SEMINAR NOTES on Hrushovski's

Stable group theory and approximate subgroups [3]

Fall 2009, by Lou van den Dries, University of Illinois at Urbana-Champaign.

For motivational background we first read two other papers:

Edgar & Miller, *Borel subrings of the reals*, PAMS 131 (2003) 1121–1129
Tao, *The sum-product phenomenon in arbitrary rings*, arXiv:0806.2497v5

Next we had talks on Hrushovski's paper above. The notes below on sections 2 and 3 are a record of talks by Henson and myself, assume only rudimentary knowledge of model theory, and give detailed proofs. If possible we choose notations and formulations similar to those in Hrushovski's paper. More to follow. Needless to say, the source for all this is [3] to which we also refer for references to other original papers and motivational comments. See also Tao [4] for discussions around [3].

Let us fix some global notational conventions:

We let m, n range over $\mathbb{N} = \{0, 1, 2, \dots\}$. Let X, Y be sets and $R \subseteq X \times Y$ a relation. For $a \in X, b \in Y$ we often write $R(a, b)$ (or even aRb) to indicate that $(a, b) \in R$. For $a \in X$ we put

$$R(a) := \{b \in Y : \ R(a, b)\}.$$

We also let $\breve{R} := \{(b, a) \in Y \times X : (a, b) \in R\}$ be the reverse of R, so for $b \in Y$ we have $\breve{R}(b) = \{a \in X : \ R(a, b)\}$. Unless specified otherwise, $|X|$ is the size (cardinality) of X.

1. Independence

This section corresponds to Section 2 in [3]. Its subsections are: *Many-sorted structures*; *The monster model*; *Finite satisfiability*; *A-invariant types*; *Dividing and forking*; *Stable separation*; *Stable relations*; *The A-topology*; *Keisler measures*; *Ideals*; *Useful global types relative to an ideal.*

Many-sorted structures. Let L be an S-*sorted* first-order language, so an L-structure $\mathcal{M} = (M; \cdots)$ has a family $M = (M_s)_{s \in S}$ of underlying sets rather than a single underlying set. The set S of sorts is part of what determines L, and we define the size $|L|$ of L to be the cardinal

$$|L| := \max\{\aleph_0, \ |\text{set of nonlogical symbols of } L|, \ |S|\}.$$

For each sort $s \in S$ we have variables of sort s that we think of as ranging over the underlying set M_s, for any L-structure $\mathcal{M}$ as above.

Let $\mathcal{M} = (M; \cdots)$ be an L-structure. Given a variable v of sort s, we put $M_v := M_s$. An S-sorted multivariable is a family $x = (x_i)_{i \in I}$ where each x_i is a variable of some sort $s_i \in S$, and x_i and x_j are different variables (possibly of the same sort) whenever $i \neq j$; for such x we set

$$M_x := \prod_i M_{x_i} \qquad (\text{the } x\text{-set of } \mathcal{M}), \qquad |x| := |I| \quad (\text{size of } x),$$

and the elements $a = (a_i) \in M_x$ are also called x-tuples in $\mathcal{M}$ (of size $|I|$). A *parameter set* (in $\mathcal{M}$) is a family $A = (A_s)_{s\in S}$ with $A_s \subseteq M_s$ for all $s \in S$, and has size $|A| := \sum_s |A_s|$. Let A be a parameter set. For a variable v of sort $s \in S$ we set $A_v := A_s$ and for an S-sorted multivariable $x = (x_i)$ we put $A_x := \prod_i A_{x_i}$. We have the extended language $L(A)$, and for any S-sorted multivariable x we have the notion of an A-definable subset of M_x. Instead of "M-definable" we also write just "definable" when the ambient structure $\mathcal{M}$ is clear from the context. For definable $X \subseteq M_x$ we let $\neg X$ be the complement of X in M_x, we let $\mathrm{Def}(X)$ be the boolean algebra of definable subsets of X, and we let $\mathrm{St}(X)$ be the Stone space of this boolean algebra, so the points of $\mathrm{St}(X)$ are the ultrafilters of $\mathrm{Def}(X)$.

Suppose A is a parameter set in $\mathcal{M}$. If $X \subseteq M_x$ is A-definable, then $\mathrm{Def}(X|A)$ is the boolean algebra of its A-definable subsets, with Stone space $\mathrm{St}(X|A)$. We also set $\mathrm{St}_x(A) := \mathrm{St}(M_x|A)$, in particular, $\mathrm{St}_x(M) = \mathrm{St}(M_x)$. Points of any of these Stone spaces are often referred to as *types*.

For more on this many-sorted set-up, see Section 3 in [1], with definitions of other basic notions such as partial elementary maps, automorphisms, saturation, homogeneity, ...

The monster model.[1] Throughout $\mathbb{U} = (U; \dots)$ is a big ambient L-structure. Here "big" means that $\mathbb{U}$ comes equipped with a certain cardinal $\kappa(\mathbb{U}) > |L|$ such that $\mathbb{U}$ is $\kappa(\mathbb{U})$-saturated and strongly $\kappa(\mathbb{U})$-homogeneous. Given our ambient $\mathbb{U}$, "small" means "of size $< \kappa(\mathbb{U})$". We let x, y, z be small disjoint S-sorted multivariables and A, B small parameter sets in $\mathbb{U}$. Also, M denotes the family (M_s) of underlying sets of a small elementary submodel of $\mathbb{U}$, which, abusing language, we also denote by M. Unless we specify otherwise, "definability" is with respect to $\mathbb{U}$. For A-definable $X \subseteq \mathbb{U}_x$ and $a \in X$,

$$\mathrm{tp}_X(a|A) := \{P \in \mathrm{Def}(X|A) : \ a \in P\} \qquad \text{(the type of } a \text{ over } A),$$

and for $X = \mathbb{U}_x$ we drop subscript X in this notation or replace it x; note that $\mathrm{St}(X|A) = \{\mathrm{tp}_X(a|A) : a \in X\}$.

Let $X \subseteq \mathbb{U}_x$ (not necessarily definable). We set $X(M) := X \cap M_x$, and we say that X is *A-invariant* if $\sigma(X) = X$ for all $\sigma \in \mathrm{Aut}(\mathbb{U}|A)$. The following are equivalent:

(1) X is A-invariant;
(2) X is a union of $\mathrm{Aut}(\mathbb{U}|A)$-orbits with respect to the usual action of $\mathrm{Aut}(\mathbb{U}|A)$ on $\mathbb{U}_x$;
(3) $X = \bigcup_{p\in E} p(\mathbb{U}_x)$ for some set $E \subseteq \mathrm{St}_x(A)$.

Let $X \subseteq \mathbb{U}_x$ be definable. A point in $\mathrm{St}(X)$ is often called a *global* type and typically denoted by a bold face letter like $\mathbf{p}$. Sometimes we want to think of $\mathbf{p}$ as the set of $L(U)$-formulas $\phi(x)$ such that $\phi(X) \in \mathbf{p}$, and in this role

[1] The term "monster" does not refer here to pathology, but to being big. However, finite structures are also "big" according to the definition of that term.

as a set of formulas we denote $\mathbf{p}$ by $\mathbf{p}(x)$. Likewise for types in $\mathrm{St}(X|A)$ if X is A-definable. For more on these matters, see Section 5 of [1].

Finite satisfiability. Let $X \subseteq \mathbb{U}_x$. For an $L(U)$-formula $\phi(x)$ we put

$$\phi(X) := \{a \in X : \models \phi(a)\}.$$

Let $\Phi = \Phi(x)$ be a set of $L(U)$-formulas $\phi(x)$. Then

$$\Phi(X) := \{a \in X : \models \phi(a) \text{ for all } \phi \in \Phi\} = \bigcap_{\phi \in \Phi} \phi(X).$$

We say that Φ is *over* A if it consists of $L(A)$-formulas. We say that Φ is *finitely satisfiable* if $\Delta(\mathbb{U}_x) \neq \emptyset$ for all finite $\Delta \subseteq \Phi$; if Φ is over A, this is equivalent to $\Phi(\mathbb{U}_x) \neq \emptyset$. We say that Φ is *finitely satisfiable in* A if $\Delta(\mathbb{U}_x) \cap A_x \neq \emptyset$ for all finite $\Delta \subseteq \Phi$. We call Φ a *partial type* if $\mathbb{U}_x \neq \emptyset$, $\phi(\mathbb{U}_x) \neq \emptyset$ for all $\phi \in \Phi$, and $\phi_1 \wedge \phi_2 \in \Phi$ whenever $\phi_1,\ \phi_2$ in Φ; note that then Φ is finitely satisfiable. A set $X \subseteq \mathbb{U}_x$ is said to *meet* A if $X \cap A_x \neq \emptyset$.

Here is a useful fact about finite satisfiability in a (small) model M:

Lemma 1.1. *Suppose $\Phi(x)$ is finitely satisfiable in M. Then $\Phi(x)$ extends to a global type in $\mathrm{St}(\mathbb{U}_x)$ that is finitely satisfiable in M.*

Proof. Let $\Psi(x)$ be the set of all $L(U)$-formulas $\psi(x)$ such that $M_x \subseteq \psi(\mathbb{U}_x)$. Then $\Phi(x) \cup \Psi(x)$ is clearly finitely satisfiable (in M), and thus extends to a global type $\mathbf{p}(x) \in \mathrm{St}(\mathbb{U}_x)$. Then $\mathbf{p}(x)$ is finitely satisfiable in M: if $\theta(x)$ is an $L(U)$-formula and $\theta(\mathbb{U}_x) \cap M_x = \emptyset$, then $\neg\theta(x) \in \Psi(x) \subseteq \mathbf{p}(x)$. □

A-invariant types. Let $X \subseteq \mathbb{U}_x$ be A-definable (and hence A-invariant). Let $\mathbf{p} \in \mathrm{St}(X)$ be a global type. We consider $\mathbf{p}$ both as an ultrafilter on the boolean algebra $\mathrm{Def}(X)$, and thus as a collection of subsets of X, and as the set $\mathbf{p}(x)$ of $L(U)$-formulas $\phi(x)$ such that $\phi(\mathbb{U}_x) \in \mathbf{p}$. We say that $\mathbf{p}$ is *A-invariant* if $\sigma(P) \in \mathbf{p}$ for all $\sigma \in \mathrm{Aut}(\mathbb{U}|A)$ and all sets $P \in \mathbf{p}$. If $\mathbf{p}$ is A-invariant, then it is B-invariant for all $B \supseteq A$. Note that $\mathbf{p}(x)$ is finitely satisfiable in A iff every $P \in \mathbf{p}$ meets A.

Lemma 1.2. *Suppose $\mathbf{p}$ is finitely satisfiable in A. Then $\mathbf{p}$ is A-invariant.*

Proof. Let $P \in \mathbf{p}$ and $\sigma(P) \notin \mathbf{p}$, $\sigma \in \mathrm{Aut}(\mathbb{U}|A)$. Then $P \cap \sigma(\neg P) \in \mathbf{p}$, so we can take $a \in P \cap \sigma(\neg P) \cap A_x$, so $a = \sigma(a) \in \sigma(P)$, a contradiction. □

In particular, if $\Phi(x)$ is finitely satisfiable in M and $\phi(\mathbb{U}_x) \subseteq X$ for all $\phi \in \Phi$, then by Lemmas 1.1 and 1.2 it extends to a global type $\mathbf{p}(x) \in \mathrm{St}(X)$ that is finitely satisfiable in M and thus M-invariant.

Lemma 1.3. *Let $\mathbf{q} \in \mathrm{St}(\mathbb{U}_y)$ be A-invariant and $a, b \in \mathbb{U}_x$ with $\mathrm{tp}(a|A) = \mathrm{tp}(b|A)$. Suppose $c \in \mathbb{U}_y$ realizes $\mathbf{q} \upharpoonright Aa$ and $d \in \mathbb{U}_y$ realizes $\mathbf{q} \upharpoonright Ab$. Then $\mathrm{tp}\big((a,c)|A\big) = \mathrm{tp}\big((b,d)|A\big)$.*

Proof. Let $\phi(u, x, y)$ be an L-formula, $e \in A_u$. Then

$$\models \phi(e,a,c) \;\Leftrightarrow\; \phi(e,a,y) \in \mathbf{q}(y) \;\Leftrightarrow\; \phi(e,b,y) \in \mathbf{q}(y) \;\Leftrightarrow \models \phi(e,b,d).$$

□

Lemma 1.4. *Let $\mathbf{q} \in \mathrm{St}(\mathbb{U}_y)$ be A-invariant, and suppose (a_n) and (b_n) are sequences in $\mathbb{U}_y$ such that for all n,*

$$a_n \models \mathbf{q}\restriction Aa_0 \dots a_{n-1}, \qquad b_n \models \mathbf{q}\restriction Ab_0 \dots b_{n-1}.$$

Then $(a_n) \equiv_A (b_n)$.

Proof. Note that $\mathrm{tp}(a_0|A) = \mathrm{tp}(b_0|A) = \mathbf{q}\restriction A$. Assume inductively that

$$\mathrm{tp}((a_0, \dots, a_n)|A) = \mathrm{tp}((b_0, \dots, b_n)|A).$$

Since a_{n+1} realizes $\mathbf{q} \restriction Aa_0 \dots a_n$ and b_{n+1} realizes $\mathbf{q} \restriction Ab_0 \dots b_n$, it follows from the previous lemma that

$$\mathrm{tp}((a_0, \dots, a_n, a_{n+1})|A) = \mathrm{tp}((b_0, \dots, b_n, b_{n+1})|A).$$

□

With A-invariant $\mathbf{q} \in \mathrm{St}(\mathbb{U}_y)$, an easy recursion yields a sequence (a_n) such that $a_n \models \mathbf{q}\restriction Aa_0 \dots a_{n-1}$ for all n, and it follows easily from Lemma 1.4, that each such sequence is indiscernible over A. We call such a sequence $\mathbf{q}$-*indiscernible over* A.

Dividing and forking. For $n > 1$, a collection $\mathcal{C}$ of subsets of a set P is said to be *n-disjoint* if all sets in $\mathcal{C}$ are nonempty and for all distinct $X_1, \dots, X_n \in \mathcal{C}$ we have $X_1 \cap \dots \cap X_n = \emptyset$. A definable set $X \subseteq \mathbb{U}_x$ is said to *divide* over A if $X \neq \emptyset$ and for some $n > 1$ there is an infinite n-disjoint collection of A-conjugates of X. As often with model-theoretic notions, it is "non-dividing" rather than "dividing" that is of most interest to us.

Lemma 1.5. *Let $\phi(x, y)$ be an $L(A)$-formula, and suppose $X = \phi(\mathbb{U}_x, b)$ is nonempty, $b \in \mathbb{U}_y$. Then X divides over A iff there is an A-indiscernible sequence $(b_i)_{i\in\mathbb{N}}$ in $\mathbb{U}_y$ with $b_0 = b$ and $\bigcap_i \phi(\mathbb{U}_x, b_i) = \emptyset$.*

Proof. Suppose X divides over A. Take $n > 1$ and an n-disjoint infinite collection $\mathcal{C}$ of A-conjugates of X. For each $Y \in \mathcal{C}$ there is an A-conjugate c of b such that $Y = \phi(\mathbb{U}_x, c)$, so we have an infinite set C of A-conjugages of b with a bijection $c \mapsto \phi(\mathbb{U}_x, c) : C \to \mathcal{C}$. By Ramsey's theorem and saturation this yields an A-indiscernible sequence $(c_i)_{i\in\mathbb{N}}$ of A-conjugates of b such that for all $i_1 < \dots < i_n$ in $\mathbb{N}$ we have

$$\phi(\mathbb{U}_x, c_{i_1}) \cap \dots \cap \phi(\mathbb{U}_x, c_{i_n}) = \emptyset.$$

It remains to apply an A-automorphism of $\mathbb{U}$ that sends c_0 to b.

