Notre Dame Journal of Formal Logic
Volume 54, Numbers 3–4, 2013

On Pseudo-Finite Dimensions

Ehud Hrushovski

Abstract We attempt to formulate issues around modularity and Zilber's trichotomy in a setting that intersects additive combinatorics. In particular, we update the open problems on quasi-finite structures from [9].

1 Introduction

This paper is based on my talk in the Oléron meeting in June 2011. It is intended as an invitation to model theorists to explore areas of common interest with additive combinatorics. Many of the streams leading up to this confluence were opened up by Anand Pillay; the paper is dedicated to him, with friendship and appreciation.

The relationship of model theory and modern combinatorics starts very near the birth of both: Ramsey's theorem [31, Theorem A], proved as a lemma for one of the first true theorems of model theory, rediscovered by Erdős, later settling down as part of the backbone of both subjects. Shelah, in his foundational work, was—and remains—keenly aware of combinatorial connections. In recent decades, with model theory turning to face group theory and geometry as well as its own intrinsic issues, the contact with combinatorics seemed limited to special areas. In reality, however, broad trends in the two subjects moved in parallel, unaware of each other but concerned with very similar questions.

Geometric stability theory was born with pseudo-finite structures. Zilber studied them under additional hypotheses: $\aleph_0$-categoricity, finite Morley rank. His main tool was a dimension theory based on Morley rank and the leading coefficient of the Zilber polynomial, counting points in finite approximations. The main themes were the dichotomy between linear and nonlinear behavior—the dividing line of modularity; the study of modularity in the almost strongly minimal setting, with the proof that totally categorical strongly minimal sets are modular; and a theory allowing information to be lifted from rank one to finite Morley rank.

Received October 7, 2012; accepted December 11, 2012
2010 Mathematics Subject Classification: Primary 03C13; Secondary 03C45, 11P99
Keywords: pseudo-finite, dimensions, stability
© 2013 by University of Notre Dame 10.1215/00294527-2143952

Zilber's theory was generalized, firstly to the superstable and stable settings, using Shelah's powerful theories of forking, domination, and regular types. A little later, pursuing conjectures of Lachlan, and influenced by the classification of the finite simple groups that seemed to have been completed at that time, the monograph Cherlin and Hrushovski [9] generalized Zilber's theory beyond the stable setting. This was the first view of geometric simplicity theory; 3-amalgamation was seen as the main principle, generalizing the canonical 2-amalgamation of stability. Once again, the theory was generalized to the supersimple and simple settings. A completely new idea was required: the compact Lascar group, as the unique obstruction to 3-amalgamation (see Kim and Pillay [23]).

A few years ago, listening to a talk by Bourgain on the sum-product phenomenon from the point of view of analysis, combinatorics, and computer science, I noted a parallel between these questions and the stabilizer theorem of model theory, closely related to 3-amalgamation. However, combinatorialists work with *arbitrary subsets*, while structural model theory was developed under restrictive assumptions of stability or simplicity. At first sight it did not seem possible to strengthen the analogy to a precise and useful bridge.

But it happens frequently in model theory that an apparently purely restrictive condition is shown to imply a positive structure, which is then the real vehicle of further results. The paradigmatic example is stability, defined in terms of *few types*. This is shown to imply a rich new structure on the algebra of formulas or the space of types; one presentation is the *relative tensor product* operation on types (with accompanying explanations regarding strong types, Lascar types, and so on). But such tensor structures can arise in other ways, without the limitation on the type space. Another example is superstability and the regular type decomposition. Once such additional structures are discovered, model theory can profitably investigate them as objects of independent interest and not work exclusively with the restrictive assumptions that gave rise to them. I have long felt so for intrinsic reasons of model-theoretic development; certainly it is essential for applications to combinatorics. To what extent can simplicity theory be stripped of the assumption that *every* type has a finite nonforking base and allow its ideas to be transposed to the pseudo-finite setting?

In Hrushovski [16], it was shown that this can be done for 3-amalgamation and the stabilizer theorem, as well as the connection to locally compact groups. It turned out that the stabilizer theorem itself was worked out independently by combinatorialists, culminating with Sanders [32].

The present paper revolves around modularity and the Zilber dichotomy. We take some tentative first steps towards a definition in the pseudo-finite setting. In Section 2, we recall the pseudo-finite dimensions that form the entry point of the model-theory/combinatorics connection considered here. Corresponding to pure model theory versus the model theory of enriched fields, we consider them both abstractly and in a setting where they are dominated by Zariski dimension.

Section 3 includes an exposition, based on [4], of a result of Breuillard, Green, and Tao and of Pyber and Szabó [29] on pseudo-finite subgroups of simple algebraic groups. Like the stabilizer theorem for groups, one could formulate a version on definability properties of hyperdefinable sets with sufficient geometric richness, reminiscent of Buechler's dichotomy; this will be done elsewhere.

Another example of the efficacy of pseudo-finite dimensions as an embedded dimension theory is included in Section 4. We deduce in somewhat generalized form a result of Breuillard, de Cornulier, Lubotzky, and Meiri [2] on escape from *many* subvarieties. From a model-theoretic viewpoint, the finiteness (i.e., boundedness) condition on the set of subvarieties is replaced by a condition of *zero-dimensionality*. This is an extremely elementary case of a potential development whose interest in general would be high.[1]

In Section 5 we consider modularity *at a single scale* and develop an analogy with the almost strongly minimal case. We leave for later consideration the study of modularity and ampleness properties and their consequences at more than one scale. Eventually one can hope to have, for example, the results of Section 3 for groups follow from a general theory (related to a line of results on definability of $\bigwedge$-definable groups, socle lemmas, and generalizations; see Blossier, Martin-Pizzaro, and Wagner [6]).

In the final section, we return to the quasi-finite setting. Our assumptions here are much stronger, but on the other hand we work only with pseudo-finite dimensions and cannot use Zariski dimension as a tool. In [9] the classification of the finite simple groups was used, via work of Kantor, Liebeck, and Macpherson [21]. It was shown, however, that the relevant part of the classification is equivalent to certain model-theoretically meaningful properties of quasi-finite structures. It remains a significant goal to find direct and conceptual proofs of these, beginning with modularity, and it seems to be the right time to take a new look. We present the list of properties, slightly reformulated, and give the new but easy proof of one of them.

Another subject addressed in the talk was higher amalgamation and the connection to combinatorics (e.g., triangle elimination). See the interesting results of the last section of Goldbring and Towsner [13]. Studying the obstruction to 3-amalgamation leads again to locally compact groups, with further connections to dynamics (e.g., weak mixing). There are many open problems here, but this direction will be taken up elsewhere. (We mention here only, answering a question in the talk, that the replacement theorem of [16] is valid with compact Lascar types for any definable measure, over any base.)

2 Fine and Coarse Pseudo-Finite Dimensions

The logarithm of the cardinality of a nonstandard finite set behaves like a dimension theory, as soon as one factors out a nontrivial convex subgroup of the nonstandard reals.

Let u be a nonprincipal ultrafilter on a set I, let $(M_i : i \in I)$ be L-structures, and let M be their ultraproduct. If D is a nonempty definable set such that $D(M_i)$ is finite for almost all i, then we have a nonstandard finite cardinality for D, an element $|D|$ of the ultrapower $\mathbb{R}^*$ of $\mathbb{R}$ along u. Let C be a nonzero convex subgroup of $\mathbb{R}^*$. We define

$$\delta_C(D) = \log |D| + C$$

the image of $\log |D|$ in $\mathbb{R}^*/C$.

In [16], various degrees of graininess are treated uniformly, with C treated as a parameter. Here we specialize to two cases. The convex hull of the standard reals is denoted C_{fin}; the corresponding dimension theory is δ_{fin}. The characteristic feature of $\delta = \delta_{\mathrm{fin}}$ is that any dimension comes with a *real-valued measure, defined up to* *a scalar multiple*. It is characterized (up to a scalar multiple) by $\mu_\alpha(X) = 0$ iff $\delta(X) < \alpha$, $\mu_\alpha(X) = \infty$ iff $\delta(X) > \alpha$, and when $\delta(X) = \delta(Y) = \alpha$,

$$\mu_\alpha(X) = \mathrm{st}\big(|X|/|Y|\big)\mu_\alpha(Y),$$

$\mathrm{st} : \mathbb{R}^*_{\geq 0} \to \mathbb{R}_\infty$ being the standard part homomorphism. If $\delta(X) = \alpha$ and $\delta(Y) = \beta$ and μ_α, μ_β are chosen, there is a unique representative measure $\mu_{\alpha+\beta}$ with $\mu_{\alpha+\beta}(X \times Y) = \mu_\alpha(X)\mu_\alpha(Y)$.

The information contained in δ_{fin} and the families μ_α can be gathered together as a single invariant δ_{C_0}, where C_0 is the convex group of infinitesimals; $\delta_{C_0}(X)$ is a more canonical version of the pair $(\delta_{\mathrm{fin}}(X), \mu_{\delta_{\mathrm{fin}}(X)})$. An element $\tilde{\alpha} \in \mathbb{R}^*/C_0$ sitting above $\alpha \in \mathbb{R}/C_{\mathrm{fin}}$ determines a measure $\mu_{\tilde{\alpha}}$, finite on definable sets X with $\delta_{\mathrm{fin}}(X) \leq \alpha$. We will, however, work with the more familiar presentation as dimensions and measures.

On the other hand, if we have in mind some definable set X, with $\log|X| = \alpha$, let $C_{<\alpha}$ be the coarsest possible dimension that does not give X dimension 0; that is, $C_{<\alpha}$ is the maximal convex subgroup of $\mathbb{R}^*$ not containing α. Let C_α be the smallest convex subgroup containing α, and restrict attention to definable sets Y with $\log|Y| \in C_\alpha$. The corresponding dimension theory can be viewed as *real valued*, using the natural isomorphism $C_\alpha/C_{<\alpha} \to \mathbb{R}$ mapping α to 1. It can be defined directly:

$$\boldsymbol{\delta}(Y) = \mathrm{st}\big(\log|Y|/\alpha\big).$$

The coarse pseudo-finite dimension $\boldsymbol{\delta} = \boldsymbol{\delta}_\alpha$ depends of course on the choice of α, but in this paper we will always treat one coarse pseudo-finite dimension at a time, and view α as fixed.

2.1 Continuity It is often desirable to have $\boldsymbol{\delta}$ and μ_α continuous in the sense of real-valued logic. Recall that if X is a definable set, T is a locally compact topological space, and $\varphi : X \to T$ is a function, we say that φ is *continuous* if for any $C \subset U \subset T$ with C compact and U open, there exists a definable set D with $\varphi^{-1}(C) \subset D \subset \varphi^{-1}(U)$. A *measure* μ on Y is said to be continuous if whenever W is a $\emptyset$-definable subset of $Y \times X$, and $W(a) = \{y : (y,a) \in W\}$, the function $a \mapsto \mu(W(a))$ is continuous. Similarly, we say that $\boldsymbol{\delta}$ is continuous if in this situation, for any $\alpha < \beta \in \mathbb{R}$ there is some $\emptyset$-definable D with

$$\big\{a : \boldsymbol{\delta}\big(W(a)\big) \leq \alpha\big\} \subset D \subset \big\{a : \boldsymbol{\delta}\big(W(a)\big) < \beta\big\}.$$

Both of these are easily seen to be true if the language is *closed under cardinality comparison quantifiers*. This means that for any formula $\varphi(x,y)$ there is some formula $\theta(x,x') =: (CCy)\varphi$ such that $M_i \models \theta(a,a')$ for almost all i iff $|\varphi(M,a)| \leq |\varphi(M,a')|$. For all combinatorial applications I am aware of, we can simply close the language under this kind of quantifier so as to obtain continuity.[2]

When the measure μ is continuous, and $\varphi : X \to \mathbb{R}$ is continuous, we can define $\int \varphi \, d\mu$. We have (with compatible normalizations, as above) $\mu_{\alpha+\beta}(W) = \int \mu_\beta(W(a))\, d\mu_\alpha(a)$ when W is a definable subset of $Y \times X$, $\delta_{\mathrm{fin}}(X) = \alpha$, and $\delta_{\mathrm{fin}}(Y) = \beta$ (or more generally for maps $\pi : W \to X$ such that each fiber has $\delta_{\mathrm{fin}} \leq \beta$).

We define $\boldsymbol{\delta}$ and μ_α for partial types by taking the infimum of their values over larger definable sets.

Continuity implies in particular that $\boldsymbol{\delta}(\varphi(x,c))$ depends only on $\mathrm{tp}(c)$. In particular, if P, Q and $I \leq P \times Q$ are complete types and $I(b) = \{a : (a,b) \in I\}$, then $\boldsymbol{\delta}(I(b))$ is a *constant*. Moving to complete types in the usual way, this will allow us at appropriate moments to ignore analytic aspects and concentrate on the geometry of the situation.

The following statements (Sections 2.2–2.6), included partly for reference, were given as exercises in a class on pseudo-finite structures; we leave them in this form. Section 2.5 was initially marked as unchecked and was worked out by Elad Levi.

2.2 Let $\delta = \delta_C$ be any pseudo-finite dimension. Let $f : X \to Y$ be a surjective definable map, and let $E = \{(x,x') \in X^2 : fx = fx'\}$. Show that $2\delta(X) \leq \delta(E) + \delta(Y)$.

Hint: The first proof is valid for $\boldsymbol{\delta}$ or any continuous dimension. For a complete type q of elements of Y, show equality:

$$2\delta\big(f^{-1}(q)\big) = \delta\big(E \cap f^{-1}(q)^2\big) + \delta(q).$$

Now $2\delta(X) = \sup_q 2\delta(f^{-1}(q))$, $\delta(E) = \sup_q \delta(E \cap f^{-1}(q)^2)$, $\delta(Y) = \sup_q \delta(q)$.

The alternative proof is valid for any pseudo-finite dimension δ_C: working in $L^2(Y_i)$, let $a(y) = |f^{-1}(y)|$, $1(y) = 1$, $(u,v) = \sum_{y \in Y} u(y)v(y)$. By Cauchy–Schwarz,

$$|X_i|^2 = (a,1)^2 \leq (a,a)(1,1) = |E_i||Y_i|.$$

2.3 Let M be any countably saturated structure in a countable language, and let Γ be a countable intersection of M-definable sets. Assume that Γ forms a subset of a definable group G of M. Show that Γ is a subgroup of G if and only if there exist definable sets X_n for $n \in \mathbb{N}$ with $1 \in X_n = X_n^{-1}$, $X_{n+1}X_{n+1} \subset X_n$, and $\Gamma = \bigcap_{n \in \mathbb{N}} X_n$.

For any set D, let $\Delta_D = \{(x,x) : x \in D\}$ (the diagonal), and for $R \leq D^2$ let $R^t = \{(x,y) : (y,x) \in R\}$, and let $R \circ S = \{(x,z) : (\exists y)((x,y) \in S, (y,z) \in R)\}$.

2.4 Let M be any countably saturated structure in a countable language, and let E be a countable intersection of M-definable sets. Assume $E \subset D^2$ for some definable set D. Show that E is an equivalence relation if and only if $E = \bigcap_{n \in \mathbb{N}} X_n$ for some sequence $(X_n : n \in \mathbb{N})$ of definable subsets of D^2 with $X_n = X_n^t$ and $X_{n+1} \circ X_{n+1} \subset X_n$. The domain of E will be $\bigcap_n D_n$, where $D_n = \{x \in D : (x,x) \in X_n\}$.

The next exercise makes sense of the dimension of a hyperimaginary set. See the proof of Theorem 3.1 for an instance where it occurs naturally.

2.5 Define $\boldsymbol{\delta}(X)$ for sets X of the form Y/E, with Y a $\bigwedge$-definable set and E a $\bigwedge$-definable equivalence relation on Y. If Y, E are defined over the countable set A, let $\boldsymbol{\delta}(Y/E) = \sup_q 2\boldsymbol{\delta}(q) - \boldsymbol{\delta}(q^2 \cap E)$; the supremum ranges over all types over A of elements of Y. Show that this does not depend on the choice of A. If f is a definable map inducing an injection $Y/E \to Y'/E'$, show that each fiber has a similar form; if each fiber F has $\boldsymbol{\delta}(F) \leq \gamma$, show that $\boldsymbol{\delta}(Y/E) \leq \boldsymbol{\delta}(Y'/E') + \gamma$. If $E = \bigcap_n X_n$ where X_n, D_n are as in Section 2.4, and if Z_n is a definable subset of D_n which is a maximal anticlique for X_n, show that $\boldsymbol{\delta}(Y/E) = \lim_n \boldsymbol{\delta}(Z_n)$. (Note that such definable sets Z_n need not exist outright but will exist in some expansion which is still an ultraproduct; so $\lim_n \boldsymbol{\delta}(Z_n)$ does not depend on the expansion.)

