

STABILITY AND STABLE GROUPS IN CONTINUOUS LOGIC

ITAI BEN YAACOV

ABSTRACT. We develop several aspects of local and global stability in continuous first order logic. In particular, we study type-definable groups and genericity.

INTRODUCTION

Continuous first order logic was introduced by A. Usvyatsov and the author in [BU], with the declared purpose of providing a setting in which classical local stability theory could be developed for metric structures. The actual development of stability theory there is fairly limited, mostly restricted to the definability of φ -types for a stable formula φ , the properties of φ -independence, and in case the theory is stable, properties of independence. Many fundamental results of classical stability theory, and specifically those related to stable groups, are missing there, and it is this gap that the present article proposes to fill.

We assume familiarity with [BU] and follow the notation used therein. Throughout T denotes a continuous theory in a language \mathcal{L} . We do *not* assume that T is complete, so various constants, such as $k(\varphi, \varepsilon)$ of Fact 2.1, are uniform across all completions of T (provided that φ is stable in T , i.e., in every completion of T separately).

By a *model* we always mean a model of T . Whenever this is convenient, we shall assume that such a model \mathcal{M} is embedded elementarily in a large monster model \mathfrak{M} , i.e., in a strongly κ -homogeneous and saturated model, where κ is much bigger than the size of any set of parameters under consideration. Notice that we may not simply choose a single monster model for T , as this would consist of choosing one completion.

1. GENERAL REMINDERS

We shall consider throughout a formula $\varphi(\bar{x}, \bar{y})$ whose variables are split in two groups. We recall from [BU] that a *definable φ -predicate* is a definable predicate $\psi(\bar{x})$, possibly with parameters, which is equivalent to an infinitary continuous combination of instances of φ :

$$\psi(\bar{x}) \equiv \theta(\varphi(\bar{x}, \bar{b}_n))_{n \in \mathbb{N}}, \quad \theta: [0, 1]^{\mathbb{N}} \rightarrow [0, 1] \text{ continuous.}$$

Equivalently, $\varphi(\bar{x})$ is a φ -predicate if it can be approximated arbitrarily well by finite continuous combinations of instances of φ , possibly restricted to the use of the connectives \neg , $\frac{1}{2}$, \div alone.

Local types, i.e., φ -types for a fixed formula φ , are discussed in [BU, Section 6]. For a model \mathcal{M} and a tuple \bar{a} in some extension $\mathcal{N} \succeq \mathcal{M}$, the φ -type of \bar{a} over \mathcal{M} , denoted $\text{tp}_\varphi(\bar{a}/\mathcal{M})$, is the partial type given by $\{\varphi(\bar{x}, \bar{b}) = \varphi(\bar{a}, \bar{b})\}_{\bar{b} \in \mathcal{M}}$. The space of all φ -types over M is denoted $S_\varphi(M)$, and it is a compact Hausdorff quotient of $S_n(M)$. If $\psi(\bar{x})$ is a φ -predicate over M then $\text{tp}_\varphi(\bar{a}/M)$ determines $\psi(\bar{a})$, so we

2000 *Mathematics Subject Classification.* 03C45 ; 03C60 ; 03C90.

Key words and phrases. stable theory ; continuous logic ; definable group.

Author supported by ANR chaire d'excellence junior THEMOMET (ANR-06-CEXC-007) and by Marie Curie research network ModNet.

Revision 986 of 29th September 2009.

may identify ψ with a mapping $\hat{\psi}: S_\varphi(M) \rightarrow [0, 1]$, sending $p \mapsto \psi^p$. Every such mapping is continuous, and conversely, every continuous mapping from $S_\varphi(M)$ to $[0, 1]$ is of this form.

For $A \subseteq M$ we define $S_\varphi(A)$ to be the quotient of $S_\varphi(M)$ where two types are identified if all A -definable φ -predicates agree on them. This is again a compact Hausdorff space, a common quotient of $S_\varphi(M)$ and of $S_k(A)$ (for the appropriate k), and the continuous mappings $S_\varphi(A) \rightarrow [0, 1]$ are precisely the A -definable φ -predicates. In particular this does not depend on the choice of \mathcal{M} .

Lemma 1.1. *Let \mathcal{M} be a structure, $K \subseteq M^\ell$ a (metrically) compact set and let $\varphi(\bar{x}, \bar{y})$ be a formula (or a definable predicate, which we may always name by a new predicate symbol without adding any structure). Then $\inf_{\bar{y} \in K} \varphi(\bar{x}, \bar{y})$ is a φ -predicate (with parameters in K) and for any tuple \bar{x} , the infimum is attained by some $\bar{y} \in K$.*

In particular, K is definable in \mathcal{M} .

Proof. Since K is compact we can find a sequence $\{\bar{c}_n\}_{n \in \mathbb{N}} \subseteq K$ such that for every $\varepsilon > 0$ there is $m = m(\varepsilon)$ such that $K \subseteq \bigcup_{n < m} B(\bar{c}_n, \varepsilon)$. Then $\inf_{\bar{y} \in K} \varphi(\bar{x}, \bar{y})$ is arbitrarily well approximated by formulae of the form $\bigwedge_{n < m} \varphi(\bar{x}, \bar{c}_n)$ as $m \rightarrow \infty$. Finally, the infimum of a continuous function on a compact set is always attained. $\blacksquare_{1.1}$

It will also be convenient to adopt the following somewhat non standard terminology:

Definition 1.2. Let \mathcal{M} be a model, $A \subseteq M$ a subset. We say that \mathcal{M} is *saturated over A* if it is strongly $(|A| + \aleph_0)^+$ -homogeneous and saturated. (In fact, for all intents and purposes it will suffice to require \mathcal{M} to be strongly \aleph_1 -homogeneous and saturated once every member of A is named.)

We say that a partial type $\pi(\bar{x})$ over \mathcal{M} is *A -invariant* if \mathcal{M} is saturated over A and π is fixed by the action of $\text{Aut}(\mathcal{M}/A)$.

An essential notion for the study of definability of types and canonical bases in a stable theory is that of imaginary elements and sorts. Let us give a brief reminder of their construction, as given in [BU]. Consider a definable predicate with parameters in some set A , let us denote it by $\varphi(\bar{x}, A)$. Then we may assume that A is countable, say $A = (a_n)_{n \in \mathbb{N}}$, and express $\varphi(\bar{x}, A)$ as a uniform limit of formulae $\varphi_n(\bar{x}, a_{<n})$. Furthermore, using a forced limit argument, we may assume that the sequence $\varphi_n(\bar{x}, y_{<n})$ converges uniformly to some infinitary definable predicate $\varphi(\bar{x}, Y)$, giving sense to $\varphi(\bar{x}, B)$ for any sequence $B = (b_n)_n$. If \mathcal{M} is any structure, we equip $M^\mathbb{N}$ with the pseudo-metric $d_\varphi(B, C) = \sup_{\bar{x}} |\varphi(\bar{x}, B) - \varphi(\bar{x}, C)|$ and define the *sort of canonical parameters for φ in \mathcal{M}* , denoted $S_\varphi^\mathcal{M}$, as the complete metric space associated to $(M^\mathbb{N}, d_\varphi)$. In other words, we divide $M^\mathbb{N}$ by the kernel $d_\varphi(Y, Z) = 0$, obtaining a true metric on the quotient, and pass to the completion. The predicate $\varphi(\bar{x}, Y)$ is uniformly continuous with respect to Y in the metric d_φ , and therefore passes first to the quotient and then to the completion, thus inducing a uniformly continuous predicate $P_\varphi(\bar{x}, z)$, where $z \in S_\varphi$. We may now add (S_φ, d_φ) as a sort to \mathcal{M} and equip the new structure with an additional predicate symbol for P_φ . This does not add structure to the original sorts of \mathcal{M} , elementary embeddings of structures commute with this construction (by which we mean, in particular, that an elementary embedding $\mathcal{M} \preceq \mathcal{N}$ extends uniquely to $(\mathcal{M}, S_\varphi^\mathcal{M}) \preceq (\mathcal{N}, S_\varphi^\mathcal{N})$), elementary classes and model completeness thereof are respected by this construction, and so. The construction can be slightly simplified when φ only uses finitely many parameters, e.g., if it is an honest formula rather than a definable predicate, but we are going to need the general case.

If $c = [A]$ is the image of A in S_φ then c is indeed a canonical parameter for $\varphi(\bar{x}, A)$, in the sense that an automorphism of \mathcal{M} or of an elementary extension thereof (and such an automorphism extends uniquely to S_φ) fixes c if and only if it fixes the predicate $\varphi(\bar{x}, A)$. By construction we have $P_\varphi(\bar{x}, c) = \varphi(\bar{x}, A)$, and by a convenient abuse of notation we shall permit ourselves to write $\varphi(\bar{x}, c)$ instead of either one.

By an *imaginary sort* we mean any sort added in this fashion, and by *imaginary elements* we mean members of such a sort. We may repeat this construction for any other definable predicate $\psi(\bar{x}', Y')$,

or for any family of predicates. A delicate point here is that even with a countable language one can construct continuum many definable predicates for whose canonical parameters imaginary sorts may be added. For the purposes of stability theory, however, no more imaginary sorts than the size of the language are truly required: we need the sort $S_{d\varphi}$ for each formula φ (see Fact 2.1 for the predicate $d\varphi$). Therefore, with some intentional ambiguity, by \mathcal{M}^{eq} we shall usually mean “ \mathcal{M} along with all the imaginary sorts we are going to need”, e.g., all the sorts $S_{d\varphi}$. By dcl^{eq} , acl^{eq} , etc., we mean the respective operations in the structure \mathcal{M}^{eq} .

Lemma 1.3. *Let A be a set of parameters and let $\varphi(\bar{x})$ be a definable predicate with parameters possibly outside A . Then φ is A -definable if and only if it is A -invariant, i.e., if and only if its canonical base belongs to $\text{dcl}^{eq}(A)$.*

Similarly, a set which is definable (type-definable) with some parameters is definable (type-definable) over A if and only if it is A -invariant.

Proof. First, let us consider an arbitrary surjective continuous mapping $\pi: X \rightarrow Y$ between two compact Hausdorff topological spaces. Then π is also closed, so $F \subseteq Y$ is closed if and only if $\pi^{-1}(F)$ is closed in X . Since π is surjective, $U \subseteq Y$ is open if and only if $\pi^{-1}(U)$ is open, and a mapping $f: Y \rightarrow [0, 1]$ is continuous if and only if $f \circ \pi$ is continuous.

The assertion now follows from applying the previous paragraph to the restriction mapping $S_n(B) \rightarrow S_n(A)$, where $B \supseteq A$ contains all the needed parameters, using the correspondence between type-definable sets and closed sets, and between definable predicates and continuous functions. For a definable set X , just argue for the definable predicate $d(\bar{x}, X)$. ■_{1.3}

Fact 1.4. [BU, Lemma 6.8] *Let $\varphi(\bar{x}, \bar{y})$ be any formula, A a set, \mathcal{M} a saturated model over A , and let $p \in S_\varphi(A)$. Then $\text{Aut}(\mathcal{M}/A)$ acts transitively on the set of extensions of p in $S_\varphi(\text{acl}^{eq}(A))$.*

The following notion and fact also appear (and are used much more extensively) in [Bena, Section 1]:

Definition 1.5. 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 1.6. *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 the zero set 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 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. ■_{1.6}

2. DEFINABILITY AND FORKING OF LOCAL TYPES

Having fixed a theory T , we shall call here a formula $\varphi(\bar{x}, \bar{y})$ *stable* if it is stable in T , that is, if it does not have the order property in any model of T . The order property was defined for continuous logic in [BU], but the reader may simply use Fact 2.1 below as the definition of a stable formula.

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

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

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

$$\text{med}_n(t_{\leq 2n}) = \bigwedge_{w \in [2n+1]^{n+1}} \bigvee_{i \in w} t_i = \bigvee_{w \in [2n+1]^{n+1}} \bigwedge_{i \in w} t_i.$$

Thus in particular $\tilde{\varphi}^0 = \tilde{\varphi}$ and every instance of $\tilde{\varphi}^n$ is a $\tilde{\varphi}$ -predicate.

