

STABLE GROUP THEORY AND APPROXIMATE SUBGROUPS

EHUD HRUSHOVSKI

1. INTRODUCTION

Stable group theory, as developed in the 1970s and 1980s, was an effective bridge between definable sets and objects of more geometric categories. One of the reasons was a body of results showing that groups can be recognized from their traces in softer categories. The first and simplest example is Zilber’s stabilizer. Working with an integer-valued dimension theory on the definable subsets of a group G , Zilber considered the dimension-theoretic stabilizer of a definable set X : this is the group S of elements $g \in G$ with $gX\Delta X$ of smaller dimension than X . Let XX be the product set $XX = \{xy : x, y \in X\}$. If X differs little from XX in the sense that $\dim(XX\Delta X) < \dim(X)$, Zilber showed that X differs little from a coset of S .

In the 1990s, Zilber’s theory was generalized to the “simple theories” of [49], again initially in a definable finite-dimensional context ([8], [23]). Here the definable sets X_t in a definable family $(X_t : t \in T)$ are viewed as “differing little from each other” if simply the pairwise intersections $X_t \cap X_{t'}$ have the same dimension as each X_t . Nevertheless it is shown that when the family of translates $(Xa : a \in X)$ satisfies this condition, there is a group H of the same dimension as X and with a large intersection with some translate of X ; this group was still, somewhat inappropriately, called the stabilizer, and we will keep this terminology.

In the present paper we prove the stabilizer theorem in a general first-order setting. A definition is given of being a “near-subgroup” (Definition 3.9), generalizing the stable and simple cases. We then prove the existence of a nearby group (Theorem 3.5). In outline, the proof remains the same as in [23]; the definability condition on the dimension was removed in [31]. The key is a general amalgamation statement for definable ternary relations, dubbed the “Independence Theorem” (see [8], p. 9 and p. 185). Roughly speaking, in maximal dimension, consistent relations among each pair of types determine consistent relations on a triple; see Theorem 2.22.

The stabilizer obtained in Theorem 3.5 is not a definable group but an \bigwedge -definable one; it is defined by a countable set of formulas in a saturated model, or alternatively as a group object in the category of projective limits of definable sets. In the finite-dimensional setting of [23], the construction of the stabilizer was complemented by a proof that \bigwedge -definable groups are limits of definable groups. This last step is not true at the level of generality considered here: the group of infinitesimals of a Lie group provide counterexamples. We show however that all counterexamples are closely associated with Lie groups; see Theorem 4.2. The proof

Received by the editors August 24, 2010 and, in revised form, May 16, 2011.

2010 *Mathematics Subject Classification*. 11P70, 03C45.

Research supported in part by Israel Science Foundation grant 1048/07.

©2011 American Mathematical Society
Reverts to public domain 28 years from publication

uses the structure theory of Gleason, Montgomery-Zippin and Yamabe for locally compact groups.

A very interesting dictionary between this part of model theory, and certain parts of finite combinatorics, can be obtained by making the model-theoretic “dimension n ” correspond to the combinatorial “cardinality of order c^n ” (cf. [8], 8.4). Near-subgroups in the above sense then correspond to asymptotic families of finite subsets X of a group (or a family of groups), with $(X \cup X^{-1})^3/|X|$ bounded. Equivalently (see [50], Lemma 3.4, and Corollary 3.11 below) $|X^k|/|X|$ is bounded for any given k . Subsets of groups with weak closure conditions were considered in combinatorics at least since [14]. An excellent survey centering on rings can be found in the first pages of [51]; see also [50] for more general noncommutative groups. The parallels to the model-theoretic development are striking. We turn now to a description of some consequences of the stabilizer theorem in this combinatorial setting.

For the sake of the introduction we consider finite subsets of G (more general situations will be allowed later.) We recall Terence Tao’s notion of an *approximate subgroup*. A finite subset $X \subseteq G$ is said to be a k -approximate group if $1 \in X$, $X = X^{-1}$, and XX is contained in k right cosets of X . Say X, Y are *commensurable* if each is contained in finitely many right cosets of the other, with the number bounded in terms of k . It is felt that approximate subgroups should be commensurable to actual subgroups, except in situations involving Abelian groups in some way. See [52] for a compelling exposition of the issue.

Gromov’s theorem [18] on finitely generated groups of polynomial growth fits into this framework, taking X to be a ball of size 2^n in the Cayley graph, for large n ; then X is a 2^d -approximate subgroup, where d is the growth exponent. Gromov shows that the group is nilpotent, up to finite index.

Theorem 4.2 says nothing about a fixed finite approximate subgroup, but it does have asymptotic consequences to the family of all k -approximate subgroups for fixed k . In particular, we obtain:

Theorem 1.1. *Let $f : \mathbb{N}^2 \rightarrow \mathbb{N}$ be any function, and fix $k \in \mathbb{N}$. Then there exist $e^*, c^*, N \in \mathbb{N}$ such that the following holds.*

Let G be any group, X a finite subset, and assume $|XX^{-1}X| \leq k|X|$.

Then there are $e \leq e^, c \leq c^*$, and subsets $X_N \subseteq X_{N-1} \subseteq \dots \subseteq X_1 \subseteq X^{-1}XX^{-1}X$ such that X, X_1 are e -commensurable, and for $1 \leq m, n < N$ we have:*

- (1) $X_n = X_n^{-1}$.
- (2) $X_{n+1}X_{n+1} \subseteq X_n$.
- (3) X_n is contained in the union of c translates of X_{n+1} .
- (4) $[X_n, X_m] \subseteq X_k$ whenever $k \leq N$ and $k < n + m$.
- (5) $N > f(e, c)$.

Roughly speaking, this is deduced as a special case of the following principle: if a sentence of a certain logic holds for all compact neighborhoods of the identity in all finite-dimensional Lie groups, then it holds for all approximate subgroups. We have not explicitly determined the relevant logic; Proposition 6.6 hints that, given further work on the first-order theory of Lie groups with distinguished closed subsets, much stronger transfer principles may be possible than what we have used.

The first three clauses of Theorem 1.1 suggest a part of a noncommutative Bourgain system as defined in [17], and conjectured by Ben Green in [52] to exist for approximate subgroups. Green’s conjecture was in part intended to show that “one

can do a kind of approximate representation theory”, which can be viewed as a description of Theorem 4.2 and the deduction between the two.

The fourth clause suggests a kind of topological nilpotence. Note that (4) implies that $[X_1, X_1] \subseteq X_1$. For a set of generators of a finite simple group, this in itself seems to be a curious property.

The use of the structure theory of locally compact groups here follows Gromov [18]. But the bridge to locally compact groups is a different one: Gromov’s is metric, while ours is measure-theoretic.

It is natural to consider a somewhat more general framework. Call a pair $(X, G, \cdot, {}^{-1}, 1)$ a *Freiman approximate group* if X is a finite subset of G , $\cdot : X^{(2)} \rightarrow G$ and ${}^{-1} : X \rightarrow X$ are functions, such that for any $(x_1, \dots, x_{12}) \in X^{12}$, the iterated products $((x_1 \cdot x_2) \cdot (x_3 \cdot \dots))$ are defined and independent of the placing of the parentheses, $xx^{-1} = x^{-1}x = 1 \in X$, $1 \cdot x = x \cdot 1 = x$, and $|XX^{-1}X|/|X| \leq k|X|$. Then Theorem 1.1 is also valid for Freiman approximate groups. In particular X has a large subset X_2 closed under $[\cdot, \cdot]$ if not under \cdot , and in fact with $[X_2, X_2]^2 \subseteq X_2$. This again suggests that approximateness can only really enter via an Abelian part of a structure. This “local” version uses local versions of the theory of locally compact groups due to Goldbring [16].

The finiteness assumption on X in the above results is really only used via the counting measure “at the top dimension”, so they remain valid in a measure-theoretic setting; see Theorem 4.15.

The remaining corollaries of Theorem 3.5 attempt to make a stronger use of finiteness. They are proved directly, without Lie theory, and go in a somewhat complementary direction. The first assumes that the group generated by an approximate subgroup X is perfect in a certain strong statistical sense. The conclusion is that X is close to an actual subgroup. We write $a^X = \{x^{-1}ax : x \in X\}$.

Corollary 1.2. *For any $k, l, m \in \mathbb{N}$, for some $p < 1$, $K \in \mathbb{N}$, we have the following statement.*

Let G be a group, X_0 a finite subset, $X = X_0^{-1}X_0$. Assume $|X_0X| \leq k|X_0|$. Also assume that with probability $\geq p$, an l -tuple $(a_1, \dots, a_l) \in X^l$ satisfies: $|a_1^X \cdots a_l^X| \geq |X|/m$.

Then there exists a subgroup S of G , $S \subseteq X^2$, such that X is contained in $\leq K$ cosets of S .

We could use $(a^X \cup (a^{-1})^X)^{(l)}$ (or a^{X_0}) in place of a^X above. See Theorem 3.12 for a weaker alternative version of the hypotheses. p can be taken to be a recursive function of k, l, m , but I have made no attempt to estimate it. As Ward Henson pointed out, the proof does give an explicit estimate for K . The proof also shows that X normalizes S . Laci Pyber remarked that with this strengthening (but not without it), the conclusion implies small tripling for X .

Here and later on, when confusion can arise between an iterated set product and a Cartesian power, we use Y^l to denote the former, and $Y^{(l)}$ for the latter.

The assumption of Corollary 1.2 may be strong in a general group-theoretic setting, but it does hold for sufficiently dense subgroups of simple linear groups. The proof uses an idea originating in the Larsen-Pink classification of large finite simple linear groups, [35], somewhat generalized and formulated as a dimension-comparison lemma in [24]. We obtain:

Theorem 1.3. *Let G be a semisimple algebraic group, defined over \mathbb{Z} , and k an integer. Then for some integer k' , the following statement holds. Let K be a field, X a finite subset of $G(K)$ with $|XX^{-1}X| \leq k|X|$. Then there exists a subgroup H of $G(K)$ such that $|X/H| \leq k'$, and either H is (the set of K -points of) a connected proper algebraic subgroup of G of degree $\leq k'$, or $H \subseteq (X^{-1}X)^2$.*

Here “degree $\leq k'$ ” means that if we view G as a subset of the $n \times n$ matrices M_n , then H_i is the intersection of G with a subvariety of M_n cut out by polynomials of degree $\leq k'$. Thus if the group generated by X is sufficiently Zariski dense, X will not be contained in such an algebraic subgroup, so that $X^{-1}X$ must be commensurable to a subgroup. A special case:

Corollary 1.4. *For any $n \in \mathbb{N}$, for sufficiently large $n' \in \mathbb{N}$, the following holds. Let X be a finite subset of $GL_n(K)$, K a field, with $|X| \geq n'$. Assume $|XX^{-1}X| \leq k|X|$, and that X generates an almost simple group S . Then $(XX^{-1})^2 = S$.*

Here S is not *assumed* to be finite. “Almost simple” means: perfect, and simple modulo a center of bounded size. The proof also shows that XX^{-1} contains 99% of the elements of S , and that $XX^{-1}X = S$; see the proof and remarks following Proposition 5.10.

For $S = SL_2(F_q)$ and $SL_3(F_p)$, Theorem 1.3 follows from results of Helfgott [20], [21] and Dinai; for $G = SL_2(\mathbb{C})$ and $G = SL_3(\mathbb{Z})$, Theorem 1.3 follows from [7] and [11]. These authors all make a much weaker assumption on a subset X of a group, namely $|XX^{-1}X| \leq |X|^{1+\epsilon}$ for a small ϵ . The combinatorial regime they work in is also meaningful model-theoretically (cf. Example 2.13), but we do not study it at present.

Stable group theory includes a family of related results; for instance, the group law may be given by a multi-valued or partial function. The partial case has antecedents in algebraic geometry, in Weil’s group chunk theorem. A version of the partial case, including the Freiman approximate groups mentioned above, will be briefly noted in the paper. It is likely that the multi-valued case too admits finite combinatorial translations along similar lines.

In §2 we introduce the model-theoretic setting and prove the independence theorem and the stabilizer theorem in a rather general context. In the presence of a σ -additive measure the stabilizer sounds close to Tao’s noncommutative Balog-Szemerédi-Gowers theorem ([50]), while the independence theorem is, in the finite setting, extremely close to the Komlós-Simonovitz corollary [32] to Szemerédi’s lemma (as I realized recently while listening to a talk by M. Malliaris.) It is thus quite possible that combinatorialists can find other proofs of the results of §2 and skip to the next section. I find the independent, convergent development of the two fields rather fascinating.

All the results we need from stability will be explicitly defined and proved. Theorem 1.1 (and the more detailed Corollary 4.15) are proved in §4. The methods here are very close to [25]; however, we do not assume NIP. This is in line with a sequence of realizations in recent years that tools discovered first in the stable setting are in fact often valid, when appropriately formulated, for first-order theories in general. Theorem 1.3 is proved in §5.

§6 contains a proof that the topology on the associated Lie group is generated by the image of a definable family of definable sets.

In §7, we use the techniques of this paper, along with Gromov’s proof of the polynomial growth theorem, to show (for any k) that if a finitely generated group is *not* nilpotent-by-finite, it has a finite set of generators contained in no k -approximate subgroup.

1.5. Basic model theory: around compactness. We recall the basic setup of model theory, directed to a large extent at an efficient use of the compactness theorem. We refer the reader to a book such as [6], [37], [43] or the lecture notes in [41] for a fuller treatment. We assume knowledge of the definition of a first-order formula, and of the compactness theorem, asserting that a finitely satisfiable set of formulas is satisfiable in some structure.

Let L be a fixed language, T a theory, M a model. We will occasionally use notation as if the language is countable (e.g. indices named n), but this will not be really assumed unless explicitly indicated. At all events for much of this paper, a language with a symbol for multiplication and an additional unary predicate will be all we need.

A will refer to a subset of M . We will assume L, A are countable (this is quite inessential, and will be used only to avoid the need for cardinal parameters in discussing saturation below). We expand L to a language $L(A)$ with an additional constant symbol for each element of A . The L -structure M is tautologically expanded to an $L(A)$ -structure, and the result is still denoted M , by abuse of notation. $T(A)$ is the $L(A)$ -theory of M . $L_x(A)$ denotes the Boolean algebra of formulas of $L(A)$ with free variables x , up to $T(A)$ -equivalence. $S_x(A) = \text{Hom}(L_x(A), 2)$ is the Stone space, or the space of *types*. A subset of $L_x(A)$ is *finitely satisfiable* if each finite subset has a common solution in M . A *type* in a variable x , over A , is a maximal finitely satisfiable subset of $L_x(A)$. For an element or tuple a over a subset A of a model M , $tp(a/A) = \{\phi(x) \in L(A) : M \models \phi(a)\}$; if $tp(a/A) = p$, we say that a realizes p . An *A-definable set* is the solution set of some $\phi \in L(A)$. It is an easy corollary of the compactness theorem that every theory T has models \mathbb{U} with the following properties holding for every small substructure B of \mathbb{U} . Here let us say that B is *small* if $2^{|B|} \leq |\mathbb{U}|$.

- (1) Saturation: Every type over B is realized in \mathbb{U} .
- (2) Homogeneity: For c, d tuples from M , $tp(c/B) = tp(d/B)$ iff there exists $\sigma \in \text{Aut}(M/B)$ with $\sigma(c) = d$.

(In fact (1) implies (2) if the generalized continuum hypothesis holds; moreover, in this case, \mathbb{U} is determined up to isomorphism by T and by $|\mathbb{U}|$, provided T is complete.)

Given a complete theory T , we fix a model \mathbb{U} of T with the above properties and with $|\mathbb{U}| \gg \aleph_0$ (if it is not finite), and interpret definable sets as subsets of \mathbb{U}^n . We will occasionally consider elementary submodels M of \mathbb{U} ; these will just be referred to as *models*. We write $A \leq M$ to mean that A is a substructure of M . As stated above, substructures A will be taken to be countable; “small” would be enough.

A partial type over A is any collection of formulas over A , in some free variable x , and closed under implication in the L_A -theory of M .

The solution sets D of partial types r (over various countable sets A) are called \wedge -definable (read: ∞ -definable) sets; so an \wedge -definable set over A is any intersection of A -definable sets. The correspondence $r \mapsto D$ is bijective, because of the saturation property (1) above. Complements of \wedge -definable sets are called \vee -definable. An

equivalence relation is called \wedge -definable if it has an \wedge -definable graph. It follows from saturation that an \wedge -definable set is either finite or has size $|\mathbb{U}|$; a \vee -definable set is either countable or has size \mathbb{U} ; an \wedge -definable equivalence relation has either $\leq 2^{\aleph_0}$ classes or $|\mathbb{U}|$ -classes. Since $|\mathbb{U}|$ is taken to be large, this gap lends sense to the notion of *bounded size* for sets and quotients at these various levels of definability.

Another consequence of countable saturation is that projections commute with countable decreasing intersections:

$$(\exists x) \bigwedge_{i=1}^{\infty} \phi_i(x, y) \iff \bigwedge_{i=1}^{\infty} (\exists x) \phi_i(x, y)$$

provided that ϕ_{i+1} implies ϕ_i for each i . The condition on the left beginning with $(\exists x)$ seems to be stronger, but compactness assures that the weaker condition on the right suffices for the existence of x in some model, and countable saturation implies that such an x exists in the given model. In particular, the projection of an \wedge -definable set is \wedge -definable. We will use this routinely in the sequel. Specifically, if Q is an \wedge -definable subset of a definable group (see below), then the product set $QQ = \{x : (\exists y, z \in Q)(x = yz)\}$ is also \wedge -definable.

By a *definable group* we mean a definable set G and a definable subset \cdot of G^3 , such that $(G(\mathbb{U}), \cdot(\mathbb{U}))$ is a group. An \wedge -definable subgroup is an \wedge -definable set which is a subgroup. It need not be an intersection of definable subgroups. We insert here a lemma that may clarify these concepts.

A subset of a set X is *relatively definable* if it has the form $X \cap Z$ for some definable Z .

Lemma 1.6. *Let G be a definable group. Let X be an \wedge -definable subset of G , Y a \vee -definable subset of G , and assume that X and $X \cap Y$ are subgroups of G , and that $X \cap Y$ has bounded index in X . Then $X \cap Y$ is relatively definable in X and has finite index in X .*

Proof. By compactness, $[X : X \cap Y] < \infty$; otherwise, one can find an infinite sequence (a_i) of elements of X such that $a_i a_j^{-1} \notin Y$ for $i \neq j$, but since these are \wedge -definable conditions, arbitrarily long sequences with the same property exist. So $X \cap Y$ has finitely many distinct cosets C_1, \dots, C_n in X . Note that $X \setminus C_i$ is \wedge -definable. Hence $C_j = \bigcap_{i \neq j} (X \setminus C_i)$ is \wedge -definable for each j . Since X_i and $C \setminus X_i$ are \wedge -definable, they are relatively definable in X . \square

A \mathbb{U} -definable set is A -definable iff it is $\text{Aut}(\mathbb{U}/A)$ -invariant. The same is true for \wedge -definable sets and for \vee -definable sets.

Types over \mathbb{U} are also called global types.

A sequence $(a_i : i \in \mathbb{N})$ of elements of \mathbb{U} is called *A -indiscernible* if for any order-preserving map $f : u \rightarrow u'$ between two finite subsets of \mathbb{N} , there exists an automorphism σ of \mathbb{U} fixing A , such that $\sigma(a_i) = a_{f(i)}$ for $i \in u$. The same applies to sequences of n -tuples. Using Ramsey's theorem and compactness, one shows that if $(b_i : i \in \mathbb{N})$ is any sequence, then there exists an indiscernible sequence $(a_i : i \in \mathbb{N})$ such that for any formula $\phi(x, y)$, if $\phi(b_i, b_j)$ holds for all $i < j$, then $\phi(a_i, a_j)$ holds for all $i < j$. A theorem of Morley's [39] asserts the same thing with the formulas $\phi(x, y)$ replaced by *types*, provided \mathbb{N} is replaced with a sufficiently large cardinal. For certain points (outside the main line), we will use Morley's theorem as follows. Let q be a global type, and construct a sequence a_i inductively, letting $A_i = \{a_j : j < i\}$, and choosing a_i such that $a_i \models q|A_i$. By

Morley’s theorem, there exists an indiscernible sequence (b_0, b_1, \dots) such that for any n , $b_n \models q_n \upharpoonright \{b_0, \dots, b_{n-1}\}$ for some $\text{Aut}(\mathbb{U})$ -conjugate q_n of q .

We will say in this situation that (b_0, b_1, \dots) are q -indiscernibles. The main case is that q is an invariant type, and then Morley’s theorem is not needed, for the original (a_j) are automatically indiscernible; see [42]. A global type finitely satisfiable in M is always M -invariant. In particular, given any type over M , this yields an M -indiscernible sequence (a_i) such that $tp(a_i/M \cup A_i)$ does not fork over M (cf. §2.1 for the definition). We remark that Morley’s theorem uses more infinite cardinals than the rest of the paper (namely, not only infinite sets but arbitrary countable iterations of the power set operation).

In all notation, if A is absent we take $A = \emptyset$. Generally a statement made for T_A over \emptyset is equivalent to the same statement for T over A , so no generality is lost.

We will occasionally refer to ultraproducts of a family M_i of L -structures. They are a specific way of constructing models M of the set of all sentences holding in all but finitely many M_i , and they have the saturation property (1). No other properties of ultraproducts will be needed.

2. INDEPENDENCE THEOREM

2.1. Stability. The material in this subsection is a presentation of [31], Lemma 3.3, here Lemma 2.3; compare also [47], §3, and the stability section in [1].

Let T be a first-order theory, \mathbb{U} a universal domain. One of the main lessons of stability is the usefulness of A -invariant types, meaning $\text{Aut}(\mathbb{U}/A)$ -invariant types. We note that if a global type p is finitely satisfiable in some $A \leq M$, then p is A -invariant; if a, a' are $\text{Aut}(\mathbb{U}/A)$ -conjugate, then $\phi(x, a) \& \neg\phi(x, a')$ cannot be satisfied in A .

We say that A is an *elementary submodel* of \mathbb{U} (written: $A \prec \mathbb{U}$) if any nonempty A -definable set has points in A . If $A \prec \mathbb{U}$, then any type over A extends to a global type, finitely satisfiable in A ([41]).

Consider two partial types $r(x, y), r'(x, y)$ over A . We say that r, r' are *stably separated* if there is no sequence $((a_i, b_i) : i \in \mathbb{N})$ such that $r(a_i, b_j)$ holds for $i < j$, and $r'(a_i, b_j)$ holds for $i > j$. Note that if arbitrarily long such sequences exist, then by compactness an infinite one exists, and in fact one can take the (a_i, b_i) to form an A -indiscernible sequence. Moreover r, r' are stably separated iff they contain formulas ϕ, ϕ' that are stably separated. By reversing the ordering one sees that stable separation is a symmetric property.

We say that r' is *equationally separated* from r if there is no sequence $((a_i, b_i) : i \in \mathbb{N})$ such that $r(a_i, b_i)$ holds for all i , and $r'(a_i, b_j)$ holds for $i < j$. This is an asymmetric condition, which implies stable separation: if stable separation fails, so that $r(a_i, b_j)$ holds for $i < j$, and $r'(a_i, b_j)$ holds for $i > j$ in some sequence (a_i, b_i) , the shifted subsequence (a_{2i}, b_{2i-1}) shows that equational separation fails too.

If r, r' are stably separated, then they are mutually inconsistent, since if $r(a, b)$ and $r'(a, b)$ we can let $a_i = a, b_j = b$. In stable theories, the converse holds.

Note that the set of stably separated pairs is open in the space S_2^2 of pairs of 2-types. Any extension of a stably separated pair to a larger base set remains stably separated.

A partial type $r'(x, b)$ is said to divide over A if there exists an indiscernible sequence b_0, b_1, \dots over A such that $\bigcup_i r'(x, b_i)$ is inconsistent, and $tp(b/A) =$

$tp(b_i/A)$. Equivalently, for some k , $\{r'(x, b_i) : i \in w\}$ is inconsistent for any k -element subset w of \mathbb{N} . By compactness, $r'(x, b)$ divides over A iff some formula $R(x, b) \in r'(x, b)$ divides over A . The ideal generated by all formulas that divide over A is called the *forking ideal*; thus $\phi(x, c)$ forks over A if it implies a disjunction of formulas that divide over A .

If $q = q(y)$ is a global type, we say that $r'(x, y)$ q -divides over A if for some n , if $b_i \models q|A(b_0, \dots, b_{i-1})$ for $i \leq n$, then $\bigcup_{i \leq n} r'(x, b_i)$ is inconsistent. This is equivalent to dividing, with the additional requirement that the indiscernible sequence be q -indiscernible.

Lemma 2.2. *Let r, r' be stably separated formulas over A . Let $q(y)$ be an A -invariant global type. Assume $r'(a, y) \in q$, $p = tp(a/A)$. Then $p(x) \cup r(x, y)$ q -divides over A .*

Proof. Suppose it does not.

Define a_1, \dots, c_1, \dots inductively: given $a_1, \dots, a_{n-1}, c_1, \dots, c_{n-1}$, choose c_n such that $c_n \models q|\{a_1, \dots, a_{n-1}, c_1, \dots, c_{n-1}\}$, and $a_n \models p$ chosen with $r(x, c_i)$ for $i < n$. The latter choice is possible since $p(x) \cup r(x, y)$ does not q -divide over A .

Then $r'(a_i, c_j)$ holds if $i < j$, but $r(a_i, c_j)$ holds when $i > j$. This contradicts the stable separation of r, r' . \square

We say that an A -invariant relation R is a *stable relation* over A if whenever $(a, b) \in R$ and $(a', b') \notin R$, $tp((a, b)/A)$ and $tp((a', b')/A)$ are stably separated. If R is stable, so is the complement of R ; but we are interested mostly in \wedge -definable R .

We will also encounter the condition of *equationality*. R is *equational* if whenever $(a, b) \in R$ and $(a', b') \notin R$, $tp((a', b')/A)$ is equationally separated from $tp((a, b)/A)$. As we have seen that equational separation implies stable separation, equationality implies stability.

We will say: “ $R(a, b)$ holds” for “ $(a, b) \in R$ ”. When q is a global type, write “ $R(a, y) \in q(y)$ ” to mean: $R(a, b)$ holds when $b \models q|A(a)$.

Lemma 2.3. *Let $p(x)$ be a type over A , and $q(y)$ be a global, A -invariant type. Let R be a stable relation over A .*

(1) *Assume $R(a, b)$ holds with $a \models p, b \models q|A(a)$. Then $R(a', b)$ holds whenever $a' \models p$ and $tp(a'/Ab)$ does not divide over A .*

(2) *Assume $tp(a/A) = tp(a'/A)$, $b \models q$, and neither $tp(a/Ab)$ nor $tp(a'/Ab)$ divides over A . Then $R(a, b)$ implies $R(a', b)$.*

(3) *Assume p too extends to a global, A -invariant type. Let $E = \{(a, b) : a \models p, b \models q|A\}$. Then the eight conditions:*

$R(a, b)$ holds for some/all pairs $(a, b) \in E$ such that $tp(a/A(b)) / tp(b/A(a))$ does not fork / divide over A

are all equivalent.

Proof. (1) Suppose $R(a', b)$ fails to hold. So $tp(a', b)$ and $tp(a, b)$ are stably separated, say by formulas r', r . By Lemma 2.2, since r holds for $b \models q|A(a)$, $r'(x, b) \cup p(x)$ divides, so $tp(a'/Ab)$ divides over A , a contradiction.

(2) Let R' be the complement of R ; it is also a stable relation. Let $c \models q|A(a)$. If $R(a, c)$ holds, then by (1) we have $R(a', b)$ and $R(a, b)$. If $R(a, c)$ holds, then similarly $R'(a', b)$ and $R'(a, b)$. In any case we have $R(a, b) \iff R(a', b)$, so the stated implication holds.

(3) Let E' be the set of pairs $(a, b) \in E$ such that $tp(a/A(b))$ does not divide over A , and E'' the set of pairs $(a, b) \in E$ such that $tp(b/A(a))$ does not divide over A . The equivalence between the four conditions for $tp(a/A(b))$ follows from (2): if $R(a, b)$ holds for some pair such that $tp(a/A(b))$ does not fork, then in particular it holds for a pair in E' (the same pair); by (2), it holds for all such pairs; hence certainly for all pairs for which $tp(a/A(b))$ does not fork over A .

