

1. PRELIMINARIES

We will consider pairs (G, X) with G a group, X a subset; for simplicity, we will assume $1 \in X$ and $X = X^{-1}$.

Definition 1.1. Let $X, Y \subset G$. Write $X \lesssim Y$ if there exist $g_1, \dots, g_k \in G$ with $X \subseteq \cup_{i=1}^k g_i Y$. X, Y are commensurable, $X \approx Y$, if $X \lesssim Y$ and $Y \lesssim X$.

We would like to understand (G, X) up to:

- (1) Commensurability.
- (2) $(G, X) \simeq (\tilde{G}, X)$ if $X \subset \tilde{G}$.
- (3) $(G, X) \simeq (\mathbf{G}, \mathbf{X})$ if $\pi : G \rightarrow \mathbf{G}$ is a surjective group homomorphism, and $X = \pi^{-1}(\mathbf{X})$.

E.g. $(G, F) \simeq (1, 1)$ for any finite $F \subset G$.

- Remark 1.2.**
- (1) Clause (2) means we can always replace G by $\langle X \rangle$, the group generated by X . In other words, only products of elements of X play a role. Model-theoretically, we can treat G as many-sorted, with sorts X, X^2, \dots .
 - (2) In fact, only products of < 100 elements of X will be used; moreover the *associativity* assumption on G will only be used that far, i.e. explicitly for such products. This state of affairs is called a *local group*.
 - (3) In the approach of [2], local groups are necessary for the proof, even if one wants the result only for actual groups; but I think we will be able to avoid them, except inasmuch as they are implicit in (1); so you are free to ignore (2).

1.3. Measures on structures. Let M be a structure; in practice it will have the form (G, X, \cdot) . Recall that $Def(M) = \cup_n Def_n(M)$ is the smallest Boolean algebra of subsets of M_n (for some n), closed under coordinate projections and pullbacks, including the basic relations of L and the diagonal on M , and all singleton sets (elements of M).

Let $D = Def_1(M)$ be the Boolean algebra of all subsets of M definable with parameters. *definable with parameters* in (G, X, \cdot) .

We will consider *finitely additive measures* on M . These are functions $\mu : D \rightarrow \mathbb{R}^{\geq 0}$, where D is a certain Boolean algebra of subsets of G . We can take D to be the algebra of all sets *definable with parameters* in (G, X, \cdot) . (Further discussion later.)

Translation invariance: the condition we really need: μ extends to an ultrapower G^* ; $\mu^{-1}(0)$ is invariant under $Aut(G^*)$ and under left, right translations.

A stronger condition: μ is definable; and for any definable set X , *either* all translates of X have measure zero, *or* all translates of X have measure bounded above zero.

2. INTRODUCTION

Consider six “categories”¹.

Near: G has a measure μ as above, with $\mu(X) > 0$ and $\mu(X^3) < \infty$. In this situation X is called a *near-subgroup* of G .

Approx: Same as NEAR; and in addition, $X^2 \simeq X$. Then X is called an *approximate subgroup* of G .

LC: G is a locally compact group.

Lie: G is a Lie group; moreover G is connected, with no normal compact subgroups other than 1. We say that G is *ncnc*.

¹we will use this term and the associated term “functor” completely informally, and will not define the class of morphisms.

bdd \wedge : $G \supset X_1 \supset X_2 \supset \dots$ is a group with a descending family of subsets, with $1 \in X_n = X_n^{-1}$ and $X_{n+1}^2 \subset X_n$. Moreover, all X_i are commensurable with X_1 .

FinApp: By a *finite near-group* we mean a sequence (G_m, X_m) of groups G_m (not necessarily finite) and *finite* subsets X_m , satisfying $1 \in X_m = X_m^{-1}$, such that for some $k \in \mathbb{R}$ independent of M , $|XXX| \leq k^2|X|$. A *finite approximate group* is a sequence satisfying the stronger condition: $X_m X_m \subset F_m X_m$ for some F_m with $|F_m| \leq k$.

Theorem 2.1 ([8]). • *The four categories bdd \wedge , Approx, LC, LIE are essentially equivalent. There are functors between any two; going back and forth yields the identity on LIE, and yields an equivalence as in Definition 1.1 on APPROX.*

- *Approx \subset NEAR; in the converse direction if $X \in$ NEAR then $XX \in$ Approx, and $X \simeq XX$.*
- *For any ultrafilter u , we obtain a functor $FinApp \rightarrow$ NEAR mapping $(G_m, X_m) \rightarrow \lim_u G_m, \lim_u X_m$.*

Remark 2.2. (1) In LIE and LC, we take X to be a compact neighborhood of 1. Note that all such neighborhoods are commensurable, so up to equivalence, only the group matters.

- (2) We gave the quick proof of $(X \in NEAR \implies XX \in Approx)$ assuming μ is left-invariant and X^4 has finite measure (Rusza). The construction of bdd \wedge from NEAR, and of Approx from bdd \wedge , use only X^3 .
- (3) The equivalence of bdd \wedge , Approx, LC can be stated more generally for approximate equivalence relations, and will be proved at this generality. The group-theoretic case follows easily, since the transitions between the three classes are canonical and preserved under the automorphism groups.

2.3. From finite near-groups to near-groups, and back.

Lemma 2.4. *Let (G_m, X_m) be a finite near (resp. approximate) group. Let μ_m be the counting measure on G_m , normalized so that $\mu_m(X_m) = 1$. Let u be an ultrafilter on ω (non-principal.) Consider the ultraproduct (G, X, μ) of the (G_m, X_m) along u . Here $\mu(Y)$ is defined for any Y of the form $\lim_u Y_m, Y_m \subset G_m$; and $\mu(Y) = \lim_u \mu_m(Y_m)$. Then (G, X, μ) is a near (respectively approximate) group. Moreover, any definable partial ordering on $X(k)$ has a maximum element.*

This allows us to move information on near groups to bear on finite approximate groups. In particular the finite statement of [2], Corollary ??, follows by a compactness argument from a statement on near groups, Corollary ??. Our structural investigation will take place entirely in the four equivalent, infinite categories.

2.5. Discussion of bdd \wedge . An \wedge -definable equivalence relation $E = \cap R_n$ is *co-bounded* if the number of classes in a model M is bounded independently of M . Equivalently, for each n there exists a finite $b = b(n)$ such that among any $b + 1$ elements a_0, \dots, a_n , two may be found (say a_i, a_j) with $R(a_i, a_j)$.

A \wedge -definable subgroup of a group G is a family of definable subsets $(X_i : i \in I)$, where $I = (I, <)$ is a directed partial ordering, with $X_n \subset X_m$ if $n > m \in I$, $1 \in X_n = X_n^{-1}$, and such that for any $n \in I$, some $m \in I$, $X_m X_m \subset X_n$. This is the same as saying that $\cap X_i(G^*)$ is a subgroup of G^* , where G^* is any ultrapower of G .

Remark 2.6. *Let (X_i) be a \wedge -definable subgroup of G , $\Gamma = \cap X_i$. Then TFAE:*

- (1) *X is co-bounded, i.e. for any $i > j$, X_i is covered by finitely many translates of X_j .*
- (2) *For any $i > j$, X_i does not contain infinitely many disjoint translates of X_j .*
- (3) *$|X(G^*)|$ is bounded independently of $G^* \succ G$.*

Lemma 2.7. *Let Γ be a $\text{bdd}\wedge$ subgroup of G . Then so is $N = \bigcap_{g \in \tilde{G}} g^{-1} X g = \bigcap_i \bigcap g \in \tilde{G} X_i = \bigcap_i \bigcap g \in X_i^2 X_i$.*

It almost always suffices to consider $I = \mathbb{N}$; e.g. any \wedge -definable subgroup of G^* is an intersection of such countably-defined ones.

2.8. From $\text{bdd}\wedge$ to LC; from LC to Lie.

2.9. From LC to $\text{bdd}\wedge$. Let G be a locally compact group, X a compact neighborhood of 1. Let (X_i) be a basis for the topology at 1, consisting of compact symmetric sets. Fix one of them, X_0 , and let $X = X_0^2$. Clearly $(G, X_i)_i$ is a co-bounded \wedge -definable group. If one wishes for a countable index set, extending slightly Lemma 2.7, it is easy to find an \wedge -definable group $N = \bigcap_{j \in J} X_j$ with J countable, such that N is normalized by $\tilde{G} = \langle X \rangle$. (For any j, m , for some $k > j$, $X_k \subset \bigcap_{y \in X^{-m}} y^{-1} X_j y$.) Note X contains $X_0 N$. Factor out N .

2.10. From LIE to NEAR. Let L be a Lie group. It is a theorem (stated by Lie and proved by F. Schur, according to Hilbert) that any Lie group admits an *analytic* structure. Thus L can be taken to have a neighborhood U of 1 which is a ball in \mathbb{R}^n , as well as a sub-ball $U_1 \subset U$, such that multiplication is given by analytic functions.