For the converse, assume that $(b_i)_{i\in\mathbb{N}}$ is an A-indiscernible sequence in $\mathbb{U}_y$ with $b_0 = b$ and $\bigcap_i \phi(\mathbb{U}_x, b_i) = \emptyset$. It follows easily that $\phi(\mathbb{U}_x, b_0) \neq \phi(\mathbb{U}_x, b_1)$, and thus $\phi(\mathbb{U}_x, b_i) \neq \phi(\mathbb{U}_x, b_j)$ for all $i \neq j$. Also, by saturation there are $i_1 < \dots < i_n$ in $\mathbb{N}$ with $n > 1$ such that

$$\phi(\mathbb{U}_x, b_{i_1}) \cap \dots \cap \phi(\mathbb{U}_x, b_{i_n}) = \emptyset.$$

Hence $\phi(\mathbb{U}_x, b_{j_1}) \cap \dots \cap \phi(\mathbb{U}_x, b_{j_n}) = \emptyset$ for all $j_1 < \dots < j_n$ in $\mathbb{N}$, and thus X divides over A. □

A definable set $P \subseteq \mathbb{U}_x$ is said to *fork over* A if $P \neq \emptyset$ and there are definable sets $X_1, \dots, X_n \subseteq \mathbb{U}_x$ that divide over A such that $n \geq 1$ and

$$P \subseteq X_1 \cup \cdots \cup X_n.$$

So if a definable set $P \subseteq \mathbb{U}_x$ divides over A, then it forks over A. If a definable set $P \subseteq \mathbb{U}_x$ meets A, then it doesn't fork over A.

Let Φ be a collection of nonempty definable sets in $\mathbb{U}_x$ such that $P \cap Q \in \Phi$ for all $P, Q \in \Phi$. We say that Φ *divides* over A if some set $P \in \Phi$ divides over A, and we say that Φ *forks* over A if some set $P \in \Phi$ forks over A. So if Φ divides over A, then it forks over A. Of course, these definitions also apply to a partial type $\Phi(x)$ by taking $\Phi := \{\phi(\mathbb{U}_x) : \phi(x) \in \Phi(x)\}$. If Φ doesn't divide over A, then it doesn't divide over any $B \supseteq A$. If all sets in Φ meet A, then Φ doesn't fork over A.

Lemma 1.6. *Let Φ be as above and let D be a* (*not necessarily small*) *parameter set, and suppose the sets in Φ are D-definable and Φ does not fork over A. Then Φ extends to some $p \in \mathrm{St}_x(D)$ that does not fork over A.*

Proof. Let $\Psi := \{X \in \mathrm{Def}(\mathbb{U}_x|D) : X \text{ forks over } A\}$. Suppose $P \in \Phi$ and $X_1, \dots, X_n \in \Psi$ with $n \geq 1$; we claim that then $P \cap (\neg X_1) \cap \cdots \cap (\neg X_n) \neq \emptyset$. Otherwise,

$$P \subseteq X_1 \cup \cdots \cup X_n,$$

so P would fork over A. This proves the claim.

It follows that Φ extends to a type $p \in \mathrm{St}_x(D)$ that contains all $\neg X$ with $X \in \Psi$, and thus p does not fork over A. □

Lemma 1.7. *Let $M \supseteq A$ be κ-saturated where κ is an infinite cardinal $> |A|$, and let $p \in \mathrm{St}_x(M)$. Then p divides over A iff p forks over A.*

Proof. Suppose p forks over A. Take $P \in p$ and definable sets $X_1, \dots, X_n \subseteq \mathbb{U}_x$ that divide over A such that $n \geq 1$ and $P \subseteq X_1 \cup \cdots \cup X_n$. Take a finite tuple a in the model M such that P is Aa-definable. After applying an automorphism of $\mathbb{U}$ over Aa we can arrange that all X_i are defined over M, and then $X_i \in p$ for some i. □

It is even simpler to show that, given a global type $\mathbf{p} \in \mathrm{St}(\mathbb{U}_x)$, we have:

$$\mathbf{p} \text{ divides over } A \iff \mathbf{p} \text{ forks over } A.$$

Lemma 1.8. *Let $p \in \mathrm{St}_x(A)$. Then p does not fork over A iff p has a global extension $\mathbf{p} \in \mathrm{St}(\mathbb{U}_x)$ that does not fork over A.*

Proof. If p does not fork over A, then a proof like that of Lemma 1.6 yields a global extension $\mathbf{p} \in \mathrm{St}_x(\mathbb{U})$ that doesn't fork over A. The converse is obvious. □

Note that if a global type $\mathbf{p} \in \mathrm{St}(\mathbb{U}_x)$ is A-invariant, then it doesn't divide over A, and hence doesn't fork over A.

Lemma 1.9. *Let $X \subseteq \mathbb{U}_x$ and $Y \subseteq \mathbb{U}_y$ and $f : X \to Y$ all be A-definable. Let $a \in X$, $A \subseteq B$, and suppose $\mathrm{tp}_X(a|B)$ doesn't fork over A. Then $\mathrm{tp}_Y(f(a)|B)$ doesn't fork over A.*

Proof. Suppose towards a contradiction that $f(a) \in P \in \mathrm{Def}(Y|B)$ and P forks over A. Then $P \subseteq P_1 \cup \cdots \cup P_n$ where $n \geq 1$ and where the definable sets $P_1, \ldots, P_n \subseteq \mathbb{U}_y$ divide over A. By shrinking P and $P_1, \ldots, P_n$ we can arrange that $P_1, \ldots, P_n \subseteq f(X)$. Then

$$a \in f^{-1}(P) \in \mathrm{Def}(X|B), \quad f^{-1}(P) \subseteq f^{-1}(P_1) \cup \cdots \cup f^{-1}(P_n),$$

and $f^{-1}(P_1), \ldots, f^{-1}(P_n)$ divide over A, contradicting the assumption. □

Let $\Phi(x)$ be a finitely satisfiable set of $L(U)$-formulas in x, not necessarily a partial type. Then $\Phi(x)$ generates a partial type $[\Phi(x)]$ consisting of the conjunctions $\phi_1(x) \wedge \cdots \wedge \phi_n(x)$ with $\phi_1, \ldots, \phi_n \in \Phi$. We say that $\Phi(x)$ *divides over A* if $[\Phi(x)]$ divides over A. (If $\Phi(x)$ is already a partial type, this agrees with the previous definition.) The following is an easy consequence of Lemma 1.5.

Lemma 1.10. *Let $\Phi(x, y)$ be a set of $L(A)$-formulas $\phi(x, y)$ and suppose $b \in \mathbb{U}_y$ is such that $\Phi(x, b)$ is finitely satisfiable. Then $\Phi(x, b)$ divides over A iff there is an A-indiscernible sequence (b_n) in $\mathbb{U}_y$ with $b_0 = b$ such that $\bigcup_n \Phi(x, b_n)$ is not finitely satisfiable.*

Stable separation. Let $\Phi(x, y)$ and $\Psi(x, y)$ be partial types over A, both consisting of formulas $\phi(x, y)$. We say that Φ, Ψ is *stably separated* if there is no A-indiscernible sequence $\{(a_n, b_n)\}$ in $\mathbb{U}_{x,y} = \mathbb{U}_x \times \mathbb{U}_y$ such that for all m, n:

$$m < n \implies \models \Phi(a_m, b_n), \qquad m > n \implies \models \Psi(a_m, b_n).$$

By Ramsey's theorem and compactness the following are equivalent:

(1) Φ, Ψ is stably separated;
(2) there is no sequence $\{(a_n, b_n)\}$ in $\mathbb{U}_{x,y}$ such that $\models \Phi(a_m, b_n)$ for all $m < n$ and $\models \Psi(a_m, b_n)$ for all $m > n$;
(3) there is $N \in \mathbb{N}^{\geq 1}$ such that there are no $(a_0, b_0), \ldots, (a_N, b_N)$ in $\mathbb{U}_{x,y}$ with $\models \Phi(a_m, b_n)$ for all $m < n$ and $\models \Psi(a_m, b_n)$ for all $m > n$;
(4) there are $N \in \mathbb{N}^{\geq 1}, \phi \in \Phi$, $\psi \in \Psi$ such that for all $(a_0, b_0), \ldots, (a_N, b_N)$ in $\mathbb{U}_{x,y}$, either $\models \neg\phi(a_m, b_n)$ for some $m < n$, or $\models \neg\psi(a_m, b_n)$ for some $m > n$.

The definition of "stably separated" mentions A, but by the equivalences above Φ, Ψ being stably separated doesn't depend on the choice of A such that Φ and Ψ are over A. If Φ, Ψ is stably separated, then clearly $\Phi \cup \Psi$ is not finitely satisfiable, and $\breve{\Phi}(y, x), \breve{\Psi}(y, x)$ is also stably separated. It follows from the equivalence with (3) that if Φ, Ψ is stably separated, so is Ψ, Φ (symmetry).

Let $\Phi(x,y), \Psi(x,y)$ be partial types over A as before, and let $\mathbf{q} \in \mathrm{St}_y(\mathbb{U})$ be A-invariant. Note that if $p \in \mathrm{St}_x(A)$, then either all $a \models p(x)$ satisfy $\Phi(a,y) \subseteq \mathbf{q}(y)$, or all $a \models p(x)$ satisfy $\Phi(a,y) \not\subseteq \mathbf{q}(y)$. We define:

$$\Phi, \Psi \text{ is } \mathbf{q}\text{-separated}$$

$$:\Longleftrightarrow$$

$$\text{for all } (a,b) \in \mathbb{U}_{x,y}, \text{ if } \Phi(a,y) \subseteq \mathbf{q}(y),\ b \models \mathbf{q}\restriction A, \text{ and } p(x) \cup \Psi(x,b) \text{ is finitely satisfiable with } p(x) := \mathrm{tp}(a|A), \text{ then } p(x) \cup \Psi(x,b) \text{ divides over } A.$$

Lemma 1.11. *Suppose Φ, Ψ is stably separated. Then Φ, Ψ is $\mathbf{q}$-separated (and thus by symmetry, Ψ, Φ is $\mathbf{q}$-separated).*

Proof. Let $(a,b) \in \mathbb{U}_{x,y}$ be such that $\Phi(a,y) \subseteq \mathbf{q}(y)$, $b \models \mathbf{q}\restriction A$, and the set $p(x) \cup \Psi(x,b)$ is finitely satisfiable where $p(x) := \mathrm{tp}(a|A)$. Our job is to show that $p(x) \cup \Psi(x,b)$ divides over A. Suppose it doesn't. Take a $\mathbf{q}$-indiscernible sequence (b_n) in $\mathbb{U}_y$ with $b_0 = b$. Then by Lemma 1.10 $p(x) \cup \bigcup_n \Psi(x,b_n)$ is finitely satisfiable. By Lemma 1.4 and subsequent remark we can choose recursively elements $a_0, a_1, \dots \in \mathbb{U}_x$ and $c_0, c_1, \dots \in \mathbb{U}_y$ such that for all n,

(1) $a_n \models p(x) \cup \bigcup_{m<n} \Psi(x,c_m)$;
(2) $c_n \models \mathbf{q}\restriction A a_0 \dots a_{n-1} c_0 \dots c_{n-1}$.

To choose a_n for $n > 0$, keep in mind that the initial segment $c_0, c_1, \dots, c_{n-1}$ begins an A-indiscernible sequence of the same type over A as $(b_i)_{i\in\mathbb{N}}$. If $m > n$, then by (1) we have $\models \Psi(a_m, c_n)$. If $m < n$, then by (1) we have $a_m \models p(x)$, so $\Phi(a_m, y) \subseteq \mathbf{q}(y)$, and hence $\models \Phi(a_m, c_n)$. Thus Ψ, Φ is not stably separated, contradicting the assumption that Φ, Ψ is stably separated. □

Stable relations. Let $R \subseteq \mathbb{U}_x \times \mathbb{U}_y = \mathbb{U}_{x,y}$ in what follows. We say that the relation R is *stable over A* if R is A-invariant and for all $a, a' \in \mathbb{U}_x$, $b, b' \in \mathbb{U}_y$ with $R(a,b)$ and $\neg R(a',b')$, the pair $\mathrm{tp}((a,b)|A)$, $\mathrm{tp}((a',b')|A)$ of (x,y)-types over A is stably separated. Note that if R is stable over A, so are $\neg R$ and $\breve{R} \subseteq \mathbb{U}_y \times \mathbb{U}_x$, and R is stable over every $B \supseteq A$. It is also clear from this definition that if I is any index set and $R_i \subseteq \mathbb{U}_x \times \mathbb{U}_y$ is stable over A for all $i \in I$, then $\bigcup_i R_i$ and $\bigcap_i R_i$ are stable over A.

Lemma 1.12. *Suppose R is A-invariant. Then the following are equivalent:*

(1) *R is not stable over A;*
(2) *there is an A-indiscernible sequence $\{(a_n, b_n)\}$ in $\mathbb{U}_{x,y}$ such that $R(a_m, b_n)$ for all $m < n$ and $\neg R(a_m, b_n)$ for all $m > n$.*

Proof. Suppose $\{(a_n,b_n)\}$ is as in (2). Put $a = a_0, b = b_1, a' = a_1, b' = b_0$. Then $R(a,b)$ and $\neg R(a',b')$, and $\{(a_n,b_n)\}$ witnesses that the pair

$$\mathrm{tp}((a,b)|A),\ \mathrm{tp}((a',b')|A)$$

is not stably separated, and so R is not stable. Conversely, let $a, a' \in \mathbb{U}_x$ and $b, b' \in \mathbb{U}_y$ be such that $R(a,b)$ and $\neg R(a',b')$, and let $\{(a_n,b_n)\}$ be an

A-indiscernible sequence in $\mathbb{U}_{x,y}$ witnessing that $\operatorname{tp}((a,b)|A)$, $\operatorname{tp}((a',b')|A)$ is not stably separated. Then this sequence is as in (2). □

An immediate consequence of Lemma 1.12 is that if $X \subseteq \mathbb{U}_x$ and $Y \subseteq \mathbb{U}_y$ are A-invariant, then $X \times Y$ is stable over A.

Lemma 1.13. *Let $p \in \operatorname{St}_x(A)$, let $\mathbf{q} \in \operatorname{St}_y(\mathbb{U})$ be A-invariant, and suppose R is stable over A. Let a, a' range over $\mathbb{U}_x$ and b, b' over $\mathbb{U}_y$.*

(1) *Assume $R(a,b)$, $a, a' \models p(x)$, $b \models \mathbf{q} \restriction Aa$, $b' \models \mathbf{q} \restriction A$ and $\operatorname{tp}(a'|Ab')$ does not divide over A. Then $R(a',b')$.*
(2) *Assume $a, a' \models p$ and $b, b' \models \mathbf{q} \restriction A$, and the types $\operatorname{tp}(a|Ab)$ and $\operatorname{tp}(a'|Ab')$ do not divide over A. Then $R(a,b) \iff R(a',b')$.*

Proof. For (1), put $\Phi(x,y) := \operatorname{tp}((a,b)|A)$, $\Psi(x,y) := \operatorname{tp}((a',b')|A)$, so

$$\Phi(a,y) = \operatorname{tp}(b|Aa), \qquad \Psi(x,b') = \operatorname{tp}(a'|Ab').$$

Suppose towards a contradiction that $\neg R(a',b')$. Then the pair Φ, Ψ is stably separated, so $\mathbf{q}$-separated by Lemma 1.11. Now $\Phi(a,y) = \operatorname{tp}(b|Aa) = \mathbf{q}\restriction Aa$, hence $\Phi(a,y) \subseteq \mathbf{q}(y)$. Also,

$$p(x) = \operatorname{tp}(a'|A) \subseteq \Psi(x,b'),$$

and hence $p(x) \cup \Psi(x,b')$ is finitely satisfied. Therefore $p(x) \cup \Psi(x,b') = \Psi(x,b') = \operatorname{tp}(a'|Ab')$ divides over A, contradicting the assumption in (1).

To prove (2), take $c \in \mathbb{U}_y$ such that $c \models \mathbf{q} \restriction Aa$. If $R(a,c)$, then by (1) applied to the pairs (a,c) and (a,b) we otain $R(a,b)$, and (1) applied to (a,c) and (a',b') gives $R(a',b')$. If $\neg R(a,c)$ we argue in the same way with the stable relation $\neg R$ to obtain $\neg R(a,b)$ and $\neg R(a',b')$. □

Lemma 1.14. *Assume R is stable over M, let $p \in \operatorname{St}_x(M)$, $q \in \operatorname{St}_y(M)$, and let (a,b) range over $p(\mathbb{U}_x) \times q(\mathbb{U}_y)$. Then the following are equivalent:*

(i) *$R(a,b)$ for some (a,b) such that $\operatorname{tp}(a|Mb)$ does not divide over M;*
(ii) *$R(a,b)$ for all (a,b) such that $\operatorname{tp}(a|Mb)$ does not divide over M.*
(iii) *$R(a,b)$ for some (a,b) such that $\operatorname{tp}(a|Mb)$ does not fork over M;*
(iv) *$R(a,b)$ for all (a,b) such that $\operatorname{tp}(a|Mb)$ does not fork over M.*

These four conditions are also equivalent to each of the four conditions obtained by replacing "$\operatorname{tp}(a|Mb)$" with "$\operatorname{tp}(b|Ma)$".

