



Approximate subgroups with bounded VC-dimension

Gabriel Conant¹ · Anand Pillay²

Received: 17 July 2020 / Accepted: 14 November 2022 / Published online: 16 December 2022
© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2022

Abstract

We combine the fundamental results of Breuillard, Green, and Tao (Publ Math Inst Hautes Études Sci 116:115–221, 2012) on the structure of approximate groups, together with “tame” arithmetic regularity methods based on work of the authors and Terry (J Eur Math Soc (JEMS) 24(2):583–621, 2022), to give a structure theorem for finite subsets A of arbitrary groups G where A has “small tripling” and bounded VC-dimension: Roughly speaking, up to a small error, A will be a union of a bounded number of translates of a coset nilprogression of bounded rank and step (see Theorem 2.1). We also prove a stronger result in the setting of bounded exponent (see Theorem 2.2). Our results extend recent work of Martin-Pizarro, Palacín, and Wolf (Selecta Math (N.S.) 27(4):Paper No. 53, 19, 2021) on finite stable sets of small tripling.

Mathematics Subject Classification 11P70 · 11B30 · 03C45 · 03C20

1 Introduction

1.1 Small doubling in abelian groups

In additive combinatorics, an “inverse theorem” is a result in which one deduces structural information for a subset A of an abelian group by studying the behavior of its iterated sumsets $A + A + \cdots + A$. If A is finite then, by comparing the size of A to the sizes of its sumsets, one can approximate the extent to which A is “group-like” in

Partially supported by NSF grants: DMS-1855503 (Conant); DMS-1665035, DMS-1790212 (Pillay).

✉ Gabriel Conant
gconant@maths.cam.ac.uk

Anand Pillay
apillay@nd.edu

¹ Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Cambridge CB3 0WB, UK

² Department of Mathematics, University of Notre Dame, Notre Dame, IN 46556, USA

the sense of being closed under addition. The philosophy is that if A is approximately structured in this way, then it can be approximated by objects that are “perfectly structured” such as subgroups or arithmetic progressions [36].

A concrete example of this philosophy is the simple exercise that if G is an abelian group and $A \subseteq G$ is a finite set satisfying $|A + A| = |A|$, then A is either empty or a coset of a subgroup of G (see [39, Proposition 2.2]). More generally, we say that a (nonempty) finite set $A \subseteq G$ has k -**doubling** if $|A + A| \leq k|A|$. The next result provides an approximate description of finite sets with small doubling in abelian groups.

Theorem 1.1 (Green and Ruzsa [15]) *Suppose G is an abelian group and $A \subseteq G$ is a finite set with k -doubling. Then there is a proper coset progression $P \subseteq 2A - 2A$ of rank $O_k(1)$ such that A is covered by $O_k(1)$ translates of P .*

Coset progressions, and related objects, are defined in Sect. 3.2. The previous theorem was first proved by Freiman [12] for $(\mathbb{Z}, +)$, and later a different proof was found by Ruzsa, who extended the result to torsion-free abelian groups and groups of bounded exponent (see [26, 27]).

1.2 Small tripling in arbitrary groups

From a distance, Theorem 1.1 says that if A is a finite set of small doubling in an abelian group, then A can be covered by a bounded number of translates of a well-structured set P of bounded complexity contained a small sumset of A . We will refer to this kind of a result as a *Freiman–Ruzsa* property. Some sources, such as [30], follow the naming convention “Bogolyubov–Ruzsa” in reference to the notable influence of [2] on Ruzsa’s work. Our terminology is chosen to be consistent with [5, Section 2].

In abelian groups, an important feature of the small doubling assumption for a finite set is that it leads to controlled growth of all iterated sumsets via the *Plunnecke–Ruzsa inequalities* (see Remark 3.3). For nonabelian groups, this is no longer the case. Instead, one needs the stronger assumption of small *tripling*. Specifically, given an arbitrary group G , we say that a finite set $A \subseteq G$ has k -**tripling** if $|A^3| \leq k|A|$. By results of Tao [37], this property is closely related to the notion of a k -approximate group (see Definition 3.1). In particular, a k -approximate group has k -tripling by definition and, conversely, if A has k -tripling then $(A \cup A^{-1})^2$ is an $O(k^{O(1)})$ -approximate group by [37]. Combining this with the main structure theorem for approximate groups, due to Breuillard, Green and Tao [5], one obtains the following analogue of Theorem 1.1 for finite sets of small tripling in arbitrary groups.

Theorem 1.2 (Breuillard–Green–Tao [5], Tao [37]) *Suppose G is a group and $A \subseteq G$ is a finite set with k -tripling. Then there is a coset nilprogression $P \subseteq (A \cup A^{-1})^8$ of rank and step $O_k(1)$, and in $O_k(1)$ -normal form, such that A is covered by $O_k(1)$ left translates of P .*

In [6], the first author modified certain parts of [5] (specifically, a combinatorial argument of Sanders [29]) to show that in Theorem 1.2, one can in fact obtain $P \subseteq A^2 A^{-2} \cap A^{-2} A^2 \cap (A A^{-1})^2 \cap (A^{-1} A)^2$ (which then amounts to an exact nonabelian

analogue of Theorem 1.1). We also refer the reader to Breuillard’s surveys [3] and [4]. The first paper points out the origins of “approximate subgroup theory” in superstrong approximation (as well as in the inverse problems from Sect. 1.1).

1.3 Bounded exponent

It is also common in additive and multiplicative combinatorics to combine the small doubling or tripling condition with further restrictions, such as a bound on the exponent of the ambient group. In the abelian case, this reflects the frequent use of groups of the form \mathbb{F}_p^n as “toy models” for problems about general abelian groups. The next result provides an analogue of Theorem 1.2 for groups of bounded exponent.

Theorem 1.3 [5, 6, 16, 27] *Suppose G is a group of exponent r and $A \subseteq G$ is a finite set with k -tripling. Then there is a subgroup $H \subseteq A^2 A^{-2} \cap A^{-2} A^2 \cap (AA^{-1})^2 \cap (A^{-1}A)^2$ such that A is covered by $O_{k,r}(1)$ left cosets of H .*

In the abelian case, this theorem is due to Ruzsa [27]. The general case was proved by Hrushovski [16], and again later by Breuillard, Green, and Tao [5] under the stronger assumption that A is a k -approximate group. Via the result of Tao [37] mentioned above, this implies Theorem 1.3 for A with k -tripling, but with $H \subseteq (A \cup A^{-1})^8$, and the final improvement is done in [6]. Also, via an observation of van den Dries [42], it suffices to only assume that every element in the set $A^2 A^{-2} \cap A^{-2} A^2 \cap (AA^{-1})^2 \cap (A^{-1}A)^2$ has order at most r (see Proposition 7.8).

2 Main results

2.1 NIP and VC-dimension

In the present article, we consider a structural assumption on subsets of groups defined using VC-dimension. To see the relevance of this assumption in the context of finding “group-like” structure, we formulate VC-dimension for subsets of groups in terms of forbidden bipartite graphs. We use the notation $\Gamma = (V, W; E)$ for bipartite graphs, where V and W are vertex sets and $E \subseteq V \times W$ is an edge set. Given a group G and a set $A \subseteq G$, let $\Gamma_G(A)$ denote the bipartite graph $(G, G; E_A)$ where $E_A = \{(x, y) \in G^2 : x \in yA\}$.

A simple exercise, in a similar spirit as the one preceding Theorem 1.1, is that if G is an *arbitrary* group and $A \subseteq G$ is a nonempty set of *any cardinality*, then A is a coset of a subgroup of G if and only if $\Gamma_G(A)$ omits $([2], [2]; \leq)$ (by which we always mean as an *induced subgraph*).¹ Thus, in addition to the notion of k -doubling or tripling, we can use omitted subgraphs in $\Gamma_G(A)$ as a test for finding “group-like” structure in A . In order to work with a numerical parameter, we say that $A \subseteq G$ is d -NIP² if $\Gamma_G(A)$ omits the graph $([d], \mathcal{P}([d]); \in)$. A good exercise is that if $\Gamma_G(A)$ omits some finite bipartite graph $(V, W; E)$, then A is d -NIP for some $d \leq |V| + \lceil \log_2 |W| \rceil$. Therefore

¹ We use $[n]$ to denote $\{1, \dots, n\}$.

² This acronym, which comes from model theory, stands for the “negation of the independence property”.

we have successfully captured the phenomenon of forbidden bipartite subgraphs in a numerical way. This measurement is also mathematically useful since it provides access to tools from VC-theory. Indeed, it follows from the definitions that $A \subseteq G$ is d -NIP if and only if the collection of left translates of A has VC-dimension strictly less than d (as a set system on G). We will also see that NIP sets small tripling are even closer to approximate groups than what is given by [37]. Specifically, if $A \subseteq G$ is d -NIP with k -tripling, then $A \cup A^{-1} \cup \{1\}$ is an $O_d(k^{O(1)})$ -approximate group (see Theorem 3.20).

The following is our main result.

Theorem 2.1 (main result) *Suppose G is a group and $A \subseteq G$ is a finite d -NIP set with k -tripling. Given $\epsilon > 0$, there is a coset nilprogression $P \subseteq G$ of rank and step $O_{d,k,\epsilon}(1)$ in $O_{d,k,\epsilon}(1)$ -normal form, and a set $Z \subseteq AP$ with $|Z| < \epsilon|A|$, satisfying the following properties.*

- (i) $P \subseteq AA^{-1} \cap A^{-1}A$ and $A \subseteq CP$ for some $C \subseteq A$ with $|C| \leq O_{d,k,\epsilon}(1)$.
- (ii) There is a set $D \subseteq C$ such that $|(A \triangle DP) \setminus Z| < \epsilon|P|$.
- (iii) If $g \in G \setminus Z$ then $|gP \cap A| < \epsilon|P|$ or $|gP \cap A| > (1 - \epsilon)|P|$.

Moreover, if G is abelian then P is a proper coset progression.

Before continuing, let us compare Theorem 2.1 to the other results mentioned above. First, condition (i) mirrors the statement of Theorem 1.2, except that we have found the coset nilprogression P inside the smaller product set $AA^{-1} \cap A^{-1}A$. Beyond this, by introducing the ϵ parameter, we obtain an even stronger structural approximation of A in terms of P . In particular, condition (ii) says that A looks approximately like a union of boundedly many translates of P . Indeed, this condition implies the simpler expression

$$|A \triangle DP| < \epsilon(|P| + |A|).$$

Here one might find it desirable to have an error bound in terms of A only. Thus it is worth noting that, since $P \subseteq AA^{-1} \cap A^{-1}A$, we can use the Plunnecke–Ruzsa inequalities (Theorem 3.2(a)) to conclude that $|P| \leq k^{O(1)}|A|$. Finally, condition (iii) says that almost all translates of P behave regularly with respect to A in the strong sense of being almost disjoint from A or almost contained in A .

If we impose a bounded exponent assumption then, as in Sect. 1.3, we obtain a stronger version of Theorem 2.1 in which the coset nilprogression is replaced by a subgroup. Given a group G , we say that a subset $X \subseteq G$ has **exponent at most r** if every element of X has order at most r .

Theorem 2.2 *Suppose G is a group and $A \subseteq G$ is a finite d -NIP set with k -tripling. Assume $AA^{-1} \cap A^{-1}A$ has exponent at most r . Given $\epsilon > 0$, there is a subgroup $H \leq G$ and a set $Z \subseteq AH$, which is a union of left cosets of H with $|Z| < \epsilon|A|$, satisfying the following properties.*

- (i) $H \subseteq AA^{-1} \cap A^{-1}A$ and $A \subseteq CH$ for some $C \subseteq A$ with $|C| \leq O_{d,k,r,\epsilon}(1)$.
- (ii) There is a set $D \subseteq C$ such that $|(A \setminus Z) \triangle DH| < \epsilon|H|$.

(iii) If $g \in G \setminus Z$ then $|gH \cap A| < \epsilon|H|$ or $|gH \cap A| > (1 - \epsilon)|H|$.

Moreover, H is a finite Boolean combination³ of bi-translates of A .

We now discuss how the two theorems above fit into previous work on NIP sets in groups. The additional “structure and regularity” properties given by conditions (i) and (iii) of Theorems 2.1 and 2.2 are based on related results concerning “tame” arithmetic regularity for subsets of finite groups. The notion of arithmetic regularity was developed by Green [13] to provide a group-theoretic analogue of Szemerédi regularity for graphs. In [40], Terry and Wolf proved a strong form of arithmetic regularity for *d-stable* subsets of \mathbb{F}_p^n (see Sect. 7.6 for details on stability). These results were then qualitatively generalized to arbitrary finite groups by the authors and Terry in [9]. A quantitative generalization for finite abelian groups was done later by Terry and Wolf [41], and then finally for arbitrary finite groups by the first author [7].

Shortly after [40] and [9], Alon, Fox, and Zhao [1] proved a similar arithmetic regularity result for NIP sets in finite abelian groups of bounded exponent. This was soon followed by work of Sisask [35] on NIP sets in finite abelian groups and, independently, work of the authors and Terry [10] on NIP sets in arbitrary finite groups. The work in [10] used model theoretic tools developed by the authors in [8], which will also play a key role in the present paper.

For infinite groups, a “tame” Freiman–Ruzsa result for *stable* sets of small tripling was recently obtained by Martin-Pizarro, Palacín, and Wolf [20] (see Theorem 7.14 below). We also note that Sisask’s earlier work in [35] contains a strong Freiman–Ruzsa component. For example, [35, Theorem 5.4] provides a result along the lines of Theorem 2.2 in the case that $G = \mathbb{F}_p^n$ for some $n \geq 1$ and fixed prime p .

2.2 Sketch of the proofs

For the most part, our proofs will rely on model theoretic tools involving ultraproducts and pseudofinite sets. Note that Theorems 2.1 and 2.2 are asymptotic statements about finite subsets of groups. Thus if we assume that such a statement is false, then we obtain an infinite sequence of counterexamples, which can be traded for a single limiting structure via an ultraproduct construction. So that we can more easily apply model theoretic methods, we will pass to a sufficiently saturated elementary extension of this ultraproduct in an appropriate first-order language. In the end, we will have a saturated group G and a pseudofinite definable set $A \subseteq G$, along with the A -normalized pseudofinite counting measure μ_A on definable subsets of G . In this setting, “small tripling” manifests as the property that $\mu_A(A^3) < \infty$. By the Plunnecke–Ruzsa inequalities, this is *equivalent* to the property that any definable subset of G contained in $\langle A \rangle$ has finite μ_A -measure.

The above setting is very similar to the ones used by Hrushovski [16] and Breuillard, Green, and Tao [5] in their work on approximate groups. So, in the same setting as above, let us further assume for the moment that A is a (pseudofinite) approximate group which, in this setting, means that $1 \in A$, $A = A^{-1}$, and A^2 is covered by finitely

³ The “complexity” of this Boolean combination can also be bounded in terms of the parameters d, k, r , and ϵ only; see Remark 7.10.

many translates of A . Then there are two main steps in the work from [16] and [5]. The first step is to show that A^4 contains a type-definable subgroup Γ of G , which is normal and bounded index in the subgroup generated by A . (This is step called the “stabilizer theorem”, as it is related to the general theory of stabilizers of types in definable groups, and is in fact close to the case of groups definable in simple theories.) Consequently, if $X \subseteq G$ is any definable set containing Γ , then A can be covered by finitely many translates of X . Since $\Gamma \subseteq A^4$, we can choose $X \subseteq A^4$ and so, altogether, this yields a pseudofinite Freiman–Ruzsa property for the set X with respect to A . It remains to find such a set X exhibiting desirable algebraic properties, and this is the second main step. The key idea is to replace X by a definable set admitting some kind of structural nilpotence, e.g., via coset nilprogressions in [5] or a more technical “Bourgain chain condition” in [16]. This part of the proof crucially relies on the underlying topological structure of $\langle A \rangle / \Gamma$ as a locally compact Hausdorff group, which can be “modeled” by Lie groups via Gleason–Yamabe type results. Actually in [5] it is shown that that the connected component of the topological group $\langle A \rangle / \Gamma$ is residually nilpotent, which explains the role of coset nilprogressions.

We now discuss some of the new aspects of the proofs in this paper, pointing out that they are more than just a direct combination of [5] and [10]. We return to the setting in which A has small tripling and is NIP. We construct the subgroup Γ as above, in a different way (compared to [16]), bearing in mind that A is NIP and giving local (formula-by-formula) versions of stabilizer theorems for definable groups in NIP theories. In particular, we show that Γ is contained in the “smaller” product set $AA^{-1} \cap A^{-1}A$. This appears in Sect. 5.

Another key tool is “generic locally compact domination” in place of the “generic compact domination” from [10], which means in the model-theoretic sense working with ind-definable groups in place of definable groups, to show that definable sets approximating Γ behave regularly with respect to A . These results appear in Sect. 6 (see Corollary 6.14 in particular) and, together with results from [5] on pseudofinite coset nilprogressions (see Theorem 7.4), they are enough to prove Theorem 2.1. In order to prove Theorem 2.2, we use a standard Gleason–Yamabe argument, which says that if $AA^{-1} \cap A^{-1}A$ has finite exponent then Γ can be written as an intersection of definable subgroups of G .

2.3 The abelian case (and a note on bounds)

Note that we have omitted all mention of explicit bounds in the results discussed above. The situation is as follows. For results on abelian groups proved using additive combinatorics, one obtains explicit (and often quite efficient) bounds. We refer the reader to the relevant sources listed above for further details. On the other hand, with the exception of [7], all of the results mentioned above involving arbitrary groups rely on an ultraproduct construction in one way or another, and thus often lead to no explicit information about bounds. This partially includes Theorem 1.2, in the sense that the bound on the number of translates of P needed to cover A is not effective. However, if one relaxes $(A \cup A^{-1})^8$ to $(A \cup A^{-1})^{24}$, then one can bound the rank and step of P by $O(k^2 \log k)$ (see [5, Theorem 2.12]).

Given the above discussion, it is natural to expect that, for the case of abelian groups, it should be possible to prove Theorems 2.1 and 2.2 via finitary methods. As a step in this direction, we will show in Sect. 8 how to prove the abelian cases of these theorems directly from the relevant regularity results for finite groups, together with a result of Green and Ruzsa [15] that a finite set of small doubling in an abelian group can be Freiman-isomorphically modeled by a dense set in a finite abelian group. Combined with quantitative results of Alon, Fox, and Zhao [1], this yields a quantitative version of Freiman–Ruzsa for NIP sets of small doubling in abelian groups of bounded exponent (i.e., the abelian case of Theorem 2.2). In order to obtain explicit bounds for NIP sets of small doubling in arbitrary abelian groups, one would need an effective version of Theorem 8.8 below. We expect that such a result could be in reach via the tools developed by Sisask in [35].

3 Combinatorial preliminaries

3.1 Notation and terminology

Given positive integers m, n , and some set X of parameters, we use the notation $m \leq O_X(n)$ to mean that $m \leq cn$ where c is a positive constant depending only on X . If $X = \emptyset$ then c is an absolute constant and we write $m \leq O(n)$.

Let G be a group. We use concatenation for the group operation in G , and we let 1 denote the identity in G . Given sets $A, B \subseteq G$, we let $AB = \{ab : a \in A, b \in B\}$. This generalizes to products of any finite number of sets in the obvious way. Given $A \subseteq G$ and $g \in G$, we use gA (resp., Ag) for $\{g\}A$ (resp., $A\{g\}$), and we call this set a **left** (resp., **right**) **translate** of A .

Given $A \subseteq G$, we let $A^0 = \{1\}$ and, by induction, $A^{n+1} = AA^n$. Similarly, we let $A^{-1} = \{a^{-1} : a \in A\}$, and set $A^{-n} = (A^{-1})^n$. Given $n \geq 1$ and $A \subseteq G$, we let $A^{\pm n}$ denote $(A \cup A^{-1})^n \cup \{1\}$. Note that $\langle A \rangle = \bigcup_{n \geq 0} A^{\pm n}$.

If G is abelian then we will switch to additive notation involving $+$ for the group operation and 0 for the identity.

The following are several properties of subsets of groups that will be used throughout the paper.

Definition 3.1 Let G be a group and fix a nonempty set $A \subseteq G$.

- (1) A is **symmetric** if $1 \in A$ and $A = A^{-1}$ (equivalently, if $A = A^{\pm 1}$).
- (2) Given $k \geq 1$, we say A has **k -tripling** if A is finite and $|A^3| \leq k|A|$.
- (3) Given $k \geq 1$, we say A is a **k -approximate group** if A is symmetric and A^2 can be covered by k left translates of A .

The following result combines Lemma 3.4 and Theorem 3.9 of [37].

Theorem 3.2 (Tao [37]) *Let G be a group, and suppose $A \subseteq G$ has k -tripling.*

- (a) *If $n \geq 1$ and $\epsilon_1, \dots, \epsilon_n \in \{1, -1\}$, then $|A^{\epsilon_1} \dots A^{\epsilon_n}| \leq k^{O_n(1)}|A|$.*
- (b) *$A^{\pm 2}$ is an $O(k^{O(1)})$ -approximate group.*

Remark 3.3 The family of inequalities in Theorem 3.2(a) are often referred to as the *Plunnecke–Ruzsa inequalities* since, in the case of abelian groups, they were first proved by Plunnecke when $\epsilon_i = 1$ for all i , and then by Ruzsa for any ϵ_i . These proofs for abelian groups yield $O_n(1) = n$ (in Theorem 3.2(a)) and, moreover, only require k -doubling for the set A . For nonabelian groups, one may take $O_n(1) = 2n$ by [28, Theorem 5.1]. See also [24] for short proofs of these inequalities (with comparable bounds) by Petridis. Up to minor changes in the bounds, the same inequalities hold under the weaker assumption that A has k -doubling and $|AaA| \leq k|A|$ for all $a \in A$.

