psi(\bar{x}, y)$ be a continuous predicate, and find formulae $\varphi_n(\bar{x}, y)$ that converge to ψ uniformly on S_{m+1} . Then the sequence $\sup_y \varphi_n$ converges uniformly on S_m to a definable predicate which will be denoted $\sup_y \psi$ (this is since we always have $|\sup_y \varphi_n - \sup_y \varphi_k| \leq |\varphi_n - \varphi_k|$). Furthermore, it follows from the uniform convergence $\varphi_n \rightarrow \psi$ that for every structure M and $\bar{a} \in M$ we have $(\sup_y \psi)^M(\bar{a}) = \sup\{\psi^M(\bar{a}, b) : b \in M\}$. The same holds with \inf , and we conclude that the family of definable predicates is closed under continuous quantification. Also, since all formulae are uniformly continuous with respect to the metric d ,

so are their uniform limits. This means, for example, that every definable predicate can be added to the language as a new actual predicate symbol.

We conclude with a nice consequence of the forced limit construction, which might not have been obvious if we had merely defined definable predicates as uniform limits of formulae.

Lemma 3.11. *Let M be a structure, and let $(\varphi_n(\bar{x}): n < \omega)$ be a sequence of formulae, or even definable predicates, such that the sequence of functions $(\varphi_n^M: n < \omega)$ converges uniformly to some $\xi: M^m \rightarrow [0, 1]$ (but need not converge at all for any other structure instead of M).*

Then there is a definable predicate $\psi(\bar{x})$ such that $\psi^M = \xi = \lim \varphi_n^M$ (i.e., ξ is definable in M).

Proof. Up to passing to a sub-sequence, we may assume that $|\varphi_n^M - \xi| \leq 2^{-n}$, so $\xi = \mathcal{F}\lim \varphi_n^M = (\mathcal{F}\lim \varphi_n)^M$. would do. ■_{3.11}

3.3. Partial types.

Definition 3.12. A *partial type* is a set of conditions, usually in a finite tuple of variables. (Thus every complete type is in particular a partial type.)

For a partial type $p(x_{<n})$ we define:

$$[p]^{S_n} = \bigcap_{s \in p} [s]^{S_n} = \{q \in S_n : p \subseteq q\}.$$

Since the sets of the form $[s]$ (where s is a condition in $x_{<n}$) form a basis of closed sets for the topology on S_n , the sets of the form $[p]$ (where $p(x_{<n})$ is a partial type) are precisely the closed subsets of S_n .

For every $n, m < \omega$ we have a natural restriction mapping $\pi: S_{n+m} \rightarrow S_n$. This mapping is continuous, and therefore closed (as is every continuous mapping between compact Hausdorff spaces). Let $p(x_{<n}, y_{<m})$ be a partial type, defining a closed subset $[p] \subseteq S_{n+m}$. Then $\pi([p]) \subseteq S_n$ is closed as well, and therefore of the form $[q(x_{<n})]$.

Definition 3.13. Let $p(\bar{x}, \bar{y})$ be a partial type. We define $\exists \bar{y} p(\bar{x}, \bar{y})$ to be any partial type $q(\bar{x})$ (say the maximal one) satisfying $[q]^{S_n} = \pi([p]^{S_{n+m}})$.

By the previous argument we have:

Fact 3.14. *For every partial type $p(\bar{x}, \bar{y})$, a partial type $\exists \bar{y} p(\bar{x}, \bar{y})$ exists. Moreover, if M is structure and $\bar{a} \in M$ then $M \models \exists \bar{y} p(\bar{a}, \bar{y})$ if and only if there is $N \succeq M$ and $\bar{b} \in N$ such that $N \models p(\bar{a}, \bar{b})$. In case M is ω -saturated (i.e., if every 1-type over a finite tuple in M is realised in M) then such \bar{b} exists in M .*

4. THEORIES

Definition 4.1. A *theory* is a set of sentential conditions, i.e., things of the form $\varphi = 0$, where φ is a sentence.

Thus in some sense a theory is an “ideal”. Since every condition of the form $\varphi \leq r$, $\varphi \geq r$ or $\varphi = r$ is logically equivalent to one of the form $\varphi' = 0$ (if r is dyadic) or to a set of such conditions (for any $r \in [0, 1]$), we may allow ourselves conditions of this form as well. It should be noted however that most theories are naturally axiomatised by conditions of the form $\varphi = 0$.

Definition 4.2. Let T be a theory. A *(pre-)model* of T is an \mathcal{L} -(pre-)structure M in which T is satisfied.

The notions of satisfaction and satisfiability of sets of conditions from Definition 2.8 apply in the special case of a theory (a set of conditions without free variables). In particular, a theory is satisfiable if and only if it has a model, and by Proposition 2.10, this is the same as having a pre-model.

A theory is *complete* if it is satisfiable and maximal as such (i.e., if it is a complete 0-type), or at least if its set of logical consequences is. The complete theories are precisely those obtained as theories of structures:

$$\begin{aligned} \text{Th}(M) &= \{\varphi = 0 : \varphi \text{ an } \mathcal{L}\text{-sentence and } \varphi^M = 0\} \\ &\equiv \{\varphi = \varphi^M : \varphi \text{ an } \mathcal{L}\text{-sentence}\}. \end{aligned}$$

(In the second line we interpret $\varphi = r$ as an abbreviation for a set of conditions as described earlier.)

4.1. Some examples of theories. Using the metric, any equational theory (in the ordinary sense) can be expressed as a theory, just replacing $x = y$ with $d(x, y) = 0$.

Example 4.3. Consider probability algebras (i.e., measure algebras, as discussed for example in [Fre04], with total measure 1). The language is $\mathcal{L} = \{0, 1, ^c, \wedge, \vee, \mu\}$, with all continuity moduli being the identity. The theory of probability algebras, denoted PrA , consists of the following axioms:

$$\begin{aligned} &\langle \text{equational axioms of Boolean algebras} \rangle \\ &\mu(1) = 1 \\ &\mu(0) = 0 \\ &\forall xy (\mu(x) + \mu(y) = \mu(x \vee y) + \mu(x \wedge y)) \\ &\forall xy (d(x, y) = \mu((x \wedge y^c) \vee (y \wedge x^c))) \end{aligned}$$

The last two axioms are to be understood in the sense of Notation 1.9. Thus $\forall xy \mu(x) + \mu(y) = \mu(x \vee y) + \mu(x \wedge y)$ should be understood as $\sup_{xy} \left| \frac{\mu(x) + \mu(y)}{2} - \frac{\mu(x \vee y) + \mu(x \wedge y)}{2} \right| = 0$, etc. In the last expression, division by two is necessary to keep the range in $[0, 1]$. As we

get used to this we will tend to omit it and simply write $\sup_{xy} |\mu(x) + \mu(y) - \mu(x \vee y) - \mu(x \wedge y)| = 0$.

Note that we cannot express $\mu(x) = 0 \rightarrow x = 0$, but we do not have to either: if M is a model, $a \in M$, and $\mu^M(a) = 0$, then the axioms imply that $d^M(a, 0^M) = 0$, whereby $a = 0^M$.

The model companion of PrA is APA , the theory of atomless probability algebras, which contains in addition the following sentence:

$$\forall x \exists y \left(\mu(y \wedge x) = \frac{\mu(x)}{2} \right),$$

Following Remark 2.11 we can express this by:

$$\sup_x \inf_y \left| \mu(y \wedge x) - \frac{\mu(x)}{2} \right| = 0$$

(We leave it to the reader to verify that this sentential condition does indeed if and only if the probability algebra is atomless.)

Example 4.4 (Convex spaces). Let us now consider a signature \mathcal{L}_{cvx} consisting of binary function symbols $c_\lambda(x, y)$ for all dyadic numbers $\lambda \in [0, 1]$, which are all, say, 1-Lipschitz in both arguments. Let T_{cvx} consist of:

$$\begin{aligned} (\forall xyz) d(z, c_\lambda(x, y)) &\leq \lambda d(z, x) + (1 - \lambda) d(z, y) \\ (\forall xyz) c_{\lambda_0 + \lambda_1} \left(c_{\frac{\lambda_0}{\lambda_0 + \lambda_1}}(x, y), z \right) &= c_{\lambda_1 + \lambda_2} \left(c_{\frac{\lambda_1}{\lambda_1 + \lambda_2}}(y, z), x \right) \quad \lambda_0 + \lambda_1 + \lambda_2 = 1 \\ (\forall xyz) d(c_\lambda(x, z), c_\lambda(y, z)) &= \lambda d(x, y). \end{aligned}$$

By [Mac73], the models of T_{cvx} are precisely the closed convex subspaces of Banach spaces of diameter ≤ 1 , equipped with the convex combination operations $c_\lambda(x, y) = \lambda x + (1 - \lambda)y$. We may in fact restrict to a single function symbol $c_{1/2}(x, y)$, since every other convex combination operator with dyadic coefficients can be expressed using this single operator. Since our structures are by definition complete, dyadic convex combinations suffice.

Example 4.5. Let us continue with the previous example. We may slightly modify our logic allowing the distance symbol to have values in the compact interval $[0, 2]$, so now models of T_{cvx} are convex sets of diameter ≤ 2 . Let us add a constant symbol 0 , introduce $\|x\|$ as shorthand for $d(x, 0)$ and λx as shorthand for $c_\lambda(x, 0)$, and add the axioms:

$$\begin{aligned} (\forall x) \|x\| &\leq 1 \\ (\forall x \exists y) d(x, y/2) \wedge (1/2 \div \|x\|) &= 0 \end{aligned}$$

The first axiom tells us that our model is a convex subset of the unit ball of the ambient Banach space. The second tells us that our model is precisely the unit ball: if $\|x\| \leq 1/2$ then $2x$ exists. One can add more structure on top of this, for example:

- (i) Multiplication by i , rendering the ambient space a complex Banach space.

- (ii) Function symbols \vee and \wedge , rendering the ambient space a Banach lattice. In order to make sure we remain inside the unit ball, we actually need to add $(x, y) \mapsto (x \wedge y)/2$ and $(x, y) \mapsto (x \vee y)/2$ rather than \vee and \wedge . In particular the model-theoretic study of independence in L_p Banach lattices carried out in [BBH] fits in this setting.

Alternatively, in order to study a Banach space E one could introduce a multi-sorted structure where there is a sort E_n for each closed ball around 0 in E of radius $n < \omega$. On each sort, all the predicate symbols have values in a compact interval, and operations such as $+$ go from $E_n \times E_m$ to E_{n+m} . However, since every sort E_n is isomorphic to E_1 up to rescaling, this boils down to the single-sorted approach described above (and in particular, re-scaled addition is indeed the convex combination operation).

It can be shown that either approach (single unit ball sort or a sort for each radius) has the same power of expression as Henson's logic, i.e., the translation from a Banach space structure à la Henson to a unit ball structure (in an appropriate signature) preserves such notions as elementary classes and extensions, type-definable subsets of the unit ball, etc. This should be intuitively clear from Remark 2.11. Thus, continuous first order logic is indeed as good a setting for the study of such properties as stability and independence in Banach space structures as Henson's logic, but as we show later it is much better adapted for such study.

The reader may find in [Benc] a treatment of unbounded continuous signatures and structures, an approach much closer in spirit to Henson's treatment of Banach space structures. It is proved there that approximate satisfaction of positive bounded formulae (which makes sense in any unbounded continuous signature) has the same power of expression as satisfaction of conditions of continuous first order logic. The equivalence mentioned in the previous paragraph follows. In addition, the single point compactification method defined there turns such unbounded structures into bounded structures as studied here without chopping them into pieces as above (and again, preserving such notions as elementarity and definability).

4.2. Type spaces of a theory. If T is a theory, we define its type spaces as in classical first order logic:

$$S_n(T) = \{p \in S_n : T \subseteq p\} = \{\text{tp}^M(\bar{a}) : M \models T \text{ and } \bar{a} \in M^n\}.$$

This is a closed subspace of S_n , and therefore compact and Hausdorff in the induced topology. We define $[\varphi \leq r]^{S_n(T)} = [\varphi \leq r]^{S_n} \cap S_n(T)$, and similarly for $\varphi = r$, etc. As before, we may omit $S_n(T)$ if the ambient type space in question is clear.

If M is a structure and $A \subseteq M$, we define $\mathcal{L}(A)$ by adding constant symbols for the elements of A and identify M with its natural expansion to $\mathcal{L}(A)$. We define $T(A) = \text{Th}_{\mathcal{L}(A)}(M)$ and $S_n(A) = S_n(T(A))$, the latter being the space of n -types over A .

This definition allows a convenient re-statement of the Tarski-Vaught Test:

Fact 4.6 (Topological Tarski-Vaught Test). *Let M be a structure, $A \subseteq M$ a closed subset. Then the following are equivalent:*

- (i) *The set A is (the domain of) an elementary substructure of M : $A \preceq M$.*
- (ii) *The set of realised types $\{\text{tp}^M(a/A) : a \in A\}$ is dense in $S_1(A)$.*

Proof. We use Fact 2.17. Assume first that the set of realised types is dense. Let $\varphi(y, \bar{a}) \in \mathcal{L}(A)$, $r = \inf_y \varphi(y, \bar{a})^M \in [0, 1]$. Then for every $\varepsilon > 0$ the set $[\varphi(y, \bar{a}) < r + \varepsilon] \subseteq S_1(A)$ is open and non-empty, so there is $b \in A$ such that $\varphi(b, \bar{a})^M < r + \varepsilon$, whence condition (ii) of Fact 2.17. Conversely, since the sets of the form $[\varphi(y, \bar{a}) < r]$ (with $\bar{a} \in A$) form a basis of open sets for $S_1(A)$, condition (ii) of Fact 2.17 implies the realised types are dense. ■_{4.6}

By previous results, a uniform limit (or forced limit) of formulae with parameters in A is the same (as functions on M , or on any elementary extension of M) as a continuous mapping $\varphi : S_n(A) \rightarrow [0, 1]$. Such a definable predicate with parameters in A is called a definable predicate *over* A , or an *A -definable* predicate.

We define κ -saturated and (strongly) κ -homogeneous structures as usual, and show that every complete theory admits a monster model, i.e., a κ -saturated and strongly κ -homogeneous model for some κ which is far larger than the cardinality of any other set under consideration. It will be convenient to assume that there is always an ambient monster model: every set of parameters we consider is a subset of a monster model, and every model we consider is an elementary substructure thereof. (Even when considering an incomplete theory, each model of the theory embeds in a monster model of its complete theory.)

If \mathfrak{M} is a monster model and $A \subseteq \mathfrak{M}$ a (small) set, we define

$$\text{Aut}(\mathfrak{M}/A) = \{f \in \text{Aut}(\mathfrak{M}) : f \upharpoonright_A = \text{id}_A\}.$$

A definable predicate with parameters (in \mathfrak{M}) is *A -invariant* if it is fixed by all $f \in \text{Aut}(\mathfrak{M}/A)$.

All type spaces we will consider in this paper are quotient spaces of $S_n(\mathfrak{M})$, where \mathfrak{M} is a fixed monster model of T , or of $S_n(T)$, which can be obtained using the following general fact:

Fact 4.7. *Let X be a compact Hausdorff space and $\mathfrak{A} \subseteq C(X, [0, 1])$ any sub-family of functions. Define an equivalence relation on X by $x \sim y$ if $f(x) = f(y)$ for all $f \in \mathfrak{A}$, and let $Y = X/\sim$. Then:*

- (i) *Every $f \in \mathfrak{A}$ factors uniquely through the quotient mapping $\pi : X \rightarrow Y$ as $f = f_Y \circ \pi$.*
- (ii) *The quotient topology on Y is precisely the minimal topology under which every such f_Y is continuous.*
- (iii) *This topology is compact and Hausdorff.*

Conversely, let Y be a compact Hausdorff space and $\pi: X \rightarrow Y$ a continuous projection. Then Y is a quotient space of X and can be obtained as above using $\mathfrak{A} = \{f \circ \pi: f \in C(Y, [0, 1])\}$.

