POSITIVE MODEL THEORY AND COMPACT ABSTRACT THEORIES

ITAY BEN-YAACOV

ABSTRACT. We develop positive model theory, which is a non first order analogue of classical model theory where compactness is kept at the expense of negation. The analogue of a first order theory in this framework is a compact abstract theory: several equivalent yet conceptually different presentations of this notion are given. We prove in particular that Banach and Hilbert spaces are compact abstract theories, and in fact very well-behaved as such.

INTRODUCTION

Trying to extend the classical model-theoretical techniques beyond the strictly first-order context seems to be a popular trend these days. In [Hru97], Hrushovski defines Robinson theories, namely universal theories whose class of models has the amalgamation property. He subsequently works in the category of its existentially closed models, which serves as an analogue of the first order model completion when this does not exist. In [Pil00], Pillay generalises this to the category of existentially closed models of any universal theory. In both cases, one works rather in an existentially universal domain for the category, which replaces the monster model of first order theories.

The present work started independently of the latter, trying to use ideas in the former in order to define a model-theoretic framework where hyperimaginary elements could be adjoined as parameters to the language, the same way we used to do it with real and imaginary ones since the dawn of time: as the type-space of a hyperimaginary sort is not totally disconnected, we need a concept of a theory who just can't say "no". In the terminology of [Hru97], this means we must no longer require the set of basic formulas Δ to be closed for boolean combinations, but only for *positive* ones. The notions of positive model theory, and in particular of positive Robinson theories, follow.

As it turns out, positive Robinson theories are but one of several alternative presentations of the same concept. We prefer therefore to make the distinction between any particular presentation and the fundamental concept itself, which we call *compact abstract theories*, or *cats*.

In the present paper we restrict ourselves to the development of the framework. General model theoretic tools, and in particular simplicity, are developed for it in [Ben02b]. Additional results, and in particular a better treatment of simplicity under the additional hypothesis of thickness, are given in [Ben02c]. These tools are applied in [Ben02a] for the treatment of the theory of lovely pairs of models of a simple theory in case that the theory of pairs is not of first order.

Date: June 11, 2002.

It was pointed out that the definition of a universal domain for a positive Robinson theory is similar to *Assumption III* from [She75, Section 2].

1. Introduction to positive model theory

We introduce positive model theory, which in particular generalises first order model theory. Although it is related to classical first order logic, its development requires a radical change in our point of view: we use at times the language of categories more than that of logic, and the usage of negation and of the universal quantifier is discouraged (not to mention unnecessary).

The basic idea is to replace the notions of elementary extensions and embeddings by that of homomorphisms: for a designated set of "positive" statements, what was true for the domain must be true for its image, but not necessarily the converse. In other words, any positive statement that's true is already decided, whereas those which are not true will not necessarily remain so: they are simply "deferred" for a later decision.

This fact, of being allowed to decide only what we want and defer everything else makes the compactness theorem almost trivial: a short and elegant proof is given below as a corollary of positive Morleyisation (which is, on the other hand, more complicated than first-order Morleyisation).

Due to this shift in point of view and language, and with an easy proof of the compactness theorem, an exposition from scratch seems reasonable, and would make this paper very much self-contained.

1.1. Language and categories of structures.

We start with the basic definitions:

Definition 1.1.

1. A *(relational) signature* $\mathcal{L}$ is a set along with a function $\nu : \mathcal{L} \to \omega$. An element $P \in \mathcal{L}$ is called a $\nu(P)$*-ary predicate symbol*. We also have a distinguished binary predicate symbol $= \in \mathcal{L}$.
2. Let $\mathcal{L}$ be a signature. A $\mathcal{L}$-structure is a set M along with a $\nu(P)$-ary predicate $P^M \subseteq M^{\nu(P)}$ for every predicate symbol $P \in \mathcal{L}$, called the *interpretation* of P in M. The symbol $=$ is always interpreted by equality.

Remark 1.2. Classically one also allows *function symbols*: however, as a function can be represented just as well by the predicate defining its graph, this is not necessary and would only serve to complicate things.

Definition 1.3. Let $X = \{x_i : i < \omega\}$ (where all the x_i are distinct) and call its elements *variables*. In fact, we could have simply taken $X = \omega$, but we follow traditional notation.

We differ somewhat from the standard definitions in the fact that we consider the set of free variables of a formula (including the dummy ones) a part of the information in the formula: for us a $\mathcal{L}$-formula is something of the form $\varphi(x_{\in I})$ for some finite $I \subseteq \omega$, where $x_{\in I}$ is shorthand for $\{x_i : i \in I\}$. If $I = n$ then we write $x_{<n}$.

We are going to define formulas by induction, and for each formula $\varphi(x_{\in I})$ and $\mathcal{L}$-structure M define the set $\varphi(M^I) \subseteq M^I$. For $a_{\in I} \in M^I$, $M \models \varphi(a_{\in I})$ is synonymous with $a_{\in I} \in \varphi(M^I)$.

1. If P is a n-ary symbol, then $P(x_{<n})$ is an *atomic formula*, and $P(M^n) = P^M$, whereby $M \models P(a_{<n}) \iff a_{<n} \in P^M$.

2. If $\varphi(x_{\in I})$ is a formula, $J \subseteq \omega$ is finite, and $f : I \to J$ is a map, then $\psi(x_{\in J}) = f_*(\varphi(x_{\in I}))$ is a formula obtained by *change of variables*:
   For $a_{\in J} \in M^J$, write $f^*(a_{\in J}) = (a_{f(i)} : i \in I) \in M^J$, and for a set $A \subseteq M^I$ define $f_*(A) = f^{*-1}(A) \subseteq M^J$: then $\psi(M^J) = f_*(\varphi(M^I))$, whereby $M \models \psi(a_{\in J}) \iff M \models \varphi(f^*(a_{\in J}))$.
   In actual notation, we may write $f_*(\varphi)$ as $\varphi(x_{f(0)}, \dots, x_{f(n-1)})$, but it must be understood that this is a formula in the variables $x_{\in J}$.
3. If $k < \omega$ and $\varphi_i(x_{\in I})$ is a formula for every $i < k$, then $\chi(x_{\in I}) = \bigwedge_{i<k} \varphi_i(x_{\in I})$ and $\rho(x_{\in I}) = \bigvee_{i<k} \varphi_i(x_{\in I})$ are formulas constructed by *positive (boolean) combinations*: *conjunction* and *disjunction*, respectively. We define $\chi(M^I) = \bigcap_{i<k} \varphi_i(M^I)$ and $\rho(M^I) = \bigcup_{i<k} \varphi_i(M^I)$.
   We sometimes denote the empty conjunction by $\top$ and the empty disjunction by $\bot$.
4. If $I \cap J = \varnothing$ and $\varphi(x_{I\cup J}) = \varphi(x_{\in I}, x_{\in J})$ is a formula then $\psi(x_{\in I}) = \exists x_{\in J}\, \varphi(x_{\in I}, x_{\in J})$ is a formula constructed by *existential quantification*, and $\psi(M^I)$ is the projection of $\varphi(M^I \times M^J)$ on M^I, whereby $M \models \psi(a_{\in I})$ if and only if there is $a_{\in J} \in M^J$ such that $M \models \varphi(a_{\in I}, a_{\in J})$.
5. If $\varphi(x_{\in I})$ is a formula, then $\psi(x_{\in I}) = \neg\varphi(x_{\in I})$ is a formula constructed by *negation*, and $\psi(M^I) = M^I \smallsetminus \varphi(M^I)$.

A formula $\varphi(x_{\in I})$ is I-ary, and the variables $x_{\in I}$ are its *free variables*. A 0-ary formula, that is without free variables, is called a *sentence*, or a *closed formula*.
A *sub-formula* of φ is any formula appearing along its construction.
$\mathcal{L}_{\omega,\omega}$ is the set of all $\mathcal{L}$-formulas.

Notation 1.4. If the set of free variables of a formula is clear from the context or is irrelevant, we just write $\varphi(\bar{x})$, $\varphi(x)$ or even φ.
Also, we may write $\bar{a} \in M$ or even $a \in M$, when it is clear that these are tuples in M^I where I is clear from the context, and we may similarly replace $\varphi(M^I)$ with $\varphi(M)$.

Definition 1.5. Two I-ary formulas φ and ψ are *equivalent* if $\varphi(M^I) = \psi(M^I)$ for every $\mathcal{L}$-structure M.

Convention 1.6. We consider equivalent formulas as equal.

Definition 1.7. An *almost atomic formula* is a change of variables on an atomic formula.

Lemma 1.8. *Every formula is equivalent to one constructed from almost atomic formulas along the same construction tree without any further changes of variables (beyond the almost atomic formulas).*

Proof. Easy. QED

Definition 1.9. 1. A set $\Delta \subseteq \mathcal{L}_{\omega,\omega}$ is a *positive fragment* of $\mathcal{L}$ if it contains all the atomic formulas in $\mathcal{L}_{\omega,\omega}$ and is closed under change of variables, sub-formulas, and positive combinations.
2. Let Δ be a positive fragment. Then $\Sigma(\Delta)$ is its closure under existential quantification, and $\Pi(\Delta) = \{\neg\varphi : \varphi \in \Sigma(\Delta)\}$.

3. The minimal positive fragment, which consists of positive combinations of almost atomic formulas (or, in a more traditional terminology, quantifier-free positive formulas), is noted Δ_0. We also note $\Sigma_1 = \Sigma(\Delta_0)$, $\Pi_1 = \Pi(\Delta_0)$.
4. Let Δ be a positive fragment. Then a map $f : M \to N$ between two $\mathcal{L}$-structures is a Δ-*homomorphism* if $M \models \varphi(a) \implies N \models \varphi(f(a))$ for every $n < \omega$, n-ary formula $\varphi(x_{<n}) \in \Delta$, and $a \in M^n$.
5. A map $f : M \to N$ between two $\mathcal{L}$-structures is a Δ-*embedding* if $M \models \varphi(a) \iff N \models \varphi(f(a))$ for every $n < \omega$, n-ary formula $\varphi(x_{<n}) \in \Delta$, and $a \in M^n$.
6. The category of $\mathcal{L}$-structure where the morphisms are the Δ-homomorphisms is noted $\mathcal{M}_\Delta$.

If the positive fragment Δ is clear from the context, we omit it.

Convention 1.10. When $\mathcal{M}_\Delta$ is clear from the context and $f : M \to N$ is a morphism, then we say that N *continues* M.
When f is clear from the context, and it is also clear that we work in N, we may sometimes omit f, identifying elements of M with their images in N. This is convenient but requires attention, as f is not necessarily injective!

Remark 1.11. If Δ is a positive fragment, then so is $\Sigma(\Delta)$ (since we only consider formulas up to equivalence). Moreover, replacing Δ with $\Sigma(\Delta)$ does not change any of $\Sigma(\Delta)$, $\Pi(\Delta)$, or the notion of Δ-homomorphism (which is always a $\Sigma(\Delta)$-homomorphism). However, it may well change the notion of a Δ-embedding, which is why we consider this an unnatural notion (see also below).

The motivation is straightforward: first of all, given the definition of an $\mathcal{L}$-structure, the natural notion of morphism is that of a map that preserves the truth (though not necessarily falsehood) of every predicate, that is every atomic formula. Clearly, if a map preserves the truth of a set of formulas it does so for every formula obtained thereof by positive combinations and change of variables, so we may pass to the generated positive fragment. We get the notion of a Δ_0-homomorphism and the category $\mathcal{M}_{\Delta_0}$, where indeed we shall work most of the time.

In fact, one sees easily that for every positive fragment Δ, a Δ-homomorphism preserves the truth of every formula in $\Sigma(\Delta)$, so we can add any such formula to Δ and pass to the generated positive fragment without changing $\mathcal{M}_\Delta$.

However, one may also desire (weird as it may seem) not only to preserve the truth of some $\varphi \in \Delta$, but its falsehood as well, which means to preserve $\neg\varphi$. Adding $\neg\varphi$ to Δ and passing to the generated positive fragment would have precisely this effect.

Doing this repeatedly (possibly an infinite number of times) we obtain any category $\mathcal{M}_\Delta$, and at the very end we find $\mathcal{M}_{\mathcal{L}_{\omega,\omega}}$, where first order model theory takes place.

We see from this that the addition to Δ of formulas constructed by negation is the only one that does not come for free, which is why we try to avoid negations. This stands in contrast to the classical exposition of model theory, where one considers Δ-embeddings rather than Δ-homomorphisms, and then it is quantification that does not come for free. Nevertheless, even in our context the truth of a $\Sigma(\Delta)$-formula is slightly more complicated to verify than that of a Δ-formula, in particular when $\Delta = \Delta_0$, which is why we keep the distinction between the two: we aim to show later on that in certain cases Δ and $\Sigma(\Delta)$ have the same power of expression, and then restrict ourselves to Δ.

Convention 1.12. Δ is a positive fragment, and a morphism is one of $\mathcal{M}_\Delta$, that is a Δ-homomorphism.

Definition 1.13. Let $(I, \leq)$ be a directed partial order, and $(M_i : i \in I)$ a Δ_0-inductive system indexed by I: for every $i \leq j$ we have a Δ_0-homomorphism $f_{ij} : M_i \to M_j$, such that $f_{jk} \circ f_{ij} = f_{ik}$ for every $i \leq j \leq k$. Let $N = \varinjlim M_i$ as sets, with maps $g_i : M_i \to N$. For $R \in \mathcal{L}$, we define $N \models R(\bar{a})$ if and only if there is $i \in I$ and $\bar{b} \in M_i$ such that $M_i \models R(\bar{b})$ and $\bar{a} = g_i(\bar{b})$. We note $\varinjlim M_i = N$, now as $\mathcal{L}$-structures.

Lemma 1.14. *With the notations of Definition 1.13, if Δ is a positive fragment and (M_i) is a Δ-inductive system, then for every $\varphi(\bar{x}) \in \Delta$: $\varinjlim M_i \models \varphi(\bar{a})$ if and only if there is $i \in I$ and $\bar{b} \in M_i$ such that $M_i \models \varphi(\bar{b})$ and $\bar{a} = g_i(\bar{b})$.*

Proof. For atomic formulas, this was the definition. We can now use the fact that a positive fragment is closed for sub-formulas and work by induction on $\varphi \in \Delta$.
Positive combinations and existential quantifiers are easy. As for negation, assume that $\neg\varphi(\bar{x}) \in \Delta$ and $\varinjlim M_i \models \neg\varphi(\bar{a})$. Then there are $i \in I$ and $\bar{b} \in M_i$ such that $g_i(\bar{b}) = \bar{a}$, and then necessarily $M_i \models \neg\varphi(\bar{b})$, since $\varphi(\bar{x}) \in \Delta$. Conversely, assume that $M_i \models \neg\varphi(\bar{b})$ but $\varinjlim M_i \models \varphi(g_i(\bar{b}))$. Then there are $j \in I$ and $\bar{c} \in M_j$ such that $M_j \models \varphi(\bar{c})$, and $g_j(\bar{c}) = g_i(\bar{b})$. Then there is $k \geq i, j$ such that $f_{ik}(\bar{b}) = f_{jk}(\bar{c}) = \bar{a}$, say, so $M_k \models \varphi(\bar{a})$ and $M_k \models \neg\varphi(\bar{a})$ (here we finally use the fact that $\neg\varphi \in \Delta$), contradiction. QED

And we conclude:

Proposition 1.15. *With the notations of Definition 1.13, if Δ is a positive fragment and (M_i) Δ-inductive system, then every g_i is a Δ-homomorphism, and $\varinjlim M_i$ is the injective limit in the sense of $\mathcal{M}_\Delta$.*

Let us imagine once more we walk along $\mathcal{M}_\Delta$. At some point we stand at the structure M, and we ask ourselves whether $\varphi(a)$ holds, where $\varphi \in \Sigma(\Delta)$ and $a \in M$. Assume indeed that $M \models \varphi(a)$: if we follow a morphism $f : M \to N$, then $N \models \varphi(f(a))$. If we follow several morphisms one after the other this is still true, and by Proposition 1.15 we can even pass to the limit of a chain of morphisms. On the other hand, if $M \not\models \varphi(a)$, we do not have a definitive answer yet on what will happen when we follow a morphism (unless $\neg\varphi \in \Delta$, of course), so we have to continue and see what happens. However, we shouldn't go too far either: for example, if $\Delta = \Delta_0$, we may end up in a structure where everything is true, in particular equality, so the structure would be reduced to a single point!