Altogether, for our purposes, one should view the Plunnecke–Ruzsa inequalities as the primary objective, and assumptions such as small doubling or tripling as means to that end.

3.2 Coset nilprogressions

In this section we recall the definition from [5] of a *coset nilprogression*. We start with the simpler notion of a *generalized progression*.

Definition 3.4 (*Generalized progression*) Let G be a group. Given real numbers $L_1, \dots, L_r > 0$, let $W(L_1, \dots, L_r)$ denote the collection of group words⁴ $w(x_1, \dots, x_r)$ in variables x_1, \dots, x_r such that, for all $1 \leq i \leq r$, the terms x_i and x_i^{-1} appear in w at most L_i times. Given group elements $u_1, \dots, u_r \in G$, define

$$P(u_1, \dots, u_r; L_1, \dots, L_r) = \{w(u_1, \dots, u_r) : w \in W(L_1, \dots, L_r)\}.$$

A **generalized progression** in G is a subset P of the form $P(u_1, \dots, u_r, L_1, \dots, L_r)$ for some choice of \bar{u} and \bar{L} . In this case, we call r the **rank** of P . By convention, $\{1\}$ is a generalized progression in G of rank 0.

The decision to allow L_1, \dots, L_r to be real numbers (rather than just integers) will be relevant in Proposition 3.14 below. In [5], generalized progressions are called *non-commutative progressions* to emphasize the setting of arbitrary groups. Indeed, for abelian groups the above definition simplifies to the more familiar notion of a *generalized arithmetic progression*.

Definition 3.5 (*Generalized arithmetic progression*) Let G be an abelian group. A **generalized arithmetic progression** in G is a subset P of the form

$$P(u_1, \dots, u_r; L_1, \dots, L_r) = \{n_1 u_1 + \dots + n_r u_r : |n_i| \leq L_i \text{ for all } 1 \leq i \leq r\}.$$

In this case, we call r the **rank** of P .

This definition of a generalized arithmetic progression differs slightly from other sources (such as [15] and [39]) where, for example, such progressions are allowed to have an affine term (see also Remark 3.9).

⁴ A *group word* is a term in the language of groups with a function symbol for inversion.

Freiman's Theorem in the integers [12, 26] says that finite sets in \mathbb{Z} with k -doubling can be approximated by generalized arithmetic progressions of rank $O_k(1)$. This fails for arbitrary abelian groups, where one may have nontrivial finite subgroups. In [15], Green and Ruzsa formulate the notion of a *coset progression* to address this issue.

Definition 3.6 (*Coset progression*) Let G be an abelian group. A **coset progression** in G is a set P of the form $P_0 + H$ where $P_0 \subseteq G$ is a generalized arithmetic progression and H is a finite subgroup of G . In this case, we say that P has **rank** r if P_0 has rank r .

A finite subgroup of an abelian group is a coset progression of rank 0. Moreover, a coset progression of rank r (in an abelian group) is a $(2^r + 1)$ -approximate group. On the other hand, it is not necessarily the case that a generalized progression in a nonabelian group is an approximate subgroup (see the discussion after [5, Remark 2.4] for precise details). Indeed, a main aspect of the results from [5] and [16] is that approximate groups exhibit a certain kind of nilpotent structure, which is made precise in [5] as follows.

Definition 3.7 (*Nilprogression*) Let G be a group. Given elements $u_1, \dots, u_r \in G$, set $c_1(\bar{u}) = \{u_1, \dots, u_r, u_1^{-1}, \dots, u_r^{-1}\}$ and, for $n > 1$, inductively define $c_n(\bar{u}) = \{[g, h] : g \in c_j(\bar{u}), h \in c_k(\bar{u}), j + k = n\}$.

A **nilprogression** in G is a generalized progression $P = P(u_1, \dots, u_r; L_1, \dots, L_r)$ such that $c_{s+1}(\bar{u}) = \{1\}$ for some $s \geq 0$. In this case, we say that P has **rank** r and **step** s .

For example, a generalized arithmetic progression in an abelian group is a nilprogression of step 1. We can now state the appropriate analogue of coset progressions for arbitrary groups.

Definition 3.8 (*Coset nilprogression*) Let G be a group. A **coset nilprogression** in G is a set P of the form $P_0 H$ where $P_0 \subseteq G$ is a nilprogression and H is a finite subgroup of G normalized by P_0 . In this case, we say that P has **rank** r and **step** s if P_0 has rank r and step s .

Note that a finite subgroup of an arbitrary group is a coset nilprogression of rank and step 0.

Remark 3.9 Any coset nilprogression in a group G is symmetric.

The rest of this section deals with tools for controlling the sizes and multiplicative behavior of nilprogressions. As before, we start with the abelian case.

Definition 3.10 (*properness*) Let G be an abelian group. A generalized arithmetic progression $P(u_1, \dots, u_r; L_1, \dots, L_r) \subseteq G$ is **proper** if the elements $n_1 u_1 + \dots + n_r u_r$ are pairwise distinct for varying choices of n_1, \dots, n_r . A coset progression $P = P_0 H \subseteq G$ is **proper** if P_0 is proper.

Note that a generalized arithmetic progression $P = P(u_1, \dots, u_r; L_1, \dots, L_r)$ in an abelian group is proper if and only if $|P| = \prod_{i=1}^r (2\lfloor L_i \rfloor + 1)$. In arbitrary

groups, this behavior is managed in an asymptotic fashion via the notion of “normal form”. Before stating this definition, we need to slightly shift the perspective on coset nilprogressions.

Remark 3.11 Suppose G is a group and $P = P_0H$ is a coset nilprogression in G , where P_0 is a nilprogression of rank r and step s and H is normalized by P_0 . Then H is a normal subgroup of $\langle P \rangle$ and, if $\pi: \langle P \rangle \rightarrow \langle P \rangle/H$ is the quotient map, then $\pi(P)$ is a nilprogression in $\langle P \rangle/H$ of rank r and step s . In particular, if $P_0 = P(u_1, \dots, u_r; L_1, \dots, L_r)$ then $\pi(P) = P(u_1H, \dots, u_rH; L_1, \dots, L_r)$.

Definition 3.12 (*c-normal form*) Let G be a group and fix an integer $c \geq 1$. A generalized progression $P = P(u_1, \dots, u_r; L_1, \dots, L_r)$ in G is in **c-normal form** if:

(i) For any $1 \leq i < j \leq r$,

$$[u_i, u_j], [u_i^{-1}, u_j], [u_i^{-1}, u_j], [u_i^{-1}, u_j^{-1}] \in P \left(u_{j+1}, \dots, u_r; \frac{cL_{j+1}}{L_iL_j}, \dots, \frac{cL_r}{L_iL_j} \right).$$

(ii) The elements $u_1^{n_1} \dots u_r^{n_r}$ are distinct for distinct choices of $n_i \leq c^{-1}|L_i|$.
 (iii) $c^{-1} \prod_{i=1}^r (2\lfloor L_i \rfloor + 1) \leq |P| \leq c \prod_{i=1}^r (2\lfloor L_i \rfloor + 1)$.

A coset nilprogression P is in **c-normal form** if the corresponding nilprogression $\pi(P)$ from Remark 3.11 is in *c-normal form*.

We will not need to delve into Definition 3.12 in any detail, except to note that a coset progression in an abelian group is proper if and only if it is in 1-normal form. Our main use of the normal form condition will be to control the sizes of “scaled” coset nilprogressions via Proposition 3.14 below.

Definition 3.13 Let G be a group, and suppose $P = P_0H$ is a coset nilprogression, where $P_0 = P(u_1, \dots, u_r; L_1, \dots, L_r)$. Given a real number $\epsilon > 0$, let $P^{(\epsilon)}$ denote the coset nilprogression $P(u_1, \dots, u_r; \epsilon L_1, \dots, \epsilon L_r)H$.

The next result follows from [5, Lemma C.1].

Proposition 3.14 Suppose G is a group and $P \subseteq G$ is a coset nilprogression of rank r in *c-normal form*. Then for any $\epsilon > 0$, one has $|P| \leq O_{r,c,\epsilon}(|P^{(\epsilon)}|)$.

3.3 Set systems and VC-dimension

Definition 3.15 Let X be a set and fix $\mathcal{S} \subseteq \mathcal{P}(X)$. Then \mathcal{S} **shatters** a subset $A \subseteq X$ if $\{A \cap S : S \in \mathcal{S}\} = \mathcal{P}(A)$. The **VC-dimension** of \mathcal{S} , denoted $\text{VC}(\mathcal{S})$, is the supremum over all $n \in \mathbb{N}$ such that \mathcal{S} shatters a subset of X of size n . Note that $\text{VC}(\mathcal{S})$ takes a value in $\mathbb{N} \cup \{\infty\}$.

The following is a corollary of the “VC Theorem for finite set systems”, proved by Vapnik and Chervonenkis in [43] (see also [32, Corollary 6.9]).

Theorem 3.16 [43] Let (X, μ) be a finite probability space, and suppose \mathcal{S} is a collection of measurable subsets of X with $\text{VC}(\mathcal{S}) \leq d < \infty$. Then, for any $0 < \epsilon < 1$, there is a sequence $(x_1, \dots, x_n) \in X^n$, with $n \leq O_{d,\epsilon}(1)$, such that for any $S \in \mathcal{S}$,

$$\left| \mu(S) - \frac{1}{n} |\{1 \leq i \leq n : x_i \in S\}| \right| \leq \epsilon.$$

Next we state Matoušek's “ (p, q) -theorem” (see [22, Theorem 4]). This result will not be necessary for the proof of our main results, and will only be used in Sect. 6.3.

Theorem 3.17 [22] Let X be a set and suppose $\mathcal{S} \subseteq \mathcal{P}(X)$ is finite with $\text{VC}(\mathcal{S}) \leq d < \infty$. Fix $p \geq q \geq 2^{d+1}$, and suppose that among any p sets in \mathcal{S} there are q sets with nontrivial intersection. Then there is some $F \subseteq X$ such that $|F| \leq O_{p,q}(1)$ and $F \cap S \neq \emptyset$ for all $S \in \mathcal{S}$.

Finally, we list some standard facts about various operations on set systems.

Fact 3.18 Let X be a set.

- (a) If $\mathcal{S} \subseteq \mathcal{P}(X)$ and $\mathcal{S}' = \{X \setminus S : S \in \mathcal{S}\}$, then $\text{VC}(\mathcal{S}) = \text{VC}(\mathcal{S}')$.
- (b) If $\mathcal{S}_1, \mathcal{S}_2 \subseteq \mathcal{P}(X)$ are such that $\text{VC}(\mathcal{S}_i) \leq d < \infty$ for $i \in \{1, 2\}$, and $\mathcal{S} = \{S_1 \cap S_2 : S_i \in \mathcal{S}_i\}$, then $\text{VC}(\mathcal{S}) < 10d$.
- (c) If $\mathcal{S} \subseteq \mathcal{P}(X)$ is such that $\text{VC}(\mathcal{S}) < \infty$ and $\mathcal{S}^* = \{S_x : x \in X\}$ where, given $x \in X$, $S_x = \{S \in \mathcal{S} : x \in S\}$, then $\text{VC}(\mathcal{S}^*) < 2^{\text{VC}(\mathcal{S})+1}$.

Proof Part (a) is easy and part (c) is a straightforward exercise (see, e.g., [32, Lemma 6.3]). Part (b) is evident from [11, Theorem 9.2.6], but we will sketch the proof to clarify the numerics. Given any $\mathcal{S} \subseteq \mathcal{P}(X)$, define the *shatter function* $\pi_{\mathcal{S}} : \mathbb{N} \rightarrow \mathbb{N}$ such that $\pi_{\mathcal{S}}(n) = \max\{|\{A \cap S : S \in \mathcal{S}\}| : A \subseteq X, |A| = n\}$. By the Sauer–Shelah Lemma (see, e.g., [32, Lemma 6.4]), if $\text{VC}(\mathcal{S}) \leq d$ then $\pi_{\mathcal{S}}(n) \leq (en/d)^d$ for any $n \geq d$. Now let $\mathcal{S}_1, \mathcal{S}_2$, and \mathcal{S} be as in part (b). One can check that $\pi_{\mathcal{S}}(n) \leq \pi_{\mathcal{S}_1}(n)\pi_{\mathcal{S}_2}(n)$ for all $n \geq 0$, and so $\pi_{\mathcal{S}}(n) \leq (en/d)^{2d}$ for all $n \geq d$. In particular, if $n \geq d$ and $(en/d)^{2d} < 2^n$, then $\pi_{\mathcal{S}}(n) < 2^n$, which implies $\text{VC}(\mathcal{S}) < n$. Setting $n = 10d$ satisfies these constraints. \square

Recall from Sect. 2.1 that a subset A of a group G is d -NIP if and only if the collection of left translates of A has VC-dimension strictly less than d . We say that $A \subseteq G$ is **NIP** if it is d -NIP for some $d \geq 1$. The previous fact has the following consequences for NIP subsets of groups.

Corollary 3.19 Let G be a group.

- (a) The collection of NIP subsets of G is a bi-invariant Boolean algebra.
- (b) If $A \subseteq G$ is d -NIP then the collection of right translates of A has VC-dimension at most 2^d .

Proof Part (a) is immediate from Fact 3.18(a, b) and the easy exercise that a translate of an NIP subset of G is NIP. For part (b), fix $A \subseteq G$ and define $\mathcal{S}_\ell = \{gA : g \in G\}$ and $\mathcal{S}_r = \{Ag : g \in G\}$. Then one can check that $\mathcal{S}_r = \mathcal{S}_\ell^*$ (using the notation from Fact 3.18(c)), which yields the result. \square

3.4 NIP sets of small tripling are essentially approximate groups

In this section, we prove a useful strengthening of Theorem 3.2(b) in the NIP setting.

Theorem 3.20 *Let G be a group, and suppose $A \subseteq G$ is d -NIP with k -tripling. Then $A^{\pm 2} \subseteq EA \cap AF$ for some $E, F \subseteq A^{\pm 3}$ of size $O_d(k^{O(1)})$. In particular, $A^{\pm 1}$ is an $O_d(k^{O(1)})$ -approximate group.*

Proof Set $X = A^{\pm 3}$, and note that X is finite. Let $\mathcal{S} = \{Ag : g \in A^{\pm 2}\}$, and note that \mathcal{S} is a collection of subsets of X . By Theorem 3.2(a), there is $\epsilon = k^{-O(1)}$ such that, for any $S \in \mathcal{S}$, we have $|S| = |A| \geq \epsilon|X|$. By Corollary 3.19(b), we have $\text{VC}(\mathcal{S}) \leq 2^d$. By Theorem 3.16, there is a set $E \subseteq X$ such that $|E| \leq O(2^d\epsilon^{-2} \log(2^d\epsilon^{-1})) \leq O_d(k^{O(1)})$ and $S \cap E \neq \emptyset$ for all $S \in \mathcal{S}$. Therefore, if $g \in A^{\pm 2}$ then $Ag^{-1} \cap E \neq \emptyset$, i.e., $g \in x^{-1}A$ for some $x \in E$. So $A^{\pm 2} \subseteq E^{-1}A$. Applying the same argument to $\mathcal{S}' = \{gA : g \in A^{\pm 2}\}$, which has VC-dimension at most $d - 1$, we obtain $F \subseteq X$ such that $|F| \leq O_d(k^{O(1)})$ and $A^{\pm 2} \subseteq AF^{-1}$. \square

Remark 3.21 In the previous proof, we only needed to bound the VC-dimension of translates of A by elements in $A^{\pm 2}$. Similar situations will arise frequently below, and so it is worth pointing out now that, for our purposes, the difference in the VC-dimensions of such set systems is negligible. Specifically, given a group G and $A \subseteq G$, let $\mathcal{S}_1 = \{gA : g \in G\}$ and $\mathcal{S}_2 = \{gA : g \in AA^{-1}\}$. Then clearly $\text{VC}(\mathcal{S}_2) \leq \text{VC}(\mathcal{S}_1)$. Conversely, we have $\text{VC}(\mathcal{S}_1) \leq \text{VC}(\mathcal{S}_2) + 1$ (this is observed by Sisask in [35]). Indeed, suppose \mathcal{S}_1 shatters a set $X \subseteq G$. After translating, we may assume $X \subseteq A$. So if $Y \subseteq X$ is nonempty, and $g \in G$ is such that $gA \cap X = Y$, then we necessarily have $g \in AA^{-1}$. Therefore \mathcal{S}_2 shatters any proper subset of X .

4 Model theoretic preliminaries

We assume familiarity with basic notions from first-order model theory, including structures, types, saturation, and ultraproducts.

4.1 Formulas and NIP

Let M be a first-order structure.

Definition 4.1 Let $\phi(\bar{x}; \bar{y})$ be a formula (possibly with parameters from M).

- (1) Given $\bar{b} \in M^{\bar{y}}$, let $\phi(M; \bar{b}) = \{\bar{a} \in M^{\bar{x}} : M \models \phi(\bar{a}; \bar{b})\}$.
- (2) A **instance** of $\phi(\bar{x}; \bar{y})$ is a formula of the form $\phi(\bar{x}, \bar{b})$ for some $\bar{b} \in M^{\bar{y}}$.
- (3) A **ϕ -formula** is a (finite) Boolean combination of instances of ϕ .
- (4) A subset of $M^{\bar{x}}$ is **ϕ -definable** if it is defined by a ϕ -formula.
- (5) Let $\phi^*(\bar{y}; \bar{x})$ denote the same formula $\phi(\bar{x}; \bar{y})$, but with the roles of object and parameter variables exchanged.

Given $B \subseteq M$, we say that a formula $\theta(\bar{x})$ is **over** B if B contains all parameters used in $\theta(\bar{x})$. A set $X \subseteq M^{\bar{x}}$ is **definable over** B if $X = \theta(M)$ for some formula $\theta(\bar{x})$

over B . If, moreover, $\theta(\bar{x})$ is a ϕ -formula for some formula $\phi(\bar{x}; \bar{y})$ over \emptyset , then we say that X is **ϕ -definable over B** .

Definition 4.2 A formula $\phi(\bar{x}; \bar{y})$ is **d -NIP (in M)** if there do not exist $(\bar{a}_i)_{i \in [d]}$ in $M^{\bar{x}}$ and $(\bar{b}_\sigma)_{\sigma \subseteq [d]}$ in $M^{\bar{y}}$ such that $M \models \phi(\bar{a}_i; \bar{b}_\sigma)$ if and only if $i \in \sigma$. We say that $\phi(\bar{x}; \bar{y})$ is **NIP** if it is d -NIP for some $d \geq 1$.

Remark 4.3 A formula $\phi(\bar{x}; \bar{y})$ is d -NIP if and only if $\{\phi(M; \bar{b}) : \bar{b} \in M^{\bar{y}}\}$ has VC-dimension strictly less than d (as a set system on $M^{\bar{x}}$). Note also that a subset A of a group G is d -NIP if and only if the “formula” $A(y \cdot x)$ is d -NIP.

As in Corollary 3.19, one can use Fact 3.18 to show that NIP formulas are closed under Boolean combinations, and that $\phi(\bar{x}; \bar{y})$ is NIP if and only if $\phi^*(\bar{y}; \bar{x})$ is NIP.

4.2 Locally compact type spaces

Definition 4.4 Given a set X , a **ring of subsets of X** is a nonempty collection $\mathcal{R} \subseteq \mathcal{P}(X)$ that is closed under finite unions, intersections, and relative complements (i.e., if $A, B \in \mathcal{R}$ then $A \setminus B \in \mathcal{R}$).

Now let M^* be a κ -saturated first-order structure for some sufficiently large cardinal κ . We say that a set I is **bounded** if $|I| < \kappa$. Given variables \bar{x} , let $\text{Def}_{\bar{x}}(M^*)$ denote the Boolean algebra of definable subsets of $(M^*)^{\bar{x}}$. Recall that a subset of $(M^*)^{\bar{x}}$ is said to be **type-definable** if it is an intersection of a bounded collection of definable sets. And a subset of $(M^*)^{\bar{x}}$ is said to be **ind-definable** (also called \vee -definable) if it is a union of a bounded number of definable sets. Note that if X is ind-definable as $\bigcup_{i \in I} X_i$, and Y is a definable subset of $(M^*)^{\bar{x}}$ which is a subset of X , then $Y \subseteq \bigcup_{i \in I_0} X_i$ for some finite $I_0 \subseteq I$ (by saturation of M^*).

We now give some refinements in the light of the above notion of a ring of subsets.

Definition 4.5 Fix \bar{x} and let \mathcal{R} be a subring of $\text{Def}_{\bar{x}}(M^*)$. Then a set $X \subseteq M^{\bar{x}}$ is **\mathcal{R} -type-definable** (resp., **countably \mathcal{R} -type-definable**) if it is an intersection of boundedly (resp., countably) many sets in \mathcal{R} .

A set $X \subseteq M^{\bar{x}}$ is **\mathcal{R} -ind-definable** (resp., **countably \mathcal{R} -ind-definable**) if it is a union of boundedly (resp., countably) many sets in \mathcal{R} .

In the previous definition, if \mathcal{R} is the Boolean algebra of ϕ -definable sets, for some formula $\phi(\bar{x}; \bar{y})$, then we replace \mathcal{R} with ϕ in the above terminology. If $\mathcal{R} = \text{Def}_{\bar{x}}(M^*)$, then we omit it. In these two cases, we will also want to be more precise with parameters.

