Strong conceptual completeness and various stability-theoretic results in continuous model theory

By

JEAN-MARTIN ALBERT, B.Sc, M.Sc

A Thesis Submitted to the School of Graduate Studies in Partial Fulfilment of the Requirements for the Degree Doctor of Philosophy

McMaster University

DOCTOR OF PHILOSOPHY (2010) | MCMASTER UNIVERSITY

Mathematical Logic | Hamilton, Ontario

TITLE: Strong conceptual completeness and various stability-theoretic results in continuous model theory

AUTHOR: Jean-Martin Albert, B.Sc. (McGill University), M.Sc. (University of Calgary)

SUPERVISOR: Professor Bradd Hart

NUMBER OF PAGES: iii, 89

Abstract

In this thesis we prove a strong conceptual completeness result for first-order continuous logic. Strong conceptual completeness was proved in 1987 by Michael Makkaï for classical first-order logic, and states that it is possible to recover a first-order theory T by looking at functors originating from the category $\mathrm{Mod}(T)$ of its models.

We then give a brief account of simple theories in continuous logic, and give a proof that the characterization of simple theories using dividing holds in continuous structures. These results are a specialization of well established results for thick cats which appear in [Ben03b] and in [Ben03a].

Finally, we turn to the study of non-archimedean Banach spaces over non-trivially valued fields. We give a natural language and axioms to describe them, and show that they admit quantifier elimination, and are $\aleph_0$-stable. We also show that the theory of non-archimedean Banach spaces has only one $\aleph_1$-saturated model in any cardinality.

Acknowledgements

First I thank my supervisor Bradd Hart for suggesting the topic of my thesis, for his help and for his brobdingnagian patience during the years it took to complete this work. I also thank the Mathematics department of McMaster University for accepting me into their Ph.D program. Finally I thank my parents for their continuing support

Contents

Contents **2**

1 Introduction **4**

2 Preliminaries and notational conventions **7**
2.1 Categories and functors 7
2.2 Generalized metric spaces 8
2.3 Ultrafilters 11
2.4 Ultraproducts 12

3 $\mathfrak{R}$-valued languages and structures **15**
3.1 $\mathfrak{R}$-valued languages 15
3.2 $\mathfrak{R}$-valued structures 17
3.3 Theories and types 19
3.4 Ultraproducts of structures 22
3.5 Metric structures 27
3.6 Definable sets 30
3.7 Imaginaries and the eq-construction 32
3.8 Conceptual completeness 37

4 Strong conceptual completeness **41**
4.1 Prelude 41
4.2 Ultrafunctors 43
4.3 Domination and covers 54

5 Simple theories in continuous logic **62**
5.1 Indiscernible sequences 62
5.2 Dividing 63
5.3 The D-ranks and the tree property 67
5.4 Morley sequences and the independence theorem 71

6 Non-archimedean Banach spaces **74**
6.1 Generalized ultra-metric spaces 75
6.2 Non-archimedean Banach spaces 77
6.3 Orthogonality and c_0 78
6.4 A language for non-archimedean Banach spaces 81

6.5 Model theory 84

Bibliography **88**

Chapter 1
Introduction

Model theory is the branch of mathematics which studies the semantics of logical systems. Traditionally, model theory deals with classical first-order predicate logic. In this logic, there are two truth values *true* and *false*. This makes first order logic ideal to study algebraic structures, such as fields and valued fields, and relational structures, such as graphs and partially ordered sets.

The concept of using more than one truth value in a logical system is not new. In its simplest incarnation, intuitionistic logic is a form of three valued logic, where the truth values are *true*, *false* and $\perp$ (notation mine), and the intended meaning of $\perp$ is "To be determined". The symbol $\perp$ is, for instance, the truth value of $p \vee \neg p$ in intuitionistic logic. It is theoretically possible to use any poset as a set of truth values. The properties of the poset dictate the properties of the logic. For example, a Boolean algebra gives rise to classical logic, whereas a Heyting algebra gives rise to intuitionistic logic.

The incarnation of continuous logic which we study in this thesis has several predecessors. In [CK66], the authors study a very general version of continuous logic, allowing any compact topological space to act as a set of truth values. This approach is too general to be used in applications.

In [HI02], the authors describe a logic which is suitable to describe Banach spaces. This is a more direct predecessor of first-order continuous logic, but is not described directly as such. Instead, the logic used behind the scene to describe Banach spaces is classical first-order logic. The languages include a special sort $\mathbb{R}$ for the real numbers. The authors of [HI02] then define a restricted class of formulae, called the *positive bounded formulae*, which roughly speaking are formulae in which all terms whose range is the special sort $\mathbb{R}$, take on a value which is bounded either above or below. There are special quantifiers $\exists x[\|x\| \leq r \wedge (\cdots)]$ and $\forall x[\|x\| \leq r \Rightarrow (\cdots)]$ for every real number r. The set of positive bounded formulae is closed under all logical connectives, except for negation. Restricting themselves to positive bounded formulae to describe Banach spaces, in [HI02], the authors recover analogues of the fundamental theorem of ultraproducts, and the Keisler-Shelah isomorphism theorem. The logic of positive bounded formulae is very much tied to the underlying class of normed spaces.

In [Ben05], Ben Yaacov discovers that a wide class of theories, called Hausdorff compact abstract theories, admit a uniformly definable metric on their structures. Compact abstract theories, also known as cats, are theories in a generalized version of first-order logic which satisfies the compactness theorem. In [BUar], Ben Yaacov and Usvyatsov show that cats can be described directly using a newly defined continuous version of first-order logic, which is a logic without negations, and in which equality is replaced by a distinguished symbol representing a distance function. They prove that this logic satisfies properly adapted versions of the compactness theorem, the Löwenheim-Skolem theorem and the fundamental theorem of ultraproducts. In the form defined in [BUar], continuous logic has allowed for the model theoretic treatment of several classes of analytic structures, such as Hilbert spaces, Banach spaces, C^*-algebras and probability algebras.

Another group of logics which are similar to the continuous logic described in [BUar] are the multi-valued logics of Lukasiewicz and Pavelka. These logics were introduced to formalize fuzzy logical systems. More specifically, Lukasiewicz logics (and continuous logic) are examples of t-norm fuzzy logics, which are $[0,1]$-valued logics in which a special operator $u : [0,1] \times [0,1] \to [0,1]$, called a t-norm, is used to implement conjunctions. While a t-norm fuzzy logic usually uses a fixed t-norm, continuous logic uses them all. The formalism of Lukasiewicz logics is used in [BP10] to establish a proof system and a completeness result for first-order continuous logic.

Chapter 3 We begin the thesis with an introduction to continuous languages and structures. To this end we take a path which is slightly different from the one in [BUar] and [BBHU08]. We first describe an *unbounded* version of the logic, in the sense that we allow truth values in the extended positive real numbers $\mathfrak{R} = \mathbb{R}^+ \cup \{\infty\}$. As an ordered set, $\mathfrak{R}$ is a compact topological space, and as such it is an adequate choice for a set of truth values. We interpret 0 as "true" and ∞ as "false". We describe this logic not with the specific purpose of studying metric spaces and analytic structures, but as a form of multivalued logic. In terms of expressive power, our description of first-order continuous logic is equivalent to the one described by Ben Yaacov in [BUar, BBHU08]. Formally though, it is closer to (and in fact also equivalent to) the continuous logic described by Hart, Farah and Sherman for C^*-algebras in [FHS]

Most of the definitions relevant to the theory of metric structures can already be formulated and proved in the more general context of $\mathfrak{R}$-valued logic and structures.

In the general context of $\mathfrak{R}$-valued languages, we obtain proofs of cornerstone theorems of model theory, namely the compactness theorem, the fundamental theorem of ultraproducts and the characterization of elementary classes and finitely axiomatizable elementary classes.

We also give a detailed description of the construction of the expansion $\mathscr{L}^{eq}$ of a language $\mathscr{L}$ in the context of $\mathfrak{R}$-valued metric languages. In this context, the definition of $\mathscr{L}^{eq}$ resembles, but is slightly different from, the definition in classical logic, in that not all imaginary elements can be obtained as canonical parameters of formulae, as is the case for first-order logic. The expansion $\mathscr{L}^{eq}$ of a language $\mathscr{L}$ and the expansion T^{eq} of a theory T that we describe were first described by Bradd Hart in [Har], where it is also proved that it is the "largest" conservative expansion of the theory T.

Chapter 4 The main contribution of this thesis is chapter 4, in which we state and prove a strong conceptual completeness result for continuous logic. This result states that one can recover, up to equivalence, a theory T by looking at functors originating from the category $\mathrm{Mod}(T)$ of its models. In the case of first-order logic, strong conceptual completeness was proved by Makkaï in [Mak88]. There, the author introduces the notion of an *ultrafunctor* from the category $\mathrm{Mod}(T)$ of models of T and elementary maps between them, to the category $\mathbb{S}$et of sets and functions between them. An ultrafunctor is a purely algebraic incarnation of the syntactic concept of a formula. It assigns to every model a set in such a way that for every ultrafilter pair (I, U), the set assigned to the ultraproduct $\prod_U \mathcal{M}_i$ is the ultraproduct of the sets assigned to the individual structures $\mathcal{M}_i$. Furthermore, ultrafunctors are required to behave well under certain maps that arise between ultraproducts.

Given a first-order language $\mathscr{L}$, the formulae of the expanded language $\mathscr{L}^{\mathrm{eq}}$ all give rise to ultrafunctors, which are called *definable functors*.

The category of sets is not the right category to express the semantics of continuous logic. We begin the chapter by finding an appropriate replacement for it, which we call $\mathbb{MET}_{\mathfrak{R}}$, and show that this category has suitable definitions for ultraproducts, and satisfies appropriate formulations of the fundamental theorem of ultraproducts. This allows us to define a notion of ultrafunctor $\mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ in the context of continuous logic.

A key result used in the proof of strong conceptual completeness for first-order logic is *subobject fullness*, which informally states that the subfunctors of a definable functor are themselves definable functors. This result is not quite true in the context of $\mathfrak{R}$-valued logic, which forces us to change the definition of definable functor slightly. In continuous logic, a definable functor is a functor $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ which can be embedded in a functor of the form $\mathfrak{ev}_\varphi$ in the sense that there is a natural transformation $\eta : \mathfrak{f} \to \mathfrak{ev}_\varphi$ such that for every $\mathcal{M} \in \mathrm{Mod}(T)$, $\eta_{\mathcal{M}}$ is injective. Strong conceptual completeness is stated in Theorem 4.2.16 and 4.2.22:

Theorem 4.2.16: For every $\varphi \in \mathscr{L}^{\mathrm{eq}}$, $\mathfrak{ev}_\varphi$ is an ultrafunctor.

Theorem 4.2.22: If $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ is a functor, then $\mathfrak{f}$ is a definable functor if and only if it is an ultrafunctor.

Chapter 5 We develop simplicity theory in the context of first-order continuous logic following the more axiomatic approach to simplicity theory in [Cas07]. Continuous theories can be construed as a special case of the concept of a compact abstract theory, for which there is already a well-established development of simplicity theory (see [Ben03b, Ben03a]). However, the greater abstraction of general cats prohibits some of the classical results of first-order simplicity theory to carry over. In a general cat, for example, having the same type may not be a type-definable condition. Consequently, types do not always have non-dividing extensions, and Morley sequences may not exist in every type. This is due to the fact that these results rely on the type-definability of indiscernible sequences.

To the knowledge of the author, no development of simplicity theory exists which is written specifically using the syntax of continuous logic. Chapter 5 is an attempt to remedy this situation. We follow [Cas07] quite closely, pointing out the places where proofs should be adapted to the context of continuous logic. The reader should note that in most cases, the continuous proofs we provide are very similar to their first-order counterpart.

Chapter 6 In this chapter we begin the study of non-archimedean Banach spaces over non-archimedean valued fields. Work in progress by Ben Yaacov in [Ben09b] deals with valued fields in the framework of continuous logic. The work done in [Ben09b] inspired chapter 6, though we do not rely on any of the results presented therein. We describe a language which allows to consider the class of non-archimedean Banach spaces as an elementary class, and find an axiomatization of the theory of this class. The language we describe is similar to the language for valued fields described in [Ben09b]. The main result of this chapter is that the theory of non-archimedean Banach spaces over a fixed non-archimedean valued field has quantifier elimination, and is $\aleph_0$-stable. Non-archimedean Banach spaces also possess a unique $\aleph_1$-saturated model of any cardinality, and are therefore unidimensional.

Chapter 2
Preliminaries and notational conventions

In this chapter we give the basic definitions and notational conventions in use throughout the thesis. We give definitions and results relevant to generalized metric spaces, which will be our main object of study. We focus on the generalized metric space we call $\mathfrak{R}$. This metric space will be used as a set of truth values, and its properties will dictate many properties of the logic we study in this thesis.

This thesis is written in ZFC. If X and Y are sets, $f : X \to Y$ is a function, and $A \subseteq X$, then we write $f[A]$ for $\{f(x) : x \in A\}$. As a convenience, we also use the notation $[x \mapsto f(x)]$ for a nameless function. Throughout the thesis, the notation $\sharp X$ will be used exclusively for set cardinality. The notation $|X|$ will usually be used for density character of topological spaces.

Section 2.1

Categories and functors

2.1.1 Definition: A *category* $\mathbb{C}$ consists of the following data:

1. A class $\mathbb{C}$, (possibly proper) of *objects*
2. Between any two objects $x, y \in \mathbb{C}$, a *set* $\mathfrak{Mor}(x, y)$ of *morphisms*
3. For any triplet $x, y, z \in \mathbb{C}$, a map $\circ_{x,y,z} : \mathfrak{Mor}(x, y) \times \mathfrak{Mor}(y, z) \to \mathfrak{Mor}(x, z)$

such that the following holds:

1. For every $x \in \mathbb{C}$, there is an element $1_x \in \mathfrak{Mor}(x, x)$ such that for any $y \in C$, $f \in \mathfrak{Mor}(x, y)$ and $g \in \mathfrak{Mor}(y, x)$, $1_x \circ_{y,x,x} g = g$ and $f \circ_{x,x,y} 1_x = f$;
2. For every $x, y, z, w \in \mathbb{C}$, every $f \in \mathfrak{Mor}(x, y)$, $g \in \mathfrak{Mor}(y, z)$ and $h \in \mathfrak{Mor}(z, w)$,

$$h \circ_{y,z,w} (g \circ_{x,y,z} f) = (h \circ_{y,z,w} g) \circ_{x,y,z} f$$

In general, we will omit the subscripts occurring in $\circ_{x,y,z}$, since there can be no confusion as to what they should be. ♣

2.1.2 Definition: A map $f \in \mathfrak{Mor}(x, y)$ is called:

1. a *monomorphism* or *monic* if whenever $g, h \in \mathfrak{Mor}(y, z)$, if $f \circ g = f \circ h$ then $g = h$;
2. an *epimorphism* or *epic* if whenever $g, h \in \mathfrak{Mor}(z, x)$, if $g \circ f = h \circ f$ then $g = h$;

3. an *isomorphism* if there is $g \in \mathfrak{Mor}(y, x)$ such that $f \circ g = 1_y$ and $g \circ f = 1_x$ ♣

2.1.3 Definition: Let $\mathbb{C}$ and $\mathbb{D}$ be categories. A *functor* $\mathfrak{F} : \mathbb{C} \to \mathbb{D}$ is a map which assigns to every object $x \in \mathbb{C}$ an object $\mathfrak{F}(x) \in \mathbb{D}$, and to every morphism $f : x \to y$ of $\mathbb{C}$ a morphism $\mathfrak{F}(f) : \mathfrak{F}(x) \to \mathfrak{F}(y)$ in $\mathbb{D}$ such that $\mathfrak{F}(f \circ h) = \mathfrak{F}(f) \circ \mathfrak{F}(h)$ whenever the composition makes sense. ♣

2.1.4 Definition: Let $\mathfrak{F}, \mathfrak{G} : \mathbb{C} \to \mathbb{D}$ be two functors. A *natural transformation* $\eta : \mathfrak{F} \to \mathfrak{G}$ consists of a family of $\mathbb{D}$-morphisms $\{\eta_x : \mathfrak{F}(x) \to \mathfrak{G}(x) : x \in \mathbb{C}\}$ such that $\eta_y \circ \mathfrak{F}(h) = \mathfrak{G}(h) \circ \eta_x$ for any $x, y \in \mathbb{C}$ and any $h \in \mathfrak{Mor}(x, y)$. A natural transformation is called:

1. *monic* if η_x is monic for every $x \in \mathbb{C}$
2. *epic* if η_x is epic for every $x \in \mathbb{C}$
3. a *natural isomorphism* if η_x is an isomorphism for every $x \in \mathbb{C}$ ♣

Section 2.2

Generalized metric spaces

The primary objects of study in continuous logic are metric spaces. In our context, they will in fact be generalized metric spaces. In this section we gather some definitions and relevant facts about these spaces.

2.2.1 Definition: A *generalized pseudo-metric* on a set X is a function $\rho : X \times X \to \mathfrak{R}$ such that

1. $\rho(x, x) = 0$
2. $\rho(x, y) = \rho(y, x)$
3. $\rho(x, y) \leq \rho(x, z) + \rho(z, y)$ for every $x, y, z \in X$.

If $\rho(x, y) = 0$ implies that $x = y$, then ρ is called a *generalized metric*.

The term "generalized" refers to the fact that nothing prevents ρ from taking on ∞ as a value. A *generalized pseudo-metric space* is a pair (X, ρ), where X is a set, and ρ is a generalized pseudo-metric on X, and a generalized metric space is a pair (X, ρ), where X is a set, and ρ is a generalized metric on X. ♣

There is an equivalence relation on X defined by $x \sim y \Leftrightarrow \rho(x, y) = 0$, and ρ induces a generalized metric d_ρ on the set $X/\sim$ of equivalence classes of $\sim$. The metric d_ρ is defined by

$$d_\rho(x/\sim, y/\sim) = \inf_{\substack{z \sim x \\ z' \sim y}} \rho(z, z').$$

The metric space $(X/\sim, d_\rho)$ will be denoted X/ρ. If (X, ρ) is a generalized pseudo-metric space, then the sets

$$B_{<\varepsilon}(x) = \{y \in X : \rho(x, y) < \varepsilon\}$$

and

$$B_{>\varepsilon}(x) = \{y \in X : \rho(x, y) > \varepsilon\}$$

form a basis of open sets for a topology on X. Throughout this thesis, generalized metric spaces and pseudo-metric spaces will be viewed as topological spaces, always with this topology.

2.2.2 Definition: If X is a generalized (pseudo-)metric space, then we define $|X|$ to be the density character of X, i.e. the cardinality of the smallest dense subset $D \subseteq X$. ♣

Throughout this thesis, whenever we refer to the size, power or cardinality of a metric space, we always mean its density character, and use the notation $|X|$. The definition of continuous and uniformly continuous map can be written for generalized pseudo-metrics as well.

2.2.3 Definition: Let X and Y be generalized pseudo-metric spaces. A function $f : X \to Y$ will be called *pseudo uniformly continuous* if and only if for every $\varepsilon > 0$, there is $\delta > 0$ such that if $\rho_X(x, y) \leq \delta$, then $\rho_Y(f(x), f(y)) \leq \varepsilon$. If f is pseudo uniformly continuous, then it induces a uniformly continuous function $\bar{f} : X/\rho_X \to Y/\rho_Y$. ♣

It is worth noting that the numbers ε and δ in the last definition, and in the definition of $B_{<\varepsilon}$ can be taken to be ∞. The generalized metric space most central to this thesis is

$$\mathfrak{R} = \{r \in \mathbb{R} : r \geq 0\} \cup \{\infty\}$$

where the generalized metric is given by $|x - y|$, with the convention that $|x - y| = \infty$ if either $x = \infty$ or $y = \infty$ but not both, and $|\infty - \infty| = 0$. We will view $\mathfrak{R}$ both as a totally ordered set and a generalized metric space. In fact, the topology induced on $\mathfrak{R}$ by this generalized metric is equal to the topology induced on it by its ordering. In particular, $\mathfrak{R}$ is compact both as a generalized metric space and as an ordered set. For every $n \in \mathbb{N}$, we have a generalized metric on $\mathfrak{R}^n$ given by

$$d(x, y) = \max\{|x_i - y_i| : 0 \leq i \leq n\}$$

and the topology induced on $\mathfrak{R}^n$ by this metric is the same as the product topology. There is also a generalized metric on $\mathfrak{R}^\omega$ defined by

$$d(x, y) = \sum_{i=0}^{\infty} \frac{|x_i - y_i|}{2^i}$$

and this metric induces the Tychonoff product topology on $\mathfrak{R}^\omega$, which means that this set is also compact as a generalized metric space. There are, of course, many possible choices of metrics on $\mathfrak{R}^\omega$ which make it a compact generalized metric space. By a *compact generalized metric*, we shall mean a metric d with respect to which $\mathfrak{R}^\omega$ is a compact metric space.

2.2.4 Definition: For $1 \leq n \leq \omega$, we define a *compact norm* on $\mathfrak{R}^n$, we shall mean a function $u : \mathfrak{R}^n \to \mathfrak{R}$ such that:

1. $u(x_1, ..., x_k, ...) = 0$ if and only if $x_1 = \cdots = x_k = \cdots = 0$
2. $u(x_1 + y_1, ..., x_k + y_k, ...) \leq u(x_1, ..., x_k, ...) + u(y_1, ..., y_k, ...)$
3. $u(x_1, ..., x_k, ...) = \infty$ if and only if $x_i = \infty$ for some $i \leq n$
4. $\mathfrak{R}^n$ is compact with respect to the topology induced by the generalized metric

$$u(|x_1 - y_1|, ..., |x_k - y_k|, ...)$$

For any $n \geq 0$, a compact norm u on $\mathfrak{R}^n$ induces a topology on $\mathfrak{R}^n$ in which the collection of sets

$$O_\varepsilon \stackrel{\text{def}}{:=} \{x \in \mathfrak{R}^n : u(x) < \varepsilon\}$$

forms a basis of open sets. If we denote by π_i the i-th coordinate projection on $\mathfrak{R}^\omega$, then we see that $\pi_i[O_\varepsilon] = \mathfrak{R}$ for all but finitely many i's. Therefore, we see that the topology on $\mathfrak{R}^n$ induced by a compact norm is exactly the Tychonoff product topology, and thus that all compact norms are topologically equivalent.

An example of a compact norm on $\mathfrak{R}^\omega$ is the function

$$u(x_1, ..., x_n, ...) = \sum_{i=0}^{\infty} \frac{x_i}{2^i}.$$

Note that the definition of the metric d above is equal to $u(|x_1 - y_1|, ..., |x_n - y_n|, ...)$. In general, if v is any compact norm, then $v(|x_1 - y_1|, ..., |x_n - y_n|, ...)$ is a generalized metric on $\mathfrak{R}^\omega$. Conversely, if $d : \mathfrak{R}^n \times \mathfrak{R}^n \to \mathfrak{R}$ is a compact generalized metric, then the function $u(x) = d(x, 0)$ is a compact norm on $\mathfrak{R}^n$. The following fact is used in the proof of proposition below, and then later in chapter 6.2.

2.2.5 Fact: *The function $[x \mapsto 1-e^{-x}]$, with the convention that $\infty \mapsto 1$ is an order preserving homeomorphism $\mathfrak{R} \to [0,1]$. Its inverse is given by $\left[y \mapsto \ln \frac{1}{1-y}\right]$, with the convention that $1 \mapsto \infty$. In fact, both these maps are uniformly continuous with respect to the metrics on $\mathfrak{R}$ and $[0,1]$*

2.2.6 Proposition: *For every $0 \leq n \leq \omega$, a function $u : \mathfrak{R}^n \to \mathfrak{R}$ is continuous if and only if it is uniformly continuous with respect to the generalized metric on $\mathfrak{R}^n$.*

Proof: Let $u : \mathfrak{R}^k \to \mathfrak{R}$ be continuous. Let $f = \left[y \mapsto \ln \frac{1}{1-y}\right]$. Then

$$g = [(x_1, ..., x_n, ...) \mapsto f^{-1}(u(f(x_1), ..., f(x_n), ...)]$$

is a continuous function $[0,1]^n \to [0,1]$. Since $[0,1]^n$ is a compact metric space for $0 \leq n \leq \omega$, g is uniformly continuous. Now note that

$$u(x_1, ..., x_n, ...) = f(g(f^{-1}(x_1), ..., f^{-1}(x_n), ...)),$$

which is a composition of uniformly continuous functions. □

Infinitary continuous functions can always be obtained as limits (in the order topology of $\mathfrak{R}$) of finitary ones: if $u : \mathfrak{R}^\omega \to \mathfrak{R}$ is a continuous function, then for every $k < \omega$, we let

$$u_k(x_1, ..., x_k) = \sup_{y_\ell : \ell > k} u(x_1, ..., x_k, y_{k+1}, ..., y_{k+i}, ...).$$

For every sequence $x_1, ..., x_n, ...$, the sequence $u_k(x_1, ..., x_k)$ is decreasing and bounded below by $u(x_1, ..., x_n, ...)$. Therefore, it converges. It is not hard to see that in fact, it converges to $u(x_1, ..., x_n, ...)$, from which we conclude that $\lim_{k\to\infty} u_k = u$.

2.2.7 Definition: For every $0 \leq n \leq \omega$, and for every continuous $u : \mathfrak{R}^n \to \mathfrak{R}$, we let $UC(u) : \mathbb{R} \to \mathbb{R}$ be defined as follows: for every $\varepsilon > 0$, $UC(u)(\varepsilon)$ is the largest $\delta > 0$ such that whenever $d_n(\bar{x}, \bar{y}) < \delta$, $|u(\bar{x}) - u(\bar{y})| < \varepsilon$, where d_n is a compact generalized metric on $\mathfrak{R}^n$. The functions $UC(u)$ are fixed throughout this thesis. ♣

2.2.8 Definition: The category MET is defined as follows:

Objects: Generalized metric spaces

Morphisms: Uniformly continuous maps

Composition: Function composition ♣

A generalized metric space X is called *complete* if every Cauchy sequence in X converges in X. Every generalized metric space X has a unique completion $\widehat{X}$ such that $X \subseteq \widehat{X}$, and whenever Y is a complete generalized metric space, every continuous function $f : X \to Y$ extends uniquely to a continuous map $\widehat{f} : \widehat{X} \to Y$. The construction of $\widehat{X}$ from X is in every way similar to the construction of the completion of an ordinary metric space. This gives rise to a functor $C : \mathrm{MET} \to \mathrm{MET}$ whose action on objects is given by $C(X) = \widehat{X}$, and if $f : X \to Y$ is continuous, then $C(f) = \widehat{i \circ f}$, where $i : Y \to \widehat{Y}$ is the inclusion map.

Section 2.3

Ultrafilters

In this section we gather up the basic definitions and facts about ultrafilters. These definitions and facts are standard, and are stated without proof. Their proofs, as well as the proofs which we do give, can be found in [CK90]. Let I be a set. Throughout we use the notation 2^I to denote the power set of I. A *filter* on I is a family $F \subseteq 2^I$ such that:

1. $\emptyset \notin F$
2. F has the finite intersection property, i.e. whenever $P, Q \in F$, $P \cap Q \in F$ as well, provided $P \cap Q \neq \emptyset$;
3. F is closed upward, i.e. if $P \in F$ and $P \subseteq Q$, then $Q \in F$.

A filter F is called *principal* if there is a set $X \subseteq I$ such that $F = \{Y \subseteq I : X \subseteq Y\}$. A *non-principal filter* is an filter that is not principal. An *ultrafilter on I* is a filter U with the additional property that for every $X \subseteq I$, either $X \in U$ or $I \backslash X \in U$.

2.3.1 Definition: An *ultrafilter pair* is a pair (I, U) with I a set and U an ultrafilter on I. ♣

2.3.2 Fact: *Let $F \subseteq 2^I$ be a family with the finite intersection property. Then there is an ultrafilter U on I with $U \supseteq F$*

2.3.3 Definition: Let J be a set, and let $f : J \to I$ be a function. Let U be an ultrafilter on J. We define $f[U]$ as follows:

$$P \in f[U] \Leftrightarrow f^{-1}(P) \in U.$$

One easily checks that $f[U]$ is indeed an ultrafilter on I, called the *ultrafilter induced by f on I*. If f is onto, then $f[U] = \{f[P] : P \in U\}$. ♣

2.3.4 Definition: Let I be any set, and let $F \subseteq 2^I$. An F-selector is a function $f : F \to I$ such that $f(x) \in x$ for every $x \in F$. ♣

2.3.5 Proposition: *Let U be an ultrafilter on I, and $f : U \to I$ be a U-selector. Then there is an ultrafilter W on U such that $f[W] = U$.*

Proof: Let $W = \{f^{-1}[P] : P \in U\}$. Then W clearly has the finite intersection property, and is closed up. Since we are assuming that f is a U-selector, $f(P) \in P$ for every $P \in U$, and therefore $\emptyset \notin W$. It is easy to see that W is in fact an ultrafilter, and that $f[W] = U$.□

2.3.6 Definition: Let (I, U) be an ultrafilter pair, and κ be a cardinal number. Then U is called *κ-regular* if and only if there is $D \subseteq U$ with $\sharp D = \kappa$, and for every $i \in I$, the set $\{P \in D : i \in P\}$ is finite. ♣

2.3.7 Proposition: *Any infinite set I admits an $|I|$-regular ultrafilter.*

Proof: Consider the set J of all finite subsets of I. For $i \in I$, let $\hat{i} = \{j \in J : i \in J\}$, and let $D = \{\hat{i} : i \in I\}$. Then $|D| = |I|$, and D has the finite intersection property. More importantly, any $j \in J$ belongs to only finitely many elements in D, since $j \in J$ means j is a finite set, and $j \in \hat{i}$ implies $i \in j$.

Since D has the finite intersection property, it extends to an ultrafilter U on J. The fact that $D \subseteq U$ witnesses $|I|$-regularity. If $f : J \to I$ is any bijection, then $f[U]$ is an $|I|$-regular ultrafilter on I, with $|I|$-regularity witnessed by $f[D]$. □

Section 2.4

Ultraproducts

Let I be a set and U be an ultrafilter on I. If $(x_i : i \in I)$ is any I-indexed sequence of elements of $\mathfrak{R}$, and $x \in \mathfrak{R}$, then we say that $\lim_{i \to U} x_i = x$ if and only if for every open set O such that $x \in O$,

$$\{i \in I : x_i \in O\} \in U.$$

Since $\mathfrak{R}$ is Hausdorff, for every I-indexed sequence $(x_i : i \in I)$, there is a *unique* $x \in \mathfrak{R}$ such that $\lim_{i \to U} x_i = x$, and we get a function $\lim_U : \mathfrak{R}^I \to \mathfrak{R}$, called the *ultra-limit modulo U*. Note that $\lim_{i \to U} x_i = \infty$ if and only if for every $r \in \mathbb{R}$, the set

$$\{i \in I : x_i \geq r\} \in U$$

and $\lim_{i \to U} x_i = x < \infty$ if and only if for every $\varepsilon > 0$,

$$\{i \in I : |x_i - x| \leq \varepsilon\} \in U.$$

This is just a restatement of the definition with a basis of open sets.

There is an equivalence relation on $\mathfrak{R}^I$ defined by $(x_i : i \in I) \sim_U (y_i : i \in I)$ if and only if $\lim_{i \to U} x_i = \lim_{i \to U} y_i$. Since $\mathfrak{R}$ is compact, for every sequence $(x_i : i \in I)$ there is $x \in \mathfrak{R}$ such that $(x_i : i \in I) \sim (x : i \in I)$. This implies that $\mathfrak{R}^I / \sim_U = \mathfrak{R}$. This simple fact, and Theorem 2.4.1 below are key to the usability of $\mathfrak{R}$ as a set of truth values.

2.4.1 Theorem: *Let $u : \mathfrak{R}^\kappa \to \mathfrak{R}$ be a continuous function. Then for every ultrafilter pair (I, U),*

$$\lim_{i \to U} u(x_{i1}, ..., x_{in}, ...) = u(\lim_{i \to U} x_{i1}, ... \lim_{i \to U} x_{in}, ...).$$

Proof: Let O be a neighbourhood of $u(\lim_{i\to U} x_{i1}, \dots \lim_{i\to U} x_{in}, \dots)$. We need to prove that the set

$$\{i \in I : u(x_{i1}, \dots, x_{in}, \dots) \in O\}$$

is in U. We have that $u^{-1}(O)$ is a neighbourhood of $(\lim_{i\to U} x_{i1}, \dots \lim_{i\to U} x_{in}, \dots)$. Write $u^{-1}(O) = \prod_{i\in\kappa} O_i$, with each O_i open in $\mathfrak{R}$. There is a finite subset $\Delta \subseteq \kappa$ such that $O_i = \mathfrak{R}$ if $i \notin \Delta$. Since $\lim_{i\to U} x_{ik} \in O_k$ for $k \in \Delta$, we have

$$P_k = \{i \in I : x_{ik} \in O_k\} \in U.$$

Let $P = \bigcap_{k\in\Delta} P_k$. Then $P \in U$, and for every $i \in P$, and $k \in \Delta$, $x_{ik} \in O_k$. This implies that for every $i \in P$, $(x_{ik} : k \leq \kappa) \in u^{-1}(O)$, so that $u(x_{i1}, \dots, x_{in}, \dots) \in O$. This implies that

$$\lim_{i\to U} u(x_{i1}, \dots, x_{in}, \dots) = u(\lim_{i\to U} x_{i1}, \dots \lim_{i\to U} x_{in}, \dots),$$

finishing the proof. □

If $\{(X_i, d_i) : i \in I\}$ is an I-indexed sequence of generalized metric spaces, then we define a generalized pseudo-metric ρ on $\prod_I X_i$ by

$$\rho(\bar{x}, \bar{y}) \stackrel{\text{def}}{:=} \lim_{i\to U} d_i(x_i, y_i).$$

The quotient of $\prod_I X_i$ by the equivalence relation $\rho(x,y) = 0$ is called the *ultraproduct* of the spaces X_i, and is denoted $\prod_U X_i$. If all the X_i's are equal, then $\prod_U X_i$ is denoted X^U, and is called an *ultrapower* of X. We have a canonical diagonal embedding $\iota : X \to X^U$, and if X is a compact metric space, then $X \cong X^U$ via ι. In general, an element of $\prod_U X_i$ will be denoted $\langle x_i \rangle_U$.

The construction of ultraproducts of metric spaces is what forces us to consider generalized pseudo-metrics instead of finite-valued metrics. Even if d_i on X_i is finite valued for every $i \in I$, it is possible that $\lim_{i\to U} d_i(x_i, y_i) = \infty$. For every set I, we define $\mathbb{MET}^I_{equiv}$ in the following way:

Objects: Sequences $(X_i : i \in I)$ of elements of $\mathbb{MET}$

Morphisms: Uniformly equicontinuous sequences $(f_i : i \in I)$ of morphisms of $\mathbb{MET}$. That is to say, sequences $(f_i : i \in I)$ such that $f_i : (X_i, d_i) \to (X'_i, d'_i)$ for every $i \in I$, and for every $\varepsilon > 0$, there is a $\delta > 0$ with the property that for every $i \in I$, if $d_i(x,y) \leq \delta$, then $d'_i(f_i(x), f_i(y)) \leq \varepsilon$.

Composition: Coordinatewise function composition

2.4.2 Proposition: *For every ultrafilter pair (I, U), there is a functor $\prod_U : (\mathbb{MET})^I_{equiv} \to \mathbb{MET}$ which assigns to every sequence $(X_i : i \in I)$ the metric space $\prod_U X_i$, and to every sequence $(f_i : i \in I)$ the function $(\prod_U f_i)(\langle x_i \rangle_U) = \langle f_i(x_i) \rangle_U$.*

Proof: That $\prod_U X_i$ is a generalized metric space was argued earlier. It remains to show that $\prod_U f_i$ is a uniformly continuous function. By the definition of $\mathbb{MET}^I_{equi}$, if $(f_i : i \in I)$ is a morphism of $\mathbb{MET}^I_{equi}$, then $(f_i : i \in I)$ is equicontinuous. Let $\varepsilon > 0$, and choose δ

so that for every $i \in I$ and every $x_i, y_i \in X_i$ $d_i(x_i, y_i) \leq \delta$ implies $d_i'(f_i(x_i), f_i(y_i)) \leq \varepsilon$. Suppose $d(\langle x_i \rangle_U, \langle y_i \rangle_U) \leq \delta$. Then

$$d\left(\left(\prod_U f_i\right)(\langle x_i \rangle_U), \left(\prod_U f_i\right)(\langle x_i \rangle_U)\right) = \lim_{i \to U} d_i(f_i(x_i), f_i(y_i))$$

by definition. By definition, since $d(\langle x_i \rangle_U, \langle y_i \rangle_U) \leq \delta$, there is $P \in U$ such that for every $i \in P$, $d(x_i, y_i) \leq \delta$. Therefore, for every $i \in P$, $d_i(f_i(x_i), f_i(y_i)) \leq \varepsilon$, showing that $\lim_{i \to U} d_i(f_i(x_i), f_i(y_i)) \leq \varepsilon$ as required. □

2.4.3 Fact: *The completion functor* $C : \mathbb{MET} \to \mathbb{MET}$ *commutes with all ultraproduct functors, i.e. for every ultrafilter pair* (I, U), $C(\prod_U X_i) = \prod_U C(X_i)$. *In fact,* $\prod_U C(X_i)$ *is always complete.*

Chapter 3

$\mathfrak{R}$-valued languages and structures

Section 3.1

$\mathfrak{R}$-valued languages

There are several paths one can take to arrive at a definition of continuous logic. Originally, in [Ben05, BUar], continuous logic was created as an internal logic for Hausdorff compact abstract theories. In [BBHU08], [BP10], [Ben09a], continuous logic is viewed more as fitting in the context of multi-valued logics. This seems more natural, as it is often easier to describe a continuous language for a class of structures from scratch than to view the class of structures as a cat and try to read the appropriate language that way. This more recent point of view also allows for a more syntactic treatment of continuous logic in [BP10], which includes a proof system and a completeness theorem.

We begin by defining the notion of $\mathfrak{R}$-valued language. Recall from chapter 2 that $\mathfrak{R}$ denotes the extended positive real numbers with the order topology. We will be using the elements of $\mathfrak{R}$ as truth values. A lot of the theory of metric structures can be discussed in the more general context of $\mathfrak{R}$-valued logic, including the very useful compactness theorem and the definition of ultraproducts. The formal definition of $\mathfrak{R}$-valued languages is very similar to the definition of classical first-order languages. A $\mathfrak{R}$-valued language $\mathscr{L}$ consists of the following data:

1. A collection $\mathbf{Sort}_{\mathscr{L}}$ of *sort symbols.*

2. A collection $\mathbf{Func}_{\mathscr{L}}$ of *function symbols.* Every $\mathsf{f} \in \mathbf{Func}_{\mathscr{L}}$ has a *domain* $\mathrm{dom}(\mathsf{f})$ consisting of an ordered tuple of elements of $\mathbf{Sort}_{\mathscr{L}}$, possibly with repetition, and a *codomain* $\mathrm{codom}(\mathsf{f})$ consisting of a single sort symbol S.

3. A collection $\mathbf{Rel}_{\mathscr{L}}$ of *relation symbols.* Every relation symbol has a domain $\mathrm{dom}(\mathsf{R})$ defined as for function symbols.

4. A collection $\mathbf{Const}_{\mathscr{L}}$ of *constant symbols.* Every constant $\mathsf{c} \in \mathbf{Const}_{\mathscr{L}}$ has a *type* consisting of a sort symbol denoted $\mathrm{type}(\mathsf{c})$

5. A collection $\mathbf{Var}_{\mathscr{L}}$ of *sorted variables* with a map *type* : $\mathbf{Var}_{\mathscr{L}} \to \mathbf{Sort}_{\mathscr{L}}^{<\omega}$ with infinite fibres.

If $x \in \mathbf{Rel}_{\mathscr{L}}$ or $x \in \mathbf{Func}_{\mathscr{L}}$, then the length of $\mathrm{dom}(x)$ is called the *arity* of x, and is denoted $|x|$. The element x will be referred to as $|x|$*-ary.* Unless more standard notation applies, the formal symbols in the disjoint union $\mathbf{Sort}_{\mathscr{L}} \coprod \mathbf{Func}_{\mathscr{L}} \coprod \mathbf{Rel}_{\mathscr{L}} \coprod \mathbf{Const}_{\mathscr{L}}$ will be typeset in a sans-serif font, with the letter c usually reserved for constants, the letter R

for relations and the letter f for functions. Symbols may contain subscripts. Variables will be typeset normally.

When there is no possible confusion, we shall drop the subscript $\mathscr{L}$. An element of $\mathbf{Sort}_{\mathscr{L}}$ (*resp.* $\mathbf{Func}_{\mathscr{L}}$, $\mathbf{Rel}_{\mathscr{L}}$ or $\mathbf{Const}_{\mathscr{L}}$) will be referred to as a sort of $\mathscr{L}$ (*resp.* function, relation, constant symbol of $\mathscr{L}$).

3.1.1 Definition: A *term* is a formal expression t built from the function and constant symbols according to the following rules, in which we also define the codomain codom(t) of a term t:

1. Any variable $x \in \mathbf{Var}_{\mathscr{L}}$ and any constant $\mathsf{c} \in \mathbf{Const}_{\mathscr{L}}$ is a term. If $x \in \mathbf{Var}_{\mathscr{L}}$, then $\text{codom}(x) = \text{type}(x)$. If $\mathsf{c} \in \mathbf{Const}_{\mathscr{L}}$, then $\text{codom}(\mathsf{c}) = \text{type}(\mathsf{c})$.

2. If $t_1, ..., t_n$ are terms of the appropriate codomains, and $\mathsf{f} \in \mathbf{Func}_{\mathscr{L}}$ is such that $\text{dom}(\mathsf{f}) = (\text{codom}(t_1), \ldots, \text{codom}(t_n))$, then the expression $t = \mathsf{f}(t_1, ..., t_n)$ is a term. The codomain of t is the codomain of f. ♣

3.1.2 Definition: A *logical connective* is a formal symbol u corresponding to a continuous function $u : \mathfrak{R}^k \to \mathfrak{R}$, where $0 \leq k \leq \omega$. The number k is the *arity* of the connective. If $k < \omega$, then the connective u is said to be *finitary*. If $k = \omega$, then u is said to be *infinitary*, and we assume that $\mathfrak{R}^k$ is endowed with the Tychonoff product topology, making it a compact set. ♣

Note that a 0-ary connective is just an element of $\mathfrak{R}$. A 0-ary connective will be denoted by $r \in \mathfrak{R}$.

3.1.3 Definition: A *formula* is a formal expression φ which is of one of the following forms:

1. $\mathsf{R}(t_1, ..., t_n)$, where t_i is a term with appropriate type for $i = 1, ..., n$, and R is an n-ary relation symbol of $\mathscr{L}$.

2. $\varphi \wedge \psi$ and $\varphi \vee \psi$, where φ and ψ are formulae.

3. For every $\mathsf{S} \in \mathbf{Sort}_{\mathscr{L}}$, $\forall x \in \mathsf{S}[\varphi]$ and $\exists x \in \mathsf{S}[\varphi]$, where φ is a formula. Since it will, in general, be possible to determine S from x, to simplify the notation, we will generally omit the sort, and write $\forall x[\varphi]$ and $\exists x[\varphi]$. We shall also write $\forall x_1 \cdots x_n[\varphi]$ instead of $\forall x_1 \cdots \forall x_n[\varphi]$, and similarly for $\exists x_1 \cdots x_n[\varphi]$

4. $u(\varphi_1, ..., \varphi_i, ...)$, where each φ_i for $i < n \leq \omega$ is a formula and $u : \mathfrak{R}^n \to \mathfrak{R}$ is an n-ary connective. ♣

We will use the notation $\varphi \leq \psi$ for the formula $\varphi \dot{-} \psi$, where $x \dot{-} y$ is the function $[(x, y) \mapsto \max\{x - y, 0\}]$.

Since every $r \in \mathfrak{R}$ can be viewed as a formula by considering a 0-ary connective, for every $r \in \mathfrak{R}$, and any formula φ, we get formulae $\varphi \leq r$ and $\varphi \geq r$.

3.1.4 Definition: The set of *free variables* of a term are defined recursively as follows:

1. $\text{FV}(x) = \{x\}$ for $x \in \mathbf{Var}_{\mathscr{L}}$;

2. $\text{FV}(\mathsf{c}) = \emptyset$ for $\mathsf{c} \in \mathbf{Const}_{\mathscr{L}}$;

3. $\text{FV}(\mathsf{f}(t_1, ..., t_n)) = \text{FV}(t_1) \cup \cdots \cup \text{FV}(t_n)$.

