

ON PERTURBATIONS OF CONTINUOUS STRUCTURES

ITAI BEN YAACOV

ABSTRACT. We give a general framework for the treatment of perturbations of types and structures in continuous logic, allowing to specify which parts of the logic may be perturbed. We prove that separable, elementarily equivalent structures which are approximately \aleph_0 -saturated up to arbitrarily small perturbations are isomorphic up to arbitrarily small perturbations (where the notion of perturbation is part of the data). As a corollary, we obtain a Ryll-Nardzewski style characterisation of complete theories all of whose separable models are isomorphic up to arbitrarily small perturbations.

INTRODUCTION

In this paper we define what we call perturbation systems and study their basic properties. These are objects which formalise the intuitive notion of allowing chosen parts of a metric structure to be perturbed by arbitrarily small amounts.

One motivation for this notion comes from an unpublished result of C. Ward Henson, consisting of a Ryll-Nardzewski style characterisation of complete continuous theories of pure Banach spaces which are separably categorical up to arbitrarily small perturbation of the norm (but not of the underlying linear structure). Seeking a general framework in which such results can be proved, we develop a general formalism for the consideration of metric structures and types up to small perturbations, which gives rise in particular to a notion of categoricity up to perturbation. In Theorem 3.5 we give a general Ryll-Nardzewski style characterisation of complete countable continuous theories which are separably categorical up to arbitrarily small perturbation, where the precise notion of perturbation is part of the given data alongside the theory. One convenient way of specifying a “perturbation system” \mathfrak{p} is via the perturbation distance $d_{\mathfrak{p}}$ between types, where $d_{\mathfrak{p}}(p, q) \in [0, \infty]$ measures by how much a model needs to be perturbed so that a realisation of p may become a realisation of q (and $d_{\mathfrak{p}}(p, q) = \infty$ if this is impossible).

Our criterion for \aleph_0 -categoricity up to perturbation bears considerable resemblance to the one used by Henson, as both criteria compare the standard logic topology on a space of types with an appropriate metric arising from the perturbation system. In Henson’s criterion, the topology is compared directly to the Banach-Mazur perturbation distance d_{BM} on the space of types of *linearly independent* tuples of a Banach space, which he calls S_n^* . In the general case considered in Theorem 3.5 we do not have an analogue of S_n^* , so the comparison must take place on the entire type space. This entails an additional complexity, not present in Henson’s criterion, in that the topology must be compared to an appropriate combination of the perturbation metric $d_{\mathfrak{p}}$ with the standard distance d . A result based on Henson’s criterion appears in a subsequent paper [Ben], where we deal with further complications caused by the

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

Key words and phrases. continuous logic, metric structures, perturbation, categoricity.

Research partially supported by NSF grant DMS-0500172, ANR chaire d’excellence junior THEMOMET (ANR-06-CEXC-007) and by Marie Curie research network ModNet.

The author would like to thank the Isaac Newton Institute and the organisers of the programme on Model Theory and Applications to Algebra and Analysis, during which this work was initiated; C. Ward Henson for many helpful discussions and comments; and Hernando Tellez for a careful reading of the manuscript.

Revision 991 of 5th October 2009.

fact that a Banach space is an unbounded structure whose unit ball is not preserved by a non trivial Banach-Mazur perturbation.

A second motivation comes from some open problems concerning the automorphism group of the separable model of an \aleph_0 -categorical continuous theory. Such problems could be addressed from a model-theoretic point of view as questions concerning the theory T_A (i.e., T with a generic automorphism, or even several non commuting ones). Just as the underlying metric of a continuous structure induces a natural metric on the space of types, it also induces one on its automorphism group, namely the metric of uniform convergence (if the structure is discrete, so are the induced metrics, so they simply do not arise as interesting objects in the classical discrete setting). The model theoretic counterpart of the consideration of small metric neighbourhoods of an automorphism is the consideration of $(M, \sigma) \models T_A$ up to small perturbations of σ . While the present paper does not contain any results in this direction, this did serve well as an example towards the general setting, and in fact was at the origin of the author's interest in perturbations.

A common feature of these two instances is that only part of the structure is allowed to be perturbed while the rest is kept untouched. In the first case, the norm is perturbed while the linear structure is untouched, while in the second it is only the automorphism that we perturb (and not the original structure). Thus a “notion of perturbation” should say what parts of the structure can be perturbed, and in what way. Also, in order to state a Ryll-Nardzewski style result concerning perturbations we need to consider on the one hand perturbations of (separable) models, and on the other perturbations of types.

In Section 1 we compare these two notions (perturbations of structures and of types): requiring them to be compatible yields the notion of a *perturbation radius*. In order to speak of “arbitrarily small perturbation” we need to consider a system of perturbation radii decreasing to the zero perturbation, which with some natural additional properties yields the notion of a *perturbation system*.

In Section 2 we study a variant of the notion of approximate \aleph_0 -saturation which takes into account a perturbation system, and show that separable models which are saturated in this sense are also isomorphic up to small perturbation.

In Section 3 we prove the main result (Theorem 3.5) and discuss various directions in which it may and may not be further generalised (Theorem 3.15 vs. Example 3.11).

In Section 4 we conclude with a few questions concerning perturbations of automorphisms.

Notation is mostly standard. We use a, b, c, \dots to denote members of structures, and use x, y, z, \dots to denote variables. Bar notation is used for (usually finite) tuples, and uppercase letters are used for sets. We also write $\bar{a} \in A$ to say that \bar{a} is a tuple consisting of members of A , i.e., $\bar{a} \in A^n$ where $n = |\bar{a}|$.

We work in the framework of continuous first order logic, as developed in [BU]. Most of the time we work within the context of a fixed continuous theory T in a language \mathcal{L} . We always assume that T is closed under logical consequences. In particular, $|T| = |\mathcal{L}| + \aleph_0$ and T is countable if and only if \mathcal{L} is.

For a general survey of the model theory of metric structures we refer the reader to [BBHU08].

1. PERTURBATIONS

1.1. Perturbation pre-radii. We start by formalising the notion of allowing structures and types to be perturbed “by this much”. We start by defining perturbation pre-radii, which tell us which types can be changed into which:

Definition 1.1. A *perturbation pre-radius* ρ (for a fixed theory T) is a family of closed subsets $\{\rho_n \subseteq S_n(T)^2\}$ containing the diagonals. If $X \subseteq S_n(T)$, then the ρ -neighbourhood around X is defined as:

$$X^\rho = \{q : (\exists p \in X) (p, q) \in \rho_n\}.$$

Notice that if X is closed then so is X^ρ .

Remark 1.2. A predecessor of sorts to this definition exists in José Iovino's notion of a *uniform structure* on the type spaces of a positive bounded theory [Iov99]. Specifically, a uniform structure in Iovino's sense can be generated by vicinities which are given by perturbation pre-radii (although the definition of a pre-radius does not appear in Iovino's work). Specific perturbation systems (see Definition 1.23 below) of importance, such as the Banach-Mazur system, occur as uniform structures in Iovino's work.

We wish to consider mappings which perturb structures: they need not be elementary, and are merely required to respect the perturbation pre-radius.

Definition 1.3. (i) Let ρ be a perturbation pre-radius, $M, N \models T$. A *partial ρ -perturbation* from M into N is a partial mapping $f: M \dashrightarrow N$ such that for every $\bar{a} \in \text{dom}(f)$:

$$\text{tp}^N(f(\bar{a})) \in \text{tp}^M(\bar{a})^\rho.$$

If f is total then it is a ρ -perturbation of M into N . The set of all ρ -perturbations of M into N is denoted $\text{Pert}_\rho(M, N)$.

- (ii) If $f \in \text{Pert}_\rho(M, N)$ is bijective, and $f^{-1} \in \text{Pert}_\rho(N, M)$, we say that $f: M \rightarrow N$ is a *ρ -bi-perturbation*, in symbols $f \in \text{BiPert}_\rho(M, N)$.
- (iii) We say that two perturbation pre-radii ρ and ρ' are *equivalent*, in symbols $\rho \sim \rho'$, if $\text{Pert}_\rho(M, N) = \text{Pert}_{\rho'}(M, N)$ for all $M, N \models T$. We say they are *bi-equivalent*, in symbols $\rho \approx \rho'$, if $\text{BiPert}_\rho(M, N) = \text{BiPert}_{\rho'}(M, N)$ for all $M, N \models T$.

Note that $\rho \sim \rho' \implies \rho \approx \rho'$.

- (iv) If ρ and ρ' are two perturbation pre-radii, we write $\rho \leq \rho'$ to mean that $\rho_n \subseteq \rho'_n$ for all n (i.e., ρ is *stricter* than ρ').

Lemma 1.4. For every perturbation pre-radius ρ there exists a minimal perturbation pre-radius equivalent to ρ , denoted $\langle \rho \rangle$, and a minimal perturbation pre-radius bi-equivalent to ρ , denoted $\llbracket \rho \rrbracket$.

If $\rho = \langle \rho \rangle$ we say that ρ is *reduced*. If $\rho = \llbracket \rho \rrbracket$ we say that ρ is *bi-reduced*.

Proof. One just verifies that $\langle \rho \rangle = \bigcap \{ \rho' : \rho' \sim \rho \}$ and $\llbracket \rho \rrbracket = \bigcap \{ \rho' : \rho' \approx \rho \}$ are perturbation pre-radii which are equivalent and bi-equivalent, respectively, to ρ . ■_{1.4}

Note that $\llbracket \rho \rrbracket \leq \langle \rho \rangle \leq \rho$, so if ρ is bi-reduced it is reduced.

Definition 1.5. Let ρ, ρ' be perturbation pre-radii. We define their *composition* as the pre-radius $\rho' \circ \rho$ defined by:

$$(\rho' \circ \rho)_n = \{ (p, q) : \exists r (p, r) \in \rho_n \text{ and } (r, q) \in \rho'_n \}.$$

It may be convenient to think of a perturbation pre-radius as the graphs of a family multi-valued mappings $\rho_n: S_n(T) \rightarrow S_n(T)$. In this case, our notion of composition above is indeed the composition of multi-valued mappings.

Notice that we also obtain a composition mapping for perturbations:

$$\circ: \text{Pert}_\rho(M, N) \times \text{Pert}_{\rho'}(N, L) \rightarrow \text{Pert}_{\rho \circ \rho'}(M, L).$$

The minimal perturbation pre-radius is $\text{id} = \{ \text{id}_n : n \in \mathbb{N} \}$, where id_n is the diagonal of $S_n(T)$, i.e., the graph of the identity mapping. It is bi-reduced, $\rho \circ \text{id} = \text{id} \circ \rho = \rho$ for all ρ , and an id-perturbation is synonymous with an elementary embedding.

1.2. Perturbation radii. A perturbation pre-radius imposes a family of conditions saying which types may be perturbed to which. We may further require these conditions to be compatible with one another:

Definition 1.6. A *perturbation radius* is a pre-radius ρ satisfying that for any two types $(p, q) \in \rho_n$ there exist models M and N and a ρ -perturbation $f: M \rightarrow N$ sending some realisation of p to a realisation of q .

Notice that the identity perturbation pre-radius is a perturbation radius.

We now try to break down the notion of a perturbation radius into several technical properties and see what each of them means.