Definition 4.6 Fix \bar{x} and $B \subseteq M^*$. Then $X \subseteq (M^*)^{\bar{x}}$ is **type-definable** (resp., **countably type-definable**) **over B** if it is an intersection of boundedly (resp., countably) many sets that are definable over B . These notions are relativized to a formula $\phi(\bar{x}; \bar{y})$ over \emptyset in the obvious way, and we also have the natural analogues with respect to ind-definability.

Definition 4.7 Fix \bar{x} and let \mathcal{R} be a subring of $\text{Def}_{\bar{x}}(M^*)$. A complete **\mathcal{R} -type** is a subset $p \subseteq \mathcal{R}$ such that:

- (i) $\emptyset \notin p$,
- (ii) $A \in p$ for some A ,
- (iii) if $A, B \in p$ then $A \cap B \in p$,
- (iv) if $A \in p$ and $B \in \mathcal{R}$, with $A \subseteq B$, then $B \in p$, and
- (v) for any $A, B \in \mathcal{R}$, if $A \in p$ and $B \notin p$ then $A \setminus B \in p$.

We let $S(\mathcal{R})$ denote the set of complete \mathcal{R} -types. (We may omit the adjective “complete”.)

Remark 4.8 $S(\mathcal{R})$ is a totally disconnected locally compact Hausdorff space, with a basis of compact-open sets given by $[X] = \{p \in S(\mathcal{R}) : X \in p\}$ for all $X \in \mathcal{R}$. Note that every $p \in S(\mathcal{R})$ is a partial type in the sense of the ambient structure M^* . On the other hand once one picks a definable set $X \in p$, then p is determined by its restriction to the Boolean algebra $\mathcal{R}^X = \{X \cap A : A \in \mathcal{R}\}$ of subsets of X . In any case both complete \mathcal{R} -types and \mathcal{R} -type definable sets will “concentrate” on some definable set $X \in \mathcal{R}$.

The previous remarks also connect the topology on $S(\mathcal{R})$ to more familiar spaces of types over Boolean algebras. Specifically, given $X \in \mathcal{R}$, the type space $S(\mathcal{R}^X)$ is a profinite space under the usual Stone topology. Given $X, Y \in \mathcal{R}$, if $\emptyset \neq X \subseteq Y$ then for any $p \in S(\mathcal{R}^X)$ there is a unique $q \in S(\mathcal{R}^Y)$ such that $p = \{X \cap A : A \in q\}$, and this induces a continuous injective map from $S(\mathcal{R}^X)$ to $S(\mathcal{R}^Y)$. These maps yield a directed system whose direct limit is naturally identified with $S(\mathcal{R})$. The final topology on $S(\mathcal{R})$ induced by this direct system coincides with the topology described above.

Given a type p and a type-definable set X , we write $p \models X$ to denote that p implies (or concentrates on) X , i.e., any realization of p (perhaps in a larger model) also realizes X when viewed as a partial type. If, moreover, $p \in S(\mathcal{R})$ for some ring \mathcal{R} , and X is \mathcal{R} -type-definable, then this means that p actually contains a small collection of \mathcal{R} -definable sets whose intersection is X .

Definition 4.9 Let X be a set and let \mathcal{R} be a ring of subsets of X . An **extended \mathcal{R} -Keisler measure** is a nontrivial finitely additive $\mathbb{R}_{\geq 0} \cup \{\infty\}$ -valued measure on \mathcal{R} . An **\mathcal{R} -Keisler measure** is an $\mathbb{R}_{\geq 0}$ -valued extended \mathcal{R} -Keisler measure.

Now let \mathcal{R} be a subring of $\text{Def}_{\bar{x}}(M^*)$ and let μ be an extended \mathcal{R} -Keisler measure. (In the situation most interesting for us, the measure will be $\mathbb{R}_{\geq 0}$ -valued.)

Definition 4.10 (1) An \mathcal{R} -type p is **μ -wide** if $\mu(X) > 0$ for any $X \in p$.
 (2) An \mathcal{R} -type-definable set $X \subseteq M$ is **μ -wide** if $\mu(D) > 0$ for any \mathcal{R} -definable set $D \subseteq (M^*)^{\bar{x}}$ containing X .

Fact 4.11 Suppose $X \subseteq (M^*)^{\bar{x}}$ is \mathcal{R} -type-definable and μ -wide. Then there is a μ -wide type $p \in S(\mathcal{R})$ such that $p \models X$.

Proof Let \mathcal{F} be the collection of \mathcal{R} -definable subsets of $(M^*)^{\bar{x}}$ containing X . Fix some $X \in \mathcal{F}$, and let $\mathcal{F}^X = \{X \cap Y : Y \in \mathcal{F}\} \subseteq \mathcal{F}$. By the usual standard exercise, there is $q \in S(\mathcal{R}^X)$ such that $\mathcal{F}^X \subseteq q$ and $\mu(Y) > 0$ for all $Y \in q$. Let $p \in S(\mathcal{R})$ be the unique lift of q . Then one easily checks that p is μ -wide and $\mathcal{F} \subseteq p$. \square

4.3 Locally compact quotients

Let G^* be a sufficiently saturated first-order expansion of a group. We work in G^* (namely take M^* from the previous section to be G^*). Suppose $\Sigma \leq G^*$ is an ind-definable subgroup of G^* , and $\Gamma \leq \Sigma$ is a type-definable bounded index subgroup of Σ . We use Σ/Γ for the set of *left* cosets of Γ in Σ . Let $\pi: \Sigma \rightarrow \Sigma/\Gamma$ be the canonical quotient map. Then Σ/Γ is a locally compact Hausdorff space under the **logic topology**, in which a subset $C \subseteq \Sigma/\Gamma$ is **closed** if and only if $\pi^{-1}(C) \cap X$ is type-definable for any definable set $X \subseteq \Sigma$. The verification of this claim is a standard exercise (see [17, Section 7] and [42, Section 4] for details in a slightly restricted setting). When Σ is actually *definable*, then Σ/Γ is compact, and this situation is covered in detail in [19] and [25], for example. The following are some other useful exercises.

Fact 4.12 Suppose $\Sigma \leq G^*$ is ind-definable, and $\Gamma \leq \Sigma$ is type-definable of bounded index.

- (a) A set $K \subseteq \Sigma/\Gamma$ is compact if and only if $\pi^{-1}(K)$ is type-definable.
- (b) A set $K \subseteq \Sigma/\Gamma$ is compact-open if and only if $\pi^{-1}(K)$ is definable.
- (c) If $X \subseteq \Sigma$ is type-definable then $\pi(X)$ is compact.
- (d) If $X \subseteq \Sigma$ is definable then $U = \{a\Gamma \in \Sigma/\Gamma : a\Gamma \subseteq X\}$ is open, and $\pi^{-1}(U) \subseteq X$.
- (e) If Σ is countably ind-definable and Γ is countably type-definable, then Σ/Γ is second countable.
- (f) If Γ is a normal subgroup of Σ , then Σ/Γ is a topological group.

We can recast the logic topology on Σ/Γ in terms of the topology on types over rings of definable sets. First, we set some notation.

Definition 4.13 Suppose \mathcal{R} is subring of $\text{Def}(G^*)$, Σ is an \mathcal{R} -ind-definable subgroup of G^* , Γ is an \mathcal{R} -type-definable bounded-index subgroup of Σ , and \mathcal{R} is left- Σ -invariant. Then every $p \in S(\mathcal{R})$ determines (i.e., implies) a unique left coset $a\Gamma$ of Γ in Σ . We let $\tau(p)$ denote this left coset $a\Gamma$, yielding a function $\tau: S(\mathcal{R}) \rightarrow \Sigma/\Gamma$.

The existence of τ above is immediate due to our convention that Γ (and every left coset $a\Gamma$) is \mathcal{R} -type definable (i.e., by a bounded collection of formulas compared to the saturation of G^*) and that p is a complete \mathcal{R} -type (over G^*). Namely, our saturation condition on G^* , and the bounded index assumption, implies that every coset of Γ in Σ (in a bigger model) has a representative in G^* , whereby if a realizes p in a bigger model, then a is in the same coset of Γ as some b in G^* and this information is part of the type p .

Proposition 4.14 Let \mathcal{R} be a subring of $\text{Def}(G^*)$. Suppose Σ is an \mathcal{R} -ind-definable subgroup of G^* , Γ is an \mathcal{R} -type-definable bounded-index subgroup of Σ , and \mathcal{R} is left- Σ -invariant. Then the logic topology on Σ/Γ is the finest topology for which $\tau: S(\mathcal{R}) \rightarrow \Sigma/\Gamma$ is continuous.

Proof We will only prove that τ is continuous (which is all we need in subsequent results), and leave the rest as an exercise.

Suppose that C is closed in the logic topology, and fix $p \notin \tau^{-1}(C)$. Then $\tau(p) \cap \pi^{-1}(C) = \emptyset$ and so, by saturation and since $\tau(p)$ is \mathcal{R} -type-definable, there is some $X \in \mathcal{R}$ such that $\tau(p) \subseteq X$ and $X \cap \pi^{-1}(C) = \emptyset$. So $[X]$ is an open set in $S(\mathcal{R})$, which contains p and is disjoint from $\tau^{-1}(C)$. \square

Our main results will require us to “localize” the topology on Σ/Γ to particular formulas. In particular, we will need the following result.

Theorem 4.15 *Let \mathcal{R} be a subring of $\text{Def}(G^*)$. Suppose Σ is an \mathcal{R} -ind-definable subgroup of G^* , Γ is an \mathcal{R} -type-definable bounded-index subgroup of Σ , and \mathcal{R} is left- Σ -invariant. Then a set $C \subseteq \Sigma/\Gamma$ is closed if and only if $X \cap \pi^{-1}(C)$ is \mathcal{R} -type-definable for any \mathcal{R} -definable set $X \subseteq \Sigma$.*

Proof Let $\pi : \Sigma \rightarrow \Sigma/\Gamma$ be the quotient map. The following claim can be proved via a straightforward generalization of [8, Lemma 4.1].

Claim If $X \subseteq \Sigma$ is type-definable then $\pi^{-1}(\pi(X))$ is \mathcal{R} -type-definable.

Now fix a definable set $X \subseteq \Sigma$, and set $K = \pi(X) \subseteq \Sigma/\Gamma$. Then K is compact by Fact 4.12. Define the \mathcal{R} -topology on K by declaring $C \subseteq K$ to be \mathcal{R} -closed if $\pi^{-1}(C)$ is \mathcal{R} -type-definable. We claim that this defines a compact Hausdorff topology on K . The verification of this nearly identical to the case of the logic topology (e.g., following [19]), and crucially relies on the claim above. Note that any \mathcal{R} -closed subset of K is closed, and so the induced logic topology on K refines the \mathcal{R} -topology. Since both topologies are compact and Hausdorff, they must coincide.

Now we prove the theorem. The left-to-right implication follows immediately from the fact that Σ is \mathcal{R} -ind-definable.

For the other direction, suppose $C \subseteq \Sigma/\Gamma$ is closed and fix an \mathcal{R} -definable set $X \subseteq \Sigma$. Then $C' = C \cap \pi(X)$ is a closed subset of $\pi(X)$, and hence is \mathcal{R} -closed by the above. So $\pi^{-1}(C')$ is \mathcal{R} -type-definable. Since $X \cap \pi^{-1}(C) = X \cap \pi^{-1}(C')$, it follows that $X \cap \pi^{-1}(C)$ is \mathcal{R} -type-definable. \square

Remark 4.16 Given the previous theorem, it is easy to verify the analogous localization of Fact 4.12 to a specific ring \mathcal{R} .

Definition 4.17 Suppose $\Sigma \leq G^*$ is ind-definable. A definable set $X \subseteq \Sigma$ is **left** (resp., **right**) **generic in** Σ if any definable subset of Σ is covered by finitely many left (resp., right) translates of X .

Our use of the word *generic* follows Definition 7.2 in [17]. In some sources, a subset X of an abstract group G is called (left) generic if G can be covered by finitely many (left) translates of X , and so the reader should make a note of this distinction.

Remark 4.18 Suppose $\Sigma \leq G^*$ is ind-definable, and $\Gamma \leq \Sigma$ is type-definable of bounded index. Then, by saturation and compactness, any definable subset of Σ containing Γ is left and right generic in Σ .

Remark 4.19 In the subsequent results, we will consider ind-definable subgroups of the form $\Sigma = \langle A \rangle$ for some definable set $A \subseteq G^*$. In this case, a definable set $X \subseteq \Sigma$ is left (resp., right) generic in Σ if and only if $A^{\pm m}$ is covered by finitely many left (resp., right) translates of X for all $m \geq 1$.

4.4 Pseudofinite sets and NIP formulas

We now review some terminology and facts about pseudofinite sets and pseudofinite counting measures. For simplicity, we work in the precise setting that will be used later to prove our main results. Specifically, let $(M_i)_{i \in I}$ be a collection of first-order structures in a common language, where I is some index set. Let \mathcal{U} be an ultrafilter over I and set $M = \prod_{\mathcal{U}} M_i$. A set $X \subseteq M$ is **internal** if $X = \prod_{\mathcal{U}} X_i$ for some sequence $(X_i)_{i \in I}$ with $X_i \subseteq M_i$ for all $i \in I$. Note that any definable subset of M is internal. An internal set $X \subseteq M$ is **pseudofinite** if $X = \prod_{\mathcal{U}} X_i$ where each X_i is finite.

Given a sequence $(r_i)_{i \in I}$ in the extended nonnegative real line $R := \mathbb{R}_{\geq 0} \cup \{\infty\}$, there is a unique $s \in R$ such that $\lim_{\mathcal{U}} r_i = s$, i.e., $\{i \in I : |r_i - s| < \epsilon\} \in \mathcal{U}$ for all $\epsilon > 0$ (here we set $\infty - \infty = 0$). If $X \subseteq M$ is nonempty, internal, and pseudofinite, then we define the **X -normalized pseudofinite counting measure**, denoted μ_X , such that, given an internal set $A = \prod_{\mathcal{U}} A_i$, $\mu_X(A) = \lim_{\mathcal{U}} |A_i|/|X_i|$. Then μ_X is an extended Keisler measure over the Boolean algebra of internal subsets of M . For the subsequent results, we will need to lift μ_X to elementary extensions of M . To do this, we expand M by a sort for $(R, +, <)$ and, for any formula $\phi(x; \bar{y})$ and any pseudofinite definable set $X \subseteq M$, a function f_{ϕ}^X from the \bar{y} sort in M to the R sort, which is interpreted as $\bar{b} \mapsto \mu_X(\phi(M; \bar{b}))$. From now on, we view M as a structure in this expanded language.⁵

Now let M^* be an elementary extension of M in this expanded language. Then in the R sort we have an elementary extension R^* of R , which can be equipped with the usual *standard part map*, denoted st , on elements in R^* of finite absolute value. We extend st to all of R^* by setting $\text{st}(x) = \infty$ for any infinite $x \in R^*$.

Given $X \subseteq M^*$, we let $X(M)$ denote $X \cap M$. A definable set $X \subseteq M^*$ is **pseudofinite** if X is definable over M and $X(M)$ is pseudofinite. Given a pseudofinite definable set $X \subseteq M^*$, we define the **X -normalized pseudofinite counting measure** μ_X on $\text{Def}(M^*)$ such that, given a formula $\phi(x; \bar{y})$ and $\bar{b} \in (M^*)^{\bar{y}}$, $\mu_X(\phi(M^*; \bar{b})) = \text{st}(f_{\phi}^{X(M)}(\bar{b}))$. Then μ_X is an extended $\text{Def}(M^*)$ -Keisler measure, and $\mu_X(X) = 1$. Moreover, $\mu_X(A) = \mu_X(A(M))$ for any $A \subseteq M^*$ definable over M (where $\mu_X(A(M))$ is computed in M). This is one of the reasons why we will take some care to keep track of definability of certain objects over M .

For the rest of the paper, when we say that a structure is sufficiently saturated, we will work in this expanded language for pseudofinite counting measures, and assume that M^* is an elementary extension of an ultraproduct M as above. When considering expansions of groups, we will use G^* and G instead of M^* and M .

The next proposition describes the definability and finite satisfiability properties of the measures μ_X with respect to NIP formulas.

Proposition 4.20 *Suppose $\phi(x; \bar{y})$ is an NIP formula, and $X \subseteq M^*$ is definable and pseudofinite. Let $Y = \{\bar{b} \in (M^*)^{\bar{y}} : \phi(M^*; \bar{b}) \subseteq X\}$, and fix $\epsilon > 0$.*

(a) *There is a finite set $F \subseteq X$ such that, if $\bar{b} \in Y$ and $\mu_X(\phi(x; \bar{b})) > \epsilon$, then $\phi(M^*; \bar{b}) \cap F \neq \emptyset$.*

⁵ To avoid ambiguity in words like “definable”, one can iterate the expansion by the symbols f_{ϕ}^X .

(b) *There is a set $D \subseteq (M^*)^{\bar{y}}$, which is countably ϕ^* -type-definable over X , such that, given $\bar{b} \in Y$, $\mu_X(\phi(x; \bar{b})) \leq \epsilon$ if and only if $\bar{b} \in D$.*

If, moreover, $\phi(x; \bar{y})$ is over M , then in part (a) we may assume $F \subseteq X(M)$, and in part (b) we may assume that D is countably ϕ^ -type-definable over $X(M)$.*

Proof This is a consequence of Theorem 3.16 and Łoś's Theorem. (See also [8, Section 2], which deals with the case that M^* is pseudofinite and $X = M^*$.) \square

Finally, we state a pseudofinite analogue of Theorem 3.17.

Proposition 4.21 *Let $\phi(x; \bar{y})$ be a d -NIP formula. Suppose $X \subseteq M^*$ is definable and pseudofinite, and $Y \subseteq (M^*)^{\bar{y}}$ is a definable set such that $\phi(M^*, \bar{b}) \subseteq X$ for all $\bar{b} \in Y$. Fix $p \geq q \geq 2^{d+1}$ and suppose that for any $\bar{b}_1, \dots, \bar{b}_p \in Y$ there is $I \subseteq [p]$ with $|I| = q$ such that $\bigcap_{i \in I} \phi(M^*; \bar{b}_i) \neq \emptyset$. Then there is a finite set $F \subseteq X$ such that, for any $\bar{b} \in Y$, $\phi(M^*; \bar{b}) \cap F \neq \emptyset$.*

Proof This can be seen by applying Theorem 3.17 to the finite sets X_i (of which X is the ultraproduct) and then again applying Łoś's Theorem. \square

5 The NIP stabilizer theorem

In this section we let G^* be a sufficiently saturated first-order structure expanding a group. So we work over a fixed ultraproduct $G = \prod_{\mathcal{U}} G_i \prec G^*$, where each G_i is an expansion of a group in a common language. The next definition provides analogues of some notions from Sect. 3.1.

Definition 5.1 Suppose $A \subseteq G^*$ is definable.

- (1) A has **finite tripling** if it is pseudofinite and $\mu_A(A^3) < \infty$.
- (2) A is an **approximate group** if A is symmetric and A^2 can be covered by finitely many left translates of A .

In our earlier definition of a subset A of a group G having k -tripling (Definition 3.1) there was an inbuilt assumption that A was finite. As we have now moved to a “pseudofinite setting”, this is reflected in the previous definition with the inbuilt assumption that A is pseudofinite (but possibly infinite). So the reader should be aware of this ambiguity, which can always be clarified by taking note of the surrounding context.

Note that a symmetric definable set $A \subseteq G^*$ is an approximate group if and only if A is left generic in $\langle A \rangle$. Recall also that if $A \subseteq G^*$ is definable and pseudofinite then $A(G) = \prod_{\mathcal{U}} A_i$, where each $A_i \subseteq G_i$ is finite. In this case, one can check that A has finite tripling if and only if there is some $k \geq 1$ such that the set of $i \in I$, for which A_i has k -tripling, is in \mathcal{U} . Similarly, A is an approximate group if and only if there is some $k \geq 1$ such that the set of $i \in I$, for which A_i is a k -approximate group, is in \mathcal{U} . Applying Theorem 3.2, we obtain the following conclusion.

Corollary 5.2 *Suppose $A \subseteq G^*$ is definable and pseudofinite with finite tripling. Then $\mu_A(A^{\pm n}) < \infty$ for all $n \geq 1$, and $A^{\pm 2}$ is an approximate group.*

Consequently, a pseudofinite definable set $A \subseteq G^*$ with finite tripling gives rise to a “locally pseudofinite” ind-definable subgroup $\langle A \rangle$. In particular, if $X \subseteq \langle A \rangle$ is definable then, by saturation, $X \subseteq A^{\pm n}$ for some $n \geq 1$, and thus $\mu_A(X) < \infty$ (which implies X is pseudofinite if X is definable over G).

In this setting, a simplified version of Hrushovski’s stabilizer theorem from [16] is as follows: If $A \subseteq G^*$ is definable and pseudofinite with finite tripling, then there is a type-definable bounded-index normal subgroup Γ of $\langle A \rangle$ such that $\Gamma \subseteq A^{\pm 8}$. Note that this statement can be viewed as a pseudofinite Freiman–Ruzsa property, in the same sense as discussed at the start of Sect. 1.2. Hrushovski’s approach is based on the general theory of stabilizers of types in definable groups, and “type-definability” of Γ is in the sense of a similarly expanded language for measures. In [5], Breuillard, Green, and Tao provide a different proof based on a combinatorial argument of Sanders [29]. This approach was later modified by Massicot and Wagner [21] in order to remove the necessity of expanding the language. In [6], the first author used similar modifications to obtain $\Gamma \subseteq A^2 A^{-2} \cap A^{-2} A^2 \cap (AA^{-1})^2 \cap (A^{-1}A)^2$.

