Fun with Fields

by

William Andrew Johnson

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Thomas Scanlon, Chair
Professor Leo A. Harrington
Professor Paul Vojta
Associate Professor Lara Buchak

Spring 2016

Fun with Fields

Abstract

Fun with Fields

by

William Andrew Johnson

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Thomas Scanlon, Chair

This dissertation is a collection of results in model theory, related in one way or another to fields, NIP theories, and elimination of imaginaries. The most important result is a classification of dp-minimal fields, presented in Chapter 9. We construct in a canonical fashion a non-trivial Hausdorff definable field topology on any unstable dp-minimal field. Using this we classify the dp-minimal pure fields and valued fields up to elementary equivalence. Furthermore we prove that every VC-minimal field is real closed or algebraically closed.

In Chapter 11, we analyze the theories of existentially closed fields with several valuations and orderings, as studied by van den Dries [16]. We show that these model complete theories are NTP_2, and analyze forking, dividing, and burden in these theories. The theory of algebraically closed fields with n independent valuation rings turns out to be an example of such a theory. This provides a new and natural example of an NTP_2 theory which is neither simple nor NIP, nor even a conceptual hybrid of something simple and something NIP.

In Chapter 8, we exhibit a bad failure of elimination of imaginaries in a dense o-minimal structure. We produce an exotic interpretable set which cannot be put in definable bijection with a definable set, after naming any amount of parameters. However, we show that these exotic interpretable sets are still amenable to some of the tools of tame topology: they must admit nice definable topologies locally homeomorphic to definable sets.

Chapter 12 proves the existence of $\mathbb{Z}/n\mathbb{Z}$-valued definable strong Euler characteristics on pseudofinite fields, which measure the non-standard "size" of definable sets, mod n. The non-trivial result is that these "sizes" are definable in families of definable sets. This could probably be proven using etale cohomology, but we give a more elementary proof relying heavily on the theory of abelian varieties.

We also present simplified and new proofs of several model-theoretic facts, including the definability of irreducibility and Zariski closure in ACF (Chapter 10), and elimination of imaginaries in ACVF (Chapter 6). This latter fact was originally proven by Haskell, Hrushovski, and Macpherson [26]. We give a proof that is drastically simpler, inspired by Poizat's proofs of elimination of imaginaries in ACF and DCF [60].

To my parents

who showed the definition of “field” to a ten-year old.

Contents

Introduction

This dissertation is a collection of results in model theory, loosely united around three recurring themes:

- Fields, and valued fields in particular
- Theories without the independence property (NIP theories)
- Elimination of imaginaries

The most important result is a classification of dp-minimal fields, presented in Chapter 9. Recall that *dp-rank* is a notion of "rank" or "dimension" that is particularly well-suited for NIP theories. Structures of dp-rank 1 are *dp-minimal*, and dp-minimality is known to generalize many of the prior notions of "minimality" such as strong minimality, o-minimality, C-minimality, VC-minimality, and p-minimality.

A number of theories of fields are known to be dp-minimal, including algebraically closed fields, real closed fields, and several theories of valued fields. Our classification result shows that every field of dp-rank 1 is algebraically closed, real closed, or a henselian valued field with a non-trivial definable valuation. Even better, we can essentially classify the dp-minimal pure fields up to elementary equivalence. Every dp-minimal field is elementarily equivalent to the underlying field of a henselian valued field whose residue field is algebraically closed, real closed, or finite. Specializing to the case of VC-minimal fields, we also prove the remarkable fact that every VC-minimal field is real closed or algebraically closed. We also obtain characterizations of which valued fields are dp-minimal.

The classification is carried out in Chapter 9. Some preliminary facts about strongly dependent valued fields are proven in Chapter 4, building off work of Kaplan, Scanlon, and Wagner [42]. The positive side of the classification—the proof that certain fields *are* dp-minimal—is carried out in Chapter 3.

The classification of dp-minimal fields is proved through a surprising method: on any unstable dp-minimal field, we produce a completely canonical field topology from scratch. This is a rare instance in model theory where a topology can be created out of purely combinatorial assumptions. More typically, a topology is exactly what's *missing*, when we try to classify things in model theory. (Hrushovski and Zilber's work on Zariski geometries [37] is a good example of this phenomenon.)

The second substantial result of this dissertation is in Chapter 11, which focuses on theories of fields with several independent valuations and orderings, in the spirit of Lou van den Dries's thesis [16]. For fixed n, m, one can look at the theory of fields with n orderings and m valuations. By [16], this theory has a model companion. We prove that these model companions are NTP_2—that is, they do not have the second tree property. NTP_2 is a common generalization of "NIP" and "simple", that has been studied extensively in recent years by Chernikov, Kaplan, and others ([9], [10], [11], [12], [71]). We also analyze forking and dividing in these theories.

In certain cases, the model companions admit very simple axioms. For example, we show that the model companion of the theory of fields with n valuations is the theory of algebraically closed fields with n pairwise independent valuations. (Interestingly, the proof uses the Riemann-Roch theorem.) The fact that this is NTP_2 implies the surprising result that $\mathbb{Q}^{alg}$ is NTP_2 after being expanded with a separate unary predicate for every valuation ring.

Chapter 11 thus provides a new, natural example of an NTP_2 theory which is neither simple nor NIP. Nor is it a conceptual hybrid of something simple and something NIP, unlike many of the prior examples (e.g. ACVF with non-standard Frobenius, bounded pseudo real-closed fields [54], and ultraproducts of $\mathbb{Q}_p$ [12]). There is, however, some overlap with the bounded PRC case: existentially closed fields with several orderings turn out to be bounded PRC fields, which were shown to be NTP_2 by Samaria Montenegro [54].

Another key part of this dissertation is Chapter 6, which presents a new and improved proof of elimination of imaginaries in ACVF, the theory of algebraically closed valued fields. Haskell, Hrushovski, and Macpherson showed in their foundational paper [26] that ACVF eliminates imaginaries after expansion by certain "geometric sorts." This result underlies much of the later work on ACVF, including the sequel paper [27] on stable domination, and the work of Hrushovski and Loeser [34] on Berkovich spaces. Chapter 6 presents a new proof of elimination of imaginaries in ACVF, that is drastically shorter than the original proof in [26]. The proof is also conceptually simpler, and more aligned with Poizat's original proofs of elimination of imaginaries in ACF and DCF [60].

Continuing in the spirit of simplification,

- Chapter 2 presents a proof of quantifier elimination and C-minimality in ACVF that attempts to minimize the amount of syntax and algebraic casework. Quantifier elimination is essentially reduced to proving two facts: (1) C-minimality for quantifier-free definable sets (which is straightforward, if tedious), and (2) an amalgamation theorem for valued fields (which admits a direct algebraic proof).

- Chapter 7 recovers some of the basic facts about generically stable types and stable domination proven in [27] and [34]. Interestingly, the proofs use very little from ACVF, and generalize to C-minimal expansions of ACVF.

- Chapter 10 gives a fast new proof that irreducibility of varieties is definable in ACF. This proof avoids the use of computational algebraic geometry, flatness, and ultraproducts, relying instead on the definability of types in stable theories, and simple dimension counting arguments.

None of these results are new, but we hope there is something novel in the proofs.

In Chapter 8, we turn away from valued fields, and focus on elimination of imaginaries in dense o-minimal structures. We give a bizarre example of a (dense) o-minimal structure containing an interpretable set which cannot be put in definable bijection with a definable set, even after naming parameters. This phenomenon came as a surprise to many, including the

author. For reference, most o-minimal structures one encounters in the wild have elimination of imaginaries and definable skolem functions. The remainder of Chapter 8 is an attempt to salvage some of the topological tools of o-minimality for interpretable sets. We show that every interpretable set admits a definable topology with finitely many connected components, that looks everywhere locally like an open definable set. This could potentially lead to an interesting theory of "interpretable manifolds."

In the final chapter, we turn away from the world of NIP structures, and focus on pseudofinite fields—the infinite fields which have all the first-order properties of finite fields. Pseudofinite fields have been extensively studied by model theorists, due to their connections with diophantine geometry and ACFA. Chapter 12 considers nonstandard "sizes" of sets, *mod* N. The main result of the chapter is that these non-standard sizes are $\mathrm{acl}(\emptyset)$-definable in families. These sizes could probably be computed by examining the traces of the action of non-standard Frobenius on ℓ-adic cohomology. We will give a more elementary proof, however. It is the hope of the author that combinatorial arguments using these non-standard "sizes" could be used to prove something non-trivial in number theory.

Acknowledgments

As a Christian, I hold that

> Christ is the image of the invisible God,
> the firstborn of all creation.
> For by him all things were created
> in the heavens and on the earth,
> visible things and invisible things,
> whether thrones or dominions
> or principalities or powers.
> All things have been created through him and for him.
> He is before all things,
> and in him all things are held together.

For the last five years, I have been blessed with the opportunity to explore part of the invisible side of creation. I am grateful to live in a universe with such an interesting provability predicate.

During this time, I have received abundant support and encouragement, and many thanks are in order. My advisor Tom Scanlon has been an invaluable resource in providing feedback and helping me focus the direction of my research. His classes on stability theory, o-minimality, and valued fields inspired my most productive research projects. I am also grateful for my friends in Veritas and roommates, who helped me stay sane, made my time in Berkeley an enjoyable experience, and were there for me in difficult times. Last but not least, I am grateful for my fellow model theory students at Berkeley, and my other model theory friends and colleagues, including but not limited to Reid Dale, Jim Freitag, Alex Kruckman, Nick Ramsey, Silvain Rideau, Adam Topaz, and Michael Wan. In my view, it is the community of mathematicians that makes mathematics fun and worthwhile.

Three years of my research were made possible by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1106400. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

It was a strange sequence of events that led me to study model theory at Berkeley. Many people played important roles in bringing me here, and I owe them a debt of gratitude. In chronological order, these include my parents, who encouraged me to learn and to pursue my own studies outside of classes and school; Josh Horowitz who told me about ultrafilters and compactness at MOSP in 2004; Jim Morrow who encouraged me to pursue mathematics in college, convinced me to apply to mathematics graduate school, and assisted greatly in the process; Jake Pardo who informed me that model theory was a subject; and Alex Kruckman, through whom I learned about the Berkeley Model Theory Seminar.

To God alone be the glory

Chapter 1

Notation and Conventions

We will write $A := B$ to define A as B. All rings will be commutative. A ring is a *domain* if it is an integral domain. If R is a domain, $\mathrm{Frac}(R)$ denotes the field of fractions. All diagrams will commute, unless explicitly stated otherwise. A map $X \to Y$ will be written as $X \twoheadrightarrow Y$ if it is surjective, $X \hookrightarrow Y$ if it is injective, and $X \xrightarrow{\sim} Y$ if it is an isomorphism. If K is a field, then K^{alg} and K^{sep} denote the algebraic closure and separable closure of K. If R is a ring, $R^\times$ denotes the group of units in R. If X is a subset of a topological space, $\overline{X}$ denotes the closure of X, X^{int} denotes the interior of X, and ∂X denotes $\overline{X} \setminus X^{int}$. If X and Y are sets, $X \triangle Y$ denotes the symmetric difference of X and Y, and X^c denotes the complement of X, in a universe that will be clear from context. We will write $X \subseteq Y$ to indicate that X is a subset of Y, and $X \subsetneq Y$ to indicate that X is a proper subset of Y (when this needs to be emphasized). A *chain* is a totally ordered subset of a poset. A collection of sets is a chain if it is totally ordered by inclusion.

The abbreviation "resp." stands for *beziehungsweise*.

1.1 Valuation theory

Valuations are allowed to be trivial. If (K, v) is a valued field, we will typically use the following notation for the various components of the valuation data:

- The value group will be denoted Γ, or $\Gamma(K)$ or vK to make the dependence on K and v clear. (Note we require the valuation map to be surjective onto Γ.)
- The valuation $K^\times \to \Gamma$ or $K \to \Gamma \cup \{\infty\}$ will be denoted $v(-)$, except in Chapter 6 where it will be denoted $\mathrm{val}(-)$.
- The residue field will be denoted k or Kv.
- The valuation ring will typically be denoted $\mathcal{O}$ or $\mathcal{O}_K$
- The maximal ideal of the valuation ring will be denoted $\mathfrak{m}$ or $\mathfrak{m}_K$

- The residue field $\mathcal{O}/\mathfrak{m}$ will be denoted k or Kv
- The residue map $\mathcal{O} \to k$ will be denoted $\mathrm{res}(-)$.
- The RV sort $K^\times/(1+\mathfrak{m})$ will be denoted rvK, and the map $K \to rvK$ will be denoted $rv(-)$.

Valuations will be written additively, so that

$$v(xy) = v(x) + v(y)$$

$$v(x+y) \geq \min(v(x), v(y))$$

If K and k are fields, a *place* $K \to k$ will mean a map

$$\nu : K \cup \{\infty\} \twoheadrightarrow k \cup \{\infty\}$$

of the form

$$\nu(x) = \begin{cases} \mathrm{res}(x) & v(x) \geq 0 \\ \infty & v(x) < 0 \\ \infty & x = \infty \end{cases}$$

for some valuation on K and some identification of k with the residue field of K.

If v_1 and v_2 are two valuations on a field K, we will say that v_1 is *coarser than* v_2, or v_1 is *a coarsening of* v_2, or v_2 is *finer than* v_1, if the following equivalent conditions hold.

- There is some convex subgroup $\Delta \subseteq v_2K$ such that v_1 is isomorphic to the composition

$$K^\times \xrightarrow{v_2} v_2K \to v_2K/\Delta$$

- The place $K \to k_2$ is a composition of $K \to k_1$ and another place $k_1 \to k_2$
- The valuation ring $\mathcal{O}_1$ contains the valuation ring $\mathcal{O}_2$
- The valuation ideal $\mathfrak{m}_1$ is contained in the valuation ideal $\mathfrak{m}_2$

If K is a valued field, K^{hens} will denote the henselization of K. An extension L/K of valued fields is *immediate* if $vL = vK$ and $Lv = Kv$. A valued field is *maximally complete* if it admits no proper immediate extensions (algebraic or transcendental). A *ball* in a valued field is a non-empty open or closed ball, possibly of radius $\pm\infty$. A valued field is *spherically complete* if every chain of balls has non-empty intersection. Spherical completeness is (non-trivially) equivalent to maximal completeness. Spherically complete fields are henselian. A valued field is *algebraically maximal* if it admits no algebraic immediate extensions. Algebraically maximal fields are henselian.

The *residue characteristic* of a valued field K means the characteristic of the residue field. We say that K has *mixed characteristic* $(0, p)$ if K has characteristic 0 and k has

characteristic p. We say that K has pure characteristic 0 or pure characteristic p if both K and k have characteristic 0 or p, respectively. If Γ is an ordered abelian group and p is a prime, $\mathrm{Int}_p \Gamma$ will denote the maximal convex p-divisible subgroup of Γ. If K is a field of mixed characteristic, we will say that K is *finitely ramified* if Γ has a minimum positive element τ and $v(p)$ is a finite multiple of τ. Equivalently, the interval $[-v(p), v(p)]$ in the value group is finite.

1.2 Model theory

If $\phi(x)$ or $\phi(x; y)$ is a formula, then x or y may be tuples of variables. Sometimes we will still use the tuple notation $\vec{x}$ or $\vec{y}$ for emphasis, especially when contrasting singletons with tuples (or tuples with tuples of tuples). We will say that $\phi(x)$ is a C-formula to indicate that there are (suppressed) constants from C. We will say that a is "in" or "from" a model M if a is a tuple of elements in M, or $a \in \mathrm{dcl}^{eq}(M)$. We will use notation like $\phi(M; b)$ to denote the set of tuples a from M such that $\phi(a; b)$ holds. We will write $\top$ and $\bot$ for true and false.

Types will be complete types unless stated otherwise. We will write $M \models T$ to indicate that M is a model of the theory T, and $a \models p$ to indicate that the tuple or imaginary a realizes the type p. We will write $p|C$ to indicate the restriction of the type p to the set C. We will write

$$A \equiv_B C$$

to indicate that the tuples A and C have the same (*-)type over B, unless A and C are structures containing B, in which case $A \equiv_B C$ will mean that A and C are elementarily equivalent over B. (This will be clear from context.) If $B = \emptyset$ in either of these settings, we will write $\equiv$ for $\equiv_B$.

We will write $M \leq N$ to indicate that M is a substructure of N, and $M \preceq N$ to indicate that M is an elementary substructure of N.

We will usually write union and concatenation of sets, elements, and tuples multiplicatively, writing Sa and aS and ST in place of $S \cup \{a\}$ or $a \cup \{S\}$ or $S \cup T$. If we want to write the concatenation ab of two tuples a and b explicitly, we write $a^\frown b$. When we need to write field-theoretic multiplication explicitly, we will usually write $a \cdot b$.

"Definable" means definable with parameters, and 0-definable means definable without parameters. Definable sets will mean interpretable sets, unless stated explicitly otherwise. If X is a definable set, $\ulcorner X \urcorner$ will denote a code for X. A *unary definable set* is a definable subset of the home sort (in a 1-sorted structure). A set is *type-definable* if it is a small intersection of definable sets, $\vee$-*definable* if it is a small union of definable sets, *pro-definable* if it is a small inverse limit of definable sets, and *ind-definable* if it is a small direct limit of definable sets. A *definable family* is a collection of the form

$$\{\{x \in X : (x, y) \in R\} : y \in Y\}$$

for some definable sets X, Y, R with $R \subseteq X \times Y$. An *ind-definable family* is defined similarly, with X, Y, R being ind-definable. A collection of sets is *uniformly definable* if it is a subcollection of a definable family.

When working in a monster, the monster model will be denoted $\mathbb{M}$. A "model" will mean a small elementary substructure of $\mathbb{M}$. A *global type* will mean a type over the monster. A C-invariant type will mean a global type which is $\mathrm{Aut}(\mathbb{M}/C)$-invariant. A C-definable type will mean a global definable type that is C-invariant. If p is a global definable type, and $\phi(x;y)$ is a formula, then $(d_p x)\phi(x;y)$ is a formula $\psi(y)$ such that

$$(\mathbb{M} \models \psi(b)) \iff \phi(x;b) \in p(x) \text{ for all } b \text{ in } \mathbb{M}$$

If p is an invariant type and f is a definable function, the pushforward f_*p is the global type such that

$$a \models p|C \implies f(a) \models f_*p|C$$

for sufficiently large sets C. If p and q are C-invariant types, then $p \otimes q$ is the C-invariant type whose realizations over a set $B \supseteq C$ are exactly the pairs (a, b) such that

$$a \models p|Bb \text{ and } b \models q|B$$

Abusing terminology slightly, we will let $p^{\otimes n}$ be the type whose realizations are sequences $a_1, \ldots, a_n$ such that $a_i \models p|Ba_{<i}$, where $a_{<i}$ denotes the set of elements $a_1, \ldots, a_{i-1}$.[1] We will also write $p^{\otimes \alpha}$ for ordinals α, defined analogously. We will use the term "Morley sequence" to refer to sequences of this sort, rather than in the sense used in simple and rosy theories (independent indiscernible sequences).

If a is a tuple and B is a set,

- $\mathrm{tp}(a/B)$ denotes the type of a over B
- $\mathrm{stp}(a/B)$ denotes the strong type of a over B, i.e., the type of a over $\mathrm{acl}^{eq}(B)$
- $SU(a/B)$ denotes the SU-rank of $\mathrm{tp}(a/B)$ in a simple theory
- $RM(a/B)$ denotes the Morley rank of $\mathrm{tp}(a/B)$
- $\text{dp-rk}(a/B)$ denotes the dp-rank of a over B
- $\dim(a/B)$ denotes the rank of a over B in an o-minimal setting
- $\mathrm{bdn}(a/B)$ denotes the burden of $\mathrm{tp}(a/B)$

Also, if X is a (type-)definable set, then $\dim(X)$, $RM(X)$, $\text{dp-rk}(X)$, $SU(X)$, and $\mathrm{bdn}(X)$ will denote the o-minimal rank, the Morley rank, the dp-rank, the SU-rank, or the burden of X.

[1] So $(p \otimes p)(x_2, x_1) = p^{\otimes 2}(x_1, x_2)$.

C-minimal and o-minimal will mean dense C-minimal and dense o-minimal. Finite structures will be considered dp-minimal. We will refer to ict-patterns as *randomness patterns*. If A, B, C are small sets, then $A \mathop{\smile\!\!\!\!\!\!|}_C B$ will indicate that $\mathrm{tp}(A/BC)$ does not fork over C, and $A \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{b}}_C B$ will indicate that $\mathrm{tp}(A/BC)$ does not thorn-fork over C.

Chapter 2

A direct proof of quantifier elimination and C-minimality in ACVF

ACVF is the theory of non-trivially valued *a*lgebraically *c*losed *v*alued *f*ields.

Theorem 2.0.1.

1. *ACVF eliminates quantifiers in the 1-sorted language consisting of the field language expanded by a binary predicate for the relation $v(x) \leq v(y)$.*
2. *ACVF is the model completion of the theory of valued fields. In particular, every valued field embeds into a model of ACVF.*
3. *ACVF is C-minimal. In other words, if $K \models ACVF$, every unary definable set $D \subseteq K$ is a boolean combination of open and closed balls (including singletons).*

These results are well-known, but we include a short proof here because they are essential for most of what follows. The history of quantifier elimination is confusing, because it was almost (but not quite) proven by Abraham Robinson in [63]. For a discussion of the history, see the proof of Theorem 2.1.1 in [26]. The C-minimality of ACVF was proven by Holly ([30], Theorem 3.26) before C-minimality was first isolated in [51]. The proofs given here are influenced by some course notes of Hrushovski that seem to have disappeared.

2.1 An abstract criterion for quantifier elimination

Definition 2.1.1. *A theory T has* prime models over substructures *if for every $M \models T_\forall$ there is some $M' \models T$ and an embedding $M \hookrightarrow M'$ such that for every $M'' \models T$ and*

embedding $M \hookrightarrow M''$, there is an embedding $M' \hookrightarrow M''$ making the diagram commute

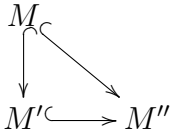

For example, the theory of fields has prime models over substructures (in the language of rings): the prime model over a domain R is the field of fractions $\mathrm{Frac}(R)$. We are abusing terminology massively, as these are not "prime models" in the standard sense.

Recall the well-known amalgamation property:

Definition 2.1.2. *A theory T has the* amalgamation property *if for any diagram of embeddings of models of T*

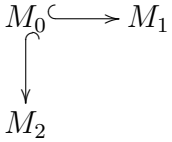

there is a model M_3 and embeddings $M_i \hookrightarrow M_3$ for $i = 1, 2$ such that the diagram commutes:

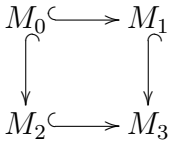

Definition 2.1.3. *Let $M \leq N$ be structures. Say that M is* 1-ec *in N if for every quantifier-free formula $\phi(x; \vec{y})$ with $|x| = 1$, and every $\vec{m}$ in M,*

$$M \models \exists x \phi(x; \vec{m}) \iff N \models \exists x \phi(x; \vec{m})$$

Equivalently, every non-empty quantifier-free M-definable subset of N intersects M.

If T is a theory, say that a model $M \models T$ is 1-ec *(as a model of T) if M is 1-ec in every model $N \models T$ extending M.*

Then, we have the following criterion for quantifier elimination:

Lemma 2.1.4. *Let T_0 be a theory with prime models over substructures and the amalgamation property. Suppose the 1-ec models of T_0 are an elementary class, axiomatized by a theory T. Then T has quantifier elimination. If T_0 is inductive, then T is the model completion of T.*

Proof. For quantifier elimination, it suffices to show that for every quantifier free formula $\phi(x; \vec{y})$, there is a quantifier-free formula $\psi(\vec{y})$ such that

$$T \vdash \psi(\vec{y}) \leftrightarrow \exists x \phi(x; \vec{y}).$$

By compactness, this amounts to the claim that the quantifier free type of a tuple $\vec{b}$ determines whether $\exists x \phi(x; \vec{b})$ holds.

Let M and N be two models of T and $\vec{b}_M$ and $\vec{b}_N$ be tuples having the same type. Let S_M and S_N be the structures generated by $\vec{b}_M$ and $\vec{b}_N$. Then $S_M \cong S_N$. Identify them as a common structure S with embeddings into both M and N. Because T_0 has prime models over substructures and the amalgamation property, the configuration

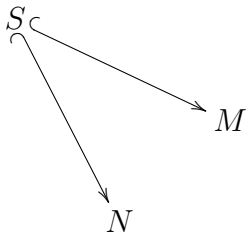

can be completed to a commutative diagram

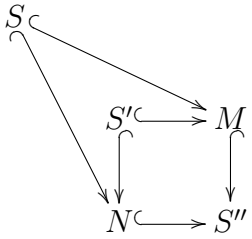

where $S', S'' \models T_0$. Then

$$M \models \exists x \phi(x; \vec{b}) \implies S'' \models \exists x \phi(x; \vec{b}) \implies N \models \exists x \phi(x; \vec{b})$$

where the second $\implies$ holds because N is 1-ec in S''. By symmetry, the converses hold. Thus

$$M \models \exists x \phi(x; \vec{b}_M) \iff N \models \exists x \phi(x; \vec{b}_N)$$

completing the proof of quantifier elimination.

Now suppose T_0 is inductive. Then every model of T_0 embeds into an existentially closed model. Existentially closed models of T_0 are 1-ec, hence models of T. Thus every model of T_0 embeds into a model of T. Conversely, every model of T already is a model of T_0. The theory T is model complete and substructure complete by quantifier elimination. □

As an example, let us use Lemma 2.1.4 to prove quantifier elimination in ACF, the theory of algebraically closed fields. Let T_0 be the theory of fields and T be ACF.

We noted above that the theory T_0 has prime models over substructures, because of fraction fields. The amalgamation property can be seen as follows. Given embeddings $K_0 \hookrightarrow K_1$ and $K_0 \hookrightarrow K_2$, form the tensor product $K_1 \otimes_{K_0} K_2$. A tensor product of two

nontrivial K_0-vector spaces is nontrivial, so $K_1 \otimes_{K_0} K_2$ is not the zero ring. It therefore admits a homomorphism to a field K_3, which in turn yields a diagram

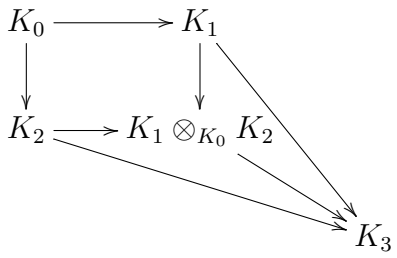

completing the amalgamation.

It remains to check the following

Claim 2.1.5. *A field K is 1-ec if and only if it is algebraically closed.*

Proof. First suppose K is algebraically closed and L is an extension field. The quantifier-free K-definable subsets of L are boolean combinations of sets

$$\{x \in L : P(x) = 0\} \text{ for some nonzero } P(X) \in K[X]$$

Factoring $P(x)$, this set is a finite subset of K. Consequently, every quantifier-free K-definable subset of L is either a finite subset of K or the complement of such a set. As K is infinite, any such set which is non-empty must intersect K.

Conversely, suppose K is not algebraically closed. Let $P(X) \in K[X]$ be an irreducible polynomial which lacks a root. Then $K[X]/(Q(X))$ is a field extending K, in which P has a root (namely X). So K is not 1-ec in the extension field $K[X]/Q(X)$. □

Consequently, all the conditions of Lemma 2.1.4 hold, so ACF has quantifier elimination.

In the process of proving quantifier elimination, we also proved strong minimality. Something similar will happen with ACVF. In fact, we have pedantically gone through the proof of quantifier elimination in ACF because the proof for ACVF will be analogous in many ways.

We would like to apply Lemma 2.1.4 where T is ACVF and T_0 is the theory of valued fields. Recall that we are working in the language of rings expanded by a binary predicate for the relation $v(x) \leq v(y)$. If R is a substructure of a valued field K, the binary predicate uniquely determines the valuation structure on $\mathrm{Frac}(R)$. Consequently, the theory of valued fields has prime models over substructures. The theory of valued fields is also inductive.

The other two conditions will require some work

1. The theory of valued fields has the amalgamation property. We will prove this in §2.2
2. The 1-ec valued fields are precisely the models of ACVF. We will prove this in §2.3. Along the way, we will prove C-minimality.

2.2 Part 1: Amalgamation

The following lemma is Theorem 3.1.1 in [19].

Lemma 2.2.1 (Chevalley)**.**

1. *Let K be a field, R be a subring, and $\mathfrak{p}$ be a prime ideal. Then there is a valuation ring $(\mathcal{O}, \mathfrak{m})$ in K, such that*

$$\begin{aligned} R &\subseteq \mathcal{O} \\ R \setminus \mathfrak{p} &= \mathcal{O}^\times \cap R \\ \mathfrak{p} &= \mathfrak{m} \cap R \end{aligned}$$

2. *Let R be a ring, and $\mathfrak{q} \subseteq \mathfrak{p}$ be two primes. Then there is a valuation ring $(\mathcal{O}, \mathfrak{m})$ and a ring homomorphism $f : R \to \mathcal{O}$ with*

$$\begin{aligned} \ker f = f^{-1}(0) &= \mathfrak{q} \\ f^{-1}(\mathfrak{m}) &= \mathfrak{p} \end{aligned}$$

Remark 2.2.2. *If $(\mathcal{O}_i, \mathfrak{m}_i)$ is a valuation ring for $i = 1, 2$, and $f : \mathcal{O}_1 \to \mathcal{O}_2$ is a ring homomorphism, then f arises from an embedding of valued fields if and only if $f^{-1}(0) = 0$ and $f^{-1}(\mathfrak{m}_2) = \mathfrak{m}_1$.*

Corollary 2.2.3. *Let L/K be an extension of valued fields. Any valuation on K extends to one on L.*

Proof. Given a valuation ring $(\mathcal{O}_K, \mathfrak{m}_K)$ on K, apply Chevalley's Lemma to the pair $(\mathcal{O}_K, \mathfrak{m}_k)$ to get a valuation ring $(\mathcal{O}_L, \mathfrak{m}_L)$ on L. □

Lemma 2.2.4. *Let R be a valuation ring.*

1. *Any finitely generated torsion-free R-module M is free.*

2. *Any torsion-free R-module M is flat.*

Proof. 1. Let $g_1, \ldots, g_n$ be a minimal set of generators. Then $g_1, \ldots, g_n$ freely generate M. If not, then there are some $r_i \in R$, not all zero, such that

$$\sum_{i=1}^{n} r_i g_i = 0.$$

Let j be such that $v(r_j)$ is minimal. Then $r_i/r_j \in R$ for each i, and

$$r_j \cdot \sum_{i=1}^{n} \frac{r_i}{r_j} g_i = \sum_{i=1}^{n} r_i g_i = 0.$$

As M is torsion-free,

$$\sum_{i=1}^{n} \frac{r_i}{r_j} g_i = 0, \text{ so that } g_j = -\sum_{i \neq j} \frac{r_i}{r_j} g_i,$$

contradicting minimality.

2. M can be written as a direct limit of its finitely-generated submodules, which are torsion-free, hence flat. A direct limit of flat modules is flat.

□

Lemma 2.2.5. *Let $f : A \to R$ be an injective ring homomorphism. Then every minimal prime $\mathfrak{p}$ of A is a pullback $f^{-1}(\mathfrak{q})$ of some prime $\mathfrak{q}$ in R.*

Proof. The set $S := f(A \setminus \mathfrak{p})$ is a submonoid of R, and $0 \notin S$ by injectivity. The localization $S^{-1}R$ is nonzero, and any prime in $S^{-1}R$ pulls back to a prime $\mathfrak{q} \lhd R$ not intersecting S. The prime ideal $f^{-1}(\mathfrak{q})$ is contained in, hence equal to $\mathfrak{p}$. □

Lemma 2.2.6. *Given a diagram of valuation rings*

$$\begin{array}{ccc} \mathcal{O}_0 & \hookrightarrow & \mathcal{O}_1 \\ \downarrow & & \\ \mathcal{O}_2 & & \end{array}$$

if $\mathfrak{q}$ is a minimal prime in $\mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2$, then $\mathfrak{q}$ pulls back to the zero ideal in $\mathcal{O}_i$ for each i.

Proof. Let K_i be $\mathrm{Frac}(\mathcal{O}_i)$. By Lemma 2.2.4.2, $\mathcal{O}_1$ and K_2 are flat $\mathcal{O}_0$-modules, so the natural map

$$\mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2 \hookrightarrow \mathcal{O}_1 \otimes_{\mathcal{O}_0} K_2 \hookrightarrow K_1 \otimes_{\mathcal{O}_0} K_2$$

is an injection. By Lemma 2.2.5, there is some prime $\mathfrak{q}_0$ in $K_1 \otimes_{\mathcal{O}_0} K_2$ which pulls back to $\mathfrak{q}$. Now for $i = 1, 2$, we have a commuting square of sets

$$\begin{array}{ccc} \operatorname{Spec} \mathcal{O}_i & \xleftarrow{f^*} & \operatorname{Spec} K_i \\ {\scriptstyle g^*}\uparrow & & \uparrow{\scriptstyle h^*} \\ \operatorname{Spec} \mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2 & \xleftarrow[i^*]{} & \operatorname{Spec} K_1 \otimes_{\mathcal{O}_0} K_2 \end{array}$$

Then

$$g^*\mathfrak{q} = g^* i^* \mathfrak{q}_0 = f^* h^* \mathfrak{q}_0 = f^*(0) = (0)$$

because $h^*\mathfrak{q}_0 \in \operatorname{Spec} K_i$ can only be the zero ideal. Pulling $g^*\mathfrak{q}$ from $\mathcal{O}_i$ to $\mathcal{O}_0$ also handles the case $i = 0$. □

Lemma 2.2.7. *The theory of valued fields has the amalgamation property.*

Proof. Suppose we are given a diagram

$$\begin{array}{ccc} K_0 & \hookrightarrow & K_1 \\ \downarrow & & \\ K_2 & & \end{array}$$

of valued fields. Let $\mathcal{O}_i$, $\mathfrak{m}_i$, and k_i be the associated data. The ring $k_1 \otimes_{k_0} k_2$ is nonzero, so it has a prime ideal $\mathfrak{p}_0$. Let $\mathfrak{p} \in \operatorname{Spec} \mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2$ be the pullback of $\mathfrak{p}_0$ along the map

$$\mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2 \to k_1 \otimes_{k_0} k_2.$$

For $i = 1, 2$, we have a commuting square of sets

$$\begin{array}{ccc} \operatorname{Spec} \mathcal{O}_i & \xleftarrow{f^*} & \operatorname{Spec} k_i \\ {\scriptstyle g^*}\uparrow & & \uparrow{\scriptstyle h^*} \\ \operatorname{Spec} \mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2 & \xleftarrow[i^*]{} & \operatorname{Spec} k_1 \otimes_{k_0} k_2 \end{array}$$

Thus

$$g^*\mathfrak{p} = g^* i^* \mathfrak{p}_0 = f^* h^* \mathfrak{p}_0 = f^*(0) = \mathfrak{m}_i$$

as $h^*\mathfrak{p}_0$ can only be the zero ideal in k_i.

Let $\mathfrak{q} \in \operatorname{Spec} \mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2$ be any minimal prime below $\mathfrak{p}$. The pullback $g^*\mathfrak{q}$ is 0, by Lemma 2.2.6. Apply Lemma 2.2.1.2 to get a homomorphism

$$m : \mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2 \to \mathcal{O}_3$$

such that $m^{-1}(0) = \mathfrak{q}$ and $m^{-1}(\mathfrak{m}_3) = \mathfrak{p}$. This in turn induces a square of valuation rings

$$\begin{array}{ccc} \mathcal{O}_0 & \longrightarrow & \mathcal{O}_1 \\ \downarrow & & \downarrow \\ \mathcal{O}_2 & \longrightarrow & \mathcal{O}_3 \end{array} \tag{2.1}$$

Pulling back 0 and $\mathfrak{m}_3$ along the maps $\mathcal{O}_i \to \mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2$, we get

$$\begin{array}{lll} \mathcal{O}_3 & \leftarrow \mathcal{O}_1 \otimes_{\mathcal{O}_0} \mathcal{O}_2 & \leftarrow \mathcal{O}_i \\ \mathfrak{m}_3 & \mapsto \mathfrak{p} & \mapsto \mathfrak{m}_i \\ (0) & \mapsto \mathfrak{q} & \mapsto (0) \end{array}$$

So by Remark 2.2.2, the maps $\mathcal{O}_i \to \mathcal{O}_3$ induce maps of valuation rings on the fraction fields. Consequently, (2.1) solves the amalgamation problem. □

2.3 Part 2: C-minimality

It remains to characterize the 1-ec valued fields. In the process, we will prove C-minimality. By a *ball* in a valued field K, we will mean a set of one of the following forms

$$\{x : v(x-c) \geq \gamma\} \text{ for some } \gamma \in vK$$
$$\{x : v(x-c) > \gamma\} \text{ for some } \gamma \in vK$$
$$\{c\}$$
$$K$$

We call c a center and γ the radius. We think of $\{c\}$ as a ball of radius $+\infty$, and K as a ball of radius $-\infty$.

Lemma 2.3.1. *Let L/K be an extension of valued fields, with $K = K^{alg}$. Every quantifier-free K-definable subset of L is a boolean combination of balls with center and radii in K and vK.*

Proof. We will use the term "K-ball" to mean a ball with center and radius in K. It suffices to consider sets cut out by $P(x) = 0$ and $v(P(x)) \leq v(Q(x))$. The set

$$\{x \in L : P(x) = 0\}$$

is easily dealt with: it is a finite union of singletons in K.

Now consider the condition $v(P(x)) \leq v(Q(x))$. Factoring P and Q, we can write the condition as

$$\sum_i n_i \cdot v(x - \alpha_i) \geq \gamma \tag{2.2}$$

for some distinct $\alpha_i \in K$ and $\gamma \in vK$.

Each condition

$$v(x - \alpha_i) < v(x - \alpha_j)$$

is an open K-ball with center α_j and radius $v(\alpha_j - \alpha_i)$.

So for each i the set

$$S_i = \{x \in L : i = \arg\max_j v(x - \alpha_j)\}$$

is a combination of K-balls. On S_i, $v(x-\alpha_j) = v(\alpha_i - \alpha_j)$ for $j \neq i$, so (2.2) is equivalent to

$$n_i \cdot v(x - \alpha_i) \geq \gamma'$$

for some $\gamma' \in vK$. As vK is divisible (because $K^\times$ is), this is equivalent to a closed K-ball or complement of an open K-ball. □

Note that Lemma 2.3.1 will give C-minimality once quantifier elimination is established. In the case of ACF, once we knew "quantifier-free strong minimality," 1-existential closedness was an easy consequence. For ACVF, we must do some annoying casework.[1]

Lemma 2.3.2. *Let K be a valued field. Any boolean combination D of balls in K can be written as a disjunction of swiss cheeses, i.e., sets of the form $C = B_0 \setminus \bigcup_{i=1}^{n} B_i$ where $B_1, \ldots, B_n$ are subballs of B_0.*

Proof. First we recall that in valued fields,

$$\text{Any two balls which intersect are comparable (one contains the other).} \tag{2.3}$$

The set D can be written as a union of intersections of balls and complements of balls. So we may assume that D is an intersection of balls and complements of balls. Write

$$D = (B_1 \cap \cdots \cap B_m) \setminus (C_1 \cup \cdots \cup C_\ell)$$

where each B_i or C_i is a ball. We may take $m > 0$ by throwing in L as one of the B's. By (2.3), $\bigcap_i B_i$ is either empty or one of the B_i's. If it is empty, then D is an empty union. Otherwise, we may write

$$D = B \setminus (C_1 \cup \cdots \cup C_\ell)$$

We may discard any C_i which is disjoint from B, so by (2.3), we may assume each $C_i \subseteq B$. □

We omit the proofs of the next three lemmas, which are straightforward but tedious.

Lemma 2.3.3. *Let L/K be an extension of valued fields. Suppose B_i is a ball in L with center and radius in K. If the value group of K is a dense linear order without endpoints (i.e., a model of DLO), then*

$$B_1 \subsetneq B_2 \implies B_1 \cap K \subsetneq B_2 \cap K$$

Lemma 2.3.4. *Let K be a valued field with $vK \models DLO$. Then every subball of $\mathfrak{m}$ is contained in a closed ball of radius $\gamma > 0$ centered at 0. Every subball of K is contained in a closed ball of radius $\gamma > -\infty$ centered at 0.*

Lemma 2.3.5. *Let K be any valued field. Every proper subball of $\mathcal{O}$ is contained in a residue class.*

[1] None of what follows is particularly ACVF-specific, with the exception of Lemma 2.3.6 below. In fact, the theory of C-structures has a model completion: a C-structure is existentially closed if an only if the following two conditions hold

- Every open ball (including the entire structure) has no maximal proper subball
- Every closed ball has infinitely many maximal proper subballs

Lemma 2.3.6 ensures that if $K \models ACVF$, then the underlying C-structure is existentially closed. Combined with Lemma 2.3.1, this immediately yields that models of ACVF are 1-ec valued fields.

Lemma 2.3.6. *Let K be a model of ACVF. Then the value group vK is a dense linear order, and the residue field Kv is infinite.*

Proof. The group $K^\times$ is divisible, as K is algebraically closed. By the short exact sequence

$$1 \to \mathcal{O}^\times \to K^\times \to vK \to 1$$

the group vK is divisible. As vK is torsion-free (being ordered), $\mathcal{O}^\times$ is divisible. The surjection $\mathcal{O}^\times \to k^\times$ then implies that $k^\times$ is divisible.

The group vK is divisible, and non-trivial (by definition of ACVF). Therefore it is a dense linear order. If k is finite, then $k^\times$ is a finite divisible group, hence trivial. Thus $k = \mathbb{F}_2$.

As $K = K^{alg}$, there is some $x \in$K such that $x(x-1) = 1$. If $v(x) < 0$, then $v(x-1) < 0$ by the ultrametric inequality, and so

$$0 = v(1) = v(x(x-1)) = v(x) + v(x-1) < 0$$

which is absurd. Thus $v(x) \geq 0$ and $x \in \mathcal{O}$. If α is the residue of x, then $\alpha(\alpha - 1) = 1$, contradicting the fact that $k = \mathbb{F}_2$. □

Lemma 2.3.7. *Let K be a model of ACVF. Then no ball can be written as a finite union of proper subballs.*

Proof. Suppose B can be written as a finite union of proper subballs. After applying an affine linear transformation, we may assume that B has center 0, and radius 0 or $\pm\infty$. Thus B is one of $0, \mathfrak{m}, \mathcal{O}$, or K.

The case $B = \{0\}$ is trivial. If $B = \mathcal{O}$, every proper subball is contained in a residue class by Lemma 2.3.5. By Lemma 2.3.6, a finite union of residue classes cannot exhaust $\mathcal{O}$.

Finally, suppose $B = \mathfrak{m}$ or $B = K$, and $B = B_1 \cup \cdots B_n$. By Lemma 2.3.4, we may assume each B_i is a closed ball of radius strictly less than the radius of B, and center 0. As all the balls have center 0, they are pairwise comparable, so $\bigcup_i B_i = B_j$ for some j. Then B_j is strictly smaller than B_i, a contradiction. □

Putting everything together, we see

Lemma 2.3.8. *Let L be a valued field extending $K \models ACVF$. Let $B_0, \cdots, B_n$ be balls in L with radii and centers in K. Suppose $B_i \subseteq B_0$ for each i. Let*

$$C = B_0 \setminus \bigcup_{i=1}^{n} B_i$$

If C is non-empty, then C intersects K.

Proof. For each i, let $B_i' = B_i \cap K$ denote the corresponding ball in K. As C is non-empty, $B_i \subsetneq B_0$ for each i. By Lemma 2.3.3, $B_i' \subsetneq B_0'$ for each i. By Lemma 2.3.7,

$$\emptyset \neq B_0' \setminus \bigcup_i B_i' = C' := C \cap K$$

□

To finish the proof of Theorem 2.0.1, it remains to show:

Lemma 2.3.9. *A valued field K is 1-ec if and only if K is algebraically closed and non-trivially valued.*

Proof. First suppose that $K \neq K^{alg}$. By Corollary 2.2.3, we can extend the valuation to K^{alg}. Because K is not 1-ec as a *field* in K^{alg}, it is not 1-ec as a valued field.

Next suppose that K is trivially valued. The field $K((t))$ of formal Laurent series is a non-trivial valued field extending K. The formula

$$\neg(v(x) \leq v(1))$$

has a solution in $K((t))$ but not K, so K is not 1-ec.

Finally, suppose K is algebraically closed and non-trivially valued. Let L/K be a field extension, and $D \subseteq L^1$ be non-empty and quantifier-free K-definable. We must show that D intersects K.

By Lemma 2.3.1, D is a finite boolean combination of balls with centers and radii in K. By Lemma 2.3.2, D is a disjunction of K-swiss cheeses, so we may assume $D = B_0 \setminus \bigcap_{i=1}^n B_i$ for some K-balls $B_0, \ldots, B_n$. Then D intersects K by Lemma 2.3.8. □

Also, the theory of valued fields is inductive, so ACVF really is the model completion of the theory of valued fields.

Chapter 3

Quantifier elimination and dp-minimality for certain valued fields

Several theories of fields and valued fields are known to be dp-minimal, including the following:

- Algebraically closed valued fields, because they are C-minimal.
- Real closed fields, because they are o-minimal.
- Characteristic 0 nonarchimedean local fields. This is proven in Corollary 7.8 of [3], among other places.[1]

Several other dp-minimal fields can be generated from these examples—for instance in [12], the following fact is proven:

> A henselian valued field (K, v) with residue characteristic 0 is dp-minimal if and only if vK and Kv are dp-minimal.

This shows that, for example, $\mathbb{Q}_p((t))$ is a dp-minimal field.

In this chapter, we will exhibit another source of dp-minimal valued fields. Let p be a prime and Γ be an ordered abelian p-divisible group such that $\Gamma/n\Gamma$ is finite for all n. We will show that the Hahn series field $\mathbb{F}_p^{alg}((t^\Gamma))$ is dp-minimal. There is also a mixed characteristic analogue, in which the p-divisibility condition on Γ is weakened slightly. Both of these results are in Theorem 3.3.7 below.

Along the way, we will manually prove quantifier elimination results for fields like $\mathbb{F}_p^{alg}((t^\Gamma))$ in §3.2. The field $\mathbb{F}_p^{alg}((t^\Gamma))$ is *tame*, hence amenable to the analysis in [47] and [48]. We will not use this, however, for the following reasons:

- The mixed characteristic fields we consider will *not* necessarily be tame

[1] The proof in [3] is stated for the p-adics, but as noted at the beginning of §7.2 of that paper, the proof generalizes to finite extensions of the p-adics.

- The analysis of tame fields uses more advanced techniques than are necessary in our case
- We would like quantifier elimination relative to the value group, rather than the RV sort.

Here is an example of a non-tame field for which we will prove dp-minimality and quantifier elimination relative to the value group. Let $\mathbb{M}$ be a mixed characteristic $(0, p)$ monster model of ACVF. Choose $t \in \mathbb{M}$ such that

$$v(t) > v(q) \qquad \forall q \in \mathbb{Q}$$

Let K be a spherical completion of the field generated by t, p^{1/p^n} for all n, and the mth roots of unity for all m prime to p. Then Theorem 3.2.16 and Theorem 3.3.7 will show that (K, v) has quantifier elimination relative to the value group and is dp-minimal.

These fields are hardly natural. Their significance lies in the fact that, in a certain sense, they are the *last remaining source of dp-minimal fields*. In Chapter 9, we will see that, together with the previously known dp-minimal fields, they generate *all* dp-minimal fields and valued fields.

3.1 Some valuation theory

The *characteristic exponent* of a valued field K is p if K has residue characteristic p, and 1 if K has residue characteristic 0.

Fact 3.1.1. *If K is henselian and L/K is a finite extension, then*

$$[L : K] = |vL/vK| \cdot [Lv : Kv] \cdot p^d$$

where p is the characteristic exponent, and $d \in \mathbb{N}$.

We will call this equation *the defect equation*. One says that the extension L/K is *defectless* if $p^d = 1$. A henselian field K is *defectless* if every finite extension is defectless.

Fact 3.1.2. *The henselian field K is defectless if any of the following three conditions hold:*

- *K has residue characteristic 0*
- *K has residue characteristic p, and p does not divide the degree of any finite extension of K*
- *K is spherically complete*

Definition 3.1.3. *If Γ is an ordered abelian group and p is prime, let $\mathrm{Int}_p \Gamma$ denote the maximal convex p-divisible subgroup of Γ.*

Definition 3.1.4. *A valuation $v : K \to \Gamma$ is* roughly p-divisible *if $[-v(p), v(p)] \subseteq p \cdot \Gamma$, where $[-v(p), v(p)]$ denotes $\{0\}$ in pure characteristic 0, denotes Γ in pure characteristic p, and denotes the usual interval $[-v(p), v(p)]$ in mixed characteristic.*

In mixed characteristic, (K, v) is roughly p-divisible if and only if $v(p) \in \mathrm{Int}_p$. In pure characteristic p, (K, v) is roughly p-divisible if and only if the value group is p-divisible.

Remark 3.1.5. *Let P be one of the following properties of valuation data:*

- *Roughly p-divisible*
- *Henselian*
- *Henselian and defectless*
- *Every countable chain of balls has non-empty intersection*

If $K_1 \to K_2$ and $K_2 \to K_3$ are places, the composition $K_1 \to K_3$ has property P if and only if each of $K_1 \to K_2$ and $K_2 \to K_3$ has property P. (In each case, this is straightforward to check.)

3.2 A quantifier elimination result

Let T_0 be the theory of henselian defectless fields (K, v) with $Kv \models ACF_p$ and with p-divisible value group. Let T be the theory of henselian defectless roughly p-divisible fields (K, v) with $Kv \models ACF_p$.

Every model of T_0 is a model of T, and the converse holds in characteristic p.

Models of T_0 are tame (in the sense of [48]), though models of T need not be. Nevertheless, we will see that many of the good properties of T_0 extend to T.

Remark 3.2.1. *If $M \models T_0$, then any finite field extension of M has degree prime to p. Indeed, if L/M is finite, then henselianity and defectlessness imply*

$$[L : M] = |vL/vM| \cdot [Lv : Mv].$$

But Mv is algebraically closed, so $[Lv : Mv] = 1$. And vM is p-divisible, so $|vL/vM|$ is prime to p.

Remark 3.2.2. *Let $(L, v)/(K, v)$ be an extension of valued fields. Suppose (L, v) is henselian and K is relatively separably closed in L. Then Kv is relatively separably closed in Lv.*

Proof. Otherwise, take $\alpha \in (Lv \cap Kv^{sep}) \setminus Kv$. Let $\overline{f}(X)$ be the monic irreducible polynomial of α over Kv. Let $f(X)$ be a lift of $\overline{f}(X)$ to $K[X]$. By henselianity of L, there is a unique root a of $f(X)$ lying over α. Moreover, a is a simple root, so $a \in K^{sep}$. Therefore a and α are in K and Kv, respectively. □

Proposition 3.2.3. *Let (M, v) be a model of T.*

1. *M is perfect*
2. *If $a \in M$ and $n \in \mathbb{N}$, then a is an nth power if and only if $v(a)$ is divisible by n.*
3. *If K is relatively algebraically closed in M, then $K \models T$.*

Proof. First suppose that $(M, v) \models T_0$.

1. If M has characteristic p, then M is perfect by Remark 3.2.1
2. One easily reduces to showing that if $v(a) = 0$, then a is an ℓth power for all primes ℓ. For $\ell \neq p$, this follows by henselianity and the fact that $\mathrm{res}(a)$ is an nth power. For $\ell = p$, this follows by Remark 3.2.1.
3. We will show $K \models T_0$. Note that K is henselian and perfect because it is relatively algebraically closed in M, which is henselian and perfect. Then M/K is regular, so $\mathrm{Gal}(M)$ surjects onto $\mathrm{Gal}(K)$. Since p is prime to $\mathrm{Gal}(M)$ (by Remark 3.2.1), p is also prime to $\mathrm{Gal}(K)$. In other words, p does not divide the degree of any finite extension of K. It follows immediately that (K, v) is defectless, Kv is perfect, and vK is p-divisible. Also, Kv is separably closed in Lv by Remark 3.2.2.

Next suppose $(M, v) \models T$ but $(M, v) \not\models T_0$. Then M has characteristic 0. Let v' be the coarsening of v with respect to the maximal convex subgroup of vM containing $v(p)$. Let v'' be the induced valuation on Mv'. Then (Mv', v'') is a model of T, and (M, v') is a henselian field of residue characteristic 0.

1. M is perfect because it has characteristic 0.
2. As before, one reduces to showing that if $v(a) = 0$, then a is an ℓth power. Because (M, v') is a henselian field of residue characteristic 0, and $v'(a) = 0$, the element a is an nth power if and only if its residue $res'(a) \in Mv'$ is an nth power. But $v''(res'(a)) = v(a) = 0$, so by the case of T_0 considered above, $res'(a)$ is an nth power.
3. Applying Remark 3.2.2 to $(M, v')/(K, v')$, we see that Kv' is relatively algebraically closed in Mv'. As Mv' is a model of T_0, so is Kv'. That is, the place $Kv \to Kv'$ is henselian and defectless, with p-divisible value group and algebraically closed residue field Kv. The place $K \to Kv'$ is henselian of residue characteristic 0 (hence defectless and roughly p-divisible). The composition $K \to Kv' \to Kv$ is therefore henselian, defectless, and roughly p-divisible. And Kv is algebraically closed.

□

Lemma 3.2.4. *Let K be a valued field. Suppose L and F are two immediate algebraic extensions of K which are models of T. Then L and F are isomorphic over K.*

Proof. We may replace K with the perfection of its henselization. We then only need to show that L and F are conjugate over K (isomorphic as fields over K).

First suppose that vK is p-divisible. Then K is Kaplansky, and the desired result follows by the uniqueness of maximal algebraic immediate extensions over Kaplansky fields (L and F are algebraically maximally complete because they are defectless).

Otherwise, K, L, and F have characteristic 0. Let v' be the coarsening with respect to the convex subgroup generated by $v(p)$. As L and F are immediate, $v'L = v'K = v'F$. As (K, v') is a henselian field with residue characteristic 0, the extensions L and F of (K, v') are unramified. By the structure theory of valued fields (Theorem 5.2.7 and Theorem 5.2.9 in [19]), L and F are isomorphic (as fields) as long as Lv' and Fv' are isomorphic extensions of Kv'. Note that Lv' and Fv' are immediate extensions of Kv'. Also, Lv' and Fv' are models of T_0. By the T_0 case considered above, Lv' and Fv' are isomorphic, and we are done. □

Definition 3.2.5. *Let M and N be valued fields. A* partial v-elementary map from M to N *is a valued field embedding $f : K \to N$ for some subfield $K \subseteq M$, such that the induced map $vf : vK \to vN$ is a partial elementary map from vM to vN. If* $\operatorname{dom} f = M$*, we call $f : M \to N$ a* v-elementary map*, or say that f is* total*.*

Lemma 3.2.6. *If $M, N \models T$ and N is $|M|^+$-saturated, and f is a maximal partial v-elementary map from M to N, then f is total.*

Proof. Let K be the domain of f.

Claim 3.2.7. *K is henselian.*

Proof. Suppose not. As M and N are henselian, both contain the henselization of K. We can extend f to an isomorphism f' between the henselizations of K and $f(K)$. The henselization of K has the same value group as K, so $vf' = vf$ is still partial elementary. Then f' is a strictly larger v-elementary map, a contradiction. □

Claim 3.2.8. *Let $P(X)$ be an irreducible polynomial over K of degree greater than 1. If $P(X)$ has a root in M, then it does* not *have a root in N.*

Proof. Otherwise, let α be a root of $P(X)$ in M and β be a root in N. By basic field theory, there is an embedding of fields $f' : K(\alpha) \to f(K)(\beta)$ extending f, sending $\alpha \mapsto \beta$. This map f' must also be a map of valued fields, because there is a unique valuation on $K(\alpha)$ extending the valuation on K, by Claim 3.2.7.

We claim that f' is v-elementary. By saturation of vN, there is some elementary embedding $g : vM \to vN$ extending vf. The group homomorphism $g - vf'$ from $vK(\alpha) \to vN$ vanishes on vK, so it factors through the finite group $vK(\alpha)/vK$. As vN is torsion-free, $g - vf'$ vanishes on $vK(\alpha)$. Thus vf' is the restriction $g|vK(\alpha)$, so vf' is partial elementary, and f' is partial v-elementary. This contradicts the maximality of f. □

Claim 3.2.9. *Every element of $\mathcal{O}_K^\times$ is a pth power (in K). Consequently Kv is perfect*

Proof. Take $a \in \mathcal{O}_K^\times$. Then $X^p - a$ has a root in both M and N, so it has one in K. □

Claim 3.2.10. *Kv is separably closed.*

Proof. If not, let $f(X) \in \mathcal{O}_K[X]$ be a monic polynomial lifting a monic separable polynomial $\overline{f}(X) \in Kv[X]$ that has no root in K. The fields Mv and Nv are algebraically closed, so $\overline{f}(X)$ has roots in both Mv and Nv. Henselianity lifts these roots to roots of $f(X)$ in M and N. This contradicts Claim 3.2.8. □

Say that an embedding of groups $A \hookrightarrow B$ is pure if B/A is torsionless. Equivalently, for every prime ℓ and every $a \in A$, if a is a multiple of ℓ in B, then a is a multiple of ℓ in A.

Claim 3.2.11. *vK is pure in vM and vN*

Proof. Suppose γ is divisible by ℓ in one of vM or vN. As vf is partial elementary, γ is divisible by ℓ in both vM and vN. Take $a \in K$ with $v(a) = \gamma$. By Proposition 3.2.3.2, the polynomial $X^\ell - a$ has a root in both M and N. By Claim 3.2.8, $X^\ell - a$ is not irreducible over K. Then a has an ℓth root in K, so $v(a) = \gamma$ is divisible by ℓ in vK. □

Claim 3.2.12. *K is relatively algebraically closed in M and N.*

Proof. Let K_M and K_N be the relative algebraic closures of K in both fields. By Proposition 3.2.3.3, K_M and K_N are models of T. The value group extension $|vK_M/vK|$ is torsion, but vK is pure in vM, so the value group extension must be trivial. Similarly, $K_Mv = Kv$ because K_Mv is algebraic over Kv, but Kv is algebraically closed.

Therefore K_M is an immediate extension of K. Similarly, K_N is an immediate extension of K. By Lemma 3.2.4, K_M and K_N are isomorphic over K. This contradicts Claim 3.2.8 unless $K_M = K = K_N$. □

Claim 3.2.13. *$Kv = Mv$*

Proof. Otherwise, let t_M be an element of M whose residue is not in Kv. By saturation of N, we can find $t_N \in N$ with residue not in Kv.

Let f' be the map $K(t_M) \to f(K)(t_N)$ sending $t_M \mapsto t_N$ and extending f. This is a map of valued fields, because there is a unique valuation on $K(t)$ making t have transcendental residue. (Modulo quantifier elimination in ACVF, this is the statement that there is a unique type $p(x)$ that lives in the closed unit ball, but not in any smaller subballs. The uniqueness of this type follows by C-minimality.) The map f' is v-elementary, because $vf' = vf$ as $vK(t_M) = vK$. Then f' contradicts maximality. □

Claim 3.2.14. *$vK = vM$.*

Proof. Otherwise, take $\gamma_M \in vK \setminus vM$. Let g be an elementary embedding $vN \to vM$ extending vf. Let $\gamma_N = g(\gamma_M)$.

Let t_M (resp. t_N) be an element of M (resp. N) having valuation γ_M (resp. γ_N). The elements t_M and t_N are transcendental over K, so there is a map of fields $f' : K(t_M) \to$

$f(K)(t_N)$ extending f and sending $t_M \mapsto t_N$. This is a map of valued fields, because γ_M and γ_N define the same cut in vK, and there is a unique valuation on $K(t)$ making $v(t)$ land in this cut (again, this follows by C-minimality and quantifier elimination in ACVF).

We claim that f' is v-elementary, and that in fact vf' is $g|vK(t_M)$. By Abhyankar's inequality,

$$\frac{vK(t_M)}{vK + \mathbb{Z} \cdot v(t_M)}$$

is torsion. So it suffices to show that vf' and g agree on vK and $v(t_M)$. The former holds because f' extends f and g extends vf, and the latter holds by choice of t_N and γ_N. □

In summary, K is relatively algebraically closed in M and N, and M/K is an immediate extension. By Proposition 3.2.3.3, K is itself a model of T. In particular, K is defectless.

Now take $a \in M \setminus K$. In $M^{alg} \models ACVF$, let $\mathcal{B}$ be the chain of K-definable balls containing a.

Claim 3.2.15. *No element of K^{alg} is in the intersection $\bigcap \mathcal{B}$.*

Proof. By the proof that spherical completeness implies maximal completeness, no element of K is in the intersection. Suppose some element of K^{alg} were in the intersection. Let a' be such an element, of minimal degree over K. By the proof that maximal completeness implies spherical completeness, the extension $K(a')/K$ is immediate. But K is defectless, so it is algebraically maximal. □

By saturation of N, we can find some $a' \in N$ living in this intersection. Let f' be the map $K(a) \to f(K)(a')$ extending f and sending $a \mapsto a'$. By C-minimality and quantifier elimination in ACVF, there is a unique valuation on $K(t)$ making t live in each of the balls in $\mathcal{B}$. Consequently, f' preserves the valuation structure. Also, $vf' = vf$ because $vM = vK(a) = vK$. So f' is a strictly bigger v-elementary map, contradicting maximality. □

Theorem 3.2.16. *Let M and N be models of T. Let f be a partial v-elementary map from M to N. Then f is a partial elementary map. In other words, if K is a common subfield of M and N, and if vM and vN are elementarily equivalent over vK, then M and N are elementarily equivalent over K. Consequently, T has quantifier elimination relative to the value group.*

Proof. If M and N are models of T, Zorn's lemma applies to partial v-elementary maps between M and N. So the previous lemma yields

Claim 3.2.17. *Let M and N be models of T, and N be $|M|^+$-saturated. Then every partial v-elementary map from M to N can be extended to a total v-elementary map from M to N.*

Now suppose M, N, and K are as in the statement of the Theorem. Build a sequence $N_1, M_2, N_3, M_4, \ldots$ where

$$M \preceq M_2 \preceq M_4 \preceq \cdots$$

$$N \preceq N_1 \preceq N_3 \preceq \cdots$$

and M_{i+1} is $|N_i|^+$-saturated and N_{i+1} is $|M_i|^+$-saturated.

By repeatedly applpying the lemma, we can find partial v-elementary maps

$$M \to N_1 \to M_2 \to N_3 \to M_4 \to \cdots$$

extending the given embedding of K into N. These combine to yield an isomorphism between $\bigcup_i M_i$ and $\bigcup_i N_i$ over K. Then by Tarski-Vaught,

$$M \equiv_K \bigcup_i M_i \cong_K \bigcup_i N_i \equiv_K N$$

□

Corollary 3.2.18. *Let Γ be an ordered abelian group. If $\Gamma = p \cdot \Gamma$, the theory of henselian defectless valued fields (K, v) of characteristic p with $Kv \models ACF_p$ and $vK \equiv \Gamma$ is complete. If $a \in \mathrm{Int}_p \Gamma$, the theory of henselian defectless mixed characteristic fields with $(vK, v(p)) \equiv (\Gamma, a)$ and $Kv \models ACF_p$ is complete.*

Proof. For pure characteristic p, let M_1 and M_2 be two models. Let K be $\mathbb{F}_p$. Then vM_1 and vM_2 are elementarily equivalent over vK, so M_1 and M_2 are elementarily equivalent.

For mixed characteristic, take K to be $\mathbb{Q}$ instead. If $(vM_1, v(p)) \equiv (vM_2, v(p))$, then vM_1 and vM_2 are elementarily equivalent over $\mathbb{Q}$. □

3.3 Dp-minimality

Lemma 3.3.1. *Let Γ be an ordered abelian group such that $\Gamma/n\Gamma$ is finite for all $n > 0$. Then every unary definable subset of Γ is a boolean combination of cosets $\gamma + n\Gamma$ and definable cuts (upward closed sets).*

Proof sketch. Consider the expansion of this structure by all constants, and unary predicates for all cuts and all cosets of the forma $a + n\Gamma$. We claim that this expansion has elimination of quantifiers. By the usual methods, one reduces to eliminating quantifiers in a formula of the form

$$\exists x \bigwedge_{i=1}^{n} (c_i x + y_i \in a_i + d_i\Gamma) \wedge \bigwedge_{i=1}^{m} (e_i x + z_i < \Xi_i)$$

where the Ξ_i are cuts in Γ. Let d be the least common multiple of the nonzero d_i's. Breaking into cases, we may assume that one of the conjuncts explicitly specifies which coset of $d\Gamma$ contains x. Modulo this, the first type of conjunct $(c_i x + y_i \in a_i + d_i\Gamma)$ is equivalent to something not involving x. So we are left with something like

$$\exists x (x \in a + d\Gamma) \wedge \bigwedge_{i=1}^{m} (e_i x + z_i < \Xi_i)$$

After the change of variables $x = a + dx'$, we reduce to an expression of the form

$$\exists x \bigwedge_{i=1}^{m} (e_i x + z_i < \Xi_i)$$

Each conjunct cuts out a downward closed set or an upwards closed set, so the intersection is non-empty if each downward set intersects each upward set. That is, the statement is equivalent to a conjunction of statements of the form

$$\exists x (a_1 x + y < \Xi_1) \wedge (-a_2 x + z < \Xi_2) \tag{3.1}$$

where $a_1, a_2 > 0$. Let a be the least common multiple of a_1 and a_2. Break into cases by which coset of $a\Gamma$ contains y and which contains z. Within each case, the truth of (3.1) depends only on how $a_2 y + a_1 z$ compares to some cut in Γ.

Therefore, quantifier elimination holds in the expanded structure. Any definable set D in one variable is therefore a boolean combination of cuts and cosets of $n\Gamma$, where n depends only on D. Then, for any $a \in \Gamma$, the definable set

$$\{x \in \Gamma : a + nx \in D\}$$

is a boolean combination of cuts. These cuts can be taken to be definable (as the non-definable ones must ultimately be irrelevant). □

Let (K, v) be a valued field. We will call sets of the following forms *round sets* with center c:

$c + a \cdot (K^\times)^n$ for some $a \in K^\times$
$v^{-1}(\Xi)$ for some definable upward-closed subset $\Xi \subseteq vK \cup \{+\infty\}$

We will call sets of the first kind *angular sets* and set of the second kind *ball-like sets*. The class of round sets is closed under affine transformations.

Proposition 3.3.2. *Let K be a henselian defectless roughly p-divisible field, such that $Kv \models ACF_p$ and $vK/n \cdot vK$ is finite for all $n \in \mathbb{N}$. Then every unary definable set in K is a finite boolean combination of round sets.*

Proof. We may replace K with an elementary extension. First pass to an extension in which every coset of $\bigcap_n n \cdot vK$ is represented. Then pass to a spherical completion (which is an elementary extension by quantifier elimination).

Now look at 1-types. It suffices to show that a 1-type is determined by which round sets contain it. Let a be a singleton from an elementary extension of K. By spherical completeness, some element of K is maximally close to a. Translating a, we may assume that element is 0. If $a = 0$, then the 1-type is determined by the assertion that $v(x) = \infty$. Otherwise, $a \notin K$, and $rv(a)$ is new (not in $rv(K)$). If $v(a)$ is new, then $\mathrm{tp}(v(a)/vK)$

implies $\mathrm{tp}(a/K)$. Indeed, if $v(a) \equiv_{vK} v(a')$, then a and a' have the same type by quantifier elimination.

But by Lemma 3.3.1, $\mathrm{tp}(v(a)/vK)$ is implied by a collection of statements of the following forms:

- $v(a) + \gamma$ is divisible by n
- $v(a)$ is greater than some cut
- $v(a)$ is less than some cut

Each of these is a round set or the complement of a round set.

Otherwise, rescale a so that $v(a) = 0$. Then $\mathrm{res}(a)$ is new (not in $\mathrm{res}(K) = Kv$), and $\mathrm{tp}(a/K)$ is the generic type of the closed unit ball, which is unique by quantifier elimination. □

Definition 3.3.3. *Fix a complete theory T. Let $\mathcal{B}$ be an ind-definable family of unary definable sets. In other words, there is a collection of formulas Φ and for any model K,*

$$\mathcal{B}(K) = \{\phi(K; \vec{a}) \ : \ \phi(x; \vec{y}) \in \Phi, \ \vec{a} \in K^{|\vec{y}|}\}$$

Say that $\mathcal{B}$ is a unary basis *if it generates the family of all unary definable sets through boolean combinations. More precisely, if $K \models T$ and $D \subseteq K$ is K-definable, then D is in the boolean algebra generated by $\mathcal{B}(K)$.*

Say that $\mathcal{B}$ is a weak unary basis *if every unary definable set is a boolean combination of traces of externally definable sets in $\mathcal{B}$. In other words, if $K \models T$, then*

$$\{K \cap D' \ : \ D' \in \mathcal{B}(K'), \ K' \succeq K\}$$

generates a boolean algebra containing all definable subsets of K.

In the setting of Proposition 3.3.2, round sets form a unary basis. Moreover, balls and angular sets form a weak unary basis, because every ball-like round set is the trace of an externally definable ball.

Lemma 3.3.4. *Let T be a complete theory with infinite models. Let $\mathcal{B}$ be a weak unary basis for T. Then T is not dp-minimal if and only if in some model of T, there are mutually indiscernible sequences*

$$\ldots, X_{-1}, X_0, X_1, \ldots$$
$$\ldots, Y_{-1}, Y_0, Y_1, \ldots$$

of sets from $\mathcal{B}$, and an element a such that

$$a \in X_0 \iff a \in X_1$$
$$a \in Y_0 \iff a \in Y_1$$

Proof. If the given configuration occurs, it directly contradicts the characterization of dp-minimality in terms of mutually indiscernible sequences (one of the two sequences of sets must be a-indiscernible).

