



G. Conant · A. Pillay · C. Terry

Structure and regularity for subsets of groups with finite VC-dimension

Received February 23, 2019; revised May 29, 2020

Abstract. Suppose G is a finite group and $A \subseteq G$ is such that $\{gA : g \in G\}$ has VC-dimension strictly less than k . We find algebraically well-structured sets in G which, up to a chosen $\epsilon > 0$, describe the structure of A and behave regularly with respect to translates of A . For the subclass of groups with uniformly fixed finite exponent r , these algebraic objects are normal subgroups with index bounded in terms of k , r , and ϵ . For arbitrary groups, we use Bohr neighborhoods of bounded rank and width inside normal subgroups of bounded index. Our proofs are largely model-theoretic, and heavily rely on a structural analysis of compactifications of pseudofinite groups as inverse limits of Lie groups. The introduction of Bohr neighborhoods into the nonabelian setting uses model-theoretic methods related to the work of Breuillard, Green, and Tao [8] and Hrushovski [28] on approximate groups, as well as a result of Alekseev, Glebskiĭ, and Gordon [1] on approximate homomorphisms.

Keywords. Arithmetic regularity, VC-dimension, pseudofinite groups

1. Introduction and statement of results

Szemerédi's Regularity Lemma [50] is a fundamental result about graphs, which has found broad applications in graph theory, computer science, and arithmetic combinatorics. Roughly speaking, the regularity lemma partitions large graphs into few pieces so that almost all pairs of pieces have uniform edge density. In 2005, Green [24] proved the first *arithmetic regularity lemma*, which uses discrete Fourier analysis to define arithmetic notions of regularity for subsets of finite abelian groups. For groups of the form $\mathbb{F}_p^n$, Green's result states that given $A \subseteq \mathbb{F}_p^n$, there is $H \leq \mathbb{F}_p^n$ of bounded index such that A is uniformly distributed in almost all cosets of H (as quantified by Fourier-analytic methods; see [24, Theorem 2.1]). Arithmetic regularity lemmas, and their higher order analogues

G. Conant: DPMMS, University of Cambridge, Cambridge CB3 0WB, UK;
gconant@maths.cam.ac.uk

A. Pillay: Department of Mathematics, University of Notre Dame, Notre Dame, IN 46656, USA;
apillay@nd.edu

C. Terry: Department of Mathematics, University of Chicago, Chicago, IL 60637, USA;
caterry@uchicago.edu

Mathematics Subject Classification (2020): Primary 03C45, 03C20, 20D60, 22C05;
Secondary 22E35

(see [23, 25]), are now important tools in arithmetic combinatorics. From a general perspective, one can view regularity lemmas as tools for decomposing mathematical objects into ingredients that are easier to study because they are either highly structured (e.g., the cosets of a subgroup) or highly random (e.g., uniformly distributed).

Recently, a large body of work has developed around strengthened regularity lemmas for classes of graphs which forbid some particular bipartite configuration. This setting is fundamental in both combinatorics and model theory, although often for very different reasons. In combinatorics, forbidden configurations can lead to significant *quantitative* improvements in results about graphs, and several well-known open problems arise in this pursuit (e.g., the Erdős–Hajnal conjecture; see [17, 1.4], and also [20]). In model theory, the focus is usually on infinite objects, and forbidden configurations are used to obtain *qualitative* results about definable sets in mathematical structures. Indeed, much of modern model theory emerged from the study of mathematical structures in which every definable bipartite graph omits a finite “half-graph” as an induced subgraph (such structures are called *stable*).

In combinatorics and model theory, the practice of forbidding finite bipartite configurations is rigorously formulated using VC-dimension. By definition, the **VC-dimension** of a bipartite graph $(V, W; E)$ is the supremum of all $k \in \mathbb{Z}^+$ such that $(V, W; E)$ contains $([k], \mathcal{P}([k]); \in)$ as an induced subgraph (where $[k] = \{1, \dots, k\}$). We call $(V, W; E)$ k -**NIP** if it has VC-dimension at most $k - 1$.¹ While this definition is based on omitting one specific bipartite graph, an illuminating exercise is that if $(V, W; E)$ omits *some* finite bipartite graph (V', W', E') as an induced subgraph, then $(V, W; E)$ is k -NIP for some $k \leq |V'| + \lceil \log_2 |W'| \rceil$. Therefore, having finite VC-dimension is equivalent to omitting *some* finite bipartite configuration.

In [2, Lemma 1.6], Alon, Fischer, and Newman proved a strengthened regularity lemma for finite graphs of bounded VC-dimension,² in which the bound on the size of the partition is polynomial in the degree of irregularity (in contrast to Szemerédi’s original work, where these bounds are necessarily tower-type [21]), and the edge density in any regular pair is close to 0 or 1. This latter condition says that, as a bipartite graph, each regular pair is almost empty or complete, and so the normally “random” ingredients of Szemerédi regularity are in fact highly structured. Graph regularity with bounded VC-dimension was also developed by Lovász and Szegedy in [35], and similar results have been found for hypergraphs (e.g., [18–20]), as well as for various model-theoretic settings inside NIP (see [5, 10, 11, 37]). The strongest conclusion is for *stable graphs*,³ where Malliaris and Shelah prove the existence of regular partitions with polynomial bounds, *no irregular pairs*, and the same “0-1” behavior of edge densities in regular pairs (see [37, Theorem 5.18]).

The goal of this article is to develop arithmetic regularity for arbitrary finite groups in the context of forbidden bipartite configurations, as quantified by VC-dimension. In

¹This terminology is from model theory, where a bipartite graph with infinite VC-dimension is said to have the *independence property*, and so NIP stands for “no independence property”.

²The VC-dimension of a graph $(V; E)$ is that of its “bipartite double cover” $(V, V; E)$.

³A bipartite graph is called k -**stable** if it omits $([k], [k]; \leq)$ as an induced subgraph; and a graph is k -**stable** if its bipartite double cover is. Note that a k -stable bipartite graph is k -NIP.

analogy to the case of graphs, we show that by forbidding finite bipartite configurations, one obtains a strengthened version of arithmetic regularity in which the normally random ingredients are instead highly structured. Moreover, our proof methods deepen the connection between model theory and arithmetic combinatorics, in that we use pseudofinite methods to extend combinatorial results for finite abelian groups to the nonabelian setting. This is in the same vein as Hrushovski's [28] celebrated work on approximate groups, and the subsequent structure theory proved by Breuillard, Green, and Tao [8]. We will use similar techniques in order to formulate arithmetic regularity in nonabelian groups using *Bohr neighborhoods*, which are fundamental objects from arithmetic combinatorics in abelian groups. Finally, our results show that arithmetic regularity for NIP sets in finite groups coincides with a certain model-theoretic phenomenon called “compact domination”. This notion was first isolated by Hrushovski, Peterzil, and the second author [29] in their proof of the so-called “Pillay conjectures” for groups definable in o-minimal theories, and later played an important role in the study of definably amenable groups definable in NIP theories [9, 29, 31, 32].

Before stating the main results of this paper, we briefly recall previous work on stable arithmetic regularity, as it provides a template for the “structure and regularity” statements we will obtain in the NIP setting. Given a group G and a subset $A \subseteq G$, we define the bipartite graph $\Gamma_G(A) = (V, W; E)$ where $V = W = G$ and $E = \{(x, y) \in G^2 : yx \in A\}$.⁴ Given $k \geq 1$, we say that a subset A of a group G is k -NIP (respectively, k -stable) if $\Gamma_G(A)$ is k -NIP (respectively, k -stable), as defined above. In [52], the third author and Wolf developed arithmetic regularity for k -stable subsets of $\mathbb{F}_p^n$. They proved that such sets satisfy a strengthened version of Green's arithmetic regularity lemma, in which there is an efficient bound on the index of H and A is uniformly distributed in *all* cosets of H . They also show that a k -stable subset of $\mathbb{F}_p^n$ is approximately a union of cosets of a subgroup of small index, which is an arithmetic analogue of “0-1 density” in regular pairs. In [14, Theorem 1.2], we generalized and strengthened the results from [52] on $\mathbb{F}_p^n$ to the setting of arbitrary finite groups, but without explicit bounds.⁵

Theorem 1.1 ([14]). *For any $k \geq 1$ and $\epsilon > 0$, there is $n = n(k, \epsilon)$ such that the following holds. Suppose G is a finite group and $A \subseteq G$ is k -stable. Then there is a normal subgroup $H \leq G$, of index at most n , satisfying the following properties:*

(i) (structure) *There is a set $D \subseteq G$, which is a union of cosets of H , such that*

$$|A \triangle D| < \epsilon |H|.$$

(ii) (regularity) *For any $g \in G$, either $|gH \cap A| < \epsilon |H|$ or $|gH \setminus A| < \epsilon |H|$.*

Moreover, H is in the Boolean algebra generated by $\{gAh : g, h \in G\}$.

⁴This is a bipartite analogue of the “Cayley sum-graph of A in G ”, as defined in [52] for abelian groups. The reason we use $yx \in A$ rather than $xy \in A$ is the model-theoretic preference for “left-invariant” formulas (see Section 2.4 for details).

⁵Quantitative results for stable sets in finite abelian groups were later proved by the third author and Wolf (see [53, Theorem 4]), and more recently for arbitrary finite groups by the first author (see [12, Theorems 1.3 and 1.4]).

We now describe the main results of this paper. Our first result on arithmetic regularity in the setting of bounded VC-dimension is for k -NIP subsets of finite groups with uniformly bounded exponent.

Theorem 3.2. *For any $k, r \geq 1$ and $\epsilon > 0$, there is $n = n(k, r, \epsilon)$ such that the following holds. Suppose G is a finite group of exponent r , and $A \subseteq G$ is k -NIP. Then there are*

- * *a normal subgroup $H \leq G$ of index at most n , and*
 - * *a set $Z \subseteq G$, which is a union of cosets of H , with $|Z| < \epsilon|G|$,*
- satisfying the following properties:*

- (i) (structure) *There is a set $D \subseteq G$, which is a union of cosets of H , such that*

$$|(A \setminus Z) \triangle D| < \epsilon|H|.$$

- (ii) (regularity) *For any $g \in G \setminus Z$, either $|gH \cap A| < \epsilon|H|$ or $|gH \setminus A| < \epsilon|H|$.*

Moreover, H is in the Boolean algebra generated by $\{gAh : g, h \in G\}$.

Thus the behavior of NIP sets in bounded exponent groups is almost identical to that of stable sets in arbitrary finite groups, where the only difference is the error set Z . This reflects similar behavior in graph regularity, where the main difference between the stable and NIP cases is the need for irregular pairs. Theorem 3.2 also qualitatively generalizes and strengthens a recent quantitative result of Alon, Fox, and Zhao [3, Theorem 1.1] on k -NIP subsets of finite abelian groups of uniformly bounded exponent. A version of Theorem 3.2 with polynomial bounds (in ϵ^{-1}), but weaker qualitative ingredients, is conjectured in [3].⁶

We then turn to k -NIP sets in arbitrary finite groups. In this case, one cannot expect a statement involving only subgroups, as in Theorem 3.2. Indeed, as noted in [3, Section 5], if $p \geq 3$ is prime and $A = \{1, 2, \dots, \lfloor p/2 \rfloor\}$, then A is 4-NIP as a subset of $\mathbb{Z}/p\mathbb{Z}$, but A cannot be approximated as in Theorem 3.2 for arbitrarily small ϵ . This example illustrates a common obstacle faced in arithmetic combinatorics when working in abelian groups with very few subgroups. In situations like this, one often works instead with certain well-structured subsets of groups called *Bohr neighborhoods*. Bohr neighborhoods in cyclic groups were used in Bourgain's improvement of Roth's Theorem [7, (0.11)], and also form the basis of Green's arithmetic regularity lemma for finite abelian groups [24, Theorem 5.2]. Results related to ours involving Bohr neighborhoods, with quantitative bounds but for abelian groups, were independently obtained by Sisask [49, Theorem 1.4].

Given a group H , an integer $r \geq 0$, and some $\delta > 0$, we define a (δ, r) -**Bohr neighborhood in H** to be a set of the form $B_{\tau, \delta}^r := \{x \in H : d(\tau(x), 0) < \delta\}$, where $\tau: H \rightarrow \mathbb{T}^r$ is a group homomorphism and d is a fixed invariant metric on the r -dimensional torus $\mathbb{T}^r$ (see Remark 4.6 and Definition 4.7). Our main structure and regularity result for NIP sets in finite groups is as follows.

⁶This weaker version of Theorem 3.2 was later proved by the first author with a bound of the form $n = \exp(c_{k,r}\epsilon^{-k})$, where $c_{k,r}$ is an ineffective constant (see [13, Theorem 1.6]).

Theorem 5.7. *For any $k \geq 1$ and $\epsilon > 0$ there is $n = n(k, \epsilon)$ such that the following holds. Suppose G is a finite group and $A \subseteq G$ is k -NIP. Then there are*

- * *a normal subgroup $H \leq G$ of index $m \leq n$,*
- * *a (δ, r) -Bohr neighborhood B in H , where $0 \leq r \leq n$ and $1/n \leq \delta \leq 1$, and*
- * *a subset $Z \subseteq G$, with $|Z| < \epsilon|G|$,*

satisfying the following properties:

- (i) (structure) *There is a set $D \subseteq G$, which is a union of at most $m(\frac{2}{\delta})^r$ translates of B , such that*

$$|(A \triangle D) \setminus Z| < \epsilon|B|.$$

- (ii) (regularity) *For any $g \in G \setminus Z$, either $|gB \cap A| < \epsilon|B|$ or $|gB \setminus A| < \epsilon|B|$.*

Moreover, H and Z are in the Boolean algebra generated by $\{gAh : g, h \in G\}$, and if G is abelian then we may assume $H = G$.

In order to prove Theorems 3.2 and 5.7, we will first prove companion theorems for these results involving definable sets in infinite pseudofinite groups (Theorems 3.1 and 5.5, respectively). We then prove the theorems about finite groups by taking ultraproducts of counterexamples in order to obtain infinite pseudofinite groups contradicting the companion theorems. To prove the companion theorems, we work with a saturated pseudofinite group G , and an invariant NIP formula $\theta(x; \bar{y})$ (see Definitions 2.1 and 2.10). In [15], the first two authors proved “generic compact domination” for the quotient group $G/G_{\theta^r}^{00}$, where $\theta^r(x; \bar{y}, u) := \theta(x \cdot u; \bar{y})$ and $G_{\theta^r}^{00}$ is the intersection of all θ^r -type-definable bounded-index subgroups of G (see Definition 2.12). In this case, $G/G_{\theta^r}^{00}$ is a compact Hausdorff group, and generic compact domination roughly states that if $A \subseteq G$ is θ^r -definable, then the set of cosets of $G_{\theta^r}^{00}$, which intersect both A and $G \setminus A$ in “large” sets with respect to the pseudofinite counting measure, has Haar measure 0 (see Theorem 2.14). This is essentially a regularity statement for A with respect to the subgroup $G_{\theta^r}^{00}$, which is rather remarkable as generic compact domination originated in [29] toward proving conjectures of the second author on the Lie structure of groups definable in o-minimal theories [41, Conjecture 1.1].

For G and $\theta(x; \bar{y})$ as above, the regularity provided by generic compact domination for θ^r -definable sets in G cannot be transferred directly to finite groups, as the statement depends entirely on type-definable data (such as $G_{\theta^r}^{00}$). Thus, much of the work in this paper focuses on obtaining definable approximations to $G_{\theta^r}^{00}$ and the other objects involved in generic compact domination. We first investigate the situation when $G/G_{\theta^r}^{00}$ is a profinite group, in which case $G_{\theta^r}^{00}$ can be approximated by definable finite-index subgroups of G . Using this, we prove Theorem 3.1 (the pseudofinite companion to Theorem 3.2 above). The connection to Theorem 3.2 is that if G is elementarily equivalent to an ultraproduct of groups of uniformly bounded exponent, then $G/G_{\theta^r}^{00}$ is a compact Hausdorff group of finite exponent, hence is profinite (see Fact 2.3(d)).

When $G/G_{\theta^r}^{00}$ is not profinite, there are not enough definable finite-index subgroups available to describe $G_{\theta^r}^{00}$, and it is for this reason that we turn to Bohr neighborhoods. This is somewhat surprising, as Bohr neighborhoods are fundamentally linked to abelian

groups, and we do not make any assumptions of commutativity. However, by a general result of the second author on “definable” compactifications of pseudofinite groups, we in fact know that the *connected component* of $G/G_{\theta^r}^{00}$ is abelian (see Theorem 2.7). It is at this point that we see the beautiful partnership between pseudofinite groups and NIP formulas. Specifically, we have generic compact domination of θ^r -definable sets by the abelian-by-profinite group $G/G_{\theta^r}^{00}$, which allows us to analyze θ^r -definable sets in G using Bohr neighborhoods in definable finite-index subgroups. In order to obtain a statement involving only *definable* objects, we use approximate homomorphisms to formulate a notion of approximate Bohr neighborhoods. This leads to Theorem 5.5 (the pseudofinite companion of Theorem 5.7 above). We then apply a result of Alekseev, Glebskii, and Gordon [1, Theorem 5.13] on approximate homomorphisms to find actual Bohr neighborhoods inside approximate Bohr neighborhoods and, ultimately, prove Theorem 5.7.

In Section 6, we prove similar results for *fsg* groups definable in *distal* NIP theories (see Theorems 6.8 and 6.7). This follows our theme, as such groups satisfy a strong form of compact domination (see Lemma 6.6). In Section 7, we discuss compact p -adic analytic groups as one concrete example of the distal *fsg* setting.

2. Preliminaries

2.1. First-order structures and definability

We start by establishing the setting involving first-order structures. See [38] for an introduction to first-order logic and model theory.

Let $\mathcal{L}$ be a first-order language. Following model-theoretic convention, we will say that an $\mathcal{L}$ -structure M^* is **sufficiently saturated** if M^* is κ -saturated and strongly κ -homogeneous for some large (e.g., strongly inaccessible) cardinal κ . Let M^* be a fixed sufficiently saturated $\mathcal{L}$ -structure. We say that a set is **bounded** if its cardinality is strictly less than the saturation cardinal κ of M^* .⁷ When $\mathcal{L}$ and M^* are fixed, we will refer to an $\mathcal{L}$ -formula with parameters from M^* as simply a **formula**. We also call an elementary substructure $M \prec M^*$ **small** if its universe is bounded. As usual, a subset $X \subseteq (M^*)^n$ is **definable** if there is a formula $\phi(x_1, \dots, x_n)$ such that $X = \phi(M^*) := \{\bar{a} \in (M^*)^n : M^* \models \phi(\bar{a})\}$; and $X \subseteq (M^*)^n$ is **type-definable** if it is an intersection of a bounded number of definable subsets of $(M^*)^n$. We will further say that $X \subseteq (M^*)^n$ is **countably-definable** if it is an intersection of countably many definable subsets of $(M^*)^n$.

