1 An Infinitary Setting

We will begin with an infinite sequence of groups G_p and subsets $X_p \subseteq X$ such that for each p, $|X_pX_pX_p| \leq K|X_p|$ for some fixed K independent of p. (In actual examples, we may modify the exact properties satisfied by X_p.)

Take an extension of the language of groups which adds predicates $m(\{\vec{x} \mid \phi(\vec{x}, \vec{a})\}) \geq r$ whenever ϕ is a formula and r a rational, and a distinguished predicate X. Then for each p, (G_p, X_p) gives a model of this language, interpreting $m(\{\vec{x} \mid \phi(\vec{x}, \vec{a})\}) \geq r$ to hold when $\frac{|\{\vec{x}|\phi(\vec{x},\vec{a})\}|}{|X^{(|\vec{x}|)}|} \geq r$.

Take an ultraproduct of these models, giving a model (G, X) with G, X infinite. We may define a measure μ on G-definable subsets of $G^{(l)}$: if $A = \{\vec{x} \mid \phi(\vec{x}, \vec{a})\}$, set $\mu(A) := \inf_r (G, X_0) \vDash m(\{\vec{x} \mid \phi(\vec{x}, \vec{a})\}) \geq r$.

If $C \subseteq G^{(n+1)}$ is a definable set, we define $C_{\vec{a}} := \{x \mid (x, \vec{a}) \in C\}$.

Note that if $A = \{\vec{x} \mid \phi(\vec{x}, \vec{a})\}$, it is not always the case that $\mu(A) \geq r$ is equivalent to $(G, X_0) \vDash m(\{\vec{x} \mid \phi(\vec{x}, \vec{a})\}) \geq r$: if $(G, X_0) \vDash m(\{\vec{x} \mid \phi(\vec{x}, \vec{a})\}) \geq s$ for every rational $s > r$ then we will have $\mu(A) = r$ even if $(G, X) \vDash \neg m(\{\vec{x} \mid \phi(\vec{x}, \vec{a})\}) \geq r$.

In particular, $\{a \mid \mu(C_a) > 0\}$ may not be definable. However, $\mu(C_a)$ is $\bigwedge$-definable: $\mu(C_a) > 0$ iff $(G, X) \vDash \neg m(C_a) \geq r$ for all $r > 0$.

Rather than work with G, we are more interested in $\tilde{G} := \bigcup_n (X^{-1}X)^n$. From here on, we will restrict to the model $(\tilde{G}, X)$. Fix a countable elementary submodel $M \prec \tilde{G}$; for the remainder, we will always allow parameters from M.

The M-definable sets give a countable, finitely additive algebra of subsets of $\tilde{G}$, and μ is a finitely additive measure on them. The M-definable sets on $\tilde{G}^{(2)}$ are also a countable, finitely additive algebra, but much bigger than the product of the collection of definable subsets of $\tilde{G}$ with itself. Since we have sufficient saturation, we could extend μ to a σ-additive measure on the σ-algebra generated by the definable sets, and the measure over $\tilde{G}^{(2)}$ would satisfy Fubini's Theorem relative to the measure over $\tilde{G}$; however we will be able to work with the following more elementary properties:

- $\mu(A \times B) = \mu(A)\mu(B)$
- If $\mu(A_a) = 0$ for all a then $\mu(A) = 0$

We have $\mu(X) = 1$, and $\mu(X^{(n)})$ is bounded.

Definition 1. *A* partial type *p over a set D is a collection of $M \cup D$-definable subsets of $\tilde{G}$ such that whenever $\mathcal{A} \subseteq p$ is finite, $\bigcap \mathcal{A}$ is non-empty. A* type *over D is a maximal partial type over D. A* global type *is a type over $\tilde{G}$.*

A type p is wide *if whenever $\mathcal{A} \subseteq p$ is finite, $\mu(\bigcap \mathcal{A}) > 0$.*

We write $a \vDash p$ if for every $A \in p$, $a \in A$. We write $p(\tilde{G}) := \bigcap p = \{a \mid a \vDash p\}$. We write $tp(a/D)$ for the collection of $M \cup D$-definable sets containing a.