This means we need some means to give definitive negative answers, without using negations. For example, by restricting ourselves to those structures where two certain positive statements cannot be true simultaneously: then saying that one is true is a definitive decision for the falsehood of the other. This leads us quite naturally to the notion of a Π-theory.

Definition 1.16. A *(Π-)theory* is a set of Π-sentences, that is a set of statements of the form "for no tuple $\bar{a}$ do we have $\varphi(\bar{a})$", for certain formulas $\varphi \in \Delta$.

If M is a structure then $\mathrm{Th}_\Pi(M)$ it its theory, that is the set of all Π-sentences true in M.
If $\mathcal{M}$ is a class of structures, then $\mathrm{Th}_\Pi(\mathcal{M}) = \bigcap_{M\in\mathcal{M}} \mathrm{Th}_\Pi(M)$.
A *model* of a theory T is a structure $M \models T$, meaning $M \models \varphi$ for every $\varphi \in T$. Two theories are *equivalent* if they have the same models.

Definition 1.17. If T is a theory (that is, a Π-theory), then $\mathcal{M}^0(T)$ is the (full) sub-category of $\mathcal{M}_\Delta$ whose objects are models of T.

Given a theory T we can give definitive positive answers, which may also imply definitive negative answers for other questions. It makes sense now to look for a model of T where every reasonable question has a definitive answer.

Definition 1.18. A model $M \models T$ is *existentially closed (e.c.)* if every Δ-homomorphism $f : M \to N$ with $N \models T$ is a Σ-embedding.
In other words, we require that for every $\varphi(\bar{x}) \in \Sigma(\Delta)$ and every $\bar{a} \in M$, if $N \models \varphi(f(\bar{a}))$ then $M \models \varphi(\bar{a})$.
The category of e.c. models of T is noted $\mathcal{M}(T)$.

Assume that $M, N \in \mathcal{M}(T)$, $f : M \to N$ is a morphism, $\bar{a} \in M$ is some tuple, and $\varphi(\bar{x}) \in \Sigma(\Delta)$. Then, by definition, $M \models \varphi(\bar{a}) \iff N \models \varphi(f(\bar{a}))$. But then, if f is actually an inclusion, then we can write rather $M \models \varphi(\bar{a}) \iff N \models \varphi(\bar{a})$. This justifies:

Notation 1.19. If $M \in \mathcal{M}(T)$, $\bar{a} \in M$ and $\varphi(\bar{x}) \in \Sigma(\Delta)$, then by slight abuse of notation we write $\models \varphi(\bar{a})$ or $\bar{a} \models \varphi$ instead of $M \models \varphi(a)$.

Next step is to show that enough e.c. models exist:

Lemma 1.20. *Every $M \models T$ continues to some e.c. model $N \in \mathcal{M}(T)$.*

Proof. Set $M_0 = M$.
Having constructed M_i, let $\{(\varphi_j, \bar{a}_j) : j < \lambda_i\}$ be an enumeration of all pairs $(\varphi, \bar{a})$ where $\varphi(x_{<n}) \in \Sigma$ and $\bar{a} \in M_i^n$. Set $M_i^0 = M_i$; for $j < \lambda_i$, let $\bar{a}'_j$ be the image of $\bar{a}_j$ in M_i^j: if there is a model $M' \models T$ and a morphism $f : M_i^j \to M'$ such that $M' \models \varphi_j(f(\bar{a}'_j))$, set $M_i^{j+1} = M'$, otherwise $M_i^{j+1} = M_i^j$; For $\delta < \lambda_i$ limit, $M_i^\delta = \varinjlim_{j<\delta} M_i^j$. In the end, take $M_{i+1} = M_i^{\lambda_i}$.
Then $M_\omega = \varinjlim_{i<\omega} M_i$ is clearly e.c.. QED

Remark 1.21. Lemma 1.20 is easily seen to be equivalent to the Axiom of Choice.

1.2. **Positive Morleyisation and compactness.**

Positive Morleyisation is a an adaptation to positive model theory of a first order construction of Michael Morley, which allows us to reduce the case of a general positive fragment to that of Δ_0. It requires more work than in the first order case, but then it is more powerful: for example, the compactness theorem follows as a trivial corollary.

Definition 1.22. Let I be any set of indices, and $\{x_i : i \in I \smallsetminus \omega\}$ new distinct symbols.
For an n-ary formula $\varphi(x_{<n})$ and a map $f : n \to I$ we define a *generalised I-ary formula* $\psi(x_{\in I}) = f_*(\varphi)$, again by change of variables: $f^*(a_{\in I}) = (a_{f(i)} : i < n) \in M^n$,

$\psi(M^I) = f_*(\varphi(M^n)) = f^{*-1}(\varphi(M^n))$, and $M \models \psi(a_{\in I}) \Longleftrightarrow M \models \varphi(f^*(a_{\in I}))$.
Note that this can only be a last step in the construction of a formula. When we say I-ary formula for an infinite I, it is understood we mean generalised.

Convention 1.23. When we say that a generalised I-ary formula $\varphi(x_{\in I})$ is in Δ, we mean that it is obtained from some $\psi \in \Delta$ by a change of variables.
Moreover, in what follows we are going to consider change of variables as done implicitly every time we consider formulas in different tuples of variables, provided that there are indeed obvious changes of variables that would make them all formulas in the same tuple (this is after all the common practice with the standard notation for formulas).

Definition 1.24. A *(partial I-ary) Δ-type* is a set $p(x_{\in I})$ of formulas $\varphi(x_{\in I}) \in \Delta$.
If $p(x_{i\in I})$ is such, and T a theory, then p is consistent with T if there is a model $M \models T$ and $\bar{a} \in M^I$ such that $M \models p(\bar{a})$, that is $M \models \varphi(\bar{a})$ for every $\varphi \in p$. If T is clear from the context, we just say that p is consistent.

We start by studying the easy case where $\Delta = \Delta_0$. We recall the notation $\Sigma_1 = \Sigma(\Delta_0)$, $\Pi_1 = \Pi(\Delta_0)$.

Lemma 1.25. *Let T be a Π_1-theory, and $\Phi(x_{\in I})$ a set of almost atomic I-ary formulas, with I possibly infinite. If every finite subset $\Phi_0 \subseteq \Phi$ is consistent with T (in fact, it suffices that it be consistent with each $\psi \in T$), then Φ is consistent with T.*

Proof. Take $M_0 = \{a_i : i \in I\}$ where these are all distinct elements, and define $\sim$ as the equivalence relation generated by $\{a_i \sim a_j : x_i = x_j \in \Phi\}$. Take $M = M/divsim$, and for every n-ary predicate symbol R take $R^M = \{([a_{i_0}], \dots, [a_{i_{n-1}}]) : R(x_{i_0}, \dots, x_{i_{n-1}}) \in \Phi\}$. Thus $M \models \Phi([a_{\in I}])$.
Assume that there is $\psi \in T$ such that $M \not\models \psi$. This means that there is something true in M that ψ forbids, and it is true in M due to a finite part $\Phi_0 \subseteq \Phi$. This means that Φ_0 is inconsistent with T, and in fact with ψ, contradicting the hypothesis. QED

We wish to make the notion of a definitive negative negation more explicit in some special case:

Lemma 1.26. *Let T be a Π_1-theory, and assume that T is closed for logical consequence: any Π_1 sentence which is true in every model of T is in T. Let $M \models T$ be e.c., R an $n+k$-ary predicate symbol, and $\bar{a} \in M^n$.*
Then $M \not\models \exists \bar{y}\, R(x_{<n}, \bar{y})$ ($|\bar{y}| = k$) if and only if there are $m < \omega$, a tuple $\bar{b} \in M^m$, a map $f : n \to m$ and $\varphi(x_{<m}) \in \Delta_0$ such that $\bar{a} = f^*(\bar{b})$, $M \models \varphi(\bar{b})$, and some $\psi \in T$ is inconsistent with $\varphi(x_{<m}) \wedge R(f^*(x_{<m}), \bar{y})$.
We say then that according to T, the instance $\varphi(\bar{b})$ (of $\varphi(\bar{x})$) forbids $R(\bar{a}, \bar{y})$.

Proof. One direction is clear, since $M \models T$.
For the other, enumerate $M = \{b_i : i \in I\}$, and let $f : n \to I$ be such that $\bar{a} = f^*(\bar{b})$. Set

$$\Phi(x_{i\in I}, \bar{y}) = \{R(f^*(\bar{x}), \bar{y})\} \cup \{R'(h^*(\bar{x})) : R' \in \mathcal{L}, h : \nu(R') \to I, M \models R'(h^*(\bar{b}))\}$$

If $\Phi(\bar{x})$ is consistent with T, say satisfied by $\bar{b}', \bar{c}$ in a model M', then the map sending $\bar{b}$ to $\bar{b}'$ is a Δ_0-homomorphism of M into M'. Since then $M' \models R(f^*(\bar{b}'), \bar{c})$ and M is e.c., we must have also $M \models \exists \bar{y}\, R(f^*(\bar{b}), \bar{y})$, that is $M \models \exists \bar{y}\, R(\bar{a}, \bar{y})$, contradiction.

Therefore Φ is inconsistent with T, and by Lemma 1.25 there is a finite $\Phi_0' \subseteq \Phi$ which is inconsistent with some $\psi \in T$. Write $\Phi_0'(\bar{x}, \bar{y}) = \Phi_0(\bar{x}) \cup \{R(f^*(\bar{x}), \bar{y})\}$. Since Φ_0 is finite it is equivalent to $\varphi(x_{\in J})$ (with the implicit change of variables) where $\varphi \in \Delta_0$ is a conjunction of almost atomic formulas and $J \subseteq I$ is finite such that $m_{\in J}$ enumerates every element of $\bar{a}$ (that is, J contains the image of f). Then $\psi \in T$ is inconsistent with $\varphi(x_{\in J}) \wedge R(f^*(x_{\in J}), \bar{y})$ and $M \models \varphi(b_{\in J})$, as required. QED

Definition 1.27. Let $\mathcal{L}$ be a given signature, and Δ a positive fragment. Define $\mathcal{L}^{\mathrm{QE}(\Delta)} = \mathcal{L} \cup \{R_\varphi : n \in \omega, \varphi(x_{<n}) \in \Delta\}$, where each R_φ is a new n-ary symbol, and $\Delta_0^{\mathrm{QE}(\Delta)}$ is the minimal positive fragment of $\mathcal{L}^{\mathrm{QE}(\Delta)}$. Similarly, $\Pi_1^{\mathrm{QE}(\Delta)} = \Pi(\Delta_0^{\mathrm{QE}(\Delta)})$.
If M is a $\mathcal{L}$-structure, expand it to a $\mathcal{L}^{\mathrm{QE}(\Delta)}$-structure $M^{\mathrm{QE}(\Delta)}$ by defining $R_\varphi^M = \varphi(M^n)$ (that is, $M \models R_\varphi(a_{<n}) \iff M \models \varphi(a_{<n})$) for every $\varphi(x_{<n}) \in \Delta$.

Remark 1.28. If M and N are any $\mathcal{L}$-structures, then $\mathrm{Hom}_{\mathcal{M}_\Delta}(M, N) = \mathrm{Hom}_{\mathcal{M}_{\Delta_0^{\mathrm{QE}(\Delta)}}}(M^{\mathrm{QE}(\Delta)}, N^{\mathrm{QE}(\Delta)})$ as sets of maps. Therefore $-^{\mathrm{QE}(\Delta)} : \mathcal{M}_\Delta \to \mathcal{M}_{\Delta_0^{\mathrm{QE}(\Delta)}}$ is a fully faithful functor.

Definition 1.29. Let T be a $\Pi(\Delta)$-theory, and define $\mathcal{M}^{\mathrm{QE}(\Delta)}(T) = \{M^{\mathrm{QE}(\Delta)} : M \in \mathcal{M}(T)\}$ and $T^{\mathrm{QE}(\Delta)} = \mathrm{Th}_{\Pi_1^{\mathrm{QE}(\Delta)}}(\mathcal{M}^{\mathrm{QE}(\Delta)}(T))$. In other words, $T^{\mathrm{QE}(\Delta)}$ is the $\Pi(\Delta_0^{\mathrm{QE}(\Delta)})$-theory of the class of all the $\mathcal{L}^{\mathrm{QE}(\Delta)}$-structures M which are, as $\mathcal{L}$-structures, e.c. models of T, and in addition interpret each R_φ as φ.
We call $T^{\mathrm{QE}(\Delta)}$ then Δ-*Morleyisation* of T.

Lemma 1.30. *Let Δ be a positive fragment, and $\Delta_0^{QE(\Delta)}$ as above. Let T be a $\Pi(\Delta)$-theory and T' a $\Pi_1^{QE(\Delta)}$-theory, such that $M = (M{\restriction}_{\mathcal{L}})^{QE(\Delta)}$ for every $M \in \mathcal{M}(T)$ (namely, $\varphi(M) = R_\varphi^M$ for every $\varphi \in \Delta$). Assume also that $M^{QE(\Delta)} \models T'$ for every $M \in \mathcal{M}(T)$ and $M{\restriction}_{\mathcal{L}} \models T$ for every $M \in \mathcal{M}(T')$.*
Then the functor $M \mapsto M^{QE(\Delta)}$ is an isomorphism $\mathcal{M}(T) \simeq \mathcal{M}(T')$. In other words, $\mathcal{M}^{QE(\Delta)}(T) = \mathcal{M}(T')$.

Proof. Let $M \in \mathcal{M}(T)$. Then we know that $M^{\mathrm{QE}(\Delta)} \models T'$, and we need to show that $M^{\mathrm{QE}(\Delta)}$ is e.c. as such. So let $N \models T'$ and $f : M^{\mathrm{QE}(\Delta)} \to N$ be a $\Delta_0^{\mathrm{QE}(\Delta)}$-homomorphism, and we need to show that is it a $\Sigma_1^{\mathrm{QE}(\Delta)}$-embedding. We may assume that $N \in \mathcal{M}(T')$, whereby $N = (N{\restriction}_{\mathcal{L}})^{\mathrm{QE}(\Delta)}$, so $f{\restriction}_{\mathcal{L}}: M \to N{\restriction}_{\mathcal{L}}$ is a Δ-homomorphism, and $N{\restriction}_{\mathcal{L}} \models T$. Since M is an e.c. model of T, $f{\restriction}_{\mathcal{L}}$ is a $\Sigma(\Delta)$-embedding, and f is a $\Sigma_1^{\mathrm{QE}(\Delta)}$-embedding as required:

$$M^{\mathrm{QE}(\Delta)} \models \exists \bar{y}\, R_\varphi(\bar{a}, \bar{y}) \iff M \models \exists \bar{y}\, \varphi(\bar{a}, \bar{y})$$
$$\iff N{\restriction}_{\mathcal{L}} \models \exists \bar{y}\, \varphi(f(\bar{a}), \bar{y}) \iff N \models \exists \bar{y}\, R_\varphi(f(\bar{a}), \bar{y})$$

So $M^{\mathrm{QE}(\Delta)} \in \mathcal{M}(T')$.
This shows that $-^{\mathrm{QE}(\Delta)}$ is a functor from $\mathcal{M}(T)$ to $\mathcal{M}(T')$. In order to show that it is an isomorphism, we still need to show that every $M \in \mathcal{M}(T')$ is in the image. We assumed that if $M \in \mathcal{M}(T')$ then $M = (M{\restriction}_{\mathcal{L}})^{\mathrm{QE}(\Delta)}$, so we only need to prove that $M{\restriction}_{\mathcal{L}} \in \mathcal{M}(T)$, with the same argument: let N be a model of T and $f : M{\restriction}_{\mathcal{L}} \to N$ a Δ-homomorphism. In order to show that f is a $\Sigma(\Delta)$-embedding, we may assume that N

is e.c., whereby $N^{\mathrm{QE}(\Delta)} \models T'$ and $f^{\mathrm{QE}(\Delta)} : M \to N^{\mathrm{QE}(\Delta)}$ is a $\Delta_0^{\mathrm{QE}(\Delta)}$-homomorphism, whereby it is a $\Sigma_1^{\mathrm{QE}(\Delta)}$-embedding and f is a $\Sigma(\Delta)$-embedding. QED

Lemma 1.31. *Let $M \in \mathcal{M}(T^{QE(\Delta)})$. Then $M = (M{\restriction}_{\mathcal{L}})^{QE(\Delta)}$.*

