

UNCOUNTABLE DENSE CATEGORICITY IN CATS

ITAY BEN-YAACOV

ABSTRACT. We prove that under reasonable assumptions, every cat (compact abstract theory) is *metric*, and develop some of the theory of metric cats. We generalise Morley’s theorem: if a countable Hausdorff cat T has a unique complete model of density character $\lambda \geq \omega_1$, then it has a unique complete model of density character λ for every $\lambda \geq \omega_1$.

INTRODUCTION

Łoś’s conjecture, subsequently known as Morley’s theorem, states that:

Let K be an elementary class in a countable language. Then K is categorical in one uncountable cardinal if and only if it is categorical in every uncountable cardinal.

This was subsequently generalised, with some variations, to uncountable languages, as well as to various kinds of non-elementary classes, and still serves as a first test-bed for many non-first-order frameworks.

The class of Hilbert spaces satisfies a variant of Łoś’s conjecture which is not covered by any previous result: A Hilbert space is uniquely characterised by its density character, provided the latter is uncountable. This “positive instance” of Łoś’s conjecture is peculiar for two reasons: First, we are dealing with a class of *complete* structures; since a countable increasing union of complete structures is not in general complete, this is not an abstract elementary class. Second, we measure the size of a structure by its density character rather than its cardinality. Thus our example does not fit in the abstract elementary classes programme, where most (if not all) of the work on uncountable categoricity to date has been done. It should rather be viewed in the framework for the model-theoretic study of Banach space structures set by Henson and Iovino in [Hen76, HI02], where the density character is indeed the “correct” measure of size.

The present paper has two main goals:

First, we set up a model theoretic framework that generalises simultaneously first order model theory and Henson’s logic for Banach space structures. In fact, the framework in

Date: January 26, 2005.

Key words and phrases. categoricity – compact abstract theories.

This is the result of research conducted at the Institut Girard Desargues (Université Lyon I), and at the University of Illinois at Urbana-Champaign supported CNRS-UIUC collaboration agreement. The author would like to thank Frank O. Wagner and C. Ward Henson for helpful discussions.

question has already been defined as *Hausdorff cats* in [Ben03a], where we were mostly interested in the properties of a big saturated universal domain. Here we show that every (countable) Hausdorff cat admits a metric, which is unique up to uniform equivalence, and has a natural class of complete sub-structures of the universal domain associated with it, which we call its *complete* models. We see that saturated Banach space structures (in the sense of Henson) are universal domains for Hausdorff cats, whose metric is the norm metric, and whose complete models are the complete Banach space structures approximately elementarily equivalent to the universal domain. The same holds for first order theories, only the metric is discrete. Using the density character as a measure for the size of a complete model we get a notion of categoricity which again generalises both the first order case and Henson’s definitions.

Second, with the above definitions, we state and prove Loś’s conjecture for arbitrary countable Hausdorff cats. Doing this we borrow ideas from Iovino [Iov96] and Shelah [She75].

In Section 1 we give a few reminders concerning the setting in which we work, and give some examples.

In Section 2 we introduce the abstract notion of *distance* and show that in most reasonable cases it corresponds to an actual definable metric.

In Section 3 we define the analogue of an elementary sub-model of the universal domain, and generalise basic first-order model-theoretic results to this context.

In Section 4 we see how the notion of positive distance allows us to refine some simplicity- and stability-related notions.

In Section 5 we put everything together and prove the main result, namely the uncountable categoricity theorem for Hausdorff cats.

1. PRELIMINARY REMARKS ABOUT THE FRAMEWORK

We aim for generality, and in particular we wish to obtain a theorem that *extends* the first-order version. Our first step therefore is to take a relatively general framework and look for natural assumptions that would allow us to carry out the argument.

The chosen framework is that of *compact abstract theories*, or *cats*. Cats were originally introduced and developed in [Ben03a, Ben03b, Ben03c] with the intention of finding a model-theoretic framework, which should be as general as possible while still allowing the development of simplicity theory: dividing, local ranks, independence, canonical bases, theory of definable groups, etc. Although this was indeed achieved using not much more than some compactness assumption, it did not give rise to any applicable theory of *supersimplicity*: classes of non-first-order structures (such as Hilbert spaces), which by every intuition should be superstable, did not seem to be so (at least not with the crude definitions we could give at that time). Also, with arbitrary cats there is no way to capture the notion of a “complete” model (one notion of a model is that of an e.c. model: indeed, Shelah proves in [She75] a variant of Morley’s theorem for this notion, but this is not what we are looking for).

On the other hand, real-life cats are not as wild as the most general case. In particular, their type-spaces are usually equipped with a Hausdorff (and compact) topology, a property which has several implications:

- First, it removes an annoying question concerning the “correct language” for a cat: given a category of structures that can be represented by a cat (a compact abstract elementary category [Ben03a]), there is always a minimal language with which this can be done, but it is not at all clear whether there is a maximal one; on the other hand, a compact Hausdorff topology is maximal among the compact topologies, so a Hausdorff cat is indeed equipped with a maximal (and therefore “correct”) language.
- Second, it implies *thickness* [Ben03c], which is required for a full development of simplicity theory.
- Third, and most importantly, we prove here that a Hausdorff cat (with a countable language) admits a definable metric which is unique up to uniform equivalence. This in turn allows us to define *complete models*, as well as to develop a satisfactory theory of supersimplicity, ω -stability, etc..

Thus we allow ourselves:

Convention 1.1. Throughout this paper we only consider Hausdorff cats, namely cats whose type-spaces are (compact and) Hausdorff.

The development of ω -stability in the presence of a metric is closely related to (and partially a generalisation of) Henson and Iovino’s model theory for Banach structures, and in particular Iovino’s development of ω -stability in [Iov96]. There are still a few essential differences in the approach:

- In Henson and Iovino’s treatment of Banach space structures, they consider the metric as an extra-logical piece of information, whereas here it is deduced from purely logical information. (In the special case of Banach space structures, all the definable metrics are uniformly equivalent to the norm.)
- Since all the definable metrics are uniformly equivalent, they all induce the same uniform structure on the space of types, whereas Henson and Iovino are more lenient and consider any uniform structure on the space of types satisfying some requirements. However, although various definitions make sense with their notion of uniform structure, the applications we need seem to require that the uniform structure be the one induced by the definable metric: Thus from our point of view there is no real loss of generality here, and there is actually a significant gain in simplicity of the exposition.
- Finally, Henson and Iovino work extensively with the *syntactic* notions of formulas and approximations, whereas what we do is independent of any particular choice of language. We find it many times convenient to ignore formulas and work instead with the purely *semantic* notions of open and closed sets in the logic topology. (This distinction is merely cosmetic.)

Given a Hausdorff cat T , there are several additional “nice” conditions that come to mind:

- (i) We say that T is *totally disconnected* if its type-spaces are. In other words, T is totally disconnected if there is a clopen basis for the type-space topology, and therefore a possible choice of language with negation (although equality may still be only type-definable). We may also say in this case that T *admits negation*.
- (ii) We say that T *has positive inequality*, or that it is *discrete*, if for every surjective map $f: n \rightarrow m$, the map $f^*: S_m(T) \rightarrow S_n(T)$ is open (for *finite* n ; we never require the infinite analogue of this condition). This is equivalent to equality being clopen, and implies that the universal domain is discrete in a sense that will be defined below, whence the terminology.
- (iii) We say that T is *open*, or that it *eliminates the universal quantifier*, if for every injective map $f: n \rightarrow m$, the map $f^*: S_m(T) \rightarrow S_n(T)$ is open. Here the finite case implies the infinite one. A cat T is open if and only if for every partial types $p(x, y)$, the property $\forall y p(x, y)$ is defined by a partial type in x .

For example, we show in [Ben03a] that first order cats are precisely those satisfying all three conditions, and Robinson theories are characterised by the first two. Hyperimaginary sorts on the other hand behave in an opposite manner: we do not expect them to preserve any of the first two conditions, but we prove in [Ben03c] that they preserve the third.

In this respect, the “analytic” examples for cats behave much like hyperimaginaries in first-order theories: they are not totally disconnected and inequality is far from being positive, but they are open.

Example 1.2. Every approximately elementary class of Banach space structures, presented as a cat (the universal domain being the unit ball of a saturated model, and the positive properties being those definable with positive bounded formulas) is (Hausdorff and) open. This is a consequence of Henson’s logic admitting universal quantification.

Example 1.3. Schrödinger’s cat \mathfrak{M} defined in [Ben] is open.

These examples, as well as the observation that hyperimaginary sorts preserve openness, lead us to the impression that openness is a natural assumption even for cats which are neither totally disconnected nor discrete; indeed, this assumption would make the statements and proofs of several of our results below somewhat simpler. Still, with some additional work (namely the introduction of the *Q-topology* in Section 3) we manage to prove our results in the non-open case as well.

Notation and terminology are pretty much standard.

For convenience, we will assume most of the time that there is a single *home sort*: in case there are several, some of the statements should be adapted accordingly. A *hyperimaginary sort* is the quotient of a possibly infinite tuples in the home sort by a type-definable equivalence relation. We use the term *sort* somewhat loosely: it can be

the home sort, a hyperimaginary sort, or in fact any (possibly infinite) tuple of such sorts: of course, a tuple of sorts can always be viewed as a hyperimaginary sort, so this is legitimate.

Accordingly, lowercase letters a, b , etc., denote elements in any sort (so in fact they may denote infinite tuples), and x, y , etc., denote variables. We consider each variable to be associated to a fixed sort, and may in fact use it to designate that sort.

A relation is *type-definable* if it is logically equivalent to a partial type. The distinction between this notion and that of a relation being *definable*, i.e., with a single formula, does not make sense here, and the latter will therefore be avoided.

2. DISTANCE AND TOPOLOGY

In this section we will ordinarily use the Hausdorff assumption through the following lemma:

Lemma 2.1. *Let $p_i(x_i, a_i)$ be partial types for $i < \alpha$, where each x_i is a sub-tuple of a big tuple of variables x . If $\bigwedge_{i < \alpha} p_i(x_i, a_i)$ is inconsistent, then there are formulas $\varphi_i(x_i, a_i)$ such that $p_i(x_i, a_i) \wedge \varphi_i(x_i, a_i)$ is contradictory for each $i < \alpha$, and $\models \bigvee_{i < \alpha} \varphi_i(x_i, a_i)$. Moreover, this can be done so that $\varphi_i(x_i, a_i)$ is the false formula \perp for all but finitely many $i < \alpha$.*

Proof. Then each $p_i(x_i, a_i)$ defines a closed set $K_i \subseteq S_x(\bar{a})$. Since $S_x(\bar{a})$ is compact and Hausdorff, and $\bigcap_i K_i = \emptyset$, we can find open neighbourhoods $U_i \supseteq K_i$ such that $\bigcap_i U_i = \emptyset$, and moreover $U_i = S_x(\bar{a})$ for all but finitely many $i < \alpha$. Each U_i is defined by $\neg q_i(x, \bar{a})$ for some partial type q_i , so each $q_i(x, \bar{a})$ contradicts $p_i(x_i, a_i)$, and $\models \bigvee_i q_i(x, \bar{a})$.

Let $x'_i = x \setminus x_i$, $r_i(\bar{z}, y_i) = \text{tp}(\bar{a}, a_i)$, and define $q'_i(x_i, y_i) = \exists x'_i, \bar{z} [r_i(\bar{z}, y_i) \wedge q_i(x, \bar{z})]$. Then $q'_i(x_i, a_i)$ still contradicts $p_i(x_i, a_i)$, and $\models \bigvee_i q'_i(x_i, a_i)$ (since $q_i(x, \bar{a}) \vdash q'_i(x_i, a_i)$). Finally, we can find a single formula $\varphi_i(x_i, a_i) \in q'_i(x_i, a_i)$ that contradicts $p_i(x_i, a_i)$, and $\models \bigvee_i \varphi_i(x_i, a_i)$ as required.

For the moreover part, recall that $U_i = S_x(\bar{a})$ for all but finitely many $i < \alpha$, and in such case we may take $q_i = q'_i = \varphi_i = \perp$. QED_{2.1}

2.1. Distance. As we do not wish to assume that a metric is given externally, we define an internal notion of abstract distance:

Definition 2.2. A *distance* is a reflexive type-definable relation $\varepsilon(x, y)$ (without parameters). The variables x and y are of course in the same sort, and we say that ε is a distance in that sort.

If $\varepsilon(x, y)$ is a distance we define $\overleftarrow{\varepsilon}(x, y) = \varepsilon(y, x)$. It is *symmetric* if $\overleftarrow{\varepsilon} = \varepsilon$. Otherwise we define $\overleftrightarrow{\varepsilon} = \varepsilon \wedge \overleftarrow{\varepsilon}$.

The minimal distance is $x = y$, and the maximal one the true formula \top . They will be denoted by 0 and ∞ , respectively.

We say that $\varepsilon \leq \varepsilon'$ if $\varepsilon \vdash \varepsilon'$, and $\varepsilon < \varepsilon'$ if $\varepsilon \subseteq \overset{\circ}{\varepsilon}'$, where we interpret ε and ε' as subsets of $S_2(T)$ and $\overset{\circ}{\varepsilon}'$ is the interior of ε' there. In other words, $\varepsilon < \varepsilon'$ if and only if there is a partial type π (without parameters) such that $\varepsilon \vdash \neg\pi \vdash \varepsilon'$, so $\pi \wedge \varepsilon$ is contradictory and $\vDash \varepsilon' \vee \pi$, and we denote this $\varepsilon <^\pi \varepsilon'$. In this case, by compactness there is a formula $\varphi \in \pi$ such that $\varepsilon <^\varphi \varepsilon'$ (partial types are always assumed to be closed under finite conjunction). Note that $<$ is neither reflexive nor anti-reflexive: $\varepsilon < \varepsilon$ holds if and only if ε is clopen (and then $\varepsilon <^{\neg\varepsilon} \varepsilon$). In particular $\infty < \infty$, and in fact $0 \leq \varepsilon < \infty$ for every distance ε . We call a distance ε *positive* if $\varepsilon > 0$.

Finally, we say that $d(a, b) \leq \varepsilon$ if $\text{tp}(a, b) \in \varepsilon$ and $d(a, b) < \varepsilon$ if $\text{tp}(a, b) \in \overset{\circ}{\varepsilon}$. The respective negations of these properties are denoted by $d(a, b) > \varepsilon$ and $d(a, b) \geq \varepsilon$. Note that the latter is type-definable.

Lemma 2.3. *Let ε be a distance, and $p(x, y)$ a partial type (without parameters) contradicting ε . Then there is a distance $\varepsilon' > \varepsilon$ such that $p(x, y)$ contradicts ε' ; if ε is symmetric, then ε' can also be taken to be symmetric.*

Proof. The asymmetric case follows from $S_2(T)$ being a normal topological space. If $\varepsilon' > \varepsilon$ and ε is symmetric, then $\overleftarrow{\varepsilon}' > \varepsilon$ as well, whence the symmetric case. QED_{2.3}

Definition 2.4. Let $(\varepsilon_i : i < \alpha)$ be distances in various sorts. Then their product is a distance in the corresponding α -tuple of sorts, defined by $\prod \varepsilon_i(x_{<\alpha}, y_{<\alpha}) = \bigwedge_{i < \alpha} \varepsilon_i(x_i, y_i)$.

Clearly, if $\varepsilon_i = \infty$ for all but finitely many $i < \alpha$, and the rest are positive, then $\prod \varepsilon_i$ is positive as well. The converse is:

Lemma 2.5. *Let ε be a distance in the sort of α -tuples. Then $\varepsilon > 0$ if and only if there are positive $(\varepsilon_i : i < \alpha)$, all but finitely many of which are equal to ∞ , and $\varepsilon \geq \prod \varepsilon_i$, and we can always take them to be symmetric.*

Proof. As we said above, if $(\varepsilon_i : i < \alpha)$ verify all the assumptions then $\prod \varepsilon_i > 0$, so $\varepsilon \geq \prod \varepsilon_i \implies \varepsilon > 0$.