In particular, $tp(a/D)$ is the unique type such that $a \vDash tp(a/D)$.

Note that being a type is relative to a particular choice of algebra: in general, a wide typeover a is only a wide partial type over a, b.

Lemma 2. *Each type is contained in $X^{(n)}$ for a fixed n, or disjoint from all $X^{(n)}$.*

Proof. This follows from the fact that $X^{(n)}$ is definable for each n. □

From here on, we always assume that a type is contained in $X^{(n)}$ for some n. By saturation, all such types are realized.

Lemma 3. *Every wide partial type can be extended to a wide type.*

Proof. Let $B_1, \ldots, B_i, \ldots$ enumerate the definable sets. Define a sequence of partial types with $q_0 := q$ and, given q_n, if for some $A \supseteq q_n$, $\mu(A \cap B_n) = 0$, $q_{n+1} := q_n \cup \{\overline{B}_n\}$, otherwise $q_{n+1} := q_n \cup \{B_n\}$.

Let $q' := \bigcup_n q_n$. If q' is not wide, there must be some n such that q_{n+1} is not wide. Let n be least such that q_{n+1} is not wide. If $B_n \in q_{n+1}$ then there is some $A \in q_n$ with $\mu(B_n \cap A) = 0$; this contradicts the definition of q_{n+1}. So $\overline{B}_n \in q_{n+1}$, and there is an $A \in q_n$ with $\mu(A \cap B_n) = 0$. But there must be some $A' \in q_n$ with $\mu(A' \cap \overline{B}_n) = 0$. Then $\mu(A \cap A') = 0$, contradicting the fact that $A, A' \in q_n$ and q_n is wide. Therefore q' is wide. □

If $A, B \subseteq \tilde{G}$ (either definable or $\bigwedge$-definable), we write $Ab = \{ab \mid a \in A\}$, $bA = \{ba \mid a \in A\}$, $AB = \{ab \mid a \in A, b \in B\}$, $A^{-1} = \{a^{-1} \mid a \in A\}$. Similarly, we write $q^{-1} := \{A^{-1} \mid A \in q\}$, $qA := \{BA \mid B \in q\}$, $Aq := \{AB \mid B \in q\}$, $pq := \{AB \mid A \in p, B \in q\}$.

Lemma 4.
- *If q is a type, so is q^{-1}*
- *If q is a type over D and $a \in D$, $qa = \{xa \mid x \in q\}$ is a type over D.*

Proof.
- For any definable B, either $B^{-1} \in q$ or $\overline{B^{-1}} \in q$, and therefore eitehr $B \in q^{-1}$ or $\overline{B} \in q^{-1}$ respectively.
- For any B, either $Ba^{-1} \in q$ or $\overline{B}a^{-1} \in q$, and therefore either $B \in qa$ or $\overline{B} \in qa$ respectively.

□

Lemma 5.
- $q^{-1}(\tilde{G}) = (q(\tilde{G}))^{-1}$
- $(qA)(\tilde{G}) = (q(\tilde{G}))A$
- $(pq)(\tilde{G}) = (p(\tilde{G}))(q(\tilde{G}))$

Proof. Only the third requires proof. If $x \in (p(\tilde{G}))(q(\tilde{G}))$ then $x \in AB$ for all $A \in p$, $B \in q$, and therefore $x \in (pq)(\tilde{G})$. Conversely, if $x \in (pq)(\tilde{G})$, so $x \in AB$ whenever $A \in p, B \in q$, it follows by saturation that there is a pair $(y, z) \in p(\tilde{G}) \times q(\tilde{G})$ with $yz = x$, and therefore $x \in p(\tilde{G})q(\tilde{G})$. □

2 Wide Types are Special

Definition 6. (b_i) *is a* sequence of indiscernibles *if whenever* $\phi(b_0, \ldots, b_n)$ *and* $m_0 < \cdots < m_n$, $\phi(b_{m_0}, \ldots, b_{m_n})$.