Thus a single truth value for R is associated with pairs $(a, b) \in E'$. Similarly, as the conditions are symmetric, a single truth value for R is associated with pairs $(a, b) \in E''$. It remains to show that these truth values are equal. Replacing R by its complement if necessary, we may assume $R(a, b)$ holds in the situation of (1), where $b \models q|A(a)$. In particular $tp(b/A(a))$ does not fork over A ; so $R(a', b')$ holds for all $(a', b') \in E''$. But (1) asserts that $R(a', b)$ holds for all $(a', b) \in E'$. Hence R holds for all pairs in $E' \cup E''$. \square

Nondividing in Lemma 2.3 (3) can be replaced by any stronger condition; non-forking was mentioned above; we will later use smaller ideals.

Remark 2.4. Let p, q, R be as in Lemma 2.3, with $R(x, y)$ equational. Let $Q = \{b : b \models q|A\}$, $P = \{a : a \models p|A\}$. If $R(a, b)$ holds with $a \models p, b \models q|A(a)$, then $P \times Q \subset R$.

Lemma 2.5. *Let $S = S_z^{nf}$ be the set of global types that do not fork over \emptyset . Define an equivalence relation $E = E_{st}$ on S : $pE_{st}p'$ iff for any stable invariant relation R , and any b , we have $R(b, z) \in p \iff R(b, z) \in p'$. Then $|S/E| \leq 2^{|T|}$.*

Proof. Let M be a model. It suffices to show that if $p|M = p'|M$, then $pE_{st}p'$. Let $R(x, z)$ be a stable relation. Let $q = tp(b/M)$, and let q^* be any M -invariant global type extending q . Let $c \models p|M$. By Lemma 2.3, since p, p' do not fork over M , $R(b, z) \in p$ iff $R(x, c) \in q^*$ iff $R(b, z) \in p'$. \square

2.6. Making measures definable. A Keisler measure μ_x is a finitely additive real-valued probability measure on the formulas (or definable sets) $\phi(x)$ over the universal domain U . See [25].

We say μ is A -invariant if for any formula $\phi(x, y)$, for some function $g : S_y(A) \rightarrow \mathbb{R}$, we have $\mu(\phi(x, b)) = g(tp(b/A))$ for all b . If in addition g is continuous, we say that μ is an A -definable measure.

Let M_i be a family of finite L -structures. We wish to expand L to a richer language $L[\mu]$, such that each $L[\mu]$ structure admits a canonical definable measure μ . For each formula $\phi(x, y)$ and $\alpha \in \mathbb{Q}$ we introduce a formula $\theta(y) = (Q_\alpha x)\phi(x, y)$ whose intended interpretation is: $\theta(b)$ holds iff $\mu_x \phi(x, b) \leq \alpha$. If we wish μ to measure new formulas as well as L -formulas, this can be iterated.

We can expand each M_i canonically to $L[\mu]$, interpreting the formulas $(Q_\alpha x)\phi(x, y)$ recursively using the counting measure.

Let N be any model of the set of sentences true in all M_i (such as the ultra-product of the M_i with respect to some ultrafilter). Define $\mu\phi(x, b) = \inf\{\alpha \in \mathbb{Q} : (Q_\alpha x)\phi(x, b)\}$. Then μ is a Keisler measure. The formulas $(Q_\alpha x)\phi$ may not have their intended interpretation with respect to μ exactly, but very nearly so: $(Q_\alpha x)\phi(x, b)$ implies $\mu_x \phi(x, b) \leq \alpha$, and is implied by $\mu_x \phi(x, b) < \alpha$. Thus μ is a definable measure on N .

We will actually only use the corollary that the 0-ideal of μ is an invariant ideal; see below.

2.7. Ideals. Let X be a definable set, over A .

$L_X(\mathbb{U})$ denotes the Boolean algebra of \mathbb{U} -definable subsets of X . An ideal I of this Boolean algebra is A -invariant if it is $\text{Aut}(\mathbb{U}/A)$ -invariant; equivalently I is a collection of formulas of the form $\{\phi(x, a) : tp(a/A) \in E_\phi\}$, where for each $\phi(x, y)$, E_ϕ is a subset of $S_y(A)$, and $\phi(x, a)$ implies $x \in X$. To emphasize the variable, we use the notation I_x .

We say that I is \wedge -definable if for any $\theta(x, y)$, the set $\{b : \theta(x, b) \in I\}$ is \wedge -definable, and similarly for \vee -definable.

We say that a partial type Q over A is I -wide if it implies no formula in I .

In case X is \vee -definable, i.e. a countable union of A -definable sets $X = \bigvee_i X_i$, we let $L_x(\mathbb{U}) = \bigcup_i L_{X_i}(\mathbb{U})$. An ideal of $L_x(\mathbb{U})$ is a subset I such that $I \cap L_{X_i}(\mathbb{U})$ is an ideal for each i ; it is called A -invariant, \wedge -definable or \vee -definable if $I \cap L_{X_i}(\mathbb{U})$ has the corresponding property, for each i .

By analogy with measures, we will sometimes denote ideals in a variable x by μ , and write $\mu(\phi) = 0$ for $\phi \in \mu$, and $\mu(\phi) > 0$ for $\phi \notin \mu$.

The following definition is the defining property of S1-rank, [23], relativized to an arbitrary ideal (so within a definable set of finite S1-rank, the definable sets of smaller S1-rank form an S1-ideal). The terms *invariant*, *formula*, *indiscernible* are understood over some fixed base set A .

Definition 2.8. An invariant ideal $I = I_x$ on X is S1 if for any formula $D(x, y)$ and indiscernible $(a_i : i \in \mathbb{N})$ with $D(x, a_i) \in L_X(\mathbb{U})$, if $D(x, a_i) \cap D(x, a_j) \in I$ for $i \neq j$, then some $D(x, a_i) \in I$.

The forking ideal is contained in any S1-ideal:

Lemma 2.9. *Let I be an invariant S1 ideal over A . If $\phi(x, b)$ forks over A , then $\phi(x, b) \in I$.*

Proof. It suffices to show that if $\phi(x, b)$ divides over A , then $\phi(x, b) \in I$. Let (b_i) be an A -indiscernible sequence, with $\{\phi(x, b_i)\}$ inconsistent; so for some k , $\phi(x, b_1) \wedge \dots \wedge \phi(x, b_k) = \emptyset$. If $\phi(x, b_1) \in I$ we are done. Otherwise let m be maximal such that $\phi(x, b_1) \wedge \dots \wedge \phi(x, b_m) \notin I$. Let $c_i = (b_1, \dots, b_{m-1}, b_{m+i})$, and let $\psi(x, c_i) = \phi(x, b_1) \wedge \dots \wedge \phi(x, b_{m-1}) \wedge \phi(x, b_{m+i})$. Then the intersection of any two $\psi(x, c_i)$ is in I , but no $\psi(x, c_i)$ is in I . This contradicts Definition 2.8. \square

The forking ideal over A is also invariant under all A -definable bijections; in particular, for subsets of a group G under left and right translations by elements of $G(A)$, i.e. by elements of G definable over A . This will not be of real use to us, however, as we will be interested in translation invariance, right and left, by elements not necessarily defined over A .

The following is a fundamental observation from [8], [23], and [31]:

Lemma 2.10. *Let I_z be an invariant S1-ideal. Let $P = P(x, z), Q = Q(y, z)$ be formulas. Define:*

$$R(a, b) \iff (P(a, z) \wedge Q(b, z)) \in I_z.$$

Then R is a stable invariant relation.

Proof. We show indeed that R is equational: if $R(a_i, b_j)$ holds for $i < j$, where $(a_i, b_i)_i$ is indiscernible, then $R(a_i, b_i)$ holds too.

Otherwise, let $C_i = \{z : P(a_i, z) \wedge Q(b_i, z)\}$. Then $C_i \notin I_z$, but $\mu_z(C_i \cap C_j) = 0$. This contradicts the S1 property of Definition 2.8. \square

Example 2.11. Let $\mu(z)$ be a Keisler measure on \mathbb{U} -definable subsets of a set Z , with $\mu(Z) = 1$. Let $e \in \mathbb{N}, \epsilon = 1/e > 0$. Let $\phi(x, z), \phi'(y, z)$ be formulas, and write $D(a, b) = \{z \in Z : \phi(a, z) \cap \phi'(b, z)\}$. Let $r(x, y), r'(x, y)$ be formulas such that if $r(a, b)$, then $\mu(D(a, b)) \geq \epsilon$, while $r'(a, b)$ implies $\mu(D(a, b)) < \epsilon^2/2$. Then r, r' are stably separated; indeed r is equationally separated from r' . For suppose $r'(a_i, b_j)$ holds for $i = 1, \dots, 2e$. Let $D_i = D(a_i, b_i)$. Then $\mu(D_i) \geq \epsilon$, but $\mu(\bigcup_{1 \leq i < j \leq 2e} D_i \cap D_j) < (2e(2e - 1)/2)(\epsilon^2/2) < 1$. So $\mu(\bigcup_i D_i) > 2e\epsilon - 1 = 1$, a contradiction.

Example 2.12. Let μ be an $\text{Aut}(\mathbb{U}/A)$ -invariant, real-valued, finitely additive measure on \mathbb{U} -definable sets. Then $I = \{\phi(x, b) : \mu(\phi(x, b)) = 0\}$ is an $\text{Aut}(\mathbb{U}/A)$ -invariant S1-ideal. It is \wedge -definable if μ is definable.

Example 2.13. Let X have nonstandard finite size α , and let I be the ideal of all definable sets with nonstandard size β , where $\log(\beta) \leq (1 - \epsilon) \log(\alpha)$ for some standard $\epsilon > 0$. (See §5 for detailed definitions.) Then I is a \vee -definable ideal. It is not S1; but the counterexamples are always families contained in a definable set of dimension $< \epsilon \log(\alpha)$ for each $\epsilon > 0$.

2.14. Wide global types. We now note the existence of useful global types relative to an ideal I , in three slightly different situations. The combinatorial applications of the present paper can be deduced from either Lemma 2.16 or Lemma 2.17; the former has a shorter, more general, but much more impredicative proof.

Lemma 2.15. *Let $I = I(x)$ be a \vee -definable ideal, defined over a model M . Then there exists a global type p , finitely satisfiable in M , such that if $b \models p|M, a \models p|M(b)$, then $tp(b/Ma)$ is I -wide. (In fact, whenever p is finitely satisfiable in M and $p|M$ is wide, then p has this property.)*

Proof. Let p_0 be any wide type over M , and let p be any extension to \mathbb{U} , finitely satisfiable in M . Let $b \models p|M, a \models p|M(b)$. If $tp(b/Ma)$ is not wide, then for some $\phi(x, y)$ we have $\phi(a, b)$ and $\phi(a, y) \in I$; by \vee -definability, for some $\theta \in tp(a/M)$, for all a' with $\theta(a')$, $\phi(a', y) \in I$. Since $tp(a/Mb)$ is finitely satisfiable in M , there exists $a' \in M$ with $\theta(a')$ and $\phi(a', b)$. It follows that $p_0 = tp(b/M)$ is not wide, a contradiction. \square

Lemma 2.16. *Let $I = I(x)$ be an A -invariant ideal. There exists a model $M \geq A$, a global M -invariant type q , finitely satisfiable in M , such that if $a \models q|M$ and $b \models q|M(a)$, then $tp(a/M(b))$ is wide.*

Proof. Let T_{sk} be a Skolemization of the theory, in an expansion L_{sk} of the language L ; so the L_{sk} -substructure $M(X)$ generated by a set X is an elementary submodel. Define a sequence of elements a_i ($i < \omega_1$), and sets $A_i = M(\{a_j : j < i\})$, with $tp_L(a_i/A_i)$ wide. By Morley's theorem [39], there exists an indiscernible sequence $(c_i : i < \omega + 2)$ such that for any n , for some $i_1 < \dots < i_n$, $tp(c_1, \dots, c_n) = tp(a_{i_1}, \dots, a_{i_n})$. In particular, $tp(c_i/\{c_j : j < i\})$ is wide. Let U be an ultrafilter on \mathbb{N} , and let q be the set of formulas $\phi(x)$ of $L(\mathbb{U})$ such that $\{i : \phi(c_i)\} \in U$. Let $M = A_\omega$. Then q is finitely satisfiable in M . Let $a = c_{\omega+1}, b = c_\omega$. Then $a \models q|M$ and $b \models q|M(a)$, and $tp(a/M(b))$ is wide. \square

Lemma 2.17. *Let $I = I(x)$ be an \wedge -definable ideal, defined over a model M with $L(M)$ countable. Assume (“Fubini”) there exists an ideal $I^2(x, y)$ on $L_{x,y}(M)$ such that: (i) if $\phi(a, y) \in I(y)$ whenever $tp(a/M)$ is I -wide, then $\phi \in I^2$; (ii) if*

$\phi(x, b) \in I(x)$ whenever $tp(b/M)$ is I -wide, then $\phi \in I^2$; (iii) if $\phi(x) \wedge \phi(y) \in I^2$, then $\phi \in I$.

Then there exists a global type p , finitely satisfiable in M , such that if $b \models p|_M$, $a \models p|M(b)$, then $tp(a/Mb)$ and $tp(b/Ma)$ are I -wide.

Proof. Let B be the Boolean algebra of formulas of M modulo I . We show that a generic ultrafilter p_0 on B (in the sense of Baire category) can be extended to a type satisfying the lemma.

Claim. Let $\phi_i(x, y)$ ($i = 1, 2, 3$) be a triple of formulas, and let $P(x) \in B \setminus I$. Assume

$$P(x) \wedge P(y) \vdash \bigvee_{i=1}^3 \phi_i(x, y).$$

Then for some $P' \in B \setminus I$ implying P , for any $a, b \in P'$, we have (*): $\phi_1(a, y) \notin I$ or $\phi_2(x, b) \notin I$ or $\phi_3(c, b)$ for some $c \in M$.

Proof. If $(P(x) \wedge \phi_3(c, x)) \notin I$ for some $c \in M$, we can let $P'(x) = P(x) \wedge \phi_3(c, x)$; then the third option in (*) is met. Otherwise, $(P(x) \wedge \phi_3(c, x)) \in I$ for all $c \in M$. It follows from the M - \wedge -definability of I that $(P(x) \wedge \phi_3(c, x)) \in I$ for all c . So $P(y) \wedge \phi_3(x, y) \in I^2$.

If for some $P' \in B \setminus I$ implying P we have: $P'(a)$ implies $\phi_1(a, x) \notin I$, then the first disjunct of (*) holds. Otherwise, using the M - \wedge -definability of I , we see that for all $a \in P$ with $tp(a/M)$ I -wide, $\phi_1(a, y) \in I$. By the Fubini assumption (i), $P(x) \wedge \phi_1(x, y) \in I^2$.

Similarly, if for some such P' , $P'(b)$ implies $\phi_2(x, b) > 0$, then the second disjunct holds. Otherwise, by Fubini (ii), $(\phi_2(x, y) \wedge P(y)) \in I^2$.

Since $P(x) \wedge P(y)$ implies the disjunction of the ϕ_i , we have $(P(x) \wedge P(y)) \in I^2$; so $P \in I$, contradicting the choice of P , and proves the claim. \square

It is now easy to construct a type p_0 over M such that, for any $\phi_1(x, y), \phi_2(x, y), \phi_3(x, y)$, if $p_0(x) \cup p_0(y) \vdash \bigvee_{i=1}^3 \phi_i(x, y)$, then (*) of the claim holds for any $a, b \models p_0$. Namely, we let $p_0 = \{P_n\}$, where $P_n \in B \setminus I$ is constructed recursively. If n is even, we choose P_{n+1} so as to imply ψ or $\neg\psi$, where ψ is the $n/2$ -nd element of some enumeration of the formulas $\psi(x)$. If $n = 2m + 1$ is odd, consider the m th triple (ϕ_1, ϕ_2, ϕ_3) in some (infinitely repetitive) enumeration of all triples of formulas over M . If $P(x) \cup P(y) \vdash \bigvee_{i=1}^3 \phi_i$, let P' be as in the claim, and let $P_{n+1} = P_n \wedge P'$.

Let $b \models p_0$, and let $\Gamma(x, b) = p_0(x) \cup \{\neg\phi_1(x, b) : \phi_1(x, b) \in I\} \cup \{\neg\phi_2(x, b) : \phi_2(b, x) \in I\} \cup \{\neg\phi_3(x, b) : (\forall c' \in M)(\phi_3(c', x) \notin p_0)\}$. If $\Gamma(x, b)$ is inconsistent, then $p_0(x) \cup p_0(y) \vdash \phi_1(x, b) \vee \phi_2(x, b) \vee \phi_3(x, b)$ for some ϕ_1, ϕ_2, ϕ_3 with $\phi_1(x, b) \in I, \phi_2(b, x) \in I, \phi_3$ such that $(\forall c' \in M)(\phi_3(c', x) \notin p_0)$. But this contradicts the construction of p_0 . Thus $\Gamma(x, b)$ is consistent, and in view of the formulas $\neg\phi_3$, finitely satisfiable in M . Let p be any extension of $\Gamma(x, b)$ to a global type finitely satisfiable in M . Let $b \models p|M, a \models p|M(b)$. Then $tp(a/Mb)$ is wide because of the formulas $\neg\phi_1$, and $tp(b/Ma)$ is wide because of the formulas $\neg\phi_2$.

We now come to the 3-amalgamation statement. It says roughly that given a triangle of types, an arbitrary replacement of one edge by another with the same vertices will not affect the wideness of the opposite vertex over the edge. To simplify notation we work over $A = \emptyset$, so “divides” means “divides over \emptyset ”.

Theorem 2.18. *Let $\mu = \mu_z$ be an invariant $S1$ -ideal. Assume $tp(c/a, b)$ is μ_z -wide, $tp(b/a)$ and $tp(b'/a)$ do not divide, $tp(a)$ extends to an invariant global type,*

and $tp(b) = tp(b')$. Then there exists c' with $tp(c'/a, b')$ wide, and $tp(c'b') = tp(cb), tp(c'a) = tp(ca)$.

Proof. Let $Q \in tp(cb), P \in tp(ca)$. By compactness, it suffices, for any such pair of formulas, to find c' with $tp(c'/a, b')$ wide, and $Q(c', b'), P(c', a)$. In other words it suffices to show that $\mu_z(Q(z, b') \wedge P(z, a)) > 0$.

Consider the relation $R(x, y)$ such that $R(d, e)$ holds iff $\mu_z(P(z, d) \wedge Q(z, e)) = 0$. By Lemma 2.10, it is a stable relation.

By assumption, $tp(b'/a)$ and $tp(b/a)$ do not divide. By Lemma 2.3, since $R(a, b)$ fails, $R(a, b')$ must fail too. Thus $\mu_z(P(a, z) \wedge Q(bz)) > 0$. \square

- Remark 2.19.*
- (1) The hypothesis that $tp(b/a)$ and $tp(b'/a)$ do not divide can be replaced by: $tp(b/a)$ and $tp(a/b')$ do not divide, using Lemma 2.3 (3).
 - (2) Over a model, the hypothesis that $tp(a)$ extends to an invariant global type holds automatically.
 - (3) If E is an \wedge -definable equivalence relation over A with boundedly many classes, and the class of a is *not* the unique wide class within $tp(a)$, then 3-amalgamation can fail; one cannot amalgamate a type $p(x, y)$ implying $\neg(xEy)$ with any types implying xEy, yEz . It is possible that this is the only obstruction, so that as in [31], Theorem 2.18 holds over any set A which is boundedly closed.

2.20. Complements. In the remainder of this section we mention a variant of Theorem 2.18 in a measured setting, bringing out the 3-amalgamation aspect, and discuss connections to NIP and to probability theory. None of this will be needed for the combinatorial applications of §§2-5.

An arbitrary triangle of 2-types cannot be expected to give a consistent 3-type, for instance since a definable linear ordering may be present; types including $x < y, y < z, z < x$ are obviously not consistent together. But in a measured setting, contrary to initial appearances, this obstruction has an effect only on a measure zero set.

Below, i ranges over elements of $\Upsilon := \{1, 2, 3\}$, while u ranges over subsets of Υ of size 2. Let x_i be a sort, and let X_i be the space of types in this sort, over a fixed base set M . We assume every type in X_i extends to an invariant type (as is the case over an elementary submodel). We also assume, for simplicity's sake, that $L(M)$ is countable. For $i \in \Upsilon$ let μ_i be an M -definable measure on $X_i = X_{x_i}$. In fact it suffices to assume that μ_i is Borel-definable over M , meaning that $\mu_i(\phi(x, b))$ is a Borel function of $tp(b/M)$.

Assume the μ_i commute, in the sense that for any $i \neq j \in \Upsilon$, for any formula $\phi(x_i, x_j)$ over M ,

$$\int \mu_j(\phi(x_i, x_j))d\mu_i = \int \mu_i(\phi(x_i, x_j))d\mu_j;$$

see [27]. Any measures obtained as ultraproducts of counting measures will certainly have this property.

The common value is denoted by $\mu_{ij}(\phi)$; this defines a measure with variables (x_i, x_j) , referred to as the tensor product of μ_i, μ_j . Similarly, for $u \subseteq \Upsilon$, let μ_u be the tensor product measure on X_u . In particular we have $\mu = \mu_{123}$ on $X_{123} = X(\Upsilon)$.

We will occasionally refer to random elements; this can be given precise set-theoretic foundations, but we will not do this here. Instead we will understand by

this an element of a type space, or a product of type spaces, avoiding a certain countable collection of measure-zero Borel sets, which can be explicitly specified by inspecting the proof. We will also omit the foundational details of the notion of conditional measures, noting only that in the context of separable totally disconnected spaces we have a canonical countable Boolean algebra, namely the clopen subsets, making things easier.

Consider the natural maps $X(\Upsilon) \rightarrow X(\{23\})$, $X(\Upsilon) \rightarrow X(2) \times X(3)$, etc. For any such map, with target Y carrying measure μ_Y , and given a random (for the pushforward measure of μ) element $y \in Y$, we let $X_{123}(y)$ denote X_{123} with the measure $\mu_{123/y}$ conditioned on y . These conditional measures concentrate on the fiber over y , and satisfy: $\mu(B) = \int (\mu_{123/y}(B)) d\mu_Y$ for any clopen B . This formula defines $\mu_{123/y}$ uniquely for a random y , in the sense that any two choices will agree for almost all y . Again we refrain from giving the foundational details, noting only that they are much easier in the present context of separable totally disconnected spaces; this is due to the availability of a canonical countable Boolean algebra generating the measure algebra, namely the clopen subsets. See [19].

We will consider formulas θ_u in variables $(x_i : i \in u)$, and let $\theta = \bigwedge_{|u|=2} \theta_u$. We interpret θ on the one hand as a clopen subset of X_Υ , on the other hand as a clopen subset of $\prod_u X_u$, namely $\prod_u \theta_u$.

Lemma 2.21. *Let $(q_1, q_2, q_3) \in \prod_i X_i$ be a random triple. Let $q_{23} = tp(a_2 a_3 / M)$, where $tp(a_3 / M(a_2))$ does not divide over M , and $tp(a_i / M) = q_i$ ($i = 1, 2$). Let $\theta_{1j}(x_1, x_j)$ be a formula of positive measure for $X_{1j}(q_1, q_j)$ (the space X_{1j} with measure μ conditioned on (q_1, q_j)). Then $\theta_{12}(x_1, x_2) \wedge \theta_{13}(x_1, x_3) \cup q_{23}$ is consistent. In fact for $(a_2, a_3) \models q_{23}$, $\theta_{12}(x_1, a_2) \wedge \theta_{13}(x_1, a_3)$ has positive μ_1 -measure.*

Proof. Choose $p_{12} \in \theta_{12}$, random in $X_{12}(q_1, q_2)$ over (q_1, q_2, q_3) . Note that p_{12} extends q_1, q_2 . Since q_2 is random over (q_1, q_3) , p_{12} is random in $X_{12}(q_1)$ over (q_1, q_3) , and in X_{12} over (q_3) . Hence (q_3, p_{12}) are random in $X_3 \times X_{12}$.

Choose $p_{13} \in \theta_{13}$, random in $X_{13}(q_1, q_3)$ over (p_{12}, q_3) . Again p_{13} extends q_1, q_3 , and (as q_3 is random over (p_{12}) in X_3), p_{13} is random in $X_{13}(q_1)$ over (p_{12}) , so (p_{12}, p_{13}) is random in $X_{12}(q_1) \times X_{13}(q_1)$ over (q_1) . Now the product measure on $X_{12}(q_1) \times X_{13}(q_1)$ coincides with the pushforward measure from $X_{123}(q_1)$. (This is best seen “over q_1 ”.) So by choosing p_{123} at random in $X_{123}(p_{12}, p_{13})$ (with the conditional measure), we find p_{123} containing p_{12}, p_{13} and random. Let p_{23} be the restriction of p_{123} to the 2, 3-variables. Let $(b_2, b_3) \models p_{23}$. Note that p_{23} is random in X_{23} , so $tp(b_3 / M(b_2))$ does not divide over M .

Now $\theta_{12}(x_1, b_2) \wedge \theta_{13}(x_1, b_3)$ has positive μ_1 -measure (otherwise p_{123} could not be random). By Theorem 2.18, $\theta_{12}(x_1, a_2) \wedge \theta_{13}(x_1, a_3)$ has positive μ_1 -measure too. \square

Theorem 2.22. *Assume $L(M)$ is countable. Let $\Upsilon = \{1, 2, 3\}$. For $i \in \Upsilon$ let μ_i be an M -definable measure on $X_i = X_{x_i}$, and assume the μ_i commute. For $u \subseteq \Upsilon$, $|u| = 2$, let μ_u be the tensor product measure on X_u . Then there exist measure-one Borel subsets $\Omega_u \subset X_u$ and $\Omega \subset X_1 \times X_2 \times X_3$ with the following amalgamation property. Assume $q_u \in \Omega_u$, $(q_1, q_2, q_3) \in \Omega$, $q_u|_i = q_i$ for $i \in u$. Then there exists $q \in X_\Upsilon$, $q|_u = q_u$.*

In fact, we can take Ω_{23} to be the set of all $tp(bc)$ such that $tp(b/c)$ does not divide over M .

Proof. It suffices to show that if (q_1, q_2, q_3) is random, in $X_1 \times X_2 \times X_3$, q_u is random in X_u for $|u| = 2$, and $q_i \subseteq q_u$ for $i \in u$, then there exists $q \in X_{\Upsilon}$, $q|_u = q_u$. Fix such q_i, q_u . By compactness, it suffices to show for any given triple of formulas $\theta_u \in q_u$ that $\theta = \bigwedge_u \theta_u$ is consistent. Fix such a θ_u . Since q_{1j} is random in X_{1j} , it is random in $X_{1j}(q_1, q_j)$ over (q_1, q_j) . Hence θ_{1j} has positive measure in $X_{1j}(q_1, q_j)$. By Lemma 2.21, even $\theta_{12}(x_1, x_2) \wedge \theta_{13}(x_1, x_3) \cup q_{23}$ is consistent. \square

Note that since Ω_w has measure 1, for a random choice of $q_i \in X_i$ ($i = 1, 2, 3$), one expects the existence of $q_w \in S_w$ ($w \subseteq \{1, 2, 3\}$ with $|w| = 2$) with $q_i \subseteq q_w$ when $i \in w$. The (obviously necessary) hypothesis of compatibility on the q_w is therefore frequently attained.

Thanks to Pierre Simon for his comments on this. This result admits a more precise numerical version, or alternatively a formulation using ideals, and a higher-dimensional generalization; this and related issues will be taken up elsewhere.

2.23. NIP and de Finetti.

Example 2.24. Let μ be an A -definable Keisler measure in a NIP theory; cf. [26]. Let $\phi(x, y), \phi'(x, y)$ be formulas. For any real α , let $R_\alpha(a, b)$ denote the relation: $\mu(\phi(x, a) \cap \phi'(x, b)) < \alpha$. Then R_α is equational. This uses the fact that for an indiscernible sequence (c_j) over A we have $\mu(\psi(x, c_i) \cap \psi(x, c_j)) = \mu(\psi(x, c_i))$, applied to $c = (a, b)$, $\psi(x, c) = \phi(x, a) \wedge \phi'(x, b)$.

When $\alpha > 0$, the relation $\mu(\phi(x, a) \cap \phi'(x, b)) = \alpha$ need not be equational, as one sees for instance by taking $\phi = \phi'$ and an indiscernible sequence (a_i, b_i) with $a_i = b_i$.

However, in *any* theory, we have:

Proposition 2.25. *For any invariant measure ν , the relation $\nu(\phi(x, a) \cap \psi(x, b)) = \alpha$ is stable. In other words, when (a_i, b_i) is an indiscernible sequence of pairs, the function $(i, j) \mapsto \nu(\phi(x, a_i) \cap \psi(x, b_j))$ is symmetric in i, j .*

It follows that for any subset Y of $[0, 1]$, the relation: $\nu(\phi(x, a) \cap \psi(x, b)) \in Y$ is stable.