A *closed term* is a term with no free variables, i.e. a term t such that $\mathrm{FV}(t) = \emptyset$. Thus a closed a term is a term formed using only variables, constants and function symbols. The free variables of a formula are defined recursively as follows:

1. $\mathrm{FV}(\mathsf{R}(t_1, ..., t_n)) = \mathrm{FV}(t_1) \cup \cdots \cup \mathrm{FV}(t_n)$
2. $\mathrm{FV}(\varphi \wedge \psi) = \mathrm{FV}(\varphi \vee \psi) = \mathrm{FV}(\varphi) \cup \mathrm{FV}(\psi)$;
3. $\mathrm{FV}(u(\varphi_1, ..., \varphi_k, ...)) = \bigcup_{k<n} \mathrm{FV}(\varphi_k)$ if u is any connective;
4. $\mathrm{FV}(\forall x \in \mathsf{S}[\varphi]) = \mathrm{FV}(\varphi) \backslash \{x\}$;
5. $\mathrm{FV}(\exists x \in \mathsf{S}[\varphi]) = \mathrm{FV}(\varphi) \backslash \{x\}$.

A formula φ such that $\mathrm{FV}(\varphi) = \emptyset$ is called a *sentence*. ♣

Note that a formula of $\mathscr{L}$ can, in general, have infinitely many free variables. However $\mathrm{FV}(\varphi)$ is always a countable set. The *arity* of a term t or a formula φ is defined to be the cardinality of $\mathrm{FV}(t)$ or $\mathrm{FV}(\varphi)$. We define the *domain* of a formula φ to be

$$\mathrm{dom}(\varphi) = (\mathrm{type}(x) : x \in \mathrm{FV}(\varphi)).$$

In order to make the notation look more familiar, instead of writing $\mathrm{dom}(x)$ as a tuple of length the arity of x, we will write it as $\mathsf{S}_1 \times \cdots \times \mathsf{S}_n \times \cdots$. We will also use the notation $\mathsf{S} \in \mathrm{dom}(\varphi)$ to indicate that S is one of the sorts listed in $\mathrm{dom}(\varphi)$.

Section 3.2

$\mathfrak{R}$-valued structures

In this section we give the definition of an $\mathfrak{R}$-valued structure. An *$\mathfrak{R}$-valued structure* $\mathcal{X}$ for a continuous language $\mathscr{L}$ consists of the following data:

1. For every sort $\mathsf{S} \in \mathbf{Sort}_{\mathscr{L}}$, a set $\mathsf{S}(\mathcal{X})$.
2. For every function symbol $\mathsf{f} \in \mathbf{Func}_{\mathscr{L}}$, a function $[\![\mathsf{f}]\!]_{\mathcal{X}} : \mathrm{dom}(\mathsf{f})(\mathcal{X}) \to \mathrm{codom}(\mathsf{f})(\mathcal{X})$;
3. For every relation symbol $\mathsf{R} \in \mathbf{Rel}_{\mathscr{L}}$, a function $[\![\mathsf{R}]\!]_{\mathcal{X}} : \mathrm{dom}(\mathsf{R})(\mathcal{X}) \to \mathfrak{R}$;
4. For every constant symbol $\mathsf{c} \in \mathbf{Const}_{\mathscr{L}}$, an element $[\![\mathsf{c}]\!]_{\mathcal{X}} \in \mathrm{type}(\mathsf{c})(\mathcal{X})$;
5. For every variable $x \in \mathbf{Var}_{\mathscr{L}}$, the identity map $[\![x]\!]_{\mathcal{X}} : \mathrm{type}(x)(\mathcal{X}) \to \mathrm{type}(x)(\mathcal{X})$;

Here we note that there is a coarsest topology on $\mathsf{S}(\mathcal{X})$ with respect to which the interpretation of all the relation symbols are continuous. Therefore, it makes sense to talk about the density character $|\mathsf{S}(\mathcal{X})|$ (which may very well be the cardinality of $\mathsf{S}(\mathcal{X})$). We define the *density character* $|\mathcal{X}|$ of an $\mathscr{L}$-structure $\mathcal{X}$ to be

$$|\mathcal{X}| = \sum_{\mathsf{S} \in \mathbf{Sort}_{\mathscr{L}}} |\mathsf{S}(\mathcal{X})|.$$

In order to avoid a huge amount of notational clutter, we use a rather informal approach to membership of elements in a multi-sorted structure. We adopt the convention that elements of multi-sorted structures "know" what sort they belong to. If $A \subseteq \mathcal{X}$, then we define the language $\mathscr{L}_A$ to consist of $\mathscr{L}$ together with a new constants a for each $a \in A$. The structure $\mathcal{X}$ can be made into an $\mathscr{L}_A$-structure by defining $[\![\mathsf{a}]\!]_{\mathcal{X}} = a$. A term (resp. formula) of $\mathscr{L}_A$ will also be referred to as a term (resp. formula) of $\mathscr{L}$ *defined over* A.

3.2.1 Definition: If t is a term, then its *interpretation* $[\![t]\!]_{\mathcal{X}}$ in $\mathcal{X}$ is defined inductively by $[\![\mathsf{f}(t_1, ..., t_n)]\!]_{\mathcal{X}} \stackrel{\text{def}}{:=} [\![\mathsf{f}]\!]_{\mathcal{X}} ([\![t_1]\!]_{\mathcal{X}}, ..., [\![t_n]\!]_{\mathcal{X}})$, where t_i for $1 \leq i \leq n$ is a term of the appropriate codomain. ♣

3.2.2 Definition: The interpretation of formulae is defined as follows. All variables present are tuples of the appropriate type and length.

1. $[\![\mathsf{R}(t_1, ..., t_n)]\!]_{\mathcal{X}} (\bar{x}) = [\![\mathsf{R}]\!]_{\mathcal{X}} ([\![t_1]\!]_{\mathcal{X}} (\bar{x}), ..., [\![t_n]\!]_{\mathcal{X}} (\bar{x}))$, where t_i for $1 \leq i \leq n$ is a term of the appropriate codomain.
2. $[\![\varphi \wedge \psi]\!]_{\mathcal{X}} (\bar{x}) = \max\{[\![\varphi]\!]_{\mathcal{X}} (\bar{x}), [\![\psi]\!]_{\mathcal{X}} (\bar{x})\}$, where φ and ψ are formulae;
3. $[\![\varphi \vee \psi]\!]_{\mathcal{X}} (\bar{x}) = \min\{[\![\varphi]\!]_{\mathcal{X}} (\bar{x}), [\![\psi]\!]_{\mathcal{X}} (\bar{x})\}$, where φ and ψ are formulae;
4. $[\![u(\varphi_1, ..., \varphi_n, ...)]\!]_{\mathcal{X}} (\bar{x}) = u([\![\varphi_1]\!]_{\mathcal{X}} (\bar{x}), ..., [\![\varphi_n]\!]_{\mathcal{X}} (\bar{x}), ...)$, where u is an infinitary connective, and φ_i is a formula for $1 \leq i \leq n$;
5. $[\![\forall x \in \mathsf{S}[\varphi]]\!]_{\mathcal{X}} (\bar{y}) = \sup_{x \in \mathsf{S}(\mathcal{X})} [\![\varphi]\!]_{\mathcal{X}} (x, \bar{y})$, where φ is a formula;
6. $[\![\exists x \in \mathsf{S}[\varphi]]\!]_{\mathcal{X}} (\bar{y}) = \inf_{x \in \mathsf{S}(\mathcal{X})} [\![\varphi]\!]_{\mathcal{X}} (x, \bar{y})$, where φ is a formula;

If $x_1, ..., x_n \in \mathcal{X}$, then we will write $\mathcal{X} \models \varphi(x_1, ..., x_n)$ if $[\![\varphi]\!]_{\mathcal{X}} (x_1, ..., x_n) = 0$. We also use the notation $\mathcal{X} \models [\varphi(x_1, ..., x_n) = r]$ for $[\![\varphi]\!]_{\mathcal{X}} (x_1, ..., x_n) = r$, $\mathcal{X} \models [\varphi(x_1, ..., x_n) \leq r]$ for $[\![\varphi]\!]_{\mathcal{X}} (x_1, ..., x_n) \leq r$ and $\mathcal{X} \models [\varphi(x_1, ..., x_n) \geq r]$ for $[\![\varphi]\!]_{\mathcal{X}} (x_1, ..., x_n) \geq r$ ♣

We pause here to note that given the definition above, the presence of $\wedge$ and $\vee$ in the formal definition of $\mathscr{L}$ is unnecessary. Since $\max : \mathfrak{R} \times \mathfrak{R} \to \mathfrak{R}$ is continuous, for any formulae φ and ψ, $\max\{\varphi(x), \psi(y)\}$ is also a formula, and for any structure $\mathcal{X}$, we have

$$\mathcal{X} \models \forall xy[(\varphi(x) \wedge \psi(y)) = \max\{\varphi(x), \psi(y)\}].$$

However, the presence of $\wedge$ and $\vee$ makes the formulae of $\mathfrak{R}$-valued languages look more like classical first-order formulae.

3.2.3 Fact: *Let $\varphi(y)$ be a formula, and let S be the sort of y. Then $\mathcal{X} \models \forall y \in \mathsf{S}[\varphi(y)]$ if and only if for every $y \in \mathsf{S}(\mathcal{X})$, $[\![\varphi]\!]_{\mathcal{X}} (y) = 0$. $\mathcal{X} \models \exists y \in \mathsf{S}[\varphi(y)]$ if and only if for every $\varepsilon > 0$, there is y_ε such that $[\![\varphi]\!]_{\mathcal{X}} (y_\varepsilon) \leq \varepsilon$*

3.2.4 Definition: A formula $\varphi(x)$ such that $\mathcal{X} \models \exists x \in \mathsf{S}[\varphi(x)]$ will be referred to as *approximately satisfiable in* $\mathcal{X}$. ♣

If Σ is a set of sentences, then we write $\mathcal{X} \models \Sigma$ if $\mathcal{X} \models \varphi$ for every $\varphi \in \Sigma$. Σ is *consistent* if there is a structure $\mathcal{X}$ such that $\mathcal{X} \models \Sigma$. A formula of the form $\varphi(x)$ is called *satisfiable* if there is a structure $\mathcal{X}$ and an element $x \in \mathcal{X}$ such that $\mathcal{X} \models \varphi(x)$. If $\Sigma(x)$ is a set of formulae of the form $\varphi(x)$, and $x \in \mathcal{X}$, then we write $\mathcal{X} \models \Sigma(x)$ to indicate that $\mathcal{X} \models \varphi(x)$ for every $\varphi \in \Sigma$. In this case, the set Σ will also be referred to as *satisfiable*, or *consistent*.

3.2.5 Definition: Let $\mathcal{X}$ and $\mathcal{Y}$ be $\mathscr{L}$-structures. A *function* $f : \mathcal{X} \to Y$ is a collection $(\mathsf{S}(f) : \mathsf{S} \in \mathbf{Sort})$ such that for every $\mathsf{S} \in \mathbf{Sort}$, $\mathsf{S}(f) : \mathsf{S}(\mathcal{X}) \to \mathsf{S}(\mathcal{Y})$. ♣

3.2.6 Definition: Let $\mathcal{X}$ and $\mathcal{Y}$ be $\mathscr{L}$-structures, and suppose $i : \mathcal{X} \to \mathcal{Y}$ is a function. Then i a *homomorphism* if for every $\mathsf{c} \in \mathbf{Const}$, $\mathsf{f} \in \mathbf{Func}$ and $\mathsf{R} \in \mathbf{Rel}$, we have:

1. $\mathsf{type}(\mathsf{c})(i)([\![\mathsf{c}]\!]_{\mathcal{X}}) = [\![\mathsf{c}]\!]_{\mathcal{Y}}$
2. For every $x \in \mathrm{dom}(\mathsf{f})(\mathcal{X})$, $[\![\mathsf{f}]\!]_{\mathcal{Y}}(\mathrm{dom}(\mathsf{f})(i)(x)) = \mathrm{codom}(\mathsf{f})(i)([\![\mathsf{f}]\!]_{\mathcal{X}}(x))$
3. For every $x \in \mathrm{dom}(\mathsf{R})(\mathcal{X})$, $[\![\mathsf{R}]\!]_{\mathcal{Y}}(\mathrm{dom}(\mathsf{R})(i)(x)) = [\![\mathsf{R}]\!]_{\mathcal{X}}(x))$

The function i will be referred to as an:

1. *embedding* if it is 1-1;
2. *elementary embedding* if it preserves the values of every formula, i.e. if φ is an $\mathscr{L}$-formula, then $[\![\varphi]\!]_{\mathcal{Y}}(\mathrm{dom}(\varphi)(i)(x)) = [\![\varphi]\!]_{\mathcal{X}}(x)$ for every $x \in \mathrm{dom}(\varphi)(\mathcal{X})$. ♣

3.2.7 Definition: If $\mathcal{X}$ and $\mathcal{Y}$ are $\mathscr{L}$-structures, then we write $\mathcal{X} \preceq Y$ in case there is an elementary embedding $h : \mathcal{X} \to \mathcal{Y}$. In this case we refer to $\mathcal{X}$ as being an *elementarily substructure of* $\mathcal{Y}$, and $\mathcal{Y}$ is an *elementary extension* of $\mathcal{X}$. ♣

3.2.8 Definition: Let $\mathcal{X}$ be an $\mathscr{L}$-structure. The *elementary diagram* of $\mathcal{X}$ is the set of all sentences $\varphi \in \mathscr{L}_{\mathcal{X}}$ such that $\mathcal{X} \models \varphi$. It is denoted $\mathrm{diag}(\mathcal{X})$. ♣

3.2.9 Fact: *There is an elementary map $f : \mathcal{X} \to \mathcal{Y}$ if and only if $\mathcal{Y}$ can be expanded to an $\mathscr{L}_{\mathcal{X}}$-structure $\mathcal{Y}'$, and $\mathcal{Y}' \models \mathrm{diag}(\mathcal{X})$.*

Now is perhaps the best time to discuss cardinality issues related to languages. Let $\mathscr{L}$ be a language, and let Σ be a set of sentences of $\mathscr{L}$. A subset $D \subseteq \mathscr{L}$ is called Σ-dense if and only if for every $\varphi \in \mathscr{L}$, there is a sequence $\varphi_n \in D$ such that for every n, $\Sigma \models |\varphi_n - \varphi| \leq 1/2^n$. The density character of $\mathscr{L}$ relative to Σ is the size of a smallest possible Σ-dense subset of $\mathscr{L}$. It will be denoted $|\mathscr{L}|_{\Sigma}$, of $|\mathscr{L}|$ if Σ is understood. As for metric spaces, every time we refer to the size, power or cardinality of a language, we mean its density character, and use the notation $|\mathscr{L}|$.

Section 3.3

Theories and types

Let $\mathscr{L}$ be a language, and let $\mathcal{X}$ be an $\mathscr{L}$-structure. The *complete theory of* $\mathcal{X}$ will consist of those sentences $\varphi \in \mathscr{L}$ such that $\mathcal{X} \models \varphi$. Formally,

$$\mathrm{Th}(\mathcal{X}) \stackrel{\text{def}}{:=} \{\varphi \in \mathscr{L} : \varphi \text{ is a sentence and } \mathcal{X} \models \varphi\}$$

A *theory* is any consistent set of sentences of $\mathscr{L}$. Clearly, if Σ is consistent, then by definition there is a structure $\mathcal{X}$ such that $\Sigma \subseteq \mathrm{Th}(\mathcal{X})$. If Δ and Σ are two sets of sentences, then we write $\Delta \models \Sigma$ to indicate that whenever $\mathcal{X} \models \Delta$, we also have $\mathcal{X} \models \Sigma$. Σ is called *closed* in case $\varphi \in \Sigma$ if and only if $\Sigma \models \varphi$. The closure of a set of sentences Δ is the smallest closed set $\Sigma \subseteq \mathscr{L}$ such that $\Delta \models \Sigma$. The theory $\mathrm{Th}(\mathcal{X})$ is an example of a closed theory.

3.3.1 Definition: Let K be a class of $\mathscr{L}$-structures. We define

$$\mathrm{Th}(K) \stackrel{\text{def}}{:=} \bigcap_{\mathcal{X}\in K} \mathrm{Th}(\mathcal{X}).$$

If T is any theory, then we define $T^+ \stackrel{\text{def}}{:=} \{\varphi \leq \delta : \varphi \in T, \delta > 0\}$. In general, if $\Sigma(x)$ is any set of formulae, then $\Sigma^+(x) \stackrel{\text{def}}{:=} \{\varphi(x) \leq \delta : \varphi \in \Sigma(x), \delta > 0\}$. ♣

3.3.2 Definition: We denote by $\mathrm{Mod}(T)$ the class of models of T. The notation $\mathrm{Mod}(T, \kappa)$ will be used to denote the class of models of T of density character κ. A class C of structures is called *elementary* if there is a theory T such that $C = \mathrm{Mod}(T)$. In this case we will say that T *axiomatizes* K. In other words, K is elementary if and only if $K = \mathrm{Mod}(\mathrm{Th}(K))$ ♣

We also use $\mathrm{Mod}(T)$ to denote the following category:

Objects: Models of T

Morphisms: Elementary maps between models.

Composition: Function composition.

3.3.3 Definition: Two structures $\mathcal{X}$ and $\mathcal{Y}$ are called *elementarily equivalent* if $\mathrm{Th}(\mathcal{X}) = \mathrm{Th}(\mathcal{Y})$. This situation is denoted by $\mathcal{X} \equiv \mathcal{Y}$. ♣

3.3.4 Definition: Let x be a tuple of variables. A *partial type* in the variable x is a consistent set $\pi(x)$ of formulae in $\mathscr{L}$. If $\mathcal{M}$ is an $\mathscr{L}$-structure, and $A \subseteq \mathcal{M}$, then we call a type in $\mathscr{L}_A$ a *partial type over* A. If $a \in \mathcal{M}$, then the *complete type of* a *over* A is the type $\{\varphi \in \mathscr{L}_A : \mathcal{M} \models \varphi(a)\}$. It will be denoted by $\mathrm{tp}_{\mathcal{M}}(a/A)$, or by $\mathrm{tp}(a/A)$ if $\mathcal{M}$ is understood. We write $a \equiv_A b$ to indicate that $\mathrm{tp}(a/A) = \mathrm{tp}(b/A)$. ♣

3.3.5 Definition: A type $\pi(x_1, ..., x_n)$ is *complete* if and only if for every formula of $\mathscr{L}_A$, there is a *unique* $r \in \mathfrak{R}$ such that "$\varphi(x_1, ..., x_n) = r$" $\in \pi$. The set of all complete types over A in the n variables $x_1, ..., x_n$ is denoted $S_n(A)$. Here we allow $n = \omega$. ♣

We note that a type is complete if and only if for every $r \in \mathfrak{R}$, and every $\varphi \in \mathscr{L}$, "$\varphi \leq r$" $\in \pi$ if and only if $\pi \models \varphi \leq r$, and "$\varphi \geq r$" $\in \pi$ if and only if $\pi \models \varphi \geq r$.

3.3.6 Definition: A partial type $\pi(x)$ over A is *realized* if and only if there is an $\mathscr{L}$-structure $\mathcal{X}$ and $x \in \mathcal{X}$ such that $\mathcal{X} \models \pi(x)$. A structure $\mathcal{X}$ is κ*-saturated* if and only if for every $A \subseteq \mathcal{X}$ such that $|A| < \kappa$, every partial type $\pi(x)$ over A is realized in $\mathcal{X}$. $\mathcal{X}$ is *saturated* if and only if it is $|\mathcal{X}|$-saturated. ♣

3.3.7 Theorem: *For every n, the sets of the form $[\pi(x_1, ..., x_n)] = \{p \in S_n(A) : \pi \subseteq p\}$, where $\pi(x_1, ..., x_n)$ is a partial type, form a basis of closed sets for a compact Hausdorff topology on $S_n(A)$.*

Proof: The definition of $[\pi(x_1, ..., x_n)]$ makes it clear that the sets of the form $[\pi(x_1, ..., x_n)]$ are closed under arbitrary intersections and finite unions. Therefore, they form a basis of closed sets for a topology. The fact that this topology is compact is a direct consequence of the compactness theorem for $\mathfrak{R}$-valued logic (Theorem 3.4.2, which we prove later). □

Let $\varphi \in \mathscr{L}$ be a formula in the free variables $x_1, ..., x_n$. Then φ induces a continuous map $f_\varphi : S_n(A) \to \mathfrak{R}$ which assigns to every complete type p the unique $r \in \mathfrak{R}$ such that $p \models \varphi = r$. The set $\mathscr{L}^* \stackrel{\text{def}}{:=} \{f_\varphi : \varphi \in \mathscr{L}\}$ is closed under the logical connectives in the following sense: if $0 \leq k \leq \omega$, $u : \mathfrak{R}^k \to \mathfrak{R}$ is a logical connective, and $(\varphi_i : i < k)$ is a sequence of formulae in $\mathscr{L}$ which all share the same free variables, then

$$u(f_{\varphi_1}, ..., f_{\varphi_n}, ...) = f_{u(\varphi_1,...,\varphi_n,...)} \tag{3.1}$$

The following theorem appears in [BUar] in the form of Fact 1.3 and Proposition 1.4. For completeness, we provide a direct proof here which is a direct instantiation of the proofs of Fact 1.3 and Proposition 1.4 in [BUar]. Though lacking the generality of the proof in [BUar], our proof highlights the construction of the formula corresponding to a continuous function $f : S_n(A) \to \mathfrak{R}$.

3.3.8 Theorem: *For every $\mathscr{L}$-formula φ, the function f_φ defined above is continuous. Furthermore, every continuous map $f : S_n(A) \to \mathfrak{R}$ is of the form f_φ for some $\mathscr{L}$-formula φ.*

Proof: First we show that f_φ defined above is continuous. Let $M > 0$. Then $f_\varphi^{-1}([0, M]) = \{p : f_\varphi(p) \leq M\}$ by definition. By definition also, $f_\varphi(p) \leq M$ if and only if "$\varphi \leq M$" $\in p$. This is true if and only if $p \in [\pi]$, where $\pi = \{\varphi \leq M\}$. Similarly, $f^{-1}([M, \infty]) = [\pi]$, where $\pi = \{\varphi \geq M\}$.

We now prove the second assertion. Note that given any two distinct types $p, q \in S_n(A)$, there are $r \neq s \in \mathfrak{R}$, and a formula φ such that "$\varphi = r \in p$ and "$\varphi = s$" $\in q$. This is because, if p and q are distinct, then there must exist by definition a formula φ such that $p \models \varphi$, and $q \models \varphi \geq 1/2^n$ for some n.

CLAIM A: *For every $r \neq s \in \mathfrak{R}$, there is a formula ψ such that $f_\psi(p) = r$ and $f_\psi(q) = s$.*

PROOF: Let $u(t) = \dfrac{s - r}{f_\varphi(q)} t + r$, and let $\psi = u \circ \varphi$. Note that u is a logical connective, and consequently, ψ is a formula. Also, it is easy to see that $f_\psi = f_{u\circ\varphi} = u \circ f_\varphi$. Therefore, $f_\psi(p) = r$ and $f_\psi(q) = s$ as required. □

By the claim, for every pair of types p and q, we can find a formula $\varphi_{p,q}$ such that $f_{\varphi_{p,q}}(p) = f(p)$ and $f_{\varphi_{p,q}}(q) = f(q)$. Let $\varepsilon > 0$, and define $V_{p,q} = \{z \in S^N(A) : f(z) - \varepsilon < f_{\varphi_{p,q}}(z)\}$. Note that $V_{p,q}$ is an open neighbourhood of q, and that the collection $C_p = \{V_{p,q} : q \in S^n(A)\}$ is an open covering of $S_n(A)$. Let F_p be a finite sub-covering of C_p, and let

$$\psi_p = \bigwedge_{V_{p,q} \in F_p} \varphi_{p,q}.$$

Now $f(z) - \varepsilon < f_{\psi_p}(z)$ for every z. The set $U_p = \{z \in S_n(A) : f_{\varphi_p}(z) < f(z) + \varepsilon\}$ is an open neighbourhood of p, and the collection $C = \{U_p : p \in S_n(A)\}$ is an open cover of $S_n(A)$. Let F be a finite subcover of C, and let

$$\psi_\varepsilon = \bigvee_{U_p \in C} \psi_p.$$

Then $f(z) - \varepsilon < f_{\psi_\varepsilon}(z) < f(z) + \varepsilon$. In other words, for every z, $|f(z) - f_{\psi_\varepsilon}(z)| < \varepsilon$. Since ε was arbitrary, this shows that f is a uniform limit of functions of the form f_φ. To see why this implies that f has the appropriate form, let φ_n be a sequence of formulas such that $f_{\varphi_n} \to f$ uniformly. There is a uniformly continuous function $u : \mathfrak{R}^\omega \to \mathfrak{R}$ such that

$$f = \lim_{n\to\infty} f_{\varphi_n} = u(f_{\varphi_1}, ..., f_{\varphi_n}, ...) = f.$$

By Equation (3.1),

$$u(f_{\varphi_1}, ..., f_{\varphi_n}, ...) = f_{u(\varphi_1,...,\varphi_n,...)}$$

. Therefore, the formula $\varphi = u(\varphi_1, ..., \varphi_n, ...)$ is the formula we are looking for. □

3.3.9 Corollary: *Let $\mathcal{M} \in \mathrm{Mod}(T)$, and let $f : \mathsf{S}(\mathcal{M}) \to \mathfrak{R}$ be a uniformly continuous map on some sort of $\mathscr{L}$. Then the following are equivalent:*

1. *There is a formula φ such that $f = [\![\varphi]\!]_{\mathcal{M}}$;*
2. *$f(a) = f(b)$ whenever $a \equiv b$.*

Proof: The facts that $f(a) = f(b)$ whenever $a \equiv b$ and that f is uniformly continuous imply that f induces a continuous function f on the type space S_1. By Theorem 3.3.8, f is of the form f_φ. □

Section 3.4

Ultraproducts of structures

We now define ultraproducts of $\mathfrak{R}$-valued structures. Let (I, U) be an ultrafilter pair, and let $\{\mathcal{X}_i : i \in I\}$ be a set of $\mathscr{L}$-structures. We put an $\mathscr{L}$-structure on $\mathcal{X} = \prod_U \mathcal{X}_i$ as follows:

1. $\mathsf{S}(\mathcal{X}) = \prod_U \mathsf{S}(\mathcal{X}_i)$ (ultraproduct of sets)
2. $[\![c]\!]_{\mathcal{X}} = \langle [\![c]\!]_{\mathcal{X}_i} \rangle_U$
3. For every tuple of terms $t_1, ..., t_n$ of the appropriate codomains,

$$[\![R]\!]_{\mathcal{X}}(t_1, ..., t_n) = \lim_{i\to U} [\![R]\!]_{\mathcal{X}_i}([\![t_1]\!]_{\mathcal{X}_i}, ..., [\![t_n]\!]_{\mathcal{X}_i})$$

4. For every tuple of terms $t_1, ..., t_n$ of the appropriate codomains,

$$[\![f]\!]_{\mathcal{X}}(t_1, ..., t_n) = \langle [\![f]\!]_{\mathcal{X}_i}([\![t_1]\!]_{\mathcal{X}_i}, ..., [\![t_n]\!]_{\mathcal{X}_i}) \rangle_U$$

This structure is called the *ultraproduct* of the structures $\mathcal{X}_i$. If all the structures $\mathcal{X}_i$ are the same, then $\prod_U \mathcal{X}_i$ is denoted $\mathcal{X}^U$, and called an *ultrapower* of $\mathcal{X}$. The fundamental theorem of ultraproducts stated below, as well as its proof, can be found in [BBHU08]. For completeness, and because of the theorem's importance in the rest of this thesis, we provide a complete direct proof.

3.4.1 Theorem (Fundamental theorem of ultraproducts): *Let $\{\mathcal{X}_i : i \in I\}$ be an I-indexed sequence of $\mathscr{L}$-structures, and let $\mathcal{X} = \prod_U \mathcal{X}_i$. Then for every formula φ, $[\![\varphi]\!]_{\mathcal{X}} = \lim_{i\to U}([\![\varphi]\!]_{\mathcal{X}})$.*

Proof: Since the logical connectives correspond to continuous functions, we get by induction that if $1 \leq \kappa \leq \omega$, $(\varphi_i : i < \kappa)$ is a sequence of formulae, and u is a κ-ary connective, then

$$\begin{aligned}
[\![u(\varphi_1, ..., \varphi_n, ...)]\!]_{\mathcal{X}} &= u([\![\varphi_1]\!]_{\mathcal{X}}, ..., [\![\varphi_n]\!]_{\mathcal{X}}, ...) && \text{(Definition)}\\
&= u(\lim_{i\to U} [\![\varphi_1]\!]_{\mathcal{X}_i}, ..., \lim_{i\to U} [\![\varphi_n]\!]_{\mathcal{X}_i}, ...) && \text{(Definition)}\\
&= \lim_{i\to U} u([\![\varphi_1]\!]_{\mathcal{X}_i}, ..., [\![\varphi_n]\!]_{\mathcal{X}_i}, ...) && \text{(Theorem 2.4.1)}\\
&= \lim_{i\to U} [\![u(\varphi_1, ..., \varphi_n, ...)]\!]_{\mathcal{X}_I} && \text{(Definition)}
\end{aligned}$$

If φ is $\forall x[\psi]$: By definition, $[\![\forall x[\psi]]\!]_{\mathcal{X}} = \sup_x [\![\psi]\!]_{\mathcal{X}}$, and by induction, $[\![\psi]\!]_{\mathcal{X}} = \lim_{i\to U} [\![\psi]\!]_{x_i}$. Suppose $\lim_{i\to U} \sup_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) = r$, and let $r \in O$. It is enough to consider the cases where O is a basic open set, either of the form $[0, M)$ or (M, ∞). If O is of the form (M, ∞), then there is $\langle x_i\rangle_U \in \mathcal{X}$ such that $\lim_{i\to U} [\![\psi]\!]_{\mathcal{X}_i}(x_i) > M$. By the definition of ultralimits, the set $P = \{i \in I : [\![\psi]\!]_{\mathcal{X}_i}(x_i) > M\} \in U$, so that for every $i \in P$, $\sup_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) > M$. Therefore, the set $\{i \in I : \sup_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) > M\} \in U$.

If O is of the form $[0, M)$, then for every $\langle x_i\rangle_U \in \mathcal{X}$, $\lim_{i\to U} [\![\psi]\!]_{\mathcal{X}_i}(x_i) < M$, so that the set $P = \{i \in I : [\![\psi]\!]_{\mathcal{X}_i}(x_i) < M\} \in U$ for every sequence $(x_i \in \mathcal{X}_i : i \in I)$. For every $i \in P$, $\sup_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) < M$, and therefore, the set $\{i \in I : \sup_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) < M\} \in U$.

If φ is $\exists x[\psi]$: By definition, $[\![\exists x[\psi]]\!]_{\mathcal{X}} = \inf_x [\![\psi]\!]_{\mathcal{X}}$, and by induction, $[\![\psi]\!]_{\mathcal{X}} = \lim_{i\to U} [\![\psi]\!]_{x_i}$. Suppose $\lim_{i\to U} \inf_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) = r$, and let $r \in O$. It is enough to consider the cases where O is a basic open set, either of the form $[0, M)$ or (M, ∞). If O is of the form $[0, M)$, then there is $\langle x_i\rangle_U \in \mathcal{X}$ such that $\lim_{i\to U} [\![\psi]\!]_{\mathcal{X}_i}(x_i) < M$. By the definition of ultralimits, the set $P = \{i \in I : [\![\psi]\!]_{\mathcal{X}_i}(x_i) < M\} \in U$, so that for every $i \in P$, $\inf_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) < M$. Therefore, the set $\{i \in I : \inf_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) < M\} \in U$.

If O is of the form (M, ∞), then for every $\langle x_i\rangle_U \in \mathcal{X}$, $\lim_{i\to U} [\![\psi]\!]_{\mathcal{X}_i}(x_i) > M$, so that the set $P = \{i \in I : [\![\psi]\!]_{\mathcal{X}_i}(x_i) > M\} \in U$ for every sequence $(x_i \in \mathcal{X}_i : i \in I)$. For every $i \in P$, $\inf_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) > M$, and therefore, the set $\{i \in I : \inf_{x_i} [\![\psi]\!]_{\mathcal{X}_i}(x_i) > M\} \in U$. □

3.4.2 Theorem (Compactness theorem): *Let K be a class of $\mathscr{L}$-structure, and let Σ be a set of $\mathscr{L}$-formulae. Suppose that for every finite $\Delta \subseteq \Sigma$ there is an $\mathscr{L}$-structure $\mathcal{X} \in K$ such that $\mathcal{X} \models \Delta$. Then there is a set I, an ultrafilter U on I and structures $\mathcal{X}_i \in K$ for $i \in I$ such that $\prod_U \mathcal{X}_i \models \Sigma$.*

Proof: The compactness theorem stated in this form appears in [BBHU08], and is key to the proof of strong conceptual completeness. Therefore, we shall give a detailed proof of it here. For every finite set $\Delta \subseteq \Sigma$, let $\mathcal{X}_\Delta \models \Delta$ be given. Let $I = \{\Delta \subseteq \Sigma : \Delta \text{ finite}\}$. For every $\varphi \in \Sigma$, let $[\varphi] \stackrel{\text{def}}{:=} \{\Delta \in I : \varphi \in \Delta\}$. Let $F = \{\varphi : \varphi \in \Sigma\}$, and note that F has the finite intersection property. Let U be any ultrafilter extending F, and let $M = \prod_U M_\Delta$. We claim that $\mathcal{X} \models \Sigma$. Let $\varphi \in \sigma$. By the fundamental theorem of ultraproducts, $[\![\varphi]\!]_{\mathcal{X}} = \lim_{\Delta\to U} [\![\varphi]\!]_{\mathcal{X}_\Delta}$. By the definition of U, $[\varphi] \in U$, and by the definition of $\mathcal{X}_\Delta$, we have $\mathcal{X}_\Delta \models \varphi$ for every Δ such that $\varphi \in \Delta$. In particular, $[\![\varphi]\!]_{\mathcal{X}_\Delta} = 0$ for every $\Delta \in [\varphi] \in U$. Therefore, $[\![\varphi]\!]_{\mathcal{X}} = \lim_{\Delta\to U} [\![\varphi]\!]_{\mathcal{X}_\Delta} = 0$ as required. □

In continuous logic, there is also an approximate form of the compactness theorem, which is sometimes more useful in practice. It follows easily from the compactness theorem above.

3.4.3 Corollary (Approximate compactness theorem): *Let Σ be a set of sentences of $\mathscr{L}$, and suppose that for every finite $\Delta \subseteq \Sigma$, and every $\varepsilon > 0$ there is an $\mathscr{L}$-structure $\mathcal{M}_{\Delta,\varepsilon}$ such that $\mathcal{M}_{\Delta,\varepsilon} \models \{\varphi \leq \varepsilon : \varphi \in \Delta\}$. Then there is an $\mathscr{L}$-structure $\mathcal{M}$ such that $\mathcal{M} \models \Sigma$.*

Proof: By compactness, for every $\varepsilon > 0$, the set $\{\varphi \leq \varepsilon : \varphi \in \Sigma\}$ has a model $\mathcal{M}_\varepsilon$. Therefore, so does the set Σ^+. Let $\mathcal{M} \models \Sigma^+$. We claim that $\mathcal{M} \models \Sigma$. Since $\mathcal{M} \models \Sigma^+$, for every $\varphi \in \Sigma$, and every $\varepsilon > 0$, we have $\mathcal{M} \models \varphi \leq \varepsilon$. Thus $[\![\varphi]\!]_{\mathcal{M}} \leq \varepsilon$ for every ε, so $[\![\varphi]\!]_{\mathcal{M}} = 0$ as required. □

3.4.4 Theorem: *Let $\Sigma(x_1, ..., x_n)$ be a set of formulae in the free variables $x_1, ..., x_n$. Suppose that for every finite $\Delta(x_1, ..., x_n) \subseteq \Sigma(x_1, ..., x_n)$, there is an $\mathscr{L}$-structure $\mathcal{M}_\Delta$ and $m_{\Delta,i} \in \mathcal{M}_\Delta$ for $i = 1, ..., n$ such that $\mathcal{M}_\Delta \models \Delta(m_{\Delta,1}, ..., m_{\Delta,n})$. Then there is an $\mathscr{L}$-structure $\mathcal{M}$ and $m_1, ..., m_n \in \mathcal{M}$ such that $\mathcal{M} \models \Sigma(m_1, ..., m_n)$.*

3.4.5 Theorem: *If (U, I) is an ultrafilter pair such that U is non-principal and I is countable, then, in $\mathcal{X}^U$, every separable type is realized.*

Proof: Let $\mathcal{X}$ be an $\mathscr{L}$-structure. Let $\pi(x)$ be a separable partial type defined which is consistent with $\mathrm{Th}(\mathcal{X})$. Since π is separable, there is no loss of generality in assuming that π is actually countable. To see this, let $D \subseteq \pi$ be a countable sense subset, and suppose $x \in \mathcal{X}^U$ realizes D. Then we claim that x realizes all of π. Let $\varphi \in \pi$, and write $\varphi = \lim_{n\to\infty} \varphi_n$, where $\varphi_n \in D$. Then it is easy to see that $\mathcal{X}^U \models \varphi(x)$.

Write $\pi = \bigcup_{n\in\mathbb{N}} \Delta_n$, with Δ_n finite for every $n \in \mathbb{N}$, and with the property that $\Delta_\ell \subseteq \Delta_n$ if $\ell \leq n$. For every n, let $\Delta_{n,m} = \{\varphi \leq 1/m : \varphi \in \Delta_n\}$. By the approximate compactness theorem, the consistency of π is equivalent to the existence of, for every $m, n \in \mathbb{N}$, an element $x_{n,m}$ realizing $\Delta_{n,m}$. Let

$$f : I \to \{\Delta_{n,m} : n, m \in \mathbb{N}\}$$

be any surjective function, and consider the element $\bar{x} = \langle x_{f(i)} : i \in I\rangle_U$. Then $\bar{x}$ realizes π in $\mathcal{X}^U$. □

3.4.6 Definition: An $\mathscr{L}$-structure $\mathcal{X}$ is called *κ-universal* if and only if for every $\mathscr{L}$-structure $\mathcal{Y}$ such that $\mathcal{Y} \equiv \mathcal{X}$ and $|\mathcal{Y}| < \kappa$, there is an elementary embedding $f : \mathcal{Y} \to \mathcal{X}$. ♣

3.4.7 Theorem: *Suppose $\alpha \geq |\mathscr{L}|$, let U be an α-regular ultrafilter on I. Then for every $\mathscr{L}$-structure $\mathcal{X}$, the ultrapower $\mathcal{X}^U$ is α^+-universal.*

Proof: Fix an $\mathscr{L}$-structure $\mathcal{X}$. Let I be a set, and U be an α-regular ultrafilter on I. Let $V \subseteq U$ be such that $|V| = \alpha$, and for every $i \in I$, the set $\{P \in V : i \in P\}$ is finite. Let $\mathcal{Y}$ be an $\mathscr{L}$-structure of density character α, and suppose $\mathcal{X} \equiv \mathcal{Y}$. We need to find an elementary embedding of $\mathcal{Y}$ into $\mathcal{X}^U$.

Consider the language $\mathscr{L}_\mathcal{Y}$, and let Δ be the $\mathscr{L}_\mathcal{Y}$-elementary diagram of $\mathcal{Y}$. Since $|\mathscr{L}| \leq \alpha$ and $|\mathcal{Y}| \leq \alpha$, we have $|\Delta| \leq |\mathscr{L}_\mathcal{Y}| \leq \alpha$, and therefore there is a injective function $h : \Delta \times \mathbb{N} \to V$. By the choice of V, for every $i \in I$, the set $\{P \in V : i \in P\}$ is finite. Therefore, the set

$$\Sigma(i, n) = \{\varphi \leq 1/2^n : i \in h(\varphi, n)\}$$

is finite as well for every $i \in I$ and every $n \in \mathbb{N}$. For every $\varphi \in \Sigma(i, n)$, let $\varphi_\mathscr{L}$ be the $\mathscr{L}$-formula obtained from φ by replacing every instance of a constant symbol of the form y for $y \in \mathcal{Y}$ by a fresh variable x_y. Then we have that $\varphi_\mathscr{L}(\mathsf{y}_1, ..., \mathsf{y}_1) = \varphi$.

Since $\mathcal{Y} \models \varphi$, $\mathcal{Y} \models \exists x_1 ... x_n[\varphi_{\mathscr{L}}(x_1, ..., x_n)]$. Since we are assuming that $\mathcal{X} \equiv \mathcal{Y}$, $\mathcal{X} \models \exists x_1, ..., x_n[\varphi_{\mathscr{L}}(x_1, ..., x_n)]$ as well. Therefore, for every k, there is a tuple $x_1, ..., x_n \in \mathcal{X}$ such that $\varphi_{\mathscr{L}}(x_1, ..., x_n) \leq 1/2^k$. Let $f_{i,k} : \mathcal{Y} \to \mathcal{X}$ be any function such that $f_{i,k}(y_j) = x_j$ for $1 \leq j \leq n$.

Let $f_k : \mathcal{Y} \to \mathcal{X}^U$ be $\prod_U f_{i,k}$. We claim that for every $\mathscr{L}$-formula $\varphi(x)$ and any $y \in \mathcal{Y}$ of the appropriate type, $|\, [\![\varphi]\!]_{\mathcal{Y}}(y) - [\![\varphi]\!]_{\mathcal{X}^U}(f_k(y))| \leq 1/2^k$. Let $\varphi(y)$ be an $\mathscr{L}$-formula, and let $y \in \mathcal{Y}$ be of the appropriate type. Suppose $[\![\varphi]\!]_{\mathcal{Y}}(y) = 0$. Then $\varphi(\mathrm{y}) \in \Delta$. By definition, $\varphi(f_i(y)) \leq 1/2^k$ for every $i \in h(\varphi, k) \in U$, so $\varphi(\langle f_i(y)\rangle_U) = \lim_{i \to U} \varphi(f_i(y)) \leq 1/2^k$. We finish the proof by noting that the sequence f_k is uniformly convergent, and take f to be its limit. □

3.4.8 Corollary (Frayne's theorem, [BS69]): *Let $\mathcal{X}$ and $\mathcal{Y}$ be $\mathscr{L}$-structures. Then $\mathcal{X} \equiv \mathcal{Y}$ if and only if there is a set I, an ultrafilter U on I, and an elementary embedding $h : \mathcal{Y} \to \mathcal{X}^U$*

Proof: If $h : \mathcal{Y} \to \mathcal{X}^U$ is an elementary embedding, then the conclusion is clear, since then $\mathcal{Y} \equiv \mathcal{X}^U \equiv \mathcal{X}$. For the converse, apply Theorem 3.4.7 with I the elementary $\mathscr{L}_{\mathcal{Y}}$-diagram of $\mathcal{Y}$, U any α-regular ultrafilter on I, where $\alpha > |\mathcal{Y}|$, to get an elementary embedding $\mathcal{Y} \to \mathcal{X}^U$. □

3.4.9 Proposition (Scott's lemma, [BS69]): *Let $\mathcal{X}$ and $\mathcal{Y}$ be $\mathscr{L}$-structures, and suppose $h : \mathcal{X} \to \mathcal{Y}$ is a map. Then h is an elementary embedding if and only if there is a set I, an ultrafilter U on I, and an elementary embedding $g : \mathcal{Y} \to \mathcal{X}^U$ making the diagram:*

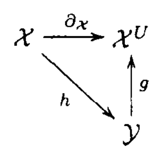

commute.

Proof: The proof we give here parallels the proof given in [BS69], with some simplification coming from the use of regular ultrafilters in the argument from [CK90]. Let $h : \mathcal{X} \to \mathcal{Y}$ be a function. Let $g : \mathcal{Y} \to \mathcal{X}^U$ be an elementary embedding such that $g \circ h = \partial_{\mathcal{X}}$, and let $\varphi(x)$ be any $\mathscr{L}$-formula. Since g is an elementary embedding, we get

$$\begin{aligned} [\![\exists x[\varphi]]\!]_{\mathcal{X}} &= [\![\exists x[\varphi]]\!]_{\mathcal{X}^U}(\partial_{\mathcal{X}}(x)) \\ &= [\![\exists x[\varphi]]\!]_{\mathcal{X}^U}(g(h(x))) \\ &= [\![\exists x[\varphi]]\!]_{\mathcal{Y}}(h(x)) \end{aligned}$$

where the first equality holds because the map $\partial_{\mathcal{X}}$ is elementary, the second by the assumption that $\partial_{\mathcal{X}} = g \circ h$, and the last equality holds by the assumption that g is elementary. This shows that h satisfies the Tarski-Vaught test for continuous logic, and therefore it is an elementary embedding.