Conversely, suppose dp-minimality fails. Let $\mathcal{B}^+$ be the closure of $\mathcal{B}$ under boolean combinations.

Claim 3.3.5. *There is randomness pattern of depth 2 made of sets from $\mathcal{B}^+$.*

Proof. Take a mutually indiscernible randomness pattern of depth 2 and stretch the two sequences of sets to have length $\kappa := |T|^+$. So we have sets X_α and Y_β and elements $a_{\alpha,\beta}$ for $\alpha, \beta < \kappa$, such that

$$a_{\alpha,\beta} \in X_{\alpha'} \iff \alpha = \alpha'$$
$$a_{\alpha,\beta} \in Y_{\beta'} \iff \beta = \beta'$$

Let M be a small model defining the X's and Y's and containing the a's. In some $|M|^+$-saturated elementary extension $M^* \succeq M$, we can find sets X'_α and Y'_α from $\mathcal{B}^+(M^*)$ such that $X_\alpha \cap M = X'_\alpha \cap M$ and $Y_\alpha \cap M = Y'_\alpha \cap M$. As the $a_{\alpha,\beta}$ are in M,

$$a_{\alpha,\beta} \in X'_{\alpha'} \iff \alpha = \alpha'$$
$$a_{\alpha,\beta} \in Y'_{\beta'} \iff \beta = \beta'$$

Because $\kappa > |T|$, some subsequence of $\langle X'_\alpha \rangle_{\alpha<\kappa}$ is uniformly definable. Passing to this subsequence, and doing the same with $\langle Y'_\beta \rangle_{\beta<\kappa}$, we get a randomness pattern of depth 2 in M^*. □

Take the randomness pattern from the claim and derive a mutually indiscernible array from it, with each row indexed by $\mathbb{Z}$. This gives mutually indiscernible sequences

$$\ldots, U_{-1}, U_0, \ldots$$
$$\ldots, V_{-1}, V_0, \ldots$$

of sets from $\mathcal{B}^+$, and an element $a = a_{00}$ such that for all i,

$$a \in U_i \iff a \in V_i \iff i = 0$$

Recall the general

Fact 3.3.6. *Let $\langle x_i \rangle_{i \in I}$ and $\langle y_j \rangle_{j \in J}$ be mutually indiscernible sequences. Given a and b, we can find a_i and b_j such that*

- *$\langle a_i x_i \rangle_{i \in I}$ and $\langle b_j y_j \rangle_{j \in J}$ are mutually indiscernible.*
- *For each i, $a_i x_i \equiv a x_0$*
- *For each j, $b_j y_j \equiv b y_0$.*

(To see this, choose a_i^0 and b_j^0 arbitrarily so that $a_i^0 x_i \equiv a x_0$ and $b_i^0 y_j \equiv b y_0$. Then derive mutually indiscernible sequences from the two sequences $\langle a_i^0 x_i \rangle_{i \in I}$ and $\langle b_j^0 y_j \rangle_{j \in J}$. Finally move the new x's and y's back to the old ones via an automorphism.)

As $U_0 \in \mathcal{B}^+$, we can write

$$U_0 = f(B^1, B^2, \ldots, B^m)$$

where the B^i are in $\mathcal{B}$ and f is some boolean operation. Similarly, we can write $V_0 = g(C^1, C^2, \ldots, C^\ell)$ for some C^i in $\mathcal{B}$.

By the Fact, we can find B_i^j and C_i^j such that

- $U_i = f(B_i^1, \ldots, B_i^m)$ for each i
- $V_j = g(C_i^1, \ldots, C_i^\ell)$ for each i
- The two sequences

$$\langle \ulcorner U_i \urcorner \ulcorner B_i^1 \urcorner \cdots \ulcorner B_i^m \urcorner \rangle_{i \in \mathbb{Z}}$$
$$\langle \ulcorner V_i \urcorner \ulcorner C_i^1 \urcorner \cdots \ulcorner C_i^\ell \urcorner \rangle_{i \in \mathbb{Z}}$$

  are mutually indiscernible.

As $a \in U_0$ and $a \notin U_1$, there must be some j such that $a \in B_0^j \iff a \in B_1^j$. Likewise, there must be some k such that $a \in C_0^k \iff a \in C_1^k$. Take $X_i = B_i^j$ and $Y_i = B_i^k$. Then the X_i's and Y_i's are mutually indiscernible, are sets in $\mathcal{B}$, and satisfy

$$a \in X_0 \iff a \in X_1$$
$$a \in Y_0 \iff a \in Y_1$$

completing the proof of the lemma. □

Theorem 3.3.7. *Let (K, v) be a henselian defectless roughly p-divisible valued field, with $vK/n \cdot vK$ finite for all $n \in \mathbb{N}$, and $Kv \models ACF_p$. Then (K, v) is dp-minimal as a valued field.*

Proof. We may take (K, v) to be a monster model.

If dp-minimality failed, then by Lemma 3.3.4 there would exist an element a and two mutually indiscernible sequences of sets

$$\ldots, X_{-1}, X_0, X_1, \ldots$$
$$\ldots, Y_{-1}, Y_0, Y_1, \ldots$$

such that

$$a \in X_0 \iff a \in X_1$$
$$a \in Y_0 \iff a \in Y_1$$

and each X_i is a ball or an angular set. (Here we are using the fact that balls and angular sets form a weak unary basis, by Proposition 3.3.2.)

As v is henselian, there is a unique extension of v to K^{alg}. Consider the map χ from balls and angular sets in K to subsets of K^{alg} defined as follows:

- $\chi(B)$ is the ball in K^{alg} with the same center and radius as B if B is a ball in K
- $\chi(B) = \{c\}$ if B is an angular set centered on c

ACF and ACVF are dp-minimal, so (K^{alg}, v) is dp-minimal and one of the two sequences

$$\ldots, \chi(X_{-1}), \chi(X_0), \chi(X_1), \ldots$$
$$\ldots, \chi(Y_{-1}), \chi(Y_0), \chi(Y_1), \ldots$$

is a-indiscernible within (K^{alg}, v). Without loss of generality, $\langle \chi(X_i) \rangle_{i \in \mathbb{Z}}$ is a-indiscernible in K^{alg}.

If the X_i's are balls, then

$$a \in X_0 \iff a \in \chi(X_0) \iff a \in \chi(X_1) \iff a \in X_1,$$

a contradiction. So the X_i's are angular sets.

Write ξ_i for the center of X_i. Then $X_i - \xi_i$ is a coset of $(K^\times)^n$. As there are only finitely many of these cosets, the indiscernible sequence $\langle X_i - \xi_i \rangle_{i \in \mathbb{Z}}$ must be constant. So $X_i - \xi_i$ is some fixed coset of $(K^\times)^n$. By Proposition 3.2.3.2, $X_i - \xi_i = v^{-1}(S)$ for some set $S \subseteq vK$. Therefore, whether $a \in X_i$ depends solely on $v(a - \xi_i)$. Consequently,

$$v(a - \xi_0) \neq v(a - \xi_1)$$

Now in K^{alg}, the sequence $\ldots, \xi_{-1}, \xi_0, \xi_1, \ldots$ is a-indiscernible. So, perhaps after reversing the sequence, we have

$$\cdots < v(a - \xi_{-1}) < v(a - \xi_0) < v(a - \xi_1) < \cdots$$

This in turn implies that

$$v(a - \xi_1) = v(\xi_2 - \xi_1) \text{ and } v(a - \xi_0) = v(\xi_2 - \xi_0)$$

Whether an element x is in X_i depend only on $v(x - \xi_i)$, so

$$\xi_2 \in X_1 \iff a \in X_1 \not\iff a \in X_0 \iff \xi_2 \in X_0$$

But X_0 and X_1 have the same type over ξ_2 (the unique center of X_2), because the sequence $\langle X_i \rangle_{i \in \mathbb{Z}}$ is indiscernible *in K itself*. So we have a contradiction.

□

Chapter 4

Some remarks on strongly dependent valued fields

In [42], it is shown that all NIP fields are Artin-Schreier closed. We will use this fact to prove several nice properties of strongly dependent valued fields. Recall that

- A type-definable (or pro-definable) set X has *rudimentarily finite weight* if there is no randomness pattern of depth $\aleph_0$ in X (see §9.2.2 below)
- A theory is *strongly dependent* if every definable set has rudimentarily finite weight.

Superstable implies strongly dependent implies NIP.

The main results on strongly dependent valued fields are Theorem 4.3.1, which establishes some divisibility conditions on the value group, and Theorem 4.3.2, which establishes defectlessness of the valuation.

In what follows, we will repeatedly use the *Shelah expansion*. If M is an NIP structure, M^{sh} denotes the expansion of M by all externally definable sets. By [67] Proposition 3.23, M^{sh} eliminates quantifiers. Using this, one sees that if M is dp-minimal or strongly dependent, then so is M^{sh}. Of course, properties like saturation will probably be lost.

4.1 Perfection

Lemma 4.1.1. *Let K be a strongly dependent field. Then K is perfect.*

Proof. If K is imperfect, then there is a definable injection $f : K \times K \hookrightarrow K$, namely $f(x, y) = x^p + b \cdot y^p$ for any $b \notin K^p$.

Let $X_0 = K$ and let $X_{i+1} = f(X_0, X_i)$. Let X_∞ be the type definable set $\bigcap_i X_i$. In the category of $*$-definable sets, there is a surjection $X_\infty \to \prod_{i=0}^{\infty} K$, roughly sending

$$f(x_0, f(x_1, f(x_2, \ldots))) \mapsto (x_0, x_1, x_2, \ldots)$$

More precisely, note that $X_\infty = f(K, X_\infty)$, and

$$f : K \times X_\infty \xleftarrow{\sim} X_\infty$$

is a bijection. Let π_1 and π_2 be the two projections

$$X_\infty \xrightarrow{\sim} K \times X_\infty \twoheadrightarrow K$$
$$X_\infty \xrightarrow{\sim} K \times X_\infty \twoheadrightarrow X_\infty$$

Then the surjection $X_\infty \to \prod_{i=0}^{\infty} K$ is the map

$$x \mapsto (\pi_1(x), \pi_1(\pi_2(x)), \pi_1(\pi_2(\pi_2(x))), \ldots)$$

Since $\prod_{i=0}^{\infty} K$ does not have rudimentarily finite weight, neither does X_∞, nor its superset K. □

Remark 4.1.2. *Let K be a strongly dependent field of characteristic p, and L/K be a finite extension. Then p does not divide $[L : K]$.*

Proof. By perfection, L/K is a separable extension, so this follows by Corollary 4.5 in [42]. □

Lemma 4.1.3. *Let K be an infinite strongly dependent field of positive characteristic p. Then* any *valuation on K has p-divisible value group, and* any *henselian valuation on K is defectless.*

Proof. The first claim follows because K is Artin-Schreier closed, or because it is perfect. For the second claim, it is a general fact that if v is a henselian valuation, if L/K is a finite extension, and if w is a prolongation of v to L, then

$$[L : K] = [Lw : Kv] \cdot |wL/vK| \cdot p^d$$

for some $d \in \mathbb{N}$. By Remark 4.1.2, p does not divide $[L : K]$, so $d = 0$, i.e., the valuation is defectless. □

4.2 Finite ramification

Lemma 4.2.1. *Let (K, v) be a strongly dependent mixed characteristic valued field. Suppose the interval $[-v(p), v(p)]$ in the value group is finite. Then the residue field Kv is finite.*

Proof. We may replace K with a sufficiently saturated elementary extension. Note that Kv is itself strongly dependent, hence perfect.

Let $\mathcal{O}$ be the valuation ring and π generate its maximal ideal. Let $\hat{\mathcal{O}}$ denote the $*$-definable set

$$\hat{\mathcal{O}} = \varprojlim \mathcal{O}/(\pi^n)$$

Then $\mathcal{O}$ surjects onto $\hat{\mathcal{O}}$ via the obvious map.

Suppose the map $\hat{\mathcal{O}} \to Kv$ had a $*$-definable section

$$s : Kv \to \hat{\mathcal{O}}$$

We would then obtain a $*$-definable bijection

$$Kv \times \hat{\mathcal{O}} \tilde{\to} \hat{\mathcal{O}}$$

$$(\alpha, x) \mapsto s(\alpha) + \pi \cdot x,$$

This would then yield definable surjections

$$\mathcal{O} \to \hat{\mathcal{O}} \to Kv \times \hat{\mathcal{O}} \to Kv \times Kv \times \hat{\mathcal{O}} \to \cdots \to Kv \times Kv \times \cdots$$

showing that $\mathcal{O}$ is not strongly dependent unless Kv is finite.

So it suffices to produce a $*$-definable section of the projection $\hat{\mathcal{O}} \to Kv$. We will use the Teichmüller character.

Claim 4.2.2. *For each n, if $m > n$ and $\operatorname{res} y_1 = \operatorname{res} y_2 \neq 0$, then $y_1^{p^m} - y_2^{p^m} \in (\pi^n)$.*

Proof. Note first that if I is any principal proper ideal of $\mathcal{O}$, then $(1 + I)^p \subseteq 1 + J$ for a strictly smaller principal ideal, namely $J = I^2 + p \cdot I$. It follows that

$$(1 + (\pi))^{p^m} \subseteq 1 + (\pi^m)$$

Then for $y_1, y_2 \in \mathcal{O}^\times$,

$$y_1 - y_2 \in (\pi) \implies \frac{y_1}{y_2} \in 1 + (\pi) \implies \frac{y_1^{p^m}}{y_2^{p^m}} \in 1 + (\pi^m) \implies y_1^{p^m} - y_2^{p^m} \in (\pi^m) \subseteq (\pi^n)$$

□

Now define a section

$$s : (Kv)^\times \to \hat{\mathcal{O}}^\times$$

as follows: given nonzero $\alpha \in Kv$, choose a sequence $y_1, y_2, \ldots$ in $\mathcal{O}$ such that $(\operatorname{res} y_n)^{p^n} = \alpha$, using perfection of Kv. Then let

$$s(\alpha) = \lim_{n \to \infty} y_n^{p^n}$$

To see that this is well-defined and $*$-definable, note that for $m > n$, the class of $y_m^{p^m}$ modulo (π^n) does not depend on y_m, by the claim, nor on m, because for $m' > m$,

$$y_{m'}^{p^{m'-m}} - y_m \in (\pi) \text{ and so } y_{m'}^{p^{m'}} - y_m^{p^m} \in (\pi^n)$$

And $\operatorname{res}(s(\alpha)) = \alpha$, by choice of the y_i's.

Therefore there is a $*$-definable section of $\hat{\mathcal{O}}^\times \to (Kv)^\times$. We can extend this to a section of $\hat{\mathcal{O}} \to Kv$ by sending 0 to 0. □

4.3 The general statements

Recall Definition 3.1.4 (rough p-divisibility) and Remark 3.1.5 from Chapter 3.

Theorem 4.3.1. *Let (K, v) be a strongly dependent valued field. If Kv is infinite, then v is roughly p-divisible. If Kv is finite, then K has characteristic zero and the interval $[-v(p), v(p)]$ is finite. If v is henselian, then v is defectless.*

Proof. All the properties described here are elementary properties, so we may replace K with a sufficiently saturated elementary extension. We break into cases by the characteristic and residue characteristic of v.

In equicharacteristic 0, Kv is infinite, rough p-divisibility is vacuous, and henselian implies defectless.

In equicharacteristic p, Kv is infinite by Proposition 5.3 in [42]. The value group is p-divisible and the valuation is defectless if henselian, by Lemma 4.1.3.

This leaves the case of mixed characteristic. Let Δ_0 be the biggest convex subgroup not containing $v(p)$, and Δ be the smallest convex subgroup containing $v(p)$. These convex subgroups decompose the place $K \to Kv$ as a composition of three places:

$$K \overset{vK/\Delta}{\to} K_1 \overset{\Delta/\Delta_0}{\to} K_2 \overset{\Delta_0}{\to} Kv \tag{4.1}$$

where each arrow is labeled by its value group. The fields K and K_1 have characteristic zero, while K_2 and Kv have characteristic p.

Note that Δ/Δ_0 embeds into $\mathbb{R}$, so is small. Because K is sufficiently saturated, we get the following chain of implications

$$\begin{aligned}\Delta_0 = 0 \implies \Delta \text{ small} \implies [-v(p), v(p)] \text{ small} \implies [-v(p), v(p)] \text{ finite} \\ \implies \Delta_0 \text{ finite} \implies \Delta_0 = 0\end{aligned}$$

so Δ_0 vanishes if and only if $[-v(p), v(p)]$ is finite.

Both Δ_0 and Δ are externally definable, hence definable in K^{sh}. So the sequence of places in (4.1) is interpretable in the strongly dependent structure K^{sh}.

In particular, Δ_0 is p-divisible, by Proposition 5.4 of [42]. So $\mathrm{Int}_p\, vK$ is non-trivial or $[-v(p), v(p)]$ is finite.

Note that we have just proven the following general fact:

> If (K, v) is a strongly dependent mixed characteristic valued field, then $\mathrm{Int}_p\, vK$ is non-trivial or $[-v(p), v(p)]$ is finite

because this depends only on the elementary equivalence class of (K, v).

Combining with Lemma 4.2.1, we have actually shown

> If (K, v) is a strongly dependent mixed characteristic valued field, then $\mathrm{Int}_p\, vK$ is non-trivial or Kv is finite

In particular, we can apply this fact to the strongly dependent place $K_1 \to K_2$ in (4.1). We see that

$$\text{Int}_p(\Delta/\Delta_0) \text{ is non-trivial, or } K_2 \text{ is finite} \tag{4.2}$$

Now we prove the three claims of the theorem.

First suppose that Kv is infinite. Then K_2 is infinite, so Δ/Δ_0 has a non-trivial p-divisible convex subgroup by (4.2). Being archimedean, Δ/Δ_0 has very few convex subgroups and must be p-divisible itself. As Δ_0 is p-divisible, it follows that Δ is p-divisible, so v is roughly p-divisible.

Next suppose that Kv is finite. If $[-v(p), v(p)]$ is infinite, then Δ_0 is non-trivial, so $K_2 \to Kv$ is an infinite strongly dependent valued field of characteristic p with a finite residue field. This contradicts Proposition 5.3 of [42].

Next suppose that v is henselian. Then all three of the places in (4.1) are henselian. By the equicharacteristic cases, $K \to K_1$ and $K_2 \to Kv$ are defectless, so it remains to show that $K_1 \to K_2$ is defectless. Because (K, v) is saturated, any countable chain of balls in (K, v) has non-empty intersection. So the place $K \to Kv$ satisfies the countable intersection property of Remark 3.1.5. Therefore, so does $K_1 \to K_2$. However, the value group of $K_1 \to K_2$ is Δ/Δ_0. This group has countable cofinality, because it embeds into $\mathbb{R}$. Consequently, $K_1 \to K_2$ is spherically complete, hence defectless. □

A valued field (K, v), not necessarily henselian, is said to be defectless if for every finite extension L/K,

$$[L : K] = \sum_w |wL/vK| \cdot [Lw : Kv]$$

where the sum ranges over the distinct extensions of v to L. This is a first order condition, because the extensions of v to L are definable, by Beth implicit definability.

It can be shown that (K, v) is defectless if and only if its henselization is defectless.

Using these facts, we can drop the henselianity assumption from the previous result:

Theorem 4.3.2. *Let (K, v) be a strongly dependent valued field. Then (K, v) is defectless.*

Proof. First suppose K has characteristic p. Then p does not divide the degree of any extension of K, nor of its henselization K^{hens}. Therefore K^{hens} is automatically defectless. Otherwise, K has characteristic 0. If K has residue characteristic 0, so does its henselization, so K^{hens} and K are automatically defectless.

Assume therefore that K has mixed characteristic $(0, p)$. We may assume (K, v) is reasonably saturated, because defectlessness is a first-order condition. By Theorem 4.3.1, we are in one of two cases:

Case 1: finite ramification Then Kv is finite by Theorem 4.3.1. Coarsening by the convex subgroup generated by $v(p)$ yields a decomposition

$$K \to K_1 \to Kv$$

where K_1 has characteristic 0, and $K_1 \to Kv$ has value group isomorphic to $\mathbb{Z}$, and is spherically complete. Note that K_1 is a local field of characteristic 0.

Let $L \to Lv$ be the henselization of K. Decomposing according to the same convex groups yields $L \to L_1 \to Lv$. As $(L \to Lv)$ is an immediate extension of $(K \to Kv)$, it follows that

- $Lv = Kv$
- The value group of $L_1 \to Lv$ equals the value group of $K_1 \to Kv$.
- The value group of $L \to L_1$ equals the value group of $K \to K_1$

By the first two points, $(L_1 \to Lv)$ is an immediate extension of $(K_1 \to Kv)$. But $K_1 \to Kv$ is spherically complete, so $L_1 = K_1$. Consequently $L_1 \to Lv$ is spherically complete, hence defectless and henselian. Meanwhile, $L \to Lv$ is henselian, so $L \to L_1$ is henselian by Remark 3.1.5. As L_1 has residue characteristic 0, $L \to L_1$ is henselian and defectless. Finally $L \to L_1 \to Lv$ is defectless by Remark 3.1.5 because it is a composition of henselian defectless places.

Case 2: infinite ramification In this case, decompose $K \to Kv$ according to the smallest (resp. largest) convex subgroup containing (resp. not containing) $v(p)$, to obtain a decomposition

$$K \to K_1 \to K_2 \to Kv$$

where each field is strongly dependent, where K and K_1 have characteristic 0, where K_2 and Kv are characteristic p, and $K_1 \to K_2$ is a spherically complete rank 1 valuation.

Claim 4.3.3. *The prime p does not divide the degree of any finite extension of K_1*

Proof. As $K_1 \to K_2$ is spherically complete (hence henselian and defectless), it suffices to show that p does not divide the degree of any extension of K_2 or of the value group of $K_1 \to K_2$. The former is clear by Remark 4.1.2. The latter follows because the value group of $K_1 \to K_2$ is a quotient of the value group of $K_1 \to Kv$, which is the convex subgroup generated by $v(p)$, and Theorem 4.3.1 ensures that this group is p-divisible. □

Now let $L \to Lv$ be the henselization of $K \to Kv$, and let $L \to L_1 \to L_2 \to Lv$ be the decomposition obtained by the two aforementioned convex subgroups. Then p does not divide the degree of any finite extension of L_1, because L_1 is algebraic over K_1. Hence $L_1 \to Lv$ is automatically defectless. Also, $L \to L_1$ has residue characteristic 0 and is defectless. By Remark 3.1.5, the composition $L \to Lv$ is also defectless.

□

Chapter 5

Unary criteria

This chapter contains three unrelated results. In some sense, each involves generalizing a condition from the home sort to the entire theory. (Also, each tells us something about ACVF.)

Definition 5.0.1. *In a one-sorted structure M, a* unary *definable set is a definable subset of M^1.*

In §5.1 we give a criterion for the elimination of $\exists^\infty$ in T^{eq}. As far as I know, this result is new.

In §5.2 we prove that, in dense C-minimal structures, no infinite definable set admits a definable total order. We will use this fact later in Chapter 9. This fact is probably known to experts, but does not appear in the literature.

Finally, in §5.3 we give a "unary" criterion for a type in an NIP theory to be generically stable. Later, in Chapter 7, we will use this generalize the characterization of generically stable types in ACVF to its C-minimal expansions.

5.1 Elimination of $\exists^\infty$ in T^{eq}

Definition 5.1.1. *Let X be a definable or interpretable set. Then $\exists^\infty$ is eliminated on X if for every definable family*

$$\{D_a \subseteq X : a \in Y\}$$

of subsets of X, the set of a such that $|D_a| < \infty$ is definable. Equivalently (by compactness), there is an integer n such that for every $a \in Y$,

$$|D_a| < n \iff |D_a| < \infty$$

Definition 5.1.2. *A theory eliminates $\exists^\infty$ if it eliminates $\exists^\infty$ on the home sort, or equivalently, on powers of the home sort.*

The equivalence follows by

Observation 5.1.3. *If $\exists^\infty$ is eliminated on X and Y, it is eliminated on $X \times Y$. In fact, $S \subseteq X \times Y$ is finite if and only if both of the projections $S \to X$ and $S \to Y$ have finite image.*

This property does not pass from T to T^{eq}. For example, $\mathbb{Q}_p$ eliminates $\exists^\infty$. But the value group, which is interpretable, is a $\mathbb{Z}$-group, hence fails to eliminate T^{eq}.

In some theories, it is difficult to precisely pin down the imaginary sorts. Even when we have an explicit description, the imaginary sorts might be unwieldy. We would like to give a criterion for determining whether T^{eq} eliminates imaginaries, that can be checked without determining T^{eq} explicitly.

Definition 5.1.4. *An interpretable set X is a* set of unary imaginaries *if there is a definable relation $R \subseteq X \times \mathbb{M}$ such that the map*

$$x \mapsto R_x = \{m \in \mathbb{M} : (x, m) \in R\}$$

is an injection.

In other words, X is a set of unary imaginaries if the elements of X are codes for unary definable sets, in some uniform way.

Our main result is the following:

Theorem 5.1.5. *Suppose that $\exists^\infty$ is eliminated on every set of unary imaginaries. Then T^{eq} eliminates $\exists^\infty$.*

Proof. Suppose not. Let M be a small submodel of the monster, containing the parameters needed to witness some failure to eliminate $\exists^\infty$. Let N be the expansion of M obtained by adding $\mathbb{N} \cup \{\infty\}$ as a new sort, and adding non-standard counting functions on all interpretable families. By resplendence we can expand the monster model to a sufficiently saturated elementary extension of N.

So, we may assume that we have a set of nonstandard natural numbers, and a nonstandard "size" associated to any set interpretable in the original language. To avoid confusion, we will let L denote the original language.

Say that an L-interpretable set is *pseudofinite* if its non-standard size is less than the symbol ∞. Say that an L-interpretable set X is *wild* if it admits an infinite pseudofinite definable family of subsets.

The assumption of the theorem says that the home sort is tame. Because every interpretable set codes subsets of some power of the home sort, it suffices to show that every power of the home sort is tame. We do this in a few steps.

Claim 5.1.6. *If X is tame, so is any definable subset of X. If X and Y are tame, then so is $X \cup Y$.*

Proof. The first statement is clear, a fortiori. For the second, let $\mathcal{D}$ be a definable family of subsets of $X \cup Y$ which is pseudofinite. Note that $\{D \cap X : D \in \mathcal{D}\}$ is

- *pseudofinite*, because $\mathcal{D}$ is pseudofinite, and
- *finite*, because X is tame

Similarly, $\{D \cap Y : D \in \mathcal{D}\}$ is finite. Finally, the map

$$D \mapsto (D \cap X, D \cap Y)$$

yields an injection from $\mathcal{D}$ into a product of two finite sets. □

Claim 5.1.7. *Let $\pi : X \to Y$ be a definable map with finite fibers. If Y is tame, then so is X.*

Proof. By saturation, there is a uniform upper bound k on the size of the fibers. We proceed by induction on k. The base case $k = 1$ is trivial. Suppose $k > 1$. Let $\mathcal{D}$ be a pseudofinite definable family of subsets of X. Let

$$\mathcal{E} = \{\pi(D) : D \in \mathcal{D}\}$$

and

$$\mathcal{F} = \{\pi(X \setminus D) : D \in \mathcal{D}\}$$

Then $\mathcal{E}$ and $\mathcal{F}$ are both pseudofinite definable families of subsets of Y. By tameness of Y, they are both finite.

Now, it suffices to show that the fibers of $\mathcal{D} \to \mathcal{E} \times \mathcal{F}$ are finite. Replacing $\mathcal{D}$ with such a fiber, we may assume that $\pi(D)$ and $\pi(X \setminus D)$ are independent of D, as D ranges over $\mathcal{D}$. Let $U = \pi(D)$ and $V = \pi(X \setminus D)$ for any/every $D \in \mathcal{D}$. Let $Y' = U \cap V$ and $X' = \pi^{-1}(Y')$. Then the map $D \mapsto D \cap X'$ is injective on $\mathcal{D}$, because every element D of $\mathcal{D}$ contains $\pi^{-1}(U \setminus V)$ and is disjoint from $\pi^{-1}(V \setminus U)$. So it suffices to show that X' is tame. Let D be some arbitrary element of $\mathcal{D}$. Then $X' \cap D$ and $X' \setminus D$ each intersect every fiber of $X' \to Y'$, by choice of X'. In particular, the two maps

$$X' \cap D \to Y'$$

$$X' \setminus D \to Y'$$

have finite fibers of size less than k. By Claim 5.1.6, Y' is tame, and by induction, $X' \cap D$ and $X' \setminus D$ are tame. By Claim 5.1.6, X' is tame. □

Claim 5.1.8. *Suppose that $\pi : X \to Y$ is a definable surjection with finite fibers. Suppose that Y is tame. Then any pseudofinite definable set of sections of the surjection π is finite.*

Proof. A section is determined by its image. □

Claim 5.1.9. *Suppose X and Y are tame. Then so is $X \times Y$.*

Proof. Let $\mathcal{D}$ be a pseudofinite definable family of subsets of $X \times Y$. For each $x \in X$, the set $Y_x := \{x\} \times Y \subseteq X \times Y$ is tame, so the collection

$$E_x := \{D \cap Y_x : D \in \mathcal{D}\}$$

is finite. So $\pi : \bigcup_{x \in X} E_x \to X$ is a definable map of definable sets, with finite fibers. Each element $D \in \mathcal{D}$ induces a section of π, namely, the map σ_D sending a point $x \in X$ to (the code for) $D \cap Y_x$. This gives a definable injection from $\mathcal{D}$ to sections of π. By Claim 5.1.8 and the fact that X is tame, it follows that $\mathcal{D}$ is finite. □

This completes the proof of the Theorem. □

As an example, we apply this to C-minimal expansions of ACVF. Let T be a C-minimal expansion of ACVF, and K be a sufficiently saturated model of T. Again, it is handy to work in a setting with nonstandard counting functions.

Observation 5.1.10. *Let $B_1, \ldots, B_n$ be pairwise disjoint balls in K. Then the union $\bigcup_{i=1}^n B_i$ cannot be written as a boolean combination of fewer than n balls.*

This requires a little thought, but boils down to the fact that the residue field is infinite.

Lemma 5.1.11. *There is no pseudofinite infinite set of pairwise disjoint balls.*

Proof. Let $\mathcal{S}$ be such a set. By compactness, there must be some sequence $\mathcal{S}_1, \mathcal{S}_2, \ldots$ such that each $\mathcal{S}_i$ is a finite set of pairwise disjoint balls, the $\mathcal{S}_i$ are uniformly definable, and $\lim_{i \to \infty} |\mathcal{S}_i| = \infty$.

The unions $U_i = \bigcup \mathcal{S}_i \subseteq K$ are uniformly definable (bounded in complexity), so there is some absolute bound on the number of balls needed to express U_i. But Observation 5.1.10 says that this number is at least $|\mathcal{S}_i|$, a contradiction. □

Because the value group Γ of K is dense o-minimal, it eliminates $\exists^\infty$. Therefore, there are no pseudofinite infinite subsets of Γ.

Lemma 5.1.12. *There is no pseudofinite infinite set of balls.*

Proof. Let $\mathcal{S}$ be such a set. Let $\mathcal{S}_0$ be the set of minimal elements of $\mathcal{S}$. For each $B \in \mathcal{S}_0$, let $\mathcal{S}_B$ denote the elements of $\mathcal{S}$ containing B. In a (pseudo)finite poset, every element is greater than or equal to a minimal elements, so

$$\mathcal{S} = \bigcup_{B \in \mathcal{S}_0} \mathcal{S}_B$$

The set $\mathcal{S}_0$ is pseudofinite, hence finite by Lemma 5.1.11. Therefore, $\mathcal{S}_B$ is infinite for some B.

Now $\mathcal{S}_B$ is a chain of balls. Therefore the map sending a ball B' to its radius is a 2-to-1 map from $\mathcal{S}_B$ into Γ. The range of this map is pseudofinite, hence finite. Therefore, so is the domain $\mathcal{S}_B$, a contradiction. □

Finally, suppose that $\exists^\infty$ is not eliminated on some set X_0 of unary codes. Then there is a pseudofinite infinite set $A \subseteq X_0$. Let D_a be the unary set associated to $a \in A$. For each a, there is a unique minimal set of balls $\mathcal{B}_a$ such that D_a can be written as a boolean combination of $\mathcal{B}_a$. The correspondence $a \mapsto \mathcal{B}_a$ is a definable finite-to-finite correspondence from A to $\mathcal{B}$. Its "image"

$$I := \bigcup_{a \in A} \mathcal{B}_a$$

is pseudofinite, hence finite. The boolean algebra generated by I is finite, and contains every D_a, so A is finite, a contradiction.

Consequently, we see that T^{eq} eliminates $\exists^\infty$ for any C-minimal expansion of ACVF. This is useful, because the expansions of ACVF by analytic functions are known to have exotic imaginaries.

5.2 Definable orders in C-minimal structures

Definition 5.2.1. *A structure M* defines no total orders *if there are no infinite definable sets admitting definable total orders.*

This condition can be checked on unary sets:

Lemma 5.2.2. *Suppose no infinite unary definable set admits a definable total ordering. Then M defines no total orders.*

Proof. Let $\mathcal{C}$ be the class of definable sets X such that there is an infinite definable set $D \subseteq X$ admitting a definable total order.

Claim 5.2.3. *If $X \times Y \in \mathcal{C}$, at least one of X and Y is in $\mathcal{C}$.*

Proof. Given $D \subseteq X \times Y$ infinite with a definable total ordering, consider the projection $\pi : D \to X$. Each fiber of π embeds definably into Y, so if some fiber of π is infinite, then $Y \in \mathcal{C}$. Otherwise, the fibers are all finite. Let $g : \pi(D) \to D$ pick out the least element of each fiber. We can pull the ordering on D back to $\pi(D)$ along g. Then the infinite subset $\pi(D)$ of X admits a definable total ordering, so $X \in \mathcal{C}$. □

Consequently, if $M^1 \notin \mathcal{C}$, then $M^n \notin \mathcal{C}$. □

In the rest of this section, we will show that dense C-minimal structures define no total orders. For the sake of contradiction, suppose that there is an infinite unary definable set X admitting a total ordering.

Definition 5.2.4. *A* tree-like set *is a non-empty finite set S such that for every ball B, $|S \cap B|$ is 0 or a power of 2.*

Note that $|S| = 2^n$ for some n, because the entirety of M is a ball. We call n the depth of S.

Lemma 5.2.5. *Let B_1 and B_2 be disjoint balls. Let $S_i \subseteq B_i$ be a tree-like set of depth n. Then $S_1 \cup S_2$ is a tree-like set of depth $n+1$.*

Proof. Let B be any ball. If B intersects only B_i for $i = 1$ or 2, then $|B \cap S| = |B \cap S_i|$ has the desired form. Otherwise, B intersects both of B_i, hence contains both. So $B \cap S = S$ and $|B \cap S| = |S_1| + |S_2| = 2^{n+1}$. □

Lemma 5.2.6. *An infinite definable set X contains arbitrarily big tree-like sets.*

Proof. The density assumption ensures that every infinite ball contains two disjoint infinite subballs. By induction on n, we see that

$$\text{If } B \text{ is an infinite ball, then } B \text{ contains a tree-like set of depth } n$$

The density assumption also ensures that X contains an infinite ball. □

For any definable set D, the characteristic function χ_D of D can be written as

$$\chi_D = \sum_{i=1}^{m} a_i \cdot \chi_{B_i}$$

where the B_i's are balls and $a_i \in \{-1, 0, 1\}$. This is an easy consequence of the swiss cheese decomposition. Call the least such m the "complexity" of D. By compactness, complexity is bounded in definable families.

If S is a tree-like set of depth n, then

$$|S \cap D| = \sum_{s \in S} \chi_D(s) = \sum_{s \in S} \sum_{i=1}^{m} a_i \cdot \chi_{B_i}(s) = \sum_{i=1}^{m} a_i \cdot |S \cap B_i| = \sum_{i=1}^{m} a_i' \cdot 2^{k_i}$$

for some $a_i' \in \{-1, 0, 1\}$ and some $k_i \in \{1, \ldots, n\}$.

In particular, as D ranges through sets of complexity m, there are only $(3n)^m$ possibilities for $|S \cap D|$.

On the other hand, as D ranges through the half-infinite intervals $(-\infty, a) \subseteq X$, there are at least 2^n possibilities for $|S \cap D|$.

Letting m bound the complexity of the half-infinite intervals, and n be large enough that $2^n > (3n)^m$, and $S \subseteq X$ be a tree-like set of depth n, we get a contradiction.

We have shown

Proposition 5.2.7. *Dense C-minimal structures never define total orders.*

Corollary 5.2.8. *Dense C-minimal structures never eliminate imaginaries.*

Proof. Any C-minimal structure interprets the set of balls. Within this, the set of balls around a given point is totally ordered, and infinite under the density assumption. So any dense C-minimal structure interprets an infinite total order, but defines no infinite total order. □

5.3 Generic stability and chain aversion

Let $\mathbb{M}$ be a monster model of an NIP complete theory T.

Definition 5.3.1. *A global invariant type $p(x)$ is* generically stable *if $p(x) \otimes p(y) = p(y) \otimes p(x)$.*

For other equivalent definitions of generic stability, see Section 3 of [36].

Definition 5.3.2. *A global invariant type $p(x)$ is* chain averse *if for every small set C, there is a cardinal $\kappa(C)$ such that for any ordinal λ and any $\vec{a} \models p^{\otimes\lambda}|C$, there is no chain of size $\kappa(C)$ of unary $C\vec{a}$-definable sets.*

Theorem 5.3.3. *Let $p(x)$ be a global invariant type. Then $p(x)$ is chain averse if and only if $p(x)$ is generically stable.*

We will prove this in several steps.

Observation 5.3.4. *If $p(x)$ is chain averse, so is $p^{\otimes n}(x_1, \ldots, x_n)$ for all $n < \omega$.*

Lemma 5.3.5. *Let $p(x)$ be a generically stable type. Then $p(x)$ is chain averse.*

Proof. Let C be a small subset. We may assume that p is C-invariant, by enlarging C. Let $\kappa = \kappa(C)$ be $(\aleph_0 + |T| + |C|)^+$. Suppose for the sake of contradiction that there is a Morley sequence $\vec{a}$ in the type p, over C, and a chain $\mathcal{U}$ of size κ of $C\vec{a}$-definable unary sets. Let λ be the length of $\vec{a}$. Note that $\lambda \geq \kappa$.

Each set $U \in \mathcal{U}$ can be written as $f(a_{\alpha_1}, \ldots, a_{\alpha_n})$ for some n, some C-definable function f, and some $\alpha_1 < \cdots < \alpha_n < \lambda$. There are at most $\aleph_0 + |T| + |C|$ possibilities for f and n. By the pigeonhole principle, we may assume that n and f are constant across all U, by passing to a subchain of $\mathcal{U}$ if necessary.

As $\mathcal{U}$ is a chain containing more than one element, we can find

$$\alpha_1 < \cdots < \alpha_n < \lambda$$
$$\beta_1 < \cdots < \beta_n < \lambda$$

such that

$$A = f(a_{\alpha_1}, \ldots, a_{\alpha_n}) \subsetneq f(a_{\beta_1}, \ldots, a_{\beta_n}) = B$$

Because p is generically stable and C-invariant, the sequence $\vec{a}$ is totally C-indiscernible; therefore so is the set

$$S = \{a_{\alpha_1}, \ldots, a_{\alpha_n}\} \cup \{a_{\beta_1}, \ldots, a_{\beta_n}\}.$$

Therefore, there is an automorphism $\sigma \in \mathrm{Aut}(\mathbb{M}/C)$ permuting S and sending $a_{\alpha_i} \mapsto a_{\beta_i}$ for each i. In particular $\sigma(A) = B$. As S is finite, there is some k such that $\sigma^k|S = id$. In particular, $\sigma^k(A) = A$. Then

$$A \subsetneq B = \sigma(A),$$

so

$$A \subsetneq \sigma(A) \subsetneq \sigma(\sigma(A)) \subsetneq \cdots \subsetneq \sigma^{k-1}(A) \subsetneq \sigma^k(A) = A$$

which is absurd. □

Lemma 5.3.6. *Let C be a small set and $\langle a_\alpha \rangle_{\alpha<\kappa}$ be a C-indiscernible sequence. Let $\phi(x;y)$ be a C-formula with $|y| = 1$. Suppose there are no $C\vec{a}$-definable chains of unary sets of length κ. Then*

$$\phi(a_0, y) \wedge \neg\phi(a_1, y) \text{ is consistent} \tag{5.1}$$

if and only if

$$\phi(a_1, y) \wedge \neg\phi(a_0, y) \text{ is consistent} \tag{5.2}$$

Proof. Note that (5.1) is inconsistent if and only if $\phi(a_0, \mathbb{M}) \subseteq \phi(a_1, \mathbb{M})$, and (5.2) is inconsistent if and only if $\phi(a_1, \mathbb{M}) \subseteq \phi(a_0, \mathbb{M})$. If exactly one of (5.1) and (5.2) is inconsistent, then

$$\phi(a_0, \mathbb{M}) \subsetneq \phi(a_1, \mathbb{M}) \text{ or } \phi(a_1, \mathbb{M}) \subsetneq \phi(a_0, \mathbb{M})$$

Either way, by indiscernibility, the sets $\phi(a_\alpha, \mathbb{M})$ form a chain of length κ. □

Lemma 5.3.7. *Let $p(X)$ be a chain averse C-invariant type. Let $a_1, \ldots, a_n$ realize $p^{\otimes n}$. Let $\phi(x;y)$ be a C-formula with $|y| = 1$. Let $\psi(y)$ be a formula over $Ca_1, \ldots, a_{i-1}, a_{i+2}, \ldots, a_n$. Then*

$$\psi(y) \wedge \phi(a_i, y) \wedge \neg\phi(a_{i+1}, y)$$

is consistent if and only if

$$\psi(y) \wedge \neg\phi(a_i, y) \wedge \phi(a_{i+1}, y)$$

is consistent.

Proof. Let $\kappa = \kappa(C)^+$. For any total order $(I, <)$ extending $\{1, \ldots, n\}$, we can find a Morley sequence $\langle a_i \rangle_{i \in I}$ in the type p over C, extending the a_i's.

Consequently, we can find $\langle b_\alpha \rangle_{\alpha<\kappa}$ such that $b_0 = a_i$, $b_1 = a_{i+1}$, and

$$a_1, \ldots, a_{i-1}, a_i = b_0, a_{i+1} = b_1, b_2, \ldots, b_\omega, b_{\omega+1}, \ldots, \ldots, a_{i+1}, \ldots, a_n$$

is a Morley sequence over C.

In particular, it is C-indiscernible, and $C\vec{a}\vec{b}$ defines no long chains of unary sets.

Now, the sequence $\langle b_\alpha \rangle_{\alpha<\kappa}$ is C'-indiscernible for

$$C' = Ca_1, \ldots, a_{i-1}a_{i+2}, \ldots, a_n$$

and there are no long $C'\vec{b}$-definable chains. The formula $\psi(y)$ is over C'. Let $\phi'(x;y)$ be $\phi(x;y) \wedge \psi(y)$. Then, applying Lemma 5.3.6 to C', ϕ', and $\vec{b}$, we see that

$$\exists y(\phi'(b_0, y) \wedge \neg\phi'(b_1, y)) \iff \exists y(\phi'(b_1, y) \wedge \neg\phi'(b_0, y))$$

But

$$\phi'(b_0, y) \wedge \neg\phi'(b_1, y) = \psi(y) \wedge \phi(a_i, y) \wedge \neg\phi(a_{i+1}, y)$$

and

$$\phi'(b_1, y) \wedge \neg\phi'(b_0, y) = \psi(y) \wedge \neg\phi(a_i, y) \wedge \phi(a_{i+1}, y)$$

completing the proof of the Lemma. □

Lemma 5.3.8. *Let $p(x)$ be a chain averse C-invariant type. Let $a_1, \ldots, a_n$ realize $p^{\otimes n}|C$. Let $\phi(x; y)$ be a C-formula with $|y| = 1$. Let $\phi^0 = \phi$ and $\phi^1 = \neg\phi$. For each word $w \in \{0, 1\}^n$, let*

$$\Sigma_w(y) = \bigwedge_{i=1}^{n} \phi^{w_i}(a_i; y)$$

If w and w' are permutations of each other, then $\Sigma_w(y)$ is consistent if and only if $\Sigma_{w'}(y)$ is consistent.

Proof. We reduce to the case where w and w' are related by a transposition, so

$$w = a01b \text{ and } w' = a10b$$

for some words a and b. Let $i = |a| + 1$ and

$$\psi(y) = \bigwedge_{1 \leq j \leq n, j \notin \{i, i+1\}} \phi^{w_j}(a_j; y)$$

Then

$$\Sigma_w(y) = \psi(y) \wedge \phi(a_i; y) \wedge \neg\phi(a_{i+1}; y)$$

and

$$\Sigma_{w'}(y) = \psi(y) \wedge \neg\phi(a_i; y) \wedge \phi(a_{i+1}; y)$$

so the result follows by Lemma 5.3.7. □

Lemma 5.3.9. *Let $p(x)$ be a chain averse C-invariant type. Let $\vec{a} \models p^{\otimes\omega}|C$. Let b be a singleton and $\phi(x; y)$ be a C-formula with $|y| = 1$. If $\phi(x; b) \in p(x)|Cb$, then $\models \phi(a_i; b)$ for all but finitely many i.*

Proof. Let $\alpha_1, \alpha_2, \ldots$ realize $p^{\otimes\omega}|Cb\vec{a}$. Then $\vec{a}\vec{\alpha} \models p^{\otimes 2\omega}|C$. Also, $\models \phi(\alpha_i; b)$ for all i.

Suppose $\phi(a_i; b)$ fails for infinitely many i. Then there are infinitely many terms a in the sequence $\vec{a}\vec{\alpha}$ such that $\phi(a; b)$ holds, and infinitely for which it fails. Let n bound the alternation number of $\phi(x; y)$. Let

$$c_1, \ldots, c_{2n}$$

be a subsequence of $\vec{a}\vec{\alpha}$ containing n elements which satisfy $\phi(x; b)$, and n elements which satisfy $\neg\phi(x; b)$.

For some word $w \in \{0, 1\}^{2n}$ containing n 1's and n 0's, the type $\Sigma_w(y)$ of Lemma 5.3.8 is realized by b. Let w' be $(01)^n$. Then w' is a permutation of w, so $\Sigma_{w'}(y)$ is consistent, contradicting the choice of n. □

Lemma 5.3.10. *Let $p(x)$ be a chain averse C-invariant type. Let κ be a cardinal of cofinality greater than $|C| + |T|$. Let $\langle a_\alpha \rangle_{\alpha<\kappa}$ be a Morley sequence in p over C of length κ. For any singleton b, some tail*

$$\langle a_\alpha \rangle_{\beta<\alpha<\kappa}$$

is a Morley sequence in p over Cb.

Proof. By Observation 5.3.4, each finite power $p^{\otimes n}$ is chain averse (and C-invariant, of course).

Claim 5.3.11. *For each C-formula $\phi(x_1, \ldots, x_n, y)$ with $|y| = 1$, there is $\beta_\phi < \kappa$ such that*

$$\models \phi(a_{\alpha_1}, \ldots, a_{\alpha_n}, b) \iff \phi(x_1, \ldots, x_n, b) \in p^{\otimes n}|Cb$$

for all $\beta_\phi < \alpha_1 < \cdots < \alpha_n$.

Proof. If no such β_ϕ exists, we can find an increasing sequence

$$\alpha_{0,1} < \cdots < \alpha_{0,n} < \alpha_{1,1} < \cdots \alpha_{1,n} < \alpha_{2,1} < \cdots < \alpha_{2,n} < \cdots$$

such that for each i,

$$\models \phi(a_{\alpha_{i,1}}, \ldots, a_{\alpha_{i,n}}, b) \not\iff \phi(x_1, \ldots, x_n, b) \in p^{\otimes n}|Cb$$

Let $c_i = a_{\alpha_{i,1}} a_{\alpha_{i,2}} \cdots a_{\alpha_{i,n}}$. Then

$$c_0, c_1, c_2, \ldots$$

is a Morley sequence in $p^{\otimes n}$ over C. Replacing ϕ with $\neg\phi$, we may assume that $\phi(\vec{x}, b) \in p^{\otimes n}|Cb$. Then $\neg\phi(\vec{c}_i, b)$ holds for infinitely many i, contradicting Lemma 5.3.9 applied to the chain-averse type $p^{\otimes n}$. □

By choice of κ, there is some $\beta < \kappa$ greater than all the β_ϕ. We claim that $\langle a_\alpha \rangle_{\beta<\alpha<\kappa}$ realizes $p^{\otimes\kappa}|Cb$. If not, there exists $n \in \mathbb{N}$ and

$$\beta < \alpha_1 < \cdots < \alpha_n < \kappa$$

such that $a_{\alpha_1} \cdots a_{\alpha_n} \not\models p^{\otimes n}|Cb$. In particular, there is some C-formula $\phi(x_1, \ldots, x_n, y)$ such that

$$\phi(x_1, \ldots, x_n, b) \in p^{\otimes n}|Cb \not\iff \models \phi(a_{\alpha_1}, \ldots, a_{\alpha_n}, b).$$

This contradicts the claim, as $\beta_\phi < \alpha_1 < \cdots < \alpha_n$. □

Lemma 5.3.12. *Let $p(x)$ be a C-invariant chain averse type. Let b be any finite tuple or imaginary. Let κ be a cardinal with cofinality greater than $|C| + |T|$. Let $\langle a_\alpha \rangle_{\alpha<\kappa}$ realize $p^{\otimes\kappa}|C$. Then some tail of $\vec{a}$ realizes $p^{\otimes\kappa}|Cb$.*

Proof. If suffices to replace b with some b' such that $b \in \mathrm{dcl}^{eq}(b')$, so we may assume that b is a tuple (not an imaginary). Let $b = b_1 b_2 \cdots b_n$. By Lemma 5.3.10, there is some tail $\vec{a}'$ realizing $p^{\otimes\kappa}|Cb_1$. As $p(x)$ is Cb_1-invariant, by Lemma 5.3.10 there is some tail $\vec{a}''$ of $\vec{a}$ realizing $p^{\otimes\kappa}|Cb_1b_2$. Continuing on in this fashion completes the proof. □

Finally, we complete the proof of Theorem 5.3.3, saying that chain aversion and generic stability are the same thing.

Proof. If $p(x)$ is generically stable, it is chain averse by Lemma 5.3.5. Conversely, suppose $p(x)$ is a C-invariant chain averse type. Let κ be a cardinality with cofinality greater than $|C| + |T|$, and let $\langle a_\alpha \rangle_{\alpha<\kappa+\kappa}$ be a Morley sequence in p over C. Applying Lemma 5.3.12 to the first half $\langle a_\alpha \rangle_{\alpha<\kappa}$ of the sequence, and the imaginary a_κ, we find some $\beta < \kappa$ such that

$$\langle a_\alpha \rangle_{\beta<\alpha<\kappa} \models p^{\otimes\kappa}|Ca_\kappa$$

In particular, $a_{\beta+1} \models p|Ca_\kappa$. Thus

$$(a_\kappa, a_{\beta+1}) \models p^{\otimes 2}|C$$

But by choice of the a's, we also have

$$(a_{\beta+1}, a_\kappa) \models p^{\otimes 2}|C$$

Consequently, $p(x_1) \otimes p(x_2)|C = p(x_2) \otimes p(x_1)|C$. We can replace C with any superset in the above argument, and so $p(x_1) \otimes p(x_2) = p(x_2) \otimes p(x_1)$. Therefore $p(x)$ is generically stable. □

In §7.2 we will apply this to ACVF and its C-minimal expansions, characterizing the generically stable types as the types orthogonal to the value group.

Chapter 6

On the proof of elimination of imaginaries in algebraically closed valued fields

ACVF is the theory of non-trivially valued *a*lgebraically *c*losed *v*alued *f*ields. This theory is the model companion of the theory of valued fields. ACVF does not have elimination of imaginaries in the home sort (the valued field sort). Nevertheless, Haskell, Hrushovski, and Macpherson in [26] were able to find a collection of "geometric sorts" in which elimination of imaginaries holds.

Let K be a model of ACVF, with valuation ring $\mathcal{O}$ and residue field k. A *lattice* in K^n is an $\mathcal{O}$-submodule $\Lambda \subseteq K^n$ isomorphic to $\mathcal{O}^n$. Let S_n denote the set of lattices in K^n. This is an interpretable set; it can be identified with $GL_n(K)/GL_n(\mathcal{O})$. For each lattice $\Lambda \subseteq K^n$, let $\operatorname{res}\Lambda$ denote $\Lambda \otimes_{\mathcal{O}} k$, a k-vector space of dimension n. Let T_n be

$$T_n = \bigcup_{\Lambda \in S_n} \operatorname{res}\Lambda = \{(\Lambda, x) : \Lambda \in S_n,\ x \in \operatorname{res}\Lambda\}.$$

This set is again interpretable.

The main result of [26] is the following:

Theorem 6.0.1 (Haskell, Hrushovski, Macpherson)**.** *ACVF eliminates imaginaries relative to the sorts K, $\{S_n : n \geq 1\}$ and $\{T_n : n \geq 1\}$.*

The proof in [26] is long and technical, and we aim to give a more straightforward proof. Our proof is a variant of Hrushovski's shorter proof in [32], except that our strategy for coding definable types is different—see 6.4.2. We also give a slightly simpler proof that finite sets of modules can be coded in the geometric sorts—see 6.4.3.

Obviously, we prove no new results. We include many details that are well-known at this point, for the sake of being self-contained. The proof of elimination of imaginaries given here is hopefully more conceptual than previous proofs. At any rate, it manages to do without internality, germs, unary codes, type-definable torsors, and upper triangular matrices.

6.1 Review of ACVF

6.1.1 Notation

In a model of ACVF, K is the home sort (the valued field), $\mathcal{O} \subseteq K$ is the valuation ring, $\mathfrak{M}$ is the maximal ideal in $\mathcal{O}$, $k = \mathcal{O}/\mathfrak{M}$ is the residue field, $\Gamma = K^\times/\mathcal{O}^\times$ is the valuation group, $\mathrm{res} : \mathcal{O} \to k$ is the residue map, and $\mathrm{val} : K \to \Gamma \cup \{+\infty\}$ is the valuation. The value group is written additively, and ordered so that

$$\mathcal{O} = \{x \in K : \mathrm{val}(x) \geq 0\}.$$

A *lattice* in K^n is an $\mathcal{O}$-submodule Λ of K^n which is free of rank n, i.e., isomorphic to $\mathcal{O}^n$. If Λ is a lattice, $\mathrm{res}\,\Lambda$ will denote $\Lambda/\mathfrak{M}\Lambda = \Lambda \otimes_{\mathcal{O}} k$. This is always an n-dimensional k-vector space. We will use the following interpretable sets:

- S_n, the set of lattices in K^n.
- T_n, the set of pairs (Λ, ξ), where $\Lambda \in S_n$ and $\xi \in \mathrm{res}\,\Lambda$
- $R_{n,\ell}$, the set of pairs (Λ, V), where $\Lambda \in S_n$ and V is an ℓ-dimensional subspace of $\mathrm{res}\,\Lambda$.

Each of these sets is easily interpretable in ACVF. Our main goal will be to prove that ACVF has elimination of imaginaries in the sorts K and $R_{n,\ell}$. In §6.5, we will note how this implies elimination of imaginaries in K, S_n, and T_n, the standard "geometric sorts" of [26]. **But until then, the term "geometric sorts" will mean the sorts K and $R_{n,\ell}$.**

When working in an abstract model-theoretic context, the monster model will be denoted $\mathbb{M}$. If a definable set or other entity X has a code in $\mathbb{M}^{eq}$, the code will be denoted $\ulcorner X \urcorner$. Unless stated otherwise, "definable" will mean "interpretable."

6.1.2 Basic facts

We assume without proof the following well-known facts about ACVF. Many of these are discussed in [49].

- Models of ACVF are determined up to elementary equivalence by characteristic and residue characteristic, which must be $(0,0)$, (p,p), or $(0,p)$ for some prime p.
- ACVF has quantifier elimination in the language with one sort K, with the ring structure on K, and with a binary predicate for the relation $\mathrm{val}(x) \geq \mathrm{val}(y)$.
- C-minimality: Every definable subset $D \subseteq K^1$ is a boolean combination of open and closed balls (including points). More precisely, D can be written as a disjoint union of "swiss cheeses," where a swiss cheese is a ball with finitely many proper subballs removed. There is a canonical minimal way of decomposing D as a disjoint union of swiss cheeses. All the balls involved in this decomposition are algebraic over the code for D.

- The theory ACVF does not have the independence property. That is, ACVF is NIP.
- The value group Γ is o-minimal, in the sense that every definable subset of Γ^1 is a finite union of points and intervals with endpoints in $\Gamma \cup \{\pm\infty\}$. (In fact, Γ is a stably embedded pure divisible ordered abelian group.)
- The residue field k is strongly minimal, hence stable and stably embedded. Moreover, every definable subset of k^n is coded by a tuple from k. (In fact, k is a stably embedded pure algebraically closed field.)

The first two points are due to Robinson [63], and the third is due to Holly [30]. The last three points are easy consequences of C-minimality[1], and the last two can also be seen from the quantifier elimination result in the three-sorted language discussed in [49] and [26].

6.1.3 Valued K-vector spaces

Let K be an *arbitrary* valued field. Following Section 2.5 of [32],

Definition 6.1.1. *A* valued K-vector space *is a K-vector space V and a set $\Gamma(V)$ together with the following structure:*

- *A total ordering on $\Gamma(V)$*
- *An action*
$$+ : \Gamma(K) \times \Gamma(V) \to \Gamma(V)$$
of $\Gamma(K) = \Gamma$ on $\Gamma(V)$, strictly order-preserving in each variable (hence free)
- *A surjective map* $\mathrm{val} : V \setminus \{0\} \to \Gamma(V)$*, such that*
$$\mathrm{val}(w + v) \geq \min(\mathrm{val}(w), \mathrm{val}(v))$$
$$\mathrm{val}(\alpha \cdot v) = \mathrm{val}(\alpha) + \mathrm{val}(v)$$
for $w, v \in V$ and $\alpha \in K$, with the usual convention that $\mathrm{val}(0) = +\infty > \Gamma(V)$.

This is merely a variation on the notion of a normed vector space over a field with an absolute value.

Remark 6.1.2. *If $\dim_K V$ is finite, then the action of $\Gamma(K)$ on $\Gamma(V)$ has finitely many orbits. In fact,*
$$|\Gamma(V)/\Gamma(K)| \leq \dim_K V.$$

[1]Modulo the fact that if T is a strongly minimal theory, in which $\mathrm{acl}(\emptyset)$ is infinite and finite sets of tuples are coded by tuples, then T eliminates imaginaries. This is Lemma 1.6 in [58].

Proof. Let $v_1, \ldots, v_n$ be non-zero vectors with $\mathrm{val}(v_n)$ in different orbits of $\Gamma(K)$. We will show that the v_i are linearly independent. If not, let $w_1, \ldots, w_m$ be a minimal subset which is linearly dependent. Then $\sum_i x_i w_i = 0$ for some $x_i \in K^\times$. But by assumption,

$$\mathrm{val}(x_i w_i) = \mathrm{val}(x_i) + \mathrm{val}(w_i) \neq \mathrm{val}(x_j) + \mathrm{val}(w_j) = \mathrm{val}(x_j w_j)$$

for any $i \neq j$. By the ultrametric inequality in V, $\sum_i x_i w_i$ cannot be zero, a contradiction. □

For the rest of this section, we will assume that all valued K-vector spaces V have value group $\Gamma(V) = \Gamma(K) = \Gamma$, since the goal is Theorem 6.1.5.

If V and W are two such valued K vector spaces, we can form a "direct sum" $V \oplus W$ by setting

$$\mathrm{val}(v, w) = \min(\mathrm{val}(v), \mathrm{val}(w)).$$

For example, $K^{\oplus n}$ is a valued K-vector space with underlying vector space K^n, with value group $\Gamma(K)$, and with valuation map given by

$$\mathrm{val}(x_1, \ldots, x_n) = \min(\mathrm{val}(x_1), \ldots, \mathrm{val}(x_n)).$$

If V and W are two subspaces of a valued K-vector space, say that V and W are *perpendicular* if $V \cap W = \emptyset$ and $V + W$ is isomorphic to $V \oplus W$. In other words, V and W are perpendicular if $\mathrm{val}(v + w) = \min(\mathrm{val}(v), \mathrm{val}(w))$ for every $v \in V$ and $w \in W$.

Recall that a valued field K is *spherically complete* if every descending sequence of balls in K has non-empty intersection. If V is a valued K-vector space, a *ball* in V is a set of the form

$$\{\mathrm{val}(x - a) \geq \gamma\} \text{ or } \{\mathrm{val}(x - a) > \gamma\}$$

for $a \in V$ and $\gamma \in \Gamma(V)$. We say that V is *spherically complete* if every descending sequence of balls in V has a non-empty intersection.

Remark 6.1.3.

1. *If V and W are spherically complete, so is $V \oplus W$, because the balls in $V \oplus W$ are of the form $B_1 \times B_2$, with B_1 a ball in V and B_2 a ball in W.*

2. *If V is a subspace of a valued K-vector space W, and $a \in W$, then the intersection of any ball in W with $a + V$ is either empty or a ball.*

3. *If V is a spherically complete subspace of W and $a \in W$, and $\mathcal{F}$ is the family of closed balls in W centered at the origin which intersect $a + V$, then $\mathcal{F} \cap V$ is a nested chain of balls in $a + V$, so it has a non-empty intersection. Equivalently, the following set has a maximum:*

$$\{\mathrm{val}(a + v) : v \in V\}.$$

That is, some element of $a + V$ is maximally close to 0.

Lemma 6.1.4. *Let W be a valued K-vector space. Let V be a subspace. Suppose that $a \in W \setminus V$ is maximally close to 0 among elements of $a + V$. Then $K \cdot a$ is perpendicular to V.*

Proof. We need to show that

$$\mathrm{val}(v + \alpha a) = \min(\mathrm{val}(v), \mathrm{val}(\alpha a)) \tag{6.1}$$

for $v \in V$ and $\alpha \in K$. Replacing v and α with $\alpha^{-1}v$ and $\alpha^{-1}\alpha$ changes both sides of (6.1) by the same amount, so we may assume that $\alpha = 0$ or $\alpha = 1$.

The $\alpha = 0$ case is trivial. Suppose that $\alpha = 1$; we want to show $\mathrm{val}(v + a) = \min(\mathrm{val}(v), \mathrm{val}(a))$. If $\mathrm{val}(v) \neq \mathrm{val}(a)$, then $\mathrm{val}(v + a) = \min(\mathrm{val}(v), \mathrm{val}(a))$ by the ultrametric inequality. In the case where $\mathrm{val}(v) = \mathrm{val}(a)$, the ultrametric inequality only implies

$$\mathrm{val}(v + a) \geq \min(\mathrm{val}(v), \mathrm{val}(a)) = \mathrm{val}(a). \tag{6.2}$$

But $\mathrm{val}(v + a) \leq \mathrm{val}(a)$, by the assumption on a. So equality holds in (6.2). □

Theorem 6.1.5. *Suppose K is spherically complete, V is an n-dimensional K-vector space, and $\Gamma(K) = \Gamma(V)$. Then V is isomorphic to $K^{\oplus n}$. In other words, there is a basis $\{v_1, \ldots, v_n\} \subseteq V$ such that*

$$\mathrm{val}(x_1 v_1 + \cdots + x_n v_n) = \min_{1 \leq i \leq n} \mathrm{val}(x_i) \text{ for every } \vec{x} \in K^n.$$

In [32], Hrushovski calls $\{v_1, \ldots, v_n\}$ a "separating basis."

Proof. Proceed by induction on $\dim_K V$. The one-dimensional case is easy. Let V' be a codimension 1 subspace. By induction, V' is isomorphic to $K^{\oplus(n-1)}$, so V' is spherically complete. Choose some $a_0 \in V \setminus V'$ and let a be an element of $a_0 + V'$ maximally close to 0. By Lemma 6.1.4, $K \cdot a$ is perpendicular to V'. Thus $V \cong V' \oplus K \cong K^{\oplus(n-1)} \oplus K = K^{\oplus n}$. □

6.1.4 Definable submodules of K^n

We now return to the setting of ACVF.

Recall that every model of ACVF is elementarily equivalent to a spherical complete one.[2]

Theorem 6.1.6. *Let K be a model of ACVF. Let V be a definable K-vector space, with $\dim_K V < \infty$. Let $N \subseteq V$ be a* definable *$\mathcal{O}$-submodule. Then N is isomorphic to $K^{n_1} \times \mathcal{O}^{n_2} \times \mathfrak{M}^{n_3}$ for some $n_1, n_2, n_3 < n$.*

[2]This is well-known, and discussed in [49]. In the pure characteristic case, one can use fields of Hahn series. In the mixed characteristic case, one can use metric ultrapowers of $\mathbb{C}_p$.

Proof. We are trying to prove a conjunction of first-order sentences, so we may replace K with an elementarily equivalent model. Therefore, we may assume K is spherically complete.

Replacing V with the K-span of N, we may assume that V is the K-span of N. Similarly, if W denotes the largest K-vector space contained in N, then by quotienting out W, we may assume that N contains no nontrivial K-vector spaces. Now $\bigcup_{\alpha\in K^\times} \alpha N = V$ and $\bigcap_{\alpha\in K^\times} \alpha N = 0$.

For any nonzero $v \in V$, let

$$\mathrm{val}(v) = \sup\{\mathrm{val}(\alpha) : v \in \alpha N\} = \inf\{\mathrm{val}(\alpha) : v \notin \alpha N\}.$$

This is well-defined by o-minimality of Γ, and one easily checks that

$$\mathrm{val}(\beta v) = \mathrm{val}(\beta) + \mathrm{val}(v). \tag{6.3}$$

$$\mathrm{val}(v) > 0 \implies v \in N \tag{6.4}$$

$$\mathrm{val}(v) < 0 \implies v \notin N \tag{6.5}$$

for all $\beta \in K$, $v \in V$. We claim that $\mathrm{val} : V \to \Gamma$ makes V into a valued K-vector space. Given (6.3), we merely need to check the ultrametric inequality

$$\mathrm{val}(v + w) \geq \min(\mathrm{val}(v), \mathrm{val}(w)).$$

If this failed, then multiplying everything by an appropriate scalar, we would get

$$\mathrm{val}(v + w) < 0 < \min(\mathrm{val}(v), \mathrm{val}(w)).$$

But then $v, w \in N$ and $v + w \notin N$, contradicting the fact that N is a module.

So $\mathrm{val} : V \to \Gamma$ makes V into a valued K-vector space. By Theorem 6.1.5, we can assume that V is $K^{\oplus n}$. Then (6.4-6.5) mean the following for $\vec{x} \in K^n$:

- If $\mathrm{val}(x_i) > 0$ for every i, then $\vec{x} \in N$. In other words, $\mathfrak{M}^n \subseteq N$.
- If $\mathrm{val}(x_i) < 0$ for some i, then $\vec{x} \notin N$. In other words, $N \subseteq \mathcal{O}^n$.

So N is sandwiched between $\mathcal{O}^n$ and $\mathfrak{M}^n$. But the possibilities for N then correspond to the submodules of $\mathcal{O}^n/\mathfrak{M}^n$, i.e., the k-subspaces of k^n. These are easy to deal with.

Specifically, note that $N/\mathfrak{M}^n$ is a k-subspace of $\mathcal{O}^n/\mathfrak{M}^n = k^n$. Let γ be an element of $GL_n(k)$ sending $N/\mathfrak{M}^n \subseteq k^n$ to $k^\ell \times 0^{n-\ell} \subseteq k^n$ for $\ell = \dim_k N/\mathfrak{M}^n$. Then γ can be lifted to some $\gamma' \in GL_n(\mathcal{O})$, because $\mathcal{O}$ is a local ring. If $N' = \gamma'(N)$, then $N'/\mathfrak{M}^n$ is $k^\ell \times 0^{n-\ell} = (\mathcal{O}^\ell \times \mathfrak{M}^{n-\ell})/\mathfrak{M}^n$. So $N' = \mathcal{O}^\ell \times \mathfrak{M}^{n-\ell} \subseteq K^n$. But N' and N are isomorphic. □

Let Mod_n denote the set of definable submodules of K^n. The theorem implies that the elements of Mod_n fall into finitely many definable families. Consequently, we get the following

Corollary 6.1.7. *The set Mod_n is interpretable.*

6.2 Generalities on Definable Types

Work in an arbitrary theory T, with monster model $\mathbb{M}$. By "C-definable type," we will mean C-definable type over the monster, as opposed to some smaller model, unless stated otherwise. By "definable type," we mean a C-definable type for some $C \subseteq \mathbb{M}$.

In this section we review some well-known facts about definable types. We omit many of the proofs, which are usually straightforward.

6.2.1 Operations on definable types

If p is a C-definable type and f is a C-definable function, there is a unique C-definable type f_*p which is characterized by the following property:

$$a \models p|B \implies f(a) \models f_*p|B, \text{ for all small } B \supseteq C \text{ and all } a.$$

The choice of C does not matter—if p and f are C'-definable for some other set C', then the resulting f_*p is the same. The type f_*p is called the pushforward of p along f.

If p and q are two C-definable types, there is a unique C-definable type $p \otimes q$ which is characterized by the following property:

$$(a,b) \models p \otimes q|B \iff (a \models p|Bb) \wedge (b \models q|B), \text{ for all small } B \supseteq C \text{ and all } a, b.$$

Again, $p(x) \otimes q(y)$ does not depend on the choice of C. The product operation is associative:

$$(p(x) \otimes q(y)) \otimes r(z) = p(x) \otimes (q(y) \otimes r(z)),$$

but commutativity

$$p(x) \otimes q(y) \stackrel{?}{=} q(y) \otimes p(x)$$

can fail.