Fact 2.1. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula. Let \mathcal{M} be a model and let $p \in S_\varphi(\mathcal{M})$ be a complete φ -type. Then*

- (i) *The type p is definable over M , i.e., there exists an M -definable $\tilde{\varphi}$ -predicate $d_p\varphi(\bar{y})$ such that $\varphi(x, \bar{b})^p = d_p\varphi(\bar{b})$ for all $\bar{b} \in M$. Moreover, this definition is uniform, in the sense that there exists an infinitary definable predicate $d\varphi(\bar{y}, X)$ which only depends on φ such that $d_p\varphi(\bar{y})$ is equal to an instance $d\varphi(\bar{y}, C)$ where $C \subseteq M$.*
- (ii) *For every $\varepsilon > 0$ there exists a number $k = k(\varphi, \varepsilon) \in \mathbb{N}$ (which depends on φ and on ε but not on p) and a tuple $\bar{c}_{\leq 2k} = \bar{c}_{\leq 2k(\varphi, \varepsilon)}$ in \mathcal{M} (which does depend on p) such that*

$$|d_p\varphi(\bar{y}) - \tilde{\varphi}^k(\bar{y}, \bar{c}_{\leq 2k})| < \varepsilon.$$

- (iii) *Assume moreover that \mathcal{M} is saturated over some subset $A \subseteq M$. Then in the previous item the tuples $\bar{c}_{\leq 2k}$ can be chosen so that each \bar{c}_n realises $p \upharpoonright_{A\bar{c}_{<n}}$.*

Proof. The first two items are taken from [BU, Lemma 7.4 and Proposition 7.6]. The third item, while not explicitly stated there, is immediate from the proof. ■_{2.1}

By abuse of notation we may sometimes write $d_p\varphi(\bar{y}) = d\varphi(\bar{y}, c)$, where $c \in M^{eq}$ is the canonical parameter for the definition. This canonical parameter is called the *canonical base* of p , denoted $\text{Cb}(p)$.

We recall that for $A \subseteq B \subseteq \mathcal{M}$, $p \in S_\varphi(B)$ *does not fork* over A if it admits an extension $p_1 \in S_\varphi(M)$ which is definable over $\text{acl}^{eq}(A)$. In this case p_1 itself does not fork over A or B . A type over a model clearly admits a unique non forking extension to any larger model (and therefore set), so this definition does not depend on the choice of \mathcal{M} .

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 . A minor enhancement of that result will be quite useful.

Lemma 2.2 (Existence of non forking extensions). *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, A a set, $\mathcal{M} \supseteq A$ a saturated model over A . Let $\pi(\bar{x})$ be a consistent A -invariant partial type over M . Then there exists $p \in S_\varphi(M)$ compatible with π which does not fork over A .*

Proof. Let $X = \{p \in S_\varphi(M) : p \cup \pi \text{ is consistent}\}$. 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. ■_{2.2}

Corollary 2.3. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, A a set, $\mathcal{M} \supseteq A$ a saturated model over A . Then $p \in S_\varphi(M)$ does not fork over A if and only if it is $\text{acl}^{eq}(A)$ -invariant.*

Proof. Left to right follows from the definition, right to left from Lemma 2.2. ■_{2.3}

Corollary 2.4. *Let A be a set, $\mathcal{M} \supseteq A$ a saturated model over A and $\pi(\bar{x})$ a consistent A -invariant partial type over M . Then there exists a complete type p such that $\pi \subseteq p \in S_n(M)$, and for every stable formula $\varphi(\bar{x}, \bar{y})$ the restriction $p|_{\varphi} \in S_{\varphi}(M)$ does not fork over A .*

Proof. We may assume that $A = \text{acl}^{eq}(A)$. Index all stable formulae of the form $\varphi_i(\bar{x}, \bar{y}_i)$ by $i < \lambda$. We define an increasing sequence of consistent A -invariant partial types π_i over M , starting with $\pi_0 = \pi$. Given π_i , by Lemma 2.2 there is $p_i \in S_{\varphi_i}(M)$ be non forking over A and compatible with π_i , so $\pi_{i+1} = \pi_i \cup p_i$ is consistent and A -invariant. For limit i we define $\pi_i = \bigcup_{j < i} \pi_j$. Finally, let $p \in S_n(M)$ be any completion of π_{λ} . Then p will do. ■_{2.4}

It follows that if the theory is stable then every complete type over a set admits non forking extensions. The same fact was proved in [BU] using a somewhat longer “gluing” argument.

Fact 2.5 (Symmetry [BU, Proposition 7.16]). *Let \mathcal{M} be a model, $p(\bar{x}) \in S_{\varphi}(M)$, $q(\bar{y}) \in S_{\tilde{\varphi}}(M)$. Then $d_p \varphi(\bar{y})^q = d_q \tilde{\varphi}(\bar{x})^p$.*

Proposition 2.6. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, \mathcal{M} a model, $A \subseteq M$. For each $\bar{b} \in M$ let $\chi_{\bar{b}}(\bar{x})$ be the definition of a non forking extension of $\text{tp}_{\tilde{\varphi}}(\bar{b}/\text{acl}^{eq}(A))$ to M .*

- (i) *Each $\chi_{\bar{b}}(\bar{x})$ is a definable φ -predicate over $\text{acl}^{eq}(A)$.*
- (ii) *A φ -type $p \in S_{\varphi}(M)$ does not fork over A if and only if $\varphi(\bar{x}, \bar{b})^p = \chi_{\bar{b}}(\bar{x})^p$ for all $\bar{b} \in M$.*
- (iii) *A φ -type over $\text{acl}^{eq}(A)$ is stationary, i.e., admits a unique non forking extension to every larger set.*
- (iv) *Let $r(\bar{x}) = \{|\varphi(\bar{x}, \bar{b}) - \chi_{\bar{b}}(\bar{x})| = 0\}_{\bar{b} \in M}$. Then the partial type $r(\bar{x})$ defines the set of φ -types which do not fork over A :*

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

- (v) *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 2.1 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 $q \in S_{\tilde{\varphi}}(M)$ be the non forking extension defined by $\chi_{\bar{b}}$. Assume $p \in S_{\varphi}(M)$ does not fork over M , so $d_p \varphi(\bar{y})$ is a $\tilde{\varphi}$ -predicate over $\text{acl}^{eq}(A)$. By Fact 2.5,

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

Conversely, assume that $\varphi(\bar{x}, \bar{b})^p = \chi_{\bar{b}}(\bar{x})^p$ for all $\bar{b} \in M$, and let $p' \in S_{\varphi}(M)$ be any non forking extension of $p|_{\text{acl}^{eq}(A)}$. Then $p = p'$, proving also the third item. The fourth item is just a re-statement of the second.

For the last item 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 . ■_{2.6}

Proposition 2.6.(iii) is the analogue of the finite equivalence relation theorem in continuous logic. It has already appeared as [BU, Proposition 7.17]. In case $p \in S_{\varphi}(A)$ is stationary, the unique non forking extension to $B \supseteq A$ will be denoted $p|_B$. Similarly, we write $d_p \varphi$ for the definition of $p|_M$ where $\mathcal{M} \supseteq A$ is any model (and this does not depend on the choice of \mathcal{M}). Thus, in hindsight, in the statement of Proposition 2.6, the definitions $\chi_{\bar{b}}$ are uniquely determined, $\chi_{\bar{b}} = d_{\bar{b}/\text{acl}^{eq}(A)} \tilde{\varphi}$.

Corollary 2.7. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, A a set, \mathcal{M} a saturated model over A . Let $p \in S_{\varphi}(A)$. Then $\text{Aut}(\mathcal{M}/A)$ acts transitively on the set of non forking extensions of p in $S_{\varphi}(M)$. If T is stable and $p \in S_n(A)$ then $\text{Aut}(\mathcal{M}/A)$ acts transitively on the set of non forking extensions of p to \mathcal{M} .*

Proof. The first assertion follows from Fact 1.4 and Proposition 2.6.(iii). For the second we need the even easier fact that $\text{Aut}(\mathcal{M}/A)$ acts transitively on the extensions of a complete type $p \in S_n(A)$ to $\text{acl}^{eq}(A)$. ■_{2.7}

Corollary 2.8. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula \mathcal{M} a model. A type $p \in S_\varphi(M)$ is definable over A if and only if it does not fork over A and $p \upharpoonright_A$ is stationary.*

Proof. We may assume that \mathcal{M} saturated over A . Let $p' \in S_\varphi(M)$ be any non forking extension of $p \upharpoonright_A$. By Corollary 2.7 there is an automorphism $f \in \text{Aut}(\mathcal{M}/A)$ sending p to p' . If p is definable over A then $p' = f(p) = p$. Conversely, if p does not fork over A and $p \upharpoonright_A$ is stationary then $\text{Aut}(\mathcal{M}/A)$ fixes p and therefore fixes $d_p \varphi$. By Lemma 1.3, the latter is over A . ■_{2.8}

Corollary 2.9. *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 $p_0 = q \upharpoonright_\varphi \in S_\varphi(A)$. Then q is compatible with every non forking extension of p_0 .*

Proof. By Lemma 2.2, q is compatible with at least one non forking extension of p to the monster model. By Corollary 2.7 it is compatible with all of them. ■_{2.9}

We pass to forking of single conditions.

Definition 2.10. Let $\varphi(\bar{x}, \bar{b})$ be an instance of a stable formula, A a set. We say that a condition $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A if there exists a φ -type $p \in S_\varphi(A\bar{b})$ non forking over A such that $\varphi(\bar{x}, \bar{b})^p \leq r$.

Proposition 2.11. *Let $\varphi(\bar{x}, \bar{b})$ be an instance of a stable formula, A a set of parameters. Then the following are equivalent:*

- (i) *The condition $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A .*
- (ii) *Every family of $\text{acl}^{eq}(A)$ -conjugates of $\varphi(\bar{x}, \bar{b}) \leq r$ is consistent.*
- (iii) *For every set $B \supseteq A, \bar{b}$ there exists a complete type $p \in S_n(B)$ such that $p \upharpoonright_\psi$ does not fork over A for any stable formula ψ (if T is stable: p does not fork over A) and $\varphi(\bar{x}, \bar{b})^p \leq r$.*

Proof. (i) \implies (ii). Let p witness that $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A . Then any non forking extension of p to a large model is $\text{acl}^{eq}(A)$ -invariant.

(ii) \implies (iii). We may assume that $B = \mathcal{M}$ is saturated over A . Let π consist of all the $\text{acl}^{eq}(A)$ -conjugates of $\varphi(\bar{x}, \bar{b}) \leq r$ in \mathcal{M} . It is consistent by assumption and $\text{acl}^{eq}(A)$ -invariant by construction so we may apply Corollary 2.4.

(iii) \implies (i). Immediate. ■_{2.11}

We may define the *non forking degree* of $\varphi(\bar{x}, \bar{b})$ over A to be

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

An easy compactness argument shows that the infimum is attained and the condition $\varphi(\bar{x}, \bar{b}) \leq \text{nf}(\varphi(\bar{x}, \bar{b})/A)$ does not fork over A . In addition, by the existence of non-forking types we have $\text{nf}(\varphi(\bar{x}, \bar{b})/A) + \text{nf}(\neg\varphi(\bar{x}, \bar{b})/A) \leq 1$.

Definition 2.12. 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 using \vee and \wedge alone is faithful (so in particular the median value connective $\text{med}_n: [0, 1]^{2n+1} \rightarrow [0, 1]$ is). Similarly, any uniform limit of faithful combinations is faithful.

Lemma 2.13. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula. Let $A = \text{acl}^{\text{eq}}(A)$ be a set of parameters, \bar{a} a tuple, $|\bar{x}| = |\bar{a}|$. Let $p = \text{tp}_\varphi(\bar{a}/A)$. Then $d_p\varphi(\bar{x}, \bar{y})$ is a faithful combination of A -conjugates of $\varphi(\bar{a}, \bar{y})$.*

Proof. By the preceding discussion and the last item of Fact 2.1. ■_{2.13}

Lemma 2.14. *Let $\varphi(\bar{x}, \bar{b})$ be an instance of a stable formula, A a set of parameters. Then there exists an A -definable predicate $\psi(\bar{x})$ such that for every tuple \bar{a} (not necessarily in A):*

$$\begin{aligned} \psi(\bar{a}) &= \inf\{\varphi(\bar{x}, \bar{b})^p : p \in S_\varphi(A\bar{b}) \text{ is a non forking extension of } \text{tp}_\varphi(\bar{a}/A)\} \\ &= \inf\{\varphi(\bar{a}, \bar{y})^q : q \in S_\varphi(A\bar{a}) \text{ is a non forking extension of } \text{tp}_{\bar{\varphi}}(\bar{b}/A)\}. \end{aligned}$$

Moreover, $\psi(\bar{x})$ can be taken to be a faithful combination of A -conjugates of $\varphi(\bar{x}, \bar{b})$.