Proof. By Lemma 1.1 we have a global M-invariant $\mathbf{q} \in \operatorname{St}_y(\mathbb{U})$ that extends q. Using $\mathbf{q}$ the implication (i) ⇒ (ii) follows from (2) of Lemma 1.13.

For the converse it is enough to produce (a,b) such that $\operatorname{tp}(a|Mb)$ does not divide over M. Take a global M-invariant $\mathbf{p} \in \operatorname{St}_x(\mathbb{U})$ that extends p, so $\mathbf{p}$ doesn't divide over M. Then for any b, $\mathbf{p}\restriction Mb$ doesn't divide over M, so if a realizes $\mathbf{p}\restriction Mb$, then $\operatorname{tp}(a|Mb)$ doesn't divide over M.

This construction of a pair (a,b) works also with "fork" instead of "divide" and thus produces a pair (a,b) such that $\operatorname{tp}(a|Mb)$ doesn't fork over M, and hence also doesn't divide over M. This yields (ii) ⇒ (iii), and (iv)⇒ (iii). The direction (iii) ⇒ (i) follows in view of "not forking" implying "not dividing". The latter also gives (iv) $\Rightarrow$ (ii). This proves the equivalence of conditions (i)–(iv).

By symmetry considerations, the four conditions obtained from these by replacing "$\mathrm{tp}(a|Mb)$" with "$\mathrm{tp}(b|Ma)$" are also pairwise equivalent; to show they are equivalent to (i)–(iv) it is enough to have a pair (a,b) such that $\mathrm{tp}(a|Mb)$ and $\mathrm{tp}(b|Ma)$ don't divide over M. To get such a pair we use Lemma 1.1 to obtain an M-invariant $\mathbf{r} \in \mathrm{St}_{x,y}(\mathbb{U})$ that extends $p \cup q$. Then $\mathbf{r}$ doesn't divide over M, so any $(a,b) \models \mathbf{r} \restriction M$ has the desired property. $\square$

Let $\mathbf{q} \in \mathrm{St}_y(\mathbb{U})$, and $a \in \mathbb{U}_x$. Then

$$(\mathbf{q} \restriction Aa)(\mathbb{U}_y) \;=\; \{b \in \mathbb{U}_y :\; b \models \mathbf{q} \restriction Aa\} \;\neq \emptyset.$$

Any two elements of $(\mathbf{q} \restriction Aa)(\mathbb{U})$ are Aa-conjugate, so for A-invariant R,

$$(\mathbf{q} \restriction Aa)(\mathbb{U}_y) \subseteq R(a) \iff (\mathbf{q} \restriction Aa)(\mathbb{U}_y) \cap R(a) \neq \emptyset.$$

It follows that if R is A-invariant and $B \supseteq A$, then

$$(\mathbf{q} \restriction Aa)(\mathbb{U}_y) \subseteq R(a) \iff (\mathbf{q} \restriction Ba)(\mathbb{U}_y) \subseteq R(a).$$

In the rest of this subsection the multivariables x, y are similar.

Lemma 1.15. *Suppose R is stable over M, $\mathbf{q}, \mathbf{r} \in \mathrm{St}_y(\mathbb{U})$ do not divide over M, and $\mathbf{q} \restriction M = \mathbf{r} \restriction M$. Then for all $a \in \mathbb{U}_x = \mathbb{U}_y$,*

$$(\mathbf{q} \restriction Ma)(\mathbb{U}_y) \subseteq R(a) \iff (\mathbf{r} \restriction Ma)(\mathbb{U}_y) \subseteq R(a).$$

Proof. Let $a \in \mathbb{U}_x$ and suppose $(\mathbf{q} \restriction Ma)(\mathbb{U}_y) \subseteq R(a)$, that is, $R(a,b)$ for all $b \models \mathbf{q} \restriction Ma$. Put $p := \mathrm{tp}_x(a|M)$, take an M-invariant $\mathbf{p} \in \mathrm{St}_x(\mathbb{U})$ that extends p. Take $c \models \mathbf{q} \restriction Ma$, so $R(a,c)$ and $\mathrm{tp}(c|Ma)$ doesn't divide over M. Then Lemma 1.14 with $q := \mathbf{q} \restriction M$ yields that $R(a',c)$ for all $a' \models \mathbf{p} \restriction Mc$. Taking such an a' and noting that $c \models \mathbf{r} \restriction M = \mathbf{q} \restriction M$ we can apply the same lemma again to conclude that $R(a,b)$ for all $b \models \mathbf{r} \restriction Ma$, that is, $(\mathbf{r} \restriction Ma)(\mathbb{U}_y) \subseteq R(a)$. $\square$

Let $S = S^{\mathrm{nd}}_{A,y}$ be the set of global types $\mathbf{q} \in \mathrm{St}_y(\mathbb{U})$ that do not divide over A. We define an equivalence relation E on S by: $\mathbf{q}\ E\ \mathbf{r}$ if and only if

for every $R \subseteq \mathbb{U}_x \times \mathbb{U}_y$ that is stable over A, and all $a \in \mathbb{U}_y$,

$$(\mathbf{q} \restriction Aa)(\mathbb{U}_y) \subseteq R(a) \iff (\mathbf{r} \restriction Aa)(\mathbb{U}_y) \subseteq R(a).$$

Corollary 1.16. *Suppose $|A|, |y| \leq |L|$. Then $|S/E| \leq 2^{|L|}$.*

Proof. Take $M \supseteq A$ with $|M| \leq |L|$. By Lemma 1.15 and the remarks preceding this lemma, we have for all $\mathbf{q}, \mathbf{r} \in S$, if $\mathbf{q} \restriction M = \mathbf{r} \restriction M$, then $\mathbf{q}\ E\ \mathbf{r}$. This yields the desired estimate since $|\mathrm{Def}_x(\mathbb{U}|M)| \leq |L|$. $\square$

The A-topology. The A-definable subsets of $\mathbb{U}_x$ form a basis for a certain topology on $\mathbb{U}_x$, the A-topology. The A-open subsets of $\mathbb{U}_x$ are the unions of A-definable subsets of $\mathbb{U}_x$; instead of *A-open* one also uses the term *$\bigvee$-definable over A*. The A-closed subsets of $\mathbb{U}_x$ are the intersections of A-definable subsets of $\mathbb{U}_x$, equivalently, the sets $\Phi(\mathbb{U}_x) \subseteq \mathbb{U}_x$, with $\Phi(x)$ a set of

$L(A)$-formulas in x; instead of A-*closed* one also uses the terms $\bigwedge$-*definable over* A and *type-definable over* A. The A-definable subsets of $\mathbb{U}_x$ are exactly the A-clopen subsets of $\mathbb{U}_x$, and $\mathbb{U}_x$ is compact (but usually not hausdorff) in the A-topology. For related basic facts on type-definability, see [2]. One use of A-closed sets (respectively, A-open sets) is the possibility of forming inverse limits (respectively, direct limits); this is also a reason for allowing x to be an infinite (but small) tuple of variables. A related viewpoint is that A-open sets are like locally compact spaces, with A-definable and A-closed sets more like compact spaces. Note that A-open and A-closed subsets of $\mathbb{U}_x$ are A-invariant.

For use in the next section we prove here a general fact, Lemma 1.17, which is implicit in Hrushovski's proof of his stabilizer theorem 3.4 in [3]. First some introductory remarks.

Let $X \subseteq \mathbb{U}_x$ be A-closed, and let $(X_i)_{i\in I}$ be a defining system for X over A in the sense of [2], that is, I is small, $X_i \in \mathrm{Def}(\mathbb{U}_x|A)$ for all $i \in I$, and $X = \bigcap_i X_i$. Then for any small collection $\Phi(x)$ of formulas $\phi(x)$ over A,

$$\Phi(X) = \bigcap_i \Phi(X_i) = \bigcap_{i\in I,\ \phi\in\Phi} \phi(X_i),$$

so $\Phi(X)$ is A-closed with defining system $(\phi(X_i))_{i\in I,\ \phi\in\Phi}$.

Assume also that $p \in \mathrm{St}_x(A)$. Then $p(X) = \emptyset$ or $p(X) = p(\mathbb{U}_x)$. This is because $p(X) \subseteq p(\mathbb{U}_x)$, $p(X)$ is an A-invariant subset of $\mathbb{U}_x$, and $p(\mathbb{U}_x)$, as an orbit, is a minimal nonempty A-invariant subset of $\mathbb{U}_x$.

Lemma 1.17. *Let $X \subseteq \mathbb{U}_x$ be M-closed and E an M-closed equivalence relation on X such that X/E is small. Let $p \in \mathrm{St}_x(M)$ and $p(X) \neq \emptyset$. Then $p(X) \subseteq C$ for a (necessarily unique) E-class $C \in X/E$, and this class C is M-closed.*

Proof. Let $(X_i), (E_i)$ be a directed defining system for X, E as defined in [2]. Note that $p(X) = p(\mathbb{M}_x)$ by the remarks above. We extend p to an M-invariant global type $\mathbf{p} \in \mathrm{St}(\mathbb{U}_x)$, and then take an elementary extension $\mathbb{U}'$ of $\mathbb{U}$ where $\mathbf{p}$ is realized by an element $a \in \mathbb{U}'_x$. This yields the definable sets $X'_i \subseteq \mathbb{U}'_x$ and $E'_i \subseteq X'_i \times X'_i$, the set $X' = \bigcap_i X'_i \subseteq \mathbb{U}'_x$ and the equivalence relation $E' = \bigcap_i E'_i$ on X'. If $C \in X/E$, then $C' \in X'/E'$, and the map

$$C \mapsto C' \ : \ X/E \to X'/E'$$

is a bijection (Lemma 3.4 of [2]). We have $a \in \mathbf{p}(\mathbb{U}'_x) \subseteq p(\mathbb{U}'_x)$. We claim that $p(\mathbb{U}'_x) \subseteq X'$: to see this, note that for each $i \in I$ we have $p(\mathbb{U}_x) = p(X) \subseteq X_i$, so we get $\phi(x) \in p(x)$ with $\phi(\mathbb{U}_x) \subseteq X_i$, hence $p(\mathbb{U}'_x) \subseteq \phi(\mathbb{U}'_x) \subseteq X'_i$. This yields the claim. It follows that we have $C \in X/E$ such that $a \in C'$. Take $c \in C$. For $i \in I$ we have $a \in C' \subseteq E_i(c)'$, so $E_i(c) \in \mathbf{p}$. On the other hand, if $d \in X$ and $(c, d) \notin E$, then there is i such that $E_i(c) \cap E_i(d) = \emptyset$ and thus $E_i(d) \notin \mathbf{p}$. In this way $\mathbf{p}$ picks out a unique E-class, namely C. But $\mathbf{p}$ is M-invariant, so C is M-invariant, and since C is type-definable (over Mc), it follows that C is M-closed.

Next we claim that $p(\mathbb{U}_x) \cap E_i(c) \neq \emptyset$ for all i. To see this, suppose $i \in I$ is such that $p(\mathbb{U}_x) \cap E_i(c) = \emptyset$. Then we have $\phi(x) \in p(x)$ such that $\phi(\mathbb{U}_x) \cap E_i(c) = \emptyset$, and since $\phi(\mathbb{U}_x) \in \mathbf{p}$, this would give $E_i(c) \notin \mathbf{p}$, a contradiction. This proves our claim, and since $C = \bigcap_i E_i(c)$, compactness yields $p(\mathbb{U}_x) \cap C \neq \emptyset$. Since $p(\mathbb{U}_x)$ is a minimal nonempty M-invariant subset of $\mathbb{U}_x$, this yields $p(X) = p(\mathbb{U}_x) \subseteq C$. □

It will also be useful to extend some earlier constructions to A-open sets.

Let $X \subseteq \mathbb{U}_x$ be A-open. Then we define:

$$\begin{aligned}\mathrm{Def}(X) &:= \{P \subseteq X : \ P \text{ is definable}\},\\ \mathrm{Def}(X|A) &:= \{P \subseteq X : \ P \text{ is } A\text{-definable}\},\end{aligned}$$

and for any $a \in X$ and any parameter set D (not necessarily small),

$$\begin{aligned}\mathrm{tp}_X(a|D) &:= \{P \in \mathrm{Def}(X) : \ a \in P\},\\ \mathrm{St}(X|D) &:= \{\mathrm{tp}_X(a|D) : \ a \in X\},\end{aligned}$$

in particular, $\mathrm{tp}_X(a|A) := \{P \in \mathrm{Def}(X|A) : \ a \in P\}$. By an X-formula we mean an L-formula $\phi(x, y)$ such that $\phi(\mathbb{U}_x, b) \subseteq X$ for all $b \in \mathbb{U}_y$ (and thus $\phi(\mathbb{U}_x, b) = \phi(X, b)$ for all $b \in \mathbb{U}_y$). Because X is A-open, every $P \in \mathrm{Def}(X)$ equals $\phi(\mathbb{U}_x, b)$ for some X-formula $\phi(x, y)$ and some $b \in \mathbb{U}_y$.

Keisler measures. Let $X \subseteq \mathbb{U}_x$ be A-open. A *Keisler measure* on X is a finitely additive measure

$$\mu : \mathrm{Def}(X) \to [0, \infty] = \mathbb{R}^{\geq 0} \cup \{\infty\}, \qquad (\text{in particular, } \mu(\emptyset) = 0)^2.$$

Let $\mu : \mathrm{Def}(X) \to [0, \infty]$ be a Keisler measure. Then we have for each X-formula $\phi(x, y)$ the function

$$\mu_\phi : M_y \to [0, \infty], \quad \mu_\phi(b) = \mu\phi(X, b).$$

We say that μ is *A-invariant* if $\mu(P) = \mu(\sigma P)$ for all $P \in \mathrm{Def}(X)$ and $\sigma \in \mathrm{Aut}(\mathbb{U}|A)$, equivalently, for each X-formula $\phi(x, y)$ we have $\mu_\phi(b) = \mu_\phi(c)$ whenever $b, c \in \mathbb{U}_y$ are A-conjugate (and so μ_ϕ induces a function $\mu_\phi : \mathrm{St}_y(A) \to [0, \infty]$ by $\mu_\phi(\mathrm{tp}(b|A)) := \mu_\phi(b)$ for $b \in \mathbb{U}_y$). We say that μ is *A-definable* (in $\mathbb{U}$) if μ is A-invariant and each function $\mu_\phi : \mathbb{U}_y \to [0, \infty]$ as above is A-continuous, equivalently, μ is A-invariant and each induced function $\mu_\phi : \mathrm{St}_y(A) \to [0, \infty]$ is continuous.

Ideals. Let $\mathcal{C}$ be a collection of subsets of $\mathbb{U}_x$. We say that $\mathcal{C}$ is *A-invariant* if $\sigma P \in \mathcal{C}$ for all $P \in \mathcal{C}$ and $\sigma \in \mathrm{Aut}(\mathbb{U}|A)$. Note that then for every L-formula $\phi(x, y)$ we have a unique set $E_{\mathcal{C},\phi} \subseteq \mathrm{St}_y(A)$ of types such that for all $b \in \mathbb{U}_y$,

$$\phi(\mathbb{U}_x, b) \in \mathcal{C} \iff \mathrm{tp}(b|A) \in E_{\mathcal{C},\phi},$$

and that this equivalence determines $\mathcal{C}$ in terms of the sets $E_{\mathcal{C},\phi}$. The collections $\mathcal{C}$ we have in mind are ideals in $\mathrm{Def}(X)$ for A-open X.

---

[2] For A-definable X one can also impose $\mu(X) = 1$.