Remark 6.2.1. *If f, g are definable functions and p, q are definable types, then $f_*p \otimes g_*q = (f \times g)_*(p \otimes q)$, where $f \times g$ sends (x, y) to $(f(x), g(y))$.*

6.2.2 Generically stable types

Now assume that T is NIP. (This includes the case of ACVF.)

Definition 6.2.2. *A definable type $p(x)$ is* generically stable *if $p(x) \otimes q(y) = q(y) \otimes p(x)$ for every definable type $q(y)$.*

For other equivalent definitions of generic stability, see Section 3 of [36].

Definition 6.2.3. *Let f be a C-definable function and p be a C-definable type. Abusing terminology significantly, say that p is* dominated along f *if*

$$f(a) \models f_*p|B \implies a \models p|B \text{ for all small } B \supseteq C \text{ and all } a.$$

Note that the converse implication holds by definition of f_*p. Unlike the previous definitions, this *does* depend on the choice of C. In the cases we care about, C will be $\emptyset$.

Remark 6.2.4. *Suppose p is dominated along f, and q is some other definable type. If B is a set over which everything is defined and over which the domination works, then*

$$(f(a), b) \models f_*p \otimes q|B \implies (a, b) \models p \otimes q|B \tag{6.6}$$

We will use the following basic facts about generically stable types:

Theorem 6.2.5.

(a) *Products of generically stable types are generically stable.*

(b) *Pushforwards of generically stable types are generically stable.*

(c) *If p is dominated along f and f_*p is generically stable, then p is generically stable.*

(d) *If p and q are generically stable, dominated along f and g, respectively, then $p \otimes q$ is dominated along $f \times g$.*

(e) *To check generic stability, it suffices to show that p commutes with itself, i.e., $p(x_1) \otimes p(x_2) = p(x_2) \otimes p(x_1)$.*

Proof.

(a) If p and q are generically stable, and r is arbitrary, then $p \otimes q \otimes r = p \otimes r \otimes q = r \otimes p \otimes q$.

(b) Suppose p is generically stable, f is a definable function, and q is arbitrary. Then $p(x) \otimes q(y) = q(y) \otimes p(x)$. Pushing both sides forwards along $(f \times id)$ and applying Remark 6.2.1, we get that $f_*p(x') \otimes q(y) = q(y) \otimes f_*p(x')$.

(c) Let q be another invariant type; we will show that $p(x) \otimes q(y) = q(y) \otimes p(x)$. Let B be a set over which p, q, f are defined. Let (b, a) realize $q \otimes p|B$. By Remark 6.2.1, $(b, f(a)) \models q \otimes f_*p|B$. Since f_*p is generically stable, $(f(a), b) \models f_*p \otimes q|B$. By (6.6), $(a, b) \models p \otimes q|B$. So $p \otimes q$ and $q \otimes p$ agree when restricted to the arbitrary set B.

(d) Let B be a sufficiently big set. Suppose that $(f(a), g(b)) \models f_*p \otimes g_*q|B$. We need to show that $(a, b) \models p \otimes q|B$. By (6.6), $(a, g(b)) \models p \otimes g_*q|B$. By generic stability of p, $(g(b), a) \models g_*q \otimes p|B$. By (6.6) again, $(b, a) \models q \otimes p|B$. By generic stability again, $(a, b) \models p \otimes q|B$.

(e) Suppose $p(x)$ commutes with itself, but $p(x) \otimes q(y) \neq q(y) \otimes p(x)$. Choose some formula $\phi(x; y; c)$ which is in $p(x) \otimes q(y)$ and not in $q(y) \otimes p(x)$. We will prove that $\phi(x; y, z)$ has the independence property. Let n be arbitrary. Let $a_1, \ldots, a_n, b, a_{n+1}, \ldots, a_{2n}$ realize $p^{\otimes n} \otimes q \otimes p^{\otimes n}$ restricted to c. Then $\models \phi(a_i; b; c) \iff i \leq n$, by choice of $\phi(x; y; c)$. The fact that p commutes with itself implies that all permutations of $(a_1, \ldots, a_{2n})$ have

the same type over c. Therefore, for each permutation π of $\{1, \ldots, 2n\}$, we can find a b_π such that $\phi(a_i; b_\pi; c)$ holds iff $\pi(i) \leq n$. It follows that for any $S \subseteq \{1, \ldots, n\}$, we can find a b_S such that $\phi(a_i; b_S; c)$ holds if and only if $i \in S$. As n was arbitrary, T has the independence property, a contradiction.

□

6.2.3 Definable types in ACVF

Now work in ACVF. Recall that ACVF is NIP. We will make use of several definable types:

- If B is an open or closed ball in the home sort, then there is a complete type $p_B(x)$ over $\mathbb{M}$ which says that $x \in B$ and x is not in any strictly smaller balls. This type is called the *generic type of* B. Completeness follows from C-minimality. This type is definable, essentially because if B' is any other ball, then the formula $x \in B'$ is in $p_B(x)$ if and only if $B' \supseteq B$. If C is any set of parameters over which B is defined, then $p_B|C$ says precisely that x is in B, and x is not in any $\mathrm{acl}^{eq}(C)$-definable proper subball of B.

- There is also a type $p_k(x)$ which says that x is in the residue field, and is not algebraic over $\mathbb{M}$. This is called the *generic type of the residue field*, and is definable because k is strongly minimal. If C is any set of parameters, $p_k|C$ says precisely that $x \in k$ and $x \notin \mathrm{acl}^{eq}(C)$.

- The valuation ring $\mathcal{O}$ is a closed ball, so it has a generic type $p_\mathcal{O}$. Over any set of parameters C, $p_\mathcal{O}(x)$ says that $x \in \mathcal{O}$, and that x is not in any $\mathrm{acl}^{eq}(C)$-definable proper subballs of $\mathcal{O}$. Every proper subball of $\mathcal{O}$ is contained in a unique one of the form $\mathrm{res}^{-1}(\alpha)$, for $\alpha \in k$. Consequently, $p_\mathcal{O}|C$ equivalently says that $x \in \mathcal{O}$ and that $x \notin \mathrm{res}^{-1}(\alpha)$ for any $\alpha \in \mathrm{acl}^{eq}(C)$. Equivalently,

$$x \models p_\mathcal{O}|C \iff \mathrm{res}(x) \models p_k|C$$

  Therefore, $p_\mathcal{O}$ is dominated along res, and $\mathrm{res}_* p_\mathcal{O} = p_k$.

The type p_k is generically stable. To see this, use (e) of Theorem 6.2.5 and stability of k.

Since p_k is generically stable, so is $p_\mathcal{O}$, by Theorem 6.2.5(c). If B is any closed ball, then there is an affine transformation $f(x) = ax + b$ sending $\mathcal{O}$ to B, and $p_B = f_* p_\mathcal{O}$. By Theorem 6.2.5(b), each p_B is generically stable.

Let $p_{\mathcal{O}^n}$ be $p_\mathcal{O}^{\otimes n}$. We think of $p_{\mathcal{O}^n}$ as the generic type of the lattice $\mathcal{O}^n$. By Theorem 6.2.5, $p_{\mathcal{O}^n}$ is generically stable, and is dominated along the map $(x_1, \ldots, x_n) \mapsto (\mathrm{res}(x_1), \ldots, \mathrm{res}(x_n))$. Also, the pushforward along this map is $p_k^{\otimes n}$, the generic type of k^n.

The generic type of k^n is stabilized by the action of $GL_n(k)$, so by domination, the generic type of $\mathcal{O}^n$ is stabilized by $GL_n(\mathcal{O})$. In light of this, the following definition does not depend on the choice of g:

Definition 6.2.6. *Let Λ be a lattice in K^n. The* generic type p_Λ of Λ *is $g_* p_{\mathcal{O}^n}$, where $g : K^n \to K^n$ is a linear map sending $\mathcal{O}^n$ to Λ.*

Moreover, p_Λ is $\ulcorner\Lambda\urcorner$-definable. Note that p_Λ is a generically stable type, because it is a pushforward of a generically stable type.

6.2.4 Left transitivity

Now return to an arbitrary theory T. Work in T^{eq}, so that acl and dcl mean acl^{eq} and dcl^{eq}.

Lemma 6.2.7. *Suppose $C \subseteq B$ are small sets and a_1, a_2 are tuples (possibly infinite, but small). If $\mathrm{tp}(a_1/B)$ is C-definable and $\mathrm{tp}(a_2/Ba_1)$ is Ca_1-definable, then $\mathrm{tp}(a_2a_1/B)$ is C-definable.*

Proof. Naming the parameters from C, we may assume $C = \emptyset$. Let $\phi(x_2, x_1; y)$ be a formula; we must produce a ϕ-definition (over $\emptyset$) for $\mathrm{tp}(a_2a_1/B)$. Since $\mathrm{tp}(a_2/Ba_1)$ is a_1-definable, the $\phi(x_2; x_1, y)$-type of a_2 over Ba_1 has a definition $\psi(x_1, y, a_1)$. In particular, for every tuple b from B,

$$\models \phi(a_2, a_1, b) \leftrightarrow \psi(a_1, b, a_1).$$

Meanwhile, since $\mathrm{tp}(a_1/B)$ is 0-definable, there is a formula $\chi(y)$ such that for every b in B,

$$\models \psi(a_1, b, a_1) \leftrightarrow \chi(b).$$

Thus, for every b in B,

$$\models \phi(a_2, a_1, b) \leftrightarrow \chi(b),$$

so $\chi(y)$ is the $\phi(x_2, x_1; y)$-definition of $\mathrm{tp}(a_2a_1/B)$. □

Lemma-Definition 6.2.8. *The following are equivalent for A, B, C small sets of parameters.*

1. *The $*$-type $\mathrm{tp}(A/BC)$ has a global C-definable extension.*
2. *For every small set of parameters D, there is a C-definable extension of $\mathrm{tp}(A/BC)$ to a $*$-type over BCD.*
3. *For every small set of parameters D, there is $D' \equiv_{BC} D$ such that $\mathrm{tp}(A/BCD')$ is C-definable.*
4. *For some small model $M \supseteq BC$, $\mathrm{tp}(A/M)$ is C-definable.*

We denote these equivalent conditions by $A \overset{def}{\underset{C}{\downarrow}} B$.

Proof. **(1 $\implies$ 2)** The restriction of a global C-definable type to BCD is C-definable.

(2 $\implies$ 3) Given D, (2) implies that there is $A' \equiv_{BC} A$ such that $\mathrm{tp}(A'/BCD)$ is C-definable. Choose D' such that $A'D \equiv_{BC} AD'$. Then $\mathrm{tp}(A/BCD')$ is C-definable and $D' \equiv_{BC} D$.

(3 $\Longrightarrow$ 4) Applying (3) to D a small model containing BC, we get a small model D' containing BC such that $\mathrm{tp}(A/BCD') = \mathrm{tp}(A/D')$ is C-definable.

(4 $\Longrightarrow$ 1) C-definable types over models have unique C-definable extensions to elementary extensions. This is true even for $*$-types.

□

Lemma 6.2.9. *If $a_1 \bigcup_C^{def} B$ and $a_2 \bigcup_{Ca_1}^{def} B$ then $a_2a_1 \bigcup_C^{def} B$, so $\bigcup^{def}$ satisfies left-transitivity.*

Proof. We use condition (3) of Lemma-Definition 6.2.8. Let D be a small set of parameters. Since $a_1 \bigcup_C^{\mathrm{def}} B$, there is $D' \equiv_{BC} D$ such that $\mathrm{tp}(a_1/BCD')$ is C-definable. As $a_2 \bigcup_{Ca_1}^{\mathrm{def}} B$ there is $D'' \equiv_{Ca_1B} D'$ such that $\mathrm{tp}(a_2/BCa_1D'')$ is Ca_1-definable. Note that $\mathrm{tp}(a_1/BCD'')$ is C-definable. By Lemma 6.2.7, it follows that $\mathrm{tp}(a_2a_1/BCD'')$ is C-definable. Since $D'' \equiv_{BC} D' \equiv_{BC} D$, we have verified $a_2a_1 \bigcup_C^{\mathrm{def}} B$ using condition (3). □

Definition 6.2.10. *If A, B, C are small sets of parameters, we will write $A \bigcup_C^{adef} B$ to mean $A \bigcup_{\mathrm{acl}(C)}^{def} B$. (Recall that $\mathrm{acl}(C)$ means $\mathrm{acl}^{eq}(C)$.)*

In other words, $a \bigcup_C^{\mathrm{adef}} B$ if $\mathrm{tp}(a/BC)$ can be extended to a type which is almost C-definable, that is, $\mathrm{acl}(C)$-definable. In a stable theory, $\bigcup^{\mathrm{adef}}$ is exactly nonforking independence.

Lemma 6.2.11. *If $A \bigcup_C^{adef} B$, then $\mathrm{acl}(AC) \bigcup_C^{adef} B$.*

Proof. Replacing C with $\mathrm{acl}(C)$, we may assume that $C = \mathrm{acl}(C)$. By condition (4) of Lemma-Definition 6.2.8, there is a small model M containing BC such that $\mathrm{tp}(A/M)$ is C-definable. We need to show that $\mathrm{tp}(\mathrm{acl}(AC)/M)$ is C-definable. This is equivalent to showing that for each $\mathrm{acl}(AC)$-definable set X, there is some C-definable set X' such that $X \cap M = X' \cap M$.

Given such an X, let $X_1, \ldots, X_n$ be the conjugates of X over AC. Let E be the equivalence relation

$$xEy \leftrightarrow \bigwedge_{i=1}^{n} (x \in X_i \leftrightarrow y \in X_i)$$

Then E is AC-definable. Since $\mathrm{tp}(AC/M)$ is C-definable, the restriction $E' = E \cap M$ of E to M is C-definable. Since E has finitely many equivalence classes, so does E', and hence each equivalence class of E' is C-definable, as $\mathrm{acl}(C) = C$. But $X \cap M$ is a union of finitely many E'-equivalence classes, so $X \cap M$ is C-definable. □

Lemma 6.2.12. *If $a_1 \bigcup_C^{adef} B$ and $a_2 \bigcup_{Ca_1}^{adef} B$, then $a_2a_1 \bigcup_C^{adef} B$, so $\bigcup^{adef}$ satisfies left-transitivity.*

Proof. By Lemma 6.2.11, we know that $\mathrm{acl}(a_1C) \bigcup_{\mathrm{acl}(C)}^{\mathrm{def}} B$. We are given $a_2 \bigcup_{\mathrm{acl}(Ca_1)}^{\mathrm{def}} B$. Combining these using Lemma 6.2.9, we conclude that $a_2 \mathrm{acl}(Ca_1) \bigcup_{\mathrm{acl}(C)}^{\mathrm{def}} B$. This easily implies $a_2a_1 \bigcup_{\mathrm{acl}(C)}^{\mathrm{def}} B$, as desired. □

6.3 An Abstract Criterion for Elimination of Imaginaries

We state a sufficient condition for a theory T to have elimination of imaginaries, extracted from [32].

Theorem 6.3.1. *Let T be a theory, with home sort K (meaning $\mathbb{M}^{eq} = \mathrm{dcl}^{eq}(K)$). Let $\mathcal{G}$ be some collection of sorts. If the following conditions all hold, then T has elimination of imaginaries in the sorts $\mathcal{G}$.*

- *For every non-empty definable set $X \subseteq K^1$, there is an $\mathrm{acl}^{eq}(\ulcorner X \urcorner)$-definable type in X.*
- *Every definable type in K^n has a code in $\mathcal{G}$ (possibly infinite). That is, if p is any (global) definable type in K^n, then the set $\ulcorner p \urcorner$ of codes of the definitions of p is interdefinable with some (possibly infinite) tuple from $\mathcal{G}$.*
- *Every finite set of finite tuples from $\mathcal{G}$ has a code in $\mathcal{G}$. That is, if S is a finite set of finite tuples from $\mathcal{G}$, then $\ulcorner S \urcorner$ is interdefinable with a tuple from $\mathcal{G}$.*

Proof. Assume the three conditions.

Claim 6.3.2. *For every non-empty definable set $X \subseteq K^n$, there is an $\mathrm{acl}^{eq}(\ulcorner X \urcorner)$-definable type in X.*

Proof. We proceed by induction on n, the base case $n = 1$ being given. Suppose $n > 1$. Take $X \subseteq K^n$. Let $C = \ulcorner X \urcorner$. Let $\pi : K^n \twoheadrightarrow K^{n-1}$ be the projection onto the first $n-1$ coordinates. Then $\pi(X)$ is C-definable, so by induction, there is an $\mathrm{acl}^{eq}(C)$-definable type in $\pi(X)$. Let a_1 realize this type. Then $a_1 \mathop{\downarrow}^{\mathrm{adef}}_{C} C$.

Let $Y = \{y \in K^1 | (a_1, y) \in X\}$, so Y is essentially $X \cap \pi^{-1}(a_1)$. Then Y is Ca_1-definable and non-empty. By assumption, there is an $\mathrm{acl}^{eq}(Ca_1)$-definable type in Y. Let a_2 realize this type; then $a_2 \in Y$ and $a_2 \mathop{\downarrow}^{\mathrm{adef}}_{Ca_1} C$. Since $a_1 \mathop{\downarrow}^{\mathrm{adef}}_{C} C$, it follows that $a_1 a_2 \mathop{\downarrow}^{\mathrm{adef}}_{C} C$ by Lemma 6.2.12. By definition of $\mathop{\downarrow}^{\mathrm{adef}}$, there is an $\mathrm{acl}^{eq}(C)$-definable type $p(x_1, x_2)$ such that $a_1 a_2 \models p|\,\mathrm{acl}^{eq}(C)$. As $a_2 \in Y$, the tuple $a_1 a_2$ is in X, so p is an $\mathrm{acl}^{eq}(C)$-definable type in X. □

Let e be any imaginary. Then there is some n and some 0-definable equivalence relation E on K^n such that e is a code for some E-equivalence class X. By the claim, there is an $\mathrm{acl}^{eq}(e)$-definable type p in X. Then $e \in \mathrm{dcl}^{eq}(\ulcorner p \urcorner)$, because X is the unique E-equivalence class in which the type $\ulcorner p \urcorner$ lives. By the second assumption, there is some small tuple $t_0 \subseteq \mathcal{G}$ such that $\ulcorner p \urcorner$ is interdefinable with t_0. Thus $e \in \mathrm{dcl}^{eq}(t_0)$ and $t_0 \in \mathrm{acl}^{eq}(e)$. By compactness, we can find some finite tuple t from $\mathcal{G}$ such that $e \in \mathrm{dcl}^{eq}(t)$ and $t \in \mathrm{acl}^{eq}(e)$. Write e as $f(t)$ for some 0-definable function f. Let S be the (finite) set of conjugates of t over e. Then S is e-definable. Moreover, $f(t') = e$ for any $t' \in S$, so e is $\ulcorner S \urcorner$-definable. Hence e and $\ulcorner S \urcorner$

are interdefinable. By the third hypothesis, $\ulcorner S \urcorner$ has a code in $\mathcal{G}$. So e has a code in $\mathcal{G}$. As e was arbitrary, T has elimination of imaginaries down to the sorts in $\mathcal{G}$. □

The conditions in the theorem are sufficient but not necessary for elimination of imaginaries to hold. Namely, the first condition has nothing to do with $\mathcal{G}$, and happens to fail in $\mathbb{Q}_p$, even if $\mathcal{G}$ is chosen to be all of $\mathbb{Q}_p^{eq}$.

6.3.1 Examples

We sketch how to use Theorem 6.3.1 to verify elimination of imaginaries in ACF and DCF_0 in the home sort K (so $\mathcal{G}$ is merely $\{K\}$). The first condition follows from stability: if X is any non-empty definable set, then the formula $x \in X$ does not fork over $\ulcorner X \urcorner$. If p is a global type which does not fork over $\ulcorner X \urcorner$ and contains this formula, then p is an $\mathrm{acl}^{eq}(\ulcorner X \urcorner)$-definable type in X.

For the second condition, one must check that every type has a code (possibly infinite) in the home sort. If p is a type in K^n, then there is a minimal Zariski-closed or Kolchin-closed set V containing p, and p and V have the same code. The second condition thus reduces to coding Zariski-closed sets or Kolchin-closed sets, respectively. So does the third condition, since any finite subset of K^n is Zariski-closed and Kolchin-closed. Now, to code a Zariski-closed or Kolchin-closed set V, we merely need to code the ideal I of polynomials or differential polynomials which vanish on V. In the ACF case, this reduces to coding, for each $d < \omega$, the intersection of I with the space of degree $\leq d$ polynomials in $K[X_1, \ldots, X_n]$. Something similar happens in DCF. So the problem reduces to coding linear subspaces of K^m for various m.

But this is doable, by the following basic and general fact:

Lemma 6.3.3. *Let K be any field. Let V be a subspace of K^n. Then V can be coded by a tuple in K, and V and K^n/V have $\ulcorner V \urcorner$-definable bases.*

Proof. Let $m = \dim V$. By linear algebra, there is some coordinate projection $\pi : K^n \to K^m$ such that the restriction of π to V is an isomorphism $V \to K^m$. Then the preimage of the standard basis under this isomorphism is an $\ulcorner V \urcorner$-definable basis for V. This basis is a code for V. Meanwhile, if we push the standard basis of K^n forward to K^n/V, then some subset of this will be a basis for K^n/V, and will be definable over the parameters (such as $\ulcorner V \urcorner$) that were used to define the set K^n/V. □

For the case of ACVF, the coding of definable types will be done similarly. But in addition to coding subspaces of K^n, we will also need to code definable ways of turning K^n into a valued K-vector space. The third condition of Theorem 6.3.1 will be verified *using* the coding of definable types.

6.4 Elimination of imaginaries in ACVF

In this section, we prove that ACVF has elimination of imaginaries in the sorts K, $R_{n,\ell}$, by applying Theorem 6.3.1. Recall that we are referring to these as the geometric sorts. We say that an object has a *geometric code* if it has a code in these sorts.

6.4.1 Coding modules

Recall that $R_{n,\ell}$ is the set of pairs (Λ, V) where Λ is a lattice in K^n and V is an ℓ-dimensional subspace of $\operatorname{res}\Lambda := \Lambda \otimes_{\mathcal{O}} k = \Lambda/\mathfrak{M}\Lambda$.

For fixed Λ, the poset of k-subspaces of $\operatorname{res}\Lambda$ is isomorphic to the poset of $\mathcal{O}$-submodules between $\mathfrak{M}\Lambda$ and Λ. Moreover, ℓ-dimensional subspaces correspond exactly to $\mathcal{O}$-submodules isomorphic to $\mathcal{O}^\ell \times \mathfrak{M}^{n-\ell}$. So we could equivalently define $R_{n,\ell}$ to be the set of all $\mathcal{O}$-submodules of K^n isomorphic to $\mathcal{O}^\ell \times \mathfrak{M}^{n-\ell}$. Under this identification,

$$\bigcup_{\ell=0}^{n} R_{n,\ell}$$

is the space of all open bounded definable $\mathcal{O}$-submodules of K^n, by Theorem 6.1.6.

In section 6.1.4, we saw that Mod_n, the set of *all* definable submodules of K^n, is interpretable. We now show that Mod_n can be embedded into the geometric sorts.

Lemma 6.4.1. *If $M \leq K^n$ is a definable submodule of K^n, then M has a geometric code.*

Proof. Let V^+ be the K-span of N, and let V^- be the maximal K-subspace of K^n contained in N. By Lemma 6.3.3, the subspaces V^+ and V^- can be coded by a tuple c from K, and the quotient V^+/V^- has a c-definable identification with K^m, for some m. Then N is interdefinable over c with the image of N/V^- in K^m. But this image will be an open bounded definable submodule, so it is an element of $R_{m,\ell}$ for some ℓ. □

6.4.2 Coding definable types

A definable type in K^n induces an ideal I in $K[X_1, \ldots, X_n]$ together with the structure of a valued K-vector space on the quotient $K[\vec{X}]/I$. By quantifier elimination in the one-sorted language, these data completely determine the type. So the problem of finding codes for definable types reduces to the (easy) problem of coding subspaces, and the problem of coding valued vector space structures on K-vector spaces.[3]

At the risk of being abstruse...

Definition 6.4.2. *Let V be a K-vector space. A* vvs structure *on V is a binary relation R on V such that there is a valued K-vector space structure $(V, \Gamma(V), \ldots)$ on V for which $xRy \iff \operatorname{val}(x) \leq \operatorname{val}(y)$.*

[3] We will be more explicit in the proof of Theorem 6.4.5 below.

The vvs structures on V are essentially the distinct ways of turning V into a valued K-vector space. Two valued K-vector spaces W and W' with the same underlying vector space V yield the same vvs structure on V iff they are isomorphic over V.

If V is a definable K-vector space, it makes sense to say that a vvs structure R is definable, meaning that R is a definable subset of $V \times V$. If R is definable, then $\Gamma(V)$ is interpretable and the map $\mathrm{val} : V \to \Gamma(V)$ and the action of $\Gamma(K)$ on $\Gamma(V)$ are all definable.

Theorem 6.4.3. *Let τ be (the code for) a definable vvs structure on K^m. Then τ is interdefinable with an element of the geometric sorts.*

Proof. Let V be the associated valued K-vector space. So K^m is the underlying vector space of V and τ is a code for the relation $\mathrm{val}(x) \leq \mathrm{val}(y)$. The set $\Gamma(V)$ is τ-interpretable, as $K^m \setminus 0$ modulo the equivalence relation $\mathrm{val}(x) \leq \mathrm{val}(y) \wedge \mathrm{val}(y) \leq \mathrm{val}(x)$.

By Remark 6.1.2, $\Gamma(V)$ consists of finitely many orbits under $\Gamma(K)$.

If there was only one orbit, and if there was a canonical identification of $\Gamma(V)$ with $\Gamma(K)$, we could proceed as follows: let B be the closed ball around 0 with valuative radius 0. The other closed balls around 0 are all of the form αB, for $\alpha \in K$. The set B is a definable $\mathcal{O}$-submodule of K^m, so it has a geometric code. It determines τ, however, because $\mathrm{val}(x) \leq \mathrm{val}(y)$ if and only if every closed ball containing 0 and x contains y. So τ and $\ulcorner B \urcorner$ would be interdefinable.

In general we have several orbits. The first order of business is finding a τ-definable element in each one:

Claim 6.4.4. *Each orbit of $\Gamma(K)$ on $\Gamma(V)$ contains a τ-definable element.*

Proof. For $x, y \in \Gamma(V)$, let $x \gg y$ indicate that $\Gamma + x > \Gamma + y$. Let $x \sim y$ indicate that $x \not\gg y$ and $y \not\gg x$. This is an equivalence relation. Each orbit of $\Gamma(K)$ is in one $\sim$-equivalence class, so there are finitely many $\sim$-equivalence classes. Each $\sim$-equivalence class is convex. Let $C_1 > C_2 > \ldots > C_\ell$ be the distinct $\sim$-equivalence classes sorted in order from most positive to most negative.

For $0 \leq i \leq \ell$, let V_i be the set of $v \in V$ such that $\mathrm{val}(v) \in C_j$ for some $j \leq i$. Each V_i is a K-vector space, yielding an ascending filtration

$$0 = V_0 \subseteq V_1 \subseteq \cdots \subseteq V_\ell = V.$$

On $V_i \setminus V_{i-1}$, the function val lands in C_i and factors through the quotient V_i/V_{i-1}, by the ultrametric inequality.

The equivalence relation $\sim$ is τ-definable. Since there are finitely many equivalence classes and they are totally ordered, each C_i is τ-definable. Consequently, each V_i is τ-definable. By Lemma 6.3.3, each quotient space V_i/V_{i-1} has a τ-definable basis. In particular, there is a τ-definable non-zero vector in V_i/V_{i-1}. Taking its valuation, we get a τ-definable element of C_i. We have shown:

$$\text{Each } C_i \text{ contains a } \tau\text{-definable element.} \tag{6.7}$$

Next, for $x, y \in \Gamma(V)$ let $x \approx y$ indicate that $x + \gamma > y > x - \gamma$ for all positive $\gamma \in \Gamma(K)$. This is again a τ-definable equivalence relation. If x and y are in the same orbit, but are not equal, then $x \not\approx y$. Consequently, each $\approx$-equivalence class contains at most one point from each orbit, so each $\approx$-equivalence class is finite. This implies that if $x \approx y$, then x and y are interalgebraic over τ. In light of the total ordering, they are actually inter*definable* over τ.

Let x be an arbitrary element of $\Gamma(V)$. We will show that the orbit $\Gamma(K) + x$ contains a τ-definable element. By (6.7), $x \sim y$ for some τ-definable element y. The set

$$\{\gamma \in \Gamma : \gamma + x \leq y\}$$

is non-empty, because $x \not\gg y$, and it is bounded above, because $y \not\gg x$. It is also definable, so it has a supremum γ_0, by o-minimality of $\Gamma(K)$. Then $\gamma_0 + x \approx y$. The element $\gamma_0 + x$ is interdefinable over τ with the τ-definable element y, so it is itself τ-definable. □

Given the claim, let $\gamma_1, \ldots, \gamma_n$ be a set of τ-definable orbit representatives. Let $B_i = \{v \in K^m : \mathrm{val}(v) \geq \gamma_i\}$. Each B_i is a τ-definable $\mathcal{O}$-submodule of K^m, i.e., an element of Mod_m. The closed balls of V containing 0 are exactly the sets of the form αB_i for $1 \leq i \leq n$ and $\alpha \in K$. The family of closed balls containing 0 is enough to determine the vvs structure, so τ is interdefinable with the tuple $(B_1, \ldots, B_n)$. But by Lemma 6.4.1, each B_i has a geometric code. □

Theorem 6.4.5. *Let $p(x)$ be a definable type in K^n. Then $p(x)$ has a code in the geometric sorts.*

Proof. For each d, let V_d be the space of polynomials in $K[X_1, \ldots, X_n]$ of degree $\leq d$. This is a finite dimensional definable K-vector space with a 0-definable basis. Let I_d be the set of $Q(X) \in V_d$ such that the formula $Q(x) = 0$ is in $p(x)$. Let R_d be the set of pairs $(Q_1(X), Q_2(X))$ in $V_d \times V_d$ such that the formula $\mathrm{val}(Q_1(x)) \leq \mathrm{val}(Q_2(x))$ is in $p(x)$. Then I_d is a subspace of V_d, and R_d induces a definable vvs structure on the quotient space V_d/I_d. Quantifier elimination in the one-sorted language implies that p is completely determined by the collection of all I_d's and R_d's for $d < \omega$. By Lemma 6.3.3, we can find codes $\ulcorner I_d \urcorner$ in the home sort for the I_d's. After naming these codes, each quotient space V_d/I_d has a definable basis, and can be definably identified with some power of K. Then each R_d is interdefinable with a definable vvs structure on a power of K. By Theorem 6.4.3, these vvs structures have codes c_d in the geometric sorts. Now the union of all the $\ulcorner I_d \urcorner$'s and c_d's is a code for p. □

6.4.3 Coding finite sets

In this section, we show that finite sets of tuples from the geometric sorts can be coded in the geometric sorts. We make use of the existence of geometric codes for definable types. For a more elementary but more complicated approach, see Proposition 3.4.1 in [26].

Definition 6.4.6. *If X is some set, $\mathrm{Sym}^n X$ will denote the n-fold symmetric product of X, that is, X^n modulo the action of the nth symmetric group. The natural map $X^n \to \mathrm{Sym}^n X$*

will be denoted σ, so that

$$\sigma(x_1, \ldots, x_n) = \sigma(y_1, \ldots, y_n)$$

if and only if there is a permutation π of n such that $x_i = y_{\pi(i)}$ for $i = 1, \ldots, n$.

Definition 6.4.7. *A 0-definable map $\pi : X \to Y$ has* definable lifting *if for every $b \in Y$ there is a b-definable type p_b in $\pi^{-1}(b)$. (In particular, π must be surjective.) Say that π has* generically stable lifting *if it has definable lifting and p_b can be taken to be generically stable.*

In both cases, we can easily modify the map $b \mapsto p_b$ to be automorphism equivariant, so that if $\sigma \in \mathrm{Aut}(\mathbb{M}/\emptyset)$, then $p_{\sigma(b)} = \sigma(p_b)$ for every b. Conversely, if there is an automorphism equivariant map $b \mapsto p_b$ from elements of Y to definable (resp. generically stable) types in X, such that p_b is in $\pi^{-1}(b)$, then π has definable (resp. generically stable) lifting—the automorphism invariance ensures that p_b is b-definable.

If $\pi : X \to Y$ has definable lifting, and q is a C-definable type in Y for some parameters C, then $q = \pi_* p$ for some C-definable type p in X. Indeed, if M is a model containing C, and b realizes $q|M$, and $a \in \pi^{-1}(b)$ realizes $p_b|Mb$, then $a \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{def}}_{Cb} M$ and $b \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{def}}_{C} M$, so $a \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{def}}_{C} M$ by left-transitivity of $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{def}}$. Thus $a \models p|M$ for some C-definable type p. By definition of pushforward, $\pi(a) = b$ realizes $\pi_* p|M$. Then the C-definable types $\pi_* p$ and q have the same restriction to a model containing C, so they are equal.

From this, we draw the following conclusion:

Observation 6.4.8. *If $\pi : X \to Y$ has definable lifting, and definable types in X have codes in the geometric sorts, then so do definable types in Y.*

Indeed, let q be a definable type in Y. Then q is $\ulcorner q \urcorner$-definable, so $q = \pi_* p$ for some $\ulcorner q \urcorner$-definable type p. But then p and q are interdefinable, and p has a geometric code.

For example, the residue map $\mathrm{res} : \mathcal{O} \to k$ has definable lifting: to each residue $\alpha \in k$ we associate the generic type of the open ball $\mathrm{res}^{-1}(\alpha)$. Since $\mathcal{O} \subseteq K$, definable types in $\mathcal{O}$ have geometric codes, so the same goes for definable types in k. This is not particularly interesting, since we already could easily see this from the fact that k is a stable stably embedded pure algebraically closed field.

But note the following

Observation 6.4.9. *If $\pi : X \to Y$ and $\pi' : X' \to Y'$ both have definable lifting, then so does the product map $\pi \times \pi' : X \times X' \to Y \times Y'$. Indeed, if f and f' are the automorphism-equivariant maps witnessing definable lifting, then $(b, b') \mapsto f(b) \otimes f'(b')$ witnesses definable lifting for the product map $\pi \times \pi'$.*

Applying this to the map $\mathcal{O} \to k$, we see that $\mathcal{O}^n \to k^n$ has definable lifting. Meanwhile, the identity map $K \to K$ has definable lifting (send each element to the associated constant type). So ultimately, we get a map $K^n \times \mathcal{O}^m \to K^n \times k^m$ with definable lifting, for every m, n. Since definable types in $K^n \times \mathcal{O}^m$ have geometric codes,

$$\text{Definable types in } K^n \times k^m \text{ have geometric codes} \tag{6.8}$$

For generically stable lifting, the analogue of Observation 6.4.9 still holds, as does the following variant:

Observation 6.4.10. *If $\pi : X \to Y$ has* generically stable *lifting, then so does* $\mathrm{Sym}^n \pi : \mathrm{Sym}^n X \to \mathrm{Sym}^n Y$.

Indeed, suppose that $b \mapsto p_b$ is the automorphism equivariant map from elements of Y to generically stable types in X. Then

$$\sigma(b_1, \ldots, b_n) \mapsto \sigma_*(p_{b_1} \otimes \cdots \otimes p_{b_n})$$

is a well-defined automorphism-equivariant map from elements of $\mathrm{Sym}^n Y$ to generically stable types in $\mathrm{Sym}^n X$, witnessing generically stable lifting for $\mathrm{Sym}^n \pi$. This is all an easy exercise; generic stability ensures that

$$\sigma_*(p_{b_1} \otimes \cdots \otimes p_{b_n}) = \sigma_*(p_{b_{\pi(1)}} \otimes \cdots \otimes p_{b_{\pi(n)}})$$

for any $b_1, \ldots, b_n$ and any permutation π of $\{1, \ldots, n\}$.

The residue map $\mathcal{O} \to k$ does not have generically stable lifting. We will use the following two maps which do have generically stable lifting:

- Let $\mathcal{B}$ denote the set of non-trivial closed balls. Let $\widetilde{\mathcal{B}}$ denote the set of $(x, y) \in K^2$ such that $x \neq y$. Consider the map

  $$\beta : \widetilde{\mathcal{B}} \to \mathcal{B}$$

  sending (x, y) to the smallest ball containing x and y. Then this map has generically stable lifting. Indeed, if B is a closed ball, with generic type p_B, then $\beta(x, y) = B$ for any $(x, y) \models p_B^{\otimes 2}$, where p_B is the generic type of B.

- Let $\widetilde{R_{n,\ell}}$ be the set of triples $(\vec{b}, \Lambda, V)$, where Λ is a lattice in K^n, $\vec{b}$ is a lattice basis, and V is an ℓ-dimensional k-subspace of $\mathrm{res}\,\Lambda := \Lambda/\mathfrak{M}\Lambda$. Then the canonical map

  $$\widetilde{R_{n,\ell}} \to R_{n,\ell}$$

  $$(\vec{b}, \Lambda, V) \mapsto (\Lambda, V)$$

  has generically stable lifting. Indeed, given Λ, any realization of $p_\Lambda^{\otimes n}$ will be a basis of Λ, where p_Λ is the generic type of Λ.

Lemma 6.4.11. *Let S be a finite C-definable subset of K. Then there is a tuple α from a power of K, such that $\alpha \overset{def}{\downarrow}_C C$ and there is an $\alpha^\frown S^\frown$-definable injection from S to a power of k.*

Proof. Let $T \subseteq \mathcal{B}$ be the set of all $\beta(x, y)$ for $x, y \in S$, and $n = |T|$. Then

$$\ulcorner T \urcorner \in \mathrm{Sym}^n \mathcal{B}$$

is C-definable. We can therefore find a C-definable type lying above this in $\mathrm{Sym}^n \widetilde{\mathcal{B}}$. Let

$$\alpha \in \mathrm{Sym}^n \widetilde{\mathcal{B}} \subseteq \mathrm{Sym}^n(K^2)$$

realize this type. Then α is a code for a subset T' of K^2. By elimination of imaginaries in ACF, we can identify α with a tuple from K. Since α was a realization of a C-definable type, $\alpha \overset{\text{def}}{\downarrow_C} C$.

It remains to produce an $\alpha^\ulcorner S \urcorner$-definable injection from S to a power of k. It suffices to show that distinct elements of S have distinct types over $\{\alpha, \ulcorner S \urcorner\} \cup k$, because definable subsets of powers of k all have codes in k.[4] Since k is stably embedded[5], it suffices to consider an automorphism σ of the monster, fixing α and k pointwise and S setwise, and show that σ fixes S pointwise.

First note that if $B \in T$ is not fixed by σ, then B and $\sigma(B)$ are disjoint. The alternative is that $B \subsetneq \sigma(B)$ or $B \supsetneq \sigma(B)$, but neither of these is possible since some power of σ acts trivially on T, as T is finite. Now suppose that $z \in S$ is not fixed by σ. Let $B = \beta(z, \sigma(z))$. Then $\sigma(B)$ and B both contain $\sigma(z)$, so $B = \sigma(B)$. Let (x, y) be the unique element of T' such that $B = \beta(x, y)$. As σ fixes $\alpha = \ulcorner T' \urcorner$ and B, it fixes x and y. The set $\mathrm{res}\, B$ of open subballs of B of the same radius as B has the structure of an affine line over k. After naming the two elements $\mathrm{res}\, x$ and $\mathrm{res}\, y$, it is in definable bijection with k. Therefore σ fixes $\mathrm{res}\, B$ pointwise. In particular, $\mathrm{res}\, z = \mathrm{res}\, \sigma(z)$. But this means that $\beta(z, \sigma(z))$ is strictly smaller than B, a contradiction. □

Lemma 6.4.12. *Definable types in* $\mathrm{Sym}^n(K^m \times k^\ell)$ *have geometric codes.*

Proof. Let p be a definable type in $\mathrm{Sym}^n(K^m \times k^\ell)$. Let M be a model over which p is defined, and let $\beta \models p|M$. Then $\beta = \sigma(a_1 b_1, \ldots, a_n b_n)$ for some $a_i \in K^m$ and $b_i \in k^\ell$. Let $S \subseteq K$ be the finite set of all elements of K occurring among the coordinates of the a_i. Thus S is β-definable, hence βM-definable. By the previous lemma, there is some tuple α of elements of K such that

$$\alpha \underset{\beta^\ulcorner p \urcorner}{\overset{\text{def}}{\downarrow}} M$$

and an $\alpha^\ulcorner S \urcorner$-definable injection f from S to a power $k^{\ell'}$ of k. For each i, let $f(a_i) \in k^{\ell' m}$ be the result of applying f coordinatewise to a_i. Now β is interdefinable over α with

$$(\sigma(a_1, \ldots, a_n), \sigma(f(a_1)b_1, \ldots, f(a_n)b_n)) \in \mathrm{Sym}^n(K^m) \times \mathrm{Sym}^n(k^{\ell' m + \ell}).$$

[4] This is because k is a stably embedded pure algebraically closed field.

[5] See the Appendix to [7].

Indeed, this element is certainly definable from α and β. Conversely, the first coordinate determines S, and S and α determine f^{-1}, which we can apply to the second coordinate to recover β.

By elimination of imaginaries in ACF, $\mathrm{Sym}^n(K^m)$ and $\mathrm{Sym}^n(k^{\ell' m+\ell})$ can be 0-definably embedded in powers of K and k, respectively. Putting everything together, we see that $\alpha\beta$ is interdefinable over $\emptyset$ with $\alpha\gamma\delta$, where $\alpha\gamma$ is a tuple from K and δ is a tuple from k.

Now $\beta \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{def}}_{\ulcorner p \urcorner} M$ and $\alpha \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{def}}_{\beta \ulcorner p \urcorner} M$, so by left-transitivity of $\mathop{\smile\!\!\!\!\!\!|}^{\mathrm{def}}$, $\alpha\beta \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{def}}_{\ulcorner p \urcorner} M$. It follows that $\mathrm{tp}(\alpha\beta/M)$ is a definable type having the same code as $\mathrm{tp}(\beta/M) = p$. Then, since $\alpha\beta$ is interdefinable with $\alpha\gamma\delta$, it follows that $\mathrm{tp}(\alpha\gamma\delta/M)$ is a definable type having the same code as p. But as we noted above (6.8), definable types in products of K and k have geometric codes. □

Theorem 6.4.13. *Let G be any geometric sort. Then elements of $\mathrm{Sym}^n G$ have geometric codes.*

Proof. We claim that there is some map $G' \to G$ with generically stable lifting, such that G' embeds (0-definably) into a product $K^m \times k^\ell$. Assuming this is true, we get a map

$$\mathrm{Sym}^n(G') \to \mathrm{Sym}^n(G)$$

with generically stable (hence definable) lifting, and $\mathrm{Sym}^n(G')$ embeds into $\mathrm{Sym}^n(K^m \times k^\ell)$. The previous Lemma ensures that definable types in $\mathrm{Sym}^n(G')$ have geometric codes, so by definable lifting, definable types in $\mathrm{Sym}^n(G)$ have geometric codes. In particular, looking at constant types in $\mathrm{Sym}^n(G)$, we see that elements of $\mathrm{Sym}^n(G)$ have geometric codes.

It remains to find $G' \to G$. The property of generically stable lifting is closed under taking products, and G is a product of K's and $R_{n,\ell}$'s, so it suffices to consider the case $G = K$ or $G = R_{n,\ell}$. For $G = K$, we take the identity map $G' = K \to K = G$. For $R_{n,\ell}$, we take the map

$$\widetilde{R_{n,\ell}} \to R_{n,\ell}$$

discussed above. It remains to embed $\widetilde{R_{n,\ell}}$ into a product of K's and k's. Let $Gr_{n,\ell}$ denote the set of ℓ-dimensional k-subspaces of k^n. Then there is a 0-definable map

$$\widetilde{R_{n,\ell}} \to K^{n^2} \times Gr_{n,\ell}$$

$$(\vec{b}, \Lambda, V) \mapsto (\vec{b}, W),$$

where W is the image of V under the identification of $\mathrm{res}\,\Lambda$ with k^n induced by the basis $\vec{b}$. This map is an injection, and $Gr_{n,\ell}$ can be embedded in a power of k by algebraic geometry, or elimination of imaginaries in ACF. □

6.4.4 Putting everything together

Theorem 6.4.14. *ACVF has elimination of imaginaries in the geometric sorts, i.e., in K and the $R_{n,\ell}$.*

Proof. This follows by Theorem 6.3.1. The second condition is Theorem 6.4.5. The third condition is Theorem 6.4.13. The first condition of Theorem 6.3.1 can be verified as follows: Let D be a one-dimensional definable set. Then D can be written as a disjoint union of $\mathrm{acl}^{eq}(\ulcorner D \urcorner)$-definable "swiss cheeses." If $B \setminus (B_1 \cup \cdots \cup B_n)$ is one of these swiss cheeses, then the generic type p_B of B is in $B \setminus (B_1 \cup \cdots \cup B_n)$, hence in D. Since B is $\mathrm{acl}^{eq}(\ulcorner D \urcorner)$-definable, so is p_B. □

6.5 Reduction to the standard geometric sorts

Haskell, Hrushovski, and Macpherson showed that ACVF has elimination of imaginaries in the sorts K, S_n, T_n. To deduce this result from Theorem 6.4.14, we need to code the $R_{m,\ell}$ sorts into the S_n and T_n. This is done in 2.6.4 of [26], but for the sake of completeness, we quickly recall the arguments here. Recall that S_n is the set of lattices in K^n, and T_n is the union of $\mathrm{res}\,\Lambda$ as Λ ranges over S_n.

First of all, we can easily code the $R_{m,\ell}$ in terms of $R_{n,0}(= S_n)$ and $R_{n,1}$ (which is roughly a projectivized version of T_n). Indeed, if Λ is a lattice in K^n, and V is an ℓ-dimensional subspace in $\mathrm{res}(V)$, then V can be coded by a one-dimensional subspace (namely $\bigwedge^\ell V$) in

$$\bigwedge^\ell \mathrm{res}(\Lambda) = \mathrm{res}(\bigwedge^\ell \Lambda),$$

and $\bigwedge^\ell \Lambda$ is a lattice in $\bigwedge^\ell K^n$. So to code an element of $R_{n,\ell}$, we can use the underlying lattice in $R_{n,0}$, and then an element in $R_{N,1}$, where $N = \dim \bigwedge^\ell K^n$.

Now to code an element of $R_{n,1}$ in terms of the S_n and T_n, we proceed by induction. Let Λ be a lattice in K^n. Let π be the projection onto the first coordinate. Then $\pi(\Lambda)$ is free[6], so we have a split exact sequence

$$0 \to \Lambda' \to \Lambda \to \pi(\Lambda) \to 0$$

where Λ' is a lattice in K^{n-1}. Since this sequence is split exact, it remains split exact after tensoring with k. So

$$0 \to \mathrm{res}(\Lambda') \to \mathrm{res}(\Lambda) \to \mathrm{res}(\pi(\Lambda)) \to 0$$

is exact. Let V be a one dimensional subspace of $\mathrm{res}(\Lambda)$. If V sits inside $\mathrm{res}(\Lambda')$, then V is interdefinable with a one-dimensional subspace of $\mathrm{res}(\Lambda')$, so can be coded in the true geometric sorts by induction.

Otherwise, V maps isomorphically onto the one-dimensional k-space $\mathrm{res}(\pi(\Lambda))$. Then to code V, it suffices to code the inverse map $\mathrm{res}(\pi(\Lambda)) \to V \hookrightarrow \mathrm{res}(\Lambda)$. But because all the $\mathcal{O}$-modules in sight are free,

$$\mathrm{Hom}_k(\mathrm{res}(\pi(\Lambda)), \mathrm{res}(\Lambda)) = \mathrm{Hom}_k(\pi(\Lambda) \otimes k, \Lambda \otimes k) = \mathrm{Hom}_{\mathcal{O}}(\pi(\Lambda), \Lambda) \otimes k.$$

And $\mathrm{Hom}_{\mathcal{O}}(\pi(\Lambda), \Lambda)$ is a lattice in $\mathrm{Hom}_K(K, K^n) \cong K^n$. So a map from $\mathrm{res}(\pi(\Lambda))$ to $\mathrm{res}(\Lambda)$ can be coded by an element of T_n.

[6]Finitely generated torsion free $\mathcal{O}$-modules are always free.

Chapter 7

Generically stable types in C-minimal expansions of ACVF

Hrushovski and Loeser's work on Berkovich spaces [34] builds off of two substantial prior papers on the model theory of ACVF by Haskell, Hrushovski, and Macpherson. The first was [26] which established elimination of imaginaries, and the second was [27] which developed the theory of stable domination. In the previous chapter, we shortened the proof of the main result of [26]. In this chapter, we focus on two important facts from [27] and [34]: the characterization of generically stable types (Proposition 2.8.1 in [34]), and the strict prodefinability of the space of generically stable types (Theorem 3.1.1 of [34]). We prove these in Theorem 7.1.2 below, while simultaneously generalizing to the setting of C-minimal expansions of ACVF.

These results were also proven by Hrushovski in some unpublished notes. The only new result here is the strict prodefinability of the space of generically stable types in C-minimal expansions of ACVF.

7.1 Definitions

Let T be some C-minimal expansion of ACVF. Let $\mathbb{M}$ be the monster model of T. Let K be the home sort, k be the residue field, and Γ be the value group. The value group Γ of $\mathbb{M}$ is an o-minimal expansion of a divisible ordered abelian group. Let $\Gamma(A)$ denote $\mathrm{dcl}^{eq}(A) \cap \Gamma$ for any subset $A \subseteq \mathbb{M}^{eq}$.

Remark 7.1.1. *Let p be a global C-invariant type. The following are equivalent:*

- *For every function f into Γ (defined with parameters from $\mathbb{M}$), the pushforward f_*p is a constant type.*
- *For every $B \supseteq C$, we have $\Gamma(Ba) = \Gamma(B)$ for a realizing $p|B$.*

We say that p is orthogonal to Γ *if these conditions hold. In particular, from the first bullet point, this is a property of p, rather than the pair (p, C).*

Proof. Suppose the first condition holds. Let $B \supseteq C$ and let a be any realization of $p|B$. For $\gamma \in \Gamma(Ba)$, we can write γ as $f(a)$ for some B-definable function. Then $\gamma \models f_*p|B$. Also, p is B-invariant and f is B-definable, so the type f_*p is B-invariant. Since it is constant, it must contain the formula $x = \gamma_0$ for some γ_0, and γ_0 must be B-definable. Therefore the formula $x = \gamma_0$ is in $f_*p|B$, and so $\gamma = \gamma_0 \in \Gamma(B)$. As γ was an arbitrary element of $\Gamma(Ba)$, we conclude that $\Gamma(Ba) = \Gamma(B)$.

Conversely, suppose that the second condition holds. Let f be an $\mathbb{M}$-definable function into Γ. Let B be a set containing C, over which f is defined. Let a realize $p|B$. Then $f(a) \in \Gamma(Ba) = \Gamma(B)$. Since $f(a) \models f_*p|B$, and $f(a)$ is B-definable, the formula $x = f(a)$ must be in $f_*p|B$, so f_*p is a constant type. □

Recall the definition of *domination* from Definition 6.2.3. Say that a type is *stably dominated* if it is dominated by a stable type.

A *pro-definable* set is an inverse limit of definable sets. A *strict pro-definable* set is an inverse limit of definable sets along surjections. An inverse limit $D = \varprojlim_\alpha D_\alpha$ is strictly pro-definable if and only if all the images $D \to D_\alpha$ have definable (rather than type-definable) image.

Theorem 7.1.2. *Let $\mathbb{M}$ be a model of a C-minimal expansion of ACVF.*

1. *Let $p(x)$ be a global invariant type in $\mathbb{M}^{eq}$. Then the following are equivalent:*
   - *a) $p(x)$ is generically stable*
   - *b) $p(x)$ is orthogonal to the value group*
   - *c) $p(x)$ is stably dominated*
2. *Let V be an interpretable set. There is a strictly prodefinable set $\hat{V}$ parametrizing the generically stable types in V.*

Compare with Proposition 2.8.1 and Theorem 3.1.1 in [34]. Part 1 appears (in more generality) in some unpublished notes of Hrushovski, though we will give a simpler and more complete proof in our setting. Part 2 is apparently new (in the generality of C-minimal expansions).

7.2 Generic stability and orthogonality to Γ

In this section, we show

Proposition 7.2.1. *Let $p(x)$ be a global invariant type. Then $p(x)$ is generically stable if and only if $p(x)$ is orthogonal to Γ.*

Proof. One direction is easy: if p is generically stable, and f is a definable function into Γ, then f_*p is a generically stable type in Γ. The Morley sequence of this type is totally indiscernible. But a totally indiscernible sequence in a totally ordered set must be constant. This ensures that f_*p is constant. The other direction will use Theorem 5.3.3 from §5.3. In fact, the result will follow almost immediately from the following

Lemma 7.2.2. *Let $A \subseteq \mathbb{M}^{eq}$ be small. If $\mathcal{C}$ is a chain of unary definable sets, then $|\mathcal{C}| \leq |\Gamma(A)| + \aleph_0$.*

Proof. For $\gamma \in \Gamma$, let B_γ and $\overline{B}_\gamma$ be the open and closed balls of radius γ. Let $\mathbb{M}_0$ be the reduct of $\mathbb{M}$ to the group language expanded by unary predicates for B_γ and $\overline{B}_\gamma$, for each $\gamma \in \Gamma(A)$. Then every A-definable unary set is $\mathbb{M}_0$-definable in $\mathbb{M}_0$, by C-minimality. In fact, if D is an A-definable unary set, then D is a boolean combination of balls with radii in $\Gamma(A)$, and each ball is a translate of some B_γ or $\overline{B}_\gamma$.

The structure $\mathbb{M}_0$ is stable, and its language has size at most $|\Gamma(A)| + \aleph_0$, so it contains no chains of definable sets of length greater than $|\Gamma(A)| + \aleph_0$. (Otherwise, by the pigeonhole principle we could find a chain of uniformly definable sets, on which an indiscernible chain could be modeled, contradicting stability). □

Now suppose $p(x)$ is orthogonal to Γ. We will show that p is chain averse in the sense of §5.3. Let C be a small set. Enlarging C, we may assume p is C-invariant. Let $\kappa(C) = (|\Gamma(C)| + \aleph_0)^+$. Let $\vec{a}$ be an arbitrarily long Morley sequence in C. By Remark 7.1.1 and induction, $\Gamma(C\vec{a}) = \Gamma(C)$. By the Lemma, there is no chain of unary $C\vec{a}$-definable sets of size greater than κ. Therefore $p(x)$ is chain averse, hence generically stable by Theorem 5.3.3. □

7.3 Stable domination

Next, we turn to proving that stable domination is equivalent to the other two conditions of Theorem 7.1.2.1.

If $A \subseteq \mathbb{M}^{eq}$, write $k(A)$ for $k \cap \mathrm{dcl}(A)$.

Remark 7.3.1. $k \cap \mathrm{acl}(A) = k \cap \mathrm{acl}(k(A))$.

Proof. Suppose $\alpha \in k$ is algebraic over A. Let S be the finite set of conjugates of α over A. Then S has a code in a power of k, so S has a code in $k(A)$. Therefore α is algebraic over $k(A)$. □

Let $p(x)$ be a generically stable type over C_0, thought of as a C_0-definable type over $\mathbb{M}$. The type $p(x)$ might live in an imaginary sort.

We will prove that there is a small set $C \supseteq C_0$ and a C-definable map f into a power of k such that p is "dominated" over C by its pushforward along f. That is, for every $D \supseteq C$ and every a, the following will be equivalent:

- $a \models p|D$

- $a \models p|C$ and $f(a) \models f_*p|D$.

Lemma 7.3.2. *Let $C = \mathrm{acl}(C)$ be a set of parameters. Suppose $\mathrm{tp}(a/C)$ is generically stable, for some $a \in \mathbb{M}^{eq}$, and suppose $b \in \mathrm{acl}(Ca)$. Then $\mathrm{tp}(b/C)$ is generically stable.*

Proof. Let p be the canonical global extension of $\mathrm{tp}(a/C)$, and let $M \supseteq C$ be a small model in which p is finitely satisfiable. On general grounds, $\mathrm{tp}(ab/C)$ is definable; let q be its canonical global extension. It suffices to show that $q(x, y)$ is finitely satisfiable in M, since this ensures that q is generically stable, hence so is its pushforward along the projection to the second coordinate.

Let d be an element of $\mathbb{M}$, and suppose $\phi(x, y, d) \in q(x, y)$. Let $\psi(x; y)$ be a C-formula such that $\psi(a; b)$ holds and $\psi(a'; \mathbb{M})$ is finite for every a'. Such a formula exists because $b \in \mathrm{acl}(Ca)$. Let $a'b' \models q|Md$. Then $a'b' \equiv_M ab$, so $\psi(a'; b')$ holds. Also, $\phi(a', b', d)$ holds. The formula

$$\exists y(\phi(x, y, d) \wedge \psi(x; y))$$

is in $p(x)$, because a' satisfies it and $a' \models p|Cd$. Because p is finitely satisfiable in M, there is some $a'' \in M$ such that

$$\exists y(\phi(a'', y, d) \wedge \psi(a''; y))$$

Choose b'' such that $\phi(a'', b'', d) \wedge \psi(a''; b'')$ holds. Then $b'' \in \mathrm{acl}(Ca'') \subseteq M$. So the pair (a'', b'') is in M, and satisfies $\phi(x, y, d)$. □

Lemma 7.3.3. *Suppose C is a small set of parameters, and B is a ball in K^1, (possibly a singleton). Suppose $\mathrm{tp}(\ulcorner B \urcorner/C)$ is generically stable. Then either B is C-definable, or there exists $A \supseteq A' \supseteq B$, where A is a C-definable closed ball of some radius, A' is an open ball of the same radius, and A' is not defined over $\mathrm{acl}(C)$.*

Proof. Let $B = B_1, B_2, \ldots$ be a Morley sequence for the type of $\ulcorner B \urcorner$ over C. Assume B is not C-definable. Then the type is not constant, so the B_i's are distinct. Since the type is generically stable, this sequence is totally indiscernible. Consequently, $B_i \cap B_j = \emptyset$ for $i \neq j$. Let A_{ij} be the smallest ball containing both B_i and B_j, for $i \neq j$. Then A_{ij} is a closed ball. The total indiscernibility of the sequence implies that $A := A_{ij}$ does not depend on i, j. As $A_{1,2}, A_{3,4}, A_{4,5}, \ldots$ is a Morley sequence of a C-definable type (a pushforward of $\mathrm{tp}(\ulcorner B \urcorner/C)^{\otimes 2}$), it follows that A is C-definable.

Let A'_i be the open subball of A of the same radius, containing B_i. Then the sequence $A'_1, A'_2, \ldots$ is a Morley sequence over C. As A is the smallest ball containing B_i and B_j, we must have $A'_i \cap A'_j = \emptyset$ for $i \neq j$. So the elements of the sequence $A'_1, A'_2, \ldots$ are pairwise distinct. As the sequence is C-indiscernible, it follows that the elements are not algebraic over C. In particular, $A' := A'_1 \supseteq B$ is not algebraic over C. □

Lemma 7.3.4. *Suppose C is a small set of parameters, and a and b are from $\mathbb{M}^{eq}$ such that $\mathrm{stp}(a/C)$ is generically stable. Then $a \mathop{\downarrow}_C b \iff b \mathop{\downarrow}_C a$.*

Proof. On general grounds, we may replace C with $\mathrm{acl}(C)$, so we may assume $C = \mathrm{acl}(C)$. Let $p(x)$ be the unique global non-forking extension of $\mathrm{tp}(a/C) = \mathrm{stp}(a/C)$. By Corollary 2.14 in [36], there is some C-invariant type $q(y)$ extending $\mathrm{tp}(b/C)$.

Suppose that $a \mathop{\smile\!\!\!\!\!\mid}_C b$. Then $\mathrm{tp}(a/Cb)$ does not fork over $\mathrm{tp}(a/C)$, so it must be $p|Cb$. Then $a \models p|Cb$ and $b \models q|C$, or equivalently, $(a, b) \models p \otimes q|C$. By one of the characterizations of generic stability, $(b, a) \models q \otimes p|C$. So $\mathrm{tp}(b/Ca) = q|Ca$. Since q does not fork over C, $b \mathop{\smile\!\!\!\!\!\mid}_C a$.

Conversely, suppose that $b \mathop{\smile\!\!\!\!\!\mid}_C a$. Then by the characterization of forking in NIP theories (Proposition 2.1(i) in [36]), $\mathrm{tp}(b/Ca)$ has some global extension $r(y)$ which is Lascar C-invariant. By Corollary 2.14 in [36], $r(y)$ is C-invariant. Then $b \models r|Ca$ and $a \models p|C$, so $(b, a) \models r \otimes p|C$. As before, this implies that $(a, b) \models p \otimes r|C$, so $a \models p|Cb$. As p is C-invariant, $a \mathop{\smile\!\!\!\!\!\mid}_C b$. □

Now fix a generically stable type $p(x)$, defined over some base set of parameters C_0. The variable x might live in an imaginary sort.

Lemma 7.3.5. *For $C \supseteq C_0$, let $r(C)$ denote the supremum of $RM(\alpha/C)$, where α is a tuple in $k(Ca)$ and a realizes $p|C$. (By Remark 7.3.1, we could even let α range over $k \cap \mathrm{acl}(Ca)$, and $r(C)$ would not change.)*

(a) *There is an integer n such that $r(C) \leq n$ for every $C \supseteq C_0$.*

(b) *If $C' \supseteq C \supseteq C_0$, then $r(C') \geq r(C)$.*

Consequently, there is some $C \supseteq C_0$ such that $r(C') = r(C)$ for every $C' \supseteq C$.

Proof. **(a)** C-minimal theories are dp-minimal, so the home sort has dp-rank 1. By additivity of dp-rank in NIP theories, every interpretable set in T has finite dp-rank. Let n be the dp-rank of the sort where the variable x lives. Suppose $C \supseteq C_0$, a realizes $p|C$, and α is a tuple in $k(Ca)$. Suppose for the sake of contradiction that $RM(\alpha/C) \geq n + 1$. As k is a strongly minimal set, we can replace α with some subtuple, and assume that α has length $n + 1$, and that it realizes the generic type of k^{n+1}, over C. Write α as $f(a)$ for some C-definable function f. Then the range of f has dp-rank at most n. But the generic type of k^{n+1} over C has dp-rank (at least) n, a contradiction.

(b) Suppose a realizes $p|C$ and $\alpha \in k(Ca)$ has $RM(\alpha/C) = m$. Moving C' over C, we may assume that a realizes $p|C'$. As p is C_0-definable, hence C-invariant, $a \mathop{\smile\!\!\!\!\!\mid}_C C'$. So $\alpha \mathop{\smile\!\!\!\!\!\mid}_C C'$. Consequently, $RM(\alpha/C) = RM(\alpha/C')$. And $\alpha \in k(C'a)$.

□

Fix some C as in the conclusion of the lemma. Let $m = r(C)$. Fix some C-definable function f into k^m such that $f_* p$ is the generic type of k^m.

For B a non-degenerate (infinite) closed ball, let $\mathrm{res}\, B$ denote the interpretable set of open subballs of the same radius.

Lemma 7.3.6. *Suppose $C' \supseteq C$. Suppose B is a C'-definable closed ball. Suppose $a \models p|C'$ and that $\alpha \in \operatorname{res} B$ is algebraic over $C'a$. Then α is algebraic over $C'f(a)$.*

Proof. Let e and d realize (independently) the generic type of B over $C'a$. Then $ed \mathop{\smile\!\!\!\!\!\!|}_{C'} a$, hence $ed \mathop{\smile\!\!\!\!\!\!|}_{C'} \alpha f(a)$. By base monotonicity on the right (which holds for forking in arbitrary theories), $ed \mathop{\smile\!\!\!\!\!\!|}_{C'f(a)} \alpha$.

Over $C'ed$, $\operatorname{res} B$ is in definable bijection with k, via the map sending the class of $x \in B$ to $\operatorname{res}((x-e)/(d-e))$, for example. So α is interdefinable over $C'ed$ with some $\alpha' \in k$. If $\alpha' \notin \operatorname{acl}(f(a)C'ed)$, then $\alpha' f(a)$ realizes the generic type of k^{m+1} over $C'ed$, so $r(C'ed) = m+1 > m = r(C)$, contradicting Remark 7.3.1. Therefore $\alpha' \in \operatorname{acl}(f(a)C'ed)$, and hence $\alpha \in \operatorname{acl}(f(a)C'ed)$. Since $ed \mathop{\smile\!\!\!\!\!\!|}_{C'f(a)} \alpha$, it follows that $\alpha \mathop{\smile\!\!\!\!\!\!|}_{C'f(a)} \alpha$. This can only happen if $\alpha \in \operatorname{acl}(C'f(a))$. □

Lemma 7.3.7. *Suppose $C' \supseteq C$. Suppose $a \models p|C'$. Suppose that b is a* singleton in the home sort*. Suppose that the type of $f(a)$ over $C'b$ is the generic type of k^m. Then $a \models p|C'b$.*

Proof. As $\operatorname{tp}(a/C')$ is stationary, it implies $\operatorname{stp}(a/C')$. So $a \models p|\operatorname{acl}(C')$. Similarly, the type of $f(a)$ over $\operatorname{acl}(C'b)$ is still generic in k^m. Replacing C' with $\operatorname{acl}(C')$, we may assume that $C' = \operatorname{acl}(C')$.

Let $\phi(x;y)$ be a C'-formula, and suppose $\phi(x;b) \in p(x)$. We will show that $\phi(a;b)$ holds. Let D be the definable set $\phi(a;\mathbb{M})$. This can be written as a boolean combination of $\operatorname{acl}(aC')$-definable balls $B_1, \ldots, B_n$. By Lemma 7.3.2, $\operatorname{tp}(\ulcorner B_i \urcorner / C')$ is generically stable for each i.

Claim 7.3.8. *For each i, either B_i is C'-definable or $b \notin B_i$.*

Proof. Suppose B_i is not C'-definable. By Lemma 7.3.3, we have the following setup: there is some C'-definable closed ball A containing B_i, and some open ball A' of the same radius, with $A \supseteq A' \supseteq B_i$, *and (the code for) A' is not algebraic over C'*. Now $\ulcorner A' \urcorner$ is an element $\alpha \in \operatorname{res} A'$, and α is definable from $\ulcorner A \urcorner$ and $\ulcorner B_i \urcorner$. As $\ulcorner A \urcorner$ is C'-definable and $\ulcorner B_i \urcorner$ is algebraic over a and C', it follows that $\alpha \in \operatorname{acl}(C'a)$. By Lemma 7.3.6, $\alpha \in \operatorname{acl}(C'f(a))$.

Since $f(a)$ realizes the generic type of k^m over $C'b$, we have $f(a) \mathop{\smile\!\!\!\!\!\!|}_{C'} b$. Consequently $\alpha \mathop{\smile\!\!\!\!\!\!|}_{C'} b$. If $b \in B_i$, then the code α for A' is algebraic over $C'b$, so we would have $\alpha \mathop{\smile\!\!\!\!\!\!|}_{C'} \alpha$. This contradicts the fact that A' is not algebraic over C'. □

Let $a^1 = a$ and $B_i^1 = B_i$. Choose $a^2, a^3, \ldots$ and B_i^j such that

$$\langle a^j \ulcorner B_1^j \urcorner \ulcorner B_2^j \urcorner \cdots \rangle_{j=2,3,\ldots}$$

is a Morley sequence over $ba\ulcorner B_1 \urcorner \cdots$ for the type

$$\operatorname{tp}(a\ulcorner B_1 \urcorner \ulcorner B_2 \urcorner \cdots / C')$$

which is generically stable by Lemma 7.3.2. Then

$$\langle a^j \ulcorner B_1^j \urcorner \ulcorner B_2^j \urcorner \cdots \rangle_{j=1,2,\ldots}$$

is a Morley sequence for this type, over C'. Also, $a^2 \models p|C'b$, so $\phi(a^2; b)$ holds if and only if $\phi(x; b) \in p(x)$. Therefore, it suffices to show for each i that

$$b \in B_i^2 \iff b \in B_i^1.$$

Note that $B_i^1, B_i^2, \ldots$ is a Morley sequence over C', and $B_i^2, B_i^3, \ldots$ is a Morley sequence over $C'b$. If $B_i = B_i^1$ is C'-definable, this sequence is constant, so $b \in B_i^1 \iff b \in B_i^2$. Otherwise, by total indiscernibility, the B_i^j are pairwise disjoint (for fixed i). So $b \notin B_i^j$ for all $j > 1$. But by the claim, $b \notin B_i^1$ either. So we are done. □

Theorem 7.3.9. *Suppose that $C' \supseteq C$ and $a \models p|C$ and $f(a)$ realizes the generic type of k^m over C'. Then $a \models p|C'$. So the (arbitrary generically stable type p) is stably dominated.*

Proof. Take some set C'' of real elements such that $C' \subseteq \mathrm{dcl}(C'')$. Moving C'' over C', we may assume that $f(a)$ realizes the generic type of k^m over C''. Replacing C' with C'', we may assume that C' is made of real elements.

Let $b_1, \ldots, b_n$ be a tuple from C', and suppose $\phi(x; b)$ is in $p(x)$. It suffices to show that $\phi(a; b)$ holds. It suffices to show that $a \models p|Cb_1b_2 \cdots b_n$.

