



Definability of groups in \aleph_0 -stable metric structures

Itaï Ben Yaacov

► To cite this version:

| Itaï Ben Yaacov. Definability of groups in \aleph_0 -stable metric structures. 2008. hal-00259318v1

HAL Id: hal-00259318

<https://hal.science/hal-00259318v1>

Preprint submitted on 27 Feb 2008 (v1), last revised 5 Sep 2009 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DEFINABILITY OF GROUPS IN \aleph_0 -STABLE METRIC STRUCTURES

ITAI BEN YAACOV

ABSTRACT. We prove that in a continuous \aleph_0 -stable theory every type-definable group is definable. The two main ingredients in the proof are:

- (i) Results concerning Morley ranks (i.e., Cantor-Bendixson ranks) from [Ben], allowing us to prove the theorem in case the metric is invariant under the group action; and
- (ii) Results concerning the existence of translation-invariant definable metrics on type-definable groups and the extension of partial definable metrics to total ones.

We conclude this paper with a development of basic stability theory for continuous logic.

INTRODUCTION

Definable sets, as well as more complex definable objects (e.g., groups) play a central and essential role in classical model theory. These are (usually) subsets of the ambient structure which are defined by a single formula of classical first order logic.

Continuous first order logic was proposed in [BU] as an extension of classical first order logic, obtained by replacing the two-element set of truth values $\{T, F\}$ with the compact interval $[0, 1]$. It allows to consider various classes of complete metric structures as elementary classes and to study definability therein. However, some things do become more complicated in continuous logic, and in particular the classical notion of a definable set splits in two. First, a set is a predicate, and a definable set is a definable predicate, i.e., a definable function into the set $\{T, F\}$ (or $\{0, 1\}$). As such, the correct analogue is a *definable continuous predicate*, i.e., a definable function to $[0, 1]$ – it is definable in the sense that it is given by a continuous first order formula, or at the very least, by a uniform limit of such. But when thinking of definable objects, such as groups, there is an essential asymmetry between what is inside (which interests us) and what is outside (about which we could hardly care less, especially if the set is stably embedded). The same asymmetry arises when we wish to quantify over a definable set. In that case the notion of a definable predicate is inadequate and we are led to the notion of a *definable set*: this is a closed set the distance to which is a definable predicate, or equivalently, over which we may quantify (see Fact 1.7 below).

Date: 28 February 2008.

2000 Mathematics Subject Classification. 03C45, 03C90.

Key words and phrases. continuous logic, definable set, definable group, definable metric, \aleph_0 -stability. Research partially supported by NSF grant DMS-0500172.

The class of definable set in a continuous structure is far less well-behaved than in classical logic. For example, the family of all definable subsets of M^n does not form a Boolean algebra, as it is not always closed under complement or intersection. Worse still, it is not at all obvious whether non trivial definable sets even exist. Indeed, examples exist of theories which do not admit enough definable sets, i.e., where there are distinct types which nonetheless agree on all definable sets (such an example is the theory of real closed fields augmented by a predicate $P(x) = st(|x - a|) \wedge 1$ where $st(\cdot)$ denotes the real standard part and a is some infinite element). As all known examples of this pathology are unstable it make sense to ask whether all stable continuous theories admit enough definable sets.

One of the beautiful aspects of stable group theory in classical logic is the proof that there are also “enough definable groups”, namely, that every type-definable group is the intersection of definable subgroups of a definable group. In the case of an \aleph_0 -stable theory, chain conditions along with the previous general fact yield that every type-definable group is definable. In continuous logic we can prove adequate analogues of the chain conditions for sequences of definable (or type-definable) groups for \aleph_0 -stable theories, but we do not know how to prove enough definable groups exist in stable theories.

In the present paper we give a direct proof of the fact that in an \aleph_0 -stable theory every type-definable group is definable, leaving open the question of the existence of definable groups in general stable theories. In the special case of the theory of probability algebras this has already been proved by Alexander Berenstein [Ber06].

In Section 1 we discuss various definability classes of sets and functions.

In Section 2 we prove the main theorem using some technical results concerning Morley ranks (i.e., Cantor-Bendixson ranks) from [Ben]. We do this under the assumption of the invariance of the metric under the group operation (Theorem 2.11). The rest of the paper aims towards the removal of this assumption.

In Section 3 we discuss stable type-definable groups. In particular we discuss the existence and basic properties generic types, as well as some easy descending chain conditions.

In Section 4 we study definable metrics other than the standard one. In particular, we study when and how partial definable metrics (on a definable or type-definable set) can be extended to total ones.

In Section 5 we prove the full version of main theorem and give some corollaries.

Finally, in Appendix A we give a more complete treatment of type-definable groups in stable continuous theories. For this purpose we first elaborate some general stability theory for the continuous context which was missing from [BU].

1. DEFINABILITY PROPERTIES

1.1. Definability classes of sets.

Definition 1.1. (i) A *type-definable set* X is the set of realisations of an arbitrary set of conditions $\{\varphi_i(x) = 0 : i < \lambda\}$.

- (ii) A *zero set* X is the set of realisations of a single condition $\psi(x) = 0$, where $\psi(x)$ is a definable predicate. We then say that X is the zero set of ψ .
- (iii) A *definable set* X is a closed set for which $d(x, X)$ is a definable predicate.

Type-definable sets and zero sets will only be considered in sufficiently saturated structures, while definable sets make sense in an arbitrary structure.

Clearly, every zero set is type-definable, and every definable set X is the zero set of $d(x, X)$. In terms of types, we know that a type-definable set X corresponds to closed sets of types $[X] \subseteq S_n(A)$ where A contains all the parameters appearing in the definition of X .

Definition 1.2. Let X and Y be two type-definable sets. We say that Y is a *logical neighbourhood* of X , in symbols $X < Y$, if there is a set of parameters A over which both X and Y are defined such that $[X] \subseteq [Y]^\circ$ in $S_n(A)$.

Notice that the interior of $[Y]$ does depend on A (i.e., if $A' \supseteq A$ then $[Y]^\circ$ calculated in $S_n(A')$ may be larger than the pullback of the interior of $[Y]$ in $S_n(A)$). We may nonetheless choose any parameter set we wish:

Lemma 1.3. *Assume that X is type-definable with parameters in B , Y type-definable possibly with additional parameters not in B . Then:*

- (i) *If $X < Y$ then $[X] \subseteq [Y]^\circ$ in $S_n(A)$ for any set A over which both X and Y are defined.*
- (ii) *If $X < Y$ then there is an intermediate logical neighbourhood $X < Z < Y$, which can moreover be taken to be a zero of a formula with parameters in B .*
- (iii) *If $Y \cap X = \emptyset$ then there is a logical neighbourhood $Z > X$ such that $Z \cap Y = \emptyset$. Moreover, we may take Z to be a zero set defined over B .*

Proof. Assume $X > Y$, where X is type-definable over B , and Y over $A \supseteq B$. Let Φ consist of all formulae $\varphi(\bar{x})$ over B which are zero on X . If $\varphi, \psi \in \Phi$ then $\varphi \vee \psi \in \Phi$, and X is defined by the partial type $p(\bar{x}) = \{\varphi(\bar{x}) \leq r : \varphi \in \Phi, r > 0\}$. By a compactness in $S_n(A)$ there is a condition $\varphi(\bar{x}) \leq r$ in $p(\bar{x})$ which already implies $\bar{x} \in Y$. Let Z be the zero set of the formula $\varphi(\bar{x}) \div r'$ where $0 < r' = \frac{k}{2^{-m}} < r$.

Then in $S_n(A)$ we have $[X] \subseteq [\varphi(\bar{x}) < r'] \subseteq [\varphi(\bar{x}) \leq r'] \subseteq [\varphi(\bar{r}) < r] \subseteq [Y]$, i.e., $[X] \subseteq [Z]^\circ \subseteq [Z] \subseteq [Y]^\circ$, proving the first two items. The third item now follows from the fact that $S_n(A)$ is a normal topological space. ■_{1.3}

Lemma 1.4. *A type-definable set X is a zero set if and only if $[X]$ is a (closed) G_δ set.*

Proof. This is just a topological statement, saying that in a compact Hausdorff space Y , a closed subset $K \subseteq Y$ is the zero set of some $f \in C(Y, [0, 1])$ if and only if it is a G_δ set. This is in fact true in an arbitrary normal space: left to right is immediate, while right to left involves a straightforward construction using ω applications Urysohn's Lemma. ■_{1.4}

It follows that finite unions and countable intersections of zero sets are zero sets. In particular, a set X which is type-definable by a countable set of conditions is a zero set.

Later on we will use the following result:

Lemma 1.5. *Let $X = \prod_{i < \omega} X_i^{m_i}$ be a type-definable set, (so each X_i is one) and $Y \supseteq X$ a zero set. Then there are zero sets $Y_i \supseteq X_i$ such that $Y \supseteq \prod_{i < \omega} Y_i^{m_i}$.*

Proof. We will only show that if $X \times X'$ is a type-definable set and $Y \supseteq X \times X'$ is a zero set then there is a zero set $Y' \supseteq X'$ such that $Y \supseteq X \times Y'$. The result then follows since the intersection of finitely (or even countably) many zero sets is a zero set.

Let $\varphi(x, y) = 0$ define Y . For $n < \omega$ consider the partial type $\{x \in X\} \cup \{y \in X'\} \cup \{\varphi(x, y) \geq 2^{-n}\}$. As it is inconsistent Y admits a logical neighbourhood $Y_n > X'$ such that $\{x \in X\} \cup \{y \in Y_n\} \cup \{\varphi(x, y) \geq 2^{-n}\}$ is inconsistent. Moreover, choosing the sets Y_n by induction on n we may arrange that $Y_n > Y_{n+1}$. Let $Y' = \bigcap Y_n$. Then $[Y'] = \bigcap [Y_n]$ is a closed G_δ set, and $Y \supseteq X \times Y'$. ■_{1.5}

In many situations we may wish to show that if a condition holds on a type-definable set then it holds on some zero set containing it. This is (tautologically) the case if the condition itself is a zero set. It is still true if the condition in question is a containment (i.e., implication) of zero sets.

Lemma 1.6. *Let $\varphi(\bar{x})$ and $\psi(\bar{x})$ be two definable predicates, X a type-definable set, and assume that for all $\bar{x} \in X$: $\varphi(\bar{x}) = 0 \implies \psi(\bar{x}) = 0$. Then there exist a zero set $Y \supseteq X$ on which $\varphi(\bar{x}) = 0 \implies \psi(\bar{x}) = 0$ holds as well.*

Proof. For every $\varepsilon > 0$ there is $\delta(\varepsilon) > 0$ such that for all $\bar{x} \in X$: $\varphi(\bar{x}) < \delta(\varepsilon) \implies \psi(\bar{x}) \leq \varepsilon$. Indeed, if not, then we can obtain a contradiction to our assumption using compactness. We can therefore take Y to be the zero set of:

$$\chi(\bar{x}) = \sum_{n < \omega} 2^{-n-1} ((\delta(2^{-n}) \div \varphi(\bar{x})) \wedge (\psi(\bar{x}) \div 2^{-n})). \quad \text{■}_{1.6}$$

Finally, when it comes to definable sets, we will use the following equivalent characterisations:

Fact 1.7. *Let M be a model, $X \subseteq M^n$ closed subset. Then the following are equivalent:*

- (i) X is definable.
- (ii) X is the zero set of a definable predicate $\varphi(\bar{x})$, and $d(\bar{x}, X) \leq \varphi(\bar{x})$.
- (iii) For every $\varepsilon > 0$, the set $\bar{B}(X, \varepsilon)$ (which is always type-definable) is a logical neighbourhood of X .
- (iv) For every definable predicate $\varphi(\bar{x}, \bar{y})$, the predicate $\sup_{\bar{x} \in X} \varphi(\bar{x}, \bar{y})$ is definable.
- (v) For every formula $\varphi(\bar{x}, \bar{y})$, the predicate $\sup_{\bar{x} \in X} \varphi(\bar{x}, \bar{y})$ is definable.

Lemma 1.8. *The product and union of two definable sets are definable.*

Proof. If X and Y are definable then $d((x, y), X \times Y) = d(x, X) \vee d(y, Y)$, where we equip the product sort with the maximum metric. If they are in the same sort then $d(x, X \cup Y) = d(x, X) \wedge d(x, Y)$. $\blacksquare_{1.8}$

A finite intersection of definable sets needs not be definable in general. When it comes to infinite unions, we propose two results.

Lemma 1.9. *Let X be a type-definable set, and assume $X = \bigcup_{i < \alpha} X_i$ where $\{X_i : i < \alpha\}$ is a possibly infinite (yet bounded) family of definable sets. Then X is definable.*

Proof. By Fact 1.7 all we need to check is that for every $\varepsilon > 0$ the set $\bar{B}(X, \varepsilon)$ is a logical neighbourhood of X . Indeed:

$$\bar{B}(X, \varepsilon) \supseteq B(X, \varepsilon) = \bigcup_{i < \alpha} B(X_i, \varepsilon) \supseteq \bigcup_{i < \alpha} \bar{B}(X_i, \varepsilon/2),$$

whereby:

$$[\bar{B}(X, \varepsilon)]^\circ \supseteq \bigcup_{i < \alpha} [\bar{B}(X_i, \varepsilon/2)]^\circ \supseteq \bigcup_{i < \alpha} [X_i] = [X]. \quad \blacksquare_{1.9}$$

Similarly, a definable union of definable sets is definable:

Lemma 1.10. *Let $X_{\bar{a}}$ be a family of uniformly definable sets with parameters in a definable set Y . That is to say that there is a definable predicate $\varphi(\bar{x}, \bar{y})$ such that $d(\bar{x}, X_{\bar{a}}) = \varphi(\bar{x}, \bar{a})$ for every $\bar{a} \in Y$. Then $Z = \bigcup_{\bar{a} \in Y} X_{\bar{a}}$ is definable.*

Proof. First, Z is a closed set by a simple compactness argument. Then we have: $d(\bar{x}, Z) = \inf_{\bar{y} \in Y} d(\bar{x}, X_{\bar{y}}) = \inf_{\bar{y} \in Y} \varphi(\bar{x}, \bar{y})$. $\blacksquare_{1.10}$

Finally, every compact set is definable, and in a particularly convenient manner:

Lemma 1.11. *Let M be a model, $K \subseteq M$ a (metrically) compact set. Then K is definable. Moreover, if $\varphi(\bar{x}, y)$ is a definable predicate then $\inf_{y \in K} \varphi(\bar{x}, y)$ is equal to a forced limit of continuous combinations of instances of φ with parameters in K (i.e., it is a φ -predicate over K) and for any tuple \bar{x} , the infimum is attained by some $y \in K$.*

Proof. As K is compact we can find a sequence $(c_i : i < \omega) \subseteq K$ such that for every $\varepsilon > 0$ there is $m(\varepsilon)$ such that $K \subseteq \bigcup_{i < m(\varepsilon)} B(c_i, \varepsilon)$. Let $\varepsilon \mapsto \delta(\varepsilon)$ be the uniform continuity modulus of $\varphi(\bar{x}, y)$ with respect to y , meaning that $d(y, y') < \delta(\varepsilon) \implies |\varphi(\bar{x}, y) - \varphi(\bar{x}, y')| \leq \varepsilon$. Then

$$\inf_{y \in K} \varphi(\bar{x}, y) = \mathcal{F}\lim_n \bigwedge_{i < m(\delta(2^{-n}))} \varphi(\bar{x}, c_i).$$

Finally, the infimum of a continuous function on a compact set is always attained. $\blacksquare_{1.11}$

1.2. Partial definable predicates and functions. We will consider objects such as predicates and functions which are only defined on some type-definable set.

Definition 1.12. Let X be a type-definable set.

- (i) A *partial type-definable predicate* $\psi(\bar{x})$ on X , is given by a continuous mapping $\psi: [X] \rightarrow [0, 1]$, where $[X]$ is the closed set of complete types corresponding to X . We call X the *domain* of ψ , denoted $\text{dom}(\psi)$, and for $\bar{a} \in X$ we denote by $\psi(\bar{a})$ the truth value $\psi(\text{tp}(\bar{a}))$.
- (ii) It is *definable* on X if it is the restriction of a definable predicate to X .
- (iii) A *partial type-definable function* $f(\bar{x})$ on X is one whose graph is given by a partial type $G_f(\bar{x}, y)$. That is to say that $G_f(\bar{x}, y) \cup G_f(\bar{x}, z) \models y = z$ and that $\text{dom}(f) = X$ is defined by the partial type $\exists y G_f(\bar{x}, y)$. With a slight abuse of notation we may write it as $f: X \rightarrow M$ or $f: X(M) \rightarrow M$, although the model M (and even its complete theory) may vary.
- (iv) It is *definable* on X if the predicate $d(f(\bar{x}), y)$ is definable on $X \times M$.

We may allow parameters in the definitions by naming them in the language. Also, when explicitly specifying the domain we will usually drop the qualifier “partial”.

It is a straightforward exercise to verify that every partial definable predicate or function is type-definable. We wish to prove the converse.

We start with a fact from general topology concerning the extensions of continuous functions.