By a theorem of Gabrielov (see Denev-Van-den-Dries), for any neighborhood of 1 in G with analytic boundary, any definable subset of (G, \cdot, X) is subanalytic, hence certainly Borel and so Haar measurable. This shows that the Haar integral gives a map from LIE to NEAR (and so to Approx).

Remark 2.11. *To avoid using Lie-Schur and Gabrielov, one could restrict the measure to a subclass of the definable sets; i.e. a slightly weaker definition of NEAR, a posteriori equivalent to NEAR, suffices.*

2.12. From NEAR to $\text{bdd}\wedge$. We need this special case of Theorem 3.1 of [8]. It follows a sequence of similar results in model theory, called "stabilizer theorems".

Theorem 2.13. *Let (G, μ, X) be a near-subgroup. Then there exist a μ -wide, \wedge -definable subgroup S of G , $S \subset X^4$. (Moreover S is normal in \tilde{G} , and $S \setminus X^2$ is not wide.)*

Note:

Corollary 2.14. *Let (G, μ, X) be a near-subgroup. Then there exist a μ -wide definable set Y with $Y^8 \subset X^4$.*

Proof. Say $S = \bigcap S_n$; then $\bigcap S_n^8 \subset S \subset X^4$, so for some n , $S_n^8 \subset X^4$. □

The corollary is in fact easily seen to be equivalent to the theorem. For (ultraproducts of) finite approximate subgroups, it was given an independent, direct proof by Sanders, following a line in combinatorics starting from Balog-Szemerédi. See [2], Theorem 5.3 for a self-contained proof in about one page.

We will nevertheless give the model-theoretic proof, in part in order to introduce the notion of a stable relation. We also give a more general formulation for equivalence relations. (I don't know if Sanders' proof can be generalized to this setting.)

Definition 2.15.

- A bipartite graph $(G_1, G_2; R)$ (i.e. a relation $R \subset G_1 \times G_2$) is *stable* if for some m , there are no $a_1, \dots, a_m \in G_1, b_1, \dots, b_m \in G_2$ with $R(a_i, b_j)$ for $i < j$ while $\neg R(a_i, b_j)$ for $i > j$.
- Let R_n be a descending sequence of relations $R_n \subset G_1 \times G_2$. We say " $\bigcap_n R_n$ is stable" if for each n , for some $n' > n$ and some m , there are no $a_1, \dots, a_m \in G_1, b_1, \dots, b_m \in G_2$ with $R_{n'}(a_i, b_j)$ while $\neg R_n(a_i, b_j)$ for $i > j$.

- Let $\rho : G_1 \times G_2 \rightarrow [0, 1] \subset \mathbb{R}$ be a function. ρ is stable if for any $\alpha < \beta$, for some m , there are no $a_1, \dots, a_m \in G_1, b_1, \dots, b_m \in G_2$ with $\rho(a_i, b_j) < \alpha$ for $i < j$ while $\rho(a_i, b_j) > \beta$ for $i > j$.

Exercise 2.16. Let μ be a probability measure on G_3 ; let $S_i \subset G_i \times G_3$ be a relation; define $R(a, b) \iff \mu(S_1(a) \cap S_2(b)) = 0$. Then R is stable. (In fact $\mu(S_1(a) \cap S_2(b)) \leq \alpha$ is stable for any α , see [[8] Prop. 2.25].)

Say I is a good ideal on $G_1 \times G_2$ if there are proper, automorphism invariant ideals I_i on definable subsets of G_i such that any element of I_i divides over \emptyset , and if $D \in I$ then $D(a, x) \in I_2$ and $D(x, b) \in I_1$ for all $a \in G_1, b \in G_2$.

If I is an ideal on definable subsets of D , and $f : D \rightarrow \mathbb{R}$ is a function, say $f(x) = \alpha$ for I -almost all $x \in D$ if for any $\epsilon > 0$, for some $D' \in I$, $|f(x) - \alpha| < \epsilon$ for all $x \in D \setminus D'$.

Theorem 2.17 ([1]). Let $f : G_1 \times G_2 \rightarrow \mathbb{R}$ be stable. Then there exists a co-bounded \wedge -definable equivalence relation E_i on G_i , and a good ideal I on definable subsets of $G_1 \times G_2$, such that (in any elementary extension), for any class X_1 of E_1 and X_2 of E_2 , for some $\alpha \in \mathbb{R}$, $f(x_1, x_2) = \alpha$ for almost all $(x_1, x_2) \in X_1 \times X_2$.

The same is true when $G_1 = G_2$ is a class of a \vee -definable equivalence relation \tilde{E} on a complete type Q , provided: $f(x, y) > 0$ implies $x\tilde{E}y$; either E will be \wedge -definable, or else $E = \tilde{E}$.

Consider relations R on a set Y ; i.e. R is a subset of Y^2 . For $n = 2, 3, \dots$, define the compositional powers of R by: $R^n = R^{n-1} \circ R$. For simplicity, as we did for groups, we will take R to be symmetric and reflexive. (In the general case we would obtain $R \circ R^{op} \circ R \circ R^{op}$ in place of R^{o4} .)

We will consider measures such that $\mu(R(a)) < \infty$ for all a . By an ϵ -slice we mean a set U such that for all a , $|R^{o3}(a) \cap U| \leq \epsilon R(a)$. We say an \wedge -definable set is *wide* if it is not contained in any 0-slice.

Theorem 2.18. Fix $k \in \mathbb{N}$, $m \in \mathbb{N}$. Let R be a symmetric, reflexive relation on a set G ; let $R^n = R^{n-1} \circ R$. Assume $R(b)$ is finite, and $|R^{o3}(a)|/|R(b)| \leq k$ for $a, b \in G$. Then there exists a symmetric, reflexive relation S such that $S^{om} \subset R^{o4}$, and for all $a \in G$ outside an ϵ -slice U , $|S(a)| \geq O_{k,m}(1)|R(a)|$.

Moreover S is 0-definable, uniformly in (G, R) , in a language with cardinality comparison quantifiers; in particular $\text{Aut}(G, R)$ leaves U, S invariant.

If $\text{Aut}(G, R)$ acts transitively on G , then of course $|S(a)| \geq O_{k,m}(1)|R(a)|$ for all $a \in G$. So we recover in this setting Corollary 2.19:

Corollary 2.19. Fix $k \in \mathbb{N}$, $m \in \mathbb{N}$. Let G be a group, X a finite subset of G , $1 \in X = X^{-1}$, and assume $|X^{\cdot 3}| \leq k|X|$. Then there exists S , $1 \in S = S^{-1} \subset G$, such that $S^m \subset X^{\cdot 4}$ and $|S| \geq O_{k,m}(1)|X|$.

Corollary 2.20. Fix $k \in \mathbb{N}$, $m \in \mathbb{N}$. Let G be a group, H a subgroup (not necessarily normalized by X), $X \subset G$ a set with $HX = X$ and $|X^{\cdot 3}/H| \leq k|X/H|$. Then there exists S , $1 \in S = S^{-1} \subset G$, such that $S^m \subset X^{\cdot 4}$ and $|S/H| \geq O_{k,m}(1)|X/H|$.

Proof of Theorem 2.18. Suppose otherwise. So for each $c \in \mathbb{N}$ there is a counterexample (G, R) , such that there are no S, U as stated with $|S(a)| \geq c^{-1}|R(a)|$ for $a \notin U$. Take a nonprincipal ultraproduct, and let μ be the ultraproduct of the counting measures, normalized so that $\mu R(a) = 1$ for some $a \in G$; hence $1/k \leq \mu R(a) \leq k$ for all $a \in G$. Then there are no definable S, U with S symmetric, reflexive, $S^{om} \subset R^{o4}$, U a 0-slice, and $\mu S(a) > 0$ for all $a \in G \setminus U$.

By Theorem 2.21 there is a cobounded 0 - \wedge -definable equivalence relation E such that $E \subset R^{\circ 4}$. We have $E = \bigcap_n S_n$ for some sequence $S_1 \supset S_2 \supset \dots$ of symmetric, reflexive definable relations. Since $E^{\circ m} = E \subset R^{\circ 4}$, some S_i (denote it by S) satisfies $S^{\circ m} \subset R^{\circ 4}$. In addition we have: $E(a)$ is μ -wide for all a realizing a wide type. Thus $\mu(S(a)) > 0$ if $tp(a)$ is wide. Let \mathcal{U} be the set of all 0 -definable thin slices. So no a falls into the \wedge -definable set: $\mu(S(x)) = 0$ intersected with all complements of sets $U \in \mathcal{U}$. By compactness, for some $U \in \mathcal{U}$, we have $\mu(S(a)) > 0$ for all $a \in G \setminus U$. This contradicts the first paragraph. \square

We consider finitely additive measures on the \mathbb{U} -definable sets in some sort D , over a base set A with $L(A)$ countable. We assume each measure μ is *Borel definable*, meaning: for any formula $\phi(x, y), \mu\phi(x, b)$ depends only on $tp(b/A)$, and is a Borel function $S_y(A) \rightarrow \mathbb{R}$. Then given μ on D and μ' on D' , we can define the iterated measure $\mu \otimes \mu'$ on $D \times D'$, by $(\mu \otimes \mu')(X) = \int (y \mapsto \mu'(X(y))) d\mu(y)$. See [?]. Note that $\mu(x) \otimes \mu'(y)$ is another Borel-definable measure, on $D' \times D$. If the natural map $D \times D' \rightarrow D' \times D, (x, y) \mapsto (y, x)$, is measure-preserving, we say that μ, μ' commute. If μ, μ commute, say that μ is *symmetric* or *self-commuting*. In general a measure need not commute with itself, but ultraproducts μ of counting measures do: $\mu(x) \otimes \mu(y) = \mu(y) \otimes \mu(x)$.