Proof. Fix a model $\mathcal{M} \supseteq A, \bar{b}$, saturated over A . Let $G = \text{Aut}(\mathcal{M}/A)$. Let $q \in S_{\bar{\varphi}}(\mathcal{M})$ be the unique non forking extension of $\text{tp}_{\bar{\varphi}}(\bar{b}/\text{acl}^{\text{eq}}(A))$. Let $\chi(\bar{x}, c) = d_q\bar{\varphi}(\bar{x})$, where $c \in \text{acl}^{\text{eq}}(A)$ is the canonical parameter for the definition. By the previous Lemma, $\chi(\bar{x}, c)$ is a faithful combination of $\text{acl}^{\text{eq}}(A)$ -conjugates of $\varphi(\bar{x}, \bar{b})$.

Let C be the set of A -conjugates of c . Since c is algebraic over A , C is (metrically) compact. By Lemma 1.1 $\psi(\bar{x}) = \inf_{c' \in C} \chi(\bar{x}, c')$ 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. Thus $\psi(\bar{x})$ is a faithful combination of A -conjugates of $\varphi(\bar{x}, \bar{b})$, and it is clearly over A .

We may assume that $\bar{a} \in M$, and let $p \in S_\varphi(M)$ be the unique non forking extension of $\text{tp}_\varphi(\bar{a}/\text{acl}^{\text{eq}}(A))$. Then

$$\begin{aligned} \psi(\bar{a}) &= \inf_{g \in G} \chi(\bar{a}, gc) = \inf_{g \in G} d_{gq}\bar{\varphi}(\bar{a}) = \dots \\ &\dots = \inf_{g \in G} d_{g^{-1}p}\varphi(\bar{y})^q = \inf_{g \in G} \varphi(\bar{x}, \bar{b})^{gp}, \\ &\dots = \inf_{g \in G} \varphi(\bar{a}, \bar{y})^{gq}. \end{aligned}$$

Since $\{gp\}_{g \in G}$ and $\{gq\}_{g \in G}$ are the sets of non forking extensions of $\text{tp}_\varphi(\bar{a}/A)$ and of $\text{tp}_{\bar{\varphi}}(\bar{b}/A)$, respectively, to M , we are done. ■_{2.14}

Theorem 2.15 (Open Mapping Theorem). *Assume T is stable, and let $A \subseteq B$ be any sets of parameters. Let $X \subseteq S_n(B)$ be the set of types which do not fork over A . Then X is compact and the restriction mapping $\rho_A: X \rightarrow S_n(A)$ sending $p \mapsto p|_A$ is an open continuous surjective mapping.*

Proof. We already know that X is compact and that ρ_A is continuous and surjective.

Consider a basic open subset $U \subseteq X$, of the form $U = X \cap [\varphi(\bar{x}, \bar{b}) < 1]$. Let $\psi(\bar{x})$ be as in Lemma 2.14 and let $V = [\psi(\bar{x}) < 1] \subseteq S_n(A)$. By Corollary 2.4 every φ -type over B which does not fork over A extends to a complete type over B which does not fork over A , whence $V = \rho_A(U)$. ■_{2.15}

Notice that a similar proof yields that if $\varphi(\bar{x}, \bar{y})$ is stable then the restriction mapping $\rho_{A, \varphi}: X_\varphi \rightarrow S_\varphi(A)$ is open, where $X_\varphi \subseteq S_\varphi(B)$ denotes the set of φ -types which do not fork over A .

It follows from Lemma 2.14 that a $\bar{\varphi}$ -type (and therefore a φ -type) over an arbitrary set A is definable over A , but of course the same definition applied to a larger set need not give a consistent complete type. This yields the following (adaptation of a) classical result:

Theorem 2.16 (Separation of variables). *Let $\varphi(\bar{x}, \bar{b})$ be an instance of a stable formula, and let X be a type-definable set in the sort of \bar{x} , say with parameters in A . Then there is a subset (at most countable) $B \subseteq X$ and a B -definable predicate $\psi(\bar{x})$ such that $\psi(\bar{x})|_X = \varphi(\bar{x}, \bar{b})|_X$.*

Moreover, $\psi(\bar{x})$ can be taken to be a faithful combination of instances $\varphi(\bar{x}, \bar{b}')$ such that $\bar{b}' \equiv_B \bar{b}$ (or even $\bar{b}' \equiv_{B'} \bar{b}$ where $B' \subseteq X$ is an arbitrary small subset).

Proof. Fix a model $\mathcal{M} \supseteq A, \bar{b}$, saturated over A , and let $C = X(\mathcal{M})$. Let $\psi(\bar{x})$ be as in Lemma 2.14. Then $\psi(\bar{x})$ is definable over C and therefore over B where $B \subseteq C$ is an appropriate countable subset. Then for all $\bar{a} \in C$ we have $\psi(\bar{a}) = \varphi(\bar{a}, \bar{y})^{\text{tp}_{\bar{x}}(\bar{b}/C)} = \varphi(\bar{a}, \bar{b})$. Now let \mathfrak{M} be the monster model and $\bar{a} \in X = X(\mathfrak{M})$. By saturation of \mathcal{M} we can find there some $\bar{a}' \equiv_{AB\bar{b}} \bar{a}$. Then $\bar{a}' \in C$ and $\varphi(\bar{a}, \bar{b}) = \varphi(\bar{a}', \bar{b}) = \psi(\bar{a}') = \psi(\bar{a})$, as desired.

The moreover part follows from the proof. ■_{2.16}

It follows that if X is an A -type-definable set and $Y \subseteq X$ is a type-definable subset, then Y is type-definable over AB for some countable $B \subseteq X$. If Y is a definable set then it is definable over AB (by Lemma 1.3, since the definable predicate $d(\bar{x}, Y)$ is AB -invariant).

Proposition 2.17. *Let $\varphi(\bar{x}, \bar{b})$ be an instance of a stable formula, A a set of parameters. 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.*

Proof. Fix a model $\mathcal{M} \supseteq A, \bar{b}$ saturated over A .

(i) \implies (ii). Let $\psi(\bar{x})$ be as in Lemma 2.14. Let also $p \in S_\varphi(A\bar{b})$ be non forking over A such that $\varphi(\bar{x}, \bar{b})^p \leq r$. Then $\psi(\bar{x})^p \leq \varphi(\bar{x}, \bar{b})^p \leq r$.

(ii) \implies (i). Let $\psi(\bar{x}) = \theta(\varphi(\bar{x}, \bar{b}_n))_{n \in \mathbb{N}}$ be definable over A as in the assumption (so $\bar{b}_n \equiv_A \bar{b}$ and θ is a faithful continuous connective).

By Lemma 2.2 there exists $p \in S_\varphi(M)$ compatible with $\psi(\bar{x}) \leq r$ and non forking over A , so in particular $\text{acl}^{\text{eq}}(A)$ -invariant. Then $\inf_n \varphi(\bar{x}, \bar{b}_n)^p \leq r$ by faithfulness, so for all $r' > r$ there exists n such that $\varphi(\bar{x}, \bar{b}_n)^p < r'$. Up to an automorphism fixing A we may assume that $\varphi(\bar{x}, \bar{b})^p < r'$, and by invariance $\varphi(\bar{x}, \bar{b}')^p < r'$ for every $\bar{b}' \equiv_{\text{acl}^{\text{eq}}(A)} \bar{b}$.

We have thus shown that for every $r' > r$, any set of $\text{acl}^{\text{eq}}(A)$ -conjugates of $\varphi(\bar{x}, \bar{b}) \leq r'$ is consistent. By compactness the same holds for $\varphi(\bar{x}, \bar{b}) \leq r$. ■_{2.17}

3. HEIRS AND CO-HEIRS

We turn to study co-heirs, and more generally, approximately realised partial types, in continuous logic. In the context of stability, approximate realisability serves as a criterion for non forking. For an earlier treatment of co-heirs in the context of metric structures see [Ben05, Section 3.2].

Definition 3.1. Let $A \subseteq B$ be two sets of parameters. We say that a partial type π over B is *approximately realised* in A if every logical neighbourhood (Definition 1.5) of π over B is realised in A .

If \mathcal{M} is a model, $B \supseteq M$, and $p \in S_n(B)$ is approximately realised in M , we may say that p is a *co-heir* of its restriction to M .

Remark 3.2. (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 type over a model \mathcal{M} is always approximately realised there. (This is essentially the Tarski-Vaught Criterion.)

Fact 3.3. *Let $A \subseteq B$ and let $\pi(\bar{x})$ be a partial type over B .*

- (i) *Let $X \subseteq S_n(B)$ consist of all types over B which are realised in A , $[\pi] \subseteq S_n(B)$ the closed set defined by π . Then π is approximately realised in A if and only if $[\pi] \cap \overline{X} \neq \emptyset$. In particular, \overline{X} is the set of all complete n -types over B which are approximately realised in A .*
- (ii) *If $C \supseteq B$ then π is approximately realised in A as a partial type over B if and only if it is approximately realised in A as a partial type over C .*

- (iii) If π is approximately realised in A then it extends to a complete type $\pi \subseteq p \in S_n(B)$ which is approximately realised in A .
- (iv) A type over a model \mathcal{M} admits extensions to arbitrary sets which are approximately realised in M .

Proof. We prove the first two items together. Clearly if π is approximately realised in A as a partial type over C then it is approximately realised in A as a partial type over B , in which case every neighbourhood of $[\pi]$ in $S_n(B)$ intersects X and by a compactness argument $[\pi]$ intersects \overline{X} . Finally, assume $[\pi] \cap \overline{X} \neq \emptyset$ and assume that $\pi \vdash \varphi(\bar{x}) > 0$. Let $Y = [\varphi = 0] \subseteq S_n(C)$ and let Z be its projection to $S_n(B)$. Then Z is compact, $Z \cap [\pi] = \emptyset$, so $U = S_n(B) \setminus Z$ is a neighbourhood of $[\pi]$. By assumption there exists $\bar{a} \in A$ such that $\text{tp}(\bar{a}/B) \in X \cap U$. Then $\text{tp}(\bar{a}/C) \notin Y$, i.e., $\varphi(\bar{a}) > 0$, as desired.

For the third item, any $p \in [\pi] \cap \overline{X}$ will do. For the fourth, use the fact that a type over a model is approximately realised there. ■_{3.3}

Fact 3.4. *Let \mathcal{N} be a model saturated over $A \subseteq N$. If $p \in S_n(N)$ or $p \in S_\varphi(N)$ is approximately realised in A then it is A -invariant.*

Proof. We only consider the case $p \in S_\varphi(N)$, since the case $p \in S_n(N)$ follows from it. Say $\bar{b}, \bar{c} \in N$, $\bar{b} \equiv_A \bar{c}$, and let $\varepsilon > 0$ be given. By assumption there is $\bar{a} \in A$ 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 $\bar{b} \equiv_A \bar{c}$ we have in particular $\varphi(\bar{a}, \bar{b}) = \varphi(\bar{a}, \bar{c})$ 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 desired. ■_{3.4}

Lemma 3.5. *Let $A \subseteq B$, $p(\bar{x}) \in S_n(B)$ approximately realised in A , and assume $\varphi(\bar{x}, \bar{y})$ is stable. Then $p \upharpoonright_\varphi \in S_\varphi(B)$ does not fork over A .*

Proof. Let $\mathcal{N} \supseteq B$ be saturated over A and let $q \in S_n(N)$ extend p , still approximately realised in A . Then q , and thus $q \upharpoonright_\varphi$, are A -invariant, so $q \upharpoonright_\varphi$ does not fork over A and neither does $p \upharpoonright_\varphi$. ■_{3.5}

Proposition 3.6. *Let $\varphi(\bar{x}, \bar{y})$ be a stable formula, \mathcal{M} a model, $A \supseteq M$. Let also $p(\bar{x}) \in S_\varphi(A)$ be 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 $p'(\bar{x}) \in S_n(A)$ extending p which is approximately realised in M and use Lemma 3.5.

(iii) \implies (i). Find $q'(\bar{x}) \in S_n(A)$ extending q which is approximately realised in M . Then $q' \upharpoonright_\varphi$ is non forking over M by Lemma 3.5, so it must be the unique non forking extension of $p \upharpoonright_M = q \upharpoonright_\varphi$. Therefore $q \cup p \subseteq q'$ is approximately realised in M . ■_{3.6}

Similarly,

Proposition 3.7. *Assume T is stable. Let \mathcal{M} be a model of T , $A \supseteq M$, $p(\bar{x}) \in S_n(A)$. Then the following are equivalent:*

- (i) p does not fork over M .
- (ii) p is approximately realised in M .

If $A = \mathcal{N} \succeq \mathcal{M}$ is saturated over M then these are further equivalent to

- (iii) p is M -invariant.