Fact 1.13 (Tietze’s Extension Theorem). *Let X be a normal space, $C \subseteq X$ closed, and let $f: C \rightarrow [0, 1]$ be a continuous function. Then f admits an extension to a continuous function $\tilde{f}: X \rightarrow [0, 1]$.*

Lemma 1.14. *Let $\psi(x_{<n}, y)$ be a definable predicate, and $f: X \rightarrow M$ a type-definable partial n -ary function. Then $\psi(\bar{x}, f(\bar{x}))$ is a partial type-definable predicate on X .*

Proof. Assume everything is defined without parameters, and let $K = [X] \subseteq S_n(T)$. Then f induces a continuous function $\hat{f}: K \rightarrow S_{n+1}(T)$ sending $\text{tp}(\bar{a}) \mapsto \text{tp}(\bar{a}, f(\bar{a}))$, and $\psi(\bar{x}, f(\bar{x}))$ is given by the composition $\psi \circ \hat{f}: K \rightarrow [0, 1]$. ■_{1.14}

Proposition 1.15. *Every partial type-definable predicate or function is definable.*

Proof. For predicates, this is just Tietze’s Extension Theorem. For functions, assume that $f: X \rightarrow M$ is a type-definable function. Then $f': X \times M \rightarrow M$ defined by $f'(\bar{x}, y) = f(\bar{x})$ is type-definable as well. Let $\psi(\bar{x}, y, z) = d(y, z)$. Then $\psi(\bar{x}, y, f'(\bar{x}, y)) = d(y, f(\bar{x}))$ is a type-definable, and therefore definable, predicate on $X \times M$. ■_{1.15}

This means that of the notions defined above we only need to retain those of partial definable functions and predicates. Moreover, while partial definable predicates are not entirely superfluous, most of the time we will avoid them, replacing any such partial predicate with an (arbitrary) total definable predicate extending it.

Notation 1.16. By the notation $\varphi(\bar{x}) \sqsupseteq \psi(\bar{x})$ we mean that $\varphi(\bar{x})$ is a definable predicate extending a partial definable predicate $\psi(\bar{x})$.

Remark 1.17. One can add to the language a sort $S_{[0,1]}$ for the interval $[0, 1]$, along with the tautological predicate $\text{id}_{[0,1]}: S_{[0,1]} \rightarrow [0, 1]$ (such a sort will usually exist in \mathcal{L}^{eq}). Then a partial predicate on a type-definable set X is definable if and only if factors through a definable function to $S_{[0,1]}$.

If $Y \subseteq X$ are type-definable, we may say that Y is definable *relative to* X if $d(x, Y)$ is a partial definable predicate on X .

Lemma 1.18. *If X is definable and $Y \subseteq X$ is definable relative to X then Y is definable.*

Proof. Let $\varphi(x) \sqsupseteq d(x, Y)$ (where $d(x, Y)$ is defined on X). Then for all x : $d(x, Y) = \inf_{y \in X} (d(x, y) \dot{+} \varphi(y))$. ■_{1.18}

Since we do not always know how to extend a partial definable function to a total one, it is worthwhile to notice the following fact:

Lemma 1.19. *Let X be a definable set and let f be a partial definable function whose domain contains X . Then $f(X)$ is definable as well.*

Proof. Let $\varphi(x, y) \sqsupseteq d(x, f(y))$. Then $d(x, f(X)) = \inf_{y \in X} \varphi(x, y)$, which is a definable predicate since X is definable, whereby $f(X)$ is. ■_{1.19}

It is usually fairly easy to reduce questions about arbitrary type-definable sets to questions about zero sets. For example, while it is not always possible to extend a partial definable function to a total one, one can always extend it to a zero set containing its domain:

Lemma 1.20. *Let X be a type-definable set, $f: X \rightarrow M$ a definable function on X . Then there is a zero set $Y \supseteq X$ such that f extends to a definable function on Y .*

Moreover, for every choice of $\varphi(\bar{x}, y) \sqsupseteq d(f(\bar{x}), y)$ there is a zero set $Y \supseteq X$ such that $\varphi|_{Y \times M}$ defines the graph of a partial definable function $f': Y \rightarrow M$ (which extends f).

Proof. Let $\varphi(\bar{x}, y) \sqsupseteq d(f(\bar{x}), y)$, and let:

$$\psi(\bar{x}, y, z) = (d(y, z) \dot{+} \varphi(\bar{x}, y) \dot{+} \varphi(\bar{x}, z)) \vee (\varphi(\bar{x}, y) \dot{+} d(y, z) \dot{+} \varphi(\bar{x}, z)) \vee \inf_t \varphi(\bar{x}, t)$$

Then ψ is zero on $X \times M^2$, and by Lemma 1.5 there is a zero set $Y \supseteq X$ such that ψ is zero on $Y \times M^2$ as well. This means that for all $\bar{x} \in Y$ there is y_0 such that $\varphi(\bar{x}, y_0) = 0$, and that for any other y one has $\varphi(\bar{x}, y) = d(y, y_0)$. Thus $\varphi(\bar{x}, y) = d(f'(\bar{x}), y)$ for some function $f': Y \rightarrow M$ extending f . ■_{1.20}

If the original type-definable set is closed under the function(s) we wish to extend, we can make sure that so is the extension:

Proposition 1.21. *Assume that we are given:*

- (i) For each $i < n$ sets $X_i \subseteq Y_i$ where X_i is type-definable and Y_i is a zero set.
- (ii) An ordinal $\alpha \leq \omega$, and for each $j < \alpha$ a partial definable function $f_j: \prod_{i < n} X_i^{m_{i,j}} \rightarrow X_{\ell_j}$.

Then there are zero sets $Y_i \supseteq Z_i \supseteq X_i$ and partial definable functions $g_j: \prod_{i < n} Z_i^{m_{i,j}} \rightarrow Z_{\ell_j}$ extending f_j .

Moreover, if we are given definable predicate $\varphi_j(\bar{x}_j, \bar{y}_j) \supseteq d(f_j(\bar{x}_j), y_j)$ then we can arrange that $\varphi_j(\bar{x}_j, y_j) \supseteq d(g_j(\bar{x}_j), y_j)$ as well.

Proof. For simplicity of notation we will consider the special case of a single function $f: X^m \rightarrow X \subseteq M$, as the general case is identical (but with a lot more indexes).

Let $\varphi(\bar{x}, y) \supseteq d(f(\bar{x}), y)$ be given (or just choose one). By Lemma 1.20, and possibly replacing the given zero set $Y \supseteq X$ with a smaller zero set, we may assume there is a partial definable function $g: Y^m \rightarrow M$ such that $\varphi(\bar{x}, y) \supseteq g(f(\bar{x}), y)$.

Let $Y_0 = Y$. Given a zero set Y_k satisfying $Y \supseteq Y_k \supseteq X$ let $W_k = g^{-1}(Y_k) \cap Y_k^m$. Then $W_k \supseteq X^m$ is a zero set and by Lemma 1.5 we can find a zero set $Y_{k+1} \supseteq X$ such that $W \supseteq Y_{k+1}^m$, i.e., such that $Y_{k+1} \subseteq Y_k$ and $g(Y_{k+1}^m) \subseteq Y_k$.

In this manner we construct a countable decreasing sequence of zero sets $Y = Y_0 \supseteq Y_1 \supseteq \dots \supseteq Y_k \supseteq \dots \supseteq X$ such that $g(Y_{k+1}^m) \subseteq Y_k$. Then $Z = \bigcap_k Y_k$ is a zero set, $Y \supseteq Z \supseteq X$ and $g(Z^m) \subseteq Z$, as required. ■_{1.21}

On the other hand, in later section we will have to consider logical neighbourhoods of domains of partial definable functions (specifically: logical neighbourhoods of type-definable groups, on which the group law is a partial definable function). While the function does not necessarily extend as such, we can extend it as a multi-valued function, which in addition will be approximately well-defined on small enough neighbourhoods of the original domain.

Lemma 1.22. *Let X be a type-definable set, $f: X \rightarrow M$ a partial definable function. Then there is a definable predicate $\varphi_f(\bar{x}, y) \supseteq d(f(\bar{x}), y)$ satisfying in addition $\sup_{\bar{x}} \inf_y \varphi_f(\bar{x}, y) = 0$. Letting $\tilde{f}(\bar{x}) = \{y: \varphi_f(\bar{x}) = 0\}$ we have:*

- (i) For all $\bar{x}: \tilde{f}(\bar{x}) \neq \emptyset$;
- (ii) If $\bar{x} \in X$ then $\tilde{f}(\bar{x}) = \{f(\bar{x})\}$;
- (iii) For every $\varepsilon > 0$ there is a logical neighbourhood $Y > X$ such that for all $\bar{x} \in Y$: $\text{diam}(\tilde{f}(\bar{x})) \leq \varepsilon$.

Proof. First choose any $\varphi_{f,0}(\bar{x}, y) \supseteq d(f(\bar{x}), y)$. Then define: $\varphi_f(\bar{x}, y) = \varphi_{f,0}(\bar{x}, y) - \inf_z \varphi_{f,0}(\bar{x}, z)$. As $\bar{x} \in X$ implies $\inf_z \varphi_{f,0}(\bar{x}, z) = 0$ we still have $\varphi_f(\bar{x}, y) \supseteq d(f(\bar{x}), y)$ whence the second item. On the other hand $\sup_{\bar{x}} \inf_y \varphi_f(\bar{x}, y) = 0$ follows from the definition and implies the first item.

Finally, consider the partial type saying that $\{\bar{x} \in X\} \cup \{y, z \in \tilde{f}(\bar{x})\} \cup \{d(y, z) \geq \varepsilon\}$. It is inconsistent by the second item, so (by compactness) $\{\bar{x} \in Y\} \cup \{y, z \in \tilde{f}(\bar{x})\} \cup \{d(y, z) \geq \varepsilon\}$ is inconsistent for some logical neighbourhood $Y > X$. ■_{1.22}

In classical logic every partial definable function on a type-definable set X can be extended to a definable set containing X . Whether this is true in continuous logic is still open, but we can show this up to an extension of the sort on which the function acts via an isometric embedding in a larger sort.

Lemma 1.23. *Let S_1 and S_2 be two (imaginary) sorts, $X \subseteq S_1$ type-definable, and $f: X \rightarrow S_2$ a partial definable function (usually we would have $S_1 = M^n$ and $S_2 = M^m$ being powers of the home sort). Then:*

- (i) *We can find $\varphi(x, y) \sqsupseteq d(f(x), y)$ satisfying:*
 - (1) $\varphi(x, y) - \varphi(x, y') \leq d(y, y') \leq \varphi(x, y) + \varphi(x, y')$.
- (ii) *There exists an imaginary sort S_3 , a definable isometric embedding $\theta: S_2 \hookrightarrow S_3$, and a total definable function $\hat{f}: S_1 \rightarrow S_3$ extending $\theta \circ f$, such that moreover $\varphi(x, y) = d(\hat{f}(x), \theta(y))$.*

Proof. Choose $\varphi_0(x, y) \sqsupseteq d(f(x), y)$ and let $\varphi(x, y) = \sup_z |\varphi_0(x, z) - d(z, y)|$. Notice that for $x \in X$ we have $\varphi(x, y) = \sup_z |d(f(x), z) - d(z, y)| = d(f(x), y)$, i.e., $\varphi(x, y) \sqsupseteq d(f(x), y)$ as well.

Let $\chi(t)$ be a formula for which 0 and 1 are possible truth values, and let

$$\psi(z, xyt) = \chi(t) \wedge \varphi_0(x, z) + (\neg\chi(t)) \wedge d(y, z).$$

Let $S_3 = \{[xyt]: x, t \in S_1, y \in S_2\}$ be the sort of canonical parameters of instances of $\psi(z, xyt)$. We recall that the metric on it is given by: $d([xyt], [x'y't']) = \sup_z |\psi(z, xyt) - \psi(z, x'y't')|$.

The embedding $\theta: S_2 \hookrightarrow S_3$ is given by $y \mapsto [x_0y\tilde{0}]$ where $\chi(\tilde{0}) = 0$ and $x_0 \in S_1$ is arbitrary. This is indeed isometric and does not depend on the choice of either $\tilde{0}$ or x_0 as $\psi(z, x_0y\tilde{0}) = d(y, z)$. Thus $d(\theta(y), \theta(y')) = \sup_z |d(y, z) - d(y', z)| = d(y, y')$.

Similarly we define $\hat{f}: S_1 \rightarrow S_3$ by $x \mapsto [xy_0\tilde{1}]$ where $\chi(\tilde{1}) = 1$, obtaining $\psi(z, xy_0\tilde{1}) = \varphi_0(x, z)$. In particular for $x \in S_1$ and $y \in S_2$ we have $d(\hat{f}(x), \theta(y)) = \varphi(x, y)$ as required, from which follows (1).

Finally, if $x \in X$ then $d(\theta \circ f(x), \theta(y)) = d(f(x), y) = \varphi(x, y) = d(\hat{f}(x), \theta(y))$, so \hat{f} extends $\theta \circ f$ and the proof is complete. ■_{1.23}

Question 1.24. Given a type-definable set X and a partial definable function $f: X \rightarrow M$, can one find a definable set $Y \supseteq X$ such that f extends to a definable $\hat{f}: Y \rightarrow M$? This is true in classical logic.

2. THE MAIN THEOREM (FIRST APPROXIMATION)

Once we have defined (type-)definable sets and functions we automatically have corresponding notions for more complex algebraic structures. For example:

Definition 2.1. By a *type-definable group* we mean a type-definable set G equipped with a definable function $\cdot: G^2 \rightarrow G$ defining a group law on G .

It is *definable* if the set G is definable.

Notice that if $\langle G, \cdot \rangle$ is a type-definable group with parameters in A then its identity e_G belongs to $\text{dcl}(A)$: indeed, every automorphism of the universal domain fixing A would fix e_G . Similarly, the function $x \mapsto x^{-1}$ is definable on G with parameters in A , its graph given by $G_{x^{-1}}(x, y) = G.(x, y, e_G)$.

The main goal of this article is to give some sufficient conditions under which type-definable groups are definable. For example, we have already essentially proved:

Proposition 2.2. *Let G be a type-definable group, $H \leq G$ a subgroup of bounded index, and assume H is definable. Then so is G .*

Proof. Say $G/H = \{g_i H : i < \lambda\}$. Each coset $g_i H$ is definable as the image of H under the partial definable function $x \mapsto g_i x$ (Lemma 1.19). Then by Lemma 1.9 $G = \bigcup_{i < \lambda} g_i H$ implies that G is definable. ■_{2.2}

2.1. Invariant metrics. The main theorem states that in an \aleph_0 -stable theory, every type-definable group is definable. As a first approximation, we will prove this under the assumption that the metric is invariant under the group operation:

Definition 2.3. A metric defined on a type-definable group is *left-invariant* (*right-invariant*) if it is invariant under left (right) translation. It is *invariant* if it is both left- and right-invariant. It is *inverse-invariant* if it is invariant under $x \mapsto x^{-1}$.

Clearly if the metric is left-invariant (or right-invariant) and inverse-invariant then it is right-invariant (or left-invariant) and therefore invariant. Conversely, if the metric is invariant then it is in particular inverse-invariant, as we have: $d(x^{-1}, y^{-1}) = d(x^{-1}y, e) = d(y, x) = d(x, y)$.

In classical first order logic it is an easy consequence of compactness that on every type-definable group, the group law and inverse can be extended to be well-defined (and make some sense) on some definable set containing the group, i.e., on some logical neighbourhood of the group. In the continuous sense things are trickier, and the best we can hope for is a logical neighbourhood on which an *approximate* product is *approximately* well-defined and (in case the metric on the group is invariant) approximately isometric.

Lemma 2.4. *Let G be a type-definable group on which the metric is invariant. Let $\varphi.(x, y, z) \supseteq d(xy, z)$ be as in Lemma 1.22, $x \tilde{\cdot} y = \{z : \varphi.(x, y, z) = 0\}$.*

Then for every $\varepsilon > 0$ and logical neighbourhood $X > G$ there is an intermediate logical neighbourhood $X > Y > G$ such that $Y \cdot Y \subseteq X$, and multiplication is almost isometric in the sense that for all $x, y, y' \in Y$ and for all $z \in x \tilde{\cdot} y$, $z' \in x \tilde{\cdot} y'$ (or $z \in y \tilde{\cdot} x$, $z' \in y' \tilde{\cdot} x$): $|d(y, y') - d(z, z')| \leq \varepsilon$.

Proof. Since $X > G$ we can find a formula $\chi_{G,X}(x)$ which is equal to 0 on G and to 1 outside X . Then all the following partial types are contradictory:

$$\begin{aligned} x, y \in G, z \in x\tilde{y}, \chi_{G,X}(z) &= 1, \\ x, y, y' \in G, z \in x\tilde{y}, z' \in x\tilde{y}', |d(y, y') - d(z, z')| &\geq \varepsilon, \\ x, y, y' \in G, z \in y\tilde{x}, z' \in y'\tilde{x}, |d(y, y') - d(z, z')| &\geq \varepsilon. \end{aligned}$$

By compactness there exist a logical neighbourhood $X > Y > G$ such that they are still contradictory when G is replaced everywhere with Y , and this Y will do. $\blacksquare_{2.4}$

2.2. Cantor-Bendixson and Morley ranks. We will use \aleph_0 -stability via Morley ranks, i.e., Cantor-Bendixson ranks in $S_n(M)$ where M is a sufficiently saturated model. Such ranks were studied in detail in [Ben] in the general setting of topometric spaces. In fact, that paper discusses several possible notions of Cantor-Bendixson ranks, of which we will use one.

Notation 2.5. In this paper, CB_ε will denote what is denoted in [Ben] by $\text{CB}_{f,\varepsilon}$, i.e., the Cantor-Bendixson rank based on the removal of open ε -finite sets.