Proof. The first item is by construction. Let \mathcal{T}_1 be the quotient topology on Y and \mathcal{T}_2 the minimal topology in which every f_Y is continuous. Then \mathcal{T}_1 is compact as a quotient of a compact topology. If $y_1, y_2 \in Y$ are distinct then there is a function f_Y separating them, whereby \mathcal{T}_2 is Hausdorff. Finally, let $V \subseteq [0, 1]$ be open and $f \in \mathfrak{A}$, so $U = f_Y^{-1}(V) \subseteq Y$ is a pre-basic open set of \mathcal{T}_2 . Then $\pi^{-1}(U) = f^{-1}(V) \subseteq X$ is open, whereby $U \in \mathcal{T}_1$. Thus \mathcal{T}_1 refines \mathcal{T}_2 . Since \mathcal{T}_1 is compact and \mathcal{T}_2 is Hausdorff they must coincide.

For the converse, the space $Y' = X/\sim$ constructed in this manner can be identified with Y . The original topology on Y refines the quotient topology by the second item, and as above they must coincide. $\blacksquare_{4.7}$

For example, let \mathfrak{M} be a monster model and $A \subseteq \mathfrak{M}$ a set. Then there is a natural projection $\pi: S_n(\mathfrak{M}) \rightarrow S_n(A)$ restricting from $\mathcal{L}(\mathfrak{M})$ to $\mathcal{L}(A)$, and let:

$$\begin{aligned}\mathfrak{A} &= \{\varphi \circ \pi: \varphi \in C(S_n(A), [0, 1])\}, \\ \mathfrak{A}' &= \{\varphi \in C(S_n(\mathfrak{M}), [0, 1]): \varphi \text{ is } A\text{-invariant}\}.\end{aligned}$$

Then clearly $\mathfrak{A} \subseteq \mathfrak{A}'$. On the other hand, \mathfrak{A} and \mathfrak{A}' separate the same types, so by Fact 4.7 every $f \in \mathfrak{A}'$ factors through $S_n(A)$ and $\mathfrak{A}' = \mathfrak{A}$. In other words we've shown:

Lemma 4.8. *Let A be a set (in the monster model) and let φ be an A -invariant definable predicate with parameters possibly outside A . Then φ is (equivalent to) an A -definable predicate A .*

Let us adapt the notions of algebraicity and algebraic closure to continuous logic:

Lemma 4.9. *Let A be a set of parameters and $p(x) \in S_1(A)$. Then the following are equivalent:*

- (i) *For every $\varepsilon > 0$ there is a condition $(\varphi_\varepsilon(x) = 0) \in p$ (with parameters in A) and $n_\varepsilon < \omega$ such that for every sequence $(a_i: i \leq n_\varepsilon)$, if $\varphi_\varepsilon(x_i) < 1/2$ for all $i \leq n_\varepsilon$ then $d(x_i, x_j) \leq \varepsilon$ for all $i < j \leq n_\varepsilon$.*
- (ii) *Every model containing A contains all realisations of p .*
- (iii) *Every indiscernible sequence in p is constant.*
- (iv) *There does not exist an infinite sequence $(a_i: i < \omega)$ of realisations of p such that $\inf\{d(a_i, a_j): i < j < \omega\} > 0$.*
- (v) *The set of realisations of p is compact.*

Proof. (i) \implies (ii). We may assume that for the choice of φ_ε , the number n_ε is minimal: we can therefore find in the universal domain elements $a_{<n_\varepsilon}$ such that $\varphi_\varepsilon(a_i) < 1/2$ for all $i < n_\varepsilon$ and yet $d(a_i, a_j) > \varepsilon$ for all $i < j < n_\varepsilon$. Then we have (in the universal domain):

$$\sup_{x < n_\varepsilon} \left(\bigwedge_{i < n_\varepsilon} \left(\frac{1}{2} \dot{-} \varphi_\varepsilon(x_i) \right) \wedge \bigwedge_{i < j < n_\varepsilon} (d(x_i, x_j) \dot{-} \varepsilon) \right) > 0.$$

Assume that $A \subseteq M$. Then the same holds in M , and we may therefore find $a_{<n_\varepsilon}$ as above inside M . Assume also that $a \models p$. Then by assumption there is some $i < n_\varepsilon$ such that $d(a, a_i) \leq \varepsilon$. Since this holds for all $\varepsilon > 0$ and M is complete we must have $a \in M$.

(ii) \implies (iii). By Downward Löwenheim-Skolem.

(iii) \implies (iv). By (compactness and) Ramsey's theorem.

(iv) \implies (i). Since the set $\{d(x_i, x_j) \geq \varepsilon : i < j < \omega\} \cup \bigcup_{i < \omega} p(x_i)$ must be inconsistent.

(iv) \iff (v). Condition (iv) is equivalent to saying that the set of realisations of p is totally bounded. Since it is in addition automatically complete (a limit of realisations of p is a realisation of p), this is the same as saying it is compact. $\blacksquare_{4.9}$

Definition 4.10. If $p \in S(A)$ satisfies any of the equivalent properties in Lemma 4.9 then it is called *algebraic*. We say that a is *algebraic* over A if $\text{tp}(a/A)$ is algebraic. We define the *algebraic closure* of A , denoted $\text{acl}(A)$, as the set of all algebraic elements over A . By Lemma 4.9 if $A \subseteq M$ then $\text{acl}(A) \subseteq M$ as well, so $\text{acl}(A)$ is the same in every model containing A . If $A = \text{acl}(A)$ then we say that A is *algebraically closed*.

4.3. The metric on types of a complete theory. Let T be complete. Since any two n -types are realised inside the monster model we can define for every $p, q \in S_n(T)$:

$$d(p, q) = \inf\{d(\bar{a}, \bar{b}) : \bar{a} \models p \text{ and } \bar{b} \models q\}$$

Here $d(a_{<n}, b_{<n}) = \bigvee_{i < n} d(a_i, b_i)$. It is trivial to verify this is a pseudo-metric. By compactness the infimum is attained, so it is in fact a metric: $d(p, q) = 0 \implies p = q$. Note also that we can construct $S_n(T)$ as a set as $\mathfrak{M}^n / \text{Aut}(\mathfrak{M})$, where \mathfrak{M} is the monster model and we divide it by the action of its automorphism group, in which case the distance between types in the one induced from \mathfrak{M}^n . (For this purpose, any ω -saturated and strongly homogeneous model of T would serve just as well.)

The metric on $S_n(T)$ refines the logic topology. Indeed, let $p \in S_n(T)$ and U a neighbourhood of p . Then there is a formula φ such that $p \in [\varphi = 0] \subseteq [\varphi < 1/2] \subseteq U$. The uniform continuity of φ implies the existence of $\delta > 0$ such that $d(p, q) \leq \delta \implies \varphi^q < 1/2$, so U contains a metric neighbourhood of p . By a theorem of Henson (for Banach space structures in positive bounded logic, but it boils down to the same thing), for a complete countable theory T , the metric on $S_n(T)$ coincides with the logic topology for all n if and only if T is separably categorical, i.e., if and only if it has a unique separable model up to isomorphism.

Also, if $F \subseteq S_n(T)$ is closed and $\varepsilon > 0$, then so is the closed ε -neighbourhood of F :

$$F^\varepsilon = \{p \in S_n(T) : d(p, F) \leq \varepsilon\}.$$

Indeed, since the set F is closed in can be written as $[p(\bar{x})]$ where p is some partial type, and then

$$F^\varepsilon = [\exists \bar{y} (d(\bar{x}, \bar{y}) \leq \varepsilon \wedge p(\bar{y}))].$$

This leads us to the following definition which will turn out to be useful later on:

Definition 4.11. A compact *topometric* space is a triplet $\langle X, \mathcal{T}, d \rangle$, where \mathcal{T} is a compact Hausdorff topology and d a metric on X , satisfying:

- (i) The metric refines the topology.
- (ii) For every closed $F \subseteq X$ and $\varepsilon > 0$, the closed ε -neighbourhood of F is closed in X as well.

When dealing with topometric spaces some care must be taken about the language. We will follow the convention that whenever we use terms which come from the realm of general topology (such as compactness, closed and open sets, etc.) we refer to the topology. When wish to refer to the metric, we will use terminology that clearly comes from the realm of metric spaces. When there may be ambiguity, we will say explicitly to which part we are referring.

We may therefore sum up the previous observations as:

Fact 4.12. *The type space $S_n(T)$ is a compact topometric space.*

We will come back to topometric spaces later. Let us now conclude with a small fact about them:

Lemma 4.13. *Let X be a compact topometric space (so by the terminological convention above, we mean to say that the topology is compact). Then it is complete (as a metric space).*

Proof. Let $(x_n : n < \omega)$ be a Cauchy sequence in X . We may assume that $d(x_n, x_{n+1}) \leq 2^{-n-1}$ for all n . For each n the set $\{x_n\}^{2^{-n}}$, (the closed 2^{-n} -ball around x_n) is closed in the topology, and $(\{x_n\}^{2^{-n}} : n < \omega)$ is a decreasing sequence of non-empty closed sets. By compactness there is some x in the intersection, and clearly $x_n \rightarrow x$ in the metric. ■_{4.13}

4.4. Quantifier elimination.

Definition 4.14. A *quantifier-free definable predicate* is a definable predicate defined by a forced limit of quantifier-free formulae.

A theory has *quantifier elimination* if every formula can be uniformly approximated over all models of T by quantifier-free formulae, i.e., if every formula is equal in models of T to a quantifier-free definable predicate.

(In order to avoid pathologies when there are no constant symbols in \mathcal{L} , we must allow that if φ is a formula without free variables, the quantifier-free definable predicate may have a free variable.)

We introduce the following criterion for quantifier elimination, analogous to the classical back-and-forth criterion:

Definition 4.15. We say that a theory T has the *back-and-forth* property if for every two ω -saturated models $M, N \models T$, non-empty tuples $\bar{a} \in M^n$ and $\bar{b} \in N^n$, and singleton $c \in M$, if \bar{a} and \bar{b} have the same quantifier-free type (i.e., $\varphi^M(\bar{a}) = \varphi^N(\bar{b})$ for all quantifier-free φ) then there is $d \in N$ such that \bar{a}, c and \bar{b}, d have the same quantifier-free type.

Theorem 4.16. *The following are equivalent for any continuous theory T (not necessarily complete):*

- (i) *The theory T admits quantifier elimination.*
- (ii) *The theory T has the back-and-forth property.*

Proof. Assume first that T admits quantifier elimination. Then under the assumptions we have $\bar{a} \equiv \bar{b}$. Let $p(x, \bar{y}) = \text{tp}(c, \bar{a})$. Then $p(x, \bar{b})$ is consistent and is realised by some $d \in N$ by ω -saturation.

For the converse we introduce an auxiliary definition: The inf_y -type of a tuple $\bar{a} \in M$ is given by the function $\varphi(\bar{x}) \mapsto \varphi^M(\bar{a})$, where φ varies over all the formulae of the form $\text{inf}_y \psi(y, \bar{x})$, ψ quantifier-free. We define $S_n^{\text{inf}_y}(T)$ as the set of all inf_y -types of n -tuples in models of T . This is a quotient of $S_n(T)$, and we equip it with the quotient topology, which is clearly compact and Hausdorff.

We claim first that if $M, N \models T$ and $\bar{a} \in M^n, \bar{b} \in N^n$ have the same quantifier-free type then they have the same inf_y -type. Since we may embed M and N elementarily in more saturated models, we may assume both are ω -saturated. Assume $\text{inf}_y \psi(y, \bar{a}) = r$. Then there are $c_m \in M$ such that $\psi(c_m, \bar{a}) \leq r + 2^{-m}$, and by ω -saturation there is $c \in M$ such that $\psi(c, \bar{a}) = r$. Therefore there is $d \in N$ such that $\psi(d, \bar{b}) = r$, whereby $\text{inf}_y \psi^M(y, \bar{a}) \geq \text{inf}_y \psi^N(y, \bar{b})$. By a symmetric argument we have equality.

We conclude that the quantifier-free formulae separate points in $S_n^{\text{inf}_y}(T)$. Since quantifier-free formulae form a family of continuous functions on $S_n^{\text{inf}_y}(T)$ which is closed under continuous connectives, the quantifier-free formulae are dense in $C(S_n^{\text{inf}_y}(T), [0, 1])$ by Corollary 1.4. In particular, every inf_y -formula can be uniformly approximated by quantifier-free formulae, and by induction on the structure of the formula, every formula can be thus approximated (on models of T , or equivalently on $S_n(T)$). ■_{4.16}

Corollary 4.17. *The theory of atomless probability algebras (described above) is complete and has quantifier elimination.*

Proof. The back-and-forth property between complete atomless probability algebras is immediate, and does not require ω -saturation (or rather, as it turns out, all complete atomless probability algebras are ω -saturated). Indeed, two n -tuples $a_{<n} \in M$ and $b_{<n} \in N$ have the same quantifier-free type ($\bar{a} \equiv^{qf} \bar{b}$) if and only if they generate isomorphic algebras. Letting $a^0 = a$, $a^1 = a^c$ and $\bar{a}^{\bar{t}} = \bigwedge_{i < n} a_i^{t_i}$ for $\bar{t} \in \{0, 1\}^n$ we have $\bar{a} \equiv^{qf} \bar{b}$ if and only if $\mu(\bar{a}^{\bar{t}}) = \mu(\bar{b}^{\bar{t}})$ for all $\bar{t} \in \{0, 1\}^n$. If $c \in M$ is any other singleton use atomlessness to find for each $\bar{t} \in \{0, 1\}^n$ an event $d_{\bar{t}} \leq \bar{b}^{\bar{t}}$ such that $\mu(d_{\bar{t}}) = \mu(c \wedge \bar{a}^{\bar{t}})$. Then $d = \bigvee d_{\bar{t}}$ will do.

Also, every two probability algebras can be embedded in a third one, which can be further embedded in a complete atomless one, whence the completeness. ■_{4.17}

Definition 4.18. A theory T is *model complete* if for every $M, N \models T$, $M \subseteq N \implies M \prec N$.

We leave the following as an exercise to the reader (see [Benb] for a complete proof):

Proposition 4.19. *A theory T is model complete if and only if every formula can be uniformly approximated on $S_n(T)$ by formulae of the form $\inf_{\bar{y}} \varphi(\bar{x}, \bar{y})$, where φ is quantifier-free.*

4.5. Continuous first order logic and open Hausdorff cats. We now show the equivalence between the framework of continuous first order logic and that of (metric) open Hausdorff cats. For this we assume familiarity with the latter framework, as exposed in [Ben05]. The reader who is not familiar with open Hausdorff cats may safely skip this part.

To every theory T we associate its type-space functor $S(T)$ in the usual manner. For every $n < \omega$ we defined $S_n(T)$ above. If $m, n < \omega$ and $f: n \rightarrow m$ is any mapping, we define $f^*: S_m(T) \rightarrow S_n(T)$ by $f^*(p(x_{< m})) = \{\varphi(x_{< n}) = \varphi(x_{f(0)}, \dots, x_{f(n-1)})^p : \varphi(x_{< n}) \in \mathcal{L}\}$, i.e., $f^*: \text{tp}(a_0, \dots, a_{m-1}) \mapsto \text{tp}(a_{f(0)}, \dots, a_{f(n-1)})$.

Fact 4.20. *Let T be a continuous first order theory. Then its type-space functor $S(T)$ is an open, compact and Hausdorff type-space functor in the sense of [Ben03].*

Since a type-space functor is one way to present a cat, this can be restated as: every continuous first order theory is an open Hausdorff cat.

