PSEUDOFINITE GROUPS AND VC-DIMENSION

GABRIEL CONANT AND ANAND PILLAY

Abstract. We develop "local NIP group theory" in the context of pseudofinite groups. In particular, given a sufficiently saturated pseudofinite structure G expanding a group, and left invariant NIP formula $\delta(x;\bar{y})$, we prove various aspects of "local *fsg*" for the right-stratified formula $\delta^r(x;\bar{y},u) := \delta(x\cdot u;\bar{y})$. This includes a δ^r-type-definable connected component, uniqueness of the pseudofinite counting measure as a left-invariant measure on δ^r-formulas, and generic compact domination for δ^r-definable sets.

1. Introduction

One of the more remarkable aspects of stable group theory is the ability to formulate useful abstract notions of tools from algebra, combinatorics, and topological dynamics. For example, given a group G definable in a (sufficiently saturated) model of a stable theory, one has at hand abstract versions of connected components, stabilizers, generic points, and invariant probability measures on definable sets (leading to the notion of *definable amenability*). As the entire field of model theory began moving outward from stability, so did the model theoretic study of groups, leading to a large body of work on groups definable in *simple* and *NIP* theories. In the case of NIP theories, connected components and invariant measures remain powerful tools for studying definable groups (e.g [1], [6], [8], [9]).

Another important aspect of stability theory is that it can be applied locally. In particular, many of the tools related to nonforking and the geometry of definable sets remain valid when one works around a single stable formula $\phi(\bar{x};\bar{y})$ (see, e.g., [7]). This is quite useful for applications to other areas of mathematics, as one would like to prove results about algebraic or combinatorial objects exhibiting good behavior related to stability and omitting half-graphs, but also have the freedom to work in an environment which is not stable (e.g. a nonstandard of model set theory).

In contrast, the study of local NIP theory is still work in progress, especially in the setting of groups. The goal of this paper is to examine NIP formulas in the context of pseudofinite groups. We will find that, in pseudofinite groups, NIP formulas exhibit many properties found in NIP groups with *finitely satisfiable generics* and *generically stable measures*. In fact, the results we obtain here could be formulated with the pseudofiniteness assumption replaced by a suitably local, albeit slightly cumbersome, assumption of generic stability for an invariant measure with respect to an NIP formula (see Remark 1.3). On the other hand, it is quite reasonable to

Date: March 24, 2018.

The second author was supported by NSF grants DMS-1360702 and DMS-1665035. This work was finished during the Model Theory, Combinatorics and Valued Fields trimester program at Institut Henri Poincaré. We thank IHP for their hospitality. The first author would like to thank Artem Chernikov for helpful conversations.

focus on pseudofinite groups, as this is a natural settings for applications to finite combinatorics and combinatorial number theory. Indeed, the work in this paper was originally motivated by generalizing the regularity lemma for stable subsets of finite groups from [2] (with C. Terry), to the NIP setting. In [3] (also with Terry), we use the work done here to obtain such a generalization.

Our setting is as follows. We work with a sufficiently saturated elementary extension G of an ultraproduct of expansions of finite groups (in some fixed language $\mathcal{L}$ expanding the language of groups), and let μ denote the pseudofinite counting measure. We also fix a formula $\delta(x;\bar{y})$, possibly with parameters, which is NIP and invariant in the sense that any left translate of an instance of δ is again an instance of δ (the canonical example of an invariant formula is something of the form $\phi(y\cdot x)$, where $\phi(x)$ is any formula). In order to prove our main results, it will be necessary to work mostly around the "right-stratified" formula $\delta^r(x;\bar{y},u)$, which we define to be $\delta(x\cdot u;\bar{y})$. The following theorem summarizes the main results of the paper.

Theorem 1.1. *Let G be as above, and assume $\delta(x;\bar{y})$ is NIP.*

(a) (Generic types) *Given a δ^r-formula $\phi(x)$, the following are equivalent:*
- (i) *$\phi(x)$ is left generic;*
- (ii) *$\phi(x)$ is right generic;*
- (iii) *$\mu(\phi(x))>0$.*

In particular, global generic δ-types and global generic δ^r-types exist.

(b) (Local G^{00}) *Let $G^{00}_{\delta^r}$ denote the intersection of all δ^r-type-definable bounded-index subgroups of G. Then:*
- (i) *$G^{00}_{\delta^r}$ is normal and δ^r-type-definable of bounded index.*
- (ii) *$G^{00}_{\delta^r}$ is the intersection of all stabilizer subgroups of the form* $\mathrm{Stab}_\mu(\phi(x)):=\{g\in G:\mu(\phi(g^{-1}x)\triangle\phi(x))=0\}$*, where $\phi(x)$ is a δ^r-formula.*
- (iii) *$G^{00}_{\delta^r}$ is the intersection of all stabilizer subgroups of the form* $\mathrm{Stab}(p):=\{g\in G:gp=p\}$*, where p is a generic δ-type over G.*
- (iv) *$G^{00}_{\delta^r}=\mathrm{Stab}(p)$ for any generic δ^r-type p over G.*

(c) (Local G^0) *Let $G^0_{\delta^r}$ denote the intersection of all δ^r-definable finite-index subgroups of G. Then $G^0_{\delta^r}$ is normal and δ^r-type-definable of bounded index. Moreover, $G^0_{\delta^r}/G^{00}_{\delta^r}$ is the connected component of the identity in $G/G^{00}_{\delta^r}$.*

(d) (Uniqueness of measure) *The measure μ is the unique left-invariant finitely additive probability measure on the Boolean algebra of δ^r-formulas.*

(e) (Generic compact domination) *Given a δ^r-formula $\phi(x)$, define $E_{\phi(x)}\subseteq G/G^{00}_{\delta^r}$ to be the (closed) set of $C\in G/G^{00}_{\delta^r}$ such that $p\models C\cap\phi(x)$ and $q\models C\cap\neg\phi(x)$ for some generic δ^r-types p,q over G. Then $E_{\phi(x)}$ has Haar measure 0.*

Remark 1.2. The reader familiar with [1], [6], [8], and [9] will notice strong similarities between the theorem above and properties of groups definable in NIP theories, especially those with *finitely satisfiable generics* (see Remark 1.3). Therefore, it is worth emphasizing that although these sources provide an invaluable guide for the proof of Theorem 1.1 (as detailed below), the results above are not simply obtained by direct translation from the global NIP setting. In particular, the formula $\delta^r(x;\bar{y})$ need not be NIP (see Example 3.7), and so we are instead forced to work with certain families of (non-uniformly) NIP formulas given by instances of δ^r. For this reason, many proof techniques from the above sources must either be redone from scratch, or replaced with a more combinatorial method that does not rely on as many model theoretic tools.

Remark 1.3. The essential use of pseudofiniteness is in our work is in deriving local analogues of definability and finite satisfiability of the pseudofinite counting measure, with respect to certain families of NIP formulas. Indeed, our results hold in a more general setting of "local fsg" for NIP formulas, and so we take the opportunity here to make this explicit. Recall from that a group G, definable in a saturated model an NIP theory, has *finitely satisfiable generics* if and only if it admits a left-invariant finitely additive measure on definable sets, which is *generically stable* (i.e. definable and finitely satisfiable) over a small model (see [15, Proposition 8.33]). Moreover, this measure is right-invariant and unique. A suitable localization of this behavior is as follows.

Suppose G is a saturated expansion of a group and $\delta(x;\bar{y})$ is an invariant NIP formula. Call a formula $\theta(x;\bar{y})$ a *stabilizing formula* if it is of the form $\phi(x \cdot y)$, $\phi(y \cdot x)$, $\phi(y_1 \cdot x) \vartriangle \phi(y_2 \cdot x)$, or $\phi(x \cdot y_1) \vartriangle \phi(x \cdot y_2)$ for some δ^r-formula $\phi(x)$. For the purposes of this remark, we say that $\delta(x;\bar{y})$ is *fsg* if there is a left and right invariant measure μ on the Boolean algebra of δ^r-definable sets such that, for any stabilizing formula $\theta(x;y_1,y_2)$ and any $\epsilon > 0$, the following holds:

- (i) *(local finite satisfiability)* there is $F \subseteq G$ finite such that, for any $b_1, b_2 \in G$, if $\mu(\theta(x;b_1,b_2)) > \epsilon$ then $\theta(x;b_1,b_2)$ is realized in F, and
- (ii) *(local definability)* the set $\{(b_1,b_2) \in G^2 : \mu(\theta(x;b_1,b_2)) \leq \epsilon\}$ is θ^{opp}-type-definable.

In this case, the measure μ satisfies the key properties demonstrated for pseudofinite counting measures in Section 2. Altogether, if G is as above and $\delta(x;\bar{y})$ is invariant, NIP, and fsg, then Theorem 1.1 holds, with μ in place of the pseudofinite counting measure.

Although we do not pursue any examples of this behavior beyond NIP formulas in pseudofinite groups, it is worth pointing out that if G is an *fsg* group definable in an NIP theory, then any invariant formula $\delta(x;\bar{y})$ in G is (locally) fsg. In other words, our work shows that the properties of NIP fsg groups can be localized around objects defined by controlled instances of a particular formula.

Remark 1.4. The initial motivation for the work in this paper was toward the subject of arithmetical regularity in finite groups (first developed by Green [4] in the abelian setting). In analogy to the various strengthened regularity lemmas for graphs definable in model theoretically tame (e.g. stable or NIP) contexts, there has been a flurry of recent interest in strengthened arithmetic regularity lemmas for similarly "tame" subsets of finite groups. In [3] (also with C. Terry), we use Theorem 1.1, and generic compact domination in particular, to prove arithmetic regularity lemmas for "VC-sets" in finite groups (i.e., sets whose family of left translates has absolutely bounded VC-dimension). Further detail is given at the end of this section.

We now give a brief summary of the paper and compare the various aspects of the above theorem to previous work on groups definable in NIP theories.

Section 2 contains preliminary observations on the pseudofinite setting above. In Section 2.2, we show that the VC-theorem transfers naturally to pseudofinite structures and, as a consequence, the pseudofinite counting measure is definable and finitely satisfiable when restricted to NIP formulas.

