THE BULLETIN OF SYMBOLIC LOGIC
Volume 11, Number 1, March 2005

COMPACTNESS AND INDEPENDENCE IN NON FIRST ORDER FRAMEWORKS

ITAY BEN-YAACOV

Abstract. This communication deals with positive model theory, a non first order model theoretic setting which preserves compactness at the cost of giving up negation. Positive model theory deals transparently with hyperimaginaries, and accommodates various analytic structures which defy direct first order treatment. We describe the development of simplicity theory in this setting, and an application to the lovely pairs of models of simple theories without the weak non finite cover property.

§1. Background. Every first order theory admitting quantifier elimination (and up to a change of language, we may always assume that it does) is the model companion of a universal theory (in fact, its model completion). The converse is known to be false, namely, not every universal theory even has a first order model companion: it has one if and only if the class of its existentially closed (e.c.) models is elementary. Nevertheless, it was observed by more than one person that much of first order model theory can be repeated in the class of e.c. models of a universal theory T, whether this class is elementary or not; how much of classical model theory can be repeated may depend on additional conditions one imposes on T. This can be viewed as studying the class of models of the model companion of T, even though such may not exist as a first order theory.

A universal theory has a model completion if and only if it first has a model companion, and second, the model companion eliminates quantifiers. The second property can be re-stated in a way which does not depend on the first: a universal theory having this property is called a *Robinson theory*. Robinson theories were defined (with a little more generality) and studied by Hrushovski in [19]. From a model-theoretic point of view, the class of e.c. models of a Robinson theory behaves very much like the class of models of a first order theory with quantifier elimination, and thus the framework of Robinson theories is just a small step outside the scope of first order theories.

Received July 10, 2004; revised October 3, 2004.

Most of the results appearing in this communication come from the first part of the author's doctoral thesis [2]. They also appear in separate articles. A survey containing results from the second part of [2] appears in [14].

1079-8986/05/1101-0002/$3.30

Simplicity and stability theory extend to Robinson theories in an obvious and straightforward manner. Robinson theories were used, among other things, to provide an example of a simple theory where the Lascar strong type differs from that of Shelah (the existence of a first order theory with this property is still open).

In [22], Pillay studies the class of e.c. model of an arbitrary universal theory (i.e., without assuming it is Robinson), and re-develops simplicity therein. Again, one may view this as working with the "model companion" of an arbitrary universal theory.

Positive model theory, with which this communication is concerned, takes Hrushovski's approach one step farther away from first order model theory, while in the same time generalising that of Pillay. The original motivation for its development is the fact that first order simple theories give rise to hyperimaginary elements and sorts which cannot be transparently adjoined as new sorts in the same manner that imaginary sorts can. The irritating thing about them is that we have a fairly precise idea of the "logic" of hyperimaginary elements, what complete and partial types are, and we even know that compactness holds for this logic. It's just that a partial type in a hyperimaginary sort is a set of formulas of a particular kind, one may call them "positive formulas", for whose negations compactness may fail, so considering hyperimaginaries in a first order language would be *wrong*. Once one accepts the fact that hyperimaginary elements are external to the structure, having a (compact!) logic of their own, one may proceed and work with them more or less as one would with ordinary imaginary elements (see [17], [28]).

Positive model theory attempts to remedy this deficiency of first order logic by repeating Hrushovski's definition of a Robinson theory in a positive language, i.e., a language without negation, resulting in the notion of a *positive Robinson theory*. This provides the sought-after model theoretic framework in which hyperimaginaries are not segregated against, while the compactness theorem still holds for positive formulas. In addition to hyperimaginary sorts, positive model theory accommodates many classes of analytic structures, and one can prove that in some sense it is the most general possible framework in which compactness holds.

Finally, it is curious to note that all of these recently studied settings turn out to be re-discoveries of settings which already appeared in some form or another in Shelah's 1970s paper [24], and remained relatively dormant since. In that paper, Shelah studies stability in various kinds of classes of structures, and in corresponding universal domains (defined by satisfying an assumption corresponding to the kind of the class). Thus the classes of Kind II are the e.c. models of a Robinson theory, and a universal domain satisfying Assumption II is a universal domain for such a theory. Similarly, Pillay's category of e.c. structures is precisely a class of Kind III, and its universal domains indeed satisfy Assumption III, but the class of *all* homogeneous structures satisfying Assumption III is actually larger, and is equal to that of universal domains of positive Robinson theories.

For the sake of simplicity, most of the development is done assuming a single-sorted language. In the case of multi-sorted language, the "length" of a tuple also contains the specific sort of each position in the tuple (which is why we sometimes refer to the length of a tuple as its *sort*). Lowercase letters a, b, c, etc., denote possibly infinite tuples of elements in structures, while x, y, z denote possibly infinite tuples of variables. When a lowercase letter denotes a singleton we say so explicitly. If we want to make the length, or index set, of a tuple explicit, we may write it as $a_{<\alpha}$ or $a_{\in I}$ (which are shorthand for $(a_i : i < \alpha)$ and $(a_i : i \in I)$, respectively).

§2. Three presentations of cats.

Although positive Robinson theories may be the most down-to-earth definition of a "theory" in the context of positive model theory, there are other, fundamentally different, ways of presenting such theories, namely compact type-space functors and compact abstract elementary categories. From a presentation of any kind (positive Robinson theory, compact type-space functor, or compact abstract elementary category) we can construct corresponding ones of the other two kinds in a manner which is unique and reversible (up to a natural equivalence). We therefore view them as manifestations of a more fundamental notion which we call a *compact abstract theory*, or *cat*:

(i) As we said earlier, positive Robinson theories are defined along the same lines as Hrushovski's Robinson theories, with the only difference being that the set of basic formulas Δ is only assumed to be closed under *positive* boolean combinations. Thus again we work in the category of e.c. models of a universal theory, or in a universal domain for such a category, but the precise meaning of these notions needs to be slightly altered so as to accommodate the fact that Δ may no longer be closed under negations.

(ii) Alternatively, we may specify the theory by specifying its type-space functor: this consists of specifying its type-spaces as pure topological spaces, with mapping between them describing restriction of types to sub-tuples etc. This is a rather abstract approach from which syntax is entirely absent.

(iii) Finally, we may start out with a mere concrete category, and observe that it satisfies some of the properties of the category of e.c. models of a positive Robinson theory. If it satisfies enough of these properties, and in particular a relatively weak instance of compactness, then in a natural sense it is equivalent to a positive Robinson theory T: under a canonical translation, it has the same universal domains as T and the same type spaces.

The last approach is quite useful when one wishes to get an initial grip on the model-theoretic structure (if such exists) that a specific notion of embedding induces on a given class of mathematical structures. It also suggests that in some sense, the cat setting is the most general one in which compactness holds.

This section contains a relatively high concentration of definitions; it is mostly extracted from [4].

2.1. Positive Robinson theories and universal domains. In order to specify a positive Robinson theory, one must first choose a language. Unlike the first order setting, when one chooses a signature $\mathcal{L}$ and then works with the entire first order language $\mathcal{L}_{\omega,\omega}$, we need to make a further choice by which we restrict the set of formulas we consider (that is, our language) to a subset $\Delta \subseteq \mathcal{L}_{\omega,\omega}$. It is only for formulas in Δ that we expect the compactness principle to hold. (One may note that this differs from most non first order approaches in which one usually extends the set of formulas under consideration to a proper *superset* of $\mathcal{L}_{\omega,\omega}$.)

The definitions here were motivated by the setting of Robinson theories in [19], and are almost identical with the exception that we do not require Δ to be closed under negations:

DEFINITION 2.1. Let $\mathcal{L}$ be a signature, and let $\mathcal{L}_{\omega,\omega}$ be the set of all first order $\mathcal{L}$-formulas written using the connectives $\neg$, $\vee$, $\wedge$ and the quantifier $\exists$.

(i) A *positive fragment* (in $\mathcal{L}$) is a subset $\Delta \subseteq \mathcal{L}_{\omega,\omega}$, containing all the atomic formulas, and closed under positive boolean combinations and for sub-formulas (the restriction on the connectives and quantifier we may use serves to fine-tune the notion of sub-formula; this is required for the existence of e.c. models and the positive Morleyisation of a negative universal theory).

(ii) Given a positive fragment Δ, we define:

$$\Sigma = \Sigma(\Delta) = \{\exists y\, \varphi(x, y) : \varphi \in \Delta\},$$
$$\Pi = \Pi(\Delta) = \{\forall y\, \neg\varphi(x, y) : \varphi \in \Delta\} = \{\neg\psi : \psi \in \Sigma\}.$$

We call members of Δ *basic positive formulas*; or rather, we should: we usually just call them formulas (that is to say that by a *formula*, without further qualification, we mean a member of Δ). Members of Σ and Π are *positive existential* and *negative universal* formulas, respectively.

