

On d -finiteness in continuous structures

by

Itai Ben Yaacov (Lyon) and Alexander Usvyatsov (Los Angeles, CA)

Abstract. We observe that certain classical results of first order model theory fail in the context of continuous first order logic. We argue that this happens since finite tuples in a continuous structure may behave as infinite tuples in classical model theory. The notion of a d -finite tuple attempts to capture some aspects of the classical finite tuple behaviour. We show that many classical results involving finite tuples are valid in continuous logic upon replacing “finite” with “ d -finite”. Other results, such as Vaught’s no two models theorem and Lachlan’s theorem on the number of countable models of a superstable theory are proved under the assumption of enough (uniformly) d -finite tuples.

The main goal of this article is to describe and study conditions under which certain results of classical model theory generalise to the model theory of metric structures, and to explain why when they do not.

We start by recalling Henson’s adaptation of the Ryll-Nardzewski theorem to metric logics (originally for the logic of positive bounded formulae, but we state and prove it for continuous first order logic). It characterises the family of countable ω -categorical (i.e., separably categorical) continuous theories in a manner analogous to the classical result. One of the equivalent characterisations is that all models of T are *approximately* ω -saturated, which is a weaker property than plain ω -saturation; in particular, the unique separable model need not be ω -saturated in the classical sense.

A good example for this phenomenon is the theory T of L^p Banach lattices [BBH] (for a fixed $1 \leq p < \infty$). Up to isomorphism, the unique separable model of this theory is $L^p[0, 1]$, which is therefore approximately ω -saturated. By quantifier elimination it embeds elementarily in $L^p[0, 2]$; however, $\text{tp}(\chi_{[1,2]}/\chi_{[0,1]})$ is a consistent type over a single parameter which is not realised in $L^p[0, 1]$, and so it is not ω -saturated in the classical sense.

2000 *Mathematics Subject Classification*: 03C90, 03C45, 03C95, 03C35.

Key words and phrases: continuous first order logic, approximately ω -saturated model, d -finite tuple, model theory of metric structures.

The first author is partially supported by NSF grant DMS-0500172

In Section 2 we explain this by arguing that “finite tuple” is not always the right notion in the setting of metric structures. Instead we define the notion of a d -finite tuple, and show (among other things) that in an approximately ω -saturated model every type over a d -finite tuple is realised. As we show that every finite tuple of events in a probability algebra is d -finite, this explains why models of the theory of atomless probability algebras *are* ω -saturated in the classical sense.

A second look at the example above might prove even more disturbing: Let now T' be the theory of $(L^p[0, 1], \chi_{[0,1]})$ in a language consisting of a new constant symbol c . Then up to isomorphism T' has precisely two separable models: $(L^p[0, 1], \chi_{[0,1]})$ and $(L^p[0, 2], \chi_{[0,1]})$ (which differ precisely on the question whether $\text{tp}(\chi_{[1,2]}/\chi_{[0,1]})$ is realised or not). This means that Vaught’s “no two models” theorem fails for continuous logic. Moreover, the theory T , and therefore T' , is superstable and indeed ω -stable: thus T' also serves as a counterexample for Lachlan’s theorem stating that a countable superstable theory has either one or infinitely many countable (or in our context, separable) models.

We explain this by observing that the theory of L^p Banach lattices does not have “enough d -finite elements”. Other continuous theories, like that of probability algebras or Hilbert spaces, do have this property. In Section 3 we prove Vaught’s theorem under the assumption of enough d -finites, and in Section 4 we prove Lachlan’s theorem under (almost) the same assumption.

We will use continuous first order logic as a framework for the model theory of metric structures. We will assume the reader is familiar with it. For general background we refer the reader to [BU, BBHU]. Much of the time we will work in T^{eq} , which is obtained from a theory T as in [BU, Section 5] (once we know how to add a single imaginary sort we can iterate this and add them all).

Most of the time we work implicitly inside a very saturated and homogeneous monster model. Thus all sets and tuples are considered to be taken inside such a model, and all models are elementary substructures of the monster model.

Given a set of parameters A and some logical property $s(x)$ defining an A -invariant set, we use $[s]^{S(A)}$ to denote the set

$$\{p \in S(A) : p(x) \text{ implies } s(x)\}.$$

If A is clear from the context we may omit the superscript. Note that $s(x)$ may be a partial type, but also something of the form $\varphi(x) < r$ (in which case $[s]$ is open).

We remind the reader that the symbols \vee and \wedge , which are used in classical logic to denote disjunction and conjunction, respectively, are also used in continuous first order logic as pointwise maximum and minimum of

formulae (i.e., join and meet, respectively, in the lattice of continuous first order formulae). This means that a statement of the form $(\varphi \wedge \psi) \leq r$ is semantically equivalent to the disjunction $(\varphi \leq r) \vee (\psi \leq r)$, and similarly for $(\varphi \vee \psi) \leq r$ and $(\varphi \leq r) \wedge (\psi \leq r)$. While in principle there should not be any ambiguity, this could turn out to be a little confusing, so in this paper we will do our best to restrict the use of the symbols \vee and \wedge to their lattice-theoretic meaning.

1. Preliminaries. Recall that every sort, be it the home sort(s) or any imaginary sort, comes equipped with an intrinsic metric. For finite tuples (of the same length, and coordinatewise in the same sorts) $a_{<n}$ and $b_{<n}$ we may define $d(\bar{a}, \bar{b}) = \max\{d(a_i, b_i) : i < n\}$. (We can view the sort of n -tuples as the sort of canonical parameters for the formula $\varphi(x_{<n}, y_{<n}) = \bigvee_{i < n} d(x_i, y_i)$: the canonical parameter for $\varphi(\bar{x}, \bar{a})$ is precisely \bar{a} , and the metric on this sort is the one given above.)

This approach is not adequate when considering infinite tuples (which we may wish to do). In the case of countable tuples we could cheat our way out by defining $d(a_{<\omega}, b_{<\omega}) = \sum 2^{-n-1}(1 \wedge d(a_i, b_i))$. A better approach, which is less arbitrary and extends well to uncountable tuples as well, is simply to define $d(a_{\in I}, b_{\in I})$ as the I -tuple $(d(a_i, b_i) : i \in I) \in [0, \infty]^I$, and redefine the way we compare I -tuples in $[0, \infty]$:

DEFINITION 1.1. By a *distance* we mean a member $\varepsilon \in [0, \infty]$. Let $\bar{\varepsilon}, \bar{\delta} \in [0, \infty]^I$.

- (i) We say that $\bar{\varepsilon} \geq \bar{\delta}$ if $\varepsilon_i \geq \delta_i$ for all $i \in I$.
- (ii) We say that $\bar{\varepsilon} > \bar{\delta}$ if $\varepsilon_i > \delta_i$ for all $i \in I$ and $\varepsilon_i = \infty$ for all but finitely many $i \in I$. For the purpose of this definition $\infty < \infty$.
- (iii) When comparing an I -tuple of distance with a single distance we treat the single distance as if it were an I -tuple whose every coordinate is that distance. (Thus $\bar{\varepsilon} > 0$, which is by far the most common instance of this rule, means that $\varepsilon_i > 0$ for all i , and $\varepsilon_i = \infty$ for all but finitely many i .)

We can now define a uniform structure on the space of I -tuples: the vicinities are given by positive tuples of distances (i.e., tuples satisfying $\bar{\varepsilon} > 0$ according to Definition 1.1). For finite and countable tuples, this uniform structure coincides with that defined by the metric in our first approach, while for any tuple length it is the inverse limit of the metric structures on the respective spaces of finite sub-tuples, justifying Definition 1.1(iii).

We recall that for all n (and sets of parameters A), the type of space $S_n(T)$ (or $S_n(A)$) is a compact Hausdorff topological space, whose closed sets are precisely the sets of the form $[p(\bar{x})]$ where p is a partial type (over A), and for which the family of sets of the form $[\varphi(\bar{x}) < 1/2]$, where φ is a

formula (with parameters in A), forms a basis of open sets. We also put a metric structure on this type space: $d(p, q)$ is the infimal distance between realisations of p and q (where distance between tuples is as above).

NOTATION 1.2.

- (i) For a (partial) type $p(\bar{x})$ and distances $\bar{\varepsilon}$, $p(\bar{x}^{\bar{\varepsilon}})$ denotes the partial type saying that p is satisfied somewhere in the $\bar{\varepsilon}$ -neighbourhood of \bar{x} , i.e.

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

Here the existential quantifier should be understood as “there exists in an elementary extension”. This is definable by a partial type by [BU, Fact 3.13].

- (ii) If $p(\bar{x}, \bar{y})$ is a (partial) type and \bar{a} a tuple of the same length as \bar{y} , then $p(\bar{x}^{\bar{\varepsilon}}, \bar{a}^{\bar{\delta}})$ is the obvious things, i.e., the result of substituting \bar{a} for \bar{y} in $p(\bar{x}^{\bar{\varepsilon}}, \bar{y}^{\bar{\delta}})$.
- (iii) Finally, if $p(\bar{x}, \bar{y}) = \text{tp}(\bar{a}, \bar{b})$, then $\text{tp}(\bar{a}^{\bar{\varepsilon}}/\bar{b}^{\bar{\delta}})$ denotes $p(\bar{x}^{\bar{\varepsilon}}, \bar{b}^{\bar{\delta}})$.