2.6 Let G be a simple algebraic group, and let Γ be a definable, Zariski-dense subgroup of G with fine pseudo-finite dimension $\delta_{\text{fin}}(G) < \infty$. Let $N \neq 1$ be a normal subgroup of Γ. Then N has finite index in Γ.

Hint: Show that if $a \neq 1$, then a^G generates G; for some even n, the map $c : G^n \to G$,

$$c(g_1, \ldots, g_n) = (g_1^{-1} a^{-1} g_1)(g_2^{-1} a g_2) \cdots (g_n^{-1} a g_n)$$

is surjective. Let U be a Zariski-open subset of G such that $\dim c^{-1}(u)$ is constant on U. If $a \in N$, show that $\delta_{\text{fin}}(c^{-1}(u) \cap \Gamma^n) \leq \delta_{\text{fin}}(\Gamma) \dim c^{-1}(u)/\dim(G)$, and conclude that $\delta_{\text{fin}}(c(\Gamma^n)) = \delta_{\text{fin}}(\Gamma)$. Hence $\delta_{\text{fin}}(N) = \delta_{\text{fin}}(G)$.

2.7 Coarse pseudo-finite dimension Lemmas 5.1–5.7 of [16] apply to $\boldsymbol{\delta}$; we recall some of the statements. We assume continuity of $\boldsymbol{\delta}$; it is used in Lemmas 2.8(4) and 2.10. We assume a normalization with $\boldsymbol{\delta}(X) = 1$.

Lemma 2.8 *We have the following.*

(1) $\boldsymbol{\delta}(Y) = 0$ *for finite* Y.
(2) $\boldsymbol{\delta}(Y' \cup Y) = \max(\boldsymbol{\delta}(Y'), \boldsymbol{\delta}(Y))$.
(3) $\boldsymbol{\delta}(Y \times Y') = \boldsymbol{\delta}(Y') + \boldsymbol{\delta}(Y)$.
(4) More generally, if f *is a definable function on* Y,

$$\boldsymbol{\delta}(Y) = \sup\big\{\alpha + \beta : \alpha \in \mathbb{R}_\infty, \beta = \delta\big\{z : \boldsymbol{\delta}\big(f^{-1}(z)\big) \geq \alpha\big\}\big\};$$

this holds for $Y \to Y/E$ *even for a* $\bigwedge$*-definable equivalence relation* E.
(5) The definable subsets φ *of* X *with* $\boldsymbol{\delta}(\varphi) < \boldsymbol{\delta}(X)$ *form an ideal; it is the ideal of [16, Example 2.13].*

Definition 2.9

- For an element $a \in X^n$ and base set A, let

$$\boldsymbol{\delta}(a/A) = \inf \boldsymbol{\delta}\big\{\varphi : \varphi \in \text{tp}(a/A)\big\}.$$

  Note that for $\varphi \in L(A)$, $\boldsymbol{\delta}(\varphi) = \max\{\boldsymbol{\delta}(a/A) : \varphi(a)\}$. (The maximum is attained because of compactness/saturation and Lemma 2.8(5).)
- Assume $\boldsymbol{\delta}(a/A) < \infty$. Write $a \downarrow_A b$ if $\boldsymbol{\delta}(a/Ab) = \boldsymbol{\delta}(a/A)$.
- Assume $\boldsymbol{\delta}(a_i/A) < \infty, i = 1, \ldots, n$. Say $a_1, \ldots, a_n$ are *independent* over A if $\sum \boldsymbol{\delta}(a_i/A) = \boldsymbol{\delta}((a_1, \ldots, a_n)/A)$.

By definition, the relation $a \downarrow_A b$ is *transitive*, in the sense that $a \downarrow_A bc$ iff $a \downarrow_A b$ and $a \downarrow_{Ab} c$. When $\delta(a/A), \delta(b/A) < \infty$, the relation $a \downarrow_A b$ is symmetric, as follows from the next lemma.

Lemma 2.10 *We have* $\boldsymbol{\delta}(a/Ab) + \boldsymbol{\delta}(b/A) = \boldsymbol{\delta}(a, b/A)$.

Proof Let $\epsilon > 0$. Let $\varphi(x, y) \in \text{tp}(a, b/A)$ be such that $\boldsymbol{\delta}(\varphi) - \boldsymbol{\delta}(a, b/A) < \epsilon$ and $\boldsymbol{\delta}(\varphi(x, b)) - \boldsymbol{\delta}(a/Ab) < \epsilon$. Let $\theta(y) \in \text{tp}(b)$ be such that $\boldsymbol{\delta}(\theta) - \boldsymbol{\delta}(b/A) < \epsilon$ and (using the continuity mentioned above) such that for any b' with $\theta(b')$, $|\boldsymbol{\delta}(\varphi(x, b')) - \boldsymbol{\delta}(\varphi(x, b))| < \epsilon$. We may replace $\varphi(x, y)$ by $\varphi(x, y) \wedge \theta(y)$. Clearly, $|\boldsymbol{\delta}(\varphi(x, y)) - (\boldsymbol{\delta}(\varphi(x, b)) + \boldsymbol{\delta}(\theta))| < \epsilon$. So $|\boldsymbol{\delta}(a, b/A) - (\boldsymbol{\delta}(a/Ab) + \boldsymbol{\delta}(b/A))| < 4\epsilon$. $\square$

From symmetry and transitivity it follows that if $\boldsymbol{\delta}(a/A), \boldsymbol{\delta}(b/A), \boldsymbol{\delta}(c/A) < \infty$, $a \downarrow_A b, c$, and $b \downarrow_A c$, then $a \downarrow_A b$ and $a, b \downarrow_A c$.

We also have an existence statement for independent extensions.

Lemma 2.11 *Assume $a \downarrow_A b$. Then for any c there exists a' with* $\mathrm{tp}(a'/Ab) = \mathrm{tp}(a/Ab)$ *and $a' \downarrow_A b, c$.*

Proof Let $\alpha = \boldsymbol{\delta}(a/A) = \boldsymbol{\delta}(a/Ab)$. We must find a' with $\mathrm{tp}(a'/Ab) = \mathrm{tp}(a/Ab)$ and $\boldsymbol{\delta}(a'/Abc) = \alpha$. In other words, a' must realize $\mathrm{tp}(a/Ab)$ along with $\{\neg\psi : \psi \in L(Abc), \boldsymbol{\delta}(\psi) < \alpha\}$. Since every formula in $\mathrm{tp}(a/Ab)$ has $\boldsymbol{\delta}$ at least α, this is consistent. □

2.12 Embedded pseudo-finite dimension We now assume that a *field structure* is included in the language, and the definable sets we are interested in are pseudo-finite subsets of algebraic varieties. (Or more generally, that some reduct of the theory has finite Morley rank.)

Define the $\boldsymbol{\delta}$-closure $\mathrm{cl}_{\boldsymbol{\delta}}(A)$ to be the set of all b with $\boldsymbol{\delta}(b/A) = 0$. Here we relax a little the usual convention that base sets are small, but there will be no difficulty translating the statements to ones about finite subsets of $\mathrm{cl}_{\boldsymbol{\delta}}(A)$. Let $d(a/A)$ be the smallest dimension of an irreducible variety W with $a \in W$, such that W is defined over $\mathrm{cl}_{\boldsymbol{\delta}}(A)$, that is, some Ab with $\boldsymbol{\delta}(b/A) = 0$. (Thus if A is $\boldsymbol{\delta}$ -closed, then $d(a/A)$ is transcendence degree.) Let $U(a/A)$ denote the variety W; it is obviously unique, so defined over Aa as well as over Ab.

$\boldsymbol{\delta}$-independence implies Zariski independence over a $\boldsymbol{\delta}$-closed A.

Lemma 2.13 *Let $A = \mathrm{cl}_{\boldsymbol{\delta}}(A)$. Then $a \downarrow_A b$ implies $d((a,b)/A) = d(a/A) + d(b/A)$.*

Proof Let V_c be an irreducible variety defined over some c with $a \downarrow_A c$, such that $a \in V_c$, and with $\dim(V_c)$ least possible. Using Lemma 2.11, find c' with $a, c \downarrow_A c'$ and with $\mathrm{tp}(c'/Aa) = \mathrm{tp}(c/Aa)$. So $a \in V_{c'}$, and hence $a \in V_c \cap V_{c'}$. But by transitivity we have $a \downarrow_A c, c'$. So $\dim(V_c \cap V_{c'}) = \dim(V_c)$, and it follows that $V_c = V_c \cap V_{c'} = V_{c'}$. Let d be a canonical parameter for V_c (i.e., Ad is a field of definition for V_c, over A). Then $d \in Ac \cap Ac'$. Since $c \downarrow_A c'$, it follows that $2\boldsymbol{\delta}(c/A) = 2\boldsymbol{\delta}(c/Ad) + \boldsymbol{\delta}(d/A)$ and also $\boldsymbol{\delta}(c/A) = \boldsymbol{\delta}(c/Ad) + \boldsymbol{\delta}(d/A)$. This forces $\boldsymbol{\delta}(d/A) = 0$. By minimality of $\dim(V_c)$ and the definition of $d(a/A)$ we have

$$d(a/Ab) \geq \dim(V_c) \geq d(a/A).$$

Clearly, $d(a/A) \geq d(a/Ab)$, so $d(a/Ab) = d(a/A)$. The equality now follows from Lemma 2.10. □

In fact, this may be generalized.

Lemma 2.14 *Let A be $\boldsymbol{\delta}$-closed. If $a \downarrow_A b$, then a, b are independent over A in the sense of any stable formula $\varphi(x, y)$ over A.*

Proof The proof is as above, using local rank for the stable formula φ in place of dimension. □

2.15 The Larsen–Pink inequality We recall the Larsen–Pink inequality, valid for any pseudo-finite dimension δ_C (see Hrushovski and Wagner [19], [16, Section 5]).

Let G be a group of finite Morley rank, and let Γ be a subgroup of G. We assume that Γ is *Zariski dense*, that is, not contained in any proper definable subgroup of G.

If G is simple, or a finite power of a simple group, and $\delta_C(G) < \infty$, then for any definable $Z \leq G$, we have

$$\delta_C(\Gamma \cap Z) \leq \frac{\dim(Z)}{\dim(G)}\delta_C(\Gamma).$$

In particular, this holds for $\boldsymbol{\delta}$.

In [16], this is generalized to the case of a $\bigwedge$-definable subgroup of strict dimension, that is, a sequence Y_n of definable sets, with $Y_{n+1}Y_{n+1}^{-1} \subset Y_n$ and $\delta_C(Y_1) = \delta_C(Y_2) = \cdots$; the statement in this case is that $\delta_C(Z \cap Y_n) \leq \delta_C(Y_1)\dim(Z)/\dim(G)$, for large enough n (see Proposition 5.6).

The strictness of dimension is an unnecessary assumption in the case of $\boldsymbol{\delta}$; it suffices to assume that $\boldsymbol{\delta}(Y_n)$ is bounded. For we can take a further ultraproduct or saturation and consider Y_{n^*} for nonstandard n^*; these all have the same $\boldsymbol{\delta}$, namely, $\boldsymbol{\delta}(Y_{n^*}) = \boldsymbol{\delta}(\bigcap_{n\in\mathbb{N}} Y_n)$; letting $\Gamma' = \bigcap_{k\in\mathbb{N}} Y_{n^*+k}$ we obtain a slightly smaller Zariski-dense subgroup, this time of strict dimension.

It is possible to restate the lemma "from below":

For $X \subset G$, $r \in \mathbb{R}$, let $X^{\cdot r} = \{x_1 \cdot \ldots \cdot x_k : x_1, \ldots, x_k \in X, k \leq r\}$.

Proposition 2.16 *Let G be a simple algebraic group, and let Y be a Zariski-dense pseudo-finite set with $0 < \boldsymbol{\delta}(Y) < \infty$. Assume $1 \in Y = Y^{-1}$. Then for any subvariety Z of G,*

$$\boldsymbol{\delta}(Y \cap Z)\dim(G) \leq \limsup_n \boldsymbol{\delta}(Y^{\cdot n})\dim(Z).$$

Proof Suppose otherwise; so for infinitely many n we have $\boldsymbol{\delta}(Y^{\cdot n}) < \frac{\dim(G)}{\dim(Z)} \times \boldsymbol{\delta}(Y \cap Z)$. In particular, $\boldsymbol{\delta}(Y^{\cdot n})$ is bounded. Let W_k be an ultraproduct (with respect to an ultrafilter on the set of $n \in \mathbb{N}$) of the $Y^{\cdot[2^{-k}n]}$. So $W_{k+1}^2 \subset W_k$, $\bigcap_k W_k$ is a Zariski-dense $\bigwedge$-definable subgroup of G, and $\boldsymbol{\delta}(W_k) = \boldsymbol{\delta}(W_{k+1}) = \cdots$. The version discussed above of [16, Proposition 5.6] applies to $\Gamma = \bigcap_k W_k$ and shows that $\boldsymbol{\delta}(\Gamma \cap Z)\dim(G) \leq \boldsymbol{\delta}(\Gamma)\dim(Z)$. We have $\boldsymbol{\delta}(Y \cap Z) \leq \boldsymbol{\delta}(\Gamma \cap Z)$. On the other hand $\boldsymbol{\delta}(Y^{\cdot n}) \to_n \boldsymbol{\delta}(\bigcap W_k)$. The proposition follows. □

3 Approximate Subgroups of Linear Groups

Breuillard, Green, and Tao [3] and Pyber and Szabó [29] proved a decisive result on approximate subgroups of linear groups. From a model-theoretic viewpoint, it falls into a line of results about definability of $\bigwedge$-definable groups. For *fine* pseudo-finite dimension, the result was first proved in [16]. The proof used the "stabilizer theorem" inspired by stability and simplicity. By utilizing more of the geometry of the situation [3] proved it for coarse pseudo-finite dimension. This is much stronger: combinatorially, it amounts to showing that for a typical subset X of size n one can expect XX to have size $n^{1+\epsilon}$, as opposed to $(1+\epsilon)n$. It turned out that the main input of the model-theoretic proof for linear groups was not the deeper, stability-inspired results but simply the pseudo-finite dimension-theoretic framework and basic properties of the dimension theories. We describe the proof in this language, restricting attention to the principal case, of simple algebraic groups G or powers of such groups.

In this section and the next (with an exception in Exercise 3.6), we use *coarse* pseudo-finite dimension. Fix α, and denote the coarse pseudo-finite dimension $\boldsymbol{\delta}_\alpha$ by $\boldsymbol{\delta}$. As explained in Section 2, we may take $\boldsymbol{\delta}$ to be continuous.

In order to bring out the precise geometry used in the proof, we state it axiomatically, in the context of groups of finite Morley rank. It is the geometry of the maximal tori, going back to (at least) Jordan's proof in [20].

Let $M = (G, \cdot, X, \ldots)$ be an expansion of a group $(G, \cdot)$, that is, a structure including a group $(G, \cdot)$, a definable set $X \subset G$, and possibly more. We assume $(G \cdot)$ has finite Morley rank g, while X is pseudo-finite, with continuous coarse dimension $\boldsymbol{\delta}$ satisfying $\boldsymbol{\delta}(X) = 1$. (This can be achieved by adding dimension quantifiers.) We write dim for Morley rank. We also assume that G is simple nonabelian, or a power of simple groups; this is in order for the Larsen–Pink inequality to be valid.

We say a subgroup H is *Zariski-dense* if it is contained in no proper $(G, \cdot)$-definable subgroup of G.

For $a \in G$, let $T_a = C_G(a)$. Let $T_a^r = \{b \in T_a : T_a = T_b\}$. Let $t = \dim(T_a)$ for a generic $a \in G$. Let $N_G(T_a)$ be the normalizer. Assume the following geometry of centralizers.