Many of our results will operate in a “local” model-theoretic setting, which is to say that we focus on a single formula $\theta(x; \bar{y})$ whose free variables are partitioned into a singleton x and a finite tuple $\bar{y}$.

Definition 2.1. Let $\theta(x; \bar{y})$ be a formula.

- (1) A $\langle \theta \rangle$ -**formula** is a formula $\phi(x; \bar{y}_1, \dots, \bar{y}_n)$ obtained as a Boolean combination of $\theta(x; \bar{y}_1), \dots, \theta(x; \bar{y}_n)$ for some $n \geq 1$ and variables $\bar{y}_1, \dots, \bar{y}_n$.

⁷This is not to be confused with later uses of the phrase “uniformly bounded” in the context of theorems about finite groups.

- (2) An **instance** of $\theta(x; \bar{y})$ is a formula of the form $\theta(x; \bar{b})$ for some $\bar{b} \in (M^*)^{|\bar{y}|}$.
- (3) A θ -**formula** $\phi(x)$ is an instance of a $\langle \theta \rangle$ -formula $\phi(x; \bar{y}_1, \dots, \bar{y}_n)$.
- (4) A set $X \subseteq M^*$ is θ -**definable** if $X = \phi(M^*)$ for some θ -formula $\phi(x)$.
- (5) A set $X \subseteq M^*$ is θ -**type-definable** (respectively, θ -**countably-definable**) if it is an intersection of a bounded (respectively, countable) number of θ -definable subsets of M^* .

Note that if $\theta(x; \bar{y})$ is a formula, and $X \subseteq M^*$ is a set that is both countably-definable and θ -type-definable, then X is θ -countably-definable by an easy saturation argument.

We end this subsection with some discussion on our conventions regarding definable groups. In general, we will work with a sufficiently saturated structure M^* and consider a group G definable in M^* . (For example, this will be the setting for the results in Sections 6 and 7.) Thus most of the preliminaries below will take place in this context. On the other hand, the results in Sections 3 and 5 are in the “local setting” discussed above, and thus we will work in the situation where $\mathcal{L}$ is an expansion of the group language⁸ and M^* expands a group. So in this case, we use G in place of M^* , and we just say that G is a *sufficiently saturated expansion of a group*.

In some situations, we may start by working in the setting of groups definable in M^* and then, for certain results, move to the latter setting where M^* is itself an expansion of a group. To clarify this, we adopt the following convention. Given a sufficiently saturated structure M^* and a group G definable in M^* , we will write $G = M^*$ to signify the additional assumption that $\mathcal{L}$ contains the language of groups, M^* is an expansion of a group, and G is defined by $x = x$. In other words, this assumption moves us to the setting where we work with a sufficiently saturated expansion of a group (as discussed above). This will be done primarily for the sake of convenience, and in order to avoid introducing further terminology and notation.

Definition 2.2. Let G be a sufficiently saturated expansion of a group. A formula $\theta(x; \bar{y})$ is **(left) invariant** if, for any $a \in G$ and $\bar{b} \in G^{|\bar{y}|}$, there is $\bar{c} \in G^{|\bar{y}|}$ such that $\theta(a \cdot x; \bar{b})$ and $\theta(x; \bar{c})$ define the same subset of G .

The typical example of an invariant formula is something of the form $\theta(x; y) := \phi(y \cdot x)$, where $\phi(x)$ is a formula in one variable. Note also that if $\theta(x; \bar{y})$ is an invariant formula then any $\langle \theta \rangle$ -formula is invariant as well.

2.2. Compact quotients

By convention, when we say that a topological space is **compact**, we mean *compact and Hausdorff*. We will frequently use the fact that a compact space is second-countable if and only if it is separable and metrizable (e.g., by Urysohn’s *Metrization Theorem*). Given a topological group K , we let K^0 denote the connected component of the identity in K ,

⁸In fact, the group language expanded by a single unary relation symbol A suffices for our main results on finite groups (Theorems 3.2 and 5.7).

which is a closed normal subgroup of K . A topological group is called **profinite** if it is a projective limit of finite groups. By a **Lie group**, we mean a finite-dimensional real Lie group. For $n \in \mathbb{N}$, let $\mathbb{T}^n$ denote the n -dimensional torus, where $\mathbb{T} = \mathbb{R}/\mathbb{Z}$ (so $\mathbb{T}^0$ is the trivial group). We view $\mathbb{T}^n$ as a compact topological group (with the product of the quotient topology on $\mathbb{R}/\mathbb{Z}$). We use $\cong$ for isomorphism of topological groups.

The structure theory for compact groups as projective limits of compact Lie groups will play a significant role in this paper. We refer the reader to [43, Section 1.1] for details on projective limits of topological spaces. Each projective limit below will be indexed by a **join semilattice** (*jsl*), i.e., a poset $I = (I, \leq)$ such that any finite set $F \subseteq I$ has a least upper bound (denoted $\sup F$). When we refer to $\mathbb{N}$ as a *jsl*, we always use the usual ordering.

The following are several important facts about compact groups that we will need at various points throughout the paper.

Fact 2.3.

- (a) If K and L are compact groups, and $\pi: K \rightarrow L$ is a surjective continuous homomorphism, then $\pi(K^0) = L^0$.
- (b) A compact group K is profinite if and only if K^0 is trivial.
- (c) A continuous homomorphic image of a profinite group is profinite.
- (d) Any compact torsion group is profinite.
- (e) Any compact second-countable group admits a bi-invariant compatible metric.
- (f) Any compact group K admits a unique left-invariant regular Borel probability measure η_K (called the **normalized Haar measure** on K).⁹
- (g) If K is a compact abelian Lie group, then

$$K \cong \mathbb{T}^n \times F$$

for some finite group F and some $n \in \mathbb{N}$.

- (h) If K is a compact group then there is a *jsl* I and a projective system $(L_i)_{i \in I}$ of compact Lie groups such that $K \cong \varprojlim L_i$. Moreover, the projection maps are surjective, and if K is second-countable then we may assume $I = \mathbb{N}$.

Proof. Part (a) is a straightforward consequence of the Open Mapping Theorem for compact groups [27, p. 704]. See also [26, Theorem 7.12] for a more general result in the locally compact setting.

Part (b) follows from the fact that a topological group is profinite if and only if it is compact and totally disconnected (see [43, Theorem 2.1.3]).

Part (c) follows from (a) and (b).

For (d), see [33, Theorem 4.5] (where torsion groups are called *periodic*).

For (e), (f), and (g), see [27, Corollary A4.19, Theorem 2.8, and Proposition 2.42].

⁹In fact, η_K is also right-invariant and is the unique right-invariant regular Borel probability measure on K .

Part (h) is part of the well-known *Peter–Weyl Theorems*. See [27, Corollary 2.43] for the first statement and surjectivity of the projection maps. The *jsl* in question is the collection I of kernels of continuous homomorphisms from K to unitary groups under reverse inclusion (see [27, Corollary 2.36 and its proof]). In particular, any open neighborhood of the identity in K contains some kernel in I (see [27, Exercise E9.1]), which yields the final claim on second-countability. ■

We now return to model theory. Let $\mathcal{L}$ be a first-order language. Given a sufficiently saturated $\mathcal{L}$ -structure M^* , a type-definable set $X \subseteq (M^*)^n$, and a complete n -type $p \in S_n(M^*)$, we say p **concentrates** on X , written $p \models X$, if p contains a bounded set $\{\phi_i(\bar{x})\}_{i \in I}$ of formulas such that $X = \bigcap_{i \in I} \phi_i(M^*)$ (note that, by saturation and since p is a complete type, we have $p \models X$ if and only if $\phi(\bar{x}) \in p$ for any formula $\phi(\bar{x})$ such that $X \subseteq \phi(M^*)$).

Fact 2.4. *Let G be a group definable in a sufficiently saturated $\mathcal{L}$ -structure M^* . Suppose $\Gamma \leq G$ is a type-definable normal subgroup of bounded index, and let $\pi: G \rightarrow G/\Gamma$ be the canonical homomorphism.*

- (a) *G/Γ is a compact topological group, where a set $X \subseteq G/\Gamma$ is **closed** if and only if $\pi^{-1}(X)$ is type-definable.*
- (b) *If $X \subseteq G$ is definable then the set $\{a\Gamma \in G/\Gamma : a\Gamma \subseteq X\}$ is open in G/Γ .*
- (c) *G/Γ is second-countable if and only if Γ is countably-definable.*
- (d) *Suppose $G = M^*$ and Γ is θ -type-definable for some invariant formula $\theta(x; \bar{y})$. Then $X \subseteq G/\Gamma$ is closed if and only if $\pi^{-1}(X)$ is θ -type-definable.*
- (e) *Suppose μ is a left-invariant finitely additive probability measure on the Boolean algebra $\mathcal{B}$ of definable subsets of G , and X is a closed subset of G/Γ . Then $\eta_{G/\Gamma}(X) = \inf \{\mu(Z) : Z \in \mathcal{B} \text{ and } \pi^{-1}(X) \subseteq Z\}$.*

Proof. For (a), see [41, Lemma 2.7], and for (b), see [41, Remark 2.4].

Concerning (c), see [34, Fact 1.3] for the right-to-left direction. Conversely, suppose G/Γ is second-countable and let $\{U_n : n \in \mathbb{N}\}$ be a neighborhood basis of the identity in G/Γ . For any $n \in \mathbb{N}$, $\pi^{-1}(U_n)$ is co-type-definable and contains Γ . By saturation, there is some definable set $D_n \subseteq G$ such that $\Gamma \subseteq D_n \subseteq \pi^{-1}(U_n)$. It follows that $\Gamma = \bigcap_{n=0}^{\infty} D_n$.

For (d), see [15, Corollary 4.2].

To prove (e), let $S_G(M^*)$ be the space of types concentrating on G . Recall that $S_G(M^*)$ is a Stone space, with a basis of clopen sets $[Z] := \{p \in S_G(M^*) : Z \in p\}$ where $Z \in \mathcal{B}$. Note that G acts on $S_G(M^*)$ by left multiplication. Moreover, there is a unique left-invariant regular Borel probability measure $\tilde{\mu}$ on $S_G(M^*)$ such that $\mu(Z) = \tilde{\mu}([Z])$ for any $Z \in \mathcal{B}$ (see, e.g., [46, Section 7.1]). In particular, if $C \subseteq S_G(M^*)$ is closed then $\tilde{\mu}(C) = \inf \{\mu(Z) : Z \in \mathcal{B} \text{ and } C \subseteq [Z]\}$. Now define the function $f: S_G(M^*) \rightarrow G/\Gamma$ where $f(p)$ is the unique coset $a\Gamma$ such that $p \models a\Gamma$. It is straightforward to show that $f^{-1}(X) = \{p \in S_G(M^*) : p \models \pi^{-1}(X)\}$ for any closed $X \subseteq G/\Gamma$. It follows that f is continuous and $\tilde{\mu}(f^{-1}(X)) = \inf \{\mu(Z) : Z \in \mathcal{B} \text{ and } \pi^{-1}(X) \subseteq Z\}$ for any closed

$X \subseteq G/\Gamma$. Finally, one checks that $\tilde{\mu} \circ f^{-1}$ is a left-invariant regular Borel probability measure on G/Γ , and thus must be $\eta_{G/\Gamma}$ by Fact 2.3(f). ■

The topology defined on G/Γ in the previous fact is called the *logic topology*, and the canonical homomorphism from G to G/Γ connects topological structure in G/Γ to definable sets in G . This is a special case of the following notion.

Definition 2.5. Suppose M^* is a sufficiently saturated $\mathcal{L}$ -structure, $X \subseteq M^*$ is definable, Y is a compact space, and $f: X \rightarrow Y$ is a function. Then f is **definable** if $f^{-1}(C)$ is type-definable for any closed $C \subseteq Y$. If X is θ -definable and each $f^{-1}(C)$ is θ -type-definable, for some fixed formula $\theta(x; \bar{y})$, then we say f is **θ -definable**.

Remark 2.6. Suppose $f: X \rightarrow Y$ is definable (as above).

- (a) Fix $C \subseteq U \subseteq Y$ with C closed and U open. Then, by saturation of M^* , there is a definable set $D \subseteq X$ such that $f^{-1}(C) \subseteq D \subseteq f^{-1}(U)$. Thus if $C \subseteq Y$ is clopen then $f^{-1}(C)$ is definable.
- (b) If $f(X)$ is finite then the fibers of f partition X into finitely many type-definable sets which, by saturation, implies that each fiber of f is definable.

Suppose further that X and f are θ -definable for some fixed formula $\theta(x; \bar{y})$. Then in part (a) we may assume D is θ -definable, and so if $C \subseteq Y$ is clopen then $f^{-1}(C)$ is θ -definable. Moreover, in part (b), we have that each fiber of f is θ -definable.

Recall that a first-order structure is **pseudofinite** if it is elementarily equivalent to an ultraproduct of finite structures. We now recall a result of the second author on definable compactifications of pseudofinite groups, which will be the key ingredient that allows us to introduce Bohr neighborhoods into the setting of possibly nonabelian finite groups.

Theorem 2.7 ([42]). *Suppose G is a sufficiently saturated pseudofinite expansion of a group, and $\Gamma \leq G$ is normal and type-definable of bounded index. Then $(G/\Gamma)^0$ is abelian.*

Proof. Suppose $M \prec G$ is a small substructure such that Γ is type-definable over M . Then $G_M^{00} \leq \Gamma$, where G_M^{00} denotes the intersection of *all* bounded-index subgroups of G that are type-definable over M . By [42, Theorem 2.2], $(G/G_M^{00})^0$ is abelian, and so $(G/\Gamma)^0$ is abelian by Fact 2.3(a) and since the canonical homomorphism from G/G_M^{00} to G/Γ is surjective and continuous. ■

Remark 2.8. The proof of [42, Theorem 2.2] uses the classification of approximate groups due to Breuillard, Green, and Tao [8]. This theorem was later generalized by Nikolov, Schneider, and Thom [39, Theorem 8], who showed that the connected component of *any* compactification of an (abstract) pseudofinite group is abelian. Their proof uses the classification of finite simple groups.

We are now ready to collect all of the previous facts and prove the main result of this subsection. In particular, let G be a group definable in a sufficiently saturated structure M^* , and let $\Gamma \leq G$ be type-definable and normal of bounded index. Since G/Γ is

a compact group, and the topology on G/Γ is controlled by type-definable objects in G , we can analyze Γ using the structure theory for compact Lie groups. The next lemma describes the main ingredients of this analysis. Given a poset I , we say that a net $(X_i)_{i \in I}$ of subsets of some fixed set X is **decreasing** if $X_j \subseteq X_i$ for any $i, j \in I$ such that $i \leq j$.

Lemma 2.9. *Let G be a group definable in a sufficiently saturated $\mathcal{L}$ -structure M^* . Suppose $\Gamma \leq G$ is type-definable and normal of bounded index. Then there is a bounded jsl I , a decreasing net $(\Gamma_i)_{i \in I}$ of countably-definable bounded-index normal subgroups of G , and a decreasing net $(H_i)_{i \in I}$ of definable finite-index normal subgroups of G , such that the following properties are satisfied:*

- (i) $\Gamma = \bigcap_{i \in I} \Gamma_i$, and $\Gamma_i \leq H_i$ for all $i \in I$.
- (ii) For all $i \in I$, G/Γ_i is a compact Lie group and $H_i/\Gamma_i = (G/\Gamma_i)^0$.
- (iii) $G/\Gamma \cong \varprojlim G/\Gamma_i$ and $(G/\Gamma)^0 \cong \varprojlim H_i/\Gamma_i$.¹⁰

Moreover:

- (a) If Γ is countably-definable then we may assume $I = \mathbb{N}$.
- (b) If $G = M^*$ and Γ is θ -type-definable for some invariant formula $\theta(x; \bar{y})$ then, for all $i \in I$, the quotient map $G \rightarrow G/\Gamma_i$ is θ -definable, and so we may assume that H_i is θ -definable and Γ_i is θ -type-definable.
- (c) If G/Γ is profinite then we may assume $H_i = \Gamma_i$ for all $i \in I$.
- (d) If $G = M^*$ and G is pseudofinite, then we may assume that for all $i \in I$, there is some $n_i \in \mathbb{N}$ such that $H_i/\Gamma_i \cong \mathbb{T}^{n_i}$.
- (e) If G is abelian then we may assume that, for all $i \in I$, there is some $n_i \in \mathbb{N}$ and some finite group F_i such that $G/\Gamma_i \cong \mathbb{T}^{n_i} \times F_i$ (and so $H_i/\Gamma_i \cong \mathbb{T}^{n_i}$).

Proof. By Fact 2.3(h), there is a projective system $(L_i)_{i \in I}$ of compact Lie groups such that $G/\Gamma \cong \varprojlim L_i$ and the projection maps $f_i: G/\Gamma \rightarrow L_i$ are surjective. Since Γ has bounded index in G , we may assume I is bounded. For $i \in I$, set $\Gamma_i = \ker(f_i \circ \pi)$, where $\pi: G \rightarrow G/\Gamma$ is the canonical homomorphism. Then each Γ_i is a type-definable bounded-index normal subgroup of G , and $\Gamma = \bigcap_{i \in I} \Gamma_i$. Moreover, $G/\Gamma_i \cong L_i$ for all $i \in I$, and so $G/\Gamma \cong \varprojlim G/\Gamma_i$ with canonical (surjective) projection maps $\tau_i: G/\Gamma \rightarrow G/\Gamma_i$. Since each G/Γ_i is a Lie group (and thus second-countable), each Γ_i is countably-definable by Fact 2.4(c).