Thus, following Definition 1.1(iii), and when dealing with types in finitely many variables, say types in $S_n(T)$, we have

$$[p(\bar{x}^{\varepsilon})] = \{q \in S_n(T) : d(p, q) \leq \varepsilon\}.$$

DEFINITION 1.3. A structure M is *approximately ω -saturated* if for every finite tuple $\bar{a} \in M$, every type $p(x, \bar{a}) \in S_1(\bar{a})$ and every $\varepsilon > 0$, there is $b' \in M$ realising $p(x^\varepsilon, \bar{a}^\varepsilon)$.

This is equivalent to the following apparently stronger condition:

FACT 1.4. *Assume that a structure M is approximately ω -saturated. Then for every finite tuple $\bar{a} \in M$, every type $p(\bar{x}, \bar{a}) \in S(\bar{a})$ in at most countably many variables, and every $\varepsilon > 0$, there is a tuple \bar{a}' in M such that $d(\bar{a}, \bar{a}') \leq \varepsilon$ and $p(\bar{x}, \bar{a}')$ is realised in M .*

Proof. STEP I. We show that the definition of approximate ω -saturation holds with any finite tuple of variables (rather than a single one). Indeed, let $p(x_{<n}, \bar{a}) \in S_n(\bar{a})$ for some finite tuple $\bar{a} \in M$. For $i \leq n$, let $p_i(x_{<i}, \bar{y}) = p(x_{<n}, \bar{y}) \upharpoonright_{(x_{<i}, \bar{y})}$.

We will choose $b_{<n} \in M$ such that for all $i \leq n$: $\models p_i(b_{<i}^{(1-2^{-i})\varepsilon}, \bar{a}^{(1-2^{-i})\varepsilon})$.

- For $i = 0$, this is true ($b_{<0}$ is the empty tuple, and $\models p_0(\bar{a})$).
- Now assume that $i < n$, and $b_{<i}$ as above are already chosen. Then in some elementary extension of M there are $b'_{<n}$ such that $\models p(b'_{<n}, \bar{a}^{(1-2^{-i})\varepsilon})$ and $d(b_{<i}, b'_{<i}) \leq (1 - 2^{-i})\varepsilon$, which yields

$$\models p_{i+1}(b_{<i}^{(1-2^{-i})\varepsilon}, b'_i, \bar{a}^{(1-2^{-i})\varepsilon}).$$

Approximate ω -saturation yields $b_i \in M$ realising $\text{tp}(b_i^{\varepsilon/2^{i+1}} / (b_{<i}\bar{a})^{\varepsilon/2^{i+1}})$. Adding up distances we see that $\models p_{i+1}(b_{\leq i}^{(1-2^{-i-1})\varepsilon}, \bar{a}^{(1-2^{-i-1})\varepsilon})$, as required.

Thus $p(x_{<n}^\varepsilon, \bar{a}^\varepsilon)$ is realised by $b_{<n} \in M$.

STEP II. We now consider the general case of $p(x_{<\omega}, \bar{a})$, and show that $p(x_{<\omega}, \bar{a}^\varepsilon)$ is realised in M . For $n < \omega$, let $p_n(x_{<n}, \bar{y}) = p(x_{<\omega}, \bar{y}) \upharpoonright_{(x_{<n}, \bar{y})}$.

For each $n < \omega$ we will choose an n -tuple $b_{<n}^\varepsilon \in M$ with the property that $\models p_n((b_{<n}^\varepsilon)^{\varepsilon/2^n}, \bar{a}^{(1-2^{-n})\varepsilon})$:

- For $n = 0$, this holds trivially.
- Given $b_{<n}^\varepsilon$, we know there are $c_{<\omega}$ in some $N \succeq M$ such that $d(c_{<n}, b_{<n}^\varepsilon) \leq \varepsilon/2^n$ and $\models p(c_{<\omega}, \bar{a}^{(1-2^{-n})\varepsilon})$. By the first step there is an $(n+1)$ -tuple $b_{<n+1}^{\varepsilon/2^{n+1}} \in M$ realising $\text{tp}(c_{<n+1}^{\varepsilon/2^{n+1}} / (b_{<n}^\varepsilon)^{\varepsilon/2^{n+1}}, \bar{a}^{\varepsilon/2^{n+1}})$. Thus in particular we have $\models p_{n+1}((b_{<n+1}^{\varepsilon/2^{n+1}})^{\varepsilon/2^{n+1}}, \bar{a}_{n+1}^{(1-2^{-n-1})\varepsilon})$, and the construction may proceed.

It follows from the construction that

$$d(b_{<n}^\varepsilon, b_{<n+1}^{\varepsilon/2^{n+1}}) \leq \frac{\varepsilon}{2^{n+1}} + \frac{\varepsilon}{2^n} + \frac{\varepsilon}{2^{n+1}} = \frac{\varepsilon}{2^{n-1}}.$$

Thus, for each $i < \omega$, the sequence $(b_i^{\varepsilon/2^{i+1}} : n < \omega)$ is a Cauchy sequence which converges to some $b_i^\omega \in M$. Clearly $\models p(b_{<\omega}^\omega, \bar{a}^\varepsilon)$.

STEP III. Given $p(x_{<\omega}, \bar{a})$, let $q(x_{<\omega}, \bar{y}, \bar{a}) := p(\bar{x}, \bar{y}) \wedge \bar{y} = \bar{a}$. Then this is a complete type in countable many variables over finitely many parameters in M , and by the second step there are $b_{<\omega}, \bar{a}' \in M$ such that $\models q(b_{<\omega}, \bar{a}', \bar{a}^\varepsilon)$, which means that $\models p(b_{<\omega}, \bar{a}')$ and $d(\bar{a}', \bar{a}) \leq \varepsilon$, as required. $\blacksquare_{1.4}$

FACT 1.5. *Any two elementarily equivalent separable approximately ω -saturated structures are isomorphic.*

Proof. This was first observed by C. Ward Henson, but no proof exists in current literature.

Let M and N be two separable approximately ω -saturated models. Let $M_0 = \{a_i : i < \omega\}$ and $N_0 = \{b_i : i < \omega\}$ be countable dense subsets of M and N , respectively.

We will construct a sequence of elementary mappings $f_i : A_i \rightarrow N$ and $g_i : B_i \rightarrow M$, where $A_i \subseteq M$ and $B_i \subseteq N$ are finite, such that:

- (i) $A_0 = B_0 = \emptyset$, and for $i > 0$,

$$A_{i+1} = a_{\leq i} \cup A_i \cup g_i(B_i), \quad B_{i+1} = b_{\leq i} \cup B_i \cup f_{i+1}(A_{i+1}).$$

- (ii) For all $c \in A_i$, $d(c, g_i \circ f_i(c)) \leq 2^{-i}$.
 (iii) For all $c \in B_i$, $d(c, f_{i+1} \circ g_i(c)) \leq 2^{-i}$.

We start with $f_0 = \emptyset$, which is elementary as we assume that $M \equiv N$.

Assume that f_i is given. Then A_i is given, and is finite by the induction hypothesis, and this determines B_i which is also finite. Fix enumerations

for A_i and B_i as finite tuples, and let $p(\bar{x}, \bar{y}) = \text{tp}^N(B_i, f(A_i))$. As f_i is elementary, $p(\bar{x}, A_i)$ is a consistent type over M , and by approximate ω -saturation there are tuples $B'_i, A'_i \subseteq M$ such that $d(A_i, A'_i) \leq 2^{-i}$ and $M \models p(B'_i, A'_i)$. Then $g_i: B_i \mapsto B'_i$ will do.

We construct f_{i+1} from g_i similarly.

We now have, for all $c \in A_i$,

$$\begin{aligned} d(c, g_i \circ f_i(c)) \leq 2^{-i} &\rightarrow d(f_{i+1}(c), f_{i+1} \circ g_i \circ f_i(c)) \leq 2^{-i} \\ &\rightarrow d(f_{i+1}(c), f_i(c)) \leq 2^{-i+1}. \end{aligned}$$

Therefore the sequence of mappings f_i converges to a mapping $f: A \rightarrow N$, where $A = \bigcup A_i$, and by uniform continuity of the language, f is elementary. As $M_0 \subseteq A$ we have $\bar{A} = M$, and as f is an isometry it extends uniquely to a mapping $\bar{f}: M \rightarrow N$. Again by uniform continuity, \bar{f} is elementary. An elementary mapping $\bar{g}: N \rightarrow M$ is constructed similarly. For $i < j < \omega$ we have

$$\begin{aligned} d(a_i, \bar{g} \circ \bar{f}(a_i)) &\leq d(a_i, \bar{g} \circ f_j(a_i)) + 2^{-j+2} \\ &\leq d(a_i, g_{j+1} \circ f_j(a_i)) + 2^{-j+1} + 2^{-j+2} \\ &\leq 2^{-j} + 2^{-j+1} + 2^{-j+2} \leq 2^{-j+3}. \end{aligned}$$

By letting $j \rightarrow \infty$ we see that $\bar{g} \circ \bar{f}$ is the identity on M_0 , and therefore on M . Similarly $\bar{f} \circ \bar{g} = \text{id}_N$. $\blacksquare_{1.5}$

DEFINITION 1.6. Let $p(\bar{x}) \in \text{S}(A)$, and let \bar{a} be a tuple of the same length as \bar{x} . Then $d(\bar{a}, p)$ is defined as $\inf\{d(\bar{a}, \bar{b}) : \models p(\bar{b})\}$, where \bar{b} varies over all tuples of appropriate length in the monster model.