We prove by induction on i that $a \models p|Cb_1 \cdots b_i$. The base case where $i = 0$ is given. Suppose that $a \models p|Cb_1 \cdots b_{i-1}$. By Lemma 7.3.7, we need only show that $\mathrm{tp}(f(a)/Cb_1 \cdots b_i)$ is the generic type of k^m. This is clear, though, since $\mathrm{tp}(f(a)/C')$ was generic in k^m, and $Cb_1 \cdots b_i \subseteq C'$. □

7.4 Strict prodefinability

For any set X, the *stable completion* of X is the set $\hat{X}$ of generically stable types in X. This can be seen as a pro-definable set, in such a way that $\hat{X}(C)$ is canonically identified with the set of C-definable generically stable types in X.

Let $X \to Y$ be a definable surjection of interpretable sets. We will show that the induced map $\hat{X} \to \hat{Y}$ on the stable completions is surjective. Using this, we can mimic the proof from Hrushovski and Loeser that $\hat{X}$ is not only pro-definable (as it would be in any NIP theory), but also *strictly pro-definable*.

7.4.1 Lifting stably dominated types

First we recall some assorted facts about chaining together definable, generically stable, and algebraic types.

Remark 7.4.1. *Recall that if M is a model and $\mathrm{tp}(a/M)$ is definable (resp. generically stable), then $\mathrm{tp}(a'/M)$ is definable (resp. generically stable), for any $a' \in \mathrm{acl}(aM)$. If $\mathrm{tp}(a/M)$ is definable (resp. generically stable) and $\mathrm{tp}(b/aM)$ has an aM-definable extension (resp.* is *generically stable), then $\mathrm{tp}(ab/M)$ is definable (resp. is generically stable).*

Lemma 7.4.2. *Let M be a model (an elementary substructure of $\mathbb{M}$) and suppose $\mathrm{tp}(a/M)$ is generically stable. Let $D \subseteq K^1$ be aM-definable and non-empty. Then there is some $b \in D$ such that $\mathrm{tp}(ab/M)$ is generically stable.*

Proof. Let c be a code for one of the swiss cheeses in the canonical decomposition of D into swiss cheeses. Then $c \in \mathrm{acl}(aM)$, so by Remark 7.4.1, $\mathrm{tp}(ac/M)$ is generically stable. Replacing a with ac and D with the swiss cheese coded by c, we may assume that D is a swiss cheese.

Suppose that there is a closed ball B' which is $\mathrm{acl}(aM)$-definable, such that the generic type of B' is in D. Then $\mathrm{tp}(\ulcorner B' \urcorner a/M)$ is generically stable (by Remark 7.4.1) and if b realizes the generic type of B', then $\mathrm{tp}(b \ulcorner B' \urcorner a/M)$ is generically stable, by Remark 7.4.1 again, so we are done.

So assume that there is no closed ball B' which is $\mathrm{acl}(aM)$-definable, such that the generic type of B' is in D.

As a swiss cheese, we can write D as $B_0 \setminus (B_1 \cup \cdots \cup B_n)$, where each B_i is a ball (open or closed, possibly K or a singleton), where $n \geq 0$, and where $B_i \subsetneq B_0$ for $i > 0$. If B_0 is a closed ball (or a singleton), then the generic type of B_0 is in D, and B_0 is aM-definable, contradicting our assumption.

Next suppose that B_0 is all of K. If $n = 0$, then $D = K$, and we can take some $b \in D \cap \mathrm{dcl}(M)$ because M is a model. Otherwise, let B be the smallest closed ball containing $B_1, \ldots, B_n$. Since M is a model, there is some $\delta < 0$ in $\Gamma(M)$. Let B' be the closed ball around B of radius δ plus the radius of B. Then the generic type of B' is in D, and B' is algebraic (in fact, definable) over aM, a contradiction.

We are left with the case that B_0 is an open ball. If $n = 1$, then we can take a closed ball between B_0 and B_1 whose (valuative) radius is halfway between the radii of B_0 and B_1 (or δ plus the radius of B_0, in the case where B_1 is a singleton). If $n > 1$, let B be the smallest closed ball containing $B_1, \ldots, B_n$. Then B is strictly smaller than B_0, and strictly bigger than each of the B_i's, so its generic type is in D, a contradiction.

So we are left with the case that $D = B_0$ is an open ball. Suppose that there is some $\mathrm{acl}(aM)$-definable subball $B \subsetneq B_0$ (open or closed or singleton) of D. Then as above, we can take a closed ball halfway between B and B_0 (or use δ), and find an $\mathrm{acl}(aM)$-definable closed subball of D, a contradiction.

So we may assume that not only is D an open ball, but that no proper subball of D is $\mathrm{acl}(aM)$-definable. From the swiss cheese decomposition, it follows that the only $\mathrm{acl}(aM)$-definable subsets of D are D and $\emptyset$.

Let b realize the generic type of D. This type is aM-definable, so $\mathrm{tp}(ab/M)$ is definable. It remains to show that $\mathrm{tp}(ab/M)$ is orthogonal to Γ. Suppose not. Then there is some aM-definable function $f : K^1 \to \Gamma^1$ such that $f(b) \notin \Gamma(M)$.

The set $f(D)$ is a definable subset of Γ. By o-minimality of Γ, it is a finite union of intervals. The endpoints of these intervals are in $\Gamma(Ma) = \Gamma(M)$. If $f(D)$ is finite, then $f(b) \in f(D) \subseteq \Gamma(M)$, a contradiction. So $f(D)$ contains an infinite interval. As M is a

model, this infinite interval contains at least three M-definable points $\gamma_1, \gamma_2, \gamma_3$. Then $f^{-1}(\gamma_1)$ and $f^{-1}(\gamma_2)$ and $f^{-1}(\gamma_3)$ are three distinct aM-definable subsets of D, a contradiction. □

Lemma 7.4.3. *Let M be a model (an elementary substructure of $\mathbb{M}$) and suppose $\mathrm{tp}(a/M)$ is generically stable. Let $D \subseteq K^n$ be aM-definable and non-empty. Then there is some $b \in D$ such that $\mathrm{tp}(ab/M)$ is generically stable.*

Proof. By induction on n. The $n = 1$ case was Lemma 7.4.2. Suppose $n > 1$. Let π be the projection $K^n \to K^{n-1}$ coordinates. By induction there is some $b_0 \in \pi(D)$ such that $\mathrm{tp}(ab_0/M)$ is generically stable. Then $\pi^{-1}(b_0) \cap D = \{b_0\} \times D'$ for some non-empty ab_0M-definable $D' \subseteq K^1$. By Lemma 7.4.2, there is some $c \in D'$ such that $\mathrm{tp}(ab_0c/M)$ is generically stable. Take $b = (b_0, c)$. □

Theorem 7.4.4. *Let $f : X \to Y$ be a definable surjection. Let p be a generically stable type in Y. Then there is a generically stable type q in X such that $f_*q = p$. In fact, if p, f, X, and Y are defined over a model M, we can take q to be defined over the same model.*

Proof. Let X' be some (M-)definable subset of K^n such that there is an (M-)definable surjection from X' onto X. If we can lift p to X' (along the composition $X' \to X \to Y$), then we can certainly lift it to X. Replacing X with X', we may assume that X is a definable subset of K^n for some n.

Let a realize $p|M$. Let D be $f^{-1}(a)$, a non-empty aM-definable set. By Lemma 7.4.3, there is some $b \in D$ such that $\mathrm{tp}(b/M)$ is generically stable. Take q to be the canonical global extension of $\mathrm{tp}(b/M)$. □

7.4.2 Proving strictness

In any NIP theory, one has uniform definability of generically stable types. That is, for every formula $\phi(x; y)$ there is some formula $\psi(y; z)$ such that for every generically stable type p, $(d_px)\phi(x; y)$ is of the form $\psi(y; c)$ for some $c \in M$. This follows from the fact that generically stable types are definable by voting in Morley sequences. So in fact we can take $\psi(y; z)$ to be of the form

$$\psi(y; z_1, \ldots, z_N) := \bigvee_{\substack{S \subseteq \{1,\ldots,2N-1\} \\ |S| \geq N}} \bigwedge_{i \in S} \phi(z_i; y)$$

where N is one or two times the alternation number of $\phi(x; y)$.

From this it follows that

Lemma 7.4.5. *If X is a definable set in* any *NIP theory, the space $\hat{X}$ of generically stable types in X is pro-definable (in T^{eq}).*

Proof. For each formula $\phi(x; y)$, choose some formula $\psi_\phi(y; z)$ which gives uniform definitions. Since we are working in T^{eq}, we may arrange that $\psi_\phi(\mathbb{M}; z) \neq \psi_\phi(\mathbb{M}; z')$ for $z \neq z'$.

Let V_ϕ be the sort where z lives. So if p is a generically stable type, then the code for the ϕ-definition of p is an element of V_ϕ.

So we have a map from generically stable types to $\prod_{\phi\in\mathcal{L}} V_\phi$. It remains to show that the range of this map is $*$-definable.

A tuple $\langle c_\phi\rangle_{\phi\in\mathcal{L}}$ will define a consistent global type if and only if it defines a type which is finitely satisfiable. This can be expressed as follows: for every $\phi_1(x;y_1),\ldots,\phi_n(x;y_n)$ in the language, the following must hold:

$$\mathbb{M}\models\forall y_1,\ldots,y_n\exists x\bigwedge_{i=1}^{n}(\phi_i(x;y_i)\leftrightarrow\psi_{\phi_i}(y_i;c_{\phi_i}))$$

So the set of tuples $\langle c_\phi\rangle_\phi$ for which we get a consistent type is $*$-definable.

Now a definable type p is generically stable iff $p(x_1)\otimes p(x_2)=p(x_2)\otimes p(x_1)$. Equivalently, for every formula $\phi(x_1;x_2;y)$,

$$(d_px_1)(d_px_2)\phi(x_1,x_2;y)=(d_px_2)(d_px_1)\phi(x_1;x_2;y). \tag{7.1}$$

We can express this as a condition in terms of the c_ϕ's. Let $\phi_1(x_2;x_1,y)$ be $\phi(x_1;x_2;y)$. Let $\phi_2(x_1;y,z)$ be $\psi_{\phi_1}(x_1,y;z)$. Let $\phi_3(y,z,w)$ be $\psi_{\phi_2}(y,z,w)$. If p is the definable type defined by the c_ϕ's, then

$$(d_px_1)(d_px_2)\phi(x_1,x_2;y)=(d_px_1)(d_px_2)\phi_1(x_2;x_1,y)=(d_px_1)\phi_2(x_1;y,c_{\phi_1})=\phi_3(y,c_{\phi_1},c_{\phi_2}).$$

Similarly, we can find some formulas ϕ_4,ϕ_5,ϕ_6 such that

$$(d_px_2)(d_px_1)\phi(x_1,x_2;y)=\phi_4(y,c_{\phi_5},c_{\phi_6}).$$

Then (7.1) is essentially the assertion that

$$\forall y\quad\phi_3(y,c_{\phi_1},c_{\phi_2})\leftrightarrow\phi_4(y,c_{\phi_5},c_{\phi_6})$$

Doing this for each $\phi(x_1,x_2,y)$ in the language, we get a small family of first order statements about the c_ϕ whose conjunction is equivalent to the condition that the resulting type is generically stable. □

The place where we must use C-minimality is to show *strict* pro-definability. It suffices to show that the image of the map

$$\hat{X}\to\prod_{i=1}^{n}V_{\phi_i}$$

is definable for all finite sets of formulas $\phi_1,\ldots,\phi_n$. But this map factors through $\hat{X}\to V_\psi$ for some ψ, so we can reduce to the case of showing that $\hat{X}\to V_\phi$ has definable image.

So we are reduced to proving the following:

Theorem 7.4.6. *(Assuming we are in a C-minimal expansion of ACVF.) Let X be a definable set. Let $\phi(x;y)$ be a formula. The set of ϕ-definitions of generically stable types in X is a small union of definable sets. (Since it is also type-definable, this implies that it is definable.)*

Proof. We will use the same argument as Hrushovski and Loeser, except using the previous section instead of metastability (which may or may not work in this setting).

Let $\psi(y;z)$ be the formula that uniformly ϕ-defines generically stable types. Let g be the generic type of k, so that $g^{\otimes n}$ is the generic type of k^n.

For each definable map $f : X \times K^m \to k^n$, and $w \in K^m$, let f_w denote the map $f(-,w) : X \to k^n$. Let W_f be the (definable) set of $w \in K^m$ such that $(d_{g^{\otimes n}}s)(s \in f_w(X))$, i.e., such that $f_w(X)$ hits the generic type of k^n. For $w \in W_f$, let $Z_{f,w}$ be the set of z such that

$$\forall y\, d_{g^{\otimes n}} s\, \forall x \in f_w^{-1}(s) \quad (\phi(x;y) \leftrightarrow \psi(y;z))$$

Note that $Z_{f,w}$ is definable uniformly in w. So the union of all the $Z_{f,w}$'s is a small union of definable sets.

We claim that this union is exactly the set of c such that $\psi(y;c)$ is the ϕ-definition of a generically stable type.

Suppose first that $\psi(y;c)$ is the ϕ-definition of some generically stable type $p(x)$. Since p is generically stable, there is some set of parameters C over which p is defined, and some C-definable map $f_0 : X \to k^n$ such that $f_{0,*}p$ is $g^{\otimes n}$ and p is "dominated" along f_0.

That is, $(p|C)(x) \cup g^{\otimes n}(f_0(x)) \vdash p(x)$.

Claim 7.4.7. *There is some finite subtype $\Sigma(x)$ of $(p|C)(x)$ such that $\Sigma(x) \cup g^{\otimes n}(f_0(x))$ implies the restriction of p to a ϕ-type.*

Proof. For each subtype $\Sigma(x)$ of $p|C$, let S_Σ denote the set of b such that

$$\Sigma(x) \cup g^{\otimes n}(f_0(x)) \vdash \phi(x;b)$$

and let S'_Σ denote the set of b such that

$$\Sigma(x) \cup g^{\otimes n}(f_0(x)) \vdash \neg\phi(x;b).$$

The set of formulas in $g^{\otimes n}(f_0(x))$ is ind-definable because $g^{\otimes n}$ is a definable type. So each of S_Σ and S'_Σ is small union of definable sets. Note that $S_{p|C} = \psi(\mathbb{M};c)$ and $S'_{p|C} = \neg\psi(\mathbb{M};c)$. In particular, $S_{p|C}$ and $S'_{p|C}$ are definable. By the most basic form of compactness,

$$S_{p|C} = \bigcup_{\substack{\Sigma \subseteq p|C \\ \Sigma \text{ finite}}} S_\Sigma$$

$$S'_{p|C} = \bigcup_{\substack{\Sigma \subseteq p|C \\ \Sigma \text{ finite}}} S'_\Sigma$$

By saturation of the monster model, it follows that $S_{p|C} = S_\Sigma$ and $S_{p|C} = S'_\Sigma$ for some finite $\Sigma \subseteq p|C$. Then for every b, if $\phi(x;b) \in p(x)$, then $b \in S_{p|C} = S_\Sigma$, so $\Sigma(x) \cup g^{\otimes n}(f_0(x)) \vdash \phi(x;b)$. And similarly, if $\neg\phi(x;b) \in p(x)$, then $\Sigma(x) \cup g^{\otimes n}(f_0(x)) \vdash \neg\phi(x;b)$. So $\Sigma(x) \cup g^{\otimes n}(f_0(x))$ implies the restriction of p to a ϕ-type. □

Let f be f_0 on $\Sigma(\mathbb{M})$, and $0 \in k^n$ off of $\Sigma(\mathbb{M})$. Since Σ is a finite type, f is still a (C-)definable function. If $a \models p|C$, then $f(a) = f_0(a) \models g^{\otimes n}|C$. So $f(X)$ still hits the generic type of k^n.

Write f as f_w. We claim that $c \in Z_{f,w}$. Let s realize $g^{\otimes n}|\mathbb{M}$, outside the monster. Suppose $a \in f^{-1}(s)$. Then $\Sigma(a)$ holds, by definition of f. Also, $f(a) = f_0(a)$ realizes $g^{\otimes n}$. By the claim, the ϕ-type of a over $\mathbb{M}$ is the restriction of p to a ϕ-type. That is, for every $b \in \mathbb{M}$, $\phi(a;b)$ holds if and only if $\psi(b;c)$ holds.

So we have shown that

$$\forall y \in \mathbb{M}\, \forall x \in f^{-1}(s) \quad (\phi(x;y) \leftrightarrow \psi(y;c))$$

Since $\mathrm{tp}(s/\mathbb{M}) = g^{\otimes n}$, this is equivalent to saying

$$\forall y \in \mathbb{M}\, d_{g^{\otimes n}} s\, \forall x \in f^{-1}(s) \quad (\phi(x;y) \leftrightarrow \psi(y;c))$$

This means that $c \in Z_{f,w}$, by definition of $Z_{f,w}$.

Conversely, suppose that $c \in Z_{f,w}$ for some f and $w \in W_f$. Then $g^{\otimes n}$ is in $f(X)$, hence is an element of the stable completion of $f(X)$. By the previous section, there is a generically stable type p in X such that $f_* p = g^{\otimes n}$. Let C be a set over which everything so far is defined.

We claim that $\psi(y;c)$ is the ϕ-definition of p. Let b be arbitrary, and let a realize $p|bC$. Then $f(a)$ realizes $g^{\otimes n}|bC$. As $c \in Z_{f,w}$, we know that

$$\models d_{g^{\otimes n}} s\, \forall x \in f^{-1}(s) \quad \phi(x,b) \leftrightarrow \psi(b;c).$$

Everything inside the $d_{g^{\otimes n}} s$ is bC-definable, and $f(a)$ realizes $g^{\otimes n}|bC$, so we can take $s = f(a)$, yielding

$$\forall x \in f^{-1}(f(a)) \quad \phi(x,b) \leftrightarrow \psi(b;c)$$

In particular, taking $x = a$, we see that $\phi(a,b) \iff \psi(b;c)$. As b was arbitrary, $\psi(-;c)$ is the ϕ-definition of p. □

Chapter 8

On o-minimal imaginaries

In this chapter, we will change notation slightly, using ∂X to denote the frontier $\overline{X} \setminus X$ of a set, and $\mathrm{bd}(X)$ to denote the boundary $\overline{X} \setminus X^{\mathrm{int}}$.

8.1 Elimination of imaginaries and o-minimality

In o-minimal expansions of real closed fields, as well as many other o-minimal theories, elimination of imaginaries holds as a corollary of definable choice. As noted in [59], some o-minimal theories fail to eliminate imaginaries. For example, elimination of imaginaries fails in the theory of $\mathbb{Q}$ with the ordering and with a 4-ary predicate for the relation $x-y = z-w$. In [62], Eleftheriou, Peterzil, and Ramakrishnan observe that in this example, elimination of imaginaries holds after naming two parameters. This leads them to pose the following question:

Question 8.1.1. *Given an o-minimal structure M and a definable equivalence relation E on a definable set X, both definable over a parameter set A, is there a definable map which eliminates X/E, possibly over $B \supseteq A$?*

They answer this question in the affirmative when X/E has a definable group structure, as well as when $\dim(X/E) = 1$. However, we will answer Question 8.1.1 negatively by giving a counterexample in §8.2. That is, we will give an o-minimal structure M and a set X/E interpretable in M, which cannot be put in definable bijection with a definable subset of M^k.

Question 8.1.1 can be reformulated in several ways, by the following observation.

Lemma 8.1.2. *Let M be a structure, and let $\mathbb{M} \succeq M$ be any elementary extension, such as a monster model. The following are equivalent:*

(a) *Every M-definable quotient can be eliminated over M.*

(b) *Every M-definable quotient can be eliminated over $\mathbb{M}$.*

(c) *Every $\mathbb{M}$-definable quotient can be eliminated over $\mathbb{M}$.*

(d) *The elementary diagram of M eliminates imaginaries.*

Proof. The implications $(a) \Rightarrow (d) \Rightarrow (c) \Rightarrow (b)$ are more or less clear. For $(b) \Rightarrow (a)$, suppose (b) holds and X/E is an M-definable quotient. By (b), X/E can be eliminated by an $\mathbb{M}$-definable function f. Since M is an elementary substructure of $\mathbb{M}$, the parameters used to define f can be moved into M, so (a) holds. □

Question 8.1.1 asks whether the equivalent conditions of Remark 8.1.2 hold in every o-minimal structure M. We will give an example in which they fail.

In a talk at the 2012 Banff meeting on Neo-Stability, Peterzil asked the following variant of Question 8.1.1:

Question 8.1.3. *Given an o-minimal structure M and an imaginary $e \in M^{eq}$, is there a set $A \subseteq M$ and a real tuple $c \in M^k$ such that $A \mathop{\smile\!\!\!\!\!\!|}^{\flat} e$ and $\mathrm{dcl}^{eq}(Ae) = \mathrm{dcl}^{eq}(Ac)$?*

Here $\mathop{\smile\!\!\!\!\!\!|}^{\flat}$ denotes thorn-non-forking, or equivalently, independence with respect to o-minimal dimension.

In contrast to the negative answer to Question 8.1.1, we anser Question 8.1.3 positively in §8.3. In some sense, this suggests that interpretable sets, while not being "globally" definable, look "locally" like definable sets. We prove this in §8.5.

8.2 The failure

Let $\mathbb{RP}^1 = \mathbb{R} \cup \{\infty\}$ be the real projective line. The group $PSL_2(\mathbb{R})$ acts on $\mathbb{RP}^1$ by fractional linear transformations, $x \mapsto \frac{ax+b}{cx+d}$, and the stabilizer of ∞ is exactly the group of affine transformations $x \mapsto ax + b$.

For $x, y_1, y_2, y_3, y_4 \in \mathbb{RP}^1$, let $P_0(x, y_1, y_2, y_3, y_4)$ indicate that $x \notin \{y_1, y_2, y_3, y_4\}$ and that

$$f(y_1) - f(y_2) = f(y_3) - f(y_4)$$

for any/every fractional linear transformation f sending x to ∞. The choice of f does not matter, because if f and f' both send x to ∞, then $f' = h \circ f$ for some affine transformation h. But in general,

$$h(z_1) - h(z_2) = h(z_3) - h(z_4) \iff z_1 - z_2 = z_3 - z_4$$

for h affine.

Remark 8.2.1. *If g is some fractional linear transformation, then g induces an automorphism on the structure $(\mathbb{RP}^1, P_0)$. In particular, if $a > 0$ and $b \in \mathbb{R}$, then the map $x \mapsto ax + b$ (fixing ∞) is an automorphism.*

Remark 8.2.2. *Write* $\cot(x)$ *for* $1/\tan(x)$*. If* $\alpha \in \mathbb{R}$*, then* $-\cot(x)$ *and* $\cot(x-\alpha)$ *are related by a fractional linear transformation not depending on* x*, sending* $-\cot(\alpha)$ *to* $\cot(0) = \infty$*. Consequently, if* $\alpha, x_1, \ldots, x_4 \in \mathbb{R}$*, then*

$$\begin{aligned} &P_0(-\cot(\alpha), -\cot(x_1), \ldots, -\cot(x_4)) \\ &\iff \cot(x_1 - \alpha) - \cot(x_2 - \alpha) = \cot(x_3 - \alpha) - \cot(x_4 - \alpha). \end{aligned}$$

Let M be the structure $(\mathbb{Z} \times \mathbb{RP}^1, <, \sigma, P)$, where

- $<$ is the lexicographic order on $\mathbb{Z} \times \mathbb{RP}^1$, where we order $\mathbb{RP}^1$ by identifying it with $[-\infty, +\infty)$.
- σ is the map $(n, x) \mapsto (n+1, x)$.
- $P(x, y_1, \ldots, y_4)$ holds if and only if

$$P_0(\pi_2(x), \pi_2(y_1), \ldots, \pi_2(y_4)) \wedge \bigwedge_{i=1}^{4} x < y_i < \sigma(x)$$

where $\pi_2 : \mathbb{Z} \times \mathbb{RP}^1 \to \mathbb{RP}^1$ is the second coordinate projection.

Remark 8.2.3. *If* $a > 0$ *and* $b \in \mathbb{R}$*, then the map* $(n, x) \mapsto (n, ax + b)$*, fixing* (n, ∞)*, is an automorphism of* M*. This uses Remark 8.2.1*

Let N be the structure $(\mathbb{R}, <, \sigma', P')$, where $<$ is the usual order on $\mathbb{R}$, $\sigma'(x) = x + \pi$, and $P'(x, y_1, \ldots, y_4)$ holds if and only if $x < y_i < x + \pi$ for each i and

$$\cot(y_1 - x) - \cot(y_2 - x) = \cot(y_3 - x) - \cot(y_4 - x).$$

Remark 8.2.4. *The structure* N *is isomorphic to the structure* M *via the map sending* x *to* $(\lfloor x/\pi \rfloor, -\cot(x))$*, using Remark 8.2.2.*

Remark 8.2.5. *For every* $\alpha \in \mathbb{R}$*, the map* $x \mapsto x + \alpha$ *is an automorphism of* N*. Consequently, the automorphism group of* N *acts transitively on* N *and the same is true for* M*.*

One thinks of the structure M as being the "universal cover" of $(\mathbb{RP}^1, P_0)$. We will show that M is o-minimal and fails condition (d) of Lemma 8.1.2.

8.2.1 O-minimality

Consider the two-sorted structure $(\mathbb{R}, \mathbb{Z}, \ldots)$ with the ring structure on $\mathbb{R}$ and the order on $\mathbb{Z}$. The inclusion $\mathbb{Z} \hookrightarrow \mathbb{R}$ is *not* definable in this structure; the two sorts $\mathbb{R}$ and $\mathbb{Z}$ have nothing to do with each other.

Remark 8.2.6. *The structure M can be interpreted in $(\mathbb{R}, \mathbb{Z}, \ldots)$, by mapping $(n, x) \in \mathbb{Z} \times \mathbb{R}$ to (n, x) and (n, ∞) to $n \in \mathbb{Z}$.*

We draw two consequences from this:

Lemma 8.2.7. *Let D be a definable subset of M^1, and suppose $a, b \in M$. Then $D \cap [a, b]$ is a finite union of points and intervals.*

Proof. In the structure $(\mathbb{R}, \mathbb{Z}, \ldots)$, the set $\mathbb{R}$ is o-minimal. Under the interpretation of M in $(\mathbb{R}, \mathbb{Z}, \ldots)$, each open interval of the form $\{n\} \times \mathbb{R} \subseteq \mathbb{Z} \times \mathbb{RP}^1 = M$ is in definable bijection with $\mathbb{R}$. Consequently, $D \cap (\{n\} \times \mathbb{R})$ is a finite union of points and intervals. More generally, each interval $[a, b] \subseteq M$ is contained in a finite union of points and open intervals of the form $\{n\} \times \mathbb{R}$, so the conclusion holds. □

Corollary 8.2.8. *Let $M + M$ be the structure obtained by laying two copies of M end-to-end. More precisely, $M + M$ is the structure $(2 \times M, <, \sigma, P)$, where*

- *$(2 \times M, <)$ is $\{1, 2\} \times M$ with the lexicographic ordering.*
- *$\sigma(i, x) = (i, \sigma(x))$ for $i = 1, 2$.*
- *$P((i_1, x_1), \ldots, (i_5, x_5))$ agrees with $P(x_1, \ldots, x_5)$ when $i_1 = i_2 = \cdots = i_5$, and is false otherwise.*

Then the two inclusion maps $\iota_1, \iota_2 : M \to M + M$ are elementary embeddings.

Proof. The two canonical inclusion maps of the ordered set $\mathbb{Z}$ into the ordered set $\mathbb{Z} + \mathbb{Z} := 2 \times \mathbb{Z}$ are both elementary embeddings. This is an easy exercise using quantifier elimination in $(\mathbb{Z}, <, \sigma)$, where $\sigma(n) = n + 1$. From this, it follows that the two canonical inclusions

$$(\mathbb{R}, \mathbb{Z}, \ldots) \hookrightarrow (\mathbb{R}, \mathbb{Z} + \mathbb{Z}, \ldots)$$

are elementary embeddings. Applying the interpretation of M in $(\mathbb{R}, \mathbb{Z}, \ldots)$ to both sides yields the desired result. □

Remark 8.2.9. *If τ_1, τ_2 are two automorpisms of M, then the map on $M + M$ which acts as τ_1 on the first copy and τ_2 on the second copy is an automorphism.*

Remark 8.2.10. *The group of automorphisms of $M + M$ which fix the first copy pointwise acts transitively on the second copy. This follows from Remarks 8.2.5 and 8.2.9. As a consequence, if D is a one-dimensional definable subset, defined over the first copy, then D or its complement contains the second copy.*

Theorem 8.2.11. *The structure M is o-minimal.*

Proof. The structure N is interpretable in $\mathbb{R}$ with the ring structure and with the trigonometric functions restricted to the interval $[0, \pi]$. This is known to be o-minimal as a consequence of Gabrielov's theorem (or the associated quantifier elimination result; see Theorem 4.6 in [14]).

Alternatively, here is a more elementary argument:

- $M \cong N$ has the order type of $\mathbb{R}$, so it suffices to show that if $D \subseteq M$ is definable, then the boundary $\mathrm{bd}(D)$ does not accumulate at any points in the extended line $\{-\infty\} \cup M \cup \{+\infty\}$.
- Lemma 8.2.7 shows that $\mathrm{bd}(D)$ cannot accumulate at any points in M.
- Suppose $\mathrm{bd}(D)$ had an accumulation point at $+\infty$. Then $\mathrm{bd}(D)$ is not bounded above. This remains true in the elementary extension $M + M$, where we identify the original M with the first copy in $M + M$. But by Remark 8.2.10, D or its complement contains the second copy, making $\mathrm{bd}(D)$ disjoint from the second copy. Then $\mathrm{bd}(D)$ is bounded above by any element from the second copy, a contradiction.
- A similar argument shows that $\mathrm{bd}(D)$ has no accumulation point at $-\infty$.

□

8.2.2 Failure of elimination of imaginaries

Consider the structure $M + M$ from Corollary 8.2.8. Call the two copies M_1 and M_2. Each is isomorphic to M, and each is an elementary substructure of $M_1 + M_2$. We will show that condition (d) of Lemma 8.1.2 fails in M_2. Suppose for the sake of contradiction that the elementary diagram of M_2 eliminates imaginaries.

Let $\alpha = (0, \infty) \in M_1$. Let X be the α-definable set

$$X = \{(x, y) : \alpha < x < y < \sigma(\alpha)\}$$

We can identify the open interval $(\alpha, \sigma(\alpha)) = \{0\} \times \mathbb{R}$ with $\mathbb{R}$. Then X is identified with $\{(x, y) \in \mathbb{R}^2 : x < y\}$. Let $\sim$ be the relation on X

$$(x, y) \sim (x', y') \iff P(\alpha, x, y, x', y').$$

Under the identification of the open interval $(\alpha, \sigma(\alpha))$ with $\mathbb{R}$, we have

$$(x, y) \sim (x', y') \iff P_0(\infty, x, y, x', y') \iff x - y = x' - y'. \tag{8.1}$$

Thus $\sim$ is an equivalence relation on X.

By Remarks 8.2.3 and 8.2.9, for each $a > 0$ and $b \in \mathbb{R}$, there is an automorphism $\tau_{a,b}$ of the structure $M_1 + M_2$ which sends (n, x) to $(n, ax + b)$ on M_1, and which fixes M_2 pointwise.

Note that $\tau_{a,b}$ fixes α, and therefore acts on the α-definable quotient $X/\sim$. Identifying X with $\{(x,y) \in \mathbb{R}^2 : x < y\}$, we see that

$$\tau_{a,b}(x,y) \sim (x,y) \iff \tau_{a,b}(x) - \tau_{a,b}(y) = x - y \iff ax - ay = x - y \iff a = 1.$$

So if $a = 1$, then $\tau_{a,b}$ acts trivially on $X/\sim$, and otherwise, $\tau_{a,b}$ has no fixed points.

Let c be any element of $X/\sim$. Under the assumption that the elementary diagram of M_2 eliminates imaginaries, c is interdefinable over $M_2\alpha$ with some subset $S \subseteq M_1$. Note that $\tau_{a,b}$ fixes $M_2\alpha$ pointwise, so $\tau_{a,b}$ fixes c if and only if it fixes S pointwise. In particular $\tau_{1,1}$ fixes c and $\tau_{2,0}$ does not, so S must be fixed pointwise by $\tau_{1,1}$, but not by $\tau_{2,0}$. This is impossible, however, since the action of $\tau_{1,1}$ on M_1 is $(n,x) \mapsto (n, x+1)$. The only fixed points are of the form (n,∞), and these are also fixed by $\tau_{2,0}$, the map sending $(n,x) \mapsto (n,2x)$. So if S is fixed pointwise by $\tau_{1,1}$, it is also fixed pointwise by $\tau_{2,0}$, a contradiction.

So the equivalent conditions of Lemma 8.1.2 fail in the o-minimal structure M.

Remark 8.2.12. *The quotient $X/\sim$ described above can be eliminated by naming parameters from M_1. This quotient is a counterexample to (d) of Lemma 8.1.2, rather than to (a). Tracing through Lemma 8.1.2, the actual quotient in M which cannot be eliminated is $Y/\approx$, where $Y \subseteq M^3$ is the set of (a,b,c) such that $a < b < c < \sigma(a)$, and where*

$$(a,b,c) \approx (a',b',c') \iff a = a' \wedge P(a,b,c,b',c').$$

8.3 The local question

Unlike Question 8.1.1, Question 8.1.3 has an easy affirmative answer.

Lemma 8.3.1. *Given an o-minimal structure M and an imaginary $e \in M^{eq}$, there is a set $A \subseteq M$ and a real tuple $c \in M^k$ such that $A \mathop{\smile\!\!\!\!\!\!|}^{b} e$ and $\mathrm{dcl}^{eq}(Ae) = \mathrm{dcl}^{eq}(Ac)$.*

Proof. Suppose e is a class of the definable equivalence relation E. Let x be some representative of this class. So x is a (real) tuple, and $e \in \mathrm{dcl}^{eq}(x)$. Consider the pregeometry on M coming from definability over e, i.e., the pregeometry where the closure of a set $S \subseteq M$ is $M \cap \mathrm{dcl}^{eq}(Se)$. Let $A \subseteq x$ be a basis for x, and let c be the remaining coordinates of x. Then $x = Ac$, and so

$$e \in \mathrm{dcl}^{eq}(x) = \mathrm{dcl}^{eq}(Ac).$$

Also, since A is a basis for x, $c \subseteq x$ is in the closure of A:

$$c \in \mathrm{dcl}^{eq}(Ae).$$

Finally, note that $\dim(x/e) = \dim(A/e) = |A|$ because A is a basis over e. Since A has size $|A|$ and is made of singletons,

$$\dim(A/\emptyset) \leq |A|.$$

On the other hand

$$\dim(A/\emptyset) \geq \dim(A/e) = |A|$$

on general grounds. So $\dim(A/\emptyset) = \dim(A/e)$, which implies $A \mathop{\smile\!\!\!\!\!\!|}^{\mathrm{b}} e$. □

The affirmative answer to Question 8.1.3 does not imply an affirmative answer to Question 8.1.1. The implication fails because the auxiliary parameters A in Lemma 8.3.1 depend too strongly on e. Lemma 8.3.1 can be vaguely interpreted as saying that interpretable sets look "locally" like definable sets.

We explicate this idea in the next two sections.

8.4 Remarks on definable topologies

If X is a subset of a topological space, $\overline{X}$ will mean the closure of X, and ∂X will mean the frontier of X, $\overline{X} \setminus X$. The interior of X will be denoted X^{int}. The boundary $\overline{X} \setminus X^{\mathrm{int}}$ will be denoted $\mathrm{bd}(X)$.

8.4.1 Definable Topologies

Let M be a structure in some language. Assume M has elimination of imaginaries. Let X be a definable set. Definable will mean "definable with parameters." By a *definable topology*, we mean a definable family of subsets $\{B_y \subseteq X\}_{y \in Y}$ which form the basis for some topology on X. The fact that these form a basis for a topology amounts to the claim that if y_1, y_2 have $B_{y_1} \cap B_{y_2} \neq \emptyset$, then for every $x \in B_{y_1} \cap B_{y_2}$ there is a $y_3 \in Y$ such that $x \in B_{y_3} \subseteq B_{y_1} \cap B_{y_2}$. This is a first-order condition, so a definable topology on X remains a definable topology in elementary extensions of M. But note that if $M \preceq M'$, the topology on $X(M)$ need not agree with the subspace topology on $X(M)$ as a subset of $X(M')$.

If D is a definable subset of X, and X has a definable topology, then the subspace topology on D is also definable. If X and Y are two sets with definable topologies, then the product topology on $X \times Y$ is definable.

If X and Y are two sets with definable topologies, and $f : X \to Y$ is a definable function, then we can express whether or not f is continuous using some first-order statement. So the continuity of f is invariant under elementary extensions, and definable in families.

We say that X is *definably connected* if there is no definable clopen set D with $\emptyset \subsetneq D \subsetneq X$. The definable connectedness of X is invariant under elementary extensions, and type-definable in families. If $f : X \to Y$ is a continuous definable function, and X is definably connected, then so is Y.

A continuous map $f : X \to Y$ between abstract topological spaces is an open map if $f(U)$ is open for every open subset $U \subseteq X$. If $\mathcal{B}$ is a basis of opens on X, it suffices to check $U \in \mathcal{B}$. If $f : X \to Y$ is a continuous definable map between two definable topological spaces, then we can express that f is an open map via a first-order statement. So elementary extensions preserve whether or not f is open, and this is definable in families.

8.4.2 The Quotient Topology

If X is an abstract topological space, and E is an equivalence relation on X, then there is a natural quotient topology on X/E. Letting $\pi : X \to X/E$ be the natural projection, a subset $U \subseteq X/E$ is open in this quotient topology if and only if $\pi^{-1}(U)$ is open. Note that π is then continuous.

A continuous surjection $f : X \to Y$ of abstract topological spaces is *identifying* if Y has the quotient topology of $X/\ker f$, where $\ker f$ is the equivalence relation $\{(x, y) : f(x) = f(y)\}$. In other words, a subset $U \subseteq Y$ is open if and only if $f^{-1}(U)$ is open.

If X is a definable topological space and E is a definable equivalence relation, the quotient topology on X/E need not be definable, as far as I know. If X and Y have definable topologies, and $f : X \to Y$ is a definable continuous function, there is not any general way of expressing that f is identifying.

Remark 8.4.1. *If $f : X \to Y$ is a continuous open surjection of abstract topological spaces, then f is identifying. Indeed, if $f^{-1}(U)$ is open for some subset $U \subseteq Y$, then $f(f^{-1}(U)) = U$ is also open. Moreover, if X_0 is an open subset of X, then the restriction of f to X_0 is a continuous open surjection to $f(X_0)$, where $f(X_0)$ has the subspace topology from Y. In particular, the quotient topology on $f(X_0)$ as a quotient of X_0 agrees with the subspace topology on $f(X)$, as a subspace of $f(X)$.*

In a definable setting, we *can* express that f is a continuous open surjection.

Definition 8.4.2. *An equivalence relation E on a set X is an* open equivalence relation *if the natural quotient map $\pi : X \to X/E$ (with the quotient topology on X/E) is an open map. Equivalently, for every open set $U \subseteq X$, the set*

$$\pi^{-1}(\pi(U)) = \{x \in X : xEu \text{ for some } u \in U\}$$

is open. If $\mathcal{B}$ is a basis of open sets on X, it suffices to check the $U \in \mathcal{B}$.

Remark 8.4.3.

(a) *If E is an open equivalence relation on X, and $U \subseteq X$ is open, then $E \restriction U = E \cap (U \times U)$ is an open equivalence relation on U, by Remark 8.4.1.*

(b) *If X is a definable topological space and E is a definable equivalence relation, then we can write a first-order sentence that expresses that E is an open equivalence relation on X. (We only need to check that $\pi^{-1}(\pi(U))$ is open for U in a basis of opens on X.) In particular, this property is definable in families and invariant under elementary extensions.*

(c) *If X is a definable topological space and E is an open definable equivalence relation on X, then the quotient topology on X/E is definable. In fact, if $\pi : X \to X/E$ denotes the natural surjection, and $\mathcal{B}$ is a definable basis of opens for X, then $\{\pi(U) : U \in \mathcal{B}\}$ is a definable basis of opens for X/E.*

8.4.3 Separation Axioms

Recall that a topological space X is *Hausdorff* if for every two distinct points $x, x' \in X$, we can find two disjoint open sets U, U' with $x \in U$, $x' \in U'$. Given a basis, we can require that U and U' be basic opens. Recall that X is T_0 if for every two distinct points $x, x' \in X$, we can find an open set U which contains exactly one of $\{x, x'\}$. Again, we can require that U be a basic open.

Recall that X is *regular* if for every $x \in X$ and every closed set $C \subseteq X$ with $x \notin C$, we can find disjoint open sets U and V with $x \in U$ and $C \subseteq V$. Equivalently, whenever D is open and $x \in D$, we can find a smaller open neighborhood $U \ni x$ with $\overline{U} \subseteq D$. In this second definition, we may assume that D and U are basic opens.

Because it suffices to check basic opens, if X is a definable topological space, then we can therefore express by a first-order statement that the topology is Hausdorff, T_0, or regular. Such properties are therefore preserved in elementary extensions and definable in families.

Note that Hausdorff implies T_0. Also, regular plus T_0 implies Hausdorff. Indeed, suppose X is regular and T_0. If x and x' are two distinct points in X, then by T_0 we can find a closed set C containing exactly one of x and x'. Now use regularity to separate C from the unique point of $\{x, x'\} \setminus C$.

If X is Hausdorff or T_0 or regular, and Y is a subset of X, then the subspace topology on Y has the same property.

8.4.4 Definable Compactness

If X is a set, a *filtered collection of subsets of* X is a collection $\mathcal{F}$ of subsets of X with the property that for every $D_1, D_2 \in \mathcal{F}$, there is some $D_3 \in \mathcal{F}$ with $D_3 \subseteq D_1 \cap D_2$. For example, if X is a topological space, $\mathcal{B}$ is a basis, and $x \in X$, then the set of basic opens neighborhoods of x is a filtered collection of non-empty subsets of X.

If X is a definable set and $\mathcal{F}$ is a definable family of subsets of X, we can express by a first-order statement that $\mathcal{F}$ is a filtered collection.

If X is a definable topological space, we shall say that X is *definably compact* if every definable filtered collection of closed non-empty subsets of X has non-empty intersection.[1] This property is preserved in elementary extensions, and type-definable in families. If D is a closed subset of a definably compact topological space, then D with the subspace topology is definably compact as well.

Lemma 8.4.4. *Let X and Y be definable topologies, and suppose X is definably compact. If $D \subseteq X \times Y$ is definable and closed, then $\pi(D)$ is closed as a subset of Y, where $\pi : X \times Y \to Y$ is the projection.*

Proof. Suppose that $y \in \overline{\pi(D)}$. Let $\mathcal{B}$ be the (definable) collection of basic open sets containing y. For each $V \in \mathcal{B}$, let C_V be the set of $x \in X$ for which there is some open neighborhood

[1] This notion has been independently studied by Fornasiero in unpublished work [21].

$U \ni x$ with $(U \times V) \cap D = \emptyset$. The family of the C_V's is definable. Let $X_V = X \setminus C_V$. Each C_V is open, so each X_V is closed. If some X_V is empty, then $C_V = X$. In particular, for every $x \in X$ we have $(\{x\} \times V) \cap D = \emptyset$, so $(X \times V) \cap D = \emptyset$ and V is disjoint from $\pi(D)$. But then $y \in V$ so $y \notin \overline{\pi(D)}$, a contradiction. Therefore each X_V is non-empty.

Note that if $V' \subseteq V$, then $C_V \subseteq C_{V'}$, so $X_{V'} \subseteq X_V$. Therefore, $\{X_V\}_{V \in \mathcal{B}}$ is a filtered collection of non-empty closed subsets of X. So there is some $x \in \bigcap_{V \in \mathcal{B}} X_V$. Now for every basic open neighborhood $U \times V \ni (x, y)$, if $(U \times V) \cap D = \emptyset$, then $x \in C_V$, a contradiction. So every basic open neighborhood of (x, y) intersects D. As D is closed, $(x, y) \in D$. Thus $x \in \pi(D)$. □

Lemma 8.4.5. *If X and Y are definably compact, then so is $X \times Y$.*

Proof. Let $\mathcal{F}$ be a definable filtered collection of non-empty closed subsets of $X \times Y$. Let $\pi : X \times Y \to Y$ be the projection. By Lemma 8.4.4, $\pi(F)$ is closed for every $F \in \mathcal{F}$. In particular, $\{\pi(F) | F \in \mathcal{F}\}$ is a definable filtered collection of non-empty closed subsets of Y. So there is some point y with $y \in \pi(F)$ for every $F \in \mathcal{F}$. Now for each $F \in \mathcal{F}$, let $G_F = \{x \in X : (x, y) \in F\}$. Then G_F is non-empty and closed for every $F \in \mathcal{F}$, and the collection $\mathcal{G} := \{G_F\}_{F \in \mathcal{F}}$ is a definable filtered collection of non-empty closed subsets. So there is some x in $\bigcap_{F \in \mathcal{F}} G_F$. Then $(x, y) \in \bigcap \mathcal{F}$. □

Lemma 8.4.6. *If $f : X \to Y$ is a continuous definable surjection of definable topological spaces, and X is definably compact, then so is Y.*

Proof. Let $\mathcal{F}$ be a definable filtered collection of closed non-empty subsets of Y. Let $f^*\mathcal{F}$ be the collection

$$\{f^{-1}(F) : F \in \mathcal{F}\}.$$

Then $f^*\mathcal{F}$ is a definable filtered collection of closed non-empty subsets of X. So there is some $x \in \bigcap_{F \in \mathcal{F}} f^{-1}(F)$. Equivalently, $f(x) \in \bigcap \mathcal{F}$. □

Lemma 8.4.7. *Let X be a Hausdorff definable topological space, and let D be a definable subset of X. Suppose that D with the subspace topology is definably compact. If $x \notin D$, then there is an open neighborhood U of x with $\overline{U}$ disjoint from D.*

Proof. Let $\mathcal{B}$ be a definable basis for the topology on X. Let $\mathcal{F}$ be the collection of sets $\{\overline{U} \cap D : x \in U \in \mathcal{B}\}$. This is a filtered definable collection of closed subsets of D.

Note that if $y \in D$, then $y \neq x$ and so by Hausdorffness we can find $x \in U$ and $y \in V$ with U and V basic opens, and $U \cap V = \emptyset$. This means that $y \notin \overline{U}$. So every point $y \in D$ is not in some $\overline{U}$, meaning that $\bigcap \mathcal{F} = \emptyset$. By definable compactness of D, some element of $\mathcal{F}$ must be empty. Therefore there is some open $U \ni x$ with $\overline{U} \cap D = \emptyset$. □

Corollary 8.4.8. *If X is a Hausdorff definable topological space, and D is a subset of X which is definably compact (with the subspace topology), then D is closed.*

Corollary 8.4.9. *If X is a Hausdorff definable topological space, and X is definably compact, then X is regular.*

Proof. If C is a closed subset of X and $x \in X \setminus C$, then C is compact (as we noted above), so by Lemma 8.4.7, there is an open neighborhood $U \ni x$ with $\overline{U}$ disjoint from C. □

8.4.5 The o-minimal Setting

Now restrict to the setting of M^{eq}, where $M = (M, <, \ldots)$ is an o-minimal structure. All o-minimal structures will be dense, i.e., expand DLO. A definable topological space X in M^{eq} will be said to be locally Euclidean if every $x \in X$ has a definable neighborhood U which is definably homeomorphic to an open subset of M^n. The dimension n might vary with x. The property of being locally Euclidean is preserved downwards but perhaps not upwards in elementary extensions.

Lemma 8.4.10. *Any closed interval $[a, b] \subseteq M$ is definably compact (in the sense defined above).*

Proof. Let $\mathcal{F}$ be a definable filtered collection of non-empty closed subsets of $[a, b]$. For $F \in \mathcal{F}$, let x_F be $\inf F$. Note $x_F \in F$, by closedness of F. The set $D = \{x_F : F \in \mathcal{F}\}$ is a definable subset of $[a, b]$; let y be $\sup D$. I claim that $y \in \bigcap \mathcal{F}$. If not, then $y \notin F_0$ for some $F_0 \in \mathcal{F}$. As F_0 is closed, there is some $z < y$ such that $[z, y] \cap F_0$ is empty. Now as $y = \sup D$, there must be some $x_{F_1} \in [z, y]$. Take $F_2 \subseteq F_1 \cap F_0$. Because $F_2 \subseteq F_1$, we have $x_{F_2} = \inf F_2 \geq \inf F_1 = x_{F_1} \geq z$. So $z \leq x_{F_2} \leq y$, and $x_{F_2} \in F_2 \subseteq F_0$. Therefore, F_0 does intersect $[z, y]$, at at least the point x_{F_2}. So we have a contradiction. □

Corollary 8.4.11. *Closed and bounded subsets of M^n are definably compact (in the sense defined above).*

Proof. Use Lemma 8.4.5 □

The converse is also true: if a subset X is definably compact, then it is closed and bounded. It is closed by Corollary 8.4.8. If it is unbounded, then $\pi(X) \subseteq M^1$ is unbounded, for some coordinate projection $\pi : M^n \to M^1$. But then $\pi(X)$ is definably compact by Lemma 8.4.6. However, an unbounded subset of M^1 is not definably compact because of the filtered collection of closed non-empty subsets obtained by intersecting with $[a, \infty)$ or $(-\infty, a]$ for $a \in M^1$. So the notion of definably compact that we are using agrees with the usual notion in o-minimal structures.

Remark 8.4.12. *If $X \subseteq M^n$ is definable, then the frontier $\partial X := \overline{X} \setminus X$ always has lower dimension than X. Also, if $X \subseteq Y \subseteq M^n$, then the relative boundary $bd_Y(X)$ of X as a subset of the topological space Y, always has* $\dim bd_Y(X) < \dim Y$.

Lemma 8.4.13. *Let X be a definable topological space in M^{eq}. Suppose that X is Hausdorff and locally Euclidean. Then M is regular.*

Proof. Let C be a closed subset of X and $x \in X \setminus C$. Take some open neighborhood U of x which is homeomorphic to an open subset of M^n. Shrinking U, we may assume that $U \cap C = \emptyset$. Suppose that U is homeomorphic to an open subset $V \subseteq M^n$ via a definable homeomorphism f. As $f(x)$ is in the interior of V, we can find a closed box B with $x \in B^{\text{int}}$ and $B \subseteq V$. Now $f^{-1}(B)$ is definably compact, so it is closed as a subset of X, by Corollary 8.4.8. Also, $f^{-1}(B^{\text{int}})$ is an open neighborhood of x. Its closure is contained in $f^{-1}(B)$ which is contained in U. In particular, $f^{-1}(B)$ is an open neighborhood of x whose closure is disjoint from C. □

Definition 8.4.14. *A* definable manifold *is a Hausdorff definable topological space X which is locally Euclidean in all elementary extensions of the model.*

This definition forces some uniformity in the manifold charts. It suffices to check the local Euclideanity in some saturated elementary extension of the original model. Moreover, because the property of being locally Euclidean uniformly must be witnessed by a definable family of charts, the property of being locally Euclidean is expressible as a small disjunction of first-order statements, and is consequently ind-definable in families.

Note by Lemma 8.4.13 that definable manifolds are always regular.

8.5 Quotients in o-minimal structures

Recall that if M is o-minimal, then M^{eq} is a supperrosy structure of finite definable rank. We will use $\dim X$ to refer to the rank of a definable set. We will use $\dim(a/b)$ to denote the rank of a over b.

8.5.1 Statement of the Theorem

Fix some o-minimal structure M.

Theorem 8.5.1. *Let $Y \subseteq M^n$ be a definable set. Let E be a definable equivalence relation on Y. Then there is some definable open subset Y' of Y such that $\dim(Y \setminus Y') < \dim Y$, and such that the natural quotient topology on Y'/E is definable, Hausdorff, regular, and locally Euclidean, and the map $Y' \to Y'/E$ is an open map. Moreover, we may arrange that these topological properties remain true in all elementary extensions.*

Corollary 8.5.2. *Let X be a definable set in M^{eq}. Then X can be put in definable bijection with a finite disjoint union of definably connected interpretable manifolds.*

Proof. Write X as Y/E, with $Y \subseteq M^n$ for some n. Proceed by induction on $\dim Y$. The base case where $Y = \emptyset$, is trivial. Assume Y is non-empty. By the Theorem, we can find Y' such that Y'/E is a definable manifold. The map $Y' \to Y'/E$ is continuous, and Y' has finitely many definably connected components (by cell-decomposition), so Y'/E also has finitely many definably connected components. Each of these is a definably connected

definable manifold. Letting $\pi : Y \to X$ be the natural quotient map, let $X' = \pi(Y')$. Then we have just put X' in definable bijection with a disjoint union of definably connected definable manifolds. Also, $X \setminus X'$ can be written as a quotient of $\pi^{-1}(X \setminus X')$. But

$$\pi^{-1}(X \setminus X') \subseteq Y \setminus Y'$$

so $\dim(\pi^{-1}(X \setminus X')) < \dim Y$. By induction, $X \setminus X'$ can be put in definable bijection with a finite disjoint union of definably connected interpretable manifolds. Writing X as $X' \coprod (X \setminus X')$, we are done. □

In proving Theorem 8.5.1, we may replace M with a saturated elementary extension—a monster model. The topological properties other than local Euclideanity are all expressible by first-order statements (because of the assumption that $Y' \to Y'/E$ is open), and uniform local Euclideanity is a small disjunction of first-order statements, so if we can find a satisfactory Y' in the monster, we can also find one in the original model M.

Henceforth, assume that M is a monster model. Work in M^{eq}. Hold Y fixed. Let $\pi : Y \to X$ be the projection. So the equivalence class of y is $\pi^{-1}(\pi(y))$, for $y \in Y$. Denote this by $E(y)$, for simplicity.

If D is a definablet set in M^{eq}, $\ulcorner D \urcorner$ will denote the code of D (as an element of M^{eq}). We will frequently use the following fact:

Remark 8.5.3. *If U is an open subset of M^n, and p is a point in U, and S is some small subset of M^{eq}, then we can find an open neighborhood V with $p \in V \subseteq U$ and with*

$$Sp\ulcorner U \urcorner \overset{b}{\downarrow} \ulcorner V \urcorner.$$

Indeed, we can take V to be an open box with generically chosen corners.

More generally, if Y is a definable set defined over some set T, and $p \in Y$ is a point, and U is an open subset of Y containing p, and S is a small subset of M^{eq}, then we can find a neighborhood V of p in Y with $p \in V \subseteq U$ and with $\ulcorner V \urcorner \overset{b}{\underset{T}{\downarrow}} Sp\ulcorner U \urcorner$. Specifically, we can take a random box, sufficiently small, and intersect it with Y.

The proof of Theorem 8.5.1 will proceed in several steps. At each step, we will replace Y with an open subset Y' such that $\dim Y \setminus Y' < \dim Y$, in such a way that certain properties will be true of Y'. These properties will be preserved in each subsequent step.

8.5.2 Step 1: Pure-dimensional Equivalence Classes

Recall that if $Z \subseteq M^n$ is a definable set, then for $z \in Z$, we can discuss the *local dimension* $\dim_z Z$ of Z at z. This is $\dim U \cap Z$ for U a sufficiently small open around z. The value $\dim_z Z$ is definable in families (as Z and z vary in a family).

Moreover, if z is a generic point in Z, then $\dim_z Z = \dim Z$. Indeed, suppose Z is defined over S, and $\dim(z/S) = \dim Z$. Take some U containing z such that $\dim U \cap Z = \dim_z Z$.

By Remark 8.5.3 we can find V independent from z over S, with $z \in V \subseteq U$. Then $\dim V \cap Z = \dim_z Z$. By the independence, $\dim(z/S) = \dim(z/\ulcorner V \urcorner S)$. But z is in the $\ulcorner V \urcorner S$-definable set $V \cap Z$, so $\dim(z/\ulcorner V \urcorner S) \leq \dim V \cap Z = \dim_z Z$. In particular,

$$\dim_z Z \leq \dim Z = \dim(z/S) = \dim(z/\ulcorner V \urcorner S) \leq \dim V \cap Z = \dim_z Z,$$

so $\dim_z Z = \dim Z$.

Say that a set Z is *pure-dimensional* if $\dim_z Z = \dim Z$ for every $z \in Z$, or equivalently, $z \mapsto \dim_z Z$ is a constant function on Z. Equivalently, every non-empty open subset of Z has the same dimension as Z. Note that any open subset of a pure-dimensional set is pure-dimensional.

Lemma 8.5.4. *In the context of Theorem 8.5.1, there is an open subset Y' of Y with $\dim Y \setminus Y' < \dim Y$, such that every equivalence class of $E \restriction Y'$ is pure-dimensional.*

Proof. Let Z be the set of $y \in Y$ such that $\dim_y E(y) = \dim E(y)$. Then Z is a definable subset of Y. Any generic point y in Y is in Z. Indeed, let S be the set over which everything is defined, and suppose $y \in Y$ with $\dim(y/S) = \dim Y$. Let $x = \pi(y)$, so $\pi^{-1}(x)$ is $E(y)$, the equivalence class of y. Then $\dim(y/Sx) = \dim E(y)$; if not then taking $y' \in E(y)$ with $\dim(y'/Sx) = \dim E(y)$, we have $\dim(y'/Sx) > \dim(y/Sx)$, and then

$$\begin{aligned}\dim(y'/S) &= \dim(y'x/S) = \dim(y'/xS) + \dim(x/S) \\ &> \dim(y/xS) + \dim(x/S) = \dim(y/S) = \dim Y,\end{aligned}$$

which is absurd. Now since y is a generic point of $E(y)$, we have $\dim_y E(y) = \dim E(y)$. So $y \in Z$.

Since every generic point of Y is in Z, it follows that $\dim Y \setminus Z < \dim Y$. Let Y' be $Y \setminus \left(\overline{Y \setminus Z}\right)$. Then

$$\dim Y \setminus Y' \leq \dim \overline{Y \setminus Z} \leq \dim Y \setminus Z < \dim Y.$$

And Y' is an open subset of Y. Let E' be the restriction of E to Y'. So if $y \in Y'$, then $E'(y) = E(y) \cap Y'$. Since Y' is open,

$$\dim E'(y) \geq \dim_y E'(y) = \dim_y E(y) = \dim E(y) \geq \dim E'(y)$$

where the second equality holds because $y \in Z$. Thus $\dim_y E'(y) = \dim E'(y)$ for every point $y \in Y'$. So the equivalence classes of E' are pure-dimensional. □

Consequently, replacing Y with Y' and X with $\pi(Y')$, we may assume that every equivalence class of E is pure-dimensional. In subsequent reductions, we will replace Y with smaller open sets. Because open subsets of pure-dimensional sets are still pure-dimensional, we will not lose the pure-dimensionality property.

8.5.3 Step 2: Open quotients

Assume that every equivalence class of E is pure-dimensional.

Let S be a set over which Y and E are defined.

Lemma 8.5.5. *Suppose $y \in Y$ is generic, i.e., $\dim(y/S) = \dim Y$. Let B be an open subset of Y, and suppose $\pi(y) \in \pi(B)$. Then there is an open neighborhood U of y in B with $\pi(U) \subseteq \pi(B)$. We can take U to be a basic open neighborhood (i.e., an intersection of an open box with Y).*

Proof. Since $\pi(y) \in \pi(B)$, there is some $y' \in B$ with $\pi(y) = \pi(y')$. By Remark 8.5.3, take some open neighborhood V of y' such that $V \subseteq B$ and

$$\ulcorner V \urcorner \overset{\text{þ}}{\underset{S}{\downarrow}} \ulcorner B \urcorner y.$$

Now $\pi(y) = \pi(y') \in \pi(V) \subseteq \pi(B)$. Also, $\pi^{-1}(\pi(V))$ is $\ulcorner V \urcorner S$-definable. If Q is $\pi^{-1}(\pi(V)) \subseteq Y$, then the boundary $bd_Y(Q)$ of Q as a subset of Y has dimension less than $\dim Y$, by Remark 8.4.12. In particular, $\dim(y/\ulcorner V \urcorner S) = \dim(y/S) > \dim bd_Y(Q)$, so $y \notin bd_Y(Q)$. As $y \in Q$, we see that some open neighborhood U of y in Y is contained in Q. Thus $y \in U$ and $\pi(U) \subseteq \pi(V) \subseteq \pi(B)$. We can take U to be basic by shrinking it. □

Let Z be the set of all $y \in Y$ with the following property: if B is a basic open set in Y and $\pi(y) \in \pi(B)$, then there is a basic open U containing y, with $\pi(U) \subseteq \pi(B)$. By the Lemma, every generic point of Y is in Z, so $\dim Y \setminus Z < \dim Y$. Let $W = Y \setminus Z$. Let $R \subseteq X$ be those x for which $\dim \pi^{-1}(x) \cap W = \dim \pi^{-1}(x)$. Let W' be $W \cup \pi^{-1}(R)$. Let $Z' = Y \setminus W'$.

Now for each $x \in X$, one of the following holds:

- $x \in R$. Then $\pi^{-1}(x) \subseteq W'$, and $x \notin \pi(Z')$. Also, $\pi^{-1}(x) \cap W' = \pi^{-1}(x)$ has the same dimension as $\pi^{-1}(x) \cap W$, by definition of R. In this case, $x \notin \pi(Z')$.
- $x \notin R$. Then $\pi^{-1}(x)$ is disjoint from $\pi^{-1}(R)$, so $\pi^{-1}(x) \cap W'$ equals $\pi^{-1}(x) \cap W$ and has lower dimension than $\pi^{-1}(x)$.

Either way, $\pi^{-1}(x) \cap W'$ and $\pi^{-1}(x) \cap W$ have the same dimension. Since this holds for every x, $\dim W' = \dim W$. Also, we see from this dichotomy that

$$\text{If } x \in \pi(Z'), \text{ then } \dim \pi^{-1}(x) \cap W' < \dim \pi^{-1}(x). \tag{8.2}$$

Claim 8.5.6. *The restriction of E to Z' is an open equivalence relation.*

Proof. Let X' be $\pi(Z')$ and let $\tau : Z' \to X'$ be the restriction of π to Z'. We need to show that for B a basic open subset of Y, $\tau^{-1}(\tau(B \cap Z'))$ is an open subset of Z'. Suppose $y \in \tau^{-1}(\tau(B \cap Z'))$. Because $y \in Z' \subseteq Z$, there is some open neighborhood $U \ni y$ of y in

Y such that $\pi(U) \subseteq \pi(B)$. It suffices to show that $U \cap Z'$ is in $\tau^{-1}(\tau(B \cap Z'))$, i.e., that $\tau(U \cap Z') = \pi(U \cap Z')$ is a subset of $\tau(B \cap Z') = \pi(B \cap Z')$.

$$\pi(U \cap Z') \stackrel{?}{\subseteq} \pi(B \cap Z').$$

Clearly

$$\pi(U \cap Z') \subseteq \pi(B) \cap \pi(Z'),$$

so it suffices to show that $\pi(B) \cap \pi(Z') \subseteq \pi(B \cap Z')$. Suppose $x \in \pi(B) \cap \pi(Z')$. Because B is open, $B \cap \pi^{-1}(x)$ is an open subset of $\pi^{-1}(x)$. By the previous section, $\pi^{-1}(x)$ is pure-dimensional. So $\dim B \cap \pi^{-1}(x) = \dim \pi^{-1}(x)$. Since $x \in \pi(Z')$, by (8.2), we have that

$$\dim \pi^{-1}(x) \cap W' < \dim \pi^{-1}(x) = \dim B \cap \pi^{-1}(x).$$

So some point $z \in B \cap \pi^{-1}(x)$ is *not* in W', i.e., it *is* in Z'. Then $z \in B \cap Z'$ and $\pi(z) = x$, so $x \in \pi(B \cap Z')$. As x was arbitrary, $\pi(B) \cap \pi(Z') \subseteq \pi(B \cap Z')$. □

Now let Y' be $Y \setminus \overline{W'}$. So

$$\dim Y \setminus Y' \leq \dim \overline{W'} \leq \dim W' \leq \dim W < \dim Y.$$

Also, Y' is an open subset of Y, and $Y' \cap Z' = Y'$ is an open subset of Z'. By Remark 8.4.3(a) applied to $E \restriction Z'$ and $E \restriction Y'$, the restriction of E to Y' is an open equivalence relation on Y'.

Therefore, replacing Y with Y' and X with $\pi(Y')$, we may assume that E is an open equivalence relation. Note that the equivalence classes of E on Y' are open subsets of the equivalence classes on Y, so they remain pure-dimensional.

So we may assume that the equivalence classes of E are pure-dimensional, and that E is an open equivalence relation. By Remark 8.4.3(b), the topology on $Y/E = X$ is now definable. Moreover, if we subsequently replace Y with an open subset Y', then $\pi(Y')$ will have the subspace topology from X, by Remark 8.4.3(c).

8.5.4 Step 3: Separation axioms

Let $F = \bigcup_{x \in X} \partial \pi^{-1}(x)$.

Claim 8.5.7. $\dim F < \dim Y$.

Proof. Let S be a set over which Y, E, F are defined. Take $z \in F$ with $\dim(z/S) = \dim F$. Then $z \in \partial \pi^{-1}(x)$ for some $x \in X$. Let $y \in \pi^{-1}(x)$ have $\dim(y/xS) = \dim \pi^{-1}(x)$. Now

$$\dim F = \dim(z/S) \leq \dim(zx/S) = \dim(z/xS) + \dim(x/S) \leq \dim \partial \pi^{-1}(x) + \dim(x/S)$$

Because the frontier of a set always has lower dimension than the set itself,

$$\begin{aligned}\dim \partial\pi^{-1}(x) + \dim(x/S) &< \dim \pi^{-1}(x) + \dim(x/S)\\ &= \dim(y/xS) + \dim(x/S)\\ &= \dim(xy/S).\end{aligned}$$

But $x = \pi(y) \in \mathrm{dcl}(Sy)$, so

$$\dim(xy/S) = \dim(y/S) \leq \dim Y.$$

Putting everything together, $\dim F < \dim Y$. □

Consequently, $\dim \overline{F} < \dim Y$. Let Y' be $Y \setminus \overline{F}$. For any $y \in Y'$, $E(y) \cap Y'$ is a closed subset of Y': if not, then there is $z \in Y' \cap \overline{E(y)} \setminus E(y) = Y' \cap \partial E(y) \subseteq Y' \cap F = \emptyset$. Replacing Y by Y' and X by $\pi(Y')$, we may therefore assume that the equivalence classes are closed. This preserves the properties obtained above, that the equivalence relation is open and that the equivalence classes are pure-dimensional.

Therefore, we may assume that the equivalence classes are closed (as subsets of Y), on top of the assumptions that the equivalence relation is open and the equivalence classes are pure dimensional. In terms of the quotient topology on X, we have arranged that singletons in X are closed.[2]

The next step is to make the topology on X be Hausdorff. Say that x, x' in X can be separated by neighborhoods if there exist open neighborhoods $V \ni x$, $V' \ni x'$ with $V \cap V' = \emptyset$. Let $H \subseteq X$ consist of those x which can be separated by neighborhoods from every $x' \neq x$. Then H is a definable set.

Claim 8.5.8. *Let S be a set over which Y, E, H are defined. If $y \in Y$ has $\dim(y/S) = \dim Y$, then $\pi(y) \in H$.*

Proof. Let $x = \pi(y)$. Suppose $x' \in X$ is not equal to x. Write x' as $\pi(y')$ for some $y' \in Y$. Then $y' \notin E(y)$. We arranged that $E(y)$ is closed (as a subset of Y). Therefore there is a basic open neighborhood $U' \ni y'$ such that $U' \cap E(y) = \emptyset$. By Remark 8.5.3, we can shrink U' a bit, and arrange that

$$\ulcorner U' \urcorner \overset{\text{þ}}{\underset{S}{\perp}} yy'.$$

Then $\dim(y/\ulcorner U' \urcorner S) = \dim(y/S) = \dim Y$. Let $Q = \pi^{-1}(\pi(U'))$; this is $\ulcorner U' \urcorner S$-definable. As $U' \cap E(y) = \emptyset$, $y \notin Q$. By Remark 8.4.12, $\dim bd_Y(Q) < \dim Y = \dim(y/\ulcorner U' \urcorner S)$, so $y \notin bd_Y(Q)$. Therefore some neighborhood of y is not in Q. Let U be this neighborhood. Then $U \cap \pi^{-1}(\pi(U')) = \emptyset$, so $\pi(U) \cap \pi(U') = \emptyset$. But $\pi(U)$ is an open neighborhood of $\pi(y) = x$, and $\pi(U')$ is an open neighborhood of $\pi(y') = x'$. So we have separated x and x'. As x' was arbitrary, $x \in H$. □

[2] In other words, the topology on X is now T_1.

It follows that $\dim Y \setminus \pi^{-1}(H) < \dim Y$. Let Y' be $Y \setminus (\overline{Y \setminus \pi^{-1}(H)})$. Then Y' is an open subset of Y, and $\dim Y \setminus Y' < \dim Y$. Also, $X' := \pi(Y')$ is a subset of H, and the quotient topology on X' as a quotient of Y' is the subspace topology from X. Since each point of X' can be separated by neighborhoods from any other point in X, it follows that X' is Hausdorff.

Replacing Y with Y' and X with X', we may therefore assume that the quotient topology is Hausdorff, on top of the assumptions that the equivalence classes are pure-dimensional and the equivalence relation is open. If we subsequently replace Y with a smaller open set, all these properties will be preserved.