For the converse, assume $h : \mathcal{X} \to \mathcal{Y}$ is an elementary embedding. Consider the languages $\mathscr{L}_{h[\mathcal{X}]} \subseteq \mathscr{L}_{\mathcal{Y}}$. In a nutshell, the strategy of the proof is to extend both $\mathcal{X}$ and $\mathcal{Y}$ to $\mathscr{L}_{\mathcal{Y}}$-structures $\mathcal{X}_{\mathcal{Y}}$ and $\mathcal{Y}_{\mathcal{Y}}$ in such a way that the $\mathscr{L}_{h[\mathcal{X}]}$-reducts of $\mathcal{X}_{\mathcal{Y}}$ and $\mathcal{Y}_{\mathcal{Y}}$ are elementarily equivalent. We can then apply Frayne's lemma to $\mathcal{X}$ and $\mathcal{Y}$ as $\mathscr{L}_{h[\mathcal{X}]}$-structures to get an embedding $g : \mathcal{Y} \to \mathcal{X}^U$, and prove that $g \circ h = \partial_{\mathcal{X}}$ by inspecting the definition of g. Let $x^+ \in \mathcal{X}$ be a fixed element. For every $y \in \mathcal{Y}$, we consider the corresponding constant symbol y of $\mathscr{L}_{\mathcal{Y}}$, and define

$$[\![\mathrm{y}]\!]'_{\mathcal{Y}} = \begin{cases} h(x^+) & \text{if } y \notin h[\mathcal{X}] \\ y & \text{if } y \in h[\mathcal{X}] \end{cases}$$

and $[\![\mathrm{y}]\!]'_{\mathcal{X}} = g^{-1}([\![\mathrm{y}]\!]'_{\mathcal{Y}})$. The interpretations $[\![\cdot]\!]'_{\mathcal{X}}$ and $[\![\cdot]\!]'_{\mathcal{Y}}$ make $\mathcal{X}$ and $\mathcal{Y}$ into $\mathscr{L}_{\mathcal{Y}}$-structure with elementary equivalent $\mathscr{L}_{h(\mathcal{X})}$-reducts. We can now assume without loss of generality that $\mathcal{X}$ is $\aleph_1$-saturated as an $\mathscr{L}_{h[\mathcal{X}]}$-structure, passing to a suitable ultrapower modulo a non-principal ultrafilter over $\mathbb{N}$ if necessary.

Let I be the elementary $\mathscr{L}_{\mathcal{Y}}$-diagram of $\mathcal{Y}$, and U be an $|I|$-regular ultrafilter on I. Let $\varphi \in I$, and let $\mathrm{y}_1, ..., \mathrm{y}_k$ enumerate those constant symbols of $\mathscr{L}_{\mathcal{Y}}$ which are not in $\mathscr{L}_{h[\mathcal{X}]}$. Let φ' be the $\mathscr{L}_{h[\mathcal{X}]}$-formula obtained by replacing every instance of y_i with a fresh variable x_i. Since $\mathcal{Y} \models \varphi$, we have $\mathcal{Y} \models \exists x_1, ..., x_k[\varphi']$. Since we are assuming $\mathcal{X}$ is $\aleph_1$-saturated, and $\mathcal{X} \equiv \mathcal{Y}$, there are $x_1, ..., x_k \in \mathcal{X}$ such that $[\![\varphi']\!]_{\mathcal{X}}(x_1, ..., x_n) = 0$. Let $f_\varphi : \mathcal{Y} \to \mathcal{X}$ be defined as follows:

$$f_\varphi(y) = \begin{cases} [\![\mathrm{y}]\!]'_{\mathcal{X}} & \text{if } y \in h[\mathcal{X}] \\ x_i & \text{if } \mathrm{y} = \mathrm{y}_i \end{cases}$$

and define $g : \mathcal{Y} \to \mathcal{X}^U$ to be $[y \mapsto \langle f_\varphi(y) \rangle_U]$. Then g is an elementary map. If $y = h(x)$, then for every $\varphi \in I$, we have $f_\varphi(y) = [\![\mathrm{y}]\!]'_{\mathcal{X}} = h^{-1}([\![\mathrm{y}]\!]'_{\mathcal{Y}}) = h^{-1}(y) = h^{-1}(h(x)) = x$. This shows that $g(h(x)) = \partial_{\mathcal{X}}(x)$ for every $x \in \mathcal{X}$ as required. □

3.4.10 Theorem: *Let K be a class of $\mathscr{L}$-structures, and suppose K is closed under elementary substructures, isomorphisms and the formation of ultraproducts. Then there is a theory T such that $K = \mathrm{Mod}(T)$.*

Proof: Let $T = \mathrm{Th}(K)$. We show that $K = \mathrm{Mod}(T)$. It is clear that $K \subseteq \mathrm{Mod}(T)$. Let $\mathcal{M} \in \mathrm{Mod}(T)$. We need to show that $\mathcal{M} \in K$. By the approximate compactness theorem, all we need to do is show that $\mathrm{Th}(\mathcal{M})^+$ is finitely satisfiable in K. Suppose not. Then there is a finite list $\varphi_1, ..., \varphi_n \in \mathscr{L}$ of sentences, and $\varepsilon > 0$ such that whenever $\mathcal{M} \models \varphi_i$ for $i = 1, ..., n$, but for every $\mathcal{N} \in K$, $\mathcal{N} \models [\varphi_i \geq \varepsilon]$ for some $1 \leq i \leq n$. Therefore, the condition $(\varphi_1 \wedge \cdots \wedge \varphi_n) \geq \varepsilon \in T$, but $\mathcal{M} \models (\varphi_1 \wedge \cdots \wedge \varphi_n)$, so $\mathcal{M} \notin \mathrm{Mod}(T)$, contradiction.

By the compactness theorem, there is a set I and an ultrafilter U on I such that for some collection $\{\mathcal{M}_i : i \in I\} \subseteq K$, $\prod_U \mathcal{M}_i \models \mathrm{Th}(\mathcal{M})$. By assumption, K is closed under ultraproducts, so we get $\prod_U \mathcal{M}_i \in K$. Also, $\mathcal{M} \equiv \prod_U \mathcal{M}_i = \mathcal{N}$, since each $\mathcal{M}_i \models \mathrm{Th}(\mathcal{M})$. By Scott's lemma, there is an embedding $\mathcal{M} \to \mathcal{N}^U$ for some ultrafilter U. Since K is closed under ultraproducts and $\mathcal{N} \in K$, we have $\mathcal{N}^W \in K$, and since $\mathcal{M} \preceq \mathcal{M}^V$ and K is closed under elementary embeddings, $\mathcal{M} \in K$, which is what we wanted. □

3.4.11 Theorem: *Let $\mathscr{L} \subseteq \mathscr{L}'$ be languages. Let T be an $\mathscr{L}$-theory, and let K be an elementary class of $\mathscr{L}'$-structures such that $\mathcal{X}^{\mathscr{L}} \models T$ for every $\mathcal{X} \in K$. Suppose that for every ultrafilter pair (I, U), and every family $\{\mathcal{X}_i : i \in I\}$ of $\mathscr{L}'$-structures such that $\mathcal{X}_i^{\mathscr{L}} \models T$, if $\prod_U \mathcal{X}_i \in K$, then for some $P \in U$, $\mathcal{X}_i \in K$ for every $i \in P$. Then there is a* finite *theory T' such that $K = \mathrm{Mod}(T \cup T')$.*

Proof: There is a theory T'' such that $K = \mathrm{Mod}(T \cup T'')$. Assume that T'' is not finitely axiomatizable. Then for every finite $\Delta \subseteq T''$, there is an $\mathscr{L}$-structure $\mathcal{X}_\Delta \notin K$ such that $\mathcal{X}_\Delta \models T \cup \Delta$. Let I be the set of all finite subsets of T, and U be an ultrafilter on I. Let

$\mathcal{M} = \prod_U \mathcal{X}_\Delta$. Then $\mathcal{M} \models T \cup T''$, so $\mathcal{M} \in \mathrm{Mod}(T \cup T'')$. However, for any $P \in U$, and any element $\Delta \in P$, we have $\mathcal{X}_\Delta \notin \mathrm{Mod}(T \cup T'')$, contradicting the hypothesis of the theorem.□

Here may be the best time to point out that any classical first-order language (*resp.* theory) can be realized as an $\mathfrak{R}$-valued language (*resp.* theory). The usual definition of a first-order language uses exactly the same data as a $\mathfrak{R}$-valued one. Therefore, any first-order language can be thought of as an $\mathfrak{R}$-valued language without any modification.

To a first-order theory we add, for every relation symbol R, the axiom

$$(\mathsf{R}(x) = 0) \vee (\mathsf{R}(x) = \infty)$$

which ensures that the relation symbols assume only two values. An easy induction on the construction of formulae shows that if $u : \mathfrak{R}^n \to \mathfrak{R}$ is a *finitary* connective, and $\varphi_1, \dots, \varphi_n$ is a sequence of n finiraty formulae, then $u(\varphi_1, \dots, \varphi_n)$ will, in any structure, assume only finitely many values.

We note that the connective $[x \mapsto 1/x] : \mathfrak{R} \to \mathfrak{R}$ has the property that $u(0) = \infty$ and $u(\infty) = 0$. It can thus act as a negation operator. In fact, any connective u with the property that $u(0) = \infty$ and $u(\infty) = 0$ can be used as a negation operator.

Since the finitary formulas form a dense subset of the set of all formulas, we see that the finite-valued formulas form a dense subset of the set of formulas of a first-order language viewed as a $\mathfrak{R}$-valued language.

3.4.12 Example (Fields of characteristic 0): This example will seem a bit artificial, but will illustrate Theorem 3.4.11. Let K be the class of fields of characteristic 0. This is an elementary class, but it is not finitely axiomatizable. Let $\{F_i : i \in \mathbb{N}\}$ be a countable collection of fields, such that $char(F_i) = p_n$, where p_n is the n-th prime number. Let U be any ultrafilter extending the Fréchet filter on $\mathbb{N}$. Then $\prod_U F_i$ is a field of characteristic 0. However, none of the fields F_i have characteristic 0. A similar argument can be used to show that the theory of algebraically closed fields cannot be finitely axiomatized. ♢

Section 3.5

Metric structures

We now define the class of $\mathfrak{R}$-valued structures we will be studying in the rest of this thesis. Informally speaking, $\mathfrak{R}$-valued languages are languages lacking a symbol for equality, and metric structures remedy this situation. We introduce metric languages and theories as expansions of $\mathfrak{R}$-valued languages by a special symbol d standing in for a metric, and metric theories as theories which have the property that the extra symbol d is a metric, and all the relations and function symbols are uniformly continuous. This all turns out to be expressible as first-order sentences. In practice, an existing symbol of $\mathscr{L}$ is singled out and used as a metric symbol.

3.5.1 Definition: A *metric language* is an $\mathfrak{R}$-valued language $\mathscr{L}$ together with the following data:

1. For every sort S of $\mathscr{L}$, a relation symbol d_S with domain $\mathsf{S} \times \mathsf{S}$;
2. For every relation symbol $\mathsf{R} \in \mathbf{Rel}_\mathscr{L}$, a function $UC(\mathsf{R}) : \mathfrak{R} \to \mathfrak{R}$;
3. For every function symbol $\mathsf{f} \in \mathbf{Func}_\mathscr{L}$, a function $UC(\mathsf{f}) : \mathfrak{R} \to \mathfrak{R}$; ♣

3.5.2 Definition: A *metric theory* T in a metric language $\mathscr{L}$ is an $\mathscr{L}$-theory with the following properties:

1. For every sort S of $\mathscr{L}$, $T \models \forall x[\mathsf{d}_\mathsf{S}(x,x)]$;
2. For every sort S of $\mathscr{L}$, $T \models \forall xyz[\mathsf{d}_\mathsf{S}(x,y) \leq (\mathsf{d}_\mathsf{S}(x,z) + \mathsf{d}_\mathsf{S}(y,z))]$;
3. For every relation symbol $\mathsf{R} \in \mathbf{Rel}_\mathscr{L}$, and every $\varepsilon > 0$, writing $\mathrm{dom}(\mathsf{R}) = \mathsf{S}_1 \times \cdots \times \mathsf{S}_n$,
$$T \models \forall x_1 \cdots x_n y_1 \cdots y_n \left[\left(\bigwedge_{1\leq i\leq n} \mathsf{d}_{\mathsf{S}_i}(x_i, y_i) \geq UC(\mathsf{R})(\varepsilon) \right) \vee (|\mathsf{R}(\bar{x}) - \mathsf{R}(\bar{y})| \leq \varepsilon) \right]$$
4. For every function symbol $\mathsf{f} \in \mathbf{Func}_\mathscr{L}$, and every $\varepsilon > 0$, writing $\mathrm{dom}(\mathsf{R}) = \mathsf{S}_1 \times \cdots \times \mathsf{S}_n$, and d for $\mathsf{d}_{\mathrm{codom}(\mathsf{f})}$,
$$T \models \forall x_1 \cdots x_n y_1 \cdots y_n \left[\left(\bigwedge_{1\leq i\leq n} \mathsf{d}_{\mathsf{S}_i}(x_i, y_i) \geq UC(\mathsf{f})(\varepsilon) \right) \vee (\mathsf{d}(\mathsf{f}(\bar{x}), \mathsf{f}(\bar{y})) \leq \varepsilon) \right].$$

3.5.3 Definition: The *uniform continuity modulus* $UC(\varphi)$ of a formula φ is a function $\mathfrak{R} \to \mathfrak{R}$ defined as follows: the uniform continuity modulus of a term is given by

$$\begin{aligned} UC(x)(\varepsilon) &= \varepsilon \quad \text{(if } x \text{ is a variable)} \\ UC(\mathsf{c})(\varepsilon) &= \varepsilon \quad \text{(if } \mathsf{c} \text{ is a constant symbol)} \\ UC(\mathsf{f}(t_1,...,t_n))(\varepsilon) &= \min\{UC(t_1)(UC(\mathsf{f})(\varepsilon)),...,UC(t_n)(UC(\mathsf{f})(\varepsilon))\} \end{aligned}$$

and $UC(\varphi)$ is defined by induction on the complexity of φ as follows:

1. $UC(\mathsf{R}(t_1,...,t_n))(\varepsilon) = \min\{UC(\mathsf{R})(UC(t_1)(\varepsilon)),...,UC(\mathsf{R})(UC(t_n)(\varepsilon))\}$
2. $UC(\varphi_1 \,\mathrm{op}\, \varphi_2)(\varepsilon) = \min\{UC(\varphi_1)(\varepsilon), UC(\varphi_2)(\varepsilon)\})$, where $\mathrm{op} \in \{\wedge, \vee\}$;
3. $UC(\forall x \in \mathsf{S}[\varphi]) = UC(\exists x \in \mathsf{S}[\varphi]) = UC(\varphi)$
4. $UC(u(\varphi_1,...,\varphi_n,...))(\varepsilon) = \min\{UC(u)(UC(\varphi_1)(\varepsilon)),...,UC(u)(UC(\varphi_n)(\varepsilon_n)),...)$, where $UC(u)$ is the function defined in definition 2.2.7 ♣

3.5.4 Theorem: *Let $\mathscr{L}$ be a metric language, and let T be a metric theory. Suppose $\mathcal{X} \models T$, then:*

1. *For every sort S of $\mathscr{L}$, $[\![\mathsf{d}_\mathsf{S}]\!]_\mathcal{X}$ is a generalized pseudo-metric;*
2. *For every term t, $[\![t]\!]_\mathcal{X}$ is a pseudo-uniformly continuous function*
$$\mathrm{dom}(t)(\mathcal{X}) \to \mathrm{codom}(t)(\mathcal{X})$$
and $UC(t)$ is a uniform continuity modulus for t;
3. *For every formula φ, $[\![\varphi]\!]_\mathcal{X}$ is a pseudo-uniformly continuous function* $\mathrm{dom}(\varphi)(\mathcal{X}) \to \mathfrak{R}$*, and $UC(\varphi)$ is a uniform continuity modulus for φ;*

Proof: The proof that the interpretation of a term or a formula yields a uniformly continuous map is done by induction on the complexity of terms and formulae. If t is of the form $\mathsf{f}(x_1, ..., x_n)$, where x_i is either a variable or a constant, then $[\![t]\!]$ is uniformly continuous by definition. Otherwise, t is of the form $\mathsf{f}(t_1, ..., t_n)$, and by induction, $[\![t_i]\!]_{\mathcal{X}}$ is uniformly continuous, and $UC(t_i)$ is a uniform continuity modulus for t_i. By definition,

$$UC(t) = \min\{UC(t_1)(UC(\mathsf{f}(\varepsilon)), ..., UC(t_n)(UC(\mathsf{f}(\varepsilon))\}$$

so if $\varepsilon > 0$, then choosing $\delta = UC(t)$, we have that $\delta \leq UC(t_i)(UC(\mathsf{F}(\varepsilon))$ for $i = 1, ..., n$. This implies that if $\bar{x}$ varies by δ, then $t_i(\bar{x})$ varies by at most $UC(\mathsf{F}(\varepsilon)$ for $i = 1, ..., n$, so that $\mathsf{f}(t_1, ..., t_n)$ varies by at most ε, proving that $[\![t]\!]_{\mathcal{X}}$ is uniformly continuous.

The proof that the interpretation of formulae is uniformly continuous is similarly done by induction on the complexity of φ. By definition relation symbols are uniformly continuous, and by an argument similar to the one in the previous paragraph, formulae of the form $u(\varphi_1, ..., \varphi_n, ...)$ are also uniformly continuous. Suppose $\varphi(y)$ is of the form $\forall x[\psi(x, y)]$. By induction, ψ is uniformly continuous, and $UC(\psi)$ is a uniform continuity modulus for it. Let $\varepsilon > 0$, and let $\delta = UC(\psi)(\varepsilon)$. Let $d(y, y') < \delta$. Then for every x, $d((x, y), (x, y')) < \delta$ as well, so $|\, [\![\psi]\!]_{\mathcal{X}} (x, y) - [\![\psi]\!]_{\mathcal{X}} (x, y')| \leq \varepsilon$ for every x. This means that $\sup_x [\![\psi]\!]_{\mathcal{X}} (x, y) - [\![\psi]\!]_{\mathcal{X}} (x, y')| \leq \varepsilon$. Now note that

$$|\sup_x [\![\psi]\!]_{\mathcal{X}} (x, y) - \sup_x [\![\psi]\!]_{\mathcal{X}} (x, y')| \leq |\sup_x [\![\psi]\!]_{\mathcal{X}} (x, y) - [\![\psi]\!]_{\mathcal{X}} (x, y')|.$$

The proof for the case where $\varphi(y) = \exists x[\psi(x, y)]$ is similar. □

3.5.5 Definition: Let $\mathscr{L}$ be a metric language, and let T be a metric theory. If $\mathcal{X} \models T$, then $\mathcal{X}$ will be called a *pre-model* of T. If $[\![\mathsf{d_S}]\!]_{\mathcal{X}}$ is a complete metric for every sort S of $\mathscr{L}$, then we call $\mathscr{L}$ a *model* of T. ♣

If $\mathcal{X}$ is a pre-model, then there is a corresponding model $\mathcal{X}/\mathsf{d}$ which we describe below. In the following, the notation $[x]$ is used to represent the equivalence class of x modulo an equivalence relation which will be clear from the context, and C is the completion functor described in section 2.2

1. For every sort S, $\mathsf{S}(\mathcal{X}/\mathsf{d}) = C(\mathsf{S}(\mathcal{M})/\, [\![\mathsf{d_S}]\!]_{\mathcal{X}})$
2. For every relation symbol R, $[\![\mathsf{R}]\!]_{\mathcal{X}/\mathsf{d}}$ is $C(f)$, where $f([x_1], ..., [x_n]) = [[\![\mathsf{R}]\!]_{\mathcal{X}} (x_1, ..., x_n)]$
3. For every function symbol f, $[\![\mathsf{f}]\!]_{\mathcal{X}/\mathsf{d}}$ is $C(g)$, where $g([x_1], ..., [x_n]) = [[\![\mathsf{f}]\!]_{\mathcal{X}} (x_1, ..., x_n)]$
4. For every constant symbol c, $[\![\mathsf{c}]\!]_{\mathcal{X}/\mathsf{d}} = [[\![\mathsf{c}]\!]_{\mathcal{X}}]$

By definition, $[\![\mathsf{d_S}]\!]_{\mathcal{X}/\mathsf{d}}$ is a metric on $\mathsf{S}(\mathcal{M})/\, [\![\mathsf{d_S}]\!]_{\mathcal{X}}$ for every $\mathsf{S} \in \mathbf{Sort}_{\mathscr{L}}$. Let d represent the unique lift of $[\![\mathsf{d_S}]\!]_{\mathcal{X}/\mathsf{d}}$ to the completion of $\mathsf{S}(\mathcal{M})/\, [\![\mathsf{d_S}]\!]_{\mathcal{X}}$. Then d is a metric. It is clearly satisfies all the properties of a pseudo-metric. Suppose $d(x, y) = 0$. By definition, $x = \lim_{n\to\infty} x_n$ and $y = \lim_{n\to\infty} y_n$, where for every n, $x_n, y_n \in \mathsf{S}(\mathcal{M})/\, [\![\mathsf{d_S}]\!]_{\mathcal{X}}$. Since d is continuous,

$$d\left(\lim_{n\to\infty} x_n, \lim_{n\to\infty} y_n\right) = \lim_{m,n\to\infty} d(x_n, y_m) = \lim_{m,n\to\infty} [\![\mathsf{d_S}]\!] (x_n, y_m) = 0$$

showing that $x = y$. This shows that $\mathcal{X}/\mathsf{d}$ is indeed a model. Note that if $\mathcal{X}$ is already a model, then $\mathcal{X}/\mathsf{d} = \mathcal{X}$.

Section 3.6

Definable sets

We now define the $\mathfrak{R}$-valued analogue of the notion of definable set. In classical first order logic, a subset $D \subseteq \mathcal{X}$ is called *definable* if and only if there is an $\mathscr{L}$-formula φ such that $D = \{x : \mathcal{X} \models \varphi(x)\}$. The direct translation of this concept yields:

3.6.1 Definition: A subset $Z \subseteq \mathcal{X}$ is called a *zero-set* if and only if there is a formula φ such that $Z = \{x : [\![\varphi]\!]_{\mathcal{X}}(x) = 0\}$. ♣

While every definable set in $\mathfrak{R}$-valued logic should be of this form, we unfortunately cannot consider all these sets as being definable. We consider the following, more general definition of the concept of definable set.

3.6.2 Definition: Let $\mathcal{X}$ be an $\mathscr{L}$-structures, and $A \subseteq \mathcal{X}$. The set $D \subseteq \mathcal{X}$ is *definable over* A if and only if for every $\mathscr{L}_A$-formula ψ, there are $\mathscr{L}_A$-formulae $\forall x \in D[\psi]$ and $\exists x \in D[\psi]$ such that

1. $[\![\forall x \in D[\psi]]\!]_{\mathcal{X}} = \sup_{x \in D} [\![\psi]\!]_{\mathcal{X}}(x)$
2. $[\![\exists x \in D[\psi]]\!]_{\mathcal{X}} = \inf_{x \in D} [\![\psi]\!]_{\mathcal{X}}(x)$ ♣

The above definition is general enough to make sense in general $\mathfrak{R}$-valued languages. However, in a general $\mathfrak{R}$-valued language, there may not exist any non-trivial definable sets. Note that by definition, all the sorts of $\mathscr{L}$ are definable sets. In [Ben], an example is given of a theory with very few definable sets. If we are given a metric language $\mathscr{L}$, and a metric theory T, then we get the following criterion to decide when a set is definable. We start by stating an easy technical lemma about uniformly continuous functions. Its use is not strictly necessary, but will greatly simplify the proof.

3.6.3 Fact: *A function $f : (X, d_X) \to (Y, d_Y)$ between two generalized metric spaces X and Y is uniformly continuous if and only if there is a logical connective u such that $u(0) = 0$, and $d_Y(f(x), f(y)) \leq u(d_X(x, y))$.*

Proof: A proof of this fact about uniformly continuous functions can be found, for the case of ordinary bounded metric spaces, in [BBHU08, Proposition 2.10]. The proof for the case of generalized metric spaces is similar, and we include it here for completeness. Let $g : \mathfrak{R} \to \mathfrak{R}$ be defined by

$$g(\varepsilon) = \sup\{\delta : d_Y(f(x), f(y)) < \varepsilon \text{ whenever } d_X(x, y) < \delta\}$$

and note that since f is uniformly continuous, $g(\varepsilon) > 0$ if $\varepsilon > 0$. Furthermore, g is increasing, $g(0) = 0$, $g(\infty) = \infty$, and $d_Y(f(x).f(y)) \leq g(d_X(x, y))$ for every $x, y \in X$. To complete the proof, we define a continuous function $u : \mathfrak{R} \to \mathfrak{R}$ such that $u(\varepsilon) \geq g(\varepsilon)$ for every ε. Let $(\varepsilon_n : n \in \mathbb{N})$ be any sequence such that $\varepsilon_0 = \infty$, and $\varepsilon_n \to 0$ as $n \to \infty$. For every $\varepsilon > 0$ there either are unique n and $t \in (0, 1)$ such that $\varepsilon = t\varepsilon_{n+1} + (1-t)\varepsilon_n$, or there is a unique n such that $\varepsilon = \varepsilon_n$ Define

$$u(\varepsilon) = \begin{cases} 0 & \text{if } \varepsilon = 0 \\ \infty & \text{if } \varepsilon = \infty \\ g(\varepsilon_{n-1}) & \text{if } \varepsilon = \varepsilon_n \\ tg(\varepsilon_{n+1}) + (1-t)g(\varepsilon_n) & \text{if } \varepsilon = t\varepsilon_{n+1} + (1-t)\varepsilon_n \end{cases}.$$

An easy argument shows that $g(\varepsilon) \leq u(\varepsilon)$ for every ε, and thus that $d_Y(f(x), f(y)) \leq u(d_X(x, y))$ for every $x, y \in X$. This completes the proof. □

3.6.4 Theorem: *Let $\mathcal{X}$ be an $\mathscr{L}$-structures, $A \subseteq \mathcal{X}$ and $D \subseteq \mathcal{X}$. Then the following are equivalent:*

1. *D is definable over A;*
2. *$\forall x \in D[\psi]$ and $\exists x \in D[\psi]$ exist for every atomic formula ψ.*
3. *$\forall x \in D[\mathsf{d}(x, y)]$ and $\exists x \in D[\mathsf{d}(x, y)]$ exist*

Proof: A proof of this theorem appears in [Ben]. Given the importance of this theorem in the rest of this thesis, we include a proof here. First note that 1⇒2⇒3 is true by our definition of definability of sets. We need to show that 3⇒1. Let $\mathcal{X}$ be an $\mathscr{L}$-structure, and $D \subseteq \mathcal{X}$. We show that we can conclude the existence of $\exists x \in D[\varphi]$ for every φ based on the existence of $\exists x \in D[\mathsf{d}(x, y)]$. The proof that $\forall x \in D[\varphi]$ exists is similar. Suppose $\forall x \in D[\mathsf{d}(x, y)]$ and $\exists x \in D[\mathsf{d}(x, y)]$ exist, and let $\varphi(x, y)$ be any formula. Since $[\![\varphi]\!]_{\mathcal{X}}$ is uniformly continuous, by Fact 3.6.3, there is a logical connective u such that

$$\mathcal{X} \models \forall xyz[|\varphi(x, z) - \varphi(y, z)| \leq u(\mathsf{d}(x, y)).$$

Since the connective u produced by Fact 3.6.3 depends only on the uniform continuity modulus of φ, and not on the particular structure used, we have that

$$T \models \forall xyz[|\varphi(x, z) - \varphi(y, z)| \leq u(\mathsf{d}(x, y)).$$

We claim that the formula

$$\exists x[\varphi(x, y) + u(\exists w \in D[\mathsf{d}(x, w)]]$$

when evaluated in any structure $\mathcal{X}$, is equal to $\inf_{x \in D} [\![\varphi]\!]_{\mathcal{X}}(x, y)$. We can then *define*

$$\exists x \in D[\varphi(x, y)] \stackrel{\text{def}}{:=} \exists x[\varphi(y, x) + u(\exists w \in D[\mathsf{d}(x, w)]].$$

By the definition of u, we have that for every x, y, z

$$[\![\varphi]\!]_{\mathcal{X}}(x, y) \leq [\![\varphi]\!]_{\mathcal{X}}(z, y) + u([\![\mathsf{d}]\!]_{\mathcal{X}}(z, x)).$$

Taking the infimum with respect to $x \in D$ on both sides of this inequality, we get

$$\inf_{x \in D} [\![\varphi]\!]_{\mathcal{X}}(x, y) \leq [\![\varphi]\!]_{\mathcal{X}}(z, y) + \inf_{x \in D} u([\![\mathsf{d}]\!]_{\mathcal{X}}(z, x))$$

and by assumption,

$$\begin{aligned} \inf_{x \in D} u([\![\mathsf{d}]\!]_{\mathcal{X}}(z, x)) &= u(\inf_{x \in D} [\![\mathsf{d}]\!]_{\mathcal{X}}(z, x)) \\ &= u([\![\exists x \in D[\mathsf{d}(z, x)]]\!]_{\mathcal{X}}) \\ &= [\![u(\exists x \in D[\mathsf{d}(z, x)])]\!]_{\mathcal{X}} \end{aligned}$$

so that

$$\begin{aligned} \inf_{x \in D} [\![\varphi]\!]_{\mathcal{X}}(x, y) &\leq [\![\varphi]\!]_{\mathcal{X}}(z, y) + [\![u(\exists x \in D[\mathsf{d}(z, x)])]\!] \\ &= [\![\varphi(z, y) + u(\exists x \in D[\mathsf{d}(z, x)])]\!] \end{aligned}$$

for every z. Therefore,

$$\begin{aligned}\inf_{x\in D}\llbracket\varphi\rrbracket_{\mathcal{X}}(x,y) &\leq \inf_{z}\llbracket\varphi(z,y)+u(\exists x\in D[\mathsf{d}(z,x)])\rrbracket(z,y)\\ &= \llbracket\exists z[\varphi(z,y)+u(\exists x\in D[\mathsf{d}(z,x)])]\rrbracket(y)\end{aligned}$$

On the other hand, if $z\in D$, then $\llbracket\exists x\in D[\mathsf{d}(z,x)]\rrbracket_{\mathcal{X}}(z)=0$, so

$$\begin{aligned}\inf_{x\in D}\llbracket\varphi\rrbracket_{\mathcal{X}}(x,y) &= \inf_{z\in D}\llbracket\varphi(z,y)+u(\exists x\in D[\mathsf{d}(z,x)])\rrbracket(z,y)\\ &\geq \inf_{z}\llbracket\varphi(z,y)+u(\exists x\in D[\mathsf{d}(z,x)])\rrbracket(z,y)\end{aligned}$$

showing that in fact,

$$\begin{aligned}\inf_{x\in D}\llbracket\varphi\rrbracket_{\mathcal{X}}(x,y) &= \inf_{z}\llbracket\varphi(z,y)+u(\exists x\in D[\mathsf{d}(z,x)])\rrbracket(z,y)\\ &= \llbracket\exists z[\varphi(z,y)+u(\exists x\in D[\mathsf{d}(z,x)])]\rrbracket(y)\end{aligned}$$

This completes the proof that $\exists x\in D[\varphi(x,y)]$ exists for every φ. □

Note that the formula $\exists x\in D[\mathsf{d}(x,y)]$, if it exists, is equal to $\inf_{x\in D}\llbracket\mathsf{d}\rrbracket_{\mathcal{X}}(x,y)$ in the structure $\mathcal{X}$, which is the definition of the "distance to D" function. We thus recover Ben Yaacov's definition of a definable set in a metric structure: D is definable if and only if there is a formula $\varphi(x)$ such that $\llbracket\varphi\rrbracket_{\mathcal{X}}(x)=d(x,D)$.

3.6.5 Definition: A *definable function* is a formula $\varphi(x,y)\in\mathscr{L}$ such that for every $\varepsilon>0$, there is $\delta>0$ such that for any structure $\mathcal{X}$

$$\mathcal{X}\models\forall xzz'\left[\varphi(x,z)\wedge\varphi(x,z')\geq\delta\right)\vee d(z,z')\leq\varepsilon\right].$$

In other words, in any structure $\mathcal{X}$, the set $\{(x,y):\mathcal{X}\models\varphi(x,y)\}$ is the graph of a function. If φ is a definable function, then we will write $[\varphi(x)=y]$ instead of $\varphi(x,y)$. ♣

Note that there there is no confusion of terminology here, as the graph of a definable function is a definable set in the sense described above. The interested reader can consult [Ben] for a proof.

Section 3.7

Imaginaries and the eq-construction

In this section we describe imaginary sorts and the eq-expansion of structures. The construction of $\mathscr{L}^{\mathrm{eq}}$ for $\mathfrak{R}$-valued logic is very close to its classical counterpart. There is, however, a key difference. In our setup, we do not obtain all the sorts of $\mathscr{L}^{\mathrm{eq}}$ as canonical parameter sorts. Whether it is even possible to achieve the construction given below using only canonical parameters is unknown to the author. The procedure we describe in this section can be found in [Har]. For this section, we fix a metric language $\mathscr{L}$, and a metric theory T. We will distinguish three types of imaginary sorts:

1. First and foremost, canonical parameters sorts of $\mathscr{L}$-formulae;
2. Finite and countable products of sorts of $\mathscr{L}$. These *can* be realized as canonical parameters, but we shall keep them separate for simplicity;

3. Finite unions of canonical parameter sorts;

In each of these cases, we show that we can expand $\mathscr{L}$ to a language $\mathscr{L}'$ sporting a new sort symbol, and the theory T to a theory T' in such a way that the forgetful functor $\mathrm{Mod}(T') \to \mathrm{Mod}(T)$ is an equivalence of categories. First, we deal with products of sorts of $\mathscr{L}$. Let $1 \leq n \leq \omega$. Recall that by a *compact norm* on $\mathfrak{R}^n$, we mean a norm which induces the compact product topology on $\mathfrak{R}^n$. All these norms are equivalent, but there is no canonical choice of one. Therefore, we can only describe a product of sorts of $\mathscr{L}$ with the data of a fixed compact norm.

3.7.1 Theorem: *Let $1 \leq n \leq \omega$, and let S_i be a sort of $\mathscr{L}$ for $i \leq n$. Let $u : \mathfrak{R}^n \to \mathfrak{R}$ be a compact norm. Let $\mathscr{L}_{\prod_{i<n} \mathsf{S}_i^u}$ be an expansion of $\mathscr{L}$ with:*

1. *A new sort symbol $\prod_{i<n}^u \mathsf{S}_i$ with metric symbol d;*
2. *For $1 \leq i \leq n$, a new function symbol $\pi_i : \prod_{i<n}^u \mathsf{S}_i \to \mathsf{S}_i$;*

Let $T_{\prod_{i<n}^u \mathsf{S}_i}$ be the following expansion of T:

P1 $\forall xy[\mathsf{d}(x,y) = u(d_{\mathsf{S}_1}(\pi_1(x), \pi_1(y)), ..., d_{\mathsf{S}_n}(\pi_n(x), \pi_n(y), ...)]$

P2$_n$ $\forall x_1 \cdots x_n \exists y \left[\bigwedge_{i \leq n} \mathsf{d}_{\mathsf{S}_i}(\pi_i(y), x_i) \right]$

Then the forgetful functor $F : \mathrm{Mod}\left(T_{\prod_{i<n}^u \mathsf{S}_i}\right) \to \mathrm{Mod}(T)$ is an equivalence of categories.

Proof: Let $u : \mathfrak{R}^n \to \mathfrak{R}$ be a compact norm, $1 \leq n \leq \omega$, and $(\mathsf{S}_i : i < n)$ be sorts of $\mathscr{L}$. To simplify the notation, let T' denote $T_{\prod_{i<n} \mathsf{S}_i^u}$, and let S' denote $\prod_{i<n}^u \mathsf{S}_i$. Define $U : \mathrm{Mod}(T) \to \mathrm{Mod}(T')$ as follows:

1. For every sort S of $\mathscr{L}$, $\mathsf{S}(U(\mathcal{M})) = \mathsf{S}(\mathcal{M})$
2. $\mathsf{S}'(U(\mathcal{M})) = \prod_{i<n} \mathsf{S}_i(\mathcal{M})$
3. For every $i < n$, $[\![\pi_i]\!]_{U(\mathcal{M})}(x_1, ..., x_n, ...) = x_i$
4. $[\![\mathsf{d}]\!]_{U(\mathcal{M})}(\bar{x}, \bar{y}) = u([\![\mathsf{d}_{\mathsf{S}_1}]\!](x_1, y_1), ..., [\![\mathsf{d}_{\mathsf{S}_n}]\!](x_n, y_n), ...)$

It is clear that $U(\mathcal{M}) \models T'$ for every $M \in \mathrm{Mod}(T)$, and that $F(U(\mathcal{M})) = \mathcal{M}$ for every $\mathcal{M} \in \mathrm{Mod}(T)$. This shows that F is essentially surjective on objects. We show that $U(F(\mathcal{M}')) \cong \mathcal{M}'$ for every $\mathcal{M}' \in \mathrm{Mod}(T')$. To do this all we need to do is show that the interpretation of S' in $\mathcal{M}'$ is the actual direct product of the interpretations of the sorts S_i in $\mathcal{M}'$. The presence of π_i for $i < n$ implies the existence of a map $f : \mathsf{S}'(\mathcal{M}') \to \prod_{i<n} \mathsf{S}_i(\mathcal{M}')$. We show that this map is in fact an isometry. Let

$$\rho(\bar{x}, \bar{y}) = u([\![\mathsf{d}_{\mathsf{S}_1}]\!](x_1, y_1), ..., [\![\mathsf{d}_{\mathsf{S}_1}]\!](x_1, y_1), ...).$$

Then ρ defines a metric on $\prod_{i \leq n} \mathsf{S}_i(\mathcal{M}')$. By **P1**,

$$\rho(f(x), f(y)) = [\![\mathsf{d}]\!]_{\mathcal{M}'}(x, y)$$

for every $x, y \in \mathsf{S}'(\mathcal{M})$ which means that f is an isometry. To show that f is onto, let $(x_1, ..., x_k, ...) \in \prod_{i<n} \mathsf{S}_i(\mathcal{M}')$. For every $k < n$, there is $y_k \in \mathsf{S}'(\mathcal{M}')$ such that $\pi_i(y_k) = x_k$ for every $i \leq k$. We claim that the sequence $(y_k : k < n)$ is Cauchy. Let $\varepsilon > 0$. Since u is a compact norm, there is $N < n$ and $\delta > 0$ such that for every sequence $t_1, ..., t_k, ...$, if $t_1, ..., t_N < \delta$, $u(t_1, ...t_k, ...) < \varepsilon$. Let $N < k < \ell$. Then for every $i \leq N$, by definition of y_k, we have $\mathsf{d}_{\mathsf{S}_i}(\pi_i(y_k), \pi_i(y_\ell)) = 0$, so that

$$\begin{aligned} \mathsf{d}(y_k, y_\ell) &= u(\mathsf{d}_{\mathsf{S}_1}(\pi_1(y_k), \pi_1(y_\ell)), ..., d_{\mathsf{S}_n}(\pi_n(y_k), \pi_n(y_\ell), ...) \\ &= u(\underbrace{0, ..., 0}_{k\text{times}}, \mathsf{d}_{\mathsf{S}_{k+1}}(\pi_1(y_k), \pi_1(y_\ell)), ...) \\ &< \varepsilon \end{aligned}$$

showing that the sequence y_k is Cauchy. Let $y = \lim_{k\to\infty} y_k$. Then it is easy to see that $f(y) = (x_1, ..., x_k, ...)$, showing that f in surjective. This proves that $U(F(\mathcal{M}')) \cong \mathcal{M}'$.

Faithfulness is a consequence of the fact that the map $f : \mathsf{S}'(\mathcal{M}') \to \prod_{i\leq n} \mathsf{S}_i(\mathcal{M}')$ defined above is in fact a definable map $f : \mathsf{S}' \to \mathsf{S}_1 \times \cdots \times \mathsf{S}_n \times \cdots$, where the product on the right represents the tuple of sorts $(\mathsf{S}_1, \cdots, \mathsf{S}_n, ...)$. Indeed, f is defined by the formula

$$f(x, x_1, ..., x_n, ...) = u(\mathsf{d}_{\mathsf{S}_1}(\pi_1(x), x_1), ..., \mathsf{d}_{\mathsf{S}_n}(\pi_n(x), x_n), ...).$$

Therefore, f defines an elementary map between the functors S' and $\prod_{i\leq n} \mathsf{S}_i$, which we view as functors $\mathrm{Mod}(T) \to \mathbb{MET}$. Since h' and $U(h)$ are both elementary maps,

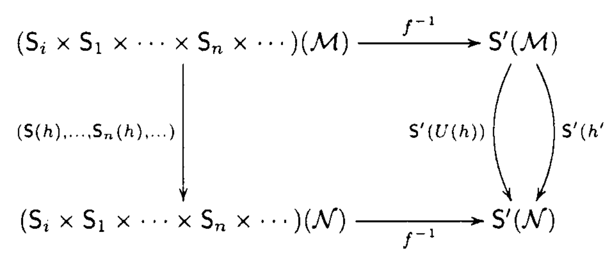

is commutative, showing that $h' = U(h)$, and that F is faithful, and completing the proof that it is an equivalence of categories. ☐

3.7.2 Proposition: *Let $\varphi(x, y)$ be a formula, and let*

$$\rho(z, z') \stackrel{\text{def}}{:=} \forall x[\varphi(x, z) - \varphi(x, z')].$$

Then for any $\mathscr{L}$-structure $\mathcal{X}$, $[\![\rho]\!]_{\mathcal{X}}(z, z')$ is a pseudo-metric.

Proof: By definition $[\![\rho]\!]_{\mathcal{X}}(z, z') = \sup_{x\in\mathrm{dom}(\varphi)} |\, [\![\varphi]\!]_{\mathcal{X}}(x, z) - [\![\varphi]\!]_{\mathcal{X}}(x, z')|$. It is clear that $[\![\rho]\!]_{\mathcal{X}}(z, z) = 0$, and that $[\![\rho]\!]_{\mathcal{X}}(z, z') = [\![\rho]\!]_{\mathcal{X}}(z', z)$. It remains to show that the triangle inequality holds. Let $z, z', z'' \in \mathrm{dom}(\varphi)(\mathcal{M})$. Then

$$|\, [\![\varphi]\!]_{\mathcal{X}}(x, z) - [\![\varphi]\!]_{\mathcal{X}}(x, z'')| \leq |\, [\![\varphi]\!]_{\mathcal{X}}(x, z) - [\![\varphi]\!]_{\mathcal{X}}(x, z')| + |\, [\![\varphi]\!]_{\mathcal{X}}(x, z') - [\![\varphi]\!]_{\mathcal{X}}(x, z'')|$$

for every x. Therefore, taking the supremum over all x, we get the desired result. ☐

3.7.3 Theorem: *Let $\mathscr{L}$ be a continuous language, and let $\varphi \in \mathscr{L}$. Write* $\mathrm{dom}(\varphi) = \mathsf{S}' \times \mathsf{S}$, *with S and S' two sorts of $\mathscr{L}$. Let $\mathscr{L}_\varphi$ be an expansion of $\mathscr{L}$ with:*

1. *A new sort symbol* S_φ;
2. *A new function symbol* $\pi_\varphi : \mathsf{S} \to \mathsf{S}_\varphi$;

Let T_φ *be the following expansion of* T*:*

1. $\forall xy\big[\mathsf{d}(\pi_\varphi(x), \pi_\varphi(y)) = \forall z|\varphi(x,z) - \varphi(y,z)|\big]$
2. $\forall x \exists y\big[d_{\mathsf{S}}(\pi_\varphi(y), x)\big]$

Then the forgetful functor $F : \mathrm{Mod}(T_\varphi) \to \mathrm{Mod}(T)$ *is an equivalence of categories.*

Proof: Let φ be given, and let $\rho(z, z')$ denote the pseudo-metric $\forall x|\varphi(x,z) - \varphi(x,z')|$. Define the functor $U : \mathrm{Mod}(T) \to \mathrm{Mod}(T_\varphi)$ as follows:

1. For every sort S of $\mathscr{L}$, $\mathsf{S}(U(\mathcal{M}) = \mathsf{S}(\mathcal{M})$
2. $\mathsf{S}_\varphi(U(\mathcal{M})) = \mathsf{S}(\mathcal{M}) / [\![\rho]\!]_{\mathcal{M}}$
3. $[\![\pi_\varphi]\!]_{U(\mathcal{M})}$ is the canonical quotient map $\mathsf{S}(\mathcal{M}) \to \mathsf{S}_\varphi(U(\mathcal{M}))$

and note that the axioms of T_φ force the definition of the interpretation of the metric symbol d. Then for every $\mathcal{M} \in \mathrm{Mod}(T)$, $U(\mathcal{M}) \in \mathrm{Mod}(T_\varphi)$ as required, and that $F(U(\mathcal{M})) = \mathcal{M}$.