For our purposes, a (*uniform*) *continuity modulus* is a mapping $\delta: (0, \infty) \rightarrow (0, \infty)$ which is increasing and left-continuous. (In other words, this is a mapping satisfying $\delta(\varepsilon) = \sup_{\varepsilon' < \varepsilon} \delta(\varepsilon')$. This additional property does not play any role at this stage, but is harmless to assume.) A mapping between metric spaces $f: (X, d) \rightarrow (X', d')$ *respects* δ if for all $\varepsilon > 0$ and all $x, y \in X$:

$$d(x, y) < \delta(\varepsilon) \implies d'(f(x), f(y)) \leq \varepsilon.$$

Such a mapping f is uniformly continuous if and only if it respects some uniform continuity modulus.

Definition 1.7. Let ρ be a perturbation pre-radius.

(i) We say that ρ *respects equality* if

$$[x = y]^\rho = [x = y].$$

(I.e., if $(p, q) \in \rho_2$, $p \models x = y$, then $q \models x = y$ as well).

(ii) We say that ρ *respects a continuity modulus* δ if every ρ -perturbation does.

(iii) We say that ρ is *uniformly continuous* if it respects some continuity modulus δ .

(iv) We say that ρ respects a continuity modulus δ *trivially* if for all $\varepsilon > 0$:

$$[d(x, y) < \delta(\varepsilon)]^\rho \subseteq [d(x, y) \leq \varepsilon].$$

Lemma 1.8. *A perturbation pre-radius ρ respects equality if and only if there exists a continuity modulus δ which ρ respects trivially.*

Proof. Assume first that ρ respects δ trivially. For every $\varepsilon > 0$ we have $\delta(\varepsilon) > 0$, whereby $[x = y] \subseteq [d(x, y) < \delta(\varepsilon)]$ and thus $[x = y]^\rho \subseteq [d(x, y) \leq \varepsilon]$. Therefore $[x = y]^\rho \subseteq \bigcap_{\varepsilon > 0} [d(x, y) \leq \varepsilon] = [x = y]$. As the other inclusion is always true, we obtain equality.

Conversely, assume that ρ respects no δ trivially. Then there exists some $\varepsilon > 0$ such that for all $\delta > 0$ there is some pair $(p_\delta, q_\delta) \in \rho_2$ such that $p_\delta \in [d(x, y) < \delta]$ and $q_\delta \in [d(x, y) > \varepsilon]$. Since $S_2(T)^2$ is compact this sequence has an accumulation point (p, q) as δ goes to 0. Since ρ_2 is closed we have $(p, q) \in \rho_2$, and clearly $(p, q) \in [x = y] \times [d(x, y) \geq \varepsilon]$ as well, so $[x = y]^\rho \neq [x = y]$. ■_{1.8}

Lemma 1.9. *Let ρ be a perturbation pre-radius and δ a continuity modulus. If ρ respects δ trivially then it respects δ . Conversely, if ρ respects δ then $\langle \rho \rangle$ respects δ trivially.*

In particular, if ρ is reduced then it respects δ if and only if it respects it trivially.

Proof. The first statement is straightforward. For the converse, let:

$$\begin{aligned} X_\delta &= \{(p, q) \in S_2(T) : (\forall \varepsilon > 0)(d(x, y)^p < \delta(\varepsilon) \rightarrow d(x, y)^q \leq \varepsilon)\} \\ &= \bigcap_{\varepsilon > 0} ([d(x, y) \geq \delta(\varepsilon)] \times S_2(T) \cup S_2(T) \times [d(x, y) \leq \varepsilon]). \end{aligned}$$

Let ρ' be obtained from ρ by replacing ρ_2 with $\rho_2 \cap X_\delta$. Notice that the identity mapping of any model of T is a ρ -perturbation and must therefore respect δ , so X_δ contains the diagonal and ρ' is a perturbation pre-radius. Clearly $\rho' \sim \rho$, so $\langle \rho' \rangle = \langle \rho \rangle$, and ρ' respects δ trivially, whereby so does $\langle \rho \rangle$. ■_{1.9}

Note that if ρ is uniformly continuous, $M, N \models T$, and $A \subseteq M$, then any partial ρ -perturbation $f: A \rightarrow M$ is uniformly continuous. It therefore extends uniquely to a mapping $\bar{f}: \bar{A} \rightarrow M$. As ρ is given by closed sets, the completion \bar{f} is also a ρ -perturbation.

Lemma 1.10. *A perturbation pre-radius ρ is a perturbation radius if and only if it is uniformly continuous and reduced.*

Proof. Left to right is easy. For right to left, consider the family $\mathcal{F} = \{(f, M, N): M, N \models T, f \in \text{Pert}_\rho(M, N)\}$. Since ρ is uniformly continuous, ultra-products of families of triplets in \mathcal{F} exist, and since ρ as a perturbation pre-radius consists of closed sets, \mathcal{F} is closed under ultra-products. Define $\rho_{\mathcal{F}}$ by:

$$\rho_{\mathcal{F},n} = \{(\text{tp}(\bar{a}), \text{tp}(\bar{b})) : (f, M, N) \in \mathcal{F}, \bar{a} \in M^n, \bar{b} = f(\bar{a})\}.$$

Since \mathcal{F} is closed under ultra-products and contains all the identity mappings, $\rho_{\mathcal{F}}$ is a perturbation pre-radius. It clearly satisfies $\rho_{\mathcal{F}} \leq \rho$, $\rho_{\mathcal{F}} \sim \rho$, and as ρ is reduced we conclude that $\rho_{\mathcal{F}} = \rho$.

On the other hand, it is clear from the construction of $\rho_{\mathcal{F}}$ that it is a perturbation radius. ■_{1.10}

Proposition 1.11. *A perturbation pre-radius ρ is equivalent to a perturbation radius if and only if it is uniformly continuous, in which case $\langle \rho \rangle$ is the unique perturbation radius equivalent to ρ .*

Proof. Immediate from Lemma 1.10. ■_{1.11}

Recall that if $p(\bar{x}, \bar{y})$ is a partial type, then the property $\exists \bar{y} p(\bar{x}, \bar{y})$ (where the existential quantifier varies over a sufficiently saturated elementary extension) is also definable by a partial type.

Definition 1.12. A perturbation pre-radius ρ *respects the existential quantifier* \exists if for every partial type $p(\bar{x}, \bar{y})$:

$$[\exists \bar{y} p(\bar{x}, \bar{y})]^\rho = [\exists \bar{y} p^\rho(\bar{x}, \bar{y})].$$

Lemma 1.13. *A perturbation pre-radius ρ respects \exists if and only if for every two sufficiently saturated models $M, N \models T$, tuples $\bar{a} \in M^n$, $\bar{b} \in N^n$, and $c \in M$:*

$$\text{tp}(\bar{b}) \in \text{tp}(\bar{a})^\rho \iff (\exists d \in N)(\text{tp}(\bar{b}d) \in \text{tp}(\bar{a}c)^\rho).$$

Proof. Easy. ■_{1.13}

If σ is an n -permutation, it acts on $S_n(T)$ by $\sigma^*(p(x_{<n})) = p(x_{\sigma^{-1}(0)}, \dots, x_{\sigma^{-1}(n-1)})$ (so $\sigma^*(\text{tp}(a_{<n})) = \text{tp}(a_{\sigma(0)}, \dots, a_{\sigma(n-1)})$).

Definition 1.14. A perturbation pre-radius ρ is *permutation-invariant* if for every n , and every permutation σ on n elements, ρ_n is invariant under the action of σ . In other words, for every $p, q \in S_n(T)$:

$$(p, q) \in \rho_n \iff (\sigma^*(p), \sigma^*(q)) \in \rho_n.$$

Proposition 1.15. *Let ρ be a perturbation pre-radius. Then the following are equivalent:*

- (i) ρ is a perturbation radius.
- (ii) ρ respects $=, \exists$ and is permutation-invariant.
- (iii) Whenever $M, N \models T$, $\bar{a} \in M^n$, $\bar{b} \in N^n$, and $\text{tp}(\bar{b}) \in \text{tp}(\bar{a})^\rho$, there exist an elementary extension $N' \succeq N$ and a ρ -perturbation $f: M \rightarrow N'$ sending \bar{a} to \bar{b} .

Proof. (i) \implies (ii). Straightforward.

(ii) \implies (iii). Let $\bar{a} \in M$ and $\bar{b} \in N$ be such that $\bar{b} \models \text{tp}(\bar{a})^\rho$. Let $N' \succeq N$ realise every type of finite tuples over finite tuples in N .

Since ρ respects \exists and N' is sufficiently saturated, for every $\bar{c} \in M$ there is $\bar{d} \in N'$ such that $\bar{b}\bar{d} \models \text{tp}(\bar{a}\bar{c})^\rho$. Since ρ respects equality, $(a_i \mapsto b_i) \cup (c_j \mapsto d_j)$ is a well-defined mapping, call it $f: M \dashrightarrow N'$. Since ρ respects \exists and is permutation-invariant, f is a partial ρ -perturbation.

Let I be the family of all partial ρ -perturbations $f: M \dashrightarrow N'$ where $\text{dom}(f)$ is finite containing \bar{a} , and $f(\bar{a}) = \bar{b}$. For any tuple $\bar{c} \in M$, let $J_{\bar{c}} = \{f \in I: \bar{c} \subseteq \text{dom}(f)\}$, and let $\mathcal{F} \subseteq \mathcal{P}(I)$ be the filter generated by the $J_{\bar{c}}$. By the argument above \mathcal{F} is a proper filter, and therefore extends to an ultra-filter \mathcal{U} .

Let $N'' = N'^{\mathcal{U}}$. Let $g: M \rightarrow N''$ be given by $g = \prod_{f \in I} f / \mathcal{U}$. In other words, for every $c \in M$ we define $g(c) \in N''$ to be $[c_f: f \in I] \in N''$, where $c_f = f(c)$ if $c \in \text{dom}(f)$: since $J_{\{c\}}$ is a large set, we

need not care about c_f for other values of f . Identifying N' with its diagonal embedding in N'' we have $N \preceq N' \preceq N''$, and clearly $g(\bar{a}) = \bar{b}$.

Finally, for every (finite) tuple $\bar{c} \in M$ we have $g(\bar{c}) = [\bar{c}_f : f \in I]$, where $\bar{c}_f = f(\bar{c})$ for every f in the large set $J_{\bar{c}}$. Since $\text{tp}(\bar{c}_f) \in \text{tp}(\bar{c})^\rho$ for all $f \in J_{\bar{c}}$, and $\text{tp}(\bar{c})^\rho$ is a closed set, we must have $\text{tp}(g(\bar{c})) \in \text{tp}(\bar{c})^\rho$.

We conclude that $g: M \rightarrow N''$ is a ρ -perturbation as required.

