Model Theory

no. 1 vol. 3 2024

$M \nvDash T$

msp

Model Theory

msp.org/mt

See inside back cover or msp.org/mt for submission instructions.

Model Theory (ISSN 2832-904X electronic, 2832-9058 printed) at Mathematical Sciences Publishers, 798 Evans Hall #3840, c/o University of California, Berkeley, CA 94720-3840 is published continuously online.

MT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

Model Theory
Vol. 3, No. 1, 2024
https://doi.org/10.2140/mt.2024.3.1

Complete type amalgamation for nonstandard finite groups

Amador Martin-Pizarro and Daniel Palacín

We extend previous work on Hrushovski's stabilizer theorem and prove a measure-theoretic version of a well-known result of Pillay–Scanlon–Wagner on products of three types. This generalizes results of Gowers on products of three sets and yields model-theoretic proofs of existing asymptotic results for quasirandom groups. We also obtain a model-theoretic proof of Roth's theorem on the existence of arithmetic progressions of length 3 for subsets of positive density in suitable definably amenable groups, such as countable amenable abelian groups without involutions and ultraproducts of finite abelian groups of odd order.

Introduction

Szemerédi [1975] answered positively a question of Erdős and Turán by showing that every subset A of $\mathbb{N}$ of upper density

$$\limsup_{n\to\infty} \frac{|A\cap\{1,\dots,n\}|}{n} > 0$$

must contain an arithmetic progression of length k for every natural number k. For $k = 3$, the existence of arithmetic progressions of length 3 (in short 3-AP) was already proven by Roth [1953] in what is now called Roth's theorem on arithmetic progressions (not to be confused with Roth's theorem on diophantine approximation of algebraic integers). There has been (and still is) impressive work done on understanding Roth's and Szemerédi's theorems, explicitly computing lower bounds for the density as well as extending these results to more general settings. In the second direction, it is worth mentioning Green and Tao's result [2008] on the existence of arbitrarily long finite arithmetic progressions among the subset of prime numbers, which however has upper density 0.

Research supported by MTM2017-86777-P as well as by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project number 2100310201 and 2100310301, part of the ANR-DFG program GeoMod.
MSC2020: primary 03C45; secondary 11B30.
Keywords: model theory, additive combinatorics, arithmetic progressions, quasirandom groups.

In the noncommutative setting, proving single instances of Szemerédi's theorem, particularly Roth's theorem, becomes highly nontrivial. Note that the sequence (a, ab, ab^2) can be seen as a 3-AP, even for noncommutative groups. Gowers [2008, Question 6.5] asked whether the proportion of pairs (a, b) in $\mathrm{PSL}_2(q)$, for q a prime power, such that a, ab and ab^2 all lie in a fixed subset A of density δ, approximately equals δ^3. Gowers's question was positively answered by Tao [2013] and later extended to arbitrary nonabelian finite simple groups by Peluse [2018]. For arithmetic progressions (a, ab, ab^2, ab^3) of length 4 in $\mathrm{PSL}_2(q)$, a partial result was obtained in [Tao 2013], whenever the element b is diagonalizable over the finite field $\mathbb{F}_q$ (which happens half of the time).

A different generalization of Roth's theorem, present in [Sanders 2009; Henriot 2016], concerns the existence of a 3-AP in finite sets of small doubling in abelian groups. Recall that a finite set A of a group has doubling at most K if the productset $A \cdot A = \{ab\}_{a,b\in A}$ has cardinality $|A \cdot A| \leq K|A|$. More generally, a finite set has tripling at most K if $|A\cdot A\cdot A| \leq K|A|$. If A has tripling at most K, the comparable set $A \cup A^{-1} \cup \{\mathrm{id}_G\}$ (of size at most $2|A| + 1$) has tripling at most $(CK^C)^2$ with respect to some explicit absolute constant $C > 0$, so we may assume that A is symmetric and contains the neutral element. Archetypal sets of small doubling are approximate subgroups, that is, symmetric sets A such that $A \cdot A$ is covered by finitely many translates of A.

The model-theoretic study of approximate subgroups originated in Hrushovski's striking paper [2012], which contained the so-called stabilizer theorem, adapting techniques of stability theory to an abstract measure-theoretic setting. Hrushovski's work has led to several remarkable applications of model theory to additive combinatorics.

In classical geometric model theory, and more generally, in a group G definable in a simple theory, Hrushovski's stabilizer of a generic type over an elementary substructure M is the connected component G^{00}_M, that is, the smallest type-definable subgroup over M of bounded index (bounded with respect to the saturation of the ambient universal model). Generic types in G^{00}_M are called *principal types*. If the theory is stable, there is a unique principal type, but this need not be the case for simple theories. However, Pillay, Scanlon and Wagner [1998, Proposition 2.2] noticed that for every three principal types p, q and r in a simple theory over an elementary substructure M, there are independent realizations a of p and b of q over M such that $a \cdot b$ realizes r. The main ingredient in their proof is a clever application of 3-complete amalgamation (also known as *the independence theorem*) over the elementary substructure M.

For the purpose of the present work, we shall not define what a general complete amalgamation problem is, but a variation of it, restricting the problem to conditions given by products with respect to the underlying group law:

Question. Fix a natural number $n \geq 2$. For each nonempty subset F of $\{1, \ldots, n\}$, let p_F be a principal generic (that is, weakly random) type over the elementary substructure M. Can we find (under suitable conditions) an independent (weakly random) tuple $(a_1, \ldots, a_n)$ of G^n such that for all $\varnothing \neq F \subseteq \{1, \ldots, n\}$, the element a_F realizes p_F, where a_F stands for the product of all a_i, with i in F, written with the indices in increasing order?

The above formulation resonates with [Green and Tao 2008, Theorem 5.3] for quasirandom groups and agrees for $n = 2$ with the aforementioned result of Pillay, Scanlon and Wagner.

In this work, we give a (partial) positive solution for $n = 2$ (Theorem 3.10) to the above question for definable groups equipped with a definable Keisler measure satisfying Fubini (e.g., ultraproducts of groups equipped with the associated counting measure localized with respect to a distinguished finite set, as in Example 1.5). As a by-product, we obtain a measure-theoretic version of the result of Pillay, Scanlon and Wagner (Theorem 3.10):

Main Theorem. Given a pseudofinite subset X of small tripling in a sufficiently saturated group G and a countable elementary substructure M, for every weakly random type q and almost all pairs (p, r) of weakly random types over M concentrated in the subgroup $\langle X \rangle$ generated by X, there is a weakly random pair (a, b) over M in $p \times q$ with $a \cdot b$ realizing r, whenever $\mathrm{Cos}(p) \cdot \mathrm{Cos}(q) = \mathrm{Cos}(r)$, where $\mathrm{Cos}(p)$ is the coset of $\langle X \rangle_M^{00}$ determined by the type p.

The result of Pillay, Scanlon and Wagner holds for all such pairs (p, r) of generic types. Unfortunately, our techniques can only prove the analogous result outside a set of measure 0. Whilst we do not know how to obtain the result for all pairs (p, r) of weakly random types over M, our results however suffice to reprove model-theoretically some known results. Using a model-theoretic analog of Croot–Sisask's almost periodicity [Croot and Sisask 2010, Corollary 1.2] (Corollary 3.2), we easily deduce a nonquantitative version of Roth's theorem (Theorem 3.14) on 3-AP for finite subsets of small doubling in abelian groups with trivial 2-torsion, which resembles previous work of Sanders [2009, Theorem 7.1] and generalizes a result of Frankl, Graham and Rödl [1987, Theorem 1].

In Section 4, we reprove model-theoretically results valid for ultra-quasirandom groups, that is, asymptotic limits of quasirandom groups, already studied in [Bergelson and Tao 2014], and later in [Palacín 2020]. In particular, in Corollary 4.8 we give nonquantitative model-theoretic proofs of [Gowers 2008, Theorems 3.3 and 5.3]. In Section 5, we explore further this analogy to extend some of the results of Gowers to a local setting, without imposing that the group is an ultraproduct of quasirandom groups (see Corollaries 5.12 and 5.13).

We assume throughout the text a certain familiarity with basic notions in model theory. Sections 1, 2 and 3 contain the model-theoretic core of the paper, whilst Sections 4 and 5 contain applications to additive combinatorics.

1. Randomness and Fubini

Most of the material in this section can be found in [Halmos 1974; Hrushovski 2012; Massicot and Wagner 2015; Simon 2015].

We work inside a sufficiently saturated model $\mathbb{U}$ of a complete first-order theory (with infinite models) in a countable language $\mathcal{L}$, that is, the model $\mathbb{U}$ is saturated and strongly homogeneous with respect to some sufficiently large cardinal κ. All sets and tuples are taken inside $\mathbb{U}$.

A subset X of $\mathbb{U}^n$ is definable over the parameter set A if there exists a formula $\varphi(x_1, \dots, x_n, y_1, \dots, y_m)$ and a tuple $a = (a_1, \dots, a_m)$ in A such that an n-tuple b belongs to X if and only if $\varphi(b, a)$ holds in $\mathbb{U}$. As usual, we identify a definable subset of $\mathbb{U}$ with a formula defining it. Unless explicitly stated, when we use the word definable, we mean definably possibly with parameters. It follows that a subset X is definable over the parameter set A if and only if X is definable (over some set of parameters) and invariant under the action of the group of automorphisms $\mathrm{Aut}(\mathbb{U}/A)$ of $\mathbb{U}$ fixing A pointwise. The subset X of $\mathbb{U}$ is type-definable if it is the intersection of a bounded number of definable sets, where bounded means that its size is strictly smaller than the degree of saturation of $\mathbb{U}$.

For our applications we mainly consider the case where the language $\mathcal{L}$ contains the language of groups and the universe of our ambient model is a group. Nonetheless, our model-theoretic setting works as well for an arbitrary definable group, that is, a group whose underlying set and its group law are both definable.

Definition 1.1. A *definably amenable pair* (G, X) consists of an underlying definable group G together with the following data:

- a definable subset X of G;
- the (boolean) ring $\mathcal{R}$ of definable sets contained in the subgroup $\langle X \rangle$ generated by X, that is, the subcollection $\mathcal{R}$ is closed under finite unions and relative set-theoretic differences;
- a finitely additive measure μ on $\mathcal{R}$ invariant under both left and right translation with $\mu(X) = 1$. (Note that we require translation invariance under both actions).

Note that the subgroup $\langle X \rangle$ generated by the subset X need not be definable, but it is *locally definable*, for the subgroup $\langle X \rangle$ is a countable union of definable sets of the form

$$X^{\odot n} = \underbrace{X_1 \cdots X_1}_{n},$$

where X_1 is the definable set $X \cup X^{-1} \cup \{\mathrm{id}_G\}$. Furthermore, every definable subset Y of $\langle X \rangle$ is contained in some finite product $X^{\odot n}$, by compactness and saturation of the ambient model.

Remark 1.2. Model-theoretic compactness implies that the finitely additive measure μ satisfies Carathéodory's criterion, so there exists a unique σ-additive measure on the σ-algebra generated by $\mathcal{R}$. On the other hand, for every definable set Y of $\mathcal{R}$ over any set of parameters C, the measure μ extends to a regular Borel finite measure on the Stone space $S_Y(C)$ of complete types over C containing the C-definable set Y; see [Simon 2015, p. 99].

We denote the above extension of μ again by μ, though there will be (most likely) Borel sets of infinite measure, as noticed by Massicot and Wagner:

Fact 1.3 [Massicot and Wagner 2015, Remark 4]. *The subgroup $\langle X \rangle$ is definable if and only if $\mu(\langle X \rangle)$ is finite.*

Throughout the paper, we always assume that the language $\mathcal{L}$ is rich enough (see [Starchenko 2017, Definition 3.19]) to render the measure μ definable without parameters.

Definition 1.4. The measure μ of a definably amenable pair (G, X) is *definable without parameters* if for every $\mathcal{L}$-formula $\varphi(x, y)$, every natural number $n \geq 1$ and every $\epsilon > 0$, there is a partition of the $\mathcal{L}$-definable set

$$\{y \in \mathbb{U}^{|y|} \mid \varphi(\mathbb{U}, y) \subseteq X^{\odot n}\}$$

into $\mathcal{L}$-formulae $\rho_1(y), \dots, \rho_m(y)$ such that whenever a pair (b, b') in $\mathbb{U}^{|y|} \times \mathbb{U}^{|y|}$ realizes $\rho_i(y) \wedge \rho_i(y')$, then

$$|\mu(\varphi(x, b)) - \mu(\varphi(x, b'))| < \epsilon.$$

The above definition is a mere adaptation of [Starchenko 2017, Definition 3.19] to the locally definable context, by imposing that the restriction of μ to every definable subset $X^{\odot n}$ is definable in the sense of [Starchenko 2017, Definition 3.19]. In particular, a definable measure of a definably amenable pair (G, X) is *invariant*, that is, its value is invariant under the action of $\mathrm{Aut}(\mathbb{U})$. Notice that whenever the measure μ is definable, given a definable subset $\varphi(x, b)$ of measure r and a value $\epsilon > 0$, the tuple b lies in some definable subset which is contained in

$$\left\{y \in \mathbb{U}^{|y|} \mid r - \epsilon \leq \mu(\varphi(\mathbb{U}, y)) \leq r + \epsilon\right\}.$$

Assuming that μ is definable, its extension to the σ-algebra generated by the definable subsets of $\langle X \rangle$ is again invariant under left and right translations, as well as under automorphisms. Indeed, every automorphism τ of $\mathrm{Aut}(\mathbb{U})$ (likewise for left and right translations) gives rise to a measure μ^τ, such that $\mu^\tau(Y) = \mu(\tau(Y))$

for every measurable subset Y of $\langle X\rangle$. Since μ^τ agrees with μ on $\mathcal{R}$, we conclude that the σ-additive measure μ^τ is μ by the uniqueness of the extension. Thus, the measure of a Borel subset Y in the space of types containing a fixed definable set Z in $\mathcal{R}$ depends solely on the type of the parameters defining Y.

Example 1.5. Let $(G_n)_{n\in\mathbb{N}}$ be an infinite family of groups, each with a distinguished finite subset X_n. Expand the language of groups to a language $\mathcal{L}$ including a unary predicate and set M_n to be an $\mathcal{L}$-structure with universe G_n, equipped with its group operation, and interpret the predicate as X_n. Following [Hrushovski 2012, Section 2.6] we can further assume that $\mathcal{L}$ has predicates $Q_{r,\varphi}(y)$ for each r in $\mathbb{Q}^{\geq 0}$ and every formula $\varphi(x, y)$ in $\mathcal{L}$ such that $Q_{r,\varphi}(b)$ holds if and only if the set $\varphi(M_n, b)$ is finite with $|\varphi(M_n, b)| \leq r|X_n|$. Note that if the original language was countable, so is the extension $\mathcal{L}$.

Consider now the ultraproduct M of the $\mathcal{L}$-structures $(M_n)_{n\in\mathbb{N}}$ with respect to some nonprincipal ultrafilter $\mathcal{U}$. Denote by G and X the corresponding interpretations in a sufficiently saturated elementary extension $\mathbb{U}$ of M. For each $\mathcal{L}$-formula $\varphi(x, y)$ and every tuple b in $\mathbb{U}^{|y|}$ such that $\varphi(\mathbb{U}, b)$ is a subset of $\langle X\rangle$, define

$$\mu(\varphi(x, b)) = \inf\{r \in \mathbb{Q}^{\geq 0} \mid Q_{r,\varphi}(b) \text{ holds}\},$$

where we assign ∞ if $Q_{r,\varphi}(b)$ holds for no value r. This is easily seen to be a finitely additive definable measure on the ring $\mathcal{R}$ of definable subsets of $\langle X\rangle$ which is invariant under left and right translation. In particular, the pair (G, X) is definably amenable.

Throughout this paper we consider two main examples:

(a) The set X equals G itself, which happens whenever the subset $X_n = G_n$ for $\mathcal{U}$-almost all n in $\mathbb{N}$. The normalized counting measure μ defined above is a definable Keisler measure [Keisler 1987] on the pseudofinite group G. Note that in this case the ring of sets $\mathcal{R}$ coincides with the Boolean algebra of all definable subsets of G.

(b) For $\mathcal{U}$-almost all n, the set X_n has *small tripling*: there is a constant $K > 0$ such that $|X_nX_nX_n| \leq K|X_n|$. The noncommutative Plünnecke–Ruzsa inequality [Tao 2008, Lemma 3.4] yields that $|X_n^{\odot m}| \leq K^{O_m(1)}|X_n|$, so the measure $\mu(Y)$ is finite for every definable subset Y of $\langle X\rangle$, since Y is then contained in $X^{\odot m}$ for some m in $\mathbb{N}$. In particular, the corresponding σ-additive measure μ is again σ-finite.

Whilst each subset X_n in the example (b) must be finite, we do not impose that the groups G_n are finite. If the set X_n has tripling at most K, the set $X^{\odot 1} = X_n \cup X_n^{-1} \cup \{\mathrm{id}_G\}$ has size at most $2|X_n| + 1$ and tripling at most $(CK^C)^2$ with respect to some explicit absolute constant $C > 0$. Thus, taking ultraproducts, both structures (G, X) and $(G, X^{\odot 1})$ will have the same sets of positive measure (or density), though the values may differ. Hence, we may assume that in a definably amenable pair (G, X) the corresponding definable set X is symmetric and contains the neutral element of G.

The above example can be adapted to consider countable amenable groups.

Example 1.6. Recall that a countable group is *amenable* if it is equipped with a sequence $(F_n)_{n\in\mathbb{N}}$ of finite sets of increasing cardinalities (so $\lim_{n\to\infty}|F_n| = \infty$) such that for all g in G,

$$\lim_{n\to\infty} \frac{|F_n \cap g\cdot F_n|}{|F_n|} = 1.$$

Such a sequence of finite sets is called a left Følner sequence. The archetypal example of an amenable group is $\mathbb{Z}$ with left Følner sequence $F_n = \{-n, \dots, n\}$.

By [Namioka 1964, Corollary 5.3], if a group is amenable, then there is a distinguished left Følner sequence where each F_n is symmetric. In particular, the sequence $(F_n)_{n\in\mathbb{N}}$ is also a right Følner sequence:

$$\lim_{n\to\infty} \frac{|F_n \cap F_n\cdot g|}{|F_n|} = 1 \quad \text{for all } g \text{ in } G.$$

Notice also that a subsequence of a Følner sequence is again Følner and so is the sequence $(F_n \times F_n)_{n\in\mathbb{N}}$ in the group $G\times G$. Given an amenable group G with a distinguished Følner sequence $(F_n)_{n\in\mathbb{N}}$ consisting of symmetric sets as well as a nonprincipal ultrafilter $\mathcal{U}$ on $\mathbb{N}$, the ultralimit

$$\mu(Y) = \lim_{n\to\mathcal{U}} \frac{|Y\cap F_n|}{|F_n|}$$

induces a finitely additive measure on the Boolean algebra of subsets of G which is invariant under left and right translation.

Starting from a fixed countable language $\mathcal{L}$ expanding the language of groups, we can render the above measure definable, similarly as in Example 1.5. Hence, we can consider every countable amenable group G as a definably amenable pair, setting $X = G$.

Example 1.7. Every stable group G has finitely satisfiable generics (is fsg) and thus is equipped with a unique left and right translation invariant Keisler measure which is generically stable (see [Hrushovski et al. 2013; Simon 2015, Example 8.34]).

Similarly, a compact semialgebraic Lie group $G(\mathbb{R})$, or more generally a definably compact group G definable in an o-minimal expansion of a real closed field, is again fsg. If the group is the $\mathbb{R}$-rational points of a compact semialgebraic Lie group, this measure coincides with the normalized Haar measure.

Hence, we can consider in these two previous cases (stable and o-minimal compact) the group G as a definably amenable pair, setting $X = G$.

If a group G is definable, so is every finite cartesian product. Moreover, the construction in Examples 1.5 and 1.6 can also be carried out for a finite cartesian product to produce for every $n \geq 1$ in $\mathbb{N}$ a definably amenable pair (G^n, X^n), where $\langle X^n \rangle = \langle X \rangle^n$, equipped with a definable σ-finite measure μ_n. Thus, the following assumption is satisfied by our Examples 1.5, 1.6 and 1.7.

Assumption 1. For every $n \geq 1$, the pair (G^n, X^n) is definably amenable for a definable σ-finite measure μ_n in a compatible fashion: the measure μ_{n+m} extends the corresponding product measure $\mu_n \times \mu_m$.

The definability condition in Definition 1.4 implies that the function

$$F^{\varphi}_{\mu_n,C} : S_m(C) \to \mathbb{R}, \quad \operatorname{tp}(b/C) \mapsto \mu_n(\varphi(x, b)),$$

is well-defined and continuous for every $\mathcal{L}_C$-formula $\varphi(x, y)$ with $|x| = n$ and $|y| = m$ such that $\varphi(x, y)$ defines a subset of $\langle X \rangle^{n+m}$. Therefore, for such $\mathcal{L}_C$-formulae $\varphi(x, y)$, consider the $\mathcal{L}_C$-definable subset $Y = \{y \in \langle X \rangle^m \mid \exists x\, \varphi(x, y)\}$ and the corresponding clopen subset $[Y]$ of $S_m(C)$. Thus, we can consider the following measure ν on $\langle X \rangle^{n+m}$:

$$\nu(\varphi(x, y)) = \int_{q \in [Y]} F^{\varphi}_{\mu_n,C}(q)\, d\mu_m = \int_{y \in Y} \mu_n(\varphi(x, y))\, d\mu_m.$$

By an abuse of notation, we write $\int_{\langle X \rangle^m} \mu_n(\varphi(x, y))\, d\mu_m$ for $\int_Y \mu_n(\varphi(x, y))\, d\mu_m$.

For the pseudofinite measures described in Example 1.5, the above integral equals the ultralimit

$$\lim_{k \to \mathcal{U}} \frac{1}{|X_k|^m} \sum_{y \in \langle X_k \rangle^m} \frac{|\varphi(x, y)|}{|X_k|^n},$$

so ν equals μ_{n+m} and consequently Fubini–Tonelli holds; see (the proof of) [Bergelson and Tao 2014, Theorem 19]. The same holds whenever the measure is given by densities with respect to a Følner sequence in an amenable group, as in Example 1.6. For arbitrary definably amenable pairs, whilst the measure ν extends the product measure $\mu_n \times \mu_m$, it need not be a priori μ_{n+m} [Starchenko 2017, Remark 3.28]. Keisler [1987, Theorem 6.15] exhibited a Fubini–Tonelli type theorem for general Keisler measures under certain conditions. These conditions hold for the unique generically stable translation invariant measure of an fsg group (see Example 1.7). We will impose a further restriction on the definably amenable pairs we consider, taking Examples 1.5, 1.6 and 1.7 as a guideline.

Assumption 2. For every definably amenable pair (G, X) and its corresponding compatible system of definable measures $(\mu_n)_{n\in\mathbb{N}}$ on the Cartesian powers of $\langle X \rangle$, the Fubini condition holds: whenever a definable subset of $\langle X \rangle^{n+m}$ is given by an

$\mathcal{L}_C$-formula $\varphi(x, y)$ with $|x| = n$ and $|y| = m$, we have

$$\mu_{n+m}(\varphi(x, y)) = \int_{\langle X\rangle^m} \mu_n(\varphi(x, y))\, d\mu_m = \int_{\langle X\rangle^n} \mu_m(\varphi(x, y))\, d\mu_n.$$

(Note that the above integrals do not run over the locally definable sets $\langle X\rangle^m$ and $\langle X\rangle^n$, but rather over definable subsets, for $\varphi(x, y)$ is itself definable).

Whilst this assumption is stated for definable sets, it extends to certain Borel sets, whenever the language $\mathcal{L}_C$ is countable.

Remark 1.8. Assume that $\mathcal{L}_C$ is countable and fix a natural number $k \geq 1$. Following [Conant et al. 2023, Definition 2.6], for every Borel subset Z of $S_{n+m}(C)$ of types $q(x, y)$ with $|x| = n$ and $|y| = m$, set

$$Z(x, b) = \{p \in S_n(\mathbb{U}) \mid \text{tp}(a, b/C) \text{ belongs to } Z \text{ for some } a \text{ realizing } p_{\restriction C,b}\}.$$

Note that $Z(x, b)$ only depends on $\text{tp}(b/C)$ by [Conant et al. 2023, Lemma 2.7]. If Z is contained in the clopen set determined by the $\mathcal{L}_C$-definable set $(X^{\odot k})^{n+m}$, we define analogously as before a function

$$F^Z_{\mu_n,C} : S_m(C) \to \mathbb{R}, \quad \text{tp}(b/C) \mapsto \mu_n(Z(x, b)).$$

This function is Borel, and thus measurable, by the definability of the measure as well as the monotone convergence theorem, for it agrees with $F^{\varphi}_{\mu_n,C}$ whenever Z is the clopen $[\varphi]$. Furthermore, the identity

$$\mu_{n+m}(Z(x, y)) = \int_{\langle X\rangle^m} \mu_n(Z(x, y))\, d\mu_m = \int_{\langle X\rangle^n} \mu_m(Z(x, y))\, d\mu_n,$$

holds by a straightforward application as in [Bergelson and Tao 2014, Theorem 20] of the monotone class theorem, using the fact that $\mu(X^{\odot k})$ is finite. In particular, the above identity of integrals holds for every Borel set of finite measure by regularity.

Remark 1.9. The examples listed in Examples 1.5, 1.6 and 1.7 satisfy both Assumptions 1 and 2.

Henceforth, the language is countable and all definably amenable pairs satisfy Assumptions 1 and 2.

Adopting some terminology from additive combinatorics, we use the word *density* for the value of the measure of a subset in $\mathcal{R}$ of a definably amenable pair (G, X). A (partial) type is said to be *weakly random* if it contains a definable subset in $\mathcal{R}$ of positive density but no definable subset in $\mathcal{R}$ of density 0. Note that every weakly random partial type $\Sigma(x)$ over a parameter set A implies a definable set $X^{\odot k}$ in $\mathcal{R}$ for some k in $\mathbb{N}$ and thus it can be completed to a weakly random complete type over any arbitrary set B containing A, since the collection of formulae

$$\Sigma(x) \cup \{X^{\odot k} \setminus Z \mid Z \text{ in } \mathcal{R} \text{ is } B\text{-definable of density } 0\}$$

is finitely consistent. Thus, weakly random types exist (yet the partial type $x = x$ is not weakly random whenever $G \neq \langle X \rangle$). As usual, we say that an element b of G is weakly random over A if $\mathrm{tp}(b/A)$ is.

Weakly random elements satisfy a weak notion of transitivity.

Lemma 1.10. *Let b be weakly random over a set of parameters C and a be weakly random over C, b. The pair (a, b) is weakly random over C.*

Proof. We need to show that every C-definable subset Z of $\langle X \rangle^{n+m}$ containing the pair (a, b) has positive density with respect to the product measure μ_{n+m}, where $n = |a|$ and $m = |b|$. Since a is weakly random over C, b, the fiber Z_b of Z over b has measure $\mu_n(Z_b) = 2r$ for some real number $0 < r$. Hence b belongs to a C-definable subset Y of

$$\{y \in \mathbb{U}^m \mid r \leq \mu_n(Z_y) \leq 3r\},$$

by the definability of the measure. In particular, the measure $\mu_m(Y)$ is strictly positive. Thus,

$$\mu_{n+m}(Z) = \int_{\langle X \rangle^m} \mu_n(Z_y)\, d\mu_m \geq \int_Y \mu_n(Z_y)\, d\mu_m \geq \mu_m(Y) r > 0,$$

as desired. □

Note that the tuple b above may not be weakly random over C, a. To remedy the failure of symmetry in the notion of randomness, we introduce *random* types, which will play a fundamental role in Section 3. Though random types already appear in [Hrushovski 2013, Subsection 2.23] (see also [Hrushovski 2012, Subsection 2.20]), we take the opportunity here to recall Hrushovski's definition of *ω-randomness*. All the ideas here until the end of this section are due to Hrushovski and we are merely writing down some of the details for the sake of the presentation.

Fix some countable elementary substructure M and some Y in $\mathcal{R}$ definable over M (so $Y \subseteq (X^{\odot k})$ for some k in $\mathbb{N}$). As in Remark 1.2, we denote by $S_{Y^m}(M)$ the compact subset of the space of types over M containing the M-definable subset Y^m.

Definition 1.11. Denote by $\mathcal{B}_M^Y$ the smallest Boolean algebra of subsets of $S_{Y^m}(M)$, as m varies, containing all clopen subsets of $S_{Y^m}(M)$ and closed under the following operations:

- The preimage of a set $W \subseteq S_{Y^m}(M)$ in $\mathcal{B}_M^Y$ under the natural continuous map $S_{Y^n}(M) \to S_{Y^m}(M)$ given by the restriction to a choice of m coordinates belongs again to $\mathcal{B}_M^Y$.
- If $Z \subseteq S_{Y^{n+m}}(M)$ belongs to $\mathcal{B}_M^Y$, then so does

$$(F_{\mu_n, M}^Z)^{-1}(\{0\}) = \{\mathrm{tp}(b/M) \in S_{Y^m}(M) \mid \mu_n(Z(x, b)) = 0\},$$

  with $Z(x, b)$ as in Remark 1.8.

Note that each element of $\mathcal{B}_M^Y$ is a Borel subset of the appropriate space of types by Remark 1.8. Furthermore, it is countable since it can be inductively built from the Boolean algebras of clopen subsets of the $S_{Y^m}(M)$'s by adding in the next step all Borel sets of the form $(F^Z_{\mu_n,M})^{-1}(\{0\})$ and closing under Boolean operations. The collection $\mathcal{B}_M^Y$ contains new sets which are neither open nor closed.

Definition 1.12. Let Y in $\mathcal{R}$ be definable over the countable elementary substructure M. An n-tuple a of elements in Y is *random* over $M \cup B$, where B is some countable subset of parameters, if $\mu_n(Z(x, b)) > 0$ for every finite subtuple b in B and every Borel subset Z in $\mathcal{B}_M^Y$ with $\mathrm{tp}(a, b/M)$ in Z.

For $B = \varnothing$, we simply say that the tuple is *random* over M.

Remark 1.13. Since $\mathcal{B}_M^Y$ contains all clopen sets given by M-definable subsets, it is easy to see that a tuple random over $M \cup B$ is weakly random over $M \cup B$, which justifies our choice of terminology (instead of using the term *wide type* from [Hrushovski 2012]).

Randomness is preserved under the group law: if a is an element of $\langle X\rangle$ random over $M \cup B$, then so are a^{-1} and $b \cdot a$ for every element b in $B \cap \langle X\rangle$.

Furthermore, note that randomness is a property of the type: if a and a' have the same type over $M \cup B$, then a is random over $M \cup B$ if and only if a' is.

Remark 1.14. Since $\mathcal{B}_M^Y$ is countable, the σ-additivity of the measure yields that every measurable subset of $S_{Y^m}(M \cup B)$, with B countable, of positive density contains a random element over $M \cup B$. In particular, every weakly random definable subset of Y^m contains random elements over M, B.

Randomness is a symmetric notion.

Lemma 1.15 [Hrushovski 2013, Exercise 2.25]. *Let Y in $\mathcal{R}$ be definable over the countable elementary substructure M. A finite tuple (a, b) of elements in Y is random over M if and only if b is random over M and a is random over M, b.*

Proof. Assume first that (a, b) is random over M. Clearly so is b by Fubini and Remark 1.8. Thus we need only prove that a is random over M, b. Suppose for a contradiction that $\mu_{|a|}(Z(x, b)) = 0$ for some $Z \subseteq S_{Y^{|a|+|b|}}(M)$ of $\mathcal{B}_M^Y$ containing $\mathrm{tp}(a, b)$. The type of the pair (a, b) belongs to

$$\begin{aligned}\widetilde{Z} &= Z \cap \pi^{-1}\big((F^Z_{\mu_{|a|},M})^{-1}(\{0\})\big)\\ &= Z \cap \big\{\mathrm{tp}(c, d/M) \in S_{Y^{|a|+|b|}}(M) \mid \mu_{|a|}(Z(x, d)) = 0\big\},\end{aligned}$$

where $\pi = \pi_{|a|+|b|,|b|}$ is the corresponding restriction map. Now, the set $\widetilde{Z}$ belongs to $\mathcal{B}_M^Y$ and contains (a, b), so it cannot have density 0. However, Remark 1.8 yields

$$0 < \mu_{|a|+|b|}(\widetilde{Z}) = \int_{Y^{|b|}} \mu_{|a|}(\widetilde{Z}(x, d))\, d\mu_{|b|} \le \int_{Y^{|b|}} \mu_{|a|}(Z(x, d))\, d\mu_{|b|} = 0,$$

which gives the desired contradiction.

Assume now that b is random over M and a is random over M, b. Suppose for a contradiction that $\mathrm{tp}(a, b/M)$ lies in some Borel $Z(x, y)$ of $\mathcal{B}_M^Y$ with $\mu_{|a|+|b|}(Z)=0$. By Remark 1.8,

$$0=\mu_{|a|+|b|}(Z)=\int_{Y^{|b|}} \mu_{|a|}(Z(x, d))\, d\mu_{|b|},$$

so $\mu_{|a|}(Z(x, d))=0$ for $\mu_{|b|}$-almost all types $\mathrm{tp}(d/M)$ in $S_{Y^{|b|}}(M)$. Hence, the set $(F^Z_{\mu_{|a|},M})^{-1}(\{0\})$ has measure $\mu_{|b|}(Y^{|b|})$. Since a is random over M, b, we have that $\mu_{|a|}(Z(x, b))>0$, so $\mathrm{tp}(b/M)$ belongs to the complement of $(F^Z_{\mu_{|a|},M})^{-1}(\{0\})$, which belongs to $\mathcal{B}_M^Y$ and has $\mu_{|b|}$-measure 0. We conclude that the element b is not random over M, which gives the desired contradiction. □

Symmetry of randomness will allow us in Sections 3 and 4 to transfer ideas arisen from the study of definable groups in simple theories to the pseudofinite context as well as to definably compact groups definable in o-minimal expansions of real closed fields. Whilst weak randomness is not symmetric, a weak form of symmetry holds (as pointed out by the anonymous referee, to whom we would like to express our sincere gratitude again).

Lemma 1.16 (the referee's lemma). *Let Y in $\mathcal{R}$ be a subset of positive density definable over the countable elementary substructure M. Given two finite tuples a and b of elements in Y with a weakly random over M and b random over M, a, then a is weakly random over M, b.*

Proof. Assuming otherwise, there is an M-definable set Z containing (a, b) such that the fiber Z_b has $\mu_{|a|}$-measure 0. Definability of the measure Definition 1.4 yields that the set

$$W=(F^Z_{\mu_{|a|},M})^{-1}(\{0\})=\{\mathrm{tp}(d/M)\in S_{Y^{|b|}}(M) \mid \mu_{|a|}(Z_d)=0\}$$

is closed and thus it can be written as a countable intersection $W=\bigcap_{m\in\mathbb{N}} W_m$ of M-definable sets with $W_{m+1}\subseteq W_m$. Now, the closed set $[Z(x, y)]\cap W(y)$ belongs to $\mathcal{B}_M^Y$ and contains $\mathrm{tp}(a, b/M)$, so $\mu_{|b|}([Z(a, y)]\cap W(y))>0$, since b is random over M, a.

Claim. *There exists some M-definable subset V containing a such that*

$$\mu_{|b|}([Z(a', y)]\cap W(y))>0$$

for all a' in V.

Note that V has positive density, for $\mathrm{tp}(a/M)$ is weakly random.

Proof of Claim. Assume for a contradiction that this is not the case. Since both the language and M are countable, we may list all M-definable subsets containing a as $\{V_n\}_{n\in\mathbb{N}}$ with $V_{n+1}\subseteq V_n$. Therefore, for every n in $\mathbb{N}$ there is some a_n in V_n with

$\mu_{|b|}([Z(a_n, y)] \cap W(y)) < 1/(n+1)$. As W is a countable intersection of the W_m's, there is some m_n in $\mathbb{N}$ such that

$$\mu_{|b|}(Z(a_n, y) \cap W_{m_n}(y)) < \frac{1}{n+1}.$$

Notice that we may construct the sequence such that $m_{n+1} > m_n$. Set

$$\theta_<(Z, W_{m_n}) = \left\{ x \in Y^{|a|} \mid \mu_{|b|}(Z(x, y) \cap W_{m_n}(y)) < \frac{1}{n+1} \right\}$$

and define $\theta_\leq(Z, W_{m_n})$ analogously. By definability of the measure, there is some M-definable subset $\theta(Z, W_{m_n})$ such that

$$\theta_<(Z, W_{m_n}) \subseteq \theta(Z, W_{m_n}) \subseteq \theta_\leq(Z, W_{m_n}).$$

In particular, we have that $\theta(Z, W_{m_{n+1}}) \subseteq \theta(Z, W_{m_n})$ for $m_{n+1} > m_n$. Now, the collection of $\mathcal{L}_M$-formulae $\{V_n(x) \wedge \theta(Z, W_{m_n})(x)\}_{n \in \mathbb{N}}$ cannot be consistent, for it would yield the existence of a tuple a' realizing $\operatorname{tp}(a/M)$ with

$$\mu_{|b|}([Z(a', y)] \cap W(y)) \leq \mu_{|b|}(Z(a', y) \cap W_{m_n}(y)) \leq \frac{1}{n+1}$$

for every n in $\mathbb{N}$, so $\mu_{|b|}([Z(a', y)] \cap W(y)) = 0 < \mu_{|b|}([Z(a, y)] \cap W(y))$, which is a contradiction. By compactness, there exists some ℓ in $\mathbb{N}$ such that no realization of V_ℓ satisfies some $\theta(Z, W_{m_j})$ with $j \leq \ell$. However, the element a_ℓ belongs to $V_\ell \cap \theta_<(Z, W_{m_\ell})$, so a_ℓ lies in every $\theta(Z, W_{m_j})$ with $j \leq \ell$, which gives the desired contradiction. $\square_{\text{Claim}}$

Consider now the closed set $W' = [V(x)] \cap [Z(x, y)] \cap W$. The Fubini condition (Remark 1.8) yields that

$$\begin{aligned} 0 &\overset{\text{Claim}}{<} \int_{\operatorname{tp}(c/M) \in [V]} \mu_{|b|}([Z(c, y)] \cap W) \, d\mu_{|b|} \\ &= \mu(W') = \int_{\operatorname{tp}(d/M) \in W} \mu_{|a|}(V(x) \cap Z(x, d)) \, d\mu_{|a|} \\ &\leq \int_{\operatorname{tp}(d/M) \in W} \mu_{|a|}(Z(x, d)) \, d\mu_{|a|} = 0. \end{aligned}$$

We deduce from the above contradiction that a lies in no definable set Z_b over M, b of density 0, so a is weakly random over M, b, as desired. $\square$

2. Forking and measures

As in Section 1, we work inside a sufficiently saturated structure and a definably amenable pair (G, X) in a fixed countable language $\mathcal{L}$ satisfying Assumptions 1 and 2, though the classical notions of forking and stability do not require the presence of a group nor of a measure.

Recall that a definable set $\varphi(x, a)$ *divides* over a subset C of parameters if there exists an indiscernible sequence $(a_i)_{i\in\mathbb{N}}$ over C with $a_0 = a$ such that the intersection $\bigcap_i \varphi(x, a_i)$ is empty. Archetypal examples of dividing formulae are of the form $x = a$ for some element a not algebraic over C. Since dividing formulae need not be closed under finite disjunctions, witnessed for example by a circular order, we say that a formula $\psi(x)$ *forks* over C if it belongs to the ideal generated by the formulae dividing over C, that is, if ψ implies a finite disjunction of formulae, each dividing over C. A type *divides* (resp. *forks*) over C if it contains an instance which does.

Remark 2.1. Since the measure is invariant under automorphisms and σ-finite, no definable subset of $\langle X\rangle$ of positive density divides, and thus no weakly random type forks over the empty set; see [Hrushovski 2012, Lemma 2.9 and Example 2.12].

Nonforking need not define a tame notion of independence. For example it need not be symmetric, yet it behaves extremely well with respect to certain invariant relations, called stable.

Definition 2.2. An A-invariant relation $R(x, y)$ is *stable* if there exists no A-indiscernible sequence $(a_i, b_i)_{i\in\mathbb{N}}$ such that $R(a_i, b_j)$ holds if and only if $i < j$.

A straightforward Ramsey argument yields that the collection of invariant stable relations is closed under Boolean combinations. Furthermore, an A-invariant relation is stable if there is no A-indiscernible sequence as in the definition of length some fixed infinite ordinal.

The following remark will be very useful in the following sections.

Remark 2.3 [Hrushovski 2012, Lemma 2.3]. Suppose that the type $\mathrm{tp}(a/M, b)$ does not fork over the elementary substructure M and that the M-invariant relation $R(x, y)$ is stable. Then the following are equivalent:

(a) The relation $R(a, b)$ holds.

(b) The relation $R(a', b)$ holds, whenever $a' \equiv_M a$ and $\mathrm{tp}(a'/Mb)$ does not fork.

(c) The relation $R(a', b)$ holds, whenever $a' \equiv_M a$ and $\mathrm{tp}(b/Ma')$ does not fork.

(d) The relation $R(a', b')$ holds, whenever $a' \equiv_M a$ and $b' \equiv_M a$ such that $\mathrm{tp}(a'/M, b')$ or $\mathrm{tp}(b'/M, a')$ does not fork.

A clever use of the Krein–Milman theorem on the locally compact Hausdorff topological real vector space of all σ-additive probability measures allowed Hrushovski to prove the following striking result (the case $\alpha = 0$ is an easy consequence of the inclusion-exclusion principle).

Fact 2.4 [Hrushovski 2012, Lemma 2.10 and Proposition 2.25]. *Given a real number α and $\mathcal{L}_M$-formulae $\varphi(x, z)$ and $\psi(y, z)$ with parameters over an elementary* *substructure M, the M-invariant relation on the definably amenable pair (G, X)*

$$R^{\alpha}_{\varphi,\psi}(a,b) \quad \Leftrightarrow \quad \mu_{|z|}(\varphi(a,z) \wedge \psi(b,z)) = \alpha$$

is stable. In particular, for any partial types $\Phi(x,z)$ and $\Psi(y,z)$ over M, the relation

$$Q_{\Phi,\Psi}(a,b) \quad \Leftrightarrow \quad \Phi(a,z) \wedge \Psi(b,z) \text{ is weakly random}$$

is stable.

Strictly speaking, Hrushovski's result in its original version is stated for arbitrary Keisler measures (in any theory). To deduce the statement above it suffices to normalize the measure $\mu_{|z|}$ by $\mu_{|z|}((X^{|z|})^{\odot k})$ for some natural number k such that $(X^{|z|})^{\odot k}$ contains the corresponding instances of $\varphi(x,z)$ and $\psi(y,z)$.

We finish this section with a summarized version of Hrushovski's stabilizer theorem tailored to the context of definably amenable pairs. Before stating it, we first need to introduce some notation.

Definition 2.5. Let X be a definable subset of a definable group G and let M be an elementary substructure. We denote by $\langle X\rangle^{00}_M$ the intersection of all subgroups of $\langle X\rangle$ type-definable over M and of bounded index.

If a subgroup of bounded index type-definable over M exists, the subgroup $\langle X\rangle^{00}_M$ is again type-definable over M and has bounded index; see [Hrushovski 2012, Lemmata 3.2 and 3.3]. Furthermore, it is also normal in $\langle X\rangle$ [Hrushovski 2012, Lemma 3.4].

Fact 2.6 [Hrushovski 2012, Theorem 3.5; Montenegro et al. 2020, Theorem 2.12]. *Let (G, X) be a definably amenable pair and let M be an elementary substructure. The subgroup $\langle X\rangle^{00}_M$ exists and equals*

$$\langle X\rangle^{00}_M = (p \cdot p^{-1})^2$$

for any weakly random type p over M, where we identify a type with its realizations in the ambient structure $\mathbb{U}$. Furthermore, the set $pp^{-1}p$ is a coset of $\langle X\rangle^{00}_M$. For every element a in $\langle X\rangle^{00}_M$ weakly random over M, the partial type $p \cap a \cdot p$ is weakly random. In particular, every weakly random element in $\langle X\rangle^{00}_M$ over M lies in $p \cdot p^{-1}$.

If the definably amenable pair we consider happens to be as in the first case of Example 1.5 or a stable group as in Example 1.7, our notation coincides with the classical notation G^{00}_M.

3. On 3-amalgamation and solutions of $xy = z$

As in Section 1, we fix a definably amenable pair (G, X) satisfying Assumptions 1 and 2. Throughout this section, we work over some fixed elementary substructure M.

We denote by $\mathrm{supp}_M(\mu)$ the *support* of μ, that is, the collection of all weakly random types over M contained in $\langle X\rangle$.

Note that each coset of the subgroup $\langle X\rangle_M^{00}$ of Definition 2.5 is type-definable over M and hence M-invariant, though it need not have a representative in M. Thus, every type p over M contained in $\langle X\rangle$ must determine a coset of $\langle X\rangle_M^{00}$. We denote by $\mathrm{Cos}(p)$ the coset of $\langle X\rangle_M^{00}$ of $\langle X\rangle$ containing some (and hence every) realization of p. The following result resonates with [Terry and Wolf 2019, Corollary 1] and [Conant et al. 2020, Theorem 1.3] beyond the definable context.

Proposition 3.1. *Consider an M-invariant subset S of $\langle X\rangle$ such that the relation $u \cdot v \in S$ is stable, as in Definition 2.2. The set S must be, up to M-definable sets of measure* 0, *a union of cosets of $\langle X\rangle_M^{00}$, that is, if an element g in $\langle X\rangle$ belongs to S with $q = \mathrm{tp}(g/M)$ in $\mathrm{supp}_M(\mu)$, then every element h in $\mathrm{Cos}(q)$ weakly random over M belongs to S as well.*

Our proof is mostly an adaptation of [Pillay et al. 1998, Proposition 2.2]. Whilst the authors used the independence theorem from simple theories, we use the stability of the M-invariant relation S instead.

Proof. Assume that the element g as above belongs to the stable M-invariant relation S. Let h be in $\mathrm{Cos}(\mathrm{tp}(g/M))$ weakly random over M and choose a realization b of $\mathrm{tp}(h/M)$ weakly random over M, g. Now, the elements g and b both lie in the same coset of $\langle X\rangle_M^{00}$, so the difference $b \cdot g^{-1}$ lies in $\langle X\rangle_M^{00}$ and is weakly random over M, g. Since weakly random types do not fork, the type $\mathrm{tp}(b \cdot g^{-1}/M, g)$ does not fork over M.

Fact 2.6 yields that the partial type $\mathrm{tp}(g/M) \cap (b \cdot g^{-1}) \cdot \mathrm{tp}(g/M)$ is weakly random. Choose therefore some element g_1 realizing $\mathrm{tp}(g/M)$ weakly random over M, g, b such that $b \cdot g^{-1} \cdot g_1 \equiv_M g$. By invariance of S, we have that $b \cdot g^{-1} \cdot g_1$ belongs to S as well.

Summarizing, the M-invariant relation $\bar{S} = \{(u, v) \in \langle X\rangle \times \langle X\rangle \mid u \cdot v \in S\}$ holds for the pair $(b \cdot g^{-1}, g_1)$ with $\mathrm{tp}(g_1/M, b \cdot g^{-1})$ weakly random and hence nonforking over M. Since the above relation is stable, for any pair (w, z) such that

$$w \equiv_M b \cdot g^{-1}, \quad z \equiv_M g_1 \quad \text{and} \quad \mathrm{tp}(w/M, z) \text{ nonforking over } M,$$

the relation $\bar{S}$ must also hold. Setting now $w = b \cdot g^{-1}$ and $z = g$, we deduce that $b = b \cdot g^{-1} \cdot g$ belongs to S. As the element h realizes $\mathrm{tp}(b/M)$, we conclude by M-invariance that h belongs to S, as desired. □

Given now two M-definable subsets A and B, the relation

$$R^{\alpha}_{A,B}(u, v) \Leftrightarrow \text{“}\mu(uA \cap vB) = \alpha\text{”}$$

is stable by Fact 2.4. So, setting $S = \{g \in \langle X\rangle \mid \mu(A \cap gB) = \alpha\}$, Proposition 3.1 yields immediately the following result, which we personally think resonates with Croot–Sisask's almost-periodicity [Croot and Sisask 2010, Corollary 1.2].

Corollary 3.2. *Given two M-definable subsets A and B, the values $\mu(A \cap gB)$ and $\mu(A \cap hB)$ agree for any two weakly random elements g and h over M within the same coset of $\langle X\rangle^{00}_M$.* □

Given now two types p_1 and p_2 over M and an element g of $\langle X\rangle$ such that the partial type $p_1 \cdot g \cap p_2$ is consistent, it follows that the type $\mathrm{tp}(g/M)$ determines the coset $\mathrm{Cos}(p_1)^{-1} \cdot \mathrm{Cos}(p_2)$, so $\mathrm{Cos}(p_1) \cdot \mathrm{Cos}(\mathrm{tp}(g/M)) = \mathrm{Cos}(p_2)$. The following result can be seen as a sort of converse. Notice that

$$S = \{g \in \langle X\rangle \mid p_1 \cdot g \cap p_2 \text{ is weakly random over } M\}$$

is M-invariant and $u \cdot v \in S$ is stable, by Fact 2.4.

Corollary 3.3. *Let p, q and r be three coset-compatible types in $\mathrm{supp}_M(\mu)$, that is,*

$$\mathrm{Cos}(p) \cdot \mathrm{Cos}(q) = \mathrm{Cos}(r).$$

If $p \cdot g \cap r$ is weakly random for some element g in $\langle X\rangle$ with $\mathrm{tp}(g/M)$ in $\mathrm{supp}_M(\mu)$, then so is $p \cdot h \cap r$ for every weakly random element h whose type over M is concentrated in $\mathrm{Cos}(q)$. □

The above result was first observed for principal generic types in a simple theory in [Pillay et al. 1998, Proposition 2.2] and later generalized to nonprincipal types in [Martin-Pizarro and Pillay 2004, Lemma 2.3]. For weakly random types with respect to a pseudofinite Keisler measure, a preliminary version was obtained in [Palacín 2020, Proposition 3.2] for ultra-quasirandom groups.

For the rest of this section, we assume that M is countable. Fix some k in $\mathbb{N}$ and consider $Y = (X^{\odot k})$. The value k should be chosen large enough to ensure that all the products and inverses of elements in the subsequent statements still belong to Y. By an abuse of language, we use the word *random* to mean a random type with respect to the corresponding class $\mathcal{B}^Y_M$ as in Definitions 1.11 and 1.12.

Remark 3.4. It follows immediately from Remark 1.14 that the Borel set of random types over M is dense in the compact Hausdorff space of weakly randoms concentrated on Y, that is, the space $[Y] \cap \mathrm{supp}_M(\mu)$, where $[Y]$ is the clopen set given by the M-definable set Y. We denote by $\mathrm{R}(\mathcal{B}^Y_M)$ the collection of random types over M concentrated on Y.

Lemma 3.5. *Given M-definable subsets A and B of Y of positive density, there exists some random element g over M with $\mu(Ag \cap B) > 0$.*

Proof. By Remark 1.14, let c be random in B over M and choose now g^{-1} in $c^{-1}A$ random over M, c. The element g is also random over M, c. By symmetry of randomness, the pair (c, g) is random over M, so c is random over M, g. Clearly the element c lies in $Ag \cap B$, so the set $Ag \cap B$ has positive density, as desired. □

Remark 3.6. Notice that the above results yields the existence of an element h random over M such that $hA \cap B$, and thus $A \cap h^{-1}B$, has positive density. Indeed, apply the statement to the definable subsets B^{-1} and A^{-1}.

For any two fixed types p and r in $\mathrm{supp}_M(\mu)$, the statement

"$p \cdot y \cap r$ is weakly random and y is weakly random"

as a property of y is finitely consistent. Indeed, given finitely many M-definable subsets $A_1, \ldots, A_n$ in p and $B_1, \ldots, B_n$ in r, the M-definable subsets $A = \bigcap_{1\leq i\leq n} A_i$ and $B = \bigcap_{1\leq i\leq n} B_i$ lie in p and r, respectively, so they both have positive density. By Lemma 3.5, there exists a random element g in $\langle X\rangle$ over M with $A_i g \cap B_j$ of positive density for all $1 \leq i, j \leq n$.

However, the condition "$p \cdot y \cap r$ is weakly random" is a G_δ-condition on y, namely

$$\bigcap_{\substack{A\in p\\ B\in r}} \{y \in A^{-1}B \mid \mu(A \cdot y \cap B) > 0\}.$$

Thus, we cannot use compactness to deduce from the above argument that we fulfill the conditions of Corollary 3.3 for all weakly random types p, q and r. We are grateful to Angus Matthews for pointing out a mistake in a previous version of this paper.

To circumvent the aforementioned issue, we use the so-called *disintegration theorem*, which allows us to fulfill the conditions of Corollary 3.3 for almost all pairs of types p and r. Whilst there are plenty of excellent references on this subject worth being named, we just refer to [Bogachev 2007; Simmons 2012].

Remark 3.7. Given n in $\mathbb{N}$ consider a set Ω and a surjective map $F : S_{Y^n}(M) \to \Omega$ such that the set $\{(p, q) \in S_{Y^n}(M) \times S_{Y^n}(M) \mid F(p) = F(q)\}$ is closed. For example, consider a type-definable equivalence relation $E(x, y)$ on $Y^n \times Y^n$ with parameters over M and set $p \sim q$ if and only if

$$p(x) \cup q(y) \cup E(x, y) \text{ is consistent.}$$

The relation $\sim$ is a closed equivalence relation on $S_{Y^n}(M)$, so set Ω to be the collection of $\sim$-equivalence classes and F the natural projection map.

We can now equip Ω with the final topology with respect to F, so a subset C of Ω is closed if and only if $F^{-1}(C)$ is closed in the topological space $S_{Y^n}(M)$. It is immediate to see that Ω with this topology becomes a compact Hausdorff separable space. Furthermore, we can define a measure on Ω, the *push-forward measure* $F_*\mu$, given by $F_*\mu(B) = \mu(F^{-1}(B))$ for every Borel subset B of Ω.

Fact 3.8 (disintegration theorem). *Consider the normalized measure μ_{Y^n} on the space of types $S_{Y^n}(M)$, so it becomes a probability space. Given $F : S_{Y^n}(M) \to \Omega$ as in Remark 3.7, there exists a* ***disintegration*** *of μ_{Y^n} by a (uniquely determined) family of Radon conditional probability measures on $S_{Y^n}(M)$ with respect to the continuous function $F : S_{Y^n}(M) \to \Omega$, i.e., there exists a mapping*

$$(Z, t) \mapsto \nu(Z, t) = \mu_t(Z),$$

where Z is a Borel set of $S_{Y^n}(M)$ and t is an element of Ω, with the following properties:

(a) *for all t in Ω, the measure μ_t is a Borel inner regular probability measure on $S_{Y^n}(M)$;*

(b) *for every measurable subset Z of $S_{Y^n}(M)$, the function $t \mapsto \mu_t(Z)$ is measurable with respect to the measure $F_*\mu_{Y^n}$;*

(c) *each measure μ_t is concentrated on the fiber $F^{-1}(t)$, that is, the measure $\mu_t(S_{Y^n}(M) \setminus F^{-1}(t)) = 0$, so $\mu_t(Z) = \mu_t(Z \cap F^{-1}(t))$ for every Borel subset Z of $S_{Y^n}(M)$;*

(d) *for every measurable function $f : S_{Y^n}(M) \to \mathbb{R}$, we have that*

$$\int_{S_{Y^n}(M)} f \, d\mu_{Y^n} = \int_{t\in\Omega} \int_{F^{-1}(t)} f \, d\mu_t \, dF_*\mu_{Y^n}.$$

In particular, setting f to be the characteristic function $\mathbb{1}_Z$ of the measurable subset Z of $S_{Y\times Y}(M)$, we have that

$$\mu_{Y^n}(Z) = \int_{t\in\Omega} \mu_t(Z) \, dF_*\mu_{Y^n}.$$

Lemma 3.9. *Consider the natural restriction map*

$$\pi : S_{Y^2}(M) \to S_Y(M) \times S_Y(M), \quad q(y_1, y_2) \mapsto (q{\restriction_{y_1}}(y_1), q{\restriction_{y_2}}(y_2)).$$

Every pair of types (p, r) of $S_Y(M) \times S_Y(M)$ outside of a $\pi_*\mu_{Y^2}$-measure 0 set can be completed to a random type of $S_{Y^2}(M)$.

Proof. Let $\mathrm{R}(\mathcal{B}_M^{Y^2})$ be the Borel set of random types on $S_{Y^2}(M)$. It follows from Remark 1.14 that $\mu_{Y^2}(\mathrm{R}(\mathcal{B}_M^{Y^2})) = 1$. Apply now the disintegration theorem (Fact 3.8) with $\Omega = S_Y(M) \times S_Y(M)$ and $F = \pi$, and deduce from

$$1 = \mu_{Y^2}(\mathrm{R}(\mathcal{B}_M^{Y^2})) = \int_{(p,r)\in S_Y(M)\times S_Y(M)} \mu_{(p,r)}(\mathrm{R}(\mathcal{B}_M^{Y^2})) \, d\pi_*\mu_{Y^2}$$

that $\mu_{(p,r)}(\mathrm{R}(\mathcal{B}_M^{Y^2})) = 1$ for $\pi_*\mu_{Y^2}$-almost all pairs (p, r), since each function $\mu_{(p,r)}$ takes values in the interval [0, 1]. In particular, the set $\pi^{-1}(p, r) \cap \mathrm{R}(\mathcal{B}_M^{Y^2})$ is nonempty by Fact 3.8(c). Every such completion yields a random pair (a, b) over M, with a realizing p and b realizing r, as desired. □

Theorem 3.10. *For every pair of types (p, r) of $S_Y(M) \times S_Y(M)$ outside of a $\pi_*\mu_{Y^2}$-measure 0 set and every weakly random type $q = \mathrm{tp}(b/M)$ concentrated on Y with $\mathrm{Cos}(p) \cdot \mathrm{Cos}(q) = \mathrm{Cos}(r)$, there is a realization a of p weakly random over M, b such that $a \cdot b$ realizes r.*

Proof. By Lemma 3.9, for every pair (p, r) of $S_Y(M) \times S_Y(M)$ outside of a $\pi_*\mu_{Y^2}$-measure 0 set there exists a random pair (c, d) over M, with c realizing p and d realizing r. By Remark 1.13 and Lemma 1.15, the pair $(c^{-1} \cdot d, d)$ is random over M, so the partial type $p \cdot (c^{-1} \cdot d) \cap r$ admits a random realization, and thus it is weakly random. The element $c^{-1} \cdot d$ is (weakly) random over M and belongs to $\mathrm{Cos}(q)$, since p, q and r are coset-compatible. We can thus apply Corollary 3.3 to deduce that $p \cdot b \cap r$ is weakly random. Choose some realization f of this partial type weakly random over M, b and notice that the element $a = f \cdot b^{-1}$ is weakly random over M, b and realizes p. By construction, the product $a \cdot b = f$ realizes r, as desired. □

Whilst Theorem 3.10 holds for almost all types (p, r), the corresponding $\pi_*\mu_{Y^2}$-measure 0 set could possibly contain all diagonal pairs (p, p), with p in $\mathrm{supp}_M(\mu)$. We conclude this section with an elementary observation, the consequences of which will be explored in detail in Section 4.

Remark 3.11. Fix a countable elementary substructure M. If there exists a random pair (a, b) over M with $a \equiv_M b$, then there exists a random type concentrated in $\langle X \rangle_M^{00}$. Indeed, the element $b^{-1} \cdot a$ is random over M by Remark 1.13 and Lemma 1.15. Clearly, the element $g = b^{-1} \cdot a$ lies in $\langle X \rangle_M^{00}$, as desired.

Question. Is there a random pair (a, b) over M with $a \equiv_M b$? More generally, is there a random type concentrated in $\langle X \rangle_M^{00}$?

A digression: Roth's theorem on arithmetic progressions

We now show how Corollary 3.2 yields solutions to the equation $x \cdot z = y^2$ in subsets of positive density for every definably amenable pair such that the squaring function $x \mapsto x^2$ preserves randomness.

Definition 3.12. The function $f : X \to G$ in the definably amenable pair (G, X) *preserves randomness* if for every element a in X and every subset C of parameters, we have that a is (weakly) random over C if and only if $f(a)$ is (weakly) random over C (so $f(a)$ must lie in $\langle X \rangle$).

Remark 3.13. Examples 1.5, 1.6 and 1.7 always have the property that the square function preserves randomness if the map $f : X \to G$ defined by $f(x) = x^2$ has finite fibers. This is always the case whenever X has distinct squares as in [Sanders 2017, Theorem 1.5] or if G is abelian and there are only finitely many involutions in $\langle X \rangle$.

Theorem 3.14. *Consider a definably amenable pair (G, X) such that the square function preserves randomness. If the definable subset A of X has positive density, then the set*

$$\{(x_1, x_2) \in A \times A \mid x_1 \cdot x_2 \in A^2\}$$

has positive μ_2-density, where $A^2 = \{x^2\}_{x \in A}$.

Assume A is definable over the countable elementary substructure M. Every pair (a, c) in the above set random over M gives raise to a *generalized* 3-AP in A. Indeed, the product $a \cdot c$ belongs to A^2 so $a \cdot c = b^2$ for some b in A. Since the square function preserves randomness, we have that b is random over M, a by Lemma 1.15. Set now $g = b^{-1} \cdot a = b \cdot c^{-1}$ and observe that the elements c, $g \cdot c$ and $g \cdot c \cdot g$ all belong to A. If the group is abelian, this is an actual 3-AP as in the introduction.

Proof. We may assume that A is definable over a countable elementary substructure M, so it contains a weakly random type p over M. Choose some weakly random element g in $\langle X \rangle^{00}_M$. By Fact 2.6, the partial type $p \cdot g \cap p$ is weakly random, so the set $A \cdot g \cap A$ has positive density. By Remark 1.14, choose an element a in $A \cdot g \cap A$ random over M, g and notice that $b = a \cdot g^{-1}$ lies in A as well.

Since squaring preserves randomness, the element a^2 is also random over M, g and hence so is $a \cdot b = a^2 \cdot g^{-1}$ by Remark 1.13. By Lemma 1.16, the element g is weakly random over $M, a \cdot b$, and hence $a^2 = (a \cdot b) \cdot g$ is weakly random over $M, a \cdot b$. We deduce that a is weakly random over $M, a \cdot b$, for squaring preserves randomness. Furthermore, multiplying on the left by $(a \cdot b)^{-1}$ we conclude that b^{-1}, and hence b, is weakly random over $M, a \cdot b$.

Note that b belongs to $A^{-1} \cdot (a \cdot b) \cap A$, so this intersection must have positive density. Corollary 3.2 yields that the set $A^{-1} \cdot a^2 \cap A$ has positive measure, for a^2 and $a \cdot b$ lie in the same coset modulo $\langle X \rangle^{00}_M$. Choose now some random element a_1 in A over M, a with $a_1^{-1} \cdot a^2 = a_2$ in A. Remark 1.13 and Lemma 1.15 yield that the pair (a_1, a_2) is random over M. Thus, the M-definable set

$$\{(x_1, x_2) \in A \times A \mid x_1 \cdot x_2 \in A^2\}$$

has positive μ_2-measure, as desired. □

Question. Consider a definably amenable pair (G, X) such that the square function preserves randomness and let M be a countable elementary substructure M. Given an M-definable subset A of $\langle X\rangle$ of positive density, does the M-definable set

$$\{(x_1,x_2)\in A\times A \mid x_2\cdot x_1^{-1}\cdot x_2\in A\}$$

have positive μ_2-density? Equivalently, is there a random pair (a,b) in $A\times A$ over M with $b\cdot a^{-1}\cdot b$ in A?

Such a pair (a,b) as above yields a 3-AP in A of the form $(a, a\cdot g, a\cdot g^2)$ with $g=a^{-1}b$. We do not currently know whether the above question has a positive answer, though it is the case for ultra-quasirandom groups, by [Tao 2013].

Remark 3.15. The proof of Theorem 3.14 in the abelian context yields immediately the existence of solutions to translation-invariant equations of the form

$$n_1x_1+\cdots+n_mx_m=ky,$$

whenever $k=\sum_{j=1}^m n_j$ and each of the maps $x\mapsto n_1x$, $x\mapsto kx$ and $x\mapsto n'x$ preserves randomness, with $n'=\sum_{j=2}^m n_j$. That is, for every M-definable subset A of X of positive density, the set

$$\mathcal{E}(A)=\Big\{(x_1,\ldots,x_m)\in A\times\overset{m}{\cdots}\times A \mid n_1x_1+\cdots+n_mx_m=kc \text{ for some } c \text{ in } A\Big\}$$

has positive μ_m-measure. Indeed, choose g, a and b as in the proof of Theorem 3.14, so $g=a-b$. If we denote by $\ell A=\{\ell d\}_{d\in A}$, we will first show that the set $n_1A+n'a\cap kA$ has positive density: By Corollary 3.2, we need only show that $n_1A+ka-n_1b\cap kA$ has positive density. Now, the element

$$ka-n_1b=n'a+n_1g$$

is random over M, g, since a is random over M, g. So g, and thus kg, is weakly random over $M, ka-n_1b$ by Lemma 1.16. Since

$$kg=ka-kb=(ka-n_1b)-n'b,$$

we deduce that $-n'b$, and hence n_1b, is weakly random over $M, ka-n_1b$. Hence, the element $ka=n_1b+(ka-n_1b)$ is weakly random over $M, ka-n_1b$ and belongs to $n_1A+ka-n_1b\cap kA$, as desired.

Choose now $a_2,\ldots,a_m$ realizations of $\mathrm{tp}(a/M)$ with each a_j weakly random over $M, a, g, a_2,\ldots,a_{j-1}$. Hence, the differences a_j-a all belong to $\langle X\rangle_M^{00}$, by Fact 2.6. Corollary 3.2 and the above paragraph yield that $n_1A+\sum_{j=2}^m n_ja_j\cap kA$ has positive density, so choose an element a_1 in A weakly random over $M, a_2,\ldots,a_m$ exemplifying that the above intersection has positive density. The weakly random type $\mathrm{tp}(a_1,\ldots,a_m/M)$ contains the M-definable set $\mathcal{E}(A)$, as desired.

4. Ultra-quasirandomness revisited

Given a definably amenable pair (G, X) with $\langle X\rangle = G$, a straightforward application of compactness yields that $X^{\odot n} = G$ for some natural number n, so X generates G in finitely many steps. Up to scaling the σ-finite measure, we may assume that $G = X$, so $\mu(G) = 1$. This observation, together with Examples 1.5(a) and 1.7, motivates the following notion.

Definition 4.1. Let (G, X) be a definably amenable pair with $X = G$. We say that the pair is *generically principal* if $G = G^{00}_M$ for some elementary substructure M.

In an abuse of notation, we simply say that the group G is generically principal.

Remark 4.2. By [Martin-Pizarro and Palacín 2019, Corollary 2.6], a group G is generically principal if and only if $G = G^{00}_M$ for every elementary substructure M, so we may assume that M is countable.

In particular, a generically principal group contains trivially random elements concentrated in $\langle X\rangle^{00}_M = G$ for every countable elementary substructure M.

Example 4.3. Three known classes of groups are generically principal:

- Connected stable groups, such as every connected algebraic group over an algebraically closed field.
- Simple definably compact groups definable in some o-minimal expansion of a real closed field, such as $\mathrm{PSL}_n(\mathcal{R})$.
- Ultra-quasirandom groups, introduced in [Bergelson and Tao 2014]. Let us briefly recall this notion. A finite group is *d-quasirandom*, with $d \geq 1$, if all its nontrivial representations have degree at least d. An ultraproduct of finite groups $(G_n)_{n\in\mathbb{N}}$ with respect to a nonprincipal ultrafilter $\mathcal{U}$ is said to be *ultra-quasirandom* if for every integer $d \geq 1$, the set $\{n \in \mathbb{N} \mid G_n \text{ is } d\text{-quasirandom}\}$ belongs to $\mathcal{U}$.

The work of Gowers [2008, Theorem 3.3] yields that every definable subset A of positive density of an ultra-quasirandom group $G(M)$ is not *product-free*, i.e., it contains a solution to the equation $xy = z$, and thus the same holds in every elementary extension. Therefore, no weakly random type over an elementary substructure is product-free and thus $G = G^{00}_N$ over any elementary substructure N by [Martin-Pizarro and Palacín 2019, Corollary 2.6], so ultra-quasirandom groups are generically principal.

Proposition 3.1 and its corollaries yield now a short proof of the result mentioned in the above paragraph.

Lemma 4.4. *The following conditions are equivalent for a definably amenable pair (G, G):*

(a) *The group G is generically principal.*

(b) *Given two definable subsets A and B of positive density, we have that $A \cdot B$ has measure* 1. *In particular, whenever the definable subset C has positive measure, so is $G = A \cdot B \cdot C$.*

(c) *There is no definable product-free set of positive density.*

Proof. For (a) $\Rightarrow$ (b): Given two subsets A and B of positive density definable over some countable elementary substructure M, we need only show that every weakly random element g lies in $A \cdot B$. Now, Lemma 3.5 yields that there exists some random element h over M with $\mu(A \cap hB^{-1}) > 0$. Corollary 3.2 gives that every element g of G weakly random over M satisfies that $\mu(A \cap gB^{-1}) > 0$ as well. So the definable set $A \cdot B$ has measure 1, as desired.

For the second assertion, given a definable set C of positive density, let g in G be arbitrary. Now,

$$\mu(A \cdot B \cap gC^{-1}) = \mu(gC^{-1}) = \mu(C) > 0,$$

so g belongs to $A \cdot B \cdot C$, as desired.

The implication (b) $\Rightarrow$ (c) is clear, taking A and B to be the same set. Thus, we are left to consider the implication (c) $\Rightarrow$ (a). Suppose that $G \neq G^{00}_M$ for some countable elementary substructure M and take a weakly random type p in a nontrivial coset $\mathrm{Cos}(p)$ of G^{00}_M. Note that $p^{-1} \cdot p \cdot p \subseteq \mathrm{Cos}(p)$. A standard compactness argument yields the existence of some M-definable set A in p such that id_G does not lie in $A^{-1} \cdot A \cdot A$, so A is product-free. Since p is weakly random, the definable subset A has positive density. □

The following result on weak mixing, already present as is in [Bergelson and Tao 2014], was implicit in [Gowers 2008]. It will play a crucial role in studying some instances of complete amalgamation of equations in a group.

Corollary 4.5 (cf. [Bergelson and Tao 2014, Lemma 33]). *Let G be a generically principal group. Given two definable subsets A and B of positive density,*

$$\mu(A \cap gB) = \mu(A)\mu(B)$$

for μ-almost all elements g.

Proof. As before, fix some countable elementary substructure M such that both A and B are M-definable. We may assume that the measure μ is also definable over M. By Corollary 3.2, set $\alpha = \mu(A \cap gB)$ for some (or equivalently, every) weakly random element g over M. Notice that $\alpha > 0$ by Remark 3.6.

The subset

$$Z = \{x \in A \cdot B^{-1} \mid \mu(A \cap xB) = \alpha\}$$

is type-definable over M and contains all weakly random elements over M. Clearly, the measure $\mu(Z) \leq \mu(AB^{-1})$ and the latter equals 1, by Lemma 4.4. If we have

$\mu(Z) < \mu(A \cdot B^{-1})$, there is an M-definable set $\widetilde{Z}$ with $Z \subseteq \widetilde{Z} \subseteq A \cdot B^{-1}$ such that $\mu(A \cdot B^{-1} \setminus \widetilde{Z}) > 0$. Thus, the set $A \cdot B^{-1} \setminus \widetilde{Z}$ has positive density and it must contain a weakly random element over M, which gives the desired contradiction, so $\mu(Z) = \mu(A \cdot B^{-1}) = 1$.

Consider now the set $W = \{(a, z) \in A \times A \cdot B^{-1} \mid z = a \cdot b^{-1} \text{ for some } b \text{ in } B\}$. Note a belongs to $A \cap z \cdot B$ and z lies in aB^{-1} if (a, z) belongs to W. If we denote by μ_2 the normalized measure in $G \times G$, an easy computation yields that

$$\mu_2(W) = \int_{z \in A \cdot B^{-1}} \mu(A \cap zB) = \alpha \mu(A \cdot B^{-1}) \overset{4.4}{=} \alpha.$$

By Fubini, we also have that

$$\alpha = \mu_2(W) = \int_{a \in A} \mu(aB^{-1}) = \int_{a \in A} \mu(B) = \mu(A)\mu(B),$$

which gives the desired conclusion. □

A standard translation using Łoś's theorem yields the following finitary version:

Corollary 4.6 (cf. [Gowers 2008, Lemma 5.1; Bergelson and Tao 2014, Proposition 3]). *For every positive δ, ϵ and η there is some integer $d = d(\delta, \epsilon, \eta)$ such that for every finite d-quasirandom group G and subsets A and B of G of density at least δ, we have that*

$$\Big|\{x \in G \mid |A \cap xB||G| < (1-\eta)|A||B|\}\Big| < \epsilon|G|.$$

Proof. Assume for a contradiction that the statement does not hold, so there are some fixed positive numbers δ, ϵ and η such that for each natural number d we find two subsets A_d and B_d of a finite d-quasirandom group G_d, each of density at least δ, such that the cardinality of the subset

$$\mathcal{X}(G_d) = \{x \in G_d \mid |A_d \cap xB_d||G_d| < (1-\eta)|A_d||B_d|\}$$

is at least $\epsilon|G_d|$.

Following the approach of Example 1.5(a), we consider a suitable expansion $\mathcal{L}$ of the language of groups and regard each group G_d as an $\mathcal{L}$-structure N_d. Choose a nonprincipal ultrafilter $\mathcal{U}$ on $\mathbb{N}$ and consider the ultraproduct $N = \prod_{\mathcal{U}} N_d$. The language $\mathcal{L}$ is chosen in such a way that the sets $A = \prod_{\mathcal{U}} A_d$ and $B = \prod_{\mathcal{U}} B_d$ are $\mathcal{L}$-definable in the ultra-quasirandom group $G = \prod_{\mathcal{U}} G_d$. Furthermore, the normalized counting measure on G_d induces a definable Keisler measure μ on G, taking the standard part of the ultralimit. By Corollary 4.5, for μ-almost all g in G, we have $\mu(A \cap gB) = \mu(A)\mu(B)$. Hence, the type-definable set

$$\Sigma = \{x \in G \mid \mu(A \cap xB) \leq (1-\eta)\mu(A)\mu(B)\}$$

does not contain any weakly random type. By compactness, it is contained in a definable set W whose density is 0, and in particular its density is strictly less than the fixed value ϵ. Since every element in the ultraproduct of the sets $\mathcal{X}(G_d)$ clearly lies in Σ, we conclude by Łoś's theorem that $|\mathcal{X}(G_d)| \leq |W(G_d)| < \epsilon|G_d|$ for infinitely many d, which yields the desired contradiction. □

The following result is a verbatim adaption of [Gowers 2008, Theorem 5.3] and may be seen as a first attempt to solve complete amalgamation problems whilst restricting the conditions to those given by products.

Theorem 4.7. *Fix a natural number $n \geq 2$. For each nonempty subset F of $\{1, \dots, n\}$, let A_F be a subset of positive density of the generically principal group G. The set*

$$\mathcal{X}_n = \big\{(a_1, \dots, a_n) \in G^n \mid a_F \in A_F \text{ for all } \varnothing \neq F \subseteq \{1, \dots, n\}\big\}$$

has measure $\prod_F \mu(A_F)$ with respect to the measure μ_n on G^n, where a_F stands for the product of all a_i with i in F written with the indices in increasing order.

Proof. We reproduce Gowers's proof of [Gowers 2008, Theorem 5.3] and proceed by induction on n. For $n = 2$, set $B = A_{\{2\}}$ and $C = A_{\{1,2\}}$. A pair (a, b) lies in $\mathcal{X}_2$ if and only if a belongs to $A_{\{1\}}$ and b to $B \cap a^{-1}C$. Thus

$$\mu_2(\mathcal{X}_2) = \int_{A_{\{1\}}} \mu(B \cap a^{-1}C)\, d\mu \overset{4.5}{=} \mu(B)\mu(C)\mu(A_{\{1\}}),$$

as desired. For the general case, for any a in $A_{\{1\}}$, set $B_{F_1}(a) = A_{F_1} \cap a^{-1}A_{1,F_1}$, for $\varnothing \neq F_1 \subseteq \{2, \dots, n\}$. Corollary 4.5 yields that $\mu(B_{F_1}(a)) = \mu(A_{F_1})\mu(A_{1,F_1})$ for μ-almost all a in $A_{\{1\}}$. A tuple $(a_1, \dots, a_n)$ in G^n belongs to $\mathcal{X}_n$ if and only if the first coordinate a_1 lies in $A_{\{1\}}$ and the tuple $(a_2, \dots, a_n)$ belongs to

$$\mathcal{X}_{n-1}(a_1) = \big\{(x_2, \dots, x_n) \in G^{n-1} \mid x_{F_1} \in B_{F_1}(a_1) \text{ for all } \varnothing \neq F_1 \subseteq \{2, \dots, n\}\big\}.$$

By induction, the set $\mathcal{X}_{n-1}(a)$ has constant μ_{n-1}-measure $\prod_{F_1} \mu(A_{F_1})\mu(A_{1,F_1})$, where F_1 now runs through all nonempty subsets of $\{2, \dots, n\}$. Thus

$$\mu_n(\mathcal{X}_n) = \int_{A_1} \mu_{n-1}(\mathcal{X}_{n-1}(a_1))\, d\mu = \mu(A_1) \prod_{F_1} \mu(A_{F_1})\mu(A_{1,F_1}) = \prod_F \mu(A_F),$$

which yields the desired result. □

A standard translation using Łoś's theorem (we refer to the proof of Corollary 4.6 to avoid repetitions) yields the following finitary version, which was already present in a quantitative form in [Gowers 2008].

Corollary 4.8 (cf. [Gowers 2008, Theorem 5.3]). *Fix a natural number $n \geq 2$. For every $\varnothing \neq F \subseteq \{1, \dots, n\}$ let $\delta_F > 0$ be given. For every $\eta > 0$ there is some integer*

$d = d(n, \delta_F, \eta)$ *such that for every finite d-quasirandom group G and subsets* A_F *of G of density at least* δ_F*, we have that*

$$|\mathcal{X}_n| \geq \frac{1-\eta}{|G|^{2^n-1-n}} \prod_F |A_F|,$$

where $\mathcal{X}_n$ *is defined as in Theorem 4.7 with respect to the group G.*

The above corollary yields in particular that

$$\big|\{(a,b,c) \in A \times B \times C \mid ab = c\}\big| > \frac{1-\eta}{|G|}|A||B||C|$$

as first proved by Gowers [2008, Theorem 3.3], which implies that the number of such triples is a proportion (uniformly on the densities and η) of $|G|^2$.

To conclude this section we answer affirmatively the question in the introduction for generically principal groups, whenever all the types are based over a common countable elementary substructure.

Theorem 4.9. *Fix a natural number* $n \geq 2$ *and a countable elementary substructure M of the generically principal definably amenable pair* (G, X)*. For each nonempty subset F of* $\{1, \dots, n\}$*, let* p_F *be a weakly random type over M. There exists a weakly random n-tuple* $(a_1, \dots, a_n)$ *in* G^n *such that* a_F *realizes* p_F *for all nonempty* $F \subseteq \{1, \dots, n\}$*, where* a_F *stands for the product of all* a_i *with i in F written with the indices in increasing order.*

Proof. Since M is countable, enumerate all the formulae occurring in each type p_F in a decreasing way, that is, write $p_F = \{A_{F,k}\}_{k\in\mathbb{N}}$ with $A_{F,k+1} \subseteq A_{F,k}$ for every natural number k. We want to show that the set

$$\mathcal{X}_n = \{(x_1, \dots, x_n) \in G^n \mid p_F(x_F) \text{ for all } \varnothing \neq F \subseteq \{1, \dots, n\}\}$$

is weakly random over M, that is, we need to prove that the partial type

$$\{\neg\psi(x_1, \dots, x_n)\}_{\psi\in\Sigma} \cup \{x_F \in A_{F,k}\}_{F\in\mathcal{P},\, k\in\mathbb{N}}$$

is consistent, where $\mathcal{P} = \mathcal{P}(\{1, \dots, n\}) \setminus \{\varnothing\}$ and Σ is the set of $\mathcal{L}_M$-formulae of μ_n-measure 0. By compactness, since the subsets $A_{F,k}$ are enumerated decreasingly, we need only consider a finite subset of the above partial type where the level k_0 is the same for each of the subsets A_{F,k_0} of positive density. By Theorem 4.7 the set

$$\mathcal{X}_{n,k_0} = \{(a_1, \dots, a_n) \in G^n \mid a_F \in A_{F,k_0} \text{ for all nonempty } F \subseteq \{1, \dots, n\}\}$$

has μ_n-measure $\prod_F \mu(A_{F,k_0}) > 0$, so we conclude the desired result. □

5. Local ultra-quasirandomness

In this final section, we adapt some of the ideas present in Section 4 to arbitrary finite groups.

Theorem 3.10 holds in any definably amenable pair for *almost all* three weakly random types, whenever their cosets modulo G^{00}_M are product-compatible. Thus, it yields asymptotic information for subsets of positive density in arbitrary finite groups satisfying certain regularity conditions, which force some completions in the ultraproduct to be in a suitable position to apply our main Theorem 3.10. We present two examples of such regularity notions. Our intuition behind these notions is purely model-theoretic and we ignore whether it is meaningful from a combinatorial perspective. We would like to express our gratitude to Julia Wolf (and indirectly to Tom Sanders) for pointing out that our previous definition of principal subsets did not extend to the abelian case.

Definition 5.1. Let A be a definable subset of $\langle X\rangle$ of positive density in a definably amenable pair (G, X). We say that A is *principal* over the parameter set B if

$$\mu(A\cap(Y\cdot Y))>0$$

whenever Y is a B-definable neighborhood of the identity (that is, the set Y is symmetric and contains the identity) such that finitely many left translates of Y cover $A\cdot A^{-1}\cdot A\cdot A^{-1}$.

Analogously, we say that A is *hereditarily principal over the parameter set B* if all of its B-definable subsets of positive density are principal.

Remark 5.2. Let A be a definable subset of $\langle X\rangle$ of positive density of a definably amenable pair (G, X) such that $\mu(A\cap(Y\cdot Y))=\mu(A)$, whenever Y is a definable neighborhood of the identity which covers $A\cdot A^{-1}\cdot A\cdot A^{-1}$ with finitely many left translates. Then the set A is hereditarily principal over any subset of parameters.

Proof. Let A_0 be a definable subset of A of positive measure. Notice that there is a maximal finite subset F of $(AA^{-1})^2$ with the property that $\mu(xA_0\cap yA_0)=0$ for any two distinct x and y in F. In particular, the set $(AA^{-1})^2$ is contained in $F\cdot A_0\cdot A_0^{-1}$. Thus, any definable neighborhood Y of the identity such that finitely many left translates of cover $A_0A_0^{-1}A_0A_0^{-1}$ also cover $AA^{-1}AA^{-1}$, so $\mu(A\cap(Y\cdot Y))=\mu(A)$ by assumption on A. Hence $\mu(A_0\cap(YY))=\mu(A_0)>0$, as desired. □

Example 5.3. If G is generically principal, every definable subset A of positive density is hereditarily principal over any parameter set. Indeed, Lemma 4.4 yields that $G=A\cdot A^{-1}\cdot A\cdot A^{-1}$. Therefore, finitely many translates of the neighborhood Y must cover G, so Y has positive measure and hence $\mu(Y\cdot Y)=1$ by Lemma 4.4.

By the previous remark, the definable subset A satisfies that $\mu(A\cap(Y\cdot Y))=\mu(A)$, so A is hereditarily principal over any subset of parameters.