(Note that by applying compactness to the partial type $p(\bar{x}) \cup \{d(\bar{a}, \bar{x}) \leq d(\bar{a}, p) + 1/n : n < \omega\}$, we see that the infimum is in fact attained in the monster model.)

DEFINITION 1.7. Let $p(\bar{x}) \in \text{S}(A)$. We say that p is *isolated* if the predicate $\bar{a} \mapsto d(\bar{a}, p)$ is definable (with parameters from A). If $\varphi(\bar{x}, A)$ is an A -definable predicate such that $d(\bar{x}, p) = \varphi(\bar{x}, A)$, we say that $\varphi(\bar{x}, A)$ *isolates* p .

FACT 1.8. Let $p(\bar{x}) \in \text{S}(A)$. Then the following are equivalent:

- (i) p is isolated.
- (ii) For every $\varepsilon > 0$, the set $[p(\bar{x}^\varepsilon)] \subseteq \text{S}(A)$ has non-empty interior.
- (iii) For every $\varepsilon > 0$, the set $[p(\bar{x}^\varepsilon)] \subseteq \text{S}(A)$ forms a neighbourhood of p .
- (iv) For every $\varepsilon > 0$ there is a formula $\varphi(\bar{x}, A)$ such that

$$p(\bar{x}) \vdash \varphi(\bar{x}, A) = 0, \quad \varphi(\bar{x}, A) \leq 1/2 \vdash p(\bar{x}^\varepsilon).$$

- (v) There is an A -definable predicate $\varphi(\bar{x}, A)$ such that $p(\bar{x}) \vdash \varphi(\bar{x}, A) = 0$, and always $d(\bar{x}, p) \leq \varphi(\bar{x}, A)$.

Proof. (i) \Rightarrow (ii). Clear.

(ii) \Rightarrow (iii). For every $\varepsilon > 0$, the set $[p(\bar{x}^\varepsilon/3)]$ contains a non-empty open set, which may be taken to be of the form $[\varphi(\bar{x}, A) < \varepsilon/2]$ where $\inf_{\bar{x}} \varphi(\bar{x}, A) = 0$. Let

$$\psi(\bar{x}, A) = \inf_{\bar{y}} (\varphi(\bar{y}, A) \vee d(\bar{x}, \bar{y})).$$

(Here $d(\bar{x}, \bar{y}) = \bigvee_i d(x_i, y_i)$.) Then

$$p \in [\psi(\bar{x}, A) < \varepsilon/2] \subseteq [p(\bar{x}^\varepsilon)].$$

(iii) \Rightarrow (iv). We have $\varphi(x, A)$ and r such that $p \in [\varphi(x, A) < r] \subseteq [p(x^\varepsilon)]$. By subtraction and re-scaling we may assume that $p \vdash \varphi(x, A) = 0$ and $r = 1/2$.

(iv) \Rightarrow (v). For all $n < \omega$ choose $\varphi_n(\bar{x}, A)$ such that $p(\bar{x}) \vdash \varphi_n(\bar{x}, A) = 0$ and

$$\varphi_n(\bar{x}, A) \leq 1/2 \vdash p(\bar{x}^{2^{-n-1}}).$$

Now let

$$\varphi(\bar{x}, A) = \sum_{n < \omega} 2^{-n} \varphi_n(\bar{x}, A).$$

(Here it is understood that the sum is truncated at 1.) Then $\varphi(\bar{x}, A)$ is an A -definable predicate and clearly $p(\bar{x}) \vdash \varphi(\bar{x}, A) = 0$. On the other hand, if \bar{a} does not realise p let n be such that

$$2^{-n} \geq d(\bar{a}, p) > 2^{-n-1}.$$

Then $\varphi_m(\bar{a}, A) > 1/2$ for all $m \geq n$, and so $\varphi(\bar{a}, A) \geq 2^{-n}$, as required.

(v) \Rightarrow (vi). Given φ as in the assumption we have

$$d(\bar{x}, p) = \inf_{\bar{y}} d(\bar{x}, \bar{y}) + \varphi(\bar{y}, A). \quad \blacksquare_{1.8}$$

The omitting type theorem has been proved in [Ben05] in a somewhat different setting, namely that of Hausdorff cats. Since every continuous first order theory is in particular an (open) Hausdorff cat, the result we will need here is a special case of [Ben05, Theorem 3.17]. However, for the benefit of the reader who is not familiar with the cat setting, we will go quickly through the proof again. In addition, there was a small mistake in the statement of the result there regarding the omission of n -types for $n > 1$, which we correct here.

LEMMA 1.9. *Let T be a countable theory, $M \models T$ and $\bar{a} = (a_i : i < \omega) \in M^\omega$. Write $\bar{a} \preceq M$ if there is an elementary submodel $M_0 \preceq M$ such that every tail $(a_i : k < i < \omega)$ is dense in M_0 . Then there is a co-meagre set $Y \subseteq S_\omega(T)$ such that if $M \models T$, $\bar{a} \in M^\omega$ and $\text{tp}^M(\bar{a}) \in Y$ then $\bar{a} \preceq M$.*

Proof. Indeed, let $\varphi(x_{<n}, y)$ be any formula, where some of the variables may be dummies, and let $0 \leq r < s \leq 1$ be rational. Define

$$Y_{\varphi,r,s} = [\inf_y \varphi(x_{<n}, y) > r] \cup \bigcup_{k < \omega} [\varphi(x_{<n}, x_{n+k}) < s] \subseteq S_\omega(T),$$

$$Y_\varphi = \bigcap \{Y_{\varphi,r,s} : 0 \leq r < s \leq 1 \text{ and } r, s \in \mathbb{Q}\},$$

$$Y = \bigcap \{Y_\varphi : \varphi(\bar{x}, y) \in \mathcal{L}\}.$$

Then $Y_{\varphi,r,s}$ is clearly open. To see that it is also dense, let $p \in S_\omega(T)$, say realised in a model M by some \bar{a} . If $\inf_y \varphi(a_{<n}) > r$ then $p \in Y_{\varphi,r,s}$, so assume $\inf_y \varphi(a_{<n}, y) \leq r$. Then there is some $b \in M$ such that $\varphi(a_{<n}, b) < s$. Let U be any neighbourhood of p ; we may assume it is of the form $[\psi(\bar{x}_{<m}) < 1]$ where $m \geq n$. Define $b_i = a_i$ for $i < m$ and $b_i = b$ for $b \geq m$. Then $\text{tp}^M(\bar{b}) \in U \cap Y_{\varphi,r,s}$, so $Y_{\varphi,r,s}$ is dense. It follows that each Y_φ is co-meagre, and since the language is countable, Y is co-meagre.

Assume now that $M \models T$, $\bar{a} \in M^\omega$ and $\text{tp}(\bar{a}) \in Y$. Let M_0 be the closure in M of the set $\{a_i : i < \omega\}$. For each formula $\varphi(x_{<n}, y) \in \mathcal{L}$ we have $\text{tp}(\bar{a}) \in Y_\varphi$, and hence

$$\inf_y \varphi(a_{<n}, y)^M = \inf \{\varphi(a_{<n}, b)^M : b \in M_0\}.$$

Thus $M_0 \preceq M$ by the Tarski–Vaught test. Moreover, by adding dummy variables we obtain, for every $k \geq n$,

$$\inf_y \varphi(a_{<n}, y)^M = \inf \{\varphi(a_{<n}, a_m)^M : k \leq m < \omega\}.$$

Applying this to the formula $d(x_{n-1}, y)$ we see that $\{a_n : k \leq n < \omega\}$ is dense in M_0 for all k . $\blacksquare_{1.9}$

FACT 1.10. *For any two ordinals (or even mere index sets) α, β , every mapping $f: \alpha \rightarrow \beta$ induces a mapping $f^*: S_\beta(T) \rightarrow S_\alpha(T)$ sending $\text{tp}(a_i : i < \beta) \mapsto \text{tp}(a_{f(i)} : i < \alpha)$. The mapping f^* is continuous, and if f is injective then f^* is open. (We then say that if T is a continuous first order theory then $\alpha \mapsto S_\alpha(T)$ is an open type space functor).*

Proof. For continuity, observe that for $i_0, \dots, i_{n-1} < \alpha$,

$$f^{*-1}([\varphi(x_{i_0}, \dots, x_{i_{n-1}}) < r]) = [\varphi(x_{f(i_0)}, \dots, x_{f(i_{n-1})}) < r].$$

If f is injective, then up to a permutation of the indices we may assume that it is the inclusion $\alpha \subseteq \beta$. Then for all $i_0, \dots, i_{n-1} < \alpha \leq j_0, \dots, j_{m-1} < \beta$,

$$f^*([\varphi(x_{i_0}, \dots, x_{i_{n-1}}, x_{j_0}, \dots, x_{j_{m-1}}) < r])$$

$$= [\inf_{x_{j_0}, \dots, x_{j_{m-1}}} \varphi(x_{i_0}, \dots, x_{i_{n-1}}, x_{j_0}, \dots, x_{j_{m-1}}) < r]. \quad \blacksquare_{1.10}$$

THEOREM 1.11 (Omitting types theorem, strong form). *Let T be a countable theory, and for each n let $X_n \subseteq S_n(T)$ be a meagre set. Then there exists a model $M \models T$ such that for each $n < \omega$ a dense subset of M^n omits every type in X_n .*

Proof. The proof consists of several steps.

First we recall that $S_\omega(T)$ is the type space of ω -tuples of elements in models of T , and that this is a compact Hausdorff space. In particular it has the Baire property, i.e., a co-meagre set is never empty.