The proof is related to a classical theorem of de Finetti, classifying the so-called *exchangeable sequences* of random variables, i.e. sequences such that the action of the symmetric group does not change joint distributions. This was subsequently generalized by [22], [34]-[33], and in a different direction by Aldous and Hoover; see [28]. Thanks to Benjy Weiss for telling me about this theory. Though the assumption is classically stated as symmetry, indiscernibility suffices for the arguments; the proof below is essentially a subset of the one in [33] (in turn a modification of [22]). The higher-dimensional case will be considered elsewhere.

Proof of Proposition 2.25. We show more generally that if $(a_i : i \in \mathbb{N})$ is an indiscernible sequence, and ψ_1, \dots, ψ_k any formulas, then $\nu(\psi_1(x, a_1) \wedge \dots \wedge \psi_n(x, a_k))$ is invariant under the action of the symmetric group on $\{a_1, \dots, a_k\}$, i.e.

$$\mu(\psi_1(x, a_1) \cap \dots \cap \psi_k(x, a_k)) = \mu(\psi_1(x, a_{\sigma 1}) \cap \dots \cap \psi_k(x, a_{\sigma k}))$$

for any $\sigma \in \text{Sym}(k)$.

Let $B(\mathbb{N})$ be the Boolean algebra generated by the formulas $\psi_i(x, a_j)$ for $i \leq k, j \in \mathbb{N}$. Let $S = S(\mathbb{N})$ be the Stone space of $B(\mathbb{N})$. Let \mathcal{M} be the space of countably additive regular Borel probability measures on $S(\mathbb{N})$. For a finite $J \subseteq \mathbb{N}$,

let $B(J)$ be the subalgebra generated by the $\psi_i(x, a_j)$ with $j \in J$, $S(J)$ the Stone space, and for $\mu \in \mathcal{M}$, let $\mu|J$ be the induced measure, i.e. the pushforward of μ under the restriction map. Let \mathcal{M}_{ind} be the subset of *indiscernible* measures, i.e. measures μ on S such that for any finite $J_1, J_2 \subset \mathbb{N}$ of the same size, with order-preserving bijection $j : J_1 \rightarrow J_2$, the induced map $j : B(J_1) \rightarrow B(J_2)$ is measure-preserving, i.e. $j_*(\mu|J_1) = \mu|J_2$.

Let \mathcal{M}_{sym} be the apparently smaller subset of *symmetric* (or *exchangeable*) measures, where we demand that $j_*(\mu|J_1) = \mu|J_2$ for *any* bijection $j : J_1 \rightarrow J_2$.

Claim 1. $\mathcal{M}_{sym} = \mathcal{M}_{ind}$.

To prove the claim, note that both sets are convex and weak-* closed subsets of the unit ball of \mathcal{M} . Hence by Krein-Milman (cf. e.g. [57]), to show equality it suffices to prove that any extreme point of \mathcal{M}_{ind} is in \mathcal{M}_{sym} . So assume μ is an extreme point of \mathcal{M}_{ind} . Now Claim 1 follows from:

Claim 2. When $\mu \in \mathcal{M}_{ind}$ is extreme, we have independence: $\mu(\phi_1(x, a_1) \wedge \cdots \wedge \phi_n(x, a_n)) = \prod_{i=1}^n \mu(\phi_i(x, a_i))$, for any $\phi_i(x, a_i) \in B(\{a_i\})$.

Let $\alpha = \mu(\phi_1(x, a_1))$. If $\alpha = 0$, then $\mu(\phi_1(x, a_j)) = 0$ for any j , and both sides of the equation vanish. If $\alpha = 1$, then ϕ_1 may be deleted on both sides, and the claim follows by induction on k . Assume therefore that $0 < \alpha < 1$. Let μ' be obtained from μ by conditioning on $\phi_1(x, a_1)$, and shifting indices:

$$\mu'(\theta(x, a_1, \dots, a_m)) = \mu(\theta(x, a_2, \dots, a_{m+1}) \wedge \phi_1(x, a_1)) / \alpha.$$

Similarly, let μ'' be obtained from μ by conditioning on $\neg\phi_1(x, a_1)$, and shifting indices. Then $\mu = \alpha\mu' + (1 - \alpha)\mu''$, and $\mu', \mu'' \in \mathcal{M}_{ind}$. As μ is extreme, we have $\mu = \mu'$. This means:

$$\mu(\phi_1(x, a_1) \wedge \theta(x, a_2, \dots, a_m)) = \mu(\phi_1(x, a_1))\mu(\theta(x, a_2, \dots, a_m)).$$

Here m, θ are arbitrary. Claim 2 follows by induction on m , letting $\theta(x, a_2, \dots, a_n) = \phi_2(x, a_2) \wedge \cdots \wedge \phi_n(x, a_n)$.

Claim 1 follows easily: the right-hand side of the formula of Claim 2 is clearly symmetric. Any formula in $B(\mathbb{N})$ is a disjoint union of set-theoretic differences of conjunctions as considered in Claim 2. The measure of the difference of two such expressions can be computed using the inclusion-exclusion formula, and of disjoint unions by additivity.

Finally note that if ν is an invariant measure, indiscernibility of the (a_i) implies indiscernibility of $\nu|B(\mathbb{N})$; hence the proposition follows from Claim 1. \square

3. THE STABILIZER

Let G be a group, X a subset, defined over some model M_0 . Let \tilde{G} be the subgroup of G generated by X (cf. [25], §7). By a definable subset of \tilde{G} , we mean a definable subset of $(X \cup X^{-1})^{\leq n}$ for some n . A subset Y of \tilde{G} is *locally definable* if $Y \cap D$ is definable for every definable subset D of \tilde{G} .

Remark 3.1. In sections 3 and 4 we will never use G , only \tilde{G} . It is thus natural to use a many-sorted reduct, whose universes consist of the sets $(X \cup X^{-1})^{\leq n}$, with the inclusion maps and multiplication maps between them, and a distinguished predicate for X . We will speak of the inclusion maps as if they were actual inclusions.

Going further, we can note that we actually use only a bounded number of multiplications. In this section we will use only elements of $(XX^{-1})^3$ and will only use associativity for products of at most twelve elements of X and their inverses.

Hence the results of this section are valid for structures (X, X', G) with $X \subset X' \subset G$, with a binary map $m : (X')^2 \rightarrow G$ and an inversion map $^{-1} : X' \rightarrow X'$, such that products of up to twelve elements of $X \cup X^{-1}$ are defined, and independent of order. We will refer to this as a “local group” situation (cf. [16]). In this case \tilde{G} -translation invariance for a measure is replaced by the condition that μ measures X , and $\mu(Y) = \mu(Ya)$ for $Y \subseteq X^{-1}X$ and $a \in X^{-1}X$. To avoid too technical a language we will state the results using the Ind-definable group $\tilde{G} = \bigcup_n (X^{-1}X)^n$, indicating occasionally how to restrict to $(X^{-1}X)^3$. The reader is welcome to ignore these refinements at a first reading.

An \wedge -definable subset of $(X^{-1}X)^3$ closed under m and $^{-1}$ will be called an \wedge -definable subgroup of \tilde{G} (though in the local setting there is a priori no group of which it is a substructure). The main case is that of countable intersections; in this case one can write $H = \bigcap_{n \in \mathbb{N}} H_n$, with H_n definable, $H_n = H_n^{-1}$, and $H_n H_n \subseteq H_{n+1}$. It is easy to see that any \wedge -definable subgroup is an intersection of such countably- \wedge -definable subgroups. \tilde{G}/H is *bounded* if for any definable subset Y of \tilde{G} , the \wedge -definable equivalence relation: $y^{-1}y' \in H$ has boundedly many classes in the sense of §1.5. (Equivalently, if \tilde{G}, H are defined over M_0 , the cardinality of $\tilde{G}(N)/H(N)$ remains bounded when N runs over all elementary extensions of M_0 .)

Let \tilde{G} be generated by the definable set X . Let μ be an ideal on \tilde{G} , invariant under right translations by elements of X (i.e. $Z \in \mu$ iff $Zb \in \mu$, $b \in X$); equivalently, μ is invariant under right translations by elements of \tilde{G} . Assume $\mu(X) > 0$. Recall that a partial type Q is called *wide* if $\mu(Q') > 0$ for any definable $Q' \supseteq Q$.

A definable subset Z of \tilde{G} is called *right generic* if finitely many right translates of Z cover any given definable subset of \tilde{G} . If Z is right generic, then clearly $\mu(Z) > 0$. In the converse direction we have the observation, due to Ruzsa in the combinatorics literature, and Newelski in the model theory literature, that if $\mu(Z) > 0$, then $Z^{-1}Z$ is right generic. We state this as a lemma for later reference.

Lemma 3.2. *Let μ be an ideal on $\tilde{G} = \langle X \rangle$, invariant under right translations by elements of X , and with $\mu(X) > 0$. If Q is a wide partial type, then so is $Q^{-1}Q$. If Z is a definable set with $\mu(Z) > 0$, then $Z^{-1}Z$ is right generic.*

Proof. The statement for partial types follows by definition from the same statement for definable sets; so consider a definable set Z with $\mu(Z) > 0$. We have to show that $Z^{-1}Z$ is right generic, and wide.

Let $X_n = (XX^{-1})^n$; say $Z \subseteq (XX^{-1})^n$, and let $\{Za_i : i \in I\}$ be a maximal collection of pairwise disjoint subsets of Z , with $a_i \in X_n$. We claim that I is finite. Otherwise, by the usual Ramsey/compactness argument on existence of indiscernibles, one can find indiscernible $(a_i : i \in \mathbb{N})$ with $a_i \in X_n$ and $Za_m \cap Za_{m'} = \emptyset$ for $m \neq m'$; by the S1 property, since $\mu(Za_i) > 0$ for each i by right invariance, while $\mu(Za_i \cap Za_j) = 0$ for $i \neq j$, I must be finite. If $a \in X_n$, then $Za \cap Za_i \neq \emptyset$ for some i ; so $a \in Z^{-1}Za_i$. This shows that $Z^{-1}Z$ is right-generic.

In particular, $X \subseteq \bigcup_{b \in B} X^{-1}Xb$, for some finite B ; since $\mu(X) > 0$ it follows that $\mu(X^{-1}Xb) > 0$ for some $b \in B$, so $\mu(X^{-1}X) > 0$. □

In the local case, we say that $Z \subseteq (X^{-1}X)^2$ is right-generic if finitely many translates Zb ($b \in X$) cover $X^{-1}X$. Again if $Z \subseteq X^{-1}X$ has positive I -measure, then $Z^{-1}Z$ is right-generic.

Lemma 3.3 (cf. [26]). *Let H be an \wedge -definable subgroup of \tilde{G} . Then \tilde{G}/H is bounded iff every definable set containing H is right-generic. For any right invariant S1-ideal μ on \tilde{G} this is also equivalent to: H is wide.*

Proof. Consider $H = \bigcap H_n$ as above. If each H_n is generic, since \tilde{G} is a countable union of definable sets, there exists a countable set C_n such that $H_n C_n = \tilde{G}$. Let $C = \bigcup_n C_n$. Let $\pi : \tilde{G} \rightarrow \tilde{G}/H$ be the natural map. We say that a sequence u_n of elements of C converges to $uH \in \tilde{G}/H$ if for each m , for all sufficiently large n , we have $H_m u_n = H_m u$. Then each sequence has at most one limit, and each point of \tilde{G}/H is the limit of some sequence from C . Hence the cardinality of \tilde{G}/H is at most continuum. (We will later define the “logic topology” on \tilde{G}/H ; in this language we have just shown that it is separable.)

Conversely if \tilde{G}/H is bounded, let X be a definable subset of \tilde{G} . The condition: $H_{k+1}x \cap H_{k+1}y = \emptyset$ is a definable relation on (x, y) , since H_{k+1} is definable. Say \tilde{G}/H is bounded by λ ; then a fortiori there cannot be more than λ distinct (a_i) with $H_{k+1}a_i$ disjoint. Compactness applies, so any such family is finite. Let (a_i) be a maximal family $H_{k+1}a_i$ of disjoint cosets of H_{k+1} , with $a_i \in X$. Then there are finitely many elements a_i in the family, and it follows that $X \subseteq \bigcup_i H_{k+1}^{-1} a_i = \bigcup_i H_k a_i$, i.e. H_k is right-generic.

Given a right invariant S1-ideal μ , if H is wide, then there can be no infinite family of disjoint cosets of H_{k+1} , so as above H_k is generic. Conversely if H_k is generic, then $\mu(\bigcup_j H_k b_j) > 0$ for some finite set b_1, \dots, b_l , so $\mu(H_k b_j) > 0$, and by right invariance $\mu(H_k) > 0$. □

If an \wedge - A -definable subgroup of bounded index exists, then there is a *minimal* one; it is denoted \tilde{G}_A^{00} . For a discussion of the dependence on A , see [25].

Lemma 3.4. \tilde{G}_A^{00} is normal in \tilde{G} .

Proof. Let $H = \tilde{G}_A^{00}$. Then H has boundedly many \tilde{G} -conjugates; their intersection is an \wedge -definable normal subgroup N of \tilde{G} . On the face of it the definition of N requires additional parameters, but N is $\text{Aut}(\mathbb{U}/A)$ -invariant, and in general if an \wedge -definable set is invariant under $\text{Aut}(\mathbb{U}/A)$, then it is an infinite intersection of A -definable sets. □

Theorem 3.5. *Let M be a model, μ an M -invariant S1-ideal on definable subsets of \tilde{G} , invariant under (left or right) translations by elements of \tilde{G} . Let q be a wide type over M (contained in \tilde{G}). Assume:*

(F) *There exist two realizations a, b of q such that $tp(b/Ma)$ does not fork over M and $tp(a/Mb)$ does not fork over M .*

Then there exists a wide, \wedge -definable over M subgroup S of G . We have $S = (q^{-1}q)^2$; the set $qq^{-1}q$ is a coset of S . Moreover, S is normal in \tilde{G} , and $S \setminus q^{-1}q$ is contained in a union of nonwide M -definable sets.

Some remarks before turning to the proof.

- (1) It follows from the statement of the theorem that S can have no proper M - \wedge -definable subgroups of bounded index. For suppose such a subgroup T exists. Then q is contained in a bounded union of cosets of T . Being a complete type over a model, it is contained in a single coset. But then $q^{-1}qq^{-1}$, a coset of S , is contained in a coset of T ; so $S = T$.
- (2) The statement about $S \setminus q^{-1}q$ can be read to say that a random element of S lies in $q^{-1}q$; for instance when M is countable, and μ is the ideal of definable measure zero sets for some finitely additive measure μ on the Boolean algebra of M -definable sets, μ extends to a Borel measure on the space of types, and almost all types of elements of S lie in $q^{-1}q$.
- (3) When μ is the zero-ideal of a measure, note that translation invariance is assumed of the ideal, not of the measure. In particular, regardless of unimodularity, this assumption is true for Haar measures on a locally compact group.
- (4) (Weakening of left invariance). Most of the proof is devoted to showing that $S = (q^{-1}q)^2$ is a subgroup of \tilde{G} , and $qq^{-1}q$ is a coset of S . For this, left-translation invariance can be replaced with existence of an f -generic extension of q , in the sense of [26], i.e. the existence of an M -invariant ideal J containing the forking ideal, and with q wide for J . We will use such a J in Claims 3' and 5' (without assuming that $\mu = J$). The statement is essentially that left generics do not fork, and involves μ but not J .

The word “wide” will refer to μ unless explicitly qualified.

Normality of S will also follow under these assumptions, but we do not obtain the final statement about $S \setminus q^{-1}q$ in this case.

- (5) In place of any form of left translation invariance, we could use a stronger Fubini-type assumption on μ itself. (In Claim 3' of the theorem, we need to find (c_1, c_2, a) with $tp(c_i/M)$ specified, $c_i \in q^{-1}q$, and with $tp(a/M(c_1, c_2))$ wide. Given a version of Fubini we can achieve this by choosing a first, then c_1, c_2 .)
- (6) (Locality). Inspection of the proof will show that for all assertions except the normality of S , we only use μ (as an S1 ideal) on definable subsets of $XX^{-1}X$. To show normality of S , we also require XaX^{-1} , where $a \in X$ or $a \in X^{-1}$. Moreover the group structure is used only up to $(X^{-1}X)^3$. This is explicitly so everywhere except in Claim 5. There, note that $qc \subseteq XX^{-1}X$. Hence $qc \cap Y \subseteq XX^{-1}X$ for any set Y , and it makes sense to say that this intersection is wide. In the proof, by the time we use qab_1 , we know that ab_1 is in $q^{-1}q$.
It is also possible to combine (4) and (6); see Example 3.8.
- (7) The theorem implies that $S \subseteq X^{-1}XX^{-1}X$, or that for an appropriate translate $Y = a^{-1}X$, we have $S = YY^{-1}Y$. Example 6.1.10 of [8] shows that this cannot be improved to $S \subseteq X^{-1}X$.
- (8) An easy Löwenheim-Skolem argument shows that the theorem reduces to the case where the language is countable, and M is countable.
- (9) We show in fact that $S \setminus St_0(q)$ is contained in a union of nonwide M -definable sets, where $St_0(q) = \{s : qs \cap q \text{ is wide}\}$. If $s \in S$ is arbitrary now, and $tp(s'/M(s))$ is wide, then $tp(s's/M)$ is wide, so $s', s's \in St_0(q)$. Hence $s = (s')^{-1}(s's) \in St_0(q)^{-1}St_0(q) = St_0(q)^2$.

- (10) The assumption that M is a model, rather than just a substructure of the universal domain, is used via the consequence that any type extends to an invariant type; thus Theorem 2.18 applies to any type $tp(a)$. See Remark 2.19.
- (11) The proof uses both the nonforking ideal and the ideal of wide sets with respect to μ (or J). The former allows Theorem 2.18 to be used for an arbitrary type, since any type has a nonforking extension. On the other hand the ideal of wide sets, unlike the nonforking ideal, enjoys translation invariance.

Proof of Theorem 3.5. We also write q to denote $\{a : tp(a/M) = q\}$, and $q^{-1} = \{a^{-1} : tp(a/M) = q\}$.

Given two subsets X, Y of \tilde{G} , let

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

Let $Q = \{a^{-1}b : (a, b) \in q \times_{nf} q\}$. Let J be as in Remark (4) (or just set $J = \mu$ for the basic statement of the theorem), and set $Q' = \{a^{-1}b : a, b \in q, tp(b/Ma) \text{ is } J\text{-wide}\}$.

Note that qq^{-1} is obviously wide by right-invariance, and similarly $q^{-1}q$ is wide assuming left-invariance. If we wish to avoid the left-invariance assumption, but are willing to use μ on X^2 instead, then wideness of $q^{-1}q$ follows from Lemma 3.2.

Throughout this proof, we will use the fact (Lemma 2.10) that wideness of $qx \cap yy^{-1}$ is a stable relation between x and y . By Lemma 2.3, or Theorem 2.18, for any two types p_1, p_2 , this relation holds for one pair $(a_1, a_2) \in p_1 \times_{nf} p_2$ iff it holds for all pairs iff it holds for one or all pairs (a_2, a_1) in $p_2 \times_{nf} p_1$.

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

Proof. Let $a, b \in q$. Using (F), find $c \models q$ such that $tp(a/Mc)$ does not fork over M , and $tp(c/Ma)$ does not fork over M . By extending $tp(c/Ma)$ to a type over $M(a, b)$ and realizing this type, we may assume that $tp(c/Mab)$ does not fork over M . So we have $(b, c) \in q \times_{nf} q$, and $(c, a) \in q \times_{nf} q$. So $b^{-1}c, c^{-1}a \in Q$; hence $b^{-1}a \in QQ$. \square

Claim 2. For all $(a, b) \in q \times_{nf} q$, $qa^{-1} \cap qb^{-1}$ is wide.

Proof. By Theorem 2.18, it suffices to show that for *some* $(a, b) \in q \times_{nf} q$, $qa^{-1} \cap qb^{-1}$ is wide. Let a_1, a_2, \dots be an M -indiscernible sequence of elements of q , such that $tp(a_i/A \cup \{a_j : j < i\})$ does not fork over M . Then $(a_i, a_j) \in q \times_{nf} q$ for any $i < j$. It suffices to show that $qa_1^{-1} \cap qa_2^{-1}$ is wide; by compactness, for any definable set D containing q , it suffices to show that $\mu(Da_1^{-1} \cap Da_2^{-1}) > 0$. This is clear since μ is an S1-ideal, and by right-invariance, $\mu(Da_i^{-1}) > 0$. \square

Claim 3'. For all $(c_1, c_2) \in (q^{-1}q) \times_{nf} Q'$, $qc_1^{-1} \cap qc_2^{-1}$ is wide.

Proof. Let $p_i = tp(c_i/M)$. As in Claim 2, it suffices to see that $qc_1^{-1} \cap qc_2^{-1}$ is wide for some $(c_1, c_2) \in p_1 \times_{nf} p_2$. Let $a_0 \models q$. Then there exists $a_1 \in q$ with $tp(a_0^{-1}a_1/M) = p_1$. Since $c_2 \in Q'$, there exists a'_2 such that $r = tp(a'_2/M(a_0))$ is J -wide and $tp(a_0^{-1}a'_2/M) = p_2$; extend r to a J -wide type r' over $M(a_0, a_1)$, and let $a_2 \models r'$. We thus have $(a_0, a_1, a_2) \in (q \times q) \times_{nf} q$, with $tp(a_0^{-1}a_i/M) = p_i$ for $i = 1, 2$. Note also, using left invariance of J , that $tp(a_0^{-1}a_2/M(a_0, a_1))$ is J -wide, hence so is $tp(a_0^{-1}a_2/M(a_0^{-1}a_1))$, so it does not fork over M .

By Claim 2 we have that $qa_1^{-1} \cap qa_2^{-1}$ is wide. By the right-invariance of μ , $qa_1^{-1}a_0 \cap qa_2^{-1}a_0$ is wide. \square

Claim 3. For all $(c, d) \in (q^{-1}q) \times_{nf} Q$, $qc^{-1} \cap qd^{-1}$ is wide.

Proof. Let $d = a^{-1}b$, with $tp(b/M(a))$ wide for the forking ideal over M . We have to show that $qc^{-1} \cap qb^{-1}a$ is wide. By Theorem 2.18, it suffices to show this for *one* instance (c, b, a) with $tp(b, a)$ specified and such that $tp(b, a/M(c))$ does not divide over M . We may thus take $tp(a/M(c))$ to be a nonforking extension of $q = tp(a/M)$, and $tp(b/M(a, c))$ to be a nonforking over M extension of $tp(b/M(a))$. The latter is possible using the assumption that $tp(b/M(a))$ does not fork over M .

By right-invariance, we need to show that $qc^{-1}a^{-1} \cap qb^{-1}$ is wide. We apply Theorem 2.18 to the pair (a, b) (viewed as a single tuple) and c . So it suffices to show that $qc^{-1}a^{-1} \cap q(b')^{-1}$ is wide, where $tp(b/M) = tp(b'/M)$ and $tp(b'/M(a, c))$ is J -wide. By left-invariance of J , the type $tp(a^{-1}b'/M(a, c))$ is J -wide, and hence $tp(a^{-1}b'/M(c))$ is J -wide; so $tp(a^{-1}b'/M(c))$ does not fork over M . Also $tp(b'/M(a))$ is J -wide, so $a^{-1}b' \in Q'$. By Claim 3', $qc^{-1} \cap q(a^{-1}b')^{-1}$ is wide. By right-invariance, $qc^{-1}a^{-1} \cap q(b')^{-1}$ is wide, as required. \square

Claim 4. Let $(b, a) \in Q \times_{nf} q^{-1}q$. Then $ab \in q^{-1}q$. In fact $qa \cap qb^{-1}$ is wide.

Proof. We have $a^{-1} \in q^{-1}q$. Since M is a model, $tp(a^{-1}/M)$ extends to a global type r finitely satisfiable type in M ; so r is M -invariant. Use Lemma 2.3 (1) and Claim 3 to conclude that $qc^{-1} \cap qb^{-1}$ is wide if $c \models r|M(b)$. Now $tp(c/M(b))$ does not divide over M , so by Theorem 2.18, since $tp(a^{-1}/M(b))$ does not divide over M either, $qa \cap qb^{-1}$ is wide. In particular, for some $d, e \in q$ we have $da = eb^{-1}$. So $ab = d^{-1}e \in q^{-1}q$. \square

Claim 5. Let $a \in q^{-1}q$, $b_1, \dots, b_n \in Q$ and assume $tp(a/M(b_1, \dots, b_n))$ is wide. Then $ab_1 \cdots b_n \in q^{-1}q$. In fact $qa \cap q(b_1 \cdots b_n)^{-1}$ is wide.

Proof. Since $tp(a/Mb_1)$ is wide, it does not fork over M (Lemma 2.9). Hence by Claim 4 we have $ab_1 \in q^{-1}q$. By right-invariance of μ , $tp(ab_1/M(b_1, \dots, b_n))$ is wide, and in particular $tp(ab_1/M(b_2, \dots, b_n))$ is wide. By induction, $qab_1 \cap q(b_2 \cdots b_n)^{-1}$ is wide. Multiplying by b_1^{-1} on the right, $qa \cap q(b_1b_2 \cdots b_n)^{-1}$ is wide. Hence as in Claim 4, $ab_1 \cdots b_n \in q^{-1}q$. \square

In view of Theorem 2.18, Claim 5 is also valid assuming $tp(a/M)$ is wide, and $tp(a/M(b_1, \dots, b_n))$ does not fork over M . To show that $qq^{-1}q$ is a coset, we will later need a variant of Claim 5, proved in the same way:

Claim 5'. Let $a \in q^{-1}q$, $b_1, \dots, b_n \in Q$ and assume $tp(a^{-1}/M(b_1, \dots, b_n))$ is J -wide. Then $ab_1 \cdots b_n \in q^{-1}q$. In fact $qa \cap q(b_1 \cdots b_n)^{-1}$ is wide.

Proof. Since $tp(a^{-1}/Mb_1)$ is J -wide, it does not fork over M , and so $tp(a/Mb_1)$ does not fork over M . Hence by Claim 4 we have $ab_1 \in q^{-1}q$. By left-invariance of J , $tp((ab_1)^{-1}/M(b_1, \dots, b_n))$ is J -wide, and in particular $tp((ab_1)^{-1}/M(b_2, \dots, b_n))$ is J -wide. By induction, $qab_1 \cap q(b_2 \cdots b_n)^{-1}$ is wide. Multiplying by b_1^{-1} on the right, $qa \cap q(b_1b_2 \cdots b_n)^{-1}$ is wide. Hence as in Claim 4, $ab_1 \cdots b_n \in q^{-1}q$. \square

Claim 6. $Q^n \subset q^{-1}qq^{-1}q$.

Proof. Let $b_1, \dots, b_n \in Q$. Let $a \in q^{-1}q$ with $tp(a/M(b_1, \dots, b_n))$ wide. Then $ab_1 \cdots b_n \in q^{-1}q$, so $b_1 \cdots b_n = a^{-1}(ab_1 \cdots b_n) \in q^{-1}qq^{-1}q$. \square

It follows from Claim 1 that Q and $q^{-1}q$ generate the same subsemigroup, which is hence a group S . By Claim 6, this group is in fact equal to the \wedge -definable set $q^{-1}qq^{-1}q$.

Since $q^{-1}q \subseteq S$, we have $q \subseteq bS$ for any $b \in q$, and so $qq^{-1}q \subseteq bS$. Conversely, choose $b \in q$. Any element x of bS can be written $x = ba_1 \cdots a_4$ with $a_i \in Q$. Let $d \in q$ be such that $tp(d/M(a_1, \dots, a_4, b))$ is J -wide. Let $e = d^{-1}b$. Then $tp(e^{-1}/M(a_1, \dots, a_4, b))$ and hence $tp(e^{-1}/M(a_1, \dots, a_4))$ are J -wide. By Claim 5' we have $ea_1 \cdots a_4 \in q^{-1}q$. So $x = ba_1 \cdots a_4 \in dq^{-1}q \subseteq qq^{-1}q$. Thus $qq^{-1}q = bS$.

We know that S is an \wedge -definable group over M . I claim any \wedge -definable over M subgroup of S of bounded index must be equal to S . For let T be such a subgroup. We have $q^{-1}q \subseteq S$, so $q \subseteq aS$ for any $a \in q$. Thus q is contained in a left translate R of S ; we have $R = qS$, so R is defined over M . Now T acts on R on the right; the equivalence relation induced is \wedge -definable over M with boundedly many orbits. Since q is a complete type over M , it has an $\text{Aut}(\mathbb{U}/M)$ -invariant extension to \mathbb{U} ; this extension must pick a specific T -orbit cT , which is hence \wedge -definable over M ; by completeness again, as the realizations of q intersect cT , q is contained in cT . But then $q^{-1}q \subseteq T$; so $S \subseteq T$.