$\diamond$ For a generic $a \in G$ and generic $b \in T_a$, we have $T_a = T_b$, and $N_G(T_a)/T_a$ is finite. In other words, let $R = \{a \in G : \dim(N_G(T_a)) = \dim(T_a) = t, \dim(T_a \setminus T_a^r) < t\}$. Then $\dim(R) = g$.

Theorem 3.1 *Let Γ be a Zariski-dense $\bigwedge$-definable (in M) subgroup of G, with $0 < \boldsymbol{\delta}(\Gamma) < \infty$. Then there exists a definable subgroup S of G containing Γ, with $\boldsymbol{\delta}(S) = \boldsymbol{\delta}(\Gamma)$.*

Proof It is clear that if $b \in T_a^r$, then $b \in R$; that is, R is the disjoint union of T_a^r over $a \in R$.

We have $\dim(G \setminus R) \leq g - 1$ and $\dim(T \setminus R) \leq t - 1$.

Recall also that for centralizers $T = T_a$, the Larsen–Pink inequality becomes an equality:

$$\delta(T \cap \Gamma) = \frac{\dim T}{\dim G} \delta(\Gamma).$$

Hence we have the same for $\boldsymbol{\delta}$.

Let $\Upsilon = \{C_G(a) : a \in R \cap \Gamma\}$. Clearly, Υ is Γ-conjugation-invariant. We will show that Υ is definable, that is, $\{b : C_G(b) \in \Upsilon\}$ is definable, using a dimension gap. Let X be a definable subset of G, $\Gamma \subseteq X$, with $\boldsymbol{\delta}(XX) - \boldsymbol{\delta}(\Gamma) < \frac{\boldsymbol{\delta}(\Gamma)}{2\dim(G)}$ (see Sections 2.3, 2.4). Let $g = \dim(G)$, and let $t = \dim(T)$ for (any) $T = C_G(a), a \in R$. Note that $\boldsymbol{\delta}(\Gamma \setminus R) \leq \frac{g-1}{g}\boldsymbol{\delta}(\Gamma) < \boldsymbol{\delta}(\Gamma)$ by the Larsen–Pink inequality; similarly, $\boldsymbol{\delta}((T \setminus R) \cap \Gamma) \leq \frac{t-1}{g}\boldsymbol{\delta}(\Gamma) = \frac{t-1}{t}\boldsymbol{\delta}(T \cap \Gamma) < \boldsymbol{\delta}(T \cap \Gamma)$.

Claim 1 *Let $a \in R$, $T = C_G(a)$. Then $\boldsymbol{\delta}(T \cap X) - \boldsymbol{\delta}(T \cap \Gamma) \leq \boldsymbol{\delta}(X\Gamma) - \boldsymbol{\delta}(\Gamma) < \frac{\boldsymbol{\delta}(\Gamma)}{2g}$.*

Proof The first inequality is obvious if we extend $\boldsymbol{\delta}$ to quotients such as T/Γ (see Section 2.4). The inclusion map $T \cap X \subset X$ induces an injective map $(T \cap X)/(T \cap \Gamma) \to X/\Gamma = X\Gamma/\Gamma$; hence

$$\boldsymbol{\delta}(T \cap X) - \boldsymbol{\delta}(T \cap \Gamma) = \boldsymbol{\delta}\big((T \cap X)/(T \cap \Gamma)\big) \leq \boldsymbol{\delta}(X\Gamma/\Gamma) = \boldsymbol{\delta}(X\Gamma) - \boldsymbol{\delta}(\Gamma).$$

But let us also give a proof independent of Section 2.4. Define a map $f : (T \cap X) \times \Gamma \to X\Gamma$ by $f(t, y) = ty$. If $f(t, y) = f(s, z) = c$, then $s^{-1}t = zy^{-1} \in T \cap \Gamma$, so the projection map $(t, y) \mapsto t$ takes $f^{-1}(c)$ to a coset of $T \cap \Gamma$, injectively; thus

$\delta(f^{-1}(c)) \leq \delta(T \cap \Gamma)$. So $(\delta(T \cap X) + \delta(\Gamma)) \leq \delta(T \cap \Gamma) + \delta(X\Gamma)$, and the first inequality follows. The second results from the assumption $\delta(XX) - \delta(\Gamma) < \frac{\delta(\Gamma)}{2\dim(G)}$. □

Claim 2 *Let $T = C_G(a)$, $a \in R$. Then*

$$T \in \Upsilon \iff \delta(T \cap X) > \frac{t - 1/2}{g}\delta(\Gamma) \iff \delta(T \cap X) \geq \frac{t}{g}\delta(\Gamma).$$

Proof If $T = C_G(a)$, $a \in R \cap \Gamma$, then $\delta(\Gamma \cap T) \geq \frac{t}{g}\delta(\Gamma)$ by Larsen–Pink. Conversely, let $T = C_G(a)$, $a \in R$, and assume $\delta(T \cap X) > \frac{t-1/2}{g}\delta(\Gamma)$. By Claim 1,

$$\delta(T \cap \Gamma) > \delta(T \cap X) - \frac{\delta(\Gamma)}{2g} > \frac{t-1}{g}\delta(\Gamma) \geq \frac{\dim(T \setminus R)}{g}\delta(\Gamma).$$

But $\delta((T \setminus R) \cap \Gamma) \leq \frac{\dim(T\setminus R)}{g}\delta(\Gamma)$ (by the Larsen–Pink inequality for the Zariski closure of $T \setminus R$). So $T \cap \Gamma \cap R \neq \emptyset$, and the claim follows by choosing any $a' \in T \cap \Gamma \cap R$, noting that $C_G(a') = C_G(a)$. □

By Claim 2 and the continuity of δ, Υ is definable. Hence the normalizer $S = N(\Upsilon)$ is a definable group, and it contains Γ. We have $\delta(\Upsilon) < \delta(R \cap \Gamma) < \infty$; we want to show that $\delta(S)$ is finite as well.

Claim 3 *$\delta(S)$ is finite.*

Proof Let Z be the intersection of all $T \in \Upsilon$, and let $Z = T_1 \cap \cdots \cap T_k$ for some $T_1, \ldots, T_k \in \Upsilon$. Then any element of $\Gamma \cap R$ commutes with any element of Z. As $\delta(\Gamma \setminus R) < \delta(\Gamma)$, any element g of Γ is a product of two elements of $\Gamma \cap R$ (i.e., $g(\Gamma \cap R) \cap (\Gamma \cap R) \neq \emptyset$). So Z commutes with every element of Γ; since Γ is Zariski-dense, Z commutes with every element of G, so Z is finite. We have a map $S \to \Upsilon^k$, $s \mapsto (s^{-1}T_1 s, \ldots, s^{-1}T_k s)$. The fibers of this map are cosets of $N(T_1) \cap \cdots \cap N(T_k)$, so they are finite. Hence $\delta(S)$ is finite. □

We have a surjective map $R \cap \Gamma \to \Upsilon, a \mapsto C_G(a)$, whose fibers are of the form $T \cap R \cap \Gamma$ for some $T \in \Upsilon$. Now $\delta(T \cap R \cap \Gamma) = \frac{t}{g}\delta(\Gamma)$ (since $\dim(T\setminus R) < \dim(T)$, and using Larsen–Pink). Hence $\frac{t}{g}\delta(\Gamma) + \delta(\Upsilon) = \delta(R \cap \Gamma) = \delta(\Gamma)$. So $\delta(\Upsilon) = (1 - \frac{t}{g})\delta(\Gamma)$.

Fix $T \in \Upsilon$. $S/(N(T) \cap S)$ embeds into Υ. $[N(T) : T]$ is finite. By the Larsen–Pink inequality for S, we have $\delta(S \cap N(T)) = \delta(S \cap T) = \frac{t}{g}\delta(S)$, and we obtain $(1 - \frac{t}{g})\delta(S) \leq \delta(\Upsilon)$. It follows that $\delta(S) \leq \delta(\Gamma)$, so equality holds. □

Corollary 3.2 (Breuillard–Green–Tao [3], Pyber–Szabó [29]) *Let $\underline{G}$ be a simple algebraic group, or a power of such a group. Let $0 < \epsilon < \epsilon'$. Then for some m, we have the following. Let F be a field, and let $G = \underline{G}(F)$. Let $X = X^{-1}$ be a subset of G. Then $|X^{\cdot m}| \geq |X|^{1+\epsilon}$, unless X is contained in a subgroup of G of size at most $|X|^{1+\epsilon'}$, or in $H(F)$ for a proper algebraic subgroup H of $\underline{G}$ of bounded complexity.*

When X is a $k = |X|^{\delta}$-approximate subgroup, that is, $XX \subset XF$ with $|F| \leq k$, we have $X^{\cdot m} \subset X \cdot F^{\cdot m-1}$, so $|X^{\cdot m}| \leq |X|^{1+m\delta}$, giving a lower bound ϵ/m for δ.

Proof Fix $0 < \epsilon < \epsilon'$, and suppose that there is no such m. Find $X_n \subset G_n = \underline{G}(F_n)$, with X_n contained in no subgroup of G_n of size at most $|X_n|^{1+\epsilon'}$ nor in any proper algebraic subgroup of complexity less than m, and $|X_n^{\cdot n}| < |X|^{1+\epsilon}$. Let $(G, F, Y_1, Y_2, \ldots)$ be a nonprincipal ultraproduct of $(G_n, F_n, X_n^{\cdot n}, X_n^{\cdot n/2}, X_n^{\cdot n/4}, \ldots)$. Then $Y_{n+1}Y_{n+1} \subset Y_n$, so $\Gamma := \bigcap_n Y_n$ is a $\bigwedge$-definable group. Normalize $\boldsymbol{\delta}$ by $\boldsymbol{\delta}(Y_1) = 1$; we have $\boldsymbol{\delta}(Y_n) \geq \boldsymbol{\delta}(Y_1)/(1+\epsilon)$. (We could easily arrange even $\boldsymbol{\delta}(Y_n) = \boldsymbol{\delta}(Y_{n+1})$.) So $0 < \boldsymbol{\delta}(\Gamma) < \infty$. Also, Y_n is contained in no definable subgroup S of G with $\boldsymbol{\delta}(S) \leq (1+\epsilon')(1+\epsilon)^{-1}\boldsymbol{\delta}(Y_n)$. And $(G, \cdot)$ has finite Morley rank and satisfies $\diamond$. Finally, Γ is Zariski-dense in $\underline{G}$.

But by Theorem 3.1, there exists a definable S with $\boldsymbol{\delta}(\Gamma) = \boldsymbol{\delta}(S)$, a contradiction. $\square$

The identity of the definable group S can be ascertained: when F has characteristic zero or sufficiently large characteristic, it is itself commensurable to a twisted algebraic group (definable over the field with a distinguished automorphism). This was first proved by Weisfeiler using the classification of the finite simple groups (see Larsen and Pink [24] for a history of this theorem and a proof by a direct interpretation, and see Hrushovsky and Pillay [18] for a model-theoretic proof in the case of prime fields).

Remark 3.3 Let X be a Zariski-dense pseudo-finite subset with $\boldsymbol{\delta}(X) = \boldsymbol{\delta}(S)$. Then by Nikolov and Pyber [27, Appendix 4], $X^3 = S$. They employ a method of Gowers, using representations of finite groups; it remains a challenge to give a model-theoretic account of this method, or another derivation of the exponent 3 here.

3.4 Sum-product phenomenon The following exercises are essentially an account of Tao [37]; but we will restrict to fields. Let K be an infinite ultraproduct of fields, as above. Let R be a $\bigwedge$-definable subring of K, with $0 < \boldsymbol{\delta}(R) < \infty$.

Exercise 3.5 If R is a subfield of K, then R is definable. Hint: Let $R^* = R \setminus (0)$ be the group of units of R. For $a \in K$, we have $a \in R^*$ iff $aR \cap R \neq (0)$ iff $\boldsymbol{\delta}(aR \cap R) = \boldsymbol{\delta}(R)$. Let X be a definable approximation to R, so that $\boldsymbol{\delta}(XX) - \boldsymbol{\delta}(R)$ is small, hence so is $\boldsymbol{\delta}(X^2) - \boldsymbol{\delta}((R^*)^2)$. As in Proposition 3.1, for $0 \neq a$ we have $a \in R$ iff $\boldsymbol{\delta}(aX \cap X)$ is large enough. Hence R is definable.

Exercise 3.6 R is definable. Hint: We already know that the field of fractions $F = \{a/b : a, b \in R, b \neq 0\}$ is definable. So we may assume that it is K and use fine pseudo-finite dimension δ_{fin}. Show that $\delta_{\text{fin}}(R) = \delta_{\text{fin}}(K)$. (Consider the map $R \times R \setminus (0) \to K$, $(a, b) \mapsto ab^{-1}$; each fiber has the same δ_{fin} as a nonzero ideal of R, and hence the same as R.) It follows that R has finite index in K, as an additive group. Let $a \in R$, $b = a^{-1}$. Then $b^{m+l} = b^m + r$ for some $m \in \mathbb{N}$, $l > 0$, $r \in R$. Multiplying by a^m we find that $b^l = 1 + ra^m \in R$, and hence $b = b^l a^{l-1} \in R$. So $R = F$ is a field, and R is definable.

4 Escape from Many Subvarieties

This section was prompted by a question from the authors of [2] and generalizes Theorem 9.1 there (see the discussion in the introduction to that paper).

For $i \in \mathbb{N}$, let G_i be a connected algebraic group over a field K_i, $\dim(G_i) = d$. Let $D = 2^{d+1}$.

Let J_i be a finite-index set, and let $(V_{ij} : j \in J_i)$ be a set of proper subvarieties of G_i. Assume that G_i, V_{ij} have bounded complexity.

Let X_i be a finite subset of $G_i(K_i)$, say, $1 \in X$ for simplicity. Let $X^{\cdot d} = \{x_1 \cdot \ldots \cdot x_d : x_1, \ldots, x_d \in X\}$. If N is a subgroup, X/N denotes the image of X in G/N.

Proposition 4.1 *Let $l \in \mathbb{N}$. Assume $X_i^{\cdot D} \subset \bigcup_{j \in J_i} V_j$, $|J_i|^i < |X_i|$. Then (for all but finitely many i) there exists a proper algebraic subgroup N_i of G_i of bounded complexity, normalized by X_i, such that $|X_i/N_i| \leq |X_i|^{1/l}$.*

In particular if G_i is simple and X_i generates a sufficiently Zariski-dense subgroup, then for large i we cannot have $X_i^{\cdot D} \subset \bigcup_{j \in J_i} V_j$, $|J|^i < |X|$.

Let (G, X, J) be an ultraproduct (over an ultrafilter u) of the (G_i, X_i, J_i). We may assume X generates a Zariski-dense subgroup of G. One could take a *definable set* to be any subset of $(X \cup J)^n$ represented by subsets Y_i of $(X_i \cup J_i)^n$. A more effective version would specify a language, including functions $\nu : X \to J$ with $x \in V_{\nu_i(x)}$, as well as relations making the dimension function $\boldsymbol{\delta}(\varphi(x, b))$ defined below continuous in $\mathrm{tp}(b)$.

For any nonempty definable set Y, represented by subsets Y_i, let

$$\boldsymbol{\delta}(Y) = \lim_u \log |Y_i| / \log |X_i| \in \mathbb{R} \cup \infty.$$

The limit is taken along the ultrafilter; equivalently, it is the standard part of the ultrapower of these numbers in $\mathbb{R}^* \cup \infty$; so $\boldsymbol{\delta}(Y)$ is a nonnegative real number or ∞.

In terms of $\boldsymbol{\delta}$, the Proposition 4.1 can be stated as follows.

Proposition 4.2 *Assume $X^{\cdot D} \subset \bigcup_{j \in J} V_j$, $\boldsymbol{\delta}(J) = 0$, $\boldsymbol{\delta}(X) < \infty$. Then there exists a definable proper subgroup S of G, normalized by X, such that $\boldsymbol{\delta}(X/S) = 0$.*

Proof Fix a base set A; recall the definition of $U(a/A)$, above Lemma 2.13. Write $U(a)$ for $U(a/A)$. Let $S_l(b) = \{x \in G : xU(b) = U(b)\}$ be the left stabilizer of $U(b)$, and similarly for S_r on the right.