Theorem 2.21. *Fix k . Let R be a symmetric, reflexive relation on a set G . Let μ be a symmetric definable measure, with $\mu(R^{\circ 3}(a))/\mu(R(b)) \leq k$ for $a, b \in G$. Then there exists a 0 - \wedge -definable equivalence relation E , such that $E \subset R^{\circ 4}$, E is co-bounded in R , and $E(a)$ is μ -wide for all a realizing a wide type.*

Remark 2.22. The statement on width in Theorem 2.21 is automatically true: the co-boundedness of E implies that $E(a)$ is μ -wide for all a realizing a wide type. Indeed let q be a wide type. So $q \cap R(a)$ is wide for some a . We have to show for any definable S' with $q \subset S'$ than $\mu(S'(a)) > 0$. Let $E \subset S$, S definable, symmetric, and with $S^{\circ 2} \subset S' \subset R^{\circ 4}$. Let a_1, \dots, a_r be a maximal subset of $q \cap R(a)$ such that $\neg S(a_i, a_j)$. This is indeed a finite set since E is co-bounded. So the sets $S^{\circ 2}(a_i)$ cover $q \cap R(a)$; since q is wide, we must have $\mu S'(a_i) > 0$ for some i . This proves that $E(a)$ is wide for $a \models q$.

2.23. Random elements. We have a continuous map $r : S_x(\mathbb{U}) \rightarrow S_x(A)$ between type spaces. Let $r_*\mu$ be the measure on $S_x(A)$ induced by μ (so the $r_*\mu$ -measure of a clopen subset of $S_x(A)$, namely of $\{q : U \in q\}$ for some A -definable set U , is just $\mu(U)$.) Then by Radon-Nykodim there exists a Borel map $p \mapsto \mu|_p$ from $S_x(A)$ to Borel-definable measures on D , such that for any definable U , $\mu(U) = \int \mu|_p(U) dr_*\mu(p)$; and $\mu|_p$ concentrates on $r^{-1}(p)$, i.e. $\mu|_p(U) = 0$ if U is A -definable and $U \not\subset r^{-1}(p)$. In general, $\mu|_p$ is not definable even if μ is definable; but it is Borel definable over A .

Lemma 2.24. *Assume μ is self-commuting. Then for $r_*\mu$ -almost all p , the measures $\mu, \mu|_p$ commute.*

Proof. It suffices to show, for all continuous functions $g(x, y)$ on S_{xy} , that $\int g d\mu(x) \otimes d\mu|_p(y) = \int g d\mu|_p(y) \otimes d\mu(x)$. The two sides of this equation can be seen as functions of p , and we need to show that these functions coincide $r_*\mu$ -a.e. Now in general to show that two functions agree a.e., it suffices to show for any Borel set U that the integral of their products with 1_U are equal. It suffices here to show for any definable D' that $\int \int 1_{D'}(p) g d\mu(x) \otimes d\mu|_p(y) dr_*\mu(p) = \int \int 1_{D'}(p) g d\mu|_p(y) \otimes d\mu(x) dr_*\mu(p)$. Let μ' be the restriction of μ to D' ; This resolves to $\int \int 1_{D'}(x) g d\mu(x) \otimes d\mu(y) = \int \int 1_{D'}(x) g d\mu(y) \otimes d\mu(x)$. But this follows from the self-commutation of μ . \square

In fact for almost all pairs (p, q) , the measures $\mu|_p, \mu|_q$ commute; though $\mu|_p$ need not commute with itself, even if μ does.

Let $Def_0(D)$ be the Boolean algebra of subsets of D generated by A -definable sets. Let $Def_{n+1}(D)$ be the Boolean algebra generated by the sets: $\{a \in D : \mu(R(a)) = 0\}$, where $R \in Def_n(D \times D')$. So each $Def_n(D)$ is countable. μ extends to a finitely additive measure on each Def_n .

Call $a \in D$ k -random over A for μ if it avoids all measure-zero sets in $Def_k(D)$. And ω -random if k -random for all k .

Note that ω -randomness is still a model-theoretic and not purely measure- or set - theoretic condition. Notably, it is possible for $tp(a/b)$ to be ω -random even though $tp(a) = tp(b)$ and $\mu(tp(b)) = 0$.

Exercise 2.25 (Fubini). *Let μ, μ' be Borel-definable measures on D, D' respectively. If a/A is $k+1$ -random for μ and $b/A(a)$ is $k+1$ -random for μ' , then $(a, b)/A$ is k -random for $\mu \otimes \mu'$. And conversely, if $(a, b)/A$ is $k+1$ -random then a/A is $k+1$ -random and $b/A(a)$ is k -random.*

Let $R^t = \{(b, a) : (a, b) \in R\}$.

In the following lemma, we assume to simplify notation that R is symmetric.

Lemma 2.26. *Let μ be a Borel-definable measure on D , over A . Assume $\mu(x) \otimes \mu(y) = \mu(y) \otimes \mu(x)$. Let R be a symmetric relation on $D \times D$, such that $R^t(c)$ has measure > 0 for all c . Then:*

- (1) *For almost all types q on D , there exist $a, b \in q$ and $c \in D$ with $R(a, c), R(b, c)$, $c/a, b$ wide, $b/A(a)$ does not divide over A .*
- (2) *Assume $c, d \in D$, $c \equiv_{l_c} d$ over A . Then there exist $a, b \in D$, $a \equiv_{l_c} b$, with $b/A(a)$ μ -wide, $tp(a/A)$ random, and $R(a, c), R(b, d)$.*

Proof. (1) Fix a random q , so that μ_q is a probability Borel-definable measure commuting with μ (Lemma 2.24) and let $a \models q$. Let c be such that $R(a, c)$, and $c/A(a)$ is ω -random. Then $a/A(c)$ is ω -random for μ (using the fact that μ, μ commute, and Lemma 2.25.) Choose b with $b/A(a, c)$ μ_q - ω -random. In particular $b/A(a)$ is μ_q - ω -random so it does not divide over A . By Lemma 2.25, (b, a, c) is random for $\mu_q \otimes \mu \otimes \mu$. Now these measures commute, so we can apply Lemma 2.25 in the opposite direction and conclude that $c/A(a, b)$ is μ -wide.

(2) Find a with $tp(a/A(c))$ random, $R(a, c)$. Let C be the \equiv_{l_c} -class of a . As $c \equiv_{l_c} d$ over A , there exists an automorphism σ fixing A and with $\sigma(C) = C$, $\sigma(c) = d$. Then $\sigma(a)/d$ is wide, and $\sigma(a) \in C$. So the partial type asserting that $x \in C$ and $x \models tp(\sigma(a)/d)$ is wide. Let $b \in C$, $bd \equiv \sigma(a)d$, and $b/a, d$ wide. \square

Remark 2.27. Let $A \leq M$, M a model. Then we have a continuous surjection from $S_x(M)$ to the set of Lascar types $S_x^{lc}(A)$; it induces a Borel measure on the space of compact Lascar types. The proof of Lemma 2.26 (1) over M provides, for a random q over M , realizations b, a with $\mu(R(a) \cap R(b)) > 0$, and such that (in particular) $tp(b/A(a))$ is 1-random, so it does not fork over A . It follows Lemma 2.26 (1) is valid for compact Lascar types over A : for almost all $q \in S_x^{lc}(A)$, for some $a, b \models q$ such that $tp(a/A(b))$ does not divide over A , $\mu(R(a) \cap R(b)) > 0$.

Proof of Theorem 2.21. We may assume the language is countable. We will show that $\equiv_{l_c} \subset R^{\circ 4}$. By Lemma 2.26 (1) and remark 2.27, for almost all compact Lascar types q over A , for some $a, b \models q$ such that $tp(a/A(b))$ does not divide over A , $\mu(R(a) \cap R(b)) > 0$. By Theorem 3.27, this is true for all such a, b . In particular, the distance between them is ≤ 2 . It follows from Lemma 2.26 (2) that for arbitrary c, d with $c \equiv_{l_c} d$, the R -distance from c to d is at most 4. \square

3. APPENDIX 1: STABILITY

We develop the basic results of stability, presented here in Theorems 3.14 and 3.27. We view them as a reduction, modulo a certain ideal, of binary relations to unary ones; thus a kind

of measurability result for binary relations for the product measure. The theory is primarily due to Shelah, and for the most part we follow standard presentations. Shelah understood the significance of having the theorem over an arbitrary base structure and not just over an elementary submodel, and introduced imaginary elements and the algebraic closure as the precise obstructions to this. In [?], the theory was extended beyond the first order setting. In [?], the main theorem was proved for arbitrary invariant stable relations *over a model*. In the same paper, for simple theories, the “bounded closure” with its compact automorphism group was recognized as the obstacle to existence of 3-amalgamation (in the finite rank case, the algebraic closure still sufficed, as shown in [?].) See [4] for a good presentation of the compact and general Lascar types; we will use it below. In [1], the theory was beautifully developed for continuous real-valued relations; 3.27 is a (less elegant) generalization for more general \wedge -definable stable relations.