Proof. We prove by induction on the structure of a formula $\varphi \in \Delta$, that $M \models \varphi(\bar{a}) \iff M \models R_\varphi(\bar{a})$ for any $\bar{a} \in M$ (of the right length) applying Lemma 1.26 to $T^{\mathrm{QE}(\Delta)}$. Of course, since we made no particular assumptions about T beyond its being a $\Pi(\Delta)$-theory, we do not know exactly what $T^{\mathrm{QE}(\Delta)}$ would forbid. However, since $T^{\mathrm{QE}(\Delta)} = \mathrm{Th}_{\Pi_1^{\mathrm{QE}(\Delta)}}(\mathcal{M}^{\mathrm{QE}(\Delta)}(T))$, we know that $T^{\mathrm{QE}(\Delta)}$ forbids something if and only if it never happens in $\mathcal{M}^{\mathrm{QE}(\Delta)}(T)$, and then we use what we know about the interpretation of R_φ in $\mathcal{M}^{\mathrm{QE}(\Delta)}(T)$:

- $\varphi = R(x_{<n})$ where $R \in \mathcal{L}$: According to $T^{\mathrm{QE}(\Delta)}$, an instance of a formula forbids $\varphi(\bar{a})$ if and only if it forbids $R_\varphi(\bar{a})$. Thus $\varphi(\bar{a})$ and $R_\varphi(\bar{a})$ are true of false together.
- $\varphi(x_{\in I}) = f_*(\psi(x_{\in J}))$: An instance forbids $R_\varphi(\bar{a})$ if and only if it forbids $R_\psi(f^*(\bar{a}))$.
- $\varphi(\bar{x}) = \bigvee_{i<k} \psi_i(\bar{x})$: An instance forbids $R_\varphi(\bar{a})$ if and only if it forbids $R_{\psi_i}(\bar{a})$ for every $i < k$.
- $\varphi(\bar{x}) = \bigwedge_{i<k} \psi_i(\bar{x})$: If $k = 0$, nothing can forbid R_φ, so we may assume that $k > 0$. Then an instance of a formula forbidding and $R_{\psi_i}(\bar{a})$ would forbid $R_\varphi(\bar{a})$. Conversely, assume that $R_\varphi(\bar{a})$ is false, say forbidden by $\chi(\bar{b})$, but $\bigwedge_{i<k-1} R_{\psi_i}(\bar{a})$ is true: then $\chi(\bar{b}) \wedge \bigwedge_{i<k-1} R_{\psi_i}(\bar{a})$ forbids $R_{\psi_{k-1}}(\bar{a})$.
- $\varphi(\bar{x}) = \exists \bar{y}\, \psi(\bar{x}, \bar{y})$: An instance forbids $R_\varphi(\bar{a})$ if and only if it forbids $R_\psi(\bar{a}, \bar{y})$.
- $\varphi(\bar{x}) = \neg\psi(\bar{x})$: On one hand, $R_\varphi(\bar{a})$ forbids $R_\psi(\bar{a})$. On the other, every instance which forbids $R_\varphi(\bar{a})$ forbids every instance which forbids $R_\psi(\bar{a})$, so precisely one of the two must be true.

QED

Proposition 1.32. *The functor $M \mapsto M^{QE(\Delta)}$ is an isomorphism $\mathcal{M}(T) \simeq \mathcal{M}(T^{QE(\Delta)})$.*

Proof. By Lemma 1.31, we have $M = (M{\restriction}_{\mathcal{L}})^{\mathrm{QE}(\Delta)}$ for every $M \in \mathcal{M}(T^{\mathrm{QE}(\Delta)})$. By definition of $T^{\mathrm{QE}(\Delta)}$ we have $M^{\mathrm{QE}(\Delta)} \models T^{\mathrm{QE}(\Delta)}$ for every $M \in \mathcal{M}(T)$. Finally, if $M \in \mathcal{M}(^{\mathrm{QE}(\Delta)})$ then necessarily $M{\restriction}_{\mathcal{L}} \models T$, using the fact that $M = (M{\restriction}_{\mathcal{L}})^{\mathrm{QE}(\Delta)}$ and that modulo the translation between Δ and $\mathcal{L}^{\mathrm{QE}(\Delta)}$, everything forbidden by T must be forbidden by $T^{\mathrm{QE}(\Delta)}$ as well.
Now apply Lemma 1.30. QED

And conversely:

Proposition 1.33. *Let Δ be a positive fragment, and T a $\Pi_1^{QE(\Delta)}$-theory, such that $M = (M{\restriction}_{\mathcal{L}})^{QE(\Delta)}$ for every $M \in \mathcal{M}(T)$ (namely, $\varphi(M) = R_\varphi^M$ for every $\varphi \in \Delta$). Let $T_0 = \mathrm{Th}_{\Pi(\Delta)}(\mathcal{M}(T){\restriction}_{\mathcal{L}})$. Then T is equivalent to $T_0^{\mathrm{QE}(\Delta)}$.*

Proof. First, if $M \models T_0$ then $M^{\mathrm{QE}(\Delta)} \models T$: modulo the translation between Δ and $\mathcal{L}^{\mathrm{QE}(\Delta)}$, everything forbidden by T must be forbidden by T_0 as well. Also, if $M \in \mathcal{M}(T)$ then $M{\restriction}_{\mathcal{L}} \models T_0$ by definition and $M = (M{\restriction}_{\mathcal{L}})^{\mathrm{QE}(\Delta)}$ by assumption. Then by Lemma 1.30 and Proposition 1.32: $\mathcal{M}(T) = \mathcal{M}^{\mathrm{QE}(\Delta)}(T_0) = \mathcal{M}(T_0^{\mathrm{QE}(\Delta)})$. It is easy to see that if two $\Pi_1^{\mathrm{QE}(\Delta)}$-theories have the same e.c. models then they have the same models. QED

Proposition 1.32 tells us that working in $\mathcal{M}(T)$ is equivalent to working in $\mathcal{M}(T^{\mathrm{QE}(\Delta)})$. But $\mathcal{M}^0(T)$ and $\mathcal{M}^0(T^{\mathrm{QE}(\Delta)})$ can be very different: $\mathcal{M}^0(T^{\mathrm{QE}(\Delta)})$ is always very simple, as is shown for example by the easy proof of Lemma 1.25. But then we see the power of Proposition 1.32 which allows us to reduce the general case to the Δ_0 one:

Theorem 1.34 (Compactness Theorem)**.** *Let $\Phi(\bar{x})$ be a set of formulas (of $\mathcal{L}_{\omega,\omega}$, without limitations), such that every finite $\Phi_0 \subseteq \Phi$ is consistent. Then Φ is consistent.*

Proof. Set $\Delta = \mathcal{L}_{\omega,\omega}$ and write $\mathcal{L}^{\mathrm{QE}} = \mathcal{L}^{\mathrm{QE}(\mathcal{L}_{\omega,\omega})}$, $\Delta_0^{\mathrm{QE}} = \Delta_0^{\mathrm{QE}(\mathcal{L}_{\omega,\omega})}$, etc. Take $T = \varnothing$, so $T^{\mathrm{QE}} = \mathrm{Th}_{\Pi_1^{\mathrm{QE}}}(\mathcal{M}^{\mathrm{QE}}(\varnothing))$.
Set $\Phi^{\mathrm{QE}} = \{R_\varphi(\bar{x}) : \varphi(\bar{x}) \in \Phi\}$. Then every finite $\Phi_0^{\mathrm{QE}} \subseteq \Phi^{\mathrm{QE}}$ is consistent with T^{QE}, thus by Lemma 1.25 it is satisfied by some tuple $\bar{a}$ in a model $M \models T^{\mathrm{QE}}$. By Lemma 1.20, there exists an e.c. model $N \models T^{\mathrm{QE}}$ and a morphism $f : M \to N$: then $N \models \Phi^{\mathrm{QE}}(f(\bar{a}))$. Finally, by Lemma 1.31, $N = (N{\restriction}_{\mathcal{L}})^{\mathrm{QE}(\Delta)}$, so $N \models \Phi(f(\bar{a}))$. QED

As a last remark, if $\Delta = \mathcal{L}_{\omega,\omega}$, then every structure is e.c. (as a model of any theory for which it is indeed a model). Thus first order model theory is a special case of positive model theory.

1.3. **Types and type-spaces.**

Convention 1.35. Δ is a positive fragment, $\Sigma = \Sigma(\Delta)$ and $\Pi = \Pi(\Delta)$ as usual, and T is a Π-theory. $\mathcal{M} = \mathcal{M}(T)$, the category of e.c. models of T.

We want to give a full description of tuples in models of T. We obtain equivalent definitions:

Definition 1.36. Let $M \in \mathcal{M}$, $a \in M^I$, and Ξ a positive fragment which is not necessarily equal to Δ: it will usually be either Δ or Σ. Then the Ξ-*type* of a (in M) is $\mathrm{tp}_\Xi^M(a) = \{\varphi(x_{\in I}) : \varphi \in \Xi, M \models \varphi(a)\}$.
The *type* of a (in M) is its Σ-type: $\mathrm{tp}^M(a) = \mathrm{tp}_\Sigma^M(a)$, *and not its Δ-type!* (at least for the time being).

In the same manner that we may write $\models \varphi(a)$ instead of $M \models \varphi(a)$, we may omit the model M: if $f : M \to N$ is an inclusion morphism between two e.c. models and $a \in M^I$, then $\mathrm{tp}^M(a) = \mathrm{tp}^N(a)$.

Lemma 1.37. *Let I be a set of indices, and $M_i \in \mathcal{M}$, $a_i \in M^I$ for $i < 2$. Then the following are equivalent:*

1. *There is $N \in \mathcal{M}$ and morphisms $f_i : M_i \to N$ for $i < 2$ such that $f_0(a_0) = f_1(a_1)$.*
2. $\mathrm{tp}^{M_0}(a_0) = \mathrm{tp}^{M_1}(a_1)$.
3. $\mathrm{tp}^{M_0}(a_0) \subseteq \mathrm{tp}^{M_1}(a_1)$.

Proof. **1 $\Longrightarrow$ 2:** If N, f_i exist then $M_0 \models \varphi(a_0) \iff N \models \varphi(f_i(a_i)) \iff M_1 \models \varphi(a_1)$ for every $\varphi \in \Sigma(\Delta)$, by e.c..
2 $\Longrightarrow$ 3: Clear.

$3 \Longrightarrow 1$**:** Assume that $M_0 \models \varphi(a_0) \Longrightarrow M_1 \models \varphi(a_1)$ for every $\varphi \in \Sigma(\Delta)$. Let $M_i = \{b_j : j \in J_i\}$ with $J_0 \cap J_1 = \varnothing$ and let $g_i : I \to J_i$ be such that $a_i = g_i^*(\bar{b}_i)$. Write $p_i(x_{\in J_i}) = \mathrm{tp}^{M_i}(b_{\in J_i})$ for $i < 2$ and $\Phi(x_{\in J_0 \cup J_1}) = p_0 \cup p_1 \cup g_0^*(x_{\in J_0}) = g_1^*(x_{\in J_1})$. Since each p_i is closed for finite conjunctions and $\mathrm{tp}^{M_0}(a_0) \subseteq \mathrm{tp}^{M_1}(a_1)$, we see that Φ is finitely satisfiable in M_1, and therefore is consistent with T by the compactness theorem. Let it be satisfied in $N \models T$, say by $c_{\in J_0 \cup J_1}$. We may assume that N is e.c., that is $N \in \mathcal{M}$. Let $f_i : M_i \to N$ be defined by $f_i(b_{\in J_i}) = c_{\in J_i}$, and then $f_0(a_0) = f_1(a_1)$ as required.

QED

Definition 1.38. Let I be a set of indices, and let $\mathrm{S}_I(T) = \{\mathrm{tp}(a) : M \in \mathcal{M}, a \in M^I\}$. For $\varphi(x_{\in I}) \in \Sigma$, let $\langle \varphi \rangle = \{p \in \mathrm{S}_I(T) : \varphi \in p\}$.
As we did for the definition of types, let Ξ be a positive fragment contained in Σ, usually either Δ or Σ. Then the Ξ-topology on $\mathrm{S}_I(T)$ is the one generated by $\{\langle \varphi \rangle : \varphi(x_{\in I}) \in \Xi\}$ as closed sets. Take for the standard topology on $\mathrm{S}_I(T)$ the Σ-topology. We call $\mathrm{S}_I(T)$ with the Σ-topology the *space of (complete pure) I-types* of T.
If $f : I \to J$ then $f^* : \mathrm{S}_J(T) \to \mathrm{S}_I(T)$ is defined by $f^*(p) = \{\varphi(x_{i \in I}) : f_*(\varphi) \in p\} = f_*^{-1}(p)$. We also define $f_* : \mathcal{P}(\mathrm{S}_I(T)) \to \mathcal{P}(\mathrm{S}_J(T))$ by $f_*(A) = f^{*-1}(A)$ ($\mathcal{P}$ denotes the power set).

Lemma 1.39. *Let $f : I \to J$ be a map, $M \in \mathcal{M}$.*

1. $f^*(\mathrm{tp}(a)) = \mathrm{tp}(f^*(a))$ *for every* $a \in M^J$.
2. $f_*(\langle \varphi \rangle) = \langle f_*(\varphi) \rangle \subseteq \mathrm{S}_J(T)$ *for every* $\varphi(x_{\in I}) \in \Sigma$.
3. $f^*(\langle \varphi \rangle) = \langle \exists y_{\in J}\, x_{\in I} = f^*(y_{\in J}) \wedge \varphi(y_{\in J}) \rangle \subseteq \mathrm{S}_I(T)$ *for every* $\varphi(x_{\in J}) \in \Sigma$ *(with our definitions, this is only meaningful for finite $I, J \subseteq \omega$, but this can be given a natural meaning for arbitrary I, J in which case it also holds).*

Proof. 1. $\varphi \in f^*(\mathrm{tp}(a)) \Longleftrightarrow f_*(\varphi) \in \mathrm{tp}(a) \Longleftrightarrow M \models \varphi(f^*(a)) \Longleftrightarrow \varphi \in \mathrm{tp}(f^*(a))$ for every $\varphi(x_{\in I}) \in \Sigma$.

2. $p \in f_*(\langle \varphi \rangle) \Longleftrightarrow \varphi \in f^*(p) \Longleftrightarrow f_*(\varphi) \in p$ for every $p \in \mathrm{S}_J(T)$.
3. By definition, $p \in f^*(\langle \varphi \rangle)$ if and only if there is $q \in \langle \varphi \rangle$ such that $p = f^*(q)$. This is equivalent to the existence of $b_{\in J}$ in some $N \in \mathcal{M}$, such that $b_{\in J} \models \varphi$ and $f^*(b_{\in J}) \models p$. If such $b_{\in J}$ exist, then $N \models \exists y_{\in J}\, f^*(b_{\in J}) = f^*(y_{\in J}) \wedge \varphi(y_{\in J})$ implies $p \in \langle \exists y_{\in J}\, x_{\in I} = f^*(y_{\in J}) \wedge \varphi(y_{\in J}) \rangle$. Conversely, if $p \in \langle \exists y_{\in J}\, x_{\in I} = f^*(y_{\in J}) \wedge \varphi(y_{\in J}) \rangle$ and p is realised by $c_{\in I}$ in N, then there are $b_{\in J} \models \varphi$ with $f^*(b_{\in J}) = c_{\in I}$ as required.

QED

Proposition 1.40. *With the given topology, every $\mathrm{S}_I(T)$ is a compact T_1 space, and every f^* is a continuous closed map. Thus, if $f : I \to J$, then f^* sends closed sets of $\mathrm{S}_J(T)$ to closed sets of $\mathrm{S}_I(T)$, and f_* sends closed sets of $\mathrm{S}_I(T)$ to closed sets of $\mathrm{S}_J(T)$. Moreover, let $S(I) = \mathrm{S}_I(T)$ and $S(f) = f^*$. Then S is a contravariant functor. We just note it by $\mathrm{S}(T)$, and call it the* type-space functor *of T.*

Proof. These are easy verifications. For the properties of f^*, use Lemma 1.39. QED

Finally, we can consider certain properties of the topology on $\mathrm{S}(T)$ as properties of T: the importance of these properties will become clearer later on.

Definition 1.41. 1. T is *Hausdorff* if every type-space of T is.

2. T is *semi-Hausdorff* if for every n, the set $\{\text{tp}^M(a,b) : M \in \mathcal{M}, a, b \in M^n, \text{tp}(a) = \text{tp}(b)\}$ is closed in $\mathrm{S}_{2n}(T)$ (we say that equality of types is definable by a partial type, that is by an infinite conjunction).