8.5.5 Step 4: Local Euclideanity

Lemma 8.5.9. *If g is a definable function $X \to M^m$, and $g \circ \pi : Y \to X \to M^m$ is continuous at some $y \in Y$, then g is continuous at $\pi(y)$.*

Proof. Given an open neighborhood U around $g(\pi(y))$, there is an open neighborhood V around y such that $g(\pi(V)) \subseteq U$, by definition of continuity. Then $V' = \pi(V)$ is an open neighborhood of $\pi(y)$, because π is open. So V' is an open neighborhood of $\pi(y)$ and $g(V') \subseteq U$. As U was arbitrary, g is continuous at $\pi(y)$. □

Lemma 8.5.10. *If Y, E are definable over some set T and if $f : X \to M^m$ is a T-definable function and if $x \in X$ is the image under π of a generic point $y \in Y$ (over T), then f is continuous at x. In fact, it is continuous on a neighborhood of x.*

Proof. By cell decomposition, $f \circ \pi$ is continuous at y. By the previous lemma, f is continuous at x. Now let $f' : X \to M^1$ be the characteristic function of where f is continuous. Applying this argument to f' in place of f, we see that f' is locally constant around x, and so f is continuous on a neighborhood of x. □

If A is a definable (i.e., interpretable) topological space, definable over a set T, say that a point $\alpha \in A$ is *nice* if every T-definable function $f : A \to M^n$ is continuous at α. By the argument in the proof of the previous lemma, this also implies that every T-definable function $f : A \to M^n$ is continuous on a neighborhood of α, and that any T-definable subset D of A has $x \notin \mathrm{bd}(D)$.

- If Y and E are T-definable, and $y \in Y$ has $\dim(y/T) = \dim Y$, then $\pi(y) \in X$ is a nice point of X, by Lemma 8.5.10.
- If Z is some T-definable subset of M^n, and $z \in Z$ has $\dim Z = \dim(z/T)$, then z is a nice point of Z. This follows by cell decomposition (z cannot be in the closure of any cell other than the top-dimensional cell which it belongs to).

Lemma 8.5.11. *Let A and B be two definable topological spaces, definable over a set T. Suppose that $f : A \to B$ is a continuous T-definable function. Suppose that $A \subseteq M^m$ for*

some m. Suppose that α is a nice point of A and $f(\alpha)$ is a nice point of B. Suppose that α and $f(\alpha)$ are inter-definable over T. Then some definable neighborhood of α in A is definably homeomorphic via f to some definable neighborhood of $f(\alpha)$ in B.

Proof. Let β be $f(\alpha)$, and write $\alpha = g(\beta)$ for some T-definable function g. Let Γ be the set of pairs $(a, b) \in A \times B$ such that $b = f(a)$, $a = g(b)$, and g is continuous at b. Note that $(\alpha, \beta) \in \Gamma$. Indeed, since $A \subseteq M^m$, the T-definable function $g : B \to A \subseteq M^m$ must be continuous at β, because β is nice. Now let A' be the projection of Γ to A, and B' be the projection of Γ to B. So $f \upharpoonright A'$ is a continuous bijection from A' to B', and the inverse is the continuous bijection $g \upharpoonright B'$. In particular, f induces a homeomorphism from A' to B'. As α is nice and A' is T-definable, α is not in the boundary of A'. So we can find some definable open neighborhood U of α such that $\alpha \in U \subseteq A'$. Similarly, we can find some definable open neighborhood V of β such that $\beta \in V \subseteq B'$. Then f induces a homeomorphism from $U \cap g(V)$ to $f(U) \cap V$. But $g(V)$ is a relatively open subset of A', so $U \cap g(V)$ is a relatively open subset of $U \cap A' = U$, hence open as a subset of A. Similarly, $f(U) \cap V$ is an open neighborhood of β in B. Therefore α and β have definably homeomorphic definable open neighborhoods, with the homeomorphism induced by f. □

Lemma 8.5.12. *If $y \in Y$ is generic (i.e., $\dim(y/S) = \dim Y$), then X is locally Euclidean around $\pi(y)$.*

Proof. Let $x = \pi(y)$. Let $k = \dim(y/Sx)$. Some k of the coordinates of y form a basis for the tuple $\vec{y}$ in the $\mathrm{dcl}_{Sx}(-)$ pregeometry. Permuting the coordinates, we may assume that $y_1, \ldots, y_k$ form a basis for $\vec{y}$. Then $y \in \mathrm{dcl}(Sxy_1, \ldots, y_k)$. As

$$k \geq \dim(y_1, \ldots, y_k/S) \geq \dim(y_1, \ldots, y_k/Sx) = k,$$

we have $y_1 \ldots y_k \mathop{\smile\!\!\!\!\!\!|}^{\mathfrak{b}}_{S} x$. Let T be $S \cup \{y_1, \ldots, y_k\}$. Then y and x are interdefinable over T.

Let Y' be the set of $y' \in Y$ whose first k coordinates agree with those of y. Note that Y' is T-definable. Note that $\pi \upharpoonright Y'$ is a T-definable continuous functoin from Y' to X, and that y and $\pi(y) = x$ are interdefinable over T. We can apply Lemma 8.5.11 with $A = Y'$, $B = X$, $f = \pi$, and $\alpha = y$ assuming we prove the following:

- x is a nice point of X (with respect to T).
- y is a nice point of Y' (with respect to T).

Take $y' \in \pi^{-1}(x) = E(y)$ with $\dim(y'/y) = \dim E(y)$. As $E(y)$ is Sx-definable,

$$\dim(y'/Sx) = \dim E(y)$$

and $y' \mathop{\smile\!\!\!\!\!\!|}^{\mathfrak{b}}_{Sx} y$. By monotonicity, $y' \mathop{\smile\!\!\!\!\!\!|}^{\mathfrak{b}}_{Sx} T$. As $x \mathop{\smile\!\!\!\!\!\!|}^{\mathfrak{b}}_{S} T$, transitivity yields $xy' \mathop{\smile\!\!\!\!\!\!|}^{\mathfrak{b}}_{S} T$. In particular, $\dim(y'/T) = \dim(y'/S)$. Now

$$\begin{aligned}\dim(y'/S) &= \dim(y'/Sx) + \dim(x/S) = \dim E(y) + \dim(x/S) \\ &\geq \dim(y/Sx) + \dim(x/S) = \dim(y/S) = \dim Y,\end{aligned}$$

so $\dim(y'/T) \geq \dim Y$. Therefore, y' is a generic point in Y as far as T is concerned. It follows by Lemma 8.5.10 that $\pi(y') = x$ is nice, with respect to T.

Meanwhile, y is a nice point of Y' because $\dim(y/T) = \dim Y'$. Indeed, take $y'' \in Y'$ with $\dim(y''/T) = \dim Y'$. Then the first k coordinates of y'' are $y_1, \ldots, y_k$, so $T \in \mathrm{dcl}(Sy'')$. Thus

$$\begin{aligned}\dim(y''/S) &= \dim(y''/T) + \dim(T/S) = \dim Y' + \dim(T/S) \\ &\geq \dim(y/T) + \dim(T/S) = \dim(y/S) = \dim Y.\end{aligned}$$

Since $y'' \in Y$ and Y is S-definable, equality must hold. Then $\dim(y/T) = \dim Y'$. It follows that y is a nice point of Y'.

So Lemma 8.5.11 applies. In particular, π induces a homeomorphism from some open neighborhood of y in Y' to some open neighborhood of y in X. It remains to show that some open neighborhood of y in Y' is Euclidean. But this is clear, using the cell-decomposition of the T-definable set Y', and the fact that $\dim(y/T) = \dim Y'$. (This ensures that y is in a top-dimensional cell, and is not in the closure of any lower-dimensional cell. Consequently, Y' looks like the interior of a cell, around the point y.) □

Now let Q be the set of points $y \in Y$ such that X is locally Euclidean around $\pi(y)$. Then Q is ind-definable (the complement of a type-definable set). We have just seen that Q contains the type-definable set $\{y \in Y : \dim(y/S) = \dim Y\}$. By compactness, there must be some definable set $D \subseteq Y$ such that

$$\{y \in Y : \dim(y/S) = \dim Y\} \subseteq D \subseteq Q$$

The first inclusion implies that $\dim Y \setminus D < \dim Y$. Letting $Y' = Y \setminus \overline{Y \setminus D}$, we get that Y' is an open subset of Y, $Y \setminus Y'$ has lower dimension than $\dim Y$, and $Y' \subseteq Q$, so that the topology on $\pi(Y')$ is locally Euclidean. As usual, we replace Y with Y' and X with $\pi(Y')$.

At this point we ensured that the topology is definable, Hausdorff, and locally Euclidean. By Lemma 8.4.13, we regularity also holds. This completes the proof of Theorem 8.5.1.

Chapter 9

VC-minimal and dp-minimal fields

9.1 Introduction

A common goal in model theory is the classification of algebraic structures that satisfy some combinatorial condition. For example, there are various conjectures and programs to classify groups of finite Morley rank, strongly minimal sets, or stable fields. In some cases, unqualified classifications are known, like Cherlin and Shelah's theorem that superstable fields are algebraically closed [8].

However, many classifications seem to require additional topological assumptions. For example, Zilber's conjecture—that sufficiently rich strongly minimal sets interpret fields—is true if the strongly minimal set is endowed with a Zariski-like topology, but false in general [37], [31].

In fact, many of model theory's greatest successes have been in settings like o-minimality and valuation theory, where there are "built-in" topologies. There are few instances where a useful topology can be produced from scratch, using nothing more than combinatorial assumptions on an algebraic structure.

This chapter is about an instance where a topology *does* appear out of thin air.

9.1.1 VC-minimality and dp-minimality

Several "minimality" properties have been defined by restricting the behavior of unary definable sets. For example, a theory is strongly minimal if every unary definable set is finite or cofinite, and o-minimal if every unary definable set is a finite union of intervals and points. These minimality notions play basic roles in classification theory, groups of finite Morley rank, the model theory of real exponentiation, and the model theoretic treatment of Berkovich spaces [34] among other things.

Adler's [2] notion of *VC-minimality* is a common weakening of many of these properties.

Definition 9.1.1. *A complete theory T is* VC-minimal *if there is a family $\mathcal{B}$ of unary definable sets such that*

- *$\mathcal{B}$ is a union of 0-definable families.*
- *Every unary definable set is a finite boolean combination of sets in $\mathcal{B}$.*
- *Whenever two members of $\mathcal{B}$ intersect, one contains the other.*

This can be seen as a generalization of C-minimality, thinking of $\mathcal{B}$ as the class of balls. In fact, all o-minimal, strongly minimal, weakly o-minimal, and C-minimal theories are VC-minimal [2].

An even weaker condition is *dp-minimality*, defined by Shelah but isolated as an interesting notion by Onshuus and Usvyatsov [66], [56].

Definition 9.1.2. *A complete theory T is **not*** dp-minimal *if there is a model $M \models T$, elements $a_{ij} \in M$ and uniformly definable unary definable sets $X_i, Y_j \subseteq M$ for $i, j < \omega$ such that*

$$a_{ij} \in X_{i'} \iff i = i'$$
$$a_{ij} \in Y_{j'} \iff j = j'$$

for all i, j, i', j'.

There is a more general notion of *dp-rank* that is somewhat well-behaved in NIP theories [41] (see 9.2.2 below). A theory is dp-minimal if the home sort has dp-rank 1 (or less).

If we change the definition of dp-minimal by requiring only the Y_j's (but not the X_i's) to be uniformly definable, then we get Guingona's notion of *dp-smallness*, which is weaker than VC-minimality and stronger than dp-minimality.

We summarize the implications between some of these "minimality" properties below:

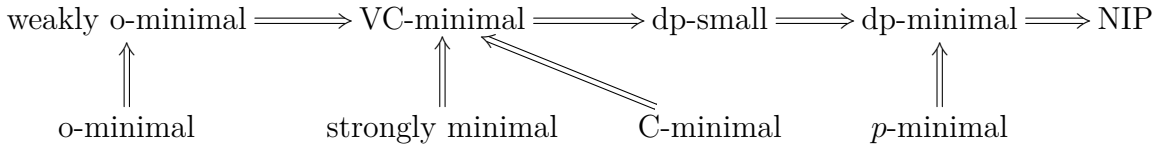

9.1.2 Main results

Our first main result puts a canonical definable topology on almost all dp-minimal fields.

Theorem 9.1.3. *Let $(K, +, \cdot, \ldots)$ be an infinite field, possibly with extra structure. Suppose K is dp-minimal but not strongly minimal. Then*

$$\{X - X : X \subseteq K \text{definable and infinite}\}$$

is a neighborhood basis of 0 for a Hausdorff non-discrete field topology on K. This topology has the following properties:

1. *If* $0 \notin \overline{X}$, *then* $0 \notin \overline{X \cdot X}$ *for* $X \subseteq K$.

2. *There is a definable basis of open sets.*

3. *Any definable set has finite boundary.*

We call this topology the *canonical topology on* K.

Condition 1 says that the field topology is a *V-topology*. By a theorem of Kowalsky and Dürbaum [43], this means that the topology arises from a non-trivial valuation or absolute value on K.

Condition 2 means that the topology is definable.[1] The valuation or absolute value defining the topology need not be definable (or even unique), but the following is true:

Theorem 9.1.4. *Let* $(K, +, \cdot, \ldots)$ *be a field, perhaps with additional structure. Suppose* K *is dp-minimal. Let* $\mathcal{O}$ *be the intersection of all 0-definable valuation rings on* K. *Then*

1. $\mathcal{O}$ *is itself a valuation ring, possibly trivial (i.e.,* $\mathcal{O}$ *might equal* K*).*

2. *Either* $\mathcal{O} = K$, *or* K *is unstable and* $\mathcal{O}$ *induces the canonical topology on* K.

3. $\mathcal{O}$ *is henselian and defectless.*

4. *The residue field of* $\mathcal{O}$ *is finite, algebraically closed, or real closed.*

Part (4) relies crucially on Jahnke and Koenigsmann's work on defining canonical p-henselian valuations [38].

With a little more work, we obtain a classification of dp-minimal pure fields up to elementary equivalence. Before stating it, we fix some notation. If (K, v) is a valued field, let Kv denote the residue field and vK denote the value group. If Γ is an ordered abelian group and p is prime, let $\mathrm{Int}_p \Gamma$ denote the maximal convex p-divisible subgroup of Γ.

Theorem 9.1.5 (Classification of dp-minimal fields)**.**

1. *Let* Γ *be an ordered abelian group such that* $\Gamma/n\Gamma$ *is finite for all* $n > 1$.

   a) *Let* k *be a local field of characteristic 0. The theory of henselian valued fields* (K, v) *with* $vK \equiv \Gamma$ *and* $Kv \equiv k$ *is complete and dp-minimal.*

   b) *If* p *is a prime and* $\Gamma = \mathrm{Int}_p \Gamma$, *then the theory of henselian defectless characteristic* p *valued fields* (K, v) *with* $vK \equiv \Gamma$ *and* $Kv \models ACF_p$ *is complete and dp-minimal.*

   c) *If* p *is a prime and* $0 < a \in \mathrm{Int}_p \Gamma$, *then the theory of characteristic 0 henselian defectless valued fields* (K, v) *with* $Kv \models ACF_p$ *and* $(vK, v(p)) \equiv (\Gamma, a)$ *is complete and dp-minimal.*

[1] We will use "definable topology" to mean a topology in which there is a definable basis of opens. In §3.2 of [34], a more general notion is defined (having an ind-definable basis), and the notion we have just given is called "definable in the sense of Ziegler," after [72]. The only definable topologies we consider will be those in the sense of Ziegler.

2. *If F is a pure field which is infinite and dp-minimal, then $Th(F)$ is the reduct to the language of rings of one of the above theories, for some Γ, k, p, a.*

Corollary 9.1.6. *An infinite field K of characteristic $p > 0$ is dp-minimal if and only if it is elementarily equivalent to a Hahn series field $\mathbb{F}_p^{alg}((t^\Gamma))$ where Γ is a p-divisible ordered abelian group with $\Gamma/n\Gamma$ finite for all $n > 0$.*

In characteristic 0, case 1a consists of exactly the fields elementarily equivalent to $K((t^\Gamma))$ where K is a characteristic 0 local field and $\Gamma/n\Gamma$ is finite for all n. However, there seems to be no clean way to describe case 1c, which includes annoyances like the spherical completions of

$$\mathbb{Q}_p^{un}(p^{1/p}, p^{1/p^2}, p^{1/p^3}, \ldots),$$

where $\mathbb{Q}_p^{un}$ is the maximal unramified extension of $\mathbb{Q}_p$.

Specializing to the case of dp-small fields, the results are much better:

Theorem 9.1.7. *Let K be an infinite field, possibly with extra structure.*

1. *If K is VC-minimal (or dp-small), then K is real closed or algebraically closed.*
2. *If K is C-minimal, then K is algebraically closed.*
3. *If K is weakly o-minimal, then K is real closed.*

In Case 2, we mean C-minimal in the sense of [28], rather than Delon's more general definition in [13] which inclues theories like RCF. Note that Cases 2 and 3 are exactly the main results of [28] and [50], *except* that we have generalized slightly: we do not assume any compatibility between the field operations and the C-predicate or ordering.

9.1.3 Previous work on dp-minimal fields

Dolich, Goodrick, and Lippel showed that $\mathbb{Q}_p$ is dp-minimal [15]. Goodrick [23] and Simon [68] proved some results concerning divisible ordered dp-minimal groups: Goodrick proved an analogue of the monotonicity theorem for o-minimal structures, and Simon proved that infinite sets have non-empty interior. Building off their work, as well as [50], Guingona proved that VC-minimal ordered fields are real closed [24].

Very recently, Walsberg, Jahnke, and Simon have classified dp-minimal *ordered* fields [39], among other things. In the process they obtained some of the results described below—they proved most of the classification result (Theorem 9.1.5) modulo Theorem 9.1.3 (see Propositions 7.4 and 8.1 in [39]).

Theorem 9.1.3 implies that stable dp-minimal infinite fields are strongly minimal. This strengthens an earlier result of Krupiński and Pillay, who proved that certain stable fields, including infinite dp-minimal stable fields, are algebraically closed (Corollary 2.4 in [46]).

9.1.4 Outline

In §9.2.1 we review some basic facts about the model theory of topological fields. In §9.2.2 we review the basics of dp-rank. With these preliminaries out of the way, we show that dp-minimal fields admit definable topologies in §9.3. Sections 9.3.1-9.3.2 show that there is a topology on the additive group, and sections 9.3.3-9.3.4 show that the topology is a definable V-topology. For some reason, the main technical step is a proof that there are a bounded number of infinitesimal types, which is done in §9.3.3.

With the topology obtained, we show in §9.4 that the valuation ring associated to the infinitesimals is henselian (equivalently, the topology is "t-henselian"). Along the way, we see that dp-rank of definable sets is definable in families, in §9.4.1.

In §9.5, we combine our results with the canonical p-henselian valuation machinery of [38], completing most of the proof of the classification theorem for dp-minimal fields. At this point, we have proven enough to classify VC-minimal, densely C-minimal, and weakly o-minimal fields. We do so in §9.6.

Finally, in §9.7, we quickly finish the proof of the classification theorem. Section 9.8 uses this to characterize the pure dp-minimal valued fields. In §9.9, we discuss open questions and potential future research directions.

9.1.5 Notations and conventions

In a topological space, closure and interior will be denoted $\overline{X}$ and X^{int}. If K is a field, K^{alg} and K^{sep} will denote the algebraic closure and separable closure, and $K^\times$ will denote the multiplicative group of K. If (K, v) is a valued field, vK will denote the value group, Kv will denote the residue field, $\mathcal{O}$ will denote the valuation ring, and $\mathfrak{m}$ will denote the maximal ideal of $\mathcal{O}$. If $\alpha \in K$, then $\text{res}(\alpha)$ will denote the residue of α, $v(\alpha)$ will denote the valuation of α, and $rv(\alpha)$ will denote the image of α in $K^\times/(1+\mathfrak{m})$.

We will use $\mathbb{M}$ to denote monster models. The symmetric difference of two sets S and T will be written $S \triangle T$. We will write expressions like aS, ST, or Ta rather than $\{a\} \cup S$, $S \cup T$, and $T \cup \{a\}$. Multiplication will always be written explicitly, like $a \cdot S$, $S \cdot T$, and so on. A code for a definable set D will be written $\ulcorner D \urcorner$. A collection of sequences is mutually indiscernible if each is indiscernible over the union of the others. We will say "a is from M" to mean that a is a tuple from M; variables will not be singletons unless stated explicitly. If φ and ψ are statements, $\varphi \not\Longleftrightarrow \psi$ will mean that exactly one of φ and ψ is true.

Valuations on fields can be trivial. Topologies need not be Hausdorff, but group topologies, ring topologies, field topologies, and V-topologies will always be Hausdorff, though possibly discrete.

9.2 Background material

We first review some background material on field topologies §9.2.1 and dp-rank §9.2.2.

9.2.1 Topologies and type-definable sets

For us, group topologies and field topologies are *Hausdorff* topologies such that the group and field operations (including inversion) are continuous on their domains. Unlike [61] we consider the discrete topology as a valid group topology or field topology. We will use $\overline{X}$ and X^{int} to denote the closure and interior of a set X in a topological space.

In a topological field K, say that a subset $X \subseteq K$ is *bounded away from 0* if $0 \notin \overline{X}$. A field topology is *of type V*, or a *V-topology*, if $X \cdot Y$ is bounded away from 0 whenever X and Y are bounded away from 0.

The topologies arising from valuations, orderings, and absolute values are all V-topologies. In fact, a theorem of Kowalsky and Dürbaum [43] says that all V-topologies arise from absolute values and valuations. A simple example of a field topology which is not a V-topology is the subspace topology on $\mathbb{Q}$ induced by the diagonal embedding $\mathbb{Q} \to \mathbb{R} \times \mathbb{Q}_2$. In this topology the sets $3^{\mathbb{Z}}$ and $2^{\mathbb{Z}}$ are bounded away from 0, but their product is not.

Fix a field K and let $\mathcal{N}$ be a family of subsets of K. Consider the following axioms on $\mathcal{N}$, adapted from [61]:

(A0) $\forall U, V \exists W \quad W \subseteq U \cap V$

(A1) $\forall U \quad 0 \in U$

(A2) $\forall U \exists V \quad V - V \subseteq U$

(A3) $\forall x \neq 0 \exists U \quad x \notin U$

(A4) $\forall U \exists x \neq 0 \quad x \in U$

(A5) $\forall U, x \exists V \quad x \cdot V \subseteq U$

(A6) $\forall U \exists V \quad V \cdot V \subseteq U$

(A7) $\forall U \exists V \quad (1+V)^{-1} \subseteq 1 + U$

(A8) $\forall U \exists V \forall x, y \quad x \cdot y \in V \implies (x \in U \vee y \in U)$

Here, uppercase variables range over $\mathcal{N}$ and lowercase variables range over K.

Then $\mathcal{N}$ is a neighborhood basis of 0 for a group topology on $(K, +)$ if and only if axioms (A0-A3) hold. If so, then...

- The topology is non-discrete if and only if axiom (A4) holds.
- The topology is a ring topology if and only if axioms (A5-A6) hold
- The topology is a field topology if and only if axioms (A5-A7) hold
- The topology is a V-topology if and only if axioms (A5-A8) hold.

Say that a family $\mathcal{N}$ of sets is (downward) *directed*[2] if it satisfies axiom (A0). Assuming directedness, the other axioms have a non-standard interpretation.

To wit, suppose that $K = (K, \cdot, +, \ldots)$ is an expansion of a field, the elements of $\mathcal{N}$ are definable subsets of K, and $\mathbb{M} \succeq K$ is a $|K|^+$-saturated elementary extension. Let I be the type-definable set of "infinitesimals"

$$I = \bigcap_{U \in \mathcal{N}} U(\mathbb{M})$$

If $\mathcal{N}$ is directed, then axioms (A1-A8) are equivalent to the following statements about I, respectively:

(A1) $0 \in I$

(A2) $I - I \subseteq I$

(A3) $I \cap K^\times = \emptyset$

(A4) $I \cap \mathbb{M}^\times \neq \emptyset$

(A5) $K \cdot I \subseteq I$

(A6) $I \cdot I \subseteq I$

(A7) $(1 + I)^{-1} \subseteq 1 + I$

(A8) $(\mathbb{M} \setminus I) \cdot (\mathbb{M} \setminus I) \subseteq \mathbb{M} \setminus I$

Many of these statements are algebraic statements about I: for instance axioms (A1) and (A2) say that I is a subgroup of $\mathbb{M}$. Taken together, the axioms say that I is the maximal ideal of a non-trivial valuation ring containing K.

We will build our topologies on dp-minimal fields by finding a directed family $\mathcal{N}$, considering the set of "infinitesimals," and showing that it has enough nice algebraic properties. Theorem 9.2.4 below is the black box we will use.

Lemma 9.2.1. *Suppose K expands a field, $\mathbb{M}$ is a $|K|^+$-saturated elementary extension, and $J_1, J_2 \subseteq \mathbb{M}$ are type-definable over K. Suppose J_1 and J_2 are closed under multiplication by K. If $a \cdot J_1 \subseteq J_2$ for some $a \in \mathbb{M}^\times$, then $J_1 \subseteq J_2$.*

Proof. Let U be any K-definable set containing J_2. We will show $J_1 \subseteq U$.

As $a \cdot J_1 \subseteq J_2 \subseteq U$, we have an implication

$$x \in J_1 \implies a \cdot x \in U$$

[2]This is not related to the notion of "directed" used to define VC-minimality in [2].

Compactness yields a K-definable neighborhood $V \supseteq J_1$ such that

$$x \in V \implies a \cdot x \in U$$

As $K \preceq \mathbb{M}$, we can find $a' \in K$ having the same property as a:

$$x \in V \implies a' \cdot x \in U$$

Then

$$J_1 = a' \cdot J_1 \subseteq a' \cdot V \subseteq U$$

As U was an arbitrary neighborhood of J_2, we conclude $J_1 \subseteq J_2$. □

Recall that a topology is *definable* if there is a (uniformly) definable basis of opens.

Lemma 9.2.2. *Suppose K is an expansion of a field and $\mathcal{N}$ is a collection of unary definable sets which determines a non-trivial V-topology. Then there is at least one set B such that $0 \in B^{int}$ and $0 \notin \overline{B^{-1}}$. If B is* any *such set, then*

$$\{a \cdot B : a \in K^\times\}$$

is a neighborhood basis of 0.

Proof. Let $\mathbb{M}$ be a $|K|^+$-saturated elementary extension, and let I be the infinitesimals. Then I satisfies axioms (A1-A8) above. Because $1 \notin I$, by axiom (A3), and $\mathbb{M} \setminus I$ is closed under multiplication, by axiom (A8), we have the implication

$$x \in I \implies x^{-1} \notin I$$

For a K-definable set B, the condition that $0 \in B^{int}$ and $0 \notin \overline{B^{-1}}$ is equivalent to the implication

$$x \in I \implies x \in B \implies x^{-1} \notin I$$

for $x \in \mathbb{M}$. We can find a B with this property by compactness (as I is type-definable). Now suppose we have such a B.

Now $I \subseteq B$, so B is a neighborhood of 0. As I is closed under multiplication by K, by axiom (A5), $I \subseteq a \cdot B$ for any $a \in K^\times$. So each set $a \cdot B$ is a neighborhood of 0.

We claim that $\{a \cdot B : a \in K^\times\}$ is a neighborhood basis of 0. Let $U \in \mathcal{N}$ be any neighborhood in the given basis $\mathcal{N}$. Let ϵ be a non-zero element of I, which exists by non-discreteness axiom (A4). Then for $x \in \mathbb{M}$, we have the following implications

$$x \in B \implies x^{-1} \notin I \overset{*}{\implies} \epsilon \cdot x \in I \implies \epsilon \cdot x \in U$$

where the starred implication holds because $\mathbb{M} \setminus I$ is closed under multiplication, by axiom (A8).

So $\epsilon \cdot B \subseteq U$. As $\mathrm{tp}(\epsilon/K)$ is finitely satisfiable in K, there is some $a \in K^\times$ such that $a \cdot B \subseteq U$. This shows that the rescalings of B form a neighborhood basis of 0. □

Lemma 9.2.3. *Suppose K is an expansion of a field and $\mathcal{N}$ is a collection of unary definable sets which determines a V-topology. Then this topology is definable.*

Proof. If the topology is non-discrete, then Lemma 9.2.2 yields a definable basis of neighborhoods around 0. Translating, there is a definable basis of neighborhoods around each $x \in K$, uniformly definable across x. This implies that the interior of any definable set is definable. We can thus convert the given neighborhoods into open neighborhoods. Now the union of all the definable bases of open neighborhoods is a basis of open sets for the topology.

If the topology is discrete, then the singletons are a definable basis of opens. □

Theorem 9.2.4. *Let K be an expansion of a field, and let $\mathcal{N}$ be a directed family of definable subsets of K. Let $\mathbb{M} \succeq K$ be a $|K|^+$-saturated elementary extension, and let*

$$I = \bigcap_{U \in \mathcal{N}} U(\mathbb{M})$$

1. *Suppose the following hold:*
   a) *I is a subgroup of $(\mathbb{M}, +)$*
   b) *Every member of $\mathcal{N}$ is infinite*
   c) *I is closed under multiplication by K*
   d) *Some member of $\mathcal{N}$ has non-empty complement.*

   Then $\mathcal{N}$ determines a non-discrete group topology on $(K, +)$.

2. *Suppose the following* additional *conditions hold:*
   a) *$1 + I$ is a subgroup of $\mathbb{M}^\times$.*
   b) *For every $a \in \mathbb{M}$, either $a \cdot I$ contains I or vice versa.*

   Then $\mathcal{N}$ determines a definable non-discrete V-topology, and I is the maximal ideal of a non-trivial valuation ring $\mathcal{O}$ on $\mathbb{M}$, with $\mathcal{O} \supseteq K$.

Proof.

1. For $\mathcal{N}$ to determine a non-discrete group topology, it must be directed (which is given), and I must satisfy axioms (A1-A4) above:
   - $0 \in I$ holds because I is a subgroup
   - $I - I \subseteq I$ holds because I is a subgroup
   - As I is closed under multiplication by K, the intersection $I \cap K$ is an ideal in K, so it is either 0 or K. By assumption, there is some $U \in \mathcal{N}$ with non-empty complement. Then $U(K) \neq K$, so $I \cap K \subsetneq K$. Thus $I \cap K = \{0\}$, and $I \cap K^\times = \emptyset$.
   - $I \supsetneq \{0\}$ because $\mathcal{N}$ is directed and every member of $\mathcal{N}$ is infinite.

2. Assume that all conditions from 1. and 2. hold.

Claim 9.2.5. *I is closed under multiplication.*

Proof. By assumption, $1 + I$ is closed under multiplication and I is closed under subtraction. So

$$x, y \in I \implies (1+x)(1+y) \in 1+I \implies x+y+xy \in I \implies xy \in I.$$

Therefore I is closed under multiplication. □

Set

$$\mathcal{O} = \{x \in \mathbb{M} : x \cdot I \subseteq I\}$$

Then

- $\mathcal{O}$ is trivially closed under multiplication, and contains 1.
- $\mathcal{O}$ is closed under addition and subtraction because I is. Therefore $\mathcal{O}$ is a subring of $\mathbb{M}$.
- $I \subseteq \mathcal{O}$ because I is closed under multiplication. So I is an ideal in $\mathcal{O}$.
- Because I satisfies axiom (A3), $1 \notin I$. So I is a proper ideal of $\mathcal{O}$.
- For every $a \in \mathbb{M}^\times$, either a or a^{-1} is in $\mathcal{O}$ by assumption 2b. So $\mathcal{O}$ is a valuation ring.
- The valuation ring $\mathcal{O}$ is non-trivial (not equal to $\mathbb{M}$), because I is a non-trivial proper ideal.
- By assumption 1c, $K \cdot I \subseteq I$, so $K \subseteq \mathcal{O}$.

Let $v : \mathbb{M} \to \Gamma$ be the valuation associated to $\mathcal{O}$. Because I is an ideal

$$v(x) \leq v(y) \implies (x \in I \implies y \in I)$$

Claim 9.2.6. *$\mathbb{M} \setminus I$ is closed under multiplication*

Proof. Let $J = \{x \in \mathbb{M} : x^2 \in I\}$. This is type-definable over K because I is. Also, J is closed under multiplication by K because I is. Note that

$$x \in J \implies x^2 \in I \implies v(x^2) \geq 0 \implies v(x) \geq 0 \implies x \in \mathcal{O}$$

so $J \subseteq \mathcal{O}$. Choosing nonzero $a \in I$, we have

$$a \cdot J \subseteq a \cdot \mathcal{O} \subseteq I$$

By Lemma 9.2.1, $J \subseteq I$.

Now suppose $x, y \in \mathbb{M} \setminus I$ but $xy \in I$. Without loss of generality, $v(x) \leq v(y)$. Then $v(xy) \leq v(y^2)$ so $y^2 \in I$, so $y \in J$, so $y \in I$, a contradiction. □

Claim 9.2.7. *I is the maximal ideal $\mathfrak{m}$ of $\mathcal{O}$*

Proof. The inclusion $I \subseteq \mathfrak{m}$ holds because $\mathfrak{m}$ is the unique maximal ideal.

Conversely, suppose $x \notin I$. By Claim 9.2.6, multiplication by x preserves $\mathbb{M} \setminus I$. Equivalently, division by x preserves I. Thus $x^{-1} \in \mathcal{O}$, or equivalently, $x \notin \mathfrak{m}$. □

So I is the maximal ideal of a non-trivial valuation ring on $\mathbb{M}$ containing K.

Now, as I is the maximal ideal of the valuation ring $\mathcal{O}$, it clearly satisfies axioms (A6-A8) above. Axioms (A1-A4) were checked in the previous point, and axiom (A5) was assumed. Therefore all the axioms on I hold, and $\mathcal{N}$ determines a non-discrete V-topology on K. The topology is definable by Lemma 9.2.3.

□

9.2.2 Dp-rank

If X is a type-definable set and κ is a cardinal, a *randomness pattern of depth κ in X* is a collection of formulas $\{\phi_\alpha(x; y_\alpha) : \alpha < \kappa\}$ and elements $\{b_{i,j} : i < \kappa, j < \omega\}$ such that for every function $\eta : \kappa \to \omega$ there is some element a_η in X such that for all i, j

$$\phi_i(a_\eta, b_{ij}) \iff j = \eta(i)$$

The dp-rank of X is defined to be the supremum of the cardinals κ such that there is a randomness pattern of depth κ in X. This definition first appears in [70].

The following fundamental facts about dp-rank are in [70] or [41], or are easily verified:

Fact 9.2.8.

1. *The formula $x = x$ has dp-rank less than ∞ if and only if the theory is NIP.*
2. *The formula $x = x$ has dp-rank at most 1 if and only if the theory is dp-minimal.*
3. *If X is type-definable over A, then* dp-rk(X) *is the supremum of* dp-rk(x/A) *for $x \in X$.*
4. dp-rk$(X) > 0$ *if and only if X is infinite.*
5. *For $n < \omega$,* dp-rk$(a/A) \geq n$ *if and only if there are sequences $I_1, \ldots, I_n$, which are mutually indiscernible over A, such that each sequence is not individually Aa-indiscernible.*
6. *Dp-rank is subadditive:* dp-rk$(ab/A) \leq$ dp-rk$(a/bA) +$ dp-rk(b/A).
7. *If X and Y are non-empty type-definable sets, then* dp-rk$(X \times Y) =$ dp-rk$(X) +$ dp-rk(Y).
8. *If* dp-rk$(a/A) = n$ *and X is an A-definable set of dp-rank 1, then there is $b \in X$ such that* dp-rk$(ab/A) = n + 1$.

9. If $X \twoheadrightarrow Y$ is a definable surjection, then $\text{dp-rk}(Y) \leq \text{dp-rk}(X)$.

Here are some basic uses of dp-rank:

Observation 9.2.9. *Let K be a field of finite dp-rank. Then K is perfect.*

Proof. The field K^p of pth powers is in definable bijection with K, so it has the same rank as K. If K is imperfect, then K is a definable K^p vector space of dimension greater than 1. It contains a two-dimensional subspace, so $K^p \times K^p$ injects definably into K. This shows

$$\text{dp-rk}(K) \geq 2 \cdot \text{dp-rk}(K^p) = 2 \cdot \text{dp-rk}(K)$$

So $\text{dp-rk}(K) = 0$, and K is finite. Finite fields are perfect. □

Observation 9.2.10. *Let K be dp-minimal field. Then K eliminates $\exists^\infty$ (in powers of the home sort).*

Proof. We may assume that K is infinite, so $\text{dp-rk}(K) = 1$. It suffices to show that a definable set $X \subseteq K$ is finite if and only if there is some $a \in K$ such that the map $(x, y) \mapsto x + a \cdot y$ is injective on $X \times X$. If X is finite, any a outside the finite set

$$\left\{ \frac{x_1 - x_2}{x_3 - x_4} : \vec{x} \in X^4 \right\}$$

will work. If X is infinite, then $\text{dp-rk}(X) \geq 1$, so $\text{dp-rk}(X \times X) = 2$ and $X \times X$ cannot definably inject into K. □

This has the following useful corollary:

Corollary 9.2.11. *Suppose K is dp-minimal. Then any infinite externally definable subset of K contains an infinite internally definable set.*

In fact, this property holds in any NIP theory which eliminates $\exists^\infty$. For lack of a reference, we recall the well-known

Proof. Suppose $S \subseteq K$ is externally definable. By honest definitions ([67] Remark 3.14), there is some formula $\phi(x; y)$ such that for every finite $S_0 \subseteq S$, there is $b \in K$ such that $S_0 \subseteq \phi(K; b) \subseteq S$. By elimination of $\exists^\infty$, there is some number n such that $\phi(K; b)$ is infinite or has size less than n, for all b. If we choose S_0 to have size greater than n, then $\phi(K; b)$ will be our desired infinite internally-definable set. □

9.3 The definable V-topology

9.3.1 Infinitesimals

Until §9.5, let $\mathbb{M}$ be a fairly saturated infinite dp-minimal field **that is not strongly minimal**.

If $X, Y \subseteq \mathbb{M}$, let $X -_\infty Y$ denote

$$\{c \in \mathbb{M} : \exists^\infty y \in Y : c + y \in X\}$$

This is a subset of $X - Y$. It is definable if X and Y are, by Observation 9.2.10.

Lemma 9.3.1. *If X and Y are infinite, so is $X -_\infty Y$.*

Proof. Suppose X and Y are A-definable. Take $(x, y) \in X \times Y$ of dp-rank 2 over A, and let $c = x - y$. By subadditivity of dp-rank, and dp-minimality,

$$2 = \text{dp-rk}(x, y/A) = \text{dp-rk}(y, c/A) \leq \text{dp-rk}(y/c, A) + \text{dp-rk}(c/A) \leq 1 + 1$$

Equality must hold, so $y \notin \text{acl}(Ac)$ and $c \notin \text{acl}(A)$. As $y \in Y \cap (X - c)$, the Ac-definable set $Y \cap (X - c)$ is infinite. Then $c \in X -_\infty Y$, so the A-definable set $X -_\infty Y$ is infinite. □

One can give a more direct proof of this fact in theories like ACVF, RCF, and pCF. In each of these theories, infinite definable sets have non-empty interior. If X and Y are infinite and definable, then X and Y contain balls $B_\epsilon(x)$ and $B_\epsilon(y)$. One checks that $B_\epsilon(x-y) \subseteq X -_\infty Y$. Note that when $X = Y$, one can take $x = y$, and so $X -_\infty X$ contains a ball around 0. Moreover, by taking X to be small enough, $X -_\infty X$ will be arbitrarily small. In other words, as X ranges over infinite definable sets, $\{X -_\infty X\}$ is a neighborhood basis of 0.

This suggests looking at the same family in the more general setting.

Proposition 9.3.2. *Let $K \preceq \mathbb{M}$ be a small model. Let*

$$\mathcal{N} = \{X -_\infty X : X \subseteq K \text{ is infinite and } K\text{-definable}\}$$

This is a directed family of infinite sets. At least one member of $\mathcal{N}$ has non-empty complement.

Proof. The members of $\mathcal{N}$ are infinite by Lemma 9.3.1.

To see $\mathcal{N}$ is directed, suppose X and Y are infinite definable sets. We want to find an infinite definable set Z such that

$$Z -_\infty Z \overset{?}{\subseteq} (X -_\infty X) \cap (Y -_\infty Y)$$

If X and Y have infinite intersection, we can simply take $Z = X \cap Y$. In the general case, it is safe to replace Y with a translate, as this does not change $Y -_\infty Y$. We can always translate Y to have infinite intersection with X, because $X -_\infty Y$ is non-empty by Lemma 9.3.1.

By failure of strong minimality and Observation 9.2.10, there is a K-definable set D which is infinite and co-infinite. Let D' be the complement of D. By the Lemma, $D -_\infty D'$ is non-empty, so there is some c such that $X := D \cap (D' + c)$ is infinite. Then $X - c \subseteq D'$ so $(X - c) \cap X = \emptyset$, and $c \notin X -_\infty X$. So $X -_\infty X$ is a non-full member of $\mathcal{N}$. □

Let I_K denote the corresponding type-definable set:

$$I_K = \{x \in \mathbb{M} : x \in X -_\infty X \text{ for all } K\text{-definable infinite } X\}$$

We call the elements of I_K the *K-infinitesimals*. Over the next few sections, we will show that I_K satisfies all the other conditions of Theorem 9.2.4.

9.3.2 Slight maps

Definition 9.3.3. *Let K be a small model. An $\mathbb{M}$-definable bijection $f : \mathbb{M} \to \mathbb{M}$ is* K-slight *if $X \cap f^{-1}(X)$ is infinite for every K-definable infinite set X.*

For example, the translation map $x \mapsto x+\epsilon$ is K-slight if and only if ϵ is a K-infinitesimal. The main goal here is to show that K-slight maps form a group under composition.

Definition 9.3.4. *Let K be a small model, $X \subseteq \mathbb{M}$ be K-definable, and f be an $\mathbb{M}$-definable bijection. Say that X is* K-displaced by f *if $X(K) \cap f^{-1}(X)$ is empty.*

Lemma 9.3.5. *Suppose $K' \succeq K$ and f' and f are $\mathbb{M}$-definable bijections such that* $\mathrm{tp}(f'/K')$ *is an heir of* $\mathrm{tp}(f/K)$*. (Here, we are identifying a bijection with its code.)*

- *If f is K-slight, then f' is K'-slight.*
- *If X is K-displaced by f, then X is K'-displaced by f'.*

For instance, the first point shows that heirs of infinitesimal types are infinitesimal.

Proof. First suppose f if K-slight. As $f' \equiv_K f$, the map f' is also K-slight. If it is not K'-slight, there is a K'-definable infinite set X such that $X \cap (f')^{-1}(X)$ is finite. As $\mathrm{tp}(K'/Kf')$ is finitely satisfiable in K, and infinity is definable, we can pull the parameters of X into K, finding a K-definable infinite set X_0 such that $X_0 \cap (f')^{-1}(X_0)$ is finite. This contradicts K-slightness of f'.

Next suppose X is K-displaced by f. Then X is K-displaced by f'. If X is not K'-displaced by f', there is some $a \in X(K')$ such that $f'(a) \in X$. As $\mathrm{tp}(a/Kf')$ is finitely satisfiable in K, there is some $a_0 \in X(K)$ such that $f'(a_0) \in X$, contradicting the fact that X is K-displaced by f'. □

Lemma 9.3.6. *No K-slight map K-displaces an infinite K-definable set.*

Proof. Suppose f_0 is a K-slight map which K-displaces an infinite K-definable set Y. Inductively build a sequence of models $K_0 = K \preceq K_1 \preceq K_2 \preceq \cdots$ and bijections $f_0, f_1, f_2, \ldots$ such that

- $\mathrm{tp}(f_i/K_i)$ is an heir of $\mathrm{tp}(f_0/K)$.
- f_i is K_{i+1}-definable.

By Lemma 9.3.5, f_i is K_i-slight, and K_i-displaces Y.

For $w \in \{0,1\}^{<\omega}$, consider the set

$$Y_w = \left\{ y \in Y : \bigwedge_{i<|w|} f_i(y) \in^{w(i)} Y \right\}$$

where $\in^0$ denotes $\notin$ and $\in^1$ denotes $\in$.

We will prove by induction on $|w|$ that Y_w is infinite. If we write f_i as f_{a_i}, this shows that the formula $f_x(y) \in Y$ has the independence property, a contradiction.

For the base case, $Y_\emptyset$ is Y which is infinite by assumption.

Now suppose that Y_w is infinite; we will show Y_{w0} and Y_{w1} are infinite. Let $n = |w|$. Then Y_w is K_n-definable. If $a \in Y_w(K_n) \subseteq Y(K_n)$, then $f_n(a) \notin Y$ because Y is K_n-displaced by f_n. This shows that the infinite set $Y_w(K_n)$ is contained in Y_{w0}.

Also, as f_n is K_n-slight and Y_w is infinite and K_n-definable, $Y_w \cap f_n^{-1}(Y_w)$ is infinite. This set is contained in $Y_w \cap f_n^{-1}(Y) = Y_{w1}$, so Y_{w1} is infinite.

So Y_w being infinite implies Y_{w0} and Y_{w1} are infinite. This ensures that all Y_w are infinite, hence non-empty, contradicting NIP. □

Proposition 9.3.7.

1. *If f is a K-slight bijection and X is K-definable, then for all but finitely many $x \in K$, we have $x \in X \iff f(x) \in X$.*
2. *The K-slight bijections form a group under composition.*
3. *If f and g are bijections, f is K-slight, and g is K-definable, then $g^{-1} \circ f \circ g$ is K-slight.*

Proof. 1. Let $S \subseteq K$ be the externally definable set of x such that $x \in X$ and $f(x) \notin X$. We claim that S is finite. Otherwise, by Corollary 9.2.11, there is some infinite K-definable set Y such that $Y(K) \subseteq S$. Then $X \cap Y$ is an infinite K-definable set which is K-displaced by f, by choice of S. This contradicts Lemma 9.3.6.

So S is finite. This means that for almost all $x \in K$, we have $x \in X \implies f(x) \in X$. Replacing X with its complement, we obtain the reverse implication (with at most finitely many exceptions).

2. Suppose f and $g \circ f$ are K-slight. We will show that g is K-slight. Let X be an infinite K-definable set. Then for almost all $x \in K$, we have

$$f(x) \in X \iff x \in X \iff g(f(x)) \in X$$

So the infinite set $f(X(K))$ is almost entirely contained in $X \cap g^{-1}(X)$. Thus $X \cap g^{-1}(X)$ is infinite, for arbitrary infinite K-definable sets X.

3. Let X be K-definable. Then $g(X)$ is infinite and K-definable, so $g(X) \cap f^{-1}(g(X))$ is infinite. Applying g^{-1}, we see that $X \cap g^{-1}(f^{-1}(g(X)))$ is infinite. Therefore $g^{-1} \circ f \circ g$ is K-slight.

□

Corollary 9.3.8. *The set I_K of K-infinitesimals is a subgroup of $(\mathbb{M}, +)$, and is closed under multiplication by K.*

Proof. The first claim follows immediately. For the second, suppose $\epsilon \in I_K$ and $a \in K$. If $a = 0$, then $a \cdot \epsilon = 0 \in I_K$ because I_K is a group. Otherwise, $x \mapsto x \cdot a^{-1}$ is a K-definable bijection, and $x \mapsto x + \epsilon$ is K-slight, so the conjugate

$$x \mapsto x \cdot a^{-1} \mapsto x \cdot a^{-1} + \epsilon \mapsto (x \cdot a^{-1} + \epsilon) \cdot a = x + a \cdot \epsilon$$

is K-slight, meaning $a \cdot \epsilon \in I_K$. □

Theorem 9.3.9. *The family*

$$\mathcal{N} = \{X -_\infty X : X \subseteq K \text{ is infinite and } K\text{-definable}\}$$

determines a non-discrete group topology on $(K, +)$. The family

$$\mathcal{N}' = \{X - X : X \subseteq K \text{ is infinite and } K\text{-definable}\}$$

determines the same topology, and the same type definable set I_K.

Proof. By Proposition 9.3.2, the family $\mathcal{N}$ is a directed family of infinite definable sets, at least one of which is non-full. Because I_K is a subgroup of $\mathbb{M}$ and is closed under multiplication by K, Theorem 9.2.4 applies and $\mathcal{N}$ determines a non-discrete group topology.

We need to show that $\mathcal{N}'$ is also a neighborhood basis for this topology, i.e., that $\mathcal{N}$ and $\mathcal{N}'$ are cofinal in each other. Every member $X - X$ of $\mathcal{N}'$ contains a member of $\mathcal{N}$, namely $X -_\infty X$.

Conversely, given $U \in \mathcal{N}$, we need to produce an infinite definable Y such that $Y - Y \subseteq U$.

As I_K is a group,

$$x, y \in I_K \implies x - y \in I_K \implies x - y \in U$$

By compactness, there is some neighborhood Y of I_K such that

$$x, y \in Y \implies x - y \in U$$

so $Y - Y \subseteq U$. But Y is infinite (because I_K is), so $Y - Y \in \mathcal{N}'$.

It follows that $\mathcal{N}$ and $\mathcal{N}'$ determine the same topology and the same type-definable set I_K. □

We call this topology the *canonical topology* on K. One can also talk about the canonical topology on $\mathbb{M}$ which is itself a dp-minimal field.

9.3.3 Germs at 0

Say that two definable sets $X, Y \subseteq \mathbb{M}$ have the same *germ at 0* if $0 \notin \overline{X \triangle Y}$. This is an equivalence relation. The main goal of this section is Theorem 9.3.16, asserting that there are only a small number of germs at 0—or equivalently, that there are only a small number of infinitesimal types over $\mathbb{M}$. Surprisingly, this technical fact easily yields the remainder of Theorem 9.1.3.

To prove Theorem 9.3.16, we would like to mimic Simon's argument in the case of ordered dp-minimal structures (Lemma 2.10 in [68]). Matters are complicated by our lack of a definable neighborhood basis.

In what follows, we will refer to sets of the form $X - X$ with X infinite and definable, as "basic neighborhoods (of 0)".

Let $\mathcal{U}$ be a 0-definable family of basic neighborhoods (of 0).

Definition 9.3.10. *Say that $\mathcal{U}$ is* good *if for every finite set $S \subseteq \mathbb{M}^\times$, there is some $U \in \mathcal{U}$ such that $U \cap S = \emptyset$.*

Definition 9.3.11. *Say that $\mathcal{U}$ is* mediocre *if for every finite tuple $(a_1, \ldots, a_n) \subseteq \mathbb{M}^\times$ of full dp-rank (of dp-rank n), there is some $U \in \mathcal{U}$ such that $U \cap \{a_1, \ldots, a_n\} = \emptyset$.*

To run Simon's argument directly, one needs a good family. However, the argument can be modified to use a mediocre family, as we will see in the proofs of Lemma 9.3.15 and Theorem 9.3.16 below. Fortunately, the next proposition provides a mediocre family.

Proposition 9.3.12. *There is a mediocre family of basic neighborhoods.*

Proof. Let $\Sigma(x)$ be the partial type over $\mathbb{M}$ saying that $x \neq 0$ and x is an $\mathbb{M}$-infinitesimal.

First suppose that $\Sigma(x)$ is not finitely satisfiable in some small model K. Then there is some $\mathbb{M}$-definable basic neighborhood $U = U_b$ such that $U_b \cap K = \{0\}$. Then for all n, we have the following chain of statements, each of which implies the next:

$$\begin{aligned}
\forall a_1, \ldots, a_n \in K^\times \quad & (U_b \cap \{a_1, \ldots, a_n\} = \emptyset) \\
\forall a_1, \ldots, a_n \in K^\times \exists y \in \mathbb{M} \quad & (U_y \cap \{a_1, \ldots, a_n\} = \emptyset) \\
\forall a_1, \ldots, a_n \in K^\times \exists y \in K \quad & (U_y \cap \{a_1, \ldots, a_n\} = \emptyset) \\
\forall a_1, \ldots, a_n \in \mathbb{M}^\times \exists y \in \mathbb{M} \quad & (U_y \cap \{a_1, \ldots, a_n\} = \emptyset)
\end{aligned}$$

By the final statement the family $\{U_b : b \in \mathbb{M}\}$ is a good family of basic neighborhoods, hence a mediocre family.

Therefore, we may assume that $\Sigma(x)$ is finitely satisfiable in any small model K. This has the following counterintuitive corollary:

Claim 9.3.13. *The canonical topology on K is the induced subspace topology from the canonical topology on $\mathbb{M}$.*

Proof. The induced subspace topology on K will have as neighborhood basis of 0, the sets of the form $N \cap K$ for N an $\mathbb{M}$-definable basic neighborhood. This already includes the K-definable basic neighborhoods on K, so it remains to show that if N is an $\mathbb{M}$-definable basic neighborhood, then there is a K-definable basic neighborhood N' such that $N' \cap K \subseteq N \cap K$.

Because the canonical topology on $\mathbb{M}$ is a group topology, there is an $\mathbb{M}$-definable basic neighborhood U such that $U - U \subseteq N$. If $U \cap K$ is finite, then by Hausdorffness there is some $\mathbb{M}$-definable neighborhood $V \subseteq U$ such that $V \cap K = \{0\}$, contradicting finite satisfiability of $\Sigma(x)$ in K.

Thus $U \cap K$ is infinite, and so it contains $Q(K)$ for some infinite K-definable set Q, by Corollary 9.2.11. Then $Q - Q$ is a K-definable basic neighborhood and

$$(Q - Q)(K) = Q(K) - Q(K) \subseteq U - U \subseteq N$$

so $(Q - Q) \cap K \subseteq N \cap K$. Then $N' := Q - Q$ is our desired K-definable basic neighborhood. This proves the claim. □

Claim 9.3.14. *There is a 0-definable family of basic neighborhoods U_b such that if $K \preceq K'$ is any inclusion of models, and $a \in K' \setminus K$, then $(a + U_b) \cap K = \emptyset$ for some $b \in K'$.*

Proof. If not, then by compactness, we would obtain a pair of models $K \preceq K'$ and an element a such that every K'-definable neighborhood of a intersects K. In other words, a is in the topological closure $\overline{K}$ of K. Embed K' into $\mathbb{M}$. Then K' has the induced subspace topology, so $a \in \overline{K}$ even within $\mathbb{M}$. Because the topology on $\mathbb{M}$ is $\mathrm{Aut}(\mathbb{M}/K)$-invariant, all the conjugates of a over K are in $\overline{K}$, so $\overline{K}$ is big. But in a Hausdorff topology, the closure of a set is bounded in terms of the size of the set (because every point in the closure can be written as an ultralimit of an ultrafilter on the set, and there are only a bounded number of ultrafilters). □

Let U_b be the family from Claim 9.3.14. We claim that U_b is mediocre. To see this, suppose $a_1, \ldots, a_n$ are elements of $\mathbb{M}^\times$ with dp-rank n over the empty set. By Fact 9.2.8.8, we can find an element $t \in \mathbb{M}$ such that $(\vec{a}, t)$ has dp-rank $n + 1$.

By subadditivity of dp-rank,

$$\begin{aligned} n + 1 &= \text{dp-rk}(t, t + a_1, \ldots, t + a_n) \\ &\leq \text{dp-rk}(t/t + a_1, \ldots, t + a_n) + \text{dp-rk}(t + a_1, \ldots, t + a_n) \\ &\leq 1 + n \end{aligned}$$

so equality holds, and $t \notin \mathrm{acl}(t - a_1, \ldots, t - a_n)$. Therefore we can find a small model K such that $t \notin K \supseteq \{t + a_1, \ldots, t + a_n\}$. By the claim there is some $b \in \mathbb{M}$ such that

$$(t + U_b) \cap \{t + a_1, \ldots, t + a_n\} \subseteq (t + U_b) \cap K = \emptyset$$

so that $U_b \cap \{a_1, \ldots, a_n\} = \emptyset$. □

Lemma 9.3.15. *Let $\mathcal{U}$ be a mediocre family of basic neighborhoods. Then given any small collection $\mathcal{C}$ of infinite definable sets, there is some $U \in \mathcal{U}$ such that $C \setminus U$ is infinite for every $C \in \mathcal{C}$.*

Proof. Because infinity is definable and $\mathcal{U}$ is a single definable family, it suffices by compactness to consider the case when $\mathcal{C}$ if a finite collection $\{C_1, \ldots, C_n\}$. By definability of infinity, there is some N (depending on $\mathcal{C}$) such that $C_i \setminus U$ will be infinite as long as it has size at least N.

Let A be a set over which $C_1, \ldots, C_n$ are all defined. The set $\prod_{i=1}^n C_i^N$ has dp-rank $N \cdot n$, so we can find some tuple in it, having dp-rank $N \cdot n$ over A, hence over $\emptyset$. By mediocrity, we can find some $U \in \mathcal{U}$ that U avoids this entire tuple. By choice of N, now each $C_i \setminus U$ is infinite. □

Theorem 9.3.16. *There are only a bounded number of germs at 0 among definable subsets of $\mathbb{M}$.*

Proof. Suppose not.

Claim 9.3.17. *There is some sequence $X_1, X_2, \ldots$ of definable subsets of $\mathbb{M}^\times$, all belonging to a single definable family, such that $0 \in \overline{X_i}$ and $0 \notin \overline{X_i \cap X_j}$ for $i \neq j$.*

Proof. By Morley-Erdős-Rado, we can produce an indiscernible sequence of sets

$$Y_1, Y_2, Y_3, \ldots \subseteq \mathbb{M}$$

having pairwise distinct germs at 0. Let $X_i = Y_{2i} \triangle Y_{2i+1}$; then $0 \in \overline{X_i}$. By indiscernibility, 0 is in every Y_i or in none; either way each $X_i \subseteq \mathbb{M}^\times$.

By NIP, the collection $\{X_i\}$ is k-inconsistent for some k. Replace X_i with $X_{2i} \cap X_{2i+1}$ until $0 \notin \overline{X_1 \cap X_2}$. This process must terminate within $\log_2 k$ steps or so. □

Fix $X_1, X_2, \ldots$ from the claim. Let K_1 be a small model over which the X_i are defined. Let $\mathcal{U}$ be a mediocre famiy from Proposition 9.3.12. Inductively build a sequence $K_1 \preceq K_2 \preceq \cdots$ and $U_1, U_2, \ldots \in \mathcal{U}$ as follows:

- U_i is chosen so that $C \setminus U_i$ is infinite for every infinite K_i-definable set $C \subseteq \mathbb{M}$. This is possible by Lemma 9.3.15.
- K_{i+1} is chosen so that U_i is K_{i+1}-definable.

Claim 9.3.18. *For any i_0, j_0, there is some a such that $a \in X_i \iff i = i_0$, and $a \in U_j \iff j < j_0$.*

Proof. By compactness, it suffices to only consider $X_1, \ldots, X_n$ and $U_1, \ldots, U_n$. Let

$$D = X_1^c \cap X_2^c \cap \cdots \cap X_{i_0-1}^c \cap X_{i_0} \cap X_{i_0+1}^c \cap \cdots \cap X_n^c$$

where S^c denotes the complement $\mathbb{M} \setminus S$ of a set S.

The set D is K-definable, and $0 \in \overline{D} \setminus D$, by choice of the X_i's. So the set

$$S = D \cap U_1 \cap \cdots \cap U_{j_0-1}$$

is infinite, as $U_1 \cap \cdots \cap U_{j_0-1}$ is a neighborhood of 0.

As S is K_{j_0} definable, it follows that $S \cap U_{j_0}^c$ is infinite, by choice of U_{j_0}. As $S \cap U_{j_0}^c$ is K_{j_0+1}-definable, it follows that $S \cap U_{j_0}^c \cap U_{j_0+1}^c$ is infinite. Continuing on in this fashion, we ultimately see that

$$S \cap U_{j_0}^c \cap \cdots \cap U_n^c$$

is infinite. If a is any element of this set, then $a \in D$, so $a \in X_i \iff i = i_0$ (for $1 \leq i \leq n$), and

$$a \in U_1 \cap \cdots \cap U_{j_0-1} \cap U_{j_0}^c \cap \cdots \cap U_n^c,$$

so $a \in U_j \iff j < j_0$ (for $1 \leq j \leq n$).

Finally, using compactness, we can send n to ∞. □

Given the claim, the sets $\{X_i\}$ and $\{U_i \setminus U_{i+1}\}$ now directly contradict dp-minimality. □

Corollary 9.3.19. *There are only a bounded number of infinitesimal types over* $\mathbb{M}$.

By Lemma 9.3.5, it follows that infinitesimal types have boundedly many heirs, and so...

Corollary 9.3.20. *Infinitesimal types are definable.*

9.3.4 Multiplying Infinitesimals

Using Theorem 9.3.16 and Corollary 9.3.19, we can complete the proof of Theorem 9.1.3.

We will repeatedly make use of the following basic observation:

Observation 9.3.21. *Let $X \subseteq \mathbb{M}$ be K-definable, and $a \in K$. Then the following are (clearly) equivalent:*

1. *There is a K-infinitesimal ϵ such that $(a + \epsilon \in X \not\Leftrightarrow a \in X)$.*
2. *The type $\Sigma(x)$ asserting that $x \in I_K$ and $(a + x \in X \not\Leftrightarrow a \in X)$ is consistent.*
3. *For every K-definable basic neighborhood U, the set $a + U$ intersects both X and $X^c := \mathbb{M} \setminus X$.*
4. *a is in the topological boundary of $X(K)$ within K.*

Note that the third of these conditions does not depend on K, in the sense that its truth is unchanged if we replace K with an elementary extension $K' \succeq K$.

First we show that definable sets have finite boundaries.

Proposition 9.3.22. *If $X \subseteq K$ is definable, then ∂X is finite, and contained in* $\mathrm{acl}(\ulcorner X \urcorner)$.

Proof. By Observation 9.3.21, we may replace K with $\mathbb{M}$—this only enlarges ∂X.

The set ∂X is type-definable, essentially by (3) of Observation 9.3.21. It is also type-definable over $\mathrm{dcl}(\ulcorner X \urcorner)$, by automorphism invariance of the topology. The proposition will therefore follow if ∂X is small.

Let $\mathbb{M}^*$ be a sufficiently saturated elementary extension of $\mathbb{M}$. By the equivalence of conditions 1 and 4 of Observation 9.3.21,

$$\partial X(\mathbb{M}) = \bigcup_{\epsilon \in I_{\mathbb{M}}} \{x \in \mathbb{M} : x + \epsilon \in X \nLeftrightarrow x \in X\} \tag{9.1}$$

Let D_ϵ denote $\{x \in \mathbb{M} : x + \epsilon \in X \nLeftrightarrow x \in X\}$. By the first part of Proposition 9.3.7, each D_ϵ is finite. Moreover, D_ϵ depends only on $\mathrm{tp}(\epsilon/\mathbb{M})$. By Corollary 9.3.19, it follows that the right hand side of (9.1) is small. □

Lemma 9.3.23. *The map $x \mapsto a \cdot x$ is K-slight if and only if $a - 1$ is a K-infinitesimal.*

Proof. Recall from Proposition 9.3.7 that the K-slight maps are closed under inversion, composition, and conjugation by K-definable maps.

First suppose $f(x) = ax$ is K-slight. As $g(x) = x+1$ is K-definable, the map $f \circ g \circ f^{-1} \circ g^{-1}$ is K-slight. But

$$(f \circ g \circ f^{-1} \circ g^{-1})(x) = \left(\frac{x-1}{a} + 1\right) \cdot a = x + (a - 1)$$

so $a - 1$ is a K-infinitesimal.

Conversely, suppose $a - 1$ is a K-infinitesimal; we will show that $x \mapsto ax$ is K-slight. Let X be an infinite K-definable set; we will show that $X \cap a^{-1}X$ is infinite. In fact, it contains $X(K) \setminus \partial X$, which is infinite by Proposition 9.3.22. To see this, suppose $x \in X(K) \setminus \partial X$. Then $(a - 1)x$ is K-infinitesimal by Corollary 9.3.8. By the equivalence of 1 and 4 in Observation 9.3.21 and the fact that $a \notin \partial X$, it follows that

$$ax = x + (a - 1)x \in X \text{ for any } x \in X(K) \setminus \partial X$$

Thus $X(K) \setminus \partial X \subseteq X \cap a^{-1}X$. □

Corollary 9.3.24. *$1 + I_K$ is a subgroup of $\mathbb{M}^\times$.*

If G is a type-definable group, G^{00} denotes the smallest type-definable group of bounded index. These are known to exist in NIP theories, by Proposition 6.1 in [35].

Lemma 9.3.25. *As a subgroup of the additive group, I_K has no type-definable proper subgroups of bounded index. In other words, $I_K^{00} = I_K$.*

Proof. Suppose for the sake of contradiction that there is $\epsilon \in I_K \setminus I_K^{00}$. Let K' be a model containing ϵ, and let ϵ' realize an heir of $\mathrm{tp}(\epsilon/K)$ to K'. By Lemma 9.3.5, ϵ' is K'-infinitesimal.

As ϵ and ϵ' have the same (Lascar strong) type over K, they are in the same coset of I_K^{00}. Then ϵ and $\epsilon - \epsilon'$ do *not* have the same type over K, because the latter is in I_K^{00} but the former is not. Choose a K-definable set X which contains ϵ but not $\epsilon - \epsilon'$. As X is K'-definable and $\epsilon \in K'$, it follows by Observation 9.3.21 that $\epsilon \in \partial X$. Then by Proposition 9.3.22, $\epsilon \in \mathrm{acl}(\ulcorner X \urcorner) \subseteq K$, which is absurd, since ϵ is a non-zero K-infinitesimal. □

Next we prove a general fact about dp-minimal groups.

Lemma 9.3.26. *Suppose G and H are type-definable subgroups of $(K, +)$, such that $G = G^{00}$ and $H = H^{00}$. Then $G \subseteq H$ or $H \subseteq G$.*

Proof. Otherwise, $G \cap H$ has unbounded index in both G and H. By Morley-Erdős-Rado we can produce an indiscernible sequence $\langle (a_i, b_i) \rangle_{i < \omega + \omega}$ of elements of $G \times H$ such that the a_i are in pairwise distinct cosets of $G \cap H$, and the b_i are in pairwise distinct cosets of $G \cap H$. The sequences $a_0, a_1, \ldots$ and $b_\omega, b_{\omega+1}, \ldots$ are mutually indiscernible. However, after naming $c := a_0 + b_\omega$, neither sequence is indiscernible. Indeed, $c - a_i \in H$ if and only if $i = 0$, and $b_{\omega+i} - c \in G$ if and only if $i = 0$. This contradicts the characterization of dp-rank 1 in terms of mutually indiscernible sequences. □

Corollary 9.3.27. *For any $a \in \mathbb{M}$, either $a \cdot I_K \subseteq I_K$ or $I_K \subseteq a \cdot I_K$.*

Combining Theorem 9.2.4, Proposition 9.3.2, Corollary 9.3.8, Corollary 9.3.24, and Corollary 9.3.27, we see that

Theorem 9.3.28. *The K-infinitesimals I_K are the maximal ideal of a valuation ring $\mathcal{O}_K$. The canonical topology on K is a non-discrete definable V-topology, and every definable set has finite boundary.*

As an exercise, one can check that the valuation topology on $\mathbb{M}$ induced by $\mathcal{O}_K$ is the canonical topology on $\mathbb{M}$. We will prove the following consequence of that fact:

Lemma 9.3.29. *I_K is open in the canonical topology on $\mathbb{M}$.*

Proof. Because I_K is a subgroup of $(\mathbb{M}, +)$, it suffices to show that I_K is a neighborhood of 0. Note that I_K is type-definable and $\mathcal{O}_K$ is $\vee$-definable, both over K. Therefore we can find a K-definable set B lying between them:

$$I_K \subseteq B \subseteq \mathcal{O}_K$$

By directedness of the family of K-definable basic neighborhoods, there is a K-definable basic neighborhood $X - X$ such that

$$I_K \subseteq X - X \subseteq B \subseteq \mathcal{O}_K$$

Now choose some non-zero $\epsilon \in I_K$. Then

$$(\epsilon \cdot X) - (\epsilon \cdot X) = \epsilon \cdot (X - X) \subseteq \epsilon \cdot \mathcal{O}_K \subseteq I_K,$$

so I_K contains a neighborhood of 0 in the canonical topology on $\mathbb{M}$. □

9.4 Definability of dp-rank and henselianity

9.4.1 Interior and Dp-rank

In this section, we show that a definable subset of $\mathbb{M}^n$ has dp-rank n if and only if it has non-empty interior (in the product topology on $\mathbb{M}^n$.

Lemma 9.4.1. *Naming infinitesimals does not algebraize anything, in the following sense:*

1. *Let $\mathbb{M}^* \succeq \mathbb{M}$ be an elementary extension, and $\epsilon \in \mathbb{M}^*$ be $\mathbb{M}$-infinitesimal. For any small $S \subseteq \mathbb{M}$, we have $\mathbb{M} \cap \operatorname{acl}(S\epsilon) = \operatorname{acl}(S)$.*

2. *Let p be an infinitesimal type over $\mathbb{M}$. Suppose $S \subseteq \mathbb{M}$ is small, $a \in \mathbb{M}$, and $\epsilon \models p|Sa$. Then $a \in \operatorname{acl}(S) \iff a \in \operatorname{acl}(S\epsilon)$.*

Proof. 1. Fix S. For $\epsilon \in I_{\mathbb{M}}$, let $X_\epsilon = \operatorname{acl}(S\epsilon) \cap \mathbb{M}$. Then X_ϵ is small and depends only on $\operatorname{tp}(\epsilon/\mathbb{M})$. By Corollary 9.3.19, it follows that $\bigcup_{\epsilon \in I_{\mathbb{M}}} X_\epsilon$ is small. It is also $\operatorname{Aut}(\mathbb{M}/S)$-invariant, so it must be contained in $\operatorname{acl}(S)$. In particular, $X_\epsilon \subseteq \operatorname{acl}(S)$ for any $\mathbb{M}$-infinitesimal ϵ.

2. Let $\mathbb{M}^* \succeq \mathbb{M}$ be an elementary extension in which p is realized by some ϵ'. Then $\epsilon' \equiv_{aS} \epsilon$, so
$$a \in \operatorname{acl}(S\epsilon) \iff a \in \operatorname{acl}(S\epsilon') \iff a \in \operatorname{acl}(S)$$
where the second equivalence follows by the previous point.