(iii) \implies (i). Clear. ■_{1.15}

It follows that the composition of perturbation radii is again one:

Lemma 1.16. *If ρ, ρ' are perturbation radii then $\rho' \circ \rho$ is a perturbation radius as well.*

Proof. Assume that $q \in p^{\rho' \circ \rho}$. Then there is a type $r \in p^\rho$ such that $q \in r^{\rho'}$. Let $\bar{a} \models p$ in M . Then there is a model N and ρ -perturbation $f: M \rightarrow N$ such that $f(\bar{a}) \models r$, and a ρ' -perturbation $g: N \rightarrow L$ such that $g \circ f(\bar{a}) \models q$. ■_{1.16}

Recall from [Ben03] that the *type-space functor* of T is a contra-variant functor from \mathbb{N} to topological spaces, sending an object $n \in \mathbb{N}$ to $S_n(T)$, and a mapping $\sigma: n \rightarrow m$ to the mapping

$$\begin{aligned} \sigma^*: S_m(T) &\rightarrow S_n(T) \\ \text{tp}(a_i : i < m) &\mapsto \text{tp}(a_{\sigma(i)} : i < n). \end{aligned}$$

We obtain the following elegant characterisation of perturbation radii:

Lemma 1.17. *A perturbation pre-radius ρ is a perturbation radius if and only if for every $n, m \in \mathbb{N}$ and mapping $\sigma: n \rightarrow m$, the induced mapping $\sigma^*: S_m(T) \rightarrow S_n(T)$ satisfies that for all $p \in S_m(T)$:*

$$\sigma^*(p^\rho) = \sigma^*(p)^\rho.$$

Viewing ρ as the family of graphs of multi-valued mappings, we could write this property more simply as $\sigma^ \circ \rho_m = \rho_n \circ \sigma^*$. Thus a perturbation pre-radius is a perturbation radius if and only if it commutes with the type-space functor structure on $\{S_n(T) : n \in \mathbb{N}\}$.*

Proof. Assume that ρ is a perturbation radius, and let $\sigma: n \rightarrow m$ be a mapping. Let $p \in S_m(T)$, $q \in S_n(T)$, and let $a_{<m} \in M$ realise p . Then each of $q \in \sigma^*(p^\rho)$ and $q \in \sigma^*(p)^\rho$ is equivalent to the existence of a ρ -perturbation $g: M \rightarrow N$ such that $q = \text{tp}(g(a_{\sigma(i)}): i < n)$.

Conversely, assume that $\sigma^*(p^\rho) = \sigma^*(p)^\rho$ for all $\sigma: n \rightarrow m$ and $p \in S_m(T)$. When restricted to the special case where $\sigma: 2 \rightarrow 1$ is the unique such mapping, this is equivalent to ρ preserving equality; when restricted to the family of inclusions $n \hookrightarrow n+1$, this is equivalent to ρ preserving \exists ; and when restricted to the permutations of the natural numbers, this is equivalent to ρ being permutation-invariant. Therefore ρ is a perturbation radius by Proposition 1.15. ■_{1.17}

Definition 1.18. We say that a perturbation radius (or pre-radius) is *symmetric* if $q \in p^\rho \iff p \in q^\rho$.

Lemma 1.19. *Assume that ρ is a symmetric perturbation radius, and let $f \in \text{Pert}_\rho(M, N)$. Then there exist elementary extensions $M' \succeq M$, $N' \succeq N$, and a bi-perturbation $f' \in \text{BiPert}(M', N')$ extending f .*

Proof. Since ρ is symmetric then $f^{-1}: f(M) \rightarrow M$ is a partial ρ -perturbation, and since ρ is a perturbation radius, we may extend f^{-1} to a ρ -perturbation $g: N \rightarrow M' \succeq M$. Proceeding this way we may thus construct two elementary chains $(M_i : i \in \mathbb{N})$ and $(N_i : i \in \mathbb{N})$ such that $M_0 = M$, $N_0 = N$, and two sequences of ρ -perturbations $f_i: M_i \rightarrow N_i$ and $g_i: N_i \rightarrow M_{i+1}$ such that $f_0 = f$, $g_i \circ f_i = \text{id}_{M_i}$, and $f_{i+1} \circ g_i = \text{id}_{N_i}$. Then at the limit we obtain $M_\omega \succeq M$ and $N_\omega \succeq N$, ρ -perturbations $f_\omega: M_\omega \rightarrow N_\omega$ and $g_\omega: N_\omega \rightarrow M_\omega$ such that $g_\omega = f_\omega^{-1}$. Thus every ρ -perturbation can be extended by a back-and-forth argument to a ρ -bi-perturbation $f_\omega \in \text{BiPert}_\rho(M_\omega, N_\omega)$. ■_{1.19}

Lemma 1.20. *A perturbation pre-radius ρ is a symmetric perturbation radius if and only if it is uniformly continuous and bi-reduced.*

Proof. If ρ is bi-reduced then it is reduced and symmetric, so one direction is by Lemma 1.10. For the other, assume ρ is a symmetric perturbation radius. Let $f \in \text{Pert}_\rho(M, N)$, and let $f' \in \text{BiPert}_\rho(M', N')$ extend it as in Lemma 1.19. Then $f' \in \text{BiPert}_{\llbracket \rho \rrbracket}(M', N')$ by definition, whereby $f' \in \text{Pert}_{\llbracket \rho \rrbracket}(M', N')$ and $f \in \text{Pert}_{\llbracket \rho \rrbracket}(M, N)$. Therefore $\llbracket \rho \rrbracket \sim \rho$, and as both are reduced they are equal. $\blacksquare_{1.20}$

Proposition 1.21. *A perturbation pre-radius ρ is bi-equivalent to a symmetric perturbation radius if and only if it is uniformly continuous, in which case $\llbracket \rho \rrbracket$ is the unique symmetric perturbation radius bi-equivalent to ρ .*

Proof. Immediate from Lemma 1.20. $\blacksquare_{1.21}$

Finally, we may find the following observation useful:

Lemma 1.22. *Let ρ^i be symmetric uniformly continuous perturbation pre-radii such that $\rho^1 \circ \rho^0 \leq \rho^2$. Then $\llbracket \rho^1 \rrbracket \circ \llbracket \rho^0 \rrbracket \leq \llbracket \rho^2 \rrbracket$.*

Proof. Let $M_0, M_1 \models T$ and $f \in \text{Pert}_{\llbracket \rho^1 \rrbracket \circ \llbracket \rho^0 \rrbracket}(M_0, M_1)$. Then every restriction of f to a finite set can be decomposed by definition into a partial $\llbracket \rho^0 \rrbracket$ -perturbation followed by a partial $\llbracket \rho^1 \rrbracket$ -perturbation. We can glue these together by an ultra-product argument to obtain $M'_i \succeq M_i$ for $i < 2$ and $M'_2 \models T$, such that f extends to $f': M'_0 \rightarrow M'_1$, which in turn decomposes into a $\llbracket \rho^0 \rrbracket$ -perturbation $g': M'_0 \rightarrow M'_2$ followed by a $\llbracket \rho^1 \rrbracket$ -perturbation $h': M'_2 \rightarrow M'_1$.

Since $\llbracket \rho^i \rrbracket$ are symmetric perturbation radii, we may use a back-and-forth argument as in the proof of Lemma 1.19 to construct extensions $M''_i \succeq M'_i$ for $i < 3$ and $g'' \in \text{BiPert}_{\llbracket \rho^0 \rrbracket}(M''_0, M''_2)$, $h'' \in \text{BiPert}_{\llbracket \rho^1 \rrbracket}(M''_2, M''_1)$. It follows that $f'' = h'' \circ g'' \in \text{Pert}_{\rho^2}(M''_0, M''_1)$ is bijective. Since ρ^2 is assumed to be symmetric, f'' is a ρ^2 -bi-perturbation and therefore a $\llbracket \rho^2 \rrbracket$ -perturbation.

This shows that $\llbracket \rho^1 \rrbracket \circ \llbracket \rho^0 \rrbracket \leq \llbracket \rho^2 \rrbracket$. $\blacksquare_{1.22}$

1.3. Perturbation systems. A single perturbation radius gives us certain leverage at perturbing types. But our goal is not to study perturbations by a single perturbation radius, but rather by “arbitrarily small” perturbation radii, where the notion of a small perturbation radius depends on the context. We formalise this through the notion of a perturbation system:

Let \mathfrak{R}^0 denote the family of perturbation pre-radii, and \mathfrak{R} denote the family of perturbation radii.

Definition 1.23. A *perturbation pre-system* is a mapping $\mathfrak{p}: \mathbb{R}^+ \rightarrow \mathfrak{R}^0$ satisfying:

- (i) *Downward continuity:* If $\varepsilon_n \searrow \varepsilon$ then $\mathfrak{p}(\varepsilon) = \bigcap \mathfrak{p}(\varepsilon_n)$.
- (ii) *Symmetry:* $\mathfrak{p}(\varepsilon)$ is symmetric for all ε .
- (iii) *Triangle inequality:* $\mathfrak{p}(\varepsilon) \circ \mathfrak{p}(\varepsilon') \leq \mathfrak{p}(\varepsilon + \varepsilon')$.
- (iv) *Strictness:* $\mathfrak{p}(0) = \text{id}$.

If in addition its range lies in \mathfrak{R} , then $\mathfrak{p}: \mathbb{R}^+ \rightarrow \mathfrak{R}$ is a *perturbation system*.

Given a perturbation (pre-)system \mathfrak{p} , we may define the *perturbation distance* between two types $p, q \in \mathbb{S}_n(T)$ as:

$$d_{\mathfrak{p},n}(p, q) = d_{\mathfrak{p}}(p, q) = \inf\{\varepsilon \geq 0: (p, q) \in \mathfrak{p}_n(\varepsilon)\}.$$

Notice that by strictness and the triangle inequality this is indeed a $[0, \infty]$ -valued metric, where infinite distance means that neither type can be perturbed into the other.

Lemma 1.24. *Let \mathfrak{p} be a perturbation pre-system. Then the family of metrics $(d_{\mathfrak{p},n}: n \in \mathbb{N})$ has the following properties:*

- (i) *For every n , the set $\{(p, q, \varepsilon) \in \mathbb{S}_n(T)^2 \times \mathbb{R}^+: d_{\mathfrak{p},n}(p, q) \leq \varepsilon\}$ is closed.*

- (ii) If \mathbf{p} is a perturbation system, then for every $n, m \in \mathbb{N}$ and mapping $\sigma: n \rightarrow m$, the induced mapping $\sigma^*: S_m(T) \rightarrow S_n(T)$ satisfies for all $p \in S_m(T)$ and $q \in S_n(T)$:

$$d_{\mathbf{p},m}(p, (\sigma^*)^{-1}(q)) = d_{\mathbf{p},n}(\sigma^*(p), q).$$

(Here we follow the convention that $d_{\mathbf{p},m}(p, \emptyset) = \inf \emptyset = \infty$.)

Conversely, given a family of metrics with values in $[0, \infty]$ satisfying the first property, and defining $\mathbf{p}_n(\varepsilon) = \{(p, q) \in S_n(T)^2 : d_{\mathbf{p},n}(p, q) \leq \varepsilon\}$, we obtain that \mathbf{p} is a perturbation pre-system, and it is a perturbation system if and only if the second property is satisfied as well.