For independent $a_1, \ldots, a_n \in X^{\cdot \leq D}$, $a_1, \ldots, a_n$ are also independent over $\mathrm{cl}_\delta(A)$ in the sense of d (Lemma 2.13). Let $a = a_1 \cdot \ldots \cdot a_n$. If $A' = Ae$ is an extension of A with $\boldsymbol{\delta}(e/A) = 0$ and such that $U(a_1), \ldots, U(a_n), U(a)$ is defined over A', then the field extensions $A'(a_i)$ are linearly disjoint. Thus $(a_1, \ldots, a_n)$ is a generic point of $U(a_1) \cdot \ldots \cdot U(a_n)$ over A'. As the product $a = a_1 \cdot \ldots \cdot a_n$ is an element of the variety $U(a)$, we have $U(a_1) \cdot \ldots \cdot U(a_n) \subseteq U(a)$.

Assume $a \in X^{\cdot \leq D}$. Then a lies in a proper subvariety of G defined with a parameter in J; as $\boldsymbol{\delta}(J) = 0$, we have $\dim U(a) < d$.

Let $m(i) = \max\{\dim(U(a/A)) : a \in X^i\}$. Then $m(i)$ is nondecreasing with i, and $m(i) < d$ for $i \leq D$. Since $D = 2^{d+1}$, there exists n with $m(n) = \cdots = m(2n+1)$, $2n+1 \leq D$. Fix such an n, and let $m = m(n)$. Let Σ_m be the set of $b \in X^{\cdot \leq n+1}$ with $d(b/A) = m$.

Let $b \in \Sigma_m \cap X^{\cdot \leq n}$. For $a \in X^{\cdot \leq n+1}$ with $a \downarrow_A b$, we have $\dim(U(a)U(b)) \leq \dim U(ab) \leq m = \dim(U(b))$; but $aU(b) \subseteq U(a)U(b)$; it follows that $aU(b) = U(a)U(b)$. So $a^{-1}U(a) \subseteq S_l(b)$.

Applying this for $a \in X^{\cdot n+1} \cap \Sigma_m$, $b \in X^n \cap \Sigma_m$, we obtain $a^{-1}U(a)b \subseteq S_l(b)b \subseteq U(b)$, but $\dim U(b) = m$ and $U(b)$ is irreducible, so $a^{-1}U(a) = S_l(b)$. In particular, $S = S_l(b)$ does not depend on $b \in \Sigma_m$: $S_l(b) = S_l(b')$ if $a, b, b' \in \Sigma_m$ are independent, hence in general. Similarly, $U(a) = S^*a$ is a right coset of the right stabilizer $S^* = S_r(b)$. Now for independent $a, b \in \Sigma_m$ we have

$U(a)U(b) = aSS^*b$, so $\dim(SS^*) = m$. Thus $S = S^*$. Since $aS = U(a) = S^*a$ it follows that S is normalized by a.

If $c \in X(A)$, then for $a \in X^n \cap \Sigma_m$ we have $U(ca) = cU(a)$. Thus $ca \in \Sigma_m$. So ca also normalizes S, and hence so does c. So any element of $X(A)$ normalizes S.

So far A was arbitrary. But A could have been chosen to be an elementary submodel; so as S is A-definable, if there exists an element of X that does not normalize S, then such an element exists in $X(A)$. It follows that every element of X normalizes S.

Let $\pi(a)$ denote the image of a in the quotient group G/S. Modulo S, the set $U(a)/S$ consists of one element. As $U(a)$ is defined over some extension Ab with $\boldsymbol{\delta}(b/A) = 0$, we have $\pi(a) \in \mathrm{dcl}(Ab)$, so $\boldsymbol{\delta}(\pi(a)/A) = 0$. Hence $\boldsymbol{\delta}(\pi(X)) = 0$. $\square$

5 Pseudo-Finite Sets on Algebraic Varieties

5.1 Erdős geometry studies finite sets within an algebraic variety (sometimes with constraints of reality, but we will not consider these here; see Elekes [10]). How does the set intersect subvarieties, or families of subvarieties? The statements are often asymptotic and can be stated in terms of a pseudo-finite dimension theory, embedded in Zariski dimension theory. We wish to take the first steps toward constructing a model-theoretic framework for Erdős geometry. The feeling is that such a framework may exist, comparable to the model theory of differentially closed fields or of quasi-finite structures, and that notions such as canonical base and versions of ampleness play a useful role in it.

We presently limit our attention to top-dimensional asymptotics, a severe restriction; see examples below for some of the phenomena left out. Even with this restriction there is more unknown than known, but we will be able to formulate some precise questions.

Most of what we say will apply to subsets of a variety V, if we assume that their intersection with proper subvarieties is bounded, uniformly in the degree. For simplicity, however, we consider only curves and indeed just subsets of the field itself. (This last restriction is purely for simplicity and, for what we do, loses no generality.)

Thus, let X_k be a subset of a field F_k, with $|X_k| = k$ and k approaching ∞. Throughout the section, we denote by X an ultraproduct of the sets X_k, lying in the ultraproduct F of the F_k.

We obtain functions on the set of subvarieties of $\mathbb{A}^n_F$; namely, if U is a subvariety of affine n-space, defined over F, we have $\delta_{\mathrm{fin}}(X \cap U)$ and the real-valued $\boldsymbol{\delta}(X \cap U)$. Normalize $\boldsymbol{\delta}$ by $\boldsymbol{\delta}(X) = 1$. A distant goal would be a full account of the possibilities for $\boldsymbol{\delta}$ on the affine varieties. At present we set our sights on a description of the top-dimensional part.

Our first construction will come in two versions, fine and coarse, corresponding to the two pseudo-finite dimensions.

Let F_0 be a base field. By a *variety*, in this section, it suffices to consider affine varieties defined over F_0; we view them as Zariski-open subsets of an irreducible Zariski-closed subset of $\mathbb{A}^n$. For a variety V over F_0, $V \leq \mathbb{A}^n$, let $V(X) = V \cap X^n$. Let $\boldsymbol{\mathcal{R}}[n] = \boldsymbol{\mathcal{R}}_{F_0}[n] = \boldsymbol{\mathcal{R}}_{F_0}(X)[n]$ be the family of all F_0-irreducible varieties $V \subset \mathbb{A}^n$ such that $\boldsymbol{\delta}(V(X)) = \dim(V)$; and let $\boldsymbol{\mathcal{R}} = \bigcup_n \boldsymbol{\mathcal{R}}[n]$. Similarly, let

$\mathcal{R} = \mathcal{R}_{F_0} = \mathcal{R}_{F_0}(X)$ be the family of all F_0-irreducible varieties $V \subset \mathbb{A}^n$ such that $\delta_{\text{fin}}(V(X)) = \dim(V)\delta_{\text{fin}}(X)$. Both conditions, $|V(X)| \geq |X|^{d-\epsilon}$ and $|V(X)| \geq c|X|^d$, occur in the combinatorial literature. Our goal is to understand $\boldsymbol{\mathcal{R}}$ and $\mathcal{R}$.

We first list some obvious properties of $\boldsymbol{\mathcal{R}}$ (the same hold for $\mathcal{R}$).

Lemma 5.2

(1) *For any variety $V \subset \mathbb{A}^n$, $\boldsymbol{\delta}(V) \leq \dim(V)$. (In fact $\delta(V) \leq \dim(V)\delta(X)$.)*
(2) *(Products) $\boldsymbol{\mathcal{R}}$ is closed under products of absolutely irreducible varieties; in general if $V, V' \in \boldsymbol{\mathcal{R}}$, then $V \times V'$ contains a component projecting onto V, V'.*
(3) *(Projections) $\boldsymbol{\mathcal{R}}$ is closed under projections.*
(4) *(Fibers) Let $V \in \boldsymbol{\mathcal{R}}$, $V \subset \mathbb{A}^{n+m}$. Let $\pi : \mathbb{A}^{n+m} \to \mathbb{A}^n$ be the projection, and let U be a Zariski-dense subvariety of πV. Then for some $a \in U(X)$, we have $V(a) = \pi^{-1}(a) \in \boldsymbol{\mathcal{R}}_{F_0(a)}(X)$.*
(5) *(Self-fiber products) Let $U \in \boldsymbol{\mathcal{R}}$, and let $f : U \to V$ be a dominant morphism of varieties. Let $U \times_V U$ be the fiber product, $\{(u, u') \in U \times U : fu = fu'\}$. Then there exists a component W of $U \times_V U$ projecting onto U in each direction, and with $W \in \boldsymbol{\mathcal{R}}$.*

Proof Assertion (2) is clear; (3) and (4) follow from Lemma 2.8(4). (For fine pseudo-finite dimension, one needs to use compactness to obtain a uniform bound on the fiber dimensions.) Assertion (1) now follows by induction on $\dim(V)$, using a finite-to-one projection of a Zariski-open subset of V to $\mathbb{A}^{\dim(V)}$. For (5) see Section 2.2. □

Now $\mathcal{R}, \boldsymbol{\mathcal{R}}$ can be viewed as sub*languages* of the language of fields (made relational). But model theory likes *structures*, and in general $\mathcal{R}, \boldsymbol{\mathcal{R}}$ do not have amalgamation. In the case when they do, we could formulate our questions about the corresponding universal domain. In general, it is a nontrivial task to formulate a good model-theoretic setting for the study of $\mathcal{R}, \boldsymbol{\mathcal{R}}$. We will suggest some possibilities. Before doing so, we pause to consider some relevant results and problems from the combinatorial literature. We state them directly in terms of pseudo-finite dimensions, leaving the translation to the reader.

5.3 Some theorems of additive combinatorics To delineate our subject we mention two conjectures that lie outside it, for different reasons. The "polynomial Freiman–Ruzsa conjecture" of Green [14] is usually stated for abelian groups A of prime exponent p, though a version for arbitrary abelian groups is expected. If X is a pseudo-finite subset of the ultrapower A^* with $X = -X$ and $\boldsymbol{\delta}(X + X) = \boldsymbol{\delta}(X)$, the conjecture predicts a *definable subgroup* $S \leq A^*$ with $\boldsymbol{\delta}(S + a \cap X) = \boldsymbol{\delta}(S) = \boldsymbol{\delta}(X)$ for some $a \in X$.[3] This uses only the abelian group structure and not properly the algebrogeometric structure; it lies already in a modular setting where the geometry is soft, and first-order interpretations cannot be expected to go far. All proofs of Freiman's theorem so far have had an analytic part, and this seems likely to be necessary and to continue in any future extensions.

At the other end, there is the Erdős–Szemerédi conjecture of [12]. It states for a pseudo-finite $X \subset \mathbb{Z}^*$ that $\boldsymbol{\delta}(X + XX) = 2\boldsymbol{\delta}(X)$. Note that the situation involves a range of different dimensions; for this reason it cannot be discerned by our $\boldsymbol{\mathcal{R}}$. It would already be interesting, and may be accessible to the methods described here, to prove it under the strong assumption that $\boldsymbol{\delta}$ or even δ_{fin} take a discrete set of values (analogous to assuming finite Morley rank in place of stability).

The next theorem is fundamental to our present concerns. It is the Szemerédi–Trotter theorem, in nonlinear versions due to many authors; of which the ultimate form has perhaps not yet been found. It concerns, in model-theoretic language, a pseudo-plane (P, L, I). Here P, L are pseudo-finite-definable or ∞-definable subsets of algebraic varieties $\underline{P}, \underline{L}$, while I can be taken to be the restriction to $P \times L$ of an algebraic variety $\underline{I}$. As far as I am aware, all proofs use real geometry; it would be extremely interesting to know why. At all events, the theorem as stated is only valid in *internal characteristic zero*, meaning that *almost each F_i has characteristic zero*, since it implies the nonexistence of finite subfields.

We cite two versions. The first, by Solymosi and Tao [36], is optimal numerically but has transversality assumptions that are not convenient at the level of generality we need. We will not state them explicitly.

Theorem 5.4 ([36] (With transversality and reality assumptions)) *Assume $\delta_{\text{fin}}(I) > \delta_{\text{fin}}(P), \delta_{\text{fin}}(L)$. Then for any real $\lambda > 2/3$,*

$$\delta_{\text{fin}}(I) \leq \lambda \delta_{\text{fin}}(P) + \frac{2}{3}\delta_{\text{fin}}(L).$$

In particular,

$$\boldsymbol{\delta}(I) \leq \frac{2}{3}\big(\boldsymbol{\delta}(P) + \boldsymbol{\delta}(L)\big).$$

As noted earlier, when L is a complete type, $I(l)$ takes a constant value e for $l \in L$. Let $a = \boldsymbol{\delta}(P), b = \boldsymbol{\delta}(L)$, so $b + e = \boldsymbol{\delta}(I)$. Using $e \geq 0$, we find that

$$\boldsymbol{\delta}(L) \leq 2\boldsymbol{\delta}(P).$$

This is intriguingly like the pseudo-modularity property that played a central role in proofs of Zilber's dichotomy for certain settings (see Hrushovski [15]). One could speculate that a failure of such pseudo-modularity (in pseudo-finite structures of any characteristic) leads to an interpretation of a field; even in characteristic zero this would give an interesting new proof. Some precise versions of this are formulated below.

Observe also that if $e = a/2$, we obtain $b \leq (1/2)a$; this is as predicted by the *fundamental rank inequality* of Cherlin, Harrington, and Lachlan [8], a strong form of modularity.

The second version, Elekes and Szabó [11, Theorem 9], is weaker numerically but with transversality assumptions that are closer to optimal. We are content to cite only Corollary 18, valid without the transversality assumptions.

Theorem 5.5 ([11, Corollary 18]) *Assume that F has internal characteristic zero. Let (P, L, I) be a pseudo-plane in (F, X), as above. If $\boldsymbol{\delta}(I(l)) = 1$ and $\boldsymbol{\delta}(P) = 2$, then $\boldsymbol{\delta}(L) \leq 1$.*

For later reference, we will also formulate a conjectural statement for small sets in positive characteristic.

Conjecture 5.6 *Assume that F has internal characteristic $p^* > 0$, so that it contains a nonstandard prime field $\mathbb{F}_{p^*}$. But assume $\boldsymbol{\delta}(\mathbb{F}_{p^*}) = \infty$. Then Theorem 5.5 holds as stated.*

The paper of Elekes and Szabó goes on to rediscover the group configuration, and the authors use it to prove that if $R \in \boldsymbol{\mathcal{R}}[3]$ has dominant, generically finite projections to any $\mathbb{A}^2 \leq \mathbb{A}^3$ via coordinate projections, then R is isogenous to addition on a one-dimensional abelian group. A certain higher-dimensional version is given ([11, Theorem 27]; here $R \subset P^3$ where P is allowed to have higher Zariski dimension, but the quantifier-free Morley dimension of $X \subset P$ in the structure $(X; R \cap X^3)$ is still assumed to be one).

Bukh and Tsimerman [5] conjecture a similar situation in positive characteristic, assuming $\boldsymbol{\delta}(P)$ sufficiently small compared to $\boldsymbol{\delta}(\mathbb{F}_{p^*})$. In their paper they prove this in certain cases where $\boldsymbol{\mathcal{R}}$ is assumed to include already a group structure (cf. Hrushovski and Pillay [17]; see also Schwartz, Solymosi, and de Zeeuw [35] and references there, where similar issues are tackled and it is δ_{fin} that dominates).

Now in positive characteristic, a linear Szemerédi–Trotter theorem is known (see Bourgain, Katz, and Tao [1]), but no nonlinear version is available. Bukh and Tsimerman reduce to the linear statement of [1] by polarization. A similar course was followed in Martin [26] for proving the trichotomy for the corresponding reducts of the theory ACF of algebraically closed fields—addition plus a polynomial. We will elaborate below on the analogy to the theory of reducts of ACFs and formulate a dichotomy, Conjecture 5.23, that would imply the main conjecture of [5, Section 9, p. 24].

Although the analogy with reducts is suggestive, reducts are assumed to be closed under first-order operations that are sensitive to individual points. We seek other model-theoretic frameworks that better reflect the statistical nature of Erdős geometry and whose operations lead to meaningful constructions in the combinatorial setting. We begin with a special case, built upon reducts at least at the quantifier-free level.