Remark 1.42. Let $\pi_n : \mathrm{S}_{2n}(T) \to \mathrm{S}_n(T)^2$ be the obvious map (just send $\text{tp}(\bar a, \bar b)$ to $(\text{tp}(\bar a), \text{tp}(\bar b))$), and let $D(\mathrm{S}_n(T)) = \{(p,p) : p \in \mathrm{S}_n(T)\} \subseteq \mathrm{S}_n(T)^2$ be the diagonal. Then T is Hausdorff if and only if every $D(\mathrm{S}_n(T))$ is closed, and it is semi-Hausdorff if and only if every $\pi_n^{-1}(D(\mathrm{S}_n(T)))$ is closed. Thus Hausdorff implies semi-Hausdorff.

1.4. **Complete theories and completions.**

Definition 1.43. 1. A Π-theory T is *complete* if it is of the form $\text{Th}_\Pi(M)$ for some M.
2. A *completion* of a theory T is a minimal complete theory containing T.

Proposition 1.44. *The following are equivalent for a consistent* Π*-theory* T*:*

1. T *is complete.*
2. *Whenever* T *implies the disjunction of two* Π*-sentences, it implies at least one of them.*
3. $|\,\mathrm{S}_0(T)| = 1$
4. $T = \text{Th}_\Pi(M)$ *for every* $M \in \mathcal{M}(T)$*.*
5. $T = \text{Th}_\Pi(M)$ *for some* $M \in \mathcal{M}(T)$*.*

Proof. $1 \Longrightarrow 2$**:** Clear.

$2 \Longrightarrow 3$**:** Assume that $p, q \in \mathrm{S}_0(T)$ are different. Then they are incomparable, and by a compactness argument we find Σ-sentences $\varphi \in p$ and $\psi \in q$ such that $T \cup \{\varphi, \psi\}$ is inconsistent. This means that $T \vdash \neg\varphi \vee \neg\psi$, whereby by hypothesis $T \vdash \neg\varphi$ or $T \vdash \neg\psi$, contradicting the assumption that p and q were types.

$3 \Longrightarrow 4$**:** Let $M \in \mathcal{M}(T)$. Then $M \models T$ so $T \subseteq \text{Th}_\Pi(M)$. For the converse, assume that $T \subsetneq \text{Th}_\Pi(M)$, so there is a Σ-sentence φ such that $M \models \neg\varphi$ but $T \cup \{\varphi\}$ are consistent. Then we can find $N \models T$ where φ holds, and we can continue N to some $N' \in \mathcal{M}$, so $N' \models \varphi$ as well. Then the types of the empty tuples in M and N' are different, contradicting $|\,\mathrm{S}_0(T)| = 1$.

$4 \Longrightarrow 5 \Longrightarrow 1$**:** Clear.

QED

Corollary 1.45. 1. *The completions of* T *are precisely the theories of e.c. models of* T*.*
2. *If* $f : M \to N$ *is a morphism between two e.c. models of* T *then their theories are equal:* $\text{Th}_\Pi(M) = \text{Th}_\Pi(N)$*.*

Proof. 1. If T' is a completion of T, it is complete, that is $T' = \text{Th}_\Pi(M)$ for some M. Then $M \models T$ so we can continue M to an e.c. model M' of T. We obtain: $T \subseteq \text{Th}_\Pi(M') \subseteq T'$, so $T' = \text{Th}_\Pi(M)$ by its minimality.

Conversely, Let M be an e.c. model of T, and $T' = \text{Th}_\Pi(M)$. Then T' is complete. Suppose that $T \subseteq T'' \subseteq T'$ and T'' is complete as well. Then M is *a fortiori* e.c. as a model of T'', thus $T'' = \text{Th}_\Pi(M) = T'$ as T'' is complete. This shows the minimality of T'.

2. In such a case we have that $\mathrm{Th}_\Pi(M) \supseteq \mathrm{Th}_\Pi(N)$ are both completions of T, and by minimality of completion we have equality.

QED

Corollary 1.46. *For each $p \in \mathrm{S}_0(T)$, let $\mathcal{M}_p(T) = \{M \in \mathcal{M}(T) : \mathrm{tp}^M(\varnothing) = p\}$ and $T_p = \mathrm{Th}_\Pi(\mathcal{M}_p(T))$. Then $\mathcal{M}(T_p) = \mathcal{M}_p(T)$, $\mathrm{S}_0(T_p) = \{p\}$, and the map $c : p \mapsto T_p$ is a bijection between $\mathrm{S}_0(T)$ and the completions of T.*

Proof. Assume that $M, M' \in \mathcal{M}_p(T)$ for some $p \in \mathrm{S}_0(T)$. Then by Lemma 1.37, there is $N \in \mathcal{M}(T)$ that continues both M and M', and by Corollary 1.45 we have $\mathrm{Th}_\Pi(M) = \mathrm{Th}_\Pi(N) = \mathrm{Th}_\Pi(M')$. This shows that $T_p = \mathrm{Th}_\Pi(\mathcal{M}_p(T)) = \mathrm{Th}_\Pi(M)$, so in particular $\mathrm{Th}_\Pi(\mathcal{M}_p(T))$ is a completion of T, and the map $c : \mathrm{S}_0(T) \to \{\text{completions of } T\}$ is well defined.
If T' is a completion of T, then $T' = \mathrm{Th}_\Pi(M)$ for some $M \in \mathcal{M}(T)$, and let $p = \mathrm{tp}^M(\varnothing)$. Then by the above, $T' = \mathrm{Th}_\Pi(M) = T_p$, so c is surjective. Finally, if $p, q \in \mathrm{S}_0(T)$ are distinct, then there is a Σ-sentence $\varphi \in p \smallsetminus q$, so $\neg\varphi \in T_q$ but $\neg\varphi \notin T_p$, whereby $T_p \neq T_q$ and c is surjective.
If $M \in \mathcal{M}_p(T)$, then it is an e.c. model of T and therefore *a fortiori* of $\mathrm{Th}_\Pi(M) = T_p$, so $M \in \mathcal{M}(T_p)$. Conversely, assume that $M \in \mathcal{M}(T_p)$, and we want to show that $M \in \mathcal{M}(T)$. Since T_p is complete we have $\mathrm{Th}_\Pi(M) = T_p$. Let $f : M \to N$ be a morphism with $N \models T$, and we need to show that it is a Σ-embedding. We may assume that $N \in \mathcal{M}(T)$, whereby $\mathrm{Th}_\Pi(N) \subseteq T_p$ are both completions of T and therefore equal. Since $M \in \mathcal{M}(T_p)$, f is a Σ-embedding, which shows that $M \in \mathcal{M}(T)$. Since every $q \in \mathrm{S}_0(T)$ different from p contradicts T_p, we must actually have $M \in \mathcal{M}_p(T)$. QED

Corollary 1.47. *Every complete theory containing T contains a completion of T.*

Proof. Assume that $T' = \mathrm{Th}_\Pi(M)$ and $T \subseteq T'$. Then $M \models T$, and we can continue it to $N \in \mathcal{M}(T)$: $\mathrm{Th}_\Pi(N)$ is a completion of T and $\mathrm{Th}_\Pi(N) \subseteq T'$. QED

Example 1.48. Set $\Delta = \mathcal{L}_{\omega,\omega}$ and $T = \varnothing$. As a set, the completions of T are all the complete first order theories in the language $\mathcal{L}$, and topology on $\mathrm{S}_0(T)$ is the classical Stone topology on this set.

2. Cats

We are going to define three objects of very different nature, that turn out to be manifestations of the same concept, which we call a *compact abstract theory*, or a *cat*.

2.1. Positive Robinson theories.

In various contexts of model theory, assuming quantifier elimination is useful as a starting point, and in abstract contexts does not cost much. In our context, this will be the assumption that Δ has the same power of expression as Σ:

Definition 2.1. A Π-theory T is a *positive Robinson theory* if whenever $M_i \in \mathcal{M}(T)$ and $a \in M_i^I$ for $i < 2$ and $\mathrm{tp}_\Delta(a_0) \subseteq \mathrm{tp}_\Delta(a_1)$ then $\mathrm{tp}(a_0) = \mathrm{tp}(a_1)$ (that is, $\mathrm{tp}_\Sigma(a_0) = \mathrm{tp}_\Sigma(a_1)$).
In other words, a Δ-type determines a Σ-type.

In particular, if $\Delta = \Sigma(\Delta)$ then every theory is positive Robinson.

This definition generalises that of Hrushovski from [Hru97], which considers the case where Δ is closed for negations as well. We refer to Hrushovski's definition as *classical* Robinson theories. Much of what follows is adaptation of results in [Hru97] to our context.

Convention 2.2. We want to restrict ourselves to what can be said using Δ alone. Therefore from now on, every formula is in Δ, unless explicitly said otherwise. If we want to consider a Σ-formula we will call it so and use the existential quantifier explicitly.

Proposition 2.3. *For a theory T, the following are equivalent:*

1. *T is positive Robinson.*
2. *Definition 2.1 restricted to* finite *tuples.*
3. *The class $\mathcal{M}$ of e.c. models of T is axiomatised by the set of all the sentences of the form:*

$$\forall x\, [\exists y\, \varphi(x,y) \longleftrightarrow \bigwedge\{\neg\psi(x) : T \vdash \neg\exists x,y\, \psi(x) \wedge \varphi(x,y)\}]$$

In other words, an $\mathcal{L}$-structure M is an e.c. model of T if and only if, for every $\varphi(x,y)$ and tuple $a \in M$ of the same length as x, $M \models \exists y\, \varphi(a,y)$ if and only if $M \not\models \psi(a)$ whenever according to T, $\psi(a)$ forbids $\varphi(a,y)$.
Informally, we would say that something exists in M if and only if there is no reason for it not to exist.

4. *Whenever $T \vdash \neg \bigwedge_{i<\lambda} \varphi_i(x,y_i)$ (namely, the set $\{\varphi_i(x,y_i) : i < \lambda\}$ is inconsistent with T), then there exist infinitary disjunctions $\Psi_i(x)$ such that $\mathcal{M} \models \bigvee \Psi_i$ (meaning that each Ψ_i is a set of formulas $\psi(x)$, and if $M \in \mathcal{M}$ and $a \in M$ is of the right length, then $M \models \psi(a)$ for some $\psi \in \bigcup \Psi_i$) and for all i and $\psi \in \Psi_i$: $T \vdash \neg\exists x, y_i\, \psi(x) \wedge \varphi_i(x,y_i)$.*
5. *As above, for x finite and $\lambda = 2$.*

Proof. $1 \Longrightarrow 2$**:** Clear.

$2 \Longrightarrow 3$**:** First, we show that these sentence are all true in $\mathcal{M}$: The $\to$ part is clear. As for the other direction, let $M \in \mathcal{M}$, $M \models \neg\exists y\, \varphi(a,y)$, and write $p(x) = \mathrm{tp}_\Delta(a)$. Then $T \cup p(x) \cup \varphi(x,y)$ is contradictory: if not, there would be $M' \models T$ with $a', b \in M'$ satisfying $p(x) \cup \varphi(x,y)$, so by the assumption and Lemma 1.37 there is N and morphisms $f : M \to N$, $f' : M' \to N$, and $f(a) = f'(a')$. But then $M' \models \exists y\, \varphi(f(a),y)$, contradicting the assumption on a and the fact that M was e.c.. Thus, there is some ψ such that $T \vdash \neg\exists x,y\, \psi(x) \wedge \varphi(x,y)$, and $M \models \psi(a)$. Thus we have shown that M satisfies all these axioms.

Conversely, suppose that M satisfies them all. Then M satisfies T, for if $\neg\exists y\, \varphi(y) \in T$, then $T \vdash \neg\exists y \top \wedge \varphi(y)$, thus we get the axiom $\exists y\, \varphi(y) \leftrightarrow \bot$, that is $\neg\exists y\, \varphi(y)$ (if one does not like this, one can always add T to the list of axioms). Now, the fact that M is e.c. is evident: if $f : M \to N$ is a morphism and $N \models \exists y\, \varphi(f(a),y)$, then all that contradicts it is false in N, therefore false in M, wherefore $M \models \exists y\, \varphi(a,y)$ as well.

$3 \Longrightarrow 4$: Note that by compactness we can easily reduce to the case where $\lambda < \omega$, but this is unnecessary. Let $\Psi_i(x) = \bigvee\{\psi(x) : T \vdash \neg\exists x, y\, \psi(x) \wedge \varphi_i(x, y_i)\}$. Then, by the assumption, they satisfy the conditions.

$4 \Longrightarrow 5$: Clear.

$5 \Longrightarrow 1$: Let $M_i \in \mathcal{M}$ and $a_i \in M_i^I$ for $i < 2$, and assume that $\mathrm{tp}_\Delta(a_0) = \mathrm{tp}_\Delta(a_1)$. We need to show that $\mathrm{tp}_\Sigma(a_0) = \mathrm{tp}_\Sigma(a_1)$. Enumerate $\mathrm{tp}_\Sigma(a_i)$ as $\{\exists y_j\, \varphi_j(x, y_j) : j \in J_i\}$, and let $\Phi_i(x, y_{\in J_i}) = \{\varphi_j(x, y_j) : j \in J_j\}$, where $J_0 \cap J_1 = \varnothing$ and all the y_j are distinct tuples of variables.
If $\Phi_0 \cup \Phi_1$ is consistent with T, let it be realised in some model N by $a, b_{\in J_0 \cup J_1}$: then $\mathrm{tp}_\Sigma(a_0), \mathrm{tp}_\Sigma(a_1) \subseteq \mathrm{tp}_\Sigma(a)$, whereby they are all equal by Lemma 1.37.
So assume that $\Phi_0 \cup \Phi_1$ is inconsistent with T: since each Φ_i is closed for conjunction, and by compactness, there are $\varphi_i(x, y_i) \in \Phi_i$ such that $T \vdash \neg\exists x, y_0, y_1\, \varphi_0(x, y_0) \wedge \varphi_1(x, y_1)$. Take Ψ_0, Ψ_1 as in the hypothesis. Then $M_0 \models \Psi_0(a_0) \vee \Psi_1(a_0)$, so $M_0 \models \psi(a_0)$ with $T \vdash \neg\exists x, y_i\, \psi(x) \wedge \varphi_i(x, y_i)$ for some $i < 2$ and for some $\psi \in \Psi_i$, whence also $M_1 \models \psi(a_1)$. But this contradicts $M_i \models \exists y_i\, \varphi_i(a_i, y_i)$.

QED

Remark 2.4. If T is positive Robinson, $M, N \models T$, M e.c. and $f : M \to N$ a partial Δ-homomorphism, then f is a partial isomorphism. Indeed, by Proposition 2.3 we can find $M' \models T$ and morphisms $g : M \to M'$ and $h : N \to M'$ such that $g = h \circ f$. Then any relation, including equality, true in N for elements of the domain of f is true in M' and therefore in M (as M is e.c.).

We defined a positive Robinson theory as one where complete Σ-types are equivalent to Δ-types. However, what about partial types?

Definition 2.5. A Π-theory T is *strongly positive Robinson* if every partial Σ-type is equivalent to a partial Δ-type.

Proposition 2.6. *Let T be a Π-theory. Then the following are equivalent:*

1. *T is strongly positive Robinson.*
2. *The Δ-topology on the type-spaces coincides with the standard one, that is the Σ-topology.*
3. *For every $f : m \to n$, the map $f^* : \mathrm{S}_n(T) \to \mathrm{S}_m(T)$ is closed in the Δ-topology.*

Proof. $1 \iff 2$ is by definition, as a closed set in the Ξ-topology corresponds to a partial Ξ-type, for $\Xi \in \{\Delta, \Sigma\}$. $2 \iff 3$ is just the definition of Σ as the closure of Δ for the existential quantifier and Lemma 1.39. QED

Remark 2.7. When considering the general case, there is no loss of generality assuming that T is strongly positive Robinson, even when $\Delta = \Delta_0$: if not, replace Δ by Σ and perform a positive Morleyisation, and $T^{\mathrm{QE}(\Sigma)}$ is strongly positive Robinson with respect to $\Delta_0^{\mathrm{QE}(\Sigma)}$.

Remark 2.8. If the Δ-topology is Hausdorff, then T is clearly strongly positive Robinson and Hausdorff, since a topology strictly stronger than a Hausdorff topology cannot be compact.
In particular, any classical Robinson theory is strongly positive Robinson.