Definition 3.8. Let \mathcal{M} be a model, $\mathcal{M} \subseteq B$. A type $p \in S_n(B)$ is said to be an *heir* of its restriction to M if for every formula $\varphi(\bar{x}, \bar{b}, \bar{m})$ with $\bar{b} \in B$ and $\bar{m} \in M$, and for every $\varepsilon > 0$, there are $\bar{b}' \in M$ such that $|\varphi(\bar{x}, \bar{b}, \bar{m}) - \varphi(\bar{x}, \bar{b}', \bar{m})|^p < \varepsilon$.

Clearly every type over a model is an heir of itself. Also, it is not difficult to check that if \mathcal{M} is a model and \bar{a}, \bar{b} are two tuples possibly outside \mathcal{M} then

$$\text{tp}(\bar{a}/M\bar{b}) \text{ is an heir of } \text{tp}(\bar{a}/M) \iff \text{tp}(\bar{b}/M\bar{a}) \text{ is a co-heir of } \text{tp}(\bar{b}/M).$$

Finally, a standard compactness argument yields that if $\mathcal{M} \subseteq B \subseteq C$ and $p \in S_n(B)$ is an heir of $p \upharpoonright_M$ then it admits an extension to C which is an heir as well.

Lemma 3.9. *Let \mathcal{M} be a model, $p(\bar{x}) \in S_n(M)$. Then p is definable if and only if it has a unique heir to every superset $B \supseteq \mathcal{M}$.*

Proof. (We follow Poizat [Poi85, Théorème 11.07].) For left to right, assume p is definable and let $q \in S_n(B)$ be an heir of p , where $B \supseteq M$. Let $\varphi(\bar{x}, \bar{b})$ be a formula over B and let $d_p\varphi(\bar{y}, c)$ be the φ -definition of p , $c \in M$. Assume that $d_p\varphi(\bar{b}, c) \neq \varphi(\bar{x}, \bar{b})^q$, i.e., $|d_p\varphi(\bar{b}, c) - \varphi(\bar{x}, \bar{b})|^q > 0$. Then there is $\bar{b}' \in M$ such that $|d_p\varphi(\bar{b}', c) - \varphi(\bar{x}, \bar{b}')|^q > 0$, a contradiction.

Conversely, assume p admits a unique heir to every structure. let \mathcal{L}' be \mathcal{L} along with a new predicate symbol $D_\varphi(\bar{y})$ for each formula $\varphi(\bar{x}, \bar{y})$ (here \bar{x} is fixed, \bar{y} may vary with φ). We define an expansion \mathcal{M}' of \mathcal{M} by interpreting $D_\varphi(\bar{b}) = \varphi(\bar{x}, \bar{b})^p$. Assume now that $\mathcal{N}' \succeq_{\mathcal{L}'} \mathcal{M}'$ and let

$$\mathcal{N} = \mathcal{N}' \upharpoonright_{\mathcal{L}}, \quad q = \{\varphi(\bar{x}, \bar{b}) = D_\varphi(\bar{b})\}_{\bar{b} \in N}.$$

It is not difficult to see that $q(\bar{x}) \in S_n(N)$ is a complete, consistent type, and that it is moreover an heir of p . Since, by assumption, the heir is unique, \mathcal{N}' is the unique expansion of \mathcal{N} which is an elementary extension of \mathcal{M}' . In other words, a model of $\text{Diag}(\mathcal{M})$ admits at most one expansion to a model of $\text{Diag}(\mathcal{M}')$. By Beth's Theorem (see [Ben09]) for each formula $\varphi(\bar{x}, \bar{y})$ there exists an \mathcal{M} -definable predicate $d_p\varphi(\bar{y})$ such that $\text{Diag}(\mathcal{M}') \vdash D_\varphi = d_p\varphi$. In particular, $\varphi(\bar{x}, \bar{m})^p = D_\varphi(\bar{m}) = d_p\varphi(\bar{m})$ for every $\bar{m} \in M$, and p is definable. ■_{3.9}

Notice that for a pair of models $\mathcal{M} \subseteq \mathcal{N}$ we could have defined a notion of a φ -type over a \mathcal{N} being an heir of its restriction to \mathcal{M} , in which case Lemma 3.9 holds, with the same proof, for local types.

Theorem 3.10. *The following are equivalent for a theory T :*

- (i) *The theory T is stable.*
- (ii) *Every type over a model has a unique co-heir to any superset.*
- (iii) *Every type over a model has a unique heir to any superset.*

Proof. (i) \implies (ii). Assume T is stable, $\mathcal{M} \subseteq B$, and $q \in S_n(B)$ is a co-heir of $p = q \upharpoonright_M$. Let $\mathcal{N} \supseteq B$ be saturated over \mathcal{M} and let $q' \in S_n(N)$ extend q , also a co-heir of p . Then q' is M -invariant and therefore the unique non forking extension of p to N . Thus q is the unique non forking extension of p to B .

(ii) \implies (iii). Let \mathcal{M} be a model, $p \in S_n(M)$. In order to show that p has a unique heir to every $B \supseteq M$ it is enough to consider the case $B = M\bar{b}$ where \bar{b} is a finite tuple. So indeed, assume that \bar{a} realises an heir of p to $M\bar{b}$. Then $\text{tp}(\bar{b}/M\bar{a})$ is a co-heir of $\text{tp}(\bar{b}/M)$ and by assumption it is uniquely determined by $\text{tp}(\bar{b}/M)$ and by \bar{a} . It follows that $\text{tp}(\bar{a}/M\bar{b})$ is uniquely determined by \bar{b} and $\text{tp}(\bar{a}/M)$, as desired.

(iii) \implies (i). The assumption and Lemma 3.9 yield that every type is definable, so T is stable. ■_{3.10}

Using the local version of Lemma 3.9 alluded to above we can prove a local version of Theorem 3.10, namely that $\varphi(\bar{x}, \bar{y})$ is stable if and only if every φ -type over a model admits a unique co-heir to larger sets if and only if every φ -type over models admits a unique heir to larger models.

4. INVARIANT TYPES, INDISCERNIBLE SEQUENCES AND DIVIDING

Fact 4.1. *Let \mathcal{M} be a model saturated over $A \subseteq M$, and let $p \in S_n(M)$ be A -invariant. Let $(\bar{a}_n)_{n \in \mathbb{N}} \subseteq M$ be a sequence constructed inductively, choosing each \bar{a}_n to realise $p \upharpoonright_{A\bar{a}_{<n}}$.*

Then the sequence $(\bar{a}_n)_{n \in \mathbb{N}}$ is A -indiscernible, and its type over A depends only on p .

Proof. Standard. ■_{4.1}

The common type over A of such sequences will be denoted by $p^{(\omega)} \upharpoonright_A$. For every finite or countable $B \subseteq M$ we may construct $p^{(\omega)} \upharpoonright_{A \cup B}$ just as well. By a gluing argument, $p^{(\omega)} = \bigcup \{p^{(\omega)} \upharpoonright_{A \cup B} : B \in [M]^{\aleph_0}\}$ is a complete type of an M -indiscernible sequence in p , and is of course A -invariant.

Lemma 4.2. *Let A be a set, $\varphi(\bar{x}, \bar{y})$ a stable formula, $p \in S_\varphi(A)$ a stationary φ -type. Let $\mathcal{M} \supseteq A$ be saturated over A , and let $p \subseteq q \in S_n(M)$, q invariant over A . Let $(\bar{c}_n)_{n \in \mathbb{N}} \models q^{(\omega)} \upharpoonright_A$ be an A -indiscernible sequence as constructed in Fact 4.1.*

Then the sequence $(\tilde{\varphi}^n(\bar{y}, \bar{c}_{\leq 2n}))_{n \in \mathbb{N}}$ converges uniformly to the definition $d_p \varphi(\bar{y})$ at a rate which only depends on φ .

Proof. Since $q \upharpoonright_\varphi$ is A -invariant, it does not fork over A , so $d_p \varphi(\bar{y}) = d_q \varphi(\bar{y})$.

Fix $\varepsilon > 0$. By Fact 2.1 there is $k = k(\varphi, \varepsilon)$ and a sequence $(\bar{c}'_n)_{n \leq 2k} \subseteq M$ such that $|d_p \varphi(\bar{y}) - \tilde{\varphi}^k(\bar{y}, \bar{c}'_{\leq 2k})| \leq \varepsilon$, and such that furthermore $\bar{c}'_n \models q \upharpoonright_{A, \bar{c}'_{<n}}$. By Fact 4.1 we have $\bar{c}_{\leq 2k} \equiv_A \bar{c}'_{\leq 2k}$. In addition, $d_p \varphi$ is over A , so $|d_p \varphi(\bar{y}) - \tilde{\varphi}^k(\bar{y}, \bar{c}_{\leq 2k})| \leq \varepsilon$.

Consider now $n > k$. First of all, by exactly the same argument as above, for every $w \in [2n+1]^{2k+1}$ we have $|d_p \varphi(\bar{y}) - \tilde{\varphi}^k(\bar{y}, \bar{c}_{\in w})| \leq \varepsilon$. In addition, for any \bar{b} there exists a subset $w \in [2n+1]^{2k+1}$ such that $\tilde{\varphi}^n(\bar{b}, \bar{c}_{\leq 2n}) = \tilde{\varphi}^k(\bar{b}, \bar{c}_{\in w})$ (from any set of $2n+1$ reals one can choose a subset of size $2k+1$ with the same median value). Thus $|d_p \varphi(\bar{y}) - \tilde{\varphi}^n(\bar{y}, \bar{c}_{\leq 2n})| \leq \varepsilon$ for all $n \geq k$, where k depends only on ε and φ , as desired. ■_{4.2}

Proposition 4.3. *Let $\varphi(\bar{x}, \bar{b})$ be an instance of a stable formula, A a set of parameters. Then the following are equivalent:*

- (i) *The condition $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A .*
- (ii) *If $(\bar{b}_n)_{n \in \mathbb{N}}$ is an A -indiscernible sequence, $\bar{b}_0 = \bar{b}$, then the set of conditions $\{\varphi(\bar{x}, \bar{b}_n) \leq r\}_{n \in \mathbb{N}}$ is consistent (i.e., the condition $\varphi(\bar{x}, \bar{b}) \leq r$ does not divide over A).*

Proof. (i) \implies (ii). If $(\bar{b}_n)_{n \in \mathbb{N}}$ is an A -indiscernible sequence and $\bar{b}_0 = \bar{b}$ then each \bar{b}_n is an $\text{acl}^{eq}(A)$ -conjugate of \bar{b} .

(ii) \implies (i). Fix models $\mathcal{N} \succeq \mathcal{M} \supseteq A$ where \mathcal{N} is saturated over M . Let $q_0 = \text{tp}(\bar{b}/\text{acl}^{eq}(A))$. By Lemma 2.2 there exists $q \in S_m(M)$ extending q_0 such that $q \upharpoonright_\varphi$ does not fork over A , i.e., such that $d_q \tilde{\varphi} = d_{q_0} \tilde{\varphi}$. Let $q_1 \in S_m(N)$ be an M -invariant extension of q . Finally, let $(\bar{b}_n)_{n \in \mathbb{N}} \models q_1^{(\omega)} \upharpoonright_M$. Then $(\bar{b}_n)_{n \in \mathbb{N}}$ is an \mathcal{M} -indiscernible sequence, and *a fortiori* A -indiscernible, in $\text{tp}(\bar{b}/A)$. Thus by assumption there exists \bar{a} such that $\varphi(\bar{a}, \bar{b}_n) \leq r$ for all n . In addition, by Lemma 4.2 we have

$$d_q \tilde{\varphi}(\bar{a}) = \lim_n \text{med}_n(\varphi(\bar{a}, \bar{b}_i))_{i \leq 2n} \leq r.$$

Let $p \in S_\varphi(M)$ be a non forking extension of $\text{tp}_\varphi(\bar{a}/\text{acl}^{eq}(A))$. Then $\varphi(\bar{x}, \bar{b})^p = d_q \tilde{\varphi}(\bar{x})^p \leq r$, witnessing that $\varphi(\bar{x}, \bar{b}) \leq r$ does not fork over A , as desired. ■_{4.3}

5. CANONICAL BASES

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

Proposition 5.1. *Assume T is stable, and let $p(\bar{x}) \in S_n(A)$ be stationary. Then:*

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

Proof. The first two items are immediate, while the third is by Corollary 2.8. Under the assumptions of the fourth we have $\text{Cb}(p) = \text{Cb}(p \upharpoonright_B) \subseteq \text{dcl}^{\text{eq}}(B)$. ■_{5.1}

The four properties listed in Proposition 5.1 determine the canonical base up to inter-definability. Indeed, if B has all four then $\text{Cb}(p) \subseteq \text{dcl}^{\text{eq}}(B)$ but also $B \subseteq \text{dcl}^{\text{eq}}(\text{Cb}(p))$, whereby $\text{dcl}^{\text{eq}}(B) = \text{dcl}^{\text{eq}}(\text{Cb}(p))$. In this case we say that B is a canonical base for p .