□

Say that a tuple $(a_1, \ldots, a_n)$ is *algebraically independent* over a set S if
$$a_i \notin \operatorname{acl}(a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n, S)$$
for each i.

Corollary 9.4.2. *Suppose $a_1, \ldots, a_n$ are algebraically independent over S, and $\epsilon \models p|S\vec{a}$ for some non-trivial global infinitesimal type p. Then $a_1, \ldots, a_n, \epsilon$ is algebraically independent over S.*

Note that $\epsilon \notin \operatorname{acl}(S\vec{a})$ because all non-zero infinitesimal types are non-algebraic.

Proposition 9.4.3. *For a K-definable set $X \subseteq \mathbb{M}^n$, the following are equivalent:*

1. *X has dp rank n*

2. *X contains a tuple which is algebraically independent over K*

3. *X has non-empty interior, in the product topology on $\mathbb{M}^n$.*

Proof. The implication $3 \implies 1$ is clear, because the product topology has a basis consisting of n-fold products of sets of dp rank 1.

For $1 \implies 2$, suppose X has dp rank n. Choose $\vec{a} = (a_1, \ldots, a_n)$ in X with dp rank n over K. Then the a_i's are algebraically independent over K. Suppose otherwise. Then, say, $a_1 \in \operatorname{acl}(a_2 a_3 \ldots S)$. By subadditivity of dp-rank,

$$\text{dp-rk}(\vec{a}/S) \leq \text{dp-rk}(a_1/a_2 a_3 \ldots S) + \text{dp-rk}(a_2 a_3 \ldots /S) \leq 0 + (n-1)$$

contradicting the choice of $\vec{a}$.

The hard part is $2 \implies 3$. Suppose $(a_1, \ldots, a_n) \in X$ is algebraically independent over K. Let p be a nonzero global infinitesimal type. Let $\epsilon_1, \ldots, \epsilon_n$ realize $p^{\otimes n}$ over $K\vec{a}$. By Corollary 9.4.2, the tuple $(a_1, \ldots, a_n, \epsilon_1, \ldots, \epsilon_n)$ is algebraically independent over K.

The type-definable set I_K^{-1} is the complement of $\mathcal{O}_K$, so $\mathcal{O}_K$ is $\vee$-definable. By compactness, we can find a K-definable set B between I_K and $\mathcal{O}_K$:

$$I_K \subseteq B \subseteq \mathcal{O}_K$$

Claim 9.4.4. *Let S be a small set containing K, let U be an S-definable neighborhood of 0, and let $\epsilon \models p|S$. Then $\epsilon \cdot B \subseteq U$.*

Proof. Let M be a small model containing S. Moving M over S we may assume $\epsilon \models p|M$. There are more M-definable basic neighborhoobds than K-definable basic neighborhoods, so $I_M \subseteq I_K$, which in turn implies $\mathcal{O}_K \subseteq \mathcal{O}_M$. Thus

$$I_M \subseteq I_K \subseteq B \subseteq \mathcal{O}_K \subseteq \mathcal{O}_M$$

Now $\epsilon \models p|M$, so ϵ is an M-infinitesimal. Thus

$$\epsilon \cdot B \subseteq \epsilon \cdot \mathcal{O}_M \subseteq I_M \subseteq U.$$

□

Claim 9.4.5. *For $0 \leq k \leq n$, we have*

$$(a_1 + \epsilon_1 \cdot B) \times \cdots \times (a_k + \epsilon_k \cdot B) \times \{(a_{k+1}, \ldots, a_n)\} \subseteq X$$

Proof. We proceed by induction on k. The base case $k = 0$ simply says $\vec{a} \in X$, which is given.

Suppose

$$(a_1 + \epsilon_1 \cdot B) \times \cdots \times (a_k + \epsilon_k \cdot B) \times \{(a_{k+1}, \ldots, a_n)\} \subseteq X$$

Let Y be the set of y such that

$$(a_1 + \epsilon_1 \cdot B) \times \cdots \times (a_k + \epsilon_k \cdot B) \times \{(y, \ldots, a_n)\} \subseteq X$$

Then $a_{k+1} \in Y$, of course. By independence of $\vec{a}\vec{\epsilon}$, the element a_{k+1} is not algebraic over the parameters used to define Y. Therefore, by Proposition 9.3.22, $a_{k+1} \in Y^{int}$. So $Y - a_{k+1}$ is

a neighborhood of 0. Since ϵ_{k+1} is infinitesimal over the parameters used to define $Y - a_{k+1}$, we see

$$\epsilon_{k+1} \cdot B \subseteq Y - a_{k+1}$$

by Claim 9.4.4. But this is equivalent to

$$(a_1 + \epsilon_1 \cdot B) \times \cdots \times (a_{k+1} + \epsilon_{k+1} \cdot B) \times \{(a_{k+2}, \ldots, a_n)\} \subseteq X$$

completing the inductive step. □

Taking $k = n$ in the claim, we see that

$$\prod_{i=1}^{n} (a_i + \epsilon \cdot B) \subseteq X$$

and so $\vec{a} \in X^{int}$, completing the proof. □

Corollary 9.4.6. *Let $f : \mathbb{M}^n \to \mathbb{M}^n$ be a finite-to-one definable map. If $X \subseteq \mathbb{M}^n$ is a set with non-empty interior, then $f(X)$ also has non-empty interior. (We are not assuming X is definable.)*

Proof. We may assume X is definable, by shrinking it (recall the basis of definable opens). Take $\vec{a}$ in X of full dp-rank. Then $f(\vec{a})$ is interalgebraic with a, so also has dp-rank n. Therefore $f(X)$ has dp-rank n, so it has interior by the proposition. □

In the next section, we will use this fact to show that $\mathcal{O}_K$ is henselian.

For completeness, we prove the following fact, which will not be used in what follows, but could be of independent interest:

Corollary 9.4.7. *If $X \subseteq \mathbb{M}^n$ is definable and $m \leq n$, then* dp-rk$(X) \geq m$ *if and only if there is some coordinate projection $\pi : \mathbb{M}^n \to \mathbb{M}^m$ such that $\pi(X)$ has non-empty interior. Consequently, dp-rank of definable sets is definable in families.*

Proof. We may assume X is K-definable. If $\pi(X)$ has non-empty interior, then dp-rk$(X) \geq$ dp-rk$(\pi(X)) = m$. Conversely, suppose dp-rk$(X) \geq m$. Let $\vec{a} \in X$ have dp-rank at least m over X. If the n-tuple $\vec{a}$ is not algebraically independent over K, then some a_i is algebraic over K and the other coordinates of $\vec{a}$. So there is some $n-1$-tuple $\vec{a}'$ which is a subtuple of $\vec{a}$, and which is interalgebraic with $\vec{a}$ over K. Iterating this process, we get a sequence

$$\vec{a} = \vec{a}^{(0)}, \vec{a}^{(1)}, \ldots, \vec{a}^{(k)}$$

for some $k \geq 0$, such that

- $\vec{a}^{(i+1)}$ is a subtuple of $\vec{a}^{(i)}$
- $\vec{a}^{(i+1)}$ is interalgebraic with $\vec{a}^{(i)}$ over k

- $|\vec{a}^{(i)}| = n - i$
- $\vec{a}^{(k)}$ is algebraically independent over K.

Then $\vec{a}^{(k)}$ is interalgebraic over K with $\vec{a}$, so

$$\text{dp-rk}(\vec{a}^{(k)}/K) = \text{dp-rk}(\vec{a}/K) \geq m$$

Let $\vec{b}$ be a subtuple of $\vec{a}$ of length m. Then $\vec{b}$ is a subtuple of $\vec{a}$, and $\vec{b}$ is algebraically independent over K. Write $\vec{b} = \pi(\vec{a})$ for some coordinate projection $\pi : \mathbb{M}^n \to \mathbb{M}^m$. Then

$$\vec{b} \in \pi(X)$$

so the K-definable set $\pi(X)$ contains an algebraically independent tuple, and therefore has non-empty interior. □

9.4.2 Henselianity

First we prove a general fact about $\vee$-definable valuation rings: their prolongations to finite extension fields are still $\vee$-definable.

Lemma 9.4.8. *Let F be a field with some structure, and L/F be a finite extension. Suppose $\mathcal{O}$ is a $\vee$-definable valuation ring on F. Then each extension of $\mathcal{O}$ to L is $\vee$-definable (over the same parameters used to define $\mathcal{O}$ and interpret L).*

Proof. Replacing L with the normal closure of L over F, we may assume L/F is a normal extension of some degree n.

Claim 9.4.9. *There is some $d = d(k, n)$ such that for $\{a_1, \ldots, a_k\} \subseteq L$, the following are equivalent:*

- *No extension of $\mathcal{O}$ to L contains $\{a_1, \ldots, a_k\}$.*
- *$1 = P(a_1, \ldots, a_k)$ for some polynomial $P(X_1, \ldots, X_k) \in \mathfrak{m}[X_1, \ldots, X_k]$ of degree less than $d(k, n)$.*

Proof. Consider the theory T_n whose models consist of degree n normal field extensions L/F with a predicate picking out a valuation ring $\mathcal{O}_L$ on L. On general valuation-theoretic grounds, the following are equivalent for $\{a_1, \ldots, a_k\} \subseteq L$

- $\{a_1, \ldots, a_k\} \not\subseteq \sigma(\mathcal{O}_L)$ for any $\sigma \in \text{Aut}(L/F)$.
- No extension of $\mathcal{O}_L \cap F$ to L contains $\{a_1, \ldots, a_k\}$.
- $1 = P(a_1, \ldots, a_k)$ for some $P(X_1, \ldots, X_k) \in \mathfrak{m}[X_1, \ldots, X_k]$.

The first condition is a definable condition on the k-tuple $(a_1, \ldots, a_k)$, so by compactness applied to T_n, there is a bound on the degree in the third condition. □

Because $\mathcal{O}$ is $\vee$-definable, $\mathfrak{m}$ is type-definable, so the second condition in the claim is type-definable.

Let $\mathcal{O}'$ be some extension of $\mathcal{O}$ to L. We can find some finite set $S \subseteq \mathcal{O}'$ such that $\mathcal{O}'$ is the unique extension of $\mathcal{O}$ containing S, because there are only finitely many extensions and they are pairwise incomparable. The claim implies type-definability of the set

$$\{x \in L : \text{no extension of } \mathcal{O} \text{ to } L \text{ contains } S \cup \{x\}\}$$

which is the complement of $\mathcal{O}'$ by choice of S. □

Recall that $\mathcal{O}_K$ denotes the valuation ring whose maximal ideal is I_K.

Proposition 9.4.10. *Let K be a small submodel of $\mathbb{M}$. Let L/K be a finite algebraic extension, and $\mathbb{L} = L \otimes_K \mathbb{M}$. (So $\mathbb{L}$ is a saturated elementary extension of L.) Then $\mathcal{O}_K$ has a unique extension to $\mathbb{L}$.*

Proof. We give the proof in characteristic $\neq 2$.

Replacing L with its normal closure over K, we may assume L/K is normal.

Let $\mathcal{O}_1, \ldots, \mathcal{O}_m$ denote the extensions of $\mathcal{O}_K$ to $\mathbb{L}$. By Lemma 9.4.8, these are all $\vee$-definable over K. Let $\mathfrak{m}_i$ be the maximal ideal of $\mathcal{O}_i$; this is type-definable over K. Let v_i be the valuation on $\mathcal{O}_i$.

Write $L = K(\alpha)$ (possible because K is perfect by Observation 9.2.9). So $\mathbb{L} = \mathbb{M}(\alpha)$ and $\{1, \alpha, \ldots, \alpha^{n-1}\}$ is a basis for $\mathbb{L}$ over $\mathbb{M}$.

Claim 9.4.11. $\bigcap_i \mathfrak{m}_i = \sum_{i=0}^{n-1} I_K \cdot \alpha^i$

Proof. Let $(F, \mathcal{O})$ be some algebraically closed valued field extending $(\mathbb{M}, \mathcal{O}_K)$, and let $\mathfrak{m}$ be the maximal ideal of $\mathcal{O}$. All the extensions of $\mathcal{O}_K$ to $\mathbb{L}$ come from embeddings of $\mathbb{L}$ into F, so if $\iota_1, \ldots, \iota_n$ denote the embeddings of $\mathbb{L}$ into F, then

$$\{\mathfrak{m}_1, \ldots, \mathfrak{m}_m\} = \{\iota_i^{-1}(\mathfrak{m}) : 1 \leq i \leq n\}$$

Thus

$$\bigcap_i \mathfrak{m}_i = \bigcap_{i=1}^{n} \iota_i^{-1}(\mathfrak{m})$$

Because $K \subseteq \mathcal{O}$, it follows that $K^{alg} \subseteq \mathcal{O}$, where K^{alg} is the algebraic closure of K inside F. Let $\alpha_1, \ldots, \alpha_n$ be the images of α under $\iota_1, \ldots, \iota_n$. These are pairwise distinct because $\mathbb{L}/\mathbb{M}$ is separable (by Observation 9.2.9 again). Let M be the Vandermonde matrix whose (i, j) entry is α_i^{j-1}. Then $M \in GL_n(K^{alg}) \subseteq GL_n(\mathcal{O})$.

It follows that multiplication by M and M^{-1} preserves $\mathfrak{m}^n \subseteq F^n$. Concretely, this means that if $(x_0, x_1, \ldots, x_{n-1}) \in F^n$, then the following are equivalent:

- Each $x_i \in \mathfrak{m}$
- $\sum_{i=0}^{n-1} x_i \alpha_j^i \in \mathfrak{m}$ for each j.

Specializing to the case where $x_0, \ldots, x_{n-1} \in \mathbb{M}$, and writing $x = \sum_{i=0}^{n-1} x_i \alpha^i$, the following are equivalent:

- Each $x_i \in I_K$
- $\iota_j(x) \in \mathfrak{m}$ for each $j \leq n$, or equivalently, $x \in \mathfrak{m}_i$ for each $i \leq m$.

□

Our goal is to show $m = 1$. Suppose for the sake of contradiction that $m > 1$. Because $\mathrm{Aut}(L/K)$ acts transitively on the $\mathcal{O}_i$'s, they are pairwise incomparable. By the approximation theorem for valuations ([6] VI.7.1 Corollaire 1), we can find an element $x \in \mathbb{L}$ such that $x \in 1 + \mathfrak{m}_1$ and $x \in -1 + \mathfrak{m}_i$ for $i > 1$.

Let $I = \bigcap_i \mathfrak{m}_i$. Then

$$\begin{aligned} x &\notin 1 + I \\ -x &\notin 1 + I \\ x^2 &\in 1 + I \end{aligned}$$

By basic valuation theory, each $1 + \mathfrak{m}_i$ is a subgroup of $\mathbb{M}^\times$. The intersection $1 + I$ is therefore also a subgroup of $\mathbb{M}^\times$. The intersection $1 + I$ is also topologically open: by the Claim

$$1 + I = (1 + I_K) + I_K \cdot \alpha + I_K \cdot \alpha^2 + \cdots + I_K \cdot \alpha^{n-1},$$

and I_K is open by Lemma 9.3.29.

The squaring map on $\mathbb{L}^\times$ is finite-to-one, so by Proposition 9.4.6, $(1 + I)^2$ has interior. Since $(1 + I)^2$ is a group, it is actually open, hence contains a neighborhood of 1:

$$(1 + I)^2 \text{ is a neighborhood of } 1 \tag{9.2}$$

Now $x \notin 1 + I$ and $-x \notin 1 + I$, and I is type-definable over K. So there is some K-definable set U containing I, such that $x \notin 1 + U$ and $-x \notin 1 + U$. By (9.2), $(1 + U)^2$ is a neighborhood of 0. It is K-definable, so it contains $1 + I$, hence x^2. Then there is $y \in 1 + U$ such that $y^2 = x^2$. Either $x \in 1 + U$ or $-x \in 1 + U$, contradicting the choice of U.

If K has characteristic 2, replace -1 and 1 with 0 and 1, replace the squaring map with the Artin-Schreier map, and replace $1 + I < \mathbb{L}^\times$ with $I < \mathbb{L}$. □

Lemma 9.4.12. *If $\mathcal{O}$ is a non-trivial K-definable valuation ring on $\mathbb{M}$, then $\mathcal{O}_K$ is a coarsening of $\mathcal{O}$. Moreover, the canonical topology on K is induced by the valuation ring $\mathcal{O}(K)$ on K.*

Proof. The maximal ideal $\mathfrak{m}$ of $\mathcal{O}$ is infinite by non-triviality of $\mathcal{O}$, and so $\mathfrak{m}$ has interior by Proposition 9.3.22. As $\mathfrak{m}$ is a subgroup of the additive group, $\mathfrak{m}$ must be open, so 0 is in the interior of the K-definable set $\mathfrak{m}$. Consequently $I_K \subseteq \mathfrak{m}$, which directly implies that $\mathcal{O} \subseteq \mathcal{O}_K$.

Moreover, because

$$I_K \subseteq \mathcal{O} \subseteq \mathcal{O}_K,$$

Lemma 9.2.2 applies and so

$$\{a \cdot \mathcal{O}(K) : a \in K^\times\}$$

is a neighborhood basis of 0 in the canonical topology on K.

□

Remark 9.4.13. *Suppose F is a field with some structure, and $\mathcal{O}_1$ and $\mathcal{O}_2$ are incomparable $\vee$-definable valuation rings on F. Then the join $\mathcal{O}_1\mathcal{O}_2$ is* definable.

Proof. The join can be written as either $\{x \cdot y : x \in \mathcal{O}_1, y \in \mathcal{O}_2\}$ (which is $\vee$-definable) or as $\{x \cdot y : x \in \mathfrak{m}_1, y \in \mathfrak{m}_2\}$, which is type-definable.[3] □

Lemma 9.4.14. *Let $\mathbb{L}/\mathbb{M}$ be a finite algebraic extension. Any two non-trivial definable valuation rings on $\mathbb{L}$ are* not *independent, i.e., they induce the same topology.*

Proof. Let w_1, w_2 be two definable valuations on $\mathbb{L}$, and let v_1 and v_2 be their restrictions to $\mathbb{M}$. Let Γ_i be the value group of w_i. Let K be a small model over which everything is defined (including the extension $\mathbb{L}/\mathbb{M}$). Let v_K be the non-definable valuation on $\mathbb{M}$ coming from $\mathcal{O}_K$ and I_K. By Lemma 9.4.12, v_K is a coarsening of v_1 and v_2. So there are convex subgroups $\Delta_i < \Gamma_i$ such that v_K is equivalent to the coarsening of v_i by Δ_i. Let w_i' be the coarsening of w_i by Δ_i. Then w_1' and w_2' are valuations on $\mathbb{L}$ extending v_K. By Proposition 9.4.10, w_1' and w_2' are equivalent (because v_K has an essentially unique extension). It follows that w_1 and w_2 have a common coarsening—the unique extension of v_K to $\mathbb{L}$. This common coarsening is non-trivial, because v_K is non-trivial. Non-trivial coarsenings induce the same topology, so w_1, w_1', and w_2 all induce the same topology. Therefore w_1 and w_2 are not independent. □

Proposition 9.4.15. *Let $\mathbb{L}$ be a finite extension of $\mathbb{M}$. Any two definable valuation rings on $\mathbb{L}$ are comparable.*

Proof. Suppose $\mathcal{O}_1$ and $\mathcal{O}_2$ are incomparable. Let $\mathcal{O} = \mathcal{O}_1 \cdot \mathcal{O}_2$ be their join, which is definable by Remark 9.4.13. Let w be the valuation corresponding to $\mathcal{O}$, and let v be its restriction to $\mathbb{M}$.

The residue field $\mathbb{L}' := \mathbb{L}w$ is a finite extension of $\mathbb{M}' := \mathbb{M}v$. Moreover, $\mathbb{L}'$ has two independent definable valuations, induced by $\mathcal{O}_1$ and $\mathcal{O}_2$. This ensures that $\mathbb{L}'$ is infinite and unstable, so $\mathbb{M}'$ is also infinite and unstable. But $\mathbb{M}'$ has dp-rank at most 1, so $\mathbb{M}'$

[3]Here, we are using the fact that if $\mathcal{O}$ is a valuation ring with maximal ideal $\mathfrak{m}$, and S is any set, then $S \cdot \mathcal{O}$ and $S \cdot \mathfrak{m}$ are closed under addition, and are equal to each other unless S has an element of minimum valuation. Incomparability of $\mathcal{O}_1$ and $\mathcal{O}_2$ ensures that e.g. $v_1(\mathcal{O}_2)$ has no minimum.

is a dp-minimal unstable field. It is also as saturated as $\mathbb{M}$, so all our results so far apply to $\mathbb{M}'$. By Lemma 9.4.14, $\mathbb{L}'$ cannot have two independent definable valuation rings, a contradiction. □

Corollary 9.4.16. *Any definable valuation ring $\mathcal{O}$ on $\mathbb{M}$ is henselian.*

Proof. Otherwise, $\mathcal{O}$ would have two incomparable extensions to some finite Galois extension of $\mathbb{M}$. □

Corollary 9.4.16 was obtained independently by Jahnke, Simon, and Walsberg (Proposition 4.5 in [39]).

Theorem 9.4.17. *The valuation ring $\mathcal{O}_K$ (whose maximal ideal is the set of K-infinitesimals)* is henselian.

Proof. Suppose not. Then $\mathcal{O}_K$ has multiple extensions to some finite algebraic extension $\mathbb{L}/\mathbb{M}$. Let $\mathcal{O}_1$ and $\mathcal{O}_2$ be two such extensions. Let $K' \succeq K$ be a larger model over which the field extension $\mathbb{L}/\mathbb{M}$ is defined. As $I_{K'} \subseteq I_K$, we see that $\mathcal{O}_{K'}$ is a coarsening of $\mathcal{O}_K$. Also, $\mathcal{O}_{K'}$ has a unique extension to $\mathbb{L}$ by Proposition 9.4.10. As in the proof of Lemma 9.4.14, this ensures that $\mathcal{O}_1$ and $\mathcal{O}_2$ are not independent. Their join $\mathcal{O}_1 \cdot \mathcal{O}_2$ is definable by Lemma 9.4.8 and Remark 9.4.13. It is also non-trivial because $\mathcal{O}_1$ and $\mathcal{O}_2$ are not independent.

So there is some definable non-trivial valuation ring on $\mathbb{M}$. The property of being a valuation ring is expressed by finitely many sentences, and $K \preceq \mathbb{M}$, so there is a K-definable non-trivial valuation ring $\mathcal{O}$. This ring is henselian by Corollary 9.4.16, and $\mathcal{O}_K$ is a coarsening, by Lemma 9.4.12. Coarsenings of henselian valuations are henselian. □

9.4.3 Summary of results so far

In what follows, we will need only the following facts from §9.3.1-§9.4.2:

Theorem 9.4.18. *Let K be a dp-minimal field.*

1. *K is perfect.*
2. *If K is infinite, sufficiently saturated, and not algebraically closed, then K admits a non-trivial Henselian valuation (not necessarily definable).*
3. *Any definable valuation on K is henselian. Any two definable valuations on K are comparable.*
4. *For any n, the cokernel of the nth power map $K^\times \to K^\times$ is finite.*

Proof.

1. Observation 9.2.9.

2. If K is strongly minimal, then K is algebraically closed by a well-known theorem of Macintyre. Otherwise, this is Theorem 9.4.17.

3. If K is not strongly minimal, this is Proposition 9.4.15 and Corollary 9.4.16. Otherwise, K is NSOP, so has only the trivial valuation.

4. If K is strongly minimal, then K is algebraically closed (Macintyre), so the cokernels are always trivial. If $K^\times/(K^\times)^n$ is infinite, we can find some elementary extension $M \succeq K$ such that $M^\times/(M^\times)^n$ is greater in cardinality than the total number of infinitesimal types over $\mathbb{M}$, by Corollary 9.3.19. By Lemma 9.3.5, heirs of infinitesimal types are infinitesimal types, so $\mathbb{M}$ has at least as many infinitesimal types as M, and therefore the cardinality of $M^\times/(M^\times)^n$ exceeds the number of infinitesimal types over M. Now for any $a \in M^\times$, and any M-infinitesimal ϵ, the element $a \cdot \epsilon^n$ is an M-infinitesimal in the same coset as a. So there are M-infinitesimals in every coset of $(M^\times)^n$, contradicting the choice of M.

□

9.5 The Canonical Valuation

We now turn to proving Theorem 9.1.4. This relies crucially on [38]'s work on defining canonical valuations.

9.5.1 Review of Jahnke-Koenigsmann

Fix a prime p. Say that a profinite group is p-nilpotent if every finite quotient is a p-group. Following [38], if K is any field, let $K(p)$ denote the compositum of all finite Galois extensions L/K with $\mathrm{Gal}(L/K)$ a p-nilpotent group. The map $K \mapsto K(p)$ is a closure operation on fields; abusing terminology we will call $K(p)$ the "p-closure" of K.

A valuation v on K is *p-henselian* if it has a unique extension to $K(p)$. This is a weakening of henselianity. On any field K there is a *canonical p-henselian* valuation v_K^p, which might be trivial. It has the following properties:

Fact 9.5.1.

1. *If the residue field Kv_K^p is not p-closed, then v_K^p is the finest p-henselian valuation on K.*

2. *Every p-henselian valuation strictly finer than v_K^p has p-closed residue field.*

3. *If K admits no orderings and contains the pth roots of unity, then the valuation ring of v_K^p is 0-definable in K from the field language.*

Say that a field K is "p-corrupted" if no finite extension is p-closed.

Lemma 9.5.2. *Let K be a perfect field which is neither algebraically nor real closed. Then some finite extension of K is p-corrupted.*

Proof. Replace K with $K(\sqrt{-1})$ in characteristic 0. Take some non-trivial finite Galois extension L/K. Take p dividing $|\mathrm{Gal}(L/K)|$. By Sylow theory there is some intermediate field $K < F < L$ such that L/F is a p-nilpotent Galois extension. Then $F(p) \neq F$. A theorem of Becker [5] says that if F is not p-closed and admits no orderings, then $[F(p) : F] = \infty$. We forced F to contain $\sqrt{-1}$ in characteristic 0, and F is not p-closed, so $[F(p) : F] = \infty$. Then no finite extension F' of F will contain $F(p)$, so F is p-corrupted.[4] □

9.5.2 Applying Jahnke-Koenigsmann

Say that a fairly saturated field L, perhaps with extra structure, is *special* if it is a finite extension of an infinite dp-minimal field.[5]

Most of Theorem 9.4.18 applies to special fields:

Remark 9.5.3. *Let L be a special field.*

1. *L is perfect*
2. *L admits a non-trivial Henselian valuation if L is not algebraically closed.*
3. *Any definable valuation on L is henselian. Any two definable valuations on L are comparable.*

All of these facts follow easily from the analogous facts for dp-minimal fields.

Also, special fields are closed under the following operations:

- Any finite extension of a special field is special.
- If L is special and w is a definable valuation on L, then Lw is finite or special.

To see the second point, let $v = w|K$, and note that Lw is a finite extension of Kw, which is dp-minimal.

We get a handle on special fields via the following trick:

Proposition 9.5.4. *Let L be a special field. Suppose L is not orderable, and contains all the pth roots of unity. Then the canonical p-henselian valuation v_L^p on L is definable, and its residue field is finite or p-closed.*

[4] It is not hard to prove a slightly weaker version of Becker's result, saying that if F is not p-closed, but contains the pth roots of unity (and $\sqrt{-1}$ when $p = 2$), then $[F(p) : F] = \infty$. The proof we have just given of Lemma 9.5.2 can be modified to use this weaker result. We leave this as an exercise to the reader.

[5] More precisely, L is special if there is an L-definable infinite subfield K with $[L : K]$ finite, and the induced structure on K is dp-minimal.

Proof. Definability of v_L^p follows by the work of Jahnke-Koenigsmann. Suppose the residue field Lv_L^p is infinite and not p-closed. Then Lv_L^p is special and not algebraically closed, so it admits a non-trivial henselian place $Lv_L^p \to L'$. The place $L \to Lv_L^p$ is henselian because it is definable, so the composition

$$L \to Lv_L^p \to L'$$

is itself a henselian place, which corresponds to a *finer* p-henselian valuation than v_L^p. But the canonical p-henselian valuation is the finest p-henselian valuation, unless its residue field is p-closed, so we have a contradiction. □

Mostly we will use the following consequence:

Corollary 9.5.5. *Let L be a special field containing all the 4pth roots of unity, and let v be a henselian valuation on L which is as fine as every definable valuation on L. Then Lv is finite or p-closed.*

Here we are using Fact 9.5.1.2.

9.5.3 The saturated case

Remark 9.5.6. *Let K be a field, $(I, <)$ a totally ordered set, and $\langle \mathcal{O}_x \rangle_{x \in I}$ be a totally ordered chain of valuation rings on K. Then the intersection*

$$\mathcal{O} = \bigcap_{x \in I} \mathcal{O}_x$$

is itself a valuation ring on K. If the intersection has residue characteristic p, then some $\mathcal{O}_x$ does: either K itself has characteristic p, or $1/p \notin \mathcal{O}$, hence $1/p \notin \mathcal{O}_x$ for some $x \in I$.

Theorem 9.5.7. *Let K be a sufficiently saturated dp-minimal field. Let $\mathcal{O}_\infty$ be the intersection of all the definable valuation rings on K. (So $\mathcal{O}_\infty = K$ if K admits no definable non-trivial valuations.)*

1. *$\mathcal{O}_\infty$ is a henselian valuation ring on K*
2. *$\mathcal{O}_\infty$ is type-definable, without parameters. In fact, it is the intersection of all 0-definable valuation rings on K.*
3. *The residue field of $\mathcal{O}_\infty$ is finite, real-closed, or algebraically closed. If it is finite, then $\mathcal{O}_\infty$ is definable.*

Proof. 1. By Theorem 9.4.18(3), the class of definable valuation rings on K is totally ordered. An intersection of a chain of valuation rings is a valuation ring. An intersection of a chain of henselian valuation rings is henselian.

2. We need to show that $\mathcal{O}_\infty$ is a small intersection. Suppose $\mathcal{O}$ is a definable valuation ring on K, defined by a formula $\phi(K;b)$. Let $\psi(x)$ be the formula asserting that $\phi(K;x)$ is a valuation ring. Then $\bigcap_{b\in\psi(K)}\phi(K;b)$ is a 0-definable valuation ring contained in $\mathcal{O}$. Thus every definable valuation ring on K contains a 0-definable valuation ring. Therefore $\mathcal{O}_\infty$ is the intersection of the 0-definable valuation rings on K. It is therefore type-definable over $\emptyset$.

3. First suppose that the residue field of $\mathcal{O}_\infty$ is finite. Let $\mathfrak{m}_\infty$ denote the maximal ideal of $\mathcal{O}_\infty$. Then $\mathfrak{m}_\infty = K \setminus \mathcal{O}_\infty^{-1}$, so $\mathfrak{m}_\infty$ is $\vee$-definable. On the other hand, $\mathcal{O}_\infty$ is a finite union of translates of $\mathfrak{m}_\infty$, so $\mathcal{O}_\infty$ is also $\vee$-definable, hence definable.

   Now suppose that the residue field is infinite. Let v_∞ denote the valuation associated with $\mathcal{O}_\infty$. Note that v_∞ is as fine as any definable valuation on K, by choice of $\mathcal{O}_\infty$.

   In particular, for every definable valuation v on K, the place $K \to Kv_\infty$ factors as a composition of two places

$$K \to Kv \to Kv_\infty$$

   We first show that Kv_∞ is perfect. If Kv_∞ has characteristic p, then Kv has characteristic p for some definable valuation v, by Remark 9.5.6. The field Kv is perfect by Observation 9.2.9, so the place $Kv \to Kv_\infty$ ensures that Kv_∞ is perfect as well (perfect equicharacteristic valued fields have perfect residue fields).

   Suppose for the sake of contradiction that Kv_∞ is not algebraically closed or real closed. As Kv_∞ is perfect, Lemma 9.5.2 applies, and some finite extension F of Kv_∞ is p-corrupted for some prime p (not necessarily the characteristic).

   Choose a finite extension L of K such that

   - L contains all the $4p$th roots of unity
   - If w_∞ denotes the (unique) extension of v_∞ to L, then Lw_∞ contains F, hence is not p-closed (nor finite).

   By Corollary 9.5.5, some definable valuation w on L is *not* a coarsening of w_∞. Let v be $w|K$. Then v is a coarsening of v_∞:

$$v(x) = v_\infty(x) + \Delta$$

   for some convex subgroup $\Delta < v_\infty K$. Coarsening w_∞ with respect to the same convex subgroup Δ, we get a coarsening w' of w_∞, whose restriction to K is v. But v is henselian, so $w = w'$, and w is coarser than w_∞, a contradiction.

□

9.5.4 The general case

We now prove Theorem 9.1.4. Let K be an arbitrary dp-minimal field. The 0-definable valuation rings on K are henselian and pairwise comparable (Theorem 9.4.18), so their intersection $\mathcal{O}$ is a henselian valuation ring, as in the previous section.

If $\mathcal{O}$ is non-trivial, then there is at least one 0-definable non-trivial valuation ring $\mathcal{O}_1$ on K. By Lemma 9.4.12, the valuation ring $\mathcal{O}_1$ induces the canonical topology on K. As $\mathcal{O}_1$ is a coarsening of $\mathcal{O}$, the valuation ring $\mathcal{O}$ also induces the canonical topology.

Let v be the valuation associated with $\mathcal{O}$.

Lemma 9.5.8. *The valuation v is defectless.*

Proof. If Kv has characteristic 0, then henselian implies defectless. So suppose Kv has characteristic p. By Remark 9.5.6, there is some 0-definable valuation w with residue characteristic p. Then $K \to Kw$ is defectless by Theorem 4.3.1 (plus Theorem 9.4.18), and $Kw \to Kv$ is defectless by Lemma 4.1.3. By Remark 3.1.5, $K \to Kv$ is defectless. □

Finally, we need to show that the residue field Kv is finite, algebraically closed, or real closed.

Let $\mathbb{M} \succeq K$ be a sufficiently saturated elementary extension, and let $\mathcal{O}_\infty$ be the type-definable subring of $\mathbb{M}$ from the previous section—the intersection of the 0-definable valuation rings on $\mathbb{M}$. Then $\mathcal{O} = \mathcal{O}_\infty \cap K$.

There are three cases:

1. If $\mathcal{O}_\infty$ has finite residue field, then $\mathcal{O}_\infty$ is definable, hence 0-definable. Then $\mathcal{O} \preceq \mathcal{O}_\infty$ so $\mathcal{O}$ has finite residue field.

2. Next, suppose $\mathcal{O}_\infty$ has algebraically closed residue field. Then for every n, we have

$$\forall a_1, \ldots, a_n \in \mathcal{O}_\infty \exists x \quad x^n + a_1 x^{n-1} + \cdots + a_n \in \mathfrak{m}_\infty$$

   By compactness, there is some 0-definable valuation ring $\mathcal{O}_n$ such that

$$\forall a_1, \ldots, a_n \in \mathcal{O}_n \exists x \quad x^n + a_1 x^{n-1} + \cdots + a_n \in \mathfrak{m}_n$$

   This remains true in K, as $K \preceq \mathbb{M}$. Because $\mathcal{O}_n \supseteq \mathcal{O}$ and $\mathfrak{m}_n \subseteq \mathfrak{m}$, we get

$$\forall a_1, \ldots, a_n \in \mathcal{O} \exists x \quad x^n + a_1 x^{n-1} + \cdots + a_n \in \mathfrak{m}$$

   which implies that all degree n monic polynomials in $\mathcal{O}/\mathfrak{m}$ have roots.

3. Next, suppose $\mathcal{O}_\infty$ has a real closed residue field. The unique extension of $\mathcal{O}_\infty$ to $\mathbb{M}(i)$ has algebraically closed residue field. Repeating the arguments we just gave with $K(i)$ and $\mathbb{M}(i)$ instead of K and $\mathbb{M}$, we see that the residue field of $K(i)$ with respect to $\mathcal{O}$ is algebraically closed. So the residue field of K is real closed or algebraically closed.

9.6 VC-minimal fields

In [24] Definition 1.4, Guingona makes the following definition:

Definition 9.6.1. *A theory T is dp-small if there does* not *exist a model $M \models T$, formulas $\phi_i(x; y_i)$ with $|x| = 1$, and a formula $\psi(x; z)$, and elements a_{ij}, b_i, c_j such that*

$$M \models \phi_i(a_{i'j}, b_i) \iff i = i'$$

$$M \models \psi(a_{ij'}, c_j) \iff j = j'$$

The combinatorial configuration here is more general than the one in the definition of dp-minimality, so dp-smallness is a stricter condition than dp-minimality.

Like dp-minimality, dp-smallness is preserved under reducts and under naming parameters. Guingona shows that VC-minimal fields are dp-small.

Theorem 9.6.2. *Let K be a dp-small field. Then K is algebraically closed or real closed.*

Proof. We can and do take K to be sufficiently saturated. By Theorem 1.6.4 of [24], the value group vK is divisible for any definable valuation v on K.

By Theorem 9.5.7, there is a henselian defectless valuation v_∞ on K whose valuation ring is the intersection of all definable valuation rings on K. The residue field of v_∞ is algebraically closed, real closed, or finite. In the finite case, v_∞ is definable, and by Theorem 4.3.1, the interval $[-v(p), v(p)]$ in the value group is finite, contradicting divisibility.

Therefore, the residue field Kv_∞ is algebraically closed or real closed. For K to be algebraically closed or real closed, it suffices to show that the value group $v_\infty K$ is divisible, by Ax-Kochen-Ershov in the real closed case, and defectlessness in the algebraically closed case.

Let ℓ be any prime. Let a be an element of $K^\times$. For each definable valuation $\mathcal{O}$ on K, the value group $K^\times/\mathcal{O}^\times$ is ℓ-divisible. So there is an element $b \in K^\times$ and $c \in \mathcal{O}^\times$ such that $a = b^\ell \cdot c$. The valuation ring $\mathcal{O}_\infty$ of v_∞ is the intersection of a small ordered set of $\mathcal{O}$'s, so by compactness, we can find $b \in K^\times$ and $c \in \mathcal{O}_\infty^\times$ such that $a = b^\ell \cdot c$. Then $v_\infty(a) = \ell \cdot v_\infty(b)$. So v_∞ has ℓ-divisible value group, for arbitrary ℓ. □

We can also specialize this result to C-minimal and (weakly) o-minimal fields.

Corollary 9.6.3. *Dense C-minimal fields are algebraically closed.*

Proof. Proposition 5.2.7 prevents dense C-minimal fields from being real-closed. □

Note that "C-minimal field" merely means "field which is C-minimal." Unlike [28], we are not assuming any compatibility between the field operations and the C-structure.

9.6.1 Weakly o-minimal fields

Weakly o-minimal fields turn out to be real closed. This will require a little bit of work. First we prove a lemma:

Lemma 9.6.4. *Let $\mathbb{M}$ be a sufficiently saturated non-strongly minimal dp-minimal field, possibly with extra structure. Then every infinitesimal type is multiplicatively stabilized by $\mathbb{G}_m^0$.*

Proof. Let p be an infinitesimal type over $\mathbb{M}$, and take $a \in \mathbb{G}_m^0$. Let $\psi(x; z)$ be a formula. We will show that $a^{-1} \cdot p$ and p have the same ψ-type.

Let $\phi(x; y, z)$ be the formula $\psi(x \cdot y; z)$. Every ψ-formula is a ϕ-formula, so it suffices to show that $a^{-1} \cdot p$ and p have the same ϕ-type. Moreover, the multiplicative group acts on ϕ-formulas and hence on ϕ-types.

For any $\alpha \in \mathbb{G}_m$, the type $\alpha \cdot p$ is an infinitesimal type, because p is infinitesimal. By Corollary 9.3.19, the orbit of p is small. Restricting to ϕ-types, we see that $p|\phi$ has a small orbit as well.

Because infinitesimal types are definable (Corollary 9.3.20), the multiplicative stabilizer of the ϕ-type $p|\phi$ is definable. Therefore the orbit is interpretable. Being bounded, it must be finite. So $p|\phi$ is stabilized by some finite-index subgroup of $\mathbb{G}_m$. As $a \in \mathbb{G}_m^0$, it follows that $a \cdot p|\phi = p|\phi$ as claimed. □

Using this, we can prove a rather strong and surprising result.

Theorem 9.6.5. *Let K be a dp-minimal algebraically closed field with extra structure. Then there is no infinite definable subset of K with a definable total ordering.*

Proof. We may replace K with a monster $\mathbb{M}$.

Suppose some infinite definable set $D \subseteq \mathbb{M}$ admits a definable total ordering $<_D$. Then D has non-empty interior by Proposition 9.3.22. Translating D, we may assume that 0 is in the interior of D. So all infinitesimal types over $\mathbb{M}$ live in D.

Let p be some nonzero infinitesimal type. Then p is multiplicatively stabilized by $\mathbb{G}_m^0$. Because $\mathbb{M}$ is algebraically closed, $\mathbb{G}_m$ is divisible. This implies that it has no proper subgroups of finite index. Therefore $\mathbb{G}_m^0 = \mathbb{G}_m$, so $a \cdot p = p$ for any $a \in \mathbb{G}_m$.

Let ω be some root of unity, other than 1. (For example, in characteristic $\neq 2$, take $\omega = -1$.) Then $\omega \cdot p = p$. As p is nonzero and $\omega \neq 1$, the type $p(x)$ must say $x \neq \omega \cdot x$. By totality of the ordering, we may assume $x <_D \omega \cdot x$ is in $p(x)$, reversing the order if necessary.

Now if a realizes p in some elementary extension of $\mathbb{M}$, then $\omega^i \cdot a \models \omega^i \cdot p = p$. In particular,

$$\omega^i \cdot a <_D \omega \cdot (\omega^i \cdot a)$$

for all i. By transitivity, the map $i \mapsto \omega^i \cdot a$ is strictly increasing, hence injective, contradicting the fact that ω is a root of unity. □

Because weakly o-minimal structures are VC-minimal and dp-small, we immediately get the following corollary, which was probably more easily proven by other means:

Corollary 9.6.6. *Weakly o-minimal fields are real-closed.*

Again, this is slightly more general than the result in [50], since we are not assuming that the weakly o-minimal ordering is a field ordering.

9.7 The classification of dp-minimal fields

At this point, only Theorem 9.1.5 remains. According to [39], Chernikov and Simon prove the following fact in [12].

Fact 9.7.1. *A henselian valued field (K, v) with residue characteristic 0 is dp-minimal if and only if vK and Kv are dp-minimal.*

Theorem 9.1.51 asserts that certain theories are complete and dp-minimal. Except in the case of positive residue characteristic, the completeness follows by the Ax-Kochen-Ershov principle, and dp-minimality follows by Fact 9.7.1, using the dp-minimality of characteristic 0 local fields, plus the characterization of dp-minimal ordered abelian groups (Proposition 5.1 in [39]). Characteristic 0 local fields are dp-minimal by Corollary 7.8 of [3] in the non-archimedean case, and by VC-minimality in the case of $\mathbb{C}$ and $\mathbb{R}$.

For the remaining case of positive residue characteristic, Corollary 3.2.18 provides completeness, and Theorem 3.3.7 establishes dp-minimality. Finally, part (2) of Theorem 9.1.5 follows from a more general fact:

Theorem 9.7.2. *Let K be a sufficiently saturated dp-minimal field. Then there is a henselian defectless valuation v on K such that*

- *The residue field Kv is algebraically closed, real closed, or a local field of characteristic 0.*
- *The value group vK satisfies $|vK/n \cdot vK| < \aleph_0$ for all $n > 0$.*
- *If Kv has characteristic p, then vK is p-divisible.*
- *If v has mixed characteristic, then $v(p) \in \mathrm{Int}_p\, vK$.*

Proof. First we note that if v is *any* valuation on K, then $vK/n \cdot vK$ is finite for all n, by Theorem 9.4.184. For the other points, we break into cases.

Let v_∞ be the valuation from Theorem 9.5.7. First suppose that Kv_∞ is finite. Then v_∞ is definable. By Theorem 4.3.1, v_∞ has mixed characteristic and the interval $[-v_\infty(p), v_\infty(p)]$ is finite. Then

$$\Delta := \mathbb{Z} \cdot v_\infty(p)$$

is the smallest convex subgroup of $v_\infty K$ containing $v_\infty(p)$.

Let v be the coarsening of v_∞ by Δ. We get a decomposition of the henselian defectless place $K \to Kv_\infty$

$$K \overset{v_\infty K/\Delta}{\to} Kv \overset{\Delta}{\to} Kv_\infty$$

Because K is saturated and v_∞ is definable, countable chains of balls in (K, v_∞) have non-empty intersection, meaning that $K \to Kv_\infty$ satisfies the countable intersection property of Remark 3.1.5. Thus $Kv \to Kv_\infty$ also satisfies this condition. But the value group of $Kv \to Kv_\infty$ is isomorphic to $\mathbb{Z}$, so the valuation $Kv \to \Delta$ is a mixed characteristic complete discrete valuation with a finite residue field. Therefore Kv is a characteristic 0 local field. So v is a henselian (and defectless) valuation on K, and its residue field Kv is local of characteristic 0. There is nothing else to show in this case, because v is equicharacteristic 0.

Otherwise, Kv_∞ is real closed or algebraically closed. In this case, we take $v = v_\infty$.

It remains to show that v_∞ is roughly p-divisible (see Definition 3.1.4) if Kv_∞ has characteristic p. By Remark 9.5.6, there is a definable valuation v_1 such that Kv_1 has characteristic p. The place $K \to Kv_\infty$ decomposes as

$$K \to Kv_1 \to Kv_\infty$$

where $K \to Kv_1$ is roughly p-divisible by Theorem 4.3.1, and $Kv_1 \to Kv_\infty$ is roughly p-divisible by Lemma 4.1.3. So the composition is roughly p-divisible by Remark 3.1.5. □

9.8 Dp-minimal valued fields

The above results easily yield a sharp characterization of dp-minimal valued fields, which we give in the next two theorems:

Theorem 9.8.1. *Let (K, v) be a valued field with infinite residue field. Then (K, v) is dp-minimal (as a pure valued field) if and only if the following conditions all hold:*

1. *The residue field Kv and value group vK are dp-minimal*
2. *The valuation v is henselian and defectless*
3. *In mixed characteristic, every element of $[-v(p), v(p)]$ is divisible by p.*
4. *In pure characteristic p, the value group vK is p-divisible.*

Proof. First suppose (K, v) is dp-minimal. Both vK and Kv are dp-minimal because they are images of the dp-minimal field K. Corollary 9.4.16 yields henselianity. Theorem 4.3.1 yields the divisibility conditions.

Conversely, suppose (K, v) satisfies conditions 1-4. These conditions are first order, so we may assume (K, v) is sufficiently saturated. As Kv is a dp-minimal field, there is a place $Kv \to k$ which is henselian, defectless, roughly p-divisible, and with k algebraically closed or

elementarily equivalent to a local field of characteristic 0. By Remark 3.1.5, the composition $K \to Kv \to k$ is also henselian, defectless, and roughly p-divisible.

Recall that an ordered abelian group Γ is dp-minimal if and only if $\Gamma/n \cdot \Gamma$ is finite for all $n > 0$. If Γ is an ordered abelian group, and Δ is a convex subgroup, then Γ is dp-minimal if and only if Δ and Γ/Δ are.

Therefore, the value group of $K \to Kv \to k$ is dp-minimal because the value groups of $K \to Kv$ and $Kv \to k$ are.

In summary, the composite place $K \to k$ is henselian, defectless, and roughly p-divisible, its residue field is local of characteristic 0, or algebraically closed, and its value group Γ has the property that $\Gamma/n\Gamma$ is finite for all n. By Theorem 9.1.5.1, the valued field $K \to k$ is dp-minimal. The original valued field $K \to Kv$ is a coarsening of $K \to k$, so it is definable in the shelah expansion of $K \to Kv$ (the expansion by all externally definable sets). The shelah expansion is still dp-minimal. Thus $K \to Kv$ is dp-minimal. □

Theorem 9.8.2. *Let (K, v) be a valued field with finite residue field. Then (K, v) is dp-minimal (as a valued field) if and only if the following conditions all hold:*

1. *The value group vK is dp-minimal*
2. *The valuation v is henselian*
3. *The valuation is finitely ramified, in the sense that $[-v(p), v(p)]$ is finite. (In particular, K has characteristic 0 if v is non-trivial.)*

Proof. First suppose (K, v) is dp-minimal. Then henselianity follows by Corollary 9.4.16, and dp-minimality of vK is immediate. Finite ramification follows by Theorem 4.3.1.

Conversely, suppose (K, v) satisfies conditions 1-3. These conditions are elementary, so we may assume K is saturated. If v is trivial, then K is finite, so it is dp-minimal. Otherwise, K has characteristic 0. Let w be the coarsening of v by the convex subgroup Δ generated by $v(p)$. As usual we get a decomposition $K \to Kw \to Kv$. The value group of $Kw \to Kv$ is Δ, which has rank 1 by finite ramification. By saturation of $K \to Kv$, the countable chain condition of Remark 3.1.5 holds in $K \to Kv$, hence in $Kw \to Kv$. Thus $Kw \to Kv$ is spherically complete. Also, Δ is isomorphic to $\mathbb{Z}$. Thus $Kw \to Kv$ makes Kw into a complete mixed characteristic DVR with finite residue field. So Kw is a local field of characteristic 0.

Now $K \to Kw$ is henselian (because $K \to Kv$ is). As $K \to Kv$ has dp-minimal value group, so does $K \to Kw$ and $Kw \to Kv$. In particular, $K \to Kw$ makes K into a henselian valued field with dp-minimal value group and residue field local of characteristic 0. So $K \to Kw$ is dp-minimal by Theorem 9.1.5.1.1a.

In characteristic 0 nonarchimedean local fields, the valuation ring is always definable from the pure field language. Consequently, the dp-minimal structure $K \to Kw$ interprets $K \to Kw \to Kv$. Thus $K \to Kv$ is also dp-minimal. □

Question 9.8.3. *Do the above theorems remain true when (K, v) is expanded by additional structure on vK and Kv (preserving the dp-minimality of each)?*

9.9 Summary and future directions

We now know exactly which *pure* fields are dp-minimal, and we know a little bit about dp-minimal expansions of fields.

Here is a summary of what can be said about a dp-minimal field $(K, +, \cdot, \ldots)$, perhaps with other structure. Either K is strongly minimal (or finite), or all of the following facts are true:

- There is a definable V-topology on K (Theorem 9.1.3)
- With respect to this topology, there are only boundedly many infinitesimal types, which are all definable (Corollaries 9.3.19 and 9.3.20).
- Any unary definable set has finite boundary (Theorem 9.1.3)
- Dp-rank of definable[6] sets is definable in families, and agrees with "geometric dimension." (Corollary 9.4.7).
- Any definable valuation ring is henselian and defectless, and any two definable valuation rings are comparable (Proposition 9.4.15, Corollary 9.4.16, Theorem 4.3.1)
- For each n, $K^\times/(K^\times)^n$ is finite (Theorem 9.4.18.4).
- $\mathbb{G}_m^{00} = \mathbb{G}_m^0 = \bigcap_n (K^\times)^n$. Indeed, $\mathbb{G}_m^0 = \bigcap_n (K^\times)^n$ by the previous point. Because each infinitesimal type lives in a specific coset of $\mathbb{G}_m^{00}$, the multiplicative stabilizer of any infinitesimal type must be $\mathbb{G}_m^{00}$, but the stabilizer is also $\mathbb{G}_m^0$ by Lemma 9.6.4.
- Any finite extension of K is dp-minimal as a pure field (but not as an expanion of K, of course). This follows by inspecting the list of dp-minimal fields.
- There is at least one definable non-trivial valuation on K, unless K is finite, real closed, or algebraically closed (Theorem 9.5.7).

There are several obvious questions we have not addressed:

Question 9.9.1. *If K is a dp-minimal field, is there always a definable valuation on K whose residue field is algebraically closed, real closed, or finite?*

Question 9.9.2. *Which unstable dp-minimal valued fields fail to define valuations?*

Question 9.9.3. *If K is a sufficiently saturated unstable dp-minimal valued field, and $\mathcal{O}_\infty$ is the valuation ring from Theorem 9.5.7, is the expansion of K by $\mathcal{O}_\infty$ still dp-minimal?*

Question 9.9.4. *Can any of the classification be extended to fields of finite dp-rank?*

[6] One cannot hope to extend this to interpretable sets. For instance, $\exists^\infty$ is not eliminated in the value group of $\mathbb{Q}_p$.

9.9.1 Defining the canonical valuation, or not

In general, the answer to Question 9.9.1 is *no*, though we can characterize the failure modes.

Proposition 9.9.5. *Let (K, v) be a sufficiently saturated valued field as in Theorem 9.1.5.1. So (K, v) is dp-minimal, and vK is elementary equivalent to $\mathbb{F}_p^{alg}$ or a characteristic 0 local field.*

Let $\mathcal{O}_\infty$ be the intersection of all valuation rings on K definable in the pure field language. Let w be the associated valuation.

- *If Kv is non-archimedean, then w is the composition of v with the canonical valuation on Kv.*
- *If Kv is real closed or algebraically closed, then w is the coarsening of v by the maximal convex divisible subgroup of vK.*

Proof. First we make a general observation.

Remark 9.9.6. *Let $K \to k$ be a place. It cannot be the case that one of K or k is finite and the other is real closed or algebraically closed. Indeed, if K is finite then k is (obviously) finite. If K is algebraically closed or real closed, then so is k.*

First suppose Kv is non-archimedean. Non-archimedean local fields define their valuation rings, so (K, v) interprets the canonical valuation on Kv. Let $K \to Kv \to Kv'$ be the composition, so v' is a definable valuation on K with finite residue field Kv'.

By Proposition 9.4.15 applied to the dp-minimal structure (K, v), the valuations v' and w must be comparable. So we either have a place map $Kw \to Kv'$ or $Kv' \to Kw$. By Theorem 9.5.7, Kw is finite, real closed, or algebraically closed. By remark 9.9.6, Kw cannot be algebraically closed or real closed, so it is finite. Then $Kw \to Kv'$ or $Kv' \to Kw$ is trivial, and $w = v'$.

Next suppose Kv is real closed or algebraically closed. Let v' be the coarsening of v by the maximal divisible convex subgroup of vK. In the sequence

$$K \to Kv' \to Kv$$

the value group of $Kv' \to Kv$ is divisible, and the value group of $K \to Kv'$ has no convex divisible subgroups. Also, $K \to Kv'$ is henselian, defectless, and roughly p-divisible. Because Kv' has a henselian defectless valuation with divisible value group and real or algebraically closed residue field, Kv' is itself real closed or algebraically closed.

So (K, v') models one of the theories from Theorem 9.1.5.1, though (K, v') need not be saturated. One can show that in (K, v'), the induced structure on Kv' is the pure field structure.

By Proposition 9.4.15, w and v' must be comparable. If v' were strictly coarser than w, we would have a non-trivial valuation $Kv' \to Kw$, definable in the structure (K, v'). But pure models of RCF and ACF do not admit non-trivial valuations.

So w is coarser than v', which is in turn coarser than v. Let Δ_w and $\Delta_{v'}$ be the convex subgroups of vK whose coarsenings yield w and v'. Then $\Delta_w \geq \Delta_{v'}$. We want to show $w = v'$, i.e., that $\Delta_w = \Delta_{v'}$. Otherwise, $\Delta_w > \Delta_{v'}$. As $\Delta_{v'}$ is the greatest convex divisible subgroup, Δ_w is *not* divisible. Neither is $\Delta_w/\Delta_{v'}$. So the place $Kw \to Kv'$ has a value group that is not divisible.

Now Kv' is real closed or algebraically closed, so Kw is not finite by Remark 9.9.6. Therefore Kw is real closed or algebraically closed. But then any valuation on Kw has divisible value group. This contradicts the non-divisibilitiy of the value group of $Kw \to Kv'$. □

Theorem 9.9.7. *Let (K, v) be a dp-minimal valued field with residue field Kv algebraically closed or elementarily equivalent to a local field of characteristic 0. Suppose K is sufficiently saturated. The following are equivalent:*

- *There is a valuation w, definable in the pure field language, such that Kw is finite, real closed, or algebraically closed.*
- *Kv is non-archimedean or the maximal convex divisible subgroup of vK is definable.*

Proof. First suppose that Kv is non-archimedean. By Proposition 9.9.5, the canonical valuation on K (in the pure field language) has finite residue field (and so is definable by Theorem 9.5.7). So there is a valuation ring on K, definable in the pure field language, with finite residue field.

Next suppose that Kv is algebraically closed or real closed. If the maximal divisible convex subgroup of vK is definable, let v' be the coarsening, which is definable in (K, v). Let w be the canonical valuation on Kv. Then $w = v'$ by the proposition. The valuation ring of w is then type definable in the pure field language, and definable in (K, v). It must then be definable in the pure field language. So w is definable in the pure field language, and Kw is finite, algebraically closed, or real closed, by Theorem 9.5.7.

Conversely, suppose that Kw is algebraically closed, real closed, or finite, for some w definable in the pure field language.

Now w is coarser than the canonical valuation on the pure field K which is coarser than v by the Proposition (applied to (K, v)). So there is a place $Kw \to Kv$. By Remark 9.9.6, Kw is algebraically closed or real closed.

Then we can apply the Proposition to (K, w), seeing that the canonical valuation on K is a coarsening of w. So w and the canonical valuation on K are coarser than each other, hence equal.

Now by the Proposition applied to (K, w), w is the coarsening of v by the maximal convex divisible subgroup of vK. This group must be definable in (K, v). As (K, v) induces the pure ordered group structure on vK, this group is definable in vK. □

Now let Γ be the lexicographic product

$$\mathbb{Z} \times \mathbb{Z}[1/2] \times \mathbb{Z}[1/2, 1/3] \times \mathbb{Z}[1/2, 1/3, 1/5] \times \cdots$$

For each number n, all but finitely many of the factors are divisible by n, and in fact $\Gamma/n\Gamma$ is finite for all n. So $\mathbb{C}((t^\Gamma))$ is dp-minimal. But in a sufficiently saturated elementary extension of Γ, the maximal divisible convex subgroup of Γ is not definable. In fact, it is the intersection of the strictly decreasing sequence of definable subgroups:

$$\mathrm{Int}_2\,\Gamma > \mathrm{Int}_3\,\Gamma > \cdots$$

9.9.2 Unstable dp-minimal valued fields that define no valuations

We know that every unstable dp-minimal field K has a V-topology. This topology need not come from a definable valuation, as exhibited by RCF. On the other hand, if K admits no definable valuation rings, then the ring $\mathcal{O}_\infty$ in Theorem 9.5.7 is trivial, so K must be algebraically closed or real closed. So most dp-minimal fields admit a definable valuation, which determines the canonical topology (by Lemma 9.4.12).

A natural open question is then:

Question 9.9.8. *Are there dp-minimal unstable expansions of ACF which define no valuation rings?*

If the answer is no, the following conjecture is true:

Conjecture 9.9.9. *Let K be an unstable dp-minimal field. Then the canonical topology on K is induced by a definable ordering or a definable valuation.*

9.9.3 Expanding by the canonical valuation

In many cases, the answer to Question 9.9.3 is *yes*, because the canonical valuation is definable. (Whether this happens is more or less characterized by Proposition 9.9.5.) In the case of pure fields, we know that the answer to Question 9.9.3 is *yes*, at least under saturation assumptions—this is essentially the content of Theorem 9.7.2.

In general, Question 9.9.3 remains open.

9.9.4 Finite dp-rank fields

Strongly minimal fields are known to be algebraically closed by a theorem of Macintyre. To prove this, one must prove a stronger statement: fields of finite Morley rank are algebraically closed.

In contrast, the tools and techniques we have used to classify dp-minimal fields do not seem to generalize in an obvious way to fields of finite dp-rank.

Unlike the case of finite Morley rank, there are (probably) pure fields of arbitrarily high finite dp-rank. Let R_n be the localization $S^{-1}\mathbb{Z}$ where S is generated by all but the first $(n-1)$ primes. Let $\pi_1, \pi_2, \ldots$ be a countable subset of a transcendence basis of $\mathbb{R}$ over $\mathbb{Q}$.

Let A_n be $\sum_{i=1}^{\infty} R_n \cdot \pi_i$. One can show that A_n has dp-rank n as an ordered abelian group, and one expects that $\mathbb{C}((t^{A_n}))$ also has dp-rank n (as a pure field).

Nevertheless, one could hope that some analogue of Theorem 9.1.5.2 holds, with Γ required only to have finite dp-rank. Several hurdles stand in the way of directly generalizing the proof.

First of all, the definition of "infinitesimal" needs to be changed. For dp-minimal fields, we defined ϵ to be a K-infinitesimal if $X \cap (X + \epsilon)$ is infinite for every infinite K-definable set ϵ.

This is no longer reasonable in rank 2 fields. For instance, consider the field of complex numbers expanded by complex conjugation. If ϵ is non-zero, then at least one of the sets $\mathbb{R} \cap (\mathbb{R} + \epsilon)$ or $i \cdot \mathbb{R} \cap (i \cdot \mathbb{R} + \epsilon)$ is *empty*. So 0 would be the only infinitesimal, which is unacceptable.

A more likely definition for "infinitesimal" in the higher rank case is the following: ϵ is K-infinitesimal if $X \cap (X + \epsilon)$ has full dp-rank whenever X is K-definable of full dp-rank.

With this definition, it becomes unclear that the set of K-infinitesimals is type-definable, however. We would need to know that the condition of being full dp-rank is definable in families (generalizing Observation 9.2.10). The proof of Observation 9.2.10 relies heavily on the rank 1 assumption.

Also, changing the definition of infinitesimals would require an analogous change in the definition of slight maps, and the proof of Proposition 9.3.7 might become more difficult.

The next big hurdle is Theorem 9.3.16, the bound on the number of germs at 0. The proof of Theorem 9.3.16 directly relies on dp-minimality. Moreover, Theorem 9.3.16 and its corollaries all fail to hold in the example of the complex numbers expanded by conjugation:

- There are unboundedly many infinitesimal types, contradicting Corollary 9.3.19
- There are infinitesimal types which are not definable, contradicting Corollary 9.3.20
- There are definable sets (for instance, the upper half plane), with infinite boundary, contradicting Proposition 9.3.22.

Intuitively, Proposition 9.3.22 should probably be changed to say that no definable set has boundary of full dp-rank. Unfortunately, all the steps used to prove this Proposition lack analogous generalizations!

The last place where dp-minimality was used in an essential way was Corollary 9.3.27, which said that the groups I_K and $a \cdot I_K$ are comparable for any $a \in \mathbb{M}$ (here, I_K denotes the K-infinitesimals). The key step to proving this Corollary was a general statement about dp-minimal abelian groups: if G and H are two subgroups which are sufficiently connected, then G and H are comparable (Lemma 9.3.26).

This is a special case of a more general "dual Baldwin-Saxl" statement

Fact 9.9.10. *Let $(K, +)$ be an abelian group of dp-rank n. Let $G_1, \ldots, G_m$ be type-definable subgroups such that $G_i = G_i^{00}$ for each i. If $m > n$, then there exist $i_1, \ldots, i_n$ such that*

$$G_1 + \cdots + G_m = G_{i_1} + \cdots + G_{i_n}$$

If we could overcome all the other hurdles (proving the existence of a nontrivial type-definable subgroup $I_K < \mathbb{G}_a$ with $(1 + I_K) < \mathbb{G}_m$ and $I_K = I_K^{00}$), this Lemma would tell us something non-trivial about the family of rescalings $a \cdot I_K$ for $a \in \mathbb{M}$.

Chapter 10

Another proof that ACF defines geometric irreducibility

The model theory of fields occasionally utilizes the "definability" of certain basic concepts from algebraic geometry, such as dimension, Zariski closure, irreducibility, and reducedness. For instance, dimension is definable in the following sense: if $\mathbb{C}$ is a model of ACF, if $\phi(\vec{x};\vec{y})$ is a formula, and if $d \in \mathbb{N}$, the set of $\vec{b}$ such that $\phi(\mathbb{C};\vec{b})$ has dimension d is definable.

This sort of statement has geometric content. Among other things, it implies that if $f : V \to W$ is a morphism of varieties, then the sets

$$W_d = \{w \in W : \dim f^{-1}(w) = d\}, \qquad d = 0, 1, 2, \ldots, \dim V$$

are constructible subsets of W.

Model-theoretically, these "definability" results are useful in axiomatizing certain theories of fields. For example, the Pseudo Algebraically Closed (PAC) fields are the fields K such that every geometrically integral K-variety contains a K-point [22]. The definability of geometric integrality ensures that the class of PAC fields is an elementary class. Similarly, the geometric axioms of existentially closed difference fields in (1.1) of [7] are first order because properties like Zariski density are definable. We will see another example of this in the next chapter, when we axiomatize the existentially closed fields with multiple valuations in section §11.2.1.

All these "definability" facts are known classically, and can be proven directly using the tools of computational algebraic geometry. Model-theoretic proofs exist as well. For instance, the definability of dimension is a consequence of the definability of Morley rank in strongly minimal sets (Proposition 1.5.16 in [57]). The other facts are less automatic. For example, the definability of irreducibility implies that ACF has the definable multiplicity property; not all strongly minimal sets have this. In his thesis [16], Lou van den Dries found model-theoretic proofs of all the definability results using ultraproducts and flatness.

In this brief chapter, we present a potentially new proof of the definability of irreducibility in algebraically closed fields. The proof will be primarily model-theoretic, using the following input from algebraic geometry:

- Krull dimension agrees with Morley rank
- The Zariski closed sets of codimension 1 are the zero-sets of irreducible polynomials
- Zariski-closed subsets of projective space are complete varieties, hence have closed images under regular maps.

Most of the complexity of the proof is probably buried in the third fact, whose proof uses valuation theory (see §II.4 of [25]).

10.1 Irreducibility in Projective Space

Let $\mathbb{C}$ be a monster model of ACF. For $\vec{x} \in \mathbb{P}^n(\mathbb{C})$, let $\mathbb{P}_{\vec{x}}$ be the $n-1$-dimensional projective space of lines through $\vec{x}$, and let $\pi_{\vec{x}} : \mathbb{P}^n \setminus \{\vec{x}\} \to \mathbb{P}_{\vec{x}}$ be the projection.

Lemma 10.1.1. *Let A be a small set of parameters, and suppose $\vec{x} \in \mathbb{P}^n(\mathbb{C})$ is generic over A. Suppose V is an A-definable Zariski closed subset of $\mathbb{P}^n$, of codimension* greater than *1. Then $\pi_{\vec{x}}(V) \subseteq \mathbb{P}_{\vec{x}}$ is Zariski closed, of codimension one less than the codimension of V. Moreover, $\pi_{\vec{x}}(V)$ is irreducible if and only if V is irreducible.*

Proof. Replacing A with $\mathrm{acl}(A)$, we may assume A is algebraically closed, implying that the irreducible components of V are also A-definable.

Note that $\vec{x} \notin V$ because $\vec{x}$ is generic, and V has codimension *at least* 1. Therefore, $\pi_{\vec{x}}$ is a regular map on V. The image $\pi_{\vec{x}}(V)$ is Zariski closed because V is a complete variety.

Claim 10.1.2. *Let C be any irreducible component of V, and let $\vec{c} \in V$ realize the generic type of C, over $A\vec{x}$. Then $\vec{c}$ is the sole preimage in V of $\pi_{\vec{x}}(\vec{c})$.*

Proof. The generic type of C is A-definable, so $\vec{c} \mathop{\smile\!\!\!\!\!\!|}_A \vec{x}$, and therefore

$$RM(\vec{x}/A\vec{c}) = RM(\vec{x}/A) = n.$$

Suppose for the sake of contradiction that there was a second point $\vec{d} \in V$, $\vec{d} \neq \vec{c}$, satisfying

$$\pi_{\vec{x}}(\vec{d}) = \pi_{\vec{x}}(\vec{c}).$$

This means exactly that the three points $\vec{c}$, $\vec{d}$, and $\vec{x}$ are collinear. Then $\vec{x}$ is on the 1-dimensional line determined by $\vec{c}$ and $\vec{d}$, so

$$RM(\vec{x}/A\vec{c}\vec{d}) \leq 1.$$

But then

$$n = RM(\vec{x}/A\vec{c}) \leq RM(\vec{x}\vec{d}/A\vec{c}) = RM(\vec{x}/A\vec{c}\vec{d}) + RM(\vec{d}/A\vec{c}) \leq 1 + RM(V) < n,$$

by the codimension assumption. □

Using the claim, we see that $\pi_{\vec{x}}(V)$ and V have the same dimension (= Morley rank). Indeed, let $\vec{v} \in V$ have Morley rank $RM(V)$ over $A\vec{x}$. Then $\vec{v}$ realizes the generic type of *some* irreducible component C, so by the claim, $\vec{v}$ is interdefinable over $A\vec{x}$ with $\pi_{\vec{x}}(\vec{v})$. But then

$$RM(\pi_{\vec{x}}(V)) \geq RM(\pi_{\vec{x}}(\vec{v})/A\vec{x}) = RM(\vec{v}/A\vec{x}) = RM(V),$$

and the reverse inequality is obvious. So the codimension of $\pi_{\vec{x}}(V)$ is indeed one less.