Example 5.4. Fix some enumeration $(q_n)_{n\in\mathbb{N}}$ of all the primes and consider the family of groups $(G_n = \mathrm{PSL}_2(q_n) \times \mathbb{Z}_2)_{n\in\mathbb{N}}$, each equipped with the distinguished subset $X_n = \mathrm{PSL}_2(q_n) \times \{\bar{0}\}$. This family produces a definably amenable pair (G, X), as in Example 1.5. Note that

$$G = \mathrm{PSL}_2(\mathbb{F}) \times \mathbb{Z}_2 \quad \text{and} \quad X = \mathrm{PSL}_2(\mathbb{F}) \times \{\bar{0}\}$$

for some infinite (pseudofinite) field $\mathbb{F}$. Over any elementary substructure M we have that G_M^{00} equals the simple group $X = \mathrm{PSL}_2(\mathbb{F}) \times \{\bar{0}\}$, which is clearly definable. The definable subset G is clearly principal yet not hereditarily principal, for the dense subset $X \cdot (0_{\mathrm{PSL}_2(\mathbb{F})}, \bar{1})$ does not intersect $X = G_M^{00}$.

Lemma 5.5. *Let M be a countable elementary substructure of a definably amenable pair (G, X).*

(a) *Principal definable sets over M contain weakly random principal types in $S_\mu(M)$, that is, types concentrated in $\langle X\rangle_M^{00}$.*
(b) *Every weakly random type over M containing a hereditarily principal definable set is principal.*

Proof. For (a), assume that the M-definable set A is principal over the model M. Note that we can write the type-definable subgroup $\langle X\rangle_M^{00}$ as a countable intersection

$$\langle X\rangle_M^{00} = \bigcap_{i\in\mathbb{N}} V_i,$$

where the decreasing chain $(V_i)_{i\in\mathbb{N}}$ consists of M-definable neighborhoods of the identity such that $V_{i+1} \cdot V_{i+1} \subseteq V_i$ for all i in $\mathbb{N}$. Since $\langle X\rangle_M^{00}$ has bounded index in the subgroup $\langle X\rangle$, compactness yields that finitely many translates of each V_i cover the subset $A \cdot A^{-1} \cdot A \cdot A^{-1}$ (yet the number of translates possibly depends on i). Hence, the type-definable subset $A \cap \langle X\rangle_M^{00}$ is weakly random, since A is principal, so A contains a weakly random type concentrated in $\langle X\rangle_M^{00}$, as desired.

For (b), suppose that the M-definable set A is hereditarily principal yet it contains a weakly random type q which does not concentrate on $\langle X\rangle_M^{00} = \bigcap_{i\in\mathbb{N}} V_i$, with the same notation as above. By compactness, this implies the existence of some i in $\mathbb{N}$ and some M-definable subset A_0 of A of positive density with $A_0 \cap V_i = \emptyset$. The subset $A_0 \cap (V_{i+1} \cdot V_{i+1})$ has in particular measure 0, so A_0 is not principal, contradicting our assumption on A. □

Proposition 5.6. *Consider a subset A of positive density definable over a countable elementary substructure M of a sufficiently saturated definably amenable pair (G, X). If A contains a weakly random type p concentrated in $\langle X\rangle_M^{00}$, then the subset*

$$\{(a, b) \in A \times A \mid a \cdot b \in A\}$$

has positive μ_2-measure. In particular, if A is principal, then the above set of pairs has positive μ_2-measure.

Notice that the definable set A above cannot be product-free, for the equation $x \cdot y = z$ has a solution in A.

Proof. The proof is an immediate application of Fact 2.6. Indeed, for every realization a of p, the partial type $p \cap a^{-1} \cdot p$ is weakly random (for the weakly random element a over M belongs to $\langle X \rangle_M^{00}$), so choose a weakly random element b over M, a realizing p such that $a \cdot b$ does as well. By Lemma 1.10, we obtain a weakly random type $\mathrm{tp}(a, b/M)$ with all three elements a, b and $a \cdot b$ in A, which yields immediately the desired result. □

Proposition 5.6 resonates with [Schur 1917, Hilfssatz] on the number of monochromatic triples $(x, y, x \cdot y)$ in any finite coloring (or cover) of the natural numbers $1, \dots, N$, for N sufficiently large. In fact, by a standard application of Łoś's theorem, the above argument yields a nonquantitative version of the following result of Sanders [2019, Theorem 1.1]:

> *For every natural number $k \geq 1$ there is some $\eta = \eta(k) > 0$ with the following property: given any coloring on a finite group G with k many colors $A_1, \dots, A_k$, there exists some color A_j, with $1 \leq j \leq k$, such that*
>
> $$\left|\{(a, b, c) \in A_j \times A_j \times A_j \mid a \cdot b = c\}\right| \geq \eta |G|^2.$$

Motivated by [Gowers 2008, Theorem 5.3] for (ultra-)quasirandom groups, we now provide a weaker version of it, taking all A_F's to be the same subset A, for $\varnothing \neq F \subseteq \{1, \dots, n\}$ as in Corollary 4.8.

Corollary 5.7. *In a sufficiently saturated definably amenable pair (G, X) with associated measure μ, consider a definable subset A of X of positive density which is hereditarily principal over the parameter set G itself. For every countable elementary substructure M of (G, X) such that both the measure and the sets A are M-definable, there is a tuple $(a_1, \dots, a_n)$ in G^n weakly random over M such that the product a_F (as in Theorem 4.9) lies in A for every subset F as above.*

An inspection of the proof shows that it suffices if the definable set A is hereditarily principal over N, where N is an $\aleph_1$-saturated elementary substructure of (G, X) containing M. This is not surprising, since an easy compactness argument shows that a set A which is hereditarily principal over an $\aleph_1$-saturated elementary substructure N of (G, X) must be hereditarily principal over the parameter set G itself.

Proof. We proceed by induction on the natural number n. Since both the base case $n = 3$ and the induction step have similar proofs, we assume that the statement of the corollary has already been shown for $n - 1$.

The set A is principal, so it contains a weakly random type concentrated in $\langle X\rangle_M^{00}$, by Lemma 5.5(a). As in the proof of Proposition 5.6, there is a weakly random element a_1 in A over M such that $A' = A \cap a_1^{-1} \cdot A$ has positive density. Notice that A' is no longer definable over M, yet it is again hereditarily principal over the parameter set G. By downwards Löwenheim–Skolem, choose some countable elementary substructure M_1 of (G, X) containing $M \cup \{a_1\}$. By induction, there is a tuple $(a_2, \ldots, a_n)$, weakly random over M_1, such that each product a_{F_1} lies in A' for every subset $\varnothing \neq F_1 \subseteq \{2, \ldots, n\}$. For $n = 3$, we obtain such a tuple by applying Proposition 5.6 to the principal M_1-definable set A'.

Lemma 1.10 yields now that the tuple $(a_1, \ldots, a_n)$ is weakly random over M. By construction, the product a_F lies in A for every subset $\varnothing \neq F \subseteq \{1, \ldots, n\}$, as desired. □

Motivated by the above result, we isolate a particular instance of a complete amalgamation problem (cf. the question in the introduction).

Question. Let M be a countable elementary substructure of a sufficiently saturated definably amenable pair (G, X) and p be a weakly random type in $\langle X\rangle_M^{00}$. Given a natural number n, is there a tuple $(a_1, \ldots, a_n)$ in G^n weakly random over M such that a_F realizes p for all $\varnothing \neq F \subseteq \{1, \ldots, n\}$, where a_F stands for the product, enumerated in an increasing order, of all a_i with i in F?

At the moment of writing, we do not have a solid guess what the answer to the above question will be. Following the lines of the proof of Corollary 5.7, the above question would have a positive answer if the following statement is true:

> *Let $p = \mathrm{tp}(a/M_0)$ be a weakly random type in $\langle X\rangle_{M_0}^{00}$, where M_0 is a countable elementary substructure of a saturated definably amenable pair (G, X). Then there are an elementary substructure M_1 containing $M_0 \cup \{a\}$ and a weakly random type q in $\langle X\rangle_{M_1}^{00}$ extending $p \cap a^{-1} \cdot p$*

Nonetheless, if the question could be positively answered, it would imply by a standard compactness argument a finitary version of Hindman's theorem [Hindman 1974].

Remark 5.8. If the above question has a positive answer, then for every natural numbers k and n there is some constant $\eta = \eta(k, n) > 0$ such that in any coloring on a finite group G with k many colors $A_1, \ldots, A_k$, there exists some color A_j, with $1 \leq j \leq k$ such that

$$\left|\{(a_1, \ldots, a_n) \in G^n \mid a_F \in A_j \text{ for all } \varnothing \neq F \subseteq \{1, \ldots, n\}\}\right| \geq \eta |G|^n,$$

where a_F stands for the product, enumerated in an increasing order, of all a_i with i in F.

We can now state the finitary versions of principal sets to provide finitary analogs of Proposition 5.6 and Corollary 5.7.

Definition 5.9. Fix $\epsilon > 0$ and k in $\mathbb{N}$. A finite subset A of a group G is (k, ϵ)-*principal* if

$$|A \cap (Y \cdot Y)| \geq \epsilon |A|$$

whenever Y is a neighborhood of the identity (that is, the set Y is symmetric and contains the identity) such that k many left translates (or equivalently, right translates) of Y cover $A \cdot A^{-1} \cdot A \cdot A^{-1}$.

We say that the finite subset A is *hereditarily* (k, ϵ)*-principal up to* ρ if all its subset of relative density at least ρ (in A) are (k, ϵ)-principal.

Example 5.10. Consider the finite group $G = \mathbb{Z}_n \times \mathbb{Z}_2$. The set G is clearly $(k, 1/k)$-principal for every natural number $k \neq 0$, yet it is not hereditarily $(2, 1/k)$-principal up to $1/2$ for any $k \neq 0$, for the subset $A = \mathbb{Z}_n \times \{\bar{1}\}$ does not intersect $Y = \mathbb{Z}_n \times \{\bar{0}\}$, which covers G in 2 steps.

Example 5.11. Given a subset A of a finite group G of density at least ϵ, the symmetric set AA^{-1} is $(k, \epsilon/k)$-principal. Indeed, if Y is a given neighborhood of the identity such that k many right translates of Y cover $(AA^{-1})^4$, then there exists some c in G such that $|Ac \cap Y| \geq |A|/k$ and so $|AA^{-1} \cap YY| \geq \epsilon |AA^{-1}|/k$, since $(Ac \cap Y)(Ac \cap Y)^{-1} \subseteq AA^{-1} \cap YY$.

Corollary 5.12. *Let $K > 0$ and $\delta > 0$ be given real numbers. There are real values $\epsilon = \epsilon(K, \delta) > 0$ and $\eta = \eta(K, \delta) > 0$ as well as a natural number $k = k(K, \delta)$ such that for every group G and a finite subset X of G of tripling at most K together with a (k, ϵ)-principal subset A of X of relative density at least δ with respect to X, the collection of triples*

$$\{(a, b) \in A \times A \mid a \cdot b \in A\}$$

has size at least $\eta |X|^2$.

Proof. Assume for a contradiction that the statement does not hold. Negating quantifiers, there are positive constants K and δ such that for each triple $\bar{\ell} = (k, n, m)$ of natural numbers there exists a group $G_{\bar{\ell}}$ and a finite subset $X_{\bar{\ell}}$ of $G_{\bar{\ell}}$ of tripling at most K as well as a $(k, 1/n)$-principal subset $A_{\bar{\ell}}$ of $X_{\bar{\ell}}$ of relative density at least δ such that the cardinality of the subset

$$\mathcal{Y}(G_{\bar{\ell}}) = \{(x, y) \in A_{\bar{\ell}} \times A_{\bar{\ell}} \mid x \cdot y \in A_{\bar{\ell}}\}$$

is bounded above by $|X_{\bar{\ell}}|^2/m$.

Following the approach of Example 1.5(b), we consider a suitable countable expansion $\mathcal{L}$ of the language of groups and regard each such group $G_{\bar{\ell}}$, with $\bar{\ell}$ of the form (k, k, k), as an $\mathcal{L}$-structure $N_{\bar{\ell}}$ in such a way that $\mathcal{L}$ contains predicates for $X_{\bar{\ell}}$ and $A_{\bar{\ell}}$. Identify now the set of such triples (k, k, k) with the natural numbers in a natural way and choose a nonprincipal ultrafilter $\mathcal{U}$ on $\mathbb{N}$. Consider the ultraproduct $N = \prod_{\mathcal{U}} N_{\bar{\ell}}$. As outlined in Example 1.5, this construction gives rise to a definable amenable pair (G, X) with respect to a measure μ equipped with an $\varnothing$-definable subset A of X of positive density (at least δ) such that $\mu_2(\mathcal{Y}(G)) = 0$. Notice that A is now principal over the parameter set N, by Łoś's theorem.

Fix a countable elementary substructure M of N. By Proposition 5.6, the set

$$\mathcal{Y}(G) = \{(x, y) \in A \times A \mid x \cdot y \in A\}$$

has positive density with respect to μ_2, which contradicts the ultraproduct construction. $\square$

The proof of the next result follows from Corollary 5.7 along the same lines as Corollary 5.12 by a standard ultraproduct argument using Łoś's theorem (and implicitly that a nonprincipal ultraproduct of finite sets is $\aleph_1$-saturated).

Corollary 5.13. *For a natural number $n \geq 3$, let real numbers $K > 0$ and $\delta_F > 0$, with $\varnothing \neq F \subseteq \{1, \ldots, n\}$, be given. There are $\epsilon = \epsilon(n, K, \delta_F) > 0$, $\rho = \rho(n, K, \delta_F)$ and $\eta = \eta(n, K, \delta_F) > 0$ as well as a natural number $k = k(n, K, \delta_F)$ such that for every group G and a finite subset X of G of tripling at most K together with a subset A of X of relative density at least δ, whenever*

$$\left|\{(a_1, \ldots, a_n) \in G^n \mid a_F \in A \textit{ for all } \varnothing \neq F \subseteq \{1, \ldots, n\}\}\right| < \eta |X|^n,$$

where a_F stands for the product, enumerated in an increasing order, of all the a_i with i in F, then A cannot be hereditarily (k, ϵ)-principal up to ρ.

In order to extend Proposition 5.6 to pairs (a, b) in the cartesian product $A \times B$ with $a \cdot b$ in C, we introduce a new notion, which we refer to as compatibility for certain subsets in a definably amenable pair.

Definition 5.14. Let A, B and C be subsets of $\langle X \rangle$ of positive density in a definably amenable pair (G, X), all three definable over the countable elementary substructure M. We say that A and B are *compatible* with respect to C over M if there exists a random pair (a, b) in $A \times B$ over M such that $a \cdot b$ lies in the same coset modulo $\langle X \rangle_M^{00}$ as some element c of C which is weakly random over M.

It is clear that every two definable subsets A and B of positive density in a generically principal group G are compatible with respect to any subset C of positive density over any countable elementary substructure M containing the parameters of definition of all three sets. More generally, we have the following observation.

Remark 5.15. Given three definable subsets A, B and C of positive density at least $\delta > 0$ in a definably amenable pair (G, X), all three defined over a countable elementary substructure M, every weakly random type of $\langle X\rangle_M^{00}$ is contained in

$$A\cdot A^{-1}\cap B\cdot B^{-1}\cap C\cdot C^{-1},$$

by Fact 2.6. Hence, the M-definable set

$$\{(x,y)\in(A\cdot A^{-1})\times(B\cdot B^{-1})\mid x\cdot y\in C\cdot C^{-1}\}$$

contains a pair (a_1,b_1) with a_1 and b_1 both in $\langle X\rangle_M^{00}$ weakly random over M such that b_1 is weakly random over M,a_1. Hence, the above set has positive density, so there exists a random pair (a,b) in $AA^{-1}\times BB^{-1}$ over M such that $a\cdot b$ belongs to CC^{-1}. Since $a\cdot b$ is (weakly) random over M, we deduce that $A\cdot A^{-1}$ and $B\cdot B^{-1}$ are compatible with respect to $C\cdot C^{-1}$.

Lemma 5.16. *Let A, B and C be subsets of $\langle X\rangle$ of positive density in a definably amenable pair (G,X), all three definable over the countable elementary substructure M.*

(a) *If for some element g in $\langle X\rangle_M^{00}$ weakly random over M, the definable subset*

$$Z_g=\{(a,b)\in A\times B\mid a\cdot b\in C\cdot g\}$$

has positive μ_2-measure, then A and B are compatible with respect to C over M.

(b) *If A and B are compatible with respect to C over M, then the M-definable set*

$$\{(a,b)\in A\times B\mid a\cdot b\in C\}$$

has positive μ_2-measure.

Proof. For (a), given a weakly random element g in $\langle X\rangle_M^{00}$, suppose that the definable set Z_g has positive density. By Remark 1.14, choose some (a,b) in Z_g random over M,g, so the element $c=a\cdot b\cdot g$ is again random over M by Remark 1.13 and Lemma 1.15. This immediately yields that A and B are compatible with respect to C over M.

For (b), suppose that A and B are compatible with respect to C over M, so by definition, there is a random pair (a,b) in $A\times B$ over M such that $a\cdot b$ lies in the same coset of $\langle X\rangle_M^{00}$ as some element c in C whose type over M is weakly random. By Lemma 1.15, the pair $(a^{-1},a\cdot b)$ is a random pair over M, so the definable set $A^{-1}\cdot(a\cdot b)\cap B$ has positive measure, for it belongs to the weakly random type $\mathrm{tp}(b/M,a\cdot b)$. By Corollary 3.2, we deduce that $A^{-1}\cdot c\cap B$ has positive measure, so choose b_1 in B weakly random over M,c such that $c=a_1\cdot b_1$. In particular, the M-definable set

$$\{(y,z)\in B\times C\mid z\cdot y^{-1}\in A\}$$

has positive μ_2-measure and so it contains a random pair (b_2, c_2) over M. The pair (a_2, b_2) of $A \times B$, with $a_2 = c_2 \cdot b_2^{-1}$ is again random over M by Lemma 1.15 and satisfies that $a_2 \cdot b_2$ belongs to C, as desired. □

Remark 5.17. If the definable set A has positive density and the pair (A, A) is compatible with respect to A over a countable elementary substructure M, then A is not product-free (cf. the corresponding comment after Proposition 5.6). On the other hand, is it the case that every principal definable set yields a compatible pair? Or are the two notions unrelated, even if they provide the same positive answer?

Lemma 5.16 yields a sufficient condition to ensure that the corresponding ultraproducts of finite subsets will be compatible. We have several candidates of finitary versions of compatibility, which will allow us to obtain a local version of [Gowers 2008, Theorem 5.3] to count the number of pairs in $A \times B$ such that the product $a \cdot b$ lies in the subset C of positive density, all within a finite subset of small tripling. However, it is unclear to us how combinatorially relevant our tentative definitions are, so we would rather leave the ultraproduct formulation as an open question: Is there a meaningful combinatorial definition (akin to Definition 5.9) of when two finite sets A and B are compatible with respect to the finite set C?

Acknowledgements

We are most indebted to Ehud Hrushovski and Julia Wolf for their helpful remarks, which have considerably improved this article. We would like to express our gratitude to Angus Matthews for pointing out a mistake in a previous version of Theorem 3.10.

We would also like to express our sincere gratitude to the anonymous referee of the previous versions of this article for the insightful comments and remarks, as well as for providing a clever way to circumvent some additional assumptions in a previous proof of Roth's theorem (Theorem 3.14). The statement and the proof of Lemma 1.16 are solely due to the referee. We are most thankful that we are allowed to include the proof here.

References

[Bergelson and Tao 2014] V. Bergelson and T. Tao, "Multiple recurrence in quasirandom groups", *Geom. Funct. Anal.* **24**:1 (2014), 1–48. MR Zbl

[Bogachev 2007] V. I. Bogachev, *Measure theory*, vol. II, Springer, 2007. MR Zbl

[Conant et al. 2020] G. Conant, A. Pillay, and C. Terry, "A group version of stable regularity", *Math. Proc. Cambridge Philos. Soc.* **168**:2 (2020), 405–413. MR Zbl

[Conant et al. 2023] G. Conant, K. Gannon, and J. Hanson, "Keisler measures in the wild", *Model Theory* **2** (2023), 1–70. MR Zbl

[Croot and Sisask 2010] E. Croot and O. Sisask, "A probabilistic technique for finding almost-periods of convolutions", *Geom. Funct. Anal.* **20**:6 (2010), 1367–1396. MR Zbl

[Frankl et al. 1987] P. Frankl, R. L. Graham, and V. Rödl, "On subsets of abelian groups with no 3-term arithmetic progression", *J. Combin. Theory Ser. A* **45**:1 (1987), 157–161. MR Zbl

[Gowers 2008] W. T. Gowers, "Quasirandom groups", *Combin. Probab. Comput.* **17**:3 (2008), 363–387. MR Zbl

[Green and Tao 2008] B. Green and T. Tao, "The primes contain arbitrarily long arithmetic progressions", *Ann. of Math.* (2) **167**:2 (2008), 481–547. MR Zbl

[Halmos 1974] P. R. Halmos, *Measure theory*, Grad. Texts in Math. **18**, Springer, 1974. MR Zbl

[Henriot 2016] K. Henriot, "Arithmetic progressions in sets of small doubling", *Mathematika* **62**:2 (2016), 587–613. MR Zbl

[Hindman 1974] N. Hindman, "Finite sums from sequences within cells of a partition of N", *J. Combinatorial Theory Ser. A* **17** (1974), 1–11. MR Zbl

[Hrushovski 2012] E. Hrushovski, "Stable group theory and approximate subgroups", *J. Amer. Math. Soc.* **25**:1 (2012), 189–243. MR Zbl

[Hrushovski 2013] E. Hrushovski, "Approximate equivalence relations", preprint, 2013, available at http://www.ma.huji.ac.il/~ehud/approx-eq.pdf.

[Hrushovski et al. 2013] E. Hrushovski, A. Pillay, and P. Simon, "Generically stable and smooth measures in NIP theories", *Trans. Amer. Math. Soc.* **365**:5 (2013), 2341–2366. MR Zbl

[Keisler 1987] H. J. Keisler, "Measures and forking", *Ann. Pure Appl. Logic* **34**:2 (1987), 119–169. MR Zbl

[Martin-Pizarro and Palacín 2019] A. Martin-Pizarro and D. Palacín, "Stabilizers, measures and IP-sets", preprint, 2019. arXiv 1912.07252

[Martin-Pizarro and Pillay 2004] A. Martin-Pizarro and A. Pillay, "Elliptic and hyperelliptic curves over supersimple fields", *J. London Math. Soc.* (2) **69**:1 (2004), 1–13. MR Zbl

[Massicot and Wagner 2015] J.-C. Massicot and F. O. Wagner, "Approximate subgroups", *J. Éc. polytech. Math.* **2** (2015), 55–64. MR Zbl

[Montenegro et al. 2020] S. Montenegro, A. Onshuus, and P. Simon, "Stabilizers, NTP_2 groups with f-generics, and PRC fields", *J. Inst. Math. Jussieu* **19**:3 (2020), 821–853. MR Zbl

[Namioka 1964] I. Namioka, "Følner's conditions for amenable semi-groups", *Math. Scand.* **15** (1964), 18–28. MR Zbl

[Palacín 2020] D. Palacín, "On compactifications and product-free sets", *J. Lond. Math. Soc.* (2) **101**:1 (2020), 156–174. MR Zbl

[Peluse 2018] S. Peluse, "Mixing for three-term progressions in finite simple groups", *Math. Proc. Cambridge Philos. Soc.* **165**:2 (2018), 279–286. MR Zbl

[Pillay et al. 1998] A. Pillay, T. Scanlon, and F. O. Wagner, "Supersimple fields and division rings", *Math. Res. Lett.* **5**:4 (1998), 473–483. MR Zbl

[Roth 1953] K. F. Roth, "On certain sets of integers", *J. London Math. Soc.* **28** (1953), 104–109. MR Zbl

[Sanders 2009] T. Sanders, "Three-term arithmetic progressions and sumsets", *Proc. Edinb. Math. Soc.* (2) **52**:1 (2009), 211–233. MR Zbl

[Sanders 2017] T. Sanders, "Solving $xz = y^2$ in certain subsets of finite groups", *Q. J. Math.* **68**:1 (2017), 243–273. MR Zbl

[Sanders 2019] T. Sanders, "Schur's colouring theorem for noncommuting pairs", *Bull. Aust. Math. Soc.* **100**:3 (2019), 446–452. MR Zbl

[Schur 1917] I. Schur, "Über die Kongruenz $x^m + y^m = z^m \pmod{p}$", *Jahresber. Dtsch. Math.-Ver.* **25** (1917), 114–117. Zbl

[Simmons 2012] D. Simmons, "Conditional measures and conditional expectation; Rohlin's disintegration theorem", *Discrete Contin. Dyn. Syst.* **32**:7 (2012), 2565–2582. MR Zbl

[Simon 2015] P. Simon, *A guide to NIP theories*, Lecture Notes in Logic **44**, Assoc. Symbolic Logic, Chicago, 2015. MR Zbl

[Starchenko 2017] S. Starchenko, "NIP, Keisler measures and combinatorics", pp. 303–334 in *Séminaire Bourbaki*, vol. 2015/2016, Astérisque **390**, Soc. Math. France, Paris, 2017. MR Zbl

[Szemerédi 1975] E. Szemerédi, "On sets of integers containing no k elements in arithmetic progression", *Acta Arith.* **27** (1975), 199–245. MR Zbl

[Tao 2008] T. Tao, "Product set estimates for non-commutative groups", *Combinatorica* **28**:5 (2008), 547–594. MR Zbl

[Tao 2013] T. Tao, "Mixing for progressions in nonabelian groups", *Forum Math. Sigma* **1** (2013), art. id. e2. MR Zbl

[Terry and Wolf 2019] C. Terry and J. Wolf, "Stable arithmetic regularity in the finite field model", *Bull. Lond. Math. Soc.* **51**:1 (2019), 70–88. MR Zbl

Received 17 Jan 2022. Revised 6 Feb 2024.

AMADOR MARTIN-PIZARRO:

pizarro@math.uni-freiburg.de
Abteilung für Mathematische Logik – Mathematisches Institut, Albert-Ludwigs-Universität Freiburg, Freiburg, Germany

DANIEL PALACÍN:

dpalacin@ucm.es
Departamento de Álgebra, Geometría y Topología, Facultad de Ciencias Matemáticas, Universidad Complutense de Madrid, Madrid, Spain

Model Theory
Vol. 3, No. 1, 2024
https://doi.org/10.2140/mt.2024.3.39

Bounded ultraimaginary independence and its total Morley sequences

James E. Hanson

We investigate the following model-theoretic independence relation: $b \mathop{\smile\!\!\!\!\!\!\;\mid}^{\mathrm{bu}}_A c$ if and only if $\mathrm{bdd}^{\mathrm{u}}(Ab) \cap \mathrm{bdd}^{\mathrm{u}}(Ac) = \mathrm{bdd}^{\mathrm{u}}(A)$, where $\mathrm{bdd}^{\mathrm{u}}(X)$ is the class of all ultraimaginaries bounded over X. In particular, we sharpen a result of Wagner to show that $b \mathop{\smile\!\!\!\!\!\!\;\mid}^{\mathrm{bu}}_A c$ if and only if $\langle \mathrm{Autf}(\mathbb{M}/Ab) \cup \mathrm{Autf}(\mathbb{M}/Ac) \rangle = \mathrm{Autf}(\mathbb{M}/A)$, and we establish full existence over hyperimaginary parameters (i.e., for any set of hyperimaginaries A and ultraimaginaries b and c, there is a $b' \equiv_A b$ such that $b' \mathop{\smile\!\!\!\!\!\!\;\mid}^{\mathrm{bu}}_A c$). Extension then follows as an immediate corollary.

We also study *total* $\mathop{\smile\!\!\!\!\!\!\;\mid}^{\mathrm{bu}}$*-Morley sequences* (i.e., A-indiscernible sequences I satisfying $J \mathop{\smile\!\!\!\!\!\!\;\mid}^{\mathrm{bu}}_A K$ for any J and K with $J + K \equiv^{\mathrm{EM}}_A I$), and we prove that an A-indiscernible sequence I is a total $\mathop{\smile\!\!\!\!\!\!\;\mid}^{\mathrm{bu}}$-Morley sequence over A if and only if whenever I and I' have the same Lascar strong type over A, I and I' are related by the transitive, symmetric closure of the relation "$J + K$ is A-indiscernible". This is also equivalent to I being "based on" A in a sense defined by Shelah (1980) in his study of simple unstable theories.

Finally, we show that for any A and b in any theory T, if there is an Erdős cardinal $\kappa(\alpha)$ with $|Ab| + |T| < \kappa(\alpha)$, then there is a total $\mathop{\smile\!\!\!\!\!\!\;\mid}^{\mathrm{bu}}$-Morley sequence $(b_i)_{i<\omega}$ over A with $b_0 = b$.

Introduction

A central theme in neostability theory is the importance of various kinds of "generic" indiscernible sequences — usually with Michael Morley's name attached to them — such as Morley sequences in stable and simple theories, strict Morley sequences in NIP and NTP_2 theories, tree Morley sequences in NSOP_1 theories, and $\mathop{\smile\!\!\!\!\!\!\;\mid}^{þ}$-Morley sequences in rosy theories. A very broad question one might ask is this: How generically can we build indiscernible sequences in *arbitrary* theories?

Over a model M, we can always extend a given type $p(x) \in S_x(M)$ to a global M-invariant type $q(x) \supset p(x)$ and then use this to generate a sequence $(b_i)_{i<\omega}$ satisfying $b_i \models q \restriction Mb_{<i}$ for each $i < \omega$. In some cases the particular choice of $q(x)$ matters, but typically these sequences are robustly generic. Sequences produced in

MSC2020: 03C45.
Keywords: ultraimaginaries, bounded closure, total Morley sequences.

this way have a certain property, which is that they are *based on M in the sense of Simon*; i.e., for any I and J with $I \equiv_M J \equiv_M b_{<\omega}$, there is a K such that $I+K$ and $J+K$ are both M-indiscernible. In NIP theories, the sequences with this property are precisely the sequences generated by an invariant type [Simon 2015, Proposition 2.38]. Over an arbitrary set of parameters A, however, there may fail to be any indiscernible sequences based on A. In the dense circular order, for instance, there are no indiscernible sequences based on $\varnothing$. Other technical issues also arise when working over arbitrary sets, such as the necessity of considering Lascar strong types over and above ordinary types.

A notion of independence $\mathop{\smile\!\!\!\!\!|}^*$ is said to satisfy *full existence* if for any A, b, and c, there is a $b' \equiv_A b$ such that $b' \mathop{\smile\!\!\!\!\!|}^*_A c$. Together with a common model-theoretic application of the Erdős-Rado theorem (Fact 1.2), this implies that for any A and b, one can build an $\mathop{\smile\!\!\!\!\!|}^*$*-Morley sequence*, an A-indiscernible sequence $(b_i)_{i<\omega}$ with $b_0 = b$ satisfying $b_i \mathop{\smile\!\!\!\!\!|}^*_A b_{<i}$ for each $i < \omega$ (assuming $\mathop{\smile\!\!\!\!\!|}^*$ also satisfies right monotonicity). Model-theoretically tame theories often have full existence for powerful independence notions, such as nonforking, but this does fail in some notable tame contexts.

One independence notion that is known to satisfy full existence in arbitrary theories is that of *algebraic independence* [Adler 2009, Proposition 1.5]: $b \mathop{\smile\!\!\!\!\!|}^{\mathrm{a}}_A c$ if $\mathrm{acl}^{\mathrm{eq}}(Ab) \cap \mathrm{acl}^{\mathrm{eq}}(Ac) = \mathrm{acl}^{\mathrm{eq}}(A)$. A natural modification of this concept is *bounded hyperimaginary independence*: $b \mathop{\smile\!\!\!\!\!|}^{\mathrm{b}}_A c$ if $\mathrm{bdd}^{\mathrm{heq}}(Ab) \cap \mathrm{bdd}^{\mathrm{heq}}(Ac) = \mathrm{bdd}^{\mathrm{heq}}(A)$. Despite perhaps sounding like an intro-to-model-theory exercise, the combinatorics necessary to prove full existence for $\mathop{\smile\!\!\!\!\!|}^{\mathrm{a}}$ are somewhat subtle. It was established in [Conant and Hanson 2022] that $\mathop{\smile\!\!\!\!\!|}^{\mathrm{a}}$ satisfies full existence in continuous logic and, relatedly, that $\mathop{\smile\!\!\!\!\!|}^{\mathrm{b}}$ satisfies full existence in discrete (and continuous) logic, answering a question of Adler [2005, Question A.8]. While the relations of $\mathop{\smile\!\!\!\!\!|}^{\mathrm{a}}$ and $\mathop{\smile\!\!\!\!\!|}^{\mathrm{b}}$ are algebraically nice,[1] they seem to lack semantic consequences outside of certain special theories (such as those with a canonical independence relation in the sense of Adler [2005, Lemma 3.2]).

While being able to build $\mathop{\smile\!\!\!\!\!|}^*$-Morley sequences is certainly good, in many applications the important property is really that of being a *total* $\mathop{\smile\!\!\!\!\!|}^*$*-Morley sequence*,[2] which is an A-indiscernible sequence satisfying $b_{\geq i} \mathop{\smile\!\!\!\!\!|}^*_A b_{<i}$ for every $i<\omega$. When $\mathop{\smile\!\!\!\!\!|}^*$ lacks the algebraic properties necessary to imply that all $\mathop{\smile\!\!\!\!\!|}^*$-Morley sequences are total $\mathop{\smile\!\!\!\!\!|}^*$-Morley sequences, it can in general be difficult to ensure their existence. Total $\mathop{\smile\!\!\!\!\!|}^{\mathrm{a}}$-Morley sequences arise in Adler's characterization of canonical independence relations. And building total $\mathop{\smile\!\!\!\!\!|}^{\mathrm{K}}$-Morley sequences, where $\mathop{\smile\!\!\!\!\!|}^{\mathrm{K}}$ is

[1] In the sense of the algebra of an independence relation, not the sense of the algebra in "algebraic closure".

[2] This use of the term "total" in the context of Morley sequences was originally introduced in [Kaplan and Ramsey 2020].

the relation of non-Kim-forking, is a crucial technical step in Kaplan and Ramsey's proofs [2020] of the symmetry of Kim-forking and the independence theorem in NSOP_1 theories.

In simple theories, Morley sequences over A are not generally based on A in the sense of Simon. They do however nearly satisfy this property. If I and J are Morley sequences over A with $I \equiv^{\mathrm{L}}_A J$,[3] then there are I' and K such that $I+I'$, $I'+K$, and $J+K$ are A-indiscernible. In an NSOP_1 theory T, if I is a tree Morley sequence over $M \models T$ and $J \equiv_M I$, then we can find K_0, K_1, and K_2 such that $I+K_0$, K_1+K_0, K_1+K_2, and $J+K_2$ are all M-indiscernible (see Proposition 4.28). These facts suggest the consideration of the following equivalence relation, originally introduced in [Shelah 1980, Definition 5.1]: Let $\approx_A$ be the transitive, symmetric closure of the relation "$I+J$ is A-indiscernible". The intuition is that what it means for an A-indiscernible sequence I to be "based on A" is that there are few $\approx_A$-classes among the realizations of $\mathrm{tp}(I/A)$. We say that I is *based on A in the sense of Shelah* if there does not exist a sequence $(I_i)_{i<\kappa}$ (with κ large) such that $I_i \equiv_A I$ for each $i<\kappa$ and $I_i \not\approx_A I_j$ for each $i<j<\kappa$. A simple compactness argument shows that I is based on A in the sense of Shelah if and only if the set of realizations of $\mathrm{tp}(I/A)$ decomposes into a bounded number of $\approx_A$-classes. Buechler [1997, Definition 2.4][4] used this relation to define a notion of canonical base. He focused on $\varnothing$-indiscernible sequences and gave the following definition: A is a *canonical base* of the $\varnothing$-indiscernible sequence I if any automorphism $\sigma \in \mathrm{Aut}(\mathbb{M})$ fixes A pointwise if and only if it fixes the $\approx_\varnothing$-class of I. One difficulty with this concept, of course, is that not all indiscernible sequences have canonical bases in this sense (even in T^{eq}, e.g., [Adler 2005, Example 3.13]).

Two of the problems we have mentioned — the lack of canonical bases for indiscernible sequences and the lack of semantic consequences of $\mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{a}}$ and $\mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{b}}$ — can both be solved by an extremely blunt move: the introduction of ultraimaginary parameters. An *ultraimaginary* is an equivalence class of an arbitrary invariant equivalence relation (as opposed to a type-definable equivalence relation, as in the definition of hyperimaginaries). Every indiscernible sequence I trivially has an ultraimaginary canonical base in the sense of Buechler, i.e., the $\approx_\varnothing$-class of I itself.

Another appealing aspect of ultraimaginaries is that they characterize Lascar strong type in the same way that hyperimaginaries characterize Kim–Pillay strong type. An ultraimaginary $[b]_E$ is *bounded over A* if it has boundedly many conjugates under $\mathrm{Aut}(\mathbb{M}/A)$. We will write $\mathrm{bdd}^{\mathrm{u}}(A)$ for the class of ultraimaginaries bounded over A. In general, it turns out that b and c have the same Lascar strong type

[3]The equivalence relation $\equiv^{\mathrm{L}}_A$ is the transitive closure of the relation "there is a model $M \supseteq A$ such that $b \equiv_M c$". If $b \equiv^{\mathrm{L}}_A c$, we say that b and c have the same *Lascar strong type* over A.

[4]This preprint is difficult to track down. The relevant ideas are developed further in [Adler 2005, Section 3.2], which is easily available.

over A if and only if they "have the same type over $\mathrm{bdd}^{\mathrm{u}}(A)$", once this concept is defined precisely.

Pure analogical thinking might lead one to consider the following independence notion: $b \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A c$ if $\mathrm{bdd}^{\mathrm{u}}(Ab) \cap \mathrm{bdd}^{\mathrm{u}}(Ac) = \mathrm{bdd}^{\mathrm{u}}(A)$. This notion is implicit in a result of Wagner [2015, Proposition 2.12], which we restate and expand slightly (Proposition 2.4): $b \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A c$ if and only if $\langle \mathrm{Autf}(\mathbb{M}/Ab) \cup \mathrm{Autf}(\mathbb{M}/Ac)\rangle = \mathrm{Autf}(\mathbb{M}/A)$ (where $\langle X\rangle$ is the group generated by X). This characterization is clearly semantically meaningful, and moreover it allows one to discuss $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$ without actually mentioning ultraimaginaries at all. One way to see why this equivalence works is the fact that ultraimaginaries are "dual" to co-small sets of automorphisms; a group $G \leq \mathrm{Aut}(\mathbb{M})$ is *co-small* if there is a small model M such that $\mathrm{Aut}(\mathbb{M}/M) \leq G$. For every co-small group G, there is an ultraimaginary a_E such that $\mathrm{Aut}(\mathbb{M}/a_E) = G$ (Proposition 1.7).

As $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$ lacks finite character, *total* $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$*-Morley sequences over* A seem to be correctly defined as A-indiscernible sequences $(b_i)_{i<\omega}$ with the property that for any $I + J \equiv^{\mathrm{EM}}_A b_{<\omega}$,[5] we have that $I \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A J$. The automorphism group characterization of $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$, together with its the nice algebraic properties and the malleability of indiscernible sequences, leads to a pleasing characterization of total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences over sets of hyperimaginary parameters (Theorem 4.8), the equivalence of the following.

- $(b_i)_{i<\omega}$ is a total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequence over A.
- For some infinite I and J, we have that $I + J \equiv^{\mathrm{EM}}_A b_{<\omega}$ and $I \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A J$.
- For any I, $I \approx_A b_{<\omega}$ if and only if there is $I' \equiv^{\mathrm{L}}_A I$ such that $b_{<\omega} + I'$ is A-indiscernible.
- $b_{<\omega}$ is based on A in the sense of Shelah; i.e., $[b_{<\omega}]_{\approx_A} \in \mathrm{bdd}^{\mathrm{u}}(A)$.

The condition in the third bullet point is a natural mutual generalization of Lascar strong type and Ehrenfeucht–Mostowski type (Definition 4.5). Theorem 4.8 also tells us that when total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences exist, they act as particularly uniform witnesses of Lascar strong type (Proposition 4.3).

Of course this all leaves two critical questions: Does $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$ always satisfy full existence? And, even if it does, can we actually build total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences in any type over any set under any theory? The bluntness of ultraimaginaries leaves us without one of the most important tools in model theory, compactness. Furthermore, $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$'s lack of finite character gives us less leeway in applying the Erdős-Rado theorem to construct indiscernible sequences with certain properties;

[5] $I \equiv^{\mathrm{EM}}_A J$ means that I and J have the same *Ehrenfeucht–Mostowski type* over A (i.e., for any increasing tuples $\bar{b} \in I$ and $\bar{c} \in J$ of the same length, $\bar{b} \equiv_A \bar{c}$). Note that I and J do not need to have the same order type.

we now need to be more concerned with the particular order types of the sequences involved.

Using some of the indiscernible tree technology from [Kaplan and Ramsey 2020], we are able to prove that $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$ does satisfy full existence over arbitrary sets of (hyperimaginary) parameters in arbitrary (discrete or continuous) theories (Theorem 3.6).[6] With regards to building total $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley sequences, Theorem 4.8 tells us that we don't need to worry *too* much about order types. All we need to get a total $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley sequence over A is an A-indiscernible sequence $(b_i)_{i<\omega+\omega}$ with $b_{\geq\omega} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A b_{<\omega}$. This is fortunate because constructing ill-ordered $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley sequences directly seems daunting. Unfortunately, $\omega+\omega$ appears to be about one ω further than we can go without a large cardinal. What we do get is this (Theorem 4.22): For any A and b in any theory T, if there is an Erdős cardinal $\kappa(\alpha)$ with $|Ab|+|T| < \kappa(\alpha)$ (for any $\alpha \geq \omega$), then there is a total $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley sequence $(b_i)_{i<\omega}$ over A with $b_0 = b$. Without a large cardinal, the best we seem to be able to do (Proposition 4.17) is a half-infinite, half-arbitrary-finite approximation of a total $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley sequence, which we call a *weakly total* $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$*-Morley sequence*. These sequences also serve as uniform witnesses of Lascar strong type without any set-theoretic hypotheses (Corollary 4.18).

1. Ultraimaginaries

Here we will set definitions and conventions, and we also take the opportunity to collect some basic facts about ultraimaginaries which are likely folklore, although we could not find explicit references.

Fix a theory T and a set-sized monster model $\mathbb{M} \models T$.

Definition 1.1. An *invariant equivalence relation of arity* κ is an equivalence relation E on $\mathbb{M}^x$ (with $|x| = \kappa$) such that for any $a, b, c, d \in \mathbb{M}^x$ with $ab \equiv cd$, aEb if and only if cEd.

An *ultraimaginary of arity* κ is a pair (E, a_E) consisting of an invariant equivalence relation E (of arity κ) and an E-equivalence class a_E of some tuple $a \in \mathbb{M}^x$. By an abuse of notation, we will write a_E for the pair (E, a_E), and we may also write $[a]_E$ if necessary for notational clarity.

Given an ultraimaginary a_E, $\mathrm{Aut}(\mathbb{M}/a_E)$ is the set of automorphisms $\sigma \in \mathrm{Aut}(\mathbb{M})$ with the property that $aE(\sigma \cdot a)$. We write $\mathrm{Autf}(\mathbb{M}/a_E)$ for the group generated by $\{\sigma \in \mathrm{Aut}(\mathbb{M}/M) : M \preceq \mathbb{M},\ \mathrm{Aut}(\mathbb{M}/M) \leq \mathrm{Aut}(\mathbb{M}/a_E)\}$.

We say that b_F is *definable over* a_E if b_F is fixed by every automorphism in $\mathrm{Aut}(\mathbb{M}/a_E)$. We write $\mathrm{dcl}^{\mathrm{u}}(a_E)$ for the class of all ultraimaginaries definable

[6] Although this result partially supersedes a result in [Conant and Hanson 2022] (full existence for $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{a}}$ in continuous logic and $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{b}}$ in discrete or continuous logic), the proof there gives more detailed numerical information which may be especially useful in the metric context.

over a_E. For any κ, we write $\mathrm{dcl}^{\mathrm{u}}_{\kappa}(a_E)$ for the set of elements of $\mathrm{dcl}^{\mathrm{u}}(a_E)$ of arity at most κ. We say that b_F and c_G are *interdefinable over* a_E if $b_F \in \mathrm{dcl}^{\mathrm{u}}(a_E c_G)$ and $c_G \in \mathrm{dcl}^{\mathrm{u}}(a_E b_F)$.

We say that b_F is *bounded over* a_E if the $\mathrm{Aut}(\mathbb{M}/a_E)$-orbit of b_F is bounded.[7] We write $\mathrm{bdd}^{\mathrm{u}}(a_E)$ for the class of all ultraimaginaries bounded over a_E. We write $\mathrm{bdd}^{\mathrm{u}}_{\kappa}(a_E)$ for the set of elements of $\mathrm{bdd}^{\mathrm{u}}(a_E)$ of arity at most κ. We say that b_F and c_G are *interbounded over* a_E if $b_F \in \mathrm{bdd}^{\mathrm{u}}(a_E c_G)$ and $c_G \in \mathrm{bdd}^{\mathrm{u}}(a_E b_F)$.

We write $a_E \equiv b_E$ to mean that there is an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M})$ with $\sigma \cdot a_E = b_E$. We write $b_F \equiv_{a_E} c_F$ to mean that $a_E b_F \equiv a_E c_F$ (i.e., there is $\sigma \in \mathrm{Aut}(\mathbb{M}/a_E)$ such that $\sigma \cdot b_F = c_F$).

Note that real elements, imaginaries, and hyperimaginaries can all be regarded as ultraimaginaries.

An easy counting argument shows that $\mathrm{bdd}^{\mathrm{u}}$ is a closure operator (i.e., for any a_E, b_F, and c_G, if $b_F \in \mathrm{bdd}^{\mathrm{u}}(a_E)$ and $c_G \in \mathrm{bdd}^{\mathrm{u}}(b_F)$, then $c_G \in \mathrm{bdd}^{\mathrm{u}}(a_E)$).

We will also sometimes define an invariant equivalence relation E on the realizations of a single type $p(x)$ over $\varnothing$. Equivalence classes of such can be thought of as ultraimaginaries by using the same trick that is commonly used with hyperimaginaries: Consider the invariant equivalence relation $E'(x, y)$ defined by $x = y \vee (E(x, y) \wedge x \models p \wedge y \models p)$.

For the sake of clarity, we will reserve the notation a_E for ultraimaginaries and write hyperimaginaries in the same way we write real elements. For the sake of cardinality issues, we will also take all hyperimaginaries to be quotients of countable tuples by countably type-definable equivalence relations. It is a standard fact that every hyperimaginary is interdefinable with some set of hyperimaginaries of this form.

Fact 1.2 [Shelah 1980].[8] *Let $(b_i)_{i<\lambda}$ be a sequence of tuples with $|b_i| < \kappa$ and let A be some set of parameters. If $\lambda \geq \beth_{(2^{\kappa+|A|+|T|})^+}$, then there is an A-indiscernible sequence $(b'_i)_{i<\omega}$ such that for every $n < \omega$, there are $i_0 < \cdots < i_n < \kappa$ such that $b'_0 \dots b'_n \equiv_A b_{i_0} \dots b_{i_n}$.*

Lemma 1.3. *Let M be a model. If $a_E \in \mathrm{bdd}^{\mathrm{u}}(M)$, then $a_E \in \mathrm{dcl}^{\mathrm{u}}(M)$.*

Proof. Assume that $a_E \notin \mathrm{dcl}^{\mathrm{u}}(M)$. Let $p(x)$ be a global M-invariant type extending $\mathrm{tp}(a/M)$. Assume that there are a_0 and a_1 realizing $\mathrm{tp}(a/M)$ such that $a_0 \not{E} a_1$. For any $i > 1$, given $a_{<i}$, let $a_i \models p{\upharpoonright}Ma_{<i}$. Since $a_i a_j \equiv_M a_i a_k$ for any $j, k < i$, we must have that $a_i \not{E} a_j$ for any $j < i$. Since we can do this indefinitely, we have that a_E is not bounded over M. □

[7]Specifically, by Proposition 1.4, this is equivalent to b_F having at most $2^{|ab|+|T|}$ conjugates over a_E.

[8]See [Tent and Ziegler 2012, Lemma 7.2.12] for a modern presentation of the result.

Proposition 1.4. *For any ultraimaginaries a_E and b_F, the following are equivalent.*

(1) $b_F \notin \mathrm{bdd}^{\mathrm{u}}(a_E)$.

(2) *There is an a-indiscernible sequence $(c_i)_{i<\omega}$ such that $c_0 \equiv_{a_E} b$ and $c_i \not F c_j$ for each $i < j < \omega$.*

(3) $|\mathrm{Aut}(\mathbb{M}/a_E) \cdot b_F| > 2^{|ab|+|T|}$.

Proof. (3)⇒(2). Let $(b^i_F)_{i<(2^{|ab|+|T|})^+}$ be an enumeration of $\mathrm{Aut}(\mathbb{M}/a_E) \cdot b_F$. Let $M \supseteq a$ be a model with $|M| \leq |a| + |T|$. Let x be a tuple of variables of the same length as b. There are at most $2^{|ab|+|T|}$ types in $S_x(M)$. Therefore, there must be $i < j < (2^{|ab|+|T|})^+$ such that $b^i \equiv_M b^j$. Let $p(x)$ be a global M-invariant type extending $\mathrm{tp}(b^i/M)$, and let $(c_i)_{i<\omega}$ be a Morley sequence generated by $p(x)$ over Mb^ib^j. Since $b^i \not F b^j$, we must have that $c_0 \not F b^i$. Therefore $c_i \not F c_j$ for any $i < j < \lambda$, and so $(c_i)_{i<\omega}$ is the required a-indiscernible sequence.

(2)⇒(1). Given an a-indiscernible sequence $(c_i)_{i<\omega}$ as in the statement of the proposition, we can extend it to an a-indiscernible sequence $(c_i)_{i<\lambda}$ for any λ. These sequences will still satisfy that $c_i \not F c_j$ for any $i < j < \lambda$, so b_F has an unbounded number of $\mathrm{Aut}(\mathbb{M}/a_E)$-conjugates and $b_F \notin \mathrm{bdd}^{\mathrm{u}}(a_E)$.

(1)⇒(3). This is immediate from the definition of $\mathrm{bdd}^{\mathrm{u}}(a_E)$. □

Corollary 1.5. *For any λ, $\mathrm{bdd}^{\mathrm{u}}_\lambda(a_E)$ has cardinality at most $2^{|a|+2^{\lambda+|T|}}$.*

Proof. For each $\alpha \leq \lambda$, $|S_{\alpha+\alpha}(T)| \leq 2^{\lambda+|T|}$. Since an invariant equivalence relation on α-tuples is specified by a subset of $S_{\alpha+\alpha}(T)$, this implies that for each $\alpha \leq \lambda$, there are at most $2^{2^{\lambda+|T|}}$ invariant equivalence relations on α-tuples. Therefore the total number of invariant equivalence relations on tuples of length at most λ is $\lambda \cdot 2^{2^{\lambda+|T|}} = 2^{2^{\lambda+|T|}}$. For each such F, the set $\{b_F : b_F \in \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E)\}$ has cardinality at most $2^{|a|+\lambda+|T|}$ by Proposition 1.4. Finally, $2^{2^{\lambda+|T|}} \cdot 2^{|a|+\lambda+|T|} = 2^{|a|+2^{\lambda+|T|}}$. □

Co-small groups of automorphisms. Here we will see that ultraimaginaries are essentially the same thing as reasonable subgroups of $\mathrm{Aut}(\mathbb{M})$.

Definition 1.6. A group $G \leq \mathrm{Aut}(\mathbb{M})$ is *co-small* if there is a small model M such that $\mathrm{Aut}(\mathbb{M}/M) \leq G$.

Clearly for any ultraimaginary a_E, $\mathrm{Aut}(\mathbb{M}/a_E)$ is co-small. The converse is true as well.

Proposition 1.7. *For any co-small G, if $\mathrm{Aut}(\mathbb{M}/M) \leq G$, then there is an ultraimaginary a_E such that $G = \mathrm{Aut}(\mathbb{M}/a_E)$ where a is some enumeration of M.*

Proof. Let M be a small model witnessing that G is co-small. Consider the binary relation defined on realizations of $\mathrm{tp}(M)$ (in some fixed enumeration) defined by $E(M_0, M_1)$ if and only if there is $\sigma \in \mathrm{Aut}(\mathbb{M})$ and $\tau \in G$ such that $\sigma \cdot M = M_0$

and $\sigma\tau \cdot M = M_1$. We need to verify that E is an invariant equivalence relation. Reflexivity is obvious.

Invariance. Suppose that $E(M_0, M_1)$, as witnessed by $\sigma \in \mathrm{Aut}(\mathbb{M})$ and $\tau \in G$. Fix $\sigma' \in \mathrm{Aut}(\mathbb{M})$. We then have that $\sigma'\sigma \cdot M = \sigma' \cdot M_0$ and $\sigma'\sigma\tau \cdot M = \sigma' \cdot M_1$, whence $E(\sigma' \cdot M_0, \sigma' \cdot M_1)$.

Symmetry. If $\sigma \cdot M = M_0$ and $\sigma\tau \cdot M = M_1$ with $\sigma \in \mathrm{Aut}(\mathbb{M})$ and $\tau \in G$, then $\sigma\tau\tau^{-1} \cdot M = M_0$ and $\sigma\tau \cdot M = M_1$. We have $\sigma\tau \in \mathrm{Aut}(\mathbb{M})$ and $\tau^{-1} \in G$, so $E(M_1, M_0)$.

Transitivity. Suppose that for $\sigma, \sigma' \in \mathrm{Aut}(\mathbb{M})$ and $\tau, \tau' \in G$, we have that $\sigma \cdot M = M_0$, $\sigma\tau \cdot M = \sigma' \cdot M = M_1$, and $\sigma'\tau' \cdot M = M_2$. This implies that $(\sigma\tau)^{-1}\sigma' = \tau^{-1}\sigma^{-1}\sigma' \in \mathrm{Aut}(\mathbb{M}/M) \leq G$. Since $\tau \in G$ as well, we have that $\sigma^{-1}\sigma' \in G$. Therefore $\sigma^{-1}\sigma'\tau' \in G$. Finally, $\sigma\sigma^{-1}\sigma'\tau' \cdot M = M_2$, so $E(M_0, M_2)$.

Consider the ultraimaginary M_E. For any $\tau \in G$, we clearly have $E(M, \tau \cdot M)$, so $G \leq \mathrm{Aut}(\mathbb{M}/M_E)$. Conversely, suppose that $\alpha \in \mathrm{Aut}(\mathbb{M}/M_E)$. By definition, this implies that $E(M, \alpha \cdot M)$, so there are $\sigma \in \mathrm{Aut}(\mathbb{M})$ and $\tau \in G$ such that $\sigma \cdot M = M$ and $\sigma\tau \cdot M = \alpha \cdot M$. Therefore $\sigma, \tau^{-1}\sigma^{-1}\alpha \in \mathrm{Aut}(\mathbb{M}/M) \leq G$. Since $\tau^{-1} \in G$, we therefore have that $\alpha \in G$. ☐

Corollary 1.8. *If $b_F \in \mathrm{bdd}^{\mathrm{u}}(a_E)$, then there is $c_G \in \mathrm{bdd}^{\mathrm{u}}(a_E)$ of arity at most $|a| + |T|$ such that b_F and c_G are interdefinable over $\varnothing$. Furthermore, c can be taken to be an enumeration of any model of size at most $|a| + |T|$ containing a.*

Proof. There is a model $M \supseteq a$ with $|M| \leq |a| + |T|$. By Lemma 1.3, we have that $\mathrm{Aut}(\mathbb{M}/M) \leq \mathrm{Aut}(\mathbb{M}/b_F)$, so by Proposition 1.7, we have that there is c_G with arity at most $|a| + |T|$ which satisfies that $\mathrm{Aut}(\mathbb{M}/c_G) = \mathrm{Aut}(\mathbb{M}/b_F)$ (i.e., c_G and b_F are interdefinable over $\varnothing$). Furthermore, we can take c to be an enumeration of M. ☐

Definition 1.9. For any co-small group G, we write $[\![G]\!]$ for some arbitrary ultraimaginary a_E of minimal arity satisfying $G = \mathrm{Aut}(\mathbb{M}/a_E)$. We will write $\mathrm{dcl}^{\mathrm{u}}[\![G]\!]$ and $\mathrm{dcl}^{\mathrm{u}}_{\lambda}[\![G]\!]$ for $\mathrm{dcl}^{\mathrm{u}}([\![G]\!])$ and $\mathrm{dcl}^{\mathrm{u}}_{\lambda}([\![G]\!])$ and likewise with $\mathrm{bdd}^{\mathrm{u}}$. (Note that $\mathrm{dcl}^{\mathrm{u}}[\![G]\!]$ and $\mathrm{bdd}^{\mathrm{u}}[\![G]\!]$ only depend on G, not on the particular choice of $[\![G]\!]$.)

It is immediate from Proposition 1.7 that for any co-small G and H, $[\![G]\!] \in \mathrm{dcl}^{\mathrm{u}}[\![H]\!]$ if and only if $G \geq H$. A similar statement for $\mathrm{bdd}^{\mathrm{u}}$ is given in Proposition 1.12.

Now we can see that intersections of $\mathrm{dcl}^{\mathrm{u}}$-closed sets (and therefore also $\mathrm{bdd}^{\mathrm{u}}$-closed sets) have semantic significance in arbitrary theories, in that intersections correspond to joins in the lattice of co-small groups of automorphisms.

Proposition 1.10. *For any a_E, b_F, c_G, and c'_G, the following are equivalent.*

(1) $c_G \equiv_{\mathrm{dcl}^{\mathrm{u}}_{\lambda}(a_E) \cap \mathrm{dcl}^{\mathrm{u}}_{\lambda}(b_F)} c'_G$ *for all* λ.

(2) *There is* $\sigma \in \langle \mathrm{Aut}(\mathbb{M}/a_E) \cup \mathrm{Aut}(\mathbb{M}/b_F) \rangle$ *such that* $\sigma \cdot c_G = c'_G$.

(3) *There is a sequence* $(a^i b^i c^i)_{i \leq n}$ *such that* $a^0 = a$, $b^0 = b$, $c^0 = c$, $c^n_G = c'_G$, *and for each* $i < n$,
   - *if* i *is even, then* $a^i = a^{i+1}$ *and* $b^i_F c^i_G \equiv_{a^i_E} b^{i+1}_F c^{i+1}_G$ *and*
   - *if* i *is odd, then* $b^i = b^{i+1}$ *and* $a^i_E c^i_G \equiv_{b^i_F} a^{i+1}_E c^{i+1}_G$.

Proof. Let $H = \langle \mathrm{Aut}(\mathbb{M}/a_E) \cup \mathrm{Aut}(\mathbb{M}/b_F) \rangle$.

Claim. $\mathrm{dcl}^{\mathrm{u}}_\lambda(a_E) \cap \mathrm{dcl}^{\mathrm{u}}_\lambda(b_F)$ and $[\![H]\!]$ are interdefinable (i.e., $\mathrm{dcl}^{\mathrm{u}}_\lambda(a_E) \cap \mathrm{dcl}^{\mathrm{u}}_\lambda(b_F) \subseteq \mathrm{dcl}^{\mathrm{u}}([\![H]\!])$ and $[\![H]\!] \in \mathrm{dcl}^{\mathrm{u}}(\mathrm{dcl}^{\mathrm{u}}_\lambda(a_E) \cap \mathrm{dcl}^{\mathrm{u}}_\lambda(b_F)))$ for all sufficiently large λ.

Proof of claim. Clearly $[\![H]\!] \in \mathrm{dcl}^{\mathrm{u}}(a_E) \cap \mathrm{dcl}^{\mathrm{u}}(b_F)$, so $[\![H]\!] \in \mathrm{dcl}^{\mathrm{u}}_\lambda(a_E) \cap \mathrm{dcl}^{\mathrm{u}}_\lambda(b_F)$ for all sufficiently large λ.

Conversely, suppose that $d_I \in \mathrm{dcl}^{\mathrm{u}}(a_E) \cap \mathrm{dcl}^{\mathrm{u}}(b_F)$. Any $\sigma \in H$ is a product of elements of $\mathrm{Aut}(\mathbb{M}/a_E)$ and $\mathrm{Aut}(\mathbb{M}/b_F)$, so it must fix d_I. Thus $\mathrm{Aut}(\mathbb{M}/d_I) \geq H$ and hence $d_I \in \mathrm{dcl}^{\mathrm{u}}[\![H]\!]$. $\square_{\text{claim}}$

So now we have that $c_G \equiv_{\mathrm{dcl}^{\mathrm{u}}_\lambda(a_E) \cap \mathrm{dcl}^{\mathrm{u}}_\lambda(b_F)} c'_G$ holds for sufficiently large λ if and only if $c_G \equiv_{[\![H]\!]} c'_G$. Also note that $c_G \equiv_{\mathrm{dcl}^{\mathrm{u}}_\lambda(a_E) \cap \mathrm{dcl}^{\mathrm{u}}_\lambda(b_F)} c'_G$ for sufficiently large λ and only if the same holds for any λ. Therefore (1) and (2) are equivalent.

There is a $\sigma \in H$ with $\sigma \cdot c_G = c'_G$ if and only if there are $\alpha_0, \ldots, \alpha_{n-1} \in \mathrm{Aut}(\mathbb{M}/a_E)$ and $\beta_0, \ldots, \beta_{n-1} \in \mathrm{Aut}(\mathbb{M}/b_F)$ such that $\sigma = \alpha_{n-1}\beta_{n-1} \ldots \beta_1 \alpha_0 \beta_0$.

For (2)⇒(3), assume that there are such $\bar{\alpha}$ and $\bar{\beta}$ for which

$$\alpha_{n-1}\beta_{n-1}\alpha_{n-2} \ldots \beta_1 \alpha_0 \beta_0 \cdot c_G = c'_G.$$

Let $a^0 b^0 c^0 = abc$, $a^1 b^1 c^1 = \alpha_{n-1} \cdot (a^0 b^0 c^0)$, $a^2 b^2 c^2 = \alpha_{n-1}\beta_{n-1} \cdot (a^0 b^0 c^0)$, and so on up to $a^{2n} b^{2n} c^{2n} = \alpha_{n-1}\beta_{n-1}\alpha_{n-2} \ldots \beta_1 \alpha_0 \beta_0 \cdot (a^0 b^0 c^0)$. Clearly we have that $c^{2n}_G = c'_G$, so we just need to verify that $(a^i b^i c^i)_{i \leq 2n}$ is the required sequence. If $i < 2n$ is even, then $\alpha_i \in \mathrm{Aut}(\mathbb{M}/a_E)$, so $a^i_E = a^{i+1}_E$. Furthermore, $b^0_F c^0_G \equiv_{a^0_E} \alpha_i \cdot (b^0_F c^0_G)$, so by invariance,

$$\alpha_{n-1}\beta_{n-1} \ldots \beta_{i+1} \cdot (b^0_F c^0_G) \equiv_{\alpha_{n-1}\beta_{n-1} \ldots \beta_{i+1} \cdot a^0_E} \alpha_{n-1}\beta_{n-1} \ldots \beta_{i+1} \alpha_i \cdot (b^0_F c^0_G),$$

which is the same as $b^i_F c^i_G \equiv_{a^i_E} b^{i+1}_F c^{i+1}_G$. If $i < 2n$ is odd, then the same argument tells us that $b^i_F = b^{i+1}_F$ and $a^i_E c^i_G \equiv_{b^i_F} a^{i+1}_E c^{i+1}_G$.

For (3)⇒(2), the above argument is reversible. Fix $(a^i_E b^i_F c^i_G)_{i \leq 2n}$ satisfying the conditions of (3). First of all we can find $\alpha_{n-1} \in \mathrm{Aut}(\mathbb{M}/a_E)$ such that $\alpha^{-1}_{n-1} \cdot (a^1_E b^1_F c^1_G) = a^0_E b^0_F c^0_G$. Then we can find $\beta_{n-1} \in \mathrm{Aut}(\mathbb{M}/b_F)$ such that $\beta^{-1}_{n-1}\alpha^{-1}_{n-1} \cdot (a^2_E b^2_F c^2_G) = a^0_E b^0_F c^0_G$. Then we can find $\alpha_{n-2} \in \mathrm{Aut}(\mathbb{M}/a_E)$ such that $\alpha^{-1}_{n-2}\beta^{-1}_{n-1}\alpha^{-1}_{n-1} \cdot (a^3_E b^3_F c^3_G) = a^0_E b^0_F c^0_G$. Continuing inductively in this way, we find $\alpha_0, \ldots, \alpha_{n-1} \in \mathrm{Aut}(\mathbb{M}/a_E)$ and $\beta_0, \ldots, \beta_{n-1} \in \mathrm{Aut}(\mathbb{M}/b_F)$ such that the same equalities as in the (2)⇒(3) proof hold. Therefore there is a $\sigma \in H$ (namely $\alpha_{n-1}\beta_{n-1}\alpha_{n-2} \ldots \beta_1 \alpha_0 \beta_0$) such that $\sigma \cdot c_G = c'_G$. $\square$

A similar statement is true for arbitrary families of ultraimaginaries: If $(a^i_{E_i})_{i\in I}$ is a (possibly large) family of ultraimaginaries, then $c_G \equiv_{\bigcap_{i\in I}\mathrm{dcl}^{\mathrm{u}}_{\lambda}(a^i_{E_i})} c'_G$ if and only if there is a $\sigma \in \left\langle \bigcup_{i\in I}\mathrm{Aut}(\mathbb{M}/a^i_{E_i})\right\rangle$ such that $\sigma\cdot c_G = c'_G$. There is also an analog of (3), but it is more awkward to state.

Lascar strong type.

Definition 1.11. For any co-small group $G \leq \mathrm{Aut}(\mathbb{M})$, let G_f be the group generated by all groups of the form $\mathrm{Aut}(\mathbb{M}/M) \leq G$ with M a small model. For any ultraimaginary a_E, let $\mathrm{Autf}(\mathbb{M}/a_E) = \mathrm{Aut}(\mathbb{M}/a_E)_f$.

We say that b_F and c_F have the same *Lascar strong type over* a_E, written $b_F \equiv^{\mathrm{L}}_{a_E} c_F$, if there is $\sigma \in \mathrm{Autf}(\mathbb{M}/a_E)$ such that $\sigma\cdot b_F = c_F$.

Proposition 1.12. *For any co-small groups G and H, $[\![G]\!] \in \mathrm{bdd}^{\mathrm{u}}[\![H]\!]$ if and only if $G \geq H_f$.*

Proof. Assume that $[\![G]\!] \in \mathrm{bdd}^{\mathrm{u}}[\![H]\!]$. Note that for a model M, by Lemma 1.3, we have that $[\![G]\!] \in \mathrm{bdd}^{\mathrm{u}}(M)$ if and only if $G \geq \mathrm{Aut}(\mathbb{M}/M)$. Therefore, for any model M with $[\![H]\!] \in \mathrm{bdd}^{\mathrm{u}}(M)$, we must have that $[\![G]\!] \in \mathrm{bdd}^{\mathrm{u}}[\![H]\!] \subseteq \mathrm{bdd}^{\mathrm{u}}(M)$ and so $G \geq \mathrm{Aut}(\mathbb{M}/M)$. Since $[\![H]\!] \in \mathrm{bdd}^{\mathrm{u}}(M)$ if and only if $H \geq \mathrm{Aut}(\mathbb{M}/M)$, we have that $G \geq H_f$.

Conversely, assume that $G \geq H_f$. This implies that for any small model M with $[\![H]\!] \in \mathrm{bdd}^{\mathrm{u}}(M)$, we have $H_f \geq \mathrm{Aut}(\mathbb{M}/M)$, so $G \geq \mathrm{Aut}(\mathbb{M}/M)$ and $[\![G]\!] \in \mathrm{dcl}^{\mathrm{u}}(M)$. Fix some such model N. Assume for the sake of contradiction that $[\![G]\!] \notin \mathrm{bdd}^{\mathrm{u}}[\![H]\!]$. For any λ, we can find $(\sigma_i)_{i<\lambda}$ in $H = \mathrm{Aut}(\mathbb{M}/[\![H]\!])$ such that $\sigma_i\cdot[\![G]\!] \neq \sigma_j\cdot[\![G]\!]$ for each $i<j<\lambda$. Since $[\![G]\!] = a_E$ for some a with $|a| \leq |N|$ by Proposition 1.7, we have that if λ is larger than $2^{|N|+|T|}$, there must be $i<j<\lambda$ such that $\sigma_i\cdot[\![G]\!] \equiv_N \sigma_j\cdot[\![G]\!]$. Let $N' = \sigma_i^{-1}\cdot N$. N' is now a model satisfying $\mathrm{Aut}(\mathbb{M}/N') \leq G$. So $[\![G]\!] \in \mathrm{dcl}^{\mathrm{u}}(N')$, but $[\![G]\!] \equiv_{N'} \sigma_i^{-1}\sigma_j\cdot[\![G]\!]$ and $[\![G]\!] \neq \sigma_i^{-1}\sigma_j\cdot[\![G]\!]$, which is a contradiction. $\square$

An important fact about ultraimaginaries is that $\mathrm{bdd}^{\mathrm{u}}$ has the same relationship with Lascar strong types that $\mathrm{bdd}^{\mathrm{heq}}$ has with Kim–Pillay strong types.

For any a_E and b_F, by an abuse of notation, we'll write $[b_F]_{\equiv^{\mathrm{L}}_{a_E}}$ for $[ab]_G$, where $G(ab, a'b')$ holds if and only if aEa' and $b_F \equiv^{\mathrm{L}}_{a_E} b'_F$. Note in particular that $[b_F]_{\equiv^{\mathrm{L}}_{a_E}} = [b'_F]_{\equiv^{\mathrm{L}}_{a_E}}$ if and only if $b_F \equiv^{\mathrm{L}}_{a_E} b'_F$.

Proposition 1.13. *For any ultraimaginaries a_E, b_F, and c_F, the following are equivalent.*

(1) $b_F \equiv^{\mathrm{L}}_{a_E} c_F$.

(2) $b_F \equiv_{\mathrm{bdd}^{\mathrm{u}}_{\lambda}(a_E)} c_F$ *for all* λ.

(3) $b_F \equiv_{\mathrm{bdd}^{\mathrm{u}}_{|a|+|T|}(a_E)} c_F$.

Proof. To see that (1) implies (3), fix a model M with $a_E \in \mathrm{bdd}^u(M)$ and some automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}/M)$. By Lemma 1.3, we have that $\mathrm{Aut}(\mathbb{M}/M) \leq \mathrm{Aut}(\mathbb{M}/\,\mathrm{bdd}^{\mathrm{u}}_{|a|+|T|}(a_E))$. Therefore $b_F \equiv_{\mathrm{bdd}^{\mathrm{u}}_{|a|+|T|}(a_E)} \sigma \cdot b_F$. By induction, we therefore have that $b_F \equiv^{\mathrm{L}}_{a_E} c_F$ implies $b_F \equiv_{\mathrm{bdd}^{\mathrm{u}}_{|a|+|T|}(a_E)} c_F$.

Corollary 1.8 implies that $\mathrm{Aut}(\mathbb{M}/\,\mathrm{bdd}^{\mathrm{u}}_{\lambda}(a_E)) \geq \mathrm{Aut}(\mathbb{M}/\,\mathrm{bdd}^{\mathrm{u}}_{|a|+|T|}(a_E))$ for all λ, so (3) implies (2).

To see that (2) implies (1), note that $[b_F]_{\equiv^{\mathrm{L}}_{a_E}} \in \mathrm{bdd}^{\mathrm{u}}_{\lambda}(a_E)$ for some sufficiently large λ (because there are a bounded number of Lascar strong types over a_E). Thus if $b_F \equiv_{\mathrm{bdd}^{\mathrm{u}}_{\lambda}(a_E)} c_F$, we must have $[b_F]_{\equiv^{\mathrm{L}}_{a_E}} = [c_F]_{\equiv^{\mathrm{L}}_{a_E}}$ or, in other words, $b_F \equiv^{\mathrm{L}}_{a_E} c_F$. □

2. Bounded ultraimaginary independence

Definition 2.1. Given sets of ultraimaginaries A, B, and C, we write $B \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A C$ to mean that $\mathrm{bdd}^{\mathrm{u}}(AB) \cap \mathrm{bdd}^{\mathrm{u}}(AC) = \mathrm{bdd}^{\mathrm{u}}(A)$.

Recall that $\mathrm{bdd}^{\mathrm{u}}$ is a closure operator (i.e., if $c_G \in \mathrm{bdd}^{\mathrm{u}}(b_F)$ and $b_F \in \mathrm{bdd}^{\mathrm{u}}(a_E)$, then $c_G \in \mathrm{bdd}^{\mathrm{u}}(a_E)$). We will ultimately show (in Proposition 2.3) that the following are equivalent: $b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$, $\mathrm{bdd}^{\mathrm{u}}_{\kappa}(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}_{\kappa}(a_E c_G) = \mathrm{bdd}^{\mathrm{u}}_{\kappa}(a_E)$ for all κ, and $\mathrm{bdd}^{\mathrm{u}}_{\kappa}(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}_{\kappa}(a_E c_G) = \mathrm{bdd}^{\mathrm{u}}_{\kappa}(a_E)$ for $\kappa = |T| + |abc|$. $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$ satisfies some of the familiar properties of $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{a}}$.

Proposition 2.2. *Fix ultraimaginaries a_E, b_F, c_G, and e_I.*

- (Invariance) *If $a_E b_F c_G \equiv a'_E b'_F c'_G$, then $b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$ if and only if $b'_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a'_E} c'_G$.*
- (Symmetry) *$b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$ if and only if $c_G \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} b_F$.*
- (Monotonicity) *If $b_F c_G \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} d_H e_I$, then $b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} d_H$.*
- (Transitivity) *If $b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$ and $d_H \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E b_F} c_G$, then $b_F d_H \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$.*
- (Normality) *If $b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$, then $a_E b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} a_E c_G$.*
- (Anti-reflexivity) *If $b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} b_F$, then $b_F \in \mathrm{bdd}^{\mathrm{u}}(a_E)$.*

Proof. Everything except transitivity is immediate. The argument for transitivity is the same as the argument for transitivity of $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{a}}$: Assume that $b_F \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$ and $d_H \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E b_F} c_G$. Let e_I be an element of $\mathrm{bdd}^{\mathrm{u}}(a_E b_F d_H) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c_G)$. This implies that it is an element of $\mathrm{bdd}^{\mathrm{u}}(a_E b_F d_H) \cap \mathrm{bdd}^{\mathrm{u}}(a_E b_F c_G)$, so by assumption it is an element of $\mathrm{bdd}^{\mathrm{u}}(a_E b_F)$. But this means that it's in both $\mathrm{bdd}^{\mathrm{u}}(a_E b_F)$ and $\mathrm{bdd}^{\mathrm{u}}(a_E c_G)$, so, by assumption again, it is an element of $\mathrm{bdd}^{\mathrm{u}}(a_E)$. □

Part of the goal of this paper is to prove full existence and therefore also extension for $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$ (although only over hyperimaginary bases).

- (Full existence over hyperimaginaries) For any set of hyperimaginaries A and ultraimaginaries b_E and c_F, there is $c'_F \equiv_A c_F$ such that $b_E \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A c'_F$.

- (Extension over hyperimaginaries) For any set of hyperimaginaries A and ultraimaginaries b_E, c_F, and d_G, if $b_E \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}}_A c_F$, then there is $b'_E \equiv_{Ac_F} b_E$ such that $b'_E \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}}_A c_F d_G$.

A fairly general argument will allow us to upgrade $\equiv_A$ to $\equiv^{\mathrm{L}}_A$ in the above two conditions, which we establish in Theorem 3.6 and Corollary 3.8.

Finite character fails very badly, of course: As considered in [Wagner 2015, Example 2.8], if E is the equivalence relation on ω-tuples of equality on cofinitely many indices, then for some sequences $(a_i)_{i<\omega}$, we will have $a_{<n} \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}} [a_{<\omega}]_E$ for all n, yet $a_{<\omega} \not\mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}} [a_{<\omega}]_E$. Given the existence of higher and higher cardinality generalizations of the previous example (e.g., equality on co-countably many indices on ω_1-tuples), local character seems unlikely except possibly in the presence of large cardinals. We do have some control over the relevant cardinalities, however.

Proposition 2.3. *For any a_E, b_F, and c_G, $b_F \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$ if and only if*

$$\mathrm{bdd}^{\mathrm{u}}_\lambda(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E c_G) = \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E),$$

where $\lambda = |ab| + |T|$.

Proof. Let $\lambda = |ab| + |T|$. If $b_F \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$, then $\mathrm{bdd}^{\mathrm{u}}_\lambda(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E c_G) = \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E)$.

Conversely, assume that $b_F \not\mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$. There is some

$$d_H \in (\mathrm{bdd}^{\mathrm{u}}(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c_G)) \setminus \mathrm{bdd}^{\mathrm{u}}(a_E).$$

By Corollary 1.8, there is e_I of arity at most λ such that d_H and e_I are interdefinable. This means that $e_I \in (\mathrm{bdd}^{\mathrm{u}}_\lambda(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E c_G)) \setminus \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E)$. Therefore $\mathrm{bdd}^{\mathrm{u}}_\lambda(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E c_G) \neq \mathrm{bdd}^{\mathrm{u}}_\lambda(a_E)$. $\square$

The following characterization of $\mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}}$ (and the manner of proof) is essentially due to Wagner [2015].

Proposition 2.4. *For any ultraimaginaries a_E, b_F, and c_G, the following are equivalent.*

1. *$b_F \mathop{\smile\!\!\!\!\!\!\mid}^{\mathrm{bu}}_{a_E} c_G$.*
2. *For any $b'_F \equiv^{\mathrm{L}}_{a_E} b_F$, there are $b^0, c^0, b^1, c^1, \dots, c^{n-1}, b^n$ such that $b^0 = b$, $c^0 = c$, $b^n = b'$, and for each $i < n$, $b^i_F \equiv^{\mathrm{L}}_{a_E c^i_G} b^{i+1}_F$ and $c^i_G \equiv^{\mathrm{L}}_{a_E b^{i+1}_F} c^{i+1}_G$ if $i < n-1$.*
3. *$\langle \mathrm{Autf}(\mathbb{M}/a_E b_F) \cup \mathrm{Autf}(\mathbb{M}/a_E c_G) \rangle = \mathrm{Autf}(\mathbb{M}/a_E)$.*

Proof. Let $H = \langle \mathrm{Autf}(\mathbb{M}/a_E b_F) \cup \mathrm{Autf}(\mathbb{M}/a_E c_G) \rangle$.

$\neg(3) \Rightarrow \neg(1)$. Assume $H \neq \mathrm{Autf}(\mathbb{M}/a_E)$, which implies that $H < \mathrm{Autf}(\mathbb{M}/a_E) = \mathrm{Autf}(\mathbb{M}/a_E)_f$. By Proposition 1.12, we have that $[\![H]\!] \notin \mathrm{bdd}^{\mathrm{u}}[\![\mathrm{Autf}(\mathbb{M}/a_E)]\!] = \mathrm{bdd}^{\mathrm{u}}(a_E)$. But since $\mathrm{Autf}(\mathbb{M}/a_E b_F) = \mathrm{Autf}(\mathbb{M}/a_E b_F)_f \leq H$ and $\mathrm{Autf}(\mathbb{M}/a_E c_G) = \mathrm{Autf}(\mathbb{M}/a_E c_G)_f \leq H$, we have that $[\![H]\!] \in \mathrm{bdd}^{\mathrm{u}}(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c_G)$ again by Proposition 1.12.

(3)⇒(1). Suppose $H = \mathrm{Autf}(\mathbb{M}/a_E)$. Fix an ultraimaginary $d_I \in \mathrm{bdd}^{\mathrm{u}}(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c_G)$. By Proposition 1.12, we have that $H \leq \mathrm{Autf}(\mathbb{M}/a_E d_I) \leq \mathrm{Autf}(\mathbb{M}/a_E)$, which implies that $\mathrm{Autf}(\mathbb{M}/a_E d_I) = H$. Hence by Proposition 1.12, $d_I \in \mathrm{bdd}^{\mathrm{u}}(a_E)$. Since we can do this for any such ultraimaginary, we have that $b_F \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_{a_E} c_G$.

(1)⇒(2). Let $b^*_{F^*} = [\![\mathrm{Autf}(\mathbb{M}/a_E b_F)]\!]$ and $c^*_{G^*} = [\![\mathrm{Autf}(\mathbb{M}/a_E c_G)]\!]$. Note that $\mathrm{bdd}^{\mathrm{u}}(a_E b_F) = \mathrm{dcl}^{\mathrm{u}}(b^*_{F^*})$ and $\mathrm{bdd}^{\mathrm{u}}(a_E c_G) = \mathrm{dcl}^{\mathrm{u}}(c^*_{G^*})$ (by Definition 1.9 and Proposition 1.12). In particular, we have that $\mathrm{dcl}^{\mathrm{u}}(b^*_{F^*}) \cap \mathrm{dcl}^{\mathrm{u}}(c^*_{G^*}) = \mathrm{bdd}^{\mathrm{u}}(a_E)$. Fix $b'_F \equiv^{\mathrm{L}}_{a_E} b_F$. By passing to a different representative of the F-equivalence class b'_F, we may assume that $b' \equiv^{\mathrm{L}}_{a_E} b$. Fix c' such that $bc \equiv^{\mathrm{L}}_{a_E} b'c'$. By Proposition 1.13, we have that $b'c' \equiv_{\mathrm{bdd}^{\mathrm{u}}_\lambda(a_E)} bc$ for all λ, so $b'c' \equiv_{\mathrm{dcl}^{\mathrm{u}}_\lambda(b^*_{F^*}) \cap \mathrm{dcl}^{\mathrm{u}}_\lambda(c^*_{G^*})} bc$ for all λ. Therefore, by Proposition 1.10, we can find a sequence $(b^{*i} c^{*i} b^i c^i)_{i \leq n}$ such that $b^{*0} = b^*$, $c^{*0} = c^*$, $b^0 c^0 = bc$, $b^n c^n = b'c'$, and for each $i < n$,

- if i is even, $b^{*i} = b^{*i+1}$ and $c^{*i} b^i c^i \equiv_{b^{*i}} c^{*i+1} b^{i+1} c^{i+1}$ and
- if i is odd, $c^{*i} = c^{*i+1}$ and $b^{*i} b^i c^i \equiv_{c^{*i}} b^{*i+1} b^{i+1} c^{i+1}$.

This implies, by induction, that $b^i c^i \equiv^{\mathrm{L}}_{a_E b^i_F} b^{i+1} c^{i+1}$ and $b^i_F = b^{i+1}_F$ for each even i and $b^i c^i \equiv^{\mathrm{L}}_{a_E c^i_G} b^{i+1} c^{i+1}$ and $c^i_G = c^{i+1}_G$ for each odd i, so $b^0 c^1 b^2 c^3 \dots c^{n-1} b^n$ is the sequence required by the proposition (after reindexing).

(2)⇒(1). Assume (2), but also assume for the sake of contradiction that (1) fails. Let d_H be an element of $(\mathrm{bdd}^{\mathrm{u}}(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c_G)) \setminus \mathrm{bdd}^{\mathrm{u}}(a_E)$. Since d_H is not bounded over a_E, there must be some $d'_H \equiv^{\mathrm{L}}_{a_E} d_H$ such that $d'_H \notin \mathrm{bdd}^{\mathrm{u}}(a_E b_E) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c_G)$. Find b'_F such that $b_F d_H \equiv^{\mathrm{L}}_{a_E} b'_F d'_H$. Let b^0, c^0, b^1, $c^1 \dots$, c^{n-1}, b^n be as in (2), with $b^n = b'$. Find $d^{1/2}, d^1, d^{3/2}, d^2, \dots, d^{n-(1/2)}, d^n$ such that $d^{1/2} = d$ and for each $i < n$,

- $b^i_F d^{i+(1/2)}_H \equiv^{\mathrm{L}}_{a_E c^i_G} b^{i+1}_F d^{i+1}$ and
- $c^i_G d^{i+1}_H \equiv^{\mathrm{L}}_{a_E b^{i+1}_F} c^{i+1}_G d^{i+(3/2)}_H$ if $i < n-1$.