Proposition 5.2. *Assume T is stable, and let $p(\bar{x}) \in S_n(A)$ be stationary. Let $q \in S_n(M)$ be the unique non forking extension of p , where \mathcal{M} is saturated over A . Then a (small) set $B \subseteq M$ is a canonical base for p if and only if, for every $f \in \text{Aut}(\mathcal{M})$: $f \upharpoonright_B = \text{id}_B \iff f(q) = q$.*

Proof. Let $C = \text{Cb}(p) = \text{Cb}(q)$. It follows directly from the definitions that an automorphism of \mathcal{M} fixes q if and only if it fixes $q \upharpoonright_\varphi$ for every formula $\varphi(\bar{x}, \dots)$, if and only if it fixes every member of C . A small set B is another canonical base for p if and only if $\text{dcl}^{\text{eq}}(B) = \text{dcl}^{\text{eq}}(C)$ which is further equivalent to B and C being fixed by the same automorphisms. ■_{5.2}

We propose an alternative characterisation of canonical bases using Morley sequences. In the case of classical first order logic it is more or less folklore. Recall that a *Morley sequence* in a (stationary) type $p(\bar{x}) \in S_m(A)$ is a sequence $I = (\bar{a}_n)_{n \in \mathbb{N}}$ of realisations of p which is independent over A , i.e., such that $\bar{a}_n \downarrow_A \bar{a}_{<n}$ for all $n \in \mathbb{N}$. It follows by standard independence calculus that $\bar{a}_{\in s} \downarrow_A \bar{a}_{\in t}$ for every two disjoint index sets $s, t \subseteq \mathbb{N}$. From stationarity of p it follows that the sequence I is indiscernible over A , and its type over A , which we may denote by $p^{(\omega)}$, is uniquely determined by p .

It is not difficult to check that if p satisfies the assumptions of Fact 4.1 then the definition of $p^{(\omega)}$ which appears thereafter agrees with the one given here. In the general case, let \mathcal{M} be saturated over A and let $q \in S_m(M)$ be the non forking extension of p . Then by construction, $p^{(\omega)} = q_A^{(\omega)}$, where the first is the type of a Morley sequence as defined here, and the second the type defined after Fact 4.1.

Definition 5.3. Let $I = (\bar{a}_n)_{n \in \mathbb{N}}$ be a sequence of tuples, or, for that matter, even of sets. Let $I^{\geq k}$ denote the tail $(\bar{a}_n)_{n \geq k}$. We define the *tail definable closure* of I as

$$\text{tdcl}^{\text{eq}}(I) = \bigcap_{k \in \mathbb{N}} \text{dcl}^{\text{eq}}(I^{\geq k}).$$

It is not difficult to see that for an indiscernible sequence I , $\text{tdcl}^{\text{eq}}(I)$ consists precisely of all $c \in \text{dcl}^{\text{eq}}(I)$ over which I is indiscernible.

Lemma 5.4. *Let $I = (\bar{a}_n)_{n \in \mathbb{N}}$ and $J = (\bar{b}_n)_{n \in \mathbb{N}}$ be indiscernible sequences such that the concatenation $I \frown J$ is indiscernible as well. Then $\text{tdcl}^{\text{eq}}(I) = \text{tdcl}^{\text{eq}}(J)$. Moreover, every automorphism which sends I to J necessarily fixes $\text{tdcl}^{\text{eq}}(I)$.*

Proof. For $k \in \mathbb{N}$ let J_k be the sequence $\bar{a}_0, \dots, \bar{a}_{k-1}, \bar{b}_k, \bar{b}_{k+1}, \dots$, namely the sequence obtained by replacing the first k elements of J with the corresponding elements from I . Since $I \frown J$ is indiscernible so is J_k for each k , and there exists an automorphism f_k sending $J \mapsto J_k$. Now let $c \in \text{tdcl}^{\text{eq}}(J)$. Since c is definable over $J^{\geq k}$ it is fixed by f_k , so $cJ \equiv cJ_k$. This holds for all k , whence $cI \equiv cJ$.

Fix an automorphism f which sends I to J (which must necessarily exist). Then $f(c)J \equiv cI \equiv cJ$, so $f(c) = c$. Thus f fixes $\text{tdcl}^{\text{eq}}(J)$. Applying f^{-1} we obtain that $\text{tdcl}^{\text{eq}}(I) = \text{tdcl}^{\text{eq}}(J)$, as desired. ■_{5.4}

Theorem 5.5. *Let $p \in S_m(A)$ be a stationary type and let $I = (\bar{a}_n)_{n \in \mathbb{N}}$ be a Morley sequence in p . Then $\text{tdcl}^{eq}(I)$ is a canonical base of p .*

Proof. First of all, we have seen that $p \upharpoonright_{\text{Cb}(p)}$ is stationary, with the same canonical base as p . It is also not difficult to check that a Morley sequence in p is also a Morley sequence in $p \upharpoonright_{\text{Cb}(p)}$. It is therefore enough to prove for $p \upharpoonright_{\text{Cb}(p)}$, i.e., we may assume that $A = \text{Cb}(p)$.

So let \mathcal{M} be saturated over $A = \text{Cb}(p)$ and let $q \in S_m(M)$ be the non forking extension of p . As pointed above, $I \models p^{(\omega)} = q_A^{(\omega)}$. By Lemma 4.2 p is definable over I , so $\text{Cb}(p) \subseteq \text{dcl}^{eq}(I)$. Also, every tail of a Morley sequence is a Morley sequence, whence $\text{Cb}(p) \subseteq \text{tdcl}^{eq}(I)$.

Conversely, let f be an automorphism fixing $A = \text{Cb}(p)$. Then f fixes p and therefore sends I to another Morley sequence in p , say J . Let K be a third Morley sequence in p , $K \perp_A I, J$. Then both $I \cap K$ and $J \cap K$ can be verified to be Morley sequences in p (of length $\omega + \omega$), and in particular indiscernible. We can decompose $f = h \circ g$ where $g(I) = K$ and $h(K) = J$. By the Lemma $\text{tdcl}^{eq}(I) = \text{tdcl}^{eq}(K) = \text{tdcl}^{eq}(J)$ and this set is fixed by g, h and therefore by f . Thus $\text{tdcl}^{eq}(I) \subseteq \text{dcl}(\text{Cb}(p))$, and the proof is complete. $\blacksquare_{5.5}$

It is also a fact, which we shall not prove here (but is proved as in classical logic), that in a stable theory every indiscernible sequence $I = (\bar{a}_n)_{n \in \mathbb{N}}$ is a Morley sequence in some type, say q . Let $A = \text{tdcl}^{eq}(I)$ and $p = \text{tp}(\bar{a}_n/A)$, which does not depend on n . By the Theorem, $A = \text{Cb}(q) = \text{Cb}(p)$ and I is a Morley sequence in p .

In the case of probability theory this is a well known fact. Indeed, in probability algebras or in spaces of random variables (say $[0, 1]$ -valued, see [Benb]), the canonical base of a type (in the real sort) can be represented by a set of real elements, so there is no need to consider imaginaries. Then Theorem 5.5 tells us that if $(X_n)_{n \in \mathbb{N}}$ is sequence of random variables which is indiscernible (i.e., exchangeable) and \mathcal{A} is its tail algebra then the sequence $(X_n)_{n \in \mathbb{N}}$ is i.i.d. over \mathcal{A} , meaning that the random variables X_n are independent over \mathcal{A} and have the same conditional distribution over \mathcal{A} .

Corollary 5.6. *Assume T is stable, and let $p(\bar{x}) \in S_m(A)$ be stationary. Let $I = (\bar{a}_n)_{n \in \mathbb{N}}$ be a Morley sequence in p , $J = I \setminus \bar{a}_0$. Then $\bar{a}_0 \perp_A J$ and $\bar{a}_0 \perp_J A$.*

Proof. The first independence is immediate and implies $\bar{a}_0 \perp_{\text{Cb}(p)} AJ$. By Theorem 5.5 we have $\text{Cb}(p) \subseteq \text{dcl}^{eq}(J)$ and the second independence follows. $\blacksquare_{5.6}$

6. STABLE TYPE-DEFINABLE GROUPS AND THEIR ACTIONS

We turn to consider groups, and more generally, homogeneous spaces, which are definable or type-definable in a stable theory.

6.1. Generic elements and types in stable group 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 shall identify G and S with their sets of realisations in a monster model \mathfrak{M} . We are 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 $(g, h) \mapsto hg^{-1}$.

Given a partial type $\pi(x)$ in the sort of S we let $\pi(S)$ denote the subset of S defined by π .

- Definition 6.1.**
- (i) A *generic set* in S is a subset $X \subseteq S$ finitely many G -translates of which cover S .
 - (ii) A *generic partial type* in S is a partial type $\pi(x)$ such that every logical neighbourhood of π (as per Definition 1.5) defines in S a generic set. Single conditions as well as complete types are generic if they are generic as partial types.

- (iii) We say that an element $s \in S$ is *generic* over a set A if $\text{tp}(s/A)$ is generic.
- (iv) A *left-generic set* in G is a subset $X \subseteq G$ which is generic under the action of G on itself on the left. We define partial types in the sort of G to be left-generic accordingly. Similarly for right-generic.

Let $\pi(x)$ be a partial type. Clearly, if $\pi(S)$ is a generic set then π is a generic partial type, but the converse is not always true. In classical logic, if π consists of a single formula (i.e., if $\pi(S)$ is a relatively definable subset of S , and so is its complement), then π is its own logical neighbourhood and the two notions coincide. Unfortunately, this will generally never happen in continuous logic (except for $\pi(S) = S$ or $\pi(S) = \emptyset$).

Lemma 6.2. *The following are equivalent for a partial type $\pi(x)$ in the sort of S , with parameters in a set A :*

- (i) *The partial type π is generic in S .*
- (ii) *For every formula $\varphi(x, \bar{a})$ over A , if the condition $\varphi(x, \bar{a}) = 0$ is a logical neighbourhood of π then it is a generic condition.*

Proof. One direction is immediate, the other follow from Lemma 1.6. ■_{6.2}

Let $S_S(A)$ denote the set of all complete types over A implying $x \in S$. Equipped with the induced topology from $S_x(A)$, it is a compact space, 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 defined by a generic partial type in S if and only if every relative logical neighbourhood of X in S defines a generic set.

For $g \in G$ and $X \subseteq S$, let $L_g[X] = gX = \{gs\}_{s \in X}$. Somewhat superfluously, we may also define $L_g^{-1}[X] = \{s \in S : gs \in X\} = L_{g^{-1}}[X]$.

Lemma 6.3. *Let A be a set of parameters, $g \in G(A) = G \cap \text{dcl}(A)$.*

- (i) *If $X \subseteq S$ is type-definable over A , say by a partial type π , then $L_g[X]$ is also type-definable over A by a partial type which will be denoted $L_g\pi$ (or $g\pi$). Moreover, π is generic if and only if $g\pi$ is.*
- (ii) *If $p = \text{tp}(s/A) \in S_S(A)$ is a complete type then $L_gp = gp = \text{tp}(gp/A)$, and $L_g : S_S(A) \rightarrow S_S(A)$ is a homeomorphism, and restricts to a homeomorphism of the set of generic types with itself.*

Proof. We only prove the parts regarding genericity. Indeed, assume that π is generic, and let $gX <^S Y$. Then $X <^S L_g^{-1}[Y]$, so $L_g^{-1}[Y]$ is a generic subset of S . It follows immediately that so is Y . Thus $g\pi$ is a generic partial type. For the converse replace g with g^{-1} . ■_{6.3}

Similarly, for $s \in S$ and $X \subseteq G$ we define $R_s[X] = Xs = \{gs\}_{g \in X}$. For $X \subseteq S$ we define $R_s^{-1}[X] = \{g \in G : gs \in X\}$.

Lemma 6.4. *Let A be a set of parameters, $s \in S(A) = S \cap \text{dcl}(A)$.*

- (i) *If $X \subseteq G$ is type-definable over A , say by a partial type π , then $R_s[X]$ is also type-definable over A by a partial type which will be denoted $R_s\pi$ (or πs). Moreover, if π is left-generic then $R_s\pi$ is generic.*
- (ii) *If $p = \text{tp}(g/A) \in S_G(A)$ is a complete type then $R_sp = ps = \text{tp}(gs/A)$, and $R_s : S_G(A) \rightarrow S_S(A)$ is a continuous surjection, sending left-generic types to generic types.*

Notice that we do not claim that every generic type in $S_S(A)$ is the image under R_s of a left-generic type in $S_G(A)$ (this is true if T is stable).