In the rest of this subsection $X \subseteq \mathbb{U}_x$ is A-open and I is an ideal of $\mathrm{Def}(X)$, so I is a collection of definable subsets of X such that $\emptyset \in I$, and for all $P, Q \in \mathrm{Def}(X)$,

$$P, Q \in I \implies P \cup Q \in I, \qquad P \subseteq Q \in I \implies P \in I.$$

We say that I is *proper* if $I \neq \mathrm{Def}(X)$. (This notion is relative to the ambient X: if $Y \subseteq X$ is A-open, then $\mathrm{Def}(Y)$ is an improper ideal of $\mathrm{Def}(Y)$ but can be proper as an ideal of $\mathrm{Def}(X)$.)

Examples. $\{P \in \mathrm{Def}(X) : \ P \text{ forks over } A\} \cup \{\emptyset\}$, is an A-invariant ideal of $\mathrm{Def}(X)$, the *forking ideal over* A (in X). If μ is an A-invariant Keisler measure on X, then

$$\{P \in \mathrm{Def}(X) : \mu(P) = 0\}$$

is an A-invariant ideal of $\mathrm{Def}(X)$, called the *zero ideal of* μ, and is a proper ideal if $\mu(P) > 0$ for some $P \in \mathrm{Def}(X)$.

We say that I is *A-definable* (respectively, *A-closed, A-open*) if for each L-formula $\phi(x, y)$ the set $\{b \in \mathbb{U}_y : \ \phi(\mathbb{U}_x, b) \in I\}$ is A-definable (respectively, A-closed, A-open). Note that if I is A-closed or A-open, then I is A-invariant. If μ is an A-definable Keisler measure on X, then its zero ideal is clearly A-closed.

Definition. *I is S1* over A if I is A-invariant, and for every L-formula $\phi(x, y)$ and A-indiscernible sequence (b_n) in $\mathbb{U}_y$, if $\phi(\mathbb{U}_x, b_0) \in \mathrm{Def}(X)$ (and thus $\phi(\mathbb{U}_x, b_n) \in \mathrm{Def}(X)$ for all n), and $\phi(\mathbb{U}_x, b_m) \cap \phi(\mathbb{U}_x, b_n) \in I$ for all $m \neq n$, then $\phi(\mathbb{U}_x, b_n) \in I$ for some n (and hence for all n).

For A-invariant I, the following are equivalent:

(1) I is S1 over A;
(2) for every A-definable relation $R \subseteq X \times \mathbb{U}_y$ and every sequence (b_n) in $\mathbb{U}_y$ with $\breve{R}(b_n) \notin I$ for all n, there are $m < n$ with $\breve{R}(b_m) \cap \breve{R}(b_n) \notin I$.

This equivalence follows as usual by Ramsey's theorem and saturation.

The zero ideal of an A-invariant Keisler measure μ on X with $\mu(X) < \infty$ is clearly S1. The forking ideal over A in X is contained in every S1-ideal over A:

Lemma 1.18. *Suppose I is S1 over A and $P \in \mathrm{Def}(X)$ forks over A. Then $P \in I$.*

Proof. We can reduce to the case that P divides over A. Suppose towards a contradiction that $P \notin I$. Take an L-formula $\phi(x, y)$ and $b \in \mathbb{U}_y$ such that $P = \phi(\mathbb{U}_x, b)$. Then by Lemma 1.5 we have an A-indiscernible sequence (b_n) in $\mathbb{U}_y$ with $b_0 = b$ and $\phi(\mathbb{U}_x, b_0) \cap \cdots \cap \phi(\mathbb{U}_x, b_n) = \emptyset$ for some $n \geq 1$. Note that $\phi(\mathbb{U}_x, b_n) \in \mathrm{Def}(X)$ for all n. Take m maximal such that $\phi(\mathbb{U}_x, b_0) \cap \cdots \cap \phi(\mathbb{U}_x, b_m) \notin I$, and put

$$P_i := \phi(\mathbb{U}_x, b_0) \cap \cdots \cap \phi(\mathbb{U}_x, b_{m-1}) \cap \phi(\mathbb{U}_x, b_{m+i}), \quad i \in \mathbb{N}.$$

Then $P_i \notin I$ for all i and $P_i \cap P_j \in I$ for all $i \neq j$. The sequence

$$\{(b_0, \dots, b_{m-1}, b_{m+i})\}_{i \in \mathbb{N}}$$

is A-indiscernible and I is S1 over A, a contradiction. □

Lemma 1.19. *Let $X \subseteq \mathbb{U}_x$ and $Y \subseteq \mathbb{U}_y$ be A-definable, let $Z \subseteq \mathbb{U}_z$ be A-open and J an ideal of* $\operatorname{Def}(Z)$ *that is S1 over A. Let $P \subseteq X \times Z$ and $Q \subseteq Y \times Z$ be A-definable, and define $R \subseteq X \times Y \subseteq \mathbb{U}_x \times \mathbb{U}_y$ by*

$$R(a, b) \iff P(a) \cap Q(b) \in J.$$

Then R is stable over A.

Proof. It is clear that R is A-invariant. Let $\{(a_n, b_n)\}$ be an A-indiscernible sequence in $\mathbb{U}_{x,y}$ with $R(a_m, b_n)$ for all $m < n$, so $a_n \in X$, $b_n \in Y$ for all n. By Lemma 1.12 it suffices to show that then $R(a_m, b_n)$ for some $m > n$. First, $R(a_n, b_n)$ for all n: otherwise, $\neg R(a_n, b_n)$ for all n, and so for $C_n := P(a_n) \cap Q(b_n)$ we have $C_n \notin J$ for all n and $C_m \cap C_n \in J$ for all $m \neq n$, which contradicts that J is S1 over A. Next, the sequence $\{(a_{2n}, b_{2n}, a_{2n+1}, b_{2n+1})\}$ is also A-indiscernible, and so is $\{(c_n, d_n)\} := \{(a_{2n+1}, b_{2n})\}$. Since $R(c_m, d_n)$ for all $m < n$, the above gives $R(c_n, d_n)$ for all n, that is, $R(a_{2n+1}, b_{2n})$ for all n. □

A set $P \subseteq X$ is said to be I-wide if $P \not\subseteq Y$ for all $Y \in I$. If $P \in \operatorname{Def}(X)$, then P is I-wide iff $P \notin I$. A partial type $\Phi(x)$ is said to be *in* X if $\phi(\mathbb{U}_x) \subseteq X$ for all $\phi \in \Phi$, and is said to be *I-wide* if Φ is in X and $\phi(\mathbb{U}_x)$ is I-wide for all $\phi \in \Phi(x)$. If $\Phi(x)$ is a partial type in X over some B, then Φ is I-wide iff $\Phi(X)$ is I-wide.

Lemma 1.20. *Let $P \in \operatorname{Def}(X)$ and suppose the nonempty partial type $\Phi(x)$ is in P and is I-wide. Then Φ extends to an I-wide global type $\mathbf{p} \in \operatorname{St}(P)$.*

Proof. If $Y \in I, Y \subseteq P$ and $\phi \in \Phi$, then $(P \setminus Y) \cap \phi(\mathbb{U}_x) \neq \emptyset$, since otherwise $\phi(\mathbb{U}_x) \subseteq Y$. This gives $\mathbf{p} \in \operatorname{St}(P)$ extending Φ such that $P \setminus Y \in \mathbf{p}$ for all $Y \in I$ with $Y \subseteq P$; such $\mathbf{p}$ is I-wide. □

Theorem 1.21. *Let $Z \subseteq \mathbb{U}_z$ be A-open and J an ideal of* $\operatorname{Def}(Z)$ *that is S1 over A. Let $a \in \mathbb{U}_x, b, b' \in \mathbb{U}_y, c \in Z$ be such that $\operatorname{tp}_Z(c|Aab)$ is J-wide, $\operatorname{tp}(b|A) = \operatorname{tp}(b'|A)$, $\operatorname{tp}(b|Aa)$ and $\operatorname{tp}(b'|Aa)$ do not divide over A, and $\operatorname{tp}(a|A)$ extends to an A-invariant global type. Then there exists $c' \in Z$ such that $\operatorname{tp}_Z(c'|Aab')$ is J-wide, and*

$$\operatorname{tp}((a, c')|A) = \operatorname{tp}((a, c)|A), \qquad \operatorname{tp}((b', c')|A) = \operatorname{tp}((b, c)|A).$$

Proof. Let $P \subseteq \mathbb{U}_x \times Z$ and $Q \subseteq \mathbb{U}_y \times Z$ be A-definable with $P(a, c)$ and $Q(b, c)$. By compactness it is enough to find for any such P, Q an element $c' \in Z$ such that $P(a, c')$, $Q(b', c')$ and $\operatorname{tp}_Z(c'|Aab')$ is J-wide, that is, it suffices to show that $P(a) \cap Q(b') \notin J$.

Let $R \subseteq \mathbb{U}_x \times \mathbb{U}_y$ be defined by $R(d, e) \iff P(d) \cap Q(e) \in J$. Then R is stable over A by Lemma 1.19. Also $P(a) \cap Q(b) \in \operatorname{tp}_Z(c|Aab)$ and $\operatorname{tp}_Z(c|Aab)$

is J-wide, so $\neg R(a,b)$. Hence by part (2) of Lemma 1.13 (interchanging the roles of x, y), we obtain $\neg R(a,b')$, that is, $P(a) \cap Q(b') \notin J$, as desired. $\square$

Useful global types relative to an ideal. Such types are provided by the next three lemmas, in three slightly different situations. We continue to assume that $X \subseteq \mathbb{U}_x$ is A-open and that I is an ideal of $\mathrm{Def}(X)$.

Lemma 1.22. *Suppose $A = M$ and I is proper and M-open. Then there is a global type $\mathbf{p} \in \mathrm{St}(X)$, finitely satisfiable in M, such that for all $a, b \in X$, if $a \models \mathbf{p}\restriction M$ and $b \models \mathbf{p}\restriction Ma$, then $\mathrm{tp}_X(a|Mb)$ is I-wide.*

Proof. Since I is proper, we have $P \in \mathrm{Def}(X)$ such that $P \setminus Y \neq \emptyset$ for all $Y \in I$. This yields a type $p \in \mathrm{St}(X|M)$ such that $P \setminus Y \in p$ for all M-definable $Y \in I$. Then Lemma 1.1 yields a global type $\mathbf{p} \in \mathrm{St}(X)$ that extends p and is finitely satisfiable in M. Let $a, b \in X$ with $a \models p$ and $b \models \mathbf{p} \restriction Ma$, and suppose towards a contradiction that $\mathrm{tp}_X(a|Mb)(X) \subseteq Y \in I$. Then by compactness we have an $L(A)$-formula $\phi(x,y)$ with x and y similar, such that $\models \phi(a,b)$ and $\phi(X,b) \subseteq Y$, and thus $\phi(X,b) \in I$. Now I is M-open, so we have an $L(M)$-formula $\theta(y)$ such that $\models \theta(b)$, and $\phi(X,b') \in I$ for all $b' \in \theta(X)$. Since $\mathrm{tp}_X(b|Ma)$ is finitely satisfiable in M we get $b' \in \theta(X) \cap M_y$ such that $\models \phi(a,b')$. Then $\phi(X,b')$ is an M-definable set in I, so $X \setminus \phi(X,b') \in p$ and thus $\models \neg\phi(a,b')$, a contradiction. $\square$

The next lemma has some Fubini-type assumptions:

Lemma 1.23. *Suppose $A = M$, L and M are countable, and I is M-closed and proper. With y similar to x, assume also that J is an ideal of $\mathrm{Def}(X^2|M)$ such that for all M-definable $P \subseteq X$ and $R \subseteq X^2$,*

(1) *if $P^2 \in J$, then $P \in I$;*
(2) *if $R(a) \in I$ for all $a \in X$ with I-wide $\mathrm{tp}(a|M)$, then $R, \breve{R} \in J$.*

Then there exists a global type $\mathbf{p} \in \mathrm{St}(X)$, finitely satisfiable in M, such that $\mathrm{tp}_X(a|Mb)$ and $\mathrm{tp}_X(b|Ma)$ are I-wide for all $a \models \mathbf{p}\restriction M$ and $b \models \mathbf{p}\restriction Ma$.

Proof. We first construct a certain I-wide $p \in \mathrm{St}(X|M)$, and then extend it to a global type $\mathbf{p}$ as required.

Claim 1. Let $P \subseteq X$ and $R_1, R_2, R_3 \subseteq X^2$ be M-definable such that $P \notin I$ and $P^2 \subseteq R_1 \cup R_2 \cup R_3$. Then there is an M-definable $Q \subseteq P$ with $Q \notin I$ such that for all $a, b \in Q$,

$$(*) \qquad R_1(a) \notin I \text{ or } \breve{R}_2(b) \notin I \text{ or } \breve{R}_3(c) \supseteq Q \text{ for some } c \in X(M).$$

Proof of Claim 1. Suppose $P \cap \breve{R}_3(c) \notin I$ for some $c \in X(M)$. Take such c, put $Q := P \cap \breve{R}_3(c)$, and note that then the third option in (*) holds for all $a, b \in Q$. So it remains to consider the case that $P \cap \breve{R}_3(c) \in I$ for all $c \in X(M)$. Now I is M-closed, so $P \cap \breve{R}_3(c) \in I$ for all $c \in X$. Hence $(P \times X) \cap R_3 \in J$, by (2).

If there is an M-definable $Q \subseteq P$ with $Q \notin I$ such that $R_1(a) \notin I$ for all $a \in Q$, then the first option of (*) holds. So we can assume that for all

M-definable $Q \subseteq P$ with $Q \notin I$ there is $a \in Q$ with $R_1(a) \in I$. Now I is M-closed, so for all $a \in P$, if $\mathrm{tp}(a|M)$ is I-wide, then $R_1(a) \in I$. Hence by (2) again we get $(P \times X) \cap R_1 \in J$.

If there is an M-definable $Q \subseteq P$ with $Q \notin I$ such that $\breve{R}_2(b) \notin I$ for all $b \in Q$, then the second option of (*) holds. We now show that there are no other possibilities: Assuming there is no such Q we obtain $(X \times P) \cap R_2 \in J$ as before. The three sets we showed to be in J yield $R_1 \cup R_2 \cup R_3 \in J$. Then from $P^2 \subseteq R_1 \cup R_2 \cup R_3$ we get $P^2 \in J$, so $P \in I$ by (1), a contradiction. This finishes the proof of Claim 1.

Claim 2. There is an I-wide $p \in \mathrm{St}(X|M)$ such that for all M-definable $R_1, R_2, R_3 \subseteq X^2$, if $p(X)^2 \subseteq R_1 \cup R_2 \cup R_3$, then there exists $Q \in p$ such that (*) holds for all $a, b \in Q$.

To prove this we use that L and M are countable. Take an enumeration (X_n) of $\mathrm{Def}(X|M)$ and an enumeration (R_{n1}, R_{n2}, R_{n3}) of $\mathrm{Def}(X^2|M)^3$ in which every $(R_1, R_2, R_3) \in \mathrm{Def}(X^2|M)^3$ occurs infinitely often. Choose recursively a descending sequence $P_0 \supseteq P_1 \supseteq P_2 \cdots$ in $\mathrm{Def}(X|M) \setminus I$ such that

(a) $P_{2n} \subseteq X_n$ or $P_{2n} \subseteq \neg X_n$;
(b) if $P_{2n}^2 \subseteq R_{n1} \cup R_{n2} \cup R_{n3}$, take Q as in Claim 1 for P_{2n} and R_{n1}, R_{n2}, R_{n3} in the role of P, R_1, R_2, R_3, and set $P_{2n+1} := Q$.

It is easy to check that $p = \{P \in \mathrm{Def}(X|M) : \ P \supseteq P_n \text{ for some } n\}$ satisfies Claim 2.

Let p be as in Claim 2, and take $a \in p(X)$. Let $\Gamma(a)$ be the collection of definable subsets of X consisting of the sets in p together with the following sets for all M-definable $R \subseteq X^2$:

(i) $\neg R(a)$ if $R(a) \in I$;
(ii) $\neg R(a)$ if $\breve{R}(a) \in I$;
(iii) $\neg R(a)$ if $\breve{R}(c) \notin p$ for all $c \in X(M)$.

Note: $\breve{R}(c) \notin p$ for all $c \in X(M)$ if and only if $X(M) \subseteq \neg R(a)$.

Claim 3. Let $X_1, \ldots, X_n \in \Gamma(a)$. Then $X_1 \cap \cdots \cap X_n \neq \emptyset$.