Definition 2.6. Let X be a type-definable set of n -tuples. We define the ε -Morley rank of X as $\text{RM}_\varepsilon(X) = \text{CB}_\varepsilon^{S_n(M)}(X)$, where M is any sufficiently saturated model containing the parameters needed for X (this does not depend on the choice of M). If $\alpha = \text{RM}_\varepsilon(X)$ then $[X] \cap S_n(M)_\varepsilon^{(\alpha)}$ is ε -finite, i.e., ε - d -finite for some d , and we define the ε -Morley degree of X to be $\text{dM}_\varepsilon(X) = d$.

Lemma 2.7. Let X be a type-definable set. For all $r > 0$ there are $0 < \varepsilon < r' < r$ such that $\text{RM}_{r'}(X) = \text{RM}_{r'-\varepsilon}(X)$.

Proof. For $n < \omega$ define $r_n = r(1 - 2^{-n-1})$. Then $(r_n : n < \omega)$ is an increasing sequence, whereby $\text{RM}_{r_n}(X)$ is a decreasing sequence of ordinals and therefore stabilises from some point onwards. Thus we may take $r' = r_{n+1}$ and $\varepsilon = r_{n+1} - r_n$ for any n large enough. $\blacksquare_{2.7}$

We recall some definitions and the main result we will use from [Ben, Section 3.3].

Notation 2.8. Let X, Y be two compact spaces, $R \subseteq X \times Y$ a closed relation. For $x \in X$ and $A \subseteq Y$ we define:

$$\begin{aligned} R_x &= \{y \in Y : (x, y) \in R\}, \\ R^{\forall A} &= \{x \in X : R_x \subseteq A\}, \\ R^{\exists A} &= \{x \in X : R_x \cap A \neq \emptyset\}. \end{aligned}$$

Fact 2.9. Let X, Y be two compact topometric spaces, $R \subseteq X \times Y$ a closed relation, and $\varepsilon, \delta > 0$ such that for all $(x, y), (x', y') \in R$: if $d_Y(y, y') \leq \delta$ then $d_X(x, x') \leq \varepsilon$. Let $K \subseteq X$ and $F \subseteq Y$ be closed sets such that $K \subseteq (R^{\exists Y})^\circ \cap R^{\forall F}$. Then $\text{CB}_\varepsilon^K(K) \leq \text{CB}_\delta^Y(F)$.

Proof. [Ben, Theorem 3.23]. ■_{2.9}

The following result contains the technical core of the proof of the (first approximation of the) main theorem. Given a logical neighbourhood of G on which product is well-behaved and a point in that neighbourhood (outside G), we can approximately translate G by that element, obtaining an approximately isometric copy of G . Using the Fact cited above we can compare the Morley ranks of G and of this approximate copy. In addition, if the element we translate by is far enough from G , then so will be the entire copy.

Lemma 2.10. *Let G be a group on which the metric is invariant.*

Then for every $\varepsilon > 0$ and logical neighbourhood $X > G$ there exists an intermediary logical neighbourhood $X > Y > G$ such that for all $r > \varepsilon$ either $Y \subseteq \bar{B}(G, r)$ or there is a type-definable subset $Z \subseteq X$ such that $d(Z, G) > r - \varepsilon$ and $\text{RM}_{r-\varepsilon}(Z) \geq \text{RM}_r(G)$.

Proof. Apply Lemma 2.4 to $X > G$ and $\varepsilon > 0$ to obtain $X > Y > G$. If $Y \subseteq \bar{B}(G, r)$ then we are done, so assume the contrary, i.e., that there exists $y_0 \in Y$ such that $d(y_0, G) > r$. Set $Z = y_0 \tilde{G} \subseteq X$, which is type-definable.

To see that $d(G, Z) > r - \varepsilon$, let $g \in G$ and $z \in Z$. Then $z \in y_0 \tilde{h}$ for some $h \in G$. Let $g' = gh^{-1} \in G$, so $d(y_0, g') > r$. Since $y_0, g', h \in Y$ we have $|d(z, g) - d(y_0, g')| \leq \varepsilon$, whereby $d(z, g) > r - \varepsilon$. By a compactness argument it follows that $d(G, Z) > r - \varepsilon$.

Let M be a fairly saturated model containing all necessary parameters, including y_0 , and let $S = S(M)$. Set

$$R = \{(\text{tp}(u/M), \text{tp}(v/M)) : u \in Y, v \in y_0 \tilde{u}\}.$$

Then $R \subseteq S^2$ is a closed relation, $[Y] = R^{\exists S}$ and $[G] \subseteq R^{\forall [Z]}$, so $[G] \subseteq (R^{\exists S})^\circ \cap R^{\forall [Z]}$. On the other, if $(p, q), (p', q') \in R$ then $|d(p, p') - d(q, q')| \leq \varepsilon$, so $d(q, q') \leq r - \varepsilon \implies d(p, p') \leq r$. Applying Fact 2.9 we get:

$$\text{RM}_r(G) = \text{CB}_r^S([G]) \leq \text{CB}_{r-\varepsilon}^S([Z]) = \text{RM}_{r-\varepsilon}(Z). \quad \text{■}_{2.10}$$

2.3. The definability proof. So far we haven't used yet the assumption of \aleph_0 -stability (the ranks in Lemma 2.10 could be infinite).

Theorem 2.11. *Let G be a type-definable group in an \aleph_0 -stable theory on which the metric is invariant. Then G is definable.*

Proof. Let $r > 0$, and we will show that $\bar{B}(G, r) > G$. There is no harm if we decrease r , so by Lemma 2.7 we may assume that $\text{RM}_r(G) = \text{RM}_{r-\varepsilon}(G)$ for some $0 < \varepsilon < r$.

By compactness we can find $X > G$ such that $\text{RM}_{r-\varepsilon}(G) = \text{RM}_{r-\varepsilon}(X)$ and $\text{dM}_{r-\varepsilon}(G) = \text{dM}_{r-\varepsilon}(X)$. Apply Lemma 2.10 to find $X > Y > G$ such that either $Y \subseteq \bar{B}(G, r)$ or there is $Z \subseteq Y$ such that $d(G, Z) > r - \varepsilon$ and $\text{RM}_{r-\varepsilon}(Z) \geq \text{RM}_r(G) = \text{RM}_{r-\varepsilon}(G)$.

In the second case we have $\text{RM}_{r-\varepsilon}(G) = \text{RM}_{r-\varepsilon}(Z) = \text{RM}_{r-\varepsilon}(X)$, so from $d(G, Z) > r - \varepsilon$ we obtain:

$$\begin{aligned} \text{dM}_{r-\varepsilon}(G) &= \text{dM}_{r-\varepsilon}(X) \geq \text{dM}_{r-\varepsilon}(G \cup Z) \\ &\geq \text{dM}_{r-\varepsilon}(G) + \text{dM}_{r-\varepsilon}(Z) \\ &> \text{dM}_{r-\varepsilon}(G). \end{aligned}$$

This is impossible, so $\bar{B}(G, r) \supseteq Y > G$, as required. ■_{2.11}

3. STABLE TYPE-DEFINABLE GROUPS

3.1. Generic elements and types in stable groups. Here we will generalise to continuous logic various classical facts concerning generic types and connected components of definable and type-definable groups in stable theories. The reader who is willing to accept that everything works more or less as in classical stable theories (see Fact 3.3 below) may safely skip it.

Let G be a type-definable group in a complete stable theory T . We will work inside a monster model $\bar{M} \models T$, and identify G with the set of elements $G(\bar{M})$. Naming the parameters used for defining G in the language we may assume G is type-definable without parameters. If $A \subseteq \bar{M}$, by $S_G(A)$ we denote the space of complete types over A of members of G , with the induced topology. The closed subsets of $S_G(A)$ correspond to subsets of G which are type-definable over A .

In order to prove the bare essentials needed in this paper we can use a general fact about the existence of generic elements and types for type-definable groups in thick simple cats. It uses the notion of *dividing-generic* types, which coincides with classical notions of genericity in the case of stable theories, and allows to extend the theory of generic elements in groups to simple ones (see Pillay [Pil98]).

Definition 3.1. Let G be a type-definable group in a simple (or stable) continuous theory, A any set of parameters, and $g \in G$. We say that g is *dividing-generic* over A , or that $\text{tp}(g/A) \in S_G(A)$ is *dividing-generic*, if whenever $h \in G$ and $g \perp_A h$ then $hg \perp h, A$.

For the rest of this section we will just say “generic” instead of “dividing-generic”.

Fact 3.2. *Let G be a type-definable group in a simple continuous theory (or in a stable one, which is a special case). Then:*

- (i) *Generic elements exist over every set.*
- (ii) *An element g is generic over A if and only if g^{-1} is. It follows that if g is generic over A and $g \perp_A h$ then $gh \perp h, A$ as well.*
- (iii) *An element g is generic over A if and only if it is generic over \emptyset and $g \perp A$.*
- (iv) *If $g, h \in G$, g is generic over A and $g \perp_A h$ then gh and hg are generic over A .*

Proof. This is shown in the more general setting of a type-definable group in a thick simple cat in [Ben03, Section 1.3]. ■_{3.2}

If M is any model, then the group $G(M)$ acts on $S_G(M)$. Indeed, if $g \in G(M)$ and $h \in G$ (i.e., $h \in G(\bar{M})$) then $\text{tp}(gh/M)$ only depends on g and on $\text{tp}(h/M)$, and we set $g \text{tp}(h/M) = \text{tp}(gh/M)$. Clearly, if $g' \in G(M)$ as well and $p \in S_G(M)$ then $g'(gp) = (g'g)p$ so this is a group action. Note that if p is generic then gp is generic as well.

Now let $N \succeq M$ be any elementary extension. If $h \models p \upharpoonright^N$ then $h \perp_M N$, whereby $g, h \perp_M N$ and finally $gh \perp_M N$, so $(gp) \upharpoonright^N = g(p \upharpoonright^N)$. In other words, G acts on the class of parallelism classes of types. For a stationary type p (say $p \in S_G(M)$) we define $\text{Stab}(p)$ as the stabiliser of the parallelism class of p under the action of G . Since p is definable, $\text{Stab}(p)$ is a type-definable subgroup of G , with parameters in $\text{Cb}(p)$. Again, the action sends parallelism classes of generic types to such. Since there is only a set of parallelism classes of generic types, $\text{Stab}(p)$ has bounded index in G for generic p .

Let $H \leq G$ be a type-definable subgroup, say over some model M , of bounded index in G . Then we can write $G = \bigcup_{i < \lambda} g_i H$ and choose $g \in G$ which is generic over $M \cup \{g_i : i < \lambda\}$. Then $g \in g_i H$ for some $i < \lambda$, so $g_i^{-1}g \in H$ is generic over M . In other words, there is some generic type p over $\text{acl}^{eq}(\emptyset)$ such that H contains a realisation of $p \upharpoonright^M$, and therefore all its realisations. It follows that H contains $\text{Stab}(p \upharpoonright^M) = \text{Stab}(p)$, which is of bounded index, so H can be written as a (bounded) union of cosets of $\text{Stab}(p)$. It follows that the collection of type-definable subgroups of bounded index over arbitrary parameters is a set (rather than a proper class), whose intersection is the smallest type-definable subgroup of bounded index, denoted G^0 and called the *connected component* of G . As G^0 is smallest among type-definable subgroups of bounded index over any set of parameters, it is automorphism invariant and thus type-definable without parameters. It is also clearly normal in G . Note however that G^0 is not necessarily the intersection of subgroups of finite index in G (some may therefore wish to use the notation G^{00} instead, and we understand them).

Let p be a stationary generic type, say over a set A , and let $M \supseteq A$ contain representatives of all cosets of G^0 . On the one hand $\text{Stab}(p) \geq G^0$ as it has bounded index. On the other hand, all realisations of $p \upharpoonright^M$ must lie in a single coset $g_0 G^0$ for some $g_0 \in M$, whereby all realisations of p must lie in the same coset. We obtain

$$G^0 \leq \text{Stab}(p) \subseteq \{gh^{-1} : g, h \models p\} \subseteq G^0,$$

so equality holds all the way. As G acts transitively on the cosets of G^0 , every such coset contains a generic type. Assume now that $p, p' \in S_G(M)$ are both generic in some coset and let $g \perp_M g'$ realise p and p' , respectively. Then $g^{-1}g' \in G^0 = \text{Stab}(p)$ and $g \perp_M g^{-1}g'$ whereby $g' \models p$. It follows that in each coset of G^0 there is a unique generic type.

Let q denote the unique generic type in G^0 over some model and let $p = q \upharpoonright_\emptyset$. Then q is a non forking extension of p , and as G^0 is automorphism-invariant, p must be stationary. If $g \models p$ then $g^{-1} \in G^0$ is also generic, whereby $g^{-1} \models p$ as well (one may say that

$p^{-1} = p$). We conclude that G^0 contains a unique generic type p and $G^0 = \{gh^{-1}: g, h \models p\} = \{gh: g, h \models p\}$.

We have been doing everything so far assuming G was type-definable over \emptyset . Allowing parameters we get:

Fact 3.3. *Let G be a type-definable group with parameters in some set A in a stable theory. Then G admits a smallest type-definable subgroup of bounded index G^0 , which contains a unique (up to parallelism) stationary generic type p of G . Moreover, p can be taken to be over A and G^0 can be recovered as the set $\{gh: g, h \models p\}$.*

It follows that G is connected (i.e., $G = G^0$) if and only if it has a unique generic type.

3.2. Chain conditions.

Definition 3.4. Let $\{X_n: n < \omega\}$ be a descending sequence of sets in a metric space. We say that the sequence *approximately stabilises* if for every $\varepsilon > 0$ there is $n_\varepsilon < \omega$ such that $X_{n_\varepsilon} \subseteq B(X_n, \varepsilon)$ for all n .

Lemma 3.5. *Assume $\{X_n: n < \omega\}$ is a descending sequence of closed subsets of a complete metric space (e.g., type-definable subsets of a continuous structure), and let $X = \bigcap_n X_n$. Then $\{X_n: n < \omega\}$ approximately stabilises if and only if for all $\varepsilon > 0$ there is n_ε such that $X_{n_\varepsilon} \subseteq B(X, \varepsilon)$.*

Proof. Right to left is immediate, so we prove left to right. By assumption for all m there is n_m such that $X_{n_m} \subseteq B(X_n, 2^{-m})$ for all n . Let m_0 be such that $2^{-m_0} < \varepsilon$. We claim that $n_\varepsilon = n_{m_0+1}$ will do. Indeed, let $x_0 \in X_{n_{m_0+1}}$. Given $x_k \in X_{n_{m_0+k+1}}$, let $x_{k+1} \in X_{n_{m_0+k+2}}$ be such that $d(x_k, x_{k+1}) < 2^{-m_0-k-1}$. Then the sequence $\{x_k\}$ converges to some point x , satisfying $d(x, x_k) < 2^{-m_0-k}$. Notice that for all n we have $x_k \in X_n$ for k large enough. Thus $x \in X_n$ for all n , whereby $x \in X$. We conclude that $d(x_0, X) < 2^{-m_0} < \varepsilon$, whence $X_{n_{m_0+1}} \subseteq B(X, \varepsilon)$ as required. $\blacksquare_{3.5}$

Lemma 3.6. *Let $\{X_n: n < \omega\}$ be a descending sequence of definable sets. If the chain approximately stabilises (in any model) then $X = \bigcap X_n$ is a definable. The converse holds in sufficiently saturated models.*

Proof. Assume first the sequence approximately stabilises. By Lemma 3.5 we have for all $\varepsilon > 0$: $X_{n_\varepsilon} \subseteq B(X, \varepsilon)$, whereby $X \subseteq X_{n_\varepsilon/2} \subseteq B(X_{n_\varepsilon/2}, \varepsilon/2) \subseteq B(X, \varepsilon)$ for all $\varepsilon > 0$. Since $X_{n_\varepsilon/2}$ is definable and this holds for all ε , X is definable.

Conversely, assume X is definable, and that the equality $X = \bigcap X_n$ holds in a sufficiently saturated model. Then by compactness, for all $\varepsilon > 0$ there is n_ε such that $\{x \in X_{n_\varepsilon}\} \cup \{d(x, X) \geq \varepsilon\}$ is contradictory, whereby $X_{n_\varepsilon} \subseteq B(X, \varepsilon)$. $\blacksquare_{3.6}$

Definition 3.7. We say that a class \mathcal{C} of sets satisfies the *metric descending chain condition (MDCC)* if every descending chain in \mathcal{C} approximately stabilises.

Theorem 3.8. *Let \mathcal{G} be a family of uniformly definable groups in a stable theory, meaning that there is a definable predicate $\varphi(x, y)$ such that for every $G \in \mathcal{G}$ the predicate $d(x, G)$ is an instance $\varphi(x, a_G)$. Then \mathcal{G} satisfies the MDCC.*

It follows that the intersection of a decreasing chain of uniformly definable groups in a stable theory is definable.

Proof. Assume $\{G_n : n < \omega\} \subseteq \mathcal{G}$ is a descending chain, and let $d(x, G_n) = \varphi(x, a_n)$. Then the sequence of definable predicates $\varphi(x, a_n)$ is increasing, and by definition of stability (no order property) it must converge uniformly.