The Hausdorff case is somewhat closer to [Hru97] than the general case, as we have finitary quantifier separation:

Fact 2.9. *Let X be a Hausdorff space, $\{K_i\}$ compact sets, and $\bigcap K_i = \varnothing$. Then there are open sets $U_i \supseteq K_i$ with $\bigcap U_i = \varnothing$. Moreover, if we have a basis for the open sets that is closed for finite union, we can take U_i to be of this basis.*

Proposition 2.10. *For a theory T, the following are equivalent:*

1. *The Δ-topology on the type spaces of T is Hausdorff (so in particular T is strongly positive Robinson).*
2. *Whenever $T \vdash \neg \bigwedge_{i<\lambda} \varphi_i(x, y_i)$ then there exist $\psi_i(x)$ such that $\mathcal{M} \models \bigvee \psi_i$ and for all i: $T \vdash \neg\exists x, y\, \psi_i(x) \wedge \varphi_i(x, y_i)$.*
3. *As above, for x finite and $\lambda = 2$.*

Proof. $1 \Longrightarrow 2$**:** Set $K_i = \langle \exists y_i\, \varphi_i(x, y_i)\rangle \subseteq \mathrm{S}_n(T)$, where x is an n-tuple. By Lemma 1.39, K_i is the image of $\langle\varphi_i\rangle$ under a continuous map, and therefore compact in the Δ-topology. The hypothesis also says that $\bigcap_{i<\lambda} K_i = \varnothing$, so we may apply Fact 2.9 to obtain open $U_i \supseteq K_i$ such that $\bigcap_{i<\lambda} U_i = \varnothing$, and by the moreover part we may assume that each $U_i = \langle\neg\psi_i\rangle$ for some formula ψ_i. Then $U_i \supseteq K_i \Longrightarrow T \vdash \neg\exists x, y\, \psi_i(x) \wedge \varphi_i(x, y_i)$, and $\bigcap_{i<\lambda} U_i = \varnothing \Longrightarrow \mathcal{M} \models \bigvee \psi_i$.

$2 \Longrightarrow 3$**:** Clear.

$3 \Longrightarrow 1$**:** Two distinct type necessarily contain two contradictory Σ-formulas. By the assumption, they are separated by two open sets in the Δ-topology, whereby it is Hausdorff.

QED

We also adapt from [Hru97] the definition of universal domains (leaving the notion of smallness a bit vague, as required by the tradition):

Definition 2.11.
1. A structure M is *homogeneous* if whenever $f : M \to M$ is a partial endomorphism with a small domain, f extends to an automorphism.
2. A structure M is *compact* if any small set of formulas over M in possibly infinitely many variables, which is finitely realised in M, is realised in M.
3. A structure M is a *universal domain* (for $\mathrm{Th}_\Pi(M)$) if it is homogeneous and compact.

This is clear:

Lemma 2.12.
1. *A universal domain is e.c..*
2. *If U is a universal domain for T and M is a small model of T, then M can be sent Δ-homomorphically into U.*
3. *Above, if $M \in \mathcal{M}$, then this Δ-homomorphism is a Δ-embedding.*

As in [Hru97], we wish to show the correspondence between complete positive Robinson theories and universal domains.

One should note first that if T is positive Robinson then every completion of T is positive Robinson: if T' is a completion of T, then any complete pure type of T' is such for T, and therefore equivalent to a Δ-type. Such properties as strong positive

Robinson, Hausdorff, semi-Hausdorff, etc., are easily verified to be also preserved under completion.

Then we prove the correspondence:

Theorem 2.13. *A complete* Π*-theory* T *is positive Robinson if and only if it has some universal domain* U*.*

Proof. **$\Longrightarrow$:** 1. As T is positive Robinson, a partial endomorphism of an e.c. model of T is a partial isomorphism, and in particular invertible. Thus one can amalgamate the model with itself over the endomorphism in both directions, thus obtaining a larger model with an endomorphism extending the original, whose domain and range contain the original model. As the original model was e.c., this endomorphism passes to an e.c. continuation. Repeating this process ω times we get an e.c. continuation on which the endomorphism extends to an automorphism. Note that being an e.c. model of T is stable for increasing unions.

2. Now, if we are given given a model M, and $\leq \lambda$ partial endomorphisms of M, we can do the same thing: we do the fundamental step for each endomorphism, and repeat this procedure ω times. Note that the endomorphisms pass to the continuations, as they were originally on e.c. models. If $\lambda \geq |M| + |T|$, the resulting model can be taken to be of cardinality λ.

3. Choose some λ, κ with $\lambda^\kappa = \lambda \geq |T|$, and some M with $\mathrm{Th}_\Pi(M) = T$ and $|M| \leq \lambda$. One can continue M to be e.c., of cardinality exactly λ. Over M there are at most λ partial types (i.e., subsets of complete types) with at most κ formulas, and at most λ partial endomorphisms whose domain is of cardinality at most κ. Thus one can extend all these to automorphisms and realise the partial types in some e.c. continuation of cardinality λ. Repeat this κ^+ times, at each step also extending automorphisms already created on the previous ones (we always have at most λ), to obtain a κ^+-universal domain for T continuing M, of cardinality λ.

$\Longleftarrow$: Say that $M, N \in \mathcal{M}$, $\bar{a} \in M$, $\bar{b} \in N$, and $\mathrm{tp}_\Delta(\bar{a}) \subseteq \mathrm{tp}_\Delta(\bar{b})$. By compactness, we can send M, N into U by morphisms. As M is e.c., it is in fact embedded into U, and the map $\bar{a} \mapsto \bar{b}$ is a partial endomorphism of U. We finish using homogeneity.

QED

We finish with a few examples.

Example 2.14. As we said before, if Δ is closed for negation then a positive Robinson theory is a classical Robinson theory, and if $\Delta = \mathcal{L}_{\omega,\omega}$ every theory is a first order theory.

Example 2.15. If $\Delta = \Sigma(\Delta)$, then every theory is strongly positive Robinson. As noted before, replacing Δ with $\Sigma(\Delta)$ does not change $\Sigma(\Delta)$, $\Pi(\Delta)$ or the notion of a morphism, so we may always assume we deal with a strongly positive Robinson theory. In particular, if Δ is the closure of Δ_0 for negation, then working in $\mathcal{M}_\Delta$ gives the context described in [Pil00], where every theory is strongly positive Robinson with respect to $\Delta' = \Sigma(\Delta)$.

Example 2.16. Let U be a universal domain for T, and $E(x, y)$ a type-definable equivalence relation on tuples of length α. The elements of U^α/E are called *hyperimaginaries*, and we wish to add them to the structure as ordinary elements (of a new sort).
For each formula $\varphi(z, x_1 \dots x_k) \in \Delta$, with each x_i of length α (of course, only finitely many variables actually appear), add a symbol R_φ of arity $|z|$ in the first sort, and k in the second, and interpret it on U' as: $R_\varphi(a, \bar{b})$ if and only if $\exists \bar{x} \in \bar{b}\, \varphi(a, \bar{x})$. Let Δ' be the positive fragment generated by Δ and atomic formulas in the new predicate symbols. One verifies that U' is a universal domain for Δ'. Moreover, if U is Hausdorff, so is U'. Equality for the sort of U' is in the Δ', but not necessarily inequality, as it is for hyperimaginaries in first-order theories.
One can similarly add several hyperimaginary sorts in one blow.

Example 2.17. Let U be a universal domain for a multi-sorted language. Let U' be the structure obtained by replacing the sorts with unary predicates. Then U' is also a universal domain. As a consequence, we obtain:
Let T be a positive Robinson theory in a multi-sorted language $\mathcal{L}$. Let $\mathcal{L}'$ be the language obtained by abolishing the sorts, adding unary predicates in their stead. Let T' be the natural image of T in $\mathcal{L}'$, with the addition of $\neg(\varphi(\bar{x}) \wedge P_s(x_i))$ whenever $\varphi \in \Delta$, $S \neq S'$ are two different sorts, and x_i would be of the sort S' in $\mathcal{L}$. Then T' is a positive Robinson theory (with Δ' being the image of Δ with the addition of the sort-predicates).

2.2. **Compact type-space functors.** We define (compact) type-space functors. These are very close to the f-spaces from [Hru97], although we use them for different ends. We show the equivalence between compact type-space functors and positive Robinson theories, namely that they are precisely the type-space functors of the those theories.

Definition 2.18.
1. We let Set be the category of sets, $FSet$ that of finite sets, and $\mathcal{T}$ that of topological spaces. We also define $\mathcal{T}_1$ as the category of T_1 topological spaces with continuous closed maps as morphisms, and $\mathcal{T}_1^c \subseteq \mathcal{T}_1$ as the full sub-category consisting of compact spaces.
2. Let $S \in \mathrm{Funct}(FSet^{op}, \mathcal{C})$, where $\mathcal{C} \in \{Set, \mathcal{T}\}$. We extend it to $\tilde{S} \in \mathrm{Funct}(Set^{op}, \mathcal{C})$ by taking inverse limits. In other words, for a set a, define $\tilde{S} = \varprojlim_{a' \in \mathcal{P}^{fin}(a)} S(a')$. Clearly, $\tilde{S}$ extends S. For $a \in Set$, we write $S_a = \tilde{S}(a)$. If $f \in \mathrm{Hom}_{Set}(a, b)$ and S is clear from the context, we write $f^* = S(f^{op}) : S_b \to S_a$. For $X \subseteq S_a$, we write $f_*(X) = (f^*)^{-1}(X) \subseteq S_b$.
3. In case that $S \in \mathrm{Funct}(FSet^{op}, \mathcal{T})$ and $\mathcal{T}' \subseteq \mathcal{T}$ is a sub-category, we say that S (resp. $\tilde{S}$) *factors through* $\mathcal{T}'$ if it factors through the inclusion $\mathcal{T}' \hookrightarrow \mathcal{T}$.
4. Let $f_{eq} : 2 \to 1$ be the only possible map, and note $eq = f_{eq}^*(S_1) \subseteq S_2$. Then S is *equality-preserving* if for every $a \in FSet$ and $h : 2 \to a$, if one notes $b = a/\{h(0) = h(1)\}$ and $f : a \to b$ the projection, then $f^*(S_b) = h_*(eq)$.
5. S has the *(finite) amalgamation property* if for every two (finite) sets a and b, the natural map $S_{a\cup b} \to S_a \times_{S_{a\cap b}} S_b$ is surjective.
6. A *type-space functor* is a functor $S : FSet^{op} \to \mathcal{C}$ which is equality-preserving and has the amalgamation property, and $\mathcal{C} \in \{Set, \mathcal{T}\}$.

S is a *set type-space functor* if $\mathcal{C} = Set$; it is a *topological type-space functor* if $\mathcal{C} = \mathcal{T}$ and $\tilde{S}$ factors through $\mathcal{T}_1$; it is a *compact type-space functor* if $\mathcal{C} = \mathcal{T}$ and $\tilde{S}$ factors through $\mathcal{T}_1^c$.
Note that we require $\tilde{S}$, and not only S, to factor through $\mathcal{T}_1$ and $\mathcal{T}_1^c$, respectively.
7. A *finite type-space functor* is a functor $S : FSet^{op} \to \mathcal{C}$ which is equality-preserving and has the finite amalgamation property, and $\mathcal{C} \in \{Set, \mathcal{T}\}$.
S is a *finite set type-space functor* if $\mathcal{C} = Set$; it is a *finite topological type-space functor* if $\mathcal{C} = \mathcal{T}$ and S factors through $\mathcal{T}_1$; it is a *finite compact type-space functor* if $\mathcal{C} = \mathcal{T}$ and S factors through $\mathcal{T}_1^c$.
In this definition we just replaced each requirement of the previous one with its finite analogue.

Remark 2.19.
1. The equality-preservation condition can be considered as an amalgamation condition. Indeed, with the notation of the definition, let $h' : 1 \to b$, $h'(0) = \{h(0), h(1)\}$, so $f \circ h = h' \circ f_{eq}$. Then b is the direct limit of $(f_{eq} : 2 \to 1, h : 2 \to a)$, and $h_*(eq)$ is the inverse limit of $(f_{eq}^* : S_1 \to S_2, h^* : S_a \to S_2)$. Then we obtain $f^* : S_b \to h_*(eq)$, and we require this to be surjective.
2. We defined type-space functors for a single sort. For multiple (say λ) sorts, one should replace Set and $FSet$ with the categories of (finite) sets coloured in λ colours, and modify the definitions appropriately.

Intuitively, the set S_a is supposed to represent the set of types of a-tuples $x_a = \langle x_i : i \in a \rangle$ in the models of some theory. Clearly, if $f : a \to b$ is a map and the type of $\langle x_i : i \in b \rangle$ is known, then so is that of $\langle x_{f(i)} : i \in a \rangle$, as it is merely a sub-tuple, possibly with repetitions. Thus we obtain the map $f^* : S_b \to S_a$. Note that in the particular case where f is injective, the effect of f^* is just taking the type of a sub-tuple, and should be surjective, whereas when f is surjective, f^* gives the types of tuples with duplicate elements, that is the sub-space corresponding to the conjunction of some equality relations. The equality-preservation condition says that the meaning of equality is sane: two elements of an a-tuple are equal if and only if they are equal as a 2-tuple. The amalgamation condition is straightforward: types which coincide on a sub-tuple can be amalgamated. Finally, the topology: we understand closed sets as definable by some possibly infinite conjunction. The T_1 requirement means that every type can be expressed in the language; the compactness is just that; the continuousness of the maps corresponds standard manipulations of the free variables of a formula: using the same variable several times, or adding dummy variables, to render it in another tuple of free variables; and finally the closedness requirement of the maps f^* corresponds both to the application of the existential quantifier (where f is the inclusion of the sub-tuple on which we do not quantify) and to the definability of equality (for example, when $f : 2 \to 1$, then $f^*(S_1) \subseteq S_2$ is the closed set corresponding to the formula $x_0 = x_1$). The modifications of the above needed when considering multiple sorts should be obvious.

As usual in model theory, equality needs a somewhat special treatment. In this context, this means we need to understand the map f^*, when f is surjective.