We know at this point that S has no proper \wedge -definable over M subgroups of bounded index. Let r be a type of elements of $X \cup X^{-1}$ over M . There cannot exist an unbounded family of cosets $a_i S$ with $a_i \in r$, for then the sets $a_i b q$ would also be disjoint for any $b \in q^{-1}$, so for some definable X' with $q \subset X' \subset X$ the sets $a_i b X'$ can be taken to be disjoint, contradicting the S1 property for μ within $r b X \subseteq (X \cup X^{-1})^3$. Thus r is contained in boundedly many left cosets of S , hence (being a complete type over a model) in one; call it C_r . So C_r is M -definable, and hence the conjugate group $S^r = C_r^{-1} S C_r$ is M -definable.

For any $c \in X \cup X^{-1} \cup \{1\}$, $r = tp(c)$, the image of qc in G/S is bounded. Otherwise there is a large collection of disjoint sets of the form $a_i c S$, with $a_i \in q$. Pick $b_0 \in q$; then $q^{-1}b_0 \subseteq S$. The sets $a_i c S b_0^{-1}$ are also disjoint, hence so are the $a_i c q^{-1}$. Thus there exists a definable $X' \subset X$ with $a_i c (X')^{-1}$ disjoint. So the sets $X c^{-1} a_i^{-1}$ are disjoint, and wide. But this contradicts the S1 property within $X c X^{-1}$. Thus qc/S is bounded. It follows that q is contained in boundedly many cosets of $c S c^{-1} = S^r$. So q is contained in a single coset $g S^r$. It follows that $q^{-1}q \subseteq S^r$, so $S \subseteq S^r$. Similarly $S \subseteq S^{r^{-1}}$, so $S^r \subseteq S$ and $S^r = S$. This shows that $X \cup X^{-1}$ normalizes S ; i.e., S is normal in \tilde{G} .

At this point we begin using left-invariance freely.

We argued above that $q^{-1}q$ is wide; in particular, S is wide. Q is also wide: suppose otherwise. So $Q \subseteq D$ for some definable D with $\mu(D) = 0$. Let $a \in q$. Then $a^{-1}q$ is wide. So $a^{-1}q \setminus D$ is wide. However $q \setminus aD$ forks over M , since if $b \in q \setminus aD$, then $a^{-1}b \notin Q$, so $tp(b/M(a))$ forks over M . Thus $D' \setminus aD$ lies in the forking ideal, for some definable D' containing q . By Lemma 2.9 we have $\mu(D' \setminus aD) = 0$; so $\mu(a^{-1}D' \setminus D) = 0$. It follows that $\mu(a^{-1}D') = 0$ and $\mu(D') = 0$, contradicting the wideness of q .

We finally show that S is contained in $q^{-1}q$ up to a union of nonwide definable sets. Let r be a wide type over M extending S ; we have to show that $r \subseteq q^{-1}q$. Pick $a_0 \in r$ and $c \in Q$ with $tp(c/M(a_0))$ wide. As $a_0^{-1} \in S$, we may write $a_0 = b_1 \cdots b_n$ with $b_i \in Q$; also as $c \in Q$ we have $c \in q^{-1}q$. Thus Claim 5 applies (with c playing the role of a of Claim 5), and we obtain that $qc \cap qa_0$ is wide. Choose $b_0 \in r$ with $tp(b_0/M(c))$ wide. In particular $tp(b_0/M(c))$ does not fork over M . By stability

of the relation and Lemma 2.3 (3), $qc \cap qb_0$ is wide too. Thus $b_0c^{-1} \in q^{-1}q$. Now $tp(b_0c^{-1}/M(c))$ is a right translate of $tp(b_0/M(c))$, so it is wide. By Claim 5 (or 3), $b_0 = (b_0c^{-1})c \in q^{-1}q$ (and $qb_0c^{-1} \cap qc^{-1}$ is wide; so $qb_0 \cap q$ is wide). So $r \subseteq q^{-1}q$, as required. In fact this shows that $r \subseteq St_0(q)$, in the notation of Remark (9). \square

Corollary 3.6. *Let μ be an invariant S1-ideal on definable subsets of \tilde{G} , invariant under translations by elements of \tilde{G} . Then there exists a model M and a wide, \wedge -definable over M subgroup S of G , with \tilde{G}/S bounded. For an appropriate complete type q over M we have $S = (q^{-1}q)^2$, and the complement $S \setminus q^{-1}q$ is contained in a union of nonwide M -definable sets.*

If μ satisfies the conditions of Lemma 2.17 over a model M_0 , or if μ is \vee -definable over M_0 , then one can take $M = M_0$.

Proof. Lemma 2.16 provides M and an M -invariant global type q^* such that if $q = q^*|M$, $a \models q|M$ and $b \models q^*|M(a)$, then $tp(a/M(b))$ is μ -wide. This implies (F). In case the assumptions of Lemma 2.17 or Lemma 2.15 hold, these lemmas provide a type over M_0 with (F) and so Theorem 3.5 applies with $M = M_0$. \square

Example 3.7. Consider the theory of divisible ordered Abelian groups $(G, +, <)$, or any \mathfrak{o} -minimal expansion, and let M be a model. We have a two-valued definable measure μ , assigning measure 0 to any bounded definable set. A two-valued invariant measure is always S1. The measure μ is translation invariant. Let q_A be the set of all measure-one M -definable formulas over A , $q = q_M$. If $a \models q_M$ and $b \models q_{M(a)}$, then $tp(a/Mb)$ does not fork over M since it is finitely satisfiable in M , and $tp(b/Ma)$ does not fork over M since it extends to an M -invariant type. Hence (F) of Theorem 3.5 holds. We can take $\tilde{G} = G$, $X = \{x \in G : x > 0\}$. The subgroup S is then G . Note that q^{-1} is *not* wide in this example.

Here is an example of the situation discussed in comments (4), (6), where the \wedge -definable group S is not normal.

Example 3.8. Consider the theory ACVF of algebraically closed valued fields, say of residue characteristic 0; the field of Puiseux series over \mathbb{C} is a model. Alternatively, let M be an ultraproduct of the p -adic fields \mathbb{Q}_p . Let K denote the valued field, \mathcal{O} the valuation ring, \mathcal{M} the maximal ideal. Let G be the semi-direct product of the additive group G_a with the multiplicative group G_m . So $G = TU$, where T, U are Abelian subgroups, $U = G_a$ normal, $T \cong G_m$. Let $t \in K$ be an element of valuation > 0 , and let g be the corresponding element of G , so that conjugation by g acts on U as multiplication by t . Let $U_0 = \{x \in K : \bigvee_{m \in \mathbb{N}} \text{val}(x) \geq -m \text{val}(g)\}$. View $\mathcal{O} \leq U_0$ as subgroups of U . Within G , let $X = g\mathcal{O}$. The group \tilde{G} generated by X is $g^{\mathbb{Z}}U_0$. Let p be a generic type of \mathcal{O} ; it avoids any coset of \mathcal{M} in \mathcal{O} . Let μ be the right-invariant ideal generated by $g\mathcal{M}$, and J the left-invariant ideal generated by $g\mathcal{M}$. These are not the same; notably \mathcal{O} is in J but not in μ . μ is not S1, but it is so when restricted to $X = XX^{-1}X$. Let $q = gp$. As in Remarks (4) and (6), the proof of Theorem 3.5 goes through to give a subgroup S , namely \mathcal{O} (it is definable in this case). But \mathcal{O} is not a normal subgroup of \tilde{G} .

Definition 3.9. We call X a *near-subgroup* of G if there exists an invariant S1-ideal μ on definable subsets of $(X \cup X^{-1})^3$, with $\mu(X) > 0$, and with $\mu(Y) = \mu(Y')$ whenever $Y, Y' \subseteq XX^{-1}X$ and $Y' = cY$ or $Y' = Yc$ for some c .

We will see in Corollary 3.11 that asking for μ defined on $\tilde{G} = \bigcup_n (X \cup X^{-1})^n$ would result in the same definition; in later sections we will work with this stronger definition.

Remark 3.10. Lou van den Dries has shown that a weaker condition suffices: $0 \in X$, and μ is defined on $XX^{-1}X$. Moreover the element c (which must by definition be in $(X \cup X^{-1})^6$) can in fact be chosen so that all products are taken within $XX^{-1}X$. This condition is essentially sharp, in view of Example 3.8. See [12].

When X is finite, any right-invariant measure must be proportional to the counting measure. Asymptotically, when (X, G) vary in some family, we have that every ultraproduct is a near-subgroup iff $|XX^{-1}X|/|X|$ is bounded in the family.

The following corollary of Theorem 3.5 is analogous to Lemma 3.4 of [50]; the point is that we do not assume a priori that $(X^{-1}X)^n$ has finite measure. The Fubini-type assumption on the ideal is much weaker here, but the conclusion is purely qualitative. We state the extension lemma for S1 ideals; a similar statement for measures is also valid, with a similar proof.

Corollary 3.11. *Let X be a near-subgroup of G . Then for any n , $(X^{-1}X)^n$ is contained in a finite union of right translates of $(X^{-1}X)^2$. μ extends to an invariant S1-ideal μ' on $\bigcup_n (X^{-1}X)^n$; μ' is the unique right-invariant ideal extending $\mu|(X^{-1}X)^2$.*

Proof. For this we may add parameters and work over a model. Let \tilde{G} be the group generated by X . By Theorem 3.5 and Remark (6) to that theorem, there exists a wide \wedge -definable normal subgroup S of \tilde{G} . The proof also shows that $S \subseteq (X^{-1}X)^2$ and that the image of X modulo S has bounded cardinality. Hence \tilde{G}/S is bounded, and in particular for any n , $(X^{-1}X)^n$ is contained in boundedly many cosets of S , and hence in boundedly many right translates of $(X^{-1}X)^2$. By compactness, finitely many right cosets of $(X^{-1}X)^2$ suffice to cover $(X^{-1}X)^n$. If D is a definable subset of $\bigcup_n (X^{-1}X)^n$, it follows that we can write $D = \bigcup_i D_i b_i$, where $D_i \subseteq (X^{-1}X)^2$ and $b_i \in (X^{-1}X)^{n+2}$.

Define μ' to be the collection of all definable sets $\bigcup_i D_i b_i$, where D_i is a definable subset of $(X^{-1}X)^2$, $b_i \in \bigcup_n (X^{-1}X)^n$, and $\mu(D_i) = 0$. This is clearly a right-invariant ideal whose restriction to a definable subset of $(X^{-1}X)^3$ is precisely μ . (If Y is a definable subset of $(X^{-1}X)^3$ and $Y = \bigcup_i D_i b_i$ as in the definition of μ' , then by invariance we have $\mu(D_i b_i) = 0$ for each i , so $\mu(Y) = 0$.)

If μ'' is any right-invariant ideal μ' extending μ , and if D is a definable subset of $(X^{-1}X)^n$, write $D = \bigcup_i D_i b_i$ where $D_i \subseteq (X^{-1}X)^2$ and $b_i \in (X^{-1}X)^{n+2}$. We have $\mu''(D) = 0$ iff $\mu''(D_i) = 0$ for each i iff $\mu(D_i) = 0$ for each i . This shows that $\mu'' = \mu'$.

To see that μ' is S1, it suffices to show for each n that the restriction to $(X^{-1}X)^n$ is S1. As above, write $(X^{-1}X)^n = \bigcup_{j=1}^l D_j b_j$. It is clear that any ideal on the union of finitely many sets must be S1 if the restriction to each of these sets is S1. So it suffices to show that $\mu'|_{(D_j b_j)}$ is S1 for each j . But $\mu'|_{(D_j b_j)}$ is isomorphic, via translation by b_j , to $\mu|_{D_j}$, which is S1. \square

This kind of characterization incidentally makes some functorialities evident, that are not so directly from the definition of a near-subgroup or an approximate subgroup; see Remark 4.10 (0),(2).

Given elements a_1, \dots, a_l and b_1, \dots, b_m of G , let $A_i = \{x^{-1}a_i x : x \in X\}$ be the set of X -conjugates of a_i , and let $W_n(a_1, \dots, a_l, b_1, \dots, b_m)$ be the set of words of length $\leq n$ in $A_1 \cup \dots \cup A_l \cup \{b_1, \dots, b_m\}$. Let $d(X; a_1, \dots, a_l)$ be the smallest integer n such that $X \subseteq W_n(a_1, \dots, a_l; b_1, \dots, b_l)$ for some $b_1, \dots, b_l \in X$, or ∞ if there is no such n .

Proposition 3.12. *For any $k, l, n \in \mathbb{N}$, for some $M, K \in \mathbb{N}$, the following holds:*

Let G be a group, X a finite subset. Assume $|XX^{-1}X| \leq k|X|$. Also assume that there exist $x_1, \dots, x_M \in X$ such that:

() for any $1 \leq i_0 < i_1 < \dots < i_l \leq M$, $d(X, x_{i_0}^{-1}x_{i_1}, \dots, x_{i_0}^{-1}x_{i_l}) \leq n$.*

Then there exists a subgroup S of G , $S \subseteq (X^{-1}X)^2$, such that X is contained in $\leq K$ cosets of S .

Proof. Fix k, l, n . Suppose there are no such M, K ; then there are groups G_M and $X = X_M \subset G_M$ such that there exist x_1, \dots, x_M with (*), and there is no subgroup S of G , $S \subseteq (X^{-1}X)^2$, such that X is contained in $\leq M$ cosets of S . Consider (G_M, X_M, \cdot) as a structure and enrich it using the Q_α -quantifiers for the normalized counting measure on X_M , as in §2.6. By compactness, there exists a countably saturated group G and a subset X such that there exists an infinite indiscernible sequence $x_1, x_2, \dots \in X$ such that

(i) (*) holds for any $1 \leq i_0 < i_1 < \dots < i_l < \infty$.

(ii) For any definable subgroup S of G with $S \subseteq (X^{-1}X)^2$, X is not contained in finitely many right translates of S .

Let \tilde{G} be the subgroup of G generated by X . By Theorem 3.5 there exists an \wedge -definable normal subgroup S of bounded index in \tilde{G} , with $S \subseteq (X^{-1}X)^2$. Since the sequence x_1, x_2, \dots is indiscernible and G/S is bounded, all x_i lie in the same coset of S . So the elements $y_i = x_i^{-1}x_i$ all lie in S . Now $d(X, y_1, \dots, y_l) \leq n$; so $X \subseteq W_n(y_1, \dots, y_l; b_1, \dots, b_l)$ for some $b_1, \dots, b_l \in X$. Let N be the normal subgroup of \tilde{G} generated by the y_i , and let \mathbf{X} be the image of X modulo N . Then $\mathbf{X} \subseteq W_n(1, \dots, 1; \bar{b}_1, \dots, \bar{b}_l)$, where $\bar{b}_i = b_i N$. Hence \mathbf{X} is finite. As $S \subseteq (X^{-1}X)^2$, it follows that the image of S modulo N is finite, i.e. $[S : N] < \infty$. Since N is \vee -definable, so is S . But S is \wedge -definable; so it is a definable group. Now $N \subseteq S$, so X is contained in finitely many translates of S , in contradiction to (ii). \square

By Ruzsa’s argument, the condition $|a_1^X \dots a_l^X| \geq |X|/m$ implies that $X \cup X^2$ is contained in the union of $\leq m$ translates $c_j a_1^X \dots a_l^X (a_l^{-1})^X \dots (a_1^{-1})^X$, with $c_j \in X$, so that $d(a_1, \dots, a_l; G) \leq \max(m, 2l)$. We can now deduce Corollary 1.2 from Proposition 3.12 using Ramsey’s theorem, but we will give a direct argument. We denote the l th Cartesian power of X by $X^{(l)}$.

Proof of Corollary 1.2. Fix k, l, m and suppose for contradiction that the conclusion fails. Then for arbitrarily large K , letting $p = 1 - 1/K$, there exists a group G_K and a finite subset $(X_0)_K, X = (X_0)_K^{-1}(X_0)_K$, such that the situation of Corollary 1.2 holds but no subgroup S of G_K with $S \subset X^2$ is contained in $\leq K$ cosets of S . Let μ_K be the counting measure on G , normalized so that $\mu_K(X) = 1$, and let μ_K^l be the counting measure on X^l , similarly normalized. Let Q_K be the set of l -tuples (a_1, \dots, a_l) such that $\mu_K(a_1^X \dots a_l^X) \geq 1/m$; then $\mu_K^l(Q_K) \geq p = 1 - 1/K$.

By compactness (as in the proof of Proposition 3.12, and in §2.6) there exists a structure including a group G , a definable subset $X = X_0^{-1}X_0 \subset G$, a definable measure μ on definable sets, with $1 = \mu(X) \leq k\mu(X_0) < \infty$. There is also as well a

definable measure μ^l on l -tuples, related to μ via Fubini's theorem; the only instance of Fubini we will need is the equality: $\mu^l(Y^l) = \mu(Y)^l$, holding for any definable Y . Further, there exists a definable set $Q \subseteq X^{(l)}$ such that if $(a_1, \dots, a_l) \in Q$, then $\mu(a_1^X \cdots a_l^X) \geq 1/m$, and $\mu^l(Q) \geq 1 - 1/K$ for any $K = 1, 2, \dots$, i.e. $\mu^l(Q) = 1$. Finally, for no definable group $S \subseteq XX$ is X contained in finitely many cosets of S .

We take G to be countably saturated. Recall that countable saturation means that any countable family of definable sets with the finite intersection property has nonempty intersection; we will actually need it for the family R_j below.

By Theorem 3.5 there exists an \bigwedge -definable normal subgroup S of bounded index in \tilde{G} . Find a countable set of definable (with parameters) equivalence relations E^j on X_0 , such that each E^j has finitely many classes, E^{j+1} refines E^j , and if $(a, b) \in E^j$ for each j , then $a^{-1}b \in S$. (For instance, say $S = \bigcap S_j$, and let C^j be a maximal subset such that $S_j x \cap S_j y = \emptyset$ for $x \neq y \in C^j$; define E^j so that $(x, y) \in E^j$ implies $\{c \in C^j : xS_j \cap cS_j = \emptyset\} = \{c \in C^j : yS_j \cap cS_j = \emptyset\}$. Alternatively note that if a, b have the same type over some countable model, then $a^{-1}b \in S$.)

Some class F_j of E^j has measure $\epsilon_j > 0$; so $\mu(F_j^{-1}F_j) \geq \epsilon_j > 0$. Thus $(F_j^{-1}F_j)^{(l)} \geq \epsilon_j^l$, and hence (as $\mu(Q) = \mu(X^l) = 1$) we have $\mu(Q \cap (F_j^{-1}F_j)^{(l)}) \geq \epsilon_j^l > 0$. Hence for each j there exist $(a_1, \dots, a_l) \in Q$ such that for each $i \leq l$, we have $a_i = b_i^{-1}c_i$ for some $(b_i, c_i) \in E^j$. As we took the E_j to refine each other, this holds for any finite set of indices j at once. In other words, the family of sets $\{R_j\}$ has the finite intersection property:

$$R_j = \{(a_1, b_1, c_1, \dots, a_l, b_l, c_l) : (a_1, \dots, a_l) \in Q, \bigwedge_{i \leq l} (a_i, b_i) \in E^j, \text{ and } a_i = b_i^{-1}c_i\}.$$

By countable saturation, $\bigcap_j R_j \neq \emptyset$; i.e., there exist $(a_1, \dots, a_l) \in Q$ and $b_1, c_1, \dots, b_l, c_l$ such that for each $i \leq l$ we have $a_i = b_i^{-1}c_i$ and $(b_i, c_i) \in \bigcap_j E^j$. By the choice of E^j , this implies $a_i \in S$.

Now S is normal in \tilde{G} , so $a_1^X \cdots a_l^X \subseteq S$. Since $\mu(a_1^X \cdots a_l^X) > 1/(m+1)$, it follows that S cannot have $\mu(X_0X)(m+1)$ disjoint cosets x_iS . So X_0/S is finite; it follows that XX/S is finite, so $XX = S \cup \bigcup_{\nu=1}^k (XX \cap c_\nu S)$ for some c_1, \dots, c_ν . Since S is \bigwedge -definable, so is each $c_i S$, and we see that the complement of S in XX is also \bigwedge -definable. When a subset of a definable set and its complement are both \bigwedge -definable, they are both definable. Hence S is a definable group. But finitely many cosets of S cover X . This contradiction proves the corollary. \square

Though we stated Proposition 3.12 for finite X , it holds with the same proof if the hypothesis $|XX^{-1}X| \leq k|X|$ is replaced by $\mu((X \cup X^{-1})^3) \leq k\mu(X)$, with μ an arbitrary right-invariant finitely additive measure on \tilde{G} , or even in the above sense on $(X \cup X^{-1})^3$.

4. NEAR-SUBGROUPS AND LIE GROUPS

Let $X \subseteq \tilde{G}$ be a near-subgroup with respect to an M -invariant, right-invariant ideal μ , as in the previous section.

Any compact neighborhood X in a Lie group L is (obviously) an approximate subgroup and a near-subgroup with respect to Haar measure. We will show that all

near-subgroups are related to these classical ones. We will use logical compactness to connect to the locally compact world, and then the Gleason-Yamabe structure theory for locally compact groups in order to find Lie groups.

4.1. Some preliminaries. We will require the following statement: every locally compact group G has an open subgroup G_1 which is isomorphic to a projective limit of Lie groups. (Gleason defines a topological group G to be a *generalized Lie group* if for every neighborhood U of the identity there is an open subgroup H of G and a compact normal subgroup C of H such that $C \subseteq U$ and H/C is a Lie group ([15], Definition 4.1). According to [55], Theorem 5', every locally compact group is a generalized Lie group. By [15], Lemma 4.5, if G is a generalized Lie group with connected component G^0 of the identity, and G/G^0 is compact, then G is a projective limit of Lie groups. Now G/G^0 is totally disconnected. So there exists an open subgroup G_1 of G containing G^0 , such that G_1/G^0 is compact. Hence G_1 is an open subgroup of G , and a projective limit of Lie groups.)

We will also use the fact that in a connected Lie group G , for any chain $C_1 \subset C_2 \subset \dots$ of compact normal subgroups, $cl(\bigcup_n C_n)$ is also a compact normal subgroup. Indeed the dimension of the Lie algebras of the C_n must stabilize, so they are locally equal, and hence the connected components C_n^0 stabilize. Factoring out the compact normal subgroup $\bigcup_n C_n^0$, we may assume the C_n are discrete, i.e. finite. Since G is connected, the C_n are contained in the center Z . The connected component Z^0 of Z has universal covering group \mathbb{R}^n , so $Z^0 \cong \mathbb{R}^k \oplus (\mathbb{R}/\mathbb{Z})^l$. The discrete group Z/Z^0 is a homomorphic image of the fundamental group of G/Z , hence is finitely generated. It has a finite torsion part A/Z^0 ; since Z^0 is divisible, A can be written as a direct sum $A_0 \oplus Z^0$. It is clear that the torsion points of Z , and hence all the C_n , are contained in the compact central subgroup $A_0 \oplus (\mathbb{R}/\mathbb{Z})^l$.

In particular any closed subgroup contains a unique maximal compact normal subgroup of G .

Further down (Lemma 6.6), we will also need to know that a compact Lie group has no infinite descending sequences of closed subgroups; this follows easily along the same lines.

The results of this section will also be valid for local groups, using the following local version of Gleason-Yamabe due to Goldbring: for a compact local group G there exist a continuous map $h : D \rightarrow L$ into a Lie group L , whose domain $D = D^{-1}$ is a smaller compact neighborhood of 1 in G , and whose image hD is a compact neighborhood of 1 in L , such that xy is defined for any $x, y \in D$, and we have: $xy \in D$ iff $h(x)h(y) \in hD$, in which case $h(xy) = h(x)h(y)$. [16] has generalized the “no-small-subgroups” theory to the local group setting; to apply it one needs to know that some neighborhood of $1 \in G$ contains a compact normal subgroup, such that the quotient has no small subgroups. This is Lemma 9.3 of [16].

Recall that we call two subsets X, X' of a group *commensurable* if each one is contained in finitely many right translates of the other. If H, H' are subgroups, and H is contained in finitely many cosets of H' , then it is contained in the same number of cosets of $H \cap H'$, so $[H : H \cap H'] < \infty$; thus for groups this coincides with the usual notion.

Theorem 4.2. *Let X be a near-subgroup of G , generating a group \tilde{G} . Then there exists a \vee -definable subgroup \check{G} contained in \tilde{G} , an \wedge -definable subgroup $K \subseteq \check{G}$, a connected, finite-dimensional Lie group L , with no nontrivial normal compact*

subgroups, and a homomorphism $h : \check{G} \rightarrow L$ with kernel K and dense image, with the following property:

If $F \subseteq F' \subseteq L$ with F compact and F' open, then there exists a definable D with $h^{-1}(F) \subseteq D \subseteq h^{-1}(F')$. Any such D is commensurable to $X^{-1}X$.

\check{G} and K are defined without parameters. The Lie group L is uniquely determined.

Let us bring out some facts implicit in the statement of the theorem (and also visible directly in the proof).

Remark 4.3. • If (G', X') is a countably saturated elementary extension of (G, X) , then h extends to $h' : \check{G}' \rightarrow L$, and h' is surjective.

- The Lie group L is determined up to isomorphism by (\tilde{G}, \cdot, X) , where \tilde{G} is the subgroup of G generated by X , in fact by the theory of (\tilde{G}, \cdot, X) , with \tilde{G} viewed as many-sorted. We call it the associated Lie group.
- Since any compact subset of a Lie group is a countable intersection of open sets, it follows that if $W \subseteq L$ is compact, then $h^{-1}(W)$ is \bigwedge -definable.
- Similarly, if $W \subseteq L$ is open, then $h^{-1}(W)$ is \bigvee -definable.
- If $W \subseteq L$ is a neighborhood of 1, then $h^{-1}(W)$ contains a definable set of the form $U^{-1}U$, with U a definable subset of (G, X) contained in \check{G} and commensurable to $X^{-1}X$.
- Any definable set containing K contains some $h^{-1}(W)$, with W a neighborhood of 1 in L .
- If L is trivial, taking $F = F' = L$ in the statement of the theorem we see that \check{G} is a definable group, commensurable to $X^{-1}X$.
- We have $K \subseteq (XX^{-1})^m$ for some m . Theorem 3.5 provides an \bigwedge -definable stabilizer contained in $(X^{-1}X)^2$, but converting it to a 0-definable one involves some (finite) enlargement.

We first show that the main statement of Theorem 4.2 holds after saturation and base change.

Lemma 4.4. *Let X be a near-subgroup of G , generating a group \tilde{G} . Assume the structure (G, X, \dots) is countably saturated. Then over parameters there exists a \bigvee -definable subgroup \check{G} contained in \tilde{G} , an \bigwedge -definable subgroup $K \subseteq \check{G}$, a connected, finite-dimensional Lie group L and a homomorphism $h : \check{G} \rightarrow L$ with kernel K and dense image, with the following property:*

If $F \subseteq F' \subseteq L$ with F compact and F' open, then there exists a definable D with $h^{-1}(F) \subseteq D \subseteq h^{-1}(F')$. Any such D is commensurable to $X^{-1}X$.

Proof. Let \tilde{G} be the subgroup of G generated by X ; let $S_0 = XX^{-1}$. Theorem 3.5 (via Corollary 3.6) provides definable subsets $S_n \subseteq (XX^{-1})^2$ of \tilde{G} such that $S = \bigcap_{n \in \mathbb{N}} S_n$ is a normal subgroup of \tilde{G} of bounded index; we may take $S_{n+1} = S_n^{-1}$ and $S_{n+1}S_{n+1} \subseteq S_n$. Define S_n for negative n too by $S_n = S_{n+1}S_{n+1}$. So S_0 is 0-definable, and $\bigcup S_n = \tilde{G}$.

We define a topology on \tilde{G}/S using the quotient map $h : \tilde{G} \rightarrow \tilde{G}/S$ by:

(*) $W \subseteq \tilde{G}/S$ is closed iff $h^{-1}(W) \cap S_n$ is \bigwedge -definable for each n .

See [25], Section 7 for a more detailed description. Let $L_0 = \tilde{G}/S$. This is easily seen to be a locally compact topological group. Compactness is an immediate consequence of saturation and logical compactness: an intersection of a small number

of \wedge -definable subsets can never be empty, unless a finite subintersection is empty. Continuity of the group operations follows from the definability of the group structure on G . The images of the sets S_n form a neighborhood basis for the identity of L_0 as noted below, so that L_0 is Hausdorff.