The goal of this section is to prove a stronger version of the above statement under the assumption that A is NIP. In particular, we will exhibit stronger definability properties for Γ , and also that $\Gamma \subseteq AA^{-1} \cap A^{-1}A$.

Definition 5.3 Suppose $A \subseteq G^*$ is definable and pseudofinite. Given $\epsilon \geq 0$ and a definable set $X \subseteq G^*$, define

$$\begin{aligned} \text{Stab}_{A,\epsilon}^\ell(X) &= \{g \in G^* : \mu_A(gX \Delta X) \leq \epsilon\}, \text{ and} \\ \text{Stab}_{A,\epsilon}^r(X) &= \{g \in G^* : \mu_A(Xg \Delta X) \leq \epsilon\}. \end{aligned}$$

Set $\text{Stab}_A^\ell(X) = \text{Stab}_{A,0}^\ell(X)$ and $\text{Stab}_A^r(X) = \text{Stab}_{A,0}^r(X)$.

Proposition 5.4 Suppose $A, X \subseteq G^*$ are definable, and A is pseudofinite.

- (a) For any $\epsilon \geq 0$, $\text{Stab}_{A,\epsilon}^\ell(X)$ and $\text{Stab}_{A,\epsilon}^r(X)$ are symmetric subsets of G^* .
- (b) $\text{Stab}_A^\ell(X)$ and $\text{Stab}_A^r(X)$ are subgroups of G^* .
- (c) If $0 \leq \epsilon < \mu_A(X)$ then $\text{Stab}_{A,\epsilon}^\ell(X) \subseteq XX^{-1}$ and $\text{Stab}_{A,\epsilon}^r(X) \subseteq X^{-1}X$.

Proof Part (a) follows from bi-invariance of μ_A , and part (b) then follows from finite additivity. For part (c), note that if $\mu_A(gX \Delta X) < \mu_A(X)$ then $\mu_A(gX \cap X) > 0$, and so $gX \cap X \neq \emptyset$, i.e. $g \in XX^{-1}$. The other inclusion is similar. \square

Next we investigate the behavior of these stabilizer subgroups when the set A in question is NIP. First, we note a pseudofinite analogue of Theorem 3.20.

Corollary 5.5 Suppose $A \subseteq G^*$ is definable, NIP, and pseudofinite with finite tripling. Then for all $m \geq 2$ there are finite $E_m, F_m \subseteq A(G)^{\pm(m+1)}$ such that $A^{\pm m} \subseteq E_m A \cap A F_m$. In particular, $A^{\pm 1}$ is an approximate group.

Proof For $m = 2$, transfer Theorem 3.20 to the pseudofinite setting. Then proceed by induction. \square

Notation 5.6 We let \mathcal{R}_A , \mathcal{R}_A^ℓ , and \mathcal{R}_A^r denote the rings of subsets of $\langle A \rangle$ generated by $\{gAh : g, h \in \langle A \rangle\}$, $\{gA : g \in \langle A \rangle\}$, and $\{Ag : g \in \langle A \rangle\}$, respectively.

Given a pseudofinite definable set $A \subseteq G^*$, we say that $X \subseteq \langle A \rangle$ is **\mathcal{R}_A -definable over G** if $X \in \mathcal{R}_A$ and X is definable over G . Note that X is \mathcal{R}_A -definable over G if and only if it is in the ring generated by $\{gAh : g, h \in \langle A(G) \rangle\}$. We formulate analogous notions with \mathcal{R}_A replaced by \mathcal{R}_A^ℓ or \mathcal{R}_A^r in the obvious way. A consequence of the previous corollary is that if $A \subseteq G^*$ is definable, NIP, and pseudofinite with finite tripling, then $\langle A \rangle$ is countably \mathcal{R}_A^ℓ -ind-definable over G and countably \mathcal{R}_A^r -ind-definable over G .

Lemma 5.7 *Suppose $A \subseteq G^*$ is definable, NIP, and pseudofinite with finite tripling. Then $\text{Stab}_A^\ell(A)$ and Stab_A^r are bounded-index subgroups of $\langle A \rangle$. Moreover, $\text{Stab}_A^\ell(A)$ is countably \mathcal{R}_A^r -type-definable over G , and $\text{Stab}_A^r(A)$ is countably \mathcal{R}_A^ℓ -type-definable over G .*

Proof We just consider $\text{Stab}_A^\ell(A)$ (the proof for $\text{Stab}_A^r(A)$ is similar). Note first that $\text{Stab}_A^\ell(A) = \bigcap_{\epsilon > 0} \text{Stab}_{A,\epsilon}^\ell(A)$, and so it suffices to fix $0 < \epsilon < 1$, and show that $S := \text{Stab}_{A,\epsilon}^\ell(A)$ is countably \mathcal{R}_A^r -type-definable over G and left generic in $\langle A \rangle$.

For any $m \geq 1$, $A^{\pm m}$ is definable and pseudofinite. If we set $t_m = \mu_A(A^{\pm m})$ then $1 \leq t_m < \infty$ by Corollary 5.2, and $\mu_A = t_m \mu_{A^{\pm m}}$. Let $\theta(x; y)$ denote the formula $A(y \cdot x) \Delta A(x)$, and note that $\theta(x; y)$ is NIP. By Corollary 5.5, there is a finite set $F \subseteq A(G)^{\pm 3}$ such that $A^{\pm 2} \subseteq AF$. So $S \subseteq AF$ by Proposition 5.4(c) and since $\epsilon < 1$. Set $Y = \{b \in G^* : \theta(G^*; b) \subseteq A^{\pm 5}\}$. Then $S \subseteq A^{\pm 2} \subseteq AF \subseteq A^{\pm 4} \subseteq Y$, and so $S = \{g \in AF : \mu_{A^{\pm 5}}(\theta(x; g)) \leq \epsilon/t_5\}$. By Proposition 4.20(b), $S = D \cap AF$ for some set $D \subseteq G^*$, which is countably θ^* -type-definable over $A(G)^{\pm 5}$. Note that any set, which is θ^* -definable over $A(G)^{\pm 5}$, is a Boolean combination of right translates of A by elements in $A(G)^{\pm 5}$. So S is countably \mathcal{R}_A^r -type-definable over G .

Finally, we fix $m \geq 1$ and show that $X := A^{\pm m}$ can be covered by finitely many left translates of S . Let $\phi(x; y_1, y_2)$ denote the formula $A(y_1^{-1} \cdot x) \Delta A(y_2^{-1} \cdot x)$. Then $\phi(x; y_1, y_2)$ is NIP, and if $b_1, b_2 \in A^{\pm m}$ then $\phi(G^*; b_1, b_2) \subseteq X$. By Proposition 4.20(a), there is a finite set $F \subseteq X$ such that if $b_1, b_2 \in A^{\pm m}$ and $\mu_X(\phi(x; b_1, b_2)) > \epsilon/t_{m+1}$, then $\phi(G^*; b_1, b_2) \cap F \neq \emptyset$. Define an equivalence relation \sim on $A^{\pm m}$ such that $g \sim h$ if $F \cap gA = F \cap hA$. Since F is finite, we can choose representatives $g_1, \dots, g_n \in A^{\pm m}$ of all \sim -classes in $A^{\pm m}$. Now fix $b \in A^{\pm m}$. Then $b \sim g_i$ for some $i \leq n$, which implies $\phi(G^*; b, g_i) \cap F = \emptyset$. So $\mu_A(g_i^{-1}bA \Delta A) = \mu_A(bA \Delta g_iA) = t_{m+1}\mu_X(\phi(x; b, g_i)) \leq \epsilon$, and thus we have $g_i^{-1}b \in S$, i.e., $b \in g_iS$. Altogether, $A^{\pm m} \subseteq g_1S \cup \dots \cup g_nS$. \square

Definition 5.8 Given a pseudofinite definable set $A \subseteq G^*$, define

$$\Gamma_A^\ell = \bigcap_{g \in \langle A \rangle} \text{Stab}_A^\ell(gA) \quad \text{and} \quad \Gamma_A^r = \bigcap_{g \in \langle A \rangle} \text{Stab}_A^r(Ag).$$

We now consider stabilizers of types. Note that if $A \subseteq G^*$ is definable, then $\langle A \rangle$ acts by left and right translation on the space $S(\mathcal{R}_A)$ of (complete) \mathcal{R}_A -types.

Definition 5.9 Suppose $A \subseteq G^*$ is definable and pseudofinite. Given $p \in S(\mathcal{R}_A)$, define the stabilizers

$$\text{Stab}_A^\ell(p) = \{g \in \langle A \rangle : gp = p\} \quad \text{and} \quad \text{Stab}_A^r(p) = \{g \in \langle A \rangle : pg = p\}.$$

Lemma 5.10 Suppose $A \subseteq G^*$ is definable and pseudofinite. Fix $p \in S(\mathcal{R}_A)$ and $\star \in \{\ell, r\}$.

- (a) If p is μ_A -wide then $\Gamma_A^\star \leq \text{Stab}_A^\star(p)$.
- (b) $\text{Stab}_A^\star(p)$ is contained in any \mathcal{R}_A -type-definable bounded-index subgroup of $\langle A \rangle$.

Proof We consider the case when $\star = \ell$ (the other case is similar).

Part (a). Assume p is μ_A -wide and fix $x \in \Gamma_A^\ell$. For a contradiction, suppose $x \notin \text{Stab}_A^\ell(p)$. Since $x \in \langle A \rangle$ (by Proposition 5.4(c)), we have $xp \neq p$, and so there are $g, h \in \langle A \rangle$ such that $xgAh \Delta gAh \in p$. So $\mu_A(xgAh \Delta gAh) > 0$, i.e., $\mu_A(xgA \Delta gA) > 0$, and thus $x \notin \text{Stab}_A^\ell(gA)$, which contradicts $x \in \Gamma_A^\ell$.

Part (b). Let Γ be an \mathcal{R}_A -type-definable bounded-index subgroup of $\langle A \rangle$. Let C be the unique right coset of Γ such that $p \models C$. Then for any $x \in \text{Stab}_A^\ell(p)$, we have $p \models xC \cap C$, which implies $xC \cap C \neq \emptyset$, and thus $x \in \Gamma$. \square

We can now state and prove the main result of this section.

Theorem 5.11 Suppose $A \subseteq G^*$ is definable, NIP, and pseudofinite with finite tripling.

- (a) $\langle A \rangle$ has a smallest \mathcal{R}_A -type-definable subgroup of bounded-index, denoted Γ_A .
- (b) Γ_A is countably \mathcal{R}_A -type-definable over G and normal in $\langle A \rangle$.
- (c) $\Gamma_A = \Gamma_A^\ell = \Gamma_A^r$ and $\Gamma_A \subseteq AA^{-1} \cap A^{-1}A$.
- (d) If $p \in S(\mathcal{R}_A)$ is μ_A -wide, then $\Gamma_A = \text{Stab}_A^\ell(p) = \text{Stab}_A^r(p)$.

Proof We first show that Γ_A^ℓ is an \mathcal{R}_A -type-definable bounded-index normal subgroup of $\langle A \rangle$, which is contained in AA^{-1} . Note first that $\Gamma_A^\ell \subseteq AA^{-1}$ by Proposition 5.4(c). Let $H = \text{Stab}_A^\ell(A)$. If $g \in G^*$ then $\text{Stab}_A^\ell(gA) = gHg^{-1}$, and so $\Gamma_A^\ell = \bigcap_{g \in \langle A \rangle} gHg^{-1}$. So it follows from Lemma 5.7 that Γ_A^ℓ is an \mathcal{R}_A -type-definable bounded-index normal subgroup of $\langle A \rangle$.

By a similar argument, Γ_A^r is an \mathcal{R}_A -type-definable bounded-index normal subgroup of $\langle A \rangle$, which is contained in $A^{-1}A$.

Now let $p \in S(\mathcal{R}_A)$ be μ_A -wide (note that such types exist by Fact 4.11). By Lemma 5.10, we have

$$\Gamma_A^\ell \leq \text{Stab}_A^\ell(p) \leq \Gamma_A^r \leq \text{Stab}_A^r(p) \leq \Gamma_A^\ell,$$

and so Γ_A^ℓ , Γ_A^r , $\text{Stab}_A^\ell(p)$, and $\text{Stab}_A^r(p)$ are all the same group. It also follows from Lemma 5.10 that this group is the smallest \mathcal{R}_A -type-definable bounded-index subgroup of $\langle A \rangle$. Thus we have the group Γ_A described in part (a). It remains to show that Γ_A is countably \mathcal{R}_A -type-definable over G . For this, we first note that, by part (a), Γ_A is automorphism invariant with respect to the group language expanded by a formula defining A . It follows that Γ_A is countably type-definable over \emptyset . A standard saturation argument then shows that Γ_A is countably \mathcal{R}_A -type-definable over G (indeed, over any model $M \prec G^*$). \square

One could also denote Γ_A by $\langle A \rangle_{\mathcal{R}_A}^{00}$ (the type-definable connected component of $\langle A \rangle$ with respect to the ring \mathcal{R}_A of formulas, or definable sets).

Remark 5.12 It follows from the proof of the previous theorem that if $A \subseteq G^*$ is definable, NIP, and pseudofinite with finite tripling, then Γ_A coincides with $\bigcap_{g \in I} g \text{Stab}_A^\ell(A)g^{-1}$ and $\bigcap_{g \in I} g \text{Stab}_A^r(A)g^{-1}$ for some countable $I \subseteq \langle A(G) \rangle$.

If $\langle A \rangle$ is abelian, then we can further conclude $\Gamma_A = \text{Stab}_A^\ell(A) = \text{Stab}_A^r(A)$. But this can fail in general. For example, suppose G^* is pseudofinite and H is a non-normal finite-index definable subgroup of G^* . Let $A = aH$ where $a \in G^*$ is such that $aHa^{-1} \neq H$. Then A is NIP and has finite tripling, but $\text{Stab}_A^\ell(A) = aHa^{-1}$ and $\text{Stab}_A^r(A) = H$.

6 Measures and domination

The goal of this section is to prove Theorem 6.12. Roughly speaking, this result says that if G^* is a sufficiently saturated group and $A \subseteq G^*$ is definable, NIP, and pseudofinite with finite tripling, then the collection of cosets C of Γ_A in $\langle A \rangle$, such that both $C \cap A$ and $C \setminus A$ are μ_A -wide, is Haar null as a subset of the locally compact group $\langle A \rangle / \Gamma_A$. This will be the origin for condition (iii) of our main result (Theorem 2.1).

In the special case that $\langle A \rangle$ is *definable* (which, in particular, implies that $\langle A \rangle$ is pseudofinite and $\langle A \rangle / \Gamma_A$ is compact), the analogue of Theorem 6.12 was proved by the authors in [8]. The proof in this case will follow a similar strategy, and most of the work will involve adapting various tools and techniques to the setting of locally compact groups.

6.1 Measures

Let $G^* \succ G = \prod_{\mathcal{U}} G_i$ be a sufficiently saturated group. Note that if $A \subseteq G^*$ is definable and pseudofinite with finite tripling, then the A -normalized pseudofinite counting measure μ_A is a bi-invariant Keisler measure on $\text{Def}(\langle A \rangle)$. In particular, the finite tripling assumption ensures that μ_A is $\mathbb{R}_{\geq 0}$ -valued on $\text{Def}(\langle A \rangle)$.

The first main result of this section (Lemma 6.2 below) will require the following standard Borel–Cantelli-type result about probability measures.

Fact 6.1 *For any $\epsilon > 0$ and $q \geq 1$, there are $\delta > 0$ and $p \geq 1$ such that the following holds. Let B be a Boolean algebra and ν be a finitely additive probability measure on B . Suppose $x_1, \dots, x_p \in B$ are such that $\nu(x_i) \geq \epsilon$ for all $i \in [p]$. Then there is a q -element set $I \subseteq [p]$ such that $\nu(\bigwedge_{i \in I} x_i) \geq \delta$.*

For the rest of this section, we fix a definable, NIP, pseudofinite set $A \subseteq G^*$ with finite tripling.

Lemma 6.2 *Given $X \in \mathcal{R}_A$, the following are equivalent:*

- (i) X is left generic in $\langle A \rangle$,
- (ii) X is right generic in $\langle A \rangle$,
- (iii) $\mu_A(X) > 0$.

So a type $p \in S(\mathcal{R}_A)$ is μ_A -wide if and only if every $X \in p$ is (left) generic in $\langle A \rangle$.

Proof Note that (i) \Rightarrow (iii) and (ii) \Rightarrow (iii) follow immediately from finite additivity and invariance of μ_A .

(iii) \Rightarrow (i). The argument is similar to Theorem 3.20. Suppose $\mu_A(X) > 0$. Fix $m \geq 1$. We want to find a finite set $F \subseteq \langle A \rangle$ such that $A^{\pm m} \subseteq FX$. Without loss of generality, we can assume $X \subseteq A^{\pm m}$. Since A has finite tripling, we have $\mu_{AA^{\pm m}}(A) > 0$ and so, for any $b \in G^*$,

$$\mu_{AA^{\pm m}}(Xb) = \mu_{AA^{\pm m}}(X) = \mu_A(X)\mu_{AA^{\pm m}}(A) > 0.$$

Let $\phi(x; y)$ be the formula $X(x \cdot y)$. Then $\phi(x; y)$ is NIP and, if $b \in A^{\pm m}$, then $\phi(G^*, b) \subseteq AA^{\pm m}$. By Proposition 4.20(a), there is a finite set $F \subseteq AA^{\pm m}$ such that if $b \in A^{\pm m}$ then $Xb^{-1} \cap F \neq \emptyset$. So $A^{\pm m} \subseteq F^{-1}X$.

(iii) \Rightarrow (ii). One can use a similar strategy as in (iii) \Rightarrow (i), except applied to the formula $X(y \cdot x)$. \square

We call a set $X \in \mathcal{R}_A$ **generic** if it satisfies the equivalent conditions of the previous corollary. We will also refer to μ_A -wide types in $S(\mathcal{R}_A)$ as **generic**, and let $S^g(\mathcal{R}_A)$ denote the set of generic \mathcal{R}_A -types.

Remark 6.3 The set $S^g(\mathcal{R}_A)$ is (topologically) closed in $S(\mathcal{R}_A)$, since any non-generic \mathcal{R}_A -type contains a non-generic set in \mathcal{R}_A . In fact, by modifying the usual argument for Boolean algebras (e.g., [23, Lemma 1.7]), one can show that if $p \in S^g(\mathcal{R}_A)$ then $S^g(\mathcal{R}_A)$ is the closure of the orbit $\{gp : g \in \langle A \rangle\}$.

Let $\mathbb{G}_A = \langle A \rangle / \Gamma_A$, and let $\pi : \langle A \rangle \rightarrow \mathbb{G}_A$ be the quotient map. Recall that \mathbb{G}_A is a second countable locally compact Hausdorff group under the logic topology defined in Sect. 4.3.

Definition 6.4 Given $X \in \mathcal{R}_A$ and $p \in S^g(\mathcal{R}_A)$, define

$$\mathbb{B}_X^p = \{a\Gamma_A \in \mathbb{G}_A : X \in ap\}.$$

Note that \mathbb{B}_X^p is well-defined since $\Gamma_A = \text{Stab}_A^\ell(p)$ by Theorem 5.11(d).

The next lemma shows that the set \mathbb{B}_X^p defined above is Borel (of low complexity). This relies on the “local NIP” version of Borel definability for invariant types, proved by Simon in [33].

Lemma 6.5 If $X \in \mathcal{R}_A$ and $p \in S^g(\mathcal{R}_A)$, then \mathbb{B}_X^p is F_σ and G_δ in \mathbb{G}_A .

Proof We first note that it suffices to show that \mathbb{B}_X^p is F_σ for all $X \in \mathcal{R}_A$ and $p \in S^g(\mathcal{R}_A)$. Indeed, suppose this holds. By Corollary 5.5, we can write $\langle A \rangle = \bigcup_{m=1}^{\infty} Y_m$, where $Y_m \in \mathcal{R}_A$. Now, given $X \in \mathcal{R}_A$ and $p \in S^g(\mathcal{R}_A)$, we have $\mathbb{G}_A \setminus \mathbb{B}_X^p = \bigcup_{m=1}^{\infty} \mathbb{B}_{Y_m \setminus X}^p$, which is F_σ , and so \mathbb{B}_X^p is G_δ .

Now fix $X \in \mathcal{R}_A$ and $p \in S^g(\mathcal{R}_A)$. Let $D \in p$ be arbitrary. If $a \in \pi^{-1}(\mathbb{B}_X^p)$ then $p \models a^{-1}X \cap D$, and so we have $\pi^{-1}(\mathbb{B}_X^p) \subseteq Y := XD^{-1}$. Let $Z = Y \cup Y^{-1}X$, and note that Z is definable. By Proposition 4.20(a), there is a countable set $E \subset G^*$ such that

for any $\epsilon > 0$, if $a_1, a_2 \in Y$ and $\mu_A(a_1^{-1}X \setminus a_2^{-1}X) > 0$, then $(a_1^{-1}X \setminus a_2^{-1}X) \cap E \neq \emptyset$. Without loss of generality, we may assume that $E \subseteq Y^{-1}X \subseteq Z$.