(iii) If M is an $\mathcal{L}$-structure and $a \in M$ is a tuple, its Δ-*type* in M is $\mathrm{tp}^M_\Delta(a) = \{\varphi(x) \in \Delta : M \models \varphi(a)\}$. The Σ-type of a in M is defined similarly.

We identify any two formulas in $\mathcal{L}_{\omega,\omega}$ which are logically equivalent, which is why we consider the two enumerations of Π above equal. This also means that Σ is a positive fragment in its own right.

Most of the time, the signature $\mathcal{L}$ and the positive fragment Δ are fixed by the context. By a *structure* we always mean an $\mathcal{L}$-structure.

Even though we claim to be working in the category of e.c. models of a theory T (to be defined more precisely below), most of the time we will just be working inside a big universal domain (just as in first order model theory, we usually work inside a big monster model). Thus, even though we discuss e.c. models later on, the definition of a universal domain would suffice for most practical purposes.

Definition 2.2. Let M and N be structures, $A \subseteq M$ a subset, and $f : A \to N$ a mapping (i.e., $f : M \dashrightarrow N$ is a partial mapping with $\mathrm{dom}(f) = A$). Then of f is a *partial* (Δ-)*homomorphism* if for every tuple $a \in A$ and every formula $\varphi(x) \in \Delta$: $M \models \varphi(a) \implies N \models \varphi(f(a))$ (in other words, $\mathrm{tp}^M_\Delta(A) \subseteq \mathrm{tp}^N_\Delta(f(A))$).

If f is total, namely $A = M$, then $f : M \to N$ is a (Δ-)*homomorphism*; if $M = N$ then f is a (partial) *endomorphism*; etc.

Definition 2.3. Let κ be a relatively big cardinal (at the very least $\kappa > |\Delta|$), and U a structure. Then U is a κ-*universal domain* if it satisfies the following properties:

(i) κ-Homogeneity: Let $f : U \dashrightarrow U$ be a partial endomorphism of U, and assume that $|\mathrm{dom}(f)| < \kappa$. Then f extends to an automorphism of U.

(ii) κ-Compactness: Let Γ be a set of formulas (from Δ!), possibly in infinitely many free variables, such that $|\Gamma| < \kappa$, and assume furthermore that every finite subset of Γ is realised in U. Then Γ is realised in U. (Whether we allow parameters from U in Γ or not is immaterial and does not affect the definition.)

Usually, however, we omit κ and simply say that U is a universal domain; set whose cardinality is less than κ are called *small*. We convene that κ is so big that every set of parameters or formulas we might consider is indeed small: for example, one might assume that κ is strongly inaccessible, although this is not strictly necessary.

The following are among the most often-used properties of universal domains:

Proposition 2.4. *Let U be a universal domain, $a \in U$ a tuple, and $T = \mathrm{Th}_\Pi(U)$ (i.e., the set of Π-sentences true in U). Then*:

(i) *Assume that $U \models \neg\exists y\, \varphi(a, y)$, where $\varphi \in \Delta$ (so $\exists y\, \varphi \in \Sigma$). Then there is a formula $\psi(x) \in \Delta$ such that $U \models \psi(a)$ and $\forall xy\, \neg(\varphi(x, y) \wedge \psi(x)) \in T$.*

(ii) *If $b \in U$ is of the same length as a and $\mathrm{tp}^U_\Delta(a) \subseteq \mathrm{tp}^U_\Delta(b)$, then $\mathrm{tp}^U_\Delta(a) = \mathrm{tp}^U_\Delta(b)$, and moreover $\mathrm{tp}^U_\Sigma(a) = \mathrm{tp}^U_\Sigma(b)$.*

(iii) *A small set of formulas Γ (without parameters) is realised in U if and only if $T \cup \Gamma$ is consistent. If Γ contains parameters, let us write it as $\Gamma(x, a)$, and let $p(y) = \mathrm{tp}^U_\Delta(a)$: then $\Gamma(x, a)$ is realised in U if and only if $T \cup \Gamma(x, y) \cup p(y)$ is consistent.*

Thus in particular, the Π-theory of a universal domain is an important invariant: if U and U' are two universal domains with the same Π-theory, then the same things "happen" in both (formally: they realise the same types), and therefore for our purposes they are essentially equivalent.

This allows us to give now the following tentative definition. A better (and more general) one appears later on.

DEFINITION 2.5. Let U be a universal domain, and $T = \mathrm{Th}_\Pi(U)$. Then we say that U is a universal domain *for* T.

A *complete positive Robinson theory* is a theory of the form $\mathrm{Th}_\Pi(U)$ where U is a universal domain.

The rest of this subsection aims for a better and more general description of positive Robinson theories and their relations with universal domains. At this point we need to divert somewhat from the classical notion of an e.c. model of a universal theory:

DEFINITION 2.6. (i) Let T be a Π-theory (i.e., consisting of sentences of the form $\forall x\, \neg\varphi(x)$, where $\varphi \in \Delta$). An *existentially closed* (*e.c.*) model of T is a model $M \models T$ satisfying that whenever $N \models T$ as well, $f : M \to N$ is a Δ-homomorphism, $a \in M$ and $\varphi(x, y) \in \Delta$, if $N \models \exists y\, \varphi(f(a), y)$ then $M \models \exists y\, \varphi(a, y)$.

(ii) The category of e.c. models of T is denoted by $\mathcal{M}(T)$.

(iii) A mapping $f : M \to N$ between two structures M and N is a (Δ-)*embedding* if for all $a \in M$ and $\varphi \in \Delta$: $M \models \varphi(a) \iff N \models \varphi(f(a))$ (i.e., $\mathrm{tp}^M_\Delta(a) = \mathrm{tp}^N_\Delta(f(a))$).

First, we should remark that if Δ is the set of all quantifier free $\mathcal{L}$-formulas then a Δ-homomorphism is just an embedding of $\mathcal{L}$-structures (and coincides with Δ-embedding), Π is the set of all universal formulas, and the definition of an e.c. model of a universal theory above coincides with the ordinary one. Thus the classical definitions extend naturally to the positive case, where the Π-theory T serves as a restraining agent which states what cannot happen, and an e.c. model of T is one where every *positive* existential property which is not forbidden is true.

Note that since the image of a homomorphism may satisfy more relations than the domain, a homomorphism needs not be injective. On the other hand, if $M, N \in \mathcal{M}(T)$ and $f : M \to N$ is a homomorphism, then it is an embedding (and in fact a Σ-embedding), so there is no ambiguity about the notion of a morphism between two e.c. models of T (in fact, we only need $N \models T$ here).

The assumption that Δ is closed for sub-formulas (and the Axiom of Choice) are used to prove:

PROPOSITION 2.7. *Let T be a Π-theory. Then for every model $M \models T$ there exist an e.c. model $N \in \mathcal{M}(T)$ and a morphism $M \to N$.*

The morphisms in the category $\mathcal{M}(T)$ play the role of elementary embeddings in first order model theory, and share some of their properties:

Proposition 2.8. *The category $\mathcal{M}(T)$ has the amalgamation property and the increasing chain property*:

(i) *Let $M, M_0, M_1 \in \mathcal{M}(T)$, and $f_i : M \to M_i$ be morphisms for $i < 2$. Then there exist N and morphisms $g_i : M_i \to N$ such that the following diagram commutes*:

$$\begin{tikzcd} & M_0 \arrow[dr, dotted, "g_0"] & \\ M \arrow[ur, "f_0"] \arrow[dr, "f_1"'] & & N \\ & M_1 \arrow[ur, dotted, "g_1"'] & \end{tikzcd}$$

(ii) *Let $(M_i : i < \alpha)$ be an increasing chain of e.c. models of a Π-theory T, where the inclusions respect Δ. Then the union $M = \bigcup_{i<\alpha} M_i$ is an e.c. models of T, and each inclusion $M_i \subseteq M$ respects Δ.*

We can now give a more direct definition of positive Robinson theories:

Definition 2.9. A *positive Robinson theory* is a Π-theory T satisfying any of the following equivalent conditions:

(i) If $M, N \in \mathcal{M}(T)$ and $a \in M$, $b \in N$ are tuples of the same length such that $\mathrm{tp}_\Delta^M(a) \subseteq \mathrm{tp}_\Delta^N(b)$, then $\mathrm{tp}_\Sigma^M(a) = \mathrm{tp}_\Sigma^N(b)$ (i.e., the Δ-type determines the Σ-type; but notice that in the assumption, only inclusion is required).