Second, using Lemma 1.9 and following its notation we have a co-meagre set $Y \subseteq S_\omega(T)$ such that for every $M \models T$ and $\bar{a} \in M^\omega$, if $\text{tp}(\bar{a}) \in Y$ then $\bar{a} \preceq M$.

Third, by Fact 1.10, if $X_n \subseteq S_n(T)$ is meagre and $f: n \hookrightarrow \omega$ is injective then $f^{*-1}(X_n) \subseteq S_\omega(T)$ is meagre as well.

We conclude that

$$Z = Y \setminus \bigcup_{n < \omega, f: n \hookrightarrow \omega} f^{*-1}(X_n) \subseteq S_\omega(T)$$

is co-meagre and therefore non-empty. Let $p \in Z$ and realise p by some $\bar{a} \in N \models T$. Let $M = \overline{\{a_i: i < \omega\}}$. Then $M \preceq N$, so $M \models T$. Also, for each $n < \omega$ let $M_n = \{(a_{i_0}, \dots, a_{i_{n-1}}): i_0 < \dots < i_{n-1} < \omega\}$. Then M_n is dense in M^n and omits every type in X_n . ■_{1.11}

REMARK 1.12. From the statement of [Ben05, Theorem 3.17] it would follow that we can have $M_n = M_1^n$ for all n , which is not true: for example, $X_2 = [d(x_0, x_1) = 0] \subseteq S_2(T)$ may be nowhere-dense, but cannot be omitted from M_1^2 .

COROLLARY 1.13. *Let T be a countable theory, and let $p \in S_n(T)$ be a non-isolated type. Then T has a model omitting p .*

Proof. If p is non-isolated then there is an $\varepsilon > 0$ such that $[p(x^\varepsilon)] \subseteq S(T)$ has empty interior, and is in particular meagre. Let $M \models T$ contain a dense subset omitting $p(x^\varepsilon)$. Then M omits p (and in fact $p(x^{\varepsilon'})$ for all $\varepsilon' < \varepsilon$). ■_{1.13}

FACT 1.14 (Ryll-Nardzewski theorem for continuous logic). *Let T be a complete countable theory. Then the following are equivalent:*

- (i) T is ω -categorical.
- (ii) Every n -type over the empty set is isolated, for all n .
- (iii) Every model of T is approximately ω -saturated.
- (iv) Every separable model of T is approximately ω -saturated.
- (v) The metric on $S_n(T)$ coincides with the logic topology for all n .
- (vi) The metric on $S_n(T)$ is compact for all n .

(This was known to C. Ward Henson for a long time.)

Proof. (i) \Rightarrow (ii). By Corollary 1.13, a non-isolated type can be omitted in a separable model, and of course can be realised in another.

(ii) \Rightarrow (iii). Let $M \models T$, $\bar{a} \in M$ finite, and $b \in N \succ M$. Then $p(x, \bar{y}) = \text{tp}(b\bar{a})$ is isolated by a definable function $\varphi(x, \bar{y}) = d(x\bar{y}, p)$. Then $N \models \inf_y \varphi(y, \bar{a}) = 0$, and hence $M \models \inf_y \varphi(y, \bar{a}) = 0$, so for every $\varepsilon > 0$ there is $c \in M$ such that $\varphi(c, \bar{a}) \leq \varepsilon$, i.e., $\models p(c^\varepsilon, \bar{a}^\varepsilon)$.

(iii) \Rightarrow (iv). Clear.

(iv) \Rightarrow (i). By Fact 1.5.

(ii) \Leftrightarrow (v). The metric always refines the logic topology, since all the formulae are uniformly continuous. By Fact 1.8, the logic topology refines the metric if and only if all the types are isolated.

(v) \Leftrightarrow (vi). Since the metric topology refines the logic topology which is compact and Hausdorff. ■_{1.14}

DEFINITION 1.15. A theory T is *small* if for all n the density character of $S_n(T)$ in the metric topology (denoted $\|S_n(T)\|$) is at most countable.

PROPOSITION 1.16. *A theory T is small if and only if it has a separable approximately ω -saturated model.*

Proof. Clearly, if T has a separable ω -saturated model then T is small. Conversely, assume that T is small. Then we may assume that the language \mathcal{L} of T is countable: otherwise, there is a countable sublanguage $\mathcal{L}_0 \subseteq \mathcal{L}$ such that any two distinct types of T differ on an \mathcal{L}_0 -formula, and we may reduce everything to \mathcal{L}_0 .

As T is small, choose for every $n < \omega$ a countable dense subset $X_n \subseteq S_n(T)$, and for every $p \in X_n$ introduce a new n -ary predicate symbol \hat{p} with the identity as uniform continuity modulus. Let $\hat{\mathcal{L}}$ be the expanded language.

Let $M \models T$ be ω -saturated. Expand it to an $\hat{\mathcal{L}}$ -structure \hat{M} by interpreting $\hat{p}(\bar{x})$ as $d(\bar{x}, p)$. As $|\hat{\mathcal{L}}| \leq \omega$, \hat{M} has a separable elementary submodel $\hat{N} \preceq \hat{M}$. Let $N = \hat{N} \upharpoonright_{\mathcal{L}}$. Then in particular $N \preceq M$ is a model of T .

Now let $\bar{a} \in N^m$ be a finite tuple and $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$. By ω -saturation there are $\bar{b} \in M$ realising p . Let $\varepsilon > 0$ be also given. Then there is $q \in X_{n+m}$ such that $d(q, p) < \varepsilon$, and so $\hat{q}(\bar{b}, \bar{a}) < \varepsilon$. As $\hat{N} \preceq \hat{M}$,

$$\left(\inf_{\bar{x}} \hat{q}(\bar{x}, \bar{a})\right)^{\hat{N}} = \left(\inf_{\bar{x}} \hat{q}(\bar{x}, \bar{a})\right)^{\hat{M}} < \varepsilon.$$

Therefore there exists $\bar{c} \in \hat{N}$ such that $\hat{q}(\bar{c}, \bar{a}) < \varepsilon$. By definition of \hat{q} we have $\models q(\bar{c}^\varepsilon, \bar{a}^\varepsilon)$ and thus $\models p(\bar{c}^{2\varepsilon}, \bar{a}^{2\varepsilon})$. This shows that N is approximately ω -saturated. ■_{1.16}

PROPOSITION 1.17. *Assume T is small. Then T has an atomic model (i.e., a model only realising isolated types).*

Proof. For every $n < \omega$ and $\varepsilon > 0$, let $X_{n,\varepsilon} \subseteq S_n(T)$ be the union of all open subsets of $S_n(T)$ of diameter smaller than ε . Let $K_{n,\varepsilon}$ be the complement of $X_{n,\varepsilon}$: this is a closed subset of $S_n(T)$.

Assume that $K_{\varepsilon,n}^\circ \neq \emptyset$ for some n, ε . Then by construction every open subset of $K_{\varepsilon,n}$ has diameter greater than ε , and by a tree argument we can find continuum many types in $K_{n,\varepsilon}$ with the distance between any two being at least $\varepsilon/2$, contradicting smallness.

Therefore $K_{\varepsilon,n}^\circ = \emptyset$ for all n, ε and we can find a model $M \models T$ which has a dense subset $M_0 \subseteq M$ omitting them all (just consider $\varepsilon = 1/m$ for $m \geq 1$). Assume that $p \in S_n(T)$ is realised in M . Then for every $\varepsilon > 0$, we know that $p(\bar{x}^\varepsilon)$ is realised in M_0 (since M_0 is dense), so $[p(\bar{x}^\varepsilon)] \cap X_{n,\varepsilon} \neq \emptyset$ and thus $[p(\bar{x}^{2\varepsilon})]$ has non-empty interior. Then p is isolated by Fact 1.8, so M only realises isolated types. ■_{1.17}

PROPOSITION 1.18. *Assume T is complete and countable and $M \models T$. Then M is prime if and only if M is atomic and separable.*

Proof. Clearly if M is prime then it must be separable and atomic. Conversely, assume M is separable and atomic. Let $\{a_i : i < \omega\}$ enumerate a dense subset. Viewing the tuple $a_{<\omega}$ with the metric as defined in the introduction we see that $\text{tp}(a_{<\omega})$ is isolated and therefore realised in every model of T . As $M = \text{dcl}(a_{<\omega})$, we obtain an elementary embedding of M into every model of T . ■_{1.18}

2. *d-finiteness*

DEFINITION 2.1. Let a and c be possibly infinite tuples, and let $p = \text{tp}(a/c)$. Here δ will denote a tuple of distances of the same length as a .

- (i) We say that a is *d-finite* over c , or that p is *d-finite*, if for every tuple b and a corresponding tuple of distances $\varepsilon > 0$ there is $\delta = \delta_{b,\varepsilon}^{a/c} > 0$ such that whenever $a' \equiv_c a$ and $d(a, a') \leq \delta$, there is b' such that $d(b, b') \leq \varepsilon$ and $a'b' \equiv_c ab$.
- (ii) We say that a is *uniformly d-finite* over c , and p is *uniformly d-finite*, if for every tuple length α and $\varepsilon > 0$ of length α there is $\delta_{\alpha,\varepsilon}^{a/c} > 0$ such that for every tuple b of length α we can take $\delta_{b,\varepsilon}^{a/c} = \delta_{\alpha,\varepsilon}^{a/c}$ and the same holds (i.e., $\delta_{b,\varepsilon}^{a/c}$ depends on $|b|$ rather than on b).