Let $C_1, \ldots, C_m$ enumerate the irreducible components of V. (Possibly $m = 1$.) Each of the images $\pi_{\vec{x}}(C_i)$ is a Zariski closed subset of $\mathbb{P}_{\vec{x}}$, for the same reason that $\pi_{\vec{x}}(V)$ is, and each image is irreducible, on general grounds. If $\pi_{\vec{x}}(C_i) \subseteq \pi_{\vec{x}}(C_j)$ for some $i \neq j$, then the generic type of C_i would have the same image under $\pi_{\vec{x}}$ as some point in C_j, contradicting the Claim. So $\pi_{\vec{x}}(C_i) \not\subseteq \pi_{\vec{x}}(C_j)$ for $i \neq j$. It follows that the images $\pi_{\vec{x}}(C_i)$ are the irreducible components of

$$\pi_{\vec{x}}(V) = \bigcup_{i=1}^{m} \pi_{\vec{x}}(C_i).$$

Therefore, $\pi_{\vec{x}}(V)$ and V have the same number of irreducible components, proving the last point of the lemma. □

Theorem 10.1.3. *Let $X_{\vec{a}} \subseteq \mathbb{P}^n$ be a definable family of Zariski closed subsets of $\mathbb{P}^n$. Then the set of $\vec{a}$ for which $X_{\vec{a}}$ is irreducible, is definable.*

Proof. Dimension is definable in families, because ACF is strongly minimal. So we may assume that all (non-empty) $X_{\vec{a}}$ have the same (co)dimension. We proceed by induction on codimension, allowing n to vary.

For the base case of codimension 1, note that

1. The family of Zariski closed subsets of $\mathbb{P}^n$ is ind-definable, i.e., a small union of definable families, because the Zariski closed subsets are exactly the zero sets of finitely-generated ideals.

2. Using 1, the family of *reducible* Zariski closed subsets of $\mathbb{P}^n$ is also ind-definable, because a definable set is a *reducible* Zariski closed set if and only if it is the union of two incomparable Zariski closed sets.

3. Whether or not a polynomial in $\mathbb{C}[x_1, \ldots, x_{n+1}]$ is irreducible, is definable in terms of the coefficients, because we only need to quantify over lower-degree polynomials.

4. A hypersurface in $\mathbb{P}^n$ is irreducible if and only if it is the zero-set of an irreducible homogeneous polynomial. It follows by 3 that the family of irreducible codimension 1 closed subsets of $\mathbb{P}^n$ is ind-definable.

5. By 2 (resp. 4), the set of $\vec{a}$ such that $X_{\vec{a}}$ is reducible (resp. irreducible) is ind-definable. Since these two sets are complementary, both are definable, proving the base case.

For the inductive step, suppose that irreducibility is definable in families of codimension one less than $X_{\vec{a}}$. By choosing an isomorphism between $\mathbb{P}_{\vec{x}}$ and $\mathbb{P}^{n-1}$, one easily verifies the definability of the set of $(\vec{x}, \vec{a})$ such that $\pi_{\vec{x}}(X_{\vec{a}})$ is irreducible and has codimension one less.

By Lemma 10.1.1, $X_{\vec{a}}$ is irreducible if and only if $(\vec{x}, \vec{a})$ lies in this set, for generic $\vec{x}$. Definability of types in stable theories then implies definability of the set of $\vec{a}$ such that $X_{\vec{a}}$ is irreducible. □

Corollary 10.1.4. *The family of irreducible closed subsets of $\mathbb{P}^n$ is ind-definable.*

Proof. The family of closed subsets is ind-definable, and by Theorem 10.1.3 we can select the irreducible ones within any definable family. □

Corollary 10.1.5. *The family of pairs $(X, \overline{X})$ with X definable and $\overline{X}$ its Zarisk-closure, is ind-definable.*

Proof. By quantifier elimination in ACF, any definable set X can be written as a union of sets of the form $C \cap U$ with C closed and U open. Replacing V with a union of irreducibles, and distributing, we can write X as a union $\bigcup_{i=1}^{m} C_i \cap U_i$, with C_i Zariski closed and U_i Zariski open. We may assume that $C_i \cap U_i \neq \emptyset$ for each i, or equivalently, that $C_i \setminus U_i \neq C_i$.

In any topological space, closure commutes with finite unions, so

$$\overline{X} = \bigcup_{i=1}^{n} \overline{C_i \cap U_i}.$$

Now $\overline{C_i \cap U_i} \subseteq \overline{C_i} = C_i$, and

$$C_i = \overline{C_i \cap U_i} \cup (C_i \setminus U_i),$$

so by irreducibility of C_i, $\overline{C_i \cap U_i} = C_i$. Therefore,

$$\overline{X} = \bigcup_{i=1}^{n} C_i.$$

Corollary 10.1.4 implies the ind-definability of the family of pairs

$$\left(\bigcup_{i=1}^{n} \overline{C_i \cap U_i}, \bigcup_{i=1}^{n} C_i \right)$$

with C_i irreducible closed, U_i open, and $C_i \cap U_i \neq \emptyset$. We have seen that this is the desired family of pairs. □

The following corollary is an easy consequence:

Corollary 10.1.6. *Let $X_{\vec{a}}$ be a definable family of subsets of $\mathbb{P}^n$. Then the Zariski closures $\overline{X_{\vec{a}}}$ are also a definable family.*

10.2 Irreducibility in Affine Space

Theorem 10.2.1. *Let $X_{\vec{a}}$ be a definable family of subsets of affine n-space.*

1. *The family of Zariski closures $\overline{X_{\vec{a}}}$ is also definable.*
2. *The set of $\vec{a}$ such that $\overline{X_{\vec{a}}}$ is irreducible is definable. More generally, the number of irreducible components of $\overline{X_{\vec{a}}}$ is definable in families (and bounded in families).*
3. *Dimension and Morley degree of $X_{\vec{a}}$ are definable in $\vec{a}$.*

Proof.

1. Embed $\mathbb{A}^n$ into $\mathbb{P}^n$. Then the Zariski closure of $X_{\vec{a}}$ within $\mathbb{A}^n$ is the intersection of $\mathbb{A}^n$ with the closure within $\mathbb{P}^n$. Use Corollary 10.1.6.
2. The number of irreducible components of the Zariski closure is the same whether we take the closure in $\mathbb{A}^n$ or $\mathbb{P}^n$. This proves the first sentence. The first sentence yields the ind-definability of the family of irreducible Zariski closed subsets of $\mathbb{A}^n$, from which the second statement is an exercise in compactness.
3. We may assume $X_{\vec{a}}$ is closed, since taking the closure changes neither Morley rank nor Morley degree. The family of d-dimensional Zariski irreducible closed subsets of $\mathbb{A}^n$ is ind-definable, making this an exercise in compactness.

□

Chapter 11

Forking and Dividing in Fields with Several Orderings and Valuations

The theory of fields with n unrelated valuations has a model companion, by the thesis of van den Dries [16]. One can also include orderings and p-valuations. More precisely, suppose that for $1 \leq i \leq n$, the theory T_i is ACVF, RCF, or pCF for some p. Arrange that for $i \neq j$, the languages of T_i and T_j overlap only in the language of rings. Then one forms the theory $\bigcup_{i=1}^{n}(T_i)_\forall$, whose models are fields[1] K with additional structure making K a model of $(T_i)_\forall$, for each $1 \leq i \leq n$. In van den Dries's notation, $\bigcup_{i=1}^{n}(T_i)_\forall$ would be denote $((T_1)_\forall, (T_2)_\forall, \ldots, (T_n)_\forall)$.

For example, if each T_i is ACVF, then $(T_i)_\forall$ is the theory of valued fields, and

$$((T_1)_\forall, (T_2)_\forall, \ldots, (T_n)_\forall)$$

is the theory of fields with n different valuations. If each T_i is RCF, the $(T_i)_\forall$ is the theory of ordered fields, and $((T_1)_\forall, (T_2)_\forall, \ldots, (T_n)_\forall)$ is the theory of fields with n orderings. The case of pCF is similar, though for technical reasons one must use the Macintyre language. The T_i can be mixed; for example

$$(ACVF_\forall, RCF_\forall, 3CF_\forall)$$

is the theory of fields with a valuation, an ordering, and a 3-valuation (+ Macintyre predicates). In all these cases, van den Dries proves that a model companion

$$\overline{((T_1)_\forall, (T_2)_\forall, \ldots, (T_n)_\forall)}$$

exists. In fact, van den Dries's result is more general than what we have stated, allowing the T_i's to be arbitrary theories with quantifier elimination such that the $(T_i)_\forall$ are "t-theories" (Definition III.1.2 in [16]).

We will only consider the case where the T_i are ACVF, RCF, or pCF, however. In these cases, we will prove the following about the model companion $\overline{((T_1)_\forall, \ldots, (T_n)_\forall)}$, which we denote T for simplicity:

[1] Or rather, domains. We will sweep this issue under the rug.

1. T is NTP_2, but is never NSOP (obviously), and not NIP if $n > 1$. See Theorems 11.5.7 and 11.5.1. If $n = 1$, then T is one of ACVF, RCF, or pCF, which are all known to be NIP.

2. In fact, T is "strong" in the sense of Adler [1], i.e., every type has finite burden. The burden of affine m-space is exactly mn, where n is the number of valuations and orderings. See Theorem 11.5.7.

3. Forking and dividing agree over sets in the home sort, so every set in the home sort is an "extension base for forking" in the sense of Chernikov and Kaplan [10]. See Theorem 11.6.5.

4. Forking in the home sort has the following characterization (Theorem 11.6.10). Suppose $K \models T$, and $A, B, C \subseteq K$ are subsets of the home sort. For $1 \leq i \leq n$, let K_i be a model of T_i extending the reduct of K to the language of T_i. For example, in the case of n orderings, K_i could be a real closure of K with respect to the i^{th} ordering. Then $A \mathop{\smile\!\!\!\!\!\!|}_C B$ if and only if $A \mathop{\smile\!\!\!\!\!\!|}_C B$ holds in K_i for every i. The choice of the K_i does not matter.

In many of the cases, (1) follows from Samaria Montenegro Guzmán's recent proof in [54] that bounded pseudo-real-closed and pseudo-p-adically-closed fields are NTP_2 (originally Conjecture 5.1 in [11]). This includes the case

$$\overline{(RCF_\forall, RCF_\forall, \ldots, RCF_\forall)}$$

of existentially closed fields with n orderings, as well as the case when ever T_i is pCF for some p. As far as I know, (3) and (4) are new. I conjecture that (3) also holds of sets of imaginaries, which would imply that Lascar strong type and compact strong type agree, by [71] Corollary 3.6. In the case of fields with n orderings, this should follow from elimination of imaginaries, which Montenegro proved.

In the case where every T_i is ACVF, the model companion T is the theory of existentially closed fields with n valuations. In this case, the above results can be expressed more cleanly. It turns out (Theorem 11.3.1) that the model companion is axiomatized by the following axioms: $(K, v_1, v_2, \ldots, v_n)$ is a model if and only if K is algebraically closed, each v_i is non-trivial, and v_i and v_j induce different topologies on K, for $i \neq j$. In this case, forking is characterized as follows: $A \mathop{\smile\!\!\!\!\!\!|}_C B$ holds if and only if it holds in the reduct (K, v_i), for every i. This holds because each (K, v_i) is already a model of ACVF, and so we can take $K_i = K$ in the statement of (4).

One can also express the axioms of T more concisely if exactly one of the T_i is not ACVF. If T_1 is RCF and $T_2, \ldots, T_n$ are all ACVF, then we are considering the model companion of the theory of ordered fields with $(n-1)$ unrelated valuations. In this case, the model companion is axiomatized by the statement that the field is real closed, the valuations are non-trivial, the valuations induce different topologies from each other, and the valuations

induce different topologies from the order topology. Something similar holds with pCF in place of RCF. See Theorem 11.3.1 for details.

As a concrete example, let K be one of the following fields: $\mathbb{F}_p(t)^{alg}$, $\mathbb{Q}^{alg}$, $\mathbb{Q}^{alg} \cap \mathbb{R}$, or $\mathbb{Q}^{alg} \cap \mathbb{Q}_p$ for some p. Let $R_1, \ldots, R_n$ be valuation rings on K. Then K with the ring structure and with a unary predicate for each R_i is a strong NTP_2 theory, and every set of real elements is an extension base. The same holds for $\bigcap_{i=1}^n R_i$ as a pure ring.

The outline of this chapter is as follows. In Section 11.1, we recall some elementary facts about ACVF, pCF, and RCF which will be needed later. In Section 11.2, we quickly reprove the main facts needed from Chapters II and III of van den Dries's thesis, arriving at a slightly different way of expressing the axioms of the model companion, and handling the case of positive characteristic, which was not explicitly considered by van den Dries. Section 11.3 is a digression aimed at proving Theorem 11.3.1, which drastically simplifies the axioms of the model companions in some cases. Theorem 11.3.1 is probably known to experts, but we include a proof here for lack of a reference. In Section 11.4, we construct some Keisler measures that will be used in the later sections. In Section 11.5, we determine where the model companion lies in terms of various classification theoretic boundaries, proving that it is NTP_2 and strong, but not NSOP and usually not NIP. In Section 11.6, we show that forking and dividing agree over sets in the home sort, and we characterize forking in terms of forking in the T_i's.

11.1 Various facts about ACVF, pCF, and RCF

Let T be one of ACVF, RCF, or pCF (p-adically closed fields). Work in the usual one-sorted languages with quantifier elimination—for pCF this would be the Macintyre language, and for RCF this would be the language of rings.

The following fact follows easily from the various quantifier-elimination results:

Fact 11.1.1. *Let $\mathfrak{M}$ be a model of T, and K be a subfield. Every K-definable set is a positive boolean combination of topologically open sets and affine varieties defined over K. In particular, any K-definable subset of $\mathfrak{M}^n$ has non-empty interior or is contained in a K-definable proper closed subvariety of $\mathbb{A}^n$.*

Let $\mathfrak{M}$ be a monster model of T.

Definition 11.1.2. *Let K be a subfield of $\mathfrak{M}$. Let $D \subseteq \mathfrak{M}^n$ be a definable set, defined over K. Define the* rank $\mathrm{rk}_K D$ *to be the supremum of* $\mathrm{tr.deg}(\alpha/K)$ *as α ranges over D.*

Lemma 11.1.3.

(a) *If $D \subseteq \mathfrak{M}^n$, then $\mathrm{rk}_K D = n$ if and only if D has non-empty interior.*

(b) *If $D \subseteq \mathfrak{M}^n$ and $1 \leq k \leq n$, then $\mathrm{rk}_K D \geq k$ if and only if $\mathrm{rk}_K \pi(D) = k$ for one of the (finitely many) coordinate projections $\pi : \mathfrak{M}^n \twoheadrightarrow \mathfrak{M}^k$.*

(c) *The rank of D does not depend on the choice of K, and rank is definable in families.*

(d) *If $D \subseteq V$ where V is absolutely irreducible, then* $\operatorname{rk} D = \dim V$ *if and only if $D(\mathfrak{M})$ is Zariski dense in $V(\mathfrak{M}^{alg})$.*

Proof. **(a)** If $\operatorname{rk}_K D < n$, then every tuple α from D lives inside an affine K-variety of dimension less than n. By compactness, D is contained in the union of finitely many affine K-varieties of dimension less than n. This union contains the Zariski closure of D, so D is not Zariski dense. This forces D to have no topological interior, because non-empty polydisks in affine space are Zariski dense. Conversely, if D has no interior, then by Fact 11.1.1, $D \subseteq V$ for some proper subvariety $V \subsetneq \mathbb{A}^n$ with V defined over K. Then $\operatorname{rk}_K D \leq \dim V < n$.

(b) Clear by properties of rank in pregeometries.

(c) Combine (a) and (b).

(d) If $\operatorname{rk} D < \dim V$, then every point in D is contained in an affine K-variety of dimension less than $\dim V$. By compactness, D is contained in the union of finitely many such varieties. This finite union contains the Zariski closure of D, and is strictly smaller than V itself. Conversely, suppose that D is not Zariski dense in V. Let $V' \subsetneq V$ be the Zariski closure of D. As V is absolutely irreducible, $\dim V' < \dim V$. Also, V' is defined over $\mathfrak{M}$ rather than $\mathfrak{M}^{alg}$, because it is the Zariski closure of a set of $\mathfrak{M}$-points. Let L be a small subfield of $\mathfrak{M}$ over which V' and D are defined. Then

$$\operatorname{rk}_K D = \operatorname{rk}_L D \leq \operatorname{rk}_L V' \leq \dim V' < \dim V.$$

□

Corollary 11.1.4. *If $K \leq L$ is an inclusion of small subfields of $\mathfrak{M}$ and α is a finite tuple, we can find $\alpha' \equiv_K \alpha$ with* $\operatorname{tr.deg}(\alpha'/L) = \operatorname{tr.deg}(\alpha'/K)$.

Proof. Let $n = \operatorname{tr.deg}(\alpha/K)$. Let $\Sigma(x)$ be the partial type asserting that $x \equiv_K \alpha$ and that x belongs to no L-variety of dimension less than n. I claim that $\Sigma(x)$ is consistent. Otherwise, there is some formula $\phi(x)$ from $\operatorname{tp}(\alpha/K)$ and some L-varieties $V_1, \ldots, V_m$ of dimension less than n, such that $\phi(\mathfrak{M}) \subseteq \bigcup_{i=1}^m V_i$. But then

$$\operatorname{rk}_K \phi(\mathfrak{M}) = \operatorname{rk}_L \phi(\mathfrak{M}) \leq \max_{1 \leq i \leq m} \dim V_i < n,$$

contradicting the fact that $\alpha \in \phi(\mathfrak{M})$ and $\operatorname{tr.deg}(\alpha/K) \geq n$.

Thus $\Sigma(x)$ is consistent. If α' is a realization, then $\alpha' \equiv_K \alpha$ and

$$\operatorname{tr.deg}(\alpha'/L) \geq n = \operatorname{tr.deg}(\alpha/K) = \operatorname{tr.deg}(\alpha'/K) \geq \operatorname{tr.deg}(\alpha'/L).$$

□

Corollary 11.1.5. *Let L and L' be two fields satisfying $T_\forall$, and suppose they share a common subfield K. Then L and L' can be amalgamated over K in a way which makes L and L' be algebraically independent over K.*

Proof. By quantifier elimination, we may as well assume that L and L' and K live inside a monster model $\mathfrak{M} \models T$. By Corollary 11.1.4 and compactness, we can extend $\mathrm{tp}(L/K)$ to L' in such a way that any realization is algebraically independent from L' over K. □

Definition 11.1.6. *Let $K \leq L$ be an inclusion of fields. Say that K is* relatively separably closed *in L if every $x \in L \cap K^{alg}$ is in the perfect closure of K.*

This is a generalization of K being relatively algebraically closed in L; in characteristic zero these two concepts are the same. Note that if we embed L into a monster model $\mathfrak{M}$ of ACF, then K is relatively separably closed in L if and only if $\mathrm{dcl}(K) = \mathrm{acl}(K) \cap \mathrm{dcl}(L)$ if and only if $\mathrm{tp}(L/K)$ is stationary. From this, one gets

Fact 11.1.7. *Let $L \geq K \leq L'$ be (pure) fields. Suppose that K is relatively separably closed in L or L'. Then there is only one way to amalgamate L and L' over K in such a way that L and L' are algebraically independent over K.*

Fact 11.1.8. *If K is relatively separably closed in L and α is a tuple from L, and V is the variety over K of which α is the generic point, then V is absolutely irreducible.*

11.1.1 Dense formulas

In this section, T continues to be one of ACVF, RCF, or pCF.

Definition 11.1.9. *Let K be a model of $T_\forall$. Let V be an absolutely irreducible affine variety defined over K. Let $\phi(x)$ be a quantifier-free formula with parameters from K, defining a subset of V in any/every model of T extending K. Say that $\phi(x)$ is* V-dense *if $\mathrm{rk}\,\phi(\mathfrak{M}) = \dim V$. Here $\mathfrak{M}$ is a monster model of T extending K.*

The choice of $\mathfrak{M}$ is irrelevant by quantifier-elimination in T and by Lemma 11.1.3(c).

Lemma 11.1.10. *Let K be a model of $T_\forall$, L be a model of T extending K, and V be an absolutely irreducible variety defined over K. For a quantifier-free K-formula $\phi(x)$, the following are equivalent:*

(a) *$\phi(x)$ is V-dense.*

(b) *$\phi(L)$ is Zariski dense in $V(L^{alg})$.*

(c) *We can extend the $T_\forall$-structure on K to the function field $K(V)$ in such a way that the generic point of V in $K(V)$ satisfies $\phi(x)$.*

Proof. **(a)** $\implies$ **(b)** Suppose $\phi(x)$ is V-dense. Let W be the Zariski closure of $\phi(L)$ in $V(L^{alg})$. Then W is defined over L rather than L^{alg}, because W is the Zariski closure of some L-points. Therefore it makes sense to think of W as a definable set. If $\mathfrak{M}$ is a monster model of T extending L, then $\dim V = \operatorname{rk}\phi(\mathfrak{M}) \leq \operatorname{rk} W \leq \dim W \leq \dim V$. Therefore $\dim W = \dim V$. As V is absolutely irreducible, $W = V$.

(b) $\implies$ **(a)** Let $\mathfrak{M}$ be a monster model of T extending L, and let $n = \dim V$. If $\phi(x)$ is not V-dense, then every element of $\phi(\mathfrak{M})$ has transcendence degree less than n over K. By compactness, $\phi(\mathfrak{M})$ is contained in a finite union of K-definable varieties of dimension less than n. We may assume these varieties are closed subvarieties of V. Of course $\phi(L)$ is also contained in this union, which is clearly a Zariski closed proper subset of V. So $\phi(L)$ is not Zariski dense.

(a) $\implies$ **(c)** Embed K into a monster model $\mathfrak{M}$. Let α be a point in $\phi(\mathfrak{M}) \subseteq V(\mathfrak{M})$ with $\operatorname{tr.deg}(\alpha/K) = \operatorname{rk}\phi(\mathfrak{M}) = \dim V$. Then α is a generic point on V, i.e., $K(\alpha) \cong K(V)$. And α satisfies $\phi(x)$.

(c) $\implies$ **(a)** Embed $K(V)$ into a monster model $\mathfrak{M}$. Let α denote the generic point of V, so that $\mathfrak{M} \models \phi(\alpha)$ holds. Clearly $\operatorname{tr.deg}(\alpha/K) = \dim V$. Thus $\operatorname{rk}_K \phi(\mathfrak{M}) \geq \operatorname{tr.deg}(\alpha/K) = \dim V$, implying V-density of $\phi(x)$.

□

Lemma 11.1.11. *Let L be a model of ACVF, and let $V \subseteq \mathbb{A}^n$ be an irreducible affine variety over L. Suppose $0 \in V$. Let $\mathcal{O}_L^n$ be the closed unit polydisk in $\mathbb{A}^n$. Then $\mathcal{O}_L^n \cap V$ is Zariski dense in V.*

This Lemma is essentially Lemma 1.1 in [18], but we will give give a more elementary proof based on the proof of Proposition 4.2.1 in [20].

Proof. Let $L(\alpha)$ be the function field of V, obtained by adding a generic point α of V to the field L. By the implication (c) $\implies$ (b) of Lemma 11.1.10 applied in the case where $\phi(x)$ is the formula defining $\mathcal{O}_L^n \cap V$, it suffices to extend the valuation on L to $L(\alpha)$ in such a way that every coordinate of α has nonnegative valuation.

Now $L[\alpha]$ is the coordinate ring of V, so the fact that $0 \in V$ implies that there is an L-algebra homomorphism $L[\alpha] \to L$ sending every coordinate of α to zero. This yields an $\mathcal{O}_L$-algebra homomorphism $f : \mathcal{O}_L[\alpha] \to \mathcal{O}_L$ sending every coordinate of α to 0. Let $\mathfrak{m}$ be the maximal ideal of $\mathcal{O}_L$, and let $\mathfrak{p} = f^{-1}(\mathfrak{m})$. Then $\mathfrak{p}$ is a prime ideal, and $\mathfrak{p} \cap \mathcal{O}_L = \mathfrak{m}$. Also, as f kills the coordinates of α, the coordinates of α live in $\mathfrak{p}$.

Since $\mathcal{O}_L[\alpha]$ is a domain, there is a valuation v' on $L(\alpha)$, the fraction field of $\mathcal{O}_L[\alpha]$, with the following properties:

- Every element of $\mathfrak{p}$ has positive valuation. In particular, the elements of $\mathfrak{m}$ and the coordinates of α have positive valuation.

- Every element of $\mathcal{O}_L[\alpha] \setminus \mathfrak{p}$ has valuation zero. In particular, the elements of $\mathcal{O}_L^\times = \mathcal{O}_L \setminus \mathfrak{m}$ have valuation zero.

(Indeed, it is a general fact that if S is a domain and $\mathfrak{p}$ is a prime ideal, then there is a valuation on the fraction field of S which assigns a positive valuation to elements of $\mathfrak{p}$ and a vanishing valuation to elements of $S \setminus \mathfrak{p}$. To find such a valuation, take a valuation ring in $\mathrm{Frac}(S)$ dominating the local ring $S_{\mathfrak{p}}$.)

The resulting valuation on $L(\alpha)$ extends the valuation on L, because it assigns positive valuation to elements in $\mathfrak{m}$, and zero valuation to elements in $\mathcal{O}_L \setminus \mathfrak{m}$. Also, the valuation of any coordinate of α is positive, hence non-negative, so α lives in the closed unit polydisk. □

Lemma 11.1.12. *Let V be an absolutely irreducible affine variety over $K \models T_\forall$, and let $\phi(x)$ be a quantifier-free K-formula. Let L be a model of T extending K. Suppose $\phi(x)$ defines an open subset of $V(L)$.*

(a) *If T is ACVF, then $\phi(x)$ is V-dense if and only if $\phi(L)$ is non-empty.*

(b) *In general, $\phi(x)$ is V-dense if $\phi(L)$ contains a smooth point of V.*

Proof. **(a)** If $\phi(x)$ is V-dense, then certainly $\phi(L)$ is non-empty. Conversely, suppose $\phi(L)$ is non-empty. Let p be a point in $\phi(L)$ and let U be an open neighborhood of p, with $U \cap V \subseteq \phi(L)$. There is some L-definable affine transformation f which sends p to the origin and moves U so as to contain the closed unit polydisk. Then $f(U \cap V) = f(U) \cap f(V)$ is Zariski dense in $f(V)$, by Lemma 11.1.11. So $\phi(L) \supseteq U \cap V$ is Zariski dense in V. Thus $\phi(x)$ is V-dense, by Lemma 11.1.10.

(b) If $\phi(x)$ is V-dense, then $\phi(L)$ contains a smooth point of V, because the smooth locus of V is a Zariski dense Zariski open. Conversely, suppose $\phi(L)$ contains a smooth point p. Note that L is perfect. The tangent space T_pV is L-definable. By Hilbert's Theorem 90, there is an L-definable basis of T_pV. Therefore, after applying an L-definable change of coordinates, we may assume T_pV is horizontal. By the implicit function theorem, V then looks locally around p like the graph of a function. In particular, the coordinate projection maps a neighborhood of p homeomorphically to an open subset of affine n-space, where $n = \dim V$. By Lemma 11.1.3, this ensures that any neighborhood of p, such as $\phi(L)$, has rank at least n. So $\phi(x)$ is V-dense.

□

Remark 11.1.13. *Here, and in Lemma 11.3.3 below, we are using the* model-theoretic *version of Hilbert's theorem 90. This folk theorem says that if $M \models ACF$, if K is a perfect subfield of M, and if V is a K-definable M-vector space, then V admits a K-definable basis. It is an easy exercise to derive this fact from standard Galois descent of vector spaces, which is part of Grothendieck's modern generalization of Hilbert's original Theorem 90. (See III.4.10 in [53].)*

Lemma 11.1.14. *Let V be an absolutely irreducible affine variety over $K \models T_\forall$, and let $\phi(x)$ be a quantifier-free K-formula that is V-dense. Then there is a quantifier-free K-formula $\psi(x)$ that is also V-dense, such that in any/every $L \models T$ extending K, $\psi(L)$ is a topologically open subset of $V(L)$, and $\psi(L) \subseteq \phi(L)$.*

Proof. Choose some monster model $\mathfrak{M} \models T$ extending K and let $\psi(\mathfrak{M})$ pick out the topological interior of $\phi(\mathfrak{M})$ inside $V(\mathfrak{M})$. By quantifier-elimination, we can take $\psi(x)$ to be quantifier-free with parameters from K. It remains to show that $\psi(\mathfrak{M})$ is V-dense. Let $\alpha \in \phi(\mathfrak{M})$ have transcendence degree n over K, where $n = \dim V$. By Fact 11.1.1, $\phi(\mathfrak{M})$ can be written as a finite union of finite intersections of K-definable opens and varieties. Let X be one of these finite intersections, containing α. So $X = W \cap U$ for some K-variety W and some K-definable open U. As $\alpha \in W$ and α is a generic point on V, we must have $V \subseteq W$. Then

$$\alpha \in V \cap U \subseteq W \cap U \subseteq \phi(\mathfrak{M}).$$

But $V \cap U$ is a relative open in $V(\mathfrak{M})$, so it must be part of $\psi(\mathfrak{M})$. In particular, $\alpha \in \psi(\mathfrak{M})$. As $\operatorname{tr.deg}(\alpha/K) = n$, we conclude that $\psi(x)$ is V-dense. □

11.1.2 Forking and Dividing

We continue to work in one of ACVF, RCF, or pCF. Recall that RCF and pCF have definable Skolem functions in the home sort. Thus if S is a subset of the home sort, then $\operatorname{acl}(S) = \operatorname{dcl}(S)$ is a model. In ACVF, $\operatorname{acl}(S)$ is the algebraic closure of S, which is a model unless $\operatorname{acl}(S)$ is trivially valued.

We will always be working in the home sort, rather than working with imaginaries.

Lemma 11.1.15. *Let S be a set (in the home sort) and let $\phi(x; b)$ be a formula. Then $\phi(x; b)$ forks over S if and only if it divides over S.*

Proof. Indiscernibility over S is the same thing as indiscernibility over $\operatorname{acl}(S)$, so $\phi(x; b)$ divides over S if and only if it divides over $\operatorname{acl}(S)$. Similarly, $\phi(x; b)$ forks over S if and only if it forks over $\operatorname{acl}(S)$. So we may assme $S = \operatorname{acl}(S)$. If T is RCF or pCF, then S is a model, and therefore forking and dividing agree over S by Theorem 1.1 of [10]. If T is ACVF, then forking and dividing agree over all sets, by Corollary 1.3 in [10]. □

We use $\downarrow$ to denote non-forking or non-dividing, and $\downarrow^{ACF}$ to denote algebraic independence.

Lemma 11.1.16. *Let $\mathfrak{M}$ be a monster model of T, and let B, C be small subsets of $\mathfrak{M}$, with B finite. Then we can find a sequence $B_0, B_1, B_2, \ldots$ in $\mathfrak{M}$ that is C-indiscernible, such that $B_0 = B$ and $B_i \downarrow_C^{ACF} B_{<i}$ for every i.*

Proof. We may assume that B is ordered as a tuple in such a way that the first k elements of B are a transcendence basis of B over C. Construct a sequence $D_0, D_1, \ldots$ of realizations

of $\mathrm{tp}(B/C)$ such that $D_i \downarrow_C^{ACF} D_{<i}$ for every i. This is possible by using Corollary 11.1.4 to extend $\mathrm{tp}(B/C)$ to a type over $CD_{<i}$ having the same transcendence degree over $CD_{<i}$ as over C. Let $B_0, B_1, B_2, \ldots$ be a C-indiscernible sequence modeled on $D_0, D_1, \ldots$. Let $\pi(X)$ pick out the first k elements of a tuple X. Then $\pi(D_0)^\frown\pi(D_1)^\frown\pi(D_2)^\frown\cdots$ is an algebraically independent sequence of singletons over C. This is part of the EM-type of the D_i over C, so it is also true that $\pi(B_0)^\frown\pi(B_1)^\frown\pi(B_2)^\frown\cdots$ is an algebraically independent sequence of singletons over C. Since $D_i \equiv_C B$ for every i, we also have $B_i \equiv_C B$ for every i. Thus $\pi(B_i)$ is a transcendence basis for B_i over C, and we conclude that $B_i \downarrow_C^{ACF} B_{<i}$ for every i. Finally, moving the B_i by an automorphism over C, we may assume that $B_0 = B$. □

Lemma 11.1.17. *$A \downarrow_C B$ implies $A \downarrow_C^{ACF} B$.*

Proof. Assume $A \downarrow_C B$. By Lemma 11.1.16, we can find a sequence $B_0, B_1, B_2, \ldots$ of realizations of $\mathrm{tp}(B/C)$, indiscernible over C, and satisfying $B_i \downarrow_C^{ACF} B_{<i}$ for every i. Suppose for the sake of contradiction that in some ambient model of ACF, $\mathrm{tp}(A/BC)$ contains a formula $\phi(X;Y)$ which divides (in the ACF sense) over C. By quantifier elimination in ACF, we may assume ϕ is quantifier-free. In stable theories such as ACF, dividing is witnessed in any Morley sequence. In particular

$$\bigwedge_i \phi(X; B_i)$$

is inconsistent in the ambient model of ACF, hence inconsistent in the original smaller structure. Thus $\phi(X;B)$ forks and divides over C in the original structure, a contradiction. □

Lastly, we show that dividing is always witnessed by an algebraically independent sequence.

Lemma 11.1.18. *If a formula $\phi(x;a)$ divides over a set A, then the dividing is witnessed by an A-indiscernible sequence $a = a_0, a_1, a_2, \ldots$ such that $a_i \downarrow_A^{ACF} a_{<i}$ for every i.*

Proof. Apply Claim 3.10 of [10] with the abstract independence relation taken to be $\downarrow$ (nonforking). Forking satisfies (1)-(7) of [10] Definition 2.9 by Fact 2.12(5) of [10]. And A is an extension base for forking by Lemma 11.1.15 above and Theorem 1.1 of [10] (or by Fact 2.14 of [10] in the cases other than pCF). So Claim 3.10 of [10] is applicable. Consequently we get a model M containing A, a global type p extending $\mathrm{tp}(a/M)$, $\downarrow$-free over A, such that any/every Morley sequence generated by p over M witnesses the dividing of $\phi(x;a)$. Because $\downarrow$ is stronger than Lascar invariance, any such Morley sequence will be M-indiscernible, hence A-indiscernible. Because $\downarrow$ is stronger than algebraic independence (Lemma 11.1.17), and p is $\downarrow$-free over A, any Morley sequence $a_0, a_1, \ldots$ generated by p will be algebraically independent over A. Specifically, $a_i \models p|_{Ma_{<i}}$, so as p is $\downarrow$-free over A, $a_i \downarrow_A Ma_{<i}$, and hence $a_i \downarrow_A^{ACF} a_{<i}$. □

11.2 The Model Companion

Now we turn our attention to fields with several valuations, several orderings, and several p-valuations. For $1 \leq i \leq n$, let T_i be one of ACVF, RCF, or pCF (in the same languages as in the previous section). Let $\mathcal{L}_i$ denote the language of T_i; assume that $\mathcal{L}_i \cap \mathcal{L}_j = \mathcal{L}_{rings}$ for $i \neq j$. Let T^0 be $\bigcup_{i=1}^n (T_i)_\forall$, the theory that would be denoted $((T_1)_\forall, (T_2)_\forall, \ldots, (T_n)_\forall)$ in van den Dries's notation. Technically speaking, models of T^0 should be allowed to be domains, rather than fields. However, we will assume that T^0 also includes the field axioms, sweeping domains under the rug.

One essentially knows that T^0 has a model companion T by Chapter III of van den Dries's thesis [16]. We will quickly reprove the existence of T in this section, expressing the axioms of the model companion in a more geometric and less syntactic form, and also including the case of positive characteristic explicitly.

11.2.1 The Axioms

Consider the following axioms that a model K of T^0 could satisfy:

A1: K is existentially closed with respect to finite extensions, i.e., if L/K is a finite algebraic extension and $L \models T^0$, then $L = K$.

A1': For every irreducible polynomial $P(X) \in K[X]$ of degree greater than 1, there is some $1 \leq i \leq n$ such that $P(x) = 0$ has no solution in any/every model of T_i extending $K \upharpoonright \mathcal{L}_i$.

A2(m): Let V be an m-dimensional absolutely irreducible variety over K. For $1 \leq i \leq n$, let $\phi_i(x)$ be a V-dense quantifier-free $\mathcal{L}_i$-formula with parameters from K. Then $\bigcap_{i=1}^n \phi_i(K) \neq \emptyset$.

A2($\leq m$): A2(m') holds for all $m' \leq m$.

A2: A2(m) holds, for all m

Remark 11.2.1. *For $K \models T^0$, A1 and A1' are equivalent.*

Proof. Suppose K satisfies A1, and $P(X) \in K[X]$ is irreducible of degree greater than 1. Suppose that for every $1 \leq i \leq n$, there is a solution α_i of $P(x) = 0$ in a model $M_i \models T_i$ extending $K \upharpoonright \mathcal{L}_i$. Then we can extend the $\mathcal{L}_i$-structure from K to $K(\alpha) \cong K[X]/P(X)$. Because this holds for every i, we can endow $K[X]/P(X)$ with the structure of a model of T^0. By A1, $K[X]/P(X)$ must be K, so $P(X)$ has degree 1.

Conversely, suppose K satisfies A1' but not A1. Let L/K be a counterexample to A1, and take some $\alpha \in L \setminus K$. Let $P(X)$ be the irreducible polynomial of α over K. This polynomial must have degree greater than 1. For each i let M_i be a model of T_i extending $L \upharpoonright \mathcal{L}_i$. Then $P(x) = 0$ has a solution in L, hence in M_i, which is a model of T_i extending K. This contradicts A1'. □

Lemma 11.2.2. *Let K be a model of T^0, and $m \geq 1$. The following are equivalent:*

(a) *For every model L of T^0 extending K, for every tuple α from L with $\mathrm{tr.deg}(\alpha/K) \leq m$, the quantifier-free type $\mathrm{qftp}(\alpha/K)$ is finitely satisfiable in K.*

(b) *K satisfies A1 and A2($\leq m$).*

Proof. (a) $\implies$ (b) For A1, suppose that L/K is a finite extension, and $L \models T$. If $\alpha \in L$, then α is algebraic over K, so $\mathrm{tr.deg}(\alpha/K) = 0 \leq m$. By (a), the quantifier-free type of α is realized in K. So the irreducible polynomial of α over K has a zero in K, implying $\alpha \in K$. As $\alpha \in L$ was arbitrary, $L = K$.

For A2(m'), let V be an m'-dimensional absolutely irreducible variety over K. For $1 \leq i \leq n$, let $\phi_i(x)$ be a V-dense quantifier-free $\mathcal{L}_i$-formula with parameters from K. By Lemma 11.1.10(c), we can extend the $\mathcal{L}_i$-structure to $K(V)$ in such a way that the generic point satisfies $\phi_i(x)$. Doing this for all i, we make $K(V)$ be a model of T^0 extending K, such that if $\alpha \in K(V)$ denotes the generic point, then $\bigwedge_{i=1}^n \phi_i(\alpha)$ holds. Now $\mathrm{tr.deg}(\alpha/K) = \dim V \leq m$, so by (a), $\mathrm{qftp}(\alpha/K)$ is finitely satisfiable in K. In particular, the formula $\bigwedge_{i=1}^n \phi(x)$ is satisfiable in K, which is the conclusion of A2(m').

(b) $\implies$ (a). Suppose L is a model of T^0 extending K and α is a tuple from L, with $\mathrm{tr.deg}(\alpha/K) \leq m$. By A1, K is relatively algebraically closed in L. Let V be the K-variety of which α is a generic point. Then V is absolutely irreducible, by Fact 11.1.8. Also, $m' := \dim V = \mathrm{tr.deg}(\alpha/K) \leq m$. Let $\psi(x)$ be a statement in $\mathrm{qftp}(\alpha/K)$. We want to show that ψ is satisfied by an element of K. We may assume that $\psi(x)$ includes the statement that $x \in V$. By Fact 11.1.1, $\psi(x)$ is a positive boolean combination of statements of the form

- $x \in W$, for some K-definable affine variety W. Since we intersected $\psi(x)$ with V, we may assume $W \subseteq V$.
- $\theta(x)$, where $\theta(x)$ is a quantifier-free $\mathcal{L}_i$-formula for some i, such that $\theta(L)$ is an open subset of the ambient affine space, for any/every $L \models T_i$ extending $K \upharpoonright \mathcal{L}_i$.

Writing $\psi(x)$ as a disjunction of conjunctions of such statements, and replacing $\psi(x)$ by whichever disjunct α satisfies, we may assume that $\psi(x)$ is a conjunction of such statements. An intersection of K-varieties is a K-variety, and an intersection of open subsets of affine space is an open subset of affine space, so we may assume

$$\psi(x) \equiv \text{“}x \in W\text{”} \wedge \bigwedge_{i=1}^{n} \phi_i(x),$$

where W is some K-variety contained in V, and where $\phi_i(x)$ is a quantifier-free $\mathcal{L}_i$-formula defining an open subset of the ambient affine space, when interpreted in any/every model of T_i extending $K \upharpoonright \mathcal{L}_i$.

Because α satisfies $\psi(x)$, and α is a generic point of V, W must be V. Rewrite ψ as $\bigwedge_{i=1}^n \phi_i'(x)$, where each $\phi_i'(x)$ asserts that $x \in V$ and $\phi_i(x)$ holds. Because K satisfies axiom

A2(m'), $\psi(x)$ will be satisfiable in K as long as $\phi_i'(x)$ is V-dense for each i. But note that L provides a way of extending the $\mathcal{L}_i$-structure from K to $K(\alpha) \cong K(V)$ in such a way that $\phi_i'(\alpha)$ holds, so ϕ_i' is V-dense by Lemma 11.1.10(c). □

Theorem 11.2.3. *The theory T^0 has a model companion T, whose models are exactly the $K \models T^0$ satisfying A1 and A2.*

Proof. It is well known that a model K is existentially closed if and only if for every model L extending K and for every tuple α from L, the quantifier-free type qftp(α/K) is finitely satisfiable in K. So by Lemma 11.2.2, a model of T^0 is existentially closed if and only if it satisfies A1 and A2. By basic facts about model companions of $\forall\exists$-theories, it remains to show that A1 and A2 are first order. For A1, this comes from Remark 11.2.1, because A1' is first order by quantifier-elimination in the T_i. Axiom A2 is first order by quantifier-elimination in the T_i, by Lemma 11.1.3(c), and by the fact that absolute irreducibility is definable by a quantifier-free formula in the language of fields (this is well-known and proven in Chapter IV of [16]). □

Henceforth, we will use T to denote the model companion. Also, we will use $T_\forall$ instead of T^0, sweeping the distinction between domains and fields under the rug.

We make several remarks about the axioms:

Remark 11.2.4. *In the case where T_i is ACVF for $i > 1$, axiom A1 merely says that $K \restriction \mathcal{L}_1$ is a model of T_1, i.e., is algebraically closed or real closed or p-adically closed.*

Remark 11.2.5. *In Axiom A2(m), it suffices to consider the case of smooth V. If V is* not *smooth, one can find an open subvariety V' of V which is smooth, and which is isomorphic to an affine variety. (Use the facts that the smooth locus of an irreducible variety is a Zariski dense Zariski open, and that the affine open subsets of a scheme form a basis for its topology.) If $\phi_i(x)$ is V-dense, then $\phi_i(x) \wedge$ "$x \in V'$" is V'-dense, essentially by Lemma 11.1.10(b). Then applying the smooth case of A2(m) to V' yields a point in V' satisfying $\bigwedge_{i=1}^n \phi_i(x)$.*

Remark 11.2.6. *In Axiom A2, it suffices to consider V-dense formulas $\phi_i(x)$ such that $\phi_i(L)$ defines an* open *subset of $V(L)$ for any/every $L \models T$ extending $K \restriction \mathcal{L}_i$. This follows by Lemma 11.1.14.*

Remark 11.2.7. *We can combine the previous two remarks. Then Lemma 11.1.12(b), yields the following restatement of A2(m): if V is an absolutely irreducible m-dimensional smooth affine variety defined over K, and if $\phi_i(x)$ is a quantifier-free $\mathcal{L}_i$-formula over K for each $1 \leq i \leq n$, and if $\phi_i(K_i)$ is a* non-empty *open subset of $V(K_i)$ for any/every $K_i \models T$ extending $K \restriction \mathcal{L}_i$, then $\bigcap_{i=1}^n \phi_i(K) \neq \emptyset$.*

Remark 11.2.8. *If every T_i is ACVF, then A1 merely says that K is algebraically closed. Consequently, in Remark 11.2.7 the K_i can be taken to be K itself. Thus A2(m) ends up being equivalent to the statement that if V is a smooth irreducible m-dimensional affine variety, and $\phi_i(x)$ is a quantifier-free $\mathcal{L}_i$-formula defining a non-empty open subset of V for*

$1 \leq i \leq n$, then $\bigcap_{i=1}^{n} \phi_i(K)$ is non-empty. Even more concisely, this means that for every smooth m-dimensional variety V, the diagonal map $V(K) \to \prod_{i=1}^{n} V(K)$ has dense image in the product topology, using the topology from the i^{th} valuation for the i^{th} entry in the product.

In fact, in Section 11.3, we will see that it suffices to check the case of $V = \mathbb{A}^1$, the affine line(!)

11.2.2 Quantifier-Elimination up to Algebraic Covers

As in the previous section, $T_\forall$ is the theory of fields with $(T_i)_\forall$ structure for each $1 \leq i \leq n$, and T is the model companion of $T_\forall$.

Lemma 11.2.9. *Let K be a model of $T_\forall$. Let L and L' be two models of $T_\forall$ extending K. Suppose that K is relatively separably closed in L or L' (Definition 11.1.6). Then L and L' can be amalgamated over K, and this can be done in such a way that L and L' are algebraically independent over K.*

Proof. For each $1 \leq i \leq n$, we can find some amalgam $M_i \models (T_i)_\forall$ of $L \restriction \mathcal{L}_i$ and $L' \restriction \mathcal{L}_i$ over $K \restriction \mathcal{L}_i$, by Corollary 11.1.5. The resulting compositums LL' must be isomorphic on the level of fields, by Fact 11.1.7. Consequently, we can endow the canonical field LL' with a $(T_i)_\forall$-structure extending those on L and L', for each i. This gives LL' the structure of a $T_\forall$-model. And L and L' are algebraically independent inside LL'. □

Corollary 11.2.10. *Let K be a model of $T_\forall$ and let L be a model of T extending K. Then K is relatively algebraically closed in L if and only if K satisfies axiom A1. (In particular, this does not depend on L.)*

Proof. If K satisfies axiom A1, then obviously K is relatively algebraically closed in L. Conversely, suppose that K is relatively algebraically closed in L but does not satisfy A1. Then there is some model L' of $T_\forall$ extending K, with L'/K finite and $L' \neq K$. By Lemma 11.2.9, we can amalgamate L and L' over K. Embed the resulting compositum LL' in a model M of T. Because T is model-complete, $L \preceq M$. Now choose some $\alpha \in L' \setminus K$. The irreducible polynomial of α over K has a root in M, and hence has a root in L, contradicting the fact that K is relatively algebraically closed in L. □

Corollary 11.2.11. *Let K be model of $T_\forall$, and suppose K satisfies A1. Then the type of K is determined, i.e., if L and L' are two models of T extending K, then K has the same type in L and L'. Equivalently, the diagram of K implies the elementary diagram of K, modulo the axioms of T.*

Proof. By Corollary 11.2.10, K is relatively algebraically closed in L and L'. So we can amalgamate L and L' over K, by Lemma 11.2.9. If M is a model of T extending LL', then by model completeness $L \preceq M \succeq L'$, ensuring that K has the same type in each. □

Corollary 11.2.12. *In models of T, field-theoretic algebraic closure agrees with model-theoretic algebraic closure.*

Proof. Let M be a model of T. Let S be a subset of M. Let K be the field-theoretic algebraic closure of S, i.e., the relative algebraic closure of S in M. By Lemma 11.2.9, we can amalgamate M and a copy M' of M over K in a way that makes M and M' be algebraically independent over K. Embedding MM' into a model N of T, and using model completeness, we get $M \preceq N \succeq M'$. Now $\mathrm{acl}(S)$ is the same when computed in M, N, or M'. In particular, $\mathrm{acl}(S) \subseteq M \cap M'$. Since M and M' are algebraically independent over K and K is relatively algebraically closed in each, $M \cap M' = K$. Thus $\mathrm{acl}(S) \subseteq K$. Obviously $K \subseteq \mathrm{acl}(S)$. □

For K a field, let $\mathrm{Abs}(K)$ denote the algebraic closure of the prime field in K.

Corollary 11.2.13. *Two models $M_1, M_2 \models T$ are elementarily equivalent if and only if $\mathrm{Abs}(M_1)$ and $\mathrm{Abs}(M_2)$ are isomorphic as models of $T_\forall$.*

Proof. If M_1 and M_2 are elementarily equivalent, we can embed them as elementary substructures into a third model $M_3 \models T$. Then $\mathrm{Abs}(M_1) = \mathrm{Abs}(M_3) = \mathrm{Abs}(M_2)$, so certainly $\mathrm{Abs}(M_1)$ is isomorphic to $\mathrm{Abs}(M_2)$.

Conversely, suppose $\mathrm{Abs}(M_1) \cong \mathrm{Abs}(M_2)$. Then, as $\mathrm{Abs}(M_1)$ is relatively algebraically closed in M_1 and in M_2, it follows by Corollaries 11.2.10 and 11.2.11 that we can amalgamate M_1 and M_2 over $\mathrm{Abs}(M_1)$. Embedding the resulting compositum into a model of T and using model completeness, we get $M_1 \equiv M_2$. □

Corollary 11.2.14. *Suppose $T_1 \neq ACVF$ and T_i is $ACVF$ for $i > 1$. Consider the expanded theory where we add in symbols for every zero-definable T_1-definable function. (This makes sense because if $M \models T$, then $M \restriction \mathcal{L}_1 \models T_1$, by Remark 11.2.4.) Then T has quantifier-elimination.*

Proof. After adding in these new symbols, a substructure is the same as a subfield K closed under all T_1-definable functions. As RCF and pCF have definable Skolem functions, this is equivalent to $K \restriction \mathcal{L}_1$ being a model of T_1, which is equivalent to K satisfying axiom A1, as noted in Remark 11.2.4. Now apply Corollary 11.2.11 to get substructure completeness, which is the same thing as quantifier-elimination. □

This probably also holds if $T_i \neq ACVF$ for more than one i, though the extra functions would become partial functions.

Without adding in extra symbols, quantifier elimination fails. But we still get quantifier-elimination up to algebraic covers, in a certain sense.

Theorem 11.2.15. *In T, every formula $\phi(x)$ is equivalent to one of the form*

$$\exists y \quad P(y, x) = 0 \wedge \psi(y, x), \tag{11.1}$$

where y is a singleton, $\psi(y, x)$ is quantifier-free, and $P(y, x)$ is a polynomial in y and the coordinates of x, with integer coefficients, monic as a poynomial in y.

Proof. Let $\Sigma(x)$ be the set of all formulas of the form (11.1). First we observe that $\Sigma(x)$ is closed under disjunction, because

$$(\exists y \quad P(y,x) = 0 \wedge \psi(y,x)) \vee (\exists y \quad Q(y,x) = 0 \wedge \psi'(y,x))$$

is equivalent to

$$\exists y \quad P(y,x)Q(y,x) = 0 \wedge \psi''(y,x),$$

where $\psi''(y,x)$ is the quantifier-free formula

$$(P(y,x) = 0 \wedge \psi(y,x)) \vee (Q(y,x) = 0 \wedge \psi'(y,x)).$$

Now given a formula $\phi(x)$, not quantifier-free, let $\Sigma_0(x)$ be the set of formulas in $\Sigma(x)$ which imply $\phi(x)$, i.e.,

$$\Sigma_0(x) = \{\sigma(x) \in \Sigma(x) : T \vdash \forall x \sigma(x) \to \phi(x)\}.$$

Of couse $\Sigma_0(x)$ is closed under disjunction. It suffices to show that $\phi(x)$ implies a finite disjunction of formulas in $\Sigma_0(x)$, because then $\phi(x)$ implies and is implied by a formula in $\Sigma_0(x)$.

Suppose for the sake of contradiction that $\phi(x)$ does not imply a finite disjunction of formulas in $\Sigma_0(x)$. Then the partial type

$$\{\phi(x)\} \cup \{\neg\sigma(x) : \sigma(x) \in \Sigma_0(x)\}$$

is consistent with T. Let M be a model of T containing a tuple α realizing this partial type. So $\phi(\alpha)$ holds in M, but not because of any formula of the form (11.1).

Let R be the ring $\mathbb{Z}[\alpha] \subseteq M$. Let $K \subseteq M$ be the smallest perfect field containing R; note that M itself is perfect so this makes sense. Indeed, if every T_i is ACVF, then M is algebraically closed by Remark 11.2.4. Otherwise, one of the T_i's is RCF or pCF, making M be characteristic zero.

Let $\overline{K}$ be the relative algebraic closure of K (or equivalently, α) inside M. By Corollaries 11.2.10 and 11.2.11, the diagram of $\overline{K}$ implies the elementary diagram of $\overline{K}$. In particular, the diagram of $\overline{K}$ implies $\phi(\alpha)$. By compactness, the diagram of L implies $\phi(\alpha)$, for some finite extension L of K. Because K is perfect, $L = K(\beta)$ for some *singleton* β. Multiplying β by an appropriate element from R, we may assume that β is integral over R. Note that L is perfect, because it is an algebraic extension of a perfect field, and in fact L is the smallest perfect field containing α and β.

As the diagram of L implies $\phi(\alpha)$, so does the diagram of $\mathbb{Z}[\alpha, \beta]$, by Lemma 11.2.16 below. By compactness, there is some quantifier-free formula $\psi(y,x)$ which is true of (β, α) such that

$$T \vdash \forall y \forall x \psi(y,x) \to \phi(x).$$

Let $P(y,x)$ be the polynomial witnessing integrality of β over R. Then clearly

$$T \vdash \forall x \, (\exists y P(y,x) = 0 \wedge \psi(y,x)) \to \phi(x),$$

so $\exists y P(y,x) = 0 \wedge \psi(y,x)$ is in $\Sigma_0(x)$, contradicting the fact that it holds of α in M. $\square$

Lemma 11.2.16. *Let M be a model of T and R be a subring of M. Let $K \subseteq M$ be the smallest perfect field containing R. Let α be a tuple from R, and $\phi(x)$ be a formula such that $M \models \phi(\alpha)$. If T and the diagram of K imply $\phi(\alpha)$, then T and the diagram of R imply $\phi(\alpha)$.*

Proof. If not, then there is a model N of T extending R, in which $\phi(\alpha)$ fails to hold. This model N must not satisfy the diagram of K. Now N certainly contains a copy of the pure field K, because the fraction field and perfect closure of a domain are unique. Consequently, there must be at least two ways to extend the T-structure from R to K, one coming from M and one coming from N. But this is absurd, because each valuation/ordering/p-valuation on R extends uniquely to K, by quantifier elimination in the T_i. □

11.2.3 Simplifying the axioms down to curves

Lemma 11.2.17. *Let K be an $\aleph_1$-saturated and $\aleph_1$-strongly homogeneous model of $T_\forall$ satisfying axioms A1 and A2(1). Let $\mathfrak{M}$ be a monster model of T extending K. Let S be a countable subset of K and α be a countable tuple from $\mathfrak{M}$. Then $\mathrm{tp}(\alpha/S)$ is realized in K.*

Proof. Consider the following statements:

- A_k: if α is a finite tuple from $\mathfrak{M}$, with $\mathrm{tr.deg}(\alpha/S) \leq k$, then $\mathrm{qftp}(\alpha/S)$ is realized in K.
- B_k: if α is a countable tuple from $\mathfrak{M}$, with $\mathrm{tr.deg}(\alpha/S) \leq k$, then $\mathrm{qftp}(\alpha/S)$ is realized in K.
- C_k: if α is a countable tuple from $\mathfrak{M}$, with $\mathrm{tr.deg}(\alpha/S) \leq k$, then $\mathrm{tp}(\alpha/S)$ is realized in K.

There are several implications between these statements:

- For each k, A_k implies B_k, by compactness.
- For each k, B_k implies C_k. Indeed, if α is as in C_k, apply B_k to $\alpha' := \mathrm{acl}(\alpha S)$ and use Corollary 11.2.11.
- C_k for all k implies the statement of the Lemma, by compactness.

Finally, observe that C_k and C_j imply C_{k+j}: if α has transcendence degree $k+j$ over S, let β be a subtuple of α with transcendence degree k. Then $\mathrm{tr.deg}(\beta/S) \leq k$ and $\mathrm{tr.deg}(\alpha/\beta S) \leq j$. By C_k, we can apply an automorphism over S to move β inside K. By C_j applied to $\mathrm{tp}(\alpha/\beta S)$, we can then find a further automorphism moving α inside K.

Lemma 11.2.2 and $\aleph_1$-saturation of K imply A_1. By the above comments, this implies C_1, which in turn implies $C_{1+1}, C_3, C_4, \ldots$. By compactness, the Lemma is true. □

Theorem 11.2.18. *A field $K \models T_\forall$ is existentially closed, i.e., a model of T, if and only if it satisfies A1 and A2(1).*

Proof. If K is existentially closed, then certainly K satisfies A1 and A2(1). Conversely, suppose K satisfies A1 and A2(1). Let K' be an $\aleph_1$-saturated $\aleph_1$-strongly homogeneous elementary extension of K. As $K \equiv K'$, it suffices to show that $K' \models T$. Let $\mathfrak{M}$ be a monster model of T, extending K'. It suffices to show that $K' \preceq \mathfrak{M}$. It suffices to show that if D is a non-empty K'-definable subset of $\mathfrak{M}$, then D intersects K'. Let S be a finite subset of K' that D is defined over, and let α be a point in D. By Lemma 11.2.17, $\mathrm{tp}(\alpha/S)$ is realized in K'. Such a realization must live in D. □

Consequently, in checking the axioms one only needs to consider curves. In fact, one only needs to consider smooth curves, by Remark 11.2.5.

11.3 A Special Case

In the case where almost every T_i is ACVF, the axioms can be drastically simplified.

Theorem 11.3.1. *Suppose $T_2, \ldots T_n$ are all ACVF. A model $K \models T_\forall$ is existentially closed (i.e., a model of T) if and only if the following three conditions hold:*

- *$K \upharpoonright \mathcal{L}_1 \models T_1$*
- *Each valuation $v_2, \ldots, v_n$ is non-trivial.*
- *T_i and T_j do not induce the same topology on K, for $i \neq j$.*

For example, if we are considering the theory of ordered valued fields, this says that a model is existentially closed if and only if the field is real closed, the valuation is non-trivial, and the ordering and valuation induce different topologies on K. A field with n valuations is existentially closed if and only if it is algebraically closed and the valuations induce distinct non-discrete topologies on the field. Using this, we can easily see that $\mathbb{Q}^{alg}$ with n distinct valuations is an existentially closed field with n valuations. This surprised me, since I expected the Rumely Local-Global principle (Theorem 1 of [64]) to be necessary in the proof.

Theorem 11.3.1 is *not* model theoretic, and is presumably known to experts in algebraic geometry or field theory.

In the proof of Theorem 11.3.1, we will use A. L. Stone's Approximation Theorem ([69], Theorem 3.4):

Fact 11.3.2. *Let K be a field. Let $t_1, \ldots, t_n$ be topologies on K arising from orderings and non-trivial valuations. Suppose that $t_i \neq t_j$ for $i \neq j$. Then the $\{t_i\}$ are independent, i.e., if U_i is a non-empty t_i-open subset of K for each i, then $\bigcap_{i=1}^n U_i$ is non-empty. Equivalently, the diagonal map $K \to \prod_{i=1}^n K$ has dense image with respect to the product topology, using the topology t_i for the i^{th} term in the product.*

Note that Fact 11.3.2 does not contradict the existence of valuations which refine each other, because two non-trivial valuations which refine each other always induce the same topology. A self-contained model-theoretic proof of Stone Approximation is given in [61], Theorem 4.1.

Also, we will need the following straightforward lemma.

Lemma 11.3.3. *Let K be a model of T. Let C be an affine smooth curve over K, absolutely irreducible. Let $\overline{C}$ be the canonical smooth projective model (as an abstract variety). For each i, let $\phi_i(x)$ be a C-dense quantifier-free $\mathcal{L}_i$-formula with parameters from K. Then we can find a K-definable rational function $f : \overline{C} \to \mathbb{P}^1$ which is non-constant, and has the property that the divisor $f^{-1}(0)$ is a sum of distinct points in $\bigcap_{i=1}^n \phi_i(K)$, with no multipliticities. (In particular, the support of the divisor contains no points from $C(K^{alg}) \setminus C(K)$ and no points from $\overline{C} \setminus C$.)*

Proof. Let g be the genus of $\overline{C}$.

Claim 11.3.4. *We can find $g+1$ distinct points $p_1, \ldots, p_{g+1}$ in $\bigcap_{i=1}^n \phi_i(K) \subseteq C(K)$.*

Proof. For each i, choose a model K_i of T_i extending $K \restriction \mathcal{L}_i$. Then $\phi_i(K_i)$ is Zariski dense in $C(K_i^{alg})$. This (easily) implies that $\phi_i(K_i)^{g+1}$ is Zariski dense in $C^{g+1}(K_i^{alg})$. If U denotes the subset of C^{g+1} consisting of $(x_1, \ldots, x_{g+1})$ such that $x_i \neq x_j$ for every i and j, then U is a Zariski dense Zariski open subset of C^{g+1}, because its complement is a closed subvariety of lower dimension. The intersection of a Zariski dense set with a Zariski dense Zariski open is still Zariski dense. So $\phi_i(K_i)^{g+1} \cap U$ is still Zariski dense in C^{g+1}. Let $\psi_i(x_1, \ldots, x_{g+1})$ be the following quantifier-free $\mathcal{L}_i$-formula over K:

$$\bigwedge_{j=1}^{g+1} \phi_i(x_j) \wedge \bigwedge_{j \neq k} x_j \neq x_k.$$

Then $\psi_i(K_i) = \phi_i(K_i)^{g+1} \cap U$ is Zariski dense in $C^{g+1}(K_i^{alg})$, so $\psi_i(-)$ is C^{g+1}-dense. By Axiom A2, it follows that some tuple $(p_1, \ldots, p_{g+1})$ satisfies

$$\bigwedge_{i=1}^{n} \psi_i(x_1, \ldots, x_{g+1}) \equiv \left(\bigwedge_{i=1}^{n} \bigwedge_{j=1}^{g+1} \phi_i(x_j) \right) \wedge \left(\bigwedge_{j \neq k} x_j \neq x_k \right).$$

Then $(p_1, \ldots, p_{g+1})$ has the desired properties. □

Now let D be the divisor $\sum_j p_j$ on the curve $\overline{C}$. By Riemann-Roch, $l(D) \geq \deg D + 1 - g = 2$. The space of global sections of $\mathcal{O}(D)$ is a K-definable vector space of dimension at least 2. Now K is either algebraically closed or has characteristic zero, so K is perfect. Therefore, by Hilbert's Theorem 90 we know that this vector space has a K-definable basis (see Remark 11.1.13). Thinking of the sections of $\mathcal{O}(D)$ as functions with poles no worse than D, we can find a non-constant meromorphic function g, with $(g) - D \geq 0$. Then the divisor of poles of g is a subset of D, so every pole of g has multiplicity 1 and is in $\bigcap_{i=1}^n \psi_i(K)$. Take $f = 1/g$. □

Proof (of Theorem 11.3.1). If $K \models T$, then K satisfies Axioms A1 and A2. Axiom A1 implies that K is algebraically closed or real closed or p-adically closed (Remark 11.2.4). As K is existentially closed, it is also reasonably clear that all the named valuations must be non-trivial. Consequently $K \restriction \mathcal{L}_1 \models T_1$ and $v_2, \ldots, v_n$ are non-trivial. Lastly, suppose T_i and T_j induce the same topology on K for some i. For notational simplicity assume $i = 1$ and $j = 2$. As the topologies are Hausdorff, we can find non-empty U_1 and U_2 with U_1 a T_1-open, U_2 a T_2-open, and $U_1 \cap U_2 = \emptyset$. Since the topologies from T_1 and T_2 have a basis of open sets consisting of quantifier-free definable sets, we can shrink U_1 and U_2 a little, and assume U_1 is quantifier-free definable in $\mathcal{L}_1$ and U_2 is quantifier-free definable in $\mathcal{L}_2$. Now U_1 and U_2 are both Zariski dense in the affine line, so the formulas defining U_1 and U_2 are $\mathbb{A}^1$-dense. Hence, by Axiom A2, U_1 must intersect U_2, a contradiction.

The other direction of the theorem is harder. We proceed by induction on n, the number of orderings and valuations. The base case where $n = 1$ is easy/trivial, so suppose $n > 1$. Suppose K satisfies the assumptions of the Theorem. By Fact 11.3.2, we know that the n different topologies on K^1 are independent. The first bullet point ensures that K satisfies axiom A1. By Theorem 11.2.18, it suffices to prove axiom A2(1). By Remark 11.2.7, we merely need to prove the following:

> Let C be an absolutely irreducible smooth affine curve defined over K. Let $\phi_1(x)$ be a quantifier-free $\mathcal{L}_1$-formula with parameters from K such that $\phi_1(K)$ is a non-empty open subset of C. For $2 \leq i \leq n$, let $\phi_i(x)$ be a quantifier-free $\mathcal{L}_i$-formula with parameters from K such that $\phi_i(x)$ defines a non-empty open subset of $C(K^{alg})$ with respect to any/every extension of the ith valuation v_i from K to K^{alg}. THEN $\bigcap_{i=1}^n \phi_i(K)$ is non-empty.

Here we are using the facts that $K \restriction \mathcal{L}_1$ is already a model of T_1, and that for $i > 1$, the field K^{alg} with any extension of v_i will be a model of $T_i = ACVF$.

For $1 < i \leq n$, choose some extension v_i' of the valuation v_i to K^{alg}.

Claim 11.3.5. *K is dense in K^{alg} with respect to the v_i'-adic topology on K^{alg}.*

Proof. The claim is trivial if all the T_i are ACVF, in which case $K = K^{alg}$. So we may assume characteristic zero. It suffices to show that K is dense in every finite Galois extension of K.[2] Let L/K be a finite Galois extension. We can write L as $K(\zeta)$ for some singleton ζ. Let $P(X) \in K[X]$ be the minimal polynomial of ζ over K. The function $x \mapsto P(x)$ from K to K is finite-to-one, so it has infinite image. As K is a model of ACVF, pCF, or RCF, we see by Fact 11.1.1 that the image $P(K)$ of this map contains an open subset of K with respect to the T_1-topology. Because the v_i-adic topology on K is independent from the T_1-topology on K, we can find elements of $P(K)$ of arbitrarily high v-valuation. By the cofinality of the value groups, for every $\gamma \in v_i'(K^{alg})$, we can find an $x \in K$ with $v_i(P(x)) > \gamma$. Let

[2]Note that the value group $v_i'(K)$ is cofinal in $v_i'(K^{alg})$, so e.g. the v_i-adic topology on K is the restriction of the v_i'-adic topology on K^{alg} to K. Various pathologies are thus avoided.

$\zeta_1, \ldots, \zeta_m \in L$ be the conjugates of ζ over K, counted with multiplicities. Then we have just seen that for any $\gamma \in v_i'(K^{alg})$, we can find an $x \in K$ with

$$\gamma < v_i'(P(x)) = \sum_{i=m}^{n} v_i'(x - \zeta_i).$$

This implies that at least one of the ζ_i's is in the topological closure of K with respect to v_i'. Consequently, the v_i'-topological closure of K in L must contain $K[\zeta_i]$ for some i. But $K[\zeta_i] = L$, so K is v_i'-dense in L. □

Now suppose we are given an absolutely irreducible smooth affine curve C defined over K, and we have $\mathcal{L}_i$-formulas $\phi_i(x)$ with parameters from K, such that $\phi_1(K)$ is a non-empty open subset of $C(K)$, and for $1 < i \leq n$, $\phi_i(K^{alg})$ is a non-empty v_i'-open subset of $C(K^{alg})$. (Here we are interpreting $\phi_i(K^{alg})$ using v_i'.) By the inductive hypothesis, $K \upharpoonright \bigcup_{i<n} \mathcal{L}_i$ is an existentially closed model of $\bigcup_{i<n}(T_i)_\forall$. Applying Lemma 11.3.3 to it, we can find a K-definable rational function $f : C \to \mathbb{P}^1$, whose divisor of zeros has no multiplicities and consists entirely of points in $\bigcap_{i<n} \phi_i(K)$ (and no points at infinity and no points in $C(K^{alg}) \setminus C(K)$). Write this divisor as $\sum_{j=1}^{m}(P_j)$, where the P_j are m distinct points in $\bigcap_{i<n} \phi_i(K)$. Note that m is the degree of f.

Claim 11.3.6. *There is a T_1-open neighborhood $U \subseteq K$ of zero such that for every $y \in U_1$, the divisor $f^{-1}(y)$ consists of j distinct points in $\phi_1(K)$. In particular, it contains no points in $C(K^{alg}) \setminus C(K)$ and no points in $\overline{C} \setminus C$.*

Proof. Because the P_j are distinct, they have multiplicity one, so f does not have a critical point at any of the P_j's. Consequently, by the implicit function theorem there is a T_1-open neighborhood $W_j \subseteq C(K)$ of P_j such that f induces a T_1-homeomorphism from W_j to an open neighborhood of 0. By shrinking W_j if necessary, we may assume that $W_j \subseteq \phi_1(K)$, and that $W_j \cap W_{j'} = \emptyset$ for $j \neq j'$. Now let $U = \bigcap_{j=1}^{m} f(W_j)$. This is an open neighborhood of 0 in the affine line K^1. And if $y \in U$, then $f^{-1}(y)$ contains at least one point in each W_j. Since the W_j are distinct, these points are distinct. Since f is a degree-m map, this exhausts the divisor $f^{-1}(y)$. □

Claim 11.3.7. *For $1 < i < n$, there is a $\gamma_i \in v_i(K)$ such that if $y \in K^{alg}$ and $v_i'(y) > \gamma_i$, then $f^{-1}(y)$ are all in $\phi_i(K^{alg})$.*

Proof. Use the same argument as Claim 11.3.6. □

By Claim 11.3.5, K is dense in K^{alg} with respect to the v_n'-adic topology. Also, by assumption, $\phi_n(x)$ interpreted in (K^{alg}, v_n') yields a non-empty v_n'-open subset $W \subseteq C(K^{alg})$. Since f is finite-to-one, the image $f(W)$ is an infinite subset of $\mathbb{P}^1(K^{alg})$, hence it has non-empty v_n'-interior. Let V be a v_n'-open subset of $\mathbb{P}^1(K^{alg})$ contained in $f(W)$. Now, as K is v_n'-adically dense in K^{alg}, V must intersect K. In particular, $V \cap K$ is a *non-empty* v_n-adic open subset of K. By independence of the topologies, we can find a y in $\mathbb{A}^1(K)$ such that

- y is in U, the T_1-open neighborhood of 0 from Claim 11.3.6.
- $v_i(y) > \gamma_i$, for $1 < i < n$, where the γ_i are from Claim 11.3.7
- y is in $V \cap K$.