We now define a new first-order structure $\mathcal{Z} = (Z; U_Y, R_X)$ with universe Z , a unary relation U_Y interpreted as Y , and a binary relation $R_X(x, y)$ interpreted as $X(x \cdot y)$. Note that \mathcal{Z} is interpretable in G^* . Working for now with respect to $\text{Th}(\mathcal{Z})$, let $\phi(x; y)$ denote the NIP formula $R_X(y, x) \wedge U_Y(y)$. Then there is a complete ϕ -type $p_0 \in S_\phi(Z)$ such that, given $a \in Z$, $\phi(x; a) \in p_0$ if and only if $a^{-1}X \in p$. Let $M \prec \mathcal{Z}$ be a countable submodel such that $E \subseteq M$. We claim that p_0 is M -invariant. For a contradiction, suppose we have $a_1, a_2 \in Z$ such that $a_1 \equiv_M a_2$ and $\phi(x; a_1), \neg\phi(x; a_2) \in p_0$. Then $a_1 \in Y$, which implies $a_2 \in Y$, and thus $a_1^{-1}X \setminus a_2^{-1}X \in p$. Since p is generic, there is some $m \in E \subseteq M$ such that $m \in a_1^{-1}X \setminus a_2^{-1}X$. So $\mathcal{Z} \models R_X(a_1, m) \wedge \neg R_X(a_2, m)$, which contradicts $a_1 \equiv_M a_2$.

Let $S = \{q \in S_y(M) : \phi(x; a) \in p_0 \text{ for some/any } a \models q\}$. By [33, Theorem 0.1], S is an F_σ set in $S_y(M)$. So if $W = \{a \in Z : a \models q \text{ for some } q \in S\}$, then we can write $W = \bigcup_{n=0}^{\infty} W_n$, where $W_n \subseteq Z$ is type-definable (over M). It is also straightforward to show that $W = \pi^{-1}(\mathbb{B}_X^p)$ (recall that $\pi^{-1}(\mathbb{B}_X^p) \subseteq Y$ by assumption).

Now we return to the structure G^* . Note that each set W_n above is type-definable (with respect to G^*) since X , Y , and Z are all definable. So each $\pi(W_n) \subseteq \mathbb{G}_A$ is closed by Fact 4.12(c), and thus $\mathbb{B}_X^p = \bigcup_{n=0}^{\infty} \pi(W_n)$ is F_σ . \square

Definition 6.6 Let η be a Haar measure on \mathbb{G}_A . Given $p \in S^g(\mathcal{R}_A)$ and $X \in \mathcal{R}_A$, define $\eta_p(X) = \eta(\mathbb{B}_X^p)$. (This is well-defined by Lemma 6.5.)

Proposition 6.7 Let η be a Haar measure on \mathbb{G}_A and suppose $p \in S^g(\mathcal{R}_A)$. Then η_p is a left- $\langle A \rangle$ -invariant \mathcal{R}_A -Keisler measure.

Proof We clearly have $\eta_p(\emptyset) = 0$. If $X, Y \in \mathcal{R}_A$ then $\mathbb{B}_{X \cup Y}^p = \mathbb{B}_X^p \cup \mathbb{B}_Y^p$ and $\mathbb{B}_{X \cap Y}^p = \mathbb{B}_X^p \cap \mathbb{B}_Y^p$, and so finite additivity of η_p is inherited from additivity of η . Also, if $X \in \mathcal{R}_A$ and $g \in \langle A \rangle$ then $\mathbb{B}_g X = g\Gamma \cdot \mathbb{B}_X^p$, and so $\eta_p(gX) = \eta_p(X)$ by left-invariance of η .

Next, we show that η_p is $\mathbb{R}_{\geq 0}$ -valued. First, fix $X \in \mathcal{R}_A$, and let $Y \in p$ be arbitrary. Then (as in the proof of Lemma 6.5) we have $\mathbb{B}_X^p \subseteq \pi(XY^{-1})$, which is compact Fact 4.12(c). Therefore $\eta_p(X) \leq \eta(\pi(XY^{-1})) < \infty$.

Finally, we show that η_p is nontrivial. Fix $X \in \mathcal{R}_A$ such that $\Gamma_A \subseteq X$. Let $U = \{a\Gamma_A \in \mathbb{G}_A : a\Gamma_A \subseteq X\}$. Then U is nonempty by assumption, and open by Fact 4.12(d). So $\eta(U) > 0$. Moreover, if $a\Gamma_A \in U$ then, since $ap \models a\Gamma_A$, it follows that $X \in ap$, i.e., $a\Gamma_A \in \mathbb{B}_X^p$. So $U \subseteq \mathbb{B}_X^p$, and thus we have $\eta_p(X) > 0$. \square

6.2 Generic locally compact domination

Before returning to the setting of the previous section, we first state a result of Simon [34], which is central to the proof of Theorem 6.12.

Definition 6.8 Let \mathbb{G} be a locally compact group and let η be a Haar measure on \mathbb{G} . A Borel set $W \subseteq \mathbb{G}$ is **pointwise large** if $\eta(W \cap U) > 0$ for any open set $U \subseteq \mathbb{G}$ such that $W \cap U \neq \emptyset$.

Note that whether a Borel set is pointwise large is independent of the choice of Haar measure. The next result is Theorem 3.7 of [34].

Theorem 6.9 (Simon [34]) Let \mathbb{G} be a locally compact second countable group, and suppose $B \subseteq \mathbb{G}$ is an NIP set such that both B and $\mathbb{G} \setminus B$ are F_σ and pointwise large. Then ∂B is Haar null.

We now continue to fix a sufficiently saturated group G^* and an NIP pseudofinite definable set $A \subseteq G^*$ with small tripling. Let \mathbb{G}_A and $\pi: \langle A \rangle \rightarrow \mathbb{G}_A$ be as above. Recall that we also have a well-defined continuous surjective map $\tau: S(\mathcal{R}_A) \rightarrow \mathbb{G}_A$ such that $p \models \tau(p)$ for any $p \in S(\mathcal{R}_A)$.

Definition 6.10 Given $X \in \mathcal{R}_A$, define \mathbb{E}_X to be the set of cosets $a\Gamma_A \in \mathbb{G}_A$ such that both $\Gamma_A \cap X$ and $\Gamma_A \setminus X$ are μ_A -wide.

Note that, by Fact 4.11, a coset $a\Gamma_A$ is in \mathbb{E}_X if and only if there are $q, q' \in S^g(\mathcal{R}_A)$ such that $q \models a\Gamma_A \cap X$ and $q' \models a\Gamma_A \setminus X$.

Lemma 6.11 Suppose $X \in \mathcal{R}_A$ and $p \in S^g(\mathcal{R}_A)$ is such that $p \models \Gamma_A$.

- (a) \mathbb{B}_{gX}^p is NIP in \mathbb{G}_A .
- (b) \mathbb{B}_X^p and $\mathbb{G}_A \setminus \mathbb{B}_X^p$ are pointwise large.
- (c) \mathbb{E}_X is compact.
- (d) $\mathbb{E}_X \subseteq \partial \mathbb{B}_X^p$.

Proof Part (a). Fix $X \in \mathcal{R}_A$ and $p \in S^g(\mathcal{R}_A)$. We want to bound the VC-dimension of $\mathcal{S} = \{\mathbb{B}_{gX}^p : g \in \langle A \rangle\}$. So suppose \mathcal{S} shatters $\{a_1\Gamma_A, \dots, a_n\Gamma_A\} \subseteq \mathbb{G}_A$. Then for any $I \subseteq [n]$ we can choose $g_I \in \langle A \rangle$ such that $a_i\Gamma_A \in \mathbb{B}_{g_I X}^p$ if and only if $i \in I$. In other words, for $i \in [n]$, we have $a_i^{-1}g_I X \in p$ if and only if $i \in I$. It follows that there is some $b \in \langle A \rangle$ such that, for $i \in [n]$, $b \in a_i^{-1}g_I X$ if and only if $i \in I$. So $\{gX : g \in \langle A \rangle\}$ shatters $\{a_1b, \dots, a_nb\}$. Altogether, $\text{VC}(\mathcal{S}) \leq \text{VC}(\{gX : g \in \langle A \rangle\})$, and so $\text{VC}(\mathcal{S}) < \infty$ since X is NIP (by Corollary 3.19).

Part (b). We first consider \mathbb{B}_X^p . Suppose $U \subseteq \mathbb{G}_A$ is open and $\mathbb{B}_X^p \cap U \neq \emptyset$. We show $\eta(\mathbb{B}_X^p \cap U) > 0$. Fix $a\Gamma_A \in \mathbb{B}_X^p \cap U$. Let $K = \tau^{-1}(\{a\Gamma_A\})$ and $V = \tau^{-1}(U)$. Then K is closed, and contained in $[Y]$ for any $Y \in \mathcal{R}_A$ containing $a\Gamma_A$. In particular, K is compact. Since V is open, and $K \subseteq V$, there is some $Y \in \mathcal{R}_A$ such that $K \subseteq [Y] \subseteq V$. Note that $X \in ap$ since $a\Gamma_A \in \mathbb{B}_X^p$, and $Y \in ap$ since $ap \models a\Gamma_A$. So $X \cap Y \in ap$. Since ap is generic, we have $\eta_p(X \cap Y) > 0$ by Proposition 6.7, which implies $\eta(\mathbb{B}_X^p \cap \mathbb{B}_Y^p) > 0$. Thus it suffices to show that $\mathbb{B}_Y^p \subseteq U$. So fix $g\Gamma_A \in \mathbb{B}_Y^p$. Then $Y \in gp$, i.e., $gp \in [Y] \subseteq V = \tau^{-1}(U)$. So $g\Gamma_A = \tau(gp) \in U$.

Now, for $\mathbb{G} \setminus \mathbb{B}_X^p$, recall from the proof of Lemma 6.5 that $\mathbb{G} \setminus \mathbb{B}_X^p = \bigcup_{m=1}^{\infty} \mathbb{B}_{Y_m \setminus X}^p$ for some $Y_m \in \mathcal{R}_A$. By the above argument, each $\mathbb{B}_{Y_m \setminus X}^p$ is pointwise large, which implies that $\mathbb{G} \setminus \mathbb{B}_X^p$ is pointwise large.

Part (c). Choose an arbitrary set $D \in \mathcal{R}_A$ such that $X\Gamma_A \subseteq D$. We claim that

$$\mathbb{E}_X = \tau(S^g(\mathcal{R}_A) \cap [X]) \cap \tau(S^g(\mathcal{R}_A) \cap [D \setminus X]).$$

To see this, first suppose $a\Gamma_A = \tau(q) = \tau(q')$ for some $q \in S^g(\mathcal{R}_A) \cap [X]$ and $q' \in S^g(\mathcal{R}_A) \cap [D \setminus X]$. Then $q \models a\Gamma_A \cap X$ and $q' \models a\Gamma_A \cap D \setminus X$. In particular, $a\Gamma_A \cap X \neq \emptyset$, which implies $a\Gamma_A \subseteq D$. So $a\Gamma_A \cap D \setminus X = a\Gamma_A \setminus X$, which yields

$a\Gamma_A \in \mathbb{E}_X$. Conversely, if $a\Gamma_A \in \mathbb{E}_X$ then there are $q, q' \in S^g(\mathcal{R}_A)$ such that $q \models a\Gamma_A \cap X$ and $q' \models a\Gamma_A \setminus X$. This again implies that $a\Gamma_A \setminus X = a\Gamma_A \cap D \setminus X$, and so $q \in S^g(\mathcal{R}_A) \cap [X]$, $q' \in S^g(\mathcal{R}_A) \cap [D \setminus X]$, and $a\Gamma_A = \tau(q) = \tau(q')$.

Now, since $S^g(\mathcal{R}_A) \cap [X]$ and $S^g(\mathcal{R}_A) \cap [D \setminus X]$ are both compact (recall Remark 6.3), and τ is continuous, it follows that \mathbb{E}_X is compact.

Part (d). Fix $a\Gamma_A \in \mathbb{E}_X$, and let $U \subseteq \mathbb{G}_A$ be an open neighborhood of $a\Gamma_A$. Fix $q, q' \in S^g(\mathcal{R}_A)$ such that $q \models a\Gamma_A \cap X$ and $q' \models a\Gamma_A \setminus X$. So $q, q' \in \tau^{-1}(\{a\Gamma_A\}) \subseteq \tau^{-1}(U)$, and thus $\tau^{-1}(U) \cap [X]$ is an open neighborhood of q and $\tau^{-1}(U) \setminus [X]$ is an open neighborhood of q' . By Remark 6.3, there are $g, g' \in \langle A \rangle$, such that $gp \in \tau^{-1}(U) \cap [X]$ and $g'p \in \tau^{-1}(U) \setminus [X]$. Since $p \models \Gamma_A$, we have $g\Gamma_A, g'\Gamma_A \in U$. Also, since $X \in gp$ and $X \notin g'p$, we have $g\Gamma_A \in \mathbb{B}_X^p$ and $g'\Gamma_A \notin \mathbb{B}_X^p$. So $U \cap \mathbb{B}_X^p \neq \emptyset \neq U \setminus \mathbb{B}_X^p$ and, altogether, it follows that $a\Gamma_A \in \partial \mathbb{B}_X^p$. \square

Theorem 6.12 *If $X \in \mathcal{R}_A$ then \mathbb{E}_X is compact and Haar null.*

Proof \mathbb{E}_X is compact by Lemma 6.11(c). Fix $X \in \mathcal{R}_A$. Choose $p \in S^g(\mathcal{R}_A)$ such that $p \models \Gamma$. Then $\partial \mathbb{B}_X^p$ is Haar null by Lemma 6.5, Theorem 6.9, and parts (a) and (b) of Lemma 6.11. So \mathbb{E}_X is Haar null by Lemma 6.11(d). \square

Toward applying Theorem 6.12 to prove Theorem 2.1, the first step is to remove the topological aspects and, in particular, relate Haar measures in \mathbb{G}_A to the pseudofinite measure μ_A .

Lemma 6.13 *Let η be a Haar measure on \mathbb{G}_A . Then there is some $\lambda > 0$ such that, if $K \subseteq \mathbb{G}_A$ is compact, then $\eta(K) = \inf\{\lambda\mu_A(X) : X \in \mathcal{R}_A, \pi^{-1}(K) \subseteq X\}$.*

Proof Since μ_A is a left- $\langle A \rangle$ -invariant finitely additive $\mathbb{R}_{\geq 0}$ -valued measure on the ring \mathcal{R}_A , it admits a unique left- $\langle A \rangle$ -invariant $\mathbb{R}_{\geq 0}$ -valued regular Borel measure $\tilde{\mu}_A$ on $S(\mathcal{R}_A)$ such that $\tilde{\mu}_A([X]) = \mu_A(X)$ for any $X \in \mathcal{R}_A$. Given a Borel set $W \subseteq \mathbb{G}_A$, define $\eta_A(W) = \tilde{\mu}_A(\tau^{-1}(W))$. Then it is straightforward to check that η_A is a Haar measure on \mathbb{G}_A , and that $\eta_A(K) = \inf\{\mu_A(X) : X \in \mathcal{R}_A, \pi^{-1}(K) \subseteq X\}$ for any compact $K \subseteq \mathbb{G}_A$. So there is some $\lambda > 0$ such that $\eta = \lambda\eta_A$, as desired. \square

We can now reformulate Theorem 6.12 in terms of the behavior of μ_A with respect to sets in \mathcal{R}_A . The proof is nearly the same as that of [10, Lemma 2.16].

Corollary 6.14 *Let $(W_n)_{n=0}^\infty$ be a decreasing sequence of definable subsets of $\langle A \rangle$, such that $\Gamma_A = \bigcap_{n=0}^\infty W_n$. Then for any $X \in \mathcal{R}_A$ and any $\epsilon > 0$, there is some $n \geq 0$ and $Z \in \mathcal{R}_A$ such that $\mu_A(Z) < \epsilon$ and if $g \in G^* \setminus Z$ then $\mu_A(gW_n \cap X) = 0$ or $\mu_A(gW_n \setminus A) = 0$. Moreover, if X is definable over G then we may assume that Z is definable over G .*

Proof Fix $X \in \mathcal{R}_A$ and $\epsilon > 0$. By Theorem 6.12 and Lemma 6.13, there is some $Z \in \mathcal{R}_A$ such that $\pi^{-1}(\mathbb{E}_X) \subseteq Z$ and $\mu_A(Z) < \epsilon$. If, moreover, X is definable over G , then $\pi^{-1}(\mathbb{E}_X)$ is invariant under automorphisms of G^* fixing G pointwise, and thus is \mathcal{R}_A -type-definable over G . Therefore, if X is definable over G then we may further assume that Z is \mathcal{R}_A -definable over G .

Toward a contradiction, suppose that for all $n \geq 0$ there is some $a_n \in G^* \setminus Z$ such that $\mu_A(a_n W_n \cap X) > 0$ and $\mu_A(a_n W_n \setminus X) > 0$. Note that, for all $n \geq 0$, $a_n W_n \cap X \neq \emptyset$, and so $a_n \in X W_n^{-1} \subseteq Y := X W_0^{-1}$. So $a_n \Gamma_A \in \pi(Y)$ for all $n \geq 0$.

Let $U = \{a\Gamma_A \in \mathbb{G}_A : a\Gamma_A \subseteq Z\}$, which is open in \mathbb{G}_A by Fact 4.12(d). Note that $\mathbb{E}_X \subseteq U$ and $\pi^{-1}(U) \subseteq Z$. Also, for any $n \geq 0$, we have $a_n \Gamma_A \notin U$ since $a_n \notin Z$. Since $\pi(Y)$ is compact, U is open, and \mathbb{G}_A is second countable, we may pass to a subsequence and assume that $(a_n \Gamma_A)_{n=0}^\infty$ converges to some $a\Gamma_A \in \pi(Y) \setminus U$. In particular, $a\Gamma_A \notin \mathbb{E}_X$.

Now, following the same argument as the claim in the proof of [10, Lemma 2.16], one can show that for all $n \geq 0$, $\mu_A(a W_n \cap X) > 0$ and $\mu_A(a W_n \setminus X) > 0$. By saturation, this yields $a\Gamma_X \in \mathbb{E}_X$, which is a contradiction. \square

6.3 Uniqueness of measure

Once again, we fix a sufficiently saturated group G^* and an NIP pseudofinite definable set $A \subseteq G^*$ with finite tripling. In this section, we observe that μ_A is the unique left $\langle A \rangle$ -invariant \mathcal{R}_A -Keisler measure up to scalar multiples. This will not be needed for the main results of the paper, but we include it as a nice application of generic locally compact domination.

Theorem 6.15 *Suppose ν is a left- $\langle A \rangle$ -invariant \mathcal{R}_A -Keisler measure. Then there is some $\lambda > 0$ such that $\nu(X) = \lambda \mu_A(X)$ for all $X \in \mathcal{R}_A$.*

Proof We first show that if $X \in \mathcal{R}_A$ and $\nu(X) > 0$, then X is generic. Fix $m \geq 1$. We find a finite set $F \subseteq \langle A \rangle$ such that $A^{\pm m} \subseteq XF$. Without loss of generality, we can assume $X \subseteq A^{\pm m}$. Let $Y = A^{\pm m}$, and note that Y is definable. By Corollary 5.5, there is a set \tilde{X} , which is \mathcal{R}_A -definable over G , such that $A^{\pm 2m} \subseteq \tilde{X}$. Note that \tilde{X} is pseudofinite. Let $\phi(x; y)$ be the formula $X(y \cdot x)$. Then $\phi(x; y)$ is NIP, and if $b \in Y$ then $\phi(G^*; b) \subseteq \tilde{X}$. Fix $d \geq 1$ such that $\phi(x; y)$ is d -NIP, and let $q = 2^{d+1}$. By assumption, we have $0 < \nu(X) \leq \nu(\tilde{X}) < \infty$. So we can define a finitely-additive probability measure $\tilde{\nu}$ on $\mathcal{R}^{\tilde{X}}$ such that $\tilde{\nu}(W) = \nu(W)/\nu(\tilde{X})$. In particular, for any $g \in Y$, we have $\tilde{\nu}(gX) = \tilde{\nu}(X) > 0$. By Fact 6.1, there is some integer $p \geq 1$ such that for any $g_1, \dots, g_p \in Y$ there is $I \subseteq [p]$ such that $|I| = q$ and $\tilde{\nu}(\bigcap_{i \in I} g_i X) > 0$. So $\phi(x; y)$, \tilde{X} , and Y satisfy the hypothesis of Proposition 4.21, witnessed by p and q , and thus there is a finite set $E \subseteq \tilde{X}$ such that, for any $g \in Y$, $g^{-1}X \cap E \neq \emptyset$. Setting $F = E^{-1}$, we have $A^{\pm m} = Y \subseteq XF$.

Now, as in the proof of Lemma 6.13, we extend ν to a left- $\langle A \rangle$ -invariant $\mathbb{R}_{\geq 0}$ -valued regular Borel measure $\tilde{\nu}$ on $S(\mathcal{R}_A)$ such that $\tilde{\nu}([X]) = \nu(X)$ for any $X \in \mathcal{R}_A$. As before, this induces a Haar measure η on \mathbb{G}_A such that $\eta(W) = \tilde{\nu}(\tau^{-1}(W))$ for any Borel $W \subseteq \mathbb{G}_A$. Let $S = S^g(\mathcal{R}_A)$. We show that for any $X \in \mathcal{R}_A$, $\nu(X) = \eta(\tau(S \cap [X]))$, which implies the desired result.