Proof. This is merely a reformulation:

- Symmetry, triangle inequality and strictness correspond to each $d_{\mathbf{p},n}$ being a metric;
- Downward continuity corresponds to the set $\{(p, q, \varepsilon) \in S_n(T)^2 \times \mathbb{R}^+ : d_{\mathbf{p},n}(p, q) \leq \varepsilon\}$ being closed; and
- Each of the $\mathbf{p}(\varepsilon)$ being a perturbation radius corresponds to $d_{\mathbf{p},m}(p, (f^*)^{-1}(q)) = d_{\mathbf{p},n}(f^*(p), q)$, by Lemma 1.17. ■_{1.24}

We say that two perturbation systems \mathbf{p} and \mathbf{p}' are *equivalent* if the perturbation metrics $d_{\mathbf{p}}$ and $d_{\mathbf{p}'}$ are uniformly equivalent on each $S_n(T)$.

We say that a perturbation pre-system \mathbf{p} *respects equality* if $\mathbf{p}(\varepsilon)$ does for all $\varepsilon > 0$. In this case, by Proposition 1.21 we can define $\llbracket \mathbf{p}(\varepsilon) \rrbracket = \llbracket \mathbf{p} \rrbracket(\varepsilon)$ to be the symmetric perturbation radius generated by $\mathbf{p}(\varepsilon)$. By Lemma 1.22, $\llbracket \mathbf{p} \rrbracket$ satisfies the triangle inequality. One can verify that $\llbracket \mathbf{p} \rrbracket$ satisfies downward continuity, and it is clearly symmetric and strict, so it is a perturbation system. As expected, we call $\llbracket \mathbf{p} \rrbracket$ the perturbation system *generated* by \mathbf{p} .

1.4. A few natural examples (and a non-example). If \mathcal{L} consists of finitely many predicate symbols, a natural perturbation system for \mathcal{L} is the one allowing to perturb all symbols by “a little”. In order to construct it we first define a perturbation pre-system \mathbf{p} by letting $\mathbf{p}(\varepsilon)$ be the (symmetric) perturbation pre-radius allowing the distance symbol d to change by a multiplicative factor of $e^{\pm\varepsilon}$, and every other symbol to change by $\pm\varepsilon$. Then \mathbf{p} respects equality, and thus generates a perturbation system $\llbracket \mathbf{p} \rrbracket$. Similarly, if \mathcal{L} is an expansion of \mathcal{L}_0 by finitely many symbols, we might want to require that all symbols of \mathcal{L}_0 be preserved precisely, while allowing the new symbols to be perturbed as in the previous case.

A particularly interesting example of the latter kind is the case of adding a generic automorphism to a stable continuous theory. Consider for example the case of infinite dimensional Hilbert spaces: If $\sigma, \sigma' \in U(H)$, then (H, σ') is obtained from (H, σ) by a small perturbation of the automorphism (which keeps the underlying Hilbert space unmodified) if and only if the operator norm $\|\sigma - \sigma'\|$ is small. Thus the notion of perturbation brings into the realm of model theory the uniform convergence topology on automorphism groups of structures. We shall say a little more about this in the last section.

In case of a classical (i.e., discrete) first order theory T in a finite language, there are no non trivial perturbation systems. Indeed, let \mathbf{p} be a perturbation system, and let P be an n -ary predicate symbol. Let

$$X_P = ([P(\bar{x})] \times [\neg P(\bar{y})]) \cup ([\neg P(\bar{x})] \times [P(\bar{y})]) \subseteq S_n(T)^2.$$

Then $X_P \cap \mathbf{p}_n(0) = \emptyset$, but X is compact, so there is $\varepsilon_P > 0$ such that $X_P \cap \mathbf{p}_n(\varepsilon_P) = \emptyset$. Replacing function symbols with their graphs we may assume the language is purely relational, and as we assumed the language to be finite we have can define $\varepsilon_0 = \min\{\varepsilon_P : P \in \mathcal{L}\} > 0$. By the construction every $\mathbf{p}(\varepsilon_0)$ -perturbation is an elementary mapping, that is to say that a small enough perturbation, according to \mathbf{p} , is not a perturbation at all. Another way of saying this is that \mathbf{p} is equivalent to the identity perturbation system. In short, structures in a finite discrete language cannot really be perturbed. The same argument holds if we have a pair of languages $\mathcal{L}_0 \subseteq \mathcal{L}$, where we only allow to perturb symbols in $\mathcal{L} \setminus \mathcal{L}_0$ which are finite in number.

Thus, the notion of perturbation is a new feature of continuous logic which essentially does not exist in discrete logic.

This last statement is of course not 100% correct, as there was a finiteness assumption. Indeed, let $\mathcal{L} = \{E_i : i \in \mathbb{N}\}$ and let T be the theory saying that each E_i is an equivalence relation with two equivalence classes, and every intersection of finitely many equivalence classes of distinct E_i 's is infinite. This is a classical example of a theory which is not \aleph_0 -categorical, but every restriction of T to a finite sub-language is. For $\varepsilon > 0$, let $\mathbf{p}(\varepsilon)$ be the symmetric perturbation radius generated by requiring E_i to be fixed for all $i < 1/\varepsilon$, and $\mathbf{p}(0) = \text{id}$. Then \mathbf{p} is a perturbation system, and ‘‘a model of T up to a small \mathbf{p} -perturbation’’ is the same as ‘‘a model of T restricted to a finite sub-language’’. Thus T is \mathbf{p} - \aleph_0 -categorical in the sense of Section 2 below.

2. SATURATION UP TO PERTURBATION

In this section, \mathbf{p} denotes a perturbation system for a theory T .

Notation 2.1. If $p(x)$ is any partial type and $\varepsilon \geq 0$ then $p(x^\varepsilon)$ denotes the partial type

$$\exists x' (p(x') \wedge d(x, x') \leq \varepsilon).$$

We define $p(x^\varepsilon, y^\delta, \dots)$ similarly. We follow the convention that the metric on finite tuples is the supremum metric, so if \bar{x} is a (finite) tuple of variables then $p(\bar{x}^\varepsilon)$ means $p(x_0^\varepsilon, x_1^\varepsilon, \dots)$.

This notation can (and will) be used in conjunction with previous notation. If $p(\bar{x})$ is a partial type and ρ a perturbation radius then $p^\rho(\bar{x})$ is also a partial type, so we can make sense of $p^\rho(\bar{x}^\varepsilon)$: $\bar{a} \models p^\rho(\bar{x}^\varepsilon)$ if and only if there exists a tuple \bar{b} (in a sufficiently saturated model containing \bar{a}) such that $d(\bar{a}, \bar{b}) \leq \varepsilon$ and $\bar{b} \models p^\rho$. Similarly, $\bar{a} \models p(\bar{x}^\varepsilon)^\rho$ if and only if there are \bar{b} and \bar{c} such that $\bar{b} \models p$, $d(\bar{b}, \bar{c}) \leq \varepsilon$, and $\text{tp}(\bar{a}) \in \text{tp}(\bar{c})^\varepsilon$. The difference between the two examples is that in the first we first perturb p and then allow the realisation to move a little, while in the second we do it the other way around. Since ρ is uniformly continuous, this does not make much difference, as for all $\varepsilon > 0$ and $\delta = \delta_\rho(\varepsilon) > 0$:

$$[p^\rho(\bar{x}^\delta)] \subseteq [p(\bar{x}^\varepsilon)]^\rho, \quad [p(\bar{x}^\delta)]^\rho \subseteq [p^\rho(\bar{x}^\varepsilon)].$$

Definition 2.2. A structure M is \mathbf{p} -approximately \aleph_0 -saturated if for every finite tuple $\bar{a} \in M$, type $p(x, \bar{a}) \in S_1(\bar{a})$ and $\varepsilon > 0$, the partial type $p^{\mathbf{p}(\varepsilon)}(x^\varepsilon, \bar{a}^\varepsilon)$ is realised in M .

Notice that when $b \in M$ realises $p^{\mathbf{p}(\varepsilon)}(x^\varepsilon, \bar{a}^\varepsilon)$, the witnesses may possibly be *outside* M . In other words, $M \models p^{\mathbf{p}(\varepsilon)}(b^\varepsilon, \bar{a}^\varepsilon)$ only means that there exist b', \bar{a}' in some elementary extension of M such that $d(b, b'), d(\bar{a}, \bar{a}') \leq \varepsilon$ and $\models p^{\mathbf{p}(\varepsilon)}(b', \bar{a}')$.

Lemma 2.3. *The definition of \mathbf{p} -approximate \aleph_0 -saturation, which was given in terms of approximate realisation of 1-types, implies the same property for n -types, for any natural n .*

Proof. Let M be \mathbf{p} -approximately \aleph_0 -saturated. We proceed by induction on n . For $n = 0$ there is nothing to prove, so we assume for n and prove for $n + 1$.

So let $p(x_{\leq n}, \bar{a}) \in S_{n+1}(\bar{a})$ for some finite tuple $\bar{a} \in M$, where $p(x_{\leq n}, \bar{y})$ is a complete type without parameters, and let $\varepsilon > 0$. We need to find in M a realisation for $p^{\mathbf{p}(\varepsilon)}(x_{\leq n}^\varepsilon, \bar{a}^\varepsilon)$.

First, find $\delta > 0$ such that $[d(x, y) \leq \delta]^{\mathbf{p}(\varepsilon)} \subseteq [d(x, y) \leq \varepsilon/2]$, so in particular, $\delta \leq \varepsilon/2$. Let $q(x_{< n}, \bar{y}) = p(x_{\leq n}, \bar{y}) \upharpoonright_{(x_{< n}, \bar{y})}$. By the induction hypothesis we can realise $q^{\mathbf{p}(\delta)}(x_{< n}^\delta, \bar{a}^\delta)$ in M . In other words, we can find $b_{< n} \in M$ and $b'_{< n}, \bar{a}'$ possibly outside M such that $d(b'_{< n}, b_{< n}) \leq \delta$ and $\models q^{\mathbf{p}(\varepsilon/2)}(b'_{< n}, \bar{a}')$. Since $\mathbf{p}(\varepsilon/2)$ is a perturbation radius we can find b'_n (still, possibly outside M) such that $\models p^{\mathbf{p}(\varepsilon/2)}(b'_{< n}, \bar{a}')$. Thus in particular:

$$\models p^{\mathbf{p}(\varepsilon/2)}(b'_n, b_{< n}^\delta, \bar{a}^\delta).$$

Let $r(x) = \text{tp}(b'_n/b_{<n}, \bar{a})$. Using \mathfrak{p} -approximate \aleph_0 -saturation, find $b_n \in M$ such that $\models r^{\mathfrak{p}(\varepsilon/2)}(b_n^{\varepsilon/2}, b_{<n}^{\varepsilon/2}, \bar{a}^{\varepsilon/2})$. That is to say that there exist $d_{\leq n}, \bar{c}$ (possibly outside M) such that

$$\begin{aligned} d(b_{\leq n} \bar{a}, d_{\leq n}, \bar{c}) &\leq \varepsilon/2, \\ \models r^{\mathfrak{p}(\varepsilon/2)}(d_{\leq n}, \bar{c}), \end{aligned}$$

From which we conclude that:

$$\begin{aligned} \models p^{\mathfrak{p}(\varepsilon/2)}(d_n, d_{<n}^{\delta}, \bar{c}^{\delta})^{\mathfrak{p}(\varepsilon/2)}, \\ \models p^{\mathfrak{p}(\varepsilon)}(d_n, d_{<n}^{\varepsilon/2}, \bar{c}^{\varepsilon/2}), \\ \models p^{\mathfrak{p}(\varepsilon)}(b_n^{\varepsilon/2}, b_{<n}^{\varepsilon}, \bar{a}^{\varepsilon}), \\ \models p^{\mathfrak{p}(\varepsilon)}(b_{\leq n}^{\varepsilon}, \bar{a}^{\varepsilon}). \end{aligned} \quad \blacksquare_{2.3}$$

Remark 2.4. Lemma 2.3 can be restated as saying that if a structure M is \mathfrak{p} -approximately \aleph_0 -saturated then it is still so after the adjunction of the sort of n -tuples (namely a sort for M^n , equipped with the supremum metric). It follows that \mathfrak{p} -approximate \aleph_0 -saturation is not affected by the adjunction of the sort of \aleph_0 -tuples (with the metric $d((a_n)_{n \in \mathbb{N}}, (b_n)_{n \in \mathbb{N}}) = \sum 2^{-n-1} d(a_n, b_n)$), or of any imaginary sort (with the natural metric). Thus the following results can be extended to \aleph_0 -tuples and imaginary sorts as well.