Proof. Clearly $S(T)$ is a Hausdorff compact type-space functor. To see it is open, let $\pi_n: S_{n+1}(T) \rightarrow S_n(T)$ consist of restriction to the n first variables (so $\pi_n = (n \hookrightarrow n+1)^*$). Let $U \subseteq S_{n+1}(T)$ be a basic open set, i.e., of the form $[\varphi(\bar{x}, y) < r]$. Then $\pi_n(U) = [\inf_y \varphi < r]$ is open as well, so π_n is an open mapping. ■_{4.20}

Recall that a *definable n -ary function* from a cat T to a Hausdorff space X is a continuous mapping $f: S_n(T) \rightarrow X$. Equivalently, this is a mapping from the models of T to X such that for every closed subset $F \subseteq X$, the property $f(\bar{x}) \in F$ is type-definable without parameters (whence *definable function*). A *definable metric* is a definable binary function which defines a metric on the models.

Note that d is indeed a definable metric, so T is a metric cat, and the models of T (in the sense of continuous first-order logic) are precisely its complete models as a metric Hausdorff cat, as defined in [Ben05].

For the converse, we will use the following property of definable functions in open cats:

Lemma 4.21. *Let T be an open cat. Let $f(\bar{x}, y)$ be an definable $n+1$ -ary function from T to $[0, 1]$ (or the reals, for that matter), and let $g(\bar{x}) = \sup_y f(\bar{x}, y)$. Then g is also a definable function.*

Proof. For every real number, we can define $g(\bar{x}) \geq r$ by the partial type $\bigwedge_{s < r} \exists y f(\bar{x}, y) \geq s$. We can also define $g(\bar{x}) \leq r$ by $\forall y f(\bar{x}, y) \leq r$, and this is expressible by a partial type since T is assumed to be open. ■_{4.21}

Remark 4.22. This can also be stated in purely topological terms: Let X and Y be compact Hausdorff spaces, and $f: X \rightarrow Y$ an open continuous surjective mapping. Let $\varphi: X \rightarrow [0, 1]$ be continuous, and let $\psi: Y \rightarrow [0, 1]$ be defined by $\psi(y) = \sup\{\varphi(x) : f(x) = y\}$. Then ψ is continuous.

Observe that if T is a metric open Hausdorff cat and d a definable metric on some sort, then by compactness the metric is bounded. Thus, up to rescaling we may always assume its range is contained in $[0, 1]$.

Theorem 4.23. *Let T be a metric open Hausdorff cat, and let d be a definable metric on the home sort with range in $[0, 1]$.*

Then there exists a metric signature $\hat{\mathcal{L}}$ whose distinguished metric symbol is \hat{d} , and an $\hat{\mathcal{L}}$ -theory \hat{T} , such that $S(\hat{T}) \simeq S(T)$, and such that the metric \hat{d} coincide with d .

Moreover, if κ is such that $S_n(T)$ has a basis of cardinality $\leq \kappa$ for all $n < \omega$, then we can arrange that $|\mathcal{L}| \leq \kappa$.

Proof. For each n , we choose a family $F_n \subseteq C(S_n(T), [0, 1])$ which separates points, are closed under $\{\neg, \div, \frac{x}{2}\}$, and such that $d \in F_2$. By Lemma 4.21 we may further assume that $\sup_y P(\bar{x}, y) \in F_n$ for each $P(\bar{x}, y) \in F_{n+1}$. We can always choose the F_n such that $|F_n| \leq \kappa$ for all n , where κ is as in the moreover part. By Corollary 1.4, F_n is dense in $C(S_n(T), [0, 1])$.

We observe that every $P \in F_n$ is uniformly continuous with respect to d . Indeed, $|P(\bar{x}) - P(\bar{y})|$ is a continuous function from $S_{2n}(T)$, and for every $\varepsilon > 0$, the following partial type is necessarily inconsistent:

$$\{|P(\bar{x}) - P(\bar{y})| \geq \varepsilon\} \cup \{d(x_i, y_i) \leq 2^{-m} : m < \omega, i < n\}.$$

Therefore there is $m < \omega$ such that $d(\bar{x}, \bar{y}) \leq 2^{-m} \implies |P(\bar{x}) - P(\bar{y})| \leq \varepsilon$.

Let $\hat{\mathcal{L}}_n = \{\hat{P} : n < \omega, P \in F_n\}$, where we associate to each n -ary predicate symbol \hat{P} the uniform continuity moduli obtained in the previous paragraph. Every universal domain of T , or closed subset thereof, is naturally a $\hat{\mathcal{L}}$ -structure, by interpreting each \hat{P} as P . In particular, all the predicates satisfy the appropriate continuity moduli.

Clearly, the family of n -ary definable functions is closed under continuous connectives. Also, if $\varphi(x_{<n}) : S_n(T) \rightarrow [0, 1]$ is a definable n -ary function and $f : n \rightarrow m$ is any mapping, then $\varphi(x_{i_0}, \dots, x_{i_{n-1}}) = \varphi \circ f^*$ is a definable m -ary function: in other words, the definable functions are closed under changes of variables. Finally, by Lemma 4.21, the definable functions are closed under quantification. Put together, every $\hat{\mathcal{L}}$ -formula $\varphi(x_{<n})$ induces a definable function $\varphi \in C(S_n(T), [0, 1])$. Since F_n is dense there, for every m we can find $P_{\varphi, m} \in F_n$ such that $|P_{\varphi, m} - \varphi| \leq 2^{-m}$. Let \hat{T} consist of:

$$\begin{array}{ll} \sup_{\bar{x}} |\hat{P}_{\varphi, m}(\bar{x}) - \varphi(\bar{x})| \leq 2^{-m} & \varphi, P_{\varphi, m} \text{ as above} \\ \sup_{\bar{x}} \hat{P} \leq 2^{-m} & P \in F_n \text{ and } P(S_n(T)) \subseteq [0, 2^{-m}] \end{array}$$

Clearly, every model of T , viewed as an $\hat{\mathcal{L}}$ -structure, is a model of \hat{T} .

We claim that $S(\hat{T}) \approx S(T)$. Indeed, for every $p \in S(T)$, define $\theta_n(p)$ by $\varphi^{\theta_n(p)} = \varphi(p)$ for every $\hat{\mathcal{L}}$ -formula $\varphi(x_{<n})$. Then $\theta_n(p)$ is just the $\hat{\mathcal{L}}$ -type of any realisation of p (again, viewing models of T as $\hat{\mathcal{L}}$ -structure), so in particular is it indeed in $S_n(\hat{T})$. Since the

predicate symbols of $\hat{\mathcal{L}}$ separate T -types, θ_n is injective. It is also clearly continuous, and therefore (as its domain is compact and its target Hausdorff) is a homeomorphism of $S_n(T)$ with a closed subspace of $S_n(\hat{T})$. If θ_n is not surjective, there is $q \in S_n(\hat{T})$, and a neighbourhood $q \in U$ which is disjoint of the image of θ_n . As usual there is a formula φ such that $q \in [\varphi = 0] \subseteq [\varphi < 1/2] \subseteq U$, and letting $\psi = \neg(\varphi \dot{+} \varphi)$: $q \in [\psi = 1] \subseteq [0 < \psi] \subseteq U$, so $\psi(S_n(T)) = \{0\}$. Then \hat{T} says that $\sup_{\bar{x}} |\psi(\bar{x}) - \hat{P}_{\psi,2}(\bar{x})| \leq 1/4$ and $\sup_{\bar{x}} \hat{P}_{\psi,2}(\bar{x}) \leq 1/4$, which in turn imply that $\psi \leq 1/2$ and therefore $\varphi \geq 1/4$, so q is inconsistent with \hat{T} .

Thus $\theta_n: S_n(T) \rightarrow S_n(\hat{T})$ is a homeomorphism for every n , and by construction it is compatible with the functor structure so $\theta: S(T) \rightarrow S(\hat{T})$ is the required homeomorphism of type-space functors.

It is clear from the way we axiomatised it that \hat{T} has quantifier elimination. ■_{4.23}

Theorem 4.23, combined Fact 4.20, says that framework of continuous first order theories coincides with that of metric open Hausdorff cats. In fact, we know from [Ben05] that if T is a non-metric Hausdorff cat then its home sort can be “split” into (uncountably many) hyperimaginary metric sorts, so in a sense every Hausdorff cat is metric. Thus, with some care, this observation can be generalised to all open Hausdorff cats.

Compact type-space functors are a structure-free and language-free way of presenting cats. By [Ben03], one can associate to each compact type-space functor a positive Robinson theory T' in some language \mathcal{L}' and talk about universal domains for that cat as \mathcal{L}' -structures. Then “complete models” mean the same thing in both settings:

Fact 4.24. *Let T be a complete continuous theory and T' the corresponding cat (i.e., positive Robinson theory in a language \mathcal{L}').*

Then there is a one-to-one correspondence between (κ) -monster models of T and (κ) -universal domains of T' over the same set of elements (for a fixed big cardinal κ), in such a way that the type of a tuple in one is the same as its type in the other.

Moreover, a closed subset of such a model is an elementary submodel in the sense of T if and only if it is a complete submodel in the sense of T' , as defined in [Ben05].

Proof. For the first condition, the identification of the type-space functors of T and of T' imposes for each monster model (or universal domain) of one an interpretation of the other language, and it is straightforward (though tedious) to verify then that the new structure is indeed a universal domain (or a monster model) for the other theory.

The moreover part is a special case of Fact 4.6. ■_{4.24}

Since T is a metric cat, its monster model is a monster metric space (momspace) as defined in [SU], and its models (in the sense of continuous logic) are precisely the class K_1^c studied there. So results proved for momspaces apply in our context.

5. IMAGINARIES

In classical first order model theory there are two common ways to view (and define) imaginaries: as canonical parameters for formulae, or, which is more common, as classes modulo definable equivalence relations. Of course, any canonical parameters for a formula can be viewed as an equivalence class, and an equivalence class is a canonical parameter for the formula defining it, so both approaches are quite equivalent in the discrete setting.

We have already observed that in the passage from discrete to continuous logic equivalence relations are replaced with pseudo-metrics. On the other hand, the notion of a canonical parameter remains essentially the same: the canonical parameter for $\varphi(\bar{x}, \bar{a})$ is something (a tuple, an imaginary...) c which an automorphism fixes if and only if it does not alter the meaning of the formula (i.e., $c = f(c) \iff \varphi(\bar{x}, \bar{a}) \equiv \varphi(x, f(\bar{a}))$) for every $f \in \text{Aut}(\mathfrak{M})$.

As in the classical setting, both approaches are essentially equivalent, but in practise the canonical parameter approach has considerable advantages. In particular, when doing stability, we would need to consider canonical parameters for definable predicates $\psi(\bar{x}, A)$, which only has finitely many free variables but may have infinitely many parameters. Canonical parameters for such definable predicates are dealt with as with canonical parameters for formulae, and the existence of infinitely many parameters introduces very few additional complications. On the other hand, if we wished to define the canonical parameter as an equivalence class modulo a pseudo-metric we would be forced to consider pseudo-metrics on infinite tuples, the logic for whose equivalence classes could become messy.

Other minor advantages include the fact that we need not ask ourselves whether a particular formula defines a pseudo-metric on every structure or only on models of a given theory, and finally the conceptually convenient fact that unlike equivalence relations which need to be replaced with pseudo-metrics, canonical parameters are a familiar notion which we leave unchanged.

Let \mathcal{L} be a continuous signature. For convenience, assume that \mathcal{L} has a single sort S . If we wanted to work with a many-sorted language we would have to keep track on which variables (in the original language) belong to which sort, but other than that the treatment is identical.

Let us start with the case of a formula $\varphi(x_{<n}, y_{<m})$, where $x_{<n}$ are the free variables, and $y_{<m}$ are the parameter variables (and to simplify notation we will consider that this partition of the variables is part of φ). We define a new signature \mathcal{L}_φ , consisting of \mathcal{L} along with a new sort S_φ for the canonical parameters for instances $\varphi(\bar{x}, \bar{a})$ of φ , and the following new symbols:

- (i) A predicate symbol $d_\varphi(z, z')$ on S_φ^2 which will be the distance symbol for S_φ .
- (ii) A predicate symbol $P_\varphi(x_{<n}, z)$ on $S^n \times S_\varphi$. Its uniform continuity moduli with respect to the first n arguments are the same as φ 's with respect to $x_{<n}$, and

with respect to the last argument it is the identity. (This symbol is not strictly necessary, but will be convenient.)

We will expand every \mathcal{L} -structure M to an \mathcal{L}_φ -structure, interpreting S_φ as the family the canonical parameters of all instances $\varphi(\bar{x}, \bar{a})$ of $\varphi(\bar{x}, \bar{y})$ in M . We first expand M into an \mathcal{L}_φ -pre-structure $M_{\varphi,0}$ by defining:

$$\begin{aligned} S_\varphi^{M_{\varphi,0}} &= M^m \\ P_\varphi^{M_{\varphi,0}}(\bar{a}, (\bar{b})) &= \varphi^M(\bar{a}, \bar{b}) \\ d_\varphi^{M_{\varphi,0}}((\bar{b}), (\bar{b}')) &= \sup_{\bar{x}} |\varphi(\bar{x}, \bar{b}) - \varphi(\bar{x}, \bar{b}')|^M \end{aligned}$$

(Here \bar{b} is an m -tuple in the home sort of M , (\bar{b}) is the corresponding element of $S_\varphi^{M_{\varphi,0}}$.) We leave the reader the verification that $d_\varphi^{M_{\varphi,0}}$ is a pseudo-metric and the uniform continuity moduli fixed above are indeed respected, so $M_{\varphi,0}$ is an \mathcal{L}_φ -pre-structure. We then define $M_\varphi = \widehat{M_{\varphi,0}}$, the structure associated to the pre-structure $M_{\varphi,0}$.

For $\bar{b} \in M^m$, let $[\bar{b}]_\varphi$ denote the image of (\bar{b}) in $S_\varphi^{M_\varphi}$. Clearly, every automorphism of M extends uniquely to an automorphism of M_φ , and it fixes $\varphi(\bar{x}, \bar{b})$ if and only if it fixes $[\bar{b}]_\varphi$, which is therefore a canonical parameter for $\varphi(\bar{x}, \bar{b})$. If $c = [\bar{b}]_\varphi$ then $P_\varphi(\bar{x}, c)$ coincides with $\varphi(\bar{x}, \bar{b})$. We therefore allow ourselves to abuse notation and denote either one simply as $\varphi(\bar{x}, c)$.

The properties of the new sort are described intuitively by the following axioms:

$$\begin{aligned} \forall z z' \left(d_\varphi(z, z') &= \sup_{\bar{x}} |P_\varphi(\bar{x}, z) - P_\varphi(\bar{x}, z')| \right), \\ \forall z \exists \bar{y} \forall \bar{x} \left(\varphi(\bar{x}, \bar{y}) &= P_\varphi(\bar{x}, z) \right), \\ \forall \bar{y} \exists z \forall \bar{x} \left(\varphi(\bar{x}, \bar{y}) &= P_\varphi(\bar{x}, z) \right). \end{aligned}$$

To get the precise axioms we apply Remark 2.11, so “ $\forall \exists \forall (\dots = \dots)$ ” should be read as “ $\sup \inf \sup | \dots - \dots | = 0$ ”, etc. We therefore define T_φ to be the following \mathcal{L}_φ -theory:

$$\begin{aligned} \sup_{zz'} \left| d_\varphi(z, z') - \sup_{\bar{x}} |P_\varphi(\bar{x}, z) - P_\varphi(\bar{x}, z')| \right| &= 0, \\ \sup_z \inf_{\bar{y}} \sup_{\bar{x}} \left| \varphi(\bar{x}, \bar{y}) - P_\varphi(\bar{x}, z) \right| &= 0, \\ \sup_{\bar{y}} \inf_z \sup_{\bar{x}} \left| \varphi(\bar{x}, \bar{y}) - P_\varphi(\bar{x}, z) \right| &= 0. \end{aligned}$$

One easily verifies that:

Proposition 5.1. *An \mathcal{L}_φ -structure is a model of T_φ if and only if it is of the form M_φ for some \mathcal{L} -structure M .*

Therefore, if T is a complete \mathcal{L} -theory then $T \cup T_\varphi$ is a complete \mathcal{L}_φ -theory. We discussed the case of a single formula φ , but we can do the same with several (all) formulae simultaneously.