The novelties here are: (i) we treat arbitrary automorphism-invariant stable relations, over any base set. We show that the fundamental theorems of stability theory hold, with strong Lascar types as the natural obstacles to both uniqueness and existence. (ii) For \wedge -definable relations, we show that compact Lascar types or Kim-Pillay types suffice. This generalizes the continuous real-valued case; different proofs are required for certain parts. (iii) We introduce a “local setting”, allowing notably to discuss stable independence over an “imaginary” element of the form a/E , where E is a \vee -definable equivalence relation.

We begin with (iii); readers interested only in (i) or (ii) can skip this, and ignore the metric later, i.e. assume it is bounded.

To ease the notation we will sometimes assume the language is countable, though the general case carries no real difficulties. We will work over a countable base denoted A , and sometimes use a countable elementary submodel M containing A .

When $R \subset X \times Y$, and $a \in X$, we let $R(a) = \{b : (a, b) \in R\}$. Define $R^t \subset Y \times X$, $R^t = \{(b, a) : (a, b) \in R\}$. When the context leaves no room for doubt, for $b \in Y$ we will write $R(b)$ for $R^t(b)$.

3.1. Local structures. Let \mathbb{U} be a structure with a metric $d : \mathbb{U}^2 \rightarrow \mathbb{N}$. If many sorts are allowed, we still assume the domain of d is the set of all pairs, belonging to the union of all sorts. We assume that any closed ball of finite radius is 0- \vee -definable.²

A typical way to obtain such a structure is to begin with an arbitrary binary relation R_0 on another structure \mathbb{U}_0 . Let \tilde{E} be the equivalence relation generated by R_0 . Then any \tilde{E} -class is naturally a local structure; the metric distance $d(x, y)$ is the length of a shortest chain $x = x_0, \dots, x_n = y$ with $R(x_i, x_{i+1})$ or $R(x_{i+1}, x_i)$ for each i . Here the balls are 0-definable.

More generally, we could take the distance along the Gaifman graph with respect to some set of definable relations.

A relation $R(x_1, \dots, x_n)$ is *local* if it implies $d(x_i, x_j) < m$ for some m . (For unary relations, this poses no constraint.) We will be concerned only with local relations. There is always a reduct generated by the local relations, which is local. This is closely related to the Gaifman graph, frequently used in finite model theory, and to Gaifman’s theorem on this subject. We will say, when only local relations are allowed, that the structure is local.

The *definable sets* are obtained by closing the basic relations under finite unions, intersections, differences, projections, and distance-bounded universal quantifiers, of the form: $(\forall x)(d(x, y) \leq 5 \rightarrow \phi(x, y))$. The complement of a definable set is only Ind-definable.

²There are natural generalizations to bigger semigroups than \mathbb{N} , both in the direction of continuous metrics and of uncountable languages, but we restrict here to the main case.

If E is a \bigvee -definable equivalence relation in a saturated structure, for simplicity a countable union of 0-definable relations, then each E -class can be presented as a local structure; the local structures setting will enable us to speak about independence over an E -class (viewed as a (generalized) imaginary element of the base.) We can present E as having the form $d(x, y) < \infty$, where d is a metric such that $d(x, y) \leq n$ is definable, for each n . Then we can take the basic relations to be the d -bounded ones (this does not depend on the choice of d .) Keeping long distance (non-local) relations would not change the automorphism group - they can be recovered as bounded unions of local relations.

If a local structure \mathbb{U} has a constant symbol, or more generally a nonempty bounded definable set D , then it can be viewed as an ind-definable, in fact piecewise-definable structure; the union of the definable sets of points at distance $\leq n$ from D . In general however, the automorphism group here need not respect any specific inductive presentation.

The metric can be extended to imaginary sorts; first to \mathbb{U}^n via: $d((x_1, \dots, x_n), (y_1, \dots, y_n)) = \max(\max_i \min_j d(x_i, y_j), \max_j \min_i d(x_i, y_j))$; then to a quotient by a bounded equivalence relation, with quotient map $\pi : \mathbb{U}^n \rightarrow \mathbb{U}^n/E$, with distance defined by $d(u, v) = \inf d(x, y) : \pi(x) = u, \pi(y) = v$.

We assume \mathbb{U} is saturated as a local structure, or locally saturated: any ball is saturated; equivalently any small family of definable sets has nonempty intersection, provided the family includes a bounded set, and that any finite subset has nonempty intersection. Local saturation can be achieved by taking an ultrapower using bounded functions only.

A remark on ultraproducts: if (N_i, d_i) are a family of local structures for the same language, an (N, d) is an ultraproduct in the usual sense, one has an equivalence relation: $d(x, y) \leq n$ for some standard n ; each equivalence class is a local structure, and Los's theorem holds. thus an ultraproduct here requires a choice of an ultrafilter along with a component, rather than just an ultrafilter. \square

3.2. Locally compact Lascar types. Call a sort S *separated* if it carries a 0- \bigwedge -definable *cobounded* local equivalence relation. If S is separated, let $\equiv_{lc} = \equiv_{lc}^S$ be the intersection of all 0- \bigwedge -definable *cobounded* local equivalence relations on S . Then \equiv_{lc} is the unique smallest such relation.

Let $\pi = \pi_S^{lc} : S \rightarrow S/\equiv_{lc}$ be the quotient map. On S/\equiv_{lc} we define a topology: Y is closed iff $\pi^{-1}Y$ is locally \bigwedge -definable.

Lemma 3.3. *The quotient by \equiv_{lc} is a locally compact space.*

Proof. See earlier notes (or [?], [4]) for the bounded case, of Kim-Pillay spaces. Let $a \in S$, and let B_n be the ball of radius $n + m$ around a , in S . Then $\pi(B_n)$ is compact (so S/\equiv_{lc} is σ -compact.) Since \equiv_{lc} is local, say $d(x, y) < m$ for $(x, y) \in S^2$ with $x \equiv_{lc} y$. Then the closed sets $\pi(S \setminus B_{n+m}), \pi(B_n)$ are disjoint. Thus $\pi(B_{m+1})$ contains a neighborhood of $\pi(a)$, the complement of $\pi(S \setminus B_{m+1})$. \square

Remark 3.4. *The local algebraic closure $\text{acl}(A)$ can be defined as the union of the locally finite definable sets. The automorphism group of \mathbb{U}/A has a quotient group acting faithfully on $\text{acl}(A)$, referred to as the automorphism group of $\text{acl}(A)$ over A is a locally profinite group (a totally disconnected locally compact group.) The stabilizer of a nonempty set is a compact group (fixing one point implies leaving invariant balls of various radii.)*

One can similarly define the local compact closure to be the union of S/\equiv_{lc} , over all sorts S such that \equiv_{lc} is defined.

On the other hand, we consider the more general setting of $Aut(\mathbb{U}/A)$ -invariant equivalence relations. Assume S has an $Aut(\mathbb{U}/A)$ -invariant cobounded local equivalence relation. Then it has a smallest one; it is denoted \equiv_{Las} . This equivalence relation is generated by $\cup_m \theta_m(a, b)$, where $\theta_m(a, b)$ holds iff a, b begin an indiscernible sequence, and $d(a, b) \leq m$. When d has the property that any two elements are connected by a chain of elements of distance 1, as is the case in the main examples, \equiv_{Las} is generated by θ_2 . At any rate, \equiv_{Las} is an F_σ relation (a countable union of \wedge -definable relations.)

3.5. Σ -compactness. The stability theory we will develop - more precisely, existence of generic extensions of a given type - requires Σ -compactness and not just local compactness. We thus assume:

(Σ): for some m_0 , for all n , any n -ball is a finite union of m_0 -balls.

Σ -compactness is true in the setting of a measure, finite on balls. More precisely assume μ is a 0-definable measure, each ball of radius 1 has nonzero measure, and each ball of radius ≤ 3 has finite measure. Then by Rusza's trick, any ball of radius 3 is a union of finitely many balls of radius 2 (consider a maximal disjoint set of radius 1- balls in the radius 3 ball; then enlarging them to radius 2 would cover the larger ball.) Assume in addition that the metric space is "geodesic" in the sense that any two points of length n are joined by a path of length n , where the successive distance is 1 (as is the case for Gaifman graphs.) Then it follows inductively that any ball of radius n is a union of finitely many balls of radius 2.