That is, for every n there is a type p_n of $n+1$-tuples such that $tp(b_{m_0}, \dots, b_{m_n}) = p_n$ whenever $m_0 < \dots < m_n$ is an increasing sequence.

Lemma 7. *Suppose $\mu(A_b) > 0$ holds and (b_i) is a sequence of indiscernibles with $tp(b_i) = tp(b)$. Then for any finite $b_1, \dots, b_n$, there is an $a \in \bigcap_{i \le n} A_{b_i}$.*

Proof. Suppose not: that is, for some k, $\bigcap_{i \le k} A_{b_i}$ is empty. Let k be least such that $\mu(\bigcap_{i \le k+1} A_{b_i}) = 0$; since $\mu(A_{b_0}) > 0$, $k \ge 0$. For $n \ge k$, let $C_n = \bigcap_{i < k} A_{b_i} \cap A_{b_n}$. Then $\mu(C_n) = \mu(\bigcap_{i \le k} A_{b_i}) > 0$ since $k < k+1$. But for $n < m$, by indiscernibility, $\mu(C_n \cap C_m) = \mu(\bigcap_{i \le k+1} A_{b_i}) = 0$. This is a contradiction since μ is a finite measure. □

Lemma 8. *If $tp(a/b)$ is wide and (b_i) is a sequence of indiscernibles with $tp(b_i) = tp(b)$ then for any finite $b_1, \dots, b_n$, there is an a' with $tp(a', b_i) = tp(a, b)$ for all $i \le n$.*

Proof. Suppose $(a, b) \in A$. Then $a \in A_b$, so by the previous lemma, there is an a' so that $a' \in \bigcap_{i \le n} A_{b_i}$, so $(a', b_i) \in A$ for each $i \le n$. By saturation, there is an a' such that $(a', b_i) \in A$ simultaneously for all $i \le n$, $A \ni (a, b)$. Therefore $tp(a', b_i) = tp(a, b)$ for all $i \le n$. □

Definition 9. *A type p is M-invariant if whenever σ is an automorphism preserving M and $A_a \in p$ then $A_{\sigma(a)} \in p$. p is M-finitely satisfiable if whenever $\mathcal{A} \subseteq p$ is finite, $\bigcap \mathcal{A} \cap M \neq \emptyset$.*

Lemma 10. *If p is finitely satisfiable in M then p is M-invariant.*

Proof. Suppose not; then there is an automorphism σ and an a so that $A_a \setminus A_{\sigma(a)} \in p$. But then there is an element $m \in M$ such that $m \in A_a \setminus A_{\sigma(a)}$, contradicting the fact that σ is an automorphism preserving M. □

Lemma 11. *If p is a type over M, there is an extension of p to a global M-finitely satisfiable (and therefore M-invariant) type.*

Proof. Let $p' := p \cup \{\overline{A} \mid A \cap M = \emptyset\}$; if p' is not finitely satisfiable in $\tilde{G}$, there is an $A \in p$ and a B with $B \cap M = \emptyset$ such that $A \cap \overline{B} = \emptyset$. But since $A \in p$ and p is a type over M, $A \neq \emptyset$, and since M is an elementary submodel of $\tilde{G}$, $A \cap M \neq \emptyset$. So $A \cap M \cap (B \cap \overline{B}) \neq \emptyset$, and since $B \cap M = \emptyset$, $A \cap \overline{B} \neq \emptyset$.

Let q be an arbitrary extension of p' to a type. Suppose q is not finitely satisfiable in M; then there is a $B \in q$ with $B \cap M = \emptyset$, which is impossible since then $\overline{B} \in p' \subseteq q$. □

Definition 12. *If S is a set and p a global type, write $p \mid S$ for the restriction of p to sets definable over S.*

Lemma 13. *Suppose p is a global M-invariant type, and recursively choose $b_n \vDash p \mid M \cup \{b_i\}_{i<n}$. Then $\{b_i\}$ is a sequence of indiscernibles.*