In Section 3, we construct $G^{00}_{\delta^r}$ by hand using stabilizers of formulas and generic types. This section contains proofs of parts (a) and (b) of Theorem 1.1 (see Proposition 3.12 and Theorem 3.15, respectively). These results are local versions, in this pseudofinite setting, of previous results on *fsg* groups in NIP theories. Namely, if G is definable in an NIP theory, then the type-definable connected component G^{00} exists, and is the intersection of all type-definable bounded-index subgroups of G. If G is also *fsg*, then generic types exist and left and right genericity coincides. Moreover, in this case there is a left-invariant Keisler measure μ on G, which is definable and finitely satisfiable in some (any) small model. See [8], [9].

In Section 4, we first “localize” the standard logic topology on G/Γ, where Γ is type-definable of bounded index, and use this to prove part (c) of Theorem 1.1 (see Corollary 4.5). This is a local analog of the fact that, for a group G definable in an NIP theory, the definable connected component G^0 exists, and is the intersection of all definable finite-index subgroups of G. Moreover, in this case, G^0/G^{00} is the connected component of the identity in G/G^{00}.

We prove part (d) of Theorem 1.1 in Section 5 (see Theorem 5.9), and part (e) in Section 6 (see Theorem 6.2). These are local analogs of the fact that, for a group G definable in an NIP theory, if G is *fsg* then there is a unique left-invariant Keisler measure on G, and if G is definably amenable (e.g. if G is *fsg*) then generic compact domination holds. This latter fact first appears in [9], with some errors in the proof, and was eventually given a correct proof by Simon in [17]. Our work in Sections 5 and 6 relies heavily on results of Simon from [16] and [17], and also involves local versions of several proofs in the work of Chernikov and Simon on definably amenable NIP groups [1].

The study of generic compact domination (and its stronger relative “compact domination”) originates from the Pillay conjectures on groups definable in o-minimal theories (see [6], [13]). It is rather remarkable that generic compact domination describes, in an infinite setting, the underlying qualitative mechanics of regularity lemmas in model theoretically tame environments, especially arithmetic regularity in the context of finite groups. In particular, given a finite group G and a suitably tame (e.g. stable or NIP) set $A \subseteq G$, the strongest kind of arithmetic regularity lemma would produce a normal subgroup H, whose index is uniformly bounded in some way, such that almost all cosets of H are “strongly regular for A”, i.e. are almost entirely contained in A or almost entirely disjoint from A (see [2] for a precise account in the stable context). In the above setting of pseudofinite groups (which arise when proving regularity for finite groups via ultraproducts), generic compact domination says that if $A \subseteq G$ is suitably NIP (e.g. defined by a δ^r-formula as above), almost all cosets of G^{00} are strongly regular for A. This generalizes the behavior of subsets of groups definable by stable formulas, in which case $G^{00} = G^0$, and *all* cosets of G^0 are strongly regular for A, leading to a structural description of A as a union of cosets of G^0 (see [2]). In the NIP setting, the use of generic compact domination, to deduce results about finite groups using ultraproduct,s requires a good deal of further work, which we carry out in [3] with C. Terry.

2. Preliminaries

2.1. **Set systems and VC-dimension.** In this section, we briefly state the basic definitions and main results on VC-dimension. Further details can be found in [15], for example.

A *set system* is a pair $(X,\mathcal{S})$ where X is a set and $\mathcal{S}\subseteq\mathcal{P}(X)$.

Definition 2.1. Let $(X,\mathcal{S})$ be a set system.

(1) The **shatter function** of $(X,\mathcal{S})$ is $\pi_{(X,\mathcal{S})}\colon\mathbb{N}\to\mathbb{N}$ such that
$$\pi_{(X,\mathcal{S})}(n)=\max\{|A\cap\mathcal{S}| : A\subseteq X,\ |A|=n\},$$
where, given $A\subseteq X$, $A\cap\mathcal{S}=\{A\cap Y : Y\in\mathcal{S}\}$.

(2) The **VC-dimension** of $(X,\mathcal{S})$ is
$$\sup\{n\in\mathbb{N} : \pi_{(X,\mathcal{S})}(n)=2^n\}\in\mathbb{N}\cup\{\infty\}.$$

Fact 2.2 (Sauer-Shelah Lemma). *For all $k\geq 1$ there is $c=c(k)$ such that, if $(X,\mathcal{S})$ is a set system of VC-dimension k, then $\pi_{(X,\mathcal{S})}(n)\leq cn^k$ for all $n\geq 0$.*

Given a set X, a unary relation U on X, and a tuple $(a_1,\ldots,a_n)\in X^n$, define
$$\mathbb{E}(a_1,\ldots,a_n;U):=\tfrac{1}{n}|\{i\in[n] : U(a_i)\text{ holds}\}|.$$
For finite sets X, we let μ_X denote the normalized counting measure on X.

Fact 2.3 (VC-Theorem). *For any $k\geq 1$ and $\epsilon>0$ there is $r=r(k,\epsilon)$ such that the following holds. Suppose X is a finite set and $(X,\mathcal{S})$ is a set system with VC-dimension at most k. Then there are (not necessarily distinct) $x_1,\ldots,x_r\in X$ such that $|\mu_X(Y)-\mathbb{E}(x_1,\ldots,x_r;Y)|<\epsilon$ for any $Y\in\mathcal{S}$.*

The sequence $(x_1,\ldots,x_r)$ in the VC-Theorem is often called an ϵ-*approximation* for the set system $(X,\mathcal{S})$, and the set $\{x_1,\ldots,x_r\}$ is an ϵ-*net* for $(X,\mathcal{S})$.

Remark 2.4. Several of the following results will yields various bounds, depending on some k and ϵ, which are explicit in terms of $r(k,\epsilon)$ in the VC-Theorem. So it is worth noting that $r(k,\epsilon)$ is $O(k\epsilon^{-2}\log(\epsilon^{-1}))$ (see [5], [10]).

2.2. **NIP formulas in pseudofinite structures.** Let $\mathcal{L}$ be a first-order language. In preparation for working with pseudofinite $\mathcal{L}$-structures, we expand $\mathcal{L}$ to a language $\mathcal{L}^+$ containing a new sort $\mathcal{I}$, on which there is a binary relation $<$ and a binary function $d(x,y)$. In any finite $\mathcal{L}$-structure, we interpret $\mathcal{I}$ as $[0,1]$ and $d(x,y)$ as the standard distance on $[0,1]$ (we will also write $|x-y|$ for $d(x,y)$). For every $\mathcal{L}$-formula $\phi(\bar{x};\bar{y})$, we add to $\mathcal{L}^+$ a $\bar{y}$-ary function symbol $\mu_\phi(\bar{y})$ into $\mathcal{I}$. In any finite $\mathcal{L}$-structure A, $\mu_\phi(\bar{y})$ is interpreted as $\mu_A(\phi(A^{\bar{x}},\bar{y}))$.

Let M be a fixed, sufficiently saturated elementary extension of an ultraproduct of finite $\mathcal{L}^+$-structures (which are canonically expanded from $\mathcal{L}$-structures as described above). By convention, *formulas* will always be in the language $\mathcal{L}$ and always allow parameters from M. We will use $\mathcal{L}$-*formula* to specify formulas with no extra parameters. We let μ denote the pseudofinite counting measure on M. Specifically, given an $\mathcal{L}$-formula $\phi(\bar{x},\bar{y})$ and $\bar{b}\in M^{\bar{y}}$, $\mu(\phi(\bar{x},\bar{b}))$ is defined as the standard part of $\mu_\phi(\bar{b})$. It is routine to verify that μ is a finitely additive probability measure on (powers of) M.

Definition 2.5. Given $k\geq 1$, a formula $\phi(\bar{x},\bar{y})$ is k-**NIP** if there do not exist sequences $(\bar{a}_i)_{i\in[k]}$ in $M^{\bar{x}}$ and $(\bar{b}_X)_{X\subseteq[k]}$ in $M^{\bar{y}}$ such that $M\models\phi(\bar{a}_i,\bar{b}_X)$ if and only if $i\in X$. A formula $\phi(\bar{x};\bar{y})$ is **NIP** if it is k-NIP for some $k\geq 1$.

Remark 2.6. A formula $\phi(\bar{x},\bar{y})$ is k-NIP if and only if the set system $(M^{\bar{x}},\{\phi(M^{\bar{x}},\bar{b}) : \bar{b}\in M^{\bar{y}}\})$ has VC-dimension at most $k-1$.

The rest of this section contains several corollaries of the VC-Theorem for pseudofinite structures. Roughly speaking, the VC-Theorem says that, restricted to set systems of finite VC-dimension, counting measures on finite sets are approximated by averages of points. We now observe that this immediately implies the same statement for the pseudofinite counting measure on M.

Corollary 2.7. *For any $k \geq 1$ and $\epsilon > 0$, there is $r = r(k,\epsilon)$ such that the following holds. Suppose $\phi(\bar{x};\bar{y})$ is a k-NIP formula. Then there are (not necessarily distinct) $\bar{a}_1,\ldots,\bar{a}_r \in M^{\bar{x}}$ such that, for any $\bar{b} \in M^{\bar{y}}$,*

$$\left|\mu(\phi(\bar{x};\bar{b})) - \mathbb{E}(\bar{a}_1,\ldots,\bar{a}_r;\phi(\bar{x};\bar{b}))\right| \leq \epsilon.$$

In particular, if $\mu(\phi(\bar{x};\bar{b})) > \epsilon$ then $\phi(\bar{x};\bar{b})$ is realized in $F(\phi,\epsilon) = \{\bar{a}_1,\ldots,\bar{a}_t\}$.