(ii) Whenever $M_0, M_1 \in \mathcal{M}(T)$ and $f : M_0 \dashrightarrow M_1$ is a partial Δ-homomorphism with domain $A \subseteq M_0$, there exists $N \in \mathcal{M}(T)$ and morphisms $g_i : M_i \to N$ such that $g_1 \circ f = g_0{\restriction_A}$:

$$\begin{tikzcd} & M_0 \arrow[dr, dotted, "g_0"] & \\ A \arrow[ur, "\subseteq"] \arrow[dr, "f"'] & & N \\ & M_1 \arrow[ur, dotted, "g_1"'] & \end{tikzcd}$$

(iii) If $M \in \mathcal{M}(T)$, $a \in M$ is a tuple, and $M \not\models \exists y\, \varphi(a, y)$ for some $\varphi \in \Delta$, then there exists $\psi(x) \in \Delta$ such that $M \models \psi(a)$ and $T \vdash \forall xy\, \neg(\varphi(x, y) \wedge \psi(x))$.

These equivalent properties stand in clear analogy with properties of universal domains, whereby it is clear that a complete positive Robinson theory, as defined in Definition 2.5, is indeed positive Robinson as defined here. The converse is not true, since not every positive Robinson theory is complete:

DEFINITION 2.10. A Π-theory T is *complete* if is the Π-theory of a structure (so a complete positive Robinson theory, as defined in Definition 2.5, is indeed complete).

If T is not complete, then a *completion* of T is a minimal (with respect to inclusion) complete Π-theory containing T. In this case, a universal domain of T is a universal domain for any of its completions, i.e., a universal domain whose Π-theory is a completion of T.

THEOREM 2.11. (i) *The completions of a Π-theory are precisely the various theories of its e.c. models.*
(ii) *A Π-theory is positive Robinson if and only if all its completions are.*
(iii) *A complete Π-theory is positive Robinson if and only if it has a universal domain (i.e., it is the Π-theory of a universal domain). Thus the definitions given here agree with Definition 2.5.*

Since a universal domain is easily verified to be an e.c. model of its Π-theory, the above can be re-stated as: T is positive Robinson if and only if every e.c. model of T embeds in a universal domain which is itself an (e.c.) model of T.

Earlier approaches appear as special cases of positive Robinson theories:
(i) In case Δ is closed under negation, positive Robinson theories coincide with Hrushovski's Robinson theories.
(ii) In case Δ is the set of all existential $\mathcal{L}$-formulas, we obtain the setting studied by Pillay in [22]. Since in this case $\Delta = \Sigma$, every universal theory is positive Robinson.
(iii) Finally, if $\Delta = \mathcal{L}_{\omega,\omega}$, we just get the classical first order setting.

2.2. Types and type spaces. Let T be a positive Robinson theory.

By a *partial Δ-type* (*Σ-type*) we mean a set of Δ-formulas (Σ-formulas), and we usually omit the Δ- prefix.

For simplicity of the exposition let add the assumption that for every partial Σ-type there is a partial Δ-type which has the same realisations in every e.c. model of T. In this case we identify the two partial types. (In [4], a theory T with this property was called *strongly* positive Robinson.) This is a harmless assumption for two main reasons:

– First, in most (if not all) of the examples we consider this is true.

– Second, replacing Δ with Σ does not modify the notions of a Π-theory, nor does it modify the category $\mathcal{M}(T)$ or the class of universal domains of T (since T is positive Robinson). On the other hand, it would make the notions of Δ-type and Σ-type agree, which is what we want.

Recall that if $M \in \mathcal{M}(T)$ and $a \in M$ is a tuple, then its type $\mathrm{tp}_\Delta^M(a)$ is the set of all Δ-formulas satisfied by a in M. We will henceforth write it simply as $\mathrm{tp}(a)$, since M should be clear from the context, and embedding it in a bigger e.c. model of T would not change the type of a. If b is any other tuple in M and $p(x, y) = \mathrm{tp}(a, b)$, then the type of *a over b* is $p(x, b) = \mathrm{tp}(a/b)$.

For every index set I we define the space of complete I-types $\mathrm{S}_I(T)$ in the natural way: its members are all the possible types of I-tuples in e.c. models of T, or equivalently the set of all maximal partial types in an I-tuple of free variables consistent with T. For a partial type $q(x)$ (where $x = x_{\in I} = (x_i : i \in I)$), define:

$$\langle q \rangle = \{p(x) \in \mathrm{S}_I(T) : q \subseteq p\}$$

We topologise $\mathrm{S}_I(T)$ by defining every set of the form $\langle q \rangle$ as closed. This is a compact T_1 topology. Similarly, if b is a tuple in an e.c. model M, then $\mathrm{S}_I(b)$ is the set of all maximal partial types with parameters in b which are finitely realised in M, or equivalently all types over b of n-tuples in e.c. models extending M. We topologise it similarly.

For infinite I and b we have:

$$\mathrm{S}_I(T) = \varprojlim_{I_0 \subseteq I \text{ finite}} \mathrm{S}_{I_0}(T) \qquad \mathrm{S}_I(b) = \varprojlim_{I_0 \subseteq I, b_0 \subseteq b \text{ finite}} \mathrm{S}_{I_0}(b_0),$$

so we may restrict our attention to finite sets of indices I, or just to natural numbers.

If $\mathrm{S}_n(T)$ is Hausdorff for every $n < \omega$, then so are $\mathrm{S}_I(T)$ and $\mathrm{S}_I(b)$ for arbitrary (possibly infinite) I and b. In this case we say that T is *Hausdorff*.

If $f : I \to J$ is any mapping, it defines a natural mapping $f^* : \mathrm{S}_J(T) \to \mathrm{S}_I(T)$ by:

$$f^*(\mathrm{tp}(a_j : j \in J)) = \mathrm{tp}(a_{f(i)} : i \in I).$$

This is a closed continuous mapping of topological spaces. All this information can be put into one object, a contra-variant functor from (finite) sets to topological spaces, denoted $\mathrm{S}(T)$, which sends a number n to $\mathrm{S}_n(T)$ and a mapping f to f^*. We call $\mathrm{S}(T)$ the *type-space functor* of T.

Note that applied to injective and surjective mappings, the type-space functor produces existential quantification and equality, respectively: If $f : I \to J$ is an inclusion, then $f^*(\langle \varphi(x_{\in I}) \rangle) = \langle \exists x_{\in J \setminus I}\, \varphi(x_{\in I}) \rangle$. On the other hand, let $f_{eq} : 2 \to 1$ be the only possibility. Then $f^*_{eq}(\mathrm{S}_1(T)) = \langle x_0 = x_1 \rangle$. Also, the completions of T are in a natural bijections with the set $\mathrm{S}_\varnothing(T)$.

FACT 2.12. *The functor* $\mathrm{S}(T)$ *has the following properties*:

(i) *Amalgamation*: *for any two sets of indices* I *and* J, *the mapping* $\mathrm{S}_{I\cup J}(T) \to \mathrm{S}_I(T) \times_{\mathrm{S}_{I\cap J}(T)} \mathrm{S}_J(T)$ *is surjective. In other words, given two complete types* $p(x_{\in I})$ *and* $p'(x_{\in J})$, *whose restrictions to the common variables* $x_{\in I\cap J}$ *coincide, there is a type* $p''(x_{\in I\cup J})$ *whose restrictions to* $x_{\in I}$ *and* $x_{\in J}$ *are* p *and* p', *respectively.*

(ii) *Preservation of equality*: *let* I *be a set of indices, and* $h : 2 \to I$ *be any mapping. Let* $J = I/\{h(0) = h(1)\}$, *and* $f : I \to J$ *the projection. Recall that* $f_{eq} : 2 \to 1$ *is the unique such mapping, and let* $eq = f^*(\mathrm{S}_1(T))$ *(so* $eq = \langle x_0 = x_1 \rangle$*). Then* $f^*(\mathrm{S}_J(T)) = (h^*)^{-1}(eq)$.

The first property can be viewed as a re-statement of the fact that the type (i.e., Δ-type) of a tuple determines its existential (i.e., Σ-) type. The second is a highly complicated way of saying an obvious thing, namely that if $p(x, y, \bar{z})$ is a complete type (here x and y are singletons, and $\bar{z}$ a tuple), and $p'(x, y) = p\restriction_{x,y}$, then $p \vdash x = y$ if and only if $p' \vdash x = y$. It is nevertheless required for the following converse result:

Definition 2.13. Let S be a contra-variant functor from the category of finite sets to that of compact T_1 topological spaces, where the morphisms are the closed continuous mappings. Assume furthermore that S satisfies the two properties of Fact 2.12. Then S is a *compact type-space functor*.