Proof. Essentially identical to that of Lemma 6.3. ■_{6.4}

Under the assumption that the theory $T = \text{Th}(\mathfrak{M})$ is stable we shall show that generic types exist and study some of their properties. We follow a path similar to that followed in [Pil96]. Toward this end we construct an auxiliary multi-sorted structure $\hat{\mathfrak{M}} = \langle G, S, \dots \rangle$ in a language $\hat{\mathcal{L}}$ (in addition to sorts G and S , $\hat{\mathcal{L}}$ consists of additional sorts which we shall describe later). We define the distance on the first two sorts by

$$d_G^{\hat{\mathfrak{M}}}(g, g') = \sup_{h \in G} d^{\mathfrak{M}}(hg, hg'), \quad d_S^{\hat{\mathfrak{M}}}(s, s') = \sup_{h \in G} d^{\mathfrak{M}}(hs, hs').$$

This coincides with the original distance in \mathfrak{M} if the latter is invariant under the action of G (on the left). In any case, $d_G^{\hat{\mathfrak{M}}}$ is a distance function, invariant under the action of G , and satisfies $d_G^{\hat{\mathfrak{M}}} \geq d^{\mathfrak{M}}$. On the other hand, if $g_n \rightarrow g$ in $d^{\mathfrak{M}}$ then $g_n \rightarrow g$ in $d_G^{\hat{\mathfrak{M}}}$ as well (if not, then by a compactness argument, for some $\varepsilon > 0$ there would exist $h \in G$ such that $d^{\mathfrak{M}}(hg, hg) \geq \varepsilon$, an absurd). It follows that $(G, d_G^{\hat{\mathfrak{M}}})$ is a complete metric space. The same observations hold for $(S, d_S^{\hat{\mathfrak{M}}})$.

Let now Φ_S consist of all \mathcal{L} -formulae of the form $\varphi(x, \bar{y})$ where x is in the sort of S . For each $\varphi \in \Phi_S$, there will be a sort C_φ , consisting of all canonical parameters of instances of φ in \mathfrak{M} . The canonical parameter of $\varphi(x, \bar{b})$ will be denoted $[\bar{b}]_\varphi$, or $[\bar{b}]$ if there is no ambiguity. We put on it the standard metric, namely

$$d_\varphi([\bar{b}]_\varphi, [\bar{b}']_\varphi) = \sup_{a \in \mathfrak{M}} |\varphi(a, \bar{b}) - \varphi(a, \bar{b}')|.$$

The only symbols in the language $\hat{\mathcal{L}}$, in addition to the distance symbols of the various sorts, are a predicate symbol $\hat{\varphi}(x_S, y_G, z_\varphi)$ for each formula $\varphi \in \Phi_S$, interpreted by

$$\hat{\varphi}(s, g, [\bar{b}])^{\hat{\mathfrak{M}}} = \varphi(g^{-1}s, \bar{b})^{\mathfrak{M}}.$$

Since φ is uniformly continuous in all its variables, so is $\hat{\varphi}$. These definitions make $\hat{\mathfrak{M}}$ a continuous $\hat{\mathcal{L}}$ -structure.

If $\langle G, S \rangle$ is definable then $\hat{\mathfrak{M}}$ is interpretable in \mathfrak{M} and $\hat{T} = \text{Th}_{\hat{\mathcal{L}}}(\hat{\mathfrak{M}})$ is stable (assuming T is). In the general case, all we know is that $\hat{\mathfrak{M}}$ 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 $\theta_h: \hat{\mathfrak{M}} \rightarrow \hat{\mathfrak{M}}$ by sending $g \in G$ to hg , $s \in S$ to hs , and fixing all the auxiliary sorts. This is easily verified to be an automorphism of $\hat{\mathfrak{M}}$. Since the action of G on S is assumed to be transitive, if $A \subseteq \bigcup_\varphi C_\varphi$ then all elements of S have the same type over A in $\hat{\mathfrak{M}}$, and similarly all elements of G .

Lemma 6.5. *Assume that $\varphi(x, \bar{y}) \in \Phi_S$ is stable. Then the following are equivalent for an instance $\varphi(x, \bar{b})$:*

- (i) *The condition $\varphi(x, \bar{b}) = 0$ is generic in S .*
- (ii) *The condition $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork in $\hat{\mathfrak{M}}$ over \emptyset .*
- (iii) *The condition $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork in $\hat{\mathfrak{M}}$ over $[\bar{b}]$.*

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 $\hat{\mathfrak{M}}$. For $\varepsilon > 0$ let $X_\varepsilon = \{s \in S: \varphi(s, \bar{b}) \leq \varepsilon\}$. By Lemma 6.2, the condition $\varphi(x, \bar{b}) = 0$ is generic if and only if X_ε is a generic set for all $\varepsilon > 0$.

(i) \implies (ii). Assume first that $\varphi(x, \bar{b}) = 0$ is generic in S , i.e., that the set X_ε is generic for every $\varepsilon > 0$. Find $g_i \in G$ such that $S = \bigcup_{i < n} g_i X_\varepsilon$, and find $s \in S$ such that $\text{tp}_{\hat{\varphi}}(s/[\bar{b}]_{g_{<n}})$ does not fork over \emptyset (in symbols $s \not\perp^{\hat{\varphi}} [\bar{b}]_{g_{<n}}$). Since $s \in \bigcup_{i < n} g_i X_\varepsilon$ 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 .

(ii) \implies (iii). Immediate.

(iii) \implies (i). Assume now that $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork over $[\bar{b}]$. By Proposition 2.17 there are $g_n \in G$ for $n \in \mathbb{N}$ and a faithful combination $\psi(x, [\bar{b}]) = \theta(\hat{\varphi}(x, g_n, [\bar{b}]))_{n \in \mathbb{N}}$ which is definable over $[\bar{b}]$ and such that $\psi(x, [\bar{b}]) = 0$ is consistent. Since $\hat{\mathfrak{M}}$ 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 $\hat{\mathfrak{M}}$, 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 $\varphi(g_n^{-1}s, \bar{b}) \geq \varepsilon$ for all n , i.e., $\hat{\varphi}(s, g_n, [\bar{b}]) \geq \varepsilon$. Since the combination above was faithful we get $\psi(s, [\bar{b}]) \geq \varepsilon > 0$, a contradiction. $\blacksquare_{6.5}$

Lemma 6.6. *Assume that $\varphi(x, \bar{y}) \in \Phi_S$ is stable and that $\varphi(x, \bar{b}) = 0$ is a generic condition in S . Then it does not fork over \emptyset .*

Proof. By Proposition 4.3 it will be enough to show that $\varphi(x, \bar{b}) = 0$ does not divide over \emptyset . For this purpose let $(\bar{b}_n)_{n \in \mathbb{N}}$ be any indiscernible sequence with $\bar{b}_0 = \bar{b}$. Since $e \in \text{dcl}(\emptyset)$, the sequence $(e, \bar{b}_n)_{n \in \mathbb{N}}$ is indiscernible as well, and thus the sequence $(e, [\bar{b}_n])_{n \in \mathbb{N}}$ is indiscernible in $\hat{\mathfrak{M}}$. On the other hand, since the condition $\varphi(x, \bar{b}) = 0$ is generic, by Lemma 6.5 the condition $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork over \emptyset , so $\{\hat{\varphi}(x, e, [\bar{b}_n])\}_{n \in \mathbb{N}}$ is consistent. Since $\hat{\mathfrak{M}}$ is saturated for such formulae, there is $s \in S$ such that $\hat{\varphi}(s, e, [\bar{b}_n]) = 0$, i.e., $\varphi(s, \bar{b}_i) = 0$, for all n , as desired. $\blacksquare_{6.6}$

From now on we assume that T is stable.

Proposition 6.7. *Let $\pi(x)$ be a partial type over A . Then π is generic if and only if it extends to a complete generic type over A , i.e., if and only if $[\pi] \subseteq S_S(A)$ 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 π is a generic partial type. Since the set of complete generic types is closed it will be enough to show that every logical neighbourhood of π contains a generic type, and we may further restrict our attention to logical neighbourhoods defined by a single condition $\varphi(x, \bar{b}) = 0$. Since π is generic in S so is $\varphi(x, \bar{b}) = 0$. By Lemma 6.5 $\hat{\varphi}(x, e, [\bar{b}]) = 0$ does not fork over \emptyset in $\hat{\mathfrak{M}}$. By Corollary 2.4 there exists a type $\hat{p} \in S_x(\hat{\mathfrak{M}})$ such that $\hat{\varphi}(x, e, [\bar{b}])^{\hat{p}} = 0$ and in addition $p \upharpoonright_{\hat{\varphi}}$ does not fork over \emptyset for every formula $\psi \in \Phi_S$. Let

$$p(x) = \{\psi(x, \bar{c}) = \hat{\psi}(x, e, [\bar{c}])^{\hat{p}}\}_{\psi \in \Phi_S, \bar{c} \in \mathfrak{M}}.$$

This type is approximately finitely realised in \mathfrak{M} (since \hat{p} is in $\hat{\mathfrak{M}}$) and therefore consistent. By Lemma 6.5 every condition in p is generic (since \hat{p} does not fork over \emptyset), and by Lemma 6.2, p is generic, and so is $p \upharpoonright_A$. We have thus found a generic type $p \upharpoonright_A \in [\varphi(x, \bar{b}) = 0]$ and the proof is complete. $\blacksquare_{6.7}$

Proposition 6.8. *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 $p \upharpoonright_A$ is generic. In particular, a generic type does not fork over \emptyset .*

Proof. First of all, the last assertion follows from Lemma 6.6 and the fact that the set of non forking types is closed.

We now prove the main assertion. For left to right, if $p \in S_S(B)$ is generic then clearly so is $p \upharpoonright_A$, and by the previous paragraph p does not fork over A . For the converse, assume that $p \in S_S(B)$ does not fork over A and $p_0 = p \upharpoonright_A$ is generic. Replacing p with a non forking extension we may assume that $B = \mathfrak{M}$. By Proposition 6.7 there is $p_1 \in S_S(\mathfrak{M})$ extending p_0 which is generic, and by what we have just shown it is also non forking over A . Since $p \upharpoonright_A = p_0 = p_1 \upharpoonright_A$ there is $f \in \text{Aut}(\mathfrak{M}/A)$ sending $p_1 \upharpoonright_{\text{acl}^{\text{eq}}(A)}$ to $p \upharpoonright_{\text{acl}^{\text{eq}}(A)}$, and therefore p_1 to p . Thus p is generic as well. $\blacksquare_{6.8}$

We can also complement Lemma 6.3:

Proposition 6.9. *The action of G on the set of generic types in $S_S(\mathfrak{M})$ is transitive.*

Proof. Let $p, q \in S_S(\mathfrak{M})$ be two generic types. Define

$$\hat{p} = \{\hat{\varphi}(x, g, [\bar{b}]) = \varphi(g^{-1}x, [\bar{b}])^p\}_{\varphi \in \Phi_S, \bar{b} \in \mathfrak{M}, g \in G},$$

and define \hat{q} similarly. Let $\hat{C} = (\text{acl}^{\text{eq}}(\emptyset))^{\mathfrak{M}}$ and let $\hat{p}_0 = \hat{p}|_{\hat{C}}$, $\hat{q}_0 = \hat{q}|_{\hat{C}}$. Since $\hat{\mathfrak{M}}$ is saturated for formulae of this form we may realise \hat{p}_0 and \hat{q}_0 in $\hat{\mathfrak{M}}$, and by transitivity there exists $h \in G$ such that $\theta_h \hat{p}_0 \cup \hat{q}_0$ is realised. Since θ_h is an automorphism of $\hat{\mathfrak{M}}$ we must have $\hat{q}_0 = \theta_h \hat{p}_0 = (\theta_h \hat{p})|_{\hat{C}}$. In addition, neither of \hat{p} , \hat{q} or $\theta_h \hat{p}$ forks over \emptyset , whereby $\theta_h \hat{p} = \hat{q}$, i.e., $hp = q$. $\blacksquare_{6.9}$

Theorem 6.10. *Let G be a type-definable group in a stable theory T , acting type-definably and transitively on a type-definable set S .*

- (i) *If $g \downarrow_A s$ (where $g \in G$, $s \in S$) and g is left-generic over A then gs is generic over A and $gs \downarrow_A s$.*
- (ii) *An element $s \in S$ is generic if and only if $g \downarrow_A s$ implies $gs \downarrow_A g$ for every $g \in G$. Moreover, in this case gs is generic over A as well.*
- (iii) *An element $g \in G$ is left-generic over A if and only if g^{-1} is.*
- (iv) *An element $g \in G$ is left-generic if and only if it is right-generic (over A). From now on we shall only speak of generic elements and types in G .*
- (v) *An element $g \in G$ is generic over A if and only if it is generic over \emptyset and $g \downarrow_A$.*

Proof. We use Proposition 6.8 repeatedly.