5.7 The Langian case Call X *Langian* if for any variety W, the family of irreducible subvarieties of W that fall into $\boldsymbol{\mathcal{R}}$ (or $\mathcal{R}$) has only finitely many maximal elements; denote the union of these subvarieties by $W^{\boldsymbol{\mathcal{R}}}$ (resp., $W^{\mathcal{R}}$). So $W^{\boldsymbol{\mathcal{R}}}$ is a Zariski-closed subset of W, and $W \in \boldsymbol{\mathcal{R}}$ iff $W = W^{\boldsymbol{\mathcal{R}}}$. Let $\mathbf{T}_X$ (resp., T_X) consist of all universal sentences in this language that hold true in F_0^{alg} (endowed with the natural $\boldsymbol{\mathcal{R}}$-structure), and in addition, the axioms $W = W^{\boldsymbol{\mathcal{R}}}$, for any variety W (i.e., $(\forall x)(x \in W \iff x \in W^{\boldsymbol{\mathcal{R}}})$). Then $\mathbf{T}_X$ is consistent, and indeed consistent with $(\exists x)(x \in U)$ for $U \in \boldsymbol{\mathcal{R}}$. To see this, consider first a single $U \in \boldsymbol{\mathcal{R}}$, and let $a = (a_1, \dots, a_n)$ be a generic point of U. Then $\{a_1, \dots, a_n\}$ is a model, since any variety in which $(a_{i_1}, \dots, a_{i_l})$ lies is the image of U under a projection, and so is in $\boldsymbol{\mathcal{R}}$. Next, if given finitely many varieties $U_1, \dots, U_k \in \boldsymbol{\mathcal{R}}$, note that a model of $\mathbf{T}_X \wedge (\exists x)(x \in \Pi_i U_i)$ is also a model of $(\exists u)(u \in U_i)$ for each i.

The universal theory $\mathbf{T}_X$ has the joint embedding property—this is clear from Lemma 5.2(2). Thus the class of existentially closed models of $\mathbf{T}_X$ fits into the setting of Shelah [33] and Pillay [28]; we refer to this class as $\widetilde{\mathbf{T}}_X$ and view it as the class of models of an ideal "theory."

The same construction works for $\mathcal{R}$ and δ_{fin}; in this case we call the universal theory T_X and call the class of existentially closed models $\widetilde{T}_X$.

Proposition 5.8 *$\widetilde{T}_X$ and $\widetilde{\mathbf{T}}_X$ are algebraically bounded, and the dimension of the Zariski closure coincides with the dimension induced by $\boldsymbol{\delta}$, up to a scalar.*

Proof Recall that the existential type of a tuple determines the orbit under the automorphism group, in an appropriately saturated model $\mathbb{U}$. Using compactness for existential types, algebraic boundedness reduces to showing that if $f : U \to V$ is a dominant map between varieties in $\boldsymbol{\mathcal{R}}$, and $U(X) \to V(X)$ has finite fibers, then $f : U \to V$ has finite fibers, at least over some Zariski-dense open V'. This is clear from Lemma 5.2(5). □

Remark 5.9 $\widetilde{T}$ is stable if $\mathcal{R}$ generates a disintegrated reduct of ACF; otherwise it is rarely simple. We have the familiar examples of ultraproducts of finite *subgroups* of a one-dimensional algebraic group, or finite fields. These and their isogenous variants are probably the *only* cases when $\widetilde{T}$ is simple. We give a heuristic argument, assuming Conjecture 5.10 (or restricting to internal characteristic zero). If $\widetilde{T}$ is not modular, then it interprets a field and we expect $\widetilde{T}$ to be isogenous to a theory of pseudo-finite fields. Assume that $\widetilde{T}$ is modular and $\mathcal{R}$ is not disintegrated. Over parameters one can find $R \in \boldsymbol{\mathcal{R}}$, $R \leq \mathbb{A}^3$, $\dim(R) = 2$, with generically finite projections to each $\mathbb{A}^2$. By modularity, $(\mathbb{A}^1, R)$ is isogenous to a one-dimensional algebraic group. Thus up to finite covers we can assume that X is an ultraproduct of approximate subgroups $X_i \leq A_i$. Now it seems that (A, X) has the strict order property, that is, a partial ordering with unbounded chains defined by an existential formula. In the fundamental case where X is an arithmetic progression $[m, n]b$, by translating we may assume $m = 0$, and then the partial ordering is $x \leq y \iff (\exists z \in X)(x + z = y)$. The general case requires more care, and a good use of Freiman's structure theorem.

Nevertheless, it seems that $\widetilde{T}$ carries a natural measure and has a form of almost everywhere 3-amalgamation for compact Lascar types; and the ideas of the simplicity theory of [28] should be of use.

As algebraic closure controls dimension, we can say that $\widetilde{T}$ (or $\widetilde{\mathbf{T}}$) is *modular* if the lattice of algebraically closed sets is modular, at least over an existentially closed substructure. We propose the following.

Conjecture 5.10 (Zilber's dichotomy for $\widetilde{T}$) *If $\widetilde{T}$ (resp., $\widetilde{\mathbf{T}}$) is not modular, then it interprets an infinite field.*

The notion of *interpretation* is a strong one here: there should exist affine algebraic varieties $V, E \leq V^2, M, P \leq V^3$, whose Zariski closures are all in $\mathcal{R}$ (resp., $\boldsymbol{\mathcal{R}}$), such that M, P respect E and $(V/E, M/E, P/E)$ is a ring, with $V(X)/E$ a subfield. The difference between these varieties and their Zariski closure should be definable in some way that we do not make precise here; a Boolean combination of existential conditions would do. See the analogous Conjecture 5.23 for comments.

We further expect that the field in the conjecture is naturally isomorphic to a subfield of the ambient field. In this form, the conjecture implies that if $\widetilde{T}$ or $\widetilde{\mathbf{T}}$ are not modular, then F (or F^n) has a pseudo-finite subfield, so that the F_i have finite subfields and must (for almost all i) have positive characteristic.

This last consequence is in fact true; it follows by first interpreting a pseudo-plane, then invoking the Szemerédi–Trotter theorem. We will see a completely analogous statement in the more generally applicable settings below, and so we omit the discussion here.

Remark 5.11 If X is Langian and in addition, $\boldsymbol{\mathcal{R}}$ (resp., $\mathcal{R}$) has amalgamation (and not only self-amalgamation and joint embedding), then $\widetilde{\mathbf{T}}$ (resp., $\widetilde{T}$) is a *reduct* of the theory ACF of algebraically closed fields. In this case the conjecture is a theorem of Eugenia Rabinovich [30].

We may expect that a "typical" X is Langian.[4] In particular, it may be that for any pseudo-finite X, and any finite subset S of varieties over $\mathbb{Q}$ that are in $\boldsymbol{\mathcal{R}}$, there exists a Langian X' with a corresponding $\boldsymbol{\mathcal{R}}' = \boldsymbol{\mathcal{R}}(X')$, with $S \subset \boldsymbol{\mathcal{R}}'$. But such things—if true—would be difficult to prove. For that matter, even for approximate subgroups of semiabelian varieties, Langianity is a nontrivial statement related to other Zilber conjectures. Despite this, the hypothesis seems to be a useful testing ground. At any rate we pass now to more general contexts.

Note that $\widetilde{T}$ and $\widetilde{\mathbf{T}}$ were formed directly out of $\mathcal{R}, \boldsymbol{\mathcal{R}}$. Without the Langian (or amalgamation) assumptions, I do not know how to make a useful structure directly out of $\boldsymbol{\mathcal{R}}$, and we return to using X. Now X may contain small subsets, say, existentially definable, that can be quite arbitrary; and quantifying over them could interpret set theory on logarithmically-sized sets, so we cannot expect any geometric properties at this level of generality. Our next framework does not replace the structure (F, X) with another but instead modifies the logic, so as to sense only top-dimensional behavior. Thus we may not ask if there exists a (single) x with a given property, only if there exist a substantial number (positive measure) of such elements. For fine dimension we use *probability logic*; one could try using *dimension logic* to deal with coarse dimension. Both take place within real-valued logic, which we proceed to review.

5.12 Real-valued logic (Yaacov and Usvayatsov [39]) A *real-valued structure* consists of a universe F and *real-valued relations*, meaning functions $g : F^n \to [0, m] \subset \mathbb{R}$. It would suffice to use functions into $[0, 1]$, but it is convenient to allow $[0, 2], \ldots$ as well. Boolean connectives are replaced by continuous functions $h : \prod_{i=1}^{k} [0, m_i] \to [0, m]$ (or a dense family of such functions). We begin with a family of *basic* real-valued relations and form new ones by composition with such "connectives." The basic relations are viewed as interpretations of basic relation symbols, and the relations generated by them are interpretations of corresponding *formulas* formed from these symbols. We will not be strict about the distinction between formulas and their interpretations.

A *discretely definable set* is a subset of F^n whose characteristic function, valued in the two-element set $\{0, 1\} \subset [0, 1]$, is a real-valued relation. For our purposes, we will assume that the diagonal on F is a discretely definable set.

Ultraproducts are formed by taking an ultraproduct of the universes as usual, while taking the limit along the ultrafilter of the real-valued relations (equivalently, taking the standard part of the ultraproduct of the relation values). Saturation is defined similarly. By a *discrete interpretation* of a structure M of a structure N we mean a discretely definable equivalence relation E on M^k and discretely definable relations R_i on M^{kn_i} respecting E, so that the interpreted structure is $N = (M^k/E, R_1/E, \ldots)$. We have a dimension function $\boldsymbol{\delta}$ induced on N.

5.13 Probability logic We use real-valued logic instead so as to retain compactness (and through it, amalgamation and some 3-amalgamation). As in Keisler [22, Section 3.4], we will use real-valued formulas and integral operators rather than relational quantifiers; they work well with real-valued logic and resolve the usual trouble with the interpretation of "there exists at least measure-β many x" at the critical boundary, that is, at the exact value β. We consider only finitary logic and use the terminology "(real-valued) formula" in place of Keisler's "term."

We thus define a *probability-logic structure* to be a real-valued structure as above, with an additional operator $\int$ taking a formula $t(x, y, \ldots)$ to a formula $\int t(x, y, \ldots)\, dx$ in variables $y, \ldots$. This operator is assumed to be linear ($\int(at + a't') = \int at + \int a't'$ for $a, a' \in \mathbb{R}_{\geq 0}$), nonnegative, and to have *norm* 1, that is, if t takes values in $[0, m]$ then so does $\int t$. We assume also that Fubini holds, inducing a unique integral operator $\int \cdots \int \cdots dx_1 \cdots dx_n$ in n variables for any n.

An n-type is by definition a real-valued function p on formulas $\varphi(x_1, \ldots, x_n)$, such that for some $(a_1, \ldots, a_n)$ in some structure A, $p(\varphi) = \varphi^A(a_1, \ldots, a_n)$. Let S_n denote the type space. It is a compact subspace of the space of functions on formulas into $\mathbb{R}$, with the Tychonoff topology.

The iterated integral extends to a norm 1 operator on the space of all continuous functions on S_n; by the Riesz representation theorem, it comes from a probability measure on S_n; these measures are compatible with the projections. Similarly, one obtains a measure on the space S_ω of types in infinitely many variables $x_1, x_2, \ldots$.

Two structures are *elementarily equivalent* if the integral operators (on formulas without parameters) act on formulas in the same way, and any two zero-place formulas have the same constant value for the two structures.

By an *elementary substructure* of M we mean an elementarily equivalent N whose universe is a subset of M, and such that for any formula t in k variables, $t^N = t^M | N^k$.

This logic gives us no way to treat, for instance, addition; within three-dimensional space, the graph of addition is zero-dimensional. In fact, typically if M is an enriched field, there exist elementary submodels whose universe is a set of algebraically independent elements. Thus as it stands, probability logic completely loses sight of the algebraic relations that are our main interest. We modify it by adding bounded existential quantifiers; equivalently, we may use relative probability quantifiers over any variety (see below).

5.14 Probability logic on varieties, at maximal fine dimension We proceed to describe the probabilistic formalism for a logic of $(F, X, \delta_{\mathrm{fin}})$ at the top dimensions. We use the associated measure at these dimensions. We will not use $\boldsymbol{\delta}$ here (see Problem 5.25).

Let $\alpha = \delta_{\mathrm{fin}}(X)$, and let $\mu = \mu_\alpha$ be the associated probability measure on definable subsets of X. We will use the real-valued probability logic described above, but *in addition to the integral operators, we use bounded quantifiers*. Namely, let $p(x, y) = p_d(y)x^d + \cdots + p_0(y)$ be a polynomial with integer coefficients, where x is a single variable and $y = y_1, \ldots, y_k$ is a tuple of variables. We introduce a quantifier, that is, an operator taking a formula $\varphi(x, y)$ to

$$\big[p_d(y) \neq 0\big]\Big(\sum_{p(x,y)=0} \varphi(x, y)\Big).$$

This is a new formula in the variables y, which vanishes if $p_d(y) = 0$ and otherwise sums $\varphi(x, y)$ over all roots of $p(x, y)$. A *structure* must correctly interpret these quantifiers. In particular, an elementary substructure is a relatively algebraically closed subfield.

As basic relations, we take the characteristic function 1_X of X, and the characteristic functions of all subvarieties of $\mathbb{A}^n$, over the prime field.

A partial type $P \subset X^n$ is *coherent* if the two natural notions of dimension coincide for P: $\delta_{\text{fin}}(P) = \dim \bar{P} \cdot \delta_{\text{fin}}(X)$, where $\bar{P}$ is the Zariski closure of P. Thus an irreducible closed subvariety $V \leq \mathbb{A}^n$ lies in $\mathcal{R}$ iff $V \cap X^n$ is coherent.

Remark 5.15

(1) By combining the bounded summation operators with the integral operators, we obtain integral operators on any variety. Conversely, the summation operators can be viewed as integral operators on (relatively) zero-dimensional varieties. The logic can thus be presented in terms of probability logic on varieties.

(2) A formula in *normal form* is formed out of the basic relations by real-valued connectives, bounded quantifier operators, and integral operators, *in this order*. We have quantifier elimination in the sense that the values of normal form formulas determine the complete type. This is related to Hoover's normal form theorem (see [22, Theorem 3.1.4]).

(3) Let S_ω be the space of types of infinite tuples of elements of X, over $\emptyset$. It is separable and carries a canonical measure as explained above. Let $(a_1, a_2, \ldots)$ be a random sequence; that is, $\text{tp}(a_1, a_2, \ldots)$ is a random element of S. Let N be the algebraic closure of the field generated by $a_1, a_2, \ldots$. Then N is an elementary submodel. Moreover, the measure of a formula in one variable with parameters in N equals the limit frequency of the solution set of the formula in $\{a_1, \ldots, a_n\}$. We refer to N as a *random elementary submodel*.

(4) In random elementary submodels, algebraic independence coincides with δ_{fin}-independence: if $a \notin F_0(b)^{\text{alg}}$ (where F_0 is the prime field and $b = (b_1, \ldots, b_m)$, $a, b_i \in N$), then any formula $\varphi(x, b)$ in $\text{tp}(a/b)$ has $\delta_{\text{fin}}(\varphi) = \alpha$ (where $\alpha = \delta_{\text{fin}}(X) = \delta_{\text{fin}}(x = x)$). This is an analogue of algebraic boundedness for $\widetilde{T}$ and leads to 3-replacement and almost everywhere 3-amalgamation. (We recall the statement of 3-replacement from [16], in this setting, over an elementary submodel. Assume that $\text{tp}(c/a, b)$ is wide, $\text{tp}(b/a)$ and $\text{tp}(b'/a)$ do not divide, and $\text{tp}(b) = \text{tp}(b')$. Then there exists c' with $\text{tp}(c'/a, b')$ wide, and $\text{tp}(c'b') = \text{tp}(cb), \text{tp}(c'a) = \text{tp}(ca)$.)

The property stated in (4) will be called *coherence*. To prove the coherence of a random elementary submodel $N = F_0(a_0, a_1, \ldots)^{\text{alg}}$, let $\delta = \delta_{\text{fin}}$. Note first that $\delta(a_1, \ldots, a_n) = n\alpha$. Say $a, b = (b_1, \ldots, b_m) \in \text{acl}(a_1, \ldots, a_n)$, with $a \notin \text{acl}(b)$. We have to show that $\delta(a/b) > 0$. Working over some of the a_i's, we may assume each $a_i \in \text{acl}(a, b)$. As $a \notin \text{acl}(b)$, $F_0(b)$ has transcendence degree at most $n - 1$ over F_0, and it follows that $\delta(b/F_0) \leq (n-1)\alpha$. If $\delta(a_i/b) = 0$ for each i, we obtain $\delta(a_1, \ldots, a_n) \leq (n-1)\alpha$, a contradiction. Thus $\delta(a_i/F_0(b)) > 0$; as $a_i \in \text{acl}(a, b)$, it follows that $\delta(a/b) > 0$. (This is superficially stronger than coherence but actually equivalent.)