Conversely, assume that $\varepsilon > 0$, and let $\varphi(x_{<\alpha}, y_{<\alpha})$ be such that $\varepsilon >^\varphi 0$. Then $\varphi(x_{<\alpha}, y_{<\alpha}) \wedge \bigwedge_{i < \alpha} x_i = y_i$ is contradictory, and by Lemma 2.1 there are formulas $\psi_i(x, y)$, all of which contradict $x = y$ and all but finitely many being the false formula \perp , such that $\vDash \neg\varphi \vee \bigvee_{i < \alpha} \psi_i(x_i, y_i)$. If $\psi_i = \perp$ then set $\varepsilon_i = \infty$, otherwise use Lemma 2.3 to find a distance $\varepsilon_i > 0$ contradicting ψ_i . Then $\prod \varepsilon_i \vdash \bigwedge_{i < \alpha} \neg\psi_i(x_i, y_i) \vdash \neg\varphi \vdash \varepsilon$, as required (and in fact, $\varepsilon >^\varphi \prod \varepsilon_i$). This still holds if we replace each ε_i with $\overleftarrow{\varepsilon}_i$. QED_{2.5}

Definition 2.6. Given two distances $\varepsilon, \varepsilon'$ we define:

$$(\varepsilon + \varepsilon')(x, y) = \exists z [\varepsilon(x, z) \wedge \varepsilon'(z, y)]$$

We write 2ε for $\varepsilon + \varepsilon$, etc..

The sum of two distances is a distance, and addition of distances is associative. On the other hand, addition is not commutative, and the sum of two symmetric distances

can be asymmetric. We could define $\overleftrightarrow{+}$ by $\varepsilon \overleftrightarrow{+} \varepsilon' = \overleftrightarrow{\varepsilon + \varepsilon'}$, which is commutative and yields symmetric distances, but would not be associative, which we find an even greater fault.

The triangle inequality is satisfied in the sense that $d(a, b) \leq \varepsilon \wedge d(b, c) \leq \varepsilon' \implies d(a, c) \leq \varepsilon + \varepsilon'$.

Lemma 2.7. *Assume that $\varepsilon_0 + \varepsilon_1 < \varepsilon$. Then there are $\varepsilon'_0 > \varepsilon_0$ and $\varepsilon'_1 > \varepsilon_1$ such that $\varepsilon'_0 + \varepsilon'_1 < \varepsilon$, and if ε_i is symmetric then so can be ε'_i .*

Proof. Assume that $\varepsilon_0 + \varepsilon_1 <^\varphi \varepsilon$, so $\varphi(x, z) \wedge \varepsilon_0(x, y) \wedge \varepsilon_1(y, z)$ is contradictory. By Lemma 2.1 there are formulas $\chi_i(x, y)$ contradicting ε_i for $i < 2$ such that $\models \neg\varphi(x, z) \vee \chi_0(x, y) \vee \chi_1(y, z)$. By Lemma 2.3 we can find $\varepsilon'_i > \varepsilon_i$ contradicting χ_i for $i < 2$, and then $\varepsilon'_0 + \varepsilon'_1 <^\varphi \varepsilon$ as well. If ε_i is symmetric, we may replace ε'_i with $\overleftrightarrow{\varepsilon'_i}$. QED_{2.7}

Corollary 2.8. *For every $\varepsilon > 0$ there is $\varepsilon' > 0$ such that $\varepsilon > 2\varepsilon'$.*

Proof. Apply Lemma 2.7 to $\varepsilon > 0 + 0$, and let $\varepsilon' = \varepsilon'_0 \wedge \varepsilon'_1$. QED_{2.8}

Lemma 2.9. *If ε is a distance then $\varepsilon = \bigwedge_{\varepsilon < \varepsilon'} \varepsilon'$; if ε is symmetric, the same holds if we only consider symmetric ε' .*

In particular, $x = y$ if and only if $d(x, y) \leq \varepsilon$ for every positive ε .

Proof. One inclusion is clear, the other is by Lemma 2.3. If ε is symmetric then $\varepsilon = \bigwedge_{\varepsilon < \varepsilon'} \overleftrightarrow{\varepsilon'}$. QED_{2.9}

2.2. Topology. Fix a sort (say the home sort), and let X be a set of elements in that sort. We define a subset $F \subseteq X$ as closed in the *logic topology* on X if it is defined in X by some partial type with parameters. Thus, an open set is one defined by the negation of some partial type. There is little importance whether we require the parameters to be taken in X or not: if $p(x, A)$ is any partial type and $q(y) = \text{tp}(A/X)$ then $p'(x) = \exists y q(y) \wedge p(x, y)$ defines the same set in X and only contains parameters in X .

For a distance $\varepsilon > 0$ and an element a we define the *closed ball of radius ε around a* in X as $B^X(a, \varepsilon) = \{b \in X : d(a, b) \leq \varepsilon\}$. (If one is bothered by balls of “asymmetric radius”, one can restrict the definition to symmetric ε , or define $B^X(a, \varepsilon) = \{b \in X : d(a, b) \leq \overleftrightarrow{\varepsilon}\}$, bearing in mind that $\varepsilon > 0 \implies \overleftrightarrow{\varepsilon} > 0$. This will make no difference.) It turns out that balls of positive radius suffice in order to define the logic topology:

Proposition 2.10. *A subset $U \subseteq X$ is open if and only if for every $a \in U$ there is $\varepsilon > 0$ such that $B^X(a, \varepsilon) \subseteq U$.*

Proof. For right to left, assume that for every $a \in U$ there is $\varepsilon_a > 0$ such that $B^X(a, \varepsilon_a) \subseteq U$, and let φ_a be such that $\varepsilon_a >^{\varphi_a} 0$. Then $X \setminus U$ is closed in the logic topology, as it defined in X by the partial type $\bigwedge_{a \in U} \varphi_a(x, a)$.

For left to right, let $U \subseteq X$ be an open set, so its complement is defined by a partial type $p(x, A)$, and let $a \in U \implies a \not\models p$. Set $q(y, Z) = \text{tp}(a, A)$ and $r(x, y) = \exists Z p(x, Z) \wedge$

$q(y, Z)$. Then $r(x, y)$ contradicts $x = y$, so by Lemma 2.3 there exist $\varepsilon > 0$ contradicting r . On the other hand $\neg r(x, a)$ defines a subset of U (since $\models q(a, A)$), and we obtain $B^X(a, \varepsilon) \subseteq U$. QED_{2.10}

It follows that, just as in the classical metric space setting, the family $\{B(a, \varepsilon) : \varepsilon > 0\}$ forms a basis to the neighbourhoods of a , and the logic topology is also a *distance topology*.

It should be remarked that in a first order theory, or more generally in a cat with positive inequality (as defined in Section 1), this topology is simply the discrete topology, which is why we also call such cats discrete. On the other hand, in natural analytic examples which come with an interesting metric topology, such as Banach space structures and probability measure algebras, this extra-logical topology turns out to coincide with the logical one.

We defined the logic topology as a relative topology, by saying when a set is closed or open in a superset. There is also an absolute notion of a set being closed, namely *completeness* (there does not seem to be an absolute notion of openness).

Recall that a *net* is something of the form $(a_i : i \in I)$ where I is a directed partially ordered set, that a net in a topological space X converges to a point a if for every neighbourhood $U \in \mathcal{N}_a$ there is $i \in I$ such that $j \geq i \implies a_j \in U$, and that if $A \subseteq X$ then its closure in X , which will be denoted by \bar{A}^X , is precisely the set of limits in X of nets in A . If X satisfies the first countability axiom (every point has a countable basis for its neighbourhoods) then we can replace “net” with “sequence”.

Definition 2.11. A *Cauchy net* is a net $(a_i : i \in I)$ in a single sort such that for every $\varepsilon > 0$ there is $i_\varepsilon \in I$ such that $j, j' \geq i_\varepsilon \implies d(a_j, a_{j'}) < \varepsilon$.

Definition 2.12. A set A in a single sort is *complete* if every Cauchy net in A converges to a point in A .

Proposition 2.13. *A set A of elements in a sort S is complete if and only if it is closed in $\text{dcl}(A) \cap S$. It is then closed in every superset.*

Proof. Assume that A is complete and $A \subseteq B$. If $a \in \bar{A}^B$, then a is the limit of a net $(a_i : i \in I)$. As every convergent net is Cauchy, the net (a_i) converges to a unique limit in A . It follows that $a \in A$.

Conversely, let $B = \text{dcl}(A) \cap S$, and assume that A is closed in B . Let $(a_i : i \in I)$ be a Cauchy net in A . For every $\varepsilon > 0$ let $i_\varepsilon \in I$ be as in Definition 2.11. Then the partial type $\bigwedge_{\varepsilon > 0} \varepsilon(x, a_{i_\varepsilon})$ is finitely consistent: for $\varepsilon_0, \dots, \varepsilon_{k-1}$ find $j \geq i_{\varepsilon_0}, \dots, i_{\varepsilon_{k-1}}$ (as I is directed) and then a_j realises it. By compactness we find a such that $\bigwedge_{\varepsilon > 0} \varepsilon(a, a_{i_\varepsilon})$. Then $a \in B$, and $a_i \rightarrow a$, so by assumption $a \in A$. QED_{2.13}

In particular, every sort of the universal domain is complete.

2.3. Metric cats.

- Definition 2.14.**
- (i) Given a sort x (more precisely, x is a variable which designates its associated sort), we define $\mathcal{E}^{>0}(x) = \{\varepsilon(x, x') : \varepsilon > 0\}$, i.e., the set of all positive distances in x .
 - (ii) A *basis of positive distances* in x is a family $\mathcal{E} \subseteq \mathcal{E}^{>0}(x)$ which is co-final in the sense that for every $\varepsilon > 0$ there is $\varepsilon' \in \mathcal{E}$ such that $\varepsilon > \varepsilon'$.
 - (iii) $\text{cfdist}(x)$ is the minimal cardinality of a basis of positive distances in x (in other words, $\text{cfdist}(x) = \text{cf}(\mathcal{E}^{>0}(x), \geq)$). If $\text{cfdist}(x) \leq \omega$, we say that the sort x is *metric*.
 - (iv) We say that T is metric if its home sort is (if T has several home sorts, then it is metric if all are).

Most of the time the sort in question is fixed by the context, so we may omit it. In the special case that x is in the home sort, we write $\text{cfdist}(T)$ instead of $\text{cfdist}(x)$ (so T is metric if $\text{cfdist}(T) \leq \omega$).

Clearly, if T is metric, then the sorts of finite and countable tuples in the home sort are also metric, as are quotients thereof by equivalence relations definable using countably many formulas.

Lemma 2.15. *Let $\mathcal{E} \subseteq \mathcal{E}^{>0}$ be such that $\bigwedge \mathcal{E} = \bigwedge_{\varepsilon \in \mathcal{E}} \varepsilon = 0$, and assume in addition that \mathcal{E} is closed under finite conjunctions. Then \mathcal{E} is a basis of positive distances.*

Proof. Assume that $\varepsilon >^\varphi 0$ for some formula φ . Since $\bigwedge \mathcal{E} = 0$, and $\varphi(x, y)$ contradicts $x = y$, there is a finite subset $\{\varepsilon_0, \dots, \varepsilon_{n-1}\} \subseteq \mathcal{E}$ such that φ contradicts $\bigwedge_{i < n} \varepsilon_i = \varepsilon' \in \mathcal{E}$. Therefore $\varepsilon \geq \varepsilon'$ (and in fact $\varepsilon >^\varphi \varepsilon'$). QED_{2.15}

Corollary 2.16. *$\text{cfdist}(x)$ is the minimal cardinality of a set $\mathcal{E} \subseteq \mathcal{E}^{<0}(x)$ such that $\bigwedge \mathcal{E} = 0$.*

Proof. The conjunction of a basis of positive distances is zero, which gives one inequality. For the other, if $\bigwedge \mathcal{E} = 0$ and $|\mathcal{E}|$ is minimal, then either $\mathcal{E} = \{0\}$ or it is infinite: in either case, closing it under finite intersection should not change its cardinality, and then conclude by Lemma 2.15. QED_{2.16}

Finally, note that if $\varepsilon > 0$, then any formula $\varphi(x, y) \in \varepsilon(x, y)$ can also be viewed as a distance, and as such $\varphi > 0$. Thus, if \mathcal{E} is the set of all $\varepsilon > 0$ on the home sort which can be defined by a single formula then $\bigwedge \mathcal{E} = 0$ and $|\mathcal{E}| \leq |T|$. It follows that $\text{cfdist}(T) \leq |T|$, and if T is countable then it is metric.

If T is metric we can find a basis of positive distances which can be enumerated as a decreasing sequence $(\varepsilon_n : n < \omega)$. We may further assume by Corollary 2.8 that the decreasingly enumerated basis satisfies in addition $\varepsilon_i > 2\varepsilon_{i+1}$ for every i , which may be quite useful (thinking of ε_i as the distance 2^{-i}). However, we can do better than this, and construct an actual definable metric on the home sort:

Definition 2.17. Let U_x denote the set of elements (of the universal domain) in the sort of x . A mapping $f : U_x \rightarrow \mathbb{R}$ is *definable* if for every r the properties $f(x) \geq r$ and

$f(x) \leq r$ are type-definable.

This is the same as saying that the mapping $S_x(T) \rightarrow \mathbb{R}$ sending $\text{tp}(a) \mapsto f(a)$ is continuous (and therefore implies that the range of f is compact).

In particular, a metric on U_x is definable if it is definable as a mapping from U_x^2 to \mathbb{R}^+ .

Assume now that T is metric, or more generally that we work in a fixed metric sort. Let $(\varepsilon_{\frac{1}{n}} : n < \omega)$ be distances on this sort such that $\varepsilon_{\frac{1}{n}} > \varepsilon_{\frac{1}{n+1}}$ and $\bigwedge \varepsilon_{\frac{1}{n}} = 0$. As for every $\varepsilon > \varepsilon'$ we can find $\varepsilon > \varepsilon'' > \varepsilon'$, we may extend our sequence to $(\varepsilon_q : q \in \mathbb{Q} \cap [0, 1])$ such that $q > r \implies \varepsilon_q > \varepsilon_r$.

Define $h(a, b) = \inf\{q : d(a, b) \leq \varepsilon_q\}$. Then $h(a, b) = \sup\{q : d(a, b) \geq \varepsilon_q\}$ as well (we convene that $\inf \emptyset = 1$ and $\sup \emptyset = 0$), and $h(x, y)$ is a definable function:

$$[h(x, y) \leq r] \equiv \bigwedge_{q > r} [d(x, y) \leq \varepsilon_q] \quad [h(x, y) \geq r] \equiv \bigwedge_{q < r} [d(x, y) \geq \varepsilon_q].$$

Since we may (and do) assume that each ε_q is symmetric, so is the function $h(x, y)$, and $h(a, b) = 0 \iff a = b$. Unfortunately, there is no reason that h should satisfy the triangle inequality.

We could define $d(a, b) = \inf_c |h(a, c) - h(c, b)|$: this would be a metric and induce the logic topology, but it is not necessarily definable (it is definable if we assume that T is open).

Instead, we can obtain a definable metric by re-scaling h . For this we need the following technical result:

Notation 2.18. Let D denote the set of dyadic numbers in $[0, 1]$.

Lemma 2.19. *Let $g : [0, 1]^2 \rightarrow [0, 1]$ be symmetric, non-decreasing, and satisfy for all $u, w, t \in [0, 1]$: $g(0, t) = t$ and if $g(u, w) < t$ then there is $u < v \leq 1$ such that $g(v, w) < t$.*

Then there is a function $f : D \rightarrow [0, 1]$ such that:

- (i) *f is strictly increasing.*
- (ii) *For all n : $f(\frac{1}{2^n}) \leq \frac{1}{2^n}$.*
- (iii) *For every $t, u \in D \cap (0, 1]$:*

$$t + u \leq 1 \implies g(f(t), f(u)) < f(t + u)$$

Proof. We will define f by steps, at the n th step extending its domain to $\{\frac{k}{2^n} : 0 \leq k \leq 2^n\}$ and verifying that the properties above hold for the part of f already defined.

For $n = 0$ we start with $f(0) = 0$, $f(1) = 1$, noting in particular that the third requirement holds vacuously.