To prove this, take $P \in p$ and M-definable $R_1, R_2, R_3 \subseteq X^2$ such that

$$X_1 \cap \cdots \cap X_n \supseteq P \cap \neg R_1(a) \cap \neg R_2(a) \cap \neg R_3(a),$$

with $R_1(a), \breve{R}_2(a) \in I$, and $\breve{R}_3(c) \notin p$ for all $c \in X(M)$. Suppose towards a contradiction that $X_1 \cap \cdots \cap X_n = \emptyset$. Then

$$P \subseteq R_1(a) \cup R_2(a) \cup R_3(a),$$

so $p(X) \times P \subseteq R_1 \cup R_2 \cup R_3$ and thus $p(X)^2 \subseteq R_1 \cup R_2 \cup R_3$. Then by Claim 2 we have $R_1(a) \notin I$ or $\breve{R}_2(a) \notin I$ or $\breve{R}_3(c) \in p$ for some $c \in X(M)$, a contradiction. This proves Claim 3.

Claim 4. Let $X_1, \ldots, X_n \in \Gamma(a)$. Then $X_1 \cap \cdots \cap X_n \cap X(M) \neq \emptyset$.

Suppose towards a contradiction that $X_1 \cap \cdots \cap X_n \cap X(M) = \emptyset$. As in the proof of Claim 3 we get $P \in p$ and M-definable $R_1, R_2, R_3 \subseteq X^2$ such that $R_1(a), \breve{R}_2(a) \in I$, and $\breve{R}_3(c) \notin p$ for all $c \in X(M)$, and

$$P(M) \subseteq R_1(a) \cup R_2(a) \cup R_3(a),$$

that is, $X(M) \subseteq \neg R(a)$ for $R \subseteq X^2$ defined by

$$R(b,c) \iff P(c) \text{ and } \neg R_1(b,c) \text{ and } \neg R_2(b,c) \text{ and } \neg R_3(b,c).$$

Thus by (iii) we have $\neg R(a) \in \Gamma(a)$, that is

$$\neg P \cup R_1(a) \cup R_2(a) \cup R_3(a) \in \Gamma(a),$$

contradicting $P, \neg R_1(a), \neg R_2(a), \neg R_3(a) \in \Gamma(a)$ in view of Claim 3.

Thus $\Gamma(a)$ extends by Claim 4 and Lemma 1.1 to a global type $\mathbf{p} \in \mathrm{St}(X)$ that is finitely satisfiable in M. Take $b \models \mathbf{p}{\restriction}Ma$. It is easy to check that $\mathrm{tp}_X(b|Ma)$ is I-wide because of (i).

It remains to show that $\mathrm{tp}_X(a|Mb)$ is I-wide. Let $P \in \mathrm{tp}_X(a|Mb)$. Then $a \in P = R(b)$ where $R \subseteq X^2$ is M-definable. Towards a contradiction, suppose that $P \in I$, that is, $R(b) \in I$. Now a and b both realize $p = \mathbf{p}{\restriction}M$, so they are M-conjugate, and thus $R(a) \in I$. Hence $\neg\breve{R}(a) \in \mathbf{p}{\restriction}Ma$ by (ii), so $b \in \neg\breve{R}(a)$, contradicting $a \in R(b)$. □

Lemma 1.24. *Suppose L and A are countable, and I is proper and A-invariant. Then there is a countable $M \supseteq A$ and a global type $\mathbf{p} \in \mathrm{St}(X)$, finitely satisfiable in M, such that for all $a, b \in X$, if $a \models \mathbf{p}{\restriction}M$, $b \models \mathbf{p}{\restriction}Ma$, then $\mathrm{tp}_X(a|Mb)$ is I-wide.*

Proof. Assume that $\mathrm{Th}(\mathbb{U})$ has definable Skolem functions and $\kappa(\mathbb{U}) > \beth_{\omega_1}$. (After proving the lemma for this case we shall reduce the general case to this special case.) For any B, let $\langle B \rangle$ be the substructure of $\mathbb{U}$ generated by B; by definability of Skolem functions, this is a (small) elementary submodel of $\mathbb{U}$. Since I is proper, there is for each M an I-wide $p \in \mathrm{St}(X|M)$. Thus by transfinite recursion we obtain a sequence $(a_i : i < \beth_{\omega_1})$ in X such that for each i, $\mathrm{tp}_X(a_i|\langle A \cup \{a_j : j < i\}\rangle)$ is I-wide. A theorem of Morley then yields an A-indiscernible sequence $(c_i)_{i<\omega+2}$ in X such that for any n, $\mathrm{tp}((c_0, \ldots, c_n)|A) = \mathrm{tp}((a_{i_0}, \ldots, a_{i_n})|A)$ for suitable $i_0 < \cdots < i_n < \beth_{\omega_1}$. In particular, $\mathrm{tp}_X(c_n|Ac_0 \ldots c_{n-1})$ is I-wide for all n.

Let F be a non-principal ultrafilter on $\mathbb{N} = \omega$ and put

$$\mathbf{p} := \{P \in \mathrm{Def}(X) : \ \{n : c_n \in P\} \in F\}.$$

Then $\mathbf{p} \in \mathrm{St}(X)$, and $\mathbf{p}$ is finitely satisfiable in $M := \langle A \cup \{c_i : i < \omega\}\rangle$. Put $a := c_{\omega+1}$ and $b := c_\omega$. Then $a \models \mathbf{p}{\restriction}M$, $b \models \mathbf{p}{\restriction}Ma$, and $\mathrm{tp}_X(a|Mb)$ is I-wide. This finishes the proof under the assumptions introduced earlier.

To reduce the general case to this special case, expand $\mathbb{U}$ to $\mathbb{U}_{\mathrm{sk}}$ while keeping the language countable and without introducing new sorts, such that $\mathrm{Th}(\mathbb{U}_{\mathrm{sk}})$ has definable Skolem functions. We cannot expect $\mathbb{U}_{\mathrm{sk}}$ to be big, so we take a big elementary extension $\mathbb{U}'$ of $\mathbb{U}_{\mathrm{sk}}$ with $\kappa(\mathbb{U}') > \beth_{\omega_1}$. If an $L(U)$-formula $\phi(x)$ defines $P \subseteq \mathbb{U}_x$ in $\mathbb{U}$, we denote the subset of $\mathbb{U}'_x$ that it defines in $\mathbb{U}'$ by P'. The A-open set $X \subseteq \mathbb{U}_x$ is likewise extended to an A-open set $X' \subseteq \mathbb{U}'_x$ (with respect to the ambient $\mathbb{U}'$) as follows:

$$\text{if } X = \bigcup_{j \in J} X_j, \quad \text{all } X_j \in \mathrm{Def}(\mathbb{U}_x|A), \quad \text{then } X' := \bigcup_{j \in J} X'_j.$$

Finally, let I' be the collection of all sets $Q \subseteq (\mathbb{U}'_x$ that are definable in $\mathbb{U}'$ and contained in P' for some $P \in I$. Then I' is a proper A-invariant ideal of $\mathrm{Def}(X')$ (with respect to the ambient $\mathbb{U}'$). It remains to apply the result of the special case with $\mathbb{U}'$, X', I' in the role of $\mathbb{U}$, X, I, and restrict suitably. □

2. The stabilizer

This section has subsections *Generic sets*, *The stabilizer theorem*, *More on the stabilizer theorem*, *Making measures definable*, and *An application.* In the last subsection *An application* we derive a result on subsets of groups in the spirit of the sum-product phenomenon.

Throughout G is a (multiplicatively written) group with a distinguished subset X such that $1_G \in X$. Put $X_1 := X \cup X^{-1}$, so $X \subseteq X_1$ and $X_1 = X_1^{-1}$. Let X_n be the set of products $g_1 \cdots g_n$ with $g_1, \dots, g_n \in X_1$, so $X_n \subseteq X_{n+1}$, and $\widehat{X} := \bigcup_n X_n$ is the subgroup of G generated by X. The set

$$X^* := XX^{-1}X$$

will play a special role. For $Y \subseteq G$ and subgroup H of G we put

$$Y/H := \{yH : y \in Y\} \subseteq G/H.$$

Generic sets. We begin with two general and elementary facts.

Lemma 2.1. *Let $Y \subseteq X_m$ and $k \in \mathbb{N}$ be such that X_{m+1} is covered by k left translates of Y. Then X_{m+n} is covered by k^{n+1} left translates of Y.*

Proof. By induction on n. Let $X_{m+n} \subseteq \bigcup_{e \in E} eY$ with $E \subseteq G$. Then

$$X_{m+n+1} = X_{m+n}X_1 \subseteq \bigcup_{e \in E} eYX_1.$$

It remains to use that $YX_1 \subseteq X_{m+1}$. □

Lemma 2.1 goes through with right translates instead of left translates.

Lemma 2.2. *Let $Y, Z \subseteq G$ and let E be a maximal subset of Z such that $eY \cap fY = \emptyset$ for all distinct $e, f \in E$. Then*

$$Z \subseteq \bigcup_{e \in E} eYY^{-1}.$$

Proof. Given $g \in Z$, the maximality of E gives $e \in E$ such that $gY \cap eY \neq \emptyset$, so $ga = eb$ with $a, b \in Y$, and thus $g \in eYY^{-1}$. □

In the rest of this subsection we work in $\mathbb{U}$ as before, and assume:

- $\kappa(\mathbb{U}) > 2^{\aleph_0}$ and x is a *finite*[3] multivariable.
- $G \subseteq \mathbb{U}_x$ is an A-definable group, that is, G is an A-definable subset of $\mathbb{U}_x$ equipped with an A-definable group operation.
- $X \subseteq G$ is A-definable. (Recall also that $1_G \in X$.)

Note that then each X_n is A-definable, so $\widehat{X}$ is A-open.

Suppose I is an ideal of $\mathrm{Def}(\widehat{X})$. We define I to be *left-invariant* if $gP \in I$ for all $P \in I$ and $g \in X_1$ (so $gP \in I$ for all $P \in I$ and $g \in \widehat{X}$). Likewise I is *right-invariant* if $Pg \in I$ for all $P \in I$ and $g \in X_1$. We set

$$I^{-1} := \{P^{-1} : \ P \in I\}.$$

Then I^{-1} is an ideal of $\mathrm{Def}(\widehat{X})$, $(I^{-1})^{-1} = I$, and I is left-invariant iff I^{-1} is right-invariant. For $Y \in \mathrm{Def}(\widehat{X})$ we define $I|Y := \{P \in I : \ P \subseteq Y\}$, which is also an ideal of $\mathrm{Def}(\widehat{X})$.

Lemma 2.3. *Let I be an ideal of $\mathrm{Def}(\widehat{X})$ such that $I|X_4$ is S1 over A and $aX \notin I$ for all $a \in X_3$. Then every set X_n is covered by finitely many left translates of XX^{-1}.*

Proof. Take a maximal $E \subseteq X_3$ such that $eX \cap fX = \emptyset$ for all distinct $e, f \in E$. Then E is finite by the S1 assumption. Given $a \in X_3$ we get $e \in E$ such that $aX \cap eX \neq \emptyset$, and so $a \in eXX^{-1}$. Thus X_3 is covered by finitely many left translates of $XX^{-1} \subseteq X_2$. It remains to apply Lemma 2.1. □

A *left-generic set* is a set $Y \in \mathrm{Def}(\widehat{X})$ such that every $Z \in \mathrm{Def}(\widehat{X})$ is covered by finitely many left translates gY with $g \in \widehat{X}$, equivalently, every set X_n is covered by finitely many such translates. The notion of *right-generic set* is defined accordingly. Note that a left-generic set cannot belong to any left-invariant proper ideal of $\mathrm{Def}(\widehat{X})$. Here is a partial converse:

Corollary 2.4. *Suppose I is a left-invariant ideal of $\mathrm{Def}(\widehat{X})$ that is S1 over A, and let $Y \in \mathrm{Def}(\widehat{X})$, $Y \notin I$. Then YY^{-1} is left-generic.*

Proof. It suffices to show that each X_n is covered by finitely many translates eYY^{-1} with $e \in X_n$. Let E be a maximal subset of X_n such that $eY \cap fY = \emptyset$ for all distinct $e, f \in E$. Since $eY \notin I$ for all $e \in E$, the assumption that I is S1 over A yields that E is finite. Now apply Lemma 2.2. □

Recall from [2] that a subset of $\mathbb{U}_x$ is said to be *countably definable* if it is a countable intersection of definable subsets of $\mathbb{U}_x$. Suppose H is a countably definable subgroup of G. Then by Lemma 4.5 of [2] we have $H = \bigcap_n H_n$ for a decreasing sequence (H_n) of definable subsets of G such that for all n,

$$H_n^{-1} = H_n, \qquad H_n \supseteq H_{n+1}H_{n+1}.$$

Note that if $H \subseteq \widehat{X}$, then $H \subseteq X_m$ for some m (by compactness) and for such m we have $H_n \subseteq X_m$ for all sufficiently large n, again by compactness.

[3]This is because we use a fact from [2] established there under this restriction.

Lemma 2.5. *Let H be a countably definable subgroup of $\widehat{X}$. Then:*

(1) $|\widehat{X}/H| \le 2^{\aleph_0}$ *iff every* $Y \in \mathrm{Def}(\widehat{X})$ *containing H is left-generic.*

(2) *Suppose I is a left-invariant proper ideal of* $\mathrm{Def}(\widehat{X})$ *and I is S1 over A. Then* $|\widehat{X}/H| \le 2^{\aleph_0}$ *iff H is I-wide.*

Proof. Take a sequence (H_n) as above with $H_n \subseteq \widehat{X}$ for all n. Let m and $n > 0$ be such that $H_n \subseteq X_m$, and let E be a maximal subset of X_m such that $eH_n \cap fH_n = \emptyset$ for all distinct $e, f \in E$. By Lemma 2.2 we have $X_m \subseteq \bigcup_{e\in E} eH_{n-1}$. If E is infinite, then by saturation $|E| > 2^{\aleph_0}$.

After these preliminary remarks, we first prove (1). If $|\widehat{X}/H| \le 2^{\aleph_0}$, then by these remarks any E as above is finite, and so by increasing m and n we see that all H_n are left-generic, and thus every $Y \in \mathrm{Def}(\widehat{X})$ containing H is left-generic.

Next, suppose that $|\widehat{X}/H| > 2^{\aleph_0}$. Take $F \subseteq \widehat{X}$ with $|F| > 2^{\aleph_0}$ such that $eH \ne fH$ for all distinct $e, f \in F$. We can arrange that for a certain m we have $F \subseteq X_m$. For any distinct $e, f \in F$ we have $n = n(e, f) \in \mathbb{N}$ such that $eH_n \cap fH_n = \emptyset$. By Erdös-Rado, we get an infinite $F' \subseteq F$ such that $n(e, f)$ takes a constant value n for distinct $e, f \in F'$. Increasing m if necessary we can assume that $H_n \subseteq X_m$. Suppose towards a contradiction that H_n is left-generic. Then $F' \subseteq X_m \subseteq e_1H_n \cup \cdots \cup e_kH_n$ with $k \in \mathbb{N}$ and $e_1, \dots, e_k \in \widehat{X}$, so we get $i \in \{1, \dots, k\}$ and distinct $e, f \in F'$ such that $e = e_i g$ and $f = e_i h$ with $g, h \in H_n$, and thus $eg^{-1} = fh^{-1} \in eH_n \cap fH_n$, a contradiction. This finishes the proof of (1).

We now prove (2). If $H \subseteq P \in I$, then $H_n \subseteq P$ for some n, and such H_n is not left-generic, so $|\widehat{X}/H| > 2^{\aleph_0}$ by (1). Conversely, assume $|\widehat{X}/H| > 2^{\aleph_0}$. Then (1) gives n such that H_n is not left-generic, so $H_{n+1}^{-1}H_{n+1}$ is not left-generic, and thus $H \subseteq H_{n+1} \in I$ by Corollary 2.4. □

With an eye towards applying this lemma we note that if L and A are countable, then every A-closed subgroup of G is countably definable.

Assuming that an ideal I of $\mathrm{Def}(\widehat{X})$ is both left-invariant and S1 over some A may be too strong for some applications. In such cases it might be enough that the restriction of the ideal to $X^* := XX^{-1}X$ satisfies S1. To be precise, let I be an ideal of $\mathrm{Def}(\widehat{X})$. Put

$$I^* \;:=\; I|X^* \;=\; \{P \in I :\; P \subseteq X^*\}.$$

Then I^* is an ideal of $\mathrm{Def}(\widehat{X})$, and if I is A-invariant, so is I^*. If $P \subseteq X^*$, then P is I-wide iff P is I^*-wide. If I is S1 over A, so is I^* (and being S1 is more realistic for I^* than for I). Note: $X, X^{-1} \subseteq X^{-1}X \subseteq X^* \subseteq X_3$.