Next, we show that for every $\mathcal{M} \in \mathrm{Mod}(T_\varphi)$, we have $\mathcal{M} \cong U(F(\mathcal{M}))$. By definition, $\mathsf{S}_\varphi(U(F(\mathcal{M})))$ is the quotient $\mathsf{S}(F(\mathcal{M}))/[\![\rho]\!]_{F(\mathcal{M})}$. We also have that $\mathsf{S}(\mathcal{M}) = \mathsf{S}(F(\mathcal{M}))$ by definition of F, since S is a sort of $\mathscr{L}$. Also, since ρ is a formula of $\mathscr{L}$, $[\![\rho]\!]_{\mathcal{M}} = [\![\rho]\!]_{F(\mathcal{M})}$. This proves that $\mathsf{S}_\varphi(\mathcal{M}) \cong \mathsf{S}(F(\mathcal{M}))/[\![\rho]\!]_{F(\mathcal{M})} = \mathsf{S}_\varphi(U(F(\mathcal{M})))$.

We finish the proof by showing that F is full and faithful. The existence of the functor U shows that F is full. To show that F is faithful, we show that U is full. Let $\mathcal{M}, \mathcal{N} \in \mathrm{Mod}(T)$, and $h : U(\mathcal{M}) \to U(\mathcal{N})$ be $\mathscr{L}_\varphi$-elementary map. Consider the diagram:

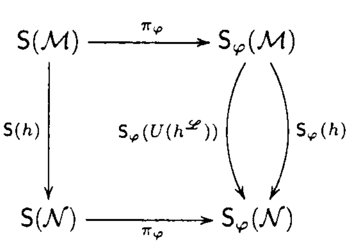

Since π_φ is definable, this diagram commutes. Therefore, $h = U(h^{\mathscr{L}})$, showing that U is full, and that F is faithful. □

The last kind of imaginary sort we introduce is one the need of which is specific to $\mathfrak{R}$-valued languages. This sort plays an important part in the proof of conceptual completeness in section 3.8 below. The next theorem shows that one can conservatively expand $\mathscr{L}$ and T to include not only canonical parameter sorts of the form S_φ, but also finite unions of such sorts. In order for the notation in Theorem 3.7.4 below to seem more natural, one should view the elements of the canonical parameter sort $\mathsf{S}_\varphi(\mathcal{M})$, where $\mathrm{dom}(\varphi) = \mathsf{S}' \times \mathsf{S}$, as functions $\mathsf{S}'(\mathcal{M}) \to \mathfrak{R}$. Each element $\pi_\varphi(c) \in \mathsf{S}_\varphi(\mathcal{M})$ corresponds to the function $[x \mapsto \varphi(x,c)]$. There is an *evaluation map* $\mathsf{eval} : \mathsf{S}_\varphi \times \mathsf{S}' \to \mathfrak{R}$ with the property that

$$\mathsf{eval}(\pi_\varphi(x), y) = \varphi(y, x)$$

for every $x \in \mathsf{S}$ and every $y \in \mathsf{S}'$.

3.7.4 Theorem: *Let $\mathscr{L}$ be a continuous language, and let $\varphi_1, ..., \varphi_n \in \mathscr{L}$. Write* $\mathrm{dom}(\varphi_i) = \mathsf{S} \times \mathsf{S}_i$*, with* S *and* S_i *two sorts of* $\mathscr{L}$*. Let* $\mathscr{L}_{\bigcup_{i \le n} \mathsf{S}_i}$ *be an expansion of* $\mathscr{L}_{\varphi_1} \cup \cdots \cup \mathscr{L}_{\varphi_n}$ *with:*

1. *A new sort symbol* $\bigcup_{i \le n} \mathsf{S}_{\varphi_i}$ *with metric symbol* d*;*
2. *For* $k = 1, ..., n$*, a function symbol* $i_{\varphi_k} : \mathsf{S}_{\varphi_i} \to \bigcup_{i \le n} \mathsf{S}_{\varphi_i}$
3. *A map* $\mathsf{eval} : \bigcup_{i \le n} \mathsf{S}_{\varphi_i} \times \mathsf{S} \to \mathfrak{R}$

Let $T_{\bigcup_{i \le n} \mathsf{S}_i}$ *be the following expansion of* $T_{\varphi_1} \cup \cdots \cup T_{\varphi_n}$*:*

1. *For every* $i \le n$, $\forall xy[\mathsf{eval}(i_{\varphi_i}(\pi_{\varphi_i}(x)), y) = \varphi_i(y, x)]$
2. *For every* $i \le n$, $\forall xy[\mathsf{d}(i_{\varphi_i}(x), i_{\varphi_i}(y)) = \mathsf{d}_{\varphi_i}(x, y)]$*, where* d_{φ_i} *is the metric symbol on the sort* S_{φ_i}
3. $\forall xy\big[\mathsf{d}(x, y) = [\forall z[\mathsf{eval}(x, z) - \mathsf{eval}(y, z)]]\big]$

Then the forgetful functor $F : \mathrm{Mod}(T_{\bigcup_{i \le n} \mathsf{S}_i}) \to \mathrm{Mod}(T)$ *is an equivalence of categories.*

Proof: Let $1 \le n < \omega$, and $(\varphi_i : i \le n)$ be formulae of $\mathscr{L}$. To simplify the notation, let T' denote $T^u_{\bigcup_{i \le n} \mathsf{S}_{\varphi_i}}$, and let S' denote $\bigcup_{i \le n} \mathsf{S}_{\varphi_i}$. Define $U : \mathrm{Mod}(T) \to \mathrm{Mod}(T')$ as follows:

1. For every sort S of $\mathscr{L}_{\varphi_1} \cup \cdots \cup \mathscr{L}_{\varphi_n}$, $\mathsf{S}(U(\mathcal{M})) = \mathsf{S}(\mathcal{M})$
2. $\mathsf{S}'(U(\mathcal{M})) = \bigcup_{i \le n} \mathsf{S}_{\varphi_i}(\mathcal{M})$
3. For every $i \le n$, $[\![i_{\varphi_i}]\!]_{U(\mathcal{M})}$ is the canonical inclusion map $\mathsf{S}_{\varphi}(\mathcal{M}) \to \bigcup_{i \le n} \mathsf{S}_{\varphi_i}(\mathcal{M})$
4. Let $[\![\mathsf{eval}]\!]_{U(\mathcal{M})}$ be the map $\bigcup_{i \le n} \mathsf{S}_{\varphi_i}(\mathcal{M}) \times \mathsf{S}(\mathcal{M}) \to \mathfrak{R}$ which is defined as follows. If $x \in \bigcup_{i \le n} \mathsf{S}_{\varphi_i}(\mathcal{M})$, then $x = \pi_{\varphi_\ell}(x')$ for some ℓ and some $x' \in \mathrm{dom}(\varphi_\ell)(\mathcal{M})$. For any $y \in \mathsf{S}(\mathcal{M})$, define
$$[\![\mathsf{eval}]\!]_{U(\mathcal{M})}(x, y) = [\![\varphi_\ell]\!]_{\mathcal{M}}(y, x')$$
and note that this map is well defined.

Note that the listed axioms for T' force a definition of $[\![\mathsf{d}]\!]_{U(\mathcal{M})}$. It is clear that $U(\mathcal{M}) \models T'$ for every $M \in \mathrm{Mod}(T)$, and that $F(U(\mathcal{M})) = \mathcal{M}$ for every $\mathcal{M} \in \mathrm{Mod}(T)$. This shows that F is essentially surjective on objects. We show that $U(F(\mathcal{M}')) \cong \mathcal{M}'$ for every $\mathcal{M}' \in \mathrm{Mod}(T')$. To do this all we need to do is show that the interpretation of S' in $\mathcal{M}'$ is the actual union of the interpretations of S_{φ_i} in $\mathcal{M}'$.

Define $f : \bigcup_{i \le n} \mathsf{S}_{\varphi_i}(\mathcal{M}) \to \mathsf{S}'(\mathcal{M}')$ as follows. Let $x \in \bigcup_{i \le n} \mathsf{S}_{\varphi_i}(\mathcal{M})$, then for some $\ell \le n$, $x \in \mathsf{S}_{\varphi_\ell}(\mathcal{M})$. Define $f(x) = i_{\varphi_i}(x)$, and note that this is well defined. The map f is injective, because each i_i is an isometry. Also, for every $x \in \mathsf{S}'(\mathcal{M}')$, there is $y \in \bigcup_{i \le n} \mathsf{S}_{\varphi_i}(\mathcal{M})$ such that $[\![\mathsf{d}]\!]_{\mathcal{M}'}(f(y), x) = 0$, showing that the map f is surjective, and that $U(F(\mathcal{M}')) \cong \mathcal{M}'$. Therefore, F is full, and injective on objects.

To finish the proof we need to show that F is faithful, or equivalently that U is full. Let $h : U(\mathcal{M}) \to U(\mathcal{M}')$ be an elementary map, and consider $U(h^{\mathscr{L}})$. We show that $h = U(h^{\mathscr{L}})$. To this end, let $x \in \mathsf{S}'(U(\mathcal{M}))$. Then there is $\ell \le n$, and $x' \in \mathrm{dom}(\varphi_\ell)(\mathcal{M})$ such that $i_{\varphi_\ell}(\pi_{\varphi_\ell}(x')) = x$. Since h is elementary, $i_{\varphi_\ell}(\pi_{\varphi_\ell}(h(x'))) = h(x)$, and since $U(h^{\mathscr{L}})$ is elementary, $i_{\varphi_\ell}(\pi_{\varphi_\ell}(U(h^{\mathscr{L}})(x'))) = U(h^{\mathscr{L}})(x)$. However, since x' is an $\mathscr{L}'$-variable, $U(h^{\mathscr{L}})(x') = h(x')$, so we get $i_{\varphi_\ell}(\pi_{\varphi_\ell}(h(x'))) = U(h^{\mathscr{L}})(x) = h(x)$, proving that $h = U(h|\mathscr{L})$ and that U is full □

We now give the definition of the expansions $\mathscr{L}^{eq}$ and T^{eq} of $\mathscr{L}$ and T. Our definitions are by no means the most efficient.

3.7.5 Definition: Let $\mathscr{L}$ be a continuous language. We form the language $\mathscr{L}^{eq}$, by taking the smallest language $\mathscr{L}'$ with the following properties:

1. $\mathscr{L} \subseteq \mathscr{L}'$
2. If $\varphi \in \mathscr{L}'$, then $\mathscr{L}'_\varphi \subseteq \mathscr{L}'$
3. If S_i is a sort of $\mathscr{L}'$ for $0 \leq i < \kappa \leq \omega$, and v is a compact norm on $\mathfrak{R}$, then $\mathscr{L}'_{\prod^u \mathsf{S}_i} \subseteq \mathscr{L}'$
4. If $\varphi_1, ..., \varphi_n \in \mathscr{L}'$, then $\mathscr{L}'_{\bigcup \mathsf{S}_{\varphi_i}} \subseteq \mathscr{L}'$

If T is an $\mathscr{L}$-theory, then we form the theory T^{eq} to be the smallest theory T' with the following properties:

1. $T \subseteq T'$
2. If $\varphi \in T'$, then $T'_\varphi \subseteq T'$
3. If S_i is a sort of T' for $0 \leq i < \kappa \leq \omega$, and v is a compact norm on $\mathfrak{R}$, then $T'_{\prod^u \mathsf{S}_i} \subseteq T'$
4. If $\varphi_1, ..., \varphi_n \in T'$, then $T'_{\bigcup \mathsf{S}_{\varphi_i}} \subseteq T'$ ♣

Given the construction of $\mathscr{L}^{eq}$ and T^{eq}, we can inductively use Theorem 3.7.3 to get the following:

3.7.6 Theorem: *The forgetful functor* $F : \mathrm{Mod}(T^{eq}) \to \mathrm{Mod}(T)$ *is an equivalence of categories.*

Section 3.8

Conceptual completeness

In the previous section we defined an expansion of a theory which we called T^{eq}. In this section, we give a proof of a result of Bradd Hart's that the T^{eq} we defined above is in a sense the right one. We give a proof of the conceptual completeness result of Makkaï's (see [MR77]) in the context of $\mathfrak{R}$-valued languages. The proof we present here is based on the argument in [Har]. We adapted the terminology to fit our own.

We introduce here a piece of terminology which will not be used elsewhere in the thesis, but which makes the statement of Theorem 3.8.3 below much cleaner. Let $\mathscr{L}$ and $\mathscr{L}'$ be languages, T be an $\mathscr{L}$-theory, and T' be an $\mathscr{L}'$-theory. An interpretation of T' in T consists of the following data:

1. For every sort S' of $\mathscr{L}'$, a T^{eq}-definable set $[\![\mathsf{S}']\!]_T$.
2. For every function symbol f' of $\mathscr{L}'$ such that $\mathsf{f}' : \mathsf{S}'_1 \times \cdots \times \mathsf{S}'_n \to \mathsf{S}'$, a T^{eq}-definable function $[\![\mathsf{f}']\!]_T : [\![\mathsf{S}'_1]\!]_T \times \cdots \times [\![\mathsf{S}'_n]\!]_T \to [\![S']\!]_T$
3. For every relation symbol $\mathbb{R}'$ of $\mathscr{L}'$ such that $\mathrm{dom}(\mathbb{R}') = \mathsf{S}'_1 \times \cdots \times \mathsf{S}'_n$, an $\mathscr{L}^{eq}$-formula $[\![\mathbb{R}']\!] : [\![\mathsf{S}'_1]\!]_T \times \cdots \times [\![\mathsf{S}'_n]\!]_T$

4. For every constant symbol c' of $\mathscr{L}'$, an element $c \in [\![\text{type}(\mathsf{c}')]\!]_T$

with the property that for every formula $\varphi' \in \mathscr{L}'$,

$$T' \models \varphi' \text{ if and only if } T \models [\![\varphi']\!]_T .$$

If there is an interpretation of T' in T, then we say that T' is *interpretable* in T.

In this section, we let $\mathscr{L}$ be a continuous language, and let $\mathscr{L}'$ expand $\mathscr{L}$. Let T' be a complete $\mathscr{L}'$-theory, and let T be the $\mathscr{L}$-reduct of T'. A notion which is central to the proof of conceptual completeness is that of stable embeddedness.

3.8.1 Definition: Let $\mathscr{L}$, $\mathscr{L}'$, T and T' be as in the previous paragraph. Let $\mathcal{M}' \in \text{Mod}(T')$, and let $\mathcal{M}$ be its $\mathscr{L}$-reduct. We say that $\mathcal{M}$ is *stably embedded in* $\mathcal{M}'$ if and only if for every $\varepsilon > 0$, and every $\mathscr{L}'_{\mathcal{M}'}$ formula $\varphi(x)$ with $\text{dom}(\varphi)$ a sort of $\mathscr{L}$, there is an $\mathscr{L}_{\mathcal{M}}$-formula $\psi(x)$ such that $\text{dom}(\psi) = \text{dom}(\varphi)$, and

$$\forall x |\varphi(x) - \psi(x)| \leq \varepsilon.$$

3.8.2 Lemma: *Suppose the forgetful functor $F : \text{Mod}(T') \to \text{Mod}(T)$ is full and faithful. Then $F(\mathcal{M})$ is stably embedded in $\mathcal{M}$ for every $\mathcal{M} \in \text{Mod}(T')$.*

Proof: Fix a model $\mathcal{M} \in \text{Mod}(T')$. Let x and y be variables of type S, where S is a sort of $\mathscr{L}$. Let $c \in \mathsf{S}'(\mathcal{M})$, where S' is a sort of $\mathscr{L}'$, $n \in \mathbb{N}$, and let $\psi(x, z)$ have domain $\mathsf{S} \times \mathsf{S}'$. and let $\Sigma(x, y, c; n, \psi)$ be the following set of formulae in $\mathscr{L}'_{\mathcal{M}}$:

$$\text{diag}(\mathcal{M}) \cup \{|\varphi(x) - \varphi(y)| \leq 1/k : \varphi \in \mathscr{L}_{F(\mathcal{M})}, k \in \mathbb{N}\} \cup \{|\psi(x, c) - \psi(y, c)| \geq 1/n\}.$$

We claim that $\Sigma(x, y, c; n, \psi)$ is inconsistent for every n and every ψ.

Assume this for the moment, and let $\psi(x, c) \in \mathscr{L}'_{\mathcal{M}}$. Since c is a fixed constant, ψ induces a function $f : S_{\mathscr{L}}(\mathcal{M}) \to \mathfrak{R}$, where $S_{\mathscr{L}}(\mathcal{M})$ denotes the set of all complete $\mathscr{L}$-types over $\mathcal{M}$. The function f is defined by $f(\text{tp}(a/\mathcal{M})) = \psi(a, c)$. We need to show that this function is continuous. Let $\varepsilon > 0$, and let $\text{tp}(a/\mathcal{M}) \in f^{-1}((0, \varepsilon))$. We will find a small open set $U \subseteq S_{\mathscr{L}}(\mathcal{M})$ such that $p \in U$ and $f[U] \subseteq (0, \varepsilon)$. Since $\Sigma(x, y, c; n, \psi)$ is inconsistent, the set

$$\Sigma_p \overset{\text{def}}{:=} \text{diag}(\mathcal{M}) \cup \{|\varphi(x) - \varphi(y)| \leq 1/k : \varphi \in p, k \in \mathbb{N}\} \cup \{|\psi(x, c) - \psi(y, c)| \geq 1/n\}$$

is inconsistent as well. By compactness, there is a number δ and a finite subset $\Delta \subseteq p$ such that if $\bigwedge_{\varphi \in \Delta} |\varphi(x) - \varphi(y)| < \delta$, then $|\psi(x, c) - \psi(y, c)| < \varepsilon$. Let $\varepsilon' = \dfrac{\varepsilon - f(\text{tp}(a/\mathcal{M}))}{2}$. By the definition of the logic topology on $S_{\mathscr{L}}(\mathcal{M})$, the type

$$[\Delta < \delta] = \{q : \varphi < \delta' \text{ for some } \varphi \in \Delta \text{ and some } 0 < \delta' \leq \delta\}$$

is a basic open set in $S_{\mathscr{L}}(\mathcal{M})$. Let $\text{tp}(b/\mathcal{M}) \in [\Delta < \delta]$. Then $|f(\text{tp}(a/\mathcal{M})) - f(\text{tp}(b/\mathcal{M}))| < \varepsilon'$, and therefore $f(\text{tp}(b/\mathcal{M})) < \varepsilon$, as is required to show that $f[[\Delta < \delta]] \subseteq (0, \varepsilon)$. This proves that f is continuous, so by Theorem 3.3.8, it is of the form f_φ for some $\mathscr{L}$-formula φ.

Suppose $\Sigma(x, y, c; n, \psi)$ is consistent, and let $\mathcal{N} \models \Sigma(a, b, c; n, \psi)$, with $a, b \in \mathsf{S}(\mathcal{N})$. Note that there is an elementary embedding $g : \mathcal{M} \to \mathcal{N}$. Since $\Sigma(a, b, c; n, \psi)$ implies that $a \equiv_{F(\mathcal{M})} b$, there is an ultrafilter pair (I, U) and an embedding $h : F(\mathcal{N}) \to F(\mathcal{N}^U)$

such that $h(a) = \partial_{F(\mathcal{N})}(b)$ and $h|F(\mathcal{M}) = \partial_{F(\mathcal{M})}$. By fullness, there is an elementary map $h' : \mathcal{N} \to \mathcal{N}^U$ such that $F(h') = h$, and $h'|\mathcal{M} = \partial_{\mathcal{M}}$. This means

$$\begin{aligned} \psi(a,c) &= \psi(h'(a), h'(c)) \\ &= \psi(h'(a), \partial_{\mathcal{N}}(c)) \\ &= \psi(\partial_{\mathcal{N}}(b), \partial_{\mathcal{N}}(c)) \\ &= \psi(b,c) \end{aligned}$$

which is impossible, since the fact that a, b, c realizes Σ implies that $|\psi(a,c) - \psi(b,c)| \geq 1/n$.□

3.8.3 Theorem: *Suppose that the forgetful functor $F : \mathrm{Mod}(T') \to \mathrm{Mod}(T)$ is an equivalence of categories. Then T' is interpretable in T^{eq}.*

Proof: For this proof we fix a sort S' of $\mathscr{L}'$. We need to find a sort S'' of $\mathscr{L}^{\mathrm{eq}}$ and a definable embedding $f : \mathsf{S}' \to \mathsf{S}''$. Fix a saturated model $\mathcal{M} \in \mathrm{Mod}(T')$, and let $c \in \mathsf{S}'(\mathcal{M})$. We first work in the expansion T'_c and find, for every $\mathscr{L}'$-formula φ with domain $\mathsf{S} \times \mathsf{S}'$, $\mathsf{S} \in \mathbf{Sort}_{\mathscr{L}}$, a sort of $\mathscr{L}^{\mathrm{eq}}$ which captures the $\mathscr{L}'_c$-definable predicate $[x \mapsto \varphi(x,c)]$. Consider the set:

$$\Sigma_n \overset{\mathrm{def}}{:=} \{\forall \bar{y}[|\varphi(\bar{x},c) - \psi(\bar{x},\bar{y})| \geq \frac{1}{2^n} : \psi \in \mathscr{L}, \mathrm{dom}(\psi) = \mathsf{S} \times \mathsf{S}_1, \mathsf{S}_1 \text{ a sort of } \mathscr{L}\}$$

We claim that Σ_n is inconsistent for every n. Indeed, if N is such that Σ_N is consistent, then let $\mathcal{M} \models \Sigma_N$. Then, in $\mathcal{M}$, for every ψ with domain $\mathsf{S} \times \mathsf{S}_1$, we have

$$|\,[\![\varphi_c]\!]_{\mathcal{M}}(\bar{x}) - [\![\psi]\!](\bar{x},\bar{y})| \geq 1/2^n$$

for every $\bar{y} \in \mathsf{S}_1(\mathcal{M})$, contradicting the fact that $F(\mathcal{M})$ is stably embedded in $\mathcal{M}$.

For every $n \in \mathbb{N}$, let $\Delta_n(\bar{x},\bar{y}) \subseteq \mathscr{L}$ be finite such that

$$\left\{\forall \bar{y}[|\varphi_c(\bar{x}) - \psi(\bar{x},\bar{y})| \geq \frac{1}{2^n} : \psi \in \Delta_n\right\}$$

witnesses the inconsistency of Σ_n. Let $u : \mathfrak{R}^\omega \to \mathfrak{R}$ be a compact norm, and let $\mathsf{S}_n = \bigcup_{\varphi \in \Delta_n} \mathsf{S}_\varphi$ for $n < \omega$. Let $\mathsf{S}^\varphi = \prod_{n \in \mathbb{N}} \mathsf{S}_n$, and note that S^φ is a sort in $\mathscr{L}^{\mathrm{eq}}$.

Now consider the set

$$\Sigma_n(c,c') \overset{\mathrm{def}}{:=} \{\forall x|\varphi(x,c) - \varphi(x,c')| : \mathrm{dom}(\varphi) = \mathsf{S} \times \mathsf{S}', \mathsf{S} \in \mathbf{Sort}_{\mathscr{L}}\} \cup \{\mathsf{d}_{\mathsf{S}'}(c,c') \geq 1/n\}$$

First we argue that this set is inconsistent for every n. This is a consequence of the faithfulness of F, as the consistency of $\Sigma_n(c,c')$ implies that $c \equiv_{F(\mathcal{M})} c'$ and that $c \neq c'$. Therefore, there is an ultrafilter pair (I,U) and an elementary map $h : \mathcal{M} \to \mathcal{M}^U$ such that $h(c) = \partial_{\mathcal{M}}(c')$. The map h has the property that $F(h) = \partial_{F(\mathcal{M})}$, and yet $h \neq \partial_{\mathcal{M}}$, contradicting the faithfulness of F.

From the inconsistency of $\Sigma_n(c,c')$, we get that there is a countable set $\{\varphi_i(x,c) : i < \omega\}$ such that if

$$[x \mapsto \varphi_i(x,c)] = [x \mapsto \varphi_i(x,c')]$$

for every $i < \omega$, then $c = c'$. Let $\mathsf{S}'' = \prod_{i<\omega}^u \mathsf{S}^{\varphi_i}$, where u is any compact norm.The sequence $f(c) = ([x \mapsto \varphi_i(x,c)] : i < \omega)$ is an element of S'', and $f(c) = f(c')$ implies that $c = c'$ by a previous argument. Therefore, f is a definable map $\mathsf{S}' \to \mathsf{S}''$.

We now use f to define the interpretation $[\![-]\!]_T$. In the following, if S is a sort of $\mathscr{L}'$, then we denote by f_{S} the function f defined in the previous paragraph.

1. If S is a sort of $\mathscr{L}'$, then $[\![\mathsf{S}]\!]_T = \mathrm{im}(f_{\mathsf{S}})$. This is a definable set by virtue of f_{S} being a definable map whose domain is a full sort. Note also that the formula defining $\mathrm{im}(f_{\mathsf{S}})$ has domain a sort of $\mathscr{L}^{\mathrm{eq}}$, is an $\mathscr{L}^{\mathrm{eq}}$-formula by stable embeddedness.

2. If $\mathsf{R}(x_1, ..., x_n)$ is a relation symbol with domain $\mathsf{S}_1 \times \cdots \times \mathsf{S}_n$, then $[\![\mathsf{R}]\!]_T\,(x_1, ..., x_n) = \mathsf{R}(f_{\mathsf{S}_1}^{-1}(x_1), ..., f_{\mathsf{S}_n}^{-1}(x_n))$. Again we note that $[\![\mathsf{R}]\!]_T\,(x_1, ..., x_n)$ is a formula with domain a sort of $\mathscr{L}^{\mathrm{eq}}$, and therefore $\mathsf{R}(f^{-1}(x_1), ..., f^{-1}(x_n))$ can be obtained as an $\mathscr{L}^{\mathrm{eq}}$-formula by stable embeddedness.

3. If $\mathsf{f}(x_1, ..., x_n) : \mathsf{S}_1 \times \cdots \times \mathsf{S}_n \to \mathsf{S}$ is a function symbol, then

$$[\![\mathsf{f}]\!]_T = f_{\mathsf{S}}(\mathsf{f}(f_{\mathsf{S}_1}^{-1}(x_1), ..., f_{\mathsf{S}_n}^{-1}(x_n))).$$

   The same comment applies.

4. If c is a constant symbol, then $[\![\mathsf{c}]\!]_T = f(\mathsf{c})$ □

Chapter 4
Strong conceptual completeness

Section 4.1
Prelude

The study of logical systems is usually divided into two fields: syntax and semantics. On the syntax side, we have formulae, theories and proofs. On the semantic side, we have structures and models.

A lot of syntactic information can be deduced by looking at classes of structures. For example, Theorem 3.4.10, which establishes when a class of structures can be characterized by a theory, or Theorem 3.4.11, which can be used to determine whether a class of structures can be isolated using a single sentence. In a sense, strong conceptual completeness takes Theorem 3.4.10 one step further. Whereas Theorem 3.4.10 requires prior knowledge of a language $\mathscr{L}$, strong conceptual completeness only requires a category $\mathbb{C}$ with appropriate structure, and produces both $\mathscr{L}$ and a theory T such that the objects of $\mathbb{C}$ are $\mathscr{L}$-structures, and $\mathbb{C} \cong \mathrm{Mod}(T)$. The result we prove here does not quite have that level of generality. We do assume prior knowledge of $\mathscr{L}$ and T, but our proof could be adapted to produce $\mathscr{L}$ without too much trouble.

If $\mathscr{L}$ is a first-order language, and $\varphi \in \mathscr{L}$ is a formula, then we can associate to φ a functor $\mathfrak{ev}_\varphi : \mathrm{Mod}(T) \to \mathbb{S}\mathrm{et}$ which is defined by

$$\mathcal{M} \mapsto \{x \in \mathcal{M} : \mathcal{M} \models \varphi(x)\}$$

on objects. Using the notation of chapter 3 for elementary maps, the action of $\mathfrak{ev}_\varphi$ on elementary maps is given by $\mathfrak{ev}_\varphi(h) = \mathrm{dom}(\varphi)(\mathcal{M})(h)$. A strong conceptual completeness result is an answer to the following question: what properties must a functor $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{S}\mathrm{et}$ have in order to be naturally isomorphic to a functor of the form $\mathfrak{ev}_\varphi$ for $\varphi \in \mathscr{L}$. The property isolated in [Mak88] is that $\mathfrak{f}$ should be an *ultrafunctor.*

The category $\mathbb{S}\mathrm{et}$ has, for every ultrafilter pair (I, U), a functor $\prod_U : \mathbb{S}\mathrm{et}^I \to \mathbb{S}\mathrm{et}$ which assigns to every I-indexed sequence of sets its ultraproduct, and to every sequence of maps the map $\langle x_i \rangle_U \mapsto \langle f_i(x_i) \rangle_U$. By Lós' theorem, the same is true of the category $\mathrm{Mod}(T)$. A *pre-ultrafunctor* $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{S}\mathrm{et}$ is a functor which commutes with $\prod_U$ for every ultrafilter pair (I, U). An easy argument using Lós' theorem shows that $\mathfrak{ev}_\varphi$ is a pre-ultrafunctor for every $\varphi \in \mathscr{L}$.

The conditions required for a pre-ultrafunctor to be called an *ultrafunctor* are more technical. Ultrafunctors must satisfy the additional property of preserving canonical relationships between ultraproducts. One example of such relationship is the canonical embedding $\partial_\mathcal{M}$ of a model $\mathcal{M}$ into its ultrapower $\mathcal{M}^U$. Another example arises in the following situation: let I and J be sets, and U be an ultrafilter on J. Suppose $f : J \to I$ is a function.

Recall from chapter 2 that there is an ultrafilter $f[U]$ on I induced by f. If $(\mathcal{M}_i : I \in I)$ is an I-indexed sequence of models of a theory T, then we get the J-indexed sequence $(\mathcal{M}_{f(j)} : j \in J)$. Let

$$\partial : \prod_{f[U]} \mathcal{M}_i \to \prod_{U} \mathcal{M}_{f(j))}$$

be the map

$$\langle x_i : i \in f[U] \rangle_U \mapsto \langle x_{f(j)} : j \in U \rangle_U.$$

We will see below that this map is well defined, and a elementary embedding of the model $\prod_{f[U]} \mathcal{M}_i$ into the model $\prod_U \mathcal{M}_{f(j))}$. The map ∂ defined above only depends on the sets I and J, ultrafilter U and the function $f : J \to I$, not on the models $\mathcal{M}_i$. It is, in this sense, "canonical". In the later sections, the tuple (I, J, U, f) will be referred to as an *ultragraph.*

If we denote by Δ an I-indexed sequence of models $\mathcal{M}_i$, and $\mathfrak{f}$ is a pre-ultrafunctor, then we get the sequence $\mathfrak{f} \circ \Delta = (\mathfrak{f}(\mathcal{M}_i) : i \in I)$. We have the canonical map $\partial_\Delta : \mathrm{Mod}(T)^I \to \mathrm{Mod}(T)$, and a similarly defined map $\partial_{\mathfrak{f}\circ\Delta} : \mathrm{Mod}(T)^I \to \mathbb{S}\mathrm{et}$. Saying that $\mathfrak{f}$ preserves the canonical embedding ∂ means that $\mathfrak{f} \circ \partial_\Delta = \partial_{\mathfrak{f}\circ\Delta}$. Ultrafunctors will be those functors which preserve *all* such canonical embeddings.

We give two examples of functors $\mathrm{Mod}(T) \to \mathbb{S}\mathrm{et}$ which are *not* ultrafunctors. These examples are set in classical first-order logic, and represent sets that are not in general definable.

4.1.1 Example (A $\bigwedge$-definable set): Let p be a complete type whose set of realization is not definable. Consider the functor $\mathfrak{f}$ defined by $\mathfrak{f}(\mathcal{M}) = \{x \in \mathcal{M} : \mathcal{M} \models p(x)\}$, and $\mathfrak{f}(h) = h|\mathfrak{f}(\mathcal{M})$. Then $\mathfrak{f}$ does not commute with ultraproducts, as the following argument shows: since the set of realizations of p is not definable, for every finite subset $\Delta \subseteq p$, there is a model $\mathcal{M}_\Delta$ which realizes Δ but omits p. Let $I = \{\Delta \subseteq p : \Delta \text{ finite}\}$, and let U be any ultrafilter containing all the sets of the form $\hat{j} = \{i \in I : j \in i\}$. Let $\mathcal{M} = \prod_U \mathcal{M}_\Delta$. Then we see that $\prod_U \mathfrak{f}(\mathcal{M}_\Delta) = \emptyset$, since none of the $\mathcal{M}_\Delta$'s realize p. However, since $\mathcal{M}_\Delta \models \Delta$, for every Δ, there is $a_\Delta \in \mathcal{M}_\Delta$ such that $\mathcal{M}_\Delta \models \Delta(a_\Delta)$. The sequence $\langle a_\Delta \rangle_U$ realizes p in $\mathcal{M}$, so $\mathfrak{f}(\mathcal{M}) \neq \emptyset$. ◊

4.1.2 Example (The algebraic closure): T is the theory of algebraically closed fields of characteristic 0. We can "define" the algebraic closure of $\emptyset$ by saying that it is the set of all algebraic elements, but $\mathrm{acl}(\emptyset)$ is not a definable set. Let $\mathbb{C}$ be the complex numbers, and let P be the set of all prime numbers > 2. For every $p \in P$, let μ_p be a primitive root of unity. Note that formulae satisfied by μ_p are satisfied by at least $p-2$ other elements. Let U be an ultrafilter on P containing all the cofinite sets, and let $F = \prod_U \mathbb{C}$. Let $\mu = \langle \mu_p : p \in P \rangle_U$. It is clear that $\mu \in \prod_U \mathrm{acl}(\emptyset)_{\mathbb{C}}$. On the other hand, if $\varphi(x)$ is any formula, and $F \models \varphi(\mu)$, then for almost every $p \in P$, $\mathbb{C} \models \varphi(\mu_p)$, so φ has at least $p-2$ realizations for cofinitely many p's. Therefore it has infinitely many realizations in F, so $\mu \notin \mathrm{acl}(\emptyset)_F$. We have shown that $(\mathrm{acl}(\emptyset)_{\mathbb{C}})^U \neq \mathrm{acl}(\emptyset)_{\mathbb{C}^U}$. ◊

When a formula φ of an $\mathfrak{R}$-valued language $\mathscr{L}$ is evaluated in a structure $\mathcal{X}$, it does not give rise to a subset of $\mathcal{X}$, but rather to a continuous function $[\![\varphi]\!]_{\mathcal{X}} : \mathrm{dom}(\varphi)(\mathcal{X}) \to \mathfrak{R}$. This indicates that the category $\mathbb{S}\mathrm{et}$ is *not* the best choice of a target category for ultrafunctors. Our first order of business is to find a suitable replacement for $\mathbb{S}\mathrm{et}$. This replacement is defined in section 4.2, and comes in the form of the category $\mathbb{MET}_{\mathfrak{R}}$ consisting of *pairs* (X, φ), where X is a generalized metric space, and φ is a uniformly continuous function $X \to \mathfrak{R}$. Section 4.2 also highlights some structural properties of $\mathbb{MET}_{\mathfrak{R}}$ in relation to the category $\mathbb{MET}$ whose objects are generalized metric spaces, and makes sense of the functors $\prod_U$ in this context.

Armed with proper definitions of ultraproduct functors, we proceed in section 4.2 with the continuous analogue of the definition of ultrafunctor given in [Mak88]. Ultrafunctors are key to the strong conceptual completeness result. Because we have a proper definition of ultraproduct functors, the definition of ultrafunctor we give is quite close to the one given in [Mak88]: ultrafunctors are functors $\text{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ (note the range change) which commute with ultraproduct functors, and preserve canonically defined embeddings of ultraproducts into one another.

A key observation in 4.2 is that if $\mathfrak{f}$ is an ultrafunctor, then $\text{dom}(\mathfrak{f})$ should be a definable set in the sense of continuous logic. In the special case where $\mathfrak{f} = \mathfrak{ev}_\varphi$ is an ultrafunctor coming from a formula, then its domain is a sort of $\mathscr{L}$, and as such is a definable set. If $\eta : \mathfrak{f} \to \mathfrak{ev}_\varphi$ is an injective natural transformation, then for every $\mathcal{M} \in \text{Mod}(T)$, $\text{im}(\eta_\mathcal{M}) \subseteq \text{dom}(\mathfrak{ev}_\varphi)$ is a *definable* subset of a sort of $\mathscr{L}$, i.e. the function $x \mapsto d(x, \text{im}(\eta_\mathcal{M}))$, which maps $\text{dom}(\varphi)(\mathcal{M})$ to $\mathfrak{R}$ is equal to a formula ψ.

The technical part of the proof is showing that in fact, the domain of an ultrafunctor *is* a definable set. In order to do this, we must find a sort S of $\mathscr{L}$, and an injective natural transformation $\eta : \mathfrak{f} \to \mathfrak{ev}_\mathsf{S}$.

For the rest of this chapter, we let $\mathscr{L}$ be a metric language as defined in chapter 3, and T is a metric $\mathscr{L}$-theory.

Section 4.2

Ultrafunctors

In this section we introduce the concept of ultrafunctor in the context of $\mathfrak{R}$-valued languages. This is the analogue of the concept described in [Mak87]. We begin by giving precise definitions for all the categories involved in this chapter. First, we recall the definition of $\text{Mod}(T)$:

Objects: Models of T

Morphisms: Elementary maps between models.

Composition: Function composition.

A formula $\varphi(x) \in T$ in the single free variable x (say), when interpreted in a model $\mathcal{M}$, gives rise to a function $[\![\varphi]\!]_\mathcal{M} : \text{dom}(\varphi)(\mathcal{M}) \to \mathfrak{R}$. This leads us to consider not the category of sets as the space on which models are based, but rather a category of such functions which we denote by $\mathbb{MET}_\mathfrak{R}$, which we define below. Recall the category $\mathbb{MET}$ from chapter 2, which is the category of generalized metric spaces and uniformly continuous maps between them.

Objects: Pairs (X, φ), where $X \in \mathbb{MET}$, and φ is a uniformly continuous map $X \to \mathfrak{R}$. The first coordinate of (X, φ) will be referred to as the *domain* of (X, φ), and denoted $\text{dom}(\varphi)$. We will use φ to denote the pair $(\text{dom}(\varphi), \varphi)$.

Morphisms: Triplets (f, φ, ψ) such that $f : \text{dom}(\varphi) \to \text{dom}(\psi)$ is an isometry, and $\psi(f(x)) = \varphi(x)$ for every $x \in \text{dom}(\varphi)$. Likewise, we will use f for both the function $f : \text{dom}(\varphi) \to \text{dom}(\psi)$, and the triple (f, φ, ψ). The metric spaces $\text{dom}(\varphi)$ and $\text{dom}(\psi)$ will be denoted $\text{dom}(f)$ and $\text{codom}(f)$ respectively.

Composition: $(g,\psi',\psi)\circ(f,\varphi,\psi')=(g\circ f,\varphi,\psi)$.

4.2.1 Definition: We let dom denote the forgetful functor dom : $\mathbb{MET}_{\mathfrak{R}}\to\mathbb{MET}$ sending a an object $\varphi\in\mathbb{MET}_{\mathfrak{R}}$ to dom(φ), and a morphism $h:\varphi\to\psi$ to the underlying isometry $h:\mathrm{dom}(\varphi)\to\mathrm{dom}(\psi)$. ♣

A key feature of $\mathbb{MET}_{\mathfrak{R}}$ is that, like $\mathbb{MET}$ and Mod(T), it has ultraproduct functors $\prod_U$ corresponding to ultrafilter pairs (I,U). In order to define the functors $\prod_U$ on $\mathbb{MET}_{\mathfrak{R}}$, we first define the domain category $\mathbb{MET}^I_{\mathfrak{R}}$. For technical reasons, we cannot let $\mathbb{MET}^I_{\mathfrak{R}}$ be the full category of I-indexed sequences of elements of $\mathbb{MET}_{\mathfrak{R}}$. $\mathbb{MET}^I_{\mathfrak{R}}$ is the following category:

Objects: Equicontinuous sequences $(\varphi_i:i\in I)$ of elements of $\mathbb{MET}_{\mathfrak{R}}$. That is to say, sequences $(\varphi_i:i\in I)$ of elements of $\mathbb{MET}_{\mathfrak{R}}$ such that for every $\varepsilon>0$, there is $\delta>0$ such that for every i, and every $x_i,y_i\in\mathrm{dom}(\varphi_i)$, if $d(x_i,y_i)<\delta$, then $|\varphi_i(x_i)-\varphi(y_i)|<\varepsilon$

Morphisms: Sequences $(f_i:i\in I)$ of morphisms of $\mathbb{MET}_{\mathfrak{R}}$. Note that by the definition of morphism of $\mathbb{MET}_{\mathfrak{R}}$, the sequences $(f_i:i\in I)$ of morphisms of $\mathbb{MET}_{\mathfrak{R}}$ are equicontinuous.

Composition: Coordinatewise function composition

4.2.2 Proposition: *For every ultrafilter pair (I,U), there is a functor $\prod_U:(\mathbb{MET}_{\mathfrak{R}})^I\to\mathbb{MET}_{\mathfrak{R}}$. The functor $\prod_U$ assigns to every sequence $(\varphi_i:i\in I)$ the element φ given by*

1. $\mathrm{dom}(\varphi)=\prod_U\mathrm{dom}(\varphi_i)$
2. $\varphi=\lim\limits_{i\to U}\varphi_i$

and to every sequence $(f_i:i\in I)$ the function $(\prod_U f_i)(\langle x_i\rangle_U)=\langle f_i(x_i)\rangle_U$.

The proof of proposition 4.2.2 is an easy application of the definition of ultraproducts of metric spaces and ultralimits of continuous functions. Note, however, that the action on maps is well-defined because all maps in $\mathbb{MET}_{\mathfrak{R}}$ are isometries, and therefore all sequences of maps are equicontinuous. This is in contrast with the similar definition of $\prod_U$ for $\mathbb{MET}$, which requires a restriction on the maps we allow in $\mathbb{MET}^I$ in order to get a functor. Proposition 4.2.2, properly modified, also holds when $\mathbb{MET}_{\mathfrak{R}}$ is replaced by $\mathbb{MET}_{iso}$, the category of generalized metric spaces and isometries, and by Mod(T). We can now define ultrafunctors. Informally, ultrafunctors are functors which commute with the operation of taking ultraproducts, and commute with certain "canonical maps" between ultraproducts.

4.2.3 Definition: An *ultragraph* is a tuple of the form (I,J,U,f), where I and J are set, U an ultrafilter on J, and $f:J\to I$ is a function. ♣

We now describe an important special case of this definition which will be used later. Recall from chapter 2 that a U-selector is a function $f:U\to I$ such that $f(P)\in P$ for every $P\in U$. Suppose (I,U) is an ultrafilter pair, and $f:U\to I$ is a U-selector. By proposition 2.3.5, the data (I,U,f) is enough to specify an ultrafilter W on U such that $f[W]=U$. This ultrafilter will be denoted $f^{-1}[U]$, and thus we get the ultragraph $(I,U,f^{-1}[U],f)$, which will be denoted by (I,U,f).