Now let $\varepsilon > 0$. Then there is some n_ε such that $\varphi(x, a_n) < \varphi(x, a_{n_\varepsilon}) + \varepsilon$ for all n , which implies that $G_{n_\varepsilon} \subseteq B(G_n, \varepsilon)$ for all n . ■_{3.8}

Lemma 3.9. *Let G be a type-definable group in an \aleph_0 -stable structure and assume G admits an invariant metric. Then the class of type-definable subgroups of G satisfies the MDCC.*

Proof. Assume $\{G_n : n < \omega\} \subseteq \mathcal{G}$ is a descending chain of type-definable groups, and let $G = \bigcap G_n$, which is type-definable as well. Let $\varepsilon > 0$, and let $(\alpha, d) = (\text{RM}, \text{dM})_\varepsilon(G)$. First, G admits a logical neighbourhood $X \supset G$ such that $(\text{RM}, \text{dM})_\varepsilon(X) = (\alpha, d)$. Second, there exists n_ε such that $G_{n_\varepsilon} \subseteq X$, in which case $(\text{RM}, \text{dM})_\varepsilon(G_{n_\varepsilon}) = (\alpha, d)$. We claim that $G_{n_\varepsilon} \subseteq \bar{B}(G, \varepsilon)$. Indeed, if $g \in G_{n_\varepsilon}$ and $d(g, G) > \varepsilon$ then $d(gG, G) = d(g, G)$ (by invariance of the metric). Thus $\text{dM}_\varepsilon(G_{n_\varepsilon}) \geq \text{dM}_\varepsilon(G \cup gG) \geq 2 \text{dM}_\varepsilon(G)$, a contradiction. ■_{3.9}

In Section 5 we will use the main theorem to do away with the invariance assumption.

4. DEFINABLE METRICS

4.1. Extensions of partial definable metrics. Metrics and pseudo-metrics are a special (and distinguished) kind of predicates.

Definition 4.1. (i) A *definable (pseudo-)metric* is a definable predicate defining a (pseudo-metric).

(ii) Let X be a type-definable set. A *partial definable (pseudo-)metric* on X is a definable partial predicate on $X \times X$ defining a (pseudo-)metric there.

Again, when specifying the domain we will drop the “partial”. On the other hand, when wishing to make explicit that a definable metric is not partial, we will say it is *total*.

All type-definable metrics on a set X are essentially the same. This has first been proved in [Ben05], but since the setting is somewhat different and the proof is short we repeat it.

Lemma 4.2. *Let d_1 and d_2 be two definable metrics on a type-definable set X . Then they are uniformly equivalent. Moreover, if d_1 is a definable total metric and d_2 is the standard metric d on that sort, then replacing d with d_1 does not change the notion of a definable set.*

Proof. Indeed, Let $\varepsilon > 0$, and consider the partial type $\{x, y \in X\} \cup \{d_1(x, y) \geq \varepsilon\} \cup \{d_2(x, y) \leq 2^{-n} : n < \omega\}$. As d_2 is a metric this partial type is inconsistent, and by compactness there exists $n < \omega$ such that for all $x, y \in X$: $d_2(x, y) \leq 2^{-n} \implies d_1(x, y) < \varepsilon$. As this works for all $\varepsilon > 0$, and when exchanging the roles of d_1 and d_2 , they are uniformly equivalent.

For the moreover part, use the characterisation of definable sets via quantification in Fact 1.7, which does not depend on the metric. $\blacksquare_{4.2}$

While we know that every partial definable metric, say on X , extends to a total definable predicate, it is not clear whether it can extend to a total definable *metric*. We will give some partial results in this direction.

The best result is under the assumption that X is definable. We start with a few lemmas that allow us to reduce the case of a metric to the case of a pseudo-metric. In fact the reduction step holds more generally for zero sets.

Lemma 4.3. *Let X be any set of points, d_1 and d_2 two pseudo-metrics on X . Then $d_1 \vee d_2$ is a pseudo-metric on X .*

Proof. Let $d_3 = d_1 \vee d_2$. Clearly $d_3(x, x) = 0$ and $d_3(x, y) = d_3(y, x)$. For the triangle inequality, given $x, y, z \in X$ we may assume that $d_3(x, y) = d_1(x, y)$ in which case $d_3(x, y) \leq d_1(x, z) + d_1(y, z) \leq d_3(x, z) + d_3(y, z)$. $\blacksquare_{4.3}$

Lemma 4.4. *Let X be a zero set. Then there exists a definable pseudo-metric d_1 with the following properties:*

- (i) *For all $x, y \in X$: $d_1(x, y) = 0$.*
- (ii) *If $x \neq y$ and $x \notin X$ then $d_1(x, y) > 0$.*

It follows that if d_2 is a definable pseudo-metric whose restriction to X is a metric then $d_1 \vee d_2$ is a definable metric which agrees with d_2 on X .

Proof. Let X be the zero set of $\varphi(x)$. Define:

$$d_1(x, y) = \sup_z |\varphi(x) \wedge d(x, z) - \varphi(y) \wedge d(y, z)|.$$

Clearly, d_1 is a definable pseudo-metric and $x, y \in X \implies d_1(x, y) = 0$. On the other hand, if $x \neq y$ and $x \notin X$ then $d_1(x, y) \geq \varphi(x) \wedge d(x, y) > 0$.

Finally, assume that d_2 is a definable pseudo-metric which is a metric on X . By Lemma 4.3, $d_1 \vee d_2$ is a (definable) pseudo-metric. The hypotheses now imply it is a metric and agrees with d_2 on X . $\blacksquare_{4.4}$

Proposition 4.5. *Let X be a definable set, d_1 a definable (pseudo-)metric on X . Then d_1 extends to a definable (pseudo-)metric.*

Proof. Choose $\psi_1(x, y) \sqsupseteq d_1(x, y)$, and define:

$$d_2(x, y) = \sup_{z \in X} |\psi_1(x, z) - \psi_1(y, z)|$$

Then d_2 is a definable pseudo-metric extending d_1 . In case d_1 is a metric, d_2 is a definable pseudo-metric whose restriction to X is a metric, so by Lemma 4.4 there is a definable metric d_3 extending d_1 . ■_{4.5}

We now move to considering the extension problem for partial definable (pseudo-)metrics on type-definable sets. We first observe this can be reduced to zero sets:

Lemma 4.6. *Assume X is a type-definable set and d_1 a definable (pseudo-)metric on X . Then there exists a zero set $Y \supseteq X$ such that d_1 extends to a definable (pseudo-)metric d_2 on Y .*

Proof. In fact we will prove something slightly stronger, namely that if $\psi_1 \sqsupseteq d_1$ is any definable predicate extending d_1 then there exists a zero set $Y \supseteq X$ such that ψ_1 defines a (pseudo-)metric on Y .

For this, let $\psi_2(x, y, z) = \psi_1(x, x) \vee (\psi_1(x, y) \div \psi_1(x, z) \div \psi_1(y, z))$. Then ψ_2 is zero on X^3 , and by Lemma 1.5 there is a zero set $Y \supseteq X$ such that ψ_2 is zero on Y^3 .

In case d_1 is a metric we have $d_1(x, y) = 0 \implies d(x, y) = 0$ on X^2 . By Lemma 1.6 (and Lemma 1.5) there is a zero set $Y' \supseteq X$ such that $d_1(x, y) = 0 \implies d(x, y) = 0$ for $x, y \in Y'$. Then ψ_1 defines a metric on the zero set $Y \cap Y' \supseteq X$. ■_{4.6}

Joining Lemma 4.6 with Lemma 4.4, we reduce the extension problems for partial metrics to the extension of pseudo-metrics on zero sets. Unfortunately, we do not know whether it is always true and suspect that in its full generality it may fail. The best approximation we have is:

Lemma 4.7. *Let X be a zero set, d_1 a definable pseudo-metric on X . Then there exists a decreasing sequence of definable pseudo-metrics $d_{2,n}$ which converges uniformly to d_1 on X .*

Proof. Let X be the zero set of φ , and let $\psi_1(x, y) \sqsupseteq d_1(x, y)$. For $n < \omega$ let $\varphi_n(x) = 1 \div 2^n \varphi(x)$: thus $\varphi(x) = 0 \implies \varphi_n(x) = 1$, and $\varphi(x) \geq 2^{-n} \implies \varphi_n(x) = 0$. Now let:

$$d_{2,n}(x, y) = \sup_z |\varphi_n(z) \wedge \psi_1(x, z) - \varphi_n(z) \wedge \psi_1(y, z)|.$$

Clearly $d_{2,n}$ is a definable pseudo-metric for every n and $d_{2,n} \geq d_{2,n+1}$. In addition, if $x, y \in X$ then choosing $z = x$ we see that $d_{2,n}(x, y) \geq d_1(x, y)$.

Assume towards a contradiction that the pseudo-metrics $d_{2,n}$ do not converge uniformly to d_1 on X . Then there is some $\varepsilon > 0$ such that for all $n < \omega$ there are $a_n, b_n \in X$ and c_n such that $|\varphi_n(c_n) \wedge \psi_1(a_n, c_n) - \varphi_n(c_n) \wedge \psi_1(b_n, c_n)| \geq d_1(a_n, b_n) + \varepsilon$. This implies in particular that $\varphi(c_n) \leq 2^{-n}$ and $|\psi_1(a_n, c_n) - \psi_1(b_n, c_n)| \geq d_1(a_n, b_n) + \varepsilon$. By compactness there are $a, b, c \in X$ such that $|\psi_1(a, c) - \psi_1(b, c)| \geq d_1(a, b) + \varepsilon$, which is impossible as ψ_1 agrees with d_1 on X and d_1 is a pseudo-metric. ■_{4.7}

Question 4.8. Let X be a zero set, d_1 a definable pseudo-metric on X . Can d_1 always be extended to a total definable pseudo-metric? Are there additional assumptions on the theory (e.g., stability or \aleph_0 -stability) under which this is true?

As it turns out, Lemma 4.7 is very useful for our purposes. We will look for an alternative weakening of the metric extension problem: given a type-definable set X equipped with a definable metric d_1 , we will look for a total definable metric d_2 which, although does not extend d_1 , does preserve its invariance properties.

For this purpose we will use a few technical results from [Ben05, Section 2.3]. There we proved that in a Hausdorff cat in which there are not too many ways for two elements to be distinct (technically: the *distance cofinality* is at most countable), there is a definable metric. For this purpose one first constructs a symmetric definable predicate $\varphi(x, y)$ (a definable function in the terminology of [Ben05], and denoted there by $h(x, y)$) satisfying $\varphi(x, y) = 0 \iff x = y$. Then one needs to replace φ with one which also satisfies the triangle inequality.

In the of an *open* Hausdorff cat, one can just define $d(x, y) = \sup_z |\varphi(x, z) - \varphi(y, z)|$, which is definable metric. In the general case the metric $\sup_z |\varphi(x, z) - \varphi(y, z)|$ needs not be definable, and a more complicated construction is required to extract a definable metric from φ directly without recourse to quantification. The same tools apply here.

Fact 4.9. *Let $g: [0, 1]^2 \rightarrow [0, 1]$ be symmetric, non decreasing, and satisfy for all $u, w, t \in [0, 1]$: $g(0, t) = t$ and if $g(u, w) < t$ then there is $u < v \leq 1$ such that $g(v, w) < t$.*

Then there is a function $f: D \rightarrow [0, 1]$, where $D = \{k2^{-n}: n < \omega, 0 \leq k \leq 2^n\}$ denotes the set of all dyadic fractions in $[0, 1]$, such that:

- (i) *f is strictly increasing.*
- (ii) *$f \leq \text{id}_D$.*
- (iii) *For every $t, u \in D \cap (0, 1]$: $t + u \leq 1 \implies g(f(t), f(u)) < f(t + u)$.*

Proof. This is [Ben05, Lemma 2.19], with the sole difference being that we require $f(t) \leq t$ for all $t \in D$ rather than only for $t = \frac{1}{2^n}$. The only modification in the proof is in the construction of $f(\frac{k}{2^n})$ for odd $3 \leq k < 2^n$. There we have $f(\frac{k-1}{2^n}) \leq \frac{k-1}{2^n} < \frac{k}{2^n}$ by the induction hypothesis, and $f(\frac{k-1}{2^n}) < \min\{s', f(\frac{k+1}{2^n})\}$ as in the original proof, so we can choose $f(\frac{k}{2^n})$ such that $f(\frac{k-1}{2^n}) < f(\frac{k}{2^n}) < \min\{s', f(\frac{k+1}{2^n}), \frac{k}{2^n}\}$. ■_{4.9}

The following result was implicit in [Ben05, Section 2.3] (modulo the slight improvement above to the requirement on f):

Proposition 4.10. *Let $\varphi(x, y)$ be a symmetric and reflexive definable predicate, by which we means that $\varphi(x, y) = \varphi(y, x)$ and $\varphi(x, x) = 0$. Then there is a continuous increasing function $h: [0, 1] \rightarrow [0, 1]$, satisfying $h(0) = 0$ and $h(x) \geq x$, such that $h \circ \varphi(x, y)$ is a definable pseudo-metric. If in addition $x \neq y \implies \varphi(x, y) > 0$ then $h \circ \varphi(x, y)$ is a definable metric.*

Proof. Define $g: [0, 1]^2 \rightarrow [0, 1]$ by:

$$g(t, u) = \sup\{\varphi(x, y) : \exists z \varphi(x, z) \leq t \wedge \varphi(y, z) \leq u\}.$$

As in the Claim following the proof of [Ben05, Lemma 2.19], g satisfies the assumptions of Fact 4.9. Thus there is a function $f: D \rightarrow [0, 1]$ such that:

- (i) f is strictly increasing.
- (ii) $f \leq \text{id}_D$.
- (iii) For every $t, u \in D \cap (0, 1]$:

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

We define $h: [0, 1] \rightarrow [0, 1]$ to be a weak inverse of f : $h(t) = \sup\{u \in D: f(u) < t\}$. As f is strictly increasing it follows that $h(t) = \inf\{u \in D: f(u) > t\}$, and that h is continuous and weakly increasing. In addition, $f \leq \text{id}_D \implies h \geq \text{id}_{[0,1]}$.

Let $\psi(x, y) = h \circ \varphi(x, y)$, which is a definable predicate by continuity of h . Clearly ψ is reflexive and symmetric as well, and it is left to show that it satisfies the triangle inequality. Indeed, assume not, that is $\psi(x, y) > \psi(x, z) + \psi(y, z) + \varepsilon$ for some x, y, z . As $\psi(x, z) + \varepsilon/2 > \psi(x, z) = h \circ \varphi(x, z)$, there is $u \in D$, $u < \psi(x, z) + \varepsilon/2$, such that $f(u) > \varphi(x, z)$. Similarly there is $w < \psi(y, z) + \varepsilon/2$ in D such that $f(w) > \varphi(y, z)$. As $u + w < \psi(x, y) \leq 1$ we have:

$$\varphi(x, y) \leq g(\varphi(x, z), \varphi(y, z)) \leq g(f(u), f(w)) < f(u + w) < f(\psi(x, y)).$$

In other words we have $\varphi(x, y) < (f \circ h)(\varphi(x, y))$, contradicting the definition of h .

We've shown that $h \circ \varphi(x, y)$ is a pseudo-metric. If φ satisfies $\varphi(x, y) = 0 \implies x = y$, so does $h \circ \varphi$ (as $h \geq \text{id}$), which is therefore a metric. ■_{4.10}

Putting it all together we get:

Theorem 4.11. *Let X be a type-definable set, d_1 a partial definable (pseudo-)metric on X . Then there exist a total definable (pseudo-)metric d_2 and a continuous increasing function $h: [0, 1] \rightarrow [0, 1]$ such that $h(0) = 0$, $h \geq \text{id}$ and $d_2 \sqsupseteq h \circ d_1$.*

Proof. First, by Lemma 4.6 we may assume X is a zero set, defined by $\varphi(x) = 0$. Let $\psi_1 \sqsupseteq d_1$, and define:

$$\begin{aligned} \psi_2(x, y) &= (\psi_1(x, y) \wedge \psi_1(y, x)) \div \psi_1(x, x) \div \psi_1(y, y), \\ \psi_3(x, y) &= d(x, y) \wedge (\varphi(x) \vee \varphi(y)), \\ \varphi_4(x, y) &= \psi_2 \vee \psi_3. \end{aligned}$$

Then $\psi_2 \sqsupseteq d_1$ as well and $\psi_3 \upharpoonright_{X^2}$ is zero, whereby $\psi_4 \sqsupseteq d_1$. In addition ψ_2 , ψ_3 and ψ_4 are all symmetric and reflexive, and we may apply Proposition 4.10 to find h as such that $d_2 = h \circ \psi_4$ is a total pseudo-metric.