The stabilizer theorem. The next theorem is a key result. It constructs a useful type-definable group from rather generic data. We follow Hrushovski in denoting this group by S. (This abuses our notation since S also names the set of sorts of our language L, but the context will prevent confusion.) The reason for using the letter S is that in more special settings this group is a *stabilizer*. (Maybe also in our setting for some natural action?) In this subsection we keep the assumptions on $\mathbb{U}, G, X$ from the previous subsection. Note that the theorem below refers to a (small) base model $M \supseteq A$, and that its hypotheses include M-invariance of I^* but not of I.

Theorem 2.6. *Let I be a left-and-right-invariant ideal of $\mathrm{Def}(\widehat{X})$ such that I^* is S1 over $M \supseteq A$. Let $q \in \mathrm{St}(X|M)$ be I-wide and assume there exist $a, b \in [q] := q(X)$ such that neither $\mathrm{tp}(b|Ma)$ nor $\mathrm{tp}(a|Mb)$ forks over M.*

Then $S := [q]^{-1}[q][q]^{-1}[q] \subseteq X^{-1}XX^{-1}X$ has the following properties:

(1) *S is an M-closed subgroup of $\widehat{X}$ and S is I-wide;*
(2) *$aS = [q][q]^{-1}[q]$ for all $a \in [q]$;*
(3) *there is no M-closed proper subgroup T of S such that S/T is small;*
(4) *$S \subseteq (X^{-1}X) \cup P$ for some M-definable $P \in I$.*

Proof. For sets $Y \subseteq \mathbb{U}_y$ and $Z \subseteq \mathbb{U}_z$ we put

$$Y \times_{\mathrm{nf}} Z := \{(a, b) \in Y \times Z : \ \mathrm{tp}(b|Ma) \text{ does not fork over } M\}.$$

Next we introduce M-invariant subsets Q and Q' of $X^{-1}X \subseteq X^*$:

$$Q := \{a^{-1}b : \ (a, b) \in [q] \times_{\mathrm{nf}} [q]\},$$
$$Q' := \{a^{-1}b : \ a, b \in [q], \ \mathrm{tp}_X(b|Ma) \text{ is } I\text{-wide}\}.$$

We have $Q \supseteq Q'$ by Lemma 1.18, and $[q][q]^{-1}$ is I-wide by right-invariance, and $[q]^{-1}[q]$ is I-wide by left-invariance. Note also that we have a type $q^{-1} \in \mathrm{St}(X^{-1}|M)$ such that $[q]^{-1} = q^{-1}(X^{-1})$. Throughout the proof we use that by Lemma 1.19 the relation $R \subseteq X_4 \times X_4$ given by

$$R(a, b) \iff [q]a^{-1} \cap [q]b^{-1} \text{ is } I^*\text{-wide},$$

is stable over M. (Note: any subset of $\widehat{X}$ not contained in X^* is I^*-wide.) By Lemma 1.14 and $\breve{R} = R$, given any types $p, p' \in \mathrm{St}(X_4|M)$, the following are equivalent:

(i) $R(a, b)$ for some $(a, b) \in p(X_4) \times_{\mathrm{nf}} p'(X_4)$;
(ii) $R(a, b)$ for all such (a, b);
(iii) $R(a, b)$ for some $(a, b) \in p'(X_4) \times_{\mathrm{nf}} p(X_4)$;
(iv) $R(a, b)$ for all such (a, b).

It follows in particular that if $Y, Z \subseteq X_4$ are M-invariant and $R(a, b)$ for all $(a, b) \in Y \times_{\mathrm{nf}} Z$, then also $R(a, b)$ for all $(a, b) \in Z \times_{\mathrm{nf}} Y$.

Claim 1. $[q]^{-1}[q] \subseteq QQ$.

To prove this, let $a, b \in [q]$. The assumption on q gives $c \in [q]$ such that neither $\mathrm{tp}(a|Mc)$ nor $\mathrm{tp}(c|Ma)$ forks over M. Use Lemma 1.6 to extend $\mathrm{tp}(c|Ma)$ to a $p \in \mathrm{St}(\mathbb{U}_x|Mab)$ that doesn't fork over M; upon replacing c by an element of $p(\mathbb{U}_x)$ we arrange that $\mathrm{tp}(c|Mab)$ doesn't fork over M. Then $(b, c), (c, a) \in [q] \times_{\mathrm{nf}} [q]$, so $b^{-1}c, c^{-1}a \in Q$, and thus $b^{-1}a \in QQ$.

Claim 2. $R(a, b)$ for all $(a, b) \in [q] \times_{\mathrm{nf}} [q]$.

By the remarks on R preceding Claim 1 it is enough to show that $R(a,b)$ for some $(a,b) \in [q] \times_{\mathrm{nf}} [q]$. By the remark following the proof of Lemma 1.2 we can extend q to an M-invariant global type $\mathbf{q} \in \mathrm{St}(X)$. Take a $\mathbf{q}$-indiscernible sequence (a_n) over M. Then $(a_m, a_n) \in [q] \times_{\mathrm{nf}} [q]$ for all $m < n$. For $P \in q$ we have $P \notin I$, so $Pa_n^{-1} \notin I^*$ for all n; since I^* is S1 over M we get $m < n$ such that $Pa_m^{-1} \cap Pa_n^{-1} \notin I^*$, and thus $Pa_0^{-1} \cap Pa_1^{-1} \notin I^*$. This holds for all $P \in q$, so $[q]a_0^{-1} \cap [q]a_1^{-1}$ is I^*-wide, that is, $R(a_0, a_1)$.

Claim 3. $R(c, c')$ for all $(c, c') \in [q]^{-1}[q] \times_{\mathrm{nf}} Q'$.

Let $(c, c') \in [q]^{-1}[q] \times_{\mathrm{nf}} Q'$, and $p := \mathrm{tp}(c|M)$ and $p' := \mathrm{tp}(c'|M)$. As before it is enough to show that $R(d, d')$ for some $(d, d') \in p(X^*) \times_{\mathrm{nf}} p'(X^*)$. Let $a_0 \in [q]$, and take $a_1 \in [q]$ such that $\mathrm{tp}(a_0^{-1}a_1|M) = \mathrm{tp}(c|M) = p$. From $c' \in Q'$ we get $a_2' \in [q]$ such that $\mathrm{tp}(a_0^{-1}a_2'|M) = \mathrm{tp}(c'|M) = p'$ and $r := \mathrm{tp}_X(a_2'|Ma_0)$ is I-wide. Extend r to an I-wide $r' \in \mathrm{St}(X|Ma_0a_1)$ and take $a_2 \in r'(X) \subseteq q(X) = [q]$. Then $\mathrm{tp}_X(a_2|Ma_0) = r = \mathrm{tp}_X(a_2'|Ma_0)$, so

$$\mathrm{tp}(a_0^{-1}a_1|M) = p, \qquad \mathrm{tp}(a_0^{-1}a_2|M) = \mathrm{tp}(a_0^{-1}a_2'|M) = p'.$$

Since $\mathrm{tp}_X(a_2|Ma_0a_1) = r'$ is I^*-wide, it doesn't fork over M by Lemma 1.18, and so $\mathrm{tp}(a_2|Ma_1)$ doesn't fork over M. Then $[q]a_1^{-1} \cap [q]a_2^{-1}$ is I^*-wide by Claim 2, so $[q]a_1^{-1}a_0 \cap [q]a_2^{-1}a_0$ is I^*-wide by right-invariance, and thus $R(a_0^{-1}a_1, a_0^{-1}a_2)$.

Claim 4. $R(c, d)$ for all $(c, d) \in [q]^{-1}[q] \times_{\mathrm{nf}} Q$.

To prove this, let $(c, d) \in [q]^{-1}[q] \times_{\mathrm{nf}} Q$. Then $d = a^{-1}b$ where $a, b \in [q]$ and $\mathrm{tp}(b|Ma)$ doesn't fork over M. We wish to show $R(c, a^{-1}b)$, which by right-invariance is equivalent to $R(ac, b)$. Since $\mathrm{tp}_X(b|M) = q$ is I^*-wide, Lemma 1.20 provides $b' \in [q]$ such that $\mathrm{tp}_X(b'|M) = \mathrm{tp}_X(b|M)$ and $\mathrm{tp}_X(b'|Mac)$ is I^*-wide. Then it follows from Lemma 1.18 that $\mathrm{tp}(b'|M(ac))$ doesn't fork over M.[4] By the remarks preceding Claim 1 it suffices to show that $R(ac, b')$, equivalently, $R(c, a^{-1}b')$. Now $\mathrm{tp}_X(b'|Ma)$ is I-wide, so $a^{-1}b' \in Q'$. Left-invariance gives that $\mathrm{tp}_{X^*}(a^{-1}b'|Mac)$ is I^*-wide, so $\mathrm{tp}_{X^*}(a^{-1}b'|Mc)$ is I^*-wide, and thus $\mathrm{tp}(a^{-1}b'|Mc)$ doesn't fork over M. Now Claim 3 yields $R(c, a^{-1}b')$, as desired.

Claim 5. Let $(b, a) \in Q \times_{\mathrm{nf}} [q]^{-1}[q]$. Then $[q]a \cap [q]b^{-1}$ is I^*-wide and (thus) $ab \in [q]^{-1}[q]$.

To see this, note that $(b, a^{-1}) \in Q \times_{\mathrm{nf}} [q]^{-1}[q]$, and thus $R(b, a^{-1})$ by Claim 4 and the symmetry property of R mentioned just before Claim 1. Hence $[q]b^{-1} \cap [q]a$ is I^*-wide, and so we can take $c, d \in [q]$ with $cb^{-1} = da$, and thus $ab = d^{-1}c \in [q]^{-1}[q]$.

[4]Here we distinguish Mac, the parameter set obtained by adjoining a and c to M, from $M(ac)$, the parameter set obtained by adjoining the product ac to M.

Claim 6. Let $a \in [q]^{-1}[q]$ and $b_1, \ldots, b_n \in Q$, $n \geq 1$, and suppose that the type $\mathrm{tp}_{X^*}(a|Mb_1 \ldots b_n)$ is I-wide. Then $ab_1 \cdots b_n \in [q]^{-1}[q]$, and the set $[q]a \cap [q](b_1 \cdots b_n)^{-1}$ is I-wide.

We prove this by induction on n. Since $\mathrm{tp}_{X^*}(a|Mb_1)$ is I^*-wide, $\mathrm{tp}(a|Mb_1)$ doesn't fork over M by Lemma 1.18, and so by Claim 5,

$$ab_1 \in [q]^{-1}[q], \qquad [q]a \cap [q]b_1^{-1} \text{ is } I\text{-wide}.$$

This gives the case $n = 1$. Let $n > 1$. Since $\mathrm{tp}_{X^*}(a|Mb_1 \ldots b_n)$ is I-wide and $ab_1 \in X^*$, right-invariance yields that $\mathrm{tp}_{X^*}(ab_1|Mb_1 \ldots b_n)$ is I-wide, and therefore $\mathrm{tp}_{X^*}(ab_1|Mb_2 \ldots b_n)$ is I-wide. Then by the inductive assumption,

$$ab_1(b_2 \cdots b_n) \in [q]^{-1}[q] \text{ and } [q]ab_1 \cap [q](b_2 \cdots b_n)^{-1} \text{ is } I\text{-wide}.$$

Then $[q]a \cap [q](b_1 \cdots b_n)^{-1}$ is I-wide by right-invariance.

To get property (2) we also need a variant of Claim 6:

Claim 7. Let $a \in [q]^{-1}[q]$ and $b_1, \ldots, b_n \in Q$, $n \geq 1$, and suppose that the type $\mathrm{tp}_{X^*}(a^{-1}|Mb_1 \ldots b_n)$ is I-wide. Then $ab_1 \cdots b_n \in [q]^{-1}[q]$, and the set $[q]a \cap [q](b_1 \cdots b_n)^{-1}$ is I-wide.

As in the proof of Claim 6 we obtain that $\mathrm{tp}(a^{-1}|Mb_1)$ and thus $\mathrm{tp}(a|Mb_1)$ doesn't fork over M, and so by Claim 5,

$$ab_1 \in [q]^{-1}[q], \qquad [q]a \cap [q]b_1^{-1} \text{ is } I\text{-wide}.$$

For $n > 1$ we note that $\mathrm{tp}_{X^*}\big((ab_1)^{-1}|Mb_1 \ldots b_n\big)$ is I-wide by left-invariance. Now proceed inductively as in the proof of Claim 6.

Claim 8. Let $n \geq 1$. Then $\{b_1 \cdots b_n : \ b_1, \ldots, b_n \in Q\} \subseteq [q]^{-1}[q][q]^{-1}[q]$.

Let $b_1, \ldots, b_n \in Q$. Now $[q]^{-1}[q] \subseteq X^*$ is I-wide, so Lemma 1.20 provides an I-wide $p \in \mathrm{St}(X^*|Mb_1 \ldots b_n)$ with $p(X^*) \subseteq [q]^{-1}[q]$. Take any $a \in p(X^*)$. Then $a \in [q]^{-1}[q]$ and $\mathrm{tp}_{X^*}(a|Mb_1 \ldots b_n)$ is I-wide, hence $ab_1 \cdots b_n \in [q]^{-1}[q]$ by Claim 6, so

$$b_1 \cdots b_n = a^{-1}(ab_1 \cdots b_n) \in [q]^{-1}[q][q]^{-1}[q].$$

Claims 1 and 8 yield that $S := [q]^{-1}[q][q]^{-1}[q]$ is indeed a subgroup of $\widehat{X}$. Clearly S is M-closed, $S \subseteq X_4 \subseteq \widehat{X}$, and $S \supseteq [q]^{-1}[q]$, so S is I-wide. Note:

$$S \subseteq X^{-1}XX^{-1}X = X^{-1}X^*.$$

Claim 9. $aS = [q][q]^{-1}[q]$ for all $a \in [q]$. To prove this, let $a \in [q]$, and note that $[q]^{-1}[q] \subseteq S$ gives $[q] \subseteq aS$, so $[q][q]^{-1}[q] \subseteq aS$. For the reverse inclusion, let $b \in aS$, so $b = ab_1b_2b_3b_4$ with $b_1, b_2, b_3, b_4 \in Q$ by Claim 1. Take $d \in [q]$ such that $\mathrm{tp}_{X^*}(d|Mab_1 \ldots b_4)$ is I-wide, and set $e := d^{-1}a$. Then $a = de$ and $\mathrm{tp}_{X^*}(e^{-1}|Mab_1 \ldots b_4)$ is I-wide by left-invariance, so $\mathrm{tp}_{X^*}(e^{-1}|Mb_1 \ldots b_4)$ is I-wide. Now $e \in [q]^{-1}[q]$, hence $eb_1 \cdots b_4 \in [q]^{-1}[q]$ by Claim 7, so

$$b = ab_1 \cdots b_4 = d(eb_1 \cdots b_4) \in [q][q]^{-1}[q].$$

Claim 10. Let T be an M-closed subgroup of S such that $|S/T| < \kappa(\mathbb{U})$. Then $S = T$.

In the proof of Claim 9 we observed that $[q]$ is contained in a single left coset of S in $\widehat{X}$. This left coset equals $[q]S$, and is therefore M-closed. We have T acting on $[q]S$ by multiplication on the right, and this action has only a small number of orbits aT with $a \in [q]S$. Then by Lemma 1.17 we have $[q] \subseteq aT$ where $a \in [q]$, and thus $[q]^{-1}[q] \subseteq T$, which gives $S = T$.

Claim 11. There exists $c \in Q$ such that $\mathrm{tp}_{X^*}(c|M)$ is I-wide.

To prove this claim, take $a \in [q]$ and use Lemma 1.20 to extend q to an I-wide type $p \in \mathrm{St}(X|Ma)$. Next, take $b \in p(X)$, so $\mathrm{tp}_X(b|Ma) = p$, and $c := a^{-1}b \in Q' \subseteq Q$. Then $\mathrm{tp}_{X^*}(c|Ma)$ is I-wide by left-invariance, and so $\mathrm{tp}_{X^*}(c|M)$ is I-wide.