For the first item, let $s \in S$, $g \in G$, and assume that $g \downarrow_A s$. If g is left-generic over A then it is left-generic over A, s . By Lemma 6.4 gs is generic over A, s . It follows that gs is generic over A and that $gs \downarrow_A s$, as desired.

For the second item, left to right, as well as the moreover part, are proved as in the previous argument, using Lemma 6.3. For right to left, assume that $s \downarrow_A g$ implies $gs \downarrow_A g$ for all g . We may choose g which is left-generic over A such that $g \downarrow_A s$. Then $g^{-1} \downarrow_A gs$ by assumption, gs is generic over A by the first item, and $s = g^{-1}gs$ is generic over A by the moreover part.

For the third item, let $g \in G$ be left-generic over A . Choose $h \in G$ left-generic over A such that $g \downarrow_A h$. By the first item gh is generic over A and $gh \downarrow_A h$. This can be re-written as $h \downarrow_A h^{-1}g^{-1}$. By the first item again, $g^{-1} = hh^{-1}g^{-1}$ is left-generic over A . Notice that g^{-1} is left-generic if and only if g is right-generic, yielding the fourth item as well.

The last item is just Proposition 6.8. $\blacksquare_{6.10}$

6.2. Stabilisers. We have already observed in Lemma 6.3 that for any set of parameters A , a group element $g \in G(A)$ induces a homeomorphism $L_g: p \mapsto gp$ on $S_S(A)$. It is also not difficult to check that $L_g \circ L_h = L_{gh}$, whence a group action of $G(A)$ on $S_S(A)$. In addition, we have seen that it restricts to an action by homeomorphism of $G(A)$ on the set of generic types in $S_S(A)$.

Specifically, we obtain an action of $G = G(\mathfrak{M})$ on $S_S(\mathfrak{M})$. The *stabiliser* of a type $p \in S_S(\mathfrak{M})$ under this action is $\text{Stab}(p) = \{g \in G: gp = p\} \leq G$. For a stationary type $p \in S_S(A)$ we define $\text{Stab}(p) = \text{Stab}(p|_{\mathfrak{M}})$.

Proposition 6.11. *Let $p \in S_S(A)$ be stationary. Then stabiliser $\text{Stab}(p)$ is a sub-group of G type-definable over $\text{Cb}(p)$.*

Moreover, assume that $s \models p$, $g \in G$ and $g \downarrow_A s$. Then $g \in \text{Stab}(p)$ if and only if $gs \models p$.

Proof. We may assume that $p \in S_S(\mathfrak{M})$.

Let $\varphi(x, \bar{z})$ be a formula, x in the sort of S . Let y be a variable in the sort of G . Then $\varphi(yx, \bar{z})$ is a definable predicate on $G \times S \times \langle \text{sort of } \bar{z} \rangle$, i.e., a continuous function $S_{G,S,\bar{z}}(T) \rightarrow [0, 1]$. By Tietze's Extension Theorem this extends to a continuous function $S_{x,y,\bar{z}}(T) \rightarrow [0, 1]$. For clarity we shall use $\varphi(yx, \bar{z})$ to denote the corresponding definable predicate.

Once this technical preliminary is taken care of we see that $\text{Stab}(p)$ is defined by the following axiom scheme:

$$\pi(y) = \left\{ \sup_{\bar{z}} |d_p \varphi(x, \bar{z}) - d_p \varphi(yx, \bar{z})| = 0 \right\}_{\varphi \in \Phi_S}.$$

The moreover part easily follows. ■_{6.11}

Lemma 6.12. *Let $H < G$ be a type-definable subgroup of bounded index, say with parameters in A , and let $g \in H$. Then g is generic over A in G if and only if it is generic over A in H .*

Proof. Naming A in the language we may assume that $A = \emptyset$. Since H has bounded index we may enumerate its cosets $\{g_i H\}_{i < \lambda}$. Let $h_0 \in G$ be generic over $\{g_i\}_{i < \lambda}$. Then $h_0 \in g_i H$ for some i , and $h_1 = g_i^{-1} h_0 \in H$ is generic in G . Now let h_2 be generic in H . Without loss of generality we may assume that $h_1 \perp h_2$. Then $h_1 h_2 \in H$ is generic both in H and in G and $h_1 h_2 \perp h_1$. Thus $h_2 = h_1^{-1} h_1 h_2$ is generic in G as well. We have thus shown that every generic of H is a generic of G . A similar argument shows that every generic of G in H is generic in H . ■_{6.12}

Proposition 6.13. *A type $p \in S_S(A)$ is generic if and only if $\text{Stab}(p)$ has bounded index in G .*

Proof. There are only boundedly many generic types over \mathfrak{M} , since they do not fork over \emptyset and therefore determined by their restriction to $\text{acl}^{\text{eq}}(\emptyset)$. In addition, the action of G on $S_S(\mathfrak{M})$ restrict to an action of G on the space of generic types, so the stabiliser of a generic type must be of bounded index.

Conversely, assume $\text{Stab}(p)$ has bounded index, and let $s \models p$. Then there exists $g \in \text{Stab}(p)$ which is generic in G over A , and we may further assume that $g \perp_A s$. Then $gs \models p$ is generic over A , i.e., p is generic. ■_{6.13}

Since G acts transitively on the generic types over \mathfrak{M} , the stabilisers of generic types are all conjugate. It is also not difficult to check that if $p \in S_S(\mathfrak{M})$ is generic, $q \in S_G(\mathfrak{M})$ is a generic type of $\text{Stab}(p)$ (and therefore of G), and $s \models p|_{\text{acl}^{\text{eq}}(\emptyset)}$, then $qs = p$. If q' is any other generic of G then (since G acts transitively on its own generic types, on the left as well as on the right) there exists $g \in G$ such that $q = q'g$ and $p = q'(gs)$. Thus the right action of S on G send each and every generic type of G onto the generic types of S , complementing Lemma 6.4.

Theorem 6.14. *Let G be a type-definable group in a stable theory, say over \emptyset . Then G admits a smallest type-definable group of bounded index (over any set of parameters), called the connected component of G , and denoted G^0 . It has the following additional properties.*

- (i) *The connected component G^0 is a normal subgroup of G , type-definable over \emptyset .*
- (ii) *The stabiliser of every generic type is equal to G^0 .*
- (iii) *Each coset gG^0 contains a unique generic type over \mathfrak{M} .*
- (iv) *The generic type of G^0 is definable over \emptyset .*
- (v) *If $p \in S_G(A)$ is any stationary generic type over a small set then $G^0 = \{g^{-1}h : g, h \models p\}$.*

Proof. We start by constructing G^0 and proving the second item. Since left generic and right generic are the same, the action of G on the generic types is transitive on either side. In particular, if $p, q \in S_G(\mathfrak{M})$ are generic then there exists $g \in G$ such that $q = pg$, and thus $\text{Stab}(p) = \text{Stab}(q)$. Let this unique stabiliser of generic types be denoted G^0 . Then G^0 is type-definable, and since $G \emptyset$ -invariant, so is G^0 , and we may conclude that G^0 is type-definable over \emptyset as well. We also already know that G^0 has

bounded index in G . Assume now that $H \leq G^0$ is another type-definable subgroup of bounded index, say over \emptyset (otherwise name the parameters in the language). Then there exists $p \in S_H(\mathfrak{M})$ generic in G , so $\text{Stab}(p) = G^0$, whereby $G^0 \subseteq H$. Thus G^0 is indeed the smallest type-definable subgroup of G of bounded index. Notice that $G^0 \cap gG^0g^{-1}$ is also type-definable of bounded index for every $g \in G$, so G^0 is normal in G . This concludes the proof of the first two items.

Let $p \in S_G(\mathfrak{M})$ be generic in G^0 . Since $G^0 = \text{Stab}(p)$ acts transitively on its generic types, p is the unique generic type in G^0 . It follows that a coset gG^0 contains a unique generic type gp . The uniqueness of the generic type of G^0 implies that it is \emptyset -invariant, and therefore definable over \emptyset .

Finally, let $p \in S_G(A)$ be any stationary generic type over a small set. Then $p|_{\mathfrak{M}}$ is the unique generic type in some coset gG^0 . It follows that gG^0 is A -invariant, so $p \vdash x \in gG^0$. Thus $\{g^{-1}h : g, h \models p\} \subseteq G^0$. Conversely, let $g \in G^0$, and let $h \models p$, $g \downarrow_A h$. Since G^0 must also be the right-stabiliser of p we have $hg \models p$ as well, and $g = h^{-1}(hg)$, as desired. ■_{6.14}

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

6.3. Global group ranks. We have seen that a type of a member of S is generic if and only if the corresponding type in \mathfrak{M} is a non forking 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 $\varphi \in \Phi_S$. 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 shall consider here the case of superstability and Lascar ranks. Morley ranks are studied in a subsequent paper [Bena], and similar results are proved.

Significant work regarding superstability has been carried out in the context of metric Hausdorff cats, and in many cases definitions and proofs transfer verbatim to continuous logic. The objects of study in this context are partial types constructed from complete types and conditions of the form $d(x, y) \leq \varepsilon$ via conjunction and existential quantification. It is a general fact that if $\pi(\bar{x}, \bar{y})$ is a partial type then the property $\exists \bar{y} \pi(\bar{x}, \bar{y})$, where the existential quantifier is interpreted in a sufficiently saturated model, is definable by a partial type as well. This mostly happens in the following form. Let $\pi(\bar{x})$ be a partial type and $\varepsilon \geq 0$ a real number. We define $\pi(\bar{x}^\varepsilon)$ to be the partial type expressing that $\exists \bar{y} (\pi(\bar{y}) \ \& \ d(\bar{x}, \bar{y}) \leq \varepsilon)$, i.e., that π is realised in the ε -neighbourhood of \bar{x} . For a tuple \bar{a} we shall use \bar{a}^ε as a notational representation for the somewhat vague concept of “ \bar{a} known up to distance ε ”, so in particular \bar{a}^0 is just another representation for \bar{a} . Accordingly, if $p(\bar{x}) = \text{tp}(\bar{a}/C)$ then we define $\text{tp}(\bar{a}^\varepsilon/C) = p(\bar{x}^\varepsilon)$, so in particular $\text{tp}(\bar{a}/C) = \text{tp}(\bar{a}^0/C) = \bigwedge_{\varepsilon > 0} \text{tp}(\bar{a}^\varepsilon/C)$.

Definition 6.15. To an arbitrary theory T we define $\kappa(T)$ to be the least infinite cardinal, if such exists, such that for every complete type $p(\bar{x})$ over a set A , and for every $\varepsilon > 0$, there is a subset $A_0 \subseteq A$ such that $p(\bar{x}^\varepsilon)$ does not divide over A_0 .

We say that T is *simple* if $\kappa(T) \leq |\mathcal{L}|^+$, and that it is *supersimple* if $\kappa(T) = \aleph_0$.

It follows from our earlier results that every stable theory is simple. It is true (but we shall not require it) that if T is not simple then $\kappa(T) = \infty$.

Definition 6.16. We say that T is λ -*stable* if for every set A , $|A| \leq \lambda$, the metric density character of $S_n(A)$ is at most λ . We define $\lambda_0(T)$ to be the least (infinite) cardinal of stability for T . We say that T is *superstable* if it is λ -stable for all λ big enough.

In the context of Henson’s logic for Banach space, this definition dates back to Iovino [Iov99]. It was shown in [BU] that T is stable if and only if it is $\lambda^{|\mathcal{L}|}$ -stable for all λ . In particular, T is stable if and only if $\lambda_0(T) < \infty$, in which case $\lambda_0(T) \leq 2^{|\mathcal{L}|}$. The following is an example for a result whose statement and proof translate word-for-word to the continuous logic setting.

Fact 6.17 ([Ben05, Theorem 4.13]). *A theory T is λ -stable if and only if $\lambda^{<\kappa(T)} = \lambda \geq \lambda_0(T)$.*

Corollary 6.18. *A theory is superstable if and only if it is stable and supersimple.*

Many of the arguments that follow are valid both for stable and for simple theories, and are stated as such, even though no development of simplicity theory for continuous logic exists in the literature. The reader may either refer to the development of simplicity in the context of cats [Ben03a, Ben03b], which encompasses that of continuous logic, or simply restrict his or her attention to the stable case.