We now have that $b'_F d'_H \equiv^{\mathrm{L}}_{a_E} b_F d_H \equiv^{\mathrm{L}}_{a_E} b'_F d^n_H$, so in particular, $d'_H \equiv^{\mathrm{L}}_{a_E b'_F} d^n_H$. For some $i < n$, consider $e_I \in \mathrm{bdd}^{\mathrm{u}}(a_E b^i_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c^i_G)$. Since $e_I \in \mathrm{bdd}^{\mathrm{u}}(a_E c^i_G)$ and since $b^i_F d^{i+(1/2)}_H \equiv_{\mathrm{bdd}^{\mathrm{u}}_\lambda(a_E c^i_G)} b^{i+1}_F d^{i+1}$ for all λ (by Proposition 1.13), we must have that $b^i_F d^{i+(1/2)}_H \equiv_{a_E e_I} b^{i+1}_F d^{i+1}$ and so $e_I \in \mathrm{bdd}^{\mathrm{u}}(a_E b^{i+1}_F)$ as well. By the reverse argument and since we can do this for any such ultraimaginary, we get that

$$\mathrm{bdd}^{\mathrm{u}}(a_E b^i_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c^i_G) = \mathrm{bdd}^{\mathrm{u}}(a_E b^{i+1}_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c^i_G).$$

Likewise, for any $i < n-1$, we get

$$\mathrm{bdd}^{\mathrm{u}}(a_E b^{i+1}_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c^i_G) = \mathrm{bdd}^{\mathrm{u}}(a_E b^{i+1}_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c^{i+1}_G).$$

Therefore $d^n_H \in \mathrm{bdd}^{\mathrm{u}}(a_E b^n_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c^{n-1}_G)$, so since $d^n_H \equiv^{\mathrm{L}}_{a_E b^n_F} d'_H$ and so $d^n_H \equiv_{\mathrm{bdd}^{\mathrm{u}}_\lambda(a_E b^n_F)} d'_H$ for every λ (by Proposition 1.13), we must also have

$$d'_H \in \mathrm{bdd}^{\mathrm{u}}(a_E b^n_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c^{n-1}_G) = \mathrm{bdd}^{\mathrm{u}}(a_E b_F) \cap \mathrm{bdd}^{\mathrm{u}}(a_E c_G),$$

which is a contradiction. □

3. Full existence

We will use the tree bookkeeping machinery from [Kaplan and Ramsey 2020], with some minor extensions (the notation $\mathcal{T}^*_\alpha$ and $\mathcal{F}_\alpha$).

Definition 3.1. For any ordinal α, $\mathcal{L}_{s,\alpha}$ is the language

$$\{\trianglelefteq, \wedge, <_{\mathrm{lex}}, P_0, P_1, \dots, P_\beta(\beta < \alpha), \dots\},$$

with $\trianglelefteq$ and $<_{\mathrm{lex}}$ binary relations, $\wedge$ a binary function, and each P_β a unary relation.

For any ordinal α, we write $\mathcal{T}^*_\alpha$ for the set of functions f with codomain ω and finite support such that $\mathrm{dom}(f)$ is an end segment of α. (For the sake of some minor edge cases, we will regard the empty functions in various $\mathcal{T}^*_\alpha$'s as distinct objects.) We write $\mathcal{T}_\alpha$ for the set of functions $f \in \mathcal{T}^*_\alpha$ with $\mathrm{dom}(f) = [\beta, \alpha)$ for a nonlimit ordinal β. We write $\mathcal{F}_{\alpha+1}$ (for forest) for $\mathcal{T}_{\alpha+1} \setminus \{\varnothing\}$.

We interpret $\mathcal{T}^*_\alpha$ and $\mathcal{T}_\alpha$ as $\mathcal{L}_{s,\alpha}$-structures by

- $f \trianglelefteq g$ if and only if $f \subseteq g$;
- $f \wedge g = f \restriction [\beta, \alpha) = g \restriction [\beta, \alpha)$, where $\beta = \min\{\gamma : f \restriction [\gamma, \alpha) = g \restriction [\gamma, \alpha)\}$ (with the understanding that $\min \varnothing = \alpha$);
- $f <_{\mathrm{lex}} g$ if and only if either $f \triangleleft g$ or f and g are $\trianglelefteq$-incomparable, $\mathrm{dom}(f \wedge g) = [\gamma, \alpha)$, and $f(\gamma) < g(\gamma)$; and
- $P_\beta(f)$ holds if and only if $\mathrm{dom}(f) = [\beta, \alpha)$.

We write $\langle i \rangle_\alpha$ for the function $\{(\alpha, i)\}$ (which is an element of $\mathcal{T}^*_{\alpha+1}$). Given $i < \omega$ and $f \in \mathcal{T}^*_\alpha$ with $\mathrm{dom}(f) = [\beta + 1, \alpha)$, we write $f \frown i$ to mean the function $f \cup \{(\beta, i)\}$ (which is an element of $\mathcal{T}^*_\alpha$). Given $i < \omega$ and $f \in \mathcal{T}^*_\alpha$, we write $i \frown f$ to mean the function $\{(\alpha, i)\} \cup f$ (which is an element of $\mathcal{T}^*_{\alpha+1}$).[9]

For $\alpha < \beta$, we define the canonical inclusion map $\iota_{\alpha\beta} : \mathcal{T}_\alpha \to \mathcal{T}_\beta$ by $\iota_{\alpha\beta}(f) = f \cup \{(\gamma, 0) : \gamma \in \beta \setminus \alpha\}$. (Note that $\iota_{\alpha,\alpha+1}(f) = 0 \frown f$.)

For $\beta \le \alpha$, we write ζ^α_β for the function whose domain is $[\beta, \alpha)$ with the property that $\zeta^\alpha_\beta(\gamma) = 0$ for all $\gamma \in [\beta, \alpha)$. (Note that ζ^α_α is $\mathcal{T}_\alpha$'s copy of the empty function.)

Given a family $(b_f)_{f \in X}$, we may refer to it briefly as $b_{\in X}$.

[9]Note that this notation is not ambiguous when f is an empty function, as we are regarding the empty functions in different $\mathcal{T}^*_\alpha$'s as distinct objects.

Definition 3.2. Given $X \subseteq \mathcal{T}_\alpha^*$, we say that a family $(b_f)_{f\in X}$ is *s-indiscernible over A* if for any tuples $f_0 \dots f_{n-1}$ and $g_0 \dots g_{n-1}$ in X with $f_0 \dots f_{n-1} \equiv^{\mathrm{qf}} g_0 \dots g_{n-1}$, $b_{f_0} \dots b_{f_{n-1}} \equiv_A b_{g_0} \dots b_{g_{n-1}}$, where quantifier-free type is in the language $\mathcal{L}_{s,\alpha}$. (Note that this does not entail that b_f's on different levels are tuples of the same sort.)

Given $f \in \mathcal{T}_\alpha$, we write $b_{\trianglerighteq f}$ to refer to some fixed enumeration of the set $\{b_g : g \in \mathcal{T}_\alpha,\ f \trianglelefteq g\}$. In particular, we choose this enumeration in a uniform way so that if $(b_f)_{f\in\mathcal{T}_\alpha}$ is s-indiscernible over A, then for any f with domain $[\beta+1,\alpha)$, the sequence $(b_{\trianglerighteq f\frown i})_{i<\omega}$ is A-indiscernible. When f is an element of $\mathcal{T}_\alpha^*$, we will also write $b_{\trianglerighteq f}$ for some fixed enumeration of the set $\{b_g : g \in \mathcal{T}_\alpha,\ f \subseteq g\}$. One particular example of this will be sequences of the form $(b_{\trianglerighteq \zeta^\alpha_{\beta+1}\frown i})_{i<\omega}$, where β is a limit ordinal. This is essentially the only situation in which we need to consider $\mathcal{T}_\alpha^*$.

Note that for a limit ordinal α, $(b_f)_{f\in\mathcal{T}_\alpha}$ is s-indiscernible over A if and only if $(b_f)_{f\in\iota_{\beta,\alpha}(\mathcal{T}_\beta)}$ is s-indiscernible over A for every $\beta<\alpha$.

We will also need the following fact.

Fact 3.3 (modeling property for s-indiscernibles [Kim et al. 2014, Theorem 4.3]). *Let X be $\mathcal{T}_\alpha$ or $\mathcal{F}_{\alpha+1}$. For any $(b_f)_{f\in X}$ and any set A of* ***hyper****imaginaries, there is a family of tuples $(c_f)_{f\in X}$ that is s-indiscernible over A and* ***locally based on*** *$b_{\in X}$ (i.e., for any finite tuple $f_0 \dots f_{n-1}$ from X and any neighborhood U of* $\mathrm{tp}(c_{f_0} \dots c_{f_{n-1}}/A)$ *(in the appropriate type space), there is a tuple $g_0 \dots g_{n-1}$ from X such that $f_0 \dots f_{n-1} \equiv^{\mathrm{qf}} g_0 \dots g_{n-1}$ and* $\mathrm{tp}(b_{g_0} \dots b_{g_{n-1}}/A) \in U$).

Note that while Fact 3.3 is normally formulated for discrete logic, the corresponding statement in continuous logic follows easily from a very soft general argument: Given a metric structure M and a tree $(b_f)_{f\in X}$ of elements of M, find α large enough that M, $\mathrm{Th}(M)$, and $b_{\in X}$ are elements of V_α and apply [Kim et al. 2014, Theorem 4.3] to V_α as a discrete structure and get some A-s-indiscernible family $(c_f^*)_{f\in X}$ of elements of an elementary extension $V_\alpha^* \succeq V_\alpha$. These elements live inside a structure $M^* \in V_\alpha^*$ that is internally a model of $\mathrm{Th}(M)$. By taking the standard parts of each real-valued predicate in M^* and then completing with regards to the metric, we get a metric structure N that is an elementary extension of M. For each $f \in X$, let c_f be the image in N of c_f^* under the canonical map from M^* to N. It is straightforward to check that $(c_f)_{f\in X}$ is the required A-s-indiscernible family.

Before proving full existence for $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$, we will need a lemma.

Lemma 3.4. *Fix α and $\gamma>\alpha$. Let $(e_f)_{f\in\mathcal{F}_{\gamma+1}}$ be an s-indiscernible family of real tuples over a set A of* ***hyper****imaginary parameters. Let $\lambda = |Ae_{\trianglerighteq\zeta_\alpha^\gamma}| + |T|$. Suppose that there is an ultraimaginary c_F such that $c_F \in \mathrm{bdd}^{\mathrm{u}}_\lambda(Ae_{\trianglerighteq\zeta_\alpha^{\gamma+1}}) \cap \mathrm{bdd}^{\mathrm{u}}_\lambda(Ae_{\trianglerighteq 1\frown\zeta_\alpha^\gamma})$. Then there is a model M with $Ac_F \subseteq \mathrm{dcl}^{\mathrm{u}}(M)$ and $|M| \le \lambda$ such that $(e_f)_{f\in\mathcal{F}_{\gamma+1}}$ is s-indiscernible over M.*

Proof. By Fact 3.3, we can find a set of real parameters B such that $|B| \le |A| + |T|$, $A \subseteq \mathrm{bdd}^{\mathrm{heq}}(B)$, and $(e_f)_{f \in \mathcal{F}_{\gamma+1}}$ is s-indiscernible over B.

Let T' be a Skolemization of T with $|T'| = |T|$. Let $\mathbb{M}'$ be the monster model of T', which we may think of as an expansion of $\mathbb{M}$. By Fact 3.3, we can find $(b'_f)_{f \in \mathcal{F}_{\gamma+1}}$ locally based on $(e_f)_{f \in \mathcal{F}_{\gamma+1}}$ which is s-indiscernible over B (in T'). By considering an automorphism of $\mathbb{M}$ (in T), we may assume that $(b'_f)_{f \in \mathcal{F}_{\gamma+1}}$ actually is $(e_f)_{f \in \mathcal{F}_{\gamma+1}}$, so that $(e_f)_{f \in \mathcal{F}_{\gamma+1}}$ is s-indiscernible over B (in T').

Find an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}'/B)$ satisfying $\sigma \cdot e_{\trianglerighteq \langle i+1 \rangle_{\gamma+1}} = e_{\trianglerighteq \langle i \rangle_{\gamma+1}}$ for every $i < \omega$. Let M be the Skolem hull of $B \cup \sigma \cdot e_{\trianglerighteq \zeta_\alpha^{\gamma+1}}$. Note that $(e_f)_{f \in \mathcal{F}_{\gamma+1}}$ is s-indiscernible over M (and therefore the same is true in T). Furthermore, note that $|M| \le \lambda$.

Let M_i be the Skolem hull of $Be_{\trianglerighteq i \frown \zeta_\alpha^\gamma}$ for both $i \in \{0, 1\}$. Note that $c_F \in \mathrm{bdd}^{\mathrm{u}}_\lambda(M_1)$ and $|M_1| \le \lambda$. Pass back to the theory T. Note that M, M_0, and M_1 are still models of T. By Corollary 1.8, there is an invariant equivalence relation G (in T) such that c_F and $[M_1]_G$ are interdefinable. Therefore we have that $[M_1]_G \in \mathrm{bdd}^{\mathrm{u}}_\lambda(Ae_{\trianglerighteq \zeta_\alpha^{\gamma+1}}) \subseteq \mathrm{bdd}^{\mathrm{u}}(M_0) = \mathrm{dcl}^{\mathrm{u}}(M_0)$. Find an automorphism $\tau \in \mathrm{Aut}(\mathbb{M}/M_1)$ such that $\tau(M_0) = M$ (which exists by indiscernibility). τ witnesses that $[M_1]_G \in \mathrm{dcl}^{\mathrm{u}}(M)$ and therefore $c_F \in \mathrm{dcl}^{\mathrm{u}}(M)$, so M is the required model. $\square$

Now we are ready to prove full existence for $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$, but we will take the opportunity to prove a certain technical strengthening which we will need later in the construction of $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley trees.

Lemma 3.5. *If $(b_f)_{f \in \mathcal{T}_\alpha}$ is a tree of real elements that is s-indiscernible over a set of **hyper**imaginaries A, then there is a $\gamma > \alpha$ and a tree $(e_f)_{f \in \mathcal{T}_{\gamma+1}}$ such that*

- *$e_{\in \mathcal{T}_{\gamma+1}}$ is s-indiscernible over A,*
- *for each $f \in \mathcal{T}_\alpha$, $b_f = e_{\iota_{\alpha,\gamma+1}(f)}$, and*
- *$e_{\trianglerighteq \zeta_\alpha^{\gamma+1}} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A e_{\trianglerighteq 1 \frown \zeta_\alpha^\gamma}$.*

(*Note that $e_{\trianglerighteq \zeta_\alpha^{\gamma+1}}$ is the original tree.*)

Proof. If $b_{\in \mathcal{T}_\alpha} \in \mathrm{acl}(A)$, then the statement is trivial, so assume that $b_{\in \mathcal{T}_\alpha} \notin \mathrm{acl}(A)$.

Fix $\lambda = |Ab_{\in \mathcal{T}_\alpha}| + |T|$. By Proposition 2.3, we have that $b_{\in \mathcal{T}_\alpha} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A c$ if and only if $\mathrm{bdd}^{\mathrm{u}}_\lambda(Ab_{\in \mathcal{T}_\alpha}) \cap \mathrm{bdd}^{\mathrm{u}}_\lambda(Ac) = \mathrm{bdd}^{\mathrm{u}}_\lambda(A)$ for any c. Let $\mu = |\mathrm{bdd}^{\mathrm{u}}_\lambda(Ab_{\in \mathcal{T}_\alpha}) \setminus \mathrm{bdd}^{\mathrm{u}}_\lambda(A)|^+$.

We will build a family $(e_f : f \in \iota_{\gamma+1,\mu}(\mathcal{T}_{\gamma+1}))$ inductively, where γ is some successor ordinal less than μ. By an abuse of notation, we will systematically conflate the sets $\iota_{\alpha,\mu}(\mathcal{T}_\alpha)$ and $\mathcal{T}_\alpha$ (and likewise for $\iota_{\alpha,\mu}(\mathcal{F}_{\alpha+1})$ and $\mathcal{F}_{\alpha+1}$) for all $\alpha < \mu$. Note that in general this will mean that $e_{\trianglerighteq \zeta_\beta^\mu}$ is the same thing as $e_{\in \mathcal{T}_\beta}$.

Let $e_f = b_f$ for all $f \in \mathcal{T}_\alpha$. Since $b_{\in \mathcal{T}_\alpha} \notin \mathrm{acl}(A)$, we can find a family $(d_f)_{f \in \mathcal{F}_{\alpha+1}}$ extending $e_{\in \mathcal{T}_\alpha}$ such that $(d_{\trianglerighteq \zeta_{\alpha+1}^\mu \frown i})_{i<\omega}$ is a nonconstant A-indiscernible sequence. By Fact 3.3, we can define e_f for all $f \in \mathcal{F}_{\alpha+1}$ in such a way that the family

$e_{\in \mathcal{F}_{\alpha+1}}$ is locally based on $d_{\in \mathcal{F}_{\alpha+1}}$. In particular, $(e_{\trianglerighteq \zeta^{\mu}_{\alpha+1} \frown i})_{i<\omega}$ will be a nonconstant A-indiscernible sequence.

At successor stage $\beta + 1 \geq \alpha$, assume that we have defined e_f for all $f \in \mathcal{F}_{\beta+1}$ and that the family $(e_f)_{f \in \mathcal{F}_{\beta+1}}$ is s-indiscernible over A. If there is no $d_E \in \operatorname{bdd}^{\mathrm{u}}_{\lambda}(Ab_{\in \mathcal{T}_\alpha}) \setminus \operatorname{bdd}^{\mathrm{u}}_{\lambda}(A)$ such that the family $(e_f)_{f \in \mathcal{F}_{\beta+1}}$ is s-indiscernible over Ad, let $e_{\zeta^{\mu}_{\beta+1}} = \varnothing$ and $\gamma = \beta$ and halt the construction. Otherwise, let $e_{\zeta^{\mu}_{\beta+1}} = d$. For later reference, let $E_{\beta+1}$ be E. Note that the family $e_{\in \mathcal{T}_{\beta+1}}$ is s-indiscernible over A. Since $d_E \notin \operatorname{bdd}^{\mathrm{u}}_{\lambda}(A)$, we can find, by Proposition 1.4, a sequence $(\sigma_i)_{i<\omega}$ of elements of $\operatorname{Aut}(\mathbb{M}/A)$ such that $(\sigma_i \cdot d)_{i<\omega}$ is an A-indiscernible sequence satisfying $(\sigma_i \cdot d) \not\!E_{\beta+1} (\sigma_j \cdot d)$ for any $i < j < \omega$. Now choose $(e_f)_{f \in \mathcal{F}_{\beta+2}}$ in such a way that $e_{\in \mathcal{F}_{\beta+2}}$ extends what was already defined, is s-indiscernible over A, and is locally based on the family $(c_f)_{f \in \mathcal{F}_{\beta+2}}$ defined by $c_{i \frown f} = \sigma_i \cdot e_f$ for all $f \in \mathcal{T}_{\beta+1}$ (which is possible by Fact 3.3). In particular, note that for any $i < j < \omega$, we still have that $(e_{\zeta^{\mu}_{\beta+2} \frown i}, e_{\zeta^{\mu}_{\beta+2} \frown j}) \equiv_A (\sigma_0 \cdot d, \sigma_1 \cdot d)$ and so, in particular, $e_{\zeta^{\mu}_{\beta+2} \frown i} \not\!E_{\beta+1} e_{\zeta^{\mu}_{\beta+2} \frown j}$.

At limit stage β, we have constructed the family $(e_f)_{f \in \mathcal{T}_\beta}$. Note that this family is automatically s-indiscernible over A. Extend it to a family $e_{\in \mathcal{F}_{\beta+1}}$ that is s-indiscernible over A. (This is always possible by Fact 3.3.)

Claim. For any $\beta < \delta < \mu$, if $E_{\beta+1} = E_{\delta+1}$, then $e_{\zeta^{\mu}_{\beta+1}} \not\!E_{\beta+1} e_{\zeta^{\mu}_{\delta+1}}$.

Proof of claim. The sequence $(e_{\zeta^{\mu}_{\beta+2} \frown i})_{i<\omega}$ is $e_{\zeta^{\mu}_{\delta+1}}$-indiscernible. Since

$$e_{\zeta^{\mu}_{\beta+2} \frown 0} \not\!E_{\beta+1} e_{\zeta^{\mu}_{\beta+2} \frown 1},$$

it must be the case that $e_{\zeta^{\mu}_{\delta+1}} \not\!E_{\beta+1} e_{\zeta^{\mu}_{\beta+2} \frown i}$ for all $i < \omega$. $\square_{\text{claim}}$

Let g be the partial function taking β to $[e_{\zeta^{\mu}_{\beta+1}}]_{E_{\beta+1}}$. By the claim, this is an injection into $\operatorname{bdd}^{\mathrm{u}}_{\lambda}(Ab_{\in \mathcal{T}_\alpha}) \setminus \operatorname{bdd}^{\mathrm{u}}_{\lambda}(A)$. By the choice of μ, g's domain cannot be cofinal in μ, so the construction must have halted at some $\gamma < \mu$.

Extend $e_{\in \mathcal{T}_\gamma}$ to $e_{\in \mathcal{F}_{\gamma+1}}$ in such a way that the resulting family is s-indiscernible over A. Set $e_{\zeta^{\mu}_{\gamma+1}} = \varnothing$.

Claim. For any $c_F \in \operatorname{bdd}^{\mathrm{u}}_{\lambda}(Ae_{\trianglerighteq \zeta^{\mu}_{\alpha}}) \setminus \operatorname{bdd}^{\mathrm{u}}_{\lambda}(A)$, $c_F \notin \operatorname{bdd}^{\mathrm{u}}_{\lambda}(Ae_{\trianglerighteq 1 \frown \zeta^{\gamma}_{\alpha}})$.

Proof of claim. Assume there is some $c_F \in (\operatorname{bdd}^{\mathrm{u}}_{\lambda}(Ae_{\trianglerighteq \zeta^{\mu}_{\alpha}}) \cap \operatorname{bdd}^{\mathrm{u}}_{\lambda}(Ae_{\trianglerighteq 1 \frown \zeta^{\gamma}_{\alpha}})) \setminus \operatorname{bdd}^{\mathrm{u}}_{\lambda}(A)$. By Lemma 3.4, we can find a model M with $Ac_F \subseteq \operatorname{dcl}^{\mathrm{u}}(M)$ and $|M| \leq \lambda$ such that $e_{\in \mathcal{F}_{\gamma+1}}$ is s-indiscernible over M. By Corollary 1.8, there is an invariant equivalence relation G such that c_F and $[M]_G$ are interdefinable. But this means that we could have chosen $[M]_G$ to be d_E at stage γ, contradicting the fact that the construction halted. Therefore no such c_F can exist. $\square_{\text{claim}}$

So, by the claim, we have that $\operatorname{bdd}^{\mathrm{u}}_{\lambda}(Ae_{\trianglerighteq \zeta^{\mu}_{\alpha}}) \cap \operatorname{bdd}^{\mathrm{u}}_{\lambda}(Ae_{\trianglerighteq 1 \frown \zeta^{\gamma}_{\alpha}}) = \operatorname{bdd}^{\mathrm{u}}_{\lambda}(A)$. Therefore, by the choice of λ, $e_{\trianglerighteq \zeta^{\mu}_{\alpha}} \mathop{\smile\!\!\!\!\!\!\mid}\nolimits^{\mathrm{bu}}_{A} e_{\trianglerighteq 1 \frown \zeta^{\gamma}_{\alpha}}$, as required. $\square$

Theorem 3.6 (full existence). *For any set of* **hyper***imaginaries A and real tuples b and c, there is $b' \equiv^{\mathrm{L}}_A b$ such that $b' \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A c$.*[10]

Proof. It is sufficient to show this in the special case that $b = c$. Specifically, given d and e, if we can find $d'e' \equiv_A de$ such that $d'e' \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A de$, then we have $d' \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A e$ by monotonicity. So fix a set of hyperimaginaries A and a real tuple b. Let B be a set containing realizations of all Lascar strong types extending $\mathrm{tp}(b/A)$. We can now apply Lemma 3.5 to the family $(B_f)_{f \in \mathcal{T}_0}$ with $B_\varnothing = B$ to get a family $(E_f)_{f \in \mathcal{T}_{\gamma+1}}$ such that $E_{\zeta_0^{\gamma+1}} = B$ for some $f \in \mathcal{T}_{\gamma+1}$, $B \equiv_A B_f$, and $B \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A B_f$. Let σ be an automorphism fixing A taking B_f to B. Let $B' = \sigma \cdot B$. B' still contains realizations of all Lascar strong types extending $\mathrm{tp}(b/A)$, so we can find $b' \in B'$ with $b' \equiv^{\mathrm{L}}_A b$, which is the required element. □

Corollary 3.7. *For any set of* **hyper***imaginaries A and any ultraimaginaries b_E and c_F, there is $b'_E \equiv^{\mathrm{L}}_A b_E$ such that $b'_E \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A c_F$.*

Proof. Apply Theorem 3.6 to b and c to get $b' \equiv^{\mathrm{L}}_A b$ such that $b' \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A c$. We then have that $\mathrm{bdd}^{\mathrm{u}}(b') \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A \mathrm{bdd}^{\mathrm{u}}(c)$, so by monotonicity, $b'_E \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A c_F$. □

Corollary 3.8 (extension). *For any set of* **hyper***imaginaries A and any ultraimaginaries b_E, c_F, and d_G, if $b_E \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A c_F$, then there is $b'_E \equiv^{\mathrm{L}}_{Ac_F} b_E$ such that $b'_E \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A c_F d_G$.*

Proof. By Corollary 3.7, we can find $b'_E \equiv^{\mathrm{L}}_{Ac_F} b$ such that $b'_E \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_{Ac_F} d_G$. By symmetry and transitivity, we have that $b'_E \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A c_F d_G$. □

Compactness is essential in the proof of Fact 3.3 and therefore also Theorem 3.6, which raises the following question.

Question 3.9. *Does Theorem 3.6 hold when A is a set of ultraimaginaries?*

4. Total $\mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}$-Morley sequences

Definition 4.1. A $\mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}$*-Morley sequence over A* is an A-indiscernible sequence $(b_i)_{i<\omega}$ such that $b_i \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A b_{<i}$ for each $i < \omega$.

A *weakly total* $\mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}$*-Morley sequence over A* is an A-indiscernible sequence $(b_i)_{i<\omega}$ such that for any finite I and any J (of any order type), if $I + J \equiv^{\mathrm{EM}}_A b_{<\omega}$, then $I \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}_A J$.[11]

[10] Anand Pillay has pointed out to us that Theorem 3.6 also follows from Theorem 615 of [Lascar 1982] (together with Wagner's characterization $\mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}$ from [Wagner 2015] given in our Proposition 2.4). Theorem 615 is stated for countable sets of parameters in a countable theory, but it is clear that the proof generalizes to the uncountable case as well.

[11] Note that if we modified this definition to allow I to be any order type and require that J be finite, the resulting sequences would be precisely the order-reversals of the weakly total $\mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}$-Morley sequences as we have defined the term here (by symmetry of $\mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}$).

A *total* $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$*-Morley sequence over* A is an A-indiscernible sequence $(b_i)_{i<\omega}$ such that for any I and J (of any order type), if $I + J \equiv^{\mathrm{EM}}_A b_{<\omega}$, then $I \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A J$.

We could write down stronger and weaker forms of the $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley condition, but we are really only interested in total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences, as they seem to be a fairly robust class (see Theorem 4.8). Weakly total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences seem to be the best we can get without large cardinals, however, which does raise the following question.

Question 4.2. *Is every weakly total* $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$*-Morley sequence a total* $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$*-Morley sequence*?

One immediate property of total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences is that they act as universal witnesses of the relation $\equiv^{\mathrm{L}}_A$ in a strong way.

Proposition 4.3. *For any* A *and* b, *if there is a total* $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$*-Morley sequence* $(b_i)_{i<\omega}$ *over* A *with* $b_0 = b$, *then for any* b', $b' \equiv^{\mathrm{L}}_A b$ *if and only if there are* $I_0, J_0, I_1, \dots, J_{n-1}, I_n$ *such that* $b \in I_0$, $b' \in I_n$, *and, for each* $i < n$, $I_i + J_i$ *and* $I_{i+1} + J_i$ *are both* A*-indiscernible and have the same EM-type as* $b_{<\omega}$.

Proof. Let $I = (b_i)_{i<\omega}$. We only need to prove that if $b' \equiv^{\mathrm{L}}_A b$, then the required configuration exists (as the required configuration is clearly sufficient to witness that $b' \equiv^{\mathrm{L}}_A b$). Choose I' so that $bI \equiv^{\mathrm{L}}_A b'I'$. Extend I to an A-indiscernible sequence $I + J$ with $I \equiv_A J$. By assumption $I \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A J$, so by Proposition 2.4, there are $I_0, J_0, I_1, J_1, \dots, J_{n-1}, I_n$ such that $I_0 = I$, $J_0 = J$, $I_n = I'$, and for each $i < n$, $I_i \equiv^{\mathrm{L}}_{AJ_i} I_{i+1}$ and $J_i \equiv^{\mathrm{L}}_{AI_{i+1}} J_{i+1}$ if $i < n$. Since $I_0 + J_0$ is A-indiscernible, we can show by induction that $I_i + J_i$ and $I_{i+1} + J_i$ are both A-indiscernible and have the same EM-type as $I_0 = b_{<\omega}$. ☐

A similar statement is true for weakly total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences, which we will state in Corollary 4.18 after we have shown that weakly total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences always exist without set-theoretic hypotheses.

Characterization of total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$*-Morley sequences.*

Definition 4.4. For any set of parameters A, we write $\approx_A$ for the transitive closure of the relation $I \sim_A J$ that holds if and only if I and J are both infinite A-indiscernible sequences (of real or hyperimaginary elements) and either $I + J$ or $J + I$ is an A-indiscernible sequence.

By an abuse of notation, we write $[I]_{\approx_A}$ for the ultraimaginary $[AI]_E$, where E is the equivalence relation on tuples of the same length as AI such that $E(AI, BJ)$ holds if and only if $A = B$ in our fixed enumeration and $I \approx_A J$.

Note that we do not in general require that I and J have the same order type. Also note that $\approx_A$ is reflexive: For any infinite A-indiscernible sequence I, we can find an infinite sequence J such that $I + J$ is also A-indiscernible. Then $I \sim_A J \sim_A I$, so $I \approx_A I$.

We will additionally need an appropriate Lascar strong type generalization of Ehrenfeucht–Mostowski type.

Definition 4.5. Given two infinite A-indiscernible sequences I and J, we say that I and J have the same *Lascar–Ehrenfeucht–Mostowski type* (or *LEM type*) *over* A, written $I \equiv^{\mathrm{LEM}}_A J$, if there is some $J' \equiv^{\mathrm{L}}_A J$ such that $I + J'$ is A-indiscernible.

To see that the name is justified, note that two infinite A-indiscernible sequences I and J have the same Ehrenfeucht–Mostowski type over A if and only if there is a $J' \equiv_A J$ such that $I + J'$ is A-indiscernible.

Lemma 4.6. *For any infinite order types O and O', $I \approx_A J$ if and only if there are $K_0, L_0, K_1, \dots, L_{n-1}, K_n$ such that*

- *$K_0 = I$ and $K_n = J$,*
- *for $0 < i < n$, K_i is a sequence of order type O,*
- *for $i < n$, L_i is a sequence of order type O', and*
- *for $i < n$, $K_i + L_i$ and $K_{i+1} + L_i$ are A-indiscernible.*

Proof. The $\Leftarrow$ direction is obvious.

For the $\Rightarrow$ direction, we will proceed by induction. First assume that $I \sim_A J$. If $I + J$ is A-indiscernible, then find L of order type O' such that $I + J + L$ is A-indiscernible. We then have that $I + L$ and $J + L$ are A-indiscernible. If $J + I$ is A-indiscernible, then find L of order type O' such that $J + I + L$ is A-indiscernible. We then have that $I + L$ and $J + L$ are A-indiscernible.

Now assume that we know the statement holds for any I and J such that there is a sequence $I'_0, \dots, I'_n$ with $I'_0 = I$, $I'_n = J$, and $I'_i \sim_A I'_{i+1}$ for each $i < n$. Now assume that there is a sequence $I'_0, \dots, I'_{n+1}$ with $I'_0 = I$, $I'_{n+1} = J$, and $I'_i \sim_A I'_{i+1}$ for each $i \le n$. Apply the induction hypothesis to get $K_0, L_0, K_1, \dots, L_{m-1}, K_m$ satisfying the properties in the statement of the lemma with $K_0 = I$ and $K_m = I'_n$. Now since $I'_n \sim_A I'_{n+1} = J$, we can apply the $n = 1$ case to get L_m such that $I'_n + L_m$ and $I'_{n+1} + L_m$ are both A-indiscernible. By compactness, we can find K^*_m of order type O such that $K^*_m + L_m$ and $K^*_m + L_{m-1}$ are both A-indiscernible. We then have that $K_0, L_0, K_1, L_1, \dots, K_{m-1}, L_{m-1}, K^*_m, L_m, K_{m+1}$ is the require sequence, where $K_{m+1} = J$. □

Proposition 4.7. *Fix a set of hyperimaginary parameters A.*

1. *$\equiv^{\mathrm{LEM}}_A$ is an equivalence relation on the class of infinite A-indiscernible sequences.*
2. *If I and J have the same order type, then $I \equiv^{\mathrm{L}}_A J$ if and only if $I \equiv^{\mathrm{LEM}}_A J$.*
3. *If $I \equiv^{\mathrm{LEM}}_A J$, then $I \equiv^{\mathrm{EM}}_A J$.*
4. *If $I \approx_A J$, then $I \equiv^{\mathrm{LEM}}_A J$.*

Proof. Recall the following fact: If I and J have the same order type and $I+J$ is A-indiscernible, then $I \equiv^{\mathrm{L}}_A J$.[12]

(1). First, to see that $\equiv^{\mathrm{LEM}}_A$ is reflexive, note that if I is an infinite A-indiscernible sequence, then any infinite A-indiscernible extension $I+I'$ will witness that $I \equiv^{\mathrm{LEM}}_A I$. To see that $\equiv^{\mathrm{LEM}}_A$ is symmetric, assume that $I \equiv^{\mathrm{LEM}}_A J$, and let J' be as in the definition of $\equiv^{\mathrm{LEM}}$. Find I' such that $IJ' \equiv^{\mathrm{L}}_A I'J$. Then extend $I'+J$ to $I'+J+I''$, where I'' has the same order type as I. We then have that $I'' \equiv^{\mathrm{L}}_A I' \equiv^{\mathrm{L}}_A I$, so $J \equiv^{\mathrm{LEM}}_A I$. To see that $\equiv^{\mathrm{LEM}}_A$ is transitive, assume that $I \equiv^{\mathrm{LEM}}_A J$ and $J \equiv^{\mathrm{LEM}}_A K$. Let this be witnessed by J' and K' such that $I+J'$ and $J+K'$ are A-indiscernible. Find K'' with the same order type as K such that $I+J'+K''$ is A-indiscernible. Then find K^* such that $J'K'' \equiv^{\mathrm{L}}_A JK^*$. Note that both $J+K^*$ and $J+K'$ are A-indiscernible. By compactness, we can find K^{**} of the same order type as K such that $K^{**}+J+K^*$ and $K^{**}+J+K'$ are both A-indiscernible. By the above fact, we then have that $K^* \equiv^{\mathrm{L}}_A K^{**} \equiv^{\mathrm{L}}_A K'$. Finally, $K' \equiv^{\mathrm{L}}_A K$ by assumption, so we have that $K'' \equiv^{\mathrm{L}}_A K$ and therefore that $I \equiv^{\mathrm{LEM}}_A K$.

(2) is immediate from the fact. (3) is obvious.

For (4), it is sufficient to show that $I \sim_A J \Rightarrow I \equiv^{\mathrm{LEM}}_A J$. This follows immediately from the fact that $I \equiv^{\mathrm{L}}_A I$ and $J \equiv^{\mathrm{L}}_A J$. $\square$

Now we will see that total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequences over A are precisely those which are "as generic as possible" in terms of $\approx_A$ (i.e., their $\equiv^{\mathrm{LEM}}_A$-equivalence class decomposes into a single $\approx_A$-equivalence class).

Theorem 4.8. *For any A-indiscernible sequence $(b_i)_{i<\omega}$ (with A a set of hyperimaginary parameters), the following are equivalent.*

1. *$b_{<\omega}$ is a total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequence over A.*
2. *There exists a pair of infinite sequences I and J (of any, possibly distinct order types) such that $I+J \equiv^{\mathrm{EM}}_A b_{<\omega}$ and $I \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A J$.*
3. *For any K, $K \approx_A b_{<\omega}$ if and only if $K \equiv^{\mathrm{LEM}}_A b_{<\omega}$.*
4. *$[b_{<\omega}]_{\approx_A} \in \mathrm{bdd}^{\mathrm{u}}(A)$.*

Proof. (1)⇒(2). This is immediate from the definition.

(2)⇒(3). First note that if $K \approx_A b_{<\omega}$, then $K \equiv^{\mathrm{LEM}}_A b_{<\omega}$ by Proposition 4.7. Let I and J be as in the statement of (2). By compactness, we may find $I' \equiv_A b_{<\omega}$ such that $I'+I+J$ is A-indiscernible. By applying an automorphism fixing A, we

[12] To see this, assume that I and J have the same order type and $I+J$ is A-indiscernible for some set of hyperimaginary parameters. Let M be a model with $A \subseteq \mathrm{bdd}^{\mathrm{heq}}(M)$. We can find an M-indiscernible sequence $I'+J'$ finitely based on $I+J$. In particular, this will have $I'+J' \equiv_A I+J$. Therefore we can find a model $M' \equiv_A M$ such that $I+J$ is M'-indiscernible. We then have that $I \equiv_{M'} J$, whereby $I \equiv^{\mathrm{L}}_A J$.

may assume that $b_{<\omega}+I+J$ is A-indiscernible. Fix K such that $K \equiv^{\mathrm{LEM}}_A b_{<\omega}$. By compactness, we can find a $K' \equiv_A K$ such that $b_{<\omega}+I+K'+J$ is A-indiscernible. We have that $K \equiv^{\mathrm{LEM}}_A b_{<\omega} \sim_A K'$ and therefore $K \equiv^{\mathrm{L}}_A K'$ by Proposition 4.7. Let $a_E \in \mathrm{bdd}^{\mathrm{u}}(AI)$ be an ultraimaginary satisfying $\mathrm{dcl}^{\mathrm{u}}(a_E)=\mathrm{bdd}^{\mathrm{u}}(AI)$. Likewise, let $b_F \in \mathrm{bdd}^{\mathrm{u}}(AJ)$ be an ultraimaginary satisfying $\mathrm{dcl}^{\mathrm{u}}(b_F)=\mathrm{bdd}^{\mathrm{u}}(AJ)$.[13] Since $\mathrm{dcl}^{\mathrm{u}}(a_F)\cap \mathrm{dcl}^{\mathrm{u}}(b_F)=\mathrm{bdd}(A)$, we have that $K \equiv_{\mathrm{dcl}^{\mathrm{u}}_\lambda(Ia_F)\cap \mathrm{dcl}^{\mathrm{u}}_\lambda(Jb_F)} K'$ for all λ. Therefore, by Proposition 1.10, we can find a sequence $(I^i a^i J^i b^i K^i)_{i\leq n}$ satisfying that $I^0a^0=Ia$, $J^0b^0=Jb$, $K^0=K'$, $K^n=K$, and for each $i<n$,

- if i is even, then $I^ia^i=I^{i+1}a^{i+1}$ and $J^ib^iK^i \equiv_A J^{i+1}b^{i+1}K^{i+1}$ and
- if i is odd, then $J^ib^i=J^{i+1}b^{i+1}$ and $I^ia^iK^i \equiv_A I^{i+1}a^{i+1}K^{i+1}$.

By induction, we have that $I^i+K^i+J^i$ is A-indiscernible for each $i \leq n$. We therefore have that

$$K'=K^0 \sim_A I^0 \sim_A J^1 \sim_A I^2 \sim_A J^3 \sim_A \cdots \sim_A L \sim_A K^n = K,$$

where L is either I^n or J^n. Therefore $K' \approx_A K$.

(3)⇒(1). Assume that for any $K \equiv^{\mathrm{LEM}}_A b_{<\omega}$, $K \approx_A b_{<\omega}$. Let I and J be infinite sequences satisfying $I+J \equiv^{\mathrm{EM}}_A b_{<\omega}$. By applying an automorphism fixing A to $I+J$, we may assume that $b_{<\omega}+I+J$ is A-indiscernible. Fix some $I' \equiv^{\mathrm{L}}_A I$. We have that $I' \equiv^{\mathrm{LEM}}_A b_{<\omega}$, which by assumption implies that $I' \approx_A b_{<\omega}$. Since $b_{<\omega} \sim_A I$, we have that $I \approx_A I'$. By Lemma 4.6, we can find $K_0, L_0, K_1, L_1, \dots, L_{n-1}, K_n$ such that $K_0=I$, $K_n=I'$, L_0 has the same order type as J, K_i has the same order type as I for each $i \leq n$, L_i has the same order type as J for each $i<n$, and K_i+L_i and $K_{i+1}+L_i$ are A-indiscernible for each $i<n$. Let $K_{-1}=I$ and $L_{-1}=J$. We now have that for each nonnegative $i<n$, $K_{i-1} \equiv^{\mathrm{L}}_{AL_{i-1}} K_i$ and $L_{i-1} \equiv^{\mathrm{L}}_{AK_i} L_i$.[14] Therefore $K_{-1}, L_{-1}, K_0, L_0, \dots, L_{n-1}, K_n$ is precisely the kind of sequence needed to apply Proposition 2.4 (with the indices shifted down by 1). Since we can do this for any $I' \equiv^{\mathrm{L}}_A I$, we have that $I \mathop{\smile\hspace{-0.9em}\mid}^{\mathrm{bu}}_A J$.

(3)⇒(4). Let x be a tuple of variables in the same sorts as $b_{<\omega}$. There are at most $2^{|Ab_{<\omega}|+|T|}$ many Lascar strong types in x over A. (3) implies therefore that there are at most $2^{|Ab_{<\omega}|+|T|}$ many $\approx_A$ classes with representatives that realize $\mathrm{tp}(b_{<\omega}/A)$. Therefore $[c_{<\omega}]_{\approx_A} \in \mathrm{bdd}^{\mathrm{u}}(A)$ for any $c_{<\omega} \equiv_A b_{<\omega}$ and so a fortiori $[b_{<\omega}]_{\approx_A} \in \mathrm{bdd}^{\mathrm{u}}(A)$.

(4)⇒(3). Let $I \equiv^{\mathrm{LEM}}_A b_{<\omega}$. Find I' such that $I \equiv^{\mathrm{L}}_A I'$ and $b_{<\omega}+I'$ is A-indiscernible. Since $[b_{<\omega}]_{\approx_A} \in \mathrm{bdd}^{\mathrm{u}}(A)$, we must have, by Proposition 1.4, that there are at most $2^{|Ab_{<\omega}|+|T|}$ conjugates of $[b_{<\omega}]_{\approx_A}$ under $\mathrm{Aut}(\mathbb{M}/A)$. For any $I'' \equiv_A I'$, we can find

[13] We can take a_E to be $[\![\mathrm{Autf}(\mathbb{M}/AI)]\!]$ and b_F to be $[\![\mathrm{Autf}(\mathbb{M}/AJ)]\!]$ by Definition 1.9 and Proposition 1.12.

[14] For $i=0$, we have that $K_{-1} \equiv^{\mathrm{L}}_{AL_{-1}} K_0$ trivially, since $K_{-1}=I=K_0$.

$c_{<\omega} \equiv_A b_{<\omega}$ such that $I'' \sim_A c_{<\omega}$. Therefore there are at most $2^{|Ab_{<\omega}|+|T|}$ conjugates of $[I']_{\approx_A}$ under $\mathrm{Aut}(\mathbb{M}/A)$ as well, and so $[I']_{\approx_A} \in \mathrm{bdd}^{\mathrm{u}}(A)$ by Proposition 1.4 again. By Proposition 1.13, there must be an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}/A, [I']_{\approx A})$ such that $\sigma \cdot I' = I$. Therefore $[I']_{\approx_A} = [I]_{\approx_A}$ and hence $I \approx_A b_{<\omega}$. □

Building ((weakly) total) $\downfree^{\mathrm{bu}}$-Morley sequences. Given that $\downfree^{\mathrm{bu}}$ satisfies full existence, an immediate, familiar Erdős-Rado argument gives that $\downfree^{\mathrm{bu}}$-Morley sequences exist, but in the end we will need a technical strengthening of this result.

Proposition 4.9. *If $(b_f)_{f\in\mathcal{T}_\alpha}$ is a family of real elements that is s-indiscernible over a set of hyperimaginaries A, then there is a family $(c_f)_{f\in\mathcal{F}_{\alpha+1}}$ such that*

- *$c_{\in\mathcal{F}_{\alpha+1}}$ is s-indiscernible over A,*
- *$c_{\iota_{\alpha,\alpha+1}(f)} = b_f$ for each $f \in \mathcal{T}_\alpha$, and*
- *the sequence $(c_{\trianglerighteq\langle i\rangle})_{i<\omega}$ is an $\downfree^{\mathrm{bu}}$-Morley sequence over A.*

Proof. Let κ be sufficiently large to apply Erdős-Rado to a sequence of tuples of the same length as $b_{\in\mathcal{T}_\alpha}$ over the set A.

Let $\gamma(0)=\alpha$. Let $c^0_f = b_f$ for all $f \in \mathcal{T}_{\gamma(0)} = \mathcal{T}_\alpha$. Let $g_0 = \varnothing$ (as an element of $\mathcal{T}_\alpha$).

At successor stage $\beta+1$, assume we have $(c^\beta_f)_{\mathcal{T}_{\gamma(\beta)}}$ which is s-indiscernible over A and which satisfies $c^\beta_{\iota_{\gamma(\delta),\gamma(\beta)}(f)} = c^\delta_f$ for all $\delta < \beta$. By Lemma 3.5, we can build a family $(c^{\beta+1}_f)_{\mathcal{T}_{\gamma(\beta+1)}}$ (for some successor ordinal $\gamma(\beta+1) > \gamma(\beta)$) such that

- $(c^{\beta+1}_f)_{f\in\mathcal{T}_{\gamma(\beta+1)}}$ is s-indiscernible over A,
- for each $f \in \mathcal{T}_{\gamma(\beta)}$, $c^\beta_f = c^{\beta+1}_{\iota_{\gamma(\beta),\gamma(\beta+1)}(f)}$, and
- $c^{\beta+1}_{\trianglerighteq\zeta^{\gamma(\beta+1)}_{\gamma(\beta)}} \downfree^{\mathrm{bu}}_A c^{\beta+1}_{\trianglerighteq 1\frown\zeta^{\gamma(\beta+1)-1}_{\gamma(\beta)}}$.

Let $g_{\beta+1} \in \mathcal{T}^*_{\gamma(\beta+1)}$ be $1 \frown \zeta_\alpha^{\gamma(\beta+1)-1}$. Note that $g_{\beta+1} \trianglerighteq h$. Also note that by construction we have that

$$c^{\beta+1}_{\trianglerighteq g_{\beta+1}} \downfree^{\mathrm{bu}}_A \{c^{\beta+1}_{\trianglerighteq\iota_{\gamma(\delta),\gamma(\beta+1)}(g_\delta)} : \delta \in (\beta+1)\setminus\lim(\beta+1)\},$$

since $\iota_{\gamma(\delta),\gamma(\beta+1)}(g_\delta) \trianglerighteq \zeta^{\gamma(\beta+1)}_{\gamma(\beta)}$ for all nonlimit $\delta < \beta+1$.

At limit stage β, let $\gamma(\beta) = \sup_{\delta<\beta}\gamma(\delta)$ and let $(c^\beta_f)_{f\in\mathcal{T}_{\gamma(\beta)}}$ be the direct limit of $(c^\delta_f)_{f\in\mathcal{T}_{\gamma(\delta)}}$ for $\delta < \beta$. Leave g_β undefined.

Stop once we have $(c^\kappa_f)_{f\in\mathcal{T}_{\gamma(\kappa)}}$. Consider the sequence $(c^\kappa_{\trianglerighteq\iota_{\gamma(\beta),\gamma(\kappa)}(g_\beta)})_{\beta\in\kappa\setminus\lim\kappa}$.[15] By our choice of κ and a standard application of the Erdős-Rado theorem, we can find a family $(c_f)_{f\in\mathcal{F}_{\alpha+1}}$ such that the sequence $(c_{\trianglerighteq\langle i\rangle_\alpha})_{i<\omega}$ is A-indiscernible and

[15] We write $\lim\alpha$ for the set of limit ordinals in α.

for every increasing tuple $\bar{\imath} < \omega$, there is $\bar{\beta} \in \kappa \setminus \lim \kappa$ such that $c_{\trianglerighteq \langle i_0 \rangle_\alpha} \dots c_{\trianglerighteq \langle i_k \rangle_\alpha} \equiv_A c^\kappa_{\trianglerighteq \iota_{\gamma(\beta_0),\gamma(\kappa)}(g_{\beta_0})} \dots c^\kappa_{\trianglerighteq \iota_{\gamma(\beta_k),\gamma(\kappa)}(g_{\beta_k})}$.

In particular, note that this implies that

$$c_{\trianglerighteq \langle i \rangle_\alpha} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A \{c_{\trianglerighteq \langle j \rangle_\alpha} : j < i\}$$

for every $i < \omega$. Clearly by applying an automorphism, we may assume that $c_{\iota_{\alpha,\alpha+1}(f)} = b_f$ for each $f \in \mathcal{T}_\alpha$, so all we need to do is show that the family $c_{\in \mathcal{F}_{\alpha+1}}$ is s-indiscernible over A.

Since the sequence $(c_{\trianglerighteq \langle i \rangle_\alpha})_{i<\omega}$ is A-indiscernible, it is sufficient, by induction, to show the following statement: For any sequence $\bar{f}_0, \bar{f}_1, \dots, \bar{f}_k, \dots, \bar{f}_\ell$ of tuples of elements of $\mathcal{F}_{\alpha+1}$ satisfying $\bar{f}_i \trianglerighteq \langle i \rangle_\alpha$ for all $i \leq \ell$ and any $\bar{h} \trianglerighteq \langle k \rangle_\alpha$ such that $\bar{f}_k$ and $\bar{h}$ realize the same quantifier-free type, we have that $c_{\bar{f}_k}$ and $c_{\bar{h}}$ realize the same type over $Ac_{\bar{f}_0} \dots c_{\bar{f}_{k-1}} c_{\bar{f}_{k+1}} \dots c_{\bar{f}_\ell}$.

So let $\bar{f}_0, \dots, \bar{f}_\ell$ and $\bar{h}$ be as in the statement. By construction, there are $\beta_0, \dots, \beta_\ell$ such that $c_{\trianglerighteq \langle i \rangle_\alpha} \equiv_A c^\kappa_{\trianglerighteq \iota_{\gamma(\beta_i),\gamma(\kappa)}(g_{\beta_i})}$ for each $i \leq \ell$. Let $\bar{f}'_0, \dots, \bar{f}'_\ell, \bar{h}'$ be the corresponding elements of $\mathcal{T}_{\gamma(\kappa)}$. (So, in particular, $\bar{f}'_i \trianglerighteq g_{\beta_i}$ for each $i \leq \ell$ and $\bar{h}' \trianglerighteq g_{\beta_k}$). We now have that $\bar{f}'_k$ and $\bar{h}'$ realize the same quantifier-free type. Therefore, by the s-indiscernible of $c^\kappa_{\in \mathcal{T}_{\gamma(\kappa)}}$, we have that $c^\kappa_{\bar{f}'_k}$ and $c^\kappa_{\bar{h}'}$ realize the same type over $Ac^\kappa_{\bar{f}'_0} \dots c^\kappa_{\bar{f}'_{k-1}} c^\kappa_{\bar{f}'_{k+1}} \dots c^\kappa_{\bar{f}'_\ell}$. From this the required statement follows, and we have that $c_{\in \mathcal{F}_{\alpha+1}}$ is s-indiscernible over A. $\square$

Corollary 4.10. *For any set of hyperimaginaries A and any real tuple b, there is an $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley sequence $(b_i)_{i<\omega}$ over A with $b_0 = b$.*

Proof. Apply Proposition 4.9 to the tree $(b_f)_{f \in \mathcal{T}_0}$ defined by $b_\varnothing = b$.[16] $\square$

The order type ω is essential, however; Erdős-Rado only guarantees the existence of sequences that satisfy the relevant condition on finite tuples. Fortunately, this is more than sufficient for the following weak "chain condition".

Lemma 4.11. *If $(b_i)_{i<\omega}$ is an $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley sequence over A that is moreover Ac-indiscernible, then $b_0 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A c$.*

Proof. Fix λ. Let $\mu = |\mathrm{bdd}^{\mathrm{u}}_\lambda(Ac) \setminus \mathrm{bdd}^{\mathrm{u}}_\lambda(A)|$. Extend $b_{<\omega}$ to $(b_i)_{i<\mu^+}$. We still have that for any $i < j < \mu^+$, $b_i \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A b_j$ (since this is only a property of $\mathrm{tp}(b_i b_j / A)$). Therefore the sets $\mathrm{bdd}^{\mathrm{u}}_\lambda(Ab_i) \setminus \mathrm{bdd}^{\mathrm{u}}_\lambda(A)$ are pairwise disjoint. Since there are μ^+ many of them, one of them must be disjoint from $\mathrm{bdd}^{\mathrm{u}}_\lambda(Ac) \setminus \mathrm{bdd}^{\mathrm{u}}_\lambda(A)$. Therefore by indiscernibility, we must have $b_0 \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A c$. $\square$

We will not use the following corollary of Lemma 4.11, but it is worth pointing out.

---

[16] This can also be proven directly by the standard argument for the existence of Morley sequences.

Corollary 4.12. *If I is a total $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequence over A that is Ac-indiscernible, then $I \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A c$.*

Proof. Extend I to an Ac-indiscernible sequence $I_0 + I_1 + I_2 + \dots$ with $I_0 = I$. Since I is totally $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley, we have that $(I_i)_{i<\omega}$ is an $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequence over A. So by Lemma 4.11, we have $I = I_0 \mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}_A c$. □

Parts (2) and (3) of following definition are equivalent to [Kim et al. 2014, Definitions 2.1, 3.4] in our context; this formulation is used implicitly in [Kaplan and Ramsey 2020] and its equivalence to the standard definition is discussed in [Kaplan and Ramsey 2020, Remark 5.8]. The rest of it is based on [Kaplan and Ramsey 2020, Definition 5.7], although we have had to modify the definition of restriction slightly in order to deal with limit ordinals more smoothly.

Definition 4.13. Fix a family $(b_f)_{f\in\mathcal{T}_\alpha}$.

(1) For $w \subseteq \alpha$, the *restriction of $\mathcal{T}_\alpha$ to the set of levels w* is given by

$$\mathcal{T}_\alpha\restriction w = \{f \in \mathcal{T}_\alpha : \min\operatorname{dom}(f) \in w,\ \ \beta \in \operatorname{dom}(f) \setminus w \Rightarrow f(\beta) = 0\}.$$

(2) A family $(b_f)_{f\in\mathcal{T}_\alpha}$ is *str-indiscernible over A* if it is s-indiscernible over A and satisfies that for any finite $w, v \subseteq \alpha \setminus \lim\alpha$ with $|w| = |v|$, $b_{\in\mathcal{T}_\alpha\restriction w}$ and $b_{\in\mathcal{T}_\alpha\restriction v}$ realize the same type over A (where we take $b_{\in\mathcal{T}_\alpha\restriction w}$ to be enumerated according to $<_{\mathrm{lex}}$, which is a well-ordering on $\mathcal{T}_\alpha\restriction w$ for finite w).

(3) We say that $b_{\in\mathcal{T}_\alpha}$ is *$\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out over A* if for any $f \in \mathcal{T}^*_\alpha$ (with $\operatorname{dom}(f) = [\beta+1, \alpha)$ for some $\beta < \alpha$), the sequence $(b_{\trianglerighteq f\frown i})_{i<\omega}$ is an $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequence over A.

(4) We say that $b_{\in\mathcal{T}_\alpha\restriction w}$ is *$\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out over A* if for any $f \in \mathcal{T}^*_\alpha$ (with $\operatorname{dom}(f) = [\beta+1, \alpha)$ for some $\beta < \alpha$ and satisfying that $(f \frown i)_{i<\omega}$ is a sequence of elements of $\mathcal{T}_\alpha\restriction w$), the sequence $(b_{\trianglerighteq f\frown i})_{i<\omega}$ is an $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequence over A (where we interpret b_f as $\varnothing$ if $f \notin \mathcal{T}_\alpha\restriction w$).

(5) $b_{\in\mathcal{T}_\alpha}$ is an *$\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-Morley tree over A* if it is $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out and str-indiscernible over A.

Note that if $b_{\in\mathcal{T}_\alpha}$ is $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out over A, then any restriction $b_{\in\mathcal{T}_\alpha\restriction w}$ is also $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out over A (even for infinite w). Also note that, by a basic compactness argument, if α is infinite and $(b_f)_{f\in\mathcal{T}_\alpha}$ is str-indiscernible over A, then for any β, we can find a tree $(c_f)_{f\in\mathcal{T}_\beta}$ which is str-indiscernible over A such that for any $w \in [\alpha]^{<\omega}$ and $v \in [\beta]^{<\omega}$ with $|w| = |v|$, $b_{\in\mathcal{T}_\alpha\restriction w} \equiv_A c_{\in\mathcal{T}_\beta\restriction v}$.

Proposition 4.14. *For any set of hyperimaginaries A, real tuple b, and κ, there is a tree $(b_f)_{f\in\mathcal{T}_\kappa}$ that is $\mathop{\smile\!\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out and s-indiscernible over A such that for each $f \in \mathcal{T}_\kappa$, $b_f \equiv_A b$.*

Proof. Let $(b^0_f)_{f \in \mathcal{T}_0}$ be defined by $b^0_\varnothing = b$. This is vacuously $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out and s-indiscernible over A.

At successor stage $\alpha + 1$, given $(b^\alpha_f)_{f \in \mathcal{T}_\alpha}$ which is $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out and s-indiscernible by Proposition 4.9, we can find an extension $(b^{\alpha+1}_f)_{f \in \mathcal{F}_{\alpha+1}}$ satisfying $b^{\alpha+1}_{\iota_{\alpha,\alpha+1}(f)} = b^\alpha_f$ for all $f \in \mathcal{T}_\alpha$ such that $b^{\alpha+1}_{\in \mathcal{F}_{\alpha+1}}$ is s-indiscernible over A and $(b^{\alpha+1}_{\trianglerighteq \langle i \rangle_\alpha})_{i<\omega}$ is an $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequence over A. By Fact 3.3, we can find $b^{\alpha+1}_\varnothing \equiv_A b$ such that the tree $(b^{\alpha+1}_f)_{f \in \mathcal{T}_{\alpha+1}}$ is s-indiscernible over A. By construction, we now have that $(b^{\alpha+1}_f)_{f \in \mathcal{T}_{\alpha+1}}$ is $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out over A.

At limit stage α, let $(b^\alpha_f)_{f \in \mathcal{T}_\alpha}$ be the direct limit of $(b^\beta_f)_{f \in \mathcal{T}_\beta}$ for $\beta < \alpha$. It is immediate from the definitions that $b^\alpha_{\in \mathcal{T}_\alpha}$ is $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out and s-indiscernible over A.

Once we have constructed $(b^\kappa_f)_{f \in \mathcal{T}_\kappa}$, let $b_f = b^\kappa_f$ for each $f \in \mathcal{T}_\kappa$. We have that $b_{\in \mathcal{T}_\kappa}$ is the required tree by induction. □

By the same argument as in [Kaplan and Ramsey 2020, Lemma 5.10], we get the following.

Fact 4.15. *Fix a set of real parameters A, and let $(b_f)_{f \in \mathcal{T}_\kappa}$ be a family of tuples of real parameters of the same length that is s-indiscernible over A. If $\kappa \geq \beth_{\lambda^+}(\lambda)$ (where $\lambda = 2^{|Ab_f|+|T|}$), then there is an str-indiscernible tree $(c_f)_{f \in \mathcal{T}_\omega}$ such that*

$$\text{for any } w \in [\omega]^{<\omega}, \text{ there is } v \in [\kappa]^{<\omega} \text{ such that } (b_f)_{f \in \mathcal{T}_\kappa \restriction v} \equiv_A (c_f)_{f \in \mathcal{T}_\omega \restriction w}. \quad (*)_A$$

Note that Fact 4.15 generalizes to continuous logic by the same soft argument as in the discussion after Fact 3.3.

Lemma 4.16. *Suppose that a family of tuples of real elements $(b_f)_{f \in \mathcal{T}_\kappa}$ is $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out and s-indiscernible over a set of hyperimaginaries A with all b_f tuples of the same length. If $\kappa \geq \beth_{\lambda^+}(\lambda)$ (where $\lambda = 2^{|Ab_f|+|T|}$), then there is an $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-Morley tree $(c_f)_{f \in \mathcal{T}_\omega}$ over A such that condition $(*)_A$ from Fact 4.15 holds.*

Proof. Find a model M with $|M| \leq |A| + \aleph_0$ such that $A \subseteq \mathrm{bdd}^{\mathrm{heq}}(M)$. Apply Fact 4.15 with M as the base to the family $(b_f)_{f \in \mathcal{T}_\kappa}$ to get a tree $(c_f)_{f \in \mathcal{T}_\omega}$ that is str-indiscernible over M and satisfies $(*)_M$. This is enough to imply that $c_{\in \mathcal{T}_\omega}$ is str-indiscernible over A and satisfies $(*)_A$. Furthermore, since the tree $c_{\in \mathcal{T}_\omega}$ has height ω and since $b_{\in \mathcal{T}_\kappa}$ is $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out over A, $(*)_A$ implies that $c_{\in \mathcal{T}_\omega}$ is $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-spread-out over A. Therefore $c_{\in \mathcal{T}_\omega}$ is an $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-Morley tree over A. □

Proposition 4.17. *If $(b_f)_{f \in \mathcal{T}_\omega}$ is a family of tuples of real elements that is an $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-Morley tree over a set of hyperimaginaries A, then $(b_{\zeta^\omega_\beta})_{\beta<\omega}$ is a weakly total $\mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}$-Morley sequence over A.*

Proof. Fix a linear order O. Let $c_\alpha = b_{\zeta^\omega_\alpha}$ for each $\alpha < \omega$.

For each positive $n < \omega$ and each $i < j < \omega$, we have that $b_{\trianglerighteq \zeta^\omega_n \frown i} \mathop{\smile\!\!\!\!\mid}^{\mathrm{bu}}_A b_{\trianglerighteq \zeta^\omega_n \frown j}$ and that the sequence $(b_{\trianglerighteq \zeta^\omega_n \frown i})_{i<\omega}$ is $Ac_{\geq n}$-indiscernible. By compactness, we can find $(c_i)_{i\in O}$ such that $(c_i)_{i\in\omega+O}$ is A-indiscernible and such that $(b_{\trianglerighteq\zeta_n^\omega\frown i})_{i<\omega}$ is $Ac_{\in[n,\omega)+O}$-indiscernible for each $n<\omega$.

Therefore, by Lemma 4.11, we have that $c_{<n} \mathop{\smile\!\!\!\!\!\!\;|}\nolimits^{\mathrm{bu}}_A c_{\in[n,\omega)+O}$. Hence, $(b_{\zeta_\beta^\omega})_{\beta<\omega}$ is a weakly total $\mathop{\smile\!\!\!\!\!\!\;|}\nolimits^{\mathrm{bu}}$-Morley sequence. $\square$

Corollary 4.18. *For any set of hyperimaginaries A and tuple of real elements b, there is an A-indiscernible sequence $(b_i)_{i<\omega}$ with $b_0=b$ such that for any $b'\equiv^{\mathrm{L}}_A b$ and $n<\omega$, there are $I_0, J_0, I_1, J_1 \dots, J_{k-1}, I_k$ with*

- *b the first element of I_0,*
- *b' the first element of I_k,*
- *$|I_i|=n$ for all $i\le k$,*
- *J_i infinite for all $i<k$, and*
- *I_i+J_i and $I_{i+1}+J_i$ realizing the same EM-type over A as $b_{<\omega}$ for all $i<k$.*

We can also arrange it so that I_i is infinite for all $i\le k$, $|J_i|=n$ for all $i<k$, and I_i+J_i and $I_{i+1}+J_i$ realize the same EM-type over A as $b_{<\omega}$ in the reverse order for all $i<k$ (with the same choice of $b_{<\omega}$ but possibly a different k).

Proof. By Lemma 4.16 and Proposition 4.17, we can find a sequence $(b_i)_{i<\omega}$ with $b_0=b$ that is a weakly total $\mathop{\smile\!\!\!\!\!\!\;|}\nolimits^{\mathrm{bu}}$-Morley sequence over A. Fix $n<\omega$, and write $b_{<\omega}$ as $I+J$ with $|I|=n$. By repeating the proof of Proposition 4.3, we get the required configuration of I_i's and J_i's.

For the final statement, by compactness, we can find an indiscernible sequence K of order type ω which has b as its first element and realizes the reverse of the EM-type of $b_{<\omega}$ over A. Fix an $n<\omega$. If we partition K as $I+J$ where $|J|=n$ and again repeat the proof of Proposition 4.3, we get the second required configuration of I_i's and J_i's. $\square$

To go further, we will need the following fact from [Silver 1971]. Recall that the statement $\kappa\to(\alpha)^{<\omega}_\gamma$ means that whenever $f:[\kappa]^{<\omega}\to\gamma$ is a function, there is a set $X\subseteq\kappa$ of order type α such that for each $n<\omega$, f is constant on $[X]^n$.

Fact 4.19 [Silver 1971, Chapter 4]. *For any limit ordinal α, if κ is the smallest cardinal satisfying $\kappa\to(\alpha)^{<\omega}_2$, then for any $\gamma<\kappa$, $\kappa\to(\alpha)^{<\omega}_\gamma$. Furthermore, κ is strongly inaccessible.*

The smallest cardinal λ satisfying $\lambda\to(\alpha)^{<\omega}_2$ is called the Erdős cardinal $\kappa(\alpha)$. In the specific case of $\alpha=\omega$, we will also need the following lemma.

Lemma 4.20. *If $\kappa\to(\omega)^{<\omega}_\gamma$, then $(\gamma^\kappa)^+\to(\omega+1)^{<\omega}_\gamma$. In particular, if $\kappa(\omega)$ exists, then $(2^{\kappa(\omega)})^+\to(\omega+1)^{<\omega}_\gamma$ for any $\gamma<\kappa(\omega)$.*

Proof. Fix a set X of cardinality $(\gamma^\kappa)^+$ and a coloring $f:[X]^{<\omega}\to\gamma$. Fix an ordering $(x_\alpha)_{\alpha<(\gamma^\kappa)^+}$ of X. Recall that a subset $Y\subseteq X$ is *end-homogeneous* if for any $\delta_0<\cdots<\delta_{n-1}<\alpha<\beta<(\gamma^\kappa)^+$, $f(\{x_{\delta_0},\dots,x_{\delta_{n-1}},x_\alpha\})=f(\{x_{\delta_0},\dots,x_{\delta_{n-1}},x_\beta\})$.

By [Erdős et al. 1984, Lemma 15.2], there is an end-homogeneous set $Y \subseteq X$ of order type $\kappa+1$. Let $(y_\alpha)_{\alpha<\kappa+1}$ be an enumeration of Y in order. Let $g(A) = f(A \cup \{y_\kappa\})$. By assumption, there is a g-homogeneous subset $Z \subseteq Y$ of order type ω. Therefore, by construction, $Z \cup \{y_\kappa\}$ is the required f-homogeneous subset of order type $\omega+1$.

The last statement follows from the fact that $\kappa(\omega)$ is strongly inaccessible and cardinal arithmetic (i.e., $2^{\kappa(\omega)} = \gamma^{\kappa(\omega)}$ for $\gamma > 1$ with $\gamma < \kappa(\omega)$). □

Lemma 4.21. *Suppose $(b_f)_{f\in\mathcal{T}_\lambda}$ is $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-spread-out and s-indiscernible over A with all b_f tuples of the same length. If $\lambda \to (\omega+1)^{<\omega}_{2^{|Ab|+|T|}}$, then there is a set $X \subseteq \lambda \setminus \lim\lambda$ with order type $\omega+1$ such that $b_{\in\mathcal{T}_\lambda\restriction X}$ is an $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley tree over A.*

Proof. Let t be the function on $[\lambda\setminus\lim\lambda]^{<\omega}$ that takes $w \in [\lambda\setminus\lim\lambda]^{<\omega}$ to $\mathrm{tp}(b_{\in\mathcal{T}_\lambda\restriction w}/A)$. By assumption, we can find $X \subset \lambda\setminus\lim\lambda$ of order type $\omega+1$ such that t is homogeneous on X. $b_{\in\mathcal{T}_\lambda\restriction X}$ is s-indiscernible over A and $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-spread-out over A, since these properties are both preserved by passing to restrictions. □

Theorem 4.22. *For any A and b in any theory T, if there is a cardinal λ satisfying $\lambda \to (\omega+1)^{<\omega}_{2^{|Ab|+|T|}}$, then there is a total $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley sequence $(b_i)_{i<\omega}$ over A with $b_0 = b$.*

In particular, it is enough if there is an Erdős cardinal $\kappa(\alpha)$ such that $|Ab|+|T| < \kappa(\alpha)$ (for any limit $\alpha \geq \omega$).

Proof. If the Erdős cardinal $\kappa(\alpha)$ exists and $|Ab|+|T| < \kappa(\alpha)$, then by Fact 4.19, we have $2^{|Ab|+|T|} < \kappa(\alpha)$ as well. Then if $\alpha = \omega$, we have that $(2^{\kappa(\omega)})^+ \to (\omega+1)^{<\omega}_{2^{|Ab|+|T|}}$ by Lemma 4.20. If $\alpha > \omega$, we clearly have $\kappa(\alpha) \to (\omega+1)^{<\omega}_{2^{|Ab|+|T|}}$ by Fact 4.19. So in any such case we have the required λ.

Let λ be a cardinal such that $\lambda \to (\omega+1)^{<\omega}_{2^{|Ab|+|T|}}$ holds. By Proposition 4.14, we can build a tree $(b_f)_{f\in\mathcal{T}_\lambda}$ that is s-indiscernible and $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-spread-out over A. By Lemma 4.21 and the choice of λ, we can extract an $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}$-Morley tree $(c_f)_{f\in\mathcal{T}_{\omega+1}}$ from this.

By compactness, we can extend this to a tree $(c_f)_{f\in\mathcal{T}_{\omega+\omega}}$ that is str-indiscernible over A. We still have that for any $i < j < \omega$,

$$c_{\trianglerighteq\zeta^{\omega+\omega}_{\omega+1}\frown i} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A c_{\trianglerighteq\zeta^{\omega+\omega}_{\omega+1}\frown j}$$

but now we also have that the $(c_{\trianglerighteq\zeta^{\omega+\omega}_{\omega+1}\frown i})_{i<\omega}$ is $A \cup \{c_{\zeta^{\omega+\omega}_{\omega+i}} : i < \omega\}$-indiscernible, by str-indiscernibility of the full tree $c_{\in\mathcal{T}_{\omega+\omega}}$. Therefore, by Lemma 4.11,

$$c_{\trianglerighteq\zeta^{\omega+\omega}_{\omega+1}\frown 0} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A \{c_{\zeta^{\omega+\omega}_{\omega+i}} : i < \omega\},$$

so in particular,

$$\{c_{\zeta^{\omega+\omega}_i} : i < \omega\} \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{bu}}_A \{c_{\zeta^{\omega+\omega}_{\omega+i}} : i < \omega\}.$$

Let $d_i = c_{\zeta_i^{\omega+\omega}}$ for each $i < \omega + \omega$. We have that $(d_i)_{i<\omega+\omega}$ is A-indiscernible. Furthermore, by Theorem 4.8, we have that $d_{<\omega}$ is a total $\mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}$-Morley sequence. By applying an automorphism, we get the required $b_{<\omega}$. □

So if we assume that for every λ, there is a κ such that $\kappa \to (\omega+1)^{<\omega}_{\lambda}$, we get that Lascar strong type is always witnessed by total $\mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}$-Morley sequences in the manner of Proposition 4.3.

The use of large cardinals in Theorem 4.22 leaves an obvious question.

Question 4.23. *Does the statement "for every A and b, there is a total $\mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}$-Morley sequence $(b_i)_{i<\omega}$ over A with $b_0 = b$" have any set-theoretic strength? What if we add cardinality restrictions, such as $|A| + |T| \le \aleph_0$ and $|b| < \aleph_0$?*

Total $\mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}$-Morley sequences in tame theories. Lemma 4.11 can be used to show that $\mathop{\smile\!\!\!\!\!\mid}^{\text{d}}$ implies $\mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}$ (where $b \mathop{\smile\!\!\!\!\!\mid}^{\text{d}}_A c$ means that $\operatorname{tp}(b/Ac)$ does not divide over A), something which was previously established for bounded hyperimaginary independence, $\mathop{\smile\!\!\!\!\!\mid}^{\text{b}}$, in [Conant and Hanson 2022, Corollary 4.13] and which was originally folklore for algebraic independence, $\mathop{\smile\!\!\!\!\!\mid}^{\text{a}}$.[17]

Proposition 4.24. *For any real elements A, b, and c, if $b \mathop{\smile\!\!\!\!\!\mid}^{\text{d}}_A c$, then $b \mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}_A c$.*

Proof. Let $(c_i)_{i<\omega}$ be an $\mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}$-Morley sequence over A with $c_0 = c$. Since $b \mathop{\smile\!\!\!\!\!\mid}^{\text{d}}_A c$, we may assume that $c_{<\omega}$ is Ab-indiscernible. Hence, by Lemma 4.11, $b \mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}_A c$. □

Corollary 4.25. *If $(b_i)_{i<\omega}$ is a (nondividing) Morley sequence over A, then it is a total $\mathop{\smile\!\!\!\!\!\mid}^{\text{bu}}$-Morley sequence over A.* □

In simple theories, we get the converse (Proposition 4.27). Recall that $B \mathop{\smile\!\!\!\!\!\mid}^{\text{b}}_A C$ means $\operatorname{bdd}^{\text{heq}}(AB) \cap \operatorname{bdd}^{\text{heq}}(AC) = \operatorname{bdd}^{\text{heq}}(A)$.

Lemma 4.26. *Let T be a simple theory. For any A, b, and c, $b \mathop{\smile\!\!\!\!\!\mid}^{\text{f}}_A C$ if and only if there is an AC-indiscernible sequence $(b_i)_{i<\omega}$ with $b_0 = b$ such that for any J and K with $J + K \equiv^{\text{EM}}_A b_{<\omega}$, $J \mathop{\smile\!\!\!\!\!\mid}^{\text{b}}_A K$.*

Proof. (The argument here is similar to the proof of [Adler 2005, Lemma 3.2], but we will give a proof for the sake of completeness.) If $b \mathop{\smile\!\!\!\!\!\mid}^{\text{f}}_A C$, then we can build an AC-indiscernible $\mathop{\smile\!\!\!\!\!\mid}^{\text{f}}$-Morley sequence $(b_i)_{i<\omega}$ over A with $b_0 = b$ (since T is simple). By some forking calculus, we have that $J \mathop{\smile\!\!\!\!\!\mid}^{\text{f}}_A K$ for any J and K with $J + K \equiv^{\text{EM}}_A b_{<\omega}$. Therefore, by [Conant and Hanson 2022, Corollary 4.13], $J \mathop{\smile\!\!\!\!\!\mid}^{\text{b}}_A K$ for any such J and K as well.

Conversely, assume that there is an AC-indiscernible sequence $(b_i)_{i<\omega}$ with $b_0 = b$ such that for any J and K with $J + K \equiv^{\text{EM}}_A b_{<\omega}$, $J \mathop{\smile\!\!\!\!\!\mid}^{\text{b}}_A K$. Let κ be a regular cardinal such that every type (in the same sort as C) does not fork over some set of cardinality less than κ. Let $(b_i)_{i<\kappa+\kappa^*}$ be an AC-indiscernible sequence

[17] There is an incorrect proof of this in the literature. To the author's knowledge, the first correct published proof of this is in [Conant and Hanson 2022, Theorem 4.11].

extending $b_{<\omega}$, where κ^* is an order-reversed copy of κ. Now we clearly have that $b_{<\kappa} \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{b}}_A b_{\in\kappa^*}$. By local character, there is a set $D \subseteq Ab_{<\kappa}$ with $|D| < \kappa$ such that $C \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_D Ab_{<\kappa}$. Since κ is regular, there is a $\lambda < \kappa$ such that $D \subseteq Ab_{<\lambda}$. Therefore, by base monotonicity, $C \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_{Ab_{<\lambda}} Ab_{<\kappa}$. Since $b_{\geq\lambda}$ is $Ab_{<\lambda}C$-indiscernible, we have that $C \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_{Ab_{<\lambda}} Ab_{\in\kappa+\kappa^*}$. Therefore, by base monotonicity again, $C \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_{Ab_{<\kappa}} Ab_{\in\kappa+\kappa^*}$. By the symmetric argument, $C \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_{Ab_{\in\kappa^*}} Ab_{\in\kappa+\kappa^*}$ as well.

In simple theories, forking is characterized by canonical bases in the following way: $E \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_D F$ (with $D \subseteq F$) holds if and only if $\mathrm{cb}(\mathrm{tp}(E/\mathrm{bdd}^{\mathrm{heq}}(F))) \in \mathrm{bdd}^{\mathrm{heq}}(D)$ [Kim 2014, Lemma 4.3.4]. Therefore, we have that $\mathrm{cb}(\mathrm{tp}(C/\mathrm{bdd}^{\mathrm{heq}}(Ab_{\in\kappa+\kappa^*}))) \in \mathrm{bdd}^{\mathrm{heq}}(Ab_{<\kappa}) \cap \mathrm{bdd}^{\mathrm{heq}}(Ab_{\in\kappa^*})$, but $\mathrm{bdd}^{\mathrm{heq}}(Ab_{<\kappa}) \cap \mathrm{bdd}^{\mathrm{heq}}(Ab_{\in\kappa^*}) = \mathrm{bdd}^{\mathrm{heq}}(A)$ by assumption. So $C \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_A b_{\in\kappa+\kappa^*}$, whence $C \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_A b_0$ and hence $b_0 \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_A C$, as required. □

Proposition 4.27. *Let T be a simple theory. For any A and A-indiscernible sequence I, the following are equivalent.*

(1) *I is an $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}$-Morley sequence over A.*

(2) *For any J and K with $J + K \equiv^{\mathrm{EM}}_A I$, $J \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{b}}_A K$.*

(3) *I is a total $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{bu}}$-Morley sequence over A.*

Proof. (1)⇒(3) is Corollary 4.25. (3)⇒(2) is obvious. For (2)⇒(1), assume that (2) holds. Fix $(b_i)_{i<\omega+\omega} \equiv^{\mathrm{EM}}_A I$. $(b_i)_{\omega\leq i<\omega+\omega}$ is $Ab_{<\omega}$-indiscernible. Therefore by Lemma 4.26, $b_\omega \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_A b_{<\omega}$, and we have that $b_{<\omega+\omega}$, and therefore I, is an $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}$-Morley sequence over A. □

On the other hand, there are easy examples in NIP theories (such as DLO) of total $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{bu}}$-Morley sequences that are not strict Morley sequences (i.e., sequences $b_{<\omega}$ satisfying that $b_i \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_A b_{<i}$ and $b_{<i} \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{f}}_A b_i$ for all $i < \omega$). Fix a model M of DLO and let $(a_i b_i)_{i<\omega}$ be a sequence of elements above M satisfying $a_i < a_{i+1} < b_{i+1} < b_i$ for all $i < \omega$. This is a total $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{bu}}$-Morley sequence since it is generated by an M-invariant type, but it is clearly not a strict Morley sequence. DLO can also be used to show that not every $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\text{þ}}$-Morley sequence in a rosy theory is a total $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{bu}}$-Morley sequence (e.g., [Adler 2005, Example 3.13] is an $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\text{þ}}$-Morley sequence since þ-forking in DLO is trivial but fails to even be an $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{b}}$-Morley sequence).

In NSOP_1 theories, we do get that tree Morley sequences are total $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{bu}}$-Morley sequences.

Proposition 4.28. *Let T be an $NSOP_1$ theory, and let $M \models T$. If I is a tree Morley sequence over M, then it is a total $\mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{bu}}$-Morley sequence over M.*

Proof. Let J be a sequence realizing the same EM-type as I over M. Find $K \equiv_M I$ such that $K \mathop{\smile\!\!\!\!\!\!\mid\,}^{\mathrm{K}}_M IJ$. Let I', J', and K' have the same order type such that $I + I'$, $J + J'$, and $K + K'$ are all M-indiscernible. Since these are tree Morley sequences, we have that $I \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{K}}_{M} I'$, $J \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{K}}_{M} J'$, and $K \mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{K}}_{M} K'$. Therefore, by the independence theorem for NSOP$_1$ theories, we can find I'' and J'' such that $I + I''$, $K + I''$, $K + J''$, and $J + J''$ are all M-indiscernible, so $I \approx_M J$.

Since we can do this for any such J, we have that I is a total $\mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}$-Morley sequence by Theorem 4.8 and the fact that Lascar strong types are types over models. $\square$

The converse is unclear. The argument in the context of simple theories relies on the existence of canonical bases for types.

Question 4.29. *If T is NSOP$_1$, is every total $\mathop{\smile\!\!\!\!\!\!\;|}^{\mathrm{bu}}$-Morley sequence over $M \models T$ a tree Morley sequence over M?*

References

[Adler 2005] H. Adler, "Explanation of independence", 2005. arXiv math/0511616

[Adler 2009] H. Adler, "A geometric introduction to forking and thorn-forking", *J. Math. Log.* **9**:1 (2009), 1–20. MR Zbl

[Buechler 1997] S. Buechler, "Canonical bases in some supersimple theories", preprint, 1997.

[Conant and Hanson 2022] G. Conant and J. Hanson, "Separation for isometric group actions and hyperimaginary independence", *Fund. Math.* **259**:1 (2022), 97–109. MR Zbl

[Erdős et al. 1984] P. Erdős, A. Hajnal, A. Máté, and R. Rado, "Trees and positive ordinary partition relations", Chapter IV, pp. 80–104 in *Combinatorial set theory*: *partition relations for cardinals*, Studies in Logic and the Foundations of Mathematics **106**, North-Holland, Amsterdam, 1984. MR Zbl

[Kaplan and Ramsey 2020] I. Kaplan and N. Ramsey, "On Kim-independence", *J. Eur. Math. Soc.* **22**:5 (2020), 1423–1474. MR Zbl

[Kim 2014] B. Kim, *Simplicity theory*, Oxford Logic Guides **53**, Oxford University Press, 2014. MR Zbl

[Kim et al. 2014] B. Kim, H.-J. Kim, and L. Scow, "Tree indiscernibilities, revisited", *Arch. Math. Logic* **53**:1-2 (2014), 211–232. MR Zbl

[Lascar 1982] D. Lascar, "On the category of models of a complete theory", *J. Symbolic Logic* **47**:2 (1982), 249–266. Zbl

[Shelah 1980] S. Shelah, "Simple unstable theories", *Ann. Math. Logic* **19**:3 (1980), 177–203. MR Zbl

[Silver 1971] J. H. Silver, "Some applications of model theory in set theory", *Ann. Math. Logic* **3**:1 (1971), 45–110. MR Zbl

[Simon 2015] P. Simon, *A guide to NIP theories*, Lecture Notes in Logic **44**, Cambridge University Press, 2015. MR Zbl

[Tent and Ziegler 2012] K. Tent and M. Ziegler, *A course in model theory*, Lecture Notes in Logic **40**, Cambridge University Press, 2012. MR Zbl

[Wagner 2015] F. O. Wagner, "Plus ultra", *J. Math. Log.* **15**:2 (2015), art. id. 1550008. MR Zbl

Received 26 Mar 2022. Revised 16 Oct 2023.