For $n \geq 1$, we only need to choose $f(\frac{k}{2^n})$ for odd $0 < k < 2^n$. For $k \geq 3$ we look for $f(\frac{k-1}{2^n}) < s < f(\frac{k+1}{2^n})$ such that for all $0 < k' \leq \frac{2^n - k + 1}{2}$:

$$(*) \quad g(s, f(\frac{k'}{2^{n-1}})) < f(\frac{k+2k'-1}{2^n})$$

First, we note that $s = f(\frac{k-1}{2^n})$ satisfies $(*)$ for all $0 < k' \leq \frac{2^n - k + 1}{2}$: by the induction hypothesis, the third property of f holds for $t = \frac{k-1}{2^n}$ and $u = \frac{k'}{2^{n-1}}$ which are positive of

sum ≤ 1 . By the assumption on g , for each $0 < k' \leq \frac{2^n - k + 1}{2}$ there is $s_{k'} > f(\frac{k-1}{2^n})$ which satisfies (*) for that value of k' . Let $s' = \min\{s_{k'} : 0 < k' \leq \frac{2^n - k + 1}{2}\}$. Then $s' > f(\frac{k-1}{2^n})$, and we may choose $f(\frac{k}{2^n})$ such that $f(\frac{k-1}{2^n}) < f(\frac{k}{2^n}) < \min\{s', f(\frac{k+1}{2^n})\}$. In particular:

$$(**) \quad g(f(\frac{k}{2^n}), f(\frac{k'}{2^{n-1}})) < f(\frac{k+2k'-1}{2^n}) \quad \text{for all } 0 < k' \leq \frac{2^n - k + 1}{2}$$

Having chosen $f(\frac{k}{2^n})$ for all $3 \leq k < 2^n$ we choose $f(\frac{1}{2^n})$. By the induction hypothesis we have

$$0 < f(\frac{1}{2^{n-1}}) \leq \frac{1}{2^{n-1}}.$$

By the assumption on g

$$g(0, \frac{1}{2}f(\frac{1}{2^{n-1}})) = \frac{1}{2}f(\frac{1}{2^{n-1}}) < f(\frac{1}{2^{n-1}}),$$

and therefore there is $0 < s_0 \leq 1$ such that

$$g(s_0, \frac{1}{2}f(\frac{1}{2^{n-1}})) < f(\frac{1}{2^{n-1}}).$$

Similarly, for $0 < k' < 2^{n-1}$:

$$g(0, f(\frac{k'}{2^{n-1}})) = f(\frac{k'}{2^{n-1}}) < f(\frac{2k'+1}{2^n}),$$

so there is $s_{k'} > 0$ such that:

$$g(s_{k'}, f(\frac{k'}{2^{n-1}})) < f(\frac{2k'+1}{2^n}).$$

Defining $f(\frac{1}{2^n}) = \min\{\frac{1}{2}f(\frac{1}{2^{n-1}}), s_k : 0 \leq k < 2^{n-1}\}$, we obtain:

- (i) $0 < f(\frac{1}{2^n}) < f(\frac{1}{2^{n-1}})$.
- (ii) $f(\frac{1}{2^n}) \leq \frac{1}{2^n}$.
- (iii) $g(f(\frac{1}{2^n}), f(\frac{1}{2^n})) < f(\frac{1}{2^{n-1}})$.
- (iv) $g(f(\frac{1}{2^n}), f(\frac{k'}{2^{n-1}})) < f(\frac{2k'+1}{2^n})$ for $0 < k' < 2^{n-1}$.

Let us now verify the properties of f for the new values. Strict monotonicity follows directly from the construction, and the second property was taken care of. We still have to show that if $k, k' > 0$ and $k + k' \leq 2^n$ then:

$$(***) \quad g(f(\frac{k}{2^n}), f(\frac{k'}{2^n})) < f(\frac{k+k'}{2^n}).$$

- If both k and k' are even: Then (***) holds by the induction hypothesis on n .
- If one is odd and the other even: As g is symmetric we may assume that k is odd and k' is even. If $k = 1$, we took care of (***) when we chose $f(\frac{1}{2^n})$. If $k \geq 3$, then we have from (**):

$$g(f(\frac{k}{2^n}), f(\frac{k'}{2^n})) < f(\frac{k+k'-1}{2^n}) < f(\frac{k+k'}{2^n})$$

- If both k and k' are odd: If $k = k' = 1$, we took care of (***) when we chose $f(\frac{1}{2^n})$. Otherwise, by symmetry of g we may assume that $k \geq 3$ and we have from (**):

$$g(f(\frac{k}{2^n}), f(\frac{k'}{2^n})) \leq g(f(\frac{k}{2^n}), f(\frac{k'+1}{2^n})) < f(\frac{k+k'}{2^n})$$

Thus everything is fine and the induction may proceed.

QED_{2.19}

We now define $g: [0, 1]^2 \rightarrow [0, 1]$ by:

$$g(t, u) = \sup\{h(a, b) : \exists c h(a, c) \leq t \wedge h(c, b) \leq u\}.$$

Claim. This g satisfies the assumptions of Lemma 2.19.

Proof of claim. Clearly, g is symmetric and non-decreasing, and $g(0, t) = t$. Also, $g(0, t) = t$. The first statement is clear. As for the second, if $g(u, w) < t$, then the following partial type is inconsistent: $[h(x, y) \geq t] \wedge [h(x, z) \leq u] \wedge [h(z, y) \leq w]$. This is in turn equivalent to $\bigwedge_{t' < t} [h(x, y) \geq t'] \wedge \bigwedge_{v > u} [h(x, z) \leq v] \wedge [h(z, y) \leq w]$, so by compactness there are $t' < t$ and $v > u$ such that $[h(x, y) \geq t'] \wedge [h(x, z) \leq v] \wedge [h(z, y) \leq w]$ is inconsistent. It follows that $g(v, w) \leq t' < t$. QED_{Claim}

It follows that there is a function f as in the conclusion of Lemma 2.19. In particular, f is strictly increasing, so we may define:

$$d(a, b) = \inf\{t : h(a, b) < f(t)\} = \sup\{t : h(a, b) > f(t)\}.$$

Observe that $d(a, b) \geq r$ if and only if: $\forall (t \in D)(t < r \rightarrow h(a, b) \geq f(t))$, whereby:

$$[d(x, y) \geq r] \equiv \bigwedge_{t < r} [h(x, y) \geq f(t)]$$

And similarly:

$$[d(x, y) \leq r] \equiv \bigwedge_{t > r} [h(x, y) \leq f(t)]$$

Thus d is definable. Also, as h is symmetric so is d , and $d(a, b) = 0 \iff a = b$ (since $\lim_{n \rightarrow \infty} f(\frac{1}{2^n}) = 0$). Finally, for all a, b, c :

$$h(a, b) \leq g(h(a, c), h(c, b))$$

Whereby:

$$\begin{aligned} d(a, b) &= \inf\{t : h(a, b) < f(t)\} \\ &\leq \inf\{t : g(h(a, c), h(c, b)) < f(t)\} \\ &\leq \inf\{t : \exists u, w g(f(u), f(w)) < f(t), h(a, c) < f(u), h(c, b) < f(w)\} \\ &\leq \inf\{t : \exists u, w f(u + w) \leq f(t), h(a, c) < f(u), h(c, b) < f(w)\} \\ &= \inf\{u + w : h(a, c) < f(u), h(c, b) < f(w)\} \\ &= d(a, c) + d(c, b) \end{aligned}$$

We have therefore proved:

Theorem 2.20. *A sort is metric if and only if it admits a definable metric, in which case the metric topology coincides with the logic topology on that sort.*

Proof. We have just shown that every metric sort admits a definable metric. Conversely, if d is a definable metric, then $([d(x, y) \leq \frac{1}{n}] : n < \omega)$ is a countable basis of positive distances on the sort in question. This also implies that the metric topology is the same as the one induced by all the positive distances. QED_{2.20}

Most of the cats we'd be interested in are metric. First order theories and their likes are even discrete, and analytic cats admit natural definable metrics, be it the norm metric in Banach space structures, or $\mu(a \oplus b)$ in the case of probability algebras.

An additional very instructive example, albeit somewhat artificial, is the cat of ω -tuples of a (say) first order theory T , denoted by $T_{\times\omega}$. The natural definable metric (or ultrametric) would be something like $d(a_{<\omega}, b_{<\omega}) = 2^{-\sup\{n : a_{<n} = b_{<n}\}}$.

If we replace ω by any greater cardinal, the resulting cat is clearly not metric, even though it is constructed from metric building blocks. We prove that this is essentially the only example of a non-metric cat, and the metric building blocks can be reconstructed:

Lemma 2.21. *Assume that $(\varepsilon_n : n < \omega)$ are symmetric distances such that $\varepsilon_n > 2\varepsilon_{n+1}$ for all $n < \omega$ (or even just $\varepsilon_n > \varepsilon_{n+1}$ and $\varepsilon_n \geq 2\varepsilon_{n+1}$). Then $E = \bigwedge_n \varepsilon_n$ is an equivalence relation, and x/E is a metric sort.*

Proof. E is symmetric and reflexive since each ε_n is, and the assumption that $\varepsilon_n \geq 2\varepsilon_{n+1}$ assures us that it is transitive as well, whereby an equivalence relation. For every $n < \omega$, the partial type $\bar{\varepsilon}_n(x/E, y/E)$ (which is a conventional notation for $\exists x'y' E(x', x) \wedge E(y', y) \wedge \varepsilon_n(x', y')$) is a distance in the sort x/E , and $\bigwedge \bar{\varepsilon}_n \equiv [x/E = y/E]$. It would therefore suffice to show that each $\bar{\varepsilon}_n$ is positive. Recall that by assumption $\varepsilon_n > \varepsilon_{n+1}$, and let $\varphi_n(x, y)$ be such that $\varepsilon_n >^{\varphi_n} \varepsilon_{n+1}$. Then $\varphi_n(x, y) \wedge \varepsilon_{n+1}(x, y)$ is contradictory, whereby $\varphi_n(x, y) \wedge E(x, y)$ and $\bar{\varphi}_n(x/E, y/E) \wedge x/E = y/E$ are contradictory as well. It follows that $\bar{\varepsilon}_n >^{\bar{\varphi}_n} 0$, as required. QED_{2.21}

Theorem 2.22. *Let x be in a non-metric sort. Then there are equivalence relations $\{E_i(x, y) : i < \text{cfdist}(T, x)\}$ such that $\bigwedge_i E_i(x, y) \equiv [x = y]$, and for each i the sort x/E_i is metric.*

Proof. This follows immediately from Lemma 2.21. QED_{2.22}

Thus, if the home sort is not metric, we may always decompose it into metric sorts, designating those as the new home sorts. We may therefore assume:

Convention 2.23. Hereafter, all the cats we consider are metric, i.e., have metric home sort(s).

3. MODELS

3.1. Pre-models. One delicate issue when working with cats is the question of which subsets of the universal domain are considered as models: just taking existentially closed subsets as was done in [Ben03a] is not a satisfactory solution, as this is very language-dependent. Alternatively, we recall that in a first order theory one can characterise an

elementary sub-model of the universal domain as a subset M such that the set of types over M realised in M (from now on we just call them *realised types*) is dense in $S(M)$. We can therefore define:

Definition 3.1. A *pre-model* of T is a subset M of a universal domain of T (or of any e.c. model of T) such that for every $n < \omega$, the realised n -types are dense in $S_n(M)$.

In other words, M is a pre-model if for every formula $\varphi(\bar{x}, \bar{m})$, where $\bar{m} \in M$, if there is \bar{a} (in the universal domain) such that $\neg\varphi(\bar{a}, \bar{m})$ then there is such \bar{a} in M .

Since pre-models have a purely topological characterisation, this is a language-independent notion. It is more general than the (language-dependent) notion of an e.c. model:

Lemma 3.2. *If M is an e.c. model in some language, then it is a pre-model.*

Proof. Assume that $m \in M$ and $\models \neg\varphi(a, m)$. Then there is ψ contradicting φ such that $\models \psi(a, m)$, and since M is e.c. there is $a' \in M$ such that $\models \psi(a', m)$, and in particular $\models \neg\varphi(a', m)$. QED_{3.2}

Moreover, if we have a language with negation (such as in first-order logic) then the notions of e.c. model and pre-model agree.

Unfortunately, a pre-model is far from being adequate to play the role of a model: in fact, it needs not even be definably closed (for example, any dense subset of an infinite dimensional Hilbert space is a pre-model for the theory of Hilbert spaces). For the time being, however, this notion will do.

The reader might be suspicious about the requirement being for every $n < \omega$, rather than just for $n = 1$ as in first order logic. This is inevitable if we want our scope to include non-open cats (such as Robinson theories which have no first order model companion). For example, one may find then a formula $\varphi(x_{<n})$ (with hidden parameters in M) such that $\neg\varphi(x_{<n})$ defines a non-empty (open) set in $S_n(M)$, but its projection to $S_1(M)$ on any coordinate $i < n$, defined by $\exists x_0, \dots, \hat{x}_i, \dots, x_{n-1} \neg\varphi(x_{<n})$, has empty interior. The following notion is meant to handle this difficulty:

Definition 3.3. Let A be any set of parameters and x (a variable designating) a sort. A *basic Q-open set* (for quantifier-open) in $S_x(A)$ is one defined by a formula $\exists y \neg\varphi(x, y, A)$, where y is any tuple of variables.

The family of basic q-open sets in $S_x(A)$ forms a basis for a topology which we call the *Q-topology*.

In other words, a basic Q-open set is the projection to $S_\alpha(A)$ of a basic open set of $S_{\alpha+\beta}(A)$, for some β , and we may always assume that $\beta < \omega$. It follows that the Q-topology on $S_\alpha(A)$ is a refinement of the ordinary topology, and it is the minimal topology such that the projection $S_{\alpha+\omega}(A) \rightarrow S_\alpha(A)$ is open when $S_{\alpha+\omega}(A)$ is taken with the ordinary topology. In fact, whenever $\alpha \leq \beta$ the projection map $S_\beta(A) \rightarrow S_\alpha(A)$ is open when both spaces are taken in Q-topology.

The Q-topology coincides with the ordinary one if and only if T is open. If not, then it is strictly stronger than a Hausdorff topology, and therefore cannot be compact. Nevertheless, we shall see that it is a quite useful notion. For example:

Definition 3.4. A set $X \subseteq S_\alpha(A)$ is *Q-dense* if it is dense in the Q-topology.

Proposition 3.5. *A set M is a pre-model if and only if the realised types are Q-dense in $S_1(M)$.*

Proof. For left to right, we assume that M is a pre-model, and prove that every non-empty basic Q-open set in $S_1(M)$ contains a realised type. Indeed, such a set is of the form $\exists \bar{y} \neg\varphi(x, \bar{y})$ with parameters in M , and $\neg\varphi(x, \bar{y})$ has a realisation somewhere. Since M is a pre-model, it has a realisation $a, \bar{b} \in M$, so in particular a is a realisation of $\exists \bar{y} \neg\varphi(x, \bar{y})$ in M .

For the converse, we assume that the realised types are Q-dense in $S_1(M)$, and prove by induction on $n < \omega$ that if $\neg\varphi(x_{<n})$ is consistent with parameters in M then it is realised in M . For $n = 0$ there is nothing to prove, so we assume for n and prove for $n + 1$. If $\neg\varphi(x_{<n+1})$ is consistent with parameters in M then $\exists x_{<n} \neg\varphi(x_{<n+1})$ defines a non-empty Q-open set in $S_1(M)$. By assumption it is realised by some $a_n \in M$. Then $\neg\varphi(x_{<n}, a_n)$ is consistent with parameters in M , and we conclude using the induction hypothesis. QED_{3.5}

3.2. Co-heirs.

Definition 3.6. Let $M \subseteq A$, and $p \in S(A)$. Then p *co-inherits* from M , or is a *co-heir* of its restriction to M , if it is in the closure in $S(A)$ of types realised in M . In other words, if for every formula $\varphi(x, a)$ over A , if $p \vdash \neg\varphi(x, a)$ then there is $m \in M$ such that $\models \neg\varphi(m, a)$.

This definition will mostly be used when M is a pre-model.

Thus by definition, a set M is a pre-model if and only if every type $p \in S(M)$ (in any tuple of variables) co-inherits from M .

Lemma 3.7. *Assume that $M \subseteq A \subseteq B$, and $p \in S(A)$ co-inherits from M . Then p has an extension to B which co-inherits from M .*

Proof. Since p co-inherits from M , there is a net $(a_\sigma: \sigma \in \Sigma)$ in M such that $\text{tp}(a_\sigma/A) \rightarrow p$. Since $S(B)$ is compact, the net $(\text{tp}(a_\sigma/B): \sigma \in \Sigma)$ has a converging sub-net $(\text{tp}(a_{\sigma'}/B): \sigma' \in \Sigma')$, say with limit $q \in S(B)$. Then q co-inherits from M . On the other hand $\text{tp}(a_{\sigma'}/A) \rightarrow p$, and since in a Hausdorff space limits are unique q is an extension of p . QED_{3.7}

Definition 3.8. Let M be a pre-model, $M \subseteq A$, and c a tuple possibly outside A . Then $\text{tp}(c/A)$ *splits* over M if there are tuples $a, b \in A$ such that $a \equiv_M b$ but $a \not\equiv_{cM} b$.