We write $(F, X)^{\text{prob}}$ for (F, X) viewed as a measure-on-varieties structure.

5.16 Interpretations in probability logic on varieties We now define *probability logic on varieties*. The sorts are algebraic varieties over some fixed field F_0. Constructible sets are viewed as basic definable sets, and morphisms between algebraic varieties are basic functions of the structure. In addition, there are real-valued relations on sorts. Each definable set D comes with a dimension $\delta(D)$; if $\delta(D) \neq \infty$, also with a measure μ_D. Real-valued relations on D can be integrated correspondingly.

A natural notion of interpretation of such structures would thus work with morphisms between varieties (rather than set morphisms between sorts); namely, a surjective constructible map $f : U_1, \ldots, U_k \to V$ of varieties, such that the pullback of any definable relation D on V^n is a definable relation on U^n ($U = \prod_{i=1}^{k} U_i$), and the dimension and measure on D are appropriately normalized pushforwards of the corresponding ones on $f^{-1}(D)$. From the point of view of U, V is an imaginary sort obtained using a *constructible* equivalence relation E, rather than an arbitrary one. Using this definition has the effect of strengthening the dichotomy conjecture Conjecture 5.23 below, while making the recognition conjecture Conjecture 5.24 almost trivial. The first conjecture would remain very interesting for more generous notions of interpretation.

5.17 Definition of modularity Call a substructure Y of F *coherent* (over A) if any m-type realized in Y over A is coherent.[5] Equivalently, for any algebraically independent $a_1, \ldots, a_n \in Y$ over A (in the sense of algebra, i.e., of ACF) and any real-valued formula φ, if $\alpha < \varphi(a_1, \ldots, a_n) < \beta$, then $\delta\{x : \alpha \leq \varphi(x) \leq \beta\} = n\delta(X)$.

Thus a random elementary submodel is coherent.

Definition 5.18 $(F, X)^{\text{prob}}$ is *modular* if for any elementary submodel F_0, if Y is a coherent (small) subset of X over F_0, then algebraic closure over F_0 is modular; that is, for $a_1, \ldots, a_n \in Y$, if $a_1 \in F_0(a_2, \ldots, a_n)^{\text{alg}} \setminus F_0(a_2, \ldots, a_k)$, then $F_0(a_1, \ldots, a_k)^{\text{alg}} \cap F_0(a_{k+1}, \ldots, a_n)^{\text{alg}} \neq F_0^{\text{alg}}$.

By a $(2,3,2)$-*pseudo-plane* we mean interpretable ∞-definable sets P, L lying on algebraic varieties $\underline{P}$, $\underline{L}$, and a constructible set $\underline{I}$ such that for any two points of $\underline{P}$, $\underline{I}(a) \cap \underline{I}(a')$ is finite, and dually; and $\delta(P) = \dim(\underline{P}) = 2$, $\delta(L) = \dim(\underline{L}) = 2, \delta(I) = \dim(I) = 3$ (where δ denotes $\alpha^{-1}\delta_{\text{fin}}$). We can now restate Theorem 5.5.

Proposition 5.19 *Assume that F has internal characteristic zero. Then $(F, X)^{\text{prob}}$ has no $(2,3,2)$-pseudo-plane.*

Now comes a sequence of statements as in almost strongly minimal model theory.

Proposition 5.20 *Assume that $(F, X)^{\text{prob}}$ interprets no $(2,3,2)$-pseudo-plane. Then $(F, X)^{\text{prob}}$ is modular.*

Proof Over parameters, if $(F, X)^{\text{prob}}$ is not modular, then one can find a coherent sequence $a_1, \ldots, a_4$ demonstrating this, via an algebraic relation R among them. View R as a binary relation on $\mathbb{A}^2$, and rename the variables writing $R = R(x, y)$, $x = (x_1, x_2)$, $y = (y_1, y_2)$. View $R(a, y)$ as a family of curves on y-space. Let R' be a Zariski-open subset, such that $R'(a, y)$ is a disjoint union of irreducible curves. Let R'' be the irreducible component of $R' \times_{\mathbb{A}^2} R'$ containing the diagonal, so that if $(a, b, b') \in R''$, then $(a, b) \in R'$, $(a, b') \in R'$ and (b, b') lie in the same component of $R'(a)$. Let $\underline{P} = R/E$ where $(a, b), (a', b') \in E$ iff $a = a'$ and $(a, b, b') \in R''$. We have a relation I_1 on $\underline{P} \times \mathbb{A}^2$, pushforward of R''; and $I_1(a)$ is an irreducible curve, for $a \in \underline{P}$.

Similarly, we may replace the right-hand side $\mathbb{A}^2$ by a finite cover $\underline{Q}$, obtaining $I \subset \underline{P} \times \underline{Q}$ such that $R_2^t(b)$ is irreducible. With a little thought one sees geometrically that $R_2(a)$ remains irreducible. Model-theoretically, this is because the canonical base b' of the strong type of a over b satisfies $b' = Cb(\mathrm{stp}(a/b)) \subset \mathrm{acl}(b) \cap \mathrm{dcl}(a, b)$, so that $\mathrm{tp}(b'/a)$ is equidefinable with $\mathrm{tp}(b/a)$ and hence remains stationary.

Now for distinct $a, a' \in \underline{P}$ we have $\underline{P}(a) \cap \underline{P}(a')$ finite (this is the intersection of two distinct irreducible curves), and dually.

Let c be the image of $((a_1, a_2), (a_3, a_4))$ modulo E, noting $\mathrm{acl}(a_1, a_2) = \mathrm{acl}(c)$. Similarly define d, and let P, Q, I be the types of $c, d, (c, d)$, respectively. □

Proposition 5.21 *Assume that $(F, X)^{\mathrm{prob}}$ is modular and that $R \in \mathfrak{R}[3]$ has dominant, generically finite projections to any $\mathbb{A}^2 \leq \mathbb{A}^3$ via coordinate projections; then R is isogenous to addition on a one-dimensional abelian group.*

The proof is standard, using the ACF group configuration. Note in particular the conditional.

Corollary 5.22 *Assume Conjecture 5.6. Then the main conjecture of [5] is valid, at least for small A: for any degree d, for some m, if $A \subset \mathbb{F}_p$ and $|A| \leq p^{1/m}$, then for any polynomial f, $|\{(a, b) \in A^2 : f(a, b) \in A\}| \leq C_d |A|^{2-1/m}$.*[6]

Conjecture 5.23 *Assume that $(F, X)^{\mathrm{prob}}$ is not modular. Then it interprets a field k with $\boldsymbol{\delta}(k) > 0$. In fact $\boldsymbol{\delta}(k) = 1$.*

Conjecture 5.24 *The field k can be embedded in F, by a discretely definable function.*

This last *recognition* conjecture is straightforward for the strict notion of interpretation used above. We state it explicitly since in generalizations of the present scenario, or weakenings of the notion of interpretation, it may well become a significant step.

Problem 5.25 Investigate a "coarse-dimension logic on varieties" for $\boldsymbol{\delta}$ analogous to the above for δ_{fin}.

(1) Define a "coarse-dimension logic" analogous to probability logic but with $\boldsymbol{\delta}$ replacing μ_α. Possibly, replace the integral operators with *dimension operators* taking a formula $t(x, y)$ to a formula $⨏_x t(x, y, r)$ with $⨏_x t(x, y, r) = \delta\{x : t(x, b) \geq r\}$ (see Lemma 2.8(4)). Is there a normal form?
(2) Define a "coarse-dimension logic on varieties" by adding bounded quantifiers, or an infimum operator over algebraically bounded finite sets, as in Section 5.14. Call an ∞-definable set P *coherent* if $\boldsymbol{\delta}(P) = \dim \bar{P}$ (where $\boldsymbol{\delta}$ is normalized with $\boldsymbol{\delta}(X) = 1$ and $\bar{P}$ is the Zariski closure of P). Thus an irreducible closed subvariety $V \leq \mathbb{A}^n$ lies in $\mathfrak{R}$ iff $V \cap X^n$ is coherent. Show that "generic" sequences realizing coherent types are the universe of coherent elementary submodels.
(3) Define modularity to mean that algebraic closure is modular on coherent substructures over an elementary submodel. Investigate the analogue of Conjecture 5.23.

Remark 5.26 We could also consider a *maximal structure* on (F, X), where any ultraproduct of subsets of F_i^m is considered a definable subset of F^m. This is a first-order structure; we can still define a coherent formula and partial type, and thus repeat the definition of modularity and the statement of Conjecture 5.23. This would be weaker than the direct transposition of Conjecture 5.25(3) to coarse logic but would still imply Conjecture 5.6.

This means that the coarse-dimension dichotomy Conjecture 5.25(3) would remain interesting if the logic is augmented somewhat beyond Conjecture 5.25(1) (within the maximal structure).

On the other hand, the corresponding conjecture for the maximal structure itself would simply be equivalent to the combinatorial statement; it would not carry real model-theoretic content or constrain the possible proofs, without a restriction of the formulas to a more geometric level.

5.27 Three examples We now recall the promised examples. The first two are standard in the additive combinatorics literature; they exhibit a distinct phenomenon intermediate between modularity and fields, when several scales are allowed, but it is a surprisingly mild addition and is the only one I am aware of.

Example 5.28 Let G be a nilpotent algebraic group; for definiteness take the group G of strict upper diagonal (3×3)-complex matrices, viewed as an extension of $\mathbb{C}^2$ by $\mathbb{C}$. Let X_N consist of the matrices in G whose top right entry is an integer of absolute value at most N^2 and whose other entries are integers of absolute value at most N. Let X be an ultraproduct over N. Then X is an approximate subgroup; if M is the graph of multiplication on G^2, we have $\boldsymbol{\delta}(M) = 2\boldsymbol{\delta}(X)$, $\dim(M) = 2\dim(G)$; thus $M \in \boldsymbol{\mathcal{R}}$. Now G is not abelian by finite, and algebraic closure is not modular in the reduct generated by M. The presence of more than one scale is easily visible: the intersection of X with a generic coset bZ of Z has $\boldsymbol{\delta}(X \cap bZ) = \frac{1}{2}\dim(X)$. In fact, the family X_N is just a subfamily of the more natural $X_{N',N} = [-N', N']^2 \times [-N^2, N^2]$, with $N' \leq N$; the ultraproduct has two scales that are only accidentally equal in the original example.

The situation is reminiscent of locally modular superstable groups, which need not be 1-based. For example, they could be a central extension of an abelian group A by an abelian group Z; the image of a commutator map is a subset of Z, not a subgroup. Such groups have infinite U-rank, with regular types of ranks 1 and ω, analogous to $N' \ll N$ above.

A similar example, presented on a one-dimensional line and reminiscent of the many "pathological" examples in o-minimality when only a bounded field structure is present, is the one giving the optimality of the Szemerédi–Trotter theorem (see the Wikipedia entry, http://en.wikipedia.org/wiki/Szemer%C3%A9di-Trotter_theorem).

Example 5.29 Let $n \in \mathbb{N}^*$, the nonstandard natural numbers. Write $[a, b]$ for intervals in $\mathbb{N}^*$, $X = [1, 2n^2]$. Let M be the graph of multiplication, restricted to X^3. Then $\boldsymbol{\delta}(M) < 3\boldsymbol{\delta}(X)$, so $M \notin \boldsymbol{\mathcal{R}}$. With an additional scale concentrated on $Y = [1, n]$, we would see multiplication as a map $Y^2 \to X$, with behavior similar to that of Example 5.28.

Let $P = [1, n] \times [1, 2n^2]$, $L = [1, n] \times [1, n^2]$, $I = \{((x, y), (m, b)) \in P \times L; y = mx + b\}$. Then $\boldsymbol{\delta}(P) = \boldsymbol{\delta}(L) = 3\boldsymbol{\delta}([1, n])$, while (since $|I(c)| = |I(m, b)| = n$) $\boldsymbol{\delta}(I(c)) = \boldsymbol{\delta}([1, n]) = \frac{1}{3}\boldsymbol{\delta}(P)$ for $c \in L$. This contradicts the fundamental rank inequality, which would have $\boldsymbol{\delta}(L) \leq \frac{2}{3}\boldsymbol{\delta}(P)$ in this situation.

Example 5.30 Let A be a simple abelian variety of dimension greater than 1, and let $a \in A$ be a nontorsion element. Let Y be an ultraproduct of $Y_n = [-n, n]a$. Then Y intersects any proper subvariety of A in a finite set, with bounds uniform in a family of subvarieties. Consider the structure (Y, P) where $P = \{(y_1, y_2, y_3) \in Y^3 : y_1 + y_2 = y_3\}$. Let A' be a Zariski-open subset of A, and let $f : A' \to \mathbb{A}^1$ be a dominant morphism of algebraic varieties. Let $X = f(Y)$. Then f is finite-to-one on Y, so the induced structure on X is not algebraically bounded in the sense of van den Dries; $f(P)$ is a nonalgebraic relation inducing a nontrivial algebraic closure on X. Note that X is obtained by a simple projection and not by the probability quantifier of the appropriate dimension $(\dim(A) - 1)$.

6 Quasi-Finite Structures

Zilber classified the $\aleph_0$-categorical strongly minimal sets. More importantly, he found the characteristic structural property of these structures: modularity. In the years surrounding his proof, the inductive classification of all finite simple groups was visualized and then consolidated. A corollary is the classification of the *large* finite simple groups (CLFSG), that is, the statement that all but finitely many finite simple groups are either alternating, cyclic, or subgroups of simple algebraic groups consisting of points fixed by Frobenius maps (possibly twisted by group automorphisms). To date, no direct proof of the CLFSG has been found, or even conjecturally outlined. Moreover, the *structural content* of the statement—the model-theoretic meaning of this statement for finite structures—has not been clearly extracted. Zilber's theorem can be viewed as achieving both goals for a natural class of structures, with automorphism groups built out of alternating groups and full projective linear groups over a fixed finite field (see [8] for further discussion of the relation between the two theorems). It remains a highly significant role for model theory to find conceptual statements and proofs of wider parts of the CLFSG.

Say that a theory T is *quasi-finite* if for some function $\nu : \mathbb{N} \to \mathbb{N}$, any finite subset T_0 of T has a finite model with at most $\nu(k)$ k-types. It follows that T is $\aleph_0$-categorical and pseudo-finite. Moreover, (T, δ) is pseudo-finite for the appropriate pseudo-finite dimension δ. However, quasi-finiteness is much stronger and, using the CLFSG via [21], can be seen to imply coordinatization by classical geometries (see [9]). Conversely, all finite simple groups over a finite field of bounded size can occur within the automorphism groups of finite approximations ("envelopes") of a quasi-finite T.[7]

In [9], most of Zilber's theory was generalized to the class of quasi-finite structures. This time the classification of the finite simple groups was used, via work of [21]. It was shown however that the classification in [21] is equivalent to the conjunction of a list of model-theoretic properties of quasi-finite structures (see [9, p. 7, Theorem 7, properties LC1–LC9]). It is plausible that such properties can be deduced from quasi-finiteness by direct means, and it becomes a significant challenge to do so. For the key property LC4 of 3-amalgamation, initially proved inductively using the CLFSG, this was already achieved in the published version of the book [9].

We present the problem here again, in slightly updated form. In this section we write δ for δ_{fin}. For short we will say that (M, δ) is *pseudo-finite* if M is an ultraproduct of finite structures M_i, $\delta(\varphi)$ is the image of $\Pi_u|\varphi(M_i)|$ in the nonstandard reals $\Pi_u\mathbb{R}$ modulo the convex hull of $\mathbb{R}$, and $\delta(\varphi(x, a))$ depends only on $\text{tp}(a)$.