JAMES E. HANSON:
jhanson9@umd.edu
Department of Mathematics, University of Maryland, College Park, MD, United States

Model Theory
Vol. 3, No. 1, 2024
https://doi.org/10.2140/mt.2024.3.71

Quelques modestes compléments aux travaux de Messieurs Mark DeBonis, Franz Delahan, David Epstein et Ali Nesin sur les groupes de Frobenius de rang de Morley fini

Bruno Poizat

Nasılsın Hoca?

Bu makala, Ali Nesin'in takımın tarafından sonlu Morley ranklı Frobenius gruplar üzerine yapılan işlere yeni bir bakış atıyor; komplemanı içinde 2-eleman olan gruplar, psödo-yerli sonlu gruplar, bağlayılan gruplar üstünde has bir dikkat veriyoruz.

В статье по-новому обозреваюстя работы Али Несина и его соратников, посвященные группам Фробениуса конечкого ранга Морли, в часности тем, что имеют в своем дополнении инволюцию, или являюстя псевдо-локально конечными или связными.

Dieser Artikel wirt ein neues Licht auf die Arbeiten von Ali Nesin und seiner Koauthoren über Frobenius-Gruppen von endlichem Morley rang, insbesondere pseudo-lokal endliche, zusammenhängende, und solche, deren Komplement eine Involution enthalten.

This paper casts a new look on the works of Ali Nesin and his team concerning Frobenius groups of finite Morley rank, in particular those which have an involution in their complement, or are pseudo-locally finite, or are connected.

Ce papier procède à un réexamen des travaux de l'équipe d'Ali Nesin sur les groupes de Frobenius de rang de Morley fini, en particulier sur ceux qui ont une involution dans leur complément, ceux qui sont pseudo-localement finis, et ceux qui sont connexes.

Nous dirons que deux sous-groupes du groupe F sont *disjoints* si leur intersection est réduite à l'élément neutre, et qu'un sous-groupe T propre de F (c'est-à-dire différent de $\{1\}$ et de F) est *malnormal* s'il est autonormalisant et disjoint de ses conjugués. Cela signifie que, pour tout a hors de T, $T \cap aTa^{-1} = \{1\}$, ou encore que, dans l'action de F sur ses classes à gauche modulo T, le fixateur de deux points distincts est toujours réduit à l'identité.

MSC2020 : 14L99, 20E32, 20F11, 20F50.
Mots-clefs : groupe de Frobenius localement fini, rang de Morley fini.

Un *groupe de Frobenius* est un groupe F possédant un sous-groupe malnormal T ; on note $U(T)$ l'ensemble des points qui n'appartiennent à aucun conjugué de T, augmenté de l'élément neutre. Le groupe T est plus traditionnellement appelé *complément* (*de Frobenius*) de F, tandis que $U(T)$ est sa *base*, surtout quand elle forme un sous-groupe de F.

Nous utiliserons tout le temps les quatre faits évidents suivants, que nous nous abstenons de démontrer pour ne pas risquer d'affecter ce que les Américains appellent le sentiment de *self-respect* de nos lecteurs.

Lemme 0. (i) *L'intersection de deux sous-groupes de F malnormaux non disjoints est aussi malnormale.*

(ii) *Si T est malnormal dans F, et si H est un sous-groupe de F qui n'est ni inclus dans T, ni disjoint de T, alors $H \cap T$ est malnormal dans H : un sous-groupe propre de F est ou bien un groupe de Frobenius, ou bien est inclus dans un conjugué de T, ou bien est inclus dans $U(T)$.*

(iii) *Soient T malnormal dans F et $a \neq 1$ un point de F. S'il est dans T, son centralisateur est inclus dans T et s'il est dans $U(T)$, son centralisateur est inclus dans $U(T)$. Tout sous-groupe normal de F qui est abélien, ou même a un centre non trivial, est inclus dans $U(T)$. Si S est un sous-groupe normal non trivial de T, ce dernier est son normalisateur.*

(iv) *Si T est un sous-groupe malnormal de H et H est un sous-groupe malnormal de F, T est malnormal dans F.*

Si F est engendré par T et par un sous-groupe V normal et disjoint de T (F est alors le produit semi-direct de V par T), on dit qu'il est *scindé*. Dans ce cas, V est inclus dans $U(T)$, mais ne lui est pas forcément égal : quand cela se produit, nous dirons que F est *scindé nettement*.

À l'autre extrême, nous dirons que T *remplit* F si $U(T)$ est réduit à l'élément neutre, et que F est *plein* s'il possède un sous-groupe malnormal le remplissant ; cette terminologie a été introduite par [Jaligot 2001]. Elle est motivée par l'apparition, dans [Cherlin 1979], d'un groupe de Frobenius plein de rang de Morley trois (en fait, c'est [Nesin 1989] qui a montré que c'était un groupe de Frobenius) ; il a fallu attendre [Frécon 2018] pour qu'on se rende compte qu'il n'existe pas. La possible existence de groupes de Frobenius pleins de rang de Morley fini reste un obstacle majeur à la conjecture d'algébricité.

Les groupes de Frobenius finis, eux, sont scindés nettement. Ils ont une structure très particulière, qui intervient lourdement dans la classification des groupes simples finis ; en effet, dans le cas fini, pour chaque T malnormal, l'ensemble $U(T)$ est un sous-groupe non trivial [Frobenius 1901], qui est nilpotent [Thompson 1959]. Ce résultat de Thompson est un prélude au théorème de [Feit et Thompson 1963]. Quand le groupe T contient une involution i, il n'est pas dur de voir que $U(T)$ est un groupe abélien inversé par i ; quand $U(T)$ est un groupe résoluble, ce n'est pas la mer à boire de montrer qu'il est nilpotent (voir notre proposition 5.4) ; mais, dans le cas général, ce sont des résultats difficiles, dont les seules démonstrations connues s'appuient sur des calculs de caractères (inventés justement par Frobenius) dont on ne connaît pas d'équivalents en rang de Morley fini.

Nous allons faire notre possible pour tenir un discours nouveau sur les groupes de Frobenius de rang de Morley fini. Il est remarquable que ces groupes occupent une place substantielle, aux pages 203–219, de [Borovik et Nesin 1994], dont la référence sera B&N dans la suite de cet article, mais que les résultats obtenus en toute généralité sont peu exploités dans le reste du livre (probablement parce qu'ils ne sont pas assez décisifs), pas plus que dans d'autres travaux sur la conjecture d'algébricité. C'est ainsi que l'expression *Frobenius group* n'apparaît pas dans [Altınel et al. 2008].

Dans la section 1, nous démontrons quelques résultats faciles sur les groupes de Frobenius finis, en faisant semblant d'ignorer les théorèmes de Frobenius et de Thompson. Nous ne faisons pas cela pour le plaisir de redécouvrir des trivialités, mais parce que nous espérons pouvoir transférer certaines démonstrations aux groupes de Frobenius *connexes* de rang de Morley fini, en remplaçant les calculs multiplicatifs sur les cardinalités par des calculs additifs sur les rangs, combinés avec des arguments de généricité.

Dans la section 2 nous renforçons [B&N, Lemma 11.20, p. 207], concernant les groupes de Frobenius de rang de Morley fini dont le complément contient une involution.

Dans la section 3 nous examinons ce qui pourrait s'opposer à la netteté des scissions dans un contexte de rang de Morley fini.

Dans la section 4, nous étendons les propriétés des groupes de Frobenius finis à ceux qui sont pseudo-localement finis et ont un rang de Morley fini, aux groupes algébriques en particulier.

La section 5 détaille la structure des groupes de Frobenius de rang de Morley fini, avec une attention particulière donnée à ceux qui sont connexes ; nous constaterons que, dans leur cas, tous les groupes utiles sont définissables et connexes (ce qu'annonce [B&N, Corollary 11.24, p. 209]).

Enfin, la section 6 est consacrée aux actions sans points fixes sur un groupe commutatif, ce qui permet de préciser les résultats de la section 4 et nous fait revisiter la problématique du comportement des corps en situation de rang de Morley fini.

Comme l'indique son titre, l'ambition de cet article est modeste, puisqu'il est consacré à une révision, une trentaine d'années après leur parution, de travaux de pionniers : ceux de [DeBonis et Nesin 1994; Delahan et Nesin 1993; 1995; Epstein et Nesin 1994; Nesin 1992; 1994], qui sont exposés dans [B&N, p. 203–219]. Je redémontre tout ce dont je fais un usage essentiel, comme la facile, mais fondamentale, proposition 4.1 de définissabilité, qui leur appartient. Ces auteurs ont généralisé des résultats algébriques dans des contextes particuliers (présence d'une involution dans T, solvabilité de $U(T)$, etc.) où les propriétés des groupes de Frobenius finis sont relativement simples à établir ; assez souvent, les groupes connexes constituent le cas facile de leurs théorèmes. À l'opposé, je resterai résolument à l'intérieur de la théorie des modèles, en exploitant l'unicité du générique pour tenter d'isoler les pathologies des groupes de Frobenius connexes (quand l'algèbre ne peut être évitée, je me reposerai sur B&N !). Il faut dire que, pour quelqu'un qui est familier des groupes finis, ou des groupes algébriques, les groupes de rang de Morley fini connexes sont les objets les plus troublants qui soient, leur propriété la plus mystérieuse étant l'existence potentielle.

Cette étude accumule des petits faits, qui pour la plupart concernent des objets hypothétiques en désaccord avec la conjecture d'algébricité ; j'ai cru utile de les rendre publics, bien qu'aucun d'entre eux ne m'a semblé mériter le nom de théorème. Ses principales nouveautés de méthode sont :

• On peut montrer directement l'unicité de l'involution dans le complément, que le groupe soit connexe ou pas, en évitant le calcul de rang de B&N (section 2).

• C'est bien de constater qu'un groupe de Frobenius est scindé, mais il faut en plus se préoccuper de la netteté de la scission ; j'ai l'impression d'être la première conscience modèle-théorique à avoir été tourmentée par elle (section 3).

• La théorie des modèles donne une explication convaincante — la pseudo-finitude locale — du transfert des propriétés des groupes de Frobenius finis aux groupes de Frobenius algébriques (section 4).

• La considération du socle d'un groupe connexe semi-simple aide à la description de la structure des groupes de Frobenius de rang de Morley fini pathogènes (section 5).

• Le beau lemme 4.1 du chapitre 1 d'[Altınel et al. 2008] (est-ce sa première application ?) permet dans les bons cas de montrer l'abélianité de la composante connexe du complément quand la base est nilpotente (section 6).

Конец этово введения — идеальное место, чтобы выразить мою самую искреннюю благодарность Александру Васильевичу, который перевернул мою научную жизнь, объяснив мне около 36 лет назад, что инволюции суть вещи, имеющие определенную значимость в теории групп.

Exemple 0.1. Soient V un groupe commutatif et F le produit semi-direct de V par une involution i qui l'inverse. Le sous-groupe $T = \{1, i\}$ est malnormal dans F si et seulement s'il est égal à son centralisateur, c'est-à-dire si V ne contient pas d'involutions. Dans ce cas, $U(T) = V$ si et seulement V est divisible par 2.

Par exemple, si V est le groupe cyclique infini $\mathbb{Z}$, F est le groupe diédral infini, isomorphe au groupe des transformations affines $(-1)^n x + m$ de l'anneau $\mathbb{Z}$ des entiers ; les conjuguées de l'involution $-x$ sont les $-x+2m$, si bien que $U(T)$ est formé des $x+m$ et des $-x+2m+1$: ce n'est pas un groupe. Il y a deux classes de conjugaison d'involutions dans F. On remarque que le quotient de ce groupe par $2\mathbb{Z}$ n'est pas un groupe de Frobenius.

Exemple 0.2. [Olchanski 1982] construit, pour chaque nombre premier p assez grand, des groupes Ol_p dénombrables dont tous les sous-groupes propres ont exactement p éléments. Chacun d'entre eux est malnormal, et ces groupes de Frobenius Ol_p ne sont pas scindés puisqu'ils sont simples ; ils sont même pleins si jamais tous leurs sous-groupes propres sont conjugués. Comme ils n'ont pas de sous-groupes abéliens infinis, ils ne peuvent être superstables.

Par compacité, on obtient à la limite des groupes Ol_∞ qui vérifient :

(i) Le centralisateur de chaque $a \neq 1$ est un groupe abélien sans torsion divisible.

(ii) Le groupe n'est pas commutatif, et si a et b ne commutent pas, leurs centralisateurs sont disjoints.

(iii) Tout sous-groupe définissable propre est le centralisateur d'un point.

L'état présent de nos connaissances ne permet pas d'exclure que certains de ces groupes de Frobenius soient pleins et de rang de Morley quatre.

Notons que, pour qu'un groupe de Ol_∞ ait un rang de Morley fini, il est nécessaire que tous les centralisateurs soient conjugués et qu'il existe un entier n tel que : (a) pour tout $x \neq 1$, chaque y est produit de n conjugués de x ou de x^{-1} ; (b) si C_1 et C_2 sont deux centralisateurs distincts, chaque y est produit de n points de $C_1 \cup C_2$. Les limites de groupes dans Ol_p ne satisfont pas à la dernière condition.

Exemple 0.3. Dans un groupe libre non commutatif, les sous-groupes commutatifs sont cycliques et malnormaux. Les groupes libres sont la source de nombreuses constructions de groupes de Frobenius exotiques.

1. Quelques résultats faciles sur les groupes de Frobenius finis

Lemme 1.1. *Soient un groupe fini F et un de ses sous-groupes propres T.*

(i) *Le nombre de points de $U(T)$ est supérieur ou égal à l'indice de T dans F ; il lui est égal si et seulement si T est malnormal.*

(ii) *Il n'y a pas de groupe de Frobenius fini plein.*

Si T est malnormal :

(iii) *F est scindé relativement à T si et seulement si $U(T)$ est un groupe.*

(iv) *Si V est un sous-groupe normal dans F strictement inclus dans $U(T)$, alors $T_1 = T/V$ est malnormal dans $F_1 = F/V$, et $U(T_1) = U(T)/V$.*

(v) *Si T contient une involution i, il n'en contient pas d'autres, et $U(T)$ est un groupe commutatif sans involutions inversé par i.*

Démonstration. (i) Notons t et f les ordres respectifs de T et de F. Si T est malnormal, on compte, en mettant de côté l'élément neutre, le nombre de points dans l'union des conjugués de T, ce qui donne $1+(t-1)f/t = f+1-f/t$; le nombre de points de $U(T)$ est donc $f-(f+1-f/t)+1 = f/t$. Si T n'est pas malnormal, ou bien ses conjugués sont en nombre inférieur à son indice, ou bien ils ne sont pas disjoints.

(ii) f/t vaut au moins 3 quand T n'est pas normal dans F.

(iii) Si $U(T)$ est un groupe, il est normal dans F et disjoint de T, si bien que le groupe G engendré par $U(T)$ et T est leur produit semi-direct. D'après le point (i), F et G ont même ordre, et sont égaux.

Si F est scindé et est produit semi-direct de V par T, V est inclus dans $U(T)$ car disjoint de T et de ses conjugués (il est normal dans F). Comme son nombre d'éléments est l'indice de T dans F, c'est $U(T)$ tout entier : on voit donc que la scission est nette.

(iv) On remarque que T est isomorphe à T_1 puisqu'il est disjoint de V. Notons f, t et v les ordres respectifs de F, T et V ; l'indice de T_1 dans F_1 vaut f/tv, et T_1 est un sous-groupe propre de F_1 puisque $v < f/t$. Comme VT est nettement scindé, $VT \cap U(T) = V$. C'est également vrai si on remplace T par un de ses conjugués, si bien que $U(T_1)$ est formé des points dont l'image réciproque est dans $U(T)$; il a donc f/tv éléments.

(v) Soit j une involution dans un autre conjugué de T ; i comme j inversent par conjugaison leur produit ij, qui ne peut être dans un conjugué de T car cela forcerait i et j à être dans ce dernier. Ce produit ij est donc dans $U(T)$. Comme en prenant une involution j dans chaque conjugué de T on a déjà le compte, il n'y a qu'une seule involution par conjugué de T, et $U(T)$ est l'ensemble des points ij inversés par i. Il ne peut pas contenir d'involutions, qui commuteraient avec i. Comme $U(T)$ est inversé par chaque involution, un produit de deux involutions le centralise, si bien que c'est un groupe commutatif, centralisateur de chacun de ses points autres que l'identité. □

Lemme 1.2. (i) *Si un groupe fini a deux sous-groupes malnormaux, chacun intersecte non trivialement un conjugué de l'autre.*

(ii) *Dans un groupe de Frobenius fini, les sous-groupes malnormaux minimaux sont conjugués.*

(iii) *Soient T malnormal dans F fini, et G un sous-groupe de F contenant T, alors tout conjugué de T qui intersecte G non trivialement est inclus dans G et conjugué de T dans G. Alors $G \cap U(T)$ est formé de* 1 *et des points de G*

qui ne sont dans aucun conjugué de T au sens de G, et si G est engendré par T et par l'une de ses parties A, il est aussi engendré par T et ses conjugués ${}^aT = aTa^{-1}$, où a parcourt A.

(iv) *Si T est malnormal dans F fini, F est engendré par $U(T)$ et T.*

Remarque. Je suis la tradition française qui note les automorphismes intérieurs comme des actions à gauche, ce qui conduit fort logiquement à placer les exposants à gauche.

Démonstration. (i) Si les conjugués de T étaient tous disjoints de ceux de T_1, le nombre d'éléments conjugués d'un point de T ou d'un point de T_1 vaudrait $1 + f - f/t + f - f/t_1 \geq 1 + 2f - \frac{1}{2}f - \frac{1}{2}f = 1 + f$.

(ii) L'intersection de deux sous-groupes malnormaux est triviale ou malnormale.

(iii) Soit T_1 un conjugué de T coupant G non trivialement. D'après (i), cette intersection est conjuguée de T à l'intérieur de G. Si un point $g \neq 1$ de G n'appartient pas à $U(T)$, il est dans un conjugué de T au sens de G. Le groupe H engendré par T et ${}^aT$ contient a : en effet, ${}^aT$ est conjugué de T par un h de H, et a est dans hT.

(iv) Le groupe G engendré par T et $U(T)$ est réunion de $U(T)$ et de conjugués de T. S'il existe a hors de G, le groupe ${}^aG = aGa^{-1}$ contient $U(T)$, qui est un ensemble normal, et les conjugués de T qui sont dans G sont disjoints de ceux qui sont dans ${}^aG$. Par conséquent l'intersection de G et de ${}^aG$ est réduite à $U(T)$, qui est un groupe ; F est donc le produit semi-direct de $U(T)$ par T, ce qui contrarie l'hypothèse. □

Lemme 1.3. *Soit F un groupe fini avec un sous-groupe malnormal T. On suppose que les centralisateurs de deux points de $U(T)$ ne sont jamais disjoints ; alors $U(T)$ est un groupe et T est malnormal maximal.*

Démonstration. Prenons $u \neq 1$ et $v \neq 1$ dans $U(T)$, et $w \neq 1$ qui commute avec chacun d'eux. Comme w commute avec u, il est dans $U(T)$, ainsi que uv qui commute avec w.

Soit G un surgroupe propre de T. Comme il n'est pas plein, il doit contenir un point u de $U(T)$ non trivial. Pour tout a hors de G, G^a contient $v = u^a$, et si G était malnormal aucun point non trivial ne pourrait commuter avec u et v. □

Remarque. D'après les théorèmes de Frobenius et de Thompson, l'hypothèse du lemme 1.3 est toujours vérifiée, si bien que, dans un groupe de Frobenius fini F, les sous-groupes malnormaux T sont tous maximaux, et donc aussi tous minimaux, et tous conjugués ; le groupe nilpotent $U(T)$ est uniquement déterminé (c'est le plus grand sous-groupe nilpotent normal dans F).

2. Une involution dans le complément

Pour ne pas lasser d'emblée la patience de nos lectrices par un exposé rébarbatif des propriétés générales des groupes de Frobenius de rang de Morley fini, nous préférons entrer dans le vif du sujet en adaptant à leur cas notre lemme 1.1(v), bien que cela nous oblige à l'occasion, pour éviter les redites, de faire appel à quelques lemmes faciles montrés par la suite. Cela précise [B&N, p. 207–208], et donne l'espoir qu'il reste quelques épis à glaner dans la relecture des œuvres de nos glorieux précurseurs. Il s'agit typiquement d'arguments déductibles d'un simple comptage dans le cas des groupes finis, mais qui sont par ailleurs susceptibles d'une analyse locale pouvant s'étendre aux groupes de rang de Morley fini, principalement lorsqu'il est question d'involutions. Le plus ancien résultat de ce genre est la conjugaison des 2-sylows, montrée dans [Borovik et Poizat 1990].

Nous suivons l'usage de surnommer *fortement réels* les produits de deux involutions.

Dans un groupe uniquement 2-divisible, comme le sont les groupes de rang de Morley fini sans involutions, le *milieu* de a et de b est l'unique point m tel que $ma^{-1}m = b$; voir [Poizat 2018].

Proposition 2.1. *Soit F un groupe de Frobenius de rang de Morley fini, avec un sous-groupe malnormal T contenant une involution i. Nous notons I l'ensemble des involutions de F.*

(i) *T ne contient pas de deuxième involution : il est le centralisateur de i. Deux involutions de F sont conjuguées par une unique involution. Tout point de F s'écrit de manière unique comme produit d'une involution et d'un élément de T, ainsi que comme produit d'un point de iI et d'un point de T. Les points fortement réels gisent dans $U(T)$, et ce sont les commutateurs d'involutions.*

(ii) *Nous avons*

$$\mathrm{RM}(T) \le \mathrm{RM}(I)\ ; \quad \mathrm{RM}(F) = \mathrm{RM}(T) + \mathrm{RM}(I)\ ;$$
$$d^\circ M(F) = d^\circ M(T) \times d^\circ M(I).$$

Si T est infini, toutes les involutions sont conjuguées sous l'action de F°, tous les points fortement réels sont dans F°, et $d^\circ M(I) = 1$.

(iii) *Les centralisateurs des points fortement réels non triviaux sont commutatifs et autocentralisants ; ce sont les sous-groupes définissables maximaux contenus dans $U(T)$ et normalisés par une involution. Si i et j sont deux involutions distinctes, le centralisateur de ij est contenu dans $iI \cap jI$.*

(iv) *S'il existe un point $a \neq 1$ inversé par toutes les involutions, $U(T)$ est un groupe commutatif inversé par chaque involution, qui est formé des points* *fortement réels, et F est le produit semi-direct de $U(T)$ et de T. Si de plus T est infini, $U(T)$ est connexe.*

(v) *Dans le cas contraire, T est infini, et aucun sous-groupe définissable non trivial normalisé par toutes les involutions n'est contenu dans $U(T)$. Le groupe engendré par les points fortement réels de F est son plus petit sous-groupe non trivial normal définissable ; c'est un groupe de Frobenius simple.*

Démonstration. (i) Soit j une involution qui n'est pas dans T, c'est-à-dire qui est dans $U(T)$ ou dans un autre conjugué de T. Autant i que j inversent leur produit ij, qui ne peut être dans un conjugué ${}^aT$ de T, car cela forcerait i et j à être dans ce conjugué : il est donc dans $U(T)$, et le plus petit sous-groupe définissable le contenant est inclus dans $U(T)$, car il est commutatif. Il est normalisé par i, qui inverse chacun de ses points, si bien qu'il ne contient pas d'involutions. Il est donc uniquement 2-divisible. La racine carrée $(ij)^{1/2}$ de ij, étant inversée par i, est de la forme ik, où k est une involution, et donc $ik \cdot ik = ij$, soit $kik = j$. Comme j est conjuguée de i (par une involution), elle ne peut être dans $U(T)$.

Si i et i' sont deux involutions de T et j est une involution dans un autre conjugué de T, il existe une involution k qui conjugue i et j, et une involution k' qui conjugue i' et j. Le produit $k'k$ conjugue i et i' ; s'il est différent de 1, il doit être dans T, et comme il est inversé par k et par k', ces derniers doivent aussi être dans T, ce qui est une situation absurde. Donc $k = k'$ et $i = i'$; i, étant la seule involution de T, est centrale dans T. Ce dernier est donc le centralisateur de i ; nous constatons qu'il est définissable, ainsi que $U(T)$.

Le produit de deux involutions est toujours dans $U(T)$, soit qu'elles soient égales, soit qu'elles appartiennent à des conjugués de T différents. Si deux involutions k et k' conjuguent i et j, leur produit kk' est dans $U(T)$ et commute avec i ; donc $kk' = 1$, $k = k'$. Comme T est le centralisateur de i, tout point a est congru modulo T à l'involution qui conjugue i et aia^{-1}.

Par conséquent $F = IT = i \cdot iIi \cdot iT = iIT$, et ces décompositions se font de manière unique.

Quel que soit a, le commutateur $[a, i] = (aia^{-1})i$ est bien produit de deux involutions ; réciproquement, toute involution j est conjuguée de i par une involution k, si bien que $ji = kik \cdot i = [k, i]$.

(ii) Aucun point non trivial de T ne commute avec une involution $j \neq i$, si bien que les conjuguées de j par les points de T sont toutes distinctes ; les deux égalités suivantes proviennent de la décomposition $F = IT$.

Si T est infini, la proposition 5.1 nous apprendra que $T \cap F^\circ = T^\circ$ est malnormal dans F°, et que F° agit transitivement sur les conjugués de T. Comme chacun d'entre eux ne contient qu'une involution, F° agit transitivement sur les involutions ; comme deux involutions sont toujours congrues modulo F°, ce dernier contient tous les points fortement réels. Comme iI est inclus dans F°, chacun de ses points s'écrit de manière unique comme produit d'un point de iI et d'un point de T°, si bien que $d^\circ M(iI) = d^\circ M(I) = 1$.

(iii) Soit A un sous-groupe définissable non trivial, normalisé par i et contenu dans $U(T)$; il n'a pas d'involutions. Pour chacun de ses points a, i commute avec le milieu de a et de iai ; ce milieu vaut donc 1, ce qui signifie que A est un groupe abélien inversé par i. Il en est de même de son centralisateur.

Le centralisateur de ij est normalisé par i comme par j.

(iv) Le point a centralise tout les éléments semi-réels, qui forment donc un groupe commutatif R, et F est le produit semi-direct de R et de T. Notre corollaire 3.2 expliquera bientôt pourquoi la commutativité de R implique la netteté de la scission, c'est-à-dire que $U(T) = R$. Si T est infini, comme $R = iI$, il est connexe.

(v) Si F est fini, le lemme 1.1(v) nous a montré que nous sommes dans le cas précédent ; il en est de même si F est infini et T est fini, car alors F° est inclus dans $U(T)$, d'après la proposition 5.1 à venir.

D'après le point (iii), un sous-groupe définissable normal, ou même seulement normalisé par toutes les involutions, et contenu dans $U(T)$, est trivial. Un sous-groupe de F définissable, normal et non trivial ne peut être fini, car sinon il commuterait avec F°. D'après la proposition 5.1, il agit transitivement sur les involutions, et contient tous les points semi-réels. Ces derniers forment donc un ensemble indécomposable au sens de Zilber, et engendrent un groupe G définissable, qui est l'unique sous-groupe définissable normal minimal de F.

Le groupe G est égal à son socle, composé du produit d'un nombre fini de groupes simples. Comme $T' = T \cap G$ est malnormal dans G, il n'est pas possible que tous soient contenus dans $U(T')$, et, comme ils commutent, il n'y en a qu'un seul (nous reprendrons ce type de raisonnement dans la proposition 5.4). $\square$

Remarques. (i) Nos involutions forment ce que j'ai appelé un *symétron* dans [Poizat 2021], où je constate que bien des propriétés connues pour les groupes s'étendent aux symétrons de rang de Morley fini ; elles ont été étendues dans la première partie de [Zamour 2022].

(ii) Nous verrons dans la section 4 qu'un groupe de Frobenius simple ne peut être algébrique. S'il existe, le groupe paradoxal décrit en (v) ou bien n'a pas d'involutions, ou bien ses 2-sylows sont des groupes de Prüfer de rang un. Dans ce deuxième cas il est engendré par un symétron d'involutions qui contredit de manière extrême un classique de la théorie des groupes finis, le théorème Z^* de Glauberman ; voir [B&N, Question B.5, p. 355].

(iii) En se laissant guider par les démonstrations du théorème 4 et de la proposition 13 de [Poizat 2018], on voit facilement que, dans n'importe quel groupe de rang de Morley fini, le sous-groupe engendré par les involutions et celui engendré par les points fortement réels sont définissables. En fait, il en est ainsi du groupe engendré par n'importe quel ensemble définissable clos par conjugaison et par carré (ou plus généralement par élévation à la puissance n, pour un $n > 1$ fixé).

3. Obstruction à la netteté des scissions

Nous examinons ici des problèmes causés par les quotients qui n'apparaissent pas chez les groupes finis, mais qui tourneront à l'obsession dans l'étude des groupes de Frobenius de rang de Morley fini. Nous sommes bien sûr incapables de donner des exemples de leur nuisance, pour la raison qu'elle ne se manifeste que dans des contextes infirmant la conjecture d'algébricité.

Étant donné un groupe d'automorphismes $T \neq \{\mathrm{id}\}$ du groupe V, nous déterminons tout d'abord les circonstances qui font que T est malnormal dans le produit semi-direct F de V par T.

Pour cela nous calculons à quelle condition $(vs)^{-1}$ conjugue ut dans T, lorsque u et v sont dans V et s et t dans T : $s^{-1}v^{-1} \cdot ut \cdot vs$ est dans T, soit encore $v^{-1} \cdot ut \cdot v$ est dans T, soit encore $v^{-1} \cdot ut \cdot v \cdot t^{-1} = v^{-1} \cdot u \cdot tvt^{-1}$ est dans T, c'est-à-dire vaut 1 puisqu'il s'agit du produit de trois éléments de V. La condition est donc que $u = vtv^{-1}t^{-1} = [v, t]$.

Dire que T est malnormal signifie que, quand $u = 1$ et $t \neq 1$, cela se produit seulement si $v = 1$; autrement dit 1 est le seul point de V qui commute avec t, le seul point fixe de l'action de t sur V par automorphisme intérieur (quand cela se produit pour tout $t \neq 1$, nous dirons que T agit *sans points fixes* sur V, ou encore que T agit *librement* sur V). Cela veut dire aussi que, à $t \neq 1$ fixé, l'application $[v, t]$ de V dans V est injective, car $[v, t] = [w, t]$ équivaut à $[t, w^{-1}v] = 1$.

Par contre, dire que tout ut a un conjugué dans T, soit encore que la scission est nette, signifie que cette application est surjective. On remarque que, dans ce cas, ut est conjugué de t par un point de V.

Si V est fini, l'injectivité implique la surjectivité (et réciproquement d'ailleurs), ce qui donne une autre explication des points (iii) et (iv) du lemme 1.1, mais l'exemple 0.1 montre que cette vérité ne franchit pas les Pyrénées ; il n'est pas certain que ce soit toujours le cas dans un contexte de rang de Morley fini, bien qu'aucun contre-exemple ne soit connu.

Par contre *c'est vrai dans un contexte localement fini*, c'est-à-dire si, pour chaque t de T, chaque partie finie de V est contenue dans un groupe fini normalisé par t.

Nous avons besoin de deux faits de pure théorie des groupes.

Si s est un endomorphisme du groupe V, l'*adjointe* de s est l'opération $a(x) = s(x)x^{-1}$; elle est soumise à la loi $a(xy) = a(x) \cdot {}^{x}a(y)$. Réciproquement, si $a(x)$ obéit à cette loi, c'est l'adjointe de l'endomorphisme $a(x)x$. L'endomorphisme s

agit sur V sans point fixe (autre que 1) si et seulement si son adjointe est injective. Le fait suivant autorise des récurrences.

Fait 3.1. *Soient s un automorphisme sans points fixes du groupe V, et W un sous-groupe normal de V normalisé par s. Si l'adjointe de s restreinte à W est surjective, l'endomorphisme s' induit par s sur V/W est sans points fixes; si de plus l'adjointe de s' est surjective, celle de s l'est aussi.*

Démonstration. Si s fixe v modulo W, $s(v) = wv = s(w')w'^{-1}v$, si bien que $s(w'^{-1}v) = w'^{-1}v$, que v est congru modulo W à un point fixe de s, c'est-à-dire est dans W.

Si v modulo W est dans l'image de s', il s'écrit

$$v = s(u)wu^{-1} = s(u)s(w')w'^{-1}u^{-1} = s(uw')(uw')^{-1}. \qquad \square$$

Corollaire 3.2. *Un groupe V résoluble par fini ne donne que des scissions nettes si le contexte est de rang de Morley fini.*

Démonstration. Notre hypothèse implique que, pour chaque t dans T, la structure formée de V et de l'automorphisme s induit par t est de rang de Morley fini. C'est clair si V est fini ; sinon V° a un sous-groupe définissable infini commutatif caractéristique A, qui est normalisé par s. L'adjointe de la restriction de s à A est un endomorphisme injectif de A dans A, et est donc surjective ; on divise par A pour conclure par induction sur le rang. $\square$

La netteté des scissions est ce qui permet de faire des quotients propres :

Fait 3.3. *Soient F un groupe de Frobenius quelconque, T un de ses sous-groupes malnormaux, et V un sous-groupe normal de F contenu dans $U(T)$. On suppose que le produit semi-direct VT est nettement scindé et strictement inclus dans F. Alors l'image T_1 de T dans F/V (qui est isomorphe à T) est malnormale dans F/V, et $U(T_1)$ est l'image de $U(T)$.*

Démonstration. Supposons que, modulo V, a conjugue $t \neq 1$ et t' dans T, et mettons en œuvre la surjectivité de l'adjointe : $ata^{-1} = vt' = [u, t'] \cdot t' = ut'u^{-1}$; on en déduit que $u^{-1}a$ est dans T, que a est dans T modulo V.

Supposons que, modulo V, x soit dans ${}^aT$: $axa^{-1} = vt = utu^{-1}t^{-1} \cdot t = utu^{-1}$; on en déduit que x est bien dans un conjugué de T. $\square$

Corollaire 3.4. *On considère un groupe de Frobenius F, un de ses sous-groupes malnormaux T, et deux de ses sous-groupes normaux V et W contenus dans $U(T)$. Alors, si VT est nettement scindé, VW est inclus dans $U(T)$; si WT est aussi nettement scindé, VWT l'est également.*

Démonstration. Si V est net, W reste disjoint de T et de ses conjugués dans le quotient F/V, et son image réciproque est dans $U(T)$. Si W est net lui aussi et

$t \neq 1$, on considère v dans V et w dans W ; $t \cdot vw = t^u \cdot w = u \cdot t(u^{-1}wu) \cdot u^{-1}$ si $v = t^{-1}utu^{-1}$, et $t(u^{-1}wu)$ est conjugué de t par un point de W. □

Remarque. La(e) rapporteuse(r) de la version préliminaire de cet article a fait l'observation suivante : si V et W sont des sous-groupes normaux contenus dans $U(T)$, et si pour un v de V et un w de W le produit vw n'est pas dans $U(T)$, l'automorphisme intérieur s associé à vw n'a pas de point fixe sur W, mais il en a dans $W/W \cap V$; son adjointe est donc injective et non surjective sur $W \cap V$.

4. Groupes de Frobenius de rang de Morley fini pseudo-localement finis

Un groupe F est *pseudo-localement fini* si tout énoncé du langage des groupes qu'il satisfait l'est aussi dans un groupe localement fini. Comme les sections d'un groupe localement fini sont localement finies, les sections définissables d'un groupe pseudo-localement fini le sont aussi.

Il est facile de constater que toute structure définissable dans un corps algébriquement clos est pseudo-localement finie ; voir par exemple [Poizat 2021]. La conjecture d'algébricité a été montrée pour les groupes localement finis dans [Thomas 1983], mais il a été ensuite remarqué par Simon Thomas lui-même que son résultat s'étendait aux groupes pseudo-localement finis. Les groupes simples de rang de Morley fini et pseudo-localement finis sont exactement les groupes algébriques simples sur un corps algébriquement clos, le travail de Thomas consistant en somme à déduire la classification de ces groupes à partir de la classification des groupes simples finis [1].

Cette pseudo-finitude locale est la clef de certains transferts immédiats de propriétés des groupes finis aux groupes algébriques (précisons-le : de corps de base algébriquement clos). Pour les groupes de Frobenius, plutôt que de reproduire des techniques de groupes algébriques comme dans [Hertzig 1961; B&N, Lemma 11.39, p. 218], nous pouvons effectuer, grâce à la proposition suivante, un transfert brutal de nature modèle-théorique.

Proposition 4.1 [B&N, Proposition 11.19, p. 206]. *Dans un groupe de Frobenius F de rang de Morley fini, tout sous-groupe malnormal est définissable* (*par une formule du langage des groupes, avec paramètres*).

Démonstration. C'est vrai si T est fini. Sinon chaque point de $T \cap F^\circ$ a un centralisateur infini [Altınel et al. 2008, Chapter 4, Corollary 4.18, p. 270] [2], qui est inclus

1. Celle-ci repose elle-même sur une bonne connaissance des groupes de Lie simples !

2. On peut se passer d'un résultat aussi délicat : si a est un point de $T \cap F^\circ$ de centralisateur fini, sa classe de conjugaison est générique dans F°, ainsi que celle de a^{-1} (différent de a) ; a et a^{-1} sont conjugués par un point b de F°, qui est dans T. Comme b^2 centralise a, l'ordre de b est fini et pair, et l'une de ses puissances est une involution, qui a un centralisateur infini ; on sait d'ailleurs que T est son centralisateur (proposition 2.1).

dans T ; le groupe engendré par les composantes connexes des centralisateurs de ses points est un groupe définissable connexe non trivial, et T est son normalisateur. $\square$

Dans un contexte pseudo-localement fini, les groupes de Frobenius définissables scindés le sont nettement : c'est un cas particulier de ce qu'on appelle le *principe de surjectivité d'Ax*.

Proposition 4.2. *Dans un groupe de Frobenius F de rang de Morley fini et pseudo-localement fini* :

(i) *Tous les sous-groupes T malnormaux sont conjugués.*

(ii) *$U(T)$ est un sous-groupe nilpotent non trivial, qui est connexe quand T est infini, et F est le produit semi-direct de $U(T)$ par T.*

(iii) *Si T contient une involution i, $U(T)$ est un groupe commutatif sans involutions inversé par i.*

Démonstration. Considérons, dans un groupe de Frobenius Φ localement fini, un sous-groupe malnormal Θ, et montrons que $U(\Theta)$ est un groupe et que Φ est le produit semi-direct de $U(\Theta)$ par Θ. En effet, si $t \neq 1$ est dans Θ, t' est dans un autre conjugué de Θ, et u et v sont dans $U(\Theta)$, ils engendrent un groupe de Frobenius fini φ, dont $\varphi \cap \Theta$ et ses conjugués dans φ sont les sous-groupes malnormaux ; ce sont aussi les traces sur φ des conjugués de Θ qui intersectent φ non trivialement, si bien que le produit uv est aussi dans $U(\Theta)$. On voit de la même façon que tout point de Φ est produit d'un point de $U(\Theta)$ et d'un point de Θ.

Si Θ' est un autre sous-groupe malnormal, une vérification locale permet de voir que $U(\Theta) = U(\Theta')$. Il existe donc un conjugué Θ'' de Θ' non disjoint de Θ ; mais alors $\Theta \cap \Theta''$ est aussi malnormal, si bien que Φ est engendré par $U(\Theta)$ et $\Theta \cap \Theta''$, par $U(\Theta)$ et Θ, et par $U(\Theta)$ et Θ'', ce qui nécessite que $\Theta = \Theta''$. Autrement dit, Θ et Θ' sont conjugués.

Le groupe $U(\Theta)$ est localement nilpotent. Il est même nilpotent d'après un théorème de [Kegel et Wehrfritz 1973], mais notre contexte, où il y a partout des bornes aux chaînes de centralisateurs, va nous permettre d'éviter l'emploi d'un résultat aussi sophistiqué.

Revenons à F. Soit $\theta(x, \underline{a})$ une formule définissant T ; c'est un énoncé élémentaire qui déclare que, si $\theta(x, \underline{y})$ définit un sous-groupe malnormal, alors $U(\theta(x, \underline{y}))$ est un groupe et chaque point de F est produit d'un point satisfaisant $U(\theta(x, \underline{y}))$ et d'un point satisfaisant $\theta(x, \underline{y})$. Comme il est vrai dans tout groupe localement fini, il l'est aussi dans F. On voit semblablement que deux sous-groupes malnormaux sont conjugués.

Reste à voir que $U(T)$ est nilpotent, alors que pour l'instant il ne l'est que pseudo-localement. On est assuré que, étant donnés $u_1, \dots, u_n$ dans $U(T)$, il existe $v \neq 1$ qui commute avec chacun d'eux. Comme $U(T)$ vérifie la condition de chaîne sur les centralisateurs, il a un centre non trivial ; $U(T)/Z(U(T))$ vérifie aussi cette condition de chaîne, si bien que $U(T)$ est abélien ou bien a un deuxième centre non trivial.

Si $U(T)$ est connexe, il n'est pas possible que son centre soit fini, car alors il serait égal à son deuxième centre. Dans ce cas, on le divise par son centre et on conclut par récurrence sur le rang.

Si $U(T)$ n'est pas connexe, d'après le raisonnement précédent $U(T)^\circ$ est nilpotent. Comme la scission de $U(T)^\circ T$ est nette, l'image T_1 de T dans le quotient $F/U(T)^\circ$ est malnormale, et $U(T_1) = U(T)/U(T)^\circ$; $U(T_1)$ est donc centralisé par la composante connexe de T_1, ce qui est impossible si T est infini. Par conséquent T est fini, $F/U(T)^\circ$ est un groupe de Frobenius fini, $U(T)/U(T)^\circ$ est nilpotent et $U(T)$ est résoluble. Pour éviter de nous fatiguer davantage, nous concluons en faisant appel à [B&N, Theorem 11.29, p. 211–214], dont le cas pénible est justement quand T est fini.

Le point (iii) se vérifie localement. □

Le résultat suivant nous permettra de vérifier, dans la section 5, qu'en fait la conjecture d'algébricité élimine les groupes de Frobenius non scindés.

Proposition 4.3. *Soient F un groupe de Frobenius de rang de Morley fini, T un sous-groupe malnormal de F, et G un sous-groupe définissable de F qui soit isomorphe à un groupe algébrique simple sur un corps algébriquement clos. Alors, ou bien G est contenu dans un conjugué de T, ou bien G, ainsi que son normalisateur, sont inclus dans $U(T)$.*

Démonstration. Comme G n'est pas un groupe de Frobenius, ou bien il est contenu dans un conjugué de T, ou bien il est disjoint de tous (lemme 0 : pas besoin de logique). S'il est inclus dans $U(T)$, on nomme Θ l'intersection de T et du normalisateur de G.

On fait alors intervenir une des conséquences les plus subtiles de la structure d'un groupe algébrique simple : si H est un groupe d'automorphismes de G telle que l'action de H sur G reste de rang de Morley fini, alors chaque point de H est un automorphisme rationnel, définissable dans le corps de base de G (et en fait dans la structure de groupe de G), et H° est formé d'automorphismes intérieurs [Altınel et al. 2008, p. 134; B&N, p. 124]. Comme les automorphismes intérieurs ont une infinité de points fixes, Θ est fini, et $G\Theta$ est algébrique ; comme il ne peut pas être un groupe de Frobenius algébrique, Θ est trivial. Le même raisonnement vaut pour les conjugués de T. □

Dans un groupe de Frobenius algébrique, T° est commutatif [B&N, Lemma 11.39, p. 218]. Dans la section 6, cette propriété sera étendue au contexte pseudo-localement fini.

5. Structure des groupes de Frobenius de rang de Morley fini

Nous allons voir dans cette section que l'étude des groupes de Frobenius de rang de Morley fini s'aborde différemment suivant que le complément est fini ou infini, et que le deuxième cas se ramène pour l'essentiel à celui où il est connexe.

Proposition 5.1. *Soit F un groupe de rang de Morley fini ayant un sous-groupe malnormal T.*

(i) *T est fini si et seulement si F° est inclus dans $U(T)$.*

(ii) *F est connexe si et seulement si T est connexe.*

(iii) *T° est l'intersection de F° et de T.*

(iv) *Tout sous-groupe V de F, normal, définissable, connexe et non inclus dans $U(T)$, agit transitivement sur les conjugués de T, si bien que $VT = F$. Cette hypothèse implique que T est infini, et alors elle s'applique à $V = F^\circ$. Si T remplit F, alors $T \cap V$ remplit V.*

(v) $2 \cdot \mathrm{RM}(T) \leq \mathrm{RM}(F)$; *si T ne remplit pas F, alors* $\mathrm{RM}(F) \leq 2 \cdot \mathrm{RM}(U(T))$; *si F est connexe, alors* $\mathrm{RM}(U(T)) < \mathrm{RM}(F)$.

Démonstration. (i) Si l'intersection de T et de F° n'est pas triviale, elle est infinie d'après la démonstration de la proposition 4.1.

(ii) Si T est connexe, F est la réunion des conjugués de T, qui sont tous inclus dans F°. Si F est connexe, T est infini, et la réunion de ses conjugués est une partie générique de F ; si T n'était pas connexe, T serait partitionné en deux sous-ensembles génériques, la réunion des conjugués de T° et son complément.

(iii) C'est vrai si T est fini d'après le point (i) ; quand T est infini, $T \cap F^\circ$ est non triviale, donc malnormale dans F°, et connexe d'après le point (ii).

(iv) Soit T_1 un conjugué de T au sens de F. La réunion des conjugués de $T \cap V$ dans V, ainsi que celle des conjugués de $T_1 \cap V$ dans V, sont génériques, et doivent avoir une intersection non triviale ; $T \cap V$ et $T_1 \cap V$ sont donc conjugués par un point de V, qui conjugue aussi T et T_1.

Tout point de F est congru modulo V a un point du normalisateur de T, qui est autonormalisant.

Comme V est connexe, il est inclus dans F° ; d'après le point (i), T est infini, et F° n'est pas inclus dans $U(T)$.

Si T remplit F, V est la réunion des conjugués de $T \cap V$ au sens de F, qui sont aussi ses conjugués au sens de V.

(v) Dans l'action de F sur les conjugués de T, ce dernier ne fixe que T, si bien qu'il agit injectivement sur ses autres conjugués.

Si a est un point non trivial de $U(T)$, son centralisateur et sa classe de conjugaison sont inclus dans $U(T)$, et $\mathrm{RM}(U(T)) \geq \mathrm{RM}(a^F) \geq \mathrm{RM}(F) - \mathrm{RM}(U(T))$. Si F est connexe, $U(T)$ n'est pas générique car son complémentaire l'est. □

Note. Si le sous-groupe malnormal T est fini, il est montré [B&N, p. 210] que F est scindé, de la forme VT ; comme alors le groupe V contient F°, il est définissable.

Proposition 5.2. *On considère un groupe de Frobenius F connexe de rang de Morley fini.*

(i) *Deux sous-groupes malnormaux de F ont des conjugués non disjoints.*

(ii) *Les sous-groupes malnormaux minimaux de F sont conjugués.*

(iii) *Quels que soient le sous-groupe malnormal T et le point a de F, le groupe engendré par T et a est le groupe* (*définissable et connexe*) *engendré par T et T^a.*

(iv) *Si H est un sous-groupe de F contenant un sous-groupe malnormal T, il est définissable, connexe et autonormalisant. Les conjugués de T dans F qui ont une intersection non triviale avec H sont inclus dans H, et conjugués de T dans H ; $H \cap U(T)$ est formé de* 1 *et des points de H qui ne sont pas dans un conjugué de T au sens de H.*

(v) *Quel que soit le sous-groupe malnormal T, le groupe engendré par $U(T)$ est définissable.*

Démonstration. (i) La réunion des conjugués du premier intersecte non trivialement la réunion des conjugués du second.

(ii) Si T et T' sont malnormaux minimaux, pour un certain a, ${}^aT \cap T'$ n'est pas trivial, et est donc malnormal ; par minimalité, ${}^aT = T'$.

(iii) C'est vrai si a est dans T. Sinon, comme T est connexe, ${}^aT = aTa^{-1}$ l'est aussi, ainsi que le groupe G engendré par T et ${}^aT$. Comme T et ${}^aT$ sont malnormaux dans G, à l'intérieur de G la réunion des conjugués de T et celle des conjugués de ${}^aT$ sont des parties génériques, qui ne peuvent être disjointes ; par conséquent T et ${}^aT$ sont conjugués par un g de G, et $g^{-1}a$ est dans G, qui est donc aussi le groupe engendré par T et a.

(iv) Comme, d'après (iii), H est engendré par des conjugués de T, il est définissable et connexe. Le reste vient de ce que, si un conjugué T' de T n'est pas disjoint de H, à l'intérieur de H la réunion des conjugués de T ne peut être disjointe de celle des conjugués de $T' \cap H$; ce dernier est conjugué de T dans H, et en fait égal à T'.

(v) Soit H le plus grand sous-groupe définissable connexe elliptiquement engendré par $U(T)$. Modulo H, $U(T)$ est un ensemble fini normal, donc central, dans le groupe connexe F/H ; comme il est clos par puissances, il engendre un groupe fini. □

Proposition 5.3. *On considère un groupe de Frobenius F, de rang de Morley fini, ayant un sous-groupe malnormal T infini.*

(i) [B&N, Proposition 11.24, p. 208] *Tout groupe V contenu dans $U(T)$ et normalisé par T° est définissable et connexe.*

(ii) *Il y a un plus grand sous-groupe normal contenu dans $U(T)$ et donnant avec T une scission nette.*

(iii) *Si F est scindé relativement à T, étant produit semi-direct de V par T, ce groupe V est définissable et connexe, et F° est le produit semi-direct de V et de T°.*

(iv) *Si V est un sous-groupe normal de F contenu dans $U(T)$, ou bien $F = VT$, ou bien l'image de T reste malnormale dans le quotient F/V. Si $U(T)$ n'est pas un groupe, F est engendré par T et $U(T)$; sinon F est scindé, ou bien le quotient $F/U(T)$ est un groupe de Frobenius plein.*

Démonstration. (i) Le groupe $G = VT^\circ$ est un produit semi-direct, et comme un point $v \neq 1$ de V ne peut normaliser ni T ni T°, ce dernier est malnormal dans G. On montre alors, comme dans la proposition 5.2(iii), que le groupe engendré par T° et v est définissable et connexe, car il est identique à celui engendré par T° et ${}^vT^\circ$; il en suit que G lui-même, étant engendré par les ${}^vT^\circ$, est définissable et connexe.

Les classes de conjugaison (au sens de G) des points de V sont incluses dans V. On considère son plus petit sous-groupe V° définissable connexe elliptiquement engendré par un nombre fini d'entre elles [Poizat 2018; 2021] ; V° est normal dans G, et on sait que chaque point de V a une classe de conjugaison finie modulo V°. Comme G/V° est connexe, V y est central modulo V°.

Soit alors H le groupe $V^\circ T^\circ$. Pour tout v de V, ${}^vT^\circ$ est inclus dans $V^\circ T^\circ = H$, et est donc conjugué de T° dans H. On en déduit que v est dans H, de la forme tu avec t dans T° et u dans V°, ou encore qu'il est dans V°. En conclusion $V = V^\circ$.

(ii) Conséquence du corollaire 3.4 et de la connexité des groupes considérés.

(iii) Le groupe V satisfait aux hypothèses du point (i), et VT° a même rang de Morley que F°.

(iv) Puisque $F = F^\circ T$ d'après la proposition 5.1(iv), si $F^\circ = VT^\circ$, alors $F = VT$; et si $F = VT$, sa composante connexe F° est VT°.

Sinon, considérons un point a de VT° tel que T° et ${}^aT^\circ$ ne soient pas disjoints modulo V. L'intersection S de ${}^aT^\circ$ et de VT° n'est donc pas triviale ; elle est par conséquent malnormale dans VT°, comme l'est T°. Comme VT° est connexe, S et T° y ont des conjugués non disjoints, et on trouve v dans V, t et t' non triviaux dans T°, tels que $ata^{-1} = vt'v^{-1}$; par conséquent $v^{-1}a$ est dans T° et a est dans T° modulo V. On voit que T° reste malnormal dans le quotient F°/V.

Si $U(T)$ n'est pas un groupe, on procède comme dans le lemme 1.2(iv) ; si c'est un groupe, il s'agit du fait 3.3, qui n'a rien à voir avec la finitude du rang. □

Remarque. La démonstration du point (iv) doit beaucoup à [B&N, Lemma 11.37, p. 217]. Il est assez troublant, car T reste malnormal dans le quotient sans que la netteté de la scission soit garantie, sans qu'on soit certain qu'aucun point de $U(T)$ n'entre dans T. On méditera également sur le corollaire 11.24 à la page 209 et sur l'exercice 5 aux pages 71 et 380.

Proposition 5.4. *On considère un groupe de Frobenius F, de rang de Morley fini, ayant un sous-groupe malnormal T infini.*

(i) *$U(T)$ contient un plus grand sous-groupe $R(F)$ normal résoluble. Il est définissable, connexe et nilpotent ; c'est le plus grand sous-groupe nilpotent normal dans F (il est indépendant du choix de T).*

(ii) *Si $F \neq R(F)T$, le quotient $F/R(F)$, ainsi que sa composante connexe $F^\circ/R(F)$, sont des groupes de Frobenius semi-simples (sans sous-groupe normal commutatif différent de $\{1\}$).*

(iii) *Quand $R(F) = \{1\}$, le socle de F° est composé soit d'un seul groupe de Frobenius simple non disjoint de T, qui engendre F avec T, soit est contenu dans $U(T)$, étant produit d'un nombre fini de groupes simples ; aucun de ces groupes simples n'est isomorphe à un groupe algébrique.*

(iv) *Sous la conjecture d'algébricité, un groupe de Frobenius de rang de Morley fini F connexe est scindé. Sa base est $R(F)$ et ses compléments sont tous conjugués.*

Démonstration. (i) Supposons que $U(T)$ contienne un sous-groupe normal R résoluble non trivial. Il contient un groupe abélien caractéristique non trivial A, qui est lui aussi normal dans F ; il est définissable et connexe d'après la proposition 5.3(i), et par conséquent infini. Comme il donne une scission nette, l'hypothèse se reproduit dans le quotient F/A (fait 3.3), et comme le rang de Morley diminue on finit par inclure R dans un groupe R_1 résoluble, normal, définissable et connexe, et contenu dans $U(T)$.

Si R' est un deuxième sous-groupe normal résoluble inclus dans $U(T)$, le passage au quotient F/A montre par induction que le groupe engendré par R et R', qui est résoluble et normal, est lui aussi inclus dans $U(T)$. D'où l'existence de $R(T)$, qui est définissable et connexe comme le sont tous les groupes normaux contenus dans $U(T)$.

Pour voir que $R(F)$ est nilpotent, nous redémontrons le cas facile de [B&N, Theorem 11.29, p. 211] : on considère un sous-groupe abélien définissable connexe infini B de T ; d'après le corollaire 3.2, pour tout $b \neq 1$ de B, le commutateur $[b,x] = bxb^{-1} \cdot x^{-1}$ définit une bijection de $R(F)$ dans $R(F)$, si bien que $R(F)$ est le dérivé du groupe résoluble connexe $R(F)B$. Selon un théorème dû à Ali Hoca Effendi, il est nilpotent [Poizat 1987, p. 94].

Enfin, si N est un sous-groupe nilpotent normal dans F, son centre est aussi normal, et contenu dans $U(T)$, ainsi que N.

(ii) Quand F n'est pas le produit semi-direct $R(F)T$, l'image de T modulo $R(F)$, qui est isomorphe à T, est malnormale dans le quotient $F/R(F)$, et celle de T° est malnormale dans sa composante connexe $F^\circ/R(F)$. Ces deux quotients n'ont pas de groupes résolubles normaux non triviaux.

(iii) Le socle (c'est-à-dire le groupe engendré par les sous-groupes normaux minimaux) du groupe connexe semi-simple F° est formé d'un produit de groupes simples normaux dans T° [Poizat 1987, p. 97]. Si l'un d'entre eux coupe T, c'est un groupe de Frobenius simple, qui n'est pas algébrique. Comme il coupe chaque conjugué de T, il est le seul groupe du socle, car il ne peut commuter avec un groupe de même espèce, ni avec un point non trivial de $U(T)$. Comme il contrôle la conjugaison de T, il engendre F avec celui-ci.

Dans le cas contraire, le socle de F est un produit de groupes simples contenus dans $U(T)$. Comme ils commutent, leur produit l'est aussi ; ce ne sont pas non plus des groupes algébriques (proposition 4.3).

(iv) Sous la conjecture d'algébricité, il ne peut pas y avoir de groupes de Frobenius de rang de Morley fini connexes semi-simples. Pour n'importe quel complément T, $F = R(F)T$. Tous les compléments sont définissablement isomorphes à $F/R(F)$; ils sont minimaux, et conjugués d'après la proposition 5.2(ii). □

Quand le socle est un groupe de Frobenius, notre analyse s'arrête là ; notons bien que nous n'avons pas affirmé que ce socle contenait la base $U(T)$. Dans le deuxième cas, nous ne savons pas si nous pouvons la poursuivre par une scission nette. Du moins avons-nous précisé les points où B&N ne peuvent éviter les groupes simples non algébriques dans leur analyse de contre-exemples minimaux aux pages 215–219, et, sous la conjecture d'algébricité, nous avons étendu les théorèmes de Frobenius et de Thompson aux groupes de Frobenius connexes de rang de Morley fini. Le sort des compléments sera réglé dans la dernière section.

En attendant, désireux de montrer tout ce que nous savons sur les groupes de Frobenius, nous ajoutons un minuscule supplément.

Proposition 5.5. *Soient F un groupe de Frobenius de rang de Morley fini*, connexe *et semi-simple*, *et T un de ses sous-groupes malnormaux. Alors* :

(i) *Tout $a \neq 1$ de $U(T)$ a une infinité de conjugués sous l'action du centralisateur d'un certain point b de $U(T)$.*

(ii) *Tout sous-groupe abélien V contenu dans $U(T)$ et normalisé par ce dernier est trivial.*

Démonstration. (i) Chaque $b \neq 1$ de $U(T)$ a un centralisateur infini $Z(b)$ — pas besoin d'[Altınel et al. 2008, Chapter 4, Corollary 4.18] : $U(T)$ n'est pas générique car son complémentaire l'est — qui est inclus dans $U(T)$. Si a n'a qu'un nombre fini de conjugués sous l'action de $Z(b)$, il est centralisé par $Z(b)^\circ$; si cela se produit pour chaque b, a est dans l'intersection Z des centralisateurs $Z(Z(b)^\circ)$, qui est un groupe définissable, contenu dans $U(T)$, et normal dans F. Ce Z est donc connexe d'après la proposition 5.3(i), et le centralisateur de a dans Z est un groupe infini, dont la composante connexe est centrale dans Z ; ce dernier a un centre non trivial, ce qui contredit la semi-simplicité de F.

(ii) Si V est un contre-exemple, on peut le supposer définissable, en le remplaçant par le centre de son centralisateur. Il ne peut être fini d'après le point (i). Si on le prend définissable minimal, il est connexe ; V est alors disjoint de chacun de ses conjugués W, au sens de F, distinct de lui-même, si bien que V et W commutent (comme ils se normalisent l'un l'autre, le commutateur d'un point de V et d'un point de W est dans leur intersection). Les conjugués de V engendrent donc un groupe commutatif, normal dans F, ce qui contredit sa semi-simplicité. □

Nous nous soucions maintenant du sort des groupes pleins.

Proposition 5.6. *Soit F un groupe connexe de rang de Morley fini, possédant des sous-groupes malnormaux remplissants.*

(i) *Si H est un sous-groupe propre de F contenant un groupe malnormal remplissant T, il est lui-même rempli par T, et c'est un sous-groupe malnormal dans F qui le remplit.*

(ii) *L'intersection de deux sous-groupes malnormaux non disjoints T et T' est remplissante si et seulement si T et T' le sont.*

(iii) *Les sous-groupes malnormaux remplissants minimaux sont conjugués ; ce ne sont pas des groupes de Frobenius pleins.*

(iv) *F n'a pas d'involutions, et chacun de ses sous-groupes finis est contenu dans un sous-groupe malnormal remplissant minimal.*

(v) *Chaque sous-groupe G définissable et connexe de F non contenu dans un groupe remplissant minimal est un groupe de Frobenius plein, et c'est vrai en particulier de son socle, qui est composé d'un seul groupe simple.*

Démonstration. (i) H est un groupe définissable connexe, qui est rempli par les conjugués de T qui l'intersectent non trivialement ; ces derniers sont tous conjugués de T dans H. Comme $U(T) \cap H = \{1\}$, H est malnormal dans F, et comme F est rempli par T il est aussi rempli par H.

(ii) Si $T \cap T'$ remplit F, nous avons vu que T et T' aussi. Réciproquement, si T' remplit F, $T \cap T'$ remplit T ; si de plus T remplit F, $T \cap T'$ aussi.

(iii) Deux sous-groupes malnormaux remplissants ont des conjugués qui se coupent non trivialement. Si T a lui même un sous-groupe malnormal le remplissant, ce dernier remplit F, et T n'est pas minimal.

(iv) Soient i et j des involutions situées dans des conjugués différents du groupe malnormal remplissant T. Le produit ij, étant inversé par i comme par j, ne peut être dans un conjugué de T ; donc $ij = 1$, $i = j$, ce qui ne se peut.

Soit T un sous-groupe remplissant minimal, et φ un groupe fini non trivial qui n'est pas contenu dans un conjugué de T. Comme φ n'est pas un groupe de Frobenius plein, les intersections non triviales de φ avec un conjugué de T se répartissent en au moins deux classes de conjugaisons, celle de θ et de θ', et le calcul fait dans le lemme 1.2(i) rend la chose impossible.

(v) Si T est malnormal remplissant et G n'est pas inclus dans T, il est rempli par toutes les intersections $G \cap T^a$ qui ne sont pas triviales ; comme il est connexe, elles sont toutes conjuguées dans G.

Comme $U(T)$ est trivial, F est semi-simple, et on est dans le premier cas de la proposition 5.4(iii). □

Nous dirons qu'un groupe de Frobenius plein, de rang de Morley fini et connexe, est *petit* si tous ses sous-groupes définissables maximaux sont conjugués (et malnormaux remplissants). Nous mettons l'emphase sur cette notion, car la carence d'un petit groupe plein en automorphismes involutifs est la partie facile, et de portée générale, de l'argumentation par contradiction de [Frécon 2018] ; la partie délicate, et de portée limitée, réside dans l'étude du symétron de ce groupe.

Proposition 5.7. (i) *Si F est un groupe connexe ayant un sous-groupe malnormal remplissant T, tout sous-groupe définissable G de F qui est minimal pour n'être pas inclus dans un conjugué de T est petit.*

(ii) *Un petit groupe est simple et n'a pas d'automorphisme définissable involutif non trivial.*

Démonstration. (i) Comme G est connexe, toutes les intersections non triviales de G avec un conjugué de T sont conjuguées dans G à l'une d'entre elles, soit T_1. Si H est un sous-groupe définissable propre de G, il est contenu dans un conjugué de T_1, qui est malnormal remplissant : c'est vrai s'il est fini, et sinon c'est vrai parce que c'est vrai pour H° par minimalité de G.

(ii) Les sous-groupes définissables propres de G, étant contenus dans un conjugué de son groupe malnormal maximal T, ne sont pas normaux.

Soit s un automorphisme involutif de G. Comme G n'a pas d'involutions, il est uniquement 2-divisible, et chaque point de G s'écrit de manière unique comme produit d'un point fixé par s et d'un point inversé par s ; voir par exemple [Poizat 2018]. Comme G n'est pas commutatif, s a des points fixes non triviaux, et si s

n'est pas l'identité il a des points inversés non triviaux. Par ailleurs s permute les conjugués de T, qui sont les sous-groupes malnormaux maximaux de G.

Si s est définissable et différent de l'identité, à conjugaison près son groupe de points fixes est contenu dans T, qui est normalisé par s. Il est nécessaire qu'il y ait d'autres points inversés que ceux de T ; si un autre conjugué T' de T en contient un, il est aussi normalisé par s, et comme il ne contient pas de points fixes c'est un groupe commutatif inversé par s. Comme T et T' sont conjugués, un point fixe non trivial a un conjugué inversé, si bien qu'il est conjugué de son inverse, ce qui produit des involutions. □

6. Où on retrouve un vieil ami

Dans cette section de conclusion, nous rappelons quelques questions, liées aux comportement des corps dans une situation de rang de Morley fini, qui surgissent de l'étude des groupes de Frobenius où $U(T)$ est un groupe commutatif. Si, dans les années 1970, était répandue la croyance optimiste, ou naïve, que tous les problèmes concernant les corps dans un environnement de rang de Morley fini étaient réglés par [Macintyre 1971], on sait aujourd'hui qu'il n'en est rien, à la lumière des travaux de Frank Wagner [2001].

Proposition 6.1. *Soit T un groupe commutatif infini agissant librement sur un groupe abélien U, de sorte que le produit semi-direct UT soit un groupe* (*de Frobenius*) *de rang de Morley fini. On peut alors définir dans ce dernier un ou plusieurs corps infinis $K_1, \dots, K_n$, des isomorphismes $f_1, \dots, f_n$ entre T et des sous-groupes de $K_1^*, \dots, K_n^*$, ainsi qu'une décomposition du sous-groupe de U engendré par ses sous-groupes normalisés par T minimaux en tant que somme directe d'espaces vectoriels V_1 sur $K_1, \dots, V_n$ sur K_n, sur laquelle T agit diagonalement comme $(f_1(t), \dots, f_n(t))$.*

Démonstration. Nous notons additivement la loi de groupe de U et multiplicativement l'action de T sur ce dernier.

Soit U_1 le plus grand sous-groupe de U définissable connexe elliptiquement engendré par la réunion d'un nombre fini de Ta. Si b est hors de U_1, Tb est fini modulo U_1, si bien que b commute avec des points de l'image de T dans le quotient UT/U_1, ce qui n'est pas conforme à la surjectivité des adjointes $tx - x$. Par conséquent $U = U_1 = Ta_1 + \cdots + Ta_m - Ta_{m+1} - \cdots - Ta_n$, ce qui permet de définir l'anneau commutatif $R = Z[T]$ des endomorphismes de U engendré par T ; en effet, tout point de U se représente dans le système générateur des a_i par une colonne de coordonnées de longueur n à valeur dans T, deux colonnes étant dites équivalentes si elles représentent le même vecteur. Ces endomorphismes sont représentés par les matrices $n \times n$ à valeur dans T qui respectent l'équivalence, propriété qui se définit en utilisant les a_i comme paramètres.

La même démonstration montre que tout sous-R-module de U, et en particulier U lui-même, est connexe (voir la proposition 5.3(i)) et finiment engendré. Un R-module minimal est engendré par chacun de ses points non nuls, si bien que R agit sur lui comme un corps. Le groupe V engendré par les R-modules minimaux est somme directe d'un nombre fini d'entre eux, si bien que la restriction de R à V n'a pas d'éléments nilpotents ; d'après [Altınel et al. 2008, Chapter 1, Lemma 4.1, p. 44], c'est un produit de corps $K_1 \times \cdots \times K_n$. Si on note V_i l'annulateur de $K_1 \times \cdots \times K_{i-1} \times \{0\} \times K_{i+1} \times \ldots K_n$, où $i = 1, \ldots, n$, on obtient la décomposition de V cherchée ; l'action de T sur V_i est celle d'un sous-groupe $f_i(T)$ de K_i^*. $\square$

Exemples. On considère trois actions de $T = K^*$ sur $V = K^+ \times K^+$:

(i) t agit comme la matrice diagonale (t, t^{-1}) ; si t_1 et t_2 sont algébriquement indépendants, $t_1 + t_2$ et $t_1^{-1} + t_2^{-1}$ le sont aussi, si bien que l'anneau R est le produit de corps $K \times K$: il y a deux corps (tous deux isomorphes à K), et deux espaces vectoriels (de dimension un).

(ii) En caractéristique p, t agit comme (t, t^p) ; les points de la forme (t, t^p) forment un corps L (isomorphe à K), et nous avons affaire en réalité à l'action diagonale de L^* sur $L^+ \times L^+$: il n'y a qu'un seul corps, et qu'un seul espace vectoriel de dimension deux.

(iii) t agit comme (t^2, t^3) : ce n'est pas une action libre.

Notation. Dans un groupe commutatif U de rang de Morley fini, l'intersection de deux sous-groupes définissables sans torsion est divisible, si bien qu'il y en a un plus grand, que nous notons U_0 ; pour chaque nombre premier p, nous notons U_p son plus grand sous-groupe définissable connexe d'exposant p.

Proposition 6.2. *Soit T un groupe agissant librement sur un groupe abélien U, de sorte que le produit semi-direct UT soit un groupe connexe de rang de Morley fini ; alors tout sous-groupe résoluble définissable connexe de T est commutatif. Plus précisément, si $U_p \neq \{0\}$ pour un p premier, la composante connexe du normalisateur d'un sous-groupe infini abélien définissable de T est toujours commutative ; si $U_0 \neq \{0\}$, T est un groupe définissablement linéaire sans unipotents sur un corps algébriquement clos de caractéristique nulle.*

Démonstration. Soit M un sous-groupe de T abélien infini définissable. Le normalisateur de M induit un groupe d'automorphismes de l'anneau R_M engendré par l'action de M sur U, et agit semi-R_M-linéairement suivant la formule $s(\lambda x) = (s\lambda s^{-1}) \cdot sx$. Il normalise le groupe V_M engendré par les R_M-modules minimaux, et permute les corps dont la restriction de R_M à V_M est le produit (car ce sont les quotients de R_M par ses idéaux maximaux). Sa composante connexe N fixe chacun de ces corps K ; elle définit un groupe d'automorphismes de K, qui est nécessairement réduit à l'identité. Donc N agit R_M-linéairement sur V_M ; il normalise l'espace vectoriel V associé à K, et agit sur lui K-linéairement.

Comme il agit sur V sans points fixes, N ne contient pas d'unipotents. Ses sous-groupes définissables résolubles sont triangularisables par fini d'après le théorème de Lie–Kolchin–Maltsev ; comme les commutateurs d'un groupe triangulaire sont unipotents, leurs composantes connexes sont commutatives. Si N n'est pas commutatif, nous notons C son plus grand sous-groupe définissable connexe normal commutatif, et N_1 et C_1 leurs clôtures de Zariski respectives. Le groupe N/C est un sous-groupe du groupe algébrique linéaire N_1/C_1 ; il peut avoir un centre fini, mais après quotient par ce dernier on obtient un groupe linéaire semi-simple dont le socle est formé de groupes simples linéaires à borels commutatifs.

D'après [Poizat 2001], c'est impossible s'il y a un corps de caractéristique p, et N est alors commutatif. Comme un groupe infini résoluble connexe normalise un sous-groupe commutatif infini connexe, T satisfait bien à la condition dite dans l'énoncé.

S'il y a un corps de caractéristique nulle, nous considérons l'action de T sur U_0. Pour chaque rationnel r on voit, en exprimant les coordonnées des ra_i dans le système générateur des a_i, que l'anneau R_M engendré par M contient la multiplication par r : l'intersection R de tous les R_M, qui est celle d'un nombre fini d'entre eux, est un anneau infini. Les R-modules sont définissables et connexes (ils sont divisibles sans torsion) ; T normalise l'espace engendré par les R-modules minimaux, et fixe chacun des corps associés à l'action de R sur ce dernier. Si K est l'un d'entre eux et si V est l'espace vectoriel associé, T agit K-linéairement sur V, et sans points fixes, et en particulier sans unipotents. Par conséquent, tout ses sous-groupes résolubles définissables connexes, qui sont triangularisables, sont diagonalisables. □

Remarques. (i) Samuel Zamour a montré dans sa thèse [2022], qui a été conduite de façon totalement indépendante des recherches exposées dans cet article, qu'un groupe de Frobenius de rang de Morley fini connexe et définissablement linéaire en caractéristique p était résoluble.

(ii) Une conséquence d'[Altınel et al. 2019, Theorem 8] est que, quand $U_p \neq \{0\}$ pour un p premier (autre que 2 !) et T contient une involution, il est commutatif (on peut aussi utiliser [Altınel et Burdges 2008, Theorem 2]). Cela éclaire notre proposition 2.1(iv).

La proposition 6.2 permet de donner à [B&N, Theorem 11.34, p. 216] une démonstration qui précise la place des sections simples non algébriques dans un contre-exemple, même s'il n'est pas minimal.

Corollaire 6.3. *Un groupe de Frobenius F connexe de rang de Morley fini qui est pseudo-localement fini, ou même qui ne contredit pas la conjecture d'algébricité* (*dans le sens où toutes ses sections définissables simples sont algébriques*), *est* *résoluble. Le groupe malnormal T est isomorphe à un sous-groupe multiplicatif de un ou plusieurs corps définissables dans F, et $U(T)$ est le groupe dérivé de F.*

Démonstration. Si $U(T)$ est trivial ou n'est pas un groupe nilpotent, on trouve des sections simples non algébriques dans le socle du quotient de F par $R(F)$ (proposition 5.4(iii)). Quand $U(T)$ est nilpotent non trivial, on obtient une action libre de T sur son centre. Si T n'est pas commutatif, on considère un sous-groupe S de T définissable connexe et minimal pour n'être pas commutatif. Le quotient de S par son radical est un groupe simple (en fait un groupe de Frobenius plein dont les compléments sont des bons tores en caractéristique p, un groupe définissablement linéaire en caractéristique nulle) qui n'est pas un groupe algébrique. □

Exemples. (i) Si F est un groupe de Frobenius algébrique connexe sur un corps algébriquement clos de caractéristique différente de 2, le tore T contient une involution, si bien que $U(T)$ est commutatif d'après la proposition 2.1.

(ii) En caractéristique 2, le groupe diagonal $T = (x \quad 1 \quad x^{-1})$ agit librement sur les matrices triangulaires unipotentes d'ordre 3, si bien que le groupe

$$\begin{pmatrix} x & u & v \\ 0 & 1 & w \\ 0 & 0 & x^{-1} \end{pmatrix}$$

est de Frobenius ; il en est de même du groupe

$$\begin{pmatrix} x & u & v \\ 0 & 1 & u \\ 0 & 0 & x^{-1} \end{pmatrix},$$

chez qui $U(T)$ est commutatif d'exposant 4.

(iii) On obtient des exemples semblables, de rang de Morley fini mais non algébriques, en prenant un corps avec un groupe multiplicatif T infini sans involutions, dont on est sûr de l'existence en caractéristique nulle grâce à [Baudisch et al. 2009].

Tout cela mène à quatre questions sur les corps de rang de Morley fini :

Question 1. *Le groupe additif d'un corps de rang de Morley fini peut-il avoir un groupe définissable connexe infini non commutatif d'automorphismes ?*

Cette question, qui ne se pose pas en caractéristique nulle, apparaît dans la proposition 3.12, p. 117 de [Poizat 1987]. Elle suppose l'existence d'un sous-groupe infini définissable propre du groupe multiplicatif d'un corps de caractéristique p, ce qui, d'après [Wagner 2003], implique qu'il n'y a qu'un nombre fini de p-nombres de Marin Mersenne (un frère mineur ayant vécu au début du XVII^e^ siècle). Cette hypothèse arithmétique semble très peu probable, et surtout hors de portée des méthodes de la théorie des modèles. Depuis [Altınel et al. 2008], tout le monde, nous compris, s'emploie à la contourner : en effet, les démonstrations ci-dessus se simplifient drastiquement si on suppose que tout sous-groupe définissable infini commutatif de T est connexe, sans sous-groupes propres définissables autres que cycliques finis.

En amalgamant des corps de caractéristique p au-dessus de leur groupe additif, on doit pouvoir obtenir des corps non définissablement isomorphes de même groupe additif.

Question 2. *S'il existe un isomorphisme définissable entre les groupes multiplicatifs de deux corps de rang de Morley fini, sont-ils définissablement isomorphes ?*

Question 3. *S'il existe un isomorphisme définissable entre des sous-groupes multiplicatifs infinis de deux corps de rang de Morley fini, sont-ils définissablement isomorphes ?*

Question 4. *La structure formée d'un corps de caractéristique nulle, avec un sous-groupe multiplicatif infini propre, peut-elle avoir un rang de Morley fini et être pseudo-localement finie ?*

Nous terminons par un bref examen du cas minimal de groupe de Frobenius de rang de Morley fini.

Proposition 6.4. *Soit F un groupe de rang de Morley fini connexe, avec un sous-groupe malnormal T, tel que* $\mathrm{RM}(F) = 2 \cdot \mathrm{RM}(U(T))$; *alors $U(T)$ est un groupe commutatif qui est d'exposant p pour un certain nombre premier, ou bien sans torsion divisible. En outre, si* $\mathrm{RM}(T) = \mathrm{RM}(U(T))$, *$F$ est scindé ; si de plus T est commutatif, F est isomorphe au produit semi-direct de K^+ par K^*, où K est un corps définissable.*

Démonstration. Considérons $a \neq 1$ dans $U(T)$. Sa classe de conjugaison C et son centralisateur $Z(a)$ sont inclus dans $U(T)$, et $\mathrm{RM}(F) = 2 \cdot \mathrm{RM}(U(T)) = \mathrm{RM}(C) + \mathrm{RM}(Z(a))$; on en déduit que $\mathrm{RM}(C) = \mathrm{RM}(Z(a)) = \mathrm{RM}(U(T))$. Comme F est connexe, C est de degré de Morley 1 ; $U(T)$ se répartit donc en un nombre fini de classes de conjugaison génériques $C = C_0, \ldots, C_d$, sans compter l'élément neutre ; $Z(a)$ doit couper génériquement au moins une classe de conjugaison C'. Le centralisateur de C est celui d'un nombre fini de ses points, qui tous commutent avec le générique de C' ; comme il est normal il doit contenir C', et sa composante connexe est égale à une classe de conjugaison C'', augmentée de l'identité.

Il ne peut pas y avoir deux groupes composés d'une classe de conjugaison, car leur intersection devrait être triviale, et ils commuteraient l'un avec l'autre, si bien que $U(T)$ devrait contenir leur produit, qui est de dimension double. Donc l'une des classes de conjugaison, augmentée de 1, est un groupe A commutatif, qui est la composante connexe du centralisateur de $U(T)$. Dans le quotient F/A, l'image de

$U(T)$ est finie, et en fait triviale car la scission est nette : autrement dit $U(T) = A$. Comme tous les points non triviaux de A sont conjugués, ils ont tous même ordre.

Si $\mathrm{RM}(T) < \mathrm{RM}(U(T))$, $F/U(T)$ est rempli ; sinon F est scindé, et quand T est commutatif, comme il a même rang que A, la proposition 6.1 nous dit que nous avons affaire à un espace vectoriel de dimension 1 sur un corps, $A = K^+$, $T = K^*$. □

Remarque. Altınel et al. [2019] montrent que T est toujours commutatif quand $\mathrm{RM}(T) = \mathrm{RM}(U(T))$, sauf peut-être en caractéristique 2 ; on trouve dans [Zamour 2022] bien d'autres précisions sur les groupes exactement deux fois transitifs de rang de Morley fini.

Remerciements

Je remercie très sincèrement Erjan Baisal pour le résumé en russe, Frank Wagner pour celui en allemand, et la rapporteuse ou le rapporteur anonyme de la première version de cet article, dont les très nombreuses remarques m'ont permis — je l'espère — d'en améliorer la présentation et l'intelligibilité.

Bibliographie

[Altınel et Burdges 2008] T. Altınel et J. Burdges, "On analogies between algebraic groups and groups of finite Morley rank", *J. Lond. Math. Soc.* (2) **78**:1 (2008), 213–232. MR Zbl

[Altınel et al. 2008] T. Altınel, A. V. Borovik et G. Cherlin, *Simple groups of finite Morley rank*, Mathematical Surveys and Monographs **145**, Amer. Math. Soc., Providence, RI, 2008. MR Zbl

[Altınel et al. 2019] T. Altınel, A. Berkman et F. O. Wagner, "Sharply 2-transitive groups of finite Morley rank", prépublication, 2019, voir https://hal.science/hal-01935537v2.

[Baudisch et al. 2009] A. Baudisch, M. Hils, A. Martin-Pizarro et F. O. Wagner, "Die böse Farbe", *J. Inst. Math. Jussieu* **8**:3 (2009), 415–443. MR Zbl

[Borovik et Nesin 1994] A. Borovik et A. Nesin, *Groups of finite Morley rank*, Oxford Logic Guides **26**, The Clarendon Press, New York, 1994. MR Zbl

[Borovik et Poizat 1990] A. V. Borovik et B. P. Poizat, "Tores et p-groupes", *J. Symbolic Logic* **55**:2 (1990), 478–491. MR Zbl

[Cherlin 1979] G. Cherlin, "Groups of small Morley rank", *Ann. Math. Logic* **17**:1-2 (1979), 1–28. MR Zbl

[DeBonis et Nesin 1994] M. DeBonis et A. Nesin, "On split Zassenhaus groups of mixed characteristic and of finite Morley rank", *J. London Math. Soc.* (2) **50**:3 (1994), 430–439. MR Zbl

[Delahan et Nesin 1993] F. Delahan et A. Nesin, "Sharply 2-transitive groups revisited", *Doğa Mat.* **17**:1 (1993), 70–83. MR Zbl

[Delahan et Nesin 1995] F. Delahan et A. Nesin, "On Zassenhaus groups of finite Morley rank", *Comm. Algebra* **23**:2 (1995), 455–466. MR Zbl

[Epstein et Nesin 1994] D. Epstein et A. Nesin, "On Frobenius groups of finite Morley rank, II", pp. 341–350 dans *Automorphisms of first-order structures*, édité par R. Kaye et D. Macpherson, Oxford Univ. Press, New York, 1994. MR Zbl

[Feit et Thompson 1963] W. Feit et J. G. Thompson, "Solvability of groups of odd order", *Pacific J. Math.* **13** (1963), 775–1029. MR Zbl

[Frécon 2018] O. Frécon, "Simple groups of Morley rank 3 are algebraic", *J. Amer. Math. Soc.* **31**:3 (2018), 643–659. MR Zbl

[Frobenius 1901] G. Frobenius, "Über auflösbare Gruppen, IV", *Berl. Ber.* (1901), 1216–1230. Zbl

[Hertzig 1961] D. Hertzig, "The structure of Frobenius algebraic groups", *Amer. J. Math.* **83** (1961), 421–431. MR Zbl

[Jaligot 2001] E. Jaligot, "Full Frobenius groups of finite Morley rank and the Feit–Thompson theorem", *Bull. Symbolic Logic* **7**:3 (2001), 315–328. MR Zbl

[Kegel et Wehrfritz 1973] O. H. Kegel et B. A. F. Wehrfritz, *Locally finite groups*, North-Holland Mathematical Library **3**, 1973. MR Zbl

[Macintyre 1971] A. Macintyre, "On ω_1-categorical theories of fields", *Fund. Math.* **71**:1 (1971), 1–25. MR Zbl

[Nesin 1989] A. Nesin, "Nonsolvable groups of Morley rank 3", *J. Algebra* **124**:1 (1989), 199–218. MR Zbl

[Nesin 1992] A. Nesin, "Notes on sharply 2-transitive permutation groups", *Doğa Mat.* **16**:1 (1992), 69–84. MR Zbl

[Nesin 1994] A. Nesin, "On Frobenius groups of finite Morley rank, I", pp. 325–339 dans *Automorphisms of first-order structures*, édité par R. Kaye et D. Macpherson, Oxford Univ. Press, New York, 1994. MR Zbl

[Olchanski 1982] A. I. Olchanski, "Groups of bounded period with subgroups of prime order", *Algebra i Logika* **21**:5 (1982), 553–618. En russe; traduit en anglais à *Algebra and Logic* **21** (1982), 369–418. MR Zbl

[Poizat 1987] B. Poizat, *Groupes stables : une tentative de conciliation entre la géométrie algébrique et la logique mathématique*, Nur al-Mantiq wal-Ma'rifah **2**, Bruno Poizat, Lyon, 1987. MR Zbl

[Poizat 2001] B. Poizat, "Quelques modestes remarques à propos d'une conséquence inattendue d'un résultat surprenant de Monsieur Frank Olaf Wagner", *J. Symbolic Logic* **66**:4 (2001), 1637–1646. MR Zbl

[Poizat 2018] B. Poizat, "Milieu et symétrie, une étude de la convexité dans les groupes sans involutions", *J. Algebra* **497** (2018), 143–163. MR Zbl

[Poizat 2021] B. Poizat, "Symétries et transvexions, principalement dans les groupes de rang de Morley fini sans involutions", *J. Symb. Log.* **86**:3 (2021), 965–990. MR Zbl

[Thomas 1983] S. Thomas, "The classification of the simple periodic linear groups", *Arch. Math.* (*Basel*) **41**:2 (1983), 103–116. MR Zbl

[Thompson 1959] J. Thompson, "Finite groups with fixed-point-free automorphisms of prime order", *Proc. Nat. Acad. Sci. U.S.A.* **45** (1959), 578–581. MR Zbl

[Wagner 2001] F. Wagner, "Fields of finite Morley rank", *J. Symbolic Logic* **66**:2 (2001), 703–706. MR Zbl

[Wagner 2003] F. O. Wagner, "Bad fields in positive characteristic", *Bull. London Math. Soc.* **35**:4 (2003), 499–502. MR Zbl

[Zamour 2022] S. Zamour, *Étude des quasi-groupes de Frobenius et des K-boucles de rang de Morley fini*, thèse de doctorat, Université Claude Bernard (Lyon 1), 2022, voir https://theses.hal.science/tel-03765160.