Lemma 3.9. *A co-heir is a non-splitting extension.*

Proof. Assume that $p \in S(N)$ does indeed split over M . Then there are $a, b \in N$ such that $a \equiv_M b$, and contradicting formulas φ, ψ , such that $p \vdash \varphi(x, a) \wedge \psi(x, b)$. As T is Hausdorff there are ρ, χ such that $\varphi \vdash \neg\rho$, $\psi \vdash \neg\chi$ and $\models \rho \vee \chi$. Therefore $p \vdash \neg(\rho(x, a) \vee \chi(x, b))$, and as it is a co-heir there is $c \in M$ such that $\models \neg(\rho(c, a) \vee \chi(c, b))$. As $a \equiv_M b$ we get $\neg\rho(c, a) \wedge \neg\chi(c, a)$, contradicting $\models \rho \vee \chi$. QED_{3.9}

Lemma 3.10. *A co-heir is a non-dividing extension.*

Proof. Let $M \subseteq A$ where M is a pre-model, and let $p \in S(A)$ co-inherit from M . Let $A \subseteq N$ where N is $|M|^+$ -saturated, and let $p' \in S(N)$ extend p co-inheriting from M . As it is a non-splitting extension, it is easy to see that p' does not divide over M , so neither does p . QED_{3.10}

Lemma 3.11. *Every type over a pre-model M is Lascar strong.*

Proof. Exactly as in the proof of [KP97, Proposition 5.4], the existence of non-splitting extensions of types over M implies that if $a \equiv_M b$ then there are $\{c_i: i < n\}$ such that a, c_0, c_1, \dots and b, c_0, c_1, \dots are both M -indiscernible. It follows that a and b have the same Lascar strong type (as c_0) over M . QED_{3.11}

These results accumulate to:

Proposition 3.12. *If T is stable, M is a pre-model, and $p \in S(M)$, then p has a unique non-dividing extension to any set, which is its unique co-heir over that set.*

3.3. Type omission. This is yet another variant of the classical omitting type theorem. We adapt the proof of [Poi85, Chapter 10.a], pointing out at the same time the usefulness of the Q-topology (in non-open cats).

We recall that a *nowhere dense* set is one whose closure has empty interior, and a *meagre* set is a countable union of nowhere dense sets. In the first order context we know that these notions are intimately related to type omission. As above, when dealing with a cat that is not open, the Q-topology yields a more adequate notion; we call the corresponding notions *nowhere Q-dense* and *Q-meagre*, respectively.

Lemma 3.13. *A set $X \subseteq S_\alpha(A)$ is nowhere Q-dense if and only if for some (every) $\beta \geq \omega$, the pull-back of X to $S_{\alpha+\beta}(A)$ is nowhere dense (in the ordinary topology).*

Proof. Fix some $\beta \geq \omega$, and let $Y \subseteq S_{\alpha+\beta}(A)$ be the pull-back of X . Let \bar{X} and \bar{Y} denote the closure of X in the Q-topology and the closure of Y in the ordinary topology, respectively. We need to prove that \bar{X} has empty Q-interior if and only if \bar{Y} has empty ordinary interior.

For this, observe at first that \bar{X} is defined by $\bigwedge \{\forall w \varphi(x, w): X \vdash \forall w \varphi(x, w)\}$, where x is a tuple of α variables and w can be any tuple, whereas \bar{Y} is defined by $\bigwedge \{\varphi(x, y): X \vdash \forall y \varphi(x, y)\}$ where y is a fixed β -tuple (all formulas here may have hidden parameters in A).

Assume now that $\neg\psi(x, z)$ is consistent, and we may assume that z is finite. Since $\beta \geq \omega$

we may assume that $z \subseteq y$, and even then there are still infinitely many variables in $y \setminus z$. Therefore $\neg\psi(x, z) \vdash \bar{Y}$ if and only if $\neg\psi(x, z) \vdash \forall w \varphi(x, w)$ for every $\varphi(x, w)$ such that $X \vdash \forall w \varphi(x, w)$, if and only if $\exists z \neg\psi(x, z) \vdash \bar{X}$, as required. QED_{3.13}

Notation 3.14. Assume that $\varphi(x, y)$ and $\psi(x, y)$ are contradictory, and define $\pi_{\varphi, \psi}(x, y_{<\omega}) = \exists y \varphi(x, y) \wedge \bigwedge_{i < \omega} \psi(x, y_i)$.

Lemma 3.15. *The partial type $\pi_{\varphi, \psi}$ defines a nowhere Q-dense set (and since it is closed, it is in particular nowhere dense).*

Proof. It is closed, so we only need to verify that its Q-interior is empty. Assume the contrary, namely that there is a formula $\chi(x, y_{<\omega}, z)$ such that $\exists z \neg\chi(x, y_{<\omega}, z) \vdash \pi_{\varphi, \psi}(x, y_{<\omega})$, and $\neg\chi$ is consistent, say $\models \neg\chi(a, b_{<\omega}, c)$. Then $\models \pi_{\varphi, \psi}(a, b_{<\omega})$, so there is b' such that $\models \varphi(a, b')$, and therefore $\neg\psi(a, b')$. Since χ has only finitely many variables actually appearing in it there is $i < \omega$ such that y_i does not appear in χ , so $\models \neg\chi(a, b_0, \dots, b_{i-1}, b', b_{i+1}, \dots, c)$ as well, contradicting $\neg\chi \vdash \pi_{\varphi, \psi}$. QED_{3.15}

Lemma 3.16. *Assume that A is countable and not a pre-model. Then there is a pair of contradicting formulas $\varphi(x, y)$, $\psi(x, y)$ and a tuple $a \in A$ in the length of x , such that if $(b_i : i < \omega)$ is an enumeration of all the tuples in A of the length of y , then $\models \pi_{\varphi, \psi}(a, b_{<\omega})$.*

Proof. If A is not a pre-model, then there are $a \in A$ and a negative formula $\neg\psi(a, y)$ which has a realisation b , but has none inside A . There is a formula $\varphi(x, y) \in \text{tp}(a, b)$ contradicting ψ , and $\models \pi_{\varphi, \psi}(a, b_{<\omega})$. QED_{3.16}

It follows that if T has a countable language, then the set of types $p \in S_\omega(T)$ whose realisations are not enumerations of pre-models is meagre (it is also Q-meagre, but we do not need that). Now the standard Baire category argument yields:

Theorem 3.17. *Let T be a cat with a countable language, and $X_n \subseteq S_n(T)$ Q-meagre sets for every $n < \omega$. Then T has a pre-model which omits every X_n .*

Proof. The set of all types $p \in S_\omega(T)$ of enumerations of pre-models omitting every X_n cannot be empty, since its complement is meagre, and a compact Hausdorff space cannot be meagre. QED_{3.17}

3.4. Complete models. We know that if A is any set (in a single sort) then $\bar{A} \subseteq \text{dcl}(A)$. In the case of pre-models we have equality in the home sort:

Lemma 3.18. *Let M be a pre-model (so in particular, it is a set of real elements). Then $\bar{M} = \text{dcl}^h(M) = \text{bdd}^h(M)$, where the superscript h means that we restrict the definable or bounded closure to the home sort.*

Proof. We already know that $\bar{M} \subseteq \text{dcl}^h(M) \subseteq \text{bdd}^h(M)$. On the other hand, Lemma 3.11 implies that $\text{dcl}(M) = \text{bdd}(M)$, so in particular $\text{dcl}^h(M) = \text{bdd}^h(M)$. In order to conclude we need to show that $a \in \text{dcl}^h(M) \implies a \in \bar{M}$.

If $a \in \text{dcl}^h(M)$ then $p(x) = \text{tp}(a/M)$ defines a . Assume that $\varepsilon >^\varphi 0$. Then $q(x) =$

$\exists y p(y) \wedge \varphi(x, y)$ is a partial type, and $a \models \neg q(x)$ (since $\neg \varphi(a, a)$). Since M is a pre-model, there is $b \in M$ such that $b \models \neg q(x) \implies \models \neg \varphi(b, a) \implies a \in B(b, \varepsilon)$. As this holds for every $\varepsilon > 0$, we obtain $a \in \bar{M}$. QED_{3.18}

We define:

Definition 3.19. A *complete model* is the closure \bar{M} of a pre-model M .

As M and \bar{M} are interdefinable, if M is a pre-model then so is \bar{M} . Therefore, a complete model is precisely a closed pre-model. By Proposition 2.13, this is the same as a complete pre-model, where completeness is taken in the sense of convergence of Cauchy nets, whence the terminology.

We leave it to the reader to verify the “converse” to this definition, namely that any dense subset of a complete model is a pre-model.

Since complete models are defined as closed sets, their size is more accurately measured by their density character than by their cardinality:

Definition 3.20. If M is a complete model then its *density character*, denoted $\|M\|$, is the minimal cardinality of a dense subset of M .

Definition 3.21. A cat T is λ -*categorical* if it has a unique complete model (up to isomorphism) of density character λ .

3.5. Principal types. Here we need to introduce some new notation:

Notation 3.22. Given a partial type $p(x)$ (with parameters) and a distance ε , we define $p(x^\varepsilon) = \exists y [\varepsilon(x, y) \wedge p(y)]$ and $p(x^{<\varepsilon}) = \forall y [d(y, x) < \varepsilon \rightarrow p(y)]$. Mind that $p(x^\varepsilon)$ is a partial type, whereas $p(x^{<\varepsilon})$ only defines an Q-closed set.

We will ordinarily use the form $x^{<\varepsilon}$ inside negations, so one should remember that $\neg p(x^{<\varepsilon}) = \exists x' [d(x, x') < \varepsilon \wedge \neg p(x')]$, which is an Q-open set.

The idea behind these definitions is that an element can contain infinitely much information (think of the cat of ω -tuples of elements of a first order theory, which is metric), but when quantifying over a ball of positive radius around an element we only retain finitely much of the information contained in that element.

For example, in a first order theory we know that if $\pi(x, B)$ is a partial type and $a \models \neg \pi$, then there is $\psi(x, b) \in \pi$ such that $a \models \neg \psi$. In particular, the parameter b is a finite tuple. The analogue of this observation is the following:

Lemma 3.23. *Assume that $\neg \varphi(a, b)$. Then there exists $\varepsilon > 0$ on the sort of b and a formula $\psi(x, y)$ such that $\models \neg \psi(a, b^\varepsilon)$ and $\neg \psi(x, y^{<\varepsilon}) \vdash \neg \varphi(x, y)$.*

Proof. Let $p(x, y) = \text{tp}(a, b)$. Then $\bigwedge_{\varepsilon > 0} p(x, y^{2\varepsilon})$ contradicts $\varphi(x, y)$, whereby $p(x, y^{2\varepsilon})$ contradicts $\varphi(x, y)$ for some $\varepsilon > 0$, which we may assume symmetric. This is equivalent to $p(x, y^\varepsilon)$ contradicting $\varphi(x, y^\varepsilon)$. Let $\psi(x, y) = \varphi(x, y^\varepsilon)$ (or, if one insists on the syntactic

distinction between a formula and a partial type, find ψ contradicting $p(x, y^\varepsilon)$ such that $\varphi(x, y^\varepsilon) \vdash \psi(x, y)$. Then $\neg\psi(a, b^\varepsilon)$, and

$$\begin{aligned} \neg\psi(x, y^{<\varepsilon}) &= \exists y' [d(y, y') < \varepsilon \wedge \neg\psi(x, y')] \\ &\vdash \exists y' [d(y, y') < \varepsilon \wedge \neg\varphi(x, y'^\varepsilon)] \\ &= \exists y' [d(y, y') < \varepsilon \wedge \neg\exists y'' [d(y', y'') \leq \varepsilon \wedge \varphi(x, y'')]] \\ &\vdash \neg\varphi(x, y). \end{aligned} \quad \text{QED}_{3.23}$$

By Theorem 3.17 on one hand and Proposition 3.5 on the other, a partial type can be omitted in a pre-model if and only if the closed set it defines has empty Q-interior (assuming that we are dealing with a countable language). As we are more interested in complete models than in pre-models, we need an appropriately modified version, which is in fact an adaptation of Henson's notion of a principal type to our context:

Definition 3.24. A complete type $p \in S(A)$ is *principal* (over A) if for every $\varepsilon > 0$ the closed set defined by $p(x^\varepsilon)$ in $S(A)$ has non-empty Q-interior.

In the definition of a principal type we did not require that p actually belong to the Q-interior of $p(x^\varepsilon)$. However, if p is principal and $\varepsilon > 0$ is given, find $\varepsilon'' > \varepsilon' > 0$ such that $\varepsilon > 2\varepsilon''$: if $q(x)$ is in the Q-interior of $p(x^{\varepsilon'})$, then there is ψ such that $q(x) \vdash \exists y \neg\psi(x, y) \vdash p(x^{\varepsilon'})$. Therefore $p(x) \vdash q(x^{\varepsilon'}) \vdash \exists y \neg\psi(x^{<\varepsilon''}, y) \vdash p(x^{\varepsilon'+\varepsilon''}) \vdash p(x^\varepsilon)$, so p is in the Q-interior of $p(x^\varepsilon)$.

Proposition 3.25. Assume that M is a complete model, $A \subseteq M$, and $p(x) \in S(A)$ is a principal type in a metric sort. Then p is realised in M .

Proof. By assumption there is a definable metric d on the sort of x . We may assume that d takes only values in $[0, 1]$. We identify a real number $r > 0$ with the positive distance defined by $d(x, y) \leq r$. For every n , let $V_n \subseteq S(A)$ denote the Q-interior of $p(x^{\frac{1}{2^n}})$. Since p is principal, $p \in V_n$ for all n .

We will construct a sequence $(a_n : n < \omega)$ in M , such that $d(a_n, a_{n+1}) < \frac{1}{2^{n-1}}$, and $\text{tp}(a_n/A) \in V_n$. Since $d \leq 1$ $V_0 = S_\alpha(A)$ and we may take a_0 to be any element in M .

Assume that a_n is already chosen. Since $\text{tp}(a_n/A) \in V_n$ we can find $b \models p$ (possibly outside M) such that $d(a_n, b) \leq \frac{1}{2^n}$. Let $U \subseteq S(Aa_n)$ denote the pull-back of V_{n+1} there, and let $W \subseteq S(Aa_n)$ denote the set defined by $\{d(a_n, x) < \frac{1}{2^{n-1}}\}$. Then U and W are Q-open and open, respectively, so $U \cap W$ is Q-open, and it is non-empty since it contains $\text{tp}(b/Aa_n)$. Therefore there exists $a_{n+1} \in M$ such that $\text{tp}(a_{n+1}/A) \in V_{n+1}$ and $d(a_n, a_{n+1}) < \frac{1}{2^{n-1}}$, as required.

Once the construction is finished, we see that $d(a_n, a_m) < \min\{\frac{1}{2^{n-2}}, \frac{1}{2^{m-2}}\}$ follows by induction from $d(a_n, a_{n+1}) < \frac{1}{2^{n-1}}$. Thus (a_n) is a Cauchy sequence, and since M is complete it has a limit $a \in M$. But then $a \models \bigwedge_n p(x^{\frac{1}{2^n}}) \implies a \models p$, as required. $\text{QED}_{3.25}$

Corollary 3.26. *Assume that T has a countable language, $|A| \leq \omega$, and p is a partial type over A with at most ω variables. Then there exists a complete model $M \supseteq A$ omitting p if and only if p is not principal.*

Proof. Since T has a countable language and p is in countably many variables, the sort of p is metric. Thus, if p is not principal, there is a distance $\varepsilon > 0$ such that $p(x^\varepsilon)$ is nowhere Q-dense, and can therefore be omitted in a pre-model $M \supseteq A$. As every realisation of p has distance at least ε from M , p is also omitted from \bar{M} . On the other hand, if p is principal, it is realised in every such model by Proposition 3.25. QED_{3.26}

Proposition 3.27. *Assume that $(a_i : i < \alpha)$ is a sequence (in possibly different sorts) such that $\text{tp}(a_i/Aa_{<i})$ is principal for every $i < \alpha$. Then $\text{tp}(a_{<\alpha}/A)$ is principal.*

Proof. We would like to prove by induction on α . If we know this for α , then $\alpha + 1$ reduces to the case where $\alpha = 2$; and if α is limit, then by Lemma 2.5 the statement follows from the induction hypothesis.

We are left with the case $\alpha = 2$. We write $b = a_0$, $a = a_1$, and may assume that $A = \emptyset$. Write $p(y) = \text{tp}(b)$ and $q(x, y) = \text{tp}(a, b)$. We need to prove that if $\varepsilon_x, \varepsilon_y > 0$ then $q(x^{\varepsilon_x}, y^{\varepsilon_y})$ has non-empty Q-interior.

By assumption, $q(x, b)$ belongs to the Q-interior (in $S(b)$) of $q(x^{\varepsilon_x}, b)$. In other words, there is a formula $\varphi(x, y, z)$ such that $q(x, b) \vdash \exists z \neg \varphi(x, b, z) \vdash q(x^{\varepsilon_x}, b)$. Let c be such that $\neg \varphi(a, b, c)$. By compactness there is $\varepsilon'_y > 0$ such that $\neg \varphi(a, b^{\varepsilon'_y}, c)$. Since we lose no generality by decreasing ε_y , we may assume that $\varepsilon_y = \varepsilon'_y$.