Remark 5.2. As we said earlier, the continuous analogue of an equivalence relation is a pseudo-metric. We can recover classes modulo pseudo-metrics from canonical parameters in very straightforward manner:

- (i) Assume that $\varphi(\bar{x}, \bar{y})$ defines a pseudo-metric on M^n . Then the pseudo-metric $d_\varphi^{M_\varphi, 0}$ (defined on $S_\varphi^{M_\varphi, 0} = M^n$) coincides with φ^M , and $(S_\varphi, d_\varphi)^{M_\varphi}$ is the completion of the set of equivalence classes of n -tuples modulo the relation $\varphi(\bar{a}, \bar{b})$, equipped with the induced metric.
- (ii) In particular, let $\xi_n(x_{<n}, y_{<n}) = \bigvee_{i < n} d(x_i, y_i)$. Then the sort S_{ξ_n} is the sort of n -tuples with the standard metric.

We now wish to define canonical parameters to definable predicates of the form $\psi(\bar{x}, B)$, i.e. to instances of $\psi(\bar{x}, Y)$, which may unavoidably have infinitely many parameters (or parameter variables). We can write $\psi(\bar{x}, Y)$ as a uniform limit of a sequence of formulae $(\varphi_n(\bar{x}, \bar{y}_n) : n < \omega)$, where (\bar{y}_n) is an increasing sequence of tuples and $Y = \bigcup_n \bar{y}_n$. We may further assume that the rate of convergence is such that $|\varphi_n - \psi| \leq 2^{-n}$. Since each \bar{y}_n is finite, we may assume that $|Y| = \omega$.

We define \mathcal{L}_ψ by adding to \mathcal{L} a new sort S_ψ and predicate symbols $d_\psi(z, z')$ and $P_\psi(\bar{x}, z)$ as before.

Given a structure M we construct $M_{\psi, 0}$ much the same as before:

$$\begin{aligned} S_\psi^{M_{\psi, 0}} &= M^\omega \\ P_\psi^{M_{\psi, 0}}(\bar{a}, (B)) &= \psi^M(\bar{a}, B) (= \lim \varphi_n^M(\bar{a}, \bar{b}_n)) \\ d_\psi^{M_{\psi, 0}}((B), (B')) &= \sup_{\bar{x}} |\psi(\bar{x}, B) - \psi(\bar{x}, B')|^M \end{aligned}$$

Again $M_{\psi, 0}$ is a pre-structure, and we define $M_\psi = \widehat{M_{\psi, 0}}$. Letting $[B]_\psi$ be the image of (B) in M_ψ , we see again that $[B]_\psi$ is a canonical parameter for $\psi(\bar{x}, B)$. If $c = [B]_\psi$ then again we allow ourselves the abuse of notation which consists of denoting either of $\psi(\bar{x}, B)$ and $P_\psi(\bar{x}, c)$ (which are equivalent) by $\varphi(\bar{x}, c)$.

Notice that \mathcal{L}_ψ is a finitary language in which the canonical parameters are singletons, and the fact that Y is an infinite tuple is indeed hidden.

The theory T_ψ , the analogue of T_φ above, consists of infinitely many axioms, according to the following schemes, which we interpret according to Remark 2.11 as above:

$$\begin{aligned} \forall z z' (d_\psi(z, z') &= \sup_{\bar{x}} |P_\psi(\bar{x}, z) - P_\psi(\bar{x}, z')|), \\ \forall z \exists \bar{y}_n \forall \bar{x} |\varphi_n(\bar{x}, \bar{y}_n) - P_\psi(\bar{x}, z)| &\leq 2^{-n}, \\ \forall y_n \exists \bar{z} \forall \bar{x} |\varphi_n(\bar{x}, \bar{y}_n) - P_\psi(\bar{x}, z)| &\leq 2^{-n} \end{aligned}$$

Again, the models of T_ψ are precisely the \mathcal{L}_ψ -structures of the form M_ψ .

Remark 5.3. It is not true that if T is model complete then so is $T \cup T_\varphi$.

In the case of canonical parameters for a single formula φ we can remedy this deficiency as in discrete logic by naming the mapping $\bar{b} \mapsto [\bar{b}]_\varphi$ with a new function symbol

$\pi_\varphi: S^m \rightarrow S_\varphi$. We leave it to the reader to verify that, if T is a model complete \mathcal{L} -theory then $T \cup T'_\varphi$ is a model complete \mathcal{L}'_φ -theory, where

$$\mathcal{L}'_\varphi = \mathcal{L}_\varphi \cup \{\pi_\varphi\}, \quad T'_\varphi = T_\varphi \cup \{\forall \bar{x}\bar{y} (\varphi(\bar{x}, \bar{y}) = P_\varphi(\bar{x}, \pi_\varphi(\bar{y})))\}.$$

The *graph* of π_φ (or of any function of continuous structures) is defined here to be the predicate $\gamma_\varphi(\bar{y}, z) = d_\varphi(\pi_\varphi(\bar{y}), z)$. This predicate is definable in \mathcal{L}_φ :

$$\gamma_\varphi(\bar{y}, z) = \sup_{\bar{x}} |\varphi(\bar{x}, \bar{y}) - P_\varphi(\bar{x}, z)|.$$

It follows that the addition of π_φ to the language does not add any structure (see [Bena, Section 1] for a more detailed discussion of definable functions in continuous logic).

In the case of canonical parameters of $\psi(\bar{x}, Y) = \lim \varphi_n(\bar{x}, \bar{y}_n)$, which are quotients of infinite tuples, we cannot add a function symbol π_φ . Instead we observe that in the case of a single formula it would have sufficed to name γ_φ by a predicate (rather than naming π_φ). While γ_ψ would depend infinitely many variables and thus still impossible to add to the language, we may add finite approximations. We add predicate symbols $\gamma_{\varphi_n, \psi}(\bar{y}_n, z)$ and add to T_ψ the axioms:

$$\forall \bar{y}_n z (\gamma_{\varphi_n, \psi}(\bar{y}_n, z) = \sup_{\bar{x}} |\varphi_n(\bar{x}, \bar{y}_n) - P_\psi(\bar{x}, z)|).$$

Call the expanded language \mathcal{L}'_ψ and the expanded theory T'_ψ . Again, we leave it to the reader to verify that if T is a model complete \mathcal{L} -theory, then so is $T \cup T'_\psi$ as an \mathcal{L}'_ψ -theory.

We leave the details to the interested reader.

6. LOCAL TYPES AND φ -PREDICATES

In this section and later we will consider formulae whose free variables are split in two groups $\varphi(\bar{x}, \bar{y})$. Following Remark 5.2, we may replace finite tuples of variables with single ones, and therefore allow ourselves to restrict our attention to formulae of the form $\varphi(x, y)$, where x and y may belong to distinct sorts.

We will associate variable letters with sorts: x, x_i , etc., belong to one sort, y, y_j , etc., to another, and so forth. Accordingly, we will denote the difference spaces of types in the variable x by $S_x(T)$, $S_x(A)$, etc.

We now fix a formula $\varphi(x, y)$. We will define spaces of φ -types as quotients of spaces we have already constructed using Fact 4.7. The notion of a φ -type over a model is fairly straightforward. A few more steps will be required in order to obtain the correct notion of a φ -type over an arbitrary set.

Definition 6.1. Let M be a structure. We define $S_\varphi(M)$ as the quotient of $S_x(M)$ given by the family of functions $\mathcal{A}_{M, \varphi} = \{\varphi(x, b) : b \in M\}$. An element of $S_\varphi(M)$ is called a (*complete*) φ -*type* over M .

Accordingly, the φ -*type of a over M* , denoted $\text{tp}_\varphi(a/M)$, is given by the mappings $\varphi(x, b) \mapsto \varphi(a, b)$ where b varies over all elements of the appropriate sort of M .

We equip $S_\varphi(M)$ with a metric structure: For $p, q \in S_\varphi(M)$ we define

$$d(p, q) = \sup_{b \in M} |\varphi(x, b)^p - \varphi(x, b)^q|.$$

Fact 6.2. *Equipped with this metric and with its natural topology (inherited as a quotient space from $S_x(M)$), $S_\varphi(M)$ is a compact topometric space as in Definition 4.11.*

Definition 6.3. Let M be a structure. A φ -predicate over M , or an M -definable φ -predicate, is a continuous mapping $\psi: S_\varphi(M) \rightarrow [0, 1]$.

Fact 6.4. *Let $\psi: S_x(M) \rightarrow [0, 1]$ be an M -definable predicate. Then the following are equivalent:*

- (i) ψ is a φ -predicate (i.e., factors through the projection $S_x(M) \rightarrow S_\varphi(M)$).
- (ii) There are formulae $\psi_n(x, \bar{b}_n)$, each obtained using connectives from several instances $\varphi(x, b_{n,j})$, where each $b_{n,j} \in M$, and in M we have $\psi(x) = \mathcal{F}\lim \psi_n(x)$.
- (iii) ψ can be written as $f \circ (\varphi(x, b_i))^M: i < \omega$ where $f: [0, 1]^\omega \rightarrow [0, 1]$ is continuous and $b_i \in M$ for all $i < \omega$.

Proof. (i) \implies (ii). Standard application of Corollary 1.4.

(ii) \implies (iii) \implies (i). Immediate. ■_{6.4}

Lemma 6.5. *Let \mathfrak{M} be a monster model, and $M \preceq \mathfrak{M}$ a model. Let $\psi(x)$ be an M -invariant φ -predicate over \mathfrak{M} . Then ψ is (equal to) a φ -predicate over M .*

Proof. We know that $\psi(x)$ is equal to a definable predicate over M , so $\psi(x) = \lim \psi_n(x, c_n)$ where each $\psi_n(x, z_n)$ is a formula and $c_n \in M$. We also know that $\psi(x) = \lim \chi_n(x, \bar{d}_n)$, where each $\chi_n(x, \bar{d}_n)$ is a combination of instances $\varphi(x, d_{n,j})$ with parameters $d_{n,j} \in \mathfrak{M}$. For all $\varepsilon > 0$ there exists $n < \omega$ such that:

$$\mathfrak{M} \models \sup_x |\psi_n(x, c_n) - \chi_n(x, \bar{d}_n)| < \varepsilon$$

In particular:

$$\mathfrak{M} \models \inf_{\bar{y}_n} \sup_x |\psi_n(x, c_n) - \chi_n(x, \bar{y}_n)| < \varepsilon$$

Since $M \preceq \mathfrak{M}$:

$$M \models \inf_{\bar{y}_n} \sup_x |\psi_n(x, c_n) - \chi_n(x, \bar{y}_n)| < \varepsilon$$

So there is $\bar{d}'_n \in M$ such that:

$$M \models \sup_x |\psi_n(x, c_n) - \chi_n(x, \bar{d}'_n)| < \varepsilon.$$

We can therefore express ψ as $\lim \chi_n(x, \bar{d}'_n)$, which is a φ -predicate over M . ■_{6.5}

This leads to the following:

Definition 6.6. Let A be a (small) set in a monster model \mathfrak{M} .

- (i) A φ -predicate over A , or an A -definable φ -predicate is a φ -predicate over \mathfrak{M} which is A -invariant.
- (ii) We define $S_\varphi(A)$ as the quotient of $S_\varphi(\mathfrak{M})$ determined by the A -definable φ -predicates. The points of this space are called (*complete*) φ -types over A .
- (iii) Accordingly, the φ -type $\text{tp}_\varphi(a/A)$ is given by the mappings $\psi(x) \mapsto \psi(a)$ where ψ varies over all A -definable φ -predicates.

By Lemma 6.5, these definitions coincide with previous ones in case that $A = M \preceq \mathfrak{M}$.

We conclude with a result about compatibility of two kinds of extensions of local types: to the algebraic closure of the set of parameters, and to more (all) formulae.

Lemma 6.7. *Let $p \in S_\varphi(A)$, and let $q \in S_\varphi(\text{acl}(A))$ and $r \in S_x(A)$ extend p . Then $q \cup r$ is consistent.*

Proof. Let $R_1 \subseteq S_x(\text{acl}(A))$ be the pullback of r (i.e., the set of all its extensions to a complete type over $\text{acl}(A)$), and $R_2 \subseteq S_\varphi(\text{acl}(A))$ the image of R_1 under the restriction projection $S_x(\text{acl}(A)) \rightarrow S_\varphi(\text{acl}(A))$. Then R_2 is the set of all extensions of p to $\text{acl}(A)$ compatible with r , and we need to show that $q \in R_2$ (i.e., that R_2 is the set of all the extensions of p).

Indeed, assume not. The sets R_1 and therefore R_2 are closed. Therefore we can separate R_2 from q by a φ -predicate $\psi(x, a)$, with parameter $a \in \text{acl}(A)$, such that $\psi(x, a)^q = 0$ and $R_2 \subseteq [\psi(x, a) = 1]$. Since $a \in \text{acl}(A)$, by Lemma 4.9 there is a sequence $(a_i : i < \omega)$ such that:

- (i) Every a_i is an A -conjugate of a .
- (ii) For every $\varepsilon > 0$ there is n such that every A -conjugate of a is in the ε -neighbourhood of some a_i for $i < n$.

Define $\psi_n(x, a_{<n}) = \bigwedge_{i < n} \psi(x, a_i)$. Then, by uniform continuity of $\psi(x, y)$ with respect to y , the sequence $(\psi_n(x, a_{<n}) : n < \omega)$ converges uniformly to the predicate $\chi(x) = \inf\{\psi(x, a') : a' \equiv_A a\}$. Thus $\chi(x)$ is a definable φ -predicate (as a limit of such) and A -invariant, so it is an A -definable φ -predicate. On the one hand we clearly have $\chi(x)^p = \chi(x)^q = 0$. On the other, as R_2 is A -invariant as well, we have $R_2 \subseteq [\psi(x, a_i) = 1]$ for all $i < \omega$, so $\chi(x)^p = \chi(x)^r = 1$. This contradiction concludes the proof. ■_{6.7}

Lemma 6.8. *Let $A \subseteq M$ where M is strongly $(|A| + \omega)^+$ -homogeneous, and $p \in S_\varphi(A)$. Then $\text{Aut}(M/A)$ acts transitively on the extensions of p to $S_\varphi(\text{acl}(A))$.*

Proof. Follows from (and is in fact equivalent to) Lemma 6.7. ■_{6.8}

7. LOCAL STABILITY

Here we answer C. Ward Henson's question mentioned in the introduction. Throughout this section T is a fixed continuous theory (not necessarily complete) in a signature \mathcal{L} .

Definition 7.1. (i) We say that a formula $\varphi(x, y)$ is ε -stable for a real number $\varepsilon > 0$ if in models of T there is no infinite sequence $(a_i b_i : i < \omega)$ satisfying for all $i < j$: $|\varphi(a_i, b_j) - \varphi(a_j, b_i)| \geq \varepsilon$.

(ii) We say that $\varphi(x, y)$ is *stable* if it is ε -stable for all $\varepsilon > 0$.

Lemma 7.2. *Let $\varphi(x, y)$ be a formula, $\varepsilon > 0$. Then the following are equivalent:*

- (i) *The formula φ is ε -stable.*
 - (ii) *It is impossible to find $0 \leq r < s \leq 1$ and an infinite sequence $(a_i b_i : i < \omega)$ such that $r \leq s - \varepsilon$ and for all $i < j$: $\varphi(a_i, b_j) \leq r$, $\varphi(a_j, b_i) \geq s$.*
 - (iii) *There exists a natural number N such that in model of T there is no finite sequence $(a_i b_i : i < N)$ satisfying:*
- (*) $\text{for all } i < j < k : \quad |\varphi(a_j, b_i) - \varphi(a_j, b_k)| \geq \varepsilon.$

Proof. (i) \iff (ii). Left to right is immediate. For the converse assume φ is not ε -stable, and let the sequence $(a_i b_i : i < \omega)$ witness this. For every $\delta > 0$ we can find using Ramsey's Theorem arbitrarily long sub-sequences $(a'_i b'_i : i < N)$ such that in addition:

$$\text{If } i < j \text{ and } i' < j' \text{ then: } |\varphi(a_i, b_j) - \varphi(a_{i'}, b_{j'})|, |\varphi(a_j, b_i) - \varphi(a_{j'}, b_{i'})| \leq \delta.$$

(For this we use the classical finite Ramsey's Theorem. We could also use the infinite version to obtain a single infinite sequence with the same properties.)