Theorem 2.14. *Let S be a compact type-space functor. Then there exists a positive Robinson theory T such that S is isomorphic* (*i.e.*, *homeomorphic*) *to* $\mathrm{S}(T)$.

More precisely, for every $n < \omega$ choose a basis of closed sets B_n for $S(n)$. Let $\mathcal{L} = \{R_X : n < \omega, X \in B_n\}$, where for each $X \in B_n$, R_X is an n-ary relations symbol. Let Δ be the set of all positive quantifier-free $\mathcal{L}$-formulas. Then in this positive fragment there exists a unique positive Robinson theory (up to logical equivalence) T such that there exists an isomorphism of functors Φ: $\mathrm{S}(T) \simeq S$, such that in addition, for every $p \in \mathrm{S}_n(T)$ and $X \in B_n$: $p \in X \iff R_X(x) \in p$.

Finally, various special cases can be discerned directly from topological properties of the type-spaces:

Definition 2.15. Let S be a compact type-space functor.

(i) It is *open* if for every n, the natural projection $S_{n+1} \to S_n$ (i.e., $(n \hookrightarrow n + 1)^*$) is open. This implies that f^* is open for every injective $f : I \to J$.
(ii) It is *totally disconnected* if S_n is totally disconnected for every n.
(iii) It has *positive inequality* if $f^*_{eq} : S_1 \to S_2$ an open mapping. This implies that f^* is open for every surjective $f : m \to n$, *for finite m and n*.

Theorem 2.16. *Let S be a compact type-space functor.*

(i) *It is* (*homeomorphic to*) *the type-space functor of a Robinson theory if and only if it is totally disconnected and has positive inequality.*
(ii) *It is* (*homeomorphic to*) *the type-space functor of a first order theory if and only if it is open, totally disconnected and has positive inequality.*

Remark 2.17. One conclusion could be made that the property identifying first order logic is the "openness". Indeed, being open and Hausdorff characterises the type-space functors of theories in a continuously-valued variant of first order logic which is under active study at the time of writing.

2.3. Compact abstract elementary categories. If earlier we described essential properties of $\mathrm{S}(T)$, the type-space functor of a positive Robinson theory T, here we would like to point out the essential properties of $\mathcal{M}(T)$, the category of e.c. models of T. The goal is slightly different: rather than trying to characterise categories which are isomorphic, in some sense, to some $\mathcal{M}(T)$, we give more relaxed criteria that characterise categories which, like $\mathcal{M}(T)$, have (compact) universal domains.

We follow a development similar to that of abstract elementary classes (see, for example, [25]); however, since our goals are different, and for the reasons mentioned above, our setting is more relaxed. In particular, since we want to *recover* a logical structure, rather than be given one in advance, our departure point is the relatively general notion of a concrete category:

Definition 2.18. 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:

(i) Whenever $A, B \in \mathcal{C}$, and $f : |A| \to |B|$ is a mapping, 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.
(ii) 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* is a functor $F : \mathcal{C} \to \mathcal{C}'$, such that $|\cdot|_{\mathcal{C}'} \circ F = |\cdot|_{\mathcal{C}}$.

In other words, a concrete category is a category of sets with additional structure.

We now impose conditions which are analogous to those of an abstract elementary class with amalgamation:

Definition 2.19. Let $\mathcal{M}$ be a concrete category. Call its objects models, and its morphisms ($\mathcal{M}$-)elementary mappings. Write $M \leq_{\mathcal{M}} N$ if $M \subseteq N$ (i.e., $|M| \subseteq |N|$) and the inclusion is elementary. Then $\mathcal{M}$ is an *abstract 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 $M_0 \leq_{\mathcal{M}} N$, $M_1 \leq_{\mathcal{M}} N$, then $M_0 \leq_{\mathcal{M}} M_1$.
Amalgamation: Whenever $f_i \in \mathrm{Hom}_{\mathcal{M}}(M, M_i)$ for $i < 2$, there are $N \in \mathcal{M}$ and $g_i \in \mathrm{Hom}_{\mathcal{M}}(M_i, N)$ such that the following diagram commutes:

$$\begin{tikzcd} & M_0 \arrow[dr, dotted, "g_0"] & \\ M \arrow[ur, "f_0"] \arrow[dr, "f_1"'] & & N \\ & M_1 \arrow[ur, dotted, "g_1"'] & \end{tikzcd}$$

Elementary chain property: Direct limits of chains exist in the category $\mathcal{M}$ (in [4] appears a slightly more relaxed yet more complicated version of this condition).

(Compare the amalgamation and chain condition with Proposition 2.8.)

We now proceed to define an abstract notion of type of a tuple in models of $\mathcal{M}$:

DEFINITION 2.20. Let $\mathcal{M}$ be an abstract elementary category with amalgamation. Let I be a (possibly infinite) set of indices, $M, N \in \mathcal{M}$. For two I-tuples $a \in M$ and $b \in N$, say that $(M, a) \equiv_{\mathcal{M}} (N, a)$ if there is $P \in \mathcal{M}$ and morphisms $f : M \to P$ and $g : N \to P$ such that $f(a) = g(b)$.

This is an equivalence relation (follows from amalgamation), and the equivalence class $(M, a)/{\equiv}$ is called the *type* (in some texts *Galois type*) of the tuple a in the model M, and is denoted $\mathrm{tp}^M_{\mathcal{M}}(a)$. As usual, we may omit $\mathcal{M}$ and even M if they are clear from the context.

Note that if T is a positive Robinson theory, then $\mathcal{M}(T)$ is an abstract elementary category with amalgamation, and the two notions of type coincide: $\mathrm{tp}^M_\Delta(a) = \mathrm{tp}^N_\Delta(b)$ if and only if $\mathrm{tp}^M_{\mathcal{M}}(a) = \mathrm{tp}^N_{\mathcal{M}}(b)$.

The collection of all I-types (i.e., types of I-tuples) is denoted $\mathrm{S}_I(M)$, and if $f : I \to J$ is any mapping then $f^* \colon \mathrm{S}_J(\mathcal{M}) \to \mathrm{S}_I(\mathcal{M})$ is defined as before. The final condition is essentially that $\mathrm{S}(\mathcal{M})$ can be rendered a type-space functor:

DEFINITION 2.21. A *compact abstract elementary category* is a abstract elementary category with amalgamation satisfying in addition the following conditions:

(i) For finite I, the collection $\mathrm{S}_I(\mathcal{M})$ is a set (rather than a proper class).

(ii) For infinite I: $\mathrm{S}_I(\mathcal{M}) = \varprojlim_{I_0 \subseteq I \text{ finite}} \mathrm{S}_{I_0}(\mathcal{M})$.

(iii) Let $x_{\in I}$ be a possibly infinite tuple of variables. Let $\Gamma(x_{\in I})$ be a set of things of the form $p(x_{i_0}, \ldots, x_{i_{n-1}})$, where $n < \omega$, $p \in \mathrm{S}_n(\mathcal{M})$, and $i_j \in I$ for all $j < n$. Say that Γ is realised in $\mathcal{M}$ if there are $M \in \mathcal{M}$ and a tuple $a_{\in I} \in M$ such that $\mathrm{tp}(a_{i_0}, \ldots, a_{i_{n-1}}) = p$ whenever $p(x_{i_0}, \ldots, x_{i_{n-1}}) \in \Gamma$.

Then we require that for every such Γ, if every finite subset $\Gamma_0 \subseteq \Gamma$ is realised in $\mathcal{M}$, then so is Γ.

The first two conditions say that $\mathrm{S}(\mathcal{M})$ is a functor to the category of sets, and as such satisfies the non-topological properties of a type-space functor: we may say is it is a *set* type-space functor (rather than a compact topological one). Modulo this, the third condition is equivalent to: $\mathrm{S}(\mathcal{M})$ admits a topology rendering it a compact type-space functor (if there is such a topology, then there is always a weakest one, and the compactness condition of Definition 2.21 consists of verifying that the would-be weakest such topology is indeed compact).