Lemma 2.20. *Let S be a type-space functor.*

1. *If a map $f : a \to b$ is surjective (injective) then f^* is injective (surjective).*
2. *Let $f : a \to b$ be surjective. Assume that $a = b \cup c$ and f is the identity on b, and let $g : b \to a$ be the inclusion. Then $f^*(X) = f^*(S_b) \cap g_*(X)$ for every $X \subseteq S_b$.*
3. *Continuing with the previous notations and assumptions, we have a strong version of equality-preservation: For $i \in c$, define $h_i : 2 \to a$ by $h_i(0) = f(i)$, $h_i(1) = i$. Then $f^*(S_b) = \bigcap_{i\in c} h_{i*}(eq)$ (in other words, many equalities hold at the same time if and only if each one holds separately).*

Proof. 1. If f is surjective, then there is (an injective) $g : b \to a$ such that $f \circ g = Id_b$. Then $g^* \circ f^* = Id_{S_b}$, so f^* must be injective. Similarly if f is injective.
2. If $y \in f^*(X)$, then $y = f^*(x) \in f^*(S_b)$ for $x \in X$ and $g^*(y) = g^* \circ f^*(x) = x \in X$, so $y \in g_*(X)$. Conversely, if $y \in f^*(S_b) \cap g_*(X)$, then $y = f^*(x)$ for some $x \in S_b$, and $x = g^* \circ f^*(x) = g^*(y) \in X$. Until now, we only used the fact that S was a functor.
3. Assume at first that a is finite, so so is b. We prove this by induction on $|c|$. When $|c| = 0$, there is nothing to show. So take $i \in c$, and write $b' = a \smallsetminus \{i\}$, $c' = c \smallsetminus \{i\}$, and define $f' : a \to b'$ as the identity on b' and f on $\{i\}$, and $f'' = f \restriction b'$. Then $f = f'' \circ f'$. By the induction hypothesis and equality-preservation:

$$f'^*(S_{b'}) = h_{i*}(eq)$$
$$f''^*(S_b) = \bigcap_{j\in c'} (f' \circ h_j)_*(eq) = f'_*(\bigcap_{j\in c'} h_{j*}(eq))$$

Thus:

$$\begin{aligned} f^*(S_b) &= f'^*(f''^*(S_b)) = f'^*(f'_*(\bigcap_{j\in c'} h_{j*}(eq))) \\ &= f'^*(f'^{*-1}(\bigcap_{j\in c'} h_{j*}(eq))) = \bigcap_{j\in c'} h_{j*}(eq) \cap f'^*(S_{b'}) \\ &= \bigcap_{j\in c} h_{j*}(eq) \end{aligned}$$

as required. For an infinite a, we prove each inclusion separately. For the first, take any $i \in c$, repeat the construction above, and define also $h' : 1 \to b'$ by $h'(0) = f(i)$. Then $f' \circ h_i = h' \circ f_{eq}$, so $h_i^*(f'^*(S_{b'})) = f_{eq}^*(h'^*(S_{b'})) = eq$, and $f^*(S_b) \subseteq f'^*(S_{b'}) \subseteq h_i^{*-1}(h_i^*(f'^*(S_{b'}))) = h_{i*}(eq)$. For the other inclusion, take $x \in \bigcap_{i\in c} h_{i*}(eq)$, and we wish to prove that $x \in f^*(S_b)$. Take $g : b \to a$ as above, and we will show that $x = f^*(g^*(x))$. By the definition of S_a as an inverse limit, it suffice to prove that for every finite $\tilde{a}' \subseteq a$ there is a finite $\tilde{a}' \subseteq \tilde{a} \subseteq a$ such that $\pi^*(x) = \pi^*(f^*(g^*(x)))$, where $\pi : \tilde{a} \to a$ is the inclusion. So given $\tilde{a}' \subseteq a$ we replace it with $\tilde{a} = \tilde{a}' \cup f(\tilde{a}')$, so that $f(\tilde{a}) \subseteq \tilde{a}$. Write $\tilde{b} = b \cap \tilde{a}$, $\tilde{c} = c \cap \tilde{a}$, so $\tilde{a} = \tilde{b} \cup \tilde{c}$, and let $\tilde{f} : \tilde{a} \to \tilde{b}$, $\tilde{g} : \tilde{b} \to \tilde{a}$ and $\tilde{h}_i : 2 \to \tilde{a}$ for $i \in \tilde{c}$ be the obvious restrictions. Then we have reduced to the finite case and $\tilde{f}^*(S_{\tilde{b}}) = \bigcap_{i\in\tilde{c}} \tilde{h}_{i*}(eq)$. On the other hand, $g \circ f \circ \pi = \pi \circ \tilde{g} \circ \tilde{f}$, and $h_i = \pi \circ \tilde{h}_i$ for $i \in \tilde{c}$. Then for $i \in c$ we have $\pi^*(x) \in \pi^*(h_{i*}(eq)) = \pi^*(\pi_*(\tilde{h}_{i*}(eq))) \subseteq \tilde{h}_{i*}(eq)$. Then $\pi^*(x) \in \tilde{f}^*(S_{\tilde{b}})$, say $\pi^*(x) = \tilde{f}^*(y)$. As $\tilde{f} \circ \tilde{g} = Id_{\tilde{b}}$, we obtain $\pi^*(x) = \tilde{f}^*(y) = \tilde{f}^*(\tilde{g}^*(\tilde{f}^*(y))) = \tilde{f}^*(\tilde{g}^*(\pi^*(x))) = \pi^*(f^*(g^*(x)))$ as required. This concludes the proof.

QED

Corollary 2.21. *A finite compact type-space functor S is a compact type-space functor.*
In other words, with the presence of compactness, the finitary versions of the requirements from a compact type-space functor imply the infinitary ones.

Proof. Clearly, S_κ is T_1 for every κ. We now verify that S_κ is compact for every infinite κ, and that for injective $f : \alpha \hookrightarrow \beta$, f^* is closed. We do this by induction on κ, assuming $|\beta| \leq \kappa$. Write $S_\kappa = \varprojlim_{i<\kappa} S_i$. Then by the assumption S_i are all compact, and the projection $f^*_{ij} : S_j \to S_i$ is closed for every $i \leq j < \kappa$. Let $\{C_\gamma : \gamma \in \Gamma\}$ be closed subsets of S_κ having the finite intersection property. We may assume that they are closed for finite intersection. By the definition of the topology on S_κ, each C_γ can be written as $\bigcap_{i<\kappa} f_{i*}(C_{\gamma,i})$ where $f_i : i \hookrightarrow \kappa$ are the injections, and $C_{\gamma,i} \subseteq S_i$ are closed subsets. We claim that we may assume that $f^*_{ij}(C_{\gamma,j}) = C_{\gamma,i}$ for every $i < j < \kappa$. Indeed, we can replace $C_{\gamma,i}$ with $C_{\gamma,i} \cap f^*_{ij}(C_{\gamma,j})$, and $C_{\gamma,j}$ with $C_{\gamma,j} \cap f_{ij*}(C_{\gamma,i})$: these sets are closed as f^*_{ij} is continuous and closed, and the changes do not influence C_γ. Iterating until the process stabilises, we obtain the required property. As a consequence we also have $f^*_i(C_\gamma) = C_{\gamma,i}$: $f^*_i(C_\gamma) \subseteq C_{\gamma,i}$ is clear, and for every $x_i \in C_{\gamma,i}$ we can construct a branch above it of $x_j \in C_{\gamma,j}$, and its limit is in C_γ.
By induction on i, choose $x_i \in \bigcap_\gamma C_{\gamma,i}$ such that $x_i = f^*_{ij}(x_j)$ for $i < j$. As S_0 is compact, just choose some $x_0 \in \bigcap C_{\gamma,0}$. For a limit i, the sequence $\langle x_j : j < i\rangle$ fixes x_i completely, and one sees it has to be in each $C_{\gamma,i} = \bigcap_{j<i} f_{ji*}(C_{\gamma,j})$. For $i = j+1$, we want to show that $\{C_{\gamma,i} : \gamma \in \Gamma\} \cup \{f_{ji*}(x_j)\}$ has the finite intersection property. Note first that

$$\bigcap_{k<n} C_{\gamma_k,i} = \bigcap_{k<n} f^*_i(C_{\gamma_k}) \supseteq f^*_i(\bigcap_{k<n} C_{\gamma_k}) = f^*_i(C_\gamma) = C_{\gamma,i}$$

for some γ, as $\{C_\gamma\}$ is closed for finite intersection. Thus it suffices to show that $C_{\gamma,i} \cap f_{ji*}(x_j) \neq \varnothing$ for every γ. However, this is equivalent to $x_j \in f^*_{ji}(C_{\gamma,i}) = C_{\gamma,j}$, which is true. Thus we may choose $x_i \in f_{ji*}(x_j) \cap \bigcap C_{\gamma,i}$, and continue the construction. The sequence $\langle x_i : i < \kappa\rangle$ fixes an element in S_κ, which is clearly in $\bigcap C_\gamma$, and we proved that S_κ is compact. Assume now that $f : a \hookrightarrow \kappa$ is an inclusion, and we want to show that f^* is closed. In fact, if a is bounded in κ, then we have already shown this. In the general case, we obtain a "ladder", where on one side we have $S_i : i \leq \kappa$ and on the other $S_{i\cap a} : i \leq \kappa$. We know by the induction hypothesis on κ that for every $i < \kappa$, the map $f^*_i : S_i \to S_{i\cap a}$ is closed. We also note the following inclusions: $g_{ij} : i \hookrightarrow j$ and $h_{ij} : i \cap a \hookrightarrow j \cap a$, and when there is an index κ we allow ourselves to omit it (except for S_κ). Let $C \subseteq S_\kappa$ be closed, and we want to prove that $f^*(C) \subseteq S_a$ is closed. Let $C_i = g^*_i(C)$ and $D_i = f^*_i(C_i)$ for $i \leq \kappa$ (so in particular $D_\kappa = D = f^*(C)$). For $i < \kappa$, C_i is closed as $i \subseteq \kappa$ is bounded, and D_i is closed as $|i| < \kappa$. We claim that $D = \bigcap_{i<\kappa} h_{i*}(D_i)$. Indeed, we clearly have $D \subseteq h_{i*}(D_i)$. On the other hand, let $x \in \bigcap_{i<\kappa} h_{i*}(D_i)$, and let $x_i = h^*_i(x) \in D_i$. Then the singleton $\{x_i\}$ is closed in $S_{i\cap a}$, thus $f_*(h_{i*}(x_i)) = g_{i*}(f_{i*}(x_i)) \subseteq S_\kappa$ is closed, and this is a decreasing sequence in i. For every i, we have $C \cap g_{i*}(f_{i*}(x_i)) \neq \varnothing$, as $x_i \in D_i = f^*_i(g^*_i(C))$. Thus, by

compactness of S_κ, there is $y \in C \cap \bigcap f_*(h_{i*}(x_i))$. This means that $f^*(y) = x$, and we are done. This concludes the proof that S_κ is compact for all κ, and that f^* is closed for every injective f.
The infinitary amalgamation property is proved in a similar manner using compactness. When $f : a \to b$ is surjective and $C \subseteq S_b$ is closed, then following notations and conventions of Lemma 2.20 we have $f^*(C) = f^*(S_b) \cap g_*(C) = \bigcap_{i \in c} h_{i*}(eq) \cap g_*(C)$ which is closed, concluding the proof. QED

Question 2.22. Is a finite set type-space functor a set type-space functor? In other words, does the finite amalgamation property implies the full one?

We can now characterise positive Robinson theories as precisely those who have compact type-spaces. On the one hand, $\mathrm{S}(T)$ is a compact type-space functor (see Lemma 1.39). Conversely, we have:

Theorem 2.23. *There is an operation* Th_{pos} *that associates canonically to every compact type-space functor* S *a strongly positive Robinson theory* $\mathrm{Th}_{pos}(S)$, *such that* $\mathrm{S}(\mathrm{Th}_{pos}(S)) \cong S$.

Proof. We start with S, and we try to obtain a positive Robinson theory $T = \mathrm{Th}_{pos}(S)$. We take $\mathcal{L}_n = \{P_R : R \subseteq S_n \text{ is closed}\}$ and $\mathcal{L} = \bigcup \mathcal{L}_n$, where the elements of $\mathcal{L}_n$ are considered as n-ary predicates, and $\Delta = \Delta_0$. For every $k, n \in \omega$, $\bar{n} \in \omega^k$, maps $f_i : n_i \to n$ and $P_{R_i} \in \mathcal{L}_{n_i}$ for $i < k$, consider $\bigcap_{i<k} f_{i*}(R_i)$. If it is empty, then T contains the sentence

$$\neg \exists x_{<n} \bigwedge_{i<k} P_{R_i}(f_i^*(x_{<n}))$$

and T contains only sentences obtained in this way. Then T is an Π_1-theory, and it is consistent as any $\mathcal{L}$-structure not satisfying any relations is a model of T. Let M be an e.c. model T, and write $M = \{a_i : i < \kappa\}$. Consider the intersection:

$$C = \bigcap \{f_*(R) : m < \omega, P_R \in \mathcal{L}_m, f : m \to \kappa, M \models P_R(f^*(\bar{a}))\}$$

This is an intersection of closed sets of S_κ, and each finite sub-intersection is non-empty, as T would forbid otherwise. Therefore $C \neq \varnothing$, and we may choose some element $p \in C$. Define now a structure M', over the same base set, such that $M' \models P_R(f^*(\bar{a}))$ if and only if $p \in f_*(R)$ for every $m < \omega$, $f : m \to \kappa$ and $P_R \in \mathcal{L}_m$. Then $M' \models T$ is a continuation of M, and as M is e.c. they must be equal. This means in particular that $C = \{p\}$, so it would make sense to define $\mathrm{tp}^S(M) = \mathrm{tp}^S(a_{<\kappa}) = p$, and similarly, for every $f : n \to \kappa$ we can define $\mathrm{tp}^S(f^*(a_{<\kappa})) = f^*(\mathrm{tp}^S(a_{<\kappa})) = f^*(p)$. We clearly have $M \models P_R(f^*(a_{<\kappa}))$ if and only if $\mathrm{tp}^S(f^*(a_{<\kappa})) \in R$.
In order to say that $\mathrm{tp}^S(\bar{b})$ determines $\mathrm{tp}^M_{\Delta_0}(\bar{b})$, there is still one predicate to deal with: equality. So let $h : 2 \to \kappa$ be $0 \mapsto i, 1 \mapsto j$, and we want to prove that $M \models P_{eq}(a_i, a_j)$ if and only if $a_i = a_j$. One direction is clear. For the other, assume that $M \models P_{eq}(a_i, a_j)$ and $i \neq j$. So let $\kappa' = \kappa \smallsetminus \{j\}$, and let $g : \kappa \to \kappa'$ send j to i and be the identity elsewhere. Then $\mathrm{tp}^S(M) \in h_*(eq) = g^*(S_{\kappa'})$, say $\mathrm{tp}^S(M) = g^*(q)$. Construct M'' on a base set $\{a_k : k \in \kappa'\}$ from q as we did M' from p, and let $\varphi : M \to M''$ send a_k to $a_{g(k)}$. Then for $f : n \to \kappa$, $M \models R(f^*(a_{<\kappa}))$ if and only if $(g \circ f)^*(q) = f^*(g^*(q)) = f^*(\mathrm{tp}^S(M)) \in R$ if and only if $M'' \models P_R((g \circ f)^*(a_{<\kappa}))$. But

$(g \circ f)^*(a_{<\kappa}) = \varphi(f^*(a_{<\kappa}))$, so φ is a homomorphism (in fact, but for equality, we have shown it to be an isomorphism), and $M'' \models \varphi(a_i) = \varphi(a_j)$. As M is e.c., we must have had $a_i = a_j$ to begin with, contradicting $i \neq j$.
We prove that $\mathrm{S}(T) \simeq S$ by sending $\mathrm{tp}(\bar{b})$ (for some $M \in \mathcal{M}(T)$ and $\bar{b} \in M^n$) to $\mathrm{tp}^S(\bar{b})$. This map is necessarily surjective, and we wish to show that it is injective. Let $M_i \in \mathcal{M}(T)$, $\bar{b}_i \in M^n$ for $i < 2$, and assume that $\mathrm{tp}^S(\bar{b}_0) = \mathrm{tp}^S(\bar{b}_1) = q \in S_n$. Since q determines equality we may assume that both $\bar{b}_i$ are without repetition. By the amalgamation property of S we can amalgamate $\mathrm{tp}^S(M_0)$ and $\mathrm{tp}^S(M_1)$ over q, and realise it in some model $N \models T$. Since $\mathrm{tp}^S(-)$ determines $\mathrm{tp}_{\Delta_0}(-)$, we obtain Δ_0-homomorphisms $g_i : M_i \to N$ with $g_0(\bar{b}_0) = g_1(\bar{b}_1)$. We may assume that $N \in \mathcal{M}(T)$ whereby $\mathrm{tp}(\bar{b}_0) = \mathrm{tp}(g_i(\bar{b}_i)) = \mathrm{tp}(\bar{b}_1)$.
Finally, the Δ_0-topology clearly coincides with the original topology on S, in which every f^* is closed. Therefore T is strongly positive Robinson. QED

We can characterise the classical subdivisions:

Theorem 2.24. *There is an operation* Th *that associates canonically to a totally disconnected compact type-space functor S a classical Robinson theory (i.e., with negation)* $\mathrm{Th}(S)$*, such that* $\mathrm{S}(\mathrm{Th}(S)) \cong S$*.*
Moreover, $\mathrm{Th}(S)$ *is first-order (i.e., has a first-order model completion) if and only if S is open as well (i.e., f^* are open maps).*

Proof. In the proof above, taking the language to be a base for the closed sets would have sufficed. If S is totally disconnected and compact, then the clopen sets form such a base. So take it as the language $\mathcal{L}$, let Δ be the set of all quantifier-free formulas, and $\Delta_0 \subseteq \Delta$ the set of all positive ones. Let T be $\mathrm{Th}_{pos}(S)$ restricted to $\mathcal{L}$: it is a positive Robinson Π_1-theory and $\mathrm{S}(T) \simeq S$. Moreover, in $\mathcal{M}(T)$, every Δ-formula is equivalent to a Δ_0-formula, so up to obvious identifications we have $\Delta_0 = \Delta_0^{\mathrm{QE}(\Delta)}$, and in can apply Proposition 1.33: in particular, if $T_0 = \mathrm{Th}_{\Pi(\Delta)}(\mathcal{M}(T))$, then $\mathcal{M}(T) = \mathcal{M}(T_0)$, and T_0 is a classical Robinson theory.
For the moreover part, we know that the type-space functor of a first order theory is clopen. For the converse, assume that S above is also open. Then $f^*(R)$ is clopen for every clopen R, thus the positive quantifier-free formulas are also closed for quantification, and in fact give all of $\mathcal{L}_{\omega,\omega}$. We can therefore apply the above argument (and in particular Proposition 1.33) with $\Delta = \mathcal{L}_{\omega,\omega}$. QED

2.3. **Compact abstract elementary categories.** Here we give yet another equivalent description of positive Robinson theories (and a method for obtaining such). We show that abstract elementary classes with amalgamation and locality of types, or a slight variant thereof, satisfying in addition a (weak) version of compactness, are indeed equivalent in a sense to positive Robinson theories (by "equivalent" we mean "having the same type-spaces"). Using this method, we can show that Hilbert and Banach spaces are such.