By compactness we can find an infinite sequence $(c_i d_i : i < \omega)$ witnessing ε -instability such that in addition, for $i < j$, $\varphi(c_i, d_j) = r$ and $\varphi(c_j, d_i) = s$ do not depend on i, j . Thus $|r - s| \geq \varepsilon$. If $r < s$ we are done. If $r > s$ we can reverse the ordering on all the finite subsequences obtained above, thus exchanging r and s , and conclude in the same manner.

(ii) \iff (iii). Now right to left is immediate. For left to write, we argue as above, using Ramsey's Theorem and compactness, that there exists an infinite sequence $(a_i b_i : i < \omega)$ such that for $i < j < k$ we have $|\varphi(a_j, b_i) - \varphi(a_j, b_k)| \geq \varepsilon$ and $\varphi(a_i, b_j) = r$ and $\varphi(a_j, b_i) = s$ do not depend on i, j . Then again $|r - s| \geq \varepsilon$ and we conclude as above. $\blacksquare_{7.2}$

It follows that stability is a symmetric property: define $\tilde{\varphi}(y, x) \stackrel{\text{def}}{=} \varphi(x, y)$; then φ is (ε) -stable if and only if $\tilde{\varphi}$ is.

Notation 7.3. If φ is ε -stable we define $N(\varphi, \varepsilon)$ to be the minimal N such that no sequence $(a_i b_i : i < N + 1)$ exists satisfying (*).

Let us define the *median value* connective $\text{med}_n : [0, 1]^{2n-1} \rightarrow [0, 1]$:

$$\text{med}_n(t_{<2n-1}) = \bigwedge_{w \in [2n-1]^n} \bigvee_{i \in w} t_i = \bigvee_{w \in [2n-1]^n} \bigwedge_{i \in w} t_i.$$

If $\varphi(x, y)$ is ε -stable define:

$$d^\varepsilon \varphi(y, x_{<2N(\varphi, \varepsilon)-1}) = \text{med}_{N(\varphi, \varepsilon)}(\varphi(x_i, y) : i < 2N(\varphi, \varepsilon) - 1).$$

Lemma 7.4. *Let M be a model and $p \in S_\varphi(M)$. Then there exist $c_{<2N(\varphi, \varepsilon)-1}^\varepsilon \in M$ such that, for every $b \in M$:*

$$|\varphi(x, b)^p - d^\varepsilon \varphi(b, c_{<2N(\varphi, \varepsilon)-1}^\varepsilon)| \leq \varepsilon.$$

Proof. We argue as in the proof of [Pil96, Lemma 2.2]. Choose a realisation $c \models p$ in the monster model: $c \in \mathfrak{M} \succeq M$. We construct by induction on n , tuples $c_n \in M$ in the sort of x , an increasing sequence of sets $K(n) \subseteq \mathbb{P}(\omega)$, and tuples $a_w \in M$ in the sort of y for each $w \in K(n)$, as follows.

At the n th step we assume we have already chosen $c_{<n}$. We define:

$$K(n) = \{w \subseteq n : \exists a \in M \text{ such that } |\varphi(c, a) - \varphi(c_i, a)| > \varepsilon \text{ for all } i \in w\}.$$

For every $w \in K(n)$ such that a_w has not yet been chosen, choose $a_w \in M$ witnessing that $w \in K(n)$. Note that if $w \subseteq m < n$ and a_w witnesses that $w \in K(m)$ then it also witnesses that $w \in K(n)$, so there is no problem keeping previously made choices. We now have:

$$\forall w \in K(n), i \in w : \quad |\varphi(c_i, a_w) - \varphi(c, a_w)| > \varepsilon,$$

whereby:

$$\sup_x \left(\bigwedge_{\substack{w \in K(n) \\ i \in w}} |\varphi(c_i, a_w) - \varphi(x, a_w)| \right) > \varepsilon.$$

This holds in \mathfrak{M} ; but since all of the parameters of the form a_w, b_w, c_i are in M , the last inequality actually holds in M . Therefore there exists $c_n \in M$ such that:

$$\forall w \in K(n), i \in w : \quad |\varphi(c_i, a_w) - \varphi(c_n, a_w)| > \varepsilon.$$

This concludes the n th step of the construction.

Note that if $w \in K(n)$ and $m < n$ then $w \cap m \in K(m)$ as well.

Claim. For all n and $w \in K(n)$: $|w| < N(\varphi, \varepsilon)$.

Proof of claim. If not there is $w = \{m_0 < \dots < m_{N-1}\} \in K(n)$ where $N \geq N(\varphi, \varepsilon)$. Define $m_N = n$ (so $m_{N-1} < m_N$), and for $j < N$, let $w_j = \{m_i : i < j\}$. Then for all $i < j < k \leq N$ we have $m_i \in w_j \in K(m_k)$, whereby:

$$|\varphi(c_{m_i}, a_{w_j}) - \varphi(c_{m_k}, a_{w_j})| > \varepsilon.$$

Thus the sequence $(c_{m_i}, a_{w_i} : i < N + 1)$ contradicts the choice of $N(\varphi, \varepsilon)$. ■ *Claim*

It follows that for every $w \in [2N(\varphi, \varepsilon) - 1]^{N(\varphi, \varepsilon)}$ and $a \in M$:

$$\bigwedge_{i \in w} \varphi(c_i, a) - \varepsilon \leq \varphi(c, a) \leq \bigvee_{i \in w} \varphi(c_i, a) + \varepsilon$$

Whereby $|\varphi(c, a) - d^\varepsilon \varphi(a, c_{<2N(\varphi, \varepsilon)-1})| \leq \varepsilon$, as required. ■ _{7.4}

Definition 7.5. Let $p(x) \in S_\varphi(M)$. A *definition* for p is an M -definable predicate $\psi(y)$ satisfying $\varphi(x, b)^p = \psi^M(b)$ for all $b \in M$. If such a definable predicate exists then it is unique (any two such definable predicates coincide on M , and therefore on every elementary extension of M), and is denoted $d_p \varphi(y)$.

Assume now that $\varphi(x, y)$ is stable, and let:

$$X = (x_i^n : n < \omega, i < 2N(\varphi, 2^{-n}) - 1),$$

$$d\varphi(y, X) = \mathcal{F}\lim_n d^{2^{-n}}\varphi(y, x_{<2N(\varphi, 2^{-n})-1}^n).$$

Proposition 7.6. *Let M be a model, and $p \in S_\varphi(M)$. Then there are parameters $C \subseteq M$ such that $d\varphi(y, C) = d_p\varphi(y)$ (so in particular, a definition $d_p\varphi$ exists). Moreover, $d_p\varphi$ is an M -definable $\tilde{\varphi}$ -predicate.*

Proof. For each $n < \omega$ choose $c_{<2N(\varphi, 2^{-n})-1}^n$ as in Lemma 7.4, and let $C = (c_i^n : n < \omega, i < 2N(\varphi, 2^{-n}) - 1)$.

Let $\xi: M \rightarrow [0, 1]$ be defined as $b \mapsto \varphi(x, b)^p$. Then $|d^{2^{-n}}\varphi(y, c_{<2N(\varphi, 2^{-n})-1}^n)^M - \xi| \leq 2^{-n}$, whereby:

$$\xi = \mathcal{F}\lim_n d^{2^{-n}}\varphi(y, c_{<2N(\varphi, 2^{-n})-1}^n)^M = d\varphi(y, C)^M.$$

This precisely means that $d\varphi(y, C) = d_p\varphi$.

That $d\varphi(x, C)$ is a $\tilde{\varphi}$ -predicate follows from its construction. ■_{7.6}

From this point onwards we assume that \mathcal{L} has a sort for the canonical parameters of instances of $d\varphi(y, X)$ for every stable formula $\varphi(x, y) \in \mathcal{L}$. If not, we add these sorts as in Section 5. It should be pointed out that if M is an \mathcal{L} -structure and $\|M\| \geq |\mathcal{L}|$, the addition of the new sorts does not change $\|M\|$: this can be seen directly from the construction, or using the Downward Löwenheim-Skolem Theorem (Fact 2.19) and the fact that we do not change $|\mathcal{L}|$.

For every stable formula and type $p \in S_\varphi(M)$ we define $\text{Cb}_\varphi(p)$ as the canonical parameter of $d_p\varphi(y)$. With the convention above we have $\text{Cb}_\varphi(p) \in M$. Notice that if $p, q \in S_\varphi(M)$, $c = \text{Cb}_\varphi(p)$ and $c' = \text{Cb}_\varphi(q)$, then $d(c, c')$ (in the sense of the sort of canonical parameters for $d\varphi$) is equal to $d(p, q)$ in $S_\varphi(M)$.

As with structures, we will measure the size of a type space $S_\varphi(M)$ by its metric density character $\|S_\varphi(M)\|$.

Proposition 7.7. *The following are equivalent for a formula $\varphi(x, y)$:*

- (i) φ is stable.
- (ii) For every $M \models T$, every $p \in S_\varphi(M)$ is definable.
- (iii) For every $M \models T$, $\|S_\varphi(M)\| \leq \|M\|$.
- (iv) There exists $\lambda \geq |T|$ such that whenever $M \models T$ and $\|M\| \leq \lambda$ then $\|S_\varphi(M)\| \leq \lambda$ as well.

Proof. (i) \implies (ii). By Proposition 7.6.

(ii) \implies (iii). Let $D \subseteq M_{d\varphi}$ be the family of canonical parameters of instances of $d\varphi(y, X)$ which actually arise as definitions of φ -types over M . Then $\|D\| \leq \|M\|$, and D is isometric to $S_\varphi(M)$.

(iii) \implies (iv). Immediate.

(iv) \implies (i). Let $\lambda \geq |T|$ be any cardinal and assume φ is unstable. It is a classical fact that there exists a linear order $(I, <)$ of cardinality λ admitting $> \lambda$ initial segments: for example, let μ be the least cardinal such that $2^\mu > \lambda$ and let $I = \{0, 1\}^{<\mu}$ equipped with the lexicographic ordering.

Assuming φ is unstable then we can find (using Lemma 7.2 and compactness) $0 \leq r < s \leq 1$ and a sequence $(a_i b_i : i \in I)$ such that $i < j$ imply $\varphi(a_i, b_j) \leq r$ and $\varphi(a_j, b_i) \geq s$. By the Downward Löwenheim-Skolem Theorem there exists a model $\{b_i : i \in I\} \subseteq M \preceq \mathfrak{M}$ such that $\|M\| \leq \lambda$. On the other hand, by compactness, for every initial segment $C \subseteq I$ there exists a_C such that $\varphi(a_C, b_i) \geq s$ if $i \in C$ and $\varphi(a_C, b_i) \leq r$ if $i \notin C$. Let $p_C = \text{tp}_\varphi(a_C/M)$.

If C, C' are two distinct initial segments of I then $d(p_C, p_{C'}) \geq s - r$, showing that $\|S_\varphi(M)\| > \lambda$ and concluding the proof. $\blacksquare_{7.7}$

Definition 7.8. Let (X, d) be a metric space. The *diameter* of a subset $C \subseteq X$ is defined as

$$\text{diam}(C) = \sup\{d(x, y) : x, y \in C\}$$

We say that a subset $C \subseteq X$ is ε -finite if it can be written as $C = \bigcup_{i < k} C_i$, where $\text{diam}(C_i) \leq \varepsilon$ for all $i < k$. In this case, its ε -degree, denoted $\text{deg}_\varepsilon(C)$, is the minimal such k .

Note that if C and C' are ε -finite then so is $C \cup C'$ and $\text{deg}_\varepsilon(C \cup C') \leq \text{deg}_\varepsilon(C) + \text{deg}_\varepsilon(C')$.

Definition 7.9. Let X be a compact topometric space.

For a fixed $\varepsilon > 0$, we define a decreasing sequence of closed subsets $X_{\varepsilon, \alpha}$ by induction:

$$\begin{aligned} X_{\varepsilon, 0} &= X \\ X_{\varepsilon, \alpha} &= \bigcap_{\beta < \alpha} X_{\varepsilon, \beta} \quad (\alpha \text{ a limit ordinal}) \\ X_{\varepsilon, \alpha+1} &= \bigcap \{F \subseteq X_{\varepsilon, \alpha} : F \text{ is closed and } \text{diam}(X_{\varepsilon, \alpha} \setminus F) \leq \varepsilon\} \\ X_{\varepsilon, \infty} &= \bigcap_{\alpha} X_{\varepsilon, \alpha} \end{aligned}$$

Finally, for any non-empty subset $C \subseteq X$ we define its ε -Cantor-Bendixson rank in X as:

$$\text{CB}_{X, \varepsilon}(C) = \sup\{\alpha : C \cap X_{\varepsilon, \alpha} \neq \emptyset\} \in \text{Ord} \cup \{\infty\}$$

If $\text{CB}_{X, \varepsilon}(C) < \infty$ we also define $\text{CBm}_{X, \varepsilon}(C) = C \cap X_{\varepsilon, \text{CB}_{X, \varepsilon}(C)}$, i.e., the set of points of maximal rank.

It is worthwhile to point out that either $X_{\varepsilon, \alpha} \neq \emptyset$ for every α (and eventually stabilises to $X_{\varepsilon, \infty}$) or there is a maximal α such that $X_{\varepsilon, \alpha} \neq \emptyset$. The same holds for the sequence $\{C \cap X_{\varepsilon, \alpha} : \alpha \in \text{Ord}\}$ if $C \subseteq X$ is closed.

Assume that $C \subseteq X$ is closed and $\alpha = \text{CB}_{X,\varepsilon}(C) < \infty$. Then by the previous paragraph C contains points of maximal rank, i.e., $\text{CBm}_{X,\varepsilon}(C) \neq \emptyset$. Moreover, $\text{CBm}_{X,\varepsilon}(C)$ is compact and admits in $X_{\varepsilon,\alpha}$ an open covering by sets of diameter $\leq \varepsilon$. By compactness, it can be covered by finitely many such, and is therefore ε -finite. This need not necessarily hold in case C is not closed.

We will use this definition for $X = \text{S}_\varphi(M)$, where M is at least ω -saturated. In this case we may write $\text{CB}_{\varphi,M,\varepsilon}$ instead of $\text{CB}_{\text{S}_\varphi(M),\varepsilon}$, etc.

Remark 7.10. In the definition of the ε -Cantor-Bendixson rank we defined $X_{\varepsilon,\alpha+1}$ by removing from $X_{\varepsilon,\alpha}$ all its “small open subsets”, i.e., its open subsets of diameter $\leq \varepsilon$. There exist other possible definitions for the ε -Cantor-Bendixson derivative, using different notions of smallness. Such notions are studied in detail in [Bend] where it is shown that in the end they all boil down to the same thing.

Proposition 7.11. *φ is stable if and only if for one (any) ω -saturated model $M \models T$: $\text{CB}_{\varphi,M,\varepsilon}(\text{S}_\varphi(M)) < \infty$ for all ε .*

Proof. If not, let $Y = \{p \in \text{S}_\varphi(M) : \text{CB}_{\varphi,M,\varepsilon} = \infty\}$. Then Y is compact, and if $U \subseteq Y$ is relatively open and non-empty then $\text{diam}(U) > \varepsilon$. We can therefore find non-empty open sets U_0, U_1 such that $\bar{U}_0, \bar{U}_1 \subseteq U$ and $d(U_0, U_1) > \varepsilon$. Proceed by induction. This would contradict stability of φ in a countable fragment of the theory.