Now, given $i \in I$, let $H_i \leq G$ be the pullback of $(G/\Gamma_i)^0$ under $G \rightarrow G/\Gamma_i$ (so $(G/\Gamma_i)^0 = H_i/\Gamma_i$). Then for any $i \in I$, H_i/Γ_i is a clopen finite-index normal subgroup of G/Γ_i .¹¹ So each H_i is a finite-index normal subgroup of G and, since $G \rightarrow G/\Gamma_i$ is definable, H_i is definable by Remark 2.6(a). If $i, j \in I$ and $i \leq j$, then the pullback of H_i/Γ_i to G/Γ_j is a clopen finite-index normal subgroup, and thus contains H_j/Γ_j ,

¹⁰Here we view $(G/\Gamma_i)_{i \in I}$ and $(H_i/\Gamma_i)_{i \in I}$ as projective systems using (i) and the assumption that $(H_i)_{i \in I}$ and $(\Gamma_i)_{i \in I}$ are decreasing.

¹¹If K is a Lie group then it is locally connected, and so K^0 is clopen. Thus, if K is also compact, then K^0 has finite index.

which implies $H_j \leq H_i$. Finally, $\tau_i((G/\Gamma)^0) = H_i/\Gamma_i$ for all $i \in I$ by Fact 2.3(a), and so $(G/\Gamma)^0 \cong \varprojlim H_i/\Gamma_i$ by [43, Corollary 1.1.8]. This finishes the proof of claims (i) through (iii).

Now we deal with the remaining claims. Claim (a) follows from Facts 2.3(h) and 2.4(c), and claim (e) follows from Fact 2.3(g). For claim (b), note that the quotient map $G \rightarrow G/\Gamma$ is θ -definable by Fact 2.4(d), and thus the remaining claims are evident from the above arguments. For claim (c), if G/Γ is profinite then each group H_i/Γ_i is trivial by Fact 2.3(b, c). Finally, for claim (d), assume $G = M^*$ and suppose G is pseudofinite. Then $(G/\Gamma)^0$ is abelian by Theorem 2.7. So, for any $i \in I$, H_i/Γ_i is a compact connected abelian Lie group, and thus $H_i/\Gamma_i \cong \mathbb{T}^{n_i}$ for some $n_i \in \mathbb{N}$ by Fact 2.3(g). ■

2.3. NIP formulas in pseudofinite groups

In this subsection, we assume that $\mathcal{L}$ expands the group language. Let G be a sufficiently saturated expansion of a group.

Definition 2.10. Let $\theta(x; \bar{y})$ be a formula. Given $k \geq 1$, $\theta(x; \bar{y})$ is k -**NIP** if there do not exist sequences $(a_i)_{i \in [k]}$ in G and $(\bar{b}_I)_{I \subseteq [k]}$ in $G^{\bar{y}}$ such that $\theta(a_i, \bar{b}_I)$ holds if and only if $i \in I$. We say $\theta(x; \bar{y})$ is **NIP** if it is k -NIP for some $k \geq 1$.

Remark 2.11. Given $k \geq 1$, a formula $\theta(x; \bar{y})$ is k -NIP if and only if the set system $\{\theta(G; \bar{b}) : \bar{b} \in G^{|\bar{y}|}\}$ on G has VC-dimension at most $k - 1$ (see, e.g., [46, Section 6.1] for details on set systems and VC-dimension).

Next, we summarize several main results from [15], which will form the basis for our work on NIP sets in finite and pseudofinite groups. We say that a set $A \subseteq G$ is **generic** if G is covered by finitely many left translates of A ; and a formula $\phi(x)$ is **generic** if $\phi(G)$ is generic. Given a formula $\theta(x; \bar{y})$, let $S_\theta(G)$ denote the space of complete θ -types over G (i.e., ultrafilters over the Boolean algebra of θ -formulas). A type $p \in S_\theta(G)$ is **generic** if every formula in p is generic.

Definition 2.12. Given a formula $\theta(x; \bar{y})$, we let $\theta^r(x; \bar{y}, u)$ denote the formula $\theta(x \cdot u; \bar{y})$, and $G_{\theta^r}^{00}$ denote the intersection of all θ^r -type-definable bounded-index subgroups of G .

Note that if $\theta(x; \bar{y})$ is invariant, then so is $\theta^r(x; \bar{y}, u)$. One can also show that if $\theta(x; \bar{y})$ is invariant and NIP, and $\phi(x)$ is a θ^r -formula, then $\phi(y \cdot x)$ is NIP.¹² We now focus on the case when G is pseudofinite, which, by saturation, implies that G is an elementary extension of an ultraproduct $\prod_{\mathcal{U}} G_i$, where each G_i is a finite $\mathcal{L}$ -structure expanding a group and $\mathcal{U}$ is a (nonprincipal) ultrafilter on some index set I . In this case, μ will always denote the **pseudofinite counting measure** on G , which is obtained by lifting (e.g., as in [29, Section 2]) the ultralimit measure $\lim_{\mathcal{U}} \mu_i$ on $\prod_{\mathcal{U}} G_i$, where μ_i is the normalized counting measure on G_i .

¹²However, $\theta^r(x; \bar{y}, u)$ itself need not be NIP (see [15, Example 3.7]).

Remark 2.13. In several proofs, we will apply Łoś's Theorem to properties of μ , which requires an expanded language $\mathcal{L}^+$ containing a sort for the ordered interval $[0, 1]$ with a distance function, and functions from the G -sort to $[0, 1]$ giving the measures of $\mathcal{L}$ -formulas. There are many accounts of this kind of formalism, and so we will omit further details and refer the reader to similar treatments in the literature, for example [15, Section 2.3], [28, Section 2.6], and [29, Section 2].

Before stating the next theorem, we “localize” the notation given before Fact 2.4. Given a formula $\theta(x; \bar{y})$, a θ -type-definable set $X \subseteq G$, and a complete θ -type $p \in S_\theta(G)$, we write $p \models X$ if p contains a bounded set $\{\phi_i(x)\}_{i \in I}$ of θ -formulas such that $X = \bigcap_{i \in I} \phi_i(G)$.

Theorem 2.14 ([15]). *Let G be a sufficiently saturated pseudofinite expansion of a group, and suppose $\theta(x; \bar{y})$ is an invariant NIP formula.*

- (a) *A θ^r -formula $\phi(x)$ is generic if and only if $\mu(\phi(x)) > 0$.*
- (b) *There are generic θ^r -types in $S_{\theta^r}(G)$.*
- (c) *$G_{\theta^r}^{00}$ is θ^r -countably-definable and normal of bounded index.*
- (d) *Suppose $A \subseteq G$ is θ^r -definable, and let $E \subseteq G/G_{\theta^r}^{00}$ be the set of $C \in G/G_{\theta^r}^{00}$ such that $p \models C \cap A$ and $q \models C \cap (G \setminus A)$ for some generic $p, q \in S_{\theta^r}(G)$. Then E is closed and $\eta_{G/G_{\theta^r}^{00}}(E) = 0$.*

Proof. See [15, Corollary 2.15, Proposition 3.12(a)] for part (a). Parts (b), (c), and (d) are Proposition 3.12(c), Theorem 3.15(a, e), and Theorem 5.2 in [15]. ■

Remark 2.15. In the setting of the previous fact, part (a) implies that if $X \subseteq G$ is θ^r -type-definable, then $p \models X$ for some generic $p \in S_{\theta^r}(G)$ if and only if $\mu(W) > 0$ for any definable (equivalently, any θ^r -definable) set $W \subseteq G$ containing X .

We now prove the main result of this subsection.

Lemma 2.16. *Let G be a sufficiently saturated pseudofinite expansion of a group. Suppose $\theta(x; \bar{y})$ is an invariant NIP formula, and let $(W_i)_{i=0}^\infty$ be a decreasing sequence of definable sets such that $G_{\theta^r}^{00} = \bigcap_{i=0}^\infty W_i$. Fix a θ^r -definable set $A \subseteq G$. Then, for any $\epsilon > 0$, there is a θ^r -definable set $Z \subseteq G$ and some $i \in \mathbb{N}$ such that $\mu(Z) < \epsilon$ and, for any $g \in G \setminus Z$, either $\mu(gW_i \cap A) = 0$ or $\mu(gW_i \setminus A) = 0$.*

Proof. To simplify notation, let $\Gamma = G_{\theta^r}^{00}$ and $K = G/\Gamma$, and let $\pi: G \rightarrow K$ be the canonical homomorphism. Let E be as in Theorem 2.14(d) (with respect to the fixed θ^r -definable set A). So E is closed and $\eta_K(E) = 0$.

Fix $\epsilon > 0$. By Fact 2.4(e), we may fix a definable set $Z' \subseteq G$ such that $\mu(Z') < \epsilon$ and $\pi^{-1}(E) \subseteq Z'$. By saturation and Fact 2.4(d), there is a θ^r -definable set $Z \subseteq G$ such that $\pi^{-1}(E) \subseteq Z \subseteq Z'$. In particular, we still have $\mu(Z) < \epsilon$. We show that Z satisfies the conclusion of the lemma. To motivate the argument, we first make a side remark. By saturation and Remark 2.15, we immediately see that for any $g \in G \setminus Z$, there is some $i \in \mathbb{N}$ such that $\mu(gW_i \cap A) = 0$ or $\mu(gW_i \setminus A) = 0$. Thus the content of the following argument is to show that we can pick one $i \in \mathbb{N}$ that works for every $g \in G \setminus Z$.

Toward a contradiction, suppose that for all $i \in \mathbb{N}$ there is some $a_i \in G \setminus Z$ such that $\mu(a_i W_i \cap A) > 0$ and $\mu(a_i W_i \setminus A) > 0$. Let $U = \{C \in K : C \subseteq Z\}$, which is open in K by Fact 2.4(b). Note that $E \subseteq U$ and $\pi^{-1}(U) \subseteq Z$. In particular, $(a_i \Gamma)_{i=0}^\infty$ is an infinite sequence in $K \setminus U$. Since U is open and K is compact and second-countable (by Fact 2.4(c) and Theorem 2.14(c)), we may pass to a subsequence and assume that $(a_i \Gamma)_{i=0}^\infty$ converges to some $a\Gamma \in K \setminus U$. In particular, $a\Gamma \notin E$.

Claim. For all $i \in \mathbb{N}$, $\mu(a W_i \cap A) > 0$ and $\mu(a W_i \setminus A) > 0$.

Proof. First, given $i \in \mathbb{N}$, let $U_i = \{C \in K : C \subseteq a W_i\}$. As above, each U_i is open in K . Moreover, for any $i \in \mathbb{N}$, since $\Gamma \subseteq W_i$, we have $a\Gamma \in U_i \subseteq a W_i / \Gamma$. Also, since $(W_i)_{i=0}^\infty$ is decreasing and $\Gamma = \bigcap_{i=0}^\infty W_i$ is a subgroup of G , it follows from saturation that, for all $i \in \mathbb{N}$, there is $n_i \in \mathbb{N}$ such that $W_{n_i} W_{n_i} \subseteq W_i$.

Now fix $i \in \mathbb{N}$. Since U_{n_i} is an open neighborhood of $a\Gamma$, there is $j \geq n_i$ such that $a_j \Gamma \in U_{n_i}$. In particular, $a_j \in \pi^{-1}(U_{n_i}) \subseteq a W_{n_i}$. Now we have $a_j W_j \subseteq a_j W_{n_i} \subseteq a W_{n_i} W_{n_i} \subseteq a W_i$. Since $\mu(a_j W_j \cap A) > 0$ and $\mu(a_j W_j \setminus A) > 0$, we have $\mu(a W_i \cap A) > 0$ and $\mu(a W_i \setminus A) > 0$. $\dashv$

Now, since $(W_i)_{i=0}^\infty$ is decreasing and $\Gamma = \bigcap_{i=0}^\infty W_i$, it follows from saturation that any definable set containing $a\Gamma \cap A$ (respectively, $a\Gamma \setminus A$) contains $a W_i \cap A$ (respectively, $a W_i \setminus A$) for some $i \in \mathbb{N}$. So $a\Gamma \in E$ by the claim (and Remark 2.15), which is a contradiction. $\blacksquare$

Note that Lemma 2.16 provides a “regularity” statement for invariant NIP formulas in pseudofinite groups. However, it does not provide any information about the shape of the definable sets W_i , other than that they approximate the subgroup $G_{\theta_r}^{00}$. Indeed, most of the remaining work in this paper is toward replacing these definable sets with ones that enjoy more meaningful algebraic properties.

Remark 2.17. Except for Sections 6 and 7, our use of finite VC-dimension is based entirely on applications of Lemma 2.16. So it is worth emphasizing that the proof of Theorem 2.14 in [15] heavily uses fundamental results about VC-dimension such as the Sauer–Shelah Lemma and the VC-Theorem (see [46, Chapter 6]), as well as work of Simon [47, 48] on NIP formulas, which builds on a large body of research on NIP theories and groups definable in such theories (e.g., [9, 29, 31]).

2.4. NIP subsets of groups

In this subsection, we briefly recall the notion of an NIP subset of an arbitrary group, and clarify the relationship to VC-dimension and NIP formulas. Recall from the introduction that if G is a group and $A \subseteq G$, then we have the bipartite graph $\Gamma_G(A) = (G, G; E)$ where $E = \{(x, y) \in G^2 : yx \in A\}$.

Definition 2.18. Given a group G and an integer $k \geq 1$, we say that A is k -NIP (in G) if $\Gamma_G(A)$ is k -NIP, i.e., it omits the bipartite graph $([k], \mathcal{P}([k]); \in)$ as an induced subgraph (where $[k] = \{1, \dots, k\}$).

Remark 2.19. Fix a group G and a subset $A \subseteq G$.

- (1) If $\Gamma_G(A)$ omits *some* finite bipartite graph (V, W, E) as an induced subgraph, then A is k -NIP for some $k \leq |V| + \lceil \log_2 |W| \rceil$.
- (2) Given $k \geq 1$, A is k -NIP if and only if the formula

$$\theta(x; y) := A(y \cdot x)$$

is k -NIP (here we view G as a structure in the group language expanded by a predicate for A).

- (3) Given $k \geq 1$, A is k -NIP if and only if the set system $\{gA : g \in G\}$ of left translates of A has VC-dimension at most $k - 1$ (see Remark 2.11).

Note that the formula $\theta(x; y)$ defined in (2) above is invariant. This is the main reason we define $\Gamma_G(A)$ in terms of the edge relation $yx \in A$ and not $xy \in A$. However, an important fact is that if a bipartite graph $(V, W; E)$ is k -NIP then the “opposite graph” $(W, V; \{(w, v) : E(v, w)\})$ is 2^k -NIP (see [46, Lemma 6.3]).¹³ So the order of the group operation when defining k -NIP sets only affects the precise value of k , and not whether the set is NIP overall.

3. Structure and regularity: the profinite case

In this section, we prove a structure and regularity theorem for θ^r -definable sets in a sufficiently saturated pseudofinite group G , where $\theta(x; \bar{y})$ is an invariant NIP formula and $G/G_{\theta^r}^{00}$ is profinite (Theorem 3.1 below). As an application, we obtain a structure and regularity theorem for NIP sets in finite groups of uniformly bounded exponent. We also view Theorem 3.1 as a warm-up to the general case. Indeed, this theorem follows almost immediately from Lemma 2.16 and the fact that profinite quotients correspond to type-definable subgroups that are intersections of definable subgroups (via Lemma 2.9(c)).

Theorem 3.1. *Let G be a sufficiently saturated pseudofinite expansion of a group, and suppose $\theta(x; \bar{y})$ is an invariant NIP formula. Assume $G/G_{\theta^r}^{00}$ is profinite. Fix a θ^r -definable set $A \subseteq G$ and some $\epsilon > 0$. Then there are*

- * a θ^r -definable finite-index normal subgroup $H \leq G$, and
- * a set $Z \subseteq G$, which is a union of cosets of H , with $\mu(Z) < \epsilon$,

satisfying the following properties:

- (i) (structure) *There is a set $D \subseteq G$, which is a union of cosets of H , such that*

$$\mu((A \setminus Z) \triangle D) = 0.$$

- (ii) (regularity) *For any $g \in G \setminus Z$, either $\mu(gH \cap A) = 0$ or $\mu(gH \setminus A) = 0$.*

¹³On the other hand, if a bipartite graph is k -stable (as defined in the introduction) then one can easily check that the opposite graph is as well. This reconciles our definitions with those in [14], where we defined k -stable subsets of groups using the relation $xy \in A$.

Proof. By Theorem 2.14(c) and Lemma 2.9(a, c), there is a decreasing sequence $(H_i)_{i=0}^\infty$ of θ^r -definable finite-index normal subgroups of G such that $G_{\theta^r}^{00} = \bigcap_{i=0}^\infty H_i$. By Lemma 2.16, there is a θ^r -definable set $Z' \subseteq G$ and some $i \in \mathbb{N}$ such that $\mu(Z') < \epsilon$ and, if $H := H_i$, then for any $g \in G \setminus Z'$, either $\mu(gH \cap A) = 0$ or $\mu(gH \setminus A) = 0$. Let $Z = \{g \in G : gH \subseteq Z'\}$. Then Z is a union of cosets of H , and thus is θ^r -definable since $[G : H]$ is finite. Note that $\mu(Z) < \epsilon$ since $Z \subseteq Z'$. Moreover, if $g \in G \setminus Z$ then $gH = g'H$ for some $g' \in G \setminus Z'$, and so we have condition (ii). Now let $D = \bigcup \{gH : g \in G \setminus Z \text{ and } \mu(gH \cap A) > 0\}$. Then D is θ^r -definable since $[G : H]$ is finite. Moreover, $\mu((A \setminus Z) \triangle D) = 0$ by condition (ii) and since $[G : H]$ is finite. So we have condition (i). ■

We now prove structure and regularity for NIP sets in finite groups of *uniformly bounded exponent*. In this case, we obtain the optimal situation where NIP sets are entirely controlled by finite-index subgroups up to small error. This is related to a similar result of Alon, Fox, and Zhao [3, Theorem 1.1] on finite *abelian* groups of bounded exponent. Our result is stronger in the sense that the abelian assumption is removed and the structural conclusions are improved, but also weaker in the sense that we do not obtain explicit bounds. This is analogous to the comparison of our stable arithmetic regularity lemma in [14] (Theorem 1.1 above) to the work of the third author and Wolf [52] on stable sets in $\mathbb{F}_p^n$.