If $c = \emptyset$ we omit it.

Note that when testing for (uniform) *d-finiteness* we may assume that b is finite and ε is a single positive distance. Indeed, by Definition 1.1(iii), if $\varepsilon > 0$ then it is equal to ∞ on all but finitely many coordinates, and we may simply restrict to these coordinates. Then we may replace the tuple ε with its minimum.

LEMMA 2.2. *Let a , b and c be tuples, and assume that $\text{tp}(a/c)$ and $\text{tp}(b/ac)$ are (uniformly) d -finite. Then so is $\text{tp}(ab/c)$.*

Proof. Let a tuple e and $\varrho > 0$ of the same length be given. Let $\varepsilon = \delta_{e,\varrho/2}^{b/ac} > 0$ be given by d -finiteness of b over ac . Let $\delta = \delta_{be,(\varepsilon/2,\varrho/2)}^{a/c} > 0$ be given by d -finiteness of a over c . We claim that $\delta_{e,\varrho}^{ab/c}$ can be taken to be $(\delta, \varepsilon/2)$, which is indeed positive.

Indeed, assume now that $a'b' \equiv_c ab$ are such that $d(ab, a'b') \leq (\delta, \varepsilon/2)$. First, by choice of δ , there are b'', e' such that $a'b''e' \equiv_c abe$ and $d(b''e', be) \leq (\varepsilon/2, \varrho/2)$. Then $b' \equiv_{a'c} b''$ and $d(b', b'') \leq \varepsilon$, so there is e'' such that $d(e', e'') \leq \varrho/2$ and $b'e'' \equiv_{a'c} b''e'$. Then $d(e, e'') \leq \varrho$, and $a'b'e'' \equiv_c abe$.

Therefore the tuple ab is d -finite over c .

The proof for uniformly d -finite is similar. ■_{2.2}

PROPOSITION 2.3. *Let M be approximately ω -saturated, and let $a \in M$ be d -finite. Then M is approximately ω -saturated as a model of $T(a) = \text{Th}(M, a)$ (i.e., the theory of M with a named).*

Proof. Let b, c be finite tuples, $b \in M$, and $p(z, b, a) = \text{tp}(c/ab) \in S(ab)$. For every real number $\varepsilon > 0$ let $\delta = \delta_{bc, \varepsilon/2}^a$ (so we view $\varepsilon/2$ as a tuple of $|bc|$ repetitions of the number $\varepsilon/2$; it is positive as a tuple since $|bc| < \omega$).

By approximate ω -saturation there exists $\tilde{c} \in M$ realising $p(z^{\varepsilon/2}, b^{\varepsilon/2}, a^\delta)$; that is to say, there are a', b', c' such that $d(a, a') \leq \delta$, $d(b\tilde{c}, b'c') \leq \varepsilon/2$ and $\models p(c', b', a')$. By choice of δ there are \tilde{b}', \tilde{c}' such that $d(b'\tilde{c}', \tilde{b}'\tilde{c}') \leq \varepsilon/2$ and $\tilde{a}\tilde{b}'\tilde{c}' \equiv a'b'c'$, so $d(b\tilde{c}, \tilde{b}'\tilde{c}') \leq \varepsilon$ and $\models p(\tilde{c}', \tilde{b}', a)$. Therefore $\tilde{c} \models p(z^\varepsilon, b^\varepsilon, a)$, as required. ■_{2.3}

COROLLARY 2.4. *Assume T is small and $a \in M \models T$ is d -finite. Then $T(a) = \text{Th}(M, a)$ is small.*

Proof. As T is small, it has a separable approximately ω -saturated model M , and we may assume that $a \in M$. Then (M, a) is a separable approximately ω -saturated model of $T(a)$, which is therefore small. ■_{2.4}

COROLLARY 2.5. *If M is an approximately ω -saturated model of T and $a \in M$ is d -finite then every type in at most countably many variables over a is realised in M .*

We can prove a converse to Proposition 2.3 under the assumption that T is small.

PROPOSITION 2.6. *Let T be small, let a be a finite or countable tuple in a model of T , and $T(a) = \text{Th}(M, a)$. Then a is d -finite if and only if every model of $T(a)$ which is approximately ω -saturated as a model of T is also approximately ω -saturated as a model of $T(a)$.*

Proof. One direction is just Proposition 2.3.

For the other, we will improve on the proof of Proposition 1.16. Assume that a is a tuple which is not d -finite. Then there are a finite tuple b and $\varepsilon > 0$ which witness this. Let $p(x, a) = \text{tp}(b/a)$.

Let M be as in the proof of Proposition 1.16, and we may assume in addition that $a \in M$ and that M is sufficiently homogeneous and saturated for what will follow. Let $\widehat{T} = \text{Th}_{\widehat{\mathcal{L}}}(\widehat{M})$, and $\widehat{T}(a) = \text{Th}_{\widehat{\mathcal{L}}}(\widehat{M}, a)$. We claim that the partial type $p(x^{\varepsilon/2}, a)$ defines in $S(\widehat{T}(a))$ a nowhere-dense set.

Indeed, assume otherwise. Then there is an $\widehat{\mathcal{L}}$ -formula $\varphi(x, y)$ such that $\varphi(x, a) < 1/2$ is consistent and implies $p(x^{\varepsilon/2}, a)$. This means that $\widehat{M} \models \inf_x \varphi(x, a) < 1/2$, so there is $c \in M$ such that $\varphi(c, a) < 1/2$, and $c \models p(x^{\varepsilon/2}, a)$. As M is a sufficiently saturated model of T , there is $c' \in M$ such that $d(c, c') \leq \varepsilon/2$ and $\models p(c', a)$: in fact, we might as well assume that $c' = b$. By uniform continuity of φ there is $\delta > 0$ such that for all a' , if $d(a, a') \leq \delta$ then $\varphi(c, a') < 1/2$ as well.

By assumption on a and b , and by saturation of M , there exists $a' \in M$ such that $a' \equiv_{\mathcal{L}} a$ and $d(a, a') \leq \delta$ and yet for no $b' \in M$ do we have $d(b, b') \leq \varepsilon$ and $\models p(b', a')$. By the homogeneity assumption there is an automorphism $f \in \text{Aut}(M)$ such that $f(a) = a'$. Then $f \in \text{Aut}(\widehat{M})$ as well, so that $a' \equiv_{\widehat{\mathcal{L}}} a$. By choice of δ we have $\varphi(c, a') < 1/2$, and since $a' \equiv_{\widehat{\mathcal{L}}} a$, this implies that $\models p(c^{\varepsilon/2}, a')$. Therefore there is $b' \in M$ such that $d(b', c) \leq \varepsilon/2$ and $\models p(b', a')$. Then $d(b, b') \leq \varepsilon$, contradicting the choice of a' .

Thus $p(x^{\varepsilon/2}, a)$ indeed defines a nowhere dense set in $S(\widehat{T}(a))$. Also, for every $q \in X_n$ (where X_n is as in the construction of $\widehat{\mathcal{L}}$) and any two rationals $0 < r < s \leq 1$, the set defined by $q(y^r) \wedge \widehat{q}(\overline{y}) \geq s$ is closed and omitted in \widehat{M} , and is therefore also nowhere dense. Thus, by Theorem 1.11, there exists a model $(\widehat{N}, a) \models \widehat{T}(a)$ in which $p(x, a)$ is omitted, and for every $q \in X_n$ and $c \in \widehat{N}^n$ we have $\widehat{q}(c) \leq d(c, q)$. A compactness argument shows that $d(c, q) \leq \widehat{q}(c)$ is simply a consequence of \widehat{T} , so we may proceed as in the proof of Proposition 1.16 to conclude that $N = \widehat{N} \upharpoonright_{\mathcal{L}}$ is approximately ω -saturated as an \mathcal{L} -structure. As it omits $p(x, a)$, it is not approximately ω -saturated once a is named. $\blacksquare_{2.6}$

Being d -finite is a property of tuples implying they are well-behaved. One can derive from it a property defining well-behaved theories:

DEFINITION 2.7. We say that T has enough d -finite elements if for every single element in the home sort a , any tuple c , and $\varepsilon > 0$, there is an imaginary $b \in \text{dcl}(ac)$ such that:

- (i) b is d -finite over c .
- (ii) $\text{tp}(a/bc) \vdash d(x, a) \leq \varepsilon$ (i.e., b “captures” a^ε over c).

Same for uniformly d -finite.

PROPOSITION 2.8. *The following are equivalent:*

- (i) *T has enough d -finite elements.*
- (ii) *For every finite or countable tuple a , and every tuple c , there is a sequence of imaginaries $(b_i : i < \omega)$ such that each b_i is d -finite over $cb_{<i}$, and a is interdefinable with $b_{<\omega}$ over c .*
- (iii) *For every finite or countable tuple a , and every tuple c , there is a sequence of imaginaries $(b_i : i < \omega)$ such that $b_{<i}$ is d -finite over c for all $i < \omega$, and a is interdefinable with $b_{<\omega}$ over c .*
- (iv) *Same for a single element a in the home sort.*

Same for uniformly d -finite.