Claim 12. $S \setminus [q]^{-1}[q]$ is contained in a union of M-definable sets in I.

To prove this, note first that the set $S \setminus [q]^{-1}[q] \subseteq Y := X^{-1}X^*$ is M-invariant, and thus a union of sets $r(Y)$ with $r \in \mathrm{St}(Y|M)$. So let any I-wide $r \in \mathrm{St}(Y|M)$ with $r(Y) \subseteq S$ be given; it suffices to show that then

$$r(Y) \subseteq [q]^{-1}[q] \quad \text{(equivalently, } r(Y) \cap [q]^{-1}[q] \neq \emptyset\text{)}.$$

Pick $s_0 \in r(Y)$, so $s_0^{-1} \in S \subseteq QQQQ$, hence $s_0^{-1} = b_1b_2b_3b_4$ with the $b_i \in Q$. Claim 11 and Lemma 1.20 yield $c \in Q$ such that $\mathrm{tp}_{X^*}(c|Mb_1b_2b_3b_4)$ is I-wide. Hence $\mathrm{tp}(c|Mb_1b_2b_3b_4)$ doesn't fork over M, and so $\mathrm{tp}(c|Ms_0)$ doesn't fork over M. Since $c \in [q]^{-1}[q]$, Claim 6 gives that $[q]c \cap [q]s_0$ is I-wide.

Pick $a \in [q]$. Since r is I-wide, Lemma 1.20 yields $s \in r(Y) \subseteq S$ such that $\mathrm{tp}_Y(s|Mac)$ is I-wide. Then $as \in X^*$ by Claim 9, and $\mathrm{tp}_{X^*}(as|Mac)$ is I-wide by left-invariance, hence $\mathrm{tp}(as|Mac)$ doesn't fork over M by Lemma 1.18, so $\mathrm{tp}(s|Mac)$ doesn't fork over M, and thus $\mathrm{tp}(s|Mc)$ doesn't fork over M. Therefore $[q]c \cap [q]s$ is I-wide as well, by remarks on R preceding Claim 1, and thus $sc^{-1} \in [q]^{-1}[q]$. Now $\mathrm{tp}_Y(s|Mc)$ is I-wide, so $\mathrm{tp}_{X^*}(sc^{-1}|Mc)$ is I-wide by right-invariance, so $\mathrm{tp}(sc^{-1}|Mc)$ doesn't fork over M. Hence $(c, sc^{-1}) \in Q \times_{\mathrm{nf}} [q]^{-1}[q]$, and so by Claim 5 we have $s = (sc^{-1})c \in [q]^{-1}[q]$, and thus $r(Y) \subseteq [q]^{-1}[q]$.

Claim 13. $S \subseteq (X^{-1}X) \cup P$ for some M-definable $P \in I$: "almost all" elements of S are in $X^{-1}X$.

This is because by Claim 12 we have

$$S \;\subseteq\; [q]^{-1}[q] \cup \bigcup_{\lambda\in\Lambda} P_\lambda \;\subseteq\; (X^{-1}X) \cup \bigcup_{\lambda\in\Lambda} P_\lambda$$

where all P_λ are M-definable and in I. Now use compactness of the M-topology to get a single P. $\square$

Given $M \supseteq A$, the hypotheses on (X, I, q) in Theorem 2.6 are inherited by (X^{-1}, I^{-1}, q^{-1}), and so we have an I^{-1}-wide M-closed subgroup

$$S' := [q][q]^{-1}[q][q]^{-1}$$

of $\widehat{X}$. Under a mild countability assumption (to be eliminated when time permits) we have $S = S'$ and S is a normal subgroup of $\widehat{X}$:

Corollary 2.7. *Assume X, I, q are as in Theorem 2.6, and L, M are countable. Then $S = S'$, $S \trianglelefteq \widehat{X}$, and $|\widehat{X}/S| \leq 2^{\aleph_0}$.*

Proof. We proceed by establishing three more claims.

Claim 14. $|X/S| \leq 2^{\aleph_0}$. Suppose otherwise. Then we have $E \subseteq X$ with $|E| > 2^{\aleph_0}$ such that $eS \neq fS$ for all distinct $e, f \in E$. Take $b \in [q]$. Then $b^{-1}[q] \subseteq S$, so $eb^{-1}[q] \cap fb^{-1}[q] = \emptyset$ for all distinct $e, f \in E$. By Erdös-Rado used as in the proof of Lemma 2.5 we get infinite $F \subseteq E$ and a definable $Y \subseteq X$ with $[q] \subseteq Y$ such that $eb^{-1}Y \cap fb^{-1}Y = \emptyset$ for all distinct $e, f \in F$. Since $Y \notin I$, this contradicts I^* being S1 over M.

In the rest of the proof we assume $a \in X$.

Claim 15. aSa^{-1} is M-closed. To prove this, let $r := \mathrm{tp}(a|M)$, so that $r(\mathbb{U}_x)$ is contained in the union of the small number of left cosets of S in G that meet X. Then it follows from Lemma 1.17 that $r(\mathbb{U}_x)$ is contained in a single such coset, which must equal aS and must be M-closed. Hence $S^r := aS \cdot S \cdot (aS)^{-1} = aSa^{-1}$ is M-closed.

Claim 16. $|X^{-1}a/S| \leq 2^{\aleph_0}$. Suppose otherwise. Then we have $E \subseteq X$ with $|E| > 2^{\aleph_0}$ such that $e^{-1}aS \neq f^{-1}aS$ for all distinct $e, f \in E$. Take $b \in [q]$. Then $[q]^{-1}b \subseteq S$, so for all distinct $e, f \in E$ we have

$$e^{-1}a[q]^{-1}b \cap f^{-1}a[q]^{-1}b = \emptyset, \quad \text{and thus } \ e^{-1}a[q]^{-1} \cap f^{-1}a[q]^{-1} = \emptyset.$$

By Erdös-Rado this yields an infinite $F \subseteq E$ and a definable $Y \subseteq X$ with $[q] \subseteq Y$ such that for all distinct $e, f \in F$ we have $e^{-1}aY^{-1} \cap f^{-1}aY^{-1} = \emptyset$, and thus $Ya^{-1}e \cap Ya^{-1}f = \emptyset$. Since $Y \notin I$ and $Y \subseteq X$, this contradicts I^* being S1 over M.

It follows that $|[q]^{-1}a/S| \leq 2^{\aleph_0}$, and so

$$\begin{aligned} |[q]^{-1}/aSa^{-1}| &= |a[q]^{-1}/aSa^{-1}| = |a([q]^{-1}a)a^{-1}/aSa^{-1}| \\ &= |[q]a/S| \leq 2^{\aleph_0}. \end{aligned}$$

But aSa^{-1} is M-closed, so $[q]^{-1}$ is contained in a single left coset of aSa^{-1}, hence $[q][q]^{-1} \subseteq aSa^{-1}$, and thus $S' \subseteq aSa^{-1}$, in particular, for $a = 1$ this gives $S' \subseteq S$, and then by symmetry, $S = S'$, and thus $S \subseteq aSa^{-1}$. Also by symmetry, $S \subseteq a^{-1}S'a$, so $S \subseteq a^{-1}Sa$, and thus $S = aSa^{-1}$. This holds for all $a \in X$, so S is normal in $\widehat{X}$. Claim 14, together with Claim 16 for $a = 1$, then gives $|\widehat{X}/S| \leq 2^{\aleph_0}$. □

Theorem 2.6 and Corollary 2.7 together are the *stabilizer theorem*.

More on the stabilizer theorem. We keep the assumptions on $\mathbb{U}, G, X$ from previous subsections. The next result is useful in realizing some hypotheses in the stabilizer theorem. In particular, Lemmas 1.22, 1.23, 1.24 help in finding a type q as required in Theorem 2.6:

Lemma 2.8. *Let I be an ideal of* $\mathrm{Def}(\widehat{X})$*.*

1. *Suppose I is M-open, $X \notin I$, and I^* is S1 over M. Then there is an I-wide $q \in \mathrm{St}(X|M)$ with $a, b \in q(X)$ such that $\mathrm{tp}(a|Mb)$ and $\mathrm{tp}(b|Ma)$ don't fork over M.*
2. *Suppose L and A are countable, $X \notin I$, and I^* is S1 over A. Then there exists a countable $M \supseteq A$ and an I-wide $q \in \mathrm{St}(X|M)$ with $a, b \in q(X)$ such that $\mathrm{tp}(a|Mb)$ and $\mathrm{tp}(b|Ma)$ don't fork over M.*

Proof. Let I be as in (1). By restricting I to X, Lemma 1.22 gives a global type $\mathbf{q} \in \mathrm{St}(X)$, finitely satisfiable in M, such that $\mathrm{tp}_X(a|Mb)$ is I-wide whenever $a, b \in X$ satisfy $a \models \mathbf{q} \restriction M$ and $b \models \mathbf{q} \restriction Ma$. Take such $a, b \in X$, and put $q := \mathbf{q} \restriction M$. Then $q \subseteq \mathrm{tp}_X(a|Mb)$, so q is I-wide. Using the S1-assumption it follows from Lemma 1.18 that $\mathrm{tp}(a|Mb)$ doesn't fork over M. Since $\mathrm{tp}(b|Ma) \subseteq \mathbf{q}$ and $\mathbf{q}$ is M-invariant, $\mathrm{tp}(b|Ma)$ doesn't fork over M.

Next, let I be as in (2). By restricting I to X, Lemma 1.24 gives a countable $M \supseteq A$ and a global type $\mathbf{q} \in \mathrm{St}(X)$, finitely satisfiable in M, such that $\mathrm{tp}_X(a|Mb)$ is I-wide whenever $a, b \in X$ satisfy $a \models \mathbf{p} \restriction M$, $b \models \mathbf{p} \restriction Ma$. Taking such $a, b \in X$ and setting $q := \mathbf{q} \restriction M$, we obtain as in the proof of (1) that q is I-wide and $\mathrm{tp}(a|Mb)$ and $\mathrm{tp}(b|Ma)$ don't fork over M. □

Next we show how the left-invariance and right-invariance conditions in the stabilizer theorem can be weakened.

Corollary 2.9. *Let $M \supseteq A$ be fixed, and let (X, I, q) be as in Theorem 2.6, except that the requirement of left-and-right invariance of I is replaced by:*

(left) *for all $P \in \mathrm{Def}(X_3)$ and $a \in X_1$: $P \in I \implies aP \in I$;*
(right) *for all $P \in \mathrm{Def}(X_3)$ and $a \in X_1$: $P \in I \implies Pa \in I$.*

Then $S := [q]^{-1}[q][q]^{-1}[q]$ satisfies the conclusions of Theorem 2.6. Also, the above assumptions on (X, I, q) are inherited by (X^{-1}, I^{-1}, q^{-1}), and so yield the subgroup $S' := [q][q]^{-1}[q][q]^{-1}$ of $\widehat{X}$. If L and M are countable, then the conclusions of Corollary 2.7 go through.

Proof. These (left) and (right) conditions can be weakened further, as shown in what follows, but this would complicate their statements. We first take a look at how right-invariance was used in proving Theorem 2.6, and show that weaker froms of right-invariance suffice.

Right-invariance. In the proof of Theorem 2.6 we can replace right-invariance of I by the following consequences of it:

(1) for all $P \in \mathrm{Def}(X)$ and $a \in X^{-1}X$: $\quad P \in I \Longleftrightarrow Pa \in I$.
(2) for all $P \in \mathrm{Def}(X^{-1}X)$ and $a \in X^{-1}X$: $\quad P \in I \Longrightarrow Pa \in I$.

Indeed, (1) can replace right-invariance as used in:

(i) showing $[q][q]^{-1}$ is I-wide;
(ii) the proofs of Claims 2, 3, 4;
(iii) proving Claim 6 (second use) and Claim 7 (corresponding use).

As an example we show how to do this in proving Claim 4, where we have $a, b \in X$ and $c \in X^{-1}X$, and want to get $R(c, a^{-1}b) \Leftrightarrow R(ac, b)$, that is,

$$[q]c^{-1} \cap [q]b^{-1}a \text{ is } I\text{-wide} \iff [q]c^{-1}a^{-1} \cap [q]b^{-1} \text{ is } I\text{-wide}.$$

For this it is enough that for all $P \in \mathrm{Def}(X)$ we have

$$Pc^{-1} \cap Pb^{-1}a \in I \iff Pc^{-1}a^{-1} \cap Pb^{-1} \in I.$$

Let $P \in \mathrm{Def}(X)$. Then $Pc^{-1}a^{-1} \cap Pb^{-1} = P_0b^{-1}$ with $P_0 \in \mathrm{Def}(X)$, and then $Pc^{-1} \cap Pb^{-1}a = P_0b^{-1}a$. Assuming (1), it follows that both sides in the last display are equivalent to $P_0 \in I$, and thus equivalent to each other.

We now show that (2) is enough to replace right-invariance in the proof of Claim 12. In that proof $\mathrm{tp}_Y(s|Mc)$ is I-wide. Assume (2); we wish to derive that $\mathrm{tp}_{X^*}(sc^{-1}|Mc)$ is I-wide. Suppose the latter is not I-wide. Since $sc^{-1} \in X^{-1}X \subseteq X^*$ we get Mc-definable $P \subseteq X^{-1}X$ such that $sc^{-1} \in P$ and $P \in I$. Then $Pc \in \mathrm{tp}_Y(s|Mc)$, and $Pc \in I$ (since $c \in X^{-1}X$), a contradiction.

The first use of right-invariance in the proof of Claim 6 can also be taken care of in this way, and here the following weaker form of (2) suffices: for all $P \in \mathrm{Def}(X^{-1}X)$ and $a \in X^{-1}X$: $\quad P \in I,\ Pa \subseteq X^{-1}X \implies Pa \in I$.

Note that (1) and (2) are both consequences of the condition (right). We now consider how left-invariance of I has been used.

Left-invariance. The proof of Theorem 2.6 uses only the following weak forms of left-invariance of I:

(3) for all $P \in \mathrm{Def}(X)$ and $a \in X$, $\quad P \notin I \implies a^{-1}P \notin I$;
(4) for all $P \in \mathrm{Def}(X^{-1}X)$ and $a \in X^{-1}X$,

$$P \in I,\ aP \subseteq X^{-1}X \implies aP \in I;$$

(5) for all $P \in \mathrm{Def}(X^*)$ and $a \in X$, $\quad P \in I \implies a^{-1}P \in I$.

Indeed, (3) suffices to replace the use of left-invariance in showing that $[q]^{-1}[q]$ is I-wide and in proving Claims 4, 8, and 10. Left-invariance as used in the proof of Claim 7 can be replaced by (4). Left-invariance as used in the proof of Claim 12 can be replaced by (5). Note that (3), (4), (5) are consequences of the condition (left).

This takes care of any kind of translation invariance of I as used in the proof of Theorem 2.6.

Translation invariance of I as used in showing that S' has properties like those claimed of S in Theorem 2.6, with X^{-1}, q^{-1}, I^{-1} instead of X, I, q, can be replaced by the left versions of (1) and (2), and the right versions of (3), (4), (5). The left version of (1) is also enough to replace the use of left-invariance in the proof of Claim 14, and the use of right-invariance in the proof of Claim 16 is taken care of by (1). All these weaker versions of translation invariance of I are implied by (left)$\wedge$(right). $\square$

Note that in the proof of Theorem 2.6 only Claim 12 seems to require the stronger forms of translation invariance given by (2) and (5).