Received 2 Jun 2022. Revised 22 Dec 2022.

BRUNO POIZAT:

poizat@math.univ-lyon1.fr
Institut Camille Jordan, Université Claude Bernard (Lyon 1), Villeurbanne, France

Model Theory
Vol. 3, No. 1, 2024
https://doi.org/10.2140/mt.2024.3.101

Nonelementary categoricity and projective locally o-minimal classes

Boris Zilber

Given a cover $\mathbb{U}$ of a family of smooth complex algebraic varieties, we associate with it a class $\mathfrak{U}$, containing $\mathbb{U}$, of structures locally definable in an o-minimal expansion of the real numbers. We prove that the class is $\aleph_0$-homogenous over submodels and stable. It follows that $\mathfrak{U}$ is categorical in cardinality $\aleph_1$. In the case when the algebraic varieties are curves we prove that a slight modification of $\mathfrak{U}$ is an abstract elementary class categorical in all uncountable cardinals.

1. Introduction

1.1. Let $\mathrm{k}_0 \subseteq \mathbb{C}$, a countable subfield, $\{\mathbb{X}_i : i \in I\}$ a collection of nonsingular irreducible complex algebraic varieties (of dim > 0) defined over k_0 and $I := (I, \geq)$ a lattice with the minimal element 0 determined by unramified k_0-rational epimorphisms $\mathrm{pr}_{i',i} : \mathbb{X}_{i'} \to \mathbb{X}_i$, for $i' \geq i$. Let $\mathbb{U}(\mathbb{C})$ be a connected complex manifold and $\{f_i : i \in I\}$ a collection of holomorphic covering maps (local biholomorphisms)

$$f_i : \mathbb{U}(\mathbb{C}) \twoheadrightarrow \mathbb{X}_i(\mathbb{C}), \quad \mathrm{pr}_{i',i} \circ f_{i'} = f_i.$$

as illustrated by

$$\begin{tikzcd}
\mathbb{U}(\mathbb{C}) \arrow[d, "f_m"'] \arrow[dr, "f_n"] & \\
\mathbb{X}_m(\mathbb{C}) \arrow[r, "\mathrm{pr}_{m,n}"] \arrow[d, "\mathrm{pr}_{m,0}"'] & \mathbb{X}_n(\mathbb{C}) \arrow[dl, "\mathrm{pr}_{n,0}"] \\
\mathbb{X}_0(\mathbb{C}) &
\end{tikzcd}$$

MSC2020: 03C75.
Keywords: categoricity, o-minimal, quasiminimal.

1.2. In a number of publications, abstract elementary classes $\mathfrak{U}$ containing structures $(\mathbb{U}, f_i, \mathbb{X}_i)$, with an abstract algebraically closed field K instead of $\mathbb{C}$ (pseudoanalytic structures) have been considered; see [Zilber 2016] for a survey. A typical result is a formulation of a "natural" $L_{\omega_1,\omega}$-axiom system Σ which holds for $(\mathbb{U}(\mathbb{C}), f_i, \mathbb{X}_i(\mathbb{C}))$ and defines a class $\mathfrak{U}$ categorical in all uncountable cardinals. The proofs, in each case, rely on deep results in arithmetic geometry, moreover one often is able to show that the fact of categoricity of Σ implies the arithmetic results.

The above raised the question of whether an uncountably categorical AEC $\mathfrak{U}$ containing $(\mathbb{U}(\mathbb{C}), f_i, \mathbb{X}_i(\mathbb{C}))$ exists under general enough assumptions on the data, leaving aside the question of axiomatisability and related arithmetic theory.

The current paper answers this question in the positive at least in the case when the $\mathbb{X}_i$ are curves. We construct $\mathfrak{U}$ as the class of structures $\mathbb{U}(\mathrm{K})$ $(\mathrm{K} = \mathsf{R} + i\,\mathsf{R})$ locally definable (in the sense of M. Edmundo and others) in models R of an o-minimal expansion of the real numbers projected (restricted) to the language L_{glob} (global) the primitives of which are given by analytic subsets of $\mathbb{U}^m$ locally defined in the o-minimal structure. The main theorem states that, for the case when the complex dimension of $\mathbb{U}(\mathbb{C})$ is equal to 1, $\mathfrak{U}$ can be extended to a class of L_{glob}-structures which is an abstract elementary class categorical in all uncountable cardinals. For the general case we only were able to prove categoricity in $\aleph_1$.

1.3. Our main technical tool is a slightly generalised theory of K-analytic sets in o-minimal expansions of the real numbers developed by Y. Peterzil and S. Starchenko [2008]. We also make an essential use of the theory of quasiminimal excellence, especially the important paper by M. Bays, B. Hart, T. Hyttinen, M. Kesälä and J. Kirby [Bays et al. 2014].

Note that our main technical results effectively prove that the structures in $\mathfrak{U}$ are *analytic Zariski* in a sense slightly weaker than in [Zilber 2017], where we proved results similar to the current ones for an analytic Zariski class.

1.4. Most of our examples, see Section 2.3, have become objects of interest in the theory of o-minimality due to the Pila–Wilkie–Zannier method of counting special points of Shimura varieties and more generally; see the survey [Scanlon 2012]. Effectively, one counts points of $\mathbb{U}(L) \cap D \cap S$ where D is an open subset of $\mathbb{U}(\mathbb{C})$ definable in the o-minimal structure, S an L_{glob}-definable analytic subsets of $\mathbb{U}(\mathbb{C})$ and L a number field relevant to the case at hand.

At the same time one should note that in representing an L_{glob}-structure as $\mathbb{U}(\mathrm{K})$, $\mathrm{K} = \mathsf{R} + i\,\mathsf{R}$, there is a remarkable degree of freedom in the choice of a model R of the underlying o-minimal theory.

This raises a lot of questions on the interaction between the theory of AEC and o-minimality, the model theory–arithmetic geometry perspective of categorical classes and the o-minimal Pila–Wilkie–Zannier method.

2. Preliminaries

2.1. Let $\mathbb{R}_{\mathrm{An}}$ be an o-minimal expansion of the real numbers, $\mathbb{C} = \mathbb{R} + i\mathbb{R}$ in the language of rings and

$$\mathrm{Mod}_{\mathrm{An}} = \{\mathsf{R} : \mathsf{R} \equiv \mathbb{R}_{\mathrm{An}}\}$$

the class of models of the complete o-minimal theory $\mathrm{Th}(\mathbb{R}_{\mathrm{An}})$ in the language L_{An}. To avoid unnecessary complications we assume that L_{An} is a countable fragment of the full language of $\mathbb{R}_{\mathrm{An}}$.

We write K for the algebraically closed field $\mathrm{K}(\mathsf{R}) := \mathsf{R} + i\mathsf{R}$.

2.2. **($\mathbb{R}_{\mathbf{An}}$, $\{f_i\}$)*-admissible open cover of* $\mathbb{U}(\mathbb{C})$.** In addition to the data and notation spelled out in Section 1.1, we assume that:

(i) There is a system of connected open subsets $D_n(\mathbb{C}) \subset \mathbb{U}(\mathbb{C})$, $n \in \mathbb{N}$, definable in $\mathbb{R}_{\mathrm{An}}$ (possibly with parameters), such that

$$\text{for any } n \in \mathbb{N}, \quad D_n \subseteq D_{n+1}, \quad \text{and} \quad \bigcup_n D_n(\mathbb{C}) = \mathbb{U}(\mathbb{C}).$$

(ii) The restriction $f_{i,n}$ of f_i on D_n is definable in $\mathbb{R}_{\mathrm{An}}$ for each $i \in I$ and $n \in \mathbb{N}$, and for each i there is n such that $f_i(D_n) = \mathbb{X}_i$.

(iii) For all $i \in I$, there is a group Γ_i of biholomorphic transformations on $\mathbb{U}(\mathbb{C})$, so that the restrictions of the transformations to the $D_n(\mathbb{C})$ are L_{An}-definable and fibres of f_i are Γ_i-orbits, that is,

$$f_i : \mathbb{U}(\mathbb{C}) \to \mathbb{X}_i(\mathbb{C}) \cong \mathbb{U}(\mathbb{C})/\,\Gamma_i.$$

Moreover, for $i > j$, Γ_i is a finite index subgroup of Γ_j, that is, the cover $\mathrm{pr}_{i,j} : \mathbb{X}_i \to \mathbb{X}_j$ is finite.

(iv) The system of maps f_i, $i \in I$ is $\mathbb{U}$-*complete*: there is a chain $I_0 \subseteq I$ such that

$$\bigcap_{l \in I_0} \Gamma_l = \{1\}.$$

2.3. ***Examples of admissible*** $\mathbb{R}_{\mathbf{An}}$**.** In all our examples $\mathbb{R}_{\mathrm{An}}$ is a L_{An}-reduct of $\mathbb{R}_{\exp,an}$, the real numbers with exponentiation and restricted analytic functions. What varies is $\mathbb{U}$, k_0 and the choice of the family $\{f_i, D_n : i \in I, n \in \mathbb{N}\}$ the members of which assumed to be L_{An}-definable.

(1) $\mathbb{U}(\mathbb{C}) = \mathbb{C}$, $I = \mathbb{N}$, $\mathbb{X}_i = \mathbb{G}_m$ for all $i \in I$, the algebraic torus,

$$D_n = \{z \in \mathbb{C} : -2\pi n < \mathrm{Im}\, z < 2\pi n\},$$

$f_k(z) = \exp\left(\frac{z}{k}\right)$, and $\mathrm{k}_0 = \mathbb{Q}$.

(2) $\mathbb{U}(\mathbb{C}) = \mathbb{C}$, $I = \mathbb{N}$, $\mathbb{X}_i = \mathrm{E}_\tau$ for all $i \in I$, an elliptic curve

$$f_k = \exp_{\tau,k} : \mathbb{C} \to \mathrm{E}_\tau \subset \boldsymbol{P}^2, \quad z \mapsto \exp_\tau\left(\tfrac{z}{k}\right),$$

the covering map for E_τ ($\exp_\tau$ is constructed from the Weierstrass $\mathfrak{P}$-function and its derivative $\mathfrak{P}'$, with period $k\Lambda_\tau = k\mathbb{Z} + \tau k\mathbb{Z}$).

D_1 is the interior of the square in $\mathbb{C}$ with vertices $(0, 1, \tau, \tau+1)$, and $D_n = n \cdot D_1$. Here k_0 is the field of definition of E_τ.

(3) $\mathbb{U}(\mathbb{C}) = \mathbb{H}$, the upper half-plane.

$$D_n = \left\{z \in \mathbb{H} : -\tfrac{1}{2}n \le \mathrm{Re}(z) < \tfrac{1}{2}n \ \&\ \mathrm{Im}(z) > \tfrac{1}{n+1}\right\}.$$

For $n = 1$ this is the interior of the fundamental domain of the j-function

$$F = \left\{z \in \mathbb{H} : -\tfrac{1}{2} \le \mathrm{Re}(z) < \tfrac{1}{2} \ \&\ \mathrm{Im}(z) > \tfrac{1}{2}\right\}$$

and the results of [Peterzil and Starchenko 2013] state that the restriction of j to F is defined in $\mathbb{R}_{\mathrm{exp,an}}$. Note that, for each n, D_n can be covered by finitely many shifts of D_1 by Möbius transformations from $\Gamma := \mathrm{PSL}_2(\mathbb{Z})$. This allows one to define j on D_n in $\mathbb{R}_{\mathrm{exp,an}}$.

Moreover, we can similarly consider more general functions

$$j_N : \mathbb{H} \to \mathbb{Y}(N) \cong \mathbb{H}/\,\Gamma(N)$$

onto level N Shimura curves. A fundamental domain for j_N is a finite union of finitely many shifts of F and the analysis of [Peterzil and Starchenko 2013] shows that the restriction of j_N on its fundamental domain is definable in $\mathbb{R}_{\mathrm{exp,an}}$. Thus we can take the family $\{j_N\}$ to be our $\{f_i\}$ $(i = N)$ and $\mathbb{Y}(N)$ to be the $\mathbb{X}_i$. It is well-known that the $\mathbb{Y}(N)$ and j_N are defined over $\mathrm{k}_0 = \mathbb{Q}^{\mathrm{ab}}$, the extension of $\mathbb{Q}$ by roots of 1.

(4) $\mathbb{U}(\mathbb{C}) = \mathbb{H}$. Let Γ is a Fuchsian subgroup of $\mathrm{PGL}_2(\mathbb{R})$ and $\{\Gamma_i : i \in I\}$ the system of all finite index subgroups of Γ (see [Katok 1992]). Then the $\mathbb{H}/\,\Gamma_i$ are biholomorphic to compact projective curves $\mathbb{X}_i(\mathbb{C})$ with bounded fundamental domains. Thus one can define D_n and f_i as in Section 2.2, with k_0 being the union of the fields of definition of the $\mathbb{X}_i$.

(5) [Peterzil and Starchenko 2013] supplies us with a plethora of other examples, in particular $\mathbb{U}(\mathbb{C}) = \mathbb{H}_g$, the Siegel half-space, and $\mathbb{X}_i$ moduli spaces of polarised algebraic varieties.

3. The K-analytic setting

3.1. ***Abstract structures definable in*** **R.** Now we extend notations of Section 2.2 and, assuming $\mathsf{R} \in \mathrm{Mod}_{\mathrm{An}}$ be given, let $\mathbb{U}$, $\mathbb{X}_i$, $(i \in I)$, D_n, Γ_i and f_i be defined as in Section 2.2 in the language L_{An}. In particular, we read $\mathbb{U} := \mathbb{U}(\mathrm{K})$, $\mathbb{X}_i := \mathbb{X}_i(\mathrm{K})$, for $\mathrm{K} = \mathrm{K}(\mathsf{R})$, when the choice of the model R does not matter.

More precisely, we define

$$\mathbb{U}(\mathrm{K}) = \bigcup_n D_n(\mathrm{K}),$$

which is an $L_{\omega_1,\omega}$ interpretation of $\mathbb{U}$ in R for each $i \in I$. Now $f_i : \mathbb{U}(\mathrm{K}) \to \mathbb{X}_i(\mathrm{K})$ is defined to be the map such that it coincides with the map $f_{i,n} : D_n(\mathrm{K}) \to \mathbb{X}_i(\mathrm{K})$ for each $n \in \mathbb{N}$. Note that the latter is K-holomorphic in the sense of [Peterzil and Starchenko 2008]. We will often say K-holomorphic (analytic) in an extended sense: the restriction $f_{i,n}$ of f_i to $D_n(\mathrm{K})$ is K-holomorphic.

We write $D_{\bar{n}} \subset \mathbb{U}^m$ meaning that $\bar{n} = \langle n_1, \dots, n_m\rangle \in \mathbb{N}^m$ and

$$D_{\bar{n}} = D_{n_1} \times \cdots \times D_{n_m}.$$

Define f_i on $D_{\bar{n}}$ as $\langle u_1, \dots, u_m\rangle \mapsto \langle f_i(u_1), \dots, f_i(u_m)\rangle$. This obviously extends to the map f_i with the domain $\mathbb{U}^m$.

We will often restrict our analysis of K-analytic sets to open neighbourhoods, where *open* always means *definable open.*

Let k_0 be a subfield of K such that $\mathrm{k}_0 \subseteq \mathrm{dcl}(\varnothing)$, that is any point of k_0 is definable in R without parameters. Note that k_0 contains any point of the form $f_i(a)$ for $i \in I$ and a definable point $a \in D_n$.

More generally, we will work with an arbitrary k such that $\mathrm{k}_0 \subseteq \mathrm{k} \subset \mathrm{K}$.

Definition 3.2. Given $S \subset \mathbb{U}^m$ we say that S is $L_{\mathrm{glob}}(\mathrm{k})$*-primitive* if there are $I_S \subseteq I$ and Zariski closed $Z_i \subseteq \mathbb{X}_i^m$, $i \in I_S$, defined over k, such that

$$S = \bigcap_{i \in I_S} f_i^{-1}(Z_i).$$

Remark 3.3. In Definition 3.2 we may assume without loss of generality that I_S is a chain and, for $i' \geq i$ in I_S,

$$\mathrm{pr}_{i'i}(Z_{i'}) = Z_i. \tag{1}$$

Proof. First, we may assume that $I_S = I$ by setting for $i \in I \setminus I_S$, $Z_i := \mathbb{X}_i^m$.

For a finite $J \subseteq I$, take a $i_J \in I$ such that $i_J \geq J$. Set, for each $k \in J$,

$$Z_{i_J,k} := \mathrm{pr}_{i_J,k}^{-1}(Z_k) \subseteq \mathbb{X}_{i_J}^m \quad \text{and} \quad Z_{i_J}^* = \bigcap_{k \in J} Z_{i_J,k}.$$

Then, since $f_k = \mathrm{pr}_{i_J,k} \circ f_{i_J}$,

$$f_{i_J}^{-1}(Z_{i_J,k}) = f_k^{-1}(Z_i) \quad \text{and} \quad \bigcap_{k \in J} f_k^{-1}(Z_i) = f_{i_J}^{-1}(Z_{i_J}^*). \tag{2}$$

Since I is a countable lattice we can represent

$$I = \bigcup_{n \in \mathbb{N}} J_n$$

where $J_n \subseteq I_S$ are finite and $J_{n+1} \supseteq J_n$ for each n.

Consider (2) with $J = J_n$ and write i_{J_n} as i_n. Clearly, $i_{n+1} \geq i_n$ and

$$S = \bigcap_{n \in \mathbb{N}} f_{i_n}^{-1}(Z_{i_n}^*). \tag{3}$$

Finally, note that in (3) $\mathrm{pr}_{i_n,i_l}(Z_{i_n}^*) \subseteq Z_{i_l}^*$ for $n \geq l$, and $\mathrm{pr}_{i_n,i_l}(Z_{i_n}^*)$ is a Zariski closed subset of $\mathbb{X}_{i_l}^m$ since pr_{i_n,i_l} is unramified (and étale). Hence, we may replace $Z_{i_l}^*$ by $\bigcap_{n \geq l} \mathrm{pr}_{i_n,i_l}(Z_{i_n}^*)$ while keeping (3). Doing this consecutively for $l = 1, 2, \ldots$ delivers us (1). $\square$

Remark. The equality relation is $L_{\mathrm{glob}}(\mathrm{k}_0)$-primitive.

3.4. K-*holomorphic maps and* K-*analytic subsets.* We refer to [Peterzil and Starchenko 2008] for definitions and basic facts on K-analyticity in open definable subsets $D_{\bar{n}}$. By slight abuse of the terminology we call a subset $S \subseteq \mathbb{U}^m$ K-analytic if $S \cap D_{\bar{n}}$ is K-analytic for each $D_{\bar{n}} \subset \mathbb{U}^m$.

Since the complex covering maps f_i are holomorphic, the maps $f_{i,n} : D_n(\mathrm{K}) \to \mathbb{X}_i(\mathrm{K})$ are K-holomorphic and locally K-biholomorphic. It follows the sets $f_i^{-1}(Z_i)$ in Definition 3.2 are K-analytic and are locally K-biholomorphically isomorphic to the Z_i.

The dimension dim is always the K-dimension of a K-analytic set. When Z is an algebraic variety, the dimension of the respective K-analytic set is $\dim Z := \dim Z(\mathrm{K})$, and this coincides with the dimension in the sense of algebraic geometry.

Lemma 3.5. *Given an $L_{\mathrm{glob}}(\mathrm{k})$-primitive S, $S \cap D_{\bar{n}}$ is* K*-analytic in $D_{\bar{n}}$. S is* K*-analytic in $\mathbb{U}^m$.*

Proof. Let S be as in Definition 3.2 with the assumption (1) and let $S_i := f_i^{-1}(Z_i)$. It follows by definition that the $S_i \cap D_{\bar{n}}$ are K-analytic. We need to prove that $\bigcap_{i \in I_S} S_i \cap D_{\bar{n}}$ is analytic.

Let $s \in S \cap D_{\bar{n}}$. For each $i \in I_S$ there is an open neighbourhood $O_{s,i}$ of s such that $S_i \cap O_{s,i}$ is irreducible. We may assume that $S_{i'} \cap O_{s,i'} \subseteq S_i \cap O_{s,i}$ for $i' \geq i$. Then there exists $i_0 \in I_S$ such that for $i' \geq i \geq i_0$, $\dim S_{i'} \cap O_{s,i'} = \dim S_i \cap O_{s,i}$.

Since $S_i \cap O_{s,i}$ is irreducible, $S_{i'} \cap O_{s,i} = S_i \cap O_{s,i}$ for all $i' \geq i \geq i_0$. Thus $S \cap O_{s,i} = S_i \cap O_{s,i}$, which proves that S is K-analytic in the neighbourhood, and hence in $D_{\bar{n}}$. $\square$

Remark 3.6. S^{sing}, the set of singular points of $L_{\mathrm{glob}}(\mathrm{k})$-primitive S, is also an $L_{\mathrm{glob}}(\mathrm{k})$-primitive since

$$S^{\mathrm{sing}} = \bigcap_{i \in I_S} f_i^{-1}(Z_i^{\mathrm{sing}}).$$

Proposition 3.7. *Let $S \subseteq \mathbb{U}^m$ be $L_{\mathrm{glob}}(\mathrm{k})$-primitive and let, for some n, $S_{j,\bar{n}} \subseteq S \cap D_{\bar{n}}$ be a* K*-analytic irreducible component of $S \cap D_{\bar{n}}$. Then*:

(i) *For any $D_{\bar{n}'} \supseteq D_{\bar{n}}$ there is unique $S_{j,\bar{n}'} \supseteq S_{j,\bar{n}}$ a* K*-analytic irreducible component of $S \cap D_{\bar{n}'}$. The set*

$$S_j := \bigcup_{D_{\bar{n}'} \supseteq D_{\bar{n}}} S_{j,\bar{n}'}$$

is well-defined. (*Call it an irreducible component of S.*)

(ii) *The number of* K*-analytic components S_j of S is at most countable.*

(iii) *The irreducible components S_j are $L_{\text{glob}}(\mathrm{k}')$-primitive for some algebraic extension* k' *of* k.

(iv) *For any i, $f_i(S_j)$ is a Zariski closed* k'*-definable geometrically irreducible subset of $\mathbb{X}_i^m$.*

Proof. By [Peterzil and Starchenko 2008, 4.12], $S_{j,\bar{n}'}$ is irreducible if and only if $S_{j,\bar{n}'} \setminus S_{j,\bar{n}'}^{\text{sing}}$ is definably connected. The union of any two irreducible extensions of $S_{j,\bar{n}} \setminus S_{j,\bar{n}}^{\text{sing}}$ will be connected, since any two points in the union can be connected by a definable path passing through $S_{j,\bar{n}} \setminus S_{j,\bar{n}}^{\text{sing}}$. Hence the extensions coincide, which gives us the first statement of proposition.

The number of such irreducible components is at most countable since the number of components in each $D_{\bar{n}'}$ is finite. This proves (i) and (ii).

Define $\dim S_j$ to be $\dim S_{j,\bar{n}}$, which does not depend on $D_{\bar{n}}$ as long as $S_j \cap D_{\bar{n}} \neq \varnothing$, since irreducible sets are of pure dimension (the proof is the same as in the complex case, see also [Peterzil and Starchenko 2008]). Define

$$\dim S := \max_j \dim S_j. \tag{4}$$

We may assume that

$$S = \bigcap_{i \in I_0} f_i^{-1}(Z_i)$$

for some chain $I_0 \subseteq I$, some Zariski closed $Z_i \subseteq \mathbb{X}_i^m$ such that $\dim Z_i = \dim S$ and $\mathrm{pr}_{i,l}(Z_i) = Z_l$ for $i > l$ in I_0.

Let $S^i := f_i^{-1}(Z_i)$ and let $S^i = \bigcup_{j \in J_i} S_j^i$ be the decomposition into irreducible analytic components with maximum dimension equal to $\dim S$. It follows that the components of S^i are also components of S^l for $i > l$, and thus S_j is a component of $f_l^{-1}(Z_l)$.

Fix l for the time being. We can represent $Z_l = \bigcup_{p \in P} Z_{l,p}$, a finite union of geometrically irreducible algebraic subvarieties $Z_{l,p}$ defined over some algebraic extension k' of k. Also, S can be represented as a finite union of $L_{\text{glob}}(\mathrm{k}')$-primitives,

$$S = \bigcup_{p \in P} T_{l,p}, \qquad \text{where } T_{l,p} = S \cap f_l^{-1}(Z_{l,p})$$

and the irreducible component S_j of S is an irreducible component of one of $T_{l,p}$.

We assume without loss of generality that Z_l is geometrically irreducible, P is a singleton and, since we are only interested in S_j, assume

$$S = f_l^{-1}(Z_l).$$

We omit the subscript l in the claim below.

Claim. $f(S_j) = Z$ and for any other component S_k of S there is $\gamma \in \Gamma$ such that $\gamma \cdot S_j = S_k$.

Proof. By Section 1.1 there is $\bar{n}$ such that $f(D_{\bar{n}}) = \mathbb{X}^m$.

By our assumption then

$$Z = f\big(\bigcup_{k\in J} S_k\big) = \bigcup_{k\in J} f(S_k \cap D_{\bar{n}}) = \bigcup_{k\in J_0} f(S_k \cap D_{\bar{n}})$$

where J lists all the components of S and J_0 lists the components S_k such that $S_k \cap D_{\bar{n}} \neq \varnothing$, so J_0 is finite.

Hence for the finite $J_1,\ J_0 \subseteq J_1 \subseteq J$, we have

$$Z = \bigcup_{k\in J_1} f(S_k).$$

Let Z^{sing} the singular points of Z and S^{sing} the singular points of S, which by the fact that f is a local biholomorphisms are related as

$$f^{-1}(Z^{\text{sing}}) = S^{\text{sing}}. \tag{5}$$

Note that if $s \in S_j \cap S_k$, a common point of two distinct components of S then $s \in S^{\text{sing}}$. That is $S \setminus S^{\text{sing}}$, the analytic subset of the open set $\mathbb{U}^m \setminus S^{\text{sing}}$, splits into nonintersecting analytic components $S_k \setminus S^{\text{sing}}$. We get from (5)

$$Z \setminus Z^{\text{sing}} = \bigcup_{k\in J_1} f(S_k \setminus S^{\text{sing}}). \tag{6}$$

The union on the right cannot be disjoint, that is, either J_1 is a singleton, or there are distinct $k_0, k_1 \in J_1$ such that $f(S_{k_0} \setminus S^{\text{sing}}) \cap f(S_{k_1} \setminus S^{\text{sing}}) \neq \varnothing$. Indeed, suppose for a contradiction that it is disjoint. Note that for a respective $D_{\bar{n}}$, $f : D_{\bar{n}} \twoheadrightarrow \mathbb{X}^m$ is a (definably) closed covering map since it is locally biholomorphisms. Hence $f(D_{\bar{n}} \cap S_k \setminus S^{\text{sing}})$, $k \in J_1$, are disjoint definably closed subsets the union of which is the definably connected algebraic set $Z \setminus Z^{\text{sing}}$, which is a contradiction.

Now we claim that

$$f(S_{k_0} \setminus S^{\text{sing}}) = Z \setminus Z^{\text{sing}} \quad \text{ for a } k_0 \in J_1. \tag{7}$$

Indeed, otherwise there are $k_0, k_1 \in J_1$ such that $f(S_{k_0} \setminus S^{\text{sing}}) \neq f(S_{k_1} \setminus S^{\text{sing}})$ but $f(S_{k_0} \setminus S^{\text{sing}}) \cap f(S_{k_1} \setminus S^{\text{sing}}) \neq \varnothing$. The latter means that there are $s_0 \in S_{k_0} \setminus S^{\text{sing}}$ and $s_1 \in S_{k_1} \setminus S^{\text{sing}}$ such that $f(s_1) = f(s_0)$, and hence $s_1 = \gamma \cdot s_0$ for some $\gamma \in \Gamma$.

It follows that the K-analytic sets S_{k_1} and $\gamma \cdot S_{k_0}$ intersect in a nonsingular point of $S \cap D_{\bar{n}}$ and thus $S_{k_1} \cap D_{\bar{n}} = \gamma \cdot S_{k_0} \cap D_{\bar{n}}$, and so

$$S_{k_1} = \gamma \cdot S_{k_0} \quad \text{and} \quad f(S_{k_1}) = f(S_{k_0}).$$

(7) follows. This finishes the proof of the claim and of the statement (iv). □

Now, for any $i \in I$ consider

$$Z_{ij} := f_i(S_j)$$

which we proved to be Zariski closed irreducible and

$$f_i^{-1}(Z_{ij}) = \bigcup_{\gamma \in \Gamma_i} \gamma \cdot S_j.$$

Since by assumption $\bigcap_{l \in I} \Gamma_l$ is trivial, for some chain $I_1 \subseteq I$ extending I_0 we have

$$S_j = \bigcap_{l \in I_1} f_l^{-1}(Z_{lj}),$$

proving (iii). □

Definitions 3.8. For an m-tuple u in $\mathbb{U}$ and a subfield $\mathrm{k} \subset \mathrm{K}$ the *locus* of u over k, written $\mathrm{loc}(u/\mathrm{k})$, is the minimum $L_{\mathrm{glob}}(\mathrm{k})$-primitive containing u.

We say an $L_{\mathrm{glob}}(\mathrm{k})$-primitive S is k-*irreducible* if S cannot be represented as $S_1 \cup S_2$ with $L_{\mathrm{glob}}(\mathrm{k})$-primitives S_1 and S_2, both distinct from S.

Remark. Note that $\mathrm{loc}(u/\mathrm{k})$ is k-irreducible.

4. L_{glob}-structures

4.1. Recall, see [Pillay and Steinhorn 1986], that an o-minimal structure R is a pregeometry, i.e., has a well-behaved dependence relation, and one can define a notion of a (combinatorial) dimension $\mathrm{cdim}\, A$ of a subset $A \subseteq \mathsf{R}$ (not to be confused with K-dimension) as the cardinality of a maximal independent subset of A.

In particular, $\mathrm{cdim}\, \mathsf{R}_0 = 0$ for the prime model R_0 of the theory $\mathrm{Th}(\mathbb{R}_{\mathrm{An}})$. And, if $\mathrm{card}\, \mathsf{R} = \kappa > \aleph_0$, then $\mathrm{cdim}\, \mathsf{R} = \kappa$.

This has the following relationship with $\dim_{\mathsf{R}} S$ (the "real" dimension in the sense of [Peterzil and Starchenko 2008]) for an R-manifold $S \subseteq \mathsf{R}^m$ defined over a set C: assuming $\mathrm{cdim}\, \mathsf{R}/C \geq m$, for any $d \in \mathbb{N}$,

$$\dim_{\mathsf{R}} S \geq d \text{ if and only if there exists } \langle s_1, \ldots, s_m \rangle \in S \text{ such that } \mathrm{cdim}(\{s_1, \ldots, s_m\}/C) \geq d. \tag{8}$$

Recall that if S is K-analytic, then

$$\dim S = \tfrac{1}{2} \dim_{\mathsf{R}} S. \tag{9}$$

Definition 4.2. Given $\mathsf{R} \in \mathrm{Mod}_{\mathrm{An}}$, define $\mathfrak{U}(\mathsf{R})$ to be the structure with universe $\mathbb{U}(\mathrm{K})$ (K the field $\mathsf{R} + i\mathsf{R}$) in the language of $L_{\mathrm{glob}}(\mathrm{k}_0)$-primitives.

Define $\mathfrak{U}$ to be the class of all structures of the form $\mathfrak{U}(\mathsf{R})$.

Fact 4.3. *For* K *an algebraically closed field, consider the structure* $\mathbb{X}(\mathrm{K})_{\mathrm{Zar},\mathrm{k}_0}$ *on an infinite algebraic variety* $\mathbb{X}(\mathrm{K})$ *over* k_0 *equipped with relations* $Z \subseteq \mathbb{X}^m$, *all Zariski closed* Z *over* k_0.

The field structure K *together with its* k_0*-points is* $\varnothing$*-interpretable in* $\mathbb{X}(\mathrm{K})_{\mathrm{Zar},\mathrm{k}_0}$.

This is well-known. A detailed proof is given in [Bays 2009, Appendix A].

Proposition 4.4. $\mathfrak{U}(\mathsf{R})$ *interprets in the first order way over* $\varnothing$ *the field* K, *points of the subfield* k_0 *and all the maps* $f_i : \mathbb{U} \to \mathbb{X}_i(\mathrm{K})$.

Proof. First note that the equivalence relations on $\mathbb{U}$,

$$E_i(u_1, u_2) :\equiv f_i(u_1) = f_i(u_2),$$

are $L_{\mathrm{glob}}(\mathrm{k})$-primitives. Thus the sets $\mathbb{X}_i(\mathrm{K})$ are $\varnothing$-interpretable as $\mathbb{U}/E_i$ together with the maps $f_i : \mathbb{U} \to \mathbb{U}/E_i$.

Given a Zariski closed $Z_i \subset \mathbb{X}_i^m$ we have $Z_i^{\mathbb{U}} := f_i^{-1}(Z_i)$, a definable subset of $\mathbb{U}^m$. Thus $Z_i = f_i(Z_i^{\mathbb{U}})$ are $\varnothing$-interpretable.

Now the structure $\mathbb{X}_0(\mathrm{K})_{\mathrm{Zar},\mathrm{k}_0}$ equipped with relations $Z \subseteq \mathbb{X}_0^m$, for all Zariski closed Z over k_0, is $\varnothing$-interpretable.

It follows from Fact 4.3 that one can interpret K and k_0-points in $\mathfrak{U}(\mathsf{R})$. $\square$

Corollary 4.5. *Any* $L_{\mathrm{glob}}(\mathrm{K})$*-primitive is type-definable in* $\mathfrak{U}(\mathsf{R})$ *using parameters.*

Below $\mathbb{U}$ is always the universe $\mathbb{U}(\mathrm{K})$ for some $\mathfrak{U}(\mathsf{R})$ in $\mathfrak{U}$.

Lemma 4.6. *If* k *is algebraically closed then* $\mathrm{loc}(u/\mathrm{k})$, *the locus of* u *over* k, *is* K*-analytically irreducible.*

If $S \subseteq \mathbb{U}^m$ *is an* $L_{\mathrm{glob}}(\mathrm{k})$*-primitive and* K*-analytically irreducible, then* $S = \mathrm{loc}(u/\mathrm{k})$, *for some* $u \in S$.

Proof. The first statement is just a corollary to Proposition 3.7(iv).

Let $\dim S = d$. By (8) and (9) there is a $u \in S$ such that $u = \langle s_1, \ldots, s_m \rangle$ with $\mathrm{cdim}(s_1, \ldots, s_m/\mathrm{k}) = 2d$. Then $\mathrm{loc}(u/\mathrm{k}) \subseteq S$ and, again by (8) and (9), $\dim \mathrm{loc}(u/\mathrm{k}) \geq d$. Since S is K-analytically irreducible, $\mathrm{loc}(u/\mathrm{k}) = S$. $\square$

Lemma 4.7. *Let* $S \subset \mathbb{U}^m$ *be an* $L_{\mathrm{glob}}(\mathrm{k})$*-primitive*, $\dim S = d$. *Assume also* $\mathrm{cdim}(\mathsf{R}/\mathrm{k}) \geq \aleph_0$. *Then, for any family* $L_{j \in J}$ *of* $L_{\mathrm{glob}}(\mathrm{k})$*-primitives such that* $\dim L_j < d$, *for all* $j \in J$,

$$S \setminus \bigcup_{j \in J} L_j \neq \varnothing. \tag{10}$$

Proof. S contains a point $u = \langle s_1, \ldots, s_m \rangle$ with $\mathrm{cdim}(s_1, \ldots, s_m/\mathrm{k}) = 2d$, which is not a point of any L_j. $\square$

Proposition 4.8 (the projection of an irreducible analytic set). *Let* k *be algebraically closed*, $\mathrm{cdim}(\mathrm{R}/\mathrm{k}) \geq \aleph_0$. *Let $T \subseteq \mathbb{U}^{m+1}$ be an $L_{\mathrm{glob}}(\mathrm{k})$-primitive* K*-analytically irreducible, and let $\boldsymbol{p} : \mathbb{U}^{m+1} \to \mathbb{U}^m$ be the projection onto the first m coordinates. Then there are an $L_{\mathrm{glob}}(\mathrm{k})$-primitive $S \subseteq \mathbb{U}^m$, an $i_0 \in I$ and a Zariski closed subset $R \subseteq \mathbb{X}^m_{i_0}$ defined over* k *such that* $\dim R < \dim S$ *and*

$$S \setminus f_{i_0}^{-1}(R) \subseteq \boldsymbol{p}(T) \subseteq S. \tag{11}$$

Moreover, for any $d \leq \dim T - \dim S$, there is a Zariski closed $R_d \subset \mathbb{X}^m_{i_0}$ defined over k *such that $R \subseteq R_d$*, $\dim R_d < \dim S$ *and*

$$\boldsymbol{p}(T) \setminus f_{i_0}^{-1}(R_d) = \boldsymbol{p}_d(T), \tag{12}$$

where

$$\boldsymbol{p}_d(T) := \{s \in \boldsymbol{p}(T) : \dim(\boldsymbol{p}^{-1}(s) \cap T) \leq d\}.$$

Proof. By Lemma 4.6,

$$T = \mathrm{loc}(\bar{u}v/\mathrm{k})$$

for some $\bar{u}v \in \mathbb{U}^{m+1}$, ($\bar{u} \in \mathbb{U}^m$, $v \in \mathbb{U}$).

Let

$$S = \mathrm{loc}(\bar{u}/\mathrm{k}).$$

By definition

$$S = \bigcap_{i \in I_0} f_i^{-1}(Z_i), \qquad T = \bigcap_{i \in I_0} f_i^{-1}(W_i)$$

for some Zariski closed $Z_i \subseteq \mathbb{X}^m_i$, $W_i \subseteq \mathbb{X}^{m+1}_i$ over k and we apply the same notation to the projection map $\boldsymbol{p} : \mathbb{X}^{m+1}_i \to \mathbb{X}^m_i$. By Proposition 3.7(iv) we may assume that all the Z_i and W_i are irreducible and of dimension equal to that of S and T respectively,

$$f_i(S) = Z_i \quad \text{and} \quad f_i(T) = W_i \qquad \text{for all } i \in I_0,$$

and $f_i(\bar{u})$ is a generic point of Z_i, $f_i(\bar{u})^\frown f_i(v)$ a generic point of W_i.

By basic algebraic geometry, $\boldsymbol{p}(W_i)$ is a constructible irreducible set and $f_i(\bar{u})$ its generic point, and thus the Zariski closure of $\boldsymbol{p}(W_i)$ is equal to Z_i. That is, there are Zariski closed $R_i \subset Z_i$ over k such that

$$Z_i = \boldsymbol{p}(W_i) \cup R_i \quad \text{and} \quad \dim R_i < \dim Z_i. \tag{13}$$

Since

$$\boldsymbol{p}\Big(\bigcap_{i \in I} f_i^{-1}(W_i)\Big) \subseteq \bigcap_{i \in I_0} \boldsymbol{p}(f_i^{-1}(W_i)) = \bigcap_{i \in I_0} f_i^{-1}(\boldsymbol{p}(W_i)),$$

we have

$$\boldsymbol{p}(T) \subseteq S.$$

Let i_0 be an element of I_0 and, for simplicity of notation, $f := f_{i_0}$, so $f(T) = W$, $f(S) = Z$ and $Z = \boldsymbol{p}(W) \cup R$ as in (13).

By the basic assumptions, given arbitrary $t \in T$, $s = \boldsymbol{p}(t)$, for some R-definable open neighbourhood $U \subset \mathbb{U}^m$ of s and open neighbourhood $U \times V \subset \mathbb{U}^{m+1}$ of t, with $V \subset \mathbb{U}$, the restriction $f_U : U \to \mathbb{X}^m$ and $f_{U\times V} : U \times V \to \mathbb{X}^{m+1}$ are injective.

Thus we obtain the commutative diagram

$$\begin{array}{ccc} T \cap (U \times V) & \xrightarrow{f_{U\times V}} & W \\ \Big\downarrow \boldsymbol{p} & & \Big\downarrow \boldsymbol{p} \\ S \cap U & \xrightarrow{\quad f_U \quad} & \boldsymbol{p}(W) \supseteq Z \setminus R \end{array} \tag{14}$$

By comparing images of the downward-pointing arrows we conclude

$$S \cap U \supseteq \boldsymbol{p}(T \cap (U \times V)) \supseteq f_U^{-1}(Z \setminus R).$$

Note that

$$f_U^{-1}(Z \setminus R) = S \cap U \setminus f^{-1}(R),$$

and the choice of R is independent on the choice of U. Hence $\boldsymbol{p}(T) \supseteq S \setminus f^{-1}(R)$ and (11) is proved.

To prove the second statement recall another basic fact of algebraic geometry: there is a Zariski closed $R_d \subset \mathbb{X}^m$ such that

$$\boldsymbol{p}(W) \setminus R_d = \boldsymbol{p}_d(W) := \{z \in \boldsymbol{p}(W) : \dim \boldsymbol{p}^{-1}(z) \cap W \leq d\}.$$

Now repeat the argument with the diagram (14) with $\boldsymbol{p}_d(W)$ in place of $\boldsymbol{p}(W)$. This proves (12). $\square$

Recall the notion of an *analytic Zariski structure*, see [Zilber 2010; 2017].

Corollary 4.9. *Assuming that* k *is algebraically closed and* $\mathrm{cdim}(\mathsf{R}/\mathrm{k}) \geq \aleph_0$, *the structure* $\mathfrak{U}(\mathsf{R})$ *in the language* $L_{\mathrm{glob}}(\mathrm{k})$ *is an analytic Zariski structure.*

Proof. The statement of Proposition 4.8 asserts that the structure on $\mathbb{U}$ determined by $L_{\mathrm{glob}}(\mathrm{k})$-primitives satisfies the key axioms (WP) and (FC) of the definition of an analytic Zariski structure. The rest of the axioms follow easily from definitions and basic algebraic geometry. $\square$

The next statements and their proofs are similar to one of the main statements of [Zilber 2017] for analytic Zariski structures. More early work of M. Gavrilovich also proves this for complex analytic Zariski structures.

Proposition 4.10. $\mathfrak{U}$ *is* $\aleph_0$*-homogeneous over algebraically closed subfields*:

Suppose $\mathfrak{U}(\mathsf{R}_1), \mathfrak{U}(\mathsf{R}_2) \in \mathfrak{U}$, $\mathsf{R}_0, \mathsf{R}_1, \mathsf{R}_2 \in \mathrm{Mod}_{\mathrm{An}}$, $\mathsf{R}_0 \subseteq \mathsf{R}_1$, $\mathsf{R}_0 \subseteq \mathsf{R}_1$.

Let $\mathrm{k} \subseteq \mathrm{K}_0 = \mathrm{K}(\mathsf{R}_0)$ *be an algebraically closed subfield such that* $\mathrm{cdim}(\mathsf{R}_1/\mathrm{k}) \geq \aleph_0$ *and* $\mathrm{cdim}(\mathsf{R}_2/\mathrm{k}) \geq \aleph_0$.

Then for any $\bar{u}_1 \in \mathbb{U}^m(\mathrm{K}_1)$, $\bar{u}_2 \in \mathbb{U}^m(\mathrm{K}_2)$, *and* $w_1 \in \mathbb{U}(\mathrm{K}_1)$ *such that*

$$\mathrm{loc}(\bar{u}_1/\mathrm{k}) = \mathrm{loc}(\bar{u}_2/\mathrm{k})$$

there is $w_2 \in \mathbb{U}(\mathrm{K}_2)$ *such that*

$$\mathrm{loc}(\bar{u}_1 w_1/\mathrm{k}) = \mathrm{loc}(\bar{u}_2 w_2/\mathrm{k}).$$

Proof. Let $S = \mathrm{loc}(\bar{u}_1/\mathrm{k})$ and $T = \mathrm{loc}(\bar{u}_1 w_1/\mathrm{k})$. Note that $\bar{u}_1$ and $\bar{u}_2$ are nonsingular points of S and $\bar{u}_1 w_1$ a nonsingular point of T, by Remark 3.6.

Let $d := \dim \boldsymbol{p}^{-1}(\bar{u}_1) \cap T$ be the dimension of the fibre over $\bar{u}_1$, and the subset $\boldsymbol{p}_d(T)$ be as defined in Proposition 4.8. Note that by the dimension theorem of algebraic geometry $\dim \boldsymbol{p}_d(T) = \dim S$, since $\dim \boldsymbol{p}_d(W) = \dim S$ (in the notation of Proposition 4.8). Note also that

$$\dim T = \dim S + d$$

since respective equality holds for the dimensions of W and Z.

It follows that $\boldsymbol{p}_d(T)$ contains all generic over k points of S, $\bar{u}_2 \in \boldsymbol{p}_d(T)$ and thus

$$\dim \boldsymbol{p}^{-1}(\bar{u}_2) \cap T = d.$$

Thus there exists w_2 such that $\bar{u}_2 w_2 \in \boldsymbol{p}^{-1}(\bar{u}_2) \cap T$ and $\dim(w_2/\bar{u}_2\mathrm{k}) = d$. Since T is k-irreducible,

$$T = \mathrm{loc}(\bar{u}_2 w_2/\mathrm{k}). \qquad \square$$

Lemma 4.11. *Let* $S \subseteq \mathbb{U}^{m+n}$ *be an* $L_{\mathrm{glob}}(\mathrm{k})$*-primitive and* $\bar{u} \in \mathbb{U}^m$*. Let*

$$S_{\bar{u}} = \{\bar{v} \in \mathbb{U}^n : \bar{u}\bar{v} \in S\}.$$

Then $S_{\bar{u}}$ *is an* $L_{\mathrm{glob}}(\mathrm{k}')$*-primitive, for* k'*, extension of* k *by coordinates of* $f_i(\bar{u})$*,* $i \in I$*.*

Proof. By definition $S = \bigcap_{i \in I} f_i^{-1}(Z_i)$ for $Z_i \subseteq \mathbb{X}_i^{m+n}$.

Let, for $z_i \in \mathbb{X}_i^m(\mathrm{K})$,

$$Z_{i,z_i} = \{x_i \in \mathbb{X}_i^n(\mathrm{K}) : z_i x_i \in Z_i\}.$$

Thus

$$S_{\bar{u}} = \Big\{\bar{v} \in \mathbb{U}^n : \bigwedge_{i \in I} f_i(\bar{u}) f_i(\bar{v}) \in Z_i\Big\} = \bigcap_{i \in I} f_i^{-1}(Z_{i, f_i(\bar{u})}). \qquad \square$$

Corollary 4.12. *Assuming* k_0 *is algebraically closed,* $\mathfrak{U}$ *is* $\aleph_0$*-homogenous over* $\varnothing$ *and over small submodels: Using the notation of Proposition 4.10, let* $V = \varnothing$ *or* $V = \mathbb{U}(\mathrm{K}_0)$ *and assume* $\mathrm{cdim}(\mathsf{R}_i/\mathrm{K}_0) \geq \aleph_0$ *for* $i = 1, 2$*.*

Then, for any $\bar{u}_1 \in \mathbb{U}^m(\mathrm{K}_1)$, $\bar{u}_2 \in \mathbb{U}^m(\mathrm{K}_2)$, $w_1 \in \mathbb{U}^m(\mathrm{K}_1)$ *such that*

$$\mathrm{tp}(\bar{u}_1/V) = \mathrm{tp}(\bar{u}_2/V)$$

there is $w_2 \in \mathbb{U}^m(\mathrm{K}_2)$ *such that*

$$\mathrm{tp}(\bar{u}_1 w_1/V) = \mathrm{tp}(\bar{u}_2 w_2/V),$$

where tp *is the quantifier-free type of the form* (10).

Proof. For the language without parameters use Proposition 4.10 with $\mathrm{k} = \mathrm{k}_0$. Over the submodel use the statement of Proposition 4.10 with $\mathrm{k} = \mathrm{K}_0$. □

Lemma 4.13. *The structure* $\mathfrak{U}(\mathsf{R}_0)$, *for* R_0 *the prime model of the o-minimal theory* $\mathrm{Th}(\mathbb{R}_{\mathrm{An}})$, *is a prime model of* $\mathfrak{U}$, *that is, there is an* L_{glob}*-embedding* $\mathfrak{U}(\mathsf{R}_0) \subseteq \mathfrak{U}(\mathsf{R})$ *for any* $\mathsf{R} \in \mathrm{Mod}_{\mathrm{An}}$.

Proof. An embedding $\mathsf{R}_0 \preccurlyeq \mathsf{R}$ induces an embedding $\mathfrak{U}(\mathsf{R}_0) \subseteq \mathfrak{U}(\mathsf{R})$. □

Theorem 4.14. *Suppose* k_0 *is algebraically closed.*

Let $\mathsf{R}_1, \mathsf{R}_2 \in \mathrm{Mod}_{\mathrm{An}}$ *and* $\aleph_0 \leq \operatorname{cdim} \mathsf{R}_1 = \operatorname{cdim} \mathsf{R}_2 \leq \aleph_1$. *Then*

$$\mathfrak{U}(\mathsf{R}_1) \cong \mathfrak{U}(\mathsf{R}_2).$$

In particular, $\mathfrak{U}$ *is categorical in cardinality* $\aleph_1$.

Proof. First consider the case when $\operatorname{cdim} \mathsf{R}_1 = \operatorname{cdim} \mathsf{R}_2 = \aleph_0$. Then $\mathfrak{U}(\mathsf{R}_1)$ and $\mathfrak{U}(\mathsf{R}_2)$ are countable and so we can construct an isomorphism via a countable back-and-forth process using Corollary 4.12, where $\mathrm{K}_0 = \mathrm{K}(\mathsf{R}_0)$, R_0 is the prime model of $\mathrm{Th}(\mathbb{R}_{\mathrm{An}})$.

In the case when $\operatorname{cdim} \mathsf{R}_1 = \operatorname{cdim} \mathsf{R}_2 = \aleph_1$, we represent by

$$\mathsf{R}_1 = \bigcup_{\alpha<\aleph_1} \mathsf{R}_{1,\alpha} \quad \text{and} \quad \mathsf{R}_2 = \bigcup_{\alpha<\aleph_1} \mathsf{R}_{2,\alpha}$$

the ascending chains of elementary extensions, $\operatorname{cdim}(\mathsf{R}_{i,\alpha+1}/\mathsf{R}_{i,\alpha}) = \aleph_0$ for $i = 1, 2$, and $\mathsf{R}_{1,0} = \mathsf{R}_{2,0}$ are prime models. Then the required isomorphism is constructed by induction on α: Assume that $\mathsf{R}_{1,\alpha} \cong \mathsf{R}_{2,\alpha}$, and even that both are equal to a R_α. Now apply Corollary 4.12 with $\mathrm{K}_0 = \mathrm{K}(\mathsf{R}_\alpha)$, $\mathrm{K}_1 = \mathrm{K}(\mathsf{R}_{1,\alpha+1})$, and $\mathrm{K}_2 = \mathrm{K}(\mathsf{R}_{2,\alpha+1})$. This again produces an isomorphism $\mathsf{R}_{1,\alpha+1} \cong \mathsf{R}_{2,\alpha+1}$ by the back-and-forth procedure.

For limit indices the extension of isomorphism is obvious. □

5. The one-dimensional case

5.1. Let $P(\mathbb{U})$ stand for the power-set of $\mathbb{U}$. Define a closure operator $\mathrm{cl} : P(\mathbb{U}) \to P(\mathbb{U})$ by the condition

$$u \in \mathrm{cl}(\overline{w}) \quad \text{if and only if} \quad \dim \mathrm{loc}(u\overline{w}/\mathrm{k}) = \dim \mathrm{loc}(\overline{w}/\mathrm{k})$$

for $\overline{w} \subset \mathbb{U}$ finite. And

$$\mathrm{cl}(W) = \bigcup \{\mathrm{cl}(\overline{w}) : w \subseteq_{\mathrm{fin}} W\}$$

for W infinite.

Lemma 5.2. *Suppose* $W \in P(\mathbb{U})$ *and* $\mathrm{cl}(W) = W$. *Then, for any* $i \in I$, *the subset* $f_i(W) \subset \mathbb{X}_i(\mathrm{K})$ *is closed under* acl, *the algebraic closure in the sense of fields.*

There is an algebraically closed subfield $L = L_W \subseteq \mathrm{K}$.

$$f_i(W) = \mathbb{X}_i(L) \quad \textit{for all } i \in I.$$

Proof. Let $\overline{w} \in W^n$ and $f_i(\overline{w}) = \bar{x} \in \mathbb{X}_i^n(\mathrm{K})$. Let $y \in \mathbb{X}_i(\mathrm{K})$ such that $y \in \mathrm{acl}(\bar{x})$, where acl is over the base field k. Thus, for

$$X = \mathrm{loc}(\bar{x}/\mathrm{k}), \qquad Y = \mathrm{loc}(\bar{x}y/\mathrm{k})$$

we have $\dim X = \dim Y$. Hence, since f_i is a local biholomorphism, for any $v \in f_i^{-1}(y)$, we have

$$\dim \mathrm{loc}(\overline{w}/\mathrm{k}) = \dim \mathrm{loc}(\overline{w}v/\mathrm{k}),$$

which implies $v \in \mathrm{cl}(\overline{w}) \subset W$. This proves that $f_i(W)$ is closed under acl and hence $f_i(W) = \mathbb{X}_i(L)$ for some algebraically closed field $L = L_{W,i}$.

We claim that $L_{W,i} = L_{W,j}$ for any $i, j \in I$. Indeed, consider the direct product $\mathbb{U} \times \mathbb{U}$ instead of $\mathbb{U}$ and

$$f_i \times f_j : \mathbb{U} \times \mathbb{U} \twoheadrightarrow X_i \times X_j$$

instead of f_i and f_j, which still are local biholomorphisms onto smooth algebraic varieties. Clearly, $\mathrm{cl}(W \times W) = W \times W$ for cl in the product structure and

$$\mathbb{X}_i(L_{W,ij}) \times \mathbb{X}_j(L_{W,ij}) = (f_i \times f_j)(W \times W) = \mathbb{X}_i(L_{W,i}) \times \mathbb{X}_j(L_{W,j}),$$

that is, $L_{W,ij} = L_{W,i} = L_{W,j} = L$. □

5.3. Recall (see [Bays et al. 2014]) that one calls $(\mathbb{U}, \mathrm{cl})$ a *quasiminimal pregeometry structure* if the following holds:

QM1 The pregeometry is determined by the language. That is, if $\mathrm{tp}(v\overline{w}) = \mathrm{tp}(v'\overline{w}')$ then $v \in \mathrm{cl}(\overline{w})$ if and only if $v' \in \mathrm{cl}(\overline{w}')$.

QM2 $\mathbb{U}$ is infinite-dimensional with respect to cl.

QM3 (Countable closure property) If $W \subset \mathbb{U}$ is finite then $\mathrm{cl}(W)$ is countable.

QM4 (Uniqueness of the generic type) Suppose that $W, W' \subseteq \mathbb{U}$ are countable subsets, $\mathrm{cl}(W) = W$, $\mathrm{cl}(W') = W'$ and W, W' enumerated so that $\mathrm{tp}(W) = \mathrm{tp}(W')$.
If $v \in \mathbb{U} \setminus W$ and $v' \in \mathbb{U} \setminus W'$ then $\mathrm{tp}(Wv) = \mathrm{tp}(W'v')$ (with respect to the same enumerations for W and W').

QM5 ($\aleph_0$-homogeneity over closed sets and the empty set) Let $W, W' \subseteq \mathbb{U}$ be countable closed subsets or empty, enumerated such that $\mathrm{tp}(W) = \mathrm{tp}(W')$, and let $\overline{w}, \overline{w}'$ be finite tuples from $\mathbb{U}$ such that $\mathrm{tp}(W\overline{w}) = \mathrm{tp}(W'\overline{w}')$, and let $v \in \mathrm{cl}(W\overline{w})$. Then there is $v' \in \mathbb{U}$ that $\mathrm{tp}(\overline{w}vW) = \mathrm{tp}(\overline{w}'v'W')$.

Proposition 5.4. *Assume that* k_0 *is algebraically closed*, $\dim \mathbb{U} = 1$ *and* $\mathrm{cdim}\, \mathsf{R} \geq \aleph_0$. *Then* $(\mathbb{U}(\mathsf{R}), \mathrm{cl})$ *is a quasiminimal pregeometry.*

Proof. QM1 is by definition.

QM2 is by the assumption on R.

QM3 follows from the fact that in the language of o-minimal structure $\mathrm{acl}(W)$ is countable and that $\mathrm{cl}(W) \subseteq \mathrm{acl}(W)$, by (8) and (9).

QM4 follows from the fact that $\mathbb{U}$ is one-dimensional irreducible and $v \notin \mathrm{cl}(W)$, $v' \notin \mathrm{cl}(W')$.

QM5. If W and W' are empty then the required follows from Proposition 4.10 when $\mathrm{k} = \mathrm{k}_0$. In the nonempty case we may assume by $\aleph_0$-homogeneity over $\varnothing$ that $W = W'$. Now Lemma 5.2 allows us to replace $\mathrm{tp}(\overline{w}W)$ and $\mathrm{tp}(\overline{w}'W')$ by $\mathrm{loc}(\overline{w}/L_W)$ and $\mathrm{loc}(\overline{w}'/L_W)$, and $\mathrm{tp}(\overline{w}vW)$ and $\mathrm{tp}(\overline{w}'v'W')$ by $\mathrm{loc}(\overline{w}v/L_W)$ and $\mathrm{loc}(\overline{w}'v'/L_W)$, respectively.

The existence of v' follows from Proposition 4.10 when $\mathrm{k} = L_W$. □

Now we recall that given a quasiminimal pregeometry structure $(\mathbb{U}, \mathrm{cl})$ one can associate with it an abstract elementary class containing the structure, see [Bays et al. 2014, 2.2–2.3], or more generally [Zilber 2017, 2.17–2.18]. Call this class $\mathfrak{U}_{\mathrm{glob}}$.

By definition, one starts with a structure $\mathbb{U} = \mathfrak{U}(\mathsf{R})$ for a R of cardinality $\aleph_1$. Define $\mathfrak{U}^-_{\mathrm{glob}}$ to be the class of all cl-closed substructures of $\mathbb{U}$ with embedding $\prec$ of structures defined as a *closed* embedding, that is, $\mathbb{U}_1 \prec \mathbb{U}_2$ if and only if $\mathbb{U}_1 \subset \mathbb{U}_2$ and, for finite $W \subset \mathbb{U}_1$,

$$\mathrm{cl}_{\mathbb{U}_1}(W) = \mathrm{cl}_{\mathbb{U}_2}(W).$$

Now define $\mathfrak{U}_{\mathrm{glob}}$ to be the smallest class which contains $\mathfrak{U}^-_{\mathrm{glob}}$ and is closed under unions of $\prec$-chains.

Lemma 5.5. $$\mathfrak{U} \subseteq \mathfrak{U}_{\mathrm{glob}}.$$

Proof. We need to show that $\mathbb{U}(\mathsf{R}) \in \mathfrak{U}_{\mathrm{glob}}$, for any $\mathsf{R} \in \mathrm{Mod}_{\mathrm{An}}$.

We prove by induction on $\kappa = \mathrm{card}\,\mathsf{R} \geq \aleph_1$ that there is a κ-chain

$$\{\mathbb{U}_\lambda \in \mathfrak{U}_{\mathrm{glob}} : \lambda \in \kappa\} \quad \text{such that} \quad \bigcup_{\lambda \in \kappa} \mathbb{U}_\lambda = \mathbb{U}(\mathsf{R}).$$

Indeed, R can be represented as

$$\mathsf{R} = \bigcup_{\lambda < \kappa} \mathsf{R}_\lambda$$

for an elementary chain

$$\{\mathsf{R}_\lambda \in: \lambda \in \kappa\}, \quad \mathrm{card}\,\mathsf{R}_\lambda = \mathrm{card}\,\lambda + \aleph_0, \quad \mathsf{R}_\lambda \prec \mathsf{R}_\mu \qquad \text{for } \lambda < \mu.$$

Hence

$$\mathbb{U}_\lambda := \mathfrak{U}(\mathsf{R}_\lambda) \in \mathfrak{U}_{\mathrm{glob}}$$

which proves the inductive step and the lemma. □

Theorem 5.6. *Assuming* $\dim_{\mathrm{K}} \mathbb{U} = 1$, *the class* $\mathfrak{U}_{\mathrm{glob}}$ *is an abstract elementary class extending* $\mathfrak{U}$. $\mathfrak{U}_{\mathrm{glob}}$ *is categorical in uncountable cardinals and can be axiomatised by an* $L_{\omega_1,\omega}(Q)$*-sentence.*

Proof. The first part is by Proposition 5.4 and Lemma 5.5. The second part is the main result, Theorem 2.3, of [Bays et al. 2014]. □

Acknowledgement

I would like to thank Martin Bays and Andres Villaveces for some useful remarks and commentaries.

References

[Bays 2009] M. Bays, *Categoricity results for exponential maps of 1-dimensional algebraic groups and Schanuel conjectures for powers and the CIT*, Ph.D. thesis, Oxford University, 2009, available at https://people.maths.ox.ac.uk/~bays/dist/thesis/thesis.pdf.

[Bays et al. 2014] M. Bays, B. Hart, T. Hyttinen, M. Kesälä, and J. Kirby, "Quasiminimal structures and excellence", *Bull. Lond. Math. Soc.* **46**:1 (2014), 155–163. MR Zbl

[Katok 1992] S. Katok, *Fuchsian groups*, University of Chicago, 1992. MR Zbl

[Peterzil and Starchenko 2008] Y. Peterzil and S. Starchenko, "Complex analytic geometry in a nonstandard setting", pp. 117–165 in *Model theory with applications to algebra and analysis*, vol. 1, edited by Z. Chatzidakis et al., London Math. Soc. Lecture Note Ser. **349**, Cambridge University Press, 2008. MR Zbl

[Peterzil and Starchenko 2013] Y. Peterzil and S. Starchenko, "Definability of restricted theta functions and families of abelian varieties", *Duke Math. J.* **162**:4 (2013), 731–765. MR Zbl

[Pillay and Steinhorn 1986] A. Pillay and C. Steinhorn, "Definable sets in ordered structures, I", *Trans. Amer. Math. Soc.* **295**:2 (1986), 565–592. MR Zbl

[Scanlon 2012] T. Scanlon, "Counting special points: logic, Diophantine geometry, and transcendence theory", *Bull. Amer. Math. Soc.* (*N.S.*) **49**:1 (2012), 51–71. MR Zbl

[Zilber 2010] B. Zilber, *Zariski geometries: geometry from the logician's point of view*, London Mathematical Society Lecture Note Series **360**, Cambridge University Press, 2010. MR Zbl

[Zilber 2016] B. Zilber, "Model theory of special subvarieties and Schanuel-type conjectures", *Ann. Pure Appl. Logic* **167**:10 (2016), 1000–1028. MR Zbl

[Zilber 2017] B. Zilber, "Analytic Zariski structures and non-elementary categoricity", pp. 299–324 in *Beyond first order model theory*, edited by J. Iovino, Taylor and Francis, 2017. MR Zbl

Received 25 Jul 2022.

BORIS ZILBER:
zilber@maths.ox.ac.uk
Mathematical Institute, University of Oxford, Oxford, United Kingdom

Model Theory
Vol. 3, No. 1, 2024
https://doi.org/10.2140/mt.2024.3.119

A Pila–Wilkie theorem for Hensel minimal curves

Victoria Cantoral Farfán, Kien Huu Nguyen,
Mathias Stout and Floris Vermeulen

Recently, a new axiomatic framework for tameness in henselian valued fields was developed by Cluckers, Halupczok, Rideau-Kikuchi and Vermeulen and termed Hensel minimality. In this article we develop Diophantine applications of Hensel minimality. We prove a Pila–Wilkie type theorem for transcendental curves definable in Hensel minimal structures. In order to do so, we introduce a new notion of point counting in this context related to dimension counting over the residue field. We examine multiple classes of examples, showcasing the need for this new dimension counting, and prove that our bounds are optimal.

1. Introduction

1.1. ***The Pila–Wilkie theorem.*** In 1989, Bombieri and Pila [1989] developed a very fruitful method to count integral and rational points on various types of geometric objects in $\mathbb{R}^2$. This method is now called the determinant method and is especially well suited for proving uniform upper bounds on points of bounded height. For a subset $X \subseteq \mathbb{R}^n$, recall that the counting function is defined as

$$N(X; B) = \#\{x \in X \cap \mathbb{Q}^n \mid H(x) \leq B\},$$

where $H(a_1/b_1, \dots, a_n/b_n) = \max(|a_i|, |b_i|)$ when $\gcd(a_i, b_i) = 1$ for all i. For example, if $f : [0, 1] \to [0, 1]$ is an analytic transcendental function and X denotes its graph, Bombieri and Pila proved that for any $\varepsilon > 0$, there is a constant c_ε such that

$$N(X; B) \leq c_\varepsilon B^\varepsilon \quad \text{for all } B > 0.$$

Cantoral Farfán was partially supported by KU Leuven IF C14/17/083. Nguyen was partially supported by FWO Flanders (Belgium) with grant number 12X3519N and by the Excellence Research Chair “FLCarPA: L-functions in positive characteristic and applications” financed by the Normandy Region. Stout was supported by KU Leuven. Vermeulen was partially supported by KU Leuven IF C14/17/083, and partially by F.W.O. Flanders (Belgium) with grant number 11F1921N.
MSC2020: primary 03C99, 14G05; secondary 03C65, 03C98, 11D88, 11G50.
Keywords: Pila–Wilkie theorem, nonarchimedean geometry, analogues to o-minimality, tame geometry on valued fields, model theory of valued fields, point counting, rational points of bounded height, parametrizations.

A vast generalization of this result is the celebrated Pila–Wilkie theorem [Pila and Wilkie 2006]. It states that, if $X \subseteq \mathbb{R}^n$ is definable in an o-minimal structure, then for any $\varepsilon > 0$, there exists a constant c_ε such that

$$N(X^{\text{trans}}; B) \leq c_\varepsilon B^\varepsilon \quad \text{for all } B > 0. \tag{1.1.1}$$

Here X^{trans} denotes the transcendental part of X, obtained from X by removing all positive-dimensional connected semialgebraic subsets of X. The proof of this result is based heavily on the existence of C^r-parametrizations, which were originally developed by Gromov and Yomdin [Yomdin 1987a; 1987b; Gromov 1987].

In the nonarchimedean setting, such parametrization results were first proved in [Cluckers et al. 2015], and a corresponding Pila–Wilkie theorem was obtained for subanalytic sets in $\mathbb{Q}_p$. These results were further improved in [Cluckers et al. 2020], where a uniform version of these bounds was proved for subanalytic sets in $\mathbb{Q}_p$ and $\mathbb{F}_p((t))$.

1.2. *Hensel minimality.* Hensel minimality, or h-minimality for short, is a recent framework for tame nonarchimedean geometry, developed by Cluckers, Halupczok and Rideau-Kikuchi in equicharacteristic zero in [Cluckers et al. 2022] and extended to mixed characteristic together with the fourth author in [Cluckers et al. 2023]. It encompasses the aforementioned analytic structure on $\mathbb{Q}_p$ as a special case, but it applies more broadly, see, e.g., [Cluckers et al. 2022, Section 6] for several examples.

Hensel minimality bears a striking resemblance to the classical theory of o-minimality. In an o-minimal structure K, each definable subset $X \subseteq K$ is a finite union of intervals and points. In other words, there is some finite tuple $(a_i)_{i \in I}$ such that X is a union of fibers of the map $x \mapsto (\operatorname{sgn}(x - a_i))_{i \in I}$. Roughly speaking, h-minimality replaces the sign map by the *leading term* map

$$\operatorname{rv} : K \to K^\times/(1 + \mathcal{M}_K) \cup \{0\},$$

where K is a valued field and $\mathcal{M}_K$ is the maximal ideal of its valuation ring.

The goal of this article is to develop an analogue of the Pila–Wilkie theorem in an h-minimal context. For this purpose, we need two important consequences of Hensel minimality: a cell decomposition statement (Theorem 3.2.5) and the Jacobian property (Theorem 3.2.1). These theorems are analogues of o-minimal cell decomposition and the monotonicity theorem, respectively. We use them to prove a T_r*-parametrization* statement for curves definable in Hensel minimal structures (Theorem 4.1.1). These T_r-parametrizations are analogues of the C^r-parametrizations used in the proof of the o-minimal Pila–Wilkie theorem, and form a key technical ingredient.

1.3. ***Counting dimension.*** Let k be a field of characteristic zero and denote by $k((t))$ the field of Laurent series over k. For any natural number s, we let $k[t]_s \subseteq k((t))$ be the set of polynomials in t of degree less than s. For (transcendental definable) curves $X \subseteq k((t))^n$ we study the number of points on $X_s := X \cap (k[t]_s)^n$ as a valued field analogue of rational points of bounded height. In particular, we are interested in bounding the growth in terms of s, similar to the Pila–Wilkie theorem.

When $k = \mathbb{C}$ and X is a transcendental curve definable in the subanalytic structure on $\mathbb{C}((t))$, then X_s is finite for each $s \in \mathbb{N}$ [Binyamini et al. 2022, Theorem 1]. However, this result does not carry over to arbitrary h-minimal structures and in Section 2.4 we give explicit examples of transcendental curves X definable in some h-minimal structure, with infinite X_s. This issue is resolved by counting relative to the residue field. More precisely, we introduce the notion of *counting dimension*, denoted by #- dim, and consider bounds of the form

$$\#\text{-}\dim(X_s) \leq (N(s), d, e(s)), \tag{1.3.1}$$

where N and e are functions $\mathbb{N} \to \mathbb{N}$ and $d \in \mathbb{N}$ is constant. Intuitively, the above inequality can be thought of as a means of expressing that

$$\#X_s \leq N(s)\, \#(k[t]/(t^{e(s)}))^d,$$

even when k is infinite. When $N = N(s)$ is constant, $d \cdot e(s)$ can be thought of as bounding the growth of the k-dimension of X_s.

For $k = \mathbb{F}_p$, note that the right-hand side becomes $N(s)p^{de(s)}$. Comparing this to the known results for transcendental definable curves in $\mathbb{F}_p((t))$ [Cluckers et al. 2020, Theorem B] leads to the following question for transcendental curves $X \subseteq k((t))^n$: if $\varepsilon > 0$ is given, can we take $N = N_\varepsilon$ constant, $d = 1$ and $e(s) = \lceil \varepsilon \cdot s \rceil$ in (1.3.1)? We stress that the importance of the counting dimension is that it makes this question meaningful when k is infinite.

We will return to the motivations for the counting dimension after precisely defining it for any henselian valued field K (and not just $k((t))$). Theorem 2.2.1 then positively answers our question above: if K is an equicharacteristic zero henselian valued field which is h-minimal, satisfying some mild extra conditions, then for any transcendental definable curve $X \subseteq K^n$ and any $\varepsilon > 0$ there exists some constant $N_\varepsilon > 0$ such that

$$\#\text{-}\dim(X_s) \leq (N_\varepsilon, 1, \lceil \varepsilon \cdot s \rceil).$$

We moreover show that this bound is optimal, by constructing certain transcendental definable curves of a specific form.

Additionally, we consider the case of algebraic curves in Theorem 2.2.3, and prove the analogue of the classical Bombieri–Pila bound here. Let us also mention that the main obstacle in extending these results to higher dimensions is that these parametrization results are only known for curves under Hensel minimality. Especially the higher-dimensional geometry under Hensel minimality has to develop further.

2. The counting dimension

2.1. ***Counting in valued fields.*** In this section we introduce the counting dimension and prove some basic results about it. We then state our main results on the counting dimension of transcendental and algebraic curves.

Let K be a nonarchimedean valued field equipped with an $\mathcal{L}$-structure, for some language $\mathcal{L}$ expanding the language of valued fields $\mathcal{L}_{\text{val}} = \{0, 1, +, \cdot, \mathcal{O}_K\}$. For $A \subset K$, a set $X \subset K^n$ is called A-definable if it is definable in $\mathcal{L}$ using parameters only from A. We call a set definable if it is K-definable. We denote by k the residue field of K, by $\mathcal{O}_K$ the valuation ring and by $\mathcal{M}_K$ the maximal ideal. Let $\Gamma_K^\times$ be the value group, where the valuation is written multiplicatively $|\cdot| : K \to \Gamma_K$.

Assume that K is henselian of equicharacteristic zero. Then there always exists a lift $\tilde{k} \subset K$ of the residue field k, i.e., a subfield of K which maps bijectively to k under the reduction map $\mathcal{O}_K \to k$. Fix also a *pseudo-uniformizer* $t \in K$; recall that this is any nonzero element of K with $|t| < 1$. For a positive integer s, define

$$\tilde{k}[t]_s = \left\{ \sum_{i=0}^{s-1} a_i t^i \;\middle|\; a_i \in \tilde{k} \right\}.$$

If X is a subset of K^n we define X_s to be $X \cap \tilde{k}[t]_s^n$. We call X_s the *set of rational points of height at most s on X*. Note that this set depends on the choice of pseudo-uniformizer t and on the lift $\tilde{k}$.

The prototypical example to keep in mind is $K = k((t))$ for some characteristic zero field k, with $\tilde{k} = k$ and t as pseudo-uniformizer. Here $\tilde{k}[t]_s$ is simply the set of polynomials over k of degree at most $s - 1$.

The set X_s is considered as a suitable analogue for the set of rational points of bounded height on X, where the height is captured by s. We will be interested in bounding the size of X_s, as s grows, for various types of subsets of K^n. We cannot simply use the number of points on X_s as a measure of size, since this set will typically be infinite. Instead we introduce the *counting dimension*, to measure the size of X_s relative to the residue field.

Definition 2.1.1. Let K be a henselian valued field of equicharacteristic zero equipped with an $\mathcal{L}$-structure, for some language $\mathcal{L}$ expanding the language of valued fields. Fix a pseudo-uniformizer t of K and a lift $\tilde{k}$ of the residue field. Let X be a subset of K^n, let d be a positive integer and let $N, e : \mathbb{N} \to \mathbb{N}$ be functions. Then we say that X has *counting dimension bounded by* (N, d, e) if there exists a definable function $f : X \to \mathcal{O}_K^d$ such that for every positive integer s, the composition

$$X_s \xrightarrow{f} \mathcal{O}_K^d \xrightarrow{\text{proj}} \left(\frac{\mathcal{O}_K}{(t^{e(s)})}\right)^d$$

has finite fibers of size at most $N(s)$. Here proj is the componentwise reduction map modulo $t^{e(s)}$.

We use the notation

$$\#\text{-}\dim X_s \le (N(s), d, e(s))$$

to mean that the counting dimension of X is bounded by (N, d, e).

This definition depends on the choice of pseudo-uniformizer t, the lift $\tilde{k}$ and the language $\mathcal{L}$. However, we suppress these in notation and always assume a fixed choice of t, $\tilde{k}$ and $\mathcal{L}$.

Our definition of the counting dimension is motivated on the one hand by counting rational points on definable subsets in local fields of mixed characteristic as in [Cluckers et al. 2015; 2020; 2023]. In that case, one has to use a different notion of rational points of bounded height, as there is no lift of the residue field. Consider for example $K = \mathbb{Q}_p$. Then, if X is a subset of $\mathbb{Q}_p^n$, define

$$X_s = \{x \in X \cap \mathbb{Z}^n \mid 0 \le x_i \le s \text{ for all } i\}.$$

Using this definition of rational points, if X has counting dimension bounded by $(N(s), d, e(s))$ then for every s, X_s contains no more than $N(s)p^{de(s)}$ points. So a bound on the counting dimension gives a corresponding bound on the number of points in X_s.

Our second motivation comes from the relation of the counting dimension to the Zariski dimension over the residue field, as in [Cluckers et al. 2015, Section 5]. Let $K = \mathbb{C}((t))$ and let X be a subset of K^n. Then for every s, X_s is a subset of $\mathbb{C}[t]_s$, which can be naturally identified with $\mathbb{C}^{ns}$. If X is algebraic, then X_s is a constructible set in $\mathbb{C}^{ns}$, and one can wonder how the Zariski dimension of X_s grows with s. In [Cluckers et al. 2015, Section 5], a bound for this quantity is provided. We obtain a similar bound using the counting dimension instead of the Zariski dimension.

2.2. ***Main results.*** Our main results concern the counting dimension of algebraic and transcendental curves definable in h-minimal structures.

By a *curve* $C \subset K^n$ we mean a set for which there exists a linear map $p : K^n \to K$ such that $p(C)$ is infinite and such that p has finite fibers on C. If the theory of K in $\mathcal{L}$ is 1-h-minimal — see [Cluckers et al. 2022, Definition 2.3.3] or Section 3.1 below — then by dimension theory [Cluckers et al. 2022, Theorem 5.3.4], a definable curve in K^n is the same as a definable set of dimension 1. We call a curve $C \subset K^n$ *transcendental* if every algebraic curve in K^n has finite intersection with C.

Let K be a henselian valued field equipped with an $\mathcal{L}$-structure and assume that $\mathrm{Th}_{\mathcal{L}}(K)$ is 1-h-minimal. Then we say that acl $=$ dcl for K if algebraic Skolem functions exist in every model of $\mathrm{Th}_{\mathcal{L}}(K)$. By this we mean that for any model K' of $\mathrm{Th}_{\mathcal{L}}(K)$ and any subset $A \subset K'$ we have that $\mathrm{acl}_{K'}(A) = \mathrm{dcl}_{K'}(A)$.

Our main result is the following analogue of the Pila–Wilkie theorem on the counting dimension of transcendental curves definable in Hensel minimal structures.

Theorem 2.2.1. *Suppose that K is a henselian valued field of equicharacteristic zero equipped with a 1-h-minimal structure. Fix a pseudo-uniformizer t and a lift of the residue field $\tilde{k}$. Suppose that* acl $=$ dcl *in K and that the subgroup of b-th powers in $k^\times$ has finite index, for some integer $b > 1$. Let $C \subset \mathcal{O}_K^n$ be a transcendental definable curve. Then for each $\varepsilon > 0$ there is a constant N such that for each integer $s \geq 0$,*

$$\#\text{-}\dim(C_s) \leq (N, 1, \lceil \varepsilon \cdot s \rceil).$$

Furthermore, the constant N can be taken to hold uniformly throughout all transcendental members of a given definable family of definable curves.

The key aspect here is the slow growth of the last component of the counting dimension, similar to the Pila–Wilkie theorem in the o-minimal setting [Pila and Wilkie 2006, Theorem 1.10]. The strategy of the proof is as follows. We use the notion of T_r-parametrizations, which form a suitable analogue for C^r-parametrizations in the nonarchimedean setting; see, e.g., [Cluckers et al. 2015; Cluckers et al. 2020].

(1) We apply cell decomposition to find a T_1-parametrization of C. The existence of such a cell decomposition follows from 1-h-minimality under the extra assumption that acl $=$ dcl in K; see [Cluckers et al. 2022, Theorem 5.2.4, Addendum 5] or Section 3.1 below.

(2) Using substitutions of the form $x \mapsto x^r$, we may even assume that we have a T_r-parametrization, for some suitably chosen integer r. For this, we need that the subgroup of b-th powers of $k^\times$ has a finite index in $k^\times$.

(3) We then use an adaptation of the Bombieri–Pila determinant method to catch all rational points of bounded height in a small ball in a single hypersurface.

(4) Finally, the fact that C is transcendental and definable in a 1-h-minimal structure then gives the desired result. Indeed, this follows from uniform finiteness in definable families; see [Cluckers et al. 2022, Lemma 2.5.3] or Section 3.1 below.

The crucial ingredient to extend Theorem 2.2.1 to higher-dimensional transcendental sets is the existence of T_r-parametrizations, which have not been proven to exist in general 1-h-minimal structures. However, if one assumes the existence of these parametrizations, then Theorem 2.2.1 follows via a similar approach based on the determinant method.

A mixed characteristic analogue of this result in $\mathbb{Q}_p$ was proven by Cluckers, Halupczok, Rideau-Kikuchi and Vermeulen [Cluckers et al. 2023, Theorem 4.1.6]. Here the notion of rational points of bounded height is defined as above. Namely, for X a subset of $\mathbb{Q}_p^n$ define

$$X_s = \{x \in X \cap \mathbb{Z}^n \mid 0 \leq x_i \leq s \text{ for all } i\}.$$

Then [Cluckers et al. 2023, Theorem 4.1.6] states that if $\mathbb{Q}_p$ carries a 1-h-minimal structure with acl $=$ dcl, and if $C \subset \mathbb{Q}_p^n$ is a transcendental definable curve, then for each $\varepsilon > 0$ there is a constant c such that for all $H \geq 1$ we have

$$\#C_s \leq cs^{\varepsilon}.$$

In this article, we restrict to equicharacteristic zero. The methods for both proofs are quite similar, with one major difference being that the residue field is no longer finite. This is the reason for introducing the counting dimension. Restricting to equicharacteristic zero has the added benefit that Hensel minimality is slightly easier to work with.

We also prove that Theorem 2.2.1 is optimal, in the sense that the last component $\lceil \varepsilon \cdot s \rceil$ of our bound cannot be improved. In more detail, one cannot replace it by a sublinear function $e(s)$, even if $N(s)$ is allowed to be completely arbitrary.

Theorem 2.2.2. *Let k be a field of characteristic zero. There exists a 1-h-minimal structure on $k((t))$ with* acl $=$ dcl *satisfying the following. Given a sublinear function $e : \mathbb{N} \to \mathbb{N}$ and any $N : \mathbb{N} \to \mathbb{N}$, there exists a definable transcendental curve $C \subseteq k((t))^2$ such that its counting dimension is not bounded by $(N(s), 1, e(s))$.*

In contrast with this result, we consider in Section 5 a specific analytic structure on $\mathbb{Q}_p((t))$ for which we are able to prove that the counting dimension of every definable transcendental curve is bounded by $(N, 1, 1)$ for some integer $N > 0$. This structure already contains many interesting examples of transcendental definable curves. For example, the graph of the exponential function $\exp : p\mathbb{Z}_p + t\mathbb{Q}_p[[t]] \to \mathbb{Q}_p((t))$ is definable.

As for algebraic curves, we prove the following theorem, generalizing the results from [Cluckers et al. 2015, Section 5].

Theorem 2.2.3. *Let K be a henselian valued field of equicharacteristic zero equipped with a 1-h-minimal structure. Assume that* acl $=$ dcl *and that the subgroup of b-th powers in $k^\times$ has finite index for some $b > 1$. Let $C \subset K^2$ be an irreducible algebraic curve of degree d for some positive integer d. Then there exists a constant c_d, depending only on d, such that*

$$\#\text{-}\dim C_s \leq (c_d s, 1, \lceil s/d \rceil).$$

This theorem can be considered the analogue of the classical Bombieri–Pila theorem [Bombieri and Pila 1989]. By considering the example $y = x^d$, we will show that one cannot improve the last component of the counting dimension in this result.