Theorem 3.2. *For any $k, r \geq 1$ and $\epsilon > 0$, there is $n = n(k, r, \epsilon)$ such that the following holds. Suppose G is a finite group of exponent r and $A \subseteq G$ is k -NIP. Then there are*

- * *a normal subgroup $H \leq G$ of index at most n , and*
- * *a set $Z \subseteq G$, which is a union of cosets of H with $|Z| < \epsilon|G|$,*

satisfying the following properties:

- (i) (structure) *There is a set $D \subseteq G$, which is a union of cosets of H , such that*

$$|(A \setminus Z) \triangle D| < \epsilon|H|.$$

- (ii) (regularity) *For any $g \in G \setminus Z$, either $|gH \cap A| < \epsilon|H|$ or $|gH \setminus A| < \epsilon|H|$.*

Moreover, H is in the Boolean algebra generated by $\{gAh : g, h \in G\}$.

Proof. Note that condition (ii) follows immediately from (i). So suppose condition (i) is false. Then we have some fixed $k, r \geq 1$ and $\epsilon > 0$ such that, for all $i \in \mathbb{N}$, there is a finite group G_i of exponent r , which is a counterexample. Specifically, there is a k -NIP subset $A_i \subseteq G_i$ such that, if $H \leq G_i$ is normal with index at most i , and $D, Z \subseteq G_i$ are unions of cosets of H with $|Z| < \epsilon|G_i|$, then $|(A_i \setminus Z) \triangle D| > \epsilon|H|$.

Let $\mathcal{L}$ be the group language with a new predicate A , and consider (G_i, A_i) as a finite $\mathcal{L}$ -structure. Let $\mathcal{U}$ be a nonprincipal ultrafilter on $\mathbb{Z}^+$, and let G be a sufficiently saturated elementary extension of $M := \prod_{\mathcal{U}} (G_i, A_i)$. Let $\theta(x; y)$ be the formula $A(y \cdot x)$. Note that $\theta(x; y)$ is invariant. Moreover, by Łoś's Theorem and since $M \prec G$, $\theta(x; y)$ is k -NIP (in G) and G has exponent r . So $G/G_{\theta^r}^{00}$ is a compact torsion group, and thus is profinite

by Fact 2.3(d). By Theorem 3.1, there is a θ^r -definable finite-index normal subgroup $H \leq G$ and sets $D, Z \subseteq G$, which are unions of cosets of H , such that $\mu(Z) < \epsilon$ and $\mu((A \setminus Z) \triangle D) = 0$.

Let $n = [G : H]$, and fix $\langle \theta^r \rangle$ -formulas $\phi(x; \bar{y})$, $\psi(x; \bar{z})$, and $\zeta(x; \bar{u})$ such that H , D , and Z are defined by instances of $\phi(x; \bar{y})$, $\psi(x; \bar{z})$, and $\zeta(x; \bar{u})$, respectively. Given $i \in \mathbb{Z}^+$, let μ_i be the normalized counting measure on G_i . Let I be the set of $i \in \mathbb{Z}^+$ such that, for some tuples $\bar{a}_i, \bar{b}_i$, and $\bar{c}_i$ from G_i ,

- (i) $\phi(x; \bar{a}_i)$ defines a normal subgroup H_i of G_i of index n ,
- (ii) $\psi(x; \bar{b}_i)$ and $\zeta(x; \bar{c}_i)$ define sets $D_i, Z_i \subseteq G_i$, respectively, which are each unions of cosets of H_i , and
- (iii) $\mu_i(Z_i) < \epsilon$ and $\mu_i((A_i \setminus Z_i) \triangle D_i) < \epsilon/n$.

Then $I \in \mathcal{U}$ by Łoś's Theorem and since $M \prec G$. So there is some $i \in I$ such that $i \geq n$, which contradicts the choice of (G_i, A_i) . ■

The statement of the previous result is almost identical to our result from [14] on *stable* arithmetic regularity in arbitrary finite groups (Theorem 1.1 above), except for the presence of the error set Z . As in [14, Corollary 3.5], we can use this result to deduce a very strong *graph regularity* statement for bipartite graphs defined by NIP subsets of finite groups of uniformly bounded exponent.

Let $\Gamma = (V, W; E)$ be a finite bipartite graph and fix nonempty sets $X \subseteq V$ and $Y \subseteq W$. The **edge density** of the pair (X, Y) is $\delta_\Gamma(X, Y) := |X \times Y \cap E|/|X \times Y|$. We say that (X, Y) is **ϵ -regular** if $|\delta_\Gamma(X, Y) - \delta_\Gamma(X_0, Y_0)| \leq \epsilon$ for any $X_0 \subseteq X$ and $Y_0 \subseteq Y$ such that $|X_0| \geq \epsilon|X|$ and $|Y_0| \geq \epsilon|Y|$. Given vertices $v \in V$ and $w \in W$, define $\deg_\Gamma(v, Y) = |\{y \in Y : E(v, y)\}|$ and $\deg_\Gamma(X, w) = |\{x \in X : E(x, w)\}|$. Following [14], we say that the pair (X, Y) is **uniformly ϵ -good for Γ** , where $\epsilon > 0$, if $|X| = |Y|$ and either

- (i) for any $x \in X$ and $y \in Y$, $\deg_\Gamma(x, Y) = \deg_\Gamma(X, y) \leq \epsilon|X|$, or
- (ii) for any $x \in X$ and $y \in Y$, $\deg_\Gamma(x, Y) = \deg_\Gamma(X, y) \geq (1 - \epsilon)|X|$.

One can show that if (X, Y) is uniformly ϵ^2 -good then it is ϵ -regular, and either $\delta_\Gamma(X, Y) \leq \epsilon$ or $\delta_\Gamma(X, Y) \geq 1 - \epsilon$. In fact, a stronger property holds: if $X_0 \subseteq X$ and $Y_0 \subseteq Y$ are nonempty and *either* $|X_0| \geq \epsilon|X|$ *or* $|Y_0| \geq \epsilon|Y|$, then $\delta_\Gamma(X_0, Y_0) \leq \epsilon$ or $\delta_\Gamma(X_0, Y_0) \geq 1 - \epsilon$ (see [14, Proposition 3.4]).

Now suppose G is a finite group and A is a subset of G . Let $\Gamma_G(A) = (G, G; E)$ where $E = \{(x, y) \in G^2 : yx \in A\}$. Given $X \subseteq G$ and $g \in G$, note that $\deg_{\Gamma_G(A)}(g, X) = |A \cap Xg|$ and $\deg_{\Gamma_G(A)}(X, g) = |A \cap gX|$. We now observe that Theorem 3.2 implies a graph regularity statement for NIP subsets of finite groups of uniformly bounded exponent, in which the partition is given by cosets of a normal subgroup and almost all pairs are uniformly good (and thus regular up to a change in ϵ). Given a group G , a normal subgroup $H \leq G$, and $C, D \in G/H$, let $C \cdot D$ denote the product of C and D in the quotient group G/H .

Corollary 3.3. *For any $k, r \geq 1$ and $\epsilon > 0$ there is $n(k, r, \epsilon)$ such that the following holds. Suppose G is a finite group of exponent r and $A \subseteq G$ is k -NIP. Then there is a normal subgroup H of index $n \leq n(k, r, \epsilon)$, and a set $\Sigma \subseteq (G/H)^2$, with $|\Sigma| \leq \epsilon n^2$, such that any $(C, D) \notin \Sigma$ is uniformly ϵ -good for $\Gamma_G(A)$.*

Proof. Fix $k, r \geq 1$ and $\epsilon > 0$ and let $n(k, r, \epsilon)$ be as in Theorem 3.2. Fix a finite group G and a k -NIP set $A \subseteq G$. Then there is a normal subgroup $H \leq G$ of index $n \leq n(k, r, \epsilon)$, and a set $\mathcal{I} \subseteq G/H$ with $|\mathcal{I}| \leq \epsilon n$, such that for any $C \notin \mathcal{I}$, either $|C \cap A| \leq \epsilon |H|$ or $|C \cap A| \geq (1 - \epsilon)|H|$. Let $\Sigma = \{(C, D) \in (G/H)^2 : C \cdot D \in \mathcal{I}\}$. Then $\Sigma = \bigcup_{C \in G/H} \{(C, C^{-1} \cdot D) : D \in \mathcal{I}\}$, and so $|\Sigma| \leq \epsilon n^2$. Finally, if $(C, D) \notin \Sigma$, then (C, D) is uniformly ϵ -good for $\Gamma_G(A)$ (this is identical to the calculation in the proof of [14, Corollary 3.5], and makes crucial use of normality of H). ■

Remark 3.4. In Theorem 3.2, the assumption of uniformly bounded exponent was used to obtain a certain profinite quotient, and so it is worth reviewing this argument from a more general perspective. Specifically, fix $k \geq 1$ and consider the following property of a class $\mathcal{G}$ of finite groups:

$(*)_k$ For any sequences $(G_i)_{i=0}^\infty$ and $(A_i)_{i=0}^\infty$, where $G_i \in \mathcal{G}$ and $A_i \subseteq G_i$ is k -NIP, and for any ultrafilter $\mathcal{U}$ on $\mathbb{N}$, if G is a sufficiently saturated elementary extension of $\prod_{\mathcal{U}} (G_i, A_i)$, then $G/G_{\theta r}^{00}$ is profinite, where $\theta(x; y) := A(y \cdot x)$.

Then, for any $\mathcal{G}$ satisfying $(*)_k$ and any $\epsilon > 0$, there is some $n = n(k, \epsilon, \mathcal{G})$ such that any group $G \in \mathcal{G}$ and k -NIP set $A \subseteq G$ satisfy the conclusions of Theorem 3.2 using n . Indeed, Theorem 3.2 only uses that for any $k, r \geq 1$, the class $\mathcal{G}_r$ of finite groups of exponent r satisfies $(*)_k$, for the rather heavy-handed reason that compact torsion groups are profinite.

Profinite quotients also arise in the stable setting (recall that $A \subseteq G$ is k -stable if $\Gamma_G(A)$ omits $([k], [k]; \leq)$ as an induced subgraph). In fact, if G is pseudofinite and saturated, and $\theta(x; \bar{y})$ is a stable invariant formula, then the group $G/G_{\theta r}^{00}$ is finite (see [15, Corollary 3.17]). Therefore, in this case, the set E in Theorem 2.14(d) is empty since it is a Haar null set in a finite group. So, if one replaces “NIP” with “stable” in Lemma 2.16 and Theorem 3.1, then one can choose the error set $Z = \emptyset$, and a similar ultraproduct argument as in Theorem 3.2 yields Theorem 1.1 (structure and regularity for stable subsets of finite groups).¹⁴

4. Bohr neighborhoods

In this section, we recall some basic definitions and facts concerning Bohr neighborhoods, and define an approximate version of Bohr neighborhoods, which we will need for later arguments involving ultraproducts.

¹⁴It is worth noting that this explanation of Theorem 1.1 is not a faster proof than what is done in [14]. In particular, [15, Corollary 3.17] relies on the same results from [30] used in [14] to directly prove Theorem 1.1. Also, pseudofiniteness is not needed to prove that $G/G_{\theta r}^{00}$ is finite.

Given a group G , 1_G denotes the identity (if G is abelian we use 0_G). We say that the pair (L, d) is a **compact metric group** if L is a compact metrizable group and d is a *bi-invariant* metric on L compatible with the topology. By Fact 2.3(e), if L is any compact second-countable group then there is a (not necessarily unique) bi-invariant metric d on L such that (L, d) is a compact metric group.

Definition 4.1. Let H be a group and let (L, d) be a compact metric group. Given some $\delta > 0$ and a homomorphism $\tau: H \rightarrow L$, define

$$B_{\tau, \delta}^L = \{x \in H : d(\tau(x), 1_L) < \delta\}.$$

A set $B \subseteq H$ is a (δ, L) -**Bohr neighborhood in H** if $B = B_{\tau, \delta}^L$ for some homomorphism $\tau: H \rightarrow L$.

Our ultimate goal is to transfer Bohr neighborhoods in pseudofinite groups to Bohr neighborhoods in finite groups. To do this, we will need to approximate Bohr neighborhoods by definable objects. This necessitates an approximate notion of a Bohr neighborhood, which involves approximate homomorphisms of groups.

Definition 4.2. Let H be a group and let (L, d) be a compact metric group.

- (1) Given $\delta > 0$, a function $f: H \rightarrow L$ is a δ -**homomorphism** if $f(1_H) = 1_L$ and, for all $x, y \in H$, $d(f(xy), f(x)f(y)) < \delta$.
- (2) Given $\epsilon, \delta > 0$, a set $Y \subseteq H$ is a δ -**approximate (ϵ, L) -Bohr neighborhood in H** if $Y = \{x \in H : d(f(x), 1_L) < \epsilon\}$ for some δ -homomorphism $f: H \rightarrow L$.

Approximate homomorphisms have been studied extensively in the literature, with a special focus on the question of when an approximate homomorphism is “close” to an actual homomorphism. For our purposes, this is what is needed to replace approximate Bohr neighborhoods with actual Bohr neighborhoods. More precisely, we will start with a definable approximate Bohr neighborhood in a pseudofinite group, and transfer it to find an approximate Bohr neighborhood in a finite group. At this point, we will be working with an approximate homomorphism from a finite group to a compact Lie group, which is a setting where one can always find a genuine Bohr neighborhood inside an approximate Bohr neighborhood, with a negligible loss in size.

Theorem 4.3 (Alekseev, Glebskii, & Gordon [1, Theorem 5.13]). *Let (L, d) be a compact metric Lie group. Then there is an $\alpha_L > 0$ such that, for any $0 < \delta < \alpha_L$, if H is a compact group and $f: H \rightarrow L$ is a δ -homomorphism, then there is a homomorphism $\tau: H \rightarrow L$ such that $d(f(x), \tau(x)) < 2\delta$ for all $x \in H$.*

An easy consequence is that in the setting of compact Lie groups, Bohr neighborhoods can be found inside approximate Bohr neighborhoods. Given a compact metric group (L, d) and $n \geq 1$, note that we have a compact metric group (L^n, d^n) , where L^n is endowed with the product topology and $d^n(\bar{x}, \bar{y}) = \max_{1 \leq i \leq n} d(x_i, y_i)$.

Corollary 4.4. *Let (L, d) be a compact metric Lie group. Then there is an $\alpha_L > 0$ such that, if H is a compact group, $n \in \mathbb{N}$, and $0 < \delta < \alpha_L$, then any δ -approximate $(3\delta, L^n)$ -Bohr neighborhood in H contains a (δ, L^n) -Bohr neighborhood in H .*

Proof. Fix $\alpha_L > 0$ from Theorem 4.3. Suppose H is a compact group, and $Y \subseteq H$ is a δ -approximate $(3\delta, L^n)$ -Bohr neighborhood in H , for some $n \in \mathbb{N}$ and $0 < \delta < \alpha_L$, witnessed by a δ -homomorphism $f: H \rightarrow L^n$. We may assume $n \geq 1$. For $1 \leq i \leq n$, let $f_i: H \rightarrow L$ be given by $f_i(x) = f(x)_i$. Then each f_i is a δ -homomorphism. Given $1 \leq i \leq n$, Theorem 4.3 provides a homomorphism $\tau_i: H \rightarrow L$ such that $d(f_i(x), \tau_i(x)) < 2\delta$ for all $x \in H$. Let $\tau: H \rightarrow L^n$ be such that $\tau(x) = (\tau_1(x), \dots, \tau_n(x))$. Then τ is a homomorphism and $d^n(f(x), \tau(x)) < 2\delta$ for all $x \in H$. Now we have $B_{\tau, \delta}^{L^n} \subseteq Y$ by the triangle inequality. ■

The next result provides a lower bound on the size of Bohr neighborhoods in finite groups. The proof is a standard averaging argument (adapted from the abelian case; see [51, Lemma 4.20] and/or [24, Lemma 4.1]). We include the details for the sake of clarity and to observe that the method works for Bohr neighborhoods in nonabelian finite groups defined using compact metric groups.

Proposition 4.5. *Let (L, d) be a compact metric group and, given $\delta > 0$, let $\ell_\delta = \eta_L(\{t \in L : d(t, 1_L) < \delta\})$. For any finite group H and $\delta > 0$, if $B \subseteq H$ is a $(2\delta, L)$ -Bohr neighborhood in H , then $|B| \geq \ell_\delta |H|$.*

Proof. Fix a finite group H , a homomorphism $\tau: H \rightarrow L$, and some $\delta > 0$. Given $x \in H$, let $f_x: L \rightarrow \{0, 1\}$ be the characteristic function of $\{t \in L : d(\tau(x), t) < \delta\}$. Then

$$\ell_\delta |H| = \sum_{x \in H} \int_L f_x d\eta_L = \int_L \sum_{x \in H} f_x d\eta_L.$$

So there must be some $t \in L$ such that $\sum_{x \in H} f_x(t) \geq \ell_\delta |H|$. In other words, if $S = \{x \in H : d(\tau(x), t) < \delta\}$ then $|S| \geq \ell_\delta |H|$. Fix $a \in S$. For any $x \in S$, we have

$$d(\tau(xa^{-1}), 1_L) = d(\tau(x), \tau(a)) \leq d(\tau(x), t) + d(\tau(a), t) < 2\delta.$$

Therefore $Sa^{-1} \subseteq B_{\tau, 2\delta}^L$, and so $|B_{\tau, 2\delta}^L| \geq |Sa^{-1}| = |S| \geq \ell_\delta |H|$. ■

Remark 4.6. As we have seen in previous results (e.g., Lemma 2.9), real tori of the form $\mathbb{T}^r$ have a distinguished role in the study of compact abelian groups. Moreover, in the setting of finite abelian groups, Bohr neighborhoods are usually defined using homomorphisms to the torus (see, e.g., [7], [24]). Thus, in order to match these definitions more explicitly, we define the metric $d_{\mathbb{T}^1}(x, y) = \min\{|x - y|, 1 - |x - y|\}$ on $\mathbb{T}^1$ (identified with $[0, 1)$) and the product metric $d_{\mathbb{T}^r} := d_{\mathbb{T}^1}^r$ on $\mathbb{T}^r$ for $r \in \mathbb{N}$. Throughout the rest of the paper, when we speak of $\mathbb{T}^r$ as a compact metric group, will always work with this choice of metric.

Definition 4.7. Given a finite group H , a homomorphism $\tau: H \rightarrow \mathbb{T}^r$, and some $\delta > 0$, we let $B_{\tau, \delta}^r$ denote $B_{\tau, \delta}^{\mathbb{T}^r}$. We call $B \subseteq H$ a (δ, r) -Bohr neighborhood in H if $B = B_{\tau, \delta}^r$ for some τ .¹⁵

¹⁵In this case, r and δ are sometimes referred to as the *rank* and *width* of B , respectively.