Proof. Fix k and ϵ and let $r(k,\epsilon)$ be as in the VC-Theorem. Let $\phi(\bar{x};\bar{y},\bar{z})$ be an $\mathcal{L}$-formula, and let $\chi(\bar{z})$ be an $\mathcal{L}$-formula expressing that $\phi(\bar{x};\bar{y},\bar{z})$ is k-NIP as a relation in $\bar{x}$ and $\bar{y}$. By the VC-theorem, if A is a finite $\mathcal{L}$-structure then

$$A \models \forall \bar{z}\,(\chi(\bar{z}) \rightarrow \exists \bar{x}_1 \ldots \bar{x}_r\, \forall \bar{y}\ |\mu_\phi(\bar{y}) - \mathbb{E}(\bar{x}_1,\ldots,\bar{x}_r;\phi(\bar{x},\bar{y},\bar{z}))| < \epsilon)$$

(where the expression on the right is an $\mathcal{L}^+$-sentence). Therefore, by Łoś's Theorem and elementarity, M satisfies this sentence, which yields the desired result. □

Corollary 2.8. *Suppose M is pseudofinite. Let $\Delta = \{\phi_i(\bar{x};\bar{y}_i) : i \in I\}$ be a collection of NIP formulas. Then there is $M_0 \preceq M$, of size at most $|I| + \aleph_0$, such that for any $i \in I$ and $\bar{b} \in M^{\bar{y}_i}$, if $\mu(\phi_i(\bar{x};\bar{b})) > 0$ then $\phi_i(\bar{x};\bar{b})$ is realized in M.*

Proof. Let $M_0 \preceq M$ be any model, of size at most $|I| + \aleph_0$, which contains the set $F(\phi_i,\epsilon)$ from Corollary 2.7 for all $i \in I$ and rational $\epsilon > 0$. □

Definition 2.9. Let $\phi(\bar{x};\bar{y})$ be a formula.

(1) An **instance of** $\phi(\bar{x};\bar{y})$ is a formula $\phi(\bar{x};\bar{b})$ or $\neg\phi(\bar{x};\bar{b})$, where $\bar{b} \in M^{\bar{y}}$.
(2) A **ϕ-formula** is a finite Boolean combination of instances of $\phi(\bar{x};\bar{y})$.
(3) A set $X \subseteq M^{\bar{x}}$ is **ϕ-definable** if it is defined by a ϕ-formula.
(4) A set $X \subseteq M^{\bar{x}}$ is **ϕ-type-definable** if it is defined by an intersection of boundedly many ϕ-formulas.
(5) Let $\phi^{\text{opp}}(\bar{y};\bar{x})$ denote $\phi(\bar{x};\bar{y})$.

Corollary 2.10. *Suppose M is pseudofinite, and fix an NIP formula $\phi(\bar{x};\bar{y})$. Then there is a countable set $A \subset M$ such that, for any closed $C \subseteq [0,1]$, the set*

$$\{\bar{b} \in M^{\bar{y}} : \mu(\phi(\bar{x};\bar{b})) \in C\}$$

is ϕ^{opp}-type-definable over A.

Proof. Given $n > 0$, we have $r_n \in \mathbb{N}$ and $\bar{a}^n_1,\ldots,\bar{a}^n_{r_n} \in M^{\bar{x}}$ such that, for all $\bar{b} \in M^{\bar{y}}$, $|\mu(\phi(\bar{x},\bar{b})) - \mathbb{E}(\bar{a}^n_1,\ldots,\bar{a}^n_{r_n};\phi(\bar{x};\bar{b}))| \leq \frac{1}{n}$. Define

$$X_n = \left\{\bar{b} \in M^{\bar{y}} : d(\mathbb{E}(\bar{a}^n_1,\ldots,\bar{a}^n_{r_n};\phi(\bar{x};\bar{b})),C) \leq \tfrac{1}{n}\right\}.$$

Then, since C is closed, it follows that $\{\bar{b} \in M^{\bar{y}} : \mu(\phi(\bar{x};\bar{b})) \in C\} = \bigcap_{n>0} X_n$. So it suffices to show that each X_n is ϕ^{opp}-definable over $A_n = \bigcup_{i=1}^{r_n} \bar{a}^n_i$. Fix $n > 0$ and, for $I \subseteq [r_n]$, define the formula

$$\theta_I(\bar{y}) := \bigwedge_{i\in I} \phi(\bar{a}^n_i;\bar{y}) \wedge \bigwedge_{i\in[r_n]\setminus I} \neg\phi(\bar{a}^n_i;\bar{y}).$$

Then $\theta_I(\bar{y})$ is a ϕ^{opp}-formula over A_n. Set $\mathcal{F} = \{I \subseteq [r_n] : d(\frac{|I|}{n}, C) < \frac{1}{n}\}$. Then X_n is defined by $\bigvee_{I \in \mathcal{F}} \theta_I(\bar{y})$. □

2.3. **NIP formulas and generic sets in pseudofinite groups.**

Definition 2.11. Let G be a group. Given $n \geq 1$, set $A \subseteq G$ is **left n-generic** (resp. **right n-generic**) if there are n left translates (resp. right translates) of A whose union is G. We say $A \subseteq G$ is **left generic** (resp. **right generic**) if it is left n-generic (resp. right n-generic) for some $n \geq 1$.

We now assume that $\mathcal{L}$ expands the language of groups, and we let G be a fixed, sufficiently saturated $\mathcal{L}$-structure which is an elementary extension of an ultraproduct of finite groups. Note that the pseudofinite counting measure μ on G is left and right invariant.

Definition 2.12. Let $\phi(x)$ be a formula.

1. Let $\phi^\ell(x;y)$ denote the formula $\phi(y \cdot x)$.
2. Let $\phi^r(x;y)$ denote the formula $\phi(x \cdot y)$.

Given a formula $\phi(x)$, note that $\phi^r(x;y) = (\phi^\ell)^{\text{opp}}(x;y)$. In particular, $\phi^\ell(x;y)$ is NIP if and only if $\phi^r(x;y)$ is NIP.

Corollary 2.13. *For any $k \geq 1$ and $\epsilon > 0$ there is $n = n(k,\epsilon)$ such that, for any formula $\phi(x)$, if $\phi^\ell(x;y)$ is k-NIP and $\mu(\phi(x)) > \epsilon$, then $\phi(x)$ is left n-generic and right n-generic.*

Proof. Fix $k \geq 1$ and $\epsilon > 0$. Let $n = \max\{r(k,\epsilon), r(2^k,\epsilon)\}$ be given by Corollary 2.7. Suppose $\phi(x)$ is a formula such that $\phi^\ell(x;y)$ is k-NIP, and assume $\mu(\phi(x)) > \epsilon$. Then $\mu(\phi(bx)) > \epsilon$ for any $b \in G$ by invariance of μ. By Corollary 2.7, there is $F \subset G$, of size at most n, such that $\phi(bx)$ is realized in F for any $b \in G$. So the right translates of $\phi(x)$ by elements in F^{-1} cover G, i.e. $\phi(x)$ is right n-generic. By choice of n and the same argument applied to $\phi^r(x;y)$ (which is 2^k-NIP), we see that $\phi(x)$ is left n-generic. □

Corollary 2.14. *Let $\phi(x)$ be a formula such that $\phi^\ell(x;y)$ is NIP. The following are equivalent:*

(i) *$\phi(x)$ is left generic;*
(ii) *$\phi(x)$ is right generic;*
(iii) *$\mu(\phi(x)) > 0$.*

Proof. $(i) \Rightarrow (iii)$ and $(ii) \Rightarrow (iii)$ are by invariance and finite additivity of μ. $(iii) \Rightarrow (i)$ and $(iii) \Rightarrow (ii)$ are by Corollary 2.13. □

In light of the previous corollary, we will just say $\phi(x)$ is *generic* (or *n-generic*), in the case that $\phi^\ell(x;y)$ is NIP and $\mu(\phi(x)) > 0$.

Corollary 2.15. *Let $\phi(x)$ be a formula such that $\phi^\ell(x;y)$ is NIP. Then at least one of $\phi(x)$ or $\neg\phi(x)$ is generic.*

Proof. At least one of $\phi(x)$ or $\neg\phi(x)$ must have positive μ-measure. □

3. Stabilizers and G^{00}

Throughout this section, and for the rest of the paper, we continue to work with a sufficiently saturated pseudofinite $\mathcal{L}$-structure G expanding a group.

3.1. **Stabilizers of formulas.**

Definition 3.1. Let $\phi(x)$ be a formula.

(1) Given $\epsilon \geq 0$, define

$$\mathrm{Stab}^{\epsilon}_{\mu}(\phi(x)) = \{g \in G : \mu(\phi(g^{-1}x) \triangle \phi(x)) \leq \epsilon\}.$$

(2) Define $\mathrm{Stab}_{\mu}(\phi(x)) = \mathrm{Stab}^{0}_{\mu}(\phi(x)) = \{g \in G : \mu(\phi(g^{-1}x) \triangle \phi(x)) = 0\}$.

Proposition 3.2. *Suppose $\phi(x)$ is a formula such that $\phi^{\ell}(x;y)$ is NIP. Then, for any $\epsilon > 0$, $\mathrm{Stab}^{\epsilon}_{\mu}(\phi(x))$ is left generic and ϕ^{r}-type-definable over a countable parameter set.*

Proof. Let $\psi(x;y_1,y_2)$ denote $\phi(y_1 \cdot x) \triangle \phi(y_2 \cdot x)$, and note that $\psi(x;y_1,y_2)$ is NIP. By Corollary 2.7, we may fix a finite set $F \subset G$ such that, for any $b_1, b_2 \in G$, if $\mu(\psi(x;b_1,b_2)) > \epsilon$ then $\psi(x;b_1,b_2)$ is realized in F. Define an equivalence relation $\sim$ on G such that $g \sim h$ if and only if $F \cap g\phi(G) = F \cap h\phi(G)$. Then $\sim$ has finitely many classes and so we may pick representatives $g_1, \ldots, g_n$. Let $X = \mathrm{Stab}^{\epsilon}_{\mu}(\phi(x))$. We show $G = g_1X \cup \ldots \cup g_nX$. Fix $h \in G$. Then $h \sim g_i$ for some $1 \leq i \leq n$. It follows that $\psi(x;h^{-1},g_i^{-1})$ is not realized in F, and so

$$\mu(\phi(h^{-1}g_i x) \triangle \phi(x)) = \mu(\psi(x;h^{-1},g_i^{-1})) \geq \epsilon.$$

Therefore $g_i^{-1}h \in X$, and so $h \in g_iX$, as desired.