Again by assumption there is $\chi(y, w)$ such that $p(y) \vdash \exists w \neg \chi(y, w) \vdash p(y^{\varepsilon_y})$. Then $q(x, y) \vdash \exists zw [\neg \varphi(x, y^{\varepsilon_y}, z) \wedge \neg \chi(y, w)]$, and:

$$\begin{aligned} \neg \varphi(x, y^{\varepsilon_y}, z) \wedge \neg \chi(y, w) &\vdash \neg \varphi(x, y^{\varepsilon_y}, z) \wedge \exists y' \varepsilon_y(y, y') \wedge p(y') \\ &\vdash \exists y' \varepsilon_y(y, y') \wedge p(y') \wedge \neg \varphi(x, y', z) \\ &\vdash \exists y' \varepsilon_y(y, y') \wedge q(x^{\varepsilon_x}, y') \\ &\equiv q(x^{\varepsilon_x}, y^{\varepsilon_y}). \end{aligned}$$

QED_{3.27}

Lemma 3.28. *Assume that $\text{tp}(a/A)$ is principal and $b \in \text{dcl}(aA)$. Then $\text{tp}(b/A)$ is principal.*

Proof. We may assume that $A = \emptyset$. Let $p(x) = \text{tp}(a)$, $q(y) = \text{tp}(b)$ and $r(x, y) = \text{tp}(a, b)$. Assume that $\varepsilon_y > 0$ on the sort of y . This means that $r(x, y) \wedge r(x, z) \wedge \varphi(y, z)$ is contradictory, whereby there are $\varepsilon_x > \varepsilon'_x > 0$ on the sort of x and a formula $\psi(x, y)$ contradicting $r(x, y)$ such that $r(x, y) \wedge \neg \psi(x^{\varepsilon'_x}, z) \wedge \varphi(y, z)$ is contradictory.

Since p is principal there is a formula $\chi(x, t)$ such that $p(x) \vdash \exists t \neg \chi(x, t) \vdash p(x^{\varepsilon'_x})$. Then clearly, $q(y) \vdash \exists xt \neg \chi(x, t) \wedge \neg \psi(x, y)$. On the other hand, assume that $\neg \chi(x, t) \wedge \neg \psi(x, y)$. Then $p(x^{\varepsilon'_x})$, which means there is x' such that $p(x') \wedge d(x, x') \leq \varepsilon'_x < \varepsilon_x$, whereby $\neg \psi(x'^{\varepsilon'_x}, y)$. On the other hand, since $p(x')$ there is y' such that $r(x', y')$. It follows that $\neg \varphi(y, y')$, so $q(y^{\varepsilon_y})$ as required. QED_{3.28}

3.6. Approximate saturation. In first order model theory we have the notions of a κ -saturated model and a κ -compact model, which coincide for $\kappa > |T|$ but may give a finer distinction for small values of κ (in particular, every model is ω -compact). On a finer inspection it turns out that when dealing with Hausdorff cats, the notion of κ -saturation has to be slightly adjusted: for $\kappa > |T|$ this changes nothing, and for $\kappa \leq |T|$ the modified notion is what we need in order to generalise classical theorems.

Definition 3.29. Let λ be an infinite cardinal. A complete model M is *approximately λ -saturated* if for every type $p(x, M) \in S(M)$, every $\varepsilon > 0$ on the sort of x , and every distance ε' on the sort of (an enumeration of) M which can be expressed as the conjunction of less than λ positive distances, there is $b \in M$ such that $\models p(b^\varepsilon, M^{\varepsilon'})$.

Proposition 3.30. For $\lambda > \omega$, approximate λ -saturation coincides with λ -saturation. (This is assuming, as we do, that the home sort is metric. Otherwise one needs to require $\lambda > \text{cfdist}(T)$.)

Proof. Let M be approximately λ -saturated. Let $A \subseteq M$, $|A| < \lambda$, and $p(x) \in S(A)$ is in the home sort. In fact, just in this proof, all variables will be in the home sort. Write $M = (m_i : i < \xi)$, and we may assume that $A = (m_i : i < \mu)$ for some $\mu < \lambda$. Let $\varepsilon'(x_{<\xi}, y_{<\xi})$ be $\bigwedge_{i < \mu} x_i = y_i$, which can be expressed as the conjunction of $\mu \times \omega < \lambda$ positive distances. By approximate λ -saturation, $p(x^\varepsilon)$ is realised in M for every $\varepsilon > 0$, and the same holds for every $q \in S(B)$ where $B \subseteq M$, $|B| < \lambda$.

Let $(\varepsilon_i : i < \omega)$ be a basis of positive distances. Construct a sequence $(b_i, p_i : i < \omega)$, where $p_i \in S(Ab_{<i})$ and $b_i \in M$ realises $p_i(x^{\varepsilon_i})$. For $i = 0$ we choose $p_0 = p$. For $i + 1$, we know that $p_i(x) \wedge \varepsilon_i(x, b_i)$ is consistent (since $b_i \models p_i(x^{\varepsilon_i})$) so p_{i+1} is any completion to $Ad_{<i}$. Once we choose p_i , we choose some $b_i \in M$ realising $p_i(x^{\varepsilon_i})$, which exists by the previous argument.

Let b realise $\bigwedge_{i < \omega} p_i$: then it realises p , and it is an accumulation point of $\{b_i : i < \text{cfdist}(T)\}$, so $b \in M$ as required. QED_{3.30}

3.7. Ehrenfeucht-Mostowski models.

Convention 3.31. Let $(I, <)$ be a totally ordered set. Although the following definitions could make sense for finite I , we usually assume that I is infinite. We denote a finite subset of I by $\mathbf{i} = \{i_0, \dots, i_{n-1}\} \in [I]^n$, and convene that this notation implies that $i_0 < \dots < i_{n-1}$.

Definition 3.32. A *generalised I -sequence* is something of the form $(a_{\mathbf{i}}^n : n < \omega, \mathbf{i} \in [I]^n)$, where each $a_{\mathbf{i}}^n$ is a possibly infinite tuple.

Notation 3.33. Given such a generalised sequence $(a_{\mathbf{i}}^n : n < \omega, \mathbf{i} \in [I]^n)$ we define by induction on n :

$$\hat{a}_{\{i_0, \dots, i_{n-1}\}}^n = \hat{a}_{\{i_1, \dots, i_{n-1}\}}^{n-1}, \dots, \hat{a}_{\{i_0, \dots, i_{j-1}, i_{j+1}, \dots, i_{n-1}\}}^{n-1}, \dots, \hat{a}_{\{i_0, \dots, i_{n-2}\}}^{n-1}, a_{\{i_0, \dots, i_{n-1}\}}^n$$

(so $\hat{a}_{\mathbf{i}}^n$ contains an awful lot of repetitions when n starts growing, but that's the way we like it.)

Definition 3.34. A generalised I -sequence $(a_{\mathbf{i}}^n : n < \omega, \mathbf{i} \in [I]^n)$ is *indiscernible* if $\text{tp}(\hat{a}_{\mathbf{i}}^n)$ depends only on n for every $n < \omega$.

Lemma 3.35. For every cardinal κ there is a cardinal λ (in fact, $\lambda = \beth_{(2^{|T|+\kappa})^+}$ will always do) such that if $(a_{\mathbf{i}}^n : n < \omega, \mathbf{i} \in [\lambda]^n)$ is a generalised λ -sequence and $|a_{\mathbf{i}}^n| \leq \kappa$ for every n, \mathbf{i} , then there is an indiscernible generalised sequence $(b_{\mathbf{i}}^n : n < \omega, \mathbf{i} \in [\omega]^n)$ such that for every $\mathbf{i} \in [\omega]^n$ there is $\mathbf{i}' \in [\lambda]^n$ such that $\hat{b}_{\mathbf{i}}^n \equiv \hat{a}_{\mathbf{i}'}^n$. Moreover, we may always take it so that $a_{\emptyset}^0 = b_{\emptyset}^0$.

Proof. Just like for ordinary sequences. QED_{3.35}

Definition 3.36. A *sequential pre-model* is a generalised I -sequence $(a_{\mathbf{i}}^n)$ such that $|a_{\mathbf{i}}^n| \leq |T|$ for every $\mathbf{i} \in [\lambda]^n$, and for every $\mathbf{j} \in [\mathbf{i}]^{n-1}$, every consistent negative formula with parameters in $\hat{a}_{\mathbf{j}}^{n-1}$ has a realisation in $a_{\mathbf{i}}^n$. (Such a sequence is in particular a pre-model.) An *Ehrenfeucht-Mostowski pre-model* is a sequential pre-model which is indiscernible as a generalised sequence.

Proposition 3.37. *Ehrenfeucht-Mostowski pre-models exist.*

Proof. Construct a very long sequential pre-model, by choosing a basic sequence of elements and then by adding witnesses at each level, such that $|a_{\mathbf{i}}^n| \leq |T|$ for each n, \mathbf{i} : since $|\hat{a}_{\mathbf{j}}^{n-1}| \leq |T|$ for all $\mathbf{j} \in [\mathbf{i}]^{n-1}$ (by the induction hypothesis), $|T|$ witnesses should indeed suffice. Then apply Lemma 3.35. QED_{3.37}

Proposition 3.38. Let α be an ordinal, $M = (a_{\mathbf{i}}^n : n < \omega, \mathbf{i} \in [\alpha]^n)$ an Ehrenfeucht-Mostowski pre-model $A \in [M]^{|T|}$. Then there are at most $|T|$ types over A realised in M .

Proof. The classical proof works. There is $J \in [\alpha]^{\leq |T|}$ such that $A \subseteq M_J = (a_{\mathbf{i}}^n : n < \omega, \mathbf{i} \in [J]^n)$. For every $n < \omega$, $\mathbf{i} \in [\alpha]^n$ and singleton (or finite tuple) $b \in \hat{a}_{\mathbf{i}}^n$, the type $\text{tp}(b/A)$ depends only on the position of b in $\hat{a}_{\mathbf{i}}^n$ and on the quantifier-free type of \mathbf{i} in $(\alpha, <, c_j : j \in J)$ (where the constant c_j is interpreted as j). There are $|T|$ possibilities for the former since $|a_{\mathbf{i}}^n| \leq |T|$, and for the latter since $|J| \leq |T|$ and α is well ordered. QED_{3.38}

4. REFINING SIMPLICITY AND STABILITY

4.1. Simplicity.

Definition 4.1. We say that $p(x)$ ε -divides over c if $p(x^\varepsilon)$ divides over c .

Proposition 4.2. A partial type $p(x, b)$ divides over c if and only if it ε -divides over c for some positive distance ε .

Proof. One direction is clear. For the other, assume that p divides over c . Since $p(x)$ is equivalent to $\bigwedge_{\varepsilon > 0} p(x^\varepsilon)$, there is a finite set of positive distances $(\varepsilon_i : i < n)$ such that $\bigwedge_{i < n} p(x^{\varepsilon_i})$ divides over c . Then $\varepsilon' = \bigwedge_{i < n} \varepsilon_i$ is positive, as the intersection of finitely many positive distances, and $p(x^{\varepsilon'})$ divides over c . QED_{4.2}

Corollary 4.3. *T is simple if and only if, for every set A , singleton a , and $\varepsilon > 0$, there is $A_0 \subseteq A$ such that $|A_0| \leq |T|$ and $\text{tp}(a/A)$ does not ε -divide over A_0 .*

Definition 4.4. Let ε be a distance on the sort of α -tuples (or on any imaginary sort, but we still designate the sort by the symbol α). Then $\kappa_\alpha^\varepsilon(T)$ is the minimal cardinal (or ∞ if non exists) such that for every a (of the appropriate sort or length) and every set A , there exists $A_0 \subseteq A$, $|A_0| < \kappa_\alpha^\varepsilon(T)$, such that $\text{tp}(a/A)$ does not ε -divide over A_0 . Define $\kappa_\alpha(T) = \sup\{\kappa_\alpha^\varepsilon(T) : \varepsilon > 0 \text{ is a distance on the appropriate sort}\}$.

Proposition 4.5. *$\kappa_\alpha(T)$ does not depend on α .*

Proof. We prove that $\kappa_\alpha(T) = \kappa_1(T)$ for every α .

We first show that $\kappa_\alpha(T) \geq \kappa_1(T)$. Let A be a set, a a singleton, and $\varepsilon > 0$ a distance on the home sort. Let $\varepsilon' = \varepsilon \times \infty^{\alpha-1}$: ε' is a positive distance ε' on the sort of α -tuples and $d(a_{<\alpha}, b_{<\alpha}) \leq \varepsilon' \iff d(a_0, b_0) \leq \varepsilon$. Let a' be the α -tuple consisting of α occurrences of a , and let $p(x) = \text{tp}(a/A)$, $p'(x') = \text{tp}(a'/A)$. Then by definition there is $A_0 \subseteq A$, $|A_0| < \kappa_\alpha(T)$, such that $p'(x'^{\varepsilon'})$ does not divide over A_0 . Therefore $p(x^\varepsilon)$ does not divide over A_0 , as required.

For $\kappa_\alpha(T) \leq \kappa_1(T)$, we may assume that T is simple, since otherwise $\kappa_1(T) = \infty$ by [HKP00, Remark 4.3]. By Lemma 2.5 it would suffice to prove that $\kappa_n^\varepsilon(T) \leq \kappa_1(T)$ for every $n < \omega$ and every $\varepsilon > 0$ of the form $\varepsilon = \prod_{i < n} \varepsilon_i$ (so $\varepsilon_i > 0$ for every $i < n$).

Let now $a_{<n}$ be an n -tuple, and A a set. By definition there is $A_0 \in [A]^{<\kappa_1(T)}$ such that $p_0(x_0, A) = \text{tp}(a_0/A)$ does not ε_0 -divide over A_0 . In other words, $p_0(x_0^{\varepsilon_0}, A)$ does not divide over A_0 , so there is b_0 such that $d(a_0, b_0) \leq \varepsilon_0$ and $b_0 \downarrow_{A_0} A$. Proceed similarly, setting at each step $p_i(x_i, Ab_{<i}) = \text{tp}(a_i/Ab_{<i})$, and finding $A_i \in [A]^{<\kappa_1(T)}$ and b_i such that $d(a_i, b_i) \leq \varepsilon_i$ and $b_i \downarrow_{A_i b_{<i}} A$.

Set $A_n = \bigcup_{i < n} A_i \in [A]^{<\kappa_1(T)}$. Then $b_i \downarrow_{A_n b_{<i}} A$ for each i whereby $b_{<n} \downarrow_{A_n} A$, witnessing that $\text{tp}(a_{<n}/A)$ does not ε -divide over A_n , as required. QED_{4.5}

We call this common value $\kappa(T)$.

Remark 4.6. In a first-order theory the zero distance is positive, so $\kappa(T) = \kappa_1^0(T)$ and our definition of $\kappa(T)$ coincides with the classical one. However, when we turn to look on infinitary sorts we see that $\kappa_\lambda^0(T) = \lambda + \kappa_1^0(T)$, since the zero distance on infinite tuples is practically never positive.

So T is simple if and only if $\kappa(T) < \infty$, and we may define:

Definition 4.7. A cat T is *supersimple* if $\kappa(T) = \omega$. In other words, if ε -dividing of is well founded, for every positive ε .

4.2. Stability. We are going to re-define the notion on λ -stability, and therefore that of superstability, in the spirit of what we did above. We only define λ -stability when $\lambda \geq |T|$ (where $|T|$ is the minimal cardinality of a basis for the topology on the set of types).

Definition 4.8. Let $\varepsilon > 0$ be symmetric, and $p, q \in S(A)$ be two types on the sort of ε . Then p and q are ε -consistent if $p(x) \wedge q(y) \wedge \varepsilon(x, y)$ is consistent. Otherwise, they are ε -inconsistent.

If $X \subseteq S(A)$ is such that every distinct $p, q \in X$ are ε -inconsistent, then X is ε -inconsistent. On the other hand, if every $p \in S(A)$ is ε -consistent with some $q \in X$ then X is ε -dense.

Notice that p and q are ε -consistent if and only if $p(x) \vdash q(x^\varepsilon)$.

Definition 4.9. T is α - λ -stable if $\lambda \geq |T|$, and for every set $|A| \leq \lambda$ and $\varepsilon > 0$ on α -tuples there is an ε -dense set $X \in [S_\alpha(A)]^{\leq \lambda}$.

T is λ -stable if it is 1- λ -stable.

T is stable if it is λ -stable for some λ .