In Theorem 3.2, we showed that NIP sets in finite groups of uniformly bounded exponent are approximated by normal subgroups of uniformly bounded index. As noted in the introduction, we cannot expect such a result for NIP sets in finite groups G of unrestricted exponent. So instead of subgroups of G , we will consider pairs (B, H) , where H is a normal subgroup of G and B is a (δ, r) -Bohr neighborhood in H . Thus, in Proposition 4.9 below, we point out some ways in which Bohr neighborhoods behave like normal subgroups of “small” index. We will need the following minor generalization of a well-known exercise, namely, if G is an amenable group and $A \subseteq G$ has positive upper density then AA^{-1} is generic.

Proposition 4.8. *Suppose G is a group, $\mathcal{B}$ is a left-invariant Boolean algebra of subsets of G , and ν is a left-invariant finitely additive probability measure on $\mathcal{B}$. Suppose $A \in \mathcal{B}$ is such that $\nu(A) > 0$. Then, for any $X \subseteq G$, there is a finite set $F \subseteq X$ such that $|F| \leq 1/\nu(A)$ and $X \subseteq FAA^{-1}$.*

Proof. We say that $Y \subseteq G$ separates A if $xA \cap yA = \emptyset$ for all distinct $x, y \in Y$. By the assumptions on ν , if $Y \subseteq G$ separates A then $|Y| \leq 1/\nu(A)$. Choose a finite set $F \subseteq X$ with maximal size among subsets of X that separate A . Fix $x \in X$. Then there is $y \in F$ such that $xA \cap yA \neq \emptyset$, and so we may fix $z \in xA \cap yA$. Then $y^{-1}z \in A$ and $z^{-1}x \in A^{-1}$, which means $y^{-1}x \in AA^{-1}$, and so $x \in FAA^{-1}$. ■

Proposition 4.9. *Let G be a finite group and $H \leq G$ be a normal subgroup of index n . Suppose B is a (δ, r) -Bohr neighborhood in H , where $r \in \mathbb{N}$ and $0 < \delta \leq 2$.*

- (a) $B = B^{-1}$, $1_G \in B$, and $gB = Bg$ for any $g \in G$.
- (b) *For any $X \subseteq G$ there is $F \subseteq X$ such that $|F| \leq n(2/\delta)^r$ and $X \subseteq FB$. Thus G is covered by at most $n(2/\delta)^r$ translates of B .*

Proof. For (a), we have $B = B^{-1}$ by bi-invariance of $d_{\mathbb{T}^r}$, and clearly $1_G \in B$. If $g \in G$ and $x \in B$ then $gxg^{-1} \in H$ (since H is normal), and so $gxg^{-1} \in B$ by bi-invariance of $d_{\mathbb{T}^r}$. It follows that $gB = Bg$ for any $g \in G$.

To prove (b), suppose $B = B_{\tau, \delta}^r$ for some $\tau: H \rightarrow \mathbb{T}^r$, and let $B_0 = B_{\tau, \delta/2}^r$. Note that $\eta_{\mathbb{T}^r}(\{t \in \mathbb{T}^r : d_{\mathbb{T}^r}(t, 0_{\mathbb{T}^r}) < \delta/4\}) = (\delta/2)^r$, and so $|B_0| \geq (\delta/2)^r |H| = n^{-1}(\delta/2)^r |G|$ by Proposition 4.5. Now fix $X \subseteq G$. By Proposition 4.8 (with ν the normalized counting measure on G), we have $X \subseteq FB_0B_0^{-1}$ for some $F \subseteq X$ with $|F| \leq n(2/\delta)^r$. Finally, note that if $x, y \in B_0$ then

$$d(\tau(xy^{-1}), 0_{\mathbb{T}^r}) = d(\tau(x), \tau(y)) \leq d(\tau(x), 0_{\mathbb{T}^r}) + d(\tau(y), 0_{\mathbb{T}^r}) < \delta.$$

So $B_0B_0^{-1} \subseteq B$, and we have $X \subseteq FB$. ■

Note that part (a) of the previous fact holds for any compact metric group (L, d) in place of $\mathbb{T}^r$ (and does not use that G is finite); and the analogue of part (b) holds for any (L, d) with $(\delta/2)^r$ replaced by $\ell_{\delta/4}$ (from Proposition 4.5). We refer the reader to [51, Section 4.4] and [24, Sections 3 & 4] for more on the role of Bohr neighborhoods in arithmetic combinatorics and discrete Fourier analysis.

5. Structure and regularity: the general case

The next goal is a result analogous to Theorem 3.1, but without the assumption that $G/G_{\theta^r}^{00}$ is profinite. For this, we need to understand more about descriptions of $G_{\theta^r}^{00}$ as an intersection of definable subsets of G . The goal is to find properties of definable sets which are both interesting algebraically, and also sufficiently first-order so that they can be transferred to finite groups in arguments with ultraproducts. In particular, we will use approximate Bohr neighborhoods.

Suppose G is a group definable in a sufficiently saturated structure, and $\Gamma \leq G$ is a type-definable normal subgroup of bounded index. By Lemma 2.9, Γ is an intersection of a bounded number of definable finite-index normal subgroups of G whenever G/Γ is profinite. The next result shows that in general, we can write $\Gamma = \bigcap_{i \in I} W_i$ where each W_i is a definable subset of a definable finite-index normal subgroup $H_i \leq G$, and there is a Bohr neighborhood B_i in H_i such that $\Gamma \subseteq B_i \subseteq W_i$. Moreover, B_i is obtained from a definable homomorphism to a compact connected Lie group (so, in particular, B_i is *co-type-definable*).

Proposition 5.1. *Let G be a group definable in a sufficiently saturated structure M^* . Suppose $\Gamma \leq G$ is type-definable and normal of bounded index. Then there is a bounded *jsl* I and a decreasing net $(W_i)_{i \in I}$ of definable subsets of G such that $\Gamma = \bigcap_{i \in I} W_i$ and, for all $i \in I$, there are*

- * a definable finite-index normal subgroup $H_i \leq G$,
- * a definable homomorphism $\pi_i: H_i \rightarrow L_i$, where (L_i, d_i) is a compact connected metric Lie group, and
- * a real number $\delta_i > 0$,

such that $\Gamma \subseteq \ker \pi_i \subseteq B_{\pi_i, \delta_i}^{L_i} \subseteq W_i \subseteq H_i$. Moreover:

- (a) If Γ is countably-definable then we may assume $I = \mathbb{N}$.
- (b) If $G = M^*$ and Γ is θ -type-definable for some invariant formula $\theta(x; \bar{y})$, then we may assume W_i , H_i , and π_i are θ -definable.
- (c) If $G = M^*$ and G is pseudofinite, then we may assume $L_i = \mathbb{T}^{n_i}$ for some $n_i \in \mathbb{N}$.
- (d) If G/Γ is abelian then we may assume $H_i = G$ and $L_i = \mathbb{T}^{n_i}$ for some $n_i \in \mathbb{N}$.

Proof. Let $(\Gamma_i)_{i \in I_0}$ and $(H_i)_{i \in I_0}$ be as in Lemma 2.9, where I_0 is a small *jsl*. For each $i \in I_0$, let $L_i = H_i/\Gamma_i$ and equip L_i with some bi-invariant metric d_i (by Fact 2.3(e)). Let $\pi_i: H_i \rightarrow L_i$ be the canonical homomorphism. Then π_i is definable since $G \rightarrow G/\Gamma_i$ is definable.

For each $i \in I_0$, let $(W_n^i)_{n=0}^\infty$ be a decreasing sequence of definable subsets of G such that $\Gamma_i = \bigcap_{n=0}^\infty W_n^i$. Let I be the set of all finite subsets of $I_0 \times \mathbb{N}$, and view I as a *jsl* under the subset ordering. Given $\sigma = \{(i_1, n_1), \dots, (i_k, n_k)\} \in I$, let $i_\sigma = \sup \{i_1, \dots, i_k\}$, and set $H_\sigma := H_{i_\sigma}$, $W_\sigma := H_\sigma \cap \bigcap_{t=1}^k W_{n_t}^{i_t}$, $\Gamma_\sigma := \Gamma_{i_\sigma}$, $L_\sigma := L_{i_\sigma}$, and $\pi_\sigma := \pi_{i_\sigma}$. Note that $\Gamma_\sigma \subseteq W_\sigma \subseteq H_\sigma$ for all $\sigma \in I$, and $\bigcap_{\sigma \in I} W_\sigma = \Gamma$. By choice of i_σ , we also see that $(W_\sigma)_{\sigma \in I}$ is decreasing.

Now, given $\sigma \in I$, the set $U_\sigma = \{a\Gamma_\sigma \in G/\Gamma_\sigma : a\Gamma_\sigma \subseteq W_\sigma\}$ is an identity neighborhood in L_σ (by Fact 2.4(b)) and, by construction, $\pi_\sigma^{-1}(U_\sigma) \subseteq W_\sigma$. So choose $\delta_\sigma > 0$ such that U_σ contains the open ball of radius δ_σ around 1_{L_σ} . Note that $\Gamma \subseteq \Gamma_\sigma = \ker \pi_\sigma \subseteq B_{\pi_\sigma, \delta_\sigma}^{L_\sigma} \subseteq W_\sigma$. This finishes the proof of the main statement.

We now deal with the remaining claims. Claims (b) and (c) follow by applying Lemma 2.9(b, d) in the above construction. For claim (a), suppose Γ is countably-definable. Then, by Lemma 2.9(a), we may assume $I_0 = \mathbb{N}$. So I is the *jsl* of finite subsets of $\mathbb{N} \times \mathbb{N}$ under the subset ordering, which contains a cofinal sub-*jsl* isomorphic to $\mathbb{N}$.

Finally, for claim (d), suppose G/Γ is abelian. Then G/Γ_i is a compact abelian Lie group for any $i \in I_0$, and thus isomorphically embeds in $\mathbb{T}^{n_i}$ for some $n_i \in \mathbb{N}$ by Fact 2.3(g). So in the argument above, we can replace each H_i and π_i with G and $G \rightarrow G/\Gamma_i \subseteq \mathbb{T}^{n_i}$, respectively. ■

Note that in the previous result, the homomorphism $\pi_i: H_i \rightarrow L_i$ is also surjective, except in part (d) where we can replace H_i with G when G/Γ is abelian. The purpose of part (d) is to note that if G is already abelian then one can obtain Bohr neighborhoods defined by tori without first passing to a subgroup H_i , and without Theorem 2.7 or an extra pseudofiniteness assumption.

One drawback of Proposition 5.1 is that the Bohr neighborhood $B_{\pi_i, \delta_i}^{L_i}$ is not necessarily definable. In order to work with definable objects, we will have to consider approximate Bohr neighborhoods.

Definition 5.2. Let G be a group definable in a sufficiently saturated structure M^* . Suppose $H \leq G$ is definable and $\pi: H \rightarrow L$ is a definable homomorphism to a compact metric group (L, d) . Given $t \geq 1$, we say that a sequence $(Y_m)_{m=0}^\infty$ of subsets of H is a **definable (t, π) -approximate Bohr chain in H** if $(Y_m)_{m=0}^\infty$ is decreasing, $\ker \pi = \bigcap_{i=0}^\infty Y_m$, and there are $(\delta_m)_{m=0}^\infty$ and $(f_m)_{m=0}^\infty$ such that

- (i) $(\delta_m)_{m=0}^\infty$ is a decreasing sequence of positive real numbers converging to 0,
- (ii) for all m , $f_m: H \rightarrow L$ is a definable δ_m -homomorphism with finite image, and
- (iii) for all m , $Y_m = \{x \in H : d(f_m(x), 1_L) < t\delta_m\}$.

Moreover, if $G = M^*$ and H , π , and f_m are all θ -definable, for some fixed formula $\theta(x; \bar{y})$, then we say “ θ -definable” in place of “definable”.

Proposition 5.3. Suppose $(Y_m)_{m=0}^\infty$ is a definable (respectively, θ -definable) (t, π) -approximate Bohr chain in $H \leq G$, as in Definition 5.2. Then each set Y_m is a δ_m -approximate $(t\delta_m, L)$ -Bohr neighborhood in H , and a definable (respectively, θ -definable) subset of G .¹⁶

Proof. The first claim is obvious. For definability, note that for any $m \in \mathbb{N}$, H is partitioned into finitely many fibers $f_m^{-1}(\lambda)$ for $\lambda \in f_m(H)$, each of which is definable

¹⁶On the other hand, we are *not* claiming that the family $\{Y_m : m \in \mathbb{N}\}$ is *uniformly* definable (which could be misconstrued from our choice of terminology).

(respectively, θ -definable) by Remark 2.6. Now Y_m is a union of the (finitely many) fibers $f_m^{-1}(\lambda)$ where $\lambda \in f(H)$ is such that $d(\lambda, 1_L) < \epsilon$. ■

The parameter t in Definition 5.2 is introduced in order to control the “width” ϵ and the “error” δ in a δ -approximate (ϵ, L) -Bohr neighborhood. Specifically, it is desirable to have ϵ be some constant multiple t of δ , and in the following results we will choose t arbitrarily. This will eventually be used to find actual Bohr neighborhoods inside approximate Bohr neighborhoods, with $L = \mathbb{T}^n$ for some $n \in \mathbb{N}$, in which case setting $t = 3$ will suffice (via Corollary 4.4).

Lemma 5.4. *Let G be a group definable in a sufficiently saturated structure M^* . Suppose H is a definable subgroup of G and $\pi: H \rightarrow L$ is a definable homomorphism to a compact metric group (L, d) . Then, for any $t \geq 1$, there is a definable (t, π) -approximate Bohr chain $(Y_m)_{m=0}^\infty$ in H . Moreover, if $G = M^*$ and H and π are θ -definable for some invariant formula $\theta(x; \bar{y})$, then $(Y_m)_{m=0}^\infty$ is a θ -definable (t, π) -approximate Bohr chain.*

Proof. Given $\lambda \in L$ and $\epsilon > 0$, let $K(\lambda, \epsilon) \subseteq L$ and $U(\lambda, \epsilon) \subseteq L$ be the closed ball of radius ϵ around λ and the open ball of radius ϵ around λ , respectively.

Fix $m \geq 1$. Choose a finite set $\Lambda \subseteq L$ such that $L = \bigcup_{\lambda \in \Lambda} K(\lambda, 1/(2m))$ and $1_L \in \Lambda$. For any $\lambda \in \Lambda$, since π is definable, there is a definable set $D_\lambda \subseteq H$ such that $\pi^{-1}(K(\lambda, 1/(2m))) \subseteq D_\lambda \subseteq \pi^{-1}(U(\lambda, 1/m))$ (see Remark 2.6(a)). Enumerate $\Lambda = \{\lambda_1, \dots, \lambda_k\}$, with $\lambda_1 = 1_L$, and, for $1 \leq i \leq k$, let $D_i = D_{\lambda_i}$. For $1 \leq i \leq k$, define $E_i = D_i \setminus \bigcup_{j < i} D_j$. Then $E_1, \dots, E_k$ are definable and partition H . This determines a definable function $f_m: H \rightarrow \Lambda$ such that $f_m(x) = \lambda_i$ if and only if $x \in E_i$. For any $x \in H$, we have $x \in D_{f_m(x)} \subseteq \pi^{-1}(U(f_m(x), 1/m))$, and so $d(\pi(x), f_m(x)) < 1/m$. Note that $f_m(1_H) = 1_L$ by definition. Also, given $x, y \in H$, we have

$$\begin{aligned} d(f_m(xy), f_m(x)f_m(y)) &\leq d(f_m(xy), \pi(xy)) + d(\pi(x)\pi(y), f_m(x)\pi(y)) + d(f_m(x)\pi(y), f_m(x)f_m(y)) \\ &= d(f_m(xy), \pi(xy)) + d(\pi(x), f_m(x)) + d(\pi(y), f_m(y)) < 3/m. \end{aligned}$$

Altogether, $f_m: H \rightarrow L$ is a $3/m$ -homomorphism.

Now fix an integer $t \geq 1$. For $m \in \mathbb{N}$, define

$$Y_m = \{x \in H : d(f_m(x), 1_L) < 3t/m\}.$$

Note that $D_1 \subseteq Y_m$, and so $\ker \pi \subseteq \pi^{-1}(K(1_L, 1/(2m))) \subseteq Y_m$. We now have a sequence $(Y_m)_{m=1}^\infty$ of definable subsets of H with $\ker \pi \subseteq Y_m$ for all $m \in \mathbb{N}$. Moreover, for any $m \in \mathbb{N}$, if $x \in Y_m$ then

$$d(\pi(x), 1_L) \leq d(\pi(x), f_m(x)) + d(f_m(x), 1_L) < (3t + 1)/m.$$

This implies $\ker \pi = \bigcap_{m=0}^\infty Y_m$. Finally, given $m \in \mathbb{N}$, we have $\pi^{-1}(K(1_L, 1/(2m))) \subseteq Y_m \subseteq \pi^{-1}(U(1_L, (3t + 1)/m))$. In particular, if $n \geq (6t + 2)m$, then $Y_n \subseteq Y_m$. So, after thinning the sequence, we may assume $Y_{m+1} \subseteq Y_m$ for all $m \in \mathbb{N}$. Altogether,

if $\delta_m = 3/m$, then $(\delta_m)_{m=1}^\infty$ and $(f_m)_{m=1}^\infty$ witness that $(Y_m)_{m=1}^\infty$ is a definable (t, π) -approximate Bohr chain in H .

For the “moreover” statement, suppose $G = M^*$, and H and π are θ -definable for some formula $\theta(x; \bar{y})$. Then one can choose each D_λ to be θ -definable (by Remark 2.6). So each f_m is θ -definable by construction, and $(Y_m)_{m=1}^\infty$ is a θ -definable (t, π) -approximate Bohr chain. ■

We now combine the above ingredients to prove a structure and regularity theorem for θ^r -definable sets in a sufficiently saturated pseudofinite group G , where $\theta(x; \bar{y})$ is an arbitrary invariant NIP formula. Recall that we use μ for the pseudofinite counting measure.