We are interested only in types of elements at finite distance from elements of \mathbb{U} . In the presence of Σ -compactness, any such type has bounded distance $\leq m_0$ from some element of \mathbb{U} . It follows that if \mathcal{X} is an $Aut(\mathbb{U}/A)$ -invariant closed set of types over \mathbb{U} , then \mathcal{X} contains a compact subset X with $Aut(\mathbb{U}/A)X = \mathcal{X}$.

3.6. Aside on continuous logic. This above use of a metric for local structures, with concern for large values, is dual to the function of the metric in the compact logic of [1], where the concern is with small values of d . They could easily be combined; this would give an unbounded real-valued logic, where the automorphism group of the bounded closure is a locally compact group.
3

3.7. Ideals of definable sets. We will work with saturated (local) structures \mathbb{U} . *Invariance* refers to the action of $Aut(\mathbb{U})$, or $Aut(\mathbb{U}/A)$ for a small substructure A . A set *divides* if for some l it has an arbitrarily large set of l -wise disjoint conjugates (i.e. any l have empty intersection).

We will consider ideals of \mathbb{U} -definable sets (of some sort S). Say I is *definably generated* if it is generated by a definable family of definable sets. Say I is \vee -definable if it is generated by some bounded family of definably generated ideals. Equivalently, for any formula definable $D \subset S \times S'$, $\{b \in S' : S(b) \in I\}$ is \vee -definable. If I is $Aut(\mathbb{U}/A)$ -invariant, then $\{b \in S' : S(b) \in I\}$ is in fact \vee -definable over A .

Dually, I determines a partial type over \mathbb{U} , generated by the complements of the definable sets in I . Any extension of this partial type is called *I -wide*. We say a/A is I -wide if a does not lie in any A -definable set lying in I . Note that $tp(a/A)$ will then extend to an I -wide complete type over \mathbb{U} .

If $f : S \rightarrow S'$ is a 0-definable surjective map, and I is a \vee -definable ideal, let $f_*I = \{D : f^{-1}D \in I\}$. This is a \vee -definable ideal on S' , proper if I is proper. If c/A is I' -wide, then $c = f(b)$ for some I -wide b/A .

If I, I' are two ideals (on S, S'), we can define an ideal $I \otimes I'$ on $S \times S'$, generated by the sets $D \subset S \times S'$ such that for some $D_1 \in I$, for all $a \in S \setminus D_1$, $D(a) \in I'$. So if a/A is I -wide

³This differs from the unbounded logic of [?], which is shown there to reduce to the bounded case functorially; structures have compact absolute Galois groups in this logic, vs. locally compact here.

and $b/A(a)$ is I' -wide, then $(a, b)/A$ is $I \otimes I'$ -wide. Conversely, if $(a, b)/A$ is $I \otimes I'$ -wide, then a/A is I -wide, and - assuming I' is \forall -definable - $b/A(a)$ is I' -wide: to see the last statement, if $b \in D(a) \in I'$, then since I' is \forall -definable, there exists $\theta(x)$ true of a such that $D(a') \in I'$ for all $a' \in \theta$; let $D' = \{(a', b') : b' \in D(a'), a' \in \theta\}$; then $D' \in I \otimes I'$; and $(a, b) \in D'$.

Inductively, we define $I^{\otimes n}$, $I^{\otimes(n+1)} = I^{\otimes n} \otimes I$. We will say $b = (b_1, \dots, b_n)$ is I -wide if it is $I^{\otimes n}$ -wide.

Let us mention here some canonical ideals, relative to a given complete type p over A . There is Shelah's non-forking ideal I_{sh} , generated by the set $Div(p)$ of formulas that divide over A . Given any invariant measure μ (such that p is wide), we have the ideal I_μ of all formulas of μ -measure zero. If μ is definable, then I_μ is \wedge -definable. We have $Div(p) \subseteq I_{sh} \subseteq I_\mu$, for any A -invariant measure μ .

If I is an ideal on S' , let $SDiv(I)$ be the family of generically I -dividing subsets of S ; i.e. the family of sets $Q(b)$, $b \in S'$, Q an A -definable subset of $S \times S'$, such that for any $I^{\otimes n}$ -wide (b_1, \dots, b_n) with $tp(b/A) = tp(b_i/A)$, $\bigcap_{i=1}^n Q(b_i) = \emptyset$. Let \check{I} be the ideal generated by $SDiv(I)$. We have $SDiv(I) \subseteq Div$ and so $\check{I} \subseteq I_{sh}$. If I is \forall -definable over A , so are $SDiv(I)$ and \check{I} .

Definition 3.8. *Let $R \subset P \times P'$ be an invariant relation, and let I be a \forall -definable ideal on P . Say R holds I -almost always if for any $c \in P'$, for any $b \in P$ with $b/A(c)$ I -wide, we have $R(b, c)$. Say R holds I -almost always in the strong sense on $P \times P'$ if R^t holds \check{I} almost always. Explicitly, if whenever $(b, c) \in P \times P' \setminus R$, there exists an A -definable local $Q \subset P \times P'$ such that $(b, c) \in Q$, and for any $I^{\otimes n}$ -wide n -tuple (b_1, \dots, b_n) , $P \cap \bigcap_{i=1}^n Q(b_i) = \emptyset$.*

If $R \subset S \times S'$ is an invariant relation, I a \forall -definable ideal on S , and $P \subseteq S, P' \subseteq S'$ invariant sets, we will also say that R holds I -almost always in the strong sense on $P \times P'$ if $R \cap (P \times P')$ does.

Lemma 3.9. *Assume R holds I -almost always in the strong sense on $S \times S'$. Then:*

- (1) R holds I -almost everywhere.
- (2) If $tp(c/A(b))$ does not divide over A , and $tp(b/A)$ is I -wide, then $R(b, c)$.

Proof. (1) Suppose not; let Q, n be as in Definition 3.8. Let $b_1 = b$. Inductively find b_k such that $Q(b_k, c)$ and b_k is wide over $A(c, b_1, \dots, b_{k-1})$; this is possible since $Q(c)$ is wide. But then $c \in \bigcap_{i=1}^n R(b_i)$, a contradiction.

(2). Suppose $\neg R(b, c)$. Let Q be a definable set as in Definition 3.8, so that for any I -wide $(b_1, \dots, b_n) \in S^n$, $\bigcap_{i=1}^n Q(b_i) = \emptyset$. As $tp(b/A)$ is I -wide, one can find $b_i \models tp(b/A)$ for $i \in \mathbb{N}$, such that $tp(b_n/A(b_1, \dots, b_{n-1}))$ is wide. Then any subsequence of length n of this infinite sequence is $I^{\otimes n}$ -wide, so the intersection of $Q(b_i)$ over any such subsequence is empty. It follows that $tp(c/A(b))$ divides over A . \square

3.10. Stable invariant local relations.

Definition 3.11. *Two definable relations $P(x, y), Q(x, y)$ are stably separated if there is no sequence of pairs $(a_i, b_i) : i \in \mathbb{N}$ with $P(a_i, b_j)$ and $Q(a_j, b_i)$ for $i < j \in \mathbb{N}$.*

Let $R \subset S \times S'$ be an $Aut(\mathbb{U}/A)$ -invariant relation.

Definition 3.12. R is stable if whenever $(a, b) \in R$ and $(c, d) \in (S \times S') \setminus R$, then there exist A -definable sets Q, Q' such that $Q(a, b), Q'(c, d)$ and Q, Q' are stably separated.

Remark 3.13. R is stable iff there is no indiscernible sequence (x_i, y_i) such that for $i \neq j$, $R(x_i, y_j)$ iff $i < j$.

Proof. If no such indiscernible sequence exists, then whenever $(a, b) \in R$ and $(c, d) \in (S \times S') \setminus R$, $tp(a, b)$ and $tp(c, d)$ must be stably separated; by compactness, for some definable P approximating $tp(a, b)$ and Q approximating $tp(c, d)$, P, Q are stably separated. Conversely if (a_i, b_i) is an indiscernible sequence as in the remark, then $tp(a_1, b_2)$ is not stably separated from $tp(a_2, b_1)$ though $R(a_2, b_1)$ and $\neg R(a_1, b_2)$. \square

Theorem 3.14. *Let \mathbb{U} be a Σ -compact local structure. Let \mathfrak{f} be a family of A -invariant stable local relations on $S \times S'$. Let $E_{\mathfrak{f}}$ be the intersection of all co-bounded A -invariant local equivalence relations on S , such that each class is a Boolean combination of a bounded number of sets $R(b) \subset S$, $R \in \mathfrak{f}$. Then for each complete type \bar{P} in S over A , there exists a proper, \bigvee -definable (over A) ideal $I(\bar{P})$ on S , satisfying:*

(*) *If $R \in \mathfrak{f}$, $P \subset \bar{P}$ is an $E_{\mathfrak{f}}$ -class, and Q is an $E_{\mathfrak{f}^t}$ -class on S' , then either R holds almost always in the strong sense for $I(\bar{P})$ on $P \times Q$, or $\neg R$ does.*

Also, symmetry holds: if for P, Q as above, if \bar{Q} is a complete type with $Q \subset \bar{Q}$, then on $P \times Q$, R holds almost always for $I(\bar{P})$ iff R^t holds almost always for $I^t(\bar{Q})$.