2.3. ***Some basic properties.*** Let us list some basic properties of the counting dimension. We will often use these implicitly in our proofs.

Proposition 2.3.1. *Let K be a nonarchimedean valued field in some language $\mathcal{L}$ expanding the language of valued fields and let X, X' be definable subsets of K^n. Assume that*

$$\#\text{-}\dim(X_s) \leq (N(s), d, e(s)), \quad \#\text{-}\dim(X'_s) \leq (N'(s), d', e'(s))$$

for some integers d, d' and some functions $N, N', e, e' : \mathbb{N} \to \mathbb{N}$. Then

(1) $\#\text{-}\dim((X \cup X')_s) \leq (N(s) + N'(s), \max\{d, d'\}, \max\{e(s), e'(s)\})$,

(2) $\#\text{-}\dim((X \times X')_s) \leq (N(s)N'(s), d + d', \max\{e(s), e'(s)\})$,

(3) *if $f : X' \to X$ is a definable map with finite fibers of size at most N'' for some integer N'', then the counting dimension of X' is bounded by $(N''N, d, e)$. If moreover f is surjective and* acl = dcl *in K, then the counting dimension of X is bounded by (N', d', e').*

Proof. The proof of the first two properties is straightforward, so let us prove the last property. Let $g : X \to \mathcal{O}_K^d$ be a definable map such that for every positive integer s, the composition

$$X_s \xrightarrow{g} \mathcal{O}_K^d \xrightarrow{\text{proj}} \left(\frac{\mathcal{O}_K}{(t^{e(s)})}\right)^d$$

has finite fibers of size at most $N(s)$. Take a similar such definable map $g' : X' \to \mathcal{O}_K^{d'}$ for X'. Then for every positive integer s, the composition

$$\text{proj} \circ g \circ f : X'_s \to \left(\frac{\mathcal{O}_K}{(t^{e(s)})}\right)^d$$

has finite fibers of size at most $N''N(s)$, so the counting dimension of X' is bounded by $(N''N, d, e)$.

Conversely, using acl = dcl there exists a section $f' : X \to X'$ for f. Then for every positive integer s, the composition

$$\text{proj} \circ g' \circ f' : X_s \to \left(\frac{\mathcal{O}_K}{(t^{e'(s)})}\right)^{d'}$$

has finite fibers of size at most $N'(s)$. $\square$

2.4. ***Some examples.*** We give some examples of bounding the counting dimension of transcendental curves, which the reader can keep in mind throughout this article. In each of these, we consider a curve C in a valued field K which is definable in some 1-h-minimal structure on K. We then give an upper bound for the counting dimension of C by simply computing the sets C_s.

Example 2.4.1. Consider the valued field $K = \mathbb{Q}_p((t))$ with valuation ring $\mathbb{Q}_p[[t]]$. In Section 5 we argue that there is a 1-h-minimal structure on K in which the exponential map

$$\exp : p\mathbb{Z}_p + t\mathbb{Q}_p[[t]] \to \mathbb{Q}_p((t)) : z \mapsto \sum_{i\geq 0} \frac{z^i}{i!}$$

is definable. Let us write $U = p\mathbb{Z}_p + t\mathbb{Q}_p[[t]]$.

Let C be the graph of this exponential function. We claim that this is a transcendental set. Indeed, suppose that $f(x, y) \in K[x, y]$ is a nonzero polynomial such that $f(x, \exp x) = 0$ for infinitely many $x \in U$. We take such an f of minimal degree. By h-minimality [Cluckers et al. 2022, Lemma 2.5.2], there is then an open ball $B \subset U$ on which this holds. Moreover, exp is differentiable on U with derivative exp. Write $f(x, y) = \sum_{i=0}^{d} f_i(x) y^i$. Define the polynomial

$$g(x, y) = \sum_{i=0}^{d} f_i'(x) y^i + \sum_{i=0}^{d-1} (i-d) f_i(x) y^i.$$

This polynomial is nonzero and has degree strictly smaller than f and for $x \in B$ we have

$$g(x, \exp x) = \frac{\mathrm{d}}{\mathrm{d}x}(f(x, \exp x)) - d f(x, \exp x) = 0.$$

This is the desired contradiction, showing that C is a transcendental set.

To compute the counting dimension, we use t as a uniformizer and $\mathbb{Q}_p \subset \mathbb{Q}_p((t))$ as a lift of the residue field. Now, if x is in $p\mathbb{Z}_p$ then $\exp x$ is in $\mathbb{Q}_p$. Hence

$$C_1 = \{(x, \exp x) \mid x \in p\mathbb{Z}_p\}.$$

In particular, C_1 is infinite. We claim that the counting dimension of C is bounded by $(1, 1, 1)$. For this purpose, consider the map $C \to \mathcal{O}_K : (x, \exp x) \mapsto x$ followed by projection $\mathcal{O}_K \to \mathcal{O}_K/\mathcal{M}_K = \mathbb{Q}_p$. This clearly has finite fibers on C_1. Even more, if x is an element of $U \cap \mathbb{Q}_p[t]$ which is not in $p\mathbb{Z}_p$ then automatically $\exp x$ is not in $\mathbb{Q}_p[t]$. Thus for any positive integer s, $C_1 = C_s$ and the counting dimension of C is bounded by $(1, 1, 1)$.

Example 2.4.2. Consider the field $\mathbb{R}((t))$ in the language of ordered valued fields. We expand the language by the full Weierstrass system $\mathcal{B}$ as in [Cluckers and

Lipshitz 2011, Section 3.1]. In more detail, let

$$A_{n,\alpha}((\mathbb{Z})) = \left\{ \sum_{i\in I} f_i t^i \;\middle|\; f_i \in A_{n,\alpha},\ I \subset \mathbb{Z} \text{ well ordered} \right\},$$

where $A_{n,\alpha}$ is the ring of real power series in $\mathbb{R}[[\xi_1, \dots, \xi_n]]$ with radius of convergence $> \alpha$, and define $B_{n,\alpha} = A_{n,\alpha}(\mathbb{Z})$. This is a real Weierstrass system in the sense of [Cluckers and Lipshitz 2011, Definition 3.1.1], and we equip $\mathbb{R}((t))$ with real analytic $\mathcal{B}$-structure. In particular, we have function symbols for all elements of the Weierstrass system $\mathcal{B}$. In [Nguyen et al. 2024], it is shown that the theory of $\mathbb{R}((t))$ in this language is 1-h-minimal. Now, by our choice of $\mathcal{B}$ the exponential

$$\exp : (-1, 1) + t\mathbb{R}[[t]] \to \mathbb{R}((t)) : z \mapsto \sum_{i\geq 0} \frac{z^i}{i!}$$

is definable. Denote by C the graph of exp, which as above is a transcendental set. We use $\mathbb{R} \subset \mathbb{R}((t))$ as a lift of the residue field and t as our choice of uniformizer. Then C_1 is infinite, since if x is in $(-1, 1)$ then $\exp(x)$ is again real. Consider the reduction map

$$C \to \mathcal{O}_K/(t) : (x, \exp x) \mapsto x \mod t.$$

Then this has finite fibers of cardinality at most 1 above C_1. In fact, if x is in $\mathbb{R}[t]_s$ but not in $\mathbb{R}$ then $\exp(x)$ is never in $\mathbb{R}[t]$. Thus for any $s \geq 1$ we have $C_s = C_1$ and so the counting dimension of C is bounded by $(1, 1, 1)$.

Example 2.4.3. Denote by $\mathcal{L}_{\text{omin}}$ the language of ordered rings. For any real number $r > 0$, let f_r denote the function

$$\mathbb{R}_{>0} \to \mathbb{R}_{>0} : x \mapsto x^r$$

and consider the expansion $\mathcal{L}$ of $\mathcal{L}_{\text{omin}}$ where we have a function symbol for every f_r. The theory of $\mathbb{R}$ in $\mathcal{L}$ is o-minimal, since it is a reduct of $\mathbb{R}_{\exp}$. Let K be a proper elementary extension of $\mathbb{R}$. Then we may turn K into a valued field by taking for $\mathcal{O}_K$ the convex closure of $\mathbb{R}$. Note that the theory $\mathcal{T}$ is power-bounded, so by [Cluckers et al. 2022, Theorem 7.2.4] the theory of K in the language $\mathcal{L} \cup \{\mathcal{O}_K\}$ is 1-h-minimal. Let t be a pseudo-uniformizer of K and denote by $\tilde{k}$ a lift of the residue field. Then $\tilde{k}$ is again a real closed field. For example, K might be the field of Hahn series $\mathbb{R}((t^{\mathbb{R}}))$ with pseudo-uniformizer t and $\tilde{k} = \mathbb{R}$.

Now let C be the graph of the function

$$K_{>0} \to K_{>0} : x \mapsto x^{\pi}.$$

This is a definable set by our choice of language. Clearly, this is also a transcendental set. But C_1 is simply the graph of $x \mapsto x^{\pi}$ on $\tilde{k}$, which certainly contains this graph on $\mathbb{R}$. In particular, this set is infinite. Also note that $C_s = C_1$ for any $s \geq 1$, so that the counting dimension of C is bounded by $(1, 1, 1)$.

3. Notation and background

3.1. ***Hensel minimality.*** In this section, we record some background material on Hensel minimality. We refer to [Cluckers et al. 2022; 2023] for further details.

Let $\mathcal{L}$ be a language containing $\mathcal{L}_{\text{val}} = \{0, 1, +, \cdot, \mathcal{O}_K\}$. Let $\mathcal{T}$ be a complete $\mathcal{L}$-theory whose models are nontrivially valued fields of equicharacteristic zero. Let K be a model of $\mathcal{T}$. We denote by $\mathcal{O}_K$ the valuation ring of K and by $\Gamma_K^\times$ the valuation group. The valuation is denoted by $|\cdot| : K \to \Gamma_K = \Gamma_K^\times \cup \{0\}$. By an open ball we mean a set of the form $B_{<\lambda}(a) = \{x \in K \mid |x - a| < \lambda\}$, where $a \in K$ and $\lambda \in \Gamma_K^\times$. Similarly, a closed ball is a set of the form $B_{\leq\lambda}(a) = \{x \in K \mid |x - a| \leq \lambda\}$. If B is an open ball as above, we denote by $\text{rad}_{\text{op}} B$ its radius λ, and similarly we use $\text{rad}_{\text{cl}} B$ for λ if B is a closed ball.

For $\lambda \leq 1$ an element of $\Gamma_K^\times$, let I_λ be the ideal $\{x \in K \mid |x| < \lambda\}$. We define $\text{RV}_\lambda^\times$ to be $K^\times/(1 + I_\lambda)$, with quotient map

$$\text{rv}_\lambda : K^\times \to \text{RV}_\lambda^\times.$$

We also consider $\text{RV}_\lambda = \text{RV}_\lambda^\times \cup \{0\}$. The map rv_λ extends to $K \to \text{RV}_\lambda$ via $\text{rv}_\lambda(0) = 0$. We write $\text{RV} = \text{RV}_1$ and $\text{rv} = \text{rv}_1$. The set RV combines information from the residue field and the value group. Indeed, there is a short exact sequence

$$1 \to (\mathcal{O}_K/\mathcal{M}_K)^\times \to \text{RV}^\times \to \Gamma_K^\times \to 1.$$

Now let $\lambda \leq 1$ be in $\Gamma_K^\times$ and let X be a subset of K. We say that a finite set C *λ-prepares* X if the following holds: for any $x, y \in K$, if

$$\text{rv}_\lambda(x - c) = \text{rv}_\lambda(y - c) \quad \text{for all } c \in C,$$

then either x and y are both in X, or they are both not in X. If $(\xi_c)_c \in \text{RV}_\lambda^{\#C}$ then the set

$$\{x \in K \mid \text{rv}_\lambda(x - c) = \xi_c \text{ for all } c \in C\}$$

is said to be a ball *λ-next to* C (if it is disjoint from C). Note that if such a set is disjoint from C, then it is indeed an open ball. We can rephrase preparing as follows. A finite set C λ-prepares X if for any ball B λ-next to C, either $B \subseteq X$ or $B \cap X = \varnothing$. Note also that the balls 1-next to a finite set C are precisely the maximal open balls disjoint from C.

3.2. ***Consequences of Hensel minimality.*** For this section, we fix a field K of equicharacteristic zero equipped with an $\mathcal{L}$-structure which is 1-h-minimal. By [Cluckers et al. 2022], we may freely add constants from K to the language $\mathcal{L}$ and preserve 1-h-minimality. Many of the results below are formulated only for $\varnothing$-definable objects, but therefore hold just as well for A-definable objects, for $A \subset K$.

Hensel minimality implies tameness results on various definable objects. For functions there is the Jacobian property and Taylor approximation.

Theorem 3.2.1 (Jacobian property [Cluckers et al. 2022, Corollary 3.2.6]). *Let $f : K \to K$ be a $\emptyset$-definable function. Then there exists a finite $\emptyset$-definable set C such that for every $\lambda \le 1$ in $\Gamma_K^\times$, every ball B λ-next to C and every $x_0, x \in B$, $x \neq x_0$, we have that*

(1) *the derivative f' (as defined in the usual way) exists on B and $\mathrm{rv}_\lambda \circ f'$ is constant on B,*

(2) $\mathrm{rv}_\lambda\big((f(x) - f(x_0))/(x - x_0)\big) = \mathrm{rv}_\lambda(f')$,

(3) *for any open ball $B' \subset B$, $f(B')$ is either a point or an open ball.*

Note in particular that (1) implies that $|f'|$ is constant on balls 1-next to C, since Γ_K is a quotient of RV_λ. We will also use the following corollary. Recall that $\tilde{k} \subset K$ is a lift of the residue field.

Corollary 3.2.2. *Let $\tilde{k} \subset K$ be a lift of the residue field of K. Let $f : \mathcal{O}_K \to \mathcal{O}_K$ be a $\emptyset$-definable function. Then for all but finitely many $a \in \tilde{k}$, the following property holds: for all $x, x_0 \in a + \mathcal{M}_K$, we have*

$$|f(x) - f(x_0)| \le |x - x_0|.$$

Proof. By the Jacobian property there exists a finite $\emptyset$-definable set $C \subseteq \mathcal{O}_K$ such that for every ball B which is 1-next to C there is a $\mu_B \in \Gamma_K$ such that for all $x, x_0 \in B$,

$$|f(x) - f(x_0)| = \mu_B |x - x_0|.$$

Moreover, by [Cluckers et al. 2022, Corollary 3.1.6] we may additionally assume that if such a B is open, then $f(B)$ is an open ball of radius $\mu_B \,\mathrm{rad}_{\mathrm{op}}(B)$. As C is finite and the balls $a + \mathcal{M}_K$ with $a \in \tilde{k}$ are pairwise disjoint, it follows that only finitely many $a + \mathcal{M}_K$ contain a point of C. Thus all but finitely many $a + \mathcal{M}_K$ are 1-next to C. Now suppose $\mu_B > 1$ for some $B = a + \mathcal{M}_K$ which is disjoint from C. This implies that $f(a + \mathcal{M}_K) = \mathcal{O}_K$ (and this is only possible if the value group is discrete). In particular we can only have $\mu_B > 1$ for finitely many such B. Indeed, otherwise $f^{-1}(y)$ would be infinite for all $y \in \mathcal{O}_K$, contradicting [Cluckers et al. 2022, Lemma 2.8.1]. Hence, the desired property holds for cofinitely many $a \in \tilde{k}$. □

The second result we need is about Taylor approximation. For a function $f : X \subset K \to K$ on an open set X which is r-fold differentiable and $x_0 \in X$ we define the *r-th order Taylor polynomial of f at x_0* to be as usual

$$T_{f,x_0}^{\le r}(x) = T_{f,x_0}^{<r+1}(x) = \sum_{i=0}^{r} \frac{f^{(i)}(x_0)}{i!}(x - x_0)^i.$$

The following result basically states that any definable function can be well approximated by its Taylor polynomial up to some fixed order, at least away from finitely many points.

Theorem 3.2.3 (Taylor approximation of order r [Cluckers et al. 2022, Theorem 3.2.2]). *Let $f : K \to K$ be a $\varnothing$-definable function and fix a positive integer r. Then there exists a finite $\varnothing$-definable set C such that for every ball B 1-next to C, f is $(r+1)$-fold differentiable on B, $|f^{(r+1)}|$ is constant on B and for $x, x_0 \in B$ we have*

$$|f(x) - T^{\leq r}_{f,x_0}(x)| \leq |f^{(r+1)}(x_0)(x - x_0)^{r+1}|.$$

We will also need results on cell decomposition. If $\mathcal{T}$ is a 1-h-minimal theory, then by [Cluckers et al. 2022, Proposition 4.3.3] there is an expansion of the language by predicates on cartesian powers of RV such that the resulting structure is still 1-h-minimal and we have $\mathrm{acl} = \mathrm{dcl}$. In particular, we can typically assume that $\mathrm{acl} = \mathrm{dcl}$ without any problems.

Definition 3.2.4. Let $A \subset K$ be a parameter set. For $n \geq m$, let $\pi_{\leq m} : K^n \to K^m$ be the projection on the first m coordinates and $X \subset K^n$ an A-definable set. Consider, for $i = 1, \dots, n$, values $j_i \in \{0, 1\}$ and A-definable functions $c_i : \pi_{<i}(X) \to K$. Fix also an A-definable set

$$R \subseteq \prod_{i=1}^{n} (j_i \cdot \mathrm{RV}^\times),$$

where $0 \cdot \mathrm{RV}^\times = \{0\}$. We say that X is an A-definable cell if

$$X = \{x \in K^n \mid \mathrm{rv}(x_i - c_i(\pi_{<i}(x)))_{i=1,\dots,n} \in R\}.$$

We call X a cell of type $(j_1, \dots, j_n)$. The functions c_i are called the *cell centers*. A *twisted box* of the cell X is a set of the form

$$\{x \in K^n \mid \mathrm{rv}(x_i - c_i(\pi_{<i}(x)))_{i=1,\dots,n} = r\},$$

for $r \in R$.

By [Cluckers et al. 2022, Theorem 5.2.4], for a $\varnothing$-definable set $X \subseteq K^n$ there always exists a $\varnothing$-definable *cell decomposition*, i.e., a partition of X into finitely many $\varnothing$-definable cells A_ℓ. We will need the following variant. Recall that a function $f : X \subset K^n \to K^m$ is said to be 1-*Lipschitz* if for all $x, x' \in X$ we have

$$|f(x) - f(x')| \leq |x - x'|,$$

where we use the maximum norm on K^n.

Theorem 3.2.5 (cell decomposition [Cluckers et al. 2022, Theorem 5.2.4, Addendum 5]). *Assume that K carries a 1-h-minimal structure with* $\mathrm{acl} = \mathrm{dcl}$. *Let $X \subset K^n$ be $\varnothing$-definable. Then there exist a partition of X into finitely many $\varnothing$-definable sets*

A_ℓ such that for every ℓ there is some coordinate permutation $\sigma_\ell : K^n \to K^n$ such that $\sigma_\ell(A_\ell)$ is a cell of type $(1, \dots, 1, 0, \dots, 0)$ and such that each component of each center is 1-Lipschitz.

3.3. *T_r-approximation.* We recall some useful definitions and results from [Cluckers et al. 2020, §4.2] about T_r-approximation. Let K be a henselian valued field of equicharacteristic zero which is 1-h-minimal in some language $\mathcal{L}$ expanding the language of valued fields.

Definition 3.3.1. Let $U \subseteq K^m$ be an open set, let $\psi = (\psi_1, \dots, \psi_n) : U \to \mathcal{O}_K^n$ be a function, and let $r > 0$ be an integer. We say that ψ satisfies *T_r-approximation* if for each $y \in U$ there is an n-tuple $T_y^{<r}$ of polynomials with coefficients in $\mathcal{O}_K$ and of degree less than r that satisfies

$$|\psi(x) - T_y^{<r}(x)| \le |x - y|^r \quad \text{for all } x \in U. \tag{3.3.1}$$

Let X be a definable subset of $\mathcal{O}_K^n$ of dimension m. We say that a definable family $(\varphi_i)_{i\in I}$ of functions $\varphi_i : U_i \to X_i \subseteq \mathcal{O}_K^n$ is a *T_r-parametrization of X* if $X = \bigcup_{i\in I} X_i$ and each φ_i is surjective and satisfies T_r-approximation.

Definition 3.3.2. Let $\alpha = (\alpha_1, \dots, \alpha_m) \in \mathbb{N}^m$ and define $|\alpha| = \alpha_1 + \cdots + \alpha_m$. We define the following sets and numbers:

$$\Lambda_m(k) := \{\alpha \in \mathbb{N}^m : |\alpha| = k\}, \qquad L_m(k) := \#\Lambda_m(k)$$
$$\Delta_m(k) := \{\alpha \in \mathbb{N}^m : |\alpha| \le k\}, \qquad D_m(k) := \#\Delta_m(k).$$

Note that $L_m(k)$ (resp. $D_m(k)$) is the number of monomials in m variables of degree exactly (resp. at most) k.

Fix an integer d and define, for all integers n and m such that $m < n$, the following integers:

$$\begin{aligned} \mu(n, d) &= D_n(d), \\ r(m, d) &= \min\{x \in \mathbb{Z} : D_m(x-1) \le \mu < D_m(x)\}, \\ V(n, d) &= \sum_{k=0}^{d} k L_n(k), \\ e(n, m, d) &= \sum_{k=1}^{r-1} k L_m(k) + r(\mu - D_m(r-1)). \end{aligned} \tag{3.3.2}$$

To apply the determinant method, we need the following lemma. The proof is a straightforward adaptation from [Cluckers et al. 2015, Lemma 3.3.1].

Lemma 3.3.3. *Let K be a henselian field of equicharacteristic zero. Let t be a pseudo-uniformizer of K. Fix integers μ, r and U an open subset of K^m which is contained in a product of m closed balls of radius $|t|^\rho$, where $\rho \ge 0$ is an integer.*

Fix $x_1, \dots, x_\mu \in U$ *and functions* $\psi_1, \dots, \psi_\mu : U \to K$. *Assume that the* ψ_i *satisfy* T_r*-approximation on* U *for some integer* r *with*

$$D_m(r-1) \le \mu < D_m(r).$$

Then

$$|\det(\psi_i(x_j))_{i,j}| \le |t|^{\rho e}.$$

4. Counting rational points on transcendental and algebraic curves

In this section we prove Theorems 2.2.1, 2.2.2 and 2.2.3. Throughout, let K be an equicharacteristic zero henselian valued field, equipped with a 1-h-minimal $\mathcal{L}$-structure, for some language $\mathcal{L}$ expanding the language of valued fields. We assume that acl $=$ dcl in K and that the subgroup of b-th powers in the residue field $k^\times$ has finite index for some $b > 1$. Fix a lift $\tilde{k}$ of the residue field and fix a pseudo-uniformizer t.

4.1. ***T_r-parametrizations.*** Crucial to our approach is the following theorem, which asserts the existence of T_r-parametrizations for definable planar curves.

Theorem 4.1.1. *Let* K *be an equicharacteristic zero valued field, equipped with a* 1*-h-minimal* $\mathcal{L}$*-structure. Assume that* acl $=$ dcl *in* K *and that the subgroup of* b*-th powers in* $k^\times$ *has finite index for some* $b > 1$. *Let* $Y \subset K^n$ *be a definable set. Fix a positive integer* r *and let* $C \subset Y \times \mathcal{O}_K^2$ *be a definable set such that for every* $y \in Y$, C_y *is a curve. Then there exist finitely many maps* $\phi_1, \dots, \phi_N : Y \times \mathcal{O}_K \to C$ *such that for every* $y \in Y$, $\phi_{1,y}, \dots, \phi_{N,y}$ *form a* T_r*-parametrization for* C_y.

To prove this theorem, we start with a T_1-parametrization of our curve, which exists because of Theorem 3.2.5. To move from a T_1-parametrization to a T_r-parametrization for curves we will use power substitutions.

Lemma 4.1.2. *Let* $X \subseteq K$ *and let* $f : X \to \mathcal{O}_K$ *be a* $\varnothing$*-definable* 1*-Lipschitz function. Fix a positive integer* r. *Then there exists a finite* $\varnothing$*-definable set* C *such that the following holds. Let* B *be a ball* 1*-next to* C *contained in* $\mathcal{O}_K$, *say* 1*-next to* $c \in C$. *For* $a, b \in \mathcal{O}_K$ *consider the map* $p_r : K \to K : x \mapsto a(x-c)^r + b$. *If* D *is any open ball not containing* 0 *with* $p_r(D) \subseteq B$, *then* $f \circ p_r$ *satisfies* T_r*-approximation on* D *with respect to its Taylor polynomial. Moreover, for* $y \in D$ *and* $j = 1, \dots, r$ *there is the bound*

$$|\partial^j (f \circ p_r)(y)| \le |y|^{r-j}.$$

Proof. Use Theorems 3.2.1 and 3.2.3 to find a finite $\varnothing$-definable set C such that f satisfies Taylor approximation up to order r on balls 1-next to C and such that the first r derivatives of f satisfy the Jacobian property on balls 1-next to C. Without loss of generality, let $c = 0$, $a = 1$ and $b = 0$, and fix $x_0 \in D$. Then $x_0^r \in B$ and

$\mathrm{rad}_{\mathrm{op}}\, B = |x_0|^r$ since B is 1-next to 0. The fact that f is 1-Lipschitz gives

$$|f'(x_0^r)| \leq 1.$$

By the Jacobian property, the first r derivatives of f all have a constant norm on B. Thus, for $i \leq r$,

$$f^{(i)}(B) \subseteq \{y \in K \mid |y| = |f^{(i)}(x_0^r)|\},$$

and hence $\mathrm{rad}_{\mathrm{op}}\, f^{(i)}(B) \leq |f^{(i)}(x_0^r)|$. On the other hand, the Jacobian property yields

$$|f^{(i)}(x_0^r)| = \frac{\mathrm{rad}_{\mathrm{op}}\, f^{(i-1)}(B)}{\mathrm{rad}_{\mathrm{op}}\, B} \leq \frac{|f^{(i-1)}(x_0^r)|}{|x_0^r|}.$$

Using induction and the fact that f is 1-Lipschitz this gives for $1 \leq i \leq r$ that

$$|f^{(i)}(x_0^r) x_0^{r(i-1)}| \leq 1. \tag{4.1.1}$$

Let $x \in D$. Then x^r is in the same ball 1-next to 0 as x_0^r and so $|x^r| = |x_0^r|$. Since both f and p_r have Taylor approximation up to order r on their respective domains B and D we can conclude by [Cluckers et al. 2015, Lemma 3.2.7]. □

Proof of Theorem 4.1.1. By enlarging r if necessary, we may assume that r is a power of b. Using cell decomposition (Theorem 3.2.5) uniformly in $y \in Y$ we obtain for every $y \in Y$ finitely many sets $P_{i,y}$ whose union is C_y such that every $P_{i,y}$ is, after a coordinate permutation, a $(1, 0)$-cell or a $(0, 0)$-cell with 1-Lipschitz centers. Since everything below works uniformly in y, we drop the subscript y from now on.

The $(0, 0)$-cells are just singletons, so let us focus on one of the $(1, 0)$-cells, say P_ℓ. After a coordinate transformation and a translation we may assume that P_ℓ is the graph of a 1-Lipschitz map

$$\phi : P \subset \mathcal{O}_K \to \mathcal{O}_K,$$

where P is a cell with center 0. The group of r-th powers in $k^\times$ has finite index in $k^\times$, since r is a power of b. Let $a_1, \ldots, a_m \in \mathcal{O}_K^\times$ reduce to representatives for the cosets of $(k^\times)^r$ in $k^\times$. For $i = 1, \ldots, m$ and $j = 0, \ldots, r-1$ let

$$D_{i,j} = \{y \in K \mid a_i t^j y^r \in P\}.$$

Then the finitely many maps

$$p_{i,j} : D_{i,j} \to P : y \mapsto a_i t^j y^r$$

cover P. Now by Lemma 4.1.2, we can find a further subdivision of P such that the $\phi \circ p_{i,j}$ are all T_r on open balls contained in $D_{i,j}$. We prove that actually $\phi \circ p_{i,j}$ even has T_r-approximation on all of $D_{i,j}$. So let $x, y \in D_{i,j}$. Since $D_{i,j}$ is a cell with center 0, if $\mathrm{rv}(x) = \mathrm{rv}(y)$ then x and y are in the same ball contained in $D_{i,j}$ and we are done. So assume that $\mathrm{rv}(x) \neq \mathrm{rv}(y)$. Then

$$|\phi(a_i t^j x^r) - T^{<r}_{\phi\circ p_{i,j},y}(x)| \\ \leq \max\bigl\{|\phi(a_i t^j x^r) - \phi(a_i t^j y^r)|, |\phi(a_i t^j y^r) - T^{<r}_{\phi\circ p_{i,j},y}(x)|\bigr\}.$$

For the first term, use that ϕ is 1-Lipschitz to obtain

$$|\phi(a_i t^j x^r) - \phi(a_i t^j y^r)| \leq |x^r - y^r| \leq \max\{|x|^r, |y|^r\} = |x-y|^r,$$

since $\mathrm{rv}(x) \neq \mathrm{rv}(y)$. For the second term, we use the bound provided by Lemma 4.1.2 to get that

$$\begin{aligned} |\phi(a_i t^j y^r) - T^{<r}_{\phi\circ p_{i,j},y}(x)| &\leq \max_{\ell=1,\dots,r-1} \left|\frac{\partial^\ell(\phi\circ p_{i,j})(y)}{\ell!}(x-y)^\ell\right| \\ &\leq \max_{\ell=1,\dots,r-1} |y|^{r-\ell}|x-y|^\ell \leq |x-y|^r. \end{aligned}$$

So $\phi\circ p_{i,j}$ satisfies T_r-approximation on all of $D_{i,j}$. In conclusion, the maps

$$\psi_{ij} : D_{ij} \to C : y \mapsto (p_{ij}(y), \phi(p_{ij}(y)))$$

all have T_r-approximation and their images cover P_ℓ. $\square$

4.2. ***Transcendental curves.*** The following lemma is an adapted version of [Cluckers et al. 2020, Lemma 5.1.3]. We use it to capture rational points of bounded height in a small ball in a single hypersurface.

Lemma 4.2.1. *Fix integers d, m, n with $m < n$ and consider r, V, e as defined in (3.3.2). Let s be a positive integer, let $U \subseteq \mathcal{O}_K^m$ and suppose a function $\psi = (\psi_1, \dots, \psi_n) : U \to \mathcal{O}_K^n$ satisfies T_r-approximation. For $\alpha > sV/e$ a positive integer, denote by $p : \mathcal{O}_K^m \to (\mathcal{O}_K/(t^\alpha))^m$ the projection map. Then for any fiber B of p, the image $\psi(B\cap U)_s$ is contained in an algebraic hypersurface of degree at most d. Moreover, V/e goes to* 0 *as d goes to infinity.*

Proof. Let $B \subseteq \mathcal{O}_K^m$ be a product of closed balls of radius $|t|^\alpha$, i.e., a fiber of the map p, and take points $P_1, \dots, P_\mu$ in $\psi(B\cap U)_s$. Take $x_i \in B\cap U$ such that $\psi(x_i) = P_i$. Consider the determinant

$$\Delta = \det((\psi(x_i))^j)_{1\leq i\leq \mu,\, j\in\Delta_n(d)}.$$

For $j \in \Delta_n(d)$, the notation $(y_1, \dots, y_n)^j$ is to be interpreted as $\prod_i y_i^{j_i}$. Since ψ satisfies T_r-approximation, Lemma 3.3.3 gives that $\mathrm{ord}_t(\Delta) \geq \alpha e$. Since the P_i are in $\tilde{k}[t]_s^n$, if Δ were nonzero then $\mathrm{ord}_t(\Delta) \leq sV$. But $\alpha > sV/e$, so that $\Delta = 0$.

Now we use the determinant method. Since $\Delta = 0$, the μ vectors $((\psi(x_i))^j)_j$ are linearly dependent. This implies that there exists some algebraic hypersurface of degree at most d passing through all of the points $\psi(x_i)$. Since this holds for any μ points in $\psi(B \cap U)_s$ we can find such a hypersurface containing all of $\psi(B \cap U)_s$. The last fact follows from an easy explicit calculation; see, e.g., [Pila 2004, p. 212]. □

We need one more projection lemma to reduce to the planar case. Let us call a set in K^2 *nonalgebraic up to degree* d if it has a finite intersection with every algebraic curve of degree at most d.

Lemma 4.2.2. *Let $C \subset K^n$ be a definable transcendental curve and fix a positive integer d. Then there exists a finite definable partition of C into sets C_i, together with coordinate projections $\pi_i : C_i \to K^2$, such that π_i is a bijection onto its image and $\pi_i(C_i)$ is nonalgebraic up to degree d.*

Proof. By the cell decomposition Theorem 3.2.5, we may partition C into finitely many definable sets C_i such that after a coordinate permutation, C_i is a $(1, 0, \dots, 0)$ cell with 1-Lipschitz centers. (We may disregard the finitely many $(0, \dots, 0)$-cells since these are just points.) In other words, C_i is the graph of a map

$$\phi : P \to K^{n-1} : x \mapsto (\phi_1(x), \dots, \phi_{n-1}(x)),$$

where all ϕ_j are 1-Lipschitz and $P \subset K$ is a 1-cell. Denote by $A_j \subset P$ the set of $x \in P$ such that in some neighborhood of x, the graph of ϕ_j is nonalgebraic up to degree d. Then the A_j are definable sets, and they cover P, since otherwise ϕ would not be transcendental. Thus we can further partition C_i into the graphs of ϕ over every A_j. Over A_j, projection onto the first and j-th coordinate gives the desired conclusion. □

Remark 4.2.3. The use of this lemma can be avoided by working with cylinders over X for each possible projection $K^n \to K^2$; see [Cluckers et al. 2015, p. 45].

With this, we can prove our main result. Recall that the strategy of the proof is as follows. Using T_r-parametrizations, we represent C as a finite union of graphs of functions satisfying T_r-approximation. Then Lemma 4.2.1 gives a hypersurface catching all rational points on C of height at most s. The fact that C is transcendental then gives the desired conclusion.

Proof of Theorem 2.2.1. Recall the definition of V, e and r from equation (3.3.2) and recall that V/e goes to zero as d goes to infinity, by Lemma 4.2.1. Take d such that

$$\frac{V}{e} < \varepsilon.$$

Up to enlarging r if necessary, we may assume that r is a power of b. By applying Lemma 4.2.2 we may assume without loss of generality that C is a planar curve in $\mathcal{O}_K^2$ which is nonalgebraic up to degree d.

By Theorem 4.1.1 there exist finitely many maps $\psi_i : U_i \subset \mathcal{O}_K \to C$ which together form a T_r-parametrization of C. Let us focus on one such ψ_i. By our construction, we may apply Lemma 4.2.1 with $\alpha = \lceil s\varepsilon \rceil$. This yields that $\psi_i(U_i \cap B)_s$ is contained in an algebraic hypersurface X of degree at most d, for any closed ball $B \subset \mathcal{O}_K$ of radius $|t|^\alpha$. Since $\psi_i(U_i)$ is contained in C, and since C is nonalgebraic up to degree d, this intersection $X \cap C$ is finite. Even more, by uniform finiteness in definable families [Cluckers et al. 2022, Lemma 2.5.2], the intersection of X with C is uniformly bounded (over all such X) by some integer N. Thus the counting dimension of C is bounded by

$$(NN', 1, \lceil \varepsilon s \rceil),$$

where N' is the total number of sets U_i required in the cell decomposition for C.

Finally, to prove a uniform upper bound on counting dimension in definable families, assume that $(C_y)_y$ is a definable family of transcendental curves, for $y \in Y \subseteq K^m$ with Y definable. Then the above proof can easily be made uniform in C_y. Indeed, by Lemma 4.2.2 we can assume that every C_y is a planar curve which is nonalgebraic up to degree d. The number of maps for a T_r-parametrization can be uniformly bounded in y since Theorem 4.1.1 is uniform in families. Similarly, the intersection of an algebraic curve of degree at most d with any C_y is finite, and thus uniformly bounded by [Cluckers et al. 2022, Lemma 2.5.2]. These two facts give the desired conclusion. □

4.3. ***Linear upper bounds are optimal.*** We now prove Theorem 2.2.2. This shows that the bound in Theorem 2.2.1 is optimal, in the sense that one cannot replace the last component $\lceil \varepsilon \cdot s \rceil$ by a sublinear function $e(s)$, even if we allow $N(s)$ to be completely arbitrary.

We recall the notion of rings of strictly convergent power series $\mathcal{O}_K\langle x_1, \dots, x_n \rangle$. By definition, their elements are those power series $\sum_{i\in\mathbb{N}} a_i x^i$ with coefficients in $\mathcal{O}_K$ such that $a_i \to 0$, when $|i| \to \infty$. Each $f \in \mathcal{O}_K\langle x_1, \dots, x_n \rangle$ can be naturally considered as a function $\mathcal{O}_K^n \to K$ and then also as a function $f : K^n \to K$, after extending by zero. By [Cluckers et al. 2022, Theorem 6.2.1] there exists a 1-h-minimal structure on K in which these functions are definable. Moreover, by [Cluckers et al. 2022, Proposition 4.3.3], there exists such a structure in which $\mathrm{acl} = \mathrm{dcl}$.

Proof of Theorem 2.2.2. We work in the structure on $K := k((t))$ as outlined above, in which all functions $\mathcal{O}_K^n \to \mathcal{O}_K$ defined by a strictly convergent power series are definable.

First fix any strictly increasing continuous function $\delta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{>0}$ such that $e(s) \leq \delta(s)$ for all $s \in \mathbb{N}$, with $\lim_{s\to\infty} \delta(s) = +\infty$ and $\lim_{s\to+\infty} \delta(s)/s = 0$. Then choose any strictly increasing function $F : \mathbb{R}_{\geq 0} \to \mathbb{R}_{>0}$ with $F(\mathbb{N}) \subseteq \mathbb{N}$ and such that for all $s \in \mathbb{N}$,

$$F(\delta(s)) > N(s).$$

Next, take any strictly increasing sequence of natural numbers $(N_n)_n$ with the property that for all $n \in \mathbb{N}$,

$$3N_{n-1}^2 F(N_{n-1}) < \frac{\delta^{-1}(N_n) - 1}{N_n}.$$

Such a sequence exists since $\lim_{u\to\infty} \delta^{-1}(u)/u = \infty$ as δ is sublinear. From these data, we construct $f \in \mathcal{O}_K\langle x\rangle$ as

$$f(x) = \sum_{n=0}^{\infty} t^{N_n} x^{N_n} \prod_{i,\ell=1}^{N_n} \prod_{j=1}^{F(N_n)} (x - i - jt^{\ell}).$$

Letting the sequence $(N_n)_n$ grow even faster, if necessary, we may assume that for each $d \in \mathbb{N}$ there is some n_d such that for all $n \geq n_d$ we have

$$N_n > d(N_{n-1} + N_{n-1}^2 F(N_{n-1})).$$

Hence, for $M = N_{n-1} + N_{n-1}^2 F(N_{n-1})$, the order of contact between f and its M-th order Taylor approximation exceeds dM. Bézout's theorem thus implies that f cannot be algebraic of any degree $d \in \mathbb{N}$ (as in [Binyamini et al. 2022, Proposition 1]). Hence the graph C of $f : \mathcal{O}_K \to K$ is a definable transcendental curve in K^2.

We now show that the counting dimension of C is not bounded by $(N, 1, e)$. Let $g : K^2 \to \mathcal{O}_K$ be any definable map. We show that the composition

$$C_s \xrightarrow{g} \mathcal{O}_K \xrightarrow{\text{proj}} \frac{\mathcal{O}_K}{(t^{e(s)})}$$

has a fiber of size strictly larger than $N(s)$ for some sufficiently large s.

Take any n and let s be such that

$$\delta(s) \leq N_n < \delta(s+1).$$

For each $i, j, \ell \in \mathbb{N}$, we have by construction that $f(i + jt^{\ell}) \in k[t]$. By our choice of s, we moreover have that the t-degree of $f(N_n + jt^{N_n})$ is strictly smaller than s when $1 \leq j \leq F(N_n)$. Indeed, it follows from the construction of $(N_n)_n$ that

$$\begin{aligned}\deg_t(f(N_n + jt^{N_n})) &\leq N_{n-1} + N_n N_{n-1} + N_n N_{n-1}^2 F(N_{n-1}) \\ &\leq 3N_n N_{n-1}^2 F(N_{n-1}) < \delta^{-1}(N_n) - 1 < s.\end{aligned}$$

Now define

$$S := \{(N_n + jt^{N_n}, f(N_n + jt^{N_n})) \mid 1 \leq j \leq F(N_n)\}$$

and note that $S \subseteq C_s$ by the above computation.

Define $h : \mathcal{O}_K \to \mathcal{O}_K : x \mapsto g(x, f(x))$. By Corollary 3.2.2 we may assume that

$$|h(N_n + jt^{N_n}) - h(N_n)| \leq |t^{N_n}| \quad \text{for all } j \in \mathbb{N},$$

possibly after increasing n (and s). As $N_n \geq \delta(s) \geq e(s)$, this implies that all points of S belong to the fiber at N_n of the composition

$$C_s \xrightarrow{g} \mathcal{O}_K \xrightarrow{\text{proj}} \frac{\mathcal{O}_K}{(t^{e(s)})}.$$

But by construction of F, the inequality $N_n \geq \delta(s)$ also implies $F(N_n) > N(s)$. As $\#S = F(N_n)$, we have found a fiber containing more than $N(s)$ elements of C_s. $\square$

4.4. ***Algebraic curves.*** In this section we prove Theorem 2.2.3, following along the lines of [Cluckers et al. 2015, Section 5]. We need some results on Hilbert functions.

For r a positive integer, denote by $K[x_0, \ldots, x_n]_r$ the space of homogeneous degree r polynomials. For I a homogeneous ideal in $K[x_0, \ldots, x_n]$, let $I_r = K[x_0, \ldots, x_n]_r \cap I$ and denote by $H_I(r) = \dim K[x_0, \ldots, x_n]_r / I_r$ the Hilbert function of I. Let $<$ be the monomial order on $K[x_0, \ldots, x_n]$ defined by $x^\alpha < x^\beta$ if $|\alpha| < |\beta|$ or $|\alpha| = |\beta|$ and $\alpha_i > \beta_i$ for some i and $\alpha_j = \beta_j$ for $j < i$. After reordering the variables, this is the graded reverse lexicographic order on monomials. Denote by $\mathrm{LT}(I)$ the ideal of leading terms of I, where the leading term of a homogeneous element $p(x)$ of $K[x_0, \ldots, x_n]$ is the monomial in $p(x)$ which is maximal for $<$. Then I and $\mathrm{LT}(I)$ have the same Hilbert functions, by [Cox et al. 1992, Chapter 9, Proposition 3.9]. For $i \in \{0, \ldots, n\}$ define

$$\sigma_{I,i}(r) = \sum_{|\alpha|=r,\, x^\alpha \notin \mathrm{LT}(I)} \alpha_i$$

and note that $r H_I(r) = \sum_i \sigma_{I,i}(r)$. Let X be an irreducible variety in $\mathbb{P}^n_K$ of degree d and dimension m, with homogeneous ideal I. The Hilbert function $H_I(r)$ of I agrees with the Hilbert polynomial $P_X(r)$ of X, for r sufficiently large. Recall that this is a degree m polynomial whose leading coefficient is $d/m!$. By [Broberg 2004], for $i = 0, \ldots, n$ there are real numbers $a_{I,i} \geq 0$ such that

$$\frac{\sigma_{I,i}(r)}{r H_I(r)} = a_{I,i} + O_{n,d}(1/r) \quad \text{for } r \to \infty.$$

Note also that $a_{I,0} + \cdots + a_{I,n} = 1$. We can now prove Theorem 2.2.3.

Proof of Theorem 2.2.3. We have an irreducible algebraic curve C in $\mathbb{A}^2_K$ of degree d. Put $C' = C(K) \cap \mathcal{O}^2_K$. Consider the embedding

$$\iota : \mathbb{A}^2_K \to \mathbb{P}^2_K : (x, y) \mapsto (1 : x : y)$$

and let I be the homogeneous ideal of the closure of $\iota(C)$ in $\mathbb{P}^2_K$. Let δ be a positive integer, which we will choose later depending on s, and define

$$M(\delta) = \{ j \in \mathbb{N}^3 \mid |j| = \delta,\ x^j \notin \mathrm{LT}(I) \}.$$

Let $\mu = \#M(\delta) = H_I(\delta)$, $\sigma_i = \sigma_{I,i}(\delta)$ for $i = 0, 1, 2$ and put $e = \mu(\mu-1)/2$. By Theorem 4.1.1 there exist finitely many maps $\phi_1, \dots, \phi_N : Y_i \subset \mathcal{O}_K \to C'$ forming a T_r-parametrization of C'. Let B_α be a closed ball in $\mathcal{O}_K$ of radius $|t|^\alpha$ for some integer α. Fix a positive integer s, take points $y_1, \dots, y_\mu$ in $(\phi_i(B_\alpha \cap Y_i))_s$ and consider the determinant

$$\Delta = \det(\iota(y_i)^j)_{j \in M(\delta),\, 1 \le i \le \mu}.$$

By Lemma 3.3.3 we have that $|\Delta| \le |t|^{\alpha e}$. On the other hand, since $\iota(y_i)$ has coordinates which are polynomials of degree $< s$, we find that Δ is in $\tilde{k}[t]$ of degree

$$\deg_t \Delta \le (s-1)(\sigma_1 + \sigma_2).$$

Therefore, if we take $\alpha > (s-1)(\sigma_1 + \sigma_2)/e$, then $\Delta = 0$. As in the proof of Lemma 4.2.1, using the determinant method, we can find a polynomial H in two variables, with coefficients in $\tilde{k}[t]$ and exponents in $M(\delta)$, which vanishes on $(\phi_i(B \cap Y_i))_s$. Since the exponents of H lie in $M(\delta)$, we also see that H does not vanish identically on C. By Bézout's theorem, the intersection of $H = 0$ and C consists of at most δd points.

We want to conclude by taking $\alpha = \lceil s/d \rceil$, so we look for a suitable δ now. Similarly to the proof of [Cluckers et al. 2015, Theorem 5.1.3], one obtains that

$$\frac{\sigma_i}{e} = \frac{2\alpha_i}{d} + O_d(\delta^{-1}).$$

By [Salberger 2007, Lemma 1.12], it follows that

$$\frac{\sigma_1 + \sigma_2}{e} \le \frac{1}{d} + O_d(\delta^{-1}).$$

Hence we may take $\delta = s\, O_d(1)$ so that

$$\frac{(s-1)(\sigma_1 + \sigma_2)}{e} < \left\lceil \frac{s}{d} \right\rceil = \alpha.$$

By Proposition 2.3.1 we can find a definable map $f : C' \to \mathcal{O}_K$ such that the composition

$$C_s \to \mathcal{O}_K \to \frac{\mathcal{O}_K}{(t^{\lceil s/d \rceil})}$$

has finite fibers of size at most $N d \delta = N s\, O_d(1)$, for every s. Now, the number of cells N required in the T_r-parametrization of C can be made uniform in definable families. Since the set of degree d curves in K^2 is a definable family, we may assume that $N = O_d(1)$. So we conclude that the counting dimension of C is bounded by

$$\left(s\, O_d(1),\, 1,\, \left\lceil \frac{s}{d} \right\rceil\right). \qquad \square$$

Example 4.4.1. We show that one cannot improve the last component in the counting dimension for algebraic curves. Let K be any equicharacteristic zero valued field equipped with a 1-h-minimal structure. Fix a pseudo-uniformizer t and a lift $\tilde{k}$ of the residue field. Denote by C the curve in $\mathcal{O}_K^2$ defined by $y = x^d$. We claim that the counting dimension of C is not bounded by $(N(s), 1, e(s))$ for any functions $N, e : \mathbb{N} \to \mathbb{N}$ for which $e(s) < \lceil s/d \rceil$ when s is sufficiently large. Let $f : C \to \mathcal{O}_K$ be any definable map.

Take $s' = sd + 1$ sufficiently large that $e(s') < \lceil s'/d \rceil = s + 1$. In particular, note that $e(s') \leq s$. Define the map $g : \mathcal{O}_K \to C : x \mapsto (x, x^d)$ and put $h = f \circ g : \mathcal{O}_K \to \mathcal{O}_K$. By Corollary 3.2.2, there exists an $a \in \tilde{k}$ such that for any $b \in \tilde{k}$ we have that

$$|h(a) - h(a + bt^s)| \leq |t^s|.$$

This implies that $h(a) \equiv h(a + bt^s)$ in $\mathcal{O}_K/(t^{e(s')})$, for all $b \in \tilde{k}$. Finally, by noting that $g(a + bt^s)$ lies in $C_{s'}$, one sees that the map

$$C_{s'} \xrightarrow{f} \mathcal{O}_K \xrightarrow{\text{proj}} \frac{\mathcal{O}_K}{(t^{e(s')})}$$

has an infinite fiber.

5. Curves with uniformly bounded counting dimension

We now consider an analytic structure on $\mathbb{Q}_p((t))$ where each $\varnothing$-definable curve C has an associated constant $N_C \in \mathbb{N}$ such that $\#\text{-}\dim(C_s) \leq (N_C, 1, 1)$. Contrast this with Theorem 2.2.2, where we produced curves, definable in some analytic structure, whose counting dimension cannot be bounded by any constant triple $(N, 1, 1)$. The stronger upper bounds in this section result from working in a more restricted analytic structure. The essential difference with the setting of Theorem 2.2.2 is that we now only add function symbols for power series whose coefficients do not involve the uniformizer t.

Throughout this section we take $t \in \mathbb{Q}_p((t))$ as our chosen pseudo-uniformizer and $\tilde{k} = \mathbb{Q}_p \subseteq \mathbb{Q}_p((t))$ as our lift of the residue field.

Definition 5.0.1. Let $\mathcal{L}_{\mathbb{Z}_p\langle x\rangle}$ be the language expanding the language of valued fields $\mathcal{L}_{\text{val}} = \{0, 1, +, \cdot, \mathcal{O}_K\}$ by

(1) a binary function symbol "$-$" and a unary function symbol $(\cdot)^{-1}$,

(2) a unary relation symbol $\mathcal{O}_{K,\text{fine}}$,

(3) n-ary function symbols for the elements of the rings of strictly convergent power series $\mathbb{Z}_p\langle x_1, \dots, x_n\rangle$, for $n \in \mathbb{N}$.

The field $K = \mathbb{Q}_p((t))$ admits a natural $\mathcal{L}_{\mathbb{Z}_p\langle x\rangle}$-structure. We interpret $\mathcal{O}_{K,\text{fine}}$ as the valuation ring $\mathbb{Z}_p + t\mathbb{Q}_p[[t]]$ and $\mathcal{O}_K$ as its equicharacteristic zero coarsening $\mathbb{Q}_p[[t]]$. Each function symbol $f \in \mathbb{Z}_p\langle x_1, \dots, x_n\rangle$ is interpreted naturally as a function $\mathcal{O}^n_{K,\mathrm{fine}} \to K$. Finally, "$-$" and $(\,\cdot\,)^{-1}$ are just subtraction and inversion on K, where the latter is extended by $0^{-1} = 0$. Note that $\mathbb{Q}_p((t))$, equipped with the valuation ring $\mathcal{O}_K$, is 1-h-minimal for this structure by [Cluckers et al. 2022, Theorem 6.2.1].

Theorem 5.0.2. *For each transcendental curve $C \subseteq \mathbb{Q}_p((t))^n$ $\varnothing$-definable in $\mathcal{L}_{\mathbb{Z}_p\langle x\rangle}$, there is some constant $N_C \in \mathbb{N}$ such that*

$$\#\text{-}\dim(C_s) \leq (N_C, 1, 1).$$

Remark 5.0.3. As before, the constant N_C can be made uniform in definable families.

Lemma 5.0.4. *For all $\lambda \in \mathbb{Z}_p^\times$ the map*

$$\tau_\lambda : \mathbb{Q}_p((t)) \to \mathbb{Q}_p((t)) : \sum_{j=k}^{+\infty} a_j t^j \mapsto \sum_{j=k}^{+\infty} a_j (\lambda t)^j$$

is an $\mathcal{L}_{\mathbb{Z}_p\langle x\rangle}$-automorphism.

Proof. Fix some $\lambda \in \mathbb{Z}_p^\times$. It is clear that τ_λ fixes the constants 0, 1 and that the relations $x \in \mathbb{Q}_p[\![t]\!]$ and $x \in \mathbb{Z}_p + t\mathbb{Q}_p[\![t]\!]$ are invariant under τ_λ. Similarly, it is straightforward to verify that τ_λ commutes with addition, multiplication and inversion.

It thus remains only to check that τ_λ commutes with all function symbols f in $\mathbb{Z}_p\langle x_1, \dots, x_n\rangle$. Take $(x_i)_{i=1}^n = \left(\sum_{j=k_i}^{+\infty} a_{ij} t^j\right)_{i=1}^n \in \mathbb{Q}_p((t))^n$. Then $f(x)$ is computed as follows:

$$f(x) = \begin{cases} 0 & \text{if } x_i \notin \mathbb{Z}_p + t\mathbb{Q}_p[\![t]\!] \text{ for some } i, \\ \sum_{s=0}^{+\infty} f_s((a_{ij})_{i,j}) t^s & \text{else,} \end{cases}$$

where each f_s is some quasihomogeneous polynomial of degree s in variables a_{ij}, each of weight j, for $1 \leq i \leq n$ and $0 \leq j \leq s$. This precisely means that $f_s((\lambda^j a_{ij})_{i,j}) = f_s((a_{i,j})_{i,j})\lambda^s$. It follows that f commutes with τ_λ. □

We will need that transcendental curves only have finite intersection with any semialgebraic curve: one-dimensional subsets of K^n definable in $\mathcal{L}_{\mathrm{val}} \cup K$. This follows from the fact that semialgebraic curves are locally algebraic, as made precise by the lemma below.

Lemma 5.0.5. *Let K be a valued field with valuation ring $\mathcal{O}_K$. Let C be a transcendental curve in K^n. Then, for any K-definable curve X in $\mathcal{L}_{\mathrm{val}}$, the intersection $C \cap X$ is finite.*

Proof. By valued field-quantifier elimination for henselian valued fields [Flenner 2011, Proposition 4.3], C is a finite union of sets $V \cap S$, where V is the vanishing locus of some polynomials with coefficients in K and S is of the form

$(\mathrm{rv}(f_i(x)))_{i=0}^n \in R$ for certain $f_i \in K[x]$ and $R \subseteq (\mathrm{RV}^\times)^n$. As X is of dimension one and S is open, only the zero- and one-dimensional irreducible components of V can have nonempty intersection with S. In particular, if $V \cap S$ met C at infinitely many points, then there would be an algebraic curve $X' \subseteq V$ containing infinitely many points of C. □

We now continue along the same lines as in the proof of [Binyamini et al. 2022, Theorem 1], using Lemma 5.0.4 instead of the quantifier elimination statement used there.

Proposition 5.0.6. *Let $C \subseteq \mathbb{Q}_p((t))^n$ be a transcendental curve which is $\varnothing$-definable in $\mathcal{L}_{\mathbb{Z}_p\langle x\rangle}$. Then $C_s \subseteq C_1$ for all integers $s > 0$. In particular, $\#\text{-}\dim(C_s) \leq (1, n, 1)$.*

Proof. Suppose $\left(\sum_{j=0}^{s-1} a_{ij} t^j\right)_i \in C_s \setminus C_1$. Consider C_s as a subset $A \subseteq \mathbb{Q}_p^{ns}$ via the identification

$$\left(\sum_{j=0}^{s-1} b_{ij} t^j\right)_i \mapsto (b_{ij})_{i,j}.$$

Let X be an algebraic curve in $\mathbb{Q}_p^{ns}$ containing all points $(\lambda^j a_{ij})_{i,j}$ for $\lambda \in \mathbb{Q}_p$. By the above Lemma 5.0.4, it follows that $X \cap A$ contains all points $(\lambda^j a_{ij})$ for $\lambda \in \mathbb{Z}_p^\times$. In particular, it is infinite.

Let Y be the image of $X(\mathbb{Q}_p[[t]])$ under $(x_{ij})_{i,j} \mapsto \left(\sum_{j=0}^{s-1} x_{ij} t^j\right)_i$. As this map and the curve $X(\mathbb{Q}_p[[t]])$ are definable in the 1-h-minimal $\mathcal{L}_{\mathrm{val}}$-structure on $\mathbb{Q}_p((t))$ (for the valuation ring $\mathbb{Q}_p[[t]]$), it follows from [Cluckers et al. 2022, Proposition 5.2.4(3,4)] that Y is an $\mathcal{L}_{\mathrm{val}} \cup \mathbb{Q}_p((t))$-definable set of dimension at most 1. Since the infinite set X injects into Y, it then follows by [Cluckers et al. 2022, Proposition 5.2.4(1)] that Y has dimension exactly 1. Now use that $X \cap A$ is infinite, whence so is $Y \cap C_s$. By Lemma 5.0.5, this contradicts the assumption that C is transcendental. □

Proof of Theorem 5.0.2. By [Cluckers et al. 2022, Theorem 5.7.3] we may assume that C is a single reparametrized cell (A, σ). As C is one-dimensional, it follows that either A is a finite collection of points (in which case we are done) or A is of type $(1, 0, \ldots, 0)$, up to a coordinate permutation of K^n. Thus A is the coordinate projection onto K^n of the graph of some $\mathcal{L}_{\mathbb{Z}_p\langle x\rangle}$-definable function $c : P \subseteq K \times \mathrm{RV}^\ell \to K^{n-1}$ (for some $\ell \in \mathbb{N}$). By [Cluckers et al. 2022, Corollary 2.6.7, Lemma 2.5.2] it holds that $\# c(x, \mathrm{RV}^\ell) \leq N$ for some $N \in \mathbb{N}$, independent of x.

Now consider the map $g : K^n \to K$, which is the projection onto the first coordinate on $\mathcal{O}_K^n$ and is identically zero on $K^n \setminus \mathcal{O}_K^n$. For any $x \in \mathcal{O}_K/(t) \cong \tilde{k} \subseteq \mathcal{O}_K$, the fiber at x of the composition

$$A_1 \xrightarrow{g} \mathcal{O}_K \xrightarrow{\text{proj}} \frac{\mathcal{O}_K}{(t)}$$

is precisely $(\{x\} \times c(x, \mathrm{RV}^{\ell})) \cap A_1$. In particular, it has size at most N. As $A_s = A_1$ for all $s \in \mathbb{N}$, by Proposition 5.0.6 it follows that $\#\text{-}\dim(A) \leq (N, 1, 1)$. $\square$

Acknowledgements

The authors would like to thank Raf Cluckers for suggesting this topic and the various discussions surrounding it. The authors also thank the anonymous referee for the many helpful comments which improved this article.

References

[Binyamini et al. 2022] G. Binyamini, R. Cluckers, and D. Novikov, "Point counting and Wilkie's conjecture for non-Archimedean Pfaffian and Noetherian functions", *Duke Math. J.* **171**:9 (2022), 1823–1842. MR Zbl

[Bombieri and Pila 1989] E. Bombieri and J. Pila, "The number of integral points on arcs and ovals", *Duke Math. J.* **59**:2 (1989), 337–357. MR Zbl

[Broberg 2004] N. Broberg, "A note on a paper by R. Heath-Brown: 'The density of rational points on curves and surfaces'", *J. Reine Angew. Math.* **571** (2004), 159–178. MR Zbl

[Cluckers and Lipshitz 2011] R. Cluckers and L. Lipshitz, "Fields with analytic structure", *J. Eur. Math. Soc.* (*JEMS*) **13**:4 (2011), 1147–1223. MR Zbl

[Cluckers et al. 2015] R. Cluckers, G. Comte, and F. Loeser, "Non-Archimedean Yomdin–Gromov parametrizations and points of bounded height", *Forum Math. Pi* **3** (2015), art. id. e5. MR Zbl

[Cluckers et al. 2020] R. Cluckers, A. Forey, and F. Loeser, "Uniform Yomdin–Gromov parametrizations and points of bounded height in valued fields", *Algebra Number Theory* **14**:6 (2020), 1423–1456. MR Zbl

[Cluckers et al. 2022] R. Cluckers, I. Halupczok, and S. Rideau-Kikuchi, "Hensel minimality, I", *Forum Math. Pi* **10** (2022), art. id. e11. MR Zbl

[Cluckers et al. 2023] R. Cluckers, I. Halupczok, S. Rideau-Kikuchi, and F. Vermeulen, "Hensel minimality, II: Mixed characteristic and a diophantine application", *Forum Math. Sigma* **11** (2023), art. id. e89. MR Zbl

[Cox et al. 1992] D. Cox, J. Little, and D. O'Shea, *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra*, Springer, 1992. MR Zbl

[Flenner 2011] J. Flenner, "Relative decidability and definability in henselian valued fields", *J. Symbolic Logic* **76**:4 (2011), 1240–1260. MR Zbl

[Gromov 1987] M. Gromov, "Entropy, homology and semialgebraic geometry", pp. 225–240 in *Séminaire Bourbaki*, vol. 1985/86, Astérisque **145-146**, Soc. Math. France, Paris, 1987. MR Zbl

[Nguyen et al. 2024] K. H. Nguyen, M. Stout, and F. Vermeulen, "Almost real closed fields with real analytic structure", preprint, 2024. arXiv 2401.10758v1

[Pila 2004] J. Pila, "Integer points on the dilation of a subanalytic surface", *Q. J. Math.* **55**:2 (2004), 207–223. MR Zbl

[Pila and Wilkie 2006] J. Pila and A. J. Wilkie, "The rational points of a definable set", *Duke Math. J.* **133**:3 (2006), 591–616. MR Zbl

[Salberger 2007] P. Salberger, "On the density of rational and integral points on algebraic varieties", *J. Reine Angew. Math.* **606** (2007), 123–147. MR Zbl

[Yomdin 1987a] Y. Yomdin, "Volume growth and entropy", *Israel J. Math.* **57**:3 (1987), 285–300. MR Zbl

[Yomdin 1987b] Y. Yomdin, "C^k-resolution of semialgebraic mappings. Addendum to *Volume growth and entropy*", *Israel J. Math.* **57**:3 (1987), 301–317. MR Zbl

Received 23 Dec 2022.

VICTORIA CANTORAL FARFÁN:

victoria.cantoralfarfan@mathematik.uni-goettingen.de
Mathematisches Institut, Georg-August Universität Göttingen, Göttingen, Germany

KIEN HUU NGUYEN:

kien.nguyenhuu@kuleuven.be
Department of Mathematics, KU Leuven, Leuven, Belgium

and

Laboratoire de Mathématiques Nicolas Oresme, Université de Caen Normandie, Caen, France

MATHIAS STOUT:

mathias.stout@kuleuven.be
Department of Mathematics, KU Leuven, Leuven, Belgium

FLORIS VERMEULEN:

floris.vermeulen@kuleuven.be
Department of Mathematics, KU Leuven, Leuven, Belgium

Model Theory
Vol. 3, No. 1, 2024
https://doi.org/10.2140/mt.2024.3.147

Model theory in compactly generated (tensor-)triangulated categories

Mike Prest and Rose Wagstaffe

We give an account of model theory in the context of compactly generated triangulated and tensor-triangulated categories $\mathcal{T}$. We describe pp formulas, pp-types and free realisations in such categories and we prove elimination of quantifiers and elimination of imaginaries. We compare the ways in which definable subcategories of $\mathcal{T}$ may be specified. Then we link definable subcategories of $\mathcal{T}$ and finite-type torsion theories on the category of modules over the compact objects of $\mathcal{T}$. We briefly consider spectra and dualities. If $\mathcal{T}$ is tensor-triangulated then new features appear, in particular there is an internal duality in rigidly-compactly generated tensor-triangulated categories.

1. Introduction and background

1A. ***Introduction.*** Model theory in a compactly generated triangulated category $\mathcal{T}$ falls within the scope of the model theory of modules via the restricted Yoneda embedding $\mathcal{T} \to \text{Mod-}\mathcal{T}^{\mathrm{c}}$ where $\mathcal{T}^{\mathrm{c}}$ denotes the subcategory of compact objects in $\mathcal{T}$. The model theory of modules over possibly many-sorted rings, such as $\mathcal{T}^{\mathrm{c}}$, is well-developed, but there are many special features of triangulated categories that make it worthwhile to directly develop model theory in the triangulated context. That is what we do here, and we also consider additional features which appear when the category is tensor-triangulated. A good number of the results appear elsewhere but we give a detailed and unified account which, we hope, will be a useful reference.

What began as the model theory of modules — the investigation of model-theoretic questions in the context of modules over a ring — has developed in scope — to much more general categories — in depth, and in purpose having for a long time been led by interests and questions coming from representation theory.

MSC2020: primary 18E45; secondary 03C60, 18G80.
Keywords: triangulated category, tensor-triangulated category, model theory, definable subcategory.

Many aspects — purity, pure-injectives, definable subcategories etc. — can be dealt with purely algebraically and, in the context of compactly generated triangulated categories, this was developed by Beligiannis [2000b] and Krause [2000] (for earlier relevant work, see [Christensen and Strickland 1998; Benson and Gnacadja 1999]). But, apart from a brief treatment in [Garkusha and Prest 2005], some use in [Arnesen et al. 2017] and a recent detailed exposition of some aspects in [Bennett-Tennenhaus 2023], there has not been much explicit appearance of model theory in triangulated categories. To some extent that is because there is a "dictionary" between model theoretic and algebraic/functor-category methods, allowing much of what can be proved with model theory to be proved by other methods. But sometimes what is obvious and natural using one language is not so easily translatable into the other. Moreover, model theory can give new insights and simpler proofs. Our main aim in this paper is to make the methods of model theory readily available to be used in compactly generated triangulated categories. Some aspects — dualities, spectra, enhancements, extensions to well-generated triangulated categories — are currently in development, so we don't aim to be comprehensive but we do present the more settled material in detail.

Some minimal acquaintance with model theory, at least with basic ideas in the model theory of modules, will be helpful for the reader but we do keep formal aspects of model theory to a minimum. Really, all that we need is the notion of a formula and its solution set in a structure.

We do need to use sorted variables. Variables in a formula are place-holders for elements from a structure; in our context these elements may belong to different *sorts*. The idea is very simple and well-illustrated by representations of the quiver A_2 which is $\bullet \to \star$. A representation of this quiver in the category of modules over a ring R consists of two R-modules $M_\bullet$, $M_\star$ and an R-linear map from $M_\bullet$ to $M_\star$. Such a structure is naturally two-sorted, with elements of the sort (labelled by) $\bullet$ being those of $M_\bullet$ and those of sort (labelled by) $\star$ being those of $M_\star$. The variables we would use in writing formulas reflect that, say with subscripts, and for this example we would use variables of two sorts (labelled respectively by $\bullet$ and $\star$). The difference between using a 2-sorted and 1-sorted language is the difference between treating (2-sorted) representations of that quiver (equivalently modules over the 2-sorted ring which is the (R-)path category of the quiver) and (1-sorted) modules over the path algebra of the quiver (the path algebra of the quiver is a normal, 1-sorted, ring). That is a matter of choice if there are only finitely many sorts but, because $\mathcal{T}^c$ is skeletally infinite, we do need to use sorted structures and take account of sorts in formulas. For more discussion, and many examples, of this, see [Prest 2019].

We suppose throughout this paper that $\mathcal{T}$ is a compactly generated triangulated category. We take this to include the requirement that $\mathcal{T}$ has infinite coproducts. We suppose that the reader knows something about these categories, but we do recall here that the derived category $\mathcal{D}(\text{Mod-}R)$ of the category Mod-R of R-modules is a basic example which is obtained from the category of chain complexes of R-modules by a type of localisation process which preserves homological information. The exact sequences of Mod-R give rise to *triangles* — certain triples of composable morphisms — in $\mathcal{D}(\text{Mod-}R)$. There is also a *shift* autoequivalence on $\mathcal{D}(\text{Mod-}R)$ which is induced by the shift operation on chain complexes. In general a *triangulated category* is an additive category equipped with a structure of triangles and a shift, subject to certain conditions which can be found in [Neeman 2001; Weibel 1994, Chapter 10], and [Stevenson 2018] for tensor-triangulated categories.

An object A of a triangulated category $\mathcal{T}$ is *compact* if the hom-functor $(A, -)$ commutes with direct sums and $\mathcal{T}$ is said to be *compactly generated* if there is, up to isomorphism, just a set of compact objects in $\mathcal{T}$ and if the compact objects of $\mathcal{T}$ see every object in the sense that, if $X \in \mathcal{T}$ and if $(A, X) = 0$ for every compact object A in $\mathcal{T}$, then $X = 0$. The restriction that $\mathcal{T}$ be compactly generated could be weakened to $\mathcal{T}$ being well-generated but, in that case, model theory using infinitary languages would be needed, so we would lose the compactness theorem of model theory and its many consequences. This is an interesting direction to follow and a start has been made, see [Krause and Letz 2023] for instance, but here we don't look any further in that direction (also cf. [Adámek and Rosický 1994, §5B]).

Let $\mathcal{T}^{\mathrm{c}}$ denote the full subcategory of compact objects of $\mathcal{T}$. Model theory for the objects of $\mathcal{T}$ is based on the key idea that the *elements* of objects of $\mathcal{T}$ are the morphisms from compact objects. That is, if X is an object of $\mathcal{T}$ and A is a compact object of $\mathcal{T}$, then an *element* of X *of sort* (indexed by) A is a morphism $A \to X$ in $\mathcal{T}$, that is, the value of the functor $(-, X) : (\mathcal{T}^{\mathrm{c}})^{\mathrm{op}} \to \mathbf{Ab}$ on A, where $\mathbf{Ab}$ denotes the category of abelian groups. This is just an extension of the fact that, if M is a (right) module over a (normal, 1-sorted) ring R, then the elements of M may be identified with the morphisms from the module R_R to M.

There is, up to isomorphism, just a set of compact objects, so we may use the objects in a small version of $\mathcal{T}^{\mathrm{c}}$ to index the sorts of the language for $\mathcal{T}$. A "small version" of $\mathcal{T}^{\mathrm{c}}$ means an equivalent category which has just a set of objects. We don't go into detail about setting up the language — for that see [Prest 2009, Appendix B] or various other background references on the model theory of modules, for instance [Prest 2019, §5; 2011a, Chapter 18] — because all we really need is that it gives us a way of writing down formulas, in particular (in our context) pp formulas. Each formula defines, for every $X \in \mathcal{T}$, a certain subset of $(A_1, X) \oplus \cdots \oplus (A_n, X)$ with $A_i \in \mathcal{T}^{\mathrm{c}}$ (the A_i label the sorts of the free variables of the formula).

Of course, for every object $X \in \mathcal{T}$, each sort (A, X), for $A \in \mathcal{T}^{\mathrm{c}}$, has an abelian group structure, and this is built into the formal language. Also built into the language is the action of (a small version of) $\mathcal{T}^{\mathrm{c}}$ on objects $X \in \mathcal{T}$ — the morphisms of $\mathcal{T}^{\mathrm{c}}$ "multiply" the "elements" of X, taking those of one sort to a possibly different sort. Explicitly, if $f : A \to B$ is a morphism of $\mathcal{T}^{\mathrm{c}}$, then this induces $b \in (B, X) \mapsto bf \in (A, X)$ — multiplication by f from sort B to sort A. Note how this generalises the action of a ring on a (1-sorted) right module. In particular, each sort (A, X) is a right module over $\mathrm{End}(A)$ but these multiplications on single sorts are only some of the multiplications that constitute the action of (the many-sorted ring) $\mathcal{T}^{\mathrm{c}}$ on objects X of $\mathcal{T}$.

In this way an object X of $\mathcal{T}$ is replaced by a (many-sorted) set-with-structure, precisely by the right $\mathcal{T}^{\mathrm{c}}$-module which is the representable functor $(-, X)$ restricted to $\mathcal{T}^{\mathrm{c}}$. This replacement is effected by the restricted Yoneda functor $y : \mathcal{T} \to \text{Mod-}\mathcal{T}^{\mathrm{c}}$ which is given on objects by $X \to (-, X) \restriction \mathcal{T}^{\mathrm{c}}$ and on morphisms $f : X \to Y$ by $f \mapsto (-, f) : (-, X) \to (-, Y)$. This functor is neither full nor faithful but, see Propositions 1.3 and 1.4 below, it loses nothing of the model theory[1] so we may do model theory directly in $\mathcal{T}$ or, equivalently, we may move to the functor/module category Mod-$\mathcal{T}^{\mathrm{c}}$, where the well-worked-out model theory of multisorted modules applies. Sometimes it is more convenient to work in the one category than the other; in any case, moving from the one context to the other is straightforward (and is detailed in this paper).

The move to Mod-$\mathcal{T}^{\mathrm{c}}$ gives us the immediate conclusion that the theory of $\mathcal{T}$ has pp-elimination of quantifiers.

Theorem 1.1. *If $\mathcal{T}$ is a compactly generated triangulated category, then every formula in the language for $\mathcal{T}$ is equivalent to the conjunction of a sentence* (*which refers to sizes of quotients of pp-definable subgroups*) *and a finite boolean combination of pp formulas.*

A pp formula (in our context) is an existentially quantified system of linear equations. A system of R-linear equations over a possibly multisorted ring R can be written in the form

$$\bigwedge_{j=1}^{m} \sum_{i=1}^{n} x_i r_{ij} = 0_j$$

(read the conjunction symbol $\bigwedge$ as "and") or, more compactly, as $\bar{x}G = 0$, where $G = (r_{ij})_{ij}$ is a matrix over R. Here x_i is a variable of sort i and r_{ij} a morphism from sort j to sort i (we are dealing with right modules, hence the contravariance). If we denote this (quantifierfree) formula as $\theta(\bar{x})$, that is, $\theta(x_1, \ldots, x_n)$, then its solution set in a module M is denoted $\theta(M)$ and is a subgroup of $M_1 \oplus \cdots \oplus M_n$, where M_i is the group of elements of M of sort i, that is, $(-, \bullet_i)(M) \simeq M(\bullet_i)$, where $\bullet_i$ is the object of R corresponding to sort i.

[1]That is because we use finitary model theory; infinitary languages would detect more, including some *phantom* morphisms, that is, morphisms f with $yf = 0$.

A projection of the solution set for such a system of equations is defined by a formula of the form

$$\exists x_{k+1}, \dots, x_n \Big(\bigwedge_{j=1}^{m} \sum_{i=1}^{n} x_i r_{ij} = 0_j \Big).$$

A formula (equivalent to one) of this form is a *pp* (for "positive primitive") *formula* (the term *regular formula* also is used). We can write a pp formula more compactly as $\exists \bar{y}\ (\bar{x}\ \ \bar{y})G = 0$, or $\exists \bar{y}\ (\bar{x}\ \ \bar{y})\binom{G'}{G''} = 0$, equivalently $\exists \bar{y}\ \bar{x}G' = \bar{y}G''$, if we want to partition the matrix G. If we denote this formula by $\phi(x_1, \dots, x_k)$ then its *solution set* $\phi(M)$ in M is the subgroup of $M_1 \oplus \cdots \oplus M_k$ obtained by projecting $\theta(M)$ to the first k components. We refer to such a solution set as a *pp-definable subgroup* of M (the terminologies "subgroup of finite definition" and "finitely matrizable subgroup" also have been used).

Example 1.2. Consider the quiver A_4 with orientation shown $1 \xrightarrow{\alpha} 2 \xleftarrow{\beta} 3 \xrightarrow{\gamma} 4$ and let $R = KA_4$ be its path algebra with coefficients from a field K. So left R-modules, equivalently K-representations of A_4 have the shape $V_1 \xrightarrow{T_\alpha} V_2 \xleftarrow{T_\beta} V_3 \xrightarrow{T_\gamma} V_4$ where the V_i are K-vector spaces and $T_\alpha, T_\beta, T_\gamma$ are K-linear maps. In order to illustrate the definitions above, we think of these structures as *right* modules over the *opposite* of the 4-object K-linear path category of A_4, that is, over the K-linear category which has objects $\bullet_i$, $i = 1, 2, 3, 4$, and with $\mathrm{End}(\bullet_i) = K \cdot 1_i$, $(\bullet_2, \bullet_1) = K\alpha$, $(\bullet_2, \bullet_3) = K\beta$, $(\bullet_4, \bullet_3) = K\gamma$ and all other morphism groups 0.

The corresponding language has four sorts, and the function symbols are, apart from the additions in each sort, the λf where $\lambda \in K$ and f is one of the identity maps or α, β or γ. An example of a system of linear equations is

$$x_2 - x_1\alpha - x_3\beta = 0_2, \quad x_3\gamma = 0_3,$$

where sorts are shown by subscripts to variables and zeroes. Note that all terms in a given equation must have the same sort.

We may quantify out the variables x_1 and x_3 to obtain the pp formula $\phi(x_2)$ which is

$$\exists x_1, x_3\ (x_2 - x_1\alpha - x_3\beta = 0_2 \wedge x_3\gamma = 0_3)$$

which, in matrix format, is

$$\exists x_1, x_3\ (x_1\ \ x_2\ \ x_3)\begin{pmatrix} -\alpha & 0 \\ 1 & 0 \\ -\beta & \gamma \end{pmatrix} = (0\ \ 0).$$

The solution set $\phi(M)$ in any module M is the set $\alpha(M) + \beta(\ker_M(\gamma))$ — a K vector subspace of $M\bullet_2$ ($= V_2$ in the representation-of-quivers notation).

All this applies to $\mathcal{T}$ since the model theory of $\mathcal{T}$ is essentially that of right $\mathcal{T}^{c}$-modules. So Theorem 1.1 follows because, if R is a (possibly many sorted) ring, then the theory of R-modules has pp-elimination of quantifiers[2] and so this applies to the theory of the image of the restricted Yoneda embedding which, as we have remarked, is the theory of $\mathcal{T}$.

It turns out, see [Garkusha and Prest 2005, 3.1, 3.2] and Section 2B, that, with this language, the theory of $\mathcal{T}$ has complete (positive) elimination of quantifiers — every (pp) formula is equivalent to a quantifier-free (pp) formula (see Theorem 2.10). There is also a dual form of this — every pp formula is equivalent to a divisibility formula (Lemma 2.9). We will also see in Section 2D that the theory of $\mathcal{T}$ has elimination of pp-imaginaries — every pp-pair is definably isomorphic to a (quantifier-free) formula.

As with any theory of modules, the initial category of sorts, in this case a small version of $(\mathcal{T}^c)^{\mathrm{op}}$, may be completed to the full category $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ of pp-definable sorts: the objects are pp-pairs and the morphisms are the pp-definable maps between these pairs (see Section 2A). In our context, this completed category of sorts has two manifestations. One is the category of coherent functors [Krause 2002] on $\mathcal{T}$. The other is a certain localisation of the category $(\text{mod-}\mathcal{T}^c, \mathbf{Ab})^{\mathrm{fp}}$ of finitely presented functors from mod-$\mathcal{T}^c$ — the category of finitely presented right $\mathcal{T}^c$-modules — to the category $\mathbf{Ab}$ of abelian groups. In fact, [Prest 2012b, 7.1, 7.2], this localisation turns out to be equivalent to the opposite of mod-$\mathcal{T}^c$ which is, in turn, equivalent to $\mathcal{T}^c$-mod. The latter equivalence, Corollary 2.4, reflects the fact that the absolutely pure = fp-injective $\mathcal{T}^c$-modules coincide with the flat $\mathcal{T}^c$-modules. We will, in Section 2A, give details of this, as well as the action of each of these manifestations of $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ on $\mathcal{T}$, respectively on $y\mathcal{T}$.

Free realisations and pp-types are used a lot in the model theory of modules and applications, so in Section 2C we point out how these look in $\mathcal{T}$.

In Section 3A we present the various types of data which can specify a definable subcategory of $\mathcal{T}$. In Section 3B we see the bijection between definable subcategories of $\mathcal{T}$ and hereditary torsion theories of finite type on Mod-$\mathcal{T}^c$ and in Section 3C we explore that connection in more detail. The category of imaginaries of a definable subcategory is described in Section 3D. Some connections between hom-orthogonal pairs in $\mathcal{T}$ and hereditary torsion theories on Mod-$\mathcal{T}^c$ are seen in Section 3E and this is continued in Section 3G with the bijection between triangulated definable subcategories and smashing subcategories of $\mathcal{T}$.

Section 3F describes spectra associated to $\mathcal{T}$ and this is continued for tensor-triangulated categories in Section 4A.

[2]For the formal statement see, for instance, [Prest 2009, A.1.1]. That is given for 1-sorted modules but the general case reduces to this, see [Kucera and Prest 1992, §1], because each formula involves only finitely many sorts, corresponding to $A_1, \ldots, A_n$ say, so is equivalent to a formula over a 1-sorted ring, namely $\mathrm{End}(A_1 \oplus \cdots \oplus A_n)$.

For definable subcategories of module categories there is a duality, elementary duality, which exists at a number of levels, in particular between definable subcategories of Mod-R and R-Mod. This carries over, at least to algebraic triangulated categories; we outline that in Section 3H. If $\mathcal{T}$ is tensor-triangulated with $\mathcal{T}^c$ rigid, then there is also an *internal* duality, induced by the duality on $\mathcal{T}^c$; that is described in Section 4B.

Tensor-closed definable subcategories are briefly considered in Section 4, and in Section 4C there is some exploration of the wider possibilities for interpreting the model-theoretic language.

Background on the model theory of modules can be found in various references; we use [Prest 2009] as a convenient compendium of results and references to the original papers. We give a few reminders in this paper. The approach in [Prest 2009] is algebraic/functor-category-theoretic; readers coming from model theory might find [Prest 1988b] or [Prest $\geq$ 2024] a more approachable introduction. For model theory of modules over many-sorted rings, see [Prest 2019].

Thanks to Isaac Bird and Jordan Williamson for a number of useful comments and for sharing their preprint [Bird and Williamson 2022].

1B. ***The restricted Yoneda functor.*** The restricted Yoneda functor $y : \mathcal{T} \to \text{Mod-}\mathcal{T}^c$, $X \to (-, X) \restriction \mathcal{T}^c$ underlies most of what we do here. Restricting its domain to the category $\mathcal{T}^c$ of compact objects gives, by the Yoneda lemma and because $\mathcal{T}$ is idempotent-complete (see [Neeman 2001, 1.6.8]), an equivalence between $\mathcal{T}^c$ and the category proj-$\mathcal{T}^c$ of finitely generated projective right $\mathcal{T}^c$-modules. The functor y is, however, neither full nor faithful and one effect of this is that the image of $\mathcal{T}$ in Mod-$\mathcal{T}^c$ is not closed under elementary equivalence, indeed it is not a definable subcategory (see Section 3A) of Mod-$\mathcal{T}^c$. We do, however, have Propositions 1.3 and 1.4 below (the second is just by the Yoneda lemma).

First we recall (see [Prest 2009, §2.1.1]) that an embedding $M \to N$ of objects in a module category, more generally in a definable additive category, is *pure* if, for every pp formula ϕ, the (image of the) solution set $\phi(M)$ is the intersection of $\phi(N)$ "with M", meaning with the product of sorts of M corresponding to the free variables of ϕ. And M is *pure-injective* if every pure embedding with domain M is split. There are many equivalent definitions; see [Prest 2009, §§4.3.1, 4.3.2].