Lemma 4.10. Assume that $\lambda \geq |T|$ and $\alpha > 0$. Then the following are equivalent:

- (i) T is λ -stable.
- (ii) T is α - λ -stable.
- (iii) For every $|A| \leq \lambda$ and $\varepsilon > 0$, if $X \subseteq S_\alpha(A)$ is ε -inconsistent then $|X| \leq \lambda$.

Proof. (i) \implies (ii). First let us prove that if T is λ -stable then it is n - λ -stable for every $0 < n < \omega$.

For $n = 1$ this is the assumption, so we assume for n and prove for $n + 1$. Assume therefore that $|A| \leq \lambda$ and $\varepsilon' > 0$ is a distance on $(n + 1)$ -tuples. One can find a distance $\varepsilon > 0$ on singletons such that $\varepsilon' > \varepsilon^{n+1}$. By assumption there are singletons $\{a_i : i < \lambda\}$ such that $\{\text{tp}(a_i/A)\}$ is ε -dense in $S_1(A)$, and for every $i < \lambda$ there are n -tuples $\{b_{i,j} : j < \lambda\}$ such that $\{\text{tp}(b_{i,j}/Aa_i) : j < \lambda\}$ is ε^n -dense in $S_n(Aa_i)$. It follows that $\{\text{tp}(b_{i,j}a_i/A) : i, j < \lambda\}$ is ε^{n+1} -dense in $S_{n+1}(A)$, and *a fortiori* ε' -dense.

The case where α is infinite reduced to the finite case by Lemma 2.5.

- (ii) \implies (iii). Assume that $\varepsilon > 0$ is given, and let $\varepsilon' > 0$ be such that $\varepsilon > 2\varepsilon'$. Then the existence of an ε' -dense subset $X \in [S_\alpha(A)]^{\leq \lambda}$ implies that every ε -contradictory set must be of cardinality at most λ .
- (iii) \implies (i). Assume not. Then there is $\varepsilon > 0$ and $X \in [S_1(A)]^{\lambda^+}$ which is ε -inconsistent. Let $X' \subseteq S_\alpha(A)$ be its image under the diagonal map, and $\varepsilon' = \varepsilon \times \infty^{\alpha-1}$. Then X' is ε' -inconsistent, even though $\varepsilon' > 0$.

QED_{4.10}

Remark 4.11. We changed (weakened) the definition of λ -stability, but not that of stability.

Proof. Since our definition of λ -stability is weaker than that of [Ben03b, Section 3] (where stability is defined by counting types, without any density considerations), stability in the sense of [Ben03b, Section 3] implies stability in our sense.

Conversely, if T is not stable in the sense of [Ben03b, Section 3], then there are contradicting formulas $\varphi_0(x, y)$ and $\varphi_1(x, y)$ such that $R(x = x, \varphi_0, \varphi_1, 2) = \infty$. By compactness

there exists $\varepsilon > 0$ such that $\varphi_0(x_0, y) \wedge \varphi_1(x_1, y) \wedge \varepsilon(x_0, x_1)$ is inconsistent. For every cardinal λ there is μ be such that $2^{<\mu} \leq \lambda < 2^\mu$, and as $R(x = x, \varphi_0, \varphi_1, 2) = \infty$ there are $\{b_\sigma : \sigma \in 2^{<\mu}\}$ such that $\{\varphi_{\sigma(i)}(x, b_{\sigma \upharpoonright i}) : i < \mu\}$ is consistent for every $\sigma \in 2^\mu$.

Let $B = \{b_\sigma : \sigma \in 2^{<\mu}\}$, and for every $\sigma \in 2^\mu$ let $p_\sigma(x) \in S_n(B)$ be any completion of $\{\varphi_{\sigma(i)}(x, b_{\sigma \upharpoonright i}) : i < \mu\}$ (here $n = |x|$). Then the set $\{p_\sigma : \sigma \in 2^\mu\}$ is an ε -inconsistent set of cardinality $2^\mu > \lambda$ over a set of parameters of cardinality $|B| = 2^{<\mu} \leq \lambda$. By Lemma 4.10 T is not λ -stable. Thus T is not stable. QED_{4.11}

Lemma 4.12. *A cat T is λ -stable if and only if $\lambda \geq |T|$ and for every $|A| \leq \lambda$ there is $B \supseteq A$, $|B| \leq \lambda$ such that every 1-type over A is realised in \bar{B} .*

Proof. First, since the home sort is assumed to be metric, it follows from Lemma 4.10 that T is λ -stable if and only if for all such A there is $X \subseteq S_1(A)$ which is ε -dense for every $\varepsilon > 0$ and $|X| \leq \lambda$.

For right to left, $X = \{\text{tp}(b/A) : b \in B\}$ is indeed ε -dense in $S_1(A)$ for all $\varepsilon > 0$.

For left to right, let $X \subseteq S_1(A)$ be ε -dense for every $\varepsilon > 0$ and $|X| \leq \lambda$. Consider the tree $t = X^{<\omega}$, and let $t_0 \subseteq t$ be the sub-tree consisting only of sequences $\bar{p} = (p_0, \dots, p_n)$ such that $p_m(x) \vdash p_{m+1}(x^{\frac{1}{2^m}})$ for every $m < n$.

Choose by induction on n elements $b_{\bar{p}}$ for every $\bar{p} = (p_0, \dots, p_n) \in t_0$ such that $b_{\bar{p}} \models p_n$ and such that $d(b_{\bar{p}, p_{n+1}}, b_{\bar{p}}) \leq \frac{1}{2^n}$. Define $B = \{b_{\bar{p}} : \bar{p} \in t_0\}$, so $|B| \leq \lambda$.

Let $q \in S_1(A)$ be given. For every $n < \omega$ choose $p_n \in X$ such that $p_n(x) \wedge q(y) \wedge d(x, y) \leq \frac{1}{2^{n+1}}$. Then $(p_0, \dots, p_n) \in t_0$ for every n , and the sequence $(b_{(p_0, \dots, p_n)} : n < \omega)$ is a Cauchy sequence whose limit is a realisation of q . Thus \bar{B} realises every 1-type over A . QED_{4.12}

Theorem 4.13 (Stability spectrum). *Let T be stable, and let $\lambda_0(T)$ be the minimal cardinality at which T is stable. Then T is λ -stable if and only if $\lambda^{<\kappa(T)} = \lambda \geq \lambda_0(T)$.*

Proof. We prove left to right first. By definition of $\lambda_0(T)$, if T is λ -stable then $\lambda \geq \lambda_0(T)$. If $\lambda^{<\kappa(T)} \neq \lambda$, then there is $\mu < \kappa(T)$ such that $\lambda^\mu > \lambda$, and we may also assume that μ is minimal as such, so $\lambda^{<\mu} = \lambda$. Since $\mu < \kappa(T)$ there exist a positive distance $\varepsilon > 0$ and an increasing sequence of types $(p_i(x, b_{<i}) : i < \mu)$, such that $p_{i+1}(x)$ ε -divides over $b_{<i}$ for all i . In other words, $p_{i+1}(x^\varepsilon)$ divides over $b_{<i}$, say with respect to a k_i -inconsistency witness $\psi_i(z_{<k_i})$. Let $I = \bigcup_{i < \mu} \lambda^{i+1}$. One can construct $(c_\sigma : \sigma \in I)$ such that $(c_{\tau \upharpoonright i+1} : i < \mu) \equiv (b_i : i < \mu)$ for every $\tau \in \lambda^\mu$, and for every $i < \mu$, $\sigma \in \lambda^i$ and $j_0 < \dots < j_{k_i-1} < \lambda : \models \psi_i(c_{\sigma, j_0}, \dots, c_{\sigma, j_{k_i-1}})$.

For each $i < \mu$ and $\sigma \in \lambda^i$ define $q_\sigma \in S(c_{\in \lambda^{<i}})$ by induction on i : for $i = j+1$ let q_σ be any extension of $q_{\sigma \upharpoonright j}(x) \wedge p_{i+1}(x, c_{\sigma \upharpoonright k+1} : k < i)$, and for i limit take $q_\sigma = \bigwedge_{j < i} q_{\sigma \upharpoonright j}$. Consider some $\sigma \in \lambda^i$, and let $J_\sigma = \{(\sigma, l) : l < \lambda\} \subseteq \lambda^{i+1}$. Then for every $\sigma' \in J_\sigma$ there is a finite bound ($k_i - 1$, if one wants to be precise) on the cardinality of $\{\sigma'' \in J_\sigma : q_{\sigma'}(x) \vdash q_{\sigma''}(x^\varepsilon)\}$ (this is just because ψ_i is a k_i -inconsistency witness for $p_{i+1}(x^\varepsilon)$). Thus, by a repeated process of choosing an element of λ and throwing finitely many away, we may in fact assume that for no two distinct $\sigma', \sigma'' \in J_\sigma$ are $q_{\sigma'}$ and $q_{\sigma''}$ ε -consistent. For every $\tau \in \lambda^\mu$ let $q_\tau(x) = \bigwedge_{i < \mu} q_{\tau \upharpoonright i} \in S(c_{\in \lambda^{<\mu}})$. Then $\{q_\tau : \tau \in \lambda^\mu\}$ are pairwise ε -inconsistent over

$\lambda = \lambda^{<\mu}$ parameters, so T is not λ -stable.

For the converse, assume that $\lambda = \lambda^{<\kappa(T)} \geq \lambda_0(T)$. Since $\lambda \geq |T|$, it would suffice to prove that if M is a pre-model, $|M| \leq \lambda$ and $\varepsilon > 0$ then there is an ε -dense set $X \in [S(M)]^{\leq \lambda}$. Since $\lambda = \lambda^{<\kappa(T)}$, we can enumerate $[M]^{<\kappa(T)} = (A_i : i < \lambda)$. Since $\kappa(T) \leq |T|^+$, we can choose for every A_i a pre-model $A_i \subseteq M_i \subseteq M$ of cardinality $|T|$. We know that $\lambda_0(T) \geq |T|$ by definition, so for every $i < \lambda$ there is $X_i \in [S(M_i)]^{\leq \lambda_0(T)}$ which is ε -dense. Since T is stable and M_i is a pre-model, every type of M_i has a unique non-dividing extension to M , which is its co-heir. Let X be the collection of all these extensions, so $|X| \leq \lambda$.

For every $p \in S(M)$ there is $i < \lambda$ such that $p(x)$ does not ε -divide over M_i . This means that there is $q \in S(M)$, ε -consistent with p , which is the unique non-dividing extension of $q_0 = q \upharpoonright_{M_i}$. We can also find $q'_0 \in X_i$ which is ε -consistent with q_0 , and let $q' \in X$ be its unique non-dividing extension to M . Let $a \models q_0$, $b \models q'_0$ such that $d(a, b) \leq \varepsilon$ and $ab \downarrow_{M_i} M$: then $a \models q$ and $b \models q'$, showing that q and q' are ε -consistent. Therefore p and q' are 2ε -consistent, so X is ε -dense. As for every $\varepsilon' > 0$ there is $\varepsilon > 0$ such that $\varepsilon' > 2\varepsilon$, this shows that T is λ -stable. QED_{4.13}

Remark 4.14. As in first order theories we may define $\kappa_r(T)$ to be $\kappa(T)$ if it is regular, or $\kappa(T)^+$ otherwise. Since for singular κ we have $\lambda^\kappa = (\lambda^{<\kappa})^{<\kappa}$, the condition $\lambda^{<\kappa(T)} = \lambda$ is equivalent to $\lambda^{<\kappa_r(T)} = \lambda$.

Theorem 4.15 (Existence of saturated models). *If T is λ -stable then it has a approximately saturated model of cardinality λ .*

Proof. Assume that T is λ -stable.

Choose a pre-model M_0 of cardinality $\leq \lambda$. Given a pre-model $|M_\alpha| \leq \lambda$, let $X_\alpha \subseteq S_1(M_\alpha)$ be distance-dense of cardinality $\leq \lambda$, and let $M_{\alpha+1}$ be a pre-model containing M_α and a realisation of every $p \in X_\alpha$, and still we may take $|M_{\alpha+1}| \leq \lambda$. For δ limit, define $M_\delta = \bigcup_{\alpha < \delta} M_\alpha$, and we stop at $M = M_\lambda$. We wish to prove that \bar{M} is approximately λ -saturated.

Claim. Let $(\varepsilon_i : i < \mu < \lambda)$ be positive distances on the sort of M , $\bar{\varepsilon} = \bigwedge \varepsilon_i$, $p \in S_1(M)$ and $\varepsilon > 0$ a positive distance for singletons. Then $p(x^{3\varepsilon}, M^{2\varepsilon})$ is realised in M .

Proof of claim. For $\alpha < \lambda$ let $p_\alpha = p \upharpoonright M_\alpha$, and find $a_\alpha \in M_{\alpha+1}$ realising $p_\alpha(x^\varepsilon)$.

By Lemma 2.5, every ε_i is only concerned with finitely many elements of M , whereby $\bar{\varepsilon}$ only is concerned with less than λ elements of M . If λ is regular, all these elements belong to some M_α , and $a_{\alpha+1} \models p(x^\varepsilon, M^\varepsilon)$.

We may therefore assume that λ is singular. We may also assume that the family $(\varepsilon_i : i < \mu)$ is closed under finite conjunctions. Since $\lambda = \lambda^{<\kappa(T)}$ we have in particular $\lambda > \text{cf}(\lambda) \geq \kappa(T)$, so there must be some $\alpha_0 < \lambda$ such that $p(x^\varepsilon)$ does not divide over M_{α_0} . Without loss of generality, we may assume that $\alpha_0 = 0$.

Let $q \vdash p(x^\varepsilon)$ be a complete type over M that does not divide over M_0 , and $q_\alpha = q \upharpoonright M_\alpha$. For each α , choose $b_\alpha \models q_\alpha$. Then $b_\alpha \downarrow_{M_0} M_\alpha$, and we may also choose it such that also

$d(b_\alpha, a_\alpha) \leq 2\varepsilon$ and $b_\alpha \downarrow_{a_\alpha M_\alpha} M b_{<\alpha}$.

We prove that $b_{<\alpha} \downarrow_{M_\alpha} M$ by induction on α . For α limit this would follow from the induction hypothesis for $\beta < \alpha$. For successors:

$$b_\alpha \downarrow_{a_\alpha M_\alpha} M b_{<\alpha} \implies b_\alpha \downarrow_{b_{<\alpha} M_{\alpha+1}} M \implies b_{<\alpha+1} \downarrow_{M_{\alpha+1}} M$$

On the other hand, we can also continue:

$$b_\alpha \downarrow_{a_\alpha M_\alpha} M b_{<\alpha} \implies b_\alpha \downarrow_M b_{<\alpha} \implies b_\alpha \downarrow_{M_\alpha} b_{<\alpha} \implies b_\alpha \downarrow_{M_0} b_{<\alpha}$$

So the sequence $(b_i : i < \lambda)$ is independent over M_0 . Consider now $r(Y) = \text{tp}(M/M_0, b_{<\lambda})$: for every $i < \mu$ there is $I_i \in [\lambda]^{<\kappa(T)}$ such that $r(Y^{\varepsilon_i})$ does not divide over $M_0 b_{\in I_i}$. Let $I = \bigcup_{i < \mu} I_i$. Since the family $(\varepsilon_i : i < \mu)$ is closed under finite conjunction, $r(Y^\varepsilon)$ does not divide over $M_0 b_{\in I}$, so there is $N \in M^\varepsilon$ such that $N \downarrow_{M_0 b_{\in I}} b_{<\lambda}$. Since $|I| \leq \kappa(T) \times \mu < \lambda$ we can choose $i \in \lambda \setminus I$, and $b_i \downarrow_{M_0 b_{\in I}} N \implies b_i \downarrow_{M_0} N$.

Find M' such that $M' \equiv_{NM_0} M$ and $M' \downarrow_{NM_0} b_i$. Then $d(M, M') \leq 2\varepsilon$, and $b_i \downarrow_{M_0} M'$. Since $M' \equiv M$ and $q(x, M)$ is the unique non-dividing extension of $q_0(x, M_0)$, $\models q(b_i, M')$. Since $q(x, Y) \vdash p(x^\varepsilon, Y)$, we sum up and obtain $\models p(a_i^{3\varepsilon}, M^{2\varepsilon})$, as promised. QED_{Claim}

In order to conclude that \bar{M} is approximately saturated we need a few slight improvements. First, we only considered types over M , but M and \bar{M} are interdefinable, so this is fine. Second, we proved for 1-types, but by induction one obtains for n -types: as for types of infinitely long tuples, since we only wish to realise something like $p(x^\varepsilon, M^{\varepsilon'})$ for $\varepsilon > 0$, this reduces to the finite case. QED_{4.15}

Lemma 4.16. *Assume that T is λ -stable, $|A| \leq \lambda$, and $(a_i : i < \lambda^+)$ is some sequence in a finite (or countable) sort. Let $\varepsilon > 0$ be on that sort.*

Then there exists a type $p(x_{<\lambda^+})$ of an A -indiscernible sequence, and a $I \subseteq \lambda^+$ of order type λ^+ , such that $a_{\in I} \models p(x_{<\lambda^+}^\varepsilon)$.