Convention 2.25. In order to put this section in the context of ZFC, a few remarks are required:

1. We shall consider many equivalence relations for which the equivalence classes are proper classes. Therefore we represent each equivalence class by the set of all its members of minimal foundation rank in this class.
2. The objects of a category may form a proper class, though the morphisms between each pair are always a set. If $\mathrm{obj}(\mathcal{C})$ is a set, we say that $\mathcal{C}$ is a set category, otherwise we say that it is a proper class category.
3. If $\mathcal{C}$, $\mathcal{C}'$ are categories, then a functor from $\mathcal{C}$ to $\mathcal{C}'$ is a class. if $\mathcal{C}$ is a set category, then such a functor is a set, and it makes sense to speak of the category $\mathrm{Funct}(\mathcal{C}, \mathcal{C}')$ (which is a set if and only if $\mathcal{C}'$ is).

Definition 2.26. Let $\mathcal{C}$ a category equipped with a functor $|\cdot|$ to the category of sets. Then $\mathcal{C}$ (or more precisely, $(\mathcal{C}, |\cdot|)$) is a *concrete category* if:

1. Whenever $A, B \in \mathcal{C}$, and $f : |A| \to |B|$ is a function, there is at most one $g \in \mathrm{Hom}_{\mathcal{C}}(A, B)$ such that $f = |g|$ (in other words, $|\cdot|$ is a faithful functor). In case g exists, we say that f is a morphism and identify it with g.
2. Whenever $A \in \mathcal{C}$, X is a set, and $f : |A| \to X$ is a bijection, there is a unique $f(A) = B \in \mathcal{C}$, such that $X = |B|$ and $f \in \mathrm{Hom}_{\mathcal{C}}(A, B)$.

A *functor of concrete categories* a functor $F : \mathcal{C} \to \mathcal{C}'$, such that $|\cdot|_{\mathcal{C}'} \circ F = |\cdot|_{\mathcal{C}}$.

Definition 2.27. Let $\mathcal{M}$ be a concrete category. Call its objects models, and its morphisms elementary embeddings. Write $M \preceq_{\mathcal{M}} N$ if $M \subseteq N$ and the inclusion is a morphism. Then $\mathcal{M}$ is an *elementary category with amalgamation* if it satisfies the following:

Injectiveness: All morphisms are injective (on the underlying sets).

Tarski-Vaught property: Whenever $M_0 \subseteq M_1$ and $M_0 \preceq_{\mathcal{M}} N$, $M_1 \preceq_{\mathcal{M}} N$, then $M_0 \preceq_{\mathcal{M}} M_1$.

Elementary chain property: A $\preceq_{\mathcal{M}}$-chain $\langle M_i \rangle$ is bounded from above, i.e. there a structure N with $M_i \preceq_{\mathcal{M}} N$ for all i. If both N_0, N_1 have this property, then there is $P \in \mathcal{M}$ and $f_i \in \mathrm{Hom}_{\mathcal{M}}(N_i, P)$ which are the same on $\bigcup |M_i|$.

Amalgamation: Whenever $f_i \in \mathrm{Hom}_{\mathcal{M}}(M, N_i)$ for $i < 2$, there are $P \in \mathcal{M}$ and $g_i \in \mathrm{Hom}_{\mathcal{M}}(N_i, P)$ such that $g_0 \circ f_0 = g_1 \circ f_1$.

Remark 2.28. Shelah's definition of an abstract elementary class (as can be found, in [She99]) does not include amalgamation, but requires more structure on the class itself. We omit these structure requirements, as we do not need them and they may just cause confusion with the structure that *we* shall put on the class later on. In any case, everything said here for elementary categories with amalgamation holds full well for abstract elementary classes with amalgamation, so the difference is not essential.
Note that the second clause in the definition of the elementary chain property just says that $\mathrm{tp}_{N_0}(\bigcup |M_i|) = \mathrm{tp}_{N_1}(\bigcup |M_i|)$ (see below), so if we ever wish to remove the amalgamation hypothesis it will make sense to weaken this clause accordingly.

Convention 2.29. When we say elementary category we mean with amalgamation.
$\mathcal{M}$ usually denotes an elementary category.
For $M \in \mathcal{M}$, we usually write M instead of $|M|$ when the meaning is clear from the context.

Definition 2.30. Let $\mathcal{M}$ be an elementary category. For $M_i \in \mathcal{M} : i < 2$, and tuples $a_i \in M_i$, write $(M_0, a_0) \equiv (M_1, a_1)$ if there is $N \in \mathcal{M}$ and $f_i \in \mathrm{Hom}_{\mathcal{M}}(M_i, N)$ such that $f_0(a_0) = f_1(a_1)$. By amalgamation, this is an equivalence relation. The class $\mathrm{S}_\alpha(\mathcal{M}) = \langle (M, a)/\equiv : M \in \mathcal{M}, a \in M^\alpha \rangle$ is the class of pure α-types in $\mathcal{M}$, and $\mathrm{S}(\mathcal{M}) = \bigcup \mathrm{S}_\alpha(\mathcal{M})$. We also write $\mathrm{tp}_M(a) = (M, a)/\equiv$, namely the type of a in M.

Definition 2.31. An elementary category $\mathcal{M}$ is *connected* if $\mathrm{S}_0(\mathcal{M})$ is a single element. More generally, for $p \in \mathrm{S}_0(\mathcal{M})$, $\mathcal{M}_p = \{M \in \mathcal{M} : \mathrm{tp}_M(\varnothing) = p\}$ is the connected component of p. One verifies easily that this is a maximal connected elementary subcategory.

In order to do first-order-like manipulation, we still need a few axioms:

Definition 2.32. We say that an elementary category is *compact* if in addition it satisfies:

Type boundedness: For all $n \leq \omega$, $\mathrm{S}_n(\mathcal{M})$ is a set.

Type locality: We have $\mathrm{tp}(a) = \mathrm{tp}(b)$ if and only if for every finite sub-tuple a' of a, $\mathrm{tp}(a') = \mathrm{tp}(b')$ where b' is the corresponding sub-tuple of b.

Weak compactness: Assume that x_I is an infinite tuple, $\mathcal{I} \subseteq \mathcal{P}(I)$ and $\Sigma(x_I) = \{p_J(x_J) : J \in \mathcal{I}\}$ is a set of of types on sub-tuples of I, such that for every finite $\Sigma_0 \subseteq \Sigma$ there are $M \in \mathcal{M}$ and $a_I \in M^I$ that realises Σ_0. Then there are $M \in \mathcal{M}$ and $a_I \in M^I$ that realise Σ.

Remark 2.33. We call this weak compactness, as it corresponds to the compactness of a somewhat poor topology, namely the minimal topology rendering $\mathrm{S}(\mathcal{M})$ a finite topological type-space functor (see below).

Remark 2.34. The following weaker version of compactness holds in any elementary category with type-locality: Let I be a set, $\mathcal{I} \subseteq \mathcal{P}(I)$, x_I an infinite tuple indexed by I, and $\{p_J(x_J) \in \mathrm{S}_J(\mathcal{M}) : J \in \mathcal{I}\}$ a set of (representatives of) pure types in sub-tuples of x_I, such that $(\mathcal{I}, \subseteq)$ is directed, and whenever $J, J' \in \mathcal{I}$ and $J \subseteq J'$, then $p_{J'}(x_{J'}) \vdash p_J(x_J)$ in the obvious sense (in the type-space functor terminology, if $f : J \hookrightarrow J'$ is the inclusion, then $f^*(p_{J'}) = p_J$). Then there is $M \in \mathcal{M}$ and $a_I \in M$ such that $\mathrm{tp}_M(a_J) = p_J$ for all $J \in \mathcal{I}$.

Proof. We prove this by induction on $\kappa = |I|$. For finite κ this is clear, so we may assume it is infinite. We reduce easily to the case where $\mathcal{I} = \mathcal{P}^{fin}(I)$. We may of course assume that $I = \kappa$ as sets. We construct by induction on $i < \kappa$ a $\preceq_{\mathcal{M}}$-chain $\langle M_i \rangle$ and choose $a_i \in M_{i+1}$ that realise the proper types, in the following manner:

1. $i = 0$: Take M_0 to be any model that realises $p_\varnothing$.
2. i limit: Take M_i to be some $\preceq_{\mathcal{M}}$-bound for $\langle M_j : j < i \rangle$.
3. $i = j + 1$: We already have constructed $a_{<j}$, and want to construct a_j. By the hypothesis of the induction on κ, as $|i| < \kappa$ there is some model M' and $b_{<i} \in M'$ that realise all p_J for $J \in \mathcal{P}^{fin}(i)$. By type locality, $\mathrm{tp}_{M'}(b_{<j}) = \mathrm{tp}_{M_j}(a_{<j})$, so we may take M_i to witness $(M_j, a_{<j}) \equiv (M', b_{<j})$. We may assume that $M_i \succeq_{\mathcal{M}} M_j$, and take a_j to be the image of b_j in M_i.

Now taking a bound for the entire sequence $\langle M_i \rangle$, we obtain the result. QED

We connect this with previous definitions:

Lemma 2.35. *An elementary category $\mathcal{M}$ has type boundedness and locality if and only if* $\mathrm{S}(\mathcal{M})$ *is a set type-space functor. It is compact if and only if* $\mathrm{S}(\mathcal{M})$ *can be topologised to be a compact type-space functor. This is further equivalent to: the minimal topology on* $\mathrm{S}(\mathcal{M})$ *rendering it a finite topological type-space functor is compact.*

Proof. The first assertion should be clear, using Remark 2.34, so we prove the second. Just during this proof, call a topology rendering $\mathrm{S}(\mathcal{M})$ a finite topological type-space functor *good*, and note that this property does not involve any $\mathrm{S}_\alpha(\mathcal{M})$ for infinite α. The discrete topology on $\mathrm{S}(\mathcal{M})$ (that is, on $\mathrm{S}_n(\mathcal{M})$ for every $n < \omega$) is clearly good, and the intersection of good topologies is also good. Thus a minimal good topology exists, as the intersection of all good topologies. On the other hand, the minimal good topology can be constructed by starting with the co-finite topology on every S_n, and adding at each step all images and inverse images of closed sets by every f^*, and taking the generated topology. Iterating this would give the minimal good topology.
We now note that in the definition of weak compactness we could have added existential quantifiers, equality, finite disjunction and infinite conjunction at no additional costs: an existential quantifier can be disposed of, making sure that the quantified variables are new, and instead of saying $x = y$, we can just replace all the occurrences of y with x. Finite disjunction and infinite conjunction are disposed of easily in standard ways so we reduce to the original compactness property.
Assume now that $\mathcal{M}$ satisfies weak compactness, and put on $\mathrm{S}(\mathcal{M})$ the topology where the closed sets on $\mathrm{S}_n(\mathcal{M})$ are defined by formulas with n free variables, generated from the type-predicates and equality by finite disjunction, infinite conjunction and existential quantification. Then this topology is clearly good, and by the previous argument it is also compact.
Conversely, assume that the minimal good topology on $\mathrm{S}(\mathcal{M})$ is compact. Then by Corollary 2.21, $\mathrm{S}(\mathcal{M})$ is a compact type-space functor. This implies that S_κ is compact for every κ, so weak compactness holds. QED

Remark 2.36. It is important to point out that if $\mathcal{M}$ is compact, then $\mathrm{S}(\mathcal{M})$ comes already with one intrinsic compact topology, namely the minimal one rendering it a topological type-space functor. Given any specific $\mathcal{M}$, one may point out richer topologies that are still compact (and we will do so in the examples we give), but one does not have to: the minimal one will always suffice. This means that in what follows, a minimal (and very poor) language can always be deduced directly from the pair $(\mathcal{M}, |\cdot|)$, without any language appearing explicitly in the input.

We thus obtain:

Proposition 2.37. *Let T be a positive Robinson theory, $\mathcal{M}$ the category of its e.c. models. Then $\mathcal{M}$ is a compact elementary category, and the notions of type in a positive Robinson theory and in an elementary category coincide.*

Proof. Easy. QED

And the converse:

Theorem 2.38. *Let $\mathcal{M}$ be a compact elementary category, with a given compact type-space functor topology on* $\mathrm{S}(\mathcal{M})$. *Then there is a signature $\mathcal{L}$, a positive Robinson* *theory T in $\mathcal{L}$ with respect to $\Delta = \Delta_0$, and a functor F from $\mathcal{M}$ to the category of $\mathcal{L}$-structures, such that:*

1. *T is the Π_1-theory of $\mathcal{M}$ (or rather, of $F(\mathcal{M})$), and $\mathrm{S}(T) \cong \mathrm{S}(\mathcal{M})$ as compact type-space functors.*
2. *F is a functor of concrete categories.*
3. *For any e.c. model M of T there is $N \in \mathcal{M}$ such that $M \subseteq_{\mathcal{L}} F(N)$.*
4. *For any $M \in \mathcal{M}$ there is an e.c. model N of T such that $F(M) \subseteq_{\mathcal{L}} N$.*
5. *For all $M, N \in \mathcal{M}$: $M \preceq_{\mathcal{M}} N \iff F(M) \subseteq_{\mathcal{L}} F(N)$.*

Proof. Take $T = \mathrm{Th}_{pos}(\mathrm{S}(\mathcal{M}))$, and F the natural interpretation of $M \in \mathcal{M}$ as an $\mathcal{L}$-structure. Then:

1. The second assertion is already known. As for the first, it suffices to show that $T \vdash \neg\exists\bar{x} \bigwedge R_i(x_{f_i})$ if and only if $\mathcal{M} \models \neg\exists\bar{x} \bigwedge R_i(x_{f_i})$, where where R_i are n_i-ary relations, and $f_i : n_i \to n$ give the relevant sub-tuple. Clearly, both are equivalent to $\bigcap f_{i*}(R_i) = \varnothing$ in $\mathrm{S}_n(\mathcal{M}) = \mathrm{S}_n(T)$.
2. Clear.
3. As $M \models T = \mathrm{Th}_\Pi(\mathcal{M})$, $\Delta(M)$ is finitely realized in $\mathcal{M}$. By weak compactness, M can be continued to a $F(N)$ for $N \in \mathcal{M}$. As $F(N) \models T$ and M is e.c., this is an embedding and $M \subseteq_{\mathcal{L}} F(N)$.
4. As $F(M) \models T$ by definition, we can continue $F(M)$ to an e.c. model N of T. As every finite tuple from $F(M)$ satisfies precisely one predicate, and every finite tuple in N satisfies at most one predicate (T would forbid more), we see that $F(M) \subseteq_{\mathcal{L}} N$.
5. One direction is evident by the construction. For the other, if $M \subseteq_{\mathcal{L}} N$ then by type locality, if m is an enumeration of M, then $\mathrm{tp}_M(m) = \mathrm{tp}_N(m)$. By definition, there is P which amalgamates M and N over m. Now by Tarski-Vaught we get that $M \preceq_{\mathcal{M}} N$.

QED

2.4. **Examples.** We start with a simple example, of an elementary abstract category with type-locality (that is, of a set type-space functor) which is not compact.

Consider first the category $\mathcal{C}$ of bipartite graphs. So a model $M \in \mathcal{C}$ is a disjoint union $U_M \cup V_M$, and there is also the graph $R_M \subseteq U_M \times V_M$, and $\mathrm{Hom}_{\mathcal{C}}(M, N)$ is the set of embeddings of M into N. It is easily verified that $\mathcal{C}$ is a compact elementary category, and $\mathrm{S}_n(\mathcal{C})$ is finite for every $n < \omega$. Thus we recover from the category structure an ω-categorical first-order theory, that of the generic bipartite graph.