Lemma 2.5. *Assume that M is \mathfrak{p} -approximately \aleph_0 -saturated. Then for every finite tuple $\bar{a} \in M$, type $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$ and $\varepsilon > 0$, $p^{\mathfrak{p}(\varepsilon)}(\bar{x}, \bar{a}^{\varepsilon})$ is realised in M .*

Proof. By Lemma 2.3 we may assume that x and a are singletons.

Let $\varepsilon_i = (1 - 2^{-i})\varepsilon$, and choose $\delta_i > 0$ small enough so that:

- (i) $\delta_i \leq 2^{-i-2}\varepsilon$.
- (ii) $[d(x, y) \leq \delta_i]^{\mathfrak{p}(\varepsilon)} \subseteq [d(x, y) \leq 2^{-i}]$.
- (iii) $[d(x, y) \leq \varepsilon_i]^{\mathfrak{p}(\delta_i)} \subseteq [d(x, y) \leq \varepsilon_i + 2^{-i-2}\varepsilon]$.

Notice that the second is possible since $\mathfrak{p}(\varepsilon)$ is uniformly continuous. The third is possible by a compactness argument using the facts that $[d(x, y) \leq \varepsilon_i + 2^{-i-2}\varepsilon]$ contains a neighbourhood of $[d(x, y) \leq \varepsilon_i]$, and

$$[d(x, y) \leq \varepsilon_i] = [d(x, y) \leq \varepsilon_i]^{\mathfrak{p}(0)} = \bigcap_{\delta > 0} [d(x, y) \leq \varepsilon_i]^{\mathfrak{p}(\delta)}.$$

Let us also agree that $\delta_{-1} = \infty$.

We now choose a sequence $b_i \in M$ such that $\models p^{\mathfrak{p}(\varepsilon_i)}(b_i^{\delta_{i-1}}, a^{\varepsilon_i})$:

- Since $\delta_{-1} = \infty$ and $p(x, a)$ is consistent, any $b_0 \in M$ will do.
- Let b_i be given. Then in an elementary extension of M there exists c such that $d(c, b_i) \leq \delta_{i-1}$ and $\models p^{\mathfrak{p}(\varepsilon_i)}(c, a^{\varepsilon_i})$. Let $q(x, y, z) = \text{tp}(c, b_i, a)$. By the saturation assumption there exists $b_{i+1} \in M$ such that $\models q^{\mathfrak{p}(\delta_i)}(b_{i+1}^{\delta_i}, b_i^{\delta_i}, a^{\delta_i})$. We know that $q(x, y, z) \vdash p^{\mathfrak{p}(\varepsilon_i)}(x, z^{\varepsilon_i})$, so:

$$\begin{aligned} q^{\mathfrak{p}(\delta_i)}(x, y, z) \vdash p^{\mathfrak{p}(\varepsilon_i + \delta_i)}(x, z^{\varepsilon_i + 2^{-i-2}\varepsilon}) \vdash p^{\mathfrak{p}(\varepsilon_{i+1})}(x, z^{\varepsilon_i + 2^{-i-2}\varepsilon}) \\ q^{\mathfrak{p}(\delta_i)}(x^{\delta_i}, y^{\delta_i}, z^{\delta_i}) \vdash p^{\mathfrak{p}(\varepsilon_{i+1})}(x^{\delta_i}, z^{\varepsilon_i + 2^{-i-2}\varepsilon + \delta_i}) \vdash p^{\mathfrak{p}(\varepsilon_{i+1})}(x^{\delta_i}, z^{\varepsilon_{i+1}}). \end{aligned}$$

Thus $\models p^{\mathfrak{p}(\varepsilon_{i+1})}(b_{i+1}^{\delta_i}, a^{\varepsilon_{i+1}})$ as required.

We also know that $q(x, y, z) \vdash d(x, y) \leq \delta_{i-1}$. It follows that $q^{\mathfrak{p}(\delta_i)}(x, y, z) \vdash d(x, y) \leq 2^{-i+1}$ (except when $i = 0$), so $d(b_i, b_{i+1}) \leq 2^{-i+1} + 2\delta_i \leq 2^{-i-1}(4 + \varepsilon)$, so $(b_i : i \in \mathbb{N})$ is a Cauchy sequence in M and therefore converges to some $b \in M$. For all $i < j \in \mathbb{N}$ we have $\models p^{\mathfrak{p}(\varepsilon)}(b_j^{\delta_i}, a^{\varepsilon})$, so $\models p^{\mathfrak{p}(\varepsilon)}(b^{\delta_i}, a^{\varepsilon})$ for all $i \in \mathbb{N}$, and as $\delta_i \rightarrow 0$ we conclude that $\models p^{\mathfrak{p}(\varepsilon)}(b, a^{\varepsilon})$, as required. $\blacksquare_{2.5}$

Proposition 2.6. *Assume that M is \mathfrak{p} -approximately \aleph_0 -saturated. Then for every finite tuple $\bar{a} \in M$, type $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$ and $\varepsilon > 0$ there are $\bar{b}, \bar{a}' \in M$ such that:*

- (i) $d(\bar{a}, \bar{a}') \leq \varepsilon$.
- (ii) $\models p^{\mathfrak{p}(\varepsilon)}(\bar{b}, \bar{a}')$.

Proof. Let

$$q(\bar{x}, \bar{y}, \bar{a}) := p(\bar{x}, \bar{y}) \wedge \bar{y} = \bar{a}.$$

By Step II there are $\bar{b}, \bar{a}' \in M$ such that $\models q^{\mathfrak{p}(\varepsilon)}(\bar{b}, \bar{a}', \bar{a}^\varepsilon)$. Since $q(\bar{x}, \bar{y}, \bar{z})$ implies that $\bar{y} = \bar{z}$, so does $q^{\mathfrak{p}(\varepsilon)}$, so $\models p(\bar{b}, \bar{a}')$ and $d(\bar{a}', \bar{a}) \leq \varepsilon$, as required. $\blacksquare_{2.6}$

Proposition 2.7. *Any two elementarily equivalent separable \mathfrak{p} -approximately \aleph_0 -saturated structures are \mathfrak{p} -isomorphic.*

Proof. Let $M \equiv N$ be two separable \mathfrak{p} -approximately \aleph_0 -saturated models, and let $\varepsilon > 0$ be given. Let $M_0 = \{a_i : i \in \mathbb{N}\}$ and $N_0 = \{b_i : i \in \mathbb{N}\}$ be countable dense subsets of M and N , respectively.

Define for convenience $\varepsilon_i = (1 - 2^{-i})\varepsilon$ for all $i \in \mathbb{N}$. As $\mathfrak{p}(\varepsilon)$ is uniformly continuous, we may also choose $\delta_i > 0$ such that $[d(x, y) \leq \delta_i]^{\mathfrak{p}(\varepsilon)} \subseteq [d(x, y) \leq 2^{-i-1}]$ (so in particular, $\delta_i \leq 2^{-i-1}$).

We construct a sequence of 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 $i > 0$:

$$\begin{aligned} A_{i+1} &= a_{\leq i} \cup A_i \cup g_i(B_i) \\ B_{i+1} &= b_{\leq i} \cup B_i \cup f_{i+1}(A_{i+1}). \end{aligned}$$

- (ii) For all $c \in A_i$: $d(c, g_i \circ f_i(c)) \leq \delta_i$.
- (iii) For all $c \in B_i$: $d(c, f_{i+1} \circ g_i(c)) \leq \delta_i$.
- (iv) For each i , f_i is a $\mathfrak{p}(\varepsilon_{2i})$ -perturbation and g_i is a $\mathfrak{p}(\varepsilon_{2i+1})$ -one.

We start with $f_0 = \emptyset$, which is 0-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 a $\mathfrak{p}(\varepsilon_{2i})$ -perturbation, there is a type $q(\bar{x}, \bar{y}) \in p^{\mathfrak{p}(\varepsilon_{2i})}$ such that $q(\bar{x}, A_i)$ is consistent. By \mathfrak{p} -approximate \aleph_0 -saturation of M there are tuples $B'_i, A'_i \subseteq M$ such that $d(A_i, A'_i) \leq \delta_i$ and $M \models q(B'_i, A'_i)^{\mathfrak{p}(2^{-2i-1}\varepsilon)}$. Then $g_i: B_i \mapsto B'_i$ is $\mathfrak{p}(\varepsilon_{2i+1})$ -elementary, so it will do.

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

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

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

Therefore the sequence of mappings f_i converges to a mapping $f: A \rightarrow N$, where $A = \bigcup A_i$. As f_i is an $\mathfrak{p}(\varepsilon)$ -perturbation for all i so is f . As $M_0 \subseteq A$ we have $\bar{A} = M$, so f extends uniquely to a $\mathfrak{p}(\varepsilon)$ -perturbation $\bar{f}: M \rightarrow N$. An $\mathfrak{p}(\varepsilon)$ -perturbation $\bar{g}: N \rightarrow M$ is constructed similarly.

Finally, for $i < j \in \mathbb{N}$ choose $k \geq j$ such that $2^{-k+2} \leq \delta_j$. Then:

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

By letting $j \rightarrow \infty$ we see that $\bar{g} \circ \bar{f}$ is the identity on M_0 , and therefore on M . Similarly $\bar{f} \circ \bar{g} = \text{id}_N$. $\blacksquare_{2.7}$

3. CATEGORICITY UP TO PERTURBATION

We now turn to the proof of a Ryll-Nardzewski style characterisation of separable categoricity up to small perturbations. As usual, T denotes a theory and \mathfrak{p} a perturbation system for T .

Definition 3.1. Let $\kappa \geq |T|$ be a cardinal (recall that $|T| = |\mathcal{L}| + \aleph_0$). We say that a theory T is \mathfrak{p} - κ -categorical if it has a model of density character κ , and in addition every two models $M, N \models T$ of density character κ are \mathfrak{p} -isomorphic.

For the purpose of this definition we consider the density character of a finite set to be \aleph_0 . In particular, the complete theory of a compact, or finite, structure, will be considered \aleph_0 -categorical, and therefore \mathfrak{p} - \aleph_0 -categorical for all \mathfrak{p} .