Proof. Since T is λ -stable: $\lambda^{<\kappa(T)} = \lambda$.

We now construct $(\bar{b}_i : i < \lambda^+) = (b_i^j : j < i < \lambda^+)$ in the following manner. At the i th step, we assume that $\bar{b}_{<i}$ are known, and set $A_i = A a_{<i} \bar{b}_{<i}$. Since $|A_i| + |T| \leq \lambda$, we can embed A_i in a pre-model of cardinality λ , whence there is an ε -dense set $X_i \in [\text{S}(\text{bdd}(A_i))]^{\leq \lambda}$. For every $j < i$ we add to X_i the (unique) non-dividing extension of every $p \in X_j$.

For $j < i$ define $\zeta_i^j = \min\{k \leq j : \text{tp}(a_i/A_j)$ does not ε -divide over $A_k\}$ (at the worst, $\zeta_i^j = j$). Then there is $c_i^j \downarrow_{A_{\zeta_i^j}} A_j$ such that $d(a_i, c_i^j) \leq \varepsilon$, and there is $p_i^j \in X_{\zeta_i^j}$ which is ε -consistent with $\text{lstp}(c_i^j/A_{\zeta_i^j})$. This means that there is $b_i^j \models p_i^j$ such that $d(b_i^j, c_i^j) \leq \varepsilon$, and we may assume that $b_i^j c_i^j \downarrow_{A_{\zeta_i^j}} A_j$. To sum up, $d(b_i^j, a_i) \leq 2\varepsilon$, and $\text{lstp}(b_i^j/A_j)$ is the unique non-dividing extension of p_i^j .

For every $i < \lambda^+$ let us define $\tau_i \in 2^{\lambda^+}$ by $\tau_i(j) = 0$ if and only if $\text{tp}(a_i/A_{j+1})$ does not ε -divide over A_j . For $j < \lambda^+$ and $\sigma \in 2^j$ define $I_\sigma = \{i < \lambda^+ : \sigma < \tau_i\}$ and $\text{otp}_1(\sigma) = \text{otp}\{k < j : \sigma(k) = 1\}$. Then define $\xi_j = \min\{\text{otp}_1(\sigma) : \sigma \in 2^j, |I_\sigma| = \lambda^+\}$. By definition of $\kappa(T)$: $I_\sigma \neq \emptyset \implies \text{otp}_1(\sigma) < \kappa(T)$, and clearly $|\{\sigma \in 2^j : \text{otp}_1(\sigma) < \kappa(T)\}| \leq \lambda^{<\kappa(T)} = \lambda$ for every $j < \lambda^+$. It follows that if $|I_\sigma| = \lambda^+$ (as is the case for $\sigma = \emptyset$), then σ has arbitrarily long extensions in $2^{<\lambda^+}$ having the same property. Therefore ξ_j is always defined, and is increasing with j . Since it is bounded by $\kappa(T)$, there must be $\xi < \kappa(T)$ and $j_0 < \lambda^+$ such that $\xi_j = \xi$ for all $j_0 \leq j < \lambda^+$.

Define an increasing sequence $(j_i \leq k_i < \lambda^+ : i < \lambda^+)$. We have already chosen j_0 ; for $i > 0$ we let $j_i = \sup\{k_l + 1 : l < i\} < \lambda^+$. Then there is $\sigma \in 2^{j_i}$ such that $\text{otp}_1(\sigma) = \xi$ and $|I_\sigma| = \lambda^+$, so let $k_i = \min I_\sigma \setminus j_i$. By definition of ξ it follows that $\text{tp}(a_{k_i}/A_{j_i})$ does not ε -divide over A_{j_0} . Therefore $\zeta_{k_i}^{j_i} \leq j_0$, so $\text{tp}(b_{k_i}^{j_i}/A_{j_i})$ does not divide over A_{j_0} .

Set $d_i = b_{k_i}^{j_i}$. Then $(d_i : i < \lambda^+)$ is an A_{j_0} -independent sequence. Moreover, $\text{lstp}(d_i/A_{j_0}) \in X_{j_0}$ which is a set of cardinality λ , so by passing to a subsequence, still of length λ^+ , we may further assume that $\text{lstp}(d_i/A_{j_0})$ is constant. Then $(d_i : i < \lambda^+)$ is a Morley sequence over A_{j_0} , and therefore in particular an A -indiscernible sequence. QED_{4.16}

Proposition 4.17. *Assume that T is λ -stable, $|A| \leq \lambda$, and $(a_i : i < \lambda^+)$ is some sequence in a metric sort.*

Then there exists a type $p(x_{<\lambda^+})$ of an A -indiscernible sequence, and a decreasing chain of subsets $\lambda^+ \supseteq I_0 \supseteq I_1 \supseteq \dots$, all of which of order type λ^+ , such that $a_{\in I_n} \models p(x_{<\lambda^+}^{\frac{1}{n}})$ for every $n < \omega$ (where $\frac{1}{n}$ is interpreted according to a fixed definable metric).

Proof. By induction we find such a decreasing sequence and types p_n such that $a_{\in I_n} \models p_n(x_{<\lambda^+}^{\frac{1}{2^n}})$. Then $\bigwedge p_n(x_{<\lambda^+}^{\frac{1}{2^n}})$ is finitely consistent and therefore consistent, and it is a complete type since $\frac{1}{2^n} \rightarrow 0$. Call this type p . Then $a_{\in I_n} \models p(x_{<\lambda^+}^{\frac{1}{2^n}})$ as required. QED_{4.17}

4.3. Morley rank. In this section we will define a variant of the Morley rank. In the classical setting of a first order theory, one defines the Morley rank of formulas with parameters (and subsequently of partial types). In a first order theory, a definable set can be characterised as a type-definable set whose complement is type-definable as well. Since in our context we do not have the luxury of clopen bases for the type-spaces, we will define the Morley rank for type-definable sets and their complements.

For the time being, we work with sets in a fixed sort.

Convention 4.18. A set (in the fixed sort) definable by a partial type with parameters (in a given set of parameters A) will be called (A) -closed. Here F, G , etc. always denote closed sets.

An (A) -open set is the complement of an (A) -closed set. Here U, V , etc. always denote open sets.

Given a set of parameters A , any A -invariant set can be viewed as a subset of $S(A)$, and then it is closed (open) by the definition above if and only if it is closed (open) in

the type-space topology. In particular, a set is A -closed (A -open) if and only if it is A -invariant and closed (open).

Since we let the set of parameters vary, we do not have a proper notion of closure or of interior: if X is an A -invariant set then there is a minimal A -closed set containing it, but this closure may decrease as A increases. Therefore when considering a given family of such sets it may be convenient to fix a set of parameters A over which they are all invariant. Then we can work in $S(A)$, where each A -invariant set is a set of points (complete A -types), and a well-defined notion of closure exists.

Definition 4.19. If U and V are open sets then V is *strongly contained* in U , in symbols $V < U$, if there is a closed set F such that $V \subseteq F \subseteq U$ (note that K is clopen if and only if $K < K$).

Fact 4.20. Assume that $F \subseteq U$ (so F is closed and U is open, by our notational conventions). Then there is V (open) such that $F \subseteq V < U$. Moreover, if both F and U are A -invariant, then we may choose V to be A -invariant as well.

Proof. This is elementary point-set topology. QED_{4.20}

Lemma 4.21. Let X be an A -invariant set: open, closed, or just an arbitrary collection of A -types, and assume that $F \subseteq X \subseteq U$, where F and U are closed and open, respectively, not necessarily A -invariant. Then there are sets F' and U' , A -closed and A -open, respectively, such that $F \subseteq F' \subseteq X \subseteq U' \subseteq U$.

Proof. Let B be a set of parameters containing A and the parameters for F and U . Consider F as a subset of $S(B)$, and let $F' = \{p \upharpoonright_A : p \in F\} \subseteq S(A)$. Then, when identified with their respective sets of realisations we have $F \subseteq F' \subseteq X$. Passing to complements we get the same result on the other side. QED_{4.21}

Definition 4.22. Let ε be a distance. Two closed sets F and G are ε -disjoint if $F(x) \wedge G(y) \wedge \varepsilon(x, y)$ is contradictory (in other words, every two realisations of F and G , respectively, are at distance greater than ε).

In particular, 0-disjointness is ordinary disjointness.

Definition 4.23. Let ε be a distance. We define the ε -Morley rank of open and closed sets simultaneously by induction. If U is an open set we define:

- (i) $\text{MR}_\varepsilon(U) \geq 0$ if $U \neq \emptyset$.
- (ii) $\text{MR}_\varepsilon(U) \geq \alpha$ for limit α if $\text{MR}_\varepsilon(U) \geq \beta$ for every $\beta < \alpha$.
- (iii) $\text{MR}_\varepsilon(U) \geq \alpha + 1$ if there is an infinite family of disjoint closed sets $(F_i : i < \omega)$ in U with $\text{MR}_\varepsilon(F_i) \geq \alpha$ for all $i < \omega$.

If F is a closed set, then $\text{MR}_\varepsilon(F) \geq \alpha$ if $\text{MR}_\varepsilon(U) \geq \alpha$ whenever $F \subseteq U$.

We set $\text{MR}_\varepsilon(U) = \sup\{\alpha : \text{MR}_\varepsilon(U) \geq \alpha\}$: if this is a bounded set then the maximum is attained, otherwise $\text{MR}_\varepsilon(U) = \infty$. Do the same for closed sets (the maximum is attained

for a closed set F since it is for every U containing it).

(MR stands for “Rang de Morley”, but the wrong way around.)

Remark 4.24. (i) If F is an A -closed set then $\text{MR}_\varepsilon(F) = \min\{\text{MR}_\varepsilon(U) : U \text{ is } A\text{-open and } F \subseteq U\}$. This is immediate from the definitions and from Lemma 4.21.

(ii) If K is a clopen set then its Morley rank as an open or as a closed set coincide.

The rank we defined is continuous in the following sense:

Fact 4.25. *For every A , ε and α , the set $\{p \in S(A) : \text{MR}_\varepsilon(p) \geq \alpha\}$ is closed.*

Proof. It is the complement of the union of all A -open sets with ε -Morley rank less than α . QED_{4.25}

One of the most important properties of any rank is that if one removes a small set from a big one, one is left with a big set:

Proposition 4.26. (i) *If $\text{MR}_\varepsilon(F) < \text{MR}_\varepsilon(U)$ then $\text{MR}_\varepsilon(U \setminus F) = \text{MR}_\varepsilon(U)$.*

(ii) *If $V \cap F \subseteq G$ and $\text{MR}_\varepsilon(G) < \text{MR}_\varepsilon(F)$ then $\text{MR}_\varepsilon(F \setminus V) = \text{MR}_\varepsilon(F)$.*

Proof. We prove simultaneously by induction on α that if $\text{MR}_\varepsilon(F) < \alpha \leq \text{MR}_\varepsilon(U)$ then $\text{MR}_\varepsilon(U \setminus F) \geq \alpha$ (we refer to this as the “open case”), and if $V \cap F \subseteq G$ and $\text{MR}_\varepsilon(G) < \alpha \leq \text{MR}_\varepsilon(F)$ then $\text{MR}_\varepsilon(F \setminus V) \geq \alpha$ (the “closed case”).

(i) We prove the open case. If $\alpha = 0$ this is immediate, and if α is limit, then this follows from the induction hypothesis for the open case. We therefore assume for α , and prove the open case for $\alpha + 1$.

Since $\text{MR}_\varepsilon(F) \leq \alpha$, there are open W', W such that $\text{MR}_\varepsilon(W) \leq \alpha$ and $F \subseteq W' < W$. Since $\text{MR}_\varepsilon(U) \geq \alpha + 1$ there are ε -disjoint $(G_i : i < \omega)$ in U such that $\text{MR}_\varepsilon(G_i) \geq \alpha$. Fix a set of parameters A over which all the mentioned sets are invariant, so we can have a proper notion of closure.

Since $\text{MR}_\varepsilon(W) \leq \alpha$, we can only have $\text{MR}_\varepsilon(G_i \cap \bar{W}') \geq \alpha$ for finitely many values of i , and we might as well assume that $\text{MR}_\varepsilon(G_i \cap \bar{W}') < \alpha$ for every i . Then $\text{MR}_\varepsilon(G_i \setminus W') \geq \alpha$ by the induction hypothesis for the closed case, and $G_i \setminus W' \subseteq U \setminus F$, so $\text{MR}_\varepsilon(U \setminus F) \geq \alpha + 1$.

(ii) We assume the open case for α , and prove the closed case for α .

We assume that $F \cap V \subseteq G$ and $\text{MR}_\varepsilon(G) < \alpha$, so we can find $G \subseteq W' < W$ such that $\text{MR}_\varepsilon(W) < \alpha$. If U is such that $F \setminus V \subseteq U$, then $F \subseteq U \cup W'$, whereby $\text{MR}_\varepsilon(\bar{W}') < \alpha \leq \text{MR}_\varepsilon(U \cup W')$, and by the open case: $\text{MR}_\varepsilon(U) \geq \text{MR}_\varepsilon((U \cup W') \setminus \bar{W}') = \text{MR}_\varepsilon(U \cup W') \geq \alpha$. Thus $\text{MR}_\varepsilon(F \setminus V) \geq \alpha$. (Here the closure \bar{W}' can be calculated in $S(A)$ for any set A over which all the sets involved are defined.)

QED_{4.26}

Lemma 4.27. *Let \mathcal{F} be a family of closed sets closed under finite intersections, and $G = \bigcap \mathcal{F}$. Then $\text{MR}_\varepsilon(G) = \min_{F \in \mathcal{F}} \text{MR}_\varepsilon(F)$.*

Proof. Clearly, $F \in \mathcal{F} \implies \text{MR}_\varepsilon(G) \leq \text{MR}_\varepsilon(F)$. On the other hand, assume that $G \subseteq U$: by compactness, $F_0 \subseteq U$ for some $F_0 \in \mathcal{F}$, so $\text{MR}_\varepsilon(U) \geq \text{MR}_\varepsilon(F_0) \geq \min_{F \in \mathcal{F}} \text{MR}_\varepsilon(F)$.

QED_{4.27}

Proposition 4.28. *Assume that F is A -closed, and for once consider it as a subset of $S(A)$. Then $\text{MR}_\varepsilon(F) = \max\{\text{MR}_\varepsilon(p) : p \in F\}$ (in particular, the maximum is attained).*

Proof. Clearly, $\text{MR}_\varepsilon(F) \geq \text{MR}_\varepsilon(p)$ for every $p \in F$. Conversely, assume that $p \in F$ and $\text{MR}_\varepsilon(p) < \text{MR}_\varepsilon(F)$. Then one can find $p \in V < U$ such that $\text{MR}_\varepsilon(U) < \text{MR}_\varepsilon(F)$. Setting $F' = F \setminus V$ we have $\text{MR}_\varepsilon(F') = \text{MR}_\varepsilon(F)$ and $p \notin F'$.

Write $F_0 = F$, $F_{i+1} = F'_i$, and at the limits $F_i = \bigcap_{j < i} F_j$, and continue as long as there is $p \in F_i$ such that $\text{MR}_\varepsilon(p) < \text{MR}_\varepsilon(F_i) = \text{MR}_\varepsilon(F)$. If the construction stops at F_α then F_α cannot be empty, and $\text{MR}_\varepsilon(p) = \text{MR}_\varepsilon(F)$ for every $p \in F_\alpha$.

QED_{4.28}

Corollary 4.29. *If $\text{MR}_\varepsilon(U) = \alpha$ then there is a finite bound on the size of a family of 2ε -disjoint closed sets in U of rank α . We call this bound the ε -Morley degree of U , denoted $\text{Md}_\varepsilon(U)$.*

Proof. Assume that for every $n < \omega$ there is a family $\{F_i^n : i \leq n\}$ of 2ε -disjoint closed subsets of rank α . Let A be a set of parameters for all of them (including U), and for every $n < \omega$ and $i \leq n$ let $p_i^n \in F_i^n$ be a type of rank α , by Proposition 4.28. We will now choose $q_n \in \{p_i^n : i \leq n\}$ by induction on $n < \omega$ such that the family $\{q_n : n < \omega\}$ is ε -disjoint. Assume that $\{q_m : m < n\}$ are already chosen. If each of p_i^n (for $i \leq n$) is not ε -disjoint of some q_m ($m < n$), then by the pigeonhole principle there are $i < j \leq n$ and $m < n$ such that both p_i^n and p_j^n are not ε -disjoint of q_m , and therefore not 2ε -disjoint of one another, contradicting the assumption. Therefore there is at least one p_i^n which is ε -disjoint of every q_m : we may therefore choose $q_n = p_i^n$ and proceed with the induction. At the end, the family $\{q_n : n < \omega\} \subseteq U$ is ε -disjoint and $\text{MR}_\varepsilon(q_n) = \alpha$ for every $n < \omega$, contradicting $\text{MR}_\varepsilon(U) = \alpha$.