There exists a canonical proper \bigvee -definable ideal I_S , such that the dichotomy () and symmetry hold I_S -almost always.*

Remark 3.15. Assume S is a complete type. Then either $E_{\mathfrak{f}}$ is local, or else for any $R \in \mathfrak{f}$, $\neg R$ holds almost always in the strong sense for $I(S)$. In the former case, \equiv_{Las} is local on S , and (*) holds for \equiv_{Las} in place of $E_{\mathfrak{f}}$. (See proof, above Lemma 3.20.)

Though the proofs go through for any \mathfrak{f} , we will assume below that $\mathfrak{f} = \{R\}$ to simplify notation. (In fact the theorem reduces easily to the case that \mathfrak{f} is finite; and then, - replacing S by $S \times \mathfrak{f}$, and considering the relation $\widehat{R}((x, R), y) \iff R(x, y)$ - to the case that \mathfrak{f} has a single element R .)

We will use the space $S_D(\mathbb{U})$ of all bounded global types on a sort D , i.e. types containing a formula implying $d(x, a) \leq n$ for some a, n . If x is a variable of sort D , we will also write $S_x(\mathbb{U})$. Let $(d_p x)R = \{b : R(x, b) \in p\}$. If $(d_p x)R = (d_{p'} x)R$, we say p, p' define the same R -type. We do not define a topology on the set of global R -types.

Lemma 3.16. *Let $S'(x, y), S(x, y)$ be definable relations (of which at least one is local.) Assume $S'(x, y)$ and $S(x, y)$ are stably separated. Then for any type p over M there exists a finite Boolean combination Y of sets $S(x, c_i)$ with $c_i \in M$, such that $d_p y S' \implies Y$ while $Y, d_p y S$ are disjoint.*

Proof. Define $a_n, b_n, c_n \in M$ recursively. Given c_1, \dots, c_n , the equivalence relation: $\bigwedge_{i \leq n} S(x, c_i) \iff S(x', c_i)$ has at most 2^{2^n} classes; if none of these classes meets both $d_p y S$ and $d_p y S'$, then some union Y of these classes contains $d_p y S$ and is disjoint from $d_p y S'$, and the lemma is proved. Otherwise, choose a_n, b_n such that $d_p y S(a_n), d_p y S'(b_n)$, while a_n, b_n lie in the same sets $S(x, c_i)$, $i \leq n$. Then, find c_{n+1} such that $S'(d, c_{n+1}) \iff S'(d, c)$, where $d \in \{a_i, b_i : i \leq n\}$.

For $n < k$ we have $S'(b_n, c_k)$. Applying Ramsey with respect to the question S and refining the sequence (a_n, b_n, c_n) , we may assume that $S(b_n, c_k)$ for all $n > k$ or for no $n > k$; but the former is impossible since S', S are stably separated. So $\neg S(b_n, c_k)$ for all $n > k$

Since a_n, b_n have the same S -type over the smaller c_i , it follows that $\neg S(a_n, c_k)$ for $n > k$. But for $n < k$ we have $S'(a_n, c_k)$; so the sequence (a_n, c_n) contradicts the stable separation of S', S . \square

Corollary 3.17. *Assume M is countable. Let S', S be stably separated local definable relations on $G_1 \times G_2$. There does not exist an uncountable set $W \subset S_x(M)$ such that for $p \neq p' \in W$, for some $b \in M$, $S'(x, b) \in p$ while $S(x, b) \in p'$.*

Proof. Let Y_p be an M -definable set such that $d_p S' \rightarrow Y \rightarrow \neg d_p S$ (Lemma 3.16). There are only countably many choices for Y_p , so there will be $p, p' \in W$ with $Y_p = Y_{p'}$. Now if $S'(x, b) \in p$ then $b \in Y_p = Y_{p'}$ so $\neg S(x, b) \in p'$. \square

It follows that there is no map f from the full binary tree $s^{<\omega}$ into G_2 , such that for each branch $\eta \in 2^\omega$, $\bigwedge S'(x, f(\eta|n+1) : \eta(n) = 0) \wedge \bigwedge S(x, f(\eta|n+1) : f(n) = 1)$ is consistent. By compactness, for some finite n , no such map exists for the height- n tree 2^n . We define the *rank* of a partial type W to be the maximum m such that there exists $f : 2^m \rightarrow G_2$, with $W \wedge \bigwedge S'(x, f(\eta|n+1) : \eta(n) = 0) \wedge \bigwedge S(x, f(\eta|n+1) : f(n) = 1)$ consistent for each $\eta \in 2^m$.

Let R be a stable invariant relation on $G_1 \times G_2$.

Lemma 3.18. *Let p, p' be types over \mathbb{U} . Assume: for any stably separated ϕ, ψ , for some $e = e_{\phi, \psi}$ we have: $e \subset p, p'$ and $rk_{\phi, \psi}(p) = rk_{\phi, \psi}(e) = rk_{\phi, \psi}(p')$. Then $p|R = p'|R$.*

Proof. Let $c \models p$ and $c' \models p'$. Suppose $p|R \neq p'|R$. Then for some $b \in \mathbb{U}$, $tp(b, c)$ implies R but $tp(b, c')$ implies $\neg R$. As R is stable, $tp(b, c)$ and $tp(b, d)$ are stably separated; hence by compactness, some $\phi(x, y) \in tp(b, c)$ and $\psi(x, y) \in tp(b, d)$ are stably separated. Let $e = e_{\phi, \psi}$, $l = rk_{\phi, \psi}(e)$. Let $[\phi(x, b)]$ be the set of types extending $\phi(x, b)$. It follows that either $rk_{\phi, \psi}(e \cap [\phi(b, x)]) < l$ or $rk_{\phi, \psi}(e \cap [\psi(b, x)]) < l$. But $rk_{\phi, \psi}(p) = rk_{\phi, \psi}(p') = l$, a contradiction. \square

In particular, if e is a partial type, and $rk_{\phi, \psi}(p) = rk_{\phi, \psi}(e) = rk_{\phi, \psi}(p')$ for all stably separated (ϕ, ψ) , then $p|R = p'|R$. This hypothesis holds if e is a type over a model M , and p, p' extend e are finitely satisfiable in M .

We can also deduce that for any global p , there are definable $d_p(\phi, \psi), d'_p(\phi, \psi)$ such that p contains $\{\neg\phi(x, b) : d_p(\phi, \psi)(b)\}$ and $\{\neg\psi(x, b) : d'_p(\phi, \psi)(b)\}$; and any type p' containing these formulas has $p|R = p'|R$.

Proposition 3.19. *Let R be a stable local A -invariant relation on $S \times S'$. Let X be a nonempty closed $Aut(\mathbb{U}/A)$ -invariant subset of $S_S(S)(\mathbb{U})$. Let $X|R = \{(d_p x)R : p \in X\}$.*

Then $1 \rightarrow 2 \rightarrow 3$:

- (1) X is minimal.
- (2) for any stably separated ϕ, ψ defined over A , $rk_{\phi, \psi}(p)$ is constant (does not depend on $p \in X$.)
- (3) $X|R$ has cardinality bounded independently of \mathbb{U} ; in fact $|X|R| \leq 2^{|A| + \aleph_0}$.

Proof. (1) implies (2) since the set of elements of X of (ϕ, ψ) -rank $\geq n$ is a closed subset of X .

Now assume (2). Fix ϕ, ψ stably separated, and say $rk_{\phi, \psi}(p) = m$ for $p \in X$. For each ball B of the metric d , the intersection of B, X and the complement of all definable sets of (ϕ, ψ) -rank $\leq m$ is empty; by (local) compactness, $B \cap X$ is covered by finitely many definable sets of (ϕ, ψ) -rank $\leq m$. Thus X is covered by countably many such definable sets, say $e(\phi, \psi, l), l \in \mathbb{N}$. Each p now determines a function $\chi_p : (\phi, \psi) \mapsto l$, where l is least such that $p \in e(\phi, \psi, l)$. But in turn $p|R$ is determined by this function. For if $p, p' \in X$ and $\chi_p = \chi_{p'}$, then by Lemma 3.18, $p|R = p'|R$. This proves (3). \square

Note - this is the place where the Σ -compactness assumption Σ - is used - that for any complete type P over A , there exists a minimal nonempty closed $Aut(\mathbb{U}/A)$ -invariant subset of $S_S(S)(\mathbb{U})$, consisting of elements compatible with P . Indeed let Z be any nonempty closed $Aut(\mathbb{U}/A)$ -invariant subset of $S_S(S)(\mathbb{U})$, consisting of elements compatible with P . Fix $b \in P$,

and let B be the ball defined by $d(x, b) \leq 2m_0$. By Σ , any type p over \mathbb{U} meets some m_0 -ball; this m_0 -ball contains a \mathbb{U} -point a ; so $d(x, a) \leq 2m_0$ is compatible with p . By invariance, $d(x, b) \leq 2m_0$ is compatible with some $p' \in Z$. Thus $\hat{B} \cap Z \neq \emptyset$ (where p is the set of all types over \mathbb{U} of elements of B .) So if Z_i is a descending chain of nonempty closed $\text{Aut}(\mathbb{U}/A)$ -invariant subset of $S_S(S)(\mathbb{U})$, consisting of elements compatible with P , then $Z_i \cap \hat{B}$ is nonempty, and as \hat{B} is compact, $\bigcap Z_i \cap \hat{B}$ is nonempty, and in particular $\bigcap Z_i$ is nonempty. Thus by Zorn's lemma a minimal element exists.