Theorem 5.5. *Let G be a sufficiently saturated pseudofinite expansion of a group, and suppose $\theta(x; \bar{y})$ is an invariant NIP formula. Fix a θ^r -definable set $A \subseteq G$ and some $\epsilon > 0$. Then there are*

- * a θ^r -definable finite-index normal subgroup $H \leq G$,
- * a θ^r -definable homomorphism $\pi: H \rightarrow \mathbb{T}^n$, for some $n \in \mathbb{N}$, and
- * a θ^r -definable set $Z \subseteq G$, with $\mu(Z) < \epsilon$,

such that, for any integer $t \geq 1$, there is

- * a θ^r -definable (t, π) -approximate Bohr chain $(Y_m)_{m=0}^\infty$ in H

satisfying the following properties, for any $m \in \mathbb{N}$:

- (i) (structure) *There is a set $D_m \subseteq G$, which is a union of finitely many left translates of Y_m , such that $\mu((A \triangle D_m) \setminus Z) = 0$.*
- (ii) (regularity) *For any $g \in G \setminus Z$, either $\mu(gY_m \cap A) = 0$ or $\mu(gY_m \setminus A) = 0$.*

Moreover, if $G/G_{\theta^r}^{00}$ is abelian then we may assume $H = G$.

Proof. Let $(W_i)_{i=0}^\infty$ be a sequence of θ^r -definable sets in G satisfying the conditions of Proposition 5.1 with $\Gamma = G_{\theta^r}^{00}$. By Lemma 2.16, there is a θ^r -definable set $Z \subseteq G$ and some $i \in \mathbb{N}$ such that $\mu(Z) < \epsilon$ and, if $W := W_i$, then for all $g \in G \setminus Z$, we have $\mu(gW \cap A) = 0$ or $\mu(gW \setminus A) = 0$. Proposition 5.1 associates to W a θ^r -definable homomorphism $\pi: H \rightarrow \mathbb{T}^n$, where H is a θ^r -definable finite-index normal subgroup of G . If $G/G_{\theta^r}^{00}$ is abelian then we may further assume $H = G$. Fix $t \geq 1$. By Lemma 5.4, there is a θ^r -definable (t, π) -approximate Bohr chain $(Y_m)_{m=0}^\infty$ in H . Recall that each Y_m is definable by Proposition 5.3. Since $\ker \pi$ is type-definable and contained in the definable set W , it follows from saturation that $Y_m \subseteq W$ for sufficiently large m . So for sufficiently large m we have, for any $g \in G \setminus Z$, either $\mu(gY_m \cap A) = 0$ or $\mu(gY_m \setminus A) = 0$. Thus, after removing finitely many sets Y_m from the sequence, we have condition (ii).

Toward proving condition (i), fix $m \in \mathbb{N}$. Since $\ker \pi$ is a subgroup of H , we may use saturation (much as in the proof of Lemma 2.16) to find some $r \geq m$ such that $Y_r Y_r^{-1} \subseteq Y_m$. Since Y_r contains a type-definable bounded-index subgroup of G (namely, $\ker \pi$), it follows that Y_r is generic and so $\mu(Y_r) > 0$. By Proposition 4.8, there is a finite set $F \subseteq G \setminus Z$ such that $G \setminus Z \subseteq F Y_m$. Let I be the set of $g \in F$ such that $\mu(gY_m \setminus A) = 0$,

and note that if $g \in F \setminus I$ then $\mu(gY_m \cap A) = 0$. Let $D_m = IY_m$. Since $G \setminus Z \subseteq FY_m$, we have

$$A \triangle D_m \subseteq Z \cup \bigcup_{g \in I} (gY_m \setminus A) \cup \bigcup_{g \in F \setminus I} (gY_m \cap A),$$

and so $\mu((A \triangle D_m) \setminus Z) = 0$. ■

The next goal is our main result for NIP sets in arbitrary finite groups (see Theorem 5.7). Roughly speaking, we will show that if A is a k -NIP set in a finite group G , then there is a normal subgroup $H \leq G$, and Bohr neighborhood B in H , such that almost all translates of B are almost contained in A or almost disjoint from A (up to some error $\epsilon > 0$). Moreover, A is approximately a union of translates of B , and the index of H and complexity of B are bounded in terms of k and ϵ .

The proof of Theorem 5.7 from Theorem 5.5 is of course in analogy to the proof of Theorem 3.2 from Theorem 3.1. However, the argument is significantly more complicated, due to certain crucial differences between Bohr neighborhoods and subgroups. Specifically, Bohr neighborhoods are not closed under the group operation, and moreover distinct translates of a Bohr neighborhood need not be disjoint.¹⁷ Thus, when using Theorem 5.5 to obtain results for finite groups, we will focus solely on the *regularity* statement. In order to then deduce *structure* from *regularity* in finite groups, we will need to argue as at the end of the proof of Theorem 5.5, while also taking into account the quantitative behavior of Bohr neighborhoods described in Propositions 4.5 and 4.9. So we first prove a lemma, which gives a rather flexible version of the *regularity* statement, and also contains the ultraproduct argument necessary to prove Theorem 5.7.

Lemma 5.6. *For any $k \geq 1$ and $\epsilon > 0$, and any function $\gamma: (\mathbb{Z}^+)^2 \times (0, 1] \rightarrow \mathbb{R}^+$, there is $n = n(k, \epsilon, \gamma)$ such that the following holds. Suppose G is a finite group and $A \subseteq G$ is k -NIP. Then there are*

- * *a normal subgroup $H \leq G$ of index $m \leq n$,*
- * *a (δ, r) -Bohr neighborhood B in H and a δ -approximate $(3\delta, r)$ -Bohr neighborhood Y in H , where $r \leq n$ and $1/n \leq \delta \leq 1$, and*
- * *a set $Z \subseteq G$ with $|Z| < \epsilon|G|$,*

such that $B \subseteq Y \subseteq H$ and, for any $g \in G \setminus Z$, either $|gY \cap A| < \gamma(m, r, \delta)|B|$ or $|gY \setminus A| < \gamma(m, r, \delta)|B|$. Moreover, H , Y , and Z are in the Boolean algebra generated by $\{gAh : g, h \in G\}$, and if G is abelian then we may assume $H = G$.

Proof. Suppose not. Then we have $k \geq 1$, $\epsilon > 0$, and $\gamma: (\mathbb{Z}^+)^2 \times (0, 1] \rightarrow \mathbb{R}^+$ witnessing this. In particular, for any $n \geq 1$, there is a finite group G_n and a k -NIP subset $A_n \subseteq G_n$ such that, for any $H, B, Y, Z \subseteq G_n$, if

- * H is a normal subgroup of G_n of index $m \leq n$, and $H = G_n$ if G_n is abelian,
- * H, Y , and Z are in the Boolean algebra generated by $\{gA_nh : g, h \in G_n\}$,

¹⁷This is also the reason why Theorem 5.7 does not yield a graph regularity statement for $\Gamma_G(A)$ involving a partition into translates of Bohr neighborhoods.

* B is a (δ, r) -Bohr neighborhood in H and Y is a δ -approximate $(3\delta, r)$ -Bohr neighborhood in H , where $r \leq n$ and $1/n \leq \delta \leq 1$,

* $|Z| < \epsilon |G_n|$ and $B \subseteq Y \subseteq H$,

then there is some $g \in G_n \setminus Z$ such that $|gY \cap A_n| \geq \gamma(m, r, \delta)|B|$ and $|gY \setminus A_n| \geq \gamma(m, r, \delta)|B|$.

Let $\mathcal{L}$ be the group language together with an extra predicate A , and consider each (G_n, A_n) as a finite $\mathcal{L}$ -structure. Let $\mathcal{U}$ be a nonprincipal ultrafilter on $\mathbb{Z}^+$, and let G be a sufficiently saturated elementary extension of $M := \prod_{\mathcal{U}} (G_n, A_n)$. We identify $A = A(G)$. Let $\theta(x; y)$ be the formula $A(y \cdot x)$. Note that $\theta(x; y)$ is invariant, and k -NIP (in G) by Łoś's Theorem and since $M \prec G$. Finally, let $\alpha := \min\{\alpha_{\mathbb{T}^1}, 1\}$, where $\alpha_{\mathbb{T}^1} > 0$ is as in Corollary 4.4. By Theorem 5.5 (with $t = 3$), there are θ^r -definable sets $Y, Z \subseteq G$ (for Y we obtain θ^r -definability via Proposition 5.3), a θ^r -definable finite-index normal subgroup $H \leq G$, and a θ^r -definable δ -homomorphism $f: H \rightarrow \mathbb{T}^r$, for some $r \in \mathbb{N}$ and $0 < \delta < \alpha$, such that

* if G is abelian then $H = G$,

* $\mu(Z) < \epsilon$,

* $f(H)$ is finite and $Y = \{x \in H : d_{\mathbb{T}^r}(f(x), 0_{\mathbb{T}^r}) < 3\delta\}$, and

* for any $g \in G \setminus Z$, either $\mu(gY \cap A) = 0$ or $\mu(gY \setminus A) = 0$.

Let $m = [G : H]$, and set $\epsilon^* = \gamma(m, r, \delta)m^{-1}\delta^r > 0$. Let $\Lambda = f(H)$ and, given $\lambda \in \Lambda$, let $F(\lambda) = f^{-1}(\lambda)$. Then each $F(\lambda)$ is θ^r -definable by Remark 2.6.

Fix (θ^r) -formulas $\phi(x; \bar{y})$, $\psi(x; \bar{z})$, $\zeta(x; \bar{u})$, and $\xi_\lambda(x; \bar{v}_\lambda)$ for $\lambda \in \Lambda$ (without parameters) such that H, Y, Z , and $F(\lambda)$ for $\lambda \in \Lambda$ are defined by instances of $\phi(x; \bar{y})$, $\psi(x; \bar{z})$, $\zeta(x; \bar{u})$, and $\xi_\lambda(x; \bar{v}_\lambda)$, respectively. Given $n \in \mathbb{N}$, let μ_n denote the normalized counting measure on G_n . Let d denote $d_{\mathbb{T}^r}$. Define $I \subseteq \mathbb{Z}^+$ to be the set of $n \in \mathbb{Z}^+$ such that, for some tuples $\bar{a}_n, \bar{b}_n, \bar{c}_n$, and $\bar{d}_{n,\lambda}$ for $\lambda \in \Lambda$,

- (i) $\phi(x; \bar{a}_n)$ defines a normal subgroup H_n of G_n of index m , and if G is abelian then so is $H_n = G_n$,
- (ii) $\zeta(x; \bar{c}_n)$ defines a subset Z_n with $\mu_n(Z_n) < \epsilon$,
- (iii) for each $\lambda \in \Lambda$, $\xi_\lambda(x; \bar{d}_{n,\lambda})$ defines a subset $F_n(\lambda)$ of H_n , and $(F_n(\lambda))_{\lambda \in \Lambda}$ forms a partition of H_n ,
- (iv) if $f_n: H_n \rightarrow \Lambda$ is defined so that $f_n(x) = \lambda$ if and only if $x \in F_n(\lambda)$, then f_n is a δ -homomorphism from H_n to $\mathbb{T}^r$,
- (v) $\psi(x; \bar{b}_n)$ defines the set $Y_n = \{x \in H_n : d(f_n(x), 0_{\mathbb{T}^r}) < 3\delta\}$, and
- (vi) for all $g \in G_n \setminus Z_n$, either $\mu_n(gY_n \cap A_n) < \epsilon^*$ or $\mu_n(gY_n \setminus A_n) < \epsilon^*$.

We claim that $I \in \mathcal{U}$ by Łoś's Theorem and since $M \prec G$. In other words, the claim is that conditions (i) through (vi) are first-order expressible (possibly using the expanded language discussed in Remark 2.13). This is clear for (i), (ii), (iii), and (vi), and the only subtleties lie in (iv) and (v). In both cases, the crucial point is that Λ is finite, and so these conditions can be described by first-order sentences (using similar ideas to those in Proposition 5.3). For instance, to express condition (iv), fix $\lambda \in \Lambda$, and let

$P_\lambda = \{(\lambda_1, \lambda_2) \in \Lambda^2 : d(\lambda, \lambda_1 + \lambda_2) < \delta\}$. Let σ_λ be a sentence expressing that for any x, y , if $x \cdot y \in F(\lambda)$ then $x \in F(\lambda_1)$ and $y \in F(\lambda_2)$ for some $(\lambda_1, \lambda_2) \in P(\lambda)$. Then the conjunction $\bigwedge_{\lambda \in \Lambda} \sigma_\lambda$ expresses precisely that f is a δ -homomorphism. The details for condition (v) are similar and left to the reader.

Since $I \in \mathcal{U}$, we may fix $n \in I$ such that $n \geq \max\{m, r, \delta^{-1}\}$. Since $0 < \delta < \alpha$ and Y_n is a δ -approximate $(3\delta, r)$ -Bohr neighborhood in H_n , it follows from Corollary 4.4 that Y_n contains a (δ, r) -Bohr neighborhood B in H_n . So, by choice of (G_n, A_n) , there must be $g \in G_n \setminus Z_n$ such that $\mu_n(gY_n \cap A_n) \geq \gamma(m, r, \delta)\mu_n(B)$ and $\mu_n(gY_n \setminus A_n) \geq \gamma(m, r, \delta)\mu_n(B)$. So, to obtain a contradiction (to (vi)), it suffices to show that $\epsilon^* \leq \gamma(m, r, \delta)\mu_n(B)$, i.e. (by choice of ϵ^*), show $m^{-1}\delta^r \leq \mu_n(B)$. To see this, note that $|B| \geq \delta^r|H_n|$ by Proposition 4.5 (applied with $L = \mathbb{T}^r$, so $\ell_{\delta/2} = \delta^r$), and so $|B| \geq \delta^r m^{-1}|G_n|$ since $[G_n : H_n] = m$. ■

We now prove the main result for NIP subsets of arbitrary finite groups.

Theorem 5.7. *For any $k \geq 1$ and $\epsilon > 0$ there is $n = n(k, \epsilon)$ such that the following holds. Suppose G is a finite group and $A \subseteq G$ is k -NIP. Then there are*

- * a normal subgroup $H \leq G$ of index $m \leq n$,
- * a (δ, r) -Bohr neighborhood B in H , where $r \leq n$ and $1/n \leq \delta \leq 1$, and
- * a subset $Z \subseteq G$, with $|Z| < \epsilon|G|$,

satisfying the following properties:

- (i) (structure) *There is a set $D \subseteq G$, which is a union of at most $m(2/\delta)^r$ translates of B , such that*

$$|(A \triangle D) \setminus Z| < \epsilon|B|.$$

- (ii) (regularity) *For any $g \in G \setminus Z$, either $|gB \cap A| < \epsilon|B|$ or $|gB \setminus A| < \epsilon|B|$.*

Moreover, H and Z are in the Boolean algebra generated by $\{gAh : g, h \in G\}$, and if G is abelian then we may assume $H = G$.

Proof. Fix $k \geq 1$ and $\epsilon > 0$. Define $\gamma : (\mathbb{Z}^+)^2 \times (0, 1] \rightarrow \mathbb{R}^+$ such that $\gamma(x, y, z) = \epsilon x^{-1}(z/2)^y$. Let $n = n(k, \epsilon, \gamma)$ be given by Lemma 5.6. Fix a finite group G and a k -NIP subset $A \subseteq G$. By Lemma 5.6, there are

- * a normal subgroup $H \leq G$ of index $m \leq n$,
- * a subset $Y \subseteq H$,
- * a (δ, r) -Bohr neighborhood B in H , where $r \leq n$ and $1/n \leq \delta \leq 1$, and
- * a set $Z \subseteq G$ with $|Z| < \epsilon|G|$,

such that H, Z are in the Boolean algebra generated by $\{gAh : g, h \in G\}$, $B \subseteq Y \subseteq G$, and for all $g \in G \setminus Z$, either $|gY \cap A| < \gamma(m, r, \delta)|B|$ or $|gY \setminus A| < \gamma(m, r, \delta)|B|$. Moreover, if G is abelian then we may assume $H = G$. Since $B \subseteq Y$ and $\gamma(m, r, \delta) \leq \epsilon$, this immediately yields condition (ii).

For condition (i), we argue as in the proof of Theorem 5.5. First, by Proposition 4.9(b), there is a set $F \subseteq G \setminus Z$ such that $|F| \leq m(2/\delta)^r$ and $G \setminus Z \subseteq FB$. Let $I = \{g \in F :$

$|gB \setminus A| < \gamma(m, r, \delta)|B|$, and note that if $g \in F \setminus I$ then $|gB \cap A| < \gamma(m, r, \delta)|B|$. Let $D = IB$. Since $G \setminus Z \subseteq FB$, we have

$$A \triangle D \subseteq Z \cup \bigcup_{g \in I} (gB \setminus A) \cup \bigcup_{g \in F \setminus I} (gB \cap A), \quad \text{and so}$$

$$|(A \triangle D) \setminus Z| \leq \sum_{g \in I} |gB \setminus A| + \sum_{g \in F \setminus I} |gB \cap A| < |F| \gamma(m, r, \delta) |B| \leq \epsilon |B|. \quad \blacksquare$$

We end this section with some remarks on NIP subsets of finite *simple* groups. To motivate this, we first consider the stable case. In particular, given $k \geq 1$ and $\epsilon > 0$, it follows immediately from Theorem 1.1 that if G is a sufficiently large (depending on k and ϵ) finite simple group, and $A \subseteq G$ is k -stable, then $|A| < \epsilon |G|$ or $|A| > (1 - \epsilon) |G|$. On the other hand, this conclusion fails for NIP subsets of abelian finite simple groups (e.g., by the examples of NIP sets in $\mathbb{Z}/p\mathbb{Z}$ mentioned in the introduction). So it is interesting to observe that for *nonabelian* finite simple groups, one recovers the same “triviality” for NIP sets.

Corollary 5.8. *Given $k \geq 1$ and $\epsilon > 0$, there is $m = m(k, \epsilon)$ such that the following holds. Suppose G is a nonabelian finite simple group, with $|G| > m$, and $A \subseteq G$ is k -NIP. Then $|A| < \epsilon |G|$ or $|A| > (1 - \epsilon) |G|$.*

Proof. Let $m = n(k, \epsilon/2)$ be as in Theorem 5.7. Suppose G is a nonabelian finite simple group, with $|G| > m$, and $A \subseteq G$ is k -NIP. Then we have a normal subgroup $H \leq G$ of index at most m , a (δ, r) -Bohr neighborhood B in H (for some δ and r), a set $Z \subseteq G$ with $|Z| < \epsilon/2 |G|$, and a union D of translates of B such that $|(A \triangle D) \setminus Z| < \epsilon/2 |B|$. So $|A \triangle D| \leq |(A \triangle D) \setminus Z| + |Z| < \epsilon |G|$. Since G is simple and $|G| > m$, we have $H = G$. Now B contains the kernel K of a homomorphism from G to $\mathbb{T}^r$. So G/K is abelian, which implies $K = B = G$ since G is simple and nonabelian. Now D is either $\emptyset$ or G , and the result follows. $\blacksquare$

The previous proof highlights the significance of obtaining Bohr neighborhoods defined from *abelian* compact Lie groups, which is ultimately possible thanks to Theorem 2.7.