Proof. (i) \Rightarrow (ii). We may assume that $a = a_{<\omega}$ is a countable tuple. Choose some enumeration $((n_i, m_i) : i < \omega)$ of ω^2 . Choose a sequence $(b_i : i < \omega)$ in $\text{dcl}(a_{<\omega}, c)$ such that $b_i \in \text{dcl}(a_{n_i}cb_{<i}) \subseteq \text{dcl}(a_{<\omega}c)$ is d -finite over $cb_{<i}$, and $\text{tp}(a_{n_i}/cb_{<i}) \vdash d(x, a_{n_i}) \leq 2^{-m_i}$. It follows that $\text{tp}(a_n/cb_{<\omega}) \vdash x = a_n$, i.e., $a_{<\omega} \in \text{dcl}(cb_{<\omega})$, as required.

(ii) \Rightarrow (iii). By Lemma 2.2.

(iii) \Rightarrow (iv). Clear.

(iv) \Rightarrow (i). Let a be a singleton, c a tuple, and $\varepsilon > 0$. By assumption there is a sequence $(b_i : i < \omega)$ of imaginaries such that $\text{dcl}(ac) = \text{dcl}(b_{<\omega}c)$ and $b_{<i}$ is d -finite over c for all $i < \omega$. For all $i < \omega$, let $p_i(x) = \text{tp}(a/b_{<i}c)$. Since $a \in \text{dcl}(b_{<\omega}c)$, it is the unique realisation of $\bigwedge_{i < \omega} p_i(x)$. Therefore the partial type $d(x, a) \geq \varepsilon \wedge \bigwedge_{i < \omega} p_i(x)$ is inconsistent, so by compactness there is some $i < \omega$ such that $p_i(x) \vdash d(x, a) < \varepsilon$. Thus $b = b_{<i}$ is the imaginary we need. $\blacksquare_{2.8}$

PROPOSITION 2.9. *Let M be a structure, let a be a tuple in M^{eq} , $T = \text{Th}(M)$ and $T(a) = \text{Th}(M, a)$. Then the following are equivalent:*

- (i) *T is ω -categorical and a is uniformly d -finite (over \emptyset).*
- (ii) *T is ω -categorical and a is d -finite (over \emptyset).*
- (iii) *$T(a)$ is ω -categorical.*

Proof. Clearly, uniform d -finiteness implies d -finiteness.

Assume now that T is ω -categorical and a d -finite. Then all separable models of T are approximately ω -saturated, and by Proposition 2.3 so are all the separable models of $T(a)$, which is thereby ω -categorical.

Finally, assume that $T(a)$ is ω -categorical. By Fact 1.14, T is ω -categorical as well: since the metric on $S_n(T(a))$ is compact for all n it is also compact on its quotient $S_n(T)$. So assume for contradiction that a is not uniformly d -finite. Then there exist n and $\varepsilon > 0$ such that for all $\delta > 0$ of the same length as a there exists an n -tuple b_δ , and $a_\delta \equiv a$, such that $d(a, a_\delta) \leq \delta$ but there is no b' satisfying $d(b_\delta, b') \leq \varepsilon$ and $ab_\delta \equiv a_\delta b'$.

Let $p_\delta(a, y) = \text{tp}(b_\delta/a)$. By compactness, there exists a complete type $p(a, y)$ which is an accumulation point for these types: for every neighbourhood $U \subseteq S(a)$ and $\delta > 0$ there is $\delta \geq \delta' > 0$ such that $p_{\delta'} \in U$ as well.

By assumption, $p(a, y)$ is isolated; there exists therefore a formula $\varphi(a, y)$ such that $\varphi(a, y)^{p(a, y)} = 0$ and

$$\varphi(a, y) \leq 1/2 \vdash p(a, y^{\varepsilon/2}).$$

Find $\delta > 0$ small enough such that $\varphi(a, y)^{p_\delta} < 1/4$ (using the fact that $[\varphi(a, y) < 1/4]$ is a neighbourhood of $p(a, y)$), and in addition if $d(a, a') \leq \delta$ then

$$\sup_y |\varphi(a, y) - \varphi(a', y)| \leq 1/4.$$

First, as $\varphi(a, b_\delta) \leq 1/2$, we have $b_\delta \models p(a, y^{\varepsilon/2})$, so there exists $b' \models p(a, y)$ such that $d(b', b_\delta) \leq \varepsilon/2$. But we also have $\varphi(a_\delta, b_\delta) \leq 1/2$, so there exists $b'' \models p(a_\delta, y)$ such that $d(b'', b_\delta) \leq \varepsilon/2$. Choose an automorphism of the universal domain sending ab' to $a_\delta b''$, and let b'_δ be the image of b_δ under this automorphism.

Then $d(b'_\delta, b'') = d(b_\delta, b') \leq \varepsilon/2$ so in all we have $d(b_\delta, b'_\delta) \leq \varepsilon$ and $ab_\delta \equiv a_\delta b'_\delta$, contradicting the hypothesis. ■_{2.9}

PROPOSITION 2.10. *Assume that T is ω -categorical. Then $\text{dcl}(\emptyset)$, restricted to the home sort (or to any other one sort), is compact (in the metric topology of the universal domain).*

Proof. Consider the mapping $\theta: \text{dcl}(\emptyset) \rightarrow S_1(T)$ sending $a \in \text{dcl}(\emptyset)$ to $\text{tp}(a)$. Since $a \in \text{dcl}(\emptyset)$ is the unique realisation of $\text{tp}(a)$, θ is an isometric embedding. As T is ω -categorical, $S_1(T)$ is compact in the metric topology (Fact 1.14), and hence totally bounded. Therefore $\text{dcl}(\emptyset)$ is totally bounded. But $\text{dcl}(\emptyset)$ is also complete (any Cauchy sequence in $\text{dcl}(\emptyset)$ converges to an element of the universal domain which must also be in $\text{dcl}(\emptyset)$).

Therefore $\text{dcl}(\emptyset)$ is compact. ■_{2.10}

The next result is an analogue of the fact that if a and b are finite tuples in a model of a classical first order theory and $\text{tp}(ab)$ is isolated, then so is $\text{tp}(a/b)$. The requirement that b be d -finite below is not redundant: as in the introduction, consider the case of the theory of L^p Banach lattices: $\text{tp}(\chi_{[0,1]}, \chi_{[1,2]})$ is isolated (since the theory is ω -categorical) but $\text{tp}(\chi_{[0,1]}/\chi_{[1,2]})$ is not.

PROPOSITION 2.11. *Assume that $\text{tp}(ab/c)$ is isolated and b is d -finite over c . Then $\text{tp}(a/bc)$ is isolated.*

Proof. Let $p(x, y) = \text{tp}(a, b/c)$. Given $\varepsilon > 0$, we want to show that $[p(x^\varepsilon, b)]^\circ \subseteq S(bc)$ is non-empty. Let $\delta = \delta_{a, \varepsilon/2}^{b/c} > 0$. We know that $[p(x^{\varepsilon/2}, y^\delta)]^\circ$

$\subseteq S(c)$ is non-empty, so let $\varphi(x, y, c)$ be a formula such that $\varphi(a, b, c) = 0$ and $[\varphi(x, y, c) < 1/2] \subseteq [p(x^{\varepsilon/2}, y^\delta)]^\circ$.

Assume that a' is such that $\varphi(a', b, c) < 1/2$. Then there are a'', b' such that $\models p(a'', b')$ and $d(a''b', a'b) \leq (\varepsilon/2, \delta)$. By choice of δ there is a''' such that $d(a''', a'') \leq \varepsilon/2$ and $a'''b \equiv_c a''b'$. Therefore $d(a', a''') \leq \varepsilon$ and $\models p(a''', b)$, so $\varphi(x, b, c) < 1/2 \vdash p(x^\varepsilon, b)$, as required. $\blacksquare_{2.11}$

Let us conclude with a few examples.

EXAMPLE 2.12. Let T be a classical first order theory in a language \mathcal{L} . We can view the class of ω -powers of models of T , $\{M^\omega : M \models T\}$, as a continuous elementary class in the following manner. For every formula $\varphi(x_{<n}) \in \mathcal{L}_{\omega, \omega}$ and $m_{<n} \in \omega^n$ let $P_{\varphi, \bar{m}}$ be an n -ary predicate symbol, 2^{m_i} -Lipschitz in the i th argument for each $i < n$. Let \mathcal{L}' be a continuous signature consisting of all these predicate symbols, plus the metric symbol d .

If M is an \mathcal{L} -structure, interpret M^ω as an \mathcal{L}' -structure by

$$P_{\varphi, \bar{m}}(a_{<\omega}, b_{<\omega}, \dots) = \begin{cases} 0 & \text{if } M \models \varphi(a_{m_0}, b_{m_1}, \dots), \\ 1 & \text{otherwise,} \end{cases}$$

$$d(a_{<\omega}, b_{<\omega}) = \inf\{2^{-\ell} : a_{<\ell} = b_{<\ell}\}.$$

Then the class of structures $\{M^\omega : M \models T\}$ is elementary, and let T^ω denote its theory.

It is easy to see that T is ω -categorical if and only if T^ω is, M is ω -saturated if and only if M^ω is approximately so, etc. Also, T is superstable if and only if T^ω is.

The definition of d -finiteness tries to capture the distinction between arbitrary elements $a_{<\omega} \in M^\omega$, which actually code infinite tuples, and ones which only code a finite tuple from M , e.g., ones which are constant from some point onwards. Indeed, let $a_{<\omega} \in M^\omega$ be constant from the n th coordinate onwards, and let $\delta = 2^{-n-1}$. If $a_{<\omega} \equiv a'_{<\omega}$ and $d(a_{<\omega}, a'_{<\omega}) \leq \delta$ then in fact $a_{<\omega} = a'_{<\omega}$. It follows that δ witnesses (quite uniformly, too) that $a_{<\omega}$ is uniformly d -finite over any set of parameters. On the other hand, we leave it to the reader to verify that if $(a_i : i < \omega)$ is a non-constant indiscernible sequence, then $a_{<\omega}$ is not d -finite.