Assume now that d_1 is a metric. Then for $x \neq y$, if $x, y \in X$ then $\psi_2(x, y) = d_1(x, y) \neq 0$, and otherwise $\psi_3(x, y) \neq 0$, so either way $\psi_4(x, y) \neq 0$. Thus d_2 is a total metric. ■_{4.11}

4.2. Existence of invariant metrics. In general type-definable groups need not be invariant. For definable groups this is easily resolved:

Proposition 4.12. *Let $\langle G, \cdot \rangle$ be a definable group. Then G admits a total definable invariant metric (i.e., a total definable metric on the sort of G which is invariant on G).*

Proof. For simplicity assume G is definable in the home sort. Since G is a definable set, the predicate $d(x, G)$ is definable, and we may quantify over members of G . Let:

$$\begin{aligned}\psi_0(u, v, x, y) &\supseteq d(uxv, uyv) \\ \psi_1(x, y) &= \sup_{u, v \in G} \psi_0(u, v, x, y) \\ d_1(x, y) &= \psi_1(x, y)|_{x, y \in G}\end{aligned}$$

This means that for $x, y \in G$ we have: $d_1(x, y) = \sup_{u, v \in G} d(uxv, uyv)$. This is easily verified to define an invariant metric on G , which can be extended to a total definable metric by Proposition 4.5. ■_{4.12}

In the case of a type-definable group we only have partial results. First, as before we can reduce any problem to the case of zero sets:

Lemma 4.13. *Let G be a type-definable group. Then there exists a type-definable supergroup $G' \supseteq G$ whose domain is a zero set. Moreover, if d_1 is an invariant definable metric on G , and $\psi_1 \supseteq d_1$, we can choose G' so that $d_2 = \psi_1|_{G'}$ is an invariant metric on G' .*

Proof. Using Proposition 1.21 we find a zero set $X_0 \supseteq G$ to which the law of G and the inverse mappings extend to definable mappings $\cdot': X_0^2 \rightarrow X_0$ and $^{-1}': X_0 \rightarrow X_0$. The satisfaction of the identity $(x \cdot' y) \cdot' z = x \cdot' (y \cdot' z)$ (i.e., $d((x \cdot' y) \cdot' z, x \cdot' (y \cdot' z)) = 0$) in X_0 translates a condition of the form $\varphi_{ass}(x, y, z) = 0$. By Lemma 1.5 there is a zero set X_1 such that $X_0 \supseteq X_1 \supseteq G$ and such that φ_{ass} is zero on X_1^3 . We take care of $x \cdot' e = x$ and $x \cdot' x^{-1'} = e$ similarly. We end up with a zero set $X_0 \supseteq X_2 \supseteq G$ on which all the identities above hold. Now X_2 needs not be closed under product and inverse, but that is taken care of by a second application of Proposition 1.21. ■_{4.13}

Remark 4.14. This argument would work for any kind of functional structure defined by a set of identities (e.g., rings). Adding Lemma 1.6 we can further extend the argument to structures whose definition involves an implication, such as integral domains. We have already seen an example of that in Lemma 4.6.

Alternatively, given a partial invariant definable metric on G we can extend it to a total definable metric up to some uniform modification which preserves the invariance.

Proposition 4.15. *Let G be a type-definable group admitting a partial invariant definable metric. Then it also admits a total one.*

Proof. Just apply Theorem 4.11 to the partial invariant metric. ■_{4.15}

Theorem 4.16. *Assume G is a connected type-definable group in a stable theory. Then there exists a total metric which is invariant on G .*

Proof. Let $p(x)$ be the unique generic type of G . Again let $\psi_0(u, v; x, y) \supseteq d(uxv, uyv)$. As $(p \otimes p)(u, v)$ is a definable stationary type over the same parameters as G , let $\psi_1(x, y)$

be its ψ_0 -definition. Thus, if $x, y \in G$ and u, v are any independent generic elements over x, y , then $\psi_1(x, y) = d(uxv, uyv)$.

Let $d_1(x, y) = \psi_1(x, y)|_{G^2}$, and we claim it is an invariant metric on G . To verify it is a metric, let $x, y, z \in G$ be any three elements. Then choosing u, v to be independent generics over x, y, z we can verify the metric axioms for this triplet. We also get that uz, v are two independent generics over x, y , whereby $d_1(zx, zy) = d(uzxv, uzyv) = d_1(x, y)$. Right-invariance is verified similarly.

Apply Proposition 4.15 to conclude. ■_{4.16}

5. THE MAIN THEOREM (FULL VERSION)

Theorem 5.1. *A type-definable group in an \aleph_0 -stable theory is definable.*

Proof. By Proposition 2.2 it will suffice to show that the connected component of G is definable, so we may assume that G is connected. Then by Theorem 4.16 there exists a total definable metric d_1 which is invariant on G . By Lemma 4.2 we may replace the ambient metric with d_1 .

Now apply Theorem 2.11. ■_{5.1}

Notice that in the proof we needed to pass to the connected component since we do not know yet (although it seems plausible) whether a general type-definable group in a stable theory admits an invariant metric. Moreover, the passage from an partial invariant definable metric on the connected component to a total one required allowing a modification to that metric. *A posteriori* we have:

Corollary 5.2. *Let G be a type-definable group in an \aleph_0 -stable theory. Then:*

- (i) *G admits an invariant metric.*
- (ii) *Every partial metric on G extends to a total one.*

Proof. By the main theorem, G is definable, so just apply Proposition 4.12 and Proposition 4.5. ■_{5.2}

We can now remove the invariance hypothesis from Lemma 3.9:

Theorem 5.3. *Let G be a type-definable group in an \aleph_0 -stable structure. Then the class of type-definable subgroups of G satisfies the MDCC.*

Proof. Notice that if we replace the ambient metric with a uniformly equivalent one we do not affect the MDCC. Thus, by Corollary 5.2 we may assume that the metric on the group is invariant and evoke Lemma 3.9. ■_{5.3}

APPENDIX A. SOME MORE CONTINUOUS STABILITY THEORY

The author feels somewhat unsatisfied with the hasty treatment of generic types given in Section 3, especially as it relies on results proved in a completely different (even though, strictly speaking, more general) context, namely that of groups definable in thick simple

cats. A proper treatment which is naturally placed in the context of stable theories in continuous logic seems in place. For that we will need to prove some results regarding local stability which are missing from [BU].

A.1. Definability and forking of local types. We will consider throughout a formula $\varphi(\bar{x}, \bar{y})$ whose variables are split in two groups. Most of the time we will assume φ is stable (i.e., does not have the order property; alternatively, use Fact A.1 below as the definition).

Let us call a φ -*predicate scheme* (or simply a φ -*scheme*) a formal expression $\psi(\bar{x}, \bar{y}_{<\omega})$ which is a forced limit of continuous combinations of $\varphi(\bar{x}, \bar{y}_i)$ (possibly restricting to $\{\neg, \div, \frac{1}{2}\}$), or equivalently, a single infinitary continuous combination $\theta(\varphi(\bar{x}, \bar{y}_i) : i < \omega)$. If z is a variable in the (imaginary) sort of canonical parameters for instances of $\psi(\bar{x}, \bar{y}_{<\omega})$ we have a bijection between instances of $\psi(\bar{x}, \bar{y}_{<\omega})$ and instances of $P_\psi(\bar{x}, z)$, and for all (or most) intents and purposes they are quite interchangeable. By a convenient abuse of notation we may write $\psi(\bar{x}, z)$ instead of $P_\psi(\bar{x}, z)$, allowing us to pretend that ψ is a finite formula. Being a φ -scheme is transitive: if $\psi(\bar{x}, z)$ is a φ -scheme then every ψ -scheme is a φ -scheme.

An instance $\psi(\bar{x}, \bar{b}_{<\omega})$ of a φ -scheme is what we called in [BU] a φ -*predicate*. It may also be written as $\psi(\bar{x}, c)$ where c is its canonical parameter. We say that it is *definable over* a set A if it is invariant over A , i.e., if $c \in \text{dcl}^{\text{eq}}(A)$. The φ -*type* of a tuple \bar{a} over a set A is given by the values of all φ -predicates which are definable over A ; in case A is a model, it suffices to consider only instances of φ with parameters in M .

Let us introduce some convenient notation. If $\varphi(\bar{x}, \bar{y})$ is any formula with two groups of variables, $\tilde{\varphi}(\bar{y}, \bar{x})$ denotes the same formula with the groups of variables interchanged. More generally, let us define

$$\tilde{\varphi}^n(\bar{y}, \bar{x}_{<2n-1}) = \text{med}_n(\varphi(\bar{x}_i, \bar{y}) : i < 2n-1),$$

where $\text{med}_n : [0, 1]^{2n-1} \rightarrow [0, 1]$ is the median value combination:

$$\text{med}_n(x_{<2n-1}) = \bigwedge_{w \in [2n-1]^n} \bigvee_{i \in w} x_i = \bigvee_{w \in [2n-1]^n} \bigwedge_{i \in w} x_i.$$

Thus in particular $\tilde{\varphi}^1 = \tilde{\varphi}$, and each $\tilde{\varphi}^n$ is a $\tilde{\varphi}$ -scheme.

Fact A.1. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula. Then for every $\varepsilon > 0$ there is $N(\varphi, \varepsilon) < \omega$ such that if M is any model and $p(\bar{x}) \in S_\varphi(M)$ any complete φ -type, then p is defined up to ε by a formula of the form $d_p^\varepsilon \varphi(\bar{y}) = \tilde{\varphi}^{N(\varphi, \varepsilon)}(\bar{y}, \bar{c}_{<2N(\varphi, \varepsilon)-1})$ where $\bar{c}_{<2N(\varphi, \varepsilon)-1} \in M$, by which we mean that for all $\bar{b} \in M$: $|\varphi(\bar{x}, \bar{b})^p - d_p^\varepsilon \varphi(\bar{b})| < \varepsilon$. Thus p is defined by $d_p \varphi(\bar{y}) = \mathcal{F}\text{lim}_{n \rightarrow \infty} d_p^{2^{-n}} \varphi(\bar{y})$ which is a $\tilde{\varphi}$ -predicate (so for all $\bar{b} \in M$: $\varphi(\bar{x}, \bar{b})^p = d_p \varphi(\bar{b})$).*

Moreover, let $A \subseteq M$ and assume that M is $|A|^+$ -saturated, and let $q \in S_n(M)$ extend p . Then we can take each \bar{c}_i to realise $q \upharpoonright_{A\bar{c}_{<i}}$.

Proof. The first assertion is [BU, Lemma 7.4]. The moreover part, while not explicitly stated there, is immediate from the proof. ■_{A.1}

Let us also recall:

Fact A.2. [BU, Lemma 6.8] *Let $\varphi(\bar{x}, \bar{y})$ be any formula, A a set, $M \supseteq A$ an $(|A| + \aleph_0)^+$ -saturated and strongly homogeneous model. Then $\text{Aut}(M/A)$ acts transitively on $S_\varphi(\text{acl}^{eq}(A))$.*

We recall that $p \in S_\varphi(M)$ does not fork over a subset $A \subseteq M$ if its definition is over $\text{acl}^{eq}(A)$. We proved in [BU, Proposition 7.15] that every φ -type over a set A admits a non forking extension to every model (and therefore every set) containing A . The exact same proof can be used to show:

Fact A.3. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, A a set, $M \supseteq A$ an $(|A| + \aleph_0)^+$ -saturated and strongly homogeneous model, and assume $q(\bar{x})$ a consistent partial type over M which is invariant under $\text{Aut}(M/A)$. Then there exists $p \in S_\varphi(M)$ compatible with q which does not fork over A .*

Proof. Let $X = \{p \in S_\varphi(M) : p \cup q \text{ is consistent}\}$. Then X is non empty and A -invariant. By [BU, Lemma 7.14], there is $Y \subseteq X$ which is A -good, i.e., which is A -invariant and metrically compact. By [BU, Lemma 7.13], any $p \in Y$ would do. ■_{A.3}

Similarly, we will slightly improve [BU, Proposition 7.17]:

Fact A.4. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, A a set, $M \supseteq A$ a model. If $p \in S_\varphi(M)$ is definable over A then $p|_A$ is stationary (i.e., p is its unique non forking extension to M).*

Proof. In case $A = \text{acl}^{eq}(A)$ this is just [BU, Proposition 7.17]. In the general case we may assume that M is strongly $(|A| + \aleph_0)^+$ -homogeneous. Let $p' \in S_\varphi(M)$ be another non forking extension of $p|_A$, and let $q = p|_{\text{acl}^{eq}(A)}$, $q' = p'|_{\text{acl}^{eq}(A)}$. By Fact A.2 there is an automorphism $f \in \text{Aut}(M/A)$ sending q to q' . As p is definable over A it is fixed by f , so $q' = q$. Thus p and p' are both non forking extensions of q , and we conclude that $p = p'$ using the first case. ■_{A.4}

Corollary A.5. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, A a set, $q(\bar{x}) \in S_n(A)$ a complete type over A , and let $p_0 = q|_\varphi \in S_\varphi(A)$. Then q is compatible with every non forking extension of p_0 .*

Proof. Let $p \in S_\varphi(\bar{M})$ be a non forking extension of p_0 . By Fact A.3 there exists $p' \in S_\varphi(\bar{M})$ non forking over A such that $p' \cup q$ is consistent. By Fact A.2 there is an automorphism $f \in \text{Aut}(\bar{M}/A)$ sending $p'|_{\text{acl}^{eq}(A)}$ to $p|_{\text{acl}^{eq}(A)}$, and by the uniqueness of the non forking extension f sends p' to p . It follows that $p \cup q$ is consistent as well. ■_{A.5}

Proposition A.6. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, M a model, $A \subseteq M$. For each $\bar{b} \in M$ let $\psi_{\bar{b}}(\bar{x})$ be the definition of $\text{tp}_{\bar{\varphi}}(\bar{b}/\text{acl}^{eq}(A))$ (i.e., of its unique non forking extension to M). Then:*

- (i) *Each $\psi_{\bar{b}}(\bar{x})$ is a φ -predicate over $\text{acl}^{eq}(A)$.*
- (ii) *For every $p \in S_\varphi(M)$ which does not fork over A : $\varphi(\bar{x}, \bar{b})^p = \psi_{\bar{b}}(\bar{x})^p$.*

- (iii) Let $r(\bar{x}) = \{|\varphi(\bar{x}, \bar{b}) - \psi_{\bar{b}}(\bar{x})| = 0 : \bar{b} \in M\}$. Then the partial type $r(\bar{x})$ defines the set of φ -types which do not fork over A :

$$\bar{a} \models r \iff \text{tp}_\varphi(\bar{a}/M) \text{ does not fork over } A.$$

- (iv) For every $B \supseteq A$, the set $\{p \in S_\varphi(B) : p \text{ does not fork over } A\}$ is closed.

Proof. The first item is by Fact A.1 and the definition of non forking.

For the second, fix $\bar{b} \in M$, let $q_0 = \text{tp}_{\tilde{\varphi}}(\bar{b}/\text{acl}^{eq}(A))$ and let $q = q_0 \upharpoonright^M$, so $\psi_{\bar{b}}$ defines q . Assume $p \in S_\varphi(M)$ does not divide over M , so $d_p\varphi(\bar{y})$ is a $\tilde{\varphi}$ -predicate over $\text{acl}^{eq}(A)$ as well. By [BU, Proposition 7.14]: $d_p\varphi(\bar{y})^q = d_q\tilde{\varphi}(\bar{x})^p = \psi_{\bar{b}}(\bar{x})^p$. We obtain:

$$\varphi(\bar{x}, \bar{b})^p = d_p\varphi(\bar{b}) = d_p\varphi(\bar{y})^{q_0} = d_p\varphi(\bar{y})^q = \psi_{\bar{b}}(\bar{x})^p.$$

We know that $\psi_{\bar{b}}(\bar{x})$ and $\varphi(\bar{x}, \bar{b})$ are φ -predicates over M (indeed, the former is over $\text{acl}^{eq}(A)$), so so is $|\varphi(\bar{x}, \bar{b}) - \psi_{\bar{b}}(\bar{x})|$. Thus the satisfaction of $r(\bar{x})$ by \bar{a} depends solely on $p = \text{tp}_\varphi(\bar{a}/M)$. If p does not fork over A (i.e., over $\text{acl}^{eq}(A)$) then $p \vdash r$. Conversely, if $p \vdash r$ then it coincides with the unique non forking extension of $p \upharpoonright_{\text{acl}^{eq}(A)}$.

Finally, we may assume that $B \subseteq M$. The set $[r] \subseteq S_\varphi(M)$ is closed, and so is its projection to $S_\varphi(B)$. This projection is precisely the set of types which do not fork over A . ■_{A.6}

Notice that the proof of Proposition A.6 contains an implicit proof of the uniqueness of the non forking extension of a type over an algebraically closed set.

Definition A.7. Let $\varphi(\bar{x}, \bar{y})$ be stable. We say that a condition of the form $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over a set A if there is a complete φ -type $p(\bar{x}) \in S_\varphi(A\bar{b})$ such that $\varphi(\bar{x}, \bar{b})^p \leq r$ and p does not fork over A .

Thus $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A if and only if for all (some) model $M \supseteq A\bar{b}$ there is $p \in S_\varphi(M)$ which is definable over $\text{acl}^{eq}(A)$. Clearly, $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A if and only if it does not fork over $\text{acl}^{eq}(A)$.

Lemma A.8. Let $\varphi(\bar{x}, \bar{y})$ be stable, $\psi(\bar{x}, c)$ a φ -predicate (whose canonical parameter is c). Then a condition $\psi(\bar{x}, c) \leq r$ does not fork over a set A if and only if there exists a model $M \supseteq A, c$ and a φ -type $p \in S_\varphi(M)$ non forking over A such that $\psi(\bar{x}, c)^p \leq r$.

Proof. First of all notice that if φ is stable then every φ -scheme is, i.e., $\psi(\bar{x}, z)$ is, so the condition makes sense. Also, we may assume that $A = \text{acl}^{eq}(A)$. Fix a strongly $|A|^+$ -homogeneous model $M \supseteq A, c$.

Assume first that $p \in S_\psi(M)$ witnesses that $\psi(\bar{x}, c) \leq r$ does not fork over A . Then p is A -invariant, and by Fact A.3 there is $\hat{p} \in S_\varphi(M)$ compatible with p and non forking over A . Then $\psi(\bar{x}, c)^{\hat{p}} = \psi(\bar{x}, c)^p \leq r$ as required.