4.2.4 Definition: Let $\mathbb{C}$ be the class of objects of either MET, $\text{MET}_{\mathfrak{R}}$ or Mod(T). Let $G = (I, J, U, f)$ be an ultragraph. A *G-ultradiagram on* $\mathbb{C}$ is an assignment $\Delta : I \to \mathbb{C}$. Given any ultradiagram Δ, we let

$$\partial_\Delta : \prod_{f[U]} \Delta(i) \to \prod_U \Delta(f(j))$$

be the map

$$\langle x_i : i \in U \rangle_U \mapsto \langle x_{f(j)} : j \in J \rangle_U.$$

The assignment $\Delta \mapsto \partial_\Delta$ will be denoted ∂, and will be called a *G-ultramorphism on* $\mathbb{C}$. ♣

This definition makes sense because all three of MET, $\text{MET}_{\mathfrak{R}}$ or Mod(T) have definitions for all the functors $\prod_U$. Let us show that ∂_Δ is well defined. We work in some unspecified sort of $\mathscr{L}^{\text{eq}}$, and write d for the metric on that sort. Suppose

$$\langle x_i : i \in U \rangle_U = \langle y_i : i \in U \rangle_U.$$

Then by definition, this means that for every $\varepsilon > 0$,

$$S = \{i : d(x_i, y_i) < \varepsilon\} \in f[U].$$

Consider

$$\langle x_{f(j)} : j \in J \rangle_U$$

and

$$\langle y_{f(j)} : j \in J \rangle_U$$

and let $\varepsilon > 0$. Since $S \in f[U]$,

$$f^{-1}[S] = \{j \in J : f(j) \in S\} \in U$$

and for every $j \in f^{-1}[S]$, we have $d(x_{f(j)}, y_{f(j)}) < \varepsilon$. Therefore, the set $\{j \in J : d(x_{f(j)}, y_{f(j)}) < \varepsilon\}$ is in U, showing that ∂_Δ is well defined.

Associated to any ultragraph G, there are two functors $[k], [\ell] : \mathbb{C}^I \to \mathbb{C}$ defined by

1. $[k](\Delta) = \prod_{f[U]} \Delta(i)$
2. $[\ell](\Delta) = \prod_U \Delta(f(j))$

and for every Δ, $\partial(\Delta) : [k](\Delta) \to [\ell](\Delta)$ is a map of $\mathbb{C}$. We make the observation that the definitions of ∂, $[k]$ and $[\ell]$ do not really depend on $\mathbb{C}$, but rather on the ultragraph G used to define them. We pause here to bring the following very important observation to the attention of the reader. In the special case where $I = \{*\}$ is a singleton, and (J, U) is an ultrafilter pair, the function $f = [x \mapsto *]$ induces the trivial ultrafilter $\{\{*\}\}$ on I. If $G = (\{*\}, J, U, [x \mapsto *])$, then a G-ultradiagram Δ is determined by the choice of a single model $\mathcal{M} \in \text{Mod}(T)$. Note that $[k](\Delta) = \mathcal{M}$, and $[\ell](\Delta) = \mathcal{M}^U$. The instance of ∂ at Δ is the diagonal embedding $\partial_{\mathcal{M}} : \mathcal{M} \to \mathcal{M}^U$. The reader can thus see that a triple of the form $(\partial, [k], [\ell])$ is a generalization of the canonical embedding of a model into its ultrapowers. In future sections, when we are considering ultragraphs of the form $(\{*\}, J, U, [x \mapsto *])$, we will identify the G-ultradiagram Δ with the model $\mathcal{M} = \Delta(*)$, and write $\partial_{\mathcal{M}}$ for $\partial_{\mathcal{M}}$ for ∂_Δ.

4.2.5 Definition: The triple $(\partial, [k], [\ell])$ corresponding to G will be denoted $(\partial_G, [k]_G, [\ell]_G)$, and its "value" in $\mathbb{C}$ will be denoted $([\![\partial_G]\!]_{\mathbb{C}}, [\![[k]_G]\!]_{\mathbb{C}}, [\![[\ell]_G]\!]_{\mathbb{C}})$. ♣

Since it will generally be clear which ultragraph and ultradiagram we are using, and in which category we are working, we often write $(\partial, [k], [\ell])$ instead of $([\![\partial_G]\!]_{\mathbb{C}}, [\![[k]_G]\!]_{\mathbb{C}}, [\![[\ell]_G]\!]_{\mathbb{C}})$.

4.2.6 Proposition: *The assignment ∂ defines a natural transformation $[k] \to [\ell]$, where $\mathbb{C}^I$ is given the obvious categorical structure.*

We now define the concept of ultrafunctor in the context of $\mathfrak{R}$-valued logic. This definition parallels the one given in [Mak87]. Let $\mathbb{C}$ and $\mathbb{D}$ be any two of MET, $\mathrm{MET}_{\mathfrak{R}}$ and Mod(T).

4.2.7 Definition: A *pre-ultrafunctor* is a functor $\mathfrak{f} : \mathbb{C} \to \mathbb{D}$ together with, for every ultrafilter U, a natural isomorphism $\iota^{\mathfrak{f},U} : \mathfrak{f} \circ \prod_U \to \prod_U \circ \mathfrak{f}^I$. Here the word "natural" is to be taken in the categorical sense of a natural transformation. A *ultratransformation* between two pre-ultrafunctors $\mathfrak{f}$ and $\mathfrak{g}$ is a natural transformation $\eta : \mathfrak{f} \to \mathfrak{g}$ such that the diagram

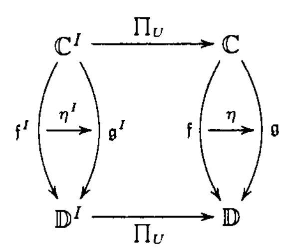

is commutative modulo the natural isomorphisms for every ultrafilter pair (I, U). ♣

In order not to clutter the notation in computations, the natural isomorphisms will not be referred to directly in the subsequent discussion. They will rather be denoted by $\cong$ when labelling arrows, and will generally be treated as identities.

4.2.8 Definition: A pre-ultrafunctor is called an *ultrafunctor* if in addition, it satisfies the following condition:

1. For every ultragraph G, and every ultradiagram $(\partial_G, [k]_G, [\ell]_G)$, the diagram

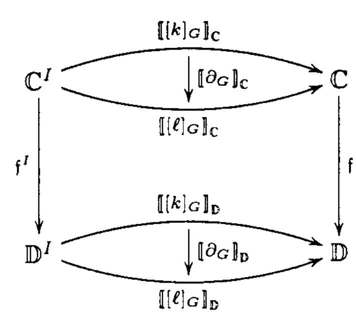

(†)

is commutative.

If $\mathfrak{f}$ is an equivalence of categories, and an ultrafunctor, and the functor $\mathfrak{f}^{-1}$ is also an ultrafunctor, then we call it an *ultra-equivalence of categories.* ♣

Behind this abstract definition hides something a bit simpler: if $\mathfrak{f} : \mathbb{C} \to \mathbb{D}$ is an ultrafunctor, $G = (I, J, U, f)$ an ultragraph and Δ an G-ultradiagram, then, in a nutshell $\mathfrak{f}(\partial_\Delta) = \partial_{\mathfrak{f}\circ\Delta}$ for every Δ. This can be most easily seen in the case where ∂ is the canonical embedding of structures into their ultrapowers, and can be read in the following diagram, which represents an instance of (†) at Δ:

$$\begin{array}{ccc} \mathfrak{f}\left(\prod_{f[U]} \Delta(i)\right) & \xrightarrow{\mathfrak{f}(\partial_\Delta)} & \mathfrak{f}(\prod_U \Delta(j)) \\ \cong\downarrow & & \downarrow\cong \\ \prod_{f[U]} \mathfrak{f}(\Delta(i)) & \xrightarrow[\partial_{\mathfrak{f}\circ\Delta}]{} & \prod_U \mathfrak{f}(\Delta(j)) \end{array} \qquad (\dagger\dagger)$$

We note that if (††) commutes for every Δ, the the diagram (†) is commutative.

4.2.9 Definition: We will denote by $\mathbb{ULT}_{\mathbb{M}}(T)$ the category defined as follows:

Objects: Ultrafunctors $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$

Morphisms: Ultra-transformations

Composition: Composition of natural transformations ♣

4.2.10 Proposition: *Let $\mathbb{C}$ and $\mathbb{D}$ be as above, and let $\mathfrak{f}, \mathfrak{g} : \mathbb{C} \to \mathbb{D}$ be ultrafunctors. Let $\nu : \mathfrak{g} \to \mathfrak{f}$ be a monic ultratransformation. Then $\mathfrak{g}$ is an ultrafunctor if and only if $\mathfrak{g} \prod_U \cong \prod_U \mathfrak{g}^I$ for every U.*

Proof: The fact that $\mathfrak{g} \prod_U \cong \prod_U \mathfrak{g}^I$ is part of the definition of $\mathfrak{g}$ being an ultrafunctor. Therefore, we only need to show the converse. What we need to do is to show that diagram (†) commutes when $\mathfrak{f}$ is replaced by $\mathfrak{g}$. Let (I, U, f) be an ultragraph, and let Δ be any ultradiagram. Consider the following cube in which, for simplicity, $V = f[U]$, and all the $\cong$ symbols represent the natural isomorphisms.

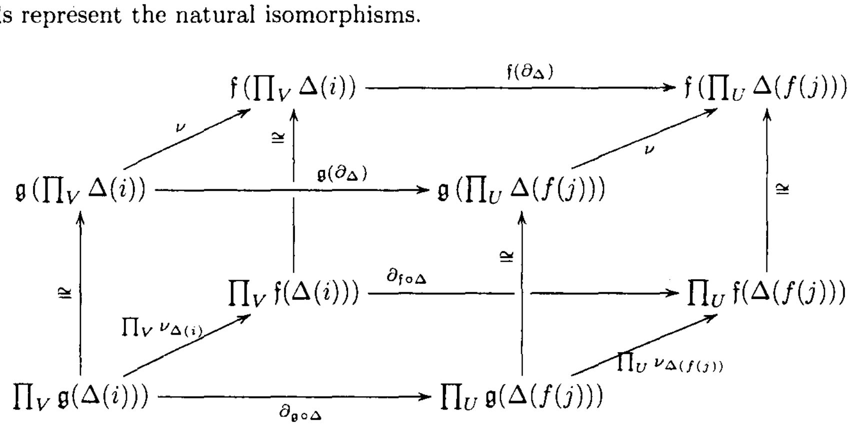

We investigate the commutativity of all faces of the cube.

1. The top face commutes by naturality of ν
2. The back face commutes because $\mathfrak{f}$ is an ultrafunctor
3. The bottom face commutes by definition of ∂.
4. The left and right faces commute, because ν is an ultratransformation

We now argue that the front face commutes as well, which is what we really want. To simplify the notation, we treat the vertical arrows as identities. Let $x \in \prod_U \mathfrak{g}(\Delta(i))$. Then

$$\nu(\mathfrak{g}(\partial_\Delta)(x)) = \mathfrak{f}(\partial_\Delta)(\nu(x))$$

because the top face commutes, so that

$$\mathfrak{g}(\partial_\Delta)(x) = \nu^{-1}(\mathfrak{f}(\partial_\Delta)(\nu(x))).$$

Note that ν^{-1} is well defined because ν is a monomorphism. Write $\mu = \prod_V \nu_{\Delta(i))}$ and $\eta = \prod_U \nu_{\Delta(f(j))}$. From the commutativity of the bottom face,

$$\mu(\partial_{\mathfrak{g}\circ\Delta}(x)) = \partial_{\mathfrak{f}\circ\Delta}(\eta(x))$$

so that

$$\partial_{\mathfrak{g}\circ\Delta}(x) = \mu^{-1}(\partial_{\mathfrak{f}\circ\Delta}(\eta(x)))$$

and again μ^{-1} is well defined because μ is a monomorphism. Now $\partial_{\mathfrak{f}\circ\Delta} = \mathfrak{f}(\partial_\Delta)$ from the back face, $\nu = \mu$ from the left face, and $\eta = \nu$ from the right face. The conclusion follows.

Since the fact that diagram (††) commutes for every Δ is equivalent to diagram (†) being commutative, we conclude that $\mathfrak{g}$ is an ultrafunctor. □

Proposition 4.2.10 will be the main tool in showing that functors are ultrafunctors. However, in order to use it, we must prove the existence of basic ultrafunctors on Mod(T). Namely, we must show that the sorts of $\mathscr{L}^{\text{eq}}$ define ultrafunctors. This is implicit in the statement that T^{eq} is a conservative expansion of T, and is made explicit in Lemma 4.2.11 and Theorem 4.2.12 below.

4.2.11 Lemma: $1 \leq \kappa \leq \omega$, $(\mathfrak{f}_i : i \in \kappa)$ *is a sequence of ultrafunctors* Mod(T) $\to$ MET, *and* $v : \mathfrak{R}^\kappa \to \mathfrak{R}$ *is a compact norm, then the functor* $\mathfrak{f} = \prod_{i\leq\kappa}^{v} \mathfrak{f}_i$ *which assigns to a model* $\mathcal{M}$ *the space* $\prod_{i\leq\kappa} \mathfrak{f}_i(\mathcal{M})$ *with metric given by* $v(d_1, ..., d_n, ...)$, *where* d_i *is the metric on* $\mathfrak{f}_i(\mathcal{M})$, *is an ultrafunctor.*

Proof: First we need to define the natural isomorphisms. Let (I, U) be an ultrafilter pair. Let $\iota_i = \iota^{\mathfrak{f}_i, U}$ be the natural isomorphism for $\mathfrak{f}_i$ and U. The sequence ι_i gives rise to a map $\iota : \prod_{j\leq\kappa} \mathfrak{f}_j(\prod_U \mathcal{M}_i) \to \prod_{j\leq\kappa} \prod_U \mathfrak{f}_j(\mathcal{M}_i)$. This map is an isometry, because each ι_i is. Now we need to define an isometry $\prod_{j\leq\kappa} \prod_U \mathfrak{f}_j(\mathcal{M}_i) \to \prod_U \prod_{j\leq\kappa} \mathfrak{f}_j(\mathcal{M}_i)$. To do so, define

$$(\langle x_{in} : i \in I\rangle_U : n < \kappa) \mapsto \langle (x_{in} : n < \kappa) : i \in I\rangle_U.$$

We show that this map is well-defined. Suppose

$$(\langle x_{in} : i \in I\rangle_U : n < \kappa) = (\langle y_{in} : i \in I\rangle_U : n < \kappa).$$

Then for every $\varepsilon > 0$, the set

$$\{i \in I : d(x_{in}, y_{in}) < \varepsilon\} \in U$$

for every n. Let $\varepsilon > 0$. Since v is a compact norm, there is $N < \omega$ and $\delta > 0$ such that if $x_1, ..., x_n < \delta$, then $v(x_1, ..., x_n, ...) < \varepsilon$. For $k \leq N$, let $P_k = \{i \in I : d(x_{in}, y_{in}) < \delta\} \in U$, and let $P = \bigcap_{k \leq N} P_k$. Then $P \in U$, and for every $i \in P$, $d(x_{in}, y_{in}) < \delta$, so that for every $i \in P$, $v(d(x_{i1}, y_{i1}), ..., (x_{in}, y_{in}), ...) < \varepsilon$. This implies that

$$\{i \in I : v(d(x_{i1}, y_{i1}), ..., (x_{in}, y_{in}), ...) < \varepsilon\} \in U.$$

By definition, this means

$$\{i \in I : d(\bar{x}, \bar{y}) < \varepsilon\} \in U$$

which in turns means

$$\langle (x_{in} : n < \kappa) : i \in I \rangle_U = \langle (x_{in} : n < \kappa) : i \in I \rangle_U$$

showing that our map is well defined. It is also easily seen to be onto. The proof that is is an isometry is a routine computation using the fact that v, being a compact norm, commutes with the taking of ultralimits. This completes the definition of the natural isomorphisms, showing that $\mathfrak{f}$ is a pre-ultrafunctor. To show that $\mathfrak{f}$ is an ultrafunctor, let G be an ultragraph, and Δ be an ultradiagram. We have to show that the diagram

$$\begin{array}{ccc}
\mathfrak{f}\left(\prod_{f[U]} \Delta(i)\right) & \xrightarrow{\mathfrak{f}(\partial_\Delta)} & \mathfrak{f}\left(\prod_U \Delta(j)\right) \\
\cong \downarrow & & \downarrow \cong \\
\prod_{f[U]} \mathfrak{f}(\Delta(i)) & \xrightarrow[\partial_{\mathfrak{f} \circ \Delta}]{} & \prod_U \mathfrak{f}(\Delta(j))
\end{array}$$

is commutative. An element of $\mathfrak{f}\left(\prod_{f[U]} \Delta(i)\right)$ is of the form $(\langle x_{in} : i \in I \rangle_{f[U]} : n < \kappa)$. By definition, we have

$$\begin{aligned}
\mathfrak{f}(\partial_\Delta)((\langle x_{in} : i \in I \rangle_{f[U]} : n < \kappa)) &= (\mathfrak{f}_n(\partial_\Delta)(\langle x_{in} : i \in I \rangle_{f[U]}) : n < \kappa) \\
&= (\partial_{\mathfrak{f}_n \circ \Delta}(\langle x_{in} : i \in I \rangle_{f[U]}) : n < \kappa) \\
&= (\langle x_{f(j)n} : j \in J \rangle_U : n < \kappa).
\end{aligned}$$

Here the passage from the first to the second line is possible because every $\mathfrak{f}_n$ is an ultrafunctor. We also have

$$\partial_{\mathfrak{f} \circ \Delta}(\langle (x_{in} : i < \kappa) : i \in I \rangle_{f[U]}) = \langle (x_{f(j)n} : n < \kappa) : j \in J \rangle_U$$

A close look at the definition of the natural isomorphism $\iota : \mathfrak{f}(\prod_U \Delta(j)) \to \prod_U \mathfrak{f}(\Delta(j))$ reveals that

$$\iota(\langle x_{f(j)n} : j \in J \rangle_U : n < \kappa)) = \langle (x_{f(j)n} : n < \kappa) : j \in J \rangle_U$$

and that

$$\iota^{-1}(\langle x_{in} : i \in I \rangle_{f[U]} : n < \kappa)) = \langle (x_{in} : n < \kappa) : i \in I \rangle_{f[U]}$$

which shows that the square commutes. □

4.2.12 Theorem: *If* S *is any sort of* $\mathscr{L}^{eq}$*, then the functor* $\mathrm{Mod}(T) \to \mathbb{MET}$ *defined by* $\mathcal{M} \mapsto \mathsf{S}(\mathcal{M})$ *and* $h \mapsto \mathsf{S}(h)$ *is an ultrafunctor.*

Proof: We use induction on the construction of $\mathscr{L}^{eq}$ sorts. By definition, an $\mathscr{L}^{eq}$ sort is:

1. A sort of $\mathscr{L}$
2. A countable product of sorts of $\mathscr{L}^{eq}$
3. A canonical parameter sort for a formula φ of $\mathscr{L}^{eq}$
4. A finite union of canonical parameter sorts

If S is a sort of $\mathscr{L}$, then as a functor $\mathsf{S} : \mathrm{Mod}(T) \to \mathbb{MET}$ it is an ultrafunctor by the very definition of ultraproduct of $\mathscr{L}$-structures, and by the comments after definition 4.2.4.

First we deal with the easier case of a canonical parameter sort. By definition, if S is a canonical parameter sort, then if is of the form S_φ for some $\varphi \in \mathscr{L}$. There is a sort S of $\mathscr{L}$, and a definable surjective map $\pi : \mathsf{S} \to \mathsf{S}_\varphi$. In order to show that S_φ is an ultrafunctor, we need to define the natural isomorphisms $\mathsf{S}_\varphi \circ \prod_U \cong \prod_U \circ \mathsf{S}_\varphi^I$, for ever ultrafilter pair (I, U). Let (I, U) be an ultrafilter pair, and write ι for the natural isomorphism $\iota^{\mathsf{S},U}$. We have the incomplete square

$$\begin{array}{ccc} \mathsf{S}(\prod_U \mathcal{M}_i) & \xrightarrow{\cong} & \prod_U \mathsf{S}(\mathcal{M}_i) \\ {\scriptstyle\pi}\downarrow & & \downarrow{\scriptstyle\pi} \\ \mathsf{S}_\varphi(\prod_U \mathcal{M}_i) & \dashrightarrow & \prod_U \mathsf{S}_\varphi(\mathcal{M}_i) \end{array}$$

in which the bottom arrow should be an instance of the natural isomorphism. Note that since π is onto, there is only one way to define this bottom arrow so that the diagram commutes. We must define it via

$$f(\pi(x)) = \langle \pi(\iota(x)_i) \rangle_U.$$

It is clear that this map is well-defined, and that it is onto. We argue that it is an isometry.

$$\begin{aligned} d(f(\pi(x)), f(\pi(y))) &= d(\langle \pi(\iota(x)_i)\rangle_U, \langle \pi(\iota(y)_i)\rangle_U) \\ &= \lim_{i \to U} d(\pi(\iota(x)_i), \pi(\iota(y)_i)) \\ &= \lim_{i \to U} \sup_{z \in \mathsf{S}'(\mathcal{M}_i)} |\varphi(z, \iota(x)_i) - \varphi(z, \iota(y)_i)| \\ &= \sup_{z \in \mathsf{S}'(\prod_U \mathcal{M}_i)} |\varphi(z, x) - \varphi(z, y)| \\ &= d(\pi(x), \pi(y)) \end{aligned}$$

This shows that the sort S_φ defines a pre-ultrafunctors. To show that S_φ in fact define an ultrafunctor, we must show that for every ultragraph $G = (I, J, U, f)$, and any G-ultradiagram Δ, the diagram

$$\begin{array}{ccc} \mathsf{S}_\varphi\left(\prod_{f[U]} \Delta(i)\right) & \xrightarrow{\mathsf{S}_\varphi(\partial_\Delta)} & \mathsf{S}_\varphi\left(\prod_U \Delta(P)\right) \\ {\scriptstyle\cong}\downarrow & & \downarrow{\scriptstyle\cong} \\ \prod_{f[U]} \mathsf{S}_\varphi(\Delta(i)) & \xrightarrow[\partial_{\mathsf{S}_\varphi \circ \Delta}]{} & \prod_U \mathsf{S}_\varphi(\Delta(P)) \end{array}$$

is commutative. To do so we consider the cube in which, for the sake of simplicity all the $\cong$ symbols represent the natural isomorphisms.

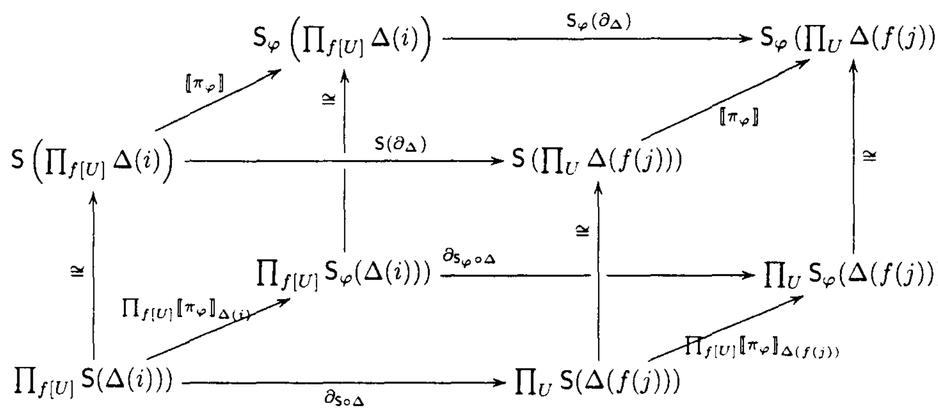

We carry out an argument similar to the one used to prove 4.2.10. The front face commutes by our assumption that S is an ultrafunctor. The top face commutes because π_φ is a definable map and ∂_Δ is elementary. The bottom face commutes by the definition of $\partial_{S\circ\Delta}$ and $\partial_{S_\varphi\circ\Delta}$. The left and right faces commute also because π_φ is a definable map. Since π_φ is surjective, this implies that the back face commutes as well, which is what we wanted to prove. The proof for the case of a finite union is similar, and the case of products is taken care of in Lemma 4.2.11. □

Theorem 4.2.12 is in fact equivalent to the following theorem relating T and T^{eq}. We omit the proof of Theorem 4.2.13 because it is identical to the proof of Theorem 4.2.12.

4.2.13 Theorem: *The forgetful functor* $F : \mathrm{Mod}(T^{eq}) \to \mathrm{Mod}(T)$ *is an ultra-equivalence of categories.*

4.2.14 Theorem: *A functor* $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ *is an ultrafunctor if and only if* $\mathrm{dom}(\mathfrak{f}) : \mathrm{Mod}(T) \to \mathbb{MET}$ *is an ultrafunctor, and* $\mathfrak{f}(\prod_U \mathcal{M}_i)(\langle x_i\rangle_U) = \lim_{i\to U} \mathfrak{f}(\mathcal{M}_i)(x_i)$ *for every ultrafilter pair* (I, U).

Proof: It is clear by definition that if $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ is an ultrafunctor, then $\mathrm{dom}(\mathfrak{f}) : \mathrm{Mod}(T) \to \mathbb{MET}$ is an utlrafunctor, and for every ultrafilter pair (I, U), $\mathfrak{f}(\prod_U \mathcal{M}_i)(\langle x : i \in I\rangle_U) = \lim_{i\to U} \mathfrak{f}(\mathcal{M}_i)(x_i)$.

For the converse, let (I, J, U, f) be an ultragraph and Δ be an ultradiagram. We investigate the commutativity of the following diagram, in which the maps $g : \prod_{f[U]} \mathfrak{f}(\Delta(i)) \to \mathfrak{R}$ and $g' : \prod_U \mathfrak{f}(\Delta(f(j))) \to \mathfrak{R}$ are defined by

$$g(\langle x_i\rangle_{f[U]}) = \lim_{i\to f[U]} \mathfrak{f}(\Delta(i))(x_i)$$

and

$$g'(\langle x_{f(j)}\rangle_U) = \lim_{j\to U} \mathfrak{f}(\Delta(f(j)))(x_{f(j)})$$

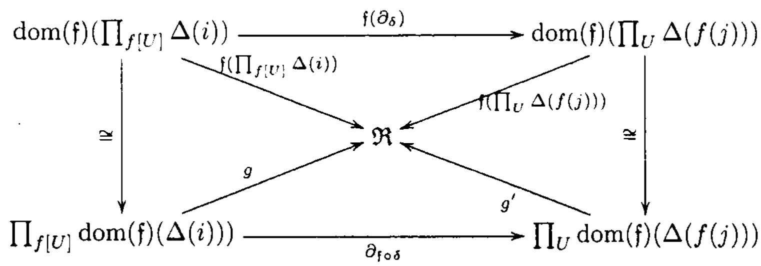

We note that the outer square commutes because dom($\mathfrak{f}$) is an ultrafunctor. The top triangle commutes because ∂_Δ is an elementary map. The left and right triangles commute by our assumption that $\mathfrak{f}(\prod_U \mathcal{M}_i)(\langle x_i \rangle_U) = \lim_{i \to U} \mathfrak{f}(\mathcal{M}_i)(x_i)$ for every ultrafilter pair (I, U). This is in fact enough to imply that the bottom triangle commutes. Therefore, the whole diagram commutes. Since the diagram above is the $\mathbb{MET}_{\mathfrak{R}}$ incarnation of (††), we conclude that $\mathfrak{f}$ is an ultrafunctor. □

4.2.15 Definition: For every formula $\varphi \in \mathscr{L}^{\mathrm{eq}}$, we define the functor $\mathfrak{ev}_\varphi : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$

1. for every model $\mathcal{M} \in \mathrm{Mod}(T)$,

   a) $\mathrm{dom}(\mathfrak{ev}_\varphi(\mathcal{M})) = \mathrm{dom}(\varphi)(\mathcal{M})$

   b) $\mathfrak{ev}_\varphi(\mathcal{M})(x) = [\![\varphi]\!]_{\mathcal{M}}(x)$.

2. if $h : \mathcal{M} \to \mathcal{N}$ is an elementary map, then we let $\mathfrak{ev}_\varphi(h) = \mathrm{dom}(\varphi)(h)$ ♣

4.2.16 Proposition: *For every $\varphi \in \mathscr{L}^{eq}$, $\mathfrak{ev}_\varphi$ is an ultrafunctor.*

Proof: By Theorem 4.2.14, this is a direct consequence of the fact that dom(φ) is an ultrafunctor (Theorem 4.2.12), and the fundamental theorem of ultraproducts. □

4.2.17 Definition: Let $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ be a functor. Then $\mathfrak{f}$ is a *definable*, or *representable* functor if and only if there are formulae $\varphi, \psi \in \mathscr{L}^{\mathrm{eq}}$ such that:

1. $\mathrm{dom}(\varphi) = \mathrm{dom}(\psi)$
2. The set $\{x \in \mathrm{dom}(\psi)(\mathcal{M}) : \mathcal{M} \models \psi(x)\}$ is definable;
3. $\mathrm{dom}(\mathfrak{f})(\mathcal{M}) = \{x \in \mathrm{dom}(\psi)(\mathcal{M}) : \mathcal{M} \models \psi(x)\}$
4. For every $x \in \mathrm{dom}(\mathfrak{f})(\mathcal{M})$, $\mathfrak{f}(\mathcal{M})(x) = \mathfrak{ev}_\varphi(\mathcal{M})(x)$ ♣

In Theorem 4.2.20 and Corollary 4.2.21, we show that $\mathfrak{f}$ being defnable is equivalent ot the existence of a formula φ and an injective $\eta : \mathfrak{f} \to \mathfrak{ev}_\varphi$.

4.2.18 Theorem: *Any definable functor is an ultrafunctor.*

Proof: By 4.2.14, all we need to show is that dom($\mathfrak{f}$) is an ultrafunctor. Since dom($\mathfrak{f}$) is a subfunctor of dom(φ), all we need to show is that dom($\mathfrak{f}$) commutes with the ultraproduct functors. Let (I, U) be an ultrafilter pair. By assumption, $\mathrm{dom}(\mathfrak{f})(\prod_U \mathcal{M}_i)$ is a definable subset of $\prod_U \mathcal{M}_i$. Therefore, for any $\langle x_i \rangle_U \in \prod_U \mathcal{M}_i$, we have

$$d(\langle x_i \rangle_U, dom(\mathfrak{f})(\prod_U \mathcal{M}_i)) = \lim_{i \to U} d(x_i, \mathrm{dom}(\mathfrak{f})(\mathcal{M}_i)).$$

Suppose $d(\langle x_i \rangle_U, \mathrm{dom}(\mathfrak{f})(\prod_U \mathcal{M}_i)) = 0$. By definition, this is true if and only if for every $\varepsilon > 0$,

$$\{i \in I : d(x_i, \mathrm{dom}(\mathfrak{f})(\mathcal{M}_i)) < \varepsilon\} \in U$$

which in turn is true if and only if

$$\langle x_i \rangle_U \in \prod_U \mathrm{dom}(\mathfrak{f})(\mathcal{M}_i))$$

showing that

$$\mathrm{dom}(\mathfrak{f})(\prod_U \mathcal{M}_i) = \prod_U \mathrm{dom}(\mathfrak{f})(\mathcal{M}_i)$$

as required to complete the proof. □

4.2.19 Lemma: *Let* $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ *be an ultrafunctor. Suppose there is a sort of* $\mathscr{L}$ *such that for every* $\mathcal{M} \in \mathrm{Mod}(T)$, $\mathrm{dom}(\mathfrak{f}(\mathcal{M})) = \mathsf{S}(\mathcal{M})$. *Then* $\mathfrak{f} \cong \mathfrak{ev}_\varphi$ *for some formula* φ.

Proof: We show that in every model $\mathcal{M}$, $\mathfrak{f}$ is constant on types, i.e. that if $\mathrm{tp}(a) = \mathrm{tp}(b)$ for $a, b \in \mathcal{M}$, then $\mathfrak{f}(\mathcal{M})(a) = \mathfrak{f}(\mathcal{M})(b)$. Let $\mathcal{M} \in \mathrm{Mod}(T)$ be any model, and $a \equiv b$ in $\mathcal{M}$. Since $a \equiv b$, by 3.4.8, there is an elementary embedding $h : \mathcal{M} \to \mathcal{M}^V$ for some ultrafilter V with the property that $h(a) = \partial_{\mathcal{M}}(b)$. Therefore, $\mathfrak{f}(\mathcal{M}^V)(h(a)) = \mathfrak{f}(\mathcal{M}^V)(\partial_{\mathcal{M}}(b))$. Since h and $\partial_{\mathcal{M}}$ are elementary, $\mathfrak{f}(\mathcal{M}^V)(h(a)) = \mathfrak{f}(\mathcal{M})(a)$ and $\mathfrak{f}(\mathcal{M}^V)(\partial_{\mathcal{M}}(b)) = \mathfrak{f}(\mathcal{M})(b)$, showing that $\mathfrak{f}(\mathcal{M})(a) = \mathfrak{f}(\mathcal{M})(b)$. By corollary 3.3.9, there is a formula φ such that $\mathfrak{f}(\mathcal{M}) = [\![\varphi]\!]_{\mathcal{M}}$. Note that even though corollary 3.3.9 is stated for a specific model, since $\mathfrak{f}$ is a functor, the same formula φ can be used for all models. By the definition of $\mathfrak{ev}_\varphi$, $\mathfrak{f} \cong \mathfrak{ev}_\varphi$. □

4.2.20 Theorem: *Let* $\mathfrak{f} \in \mathbb{ULT}_{\mathbb{M}}(T)$, *and let* $\eta : \mathfrak{f} \to \mathfrak{ev}_\varphi$ *be a monomorphism. Then* $\mathrm{im}(\eta)$ *is a definable set in the following sense: there is a formula* ψ *such that in any model* $\mathcal{M} \in \mathrm{Mod}(T)$, $[\![\psi]\!]_{\mathcal{M}}(x) = d(x, \mathrm{im}(\eta_{\mathcal{M}}))$.

Proof: Let $\mathfrak{g}$ be the functor defined by $\mathrm{dom}(\mathfrak{g}(\mathcal{M})) = \mathrm{dom}(\varphi)(\mathcal{M})$, and

$$\mathfrak{g}(\mathcal{M})(x) = d(x, \mathrm{im}(\varphi_{\mathcal{M}}))$$

We claim that $\mathfrak{g}$ is a ultrafunctor, which will make $d(x, \mathrm{im}(\varphi_{\mathcal{M}}))$ into a formula by lemma 4.2.19. Since $\mathrm{dom}(\mathfrak{g}) = \mathrm{dom}(\varphi)$ is an ultrafunctor by Theorem 4.2.14, it is enough to show that $\mathfrak{g}(\prod_U \mathcal{M}_i)(\langle x_i \rangle_U) = \lim_{i \to U} \mathfrak{g}(\mathcal{M}_i)(x_i)$.

Let U be an ultrafilter on I. We want to show that

$$d(x, \mathrm{im}(\eta_{\prod_U \mathcal{M}_i})) = \lim_{i \to U} d(x_i, \mathrm{im}(\eta_{\mathcal{M}_i})).$$

Suppose $r < \lim_U d(x_i, \mathrm{im}(\eta_{\mathcal{M}_i})) < s$. Then by definition,

$$J = \{j \in I : r < d(x_j, \mathrm{im}(\eta_{\mathcal{M}_j})) < s\} \in U$$

For every $j \in J$ and $z_i \in \mathrm{im}(\eta_{\mathcal{M}_i})$, we have $r < d(x_i, z_i)$. Fix any $\langle z_i \rangle_U \in \prod_U \mathrm{im}(\eta_{\mathcal{M}_i})$. Then since $z_i \in \mathrm{im}(\eta_{\mathcal{M}_i})$, we have $d(x, \langle z_i \rangle) > r$. Therefore $d(x, \prod_U \mathrm{im}(\eta_{\mathcal{M}_i})) > r$ as required.

For every $j \in J, d(x_i, \mathrm{im}(\eta_{\mathcal{M}_i})) < s$. Therefore, there is $z_i \in \mathrm{im}(\eta_{\mathcal{M}_i})$ such that $d(x_i, z_i) \leq s$. Let $z_i' \in (\eta_{\mathcal{M}_i})$ for $i \in I \backslash J$, and let $\langle z_i \rangle_U$ be the resulting sequence. Then $d(x, \langle z \rangle) < s$. This completes the proof. □

4.2.21 Corollary: *An ultrafunctor* $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ *is definable if and only if there is a formula* φ *of* $\mathscr{L}^{eq}$ *and a monomorphism* $\eta : \mathfrak{f} \to \mathfrak{ev}_\varphi$.

Proof: If $\eta : \mathfrak{f} \to \mathfrak{ev}_\varphi$ is a monomorphism, then $\mathfrak{f}$ is a definable functor by Theorem 4.2.20, letting ψ be the formula defining $\mathrm{im}(\eta)$. If $\mathfrak{f}$ is a definable ultrafunctor, then by definition there are formulae ψ and φ such that

1. $\mathrm{dom}(\varphi) = \mathrm{dom}(\psi)$
2. $\mathrm{dom}(\mathfrak{f})(\mathcal{M}) = \{x \in \mathrm{dom}(\psi)(\mathcal{M}) : \mathcal{M} \models \psi(x)\}$
3. $\mathfrak{f}(\mathcal{M})(x) = \mathfrak{ev}_\varphi(\mathcal{M})(x)$

For $\mathcal{M} \in \mathrm{Mod}(T)$, let $\eta_\mathcal{M}$ be the inclusion map $\mathrm{dom}(\mathfrak{f})(\mathcal{M}) \to \mathrm{dom}(\varphi)$. The map $\eta_\mathcal{M}$ is a monomorphism by definition. The fact that it is an ultratransformation follows from the fundamental theorem of ultraproducts. □

We can put proposition 4.2.10 and corollary 4.2.21 together in the following way, which is a complete description of which functors $\mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ are definable. A functor $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ is a definable functor if and only if

1. for every ultrafilter pair (I, U), $\mathfrak{f} \circ \prod_U \cong \prod_U \circ \mathfrak{f}^I$
2. there is an $\mathscr{L}^{eq}$ formula φ, and a monic ultratransformation $\eta : \mathfrak{f} \to \mathfrak{ev}_\varphi$

By 4.2.10, any functor satisfying these two conditions is an ultrafunctor, and therefore a definable functor by 4.2.21. The rest of this chapter is devoted to the proof of the following theorem:

4.2.22 Theorem (Strong conceptual completeness): *If* $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ *is an ultrafunctor, then* $\mathfrak{f}$ *is a definable functor.*

Section 4.3

Domination and covers

Let us briefly outline the proof of Theorem 4.2.22. For this section, we fix an ultrafunctor $\mathfrak{f} \in \mathbb{ULT}_\mathbb{M}(T)$. We will define an expansion $\mathscr{L}_\mathfrak{f}$ of $\mathscr{L}$ which will consist of a new sort symbol $\mathsf{S}_\mathfrak{f}$ standing in for $\mathrm{dom}(\mathfrak{f})$, and predicate symbols for all the ultrafunctors $\mathfrak{g}$ such that $\mathrm{dom}(\mathfrak{g}) = \mathsf{S} \times \mathrm{dom}(\mathfrak{f})$, where S is a sort of $\mathscr{L}$. We then show that any model $\mathcal{M}$ can be expanded to an $\mathscr{L}_\mathfrak{f}$-structure $\mathcal{M}_\mathfrak{f}$ in an easy way, and that in fact $\mathcal{M}$ is stably embedded in this expansion $\mathcal{M}_\mathfrak{f}$. In fact, the class $K = \{\mathcal{N} : \mathcal{N} \cong \mathcal{M}_\mathfrak{f}$ for some $\mathcal{M} \in \mathrm{Mod}(T)\}$ is an elementary class, and as a category of structures, it is equivalent to $\mathrm{Mod}(T)$, and its inverse is the forgetful functor $F : K \to \mathrm{Mod}(T)$.

Recall that attached to every ultrafunctor $\mathfrak{f}$, there is an ultrafunctor $\mathrm{dom}(\mathfrak{f}) : \mathrm{Mod}(T) \to \mathbb{MET}$ defined by $\mathcal{M} \mapsto \mathrm{dom}(\mathfrak{f})(\mathcal{M})$. First we show that when an ultrafunctor $\mathfrak{f}$ is fixed, then there is a *set* of ultrafunctors $\mathfrak{g}$ with the property that as functors, $\mathrm{dom}(\mathfrak{g}) = \mathrm{dom}(\mathfrak{f})$. The Löwenheim-Skolem number of $\mathscr{L}$ is the least cardinal κ with the property that every $\mathscr{L}$-structure $\mathcal{X}$ has an elementary substructure $\mathcal{Y}$ such that $|\mathcal{Y}| \le \kappa$.

4.3.1 Theorem: *Let* $\mathfrak{f}, \mathfrak{g} \in \mathbb{ULT}_{\mathbf{M}}(T)$ *be such that* $\mathrm{dom}(\mathfrak{f}) = \mathrm{dom}(\mathfrak{g})$, *and let* κ *be the Löwenheim-Skolem number of* T. *Suppose* $\mathfrak{f}|\mathrm{Mod}(T,\kappa) = \mathfrak{g}|\mathrm{Mod}(T,\kappa)$. *Then* $\mathfrak{f} = \mathfrak{g}$.

Proof: Let $\mathcal{M} \in \mathrm{Mod}(T)$, $x \in \mathcal{M}$. Choose $\mathcal{M}_0 \in \mathrm{Mod}(T,\kappa)$ with $\mathcal{M}_0 \preceq \mathcal{M}$. Let $g : \mathcal{M}_0 \to \mathcal{M}$ be an elementary embedding. Apply Theorem 3.4.9 to get an ultrafilter U and an embedding $h : \mathcal{M} \to \mathcal{M}_0^U$ such that $h \circ g = \partial_{\mathcal{M}_0}$.

Since we are assuming that $\mathfrak{f}$ and $\mathfrak{g}$ are ultrafunctors, we also get that for every $\langle x_i \rangle_U \in \mathcal{M}_0^U$,

$$\mathfrak{f}(\mathcal{M}_0^U)(\langle x_i \rangle_U) = \lim_{i\to U} \mathfrak{f}(\mathcal{M}_0)(x_i) = \lim_{i \to U} \mathfrak{g}(\mathcal{M}_0)(x_i) = \mathfrak{g}(\mathcal{M}_0^U)(\langle x_i \rangle_U).$$

Since $h : \mathcal{M} \to \mathcal{M}_0^U$ is elementary, for every $x \in \mathcal{M}$,

$$\mathfrak{f}(\mathcal{M})(x) = \mathfrak{f}(\mathcal{M}_0^U)(h(x)) = \mathfrak{g}(\mathcal{M}_0^U)(h(x)) = \mathfrak{g}(\mathcal{M})(x)$$

showing that $\mathfrak{f}(\mathcal{M}) = \mathfrak{g}(\mathcal{M})$ for every $\mathcal{M}$, the thus that $\mathfrak{f} = \mathfrak{g}$. □

We use Theorem 4.3.1 to argue that there is, up to equivalence, a *set* of ultrafunctors whose domains are of the form $\mathsf{S} \times \mathrm{dom}(\mathfrak{f})$ for S a sort of $\mathscr{L}^{\mathrm{eq}}$. Let κ be the Löwenheim-Skolem number of T. There are set many isomorphism classes of objects in the category $\mathrm{Mod}(T,\kappa)$ form a set. Therefore, if $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ is any functor, since $\mathfrak{f}$ sends isomorphisms to isomorphisms, then there is set many isomprphism classes in the image of $\mathrm{dom}(\mathfrak{f})|\mathrm{Mod}(T,\kappa)$ in $\mathbb{MET}$. This means that there is a cardinal λ such that every $X \in \mathrm{im}(\mathrm{dom}(\mathfrak{f})|\mathrm{Mod}(T,\kappa))$ has density character $\le \lambda$. By Theorem 4.3.1, an ultrafunctor $\mathfrak{g}$ such that $\mathrm{dom}(\mathfrak{g}) = \mathrm{dom}(\mathfrak{f})$ is determined by the choice of a function $\mathfrak{g}(\mathcal{M}) : \mathrm{dom}(\mathfrak{g})(\mathcal{M}) \to \mathfrak{R}$ for every $\mathcal{M} \in \mathrm{Mod}(T,\kappa)$. Since there is an upper bound on the density character of $\mathrm{dom}(\mathfrak{g})(\mathcal{M})$, the collection of such choices forms a set, and therefore the collection of ultrafunctors $\mathfrak{g}$ such that $\mathrm{dom}(\mathfrak{g}) = \mathrm{dom}(\mathfrak{f})$ forms a set. Let Σ be the following set:

$$\Sigma = \{\mathfrak{g} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}} : \mathrm{dom}(\mathfrak{g}) = \mathsf{S} \times \mathrm{dom}(\mathfrak{f})^n, \mathsf{S} \text{ a sort of } \mathscr{L}^{\mathrm{eq}}, n \in \mathbb{N}\}$$

Let $\mathscr{L}_{\mathfrak{f}}$ be the following expansion of $\mathscr{L}$:

1. Add a new sort symbol $\mathsf{S}_{\mathfrak{f}}$
2. For every $\mathfrak{g} \in \Sigma$, add a predicate symbol $P_{\mathfrak{g}}$ with domain $\mathsf{S} \times \mathsf{S}_{\mathfrak{f}}^n$.