Let $R \subset S \times S'$ be A -invariant, stable.

Let $\text{Gen}_A^R(S)$ be the set of all restrictions $p|R$, where p is a global type of S and $p|R$ has a small orbit under $\text{Aut}(\mathbb{U}/A)$. (The total number of orbits is small, say by Lemma 3.18, so $\text{Gen}_A^R(S)$ is small.) When A or R do not vary, we omit them from the notation. Any type P on S extends to some element of $\text{Gen}(S)$, by Proposition 3.19 1 \rightarrow 3, and the comment below it. It follows that for any \equiv_{Las} -class X on S there exists an element q_X of $\text{Gen}_A^R(S)$ such that for any small N , $q_X|N$ is realized in X . Indeed some \equiv_{Las} -class of P has this property; since all \equiv_{Las} -classes in P are conjugate, all have it.

Similarly define $\text{Gen}^{R^t}(S')$. But for short we will write $\text{Gen}(S), \text{Gen}(S')$.

Define an equivalence relation E_f on S by: $(a, a') \in E_f$ iff for all $p \in \text{Gen}(S')$ and $R \in \mathfrak{f}$, $(d_p y)R(a, y) \iff (d_p y)R(a', y)$; and dually define E_{R^t} on S' . E_f is co-bounded since $\text{Gen}(S')$ is bounded. E_f is local since R is local: if $aE_f b$ then for some c , $R(a, c)$ and $R(b, c)$; so $d(a, b) \leq d(a, c) + d(b, c)$.

We say that $q|R$ is consistent with an invariant set Z if any small subset q_0 of $q|R$ is realized by some element of Z .

Lemma 3.20 (symmetry and uniqueness). *Any E_f -class on S is consistent with a unique $q \in \text{Gen}(S)$. If $q \in \text{Gen}(S), q' \in \text{Gen}(S'), a \in S, a' \in S'$, and q is consistent with $E_f(a)$, and q' with $E_{R^t}(a')$, then $d_{q'} y R(a, y) \iff d_{q_x} R(x, b)$.*

Proof. We prove the symmetry statement first, following the standard route. Suppose for contradiction that it fails for q, q', a, a' . Say $d_{q'} y R(a, y)$ holds but $d_{q_x} R(x, b)$ fails. Construct a_n, a'_n so that $a_n \models q|A(a'_i : i < n)$, $a_n E_f a$, and $a'_n \models q'|A(a_i : i < n)$, $a'_n E_{R^t} a'$. Then since $a_n E_f a$, $d_{q'} y R(a_n, y)$ holds, and similarly $d_{q_x} R(x, a'_n)$ fails. Thus if $i > n$ then $R(a_n, a'_i)$ holds but $R(a_i, a'_n)$ fails. This contradicts the stability of R .

We have already shown that there exists $q' \in \text{Gen}(S')$ consistent with $E_{R^t}(a')$. Now if $q_1, q_2 \in \text{Gen}(S)$ are both consistent with $E_f(a)$, then by symmetry we have $d_{q_1} x R(x, b) \iff d_{q'} y R(a, y) \iff d_{q_2} x R(x, b)$. Thus $q_1 = q_2$. \square

Because of this lemma, if χ is an E_f -class and q is the unique element of $\text{Gen}(S)$ consistent with it, we can write $(d_\chi x)(R(x, y))$ for $(d_q x)R(x, y)$.

Let χ be an E_f -class, consistent with q . Let M be a substructure such that for any two elements $q_1 \neq q_2 \in \text{Gen}(S)$, there exists $b \in M$ with $R(x, b) \in q_1$ but $R(x, b) \notin q_2$, or vice versa. Let E_M^f be the equivalence relation: $a E_M^f b$ iff for any $R \in \mathfrak{f}$ and $b \in M$, $R(a, b) \iff R(a, b')$. Then χ is a cobounded equivalence relation, each class is a bounded Boolean combination of sets $R^t(b)$, and E_M^f refines E_f . Indeed by construction a unique element $q \in \text{Gen}(S)$ will be consistent with a given E_M^f -class χ . So for any $q' \in \text{Gen}(S')$, let d be such that $tp(d/M)$ is consistent with q' ; then for $a \in \chi$, $R(a, y) \in q'$ iff $R(x, d) \in q$.

Since all E_f classes of a complete type P over A are $\text{Aut}(\mathbb{U}/A)$ -conjugate, it follows from uniqueness that all elements q of $\text{Gen}(S)$ consistent with P are $\text{Aut}(\mathbb{U}/A)$ -conjugate.

We choose a minimal nonempty closed $\text{Aut}(\mathbb{U}/A)$ -invariant set $X = X_P$ of global types extending P , as in Lemma 3.19. By this lemma, for any ϕ, ψ , $\beta_p(\phi, \psi) = rk_{\phi, \psi}(p)$ does not

depend on the choice of $p \in X$. Let $I(X_p)$ be the ideal generated by all definable sets D such that for some ϕ, ψ , $rk_{\phi, \psi}(D) < \beta_p(\phi, \psi)$.

Lemma 3.21 (dividing). *Let q' be a global type of elements of S' , Assume $q'|R^t \in \text{Gen}(S')$, P is an E_f -class, and $R(a, y) \in q'$ for $a \in P(\mathbb{U})$. For $i \in \omega_1$, let $b_i \models q'|A(b_j : j < i)$. Then for any $a \in P(\mathbb{U})$, for cofinally many $\alpha \in \omega_1$ we have $R(a, b_\alpha)$.*

Proof. Re-define b_i (without changing the type of the sequence) as follows: let $M_i \prec \mathbb{U}$ be a small model containing a_j for $j < i$, and let $b_i \models q'|M_i$. Let $M = \cup_{i < \omega_1} M_i$. For any pair (ϕ, ψ) , for some $i < \omega_1$, we have $rk_{\phi, \psi}(tp(a/M_i)) = rk_{\phi, \psi}(tp(a/M))$. Since ω_1 has uncountable cofinality, for some $\alpha < \omega_1$, for any ϕ, ψ , $rk_{\phi, \psi}(tp(a/M_\alpha)) = rk_{\phi, \psi}(tp(a/M))$. Since $M_\alpha \prec \mathbb{U}$, there exists a global type q extending $tp(a/M_\alpha)$ such that $rk_{\phi, \psi}(tp(a/M_\alpha)) = rk_{\phi, \psi}(q)$. By Lemma 3.18, $q|R$ is uniquely determined. On the other hand since $q'|R^t \in \text{Gen}_A^R(S')$, it is clear that $q'|R^t \in \text{Gen}_M^R(S')$. Since $R(a, y) \in q'$, by Lemma 3.20, $R(x, b) \in q$ if $tp(b/M_\alpha)$ is consistent with q' . Hence $R(x, b_i) \in q$ for $i \geq \alpha$. But we can also construct a global type q^+ extending $tp(a/M_{\alpha+1})$ with $rk_{\phi, \psi}(tp(a/M_{\alpha+1})) = rk_{\phi, \psi}(q^+)$. As $rk_{\phi, \psi}(tp(a/M_{\alpha+1})) = rk_{\phi, \psi}(tp(a/M_\alpha))$, it follows that $q = q^+$; as $R(x, b_\alpha) \in q$ we have $R(x, b_\alpha) \in q^+$, i.e. $R(a, b_\alpha)$. \square

It follows from Lemma 3.21 (as well as from Lemma 3.20, as we saw before) that $(d_p y)R(x, y)$ is a bounded (but infinitary) Boolean combination of instances of $R(x, b)$; namely $(d_p y)R(a, y)$ iff $R(a, b_j)$ holds for cofinally many j , where (b_j) is a sufficiently long sequence as in the lemma.

Proof of Theorem 3.14. We will use the equivalence relation E_f and the ideals $I(X_p)$ defined above. We have to show

(*) If $R \in \mathfrak{f}$, $P \subset \bar{P}$ is an E_f -class, and Q is an E_{f^t} -class on S' , then either R holds almost always in the strong sense for $I(p)$ on $P \times Q$, or $\neg R$ does.