Let $\pi_S : \tilde{G} \rightarrow \tilde{G}/S$ be the projection. Note that $\pi_S^{-1}\pi_S(S_n) = \bigcap_m S_n S_m \subseteq S_n S_n$. In particular $\pi_S^{-1}\pi_S(S_n)$ is contained in a definable subset of \tilde{G} . In fact for any definable set $D \subseteq S_n$, $\pi_S^{-1}\pi_S(D) = SD$ is an \wedge -definable subset of S_{n-1} . More generally for any locally definable¹ subset D of \tilde{G} , $\pi_S(D)$ is closed. Indeed $\pi_S^{-1}\pi_S(D) \cap S_n = \pi_S^{-1}(\pi_S(D \cap S_n S_n))$.

In particular, the image of $\tilde{G} \setminus S_n S_n$ in \tilde{G}/S is closed, and disjoint from $\pi_S(S_n)$ since $\pi_S(S_n S_n) \cup \pi_S(\tilde{G} \setminus S_n S_n) = \tilde{G}/S$, $\pi_S(S_n)$ lies in the interior of $\pi_S(S_n S_n)$. In particular, each $\pi_S(S_n S_n)$ is a neighborhood of 1, as is therefore $\pi_S(S_{n+1})$.

By Yamabe, L_0 has an open subgroup \check{G}/S , isomorphic to a projective limit of Lie groups. \check{G}/S is also closed, so both $\check{G} \cap D$ and $D \setminus \check{G}$ are \wedge -definable, for any definable D contained in \tilde{G} . Thus \check{G} is locally definable in \tilde{G} ; i.e., it has a definable intersection with any definable subset of \tilde{G} .

The topology of a projective limit $\varprojlim L_i$ is generated by pullbacks of open subsets of individual factors L_i . So there exist a Lie group L , a neighborhood U_1 of the identity in L , and a homomorphism $h : \check{G}/S \rightarrow L$, such that $h^{-1}(U_1) \subseteq \pi_S(S_1)$. By shrinking \check{G} down further to the pullback of the (open) connected component of 1 in L , we can take L to be connected. Let $\pi : \check{G} \rightarrow \check{G}/S \rightarrow L$ be the composition.

Now (*) holds for L : the morphism from a projective limit to one of the factors is closed; so $Y \subseteq L$ is closed iff $h^{-1}(Y) = \pi^{-1}(Y)/S$ is closed iff $\pi^{-1}(Y)$ meets every definable set in an \wedge -definable set.

We also have:

(**) For any compact neighborhood U of 1 in L , $\pi^{-1}(U)$ is commensurable to $X^{-1}X$.

This is the case because any two compact neighborhoods of 1 in L are commensurable, each one being contained in a union of translates of the other, which can be reduced by compactness to a finite union. This comparability is preserved by π^{-1} . So it suffices to show that $\pi^{-1}(U)$ contains $X^{-1}X$ for some U , and that $\pi^{-1}(U')$ is contained in finitely many translates of $X^{-1}X$ for some U' . On the other hand by the Ruzsa argument (above Lemma 3.3), any S_n is commensurable to $X^{-1}X$. We saw that $\pi_S(X^{-1}X)$ is compact; hence $\pi(X^{-1}X)$ is compact, so it is contained in some compact open neighborhood U , and thus $X^{-1}X \subseteq \pi^{-1}(U)$. By construction, $\pi^{-1}(U_1) = \pi_S^{-1}h^{-1}(U_1) \subseteq S_1$, giving the second direction.

If F is a compact subset of L and F' is an open subset, with $F \subset F'$, then there exists a definable D with $h^{-1}(F) \subset D \subset h^{-1}(F')$. Indeed $\pi^{-1}(F)$ is an \wedge -definable set contained in the \vee -definable set $\pi^{-1}(F')$, so there exists a definable D with $\pi^{-1}(F) \subseteq D \subseteq \pi^{-1}(F')$. \square

We now begin to address the issue of parameters.

Lemma 4.5. *With the assumptions and notation of Lemma 4.4, there exists an \wedge -definable subgroup S of \tilde{G} without parameters, with \tilde{G}/S bounded.*

¹See the definition in the first lines of §3.

Proof. We may work in a homogeneous elementary extension \mathbb{U} of (G, X, \cdot) , so that \wedge -definable sets are \wedge -definable without parameters as soon as they are $\text{Aut}(\mathbb{U})$ -invariant.

Let α be the set of pairs (H, Γ) such that $\Gamma \leq H \leq \tilde{G}$, and for some small base A , Γ is a normal subgroup of H , Γ is A - \wedge -definable, H is a locally definable subgroup of \tilde{G} over A , and \tilde{G}/Γ is bounded. Let β be the set of pairs $(H, \Gamma) \in \alpha$ such that if $(H', \Gamma') \in \alpha$ and $\Gamma \leq \Gamma' \leq H' \leq H$, then $H = H'$ and $\Gamma = \Gamma'$. Equivalently, the locally compact group H/Γ is connected, with no nontrivial compact normal subgroups. (Hence, by Yamabe, H/Γ is a Lie group.)

For $(H, \Gamma) \in \beta$ it is clear that H determines Γ , since if $(H, \Gamma') \in \beta$ then $\Gamma = \Gamma' = \Gamma$.

Claim 1. β is nonempty.

Proof. We saw above that there exists $(H, \Gamma) \in \alpha$ with H/Γ a connected Lie group. In the preliminaries to this section we saw that H/Γ has a maximal compact normal subgroup; it has the form H/Γ' with Γ' \wedge -definable. Then (H, Γ') is in β . \square

Claim 2. Let $(H, \Gamma), (H', \Gamma') \in \beta$. Then $(H \cap H', \Gamma \cap H') \in \beta$.

Proof. Since H' is locally definable, while Γ is contained in a definable set, it is clear that $H' \cap \Gamma$ is \wedge -definable. Since \tilde{G}/Γ and \tilde{G}/H' are bounded, so is $\tilde{G}/(\Gamma \cap H')$. Also $H \cap H'$ is locally definable. Thus $(H \cap H', \Gamma \cap H') \in \alpha$.

Now $\Gamma'/(\Gamma' \cap H)$ is bounded (as it embeds into \tilde{G}/H). By Lemma 1.6, $\Gamma' \cap H$ has finite index in Γ' .

Similarly, Γ is contained in finitely many cosets of H' , hence of $H' \cap H$. So $\Gamma(H' \cap H)$ is a finite union of cosets of $H' \cap H$, and hence is a locally definable subgroup of H . We saw that for any definable set D containing Γ , the image of $D^{-1}D$ contains an open neighborhood of the identity. Hence the image of $\Gamma(H' \cap H)$ in H/Γ is open.

Now the natural map $(H \cap H')/(\Gamma \cap H') \rightarrow H/\Gamma$ is injective. But it has open image and the group H/Γ is connected, so the map is surjective. Thus $(H \cap H')/(\Gamma \cap H') \cong H/\Gamma$ and hence has no nontrivial compact normal subgroups. \square

Similarly $(H \cap H', \Gamma' \cap H) \in \beta$. So $\Gamma' \cap H = \Gamma \cap H'$ and thus $\Gamma' \cap H = \Gamma \cap \Gamma'$.

We noted that $\Gamma' \cap H$ has finite index in Γ' ; moreover since this holds for any pair from α , in particular it holds for Γ' and any $\text{Aut}(\mathbb{U})$ -conjugate $\sigma(H)$ of H , so $\Gamma' \cap \sigma(H)$ has index bounded independently of σ .

So $\Gamma \cap \sigma(\Gamma')$ has finite index in Γ' , bounded independently of $\sigma \in \text{Aut}(\mathbb{U})$. By symmetry, Γ, Γ' are commensurable, and all conjugates of Γ are uniformly commensurable.

Pick $\Gamma_1 \in \beta$. By [2], there exists an $\text{Aut}(\mathbb{U})$ -invariant group S_1 commensurable to each conjugate of Γ_1 . The proof of [2] shows that S_1 contains a finite intersection of conjugates of Γ_1 as a subgroup of finite index; so S_1 is \wedge -definable, and of bounded index in \tilde{G} . Being $\text{Aut}(\mathbb{U})$ -invariant, it is \wedge -definable over \emptyset . Let S be the intersection of all \tilde{G} -conjugates of S_1 ; then S is normal in \tilde{G} , \wedge -definable over \emptyset , and of bounded index. \square

Proof of Theorem 4.2. We may assume that (G, X) is countably saturated, since the statements descend from a saturated extension of (G, X) to (G, X) by restriction, using the same (0-definable) \check{G}, h, L . Let S be the 0- \wedge -definable group given

by Lemma 4.4. Since $\beta \neq \emptyset$ in Lemma 4.4, we know that \check{G}, h exist over parameters, and it remains only to show that \check{G} and $\ker(h)$ can be chosen to be $\check{\vee}$ -definable and $\check{\wedge}$ -definable (respectively) without parameters.

We begin with \check{G} . We may replace \check{G} by the pre-image of any open subgroup of \check{G}/S (the “connected-by-compact” condition will remain valid). Let G_c be the group generated by $\check{G} \cap S_1$. Note that G_c is locally definable, and is generated by $Z_c = G_c \cap S_1$; so each of G_c, Z_c can be used to define the other. Now Z_c is a definable set, with parameter c say. Let Q be the set of realizations of $tp(c)$. If $c' \in Q$, then $Z(c')$ generates a group $G_{c'}$, and $G_{c'} \cap S_1 = Z(c')$. Thus for $c'', c' \in Q$, $G_{c''} = G_{c'}$ iff $Z_{c''} = Z_{c'}$; this is a definable equivalence relation.

We have $h : \check{G} \rightarrow L$. Note that if C is a compact normal subgroup of L , the composition of h with the quotient map $L \rightarrow L/C$ has the same properties as $h : \check{G} \rightarrow L$. Replacing L by L/C for a maximal compact normal subgroup C of L , we may assume that L has no compact normal subgroups.

Let K be the kernel of h . Then K is $\text{Aut}(\mathbb{U})$ -invariant. For if K' is an $\text{Aut}(\mathbb{U})$ -conjugate of K , then K, K' are $\check{\wedge}$ -definable normal subgroups of \check{G} . $K'K/K$ is a compact normal subgroup of \check{G}/K , hence it is trivial, and similarly $K'K/K'$ is trivial, so $K = K'$. Thus K is $\check{\wedge}$ -definable without parameters.

It remains to prove the uniqueness of L . Let us compare L to the locally compact group $\tilde{H} := \tilde{G}/S$, where $S = \tilde{G}_0^{00}$ is the smallest 0 - $\check{\wedge}$ -definable subgroup of \tilde{G} of bounded index. Let H be the image of \check{G} in \tilde{H} . Then H is an open subgroup of \tilde{H} , so the connected component of the identity \tilde{H}^0 is contained in H , and equals H^0 . Let C be the image of K in \tilde{H} . So C is a normal subgroup of H . Since H/C is connected, we have that $H/(CH^0)$ is both connected and totally disconnected. (Unlike the situation in the category of topological spaces, in the category of topological groups the image of a totally disconnected group is still totally disconnected. Indeed it has a pro-finite open subgroup, and this remains the case for a quotient group.) So $CH^0 = H$. Both C and H^0 are normal in H , so letting $C_0 = C \cap H^0$ we have $H/C_0 \cong C/C_0 \times H^0/C_0$. Thus the action of C by conjugation on H^0 is trivial modulo C_0 . Now C is a maximal normal compact subgroup of H ; C_0 is a compact normal subgroup of H_0 , maximal with respect to being normalized by C too, but we have just shown that this last condition is trivial, so C_0 is a maximal compact normal subgroup of H_0 . We have $L = H/C \cong H^0/C_0$ canonically. Now it is clear that C_0 is the unique maximal normal compact subgroup of H^0 . (If C_1 were another, C_0C_1 would be still bigger.) This proves the uniqueness of L . \square

If we expand $(\tilde{G}, X, \cdot, \dots)$ to a structure $\tilde{\underline{G}} = (\tilde{G}, X, \cdot, \dots, R_{new}, \dots)$ with more definable sets, the smallest 0 - $\check{\wedge}$ -definable subgroup of bounded index may become smaller: $\tilde{\underline{G}}_0^{00} \subset \tilde{G}_0^{00}$. Thus $H = \tilde{G}/S, H^0, C_0$ will change with the added structure. Nevertheless the isomorphism proved in the last paragraph of the proof remains valid; hence *the associated Lie group $L = H^0/C_0$ does not change if the structure is enriched.*

Definition 4.6. Let \tilde{G} be a $\check{\vee}$ -definable group, X a definable near-subgroup of \tilde{G} , generating \tilde{G} . Let $M = (\tilde{G}, X, \cdot, \dots)$, where \dots indicates possible additional structure.

- $LC(M) = \tilde{G}/S$, where S is the smallest $\check{\wedge}$ -definable subgroup of \tilde{G} , without parameters, of bounded index.

- $L(X)$ is the Lie group associated to X ; so $L(X) = LC(M)^0/C_0$, with C_0 a maximal normal compact subgroup of $LC(M)^0$. Let $\widehat{L}(X) = LC(M)/C_0$; then $L(X) = \widehat{L}(X)^0$, the connected component.
- $l(X) = \dim L(X)$.

We refer to $l(X)$ as the Lie rank of X , or of \widetilde{G} .

Example 4.7. If a near-subgroup X has $l(X) = 0$, then there exists a definable group S with X, S commensurable. Indeed in this case kernel S of the homomorphism $\check{G} \rightarrow L$ is equal to \check{G} ; but S is \wedge -definable and \check{G} is \vee -definable, so they are definable.

Lemma 4.8. *In the situation of Theorem 4.2, assume that the $S1$ -ideal arises from an invariant, translation invariant measure μ . Let $k_5 = \mu(XX^{-1}XX^{-1}X)/\mu(X)$. Extend μ to the σ -algebra generated by the ∞ -definable subsets of \widetilde{G} , and let λ be the pushforward of μ to $\widehat{L} = \widehat{L}(X)$, i.e. $\lambda(U) = \mu(\pi^{-1}(U)) \in \mathbb{R}_\infty$, where π is the quotient map. Then λ is a Haar measure on \widehat{L} . We have $\lambda((\pi X)(\pi X)^{-1}(\pi X)) \leq k_5\lambda(\pi X)$. Moreover, there exists a compact subset W of $L = L(X)$ with $\lambda(W) > 0$ and $\lambda(WW^{-1}W) \leq k_5\lambda(W)$. We can take $1 \in W$.*

Proof. It is clear that λ is a nonzero, translation invariant measure, hence a Haar measure. We have $XX^{-1}X \subseteq \pi^{-1}((\pi X)(\pi X)^{-1}(\pi X)) \subseteq (XX^{-1}XX^{-1}X)$, since $\pi^{-1}(1) \subseteq X^{-1}X$. This implies the first inequality, by definition of the pushforward measure. Moving to L , recall that we have $h : \check{G} \rightarrow L$ with kernel K (Theorem 4.2), with \check{G} a \vee -definable subgroup of \widetilde{G} , and \widetilde{G}/\check{G} bounded. In particular X/\check{G} is bounded, so X intersects finitely many cosets of \check{G} ; say $X = \bigcup_{i=1}^r X_i$, with $X_i \subseteq c_i\check{G}$, and c_i lying in distinct cosets of \check{G} . Let $k_3 = \mu(XX^{-1}X)/\mu(X)$. Then, noting that $X_iX_i^{-1}X_i \subseteq c_i\check{G}$, and the $c_i\check{G}$ are disjoint, we have:

$$\sum_i \mu(X_iX_i^{-1}X_i) \leq \mu(XX^{-1}X) \leq k_3\mu(X) = \sum_i k_3\mu(X_i).$$

The sum being extended over all $i \leq r$ such that $\mu(X_i) > 0$, it follows that for at least one i with $\mu(X_i) > 0$, we have $\mu(X_iX_i^{-1}X_i) \leq k_3\mu(X_i)$. Similarly, for at least one i with $\mu(X_i) > 0$, we have $\mu(X_iX_i^{-1}X_iX_i^{-1}X_i) \leq k_5\mu(X_i)$. Let $Y = c_i^{-1}X_i$. Then $h(Y)$ is a compact subset of L ; $\lambda(h(Y)) = \mu(h^{-1}h(Y)) \geq \mu(Y) > 0$ and $\lambda(Y Y^{-1} Y) \leq k_5\lambda(Y)$ by the same argument as for \widehat{L} above. By translating W , we can arrange $1 \in W$. \square

Can Lemma 4.8 be used to bound $l(X) = \dim(L)$ in terms of doubling constants of X ? When G is nilpotent, we have: $l(X) \leq \log_2(k_5)$. This follows from Lemma 4.8 and Lemma 4.9, due (with a different proof) to Tsachik Gelander; thanks for allowing me to include it here. Use $1 \in W$ to obtain $WW \subseteq WW^{-1}W$ in order to apply the lemma.

Lemma 4.9 (Gelander). *Let X be a compact subset of \mathbb{R}^d , or more generally of a connected, simply connected Lie group, and let λ be Haar measure. Then $\lambda(XX) \geq 2^d\lambda(X)$.*

Proof. In fact we have $\lambda(s(X)) \geq 2^d\lambda(X)$, where $s(x) = x^2$. The ambient group H is isomorphic to a subgroup of the strict upper triangular matrices over \mathbb{R} , of some dimension; the map s is hence injective. Moreover H is diffeomorphic to \mathbb{R}^d , and

the differential ds of s at any point is a linear transformation of the form $2 + M$, with M nilpotent. It follows that the Jacobian determinant has value 2^d , so by the change of variable formula for integration, the diffeomorphism s expands volume by exactly 2^d . \square

Remark 4.10. (Compare [50], Lemma 7.7 and Theorem 7.12.)

Let Γ be a \wedge -definable subgroup of bounded index in the \vee -definable group \tilde{G} . Let \tilde{N} be a locally definable normal subgroup of \tilde{G} , and let $\pi : \tilde{G} \rightarrow \tilde{G}/\tilde{N}$ be the quotient map. The main case is that \tilde{N} is the intersection with \tilde{G} of a definable normal subgroup N of G .

(0) The image $\mathbf{\Gamma}$ of Γ has bounded index in the image \mathbf{G} of \tilde{G} modulo \tilde{N} , and also $\Gamma \cap \tilde{N}$ has bounded index in \tilde{N} . Conversely in this situation the boundedness of \tilde{G}/Γ follows from that of $\mathbf{G}/\mathbf{\Gamma}$ and of $\Gamma \cap \tilde{N}$ in \tilde{N} .

(1) View \tilde{G}/Γ , $\mathbf{G}/\mathbf{\Gamma}$ and $\tilde{N}/(\Gamma \cap \tilde{N})$ as locally compact groups. Then $\tilde{N}/(\Gamma \cap \tilde{N})$ with the logic topology is homeomorphic to the image of \tilde{N} in \tilde{G}/Γ , with the subspace topology. Indeed the natural map $\tilde{N}/(\Gamma \cap \tilde{N}) \rightarrow \tilde{G}/\Gamma$ is a continuous injective homomorphism. To see that it is also a closed map, since \tilde{G}/Γ is covered by the interiors of sets of the form $\pi(D)$, with D definable, we may restrict attention to the inverse image of such a set. But then we are looking at an injective continuous map between compact Hausdorff spaces, hence an isomorphism.

Similarly, $\mathbf{G}/\mathbf{\Gamma} \cong (\tilde{G}/\Gamma)/(\tilde{N}/(\Gamma \cap \tilde{N}))$ as topological groups.

(2) Γ is definable iff the topology on \tilde{G}/Γ is discrete. This makes it plain that Γ is definable iff $\pi(\Gamma)$ and $\Gamma \cap \tilde{N}$ are.

(3) If \tilde{G}/Γ is a Lie group, then so are $\tilde{N}/(\Gamma \cap \tilde{N})$ and $\mathbf{G}/\mathbf{\Gamma}$, and we have an exact sequence

$$1 \rightarrow \tilde{N}/(\Gamma \cap \tilde{N}) \rightarrow \tilde{G}/\Gamma \rightarrow \mathbf{G}/\mathbf{\Gamma} \rightarrow 1.$$

This in turn induces an exact sequence of homomorphisms among the Lie algebras. It follows that $\dim(\tilde{G}/\Gamma) = \dim(\mathbf{G}/\mathbf{\Gamma}) + \dim(\tilde{N}/(\Gamma \cap \tilde{N}))$.

(4) From (3) it follows that

$$l(\tilde{G}) \geq l(\tilde{G}/\tilde{N}) + l(\tilde{N}).$$

Indeed we may move from \tilde{G} to \mathbf{G} , changing none of the three numbers. Then we may enlarge Γ so that \mathbf{G}/Γ has no nontrivial normal compact subgroups. By (3) we obtain in this situation: $l(\tilde{G}) = \dim(\mathbf{G}/\mathbf{\Gamma}) + \dim(\tilde{N}/(\Gamma \cap \tilde{N}))$. Now $\tilde{N}/(\Gamma \cap \tilde{N})$ may have nontrivial compact subgroups, but we have at all events $l(\tilde{N}) \leq \dim(\tilde{N}/(\Gamma \cap \tilde{N}))$ (the inequality may be strict). Similarly $l(\tilde{G}/\tilde{N}) \leq \dim(\tilde{G}/\tilde{N})$, and (4) follows.

See §7 for an inductive use of this invariant, similar to Gromov's use of the growth rate in the case of his polynomial growth assumption.

Remark 4.11. The canonicity of L in Theorem 4.2 is achieved at a price. We noted already that it requires moving from $(X^{-1}X)^2$ to $(X^{-1}X)^m$, where m is difficult to control. In addition, factoring out the maximal compact normal subgroup can lead to substantial loss of information.

In some cases there will exist a largest \wedge -definable normal subgroup Δ of \tilde{G} with $\Delta \subseteq X^{-1}X$. By Yamabe, \tilde{G}/Δ is a Lie group \tilde{L} . In this case \tilde{L} too is an invariant of (\tilde{G}, X) and is superior in both respects. When it exists, we may call \tilde{L}

the directly associated Lie group. More generally we may need to look at a number of \tilde{G}/Δ , differing by compact isogenies.

For example, let $\alpha > 10$ be an irrational real number, and let

$$X[n] = X[n, \alpha] = \{[m\alpha] : m \in \mathbb{Z}, -n \leq m \leq n\},$$

where $[m\alpha]$ is the integer part of $m\alpha$. $X[n]$ is symmetric, and satisfies $|X[n]X[n]|/|X[n]| \leq 4$. Let (G, X, n^*) be a nonprincipal ultraproduct of $(\mathbb{Z}, X[n], n)$. Then the directly associated Lie group is the product of the circle $\mathbb{R}/\alpha\mathbb{Z}$ with \mathbb{R} . The map $\tilde{G} \rightarrow \mathbb{R}$ takes x to the standard part of x/n^* . The map $\tilde{G} \rightarrow \mathbb{R}/\alpha\mathbb{Z}$ takes x to the standard part of the image of x in the nonstandard circle \mathbb{R}^*/α . The image of X in the cylinder $\mathbb{R} \times \mathbb{R}/\alpha\mathbb{Z}$ is the image of the square $[-\alpha, \alpha] \times [-1, 0]$. The image of the element $[\alpha]$ is $(0, m + \alpha\mathbb{R})$ for some nonzero integer m ; it follows that $(0) \times \mathbb{R}/\alpha\mathbb{Z}$ is contained in the image of \tilde{G} , so that $\tilde{G} \rightarrow \mathbb{R} \times \mathbb{R}/\alpha\mathbb{Z}$ is surjective. The doubling of this square within the cylinder is similar to the doubling of the $X[n]$ within \mathbb{Z} . By contrast the associated Lie group without compact subgroups is \mathbb{R} , which does not account for the doubling of X or XX very well, and only begins to work around the $[\alpha]$ th set power of X .

It is also interesting to note here that if one takes $X[n]' = \{[m\alpha_n] : m \in \mathbb{Z}, -n \leq m \leq n\}$, where α_n approaches ∞ , the associated Lie group will be \mathbb{R}^2 ; this limit is natural for the directly associated Lie group but not for the reduced one.

For the record we state a version of Theorem 4.2 waiving canonicity but gaining more control of the location of the kernel.

Lemma 4.12. *Let X generate a \vee -definable group \tilde{G} , and assume an ideal on \tilde{G} exists satisfying the assumption of Lemma 2.17. Then there exists a \vee -definable subgroup \check{G} contained in the group generated by X , a \wedge -definable subgroup $K \subseteq \check{G}$, a connected, finite-dimensional Lie group L and a homomorphism $h : \check{G} \rightarrow L$ with kernel $K \subseteq (X^{-1}X)^2$ and dense image, such that:*

If $F \subseteq F' \subseteq L$ with F compact and F' open, then there exists a definable D with $h^{-1}(F) \subseteq D \subseteq h^{-1}(F')$. Any such D is commensurable to $X^{-1}X$.

\check{G}, K may be defined with parameters in any given model.

Proof. By Lemma 2.17 and Theorem 3.5, one obtains an \wedge -definable stabilizer S defined over a given model, and with $S \subseteq (X^{-1}X)^2$. It follows that the image U of $(X^{-1}X)^2$ in \tilde{G}/S contains the identity in its interior. We follow the proof of Theorem 4.2, taking care to factor out only by a compact subgroup contained in the given neighborhood U . □

In the local group setting, the conclusion reads: there exists a homomorphism $h : W \rightarrow L$ of local groups, W a subset of X commensurable to $X^{-1}X$, such that $\mathbf{X} = h(X)$ is a compact neighborhood of $1 \in L$; and if $F \subseteq F' \subseteq \mathbf{X}$ with F compact and F' open, then there exists a definable D with $h^{-1}(F) \subseteq D \subseteq h^{-1}(F')$. Any such D is again commensurable to $X^{-1}X$. I have not checked the question of parameters for local groups.

Corollary 4.13. *Let X be a near-subgroup of a group G_0 , generating \tilde{G} . Then there exist 0-definable subsets X_1, X_2, \dots of \tilde{G} , commensurable to $X^{-1}X$, and $c \in \mathbb{N}$, with:*

- (1) $1 \in X_n = X_n^{-1}$.
- (2) $X_{n+1}X_{n+1} \subseteq X_n$.

- (3) X_n is contained in $\leq c$ translates of X_{n+1} .
- (4) $aX_{n+1}a^{-1} \subseteq X_n$ for $a \in X_1$.
- (5) $[X_n, X_m] = \{xyx^{-1}y^{-1} : x \in X_n, y \in X_m\} \subseteq X_k$ whenever $k < n + m$. In particular each X_n is closed under the commutator bracket.
- (6) $X_{n+1} = \{x \in X_1 : x^4 \in X_n\}$.
- (7) Let $x, y \in X_m, m \geq 2$ and suppose $x^2 = y^2$. Then $xy^{-1} \in \bigcap_n X_n$.

Proof. We may assume (G_0, X) is \aleph_0 -saturated. Let h, L be as in Theorem 4.2. We first show that L has a system U_n of compact neighborhoods of the identity with properties (1)–(3).

Let \mathfrak{L} be the Lie algebra of L , $\exp : \mathfrak{L} \rightarrow L$ the exponential map, and fix a Euclidean inner product on \mathfrak{L} . Let V be a simply connected open neighborhood of $0 \in \mathfrak{L}$ such that \exp is a diffeomorphism $V \rightarrow \exp(V) = U$, and such that the image of $(X^{-1}X)^2$ in L contains U in its interior. Let V_n be the ball of radius $r_0 2^{-n}$ around 0 . Here $r_0 > 0$ is chosen small enough so that V_0 is contained in V ; some further constraints on r_0 will be specified later. Viewing \mathfrak{L} as the tangent space at 1 of L , fix on L the unique left-invariant Riemannian metric extending the given inner product at 1 . Let $U_n = \exp(V_n)$. Note that U_n is the set of points of U at Riemannian distance $\leq r_0 2^{-n}$ from the identity element (cf. e.g. [36], Prop. 6.10). It follows that (1)–(2) hold: $1 \in U_n = U_n^{-1}$ and $U_{n+1}U_{n+1} \subseteq U_n$.

Fix an invariant volume form ω on L . We claim that for some constant $c' > 0$, we have $\text{vol}(U_{n+1}) \geq c' \text{vol}(U_n)$ for large enough n . We have $\text{vol}(U_n) = \int_{V_n} \exp^* \omega$ where $\exp^* \omega$ is the pullback. Now V_n has volume proportional to 2^{-nd} , with respect to the standard Euclidean volume form ω_1 . We have $\exp^* \omega = f \omega_1$ for some non-vanishing smooth function f , which we can take to be positive. On V we have $(c'')^{-1} \leq f \leq c''$ for some $c'' > 0$, so $\text{vol}(U_n) \leq c'' \text{vol}(V_n) \leq 2^d c'' \text{vol}(V_{n+1}) \leq 2^d (c'')^2 \text{vol}(U_{n+1})$.