In light of Corollary 6.18 and of the definition of $\kappa(T)$ we wish to define global forking ranks SU_ε for $\varepsilon > 0$ such that $SU_\varepsilon(\bar{a}/C) > SU_\varepsilon(\bar{a}/C\bar{b})$ if and only if \bar{a} “ ε -depends” on \bar{b} over C . We have quite a bit of liberty in choosing what “ ε -depends” should mean (i.e., different choices can give rise to ranks which have all the properties we seek). For example, we could say that this happens if $\text{tp}(\bar{a}^\varepsilon/C\bar{b})$ forks over C , or if $\text{tp}(\bar{a}^\varepsilon/C\bar{b}) \wedge \text{tp}(\bar{a}/\text{acl}^{\text{eq}}(C))$ forks over C . For consistency with earlier work we shall opt for a slightly more complex definition which was given in [Ben06] and which has some advantages over other definitions for the purposes of that paper.

Definition 6.19. Let \bar{a} and \bar{b} be tuples, C a set and $\varepsilon > 0$. We keep in mind that \bar{a}^ε represents “ \bar{a} known up to distance ε ”.

- (i) We say that an indiscernible sequence $(\bar{b}_n)_{n \in \mathbb{N}}$ could be in $\text{tp}(\bar{b}/\bar{a}^\varepsilon C)$ if there are C' and a sequence $(\bar{a}_n)_n$ such that $(\bar{a}_n \bar{b}_n)_n$ is C' -indiscernible, $\bar{a}_n \bar{b}_n C' \equiv \bar{a} \bar{b} C$ and $d(\bar{a}_0, \bar{a}_1) \leq \varepsilon$.
- (ii) We say that $\bar{a}^\varepsilon \downarrow_C \bar{b}$ if every indiscernible sequence in $\text{tp}(\bar{b}/C)$ could be in $\text{tp}(\bar{b}/\bar{a}^\varepsilon C)$.
- (iii) We define $SU_\varepsilon(\bar{a}/C)$ as may be expected: $SU_\varepsilon(\bar{a}/C) \geq \alpha + 1$ if and only if there is \bar{b} such that $\bar{a}^\varepsilon \not\downarrow_C \bar{b}$ and $SU_\varepsilon(\bar{a}/C\bar{b}) \geq \alpha$. (In [Ben06] the notation $SU(\bar{a}^\varepsilon/C)$ was used.)

A few justifications for these definitions may be in place. Indeed, for $\varepsilon = 0$ it is easy to see that a sequence (\bar{b}_n) could be in $\text{tp}(\bar{b}/\bar{a}^0 C)$ if and only if it admits a conjugate which is in $\text{tp}(\bar{b}/\bar{a} C)$ in the ordinary sense. By a compactness argument, this is further equivalent to the property that (\bar{b}_n) could be in $\text{tp}(\bar{b}/\bar{a}^\varepsilon C)$ for every $\varepsilon > 0$.

Next, we wish to justify the notation \downarrow . The result in [Ben06] on which this justification relies does not immediately make sense in continuous logic, since at some point it uses a local ranks argument which is specific to the formalism of compact abstract theories. The core of that argument resides in the following technical result, which indeed can be proved very quickly in many different contexts (simple or stable theories, classical logic, continuous logic, cats) using the appropriate local ranks. Since the notion of local forking ranks varies drastically between contexts we shall provide here a more combinatorial and therefore more universal argument.

Lemma 6.20. *Let a, b and c be three tuples, possibly infinite, in a stable or even simple theory (classical, continuous, or any other setting in which basic simplicity or stability hold). Assume moreover that $b \equiv_a c$. Then $a \downarrow_b c$ if and only if $a \downarrow_c b$.*

Proof. Assume not, say $a \downarrow_c b$ but $a \not\downarrow_b c$. We construct by induction a sequence $(b_n)_{n \in \mathbb{N}}$ such that $b_n b_{n+1} \equiv_a bc$ and $ab_{n+1} \downarrow_{b_n} b_{<n}$ for all n . We start with $b_0 = b, b_1 = c$. Then for each n we can choose b_{n+2} such that $b_{n+1} b_{n+2} \equiv_a bc$ and $b_{n+2} \downarrow_{ab_{n+1}} b_{\leq n}$. Our induction hypothesis tells us that $ab_{n+1} \downarrow_{b_n} b_{<n}$ and $a \downarrow_{b_{n+1}} b_n$. By standard independence calculus we obtain $a \downarrow_{b_{n+1}} b_{\leq n}$, whence $ab_{n+2} \downarrow_{b_{n+1}} b_{\leq n}$, as desired.

Since $a \not\downarrow_b c$, there exists a b -indiscernible sequence $(c^k)^k$, starting with $c^0 = c$, such that there exists no a' satisfying $a'c^k \equiv_b ac$ for all k . For each n we may choose a copy $(c_n^k)_k$ such that $(c_n^k)_k, b_{n+1}, b_n \equiv (c^k)_k, c, b$, so in particular this copy starts with $c_n^0 = b_{n+1}$ and is indiscernible over b_n . Since $b_{n+1} \downarrow_{b_n} b_{<n}$, we may choose $(c_n^k)_k$ to be indiscernible over $b_{\leq n}$.

Let us now define $c_n = b_{n+1}$ for all n and consider the sequence $(b_n c_n)_n$. Applying compactness, we can find for arbitrarily big λ a sequence $(b_i c_i)_{i < \lambda}$ in $\text{tp}(bc/a)$, such that for each i there exists a sequence

$(c_i^k)_k$ which is indiscernible over $b_{\leq i}c_{< i}$ such that $(c_i^k)_k, c_i, b_i \equiv (c^k)_k, c, b$, so in particular $c_i^0 = c_i$. Each of these sequences witnesses that $a \not\downarrow_{b_{< i}c_{< i}} b_i c_i$, contradicting the local character. $\blacksquare_{6.20}$

We can now prove:

Fact 6.21 ([Ben06, Lemma 1.13]). *Assume that T is simple (or stable). For \bar{a}^ε , C and \bar{b} , the following imply one another from top to bottom:*

- (i) $\bar{a}^\varepsilon \downarrow_C \bar{b}$.
- (ii) *There is a Morley sequence for \bar{b} over C which could be in $\text{tp}(\bar{b}/\bar{a}^\varepsilon C)$.*
- (iii) $\text{tp}(\bar{a}^\varepsilon/\bar{b}C)$ *does not fork over C .*
- (iv) $\bar{a}^{2\varepsilon} \downarrow_C \bar{b}$.

Proof. We only need to provide a proof for (ii) \implies (iii), for which the original proof used local ranks.

Indeed let $(\bar{a}_n)_n$ and C' witness that a Morley sequence (\bar{b}_n) could be in $\text{tp}(\bar{b}/\bar{a}^\varepsilon C)$. We may assume that $(\bar{a}_n \bar{b}_n)_n$ is indiscernible over CC' , and we may further assume that for some \bar{a} , the pair $\bar{a}\bar{b}$ continues the sequence $(\bar{a}_n \bar{b}_n)_n$ indiscernibly over CC' . Standard arguments regarding indiscernibility provide that $\bar{b} \downarrow_{(\bar{b}_n)_n} CC'a_0$. Since $\bar{b} \downarrow_C (\bar{b}_n)_n$ we obtain $\bar{b} \downarrow_C C'a_0$. By Lemma 6.20 we have $\bar{b} \downarrow_{C'} C$ and thus $\bar{b} \downarrow_{C'} \bar{a}_0$. Let f be an automorphism sending $\bar{a}\bar{b}C'$ to $\bar{a}\bar{b}C$. Then $\bar{b} \downarrow_C f(\bar{a}_0)$ and $d(f(\bar{a}_0), \bar{a}) = d(\bar{a}_0, \bar{a}) \leq \varepsilon$. Thus $\text{tp}(\bar{a}^\varepsilon/\bar{b}C)$ does not fork over C , as desired. $\blacksquare_{6.21}$

Thus in particular $\bar{a}^0 \downarrow_C \bar{b}$ if and only if $\text{tp}(\bar{a}/\bar{b}C)$ does not fork over C , i.e., if and only if $\bar{a} \downarrow_C \bar{b}$, justifying our notation. By earlier observations, this is further equivalent to $\bar{a}^\varepsilon \downarrow_C \bar{b}$ holding for all $\varepsilon > 0$. It is further 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 $\varepsilon > 0$. Moreover, in a supersimple theory T , SU_ε ranks characterise independence: $\bar{a} \downarrow_C B$ if and only if $\text{SU}_\varepsilon(\bar{a}/C) = \text{SU}_\varepsilon(\bar{a}/BC)$ for all $\varepsilon > 0$.

Since our global ranks depend (inevitably) on a metric resolution parameter ε we may 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.

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

Lemma 6.22. *Assume $\langle G, S \rangle$ is a type-definable transitive group action in a stable theory T , $s \in S$ generic over a set A , $t \in S$ satisfying $t \downarrow_A s$. Then there is $g \in G$, $g \downarrow_A t$ 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, $g \downarrow s, t$. Then s is generic over g, t , so $gs \downarrow g, t$. By standard independence calculus we obtain $g \downarrow gs, t$. Since the action is transitive we can find $h \in G$ such that $hgs = t$, and we may take it so that $h \downarrow_{gs, t} g$. Then g is generic over t, gs, h , and so is gh , and in particular $hg \downarrow t$. Then $g' = hg$ is generic over t as required. $\blacksquare_{6.22}$

Theorem 6.23. *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 $\varepsilon > 0$. In particular, types of maximal SU_ε -rank exist.*

Proof. We may assume that $A = \emptyset$. We shall use the fact that if $p \in S_n(B)$, $q \in S_m(B)$ and $f: p(\mathfrak{M}) \rightarrow q(\mathfrak{M})$ is B -definable and isometric then $\text{SU}_\varepsilon(p) = \text{SU}_\varepsilon(q)$ for all $\varepsilon > 0$. The proof of this fact is left as an exercise to the reader.

Let $s \models p$, and assume first that p is generic. Let $t \in S$ realise an arbitrary type over \emptyset . We may nonetheless assume that $t \downarrow s$. By the Lemma there exists $g \downarrow t$ such that $gs = t$. Since multiplication by g is isometric we obtain $\text{SU}_\varepsilon(s) \geq \text{SU}_\varepsilon(s/g) = \text{SU}_\varepsilon(t/g) = \text{SU}_\varepsilon(t) = \text{SU}_\varepsilon(q)$.

Conversely, let $s \in S$ and assume that $\text{SU}_\varepsilon(s) \geq \text{SU}_\varepsilon(q)$ for all $q \in S_S(\emptyset)$ and all $\varepsilon > 0$. Let $g \in G$, $g \downarrow_A s$. Then $\text{SU}_\varepsilon(gs/g) = \text{SU}_\varepsilon(s/g) = \text{SU}_\varepsilon(s) \geq \text{SU}_\varepsilon(gs) \geq \text{SU}_\varepsilon(gs/g)$. Thus equality holds all the way for all $\varepsilon > 0$, whereby $gs \downarrow g$, so s is generic. ■_{6.23}

REFERENCES

- [Bena] Itai Ben Yaacov, *Definability of groups in \aleph_0 -stable metric structures*, Journal of Symbolic Logic, to appear, arXiv:0802.4286.
- [Benb] ———, *On theories of random variables*, submitted, arXiv:0901.1584.
- [Ben03a] ———, *Simplicity in compact abstract theories*, Journal of Mathematical Logic **3** (2003), no. 2, 163–191.
- [Ben03b] ———, *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, arXiv:0902.0118.
- [Ben09] ———, *Modular functionals and perturbations of Nakano spaces*, Journal of Logic and Analysis **1:1** (2009), 1–42, arXiv:0802.4285.
- [BU] Itai Ben Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society, to appear, arXiv:0801.4303.
- [Iov99] José Iovino, *Stable Banach spaces and Banach space structures, I and II*, Models, algebras, and proofs (Bogotá, 1995), Lecture Notes in Pure and Appl. Math., vol. 203, Dekker, New York, 1999, pp. 77–117.
- [Pil96] Anand Pillay, *Geometric stability theory*, Oxford Logic Guides, vol. 32, The Clarendon Press Oxford University Press, New York, 1996, Oxford Science Publications.
- [Poi85] Bruno Poizat, *Cours de théorie des modèles*, Nur al-Mantiq wal-Ma’rifah, 1985.

ITAÏ BEN YAACOV, UNIVERSITÉ CLAUDE BERNARD – LYON 1, INSTITUT CAMILLE JORDAN, CNRS UMR 5208, 43 BOULEVARD DU 11 NOVEMBRE 1918, 69622 VILLEURBANNE CEDEX, FRANCE

URL: <http://math.univ-lyon1.fr/~begnac/>