For $a \in U_M$ note $R_M(a) = \{b \in V_M : (a, b) \in R_M\}$. For $M \in \mathcal{C}$ and $a \neq a' \in U_M$ (that is, $a, a' \in U_M$ such that $a \neq a'$), write:

$$\mathrm{acl}_M(aa') = \begin{cases} |R_M(a) \cap R_M(a')| < \omega & \{a, a'\} \cup (R_M(a) \cap R_M(a')) \\ \text{otherwise} & \{a, a'\} \end{cases}$$

And for $A \subseteq M$:

$$\mathrm{acl}_M(A) = A \cup \bigcup_{\{a,a'\} \in [U_M \cap A]^2} \mathrm{acl}_M(aa')$$

In order to add structure to $\mathcal{C}$, we consider the sub-category $\mathcal{C}'$ where the objects are the same, but morphisms are now required to preserve acl. Thus, working in $\mathcal{C}'$, we do not need to specify in which model we take the acl. One verifies that $\mathcal{C}'$ is an elementary category (with amalgamation). We claim then that for a set A, $\mathrm{tp}(A)$ is just the ismorphism class of $\mathrm{acl}(A)$ (as a bipartite graph, over A). This is clear, as $A \subseteq M$ implies that $\mathrm{acl}(A) \leq_{\mathcal{C}'} M$, so just apply amalgamation. This shows type locality and boundedness. In particular, this means that $\mathrm{S}_2(\mathcal{C})$ contains the following types: $x = y \in U$, $x = y \in V$, $x\ R\ y$, $y\ R\ x$, $\neg(x\ R\ y)$ $\neg(y\ R\ x)$ (in the last two we unterstand implicitly that the left-hand side is in U and the right-hand in V), $x \neq y \in V$, and for $\alpha \in \omega \cup \{\infty\}$: $x \neq y \in U \wedge |R(x) \cap R(y)| = \alpha$. It can be verified that $\mathcal{C}'$ is also compact.

Our goal however is $\mathcal{C}'' \subseteq \mathcal{C}'$, which is the full sub-category consisting of all objects $M \in \mathcal{C}'$ such that $|R_M(a) \cap R_M(a')| < \omega$ for all $a \neq a' \in U_M$. Then everything we said about $\mathcal{C}'$ holds, except the following: the 2-type $x \neq y \in U \wedge |R(x) \cap R(y)| = \infty$ no longer exists, and compactness is lost. Indeed, the following set is finitely consistent but not consistent:

$$\{xRz_i : i < \omega\} \cup \{yRz_i : i < \omega\} \cup \{z_i \neq z_j \in V : i < j < \omega\} \cup \{xRw, \neg(yRw)\}$$

So in a sense, $\mathcal{C}'$ is a compactification of $\mathcal{C}''$.

We continue with very natural examples, Hilbert and Banach spaces. Note that similar examples exist also in [BL00].

Example 2.39. Let $\mathbb{F}$ be $\mathbb{R}$ or $\mathbb{C}$. Then category $\mathscr{B}$ of normed vector spaces over $\mathbb{F}$, with isometries as morphisms, is a connected compact elementary category, and the sub-category $\mathscr{H}$ of inner product spaces is a connected compact elementary sub-category. In both, two tuples have the same type if and only if there is an isometry of the subspaces they generate sending one onto the other.

Proof. We give the proof for $\mathscr{B}$, unless where it is different for $\mathscr{H}$.
This is clearly a concrete category. Let us verify that it is elementary:

Injectiveness: Isometries are injective.

Tarski-Vaught property: Clear.

Elementary chain property: Clear.

Amalgamation: This should probably be in every textbook of functional analysis. However, as I know none, I give it:

The normed case: Suppose that $f_i : A \to B_i$ are isometries for $i < 2$. Let $\bar{A}, \bar{B}_i$ be the completions of A, B_i respectively. Then we can extend f_i to $f_i' : \bar{A} \to \bar{B}_i$ in a unique way, and we may define $C = \bar{B}_0 \oplus_{\bar{A}} \bar{B}_1$, the fibred sum as vector spaces. Define for $b_i \in \bar{B}_i$: $\|b_0 - b_1\|_C = \inf_{a \in \bar{A}}(\|b_0 - a\|_{\bar{B}_0} + \|b_1 - a\|_{\bar{B}_1})$. Then it is a well defined function from C to $\mathbb{R}_+$. It is a norm:

$\|x\|_C \geq 0$: Clear.

$\|x\|_C = 0 \Longrightarrow x = 0$: Suppose that $b_i \in \bar{B}_i$, and $\|b_0 - b_1\|_C = 0$. Then for all n there is $a_n \in \bar{A}$ such that $\|b_i - a_n\| < 1/n$. Thus $b_0 = b_1 \in \bar{A}$, and $b_0 - b_1 = 0$ in C.

$\|\alpha a\| = |\alpha|\|a\|$: Clear.

Triangle inequality:

$$
\begin{aligned}
\|b_0 + b_0' - b_1 - b_1'\|_C &= \inf_{a\in\bar{A}}(\|b_0 + b_0' - a\|_{\bar{B}_0} + \|b_1 + b_1' - a\|_{\bar{B}_1}) \\
&= \inf_{a,a'\in\bar{A}}(\|b_0 + b_0' - a - a'\|_{\bar{B}_0} + \|b_1 + b_1' - a - a'\|_{\bar{B}_1}) \\
&\leq \inf_{a,a'\in\bar{A}} (\|b_0 - a\|_{\bar{B}_0} + \|b_0' - a'\|_{\bar{B}_0} + \\
&\qquad\qquad \|b_1 - a\|_{\bar{B}_1} + \|b_1' - a'\|_{\bar{B}_1}) \\
&= \inf_{a\in\bar{A}}(\|b_0 - a\|_{\bar{B}_0} + \|b_1 - a\|_{\bar{B}_1}) \\
&\quad + \inf_{a'\in\bar{A}}(\|b_0' - a'\|_{\bar{B}_0} + \|b_1' - a'\|_{\bar{B}_1}) \\
&= \|b_0 - b_1\|_C + \|b_0' + b_1'\|_C
\end{aligned}
$$

Finally, the natural mappings of B_i into C are isometries: For $b \in B_0$:

$$
\begin{aligned}
\|b\|_{\bar{B}_0} &\geq \inf_{a\in\bar{A}}(\|b - a\|_{\bar{B}_0} + \|a\|_{\bar{B}_1}) \\
&= \inf_{a\in\bar{A}}(\|b - a\|_{\bar{B}_0} + \|a\|_{\bar{B}_0}) \\
&\geq \|b\|_{\bar{B}_0}
\end{aligned}
$$

and similarly for $b \in B_1$. Restricting back to A, B_i, we get the required amalgamation result. As a matter of notation, if $C' \subseteq C$ is the subspace spanned by the images of B_i, we note: $C' = B_0 \oplus_A^{\mathscr{B}} B_1$.

The inner product case is easier: Write $C = \bar{A} \oplus (\bar{B}_0 \cap A^\perp) \oplus (\bar{B}_1 \cap A^\perp)$, the direct sums being orthogonal, and take $C' = B_0 \oplus_A^{\mathscr{H}} B_1$ to be the subspace spanned in C by the B_i.

The characterisation of equality of types should be clear. As there is a unique type for the empty tuple, we have connectedness. We have yet to show that this is compact:

Type boundedness: Clear.

Type locality: If a, b are infinite tuples, and every two corresponding finite subtuples have the same type, then we have an isomorphism of the vector spaces generated by a and b sending a onto b. But then it preserves norms as well.

Weak compactness for Banach spaces: We shall prove a stronger result, that is the compactness of a space with a richer topology (or language). For every $n < \omega$, $\bar{\lambda} \in \mathbb{F}^n$ and $K \subseteq \mathbb{R}^+$ compact, we define the n-ary predicate $p^{\mathscr{B}}_{n,\bar{\lambda},K}$ by:

$$
p^{\mathscr{B}}_{n,\bar{\lambda},K}(\bar{x}) \iff \|\sum_{i<n} \lambda_i x_i\| \in K
$$

and we take $\mathcal{L}_{\mathscr{B}}$ to be the set of all such predicates. First, every complete n-type is the conjunction of (infinitely many) such predicates, so it suffices to show compactness in this language. We claim then that if $\Sigma(\bar{x})$ is a set of positive quantifier-free formulas in this language in some infinite tuple of variables, and it is finitely realized in $\mathscr{B}$, then it is realized in $\mathscr{B}$:

First, we may clearly assume that Σ is in fact a set of predicates. Let $V = \bigoplus \mathbb{F}x_i$ the space of linear combinations in $\bar{x}$, and let $W \subseteq V$ be the subspace generated by those linear combinations which are present in the predicates which appear in Σ. Choose a basis $\bar{y}$ for W, and find W' such that $V = W \oplus W'$ as vector spaces. The compactness of the sets K gives us for every y_i in this basis a bound $M_i \in \mathbb{R}_+$ and some finite $\Sigma^i \subseteq \Sigma$ such that $\Sigma^i \vdash \|y_i\| \leq M_i$. Now, for finite $\Sigma_0 \subseteq \Sigma$, the linear combinations of $\bar{x}$ that appear in Σ_0 generate a finite-dimensional subspace of W, so a finite subset $\{y_i : i \in I\}$ of the base suffices to generate it. Let $\Sigma_0' = \Sigma_0 \cup \bigcup_{i \in I} \Sigma^i \subseteq \Sigma$, so it is still finite, and by assumption there are $\bar{a} \in U \in \mathscr{B}$, $U \models \Sigma_0'(\bar{a})$. Let $\pi : V \to U$ send y_I onto a_I, and the rest of $\bar{y}$ along with W' to 0, and define $\rho_{\Sigma_0}(w) = \|\pi(w)\|_U$. Then ρ_{Σ_0} is a semi-norm on V such that $(V, \rho_{\Sigma_0}) \models \Sigma_0$, $\rho_{\Sigma_0}(y_i) \leq M_i$ for all i, and it is zero on W'. This means that as elements of $(\mathbb{R}^+)^V$, all the ρ_{Σ_0} belong to some compact, and therefore $\langle \rho_{\Sigma_0} : \Sigma_0 \subseteq \Sigma \text{ finite} \rangle$ has an an accumulation point ρ. ρ is a semi-norm as the limit of such, and as each predicates is a closed condition which is realized by ρ_{Σ_0} from some point on, we have $(V, \rho) \models \Sigma$. Dividing out by the elements of semi-norm zero gives the desired normed space.

Weak compactness for Hilbert spaces: We do a similar thing using the previous result, translating conditions on inner products to conditions on norms. For all $n < \omega$, $\bar{\lambda}, \bar{\mu} \in \mathbb{F}^n$, $X \subseteq \mathbb{F}$ closed and $N \in \mathbb{R}^+$, we define the n-ary predicate $p^{\mathscr{H}}_{n,\bar{\lambda},\bar{\mu},X,N}$ by:

$$p^{\mathscr{H}}_{n,\bar{\lambda},\bar{\mu},X,N}(\bar{x}) \iff \left\langle \sum_{i<n} \lambda_i x_i, \sum_{i<n} \mu_i x_i \right\rangle \in X \wedge \bigwedge_{i<n} \|x_i\| \leq N$$

and we take $\mathcal{L}_{\mathscr{B}}$ to be the set of all such predicates. Again, it suffices to prove compactness in this language. We claim that if $\Sigma(\bar{x})$ is a set of predicates of this form in some infinite tuple of variables, and it is finitely realized in $\mathscr{H}$, then it is realized in $\mathscr{H}$. This can be proved directly in a similar fashion as for normed spaces, but we can in fact reduce to that case. Recall that in an inner product space over $\mathbb{R}$:

$$\langle u, v \rangle = \frac{\|u+v\|^2 - \|u-v\|^2}{4}$$

And over $\mathbb{C}$:

$$\langle u, v \rangle = \frac{\|u+v\|^2 - \|u-v\|^2}{4} + i\frac{\|u+iv\|^2 - \|u-iv\|^2}{4}$$

So we see that the condition on the norm of being induced by an inner product is a closed one (that is, can be expressed by a conjunction of predicated of the form $p^{\mathscr{B}}_{n,\bar{\lambda},F}$, where $F \subseteq \mathbb{R}^+$ is closed), and so are the conditions $\left\langle \sum_{i<n} \lambda_i a_i, \sum_{i<n} \mu_i a_i \right\rangle \in X$. As they are not compact, we also need the bounds on the norms of the participating variables. Combining the closed conditions with these bounds we obtain compact conditions, that is a set $\Sigma'(\bar{x}) \subseteq \mathcal{L}_{\mathscr{B}}$, which is finitely realized in $\mathscr{H}$ and therefore in $\mathscr{B}$. Then by the compactness for Banach spaces it is realized in $\mathscr{B}$. But as Σ' contains predicates that say that the norm is induced by an inner product, it is in fact realized in $\mathscr{H}$ and we are done.

QED

Remark 2.40. Note that we just gave a richer topology than the minimal one. We did not change the type spaces as sets.

Remark 2.41. With the languages we gave for $\mathscr{B}$ and $\mathscr{H}$, they are not Hausdorff, nor indeed semi-Hausdorff. In fact, not only is the language for $\mathscr{B}$ not semi-Hausdorff, but we also cannot expand it to be (while keeping compactness). Indeed, if we add predicates for type equality, then the set $\{x_i \equiv x_{i+1} : i < \omega\} \cup \{\|x_{i+1} - x_i\| = i : i < \omega\}$ is finitely realised but not realised.

Remark 2.42. Replace the language we gave for $\mathscr{B}$ with the somewhat weaker language consisting of the predicates:

$$q^{\mathscr{B}}_{n,\bar{\lambda},X,N}(\bar{x}) \iff \|\sum_{i<n} \lambda_i x_i\| \in X \wedge \bigwedge_{i<n} \|x_i\| \leq N$$

and add predicates for equality of types. Then $\mathrm{S}(\mathscr{B})$ is compact with the induced topology, and in particular $\mathscr{B}$ is semi-Hausdorff in this language. Indeed, assume we are given a finitely consistent set Σ of predicates in this language. Then we have two kinds of predicates: norm predicates and type-equality predicates. The norm predicates (in the modified language) give us bounds on the norm of every variable they apply to. Then, if we have a bound on $\|x\|$ and we know that $x \equiv y$ then we also have a bound on $\|y\|$. Let $x_i : i < \lambda$ be the variables on whose norms we have bounds, and $x'_i : i < \lambda'$ the others. Let W, W' be the spaces of linear combinations of these sets of variables respectively, and $V = W \oplus W'$. For every finite $\Sigma_0 \subseteq \Sigma$ we have a semi-norm ρ on V such that $(V, \rho) \models \Sigma_0$. Define $\rho'(w + w') = \rho(w)$ (where $w + w'$ is the decomposition of $v \in V = W \oplus W'$). One verifies that then $(V, \rho') \models \Sigma_0$. This means that $\Sigma' = \Sigma \cup \{\|x'_i\| = 0 : i < \lambda'\}$ is also finitely consistent, and now we have bounds on all the variables. We can then replace the type-equality conditions, which are closed, by compact ones and reduce to the language we gave in the beginning.
A similar argument shows that we can add type-equality predicates to the language we gave for $\mathscr{H}$ preserving compactness.
Thus, when needed, both $\mathscr{B}$ and $\mathscr{H}$ can be assumed to be semi-Hausdorff.

References

[Ben02a] Itay Ben-Yaacov, *Lovely pairs of models: the non first order case*, Preprint, 2002.
[Ben02b] Itay Ben-Yaacov, *Simplicity in compact abstract theories*, Preprint, 2002.
[Ben02c] Itay Ben-Yaacov, *Thickness, and a categoric view of type-space functors*, Preprint, 2002.
[BL00] Steven Buechler and Olivier Lessmann, *Simple homogeneous models*, Preprint, 2000.
[Hru97] Ehud Hrushovski, *Simplicity and the Lascar group*, Preprint, 1997.
[Pil00] Anand Pillay, *Forking in the category of existentially closed structures*, Connections between Model Theory and Algebraic and Analytic Geometry (Angus Macintyre, ed.), Quaderni di Matematica, vol. 6, University of Naples, 2000.
[She75] Saharon Shelah, *The lazy model-theoretician's guide to stability*, Logique et Analyse **71–72** (1975).
[She99] Saharon Shelah, *Categoricity for abstract classes with amalgamation*, Annals of Pure and Applied Logic **98** (1999), 261–294.

Itay Ben-Yaacov, Équipe de Logique Mathématique, UFR de Mathématiques (case 7012), Université Paris 7, 2 place Jussieu, 75251 Paris Cedex 05, France
E-mail address: pezz@logique.jussieu.fr
URL: http://www.logique.jussieu.fr/www.pezz