Remark 5.9. Corollary 5.8 is one of the few results on NIP sets in *nonabelian* finite groups for which we know an explicit bound. Specifically, in [13], the first author used techniques of Alon, Fox, and Zhao [3] and a result of Gowers [22] on “quasirandom” groups to obtain the bound $m = \exp(c(90/\epsilon)^{6k-6})$ in Corollary 5.8, where c is an absolute constant (see [13, Remark 8.7]).

As for stable sets, the proofs of Corollaries 1 and 3 in [53] yield an explicit lower bound of the form $|G| \geq \exp(c_k(\epsilon^{-d_k}))$ for the analogue of Corollary 5.8 in which A is k -stable and G is an *abelian* finite simple group (i.e., a cyclic group of prime order). The work in [12] removes the exponential and yields improved values for the constants c_k, d_k (see [12, Theorem 1.3]).

6. Distal arithmetic regularity

In this section, we adapt the preceding results to the case of NIP *fsg* groups with smooth left-invariant measures (e.g., *fsg* groups definable in *distal* theories). In contrast to previous results, where we focused on a single NIP formula, here we will operate in the setting of groups definable in an NIP theory.

In this section, we let T be a complete theory and we work in a saturated model M^* . Given $M \preceq M^*$ and an M -definable set X in M^* , a **Keisler measure on X over M** is a finitely additive probability measure defined on the Boolean algebra of M -definable subsets of X . If $M^* = M$ then μ is a **global Keisler measure on X** . If T is NIP, then a global Keisler measure μ on a definable set X is **generically stable** if there is a small model $M \prec M^*$ such that X is M -definable, μ is *M -definable* (i.e., for any formula $\phi(x; \bar{y})$ with x in the same sort as X , the map $\bar{b} \mapsto \mu(\phi(x; \bar{b}) \cap X)$ is M -definable) and μ is *finitely satisfiable in M* (i.e., if $Y \subseteq X$ is definable and $\mu(Y) > 0$ then $Y \cap M \neq \emptyset$). See [46, Section 7.5] for details.

Definition 6.1. (T is NIP.) A definable group G is *fsg* if it admits a generically stable left-invariant (global) Keisler measure.

We note that this is not the original definition of *fsg* (which is given in [29]), but rather the right characterization (in NIP theories) for our purposes (see [29, Proposition 6.2] and [32, Remark 4.2]). The significance of this notion in the context of our work is illustrated by the following example.

Example 6.2. If T is NIP and G is a definable pseudofinite group, then G is *fsg* since the pseudofinite counting measure on G is generically stable. This follows directly from the VC-Theorem (see, e.g., [46, Example 7.32], [15, Section 2]).

The *fsg* property for a definable group in an NIP theory has strong consequences. For example, we will use the following result from [31].

Theorem 6.3 (Hrushovski & Pillay [31, Theorem 7.7]). *An *fsg* group definable in an NIP theory admits a unique left-invariant Keisler measure μ , which is also the unique right-invariant Keisler measure.*

In NIP theories, definable *fsg* groups also satisfy a generic compact domination statement similar to Theorem 2.14(d) (see [48, Corollary 4.9]). However, in this section, we focus on a certain strengthening of generic compact domination for definable *fsg* groups in NIP theories whose unique left-invariant Keisler measure is *smooth*.

Definition 6.4. Let X be definable in M^* . A global Keisler measure μ on X is **smooth** if there is a small model $M \prec M^*$ such that X is M -definable and μ is the unique global Keisler measure on X extending $\mu|_M$.

If T is NIP, then any smooth measure is generically stable (see [46, Proposition 7.10 and Theorem 7.29]), and so any definable group with a smooth left-invariant Keisler measure is *fsg*. A place to find smooth measures is in the setting of *distal theories*, which were

introduced by Simon in [45]. For our purposes we will take the following characterization as a *definition* of distality (see [45, Theorem 1.1]).

Definition 6.5. T is **distal** if it is NIP and every global generically stable Keisler measure is smooth.

So, in particular, if T is distal and G is a definable fsg group, then the unique left-invariant Keisler measure on G is smooth. In [46, Proposition 8.41], Simon shows that if T is countable and NIP, and G is a definable group with a smooth left-invariant Keisler measure, then generic compact domination for G can be strengthened to outright “compact domination”. We give a formulation of this result suitable for our applications. First, recall that if T is NIP and G is a definable group, then G has a smallest type-definable subgroup of bounded index, denoted G^{00} , which is an intersection of at most $|T|$ definable sets (see [44] and/or [29, Proposition 6.1]). The reader should compare the following result to Lemma 2.16.

Lemma 6.6. *Assume T is countable and NIP, and G is a definable fsg group. Let μ be the unique left-invariant Keisler measure on G , and assume μ is smooth. Let $(W_i)_{i=0}^\infty$ be a decreasing sequence of definable sets such that $G^{00} = \bigcap_{i=0}^\infty W_i$. Fix a definable set $A \subseteq G$. Then, for any $\epsilon > 0$, there is a definable set $Z \subseteq G$ and some $i \in \mathbb{N}$ such that $\mu(Z) < \epsilon$ and, for any $g \in G \setminus Z$, either $gW_i \cap A = \emptyset$ or $gW_i \subseteq A$.*

Proof. Let $F = \{C \in G/G^{00} : C \cap A \neq \emptyset \text{ and } C \cap (G \setminus A) \neq \emptyset\}$. Then F is closed by Fact 2.4(b), and $\eta_{G/G^{00}}(F) = 0$ by [46, Proposition 8.41]. Now fix $\epsilon > 0$. By Fact 2.4(e), there is a definable set $Z \subseteq G$ such that $\mu(Z) < \epsilon$ and $\{a \in G : aG^{00} \in F\} \subseteq Z$. By saturation of G and since $(W_i)_{i=0}^\infty$ is decreasing, there is some $i \in \mathbb{N}$ such that for any $g \in G \setminus Z$, either $gW_i \cap A = \emptyset$ or $gW_i \subseteq A$. ■

We now prove analogues of Theorems 3.1 and 5.5 for an fsg group G , definable in an NIP theory, such that the unique left-invariant Keisler measure on G is smooth. In these results, the assumptions are stronger in the sense that the whole theory is assumed to be NIP. Moreover, in the conclusions we have definability of the data, but no claims about definability in a certain Boolean fragment (see Remark 6.9). On the other hand, we have outright inclusion or disjointness, rather than up to ϵ , which yields stronger structure and regularity statements.

Theorem 6.7. *Assume T is NIP. Let G be a definable fsg group, and let μ be the unique left-invariant Keisler measure on G . Suppose μ is smooth (e.g., if T is distal) and G/G^{00} is profinite. Fix a definable set $A \subseteq G$ and some $\epsilon > 0$. Then there are*

- * *a definable finite-index normal subgroup H of G , and*
- * *a set $Z \subseteq G$, which is a union of cosets of H with $\mu(Z) < \epsilon$,*

satisfying the following properties:

- (i) (structure) *$A \setminus Z$ is a union of cosets of H .*
- (ii) (regularity) *For any $g \in G \setminus Z$, either $gH \cap A = \emptyset$ or $gH \subseteq A$.*

Proof. The proof is similar to that of Theorem 3.1. First, we may restrict to a countable language in which G and A are definable and μ is still smooth (see [46, Lemma 7.8] and subsequent remarks). By Fact 2.3(c), we also preserve the hypothesis that G/G^{00} is profinite.

By Lemma 2.9, we may fix a decreasing sequence $(H_i)_{i=0}^\infty$ of definable finite-index normal subgroups of G such that $G^{00} = \bigcap_{i=0}^\infty H_i$. By Lemma 6.6, we have a definable set Z and some $H = H_i$ satisfying condition (ii). As in the proof of Theorem 3.1, we may assume Z is a union of cosets of H . For condition (i), set $D = \{g \in G : gH \subseteq A \text{ and } gH \cap Z = \emptyset\}$. Then D is a union of cosets of H , and $D \subseteq A \subseteq D \cup Z$. Since $D \cap Z = \emptyset$, we have $A \setminus Z = D$. ■

Now we prove the general statement.

Theorem 6.8. *Assume T is NIP. Let G be a definable fsg group, and let μ be the unique left-invariant Keisler measure on G . Suppose μ is smooth (e.g., if T is distal). Fix a definable set $A \subseteq G$ and some $\epsilon > 0$. Then there are*

- * a definable finite-index normal subgroup H of G ,
- * a compact connected metric Lie group (L, d) ,
- * a definable homomorphism $\pi: H \rightarrow L$, and
- * a definable set $Z \subseteq G$ with $\mu(Z) < \epsilon$,

such that, for any integer $t \geq 1$, there is

- * a definable (t, π) -approximate Bohr chain $(Y_m)_{m=0}^\infty$ in H ,

satisfying the following properties, for any $m \in \mathbb{N}$:

- (i) (structure) *There is $D_m \subseteq G$, which is a union of finitely many left translates of Y_m , such that $D_m \subseteq A \subseteq D_m \cup Z$.*
- (ii) (regularity) *For any $g \in G \setminus Z$, either $gY_m \cap A = \emptyset$ or $gY_m \subseteq A$.*

Moreover, if G is pseudofinite then we may assume $L = \mathbb{T}^n$ for some $n \in \mathbb{N}$; and if G/G^{00} is abelian then we may assume $L = \mathbb{T}^n$ for some $n \in \mathbb{N}$ and $H = G$.

Proof. The proof is similar to that of Theorem 5.5. First, as in the proof of Theorem 6.7, we may assume T is in a countable language. Let $(W_i)_{i=0}^\infty$ be a decreasing sequence of definable subsets of G satisfying the conditions of Proposition 5.1, with $\Gamma = G^{00}$. By Lemma 6.6, there is a definable set $Z \subseteq G$ and some $i \in \mathbb{N}$ such that $\mu(Z) < \epsilon$ and, if $W := W_i$, then for all $g \in G \setminus Z$, either $gW \cap A = \emptyset$ or $gW \subseteq A$. Proposition 5.1 associates to W a definable homomorphism $\pi: H \rightarrow L$, with $\ker \pi \subseteq W$, where H is a definable finite-index normal subgroup of G and (L, d) is a compact connected metric Lie group. Note also that the “moreover” statement is given by Proposition 5.1. By Lemma 5.4, there is a definable (t, π) -approximate Bohr chain $(Y_m)_{m=0}^\infty$ in H . By saturation, if m is sufficiently large then $Y_m \subseteq W$ and so if $g \in G \setminus Z$ then either $gY_m \cap A = \emptyset$ or $gY_m \subseteq A$. This yields condition (ii). For condition (i), mimic the end of the proof of Theorem 5.5 to find $D_m \subseteq G$, which is a union of finitely many left translates of Y_m , such

that $D_m \subseteq A \subseteq D_m \cup Z$ (replace each occurrence of $\mu(A \cap gY_m) = 0$ with $A \cap gY_m = \emptyset$, and each occurrence of $\mu(A \setminus gY_m) = 0$ with $gY_m \subseteq A$). ■

Since Theorems 6.7 and 6.8 use global assumptions on the theory, a fully general statement for applications to finite groups would be rather cumbersome to state. However, one concludes, in a routine fashion, structure and regularity for suitable families of finite groups, for example a collection $\mathcal{G}$ of finite $\mathcal{L}$ -structures expanding groups, such that any completion of $\text{Th}(\mathcal{G})$ is distal. In the next section, we will discuss an application to the family of finite groups obtained as quotients of a compact p -adic Lie group by its open normal subgroups.

Remark 6.9. Results along the lines of Theorems 3.1 and 5.5 can also be shown for fsg groups definable in NIP theories, but without the smoothness assumption. Indeed, such groups satisfy generic compact domination just as in Theorem 2.14(d), but with $G_{\theta^r}^{00}$ replaced by G^{00} (see [48, Corollary 4.9]). However, as discussed in [15, Remark 1.3], since the unique left-invariant measure on an NIP fsg group G is generically stable, one could fix an invariant formula $\theta(x; \bar{y})$ and construct $G_{\theta^r}^{00}$ as in Theorem 2.14. This yields structure and regularity theorems as before with additional information about the definability of the data. Precisely:

Assume T is NIP. Suppose G is definable and fsg , and let μ be the unique left-invariant Keisler measure on G . Without loss of generality, assume $G = M^*$. Fix an invariant formula $\theta(x; \bar{y})$. Then, for any θ^r -definable $A \subseteq G$ and any $\epsilon > 0$, we have the conclusion of Theorem 5.5, except with $\mathbb{T}^n$ replaced by some compact connected metric Lie group (L, d) . If $G/G_{\theta^r}^{00}$ is profinite, then the conclusion of Theorem 3.1 holds exactly as stated.

It would be interesting to pursue notions of smoothness for local measures, or “local distality” for formulas, and recover local versions of Theorems 6.8 and 6.7. For instance, one might consider an NIP formula $\theta(x, \bar{y})$ such that every generically stable global Keisler measure on the Boolean algebra of θ -formulas is smooth.

7. Compact p -adic analytic groups

In this section, we apply Theorem 6.7 to the setting of compact p -adic analytic groups (see Theorems 7.3 and 7.5). We assume some familiarity with the p -adic field $\mathbb{Q}_p$ and p -adic model theory. See [6] for further reading. The topology on $\mathbb{Q}_p$ is given by the valuation v where open neighborhoods of a point a are defined by $v(x - a) \geq n$ for $n \in \mathbb{Z}$. The topology on $\mathbb{Q}_p^n$ is the product topology. A **p -adic analytic function** is a function f , from some open $V \subseteq \mathbb{Q}_p^n$ to $\mathbb{Q}_p$, such that for every $a \in V$, there is an open neighborhood of $a \in V$ in which f is given by a convergent power series. We obtain the notions of a p -adic analytic manifold and a p -adic analytic (or Lie) group. Recall that any compact p -adic analytic group is profinite (see also [43, Section 2.12.1]).

We let $\mathbb{Q}_p^{\text{an}}$ denote the expansion of the field $(\mathbb{Q}_p, +, \cdot)$ by symbols for all convergent (in $\mathbb{Z}_p$) power series in $\mathbb{Z}_p[[X_1, \dots, X_n]]$ for all n . Then any compact p -adic analytic manifold or group is seen to be naturally definable in the structure $\mathbb{Q}_p^{\text{an}}$ (we conflate definable and interpretable at this point). For example, given a compact p -adic analytic group,

we can find an atlas consisting of finitely many open definable (even semialgebraic) sets with analytic transition functions and such that group operation is analytic when read in the charts. When we talk about a compact p -adic analytic group being definable in $\mathbb{Q}_p^{\text{an}}$ we will mean definable in such a manner.

We note that $\text{Th}(\mathbb{Q}_p^{\text{an}})$ is distal. Indeed, this follows from several results in the literature: distality of $\text{Th}(\mathbb{Q}_p, +, \cdot)$ [46, Example 9.20], dp -minimality of $\text{Th}(\mathbb{Q}_p^{\text{an}})$ (see [4, Corollary 7.9 and Remark 7.10]), and the fact that a dp -minimal expansion of a distal and dp -minimal theory is distal [10, Remark 6.7]. It is also well known that distality passes from T to T^{eq} (see [46, Exercise 9.12]).

The next lemma follows from [36], and we will give an explanation afterwards.

Lemma 7.1. *Let K be a compact p -adic analytic group. Then the family of open normal subgroups is uniformly definable in $\mathbb{Q}_p^{\text{an}}$.*

Explanation. First, by [36, Theorem 2.1], K has an open normal subgroup G (so of finite index) which is a “uniformly powerful pro- p group of finite (topological) rank d ” (see [36] for the definitions). Proposition 1.2 of [36], and its proof, states that G is isomorphic (as a topological group) to $\mathbb{Z}_p^d$ equipped with a certain analytic group structure $*$ (and note $(\mathbb{Z}_p^d, *)$ is definable in $\mathbb{Q}_p^{\text{an}}$) such that moreover the collection of open subgroups of $(\mathbb{Z}_p^d, *)$ is (uniformly) definable in $\mathbb{Q}_p^{\text{an}}$. By [16, Theorem 10.5], there is an analytic (so definable in $\mathbb{Q}_p^{\text{an}}$) isomorphism between G and $(\mathbb{Z}_p^d, *)$, whereby the family of open subgroups of G is uniformly definable in $\mathbb{Q}_p^{\text{an}}$. As G has finite index and is normal in K , it follows that the collection of open subgroups of K is uniformly definable in $\mathbb{Q}_p^{\text{an}}$. (See [36, last part of the proof of Theorem 1.1((1) $\Rightarrow$ (2)) on page 1043 of that paper].)

In particular, the collection of quotients of a compact p -adic analytic group by its open normal subgroups is a family of uniformly definable (in $\mathbb{Q}_p^{\text{an}}$) finite groups. Let us fix a compact p -adic analytic group K (so definable in $\mathbb{Q}_p^{\text{an}}$). If $M^* > \mathbb{Q}_p^{\text{an}}$ is sufficiently saturated, then $K(M^*)$ denotes the group definable in the structure M^* by the same formula as the one defining K in $\mathbb{Q}_p^{\text{an}}$. The next result combines the “Claim” in [29, Section 6] with Corollaries 2.3 and 2.4 of [40].

Fact 7.2 ([29, 40]). *Let $M^* > \mathbb{Q}_p^{\text{an}}$ be sufficiently saturated. Then $K(M^*)$ is fsg. Moreover, $K(M^*)/K(M^*)^{00}$ is isomorphic to K , and thus is profinite. In particular, $K(M^*)^{00} = \bigcap_{i=0}^{\infty} H_i(M^*)$ where $\{H_i : i \in \mathbb{N}\}$ is the neighborhood basis at the identity consisting of open normal subgroups of K .*

Our first application of Theorem 6.7 is the following “structure and regularity” statement for definable subsets of compact p -adic analytic groups. In fact, as we point out below, this can also be seen as a fairly direct application of [40, Proposition 2.8], and an extension of certain results in that paper.