Conversely, assume $p \in S_\varphi(M)$ does not fork over A and $\psi(\bar{x}, c)^p \leq r$. As $\psi(\bar{x}, z)$ is a φ -scheme, the φ -type p determines a ψ -type $p \upharpoonright_\psi \in S_\psi(M)$. Since p is A -invariant so is $p \upharpoonright_\psi$. This means that the definition of $p \upharpoonright_\psi$ is over A , so it does not fork over A either. ■_{A.8}

Definition A.9. Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, $A, \bar{b} \subseteq M$. We define *non forking degree* of $\varphi(\bar{x}, \bar{b})$ over A to be

$$\text{nf}(\varphi(\bar{x}, \bar{b})/A) = \inf\{r : \varphi(\bar{x}, \bar{b}) \leq r \text{ does not fork over } A\}.$$

As the set of non forking types is closed we see that $\varphi(\bar{x}, \bar{b}) \leq \text{nf}(\varphi(\bar{x}, \bar{b})/A)$ does not fork over A . In addition, by the existence of non-forking types we have $\text{nf}(\varphi(\bar{x}, \bar{b})/A) + \text{nf}(\neg\varphi(\bar{x}, \bar{b})/A) \leq 1$.

Proposition A.10. Let $A \subseteq B$ be two sets, $\varphi(\bar{x}, \bar{y})$ a stable formula, and $p \in S_\varphi(B)$. Then the following are equivalent:

- (i) The φ -type p does not fork over A .
- (ii) For every φ -predicate $\psi(\bar{x}, c)$ over B : $\psi(\bar{x}, c)^p \geq \text{nf}(\psi(\bar{x}, c)/A)$.
- (iii) For every φ -predicate $\psi(\bar{x}, c)$ over B , if $\psi(\bar{x}, c)^p = 0$ then for all $\varepsilon > 0$ the condition $\psi(\bar{x}, c) \leq \varepsilon$ does not fork over A .

Proof. (i) \implies (ii). By Lemma A.8: $\psi(\bar{x}, c) \leq \psi(\bar{x}, c)^p$ does not fork over A .

(ii) \implies (iii). By definition of the non forking degree.

(iii) \implies (i). Let $X \subseteq S_\varphi(B)$ be the set of types which do not fork over A , which is closed by Proposition A.6. By assumption and Lemma A.8 every neighbourhood of p intersects X , whereby $p \in X$. ■_{A.10}

A.2. Approximately realised types, indiscernible sequences and dividing.

Definition A.11. Let M be a model, $A \supseteq M$ a set containing M .

- (i) Let p be a partial type over A . We say that p is *approximately realised* in M if every logical neighbourhood of p is realised in M .
- (ii) Let p be a complete type over A . We say that p is *non splitting* over M if for every two tuples $\bar{b}, \bar{c} \in A$, if $\bar{b} \equiv_M \bar{c}$ then $\varphi(\bar{x}, \bar{b})^p = \varphi(\bar{x}, \bar{c})^p$ for every formula φ with the right number of variables.

Remark A.12. (i) The classical logic analogue of an approximately realised type is a *finitely realised* one, but this terminology would be misleading in the continuous setting.

- (ii) A complete approximately M -realised type is sometimes called a *co-heir* of its restriction to M .
- (iii) A complete type over a model M is always approximately realised there.

Fact A.13. Let $M \subseteq N$ be two models. Then every approximately M -realised partial type p over N extends to a complete one.

Proof. Let $A = \{\text{tp}(\bar{b}/N) : \bar{b} \in M^n\} \subseteq S_n(N)$, and let $[p] \subseteq S_n(N)$ be the closed set defined by p . Then p is approximately realised in M if and only if $[p] \cap \bar{A} \neq \emptyset$, in which case any q in the intersection would do. ■_{A.13}

Fact A.14. Assume $A \supseteq M$ and $p \in S_n(A)$ is approximately realised in M . Then it is non splitting over M .

Proof. Say $\bar{b}, \bar{c} \in A$, $\bar{b} \equiv_M \bar{c}$, and let $\varphi(\bar{x}, \bar{y})$ and $\varepsilon > 0$ be given. By assumption there is $\bar{a} \in M$ such that

$$|\varphi(\bar{a}, \bar{b}) - \varphi(\bar{x}, \bar{b})^p| < \varepsilon/2, \quad |\varphi(\bar{a}, \bar{c}) - \varphi(\bar{x}, \bar{c})^p| < \varepsilon/2.$$

As we assumed that $\bar{b} \equiv_M \bar{c}$ we get $\varphi(\bar{a}, \bar{b}) = \varphi(\bar{a}, \bar{c})$ and thus $|\varphi(\bar{x}, \bar{b})^p - \varphi(\bar{x}, \bar{c})^p| < \varepsilon$, for every $\varepsilon > 0$. We conclude that $\varphi(\bar{x}, \bar{b})^p = \varphi(\bar{x}, \bar{c})^p$, as required. $\blacksquare_{A.14}$

Fact A.15. *Let $M \subseteq N$ be two models. Assume that N is $|M|^+$ -saturated, and that $p \in S_n(N)$ is non splitting over M . Let $(\bar{a}_i : i < \omega) \subseteq N$ be a sequence constructed inductively, choosing each \bar{a}_i to realise $p \upharpoonright_{M\bar{a}_{<i}}$.*

Then the sequence $(\bar{a}_i : i < \omega)$ is M -indiscernible, and its type over M depends only on p .

Proof. Standard. $\blacksquare_{A.15}$

Lemma A.16. *Let $M \subseteq N$ be two models, $p(\bar{x}) \in S_n(N)$ approximately realised in M , and assume $\varphi(\bar{x}, \bar{y})$ is stable. Then $p_\varphi = p \upharpoonright_\varphi$ is non forking over M .*

Proof. If $N' \supseteq N$ is a larger model and $p' \in S_n(N')$ extends p , it will suffice to show that $p' \upharpoonright_\varphi$ does not fork over M . Moreover, we may always choose such p' which is approximately realised in M . We may therefore assume that N is $|M|^+$ -saturated and strongly homogeneous.

Now let $\psi(\bar{y}, c)$ be the φ -definition of p , and let $f \in \text{Aut}(N/M)$. As p does not split over M we have $f(p) = p$, whereby $\psi(\bar{y}, c) = \psi(\bar{y}, f(c))$. This shows that $\psi(\bar{x}, \bar{c})$ is over M , as required. $\blacksquare_{A.16}$

Proposition A.17. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, $M \subseteq N$ models, $p(\bar{x}) \in S_\varphi(N)$ a complete φ -type, and $q(\bar{x}) \in S_n(M)$ a complete type over M such that $p \upharpoonright_M = q \upharpoonright_\varphi \in S_\varphi(M)$. Then the following are equivalent:*

- (i) $p \cup q$ is approximately realised in M .
- (ii) p is approximately realised in M .
- (iii) p does not fork over M .

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

(ii) \implies (iii). Find $p'(\bar{x}) \in S_n(N)$ extending p which is approximately realised in M and use Lemma A.16.

(iii) \implies (i). Find $q'(\bar{x}) \in S_n(N)$ extending q which is approximately realised in M . Then $q' \upharpoonright_\varphi$ is non forking over M by Lemma A.16, so it must be the unique non forking extension of $p \upharpoonright_M = q \upharpoonright_\varphi$. It follows that $p = q' \upharpoonright_\varphi$, so $q' \models q \cup p$. As q' is approximately realised in M , *a fortiori* so is $q \cup p$. $\blacksquare_{A.17}$

Lemma A.18. *Let M be a model, $\varphi(\bar{x}, \bar{y})$ a stable formula, $p \in S_\varphi(M)$ a complete φ -type.*

Let $q \in S_n(M)$ extend p to a complete type, and let $q \subseteq q_1 \in S_n(N)$ be approximately realised in M , i.e., a co-heir of q , where $N \supseteq M$ is an $|M|^+$ -saturated extension.

Construct $(\bar{c}_i: i < \omega) \subseteq N$ letting each \bar{c}_i realise $q_1 \upharpoonright_{M, \bar{c}_{<i}}$. Then the sequence $(\tilde{\varphi}^n(\bar{y}, \bar{c}_{<2n-1}): n < \omega)$ converges uniformly to the definition of p at a rate which only depends on φ .

Proof. Let $\psi(\bar{y})$ denote the definition of p . Then it also defines $p_1 = q_1 \upharpoonright_{\varphi} \in S_{\varphi}(N)$, as the latter is approximately realised in M and is thus a non forking extension of p .

Fix $\varepsilon > 0$, and let $N_0 = N(\varphi, \varepsilon)$. By Fact A.1 there is a sequence $(\bar{c}'_i: i < 2N_0 - 1) \subseteq N$ such that $|\psi(\bar{y}) - \tilde{\varphi}^{N_0}(\bar{y}, \bar{c}'_{<2N_0-1})| \leq \varepsilon$, and such that furthermore $\bar{c}'_i \models q_1 \upharpoonright_{M, \bar{c}'_{<i}}$. By Fact A.15 we have $\bar{c}_{<2N_0-1} \equiv_M \bar{c}'_{<2N_0-1}$, and since in addition ψ is over M : $|\psi(\bar{y}) - \tilde{\varphi}^{N_0}(\bar{y}, \bar{c}_{<2N_0-1})| \leq \varepsilon$.

Moreover, for any $n \geq N_0$ and any \bar{b} (say in N) the median value of $(\varphi(\bar{c}_i, \bar{b}): i < 2n-1)$ equals the median value of $(\varphi(\bar{c}_{i_j}, \bar{b}): j < 2N_0 - 1)$ for some $j_0 < \dots < j_{2N_0-2} < 2n-1$, whereby $\tilde{\varphi}^n(\bar{b}, \bar{c}_{<2n-1}) = \tilde{\varphi}^{N_0}(\bar{b}, \bar{c}_{i_{<2N_0-1}})$. As before $\bar{c}_{i_{<2N_0-1}} \equiv_M \bar{c}'_{<2N_0-1}$ and $|\psi(\bar{y}) - \tilde{\varphi}^{N_0}(\bar{y}, \bar{c}_{i_{<2N_0-1}})| \leq \varepsilon$, whereby $|\psi(\bar{b}) - \tilde{\varphi}^n(\bar{b}, \bar{c}_{<2n-1})| \leq \varepsilon$. Thus $|\psi(\bar{y}) - \tilde{\varphi}^n(\bar{y}, \bar{c}_{<2n-1})| \leq \varepsilon$ for all $n \geq N_0$, where N_0 depends only on ε and φ . ■_{A.18}

We now seek an analogue of [Pil96, Lemma 1.2.16]. The notion of a positive Boolean combination will be replaced here with that of a *faithful combination*:

Definition A.19. A *faithful continuous connective* in α variables is a continuous function $\theta: [0, 1]^\alpha \rightarrow [0, 1]$ satisfying $\inf \bar{a} \leq \theta(\bar{a}) \leq \sup \bar{a}$.

If $\theta: [0, 1]^\alpha \rightarrow [0, 1]$ is a faithful continuous connective and $(\varphi_i: i < \alpha)$ a sequence of definable predicates, then the definable predicate $\theta(\varphi_i: i < \alpha)$ is called a *faithful combination* of $(\varphi_i: i < \alpha)$.

Since a continuous function to $[0, 1]$ can only take into account countably many arguments, we may always assume that $\alpha \leq \omega$. Notice that any connective constructed with \vee and \wedge is faithful (so in particular the median value function $\text{med}_n: [0, 1]^{2n-1} \rightarrow [0, 1]$ is), as is the forced limit function $\mathcal{F}\text{lim}: [0, 1]^\omega \rightarrow [0, 1]$. In fact these connectives are also positive (i.e., increasing in every argument) but this will not be needed here.

Lemma A.20. Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, $r \in [0, 1]$ and $\varphi(\bar{x}, \bar{b})$ an instance of φ . Let A be some set of parameters (so $A, \bar{b} \subseteq \bar{M}$). Then the following are equivalent:

- (i) The condition $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A .
- (ii) There is an A -definable predicate $\psi(\bar{x})$ which is a faithful combination of A -conjugates of $\varphi(\bar{x}, \bar{b})$ such that $\psi(\bar{x}) \leq r$ is consistent.
- (iii) Every set of $\text{acl}^{\text{eq}}(A)$ -conjugates of the condition $\varphi(\bar{x}, \bar{b}) \leq r$ is consistent.
- (iv) If $(\bar{b}_i: i < \omega)$ is an A -indiscernible sequence, $\bar{b}_0 = \bar{b}$, then $\{\varphi(\bar{x}, \bar{b}_i) \leq r: i < \omega\}$ is consistent (i.e., the condition $\varphi(\bar{x}, \bar{b}) \leq r$ does not divide over A).

Proof. Fix $M \supseteq A, \bar{b}$ which is $|A|^+$ -saturated. Let $q_0 = \text{tp}(\bar{b}/\text{acl}^{\text{eq}}(A))$. By Corollary A.5 there is $q \in S_m(M)$ extending q_0 such that $q \upharpoonright_{\varphi}$ does not fork over A , and notice that this determines $q \upharpoonright_{\varphi}$ uniquely. We let $\chi(\bar{x}, c) = d_q \tilde{\varphi}(\bar{x})$ be its $\tilde{\varphi}$ -definition: it is a φ -predicate and $c \in \text{acl}^{\text{eq}}(A)$.

(i) \implies (ii). Let $p \in S_\varphi(M)$ witness that $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A . By Proposition A.6: $\chi(\bar{x}, c)^p = \varphi(\bar{x}, \bar{b})^p \leq r$.

Let $C = \{c' : c' \equiv_A c\}$ be the set of A -conjugates of c . Since c is algebraic over A , C is (metrically) compact. By Lemma 1.11 $\psi(\bar{x}) = \inf_{z \in C} \chi(\bar{x}, z)$ is a continuous combination of instances $\chi(\bar{x}, c')$ with $c' \in C$, i.e., of A -conjugates of $\chi(\bar{x}, c)$, and it is clearly a faithful combination. It follows that $\psi(\bar{x})$ is a faithful combination of A -conjugates of $\varphi(\bar{x}, \bar{b})$. It is clearly definable over A . Finally, $\psi(\bar{x})^p \leq \chi(\bar{x}, c)^p \leq r$.

(ii) \implies (iii). Let $\psi(\bar{x}) = \theta(\varphi(\bar{x}, \bar{b}_i) : i < \omega)$ be as in the assumption (so $\bar{b}_i \equiv_A \bar{b}$ and θ is a faithful continuous connective). We need to show that if $\bar{b}'_j \equiv_{\text{acl}^{eq}(A)} \bar{b}$ for $j < \lambda$ then $\{\varphi(\bar{x}, \bar{b}'_j) \leq r : j < \lambda\}$ is consistent.

By assumption there is $p_0 \in S_\varphi(A)$ such that $\psi(\bar{x})^{p_0} \leq r$. Let $p \in S_\varphi(\bar{M})$ be a non forking extension of p_0 . As θ is faithful we have $\inf\{\varphi(\bar{x}, \bar{b}_i)^p : i < \omega\} \leq r$. Thus for any $\varepsilon > 0$ there is $i < \omega$ such that $\varphi(\bar{x}, \bar{b}_i)^p < r + \varepsilon$. Let $f \in \text{Aut}(M/A)$ be such that $f(\bar{b}_i) = \bar{b}$, and let $p' = f(p)$. Then p' is also a non forking extension of p_0 , and $\varphi(\bar{x}, \bar{b})^{p'} < r + \varepsilon$. Since p' is definable over $\text{acl}^{eq}(A)$ we have $\varphi(\bar{x}, \bar{b}'_j)^{p'} = \varphi(\bar{x}, \bar{b})^{p'} < r + \varepsilon$ for all $j < \lambda$, showing that the set of conditions $\{\varphi(\bar{x}, \bar{b}'_j) \leq r + \varepsilon : j < \lambda\}$ is consistent for all $\varepsilon > 0$. We conclude using compactness.

(iii) \implies (iv). If $(\bar{b}_i : i < \omega)$ is an A -indiscernible sequence and $\bar{b}_0 = \bar{b}$ then each \bar{b}_i is an $\text{acl}^{eq}(A)$ -conjugate of \bar{b} .

(iv) \implies (i). Use Lemma A.18 to construct an M -indiscernible sequence $(\bar{b}_i : i < \omega)$ such that each \bar{b}_i realises q and $d_q \tilde{\varphi}(\bar{x})$ is a faithful combination of $\{\varphi(\bar{x}, \bar{b}_i) : i < \omega\}$. By assumption there exists \bar{a} such that $\varphi(\bar{a}, \bar{b}_i) \leq r$ for all $i < \omega$, whereby $d_q \tilde{\varphi}(\bar{x}) \leq r$.

Let $p_0 = \text{tp}_\varphi(\bar{a} / \text{acl}^{eq}(A))$ and let $p \in S_\varphi(M)$ be a non forking extension. Then $p \varphi(\bar{x}, \bar{b})^p = d_q \tilde{\varphi}(\bar{x})^p \leq r$. ■_{A.20}

Let us conclude this discussion on a somewhat different tone, showing that in continuous \aleph_0 -stable theories independence is (as expected) characterised via Morley ranks.

Lemma A.21. *Assume T is \aleph_0 -stable, $A \subseteq B$, $q \in S_n(B)$, $p = q \upharpoonright_A \in S_n(A)$. Then the following are equivalent:*

- (i) *The type q is a non forking extension of p .*
- (ii) *For all $\varepsilon > 0$: $\text{RM}_\varepsilon(p) = \text{RM}_\varepsilon(q)$.*
- (iii) *For all $\varepsilon > \varepsilon' > 0$: $\text{RM}_\varepsilon(p) \leq \text{RM}_{\varepsilon'}(q)$.*
- (iv) *For arbitrarily small $\varepsilon > 0$: $\text{RM}_\varepsilon(p) = \text{RM}_\varepsilon(q)$.*

Proof. There is no harm in assuming that $B = \bar{M}$.