Proof. By induction on n, we show $tp(b_{m_0}, \dots, b_{m_n})$ is constant whenever $m_0 < \dots < m_n$. When $n = 1$, this follows since each $b_i \in p(\tilde{G})$. Suppose $(b_{n+1}, b_n, \dots, b_0) \in A$. Then also $(b_{m_{n+1}}, b_n, \dots, b_0) \in A$ holds because $tp(b_{m_{n+1}}/b_0, \dots, b_n) = tp(b_{n+1}/b_0, \dots, b_n)$. By IH, $tp(b_0, \dots, b_n) = tp(b_{m_0}, \dots, b_{m_n})$, and therefore there is an automorphism σ fixing M such that $\sigma(b_i) = b_{m_i}$ for all $i \leq n$. Since $A_{b_n,\dots,b_0} \in p$, also $A_{b_{m_n},\dots,b_{m_0}} \in p$, so $(b_{m_{n+1}}, b_{m_n}, \dots, b_{m_0}) \in A$. □

Theorem 14. *Let p, q be types (over M), and let $a, a' \in p(\tilde{G})$, $b, b' \in q(\tilde{G})$ with $tp(a/b), tp(a'/b')$ wide. Then $\mu(A_a \cap B_b) > 0$ iff $\mu(A_{a'} \cap B_{b'}) > 0$.*

Proof. Suppose $\mu(A_a \cap B_b) \geq \delta > 0$ but $\mu(A_{a'} \cap B_{b'}) = 0$.

Fix an extension q' of q to an invariant global type. We consider two cases; first, suppose that there is some $\epsilon > 0$ such that whenever $\hat{a} \vDash p$ and $\hat{b} \vDash q' \mid \hat{a}$, $\mu(A_{\hat{a}} \cap B_{\hat{b}}) \geq \epsilon$. Choose $a_0 \vDash p$ and $b_0 \vDash q' \mid a_0$. Given $\{a_i, b_i\}_{i \leq n}$, since $tp(a'/b)$ is wide, $\mu(\{x \mid \mu(A_x \cap B_b) = 0\})$ is $\bigwedge$-definable, and therefore there is an a_{n+1} with $tp(a_{n+1}) = tp(a)$ such that $\mu_x(A_{a_{n+1}} \cap B_{b_i}) = 0$ for $i \leq n$, and let $b_{n+1} \vDash q \upharpoonright \{a_i, b_i\}_{i \leq n}, \{a_{n+1}\}$. For each i, let $C_i = A_{a_i} \cap B_{b_i}$. Since each $b_n \vDash q \upharpoonright a_n$, $\mu(C_i) \geq \epsilon$. This is a contradiction.

In the second case, whenever $\hat{a} \vDash p$ and $\hat{b} \vDash q' \mid \hat{a}$, $\mu(A_{\hat{a}} \cap B_{\hat{b}}) = 0$. In this case, given $\{a_i, b_i\}_{i<n}$ (including the empty case where $n = 0$), choose $b_0 \vDash q' \mid \{a_i, b_i\}_{i<n}$. Since $\mu(\{x \mid \mu(A_x \cap B_b) \geq \delta\})$ is $\bigwedge$-definable, there is an $a_n \vDash p$ such that $\mu(A_{a_n} \cap B_{b_n}) \geq \delta$. Then, again setting $C_i := A_{a_i} \cap B_{b_i}$, $\mu(C_i) \geq \delta$, and when $i < j$, $\mu(C_i \cap C_j) = 0$ since $b_j \vDash q' \mid a_i$, so $\mu(A_{a_i} \cap B_{b_j}) = 0$. □

Lemma 15. *If q is a wide type, p is a wide type, and $a, b \in q$, $tp(a/b)$ wide, then $pa \cap pb$ is wide.*

Proof. Suppose $pa \cap pb$ is not wide; then there is a definable $A \in p$ with $\mu(A) > 0$ such that $\mu(Aa \cap Ab) = 0$. Let $a_0, \dots, a_n$ be a sequence of elements from q with $tp(a_n/a_0, \dots, a_{n-1})$ wide for each n. Then $\mu(Aa_i \cap Aa_j) = 0$ for all $i \neq j$; but this is impossible since $\mu(Aa_i)$ is constantly non-zero. □