First, since $\nu(X) = 0$ for any non-generic $X \in \mathcal{R}_A$, and $\tilde{\nu}$ is regular, it follows that $\tilde{\nu}(S(\mathcal{R}_A) \setminus S) = 0$, and so $\tilde{\nu}(W) = \tilde{\nu}(S \cap W)$ for any Borel $W \subseteq S(\mathcal{R}_A)$. Now fix $X \in \mathcal{R}_A$. Let $W = \tau^{-1}(\tau(S \cap [X])) \setminus (S \cap [X])$. Then $\tau(S \cap Z) \subseteq \mathbb{E}_X$ and so, by

Theorem 6.12, we have $\tilde{\nu}(Z) = \tilde{\nu}(S \cap Z) \leq \eta(\tau(S \cap Z)) = 0$. Therefore

$$\nu(X) = \tilde{\nu}(S \cap X) + \tilde{\nu}(Z) = \tilde{\nu}(\tau^{-1}(\tau(S \cap [X]))) = \eta(\tau(S \cap [X])).$$

□

Remark 6.16 Using a similar strategy, one can show that the conclusion of Theorem 6.15 also holds for any *right*- $\langle A \rangle$ -invariant \mathcal{R}_A -Keisler measure.

7 Proof of the main results

In this section we will prove Theorems 2.1 and 2.2.

7.1 Main result: pseudofinite version

Let $G = \prod_{\mathcal{U}} G_i$, where each G_i is a group and \mathcal{U} is a nonprincipal ultrafilter on some index set I . We call a subset W of G an *internal approximate subgroup* if it is of the form $\prod_{\mathcal{U}} W_i$ where W_i is a k -approximate subgroup of G_i for some fixed k .

The next result transfers Corollary 6.14 to G , and reformulates the regularity statement in terms of approximate groups.

Lemma 7.1 *Suppose $A \subseteq G$ is internal, NIP, and pseudofinite with finite tripling. Then for any $\epsilon > 0$, there is an \mathcal{R}_A -definable internal approximate group $W \subseteq G$, and an \mathcal{R}_A -definable set $Z \subseteq G$, such that $W^4 \subseteq AA^{-1} \cap A^{-1}A$, A can be covered by finitely many left translates of W , $\mu_A(Z) < \epsilon$, and if $g \in G \setminus Z$ then $\mu_A(gW^4 \cap A) = 0$ or $\mu_A(gW^4 \setminus A) = 0$.*

Proof View G as a first-order structure in the group language expanded by a unary predicate naming A . Let $G^* \succ G$ be a sufficiently saturated elementary extension, and let $A_* = A(G^*)$. By Theorem 5.11, we may choose a decreasing sequence $(W_n)_{n=0}^\infty$ of subsets of $A_* A_*^{-1} \cap A_*^{-1} A_*$ such that each W_n is \mathcal{R}_{A_*} -definable over G and $\Gamma_{A_*} = \bigcap_{n=0}^\infty W_n$. Since $(\Gamma_{A_*})^4 = \Gamma_{A_*}$, it follows from saturation that $\Gamma_{A_*} = \bigcap_{n=0}^\infty W_n^4$. For any $n \geq 0$, we have $\Gamma_{A_*} \subseteq W_n$, and so W_n is generic in $\langle A_* \rangle$. It follows that for any $n \geq 0$, W_n is an approximate group and A_* can be covered by finitely many left translates of W_n . Now fix $\epsilon > 0$. By Corollary 6.14, there is $Z_* \subseteq G^*$ and $n \geq 0$ such that Z_* is \mathcal{R}_A -definable over G , $\mu_A(Z_*) < \epsilon$, and if $g \in G^* \setminus Z_*$ then $\mu_A(gW_n^4 \cap A_*) = 0$ or $\mu_A(gW_n^4 \setminus A_*) = 0$.

Now let $Z = Z_*(G)$ and $W = W_n(G)$. Since W_n and Z_* are definable over G , we have the desired conclusions (recall the remarks made after defining normalized pseudofinite counting measures in Sect. 4.4). □

Next, we state a pseudofinite version of *Ruzsa's Covering Lemma*.

Lemma 7.2 *Suppose $X, Y \subseteq G$ are internal and pseudofinite, with $\mu_Y(XY) < \infty$. Then there is $F \subseteq X$ such that $|F| \leq \mu_Y(XY)$ and $X \subseteq FY^2$.*

Proof See [5, Lemma 5.1] for the finite version, which implies the pseudofinite version via Łoś's Theorem. (Or one can directly adapt the proof to pseudofinite counting measures, e.g., as in [10, Proposition 4.9].) \square

We are now ready to bring coset nilprogressions into the picture.

Definition 7.3 A subset $P \subseteq G$ is an **internal coset nilprogression** if there are $r, s \in \mathbb{N}$ such that $P = \prod_{\mathcal{U}} P_i$, where each P_i is a coset nilprogression in G_i of rank r and step s . Moreover, P is in **normal form** if each P_i is in c -normal form for some fixed $c \geq 1$.

Now we state the pseudofinite version of the main structure theorem for approximate groups (see [5, Theorem 4.2]). The paper [5] uses the expressions “ultra-approximate subgroups” and “ultra-coset nilprogressions”, but in our context of $G = \prod_{\mathcal{U}} G_i$, and the definition above, we can say “internal” in place of “ultra”.

Theorem 7.4 (Breuillard, Green, and Tao [5]) *Suppose $W \subseteq G$ is an internal approximate group. Then there is an internal coset nilprogression $P \subseteq G$ in normal form such that $P \subseteq W^4$ and $\mu_P(W) < \infty$.*

Remark 7.5 We will apply the previous theorem in the setting where W is in fact NIP. Following the sketch in Sect. 2.2, the “stabilizer theorem step” of Theorem 7.4 (which is also called “Sanders–Croot–Sisask theory” in [5]) could be replaced by Theorem 5.11 above. On the other hand, we still need the Gleason–Yamabe steps and “escape norm” technology underlying [5, Theorem 4.2] in order recover coset nilprogressions.

We can now prove a pseudofinite version of Theorem 2.1.

Theorem 7.6 *Suppose $A \subseteq G$ is internal, NIP, and pseudofinite with finite tripling. Then for any $\epsilon > 0$, there is an internal coset nilprogression $P \subseteq G$ in normal form and an internal set $Z \subseteq AP$, with $\mu_A(Z) < \epsilon$, satisfying the following properties.*

- (i) $P \subseteq AA^{-1} \cap A^{-1}A$ and $A \subseteq CP$ for some finite $C \subseteq A$.
- (ii) There is a set $D \subseteq C$ such that $\mu_A((A \triangle DP) \setminus Z) = 0$.
- (iii) If $g \in G \setminus Z$ then $\mu_A(gP \cap A) = 0$ or $\mu_A(gP \setminus A) = 0$.

Proof Fix $\epsilon > 0$. By Lemma 7.1, there is an internal approximate group $W \subseteq G$, and a set $Z \in \mathcal{R}_A$, such that $W^4 \subseteq AA^{-1} \cap A^{-1}A$, A can be covered by finitely many left translates of W , $\mu_A(Z) < \epsilon$, and if $g \in G \setminus Z$ then $\mu_A(gW^4 \cap A) = 0$ or $\mu_A(gW^4 \setminus A) = 0$. By Theorem 7.4, there is an internal coset nilprogression $P \subseteq G$ in normal form such that $P \subseteq W^4$ and $\mu_P(W) < \infty$. In particular, since $P \subseteq W^4$, we have $\mu_A(gP \cap A) = 0$ or $\mu_A(gP \setminus A) = 0$ for any $g \in G \setminus Z$.

By assumption, we have $P = \prod_{\mathcal{U}} P_i$ where each $P_i \subseteq G_i$ is a coset nilprogression of rank r in c -normal form, for some fixed $c, r \geq 0$. For $i \in \mathbb{N}$, let $Q_i = P_i^{(1/2)}$ (recall Definition 3.13), and set $Q = \prod_{\mathcal{U}} Q_i$. Note that $Q^2 \subseteq P$. By Proposition 3.14, we have $|P_i| \leq O_{r,c}(|Q_i|)$ for all i , and so $\mu_Q(P) < \infty$.

We claim that $\mu_Q(AQ) < \infty$. Indeed, we have shown $\mu_Q(P) < \infty$, and we have $\mu_P(W) < \infty$ by assumption. Also, $\mu_W(A) < \infty$ since A can be covered by finitely

many left translates of W . Finally, $\mu_A(AQ) < \infty$ since $Q \subseteq P \subseteq W \subseteq AA^{-1}$ and A has finite tripling. Altogether,

$$\mu_Q(AQ) = \mu_Q(P)\mu_P(W)\mu_W(A)\mu_A(AQ) < \infty.$$

Now, by Lemma 7.2, there is a finite set $C \subseteq A$ such that $A \subseteq CQ^2 \subseteq CP$, and a finite set $F \subseteq A \setminus Z$ such that $A \setminus Z \subseteq FQ^2 \subseteq FP$. Without loss of generality, assume $F \subseteq C$. Define $D = \{g \in F : \mu_A(gP \setminus A) = 0\}$. Note that, since $F \cap Z = \emptyset$, we have $\mu_A(gP \cap A) = 0$ for all $g \in E := F \setminus D$. So $\mu_A(DP \setminus A) = 0$ and $\mu_A(EP \cap A) = 0$. Therefore, to show $\mu_A((A \triangle DP) \setminus Z) = 0$, it suffices to show $A \setminus DP \subseteq Z \cup (EP \cap A)$. So fix $x \in A \setminus DP$, and suppose $x \notin Z$. Then $x \in A \setminus Z \subseteq FP$, and so there is some $g \in F$ such that $x \in gP$. Since $x \notin DP$, it follows that $g \in E$, and so $x \in EP \cap A$. \square

Remark 7.7 In the previous proof, by employing [5] to trade the approximate group W for the coset nilprogression P , we lose “definability” of P in terms of \mathcal{R}_A . However, we do maintain $Z \in \mathcal{R}_A$. Therefore, in the statement of Theorem 2.1, one could add that the error set Z is a Boolean combination of bi-translates of A (of “bounded complexity” as detailed in Remark 7.10).

7.2 Proof of Theorem 2.1

We will postpone discussion of the abelian case to the end of the proof. So suppose the main statement of Theorem 2.1 fails. Then we have fixed $d, k \geq 1$ and $\epsilon > 0$ such that, for any integer $i \geq 1$, there is a group G_i and a finite d -NIP set $A_i \subseteq G_i$ with k -tripling such that, if $P \subseteq G_i$ is a coset nilprogression of rank and step i , and in i -normal form, and $Z \subseteq A_i P$ satisfies $|Z| < \epsilon |A_i|$, then it is not the case that all three of the following properties hold:

- (i) $P \subseteq A_i A_i^{-1} \cap A_i^{-1} A_i$ and $A_i \subseteq CP$ for some $C \subseteq A_i$ with $|C| \leq i$.
- (ii) There is a set $D \subseteq C$ such that $|(A_i \triangle DP) \setminus Z| < \epsilon |P|$.
- (iii) For any $g \in G_i \setminus Z$, either $|gP \cap A_i| < \epsilon |P|$ or $|gP \setminus A_i| < \epsilon |P|$.

Let \mathcal{U} be a nonprincipal ultrafilter on \mathbb{Z}^+ , and let $G = \prod_{\mathcal{U}} G_i$ and $A = \prod_{\mathcal{U}} A_i \subseteq G$. Then A is pseudofinite and d -NIP, and we have $\mu_A(A^3) \leq k < \infty$. By Theorem 7.6, there is an internal coset nilprogression $P \subseteq G$ in normal form and an internal set $Z \subseteq G$, with $\mu_A(Z) < \epsilon$, satisfying the following properties:

- * $P \subseteq AA^{-1} \cap A^{-1}A$ and $A \subseteq CP$ for some finite $C \subseteq A$.
- * There is a set $D \subseteq C$ such that $\mu_A((A \triangle DP) \setminus Z) = 0$.
- * If $g \in G \setminus Z$ then $\mu_A(gP \cap A) = 0$ or $\mu_A(gP \setminus A) = 0$.

Write $Z = \prod_{\mathcal{U}} Z_i$ where each Z_i is a subset of G_i . By definition, there are $c, r, s \in \mathbb{N}$ such that $P = \prod_{\mathcal{U}} P_i$, where each $P_i \subseteq G_i$ is a coset nilprogression of rank r and step s in c -normal form. Set $n = |C|$. Let $I \subseteq \mathbb{Z}^+$ be the set of $i \in \mathbb{Z}^+$ such that the following properties hold:

- (1) $P_i \subseteq A_i A_i^{-1} \cap A_i^{-1} A_i$ and $A_i \subseteq C_i P_i$ for some $C_i \subseteq A_i$ with $|C_i| \leq n$.
- (2) There is a set $D_i \subseteq C_i$ such that $|(A_i \triangle D_i P_i) \setminus Z_i| / |A_i| < \epsilon/n$.
- (3) If $g \in G_i \setminus Z_i$ then $|gP_i \cap A_i| / |A_i| < \epsilon/n$ or $|gP_i \setminus A_i| / |A_i| < \epsilon/n$.

(4) $Z_i \subseteq A_i P_i$ and $|Z_i|/|A_i| < \epsilon$.

Then $I \in \mathcal{U}$ and so, since \mathcal{U} is nonprincipal, we may choose some $i \in I$ such that $i \geq \max\{c, n, r, s\}$. Note that $|A_i| \leq n|P_i|$ by condition (1), and so $\epsilon/n \leq \epsilon|P_i|/|A_i|$. Altogether, conditions (1) through (4) contradict the choice of G_i and A_i .

Finally, we explain how to get proper coset progressions when restricting to abelian groups. From the above proof, it is clear that we only need to ensure that, in the statement of Theorem 7.6, if $G = \prod_{\mathcal{U}} G_i$ is abelian then P can be assumed to be an ultraproduct of proper coset progressions (i.e., we may take $c = 1$ in the proof). For this, one can transfer Theorem 1.1 to the ultraproduct G , and use this in place of Theorem 7.4. Alternatively, it is easy to check that if G is abelian and $P \subseteq G$ is an internal coset progression in normal form, then there is an internal proper coset progression $P' \subseteq P$ such that $\mu_{P'}(P) < \infty$. So Theorem 7.4 also recovers proper coset progressions in abelian groups. \square

7.3 Finite exponent: pseudofinite version

The goal of this subsection is to prove a pseudofinite analogue of Theorem 2.2, which is our main result for NIP sets of small tripling in the setting of bounded exponent. The key use of bounded exponent is essentially the fact that a locally compact Hausdorff torsion group is totally disconnected, which is a consequence of the Gleason–Yamabe theorem. In the model-theoretic setting, this fact can be used to show that if G is a sufficiently saturated group and Σ is an ind-definable subgroup of finite exponent, then any type-definable bounded-index normal subgroup of Σ is in fact an intersection of definable subgroups of Σ . This is a key ingredient in the proof of Theorem 1.3 in [5] and [16]. Our version follows [42, Corollary 5.6], which relaxes the finite exponent assumption somewhat. We also add extra control on the definability properties of the objects involved.

Proposition 7.8 *Let G^* be a sufficiently saturated group and let \mathcal{R} be a subring of $\text{Def}(G^*)$. Suppose Σ is an \mathcal{R} -ind-definable subgroup of G^* , Γ is an \mathcal{R} -type-definable bounded-index normal subgroup of G , and \mathcal{R} is left- Σ -invariant. If $X \subseteq \Sigma$ is a definable set of finite exponent containing Γ , then there is an \mathcal{R} -definable subgroup $H \leq \Sigma$ such that $\Gamma \leq H \leq X$.*

Proof Let $\pi : \Sigma \rightarrow \Sigma/\Gamma$ be the quotient map. By Fact 4.12, there is an open identity neighborhood $U \subseteq \Sigma/\Gamma$ such that $\pi^{-1}(U) \subseteq X$. By the Gleason–Yamabe-Theorem (see, e.g., [5, Theorem B.17]), there is an open subgroup $K \leq \Sigma/\Gamma$, and a compact normal subgroup $N \leq K$ such that $N \subseteq U$ and K/N is a connected Lie group. Since X has finite exponent in Σ , it follows that K/N is a connected Lie group with an identity neighborhood of finite exponent. Therefore K/N is trivial (e.g., by considering 1-parameter subgroups). So K is a compact-open subgroup of Σ/Γ contained in U . Let $H = \pi^{-1}(K)$. Then $H \in \mathcal{R}$ by Theorem 4.15 (and Remark 4.16). Moreover, H is a subgroup of Σ and $\Gamma \leq H \leq \pi^{-1}(U) \subseteq X$. \square

We can now prove a pseudofinite version of Theorem 2.2.

Theorem 7.9 Let $G = \prod_{\mathcal{U}} G_i$ where each G_i is a group and \mathcal{U} is an ultrafilter on some index set I . Suppose $A \subseteq G$ is internal, NIP, and pseudofinite with finite tripling, and assume that $AA^{-1} \cap A^{-1}A$ has finite exponent. Then for any $\epsilon > 0$, there is an \mathcal{R}_A -definable subgroup H of G and a set $Z \subseteq AH$, which is a union of left cosets of H with $\mu_A(Z) < \epsilon$, satisfying the following properties.

- (i) $H \subseteq AA^{-1} \cap A^{-1}A$ and $A \subseteq CH$ for some finite $C \subseteq A$.
- (ii) There is a set $D \subseteq C$ such that $\mu_A((A \setminus Z) \triangle DH) = 0$.
- (iii) If $g \in G \setminus Z$ then $\mu_A(gH \cap A) = 0$ or $\mu_A(gH \setminus A) = 0$.

Proof We start as in the proof of Lemma 7.1 and move to a saturated extension G^* . Let $A_* = A(G^*)$. By Proposition 7.8, there is a decreasing sequence $(H_n)_{n=0}^{\infty}$ of \mathcal{R}_{A_*} -definable subgroups of $\langle A_* \rangle$ such that $\Gamma_{A_*} = \bigcap_{n=0}^{\infty} H_n$. Since Γ_{A^*} is countably \mathcal{R}_{A_*} -type-definable over G , we may assume that each H_n is \mathcal{R}_{A_*} -definable over G . Now fix $\epsilon > 0$. By Corollary 6.14, there is a definable set $Z_* \subseteq G^*$ and some $n \geq 0$ such that $\mu_A(Z_*) < \epsilon$, and if $g \in G^* \setminus Z_*$ then $\mu_A(gH_n \cap A_*) = 0$ or $\mu_A(gH_n \setminus A_*) = 0$. Without loss of generality, we may assume that Z_* is a union of left cosets of H_n .

Let $Z = Z_*(G)$ and $H = H_n(G)$. Then, working in G , we have $H \subseteq AA^{-1} \cap A^{-1}A$, $A \subseteq CH$ for some finite $C \subseteq G$, $\mu_A(Z) < \epsilon$, and if $g \in G \setminus Z$ then $\mu_A(gH \cap A) = 0$ or $\mu_A(gH \setminus A) = 0$. We may assume $Z \subseteq AH$ since if $g \notin AH$ then $gH \cap A = \emptyset$. Similarly, we may assume that $gH \cap A \neq \emptyset$ for all $g \in C$ and thus, after changing coset representatives, assume $C \subseteq A$. Now let $D = \{g \in C \setminus Z : \mu_A(gH \setminus A) = 0\}$ and $E = \{g \in C \setminus Z : \mu_A(gH \cap A) = 0\}$. By construction,

$$(A \setminus Z) \triangle DH \subseteq Z \cup (DH \setminus A) \cup (EH \cap A),$$

and so $\mu_A((A \setminus Z) \triangle DH) = 0$. □

7.4 Proof of Theorem 2.2

The proof involves a similar ultraproduct construction as in the proof of Theorem 2.1 in Sect. 7.2. Thanks to the bounded exponent assumption, one can use Theorem 7.9 in place of Theorem 7.6 to replace coset nilprogressions by subgroups. Details are left to the reader. □

Remark 7.10 In the proof of Theorem 2.2, one can also control the “complexity” of H as a Boolean combination of bi-translates of A in the following way. Given a Boolean function $f(x_1, \dots, x_n)$ in finitely many variables, a set $A \subseteq G$, and n -tuples \bar{g}, \bar{h} of elements in G , let $f^A(\bar{g}, \bar{h})$ be the resulting Boolean combination of g_1Ah_1, \dots, g_nAh_n . Now let $(f_t)_{t=0}^{\infty}$ be an enumeration of all such Boolean functions. Then, in the conclusion of Theorem 2.2, we can say more precisely that there are $m, t \leq O_{d,k,r,\epsilon}(1)$ and n -tuples \bar{g}, \bar{h} of elements in $A^{\pm m}$, where f_t is in n variables, such that $H = f_t^A(\bar{g}, \bar{h})$. In fact, since we work with the ring \mathcal{R}_A , we may restrict to Boolean functions in disjunctive normal form, in which each disjunct contains at least one positive literal.

7.5 Theorem 2.2 via NIP arithmetic regularity in finite groups

As explained at the start of Sect. 6, Theorem 6.12 was obtained as a generalization of a related statement for pseudofinite groups from [8]. This result from [8] is used in [10, Theorem 3.1] to show that if G is a pseudofinite group of finite exponent, and $A \subseteq G$ is internal and NIP, then A can be approximated by (cosets of) a definable finite-index normal subgroup of G . Now, in the proof of Theorem 7.9, we obtained an internal *pseudofinite group* H contained in $AA^{-1} \cap A^{-1}A$, such that $A \subseteq CH$ for some finite set C . Therefore, one can apply the results from [10] to directly to H , in order to approximate any set of the form $aH \cap A$ by a finite-index subgroup H_a of H . This leads to an approximation of A by $\bigcap_{a \in C} H_a$. We leave the details (which involve some careful calculation) to the interested reader.