Theorem 6.1 *Let $T = Th(M)$, and assume that (M, δ) is pseudo-finite. Assume the five conditions below hold.*

(1) T is $\aleph_0$-categorical.
(2) Modularity for δ-dependence: If A, B are algebraically closed subsets of M^{eq}, and $a \in A$, then $\delta(a/B) = \delta(a/A \cap B)$.
(3) Every definable subset of an abelian group is a Boolean combination of cosets and an A_0-definable set, for some finite A_0.
(4) T does not interpret the theory T_{RND} of the bipartite random graph.
(5) For every definable $\mathbb{F}_p$-vector space V, any definable family of linear maps $V \to \mathbb{F}_p$ is contained in a definable vector space of such maps.

Then T is quasi-finite (and coordinatized by classical geometries).

We will also sketch in Section 6.8 how (5) can be replaced with an additional variant of (4).

Proof The conclusion is the same as that of [9, Theorem 7.5.6, p. 167]. The hypotheses LC1–LC9 of that theorem will be recalled and related to the assumptions of Theorem 6.1. LC1, 2, 5, 6, 8 are included among our assumptions: LC1 = (1), LC2 is pseudo-finiteness, LC5 is modularity, LC6 is (3), and LC8 is (4). Property LC4 is discussed immediately below; it follows from pseudo-finiteness and $\aleph_0$-categoricity of (M, δ). We will show in Lemma 6.15 that LC7 follows from the other conditions. Properties LC3 and LC9 are recalled further down; we will deduce LC3 in Section 6.13 and show that (5) is equivalent to LC9 in Lemma 6.9. □

Call a complete theory T (and any model of T) *almost quasi-finite* if it satisfies the properties listed in Theorem 6.1. In [9] it is proved (directly and model-theoretically) that almost quasi-finite structures are quasi-finite.

Problem 6.2 Find a direct, conceptual proof that quasi-finite structures are almost quasi-finite.

For condition (5), we will solve the problem below (see Lemma 6.12).

We now discuss the conditions of Theorem 6.1, prove the lemmas used in the proof of Theorem 6.1 above, and comment on Problem 6.2.

6.3 $\aleph_0$-categoricity and pseudo-finiteness of (M, δ) $\aleph_0$-categoricity, and the fact that δ is automorphism-invariant, follow immediately from the definition of quasi-finiteness.

6.4 3-amalgamation (LC4). Let $p_i(x_i), p_{ij}(x_i, x_j)$ $(i < j = 1, 2, 3)$ be types over an algebraically closed set A. Assume $p_i(x_i) \subset p_{ij}(x_i, x_j)$ and $\delta(p_{ij}) = \delta(p_i) + \delta(p_j)$. Then there exists $p(x_1, x_2, x_3)$ containing all p_{ij}, with $\delta(p) = \delta(p_1) + \delta(p_2) + \delta(p_3)$.

Property LC4 is already proved in [9, Proposition 8.4.2] for quasi-finite structures. The proof goes through with the invariance and pseudo-finiteness assumptions of Theorem 6.1. (Note that since δ is invariant and there are finitely many types, there are only finitely many possible values for δ.)

6.5 Modularity Condition (2) is clearly the key structural property. It takes a strong form here, of 1-basedness. Note that $\delta(a/B) > 0$ if $a \notin \mathrm{acl}(B)$ by definition of δ. A more local version of modularity would require B to be closed in the stronger sense that if $\delta(c/B) \ll \delta(a/B)$, then $c \in B$.

Could one more ambitiously prove a trichotomy result—either modularity, or an interpretable field of unbounded size—under weaker assumptions, including the analogue of finite Morley rank, $\delta(\mathrm{Def}) \cong \mathbb{Z}$ (see Macpherson and Steinhorn [25])?

Modularity of algebraic closure in definable duals of definable vector space is shown in Shelah [34] for 2-dependent theories, and in particular for theories with a simply exponential growth rate of local types over finite sets (see remark on p. 2 of [34]). In fact all quasi-finite theories are 2-dependent, but this is presently known only using the CLFSG.

6.6 Definable subgroups of abelian groups Condition (3), in the stable case, is an old theorem of Anand Pillay and myself. In that case, the exceptional family of A_0-definable sets does not intervene. A possible alternative goal would be to prove *quadraticity* rather than linearity. Discrete harmonic analysis is a tempting tool in this connection.

6.7 The random graph Item (4) concerns the noninterpretability of a specific theory. A bipartite graph is called *k-random* if for any two disjoint sets B, B' on one side of the graph, with $|B| \cup |B'| \leq k$, there exists a vertex c on the other side, adjacent to each element of B and to no element of B'. T_{RND} is the theory of bipartite graphs that are k-random for all k. One really needs to know that no such graphs are interpreted on a set of rank 1. If one were, by modularity it could be interpreted on a quotient of M^2; so assuming that M has few 4-types, the graph would have few 2-types. It seems plausible that for a given k, there is no finite, $2k$ (or even $k+4$) random graph with at most k 2-types.

See Cherlin [7] for the difficulty of proving non-pseudo-finiteness of a very simple specific theory, the model completion of the theory of triangle-free graphs, with present ideas. However non-quasi-finiteness can be proved rapidly. The unique 1-type p_0 over $\emptyset$ extends to an invariant global type p, as it happens a unique one: the type of an element x not connected in the graph to any existing element. It follows that $p|\,\mathrm{acl}^{\mathrm{eq}}(\emptyset)$ is invariant under automorphisms of $\mathrm{acl}^{\mathrm{eq}}(\emptyset)$; on the other hand all extensions of p_0 to $\mathrm{acl}^{\mathrm{eq}}(\emptyset)$ are conjugate by such automorphisms, so p_0 implies a complete type over $\mathrm{acl}^{\mathrm{eq}}(\emptyset)$.

This also shows that if b, b' are distinct and not connected by edge, then $\mathrm{acl}^{\mathrm{eq}}(b) \cap \mathrm{acl}^{\mathrm{eq}}(b') = \mathrm{acl}^{\mathrm{eq}}(\emptyset)$. (For instance, given a $\emptyset$-definable function g, find b'' such that $g(b') \neq g(b'')$; find $b''' \models p|\{b', b''\}$; then $g(b'') \neq g(b')$ or $g(b''') \neq g(b'')$; since $\mathrm{tp}(b, b') = \mathrm{tp}(b''', b') = \mathrm{tp}(b''', b'')$, it follows in either case that $g(b) \neq g(b')$.) Now if $a \neq b$, one can find b' with $\mathrm{tp}(b'/\,\mathrm{acl}^{\mathrm{eq}}(a)) = \mathrm{tp}(b/\,\mathrm{acl}^{\mathrm{eq}}(a))$ and such that (b, b') is not an edge. Hence $\mathrm{acl}^{\mathrm{eq}}(a) \cap \mathrm{acl}^{\mathrm{eq}}(b) \subset \mathrm{acl}^{\mathrm{eq}}(b) \cap \mathrm{acl}^{\mathrm{eq}}(b') = \mathrm{acl}^{\mathrm{eq}}(\emptyset)$. By 1-basedness (2), any two distinct points are independent. By 3-amalgamation, there exist a_1, a_2, a_3 with (a_i, a_j) an edge. This contradicts the no-triangles assumption.

6.8 Definable duals We first move from (5) to the formulation (LC9) given in [9].

Lemma 6.9 *Let T be almost quasi-finite. Let V be a vector space over $\mathbb{F}_p$, interpreted in T. Then there exists a definable group V^* and a definable bilinear map $\rho : V \times V^* \to \mathbb{F}_p$, such that any M-interpretable linear map $V \to \mathbb{F}_p$ has the form $x \mapsto \rho(x, b)$ for a unique $b \in V^*$. V^* is unique up to a unique definable isomorphism.*

Proof Uniqueness is clear; in fact V^* is clearly a piecewise-definable group, defined uniquely up to a unique definable isomorphism. Write $V^* = \cup Y_n$ with Y_n $\emptyset$-definable, and write $Y_0 \subset Y_1 \subset \cdots$. By condition (5), we can take Y_n to be a definable subspace of V^*. Let $Y_n \perp = \{v \in V : (\forall y \in Y_n)\rho(v, y) = 0\}$. So $Y_1 \supseteq Y_2 \supseteq \cdots$. By $\aleph_0$-categoricity, the kernels $Y_n \perp$ stabilize; $Y = Y_n = Y_{n+1} = \cdots$. Replacing V by V/Y, we may assume $Y = (0)$, and the pairing $\rho : V \times Y_m \to \mathbb{F}_p$ is nondegenerate for any $m \geq n$. In any finite model it follows that $|V| = |Y_m|$, so $Y_m = Y_{m+1}$. By pseudo-finiteness, in fact, $Y_m = Y_{m+1}$ in T. Thus $V^* = Y_n$ is definable. □

For condition (5), Problem 6.2 can be solved using the following observation made to me by David Kazhdan. For any set X, let ${}^X\mathbb{C}$ denote the space of functions $X \to \mathbb{C}$. For any representation U of G, U^G denotes the subspace of vectors $u \in U$ fixed by G.

Remark 6.10 Let V be a finite-dimensional representation over a finite field F of a group G. Let V^* be the dual vector space, with the natural G action; then G has as many orbits on V as on V^*. Indeed, the number of orbits of G on any finite set X equals $\dim_{\mathbb{C}}({}^X\mathbb{C})^G$. But we have a G-invariant isomorphism ${}^V\mathbb{C} \to {}^{V^*}\mathbb{C}$, namely, the Fourier transform. So $\dim_{\mathbb{C}}({}^V\mathbb{C})^G = \dim_{\mathbb{C}}({}^{V^*}\mathbb{C})^G$.

Remark 6.11 It follows directly from Remark 6.10 that if T is quasi-finite, and V is a definable vector space, and if T' is obtained from T by adding a sort for the dual of V, then T' is quasi-finite. This was already used in [9] at least for one-dimensional V, but the proof there used [21].

Lemma 6.12 *Assume that T is quasi-finite. Then Theorem 6.1(5) holds; indeed for every definable $\mathbb{F}_p$-vector space V, there exists a definable dual V^* containing all parametrically definable linear maps $V \to \mathbb{F}_p$.*

Proof Let V be a $\emptyset$-definable vector space over $\mathbb{F}_p$. View the space of all definable (with parameters) homomorphisms $V \to \mathbb{F}_p$ as a piecewise-definable vector space, as above. Let I be any definable subset of V. We have to show that the linear span of I is definable, that is, contained in a definable piece of V^*.

To see this, let $I_n = \{\sum_{i=1}^n \alpha_i a_i : \alpha_i \in \mathbb{F}_p, a_i \in I\}$. If $I_n = I_{n+1}$, it is easy to see that I_n is the span of I and is definable. Let $T_{2,k}$ be the subset of T asserting that V is a vector space, that $I \subset V^*$, and that $I_k \neq I_{k+1}$. Hence if $M \models T_{2,k}$, then $\mathrm{Aut}(M)$ has at least k orbits on V^* (subsets of $I_1, I_2 \setminus I_1, \ldots, I_k \setminus I_{k-1}$). Hence by Remark 6.10, for any model M of $T_{2,k}$, the number of orbits of $\mathrm{Aut}(M)$ on V is at least k. Thus T is not quasi-finite. □

6.13 Finite rank (LC3) By Lemma 2.2.3 of [9], $\mathrm{rk}(a/\emptyset) \geq 2$ iff for some $C = \mathrm{acl}(C)$, there exists $F = \mathrm{acl}(F)$ strictly intermediate between C and $\mathrm{acl}(C, a)$. Let $F' = \mathrm{acl}(a) \cap F$. It follows from (2) that F' is strictly intermediate between $\mathrm{acl}(\emptyset)$ and $\mathrm{acl}(a)$: we have $a \notin F'$ since $a \notin F$, but $F' \neq \mathrm{acl}(\emptyset)$ since

$\delta(a/F') = \delta(a/F) < \delta(a/C) \leq \delta(a/\emptyset)$. Thus $\mathrm{rk}(a/\emptyset) = 1$ if there exists no F' strictly between $\mathrm{acl}(\emptyset), \mathrm{acl}(a)$. Now by [9, Lemma 7.5.1], there are finitely many $F' = \mathrm{acl}(F')$ between $\mathrm{acl}(\emptyset)$ and $\mathrm{acl}(a)$. From Lemma 2.2.4 it follows that $\mathrm{rk}(a) = n$, where $\emptyset \subset F_1 \subset \cdots \subset F_n = \mathrm{acl}(a)$ is a maximal chain. So the rank is finite.

6.14 General position of large $\emptyset$-definable sets (LC7)

Lemma 6.15 *Let M be almost quasi-finite, and let A, B be $\emptyset$-definable groups admitting a definable bilinear map $\rho : A \times B \to \mathbb{F}_p$, nondegenerate on the right. Let $D \subset A$ be $\emptyset$-definable with $\delta(D) = \delta(A) > 0$; for $b \in B \setminus \mathrm{acl}(\emptyset), \alpha \in \mathbb{F}$, let $D(b, \alpha) = \{d \in D : \rho(d, b) = \alpha\}$. Then $\delta(D(b, \alpha)) = \delta(A)$.*

Proof We work over $\mathrm{acl}(\emptyset)$ (so $\mathrm{tp}(c) = \mathrm{tp}(c/\mathrm{acl}^{\mathrm{eq}}(\emptyset))$). Recall that $a \downarrow b$ means $\delta(a/b) = \delta(a)$. Let p be a type of D over $\mathrm{acl}(\emptyset)$ with $\delta(p) = \delta(D)$.

The theory of stabilizers goes through here. Let $S_0 = \{a - b : a \models p, b \models q, a \downarrow b\}$. Let S be the group generated by S_0. Then $\delta(S) = \delta(D) = \delta(A)$, so A/S is finite. By nondegeneracy of ρ on the right, it follows that $\mathrm{Ann}(S_0) = \mathrm{Ann}(S) = \{c : (\forall s \in S) \rho(s, c) = 0\}$ is finite. Now if $\rho(a, c) = 0$ for any $a \models p$ with $a \downarrow c\}$, then $c \in \mathrm{Ann}(S_0)$: to see this let $b \models S_0$; pick $a \models p$ with $a \downarrow c, b$ and such that $a + b \models p$; then $(a, c) = (a + b, c) = 0$ so $(a, c) = 0$. Hence the set of such c is finite.

For any nonalgebraic type q of B over $\mathrm{acl}(\emptyset)$, let $W(q) = \{\rho(a, b) : a \models p, b \models q, a \downarrow b\}$. If $q_i = \mathrm{tp}(c_i/\emptyset)$ for $i = 1, 2, 3$, $c_1 \downarrow c_2$, and $c_3 = c_1 + c_2$, and $\alpha_i \in W(q_i)$, then by 3-amalgamation there exists $a \models p$ with $a \downarrow c_1, c_2$ and $\rho(a, c_i) = \alpha_i$ $(i = 1, 2)$; so $\alpha_1 + \alpha_2 \in W(q_3)$. Similarly for $c_3 = c_1 - c_2$.

Let $q = \mathrm{tp}(b)$. If $\alpha_1, \alpha_2, \alpha_3 \in W(q)$, let $b_1, b_2 \models q$ be independent, $b' = b_1 - b_2$, $q' = \mathrm{tp}(b')$. Then by the above, $\alpha_1 - \alpha_2 \in W(q')$; and by a second application, $(\alpha_1 - \alpha_2) + \alpha_3 \in W(q)$. Thus $W(q)$ is a coset of a subgroup of $\mathbb{F}_p$. Since clearly $W(q) \neq \emptyset$, we have $W(q) = \mathbb{F}_p$. In particular $\alpha \in W(q)$. So there exists $a' \models p$, $a' \downarrow b$, $\rho(a', b) = \alpha$. This proves the lemma. □

In the statement of [9, Lemma 6.4.1], one should exclude the case where A^* is finite, since otherwise any element (including 0) is generic. With this correction, the lemma follows immediately (using induction on n) from Lemma 6.15.

6.16 Definable duals and the theory of the dual pseudo-basis Two specific theories have already been seen to play an important role in delineating the class of quasi-finite theories. The first was the theory T_{RND} of the random bipartite graph, just outside the class; the second was the analogous linear theory, of a bilinear map between two vector spaces or a polarity between two projective spaces, which *is* quasi-finite and determines much of the flavor of quasi-finite theories. We pause for a moment to consider a mix of the two: a relation between a projective space over a finite field, and a pure set. This again falls outside the quasi-finite class and is closely related to (5) of Theorem 6.1, or LC9.