Theorem 7.3. *Let K be a compact p -adic analytic group (so definable in the structure $\mathbb{Q}_p^{\text{an}}$). Let $A \subseteq K$ be definable in $\mathbb{Q}_p^{\text{an}}$, and let $\epsilon > 0$. Then there are*

- * *an open (so finite-index) normal subgroup H of K , and*
- * *a set $Z \subseteq K$, which is a union of cosets of H , with $\eta_K(Z) < \epsilon$,*

satisfying the following properties:

- (i) (structure) $A \setminus Z$ is a union of cosets of H .
- (ii) (regularity) For any $g \in K \setminus Z$, either $gH \cap A = \emptyset$ or $gH \subseteq A$.

Proof. Let $M^* \succ \mathbb{Q}_p^{\text{an}}$ be sufficiently saturated. By distality, we may apply Theorem 6.7 to the group $K(M^*)$ and the definable set $A(M^*)$. Since $K(M^*)^{00}$ is contained in any definable finite-index subgroup of $K(M^*)$, we may use Fact 7.2 to further assume that the finite-index normal subgroup of $K(M^*)$ given by Theorem 6.7 is of the form $H(M^*)$ for some open normal subgroup H of K . So the error set is of the form $Z(M^*)$, where Z is definable in $\mathbb{Q}_p^{\text{an}}$. Note also that if μ is the (unique) left-invariant Keisler measure on $K(M^*)$, then $\mu(H(M^*)) = \eta_K(H) = 1/[K : H]$ since H is a finite index subgroup of K , and thus $\mu(Z(M^*)) = \eta_K(Z)$. Altogether, by elementarity, we have (i) and (ii). ■

Remark 7.4.

- (1) The proof of Theorem 7.3 only requires definability of the open normal subgroups of K (and not *uniform* definability).
- (2) Proposition 2.8 of [40] states that $K(M^*)$ is compactly dominated via the map $K(M^*) \rightarrow K(M^*)/K(M^*)^{00}$. So we could also have deduced Theorem 7.3 from this result, together with the standard methods.
- (3) It would be interesting to prove Theorem 7.3 from the cell decomposition results of Denef and others (at least when $K = \mathbb{Z}_p^n$ for some n).

Our second application of Theorem 6.7 is to the family of quotients of a compact p -adic analytic group by its open normal subgroups.

Theorem 7.5. *Let K be a compact p -adic analytic group. Let $(G_i)_{i \in I}$ be the family of finite groups obtained as quotients of K by open normal subgroups. Let $A \subseteq K$ be definable in $\mathbb{Q}_p^{\text{an}}$ and for $i \in I$, let $A_i \subseteq G_i$ be the image of A under the quotient map. Fix $\epsilon > 0$. There is some $n = n(K, A, \epsilon)$ such that for any $i \in I$, there are*

- * a normal subgroup $H_i \leq G_i$ of index at most n , and
- * a set $Z_i \subseteq G_i$, which is a union of cosets of H_i with $|Z_i| < \epsilon |G_i|$,

satisfying the following properties:

- (i) (structure) $A_i \setminus Z_i$ is a union of cosets of H_i .
- (ii) (regularity) For any $g \in G_i \setminus Z_i$, either $gH_i \cap A_i = \emptyset$ or $gH_i \subseteq A_i$.

Proof. It suffices to prove condition (i). Toward a contradiction, suppose we have $\epsilon > 0$ such that for any $n \geq 1$ there is some $i_n \in I$ such that if $H \leq G_{i_n}$ is a normal subgroup of index at most n , and $Z \subseteq G_{i_n}$ is a union of cosets of H with $|Z| < \epsilon |G_{i_n}|$, then $A_{i_n} \setminus Z$ is not a union of cosets of H .

Let $G_i = K/U_i$, where $(U_i)_{i \in I}$ lists the open normal subgroups of K . By Lemma 7.1, we may view U_t as a formula (over $\mathbb{Q}_p$) in the variable t , and I as a definable set in $\mathbb{Q}_p^{\text{an}}$ in the sort for t . Let $S_I(\mathbb{Q}_p)$ be the space of types with parameters from $\mathbb{Q}_p$ concentrating on I . Let $M^* \succ \mathbb{Q}_p^{\text{an}}$ be sufficiently saturated, and let $i_* \in I(M^*)$ realize an accumulation

point of $\{\text{tp}(i_n/\mathbb{Q}_p) : n \geq 1\}$ in $S_I(\mathbb{Q}_p)$. Let $G = K(M^*)/U_{i*}$. Then G is a pseudofinite group definable in M^* , and thus is *fsg* (see Example 6.2). Since G is a definable quotient of $K(M^*)$, and $K(M^*)/K(M^*)^{00}$ is profinite by Fact 7.2, it follows (using Fact 2.3(c)) that G/G^{00} is also profinite. By Theorem 6.7, there are $m, n, r \in \mathbb{N}$, with $n \geq 1$, $m < \epsilon n$, and $r \leq n$, and a formula $\phi(x; \bar{y})$ (with no parameters) such that:

- (*) For some $\bar{b} \in (M^*)^{\bar{y}}$, $H := \phi(M^*; \bar{b})$ is a normal subgroup of $G = K(M^*)/U_{i*}$ of index n , and there is $Z \subseteq G$, which is a union of m cosets of H , such that $(A(M^*)/U_{i*}) \setminus Z$ is a union of r cosets of H .

In particular, (*) is a property of i^* that can be expressed using a formula $\zeta(t)$ over $\mathbb{Q}_p$. So there is $n' \geq n$ such that $\mathbb{Q}_p^{\text{an}} \models \zeta(i_{n'})$, which is a contradiction. ■

Remark 7.6. Suppose K is a compact p -adic analytic group such that the family of open normal subgroups is (eventually) linearly ordered by inclusion (e.g., $K = \mathbb{Z}_p$). In this case, Theorem 7.5 can be deduced easily from Theorem 7.3. Indeed, given a definable set $A \subseteq G$ and some $\epsilon > 0$, let $H \leq K$ and $Z \subseteq K$ be as in Theorem 7.3. Then all but finitely many open normal subgroups of K are contained in H and, given such a subgroup $U_i \leq H$, if $H_i = H/U_i$ and $Z_i = Z/U_i$, then H_i and Z_i satisfy conditions (i) and (ii) in Theorem 7.5. Choosing n sufficiently large, we can then let H_i be trivial for any U_i not contained in H .

It is not clear whether such an argument can be given for arbitrary K (or if there is a proof of Theorem 7.5 not relying on Lemma 7.1). Note that the above assumption does not hold for every compact p -adic analytic group; for example $(\mathbb{Z}_p^2, +)$ does not have eventually linearly ordered open (normal) subgroups.

Acknowledgments. Much of the work in this paper was carried out during the 2018 Model Theory, Combinatorics and Valued Fields trimester program at Institut Henri Poincaré. We thank IHP for their hospitality. We also thank Julia Wolf for helpful conversations, and the anonymous referee for their careful reading and for providing many valuable revisions, corrections, and suggestions.

Funding. The authors were partially supported by NSF grants: DMS-1855503 (Conant); DMS-136702, DMS-1665035, DMS-1790212 (Pillay); DMS-1855711 (Terry).

References

- [1] Alekseev, M. A., Glebskiĭ, L. Y., Gordon, E. I.: On approximations of groups, group actions and Hopf algebras. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov.* **256**, 224–262, 268 (1999) Zbl [1130.20306](#) MR [1708567](#)
- [2] Alon, N., Fischer, E., Newman, I.: Efficient testing of bipartite graphs for forbidden induced subgraphs. *SIAM J. Comput.* **37**, 959–976 (2007) Zbl [1141.05073](#) MR [2341924](#)
- [3] Alon, N., Fox, J., Zhao, Y.: Efficient arithmetic regularity and removal lemmas for induced bipartite patterns. *Discrete Anal.* **2019**, art. 3, 14 pp. Zbl [07152902](#) MR [3943117](#)
- [4] Aschenbrenner, M., Dolich, A., Haskell, D., Macpherson, D., Starchenko, S.: Vapnik–Chervonenkis density in some theories without the independence property, I. *Trans. Amer. Math. Soc.* **368**, 5889–5949 (2016) Zbl [1423.03119](#) MR [3458402](#)

- [5] Basu, S.: Combinatorial complexity in o-minimal geometry. *Proc. London Math. Soc.* (3) **100**, 405–428 (2010) Zbl [1186.52017](#) MR [2595744](#)
- [6] Bélair, L.: Panorama of p -adic model theory. *Ann. Sci. Math. Québec* **36**, 43–75 (2012) Zbl [1362.11107](#) MR [3113291](#)
- [7] Bourgain, J.: On triples in arithmetic progression. *Geom. Funct. Anal.* **9**, 968–984 (1999) Zbl [0959.11004](#) MR [1726234](#)
- [8] Breuillard, E., Green, B., Tao, T.: The structure of approximate groups. *Publ. Math. Inst. Hautes Études Sci.* **116**, 115–221 (2012) Zbl [1260.20062](#) MR [3090256](#)
- [9] Chernikov, A., Simon, P.: Definably amenable NIP groups. *J. Amer. Math. Soc.* **31**, 609–641 (2018) Zbl [06870167](#) MR [3787403](#)
- [10] Chernikov, A., Starchenko, S.: Regularity lemma for distal structures. *J. Eur. Math. Soc.* **20**, 2437–2466 (2018) Zbl [06966829](#) MR [3852184](#)
- [11] Chernikov, A., Starchenko, S.: Definable regularity lemmas for NIP hypergraphs. *Quart. J. Math.* ([online](#), 2021)
- [12] Conant, G.: Quantitative structure of stable sets in arbitrary finite groups. [arXiv:2004.02819](#) (2020)
- [13] Conant, G.: On finite sets of small tripling or small alternation in arbitrary groups. *Combin. Probab. Comput.* **29**, 807–829 (2020) MR [4173133](#)
- [14] Conant, G., Pillay, A., Terry, C.: A group version of stable regularity. *Math. Proc. Cambridge Philos. Soc.* **168**, 405–413 (2020) Zbl [07167602](#) MR [4064112](#)
- [15] Conant, G., Pillay, A.: Pseudofinite groups and VC-dimension. *J. Math. Logic* ([online](#), 2020)
- [16] Dixon, J. D., du Sautoy, M. P. F., Mann, A., Segal, D.: *Analytic pro- p Groups*. 2nd ed., Cambridge Stud. Adv. Math. 61, Cambridge Univ. Press, Cambridge (1999) Zbl [0934.20001](#) MR [1720368](#)
- [17] Erdős, P., Hajnal, A.: Ramsey-type theorems. *Discrete Appl. Math.* **25**, 37–52 (1989) Zbl [0715.05052](#) MR [1031262](#)
- [18] Fox, J., Gromov, M., Lafforgue, V., Naor, A., Pach, J.: Overlap properties of geometric expanders. *J. Reine Angew. Math.* **671**, 49–83 (2012) Zbl [1306.05171](#) MR [2983197](#)
- [19] Fox, J., Pach, J., Suk, A.: A polynomial regularity lemma for semialgebraic hypergraphs and its applications in geometry and property testing. *SIAM J. Comput.* **45**, 2199–2223 (2016) Zbl [1353.05090](#) MR [3585030](#)
- [20] Fox, J., Pach, J., Suk, A.: Erdős–Hajnal conjecture for graphs with bounded VC-dimension. *Discrete Comput. Geom.* **61**, 809–829 (2019) Zbl [1411.05179](#) MR [3943496](#)
- [21] Gowers, W. T.: Lower bounds of tower type for Szemerédi’s uniformity lemma. *Geom. Funct. Anal.* **7**, 322–337 (1997) Zbl [0876.05053](#) MR [1445389](#)
- [22] Gowers, W. T.: Quasirandom groups. *Combin. Probab. Comput.* **17**, 363–387 (2008) Zbl [1191.20016](#) MR [2410393](#)
- [23] Gowers, W. T., Wolf, J.: Linear forms and higher-degree uniformity for functions on $\mathbb{F}_p^n$. *Geom. Funct. Anal.* **21**, 36–69 (2011) Zbl [1262.11012](#) MR [2773103](#)
- [24] Green, B.: A Szemerédi-type regularity lemma in abelian groups, with applications. *Geom. Funct. Anal.* **15**, 340–376 (2005) Zbl [1160.11314](#) MR [2153903](#)
- [25] Green, B., Tao, T.: An arithmetic regularity lemma, an associated counting lemma, and applications. In: *An Irregular Mind*, Bolyai Soc. Math. Stud. 21, János Bolyai Math. Soc., Budapest, 261–334 (2010) Zbl [1222.11015](#) MR [2815606](#)

- [26] Hewitt, E., Ross, K. A.: Abstract Harmonic Analysis. Vol. I. 2nd ed., Grundlehren Math. Wiss. 115, Springer, Berlin (1979) Zbl [0837.43002](#) MR [551496](#)
- [27] Hofmann, K. H., Morris, S. A.: The Structure of Compact Groups. De Gruyter Stud. Math. 25, De Gruyter, Berlin (2013) Zbl [1441.22001](#) MR [3114697](#)
- [28] Hrushovski, E.: Stable group theory and approximate subgroups. J. Amer. Math. Soc. **25**, 189–243 (2012) Zbl [1259.03049](#) MR [2833482](#)
- [29] Hrushovski, E., Peterzil, Y., Pillay, A.: Groups, measures, and the NIP. J. Amer. Math. Soc. **21**, 563–596 (2008) Zbl [1134.03024](#) MR [2373360](#)
- [30] Hrushovski, E., Pillay, A.: Groups definable in local fields and pseudo-finite fields. Israel J. Math. **85**, 203–262 (1994) Zbl [0804.03024](#) MR [1264346](#)
- [31] Hrushovski, E., Pillay, A.: On NIP and invariant measures. J. Eur. Math. Soc. **13**, 1005–1061 (2011) Zbl [1220.03016](#) MR [2800483](#)
- [32] Hrushovski, E., Pillay, A., Simon, P.: Generically stable and smooth measures in NIP theories. Trans. Amer. Math. Soc. **365**, 2341–2366 (2013) Zbl [1294.03023](#) MR [3020101](#)
- [33] Iltis, R.: Some algebraic structure in the dual of a compact group. Canad. J. Math. **20**, 1499–1510 (1968) Zbl [0169.34801](#) MR [232892](#)
- [34] Krupiński, K., Newelski, L.: On bounded type-definable equivalence relations. Notre Dame J. Formal Logic **43**, 231–242 (2003) (2002) Zbl [1050.03026](#) MR [2034748](#)
- [35] Lovász, L., Szegedy, B.: Regularity partitions and the topology of graphons. In: An Irregular Mind, Bolyai Soc. Math. Stud. 21, János Bolyai Math. Soc., Budapest, 415–446 (2010) Zbl [1242.05188](#) MR [2815610](#)
- [36] Macpherson, D., Tent, K.: Profinite groups with NIP theory and p -adic analytic groups. Bull. London Math. Soc. **48**, 1037–1049 (2016) Zbl [1430.03054](#) MR [3608949](#)
- [37] Malliaris, M., Shelah, S.: Regularity lemmas for stable graphs. Trans. Amer. Math. Soc. **366**, 1551–1585 (2014) Zbl [1283.05234](#) MR [3145742](#)
- [38] Marker, D.: Model Theory. Grad. Texts in Math. 217, Springer, New York (2002) Zbl [1003.03034](#) MR [1924282](#)
- [39] Nikolov, N., Schneider, J., Thom, A.: Some remarks on finitarily approximable groups. J. École Polytech. Math. **5**, 239–258 (2018) Zbl [1452.20025](#) MR [3749196](#)
- [40] Onshuus, A., Pillay, A.: Definable groups and compact p -adic Lie groups. J. London Math. Soc. (2) **78**, 233–247 (2008) Zbl [1153.03015](#) MR [2427062](#)
- [41] Pillay, A.: Type-definability, compact Lie groups, and o-minimality. J. Math. Logic **4**, 147–162 (2004) Zbl [1069.03029](#) MR [2114965](#)
- [42] Pillay, A.: Remarks on compactifications of pseudofinite groups. Fund. Math. **236**, 193–200 (2017) Zbl [1420.03063](#) MR [3591278](#)
- [43] Ribes, L., Zalesskii, P.: Profinite Groups, 2nd ed., Ergeb. Math. Grenzgeb. 40, Springer, Berlin (2010) Zbl [1197.20022](#) MR [2599132](#)
- [44] Shelah, S.: Minimal bounded index subgroup for dependent theories. Proc. Amer. Math. Soc. **136**, 1087–1091 (2008) Zbl [1144.03026](#) MR [2361885](#)
- [45] Simon, P.: Distal and non-distal NIP theories. Ann. Pure Appl. Logic **164**, 294–318 (2013) Zbl [1269.03037](#) MR [3001548](#)
- [46] Simon, P.: A Guide to NIP Theories. Lecture Notes in Logic 44, Assoc. Symbolic Logic, Chicago, IL, and Cambridge Sci. Publ., Cambridge (2015) Zbl [1332.03001](#) MR [3560428](#)
- [47] Simon, P.: Rosenthal compacta and NIP formulas. Fund. Math. **231**, 81–92 (2015) Zbl [1372.03068](#) MR [3361236](#)

- [48] Simon, P.: VC-sets and generic compact domination. *Israel J. Math.* **218**, 27–41 (2017) Zbl [1367.03068](#) MR [3625123](#)
- [49] Sisask, O.: Convolutions of sets with bounded VC-dimension are uniformly continuous. *Discrete Anal.* ([online](#), 2018)
- [50] Szemerédi, E.: Regular partitions of graphs. In: *Problèmes combinatoires et théorie des graphes* (Orsay, 1976), *Colloq. Internat. CNRS* 260, CNRS, Paris, 399–401 (1978) Zbl [0413.05055](#) MR [540024](#)
- [51] Tao, T., Vu, V.: *Additive Combinatorics*. *Cambridge Stud. Adv. Math.* 105, Cambridge Univ. Press, Cambridge (2006) Zbl [1127.11002](#) MR [2289012](#)
- [52] Terry, C., Wolf, J.: Stable arithmetic regularity in the finite field model. *Bull. London Math. Soc.* **51**, 70–88 (2019) Zbl [07040809](#) MR [3919562](#)
- [53] Terry, C., Wolf, J.: Quantitative structure of stable sets in finite abelian groups. *Trans. Amer. Math. Soc.* **373**, 3885–3903 (2020) Zbl [07207622](#) MR [4105513](#)