Now U_{n-1} contains at most $\text{vol}(U_{n-1}) / \text{vol}(U_{n+2})$ disjoint U_n -translates of U_{n+2} ; hence U_n is contained in that many translates of $U_{n+2}^{-1}U_{n+2} \subseteq U_{n+1}$. This gives the analogue of (3).

To obtain (4), we may begin with r_1 small enough so that for $x \in U_1$, $1 - ad_x$ has operator norm $< 1/2$. Then $ad_x(V_{n+1}) \subseteq V_n$, so $x^{-1}U_{n+1}x \subseteq U_n$.

(5) Let $c(x, y) = \log(\exp(x)\exp(y)\exp(-x)\exp(-y))$. We have to show that $c(V_n, V_m) \subseteq V_k$ when $k \leq N, k < n + m$. Now if $|u| < 2^{-n}$ and $|v| < 2^{-m}$, then $|c(u, v) - [u, v]| O(2^{-m-n-\min(m,n)})$, where $[u, v]$ is the Lie algebra bracket. This can be seen by looking at the power series expansion of c ; it begins with $[u, v]$, followed by higher-order terms. So the statement holds for large enough m, n ; by renormalizing (replacing V_n by V_{n+k}) we obtain the result.

Finally note that $U_{n+1} = \{u \in U_1 : u^2 \in U_n\}$ since for $u = \exp(v)$ we have $u^2 = \exp(2v)$ and $u \in U_{n+1}$ iff $v \in V_{n+1}$ iff $2v \in V_n$ iff $u^2 \in U_n$.

Since $h^{-1}(U_2)$ is an \wedge -definable set contained in the definable set $h^{-1}(U_1)$, there exists a definable set Y_1 with $h^{-1}(U_2) \subseteq Y_1 \subseteq h^{-1}(U_1)$. Define Y_n inductively by $Y_{n+1} = \{y \in Y_1 : y^2 \in Y_n\}$. It follows that $h^{-1}(U_{n+1}) \subseteq Y_n \subseteq h^{-1}(U_n)$. (If $h(x) \in U_{n+1}$, then $h(x^2) \in U_n$. By induction $x^2 \in Y_{n-1}$; so $x \in Y_n$. If $x \in Y_n$, then $x^2 \in Y_{n-1} \subseteq h^{-1}U_{n-1}$, so $h(x)^2 \in U_{n-1}$ and $h(x) \in U_n$.) Clearly $Y_n = Y_n^{-1}$. It follows from the intertwining of the Y_n in the $h^{-1}U_n$ that $Y_{n+2}Y_{n+2} \subseteq Y_n$, that Y_n is contained in at most c^2 translates of Y_{n+1} , $aY_{n+2}a^{-1} \subseteq Y_n$, and $[Y_n, Y_m] \subseteq Y_k$ whenever $k + 1 < n + m$.

Let $X_n = Y_{2n}$. Then it is clear that (1)–(6) hold. (7) follows from the fact that squaring is injective on U_1 (if one chooses U_1 small enough). \square

Remark 4.14. (1) Let (X_n) be as in Lemma 4.13. Let (G, X) be a nonprincipal ultraproduct of (G_0, X_n) , and let \tilde{G} be the subgroup of G generated by X . Then there exists a locally definable subgroup $\check{G} \leq \tilde{G}$, a connected, finite-dimensional Lie group U , and a homomorphism $h : \check{G} \rightarrow U$ as in Theorem 4.2, such that in addition, U is Abelian.

Proof. Let h, L, X_n be as in Corollary 4.13. For $k \in \mathbb{Z}$, for all $n \geq k$, define $X_n[k] = X_{n+k}$. This carries over to the ultraproduct, so $X[k]$ is defined for all $k \in \mathbb{Z}$, and has similar properties: $X[k]X[k] \subseteq X[k+1]$. Since the ultrafilter is nonprincipal, it concentrates on $n > k$, so by (5) of Corollary 4.13 we have $[X, X[k]] \subseteq X[k+n-1] \subseteq X[k']$ for all $k' \in \mathbb{N}$. Factoring out $\bigcap_k hX[k]$ we obtain a commutative locally compact group. As in Theorem 4.2 we may replace it with a commutative Lie group. \square

We now deduce a version in the asymptotic setting. Here we do not obtain an infinite chain, but the function f serves to say that the length of the chain is arbitrarily large compared to e, c, k . We say that two sets are e -commensurable if each is contained in the union of $\leq e$ cosets of the other. Taking ν to be the counting measure, we obtain (a strengthening of) Theorem 1.1.

Corollary 4.15. *Let $f : \mathbb{N}^2 \rightarrow \mathbb{N}$ be any function, and fix $k \in \mathbb{N}$. Then there exist $e^*, c^*, N \in \mathbb{N}$ such that the following holds.*

Let G be any group, X a subset, and assume there exists a translation invariant finitely additive real-valued measure ν on the definable subsets of G contained in some power of XX^{-1} , with $\nu(XX^{-1}X) \leq k\nu(X)$.

Then there exist $e \leq e^, c \leq c^*$ and 0-definable subsets $X_N \subseteq X_{N-1} \subseteq \dots \subseteq X_1$, $N > f(e, c)$ such that $X^{-1}X$ and X_1 are e -commensurable and for $1 \leq m, n < N$ we have*

- (1) $X_n = X_n^{-1}$.
- (2) $X_{n+1}X_{n+1} \subseteq X_n$.
- (3) X_n is contained in $\leq c$ translates of X_{n+1} .
- (4) $aX_{n+1}a^{-1} \subseteq X_n$ for $a \in X_1$.
- (5) $[X_n, X_m] \subseteq X_k$ whenever $k \leq N$ and $k < n + m$. In particular each X_n is closed under the commutator bracket.
- (6) $X_{n+1} = \{x \in X_1 : x^4 \in X_n\}$.

Proof. Fix f, k . We consider groups G and subsets X admitting a measure as above, with $\nu(X) = 1, \nu(XX^{-1}X) \leq k$ (as we can always arrange by renormalizing ν). Consider integers c , and formulas ϕ of one free variable. Given c, ϕ and X , let X_1 be the subset defined by ϕ . Let $N = f(e, c) + 1$, and define X_n using (6) for $2 \leq n \leq N$. Let us say that (c, e, ϕ) works for X if properties (1)–(5) hold for the sets X_n defined in this way.

We will show that for some finite set $(c_1, e_1, \phi_1), \dots, (c_n, e_n, \phi_n)$, for any G and any k -near-subgroup X of G , some (c_i, e_i, ϕ_i) works for X . Suppose this is false. Then by the compactness theorem there exists G , a measure μ on the definable subsets of the group \tilde{G} generated by X with $\mu(X) = 1$, and a definable subset X of G with $\mu(XX^{-1}X) \leq k\mu(X)$ such that no (c, ϕ) works. But let X_1 be the definable set provided by Corollary 4.13. Let c be the integer of Corollary 4.13 (3), and e the

number of translates of X_1 needed to cover X . Then (c, e, ϕ) works for X (indeed the X_n have the required properties beyond any bound). This contradiction proves the statement and the theorem. \square

Remark 4.16. (1) Again we can also add (7): if $x, y \in X_2$ and $x^2 = y^2$, then $xy^{-1} \in X_N$.

(2) We could add that $X_1 \subseteq (X^{-1}X)^2$ (as we do in the statement of Theorem 1.1), if we waive the 0-definability of the X_n . They remain definable over parameters from the given structure. See Lemma 4.12.

(3) This type of proof is always effective in the sense of Gödel. Note that if f is recursive, then e, c, N are automatically given by a recursive function (it suffices to search for e, c, N, X_1 such that (1)–(6) hold).

The sequence of subsets X_n in Corollary 4.15 is recursively determined by X_1 , via (6). The question is thus how to describe X_1 . Studying the proof of the fundamental theorems on locally compact groups should provide detailed information; for now we state what is clear a posteriori when they are treated as a black box.

Regarding X_1 , we have:

Corollary 4.17. *Fix $k \in \mathbb{N}$, and f as above. Then there exists m and an algorithm that accepts as input the multiplication table of a finite near-subgroup X up to $(X^{-1}X)^m$ and yields the set X_1 in polynomial time.*

Proof. A formula in a logic with measure quantifiers Q_ϵ can be computed in polynomial time. \square

Note that we do not assume that G itself is finite, and even if finite, it is not available to the algorithm, beyond $(X^{-1}X)^m$. Indeed the algorithm can be made to work for local groups.

We can improve the m to 3 if, as in Remark 4.16 (2), we use a formula with a parameter from X . In this case the algorithm will first search for a parameter satisfying an appropriate auxiliary formula, then compute X_1 using this parameter.

To illustrate Theorem 4.12 we recover an easy version of a theorem of Freiman’s (see [54]) (generalized to the noncommutative case).

Corollary 4.18. *Fix m, k . Then there exists $e = e(k, m)$ with the following property. Let G be a group of exponent m , i.e. $x^m = 1$ for any $x \in G$. Let X be a finite k -approximate subgroup of G . Then there exists a subgroup S of G such that S, X are e -commensurable.*

Proof. By compactness it suffices to show that if X is a near-subgroup of a group G of bounded exponent, then there exists a definable subgroup S of G such that S, X are commensurable.

By a theorem of Schur’s ([9], 36.14), a periodic subgroup P of $GL_n(\mathbb{C})$ has an abelian normal subgroup of finite index. When the period is bounded, the abelian subgroup and hence P must be finite. If L is a connected Lie group with center Z , by considering the action of L on its Lie algebra we see that L/Z is linear. Hence a periodic subgroup P of L of bounded period must be contained in Z up to finite index, and again it follows that P is finite.

Thus the image of \hat{G} in the Lie group L associated to X is finite. Since this image is dense in L , and L is connected, it follows that L is trivial. The conclusion follows from Example 4.7. \square

Actually it is easy to see in the same way (using the fact that L has no compact normal subgroups) that if G is a periodic group and X is a near-subgroup, then there exists a subgroup S of G such that S, X are commensurable. This does not extend to families of finite approximate groups, since without a uniform bound taking an ultraproduct will not preserve periodicity.

5. LINEAR GROUPS

Up to this point the hypotheses in this paper were purely measure-theoretic, at the top dimension as it were. We will now look at lower dimensions as well. Numerically this means that if a subset Y of X has about $c|X|^\alpha$ elements, we pay attention to $\alpha < 1$ and not only to c when $\alpha = 1$. Our main tool is a cardinality estimate due in its original form to Larsen and Pink. In [24], it was presented as a dimension comparison lemma and slightly generalized in a number of directions; one of these will be needed here. We will first define quasi-finite dimension in general, and specifically for ultraproducts of finite approximate subgroups. Then, assuming the group is linear (or indeed densely embedded in a group with a nice dimension theory) we show that the ambient dimension \dim constrains strongly the quasi-finite dimension. Finally, knowing that the group looks sufficiently noncommutative by certain measures using quasi-finite dimension, we can conclude using the stabilizer that it is in fact definable.

5.1. The semigroup of dimensions. Let K be an ultraproduct of structures K_i for some language L . For each i we consider, along with K_i , the counting measure on definable sets, as a map from the class of definable sets into \mathbb{R} . Taking the ultraproduct of these maps as well, we obtain a map from the class of definable sets of K into the ultrapower \mathbb{R}^* of \mathbb{R} . This is a countably saturated real closed field. For nonempty definable X (represented by a sequence $X(K_i)$), we have a nonstandard real number $\log|X|$ (represented by the sequence $\log|X(K_i)|$).

Let C be a convex subgroup of \mathbb{R}^* . We assume C is a countable union or a countable intersection of definable subsets of \mathbb{R}^* . Then \mathbb{R}^*/C is an ordered \mathbb{Q} -vector space. We define $\delta(X)$ to be the image of $\log|X|$ in \mathbb{R}^*/C . We view δ as a (nonintegral) dimension.

Subadditivity: Let $f : X \rightarrow Y$ be a definable map; assume $\delta(f^{-1}(y)) \leq \alpha = a + C$ for each $y \in Y$, and $\delta(Y) \leq \beta = b + C$. Then $\delta(X) \leq \alpha + \beta$. To see this, if $C = \bigcup_n C_n$ is a countable union, we may take $C_1 \subset C_2 \subset \dots$. We have $\log|Y| \leq b + c$ with $c \in C_n$ for some n , and by compactness, $\log|f^{-1}(y)| \leq a + c'$, with $c' \in C_{n'}$ for some n' . It follows that $\log|X| \leq a + b + c + c'$, so $\delta(X) \leq \alpha + \beta$. If $C = \bigcap_n C_n$, then $\log|Y| - b \in C_n$, and $\log|f^{-1}(y)| - a \in C_n$ for each n ; hence $\log|X| - a - b \in C_n$ for each n .

We would like to extend the dimension to \wedge -definable sets. Fix $\delta_0 \in \mathbb{R}^*$, $\delta_0 > C$. Let $V_0 = V_0(\delta_0)$ be the group of elements $a \in \mathbb{R}^*/C$ such that $-n\delta_0 + C \leq a \leq n\delta_0 + C$ for some $n \in \mathbb{N}$. Let $V = V(\delta_0)$ be the set of cuts of V_0 , i.e. subsets $s \subset V_0$ that are nonempty, bounded above, and closed downwards. This is a semigroup under set addition, linearly ordered by set inclusion. V_0 embeds into V , by $a \mapsto \{v : v \leq a\}$. We identify V_0 with its image in V . Any subset of V that is bounded below has a greatest lower bound, namely the intersection. We note that V_0 consists of invertible elements of V , and that it is semi-dense in V , in the sense that if $u < v \in V$, then there exists $z \in V_0$ with $u < z \leq v$.

It will suffice for our purposes to use the intermediate subsemigroup V_1 consisting of infima of bounded countable subsets of V_0 .

We could also form the linearly ordered semigroup V' of cuts in $V'_0 := \{a \in \mathbb{R}^* : \text{for some } n \in \mathbb{N}, -n\delta_0 \leq a \leq n\delta_0\}$. The natural map $V'_0 \rightarrow V_0$ maps cuts to cuts and respects addition and \leq . V can be identified with the ordered subsemigroup of cuts $I \in V'$ with $C + I = I$. Note that for a subset of V , the infimum (in the sense of V') lies in V and agrees with the infimum in V .

When α_n, β_n are descending sequences of cuts, $\inf_n(\alpha_n + \beta_n) = \inf_n \alpha_n + \inf_n \beta_n$ holds in V . By the above remark, it suffices to check this in V' . The inequality \geq is clear, since $\alpha_n + \beta_n \geq \inf_n \alpha_n + \inf_n \beta_n$ for each n . For the other inequality, let $\alpha'_n \in \alpha_n \setminus \alpha_{n+1}$, $\beta'_n \in \beta_n \setminus \beta_{n+1}$. Suppose $c \in \mathbb{R}^*$ and $c \leq \inf_n(\alpha_n + \beta_n)$. Then by countable saturation there exist $(\alpha, \beta) \in \mathbb{R}^*$ with $\alpha \leq \alpha'_n, \beta \leq \beta'_n$ for each n , and $c \leq \alpha + \beta$. Hence $c \leq \inf_n \alpha_n + \inf_n \beta_n$.

Let us also point out that if $\alpha < \alpha'$ and $\beta < \beta'$ are cuts, then $\alpha + \beta < \alpha' + \beta'$. This holds for any semigroup of cuts in a dense linearly ordered group; to prove it we may consider the semigroup of all cuts. Let $a \in \alpha' \setminus \alpha, b \in \beta' \setminus \beta$. Let $a^- = \{x : x < a\}$, and similarly for b^- . We have $a^- + b^- \leq a + b$, and $a^- + b^- \neq a + b$ since the cut $a^- + b^-$ has no maximal point. Thus $\alpha + \beta \leq a^- + b^- < a + b \leq \alpha' + \beta'$. (One strict inequality and one weak inequality would not suffice for the same.)

The multiplicative group $\mathbb{Q}^{>0}$ acts on V_0 , and hence on V .

5.2. Quasi-finite dimension. For an \wedge -definable set X define:

$$\delta(X) = \inf \delta(D),$$

where D ranges over all definable sets containing X . Note a continuity property of the dimension: If $X = \bigcap X_n$ with $X_1 \supset X_2 \supset \dots$ \wedge -definable, then $\delta(X) = \inf_n \delta(X_n)$.

The subadditivity property holds for \wedge -definable sets X : let f be a definable map, let $\gamma \in V_1$, and assume $\delta(f^{-1}(a) \cap X) \leq \gamma$ for all a . Then $\delta(X) \leq \delta(f(X)) + \gamma$. Indeed if $X = \bigcap X_n$ with X_n a descending sequence of definable sets, then $f(X) = \bigcap_n f(X_n)$ by compactness (saturation); say $\gamma = \inf \gamma_k$; then for each k , for some $n(k)$, we have $\delta(f^{-1}(a) \cap X_{n(k)}) \leq \gamma_k$, again by compactness. So $\delta(X_{n(k)}) \leq \delta(f(X_{n(k)})) + \gamma_k$. Thus $\inf_n \delta(X_n) \leq \inf_k \delta(X_{n(k)}) \leq \inf_k \delta(f(X_{n(k)})) + \inf_k \gamma_k = \delta(f(X)) + \gamma$.

As a very special case of subadditivity, noting that $\delta(F) = 0$ for finite F , we have $\delta(D_1 \cup D_2) = \max(\delta(D_1), \delta(D_2))$.

Also, $\delta(D_1 \times D_2) = \delta(D_1) + \delta(D_2)$: let \mathcal{E}_i be the family of definable sets containing D_i . For any definable set E with $D_1 \times D_2 \subseteq E$ there exist (by compactness) $E_i \in \mathcal{E}_i$ with $D_i \subseteq E_i$ ($i = 1, 2$) and $E_1 \times E_2 \subseteq E$. Thus $\delta(D_1 \times D_2) = \inf_{E_1 \in \mathcal{E}_1, E_2 \in \mathcal{E}_2} \delta(E_1 \times E_2) = \inf_{E_1, E_2} (\delta(E_1) + \delta(E_2)) = \inf_{E_1} \delta(E_1) + \inf_{E_2} \delta(E_2) = \delta(D_1) + \delta(D_2)$.

If X is \wedge -definable over a set A , there exists a complete type P over A containing X with $\delta(X) = \delta(P)$. To see this it suffices to check that $X \setminus \bigcup \{D : \delta(D) < \delta(X)\}$ is nonempty, since any type extending this will do. By compactness it suffices to see that X is not contained in a finite union of sets D with $\delta(D) < \delta(X)$. This is clear using $\delta(D_1 \cup D_2) = \max(\delta(D_1), \delta(D_2))$.

If $\delta(X) \in V_0$, we say that X has strict quasi-finite dimension. Note in this case, by saturation of \mathbb{R}^* , that if $\delta(X) = \inf_{n \in \mathbb{N}} \alpha_n$, then $\delta(X) = \alpha_n$ for large enough n .

5.3. Examples. The best behaved case is of totally categorical theories, i.e. theories T with a unique model in each power; a basic example is the theory of vector

spaces over a fixed finite field. In this case, any finite subset T_0 of T has finite models, described by a single integer parameter; the cardinality is precisely given by a polynomial P in this dimension parameter. In this case, regardless of the choice of convex subgroup, the quasi-finite dimension equals the degree of the Zilber polynomial (times a fixed scalar) and recovers the Morley dimension. This was an essential ingredient of Zilber's theory of totally categorical structures; see [8] for generalizations and converses.

In general, the two most natural choices for a convex subgroup C are the smallest nontrivial convex subgroup, the convex hull C_{min} of \mathbb{Z} , and the largest convex subgroup C_{max} with $\delta_0 \notin C_{max}$.

C_{max} is a countable *intersection* of definable subsets of \mathbb{R}^* . The corresponding group of dimensions is canonically isomorphic to \mathbb{R} (with δ_0 mapping to 1). Each dimension α induces an ideal I_α (Example 2.13), which is not S1. Asymptotically, a set X represented by a sequence of finite sets X_i has the same dimension as $Y \subset X$ (represented by Y_i) if for any ϵ , for almost all i , $|Y_i| \geq |X_i|^{1-\epsilon}$.

C_{min} is a countable *union* of definable subsets of \mathbb{R}^* . The group of dimensions is more complicated, but when $\delta(X) = \alpha$, the ideal of lower-dimensional subsets of X is an S1-ideal. Here $Y \subset X$ has the same dimension as X if for some k , $|X_i| \leq k|Y_i|$ for almost all i .

5.4. Minimality. Now assume that each K_i is a field, possibly with additional structure. There will be no loss of generality in assuming that K_i is algebraically closed. Let K be an ultraproduct of the structures K_i . Constructible sets and varieties will be assumed to be defined over K . Here the word “constructible” means: definable in $K = K^{alg}$ as a field, whereas “definable” means: definable in (K, \dots) as an L -structure. For a constructible set S , aside from the pseudo-finite dimension δ constructed above, we have the dimension in the sense of algebraic varieties. This can be defined as the Morley rank of the set, viewed as definable in $(K, +, \cdot)$, or it can be defined as the dimension of the Zariski closure of S , as in [56]; see [44].

Let G be a simple algebraic group over the ultraproduct K . G can be viewed as a group subvariety of the group GL_n of invertible matrices. We write G when we think of the defining equations, and $G(K)$ when we think of the set of points of K .

Let Γ_0 be a Zariski dense subset of $G(K)$. Consider the functions $F_c(x, y) = cx^{-1}c^{-1}y$, $c \in \Gamma_0$. Any subvariety H of $G(K)$ closed under all the F_c must be a group subvariety of G , normalized by Γ_0 , hence by the Zariski closure of this group, i.e. by G . Since G is simple, we must have $H = 1$ or $H = G$.

It follows that if Y, Z are constructible subsets of G , defined over a subfield A of K , and if $0 < \dim(Y) \leq \dim(Z) < \dim(G)$, then $\dim(F_c(Y \times Z)) > \dim(Z)$ for some $c \in \Gamma_0$. Moreover, let $Y \times' Z = Y \times Z \setminus \bigcup_j W_j$, where W_j ranges over all A -definable constructible subsets W of $Y \times Z$ with $\dim(W) < \dim(Y) + \dim(Z)$. (This is the same, for the theory ACF of algebraically closed fields, as the product \times_{nf} encountered in the proof of Theorem 3.5. Note that this product depends on the base set A , but as A will be fixed we will omit it from the notation.) Then $\dim(F_c(Y \times' Z)) > \dim(Z)$ for some c . This is a typical application of Zilber's stabilizer, and in itself an instance of the “sum-product” phenomenon in a constructible setting: we may assume that Z is irreducible. If $\dim(F_c(Y \times' Z)) = \dim(Z)$, we find that Y and all Γ_0 -conjugates of Y are contained in finitely many cosets of the Zilber stabilizer $H = \{y : \dim(yZ\Delta Z) < \dim(Z)\}$, up to smaller

dimension. But then $\bigcap_{x \in \Gamma_0} x^{-1}Hx$ (a finite intersection of conjugates of H) is closed under all the F_c , so by the first paragraph it equals G , i.e. contradicts the previous paragraph.

This property of $(G, F_c)_{c \in \Gamma_0}$ is referred to as *minimality*. See [24], Example 2 for details and generalizations.

5.5. The dimension inequality. Let $\Gamma \subseteq X$ be an \wedge -definable subgroup, of strict quasi-finite dimension $\delta(\Gamma) = \delta(X)$; and with $\Gamma_0 \leq \Gamma$. Let $\delta_0 = \delta(\Gamma)$, $V = V(\delta_0)$. For $n \in \mathbb{Q}$ and $v \in V$, nv is defined; we write $\gamma_0 n$ for $n\gamma_0$. Let $\gamma_0 = \delta(\Gamma) / \dim(G)$.

For a constructible $Z \subseteq G(K)^n$, define $\delta_\Gamma(Z) = \delta(Z \cap \Gamma^n)$. For any $W \subset G(K)^n$, let $\dim(W)$ denote the dimension of the Zariski closure of W .

Proposition 5.6. *For any constructible $Z \subseteq G^n$, we have $\delta_\Gamma(Z) \leq \gamma_0 \dim(Z)$.*

Proof. This is a special case of Corollary 1.12 of [24], as generalized in Remark 1.11. We give the proof in the present case, for $Z \subseteq G$. For G^n , see the remark below.

Let W be an \wedge -definable subset (over A) of Γ^n . Call W *unbalanced* if $\delta(W) > \gamma_0 \dim(W)$. There exists a complete type $W' \subset W$ defined over A with $\delta_\Gamma(W') = \delta_\Gamma(W)$. As $\dim(W') \leq \dim(W)$, W' is unbalanced if W is.

We must show that no unbalanced sets exist. Otherwise, let Y, Z be unbalanced \wedge -definable sets with $\dim(Y)$ minimal, and $\dim(Z)$ maximal possible. Clearly $0 < \dim(Y) \leq \dim(Z) < \dim(G)$. Say Y, Z, c are defined over the countable $A \leq K$. By the above, we may take Y, Z to be complete types over A . Form $Y \times' Z$. By minimality of $(G, F_c)_{c \in X}$ there exists $c \in X$ with $\dim(F_c(Y \times' Z)) > \dim(Z)$.

We note first that $Y \times' Z$ is balanced: Let f be the restriction of F_c to $Y \times' Z$. Then since $Y \times' Z$ implies a complete quantifier-free type over A in the language of fields, the fiber dimension $\dim f^{-1}(a)$ is constant ($= b$) for $a \in F_c(Y \times' Z)$, and from $\dim(f(Y \times' Z)) > \dim(Z)$ it follows that $b < \dim(Y)$. So the fibers are not unbalanced, i.e. $\delta_\Gamma(f^{-1}(a)) \leq b\gamma_0$. On the other hand since $\dim(f(Y \times' Z)) > \dim(Z)$, $f(Y \times' Z)$ is not unbalanced either, so $\delta_\Gamma(f(Y \times' Z)) \leq \dim(f(Y \times' Z))\gamma_0$. By subadditivity we obtain $\delta_\Gamma(Y \times' Z) \leq (b + \dim(f(Y \times' Z)))\gamma_0 = \dim(Y \times' Z)\gamma_0$.

Now there exists a complete type Q over A with $Q \subseteq Y \times Z$ and $\delta_\Gamma(Q) = \delta_\Gamma(Y) + \delta_\Gamma(Z)$. I claim that $Q = Y \times' Z$ (formed over A). For if Q is any other type, the fibers $Q_a = \{w : (w, a) \in Q\}$ have dimension $\dim(Q_a) = b' < \dim(Y)$ for $a \in Z$ (the dimension is constant on Z since Z is a complete type). By minimality of $\dim(Y)$ we have $\delta_\Gamma(Q_a) \leq b'\gamma_0$. By subadditivity it follows that

$$\delta_\Gamma(Y) + \delta_\Gamma(Z) = \delta_\Gamma(Q) \leq b'\gamma_0 + \delta_\Gamma(Z) < \dim(Y)\gamma_0 + \delta_\Gamma(Z) \leq \delta_\Gamma(Y) + \delta_\Gamma(Z).$$

So $Q = Y \times' Z$.

Since $\gamma_0 \dim(Y) < \delta_\Gamma(Y)$ and $\gamma_0 \dim(Z) < \delta_\Gamma(Z)$, we have $\gamma_0 \dim(Y \times' Z) = \gamma_0(\dim(Y) + \dim(Z)) < \delta_\Gamma(Y) + \delta_\Gamma(Z) = \delta_\Gamma(Y \times' Z)$. So $Y \times' Z$ is unbalanced, a contradiction. \square

Remark 5.7. Let $f : X \rightarrow X'$ be a constructible map, Γ an \wedge -definable subset of X , $\Gamma' = f(\Gamma)$. If the inequality of Proposition 5.6 holds for $\Gamma' \subseteq X'$ and for each fiber $f^{-1}(a) \subseteq X$, $a \in \Gamma'$, all with the same value of γ_0 , then it holds for $\Gamma \subseteq X$. This is an easy consequence of subadditivity and definability of Zariski dimension; cf. [24].

Recall that a morphism $f : U \rightarrow V$ of algebraic varieties is *dominant* if there exists no proper subvariety V' of V such that the image of U , over any field, is contained in V' .

Lemma 5.8. *Let U be a Zariski open subset of G^m . Let $f : U \rightarrow W \subseteq G^n$ be a dominant morphism of varieties. Then $\delta(f(U \cap \Gamma^m)) = \dim(W)\gamma_0$.*

Proof. We first show that if U is Zariski open in G^m , then $\delta_\Gamma(U) = m \dim(G)\gamma_0$. We have $\delta(\Gamma^m) = m\delta_\Gamma(G)$, On the other hand, if V is a proper Zariski closed subset of Γ^m , then $\dim(V) \leq m \dim(G) - 1$, and by Proposition 5.6, $\delta(V \cap \Gamma^m) \leq \dim(V)\gamma_0 < m\delta_\Gamma(G)$. It follows that $\delta(\Gamma^m \setminus V) = m\delta_\Gamma(G) = m \dim(G)\gamma_0$.