It follows that T^ω admits enough uniformly d -finites: Indeed, let c be any tuple of parameters, and $a = a_{<\omega}$. Define $a^m = a_{<\omega}^m$ by $a_{<m}^m = a_{<m}$, and $a_k^m = a_m$ for $k \geq m$. Then $\text{tp}(a^m/c, a^{<m})$ is uniformly d -finite, and a is interdefinable with $a_{<\omega}$.

EXAMPLE 2.13. *The theory of atomless probability algebras admits enough uniform d -finites. In fact, every finite tuple of vectors is uniformly d -finite over any tuple of parameters.*

Proof. Let us first consider the case of a single event without parameters. Let $a, a' \in \mathcal{A}$, where \mathcal{A} is the unique separable complete atomless probability algebra, and assume that $a \equiv a'$ and $d(a, a') = \mu(a \oplus a') \leq \varepsilon$. Then $\mu(a) = \mu(a')$, which yields $\mu(a \setminus a') = \mu(a' \setminus a)$, and hence there is an automorphism $\sigma \in \text{Aut}(\mathcal{A})$ exchanging $a \setminus a'$ and $a' \setminus a$, and fixing every $b \in \mathcal{A}$ which is disjoint from $a \oplus a'$. In particular $\sigma(a \wedge a') = a \wedge a'$, implying that $\sigma(a) = a'$. On the other hand, for all $b \in \mathcal{A}$ we have $b \oplus \sigma(b) \leq a \oplus a'$, so $d(b, \sigma(b)) \leq \varepsilon$.

We conclude that there is an automorphism sending a to a' while moving nothing by more than ε , so $\text{tp}(a)$ is uniformly *d*-finite.

The same argument can be generalised to a finite tuple of events (generating an algebra with finitely many atoms, and we assume none of the atoms moves much), and replacing probabilities with conditional probabilities also to types over parameters. ■_{2.13}

EXAMPLE 2.14. *The theory of (the closed unit ball of) Hilbert spaces admits enough uniform d-finites. In fact, every finite tuple of events is uniformly d-finite over any tuple of parameters.*

Proof. By moving to orthogonal components we may always assume there are no parameters. Also, we may restrict our consideration to tuples of orthogonal vectors of norm 1. Let $v_{<n}, w_{<n}$ be two such tuples (so $v_{<n} \equiv w_{<n}$) and assume that $d(v_{<n}, w_{<n}) \leq \delta$ is small. Let V, W and U be the spans of \bar{v}, \bar{w} and $\bar{v}\bar{w}$, respectively. Let V^\perp and W^\perp be the orthogonal complements of V and W , respectively, in U . Let $u_{<m}$ be an orthonormal base for V^\perp . Write $u_i = u'_i + u''_i$ where $u'_i \in W^\perp, u''_i \in W$. Then $\|u''_i\|^2 \leq n\delta^2$, so if we assume δ to be small enough the tuple $u'_{<m}$ is close to being orthonormal. Let $\hat{u}_{<m}$ be the result of applying Gram–Schmidt to $u'_{<m}$. Given any $\varepsilon > 0$ we may choose $\delta > 0$ small enough so that $u'_{<m}$ suffices to span W^\perp , so $\hat{u}_{<m}$ is an orthonormal base for W^\perp , and $d(u_{<m}, \hat{u}_{<m}) \leq \varepsilon$.

Let T be the automorphism of U sending the orthonormal base $v_{<n}u_{<m}$ to $w_{<n}\hat{u}_{<m}$. Then $\|Tt - t\| \leq \varepsilon\|t\|$ for every $t \in U$. We can now extend T to an automorphism of any ambient Hilbert space by setting it to be the identity on the orthogonal complement. We conclude there is an automorphism sending \bar{v} to \bar{w} and moving nothing in the unit ball by more than ε . It follows that \bar{v} is uniformly *d*-finite. ■_{2.14}

EXAMPLE 2.15. *Let $1 \leq p < \infty$, and let T be the theory of atomless L^p Banach lattices [BBH]. Then the type of any non-zero function over \emptyset is non-d-finite. It follows that T does not admit enough d-finites.*

Proof. Let us start with the specific example of $f = \chi_{[0,1]}$ in $L^p(\mathbb{R})$. Let $\delta > 0$ and let $g = (1 + \delta)^{-1/p}\chi_{[0,1+\delta]}$, $h = \delta^{-1/p}\chi_{[1,1+\delta]}$. Then $\|g\| = \|f\| = \|h\| = 1$ and all are positive so $f \equiv g$ by quantifier elimination in T . Also,

$\|f - g\|_p \leq \delta/p + \delta \leq 2\delta$. On the other hand, if h' is such that $fh \equiv gh'$ then necessarily $\|h - h'\|_p = 2^{1/p} > 1$. Thus no $\delta > 0$, however small, is good enough for $\varepsilon = 1$. A similar argument shows that no non-zero function is d -finite (over \emptyset). ■_{2.15}

3. Vaught's theorem. Assume in this section that T is complete in a countable language and has enough d -finite elements.

LEMMA 3.1. *Assume T is not ω -categorical. Then there is a d -finite type $p(x) \in S(T)$ which is not isolated.*

Proof. As T is not ω -categorical, there is a type $q(y) \in S(T)$, where y is a finite tuple of variables, which is not isolated. Therefore there exists $\varepsilon > 0$ such that $[q(y^\varepsilon)]$ is nowhere-dense in $S(T)$. Let $b \models q$. As T has enough d -finite elements, there is $a \in \text{dcl}(b)$ which is d -finite, and $\text{tp}(b/a) \vdash d(y, b) < \varepsilon$. We claim that $p = \text{tp}(a)$ is not isolated.

Assume, towards a contradiction, that p is isolated. By a compactness argument there is a formula $\varphi(x, y)$ such that $\varphi(a, b) = 0$ and

$$(*) \quad (\varphi(x, y) \vee \varphi(x, y')) \leq 1/2 \vdash d(y, y') < \varepsilon.$$

By uniform continuity there is also $\delta > 0$ such that

$$(**) \quad d(x, x') \leq \delta \vdash |\varphi(x, y) - \varphi(x', y)| \leq 1/4.$$

As p was assumed to be isolated, there is a formula $\psi(x)$ such that $\psi(a) = 0$ and

$$(***) \quad \psi(x) \leq 1/2 \vdash p(x^\delta).$$

Let $\chi(y) = \inf_x (\varphi(x, y) \vee \psi(x))$, so clearly $\chi(b) = 0$. Also, assume that b' is any tuple such that $\chi(b') < 1/4$. Then there is a' such that $\varphi(a', b') \vee \psi(a') < 1/4$. In particular $\psi(a') < 1/2$, so by (***) there is $a'' \models p$ such that $d(a', a'') \leq \delta$, and up to applying an automorphism we may assume that $a'' = a$. By (**), $\varphi(a, b') < 1/2$, and by (*), $d(b', b) < \varepsilon$, so $b' \models q(y^\varepsilon)$. We have thus shown that $[\chi(y) < 1/4]$ defines a non-empty open subset of $[q(y^\varepsilon)]$, contradicting the choice of ε . ■_{3.1}

THEOREM 3.2 (Vaught's theorem for continuous logic). *Assume T has enough d -finite elements. Then T cannot have precisely two non-isomorphic separable models.*

Proof. Assume for contradiction that T has precisely two separable models. Then T is not ω -categorical, so it has a non-isolated d -finite type $p(x)$, which is therefore omitted in some separable model of T . As T has only countably many separable models it is necessarily small, so it has a separable model which is approximately ω -saturated, and therefore realises p . Thus a separable model of T is approximately ω -saturated if and only if it realises p .

Let $a \models p$. Then every separable model of $T(a)$ is a model of T realising p , and is therefore approximately ω -saturated as a model of T . By Proposition 2.3, a separable model of $T(a)$ is also approximately ω -saturated as a model of $T(a)$. By Fact 1.5, $T(a)$ is ω -categorical, and therefore so must be T , by Proposition 2.9. Contradiction. ■_{3.2}

4. Lachlan’s theorem. In this section we will need the somewhat strong notion of a uniformly d -finite tuple. We adapt the proof of Lachlan’s theorem from [Pil83] to our setting.

DEFINITION 4.1. Let a and b be finite tuples, and A a set. We say that a *semi-isolates* b over A if there exists an A -definable predicate $\varphi(x, y)$ such that $\varphi(a, b) = 0$, and for all c ,

$$d(c, \text{tp}(b/A)) \leq \varphi(a, c).$$

We then say that $\varphi(a, y)$ *witnesses* that a semi-isolates b over A .