This is the first step towards the proof of the following, which is a weakened form of Theorem 3.4:

Theorem 16. *If q is a wide type and for every $a \in q$, there is a $b \in q$ such that $tp(b/a)$ and $tp(a/b)$ are both wide, then $q^{-1}qq^{-1}q$ is a group, normal in $\bigcup X^n$.*

Lemma 17. *If q is wide, q has index $\leq 2^{\aleph_0}$ in $\tilde{G}$.*

Proof. Suppose not; let $\{Sa_i\}_{i<(2^{\aleph_0})^+}$ be distinct cosets. By the pigeonhole principle, we may assume $a_i \in (X^{-1}X)^n$ for some fixed n.

We may write $q = \bigcap_{k<\omega} S_k$ with each S_k definable and $\mu(S_k) > 0$. By compactness, for each $i \neq j$, there is a $k(i, j)$ such that $S_k a_i \cap S_k a_j = \emptyset$ for each i, j. By Erdös-Rado, there is a countable set I such that k is constant on $[I]^2$; but this gives countable many disjoint copies of S_k, a contradiction. □

The following two lemmas essentially give the result of Lemma 2.15. If B is a set of pairs, write $B^\top := \{(y, x) \mid (x, y) \in B\}$.

Lemma 18. *Let C, D be definable sets of pairs, and let P be a definable set with $\mu(P) > 0$. Assume $P \times P \subseteq C \cup D$. Then there is a definable $P' \subseteq P$ with $\mu(P') > 0$ and either:*

- $\mu(C_a) > 0$ *for every $a \in P'$, or*
- $\mu(D_a^{\top}) > 0$ *for every $a \in P'$*

Proof. Let $Q_1 \subseteq Q_2 \subseteq \cdots$ be a sequence of subsets of P such that $a \in Q_n$ implies $\mu(C_a) > 0$[1]. If for some n, $\mu(Q_n) > 0$, we may take $P' := Q_n$. If $\mu(Q_n) = 0$ for all n, it follows by the Fubini property of μ that $\mu(C \cap (\tilde{G} \times P)) = 0$.

Similarly, either there is a $P' \subseteq P$ $\mu(P') > 0$ and $\mu(D_a^{\top}) > 0$ for every $a \in P'$, or $\mu(D \cap (P \times \tilde{G})) = 0$.

If neither of these possibilities for P' exist, we have

$$\mu(P)^2 = \mu(P \times P) = \mu((C \cup D) \cap (P \times P)) \leq \mu(C \cap (\tilde{G} \times P)) + \mu(D \cap (P \times \tilde{G})) = 0$$

contradicting the assumption. □

Lemma 19. *Let B be a definable set with $\mu(B) > 0$. Then there is a wide, global, finitely satisfiable $q \ni B$ such that for each $a \models q$, there is a $b \models q$ with $tp(a/b)$ and $tp(b/a)$ wide.*

Proof. By choosing an ordering of the definable sets, and an ordering of the pairs of definable sets in which each pair appears infinitely often, we may choose a wide type $q \ni B$ such that whenever $P \in q$, $P \times P \subseteq C \cup D$, either for every $a \in P$ $\mu(C_a) > 0$, or for every $a \in P$ $\mu(D_a^{\top}) > 0$.

Let $a \in q$, and let $q_0 := q \cup \{\overline{C_a} \mid \mu(C_a) = 0\} \cup \{\overline{D_a} \mid \mu(D_a^{\top}) = 0\}$. First, we must check that q_0 is consistent: if not, we have $P \cap \overline{C_a} \cap \overline{D_a} = \emptyset$ for some $P \in q$ and some C, D with $\mu(C_a) = \mu(D_a^{\top}) = 0$. Then $P \subseteq C_a \cup D_a$; since q is a type, $P' := \{x \mid P \subseteq C_x \cup D_x\} \in q$, and therefore setting $P'' := P \cap P'$, we have $P'' \in q$ and whenever $x, y \in P''$, $y \in C_x \cup D_x$. Equivalently, $P'' \times P'' \subseteq C \cup D$. By the construction of q, there is a $Q \in q$ such that either for every $x \in Q$, $\mu(C_x) > 0$, or for every $x \in Q$, $\mu(D_x^{\top}) > 0$. In particular, $a \in Q$, contradicting the fact that $\mu(C_a) = \mu(D_a^{\top}) = 0$.