The converse is not really important, and is pretty standard. ■_{7.11}

From now on we assume that φ is stable.

Definition 7.12. Let $M \models T$ and $A \subseteq M$, and assume that M is $(|A| + \omega)^+$ -saturated and strongly homogeneous. A subset $F \subseteq \text{S}_\varphi(M)$ is *A-good* if it is:

- (i) Metrically compact.
- (ii) Invariant under automorphisms of M fixing A .

Recall the notions of algebraicity and algebraic closure from Definition 4.10.

Lemma 7.13. *Assume that $F \subseteq \text{S}_\varphi(M)$ is A-good. Then every $p \in F$ is definable over $\text{acl}(A)$.*

Proof. We know that p is definable, so let $d\varphi(y, C)$ be its definition, and $c = \text{Cb}_\varphi(p)$ the canonical parameter of the definition. We may write $d\varphi(x, C)$ as $d_p\varphi(y, c)$.

Assume that $c \notin \text{acl}(A)$. Then there exists an infinite sequence $(c_i : i < \omega)$ in $\text{tp}(c/A)$ such that $d(c_i, c_j) \geq \varepsilon > 0$ for all $i < j$, and we can realise this sequence in M by the saturation assumption. By the homogeneity assumption, each $d_p\varphi(y, c_i)$ defines a type p_i which is an A -conjugate of p . Therefore $p_i \in F$ for all $i < \omega$, and $d(p_i, p_j) \geq \varepsilon$ for all $i < j < \omega$, contradicting metric compactness. ■_{7.13}

Lemma 7.14. *Assume that $A \subseteq M \models T$, M is $(|A| + \omega)^+$ -saturated and strongly homogeneous, and $F \subseteq \text{S}_\varphi(M)$ is closed, non-empty, and invariant under $\text{Aut}(M/A)$. Then F contains an A-good subset.*

Proof. Define by induction on n : $F_0 = F$, and $F_{n+1} = \text{CBm}_{\varphi, M, 2^{-n}}(F_n)$. Then $(F_n : n < \omega)$ is a decreasing sequence of non-empty closed subsets of $S_\varphi(M)$, and so $F_\omega = \bigcap_n F_n \neq \emptyset$. The limit set F_ω is ε -finite for every $\varepsilon > 0$, i.e., it is totally bounded. Since the metric refines the topology, F_ω is also metrically closed in $S_\varphi(M)$ and thus complete. We see that f_ω is a totally bounded complete metric space and therefore metrically compact. Also, each of the F_n is invariant under $\text{Aut}(M/A)$, and so is F_ω . $\blacksquare_{7.14}$

Proposition 7.15. *Let $A \subseteq M \models T$, and let $p \in S_\varphi(A)$. Then there exists $q \in S_\varphi(M)$ extending p which is definable over $\text{acl}(A)$.*

Proof. We may replace M with a larger model, so we might as well assume that M is $(|A| + \omega)^+$ -saturated and strongly homogeneous. Let $P = \{q \in S_\varphi(M) : p \subseteq q\}$. By Lemma 7.14, there is an A -good subset $Q \subseteq P$, which is non-empty by definition. By Lemma 7.13, any $q \in Q$ is an $\text{acl}(A)$ -definable extension of p . $\blacksquare_{7.15}$

Proposition 7.16. *Let $M \models T$, $p(x) \in S_\varphi(M)$ and $q(y) \in S_{\tilde{\varphi}}(M)$. Let $d_p\varphi(y)$ and $d_q\tilde{\varphi}(x)$ be their respective definitions, and recall that these are a $\tilde{\varphi}$ - and a φ -predicate, respectively. Then $d_p\varphi(y)^q = d_q\tilde{\varphi}(x)^p$.*

Proof. Let $M_0 = M$. Given $M_n \succeq M$, obtain $p_n \in S_\varphi(M_n)$ and $q_n \in S_{\tilde{\varphi}}(M_n)$ by applying the definition of p and q , respectively, to M_n (these are indeed complete satisfiable φ - and $\tilde{\varphi}$ -types). Realise them by a_n and b_n , respectively, in some extension $M_{n+1} \succeq M_n$. Repeat this for all $n < \omega$.

We now have for all $i < j$: $\varphi(a_j, b_i) = d_p\varphi(y)^q$ and $\varphi(a_i, b_j) = d_q\tilde{\varphi}(x)^p$, and if these differ we get a contradiction to the stability of φ . $\blacksquare_{7.16}$

Proposition 7.17. *Assume that $A \subseteq M$ is algebraically closed, $p, p' \in S_\varphi(M)$ are both definable over A , and $p \upharpoonright_A = p' \upharpoonright_A$. Then $p = p'$.*

Proof. Let $b \in M$, $q = \text{tp}_{\tilde{\varphi}}(b/A)$. By Proposition 7.15 there is $\hat{q} \in S_{\tilde{\varphi}}(M)$ extending q which is definable over $\text{acl}(A) = A$, and let $d_{\hat{q}}\tilde{\varphi}(x)$ be this definition. Recalling that $d_p\varphi$ and $d_{p'}\varphi$ are $\tilde{\varphi}$ -predicates, $d_{\hat{q}}\tilde{\varphi}$ is a φ -predicate, and all of them are over A , we have:

$$\begin{aligned} \varphi(x, b)^p &= d_p\varphi(b) = d_p\varphi(y)^q = d_p\varphi(y)^{\hat{q}} \\ &= d_{\hat{q}}\tilde{\varphi}(x)^p = d_{\hat{q}}\tilde{\varphi}(x)^{p'} \\ &= d_{p'}\varphi(y)^{\hat{q}} = d_{p'}\varphi(y)^q = d_{p'}\varphi(b) \\ &= \varphi(x, b)^{p'}. \end{aligned}$$

Therefore $p = p'$. $\blacksquare_{7.17}$

Given $A \subseteq M \models T$ and $p \in S_\varphi(\text{acl}(A))$, we denote the unique $\text{acl}(A)$ -definable extension of p to M by $p \upharpoonright^M$. The definition of $p \upharpoonright^M$ is an $\text{acl}(A)$ -definable $\tilde{\varphi}$ -predicate which does not depend on M , and we may therefore refer to it unambiguously as $d_p\varphi(y)$ (so far we only used the notation $d_p\varphi(y)$ when p was a φ -type over a model).

If $p \in S_\varphi(A)$, we define (with some abuse of notation) $p \upharpoonright^M = \{q \upharpoonright^M : p \subseteq q \in S_\varphi(\text{acl}(A))\}$.

Proposition 7.18. *If $A \subseteq M \models T$ and $p \in S_\varphi(A)$ as above, then $p \upharpoonright^M$ is closed in $S_\varphi(M)$.*

Assume moreover that M is $(|A| + \omega)^+$ -saturated and strongly homogeneous, and let $P = \{q \in S_\varphi(M) : p \subseteq q\}$. Then $p \upharpoonright^M$ is the unique A -good set contained in P .

Also, we have $p \upharpoonright^M = \bigcap_{\varepsilon > 0} \text{CBm}_{\varphi, M, \varepsilon}(P)$, and in fact $p \upharpoonright^M = \bigcap_{\varepsilon \in E} \text{CBm}_{\varphi, M, \varepsilon}(P)$ for any $E \subseteq (0, \infty)$ such that $\inf E = 0$. In other words, $q \in p \upharpoonright^M$ if and only if $q \in P$, and it has maximal $\text{CB}_{\varphi, M, \varepsilon}$ -rank as such for every $\varepsilon > 0$.

Proof. If $M \preceq M'$, then $p \upharpoonright^M = \{q \upharpoonright_M : q \in p \upharpoonright^{M'}\}$, so we may assume that M is $(|A| + \omega)^+$ -saturated and strongly homogeneous.

Let $Q \subseteq P$ be any A -good subset. If $q \in Q$, then $q = (q \upharpoonright_{\text{acl}(A)}) \upharpoonright^M \in p \upharpoonright^M$. By Lemma 6.8 it follows that $Q = p \upharpoonright^M$. Therefore $p \upharpoonright^M$ is closed.

It follows by Lemma 7.14 that $p \upharpoonright^M \subseteq \text{CBm}_{\varphi, M, \varepsilon}(P)$ for every $\varepsilon > 0$, whereby $\bigcap_{\varepsilon > 0} \text{CBm}_{\varphi, M, \varepsilon}(P) \neq \emptyset$. The other requirements for $\bigcap_{\varepsilon > 0} \text{CBm}_{\varphi, M, \varepsilon}(P)$ to be A -good follow directly from its definition, and we conclude that $p \upharpoonright^M = \bigcap_{\varepsilon > 0} \text{CBm}_{\varphi, M, \varepsilon}(P)$. $\blacksquare_{7.18}$

Proposition 7.19. *Assume that $M \preceq N \models T$ are both ω -saturated. Let $p \in S_\varphi(M)$, and let $q \in S_\varphi(N)$ extend it. Then $\text{CB}_{\varphi, M, \varepsilon}(p) \geq \text{CB}_{\varphi, N, \varepsilon}(q)$, and equality holds for all $\varepsilon > 0$ if and only if $q = p \upharpoonright^N$.*

Proof. Assume first that N is ω_1 -saturated and strongly homogeneous. Let:

$$X_{\varepsilon, \alpha} = \{p \in S_\varphi(M) : \text{CB}_{\varphi, M, \varepsilon}(p) \geq \alpha\} \quad Y_{\varepsilon, \alpha} = \{q \in S_\varphi(N) : \text{CB}_{\varphi, N, \varepsilon}(q) \geq \alpha\}$$

We first prove by induction on α that if $\text{CB}_{\varphi, M, \varepsilon}(p) \leq \alpha$ and $p \subseteq q \in S_\varphi(N)$ then $\text{CB}_{\varphi, N, \varepsilon}(q) \leq \alpha$. Given a φ -predicate $\psi(x, a)$ with parameters $a \in M$ and $r \in [0, 1]$, let us write:

$$[\psi(x, a) < r]^M = [\psi(x, a) < r]^{S_\varphi(M)} = \{p' \in S_\varphi(M) : \psi(x, a)^{p'} < r\}.$$

Sets of this form form a basis of open sets for $S_\varphi(M)$. Since $\text{CB}_{\varphi, M, \varepsilon}(p) \leq \alpha$, there are such $\psi(x, a)$ and r such that $p \in [\psi(x, a) < r]^M$ and $\text{diam}([\psi(x, a) < r]^M \cap X_{\varepsilon, \alpha}) \leq \varepsilon$.

Clearly, $q \in [\psi(x, a) < r]^N$, so we'll be done if we prove that $\text{diam}([\psi(x, a) < r]^N \cap Y_{\varepsilon, \alpha}) \leq \varepsilon$ as well. Indeed, assume that there are $q', q'' \in [\psi(x, a) < r]^N$ such that $d(q', q'') > \varepsilon$. Let $d\varphi(y, e')$ and $d\varphi(y, e'')$ be their respective definitions, where e' and e'' are the canonical parameters. Then we can find $f', f'', b \in M$ such that $f'f''b \equiv e'e''a$ and $d(a, b)$ is as small as we want (in fact using ω -saturation we can actually have $a = b$, but the argument goes through even if we can only have b arbitrarily close to a ; therefore the result is true even if M is merely *approximately* ω -saturated, as defined in [Ben05] or [BU07]).

Let $p', p'' \in S_\varphi(M)$ be defined by $d\varphi(y, f')$ and $d\varphi(y, f'')$, respectively. Then $p', p'' \in [\psi(x, b) < r]^M$, and having made sure that b is close enough to a , we can get $p', p'' \in [\psi(x, a) < r]^M$. Also, we still have $d(p', p'') > \varepsilon$. Therefore at least one of $p' \notin X_{\varepsilon, \alpha}$ or

$p'' \notin X_{\varepsilon, \alpha}$ must hold, so let's say it's the former. In this case, by the induction hypothesis, $p' \upharpoonright^N \notin Y_{\varepsilon, \alpha}$. Since N is ω_1 -strongly homogeneous, $p' \upharpoonright^N$ and q' are conjugates by $\text{Aut}(N)$, so $q' \notin Y_{\varepsilon, \alpha}$, as required.

Now let $q = p \upharpoonright^N$, and let $d\varphi(y, e)$ be the common definition. Assume that $\text{CB}_{\varphi, N, \varepsilon}(q) \leq \alpha$, so there are a φ -predicate $\psi(x, a)$ with $a \in N$, and r , such that $q \in [\psi(x, a) < r]^N$ and $\text{diam}([\psi(x, a) < r]^N \cap Y_{\alpha, \varepsilon}) \leq \varepsilon$. We may find $b, f \in M$ such that $bf \equiv ae$, and such that f is as close as we want to e (again: if we take the plain definition of ω -saturation we can even have $e = f$, but we want an argument that goes through if M is only approximately ω -saturated). If $p' \in S_\varphi(M)$ is defined by $d\varphi(y, f)$, then $p' \in [\psi(x, b) < r]^M$, and assuming f and e are close enough we also have $p \in [\psi(x, b) < r]^M$. By the homogeneity assumption for N we get $\text{diam}([\psi(x, b) < r]^N \cap Y_{\alpha, \varepsilon}) \leq \varepsilon$.

Assume now that $p'', p''' \in [\psi(x, b) < r]^M$, and $d(p'', p''') > \varepsilon$. Let $q'' = p'' \upharpoonright^N$ and $q''' = p''' \upharpoonright^N$. Then $d(q'', q''') > \varepsilon$, so either $q'' \notin Y_{\varepsilon, \alpha}$ or $q''' \notin Y_{\varepsilon, \alpha}$ (or both), so let's say it's the former. By the induction hypothesis we get $p'' \notin X_{\varepsilon, \alpha}$, showing that $\text{diam}([\psi(x, b) < r]^M \cap X_{\alpha, \varepsilon}) \leq \varepsilon$. Therefore $\text{CB}_{\varphi, M, \varepsilon}(p) \leq \alpha$.

Since $p \upharpoonright^N$ is the unique extension of p to N having maximal $\text{CB}_{\varphi, N, \varepsilon}$ rank for all $\varepsilon > 0$, this proves what we wanted.

If N is not ω_1 -strongly homogeneous, take a common elementary extension for M and N which is. ■ 7.19

Corollary 7.20. *Again, let $M \preceq N$ be ω -saturated. Let $X \subseteq S_\varphi(M)$ be any set of φ -types (without any further assumptions), and let $Y \subseteq S_\varphi(N)$ be its pre-image under the restriction mapping. Then $\text{CB}_{\varphi, M, \varepsilon}(X) = \text{CB}_{\varphi, N, \varepsilon}(Y)$.*

In particular, this gives us an absolute notion of $\text{CB}_{\varphi, \varepsilon}(p)$ where p is a partial φ -type, without specifying over which (ω -saturated) model we work: just calculate it in any ω -saturated model containing the parameters for p .

For example, assume that $A \subseteq B \subseteq M \models T$, $p \in S_\varphi(A)$ and $p \subseteq q \in S_\varphi(B)$. Then q is definable over $\text{acl}(A)$ if and only if $\text{CB}_{\varphi, \varepsilon}(p) = \text{CB}_{\varphi, \varepsilon}(q)$ for all $\varepsilon > 0$.

8. GLOBAL STABILITY AND INDEPENDENCE

In the previous section we only considered local stability, i.e., stability of a single formula $\varphi(x, y)$. In this section we will use those results to deduce a global stability theory.

8.1. Gluing local types. Let A be an algebraically closed set, and say $A \subseteq M$. Let $\varphi(x, y)$ and $\psi(x, z)$ be two stable formulae, $p_\varphi \in S_\varphi(A)$, $p_\psi \in S_\psi(A)$. Then we know that each of p_φ and p_ψ have unique extensions $q_\varphi \in S_\varphi(M)$ and $q_\psi \in S_\psi(M)$, respectively, which are A -definable.