Pick $p \in X(\bar{P})$, and $p' \in X(\bar{Q})$ (with respect to ${}^t R$.) By definition of E_f , for any $a \in P$, $p'(y)$ implies $R(a, y)$, or else for any $a \in P$, $p'(y)$ implies $\neg R(a, y)$. Without loss of generality the latter holds. Now suppose $\neg R(c, b)$ holds with $c \in P, b \in Q$. As $p'(y)$ implies $R(a, y)$ and $E_{f^t}(a, c)$, $p'(y)$ also implies $R(c, y)$. Let $r = tp(c, b/A)$. We have to show that the condition in Definition 3.8 holds, i.e. that for some n , and some $D \in r$, $\cup D(x, y_j) \cup \neg I_{f^t}^{\otimes n}(y_1, \dots, y_n)$ is inconsistent. Otherwise, there exists a sequence c, b_1, b_2, \dots with $b_k/A(b_1, \dots, b_{k-1})$ wide for I_{f^t} for each k , and $r(c, b_i)$ holds for each i . Let σ be an automorphism taking (c, b) to (c, b_1) . Then $q' = \sigma(p')$ is a global type, $q'|R^t \in \text{GEN}$, consistent with E_{f^t} -class of $\sigma(b_1)$, and $q'(y)$ implies $R(c, y)$ (since $\sigma(c) = c$.) By Lemma 3.21, $R(c, b_i)$ holds for some i . But r is a complete type, and cannot be consistent with both $\neg R(c, b)$ and $R(c, b_i)$. This shows that $\cup D(x, y_j) \cup \neg I_{f^t}^{\otimes n}(y_1, \dots, y_n)$ is indeed inconsistent.

We saw that $(d_p y)R(x, y)$ is a bounded Boolean combination of instances of $R(x, b)$; hence any E_f -class can be expressed as Boolean combination of a bounded number of sets $R(b) \subset S$, $R \in \mathfrak{f}$. Given this, the finest co-bounded equivalence relation with this property refines E_f , and so also satisfies (*). \square

Remark 3.22. *Let $p(x, y)$ be a type (or partial type) over A . Then there exists a unique smallest stable (respectively equational) A -invariant relation P , containing p . (I.e. p implies P .) P is F_σ .*

Proof. (We omit A from the notation, and prove the stable case; the equational case is the same, with $a_0 = a, b_0 = b$ below.) For any invariant relation $P(x, y)$, let $P'(a, b)$ hold iff there exists an indiscernible sequence of pairs (a_i, b_i) with $a_1 = a, b_0 = b$, and $P(a_0, b_1)$. Clearly P'

is ∞ -definable if P is; and P is stable if $P = P'$. Also if $P = \bigvee_j P_j$ then $P' = \bigvee_j P'_j$. So let $P_0 = p$, $P_{n+1} = P'_n$ and $P = \bigcup_{n \in \mathbb{N}} P_n$. Then P is F_σ and stable. \square

Presumably P is usually not \bigwedge -definable. (For instance when p implies \equiv_{Las} and \equiv_{Las} is not ∞ -definable.) Note that \equiv_{Las} is itself a stable invariant relation.

3.23. \bigwedge -definable stable relations. Assume R is ∞ -definable, stable.

First we note that the p -definition of R is ∞ -definable, for any type p .

Lemma 3.24 (definability). *Let $p \in S_x(M)$. Let $R = \bigcap_n R_n$, with R_n definable. Then $d_p R$ is \bigwedge -definable; it is an intersection of Boolean combinations of the definable sets R_n .*

In more detail, the lemma states: fix $m \in \mathbb{N}$. Then there exists n_0 and a finite Boolean combination Y of sets $R_{n_0}(x, c_i)$, $c_i \in M$, such that $d_p R \rightarrow Y \rightarrow d_p R_m$.

Proof. By stability, there is no sequence d_n, e_k with $\neg R_m(d_n, e_k)$ for $k > n$ and $R(d_n, e_k)$ for $k < n$. By compactness, for some n_0 , there is no sequence with $\neg R_m(d_n, e_k)$ for $k > n$ and $R_{n_0}(d_n, e_k)$ for $k < n < n_0$. Thus $\neg R_m, R_{n_0}$ are stably separated. By Lemma 3.16, there exists a finite Boolean combination Y of sets $R_{n_0}(x, c_i)$, $c_i \in M$, such that $d_p R_{n_0} \rightarrow Y \rightarrow d_p R_m$. \square

Lemma 3.25. *Any $E_{\mathfrak{f}}$ -class is \bigwedge -definable. (Over parameters, it is cut out by certain sets of the form $(d_q y)R(x, y)$.)*

Proof. Let P be a complete type of G_1 .

We can find $a \in P$ such that $Q(a) = \{q \in \text{Gen}(S') : a \in (d_q y)R(x, y)\}$ is maximal. This uses Zorn's lemma, and the fact that $(d_q y)R(x, y)$ is ∞ -definable, so if $(d_q y)R(a_i, y)$ and $a_i \rightarrow a$ then $(d_q y)R(a, y)$ (working with types over M .)

Let $Q = Q(a)$. Now $aE_{\mathfrak{f}}b$ iff for each $q \in Q$, $(d_q y)R(b, y)$. So the $E_{\mathfrak{f}}$ -class of a is ∞ -definable.

Since all $E_{\mathfrak{f}}$ -classes in P are conjugate, all $E_{\mathfrak{f}}$ -classes in P are \bigwedge -definable. As P was arbitrary, the lemma follows. \square

Corollary 3.26. *If $a \equiv_{lc} b$ then $(a, b) \in E_{\mathfrak{f}}$.*

Proof. Define: aEb iff $tp(a/c) = tp(b/c)$ for any $E_{\mathfrak{f}}$ -class c (i.e. there exists an automorphism fixing c and taking a to b .) Clearly $E \subset E_{\mathfrak{f}}$. Let $\{C_i : i \in I\}$ list all the classes. then aEb iff for each i , $(\exists c)(\exists d)(c, d \in C_i \& ac \equiv bd)$. Since each C_i is \bigwedge -definable by Lemma 3.25, E is \bigwedge -definable. Since the number of classes C_i is bounded, and elements with the same type over some representative $c_i \in C_i$ also have the same type over C_i , it is clear that E is cobounded. Hence $\equiv_{lc} \subset E$, so $\equiv_{lc} \subset E_{\mathfrak{f}}$. \square

From this and Theorem 3.14 we obtain:

Theorem 3.27 (locally compact equivalence relation theorem). *Let \mathfrak{f} be a nonempty family of ∞ -definable stable local relations on $S \times S'$. Let P, Q be classes of \equiv_{lc} on S, S' respectively. There exists a proper \bigvee -definable ideal I' of definable subsets of S' , such that if $R \in \mathfrak{f}$, then R holds almost always on $P \times Q$ in the strong sense for I' , or $\neg R$ does. Symmetry holds as in Theorem 3.14. Also, the analogue of Remark 3.15 is valid.*

In particular, fix a and assume $tp(a/A)$ forms a single \equiv_{lc} -class; then for b such that $tp(a/Ab)$ or $tp(b/Aa)$ does not divide over A , the truth value of $R(a, b)$ depends only on $tp(b)$.

REFERENCES

- [1] Ben Yaacov, Itai, Usvyatsov, Alexander Continuous first order logic and local stability. *Trans. Amer. Math. Soc.* 362 (2010), no. 10, 52135259.
- [2] Emmanuel Breuillard, Ben Green, Terence Tao, Approximate subgroups of linear groups, arXiv:1005.1881
- [3] The structure of approximate groups. Emmanuel Breuillard, Ben Green, Terence Tao. arXiv:1110.5008
- [4] Casanovas, E.; Lascar, D.; Pillay, A.; Ziegler, M. Galois groups of first order theories. *J. Math. Log.* 1 (2001), no. 2, 305319.
- [5] Cherlin, Gregory, Two problems on homogeneous structures, revisited in *Model Theoretic Methods in Finite Combinatorics*, M. Grohe and J.A. Makowsky eds., Contemporary Mathematics, 558, American Mathematical Society, 2011
- [6] Green, Ben; Ruzsa, Imre Z. Freiman's theorem in an arbitrary abelian group. *J. Lond. Math. Soc. (2)* 75 (2007), no. 1, 163175.
- [7] Joram Hirschfeld, The Nonstandard Treatment of Hilbert's Fifth Problem, *Transactions of the American Mathematical Society*, Vol. 321, No. 1 (Sep., 1990), pp. 379-400
- [8] Hrushovski, Ehud Stable group theory and approximate subgroups. *J. Amer. Math. Soc.* 25 (2012), no. 1, 189243
- [9] Tom Sanders, On a non-abelian Balog-Szemerédi-type lemma, arXiv:0912.0306
- [10] Terence Tao, <http://terrytao.wordpress.com/2007/03/02/open-question-noncommutative-freiman-theorem/>
- [11] Yamabe, Hidehiko A generalization of a theorem of Gleason. *Ann. of Math. (2)* 58, (1953). 351–365.
- [12] Van den Dries, Goldbring, Seminar Notes on Hilbert's 5th problem