The theory of purity — intimately connected with solution sets of pp formulas and so with the model theory of additive structures — was developed, in algebraic terms, in compactly generated triangulated categories in [Beligiannis 2000b; Krause 2000]. Essentially, it is the theory of purity in Mod-$\mathcal{T}^c$, more precisely, in the definable subcategory generated by $y\mathcal{T}$, pulled back to $\mathcal{T}$. For example, $X \in \mathcal{T}$ is pure-injective if and only if yX is a pure-injective $\mathcal{T}^c$-module. Since yX is absolutely pure [Krause 2000, Lemma 1.6], that is equivalent to it being an injective

$\mathcal{T}^c$-module. The pure-injective objects of $\mathcal{T}$ play the same key role that they do in the model theory of modules. For instance every ($\varnothing$-)saturated module is pure-injective and the pure-injective modules are exactly the direct summands of saturated modules (see [Prest 2011a, Proposition 21.1, Theorem 21.2] or [Prest 1988b, 2.9]); this is equally true in compactly generated triangulated categories.[3]

Proposition 1.3 [Krause 2000, 1.8]. *If $X \in \mathcal{T}$ is pure-injective then, for every $Y \in \mathcal{T}$, the restricted Yoneda map $y : (Y, X) \to (yY, yX)$ is bijective.*

Proposition 1.4. *If $A \in \mathcal{T}$ is a compact object then, for every $X \in \mathcal{T}$, the restricted Yoneda map $y : (A, X) \to (yA, yX)$ is bijective.*

In fact there is symmetry here in that Proposition 1.4 holds more generally for A pure-projective (that is, a direct summand of a direct sum of compact objects).

We will use the fact that the restricted Yoneda functor induces an equivalence between the category $\mathrm{Pinj}(\mathcal{T})$ of pure-injective objects in $\mathcal{T}$ and the category $\text{Inj-}\mathcal{T}^c$ of injective right $\mathcal{T}^c$-modules.

Theorem 1.5 [Krause 2000, 1.9]. *The restricted Yoneda functor $y : \mathcal{T} \to \text{Mod-}\mathcal{T}^c$ induces an equivalence*

$$\mathrm{Pinj}(\mathcal{T}) \simeq \text{Inj-}\mathcal{T}^c.$$

1C. *Definable subcategories of module categories.* Very briefly, we recall the context of the model theory of modules and the principal associated structures. Some of this is defined more carefully later in the paper but see the references for more detail.

In model theory in general, the context is typically the category of models of some complete theory, with elementary embeddings. In the context of modules, it turns out to be more natural to work with *definable subcategories*, meaning full subcategories of module categories which are closed under elementary equivalence and which are *additive*, meaning closed under direct sums and direct summands. These subcategories are equivalently characterised, without reference to model theory, as follows (see [Prest 2009, §3.4] for this and various other characterisations by closure conditions).

Theorem 1.6. *A subcategory $\mathcal{D}$ of a module category is a definable subcategory if and only if $\mathcal{D}$ is closed under direct products, directed colimits and pure submodules.*

If $\mathcal{X}$ is a set of modules, then we denote by $\langle \mathcal{X} \rangle$ the *definable subcategory generated by $\mathcal{X}$*. It is the closure of $\mathcal{X}$ under the above operations, equally it is the smallest additive subcategory containing $\mathcal{X}$ and closed under elementary equivalence.

[3]This comment, like a few others, is particularly directed to those coming from model theory.

It is the case, see [Prest 2009, 3.4.8], that every definable subcategory is closed under pure-injective hulls where, if M is a module, its *pure-injective hull* $H(M)$ is a minimal pure, pure-injective extension of M.[4] It follows that every definable subcategory is determined by the pure-injective modules in it. If $\mathcal{T}$ is a compactly generated triangulated category and $X \in \mathcal{T}$, then the *pure-injective hull* of X may be defined to be the (unique-to-isomorphism over X, by Theorem 1.5) object $H(X)$ of $\mathcal{T}$ such that $yH(X) = E(yX)$, where E denotes injective hull in the module category Mod-$\mathcal{T}^{\mathrm{c}}$.

To each *definable category* $\mathcal{D}$ — meaning a category equivalent to a definable subcategory of a module category — there is associated a skeletally small abelian category, fun($\mathcal{D}$), of functors on $\mathcal{D}$. This can be defined as the category of pp-imaginaries (see Section 2A) for $\mathcal{D}$, or as a localisation of the free abelian category on R where $\mathcal{D}$ is a definable subcategory of Mod-R (R a possibly many-sorted ring), or as the category of *coherent functors* — those that commute with direct products and directed colimits — from $\mathcal{D}$ to **Ab**. Each definable subcategory[5] $\mathcal{C}$ of $\mathcal{D}$ is determined by the Serre subcategory $\mathcal{S}_{\mathcal{C}}$ of fun($\mathcal{D}$) which consists of those functors which are 0 on $\mathcal{C}$, and then fun($\mathcal{C}$) is the (abelian) quotient category fun($\mathcal{D}$)$/\mathcal{S}_{\mathcal{C}}$ — the Serre localisation (see [Krause 2022, p. 30ff.]) of fun($\mathcal{D}$) at $\mathcal{S}_{\mathcal{C}}$.

Also associated to a definable category $\mathcal{D}$ is its *Ziegler spectrum* Zg($\mathcal{D}$) ([Ziegler 1984], see [Prest 2009, Chapter 5]) — a topological space whose points are the isomorphism classes of indecomposable pure-injective objects in $\mathcal{D}$ and whose open subsets are the complements of zero-sets of sets of coherent functors on $\mathcal{D}$. The closed subsets of Zg($\mathcal{D}$) are in natural bijection with the definable subcategories of $\mathcal{D}$; see [Prest 2009, 5.1.6]. See Section 3F for more on this.

2. Model theory in compactly generated triangulated categories

We use formulas to specify the definable subsets of objects of $\mathcal{T}$. In order to set these up, we choose a sub*set* $\mathcal{G}$ of $\mathcal{T}^{\mathrm{c}}$ which we will assume to be generating in the sense that, if $X \in \mathcal{T}$, then $(G, X) = 0$ for every $G \in \mathcal{G}$ implies $X = 0$, and we take the (opposite of the) full subcategory on $\mathcal{G}$ to be the category of sorts. For convenience, we will assume that $\mathcal{G}$ is equivalent to $\mathcal{T}^{\mathrm{c}}$, that is, contains at least one isomorphic copy of each compact object of $\mathcal{T}$. By $\mathcal{L}_{\mathcal{G}}$ we denote the resulting language, meaning the resulting set of formulas.

We could take a smaller category of sorts, for instance, if $\mathcal{T}$ is monogenic, generated by a single compact object S, then we could consider the 1-sorted language based on S. The obvious question is whether this would suffice, in the sense that every set definable in the larger language would also be definable in the

[4] In fact, M is an elementary submodule of $H(M)$, [Sabbagh 1970, corollaire 4 de théorème 4].

[5] The containing module category in Theorem 1.6 may be replaced by any definable category.

1-sorted language. We don't pursue this here, but the relative approach and results in [Garkusha and Prest 2004; 2005] should be helpful in answering this question.

In the other direction, we could make the maximal choice of sorts and use a language with the category $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ of pp-imaginaries (see Section 2A) for the sorts. Since pp-imaginaries are already definable, this does not increase the collection of definable subsets. For most purposes the choice of category of sorts does not matter provided the definable subsets are the same. However, elimination of quantifiers and elimination of imaginaries are language-dependent, rather than structure-dependent. Our choice of $\mathcal{G}$ as (essentially) $\mathcal{T}^{\mathrm{c}}$ is exactly analogous to basing a language for the model theory of R-modules (R a 1-sorted ring) on the category mod-R of finitely presented modules, rather than using the 1-sorted language based on the single module R_R (see [Prest 2011b] for more on choices of languages for additive categories).

Having chosen $\mathcal{G}$ we introduce a sort s_A for each $A \in \mathcal{G}$ and a symbol for addition (and a symbol for the zero) on each sort and, for each $f : A \to B$ in $\mathcal{G}$, a corresponding function symbol from sort B to sort A to represent multiplication by f (= composition with f). Note that the morphisms of $\mathcal{G}$ are the "elements of the ring-with-many-objects $\mathcal{G}$".

Each object $X \in \mathcal{T}$ then becomes a structure for this language by taking its elements of sort s_A to be the elements of (A, X) and then interpreting the function symbols in the usual/obvious way.

Remark 2.1. If $\mathcal{T}$ is tensor-triangulated and has an internal hom functor right adjoint to $\otimes$, then these sorts, which by definition are abelian groups, can be taken instead to be objects of $\mathcal{T}$, in the sense that we could interpret the sort $s_A(X)$ to be the internal hom object $[A, X] \in \mathcal{T}$. In this "internal" interpretation of the language, we have, since $(A, X) \simeq (\mathbb{1}, [A, X])$ where $\mathbb{1}$ is the tensor-unit, the (usual) elements of X of sort A identified with the morphisms $\mathbb{1} \to [A, X]$.

We will write $\mathcal{L}(\mathcal{T})$, or just $\mathcal{L}$ for the language. Since we assume that $\mathcal{G}$ is equivalent to $\mathcal{T}^{\mathrm{c}}$, the $\mathcal{L}(\mathcal{T})$-structure $X \in \mathcal{T}$, which is literally a right $\mathcal{G}$-module, may be identified with the image, $yX = (-, X) \upharpoonright \mathcal{T}^{\mathrm{c}}$, of X under the restricted Yoneda functor $y : \mathcal{T} \to \text{Mod-}\mathcal{T}^{\mathrm{c}}$. Therefore the model theory of X as an object of $\mathcal{T}$ is exactly that of yX as a right $\mathcal{T}^{\mathrm{c}}$-module. Indeed, $\mathcal{L}(\mathcal{T})$ is equally a language for $\mathcal{T}$ and for the module category Mod-$\mathcal{T}^{\mathrm{c}}$, but bear in mind that there are more $\mathcal{T}^{\mathrm{c}}$-modules than those which are in the image of $\mathcal{T}$ in Mod-$\mathcal{T}^{\mathrm{c}}$, more even than in the definable subcategory of Mod-$\mathcal{T}^{\mathrm{c}}$ which is generated by that image.

Indeed, the definable subcategory, $\langle y\mathcal{T} \rangle$, of Mod-$\mathcal{T}^{\mathrm{c}}$ generated by the image of $\mathcal{T}$ is exactly the subcategory, Flat-$\mathcal{T}^{\mathrm{c}}$ = Abs-$\mathcal{T}^{\mathrm{c}}$, consisting of the flat (= absolutely pure[6]) $\mathcal{T}^{\mathrm{c}}$-modules.

[6] In "most" module categories the flat and absolutely pure modules have little overlap; the fact that they are equal over the ring $\mathcal{T}^{\mathrm{c}}$ is a very characteristic feature here.

Theorem 2.2 [Beligiannis 2000a, 8.11, 8.12; Krause 2000, 2.7]. *If $\mathcal{T}$ is a compactly generated triangulated category and $y : \mathcal{T} \to \text{Mod-}\mathcal{T}^{\text{c}}$ is the restricted Yoneda functor, then $\langle y\mathcal{T}\rangle = \text{Abs-}\mathcal{T}^{\text{c}} = \text{Flat-}\mathcal{T}^{\text{c}}$*

Therefore the model theory of $\mathcal{T}$ is the same as the model theory of the flat (= absolutely pure) right $\mathcal{T}^{\text{c}}$-modules.[7] The one difference is that some structures are missing from $\mathcal{T}$: except in the case that $\mathcal{T}$ is pure semisimple [Beligiannis 2000b, Theorem 9.3], there are structures in $\langle y\mathcal{T}\rangle$ which are not in $y\mathcal{T}$. However, the equivalence, Theorem 1.5, of categories $\text{Pinj}(\mathcal{T}) \simeq \text{Inj-}\mathcal{T}^{\text{c}}$ between the pure-injective objects of $\mathcal{T}$ and the injective $\mathcal{T}^{\text{c}}$-modules, implies that $y\mathcal{T}$ does contain all the pure-injective models, in particular all the saturated models, of its theory. It follows from Theorem 2.2 that implications and equivalences of pp-formulas on $\mathcal{T}$ and on $\text{Flat-}\mathcal{T}^{\text{c}} = \text{Abs-}\mathcal{T}^{\text{c}}$ are the same.

For convenience we will sometimes write $(-, X)$ instead of $(-, X) \restriction \mathcal{T}^{\text{c}} = yX$ when $X \in \mathcal{T}$.

2A. ***The category of pp-sorts.*** Let R be a, possibly multisorted, ring and let $\mathcal{D}$ be a definable subcategory of Mod-R. We recall how to define the *category* $\mathbb{L}(\mathcal{D})^{\text{eq}+}$ *of pp sorts* (or *pp-imaginaries*) for $\mathcal{D}$.

First, for $\mathcal{D} = \text{Mod-}R$, the category $\mathbb{L}(\text{Mod-}R)^{\text{eq}+}$, more briefly denoted $\mathbb{L}_R^{\text{eq}+}$, has, for its objects, the *pp-pairs* ϕ/ψ, that is pairs (ϕ, ψ) of pp formulas for R-modules with $\phi \geq \psi$, meaning $\phi(M) \geq \psi(M)$ for all $M \in \text{Mod-}R$. For its arrows, we take the pp-definable maps between these pairs. See [Herzog 1993, §1] or [Prest 2009, §3.2.2] for details and the fact that this category is abelian. Each such pp-pair defines a coherent functor $M \mapsto \phi(M)/\psi(M)$ from Mod-R to **Ab** and every coherent functor has this form, see, for instance, [Prest 2009, §10.2].

For general $\mathcal{D}$, a definable subcategory of Mod-R, we let $\Phi_{\mathcal{D}}$ be the Serre subcategory of $\mathbb{L}_R^{\text{eq}+}$ consisting of those pp-pairs ϕ/ψ which are *closed on*, that is 0 on, every $M \in \mathcal{D}$ (that is, $\phi(M) = \psi(M)$ for every $M \in \mathcal{D}$). Then $\mathbb{L}(\mathcal{D})^{\text{eq}+}$ is defined to be the quotient = Serre-localisation $\mathbb{L}_R^{\text{eq}+}/\Phi_{\mathcal{D}}$. So $\mathbb{L}(\mathcal{D})^{\text{eq}+}$ has the same objects as $\mathbb{L}_R^{\text{eq}+}$ — the pp-pairs — and the morphisms in $\mathbb{L}(\mathcal{D})^{\text{eq}+}$ are given by pp formulas which on every $M \in \mathcal{D}$ define a function. In particular the pp-pairs closed on $\mathcal{D}$ are isomorphic to 0 in $\mathbb{L}(\mathcal{D})^{\text{eq}+}$. The localised category $\mathbb{L}(\mathcal{D})^{\text{eq}+}$ also is abelian; in fact, see [Prest and Rajani 2010, 2.3], every skeletally small abelian category arises in this way.

An equivalent [Prest 2011a, 12.10], but less explicit, definition is that $\mathbb{L}(\mathcal{D})^{\text{eq}+} = (\mathcal{D}, \mathbf{Ab})^{\prod\to}$ — the category of functors[8] from $\mathcal{D}$ to **Ab** which commute with direct

[7] $\mathcal{T}^{\text{c}}$ is both right and left coherent as a ring with many objects (see [Oberst and Röhrl 1970, §4]), which is why the flat and the absolutely pure objects form definable subcategories (see [Prest 2009, Theorem 3.4.24]).

[8] additive, as always assumed in this paper

products and directed colimits (that is, coherent functors, equivalently [Prest 2011a, 25.3] interpretation functors in the model-theoretic sense).

It is well-known, see [Prest 2009, 10.2.37, 10.2.30], and much-used, that, for $\mathcal{D} = \text{Mod-}R$, the category of pp-pairs is equivalent to the free abelian category on R^{op} and, also, that it can be realised as the category $(\text{mod-}R, \mathbf{Ab})^{\text{fp}}$ of finitely presented functors on finitely presented modules (see [Prest 2009, 10.2.30, 10.2.37]) equivalently, as just said, it is equivalent to the category of coherent functors on all modules (see [Prest 2009, §10.2.8]). Then, for a general definable subcategory $\mathcal{D}$ of Mod-R, we obtain $\mathbb{L}(\mathcal{D})^{\text{eq}+}$ as the Serre-quotient $(\text{mod-}R, \mathbf{Ab})^{\text{fp}}/\mathcal{S}_{\mathcal{D}}$ where $\mathcal{S}_{\mathcal{D}}$ is the Serre subcategory of those functors $F \in (\text{mod-}R, \mathbf{Ab})^{\text{fp}}$ with $\vec{F}\mathcal{D} = 0$. Here $\vec{F}$ is the unique extension of (a finitely presented) $F : \text{mod-}R \to \mathbf{Ab}$ to a (coherent) functor from Mod-R to $\mathbf{Ab}$ which commutes with directed colimits. Often we simplify notation by retaining the notation F for this extension $\vec{F}$.

Under the identification of $\mathbb{L}_R^{\text{eq}+}$ and $(\text{mod-}R, \mathbf{Ab})^{\text{fp}}$ the Serre subcategory $\Phi_{\mathcal{D}}$ is identified with $\mathcal{S}_{\mathcal{D}}$.

In applying this in our context, we use the following result, where Flat-R denotes the category of flat right R-modules and Abs-R denotes the category of absolutely pure (= fp-injective) right R-modules. For the notion of a left coherent multisorted ring — one whose category of left modules is locally coherent — see [Oberst and Röhrl 1970, 4.1].

Theorem 2.3 [Prest 2012b, 7.1/7.2]. *If R is any left coherent (multisorted) ring, then* Flat-R *is a definable subcategory of* Mod-R *and*

$$\mathbb{L}(\text{Flat-}R)^{\text{eq}+} \simeq R\text{-mod}.$$

If R is a right coherent ring, then Abs-R *is a definable subcategory of* Mod-R *and*

$$\mathbb{L}(\text{Abs-}R)^{\text{eq}+} \simeq (\text{mod-}R)^{\text{op}}.$$

Because $\mathcal{T}^{\text{c}}$ is right and left coherent, [Beligiannis 2000a, 8.11, 8.12], and since Abs-$\mathcal{T}^{\text{c}}$ = Flat-$\mathcal{T}^{\text{c}}$, we have the following corollary.

Corollary 2.4. *If $\mathcal{T}$ is a compactly generated triangulated category, then there is an equivalence*

$$d : \mathcal{T}^{\text{c}}\text{-mod} \simeq (\text{mod-}\mathcal{T}^{\text{c}})^{\text{op}}$$

and this category is equivalent to the category $\mathbb{L}(\mathcal{T})^{\text{eq}+}$ of pp-imaginaries for $\mathcal{T}$.

We write d for the (anti)equivalence in each direction.

There is another description of the category appearing in Corollary 2.4. We say that a *coherent* functor on $\mathcal{T}$ is one which is the cokernel of a map between representable functors $(A, -) : \mathcal{T} \to \mathbf{Ab}$ with $A \in \mathcal{T}^{\mathrm{c}}$. Explicitly, if $f : A \to B$ is in $\mathcal{T}^{\mathrm{c}}$ then we obtain an exact sequence of functors on $\mathcal{T}$:

$$(B, -) \xrightarrow{(f,-)} (A, -) \to F_f \to 0;$$

and the cokernel F_f is a typical coherent functor on $\mathcal{T}$.

In module categories having a presentation of this form, with A and B finitely presented, is equivalent to commuting with products and directed colimits but triangulated categories don't have directed colimits. There is the following analogous characterisation.

Theorem 2.5 [Krause 2002, 5.1]. *Suppose that $\mathcal{T}$ is a compactly generated triangulated category. Then $F : \mathcal{T} \to \mathbf{Ab}$ is a coherent functor if and only if F commutes with products and sends homology colimits to colimits.*

We denote the category of coherent functors on $\mathcal{T}$, with the natural transformations between them, by $\mathrm{Coh}(\mathcal{T})$. This category is abelian; in fact we have the following.

Theorem 2.6 [Krause 2002, 7.2]. *There is a duality*

$$(\text{mod-}\mathcal{T}^{\mathrm{c}})^{\mathrm{op}} \simeq \mathrm{Coh}(\mathcal{T})$$

and hence

$$\mathrm{Coh}(\mathcal{T}) \simeq \mathcal{T}^{\mathrm{c}}\text{-mod}.$$

Indeed, to go from $\mathrm{Coh}(\mathcal{T})$ to $\mathcal{T}^{\mathrm{c}}$-mod we just restrict the action of $F \in \mathrm{Coh}(\mathcal{T})$ to $\mathcal{T}^{\mathrm{c}}$ and, in the other direction, we apply the projective presentation $(B, -) \to (A, -) \to G \to 0$ of a finitely presented left $\mathcal{T}^{\mathrm{c}}$-module in $\mathcal{T}$ and we get a coherent functor. The category $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ of pp-definable sorts and pp-definable maps for $\mathcal{T}$ is defined just as for a module category. Since the model theories of $\mathcal{T}$ and Flat-$\mathcal{T}^{\mathrm{c}}$ are identical, it is exactly $\mathbb{L}(\text{Flat-}\mathcal{T}^{\mathrm{c}})^{\mathrm{eq}+}$.

Corollary 2.7. *The category of pp-imaginaries for a compactly generated triangulated category $\mathcal{T}$ can be realised in the forms*

$$\mathbb{L}(\mathcal{T})^{\mathrm{eq}+} \simeq \mathrm{Coh}(\mathcal{T}) \simeq \mathcal{T}^{\mathrm{c}}\text{-mod}.$$

The duality in Theorem 2.6 respects the actions of those categories of functors on $\mathcal{T}$. We give the details.

The action of $\mathrm{Coh}(\mathcal{T})$ on $\mathcal{T}$ is given by the exact sequence above presenting F_f: if $X \in \mathcal{T}$, then $F_f(X)$ is defined by exactness of the sequence

$$(B, X) \to (A, X) \to F_f(X) \to 0.$$

The action of mod-$\mathcal{T}^c$ on $\mathcal{T}$ is given by Hom applied after the restricted Yoneda functor y. Explicitly: the typical finitely presented right $\mathcal{T}^c$-module G_f is given by an exact sequence (a projective presentation)

$$yA \xrightarrow{yf} yB \to G_f \to 0,$$

that is,

$$(-, A) \xrightarrow{(-,f)} (-, B) \to G_f \to 0,$$

where $A \xrightarrow{f} B \in \mathcal{T}^c$. The action of G_f on $X \in \mathcal{T}$ is induced by the action of $(-, yX)$ on it: we have the exact sequence

$$0 \to (G_f, (-, X)) \to ((-, B), (-, X)) \xrightarrow{((-,f),(-,X))} ((-, A), (-, X)),$$

that is,

$$0 \to G_f(X) \to (B, X) \xrightarrow{(f,X)} (A, X),$$

defining the value of G_f on the typical object $X \in \mathcal{T}$. So, if $G \in$ mod-$\mathcal{T}^c$ and $X \in \mathcal{T}$, then the action of G on X is defined by

$$G(X) = (G, yX).$$

Notice that the morphism $f : A \to B$ in $\mathcal{T}^c$ has given rise to the exact sequence of abelian groups:

$$0 \to G_f(X) \to (B, X) \xrightarrow{(f,X)} (A, X) \to F_f(X) \to 0. \tag{1}$$

The duality $(\text{mod-}\mathcal{T}^c)^{\text{op}} \simeq \text{Coh}(\mathcal{T})$ in Theorem 2.6 takes a finitely presented right $\mathcal{T}^c$-module G to the coherent functor

$$G^\circ : X \mapsto (G, yX) = (G, (-, X))$$

for $X \in \mathcal{T}$—that is, the action we defined just above. This takes the representable functor $G = (-, A)$ where $A \in \mathcal{T}^c$, to the representable coherent functor $(A, -) : \mathcal{T} \to \mathbf{Ab}$. Therefore, the 4-term exact sequence (1) above can be read as the application of the following exact sequence of functors in $\text{Coh}(\mathcal{T})$ to X:

$$0 \to G_f^\circ \to (B, -) \xrightarrow{(f,-)} (A, -) \to F_f \to 0. \tag{2}$$

In the other direction, the duality $(\text{mod-}\mathcal{T}^c)^{\text{op}} \simeq \text{Coh}(\mathcal{T})$ takes $F \in \text{Coh}(\mathcal{T})$ to the finitely presented $\mathcal{T}^c$-module

$$F^\diamond : C \mapsto (F, (C, -))$$

for $C \in \mathcal{T}^c$. So $(A, -)^\diamond = (-, A)$. If $F = F_f$, then applying $(-, (C, -))$ to the presentation (2) of F_f and using that $(C, -)$ is injective in $\text{Coh}(\mathcal{T})$ (by Theorem 2.6

and since $(-, C)$ is projective in Mod-$\mathcal{T}^{\mathrm{c}}$) allows us to read the resulting 4-term exact sequence as the application, to $C \in \mathcal{T}^{\mathrm{c}}$, of the following exact sequence of functors in mod-$\mathcal{T}^{\mathrm{c}}$:

$$0 \to F_f^{\diamond} \to (-, A) \xrightarrow{(-,f)} (-, B) \to G_f \to 0. \tag{3}$$

Applying the duality-equivalences

$$(-)^{\diamond} : (\mathrm{Coh}(\mathcal{T}))^{\mathrm{op}} \to \text{Mod-}\mathcal{T}^{\mathrm{c}} \quad \text{and} \quad (-)^{\circ} : (\text{Mod-}\mathcal{T}^{\mathrm{c}})^{\mathrm{op}} \to \mathrm{Coh}(\mathcal{T})$$

interchanges (2) — an exact sequence in $\mathrm{Coh}(\mathcal{T})$ — and (3) — an exact sequence in mod-$\mathcal{T}^{\mathrm{c}}$.

The equivalences of these functor categories with the category $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ of pp-pairs for $\mathcal{T}$ are given explicitly on objects as follows. Let $f : A \to B$ be a morphism in $\mathcal{T}^{\mathrm{c}}$, so F_f is a typical coherent functor:

$$\begin{tikzcd} A \arrow[r, "f"] \arrow[d, "x_A"'] & B \arrow[dl, dotted, "y_B"] \\ X & \end{tikzcd}$$

We have that $F_f X = (A, X)/\mathrm{im}(f, X)$ and hence F_f is the functor given by the pp-pair $(x_A = x_A)/(\exists y_B\ x_A = y_B f)$, that is,

$$F_f = (x_A = x_A)/(f | x_A).$$

We use subscripts on variables to show their sorts but might sometimes drop them for readability. We also use variables (which really belong in formulas) to label morphisms (for which they are place-holders) in what we hope is a usefully suggestive way.

Also, from the exact sequence (1), we see that $G_f^{\circ}(-) = \ker(f, -)$ and so is the functor given by the pp-pair

$$G_f^{\circ} = (x_B f = 0)/(x_B = 0).$$

Since the duality $\mathrm{Coh}(\mathcal{T}) \simeq (\text{mod-}\mathcal{T}^{\mathrm{c}})^{\mathrm{op}}$ preserves the actions on $\mathcal{T}$, these pp-pairs also give the actions of, respectively, $F_f^{\diamond}$ and G_f on $\mathcal{T}$.

To go from pp-pairs to functors, we may use Theorem 2.15 below, which says that every pp-pair is isomorphic to one of a form seen above, namely $xf = 0/x = 0$.

2B. ***Elimination of quantifiers.*** If a ring R is right coherent then every pp formula is equivalent on Abs-R to an annihilator formula and, if R is left coherent, then every pp formula on Flat-R is equivalent to a divisibility formula (see [Prest 2009, 2.3.20, 2.3.9, 2.3.19]). These results are equally valid for rings with many objects (because any formula involves only finitely many sorts, so is equivalent to a formula over a ring with one object). It follows that the theory of $\mathcal{T}$ has elimination of quantifiers, indeed it has the stronger property elim-q$^+$, meaning that each pp formula is equivalent to a quantifier-free pp formula, that is, to a conjunction of equations.[9] Also $\mathcal{T}$ has the elementary-dual elimination of pp formulas to divisibility formulas. But it is instructive to see exactly how this works when the ring is the category $\mathcal{T}^{\mathrm{c}}$ of compact objects of a compactly generated triangulated category $\mathcal{T}$. This is an expansion of [Garkusha and Prest 2005, 3.1, 3.2]. We write 0 for any n-tuple $(0, \dots, 0)$.

Given $f : A \to B$ in $\mathcal{T}^{\mathrm{c}}$, form the distinguished triangle[10] as shown:

$$A \xrightarrow{f} B \xrightarrow{g} C \to \Sigma A.$$

Since $\mathcal{T}^{\mathrm{c}}$ is triangulated, $C \in \mathcal{T}^{\mathrm{c}}$. Since representable functors on a triangulated category are exact (meaning that they take triangles to (long) exact sequences), for every $X \in \mathcal{T}$, $(C, X) \xrightarrow{(g,X)} (B, X) \xrightarrow{(f,X)} (A, X)$ is exact so, for $x_B \in (B, X)$, we have $x_B \in \ker(f, X)$ if and only if $x_B \in \mathrm{im}(g, X)$, that is, $x_B f = 0$ if and only if $g \mid x_B$ that is, if and only if $\exists y_C\, (x_B = y_C g)$. Thus

$$x_B f = 0 \iff g \mid x_B.$$

Since $\mathcal{T}^{\mathrm{c}}$ has finite direct sums, tuples of variables may be wrapped up into single variables (we do this explicitly below), so these formulas are general annihilator and divisibility formulas. Therefore every annihilator formula is equivalent to a divisibility formula and *vice versa*. We record this.

Proposition 2.8. *If $A \xrightarrow{f} B \xrightarrow{g} C \to \Sigma A$ is a distinguished triangle, then the formula $x_B f = 0$ is equivalent to $g \mid x_B$.*

Before continuing, note that, because $\mathcal{T}^{\mathrm{c}}$ is closed under finite direct sums, a finite sequence $(x_1, \dots, x_n)$ of variables, with x_i of sort A_i, may be regarded as a single variable of sort $A_1 \oplus \cdots \oplus A_n$. That simplifies notation and allows us to treat a general pp formula as one of the form $\exists x_{B'}\, (x_B f = x_{B'} f')$, that is, $f' | x f$ for short.

$$\begin{tikzcd}
& B \arrow[dr, "x_B"] & \\
A \arrow[ur, "f"] \arrow[dr, "f'"'] & & (-) \\
& B' \arrow[ur, "x_{B'}"'] &
\end{tikzcd}$$

[9] Indeed, since our sorts are closed under finite direct sums, every pp formula is equivalent to a single equation

[10] We will often write "triangle" meaning distinguished triangle.

That is equivalent to

$$\exists x_{B'}\Big((x_B, x_{B'})\big(\begin{smallmatrix} f \\ f' \end{smallmatrix}\big) = 0\Big) \qquad \begin{tikzcd} A \arrow[r, "\binom{f}{f'}"] & B \oplus B' \arrow[d, "(x_B, x_{B'})"] \\ & (-) \end{tikzcd}$$

So form the triangle

$$A \xrightarrow{\binom{f}{f'}} B \oplus B' \xrightarrow{\bar{g}=(g,g')} C \to \Sigma A.$$

By Proposition 2.8 above, the formula $\exists\, x_{B'}\,\big((x_B, x_{B'})\big(\begin{smallmatrix} f \\ f' \end{smallmatrix}\big) = 0\big)$ is equivalent to $\exists\, x_{B'}\, \exists x_C\, \big((x_B, x_{B'}) = x_C\, \bar{g}\big)$, that is, to

$$\exists x_{B'}\, \exists x_C\, (x_B = x_C\, g \;\wedge\; x_{B'} = x_C\, g'),$$

and the $x_{B'}$ is irrelevant now (set $x_{B'} = x_C\, g'$). So the original formula is equivalent to $g \mid x_B$ where g is, up to sign, the map which appears in the weak pushout

$$\begin{tikzcd} A \arrow[r, "f"] \arrow[d, "f'"'] & B \arrow[d, "g"] \\ B' \arrow[r, "g'"'] & C \end{tikzcd}$$

Lemma 2.9. *Given morphisms $f, f' : A \to B$ in $\mathcal{T}^{\mathrm{c}}$, the (typical pp) formula*

$$\exists\, x_{B'}\, (x_B\, f = x_{B'}\, f')$$

is equivalent to the divisibility formula $g | x_B$, where g is as in the distinguished triangle

$$A \xrightarrow{\binom{f}{f'}} B \oplus B' \xrightarrow{(g,g')} C \to \Sigma A,$$

and hence is also equivalent to the annihilation formula $x_B\, f'' = 0$, where

$$A' \xrightarrow{f''} B' \xrightarrow{g} C \to \Sigma A'$$

is a distinguished triangle.

Thus every pp formula is equivalent on $\mathcal{T}$ to a divisibility formula and hence also to an annihilator formula. In particular:

Theorem 2.10 [Garkusha and Prest 2005, 3.1, 3.2]. *If $\mathcal{T}$ is a compactly generated triangulated category and $\mathcal{L}$ is the language for $\mathcal{T}$ based on $\mathcal{T}^{\mathrm{c}}$, then (the theory of*[11]*) $\mathcal{T}$ has elimination of quantifiers, indeed has elim-q^+.*

2C. ***Types and free realisations.*** We start with a little model theory but soon come back to the algebra.

If $A_1, \ldots, A_n$ are compact objects of $\mathcal{T}$ and if $a_i : A_i \to X \in \mathcal{T}$ are elements of $X \in \mathcal{T}$, then the *type* of $\bar{a} = (a_1, \ldots, a_n)$ (in X) is the set of formulas χ such that $\bar{a} \in \chi(X)$. The *pp-type* of $\bar{a} \in X$ is

$$\mathrm{pp}^X(\bar{a}) = \{\phi \text{ pp} : \bar{a} \in \phi(X)\}.$$

Since we have pp-elimination of quantifiers (Theorem 1.1) the type of $\bar{a}$ in X is determined by its subset $\mathrm{pp}^X(\bar{a})$. Indeed it is equivalent, modulo the theory of $\mathcal{T}$ (equivalently, the theory of absolutely pure (= flat) $\mathcal{T}^{\mathrm{c}}$-modules) to the set $\mathrm{pp}^X(\bar{a}) \cup \{\neg\psi : \psi \text{ pp and } \psi \notin \mathrm{pp}^X(\bar{a})\}$.[12]

As remarked already, because $\mathcal{T}^{\mathrm{c}}$ has finite direct sums, we can replace a tuple $(x_1, \ldots, x_n)$ of variables x_i of sort A_i by a single variable of sort $A_1 \oplus \cdots \oplus A_n$ (and, similarly, tuples of elements may be replaced by single elements). So any pp-definable subgroup of an object $X \in \mathcal{T}$ — that is, the solution set $\phi(X)$ in X of some pp formula ϕ — can be taken to be a subgroup of (A, X) for some $A \in \mathcal{T}^{\mathrm{c}}$.

We say that two formulas are *equivalent* (on $\mathcal{T}$) if they have the same solution set in every $X \in \mathcal{T}$. There is an ordering on the set of (equivalence classes of) pp formulas: if ϕ, ψ are pp formulas in the same free variables, then we set $\phi \leq \psi$ if and only if for all $X \in \mathcal{T}$, $\phi(X) \leq \psi(X)$. This (having fixed the free variables) is a lattice with meet given by conjunction $\phi \wedge \psi$ (defining the intersection of the solution sets) and join given by sum $\phi + \psi$ (defining the sum of the solution sets).

By a *pp-type* (without parameters) we mean a deductively closed set of pp formulas, equivalently a filter (i.e., meet- and upwards-closed) in the lattice of (equivalence classes of) pp formulas (always with some fixed sequence of free variables). We note the following analogue of the module category case (see [Prest 2009, 1.2.23]).

Lemma 2.11. *Suppose that $\mathcal{T}$ is a compactly generated triangulated category and ϕ, ψ are pp formulas with the same free variables. Then $\phi \leq \psi$ if and only if for all $A \in \mathcal{T}^{\mathrm{c}}$ we have $\phi(A) \leq \psi(A)$.*

[11] Meaning that every completion of the theory of $\mathcal{T}$ has elimination of quantifiers and the elimination is uniform over these completions.

[12] This is also true for types with parameters but we don't use these in this paper. For more on this see, for instance, [Prest 1988b, 2.20].

Proof. Suppose that for all $A \in \mathcal{T}^{\mathrm{c}}$ we have $\phi(A) \leq \psi(A)$ and let $X \in \mathcal{T}$. Since yX is a flat object of Mod-$\mathcal{T}^{\mathrm{c}}$, it is the direct limit of some directed diagram of finitely generated projective $\mathcal{T}^{\mathrm{c}}$-modules. The latter all have the form yA for some $A \in \mathcal{T}^{\mathrm{c}}$. Since, for any pp formula ϕ, $\phi(-)$ commutes with direct limits (see [Prest 2009, 1.2.31]), we conclude that $\phi(yX) \leq \psi(yX)$, and hence that $\phi(X) \leq \psi(X)$, as required. □

In the above proof we made the (harmless and useful) identification of pp formulas for objects of $\mathcal{T}$ and for right $\mathcal{T}^{\mathrm{c}}$-modules.

Suppose that p is a pp-type, consisting of pp formulas with free variables $x_1, \dots, x_n$, where x_i has sort (labelled by) $A_i \in \mathcal{T}^{\mathrm{c}}$. Then, by [Prest 2009, 3.3.6, 4.1.4], p has a *realisation* in some object M in the definable subcategory $\langle y\mathcal{T}\rangle$ of Mod-$\mathcal{T}^{\mathrm{c}}$, meaning there is a tuple $\bar{b}$ of elements in M with $\mathrm{pp}^M(\bar{b}) = p$. Pp-types are unchanged by pure embeddings and every such module M is a pure, indeed elementary, subobject of its pure-injective (= injective) hull, which has the form yX for some $X \in \mathcal{T}$. So we obtain a realisation of p in some object $X \in \mathcal{T}$: there is $\bar{a} = (a_1, \dots, a_n)$ with $a_i : A_i \to X$ such that $\mathrm{pp}^X(\bar{a}) = p$. The object X is pure-injective in $\mathcal{T}$ 1.5 and, moreover, may be chosen to be minimal such,[13] in which case it is denoted $H(p)$ — the *hull* of p. This is unique up to isomorphism in the sense that if N is a pure-injective object of $\mathcal{T}$ and if $\bar{c}$ is a tuple from N with $\mathrm{pp}^N(\bar{c}) = p$, then there is an embedding of $H(p)$ into N as a direct summand, taking $\bar{a}$ to $\bar{c}$ and this will be an isomorphism if N also is minimal over $\bar{c}$. See [Prest 2009, §4.3.5] for this and related results — these all apply to any compactly generated triangulated category $\mathcal{T}$ because its model theory is really just that of a definable subcategory of Mod-$\mathcal{T}^{\mathrm{c}}$, and because all the pure-injective objects of that definable subcategory are images of objects of $\mathcal{T}$.

If ϕ is a pp formula, then we have the pp-type it generates:

$$\langle\phi\rangle = \{\psi : \phi \leq \psi\}.$$

We say that a pp-type is *finitely generated* (*by* ϕ) if it has this form for some ϕ.

If ϕ is a pp formula with free variable of sort A (without loss of generality we may assume that there is just one free variable) then a *free realisation* of ϕ is a pair (C, c_A) where $C \in \mathcal{T}^{\mathrm{c}}$ and $c_A : A \to C$ is an element of C of sort A with $\mathrm{pp}^C(c_A) = \langle\phi\rangle$. We have the following analogue to [Prest 2009, 1.2.7]. In the statement of this result, we continue to overuse notation by allowing x_A to denote an element of sort A (in addition to our use of x_A to denote a variable of sort A).

Lemma 2.12. *Suppose ϕ is a pp formula with free variable x_A (for some $A \in \mathcal{T}^{\mathrm{c}}$). Let $C \in \mathcal{T}^{\mathrm{c}}$ and suppose $c_A \in (A, C)$ with $c_A \in \phi(C)$. Then (C, c_A) is a free*

[13] Corresponding to the injective hull of the submodule of M generated by the entries of $\bar{b}$.

realisation of ϕ if and only if for every $x_A : A \to X \in \mathcal{T}$ such that $x_A \in \phi(X)$, there is a morphism $h : C \to A$ with $hc_A = x_A$.

Proof. Existence of free realisations in $\mathcal{T}$ (Corollary 2.14 below) gives the direction ($\Leftarrow$) since, if (B, b) is a free realisation of ϕ, then there is a morphism $g : C \to B$ with $gc_A = b$, so $\mathrm{pp}^C(c_A) = \langle\phi\rangle$ (because morphisms are nondecreasing on pp-types — see [Prest 2009, 1.2.8]). For the converse, if $a \in \phi(X)$, then $ya \in \phi(yX)$.[14] Since the pp-type of yc_A in yC is exactly that of c_A in C, it is generated by ϕ and hence, since $ya \in \phi(yX)$, there is, by [Prest 2009, 1.2.7], a morphism $f' : yC \to yX$ with $f' \cdot yc_A = ya$. Because $C \in \mathcal{T}^c$, there is, by Proposition 1.4, $f : C \to X$ with $f' = yf$. Therefore $y(fc_A) = ya$ so, again by Proposition 1.4, $fc_A = a$, as required. □

We show that every pp formula in the language for $\mathcal{T}$ has a free realisation in $\mathcal{T}$. We use the fact that every formula is equivalent to a divisibility formula.

If a morphism f *factors initially through* a morphism g — that is, $f = hg$ for some h — then write $g \geq f$.

Lemma 2.13. *If $f : A \to B$ is a morphism in $\mathcal{T}^c$ then the pp-type, $\langle f | x_A\rangle$, generated by the formula $f | x_A$ is, up to equivalence of pp formulas, $\{g | x_A : g \geq f\}$.*

Proof. By Lemma 2.9 every pp formula is equivalent to a divisibility formula, so we need only consider formulas of the form $g | x_A$.

If $g \geq f$, say $g : A \to C$ and $f = hg$ with $h : C \to B$, then, for any $x_A : A \to X \in \mathcal{T}$ with $f | x_A$, say $x_A = x_B f$, we have $x_A = x_B hg = x_C g$ with $x_C = x_B g$, so we have $g | x_A$. That is, $g | x_A \in \langle f | x_A\rangle$.

For the converse, if $g : A \to C$ is in $\mathcal{T}^c$ and $g | x_A \in \langle f | x_A\rangle$, then, applying this with $X = B$ and $x_A = f$, we obtain that there is $h : C \to B$ such that $hg = f$, and $g \geq f$, as required. □

Corollary 2.14. *Suppose that $\phi(x_A)$ is a pp formula for the language of $\mathcal{T}$. Choose (by Lemma 2.9) a morphism $f : A \to B$ in $\mathcal{T}^c$ such that ϕ is equivalent to $f | x_A$. Then (B, f) is a free realisation of ϕ.*

2D. *Elimination of imaginaries.* Next we prove elimination of pp-imaginaries: we show that every pp-pair is isomorphic, in the category $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ of pp-pairs, to a pp formula, indeed by Theorem 2.10, to a quantifier-free pp formula if we identify a pp formula $\phi(\bar{x})$ with the pp-pair $\phi(\bar{x})/(\bar{x} = 0)$ in $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$.

Recall (Corollary 2.7) that the category of pp-imaginaries is equivalent to the category $\mathrm{Coh}(\mathcal{T})$ of coherent functors on $\mathcal{T}$. So let us take a coherent functor F_g defined by the exact sequence $(C, -) \xrightarrow{(g,-)} (B, -) \to F_g \to 0$ for some $g : B \to C$

[14] For clarity, the language for $\mathcal{T}$ is exactly the language for Mod-$\mathcal{T}^c$ and the definition of the solution set $\phi(X)$ is identical to the definition of the solution set of $\phi(yX)$.

in $\mathcal{T}^{\mathrm{c}}$. We have the distinguished triangle $A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} \Sigma A$ and extend it to $\Sigma^{-1}C \xrightarrow{\Sigma^{-1}h} A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} \Sigma A$ then consider the exact sequence of functors on $\mathcal{T}$:

$$(\Sigma A, -) \xrightarrow{(h,-)} (C, -) \xrightarrow{(g,-)} (B, -) \xrightarrow{(f,-)} (A, -) \xrightarrow{(\Sigma^{-1}h,-)} (\Sigma^{-1}C, -)$$

where we have the factorisation

$$\begin{tikzcd} (B,-) \arrow[rr, "(f{,}-)"] \arrow[dr, two heads] & & (A,-) \\ & F_g \arrow[ur, hook] & \end{tikzcd}$$

So $F_g \simeq \mathrm{im}(f, -)$ in $(A, -)$ and therefore F_g is isomorphic to the functor given by the pp formula $f \mid x_A$ which, by Proposition 2.8, is equivalent to the quantifier-free pp formula $x_A \cdot \Sigma^{-1}h = 0$; that is $F_g \simeq G^{\circ}_{\Sigma^{-1}h}$ (this is also clear from the above exact sequence). Thus we have the following.

Theorem 2.15 [Garkusha and Prest 2005, 4.3]. *Every pp-pair is pp-definably isomorphic to a pp formula which may be taken to be quantifier-free* (*alternatively a divisibility formula*). *Thus*, (*the theory of*) $\mathcal{T}$ *has elimination of pp imaginaries.*

Explicitly, *if* $g : B \to C$ *is in* $\mathcal{T}^{\mathrm{c}}$ *then the* (*typical*) *pp-pair*

$$F_g = \mathrm{coker}((g, -) : (C, -) \to (B, -))$$

is equivalent to the divisibility formula $f|x_A$ *and to the annihilation formula* $x_A \Sigma^{-1}h = 0$, *where* f *and* h *are such that*

$$\Sigma^{-1}C \xrightarrow{\Sigma^{-1}h} A \xrightarrow{f} B \xrightarrow{g} C (\xrightarrow{h} \Sigma A)$$

is a distinguished triangle.

2E. ***Enhancements, ultraproducts.*** Arguments using reduced products, in particular ultraproducts, are often used in model theory. In many cases their use can be replaced by arguments involving realising types in elementary extensions but in some cases the more algebraic and "explicit" (modulo use of the axiom of choice[15]) ultraproduct construction is better. At first sight we can't use ultraproducts in compactly generated triangulated categories because, even though typically they have direct products, they almost never have all directed colimits (recall, e.g., [Prest 2009, §3.3.1], that an ultraproduct is a directed colimit of direct products of its component structures). Homotopy colimits along a countably infinite directed set are available but that is not enough to form ultraproducts.

[15] needed to extend a filter to a nonprincipal ultrafilter

Laking [2020] introduced ultraproducts in this context by using Grothendieck derivators. We don't go into the details here but see [Laking 2020, §2] for the construction of coherent reduced products for derivators. In [Laking and Vitória 2020] a different approach, using dg-categories and model categories, is taken. This gives, for algebraic compactly generated triangulated categories, a characterisation of definable subcategories (see Section 3A) which is analogous to Theorem 1.6. This extends to any triangulated category with a suitable enhancement, see [Saorín and Šťovíček 2023, 8.8; Bird and Williamson 2022, 6.8] which has the following formulation.

Theorem 2.16 [Laking 2020, 3.11; Laking and Vitória 2020, 4.7; Saorín and Šťovíček 2023, 8.8; Bird and Williamson 2022, 6.8]. *If $\mathcal{D}$ is a subcategory of a compactly generated triangulated category $\mathcal{T}$ which is the underlying category of a strong and stable derivator, then the following are equivalent*:

(i) *$\mathcal{D}$ is a definable subcategory of $\mathcal{T}$.*

(ii) *$\mathcal{D}$ is closed in $\mathcal{T}$ under pure subobjects, products and directed homotopy colimits.*

(iii) *$\mathcal{D}$ is closed in $\mathcal{T}$ under pure subobjects, products and pure quotients.*

Derived categories, derivators, dg-categories, model categories (in the sense of, say, [Hovey 1999]) and ∞-categories all provide ways of representing triangulated categories as the result of applying a process to a somewhat more amenable type of category. In those additive categories with extra structure one can expect the model theory of (multisorted) modules to be directly applicable to the objects. This gives the possibility of approaching the model theory of a triangulated category by developing model theory in such an enhancement and then passing this through a localisation-type functor to the triangulated category. Examples include setting up elementary duality as done in [Angeleri Hügel and Hrbek 2021; Bird and Williamson 2022], see Section 3H. We don't pursue this, so far relatively undeveloped, direction here.

3. Definable subcategories

3A. ***Definable subcategories of $\mathcal{T}$.*** A full subcategory $\mathcal{D}$ of $\mathcal{T}$ is *definable* if its objects form the zero-class of a set of coherent functors, that is, if there is $\mathcal{A} \subseteq \mathrm{Coh}(\mathcal{T})$ such that

$$\mathcal{D} = \{X \in \mathcal{T} : FX = 0 \ \forall F \in \mathcal{A}\}.$$

We will write $\mathcal{D} = \mathrm{Ann}(\mathcal{A}) = \mathrm{Ann}_{\mathcal{T}}(\mathcal{A})$.[16] We will see in Section 3B how this is a natural extension of the notion of definable subcategory of a module category.

[16] We will also use this notation with a set of morphisms replacing $\mathcal{A}$ and hope this will not give rise to confusion.

Also, if $\mathcal{X}$ is a subcategory of $\mathcal{T}$, set

$$\mathrm{Ann}_{\mathrm{Coh}(\mathcal{T})}(\mathcal{X}) = \{F \in \mathrm{Coh}(\mathcal{T}) : FX = 0\ \forall X \in \mathcal{X}\}.$$

As for module categories, we denote by $\langle \mathcal{X} \rangle$ the definable subcategory of $\mathcal{T}$ *generated by* $\mathcal{X}$ — that is, the smallest definable subcategory of $\mathcal{T}$ containing $\mathcal{X}$.

Given a set Φ of morphisms in $\mathcal{T}^{\mathrm{c}}$ we have its annihilator

$$\mathrm{Ann}_{\mathcal{T}}\Phi = \{X \in \mathcal{T} : \forall A \xrightarrow{f} B \in \Phi,\ \forall B \xrightarrow{b} X \text{ we have } bf = 0\}.$$

We write the condition $\forall B \xrightarrow{b} X$ $(bf = 0)$ succinctly as $Xf = 0$ (this being directly analogous to the relation $Mr = 0$ for a right module M and ring element r). Of course we can equally write this condition as $(f, X) = 0$ or $(-, X)f = 0$, according to our viewpoint. Then, [Krause 2002, §7], $\mathrm{Ann}_{\mathcal{T}}\Phi$ is a (typical) definable subcategory of $\mathcal{T}$.

In the other direction, if $\mathcal{X}$ is a subcategory of $\mathcal{T}$, then we may set

$$\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}\mathcal{X} = \{A \xrightarrow{f} B \in \mathcal{T}^{\mathrm{c}} : Xf = 0\ \forall X \in \mathcal{X}\}.$$

The classes of morphisms of the form $\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}\mathcal{X}$ are what Krause calls the *cohomological ideals* of $\mathcal{T}^{\mathrm{c}}$; we will refer to them simply as *annihilator ideals* in $\mathcal{T}^{\mathrm{c}}$.

Lemma 3.1 [Krause 2002, §7]. *If Φ is a set of morphisms in $\mathcal{T}^{\mathrm{c}}$, then $\mathrm{Ann}_{\mathcal{T}}\Phi$ is a definable subcategory of $\mathcal{T}$. If $\mathcal{X}$ is any subcategory of $\mathcal{T}$, then $\mathrm{Ann}_{\mathcal{T}}(\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}\mathcal{X})$ is equal to $\langle \mathcal{X} \rangle$, the definable subcategory of $\mathcal{T}$ generated by $\mathcal{X}$. In particular there is a natural bijection between the definable subcategories of $\mathcal{T}$ and the cohomological (= annihilator) ideals in $\mathcal{T}^{\mathrm{c}}$.*

We have seen already that if

$$A \xrightarrow{f} B \xrightarrow{g} C \to \Sigma A$$

is a triangle, then

$$bf = 0 \iff g \mid b.$$

So we consider, given a set Ψ of morphisms in $\mathcal{T}^{\mathrm{c}}$,

$$\mathrm{Div}_{\mathcal{T}}\Psi = \{X \in \mathcal{T} : \forall\, B \xrightarrow{g} C \in \Psi,\ \forall\, B \xrightarrow{b} X,\ \exists\, C \xrightarrow{c} X \text{ such that } b = cg\}$$

$$\begin{tikzcd} B \arrow[r, "g"] \arrow[d, "\forall"'] & C \arrow[dl, dotted, "\exists"] \\ X & \end{tikzcd}$$

— the class of Ψ-divisible objects of $\mathcal{T}$. We write $g|X$ as a succinct expression of the condition "$\forall\, B \xrightarrow{b} X\ \exists\, C \xrightarrow{c} X$ such that $b = cg$" (being the analogue of the condition that every element of a module M be divisible by an element r of the ring[17]). Then $\mathrm{Div}_{\mathcal{T}}\Psi$ is a (typical) definable subcategory of $\mathcal{T}$.

And, in the other direction, given a subcategory $\mathcal{X}$ of $\mathcal{T}$, we define[18]

$$\mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}\mathcal{X} = \{B \xrightarrow{g} C \in \mathcal{T}^{\mathrm{c}} : g|X \ \forall X \in \mathcal{X}\}.$$

Lemma 3.2 [Angeleri Hügel and Hrbek 2021, 2.2]. *If Ψ is a set of morphisms in $\mathcal{T}^{\mathrm{c}}$, then $\mathrm{Div}_{\mathcal{T}}\Psi$ is a definable subcategory of $\mathcal{T}$. If $\mathcal{X}$ is any subcategory of $\mathcal{T}$, then $\mathrm{Div}_{\mathcal{T}}(\mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}\mathcal{X}) = \langle \mathcal{X} \rangle$.*

Proof. Take $Y \in \mathrm{Div}_{\mathcal{T}}(\mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}\mathcal{X})$. If $g \in \mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}\mathcal{X}$ then $g|Y$ so, if f is as above, $Yf = 0$. This is so for all such f (as g varies) so, by Lemma 3.1, $Y \in \langle \mathcal{X} \rangle$, as required. □

Corollary 3.3. (1) *If $\mathcal{D} = \mathrm{Ann}_{\mathcal{T}}\Phi$ is a definable subcategory of $\mathcal{T}$ then also*

$$\mathcal{D} = \mathrm{Div}_{\mathcal{T}}\{g : A \xrightarrow{f} B \xrightarrow{g} C \to \Sigma A \textit{ is a distinguished triangle and } f \in \Phi\}.$$

(2) *If $\mathcal{D} = \mathrm{Div}_{\mathcal{T}}\Psi$ is a definable subcategory of $\mathcal{T}$ then also*

$$\mathcal{D} = \mathrm{Ann}_{\mathcal{T}}\{f : A \xrightarrow{f} B \xrightarrow{g} C \to \Sigma A \textit{ is a distinguished triangle and } g \in \Psi\}.$$

Definable subcategories are so-called because they can be defined by closure of certain pairs of pp formulas, that is, by requiring that certain quotients of pp-definable subgroups be 0. For each of the annihilation and divisibility methods of specifying these subcategories, the pp-pairs needed are obvious, being respectively $\{(x_B = x_B)/(x_B f = 0) : f : A \to B \in \Phi\}$ and $\{(x_B = x_B)/(g|x_B) : g : B \to C \in \Psi\}$ with Φ, Ψ as above.

We have used that pp-pairs can be given in both annihilation and divisibility forms, but there is another, "torsionfree" form that is not so obvious if we consider only formulas and their reduction to divisibility or annihilator forms, rather than pp-pairs. Let us consider an extended triangle as before:

$$\Sigma^{-1}C \xrightarrow{\Sigma^{-1}h} A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} \Sigma A.$$

If $X \in \mathcal{T}$ then we obtain an exact sequence of abelian groups

$$(\Sigma A, X) \xrightarrow{(h,X)} (C, X) \xrightarrow{(g,X)} (B, X) \xrightarrow{(f,X)} (A, X) \xrightarrow{(\Sigma^{-1}h,X)} (\Sigma^{-1}C, X).$$

Then

$$X \in \mathrm{Div}_{\mathcal{T}}(g) \iff (g, X) \text{ is epi} \iff (f, X) = 0 \iff (\Sigma^{-1}h, X) \text{ is monic.}$$

[17] But the corresponding notation $Xg = X$ would be less appropriate than in the usual module case because X has many sorts and that equation applies only to the B-sort of X.

[18] We are overworking the notations Ann and Div but they are useful.

If we denote by $\mathrm{ann}_X(\Sigma^{-1}h)$ the set $\{a : A \to X : a.\Sigma^{-1}h = 0\}$, then we have

$$Xf = 0 \iff g|X \iff \mathrm{ann}_X(\Sigma^{-1}h) = 0. \tag{4}$$

That is,

$$X \in \mathcal{T} \text{ annihilates } f \iff X \text{ is } g\text{-divisible} \iff X \text{ is } \Sigma^{-1}h\text{-torsionfree.}$$

This gives us a third way of using morphisms in $\mathcal{T}^{\mathrm{c}}$ to cut out definable subcategories of $\mathcal{T}$. We set, given $\mathcal{X} \subseteq \mathcal{T}$,

$$\mathcal{X}\text{-Reg} = \{\ell \in \mathcal{T}^{\mathrm{c}} : \mathrm{ann}_X(\ell) = 0 \ \forall X \in \mathcal{X}\}$$

and call such classes, for want of a better word, *regularity* classes (of morphisms of $\mathcal{T}^{\mathrm{c}}$).

In the other direction, given a set Ξ of morphisms in $\mathcal{T}^{\mathrm{c}}$, we define

$$\Xi\text{-TF} = \{X \in \mathcal{T} : \mathrm{ann}_X(\ell) = 0 \ \forall \ell \in \Xi\}.$$

Lemma 3.4. *If Ξ is a set of morphisms in $\mathcal{T}^{\mathrm{c}}$, then Ξ-TF is a definable subcategory of $\mathcal{T}$. If $\mathcal{X}$ is any subcategory of $\mathcal{T}$, then $(\mathcal{X}\text{-Reg})\text{-TF} = \langle \mathcal{X} \rangle$.*

The argument is as for Lemma 3.2.

The set of pp-pairs corresponding to Ξ is $\{(x_A\ell = 0)/(x_A = 0) : D \xrightarrow{\ell} A \in \Xi\}$.

The next result summarises some of this; see [Saorín and Šťovíček 2023, 8.6] and, for the case where $\mathcal{T}$ is the derived category of modules over a ring, [Angeleri Hügel and Hrbek 2021, 2.2].

Theorem 3.5. *A definable subcategory $\mathcal{D}$ of $\mathcal{T}$ may be specified by any of the following means*:

- $\mathcal{D} = \{X \in \mathcal{T} : \phi(X)/\psi(X) = 0 \ \forall \phi/\psi \in \Phi\}$ *where Φ is a set of pp-pairs in* $\mathcal{L}(\mathcal{T})$;
- $\mathcal{D} = \mathrm{Ann}_{\mathcal{T}}(\mathcal{A})$ *where* $\mathcal{A} \subseteq \mathrm{Coh}(\mathcal{T})$;
- $\mathcal{D} = \mathrm{Ann}_{\mathcal{T}}\Phi$ *where Φ is a set of morphisms in* $\mathcal{T}^{\mathrm{c}}$;
- $\mathcal{D} = \mathrm{Div}_{\mathcal{T}}\Psi$ *where Ψ is a set of morphisms in* $\mathcal{T}^{\mathrm{c}}$;
- $\mathcal{D} = \Xi\text{-TF}$ *where Ξ is a set of morphisms in* $\mathcal{T}^{\mathrm{c}}$.

The subcategories of $\mathrm{Coh}(\mathcal{T})$ *of the form* $\mathrm{Ann}_{\mathrm{Coh}(\mathcal{T})}(\mathcal{D})$ *are the Serre subcategories, the classes of morphisms of $\mathcal{T}^{\mathrm{c}}$ of the form* $\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ *are the annihilator = cohomological ideals.*[19]

Moving between the last three specifications is described by (4) above.

[19]The classes $\mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}\mathcal{D}$ and $\mathcal{D}$-Reg are described indirectly, in terms of the functors they present, at the end of Section 3C.

In Section 3C we will say this in torsion-theoretic terms with mod-$\mathcal{T}^c$ in place of $\mathrm{Coh}(\mathcal{T})$. In Section 3B we give the relevant background.

3B. ***Torsion theories on* Mod-$\mathcal{T}^c$.** A *torsion pair* in a Grothendieck category, such as Mod-$\mathcal{T}^c$, consists of two classes: $\mathcal{G}$ — the *torsion* class, and $\mathcal{F}$ — the *torsionfree* class, with $(\mathcal{G}, \mathcal{F}) = 0$ and with $\mathcal{G}$, $\mathcal{F}$ maximal such. Such a torsion pair, or *torsion theory*, is *hereditary* if $\mathcal{G}$ is closed under subobjects, equivalently if $\mathcal{F}$ is closed under injective hulls and, if so, it is *of finite type* if $\mathcal{G}$ is generated, as a hereditary torsion class, by finitely presented objects, equivalently if $\mathcal{F}$ is closed under directed colimits (see, for instance, [Prest 2009, 11.1.12, 11.1.14]). We also use without further comment that, for a hereditary torsion theory, if F is a torsionfree module then the injective hull $E(F)$ of F is torsionfree (and conversely, since the torsionfree class is closed under subobjects). For background on torsion theories, see [Stenström 1975].

The restricted Yoneda functor from $\mathcal{T}$ to Mod-$\mathcal{T}^c$ allows us to realise the definable subcategories of $\mathcal{T}$ as the inverse images of finite-type torsionfree classes on Mod-$\mathcal{T}^c$, as follows.

Suppose that $\mathcal{D}$ is a definable subcategory of $\mathcal{T}$. Then $\mathcal{D}$ is determined by the class $\mathcal{D} \cap \mathrm{Pinj}(\mathcal{T})$ of pure-injectives in it, being the closure of that class under pure subobjects (by the comments after Theorem 1.6). By Theorem 1.5 the image $\mathcal{E} = y(\mathcal{D} \cap \mathrm{Pinj}(\mathcal{T}))$ is a class of injective $\mathcal{T}^c$-modules which is closed under direct products and direct summands, hence (e.g., [Prest 2009, 11.1.1]) which is of the form $\mathcal{F} \cap \text{Inj-}\mathcal{T}^c$ for some hereditary torsionfree class $\mathcal{F} = \mathcal{F}_{\mathcal{D}}$ of $\mathcal{T}^c$-modules.

We recall, [Prest 1979, 3.3] see [Prest 2009, 11.1.20], that a hereditary torsionfree class of modules is of finite type exactly if it is definable. So we have to show that definability of $\mathcal{D}$ corresponds to definability of $\mathcal{F}_{\mathcal{D}}$, equivalently to definability of the class of absolutely pure objects in $\mathcal{F}_{\mathcal{D}}$ ("equivalently" because Mod-$\mathcal{T}^c$ is locally coherent, so the absolutely pure objects form a definable subcategory, see [Prest 2009, 3.4.24], hence so is their intersection with any other definable subcategory; in the other direction, if $\mathcal{F}_{\mathcal{D}} \cap \text{Abs-}\mathcal{T}^c$ is definable then so also, by, e.g., Theorem 1.6, is its class of subobjects, which is precisely $\mathcal{F}_{\mathcal{D}}$). So we have to show that the torsionfree class $\mathcal{F}_{\mathcal{D}}$ above is of finite type and that every finite type torsionfree class arises in this way.

To see, this, note that, if $X \in \mathcal{T}$ and $F \in \mathrm{Coh}(\mathcal{T})$, then (Section 2A) $FX = 0$ if and only if $(F^\diamond, yX) = 0$. Set $\mathcal{A} = \mathrm{Ann}_{\mathrm{Coh}(\mathcal{T})}(\mathcal{D})$. We have the duality from Section 2A between $\mathrm{Coh}(\mathcal{T})$ and mod-$\mathcal{T}^c$, so consider the corresponding set $\mathcal{A}^\diamond = \{F^\diamond : F \in \mathcal{A}\}$ of finitely presented $\mathcal{T}^c$-modules. Since $\mathcal{A}$ is a Serre subcategory of $\mathrm{Coh}(\mathcal{T})$, this is a Serre subcategory of mod-$\mathcal{T}^c$; we set $\mathcal{S}_{\mathcal{D}} = \mathcal{A}^\diamond$. The $\varinjlim$-closure,[20] $\overrightarrow{\mathcal{S}_{\mathcal{D}}}$, in Mod-$\mathcal{T}^c$ of $\mathcal{S}_{\mathcal{D}}$ is a typical hereditary torsion class of

[20]If $\mathcal{S}$ is a subcategory of a module category, then we will denote its $\varinjlim$-*closure* — its closure under directed colimits — by $\vec{\mathcal{S}}$.

finite type in Mod-$\mathcal{T}^c$ (see [Prest 2009, 11.1.36]). The corresponding hereditary torsionfree class $\mathcal{F} = \{M \in \text{Mod-}\mathcal{T}^c : (\overrightarrow{\mathcal{S}_\mathcal{D}}, M) = 0\}$ is just the hom-perp of $\mathcal{S}_\mathcal{D} : \mathcal{F} = \{M \in \text{Mod-}\mathcal{T}^c : (\mathcal{S}_\mathcal{D}, M) = 0\}$. If $M \in \mathcal{F}$ is injective, hence (Theorem 1.5) of the form yN for some pure-injective $N \in \mathcal{T}$, then the condition $(\mathcal{S}_\mathcal{D} = \mathcal{A}^\diamond, M) = 0$ is exactly the condition $FN = 0$ for every $F \in \mathcal{A}$, that is, the condition that N is in $\mathcal{D}$. Thus $\mathcal{F} = \mathcal{F}_\mathcal{D}$ and we have the correspondence between classes of pure-injectives in $\mathcal{T}$ of the form $\mathcal{D} \cap \text{Pinj}(\mathcal{T})$ and classes of injectives in Mod-$\mathcal{T}^c$ of the form $\mathcal{F} \cap \text{Inj-}\mathcal{T}^c$ for some hereditary torsionfree class $\mathcal{F}$. (For, note that given such a class $\mathcal{E}$ of injectives, the class of pure submodules of modules in $\mathcal{E}$ is the class of absolutely pure modules in $\mathcal{F}$ which, by finite type, is definable and hence has definable inverse image in $\mathcal{T}$). Therefore we have shown the following.

Theorem 3.6. *A subcategory $\mathcal{D}$ of a compactly generated triangulated category $\mathcal{T}$ is definable if and only if it has any of the following equivalent forms, where $y : \mathcal{T} \to \text{Mod-}\mathcal{T}^c$ is the restricted Yoneda functor*:

- $\mathcal{D} = y^{-1}\mathcal{F}$, *where $\mathcal{F}$ is a finite-type hereditary torsionfree class in* Mod-$\mathcal{T}^c$;
- $\mathcal{D} = y^{-1}\mathcal{E}$, *where $\mathcal{E}$ is the class of absolutely pure objects in a hereditary torsionfree class of finite type*;
- $\mathcal{D} = y^{-1}\mathcal{E}$, *where $\mathcal{E}$ is a definable class of absolutely pure objects in* Mod-$\mathcal{T}^c$.

We denote by $\tau_\mathcal{D} = (\mathscr{T}_\mathcal{D}, \mathcal{F}_\mathcal{D})$ the finite-type hereditary torsion theory on Mod-$\mathcal{T}^c$ corresponding to $\mathcal{D}$.

Corollary 3.7. *The definable subcategories $\mathcal{D}$ of $\mathcal{T}$ are in natural bijection with the definable* (= *finite-type*) *hereditary torsionfree classes in* Mod-$\mathcal{T}^c$ *and also with the definable subcategories of* Abs-$\mathcal{T}^c$.

Explicitly, to $\mathcal{D}$ correspond respectively the closure $\mathcal{F}_\mathcal{D}$ of $\langle y\mathcal{D}\rangle$ under submodules, and $\mathcal{F}_\mathcal{D} \cap$ Abs-$\mathcal{T}^c$. *In the other direction, we simply apply y^{-1}, where y is the restricted Yoneda functor.*

Note the almost complete analogy of this with the bijection (see [Prest 2009, 12.3.2]) between definable subcategories of a module category Mod-R and the finite type (= definable) hereditary torsionfree classes in (R-mod)-Mod = (R-mod, **Ab**), equivalently with the definable classes of absolutely pure objects in (R-mod)-Mod = (R-mod, **Ab**). One notable difference is that the image of a definable subcategory of a triangulated category is "most" of the definable subcategory $\langle y\mathcal{D}\rangle \subseteq \text{Abs-}\mathcal{T}^c$ of modules, whereas in the module case it is all of the corresponding class of modules. This reflects the lack of directed colimits in triangulated categories, but see [Laking 2020; Laking and Vitória 2020] for some replacement using Grothendieck derivators for the triangulated case.

The other notable difference is that the module case uses tensor product to embed (fully and faithfully) Mod-R in (R-mod, **Ab**). Here we have somehow avoided that.

We also record the equivalence at the level of pure-injectives.

Corollary 3.8. *If $\mathcal{D}$ is a definable subcategory of $\mathcal{T}$ and $\mathcal{F}_{\mathcal{D}}$ is the corresponding hereditary torsionfree class in* Mod-$\mathcal{T}^{\mathrm{c}}$, *then the restricted Yoneda functor y induces an equivalence*

$$\mathrm{Pinj}(\mathcal{D}) \simeq \mathcal{F} \cap \text{Inj-}\mathcal{T}^{\mathrm{c}}$$

between the category $\mathrm{Pinj}(\mathcal{D})$ *of pure-injective objects of $\mathcal{T}$ which lie in $\mathcal{D}$ and the category $\mathcal{F} \cap$* Inj-$\mathcal{T}^{\mathrm{c}}$ *of $\mathcal{T}^{\mathrm{c}}$-injective modules which lie in $\mathcal{F}$.*

This gives some justification for our saying that the Yoneda image of a definable subcategory $\mathcal{D}$ in Mod-$\mathcal{T}^{\mathrm{c}}$ constitutes "most of" the flat (= absolutely pure) objects of the corresponding hereditary torsionfree class of finite type. For, every injective in the class is in the image and every absolutely pure object in the class is a pure (even elementary) submodule of an object in the image.

Note that the fact that the objects of $\mathcal{D}$ are the pure subobjects of the pure-injectives in $\mathcal{D}$ exactly corresponds to the fact that the absolutely pure modules in $\mathcal{F}$ are the pure submodules of the injective modules in $\mathcal{F}$.

3C. ***Definable subcategories of* Abs-$\mathcal{T}^{\mathrm{c}}$.** In Section 3A we associated to a definable subcategory $\mathcal{D}$ of $\mathcal{T}$ three sets of morphisms, $\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$, $\mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ and $\mathcal{D}$-Reg, each of which determines $\mathcal{D}$. In this section we identify the corresponding sets of morphisms in mod-$\mathcal{T}^{\mathrm{c}}$ and the ways in which they cut out the hereditary finite type torsion theory $\tau_{\mathcal{D}}$ cogenerated by $\langle y\mathcal{D}\rangle$ in Mod-$\mathcal{T}^{\mathrm{c}}$.

We have the following from Section 3B.

Corollary 3.9. *If $\mathcal{T}$ is a compactly generated triangulated category, then the following are in natural bijection*:

(i) *The definable subcategories of $\mathcal{T}$.*

(ii) *The definable subcategories of* Mod-$\mathcal{T}^{\mathrm{c}}$ *which are contained in* (*so are definable subcategories of*) Abs-$\mathcal{T}^{\mathrm{c}}$ = Flat-$\mathcal{T}^{\mathrm{c}}$.

(iii) *The hereditary torsion theories on* Mod-$\mathcal{T}^{\mathrm{c}}$ *of finite type.*

(iv) *The Serre subcategories of* mod-$\mathcal{T}^{\mathrm{c}}$.

Given a definable subcategory $\mathcal{D}$ of $\mathcal{T}$, let

$$\mathcal{S}_{\mathcal{D}} = \{G \in \text{mod-}\mathcal{T}^{\mathrm{c}} : (G, yX) = 0 \ \forall\, X \in \mathcal{D}\}$$

be the corresponding Serre subcategory of mod-$\mathcal{T}^{\mathrm{c}}$. As noted in Section 3B, this is the Serre subcategory $(\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D}))^{\diamond}$ of mod-$\mathcal{T}^{\mathrm{c}}$, it $\varinjlim$-generates the finite type hereditary torsion class $\mathscr{T}_{\mathcal{D}}$ and $\tau_{\mathcal{D}} = (\mathscr{T}_{\mathcal{D}}, \mathcal{F}_{\mathcal{D}})$ is the torsion theory corresponding to $\mathcal{D}$ under (i)$\leftrightarrow$(iii) of Corollary 3.9.

If τ is any hereditary torsion theory then a submodule L of a module M is *τ-dense in M* if M/L is torsion. Also, the *τ-closure*, $\mathrm{cl}_{\tau}^{M}(L)$, of a submodule L

of a module M is the maximal submodule of M in which L is τ-dense, also characterised as the smallest submodule L' of M which contains L and is such that M/L' is τ-torsionfree. See [Stenström 1975] or [Prest 2009, §11.1] for details.

First we see that the annihilation, divisibility and regularity conditions with respect to $\mathcal{D}$ translate directly to Mod-$\mathcal{T}^{\mathrm{c}}$.

Proposition 3.10. *Suppose that $\mathcal{D}$ is a definable subcategory of $\mathcal{T}$ and $f : A \to B$ is in $\mathcal{T}^{\mathrm{c}}$. Then*

(1) $f \in \mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ *if and only if* $yX.yf = 0$ *for all* $X \in \mathcal{D}$;

(2) $f \in \mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ *if and only if, for every* $X \in \mathcal{D}$, yX *is* yf*-divisible*;

(3) $f \in \mathcal{D}$-Reg *if and only if, for every* $X \in \mathcal{D}$, *if* $b' : yB \to yX$ *is such that* $b'.yf = 0$ *then* $b' = 0$.

Proof. First we note that, in all three cases, it is enough for the direction ($\Leftarrow$) to prove that f has the property (annihilation, divisibility, regularity) for $X \in \mathcal{D}$ pure-injective. That is because, if $X \in \mathcal{D}$, then f satisfies, say, $Xf = 0$ if (indeed if and only if) $H(X)f = 0$, where $H(X)$ is the pure-injective hull of X. That is because X is pure in (indeed is an elementary substructure of) its pure-injective hull so, if a pp-pair is closed on $H(X)$, then it will be closed on X (and *vice versa*).

(1) The defining condition for f to be in $\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$, namely that $Xf = 0$ for all $X \in \mathcal{D}$, certainly implies $yX.yf = 0$ for all $X \in \mathcal{D}$. If, conversely, $yX.yf = 0$ for all $X \in \mathcal{D}$, then take $X \in \mathcal{D}$ and suppose we have $b : B \to X$. Then $y(bf) = yb.yf = 0$ so, by Proposition 1.4, $bf = 0$. Therefore $Xf = 0$, as required.

(2) If $f \in \mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ and we have $a' : yA \to yX$, then we compose with the inclusion of yX into its injective hull $E(yX) = yH(X)$ (by Theorem 1.5) to get a morphism $a'' : yA \to yH(X)$ which, by Proposition 1.3, has the form ya for some $a : A \to H(X)$. By assumption, and since $H(X) \in \mathcal{D}$, a factors through f, say $a = bf$ with $b : B \to H(X)$; therefore $a'' = yb.yf$. Thus $\exists x_{yB}(a'' = x_{yB}.yf)$ is true in $yH(X)$. Since yX is a pure submodule of $yH(X)$ we deduce that $\exists x_{yB}(a' = x_{yB}.yf)$ is true in yX, that is, yX is yf-divisible. This gives ($\Rightarrow$).

For the converse, suppose that, for every $X \in \mathcal{D}$, yX is yf-divisible and take $X \in \mathcal{D}$ pure-injective and $a : A \to X$. Then we have $ya : yA \to yX$ so, by hypothesis, there is $b' : yB \to yX$ with $b'.yf = ya$. Since X is pure-injective, by Proposition 1.3 there is $b : B \to X$ such that $b' = yb$, giving $y(bf) = ya$. By Proposition 1.4 it follows that $bf = a$, showing that every pure-injective object in $\mathcal{D}$ is f-injective. By the comments at the beginning of the proof and the fact that the divisibility condition is expressed by closure of a pp-pair, it follows that every object of $\mathcal{D}$ is f-injective, as required.

(3) The direction ($\Leftarrow$) follows immediately from Proposition 1.4. For the converse, if $f \in \mathcal{D}$-Reg then take $X \in \mathcal{D}$ to be pure-injective, and suppose $b' : yB \to yX$

is such that $b'.yf = 0$. By Proposition 1.3, $b' = yb$ for some $b : B \to X$. That gives $y(bf) = 0$ hence, by Proposition 1.4, $bf = 0$, hence, by assumption, $b = 0$, so that $b' = 0$. Thus f is regular on every pure-injective in $\mathcal{D}$ and so, since that is expressed by closure of a pp-pair, f is regular on every $X \in \mathcal{D}$, as required. □

Set $\mathcal{S}_{\mathcal{D}}^{\circ} = \{G^{\circ} : G \in \mathcal{S}_{\mathcal{D}}\}$ to be the image of $\mathcal{S}_{\mathcal{D}} \subseteq \text{mod-}\mathcal{T}^{\mathrm{c}}$ in $\mathrm{Coh}(\mathcal{T})$ under the antiequivalence 2.6. Note that, by definition of $G \mapsto G^{\circ}$, $\mathcal{S}_{\mathcal{D}}^{\circ}$ consists exactly of the coherent functors F such that $FX = 0$ for every $X \in \mathcal{D}$, that is $(\mathcal{S}_{\mathcal{D}})^{\circ} = \mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$.

Proposition 3.11. *Suppose that $\mathcal{D}$ is a definable subcategory of $\mathcal{T}$, let $\mathcal{S}_{\mathcal{D}}$ be the corresponding Serre subcategory of* mod-$\mathcal{T}^{\mathrm{c}}$*. Denote by $\tau_{\mathcal{D}}$ the corresponding hereditary (finite-type) torsion theory in* Mod-$\mathcal{T}^{\mathrm{c}}$*. Let $f : A \to B$ be a morphism in $\mathcal{T}^{\mathrm{c}}$. Then*

(1) $f \in \mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ *if and only if* $\mathrm{im}(yf) \in \mathcal{S}_{\mathcal{D}}$,

(2) $f \in \mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ *if and only if* $\ker(yf) \in \mathcal{S}_{\mathcal{D}}$ *if and only if* $F_f \in \mathcal{S}_{\mathcal{D}}^{\circ}$,

(3) $f \in \mathcal{D}$-Reg *if and only if* $G_f = \mathrm{coker}(yf) \in \mathcal{S}_{\mathcal{D}}$*, that is, if and only if* $\mathrm{im}(yf)$ *is $\tau_{\mathcal{D}}$-dense in yB.*

Proof. We use that $X \in \mathcal{D}$ if and only if yX is ($\tau_{\mathcal{D}}$-)torsionfree, that is, if and only if $(\mathcal{S}_{\mathcal{D}}, yX) = 0$.

(1) If the image $\mathrm{im}(yf)$ is in $\mathcal{S}_{\mathcal{D}}$ then, for every $X \in \mathcal{D}$, we have $(\mathrm{im}(yf), yX) = 0$ because yX is torsionfree. Therefore $yX.yf = 0$, for all $X \in \mathcal{D}$ giving, by Proposition 3.10, the implication ($\Leftarrow$). For the other direction, first note that any morphism from $\mathrm{im}(yf)$ to yX extends to a morphism from yB to yX by absolute purity (= fp-injectivity) of yX. If $\mathrm{im}(yf)$ were not torsion, there would be a nonzero morphism from $\mathrm{im}(yf)$ to some torsionfree object which, for instance replacing the object by its injective hull, we may assume to be of the form yX with $X \in \mathcal{D}$. This would give a morphism $a : yB \to yX$ with $af \neq 0$, contradicting Proposition 3.10.

(2) ($\Rightarrow$) By Proposition 3.10 we have that yX is yf-divisible for every $X \in \mathcal{D}$. If $\ker(yf)$ were not torsion (that is, since, by local coherence of Mod-$\mathcal{T}^{\mathrm{c}}$, it is finitely presented, not in $\mathcal{S}_{\mathcal{D}}$) then it would have a nonzero torsionfree quotient M. The (torsionfree) injective hull of M would have the form yX for some pure-injective $X \in \mathcal{D}$, yielding a morphism $yA \to yX$ which is not zero on the kernel of yf, hence which cannot factor through yf — a contradiction.

For the converse, assume that $\ker(yf) \in \mathcal{S}_{\mathcal{D}}$. Then any morphism $a' : yA \to yX$ with $X \in \mathcal{D}$ must be zero on $\ker(yf)$, since yX is torsionfree. Therefore a' factors through $\mathrm{im}(yf)$. But yX is absolutely pure so, since $\mathrm{im}(yf)$ is a finitely generated subobject of yB, that factorisation extends to a morphism $b' : yB \to yX$. Thus we have a factorisation of a' through yf, and so yX is yf-divisible. By Proposition 3.10 that is enough.

For the part involving $\mathcal{S}_{\mathcal{D}}^{\circ}$, we have $f \in \mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ if and only if $(f, X) : (B, X) \to (A, X)$ is epi for every $X \in \mathcal{D}$ if and only if $\mathrm{coker}(f, X) = 0$ for every $X \in \mathcal{D}$, that is, if and only if $F_f X = 0$ for every $X \in \mathcal{D}$ and that, as noted above, is the case if and only if $F_f \in \mathcal{S}_{\mathcal{D}}^{\circ}$.

(3) If $\mathrm{im}(yf)$ is not $\tau_{\mathcal{D}}$-dense in yB, there will be a nonzero morphism from yB and with kernel containing $\mathrm{im}(yf)$ to a torsionfree object, hence to an object of the form yX with $X \in \mathcal{D}$. Therefore yf is not $y\mathcal{D}$-regular and so, by Proposition 3.10, f is not $\mathcal{D}$-regular.

For the converse, suppose that $\mathrm{im}(yf)$ is $\tau_{\mathcal{D}}$-dense in yB. Then, if b' is a morphism from yB to a torsionfree object and the kernel of b' contains $\mathrm{im}(yf)$ then, since the image of b' is torsion, we have $b' = 0$. Therefore every object in $y\mathcal{D}$ is yf-torsionfree which, by Proposition 3.10, is as required. □

From this, Theorem 3.5 and the equivalences (4), we have the following, where we apply the notations Ann, Div and Reg and their definitions to Mod-$\mathcal{T}^{\mathrm{c}}$ with, of course, mod-$\mathcal{T}^{\mathrm{c}}$ replacing $\mathcal{T}^{\mathrm{c}}$ as the subcategory of "small" objects. This is mostly [Wagstaffe 2021, 5.1.4].

Theorem 3.12. *Suppose that $\mathcal{D}$ is a definable subcategory of $\mathcal{T}$, let $\tau_{\mathcal{D}}$ be the corresponding finite-type hereditary torsion theory in* Mod-$\mathcal{T}^{\mathrm{c}}$ *and let $\mathcal{S}_{\mathcal{D}}$ denote the Serre subcategory of $\tau_{\mathcal{D}}$-torsion finitely presented $\mathcal{T}^{\mathrm{c}}$-modules.*

Suppose that

$$A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} \Sigma A$$

is a distinguished triangle. Then

(i) $f \in \mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ *if and only if* $yf \in \mathrm{Ann}_{\text{mod-}\mathcal{T}^{\mathrm{c}}}(y\mathcal{D})$ *if and only if* $\mathrm{im}(yf) \in \mathcal{S}_{\mathcal{D}}$;

(ii) $g \in \mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ *if and only if* $yg \in \mathrm{Div}_{\text{mod-}\mathcal{T}^{\mathrm{c}}}(y\mathcal{D})$ *if and only if* $\ker(yg) \in \mathcal{S}_{\mathcal{D}}$, *that is, if and only if* $F_g \in \mathcal{S}_{\mathcal{D}}^{\circ}$;

(iii) $\Sigma^{-1}h \in \mathcal{D}$-Reg *if and only if the image of $y(\Sigma^{-1}h)$ is $\tau_{\mathcal{D}}$-dense in $y(\Sigma^{-1}C)$, that is, if and only if $G_{\Sigma^{-1}h} \in \mathcal{S}_{\mathcal{D}}$.*

Furthermore, the conditions (i), (ii) *and* (iii) *are equivalent.*

3D. ***Model theory in definable subcategories.*** If $\mathcal{D}$ is a definable category, meaning a category equivalent to a definable subcategory of a module category (over a ring possibly with many objects), then the model theory of $\mathcal{D}$ is intrinsic to $\mathcal{D}$, in the following senses.