Let $\varepsilon > 0$. For an ordinal α recall that $S_n(\bar{M})_\varepsilon^{(\alpha)}$ denotes the α th ε -Cantor-Bendixson derivative of $S_n(\bar{M})$ (i.e., (ε, f) -Cantor-Bendixson derivative in the terminology of [Ben]). Let $\alpha_\varepsilon = \text{RM}_\varepsilon(p)$, i.e., the maximal α such that $S_n(\bar{M})_\varepsilon^{(\alpha)} \cap [p] \neq \emptyset$. So let $X_\varepsilon = S_n(\bar{M})_\varepsilon^{(\alpha_\varepsilon)} \cap [p]$, the set of extensions of p of maximal RM_ε rank. It is compact, and since α_ε is maximal it admits a cover by relatively open ε -finite subsets of $S_n(\bar{M})_\varepsilon^{(\alpha_\varepsilon)}$. This cover admits a finite sub-cover so X_ε is ε -finite itself.

(i) \implies (ii). Observe that each X_ε is invariant under the action of automorphisms of \bar{M} which fix A , and thus contains a non forking extension of p . As all non forking extensions of p are conjugate over A they all belong to X_ε , i.e., they all satisfy $\text{RM}_\varepsilon(q) = \alpha_\varepsilon = \text{RM}_\varepsilon(p)$.

(ii) \implies (iii). As $\varepsilon > \varepsilon' \implies \text{RM}_\varepsilon(p) \leq \text{RM}_{\varepsilon'}(p) = \text{RM}_{\varepsilon'}(q)$.

(iii) \implies (iv). By Lemma 2.7 there are arbitrarily small pairs $\varepsilon > \varepsilon' > 0$ such that $\text{RM}_\varepsilon(q) = \text{RM}_{\varepsilon'}(q)$, whereby $\text{RM}_\varepsilon(p) \leq \text{RM}_\varepsilon(q)$. As the inverse inequality is immediate we have equality for arbitrarily small $\varepsilon > 0$.

(iv) \implies (i). Let $\mathcal{E} = \{\varepsilon > 0 : \text{RM}_\varepsilon(p) = \text{RM}_\varepsilon(q)\}$, so $\inf \mathcal{E} = 0$. Let $X = \bigcap_{\varepsilon \in \mathcal{E}} X_\varepsilon \subseteq [p] \subseteq S_n(\bar{M})$. On the one hand X is compact and therefore metrically complete (see [Ben]). On the other hand it is ε -finite for every ε , i.e., totally bounded. Thus X is metrically compact. It follows that for every $\varphi(\bar{x}, \bar{y})$, the image of X in $S_\varphi(\bar{M})$ is metrically compact (since φ is uniformly continuous), so each $q \in X$ is definable over $\text{acl}^{\text{eq}}(A)$. Thus every $q \in X$ is a non forking extension of p . $\blacksquare_{\text{A.21}}$

Notice in passing that we showed that in a continuous \aleph_0 -stable theory every type has “compact multiplicity”, in analogy with the finite multiplicity of types in a classical \aleph_0 -stable theory.

A.3. Canonical bases. Recall that the *canonical base* of a stationary type $p \in S_n(A)$ in a stable theory is $\text{Cb}(p) = \{\text{Cb}(p|_\varphi) : \varphi(\bar{x}, \dots) \in \mathcal{L}\}$, namely the set of all canonical parameters of φ -definitions of p .

Proposition A.22. *Assume T is stable, and let $p(\bar{x}) \in S_n(A)$ be stationary. Then:*

- (i) $\text{Cb}(p) \subseteq \text{dcl}(A)$.
- (ii) p does not fork over $\text{Cb}(p)$.
- (iii) $p|_{\text{Cb}(p)}$ is stationary.
- (iv) $\text{Cb}(p)$ is minimal for the three previous properties, meaning that if $B \subseteq \text{dcl}(A)$ and $p|_B$ is a stationary non forking restriction then $\text{Cb}(p) \subseteq \text{dcl}(B)$.

Proof. The first two items are immediate, while the third is by Fact A.4. Under the assumptions of the fourth we have $\text{Cb}(p) = \text{Cb}(p|_B) \subseteq \text{dcl}(B)$. $\blacksquare_{\text{A.22}}$

The four properties listed in Proposition A.22 determine the canonical base up to inter-definability. Indeed, if B has all four then $\text{Cb}(p) \subseteq \text{dcl}(B)$ but also $B \subseteq \text{dcl}(\text{Cb}(p))$ whence $\text{dcl}(B) = \text{dcl}(\text{Cb}(p))$. In this case we say that B is a canonical base for p .

Proposition A.23. *Assume T is stable, and let $p(\bar{x}) \in S_n(A)$ be stationary. Let $q \in S_n(\bar{M})$ be the unique non forking extension of p to the monster model. Then a (small) set $B \subseteq \bar{M}$ is a canonical base for p if and only if, for every $f \in \text{Aut}(\bar{M})$: $f|_B = \text{id}_B \iff f(q) = q$.*

Proof. Let $C = \text{Cb}(p) = \text{Cb}(q)$. On the one hand q determines each member of C , while on the other q is the unique non forking extension of $q|_C$. Thus an automorphism fixes q if and only if it fixes each member of C . Finally, a small set B is another canonical base

for p if and only if $\text{dcl}(B) = \text{dcl}(C)$ which is further equivalent to B and C being fixed by the same automorphisms. ■_{A.23}

We propose an alternative characterisation of canonical bases using Morley sequences. Recall that a *Morley sequence* in p is a sequence $(\bar{a}: i < \omega)$ of realisations of p which is independent over A , i.e., such that $\bar{a}_i \perp_A \bar{a}_{<i}$ for all $i < \omega$. It follows by standard independence calculus that $\bar{a}_{\in I} \perp_A \bar{a}_{\in J}$ for every two disjoint index sets $I, J \subseteq \omega$. Since p is stationary it follows that this is an A -indiscernible sequence, and moreover, an A -indiscernible set.

Theorem A.24. *Assume T is stable, and let $p(\bar{x}) \in S_n(A)$ be stationary. Let $(\bar{a}_i: i < \omega)$ be a Morley sequence in p . Then:*

- (i) $\text{Cb}(p) \subseteq \text{dcl}(A) \cap \text{dcl}(\bar{a}_{<\omega})$.
- (ii) p does not fork over $\text{Cb}(p)$.
- (iii) $p|_{\text{Cb}(p)}$ is stationary.

Conversely, if $B \subseteq \text{dcl}(A) \cap \text{dcl}(\bar{a}_{<\omega})$ and $p|_B$ is a stationary non forking restriction of p then B is also a canonical base for p .

Proof. For the main assertion we only need to show that $\text{Cb}(p) \subseteq \text{dcl}(\bar{a}_{<\omega})$. Since p is stationary, all Morley sequences in p have the same type over A , so it will be enough show that there *exists* a Morley sequence $(\bar{a}'_i: i < \omega)$ in p such that $\text{Cb}(p) \subseteq \text{dcl}(\bar{a}'_{<\omega})$.

Choose models $N \supseteq M \supseteq A$ where N is $|M|^+$ -saturated. Let $p_2 = p|_N$ and $p_1 = p|_M$ be the unique non forking extensions to N and M , respectively. Then p_1 is approximately realised in M . Construct a sequence $(\bar{a}'_i: i < \omega)$ such that $\bar{a}'_i \models p_1|_{N, \bar{a}'_{<i}}$. Then $(\bar{a}'_i: i < \omega)$ is a Morley sequence in p_2 by construction, and by transitivity of independence it is a Morley sequence in p . On the other hand, by Lemma A.18 we have $\text{Cb}(p) = \text{Cb}(p_2) \subseteq \text{dcl}(\bar{a}'_{<\omega})$.

For the converse, let B be as assumed there. Then we already know that $\text{Cb}(p) \subseteq \text{dcl}(B)$, and we are left with showing the converse inclusion.

For any two Morley sequences in p there is an automorphism sending one to the other while fixing A , and thus B , pointwise. Since B is in the definable closure of each sequence, it follows that every automorphism sending one to the other (fixing A or not) fixes B . In particular, if we have two independent such sequences (i.e., independent as sets over A) then every automorphism sending one to the other fixes B . However, such a pair of independent Morley sequences in p is also a pair of Morley sequences in $q = p|_{\text{Cb}(p)}$, which are independent as such, i.e., over $\text{Cb}(p)$. Since all pairs of independent Morley sequence in q have the same type, any automorphism sending one Morley sequence in q to another, independent over $\text{Cb}(p)$ from the first, must fix B , as long as at least one of the two is in fact a Morley sequence in p . Now let $\bar{b}_{<\omega}$ and $\bar{c}_{<\omega}$ be any two Morley sequences in q and let $\bar{a}_{<\omega}$ be a Morley sequence in $p|_{A, \bar{b}_{<\omega}, \bar{c}_{<\omega}}$. Then $\bar{a}_{<\omega}$ is a Morley sequence in p and $\bar{a}_{<\omega} \perp_{\text{Cb}(p)} A, \bar{b}_{<\omega}, \bar{c}_{<\omega}$. Since in addition there does exist an automorphism sending $\bar{b}_{<\omega} \mapsto \bar{a}_{<\omega}$, we conclude by composition that every automorphism sending $\bar{b}_{<\omega}$

to $\bar{c}_{<\omega}$ must fix B . In other words, every automorphism sending one Morley sequence in q to another must fix B and in particular every automorphism fixing $\text{Cb}(p)$ must fix B , whence $B \subseteq \text{dcl}(\text{Cb}(p))$. ■_{A.24}

Corollary A.25. *With the assumptions of Theorem A.24:*

- (i) $\bar{a}_0 \downarrow_A \{\bar{a}_i : 0 < i < \omega\}$.
- (ii) $\bar{a}_0 \downarrow_{\{\bar{a}_i : 0 < i < \omega\}} A$.

Proof. The first is immediate and implies $\bar{a}_0 \downarrow_{\text{Cb}(p)} A, \{\bar{a}_i : 0 < i < \omega\}$. By Theorem A.24 we have $\text{Cb}(p) \subseteq \text{dcl}(\{\bar{a}_i : 0 < i < \omega\})$ and the second follows. ■_{A.25}

A.4. Stable type-definable groups and their actions. Let $\langle G, S \rangle$ be a homogeneous space, type-definable in models of a stable theory T . This is to say that G is a type-definable group and S a type-definable set, equipped with a type-definable (and therefore definable) transitive group action $G \times S \rightarrow S$. For convenience let us assume that both are defined without parameters. We will identify G and S with their sets of realisations in a monster model \bar{M} . We will be particularly interested in the case where $S = G$ where G acts on itself either on the left $(g, h) \mapsto gh$ or on the right $(g, h) \mapsto hg^{-1}$.

Definition A.26. (i) A subset $X \subseteq S$ is *generic* if finitely many G -translates of X cover S .
(ii) A partial type $p(x)$ is *generic as a partial type* in S if every logical neighbourhood of p defines in S a generic set. A condition $\varphi(x, \bar{b}) \leq r$ is generic in S if it is as a partial type.
(iii) A subset $X \subseteq G$ is *left-generic* if it is generic under the action of G on itself on the left. We define (partial) types and conditions in the sort of G to be left-generic accordingly. Similarly for right-generic.

Observe that a generic partial type (or condition) does not necessarily define a generic subset of S : we only know that its logical neighbourhoods do.

Lemma A.27. *Let $p(x)$ be a partial type with parameters in A . Then the following are equivalent:*

- (i) *The partial type p is generic in S .*
- (ii) *Every zero set over A which is a logical neighbourhood of p defines in S a generic set.*
- (iii) *Every zero set over A containing p is generic (as a partial type) in S .*

Proof. By Lemma 1.3 every logical neighbourhood of p is a logical neighbourhood of a zero set over A containing p , which is itself a neighbourhood of p . This shows that the last item implies the first. The rest is immediate. ■_{A.27}

Let us denote by $S_S(A)$ the set of all complete types over A implying $x \in S$. Then $S_S(A)$ is compact, and the set of all generic complete types over A is closed. Closed subsets of $S_S(A)$ are in bijection with partial types over A implying $x \in S$, i.e., with

type-definable subsets of S using parameters in A . If $X, Y \subseteq S$ are two such sets, say that Y is a logical neighbourhood of X *relative to S* , in symbols $Y >^S X$, if $[X] \subseteq [Y]^\circ$ where the interior is calculated in $S_S(A)$. This is equivalent to saying that there exists a true logical neighbourhood $Y' > X$ such that $Y = Y' \cap S$. Thus a type-definable set $X \subseteq S$ is generic in S if and only if every relative logical neighbourhood of X in S defines a generic set.

Given any set of parameters A and $g \in G \cap A$, the mapping $s \mapsto gs$ is A -definable inducing a homeomorphism $S_S(A) \rightarrow S_S(A)$, sending $p \mapsto gp$. It sends the closed set defining a type-definable set $X \subseteq S$ to that defining gX , and sends relative logical neighbourhoods of X (over A) to relative logical neighbourhoods of gX . Thus (the partial type defining) X is generic in S if and only if (the partial type defining) gX is. In particular, the action of G on $S_S(\bar{M})$ acts on the space of generic types. Similarly, the mapping $g \mapsto g^{-1}$ induces an automorphism of $S_G(A)$, and a type p is left-generic if and only if p^{-1} is right-generic.

Similarly, if $s \in S \cap A$ then we have a surjective continuous mapping $\rho_s: S_G(A) \rightarrow S_S(A)$ sending $p \mapsto ps$ and $X \mapsto Xs$. If $X \subseteq G$ is type-definable and Y is a relative logical neighbourhood of Xs in S then $\rho_s^{-1}(Y)$ is a relative logical neighbourhood of X . Thus, if X is left-generic (as a partial type) in G , so is Xs in S . In particular, $p \mapsto ps$ sends generic types to generic types.

Our aim is to show the existence of generic types under the assumption that the theory is stable. We follow a path similar to that followed in [Pil96]. Toward this end we will construct an auxiliary multi-sorted structure $\hat{M} = \langle G, S, \dots \rangle$ in a language $\hat{\mathcal{L}}$. If the metric on G is not invariant we may modify it a little, defining $d_G^{\hat{M}}(g, g') = \sup\{d^{\bar{M}}(hg, hg') : h \in G\}$. This is indeed a metric and $d_G^{\hat{M}} \geq d^{\bar{M}}$, so $G^{\hat{M}}$ is complete. Similarly we define $d_S^{\hat{M}}(s, s') = \sup\{d^{\bar{M}}(hs, hs') : h \in G\}$. In addition, for every \mathcal{L} -formula $\varphi(x, \bar{y})$ where x is in the sort of S there will be an additional sort S_φ , consisting of all canonical parameters of instances of φ . The canonical parameter of $\varphi(x, \bar{b})$ will be denote $[\bar{b}]_\varphi$, or $[\bar{b}]$ if there is no ambiguity. We put on it the standard metric, namely $d_\varphi([\bar{b}]_\varphi, [\bar{b}']_\varphi) = \sup\{|\varphi(a, \bar{b}) - \varphi(a, \bar{b}')| : a \in M\}$. The only other symbols in the language are a predicate symbol $\hat{\varphi}(x_S, y_G, z_\varphi)$ for each formula φ as above. We interpret it by $\hat{\varphi}(s, g, [\bar{b}]) = \varphi(g^{-1}s, \bar{b})$. This is uniformly continuous in all variables, since φ is. These definitions make \hat{M} a continuous structure.

If $\langle G, S \rangle$ is definable then \hat{M} is interpretable in \bar{M} and $\hat{T} = \text{Th}_{\hat{\mathcal{L}}}(\hat{M})$ is stable (assuming T is). In the general case, all we know is that \hat{M} is saturated for quantifier-free types in which only $\hat{\varphi}$ appear. It follows from stability in T that each formula $\hat{\varphi}(x, y, z)$, with any partition of the variables, is stable.

For $h \in G$ define a mapping $\theta_h: \hat{M} \rightarrow \hat{M}$ by sending $g \in G$ to hg , $s \in \hat{S}$ to ha , and fixing all the auxiliary sorts. This is easily verified to be an automorphism of \hat{M} . As the action is assumed to be transitive, if $A \subseteq \bigcup_\varphi S_\varphi$ then all elements of S have the same type over A in \hat{M} , and similarly all elements of G .

Lemma A.28. *A condition of the form $\varphi(x, \bar{b}) = 0$ is generic in S if and only if $\hat{\varphi}(x, e, [\bar{b}])$ does not fork over \emptyset .*

Proof. Recall that the $\hat{\mathcal{L}}$ -formula $\hat{\varphi}(x_S, y_G z_\varphi)$ with this (or any other) partition of the variables is stable in \hat{T} .

For left to right we assume that for every $\varepsilon > 0$ the set $X_\varepsilon = \{s \in S : \varphi(s, \bar{b}) \leq \varepsilon\}$ is generic. Find $g_i \in G$ such that $S = \bigcup_{i < \omega} g_i X_\varepsilon$, and find $s \in S$ such that $\text{tp}_{\hat{\varphi}}(s/[\bar{b}]g_{<n})$ does not fork over \emptyset (in symbols $s \not\downarrow^{\hat{\varphi}} [\bar{b}]g_{<n}$). As $s \in \bigcup_{i < \omega} g_i X_\varepsilon$ we may assume that $s \in g_0 X_\varepsilon$, so $\hat{\varphi}(s, g_0, [\bar{b}]) = \varphi(g_0^{-1}s, \bar{b}) \leq \varepsilon$. Thus $\hat{\varphi}(x, g_0, [\bar{b}]) \leq \varepsilon$ does not fork over \emptyset . Applying $\theta_{g_0^{-1}}$ we see that $\hat{\varphi}(x, e, [\bar{b}]) \leq \varepsilon$ does not fork over \emptyset either. It follows that $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork over \emptyset .