This equivalent condition is sometimes easier to verify: in particular, if T is a positive Robinson theory, then $\mathrm{S}(T) \simeq \mathrm{S}(\mathcal{M}(T))$ as functors to sets, whereby $\mathcal{M}(T)$ is a compact abstract elementary category. Conversely, if $\mathcal{M}$ is a compact abstract elementary category, then by Theorem 2.14 there is a positive Robinson theory T such that $\mathrm{S}(T) \simeq \mathrm{S}(\mathcal{M})$. It is not true that $\mathcal{M}(T) \simeq \mathcal{M}$, but a sufficient approximation for that does hold:

THEOREM 2.22. *Let $\mathcal{M}$ be a compact abstract elementary category, and choose a topology rendering $\mathrm{S}(\mathcal{M})$ a compact type-space functor. Let $\mathcal{L}$, Δ and T be as in Theorem 2.14. For every $M \in \mathcal{M}$, make it an $\mathcal{L}$-structure in the natural manner: if $a \in M^n$, $X \subseteq \mathrm{S}_n(\mathcal{M})$ and $R_X \in \mathcal{L}$, say that $M \models R_X(a)$ if and only if $\mathrm{tp}^M(a) \in X$.*

Then:

(i) *If $M, N \in \mathcal{M}$ and $f : M \to N$ is a mapping of the underlying sets, then $f \in \mathrm{Hom}_{\mathcal{M}}(M, N)$ if and only if it is a Δ-homomorphism.*
(ii) *$T = \mathrm{Th}_{\Pi}(\mathcal{M})$, where $\mathcal{M}$ is viewed as a class of $\mathcal{L}$-structures.*
(iii) *Every $M \in \mathcal{M}$ embeds in an $N \in \mathcal{M}(T)$, and* vice versa.

Thus, the same things "happen" in $\mathcal{M}$ and in $\mathcal{M}(T)$, and we conclude that for our purposes they are equivalent.

Alternatively, Theorem 2.22 implies that every $M \in \mathcal{M}$ embeds in a universal domain of T, and conversely, if U is a universal domain of T and $A \subseteq U$ is any (small) set, then there is a (small) $M \in \mathcal{M}$ such that $A \subseteq M \subseteq U$ (this can stated more precisely, and be viewed as a downward Löwenheim-Skolem property). Thus the universal domains of $\mathcal{M}(T)$ are in a sense the universal domains of $\mathcal{M}$.

2.4. Examples. As we said above, if Δ is either closed under negations we obtain the framework of Robinson theories from [19]; if Δ is the set of all existential $\mathcal{L}$-formulas then we get the framework described in [22]; and if $\Delta = \mathcal{L}_{\omega,\omega}$ then a universal domain is just a monster model, and we get classical first order model theory.

At the other extreme is the case where Δ is minimal, i.e., the set of positive quantifier-free $\mathcal{L}$-formulas. Most of the examples we give will be of this form. Also, every other example can be transformed into this form via *positive Morleyisation*:

PROPOSITION 2.23. *Let $\Delta \subseteq \mathcal{L}$ be any positive fragment, and T a positive Robinson theory. Let $\mathcal{L}'$ consist of $\mathcal{L}$ along with an n-ary predicate R_φ for every n-ary $\varphi \in \Delta$, and let Δ' be the minimal positive $\mathcal{L}'$-fragment. For a $\mathcal{L}$-structure M, let $\theta(M)$ be the $\mathcal{L}'$-expansion obtained by defining $R_\varphi^{\theta(M)} = \varphi(M)$ for each $\varphi \in \Delta$. Let $T' = \mathrm{Th}_{\Pi'}(\{\theta(M) \colon M \in \mathcal{M}(T)\})$ (where $\Pi' = \forall\neg\Delta'$). Then T' is a positive Robinson theory (in Δ'), and θ induces a bijection between $\mathcal{M}(T)$ and $\mathcal{M}(T')$.*

(Thus T' eliminates both quantifiers and negation from T.)

The original motivation for studying positive Robinson theories was to find a way to adjoin hyperimaginary sorts the same way that we adjoin imaginary sorts to a first order theory:

PROPOSITION 2.24. *Let T be a positive Robinson theory in $\Delta \subseteq \mathcal{L}_{\omega,\omega}$ (so in particular T could be a first order theory). Let $E(\bar{x}, \bar{y})$ be a partial type which defines an equivalence relation on α-tuples in e.c. models of T.*

For every Δ-formula $\varphi(x_0, \ldots, x_{k-1}, \bar{y}_0, \ldots, \bar{y}_{l-1})$, where each x_i is a singleton and $\bar{y}_j$ is an α-tuple (there may be infinitely many dummy variables in this way of writing φ), define a new predicate symbol $R_\varphi(x_0, \ldots, x_{k-1}, z_0, \ldots, z_{l-1})$ where the z_j are singletons in a new sort. Let $\mathcal{L}'$ consist of $\mathcal{L}$ along with the new predicates, and Δ' the minimal positive fragment.

For every $M \in \mathcal{M}(T)$, let $\theta(M)$ be the $\mathcal{L}'$-structure obtained by adding M^α/E to M as the new sort, and defining $\theta(M) \models R_\varphi(a_0, \ldots, a_{k-1}, c_0, \ldots, c_{l-1})$ if and only if each equivalence classes c_i has a representative $\bar{b}_i$ such that $M \models \varphi(a_0, \ldots, a_{k-1}, \bar{b}_0, \ldots, \bar{b}_{l-1})$. Then $T' = \mathrm{Th}_{\Pi'}(\{\theta(M) \colon M \in \mathcal{M}(T)\})$ is a positive Robinson theory, and θ is a bijection between $\mathcal{M}(T)$ and $\mathcal{M}(T')$.

Of course, we can do the same for several type-definable equivalence relations, or even for all of them at the same time. In case T is a first order theory, the logic we obtain for hyperimaginary sorts coincides with the logic appearing in [17], [28].

An important class of examples consists of various analytic structures.

For example, Banach structure theories in the sense of C. Ward Henson's logic of positive bounded formulas and approximate satisfaction [18] can be viewed as a special case of cats. If T is such a theory in a positive bounded language $\mathcal{L}$, then the class of approximate models of T, equipped with approximate-elementary embeddings, is a compact abstract elementary category and therefore defines a cat in an appropriate language. To obtain the language, we follow a procedure similar to that used above: for each positive bounded formula we create a new predicate symbol, and obtain a language $\mathcal{L}'$, a minimal positive fragment $\Delta' \subseteq \mathcal{L}'_{\omega,\omega}$, and a positive Robinson theory T' which is the Π'-theory of the class of approximate models of T viewed as $\mathcal{L}'$-structures. (If T has quantifier elimination, then it suffices to create predicates for quantifier-free formulas only.) Then T and T' have the same type-spaces (in the sense of the respective logics), and therefore share the same universal domains. It is not true, though, that the Banach spaces which are approximate models of T are precisely the e.c. models of T'. This deficiency of the notion of e.c. model is addressed in [10] in the form of *complete model*, and exceeds the scope of this communication. Numerous examples of this kind exist, and let us just mention the theories of Hilbert spaces whose hyperimaginaries are characterised in [11], and L_p Banach lattices [12].

Probability measure algebras form an important analytic example which is not a Banach structure:

EXAMPLE 2.25. Let U be the measure algebra of $[0,1]^\kappa$, i.e., the boolean algebra of Borel sets modulo null-measure sets, equipped with the induced measure function μ. Consider it as a structure in the language $\mathcal{L} = \{0, 1, \wedge, \vee, \cdot^c, \text{“}\mu(\cdot) \leq r\text{”}, \text{“}\mu(\cdot) \leq r\text{”}\}$ where r varies over $\mathbb{Q} \cap [0,1]$. Let Δ be the minimal positive fragment. Provided κ is large enough, U is a κ-universal domain.

This example is the topic of [8].

§3. **Simplicity.** Most of the results discussed here are presented in fuller detail in [5]. For a general reference on first order simplicity theory, one may refer to [27].

In this section we work with a cat T: tuples of elements are assumed to be taken from a κ-universal domain of T, where κ is bigger than the cardinality of any set or tuple we may consider (it is convenient, although not necessary, to assume that κ is strongly inaccessible).

Let $\langle a_i : i < \omega \rangle$ be a sequence of tuples of the same length, and c another tuple. We recall that the sequence $\langle a_i : i < \omega \rangle$ is *indiscernible* over c, or c-*indiscernible*, if for every $n < \omega$ and $i_0 < i_1 < \cdots < i_{n-1} < \omega$, $\operatorname{tp}(a_{i_0}, \ldots, a_{i_{n-1}}/c)$ depends only on n. Let $p(x, b)$ be a partial type with parameters in b. We say that $p(x, b)$ *divides* over another tuple c if there exists a c-indiscernible sequence $\langle b_i : i < \omega \rangle$, with $b_0 = b$, such that $\bigcup p(x, b_i)$ is inconsistent (with T). In case that a, b, c are tuples such that $\operatorname{tp}(a/bc)$ does not divide over c, we write this down as $a \mathop{\smile\!\!\!\!\!\!\mid}_c b$ (read: a is independent from b over c).