Altogether, this yields a proof of Theorem 2.2 that only requires Theorem 5.11 and Proposition 7.8 (in addition to [8, 10]). From a broad perspective, this strategy is analogous to the proof of the *stable* Freiman–Ruzsa result of from [20], which we discuss in the next section (see also Remark 7.16).

7.6 The stable case

In this section, we discuss the *stable* Freiman–Ruzsa result of Martin-Pizarro, Palacín, and Wolf [20].

Definition 7.11 Let G be a group. Then a set $A \subseteq G$ is **d -stable** if there do not exist $a_1, \dots, a_d, b_1, \dots, b_d \in G$ such that $a_i b_j \in A$ if and only if $i \leq j$.

Remark 7.12 Note that stability of $A \subseteq G$ can be rephrased in terms of forbidden subgraphs in the bipartite graph $\Gamma_G(A) = (G, G; yx \in A)$. Indeed, A is d -stable if and only if $\Gamma_G(A)$ omits $([d], [d]; \leq)$. One can easily check that a d -stable subset of a group is d -NIP. As we observed at the start of Sect. 2.1, a nonempty set $A \subseteq G$ is a coset of subgroup of G if and only if A is 2-stable.

As with NIP, we call a set $A \subseteq G$ **stable** if it is d -stable for some $d \geq 1$. More generally, a formula $\phi(\bar{x}; \bar{y})$ is **stable** with respect to a first-order structure M if, for some $d \geq 1$, there do not exist $\bar{a}_1, \dots, \bar{a}_d \in M^{\bar{x}}$ and $\bar{b}_1, \dots, \bar{b}_d \in M^{\bar{y}}$ such that $M \models \phi(\bar{a}_i, \bar{b}_j)$ if and only if $i \leq j$. One of the most important properties of stable formulas is the following result of Shelah (see Theorem II.2.2 in [31]).

Theorem 7.13 (Shelah [31]) *Let M be a first-order structure. Suppose $\phi(\bar{x}; \bar{y})$ is a stable formula and $p \in S_\phi(M)$ is a complete ϕ -type over M . Then the set $\{\bar{b} \in M^{\bar{y}} : \phi(\bar{x}; \bar{b}) \in p\}$ is definable.*

By Remark 7.12, one could approximate finite stable sets of small tripling by coset nilprogressions as in Theorem 2.1. However, stable sets are much more well-behaved for two reasons. First, in the setting of Theorem 5.11, if A is stable then Γ_A is an intersection of definable subgroups of $\langle A \rangle$. Second, in the setting of Theorem 6.12, if A is stable then $\mathbb{E}_X = \emptyset$ for any $X \in \mathcal{R}_A$. Both of these statements rely on Theorem 7.13, and will be shown in the proof of Theorem 7.15 below.

Let us now state the stable Freiman–Ruzsa result, which is nearly identical to [20, Theorem A], but with some small additions.

Theorem 7.14 (Martin-Pizarro, Palacín, Wolf) *Suppose G is a group and $A \subseteq G$ is a finite d -stable set with k -tripling. Given $\epsilon > 0$, there is a subgroup $H \leq G$ satisfying the following properties.*

- (i) $H \subseteq AA^{-1} \cap A^{-1}A$ and $A \subseteq CH$ for some $C \subseteq A$ with $|C| \leq O_{d,k,\epsilon}(1)$.
- (ii) There is a set $D \subseteq C$ such that $|A \Delta DH| < \epsilon|H|$.
- (iii) For any $g \in G$, either $|gH \cap A| < \epsilon|H|$ or $|gH \cap A| > (1 - \epsilon)|H|$.

Moreover, H is a finite Boolean combination of bi-translates of A .

We will give a proof of the previous theorem, which only requires Lemma 6.11(d), Theorem 7.13, and the features of Γ_A proved in Sect. 5. See Remark 7.16 for a discussion of how this compares to the proof in [20]. As was the case with Theorems 2.1 and 2.2, it suffices to prove the following pseudofinite version.

Theorem 7.15 *Let $G = \prod_{\mathcal{U}} G_i$, where each G_i is a group and \mathcal{U} is a nonprincipal ultrafilter on some index set I . Suppose $A \subseteq G$ is internal, stable, and pseudofinite with finite tripling. Then there is an \mathcal{R}_A -definable subgroup H of G satisfying the following properties.*

- (i) $H \subseteq AA^{-1} \cap A^{-1}A$ and $A \subseteq CH$ for some finite $C \subseteq A$.
- (ii) There is a set $D \subseteq C$ such that $\mu_A(A \Delta DH) = 0$.
- (iii) If $g \in G$ then $\mu_A(gH \cap A) = 0$ or $\mu(gH \setminus A) = 0$.

Proof Note that (ii) and (iii) are equivalent. We prove (i) and (iii). Let $G^* \succ G$ be sufficiently saturated and set $A_* = A(G^*)$. Let \mathcal{B} , \mathcal{B}^ℓ , and \mathcal{B}^r denote the Boolean algebras of subsets of G^* generated by all bi-translates, left translates, and right translates of A_* , respectively. Let $\Gamma = \text{Stab}_A^\ell(A_*)$.

Claim 1 Γ is $\mathcal{R}_{A_*}^r$ -definable over G .

Proof Recall that Γ is $\mathcal{R}_{A_*}^r$ -type-definable over G by Lemma 5.7. Fix a μ_A -wide type $p \in S(\mathcal{B}^r)$ such that $p \models \Gamma$. Then p is a complete ϕ -type over G for the stable formula $\phi(x; y) := A(x \cdot y)$, and so $H := \text{Stab}_{G^*}^r(p)$ is definable by Theorem 7.13. Note that $H \leq \Gamma$ since $p \models \Gamma$. Note also that $H = \text{Stab}_{A_*}^r(p \upharpoonright \mathcal{R}_{A_*}^r)$, and so $\Gamma_{A_*} \leq H$ by Lemma 5.10(a). Altogether, H is a definable subgroup of the type-definable group Γ , and $[\Gamma : H]$ is bounded since $\Gamma_{A_*} \leq H$. By a straightforward compactness exercise, H has finite index in Γ (e.g., set $X = \Gamma$ and $Y = H$ in [16, Lemma 1.6]). So Γ is definable, and therefore $\mathcal{R}_{A_*}^r$ -definable over G by saturation. $\neg\text{claim}$

By Claim 1 and Remark 5.12, $\Gamma_{A_*} = \bigcap_{n=0}^{\infty} H_n$, where each $H_n \leq G^*$ is \mathcal{R}_{A_*} -definable over G and contained in $A_* A_*^{-1} \cap A_*^{-1} A_*$. Note also that, for all $n \geq 0$, A_* is covered by finitely many left cosets of H_n .

Claim 2 $\mathbb{E}_{A_*} = \emptyset$.

Proof Let $p_* \in S(\mathcal{B})$ be a μ_A -wide type such that $p_* \models \Gamma_{A_*}$, and let $p = p_* \upharpoonright \mathcal{R}_{A_*}$. Then $\mathbb{E}_{A_*} \subseteq \partial \mathbb{B}_{A_*}^p$ by Lemma 6.11(d), and so it suffices to show that $\mathbb{B}_{A_*}^p$ is clopen in $\langle A_* \rangle / \Gamma_{A_*}$. Note that $\pi^{-1}(\mathbb{B}_{A_*}^p) = \{a \in G^* : A_* \in ap\}$. Moreover, if $p_0 = p_* \upharpoonright \mathcal{B}^\ell$, then p_0 is a complete ϕ -type with respect to the stable formula $\phi(x; y) := A(y \cdot x)$, and $\phi(x; a) \in p_0$ if and only if $A_* \in ap$. So $\pi^{-1}(\mathbb{B}_{A_*}^p)$ is definable by Theorem 7.13, and thus $\mathbb{B}_{A_*}^p$ is compact-open by Fact 4.12(b). \dashv_{claim}

By Claim 2, we can carry out the proof of Corollary 6.14 with $Z = \emptyset$. So there is some $n \geq 0$ such that if $g \in G^*$ then $\mu_A(gH_n \cap A_*) = 0$ or $\mu_A(H_n \setminus A_*) = 0$. This yields the desired result in G as usual. \square

In the context of the previous proof, one can similarly show that $\text{Stab}_A^r(A_*)$ is in $\mathcal{R}_{A_*}^\ell$, and that $\mathbb{E}_X = \emptyset$ for any $X \in \mathcal{R}_{A_*}$. Note that we also obtain an explicit description of the subgroup H in the statement of Theorem 7.15, namely, $H = \bigcap_{g \in F} g \text{Stab}_A^\ell(A)g^{-1}$ for some finite $F \subset G$.

Remark 7.16 The proof of Theorem 7.14 in [20] follows the same strategy with ultraproducts, and a similar argument involving definability of types is employed to find a definable subgroup $H \subseteq AA^{-1}$ such that A is covered by finitely many left cosets of H . They then apply the stable arithmetic regularity lemma of the authors and Terry from [9] to obtain the remaining properties. We also note that a quantitative version of Theorem 7.14 was proved recently by the first author [7].

8 The abelian case

The goal of this section is to prove Freiman–Ruzsa results for finite NIP sets with small doubling in abelian groups *directly* from arithmetic regularity results for dense NIP sets in finite groups. The reason for doing this is that these regularity results for finite abelian groups are often effective (although sometimes qualitatively weaker; see Remark 8.3 below). For example, using work of Alon, Fox, and Zhao [1], we will obtain an effective Freiman–Ruzsa result for finite NIP sets with small doubling in abelian groups of bounded exponent. Altogether, the material in this section does not require the theorems proved above (although Theorem 8.1 crucially relies on [8] and [10]).

Let us now state the two main results that will be proved in this section. The first is simply a restatement of Theorem 2.1 for abelian groups.

Theorem 8.1 *Suppose G is an abelian group and $A \subseteq G$ is a finite d -NIP set with k -doubling. Given $\epsilon > 0$, there is a proper coset progression $P \subseteq G$ of rank $O_{d,k,\epsilon}(1)$, and a set $Z \subseteq A + P$ with $|Z| < \epsilon|A|$, satisfying the following properties.*

- (i) $P \subseteq A - A$ and $A \subseteq C + P$ for some $C \subseteq A$ with $|C| \leq O_{d,k,\epsilon}(1)$.
- (ii) There is a set $D \subseteq C$ such that $|(A \Delta (D + P)) \setminus Z| < \epsilon|P|$.
- (iii) For any $g \in G \setminus Z$, either $|(g + P) \cap A| < \epsilon|P|$ or $|(g + P) \setminus A| < \epsilon|P|$.

We still do not obtain an effective bound on $O_{d,k,\epsilon}(1)$. However, the only barrier to obtaining such bounds would be an effective version of Theorem 8.8 below (see

Remark 8.10). On the other hand, using work of Alon, Fox, and Zhao [1], we will obtain an effective result in the setting of bounded exponent. The statement uses the following notation (which is explained by Lemma 8.5). Given integers $k, r \geq 1$ define

$$c_r(k) = \begin{cases} (r-1)k^{12} & \text{if } r \text{ is prime,} \\ k^2 r^{2k^2-2} & \text{otherwise.} \end{cases}$$

We now state a quantitative version of Theorem 2.2 for abelian groups.

Theorem 8.2 *Suppose G is an abelian group of exponent r , and $A \subseteq G$ is a finite d -NIP set with k -doubling. Fix $0 < \delta \leq 1$ and let $\epsilon = \delta/c_r(k)$. Then there is a subgroup $H \leq G$ satisfying the following properties.*

- (i) $H \subseteq A - A$ and $A \subseteq C + H$ for some $C \subseteq A$ with $|C| \leq \exp(O_r(d^8))\epsilon^{-d}$.
- (ii) There is a set $D \subseteq C$ such that $|A \Delta (D + H)| < \delta|A|$.
- (iii) There is a set $Z \subseteq A + H$, which is a union of cosets of H with $|Z| < \delta^{1/2}|A|$, such that if $g \in G \setminus Z$ then $|(g + H) \cap A| < \epsilon^{1/4}|H|$ or $|(g + H) \setminus A| < \epsilon^{1/4}|H|$.

The implied constant $O_r(1)$ in the bound on $|C|$ from Theorem 8.2 is the same as in [30, Theorem 11.1] up to an absolute multiplicative factor.

Remark 8.3 Although Theorem 8.2 is quantitative, the structural approximation of the set A in condition (ii) is qualitatively weaker than in the non-abelian counterparts proved using ultraproducts. Indeed, Theorem 2.2(ii) implies Theorem 8.2(ii) modulo replacing δ with δ/k^2 (as explained after Theorem 2.1). On the other hand, Theorem 2.2(ii) provides a finer description of the error between A and $D + H$, which cannot be obtained directly from Theorem 8.2(ii) alone.

Theorems 8.1 and 8.2 will both be proved using arithmetic regularity results for subsets of *finite groups*. The key ingredient that makes this reduction possible is Green and Ruzsa’s “modeling lemma”, which lies at the heart of their proof of Freiman’s Theorem for arbitrary abelian groups (Theorem 1.1 above). Roughly speaking, this lemma says that a finite set of small doubling in an abelian group can be algebraically modeled by a dense set in a finite group. In order to state the result precisely, we first recall the definition of *Freiman isomorphism*.

Definition 8.4 Suppose G and G' are abelian groups, and fix $A \subseteq G$, $A' \subseteq G'$, and $s \geq 1$. A function $\phi : A \rightarrow A'$ is a **Freiman s -isomorphism** if ϕ is bijective and for all $a_1, \dots, a_s, b_1, \dots, b_s \in A$,

$$a_1 + \dots + a_s = b_1 + \dots + b_s \Leftrightarrow \phi(a_1) + \dots + \phi(a_s) = \phi(b_1) + \dots + \phi(b_s).$$

It is easy to check that a Freiman s -isomorphism is a Freiman s' -isomorphism for any $1 \leq s' \leq s$. Next we state the “modeling lemma”.

Lemma 8.5 (Green and Ruzsa [14, 15]) *Suppose G is an abelian group and $A \subseteq G$ is a finite set with k -doubling. Then, for any $s \geq 1$, there is an abelian group G' , with $|G'| \leq (10sk)^{10k^2}|A|$ and a Freiman s -isomorphism $\phi : A \rightarrow A'$ for some*

$A' \subseteq G'$ with $0 \in A'$. Moreover, if G has exponent r then one may assume that G' has exponent r and size at most $k^2 r^{2k^2-2} |A|$. If r is prime, then one may further assume $|G'| \leq (r-1)k^{2s} |A|$.

Proof The main claim is [15, Proposition 1.2] (one can always ensure $0 \in A'$ by translating the isomorphism). If G has exponent r then, by [14, Theorem 6.1], there is $G' \leq G$ of size $k^2 r^{2k^2-2} |A|$, and some $x \in G$, such that $A \subseteq x + G'$. So $\phi: A \rightarrow A - x$ such that $\phi(a) = a - x$ is a Freiman s -isomorphism for all $s \geq 1$. The final claim when r is prime can be shown by directly generalizing [15, Proposition 6.1] (which deals with $r = s = 2$). \square

The last ingredient we need is the fact that NIP sets are preserved by Freiman isomorphism, which was first shown by Sisask [35, Lemma 4.4] (using a slightly different setup). Translated to our framework, Sisask's result implies that if $A \subseteq G$ is d -NIP in G and Freiman 2-isomorphic to $A' \subseteq G'$, then A' is at worst $(d+1)$ -NIP in G' . (This discrepancy is necessary in general, e.g., consider $A = A' = G$ and let G' be a group properly containing G .) Although this is only a minor loss, the bounds in [1] are quite sharp in terms of d . Thus it is worth noticing that the loss can be avoided when applying the modeling lemma.

Given a group G , a set $A \subseteq G$, and $\mathcal{S} \subseteq \mathcal{P}([d])$ for some $d \geq 1$, we say that A cuts out \mathcal{S} in G if there are $x_1, \dots, x_d \in G$ and $g_S \in G$ for $S \in \mathcal{S}$ such that, given $i \in [d]$ and $S \in \mathcal{S}$, $x_i \in g_S + A$ if and only if $i \in S$. Note that A is d -NIP in G if and only if it does not cut out $\mathcal{P}([d])$ in G .

Proposition 8.6 *In the statement of Lemma 8.5, if $s \geq 4$ and A is d -NIP in G then one may assume A' is d -NIP in G' .*

Proof Assume A is d -NIP in G . Note that any translate of A is d -NIP and Freiman s -isomorphic to A for any s . So we may assume $0 \in A$. Let $\phi: A \rightarrow A' \subseteq G'$ be a Freiman s -isomorphism as in Lemma 8.5, with $s \geq 4$. Without loss of generality, we may assume G' is generated by A' . In this case, we show that A' is d -NIP in G' . For a contradiction, suppose A' cuts out $\mathcal{P}([d])$ in G' , witnessed by $X' = \{x'_1, \dots, x'_d\} \subseteq G'$ and $\{g_S : S \subseteq [d]\} \subseteq G'$. After translating by $-g_{[d]}$, we may assume $X' \subseteq A'$.

Let $\mathcal{S} = \mathcal{P}([d]) \setminus \{\emptyset\}$. Then, for any $S \in \mathcal{S}$, we have $(g_S + A') \cap X' \neq \emptyset$, and so $g_S \in A' - A'$. Let $g_S = a_S - b_S$, where $a_S, b_S \in A'$, and set $h_S = \phi^{-1}(a_S) - \phi^{-1}(b_S) \in G$. We claim that, for any $S \in \mathcal{S}$, we have $\phi^{-1}(x'_i) \in h_S + A$ if and only if $i \in S$, and so A cuts out \mathcal{S} in G . So fix $S \in \mathcal{S}$ and $i \in [d]$. Suppose $\phi^{-1}(x'_i) = h_S + a$ for some $a \in A$, and let $a' = \phi(a)$. Then $\phi^{-1}(x'_i) + \phi^{-1}(b_S) = \phi^{-1}(g_S) + \phi^{-1}(a')$, and so $x'_i = g_S + a' \in g_S + A'$, which implies $i \in S$. Conversely, if $i \in S$, then there is some $a' \in A'$ such that $x'_i + b_S = a_S + a'$, and so $\phi^{-1}(x'_i) = h_S + \phi^{-1}(a') \in h_S + A$.

Now, since A does not cut out $\mathcal{P}([d])$, it follows that $(h + A) \cap X \neq \emptyset$ for all $h \in G$, i.e., $G = X - A$. Define $\psi: A - A \rightarrow A' - A'$ such that, given $a, b \in A$, $\psi(a - b) = \phi(a) - \phi(b)$. Then ψ is a well-defined Freiman $\lfloor s/2 \rfloor$ -isomorphism with domain G . Moreover, $\psi(G)$ is a subgroup of G' containing $A' = \psi(A)$ (recall that A contains 0). Therefore $G' = \psi(G) = \psi(X - A) = \phi(X) - \phi(A) = X' - A'$. But this is a contradiction, since $g_\emptyset \notin X' - A'$. \square

The proofs of Theorems 8.1 and 8.2 will follow a common strategy. Given an abelian group G and a finite NIP set $A \subseteq G$ with small doubling, we will use Lemma 8.5 (and Proposition 8.6) to move to a finite abelian group and an NIP dense set $A' \subseteq G'$. We will then apply an appropriate arithmetic regularity result in G' to show that A' can be approximated by algebraically well-structured objects that are preserved by Freiman isomorphism, and thus can be pulled back to approximate the original set A . In Theorem 8.2, these algebraic objects will be subgroups, and will transfer easily. However, for Theorem 8.1, we will start with Bohr sets (defined below) in G' , which themselves will need to be approximated by coset progressions before being pulled back to G . In order to do this secondary approximation, we will need a qualitatively sharper arithmetic regularity result (Theorem 8.8) than what is available in the literature of quantitative tools in additive combinatorics. It is for this reason that Theorem 8.1 remains ineffective. In order to state Theorem 8.8, we need the following definition.

Definition 8.7 Let G be a group, and suppose $\tau : G \rightarrow (\mathbb{R}/\mathbb{Z})^r$ is a group homomorphism, where we view $(\mathbb{R}/\mathbb{Z})^r$ as an additive group with identity $\mathbf{0}$. Given $\delta > 0$, define

$$B_\tau(\delta) = \{x \in G : d(\tau(x), \mathbf{0}) < \delta\},$$

where d is the r -fold product of the arclength metric on $\mathbb{R}/\mathbb{Z} = S^1$. A subset of G obtained in this way is called a (δ, r) -Bohr set in G .

Theorem 8.8 (C., P., Terry) *Given $d \in \mathbb{Z}^+$, $\alpha, \epsilon \in \mathbb{R}^+$, and $f : (0, 1] \times \mathbb{N} \rightarrow \mathbb{R}$, there is an integer $n = n(d, \alpha, \epsilon, f)$ such that the following holds. Suppose G is a finite abelian group and $A \subseteq G$ is d -NIP, with $|A| \geq \alpha|G|$. Then there is a (δ, r) -Bohr set $B \subseteq G$ and a set $Z \subseteq G$ such that:*

- (i) $\delta^{-1}, r \leq n$, $|Z| < \epsilon|G|$, and $B \subseteq A - A$,
- (ii) for any $g \in G \setminus Z$, either $|(g + B) \cap A| < f(\delta, r)|B|$ or $|(g + B) \setminus A| < f(\delta, r)|B|$.