Remark 3.2. It is not difficult to verify a general converse to Proposition 2.7, i.e., that if \mathfrak{p} is a perturbation system and M and N are \mathfrak{p} -isomorphic then $M \equiv N$. Thus Vaught's Test holds just as well for perturbed categoricity: if T has no compact models and is \mathfrak{p} - κ -categorical for some $\kappa \geq |T|$ then T is complete.

Convention 3.3. For the rest of this section we assume that T admits non compact models.

Lemma 3.4. *A complete countable theory T is \mathfrak{p} - \aleph_0 -categorical if and only if all separable models of T are \mathfrak{p} -approximately \aleph_0 -saturated.*

Proof. Right to left follows from 2.7.

Conversely, assume that T is \mathfrak{p} - \aleph_0 -categorical, and let $M \models T$ be separable. Let $\bar{a} \in M^n$, $q(\bar{x}) = \text{tp}(\bar{a})$, and let $q(y, \bar{a}) \in S_1(\bar{a})$, where $q(y, \bar{x}) \in S_{n+1}(T)$ is a complete pure type. Let also $\varepsilon > 0$, and $\delta = \delta_\rho(\varepsilon) > 0$.

By the downward Löwenheim-Skolem theorem there exists a separable model $N \models T$ such that for every n -tuple $\bar{b} \in N$, if \bar{b} satisfies $p^{\mathfrak{p}(\varepsilon/2)}$ then $q^{\mathfrak{p}(\varepsilon/2)}(y, \bar{b}^\delta)$ is realised in N . By assumption there exists a $\mathfrak{p}(\varepsilon/2)$ -isomorphism $f: M \rightarrow N$. Then $\bar{b} = f(\bar{a}) \models p^{\mathfrak{p}(\varepsilon/2)}$, and let $c \in N$ be such that $\models q^{\mathfrak{p}(\varepsilon/2)}(c, \bar{b}^\delta)$. Letting $d = f^{-1}(c)$ we get $\models q^{\mathfrak{p}(\varepsilon)}(d, \bar{a}^\varepsilon)$.

Thus M is \mathfrak{p} -approximately \aleph_0 -saturated. ■_{3.4}

We observe that if \mathfrak{p} is a perturbation system, then the topology on $S_n(T)$ induced by $d_{\mathfrak{p}}$ is finer than the logic topology. Indeed, if U is a neighbourhood of p , then $\bigcap_{\varepsilon > 0} p^{\mathfrak{p}(\varepsilon)} = \{p\} \subseteq U$, and by compactness we must have $p^{\mathfrak{p}(\varepsilon)} \subseteq U$ for some $\varepsilon > 0$. In fact, $d_{\mathfrak{p}}$ is usually too fine to be used directly for characterising \aleph_0 -categoricity. For example, in case of the identity perturbation (i.e., no perturbation allowed at all), $d_{\mathfrak{p}, n}$ is a discrete metric (with values in $\{0, \infty\}$), while the standard Ryll-Nardzewski theorem for continuous logic does consider a much coarser topology, namely that induced by the metric d . We therefore need to take both metrics into account.

Theorem 3.5. *Let T be a complete countable theory, \mathfrak{p} a perturbation system for T . Then the following are equivalent:*

- (i) *The theory T is \mathfrak{p} - \aleph_0 -categorical.*
- (ii) *For every $n \in \mathbb{N}$, finite \bar{a} , $p \in S_n(\bar{a})$ and $\varepsilon > 0$, the set $[p^{\mathfrak{p}(\varepsilon)}(\bar{x}^\varepsilon, \bar{a}^\varepsilon)]$ has non empty interior in $S_n(\bar{a})$.*
- (iii) *Same restricted to $n = 1$.*

Proof. (i) \implies (ii). Assume there is some finite tuple \bar{a} , $n \in \mathbb{N}$ and $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$, such that for some $\varepsilon > 0$ the set $[p^{\mathfrak{p}(\varepsilon)}(\bar{x}^\varepsilon, \bar{a}^\varepsilon)]$ has empty interior in $S_n(\bar{a})$. Then it is nowhere dense in $S_n(\bar{a})$, and can be omitted in a dense subset of some separable model $(M, \bar{a}) \models T_{\bar{a}}$. Therefore $[p^{\mathfrak{p}(\varepsilon)}(\bar{x}, \bar{a}^\varepsilon)]$ is omitted in M , which is therefore not \mathfrak{p} -approximately \aleph_0 -saturated. Therefore T cannot be \mathfrak{p} - \aleph_0 -categorical.

(ii) \implies (iii). Clear.

(iii) \implies (i). We show that every $M \models T$ is \mathfrak{p} -approximately \aleph_0 -saturated. Indeed, let $\bar{a} \in M$ be a finite tuple and $p(x, \bar{a}) \in S_1(T)$. As $[p^{\mathfrak{p}(\varepsilon)}(x^\varepsilon, \bar{a}^\varepsilon)]$ has non empty interior in $S_n(\bar{a})$ it must be realised in M . ■_{3.5}

We may wish to combine the two metrics in a single one. While one may try to achieve this through various general approaches for the combination of two metrics, the specific situation in which we find ourselves suggests a specific construction as the “natural” one.

Fix $m \in \mathbb{N}$, and let $\mathcal{L}_{\bar{c}} = \mathcal{L} \cup \{c_i : i < m\}$, where each c_i is a new distinct constant symbol. Let $T_{\bar{c}}$ be the (incomplete) $\mathcal{L}_{\bar{c}}$ -theory generated by T . We extend \mathfrak{p} into a perturbation system $\mathfrak{p}_{\bar{c}}$ for $T_{\bar{c}}$ by allowing the new constant symbols to move a little. It is more convenient to think in terms of a relational language, in which each of the constants c_i is represented by a unary predicate giving the distance to c_i . We therefore define $\mathfrak{p}_{\bar{c}}^0(\varepsilon)$ to be the perturbation pre-radius which, for $p, q \in S_n(T_{\bar{c}})$, allows to perturb p to q if and only if:

- (i) $q \upharpoonright_{\mathcal{L}} \in (p \upharpoonright_{\mathcal{L}})^{\mathfrak{p}(\varepsilon)}$; and:
- (ii) For all $i < m$ and $j < n$: $|d(x_j, c_i)^p - d(x_j, c_i)^q| \leq \varepsilon$.

It is easy to verify that $\mathfrak{p}_{\bar{c}}^0$ is a uniformly continuous perturbation pre-system, which generates a perturbation system $\mathfrak{p}_{\bar{c}} = \llbracket \mathfrak{p}_{\bar{c}}^0 \rrbracket$. Thus $\mathfrak{p}_{\bar{c}}$ can be roughly described as allowing to perturb models of T according to \mathfrak{p} , and to move the new constants (i.e., change the distance to them) a little as well. If $\bar{c} = \emptyset$ we changed nothing: $\mathfrak{p}_{\emptyset} = \mathfrak{p}$.

By definition of the bi-reduct $\llbracket \cdot \rrbracket$, we have for all $\varepsilon > 0$, and $(M, \bar{a}), (N, \bar{b}) \models T_{\bar{c}}$:

$$\begin{aligned} \text{BiPert}_{\mathfrak{p}_{\bar{c}}^0(\varepsilon)}((M, \bar{a}), (N, \bar{b})) &= \text{BiPert}_{\mathfrak{p}_{\bar{c}}(\varepsilon)}((M, \bar{a}), (N, \bar{b})) = \\ &= \{f \in \text{BiPert}_{\mathfrak{p}(\varepsilon)}(M, N) : (\forall e \in M, i < m) (|d^M(e, a_i) - d^N(f(e), b_i)| \leq \varepsilon)\}. \end{aligned}$$

The space $S_0(T_{\bar{c}})$ is the set of completions of $T_{\bar{c}}$ and can be naturally identified with $S_m(T)$. We define a metric $\tilde{d}_{\mathfrak{p}, m}$ on $S_m(T)$ as the image of $d_{\mathfrak{p}_{\bar{c}}}$ under this identification. Equivalently:

Definition 3.6. For $p, q \in S_n(T)$, we define $\tilde{d}_{\mathfrak{p}}(p, q)$ as the infimum of all ε for which there exist models $M, N \models T$, $\bar{a} \in M^n$ and $\bar{b} \in N^n$ and a mapping $f: M \rightarrow N$ such that:

- (i) $\bar{a} \models p$ and $\bar{b} \models q$.
- (ii) $f \in \text{BiPert}_{\mathfrak{p}(\varepsilon)}(M, N)$.
- (iii) For all $i < n$ and $c \in M$: $|d^M(c, a_i) - d^N(f(c), b_i)| \leq \varepsilon$.

Alternatively, we may wish to restrict $\mathfrak{p}_{\bar{c}}$ to a specific completion of $T_{\bar{c}}$. Any such completion is of the form $T_{\bar{a}} = \text{Th}(M, \bar{a})$, where $M \models T$ and $\bar{a} \in M^n$. Let us denote the restriction of $\mathfrak{p}_{\bar{c}}$ to $T_{\bar{a}}$ by $\mathfrak{p}_{\bar{a}}$.

Of course, once we have constructed $\mathfrak{p}_{\bar{a}}$, we can construct $\tilde{d}_{\mathfrak{p}_{\bar{a}}}$ as above, and it follows immediately from the definitions that:

Lemma 3.7. *The construction $\mathfrak{p} \mapsto \tilde{d}_{\mathfrak{p}}$ commutes with the addition of parameters, in the sense that for all \bar{a}, \bar{b} and \bar{c} , if $|\bar{b}| = |\bar{c}|$ then:*

$$\tilde{d}_{\mathfrak{p}_{\bar{a}}}(\text{tp}(\bar{b}/\bar{a}), \text{tp}(\bar{c}/\bar{a})) = \tilde{d}_{\mathfrak{p}}(\text{tp}(\bar{b}, \bar{a}), \text{tp}(\bar{c}, \bar{a})).$$

In an arbitrary metric space (X, d) , let $B_d(x, \varepsilon)$ denote the closed ε -ball around a point x . The following result characterise the topology defined by $\tilde{d}_{\mathfrak{p}}$:

Lemma 3.8. *Fix $n \in \mathbb{N}$ and a finite tuple $\bar{a} \in M \models T$. The metric $\tilde{d}_{\mathfrak{p}_{\bar{a}}}$ is coarser (i.e., smaller) on $S_n(\bar{a})$ than both d and $d_{\mathfrak{p}_{\bar{a}}}$, and finer than the logic topology.*

Also, for every $p(\bar{x}, \bar{a}) \in S_n(\bar{a})$, the family $\{[p^{\mathfrak{p}(\varepsilon)}(\bar{x}^\varepsilon, \bar{a}^\varepsilon)] : \varepsilon > 0\}$ forms a base of $\tilde{d}_{\mathfrak{p}_{\bar{a}}}$ -neighbourhoods for $p(\bar{x}, \bar{a})$.