We say that T is *simple* if for any tuple b and any finite tuple (or equivalently, singleton) a, there is a sub-tuple $b' \subseteq b$ such that $|b'| \leq |T|$ and $a \mathop{\smile\!\!\!\!\!\!\mid}_{b'} b$. This property of non-dividing is called the *local character*. In a first order theory, one consequence of simplicity is the *extension axiom*: For every type $p \in \mathrm{S}(c)$ and tuple b there is $p \subseteq q \in \mathrm{S}(cb)$ such that q does not divide over c. This is not true in cats [5, Example 4.3], so Morley sequences do not necessarily exist, and the classical development of simplicity breaks down.

In [22] Pillay re-develops some of the theory of simplicity in the framework of the e.c. models of arbitrary universal theories. Even though his framework is a special case of cats, for the development of simplicity it presents the full problematic of the general case, and in particular the one mentioned above. This led Pillay to define a simple theory as one where both the local character *and* the extension axiom hold: this patches the hole in the classical argument, and with some improvements the rest follows through. However, there are many reasons why one may wish to drop the second assumption (i.e., extension): for example, [5, Example 4.3] mentioned above is stable, so it *should* be simple, although the extension axiom fails there.

In order to develop simplicity in theories where extension fails, we take recourse to the technical notion of *array-dividing*. Let us give a few details:

DEFINITION 3.1. (i) A *d-dimensional array* is something of the form $\langle a_\sigma : \sigma \in \omega^d \rangle$. It is b-indiscernible if for every $k < d$, the sequence $\langle \bar{a}_i^k : i < \omega \rangle$ is b-indiscernible, where $\bar{a}_i^k = \langle a_\sigma : \sigma \in \omega^d, \sigma(k) = i \rangle$ (in other words, if the sequence of hyperplanes in each direction is b-indiscernible).

(ii) A partial type $p(x, b)$ *array-divides* over c if there exist $d < \omega$ and a c-indiscernible array $\langle b_\sigma : \sigma \in \omega^d \rangle$ such that $b = b_{\bar{0}}$ and $\bigcup_\sigma p(x, b_\sigma)$ is inconsistent.

(iii) The notation $a \mathop{\smile\!\!\!\!\!\!\,|}^a_c b$ means that $\mathrm{tp}(a/bc)$ does not array-divide over c.

(To the best of our knowledge, indiscernible arrays were introduced in [20], where it was proved that in a simple first order theory array-dividing coincides with dividing.)

As dividing is a special case ($d = 1$) of array-dividing, it is easier to array-divide than to divide: thus, the local character of array-dividing seems *a priori* stronger than that of dividing. Luckily enough, though, they turn out to be equivalent:

PROPOSITION 3.2. *In a simple cat, array-dividing satisfies the local character.*

From now on, assume that T is simple.

As in the case of dividing, the local character of array-dividing is equivalent to the finiteness of a certain class of local ranks, $D(p, \varphi, \psi, d)$, which measure how many times we can still add to a partial type p instances of φ, which array-divide over prior parameters with a d-dimensional array in a way that is witnessed by ψ (here ψ plays the role of k in the $D(-, \varphi, k)$-ranks used in the first order setting; d has no counterpart, of course). Unlike the local D-ranks of dividing, the local D-ranks of array-dividing allow us to prove directly:

PROPOSITION 3.3. *Let a, b, c be tuples. Then the following imply one another from top to bottom*:

(i) $D(\mathrm{tp}(a/bc), \varphi, \psi, d) = D(\mathrm{tp}(a/c), \varphi, \psi, d)$ *for every* φ, ψ, d.

(ii) $a \mathop{\smile\!\!\!\!\!\!\,|}^a_c b$.

(iii) $D(\mathrm{tp}(b/ac), \varphi, \psi, d) = D(\mathrm{tp}(b/c), \varphi, \psi, d)$ *for every* φ, ψ, d.

In the classical development the analogue of the implication (ii) $\Longrightarrow$ (iii) is inevitably proved at a much later stage using transitivity and symmetry of non-dividing [21, Section 6]. Here we turn things around, and symmetry and transitivity are its (immediate) consequences:

COROLLARY 3.4. *By symmetry, the conditions of Proposition* 3.3 *are equivalent. Furthermore*:

(i) $a \mathop{\smile\!\!\!\!\!\!\,|}^a_c b \Longleftrightarrow b \mathop{\smile\!\!\!\!\!\!\,|}^a_c a$.

(ii) $a \mathop{\smile\!\!\!\!\!\!\,|}^a_c bd \Longleftrightarrow a \mathop{\smile\!\!\!\!\!\!\,|}^a_c b \wedge a \mathop{\smile\!\!\!\!\!\!\,|}^a_{bc} d$.

A somewhat less immediate consequence is:

PROPOSITION 3.5. *Dividing and array dividing define the same notion of independence*: $a \mathop{\smile\!\!\!\!\!\!|}^{a}_{c} b \iff a \mathop{\smile\!\!\!\!\!\!|}_{c} b$ *for every* a, b, c.

This ends the "detour", and from now on we forget array-dividing and only mention dividing. The next natural step is the amalgamation (or independence) theorem for Lascar strong types. Compared with the proofs in [21], [22], there are the added complications due to the fact that not all types have non-dividing extensions:

DEFINITION 3.6. (i) A type $p \in \mathrm{S}(c)$ is *extendible* if for every b there is $p \subseteq q \in \mathrm{S}(cb)$ which does no divide over c.

(ii) Let a, b be two tuples of the same length, and c any tuple. We say that a and b have the same *Lascar strong type* over c, in symbols $a \equiv^{\mathrm{Ls}}_{c} b$, if they are equivalent according to every c-invariant equivalence relation with a bounded ("small") set of equivalence class.

We can then prove the amalgamation theorem for *extendible* Lascar strong types (see below for precise statement). For this to be interesting, extendible types must exist:

PROPOSITION 3.7. *A type p is extendible if and only if all (one) of its non-dividing extensions are. Every type over a sufficiently saturated sub-model of the universal domain is extendible.*

Thus, there is a relatively large class of tuples every type over which is extendible. We formalise this notion by saying that the class of enumerations of sufficiently saturated models is co-final:

DEFINITION 3.8. A *co-final* class of tuples $\mathcal{A}$ is one which is invariant under permutations of the tuples and automorphisms of the universal domain, and such that for every tuple a there is a tuple $b \in \mathcal{A}$ extending a.

We can now sum it all up in:

THEOREM 3.9. *Let T be simple, and let $\mathcal{A}$ be the a co-final class of enumerations of sufficiently saturated models. Then*:

(i) *Invariance*: $a \mathop{\smile\!\!\!\!\!\!|}_{c} b$ *depends only on* $\mathrm{tp}(a, b, c)$, *and is invariant under permutations of each tuple.*
(ii) *Symmetry*: $a \mathop{\smile\!\!\!\!\!\!|}_{c} b \iff b \mathop{\smile\!\!\!\!\!\!|}_{c} a$.
(iii) *Transitivity*: $a \mathop{\smile\!\!\!\!\!\!|}_{c} bd \iff a \mathop{\smile\!\!\!\!\!\!|}_{c} b \wedge a \mathop{\smile\!\!\!\!\!\!|}_{bc} d$.
(iv) *Finite character*: $a \mathop{\smile\!\!\!\!\!\!|}_{c} b$ *if and only if* $a' \mathop{\smile\!\!\!\!\!\!|}_{c} b$ *for every finite* $a' \subseteq a$.
(v) *Existence*: $a \mathop{\smile\!\!\!\!\!\!|}_{c} c$.
(vi) *Local character*: *For every two tuples a and b there exists $b' \subseteq b$ such that* $|b'| \leq |T| + |a|$ *and* $a \mathop{\smile\!\!\!\!\!\!|}_{b'} b$.
(vii) *$\mathcal{A}$-Extension*: *If* $a \mathop{\smile\!\!\!\!\!\!|}_{c} b$ *and* $c \in \mathcal{A}$, *then for every d there exists a' such that* $\mathrm{tp}(a'/bc) = \mathrm{tp}(a/bc)$ *and* $a' \mathop{\smile\!\!\!\!\!\!|}_{c} bd$.