There exists a relatively Zariski open $W' \subseteq W$, $\dim(W') = \dim(W)$, such that $\dim f^{-1}(b)$ is constant for $b \in W'$. Replacing W by W' and U by $f^{-1}(W')$, we may assume that $\dim f^{-1}(b) = d$ is constant for $b \in W$. So $\dim(U) = d + \dim(W)$. By Proposition 5.6, for any $b \in W \cap \Gamma^n$, we have $\delta_\Gamma(f^{-1}(b)) \leq d\gamma_0$. Hence if $\delta(f(U \cap \Gamma^m)) = \gamma < \dim(W)\gamma_0$, then by subadditivity of δ we have $\delta(U \cap \Gamma^m) \leq \gamma + d\gamma_0 < (\dim(W) + d)\gamma_0 = \dim(U)\gamma_0$; this contradicts the first paragraph. Note that adding the invertible element $d\gamma_0$ preserves strict inequalities. \square

In case $f(U \cap \Gamma^m) \subseteq \Gamma^n$, it follows that $\delta_\Gamma(W) = \dim(W)\gamma_0$. The proof shows more generally that the class of subvarieties U of G^m satisfying $\delta_\Gamma(U) = \dim(U)\gamma_0$ for all Zariski dense open U' is closed under forward images of such morphisms.

Note that an \wedge -definable subgroup Γ of $G(K)$ has strict quasi-finite dimension iff for some \vee -definable \tilde{G} containing Γ , \tilde{G}/Γ is bounded.

5.9. From now on we assume C is the convex hull of \mathbb{R} in \mathbb{R}^* , a \vee -definable convex subgroup. For $Y \subseteq X$, let $\mu(Y)$ be the unique real number r such that for any rational α , $\alpha|X| > |Y|$ if $\alpha > r$ and $\alpha|X| < |Y|$ if $\alpha < r$. Then μ is a definable measure on definable subsets of X , and we have $\mu(Y) > 0$ iff $\delta(Y) = \delta(\Gamma) = \delta(X)$.

Proposition 5.10. *Let Γ be a Zariski dense \wedge -definable subgroup of $G(K)$, G a semisimple algebraic group over K . Assume Γ has strict quasi-finite dimension. Then Γ is definable.*

Proof. Assume first that G is simple. Let K^a be the algebraic closure of K . Let $\Gamma_0 \leq \Gamma$ be any Zariski dense set of points, so that the previous lemmas apply. Since G is a simple group, any noncentral conjugacy class C of $G(K^a)$ generates G in a finite number $d \leq 2 \dim(G)$ of steps. Thus for any noncentral b , the morphism of varieties $f_b : G^d \rightarrow G$, $f_b(x_1, \dots, x_d, b) = x_1^{-1}bx_1x_2^{-1}bx_2 \cdots x_d^{-1}bx_d$ is surjective on K^a -points. By Lemma 5.8, $\delta(f_b(\Gamma^d)) \geq \dim(G)\gamma_0 = \delta(G)$. Let X be a definable set containing Γ with $\delta(X) = \delta(\Gamma)$, and let \tilde{G} be the group generated by X . Then \tilde{G}/Γ is bounded. Let $S \subseteq \Gamma$ be an \wedge -definable normal subgroup of \tilde{G} with \tilde{G}/S bounded (Lemma 3.4), and choose a noncentral $b \in S$. Let Y be the definable set $Y = f_b(X^d)$. Since S is normal, $Y \subseteq S$. We have $\delta(Y) \geq \delta(f_b(\Gamma^d)) = \delta(G)$, so $\mu(Y) > 0$. Hence S contains a bounded finite number of disjoint translates $s_i Y$ of Y , so any $s \in S$ lies in $s_i Y Y^{-1}$ for some i (Ruzsa's argument). Hence $S = \bigcup_i Y Y^{-1}$ is definable. Since Γ/S is bounded and \wedge -definable, it must be finite, so Γ is definable too.

When G is semisimple, we proceed by induction on $\dim(G)$. Let N be a normal algebraic subgroup, $\pi : G \rightarrow G/N$ the natural homomorphism. Since Γ has strict quasi-finite dimension, for some \vee -definable \tilde{G} we have that \tilde{G}/Γ is bounded. It follows that $(N \cap \tilde{G})/(N \cap \Gamma)$ and $\pi(\tilde{G})/\pi(\Gamma)$ are bounded, so $N \cap \Gamma$ and $\pi(\Gamma)$ have strict quasi-finite dimension; by induction they are definable. By Remark 4.10, Γ is definable. \square

Proof of Theorem 1.3, and Corollary. Suppose not. Then there exists an ultraproduct (K, X) of (K_i, X_i) such that for no definable subgroup H of $G(K)$ do we have $H \subseteq (X^{-1}X)^2$ and X contained in finitely many cosets of H . However X is a near-subgroup of $G(K)$. Let \tilde{G} be the subgroup generated by X . By Theorem 3.5 there exists an \wedge -definable group $\Gamma \subseteq (X^{-1}X)^2$, normal in \tilde{G} , with X/Γ bounded. By Proposition 5.10, Γ is definable. By compactness, X/Γ is finite. This contradiction proves the theorem.

The corollary easily follows, and can also be quickly proved directly in the same way: if it fails, we obtain an ultraproduct (K, X) with X an infinite near-subgroup, generating a \vee -definable group \tilde{G} strictly bigger than $(X^{-1}X)^2$, and such that no infinite definable proper subgroup of \tilde{G} is normalized by X . Let Γ be as above. Again Γ is definable; hence (being normalized by X) either $\Gamma = \tilde{G}$ or Γ is finite. If Γ is finite, then since X/Γ is bounded it is finite, contradicting the assumption that X is infinite. If $\Gamma = \tilde{G}$, then since $\Gamma \subseteq (X^{-1}X)^2$ we must have $\tilde{G} = (X^{-1}X)^2$, again a contradiction. \square

Similarly we can obtain $|S|/|X^{-1}X| \geq .9$ in Corollary 1.4; otherwise we obtain (K, X) as above and also a measure μ on \tilde{G} with no infinite definable subgroup H , contained in \tilde{G} and normalized by X , satisfying $\mu(H)/\mu(X^{-1}X) \geq .99$. But again Γ is definable, and by Theorem 3.5, $\Gamma \setminus X^{-1}X$ is contained in a union of non- μ -wide sets; by saturation and definability of Γ it is contained in finitely many such sets, so $\mu(\Gamma \setminus X^{-1}X) = 0$, a contradiction. One can also get $X_i X_i^{-1} X_i = S_i$ from the fact that $q q^{-1} q$ is a coset of S in Theorem 3.5, and that S_i has no subgroups of bounded index. I noted this stronger statement after László Pyber pointed out that the statement of Corollary 1.4 implies $X_i X_i^{-1} X_i = S_i$, using [40].

One can immediately deduce a version for arbitrary linear groups:

Corollary 5.11. *Let $k, n \in \mathbb{N}$. Then there exist $k' \in \mathbb{N}$, such that if X is a k -approximate subgroup of $GL_n(K)$ for some field K , then there exist algebraic subgroups $H \leq G$ of GL_n with H solvable and normal in G , and a subgroup Δ of G (normalized by X) with $\Delta \subseteq (X^{-1}X)^2 H$ and such that X is contained in $\leq k'$ cosets of Δ .*

The groups H, N are defined by polynomial equations in the matrix entries; these equations can be taken to have degree bounded by a function of k, n alone.

Jordan has shown that finite subgroups of linear groups are bounded, up to an Abelian part, provided they contain no nontrivial unipotent elements. (Jordan’s beautiful proof occupies some 13 pages of [29]. [9] contains a different proof in characteristic 0, due to Frobenius.) We may now extend this to say that approximate subgroups of connected Lie groups are bounded, up to a (connected, closed) solvable subgroup.

Corollary 5.12. *Let $k \in \mathbb{N}$, and let L be a connected Lie group of dimension d . Then there exist $k'' \in \mathbb{N}$, such that if X is a (finite) k -approximate subgroup of L , then there exists a $(d + 2)$ -solvable subgroup S of L such that X is contained in $\leq k''$ cosets of S .*

Proof. Let X be a k -approximate subgroup of L . Assume first that L embeds into $GL_d(\mathbb{R})$. In this case, let H, G be the subgroups provided by Corollary 5.11; so H is d -solvable. So X is contained in boundedly many cosets of a subgroup Δ of

G/H , with Δ/H finite (as X is finite). By [29], Δ/H contains a normal Abelian subgroup S/H of bounded index. Then S is $(d+1)$ -solvable, and X is contained in boundedly many cosets of S .

In general, let Z be the center of L . Then L/Z acts faithfully on the Lie algebra of L by conjugation, so it embeds into $GL_d(\mathbb{R})$. By the linear case, if X is a k -approximate subgroup of L , then the image of X in L/Z is a k -approximate subgroup of L/Z , so by the linear case it is contained in boundedly many cosets of a solvable subgroup S/Z . The pullback S of S/Z to L is $(d+2)$ -solvable, and finitely many cosets of S cover X . \square

Remarks.

- (1) Once the definability of Γ is established, it is known to be definable in the field language, possibly expanded by an automorphism, and indeed to be a simple group of (possibly twisted) Lie type; see [35].
- (2) We could also deduce Theorem 1.3 from Corollary 1.2; the proof of Proposition 5.10, together with saturation, shows that for some $m \in \mathbb{N}$ we have $\mu(C_b^d) \geq 1/m$ for all noncentral b . It follows that with probability very close to 1 (in b), $\mu(C_b^d) \geq 1/m$; so the hypothesis of Corollary 1.2 holds.
- (3) Using another direction of generalization taken in [24], results of this section are valid for near-subgroups of groups G of finite Morley rank, in place of algebraic groups.

We further remark that Theorem 1.1 of [5] in the sum-product setting, as well as the theorem of [20] for subsets of $SL_2(\mathbb{F}_p)$, can be put in the framework of Proposition 5.10 if one takes the $C = C_{max}$ to be the largest convex subgroup of \mathbb{R}^* not containing δ_0 (in place of the smallest nonzero convex subgroup, as we took it to be).²

6. UNIFORM DEFINABILITY OF THE TOPOLOGY

We prove a stronger version of the stabilizer theorem for arbitrary S1-ideals on an Ind-definable group, with more uniform control of the topology of the Lie group. It follows that the Lie group associated to a near-subgroup is always associated already to the reduct to a finite sublanguage. Stronger uniformity statements in this direction may give a more powerful means for finitization of results about saturated models.

Remark 6.1. Let T be a simple theory, or a NIP theory. Then the forking ideal is an S1-ideal.

Proof. Let (a_i) be an A -indiscernible sequence, and suppose $\phi(x, a_i)$ does not fork over A . We have to show that $\phi(x, a_i) \wedge \phi(x, a_j)$ does not fork over A , for some $i \neq j$.

Simple case: the a_i are independent over some M . Let c_i be such that $\phi(c_i, a_i)$ with c_i, a_i independent; choose c_i so that c_i, M are independent over a_i . Then c_i, Ma_i are independent over A . The sequence (a_i) could be taken to be long; by refining it we can assume that $tp(a_i/M)$ is constant. By 3-amalgamation we can find c independent over M from $(a_i)_i$, with $tp(c, a_i/M) = tp(c_i, a_i/M)$. Since

²In fact, after these lines were written, Breuillard, Green and Tao essentially took this route; using a beautiful analysis of the geometry of tori, continuing a line started in [29], they obtain an *effective, polynomial* version of Theorem 1.3. See [3].

$tp(c_i/M) = tp(c/M)$, c, M are independent over A , so c is independent from a_1, a_2 over A . Hence $tp(c'/a_1a_2)$ does not fork over A .

NIP case: Let q_i be a global type with $\phi(x, a_i) \in q_i$, such that q_i does not fork over A . Let M be a model containing A . Then q_i does not fork over M . So q_i is M -invariant. There are few choices for M -invariant types, so $q_i = q_j$ for some $i \neq j$. Since q_i does not fork over A , $\phi(x, a_i) \wedge \phi(x, a_j)$ does not fork over A . \square

Let X be a topological space, $p \in X$. We say that a collection C of sets strongly generates the topology at p if p is in the interior of each set in C , and any open neighborhood of p contains some element of C .

Let M be a Riemannian manifold. Let $\rho(p, q)$ denote the Riemannian distance, and $B(p, r)$ (respectively $\bar{B}(p, r)$) the open (resp. closed) ball of radius r . A geodesic ball around p is the image under the exponential map exp_p of a ball b of radius r around 0 in the tangent space to p , where r is small enough that exp_p is a diffeomorphism. We have $\rho(p, exp_p(v)) = |v|$ if $v \in b$ ([36], Proposition 6.10, p. 105). A subset U is called convex if for each $p, q \in U$ there is a unique geodesic x from p to q of length $\rho(p, q)$, contained entirely in U . Any point has a convex neighborhood ([36], 6-4, p.112).

Lemma 6.2. *Let M be a Riemannian manifold, G a topological group acting isometrically and transitively on M (the action $G \times M \rightarrow M$ is assumed continuous). Let $B(p_0, r)$ be a geodesic ball of M , contained in a convex set W . Assume there exists a compact $Y \subseteq G$ such that if $x, x' \in B(p_0, r)$, then for some $g \in Y$ we have $gx = x'$ and $\rho(x, g^2x) = 2\rho(x, x')$.*

Let U be any open set of diameter $< r$. Let C be the collection of neighborhoods of p_0 of the form $cl(g_1U \cap g_2U)$. Then C strongly generates the topology at p_0 .

Proof. It suffices to show that there are nonempty sets of the form $g_1U \cap g_2U$, of arbitrarily small diameter. For then by translation we may take these sets to contain p_0 , and their closures will still have small diameter, and will strongly generate the topology at p_0 .

Let \bar{U} be the closure of U , and let $\delta < r$ be the diameter of U .

Find $p_n, q_n \in U$ with $\rho(p_n, q_n) \geq \delta - 1/n$, and find $g_n \in G$ with $g_n p_n = q_n$ and $\rho(p_n, g_n^2 p_n) = 2\rho(p_n, q_n)$. By assumption, we may choose g_n in a compact set, and all p_n, q_n lie within a compact set (a closed ball of radius r). Refining the sequence (p_n, q_n, g_n) , we may thus assume it converges to a point $(p, q, g) \in \bar{U}^2 \times G$, and we have $\rho(p, q) = \delta$, $\rho(p, g^2 p) = 2\rho(p, q) = 2\delta$. It follows from uniqueness of the minimizing geodesic between p and $g^2 p$ that $\bar{B}(p, \delta) \cap \bar{B}(g^2 p, \delta) = \{q\}$. By definition of δ we have $\bar{U} \subseteq \bar{B}(x, \delta)$ for any $x \in \bar{U}$. In particular, $\bar{U} \subseteq \bar{B}(p, \delta)$, and $\bar{U} \subseteq \bar{B}(q, \delta)$. From the latter we obtain $g\bar{U} \subseteq \bar{B}(g^2 p, \delta)$. So $\bar{U} \cap g\bar{U} = \{q\}$.

The set $U \cap g_n U$ is nonempty, since $q_n \in U \cap g_n U$. It remains only to show that the diameter of $U \cap g_n U$ approaches 0 as $n \rightarrow \infty$.

Suppose otherwise; then there exist $\gamma > 0$, and $a_n, b_n \in U \cap g_n U$ such that $\rho(a_n, b_n) \geq \gamma$ for infinitely many n . We can refine the sequence again to assume $a_n \rightarrow a, b_n \rightarrow b$; we have $\rho(a, b) \geq \gamma$, so $a \neq b$, and $a, b \in \bar{U} \cap g\bar{U}$. But we have seen that $\bar{U} \cap g\bar{U}$ consists of a single point, a contradiction. \square

The hypotheses of Lemma 6.2 are satisfied when G is a Lie group, acting on itself by left translation, M is G with a left-invariant Riemannian metric, and $B(p_0, 2r)$ is a geodesic ball. For then $Y = \bar{B}(p_0, r)\bar{B}(p_0, r)^{-1}$ is compact. For $x, x' \in B(p_0, r)$, let $g = x'x^{-1}$, $h = x^{-1}gx = x^{-1}x'$, and let $|u| = \rho(1, u)$. Then $gx = x'$. We have

$\rho(x, x') = \rho(x', g^2x) = \rho(1, x^{-1}gx) = |h|$, $\rho(x, g^2x) = \rho(1, x^{-1}g^2x) = |h^2|$, so we have to show that $|h^2| = 2|h|$. We have $h = \exp(v)$ for some v , where \exp is the the exponential map at 1, $h^2 = \exp(2v)$, and $|h^2| = |2v| = 2|v| = 2|h|$.

Corollary 6.3 (Stabilizer theorem). *Let X be a near-subgroup of G .*

- *There exist a \forall -definable \check{G} and an \wedge -definable normal subgroup $\Gamma \subseteq \check{G}$, both defined without parameters, such that \check{G}/Γ is bounded, and any definable D with $\Gamma \leq D \leq \check{G}$ is commensurable to $X^{-1}X$.*
- *There exist a connected Lie group L and a homomorphism $\pi : \check{G} \rightarrow L$ with dense image, and kernel Γ . If D is a definable subset of G , write πD for the closure of $\pi(D)$. π intertwines the definable sets containing Γ , contained in \check{G} with the compact neighborhoods of L .*
- *There exist a uniformly definable family of definable sets D_a , and a definable set E , with $\partial\pi(D_a) \cap \pi E \subset \text{int}(\pi(E))$ such that the neighborhoods of 1 of the form $\pi E \setminus \pi D_a$ generate the topology of L at 1.*

Proof. The first two parts follow from Theorem 4.2.

There remains to prove the uniform generation of the topology of L . Fix a left-invariant Riemannian metric on L , and view $M = L$ as a Riemannian manifold. Let $p_0 = 1$ and let r be as in Lemma 6.2; renormalizing, we may assume $r = 4$. Write B_s for $B(1, s)$. By the above remark there exists a definable E with $\bar{B}_5 \subseteq \pi E \subseteq B_6$. Similarly there exists a definable D such that πD contains $\bar{B}_9 \setminus B_2$ (so $\partial(\pi(D)) \cap E \subset \text{int}(E)$) and is disjoint from \bar{B}_1 . Then $U = \pi(E) \setminus \pi(D) = B_2 \setminus \pi(D)$ is an open neighborhood of 1. By Lemma 6.2, there exists $g, g' \in B_3$ with $gU \cap g'U$ of arbitrarily small diameter, and containing 1. We compute $U \cap gU = (\pi(E) \cap \pi(gE)) \setminus \pi(D \cup gD) = B_2 \setminus \pi(D \cup gD) = \pi(E) \setminus \pi(D \cup gD)$, and similarly for $gU \cap g'U$. The uniformly definable family is the family of unions $D \cup gD$. \square

6.4. The locally compact Lascar group. Let T be a theory, \mathbb{U} a universal domain, \tilde{E} a \forall -definable equivalence relation, Σ an \wedge -definable equivalence relation, such that Σ implies \tilde{E} . Let P be a complete type. Let \tilde{a} be a class of \tilde{E} restricted to P , such that $\tau = \tilde{a}/\Sigma$ is bounded. Let $\pi : \tilde{a} \rightarrow \tilde{a}/\Sigma$ be the quotient map. Then \tilde{a}/Σ admits a natural locally compact topology, generated by the complements of the images $\pi(D)$ of definable sets. $G = \text{Aut}(\mathbb{U}/\tilde{a})$ acts on τ . Let K be the kernel of this action, and $L = G/K$. Then L admits a natural locally compact group structure; we call it the locally compact Lascar group of (\tilde{a}, Σ) .

We have transposed from definable groups (as in Theorem 6.3) to automorphism groups. In both cases, the set of conjugates of a definable set lie in a uniformly definable family. We will use this in Lemma 6.6 below.

6.5. The compact Lascar group. So far, the case where \tilde{a} is a definable set and L is compact has been useful. For simplicity, we will also restrict to this case in the statement below. *For the rest of this section we assume \tilde{E} is the indiscrete equivalence relation, so $\tilde{a} = P$ and $\tau = P/\Sigma$ is compact.* We do not expect any trouble in generalizing to the locally compact case.

Lemma 6.6. *Let $L' = L/N$ be a finite-dimensional quotient of L , so N is a compact normal subgroup and L' is a compact Lie group. For large enough k , L' has a regular orbit on $\tau^k/N = P^k/N$. Let τ' be such an orbit. There exists a uniformly definable family of definable sets D_a , such that the sets $\tau' \setminus \pi(D_a)$ strongly generate the topology on τ' at every point.*

Proof. For $x \in \tau$, let S_x be the stabilizer of x . Let Ξ be the set of finite subsets of τ . For $u \in \Xi$, let $S_u = \bigcap_{x \in u} S_x$. We have $\bigcap_{x \in \tau} S_x = K$. Since N is compact, $\bigcap_{u \in \Xi} S_u N = KN = N$. (Let $a \in \bigcap_{u \in \Xi} S_u N$. Pick an ultrafilter on Ξ including all sets of the form $\{u : x \in u\}$. Write $a = s_u n_u$ with $n_u \in N, s_u \in S_u$. Then $s_u \rightarrow s$ and $n_u \rightarrow n$ for some s, n . We have $s \in \bigcap_x S_x = K$ and $n \in N$.)

Now L' is a compact Lie group, so it has no infinite descending sequences of closed subgroups. Thus for some finite tuple $u = (x_1, \dots, x_k)$ we have $S_u N = N$. It follows that $\tau' = L'x$ is a regular orbit in τ^k/N . The uniformity statement follows from Lemma 6.2, as in the proof of Theorem 6.3; since compactness is assumed, we can take E to be the entire ambient sort. \square

An earlier version of this section attempted an application to SOP theories, but in this, Krzysztof Krupinski found a gap.

7. GROUPS WITH LARGE APPROXIMATE SUBGROUPS

In this section we aim to prove:

Theorem 7.1. *Let G_0 be a finitely generated group, $k \in \mathbb{N}$. Assume G_0 has a cofinal family of k -approximate subgroups (i.e., any finite $F_0 \subset G_0$ is contained in one). Then G_0 is nilpotent-by-finite.*

This generalizes Gromov’s theorem [18], asserting the same conclusion if G_0 has polynomial growth. There is by now a small family of proofs of Gromov’s theorem and extensions, descending from either Gromov’s original proof or Kleiner’s; the first may have been [13], and the most recent, [46]. I believe all view the group as a metric space, via the Cayley graph, and analyze it either geometrically or analytically.

We will consider an arbitrary sequence of approximate subgroups, rather than balls in the Cayley graph. A Lie group L lies at the heart of the proof, as in the case of Gromov’s. While Gromov’s group arises in the automorphism group of the Cayley graph “viewed from afar”, we find L and a homomorphism $h : G_0 \rightarrow L$ using the model-theoretic/measure-theoretic construction Theorem 3.5, which has no metric aspect.

Beyond this point, our proof will adhere very closely to the outline of Gromov’s. If the homomorphism into L is trivial, we conjugate it to a nontrivial one in exactly the way taken by Gromov, succeeding unless G_0 is already virtually Abelian³ (in which case we are already done). We now use the earlier Corollary 5.12 covering the linear case to show that the image is essentially solvable, and hence a non-trivial homomorphism into an Abelian group can be obtained. Gromov used the Tits alternative at the parallel point. We show that the kernel satisfies the same assumptions as G_0 ; here we make some further use of Lie theory. Induction is carried out on the Lie dimension, rather than the growth rate exponent which is not available to us; we conclude that the group is polycyclic-by-finite, and in particular virtually solvable. To pass from the polycyclic solvable to the nilpotent case, we quote Tao [53] or Breuillard-Green [4] where Gromov cited Milnor-Wolf.

We will see along the way that G_0 is polycyclic-by-finite with d infinite cyclic factors, where d is the dimension of the associated Lie group.

³We say that a group G_0 is *virtually* P if some finite index subgroup is P .

An alternative statement is that when G_0 is *not* nilpotent-by-finite, then for some finite $F_0 \subset G_0$, G_0 has *no* k -approximate subgroups containing F_0 . If one wishes to seriously use the ambient group G_0 , some hypothesis on containing sets of generators is necessary (e.g. since any countable family of finitely generated groups embeds jointly in a single one).

The strongest possible general conjecture on the structure of k -approximate subgroups would be this: for some k', k'' , any k -approximate subgroup of a group G is k' -commensurable with one induced by a map into a k'' -nilpotent group. Here we say that X is induced by h if h is a homomorphism on some subgroup A of G into a group N , and $X = h^{-1}h(X)$. Statements in this vein, possibly restricted to approximate subgroups of a fixed group, have been suggested by Helffgott, E. Lindenstrauss, Breuillard and Tao.

The conjugation method used in the present section would be powerless in the following scenario: X_n is a k -approximate subgroup of the alternating group A_n , and X_n is conjugation-invariant.

Towards the proof of Theorem 7.1, we will study the following situation \diamond :

- A language with two sorts G, Φ ; G carries a group structure; a relation on $G \times \Phi$ defines a family of definable subsets of G , $(X_c : c \in \Phi)$. Additional structure is allowed.
- M^* is a saturated structure, with an elementary submodel M .
- $G_0 = G(M)$ is finitely generated.
- $X = X_{c^*}$ is a c^* -definable subset X with $G(M) \subset X$ (c^* is an element of $\Phi(M^*)$).
- For all $c \in \Phi(M)$, X_c is finite.
- There is an \wedge -definable subgroup Γ of G , and a \vee -definable subgroup \tilde{G} , with $\Gamma \subseteq X \subseteq \tilde{G}$, and \tilde{G}/Γ bounded. Γ, \tilde{G} are defined over some small subset of M^* .
- Any subgroup of $G(M)$ has the form $S(M)$ for some 0-definable subgroup S of G .

In this situation, note:

- (1) We may replace \tilde{G} by the group generated by X , without disturbing the hypotheses.
- (2) Let G' be a 0-definable subgroup of G ; $X'_c = X_c \cap G'$; $\Gamma' = \Gamma \cap G'$, $\tilde{G}' = \tilde{G} \cap G'$. Then \diamond holds with the new data, except possibly for the finite generation of $G'(M)$. When G' has finite index in G , this too holds.
- (3) There exists a \vee -definable $\check{G} \leq \tilde{G}$ and a normal \wedge -definable subgroup Γ' of \check{G} containing $\check{G} \cap \Gamma$, such that \check{G}/Γ' is a connected Lie group. (This is Theorem 4.2. We have $\check{G} \cap \Gamma \subseteq \Gamma'$ since the image of $\check{G} \cap \Gamma$ in \check{G}/Γ' is a compact normal subgroup, hence trivial.)
- (4) X is contained in finitely many cosets of \check{G} (the image of X modulo \check{G} is a compact subset of the discrete space \tilde{G}/\check{G}).
- (5) $G_0 \cap \check{G}$ has finite index in G_0 (since $G_0 \subseteq X$, by (4), G_0 is contained in finitely many cosets of \check{G} , equivalently of $G_0 \cap \check{G}$).
- (6) Let $H_0 = G_0 \cap \check{G}$; let H be a 0-definable group, with $H_0 = H(M)$. So H has finite index in G . Let $\tilde{H} = \tilde{G} \cap H$, let Y be a definable subset of \tilde{H} commensurable with X , and containing $X \cap \tilde{H}$, with corresponding family $(Y_c : c \in \Phi')$. We choose $\Phi' \subset \Phi$ so that Y_c is commensurable to X_c for

$c \in \Phi'$: in particular, Y_c is finite for $c \in \Phi'(M)$. So $H(M) \subset Y$. Now the hypotheses \diamond hold for $(H, Y, \check{H}, \Gamma \cap H)$. Let $\check{H} = \check{G} \cap H$ and $\Gamma'' = \Gamma' \cap H$. Then \check{H}/Γ'' has finite index in \check{G}/Γ' , but the latter is connected, so they are equal. Hence \check{H}/Γ'' is connected, and $H_0 = G_0 \cap \check{G} \subseteq \check{G} \cap H = \check{H}$. Now we are in the same situation \diamond , but have in addition $H_0 \leq \check{H}$.

Before entering the proof proper, we can clarify the meaning of this setup by looking at the Lie rank zero case.