For right to left our assumption implies that $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork over $[\bar{b}]$ (as it does not fork over \emptyset). By Lemma A.20 there are $g_i \in G$ and a faithful combination $\psi(x, [\bar{b}]) = \theta(\hat{\varphi}(x, g_i, [\bar{b}]))$ which is definable over $[\bar{b}]$ and such that $\psi(x, [\bar{b}]) = 0$ is consistent. As \hat{M} is saturated for quantifier-free types involving only $\hat{\varphi}$, there is $s \in S$ such that $\psi(s, [\bar{b}]) = 0$. Since all elements of S have the same type over $[\bar{b}]$ in \hat{M} , we see that $\psi(s, [\bar{b}]) = 0$ for all $s \in S$. Assume (toward a contradiction) that there exists $\varepsilon > 0$ such that $\varphi(x, \bar{b}) \leq \varepsilon$ is not generic. By compactness we can find $s \in S$ such that $\varphi(g_i s, \bar{b}) \geq \varepsilon$ for all $i < \omega$, i.e., $\hat{\varphi}(s, g_i, [\bar{b}]) \geq \varepsilon$. Since the combination above was faithful we get $\psi(s, [\bar{b}]) \geq \varepsilon > 0$, a contradiction. ■_{A.28}

Lemma A.29. *Let $p_0(x)$ be a partial type over A . Then p_0 is generic if and only if it extends to a complete generic type over A , i.e., if and only if $[p]$ in $S_S(A)$ contains a generic type. In particular, generic types exist over every set.*

Proof. Right to left is clear, so let us prove left to right. Assume therefore that $[p_0]$ contains no generic type. As the set of generic types is closed, p_0 admits a logical neighbourhood which contains no generic types, and we may assume it to be the zero set of a single condition $\varphi(x, \bar{b}) = 0$. As p_0 is generic in S so is $\varphi(x, \bar{b}) = 0$. By Lemma A.28 $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork over \emptyset in \hat{M} . Let \hat{A} consist of all canonical parameters of the form $[\bar{c}]_\psi$ for $\bar{c} \in A$. Let Ψ consist of all $\hat{\mathcal{L}}$ -formulae of the form $\hat{\psi}(x_S, y_G z_\psi)$. Since all formulae in Ψ are stable there is $\hat{p} \in S_\Psi(\hat{A}, e)$ which does not fork over \emptyset , and such that $\hat{\varphi}(x, e, [\bar{b}])^{\hat{p}} = 0$. Since \hat{M} is saturated for formulae of this form, \hat{p} admits a realisation $s_0 \in S$. Let $p(x)$ be given by the conditions $\psi(x, \bar{c}) = \hat{\psi}(x, e, [\bar{c}])^{\hat{p}}$ for all ψ and $\bar{c} \in A$. Then $s_0 \models p$, so p is a consistent complete type over A . It follows that $p(x) \vdash x \in S$. In addition, p is generic, as can be seen using the fact that \hat{p} does not fork over \emptyset and Lemma A.28. Thus we have found a generic type satisfying $\varphi(x, \bar{b}) = 0$ after all. ■_{A.29}

Lemma A.30. *Assume $A \subseteq B$. Then a type $p \in S_S(B)$ is generic if and only if it does not fork over A and $p|_A$ is generic. In particular, a generic type does not fork over \emptyset .*

Proof. Let us show the last assertion first. Indeed, let $p \in S_S(A)$ be generic. If it forks over \emptyset then there is some condition $\varphi(x, \bar{b}) = 0$ in p which forks, and therefore

divides, over \emptyset . Thus there is an indiscernible sequence $(\bar{b}_i : i < \omega)$ with $\bar{b}_0 = \bar{b}$ such that $\{\varphi(x, \bar{b}) = 0 : i < \omega\}$ is inconsistent. Since $e \in \text{dcl}(\emptyset)$ the sequence $(e, \bar{b}_i : i < \omega)$ is indiscernible as well, and thus the sequence $(e, [\bar{b}_i] : i < \omega)$ is indiscernible in \bar{M} . By Lemma A.28 the condition $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork over \emptyset , so $\{\hat{\varphi}(x, e, [\bar{b}_i]) : i < \omega\}$ is consistent. Since \hat{M} is saturated for such formulae, there is $s \in S$ such that $\hat{\varphi}(s, e, [\bar{b}_i]) = 0$, i.e., $\varphi(s, \bar{b}_i) = 0$, for all $i < \omega$, a contradiction.

This shows that if $p \in S_S(B)$ is generic then it does not fork over any subset $A \subseteq B$, and clearly $p|_A$ is generic. For the converse, assume that $p \in S_S(B)$ does not fork over A and $p_0 = p|_A$ is generic. Replacing p with a non forking extension we may assume that $B = M$ is a very homogeneous and saturated model. By Lemma A.29 there is $p' \in S_S(M)$ extending p_0 which is generic, and by what we have just shown it is also non forking over A . As $p|_A = p'|_A$ there is $f \in \text{Aut}(M/A)$ sending $p'|_{\text{acl}^{eq}(A)}$ to $p|_{\text{acl}^{eq}(A)}$. By uniqueness of the non forking extension it sends p' to p . Thus p is generic. $\blacksquare_{A.30}$

We can now give a proper proof to Fact 3.2.

Theorem A.31. *Let G be a type-definable group in a stable theory T , possibly acting type-definably and transitively on a type-definable set S . Assume all definitions are without parameters. Say that $s \in S$ is dividing-generic over a set A if whenever $s \downarrow_A g$: $gs \downarrow A, g$.*

- (i) *An element $s \in S$ is dividing-generic over a set A if and only if it is generic over A .*
- (ii) *An element $g \in G$ is left-generic over A if and only if g^{-1} is.*
- (iii) *An element $g \in G$ is left-generic if and only if it is right-generic (over A). From now on we will only speak of generic elements and types in G .*
- (iv) *An element $g \in G$ is generic over A if and only if it is generic over \emptyset and $g \downarrow A$.*
- (v) *If $g \in G$, $s \in S$, $g \downarrow_A s$, and at least one of g and s is generic over A , then gs is generic over A .*

Proof. We will use Lemma A.30 and the discussion following Lemma A.27 repeatedly.

Let us first prove the first and last items. If $s \in S$ is generic over A and $s \downarrow_A g$ then s is generic over A, g , whereby gs is generic over A, g . It follows that gs is generic over A and that $gs \downarrow A, g$, so s is dividing-generic.

Similarly, if $g \in G$ left-generic over A and $g \downarrow_A s$ then gs is generic over A, s and thus over A , and $gs \downarrow A, s$. Assume that s is dividing-generic and choose such g . We get $g^{-1} \downarrow_A gs$, and by the previous argument $s = g^{-1}gs$ is generic over A, g and thus over A .

Assume $g \in G$ is left-generic over A . Choose $h \in G$ left-generic over A such that $g \downarrow_A h$. As g is left-generic, gh is left-generic over A and $h \downarrow_A gh$, whereby $h \downarrow_A h^{-1}g^{-1}$. As h is left-generic over A $g^{-1} = hh^{-1}g^{-1}$ is. It follows that g is right-generic as well. Similarly if g is right-generic then g^{-1} is left-generic whereby g is.

The fourth item is just Lemma A.30. $\blacksquare_{A.31}$

A.5. Characterising genericity via global ranks. We have seen that a type of a member of S is generic if and only if the corresponding type in \hat{M} is a non-dividing extension of the unique type over \emptyset , i.e., if its $\hat{\varphi}$ -type has the same Cantor-Bendixson ranks as all of S for every $\varphi(x, \dots)$. Thus the various ε - $\hat{\varphi}$ -Cantor-Bendixson ranks play the role of stratified local ranks characterising genericity. In a superstable (and even more so in an \aleph_0 -stable) theory one would expect a similar characterisation via global Lascar and/or Morley ranks. We will do so in an intentionally brief and sketchy manner.

Morley ranks were defined above: $\text{RM}_\varepsilon(X)$ is the ε -Cantor-Bendixson rank of $[X] \subseteq S_n(\bar{M})$, where \bar{M} is the monster model. The role of the Lascar ranks will be played by the ranks $\text{SU}_\varepsilon(\bar{a}/B)$ defined in [Ben06] (denoted there by $\text{SU}(\bar{a}^\varepsilon/B)$):

Definition A.32. (i) We say that an indiscernible sequence $(\bar{c}_i: i < \omega)$ *could be in* $\text{tp}(\bar{c}/\bar{a}^\varepsilon B)$ if there is a B -indiscernible sequence $(\bar{a}'_i \bar{c}'_i: i < \omega)$ such that $\bar{a}'_0 \bar{c}'_0 = \bar{a} \bar{c}$, $\bar{c}'_{<\omega} \equiv \bar{c}_{<\omega}$ (not necessarily over B !) and $d(\bar{a}'_0, \bar{a}'_1) \leq \varepsilon$.
(ii) We say that $\bar{a} \perp_B \bar{c}$ if every indiscernible sequence in $\text{tp}(\bar{c}/B)$ could be in $\text{tp}(\bar{c}/\bar{a}^\varepsilon B)$.
(iii) We define $\text{SU}_\varepsilon(\bar{a}/B)$ as may be expected: $\text{SU}_\varepsilon(\bar{a}/B) \geq \alpha + 1$ if and only if there is \bar{c} such that $\bar{a}^\varepsilon \not\perp_B \bar{c}$ and $\text{SU}_\varepsilon(\bar{a}/B\bar{c}) \geq \alpha$.

It was shown in [Ben06] that T is supersimple if and only if $\text{SU}_\varepsilon(\bar{a}/B)$ is ordinal for every finite tuple \bar{a} and $\varepsilon > 0$ (and T is superstable if and only if it is stable and supersimple). Moreover, in a supersimple theory T SU_ε ranks characterise independence: $\bar{a} \perp_B C$ if and only if $\text{SU}_\varepsilon(\bar{a}/B) = \text{SU}_\varepsilon(\bar{a}/C)$ for all $\varepsilon > 0$.

Both notions of rank depend inevitably on a metric parameter ε . We may therefore only hope to characterise genericity in case the metric is invariant under the group action, i.e., if the action of each $g \in G$ on S is an isometry. In the case of an \aleph_0 -stable theory this assumption is hardly restrictive. Indeed, G is definable by Theorem 5.1, and as the action is transitive so is S . On S we may define $d'(s, s') = \sup_{g \in G} d(gs, gs')$ and as S is definable we may extend this to a global metric on the sort of \hat{S} .

We have seen that if g is generic over A, t then gt is generic over A . We now prove a converse:

Fact A.33. *Assume $\langle G, S \rangle$ is a type-definable transitive group action with in a stable theory T , $s \in S$ generic over a set A , $t \in S$ satisfying $t \perp_A s$. Then there is $g \in G$, $g \perp_A t$ such that $gs = t$. Moreover, g can be chosen generic over A (i.e., over At).*

Proof. We may assume $A = \emptyset$. First choose $g \in G$ generic, $g \perp s, t$. Then $s \perp_t g$ and as s is generic: $gs \perp g, t$. By standard independence calculus: $g \perp gs, t$. Since the action is transitive we can find $h \in G$ such that $hgs = t$, and we may take it so that $g \perp gs, t, h$. Now $g \perp_t h$ and g is generic over t , so $g' = hg$ is generic over t as required. ■_{A.33}

(The same holds in a simple theory replacing everywhere generic with dividing-generic.)

Theorem A.34. *Assume $\langle G, S \rangle$ is a type-definable transitive group action with an invariant metric in a superstable continuous theory T , $p \in S_S(A)$. Then p is generic if and only if $\text{SU}_\varepsilon(p) = \text{SU}_\varepsilon(S) = \sup\{\text{SU}_\varepsilon(q) : q \in S_S(\emptyset)\}$ for all $\varepsilon > 0$. The same holds in a supersimple theory replacing “generic” with “dividing-generic”. In particular, types of maximal SU_ε -rank exist.*

Proof. It will be enough to prove the supersimple case, using the fact that dividing-generic types exist. We will also use the fact that if $p \in S_n(A)$, $q \in S_m(A)$ and $f : p(\bar{M}) \rightarrow q(\bar{M})$ is A -definable and isometric then $\text{SU}_\varepsilon(p) = \text{SU}_\varepsilon(q)$ for all $\varepsilon > 0$, which we leave as an exercise to the reader.

Assume first that $s \in S$ is dividing-generic over A , $q \in S_S(\emptyset)$. Let $t \models q$, $t \perp_A s$. Then there is $g \perp_A t$ such that $gs = t$. We get $\text{SU}_\varepsilon(s/A) \geq \text{SU}_\varepsilon(s/A, g) = \text{SU}_\varepsilon(t/A, g) = \text{SU}_\varepsilon(t) = \text{SU}_\varepsilon(q)$. In particular $\text{SU}_\varepsilon(s/A) \geq \text{SU}_\varepsilon(s) \geq \text{SU}_\varepsilon(s/A)$, i.e., $\text{SU}_\varepsilon(s/A) = \text{SU}_\varepsilon(s)$. Conversely, let $s \in S$ and assume that $\text{SU}_\varepsilon(s/A) \geq \text{SU}_\varepsilon(q)$ for all $q \in S_S(\emptyset)$ and all $\varepsilon > 0$. Let $g \in G$, $g \perp_A s$. Then $\text{SU}_\varepsilon(gs/A, g) = \text{SU}_\varepsilon(s/A, g) = \text{SU}_\varepsilon(s/A) \geq \text{SU}_\varepsilon(gs) \geq \text{SU}_\varepsilon(gs/A, g)$. Thus equality holds all the way for all $\varepsilon > 0$, whereby $gs \perp_A g$, so s is dividing-generic. ■_{A.34}

In the \aleph_0 -stable case we run into complications since Morley ranks take into account topological neighbourhoods of the types in question which may extend outside the group, unless of course the group is the entire ambient structure. In the latter case there is no problem. In the former, the group action can be extended to an approximate multi-valued operation which is approximately isometric as in the proof of the main theorem.

One can show:

Lemma A.35. *Assume $f : M \rightarrow M$ is an isometric bijection definable with parameters in $A \subseteq M$. Then for all $\varepsilon > 0$ and $a \in M$: $\text{RM}_\varepsilon(a/A) = \text{RM}_\varepsilon(f(a)/A)$. (The same holds if f is an isometric bijection between distinct sorts of M .)*

If $f : X \rightarrow f(X)$ is a mere partial isometry (where X and $f(X)$ are type-definable) definable over A then for all $\varepsilon > \varepsilon' > 0$ we have $\text{RM}_\varepsilon(a/A) \leq \text{RM}_{\varepsilon'}(f(a)/A)$ and vice versa.

Proof. The first assertion is immediate. The second is proved using similar tools to those used to prove Lemma 2.10. ■_{A.35}

In case X is definable, rather than type-definable, so is its image $f(X)$. One may view $\langle X, f(X), \dots \rangle$ with the entire induced structure as a new structure: there is no new structure since X and $f(X)$ are definable (so restricted quantifiers are understood by M) and no loss of structure by separation of variables. In particular the new structure is \aleph_0 -stable and the first case of Lemma A.35 applies. If we knew that the passage to the structure $\langle X, f(X), \dots \rangle$ preserves Morley ranks (as it would in classical logic) we would obtain the equality $\text{RM}_\varepsilon(a/A) = \text{RM}_\varepsilon(f(a)/A)$ in M .

Question A.36. In classical logic types in the structure $\langle X, f(X), \dots \rangle$ have the same Morley ranks as the corresponding types in M . Is it true in continuous logic?

Nonetheless one can use Lemma A.35 to prove:

Theorem A.37. *Assume $\langle G, S \rangle$ is a type-definable transitive group action with an invariant metric in an \aleph_0 -stable continuous theory T , $p \in S_S(A)$. Then p is generic if and only if $\text{RM}_{\varepsilon'}(p) \geq \text{RM}_{\varepsilon}(S)$ for all $\varepsilon > \varepsilon' > 0$. If S occupies an entire sort then this is further equivalent to $\text{RM}_{\varepsilon}(p) = \text{RM}_{\varepsilon}(S)$ for all $\varepsilon > 0$.*

Proof. One follows the same lines as in the superstable/supersimple case, using Lemma A.21 and Lemma A.35. Again we leave the details to the reader. ■_{A.37}

If the answer to Question A.36 is positive then the assumption that S occupies an entire sort is superfluous (as in any case S is definable).

REFERENCES

- [Ben] Itai Ben Yaacov, *Topometric spaces and perturbations of metric structures*, submitted.
- [Ben03] ———, *Thickness, and a categoric view of type-space functors*, *Fundamenta Mathematicæ* **179** (2003), 199–224.
- [Ben05] ———, *Uncountable dense categoricity in cats*, *Journal of Symbolic Logic* **70** (2005), no. 3, 829–860.
- [Ben06] ———, *On supersimplicity and lovely pairs of cats*, *Journal of Symbolic Logic* **71** (2006), no. 3, 763–776.
- [Ber06] Alexander Berenstein, *Definable subgroups of measure algebras*, *Mathematical Logic Quarterly* **52** (2006), no. 4, 367–374.
- [BU] Itai Ben Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, *Transactions of the AMS*, to appear.
- [Pil96] Anand Pillay, *Geometric stability theory*, *Oxford Logic Guides*, vol. 32, Clarendon Press, 1996.
- [Pil98] ———, *Definability and definable groups in simple theories*, *Journal of Symbolic Logic* **63** (1998), 788–796.

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