(viii) *$\mathcal{A}$-Amalgamation: Assume that $c \in \mathcal{A}$, $a_0 \mathop{\smile\!\!\!\!\!|}_c a_1$, $b_0 \equiv_c^{\mathrm{Ls}} b_1$, and for $i < 2$ we have $b_i \mathop{\smile\!\!\!\!\!|}_c a_i$. Then there exists b such that $b \mathop{\smile\!\!\!\!\!|}_c a_0a_1$ and $b_i \equiv_{ca_i}^{\mathrm{Ls}} b_i$ for $i < 2$.*

Conversely, for any cat T, if there exist a co-final class $\mathcal{A}$ and a notion of independence $\mathop{\smile\!\!\!\!\!|}'$ satisfying the axioms above, then T is simple and $\mathop{\smile\!\!\!\!\!|}'$ coincides with $\mathop{\smile\!\!\!\!\!|}$.

(The converse part for first order theories appears in [21].)

In fact, the amalgamation theorem is slightly stronger than $\mathcal{A}$-amalgamation as stated above: instead of assuming that $c \in \mathcal{A}$, it suffices to assume that $\mathrm{tp}(b_i/c)$ is extendible. It follows that restricted to a single extendible type over c, the relation $\equiv_c^{\mathrm{Ls}}$ is type-definable.

DEFINITION 3.10. A type $p \in \mathrm{S}(c)$ is called an *amalgamation base* if it is an extendible Lascar strong type (i.e., $a, b \models p \Longrightarrow a \equiv_c^{\mathrm{Ls}} b$).

An amalgamation base satisfies the amalgamation theorem: whenever $a_0 \mathop{\smile\!\!\!\!\!|}_c a_1$, and $p \subseteq p_i \in \mathrm{S}(ca_i)$ do not divide over c for $i < 2$, then there is $p_0 \cup p_1 \subseteq q \in \mathrm{S}(ca_0a_1)$ which does not divide over c.

Types over sufficiently saturated sub-models of U are amalgamation bases, and every extendible type has a non-dividing extension which is. In fact, if we allow hyperimaginary elements, then every extension of an extendible type $p \in \mathrm{S}(c)$ to $\mathrm{bdd}(c)$ is (non-dividing and) an amalgamation base. (The adjunction of hyperimaginary sorts does not alter in the least the simplicity of T, or the notion of non-dividing, so everything we said above extends to them smoothly). In [17], hyperimaginary canonical bases for amalgamation bases (or parallelism classes) were constructed in first order simple theories. Although the proof there cannot be translated to the positive setting, its underlying spirit still holds and yields the same result:

DEFINITION 3.11. Let $p \in \mathrm{S}(a)$ be an amalgamation base. A (possibly hyperimaginary) element $c \in \mathrm{dcl}(a)$ is a *canonical base* for p, denoted $c = \mathrm{Cb}(p)$, if the following hold:

(i) p does not divide over c.
(ii) $p{\restriction}_c$ is an amalgamation base.
(iii) If $b \in \mathrm{dcl}(a)$ and p does not divide over b, then $c \in \mathrm{bdd}(b)$; if in addition $p{\restriction}_b$ is an amalgamation base, then $c \in \mathrm{dcl}(b)$.

Note that this does entirely characterise the canonical base of p, but rather characterises it up to interdefinability.

THEOREM 3.12. *Let $p \in \mathrm{S}(a)$ be an amalgamation base. Then there exists a type-definable equivalent relation $\sim$ on the sort of a such that $a_\sim = \mathrm{Cb}(p)$, where $a_\sim$ is the class of a modulo $\sim$, viewed as a hyperimaginary.*

There is very little question on what the correct definition of stability should be: counting types, definability of types, finiteness of local R-ranks and no ordered indiscernible sequence all agree (for matters of stability in several non first order contexts one may refer to [24]). Since we define simplicity without requiring that every type be extendible, it is (as one would expect) a generalisation of stability:

THEOREM 3.13. *A cat T is stable if and only if it is simple, and every amalgamation base is stationary (i.e., has a unique non dividing extension to every superset of its domain).*

§**4. Thickness.** Simplicity in cats, as described in the previous section, is not entirely satisfactory. For example, the following are some extensively used properties of simple first order theories which could fail in simple cats:

FACT 4.1. *Let T be a simple first order theory. Then:*

(i) *Every complete type is extendible.*
(ii) *For every complete type $p(x) \in \mathrm{S}(c)$ and tuple y there exists a partial type $q(x, y, c)$ such that $a, b \models q$ if and only if $a \models p$, and $b \mathop{\smile\!\!\!\!\!\!|}_c a$ (in the terminology of* [1], *we would say that complete types have* definable independence).
(iii) *Equality of Lascar strong types (over a tuple c) is type-definable (with parameters in c).*
(iv) *Every type-definable (or even hyperdefinable* [28]*) groups have generic elements (and various other properties of such groups hold).*

For an example of a definable group without generic elements in a cat, see [6, Proposition 1.22] *and the discussion which precedes it.*

Surprisingly (or not), the possible failure of each of the items of Fact 4.1 can be traced to a single pathology, namely the failure of *thickness*. Thick cats are discussed in [6]:

DEFINITION 4.2. Let a, b be two tuples of the same length. Say that $d(a, b) \leq 1$ if a, b are the first two tuples in an infinite indiscernible sequence.

Say that T is *thick* if for every tuple length, the property $d(x, y) \leq 1$ is type-definable.

The terminology comes from the notion of a *thick* formula, originally defined by Ziegler [16]. A formula $\varphi(x, y)$ (where x and y are of the same length) is *thick* (*thin*) if whenever $\langle a_i : i < \omega \rangle$ is an indiscernible sequence then $\models \varphi(a_0, a_1)$ ($\not\models \varphi(a_0, a_1)$). It follows that $d(a, b) \leq 1$ if and only if $\not\models \varphi(a_0, a_1)$ for every thin φ. In first order, a formula φ is thick if and only if $\neg\varphi$ is thin, so $d(x, y) \leq 1$ is defined by the partial type consisting of all thick formulas in x, y. A thick cat is thus one in which there are "sufficiently many" thick formulas, so that the same holds:

PROPOSITION 4.3. *For a cat T the following are equivalent:*

(i) *T is thick (i.e., $d(x, y) \leq 1$ is type-definable).*

(ii) *For every thin formula $\varphi(x, y)$ there is a thick formula $\psi(x, y)$ contradicting it.*

(iii) *For every tuple of variables x there is a partial type $\pi(x_{<\omega})$ (where each x_i is of the same length as x) such that $a_{<\omega} \models \pi$ if and only if $\langle a_i : i < \omega \rangle$ is indiscernible (in other words: indiscernibility is type-definable).*

THEOREM 4.4. *In a thick simple cat, the conclusions of Fact 4.1 hold.*

Conversely, if equality Lascar strong types over $\varnothing$ is type-definable (a special case of Fact 4.1(iii)), then T is thick: a sequence $\langle a_i : i < \omega \rangle$ is indiscernible if and only if, for every n and $i_0 < \cdots < i_{n-1}$, the Lascar strong type of $a_{i_0} \ldots a_{i_{n-1}}$ (over $\varnothing$) depends only on n.

Almost all the examples for cats are thick (and the ones which aren't were constructed specifically for this purpose). Moreover, every Hausdorff cat is thick, and more generally, if equality of types is a type-definable relation in T then T is thick. Thus thickness is a very mild assumption which one may feel free to use.

Moreover, many unrelated results which were proved in the context of first order theories can be seen, upon careful reading, to use only thickness assumption, and therefore hold in any thick simple cat. One example for this are the series of articles on group constructions in simple theories [1, 15, 3, 9]. Another is the theory of lovely pairs discussed below.

§5. An application: lovely pairs. The notion of a *lovely pair* of models of a simple first order theory was defined in [13]:

DEFINITION 5.1. A *pair* is something of the form (A, P) where A is a subset of a universal domain of T, and P a unary predicate on A, defining an acl-closed (or bdd-closed, in the case of cats) subset.

A pair (M, P) is κ-*lovely* if for every $A \subseteq M$, $|A| < \kappa$ and $p \in S_n(A)$:

(i) There exists $a \models p$ in M such that $a \mathop{\smile\!\!\!\!\!\!\mid}_A P(M)$.

(ii) If p does not divide over $P(A)$, then there exists $a \models p$ in $P(M)$.

It follows that both M and $P(M)$ are κ-saturated models of T.

This generalises (almost) the notion of a beautiful pair of stable structures from [23], and that of a generic pairs of models of an SU-rank 1 theory from [26].

If T is complete then all the lovely pairs are elementarily equivalent as $\mathcal{L}_P$-structures, where $\mathcal{L}_P = \mathcal{L} \cup \{P\}$. The common theory T_P could, or not, be a means for the study of lovely pairs:

DEFINITION 5.2. Let $(A, P) \subseteq (B, P)$ be an extension of pairs. It is *free* if $A \mathop{\smile\!\!\!\!\!\!\mid}_{P(A)} P(B)$.