4.3.2 Theorem: *Every formula* $\varphi \in \mathscr{L}_{\mathfrak{f}}$ *corresponds to an ultrafunctor* $\overline{\varphi} : \mathrm{Mod}(T) \to \mathbb{MET}_{\mathfrak{R}}$ *such that for every* $\mathcal{M} \in \mathrm{Mod}(T_{\mathfrak{f}})$ $[\![\mathsf{P}_{\overline{\varphi}}]\!]_{\mathcal{M}} = [\![\varphi]\!]_{\mathcal{M}}$.

Proof: We show that it is possible to combine ultrafunctors using continuous connectives and quantification to obtain new ultrafunctors. Recall that if I is a set, then the category $\mathbb{MET}_{\mathfrak{R}}^I$ has I-indexed sequences of elements of $\mathbb{MET}_{\mathfrak{R}}$ as objects, and I-indexed sequences of morphisms in $\mathbb{MET}_{\mathfrak{R}}$ as morphisms.

CLAIM B: *Let* $0 \le \kappa \le \omega$, *and let* $u : \mathfrak{R}^\kappa \to \mathfrak{R}$ *be continuous, and let* v *be a compact norm on* $\mathfrak{R}^\kappa$. *Then* u *induces a functor* $u : \mathbb{MET}_{\mathfrak{R}}^\kappa \to \mathbb{MET}_{\mathfrak{R}}$ *defined as follows:*

1. $\mathrm{dom}(u)(\varphi_1, ..., \varphi_n, ...) = \prod_{i \leq \kappa} \mathrm{dom}(\varphi_i)$, *where the metric on* $\prod_{i\leq\kappa} \mathrm{dom}(\varphi_i)$ *is given by* $v(d_1, ..., d_n, ...)$, *with* d_i *the metric on* $\mathrm{dom}(\varphi_i)$
2. $u(\varphi_1, ..., \varphi_n, ...)(x_1, ..., x_n, ...) = u(\varphi_1(x_1), ..., \varphi_n(x_n), ...)$
3. *If* $(f_i : i \in I)$ *is a sequence of maps, then*

$$u(f_1, ..., f_n, ...)(x_1, ..., x_n, ...) = (f_1(x_1), ..., f_n(x_n), ...)$$

PROOF: All we need to prove is that $u(f_1, ..., f_n, ...)$ is indeed a map in $\mathbb{MET}_{\mathfrak{R}}$, i.e. that it defines an isometry on $\prod_{i\leq\kappa} \mathrm{dom}(\varphi_i)$. By definition, if $x = (x_i : i \in I)$ and $y = (y_i : i \in I)$ are elements of $\prod_{i\leq\kappa} \mathrm{dom}(\varphi_i)$, then $d(x, y) = v(d_1(x_1, y_1), ..., d_n(x_n, y_n), ...)$. Since f_i is assumed to be an isometry for every $i \in I$, we get that $d_i(x_i, y_i) = d_i(f_i(x_i), f_i(y_i))$, so that

$$\begin{aligned} v(d_1(x_1, y_1), ..., d_n(x_n, y_n), ...) &= v(d_1(f_1(x_1), f_1(y_1)), ..., d_n(f_n(x_n), f_n(y_n)), ...) \\ &= d(f(x), f(y)) \end{aligned}$$

completing the proof. □

CLAIM C: *If* $0 \leq \kappa \leq \omega$, $u : \mathfrak{R}^\kappa \to \mathfrak{R}$ *is a logical connective,* $v : \mathfrak{R}^\kappa \to \mathfrak{R}$ *is a compact norm, and* $(\mathfrak{f}_i : i \leq \kappa)$ *is a sequence of ultrafunctors, then the functor* $u(\mathfrak{f}_1, ..., \mathfrak{f}_n, ...)$ *defined by*

1. $\mathrm{dom}(u)(\mathfrak{f}_1, ..., \mathfrak{f}_n, ...)(\mathcal{M}) = \prod_{i\leq\kappa} \mathrm{dom}(\mathfrak{f}_i)(\mathcal{M})$, *where the metric on* $\prod_{i\leq\kappa} \mathrm{dom}(\varphi_i)(\mathcal{M})$ *is given by* $v(d_1, ..., d_n, ...)$, *with* d_i *the metric on* $\mathrm{dom}(\mathfrak{f}_i)(\mathcal{M})$
2. $u(\mathfrak{f}_1, ..., \mathfrak{f}_n, ...)(\mathcal{M})(x_1, ..., x_n, ...) = u(\mathfrak{f}_1(\mathcal{M})(x_1), ..., \mathfrak{f}_n(\mathcal{M})(x_n), ...)$

is an ultrafunctor.

PROOF: This is a direct application of Proposition 4.2.10 and the fact that ultralimits commute with continuous functions. □

CLAIM D: *Let* $\mathfrak{f}$ *be an ultrafunctor with domain* $\mathrm{dom}(\mathfrak{f}) = \mathsf{X} \times \mathsf{Y}$, *where* $\mathsf{X}, \mathsf{Y} : \mathrm{Mod}(T) \to \mathbb{MET}$ *are ultrafunctors. Then the functor* $\forall x[\mathfrak{f}(x, y)]$ *defined by*

1. $\mathrm{dom}(\forall x[\mathfrak{f}]) = \mathsf{Y}$
2. $\forall x[\mathfrak{f}](\mathcal{M})(y) = \sup_{x\in\mathsf{X}(\mathcal{M})} [\mathfrak{f}(\mathcal{M})(x, y)]$

is an ultrafunctor.

PROOF: We need to show that $\forall x[\mathfrak{f}]$ commutes with ultraproduct functors. It is clear that $\mathrm{dom}(\forall x[\mathfrak{f}])(\prod_U \mathcal{M}_i) = \prod_U \mathrm{dom}(\forall x[\mathfrak{f}](\mathcal{M}_i))$, since $\mathrm{dom}(\forall x[\mathfrak{f}]) = \mathsf{Y}$ and Y is an ultrafunctor. □

There is a similarly given ultrafunctor $\exists x[\mathfrak{f}]$, defined by $\exists x[\mathfrak{f}](\mathcal{M})(y) = \inf_{x\in\mathsf{X}(\mathcal{M})} [\mathfrak{f}(\mathcal{M})(x, y)]$. We omit the proof that $\exists x[\mathfrak{f}]$ is an ultrafunctor, as it is identical to the proof given above. We define $\overline{\varphi}$ as follows:

1. $\overline{\mathsf{P}_{\mathfrak{g}}} = \mathfrak{g}$. This is an ultrafunctor by definition of $\mathsf{P}_{\mathfrak{g}}$. Note that the case of an atomic $\mathscr{L}$-formula is covered by ultrafunctors $\mathfrak{g}$ with domain S for S a sort of $\mathscr{L}$.

2. If $u : \mathfrak{R}^\kappa \to \mathfrak{R}$ is a connective, then $\overline{u(\varphi_1, ..., \varphi_n, ...)} = u(\overline{\varphi_1}, ..., \overline{\varphi_n}, ...)$. This is an ultrafunctor by claim C
3. $\overline{\forall x[\varphi]} = \forall x[\overline{\varphi}]$. This is an ultrafunctor by claim D
4. $\overline{\exists x[\varphi]} = \exists x[\overline{\varphi}]$.

and the proof is complete. □

4.3.3 Theorem: *An $\mathscr{L}$-structure $\mathcal{M}$ can be expanded to an $\mathscr{L}_\mathfrak{f}$-structure $\mathcal{M}_\mathfrak{f}$ by defining $\mathsf{S}_\mathfrak{f}(\mathcal{M}_\mathfrak{f}) = \mathrm{dom}(\mathfrak{f})(\mathcal{M})$, and $[\![\mathsf{P}_\mathfrak{g}]\!]_{\mathcal{M}_\mathfrak{f}} = \mathfrak{g}(\mathcal{M})$. An elementary map $h : \mathcal{M} \to \mathcal{N}$ in* Mod(T) *can be expanded to an $\mathscr{L}_\mathfrak{f}$-elementary map $h_\mathfrak{f} : \mathcal{M}_\mathfrak{f} \to \mathcal{N}_\mathfrak{f}$ by defining $\mathsf{S}_\mathfrak{f}(h_\mathfrak{f}) = \mathfrak{f}(h)$.*

Proof: We only really need to show that if $h : \mathcal{M} \to \mathcal{N}$ is an elementary map in Mod(T), then the map $h_\mathfrak{f} : \mathcal{M}_\mathfrak{f} \to \mathcal{N}_\mathfrak{f}$ defined by extending h to $\mathscr{L}_\mathfrak{f}$ via $\mathsf{S}_\mathfrak{f}(h) = \mathfrak{f}(h|\mathscr{L})$ is an elementary map. By the definition of being an ultrafunctor, for every $\mathcal{M} \in \mathrm{Mod}(T)$, we have for every $x \in \mathrm{dom}(\mathfrak{f})(\mathcal{M})$, $\mathfrak{f}(\mathcal{N})(\mathfrak{f}(h)(x)) = f(\mathcal{M})(x)$. Let φ be any $\mathscr{L}_\mathfrak{f}$ formula, and write $\mathrm{dom}(\varphi) = \mathsf{S} \times \mathsf{S}_\mathfrak{f}^\ell$, where $0 \leq k, \ell \leq \omega$ and S is a sort of $\mathscr{L}^{\mathrm{eq}}$. By Theorem 4.3.2, φ corresponds to an ultrafunctor $\overline{\varphi}$ with $\mathrm{dom}(\overline{\varphi}) = \mathsf{S} \times \mathrm{dom}(\mathfrak{f})^\ell$. By the definition of ultrafunctor, we have

$$\overline{\varphi}(\mathcal{N})(\overline{\varphi}(h)(\bar{x}, \bar{y})) = \overline{\varphi}(\mathcal{M})(\bar{x}, \bar{y})$$

and by assumption,

$$\overline{\varphi}(h)(\bar{x}, \bar{y}) = (\mathsf{S}(h)(\bar{x}), \mathfrak{f}(h)(\bar{y})).$$

Putting these two equalities together completes the proof that $\mathfrak{h}_\mathfrak{f}$ is an elementary map. □

4.3.4 Definition: Let $\mathfrak{f} : \mathrm{Mod}(T) \to \mathbb{MET}_\mathfrak{R}$ be an ultrafunctor, and let $\mathcal{M} \in K$. Let $x \in \mathrm{dom}(\mathfrak{f})(\mathcal{M})$. A *cover of* x is an element $\bar{a} \in \mathcal{M}$, such that for every $\mathcal{N} \in \mathrm{Mod}(T)$, and every pair of morphisms of Mod(T), $h_1, h_2 : \mathcal{M} \to \mathcal{N}$, if $h_1(\bar{a}) = h_2(\bar{a})$, then $\mathfrak{f}(h_1)(x) = \mathfrak{f}(h_2)(x)$. If $\bar{a}$ is a coverof x, then we will say that $\bar{a}$ *dominates* x. ♣

4.3.5 Theorem: *Let $\mathfrak{f} \in \mathbb{ULT}_\mathbb{M}(T)$, $\mathcal{M} \in \mathrm{Mod}(T)$, and let $x \in \mathrm{dom}(\mathfrak{f})(M)$. Then there is a sort $\mathsf{S} \in \mathscr{L}^{eq}$ and an $a \in \mathsf{S}(\mathcal{M})$ such that a dominates x.*

Proof: If for every $\varepsilon > 0$, there is a sort S_ε of $\mathscr{L}^{\mathrm{eq}}$ and $a_\varepsilon \in \mathsf{S}_\varepsilon(\mathcal{M})$ such that for every $\mathcal{N} \in K$, and every pair of morphisms of K, $h_1, h_2 : \mathcal{M} \to \mathcal{N}$, if $h_1(\bar{a}) = h_2(\bar{a})$, then $d(\mathfrak{f}(h_1)(x), \mathfrak{f}(h_2)(x)) \leq \varepsilon$, then the theorem holds with $\mathsf{S} = \prod_{n\in\mathbb{N}} \mathsf{S}_{1/2^n}$ and $a = (a_{1/2^n} : n \in \mathbb{N})$, since $\mathsf{S} = \prod_{n\in\mathbb{N}} \mathsf{S}_{1/2^n}$ is a sort of $\mathscr{L}^{\mathrm{eq}}$.

Suppose that the theorem does not hold. Let I be the set of all finite sets of the form $i = \{(\mathsf{S}_1, a_1), ..., (\mathsf{S}_n, a_n)\}$, where $a_k \in \mathsf{S}_k(\mathcal{M})$, S_k a sort of $\mathscr{L}^{\mathrm{eq}}$. For every $i \in I$, there are models $\mathcal{N}_i$ and maps $h_1^i, h_2^i : \mathcal{M} \to \mathcal{N}_i$ such that $h_1^i(a_j) = h_2^i(a_j)$ for $j \in i$, but $d(\mathfrak{f}(h_1^i)(x), \mathfrak{f}(h_2^i)(x)) \geq \varepsilon$.

There is an ultrafilter U on I containing all sets of the form $\{i' \in I : i' \supseteq i\}$. Let $h_1 = \prod_U h_1^i$ and $h_2 = \prod_U h_2^i$. Let S be any sort of $\mathscr{L}^{\mathrm{eq}}$, and $a \in \mathsf{S}(\mathcal{M})$, and consider $i = \{(\mathsf{S}, a)\}$. Then $h_1^j(a) = h_2^j(a)$ for every $j \supseteq i$, implying that $h_1(\partial_\mathcal{M}(a)) = h_2(\partial_\mathcal{M}(a))$, where $\partial_\mathcal{M} : \mathcal{M} \to \mathcal{M}^U$ is the canonical embedding. Since S and a were arbitrary, we can conclude that the diagram:

$$\mathcal{M} \xrightarrow{\partial} \mathcal{M}^U \underset{h_2}{\overset{h_1}{\rightrightarrows}} \prod_U \mathcal{N}_i$$

is commutative, so $h_1 \circ \partial = h_2 \circ \partial$. Now consider the following diagram, in which the vertical arrows are the canonical isomorphisms:

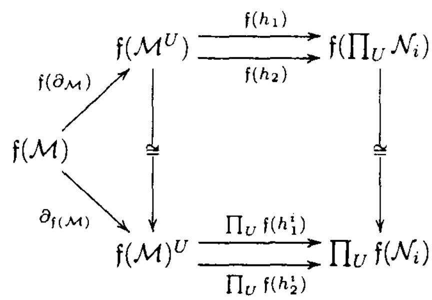

The square with the top arrows and the isomorphisms commutes, and so does the square with the bottom arrows and the isomorphisms. The triangle on the left commutes as well, since $\mathfrak{f}$ is an ultrafunctor. This implies that the entire diagram is commutative. The top row of the diagram indicates that

$$\mathfrak{f}(h_1) \circ \mathfrak{f}(\partial_{\mathcal{M}}) = \mathfrak{f}(h_2) \circ \mathfrak{f}(\partial_{\mathcal{M}}).$$

On the other hand, looking at the bottom row and remembering the definition of the individual models $\mathcal{N}_i$, we get that

$$d(\prod_U \mathfrak{f}(h_1^i)(\partial_{\mathfrak{f}(\mathcal{M})}(x), \prod_U \mathfrak{f}(h_2^i)(\partial_{\mathfrak{f}(\mathcal{M})}(x)) \geq \varepsilon > 0$$

so that

$$\prod_U \mathfrak{f}(h_1^i)(\partial_{\mathfrak{f}(\mathcal{M})}(x) \neq \prod_U \mathfrak{f}(h_2^i)(\partial_{\mathfrak{f}(\mathcal{M})}(x).$$

But by commutativity of the diagram,

$$\prod_U \mathfrak{f}(h_1^i)(\partial_{\mathfrak{f}(\mathcal{M})}(x) = \mathfrak{f}(h_1) \circ \mathfrak{f}(\partial_{\mathcal{M}})$$

and

$$\prod_U \mathfrak{f}(h_2^i)(\partial_{\mathfrak{f}(\mathcal{M})}(x) = \mathfrak{f}(h_2) \circ \mathfrak{f}(\partial_{\mathcal{M}})$$

and this contradiction finishes the proof. □

The following theorem is the keystone of the proof of strong conceptual completeness. In the context of first-order logic, a proof can be found in [Bac74]. Recall the notation $[\alpha(x) = y]$ for the formula defining a definable function $\alpha : \mathsf{S} \to \mathsf{S}'$. If $h : \mathcal{M} \to \mathcal{N}$ is an $\mathscr{L}_{\mathfrak{f}}$-elementary map, then we will use the notation $h|\mathscr{L}$ to represent the $\mathscr{L}$-reduct of h. If $h : \mathcal{M} \to \mathcal{N}$ is an $\mathscr{L}_{\mathfrak{f}}$-elementary map, then we will use the notation $h|\mathscr{L}$ to represent the $\mathscr{L}$-reduct of h.

4.3.6 Theorem: *Let $\mathcal{M}$ be an $\mathscr{L}_{\mathfrak{f}}$-structure, $x \in \mathsf{S}_{\mathfrak{f}}(\mathcal{M})$ and $\bar{a} \in \mathsf{S}(\mathcal{M})$. Then $\bar{a}$ dominates x if and only if there is an $\mathscr{L}_{\mathfrak{f}}$-definable function α such that $\mathcal{M} \models [\alpha(\bar{a}) = x]$.*

Proof: Suppose α is a definable function $\mathsf{S} \to \mathsf{S}_\mathfrak{f}$, and let $h_1, h_2 : \mathcal{M} \to \mathcal{N}$ be elementary maps with the property that $h_1(a) = h_2(a)$. We have, for every a and x, $[\alpha(a) = x] = [\alpha(h_1(a)] = \mathfrak{f}(h_1)(x)] = [\alpha(h_2(a)) = \mathfrak{f}(h_2)(x)]$ because h_1 and h_2 are elementary maps. By assumption, $h_1(a) = h_2(a)$, so that

$$[\alpha(h_2(a)) = \mathfrak{f}(h_2)(x)] = [\alpha(h_1(a)) = \mathfrak{f}(h_2)(x)] = [\alpha(h_1(a)) = \mathfrak{f}(h_1)(x)].$$

Since α is a definable function, $\mathfrak{f}(h_1)(x) = \mathfrak{f}(h_2)(x)$.

To prove the converse, we use Theorem 3.4.11. Let $\mathcal{M}_0 \in \mathrm{Mod}(T)$, and let $a_0 \in \mathsf{S}(\mathcal{M}_0)$ dominate $x_0 \in \mathsf{S}_\mathfrak{f}(\mathcal{M}_0)$. Let $T^* = \mathrm{Th}(\mathcal{M}_0, a_0, x_0)$, and $K = \mathrm{Mod}(T^*_{\mathcal{M}_0})$. We let K' denote the class of all structures of the form $(\mathcal{M}, a, x)$, where $\mathcal{M} \in \mathrm{Mod}(T_{\mathcal{M}_0})$, $a \in \mathsf{S}(\mathcal{M})$ and $x \in \mathsf{S}_\mathfrak{f}(\mathcal{M})$. It is clear that $K \subseteq K'$, and that K' is an elementary class.

CLAIM E: *Let $h : (\mathcal{M}, a, x) \to (\mathcal{M}', a', x')$ be an elementary map in K. Then* $\mathsf{S}_\mathfrak{f}(h)(x) = \mathfrak{f}(h|\mathscr{L})(x) = x'$

PROOF: This is a consequence of our assumption that a_0 dominates x_0. Composing h with an elementary embedding $g : (\mathcal{M}_0, a_0, x_0) \to (\mathcal{M}, a, x)$ gives

$$\mathsf{S}_\mathfrak{f}(h)(x) = \mathsf{S}_\mathfrak{f}(h)(g(x_0)) = \mathsf{S}_\mathfrak{f}(h \circ g)(x_0)$$

and

$$\mathfrak{f}(h|\mathscr{L})(x) = \mathfrak{f}(h|\mathscr{L})(g(x_0)) = \mathfrak{f}(h|\mathscr{L} \circ g|\mathscr{L})(x_0).$$

Since a_0 dominates x_0, we get $\mathfrak{f}(h|\mathscr{L} \circ g|\mathscr{L})(x_0) = \mathsf{S}_\mathfrak{f}(h \circ g)(x_0)$, so that $\mathsf{S}_\mathfrak{f}(h)(x) = \mathfrak{f}(h|\mathscr{L})(x)$ □

CLAIM F: *The class K satisfies the additional property required in 3.4.11, i.e. that for every set I, and any ultrafilter U on I, if $\prod_U(\mathcal{M}_i, a_i, x_i) \in K$, then there is $P \in U$ such that for every $i \in P$, $(\mathcal{M}_i, a_i, x_i) \in K$.*

Claim finishes the proof of the theorem. For clarity, we delay its proof until after claim G. By claim 40, $\mathrm{Th}(K)$ consists of $T_{\mathcal{M}_0}$ and a single formula α with domain $\mathsf{S} \times \mathsf{S}_\mathfrak{f}$. Then $\mathcal{M}_0 \models \alpha(a_0, x_0)$. If $\mathcal{M}_0 \models \alpha(a_0, y)$, then $x_0 \equiv_{a_0} y$. Suppose $d(x_0, y) \geq \varepsilon > 0$, then since $x_0 \equiv_{a_0} y$, there is an ultrafilter U on a set I and a map $h : \mathcal{M}_0 \to \mathcal{M}_0^U$ such that $h(x_0) = y$, and $h(a_0) = \langle a_0 \rangle_U$. Comparing h to the canonical embedding $\partial : \mathcal{M}_0 \to \mathcal{M}_0^U$ contradicts the fact that a_0 dominates x_0.

CLAIM G: *Suppose $(\mathcal{N}, a, x) \notin K$. Then there is $\varepsilon > 0$ such that for every $x' \in \mathsf{S}_\mathfrak{f}(\mathcal{N})$ such that $(\mathcal{N}, a, x') \in K$, we have $d(x', x) \geq \varepsilon$.*

PROOF (OF CLAIM G): Suppose not, then for every n, there is $x_n \in \mathsf{S}_\mathfrak{f}(\mathcal{N})$ such that $d(x_n, x) \leq 1/2^n$, and $(\mathcal{N}, a, x_n) \in K$. Note that all the structures $(\mathcal{N}, a, x_n)$ and $(\mathcal{N}, a, x)$ have identical underlying sets. Therefore, $x_n \to x$ in $\mathsf{S}_\mathfrak{f}(\mathcal{N})$. Let $(\mathbb{N}, U)$ be an ultrafilter pair, and consider the ultraproduct $\mathcal{M} = \prod_U(\mathcal{N}, a, x_n)$. We can embed $(\mathcal{N}, a, x)$ into $\mathcal{M}$ by sending x to $\langle x_n \rangle_U$. Since K is elementary, this means $(\mathcal{N}, a, x) \in K$, contradicting our assumption. □

We now prove claim 40. Let $\varepsilon > 0$ be the number given in claim G, and assume that the conclusion of the claim fails. Let $\Delta_1 : I \to K'$ be an ultradiagram such that $\Delta(f(P)) \notin K$ for $P \in U$, but nevertheless $\prod_U \Delta_1(i) \in K$. Write $\Delta_1(i) = (\mathcal{M}_i, a_i, x_i)$, with

$\mathcal{M}_i \in \mathrm{Mod}(T_{\mathcal{M}_o})$, $a_i \in \mathsf{S}(\mathcal{M}_i)$ and $x_i \in \mathsf{S}(\mathcal{M}_i)$. By definition of K', we have an elementary map $h_i : \mathcal{M}_0 \to \mathcal{M}_i$ for every $i \in I$. Since for every $i \in I$, $a_i \in \mathsf{S}(\mathcal{M}_i)$ belongs to a sort of $\mathscr{L}$, we have $h_i(a_0) = a_i$. Note also that if it happens to be the case that $(\mathcal{M}_i, a_i, x_i) \in K$, then $x_i = \mathfrak{f}(h_i)(x_0)$, and that $d(x_{f(P)}, \mathfrak{f}(h_{f(P)})(x_0)) \geq \varepsilon$. Define $\Delta_2(i) = (\mathcal{M}_i, a_i, \mathfrak{f}(h_i)(x_0))$. For $\epsilon = 1, 2$, we consider the maps:

$$\partial_{\Delta_\epsilon} : \prod_U \Delta_\epsilon(i) \to \prod_{f^{-1}[U]} \Delta_\epsilon(f(P)).$$

By the definition of Δ_2, we get an elementary map $h_2 : (\mathcal{M}_0, a_0, x_0) \to \prod_U \Delta_2(i)$ defined via the composition

$$\mathcal{M}_0 \xrightarrow{\partial_{\mathcal{M}_0}} \mathcal{M}_0^U \xrightarrow{\prod_U h_i} \textstyle\prod_U \mathcal{M}_i .$$

Since $\mathfrak{f}$ is an ultrafunctor, we have:

$$\mathfrak{f}(\mathfrak{h}_2|\mathscr{L}) = \mathfrak{f}(\prod_U h_i \circ \partial_{\mathcal{M}_o}) = \mathfrak{f}(\prod_U h_i) \circ \mathfrak{f}(\partial_{\mathcal{M}_o}) = \prod_U \mathfrak{f}(h_i) \circ \partial_{\mathfrak{f}\circ\mathcal{M}_o}$$

so that $\mathfrak{f}(\mathfrak{h}_2|\mathscr{L})(x_0) = \langle \mathfrak{f}(h_i|\mathscr{L})(x_0)\rangle_U$. Also, by assumption, $\prod_U \Delta_1(i) \in K$, so there is an elementary map $h_1 : (\mathcal{M}_0, a_0, x_0) \to \prod_U \Delta_1(i)$ such that $h_1(a_0) = \langle a_i \rangle_U$ and $\mathfrak{f}(h_1|\mathscr{L})(x_0) = \langle x_i \rangle_U$ by claim E. A key observation here is $h_1(a_0) = h_2(a_0)$, which is true because a_0 belongs to one of the original sorts of $\mathscr{L}$. The relationship between $\mathfrak{f}(h_1|\mathscr{L})(x_0)$ and $\mathfrak{f}(h_2|\mathscr{L})(x_0)$ is yet to be determined. Since $h_1(a_0) = h_2(a_0)$, we have

$$\partial_{\Delta_1} \circ h_1(a_0) = \partial_{\Delta_2} \circ h_2(a_0).$$

On the other hand, we have the following calculation, in which the definition of ultrafunctor plays a central rôle:

$$\begin{aligned} d(\mathfrak{f}(\partial_{\Delta_1} \circ h_1)(x_0), \mathfrak{f}(\partial_{\Delta_2} \circ h_2)(x_0)) &= d(\mathfrak{f}(\partial_{\Delta_1})(\mathfrak{f}(h_1)(x_0)), \mathfrak{f}(\partial_{\Delta_2})(\mathfrak{f}(h_2)(x_0))) \\ &= d(\partial_{\mathfrak{f}\circ\Delta_1}(\langle x_i\rangle_U), \partial_{\mathfrak{f}\circ\Delta_2}(\langle \mathfrak{f}(h_i|\mathscr{L})(x_0)\rangle_U)) \\ &= d(\langle x_{f(P)}\rangle_{f^{-1}[U]}, \langle \mathfrak{f}(h_{f(P)}|\mathscr{L})(x_0)\rangle_{f^{-1}[U]})) \\ &\geq \varepsilon \end{aligned}$$

Therefore, the two maps

$$\mathcal{M}_0 \underset{h_2}{\overset{h_1}{\rightrightarrows}} \textstyle\prod_U \mathcal{M}_i$$

bear witness to the fact that a_0 does not dominate x_0, contradicting the assumption of the theorem. This proves the claim, and finishes the proof of the theorem. □

4.3.7 Theorem: *The class $K = \{\mathcal{N} : \mathcal{N} \cong \mathcal{M}_{\mathfrak{f}}$ for some $\mathcal{M} \in \mathrm{Mod}(T)\}$ is elementary. There is a functor $E : \mathrm{Mod}(T) \to \mathrm{Mod}(\mathrm{Th}(K))$ whose action on objects is given by $\mathcal{M} \mapsto \mathcal{M}_{\mathfrak{f}}$ and whose action on maps is given by $h \mapsto h_{\mathfrak{f}}$. The functor E is an ultra-equivalence of categories, and its inverse is the forgetful functor $F : \mathrm{Mod}(\mathrm{Th}(K)) \to \mathrm{Mod}(T)$*

Proof: The class K is closed under ultraproducts by virtue of $\mathfrak{f}$ being an ultrafunctor, and under isomorphisms by definition. To show that it is closed under elementary substructures, let $\mathcal{M} \in K$, and $\mathcal{N} \preceq \mathcal{M}$. By the definition of K, $\mathsf{S}_{\mathfrak{f}}(\mathcal{M}) \cong \mathrm{dom}(\mathfrak{f})(\mathcal{M}^\ell)$, where $\mathcal{M}|\mathscr{L}$

denotes the $\mathscr{L}$-reduct of $\mathcal{M}$. We need to show that $\mathsf{S}_{\mathfrak{f}}(\mathcal{N}) \cong \mathrm{dom}(\mathfrak{f})(\mathcal{N}^{\ell})$. Denote by h the elementary embedding $\mathcal{N} \to \mathcal{M}$. Let $x \in \mathsf{S}_{\mathfrak{f}}(\mathcal{N})$. By Theorem 4.3.6, there is a sort S of $\mathscr{L}$ and a definable map $\alpha : \mathsf{S} \to \mathsf{S}_{\mathfrak{f}}$ such that for some $y \in \mathsf{S}(\mathcal{N})$, $\mathcal{N} \models [\alpha(y) = x]$. Let p be $\mathrm{tp}(x/\mathcal{N}|\mathscr{L})$. We claim that p has a realization in the structure $(\mathcal{N}|\mathscr{L}, \mathrm{dom}(\mathfrak{f})(\mathcal{M}))$. Indeed, p is determined by the definable map α, and since $h : \mathcal{N} \to \mathcal{M}$ is elementary, $\mathcal{M} \models [\alpha(h(y)) = h(x)]$. By Theorem 4.3.3, the map $(h|\mathscr{L}, \mathfrak{f}(h)) : \mathcal{N}|\mathscr{L} \to \mathcal{M}$ is an elementary map, and $\mathcal{M} \models \exists x'[\alpha(x) = x']$, a formula with parameters from $\mathcal{N}|\mathscr{L}$. Therefore, there is $x' \in \mathrm{dom}(\mathfrak{f})(\mathcal{N})$ realizing p. Let $f = [x \mapsto x']$, and note that this map is well defined. We claim that it is an elementary map. To see this, let $x_1 x_1 \cdots x_n \in \mathsf{S}_{\mathfrak{f}}(\mathcal{N})$. By applying Theorem 4.3.6 to the whole tuple $x_1 \cdots x_n$, we get a tuple $y_1 \cdots y_n$ of elements of $\mathrm{dom}(\mathfrak{f})(\mathcal{N})$ which realizes $\mathrm{tp}(x_1 \cdots x_n/\mathcal{N}|\mathscr{L})$. Now we note that $y_i \equiv_{\mathcal{N}|\mathscr{L}} f(x_i)$ for $1 \leq i \leq n$, and therefore $y_i = f(x_i)$, showing that f is elementary.

The fact that E is an ultrafunctor follows directly from the fact that $\mathfrak{f}$ is an ultrafunctor. It is also easy to verify that E and F are mutual inverses. The proof is similar to the proof that K is elementary. □

We now conclude the proof of Theorem 4.2.22. Recall the following definitions and notations:

1. $\mathfrak{f}$ is a fixed ultrafunctor;
2. $\mathscr{L}_{\mathfrak{f}}$ is the expansion of $\mathscr{L}$ by a new sort symbol $\mathsf{S}_{\mathfrak{f}}$ and new predicate symbols $\mathsf{P}_{\mathfrak{g}}$ for every ultrafunctor $\mathfrak{g}$ such that $\mathrm{dom}(\mathfrak{g}) = \mathsf{S} \times \mathrm{dom}(\mathfrak{f})$, where S is a sort of $\mathscr{L}$.
3. $\mathcal{M}_{\mathfrak{f}}$ is the expansion of $\mathcal{M} \in \mathrm{Mod}(T)$ to an $\mathscr{L}_{\mathfrak{f}}$-structure obtained by letting $\mathsf{S}_{\mathfrak{f}}(\mathcal{M}_{\mathfrak{f}}) = \mathrm{dom}(\mathfrak{f})(\mathcal{M})$ and $[\![\mathsf{P}_{\mathfrak{g}}]\!]_{\mathcal{M}_{\mathfrak{f}}}(x) = \mathfrak{g}(\mathcal{M})(x)$
4. $K = \{\mathcal{N} : \mathcal{N} \cong \mathcal{M}_{\mathfrak{f}}$ for some $\mathcal{M} \in \mathrm{Mod}(T)\}$. It is an elementary class by Theorem 4.3.7, and is equivalent to $\mathrm{Mod}(T)$ via the ultra-equivalence E.

We restate the conceptual completeness theorem with minor changes to accommodate the categorical setup:

4.3.8 Theorem (Conceptual completeness, [Har]): *Let $\mathscr{L}$ be a continuous language, and let $\mathscr{L}'$ expand $\mathscr{L}$. Let T' be a complete $\mathscr{L}'$-theory, and let T be the $\mathscr{L}$-reduct of T'. Suppose that the forgetful functor $F : \mathrm{Mod}(T') \to \mathrm{Mod}(T)$ is an equivalence of categories. Then for every formula ψ of $\mathscr{L}'$, there is a formula $\varphi \in \mathscr{L}^{eq}$ and a monomorphism $\eta : \mathfrak{ev}_{\psi} \to \mathfrak{ev}_{\varphi}$.*

We apply this theorem to $\mathscr{L}' = \mathscr{L}_{\mathfrak{f}}$ and $T' = \mathrm{Th}(K)$ to get an $\mathscr{L}^{eq}$ formula φ such that $\mathfrak{ev}_{\mathsf{P}_{\mathfrak{f}}} = \mathfrak{ev}_{\varphi}$. By the definition of $\mathscr{L}_{\mathfrak{f}}$, $\mathfrak{ev}_{\mathsf{P}_{\mathfrak{f}}} = \mathfrak{f}$. This completes the proof.

Chapter 5

Simple theories in continuous logic

In this chapter we develop simplicity theory in the context of first-order continuous logic. We follow the more axiomatic approach to simplicity theory in [Cas07]. Continuous theories can be construed as a special case of the concept of a compact abstract theory, for which there is already a well-established development of simplicity theory (see [Ben03b, Ben03a]). However, the greater abstraction of general cats prohibits some of the classical results of first-order simplicity theory to carry over. In a general cat, for example, having the same type may not be a type-definable condition. Consequently, types do not always have non-dividing extensions, and Morley sequences may not exist in every type.

Continuous theories correspond to the concept of a Hausdorff cat. There, simplicity can be developed in all its glory. In a Hausdorff cat, indiscernible sequences are type-definable. This fact alone allows to prove many of the results of simplicity theory, including the existence of non-dividing extensions and Morley sequences. The development of simplicity in theories where indiscernibility is type-definable is done in [Ben03b].

To the knowledge of the author, no such development exists which is written specifically in continuous logic. This chapter is an attempt to remedy this situation. We will follow [Cas07] quite closely, pointing out the places where proofs should be adapted to the context of continuous logic. In the cases where the proofs for continuous logic would only be a minor modification of the classical proof, they will be omitted, unless the similarity between the two proofs needs to be highlighted.

For this chapter we fix a metric $\mathfrak{R}$-valued language $\mathscr{L}$. We let T be a complete metric $\mathscr{L}$-theory with infinite models, and we let $\mathcal{C} \in \mathrm{Mod}(T)$ be very large, κ-saturated and κ-universal, where κ is a big cardinal. We take as a convention that all parameters come from $\mathcal{C}$. A set $X \subseteq \mathcal{C}$ will be called *small* if $|X| < |\mathcal{C}|$. Note that every small model of T can be elementarily embedded in $\mathcal{C}$. For convenience, we shall not make any notational distinction between elements of $\mathcal{C}$ and tuples of elements of $\mathcal{C}$ of small length. We allow for the possibility that a tuple of small length be an enumeration of a small model of T.

Section 5.1

Indiscernible sequences

We begin by recalling the notion of indiscernible sequence for metric structures.

5.1.1 Definition: Let $\mathcal{X}$ be an $\mathscr{L}$-structure and $\varepsilon > 0$. We say $\mathcal{X}$ is *ε-finite* in case there does not exist an infinite set $X \subseteq \mathcal{X}$ such that $[\![\mathsf{d}]\!]_{\mathcal{X}}(x, y) \geq \varepsilon$ for every $x, y \in X$. In other words, $\mathcal{X}$ is ε-finite if and only if given any infinite $X \subseteq \mathcal{X}$ one can find $x, y \in X$ such that

$[\![\mathsf{d}]\!]_{\mathcal{X}}(x,y) < \varepsilon$. If $\mathcal{X}$ is not ε-finite, then it is called *ε-infinite*. A structure is *finite* if and only if it is ε-finite for every ε, and infinite if it is ε-infinite for some $\varepsilon > 0$. ♣

5.1.2 Definition: Let $\mathcal{X}$ be an $\mathscr{L}$-structure, and $A \subseteq \mathcal{X}$ be a set. A sequence $I = (a_i : i < \omega) \subseteq \mathcal{X}$ is *A-indiscernible* if and only if for any two tuples $a, b \in I$ of the same order type, we have $a \equiv_A b$. ♣

It is worth pointing out there that the relation $a \equiv_A b$ is type-definable. As in first-order logic, it is defined by the type $\{|\varphi(x) - \varphi(y)| : \varphi \in \mathscr{L}\}$, where x and y are tuples of variables representing as a and b. In [Ben03b], it is pointed out that this property allowing to derive many of the results of simplicity theory.

5.1.3 Theorem: *Let $\kappa > |T|$, and let $\lambda = (2^{2^\kappa})^+$, $|A| \leq \kappa$, and $(a_i : i < \lambda)$ is a sequence of sequences such that $|a_i| = \alpha < \kappa^+$. Then there is an A-indiscernible sequence $(b_i : i < \omega)$ such that for each $n < \omega$, there is an increasing function $f : n \to \lambda$ such that $b_1, ..., b_n \equiv_A a_{f(0)}, ..., a_{f(n)}$.*

Proof: See [Ben03a] for a proof in the context of cats. The proof for continuous logic is identical. □

Section 5.2

Dividing

We give a direct definition of dividing in continuous logic. As we shall see, the definition is quite similar to the usual definition in first-order logic. The main difficulty in the formulation of dividing comes from the fact that continuous logic lacks a negation operator. In order to obtain the behaviour of $\neg\varphi$, we must exhibit $\varepsilon > 0$ such that $\varphi \geq \varepsilon$. As we saw in section 5.1, this impacts the definition of "infinite", which must be written as ε-infinite. It also impacts the definition of "inconsistent": a syntactic notion of inconsistency for a set of formulae in continuous logic looks like this:

5.2.1 Definition: Let $\varepsilon > 0$, and let Δ be a finite set of formulae. Then Δ is *ε-inconsistent* if and only if

$$\mathcal{C} \models \forall x \left[\left(\bigwedge_{\varphi \in \Delta} \right) \varphi(x) \geq \varepsilon \right].$$

An infinite set Σ is k, ε-inconsistent if and only if every subset $\Delta \subseteq \Sigma$ such that $|\Delta| = k$ is ε-inconsistent. ♣

We note that the extra parameter ε in the definition of "inconsistent" can be done away with. Semantically, a set of formulae $\Sigma(x)$ is inconsistent if and only if no structure $\mathcal{M}$ realizes it. This in fact *forces* the existence of $\varepsilon > 0$ such that $\Sigma(x)$ is ε-inconsistent, for suppose such an ε did not exist. This means that for every $n \in \mathbb{N}$, there is a structure $\mathcal{M}_n$ such that

$$\mathcal{M}_n \models \exists x[\varphi_1(x) \wedge \cdots \wedge \varphi_k(x) \leq 1/n]$$

for every finite $\{\varphi_1, ..., \varphi_k\} \subseteq \Sigma(x)$, and a suitable ultraproduct of the $\mathcal{M}_n$'s would produce a structure $\mathcal{M}$ in which $\Sigma(x)$ is realized. However, the use of the parameter ε cannot be avoided in the definition of k, ε-inconsistent: for every $\Delta \subseteq \Sigma(x)$ of size k, if we assume that Δ is inconsistent, then there is $\varepsilon_\Delta > 0$ such that Δ is ε_Δ-inconsistent. The set $\{\varepsilon_\Delta : \Delta \subseteq \Sigma(x), |\Delta| = k\}$, however, may not be bounded away from zero.

5.2.2 Definition: Let $\pi(x, a)$ be a partial type. We will say that π *(k, ε, δ)-divides over A* if and only if there is a δ-infinite sequence $(b_i : i < \omega)$ such that

1. $b_i \equiv_A a$ for every $i < \omega$, and
2. The set $\bigcup_{i \leq \omega} \pi(x, b_i)$ is k, ε-inconsistent.

We will say that π *k-divides* over A if there are $\varepsilon, \delta > 0$ such that $\pi(x, a)$ k, ε, δ-divides, and that π *divides over A* if and only if there is k such that π k-divides over A. ♣

5.2.3 Proposition: *A partial type $\pi(x, a)$ divides over A if and only if there is an infinite A-indiscernible sequence $(b_i : i < \omega)$ in $\mathrm{tp}(a/A)$ such that the set $\bigcup_{i<\omega} \pi(x, b_i)$ is inconsistent.*

Proof: If $\pi(x, a)$ divides, then an infinite indiscernible sequence exists by Theorem 5.1.3. To see this, note that by compactness, there is a sequence $J = (a_j : j < \lambda)$ witnessing dividing of length $\lambda = (2^{2^\kappa})^+$, with $\kappa = |T|^+$. By Theorem 5.1.3, there is an indiscernible sequence $(b_i : i \in \omega)$ with the property that for every n, there is an increasing function $f : n \to \lambda$ such that $b_1 \cdots b_n \equiv_A a_{f(0)} \cdots a_{f(n)}$. This property guarantees that $(b_i : i \in \omega)$ also witnesses dividing for $\varphi(x, a)$.

For the converse, let $(b_i : i < \omega)$ be an infinite A-indiscernible sequence in the type $\mathrm{tp}(a/A)$, and suppose $\Pi = \bigcup_{i<\omega} \pi(x, b_i)$ is inconsistent. By the approximate compactness theorem, there is $\varepsilon > 0$ and a finite subset $\{\varphi_1, ..., \varphi_n\} \subseteq \Pi$ such that

$$\mathcal{C} \models \forall x[(\varphi_1(x, b_1) \wedge \cdots \wedge \varphi_n(x, b_n)) \geq \varepsilon].$$

Then $\pi \models \varphi_1(x, a) \wedge \cdots \wedge \varphi_n(x, a)$. We claim that this formula divides. Let B be the set of all n-tuples of elements of $(b_i : i < \omega)$ whose order type is that of $(b_1, ..., b_n)$. Since $(b_i : i < \omega)$ is indiscernible, it is δ-infinite for some δ, and therefore B is δ-infinite as well. Since $\mathcal{C} \models \forall x[\varphi_1(x, b_1) \wedge \cdots \wedge \varphi_n(x, b_n) \geq \varepsilon]$, we get that the set

$$\Sigma = \{\varphi_1(x, b_1) \wedge \cdots \wedge \varphi_n(x, b_n) : (b_1, ..., b_n) \in B\}$$

is ε-inconsistent. Therefore, there is a finite $\Delta \subseteq \Sigma$ which is ε-inconsistent. Let $k = |\Delta|$. By definition, $\varphi_1(x, a) \wedge \cdots \wedge \varphi_n(x, a)$ k-divides over A. This completes the proof. □

5.2.4 Corollary: *A partial type $\pi(x, a)$ divides over A if and only if there is a finite conjunction of formulae in π which divides over A.*

Proof: The formulae $\varphi_i(x, y)$ exhibited in the proof of proposition 5.2.3 are in $\pi(x, y)$. □

We begin by establishing some results concerning the behaviour of non-dividing with respect to the extra structure on types and formulae that is present in continuous logic but not in classical first order logic. Let x be a tuple of variables, and let $\Sigma(x)$ be the set of all the formulae of $\mathscr{L}$ in the free-variable x. Given the large model $\mathcal{C}$, we can define a pseudo-metric space structure on $\Sigma(x)$ as follows:

$$d_{\mathcal{C}}(\varphi, \psi) = \sup_{x \in \mathcal{C}} [\![\varphi - \psi]\!]_{\mathcal{C}}(x).$$

Note that $d_\mathcal{C}(\varphi, \psi) = 0$ if and only if φ and ψ are logically equivalent. Therefore, $d_\mathcal{C}$ induces a metric on the equivalence classes of $\Sigma(x)$ modulo logical equivalence. When we refer to a metric space of the form $(\Sigma(x), d_\mathcal{C})$, we will be referring to $(\Sigma(x)/\sim, d_\mathcal{C})$, where $\sim$ is logical equivalence.