First the vector version T_{VI}; it can be described as the model completion of the theory of a vector space over a finite field F with p elements, and an infinite, linearly independent subset I of the *dual space* to V.

Thus T_{VI} has two sorts, one of a vector space over a finite field F, the other a pure set, and we have in addition a function $b : V \times I \to F$, linear in the first variable, with the following axioms.

VI_n The natural images in V^* of the elements of I are linearly independent: for any $n \in \mathbb{N}$ and any distinct $a_1, \ldots, a_n \in I$ and $\alpha_1, \ldots, \alpha_n \in F$ there exists $v \in V$ with $b(v, a_i) = \alpha_i, i = 1, \ldots, n$.

VI^n The subspace of V^* generated by the image of I is dense: for any linearly independent $v_1, \ldots, v_n \in V$, and $\alpha_1, \ldots, \alpha_n \in F$, for some $a \in I$ we have $b(v, a_i) = \alpha_i, i = 1, \ldots, n$.

We will need a slight variant, the reduct T'_{VI} of T_{VI} where I, viewed as a subset of V^*, is replaced by the set F^*I. With $I' = F^*I$, the axioms VI^n remain the same, while the axioms VI_n must be restricted to *pairwise linearly independent* elements $a_1, \ldots, a_n$ and augmented by the axiom that $F^*I' = I'$.

T'_{VI} interprets, and is parametrically interpretable in, the theory T_{PI} induced on the sorts P, I from T_{VI}, where P is the projectivization of V.

We call I a *dual pseudo-basis*. In infinite models, it is only a linearly independent, dense subset of the dual space. In finite models of approximations to the theory, it spans the dual space but is only k-linearly independent for some k; it cannot be a basis since the formula defining the scalar multiples of the dual basis $(\exists! y \in I)\rho(x, y) \neq 0$ in fact has no solutions.

It is clear that T_{VI} is complete (with quantifier elimination), $\aleph_0$-categorical, modular, and satisfies Theorem 6.1(3).

Lemma 6.17 *T_{VI} is pseudo-finite, but not quasi-finite.*

Proof If T_{VI} were quasi-finite, by Lemma 6.12, V has a definable dual V^*. But then I embeds definably into V^*, and the distinct sets $I, I + I, I + I + I, \ldots$ show that the number of 1-types of V^* is unbounded, a contradiction.

To prove pseudo-finiteness, fix n. Let V be a vector space over F of large finite dimension, $|V| = N$. Let k be such that $Nn \log_2(p) < k < p^{(N/n-3)}$. Choose successively and at random elements $a_1, \ldots, a_k \in V^*$, subject only to the requirement that a_i not be in the span of any n elements a_j with $j < i$. This rules out only $p^n i^n \leq p^n k^n \leq p^{N-2n}$ elements, so at least $N/(2p^n)$ choices remain available within any subspace of codimension n. Now the linear independence requirement (VI_n) is met by construction, and (VI^n) will hold with high probability by the usual argument. □

If T does not interpret (with parameters) T_{PI} for any finite field F, it can be shown that Theorem 6.1(5) holds in dimension 1. We sketch the proof. First, T does not interpret (V, I), where the image of I is an infinite independent subset of V^*, and $\dim(V) = 1$. Otherwise, replacing V by $V/\{v \in V : (\forall x \in I)(\rho(v, x) = 0)\}$, we may assume that ρ is nondegenerate on the left, that is, $(\forall x \in I)(\rho(v, x) = 0)$ implies $v = 0$. Also we may assume $F^*I = I$. Thus the linear independence axioms VI_n hold in (V, I). The density axioms are then proved using 3-amalgamation. The rest of the proof no longer uses one-dimensionality; given a definable subset I of V^*, we know by the above that I is not linearly independent, and similarly that the image of I in projective space PV^* is not linearly independent. A linear relation among pairwise independent elements of I is then used, with the methods of [9], to find a definable group containing I.

It would be interesting to see an explicit list of theories whose noninterpretation is equivalent to condition (2).

Notes

1. The most appealing—and presently quite mysterious—example would be a generalization of the notion of a *compact group*. Model-theoretically, this is the same as a *hyper-definable group internal to a finite set*. *Connectedness* can be formulated as admitting no definable homomorphism to a finite group. This suggests a generalization to hyper-definable groups internal to a *small* set J, admitting no homomorphism to a J-internal group. Even when J is stably embedded (so very small indeed), it would be valuable to study such groups.

2. In [9] unrestricted CC-quantifiers were seen to lead to undecidability; more precise—and decidable—dimension quantifiers could be used to the same effect and were preferred. But at the level of generality of Sections 2–5 there is no harm in closing under CC.

3. The conjecture affirms this for any pseudo-finite dimension δ; the case of δ_{fin} is Freiman's theorem.

4. "Typical" is not meant in any statistical sense here. We could nevertheless look at a statistical sense: X_i a randomly chosen set of n_i points from $\mathbb{F}_{p_i}$. If n_i is sufficiently small compared to p_i, then Langianity is true but uninteresting, as $\mathcal{R}$ reduces to diagonals. If n_i is large compared to p_i, then $\mathcal{R}$ should consist of all absolutely irreducible F-varieties, and is still Langian. Some intermediate range may be more difficult to call.

5. We say *substructure* where usually one just says *subset* by abuse of language; this is to avoid any conflict between the definitions of coherence for substructures and for definable sets. The definitions are compatible via a certain duality, not by viewing a $\bigwedge$-definable set as a substructure.

6. Added in proof: Tao's recent preprint [38] proves a strong form of the conjecture for large A, that is, $|A| > p^{1-1/m}$ for appropriate m; his methods are highly suggestive from the point of view expounded in this paper.

7. Using the CLFSG one can show that if one restricts the hypothesis to values of k below 4, it still implies quasi-finiteness (see [9, Theorem 3]). We prefer to state the problem in a setting invariant under adding constants, but to fully recover the results of [9] one would have to work with $k = 4$.

References

[1] Bourgain, J., N. H. Katz, and T. Tao, "A sum-product estimate in finite fields, and applications," *Geometric and Functional Analysis*, vol. 14 (2004), pp. 27–57. Zbl 1145.11306. MR 2053599. DOI 10.1007/s00039-004-0451-1. 478

[2] Breuillard, E., Y. Cornulier, A. Lubotzky, and C. Meiri, "On conjugacy growth of linear groups," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 154 (2013), pp. 261–77. Zbl pre06144827. MR 3021813. DOI 10.1017/S030500411200059X. 465, 473

[3] Breuillard, E., B. Green, and T. Tao, "Approximate subgroups of linear groups," *Geometric and Functional Analysis*, vol. 21 (2011), pp. 774–819. MR 2827010. DOI 10.1007/s00039-011-0122-y. 470, 472

[4] Breuillard, E., B. Green, and T. Tao, "Linear approximate groups," *Electronic Research Announcements in Mathematical Sciences*, vol. 17 (2010), pp. 57–67. MR 2718104. DOI 10.3934/era.2010.17.57. 464

[5] Bukh, B., and J. Tsimerman, "Sum-product estimates for rational functions," *Proceedings of the London Mathematical Society (3)*, vol. 104 (2012), pp. 1–26. MR 2876962. DOI 10.1112/plms/pdr018. 478, 484

[6] Blossier, T., A. Martin-Pizarro, and F. O. Wagner, "On variants of CM-triviality," *Fundamenta Mathematicae*, vol. 219 (2012), pp. 253–62. Zbl pre06113309. MR 3001242. 465

[7] Cherlin, G., "Two problems on homogeneous structures, revisited," pp. 319–416 in *Model Theoretic Methods in Finite Combinatorics*, edited by M. Grohe and J. A. Makowsky, vol. 558 of *Contemporary Mathematics*, American Mathematical Society, Providence, 2011. 488

[8] Cherlin, G., L. Harrington, and A. H. Lachlan, "$\aleph_0$-categorical, $\aleph_0$-stable structures," *Annals of Pure and Applied Logic*, vol. 28 (1985), pp. 103–35. MR 0779159. DOI 10.1016/0168-0072(85)90023-5. 477, 486

[9] Cherlin, G., and E. Hrushovski, *Finite Structures with Few Types*, vol. 152 of *Annals of Mathematics Studies*, Princeton University Press, Princeton, 2003. MR 1961194. 463, 464, 465, 486, 487, 488, 489, 490, 491, 492

[10] Elekes, G., "SUMS versus PRODUCTS in number theory, algebra and Erdős geometry," pp. 241–90 in *Paul Erdős and his Mathematics, II (Budapest, 1999)*, vol. 11 of *Bolyai Society Mathematical Studies*, János Bolyai Mathematical Society, Budapest, 2002. MR 1954729. 475

[11] Elekes, G., and E. Szabó, "How to find groups? (and how to use them in Erdős geometry?)," *Combinatorica*, vol. 32 (2012), pp. 537–71. MR 3004808. DOI 10.1007/s00493-012-2505-6. 477, 478

[12] Erdős, P., and E. Szemerédi, "On sums and products of integers," pp. 213–18 in *Studies in Pure Mathematics*, Birkhäuser, Basel, 1983. MR 0820223. 476

[13] Goldbring, I., and H. Towsner, "An approximate logic for measures," preprint, arXiv:1106.2854v2 [math.LO]. 465

[14] Green, B., "Finite field models in additive combinatorics," pp. 1–27 in *Surveys in Combinatorics 2005*, vol. 327 of *London Mathematical Society Lecture Notes Series*, Cambridge University Press, Cambridge, 2005. MR 2187732. DOI 10.1017/CBO9780511734885.002. 476

[15] Hrushovski, E., "Unimodular minimal structures," *Journal of the London Mathematical Society (2)*, vol. 46 (1992), pp. 385–96. Zbl 0804.03023. MR 1190425. DOI 10.1112/jlms/s2-46.3.385. 477

[16] Hrushovski, E., "Stable group theory and approximate subgroups," *Journal of the American Mathematical Society*, vol. 25 (2012), pp. 189–243. MR 2833482. DOI 10.1090/S0894-0347-2011-00708-X. 464, 465, 468, 469, 470, 482

[17] Hrushovski, E. [U. Hrushovski], and A. Pillay, "Weakly normal groups," pp. 233–44 in *Logic Colloquium '85 (Orsay, 1985)*, vol. 122 of *Studies in Logic and the Foundations of Mathematics*, North-Holland, Amsterdam, 1987. MR 0895647. DOI 10.1016/S0049-237X(09)70556-7. 478

[18] Hrushovski, E., and A. Pillay, "Groups definable in local fields and pseudo-finite fields," *Israel Journal of Mathematics*, vol. 85 (1994), pp. 203–62. Zbl 0804.03024. MR 1264346. DOI 10.1007/BF02758643. 473

[19] Hrushovski, E., and F. Wagner, "Counting and dimensions," pp. 161–76 in *Model Theory*

with Applications to Algebra and Analysis, Vol. 2, vol. 350 of *London Mathematical Society Lecture Note Series*, Cambridge University Press, Cambridge, 2008. MR 2436141. DOI 10.1017/CBO9780511735219.005. 469

[20] Jordan, C., "Mémoire sur les équations différentielles linéaires à intégrale algébrique," *Crelle*, vol. 84 (1878), pp. 89–215. 471

[21] Kantor, W. M., M. W. Liebeck, and H. D. Macpherson, "$\aleph_0$-categorical structures smoothly approximated by finite substructures," *Proceedings of the London Mathematical Society (3)*, vol. 59 (1989), pp. 439–63. Zbl 0649.03018. MR 1014866. DOI 10.1112/plms/s3-59.3.439. 465, 486, 489

[22] Keisler, H. J., "Probability quantifiers," pp. 509–56 in *Model-Theoretic Logics*, edited by J. Barwise and S. Feferman, Springer, New York, 1985. MR 0819545. 480, 482

[23] Kim, B., and A. Pillay, "Simple theories," pp. 149–64 in *Joint AILA-KGS Model Theory Meeting (Florence, 1995)*, vol. 88 of *Annals of Pure and Applied Logic*, Elsevier, Amsterdam, 1997. MR 1600895. DOI 10.1016/S0168-0072(97)00019-5. 464

[24] Larsen, M. J., and R. Pink, "Finite subgroups of algebraic groups," *Journal of the American Mathematical Society*, vol. 24 (2011), 1105–58. Zbl 1241.20054. MR 2813339. DOI 10.1090/S0894-0347-2011-00695-4. 473

[25] Macpherson, D., and C. Steinhorn, "One-dimensional asymptotic classes of finite structures," *Transactions of the American Mathematical Society*, vol. 360 (2008), pp. 411–48. MR 2342010. DOI 10.1090/S0002-9947-07-04382-6. 488

[26] Martin, G. A., "Definability in reducts of algebraically closed fields," *Journal of Symbolic Logic*, vol. 53 (1988), pp. 188–99. Zbl 0653.03021. MR 0929384. DOI 10.2307/2274437. 478

[27] Nikolov, N., and L. Pyber, "Product decompositions of quasirandom groups and a Jordan type theorem," *Journal of the European Mathematical Society (JEMS)*, vol. 13 (2011), 1063–77. MR 2800484. DOI 10.4171/JEMS/275. 473

[28] Pillay, A., "Forking in the category of existentially closed structures," pp. 23–42 in *Connections Between Model Theory and Algebraic and Analytic Geometry*, vol. 6 of *Quaderni di Matematica*, Department of Mathematics Seconda University of Napoli, Caserta, 2000. Zbl 0986.03031. MR 1930681. 478, 479

[29] Pyber, L., and E. Szabó, "Growth in finite simple groups of Lie type of bounded rank," preprint, arXiv:1005.1858v2 [math.GR]. 464, 470, 472

[30] Rabinovich, E. D., *Definability of a Field in Sufficiently Rich Incidence Systems*, with an introduction by Wilfrid Hodges, vol. 14 of *QMW Maths Notes*, Queen Mary and Westfield College, School of Mathematical Sciences, London, 1993. MR 1213456. 480

[31] Ramsey, F. P., "On a problem of formal logic," *Proceedings of the London Mathematical Society*, vol. S2-30 (1930), pp. 264–85. Zbl 55.0032.04. MR 1576401. DOI 10.1112/plms/s2-30.1.264. 463

[32] Sanders, T., "On a nonabelian Balog–Szemerédi-type lemma," *Journal of the Australian Mathematical Society*, vol. 89 (2010), pp. 127–32. MR 2727067. DOI 10.1017/S1446788710000236. 464

[33] Shelah, S., "The lazy model-theoretician's guide to stability," pp. 241–308 in *Comptes Rendus de la Semaine d'Étude en Théorie des Modèles (Louvain-la-Neuve, 1975)*, vol. 18 of *Logique et Analyse (N.S.)*, Centre National Belge Recherches de Logique, Brussels, 1975. MR 0539969. 478

[34] Shelah, S., "Definable groups for dependent and 2-dependent theories," preprint, arXiv:math/0703045v4 [math.LO]. 488

[35] Schwartz, D., J. Solymosi, and F. de Zeeuw, "Extensions of a result of Elekes and Rónyai," preprint, arXiv:1206.2717v1 [math.CO]. 478

[36] Solymosi, J., and T. Tao, "An incidence theorem in higher dimensions," *Discrete and Computational Geometry*, vol. 48 (2012), pp. 255–80. MR 2946447. DOI 10.1007/s00454-012-9420-x. 477

[37] Tao, T., "The sum-product phenomenon in arbitrary rings," *Contributions to Discrete Mathematics*, vol. 4 (2009), pp. 59–82. MR 2592424. 473

[38] Tao, T., "Expanding polynomials over finite fields of large characteristic, and a regularity lemma for definable sets," preprint, arXiv:1211.2894v4 [math.CO]. 492

[39] Yaacov, I. B., and A. Usvyatsov, "Continuous first order logic and local stability," *Transactions of the American Mathematical Society*, vol. 362 (2010), pp. 5213–59, arXiv:0801.4303. MR 2657678. DOI 10.1090/S0002-9947-10-04837-3. 480

Acknowledgments

I would like to thank Zoé Chatzidakis and Pierre Simon, and especially the anonymous referees, for many precise and useful comments. Research was supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013) / ERC grant agreement no. 291111.

Department of Mathematics
The Hebrew University
Jerusalem
Israel
ehud@math.huji.ac.il