Finally, let $\theta(x;y)$ denote $\phi(y \cdot x) \triangle \phi(x)$, which is NIP. We have

$$\mathrm{Stab}^{\epsilon}_{\mu}(\phi(x)) = \{g \in G : \mu(\theta(x;g)) \leq \epsilon\},$$

and so $\mathrm{Stab}^{\epsilon}_{\mu}(\phi(x))$ is θ^{opp}-type-definable over a countable parameter set by Corollary 2.10. Since any instance of $\theta^{\mathrm{opp}}(y;x)$ is equivalent to an instance of $\phi^{r}(x;y)$, we have the desired result. ☐

Remark 3.3. Let $\phi(x)$ and $\epsilon > 0$ be as in the proof of Proposition 3.2. Note that if $\phi^{\ell}(x;y)$ is k-NIP, and π denotes the shatter function for $(G, \{\phi(gx) : g \in G\})$, then $\sim$ has at most $\pi(r(k,\epsilon))$ classes, where $r(k,\epsilon)$ is given by Corollary 2.7. By the Sauer-Shelah Lemma and Remark 2.4, $\mathrm{Stab}^{\epsilon}_{\mu}(\phi(x))$ is n-generic with $n \leq \epsilon^{-O_k(1)}$.

Corollary 3.4. *Suppose $\phi(x)$ is a formula such that $\phi^{\ell}(x;y)$ is NIP. Then $\mathrm{Stab}_{\mu}(\phi(x))$ is a subgroup of G of bounded index, which is ϕ^{r}-type-definable over a countable parameter set.*

Proof. Using invariance and finite additivity of μ, it is straightforward to check that $\mathrm{Stab}_{\mu}(\phi(x))$ is a subgroup of G. By definition, $\mathrm{Stab}_{\mu}(\phi(x)) = \bigcap_{\epsilon \in \mathbb{Q}^{+}} \mathrm{Stab}^{\epsilon}_{\mu}(\phi(x))$. By Proposition 3.2, each set in this intersection is generic and ϕ^{r}-type-definable over a countable parameter set. Therefore $\mathrm{Stab}_{\mu}(\phi(x))$ has bounded index and is ϕ^{r}-type-definable over a countable parameter set. ☐

Given a formula $\phi(x)$, the formula $\phi^{\ell}(x;y)$ is *invariant* in the sense that any left translate of an instance of $\phi^{\ell}(x;y)$ is also an instance of $\phi^{\ell}(x;y)$. We want to work with the general class of formulas satisfying this property.

Definition 3.5. A formula $\delta(x;\bar{y})$ is **(left) invariant** if, for any $a, \bar{b} \in G$, there is $\bar{c} \in G$ such that $\delta(ax;\bar{b})$ is equivalent to $\delta(x;\bar{c})$.

The main reason to work with invariant $\mathcal{L}$-formulas is so that we have a well-defined action by G on the space of δ-types (defined below). However, given a formula $\delta(x;\bar{y})$, which is invariant and NIP, it will be necessary to consider right translates of δ-formulas in order to pinpoint type-definability at various steps of the subsequent work (as suggested by Proposition 3.2). Therefore, we set the following notation.

Definition 3.6. Given a formula $\delta(x;\bar{y})$, let $\delta^r(x;\bar{y},u)$ denote the formula $\delta(x{\cdot}u;y)$.

Note that if an invariant $\mathcal{L}$-formula $\delta(x;\bar{y})$ is also *right* invariant (e.g. if G is abelian), then $\delta^r(x;\bar{y},u)$ is essentially the same as $\delta(x;\bar{y})$. However, in general, $\delta^r(x;\bar{y},u)$ may behave quite differently. Most importantly, $\delta^r(x;\bar{y})$ may be NIP, while $\delta^r(x;\bar{y},u)$ is not, as demonstrated by the following example.

Example 3.7. Given $k\in\mathbb{N}$, let G_k be the group of permutations of $\{1,\dots,k+1\}$, and let H_k be the subgroup of permutations fixing 1. Then, with G_k as the ambient structure, the formula $yx\in H_k$ is 2-stable (and thus 2-NIP) since H_k is a subgroup. But $yxy\in H_k$ is not k-NIP. To see this, let $X=\{2,\dots,k+1\}$. Given $n\in X$ and $I\subseteq X$, let $a_n\in G_k$ be the transposition $(1\ n)$, and let $b_I\in G_k$ be a permutation whose set of fixed points in X is precisely I (such a permutation always exists since $1\notin X$). Then, given $n\in X$ and $I\subseteq X$, $a_nb_Ia_n\in H_k$ if and only if $n\in I$.

Now let $\mathcal{U}$ be a nonprincipal ultrafilter on $\mathbb{N}$ and let $G=\prod_{\mathcal{U}}G_k$. If $A=\prod_{\mathcal{U}}H_k$, then $yx\in A$ is stable, while $yxy\in A$ has the independence property.

Despite the behavior seen in the last example, we will still recover sufficiently good behavior for instances of the formula $\delta^r(x;\bar{y},u)$ (see, e.g., Proposition 3.12).

The next goal is to define the local analog of G^{00}. We will first give an explicit construction using "measure-stabilizers" of formulas, and then show that the object obtained behaves as expected (see Theorem 3.15).

Definition 3.8. Let $\delta(x;\bar{y})$ be a formula. Define

$$G^*_\delta=\bigcap_{\bar{a}\in G^{\bar{y}}}\mathrm{Stab}_\mu(\delta(x;\bar{a})).$$

Lemma 3.9. *If $\delta(x;\bar{y})$ is invariant and NIP, then there is a bounded set $A\subseteq G^{\bar{y}}$ such that $G^*_\delta=\bigcap_{\bar{a}\in A}\mathrm{Stab}_\mu(\delta(x;\bar{a}))$.*

Proof. Define an equivalence relation $\sim$ on $G^{\bar{y}}$ such that $\bar{a}\sim\bar{b}$ if and only if $\mu(\delta(x;\bar{a})\triangle\delta(x;\bar{b}))$. To find the desired set A, it suffices to show that $\bar{a}\sim\bar{b}$ implies $\mathrm{Stab}_\mu(\delta(x;\bar{a}))=\mathrm{Stab}_\mu(\delta(x;\bar{b}))$, and that $\sim$ has a bounded number of classes.

For the first claim, fix $\bar{a},\bar{b},g\in G$. The formula $\delta(g^{-1}x;\bar{a})\triangle\delta(x;\bar{b})$ implies

$$(\delta(g^{-1}x;\bar{b})\triangle\delta(g^{-1}x;\bar{a}))\vee(\delta(g^{-1}x;\bar{a})\triangle\delta(x;\bar{a}))\vee(\delta(x;\bar{a})\triangle\delta(x;\bar{b})).$$

So if $\bar{a}\sim\bar{b}$ and $g\in\mathrm{Stab}_\mu(\delta(x;\bar{a}))$, then $g\in\mathrm{Stab}_\mu(\delta(x;\bar{b}))$ by invariance and finite additivity of μ.

The second claim is standard fact about NIP formulas (details are included for the sake of clarity). If $\sim$ has unboundedly many classes then by Erdős-Rado there is an indiscernible sequence $(\bar{b}_i)_{i<\omega}$, and some $\epsilon>0$, such that $\mu(\delta(x;\bar{b}_i)\triangle\delta(x;\bar{b}_j))\geq\epsilon$ for all $i\neq j$. Then $\{\delta(x;\bar{b}_{2i})\triangle\delta(x;\bar{b}_{2i+1}):i<\omega\}$ is consistent by [6, Lemma 2.8]. This contradicts that $\delta(x;\bar{y})$ is NIP and thus has finite alternation number (e.g. [14, Theorem 12.17]). $\square$

From Corollary 3.4 and Lemma 3.9, we immediately obtain the following result.

Corollary 3.10. *If $\delta(x;\bar{y})$ is invariant and NIP then G^*_δ is a δ^r-type-definable bounded-index subgroup of G.*

3.2. **Stabilizers of types.**

Definition 3.11. Fix an invariant formula $\delta(x;\bar{y})$.

(1) Given $A \subseteq G$, let $S_\delta(A)$ denote the space of **complete δ-types** (i.e. maximal consistent sets of instances of δ) with parameters from A.
(2) Given $p \in S_\delta(G)$, let $\operatorname{Stab}(p) = \{g \in G : gp = p\}$ (where $gp = \{\phi(g^{-1}x) : \phi(x) \in p\}$).
(3) A δ-type p is **left generic** (resp. **right generic**) if every formula in p is left generic (resp. right generic).

Proposition 3.12. *Suppose $\delta(x;\bar{y})$ is invariant and NIP.*

(a) If $\phi(x)$ is a δ^r-formula then $\phi^\ell(x;y)$ is NIP.
(b) Given $p \in S_{\delta^r}(G)$, the following are equivalent:
  (i) p is left generic;
  (ii) p is right generic;
  (iii) $\mu(\phi(x)) > 0$ for all $\phi(x) \in p$.
(c) The space of (left) generic types in $S_{\delta^r}(G)$ is nonempty and invariant under left and right multiplication.

Proof. Part (a). Fix $k \geq 1$ such that $\delta(x;\bar{y})$ is k-NIP. We first claim that, for any $\bar{b}, c \in G$, if $\phi(x)$ denotes $\delta^r(x;\bar{b},c)$, then $\phi^\ell(x;y)$ is k-NIP. To see this, suppose we have $(r_i)_{i\in[n]}$ and $(s_I)_{I\subseteq[n]}$ such that $\delta^r(s_I r_i;\bar{b},c)$ holds if and only if $i \in I$. For any I, there is $\bar{a}_I$ such that $\delta(s_I \cdot x;\bar{b})$ is equivalent to $\delta(x;\bar{a}_i)$. So, setting $g_i = r_i c$, we have $\delta(g_i;\bar{a}_I)$ if and only if $i \in I$. So $n < k$. Part (a) now follows by induction on the construction of δ^r-formulas.

Part (b). This follows from part (a) and Corollary 2.14.

Part (c). By finite additivity of μ, the measure 0 sets form an ideal, and so there are types $p \in S_{\delta^r}(G)$ satisfying condition (iii) of part (b). So the claims follow from parts (a) and (b). $\square$

Given an invariant NIP formula $\delta(x;\bar{y})$ and a δ^r-type p, we will just call p *generic* in case it is left generic (equivalently right generic).