Since q_0 is consistent, it may be extended to a type q' over a. Let $b \in q'$; if $C(b, a)$ holds, $b \in C_a$, and if $\mu(C_a) = 0$, we would have $\overline{C_a} \in q_0 \subseteq q'$. So $\mu(C_a) > 0$. Therefore $tp(b/a)$ is wide.

On the other hand, suppose $D(b, a)$ holds, so $b \in D_a$, and therefore $\mu(D_a^{\top}) > 0$. Since $tp(b) = tp(a)$, it follows that $\mu(D_b^{\top}) > 0$ as well, so $tp(a/b)$ is wide also. □

3 Corollary 1.2

Write $a^X := \{x^{-1}ax \mid x \in X\}$, $A^{(l)}$ for l-tuples from A.

Theorem 20 (Corollary 1.2)**.** *For any k, l, m, there is a $p < 1$ and a K such that whenever G is a group, $X_0 \subseteq G^{p,k}$ finite, and $X = X_0^{-1} X_0$:*

[1] Such a sequence exists by the $\bigwedge$-definability of μ; for instance, take $Q_n := \{a \in P \mid m(C_a) > 1/n\}$.

- $|X_0X| \leq k|X_0|$
- *With probability* $\geq p$*, an* l*-tuple* $(a_1, \ldots, a_l) \in (X^2)^{(l)}$ *satisfies* $|a_1^X \cdots a_l^X| \geq |X|/m^2$

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

Suppose not; then for some k, l, m and every p, K, there is a $G^{p,k}$, an $X_0^{p,k} \subseteq G^{p,k}$ finite, such that the conditions hold, but there is no definable subgroup S with X contained in K cosets of S.

So $\{(a_1, \ldots, a_l) \in X^{(l)} \mid \mu(a_1^X \cdots a_l^X) \geq 1/m\}$ is not definable. However, if we choose $r \in (1/(m+1), 1/m)$, and set $Q = \{(a_1, \ldots, a_l) \in X^{(l)} \mid m(a_1^X \cdots a_l^X) \geq r\}$, Q is definable, and we will have $\mu(Q) = 1$ and $\mu(a_1^X \cdots a_l^X) \geq 1/m+1$ whenever $(a_l, \ldots, a_l) \in Q$.

We may refine X_0 to a type q to which 3.4 applies, so $S := q^{-1}qq^{-1}q$ is a normal subgroup of $\tilde{G}$.

Since S is wide, for each k, $\mu(S_k) > 0$, so there is $(a_1, \ldots, a_l) \in S_k^{(l)} \cap Q$[3]. By saturation, there is an $a_1, \ldots, a_l \in S \cap Q$. Since S is normal, $a_1^X \cdots a_l^X \subseteq S$, and therefore S has fewer than $\mu(X_0X)(m+1)$ distinct cosets.

So X_0/S is finite and since $\mu(XX) \leq k|\mu(X_0)|$, XX/S is finite as well. This means that XX is the union of finitely many translations of S: $XX = S \cup \bigcup_i c_iS$.

But this means that both S and the complement of S are $\bigwedge$-definable. If there were no j with $S_j = S$ then we would have $S_j \cap \overline{S}$ non-empty for all j, and therefore $S \cap \overline{S}$ non-empty by saturation; this is clearly impossible, so $S = S_j$ for some j. Therefore S is definable and finitely many cosets of S cover X. This is a contradiction, so there must have been some p, K satisfying the theorem.

[2] This differs slightly from the statement of Corollary 1.2, since we take tuples from X^2 rather than X.

[3] If Q were a subset of $X^{(l)}$, instead of $(X^2)^{(l)}$, a slightly more complicated argument is required here.