5.2.5 Lemma: *Let $\varphi(x, a)$ be a formula. If $\varphi(x, a)$ divides over A, then there is $\varepsilon > 0$ such that for every formula ψ, if $d_\mathcal{C}(\varphi, \psi) \leq \varepsilon$, then $\psi(x, a)$ divides over A as well.*

Proof: Suppose φ divides over A, and let I be an indiscernible sequence such that $\Sigma = \{\varphi(x, c) : c \in I\}$ is inconsistent. Let $\delta > 0$ and $c_1, ..., c_n \in C$ be such that

$$\mathcal{C} \models \forall x[\varphi(x, c_1) \wedge \cdots \wedge \varphi(x, c_n) \geq \delta].$$

Let $\varepsilon = \delta/3$, and let ψ be such that $\mathcal{C} \models \forall x[|\psi(x, a) - \varphi(x, a)| \leq \varepsilon]$. We show that $\{\psi(x, c) : c \in C\}$ is inconsistent. Let $x \in \mathcal{C}$, and i be such that $\varphi(x, c_i) \geq \delta$. Since $|\varphi(x, c_i) - \psi(x, c_i)| \leq \delta/3$, is must be the case that $\psi(x, c_i) \geq \delta/2$. Therefore

$$\mathcal{C} \models \forall x[\psi(x, c_1) \wedge \cdots \wedge \psi(x, c_n) \geq \delta/2]$$

showing that $\{\psi(x, c) : c \in C\}$ is inconsistent. □

5.2.6 Corollary: *If $\varphi_n(x, a)$ does not divide over A, and $\varphi_n \to \varphi$ in $\Sigma(x, a)$, then $\varphi(x, a)$ does not divide over A. Consequently, the set*

$$\{\varphi(x, a) : \varphi \textit{ does not divide over } A\}$$

is closed in the metric space $(\Sigma(x, a), d_\mathcal{C})$.

5.2.7 Definition: Let $a, b, c \in \mathcal{C}$ be small tuples. We write $a \underset{c}{\perp} b$ if and only if $\mathrm{tp}(a/bc)$ does not divide over c. ♣

Note that if $a, b, c \in \mathcal{C}$ are small tuples, then $a \underset{c}{\perp} b$ if and only if for every finite sub-tuple $a' \subseteq a$, $a' \underset{c}{\perp} b$. This requires a bit more of an argument in continuous logic, since we are allowing for infinitary connectives. It is clear that if $a \underset{c}{\perp} b$ then for every finite sub-tuple $a' \subseteq a$, $a' \underset{c}{\perp} b$. For the converse, assuming that no formula in finitely many variables divides, we need to show that no formula $\varphi \in \mathrm{tp}(a/bc)$ with infinitely many free variables divides. By definition, a formula φ in infinitely many free variables can be written as $u(\varphi_1, ..., \varphi_n, ...)$, where u is an infinitary connective, and φ_i is a formula in finitely many variables, and therefore is a limit $\lim_{n\to\infty} u_k(\varphi_1, ..., \varphi_k)$. Note that $u_k(\varphi_1, ..., \varphi_k)$ is a formula in finitely many free-variables, and we are assuming those formulae do not divide. Therefore, their limit does not divide either by Lemma 5.2.5.

We also have the following properties whose proofs are immediate from the definition. The following is stated as Remark 5.1 in [Cas07]. Let $a, b, c \in \mathcal{C}$ be tuples of small length:

1. For any automorphism $\sigma : \mathcal{C} \to \mathcal{C}$, $a \underset{c}{\perp} b$ if and only if $\sigma(a) \underset{\sigma(c)}{\perp} \sigma(b)$

2. $a \underset{c}{\perp} b$ if and only if $a \underset{c}{\perp} bc$

3. If $a \underset{c}{\perp} b$, and $b' \subseteq b$, then $a \underset{cb'}{\perp} b$

4. $a \underset{c}{\perp} b$ if and only if for every finite $a' \subseteq a$ and $b' \subseteq b$, $a' \underset{c}{\perp} b'$

If π is a partial type, we write $\pi \leq \varepsilon$ for the type $\{\varphi \leq \varepsilon : \varphi \in \pi\}$. The proof of the following proposition is identical to its classical counterpart, and can be found in [Wag00].

5.2.8 Proposition: *Let $\varphi(x, a)$ and $\psi(x, a)$ be formulae and suppose $\varphi(x, a) \models \psi(x, a)$. If $\psi(x, a)$ divides, then so does $\varphi(x, a)$.*

Without further ado, we prove the finite character of non-dividing in continuous logic. The proof is actually identical to the first order proof.

5.2.9 Proposition (Lemma 2.10 of [She80]): *The following statements are equivalent:*

1. $\mathrm{tp}(a/Ab)$ *does not divide over* A
2. *For every A-indiscernible sequence I such that $b \in I$, there is $a' \equiv_A a$ such that I is Aa'-indiscernible*
3. *For every A-indiscernible sequence I, there is $J \equiv_{Ab} I$ such that J is Aa-indiscernible.*

The following propositions follows immediately from 5.2.9

5.2.10 Proposition: *Let $A \subseteq B$ be small subsets of $\mathcal{C}$. If $a \underset{A}{\perp} B$ and $b \underset{Aa}{\perp} B$, then $ab \underset{A}{\perp} B$.*

5.2.11 Proposition: *If $\varphi(x, a)$ k-divides over A, and* $\mathrm{tp}(b/Aa)$ *does not divide over A, then $\varphi(x, a)$ k-divides over Ab.*

5.2.12 Definition: A theory T is called *simple* if and only if, given a very large saturated model $\mathcal{C}$ of T, for every $a, b \in \mathcal{C}$, there is a subset $c \subseteq b$ of length $\leq |T|$ such that $a \underset{c}{\perp} b$. This property is called *local character.* ♣

In [Ben03a, Introduction], it is pointed out that a fair amount of simplicity theory can be carried out using only local character, at the cost of having to write more technical proofs. Local character is enough to prove symmetry, transitivity and the characterization of dividing via the local D-ranks.

The following proposition establishes the behaviour of dividing with respect to the metric space structure on the type space $S_{\mathcal{C}}(B)$ for a small set $B \subseteq \mathcal{C}$. Recall that the metric d_S on $S_{\mathcal{C}}(B)$ is defined by

$$d(p,q) = \inf \{ [\![\mathsf{d}]\!]_{\mathcal{C}} (a,b) : a \models p, b \models q \}$$

where d is the symbol representing the metric on the domain sort of $S_{\mathcal{C}}(B)$.

5.2.13 Proposition: *Let $A \subseteq B \subseteq \mathcal{C}$ be small sets, and let $p = \mathrm{tp}(a/B)$. Suppose p divides over A. There is $\varepsilon > 0$ such that if $d(a,b) \leq \varepsilon$, then $\mathrm{tp}(b/B)$ divides over A as well. In other words, if p divides, then there is $\varepsilon > 0$ such that for every q, if $d_S(p,q) \leq \varepsilon$, then q divides as well.*

Proof: There is ε such that $p \leq \varepsilon$ divides. Using Lemma 5.2.4, choose $\psi(x,b) \in p \leq \varepsilon$ such that $\psi(x,b)$ divides over A. By the definition of $p \leq \varepsilon$, ψ is of the form $\varphi(x,b) \leq \varepsilon$ for $\varphi \in p$.

Since φ is uniformly continuous, there is δ such that $d(a,c) \leq \delta$ implies $|\varphi(a,b) - \varphi(c,b)| \leq \varepsilon$. Since $\varphi \in p$, $\mathcal{C} \models \varphi(a,b) = 0$, so that it must be the case that $\varphi(c,b) \leq \varepsilon$. Therefore, "$\varphi(x,b) \leq \varepsilon$" $\in \mathrm{tp}(c/B)$ for every $c \in B_\delta(a)$, and it divides. Therefore, for every $c \in B_\delta(a)$, $\mathrm{tp}(c/B)$ divides. □

Section 5.3

The D-ranks and the tree property

A discussion of simple theories would not be complete without a mention of the tree property and of the D-ranks. In this section we also show that a partial type has a *non-dividing* extension to a complete type over any set. Consequently, a complete type $p \in S_{\mathcal{C}}(A)$ has a non-dividing extension to any small $B \supseteq A$. We use the approach of [Ben03b, section 1] to avoid mentioning forking.

The definition of the tree property is identical in the continuous context as it is in the first-order context; note the presence of the extra parameter ε in the definition. It is needed in the definition of "inconsistent"

5.3.1 Definition: We say that a formula $\varphi(x,y)$ has the *tree property* with respect to $k < \omega$, $\varepsilon > 0$ and $\delta > 0$ if and only if there is a tree $(a_\eta : \eta \in \omega^{<\omega})$ such that for every branch η, the set $\pi(x,a) \cup \{\varphi(x, a_{\eta|n}) : n \in \omega\}$ is consistent, and for every node $s \in \omega^{<\omega}$, the set $\{a_{s\frown j} : j \in \omega\}$ is δ-infinite, and the set $\{\varphi(x, a_{s\frown i}) : i \in \omega\}$ is k, ε-inconsistent. ♣

The proof of the following lemma in the context of continuous logic is identical to the proof given in [Cas07].

5.3.2 Lemma: *Let α be an ordinal number, $\pi(x,a)$ be a partial type, and $\varphi_i(x,y_i)$ be a formula for $i < \alpha$. Let $(k_i : i < \alpha)$ be a sequence of natural numbers, and $(\varepsilon_i : i < \alpha)$ and $(\delta_i : i < \alpha)$ be sequences of real numbers in $\mathfrak{R}$. Then the following two statements are equivalent:*

1. *There is a tree $(a_\eta : \eta \in \omega^{<\alpha})$ and $\delta > 0$ such that for every branch $\eta \in \omega^\alpha$, the set $\pi(x,a) \cup \{\varphi(x, a_{\eta|i}) : i < \alpha\}$ is consistent, and for every node $s \in \omega^{<\alpha}$, the set $\{a_{s\frown i} : i \in \alpha\}$ is δ-infinite, and $\{\varphi(x, a_{s\frown i}) : i \in \alpha\}$ is k_i, ε_i-inconsistent.*

2. *There is a sequence $(a_i : i < \alpha)$ such that $\pi(x,a) \cup \{\varphi_i(x,a_i) : i : \alpha\}$ is consistent and for every $i > \alpha$, $\varphi_i(x,a_i)$ $(k_i, \varepsilon_i, \delta_i)$-divides over $\{a\} \cup \{a_j : j < i\}$*

5.3.3 Definition: Let $(k_i : i < \alpha)$ be a sequence of natural numbers, and $(\varepsilon_i : i < \alpha)$ and $(\delta_i : i < \alpha)$ be bounded sequences of real numbers. A *dividing chain with respect to* $(k_i : i < \alpha)$, $(\varepsilon_i : i < \alpha)$ *and* $(\delta_i : i < \alpha)$ for $\varphi(x,y)$ is a sequence $(a_i : i < \alpha)$ such that $\pi(x,a) \cup \{\varphi(x,a_i) : i < \alpha\}$ is consistent, and for every $i < \alpha$, $\varphi(x,a_i)$ $(k_i, \varepsilon_i, \delta_i)$-divides over $\{a_j : j < i\}$. We say that φ *divides α times* with respect to $(k_i : i < \alpha)$, $(\varepsilon_i : i < \alpha)$ and $(\delta_i : i < \alpha)$ if there is a dividing chain of length α with respect to $(k_i : i < \alpha)$, $(\varepsilon_i : i < \alpha)$ and $(\delta_i : i < \alpha)$. ♣

5.3.4 Lemma: *1. $\varphi(x,y)$ divides ω times with respect to k, ε and δ if and only if it has the tree property with respect to k, ε and δ;*

2. *The following two statements are equivalent:*

   a) *For every n, $\varphi(x,y)$ divides n times with respect to k, ε and δ*

   b) *For every ordinal α, $\varphi(x,y)$ divides α times with respect to k, ε and δ.*

3. *If φ divides ω_1 times with respect to some sequence $(k_i : i < \omega_1)$, $(\varepsilon_i : i < \omega_1)$ and $(\delta_i : i < \omega_1)$, then there are $k \in \mathbb{N}$, $\varepsilon > 0$ and $\delta > 0$ such that φ divides ω times with respect to k, ε and δ.*

Proof: The first statement follows directly from the definition. The second item follows from the definition and compactness. For the third item, if $\varphi(x,a)$ divides ω_1 times with respect to the sequences $(k_i : i < \omega_1)$, $(\varepsilon_i : i < \omega_1)$ and $(\delta_i : i < \omega_1)$, then there is $i_0 < \omega_1$ such that $\{j < \omega_1 : k_j = k_{i_0}\}$ has size ω_1. This implies that φ divides ω_1 times with respect to k_{i_0} and the sequences $(\varepsilon_i : i < \omega_1)$ and $(\delta_i : i < \omega_1)$. For every $i \in \omega_1$, let n_i and m_i be such that $1/n_i < \varepsilon_i$ and $1/m_i < \delta_i$, and note that $\varphi(x,a)$ divides ω_1 times with respect to the sequences k_{i_0} and $(1/n_i : i < \omega_1)$ and $(1/m_i : i < \omega_1)$. There is $j_0 < \omega_1$ such that $\{j < \omega_1 : n_j = n_{j_0}$ and $m_j = m_{j_0}\}$ has size ω_1. It is easy to see that $\varphi(x,a)$ then divides ω_1-times with respect to k_{i_0}, $1/n_{j_0}$ and $1/m_{j_0}$. □

5.3.5 Definition: A theory T *does not have the tree property* or is *NTP* if and only if no formula has the tree property. ♣

Given the combinatorial nature of the tree property, notwithstanding the extra parameter ε, the proof of the following proposition in the continuous context is identical to the proof in [Cas07, remark 4.1 and proposition 4.4]

5.3.6 Proposition: *The following are equivalent:*

1. *T is NTP*

2. *For every $\varepsilon > 0$, k and δ, no formula divides ω_1 times with respect to k, ε and δ*

3. *For every $\delta > 0$, $\varepsilon > 0$ and k, no formula divides ω times with respect to k, ε and δ;*

4. *For every B, and for every complete type $p \in S(B)$ in finitely many variables, and there is a set $A \subseteq B$ such that $|A| \leq |T|$, and p does not divide over A.*

The equivalence of 1 and 4 can be restated as follows:

5.3.7 Theorem: *T is simple if and only if it has NTP*

5.3.8 Definition: Let $\pi(x, a)$ be a partial type over A, Δ be a finite set of formulae in the free variables x and y, where y is a tuple of variables of the same order type as a. Let $\varepsilon, \delta > 0$. The *D-rank* $D(\varphi, \Delta, k, \varepsilon, \delta)$ is defined to be the largest ordinal α such that $D(\varphi, \Delta, k, \varepsilon, \delta) \geq \alpha$ as per the following recursive definition:

1. $D(\pi, \Delta, k, \varepsilon, \delta) \geq 0$ if and only if π is consistent;
2. If λ is a limit ordinal, then $D(\pi, \Delta, k, \varepsilon, \delta) \geq \lambda$ if and only if $D(\pi, \Delta, k, \varepsilon, \delta) \geq \beta$ for every $\beta < \lambda$
3. If $\alpha = \beta + 1$, then $D(\pi, \Delta, k, \varepsilon, \delta) \geq \alpha$ if and only if there is $\psi \in \Delta$, and a b such that $D(\pi \cup \{\psi(x, b)\}, \Delta, k, \varepsilon, \delta) \geq \beta$, and $\psi(x, b)$ k, ε, δ-divides over A.

If $D(\pi, \Delta, k, \varepsilon, \delta) \geq \alpha$ for every α, then we write $D(\pi, \Delta, k, \varepsilon, \delta) = \infty$. ♣

To simplify the notation, we let $\bar{s}$ denote tuples $(\Delta, k, \delta, \varepsilon)$, and we write $D(\pi(x, a), \bar{s})$ for the corresponding $D(\pi(x, a), \Delta, k, \varepsilon, \delta)$. If $\bar{s}_i = (\Delta_i, k_i, \varepsilon_i, \delta_i)$ for $1 \leq i \leq n$, then we let

$$\bigcup_{i \leq n} \bar{s}_i = \left(\bigcup_{1 \leq i \leq n} \Delta_i, \max\{k_1, ..., k_n\}, \min\{\varepsilon_1, ..., \varepsilon_n\}, \min\{\delta_1, ..., \delta_n\} \right).$$

Note that since we are not allowing the numbers δ_i and ε_i to be 0, $\bigcup_{i \leq n} \bar{s}_i$ is still a tuple of the appropriate form. If $\bar{s} = (\Delta, k, \varepsilon, \delta)$, then we let $\Delta(\bar{s}) = \Delta$, $k(\bar{s}) = k$, $\varepsilon(\bar{s}) = \varepsilon$ and $\delta(\bar{s}) = \delta$.

The following lemma establishes the correspondence between the D-ranks and the tree-property. Intuitively, the D-rank of a formula φ is the height of a dividing tree for φ. The fact that this correspondence between the tree property and the D-rank holds in continuous logic allows us to type-define the property $D(\pi(x, a), \bar{s}) \geq n$.

5.3.9 Lemma: *$D(\pi(x, a), \bar{s}) \geq n$ if and only if there is a sequence $(\varphi_i(x, x_i) : i < n)$ of formulae in Δ and a sequence $(a_i : i < n)$ such that $\pi(x, a) \cup \{\varphi_i(x, a_i) : i < n\}$ is consistent, and $\varphi_i(x, a_i)$ $\bar{s}$-divides over $\{a\} \cup \{a_j : j < i\}$.*

5.3.10 Lemma: *For any partial type $\pi(x, a)$, n, k,$\varepsilon > 0$ and $\delta > 0$, there is a type $\Theta_{\pi, k, \varepsilon, \delta}$ in the variables $x_1, ..., x_n$ such that $b_1, ..., b_n \models \Theta_{\varphi, k, \varepsilon, \delta}(x_1, ..., x_n)$ if and only if $b_1, ..., b_n$ is a dividing chain of length n with respect to k and ε for π.*

5.3.11 Theorem: *T is NTP if and only if for every finite Δ, k, $\varepsilon, \delta > 0$, $D(\varphi, \Delta, k, \varepsilon, \delta)$ is finite.*

Proof: This follows from 5.3.9, since if $D(\pi, \Delta, k, \varepsilon, \delta) > n$ for every n, then $\pi(x, a)$ has arbitrarily long dividing chains, and therefore, it has a dividing chain of length ω_1. By Lemma 5.3.2, this implies the existence of a dividing tree of length ω_1 for π. □

5.3.12 Lemma: *Assume T is NTP. Let $\pi_1(x, a) \subseteq \pi_2(x, a)$ be partial types. If π_2 is a dividing extension of π_1, then for some $\bar{s}$, $D(\pi_1, \bar{s}) > D(\pi_2, \bar{s})$.*

Proof: Without loss of generality, we can assume π_1 and π_2 are closed under conjunctions. Since $\pi_1(x,a) \subseteq \pi_2(x,a)$, any dividing chain for π_2 is a dividing chain for π_1. Therefore, for every $\bar{s}$,

$$D(\pi_2, \bar{s}) \leq D(\pi_1, \bar{s}).$$

Suppose that π_2 divides. Then by corollary 5.2.4, there is a formula $\varphi \in \pi_2$ which (k, ε, δ)-divides over A for some k, $\varepsilon > 0$ and $\delta > 0$. We have that $\pi_1(x,a) \cup \{\varphi(x,a)\} \subseteq \pi_2(x,a)$. Therefore, letting $\bar{s} = (\{\varphi\}, k, \varepsilon, \delta)$, we get

$$D(\pi_2, \bar{s}) \leq D(\pi_1 \cup \{\varphi(x,a)\}, \bar{s}).$$

Let $\alpha = D(\pi_1 \cup \{\varphi(x,a)\}, \bar{s})$. Then there is a dividing chain of length α for $\pi_1 \cup \{\varphi(x,a)\}$, which means that there is a dividing chain of length $\alpha + 1$ for π_1, thus showing that

$$D(\pi_1 \cup \{\varphi(x,a)\}, \bar{s}) < D(\pi_1, \bar{s}).$$

A proof of the following lemma for first-order logic can be found in [Cas07]. The proof for continuous logic is identical

5.3.13 Lemma: *Let $\pi(x,a)$ be a partial type, and let $\varphi_1, ..., \varphi_n$ be formulae. Then for every $\bar{s} = (\Delta, k, \varepsilon, \delta)$,*

$$D\left(\pi \cup \left\{\bigvee_{1 \leq i \leq n} \varphi_i(x,a)\right\}, \bar{s}\right) = \max\{D(\pi \cup \{\varphi_i(x,a)\}, \bar{s})\}.$$

5.3.14 Lemma: *Suppose T with NTP. Suppose $\pi(x,a)$ is a partial type over A. Then π has a non-dividing extension to a complete type p over A.*

Proof: By lemma 5.2.5, for every φ such that φ divides, there is δ such that $\varphi(x,a) \leq \delta$ divides. Let δ_φ be any such δ, and consider the set

$$\Sigma = \pi(x,a) \cup \{\varphi(x,a) \geq \delta_\varphi : \varphi \text{ divides}\}.$$

We need to show that this set is consistent. If not, then there is a finite set $\varphi_1, ..., \varphi_n$ such that $\pi(x,a) \models \bigvee_{i \leq n}(\varphi_i(x,a) < \delta_\varphi)$. Now the formula $\varphi \leq \delta_{\varphi_i}$ divides. By lemma 5.3.12, there is $\bar{s}_i$ such that

$$D(\pi(x,a) \cup \{\varphi_i(x,a) \leq \delta_\varphi\}, \bar{s}_i) < D(\pi(x,a), \bar{s}_i).$$

Let $\bar{t} = \bigcup \bar{s}_i$. Then

$$D\left(\pi \cup \left\{\bigvee_{1 \leq i \leq n} \varphi_i(x,a) \leq \delta_{\varphi_i}\right\}, \bar{t}\right) = \max\{D(\pi \cup \{\varphi_i(x,a) \leq \delta_{\varphi_i}\}, \bar{t})\}$$

from which we get

$$D\left(\pi \cup \left\{\bigvee_{1 \leq i \leq n} \varphi_i(x,a) \leq \delta_{\varphi_i}\right\}, \bar{t}\right) < D(\pi, \bar{t}).$$

On the other hand, since $\pi(x,a) \models \bigvee_{i\leq n}(\varphi_i(x,a) < \delta_\varphi)$, any dividing sequence for π is also a dividing sequence for $\pi \cup \{\varphi_i(x,a) \leq \delta_{\varphi_i}\}$, so we get that

$$D\left(\pi \cup \left\{ \bigvee_{1\leq i\leq n} \varphi_i(x,a) \leq \delta_{\varphi_i} \right\}, \bar{t}\right) \geq D(\pi, \bar{t}),$$

which is a contradiction. Therefore, it must be the case that Σ is consistent. If p is any complete type extending Σ, then p does not divide. This completes the proof. □

Using lemma 5.3.14, one can carry the proof of the following corollary over from [Wag00].

5.3.15 Corollary: *Let p be a partial type over B, and let $A \subseteq B$, Then p is a non-dividing extension of $p|A$ if and only if for every $\bar{s}$, $D(p,\bar{s}) = D(p|A,\bar{s})$.*

5.3.16 Theorem (Extension theorem): *If T is simple, and $A \subseteq B \subseteq C$ are small sets, then any complete type $p \in S_C(A)$ has non-dividing extensions to B.*

Proof: By 5.3.7, p does not divide over A as a partial type over B. Therefore, there is a complete type $q \in S_C(B)$ extending p which does not divide over A by lemma 5.3.14. □

Section 5.4

Morley sequences and the independence theorem

Let I be a linearly ordered set. A sequence $(a_i : i \in I)$ is *A-independent* if and only if for every $i \in I$, $a_i \underset{A}{\perp} \{a_j : j < i\}$. A *Morley sequence* in a type $p(x)$ is a sequence $(a_i : i \in I)$ which is A-independent and A-indiscernible, and such that $\mathcal{C} \models p(a_i)$ for every $i \in I$.

The goal of this section is to state the analogues of the final three properties of non-dividing, namely symmetry, transitivity and the independence theorem. At this point, the proofs of these properties in the first-order context are purely combinatorial, and therefore proofs in the continuous context are identical.

5.4.1 Proposition: *Let B be a small set, and let $A \subseteq B$. Let $p(x)$ be a complete type over B, and suppose $p(x)$ does not divide over A. Then there is a Morley sequence $(a_i : i < \omega)$ in p which is B-indiscernible.*

5.4.2 Proposition: *Suppose T is simple. Then a formula $\varphi(x,a)$ divides over A if and only if for every infinite Morley sequence $(a_i : i \in I)$ in $\mathrm{tp}(a/A)$, the set $\{\varphi(x,a_i) : i \in I\}$ is inconsistent.*

5.4.3 Theorem (Propositions 5.5 and 5.6 in [Cas07]): *If T is a simple theory, then:*

Symmetry: *For any small tuples $a,b,c \in \mathcal{C}$, $a \underset{c}{\perp} b$ if and only if $b \underset{c}{\perp} a$;*

Transitivity: *For any small tuples $a,b,c,d \in \mathcal{C}$, if $a \underset{c}{\perp} b$ and $a \underset{b}{\perp} d$, then $a \underset{c}{\perp} d$*

We note here that transitivity follows from symmetry and 5.2.10.

5.4.4 Theorem (Independence theorem): *Let T be a simple theory, $\mathcal{M} \models T$, $p \in S(\mathcal{M})$ and $a, b \supset \mathcal{M}$ be such that* $a \underset{\mathcal{M}}{\perp} b$. *Let $p \in S(a)$ and $q \in S(b)$ be non-dividing extensions of p to a and b. Then p and q have a common non-dividing extension to a complete type over ab.*

5.4.5 Theorem: *Let $R(a, b, c)$ be a relation defined on the small subsets of $\mathcal{C}$. Suppose R satisfies the following:*

Invariance: *R is invariant under automorphisms of $\mathcal{C}$*

Existence: *For every a and every $c \subseteq b$, there is $a' \equiv_c a$ such that $R(a', b, c)$;*

Symmetry: *For every a and every $c \subseteq b$, $R(a, b, c) \iff R(b, a, c)$*

Transitivity: *For any a, and any $b \subseteq c \subseteq d$, $R(a, d, b)$ if and only if $R(a, d, c)$ and $R(a, c, b)$*

Local character: *For any finite a, and any b, there is a $c \subseteq b$ such that $|c| \leq |T|$, and $R(a, b, c)$*

Finite character: *$R(a, b, c)$ if and only if $R(a', b', c)$ for every finite $a' \subseteq a$ and $b' \subseteq b$*

Independence theorem: *Let $\mathcal{M} \models T$, and let $a, b \supset \mathcal{M}$ be such that $R(a, b, \mathcal{M})$. Suppose $R(c, a, \mathcal{M})$ and $R(c', b, \mathcal{M})$, and $c \equiv_{\mathcal{M}} c'$. Then there is $c'' \equiv_a c$ such that $c'' \equiv_b c'$ and $R(c'', ab, \mathcal{M})$*

Then $R(a, b, c)$ if and only if $a \underset{c}{\perp} b$, *and consequently, the theory T is simple.*

This proof relies on the following statement from set theory.

5.4.6 Lemma (Fodor's lemma): *Let κ be a regular uncountable cardinal, and let $S \subseteq \kappa$ be stationary. If $f : S \to \kappa$ has the property that $f(x) < x$ for every $x \in S$, then there is a subset $S' \subseteq S$, stationary in κ such that f is constant on S'.*

Proof: This proof was communicated to us by Bradd Hart. Let a, b, c be small sets such that $c \subseteq b$, and suppose $R(a, b, c)$. We want to show that $a \underset{c}{\perp} b$. Let $p(x, b) = \text{tp}(a/b)$, and let $I = (b_i : i < |T|^{++})$ be a c-indiscernible sequence in $\text{tp}(b/c)$. We need to show that $\bigcup_{b_i \in I} p(x, b_i)$ is consistent. Let $S = \{\alpha : cf(\alpha) = |T|^+\}$, and note that this is a stationary subset of $|T|^{++}$. By local character, for every $\alpha \in S$, there is $\beta < \alpha$ such that $R(b_\alpha, b_{<\alpha}, cb_{<\beta})$, where $b_{<\alpha} = (b_i : i < \alpha)$. Define $f : S \to |T|^{++}$ so that $f(\alpha)$ is the smallest β with this property. Then $f(\alpha) < \alpha$ for every α. By Fodor's lemma, there is a stationary subset $S' \subseteq S$ such that $f|S'$ is constant. Let β be the value of f on S', and let $I_{S'} = (b_i : i \in S')$. Note that by the definition of f and β, for any $\alpha' \in S'$, we have $R(b_{\alpha'}, b_{<\alpha'} \cap I_{S'}, cb_{<\beta})$, and $I_{S'}$ is indiscernible over $cb_{<\beta}$.

Let $\mathcal{M} \in \text{Mod}(T)$ be such that $R(\mathcal{M}, I_{S'}, cb_{<\beta})$. Using compactness, one can find an $\mathcal{M}$-indiscernible sequence J such that for every $a_1, ..., a_n \in J$, there are $a'_1, ..., a'_n \in I_{S'}$

such that $a_1 \cdots a_n \equiv_{cb_{<\beta}} a'_1 \cdots a'_n$. Note that this makes J an R-independent sequence by invariance.

Let $b' \in J$, and consider $p(x, b')$. Let $p'(x, b', \mathcal{M})$ be an R-extension of $p(x, b')$ to $\mathcal{M}b'$ (such a type exists by the extension property). There is $a' \equiv_{b'} a$ such that $R(a', \mathcal{M}b', b')$ by the existence property. By invariance, $R(a', b', c)$, so by transitivity and weakening, $R(a', b', \mathcal{M})$. There is a type q over $\mathcal{M}$ such that $p'(x, b', \mathcal{M})$ is an R-extension of q. Note that $p'(x, b'', \mathcal{M})$ extends q for every $b'' \in J$ by the $\mathcal{M}$-indiscernibility of J. We can therefore use the independence theorem and conclude that $\bigcup_{b'' \in J} p'(x, b'', \mathcal{M})$ is consistent.

Now, $\bigcup_{b'' \in J} p'(x, b'', \mathcal{M})$ contains the type $\bigcup_{b'' \in J} p(x, b'')$, so this type is consistent as well. For any $b''_1, ..., b''_n \in J$, there are $b_1, ..., b_n \in I_{S'}$ such that $b''_1 \cdots b''_n \equiv_c b_1 \cdots b_n$, therefore $\bigcup_{b \in T_{S'}} p(x, b)$ is consistent. By indiscernibility of $I_{S'}$ and the fact that $I_{S'} \subseteq I$, we get that $\bigcup_{b' \in I} p(x, b')$ is consistent. □

Chapter 6
Non-archimedean Banach spaces

Important recent developments in algebraic model theory deal with various theories of valued fields. The general theory of fields is almost absent from the continuous framework, because general fields are intrinsically discrete. Any attempt to put a metric d on a field K would result in a discrete metric, as the predicate $d(x, y) > 0$ would be definable (it is equivalent to $\exists z[z(x-y) = 1]$). Therefore, continuous logic cannot really shed any more light on fields than classical logic. Work in progress by Ben Yaacov in [Ben09b] deals with valued fields in the framework of continuous logic.

One of the principal difficulties put forward in [Ben09b] is finding a suitable language to describe valued fields in continuous logic. Because they are unbounded structures, they do not fit within the confines of the bounded continuous logic described in [BBHU08], and because they are ultrametric spaces, the usual trick of working only inside the unit ball (which in this case is just the valuation ring) does not work.

Banach spaces over non-archimedean valued fields have the same problem. One possible solution is to use $\mathfrak{R}$-valued languages from chapter 3 to define a language, using norms of linear combinations as relation symbols. However, this approach has the unwanted consequence of introducing points of infinite norm. While points of infinite norm do not cause any problems as far as model theory is concerned, they do limit our ability to transfer results from functional analysis, as these results usually do not account for points of infinite norm.

We begin the chapter by stating some basic facts about the theory of generalized ultrametric spaces, viewed as $\mathfrak{R}$-valued structures. Proposition 6.1.2 shows why it is impossible to restrict ourselves to a proper ball in an ultrametric space.

We then recall the definitions of non-archimedean valued field and non-archimedean Banach spaces over them. What really sets non-archimedean Banach spaces apart from their real or complex counterparts is the presence of a notion of orthogonality which can be defined using only the norm, a property which in the real or complex world is shared only with Hilbert spaces. Also, *every* non-archimedean Banach space contains a copy of c_0. In fact, every non-archimedean Banach space is an immediate extension of a c_0. These results, which are folklore results of non-archimedean functional analysis, suggest that non-archimedean Banach spaces behave a lot like real or complex Hilbert spaces, and hint towards the fact that the theory of non-archimedean Banach spaces (if such a theory can be defined) should have properties similar to those of the theory of real Hilbert spaces.

We then define a language $\mathscr{L}_K$ for non-archimedean Banach spaces based on norms of linear combinations, and define a theory $T_{B,K}$ of "Banach space-like" structures. We then argue that though it does nothing to prevent points of infinite norm, $T_{B,K}$ is very close to the usual category of Banach spaces. The points of infinite norm in models of $T_{B,K}$ are completely indistinguishable from one another, and thus can be collected into a single point via the *emboundment process* described in [Ben08]. This allows us to define a functor $F : \mathbb{NAB}_K \to \mathrm{Mod}_=(T_{B,K})$, which turns out to be an equivalence of categories. Therefore, the results of non-archimedean functional analysis, which are stated in $\mathbb{NAB}_K$, can be transferred to $\mathrm{Mod}_=(T_{B,K})$.

In section 6.5, we state and prove some of these results. We show that $T_{B,K}$ has quantifier elimination and is λ-stable for every cardinal λ such that $\lambda^{\aleph_0} = \lambda$. Under the assumptions that K be either locally compact and densely valued, or that the value group $val(K)$ is the same as the space of norms of the Banach space V, we show that for any κ, there is a unique $\aleph_1$-saturated non-archimedean Banach space of dimension κ.

Section 6.1

Generalized ultra-metric spaces

Throughout, we will assume that $\Gamma \subseteq \mathfrak{R}$. A *generalized ultra-metric space* is a set X together with a generalized metric function $d : X \times X \to \Gamma$ satisfying the strong triangle inequality $d(x,y) \leq \max\{d(x,z), d(z,y)\}$ for every $x, y, z \in X$. If Γ is a dense ordering, then we will call X a *densely valued* space. If Γ is discrete, the X is called *discretely valued.* The epithet "generalized" refers to the fact that d is allowed to assume ∞ as a value.

6.1.1 Fact: *In an ultra-metric space, if $y \in B(x,\varepsilon)$, then $B(x,\varepsilon) = B(y,\varepsilon)$. Therefore, two balls in an ultrametric space are either disjoint or comparable via inclusion.*

In the following proposition, we show that whereas generalized ultrametric spaces clearly form an elementary class, the closed balls of an ultrametric space are in general not definable. Recall that a subset $D \subseteq X$ is definable if and only if there is a formula $\varphi(x)$ such that $\varphi(x) = d(x, D)$.

6.1.2 Proposition: *Assume Γ is dense and let (I, U) be an ultrafilter pair. Le (X, d) be a generalized ultrametric space with $x_0 \in X$. Let $\gamma \in \Gamma$. Then there is $\langle x_i \rangle_U \in X^U$ with the following properties:*

1. $d_{X^U}(\langle x_i \rangle_U, \langle x_0 \rangle_U) = \gamma$
2. *there is no sequence $(y_i : i \in I) \in X^I$ such that $\langle y_i \rangle_U = \langle x_i \rangle_U$ with the property that $d(y_i, x_0) \leq \gamma$ for every $i \in I$.*

Proof: Without loss of generality, $I = \mathbb{N}$. Pick $y_i \in X$ such that $\gamma < d(y_i, x_0) \leq \gamma + 1/i$. This is possible because we are assuming that Γ is dense. Then $\lim_{i \to U} d(y_i, x_0) = \gamma$, so we have that $\langle y_i \rangle_U$ satisfies $d(\langle y_i \rangle_U, \langle x_0 \rangle_U) = \gamma$.

However, since we are in an ultrametric space, if $(x_i : i \in \mathbb{N})$ is any sequence such that $d(x_i, x_0) \leq \gamma$, then $d(x_i, y_i) = d(y_i, x_0)$, so $\lim_{i \to U} d(x_i, y_i)$ cannot be 0. □

Proposition 6.1.2 also implies that closed balls of finite radii in a generalized ultrametric space cannot be definable. As a consequence, we cannot quantify over a ball of diameter strictly smaller than the diameter of the whole space.

6.1.3 Definition: Two decreasing chains of closed balls C and C' in a banach space $\mathcal{V}$ are called *equivalent*, denoted $C \sim C'$ if there is a chain D such that both C and C' can be cofinally embedded in D. A *sphere* is an equivalence class of $\sim$. ♣

Note that because we are assuming that $\Gamma \subseteq \mathfrak{R}$, every sphere has a countable representative. The collection of all spheres in X is denoted $\mathrm{Sph}(X)$. If C is a chain of balls, then we define

$$\mathrm{rad}(C) \stackrel{\text{def}}{:=} \inf\{\mathrm{rad}(B) : b \in C\}.$$

Clearly two equivalent chains have the same radius, so we can put

$$\mathrm{rad}(S) = \mathrm{rad}(C)$$

for any $C \in S$. We will say that a sphere S *contains a point* if and only if any representative for S has non-empty intersection.

6.1.4 Definition: A ultrametric space is *complete* if and only if every sphere of radius 0 contains a point, and it is *spherically complete* if and only if every sphere contains a point. Note that an ultrametric space X is complete if and only if every Cauchy sequence in X converges. ♣

When considering ultrametric spaces as $\mathfrak{R}$-valued structures, we do so using a language $\mathscr{L}$ with a single sort symbol S, and a single relation symbol d standing in for the metric. In this language, we have the following theorem:

6.1.5 Theorem: *Let X be an $\aleph_1$-saturated ultra-metric space. Then $\mathcal{V}$ is spherically complete. In particular, every κ-saturated ultra-metric space is spherically complete.*

Proof: Let $C = \{B(\mathbf{a}_i, \varepsilon_i) : i \in \mathbb{N}\}$ be a decreasing sequence of balls. The type $\pi(x) = \{\|x - \mathbf{a}_i\| \leq \varepsilon_i : i \in \mathbb{N}\}$ is finitely consistent, since every finite subset of it states the existence of an element in the intersection of finitely many of the balls $B(\mathbf{a}_i, \varepsilon_i)$. Since $\mathcal{V}$ is $\aleph_1$-saturated, it contains a realization $\mathbf{x}$ of $\pi(x)$. □

6.1.6 Theorem: *A increasing union of spherically complete ultrametric spaces $\{X_i : i \in \mathbb{N}\}$ is spherically complete.*

Proof: Let $\{X_i : i \in \mathbb{N}\}$ be an increasing list of spherically complete spaces, and let $X = \bigcup X_i$. Let $\{B_X(a_i, \gamma_i) : i \in \mathbb{N}\}$ be a decreasing sequence of balls. For every i such that $a \in X_i$, we have $B_X(a, \gamma) \cap X_i = B_{X_i}(a, \gamma)$. Since we are working in a ultra-metric space, this means that for every i, either $B_X(a, \gamma) \cap X_i = \emptyset$ or $B_X(a, \gamma) \cap X_i = B_{X_i}(b, \gamma)$ for some $b \in X_i$. In other words, the intersection of any X_i with a ball of X is either empty or a ball. Therefore, if for some i, $B_X(a_j, \gamma_j) \cap X_i$ is not empty for every j, then $\{B_X(a_j, \gamma_j) \cap X_i : j \in J\}$ is a decreasing sequence of balls in X_i, and therefore the intersection is non-empty since X_i is spherically complete We now show that this always happens. Let $f : \mathbb{N}^{[2]} \to \{0, 1\}$ be defined by, for $i < j$:

$$f(i,j) = \begin{cases} 0 & \text{if } B_X(a_j, \gamma_j) \cap X_i = \emptyset \\ 1 & \text{otherwise} \end{cases}$$

By Ramsey's theorem, there is an infinite $I \subseteq \mathbb{N}$ such that f is constant on $I^{[2]}$. Suppose $f(I^{[2]}) = 0$. Then for every $i, j \in I$, we have $B_X(a_j, \gamma_j) \cap X_i = \emptyset$. Since I is infinite, it is cofinal. Therefore, for every $j \in I$, the ball $B_X(a_j, \gamma_j)$ has empty intersection with cofinally many X_i's, so $B_X(a_j, \gamma_j) = \emptyset$, which is absurd. Therefore, $f(i^{[2]}) = 1$, which means that for some j, X_j intersects cofinitely many of the balls $B_X(a_i, \gamma_i)$. That is to say, there is an infinite $I \subseteq \mathbb{N}$ such that $X_j \cap B_X(a_i, \gamma_i)$ for every $i \in I$. Since X_j is spherically complete the sequence $\{B_X(a_i, \gamma_i) : i \in I\}$ has non-empty intersection in X_j, so the whole sequence has non-empty intersection in X, which is what we wanted. □

Section 6.2

Non-archimedean Banach spaces

We list some definitions and properties relevant to the theory of valued fields and non-archimedean Banach spaces. The results of this section are standard results of non-archimedean functional analysis, and are thus stated without proof. Let G be an ordered multiplicative abelian group. By a *valued field with value group* G, we shall mean a structure $(K, |\cdot|)$ with the following properties:

1. K is a field;
2. $|\cdot| : K^\times \to G$ is a multiplicative group homomorphism;
3. $|0| = 0$
4. $|x+y| \leq \max\{|x|, |y|\}$

Aside from writing the valuation multiplicatively, the notation is standard. We let $\mathrm{val}(K) \overset{\text{def}}{:=} \{|x| : x \in K\}$ denote the value group of K. We will assume throughout that $\mathrm{val}(K) = G$. We use the notation $\mathcal{O}_K$ for the valuation ring, $\mathfrak{m}_K$ for the maximal ideal, and $k = \mathcal{O}_K/\mathfrak{m}_K$ for the residue field. We will assume that the value group $G \subseteq \mathbb{R}$ is a multiplicative subgroup.

6.2.1 Definition: K is *maximal* if there is no valued field extension L/K with $\mathrm{val}(L) = \mathrm{val}(K)$ and $\mathcal{O}_L/\mathfrak{m}_L = \mathcal{O}_K/\mathfrak{m}_K$. An extension L/K with $\mathrm{val}(L) = \mathrm{val}(K)$ and $\mathcal{O}_L/\mathfrak{m}_L = \mathcal{O}_K/\mathfrak{m}_K$ is called an *immediate extension*. Thus K is maximal if and only if it has no proper immediate extension. ♣

6.2.2 Definition: Let G be an ordered abelian group. A *G-module* Γ is a partially ordered set with a least element $0 \in \Gamma$ on which G acts in an order-preserving manner, i.e.

$$g\gamma \leq g\gamma' \text{ if } \gamma \leq \gamma'$$

and

$$g\gamma \leq g'\gamma \text{ if } g \leq g'$$

To achieve greatest generality, one should consider any poset acted upon by G. However, for simplicity, unless we are working with a discrete Γ, we will always work with $G = \mathbb{R}^+_\infty$ and Γ is dense in G, where the action is given by multiplication. The variable γ will range over elements of Γ. A *Banach space over K with value G-module* Γ is a structure $(\mathcal{V}, \|\cdot\|)$ which satisfies the following:

B1: $\mathcal{V}$ is a K-vector-space;

B2: $\|\cdot\| : \mathcal{V} \to \Gamma$ is a function

B3: $\|\mathbf{u} + \mathbf{v}\| \leq \max\{\|\mathbf{u}\|, \|\mathbf{v}\|\}$;

B4: $\|r\mathbf{u}\| = |r|\|\mathbf{u}\|$;

B5: $\mathcal{V}$ is complete with respect to $\|\cdot\|$.