Assume now that p_φ and p_ψ are compatible, i.e., that $p_\varphi(x) \cup p_\psi(x)$ is satisfiable. We would like to show that q_φ and q_ψ are compatible as well. For this purpose there is no

harm in assuming that M is strongly $(|A| + \omega)^+$ -homogeneous, or even that $M = \mathfrak{M}$ is our monster model.

Let t, w be any variables in a single sort, say the home sort, and $e \neq e' \in M$ in that sort. We may assume that $d(e, e') = 1$: even if not, everything we do below would work when we replace $d(t, w)$ with $d(t, w) \dot{+} \dots \dot{+} d(t, w)$. Define:

$$\chi_{\varphi, \psi}(x, yztw) = \varphi(x, y) \wedge d(t, w) \dot{+} \psi(x, z) \wedge \neg d(t, w).$$

Since we assume that φ and ψ are stable so is $\chi_{\varphi, \psi}(x, yztw)$ by the following easy result:

Lemma 8.1. *Assume $\varphi_i(x, y)$ are stable formulae for $i < n$ and f is an n -ary continuous connective. Then $(f \circ \varphi_{<n})(x, y)$ is stable as well.*

Let a, b and c be in the appropriate sorts. Then:

$$\begin{aligned} \varphi(a, b) &= \chi_{\varphi, \psi}(a, bcee'), \\ \psi(a, c) &= \chi_{\varphi, \psi}(a, bcee'). \end{aligned}$$

Thus every instance of φ and ψ is an instance of $\chi_{\varphi, \psi}$, so every φ -predicate or ψ -predicate (with parameters in M) is a $\chi_{\varphi, \psi}$ -predicate. Moreover, if $B \subseteq M$ and $\rho(x)$ is B -definable as a φ -predicate (or ψ -predicate) then it is B -invariant and therefore a B -definable $\chi_{\varphi, \psi}$ -predicate. Notice that there is no need to assume here that $e, e' \in B$.

We therefore obtain for every set $B \subseteq M$ quotient mappings $\theta_\varphi: S_{\chi_{\varphi, \psi}}(B) \rightarrow S_\varphi(B)$ and $\theta_\psi: S_{\chi_{\varphi, \psi}}(B) \rightarrow S_\psi(B)$, and if $B \subseteq C$ then the following diagram commutes:

$$\begin{array}{ccc} S_{\chi_{\varphi, \psi}}(C) & \longrightarrow & S_\varphi(C) \times S_\psi(C) \\ \downarrow & & \downarrow \\ S_{\chi_{\varphi, \psi}}(B) & \longrightarrow & S_\varphi(B) \times S_\psi(B) \end{array}$$

Let us now return to the situation we started with, namely $A \subseteq M$ algebraically closed and a pair of compatible $p_\varphi \in S_\varphi(A)$ and $p_\psi \in S_\psi(A)$. Since they are compatible, there exists $p_{\chi_{\varphi, \psi}} \in S_{\chi_{\varphi, \psi}}(A)$ such that $p_\varphi = \theta_\varphi(p_{\chi_{\varphi, \psi}})$, $p_\psi = \theta_\psi(p_{\chi_{\varphi, \psi}})$ (p_φ and p_ψ actually determine $p_{\chi_{\varphi, \psi}}$, but we do not need this fact).

Let $q_{\chi_{\varphi, \psi}} \in S_{\chi_{\varphi, \psi}}(M)$ be the unique extension of $p_{\chi_{\varphi, \psi}}$ which is A -definable, and let $q'_\varphi = \theta_\varphi(q_{\chi_{\varphi, \psi}}) \in S_\varphi(M)$, $q'_\psi = \theta_\psi(q_{\chi_{\varphi, \psi}})$. Then q'_φ is definable over M (by stability of φ) and is invariant under $\text{Aut}(M/A)$ (since $q_{\chi_{\varphi, \psi}}$ is). Since M is strongly $(|A| + \omega)^+$ -homogeneous, it follows that the definition $d_{q'_\varphi}(\varphi(y))$ is A -invariant, and therefore A -definable. By the commutativity of the diagram above (with $B = A$, $C = M$) q'_φ extends p_φ . Therefore, by uniqueness, $q_\varphi = q'_\varphi$. Similarly, $q_\psi = q'_\psi$, so $q_\varphi \cup q_\psi \subseteq q_{\chi_{\varphi, \psi}}$ is satisfiable.

8.2. Global stability.

Definition 8.2. A theory T is stable if all formulae are stable in T .

Definition 8.3. A theory T is λ -stable if for all $n < \omega$ and all sets A such that $|A| \leq \lambda$: $\|S_n(A)\| \leq \lambda$.

Definition 8.4. Let M be a model and $p \in S_n(M)$. We say that p is *definable* if $p \upharpoonright_\varphi$ is for every formula of the form $\varphi(x_{<n}, \bar{y})$, i.e., if for every such formula there is an M -definable predicate $d_p\varphi(\bar{y})$, called the φ -*definition* of p , such that for all $\bar{b} \in M$:

$$\varphi(\bar{x}, \bar{b})^p = d_p\varphi(\bar{b}).$$

Theorem 8.5. *The following are equivalent for a theory T :*

- (i) T is stable.
- (ii) All types over models are definable.
- (iii) T is λ -stable for all λ such that $\lambda = \lambda^{|T|}$.
- (iv) T is λ -stable for some $\lambda \geq |T|$.

Proof. (i) \iff (ii). By Proposition 7.7.

(i) \implies (iii). Assume T is stable $\lambda = \lambda^{|T|}$, $n < \omega$, and $|A| \leq \lambda$. Then by Downward Löwenheim-Skolem we can find $M \supseteq A$ such that $\|M\| \leq \lambda$. Let $\varphi(x_{<n}, \bar{y})$ be any formula. Then by Proposition 7.7 we have $\|S_\varphi(M)\| \leq \lambda$, whereby $|S_\varphi(M)| \leq \lambda^\omega = \lambda$. Let $\{\varphi_i(x_{<n}, \bar{y}) : i < |T|\}$ enumerate all formulae of this form. Then $|S_n(M)| \leq \prod |S_{\varphi_i}(M)| \leq \lambda^{|T|} = \lambda$, and *a fortiori* $\|S_n(M)\| \leq \lambda$.

(iii) \implies (iv). Let $\lambda = 2^{|T|}$.

(iv) \implies (i). Let T be λ -stable ($\lambda \geq |T|$) and $\varphi(x_{<n}, \bar{y})$ be any formula, and we will show that φ is stable in T . Let M be any model such that $\|M\| \leq \lambda$, so let $M_0 \subseteq M$ be a dense subset such that $|M_0| = \lambda$. Then $S_n(M) = S_n(M_0)$ (i.e., the quotient mapping $S_n(M) \rightarrow S_n(M_0)$ is a homeomorphism and an isometry). Since φ is uniformly continuous, the quotient mapping $S_n(M) \rightarrow S_\varphi(M)$ is uniformly continuous as a mapping between metric spaces. Therefore $\|S_\varphi(M)\| \leq \|S_n(M)\| = \|S_n(M_0)\| \leq \lambda$.

Since this holds for all M such that $\|M\| \leq \lambda$, we conclude by Proposition 7.7 that φ is stable. ■_{8.5}

Convention 8.6. From now on we assume T is stable.

Proposition 8.7. *Let $A \subseteq M$, where A is algebraically closed, and let $p(x) \in S_x(A)$. Then p has a unique extension to M , denoted $p \upharpoonright^M$, which is A -definable. Moreover, the A -definable definitions of such extensions do not depend on M , and will be denoted as usual by $d_p\varphi$.*

Proof. For every formula $\varphi(x, y)$ let $d_p\varphi = d_{p \upharpoonright_\varphi}\varphi$. Then uniqueness and moreover part are already a consequence of Proposition 7.17. Thus all that is left to show is that the following set of conditions is satisfiable (and therefore a complete type):

$$\begin{aligned} p \upharpoonright^M &= \{\varphi(x, b) = d_p\varphi(x, b)^M : \varphi(x, y) \in \mathcal{L}, b \in M \text{ in the sort of } y\} \\ &= \bigcup_{\varphi(x, y) \in \mathcal{L}} (p \upharpoonright_\varphi) \upharpoonright^M. \end{aligned}$$

(Here x is fixed but y varies with φ .)

By compactness it suffices to show this for unions over finitely many formulae φ . For two formulae this was proved in the previous subsection, by coding both formulae in a single one. But we can repeat this process encoding any finite set of formulae in a single one, whence the required result. $\blacksquare_{8.7}$

Definition 8.8. Let $A \subseteq B$, $p \in S(B)$. We say that p *does not fork* over A if there exists an extension $p \subseteq q \in S_n(\text{acl}(B))$ such that all the definitions $d_q\varphi$ are over $\text{acl}(A)$.

If \bar{a} is a tuple, A and B sets, and $\text{tp}(\bar{a}/AB)$ does not fork over A , we say \bar{a} is *independent* from B over A , in symbols $\bar{a} \downarrow_A B$.

Corollary 8.9. Let $A \subseteq B$, where A is algebraically closed, and let $p \in S_n(A)$. Then there exists a unique $q \in S_n(B)$ extending p and non-forking over A . This unique non-forking extension is denoted $p \upharpoonright^B$, and is given explicitly as

$$p \upharpoonright^B = \{\varphi(\bar{x}, \bar{b}) = d_p\varphi(\bar{b}) : \varphi(\bar{x}, \bar{y}) \in \mathcal{L}, \bar{b} \in B\}.$$

Proof. Let M be any model such that $B \subseteq M$. Then $\text{acl}(B) \subseteq M$, and $(p \upharpoonright^M) \upharpoonright_{\text{acl}(B)} = p \upharpoonright^{\text{acl}(B)}$ is A -definable, so $p \upharpoonright^B$ is a non-forking extension of p .

Conversely, let $q \in S_n(B)$ be a non-forking extension of p . Then there exists $q' \in S_n(\text{acl}(B))$ which is A -definable. Then $q' \upharpoonright^M$ is an A -definable extension of p , so $q' \upharpoonright^M = p \upharpoonright^M$, whereby $q = p \upharpoonright^B$. $\blacksquare_{8.9}$

Corollary 8.10. Let A and B be sets, \bar{a} a tuple, and let $p = \text{tp}(\bar{a}/\text{acl}(A))$ and $q = \text{tp}(\bar{a}/\text{acl}(AB))$. Then $\bar{a} \downarrow_A B$ if and only if $d_p\varphi = d_q\varphi$ for every formula $\varphi(\bar{x}, \bar{y})$.

Proof. Right to left is immediate from the definition. So assume $\bar{a} \downarrow_A B$. This means there is a type $q' \in S_n(\text{acl}(AB))$ extending $\text{tp}(\bar{a}/AB)$ such that $d_{q'}\varphi$ is $\text{acl}(A)$ -definable for all φ . Then q and q' are conjugates by an automorphism fixing AB . Such an automorphism would fix $\text{acl}(A)$ setwise, so $d_q\varphi$ is $\text{acl}(A)$ -definable for all φ . Now let $M \supseteq AB$ be a model, and $r = q \upharpoonright^M$. Then $d_q\varphi = d_r\varphi$ by definition, and r is $\text{acl}(A)$ -definable and extends p , whereby $d_p\varphi = d_r\varphi$. $\blacksquare_{8.10}$

We conclude:

Theorem 8.11. Assume T is stable. Then:

- (i) Invariance: The relation \downarrow is automorphism-invariant.
- (ii) Symmetry: $\bar{a} \downarrow_A \bar{b} \iff \bar{b} \downarrow_A \bar{a}$
- (iii) Transitivity: $\bar{a} \downarrow_A BC$ if and only if $\bar{a} \downarrow_A B$ and $\bar{a} \downarrow_{AB} C$.
- (iv) Existence: For all \bar{a} , A and B there is $\bar{b} \equiv_A \bar{a}$ such that $\bar{b} \downarrow_A B$.
- (v) Finite character: $\bar{a} \downarrow_A B$ if and only if $\bar{a} \downarrow_A \bar{b}$ for all finite tuples $\bar{b} \in B$.
- (vi) Local character: For all \bar{a} and A there is $A_0 \subseteq A$ such that $|A_0| \leq |T|$ and $\bar{a} \downarrow_{A_0} A$.

- (vii) Stationarity: Assume A is algebraically closed, and $B \supseteq A$. If $\bar{a} \equiv_A \bar{b}$ and $\bar{a} \downarrow_A B$, $\bar{b} \downarrow_A B$ then $\bar{a} \equiv_{AB} \bar{b}$.

Proof. Invariance is clear.

Symmetry follows from Proposition 7.16.

Transitivity is immediate from Corollary 8.10.

For existence, we may replace B with any model containing B . Then let p' be any extension of $p = \text{tp}(\bar{a}/A)$ to $\text{acl}(A)$, and then let \bar{b} realise the unique non-forking extension of p' to M .

For finite character, we may replace A with $\text{acl}(A)$ without changing the statement. But then $\bar{a} \downarrow_A B$ if and only if $\text{tp}(\bar{a}/AB) = \text{tp}(\bar{a}/A) \upharpoonright^{AB}$, and if this fails it is due to some finite tuple in B .

Let $p = \text{tp}(\bar{a}/\text{acl}(A))$. Recall we defined $\text{Cb}_\varphi(p)$ as the canonical parameter of $d_p\varphi$. Let $\text{Cb}(p) = \{\text{Cb}_\varphi(p) : \varphi(\bar{x}, \dots) \in \mathcal{L}\}$ be the *canonical base* of p . Then $|\text{Cb}(p)| = |T|$, and p is definable over its canonical base so $\bar{a} \downarrow_{\text{Cb}(p)} \text{acl}(A)$. For each $c \in \text{Cb}(p)$ we know that $\text{tp}(c/A)$ is algebraic, and going back to the definition of algebraicity in Lemma 4.9 we see that countably many parameters in A suffice to witness this. Let A_0 be the union of all these witness sets for all $c \in \text{Cb}(p)$. Then $A_0 \subseteq A$, $|A_0| \leq |T|$, and $\text{Cb}(p) \subseteq \text{acl}(A_0)$, so $\bar{a} \downarrow_{\text{acl}(A_0)} \text{acl}(A)$, or equivalently, $\bar{a} \downarrow_{A_0} A$.

Stationarity is just Corollary 8.9. ■_{8.11}

APPENDIX A. A REMARK ON CONTINUITY MODULI ON BOUNDED SPACES

The usual definition of (uniform) continuity in the metric setting goes “for all $\varepsilon > 0$ there is $\delta > 0$ such that...”, whence our definition of a continuity modulus as a function $\delta: (0, \infty) \rightarrow (0, \infty)$, mapping each ε to a corresponding δ . We would like to present here an alternative definition, which rather goes the other way around.

In the cases which interest us all metric spaces (structures and type spaces) are bounded, usually of diameter ≤ 1 . We may therefore allow ourselves the following simplification:

Convention A.1. Hereafter, all metric spaces are bounded of diameter ≤ 1 .

Definition A.2. An *inverse continuity modulus* is a continuous monotone function $\mathbf{u}: [0, 1] \rightarrow [0, 1]$ such that $\mathbf{u}(0) = 0$.

We say that a mapping $f: (X, d) \rightarrow (X', d')$ *respects* \mathbf{u} , or that it is *uniformly continuous with respect to* \mathbf{u} , if for every $x, y \in X$:

$$d'(f(x), f(y)) \leq \mathbf{u}(d(x, y)).$$

In other words, an inverse uniform continuous modulus maps a δ to an ε . (In case that the destination space is not bounded we may still consider inverse continuity moduli, but then we need to allow the range of \mathbf{u} to be $[0, \infty]$.)