Proof This is essentially the abelian case of [10, Lemma 5.6], except we have added an extra assumption that $|A| \geq \alpha|G|$ and the extra conclusion $B \subseteq A - A$. To obtain this modification, note that in the proof of [10, Lemma 5.6] if we have $G = \prod_{t \in \mathcal{U}} G_t$ and $A = \prod_{t \in \mathcal{U}} A_t$ with $|A_t| \geq \alpha|G_t|$ for all $t > 0$, then we have the extra property that $\mu(A) \geq \alpha > 0$, where μ is the G -normalized pseudofinite counting measure. Let $G^* \succ G$ be sufficiently saturated. Following the proof of [10, Lemma 5.6], we obtain a definable approximate Bohr neighborhood Y whose translates are regular for A_* (outside a set of measure less than ϵ). From the rest of the proof, one can see that it suffices to show $Y \subseteq A_* - A_*$. At this point, we turn to the proof of [10, Theorem 5.5], where Y is obtained inside some definable set W taken from a descending chain $(W_i)_{i=0}^\infty$ of definable sets whose intersection is $G_{\theta^r}^{00}$, where $\theta^r(x; y, z)$ is the formula $A(y \cdot x \cdot z)$. In this situation, we can replace W by W_i for any large enough i , without affecting the rest of the proof. Therefore, it suffices to show $G_{\theta^r}^{00} \subseteq A_* - A_*$. This follows from the fact that there is a generic θ^r -type p containing A_* by [8, Proposition 3.12], and $G_{\theta^r}^{00} = \text{Stab}_G^\ell(p)$ by [8, Theorem 3.12]. \square

The next proposition collects several standard facts from additive combinatorics concerning coset progressions and Bohr sets.

Proposition 8.9 *Let G be an abelian group.*

- (a) *Suppose $P \subseteq G$ is a coset progression, and $\phi: P \rightarrow G'$ is a Freiman 2-isomorphism from P to a subset of an abelian group G' , with $\phi(0) = 0$. Then $\phi(P)$ is a coset progression of the same rank as P . Moreover, if P is proper then so is $\phi(P)$.*
- (b) *Suppose G is finite and $B = B_\tau(\delta)$ is a (δ, r) -Bohr set in G .*
 - (i) *There is a proper coset progression $P \subseteq G$ of rank $s \leq r$ such that $B_\tau(s^{-2s}\delta) \subseteq P \subseteq B_\tau(\delta)$.*
 - (ii) *For any $A \subseteq G$ there is $X \subseteq A$ such that $|X| \leq (2/\delta)^r$ and $A \subseteq X + B$.*

Proof See [39, Proposition 5.24] for part (a),⁶ and [39, Lemma 4.22] for part (b)(i). Part (b)(ii) follows from Lemmas 2.14⁷ and 4.20 of [39], together with the fact that $B_\tau(\delta/2) + B_\tau(\delta/2) \subseteq B_\tau(\delta)$. (See also [10, Proposition 4.9].) \square

We can now prove the main theorems stated at the beginning of this section.

Proof of Theorem 8.1 Fix an abelian group G and a finite d -NIP set $A \subseteq G$ with k -doubling. By Lemma 8.5 and Proposition 8.6, there is an abelian group G' of size at most $(100k)^{10k^2}|A|$, a d -NIP subset $A' \subseteq G'$ containing 0, and a Freiman 10-isomorphism $\phi: A \rightarrow A'$. Note that it is enough to prove the desired conditions for a translate of A , and so after shifting ϕ and A , we can further assume $0 \in A$ and $\phi(0) = 0$.

Fix $\epsilon > 0$, and set $f(u, v) = \epsilon(u/2v^{2v})^{2v}$ and $\alpha = (100k)^{-10k^2}$. Let $n = n(d, \alpha, \epsilon\alpha, f)$ be as in Theorem 8.8. Then, by Theorem 8.8, there is a (δ, r) -Bohr set $B = B_\tau(\delta) \subseteq G'$, with $\delta^{-1}, r \leq n$, and a set $Z' \subseteq G'$, with $|Z'| < \epsilon\alpha|G'| \leq \epsilon|A'|$, such that $B \subseteq A' - A'$ and, for all $x \in G' \setminus Z'$, either $|(x + B) \cap A'| < f(\delta, r)|B|$ or $|(x + B) \setminus A'| < f(\delta, r)|B|$. Set $\gamma = (\delta/2r^{2r})^r$, and note that $f(\delta, r) = \epsilon\gamma^2$.

By Proposition 8.9(b)(i), we may fix a proper coset progression P' in G' of rank $s \leq r$ such that $B' := B_\tau(s^{-2s}\delta) \subseteq P' \subseteq B$. Note also that $|B| \leq \gamma|B'|$ by Proposition 8.9(b)(ii), and so $f(\delta, r)|B| = \epsilon\gamma^2|B| \leq \epsilon\gamma|B'|$. It follows that for any $x \in G' \setminus Z'$, we have $|(x + P') \cap A'| < \epsilon\gamma|B'|$ or $|(x + P') \setminus A'| < \epsilon\gamma|B'|$. Without loss of generality, we may assume $Z' \subseteq A' + P'$ while still maintaining the previous property and $|Z'| < \epsilon|A'|$. By Proposition 8.9(b)(ii), there are $C' \subseteq A'$ and $F' \subseteq A' \setminus Z'$ such that $|C'|, |F'| \leq \gamma^{-1}$, $A' \subseteq C' + B'$, and $A' \setminus Z' \subseteq F' + P'$.

Let us summarize what has been proved. We have a proper coset progression $P' \subseteq G'$ of rank $s \leq O_{d, k, \epsilon}(1)$, and a set $Z' \subseteq A' + P'$, such that:

- (1) $|Z'| < \epsilon|A'|$ and $P' \subseteq A' - A'$,
- (2) if $x \in G \setminus Z'$ then either $|(x + P') \cap A'| < \epsilon\gamma|P'|$ or $|(x + P') \setminus A'| < \epsilon\gamma|P'|$,
- (3) there is $C' \subseteq A'$ such that $|C'| \leq \gamma^{-1}$ and $A' \subseteq C' + P'$, and
- (4) there is $F' \subseteq A' \setminus Z'$ such that $|F'| \leq \gamma^{-1}$ and $A' \setminus Z' \subseteq F' + P'$.

Recall that we assumed $0 \in A'$, and so the sets $A', P', A' + P'$, and $A' + 2P'$ are all contained in $W' := 3A' - 2A'$. Let $W = 3A - 2A$. Then we have a well-defined

⁶ In [39], coset progressions are not assumed to be symmetric, but it is easy to check that a Freiman 2-isomorphism preserving the identity also preserves symmetric sets.

⁷ This is the abelian case of Ruzsa's Covering Lemma (Lemma 7.2).

Freiman 2-isomorphism $\psi: W' \rightarrow W$ such that, given $a, b, c, x, y \in A'$,

$$\psi(a + b + c - x - y) = \phi^{-1}(a) + \phi^{-1}(b) + \phi^{-1}(c) - \phi^{-1}(x) - \phi^{-1}(y).$$

Note that if $x \in A'$, then $\psi(x) = \psi(x + 0 + 0 - 0 - 0) = \phi^{-1}(x)$, and so $\psi(A') = \phi^{-1}(A') = A$. Moreover, since ψ is a Freiman 2-isomorphism and $\psi(0) = 0$, it follows that for any $X, Y \subseteq W'$, if $X + Y \subseteq W'$ then $\psi(X + Y) = \psi(X) + \psi(Y)$, and if $X - Y \subseteq W'$ then $\psi(X - Y) = \psi(X) - \psi(Y)$.

Let $P = \psi(P')$. By Proposition 8.9(a), P is a proper coset progression of rank $s \leq O_{d,k,\epsilon}(1)$. Let $Z = \psi(Z')$. Then $|Z| = |Z'| < \epsilon|A'| = \epsilon|A|$. We show that P and Z satisfy conditions (i), (ii), and (iii) in the statement of the theorem. Note first that $P \subseteq \psi(A' - A') = \psi(A') - \psi(A') = A - A$. Similarly, if we set $C = \psi(C')$ then using property (3) above, together with the fact that ψ is a bijection and $A = \psi(A')$, we have $|C| \leq \gamma^{-1} \leq O_{d,k,\epsilon}(1)$ and $C \subseteq A \subseteq C + P$. So we have (i).

Next we show (iii) in a stronger form. Fix $g \in G \setminus Z$. If $g \notin A + P$ then $(g + P) \cap A = \emptyset$, so assume $g \in A + P$. Thus we can write $g = \psi(x)$ for some $x \in (A' + P') \setminus Z'$. By property (2), either $|(x + P') \cap A'| < \epsilon\gamma|P'| = \epsilon\gamma|P|$ or $|(x + P') \setminus A'| < \epsilon\gamma|P'| = \epsilon\gamma|P|$. Note that $x + P' \subseteq A' + 2P' \subseteq W'$, and so $\psi((x + P') \cap A') = (g + P) \cap A$ and $\psi((x + P') \setminus A') = (g + P) \setminus A$. So either $|(g + P) \cap A| < \epsilon\gamma|P|$ or $|(g + P) \setminus A| < \epsilon\gamma|P|$. Since $\epsilon\gamma \leq \epsilon$, this yields (iii).

Finally, we prove (ii). Let $F = \psi(F')$. So $|F| \leq \gamma^{-1}$ and $A \setminus Z \subseteq F + P$ by property (4). Set $D = \{g \in F : |(g + P) \setminus A| < \epsilon\gamma|P|\}$, and note that if $g \in E := F \setminus D$ then $|(g + P) \cap A| < \epsilon\gamma|P|$. Since $A \setminus Z \subseteq F + P$, it follows (as in the proof of Theorem 7.6) that

$$A \triangle (D + P) \subseteq Z \cup ((D + P) \setminus A) \cup ((E + P) \cap A).$$

Therefore $|(A \triangle (D + P)) \setminus Z| \leq |F| \epsilon\gamma|P'| \leq \epsilon|P'|$. \square

Remark 8.10 As previously stated, in order to obtain explicit bounds in Theorem 8.1 via the previous proof, one would need a quantitative version of Theorem 8.8 for $f(u, v) = \epsilon(u/2v^{2v})^{2v}$. In particular, given explicit dependences $\delta(d, \alpha, \epsilon)$ and $r(d, \alpha, \epsilon)$ in Theorem 8.8, the previous proof would yield $|C| \leq (2r^{2r}/\delta)^r$ in the statement of Theorem 8.1, where $\alpha = (100k)^{-10k^2}$, $\delta = \delta(d, \alpha, \epsilon)$ and $r = r(d, \alpha, \epsilon)$. The tools developed by Sisask in [35] could be one avenue toward such a result. On the other hand, with Remark 8.3 in mind, it is perhaps more reasonable to expect a version of Theorem 8.1 with better dependence of $|C|$ on δ and r , but with the weaker statement $|A \triangle (D + P)| < \epsilon\alpha|A|$ in condition (ii).

Proof of Theorem 8.2 The strategy is the same as in Theorem 8.1, and so we will only sketch the important changes. We again fix an abelian group G and a finite d -NIP set $A \subseteq G$ with k -doubling. By Lemma 8.5 and Proposition 8.6, there is an abelian group G' of exponent r , with $|G'| \leq c_r(k)|A|$, and a Freiman 6-isomorphism $\phi': A \rightarrow A'$ where $0 \in A' \subseteq G'$ and A' is d -NIP in G' . As before, we assume without loss of generality that $0 \in A$ and $\phi(0) = 0$.

Now fix $0 < \delta \leq 1$ and let $\epsilon = \delta/c_r(k)$. Following the proof of [1, Lemma 2.4], there is a subgroup $H' \leq G'$ of index $n \leq \exp(O_r(d^8))\epsilon^{-d}$ such that $|(x + A') \triangle A'| < \epsilon|G'|$

for all $x \in H'$. Let $D' = \{x \in G' : |(x + H') \cap A'| > |H'|/2\}$. Then, as in [1, Lemma 2.4], it follows that $|A' \Delta (D' + H')| < \epsilon|G'| \leq \delta|A|$. Using similar methods (which are detailed explicitly in [6, Lemma 8.2]), it also follows that there is some $Z' \subseteq A' + H'$, which is a union of cosets of H' with $|Z'| < \epsilon^{1/2}|G'| \leq \delta^{1/2}|A|$, such that for all $x \in G' \setminus Z'$ either $|(x + H') \cap A'| < \epsilon^{1/4}|H'|$ or $|(x + H') \setminus A'| < \epsilon^{1/4}|H'|$.

Since $[G' : H'] \leq n$, we can choose $C' \subseteq G'$ with $|C'| \leq n$ such that $A' \subseteq C' + H'$. By choosing C' of minimal size, we have $(x + H') \cap A' \neq \emptyset$ for all $x \in C'$, and so after possibly changing coset representatives we may assume $C' \subseteq A'$. Similarly, we may assume without loss of generality that $D' \subseteq C'$. So the last thing we need before transferring back to G is $H' \subseteq A' - A'$. To prove this, first note that $D' \neq \emptyset$ since $\delta \leq 1$. So fix $x_0 \in D'$. For any $h \in H'$, we have $|(x_0 + H') \cap (h + A')| = |(x_0 + H') \cap A'| > |H'|/2$, which implies $(h + A') \cap A' \neq \emptyset$, i.e., $h \in A' - A'$.

Now one proceeds exactly as in the proof of Theorem 8.1 to transfer the above situation back to G via the Freiman 2-isomorphism $\psi : 2A' - A' \rightarrow 2A - A$ induced by ϕ^{-1} . (We can make do with $2A' - A'$ since $2H' = H'$.) The only extra detail required is that $\psi(H')$ is a subgroup of G , which is easily verified. \square

Remark 8.11 Given the discussion of stable sets in Sect. 7.6, a natural question is whether a quantitative version of Theorem 7.14 can be obtained for abelian groups using a similar strategy with Freiman isomorphism, together with the (effective) arithmetic regularity results of Terry and Wolf [40] for stable subsets of finite abelian groups. We leave it as an exercise to verify that this is indeed the case. In particular, using similar arguments, one can show that if $A \subseteq G$ is k -stable and Freiman s -isomorphic to $A' \subseteq G'$, for some $s \geq 2$, then A' is k -stable (in G'). On the other hand, this method is unnecessary in the stable case due to recent work of the first author [7] which directly gives a quantitative version of Theorem 7.14, with improved bounds compared to what would be obtained from applying [40].

References

1. Alon, N., Fox, J., Zhao, Y.: Efficient arithmetic regularity and removal lemmas for induced bipartite patterns. *Discrete Anal.*, Paper No. 3, 14 (2019)
2. Bogoliuboff, N.: Sur quelques propriétés arithmétiques des presque-périodes. *Ann. Chaire Phys. Math. Kiev* **4**, 185–205 (1939)
3. Breuillard, E.: A brief introduction to approximate groups, *Thin groups and superstrong approximation*, Math. Sci. Res. Inst. Publ., vol. 61. Cambridge University Press, Cambridge, pp. 23–50 (2014)
4. Breuillard, E.: Lectures on approximate groups and Hilbert’s 5th problem, *Recent trends in combinatorics*, IMA Vol. Math. Appl., vol. 159. Springer, Cham, pp. 369–404 (2016)
5. Breuillard, E., Green, B., Tao, T.: The structure of approximate groups. *Publ. Math. Inst. Hautes Études Sci.* **116**, 115–221 (2012)
6. Conant, G.: On finite sets of small tripling or small alternation in arbitrary groups. *Combin. Probab. Comput.* **29**(6), 807–829 (2020)
7. Conant, G.: Quantitative structure of stable sets in arbitrary finite groups. *Proc. Am. Math. Soc.* **149**(9), 4015–4028 (2021)
8. Conant, G., Pillay, A.: Pseudofinite groups and VC-dimension. *J. Math. Log.* **21**(2), Paper No. 2150009 (2021)
9. Conant, G., Pillay, A., Terry, C.: A group version of stable regularity. *Math. Proc. Camb. Philos. Soc.* **168**(2), 405–413 (2020)

10. Conant, G., Pillay, A., Terry, C.: Structure and regularity for subsets of groups with finite VC-dimension. *J. Eur. Math. Soc. (JEMS)* **24**(2), 583–621 (2022)
11. Dudley, R.M.: A course on empirical processes, École d’été de probabilités de Saint-Flour, XII—1982, Lecture Notes in Math., vol. 1097. Springer, Berlin, pp. 1–142 (1984)
12. Freiman, G.A.: Foundations of a Structural Theory of Set Addition. American Mathematical Society, Providence (1973). Translated from the Russian, Translations of Mathematical Monographs, vol. 37
13. Green, B.: A Szemerédi-type regularity lemma in abelian groups, with applications. *Geom. Funct. Anal.* **15**(2), 340–376 (2005)
14. Green, B., Ruzsa, I.Z.: Sets with small sumset and rectification. *Bull. Lond. Math. Soc.* **38**(1), 43–52 (2006)
15. Green, B., Ruzsa, I.Z.: Freiman’s theorem in an arbitrary abelian group. *J. Lond. Math. Soc. (2)* **75**(1), 163–175 (2007)
16. Hrushovski, E.: Stable group theory and approximate subgroups. *J. Am. Math. Soc.* **25**(1), 189–243 (2012)
17. Hrushovski, E., Peterzil, Y., Pillay, A.: Groups, measures, and the NIP. *J. Am. Math. Soc.* **21**(2), 563–596 (2008)
18. Hrushovski, E., Pillay, A.: On NIP and invariant measures. *J. Eur. Math. Soc. (JEMS)* **13**(4), 1005–1061 (2011)
19. Lascar, D., Pillay, A.: Hyperimaginaries and automorphism groups. *J. Symb. Logic* **66**(1), 127–143 (2001)
20. Martin-Pizarro, A., Palacín, D., Wolf, J.: A model-theoretic note on the Freiman–Ruzsa theorem. *Selecta Math. (N.S.)* **27**(4), Paper No. 53, 19 (2021)
21. Massicot, J.-C., Wagner, F.O.: Approximate subgroups. *J. Éc. Polytech. Math.* **2**, 55–64 (2015)
22. Matoušek, J.: Bounded VC-dimension implies a fractional Helly theorem. *Discrete Comput. Geom.* **31**(2), 251–255 (2004)
23. Newelski, L.: Topological dynamics of definable group actions. *J. Symb. Logic* **74**(1), 50–72 (2009)
24. Petridis, G.: New proofs of Plünnecke-type estimates for product sets in groups. *Combinatorica* **32**(6), 721–733 (2012)
25. Pillay, A.: Type-definability, compact Lie groups, and o-minimality. *J. Math. Log.* **4**(2), 147–162 (2004)
26. Ruzsa, I.Z.: Generalized arithmetical progressions and sumsets. *Acta Math. Hungar.* **65**(4), 379–388 (1994)
27. Ruzsa, I.Z.: An analog of Freiman’s theorem in groups. *Astérisque* **258**, 323–326 (1999)
28. Ruzsa, I.Z.: Towards a noncommutative Plünnecke-type inequality, An irregular mind, Bolyai Soc. Math. Stud., vol. 21. János Bolyai Math. Soc., Budapest, pp. 591–605 (2010)
29. Sanders, T.: On a nonabelian Balog–Szemerédi-type lemma. *J. Aust. Math. Soc.* **89**(1), 127–132 (2010)
30. Sanders, T.: On the Bogolyubov–Ruzsa lemma. *Anal. PDE* **5**(3), 627–655 (2012)
31. Shelah, S.: Classification theory and the number of nonisomorphic models, 2nd edn. Studies in Logic and the Foundations of Mathematics, vol. 92. North-Holland Publishing Co., Amsterdam (1990)
32. Simon, P.: A guide to NIP theories. Lecture Notes in Logic, vol. 44. Cambridge Scientific Publishers, Cambridge; Association for Symbolic Logic, Chicago (2015)
33. Simon, P.: Rosenthal compacta and NIP formulas. *Fundam. Math.* **231**(1), 81–92 (2015)
34. Simon, P.: VC-sets and generic compact domination. *Isr. J. Math.* **218**(1), 27–41 (2017)
35. Sisask, O.: Convolutions of sets with bounded VC-dimension are uniformly continuous. *Discrete Anal.*, Paper No. 1, 25 (2021)
36. Tao, T.: The dichotomy between structure and randomness, Presented at the 2006 International Congress of Mathematicians, Madrid. <https://www.math.ucla.edu/~tao/preprints/Slides/icmslides2.pdf>
37. Tao, T.: Product set estimates for non-commutative groups. *Combinatorica* **28**(5), 547–594 (2008)
38. Tao, T.: Hilbert’s fifth problem and related topics, Graduate Studies in Mathematics, vol. 153. American Mathematical Society, Providence (2014)
39. Tao, T., Van, V.: Additive combinatorics, Cambridge Studies in Advanced Mathematics, vol. 105. Cambridge University Press, Cambridge (2006)
40. Terry, C., Wolf, J.: Stable arithmetic regularity in the finite field model. *Bull. Lond. Math. Soc.* **51**(1), 70–88 (2019)
41. Terry, C., Wolf, J.: Quantitative structure of stable sets in finite abelian groups. *Trans. Am. Math. Soc.* **373**(6), 3885–3903 (2020)

42. van den Dries, L.: Approximate groups [according to Hrushovski and Breuillard, Green, Tao]. *Astérisque* **367–368**, Exp. No. 1077, vii, 79–113 (2015)
43. Vapnik, V.N., Červonenkis, A.Ja.: The uniform convergence of frequencies of the appearance of events to their probabilities. *Teor. Verojatnost. i Primenen.* **16**, 264–279 (1971)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.