The following are some technical observations that will be needed in the proof of Theorem 3.15.

Proposition 3.13. *Suppose $\delta(x;\bar{y})$ is invariant.*

(a) $G^*_\delta \subseteq \operatorname{Stab}_\mu(\phi(x))$ for any δ^r-formula $\phi(x)$.
(b) $G^*_\delta = G^*_{\delta^r}$.
(c) For any $p \in S_{\delta^r}(G)$, there is a unique right coset C of G^*_δ such that $p \models C$.

Proof. Part (a). Given δ^r-formulas $\phi(x)$ and $\psi(x)$, we have $\operatorname{Stab}_\mu(\neg\phi(x)) = \operatorname{Stab}_\mu(\phi(x))$ and $\operatorname{Stab}_\mu(\phi(x)) \cap \operatorname{Stab}_\mu(\psi(x)) \subseteq \operatorname{Stab}_\mu(\phi(x) \wedge \psi(x))$. So the claim follows by induction on the construction of δ^r-formulas.

Part (b). By definition, $G^*_{\delta^r} \subseteq G^*_\delta$. For the other containment, fix $g \in G^*_\delta$ and $\bar{b}, c \in G$. By right invariance of μ, and since $g \in \operatorname{Stab}_\mu(\delta(x;\bar{b}))$, we have

$$\mu(\delta^r(g^{-1}x;\bar{b},c) \triangle \delta^r(x;\bar{b},c)) = \mu(\delta(g^{-1}x;\bar{b}) \triangle \delta(x;\bar{b})) = 0,$$

and so $g \in \operatorname{Stab}_\mu(\delta^r(x;\bar{b},c))$.

Part (c). Note that all right cosets of G^*_δ are δ^r-type-definable, since G^*_δ is δ^r-type-definable and δ^r-formulas are right invariant. Since any complete δ^r-type concentrates on at most one right coset of G^*_δ, it suffices to show that every complete δ^r-type concentrates on some right coset of G^*_δ. Since G^*_δ is δ^r-type-definable of bounded index, we may fix a small model $M \prec G$ such that all right cosets of G^*_δ are δ^r-type-definable over M. Now, given $p \in S_{\delta^r}(G)$, if $a \in G$ realizes $p|_M$, then p concentrates on $G^*_\delta a$. □

Definition 3.14. Fix a formula $\delta(x;\bar{y})$.

(1) Let $S^g_\delta(G)$ denote the set of generic δ-types in $S_\delta(G)$.
(2) Given $p \in S_\delta(G)$, define $\mathrm{Stab}(p) = \{g \in G : gp = p\}$.
(3) Let G^{00}_δ denote the intersection of all δ-type-definable bounded-index subgroups of G.

Note that, for any invariant formula $\delta(x;\bar{y})$, the class of δ^r-type-definable bounded-index subgroups of G is closed under conjugation, and so $G^{00}_{\delta^r}$ is always a normal subgroup of G. The next theorem is the main result on G^*_δ, for $\delta(x;\bar{y})$ invariant and NIP.

Theorem 3.15. *Suppose $\delta(x;\bar{y})$ is invariant and NIP.*

(a) *G^*_δ is a δ^r-type-definable bounded-index subgroup of G.*

(b) $G^*_\delta = \displaystyle\bigcap_{p \in S^g_\delta(G)} \mathrm{Stab}(p) = \bigcap_{p \in S^g_\delta(G)} \bigcap_{\phi(x)\in p} \mathrm{Stab}_\mu(\phi(x))$.

(c) *If $p \in S_{\delta^r}(G)$ is generic then*

$$G^*_\delta = \mathrm{Stab}(p) = \bigcap_{\phi(x)\in p} \mathrm{Stab}_\mu(\phi(x)).$$

(d) $G^*_\delta = G^{00}_{\delta^r}$.

Proof. Part (a). This is Corollary 3.10.

Part (b). We show

$$G^*_\delta \subseteq \bigcap_{p\in S^g_\delta(G)} \bigcap_{\phi(x)\in p} \mathrm{Stab}_\mu(\phi(x)) \subseteq \bigcap_{p\in S^g_\delta(G)} \mathrm{Stab}(p) \subseteq G^*_\delta.$$

The first containment is immediate from Proposition 3.13(a). For the second containment, we fix a generic type $p \in S_\delta(G)$ and show $\bigcap_{\phi(x)\in p}\mathrm{Stab}_\mu(\phi(x)) \subseteq \mathrm{Stab}(p)$. Indeed, suppose $g \in \bigcap_{\phi(x)\in p}\mathrm{Stab}_\mu(\phi(x))$ and fix $\phi(x) \in p$. If $\phi(g^{-1}x) \notin p$ then $\phi(g^{-1}x) \triangle \phi(x) \in p$, which contradicts that p is generic and $g \in \mathrm{Stab}_\mu(\phi(x))$. So $\phi(g^{-1}x) \in p$, and thus we have $g \in \mathrm{Stab}(p)$.

For the third containment, suppose $g \notin G^*_\delta$. Then there is a δ-formula $\phi(x)$ such that $\mu(\phi(g^{-1}x) \triangle \phi(x)) > 0$, and so there is a generic type $p \in S_\delta(G)$ containing the formula $\phi(g^{-1}x) \triangle \phi(x)$. So $g \notin \mathrm{Stab}(p)$.

Part (c). Fix a generic type $p \in S_{\delta^r}(G)$. We show

$$G^*_\delta \subseteq \bigcap_{\phi(x)\in p} \mathrm{Stab}_\mu(\phi(x)) \subseteq \mathrm{Stab}(p) \subseteq G^*_\delta.$$

The first containment is immediate from parts (a) and (b) of Proposition 3.13, and the second containment is similar to part (b).

For the third containment, first fix $a \in M$ such that p concentrates on $G^*_\delta a$ (such an a exists by Proposition 3.13(c)). Now fix $g \in \mathrm{Stab}(p)$. Then $p \models gG^*_\delta a$, and so $G^*_\delta a \cap gG^*_\delta a$ is a consistent type, which is therefore realized in G. So there are $x, y \in G^*_\delta$ such that $xa = gya$, and so $g = xy^{-1} \in G^*_\delta$, as desired.

Part (d). Since any conjugate of $G^*_{\delta^r}$ is a δ^r-type-definable bounded index subgroup of G, it suffices to prove the second claim. So suppose Γ is a δ^r-type-definable subgroup of bounded index. We want to show $G^*_\delta \subseteq \Gamma$. Let $p \in S_{\delta^r}(G)$ be a generic δ^r-type concentrating on Γ, and fix $a \in G^*_\delta$. Since $G^*_\delta = \mathrm{Stab}(p)$, it follows that $ap \models \Gamma$, and so $a\Gamma = \Gamma$. $\square$

We end this section by analyzing the situation when $\delta(x; \bar{y})$ is stable.

Definition 3.16. Given a formula $\delta(x; \bar{y})$, let G^0_δ denote the intersection of all δ-definable finite-index subgroups of G.

For stable $\delta(x; \bar{y})$, the group G^0_δ is δ-definable of finite index (this follows from [7], with further detail in [2]). The next corollary explains the relationship between G^0_δ, G^{00}_δ, $G^0_{\delta^r}$, and $G^{00}_{\delta^r}$ in this case (note that $\delta^r(x; \bar{y}, u)$ need not be stable, as demonstrated by Example 3.7).

Corollary 3.17. *Assume $\delta(x; \bar{y})$ is invariant and stable. Then $G^{00}_\delta = G^0_\delta$ and $G^{00}_{\delta^r} = G^0_{\delta^r}$. Moreover, $G^0_{\delta^r}$ is the normal core of G^0_δ, and thus is δ^r-definable of finite index.*

Proof. We first claim that, for any generic δ-type $p \in S_\delta(G)$, if $p \models aG^0_\delta$ then $\mathrm{Stab}(p) = aG^0_\delta a^{-1}$. Indeed, fix $p \in S_\delta(G)$ generic and let $p \models aG^0_\delta$. If $g \in \mathrm{Stab}(p)$ then $p = gp \models gaG^0_\delta$, and so $gaG^0_\delta = aG^0_\delta$, i.e. $g \in aG^0_\delta a^{-1}$. Conversely, if $g \in aG^0_\delta a^{-1}$ then $gaG^0_\delta = aG^0_\delta$, and so $gp = p$.

By the previous claim, and parts (b) and (d) of Theorem 3.15, we conclude that $G^{00}_{\delta^r}$ is the normal core of G^0_δ, and therefore $G^{00}_{\delta^r}$ is δ^r-definable of finite index. This further implies that $G^{00}_{\delta^r} = G^0_{\delta^r}$.

It remains to show $G^{00}_\delta = G^0_\delta$. So suppose H is a δ-type-definable subgroup of G. Then $G^0_{\delta^r} = G^{00}_{\delta^r} \subseteq H$, and so H is a union of cosets of $G^0_{\delta^r}$. Since $G^0_{\delta^r}$ has finite index, H is definable. By compactness, H is δ-definable, and so $G^0_\delta \subseteq H$. $\square$

4. The local logic topology and G^0

Recall from [13, Lemma 2.7] that, if Γ is a type-definable bounded-index subgroup of G, then we have the *logic topology* on G/Γ under which G/Γ is a compact (Hausdorff) space, and a topological group when Γ is normal. In particular, $X \subseteq G/\Gamma$ is closed if $\{x \in G : x\Gamma \in X\}$ is type-definable. In this section we show that if Γ is δ-type-definable, for some invariant $\delta(x, \bar{y})$, then it suffices to consider δ-type-definable sets in the construction of the logic topology on G/Γ. Certain aspects of this are likely in the folklore, and so some proofs will be brief.

Lemma 4.1. *Fix an invariant formula $\delta(x, \bar{y})$ and suppose $\Gamma \leq G$ is δ-type-definable of bounded index. Then, for any $\mathcal{L}(G)$-formula $\phi(x)$, the set*

$$X = \{a \in G : a\Gamma \cap \phi(G) \neq \emptyset\}$$

is δ-type-definable.