Proof. Let us start by showing that for every $\varepsilon > 0$ there is $\varepsilon' > 0$ such that:

$$[p^{\mathfrak{p}(\varepsilon')}(\bar{x}^{\varepsilon'}, \bar{a}^{\varepsilon'})] \subseteq B_{\tilde{d}_{\bar{p}\bar{a}}}(p(\bar{x}, \bar{a}), \varepsilon) \subseteq [p^{\mathfrak{p}(\varepsilon)}(\bar{x}^{\varepsilon}, \bar{a}^{\varepsilon})].$$

Let us first consider the case without parameters. The set $\{(q(x, y), q'(x, y)) \in S_2(T) : |d(x, y)^q - d(x, y)^{q'}| \geq \varepsilon/2\}$ is closed and disjoint of the diagonal, so by compactness there is $\varepsilon' > 0$ such for all $(q, q') \in \mathfrak{p}_2(\varepsilon')$: $|d(x, y)^q - d(x, y)^{q'}| \leq \varepsilon/2$. We may of course assume that $\varepsilon' \leq \varepsilon/2$, and the first inclusion follows. The second inclusion is immediate from the definition of $\tilde{d}_{\bar{p}}$.

The case over parameters \bar{a} follows from the case without parameters and the fact that by Lemma 3.7:

$$B_{\tilde{d}_{\bar{p}\bar{a}}}(p(\bar{x}, \bar{a}), \varepsilon) = \{q(\bar{x}, \bar{a}) \in S_n(\bar{a}) : q(\bar{x}, \bar{y}) \in B_{\tilde{d}_{\bar{p}}}(p, \varepsilon)\}.$$

Finally, let $K \subseteq S_n(T)$ be closed in the logic topology and $q \notin K$. Then there is $\varepsilon > 0$ such that $q^{\mathfrak{p}(\varepsilon)} \cap K = \emptyset$, and since $q^{\mathfrak{p}(\varepsilon)}$ is closed in the logic topology there is also $\varepsilon' > 0$ such that $[q^{\mathfrak{p}(\varepsilon)}(\bar{x}^{\varepsilon'})] \cap K = \emptyset$. Letting $\varepsilon'' = \min\{\varepsilon, \varepsilon'\}$ we see that $[q^{\mathfrak{p}(\varepsilon'')}(\bar{x}^{\varepsilon''})] \cap K = \emptyset$. Therefore K is $\tilde{d}_{\bar{p}}$ -closed. This shows that $\tilde{d}_{\bar{p}}$ refines the logic topology. It is clearly coarser than both d and $d_{\bar{p}}$. Substituting $T_{\bar{a}}$ for T in the last argument we get the case with parameters. $\blacksquare_{3.8}$

Thus we can restate Theorem 3.5 as:

Theorem 3.9. *Let T be a complete countable theory. Then the following are equivalent:*

- (i) *The theory T is $\mathfrak{p}\text{-}\aleph_0$ -categorical.*
- (ii) *For every $n \in \mathbb{N}$, finite \bar{a} , $p \in S_n(\bar{a})$ and $\varepsilon > 0$, the ε -ball $B_{\tilde{d}_{\bar{p}\bar{a}}}(p, \varepsilon)$ has non empty interior in the logic topology on $S_n(\bar{a})$.*
- (iii) *Same restricted to $n = 1$.*

The statement of the result in terms of non empty interior may sound a little weird, as the non perturbed Ryll-Nardzewski theorem tells us that T is \aleph_0 -categorical if and only if the metric d coincides with the logic topology. In order to explain this apparent discrepancy let us make a few more observations.

First, the coincidence of the logic topology with the metric $\tilde{d}_{\bar{p}}$ is a sufficient condition for T to be $\mathfrak{p}\text{-}\aleph_0$ -categorical. In this case it suffices to check $S_n(T)$ alone (i.e., no need to consider parameters).

Proposition 3.10. *Assume that T is countable and complete, and $\tilde{d}_{\bar{p}}$ coincides with the logic topology on $S_n(T)$ for all T . Then T is $\mathfrak{p}\text{-}\aleph_0$ -categorical.*

Proof. Let $M \models T$, and let $\bar{a} \in M$, $p(x, \bar{a}) \in S_1(\bar{a})$, and $\varepsilon > 0$. By assumption $B_{\tilde{d}_{\bar{p}}}(p(x, \bar{y}), \varepsilon)$ is a neighbourhood of p , so there is a formula $\varphi(x, \bar{y})$ such that

$$p \in [\varphi = 0] \subseteq [\varphi < 1/2] \subseteq B_{\tilde{d}_{\bar{p}}}(p(x, \bar{y}), \varepsilon) \subseteq [p^{\mathfrak{p}(\varepsilon)}(x^\varepsilon, \bar{y}^\varepsilon)].$$

Therefore $[\varphi(x, \bar{a}) < 1/2]$ is a non empty open subset of $S_1(\bar{a})$ (as it contains $p(x, \bar{a})$), and is therefore realised in M . Thus $p^{\mathfrak{p}(\varepsilon)}(x^\varepsilon, \bar{a}^\varepsilon)$ is realised in M . $\blacksquare_{3.10}$

We should point out that the consideration of parameters in Theorem 3.5 is unavoidable. Indeed, if $p(x, \bar{a}) \in S_1(\bar{a})$ and we only assume that $B_{\tilde{d}_{\bar{p}}}(p(x, \bar{y}), \varepsilon)$ has non empty interior in $S_{n+1}(T)$, which need not necessarily contain $p(x, \bar{y})$, it may happen that no type $q(x, \bar{y})$ in this interior is consistent with $r(\bar{y}) = \text{tp}(\bar{a})$, so pulling up to $S_1(\bar{a})$ we may end up with an empty set.

The sufficient condition in Proposition 3.10 seems far more convenient and natural than the one in Theorem 3.5, and one might hope to show that it is also necessary. The following example shows that this is impossible. Roughly speaking, this example says that if the ‘‘if and only if’’ variant of Proposition 3.10 were true, we could prove Vaught’s no-two-models theorem, which fails in continuous first order logic.

Example 3.11. Let T be the theory of atomless L^p -Banach lattices for some fixed $p \in [1, \infty)$, studied in [BBH]. It is known that T is \aleph_0 -categorical, all of its separable models being isomorphic to $L^p[0, 1]$.

Let f be any positive function of norm 1 (this determines $\text{tp}(f)$), say $f = \chi_{[0,1]}$, then T_f has precisely two non isomorphic separable models, namely $(L^p[0, 1], \chi_{[0,1]})$ and $(L^p[0, 2], \chi_{[0,1]})$. (The theory T_f is \aleph_0 -categorical up to perturbations of the new constant f , but that's not what we are looking for). Let g be another positive function of norm 1 such that $f \wedge g = 0$ (this determines $\text{tp}(f, g)$). Then again, $T_{f,g}$ has precisely two separable models, $(L^p[0, 2], \chi_{[0,1]}, \chi_{[1,2]})$ and $(L^p[0, 3], \chi_{[0,1]}, \chi_{[1,2]})$.

Let \mathfrak{p} be the identity perturbation system for T_f , and thus \mathfrak{p}_g is the perturbation system for $T_{f,g}$ that allows to perturb g while preserving all the rest untouched. Then the two models above are \mathfrak{p}_g -isomorphic, so $T_{f,g}$ is \mathfrak{p}_g - \aleph_0 -categorical. Let $\pi_n: S_n(T_{f,g}) \rightarrow S_n(T_f)$ be the reduct projection. As \mathfrak{p} is the identity perturbation on T_f , $\tilde{d}_{\mathfrak{p}} = d$ on $S_n(T_f)$. Therefore, if $U \subseteq S_n(T_f)$ is d -open then $\pi_n^{-1}(U) \subseteq S_n(T_{f,g})$ is $\tilde{d}_{\mathfrak{p}_g}$ -open.

But T_f is not \aleph_0 -categorical, so the metric d defines a non compact topology on $S_n(T)$, whereby $\tilde{d}_{\mathfrak{p}_g}$ defines a non compact topology on $S_n(T_{f,g})$, which in particular cannot coincide with the logic topology, even though $T_{f,g}$ is \mathfrak{p}_g - \aleph_0 -categorical.

One last point arises from a comparison of Theorem 3.5 with the unperturbed Ryll-Nardzewski Theorem for continuous logic. The latter characterises unperturbed \aleph_0 -categoricity by the coincidence of the logic topology with the metric, and thus does not seem to follow as a special case of Theorem 3.5. To see that it actually does, we need to explore some further properties perturbation metrics may have.

Let us start by recalling properties of the standard metric d on $S_n(T)$. We observe in [BU] that the metric d has the following properties:

- (i) It refines the logic topology.
- (ii) If $F = [p(\bar{x})] \subseteq S_n(T)$ is closed, then so is $F^\varepsilon = \{p: d(p, F) \leq \varepsilon\} = [p(\bar{x}^\varepsilon)]$.
- (iii) For every injective $\sigma: n \rightarrow m$, $p \in S_n(T)$, $q \in S_m(T)$:

$$d(p, \sigma^*(q)) = d(\sigma^{*-1}(p), q).$$

A perturbation metric has all these properties as well, and in fact satisfies the last one also for σ which is not injective. One last interesting property of $(S_n(T), d)$ is analogous to the second property:

Lemma 3.12. *If $U \subseteq S_n(T)$ is open, then so is $U^{<\varepsilon} = \{p: d(p, U) < \varepsilon\}$.*

Proof. It suffices to show this for a basis of open sets, i.e., for sets of the form $U = [\varphi(\bar{x}) < \varepsilon]$. But then $U^{<\varepsilon} = [\inf_{\bar{y}}(\varphi(\bar{y}) \vee d(\bar{x}, \bar{y})) < \varepsilon]$ is open. ■_{3.12}

For lack of a better name, let us call provisionally a metric on a topological space *open* if it satisfies the property of Lemma 3.12.

Definition 3.13. Let \mathfrak{p} be a perturbation system for T .

- (i) We say that \mathfrak{p} is *open* if $d_{\mathfrak{p}}$ is open on $S_n(T)$ for all n .
- (ii) We say that \mathfrak{p} is *weakly open* if for all $\varepsilon > 0$ and $n \in \mathbb{N}$ there is $\delta > 0$ such that for every open set $U \subseteq S_n(T)$:

$$U^{\tilde{d}_{\mathfrak{p}} < \delta} \subseteq (U^{\tilde{d}_{\mathfrak{p}} < \varepsilon})^\circ.$$

(Where $U^{\tilde{d}_{\mathfrak{p}} < \delta} = \{p: \tilde{d}_{\mathfrak{p}}(p, U) < \delta\}$.)

Lemma 3.14. *Let \mathfrak{p} be a perturbation system.*

- (i) \mathfrak{p} is *weakly open* if and only if for every $\varepsilon > 0$ and $n \in \mathbb{N}$ there is $\delta > 0$ such that for every open $U \subseteq S_n(T)$:

$$U^{\mathfrak{p}(\delta)} \subseteq (U^{\tilde{d}_{\mathfrak{p}} < \varepsilon})^\circ.$$

- (ii) If \mathfrak{p} is open then it is weakly open.
- (iii) If $\tilde{d}_{\mathfrak{p}}$ is open on $S_n(T)$ for all n then \mathfrak{p} is weakly open.