First, the notion of pure-exact sequence is intrinsic to $\mathcal{D}$ because an exact sequence is pure-exact if and only if some ultraproduct of it is split-exact; see [Prest 2009, 4.2.18]. Ultraproducts are obtained as directed colimits of products, so definable categories have ultraproducts. Definable subcategories of compactly generated triangulated categories do not in general have directed colimits, so they are not (quite) "definable categories" in this sense, though they are quite close; see Theorem 2.16. Nevertheless, as we have seen, the restricted Yoneda functor associates, to a definable subcategory $\mathcal{D}$ of a compactly generated triangulated category, a definable subcategory of a module category which has the same model theory.

Question. Is the model theory of a definable subcategory $\mathcal{D}$ of a compactly generated triangulated category intrinsic, meaning definable just from the structure of $\mathcal{D}$ as a category?

Second, the category $\mathbb{L}(\mathcal{D})^{\mathrm{eq}+}$ of pp-imaginaries for a definable subcategory $\mathcal{D}$ of a module category Mod-R is equivalent to the Serre localisation $\mathbb{L}_R^{\mathrm{eq}+}/\mathcal{S}_{\mathcal{D}}$, where $\mathcal{S}_{\mathcal{D}}$ is the Serre subcategory of coherent functors which annihilate $\mathcal{D}$. We have the same description for a definable subcategory of a compactly generated triangulated category, via the restricted Yoneda functor. But neither of those descriptions is intrinsic because both refer to a containing (module, or triangulated) category. In the module case, there is an intrinsic description of $\mathbb{L}(\mathcal{D})^{\mathrm{eq}+}$ as the category $(\mathcal{D}, \mathbf{Ab})^{\prod\rightarrow}$ of functors from $\mathcal{D}$ to $\mathbf{Ab}$ which commute with direct products and directed colimits. For $\mathcal{T}$ itself, there is a similar description in [Krause 2002, 5.1] but we may ask whether this extends to definable subcategories.

In any case, if $\mathcal{D}$ is a definable subcategory of a compactly generated triangulated category $\mathcal{T}$, then the category, $\mathbb{L}(\mathcal{D})^{\mathrm{eq}+}$, of pp-imaginaries for $\mathcal{D}$ is the quotient of $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ by its Serre subcategory consisting of those pp-pairs which are closed on $\mathcal{D}$. In terms of the other forms of the category of pp-imaginaries given by Corollary 2.7, $\mathbb{L}(\mathcal{D})^{\mathrm{eq}+}$ also has the following descriptions.

Proposition 3.13. *If $\mathcal{D}$ is a definable subcategory of a compactly generated triangulated category $\mathcal{T}$, then the following categories are equivalent*:

(i) *The category*, $\mathbb{L}(\mathcal{D})^{\mathrm{eq}+}$, *of pp-imaginaries for $\mathcal{D}$.*

(ii) $\mathrm{Coh}(\mathcal{T})/\mathrm{Ann}_{\mathrm{Coh}(\mathcal{T})}(\mathcal{D})$.

(iii) mod-$\mathcal{T}^{\mathrm{c}}/\mathcal{S}_{\mathcal{D}}$.

Note that the contravariant action of $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ via $(\mathbb{L}(\mathcal{T})^{\mathrm{eq}+})^{\mathrm{op}} \simeq \text{mod-}\mathcal{T}^{\mathrm{c}}$ acting by $G(X) = (G, yX)$ for $G \in \text{mod-}\mathcal{T}^{\mathrm{c}}$ and $X \in \mathcal{T}$ localises as the action of mod-$\mathcal{T}^{\mathrm{c}}/\mathcal{S}_{\mathcal{D}}$ on $\langle Q_{\mathcal{D}}(y\mathcal{D})\rangle = \langle Q_{\mathcal{D}}(y\mathcal{T})\rangle$, where $Q_{\mathcal{D}} : \text{Mod-}\mathcal{T}^{\mathrm{c}} \to \text{Mod-}\mathcal{T}^{\mathrm{c}}/\overrightarrow{\mathcal{S}_{\mathcal{D}}}$ is the corresponding Gabriel localisation and the action is given by the same formula. This places both the category of models and the category of imaginaries (the latter contravariantly) into the same Grothendieck abelian category, just as in the module case where we can use the tensor embedding; see [Prest 2009, §12.1.1].

3E. *Hom-orthogonal pairs on $\mathcal{T}$ and torsion theories on* Mod-$\mathcal{T}^{\mathrm{c}}$. A *hom-orthogonal pair*[21] on $\mathcal{T}$ is a pair $(\mathcal{U}, \mathcal{V})$ of subcategories with $\mathcal{U} = {}^{\perp}\mathcal{V}$ the *torsion* class and $\mathcal{V} = \mathcal{U}^{\perp}$ the *torsionfree* class. Such a pair $(\mathcal{U}, \mathcal{V})$ is said to be *compactly generated* if there is $\mathcal{A} \subseteq \mathcal{T}^{\mathrm{c}}$ such that $\mathcal{V} = \mathcal{A}^{\perp} = \{Y \in \mathcal{T} : (A, Y) = 0 \ \forall A \in \mathcal{A}\}$, in which case $\mathcal{U} = {}^{\perp}(\mathcal{A}^{\perp}) = \{Z \in \mathcal{T} : (Z, \mathcal{A}^{\perp}) = 0\}$; we say that $\mathcal{A}$ *generates* the hom-orthogonal pair. Note that $\mathcal{V}$ is in this case definable, being given by the conditions that each sort $(A, -)$ for $A \in \mathcal{A}$ is 0, that is, all the pp-pairs $x_A = x_A / x_A = 0$ for $A \in \mathcal{A}$ are closed on $\mathcal{V}$.

Proposition 3.14. *Suppose that $(\mathcal{U}, \mathcal{V})$ is a hom-orthogonal pair in $\mathcal{T}$, compactly generated by $\mathcal{A} \subseteq \mathcal{T}^{\mathrm{c}}$. Let $\tau_{\mathcal{V}} = (\mathscr{T}_{\mathcal{V}}, \mathcal{F}_{\mathcal{V}})$ denote the finite-type hereditary torsion theory on* Mod-$\mathcal{T}^{\mathrm{c}}$ *corresponding (Corollary 3.7) to the definable subcategory $\mathcal{V}$. Let* $\mathrm{Ser}(y\mathcal{A})$ *denote the Serre subcategory of* mod-$\mathcal{T}^{\mathrm{c}}$ *generated by $y\mathcal{A}$.*

Then $\mathscr{T}_{\mathcal{V}} = \overrightarrow{\mathrm{Ser}(y\mathcal{A})}$ and $\mathcal{F}_{\mathcal{V}} = (y\mathcal{A})^{\perp} = \{M \in \text{Mod-}\mathcal{T}^{\mathrm{c}} : (yA, M) = 0 \ \forall A \in \mathcal{A}\}$.

Proof. This follows from what we have seen already; we give the details. Since $(\mathcal{A}, \mathcal{V}) = 0$, it follows by Proposition 1.4 that $(y\mathcal{A}, y\mathcal{V}) = 0$, so $\overrightarrow{\mathrm{Ser}(y\mathcal{A})} \subseteq \mathscr{T}_{\mathcal{V}}$. Hence $\mathcal{F}_{\mathcal{V}} = (\mathscr{T}_{\mathcal{V}})^{\perp} \subseteq (\overrightarrow{\mathrm{Ser}(y\mathcal{A})})^{\perp} = (y\mathcal{A})^{\perp}$ (equality since $\tau_{\mathcal{V}}$ is of finite type). If, conversely, $M \in (y\mathcal{A})^{\perp}$, then so is $E(M)$, which has the form yN for some pure-injective $N \in \mathcal{T}$. By Proposition 1.3 (or Proposition 1.4), $(\mathcal{A}, N) = 0$ and hence $N \in \mathcal{V}$, so $E(M)$, and hence M is in $\mathcal{F}_{\mathcal{V}}$. Thus $\mathcal{F}_{\mathcal{V}} = (y\mathcal{A})^{\perp}$ and hence also $\mathscr{T}_{\mathcal{V}} = \overrightarrow{\mathrm{Ser}(y\mathcal{A})}$. □

By Corollary 3.7, every finite-type hereditary torsion theory $(\mathscr{T}, \mathcal{F})$ on Mod-$\mathcal{T}^{\mathrm{c}}$ gives rise to a hom-orthogonal pair in $\mathcal{T}$, namely $({}^{\perp}\mathcal{D}, ({}^{\perp}\mathcal{D})^{\perp})$ where $\mathcal{D} = y^{-1}\mathcal{F}$. If this hom-orthogonal pair is compactly generated, by $\mathcal{A}$ say, so $({}^{\perp}\mathcal{D})^{\perp} = \mathcal{A}^{\perp}$ is definable, then it follows from the above that $\mathcal{F} = (y\mathcal{A})^{\perp}$ and hence $\mathcal{D} = y^{-1}\mathcal{F} = y^{-1}((y\mathcal{A})^{\perp}) = \mathcal{A}^{\perp}$ (by the bijection 3.7) $= ({}^{\perp}\mathcal{D})^{\perp}$. But in general not every finite-type hereditary torsion class in Mod-$\mathcal{T}^{\mathrm{c}}$ arises from a hom-orthogonal pair in $\mathcal{T}$ in this way. Indeed, since, for $A \in \mathcal{T}^{\mathrm{c}}$, yA is a projective $\mathcal{T}^{\mathrm{c}}$-module, and all of the finitely generated projectives in Mod-$\mathcal{T}^{\mathrm{c}}$ are of this form, we have the following, where we denote by $\gamma_{\mathcal{X}}$ the hereditary (finite type) torsion theory generated by (that is, with torsion class generated by) $y\mathcal{X}$.

Corollary 3.15. *There is a natural injection $(\mathcal{U}, \mathcal{V}) \mapsto \gamma_{\mathcal{U}}$ from the set of compactly generated hom-orthogonal pairs in $\mathcal{T}$ to the set of hereditary torsion theories of finite type on* Mod-$\mathcal{T}^{\mathrm{c}}$.

The image is the set of hereditary torsion theories where the torsion class is generated by a set of finitely generated projectives.

[21] In the context of triangulated categories, the term "torsion pair" is used for a stronger concept; see [Šťovíček and Pospíšil 2016, §3].

Thus we have an embedding of the lattice of compactly generated hom-orthogonal pairs in $\mathcal{T}$ into the lattice of finite type hereditary torsion theories on Mod-$\mathcal{T}^{\mathrm{c}}$ (the ordering in each case being by inclusion of torsion classes), and the latter is isomorphic to the lattice of definable subcategories of $\mathcal{T}$. The definable subcategories, $\mathcal{D}$, of $\mathcal{T}$ occurring as $\mathcal{V}$ in a compactly generated hom-orthogonal pair $(\mathcal{U}, \mathcal{V})$, are, by Proposition 3.11(1), those for which the corresponding annihilator ideal $\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ of $\mathcal{T}^{\mathrm{c}}$ is generated as such by objects (that is, by identity morphisms of some compact objects).

Note also that, if $\mathcal{D}$ is a definable subcategory of $\mathcal{T}$ which occurs as $\mathcal{V}$ in a compactly generated hom-orthogonal pair $(\mathcal{U}, \mathcal{V})$, and if $(\mathscr{T}, \mathcal{F})$ is the corresponding, in the sense of Corollary 3.7, torsion theory $\tau_{\mathcal{D}}$, then we always have $\mathcal{U} \subseteq y^{-1}\mathscr{T}$. That is because $\mathscr{T} = {}^{\perp}(\mathcal{F} \cap \text{Inj-}\mathcal{T}^{\mathrm{c}})$ and because each object of $\mathcal{F} \cap \text{Inj-}\mathcal{T}^{\mathrm{c}}$ has the form yN for some pure-injective $N \in \mathcal{V}$ and then $(\mathcal{U}, N) = 0$ implies, by Proposition 1.3, that $(y\mathcal{U}, yN) = 0$, so $y\mathcal{U} \subseteq \mathscr{T}$. For equality, $\mathcal{U} \subseteq y^{-1}\mathscr{T}$ — that is, $\gamma_{\mathcal{U}} = \tau_{\mathcal{D}}$ — we need, by the argument just given, that $\mathcal{U} = {}^{\perp}(\mathcal{V} \cap \mathrm{Pinj}(\mathcal{T}))$. That is, equality holds if and only if the hom-orthogonal pair $(\mathcal{U}, \mathcal{V})$ is *cogenerated by pure-injectives*. For instance, if $(\mathcal{U}, \mathcal{V})$ is a t-structure with $\mathcal{V}$ definable, then this will be the case; see [Angeleri Hügel and Hrbek 2021, 2.10; Saorín and Šťovíček 2023, 8.20] and also Proposition 3.19 below.

For more about this and TTF-classes in compactly generated triangulated categories, see [Wagstaffe 2021, Chapter 8].

3F. ***Spectra.*** By a definable (additive) category we mean a category which is equivalent to a definable subcategory of the category of modules over some (possibly multisorted) ring. Every definable additive category $\mathcal{C}$ is determined by its full subcategory of pure-injective objects (by [Prest 2009, 5.1.4] or, more intrinsically, by [Prest 2012a, §3.2]). Indeed, every definable category is determined by the indecomposable pure-injective objects in it (e.g., see [Prest 2009, 5.3.50, 5.3.52]). The Ziegler spectrum, $\mathrm{Zg}(\mathcal{C})$, also written Zg_R in the case $\mathcal{C} = \text{Mod-}R$, is the set, $\mathrm{pinj}(\mathcal{C})$, of isomorphism classes of indecomposable pure-injectives in $\mathcal{C}$ endowed with the topology which has, for a basis of open sets, the

$$(\phi/\psi) = \{N \in \mathrm{pinj}(\mathcal{C}) : \phi(N) > \psi(N)\}$$

as ϕ/ψ ranges over pp-pairs (in any suitable language for $\mathcal{C}$). These are exactly the compact open sets in $\mathrm{Zg}(\mathcal{C})$; see [Prest 2009, 5.1.22].

Every definable subcategory $\mathcal{D}$ of a definable category $\mathcal{C}$ is determined by the set $\mathrm{pinj}(\mathcal{D}) = \mathcal{D} \cap \mathrm{pinj}(\mathcal{C})$ of indecomposable pure-injectives in $\mathcal{D}$, hence by the closed subset $\mathrm{Zg}(\mathcal{D}) = \mathcal{D} \cap \mathrm{Zg}(\mathcal{C})$ of $\mathrm{Zg}(\mathcal{C})$, and every closed set in $\mathrm{Zg}(\mathcal{C})$ is of the form $\mathrm{Zg}(\mathcal{D})$ for some definable subcategory $\mathcal{D}$ of $\mathcal{C}$; see [Prest 2009, 5.1.1].

Krause [2000] showed how this carries over to compactly generated triangulated categories $\mathcal{T}$. The *Ziegler spectrum*, $\mathrm{Zg}(\mathcal{T})$, of $\mathcal{T}$ is defined to have, for its points, the (isomorphism classes of) indecomposable pure-injectives. As for definable subcategories of module categories, there are many equivalent ways of specifying a basis of (compact) open sets on this set of points, including the following (the second by Theorem 2.15):

$$(\phi/\psi) = \{N \in \mathrm{pinj}(\mathcal{T}) : \phi(N)/\psi(N) \neq 0\} \text{ for } \phi/\psi \text{ a pp-pair;}$$
$$\{N \in \mathrm{pinj}(\mathcal{T}) : \mathrm{ann}_N(f) \neq 0\} \text{ for } f \text{ a morphism in } \mathcal{T}^{\mathrm{c}};$$
$$(F) = \{N \in \mathrm{pinj}(\mathcal{T}) : FN \neq 0\} \text{ for } F \in \mathrm{Coh}(\mathcal{T}).$$

There are other topologies of interest here. First consider the case where R is commutative noetherian. Then the subcategory, Inj-R, of injectives in Mod-R is definable (see [Prest 2009, 3.4.28]) and the corresponding closed subset of Zg_R is just the set, inj_R, of indecomposable injective R-modules. For such a ring the set inj_R may be identified [Gabriel 1962], see [Prest 2009, §14.1.1], with $\mathrm{Spec}(R)$ via $P \mapsto E(R/P)$, where P is any prime ideal of R and $E(-)$ denotes injective hull. However, the Ziegler topology restricted from Zg_R to inj_R induces, via the above bijection, not the Zariski topology on $\mathrm{Spec}(R)$ but its Hochster dual [Prest 1988b, pp. 104–105]. Recall that the *Hochster dual* of a topology has, as a basis (on the same set of points), the complements of the compact open sets in the original topology.

That fact inspired the general definition [Prest 1993, pp. 200–202] of the *dual-Ziegler* (or "rep-Zariski") topology on $\mathrm{pinj}(\mathcal{C})$ for any definable category $\mathcal{C}$, as the Hochster-dual of the Ziegler topology.[22] So this dual topology has the same underlying set, $\mathrm{pinj}(\mathcal{C})$, and has, for a basis of open sets, the complements

$$[\phi/\psi] = \mathrm{Zg}(\mathcal{C}) \setminus (\phi/\psi)$$

of the compact Ziegler-open sets.

If $\mathcal{C}$ is a locally coherent category, in particular if it is Mod-R for a right coherent ring (possibly with many objects), then[23] the absolutely pure objects form a definable subcategory with corresponding closed subset of $\mathrm{Zg}(\mathcal{C})$ again being the set $\mathrm{inj}(\mathcal{C})$ of (isomorphism types of) indecomposable injectives in $\mathcal{C}$. This set carries a (*Gabriel–*)*Zariski* topology which has, for a basis of open sets, those of the form

$$[A] = \{E \in \mathrm{inj}(\mathcal{C}) : (A, C) = 0\}$$

[22] These spaces are, however, unlike those in Hochster's original definition, not spectral, and it is not always that case that the Ziegler topology is returned as the dual of the dual-Ziegler topology [Burke and Prest 2002, 3.1]

[23] For module categories, this goes back to [Eklof and Sabbagh 1971], see [Prest 2009, 3.4.24]; the general case is proved the same way and also follows from, for example, [Prest 2011a, Chapter 6].

for A a finitely presented object of $\mathcal{C}$. Thus we extend the domain of applicability of the category-theoretic reformulation [Gabriel 1962; Roos 1961] of the definition of the Zariski topology on a commutative coherent ring. For such a category $\mathcal{C}$ the Gabriel–Zariski topology coincides with the dual-Ziegler topology restricted to $\mathrm{inj}(\mathcal{C})$ [Prest 2009, 14.1.6].

We may compare these topologies over a commutative coherent ring R where, in general, the map $P \mapsto E(R/P)$ is only an inclusion of $\mathrm{Spec}(R)$ into inj_R, because there may be indecomposable injectives not of the form $E(R/P)$, e.g., [Prest 2009, 14.4.1]. The inclusion, nevertheless, is a topological equivalence — an isomorphism of frames of open subsets: every indecomposable injective is elementarily equivalent to, hence topologically equivalent to, a module of the form $E(R/P)$ with P a prime; see [Prest 2009, 14.4.5]. So, for commutative coherent rings, we may consider these various topologies as topologies on $\mathrm{Spec}(R)$ and, so considered, the Ziegler topology coincides with the *Thomason* topology, which is defined to be the Hochster-dual of the Gabriel–Zariski topology [Garkusha and Prest 2008]. That is, the Ziegler topology has, for its open sets, those of the form $\bigcup_\lambda (R/I_\lambda)$ with the I_λ finitely generated ideals of R, where

$$(R/I_\lambda) = \{N \in \mathrm{pinj}_R : (R/I_\lambda, N) \neq 0\} = (xI_\lambda = 0/x = 0).$$

In terms of sets of primes, the Ziegler-open sets have the form $\bigcup_\lambda V(I_\lambda)$ with the I_λ finitely generated.[24] These various topologies are compared in [Prest 2012c, §6].

The discussion above applies to the locally coherent category $\text{Mod-}\mathcal{T}^{\mathrm{c}}$. As we have seen in Theorem 1.5, the restricted Yoneda functor y induces an equivalence between the category, $\mathrm{Pinj}(\mathcal{T})$, of pure-injective objects of $\mathcal{T}$ and the category, $\text{Inj-}\mathcal{T}^{\mathrm{c}}$, of injective right $\mathcal{T}^{\mathrm{c}}$-modules. Indeed, this gives a homeomorphism of spectra.

Theorem 3.16. *Suppose that $\mathcal{T}$ is a compactly generated triangulated category. Then $y : \mathcal{T} \to \text{Mod-}\mathcal{T}^{\mathrm{c}}$ induces a bijection between $\mathrm{pinj}(\mathcal{T})$ and $\mathrm{inj}\mathcal{T}^{\mathrm{c}}$. This is a homeomorphism between $\mathrm{Zg}(\mathcal{T})$ and $\mathrm{Zg}(\text{Abs-}\mathcal{T}^{\mathrm{c}} = \text{Flat-}\mathcal{T}^{\mathrm{c}})$ (the latter can also be regarded as $\mathrm{inj}_{\mathcal{T}^{\mathrm{c}}}$ with the Thomason topology) and is also a homeomorphism between the dual-Ziegler spectrum $\mathrm{Zar}(\mathcal{T})$ of $\mathcal{T}$ and $\mathrm{inj}^{\mathrm{c}}_{\mathcal{T}}$ if the latter is equipped with the Gabriel–Zariski topology which has, for a basis of open sets, the sets $[G] = \{E \in \text{Inj-}\mathcal{T}^{\mathrm{c}} : (G, E) = 0\}$ for $G \in \text{mod-}\mathcal{T}^{\mathrm{c}}$.*

Since closed subsets of the Ziegler spectrum are in natural correspondence with definable subcategories, this homeomorphism underlies the bijection (Corollary 3.7) between definable subcategories of $\mathcal{T}$ and finite-type hereditary torsionfree classes in $\text{Mod-}\mathcal{T}^{\mathrm{c}}$. That also reflects the fact that a finite-type hereditary torsion theory is

[24]For a general commutative ring, the Ziegler topology on inj_R is finer, having open sets of a similar form but with pp-definable ideals replacing finitely generated ideals; in coherent rings the pp-definable ideals coincide with the finitely generated ideals; see [Prest 2012c, §6].

determined by (it is the torsionfree class cogenerated by) the set of indecomposable torsionfree injectives (see [Prest 2009, 11.1.29]). We have already, in Section 3E, considered the part of this correspondence coming from compactly generated hom-orthogonal pairs in $\mathcal{T}$, and we will also, in Section 4A, look at how the Balmer spectrum fits into this picture in the case that $\mathcal{T}$ is tensor-triangulated.

3G. ***Triangulated definable subcategories.*** In this section we consider the definable subcategories $\mathcal{D}$ of $\mathcal{T}$ which are *triangulated*, that is, *shift-closed* (if $X \in \mathcal{D}$, then $\Sigma^{\pm} X \in \mathcal{D}$) and extension-closed, where by *extension-closed* we mean that, if $X \to Y \to Z \to \Sigma X$ is a distinguished triangle with both X and Z in $\mathcal{D}$, then also $Y \in \mathcal{D}$. First, some remarks on extending definable subcategories to shift-closed definable subcategories.

If $\mathcal{D}$ is a definable subcategory of $\mathcal{T}$ then each shift $\Sigma^i \mathcal{D}$ is definable (e.g., see [Wagstaffe 2021, 6.1.1]). We can define the shift-closure of $\mathcal{D}$ to be the definable closure of $\bigcup_{i \in \mathbb{Z}} \Sigma^i \mathcal{D}$. That this is, in general, larger than $\mathrm{Add}^+\big(\bigcup_{i \in \mathbb{Z}} \Sigma^i \mathcal{D}\big)$ ($^+$ denoting closure under pure submodules) is shown by the following example.

Example 3.17. Consider the derived category $\mathcal{D}_{k[\epsilon]} = \mathcal{D}(\text{Mod-}k[\epsilon])$, of the category of modules over $k[\epsilon] = k[x]/(x^2)$. Let $\mathcal{D}$ be the subcategory of $\mathcal{D}_{k[\epsilon]}$ consisting of complexes which are 0 in every degree $i < 0$. Then $\mathcal{D}$ is a definable subcategory, defined by the conditions $(k[\epsilon][i], -) = 0$ $(i < 0)$, where $k[\epsilon]$ here denotes the complex with $k[\epsilon]$ in degree 0 and zeroes elsewhere.

The union of the (left) shifts of $\mathcal{D}$ contains only complexes which are bounded below and so the additive closure of the union $\bigcup_i \mathrm{Zg}(\Sigma^i \mathcal{D})$ of the Ziegler-spectra of these shifts does not contain, for example, the doubly infinite complex which has $k[\epsilon]$ in each degree and multiplication by ϵ for each of its maps. But that indecomposable pure-injective complex belongs to the Ziegler-closure of $\bigcup_i \mathrm{Zg}(\Sigma^i \mathcal{D})$, indeed it is in the Ziegler-closure of the set of complexes obtained from it by replacing $k[\epsilon]$ by 0 in every degree $\leq i$ for some i; this is proved in [Han 2013, §3.4] and, in greater generality, in [Arnesen et al. 2017, §6, §4].

In contrast, if we were to take $\mathcal{D}$ to be the image of Mod-$k[\epsilon]$ consisting of complexes concentrated in degree 0, then the additive closure of the union of the shifts of $\mathcal{D}$ is definable. That follows because every object in the definable category generated by that union is finite endolength, so the Ziegler closure contains no new indecomposable pure-injectives (e.g., see [Prest 2009, 4.4.30]).

Thus, if X is a closed subset of the Ziegler spectrum of $\mathcal{T}$, it may be that $\bigcup_i \Sigma^i X$ is not Ziegler-closed.

It is the case, see [Wagstaffe 2021, 6.1.10], that, if points of $\mathrm{Zg}(\mathcal{T})$ are identified with their shifts and the set of equivalence classes is given the quotient topology, then this is topologically equivalent to the space based on $\mathrm{pinj}(\mathcal{T})$ which has, for its closed sets, those of the form $\mathcal{D} \cap \mathrm{pinj}(\mathcal{T})$ where $\mathcal{D}$ is a shift-closed definable subcategory of $\mathcal{T}$. The first example in Example 3.17 shows that the projection map taking a point of the Ziegler spectrum of $\mathcal{T}$ to its shift equivalence class need not be closed (the complexes in that example are endofinite, hence Ziegler-closed points).

Further Ziegler-type topologies on $\mathrm{pinj}(\mathcal{T})$ are obtained by using positively-(alternatively, negatively-) shift-closed definable subcategories of $\mathcal{T}$; see [Wagstaffe 2021, §6.1]).

A triangulated subcategory $\mathcal{B}$ of $\mathcal{T}$ is *smashing* if it is the kernel of a Bousfield localisation $q : \mathcal{T} \to \mathcal{T}'$ for which the left adjoint to q, including $\mathcal{T}' = \mathcal{T}/\mathcal{B}$ into $\mathcal{T}$, preserves coproducts. Hom-orthogonality gives a bijection between the definable subcategories which are triangulated and the smashing subcategories of $\mathcal{T}$.

Theorem 3.18 ([Krause 2005], see [Wagstaffe 2021, 5.2.10]). *If $\mathcal{D}$ is a triangulated definable subcategory of the compactly generated triangulated category $\mathcal{T}$, then $\mathcal{B} = {}^{\perp}\mathcal{D}$ is a smashing subcategory of $\mathcal{T}$ and $\mathcal{D} = \mathcal{B}^{\perp}$, so $(\mathcal{B}, \mathcal{D})$ is a torsion pair. Every smashing subcategory of $\mathcal{T}$ arises in this way.*

Proposition 3.19 [Krause 2000, 3.9, Theorem C]. *Suppose that $\mathcal{B}$ is a smashing subcategory of $\mathcal{T}$ and $\mathcal{D} = \mathcal{B}^{\perp}$ is the corresponding triangulated definable subcategory. Then $\mathcal{B} = y^{-1}\mathscr{T}_{\mathcal{D}}$, where $\mathscr{T}_{\mathcal{D}} = \overrightarrow{\mathcal{S}_{\mathcal{D}}}$ is the torsion class for the torsion theory $\gamma_{\mathcal{B}} = \tau_{\mathcal{D}}$ generated by $y\mathcal{B}$, equivalently cogenerated by $y\mathcal{D}$.*

Corollary 3.20. *If $\mathcal{D}$ is a triangulated definable subcategory of $\mathcal{T}$, and $\mathscr{T}_{\mathcal{D}}$ is the corresponding hereditary torsion class in* $\text{Mod-}\mathcal{T}^{\mathrm{c}}$, *then $y^{-1}\mathscr{T}_{\mathcal{D}} = {}^{\perp}\mathcal{D}$ is a* (*typical*) *smashing subcategory of $\mathcal{T}$.*

One says that $\mathcal{T}$ has the *telescope property* if, for each smashing subcategory $\mathcal{B}$, the torsion pair $(\mathcal{B}, \mathcal{D})$ is compactly generated, equivalently, Corollary 3.15, if the Serre subcategory $\mathcal{S}_{\mathcal{D}} = \mathscr{T}_{\mathcal{D}} \cap \text{mod-}\mathcal{T}^{\mathrm{c}}$ is generated by projective (= representable) objects; see [Krause 2000, Introduction].

3H. *Elementary duality.* If R is any skeletally small preadditive category (= multisorted ring), then there is a duality — *elementary duality*, [Prest 1988a; Herzog 1993], see [Prest 2009, §§1.3, 10.3] — between the category of pp-pairs for right R-modules and the category of pp-pairs for left R-modules. This duality induces a natural bijection between the definable subcategories of Mod-R and R-Mod, [Herzog 1993, 6.6] see [Prest 2009, §3.4.2].

In particular this applies with $R = \mathcal{T}^{\mathrm{c}}$. Because the model theory of $\mathcal{T}$ is essentially that of $\text{Flat-}\mathcal{T}^{\mathrm{c}} = \text{Abs-}\mathcal{T}^{\mathrm{c}}$ inside $\text{Mod-}\mathcal{T}^{\mathrm{c}}$, it follows that we have a version of elementary duality between $\mathcal{T}$ and the definable subcategory $\mathcal{T}^{\mathrm{c}}\text{-Abs} = \mathcal{T}^{\mathrm{c}}\text{-Flat}$ of $\mathcal{T}^{\mathrm{c}}$-Mod. In particular, elementary duality gives a natural bijection between the definable subcategories of $\mathcal{T}$ and those of $\mathcal{T}^{\mathrm{c}}$-Flat.

With the module situation in mind, it is natural to ask whether there is a compactly triangulated category $\mathcal{T}_1$ such that $\mathcal{T}_1^{\mathrm{c}} \simeq (\mathcal{T}^{\mathrm{c}})^{\mathrm{op}}$ and hence an elementary duality between the model theory of $\mathcal{T}$ and the model theory of $\mathcal{T}_1$ via Mod-$\mathcal{T}_1^{\mathrm{c}} \simeq \mathcal{T}^{\mathrm{c}}$-Mod. This situation is considered in [Garkusha and Prest 2005, §7]. In particular, if $\mathcal{T}$ is the derived category of modules over a ring then this is so, [Garkusha and Prest 2005, 7.5], see also [Angeleri Hügel and Hrbek 2021]; more generally it is so if $\mathcal{T}$ is an algebraic triangulated category, [Bird and Williamson 2022].

Question. If $\mathcal{T}$ is a compactly generated triangulated category, is there a triangulated category $\mathcal{T}_1$ and an elementary duality between $\mathcal{T}$ and $\mathcal{T}_1$? If such a category $\mathcal{T}_1$ exists, is it essentially unique?

By "an elementary duality" we mean at least a natural bijection between definable subcategories, probably also an antiequivalence between the respective categories of pp-sorts, perhaps also a duality at the level of pp formulas. See the remarks in Section 2E about enhancements.

This also raises some further general questions.

Questions. What is a characterisation of the categories which arise as $\mathcal{T}^{\mathrm{c}}$ where $\mathcal{T}$ is compactly generated triangulated? Given such a category, does it come from a unique compactly generated triangulated category $\mathcal{T}$? and, if so, how can $\mathcal{T}$ be constructed from it? In particular is $(\mathcal{T}^{\mathrm{c}})^{\mathrm{op}}$ of the form $\mathcal{T}_1^{\mathrm{c}}$ for some compactly generated triangulated category $\mathcal{T}_1$?

These seem to be hard questions to answer; they include the, only partly resolved, Margolis conjecture in the case that $\mathcal{T}$ is the stable homotopy category of spectra.

If $\mathcal{T}$ is the derived category $\mathcal{D}_R = \mathcal{D}(\text{Mod-}R)$ of some ring R, we do get a good elementary duality between $\mathcal{D}_R$ and $\mathcal{D}_{R^{\mathrm{op}}} = \mathcal{D}(R\text{-Mod})$. This follows because the duality $(\text{proj-}R)^{\mathrm{op}} \to \text{proj-}R^{\mathrm{op}}$ between the categories of finitely generated projectives given by $P \mapsto (P, R)$ extends to the respective categories of perfect complexes, that is, to a duality $(-)^{\mathrm{t}} : (\mathcal{D}_R^{\mathrm{c}})^{\mathrm{op}} \simeq \mathcal{D}_{R^{\mathrm{op}}}^{\mathrm{c}}$, see [Garkusha and Prest 2005, §7; Angeleri Hügel and Hrbek 2021, §2.2]. In these papers, R is a 1-sorted ring but the arguments also apply if R is a skeletally small preadditive category. In [Bird and Williamson 2022, §3.2] this is extended to algebraic triangulated categories via dg-enhancements. We will, in Section 4B, describe an internal duality, from [Wagstaffe 2021, Chapter 7] in the tensor-triangulated case. If R is commutative, so $\mathcal{D}_R \simeq \mathcal{D}_{R^{\mathrm{op}}}$, the duality in [Angeleri Hügel and Hrbek 2021] does coincide ([Wagstaffe 2021, 7.3.5]) with the internal duality described in Section 4B.

For details, we refer the reader to those papers; in particular, the generalisation in [Bird and Williamson 2022] to algebraic triangulated categories uses enhancements (see Section 2E), which we don't go into here (also see [Laking and Vitória 2020] for related use of enhancements). For an abstract approach to dualities between triangulated categories, see [Bird and Williamson 2022].

We continue a little further in the case that $\mathcal{T}$ is the derived category $\mathcal{D}_R$ of a module category. If $\mathcal{D}$ is a definable subcategory of $\mathcal{D}_R$, then we have the corresponding annihilator ideal $\mathrm{Ann}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D})$. Set $(\mathrm{Ann}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D}))^{\mathrm{t}} = \{f^{\mathrm{t}} : f \in \mathrm{Ann}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D})\}$, where $(-)^{\mathrm{t}} : (\mathcal{D}_R^{\mathrm{c}})^{\mathrm{op}} \simeq \mathcal{D}_{R^{\mathrm{op}}}^{\mathrm{c}}$ is the duality from the previous paragraph. Then, [Angeleri Hügel and Hrbek 2021, 2.3], $(\mathrm{Ann}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D}))^{\mathrm{t}}$ is an annihilator ideal of $\mathcal{D}_{R^{\mathrm{op}}}^{\mathrm{c}}$. We set $\mathcal{D}^{\mathrm{d}} = \mathrm{Ann}_{\mathcal{D}_{R^{\mathrm{op}}}}((\mathrm{Ann}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D}))^{\mathrm{t}})$ and refer to this as the definable subcategory of $\mathcal{D}_{R^{\mathrm{op}}}$ *elementary dual* to $\mathcal{D}$. The terminology is further justified by the following, which refers, using the obvious notations, to the other ways of specifying definable subcategories.

Proposition 3.21 [Angeleri Hügel and Hrbek 2021, 2.2–2.5]. *If $\mathcal{D}$ is a definable subcategory of $\mathcal{D}_R$ and $\mathcal{D}^{\mathrm{d}}$ is its elementary dual definable subcategory of $\mathcal{D}_{R^{\mathrm{op}}}$, then*:

$$\mathrm{Ann}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D}^{\mathrm{t}}) = (\mathrm{Ann}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D}))^{\mathrm{t}}, \quad \mathrm{Div}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D}^{\mathrm{t}}) = (\mathcal{D}\text{-TF})^{\mathrm{t}}, \quad \mathcal{D}^{\mathrm{t}}\text{-TF} = (\mathrm{Div}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D}))^{\mathrm{t}}.$$

Proof. The first is by definition and [Angeleri Hügel and Hrbek 2021, 2.3]. For the others consider $f \in \mathrm{Ann}_{\mathcal{D}_R^{\mathrm{c}}}(\mathcal{D})$ and form the extended triangle

$$\Sigma^{-1}B \xrightarrow{\Sigma^{-1}g} \Sigma^{-1}C \xrightarrow{\Sigma^{-1}h} A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} \Sigma A$$

then dualise it:

$$(\Sigma A)^{\mathrm{t}} = \Sigma^{-1}A^{\mathrm{t}} \xrightarrow{h^{\mathrm{t}}} C^{\mathrm{t}} \xrightarrow{g^{\mathrm{t}}} B^{\mathrm{t}} \xrightarrow{f^{\mathrm{t}}} A^{\mathrm{t}} \xrightarrow{\Sigma h^{\mathrm{t}}} \Sigma\, C^{v} \xrightarrow{\Sigma g^{\mathrm{t}}} \Sigma\, B^{\mathrm{t}}.$$

Then we use the equivalences (4) from Section 3A, namely:

$$Xf = 0 \iff g|X \iff \mathrm{ann}_X(\Sigma^{-1}h) = 0.$$

From that we directly obtain the other two equalities. □

We also have, just as for definable subcategories of module categories, that the category of pp-pairs for $\mathcal{D}^{\mathrm{d}}$ is the opposite to that for $\mathcal{D}$. The latter is equivalent to mod-$\mathcal{D}_R^{\mathrm{c}}/\mathcal{S}_{\mathcal{D}}$, where $\mathcal{S}_{\mathcal{D}} = \{G : (G, yX) = 0\ \forall X \in \mathcal{D}\}$. We set $d\mathcal{S}_{\mathcal{D}} = \{dG : G \in \mathcal{S}_{\mathcal{D}}\}$, where d is the duality of Corollary 2.4.[25]

Proposition 3.22. *If $\mathcal{D}$ is a definable subcategory of $\mathcal{D}_R$ and $\mathcal{D}^{\mathrm{d}}$ is its elementary dual definable subcategory of $\mathcal{D}_{R^{\mathrm{op}}}$, then*

$$\mathcal{S}_{\mathcal{D}^{\mathrm{d}}} = d\mathcal{S}_{\mathcal{D}}.$$

Hence

$$\mathbb{L}^{\mathrm{eq}+}(\mathcal{D}^{\mathrm{d}}) = (\mathcal{D}_{R^{\mathrm{op}}}^{\mathrm{c}})\text{-mod}/\mathcal{S}_{\mathcal{D}^{\mathrm{d}}} \simeq (\text{mod-}\mathcal{D}_R^{\mathrm{c}}/\mathcal{S}_{\mathcal{D}})^{\mathrm{op}} = (\mathbb{L}^{\mathrm{eq}+}(\mathcal{D}))^{\mathrm{op}}.$$

[25] One can set up duality at the level of pp formulas but it's duality of pp-pairs which we really need. Also see Section 4B for the issues re well-definedness/independence of enhancements which arise.

This is a special case of [Garkusha and Prest 2005, 7.4] which deals with the general case of pairs, $\mathcal{T}$, $\mathcal{T}_1$, of compactly generated triangulated categories with $\mathcal{T}_1^{\mathrm{c}} \simeq (\mathcal{T}^{\mathrm{c}})^{\mathrm{op}}$, also showing that, in this situation, we have a frame isomorphism between $\mathrm{Zg}(\mathcal{T})$ and $\mathrm{Zg}(\mathcal{T}_1)$.

It is shown in [Angeleri Hügel and Hrbek 2021] that, for derived categories of module categories, elementary duality has the same relation to algebraic Hom-dualities as in the case of definable subcategories of module categories. In [Bird and Williamson 2022] this is treated in a very general way and a variety of specific examples, from algebra and topology, are given.

4. Tensor-triangulated categories

Suppose now that the compactly generated triangulated category $\mathcal{T}$ has a monoidal, that is a tensor, structure. So we have $\otimes : \mathcal{T} \times \mathcal{T} \to \mathcal{T}$, which we assume to be commutative as well as associative, for which we have a tensor-unit $\mathbb{1}$ — so $\mathbb{1} \otimes X \simeq X$ for every $X \in \mathcal{T}$. *We assume $\otimes$ to be exact in each variable.* We drop explicit mention of associators et cetera; see for instance [Levine 1998, Part II] for more background.

We suppose that $\mathcal{T}$ is *rigidly-compactly generated*. That is, we assume in addition

- that the tensor structure is *closed*, meaning that there is an internal hom $[-,-] : \mathcal{T} \times \mathcal{T} \to \mathcal{T}$ which is right adjoint to $\otimes$: $(X \otimes Y, Z) \simeq (X, [Y, Z])$ for $X, Y, Z \in \mathcal{T}$, in particular $(Y, Z) \simeq (\mathbb{1}, [Y, Z])$; and,
- writing $X^\vee = [X, \mathbb{1}]$ for the *dual* of an object $X \in \mathcal{T}$, we assume that every compact object A is *rigid*, meaning that the natural map $A^\vee \otimes B \to [A, B]$ is an isomorphism for every $B \in \mathcal{T}^{\mathrm{c}}$.

It follows that $\mathcal{T}^{\mathrm{c}}$ is a *tensor-subcategory* of $\mathcal{T}$ (i.e., is closed under $\otimes$), that $(A^\vee)^\vee \simeq A$, that $A^\vee \otimes X \simeq [A, X]$ for $X \in \mathcal{T}$ and $A \in \mathcal{T}^{\mathrm{c}}$, and that the duality functor $(-)^\vee$ is exact (e.g., see [Stevenson 2018, §1, 2.12]).

The monoidal structure on $\mathcal{T}^{\mathrm{c}}$ induces, by Day convolution (see [Balmer et al. 2020, Appendix]), a right-exact monoidal structure on $\text{mod-}\mathcal{T}^{\mathrm{c}}$ and hence on $\text{Mod-}\mathcal{T}^{\mathrm{c}}$. By definition we have $y(A \otimes B) \simeq yA \otimes yB$ for $A, B \in \mathcal{T}^{\mathrm{c}}$ and, see [Balmer et al. 2020, A.14], the restricted Yoneda functor $y : \mathcal{T} \to \text{Mod-}\mathcal{T}^{\mathrm{c}}$ is monoidal. The duality (Theorem 2.6) between $\text{mod-}\mathcal{T}^{\mathrm{op}}$ and $\mathrm{Coh}(\mathcal{T})$ is monoidal if the latter is given the natural tensor structure (see [Wagstaffe 2021, §5.1]).

We say that a definable subcategory $\mathcal{D}$ of $\mathcal{T}$ is *tensor-closed* if, for every $X \in \mathcal{D}$ and $Y \in \mathcal{T}$, we have $X \otimes Y \in \mathcal{D}$. It is sufficient, see below, that this be so for every $Y \in \mathcal{T}^{\mathrm{c}}$. The theorem below says that this tensor-closed condition is equivalent to corresponding requirements on the associated data. We write $f \otimes A$ for $f \otimes \mathrm{id}_A$ if f is a morphism and A an object.

Theorem 4.1 [Wagstaffe 2021, 5.1.8]. *Suppose that $\mathcal{T}$ is a rigidly-compactly generated tensor-triangulated category. Then the following conditions on a definable subcategory $\mathcal{D}$ are equivalent*:

(i) *$\mathcal{D}$ is tensor-closed.*

(ii) *$X \in \mathcal{D}$ and $A \in \mathcal{T}^{\mathrm{c}}$ implies $X \otimes A \in \mathcal{D}$.*

(iii) *If $f \in \mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ and $A \in \mathcal{T}^{\mathrm{c}}$, then $f \otimes A \in \mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$.*

(iv) *The corresponding Serre subcategory $\mathcal{S}_{\mathcal{D}}$ of $\text{mod-}\mathcal{T}^{\mathrm{c}}$ is a tensor-ideal of $\text{mod-}\mathcal{T}^{\mathrm{c}}$ (it is enough that it be closed under tensoring with representable functors yA with $A \in \mathcal{T}^{\mathrm{c}}$).*

(v) *The corresponding Serre subcategory $\mathrm{Ann}_{\mathrm{Coh}(\mathcal{T})}(\mathcal{D}) = \mathcal{S}_{\mathcal{D}}^{\circ}$ of $\mathrm{Coh}(\mathcal{T})$ is a tensor-ideal of $\mathrm{Coh}(\mathcal{T})$ (it is enough that it be closed under tensoring with representable functors $(A, -)$ with $A \in \mathcal{T}^{\mathrm{c}}$).*

A stronger condition on a definable subcategory $\mathcal{D}$ of $\mathcal{T}$ is that it be a *tensor-ideal* of $\mathcal{T}$, meaning that it is tensor-closed and triangulated. The corresponding, in the sense of Theorem 4.1, annihilator ideals and Serre subcategories are characterised in [Wagstaffe 2021, 5.2.14]. The additional condition on $\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ is that it be exact and the additional condition on $\mathcal{S}_{\mathcal{D}}$ is that it be perfect; these conditions come from [Krause 2005]; see [Wagstaffe 2021, §5.2] for the detailed statements. Furthermore, the tensor version of Theorem 3.18 is true: the triangulated tensor-closed definable subcategories of $\mathcal{T}$ are in bijection, via torsion pairs, with the smashing tensor-ideals of $\mathcal{T}$ [Wagstaffe 2021, 5.2.14].

Wagstaffe [2021, Chapter 6] defined and investigated various coarsenings of the Ziegler topology on $\mathrm{pinj}(\mathcal{T})$, in particular, the tensor-closed Ziegler spectrum, $\mathrm{Zg}^{\otimes}(\mathcal{T})$, which is obtained by taking the closed subsets to be those of the form $\mathcal{D} \cap \mathrm{pinj}(\mathcal{T})$, where $\mathcal{D}$ is a tensor-closed definable subcategory of $\mathcal{T}$.

4A. ***Spectra in tensor-triangulated categories.*** A *prime* of the tensor-triangulated category $\mathcal{T}$ is a (thick) tensor-ideal $\mathcal{P}$ of $\mathcal{T}^{\mathrm{c}}$ such that if $A, B \in \mathcal{T}^{\mathrm{c}}$ and $A \otimes B \in \mathcal{P}$, then A or B is in $\mathcal{P}$. The *Balmer spectrum* [2005], $\mathrm{Spc}(\mathcal{T}^{\mathrm{c}})$ or just $\mathrm{Spc}(\mathcal{T})$, consists of these primes, with the topology which has, for a basis of open sets, those of the form

$$U(A) = \{\mathcal{P} \in \mathrm{Spc}(\mathcal{T}) : A \in \mathcal{P}\}$$

for $A \in \mathcal{T}^{\mathrm{c}}$. This is a spectral space and we may also consider, as in Section 3F, the Hochster-dual, or *Thomason*, topology on the same set, which is defined by declaring that the $U(A)$ generate, under finite union and arbitrary intersection, the *closed* sets. Both these topologies are natural and have their uses in various contexts; see, for instance, [Balmer 2020a].

There are various routes by which $\mathrm{Spc}(\mathcal{T})$ and inj-$\mathcal{T}^{\mathrm{c}}$, and also the homological spectrum, $\mathrm{Spc}^{\mathrm{h}}(\mathcal{T})$, from [Balmer 2020b], with their various topologies, may be connected; see in particular [Bird and Williamson 2023] and references therein. We also have the following.

To a point $\mathcal{P}$ of $\mathrm{Spc}(\mathcal{T})$ we can associate the finite type hereditary torsion theory $\gamma_{\mathcal{P}} = (\overrightarrow{\mathcal{S}_{y\mathcal{P}}}, (y\mathcal{P})^{\perp})$ on Mod-$\mathcal{T}^{\mathrm{c}}$(see Section 3E) whose torsion class is generated as such by $y\mathcal{P}$, that is, the torsion class is the $\varinjlim$-closure of the Serre subcategory $\mathcal{S}_{y\mathcal{P}}$ generated by $y\mathcal{P}$.

By [Balmer 2020b, 3.9] this gives an injection of the lattice of Balmer primes into the lattice of finite-type hereditary torsion theories, the latter ordered by inclusion of torsion classes. For, if $\mathcal{P} \subset \mathcal{Q}$ is a proper inclusion of Balmer primes, then, by Balmer's result, there is a maximal Serre tensor-ideal $\mathcal{B}$ of mod-$\mathcal{T}^{\mathrm{c}}$ such that $\mathcal{P} = y^{-1}\mathcal{B}$. Certainly $\mathcal{S}_{y\mathcal{P}} \subseteq \mathcal{B}$ so, if we had $\mathcal{S}_{y\mathcal{P}} = \mathcal{S}_{y\mathcal{Q}}$, then we would have $y\mathcal{Q} \subseteq \mathcal{B}$ and hence a contradiction.

Further, each finite type hereditary torsionfree class $\mathcal{F}$ is determined by its intersection with $\mathrm{inj}_{\mathcal{T}^{\mathrm{c}}}$, see [Prest 2009, 11.1.29], and the resulting sets $\mathcal{F} \cap \mathrm{inj}_{\mathcal{T}^{\mathrm{c}}}$ are the closed sets in the Ziegler topology on $\mathrm{inj}_{\mathcal{T}^{\mathrm{c}}}$ (see [Prest 2009, §14.1.3]). So, to a Balmer prime $\mathcal{P}$, we also have the associated Ziegler-closed set $(y\mathcal{P})^{\perp} \cap \mathrm{inj}_{\mathcal{T}^{\mathrm{c}}}$. Note that this association is inclusion-reversing.

If $A \in \mathcal{T}^{\mathrm{c}}$ then we have

$$\mathcal{P} \in U(A) \iff A \in \mathcal{P} \iff yA \in \overrightarrow{\mathcal{S}_{y\mathcal{P}}} \iff (y\mathcal{P})^{\perp} \subseteq (yA)^{\perp}.$$

The second equivalence is by the argument just made. Note that $(yA)^{\perp} \cap \mathrm{inj}_{\mathcal{T}^{\mathrm{c}}}$ is the complement of the basic Ziegler-open subset of $\mathrm{inj}_{\mathcal{T}^{\mathrm{c}}}$ that is defined by $(yA, -) \neq 0$, hence it is basic open in the dual-Ziegler topology.

For instance, if R is commutative noetherian, then the above essentially gives the embedding (see [Balmer 2005; Garkusha and Prest 2008]) of $\mathrm{Spc}(\mathcal{D}_R^{\mathrm{perf}})$ with the Thomason topology into the frame of Ziegler-open subsets of $\mathrm{Spec}(R)$, the latter being isomorphic, as a lattice, to the opposite of the lattice of finite type hereditary torsionfree classes of R-modules.

4B. ***Internal duality in tensor-triangulated categories.*** In [Wagstaffe 2021, Chapter 7] an *internal* duality for rigidly-compactly generated tensor-triangulated categories $\mathcal{T}$ is defined. In this respect it is somewhat similar to elementary duality in the case that R is a commutative ring, since then the categories of right and left R-modules are naturally identified and so, in that particular context, elementary duality is an internal duality on Mod-R. Indeed, for a commutative ring R and the derived-tensor structure on the derived category $\mathcal{D}_R$, this internal duality coincides with elementary duality, [Wagstaffe 2021, 7.3.5].

The internal duality for rigidly-compactly generated tensor-triangulated $\mathcal{T}$ comes from the second author's thesis [Wagstaffe 2021] and it was also discovered independently by Bird and Williamson [2022]. In [Wagstaffe 2021] it is defined in terms of cohomological ideals, Serre subcategories and definable subcategories; here we note that it can also be defined at the level of formulas and pp-pairs. We continue to assume that $\mathcal{T}$ is a rigidly-compactly generated tensor-triangulated category.

Just as for the "external" duality, we can define the duality using a hom functor to an object but, in this case, we use the internal hom functor: for $A \in \mathcal{T}^{\mathrm{c}}$, consider $A \mapsto [A, \mathbb{1}] \simeq A^\vee \otimes \mathbb{1} \simeq A^\vee$. Similarly, internal duality $(-)^\vee = [-, \mathbb{1}]$ applied to a morphism $A \xrightarrow{f} B$ in $\mathcal{T}^{\mathrm{c}}$ gives the morphism $B^\vee \xrightarrow{f^\vee} A^\vee$ in $\mathcal{T}^{\mathrm{c}}$. Since $\mathcal{T}$ is rigidly-compactly generated, we have that $(-)^\vee$ is an antiequivalence $(\mathcal{T}^{\mathrm{c}})^{\mathrm{op}} \simeq \mathcal{T}^{\mathrm{c}}$ with $(-)^{\vee\vee}$ naturally equivalent to the identity functor on $\mathcal{T}^{\mathrm{c}}$ (see [Stevenson 2018, 1.4]). We also apply these notations to arbitrary objects and morphisms of $\mathcal{T}$.

Given a definable subcategory $\mathcal{D}$ of $\mathcal{T}$, with associated annihilator ideal $\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$, we define its *internal dual* definable subcategory of $\mathcal{T}$ to be $\mathcal{D}^\vee = \mathrm{Ann}_{\mathcal{T}}(\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})^\vee)$, where we set $\mathcal{A}^\vee = \{f^\vee : f \in \mathcal{A}\}$ for $\mathcal{A}$ a collection of morphisms in $\mathcal{T}^{\mathrm{c}}$.

Proposition 4.2 (mostly [Wagstaffe 2021, §7.1]). *Suppose that $\mathcal{T}$ is a rigidly-compactly generated tensor-triangulated category, let $\mathcal{D}$ be a definable subcategory and consider its elementary dual definable subcategory $\mathcal{D}^\vee$. Then $(\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D}))^\vee$ is an annihilator ideal, $(\mathcal{D}^\vee)^\vee = \mathcal{D}$ and*

$$\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D}^\vee) = (\mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})^\vee), \quad \mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D}^\vee) = (\mathcal{D}\text{-TF})^\vee, \quad \mathcal{D}^\vee\text{-TF} = (\mathrm{Div}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D}))^\vee.$$

Proof. The proof is very similar to that of Proposition 3.21, using [Garkusha and Prest 2005, §7] to get the first statements. For the last two, consider $f \in \mathrm{Ann}_{\mathcal{T}^{\mathrm{c}}}(\mathcal{D})$ and form the extended triangle

$$\Sigma^{-1}B \xrightarrow{\Sigma^{-1}g} \Sigma^{-1}C \xrightarrow{\Sigma^{-1}h} A \xrightarrow{f} B \xrightarrow{g} C \xrightarrow{h} \Sigma A$$

then dualise it:

$$(\Sigma A)^\vee = \Sigma^{-1}A^\vee \xrightarrow{h^\vee} C^\vee \xrightarrow{g^\vee} B^\vee \xrightarrow{f^\vee} A^\vee \xrightarrow{\Sigma h^\vee} \Sigma C^\vee \xrightarrow{\Sigma g^\vee} \Sigma B^\vee.$$

Then apply equation (4) from Section 3A. □

This internal duality can also be given by a duality operation on pp formulas and pp-pairs. This is defined exactly as one would expect from the abelian/modules case. Namely, if $\phi(x_B)$, being $\exists x_{B'}\,(x_B f = x_{B'} f')$, is a typical pp formula, where $f : A \to B$ and $f' : A \to B'$ are in $\mathcal{T}^{\mathrm{c}}$, then we define the *dual* pp formula, $\phi^\vee(x_{B^\vee})$ to be $\exists y_{A^\vee}\,(y_{A^\vee} f^\vee = x_{B^\vee} \wedge y_{A^\vee} f'^\vee = 0_{B'^\vee})$. In particular, the dual of the pp formula $x_B f = 0$, where $f : A \to B$, is $f^\vee | x_{B^\vee}$ and the dual of $f' | x_B$ is $x_{B^\vee} f'^\vee = 0$.

The *dual* of a pp-pair ϕ/ψ is then defined to be $\psi^\vee/\phi^\vee$.

Note that what we have defined here is an internal duality on pp formulas in the language for (right) $\mathcal{T}^{\mathrm{c}}$-modules. There is a subtlety, which is pointed out in [Wagstaffe 2021]. Namely, two pp formulas might be equivalent on $\mathcal{T}$ — that is, have the same solution set on every object of $\mathcal{T}$ — yet their duals might not be equivalent. Indeed, we might have pp formulas ϕ, ϕ_1 with $\phi(X)=\phi_1(X)$ for every $X\in\mathcal{T}$, yet with $\phi^\vee(X)\neq\phi_1^\vee(X)$ perhaps even for every $X\in\mathcal{T}$ since these might be definable subgroups of distinct sorts — see [Wagstaffe 2021, Example 7.1.4]. Nevertheless $\phi^\vee$ and $\phi_1^\vee$ will define isomorphic coherent functors, meaning that the pairs $\phi^\vee(x)/(x=0)$ and $\phi_1^\vee(x_1)/(x_1=0)$ will be isomorphic in the category $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ of pp-imaginaries for $\mathcal{T}$. More generally, if ϕ/ψ is a pp-pair with ϕ_1 equivalent to ϕ and ψ_1 equivalent to ψ, then the pp-pairs $\psi^\vee/\phi^\vee$ and $\psi_1^\vee/\phi_1^\vee$ might be distinct but they will be isomorphic; in particular for every $X\in\mathcal{T}$, we will have $\psi^\vee(X)/\phi^\vee(X)=0$ if and only if $\psi_1^\vee(X)/\phi_1^\vee(X)=0$. That follows from [Garkusha and Prest 2005, 7.4], cf. Proposition 3.22, indeed it follows that there is an induced anti-isomorphism of the category $\mathbb{L}(\mathcal{T})^{\mathrm{eq}+}$ with itself.

We give some more detail; see also [Wagstaffe 2021, Chapter 7]. Since we have a duality $(-)^\vee:(\mathcal{T}^{\mathrm{c}})^{\mathrm{op}}\to\mathcal{T}^{\mathrm{c}}$ we have, by [Garkusha and Prest 2005, 7.4], an equivalence mod-$\mathcal{T}^{\mathrm{c}}\to\mathcal{T}^{\mathrm{c}}$-mod which is given by taking

$$G_f=\operatorname{coker}((-,f):(-,A)\to(-,B)),$$

where $f:A\to B$, to

$$F_{f^\vee}=\operatorname{coker}((f^\vee,-):(A^\vee,-)\to(B^\vee,-)).$$

We also have the duality $\mathcal{T}^{\mathrm{c}}$-mod$(\simeq\operatorname{Coh}(\mathcal{T}))\to(\text{mod-}\mathcal{T}^{\mathrm{c}})^{\mathrm{op}}$ which takes $F_{f^\vee}$ to $(F_{f^\vee})^\diamond:C\mapsto(F_{f^\vee},(C,-))$ for $C\in\mathcal{T}^{\mathrm{c}}$.

Composing these, we have a duality mod-$\mathcal{T}^{\mathrm{c}}\to$ mod-$\mathcal{T}^{\mathrm{c}}$ which takes G_f to $(F_{f^\vee})^\diamond$. In view of the exact sequence (3)

$$0\to(F_{f^\vee})^\diamond\to(-,B^\vee)\xrightarrow{(-,f^\vee)}(-,A^\vee)\to G_{f^\vee}\to 0$$

we can formulate this as follows.

Proposition 4.3. *Suppose $\mathcal{T}$ is a rigidly-compactly generated tensor-triangulated category. Then there is a duality on* mod-$\mathcal{T}^{\mathrm{c}}$ *which is given on objects by $G_f\mapsto$* $\ker(-,f^\vee)$, *where $(-)^\vee$ is the duality on $\mathcal{T}^{\mathrm{c}}$.*

The next result follows directly from [Bird and Williamson 2022, 6.12] (also [Angeleri Hügel and Hrbek 2021, 2.3] in the case $\mathcal{T}=\mathcal{D}_R$, R commutative).

Proposition 4.4. *Suppose $\mathcal{T}$ is a rigidly-compactly generated tensor-triangulated category and let $\mathcal{D}$ be a definable subcategory. Then the definable subcategory of $\mathcal{T}$ generated by the collection of objects $\{X^\vee : X\in\mathcal{D}\}$ is exactly the dual definable subcategory $\mathcal{D}^\vee$.*

There is potential ambiguity in the notation $\mathcal{D}^\vee$ — we have defined it to be the dual definable subcategory but it would also be a natural notation for $\{X^\vee : X \in \mathcal{D}\}$ but the latter, a subclass of $\mathcal{D}^\vee$, is not in general all of the definable category $\mathcal{D}^\vee$ (it might not be closed under pure subobjects).

Tensor-closed definable subcategories are self-dual.

Theorem 4.5 [Wagstaffe 2021, 7.2.2]. *If $\mathcal{D}$ is a tensor-closed definable subcategory of a rigidly-compactly generated tensor-triangulated category, then $\mathcal{D}$ is self-dual: $\mathcal{D}^\vee = \mathcal{D}$.*

4C. *Internal Hom interpretation.* We finish by pointing out some more ideals of $\mathcal{T}^c$ associated to a definable category $\mathcal{D}$ in the rigidly-compactly generated tensor-triangulated context. They appear (along with their rather provisional names) in the statement of the next result.

Proposition 4.6. *Suppose $\mathcal{T}$ is a rigidly-compactly generated tensor-triangulated category and let $\mathcal{X} \subset \mathcal{T}$. We define the **tensor-annihilator** of $\mathcal{X}$:*

$$\otimes\text{-ann}_{\mathcal{T}^c}\mathcal{X} = \{f : a \to b \in \mathcal{T}^c : f \otimes X = 0 : a \otimes X \to b \otimes X \ \forall X \in \mathcal{X}\},$$

*the **internal-hom-annihilator** of $\mathcal{X}$:*

$$[\text{ann}]_{\mathcal{T}^c}\mathcal{X} = \{f : a \to b \in \mathcal{T}^c : [f, X] = 0 : [b, X] \to [a, X] \ \forall X \in \mathcal{X}\},$$

*the **tensor phantomiser** of $\mathcal{X}$:*

$$\otimes\text{-phan}_{\mathcal{T}^c}\mathcal{X} = \{f : a \to b \in \mathcal{T}^c : f \otimes X : a \otimes X \to b \otimes X \textit{ is phantom } \forall X \in \mathcal{X}\},$$

*and the **internal-hom-phantomiser** of $\mathcal{X}$:*

$$[\text{phan}]_{\mathcal{T}^c}\mathcal{X} = \{f : a \to b \in \mathcal{T}^c : [f, X] : [b, X] \to [a, X] \textit{ is phantom } \forall X \in \mathcal{X}\}.$$

All these are ideals of $\mathcal{T}^c$ and the tensor-annihilator and internal-hom-annihilator are dual ideals:

$$(\otimes\text{-ann}_{\mathcal{T}^c}\mathcal{X})^\vee = [\text{ann}]_{\mathcal{T}^c}\mathcal{X}.$$

*Moreover, the tensor phantomiser and internal-hom-phantomiser coincide (we could call this the **phantomiser**) and are equal to the annihilator ideal of the smallest tensor-closed definable subcategory $\langle \mathcal{X} \rangle^\otimes$ containing $\mathcal{X}$:*

$$\otimes\text{-phan}_{\mathcal{T}^c}\mathcal{X} = [\text{phan}]_{\mathcal{T}^c}\mathcal{X} = \text{Ann}_{\mathcal{T}^c} \langle \mathcal{X} \rangle^\otimes.$$

Thus this is also the annihilator ideal generated by each of $\otimes\text{-ann}_{\mathcal{T}^c}\mathcal{X}$ and $[\text{ann}]_{\mathcal{T}^c}\mathcal{X}$.

Proof. For every $X \in \mathcal{T}$, $A \otimes X \xrightarrow{f \otimes X} B \otimes X$ is (isomorphic to) $A^{\vee\vee} \otimes X \xrightarrow{f^{\vee\vee}} B^{\vee\vee} \otimes X$ and hence is $[A^\vee, X] \xrightarrow{[f^\vee, X]} [B^\vee, X]$. Thus, the condition $f \otimes X = 0 : A \otimes X \to B \otimes X$ is equivalent to the condition $[f^\vee, X] = 0 : [A^\vee, X] \to [B^\vee, X]$ and we have $\otimes\text{-ann}_{\mathcal{T}^c}\mathcal{X} = ([\text{ann}]_{\mathcal{T}}\mathcal{X})^\vee$.

For the other parts, we have $f \in \otimes\text{-phan}_{\mathcal{T}^c}\mathcal{X}$ if and only if for every $c \in \mathcal{T}^c$ we have $(c, f \otimes X) = 0$, that is $(f^\vee, c^\vee \otimes X) = 0$ which, since every compact object is a dual, is equivalent to $f^\vee \in \otimes\text{-ann}_{\mathcal{T}^c} \langle \mathcal{X} \rangle^\otimes$. By Theorem 4.5, $f^\vee \in \otimes\text{-ann}_{\mathcal{T}^c} \langle \mathcal{X} \rangle^\otimes$ if and only if $f \in \otimes\text{-ann}_{\mathcal{T}^c} \langle \mathcal{X} \rangle^\otimes$. Therefore $\otimes\text{-phan}_{\mathcal{T}^c}\mathcal{X} = \otimes\text{-ann}_{\mathcal{T}^c} \langle \mathcal{X} \rangle^\otimes$.

Also, $f \in [\text{phan}]_{\mathcal{T}^c}\mathcal{X}$ if and only if for every $c \in \mathcal{T}^c$ we have $(c, [f, X]) = 0$, equivalently $(f, c^\vee \otimes X) = 0$ which, since every compact object is a dual, is equivalent to $f \in \text{ann}_{\mathcal{T}^c} \langle \mathcal{X} \rangle^\otimes$. Therefore $[\text{phan}]_{\mathcal{T}^c}\mathcal{X} = \text{ann}_{\mathcal{T}^c} \langle \mathcal{X} \rangle^\otimes = \otimes\text{-phan}_{\mathcal{T}^c}\mathcal{X}$, as claimed. □

Note that the condition $f^\vee \in [\text{ann}]_{\mathcal{T}}\mathcal{X}$ is expressed by the condition "$Xf^\vee = 0$" with $B^\vee \xrightarrow{f^\vee} A^\vee$. This looks like an annihilator sentence but it is for internal hom, rather than actual hom, groups. This suggests an alternative, internal-hom, interpretation of the model-theoretic language (Remark 2.1) when $\mathcal{T}$ is a rigidly-compactly generated tensor-triangulated category. In this interpretation the value of $X \in \mathcal{T}$ at sort $A \in \mathcal{T}^c$ is $[A, X]$, rather than (A, X), and the interpretation of $A \xrightarrow{f} B \in \mathcal{T}^c$ in X is $[f, X] : [B, X] \to [A, X]$ rather than $(f, X) : (B, X) \to (A, X)$. In this interpretation of the language the values of sorts at objects of $\mathcal{T}$ are again objects of $\mathcal{T}$, not abelian groups.

This also constitutes an alternative "internal restricted Yoneda" functor from $\mathcal{T}$ to the "$\mathcal{T}$-valued-module category" $\text{Mod}_{\mathcal{T}}\text{-}\mathcal{T}^c = ((\mathcal{T}^c)^{\text{op}}, \mathcal{T})$, which takes $X \in \mathcal{T}$ to the functor $[-, X] : (\mathcal{T}^c)^{\text{op}} \to \mathcal{T}$ and takes $f : X \to Y$ to $[-, f] : [-, X] \to [-, Y]$. In this internal-hom interpretation, the language for $\mathcal{T}$ stays the same but the interpretation has changed: instead of $(-, X)$ we use $[-, X]$.

Similarly, the tensor-annihilator that we defined above belongs to a third (in this case, covariant) interpretation of the same language, based on $- \otimes X$, rather than $(-, X)$ or $[-, X]$.

In both these new interpretations the sorts belong to $\mathcal{T}$ rather than to **Ab**, so we cannot immediately make sense of "elements" of a sort. But, using the idea of an "element" being an arrow from the tensor-unit $\mathbb{1}$, we can move back to the category of $\mathcal{T}^c$-modules. If we do that, we recover the usual interpretation (from the internal-hom interpretation) and an "internal dual" interpretation (from the tensor interpretation). That is, we have:

$$\begin{aligned} y &: \mathcal{T} \to \text{Mod-}\mathcal{T}^c && \text{given by} \quad X \mapsto (-, X); \\ [y] &: \mathcal{T} \to ((\mathcal{T}^c)^{\text{op}}, \mathcal{T}) && \text{given by} \quad X \mapsto [-, X]; \\ \epsilon &: \mathcal{T} \to (\mathcal{T}^c, \mathcal{T}) && \text{given by} \quad X \mapsto (- \otimes X). \end{aligned}$$

The latter two can then be composed with $(\mathbb{1}, -)$:

$$(\mathbb{1}, -)[y] = y : \mathcal{T} \to ((\mathcal{T}^c)^{\text{op}}, \mathcal{T}) \to \text{Mod-}\mathcal{T}^c$$
$$\text{given by} \quad X \mapsto [-, X] \mapsto (\mathbb{1}, [-, X]) \simeq (-, X);$$

and

$$(\mathbb{1}, -)\epsilon : \mathcal{T} \to (\mathcal{T}^{\mathrm{c}}, \mathcal{T}) \to \mathcal{T}^{\mathrm{c}}\text{-Mod}$$

$$\text{given by} \quad X \mapsto (- \otimes X) \mapsto (\mathbb{1}, - \otimes X) \simeq (\mathbb{1}, [(-)^{\vee}, X]) \simeq ((-)^{\vee}, X)$$

Also, essentially following [Bird and Williamson 2023, 4.13], note that if $A \in \mathcal{T}^{\mathrm{c}}$ and $X \in \mathcal{T}$, then $[A, X] = 0$ if and only if, for all $C \in \mathcal{T}^{\mathrm{c}}$, we have $(C, [A, X]) = 0$ if and only if, for all $C \in \mathcal{T}^{\mathrm{c}}$, we have $(C \otimes A, X) = 0$. In particular

$$\{N \in \mathrm{Zg}(\mathcal{T}) : [A, N] = 0\} = \bigcap_{C \in \mathcal{T}^{\mathrm{c}}} \{N \in \mathrm{Zg}(\mathcal{T}) : (C \otimes A, N) = 0\}$$

is an intersection of Ziegler-closed sets, hence is itself Ziegler-closed.

Furthermore, continuing the above computation, we have $[A, X] = 0$ if and only if, for all $C \in \mathcal{T}^{\mathrm{c}}$, we have $(A \otimes C, X) = 0$ if and only if, for all $C \in \mathcal{T}^{\mathrm{c}}$, we have $(A, [C, X]) = 0$ if and only if, for all $C \in \mathcal{T}^{\mathrm{c}}$, we have $(A, C^{\vee} \otimes X) = 0$, if and only if, for all $C \in \mathcal{T}^{\mathrm{c}}$, we have $(A, C \otimes X) = 0$. So if $\mathcal{D}$ is the definable subcategory of $\mathcal{T}$ cut out by the condition $(A, -) = 0$, then the condition $[A, -] = 0$ cuts out the smallest tensor-closed definable subcategory of $\mathcal{T}$ containing $\mathcal{D}$.

References

[Adámek and Rosický 1994] J. Adámek and J. Rosický, *Locally presentable and accessible categories*, London Mathematical Society Lecture Note Series **189**, Cambridge University Press, 1994. MR Zbl

[Angeleri Hügel and Hrbek 2021] L. Angeleri Hügel and M. Hrbek, "Parametrizing torsion pairs in derived categories", *Represent. Theory* **25** (2021), 679–731. MR Zbl

[Arnesen et al. 2017] K. K. Arnesen, R. Laking, D. Pauksztello, and M. Prest, "The Ziegler spectrum for derived-discrete algebras", *Adv. Math.* **319** (2017), 653–698. MR Zbl

[Balmer 2005] P. Balmer, "The spectrum of prime ideals in tensor triangulated categories", *J. Reine Angew. Math.* **588** (2005), 149–168. MR Zbl

[Balmer 2020a] P. Balmer, "A guide to tensor-triangular classification", pp. 145–162 in *Handbook of homotopy theory*, edited by H. Miller, CRC Press, Boca Raton, FL, 2020. MR Zbl

[Balmer 2020b] P. Balmer, "Nilpotence theorems via homological residue fields", *Tunis. J. Math.* **2**:2 (2020), 359–378. MR Zbl

[Balmer et al. 2020] P. Balmer, H. Krause, and G. Stevenson, "The frame of smashing tensor-ideals", *Math. Proc. Cambridge Philos. Soc.* **168**:2 (2020), 323–343. MR Zbl

[Beligiannis 2000a] A. Beligiannis, "On the Freyd categories of an additive category", *Homology Homotopy Appl.* **2** (2000), 147–185. MR Zbl

[Beligiannis 2000b] A. Beligiannis, "Relative homological algebra and purity in triangulated categories", *J. Algebra* **227**:1 (2000), 268–361. MR Zbl

[Bennett-Tennenhaus 2023] R. Bennett-Tennenhaus, "Characterisations of Σ-pure-injectivity in triangulated categories and applications to endocoperfect objects", *Fund. Math.* **261**:2 (2023), 133–155. MR Zbl

[Benson and Gnacadja 1999] D. J. Benson and G. P. Gnacadja, "Phantom maps and purity in modular representation theory, I", pp. 37–91 in *Algebraic topology* (Kazimierz Dolny, 1997), vol. 161, edited by W. Dwyer and S. Jackowski, Fund. Math. **1-2**, 1999. MR Zbl

[Bird and Williamson 2022] I. Bird and J. Williamson, "Duality pairs, phantom maps, and definability in triangulated categories", 2022. arXiv 2202.08113

[Bird and Williamson 2023] I. Bird and J. Williamson, "The homological spectrum via definable subcategories", 2023. arXiv 2304.05179

[Burke and Prest 2002] K. Burke and M. Prest, "The Ziegler and Zariski spectra of some domestic string algebras", *Algebr. Represent. Theory* **5**:3 (2002), 211–234. MR Zbl

[Christensen and Strickland 1998] J. D. Christensen and N. P. Strickland, "Phantom maps and homology theories", *Topology* **37**:2 (1998), 339–364. MR Zbl

[Eklof and Sabbagh 1971] P. Eklof and G. Sabbagh, "Model-completions and modules", *Ann. Math. Logic* **2**:3 (1971), 251–295. MR Zbl

[Gabriel 1962] P. Gabriel, "Des catégories abéliennes", *Bull. Soc. Math. France* **90** (1962), 323–448. MR Zbl

[Garkusha and Prest 2004] G. Garkusha and M. Prest, "Injective objects in triangulated categories", *J. Algebra Appl.* **3**:4 (2004), 367–389. MR Zbl

[Garkusha and Prest 2005] G. Garkusha and M. Prest, "Triangulated categories and the Ziegler spectrum", *Algebr. Represent. Theory* **8**:4 (2005), 499–523. MR Zbl

[Garkusha and Prest 2008] G. Garkusha and M. Prest, "Classifying Serre subcategories of finitely presented modules", *Proc. Amer. Math. Soc.* **136**:3 (2008), 761–770. MR Zbl

[Han 2013] Z. Han, *The homotopy categories of injective modules of derived discrete algebras*, Ph.D. thesis, Bielefeld University, 2013, https://pub.uni-bielefeld.de/record/2610302.

[Herzog 1993] I. Herzog, "Elementary duality of modules", *Trans. Amer. Math. Soc.* **340**:1 (1993), 37–69. MR Zbl

[Hovey 1999] M. Hovey, *Model categories*, Mathematical Surveys and Monographs **63**, American Mathematical Society, Providence, RI, 1999. MR Zbl

[Krause 2000] H. Krause, "Smashing subcategories and the telescope conjecture — an algebraic approach", *Invent. Math.* **139**:1 (2000), 99–133. MR Zbl

[Krause 2002] H. Krause, "Coherent functors in stable homotopy theory", *Fund. Math.* **173**:1 (2002), 33–56. MR Zbl

[Krause 2005] H. Krause, "Cohomological quotients and smashing localizations", *Amer. J. Math.* **127**:6 (2005), 1191–1246. MR Zbl

[Krause 2022] H. Krause, *Homological theory of representations*, Cambridge Studies in Advanced Mathematics **195**, Cambridge University Press, 2022. MR Zbl

[Krause and Letz 2023] H. Krause and J. C. Letz, "The spectrum of a well-generated tensor-triangulated category", *Bull. Lond. Math. Soc.* **55**:2 (2023), 680–705. MR Zbl

[Kucera and Prest 1992] T. G. Kucera and M. Prest, "Imaginary modules", *J. Symbolic Logic* **57**:2 (1992), 698–723. MR Zbl

[Laking 2020] R. Laking, "Purity in compactly generated derivators and t-structures with Grothendieck hearts", *Math. Z.* **295**:3-4 (2020), 1615–1641. MR Zbl

[Laking and Vitória 2020] R. Laking and J. Vitória, "Definability and approximations in triangulated categories", *Pacific J. Math.* **306**:2 (2020), 557–586. MR Zbl

[Levine 1998] M. Levine, *Mixed motives*, Mathematical Surveys and Monographs **57**, American Mathematical Society, Providence, RI, 1998. MR Zbl

[Neeman 2001] A. Neeman, *Triangulated categories*, Annals of Mathematics Studies **148**, Princeton University Press, 2001. MR Zbl

[Oberst and Röhrl 1970] U. Oberst and H. Röhrl, "Flat and coherent functors", *J. Algebra* **14** (1970), 91–105. MR Zbl

[Prest 1979] M. Prest, *Applications of logic to torsion theories in abelian categories*, Ph.D. thesis, University of Leeds, 1979.

[Prest 1988a] M. Prest, "Duality and pure-semisimple rings", *J. London Math. Soc.* (2) **38**:3 (1988), 403–409. MR Zbl

[Prest 1988b] M. Prest, *Model theory and modules*, London Mathematical Society Lecture Note Series **130**, Cambridge University Press, 1988. MR Zbl

[Prest 1993] M. Prest, "Remarks on elementary duality", *Ann. Pure Appl. Logic* **62**:2 (1993), 185–205. MR Zbl

[Prest 2009] M. Prest, *Purity, spectra and localisation*, Encyclopedia of Mathematics and its Applications **121**, Cambridge University Press, 2009. MR Zbl

[Prest 2011a] M. Prest, *Definable additive categories: purity and model theory*, Mem. Amer. Math. Soc. **987**, Amer. Math. Soc., Providence, RI, 2011. MR Zbl

[Prest 2011b] M. Prest, "Model theory in additive categories", pp. 231–244 in *Models*, *logics*, *and higher-dimensional categories*, edited by B. Hart et al., CRM Proc. Lecture Notes **53**, Amer. Math. Soc., Providence, RI, 2011. MR Zbl

[Prest 2012a] M. Prest, "Abelian categories and definable additive categories", 2012. arXiv 1202.0426

[Prest 2012b] M. Prest, "Categories of imaginaries for definable additive categories", 2012. arXiv 1202.0427

[Prest 2012c] M. Prest, "Spectra of small abelian categories", 2012. arXiv 1202.0431

[Prest 2019] M. Prest, "Multisorted modules and their model theory", pp. 115–151 in *Model theory of modules*, *algebras and categories*, edited by A. Facchini et al., Contemp. Math. **730**, Amer. Math. Soc., Providence, RI, 2019. MR Zbl

[Prest $\geq$ 2024] M. Prest, "Model theory of modules", https://personalpages.manchester.ac.uk/staff/mike.prest/publications.html. To appear in *Research Trends in Contemporary Logic*.

[Prest and Rajani 2010] M. Prest and R. Rajani, "Structure sheaves of definable additive categories", *J. Pure Appl. Algebra* **214**:8 (2010), 1370–1383. MR Zbl

[Roos 1961] J.-E. Roos, "Sur les foncteurs dérivés de $\varprojlim$, applications", *C. R. Acad. Sci. Paris* **252** (1961), 3702–3704. MR Zbl

[Sabbagh 1970] G. Sabbagh, "Aspects logiques de la pureté dans les modules", *C. R. Acad. Sci. Paris Sér. A-B* **271** (1970), A909–A912. MR Zbl

[Saorín and Šťovíček 2023] M. Saorín and J. Šťovíček, "t-structures with Grothendieck hearts via functor categories", *Selecta Math.* (*N.S.*) **29**:5 (2023), art. id. 77. MR Zbl

[Stenström 1975] B. Stenström, *Rings of quotients*: *an introduction to methods of ring theory*, Grundl. Math. Wissen. **217**, Springer, 1975. MR Zbl

[Stevenson 2018] G. Stevenson, "A tour of support theory for triangulated categories through tensor triangular geometry", pp. 63–101 in *Building bridges between algebra and topology*, edited by D. Herbera et al., Springer, 2018. MR Zbl

[Šťovíček and Pospíšil 2016] J. Šťovíček and D. Pospíšil, "On compactly generated torsion pairs and the classification of co-t-structures for commutative noetherian rings", *Trans. Amer. Math. Soc.* **368**:9 (2016), 6325–6361. MR Zbl

[Wagstaffe 2021] R. Wagstaffe, *Definability in monoidal additive and tensor triangulated categories*, Ph.D. thesis, University of Manchester, 2021, https://personalpages.manchester.ac.uk/staff/mike.prest/publications.html.

[Weibel 1994] C. A. Weibel, *An introduction to homological algebra*, Cambridge Studies in Advanced Mathematics **38**, Cambridge University Press, 1994. MR Zbl

[Ziegler 1984] M. Ziegler, "Model theory of modules", *Ann. Pure Appl. Logic* **26**:2 (1984), 149–213. MR Zbl

Received 17 May 2023. Revised 8 Feb 2024.

MIKE PREST:

mprest@manchester.ac.uk
Department of Mathematics, University of Manchester, Manchester, United Kingdom

ROSE WAGSTAFFE:

rose.wagstaffe@manchester.ac.uk
Department of Mathematics, University of Manchester, Manchester, United Kingdom

Guidelines for Authors

Authors may submit manuscripts in PDF format on-line at the submission page.

Originality. Submission of a manuscript acknowledges that the manuscript is original and and is not, in whole or in part, published or under consideration for publication elsewhere. It is understood also that the manuscript will not be submitted elsewhere while under consideration for publication in this journal.

Language. Articles are usually in English or French, but articles written in other languages are welcome.

Required items. A brief abstract of about 150 words or less must be included. It should be self-contained and not refer to bibliography keys. If the article is not in English, two versions of the abstract must be included, one in the language of the article and one in English. Also required are keywords and a Mathematics Subject Classification for the article, and, for each author, affiliation (if appropriate) and email address.

Format. Authors are encouraged to use LaTeX and the standard amsart class, but submissions in other varieties of TeX, and exceptionally in other formats, are acceptable. Initial uploads should normally be in PDF format; after the refereeing process we will ask you to submit all source material.

References. Bibliographical references should be complete, including article titles and page ranges. All references in the bibliography should be cited in the text. The use of BibTeX is preferred but not required. Tags will be converted to the house format, however, for submission you may use the format of your choice. Links will be provided to all literature with known web locations and authors are encouraged to provide their own links in addition to those supplied in the editorial process.

Figures. Figures must be of publication quality. After acceptance, you will need to submit the original source files in vector graphics format for all diagrams in your manuscript: vector EPS or vector PDF files are the most useful.

Most drawing and graphing packages — Mathematica, Adobe Illustrator, Corel Draw, MATLAB, etc. — allow the user to save files in one of these formats. Make sure that what you are saving is vector graphics and not a bitmap. If you need help, please write to graphics@msp.org with as many details as you can about how your graphics were generated.

Bundle your figure files into a single archive (using zip, tar, rar or other format of your choice) and upload on the link you been provided at acceptance time. Each figure should be captioned and numbered so that it can float. Small figures occupying no more than three lines of vertical space can be kept in the text ("the curve looks like this:"). It is acceptable to submit a manuscript with all figures at the end, if their placement is specified in the text by means of comments such as "Place Figure 1 here". The same considerations apply to tables.

White Space. Forced line breaks or page breaks should not be inserted in the document. There is no point in your trying to optimize line and page breaks in the original manuscript. The manuscript will be reformatted to use the journal's preferred fonts and layout.

Proofs. Page proofs will be made available to authors (or to the designated corresponding author) at a Web site in PDF format. Failure to acknowledge the receipt of proofs or to return corrections within the requested deadline may cause publication to be postponed.

Model Theory

no. 1 vol. 3 2024