Proof. First, since Γ is δ-type-definable, it follows from saturation of G that X is type-definable (*a priori*, using $\phi(x)$ and existential quantification over δ-formulas). We need to show that X is type-definable by δ-formulas. By a δ^{-1}-*formula* we mean a formula of the form $\phi(x^{-1})$ where $\phi(x)$ is a δ-formula. A δ^{-1}*-type* is a small consistent set of δ^{-1}-formulas.

Note that Γ is δ^{-1}-type-definable since $\Gamma^{-1} = \Gamma$. Since Γ has bounded index, we may fix a sequence $(p_i)_{i<\kappa}$ of left translates of δ^{-1}-types, with κ small, such that any coset of Γ is definable by some p_i, and p_0 is a δ^{-1}-type defining Γ. Let $A \subset G$ be a small set such that each p_i is over A and $\phi(x)$ is over A. Let $M \prec G$ be a small $|A|^+$-saturated model. Now let $S = \{p \in S_\delta(M) : X \cap p(G) = \emptyset\}$. We show

$$(\dagger) \qquad X = \bigcap_{p \in S} \bigcup_{\psi(x) \in p} \neg\psi(G).$$

By saturation and type-definability of X, this will show that X is δ-type-definable.

By choice of S, the left-to-right containment in $(\dagger)$ is clear. So suppose $a \notin X$ and let $p = \mathrm{tp}_\delta(a/M)$. It suffices to show $p \in S$. So suppose, toward a contradiction, that we have $b \in X \cap p(G)$. Then $b\Gamma \cap \phi(G) \neq \emptyset$. In particular, if p_i is the type-definition of $b\Gamma$, then $p_i(x) \wedge \phi(x)$ is consistent, and thus realized by some $m \in M$. Then $p_0(b^{-1}m)$ holds and so, since p_0 is a δ^{-1}-type, $m^{-1} \in M$, and $b \models p$, it follows that $p_0(a^{-1}m)$ holds. But then $m \in a\Gamma \cap \phi(G)$, contradicting that $a \notin X$. $\square$

Corollary 4.2. *Fix an invariant formula $\delta(x;\bar{y})$ and suppose $\Gamma \leq G$ is a δ-type-definable of bounded index. Then $X \subseteq G/\Gamma$ is closed in the logic topology if and only if $\{a \in G : a\Gamma \in X\}$ is δ-type-definable.*

Proof. Call $X \subseteq G/\Gamma$ *δ-closed* if $\{a \in G : a\Gamma \in X\}$ is δ-type-definable. It suffices to show that the δ-closed sets define a compact Hausdorff topology on G/Γ. Indeed, given this, since the logic topology clearly refines the δ-topology, it will follow that the two topologies are the same.

The verification that the δ-closed sets generate a compact topology is exactly as in the usual case of the logic topology ([11, Lemma 3.3] or [13, Lemma 2.5]). Moreover, Lemma 4.1 is precisely what is needed to show that the standard argument of Hausdorff separation goes through. $\square$

Now, if $\delta(x;\bar{y})$ is a formula and $\Gamma \leq G$ is δ-definable of bounded index then, for any $p \in S_\delta(G)$ there is a unique left coset C of Γ such that $p \models C$. So we have a well-defined function $\pi_\Gamma \colon S_\delta(G) \to G/\Gamma$ such that $p \models \pi_\Gamma(p)$. The following conclusion is a straightforward from Corollary 4.2.

Corollary 4.3. *Fix an invariant formula $\delta(x;\bar{y})$ and suppose $\Gamma \leq G$ is δ-type-definable of bounded index. Then π_Γ is continuous.*

Remark 4.4. Lemma 4.1, Corollary 4.2, and Corollary 4.3 hold for any (sufficiently saturated) $\mathcal{L}$-structure G expanding a group (i.e. G need not be pseudofinite).

As a special case of the above situation, we can work with $\Gamma = G^{00}_{\delta^r}$, where $\delta(x;\bar{y})$ is invariant and NIP. We have already shown that, for $\delta(x,\bar{y})$ invariant and NIP, $G^{00}_{\delta^r}$ behaves like the "type-definable connected component" of G localized at the formula δ^r. Next, we show that $G^0_{\delta^r}$ fits into this picture the way one would expect from known facts in the global NIP setting.

Corollary 4.5. *Suppose $\delta(x;\bar{y})$ is invariant and NIP. Then $G^0_{\delta^r}$ is a normal δ^r-type-definable subgroup of G of bounded index, and $G^0_{\delta^r}/G^{00}_{\delta^r}$ is the connected component of the identity in $G/G^{00}_{\delta^r}$.*

Proof. Let $C \subseteq G/G^{00}_{\delta^r}$ be the connected component of the identity, and recall that C is a closed subgroup of $G/G^{00}_{\delta^r}$. Let K be the pullback of C to G. Then K is a normal δ^r-type-definable bounded-index subgroup of G containing $G^{00}_{\delta^r}$. We also have $C = K/G^{00}_{\delta^r}$. Altogether, to prove the result, it suffices to show $K = G^0_{\delta^r}$.

We first show $K \subseteq G^0_{\delta^r}$. Let $H \leq G$ be δ^r-definable of finite index. We have $G^{00}_{\delta^r} \leq H$, and $H/G^{00}_{\delta^r}$ is a clopen subgroup of $G/G^{00}_{\delta^r}$ containing the identity. So $C \subseteq H/G^{00}_{\delta^r}$, i.e., $K \subseteq H$.

Now, to prove $G^0_{\delta^r} \subseteq K$, fix $a \notin K$. There is $X \subseteq G/G^{00}_{\delta^r}$ clopen such that $X = X^{-1}$, $aG^{00}_{\delta^r} \notin X$, and $G^{00}_{\delta^r} \in X$. Let A be the pullback of X to G, and note that $a \notin A$ and A is definable. Let $H = \{g \in G : gA = A\}$. Then H is a definable subgroup of G, and $G^{00}_{\delta^r} \subseteq H \subseteq A$. In particular, H has finite index, and is δ^r-definable by Corollary 4.2. Since $a \notin H$, we have shown $a \notin G^0_{\delta^r}$. □

Finally, we set some notation that will be used in later sections.

Definition 4.6. Suppose $\delta(x;\bar{y})$ is invariant and NIP.

(1) Let H_δ denote the compact Hausdorff group $G/G^{00}_{\delta^r}$.
(2) Let η_δ denote the normalized Haar measure on H_δ.
(3) Given $a \in G$, let $[a]_\delta$ denote the element $aG^{00}_{\delta^r}$ in H_δ.

5. Uniqueness of Measure

The goal of this section is to show that if $\delta(x;\bar{y})$ is invariant and NIP, then the pseudofinite counting measure is the unique left-invariant finitely additive probability measure on the Boolean algebra of δ^r-formulas. The proof of this will closely follow results about globally NIP groups from [1] and [8]. Due to the nature of the arguments, it will be cleaner to work around formulas $\phi(x)$ such that $\phi^\ell(x;y)$ is NIP, rather than fixing $\delta(x;\bar{y})$ outright.

Definition 5.1. Given a formula $\delta(x;\bar{y})$, let $\mathcal{B}_\delta(G)$ be the Boolean algebra of δ-formulas, and let $\mathfrak{M}_\delta(G)$ be the space of left invariant finitely additive probability measures defined on $\mathcal{B}_\delta(G)$.

Proposition 5.2. *Fix a formula $\phi(x)$ and a measure $\nu \in \mathfrak{M}_{\phi^\ell}(G)$. Suppose $\phi^\ell(x;y)$ is NIP. Then, for any ϕ^ℓ-formula $\psi(x)$, $\nu(\psi(x)) > 0$ if and only if $\psi(x)$ is generic.*

Proof. First, note that $\psi^\ell(x;y)$ is NIP for any ϕ^ℓ-formula $\psi(x)$. Now fix a ϕ^ℓ-formula $\psi(x)$, and suppose $\nu(\psi(x)) > 0$. By local versions of results from [1], essentially relying on Matoušek's (p,q)-theorem for set systems of finite VC-dimension, and since $\psi^\ell(x;y)$ is NIP, it follows that $\psi(x)$ is *weakly generic*, i.e., there are $g_1,\ldots,g_n \in G$ such that, if $\theta(x) := \bigvee_{i=1}^n \psi(g_i x)$, then $\neg\theta(x)$ is not generic. Since $\theta^\ell(x;y)$ is NIP, it follows from Corollary 2.15 that $\theta(x)$ is generic, which implies $\psi(x)$ is generic. □

Lemma 5.3. *Fix a formula $\phi(x)$ and a measure $\nu \in \mathfrak{M}_{\phi^\ell}(G)$. Suppose $\phi^\ell(x;y)$ is NIP. Then, for any $\epsilon > 0$, there are generic $p_1,\ldots,p_n \in S_{\phi^\ell}(G)$ such that, for any $g \in G$,*

$$|\nu(\phi(x)) - \mathbb{E}(p_1,\ldots,p_n;\phi(gx))| \leq \epsilon.$$

Proof. The strategy is to follow Section 4 of [8], in particular Lemma 4.8. We work over parameters in $\phi(x)$. Note that ν extends to a regular Borel probability measure on $S_{\phi^\ell}(G)$, which we also denote ν. Applying the *-version in [8] of the

VC Theorem, as in [8, Lemma 4.8], there are $n \geq 1$ and $B \subseteq (S_{\phi^\ell}(G))^n$ such that $\nu^n(B) > 0$ and, for any $(p_1, \ldots, p_n) \in B$ and for any $g \in G$,

$$|\nu(\phi(gx)) - \mathbb{E}(p_1, \ldots, p_n; \phi(gx))| \leq \epsilon.$$

By invariance of ν, the only remaining thing to show is that B contains a tuple $(p_1, \ldots, p_n)$ with each p_i generic. To do this, it suffices to show $\nu(X) = 0$, where $X := S_{\phi^\ell}(G) \backslash S^g_{\phi^\ell}(G)$ is the (open) set of non-generic ϕ^ℓ-types. Suppose $Y \subseteq X$ is compact and ν-measurable. Then, by compactness, there are non-generic ϕ^ℓ-formulas $\psi_1(x), \ldots, \psi_m(x)$ such that $Y \subseteq \bigcup_{i=1}^m [\psi_i(x)]$. Since $\nu(\psi_i(x)) = 0$ for all $1 \leq i \leq m$, this implies $\nu(Y) = 0$. By regularity of ν, we have $\nu(X) = 0$. □

For the rest of this section, fix a formula $\phi(x)$ such that $\delta(x; y) := \phi^\ell(x; y)$ is NIP. We will apply results of the previous section to $\delta(x; y)$.