QED_{4.29}

The usage of 2ε instead of ε in the definition of the Morley degree may be bothering, but it seems inevitable, and in any case does not prevent the notion from being quite useful. For closed sets we define the Morley degree much like we did for the Morley rank.

Definition 4.30. If F is a closed set and $\text{MR}_\varepsilon(F) = \alpha$ then $\text{Md}_\varepsilon(F)$ is defined as $\min\{\text{Md}_\varepsilon(U) : F \subseteq U, \text{MR}_\varepsilon(U) = \alpha\}$.

Note that by Lemma 4.21, if F is A -closed, then we may restrict ourselves to sets U which are A -open.

Notation 4.31. If U is an A -open set and $\text{MR}_\varepsilon(U) < \infty$, then $\text{mMR}_{\varepsilon,A}(U) = \{p \in U : \text{MR}_\varepsilon(p) = \text{MR}_\varepsilon(U)\}$, where U is considered as a subset of $S(A)$.

Lemma 4.32. *Let $p \in S(A)$ and assume that $\text{MR}_\varepsilon(p) < \infty$. Then there is an open set $U \subseteq S(A)$ such that $p \in U$, $\text{MR}_\varepsilon(U) = \text{MR}_\varepsilon(p)$ and $\text{Md}_\varepsilon(U) = \text{Md}_\varepsilon(p)$, and given a basis for the topology we may assume that U is basic.*

Such a U determines p over A up to 2ε , in the sense that $p \in \text{mMR}_{\varepsilon,A}(U) \subseteq p(x^{2\varepsilon})$, where the latter is interpreted as a subset of $S(A)$.

Proof. By definition of $\text{Md}_\varepsilon(p)$ there is an open set $U \subseteq S(A)$ containing p having the same Morley rank and degree as p . Given a basis for the topology on $S(A)$, we can find a basic open set $p \in U' \subseteq U$ which would necessarily have the same Morley degree and rank as p and U . By reason of Morley degree, if q is another type of the same ε -Morley rank as p in U (or U'), then it cannot be 2ε -disjoint of p . QED_{4.32}

Definition 4.33. T is *totally transcendental* if for every $\varepsilon > 0$, every closed or open set has an ordinal ε -Morley rank.

Proposition 4.34. *A cat with a countable language is totally transcendental if and only if it is ω -stable.*

Proof. Assume that T is totally transcendental. If A is a countable set, since the language is also countable there is a countable basis of open sets for $S(A)$. Assume that $\varepsilon' > 0$ is given, and find $\varepsilon > 0$ such that $\varepsilon' > 2\varepsilon$. Let $X \subseteq S(A)$ be a countable set intersecting $\text{mMR}_{\varepsilon,A}(U)$ for every basic open $U \subseteq S(A)$ such that $\text{mMR}_{\varepsilon,A}(U) \neq \emptyset$. By Lemma 4.32 for every $p \in S(A)$ there is a basic open set U such that $p \in \text{mMR}_{\varepsilon,A}(U) \subseteq p(x^{2\varepsilon}) \subseteq p(x^{\varepsilon'})$, so $p(x^{\varepsilon'}) \cap X \neq \emptyset$. It follows that X is ε' -dense, so T is ω -stable.

Conversely, assume that T is not totally transcendental, so $\text{MR}_\varepsilon(S(T)) = \infty$ for some $\varepsilon > 0$. Since there are, up to automorphism, boundedly many basic closed sets (in Hebrew: formulas with parameters), there is a bound α on the ranks of all *ranked* basic closed sets, and therefore on the ranks of all ranked closed or open sets.

Assume that U is an open set, with $\text{MR}_\varepsilon(U) = \infty > \alpha$. Then there is an infinite family of ε -disjoint unranked closed subsets in U , and let F, F' be two of them. By compactness we may now find open sets V, V' and closed sets \hat{F}, \hat{F}' such that $F \subseteq V \subseteq \hat{F} \subseteq U$, $F' \subseteq V' \subseteq \hat{F}' \subseteq U$ and \hat{F}, \hat{F}' are ε -disjoint. It follows that $\text{MR}_\varepsilon(V) = \text{MR}_\varepsilon(V') = \infty$.

Iterating this process we obtain $(V_\sigma : \sigma \in 2^{<\omega})$ such that $V_{\sigma i} < V_\sigma$ for every $\sigma \in 2^{<\omega}$ and $i < 2$, and $V_{\sigma 0} \cap V_{\sigma 1} = \emptyset$. Since $V_{\sigma i} < V_\sigma$ we can find a basic set $F_{\sigma i}$ such that $V_{\sigma i} \subseteq F_{\sigma i} \subseteq V_\sigma$. Since a basic set requires only finitely many parameters, we only need countably many for all of the sets $(F_\sigma : 2^{<\omega})$, and over this countable set there are 2^ω ε -inconsistent types, contradicting ω -stability. QED_{4.34}

4.4. Prime models. We assume as usual that T is metric.

Lemma 4.35. *If U is A -open and $\text{MR}_\varepsilon(U) < \infty$ for every $\varepsilon > 0$ then U contains a principal type over A .*

Proof. Choose a sequence $\varepsilon_n \rightarrow 0$ of positive distances. We construct by induction a decreasing sequence of A -open sets $U = U_0 > U_1 > \dots$ such that $U_{i+1} \subseteq p(x^{4\varepsilon_i})$ for every i and $p \in U_{i+1}$.

We start with $U = U_0$. Once U_i is known, we choose $p_i \in U_i$ of minimal $\text{MR}_{\varepsilon_i}$ -rank. By Lemma 4.32 there is an open set U_{i+1} such that $p_i \in \text{mMR}_{\varepsilon_i,A}(U_{i+1}) \subseteq p_i(x^{2\varepsilon_i})$, and we

may always assume that $U_{i+1} < U_i$ (otherwise find $p_i \in W < U_i$ and replace U_{i+1} with $W \cap U_{i+1}$). Since p_i is of minimal $\text{MR}_{\varepsilon_i}$ -rank in U_i , and *a fortiori* in U_{i+1} , we get in fact in fact $U_{i+1} = \text{mMR}_{\varepsilon_i, A}(U_{i+1}) \subseteq p_i(x^{2\varepsilon_i})$.

It follows that $U_{i+1} \subseteq p(x^{4\varepsilon_i})$ for every $p \in U_{i+1}$.

Let $X = \bigcap_{i < \omega} U_i$. Then $X = \bigcap_{i < \omega} \bar{U}_i \neq \emptyset$, and we may choose $p \in X$. Then for every $\varepsilon > 0$ there is i such that $\varepsilon > 4\varepsilon_i$, so $p \in U_{i+1} \subseteq p(x^\varepsilon)$. This shows that $p \in U$ is principal. QED_{4.35}

Definition 4.36. Let $A \subseteq B$. Then B is *constructed* over A if there is an enumeration $B = \{b_i : i < \alpha\}$ such that $\text{tp}(b_i/Ab_{<i})$ is principal for every $i < \alpha$.

A *constructed model* over A is a complete model $\bar{M} \supseteq A$ such that the pre-model M is constructed over A .

Proposition 4.37. *If T is totally transcendental then constructed models exist over any set.*

Proof. Let A be a set, and let $\{\exists y_i \neg \varphi_i(x, y_i, a_i) : i < \lambda\}$ enumerate all non-empty basic Q-open sets in $S_1(A)$ (so x is a singleton, but y_i may be a finite tuple). Construct a sequence $(b_i : i < \lambda)$ as follows. By Lemma 4.35, for every i there are b_i, c_i such that $\neg \varphi_i(b_i, c_i, a_i)$ and $\text{tp}(b_i, c_i/Ab_{<i})$ is principal. By Lemma 3.28 $\text{tp}(b_i/Ab_{<i})$ is principal. Thus $A' = Ab_{<\lambda}$ is constructed over A . Now let $A_0 = A$ and $A_{n+1} = A'_n$. Then $M = \bigcup_{n < \omega} A_n$ is a pre-model, and \bar{M} is a constructed model over A . QED_{4.37}

By Proposition 3.25 if M is a constructed model over A then it is prime over A , in the sense that if M' is a complete model containing A then M embeds in M' over A . (We are not interested in the issue of uniqueness of prime model).

Proposition 4.38. *A constructed model (over A) is atomic (over A), by which we mean that the type of every element over A is principal.*

Proof. Let \bar{M} be a constructed model over A , where M is constructed as a pre-model over A . By Proposition 3.27, $\text{tp}(M/A)$ is principal. By Lemma 3.28, $\text{tp}(\bar{M}/A)$ is principal, as is $\text{tp}(a/A)$ for every $a \in \bar{M}$. QED_{4.38}

5. THE MAIN THEOREM

We assume throughout that T is countable.

Proposition 5.1. *If T is categoric in an uncountable cardinal λ then T is ω -stable.*

Proof. Assume that T is not ω -stable, so there is a countable set A , a positive distance ε and a sequence $(a_i : i < \omega_1)$ such that $\{\text{tp}(a_i/A)\}$ are ε -inconsistent. Since $|Aa_{<\omega_1}| = \omega_1 \leq \lambda$ there is a complete model $N \supseteq Aa_{<\omega_1}$ of density character λ .

Let M be an Ehrenfeucht-Mostowski pre-model of length λ . Then \bar{M} is the unique complete model of density character λ , and we may therefore assume that $N = \bar{M}$.

Let $\varepsilon' > 0$ be such that $\varepsilon > 3\varepsilon'$. Let $B \subseteq M$ be countable such that $A \subseteq \bar{B}$, and for

every $i < \omega_1$ find $b_i \in M$ such that $d(b_i, a_i) < \varepsilon'$. Then $\{\text{tp}(b_i/B)\}$ are ε' -inconsistent, and therefore distinct, contradicting Proposition 3.38. QED_{5.1}

We prove:

Theorem 5.2. *A countable cat T is categoric in some uncountable cardinal, then it is categoric in every uncountable cardinal.*

Proof. We will prove this along the lines of [She75]. For the purposes of the proof we may assume that T is categoric in some uncountable cardinal, whereby it is ω -stable. It follows that T has a approximately saturated model in every cardinal. For $\lambda > \omega$ approximate saturation coincides with saturation, so T actually has a saturated model in every uncountable cardinality, which is unique by the usual back-and-forth argument (in fact the separable approximately saturated model is also unique, but this requires a more elaborate argument, and in any case unnecessary for our purposes).

It is left to show that if there exists a non-saturated model in $\lambda > \omega$ then there exists one in every $\lambda' > \omega$. Assume therefore that M is a complete model, $\|M\| = \lambda$, but M is not λ -saturated. Then M is not even approximately λ -saturated, so there are $A \subseteq M$ with $|A| = \mu < \lambda$, a distance $\varepsilon_1 > 0$, and a type $p \in S(A)$ such that $p(x^{\varepsilon_1})$ is omitted in M .

Since $\|M\| = \lambda > \mu$ there is $\varepsilon_0 > 0$ and a sequence $(a_i : i < \mu^+)$ in M such that $d(a_i, a_j) > \varepsilon_0$ for all $i < j < \mu^+$.

By Proposition 4.17 there is an A -indiscernible sequence $(b_i : i < \mu^+)$ such that, if $q(x_{<\mu^+}) = \text{tp}(b_{<\mu^+}/A)$, then for every $\varepsilon > 0$ there is a subsequence $(a_{i_j} : j < \mu^+)$ satisfying $q(x_{<\mu^+}^\varepsilon)$.

Let $\varphi(x, t, y_{<n}, z)$ be any formula, and assume that $\exists t \neg \varphi(x, t, b_{<n}, a)$ is consistent where $a \in A$. Then we claim that $\exists t \neg \varphi(x, t, b_{<n}, a) \not\vdash p(x^{\varepsilon_1})$. Indeed, let c, d be such that $\neg \varphi(c, d, b_{<n}, a)$. By Lemma 3.23 there exist a formula $\psi(x, t, y_{<n}, A)$ and $\varepsilon_2 > 0$ such that $\neg \psi(c, d, b_{<n}^{\leq \varepsilon_2}, A)$ and $\neg \psi(x, t, y_{<n}^{\leq \varepsilon_2}, A) \vdash \neg \varphi(x, t, y_{<n}, A)$. But then, for some $I \in [\mu^+]^n$, we may actually assume that $d(b_{<n}, a_{\in I}) < \varepsilon_2$, in which case $\neg \psi(c, d, a_{\in I}, A)$. Then there must $c', d' \in M$ such that $\neg \psi(c', d', a_{\in I}, A)$. But: $\neg \psi(c', d', a_{\in I}, A) \vdash \neg \psi(c', d', b_{<n}^{\leq \varepsilon_2}, A) \vdash \neg \varphi(c', d', b_{<n}, A)$, and we assumed that M omits $p(x^{\varepsilon_1})$.

Let $A_0 = \emptyset$. Assuming that $A_n \subseteq A$ is countable consider all consistent formulas of the form $\exists t \neg \varphi(x, t, b_{<n}, a)$ where $a \in A_n$. For each one, finitely many elements in A suffice in order to show that some realisation of $\exists t \neg \varphi(x, t, b_{<n}, a)$ does not satisfy $p(x^{\varepsilon_1})$. In all we need countably many elements from A , and their addition to A_n result in A_{n+1} which is still countable. Let $A' = \bigcup A_n$, and $p' = p \upharpoonright A'$. Therefore if $\exists t \neg \varphi(x, t, b_{<n}, a)$ is consistent and $a \in A'$ then $\exists t \neg \varphi(x, t, b_{<n}, a) \not\vdash p'(x^{\varepsilon_1})$.

It follows that if $b_{<\lambda'}$ is any sequence similar to b_{μ^+} over A' , then any completion of p' to a type over $A'b_{<\lambda'}$ cannot be principal, and therefore cannot be realised in a constructed model over $A'b_{<\lambda'}$. It follows that if $\lambda' > \omega$ then the constructed model over $A'b_{<\lambda'}$ is not saturated (it is in fact not even ω_1 -saturated), and T is not λ' -categoric. QED_{5.2}

REFERENCES

- [Ben] Itay Ben-Yaacov, *Schrödinger's cat*, submitted.
- [Ben03a] ———, *Positive model theory and compact abstract theories*, *Journal of Mathematical Logic* **3** (2003), no. 1, 85–118.
- [Ben03b] ———, *Simplicity in compact abstract theories*, *Journal of Mathematical Logic* **3** (2003), no. 2, 163–191.
- [Ben03c] ———, *Thickness, and a categoric view of type-space functors*, *Fundamenta Mathematicae* **179** (2003), 199–224.
- [Hen76] C. Ward Henson, *Nonstandard hulls of banach spaces*, *Israel Journal of Mathematics* **25** (1976), 108–144.
- [HI02] C. Ward Henson and José Iovino, *Ultraproducts in analysis*, *Analysis and Logic* (Catherine Finet and Christian Michaux, eds.), London Mathematical Society Lecture Notes Series, no. 262, Cambridge University Press, 2002.
- [HKP00] Bradd Hart, Byunghan Kim, and Anand Pillay, *Coordinatisation and canonical bases in simple theories*, *Journal of Symbolic Logic* **65** (2000), 293–309.
- [Iov96] José Iovino, *The morley rank of a banach space*, *Journal of Symbolic Logic* **61** (1996), no. 3, 928–941.
- [KP97] Byunghan Kim and Anand Pillay, *Simple theories*, *Annals of Pure and Applied Logic* **88** (1997), 149–164.
- [Poi85] Bruno Poizat, *Cours de théorie des modèles*, Nur al-Mantiq wal-Ma'rifah, 1985.
- [She75] Saharon Shelah, *The lazy model-theoretician's guide to stability*, *Logique et Analyse* **71–72** (1975), 241–308.

ITAY BEN-YAACOV, UNIVERSITY OF WISCONSIN – MADISON, DEPARTMENT OF MATHEMATICS,
480 LINCOLN DRIVE, MADISON, WI 53706-1388, USA

URL: <http://www.math.wisc.edu/~pezz>