Lemma 7.2. *Assume \diamond , and further assume that \check{G}/Γ is totally disconnected. Then $G = G_0$ is finite.*

Proof. Being totally disconnected, \check{G}/Γ contains a compact open subgroup C . If $\psi : \check{G} \rightarrow L$ is the canonical map, then $H = \psi^{-1}(C)$ is a definable group by compactness of C and is commensurable with X by openness. Since X contains $G_0 = G(M)$, H is covered by finitely many cosets of G_0 , so $H \cap G_0$ has finite index in G_0 . In particular it is finitely generated. Let F_1 be a finite set of generators for $H \cap G_0$. Since $M \prec M^*$, there exists a definable group H_c containing F_1 and commensurable with X_c , for some $c \in \Phi(M)$; so H_c is finite. It follows that the group generated by F_1 is finite, i.e. $H \cap G_0$ is finite; thus G_0 is finite. \square

We will need some lemmas on finite generation. First, if E is a finitely generated group, N a normal subgroup with E/N finitely presented, then N is finitely generated as a normal subgroup. (In particular when E/N is finite, this implies the finite generation of N , a well-known statement used above.) This in fact is valid for any equational class: *If E is finitely generated and N is a congruence with E/N finitely presented, then N is finitely generated as a congruence.* Indeed let F be a finitely generated free algebra in this equational class, and $h : F \rightarrow E$ a surjective homomorphism. Let $g : F \rightarrow E/N$ be the composition $F \rightarrow E \rightarrow E/N$. Since E/N is finitely presented, g has a finitely generated kernel K . Thus $N = h(K)$ is finitely generated.

A $\check{\vee}$ -definable subgroup is called *definably generated* if it is generated by a definable subset. If G is a topological group, let G^0 denote the connected component of 1; it is a closed normal subgroup of G .

Let H be a sufficiently saturated group (with possible additional structure), \check{H} be a $\check{\vee}$ -definable subgroup, Γ an $\check{\wedge}$ -definable subgroup, with $\Gamma \leq \check{H}$. Let $\pi : \check{H} \rightarrow \check{H}/\Gamma$ be the quotient map. Recall the logic topology on \check{H}/Γ from §4. In particular, a subset Z of the quotient is compact iff $\pi^{-1}(Z)$ is contained in a definable set.

Lemma 7.3. *Let $H, \check{H}, \Gamma, \pi$ be as above, and assume $A = \check{H}/\Gamma$ is locally compact. Assume A/A^0 is finitely generated. Then \check{H} is definably generated.*

Proof. Let U be a compact neighborhood of 1 in A . Then $\pi^{-1}(U)$ is contained in a definable subset D of \check{H} . U generates an open subgroup of A ; this open subgroup is also closed, and must contain A^0 . On the other hand, A/A^0 is generated by finitely many elements $\pi(h_1), \dots, \pi(h_r)$. Let $D' = D \cup \{h_1, \dots, h_r\}$. Since D contains $\ker \pi$ and $\pi\check{H}$ is generated by $\pi(D')$, it follows that \check{H} is generated by D' . \square

In fact if A/A^0 is m -generated, the proof shows that \check{H} is generated by Y along with m additional elements, whenever Y is a definable set containing Γ .

Lemma 7.4. *Let G be a sufficiently saturated group (with additional structure), H a definable normal subgroup with G/H Abelian. Let \check{G} be a $\check{\vee}$ -definable subgroup of G , Γ an $\check{\wedge}$ -definable subgroup, with $\Gamma \trianglelefteq \check{G}$ and $E = \check{G}/\Gamma$ a connected Lie group. Then $H \cap \check{G}$ is definably generated.*

Proof. Let $\check{H} = H \cap \check{G}$, and $\pi : \check{G} \rightarrow E$ be the canonical map. By Remark 4.10, $\pi|_{\check{H}}$ induces an isomorphism of topological groups $\check{H}/(\check{H} \cap \Gamma) \cong \pi(\check{H})$.

Since H is normal in G , \check{H} is normal in \check{G} , so $\pi(\check{H})$ is normal in E , and hence so is $\pi(\check{H})^0$. The commutator subgroup $[E, E]$ is contained in $\pi(\check{H})^0$, so the quotient $E/\pi(\check{H})^0$ is isomorphic to $\mathbb{R}^n \times \mathbb{R}^m/\mathbb{Z}^m$. The image of $\pi(\check{H})$ in $E/\pi(\check{H})^0$ contains no nontrivial connected groups, so it is discrete. Now it is well known that a discrete subgroup of $\mathbb{R}^n \times \mathbb{R}^m/\mathbb{Z}^m$ is finitely generated, indeed admits a generating set with at most $n + m$ elements. By Lemma 7.3, $H \cap \check{G}$ is definably generated. \square

The next lemma will play an essential role in the proof, allowing the key Lemma 7.6 to be propagated. It continues to hold if G/N is assumed to be nilpotent, rather than Abelian; indeed it suffices to find a sequence of definable subgroups $G = H_1 \supset \dots \supset H_k = H$ with H_{i+1} normal in H_i , and H_i/H_{i+1} Abelian, and apply the lemma inductively.

Lemma 7.5. *Assume \diamond holds, and let N be a 0-definable normal subgroup of G .*

Then \diamond holds if G is replaced by G/N , and X, \check{G}, Γ by their images in G/N .

If G/N is finite or Abelian, then \diamond holds if G, X, \check{G}, Γ are replaced by $N, X \cap N, \check{G} \cap N, \Gamma \cap N$.

Proof. The first statement is straightforward; so is the second, except for the finite generation of $N_0 = N(M)$. We proceed to show this. In case G/N is finite, so is G_0/N_0 since M is an elementary submodel, so finite generation is clear. Assume therefore that G/N is Abelian.

Let $G_1 = G_0 \cap \check{G}$. As G_1 has finite index in G_0 by (5), it is a finitely generated group.

Let $N_1 = N \cap G_1$. Since $G_1/(N \cap G_1)$ is finite or Abelian, it is a finitely presented group. By the remarks above, N_1 is finitely generated as a normal subgroup of G_1 .

Let g_1, \dots, g_r be generators for G_1 , and let $T_i(x) = g_i^{-1}xg_i$. Then N_1 is finitely generated as a group with these operators. Let Y be a finite subset of N_1 such that N_1 is generated by Y under multiplication and the operators T_i .

By Lemma 7.4, $\check{G} \cap N$ is generated by an M^* -definable set U . We may take $Y \subset U = U^{-1}$. Since \check{G} and N are closed under the operators T_i (as $g_i \in \check{G}$ and N is normal), we have $T_i(U) \subset U \dots U = U^m$ for some m . Since U is a definable subset of \check{G} , it is contained in finitely many translates of X . Now M is an elementary submodel of M^* . So there exists an M -definable set $U' \subset N$ containing Y , with $T_i(U') \subset U' \dots U'$, and U' contained in finitely many translates of some $X_c, c \in \Phi(M_0)$. From the last property it follows that U' is finite; so $U' \subset M$ and hence $U' \subset G_0$. Thus $U' \subset N_1$. Moreover the group generated by U' is closed under the operators T_i , and contains Y . So it equals N_1 . This shows that N_1 is a finitely generated group. Since it has finite index in $N_0 = N \cap G_0$, it follows that N_0 too is a finitely generated group. \square

Lemma 7.6. *Assume \diamond holds, and G_0 is infinite. Then there exists a normal subgroup N_0 of G_0 with G_0/N_0 virtually Abelian, and infinite.*

Proof. Let $G'_0 = G_0 \cap \check{G}$. By note (5) above, G'_0 has finite index in G_0 . If N'_0 is normal in G'_0 with infinite virtually Abelian quotient, let N_0 be the intersection of the finitely many G_0 -conjugates of N'_0 ; then G_0/N_0 is infinite and virtually Abelian. Thus proving the lemma for G'_0 would imply it for G_0 . By note (6), the hypotheses hold for G'_0 ; so we may assume $G_0 \leq \check{G}$.

Let $L = \check{G}/\Gamma'$ as in note (3), $d = \dim(L)$, and consider the natural homomorphism $\psi : \check{G} \rightarrow L$. Note that $\psi(G_0)$ has a cofinal system of k -approximate subgroups. By Corollary 5.12, any finite subset w of $\psi(G_0)$ is contained in at most k'' cosets of a $(d+2)$ -solvable subgroup S_w of L . Taking an ultraproduct, L embeds in an ultraproduct of itself, in such a way that the image of $\psi(G_0)$ is contained in at most k'' cosets of a $(d+2)$ -solvable group S . Thus $\psi(G_0)$ has a solvable subgroup S' of finite index. If S' is infinite, then it contains a subgroup S'' of finite index, such that $S''/[S'', S'']$ is infinite. Thus $\psi^{-1}(S'')$ is a finite index subgroup of G_0 , and $\psi^{-1}([S'', S''])$ is a normal subgroup with infinite Abelian quotient. So we are done unless S' above is finite, i.e. $\psi(G_0)$ is finite, so that a finite index subgroup H_0 of G_0 is contained in Γ' . We have $H_0 = H(M)$ for some 0-definable subgroup H of G .

For $g \in G$, let $ad_g(x) = g^{-1}xg$, and let $\tau_g = ad_g|_{H_0}$. Let $J = \{g \in G : \tau_g(H_0) \leq \check{G}\}$. If $g \in J$, we may repeat the previous paragraph with $\psi \circ \tau_g$ in place of ψ . Thus again we are done unless $\psi \circ \tau_g(H_0)$ is finite for any $g \in J$. We thus assume this is the state of affairs.

The rest of the proof is a straightforward transcription of the corresponding part of [18]. By Jordan's theorem [29], since $\psi \circ \tau_g(H_0)$ is a finite subgroup of the Lie group L , it has an Abelian subgroup S_g of index $\leq \mu$, with μ independent of g . If $\psi \circ \tau_g$ can have arbitrarily large finite size for $g \in J$, taking an ultraproduct, we obtain a homomorphism to a group with an infinite Abelian subgroup of index $\leq \mu$. Thus in this case too the lemma is proved, and we may assume that $\psi \circ \tau_g(H_0)$ has size $\leq \mu'$ for some fixed μ' .

Let F_1 be a finite set of generators for H_0 . Let U be a neighborhood of the identity in the Lie group L , such that if $u \in U$ is an element of order $\leq \mu'$, then $u = 1$. (For instance we can take a neighborhood V of the Lie algebra on which the exponential map is injective, and then let $U = \exp((1/\mu')V)$.) Since $\{1\}$ is closed and U is open, there exists a definable set $D_2 \subset G$ with $\Gamma' \subset D_2 \subset \psi^{-1}(U)$. Since $F_1^{-1}\Gamma'F_1 = \Gamma' \subset D_2$, we can find a definable set D with $\Gamma' \subset D$ and $F_1^{-1}DF_1 \subset D_2$. Any subgroup of D_2 of size $\leq \mu'$ is trivial. Now if $\tau_g(F_1) \subset D_2$, then $\tau_g(H_0) \leq \check{G}$, so $g \in J$. Hence $\psi \circ \tau_g(H_0)$ has size $\leq \mu'$, but $\psi \circ \tau_g(F_1)$ is a set of elements of U , and any such nonidentity element has order $> \mu'$; so $\psi \circ \tau_g(F_1)$ must reduce to the identity element of L . Hence if $\tau_g(F_1) \subset D_2$, then $\tau_g(H_0) \subset \Gamma'$, and in particular $\tau_g(F_1) \subset D$.

Let $W = \{g : \tau_g(F_1) \subset D\} = \{g : \tau_g(F_1) \subset D_2\}$. If $g \in W$ and $f \in F_1$, then $gf \in W$, since $(gf)^{-1}F_1gf \subseteq f^{-1}Df \subseteq D_2$. So W is a definable, right F_1 -invariant set. Now in the model M , any definable, right F_1 -invariant set is empty or contains H_0 . Since $M \prec M^*$, it follows that $W = \emptyset$ or W contains H . We have $1 \in W$, as $H_0 \leq \Gamma' \leq D$. So all H -conjugates of F_1 are contained in D . Note that D is contained in the union of finitely many translates of X . It follows that all H_0 -conjugates of F_1 are contained in finitely many translates of some X_c , $c \in \Phi(M)$. In this case each element of F_1 has centralizer of finite index in H_0 , so H_0 has a center of finite index; we may take G'_0 to be this center, and $N = 1$. \square

We will need some elementary group-theoretic discussion before proceeding. We define a group H to be 0-polycyclic if it is trivial, and to be $(d + 1)$ -polycyclic if it has a d -polycyclic normal subgroup N , with H/N a finitely generated Abelian group of rank 1. In particular, H is d -solvable.

For $d \geq 0$, we say that a finitely generated group H is almost d -polycyclic if it has subgroups $H_{2d+2} \trianglelefteq H_{2d+1} \trianglelefteq H_{2d} \trianglelefteq \cdots \trianglelefteq H_1 = H$ with H_{2i}/H_{2i+1} finite ($0 \leq i \leq d$), and $H_{2i+1}/H_{2i+2} \cong \mathbb{Z}$ ($0 \leq i \leq d$), and $H_{2d+2} = 0$. By reverse induction on i we see that each H_i is finitely generated in this situation.

These definitions differ in that the quotients H_{2i}/H_{2i+1} are not required to be Abelian, but if H is almost d -polycyclic, then it does have a d -polycyclic normal subgroup of finite index. To show this we may pass to a finite index subgroup, so we may assume H has an almost $(d - 1)$ -polycyclic normal subgroup N with $H/N \cong \mathbb{Z}$. Like all almost polycyclic groups, N is finitely generated. Using the induction hypothesis, let N_1 be a $(d - 1)$ -polycyclic subgroup of N , with $[N : N_1] = r < \infty$. As N is finitely generated, it has only finitely many subgroups of index r . Let N_2 be their intersection. Then N_2 is $(d - 1)$ -polycyclic and is characteristic in N , hence normal in H . H/N_2 contains the finite group N/N_2 as a normal subgroup; within H/N_2 , the centralizer of N/N_2 has the form H_1/N_2 , with H_1 a finite index subgroup of H . Now $H_1/(N \cap H_1) \cong \mathbb{Z}$, while $(N \cap H_1)/N_2$ is a finite central subgroup of H_1/N_2 ; so H_1/N_2 is a finitely generated Abelian group of rank 1. Thus H_1 is d -polycyclic.

Lemma 7.7. *Assume \diamond . Then $G_0 = G(M)$ is polycyclic-by-finite (and in particular solvable-by-finite).*

Proof. We use induction on $d = \dim(L)$, $L = \check{G}/\Gamma'$. If $d = 0$, then \check{G}/Γ is totally disconnected; hence G_0 is finite by Lemma 7.2. For higher d , we use Lemma 7.6. By note (2) to \diamond we may pass to a finite index subgroup; so we may assume there exists a 0-definable normal subgroup N of G , with G/N infinite Abelian. By Lemma 7.5, the hypotheses \diamond hold for $N, X \cap N, \check{G} \cap N, \Gamma \cap N$, and also for the images in G/N . By Lemma 7.2 applied to G/N , we see that the image of \check{G} in G/N has Lie rank ≥ 1 . By Remark 4.10 (3) it follows that $\check{G} \cap N$ has Lie rank $< d$. So the inductive hypothesis applies, and $N_0 = N(M)$ is polycyclic-by-finite. Thus G_0 is almost polycyclic and hence also polycyclic by finite. \square

We are now essentially in the solvable case, and can quote either [53], or [4]. Polycyclicity is a strong additional tool, and with it one may be able to reduce to the sum-product phenomenon for fields somewhat more rapidly; a model-theorist is reminded here of Zilber's arguments in the 1970s, connecting solvable groups of finite Morley rank with definable fields. We will simply invoke [4]; thanks to Emmanuel Breuillard for pointing out a nicer path to this paper than we had initially.

Lemma 7.8. *Assume \diamond . Then $G_0 = G(M)$ is nilpotent-by-finite.*

Proof. By Lemma 7.6, G_0 has a normal subgroup N_0 with G_0/N_0 virtually Abelian; as we just saw, N_0 satisfies \diamond , has lower Lie rank, and so inductively is nilpotent by finite. Passing to a finite index subgroup, we may assume G_0 is solvable, as well as polycyclic. Now any polycyclic group is linear over \mathbb{C} ; see chapter 4 of [45] for a stronger result, due to Auslander-Swan. Hence G_0 can be viewed as a solvable subgroup of $GL_n(\mathbb{C})$; according to [4], every k -approximate subgroup of

G_0 is covered by a bounded number of cosets of an $(n - 1)$ -step nilpotent subgroup of $GL_n(C)$. Taking the ultraproduct of the cofinal family of approximate groups associated with G_0 , we see that X and hence G_0 are covered by finitely many cosets of an $(n - 1)$ -step nilpotent group; hence G_0 is itself virtually nilpotent. \square

Proof of Theorem 7.1. Let $(X_c : c \in \Phi_0)$ be the given family of k -approximate subgroups of G_0 . Consider the two-sorted structure (G, Φ_0, \cdot, E) where $(x, c) \in E$ if $x \in X_c$. Enrich it by adding a predicate for each subgroup of G_0 . Further enrich the language by closing under probability quantifiers as in §2.6. Let M be the resulting structure, and let M^* be a saturated elementary extension. By saturation and by the cofinality of the X_c , there exists $c^* \in \Phi(M^*)$ with $G_0 \subset X_{c^*}$. All clauses of \diamond are now clear, so by Lemma 7.8, G_0 is nilpotent-by-finite. \square

ACKNOWLEDGMENT

This paper was prepared for a talk at the Cherlin Bayramı in Istanbul in June 2009. The author is grateful to Dugald Macpherson for a valuable conversation on this subject in Leeds. Thanks also to Zoé Chatzidakis, Lou van den Dries, Ward Henson, Itay Kaplan, Krzysztof Krupinski, Elon Lindenstrauss, Dugald, Anand Pillay, Françoise Point, Laci Pyber, Tom Scanlon, Pierre Simon, Terry Tao, Alex Usvyatsov, and the two anonymous referees, for many useful comments.

REFERENCES

1. Ben Yaacov, Itai; Usvyatsov, Alexander, Continuous first order logic and local stability. *Trans. Amer. Math. Soc.* 362 (2010), no. 10, 5213–5259. MR2657678
2. Bergman, George M.; Lenstra, Hendrik W., Jr., Subgroups close to normal subgroups. *J. Algebra* 127 (1989), no. 1, 80–97. MR1029404 (91a:20004)
3. Breuillard, Emmanuel; Green, Ben; Terence, Tao, Approximate subgroups of linear groups, arXiv:1005.1881.
4. Breuillard, Emmanuel; Green, Ben, Approximate groups, II: the solvable linear case, arXiv:0907.0927.
5. Bourgain, J.; Katz, N.H.; Tao, T.C., A sum-product estimate in finite fields and applications, *Geom. Funct. Anal.* 14 (2004), no. 1, 27–57. MR2053599 (2005d:11028)
6. Chang, C. C.; Keisler, H. J., *Model theory*. Third edition. *Studies in Logic and the Foundations of Mathematics*, 73. North-Holland Publishing Co., Amsterdam, 1990. MR1059055 (91c:03026)
7. Chang, Mei-Chu, Product theorems in SL_2 and SL_3 . *J. Inst. Math. Jussieu* 7 (2008), no. 1, 1–25. MR2398145 (2009d:20112)
8. Cherlin, Gregory; Hrushovski, Ehud, *Finite structures with few types*. *Annals of Mathematics Studies*, 152. Princeton University Press, Princeton, NJ, 2003, vi+193 pp. MR1961194 (2004c:03037)
9. Curtis, Charles W.; Reiner, Irving, *Representation theory of finite groups and associative algebras*. Reprint of the 1962 original. AMS Chelsea Publishing, Providence, RI, 2006. xiv+689 pp. MR2215618 (2006m:16001)
10. Dinai, Oren, *Expansion properties of finite simple groups*. Ph.D. thesis. The Hebrew University of Jerusalem (2009). arXiv:1001.5069.
11. Elekes, György; Kirély, Zoltán, On the combinatorics of projective mappings. (English summary) *J. Algebraic Combin.* 14 (2001), no. 3, 183–197. MR1869409 (2003e:52034)
12. van den Dries, L., notes in <http://www.math.uiuc.edu/~vddries/>.
13. van den Dries, L.; Wilkie, A. J., Gromov’s theorem on groups of polynomial growth and elementary logic. *J. Algebra* 89 (1984), no. 2, 349–374. MR751150 (85k:20101)
14. Erdős, P.; Szemerédi, E., On sums and products of integers. *Studies in pure mathematics*, 213–218, Birkhäuser, Basel, 1983. MR820223 (86m:11011)
15. Gleason, A. M., The structure of locally compact groups. *Duke Math. J.* 18 (1951), 85–104. MR0039730 (12:589a)

16. Goldbring, Isaac, Hilbert's fifth problem for local groups. *Ann. of Math. (2)* 172 (2010), no. 2, 1269–1314. MR2680491 (2011f:22010)
17. Green, Ben; Sanders, Tom, A quantitative version of the idempotent theorem in harmonic analysis, *Annals of Mathematics*, 168 (2008), 1025–1054. MR2456890 (2010c:11013)
18. Gromov, Mikhael, Groups of polynomial growth and expanding maps. *Inst. Hautes Études Sci. Publ. Math. No. 53* (1981), 53–73. MR623534 (83b:53041)
19. Halmos, P., "Measure Theory," Van Nostrand, Princeton, NJ, 1950. MR0033869 (11:504d)
20. Helfgott, H. A., Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$. *Ann. of Math. (2)* 167 (2008), no. 2, 601–623. MR2415382 (2009i:20094)
21. Helfgott, H. A., Growth in $SL_3(\mathbb{Z}/p\mathbb{Z})$, arXiv:0807.2027.
22. Hewitt, E.; Savage, L.J., Symmetric measures on Cartesian products. *Trans. Amer. Math. Soc.*, 80 (1955), 470–501. MR0076206 (17:863g)
23. Hrushovski, Ehud, Pseudo-finite fields and related structures. *Model theory and applications*, 151–212, *Quad. Mat.*, 11, Aracne, Rome, 2002. MR2159717 (2006d:03059)
24. Hrushovski, Ehud; Wagner, Frank, Counting and dimensions. *Model theory with applications to algebra and analysis. Vol. 2*, 161–176, *London Math. Soc. Lecture Note Ser.*, 350, Cambridge Univ. Press, Cambridge, 2008. MR2436141 (2009k:03042)
25. Hrushovski, Ehud; Peterzil, Ya'acov; Pillay, Anand, Groups, measures, and the NIP. *J. Amer. Math. Soc.* 21 (2008), no. 2, 563–596 (section 7). MR2373360 (2008k:03078)
26. Hrushovski, Ehud; Pillay, Anand, On NIP and invariant measures, arXiv:0710.2330.
27. Hrushovski, Ehud; Pillay, Anand; Simon, Pierre, Generically stable and smooth measures in NIP theories, arXiv:1002.4763.
28. Kallenberg, Olav, On the representation theorem for exchangeable arrays. *J. Multivariate Anal.* 30 (1989), no. 1, 137–154. MR1003713 (90j:60034)
29. Jordan, Camille, Mémoire sur les équations différentielles linéaires intégrales algébriques, *Crelle* 84 (1878) pp. 89–215.
30. Kaplansky, Irving, Lie algebras and locally compact groups. Reprint of the 1974 edition. *Chicago Lectures in Mathematics*. University of Chicago Press, Chicago, IL, 1995. xii+148 pp. MR1324106 (96a:22008)
31. Kim, Byunghan; Pillay, Anand, Simple theories. *Joint AILA-KGS Model Theory Meeting (Florence, 1995)*. *Ann. Pure Appl. Logic* 88 (1997), no. 2-3, 149–164. MR1600895 (99b:03049)
32. Komlós, J.; Simonovits, M., Szemerédi's regularity lemma and its applications in graph theory. *Combinatorics, Paul Erdős is eighty, Vol. 2* (Keszthely, 1993), 295–352, *Bolyai Soc. Math. Stud.*, 2, János Bolyai Math. Soc., Budapest, 1996. MR1395865 (97d:05172)
33. Krauss, Peter H., Representation of Symmetric Probability Models, *J. Symbolic Logic* 34, No. 2 (June 1969), pp. 183–193. MR0275482 (43:1236)
34. Scott, D.; Krauss, P., Assigning probabilities to logical formulas, in: *Aspects of inductive logic*, ed. J. Hintikka and P. Suppes, North-Holland, Amsterdam, 1966, pp. 219–259. MR0201269 (34:1153)
35. Larsen, Michael J.; Pink, Richard, Finite subgroups of algebraic groups, 1998. Available from <http://www.math.ethz.ch/~pink/publications.html>.
36. Lee, John M., *Riemannian manifolds: An introduction to curvature*. Graduate Texts in Mathematics, 176. Springer-Verlag, New York, 1997. MR1468735 (98d:53001)
37. Marker, David, *Model theory. An introduction*. Graduate Texts in Mathematics, 217. Springer-Verlag, New York, 2002. MR1924282 (2003e:03060)
38. Marker, David, Semialgebraic Expansions of \mathbb{C} , *Transactions of the American Mathematical Society*, Vol. 320, No. 2 (Aug., 1990), pp. 581–592. MR964900 (90k:03034)
39. Morley, Michael, The Löwenheim-Skolem theorem for models with standard part. 1971 *Symposia Mathematica*, Vol. V (INDAM, Rome, 1969/70), Academic Press, London, pp. 43–52. MR0284322 (44:1551)
40. Nikolov, Nikolay; Pyber, László, Product decompositions of quasirandom groups and a Jordan type theorem, arXiv:math/0703343.
41. Pillay, Anand, <http://www.amsta.leeds.ac.uk/~pillay/>.
42. Pillay, Anand, *An introduction to stability theory*. Oxford Logic Guides, 8. The Clarendon Press, Oxford University Press, New York, 1983. xii+146 pp. MR719195 (85i:03104)
43. Poizat, Bruno, *Cours de théorie des modèles*. Bruno Poizat, Lyon, 1985. vi+584 pp. MR817208 (87f:03084)

44. Poizat, Bruno, An introduction to algebraically closed fields and varieties. in: The model theory of groups (Notre Dame, IN, 1985–1987), 41–67, Notre Dame Math. Lectures, 11, Univ. Notre Dame Press, Notre Dame, IN, 1989. MR985339
45. Raghunathan, M. S., Discrete subgroups of Lie groups. Math. Student 2007, Special Centenary Volume, 59–70 (2008). MR2527560 (2010d:22016)
46. Shalom, Yehuda; Tao, Terence, A finitary version of Gromov’s polynomial growth theorem, arXiv:0910.4148v2 [math.GR].
47. Shelah, Saharon, The lazy model-theoretician’s guide to stability. Comptes Rendus de la Semaine d’étude en Théorie des Modules (Inst. Math., Univ. Catholique Louvain, Louvain-la-Neuve, 1975). Logique et Analyse (N.S.) 18 (1975), no. 71-72, 241–308. MR0539969 (58:27447)
48. Shelah, S., Classification theory and the number of nonisomorphic models. Second edition. Studies in Logic and the Foundations of Mathematics, 92. North-Holland Publishing Co., Amsterdam, 1990. xxxiv+705 pp. MR1083551 (91k:03085)
49. Shelah, Saharon, Simple unstable theories. Ann. Math. Logic 19 (1980), no. 3, 177–203. MR595012 (82g:03055)
50. Tao, Terence, Product set estimates for non-commutative groups, arXiv math/0601431, Combinatorica 28 (2008), 547–594. MR2501249 (2010b:11017)
51. Tao, Terence, The sum-product phenomenon in arbitrary rings, arXiv:0806.2497, Contrib. Discrete Math. 4 (2009), 59–82. MR2592424 (2011a:11020)
52. Tao, Terence, <http://terrytao.wordpress.com/2007/03/02/open-question-noncommutative-freiman-theorem/>.
53. Tao, Terence, Freiman’s theorem for solvable groups, arXiv:0906.3535.
54. Tao, Terence; Vu, Van, Additive combinatorics. Cambridge Studies in Advanced Mathematics, 105. Cambridge University Press, Cambridge, 2006. xviii+512 pp. MR2289012 (2008a:11002)
55. Yamabe, Hidehiko, A generalization of a theorem of Gleason. Ann. of Math. (2) 58 (1953), 351–365. MR0058607 (15:398d)
56. Weil, André, Foundations of algebraic geometry. American Mathematical Society, Providence, R.I., 1962, xx+363 pp. MR0144898 (26:2439)
57. Zimmer, Robert J., Essential results of functional analysis. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 1990. x+157 pp. MR1045444 (91h:46002)

INSTITUTE OF MATHEMATICS, HEBREW UNIVERSITY AT JERUSALEM, GIV’AT RAM, 91904 JERUSALEM, ISRAEL

E-mail address: ehud@math.huji.ac.il