In analogy with valued fields, we define, for every $\gamma \in \Gamma$,

1. $\mathcal{O}_{\mathcal{B}}(\gamma) \stackrel{\text{def}}{:=} \{\mathbf{x} \in \mathcal{B} : \|\mathbf{x}\| \leq \gamma\}$

2. $\mathfrak{m}_{\mathcal{B}}(\gamma) \stackrel{\text{def}}{:=} \{\mathbf{x} \in \mathcal{B} : \|\mathbf{x}\| < \gamma\}$

The set $\{\|\mathbf{x}\| : \mathbf{x} \in \mathcal{V}\}$ will be denoted $\|\mathcal{V}\|$. We shall not, in general, assume that $\|\mathcal{V}\| = \Gamma$. For every $\gamma \in \Gamma$, $\mathcal{O}_{\mathcal{B}}(\gamma)$ and $m_{\mathcal{B}}(\gamma)$ are naturally $\mathcal{O}_K$-modules, and the quotient $\mathcal{O}_{\mathcal{B}}(\gamma)/\mathfrak{m}_{\mathcal{B}}(\gamma)$ is a k-vector space. We will denote this k-vector space by $\overline{\mathcal{B}(\gamma)}$, and call it the *residue vector space at* γ. This is the terminology and notation used in [Roo78]. When $\gamma = 1$, then we write $\mathcal{O}_{\mathcal{B}}$ instead of $\mathcal{O}_{\mathcal{B}}(1)$ and $\mathfrak{m}_{\mathcal{B}}$ instead of $\mathfrak{m}_{\mathcal{B}}(1)$. In analogy with the notation used for valued fields, we write $\mathfrak{b}$ for the residue vector space $\mathcal{O}_{\mathcal{B}}/\mathfrak{m}_{\mathcal{B}}$. There is a topology on $\mathcal{V}$ obtained by taking all sets of the form $B_\varepsilon(x) = \{y \in \mathcal{V}, \|y - x\| < \varepsilon\}$ as basic open sets.

6.2.3 Definition: Let $X \subseteq \mathcal{V}$ be any set. We denote by $\langle X \rangle$ the closure of the smallest subspace of $\mathcal{V}$ containing X. $\langle X \rangle$ is called the *closed subspace of* $\mathcal{V}$ *generated by* X ♣

6.2.4 Definition: Let $T : \mathcal{V} \to \mathcal{W}$ be linear. We define $\|T\| \stackrel{\text{def}}{:=} \sup_{\mathbf{x} \in \mathcal{V}} \frac{\|T(\mathbf{x})\|}{\|\mathbf{x}\|}$. T is called *linear homeomorphism* if it is a linear bijection, and $\|T\|$ and $\|T^{-1}\|$ are both finite. It is called a *similarity* if there is $c \in \mathbb{R}$ such that $\|T(\mathbf{x})\| = c\|\mathbf{x}\|$ for every $\mathbf{x} \in \mathcal{V}$. If $c = 1$, then T is called an *isometry.* ♣

6.2.5 Definition: Let $X \subseteq \mathcal{W}$, where $\mathcal{W}$ is a K-Banach space. $\mathcal{W}$ is called an *immediate extension of* X if and only if for every $\mathbf{x} \in \mathcal{W}\backslash\{0\}$, $d(\mathbf{x}, X) < \|\mathbf{x}\|$. Here $d(x, X) = \inf\{\|x - y\| : y \in X\}$. ♣

6.2.6 Definition: The *spherical completion* $\mathcal{V}^{Sph}$ of a Banach space $\mathcal{V}$ is the smallest spherically complete extension of $\mathcal{V}$. That is to say, $\mathcal{V}$ embeds isometrically in $\mathcal{V}^{Sph}$, and if $\mathcal{W}$ is any other spherically complete extension of $\mathcal{V}$, then there is an isometric embedding $\mathcal{V}^{Sph} \to \mathcal{W}$. ♣

Note that the previous definition can be relativised. If $\mathcal{V} \subseteq \mathcal{W}$, then there is, in $\mathcal{W}$, a maximal subspace which is an immediate extension of $\mathcal{V}$. Such a subspace will be referred to as a *relative spherical completion* of $\mathcal{V}$ in $\mathcal{W}$.

6.2.7 Theorem (Fleischer, see [Roo78]): *Any Banach space* $\mathcal{V}$ *has a spherical completion* $\mathcal{V}^{Sph}$*, which is unique up to isomorphism.* $\mathcal{V}^{Sph}$ *is an immediate extension of* $\mathcal{V}$*, and any spherically complete immediate extension of* $\mathcal{V}$ *is isomorphic to* $\mathcal{V}^{Sph}$.

6.2.8 Theorem (Ingleton): *A Banach space* $\mathcal{I}$ *is spherically complete if and only if the following holds: for every* $\mathcal{D}$*, every injective linear map* $S : \mathcal{D} \to \mathcal{E}$ *and every linear map* $T : \mathcal{D} \to \mathcal{I}$*, there is a linear map* $\tilde{T} : \mathcal{E} \to \mathcal{I}$*, such that* $\tilde{T} \circ S = T$*, and* $\|\tilde{T}\| = \|T\|$.

Section 6.3

Orthogonality and c_0

What sets non-archimedean Banach spaces apart from their archimedean counterparts is the presence of a notion of orthogonality. In this respect, a Banach space over a non-archimedean field has a structure that is cleaner than a real banach space, and almost resembles a real Hilbert space.

6.3.1 Definition: Let $\mathcal{B}$ be a Banach space, and let $\mathbf{v}, \mathbf{w} \in \mathcal{B}$. Then $\mathbf{v}$ and $\mathbf{w}$ are *orthogonal*, denoted $\mathbf{v} \perp \mathbf{w}$, if and only if for every $k, \ell \in K$,

$$\|k\mathbf{v} + \ell\mathbf{w}\| = \max\{\|k\mathbf{v}\|, \|\ell\mathbf{w}\|\}.$$

A finite set $\{\mathbf{v}_i : 1 \leq i \leq n\}$ is orthogonal if and only if for every $k_1, ..., k_n \in K$, we have

$$\left\|\sum_{i=1}^{n} k_i\mathbf{v}_i\right\| = \max\{\|k_i\mathbf{v}_i\| : 1 \leq i \leq n\}.$$

If $X = \{\mathbf{v}_i : i \in I\}$ is any set of vectors in $\mathcal{B}$, then X is orthogonal if all its finite subsets are orthogonal. ♣

We now show the existence of maximal orthogonal sets. We follow the argument given in [Roo78]. Let $\mathcal{B}$ be a Banach space.

6.3.2 Proposition: *Let X be an orthogonal subset of $\mathcal{V}$. If val(K) is dense, then for every $\varepsilon > 0$, there is an orthogonal subset Y of $\mathcal{V}$ such that $\sharp Y = \sharp X$ and for every $\mathbf{y} \in Y$, $1 - \varepsilon \leq \|\mathbf{y}\| \leq 1$.*

If val(K) is discrete, then there is an orthogonal subset Y of $\mathcal{V}$ such that $\sharp Y = \sharp X$ and for every $\mathbf{y} \in Y$, $\pi \leq \|\mathbf{y}\| \leq 1$, where π is a generator for val(K) such that $\pi < 1$.

Proof: For every $\mathbf{x} \in X$, choose $\ell_{\mathbf{x}} \in K$ with $|\ell| \leq 1/\|\mathbf{x}\|$, and let $Y = \{\ell_{\mathbf{x}}\mathbf{x} : \mathbf{x} \in X\}$. The map $\mathbf{x} \mapsto \ell_{\mathbf{x}}\mathbf{x}$ is a bijection, since $\ell_{\mathbf{x}}\mathbf{x} = \ell_{\mathbf{y}}\mathbf{y}$ for $\mathbf{x} \neq \mathbf{y}$ would contradict the fact that $\mathbf{x} \perp \mathbf{y}$. □

Recall from section 6.2 that $\mathcal{O}_{\mathcal{B}}$ is an $\mathcal{O}_K$-module, and that $\mathfrak{b} = \mathcal{O}_{\mathcal{B}}/\mathfrak{m}_{\mathcal{B}}$ is a k-vector space.

6.3.3 Theorem: *Let $X \subseteq \mathcal{O}_{\mathcal{B}}$. Let $\pi : \mathcal{O}_{\mathcal{B}} \to \mathfrak{b}$ be the canonical projection. Then X is orthogonal if and only if $\pi[X]$ is k-linearly independent. It is a maximal orthogonal set if and only if $\pi[X]$ is a basis of $\mathcal{O}_{\mathcal{B}}/\mathfrak{m}_{\mathcal{B}}$ over k.*

Proof: The second assertion follows readily from the first. For the first assertion, suppose X is orthogonal. Then for every finite $F \subseteq X$, we have

$$\left\|\sum_{\mathbf{x} \in F} k_{\mathbf{x}}\mathbf{x}\right\| = \max\{\|k_{\mathbf{x}}\mathbf{x}\| : \mathbf{x} \in F\}$$

for every $(k_{\mathbf{x}} : \mathbf{x} \in F)$. In particular, this is true of the elements of $\mathcal{O}_K$. Suppose $\sum_{\mathbf{x} \in F} k_{\mathbf{x}}\pi(\mathbf{x}) = 0$ is a finite linear combination of elements of $\pi[X]$, where $k_{\mathbf{x}} \in k$. By definition, this means

$$\sum_{\mathbf{x} \in \mathcal{B}} k_{\mathbf{x}}\pi(\mathbf{x}) \in \mathcal{M}_{\mathrm{B}}, \text{ so that } \left\|\sum_{\mathbf{x} \in F} k_{\mathbf{x}}\pi(\mathbf{x})\right\| < 1.$$

Since X is orthogonal,

$$\left\|\sum_{\mathbf{x}\in F} k_{\mathbf{x}}\pi(\mathbf{x})\right\| = \max\{\|k_{\mathbf{x}}\mathbf{x}\| : \mathbf{x} \in F\}$$

so $\max\{\|k_{\mathbf{x}}\mathbf{x}\| : \mathbf{x} \in F\} < 1$. Since $\|\mathbf{x}\| = 1$ for every $\mathbf{x} \in X$, we conclude that it must be the case that $|k_{\mathbf{x}}| < 1$ for every $\mathbf{x} \in F$, which means that $k_{\mathbf{x}} = 0$ in k for every $\mathbf{x} \in F$.

Conversely, suppose X is not orthogonal. Let $F \subseteq X$ be a finite set such that

$$\left\|\sum_{\mathbf{x}\in F} k_{\mathbf{x}}\mathbf{x}\right\| < \max\{\|k_{\mathbf{x}}\mathbf{x}\| : \mathbf{x} \in F\}.$$

We can pick $k_{\mathbf{x}}$ of norm 1, so $\overline{k_{\mathbf{x}}} \neq 0$ in k. Then $\sum_{\mathbf{x}\in F} k_{\mathbf{x}}\mathbf{x} = 0$ in $\mathfrak{b}$. □

6.3.4 Corollary: *1. Any Banach space contains a maximal orthogonal set. Any two maximal orthogonal sets X and Y in $\mathcal{B}$ have the same cardinality, and $\langle X\rangle$ and $\langle Y\rangle$ are isomorphic.*

2. With the same notation as proposition 6.3.2, the following is true. If val(K) is dense, then for every $\varepsilon > 0$, there is a maximal orthogonal subset X of $\mathcal{V}$ such that for every $\mathbf{x} \in X$, $1 - \varepsilon \leq \|\mathbf{x}\| \leq 1$. If val($K$) is discrete, then there is a maximal orthogonal subset X of $\mathcal{V}$ such that for every $\mathbf{x} \in X$, $\pi \leq \|\mathbf{x}\| \leq 1$.

6.3.5 Definition: The cardinality of a maximal orthogonal set in $\mathcal{B}$ will be denoted $\dim(\mathcal{B})$, and referred to as the *dimension* of $\mathcal{B}$. ♣

Note that in the non-archimedean framework, there can be non-isomorphic spaces of the same dimension. This is because if $\mathcal{W}$ is an immediate extension of $\mathcal{V}$, then $\dim(\mathcal{V}) = \dim(\mathcal{W})$. However, this is the only case in which this happens.

The simplest examples of Banach spaces over the real or complex fields are c_0 and ℓ^∞. It is also possible to define these spaces over non-archimedean fields, and in this setup, they are very important. The material of this section comes from [Roo78], in which the reader can find proofs for all the results.

6.3.6 Definition: Let $\mathcal{V}$ be a Banach space, and let I be a set. Let $S = \{x_i : i \in I\} \subseteq \mathcal{V}$. Then S is called *summable to s* if and only if for every ε there is a finite set $J_\varepsilon \subseteq I$ such that $\left\|s - \sum_{j\in J} \mathbf{x}_j\right\| < \varepsilon$ for any finite set $J \supseteq J_\varepsilon$. The element s, if it exists, is unique, and denoted $\sum_{i\in I} \mathbf{x}_i$. ♣

6.3.7 Proposition: *1. A countable set $\{x_n : n \in \mathbb{N}\}$ is summable if and only if*

$$\lim_{n\to\infty} \|\mathbf{x}_n\| = 0$$

2. If $\{x_n : n \in \mathbb{N}\}$, and $\rho : \mathbb{N} \to \mathbb{N}$ is any injective map, then $\{x_{\rho(n)} : n \in \mathbb{N}\}$ is summable.

3. If ρ is a permutation, then $\sum_{i=0}^{\infty} \mathbf{x}_i = \sum_{i=0}^{\infty} \mathbf{x}_{\rho(i)}$.

4. *If I is not countable, and $\{\mathbf{x}_i : i \in I\}$ is summable, then the set $\{i \in I : \mathbf{x}_i \neq 0\}$ is countable.*

6.3.8 Definition: Let $\mathcal{B}$ be a non-archimedean Banach space over K, and let $X \subseteq \mathcal{B}$ be a set. The Banach space $c_0(X)$ consists of those elements $\mathbf{x} \in K^X$ such that for every $\varepsilon > 0$, $\{x \in X : |\mathbf{x}_x|\|x\| > \varepsilon\}$ is finite. Note that $c_0(X)$ consists of those sequences in K^X which are summable in the sense of definition 6.3.6. ♣

6.3.9 Proposition: *If $X \subseteq \mathcal{V}$ is orthogonal, then $\langle X\rangle \cong c_0(X)$.*

6.3.10 Theorem: *Every non-archimedean Banach space $\mathcal{V}$ has a closed subspace $\mathcal{U}$ such that $\mathcal{U} \cong c_0(X)$ for some $X \subseteq \mathcal{V}$. In fact, every non-archimedean Banach space $\mathcal{V}$ is an immediate extension of a space of the form $c_0(X)$.*

Section 6.4

A language for non-archimedean Banach spaces

Non-archimedean Banach spaces are unbounded structures. Proposition 6.1.2, when applied to a non-archimedean Banach space over a valued field implies that we cannot use the usual trick of considering the class of unit balls of Banach spaces as our primary object of study. This class is not elementary. We must therefore consider Banach spaces as a whole, and describe a language in full $\mathfrak{R}$ logic.

This unfortunately separates Banach spaces as models of a continuous theory from actual Banach spaces. In $\mathfrak{R}$-valued logic, general predicates can (and often will) assume the value ∞. Therefore, any language which does not impose specific finite bounds on the norm predicate on Banach spaces will give rise to elements of infinite norm. Also, the presence of elements of infinite norm forces us to consider a *relational* language for Banach spaces, since only a relational language can make sense of $a + b$ if both a and b are elements of infinite norm. The language $\mathscr{L}_K$ is defined as follows:

1. A sort symbol B;

2. For every finite tuple $r_1, ..., r_n \in K$, a relation symbol $\left\|\sum_{i=1}^{n} r_i x_i\right\|$. We write $\|x\|$ for $\left\|\sum_{i=1}^{n} x_i\right\|$;

3. A relation symbol d with domain $\mathsf{B} \times \mathsf{B}$.

We now state the axioms for non-archimedean Banach spaces in $\mathscr{L}_K$. This list of axioms has a lot of redundancy.

B0: $\forall xy[\mathsf{d}(x, y) = \|x - y\|]$

B1: For every n, and every $r_1, ..., r_n \in K$, $\forall x_1 ... x_n \left[\left\|\sum_{i=1}^{n} r_i x_i\right\| \leq \max\{\|r_1 x_1\|, ..., \|r_n x_n\|\}\right]$

B2: For every n, and every $r_1, ..., r_n \in K$, and every permutation $\rho : \{1, ..., n\} \to \{1, ..., n\}$.

$$\forall x_1 ... x_n \left[\left\| \sum_{i=1}^{n} r_i x_i \right\| = \left\| \sum_{i=1}^{n} r_{\rho(i)} x_{\rho(i)} \right\| \right].$$

B3: For every n, and every $r_1, ..., r_n \in K$, $\forall x_1 ... x_n \left[\left\| \sum_{i=1}^{n} r_i x_i \right\| = \left\| \sum_{r_j \neq 0} r_j x_j \right\| \right]$, provided that $\{i : r_i \neq 0\} \neq \emptyset$.

B4: For every n, and every $s_1, ..., s_n \in K$, $\forall y_1 ... y_n \exists x \left[\left\| x - \sum_{i=1}^{n} s_i y_i \right\| \right]$

B5: For every n, and every $r_1, ..., r_n \in K$, $\forall x \left[\left\| \sum_{i=1}^{n} r_i x \right\| = \left\| (\sum_{i=1}^{n} r_i) x \right\| \right]$

B6: For every n, and every $r \in K$, $\forall x_1, ..., x_n \left[\left\| \sum_{i=1}^{n} r x_i \right\| = |r| \left\| \sum_{i=1}^{n} x \right\| \right]$

6.4.1 Lemma: *The relation symbol $\|y - x\|$, when interpreted in $\mathcal{B} \in \mathrm{Mod}(T_{B,K})$, defines a generalized pseudo-ultrametric on $\mathcal{B}$, which restricts to a pseudo-ultrametric on $\mathcal{B}^{<\infty}$*

Proof: $\|x - x\| = 0$ by axiom **B5** and **B6**. If $x, y, z \in \mathcal{B}$, then by **B2**, **B3** and **B5**,

$$\|x - y\| = \|x + z - z + y\|$$

and by **B1** and **B2**,

$$\|x + z - z + y\| \leq \max\{\|x - z\|, \|z - y\|\}$$

as required. For the second assertion, note that if $\|x\| < \infty$ and $\|y\| < \infty$, then by **B1**, $\|x - y\| \leq \max\{\|x\|, \| - y\|\} = \max\{\|x\|, |-1|\|y\|\} = \max\{\|x\|, \|y\|\} < \infty$, thus completing the proof. □

6.4.2 Lemma: *The formula*

$$\rho(x, y) = \frac{1 - \dfrac{1}{e^{\|y-x\|}}}{e^{\|y-x\|} e^{\|x\| \vee \|y\|}}$$

defines a pseudo-metric on $\mathcal{B}$. The formula ρ has the property that $\rho(x, y) = 0$ if and only if either $\|x - y\| = 0$, or both x and y have infinite norm.

Proof: The proof that ρ is a pseudo-metric is written in [Ben09b, page 14]. The numerator in the definition of ρ is always a number in $\mathfrak{R}$. Therefore, $\rho(x, y) = 0$ if and only if either the numerator is 0, or the denominator is infinite. If $e^{\|y-x\|} - 1 = 0$, then $\|x - y\| = 0$. The denominator $e^{\|y-x\|} e^{\|x\| \vee \|y\|}$ is infinite if and only if $e^{\|x\| \vee \|y\|} = \infty$, if and only if $\|x\| \vee \|y\| = \infty$, which happens if and only if both x and y have infinite norm. □

6.4.3 Corollary: *$\mathcal{B}/\rho$ is an $\mathscr{L}_K$-structure, and contains a unique point of infinite norm. The relation symbol $\|y - x\|$ defines a generalized ultrametric on $\mathcal{B}/\rho$.*

Proof: Since ρ is a metric on $\mathcal{B}/\rho$, Lemma 6.4.2 implies that any two points of infinite norm are equal in $\mathcal{B}/\rho$. For the second assertion, by Lemma 6.4.1, $\|x-y\|$ defines a generalized pseudo-ultrametric on $\mathcal{B}$, and it is compatible with ρ by Lemma 6.4.2. It is therefore still a generalized pseudo-ultrametric after passing to the quotient $\mathcal{B}/\rho$. If $\|x-y\|=0$, then by $\rho(x,y)=0$, so that $x=y$ in $\mathcal{B}/\rho$. □

We let $\mathscr{L}_{K,\infty}$ be an expansion of $\mathscr{L}_K$ by a relation symbol d_∞ of domain $\mathsf{B}\times\mathsf{B}$. We let $T_{B,K,\infty}$ expand $T_{B,K}$ with the axioms

$$\mathsf{d}_\infty(x,y)=\frac{1-\frac{1}{e^{\|y-x\|}}}{e^{\|y-x\|}e^{\|x\|\vee\|y\|}}.$$

Consider the category $\mathbb{NAB}_K$ defined as follows:

Objects: Non-archimedean Banach spaces over K

Morphisms: Isometries between Banach spaces

Composition: Function composition

For every ultrafilter pair (I,U), there is an ultraproduct functor $\prod_U:\mathbb{NAB}_K^I\to\mathbb{NAB}_K$ defined as follows for an I-indexed sequence $(\mathcal{B}_i:i\in I)$:

1. The underlying set of $\prod_U\mathcal{B}_i$ is $\{\langle x_i\rangle_U\in\prod_U^*\mathcal{B}_i:\|x_i\|<\infty \text{ for every } i\in I\}$, where $\prod_U^*\mathcal{B}_i$ denotes the ultraproduct of the topological spaces $(\mathcal{B}_i:i\in I)$;
2. $\langle x_i\rangle_U+\langle y_i\rangle_U=\langle x_i+y_i\rangle_U$
3. $r\langle x_i\rangle_U=\langle rx_i\rangle_U$
4. $0=\langle 0\rangle_U$

Note that since sequences $(f_i:i\in I)$, where each $f_i:\mathcal{B}_i\to\mathcal{B}_i'$ is an isometry are equicontinuous in the usual topology on Banach spaces, $\prod_U$ is indeed a functor on $\mathbb{NAB}_K^I$.

6.4.4 Theorem: *There is an ultraproduct preserving equivalence of categories between the category* $\mathrm{Mod}(T_{B,K,\infty})$ *and the category* $\mathbb{NAB}_K$.

Proof: Let $\mathcal{B}\in\mathrm{Mod}(T_{B,K,\infty})$. First we put a normed vector space "structure" on $\mathcal{B}^{<\infty}$ by defining the operations of addition and scalar multiplication up to elements of norm 0. If $w,w'\in\mathcal{B}$ both satisfy $\|x-u-v\|=0$, then $\|w-w'\|=0$, and therefore $\mathsf{d}(x,y)=0$. This is a consequence of **B1** and **B5**, as

$$\|w-w'\|=\|w-u-v-w'+u+v\|\le\max\{\|w-u-v\|,\|w'-u-v\|\}=0$$

thus showing that if a solution exists, then it is unique up to $\|w-w'\|=0$, and therefore up to $\mathsf{d}(x,y)=0$ as required. Existence is a consequence of **B4**.

If w is a solution of $\|x-u-v\|=0$, then

$$\|w\|=\|w+u-u+v-v\|\le\max\{\|w-u-v\|,\|u+v\|\}=\|u+v\|$$

and

$$\|u+v\|\le\|u+v+w-w\|\le\max\{\|u+v-w\|,\|w\|\}=\|w\|$$

showing that $\|w\| = \|u+v\|$ as should be expected. By **B1**, if $\|u\| < \infty$ and $\|v\| < \infty$, then $\|u+v\| \leq \max\{\|u\|, \|v\|\} < \infty$. This therefore defines an addition operation on $\mathcal{B}^{<\infty}/\llbracket \mathsf{d} \rrbracket_{\mathcal{B}}$. Similarly, we can define scalar multiplication, and the element 0, thus making $\mathcal{B}^{<\infty}/\llbracket \mathsf{d} \rrbracket_{\mathcal{B}}$ into a normed vector space, the norm being given by $\|x\|$. Taking its completion makes is into a Banach space. We let $F(\mathcal{B})$ be this Banach space. It is easy to see that if $h : \mathcal{B} \to \mathcal{B}'$ is an elementary map, then it induces an isometry $F(h) : F(\mathcal{B}) \to F(\mathcal{B}')$, thus showing that F is indeed a functor.

We must show that F commutes with all the ultraproduct functors. Let (I, U) be an ultrafilter pair, and $\{\mathcal{B}_i : i \in I\}$ be an I-indexed sequence of models of $T_{B,K,\infty}$, and let $\mathsf{B} = \prod_U \mathcal{B}_i$. By definition, $\|\langle x_i \rangle_U\| = \lim_{i\to U} \|x_i\|$, and $\lim_{i \to U} \|x_i\| \leq r < \infty$ if and only if for every $R > r$, $\{i \in I : \|r_i\| \leq R\} \in U$. Therefore, any $\langle x_i \rangle_U$ of finite norm is U-equivalent to a *bounded* sequence $(x_i : i \in I)$, and is therefore an element of $\prod_U F(\mathcal{B}_i)$. Conversely, by definition, an element of $\prod_U F(\mathcal{B}_i)$ is a bounded sequence $(x_i \in F(\mathcal{B}_i) : i \in I)$, and is therefore easily seen to be an element of $F(\prod_U \mathcal{B}_i)$.

Let $\mathcal{B}$ be a Banach space, and define the model $\mathcal{B}^\infty$ as follows:

1. $\mathsf{B}(\mathcal{B}^\infty) = \mathcal{B} \cup \{\infty\}$
2. $\llbracket \|\sum_{i=1}^n r_i x_i\| \rrbracket_{\mathcal{B}^\infty} (v_1, \ldots, v_n) = \|\sum_{i=1}^n r_i v_i\|$ for every $v_1, \ldots, v_n \in \mathcal{B}$

Since ∞ is supposed to be the unique point to which all unbounded sequences converge, the definition of $\llbracket \|\sum_{i=1}^n r_i x_i\| \rrbracket_{\mathcal{B}^\infty}$ extends by continuity to ∞ for every relation symbol. It is easy to see that $F(\mathcal{B}^\infty) = \mathcal{B}$ for any Banach space $\mathcal{B}$, thus showing that F is essentially surjective on objects. F is also faithful, since two distinct elementary maps $g, h : \mathcal{B} \to \mathcal{B}'$ have to disagree on the set of points of finite norm of $\mathcal{B}$. Also, a map $f : F(\mathcal{B}) \to F(\mathcal{B}')$ can be lifted to a map $\mathcal{B} \to \mathcal{B}'$ by defining $f(\infty) = \infty$. This finished the proof that F is an equivalence of categories. □

Section 6.5

Model theory

In this section we state and prove the model theoretic results about non-archimedean Banach spaces which were advertised in the introduction, namely that over any K, $T_{B,K,\infty}$ has quantifier elimination and is λ-stable for every cardinal λ satisfying $\lambda^{\aleph_0} = \lambda$. We also show that under some assumptions on K, $T_{B,K,\infty}$ has a unique $\aleph_1$-saturated model of any dimension. First, the definition of $\mathscr{L}_K$ and of $T_{B,K,\infty}$ makes the following obvious by quantifier elimination:

6.5.1 Theorem: *A model $\mathcal{V} \models T_{B,K,\infty}$ is spherically complete if and only if it is $\aleph_1$-saturated.*

Consequently, we get the following model theoretic proof of the existence of spherical completions: for any Banach space $\mathcal{V}$, and any non-principal ultrafilter on $\mathbb{N}$, the ultrapower $\mathcal{V}^U$ is spherically complete. Therefore, it contains a spherical completion of $\mathcal{V}$.

6.5.2 Lemma: *Let $\mathcal{V}$ be a Banach space, and let $\mathcal{W}$ be an immediate extension of $\mathcal{V}$. Let $\mathcal{X}$ be a spherically complete space, and suppose $T : \mathcal{V} \to \mathcal{X}$ is an isometric embedding. Then T extends to an isometric embedding $\bar{T} : \mathcal{W} \to \mathcal{X}$.*

Proof: This proof is taken from [Roo78]. Since $\mathcal{X}$ is spherically complete, T extends to a map $\bar{T} : \mathcal{W} \to \mathcal{X}$ such that $\|T\| = \|\bar{T}\|$. We show that this map is an isometry. Since T is an isometry, $\|T\| = 1$, and therefore $\|\bar{T}\| = 1$. Let $\mathbf{a} \in \mathcal{W}$. There is $\mathbf{b} \in \mathcal{V}$ such that $\|\mathbf{b} - \mathbf{a}\| < \|\mathbf{a}\|$. Note that $\|\mathbf{a}\| = \|\mathbf{b}\|$ Then we have: .

$$\|T(\mathbf{a}) - T(\mathbf{b})\| = \|T(\mathbf{a} - \mathbf{b})\| \leq \|T\|\|\mathbf{b} - \mathbf{a}\| < \|\mathbf{a}\| = \|\mathbf{b}\| = \|\bar{T}(\mathbf{b})\|$$

The last equality holds because $\mathbf{b} \in \mathcal{V}$, and $\hat{T}|\mathcal{V} = T$, which is an isometry. From this we get that $\|T(\mathbf{b}) - T(\mathbf{a})\| < \|T(\mathbf{b})\|$, which implies $\|T(\mathbf{a})\| = \|T(\mathbf{b})\| = \|\mathbf{b}\| = \|\mathbf{a}\|$. □

Since the spherical completion of $\mathcal{V}$ is an immediate extension, we get the following as a direct corollary of Lemma 6.5.2:

6.5.3 Theorem: *Let $\mathcal{X}$ be spherically complete, and let $T : \mathcal{V} \to \mathcal{X}$ be an isometric embedding. Then T extends to an isometric embedding $T : \mathcal{V}^{Sph} \to \mathcal{X}$. The same conclusion is true of any relative spherical completion $\mathcal{U}$ of $\mathcal{V}$ in $\mathcal{W}$, where $\mathcal{V} \subseteq \mathcal{W}$.*

6.5.4 Lemma: *Let $\dim(\mathcal{U}) \leq \dim(\mathcal{V})$. Let $M_{\mathcal{U}}$ and $M_{\mathcal{V}}$ be maximal orthogonal sets in $\mathcal{U}$ and $\mathcal{V}$ respectively. Let*

$$c = \inf\left\{\frac{\|\mathbf{x}\|_{\mathcal{V}}}{\|\mathbf{y}\|_{\mathcal{U}}} : \mathbf{x} \in M_{\mathcal{V}}, \mathbf{y} \in M_{\mathcal{U}}\right\} \text{ and } C = \sup\left\{\frac{\|\mathbf{x}\|_{\mathcal{V}}}{\|\mathbf{y}\|_{\mathcal{U}}} : \mathbf{x} \in M_{\mathcal{V}}, \mathbf{y} \in M_{\mathcal{U}}\right\}$$

If $0 < c \leq C < \infty$, then there is a topological embedding $T : \langle M_{\mathcal{U}}\rangle \to \langle M_{\mathcal{V}}\rangle$ such that for every $\mathbf{x}$, $c \leq \|T\mathbf{x}\|/\|\mathbf{x}\| \leq C$. If in addition $\mathcal{V}$ is spherically complete, then T extends to a map $\mathcal{U} \to \mathcal{V}$.

Proof: Let $f : M_{\mathcal{U}} \to M_{\mathcal{V}}$ be any injective map. Then for every $\mathbf{x} \in \mathcal{U}$,

$$c \leq \frac{\|f(\mathbf{x})\|_{\mathcal{V}}}{\|\mathbf{x}\|_{\mathcal{U}}} \leq C.$$

Extend f to an embedding $F : \langle X_{\mathcal{U}}\rangle \to \langle X_{\mathcal{V}}\rangle$ by linearity. Then T has the required property. If $\mathcal{V}$ is spherically complete, then T can be extended to all of $\mathcal{U}$, completing the proof. □

6.5.5 Proposition: *Let $\mathcal{V}^{Sph} \subseteq \mathcal{W}$, and let $\mathbf{w} \in \mathcal{W}$. Then $\mathbf{w} \perp \mathcal{V}^{Sph}$ if and only if $\mathbf{w} \perp \mathcal{V}$.*

Proof: If $\mathbf{w} \perp \mathcal{V}^{Sph}$, then clearly $\mathbf{w} \perp \mathcal{V}$ as well. Conversely, suppose $\mathbf{w} \not\perp \mathcal{V}^{Sph}$. There is $\mathbf{v} \in \mathcal{V}^{Sph}$ such that $\|\mathbf{w} - \mathbf{v}\| = \|\mathbf{w}\|$. Note that $\|\mathbf{v}\| = \|\mathbf{w}\|$. Let $\mathbf{v}' \in \mathcal{V}$ be such that $\|\mathbf{v} - \mathbf{v}'\| < \|\mathbf{v}\|$. We then have:

$$\begin{aligned}
\|\mathbf{w} - \mathbf{v}'\| &= \|\mathbf{w} - \mathbf{v} + \mathbf{v} - \mathbf{v}'\| \\
&\leq \max\{\|\mathbf{w} - \mathbf{v}\|, \|\mathbf{v} - \mathbf{v}'\|\} \\
&< \max\{\|\mathbf{w}\|, \|\mathbf{v}\|\} \\
&= \|\mathbf{w}\| = \|\mathbf{v}\|
\end{aligned}$$

Thus showing that $d(\mathbf{w}, \mathcal{V}) < \|\mathbf{w}\|$ as well, completing the proof. □

6.5.6 Proposition: *Let $\mathcal{U}$ be a relative spherical completion of $\mathcal{V}$ in $\mathcal{W}$. Then $\mathcal{U}$ has an orthogonal complement in $\mathcal{W}$.*

Proof: Otherwise $\mathcal{W}$ would be an immediate extension of $\mathcal{V}^{Sph_\mathcal{W}}$. □

A metric theory T has *quantifier elimination* if and only if for every formula $\varphi(x)$, there is a quantifier-free formula $\psi(x)$ such that $T \models \forall x[|\varphi(x) - \psi(x)|]$. The following theorem states the criterion for quantifier elimination which we will be using on non-archimedean Banach spaces. A proof can be found in [BBHU08].

6.5.7 Theorem: *A metric theory T has quantifier elimination if and only if for every $\mathcal{M} \in \mathrm{Mod}(T)$, every substructure $\mathcal{M}_0 \subseteq \mathcal{M}$, and every $|\mathcal{M}|^+$-saturated model $\mathcal{N}$, every embedding $\mathcal{M}_0 \to \mathcal{N}$ can be extended to an embedding $\mathcal{M} \to \mathcal{N}$.*

We will also need the following definition:

6.5.8 Definition: A theory T is λ-stable if and only if for every set A of density character $|A| \leq \lambda$, the type space $S(A)$ has density character $|S(A)| \leq \lambda$ ♣

6.5.9 Theorem: *If val(K) $= \Gamma$, or if K is a locally compact non-archimedean valued field with a dense valuation, then $T_{B,K,\infty}$ has quantifier elimination and is $\aleph_0$-stable in $\mathscr{L}_{K,\infty}$.*

Proof: Let $\mathcal{V} \models T_{B,K,\infty}$, and let $\mathcal{U} \subseteq \mathcal{V}$ be a substructure. Let $\mathcal{W}$ be $\dim(\mathcal{V})^+$-saturated. Suppose $T : \mathcal{U} \to \mathcal{W}$ is an isometric embedding. By Lemma 6.5.2, T extends to an isometric embedding $T : \mathcal{U}' \to \mathcal{W}$ for any immediate extension $\mathcal{U}'$ of $\mathcal{U}$. Therefore, it extends to an isometric embedding $T : \mathcal{U}' \to \mathcal{W}$, where $\mathcal{U}'$ is a maximal immediate extension of $\mathcal{U}$.

Let X be a maximal orthogonal subset of $\mathcal{U}$, and let Y be a maximal orthogonal subset of $\mathcal{V}$ extending X. Note that X and Y can be chosen so that each of their elements has norm 1. If $y \in Y\backslash X$, then $y \perp \mathcal{U}$ by the maximality of X. Since $\mathcal{W}$ is $\dim(\mathcal{V})^+$-saturated, and $|Y\backslash X| < \dim(\mathcal{V})^+$, there is, in $\mathcal{W}$, an orthogonal subset Z extending the set $T[X]$, each of whose element has norm 1, and such that $|Y\backslash X| \leq |Z|$. By Lemma 6.5.4, an injective map $U : Y\backslash X \to Z$ extends to an isometric embedding $U : c_0(Y\backslash X) \to \mathcal{W}$. Since every $y \in Y\backslash X)$ is orthogonal to $\mathcal{U}$, we get an isometric embedding $U^* : \mathcal{U} \oplus c_0(\mathcal{Y}\backslash X) \to \mathcal{W}$. Now note that $\mathcal{V}$ is an immediate extension of $\mathcal{U} \oplus c_0(\mathcal{Y}\backslash X)$, so U^* can be extended to an isometric embedding $\mathcal{V} \to \mathcal{W}$, thus proving quantifier elimination.

We thank Bradd Hart for the following argument for $\aleph_0$-stability. Since we are assuming real norms, Γ has a countable dense subset $X \subseteq \Gamma$. Let A be a countable set, and let $f : A \times X \to S(A)$ be any function such that "$\|a - x\| \leq \varepsilon$" $\in f(a, \varepsilon)$. We claim that $f[A \times X]$ is dense in $S(A)$. Let $p \in S(A)$ be any complete type. By quantifier elimination and the fact that X is dense in Γ, p is determined by a sphere

$$S = \{B(a_n, \varepsilon_n) : n \in \omega, \varepsilon_n \in X, a_n \in A\}.$$

It is easy to see that in fact, $p = \lim_{n\to\infty} f(a_n, \varepsilon_n)$ in the logic topology of $S(A)$. Since $A \times X$ is countable, $f[A \times X]$ is countable, which means $f[A \times X]$ is a countable dense subset of $S(A)$, proving $\aleph_0$-stability. □

6.5.10 Theorem: *Let κ be a cardinal. There is, up to linear homeomorphism, only one spherically complete Banach space of dimension κ. Moreover, if we let T denote the linear homeomorphism, then the following is true*

1. *If Γ = val(K), then T can be chosen to be an isometry.*

2. *If Γ and val(K) are both dense and K is locally compact, then T can be chosen to be an isometry.*

Consequently, in both these cases, there is, up to isomorphism, only one $\aleph_1$-saturated model of $T_{B,K,\infty}$ of dimension κ for any cardinal κ.

Proof: If $\Gamma = \text{val}(K)$, then the conclusion is trivial, since then every Banach space has a maximal orthogonal set consisting of elements of norm 1. Apply lemma 6.5.4 to any such sets.

Suppose Γ and val(K) are both dense, and K is locally compact. If $\Gamma = \text{val}(K)$, then there is nothing to prove. We claim that if K is locally compact, then every $\mathbf{x} \in \mathcal{V}$ has a scalar multiple of norm 1. We can get the required isometry by choosing orthogonal sets of norm 1. Since val(K) is dense, for every $\varepsilon > 0$, there is $k_\varepsilon \in K$ such that $||k_\varepsilon| - 1/\|x\|| \leq \varepsilon$. Note that $\{k_\varepsilon : \varepsilon > 0\}$ is a bounded subset of K. Since K is locally compact, there is a sequence $(\varepsilon_n : n \in \mathbb{N})$ such that $\varepsilon_n \to 0$, and $(k_{\varepsilon_n} : n \in \mathbb{N})$ is Cauchy. Since K is complete, $k_{\varepsilon_n} \to k$ in K. Now $|k| = 1/\|\mathbf{x}\|$, so $\|k\mathbf{x}\| = 1$, and $\mathbf{x}$ can be normalized. □

Bibliography

[Bac74] Paul D. Bacsich. Model theory of epimorphisms. *Canadian Mathematical Bulletin*, 17(4):471–477, 1974.

[BBHU08] Itaï Ben Yaacov, Alexander J. Berenstein, C. Ward Henson, and Alexander Usvyatsov. Model theory for metric structures. In Anand Pillay Zo Chatzidakis, Dugald Macpherson and Alex Wilkie, editors, *Model theory with Applications to Algebra and Analysis, vol. 2*, volume 350 of *London Math Society Lecture Note Series*, pages 315–427. Cambridge University Press, 2008.

[Ben] Itaï Ben Yaacov. Definability of groups in $\aleph_0$-stable metric structures. *Journal of Symbolic Logic*. arXiv:0802.4286.

[Ben03a] Itaï Ben Yaacov. Simplicity in compact abstract theories. *Journal of Mathematical Logic*, 3(2):163–191, 2003.

[Ben03b] Itaï Ben Yaacov. Thickness, and a categoric view of type-space functors. *Fundamenta Mathematicæ*, 179:199–224, 2003.

[Ben05] Itaï Ben Yaacov. Uncountable dense categoricity in cats. *Journal of Symbolic Logic*, 70(3):829–860, 2005.

[Ben08] Itaï Ben Yaacov. Continuous first order logic for unbounded metric structures. *Journal of Mathematical Logic*, 8(2):225–249, 2008. arXiv:0903.4957.

[Ben09a] Itaï Ben Yaacov. Continuous and random vapnik-chervonenkis classes. *Israel Journal of Mathematics*, 173(1):309–333, 2009. doi:10.1007/s11856-009-0094-x, arXiv:0802.0068.

[Ben09b] Itaï Ben Yaacov. Model theoretic properties of metric valued fields, 2009. arXiv:0907.4560.

[BP10] Itaï Ben Yaacov and Arthur Paul Pedersen. A proof of completeness for continuous first-order logic. *The Journal of Symbolic Logic*, 75(1):168–190, March 2010. doi:10.2178/jsl/1264433914, arXiv:0903.4051.

[BS69] J.L. Bell and A.B. Slomson. *Models and Ultraproducts: An Introduction*. North-Holland Publishing Company, 1969.

[BUar] Itaï Ben Yaacov and Alexander Usvyatsov. Continuous first order logic and local stability. *Transactions of the American Mathematical Society*, To appear.

[Cas07] Enrique Casanovas. Simplicity simplified. *Revista Colombiana de Matemáticas*, 41:263–277, 2007.

[CK66] C. C. Chang and H. J. Keisler. *Continuous Model Theory.* Princeton University Press, 1966.

[CK90] C. C. Chang and H. J. Keisler. *Model Theory*, volume 73 of *Studies in Logic and the Foundations of Mathematics.* Elsevier, 1990.

[FHS] Ilijas Farah, Bradd Hart, and David Sherman. Model theory of operator algebras ii: Model theory. *arxiv.org.* arXiv:1004.0741v1.

[Har] Bradd Hart. Conceptual completeness for continuous logic. Slides presented at Makkaifest 2009.

[HI02] Ward Henson and José Iovino. Ultraproducts in analysis. In *In Analysis and Logic, volume 262 of London Mathematical Society Lecture Notes*, pages 1–115. Cambridge University Press, 2002.

[Mak87] Michael Makkai. Stone duality for first order logic. *Advances in Mathematics*, 65:97–170, 1987.

[Mak88] Michael Makkai. Strong conceptual completeness for first order logic. *Annals of Pure and Applies Logic*, 40:167–215, 1988.

[MR77] Michael Makkai and Gonzalo Reyes. *First Order Categorical Logic*, volume 611 of *Lecture Notes In Mathematics.* Springer-Verlag, 1977.

[Roo78] A.C.M Van Rooij. *Non-archimedean functional analysis.* Marcel Dekker, 1978.

[She80] Saharon Shelah. Simple unstable theories. *Annals of Mathematical Logic*, 19(3):177–203, December 1980. "doi:10.1016/0003-4843(80)90009-1".

[Wag00] Frank Wagner. *Simple Theories*, volume 240 of *Lecture Notes of the London Mathematical Society.* Cambridge University Press, 2000.