Proof. (i) For one direction use the fact that $U^{\mathfrak{p}(\delta/2)} \subseteq U^{\tilde{d}_{\mathfrak{p}} < \delta}$. For the other, assume that $U^{\mathfrak{p}(\delta)} \subseteq (U^{\tilde{d}_{\mathfrak{p}} < \varepsilon/2})^\circ$ and $\delta < \varepsilon/2$. Then since the metric d is open:

$$U^{\tilde{d}_{\mathfrak{p}} < \delta} \subseteq (U^{\mathfrak{p}(\delta)})^{d < \delta} \subseteq \left((U^{\tilde{d}_{\mathfrak{p}} < \varepsilon/2})^\circ \right)^{d < \delta} \subseteq \left((U^{\tilde{d}_{\mathfrak{p}} < \varepsilon/2})^{d < \delta} \right)^\circ \subseteq (U^{\tilde{d}_{\mathfrak{p}} < \varepsilon})^\circ.$$

- (ii) We use the criterion from the previous item:

$$U^{\mathfrak{p}(\varepsilon/2)} \subseteq U^{d_{\mathfrak{p}} < \varepsilon} = (U^{d_{\mathfrak{p}} < \varepsilon})^\circ \subseteq (U^{\tilde{d}_{\mathfrak{p}} < \varepsilon})^\circ.$$

- (iii) Immediate from the definition. ■_{3.14}

Theorem 3.15. *Let T be a complete countable theory, \mathfrak{p} a weakly open perturbation system. Then T is \mathfrak{p} - \aleph_0 -categorical if and only if for every n , $\tilde{d}_{\mathfrak{p}}$ coincides with the logic topology on $S_n(T)$.*

Proof. Right to left is by Proposition 3.10, so we prove left to right.

Assume that T is \mathfrak{p} - \aleph_0 -categorical. Fix $p \in S_n(T)$ and $\varepsilon > 0$. Then by definition there is $\delta > 0$ such that for every open set $U \subseteq S_n(T)$: $U^{\tilde{d}_{\mathfrak{p}} < \delta} \subseteq (U^{\tilde{d}_{\mathfrak{p}} < \varepsilon/2})^\circ$. We may also assume that $\delta < \varepsilon$.

Let $U = B_{\tilde{d}_{\mathfrak{p}}}(p, \delta/2)^\circ$, so $U \neq \emptyset$ by Theorem 3.5. Then:

$$p \in U^{\tilde{d}_{\mathfrak{p}} < \delta} \subseteq (U^{\tilde{d}_{\mathfrak{p}} < \varepsilon/2})^\circ \subseteq \left(B_{\tilde{d}_{\mathfrak{p}}}(p, \delta/2)^{\tilde{d}_{\mathfrak{p}} < \varepsilon/2} \right)^\circ \subseteq (B_{\tilde{d}_{\mathfrak{p}}}(p, \varepsilon))^\circ$$

Therefore $B_{\tilde{d}_{\mathfrak{p}}}(p, \varepsilon)$ is a logic neighbourhood of p for all p and $\varepsilon > 0$. Since $\tilde{d}_{\mathfrak{p}}$ refines the logic topology, they must coincide. ■_{3.15}

Example 3.16. The identity perturbation system is open.

Corollary 3.17 (Henson's unperturbed Ryll-Nardzewski Theorem). *A complete countable theory T is \aleph_0 -categorical if and only if the standard metric d coincides with the logic topology on $S_n(T)$, for all $n \in \mathbb{N}$.*

Proof. Since $\tilde{d}_{\text{id}} = d$. ■_{3.17}

Example 3.18. Let $\bar{a} \in M \models T$, and let $p = \text{tp}(\bar{a})$ be isolated (i.e., $d(\bar{x}, p)$ is a definable predicate). Let \mathfrak{p} be the identity perturbation for T , and $\mathfrak{p}_{\bar{a}}$ as above be a perturbation system for $T_{\bar{a}}$ allowing to move the named parameter. Then $\mathfrak{p}_{\bar{a}}$ is open.

Proof. Exercise. ■_{3.18}

Of course, in Example 3.11 the type of the new parameter $\text{tp}(g/f)$ was not isolated.

4. PERTURBATIONS OF AUTOMORPHISMS

We conclude with a few problems concerning perturbations of automorphisms which motivated the author's initial interest in perturbations, and which the author therefore finds worthy of future study.

One such problem comes from the study of the properties of automorphism groups of classical (i.e., discrete) countable structures, and in particular of ones whose first order theory is \aleph_0 -categorical, viewed as topological groups. Model-theoretic questions of this kind are treated, for example, in [HHLS93], while more topologically profound questions are studied by Kechris and Rosendal [KR07]. It is natural to ask whether such of these results can be generalised to the automorphism groups of separable continuous structures (with a separably categorical theory). A very simple instance would to consider the unitary group $U(H)$ where H is a separable infinite dimensional Hilbert space. Indeed, $U(H)$ is a polish group in the point-wise convergence topology, also known as the strong operator topology, but it is quickly

revealed that $U(H)$ is just way too big for any of the properties that Kechris and Rosendal were looking for (e.g., existence of ample generics) to hold.

This is definitely not a new phenomenon. We already know that the type space of a continuous theory, viewed as a pure topological space, is too big. In order to study notions such as superstability, \aleph_0 -stability, or even local φ -stability, one needs to take an additional metric structure into account, considering points (types) up to small distance. Such considerations date as far back as Iovino's definition of λ -stability and of superstability in Banach space structures [Iov99]. Similarly, the automorphism group of a metric structure admits a natural metric, namely the metric of uniform convergence (for example, the operator norm on $U(H)$ is the metric of uniform convergence on the unit ball). In the terminology of [Ben08], $\text{Aut}(M)$ is a *topometric group*, namely a topological group (in the point-wise convergence topology) which is at the same time a metric group (in the uniform convergence metric), such that in addition the distance function $d: G^2 \rightarrow \mathbb{R}^+$ is lower semi-continuous in the topology. One can then restate the question of the existence of ample generics as follows:

Question 4.1. Let M be a separable metric structure, $G = \text{Aut}(M)$. Under what assumptions on M can we find, for each $n \in \mathbb{N}$, a tuple $\bar{g} \in G^n$ such that for every $\varepsilon > 0$, the G -conjugacy class of the (metric) ε -ball around \bar{g} is (topologically) co-meagre? In other words can we find \bar{g} such that the metric closure of the orbit of \bar{g} is co-meagre? In particular, can one prove this is the case if $\text{Th}(M)$ is \aleph_0 -categorical and \aleph_0 -stable?

Considering an automorphism $\tau \in \text{Aut}(M)$ up to small distance in uniform convergence is essentially the same as considering the structure (M, τ) up to a small perturbation of τ , whence the connection with the topic of the present paper. For the special case of $U(H)$, a positive answer essentially follows from [Dav96, Theorem II.5.8]. What about the automorphism group of the unique separable atomless probability algebra?

Another question leading to similar considerations is raised by Berenstein and Henson [BH]. In this paper they consider the theory of probability algebras with a generic automorphism, and ask whether it is superstable (equivalently, supersimple, since they showed that the theory is stable). In classical first order logic the answer would be positive, by a theorem of Chatzidakis and Pillay [CP98]. Henson's and Berenstein's question was nonetheless answered negatively by the author, raising the following natural "second best" question, namely whether the theory of probability algebras with a generic automorphism is superstable *up to small perturbations of the automorphism*. This was subsequently answered positively by the author and Berenstein [BB09], where we show moreover that up to perturbations of the automorphism, the theory is \aleph_0 -stable.

Question 4.2. Let T be a superstable continuous theory. Let

$$T_\sigma = T \cup \{\text{"}\sigma \text{ is an automorphism"}\}.$$

Assume furthermore that T_σ has a model companion T_A . Is T_A supersimple up to small perturbations of σ ?

And in fact,

Question 4.3. What should it mean precisely for a theory to be supersimple up to small perturbations?

Regarding the last question it should be pointed out that there are several natural candidates for the definition of " a is independent up to distance ε from B over A " (denoted usually $a^\varepsilon \downarrow_A B$). While these notions of approximate independence are not equivalent, they all give rise to the same notion of supersimplicity (see for example in [Ben06]), and in a stable theory they are further equivalent to superstability. Superstability and λ -stability up to perturbation are introduced by the author in [Ben08], and one should seek a notion of supersimplicity up to perturbation which, in stable theories, coincides with superstability up to perturbation.

REFERENCES

- [BB09] Itaï Ben Yaacov and Alexander Berenstein, *On perturbations of Hilbert spaces and probability algebras with a generic automorphism*, Journal of Logic and Analysis **1:7** (2009), 1–18, arXiv:0810.4086.
- [BBH] Itaï Ben Yaacov, Alexander Berenstein, and C. Ward Henson, *Model-theoretic independence in the Banach lattices $L_p(\mu)$* , Israel Journal of Mathematics, to appear, arXiv:0907.5273.
- [BBHU08] Itaï Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov, *Model theory for metric structures*, Model theory with Applications to Algebra and Analysis, volume 2 (Zoé Chatzidakis, Dugald Macpherson, Anand Pillay, and Alex Wilkie, eds.), London Math Society Lecture Note Series, vol. 350, Cambridge University Press, 2008, pp. 315–427.
- [Ben] Itaï Ben Yaacov, *Continuous first order logic for unbounded metric structures*, submitted, arXiv:0903.4957.
- [Ben03] ———, *Positive model theory and compact abstract theories*, Journal of Mathematical Logic **3** (2003), no. 1, 85–118.
- [Ben06] ———, *On supersimplicity and lovely pairs of cats*, Journal of Symbolic Logic **71** (2006), no. 3, 763–776, arXiv:0902.0118.
- [Ben08] ———, *Topometric spaces and perturbations of metric structures*, Logic and Analysis **1** (2008), no. 3–4, 235–272, arXiv:0802.4458.
- [BH] Alexander Berenstein and C. Ward Henson, *Model theory of probability spaces with an automorphism*, submitted.
- [BU] Itaï Ben Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society, to appear, arXiv:0801.4303.
- [CP98] Zoé Chatzidakis and Anand Pillay, *Generic structures and simple theories*, Annals of Pure and Applied Logic **95** (1998), 71–92.
- [Dav96] Kenneth R. Davidson, *C^* -algebras by example*, Fields Institute Monographs, vol. 6, American Mathematical Society, Providence, RI, 1996.
- [HHLS93] Wilfrid Hodges, Ian Hodkinson, Daniel Lascar, and Saharon Shelah, *The small index property for ω -stable ω -categorical structures and for the random graph*, Journal of the London Mathematical Society. Second Series **48** (1993), no. 2, 204–218.
- [Iov99] José Iovino, *Stable Banach spaces and Banach space structures, I and II*, Models, algebras, and proofs (Bogotá, 1995), Lecture Notes in Pure and Appl. Math., vol. 203, Dekker, New York, 1999, pp. 77–117.
- [KR07] Alexander S. Kechris and Christian Rosendal, *Turbulence, amalgamation, and generic automorphisms of homogeneous structures*, Proceedings of the London Mathematical Society. Third Series **94** (2007), no. 2, 302–350.

ITAÏ BEN YAACOV, UNIVERSITÉ CLAUDE BERNARD – LYON 1, INSTITUT CAMILLE JORDAN, CNRS UMR 5208, 43 BOULEVARD DU 11 NOVEMBRE 1918, 69622 VILLEURBANNE CEDEX, FRANCE
 URL: <http://math.univ-lyon1.fr/~begnac/>