Definition 5.4. Given a δ-formula $\psi(x)$ and a generic type $p \in S_\delta(G)$, define

$$U^p_{\psi(x)} = \{[a]_\delta \in H_\delta : \psi(x) \in ap\}.$$

Note that, in the previous definition, $U^p_{\psi(x)}$ is well-defined since $G^*_\delta \subseteq \mathrm{Stab}(p)$ by Theorem 3.15(b).

Proposition 5.5. *For any δ-formula $\psi(x)$ and generic type $p \in S_\delta(G)$, both $U^p_{\psi(x)}$ and its complement are F_σ subsets of H_δ.*

Proof. Fix a δ-formula $\psi(x)$ and generic $p \in S_\delta(G)$. Let $\theta(x; y_1, y_2)$ be the formula $\psi(y_1 \cdot x) \wedge \neg\psi(y_2 \cdot x)$, and note that $\theta(x; y_1, y_2)$ is NIP. By Corollary 2.8, we may find a countable model $M \prec G$ such that $\theta(x; y_1, y_2)$ is over M and, for any $a_1, a_2 \in G$, if $\mu(\theta(x; a_1, a_2)) > 0$ then $\theta(x; a_1, a_2)$ is realized in M.

We are going to use [16], which requires a countable theory. So let T be the complete theory of G in the language containing the group operation, $\psi(x)$, and constants for parameters in $\psi(x)$. For the rest of the proof we work in T. Since M contains the parameters in $\psi(x)$, we still have $M \prec G$ and we just treat $\psi(x)$ as a formula with no parameters. Let p_0 be the global ψ^ℓ-type obtained by restricting p to instances of $\psi^\ell(x; y)$. Then p_0 is M-invariant. Indeed, if $a_1 \equiv_M a_2$ and $\psi(a_1 x) \wedge \neg\psi(a_2 x) \in p_0$, then $\mu(\theta(x; a_1, a_2)) > 0$ since p_0 is generic, and so $\theta(x; a_1, a_2)$ is realized in M, a contradiction. By the main result of [16], if $\Sigma = \{q \in S_y(M) : \psi(ax) \in p_0 \text{ for } a \models q\}$ then both Σ and its complement are F_σ subsets of $S_y(M)$. Finally,

$$\{a \in G : [a]_\delta \in U^p_{\psi(x)}\} = \{a \in G : a \models q \text{ for some } q \in \Sigma\},$$

and so we have the desired result by Corollary 4.3. □

Definition 5.6. Given a generic type $p \in S_{\delta^r}(G)$, define $\mu_p \colon \mathcal{B}_\delta(G) \to [0, 1]$ such that

$$\mu_p(\psi(x)) = \eta_\delta(U^p_{\psi(x)}).$$

It is straightforward to check that $\mu_p \in \mathfrak{M}_{\phi^\ell}(G)$ for any generic $p \in S_\delta(G)$.

Definition 5.7. Given a δ-formula $\psi(x)$, define $\mathcal{S}_\psi = \{U^p_{\psi(gx)} : g \in G,\ p \in S^g_\delta(G)\}$.

We view each $\mathcal{S}_\psi$ as a set system with base set H_δ. Since $\psi^\ell(x; y)$ is NIP for any δ-formula $\psi(x)$, it follows that $\mathcal{S}_\psi$ has finite VC-dimension (see, e.g., [1, Lemma 3.19]). The following is the main technical lemma for proving the desired uniqueness properties of the pseudofinite measure μ. The proof is largely a local adaptation of work in [1, Section 3], and so we will sketch the strategy.

Lemma 5.8. *Given a generic type $p \in S_\delta(G)$ and $\nu \in \mathfrak{M}_{\phi^\ell}(G)$, we have $\mu_p(\psi(x)) = \nu(\psi(x))$ for any δ-formula $\psi(x)$.*

Proof. First, since G^*_δ is type-definable by δ^r-formulas, we may assume $\mathcal{L}$ is countable. So H_δ is a Polish space with respect to the logic topology and we may assume η_δ is complete (see [1, Remark 2.15]).

Claim 1: If $p \in S_\delta(G)$ is generic then $S^g_\delta(G) = \overline{\{gp : g \in G\}}$.
Proof: (See also [12].) Fix a generic type $p \in S_\delta(G)$. Note that $S^g_\delta(G)$ is closed, and clearly contains gp for any $g \in G$. For the other containment, suppose $q \in S_\delta(G)$ is generic, and let $\psi(x) \in q$. We want to find $g \in G$ such that $\psi(x) \in gp$. Since $\psi(x)$ is generic there are $g_1, \ldots, g_n \in G$ such that $G = g_1\psi(G) \cup \ldots \cup g_n\psi(G)$, and so $\psi(x) \in g_i^{-1}p$ for some i. $\dashv_{\text{Claim 1}}$

Claim 2: For any δ-formula $\psi(x)$, $\epsilon > 0$, and countable $S \subseteq S^g_\delta(G)$, there are $g_1, \ldots, g_m \in G$ such that, for any $a, a' \in G$ and $p \in S$,

$$|\mu_{ap}(\psi(x)) - \mathbb{E}(ag_1a'p, \ldots, ag_ma'p; \psi(x))| \leq \epsilon.$$

Proof: This is a direct translation of [1, Proposition 3.23], which involves an application of the full VC Theorem to $\mathcal{S}_\psi$ (see [1, Lemma 3.21], which uses the Polish structure on H_δ and completeness of η_δ). $\dashv_{\text{Claim 2}}$

Claim 3: If $p, q \in S_\delta(G)$ are generic, then $\mu_p = \mu_q$.
Proof: This is a direct translation of [1, Proposition 3.24], which relies on Claims 1 and 2. $\dashv_{\text{Claim 3}}$

We now combine everything to prove the lemma, which essentially amounts to translating Lemma 3.26 and Corollary 3.27 of [1]. Fix $p \in S^g_\delta(G)$, $\nu \in \mathfrak{M}_{\phi^\ell}(G)$, $\psi(x) \in \mathcal{B}_\delta(G)$, and $\epsilon > 0$. We show $|\nu(\psi(x)) - \mu_p(\psi(x))| \leq \epsilon$. By Claim 3, it is enough to find generic $p_1, \ldots, p_n \in S_\delta(G)$ such that

$$\left|\nu(\phi(x)) - \tfrac{1}{n}\textstyle\sum_{i=1}^n \mu_{p_i}(\phi(x))\right| \leq \epsilon.$$

To do this, one directly translates [1, Lemma 3.26]. The proof in [1] cites [1, Fact 2.9] and [1, Proposition 3.23], which we replace with Lemma 5.3 and Claim 2, respectively. $\square$

From Lemma 5.8, and the fact that μ is a left invariant finitely additive probability measure on all formulas, we immediately have the following result.

Theorem 5.9. *Let $\phi(x)$ be a formula such that $\phi^\ell(x; y)$ is NIP. Suppose ν is a left invariant finitely additive probability measure defined on some Boolean algebra containing all left translates of $\phi(x)$. Then $\nu(\phi(x)) = \mu(\phi(x))$.*

6. Generic compact domination

Throughout this section, we fix an invariant NIP formula $\delta(x; \bar{y})$.

Definition 6.1.

(1) Let π_δ denote the map $\pi_{G^{00}_{\delta^r}} : S_{\delta^r}(G) \to H_\delta$ defined before Corollary 4.3.
(2) Given $\alpha \in H_\delta$, define

$$S^\alpha_{\delta^r}(G) := \pi_\delta^{-1}(\alpha) \cap S^g_{\delta^r}(G),$$

i.e. $p \in S^{\alpha}_{\delta^r}(G)$ if and only if p is a global generic δ^r-type containing the δ^r-type-definition of α (as a coset of $G^{00}_{\delta^r}$).

(3) Given a δ^r-formula $\phi(x)$, define

$$E_\phi = \{\alpha \in H_\delta : S^{\alpha}_{\delta^r}(G) \cap \phi(x) \neq \emptyset \text{ and } S^{\alpha}_{\delta^r}(G) \cap \neg\phi(x) \neq \emptyset\},$$

where we identify a δ^r-formula with a clopen set of types in $S_{\delta^r}(G)$.

(4) Given a δ^r-formula $\phi(x)$ and a generic type $p \in S_{\delta^r}(G)$, define

$$U^p_\phi = \{[a]_\delta \in H_\delta : \phi(x) \in ap\}.$$

Note that, in the previous definition, $U^p_{\phi(x)}$ is well-defined since $G^{00}_{\delta^r} = \mathrm{Stab}(p)$ by Theorem 3.15(c).

Theorem 6.2. *For any δ^r-formula $\phi(x)$, E_ϕ is closed and $\eta_\delta(E_\phi) = 0$.*

Proof. The proof adapts parts of [1] and [17], and relies on the main results of [16] and [17]. Throughout the proof we will use the fact that, for any δ^r-formula $\phi(x)$, the formula $\phi(x \cdot y)$ is NIP (see Proposition 3.12(a)). In particular, for any generic $p \in S_{\delta^r}(G)$, and any δ^r-formula $\phi(x)$, the family of left translates of U^p_ϕ has finite VC-dimension (as a set system on H_δ).

First, we observe that E_ϕ is the intersection of $\pi_\delta(S^g_{\delta^r}(G) \cap \phi(x))$ and $\pi_\delta(S^g_{\delta^r}(G) \cap \neg\phi(x))$. Thus E_ϕ is closed since $S^g_{\delta^r}(G)$, $\phi(x)$, and $\neg\phi(x)$ are closed, and π_δ is a continuous map between compact Hausdorff spaces.