FACT 5.3. *The following conditions are equivalent (for a first order simple theory T):*

(i) *The κ-saturated models of T_P are precisely the κ-lovely pairs for every $\kappa > |T|$.*
(ii) *There is a $|T|^+$-saturated model of T_P which is a lovely pair.*
(iii) *The notion of elementary extension of models of T_P coincides with that of a free extension.*
(iv) *Every model of T_P embeds elementarily in a lovely pair.*

If either condition is true, we say that T has the weak non finite cover property (wnfcp).

If T has the wnfcp then the first order theory T_P "captures" the lovely pairs of T and can be used to study them; in the contrary case first order model theory is simply not equipped for the study of the lovely pairs of T.

A more general approach is proposed in [7]. From Fact 5.3(iii) we conclude that we need a "theory" for whose models an "elementary extension" is a free extension. This information is captured by letting $\mathfrak{P}$ be the category whose objects are the pairs, and whose morphisms are the free embeddings. It is a (long but) straightforward process to verify that $\mathfrak{P}$ is a compact abstract elementary category (Definition 2.21), during which we also learn the correct notion of types of tuples in lovely pairs. In fact, the assumption that T is a first order theory is nowhere used, and all we need to assume is that T is a thick cat. We then apply Theorem 2.22 to conclude:

THEOREM 5.4. *Let T be any simple thick cat. Then*:
(i) *$\mathfrak{P}$ is a compact abstract elementary category.*
(ii) *As a matter of notation, if $(A, P) \in \mathfrak{P}$ and $a \in A$, then $a^c = \mathrm{Cb}(a/P(A))$ (a^c is preserved by free embeddings). Then $\mathrm{tp}^{\mathfrak{P}}(a)$ determines and is determined by $\mathrm{tp}^T(a, a^c)$.*
(iii) *There is a natural language $\Delta^{\mathfrak{P}} \subseteq \mathcal{L}^{\mathfrak{P}}_{\omega,\omega}$, extending $\Delta \cup \{P\}$, which corresponds to a natural compact topology on $\mathrm{S}(\mathfrak{P})$. Let $T^{\mathfrak{P}} = \mathrm{Th}_{\Pi^{\mathfrak{P}}}(\mathfrak{P})$. Then $T^{\mathfrak{P}}$ is a thick positive Robinson theory, $\mathrm{S}(T^{\mathfrak{P}}) \simeq \mathrm{S}(\mathfrak{P})$, and for $\kappa > |T|$ the κ-saturated models of $T^{\mathfrak{P}}$ are precisely the κ-lovely pairs of T.*

The same independence-theoretic results proved for T_P when T is a first order theory with wnfcp can be proved for $T^{\mathfrak{P}}$ without any assumptions beyond T being thick and simple:

THEOREM 5.5. *We still assume that T is a thick and simple cat.*
(i) *$T^{\mathfrak{P}}$ is also (thick and) simple.*
(ii) *Independence is characterised as follows: if $U^{\mathfrak{P}}$ is a universal domain for $T^{\mathfrak{P}}$ and $a, b, c \in U^{\mathfrak{P}}$, then $a \mathop{\smile\!\!\!\!\!\!|}^{\mathfrak{P}}_c b$ (independent in the sense of $T^{\mathfrak{P}}$) if and only if $a \mathop{\smile\!\!\!\!\!\!|}_{cP(U)} b$ and $(ac)^c \mathop{\smile\!\!\!\!\!\!|}_{c^c} (ab)^c$ (here in the sense of T). (If the reader is worried about $P(U)$ not being a "small" set, the first condition can be equivalently be stated as $a \mathop{\smile\!\!\!\!\!\!|}_{c,(abc)^c} b$.)*
(iii) *If T is stable, supersimple, or one-based, then so is $T^{\mathfrak{P}}$.*

The characterisation of independence can be understood as: a is independent from b over c in the sense of $T^{\mathfrak{P}}$ if and only if, in the sense of T, they are independent both *over* P and *inside* P.

REFERENCES

[1] Itay Ben-Yaacov, *Group configurations and germs in simple theories*, ***The Journal of Symbolic Logic***, vol. 67 (2002), no. 4, pp. 1581–1600.

[2] ———, *Théories simples: constructions de groupes et interprétabilité généralisée*, ***Ph.D. thesis***, Université Paris VII – Denis Diderot, 2002.

[3] ———, *On the fine structure of the polygroup blow-up*, ***Archive for Mathematical Logic***, vol. 42 (2003), pp. 649–663.

[4] ———, *Positive model theory and compact abstract theories*, ***Journal of Mathematical Logic***, vol. 3 (2003), no. 1, pp. 85–118.

[5] ———, *Simplicity in compact abstract theories*, ***Journal of Mathematical Logic***, vol. 3 (2003), no. 2, pp. 163–191.

[6] ———, *Thickness and a categoric view of type-space functors*, ***Fundamenta Mathematicae***, vol. 179 (2003), pp. 199–224.

[7] ———, *Lovely pairs of models: the non first order case*, ***The Journal of Symbolic Logic***, vol. 69 (2004), no. 3, pp. 641–662.

[8] ———, *Schrödinger's cat*, submitted.

[9] ———, *Simple almost hyperdefinable groups*, submitted.

[10] ———, *Uncountable dense categoricity in cats*, submitted.

[11] Itay Ben-Yaacov and Alexander Berenstein, *Imaginaries in Hilbert spaces*, ***Archive for Mathematical Logic***, vol. 43 (2004), no. 4, pp. 459–466.

[12] Itay Ben-Yaacov, Alexander Berenstein, and C. Ward Henson, *Model-theretic independence in L_p Banach lattices*, in preparation.

[13] Itay Ben-Yaacov, Anand Pillay, and Evgueni Vassiliev, *Lovely pairs of models*, ***Annals of Pure and Aplied Logic***, vol. 122 (2003), pp. 235–261.

[14] Itay Ben-Yaacov, Ivan Tomašić, and Frank O. Wagner, *The group configuration in simple theories and its applications*, this Bulletin, vol. 8 (2002), no. 2, pp. 283–298.

[15] ———, *Constructing an almost hyperdefinable group*, to appear in the ***Journal of Mathematical Logic***.

[16] Enrique Casanovas, Daniel Lascar, Anand Pillay, and Martin Ziegler, *Galois groups of first order theories*, ***Journal of Mathematical Logic***, vol. 1 (2001), no. 2, pp. 305–319.

[17] Bradd Hart, Byunghan Kim, and Anand Pillay, *Coordinatisation and canonical bases in simple theories*, ***The Journal of Symbolic Logic***, vol. 65 (2000), pp. 293–309.

[18] C. Ward Henson, *Nonstandard hulls of Banach spaces*, ***Israel Journal of Mathematics***, vol. 25 (1976), pp. 108–144.

[19] Ehud Hrushovski, *Simplicity and the Lascar group*, unpublished, 1997.

[20] Byunghan Kim, *Forking in simple unstable theories*, ***Journal of the London Mathematical Society***, vol. 57 (1998), no. 2, pp. 257–267.

[21] Byunghan Kim and Anand Pillay, *Simple theories*, ***Annals of Pure and Applied Logic***, vol. 88 (1997), pp. 149–164.

[22] Anand Pillay, *Forking in the category of existentially closed structures*, ***Connections between model theory and algebraic and analytic geometry*** (Angus Macintyre, editor), Quaderni di Matematica, vol. 6, University of Naples, 2000.

[23] Bruno Poizat, *Paires de structures stables*, ***The Journal of Symbolic Logic***, vol. 48 (1983), no. 2, pp. 239–249.

[24] Saharon Shelah, *The lazy model-theoretician's guide to stability*, ***Logique et Analyse***, vol. 71–72 (1975), pp. 241–308.

[25] ———, *Categoricity for abstract classes with amalgamation*, ***Annals of Pure and Applied Logic***, vol. 98 (1999), pp. 261–294.

[26] Evgueni Vassiliev, *Generic pairs of SU-rank 1 structures*, ***Annals of Pure and Applied Logic***, vol. 120 (2002), pp. 103–149.

[27] Frank O. Wagner, ***Simple theories***, Kluwer Academic Publishers, 2000.

[28] ———, *Hyperdefinable groups in simple theories*, ***Journal of Mathematical Logic***, vol. 1 (2001), pp. 152–172.

UNIVERSITY OF WISCONSIN – MADISON
DEPARTMENT OF MATHEMATICS
480 LINCOLN DRIVE
MADISON, WI 53706–1388, USA
URL: http://www.math.wisc.edu/~pezz