We call X a *near-subgroup* of G over A if there is an ideal I on $\mathrm{Def}(\widehat{X})$ such that $I|X^*$ and $I|(X^*)^{-1}$ are S1 over A, $X \notin I$, and for all $P \in \mathrm{Def}(X_3)$ and $a \in X_1$: $P \in I \iff aP \in I \iff Pa \in I$. (This requirement of translation invariance appears to be slighly stronger than the condition (left)$\wedge$(right) of Corollary 2.9.) The following is a variant of Lemma 2.3

Corollary 2.10. *Let X be a near-subgroup of G over A with countable L and A. Then both XX^{-1} and $X^{-1}X$ are left-generic and right-generic.*

Proof. Take an ideal I as in the definition of *near-subgroup*. By Lemma 2.8 we have a countable $M \supseteq A$ and a type q such that (X, I, q) satisfies the hypotheses of Corollary 2.9. Pick $a \in [q]$ and set $S = S' := a[q]^{-1}[q][q]^{-1}$, an M-closed I-wide normal subgroup of $\widehat{X}$ with $|\widehat{X}/S| \le 2^{\aleph_0}$. For $X^{-1}X$ to be left-generic it suffices by part (1) of Lemma 2.5 that S can be covered by finitely many left translates of $X^{-1}X$. Take a maximal $E \subseteq [q]^{-1}[q][q]^{-1}$ such that $eX \cap fX = \emptyset$ for all distinct $e, f \in E$. Then

$$a^{-1}S \;=\; Sa^{-1} \;=\; [q]^{-1}[q][q]^{-1} \;\subseteq\; \bigcup_{e\in E} eX^{-1}X,$$

so $X^{-1}X$ is indeed left-generic if E is finite. Suppose towards a contradiction that E is infinite. Then $|E| > 2^{\aleph_0}$, and we have $e[q] \cap f[q] = \emptyset$ for all distinct $e, f \in E$. For $e \in E$ we have $e[q] \subseteq S$, so

$$e[q]a^{-1} \subseteq Sa^{-1} = [q]^{-1}[q][q]^{-1} \subseteq X^{-1}XX^{-1}.$$

Via Erdös-Rado we obtain infinite $F \subseteq E$ and $P \in \mathrm{Def}(X)$ with $[q] \subseteq P$ such that $ePa^{-1} \subseteq X^{-1}XX^{-1}$ for all $e \in F$, and $ePa^{-1} \cap fPa^{-1} = \emptyset$ for all distinct $e, f \in F$. In particular, $P \notin I$, and for $e \in F$ we have $ePa^{-1} \subseteq X_3$, so $ePa^{-1} \notin I$ by several applications of the translation invariance required of a near-subgroup. As $ePa^{-1} \cap fPa^{-1} = \emptyset$ for all $e \in F$, this contradicts the assumption that $I|X^{-1}XX^{-1}$ is S1.

Thus $X^{-1}X$ is indeed left-generic, and since $X^{-1}X$ is symmetric, it is also right-generic. Using (X^{-1}, I^{-1}, q^{-1}) instead of (X, I, q), we obtain that XX^{-1} is left-generic and right-generic. $\square$

Making measures definable. In a big model, the zero ideal of an A-invariant Keisler measure μ on an A-definable set Y with $\mu(Y) < \infty$ is S1 over A, as was mentioned earlier. Here we indicate a first-order setting from which such Keisler measures emerge in a *definable* way. This will be used at the end of this section. In this subsection the language L is *one-sorted*, with a fixed countably infinite supply of variables. (This restriction simplifies definitions and is enough for the applications in these notes.)

We single out one particular variable x. We extend L to a (still one-sorted) language $\mathbf{p}L$ by a recursive clause that assigns to each $\mathbf{p}L$-formula

$$\phi = \phi(x; y) \qquad \text{with } y = (y_1, \dots, y_n)$$

and rational $r > 0$ an n-ary relation symbol $\mathbf{p}^r_\phi$. One way to interpret these new symbols is as follows. Let $\mathcal{M} = (M; \dots)$ be an L-structure (not necessarily small) and F a nonempty finite subset of M. Then $\mathcal{M}$ has a unique $\mathbf{p}L$-expansion $\mathcal{M}[F]$ such that for each $\phi = \phi(x; y)$ and r as above

$$\mathcal{M}[F] \models \mathbf{p}^r_\phi(b) \iff |\phi(M; b)| < r|F| \qquad (b \in M^n).$$

The $\mathbf{p}L$-sentences in (i)–(v) below are clearly true in such $\mathbf{p}L$-structures $\mathcal{M}[F]$. In (i)–(v) we let r, s, t range over $\mathbb{Q}^{>0}$, and ϕ, ψ over $\mathbf{p}L$-formulas with free variables as indicated, with $y = (y_1, \dots, y_n)$ and $z = (z_1, \dots, z_p)$ disjoint tuples of distinct variables, all distinct from x.

(i) the universal closure $\forall y \forall z \theta(y, z)$ of each formula $\theta(y, z)$ of the form

$$[\mathbf{p}^r_\psi(y) \wedge \forall x\big(\phi(x; y) \to \psi(x; z)\big)] \to \mathbf{p}^s_\phi(z) \qquad (r < s);$$

(ii) the sentences $\mathbf{p}^r_\phi$, for the formula $\phi(x) = \perp$ defining $\emptyset \subseteq M$;

(iii) the universal closure of $\mathbf{p}^r_\phi(y) \to \mathbf{p}^s_\phi(y)$, for $\phi = \phi(x; y)$ and $r < s$;

(iv) the universal closure of each formula

$$[\mathbf{p}^r_\phi(y) \wedge \mathbf{p}^s_\psi(y) \wedge \neg\exists x(\phi(x; y) \wedge \psi(x; y))] \to p^t_\Phi(y)$$

for $\Phi(x; y) := \phi(x; y) \vee \psi(x; y)$ and $r + s < t$;

(v) the universal closure of each formula

$$[\neg\mathbf{p}^r_\phi(y) \wedge \neg\mathbf{p}^s_\psi(y) \wedge \neg\exists x(\phi(x, y) \wedge \psi(x, y))] \to \neg\mathbf{p}^t_\Phi(y)$$

for Φ as in (vi) and $r + s > t$;

Let $\bar{\mathbf{p}}L$ be the set of $\mathbf{p}L$-sentences described in (i)–(v) above. It depends only on L. Let $\mathcal{M} = (M; \dots)$ be an L-structure and $\mathbf{p}\mathcal{M}$ a $\mathbf{p}L$-expansion of $\mathcal{M}$ such that $\mathbf{p}\mathcal{M} \models \bar{\mathbf{p}}L$. To distinguish definability with respect to $\mathcal{M}$ and $\mathbf{p}\mathcal{M}$, we introduce some notation: if $Y \subseteq M^n$ is definable in $\mathcal{M}$ we put

$$\mathrm{Def}(Y) := \{P \subseteq Y : \ P \text{ is definable in } \mathcal{M}\},$$

and if $Y \subseteq M^n$ is definable in $\mathbf{p}\mathcal{M}$ we put

$$\mathbf{p}\,\mathrm{Def}(Y) := \{P \subseteq Y : \ P \text{ is definable in } \mathbf{p}\mathcal{M}\}.$$

Then we have a finitely additive measure [5] $\mu : \mathbf{p}\,\mathrm{Def}(M) \to [0, \infty]$ such that for $\phi = \phi(x; y)$ as before,

$$\mu\big(\phi(M; b)\big) = \inf\{r \in \mathbb{Q}^{>0} : \ \mathbf{p}\mathcal{M} \models \mathbf{p}^r_\phi(b)\} \qquad (b \in M^n).$$

[5] This includes the requirement $\mu(\emptyset) = 0$.

Thus for $\phi(x;y)$ as before and $b \in M^n,\ s \in \mathbb{R}^{\geq 0}$,

$$\mu\big(\phi(M;b)\big) = s \iff \mathbf{p}\mathcal{M} \models \neg \mathbf{p}^r_\phi(b) \text{ for all } r \in \mathbb{Q}^{>0} \text{ with } r < s, \text{ and}$$
$$\mathbf{p}\mathcal{M} \models \mathbf{p}^t_\phi(b) \text{ for all } t \in \mathbb{Q}^{>0} \text{ with } s < t.$$

It follows that for $\phi(x;y)$ as above and $s \in \mathbb{R}^{\geq 0}$ the set

$$\{b \in M^n :\ \mu\big(\phi(M;b)\big) = s\}$$

is countably definable over $\emptyset$ in $\mathbf{p}\mathcal{M}$ (that is, an intersection of countably many subsets of M^m that are 0-definable in $\mathbf{p}\mathcal{M}$). This fact plays a role for big $\mathbf{p}\mathcal{M}$, and we now turn to this case.

Suppose that $\mathbf{p}\mathcal{M}$ is big. Then the induced finitely additive measure μ on $\mathbf{p}\,\mathrm{Def}(M)$ is clearly a 0-invariant Keisler measure. Therefore, if $Y \subseteq M$ is A-definable in $\mathbf{p}\mathcal{M}$ and $\mu(Y) < \infty$, then the zero ideal

$$I := \{P \in \mathbf{p}\,\mathrm{Def}(Y) :\ \mu(P) = 0\}$$

of the restriction of μ to $\mathbf{p}\,\mathrm{Def}(Y)$ is S1 over A.

An application. Recall that G denotes a group and X a subset of G containing the group identity. For $a \in G$ we put

$$a^X := \{x^{-1}ax :\ x \in X\}.$$

Theorem 2.11. *Let $k, l, m \in \mathbb{N}^{\geq 1}$ be given. Then there is an $n \geq 1$ such that for any G, X with finite X and $|X^*| \leq k|X|$, if there is no $D \subseteq X^{-1}X$ with $|D| \geq |X|/n$ and $|a_1^X \cdots a_l^X| < |X|/m$ for all $(a_1, \dots, a_l) \in D^l$, then there is a normal subgroup H of $\widehat{X}$ such that $H \subseteq a^{-1}X^*$ for some $a \in X$, and $X \subseteq \bigcup_{i=1}^{km} a_i H$ for some $a_1, \dots, a_{km} \in X$.*

The explicit bound km for the number of cosets of H enough to cover X was noticed by Henson. As to the proof of Theorem 2.11, if there is no n as claimed, then there is a sequence of counterexamples with n tending to infinity. This suggests working in a nonstandard G with a hyperfinite X such that $|X^*| \leq k|X|$ and there is no internal $D \subseteq X^{-1}X$ such that $|D|/|X|$ is not infinitesimal, and $|a_1^X \cdots a_l^X| < |X|/m$ for all $(a_1, \dots, a_l) \in D^l$. Using the internal counting measure μ on G normalized such that $\mu(X) = 1$, and taking its zero ideal restricted to $\widehat{X}$ we can then try to apply the stabilizer theorem to get an H as required.

What we do below can be viewed in that way, but we prefer to avoid the formalism of non-standard analysis, and aim instead for a result that is more general in not being tied to counting measures, and in a more constructive spirit in showing that n can be taken as a recursive function of (k, l, m).

To describe this result, let L be the one-sorted language of groups with an extra unary predicate symbol X. We construe pairs (G, X) as L-structures in the obvious way.

Proposition 2.12. *Let $\mathbf{p}(G,X)$ be a $\mathbf{p}L$-expansion of a pair (G,X) such that $\mathbf{p}(G,X) \models \bar{\mathbf{p}}L$, and let μ be the induced measure on $\mathbf{p}\,\mathrm{Def}(G)$. Suppose $\mu(X) = 1$, and $\mu(P) = \mu(gP) = \mu(Pg)$ for all $P \subseteq G$ definable in (G,X) and all $g \in G$. Let $k,l,m \in \mathbb{N}^{\geq 1}$, and suppose $\mu(X^*) \leq k$ and there is no $D \subseteq X^{-1}X$ definable in $\mathbf{p}(G,X)$ with $\mu(D) > 0$ and $\mu(a_1^X \cdots a_l^X) < 1/m$ for all $(a_1,\ldots,a_l) \in D^l$.*

Then there is a normal subgroup H of $\widehat{X}$, definable in $\mathbf{p}(G,X)$, such that $H \subseteq a^{-1}X^*$ for some $a \in X$ and $X \subseteq \bigcup_{i=1}^{km} a_i H$ for some $a_1,\ldots,a_{km} \in X$.

Proof. We can assume that $\mathbf{p}(G,X)$ is big. Consider the set

$$Q := \{(a_1,\ldots,a_l) \in (X^{-1}X)^l : \ \mu(a_1^X \cdots a_l^X) \geq 1/m\}.$$

Note that $Q = \Phi(G^l)$ for some countable set Φ of $\mathbf{p}L$-formulas $\phi(y_1,\ldots,y_l)$ where Φ is independent of $\mathbf{p}(G,X)$. The complement $(X^{-1}X)^l \setminus Q$ is *sparse*: it has no subset D^l with $D \subseteq X^{-1}X$ definable in $\mathbf{p}(G,X)$ and $\mu(D) > 0$.

Let I be the restriction to $\widehat{X}$ of the zero ideal of μ, so I is left-and-right-invariant, $I|X^*$ is S1 over $\emptyset$, and $X \notin I$. Then by the stabilizer theorem and Lemma 2.8 we obtain a countable $M \preceq \mathbf{p}(G,X)$ and a normal I-wide subgroup H of $\widehat{X}$ that is countably definable over M in $\mathbf{p}(G,X)$, such that $H \subseteq a^{-1}X^*$ for some $a \in X$ and

$$H \subseteq (X^{-1}X) \cup Y \quad \text{with } Y \in I.$$

To get H definable in $\mathbf{p}(G,X)$, it suffices to show that its relative complement $X_4 \setminus H$ is also type-definable.

Let (H_n) be a descending sequence of subsets of G, each definable in $\mathbf{p}(G,X)$, such that $H = \bigcap_n H_n$. Since $H \subseteq (X^{-1}X) \cup Y$ we can arrange that $H_n \subseteq (X^{-1}X) \cup Y$ for all n. Now for all n we have $\mu(H_n) > 0$, so $\mu(H_n \cap X^{-1}X) > 0$, and so $H_n^l \cap Q \neq \emptyset$. Therefore $H^l \cap Q \neq \emptyset$, say, $(a_1,\ldots,a_l) \in H^l \cap Q$, so $\mu(a_1^X \cdots a_l^X) \geq 1/m$. Since H is normal in $\widehat{X}$, we have $a_1^X \cdots a_l^X \subseteq H$. If $b \in X$, then $bH \subseteq X^* \cup bY$ with $\mu(bY) = 0$, so at most km cosets of H can meet X, that is, $|X/H| \leq km$. In particular, X/H is finite, and so X_4/H is finite as well. Then $X_4 \setminus H$ is a union of finitely many sets of the form $X_4 \cap aH$ with $a \in X_4$, so $X_4 \setminus H$ is type-definable, and so H is definable. □

Let now $k,l,m \in \mathbb{N}^{\geq 1}$ given. It is easy to specify an infinite set $\mathrm{Hyp}(k,l,m)$ of $\mathbf{p}L$-sentences, including all sentences of $\bar{\mathbf{p}}L$, such that (with the present values of k,l,m) the *hypothesis* of Proposition 2.12 can be expressed as:

$$\mathbf{p}(G,X) \models \mathrm{Hyp}(k,l,m).$$

Likewise, one can specify an infinite set $\mathrm{Con}(k,l,m)$ of $\mathbf{p}L$-sentences such that the the *conclusion* of Proposition 2.12 can be expressed as:

$$\mathbf{p}(G,X) \models \sigma \quad \text{for some } \sigma \in \mathrm{Con}(k,l,m).$$

Of course, these sets $\mathrm{Hyp}(k,l,m)$ and $\mathrm{Con}(k,l,m)$ should depend only on k,l,m, not on $\mathbf{p}(G,X)$. (To express normality, use that a subgroup of $\widehat{X}$ is normal in $\widehat{X}$ iff it is normalized by each $x \in X$.) Compactness gives *finite* sets $\Delta \subseteq \text{Hyp}(k,l,m)$ and $\Sigma \subseteq \text{Con}(k,l,m)$ such that

$$\bigwedge_{\delta\in\Delta} \delta \longrightarrow \bigvee_{\sigma\in\Sigma} \sigma$$

is a logical truth. This yields in particular a number $n = n(k,l,m) \geq 1$ such that Proposition 2.12 holds with "$\mu(D) > 0$" replaced by "$\mu(D) \geq 1/n$". In the situation of Theorem 2.11 we take the counting measure on G normalized so that X has measure 1 to obtain the validity of this theorem with the above $n = n(k,l,m)$.

References

[1] L. van den Dries, Introduction to model-theoretic stability, available at `http://www.math.uiuc.edu/~vddries/`

[2] L. van den Dries, Type-definable sets and their quotients, available at `http://www.math.uiuc.edu/~vddries/`

[3] E. Hrushovski, Stable group theory and approximate subgroups, arXiv:0909.2190v2

[4] T. Tao, Reading seminar: "Stable group theory and approximate subgroups", by Ehud Hrushovski, available at `http://terrytao.wordpress.com/category/teaching/logic-reading-seminar/`