Claim 1: If $\phi(x)$ is a δ^r-formula and $p \in S_{\delta^r}(G)$ is generic, then both U^p_ϕ and its complement are F_σ subsets of H_δ.
Proof: Let p_0 be the restriction of p to instances of $\phi(gx)$ or $\neg\phi(gx)$ for $g \in G$. Then p_0 is a generic ϕ^ℓ-type and $\{a \in G : \phi(ax) \in p\} = \{a \in G : \phi(ax) \in p_0\}$. So the result follows as in Proposition 5.5 (using [16]). $\dashv_{\text{Claim 1}}$

By Claim 1, each generic $p \in S_{\delta^r}(G)$ induces a left-invariant finitely additive probability measure on δ^r-formulas, by assigning the measure of a δ^r-formula $\phi(x)$ to be $\eta_\delta(U^p_\phi)$. By Theorem 5.9, $\eta_\delta(U^p_\phi) = \mu(\phi(x))$ for any δ^r-formula $\phi(x)$. For the rest of the proof, fix a generic type $p \in S_{\delta^r}(G)$ concentrating on $G^{00}_{\delta^r}$.

Claim 2: For any δ^r-formula $\phi(x)$ and any $\alpha \in U^p_\phi$, if $V \subseteq H_\delta$ is an open neighborhood of α then $\eta_\delta(U^p_\phi \cap V) > 0$.
Proof: We follow the proof of Claim 2 of [17, Theorem 3.2]. Since $\pi^{-1}_\delta(\alpha)$ and $\pi^{-1}_\delta(\neg V)$ are disjoint closed subsets of $S_{\delta^r}(G)$, there is some δ^r-formula $\psi(x)$ such that $\pi^{-1}_\delta(\alpha) \subseteq \psi(x) \subseteq \pi^{-1}_\delta(V)$. Fix $a \in \alpha$. Then $\psi(x) \in ap$, and so $\phi(x) \wedge \psi(x) \in ap$, which implies $\mu(\phi(x) \wedge \psi(x)) > 0$. Since $U^p_{\phi(x)\wedge\psi(x)} = U^p_\phi \cap U^p_\psi$, we have $\eta_\delta(U^p_\phi \cap U^p_\psi) > 0$. Now suppose $[g]_\delta \in U^p_\psi$. Then $\psi(x) \in gp$ and so, since p concentrates on $G^{00}_{\delta^r}$, we have $gG^{00}_{\delta^r} \subseteq \psi(G)$, and so $\pi^{-1}_\delta([g]_\delta) \in \psi(x) \subseteq \pi^{-1}_\delta(V)$. So $U^p_\psi \subseteq V$ and thus $\eta_\delta(U^p_\phi \cap V) > 0$. $\dashv_{\text{Claim 2}}$

By Claim 2, Proposition 5.5, and [17, Theorem 2.7], we have $\eta_\delta(\partial U^p_\phi) = 0$. So, to prove the result, it suffices to show $E_\phi \subseteq \partial U^p_\phi$. The argument follows the proof of [1, Theorem 5.3]. Fix $\alpha \in E_\phi$, and let $V \subseteq H_\delta$ be open, with $\alpha \in V$. Since $\alpha \in E_\phi$, there are $q, q' \in S^{\alpha}_{\delta^r}(G)$ such that $\phi(x) \in q$ and $\neg\phi(x) \in q'$. Let $S = \pi^{-1}_\delta(V)$, and note that $S \subseteq S_{\delta^r}(G)$ is open, with $\pi^{-1}_\delta(\alpha) \subseteq S$. In particular, $q \in S \cap \phi(x)$ and

$q' \in S \cap \neg\phi(x)$, and so these are are nonempty open sets in $S_{\delta^r}(G)$. As in Claim 1 in the proof of Lemma 5.8, we have $S^g_{\delta^r}(G) = \overline{\{gp : g \in G\}}$. Therefore, there are $g, g' \in G$ such that $gp \in S \cap \phi(x)$ and $g'p \in S \cap \neg\phi(x)$. Since p concentrates on $G^{00}_{\delta^r}$, we have $\pi_\delta(gp) \in V \cap U^p_\phi$ and $\pi_\delta(g'p) \in V \cap \neg U^p_\phi$. Altogether, $\alpha \in \partial U^p_\phi$. □

Corollary 6.3. *Suppose $C \subseteq H_\delta$ is closed, with $\eta_\delta(C) = 0$, and let $X = \{a \in G : [a]_\delta \in C\}$. Then, for any $\epsilon > 0$, there is a δ^r-definable set $Z \subseteq G$ such that $X \subseteq Z$ and $\mu(Z) \leq \epsilon$.*

Proof. We follow [6, Section 9] and [9, Section 5]. First, we use η_δ to define a left-invariant finitely additive probability measure ν on δ^r-definable sets. In particular, given a δ^r-formula $\phi(x)$, define

$$C_\phi = \{\alpha \in H_\delta : S^\alpha_{\delta^r}(G) \cap \phi(x) \neq \emptyset\}.$$

We have already observed that C_ϕ is closed. Define $\nu(\phi(x)) = \eta_\delta(C_\phi)$. Given a δ^r-formula $\phi(x)$ and $g \in G$, we have $[g]_\delta C_\phi = C_{\phi(g^{-1}x)}$, and so ν is left invariant by left invariance of η_δ. It remains to show finite additivity, so fix disjoint δ^r-formulas $\phi(x)$ and $\psi(x)$. Since $C_{\phi\vee\psi} = C_\phi \cup C_\psi$, it suffices, by finite additivity of η_δ, to show $\eta_\delta(C_\phi \cap C_\psi) = 0$. Since $C_\psi \subseteq C_{\neg\phi}$, we have $C_\phi \cap C_\psi \subseteq E_\phi$, and so this follows from Theorem 6.2.

By Theorem 5.9, we have $\nu(\phi(x)) = \mu(\phi(x))$ for any δ^r-formula $\phi(x)$. Fix a closed set $C \subseteq H_\delta$, with $\eta_\delta(C) = 0$, and let $X = \{a \in G : [a]_\delta \in C\}$. Then X is type-definable and so we may fix a small family $\{\phi_i(x) : i \in I\}$ of δ^r-formulas, which is closed under finite conjunctions, such that $X = \bigcap_{i\in I} \phi_i(G)$. To prove the result, it suffices to show $\inf_{i\in I} \nu(\phi_i(x)) = 0$. Given $i \in I$, let $C_i = C_{\phi_i}$ and define the closed set $D = \bigcap_{i\in I} C_i$. Note that $D \subseteq C$ since, if $\alpha \in D$ then $\alpha \cap \bigcap_{i\in I} \phi_i(G)$ is finitely satisfiable, and so $\alpha \cap X \neq \emptyset$, which implies $\alpha \in C$. In particular, $\eta_\delta(D) = 0$. By compactness of H_δ, if $U \subseteq H_\delta$ is open and $D \subseteq U$, then there is some $i \in I$ such that $C_i \subseteq U$. Altogether, by regularity of η_δ, it follows that

$$0 = \eta_\delta(D) = \inf_{i\in I} \eta_\delta(C_i) = \inf_{i\in I} \nu(\phi_i(x)),$$

as desired. □

References

[1] Artem Chernikov and Pierre Simon, *Definably amenable NIP groups*, arXiv:1502.04365, 2015, to appear in J. Amer. Math. Soc.

[2] Gabriel Conant, Anand Pillay, and Caroline Terry, *A group version of stable regularity*, arXiv:1710.06309, 2017.

[3] ______, *Structure and regularity for subsets of groups with finite VC-dimension*, arXiv 1802.04246, 2018.

[4] B. Green, *A Szemerédi-type regularity lemma in abelian groups, with applications*, Geom. Funct. Anal. **15** (2005), no. 2, 340–376. MR 2153903

[5] David Haussler and Emo Welzl, *ϵ-nets and simplex range queries*, Discrete Comput. Geom. **2** (1987), no. 2, 127–151. MR 884223

[6] Ehud Hrushovski, Ya'acov Peterzil, and Anand Pillay, *Groups, measures, and the NIP*, J. Amer. Math. Soc. **21** (2008), no. 2, 563–596.

[7] Ehud Hrushovski and Anand Pillay, *Groups definable in local fields and pseudo-finite fields*, Israel J. Math. **85** (1994), no. 1-3, 203–262.

[8] ______, *On NIP and invariant measures*, J. Eur. Math. Soc. (JEMS) **13** (2011), no. 4, 1005–1061.

[9] Ehud Hrushovski, Anand Pillay, and Pierre Simon, *Generically stable and smooth measures in NIP theories*, Trans. Amer. Math. Soc. **365** (2013), no. 5, 2341–2366.

[10] János Komlós, János Pach, and Gerhard Woeginger, *Almost tight bounds for ϵ-nets*, Discrete Comput. Geom. **7** (1992), no. 2, 163–173.
[11] D. Lascar and A. Pillay, *Hyperimaginaries and automorphism groups*, J. Symbolic Logic **66** (2001), no. 1, 127–143.
[12] Ludomir Newelski, *Topological dynamics of definable group actions*, J. Symbolic Logic **74** (2009), no. 1, 50–72.
[13] Anand Pillay, *Type-definability, compact Lie groups, and o-minimality*, J. Math. Log. **4** (2004), no. 2, 147–162.
[14] Bruno Poizat, *A course in model theory*, Universitext, Springer-Verlag, New York, 2000, An introduction to contemporary mathematical logic, Translated from the French by Moses Klein and revised by the author. MR 1757487 (2001a:03072)
[15] Pierre Simon, *A guide to NIP theories*, Lecture Notes in Logic, vol. 44, Association for Symbolic Logic, Chicago, IL; Cambridge Scientific Publishers, Cambridge, 2015.
[16] Pierre Simon, *Rosenthal compacta and NIP formulas*, Fund. Math. **231** (2015), no. 1, 81–92.
[17] ———, *VC-sets and generic compact domination*, Israel J. Math. **218** (2017), no. 1, 27–41.

Department of Mathematics, University of Notre Dame, Notre Dame, IN, 46656, USA
E-mail address: gconant@nd.edu

Department of Mathematics, University of Notre Dame, Notre Dame, IN, 46656, USA
E-mail address: apillay@nd.edu