LEMMA 4.2. *Let a, b and A be as above, and $p(y) = \text{tp}(b/A)$. Then a semi-isolates b over A if and only if for every $\varepsilon > 0$ there is a formula $\varphi(x, y)$ with parameters in A such that $\varphi(a, b) = 0$ and $\varphi(a, y) \leq 1/2 \vdash p(y^\varepsilon)$.*

Proof. Left to right is easy. For right to left, notice that the right condition can be read as: for every $\varepsilon > 0$ there is a neighbourhood of $\text{tp}(b/a)$ whose pull-back to $S_n(aA)$ (where $|b| = n$) contains the pull-back of $p(y^\varepsilon)$ there. By an Urysohn-lemma-style argument, there is a continuous mapping $\varphi: S_n(a) \rightarrow [0, 1]$ which is 0 at $\text{tp}(b/a)$, and for every $\varepsilon > 0$, the pull-back of $[\varphi \leq \varepsilon]$ to $S_n(aA)$ contains the pull-back of $p(y^\varepsilon)$ there. But this is just a re-statement of the left condition. ■_{4.2}

LEMMA 4.3. *Assume that a semi-isolates b over A and b semi-isolates c over A . Then a semi-isolates c over A .*

Proof. Let $\varphi(a, y)$ witness that a semi-isolates b over A , and let $\psi(b, z)$ witness that b semi-isolates c over A . Let $u_{\psi, y}$ be the inverse continuous uniform continuity modulus, as defined in [BU, Appendix A], of $\psi(y, z)$ with respect to y , and let

$$\varrho(x, z) = \inf_y (u_{\psi, y} \circ \varphi(x, y) + \psi(y, z)).$$

Then $\varrho(a, z)$ witnesses that a semi-isolates c over A . ■_{4.3}

NOTATION 4.4. Let $a^\varepsilon \downarrow_c b$ mean there is a tuple $a' \equiv_c a$ such that $d(a, a') \leq \varepsilon$ and $a' \downarrow_c b$ (i.e., $\text{tp}(a^\varepsilon/bc) \cup \text{tp}(a/c)$ does not divide over c).

It is easy to see that $a \downarrow_c b$ if and only if $a^\varepsilon \downarrow_c b$ for all $\varepsilon > 0$.

LEMMA 4.5. *Assume T is stable. Let a and b be finite tuples, A a set, $p(y) = \text{tp}(b/A)$, $\varepsilon > 0$ and $\varphi(x, y)$ a formula such that:*

- (i) $\varphi(a, b) = 0$.
- (ii) $\varphi(a, y) < 1/2 \vdash p(y^\varepsilon)$.
- (iii) $[p(y^{3\varepsilon})]$ is nowhere dense in $S_y(A)$.

Then $b^\varepsilon \not\downarrow_A a$. Therefore, in particular, if a semi-isolates b over A , but $p(y) = \text{tp}(b/A)$ is not isolated, then $a \not\downarrow_A b$.

Proof. Assume for contradiction that $b^\varepsilon \downarrow a$, so there is c such that $d(b, c) \leq \varepsilon$ and $c \downarrow_A a$. Let

$$X = \{q \in S_n(aA) : q \text{ does not divide over } A\} \subseteq S_n(aA).$$

Then the restriction mapping $\theta : X \rightarrow S_n(A)$ is open by the open mapping theorem.

Let

$$\psi(x, y) = \inf_z \left(\varphi(x, z) + \frac{d(z, y)}{4\varepsilon} \right).$$

Then $\psi(a, y) < 1/2 \vdash p(y^{3\varepsilon})$. Let

$$Y = X \cap [\psi(a, y) < 1/2].$$

Then $Y \subseteq X$ is relatively open, and therefore so is $\theta(Y) \subseteq [p(y^{3\varepsilon})]$. On the other hand, $\psi(a, c) \leq 1/4 < 1/2$, and $c \downarrow_A a$, so $\text{tp}(c/Aa) \in Y \neq \emptyset$.

Thus $[p(y^{3\varepsilon})]$ contains a non-empty open set, contrary to the assumption. $\blacksquare_{4.5}$

LEMMA 4.6. *Assume that T is superstable. Then there are no finite (imaginary) tuples a, a', b' and c' , and $\varepsilon > 0$, such that if $p(z) = \text{tp}(c')$ then:*

- (i) $a \equiv a'$, and they are uniformly d -finite.
- (ii) $\text{tp}(a'b'/a)$ is isolated.
- (iii) $\text{tp}(c'/a'b') \vdash d(z, c') < \varepsilon$.
- (iv) $a' \downarrow c'$.
- (v) $[p(z^{3\varepsilon})]$ is nowhere dense in $S(\emptyset)$.

Proof. Assume that such tuples exist. Construct a sequence $(a_i b_i c_i : i < \omega)$ by induction as follows. Start with $a_0 = a$. At the n th step, given $a_{<n}, b_{<n}$ and $c_{<n}$, choose $a_{n+1} b_n c_n$ such that:

- (i) $a_n a_{n+1} b_n c_n \equiv a a' b' c'$.
- (ii) $a_{n+1} b_n c_n \downarrow_{a_n} a_{<n} b_{<n} c_{<n}$.

Then $a_{n+1} \downarrow c_n$ for all n , and by induction on n one proves that $a_{n+1} c_n \downarrow c_{<n}$. In particular, $(c_n : n < \omega)$ is an independent sequence. Using Lemma 4.3 one can also prove that a ($= a_0$) semi-isolates a_n , and therefore $a_{n+1} b_n$, for all n .

By compactness there is a formula $\psi(xy, z)$ such that $\psi(a'b', c') = 0$ and $\psi(a'b', z) \leq 1/2 \vdash d(c', z) < \varepsilon$. Let $\mathbf{u}_{\psi, xy}$ be its inverse uniform continuity

modulus with respect to the first group of variables. Let $\chi_n(a, xy)$ witness that a semi-isolates $a_{n+1}b_n$. Let

$$\varphi_n(a, z) = \inf_{xy} \mathbf{u}_{\psi, xy}(\chi_n(a, xy)) + \psi(xy, z).$$

Then $\varphi_n(a, c_n) = 0$, and $\varphi_n(a, z) < 1/2 \vdash p(z^\varepsilon)$. By Lemma 4.5, $c_n^\varepsilon \not\downarrow a$ for all n .

Now use the assumption that a is uniformly d -finite, and let $\delta = \delta_{|c'|, \varepsilon}^a > 0$. By superstability there is n such that $a^\delta \downarrow_{c_{<n}} c_{<\omega}$. Therefore there is $a' \equiv a$ such that $d(a, a') \leq \delta$ and $a' \downarrow_{c_{<n}} c_{<\omega}$. By independence of $(c_i : i < \omega)$ we get $a' \downarrow c_n$. By choice of δ there is c'_n such that $d(c_n, c'_n) \leq \varepsilon$ and $c'_n a \equiv c_n a'$. But then $a \downarrow c'_n$, so c'_n witnesses that $c_n^\varepsilon \downarrow a$. Contradiction. $\blacksquare_{4.6}$

THEOREM 4.7. *Assume T is superstable and has enough uniformly d -finite elements. Then T is either ω -categorical or has infinitely many separable models.*

Proof. Assume that T has finitely many countable models, but more than one, and we will prove it cannot be superstable.

Since T has few separable models it is small. Therefore it has a separable and approximately ω -saturated model M . By Proposition 2.8, M is interdefinable with a sequence $(a_i : i < \omega)$ such that $\text{tp}(a_i/a_{<i})$ is uniformly d -finite for all i .

By Lemma 2.2, $a_{<n}$ is uniformly d -finite for all n . It follows that $T(a_{<n})$ is small by Corollary 2.4, and therefore has a prime model. In other words, T has a prime model over $a_{<n}$. Since there are only finitely many possibilities for the prime model over $a_{<n}$, there is n such that the prime model over $a_{<n}$ realises $\text{tp}(a_{<m})$ for all m . Let $a = a_{<n}$.

As T is also assumed to be non- ω -categorical, there is c such that $p(z) = \text{tp}(c)$ is non-isolated. Therefore there is $\varepsilon > 0$ such that $[p(z^{3\varepsilon})]$ is nowhere dense. Since M is approximately ω -saturated over a , we may assume that $c \in M$ and $c \downarrow a$. Since $c \in M = \text{dcl}(a_{<\omega})$, there is $m < \omega$ such that $\text{tp}(c/a_{<m}) \vdash d(z, c) < \varepsilon$. Let $b = a_{<m}$.

Then $ab = a_{<n}a_{<m}$, so by choice of n there are $a'b'$ in the prime model over a realising $\text{tp}(ab)$. Find c' such that $abc \equiv a'b'c'$. Then the existence of $aa'b'c'$ shows that T cannot be superstable, by Lemma 4.6. $\blacksquare_{4.7}$

References

- [Ben05] I. Ben Yaacov, *Uncountable dense categoricity in cats*, J. Symbolic Logic 70 (2005), 829–860.
- [BBH] I. Ben Yaacov, A. Berenstein, and C. W. Henson, *Model-theoretic independence in the Banach lattices $L^p(\mu)$* , submitted.

- [BBHU] I. Ben Yaacov, A. Berenstein, C. W. Henson, and A. Usvyatsov, *Model theory for metric structures*, expanded lecture notes for a workshop given in March/April 2005, Isaac Newton Institute, Univ. of Cambridge, 2005.
- [BU] I. Ben Yaacov and A. Usvyatsov, *Continuous first order logic and local stability*, submitted.
- [Pil83] A. Pillay, *Countable models of stable theories*, Proc. Amer. Math. Soc. 89 (1983), 666–672.

Université de Lyon
Université Lyon 1
Institut Camille Jordan, CNRS, UMR 5208
43 boulevard du 11 novembre 1918
69622 Villeurbanne Cedex, France
<http://math.univ-lyon1.fr/~begnac/>

Mathematics Department
UCLA
Box 951555
Los Angeles, CA 90095-1555, U.S.A.
<http://www.math.ucla.edu/~alexus>

*Received 17 January 2006;
in revised form 2 August 2006*