Lemma A.3. *Let \mathbf{u} be an inverse continuity modulus. For $\varepsilon > 0$ define $\delta(\varepsilon) = \sup\{t \in [0, 1] : \mathbf{u}(t) \leq \varepsilon\}$. Then δ is a continuity modulus, and every function which respects \mathbf{u} (as an inverse uniform continuity modulus) respects δ (as a uniform continuity modulus).*

Proof. That $\varepsilon > 0 \implies \delta(\varepsilon) > 0$ follows from the fact that $\mathbf{u}(t) \rightarrow 0$ as $t \rightarrow 0$. Assume now that $f: (X, d) \rightarrow (X', d')$ respects \mathbf{u} , $\varepsilon > 0$, and $d(x, y) < \delta(\varepsilon)$. By monotonicity of \mathbf{u} and definition of δ :

$$d'(f(x), f(y)) \leq \mathbf{u}(d(x, y)) \leq \varepsilon.$$

■_{A.3}

The converse is not much more difficult:

Lemma A.4. *Let δ be a continuity modulus. Then there exists an inverse continuity modulus \mathbf{u} such that every function respecting \mathbf{u} respects δ .*

Proof. For $r, r' \in [0, 1]$ define:

$$\begin{aligned} \mathbf{u}_0(r) &= \inf\{\varepsilon > 0 : \delta(\varepsilon) > r\} && (\inf \emptyset = 1) \\ \mathbf{u}_1(r, r') &= \begin{cases} \mathbf{u}_0(r') & r \geq r' \\ \mathbf{u}_0(r') \cdot \left(\frac{2r}{r'} - 1\right) & \frac{r'}{2} \leq r < r' \\ 0 & r < \frac{r'}{2} \end{cases} \\ \mathbf{u}(r) &= \sup_{r' \in [0, 1]} \mathbf{u}_1(r, r'). \end{aligned}$$

Then $\mathbf{u}_0: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is an increasing function, not necessarily continuous. It is however continuous at 0: $\lim_{r \rightarrow 0^+} \mathbf{u}_0(r) = 0 = \mathbf{u}_0(0)$ (since for every $\varepsilon > 0$, $\mathbf{u}_0(\delta(\varepsilon)) \leq \varepsilon$). For every $r_0 > 0$, the family of function $r \mapsto \mathbf{u}_1(r, r')$, indexed by r' , is equally continuous on $[r_0, 1]$, so \mathbf{u} is continuous on $(0, 1]$. We also have $r \leq 1/2 \implies \mathbf{u}(r) \leq \mathbf{u}_0(2r)$ (since for $r' \geq 2r$, $\mathbf{u}_1(r, r')$ contributes nothing to $\mathbf{u}(r)$), whereby $\lim_{r \rightarrow 0^+} \mathbf{u}(r) = 0 = \mathbf{u}(0)$. Therefore \mathbf{u} is continuous on $[0, 1]$, and therefore an inverse continuity modulus. Observe also that:

$$\mathbf{u}(r) \geq \mathbf{u}_1(r, r) = \mathbf{u}_0(r) \geq \sup\{\varepsilon \leq 1 : (\forall 0 < \varepsilon' < \varepsilon)(\delta(\varepsilon') \leq r)\}.$$

Assume now that $f: (X, d) \rightarrow (X', d')$ respects δ . If $x, y \in X$ and $\varepsilon > 0$ satisfy $d'(f(x), f(y)) > \varepsilon$, Then $d(x, y) \geq \delta(\varepsilon')$ for all $0 < \varepsilon' < \varepsilon$, whereby $\mathbf{u}(d(x, y)) \geq \varepsilon$. Therefore $d'(f(x), f(y)) \leq \mathbf{u}(d(x, y))$, and f respects \mathbf{u} , as required. ■_{A.4}

Together we obtain:

Theorem A.5. *A mapping between bounded metric spaces $f: (X, d) \rightarrow (X', d')$ is uniformly continuous with respect to a (standard) continuity modulus if and only if it is uniformly continuous with respect to an inverse one. In other words, the two distinct definitions of continuity moduli give rise to the same notion of uniform continuity.*

Inverse continuity moduli give us (continuously) a direct answer to the question “how much can the value of f change from x to y ?” For example, if we attached to symbols in a signature inverse continuity moduli, rather than usual ones, the axiom scheme $\text{UC}_{\mathcal{L}}$ would take the more elegant form:

$$\begin{aligned} (\text{UC}_{\mathcal{L}}^u) \quad & \sup_{x < i, y < n-i-1, z, w} d(f(\bar{x}, z, \bar{y}), f(\bar{x}, w, \bar{y})) \dot{-} \mathbf{u}_{f,i}(d(z, w)) = 0 \\ & \sup_{x < i, y < n-i-1, z, w} |P(\bar{x}, z, \bar{y}) - P(\bar{x}, w, \bar{y})| \dot{-} \mathbf{u}_{P,i}(d(z, w)) = 0. \end{aligned}$$

Since the continuity moduli $\mathbf{u}_{s,i}$ are continuous functions, they can be admitted as connectives in the language, so the above can indeed be viewed as sentences. In fact, in almost all actual cases, the inverse continuity moduli can be directly constructed from the standard connectives $\{\neg, \dot{-}, \frac{x}{2}\}$, so there is no need to introduce new connectives (for example, in the case of probability algebras, all the inverse continuity moduli can be taken to be the identity).

APPENDIX B. ON STABILITY INSIDE A MODEL

This appendix answers a question posed by C. Ward Henson to the first author concerning stability of a formula inside a specific structure (in contrast with stability of a formula in all models of a theory, discussed in Section 7 above). The notion of stability inside a model appears for example in the work of Krivine and Maurey on stable Banach spaces [KM81]: a Banach space E is stable in the sense of Krivine and Maurey precisely if the formula $\|x + y\|$ is stable in the unit ball of E (viewed as a continuous structure in an appropriate language) in the sense defined below.

Definition B.1. Let M be a structure, $\varphi(x, y)$ be a formula and $\varepsilon > 0$. We say that φ is ε -stable in M if there is no sequence $(a_i b_i : i < \omega)$ in M such that $|\varphi(a_i, b_j) - \varphi(a_j, b_i)| \geq \varepsilon$ for all $i < j < \omega$. We say that φ is stable in M if it is ε -stable in M for all $\varepsilon > 0$.

Note that $\varphi(x, y)$ is (ε) -stable in M if and only if $\tilde{\varphi}(y, x)$ is. Also, φ is (ε) -stable in a theory T if and only if it is in every model of T .

Lemma B.2. Assume that $\varphi(x, y)$ is ε -stable in M . Then for every $p \in S_{\varphi}(M)$ there exists a finite sequence $(c_i : i < n)$ in M such that for all $a, b \in M$:

$$(\forall i < n)(\varphi(c_i, a) \leq \varphi(c_i, b) + \varepsilon) \implies \varphi(x, a)^p \leq \varphi(x, b)^p + 3\varepsilon.$$

Proof. Assume not. We will choose by induction on n elements $a_n, b_n, c_n \in M$ and $r_n, s_n \in [0, 1]$ as follows. At each step, there are by assumption $a_n, b_n \in M$ such that $\varphi(c_i, a_n) \leq \varphi(c_i, b_n) + \varepsilon$ for all $i < n$, and yet $\varphi(x, a_n)^p > \varphi(x, b_n)^p + 3\varepsilon$. Choose r_n, s_n such that $\varphi(x, b_n)^p < r_n < r_n + 3\varepsilon < s_n < \varphi(x, a_n)^p$. Once these choices are made we have $\varphi(x, a_i)^p > s_i$ and $\varphi(x, b_i)^p < r_i$ for all $i \leq n$, and we may therefore find $c_n \in M$ such that $\varphi(c_n, a_i) > s_i$ and $\varphi(c_n, b_i) < r_i$ for all $i \leq n$.

Once the construction is complete, for every $i < j < \omega$ colour the pair $\{i, j\}$ as follows: if $s_i > \varphi(c_i, a_j) + \varepsilon$, colour the pair $\{i, j\}$ yellowish maroon; otherwise, colour it fluorescent pink. Notice that if $\{i, j\}$ is fluorescent pink then $\varphi(c_i, b_j) - \varepsilon > r_i$. By Ramsey's Theorem there is an infinite monochromatic subset $I \subseteq \omega$, and without loss of generality $I = \omega$. If all pairs are fluorescent pink then we have for all $i < j < \omega$: $\varphi(c_j, a_i) - \varphi(c_i, a_j) > s_i - (s_i - \varepsilon) = \varepsilon$. If all are yellowish maroon we get $\varphi(c_i, b_j) - \varphi(c_j, b_i) > (r_i + \varepsilon) - r_i = \varepsilon$. Either way, we get a contradiction to ε -stability in M . $\blacksquare_{\text{B.2}}$

Lemma B.3. *Assume that $\varphi(x, y)$ is ε -stable in M . Then for every $p \in S_\varphi(M)$ there exists a finite sequence $(c_i: i < n)$ in M and a continuous increasing function $f: [0, 1]^n \rightarrow [0, 1]$, such that for all $a \in M$:*

$$|\varphi(x, a)^p - f(\varphi(c_i, a): i < n)| \leq 3\varepsilon.$$

Proof. Let $(c_i: i < n)$ be chosen as in the previous Lemma. As a first approximation, let:

$$f(\bar{u}) = \sup\{\varphi(x, a)^p: a \in M \text{ and } \varphi(c_i, a) \leq u_i \text{ for all } i < n\}.$$

This function is increasing, but not necessarily continuous. We define a family of auxiliary functions $h_{\bar{u}}: [0, 1]^n \rightarrow [0, 1]$ for $\bar{u} \in [0, 1]^n$:

$$h_{\bar{u}}(\bar{v}) = \frac{1}{\varepsilon} \bigwedge_{i < n} (((v_i + \varepsilon) \dot{-} u_i) \wedge \varepsilon).$$

In other words, $h_{\bar{u}}(\bar{v})$ is a linear-by-pieces function which is equal to 1 if $v_i \geq u_i$ for all i , to 0 if $v_i \leq u_i - \varepsilon$ for some i , and is linear in between. We now define:

$$g(\bar{v}) = \sup_{\bar{u} \in [0, 1]^n} h_{\bar{u}}(\bar{v}) f(\bar{u}).$$

Since all the functions $h_{\bar{u}}$ are equally continuous, g is continuous. It is clearly increasing, and also satisfies for all $\bar{v} \in [0, 1]^n$:

$$f(\bar{v}) \leq g(\bar{v}) \leq f(\bar{v} \dot{+} \varepsilon).$$

Indeed, the first inequality is clear from the definition, and the second follows from the fact that f is increasing, and every tuple \bar{u} such that $f(\bar{u})$ contributes to $g(\bar{v})$ must be smaller in every coordinate than $\bar{v} \dot{+} \varepsilon$.

Now let $a \in M$ and $v_i = \varphi(c_i, a)$ for $i < n$. Then by choice of \bar{c} :

$$\begin{aligned} g(\bar{v}) &\leq f(\bar{v} \dot{+} \varepsilon) \\ &= \sup\{\varphi(x, b)^p: b \in M \text{ and } \varphi(c_i, b) \leq \varphi(c_i, a) + \varepsilon \text{ for all } i < n\} \\ &\leq \varphi(x, a)^p + 3\varepsilon, \end{aligned}$$

and

$$\begin{aligned} g(\bar{v}) &\geq f(\bar{v}) \\ &= \sup\{\varphi(x, b)^p : b \in M \text{ and } \varphi(c_i, b) \leq \varphi(c_i, a) \text{ for all } i < n\} \\ &\geq \varphi(x, a)^p. \end{aligned}$$

Therefore $|g(\varphi(c_i, a) : i < n) - \varphi(x, a)^p| \leq 3\varepsilon$. ■_{B.3}

Theorem B.4. *Assume φ is stable in M . Then every $p \in S_\varphi(M)$ is definable. Moreover, for every such p there is a sequence $(c_i : i < \omega)$ and a continuous increasing function $f : [0, 1]^\omega \rightarrow [0, 1]$ such that $d_p\varphi(y) = f \circ (\varphi(c_i, y) : i < \omega)$.*

Proof. For all $m < \omega$ choose a sequence $(c_{m,i} : i < n_m)$ and function $f_m : [0, 1]^{n_m} \rightarrow [0, 1]$ as in the previous Lemma corresponding to $\varepsilon = 2^{-m-17}$. Let:

$$\begin{aligned} N_m &= \sum_{k < m} n_k, \\ d_{N_m+i} &= c_{m,i} && i < n_m, \\ f(u_{<\omega}) &= \mathcal{F}\lim_m f_m(u_{N_m}, \dots, u_{N_{m+1}-1}). \end{aligned}$$

Since each f_m is increasing and continuous, as is $\mathcal{F}\lim : [0, 1]^\omega \rightarrow [0, 1]$, we have that f is increasing and continuous. Also, by the choice of the parameters we have for all $a \in M$:

$$f(\varphi(d_i, a) : i < \omega) = \mathcal{F}\lim_m f_m(\varphi(c_i, a) : i < n_m) = \varphi(x, a)^p. \quad \blacksquare_{B.4}$$

Notice that we get almost the same result as for a formula which is stable in a theory: the definition is still a limit of positive (i.e., increasing) continuous combinations of instances of φ with parameters in M . However, these combinations are not necessarily the particularly elegant median value as in Section 7.

REFERENCES

- [BBH] Itai Ben Yaacov, Alexander Berenstein, and C. Ward Henson, *Model-theoretic independence in the Banach lattices $L^p(\mu)$* , submitted.
- [Bena] Itai Ben Yaacov, *Definability of groups in \aleph_0 -stable metric structures*, submitted.
- [Benb] ———, *Modular functionals and perturbations of Nakano spaces*, submitted.
- [Benc] ———, *On perturbations of continuous structures*, submitted.
- [Bend] ———, *Topometric spaces and perturbations of metric structures*, submitted.
- [Ben03] ———, *Positive model theory and compact abstract theories*, *Journal of Mathematical Logic* **3** (2003), no. 1, 85–118.
- [Ben05] ———, *Uncountable dense categoricity in cats*, *Journal of Symbolic Logic* **70** (2005), no. 3, 829–860.
- [BU07] Itai Ben Yaacov and Alexander Usvyatsov, *On d -finiteness in continuous structures*, *Fundamenta Mathematicæ* **194** (2007), 67–88.
- [CK66] C. C. Chang and H. Jerome Keisler, *Continuous model theory*, Princeton University Press, 1966.

- [Fre04] D. H. Fremlin, *Measure theory volume 3: Measure algebras*, Torres Fremlin, 25 Ireton Road, Colchester CO3 3AT, England, 2004, <http://www.essex.ac.uk/maths/staff/fremlin/mt3.2004/index.htm>.
- [HI02] C. Ward Henson and José Iovino, *Ultraproducts in analysis*, Analysis and Logic (Catherine Finet and Christian Michaux, eds.), London Mathematical Society Lecture Notes Series, no. 262, Cambridge University Press, 2002.
- [KM81] Jean-Louis Krivine and Bernard Maurey, *Espaces de Banach stables*, Israel Journal of Mathematics **39** (1981), no. 4, 273–295.
- [Mac73] Hilton Vieira Machado, *A characterization of convex subsets of normed spaces*, Kōdai Mathematical Seminar Reports **25** (1973), 307–320.
- [Pil96] Anand Pillay, *Geometric stability theory*, Clarendon Press, 1996.
- [SU] Saharon Shelah and Alexander Usvyatsov, *Model theoretic stability and categoricity for complete metric spaces*, preprint.

ITAÏ BEN YAACOV, UNIVERSITÉ DE LYON, UNIVERSITÉ LYON 1, INSTITUT CAMILLE JORDAN, UMR 5208 CNRS, 43 BOULEVARD DU 11 NOVEMBRE 1918, F-69622 VILLEURBANNE CEDEX, FRANCE
URL: <http://math.univ-lyon1.fr/~begnac/>

ALEXANDER USVYATSOV, UCLA MATHEMATICS DEPARTMENT, BOX 951555, LOS ANGELES, CA 90095-1555, USA
URL: <http://www.math.ucla.edu/~alexus>