Having chosen such a y, we know by Claim 11.3.6 that $f^{-1}(y)$ consists of j distinct points in $\phi_1(K)$. In particular, each point in $f^{-1}(y)$ is a point of $C(K)$. And by Claim 11.3.7, each of these points also belongs to $\phi_i(K^{alg})$, hence satisfies $\phi_i(-)$, for $i < n$. Finally, because y is in $V \cap K$, y is in the image of $\phi_n(K^{alg})$ under f. So there is some $x \in \phi_n(K^{alg})$ mapping to y. But we said that every point in $C(K^{alg})$ mapping to y is already in $C(K)$ and even in $\bigcap_{i<n} \phi_i(K)$. Thus

$$x \in \phi_n(K^{alg}) \cap \bigcap_{i<n} \phi_i(K) = \bigcap_{i=1}^{n} \phi_i(K).$$

In particular some point in $C(K)$ satisfies $\bigwedge_{i=1}^{n} \phi_i(x)$, and the theorem is proven. □

11.4 Keisler Measures

To establish NTP_2 and analyze forking and dividing in T, we need the following tool.

Theorem 11.4.1. *Let T be one of the model companions from §11.2. For each $K \models T_\forall$ that is a perfect field, each formula $\phi(x)$ and each tuple a from K, we can assign a number $P(\phi(a), K) \in [0, 1]$ such that the following conditions hold:*

- *If K is held fixed, the function $P(-, K)$ is a Keisler measure on the space of completions of the quantifier-free type of K. Thus*

$$P(\phi(a), K) + P(\psi(b), K) = P(\phi(a) \wedge \psi(b), K) + P(\phi(a) \vee \psi(b), K)$$

$$P(\neg\phi(a), K) = 1 - P(\phi(a), K)$$

  for sentences $\phi(a)$ and $\psi(b)$ over K. And if $\phi(a)$ holds in every model of T extending K, then $P(\phi(a), K) = 1$. For example, if $\phi(x)$ is quantifier-free, then $P(\phi(a), K)$ is 1 or 0 according to whether or not $K \models \phi(a)$. And if K is satisfies axiom A1 of §11.2.1, then $P(\phi(a), K) \in \{0, 1\}$ for every $\phi(a)$, by Corollary 11.2.11.

- *Isomorphism invariance: if K, L are two perfect fields satisfying $T_\forall$, and $f : K \to L$ is an isomorphism of structures, then $P(\phi(a), K) = P(\phi(f(a)), L)$ for every K-sentence $\phi(a)$.*

- *Extension invariance: if $K_0 \subseteq K$ are perfect fields satisfying $T_\forall$, and K_0 is relatively algebraically closed in K, and $\phi(a)$ is a formula with parameters from K_0, then $P(\phi(a), K_0) = P(\phi(a), K)$.*

- *Density: if $K \models T_\forall$ is a perfect field and $\phi(a)$ is a K-formula, and if $M \models \phi(a)$ for at least one $M \models T$ extending K, then $P(\phi(a), K) > 0$. In other words, the associated Keisler measure is spread out throughout the entire Stone space of completions of* $\mathrm{qftp}(K)$.

11.4.1 The Algebraically Closed Case

We first prove Theorem 11.4.1 in the case where every T_i is a model of ACVF, i.e., the case of existentially closed fields with n valuations. Define $P(\phi(a), K)$ as follows. Fix some algebraic closure K^{alg} of K. For each $1 \leq i \leq n$, let v_i' be an extension to K^{alg} of the i^{th} valuation v_i on K. Choose automorphisms $\sigma_1, \ldots, \sigma_n \in \mathrm{Gal}(K^{alg}/K)$ randomly with respect to Haar measure on $\mathrm{Gal}(K^{alg}/K)$. Then

$$K_{\sigma_1,\ldots,\sigma_n} := (K^{alg}, v_1' \circ \sigma_1, v_2' \circ \sigma_2, \ldots, v_n' \circ \sigma_n)$$

is a model of $T_\forall$ satisfying axiom A1 of §11.2.1. In particular, whether or not $\phi(a)$ holds in a model of T extending $K_{\sigma_1,\ldots,\sigma_n}$ does not depend on the choice of the model, by Corollary 11.2.11. Define $P(\phi(a), K)$ to be the probability that $\phi(a)$ holds in any/every model of T extending $K_{\sigma_1,\ldots,\sigma_n}$. This probability exists, i.e., the relevant event is measurable, because whether or not $\phi(a)$ holds is determined by the behavior of the valuations on some finite Galois extension L/K, by virtue of Theorem 11.2.15.

Note that the choice of the v_i' does not matter. If v is a valuation on K and w_1 and w_2 are two extensions of v to K^{alg}, then there is a τ in $\mathrm{Gal}(K^{alg}/K)$ such that $w_1 = w_2 \circ \tau$. I believe this is well-known, and at any rate it is an easy consequence of quantifier-elimination in ACVF. From this, it follows that if σ is a randomly chosen element of $\mathrm{Gal}(K^{alg}/K)$, then $w_1 \circ \sigma$ and $w_2 \circ \sigma$ have the same distribution. Consequently the choice of the valuations v_i' does not effect the resulting value of $P(\phi(a), K)$.

So we have a well-defined number $P(\phi(a), K)$, and it is defined canonically. The first two bullet points of Theorem 11.4.1 are therefore clear. The density part can be seen as follows: suppose $M \models \phi(a)$ for some $M \models T$ extending K. Let K^{alg} be the algebraic closure of K in M. For the v_i''s, take the restrictions of the valuations on M to K^{alg}. By Theorem 11.2.15, there is a field $K \leq L \leq K^{alg}$ with L/K a finite Galois extension, such that $\phi(a)$ is implied by T and the diagram of L. Specifically, write $\phi(a)$ as $\exists y Q(y; a) = 0 \wedge \psi(y; a)$, and let L be the splitting field of the polynomial $Q(X; a) \in K[X]$. Now with probability $1/[L : K]^n$, every σ_i will restrict to the identity on L. Consequently, $K_{\sigma_1,\ldots,\sigma_n}$ will be a model of $T_\forall$ extending L, so in any model M of T extending $K_{\sigma_1,\ldots,\sigma_n}$, $\phi(a)$ will hold. So $\phi(a)$ holds with probability at least $1/[L : K]^n$, and consequently $P(\phi(a), K) \geq 1/[L : K]^n$.

It remains to verify the extension invariance part of Theorem 11.4.1. Let $K_0 \leq K$ be an inclusion of perfect fields, with K_0 relatively algebraically closed in K. Let $\phi(a)$ be a formula with parameters a from K_0. As in the previous paragraph, write $\phi(a)$ as $\exists y Q(y; a) = 0 \wedge \psi(y; a)$ and let L_0 be the splitting field of $Q(y; a)$ over K_0. At present L_0 is nothing but a pure field. Write $L_0 = K_0(\beta)$ for some singleton $\beta \in L_0$, and let $Q(X)$ be

the irreducible polynomial of β over K_0. Let $L = L_0K = K(\beta)$; this is a Galois extension of K. There are only finitely many ways of factoring $Q(X)$ in K^{alg}, so in each way of factoring $Q(X)$, the coefficients come from K_0^{alg}. In particular, if $Q(X)$ can be factored over K, the coefficients would belong to $K_0^{alg} \cap K = K_0$. So $Q(X)$ is still irreducible over K. Consequently $[L : K] = \deg Q(X) = [L_0 : K_0]$. Now there is a natural restriction map $\mathrm{Gal}(L/K) \to \mathrm{Gal}(L_0/K_0)$. It is injective because an element of $\mathrm{Gal}(K(\beta)/K)$ is determined by what it does to β. Since $\mathrm{Gal}(L/K)$ has the same size as $\mathrm{Gal}(L_0/K_0)$, the restriction map must be an isomorphism. Consequently, if τ is chosen from $\mathrm{Gal}(L/K)$ randomly, its restriction to L_0 is a random element of $\mathrm{Gal}(L_0/K_0)$. Consequently, if σ is a random element of $\mathrm{Gal}(K^{alg}/K)$ and σ_0 is a random element of $\mathrm{Gal}(K_0^{alg}/K_0)$, then $\sigma \restriction L_0$ and $\sigma_0 \restriction L_0$ have the same distribution, namely, the uniform distribution on $\mathrm{Gal}(L_0/K_0)$. From this, it follows easily that $P(\phi(a), K) = P(\phi(a), K_0)$.

This completes the proof of Theorem 11.4.1 when every T_i is ACVF. The other cases are more complicated, though as a consolation all fields are characteristic zero, hence perfect.

A first attempt at defining $P(\phi(a), K)$ is as follows: fix some algebraic closure K^{alg} of K. For each i such that T_i is RCF, let K_i be a real closure of $(K, <_i)$ inside K^{alg}. For each i such that T_i is pCF, let K_i be a p-adic closure of (K, v_i) inside K^{alg}. For each i such that T_i is ACVF, let K_i be K^{alg} with some valuation extending v_i. In each case, there is a choice, but any two choices are related by an element of $\mathrm{Gal}(K^{alg}/K)$. Now choose $\sigma_1, \ldots, \sigma_n \in \mathrm{Gal}(K^{alg}/K)$ randomly. For each i, consider $\sigma_i(K_i)$, which is (usually) a model of T_i extending K. Let K' be the field

$$K' = \bigcap_{i=1}^{n} \sigma_i(K_i).$$

There is an obvious way to give K' the structure of a $T_\forall$-model. If we knew that K' satisfies condition A1 of §11.2.1 with high probability, we could define $P(\phi(a), K)$ to be the probability that $\phi(a)$ holds in any/every model of T extending K'. Unfortunately, K' usually satisfies condition A1 with probability *zero*. Instead, we will proceed by repeating the above procedure with K' in place of K, getting a third field K''. Iterating this, we get an increasing sequence $K \subseteq K' \subseteq K'' \subseteq \cdots \subseteq K^{(n)} \subseteq \cdots$ of $T_\forall$-structures on subfields of K^{alg}. The union $K^\infty = \bigcup_{n=1}^\infty K^{(n)}$ does actually turn out to satisfy axiom A1 with probability 1, and we let $P(\phi(a), K)$ be the probability that $\phi(a)$ holds in any/every model of T extending K^∞.

The rest of this section will make this construction more precise, and verify that it satisfies the requirements of Theorem 11.4.1.

11.4.2 The General Case

All fields will be perfect, unless stated otherwise. All models of $T_\forall$ and $(T_i)_\forall$ will be (perfect) fields, unless stated otherwise. Galois extensions need not be finite Galois extensions.

We start off with some easy but confusing facts that will be needed later.

Lemma 11.4.2. *Let L/K be a Galois extension of fields, and suppose K has the structure of a $(T_i)_\forall$ model (but L does not). The following are equivalent*

(a) *For every F, if F is a model of $(T_i)_\forall$ extending K, and F is a subfield of L, then $F = K$.*

(b) *There is a model $M \models T_i$ extending K, such that $M \cap L = K$.*

(c) *For every model $M \models T_i$ extending K, $M \cap L = K$.*

Note that it makes sense to talk about whether $M \cap L = K$, because L/K is Galois.

Proof. The equivalence of (b) and (c) follows from quantifier elimination in T_i. Indeed, the statement that $M \cap L = K$ is equivalent to the statement that for each $x \in L \setminus K$, the irreducible polynomial of x over K has no zeros in M. This is a conjunction of first order statements about K, so it holds in one choice of M if and only if it holds in another choice of M.

Suppose (a) holds. Let M be a model of T_i extending K. Taking $F = M \cap L$, (a) implies that $M \cap L = K$. So (a) implies (c).

Conversely, suppose (a) does not hold. Let F witness a contradiction to (a), so $K \subsetneq F \subseteq L$, and F is a model of $(T_i)_\forall$ extending K. Let M be a model of T_i extending F and hence K. Then $M \cap L$ contains F, contradicting (c). □

Definition 11.4.3. *Say that K is* locally T_i-closed in L *if it satisfies the equivalent conditions of the previous lemma.*

Definition 11.4.4. *Let L/K be a Galois extension of fields, and suppose K has the structure of a $(T_i)_\forall$-model (but L does not). Let $\mathfrak{C}_i(L/K)$ denote the set of models of $(T_i)_\forall$ which extend K, are subfields of L, and are locally T_i-closed in L.*

The subscript on $\mathfrak{C}_i$ is present so that $\mathfrak{C}_i(L/K)$ will be unambiguous when K is a model of $T_\forall$, in addition to being a model of $(T_i)_\forall$.

There is a natural action of $\mathrm{Gal}(L/K)$ on $\mathfrak{C}_i(L/K)$.

Lemma 11.4.5. *Suppose L/K is a Galois extension of fields, and $K \models (T_i)_\forall$.*

(a) *The action of $\mathrm{Gal}(L/K)$ on $\mathfrak{C}_i(L/K)$ has exactly one orbit.*

(b) *Suppose K' is a model of $(T_i)_\forall$ extending K, and L' is a field extension of L and K', with L' Galois over K'. If $F \in \mathfrak{C}_i(L'/K')$, then $F \cap L \in \mathfrak{C}_i(L/L \cap K')$.*

Proof. **(a)** Note that $\mathfrak{C}_i(L/K)$ is non-empty by a Zorn's lemma argument and condition (a) of Lemma 11.4.2. Now suppose F and F' are two elements of $\mathfrak{C}_i(L/K)$. By quantifier elimination in T_i, we can amalgamate F and F' over K. Thus, we can find a model $M \models T$ extending F, and an embedding $\iota : F' \to M$ which is the identity on K. Choosing some way of amalgamating M and L as fields, we get that $\iota(F') \subseteq L \supseteq F$,

because of L/K being Galois. The compositum $\iota(F')F$ is a subfield of L with a $(T_i)_\forall$-structure extending that on F and $\iota(F')$, so by local T_i-closedness of $\iota(F')$ and F in L, $\iota(F') = \iota(F')F = F$. It follows that F' and F are isomorphic over K. This isomorphism must extend to an automorphism of L, because L/K is Galois. So some automorphism on L/K maps F' to F (as $(T_i)_\forall$-structures).

(b) Let M be a model of T_i extending F. Choose some way of amalgamating M with L'. Then $M \cap L' = F$ by (c) of Lemma 11.4.2. Therefore, $M \cap L = M \cap L' \cap L = F \cap L$. So by (b) of Lemma 11.4.2, $F \cap L$ is locally T_i-closed in L. Therefore it is in $\mathfrak{C}_i(L/L \cap K')$. □

Now we turn our attention from T_i to T.

Definition 11.4.6. *Let $K \models T_\forall$ and let L be a pure field that is a Galois extension of K. Let $\mathfrak{S}(L/K)$ be the set of all $K' \models T_\forall$ extending K, with K' a subfield of L. In other words, an element of $\mathfrak{S}(L/K)$ is a subfield F of L, endowed with a $T_\forall$-structure, such that $F \supseteq K$ and the structure on F extends the structure on K.*

There is a natural partial order on $\mathfrak{S}(L/K)$ coming from inclusion of substructures. There is also a natural action of $\mathrm{Gal}(L/K)$ on $\mathfrak{S}(L/K)$. One should think of $\mathfrak{S}(L/K)$ as the set of states in a Markov chain, specifically the random process described at the end of the previous section.

Definition 11.4.7. *Suppose $K \models T_\forall$ and L/K is a Galois extension of K. For $1 \leq i \leq n$, choose some $L_i \in \mathfrak{C}_i(L/K)$. Choose $\sigma_1, \ldots, \sigma_n \in \mathrm{Gal}(L/K)$ independently and randomly, using Haar measure on $\mathrm{Gal}(L/K)$. Let F be $\bigcap_{i=1}^n \sigma_i(L_i)$, with the obvious choice of a $T_\forall$ structure. So F is a random variable with values in $\mathfrak{S}(L/K)$. Let $\mu^1_{L/K}$ be the probability distribution on $\mathfrak{S}(L/K)$ obtained in this way. The choice of the L_i's is irrelevant, by Lemma 11.4.5(a).*

The superscript 1 is to indicate that this is the first step of the Markov chain.

Lemma 11.4.8. *Suppose L/K is finite. Then every event (subset of $\mathfrak{S}(L/K)$) which has positive probability with respect to $\mu^1_{L/K}$ has probability at least $1/m^n$, where $m = [L : K]$.*

Proof. The only randomness comes from the σ_i's. Each element of $\mathrm{Gal}(L/K)$ has an equal probability under Haar measure, and this probability is $1/m$. Since the σ_i's are chosen independently, each choice of the σ_i's has probability $1/m^n$ of occurring. □

Lemma 11.4.9. *Suppose L/K is finite, and F is a maximal element of $\mathfrak{S}(L/K)$. Then $\mu^1_{L/K}(\{F\}) > 0$.*

Proof. For each i, let M_i be a model of T_i extending $F \restriction \mathcal{L}_i$, and choose some way of amalgamating M_i and L as fields over F. Let $F_i = L \cap M_i$. Of course $F_i \supseteq F$. By Lemma 11.4.2(b), $F_i \in \mathfrak{C}_i(L/K)$. Let $F' = \bigcap_{i=1}^n F_i$. Then $F' \in \mathfrak{S}(L/K)$ and F' extends

F, so $F = F'$ by maximality of F. Now if we choose $\sigma_1, \ldots, \sigma_n \in \mathrm{Gal}(L/K)$ randomly, then $\bigcap_{i=1}^n \sigma_i(F_i)$ is distributed according to $\mu^1_{L/K}$. Since L/K is finite, there is a positive probability that $\sigma_i = 1$ for every i, in which case $\bigcap_{i=1}^n \sigma_i(F_i) = F' = F$. □

Lemma 11.4.10. *Let L/K be a Galois extension, and K be a model of $T_\forall$. Let K' be a model of $T_\forall$ extending K. Let L' be a field extending L and K', with L' Galois over K'. If $F \in \mathfrak{S}(LK'/K')$ is distributed randomly according to $\mu^1_{L'/K'}$, then $F \cap L$ is distributed randomly according to $\mu^1_{L/L\cap K'}$.*

Proof. For $1 \leq i \leq n$, choose some $F_i \in \mathfrak{C}_i(L'/K')$. By Lemma 11.4.5(b), $F_i \cap L$ is in $\mathfrak{C}_i(L/L \cap K')$.

Claim 11.4.11. *If we choose σ from $\mathrm{Gal}(L'/K')$ randomly using Haar measure, then $\sigma \restriction L$ is also randomly distributed in $\mathrm{Gal}(L/L \cap K')$ with respect to Haar measure.*

Proof. It probably suffices to show that if $x \in L$, then the probability that $\sigma \restriction L$ fixes x is the same as the probability that a random element of $\mathrm{Gal}(L/L \cap K')$ fixes x. This is equivalent to showing that x has the same number of conjugates under the action of $\mathrm{Gal}(L/L \cap K')$ as under the action of $\mathrm{Gal}(L'/K')$. If x and y are conjugate over K', then clearly they are conjugate over $L \cap K'$. Conversely, if they are conjugate over $K' \cap L$ but not over K', then let $S \subseteq L$ be the orbit of x under $\mathrm{Gal}(L'/K')$. The code for the finite set S is definable over K'. It is also definable over L, as $S \subseteq L$. Thus the code for S is in $L \cap K'$. As x and y are conjugate over $L \cap K'$, we conclude that $y \in S \iff x \in S$, a contradiction. □

From the Claim, we conclude that if the σ_i are distributed randomly from $\mathrm{Gal}(L'/K')$, then $\sigma_i \restriction L$ are distributed randomly from $\mathrm{Gal}(L/L \cap K')$. Taking $F = \bigcap_{i=1}^n \sigma_i(F_i)$, we get F distributed according to $\mu^1_{L'/K'}$. But

$$F \cap L = \bigcap_{i=1}^n (\sigma_i \restriction L)(F_i \cap L)$$

is then distributed according to $\mu^1_{L/L\cap K'}$, because $F_i \cap L \in \mathfrak{C}_i(L/(K' \cap L))$ and $\sigma_i \restriction L$ is distributed according to Haar measure on $\mathrm{Gal}(L/L \cap K')$. □

Definition 11.4.12. *Let L/K be a Galois extension, and K be a model of $T_\forall$. Define a series of distributions $\{\mu^i_{L/K}\}_{i<\omega}$ on $\mathfrak{S}(L/K)$ as follows:*

- *$\mu^0_{L/K}$ assigns probability 1 to $\{K\} \subseteq \mathfrak{S}(L/K)$.*
- *$\mu^1_{L/K}$ is as above.*
- *For $i > 0$, if we choose $F \in \mathfrak{S}(L/K)$ randomly according to $\mu^i_{L/K}$, and then choose $F' \in \mathfrak{S}(L/F) \subseteq \mathfrak{S}(L/K)$ randomly according to $\mu^1_{L/F}$, then F' is distributed according to $\mu^{i+1}_{L/K}$.*

In other words, we are running some kind of Markov chain whose states are the elements of $\mathfrak{S}(L/K)$. The transition probabilities out of the state F are given by $\mu^1_{L/F}$, and $\mu^n_{L/K}$ is the distribution of the Markov chain after n steps.

Lemma 11.4.13. *Let L/K be a finite Galois extension, and K be a model of $T_\forall$. Then $\lim_{i\to\infty} \mu^i_{L/K}$ exists, and the corresponding distribution on $\mathfrak{S}(L/K)$ is concentrated on the maximal elements of $\mathfrak{S}(L/K)$.*

Proof. The fact that the limit distribution exists is a general fact about Markov chains with finitely many states such that the graph of possible transitions has no cycles other than self-loops.

It remains to check that in the limit, we land in a maximal element of $\mathfrak{S}(L/K)$ with probability one. Let $m = [L : K]$. If $F \in \mathfrak{S}(L/K)$ is not maximal, then the probability of moving from F to some bigger element is positive by Lemma 11.4.9, and at least $1/m^n$, by Lemma 11.4.8. The probability of getting stuck at F is therefore bounded above by $\lim_{k\to\infty}(1 - 1/m^n)^k = 0$. As there are finitely many non-maximal F, we conclude that the probability of getting stuck at any of them is zero. □

We let $\mu^\infty_{L/K}$ denote the limit distribution on $\mathfrak{S}(L/K)$.

Lemma 11.4.14. *Let L/K be a finite Galois extension, and K be a model of $T_\forall$. Then every maximal element of $\mathfrak{S}(L/K)$ has a positive probability with respect to $\mu^\infty_{L/K}$.*

Proof. This follows immediately from Lemma 11.4.9, and the fact that once the Markov chain reaches a maximal element of $\mathfrak{S}(L/K)$, it must remain there. □

Lemma 11.4.15. *Let L/K be a Galois extension, with K a model of $T_\forall$. Let K' be a model of $T_\forall$ extending K. Let L' be a field extending K' and L, Galois over K'. If F is a random element of $\mathfrak{S}(L'/K')$ distributed according to $\mu^i_{LK'/K'}$, then $F \cap L$ is distributed according to $\mu^i_{L/L\cap K'}$.*

Proof. We proceed by induction on i. For $i = 0$, F is guaranteed to be K', and $F \cap L$ is guaranteed to be $K' \cap L$, which agrees with $\mu^0_{L/L\cap K'}$.

For the inductive step, suppose we know the statement of the lemma for μ^i, and prove it for μ^{i+1}. If we let $F \in \mathfrak{S}(L'/K')$ be chosen according to $\mu^i_{L'/K'}$, and we then choose $F' \in \mathfrak{S}(L'/F) \subseteq \mathfrak{S}(L'/K')$ according to $\mu^1_{L'/F}$, then F' is randomly distributed according to $\mu^{i+1}_{L'/K'}$, by definition of μ^{i+1}. Also, $F\cap L$ is distributed according to $\mu^i_{L/L\cap K'}$, by the inductive hypothesis. By Lemma 11.4.10 we know that $F' \cap L$ is distributed according to $\mu^1_{L/L\cap F}$. In particular, the distribution of $F' \cap L$ only depends on $F \cap L$. So if we want to sample $F' \cap L$, we can simply choose $F \cap L$ using $\mu^i_{L/L\cap K'}$, and can then choose $F' \cap L$ using $\mu^1_{L/F\cap L}$. This is the recipe for sampling the distribution $\mu^{i+1}_{L/K'\cap L}$. So $F' \cap L$ is indeed distributed according to $\mu^{i+1}_{L/K'\cap L}$. □

Corollary 11.4.16. *When L/K and L'/K' are finite Galois extension, Lemma 11.4.15 holds for $i = \infty$.*

Definition 11.4.17. *Let $K \models T_\forall$ be a perfect field, $\phi(a)$ be a formula in the language of T with parameters a from K. Say that a finite Galois extension L/K* determines the truth of *$\phi(a)$ if the following holds: whenever M and M' are two models of T extending K, if $M \cap L$ is isomorphic as a model of $T_\forall$ to $M' \cap L$, then $[M \models \phi(a)] \iff [M' \models \phi(a)]$. (Note that the isomorphism class of $M \cap L$ does not depend on how we choose to form the compositum ML.)*

For every formula $\phi(a)$, there *is* some finite Galois extension L/K which determines the truth of $\phi(a)$. Namely, use Theorem 11.2.15 to write $\phi(a)$ in the form $\exists y Q(y; a) = 0 \wedge \psi(y; a)$, and take L to be the splitting field over K of $Q(X; a) \in K[X]$.

Lemma 11.4.18. *Let K be a model of $T_\forall$, M be a model of T extending K, and let L/K be a Galois extension of K. Assume M and L are embedded over K into some bigger field. Then $M \cap L$ is a maximal element of $\mathfrak{S}(L/K)$.*

Proof. Suppose not. Let F be an element of $\mathfrak{S}(L/K)$, strictly bigger than $M \cap L$, and finitely generated over $M \cap L$. Let x be a generator of F over $M \cap L$. If S denotes the set of algebraic conjugates of x over M, then the code for the finite set S is in M, and also in L because $S \subseteq L$. So the code for S is in $M \cap L$, implying that S is also the set of algebraic conjugates of x over $M \cap L$. Since we are assuming that all fields are perfect, this implies that the degree of x over M is the same as the degree of x over $M \cap L$. In particular, the irreducible polynomial $Q(X)$ of x over $M \cap L$ remains irreducible over M. For $1 \leq i \leq n$, let M_i be a model of T_i extending $M \restriction \mathcal{L}_i$. Let N_i be a model of T_i extending $F \restriction \mathcal{L}_i$. The polynomial $Q(X)$ has a zero in F, namely x. Hence it has a zero in $N_i \supseteq F$. As M_i and N_i are two models of T_i extending $M \cap L$ and $Q(X)$ is defined over $M \cap L$, it follows from quantifier-elimination in T_i that $Q(X)$ also has a zero in M_i.

Now we have a polynomial $Q(X)$ of degree > 1, irreducible over M, such that $Q(X)$ has a root in M_i for every i. This contradicts condition A1' of §11.2.1. □

Definition 11.4.19. *Let L/K be a Galois extension, with K a model of $T_\forall$. Let $\mathfrak{F}(L/K)$ be the set of maximal elements of $\mathfrak{S}(L/K)$.*

By Zorn's lemma, it is clear that every element of $\mathfrak{S}(L/K)$ is bounded above by an element of $\mathfrak{F}(L/K)$, even if L/K is infinite. When L/K is a finite extension, $\mu^\infty_{L/K}$ induces a probability distribution on $\mathfrak{F}(L/K)$.

Remark 11.4.20. *$\mathfrak{F}(L/K)$ is exactly the set of F of the form $L \cap M$, where M is a model of T extending K. One inclusion is Lemma 11.4.18. The other inclusion is obvious: if F is a maximal element of $\mathfrak{S}(L/K)$, then letting M be a model of T extending F, and combining M and L into a bigger field in any way we like, we have $F \subseteq M \cap L \in \mathfrak{S}(L/K)$, so maximality of F forces $M \cap L = F$.*

Suppose that L/K determines the truth of $\phi(a)$. Then by Remark 11.4.20, there must be a uniquely determined map $f_{\phi(a),L}$ from $\mathfrak{F}(L/K)$ to $\{\bot, \top\}$ such that for every $M \models T$ extending K, and every way of forming the compositum ML, the truth of $M \models \phi(a)$ is given by $f_{\phi(a),L}(M \cap L)$.

Another corollary of Remark 11.4.20 is that if $K \leq L \leq L'$, with L' and L Galois extensions of $K \models T_\forall$, and if $F \in \mathfrak{F}(L'/K)$, then $F = M \cap L'$ for some model M, and hence $F \cap L = M \cap L' \cap L = M \cap L$ is in $\mathfrak{F}(L/K)$.

Finally, we define $P(\phi(a), K)$ to be $\mu^\infty_{L/K}(\{F : f_{\phi(a),L}(F) = \top\})$.

Lemma 11.4.21. *The choice of L does not matter.*

Proof. If L and L' are two finite Galois extensions of K which determine the truth of $\phi(a)$, then so does their compositum LL'. So we may assume $L \subseteq L'$. Let $r : \mathfrak{F}(L'/K) \to \mathfrak{F}(L/K)$ be the restriction map, $F \mapsto F \cap L$.

Claim 11.4.22. $f_{\phi(a),L'} = f_{\phi(a),L} \circ r$.

Proof. For $F \in \mathfrak{F}(L'/K)$, we will show $f_{\phi(a),L'}(F) = f_{\phi(a),L}(r(F))$. Write F as $M \cap L'$, with M a model of T extending F. Then $f_{\phi(a),L'}(M \cap L') = f_{\phi(a),L'}(F)$ is the truth value of $M \models \phi(a)$. But $M \cap L = F \cap L$, so by definition of $f_{\phi(a),L}$, we also know that $f_{\phi(a),L}(M \cap L) = f_{\phi(a),L}(r(F))$ is the truth value of $M \models \phi(a)$, which is the same thing. So $f_{\phi(a),L}(r(F)) = f_{\phi(a),L'}(F)$. □

By Corollary 11.4.16 applied in the $K = K'$ case, if F is a random element of $\mathfrak{F}(L'/K)$ chosen according to $\mu^\infty_{L'/K}$, then $r(F) = F \cap L$ is distributed according to $\mu^\infty_{L/K}$. In particular, the probability of $f_{\phi(a),L'}(F)$ or equivalently of $f_{\phi(a),L}(r(F))$ is the same as the probability of $f_{\phi(a),L}(F')$, with F' chosen directly from $\mu^\infty_{L/K}$. But the former probability is $P(\phi(a), K)$ computed using L', while the latter is $P(\phi(a), K)$ computed using L. □

So $P(\phi(a), K)$ is at least a well-defined number. The "isomorphism invariance" part of Theorem 11.4.1 is clear from the definitions. We need to prove the other conditions of Theorem 11.4.1.

Lemma 11.4.23. *For any fixed K, the function $P(-, K)$ is a Keisler measure on the space of completions of the quantifier-free type of K.*

Proof. It suffices to prove the following:

- If $\phi(a)$ and $\psi(b)$ are forced to be logically equivalent by T and the diagram of K, then $P(\phi(a), K) = P(\psi(b), K)$. This is easy/trivial, because if we choose a finite Galois extension L determining the truth of both $\phi(a)$ and $\psi(b)$, we see that $f_{\phi(a),L} = f_{\psi(b),L}$ by unwinding the definitions.
- $P(\phi(a), K) = 1 - P(\neg\phi(a), K)$, which follows similarly, though it uses the fact that $\mu^\infty_{L/K}$ is concentrated on $\mathfrak{F}(L/K)$.

- If $\phi(a) \wedge \psi(b)$ contradicts $T \cup \mathrm{diag}(K)$, then $P(\phi(a) \vee \psi(b), K) = P(\phi(a), K) + P(\psi(b), K)$. Again, this is not difficult: if L is a field determining the truth of both $\phi(a)$ and $\psi(b)$, then it is clear that

$$f_{\phi(a),L} \wedge f_{\psi(b),L} = f_{\phi(a)\wedge\psi(b),L} = \bot$$

$$f_{\phi(a),L} \vee f_{\psi(b),L} = f_{\phi(a)\vee\psi(b),L}.$$

  Consequently, $\{F : f_{\phi(a)\vee\psi(b),L}(F) = \top\}$ is a disjoint union of $\{F : f_{\phi(a),L}(F) = \top\}$ and $\{F : f_{\psi(b),L}(F) = \top\}$, so we reduce to the fact that μ^∞ is a probability distribution.

- $0 \leq P(\phi(a), K) \leq 1$, which is clear from the definition.

This finishes the lemma. □

Lemma 11.4.24. *If $K \subseteq K'$ are models of $T_\forall$ and K is relatively algebraically closed in K', and $\phi(a)$ is a formula with parameters from K, then $P(\phi(a), K) = P(\phi(a), K')$.*

(This is the "extension invariance" part of Theorem 11.4.1.)

Proof. Let L be a finite Galois extension of K determining the truth of $\phi(a)$. Let L' be a finite Galois extension of K' determining the truth of $\phi(a)$; we may assume $L' \supseteq L$. (In fact, we can take $L' = LK'$.) Because K is relatively algebraically closed in K', $L \cap K' = K$. So by Corollary 11.4.16, if $F \in \mathfrak{F}(L'/K')$ is distributed according to $\mu^\infty_{L'/K'}$, then $F \cap L$ is distributed according to $\mu^\infty_{L/K}$. Using Lemma 11.4.14, this implies that $F \cap L \in \mathfrak{F}(L/K)$ for any $F \in \mathfrak{F}(L'/K')$. Let $r : \mathfrak{F}(L'/K') \to \mathfrak{F}(L/K)$ be the map $F \mapsto F \cap L$. By unwinding the definitions (as in the claim in the proof of Lemma 11.4.21), one sees that $f_{\phi(a),L'/K'} = f_{\phi(a),L/K} \circ r$. As in the proof of Lemma 11.4.21, we see that for F' choosen randomly from $\mathfrak{F}(L'/K')$ and F chosen randomly from $\mathfrak{F}(L/K)$, the distribution of F and $r(F')$ is the same, and therefore so too is the distribution of

$$f_{\phi(a),L'/K'}(F') = f_{\phi(a),L/K}(r(F'))$$

and

$$f_{\phi(a),L/K}(F).$$

This ensures that $P(\phi(a), K) = P(\phi(a), K')$. □

Lemma 11.4.25. *If $K \models T_\forall$ and $\phi(a)$ is a K-formula which holds in some model of T extending K, then $P(\phi(a), K) > 0$.*

(This is the "density" part of Theorem 11.4.1.)

Proof. Let M be the model where $\phi(a)$ holds. Let L be a Galois extension of K determining the truth of $\phi(a)$. Then $L \cap M \in \mathfrak{F}(L/K)$ and $f_{\phi(a),L}(L \cap M) = \top$. By Lemma 11.4.14, $P(\phi(a), K) > 0$. □

We have verified each condition of Theorem 11.4.1, which is now proven.

11.5 NTP_2 and the Independence Property

We show that the model companion T (usually) fails to be NIP, but is always NTP_2, the next best possibility. In many cases, perhaps all, these results are already known, though they may not have been written down yet.

11.5.1 Failure of NIP

If $n = 1$, then $T = T_i$ is one of ACVF, RCF, or pCF, which are all known to be NIP. On the other hand,

Theorem 11.5.1. *Suppose $n > 1$. Then T has the independence property.*

Other people have already pointed this out, but here is a proof.

Proof. We give a proof which works in characteristic $\neq 2$. It is not hard to modify it to work in characteristic 2.

Claim 11.5.2. *For each i, we can produce quantifier-free $\mathcal{L}_i$-formulas $\phi_i(x, y)$ and $\chi_i(y)$ without parameters such that x, y are singletons, and such that if $K_i \models T_i$, then $\chi_i(K_i)$ is a non-empty open set and for every $b \in \chi_i(K_i)$, both square roots of b are present in K_i, and exactly one of them satisfies $\phi_i(x, b)$.*

Proof. If T_i is RCF, let $\chi_i(y)$ say that $y > 0$ and $\phi_i(x, y)$ say that $x > 0$. If T_i is ACVF, let $\chi_i(y)$ say that $v(y-1/4) > 0$, and $\phi_i(x, y)$ say that $v(x-1/2) > 0$. Note that if $v(y-1/4) > 0$ and $x^2 = y$, then $t = x - 1/2$ satisfies

$$t^2 + t + 1/4 - y = (t + 1/2)^2 - y = 0.$$

By Newton polygons, one of the possibilities for t has valuation zero, and the other has valuation $v(y - 1/4) > 0$. If T_i is pCF, the same formulas work as in the case of ACVF. The only thing to check is that if $v(y - 1/4) > 0$ for some $y \in K \models pCF$, then the two roots of $T^2 + T + (1/4 - y) = 0$ are present in K. If not, then since the two roots have different valuations (in an ambient model of ACVF), there are two different ways to extend the valuation from K to $K[T]/(T^2 + T + (1/4 - y))$, contradicting Henselianity of K. □

Given the ϕ_i and χ_i from the Claim, let $\chi(y) = \bigwedge_{i=1}^{n} \chi_i(y)$. Note that $\chi(y)$ defines an infinite subset of any model of T, by condition A2 of §11.2.1. (Each $\chi_i(-)$ is $\mathbb{A}^1$-dense.) If $K \models T$ and $b \in \chi(K)$, then $X^2 - b$ has roots in K by choice of $\chi_i(-)$ and condition A1' of §11.2.1. So each element of $\chi(K)$ is a square.

Let $\psi(y)$ assert that $\chi(y)$ holds and there is a square root of y which satisfies exactly one of ϕ_1 and ϕ_2. Note that if $\chi(b)$ holds, then both square roots of b are present in K, exactly one of them satisfies ϕ_1, and exactly one of them satisfies ϕ_2. Letting, $\oplus$ denote exclusive-or, we can write $\psi(y)$ as $\phi_1(\sqrt{y}) \oplus \phi_2(\sqrt{y})$, where the choice of $\sqrt{y}$ is unimportant.

Let K be a model of T. We will show that $\psi(x+y)$ has the independence property in K. Let $a_1, \ldots, a_m$ be any m elements in $\chi(K)$, which as we noted above is an infinite set. We will show that for any subset $S_0 \subseteq \{1, \ldots, m\}$, there is a b in K such that $j \in S_0 \iff K \models \psi(b + a_j)$. It suffices to find such a b in an elementary extension of K, rather than K itself. Let $K' \succeq K$ be an elementary extension containing an element ϵ which is infinitesimal compared to K, with respect to every one of the valuations. That is, for each i such that T_i is valuative, we want $v_i(\epsilon) > v_i(K)$, and for each i such that T_i is RCF, we want $-\alpha <_i \epsilon <_i \alpha$ for every $\alpha >_i 0$ in K. The fact that such an ϵ exists follows by our axiom A2, and can be shown directly.

Note that for $1 \leq j \leq m$, $a_j + \epsilon \in \chi(K')$. (Indeed, for every i, $K' \models \chi_i(a_i + \epsilon)$, because $\chi_i(-)$ defines an open set in a model of T_i, and ϵ is infinitesimal with respect to the prime model of T_i over $K \upharpoonright \mathcal{L}_i$.) Consequently, $\sqrt{a_j + \epsilon} \in K'$ for every $1 \leq j \leq m$. Let L be $K(\sqrt{a_j + \epsilon} : 1 \leq j \leq m) \subseteq K'$, as a model of $T_\forall$. Since ϵ is transcendental over K, $\mathrm{Gal}(L/K(\epsilon)) \cong (\mathbb{Z}/2\mathbb{Z})^m$. In particular, for every $S \subseteq \{1, \ldots, m\}$, there is a field automorphism $\sigma_S \in \mathrm{Gal}(L/K(\epsilon))$ which swaps the square roots of $a_j + \epsilon$ if and only if $j \in S$. Let L_S be the $T_\forall$-model with underlying field L, with the same $\mathcal{L}_i$-structure as L for $i > 1$, and with the $\mathcal{L}_1$-structure obtained by pulling back the $\mathcal{L}_1$-structure of L along σ_S. If Δ denotes symmetric difference of sets, then

$$\begin{aligned}
\{j : L_S \models \phi_2(\sqrt{a_j + \epsilon})\} &= \{j : L \models \phi_2(\sqrt{a_j + \epsilon})\} \\
\{j : L_S \models \phi_1(\sqrt{a_j + \epsilon})\} &= \{j : L \models \phi_1(\sigma_S(\sqrt{a_j + \epsilon}))\} \\
&= \{j : L \models \phi_1(\sqrt{a_j + \epsilon})\} \ \Delta \ S,
\end{aligned}$$

where the last equality holds because $L \models \phi_1(\sqrt{a_j + \epsilon}) \iff \neg\phi_1(-\sqrt{a_j + \epsilon})$. Now let K_S be a model of T extending L_S. Since L_S is a model of $T_\forall$ extending K, $K_S \succeq K$. Also,

$$\begin{aligned}
\{j : K_S \models \psi(a_j + \epsilon)\} &= \{j : L_S \models \phi_1(\sqrt{a_j + \epsilon})\} \Delta \{j : L_S \models \phi_2(\sqrt{a_j + \epsilon})\} \\
&= \{j : L \models \phi_1(\sqrt{a_j + \epsilon})\} \Delta \{j : L \models \phi_2(\sqrt{a_j + \epsilon})\} \Delta S \\
&= \{j : K' \models \psi(a_j + \epsilon)\} \Delta S
\end{aligned}$$

Therefore, by choosing $S = S_0 \ \Delta \ \{j : K' \models \psi(a_j + \epsilon)\}$, we can arrange that

$$\{j : K_S \models \psi(a_j + \epsilon)\} = S_0,$$

i.e., $K_S \models \psi(a_j + \epsilon)$ if and only if $j \in S_0$. Taking b to be $\epsilon \in K_S$, this completes the proof. □

Because T has the independence property and clearly has the strict order property, the best classification-theoretic property we could hope for T to have is NTP_2.

11.5.2 NTP$_2$ holds

First we make some elementary remarks about relative algebraic closures.

Lemma 11.5.3. *Let M be a pure field. Let K be a subfield of M which is relatively separably closed in M (in the sense of Definition 11.1.6). Let a and b be two tuples from M such that $a \mathop{\smile\!\!\!\!\!\!|}_K^{ACF} b$, i.e., a and b are algebraically independent from each other over K. Then $K(a)$ is relatively separably closed in $K(a,b)$.*

Proof. Embed M into a monster model $\mathfrak{M} \models ACF$. By the remarks after Definition 11.1.6, $\text{tp}(a/K)$ and $\text{tp}(b/K)$ are stationary. Since $a \mathop{\smile\!\!\!\!\!\!|}_K b$, the type of b over $\text{acl}(K(a))$ is K-definable. Now suppose that some singleton $c \in K(a,b)$ is algebraic over $K(a)$. Write c as $f(a,b)$, for some rational function $f(X,Y) \in K(X,Y)$. Note that $\text{stp}(b/K(a))$ includes the statement $f(a,x) = c$. On the other hand, it does not include $f(a,x) = c'$ for any conjugate $c' \neq c$ of c over $K(a)$. As $\text{stp}(b/K(a))$ is definable over K, ac and ac' cannot have the same type over K. But if c and c' are conjugate over $K(a)$, then ac and ac' have the same type. So c' does not exist, and c has no other conjugates over $K(a)$. Thus $c \in \text{dcl}(K(a))$. So $c^{p^k} \in K(a)$ for some k. As c was an arbitrary element of $K(a,b) \cap K(a)^{alg}$, we see that $K(a)$ is relatively separably closed in $K(a,b)$. □

Lemma 11.5.4. *Let M be a pure field. Suppose $K_0 \subseteq K_1 \subseteq K_2$ are three subfields of M, each relatively separably closed in M. Let c be tuple from M, possibly infinite. Suppose that $c \mathop{\smile\!\!\!\!\!\!|}_{K_0}^{ACF} K_2$, i.e., K_2 and c are algebraically independent over K_0. Then $K_1(c)$ is relatively separably closed in $K_2(c)$.*

Proof. As in the previous lemma, embed M into a monster model $\mathfrak{M}$ of ACF. Then $c \mathop{\smile\!\!\!\!\!\!|}_{K_0} K_2$, and by properties of forking, $K_1(c) \mathop{\smile\!\!\!\!\!\!|}_{K_1} K_2$. By the previous lemma, $K_1(c)$ is relatively separably closed in $K_2K_1(c) = K_2(c)$. □

Now we return to existentially closed fields with valuations and orderings. As always, T is the model companion.

Lemma 11.5.5. *In a monster model of T, let B be a small set of parameters and $a_1, a_2, \ldots$ be a B-indiscernible sequence. Suppose that $B = \text{acl}(B)$ and $a_i = \text{acl}(Ba_i)$ for any/every i. Suppose also that $a_j \mathop{\smile\!\!\!\!\!\!|}_B^{ACF} a_{<j}$ for every j, i.e., the sequence is algebraically independent over B. Let c be a finite tuple and suppose that $a_1, a_2, \ldots$ is quantifier-free indiscernible over cB, i.e., if $i_1 < \cdots < i_m$ and $j_1 < \cdots < j_m$, then*

$$\text{qftp}(a_{i_1} a_{i_2} \cdots a_{i_m}/cB) = \text{qftp}(a_{j_1} a_{j_2} \cdots a_{j_m}/cB).$$

Let $\phi(x;y)$ be a formula over B such that $\phi(c;a_1)$ holds. Then $\bigwedge_{j=1}^{\infty} \phi(x;a_j)$ is consistent.

Proof. Because $a_1, a_2, \ldots$ is B-indiscernible, it suffices to show for each k that $\bigwedge_{j=1}^{\infty} \phi(x;a_j)$ is not k-inconsistent.

First observe that whether or not $c \mathop{\smile\!\!\!\!\!\!|}_B^{ACF} a_j$ holds depends only on the quantifier-free type of c and a_j over B. In particular, it does not depend on j, by quantifier-free indiscernibility of $a_1, a_2, \ldots$ over cB. If $c \not\mathop{\smile\!\!\!\!\!\!|}_B^{ACF} a_j$ for one j, then this holds for all j. As the a_j are an algebraically independent sequence over B, this contradicts the fact that finite tuples have finite preweight in ACF. So $c \mathop{\smile\!\!\!\!\!\!|}_B^{ACF} a_j$ for each j. The same argument applied to the sequence $a_1a_2, a_3a_4, \ldots$ shows that $c \mathop{\smile\!\!\!\!\!\!|}_B^{ACF} a_1a_2$. Similarly $c \mathop{\smile\!\!\!\!\!\!|}_B^{ACF} a_1a_2a_3$, and so on, and so $c \mathop{\smile\!\!\!\!\!\!|}_B^{ACF} a_1a_2a_3 \ldots$.

Let M be the monster model. Any subset of M closed under $\mathrm{acl}(-)$ is relatively algebraically closed in M, hence relatively separably closed in M. In particular, if we let $K_0 = B$, $K_1 = B(a_j) = a_j$, and $K_2 = \mathrm{acl}(Ba_1a_2\ldots)$, then each of K_0, K_1, K_2 is relatively separably closed in M, and $K_0 \subseteq K_1 \subseteq K_2$. By the previous paragraph, $c \mathop{\smile\!\!\!\!\!\!|}_B^{ACF} K_2$, so by Lemma 11.5.4, we conclude that $K_1(c)$ is relatively separably closed in $K_2(c)$, i.e., $B(a_j, c)$ is relatively separably closed in $K_2(c)$. Using bars to denote perfect closures, this means that $\overline{B(a_j, c)}$ is relatively algebraically closed in $\overline{K_2(c)}$.

Recall the function $P(-,-)$ from Theorem 11.4.1. By the "extension invariance" part of that theorem,

$$P(\phi(c; a_j); \overline{B(a_j, c)}) = P(\phi(c; a_j); \overline{K_2(c)}).$$

Now by quantifier-free indiscernibility of $a_1, a_2, \ldots$ over cB, we see that $B(a_j, c) \cong B(a_{j'}, c)$ for all j, j'. By the isomorphism-invariance part of Theorem 11.4.1,

$$P(\phi(c; a_j); \overline{B(a_j, c)}) = P(\phi(c; a_{j'}); \overline{B(a_{j'}, c)})$$

for all j, j'. Consequently, $P(\phi(c; a_j); \overline{K_2(c)})$ does not depend on j.

Now M is a model of T extending $\overline{K_2(c)}$, and in M, $\phi(c; a_1)$ holds. So by the "density" part of Theorem 11.4.1, $P(\phi(c; a_1); \overline{K_2(c)})$ is some positive number $\epsilon > 0$. Consequently, $P(\phi(c; a_j); \overline{K_2(c)}) = \epsilon > 0$ for every j.

Suppose for the sake of contradiction that $\bigwedge_{j=1}^{\infty} \phi(x; a_j)$ is k-inconsistent for some k. Let N be big enough that $N\epsilon > k$. Let $\psi(x)$ be the statement over K_2 asserting that at least k of $\phi(x; a_1), \ldots, \phi(x; a_N)$ hold. By the Keisler measure part of Theorem 11.4.1, $P(\psi(c); \overline{K_2(c)}) > 0$, and there is a model M' of T extending $\overline{K_2(c)}$ in which ψ holds. In particular, $M' \models \exists x \psi(x)$. But K_2 is relatively algebraically closed in M, hence satisfies axiom A1 of §11.2.1 by Corollary 11.2.10. By Corollary 11.2.11, the statement $\exists x \psi(x)$ holds in M if and only if it holds in M'. Consequently, it holds in M, and therefore $\bigwedge_{j=1}^{n} \phi(x; a_j)$ is not k-inconsistent. □

Recall from [1] or [9] that the burden of a partial type $p(x)$ is the supremum of κ such that there is an inp-pattern in $p(x)$ of depth κ, that is, an array of formula $\phi_i(x; a_{ij})$ for $i < \kappa$ and $j < \omega$, and some $k_i < \omega$ such that the ith row $\bigwedge_{j<\omega} \phi_i(x; a_{ij})$ is k_i-inconsistent for each i, and such that for any $\eta : \kappa \to \omega$, the corresponding downwards path $\bigwedge_{i<\kappa} \phi_i(x; a_{i,\eta(i)})$ is consistent with $p(x)$. A theory is NTP_2 if every partial type has burden less than ∞. A theory is strong if every partial type has burden less than $\aleph_0$, roughly. (See [1] for a more precise statement.) At any rate, if every partial type has burden less than $\aleph_0$, then the

theory is strong. By the submultiplicativity of burden (Theorem 11 in [9]), it suffices to check the burden of the home sort.

Fact 11.5.6. *If D and E are definable sets,* $\mathrm{bdn}(D \times E) \geq \mathrm{bdn}(D) + \mathrm{bdn}(E)$.

In fact, if $\phi_i(x; a_i)$ is an inp-pattern for D and $\psi_j(y; b_j)$ is an inp-pattern for E, then $\{\phi_i(x; a_i)\} \cup \{\phi_j(y; b_j)\}$ is an inp-pattern for $D \times E$.

In NIP theories, burden is the same thing as dp-rank, which is known to be additive [41]. The theories ACVF, pCF, and RCF are all known to be dp-minimal, i.e., to have dp-rank 1 [15]. One of the descriptions of dp-rank is that a partial type $\Sigma(x)$ over a set C has dp-rank $\geq \kappa$ if and only if there are κ-many mutually indiscernible sequences over C and a realization a of $\Sigma(x)$ such that each sequence is not indiscernible over Ca.

Theorem 11.5.7. *The model companion T is NTP_2, and strong. In fact, the burden of affine m-space is exactly mn, where n is the number of valuations and orderings.*

Proof. To show that the burden of $\mathbb{A}^m$ is at least mn, it suffices by Fact 11.5.6 to show that $\mathrm{bdn}(\mathbb{A}^1) \geq n$. In the case where every T_i is ACVF, one can take $\phi_i(x; y)$ to assert that $v_i(x) = y$, for $1 \leq i \leq n$, and take $a_{i,0}, a_{i,1}, \ldots$ to be an increasing sequence in the ith valuation group. Variations on this handle the remaining cases. We leave the details as an exercise to the reader.

For the upper bound, suppose for the sake of contradiction that there is an inp-pattern $\{\phi_i(x; a_{ij})\}_{i<mn+1;\ 0\leq j<\omega}$ of depth $mn + 1$, with x a tuple of length m. We may assume that the a_{ij} form a mutually $\emptyset$-indiscernible array. Extend the array to the left, i.e., let j range over negative numbers. Let B be $\mathrm{acl}(a_{ij} : j < 0)$. From stability theory, one knows that $a_{ij} \mathop{\downarrow}^{ACF}_{B} a_{i0}a_{i1}\cdots a_{i,j-1}$ for every j. By mutual indiscernibility, each sequence $a_{i0}, a_{i1}, \ldots$ is indiscernible over $\{a_{ij} : j < 0\}$, hence over B. In particular, $a_{ij} \equiv_B a_{ij'}$ for $j \neq j'$. For each $i < mn + 1$, let b_{i0} be an enumeration of $\mathrm{acl}(Ba_{i,0})$. For $j > 0$, choose $b_{i,j}$ such that $a_{i,j}b_{i,j} \equiv_B a_{i,0}b_{i,0}$. Then $b_{i,j}$ is an enumeration of $\mathrm{acl}(Ba_{i,j})$ for every i and every $j \geq 0$. Let $c_{i,j}d_{i,j}$ be a mutually B-indiscernible array modeled on the array $a_{i,j}b_{i,j}$. Then $c_{i,j}d_{i,j} \equiv_B a_{i,0}b_{i,0}$, so $d_{i,j}$ is an enumeration of $\mathrm{acl}(Bc_{i,j})$. Also, because $a_{i,0}, a_{i,1}, \ldots$ was already B-indiscernible, we must have

$$c_{i,0}c_{i,1}\cdots \equiv_B a_{i,0}a_{i,1}\cdots$$

for each i. Consequently, $c_{i,j} \mathop{\downarrow}^{ACF}_{B} c_{i,0}\cdots c_{i,j-1}$. And since $d_{i,j} \subseteq \mathrm{acl}(Bd_{i,j})$, we also have

$$d_{i,j} \mathop{\underset{B}{\downarrow}}^{ACF} d_{i,0}d_{i,1}\cdots d_{i,j-1},$$

using Corollary 11.2.12. As $b_{i,0}$ is an enumeration of $\mathrm{acl}(Ba_{i,0})$, the elements of $a_{i,0}$ must actually appear somewhere in $b_{i,0}$. Let π_i be the coordinate projection such that $\pi_i(b_{i,0}) = a_{i,0}$. Hence $c_{i,j} = \pi_i(d_{i,j})$.

Because the a_{ij} formed a mutually $\emptyset$-indiscernible array, the collective type of all the $c_{i,j}$'s must agree with that of all the $a_{i,j}$'s. Hence $\phi_i(x; c_{i,j})$ is still an inp-pattern of depth $mn+1$. Let $\psi_i(x; y)$ be $\phi_i(x; \pi_i(y))$. Then $\psi_i(x; d_{i,j})$ is an inp-pattern of depth $mn + 1$. Let c be a realization of $\bigwedge_{i<mn+1} \psi_i(x; d_{i,0})$. Note that c is a tuple of length m.

Let $\mathfrak{M}$ be the ambient monster. For each $1 \leq k \leq n$, let $\mathfrak{M}_k$ be a model of T_k extending $\mathfrak{M} \upharpoonright \mathcal{L}_k$. By quantifier-elimination, the array $\{d_{i,j}\}$ is still mutually B-indiscernible in $\mathfrak{M}_i$. By additivity of dp-rank and by dp-minimality of the home sort in $\mathfrak{M}_k$, we know that the dp-rank of $\mathrm{tp}(c/B)$ in $\mathfrak{M}_k$ is at most m. In particular, for each $1 \leq k \leq n$, at most m of the rows in the array $\{d_{i,j}\}$ can fail to be Bc-indiscernible in $\mathfrak{M}_k$. By the pigeonhole principle, there must be some value of i such that the sequence $d_{i,0}, d_{i,1}, \ldots$ is Bc-indiscernible in each of $\mathfrak{M}_1, \mathfrak{M}_2, \ldots, \mathfrak{M}_n$. Back in $\mathfrak{M}$, this means that $d_{i,0}, d_{i,1}, \ldots$ is quantifier-free Bc-indiscernible. Since $d_{i,0}, d_{i,1}, \ldots$ is also B-indiscernible and B-independent, Lemma 11.5.5 applies. Consequently, $\bigwedge_{j=0}^{\infty} \psi_i(x; d_{i,j})$ is consistent, because $\psi_i(c; d_{i,0})$ holds. This contradicts the fact that $\{\psi_i(x; d_{i,j})\}$ is an inp-pattern. □

11.6 Forking and Dividing

We will make use of the following general fact, which is the implication (ii) $\Longrightarrow$ (i) in Proposition 4.3 of [36].[3]

Fact 11.6.1. *Let $\mathfrak{M}$ be a monster model of some theory, let $S \subseteq \mathfrak{M}$ be a small set, and let $\phi(x)$ be a formula with parameters from $\mathfrak{M}$. Suppose there is a global Keisler measure μ which is Lascar-invariant over S, and suppose $\mu(\phi(x)) > 0$. Then $\phi(x)$ does not fork over S.*

Now we specialize to the theory T under consideration.

Lemma 11.6.2. *Let $\mathfrak{M}$ be a monster model of T. Let S be a small subset of $\mathfrak{M}$, and let p be a complete quantifier-free type over $\mathfrak{M}$ which is Lascar-invariant over S. Then there is a Keisler measure μ on $S(\mathfrak{M})$, Lascar-invariant over S, whose support is exactly the set of completions of p.*

This is nothing but a restatement or special case of Theorem 11.4.1.

Proof. Let a be a realization of p in some bigger model, and consider the structure $\mathfrak{M}[a]$ generated by $\mathfrak{M}$ and a. The structure of $\mathfrak{M}[a]$ is determined by p. Also, if σ is any Lascar strong automorphism of $\mathfrak{M}$ over S, then $p = \sigma(p)$. This implies that there is a uniquely determined automorphism σ' of $\mathfrak{M}[a]$ extending σ on $\mathfrak{M}$ and fixing a.

Let $\overline{\mathfrak{M}[a]}$ denote the perfect closure of the field of fractions of $\mathfrak{M}[a]$. This is uniquely determined (as a model of $T_\forall$) by $\mathfrak{M}[a]$, and hence is determined by p. Let μ be the Keisler measure on $\mathfrak{M}$ which assigns to an $\mathfrak{M}$-formula $\phi(x; b)$ the value

$$P(\phi(a; b); \overline{\mathfrak{M}[a]}),$$

[3] Hrushovski and Pillay assume NIP, but the assumption is unused for the implication (ii) $\Longrightarrow$ (i).

where P is as in Theorem 11.4.1. By the Keisler measure part of Theorem 11.4.1, this is a Keisler measure on the space of completions of qftp($\mathfrak{M}[a]$). By model completeness, any extension of qftp($\mathfrak{M}[a]$) to a complete type must satisfy tp($\mathfrak{M}$), so we have a legitimate Keisler measure on the space of extensions of p to complete types over $\mathfrak{M}$. And if σ is any Lascar strong automorphism over S, then by the "isomorphism invariance" part of Theorem 11.4.1,

$$P(\phi(a;\sigma(b));\overline{\mathfrak{M}[a]}) = P(\phi(\sigma'(a);\sigma'(b));\overline{\mathfrak{M}[a]}) = P(\phi(a,b);\overline{\mathfrak{M}[a]})$$

where σ' is the aforementioned automorphism of $\mathfrak{M}[a]$ extending σ and fixing a. Thus $\mu(\phi(x;b)) = \mu(\phi(x;\sigma(b)))$. We conclude that $\mu(\phi(x;b)) = \mu(\phi(x;b'))$ for any formula $\phi(x;y)$ and any $b, b' \in \mathfrak{M}$ having the same Lascar strong type over S. Finally, if b is a tuple from $\mathfrak{M}$ and $\phi(a;b)$ is a formula which is consistent with p, then $\phi(a;b)$ is also consistent with the diagram of $\mathfrak{M}[a]$, hence has positive probability by the "density" part of Theorem 11.4.1. □

Corollary 11.6.3. *Let $\mathfrak{M}$ be a monster model of T and S be a small subset of $\mathfrak{M}$. Suppose q is a complete quantifier-free type on $\mathfrak{M}$ which is Lascar invariant over S. Then every complete type on $\mathfrak{M}$ extending q does not fork over S.*

Proof. Let $p(x)$ be a complete type extending $q(x)$. Let $\phi(x)$ be any formula from $p(x)$. Let μ be the Keisler measure from Lemma 11.6.2. Then μ is Lascar invariant over S, and $\mu(\phi(x))$ is positive because $\phi(x)$ is consistent with $q(x)$. By Fact 11.6.1, $\phi(x)$ does not fork over S. □

If M is a model of T and A, B, C are subsets of M, let $A \mathop{\smile\!\!\!\!\!\!|}^{T_i}_C B$ indicate that $A \mathop{\smile\!\!\!\!\!\!|}_C B$ holds in any/every model of T_i extending $M \restriction \mathcal{L}_i$.

Lemma 11.6.4. *Work in a monster model $\mathfrak{M}$ of T. Let a be a finite tuple, and B and C be sets (in the home sort, as always). Suppose $C = \mathrm{acl}(C)$. Suppose $a \mathop{\smile\!\!\!\!\!\!|}^{T_i}_C B$ holds for every $1 \leq i \leq n$. Then* qftp(a/BC) *can be extended to a quantifier-free type $q(x)$ on $\mathfrak{M}$ which is Lascar invariant over C.*

Proof. Let V be the variety over C of which a is a generic point. By Fact 11.1.8, V is absolutely irreducible.

Let $\mathfrak{M}_i$ be a model of T_i extending $\mathfrak{M} \restriction \mathcal{L}_i$. Within $\mathfrak{M}_i$, $a \mathop{\smile\!\!\!\!\!\!|}^{T_i}_C B$. By Adler's characterization of forking in NIP theories (Proposition 2.1 in [36]), there is an $\mathcal{L}_i$-type $p_i(x)$ on $\mathfrak{M}_i$ which extends the type of a over BC and which is Lascar-invariant over C. The restriction of this $\mathcal{L}_i$-type to a quantifier-free $\mathcal{L}_{rings}$-type must say that x lives on V and on no $\mathfrak{M}_i$-definable proper subvarieties of V. This follows from Lemma 11.1.17. Let $q_i(x)$ be the set of quantifier-free $\mathcal{L}_i$-statements in $p_i(x)$ with parameters from $\mathfrak{M}$. Then $q_i(x)$ is a complete quantifier-free $\mathcal{L}_i$-type on $\mathfrak{M}$. Let $q(x)$ be $\bigcup_{i=1}^n q_i(x)$. This is a complete quantifier-free type on $\mathfrak{M}$; it is consistent because the $q_i(x)$ all have the same restriction to the language of rings, namely, the generic type of V. Also, $q(x)$ extends qftp(a/BC), because the $\mathcal{L}_i$-part of qftp(a/BC) is present in $p_i(x)$ and $q_i(x)$.

To show Lascar-invariance of $q(x)$ over C, it suffices to show that if I is a C-indiscernible sequence in $\mathfrak{M}$, a and a' are two elements of I, and $\phi(x;y)$ is a quantifier-free formula, then $\phi(x;a) \in q(x)$ if and only if $\phi(x;a') \in q(x)$. In fact, we only need to consider the case where $\phi(x)$ is a quantifier-free $\mathcal{L}_i$-formula, for some i. But then

$$\phi(x;a) \in q(x) \iff \phi(x;a) \in p_i(x) \iff \phi(x;a') \in p_i(x) \iff \phi(x;a') \in q(x)$$

where the middle equivalence follows from the fact that $p_i(x)$ is Lascar-invariant, and I is C-indiscernible within $\mathfrak{M}_i$ (by quantifier-elimination in T_i). Thus $q(x)$ is Lascar-invariant over C, as claimed. □

Theorem 11.6.5. *Forking and dividing agree over every set (in the home sort).*

Proof. First we show that if a is a finite tuple and B is a set, then $\mathrm{qftp}(a/B)$ does not fork over B. By Lemma 11.6.4, there is a global quantifier-free type $q(x)$ which is Lascar-invariant over B. By Corollary 11.6.3, any extension of $q(x)$ to a complete global type does not fork over B. So $\mathrm{qftp}(a/B)$ has a global non-forking extension. Now if a is any small tuple, and B is a set, then $\mathrm{qftp}(a/B)$ does not fork over B, by compactness. Consequently, if a is a small tuple and B is a (small) set, then $\mathrm{qftp}(a'/B)$ does not fork over B, where a' enumerates $\mathrm{acl}(aB)$. By Corollary 11.2.11, $\mathrm{qftp}(a'/B)$ implies $\mathrm{tp}(a'/B)$, so $\mathrm{tp}(a'/B)$ does not fork over B. By monotonicity, $\mathrm{tp}(a/B)$ does not fork over B. As a and B are arbitrary, every set in the home sort is an extension base for forking in the sense of [10], so by Theorem 1.2 in [10], forking and dividing agree over every set in the home sort. □

Lemma 11.6.6. *Let $\mathfrak{M}$ be a monster model of T and $C = \mathrm{acl}(C)$ be a small subset of M. Suppose $p(x)$ is a complete type on C and $q(x)$ is a complete quantifier-free type on M, with $q(x)$ extending the quantifier-free part of $p(x)$. Suppose $q(x)$ is Lascar-invariant over C. Then $q(x) \cup p(x)$ is consistent.*

Proof. Let $\mathfrak{M}[a]$ be the structure obtained by adjoining a realization a of $q(x)$ to $\mathfrak{M}$. Let W be the variety over $\mathfrak{M}$ of which a is the generic point. By Fact 11.1.8, W is absolutely irreducible. Moreover, the ACF-theoretic code $\ulcorner W \urcorner$ for W must lie in $\mathfrak{M}$. By Lascar invariance of $q(x)$, one sees that W is Lascar invariant over C. Consequently, the finite tuple $\ulcorner W \urcorner$ is fixed by every Lascar strong automorphism over C. So $\ulcorner W \urcorner \subseteq \mathrm{acl}(C) = C$. Consequently, in an ambient model of ACF we have $Cb(\mathrm{stp}(a/\mathfrak{M})) \subseteq C$, and so $a \mathop{\smile\!\!\!\!\!\!|}^{ACF}_{C} \mathfrak{M}$. By Lemma 11.5.3, $C(a)$ is relatively algebraically closed in $\mathfrak{M}(a)$.

Because the quantifier-free type of a over C is consistent with $p(x)$, there is a model $N \models T$ extending $C[a]$ such that within N, $\mathrm{tp}(a/C) = p(x)$. By Lemma 11.2.9, we can amalgamate $\mathfrak{M}(a)$ and N over $C(a)$. So there is a model N' of T extending N and $\mathfrak{M}(a)$. In N, $\mathrm{tp}(a/C) = p(x)$. As $N \preceq N'$, $\mathrm{tp}(a/C) = p(x)$ holds in N' as well. And as $N' \succeq \mathfrak{M}(a)$, $\mathrm{qftp}(a/\mathfrak{M}) = q(x)$. So $q(x) \cup p(x)$ is consistent. □

Lemma 11.6.7. *Work in a monster model $\mathfrak{M}$ of T. Let a be a finite tuple, and B and C be sets (in the home sort, as always). Suppose $a \mathop{\smile\!\!\!\!\!\!|}^{T_i}_{C} B$ holds for every $1 \leq i \leq n$. Then $a \mathop{\smile\!\!\!\!\!\!|}_{C} B$.*

Proof. A type forks/divides over C if and only if it forks/divides over $\mathrm{acl}(C)$, so it suffices to show that $\mathrm{tp}(a/BC)$ does not fork over $\mathrm{acl}(C)$. By monotonicity, it suffices to show that $\mathrm{tp}(a/\mathrm{acl}(BC))$ does not fork over $\mathrm{acl}(C)$. By Claim 3.6 in [10] and Lemma 11.1.15 above, $a \downarrow^{T_i}_{\mathrm{acl}(C)} \mathrm{acl}(CB)$ for every i. So we may assume that $C = \mathrm{acl}(C) \subseteq B = \mathrm{acl}(B)$.

Now by Lemma 11.6.4, there is a global quantifier-free type $q(x)$ extending $\mathrm{qftp}(a/BC) = \mathrm{qftp}(a/B)$, with $q(x)$ Lascar-invariant over C. Clearly $q(x)$ is also Lascar-invariant over B, so by Lemma 11.6.6, $q(x)$ is consistent with $\mathrm{tp}(a/B)$. Let $p(x)$ be a global complete type extending $q(x) \cup \mathrm{tp}(a/B)$. Then $p(x)$ does not fork over C by Corollary 11.6.3. □

Let $\mathrm{qftp}^i(a/B)$ denote the quantifier-free $\mathcal{L}_i$-type of a over B, and let $\mathrm{qftp}^{ACF}(a/B)$ denote the field-theoretic quantifier-free type of a over B.

Lemma 11.6.8. *Let $\mathfrak{M}$ be a monster model of T, and let $C = \mathrm{acl}(C)$ be a small subset. For each i, let M_i be a model of T_i extending $\mathfrak{M} \restriction \mathcal{L}_i$. For each i, let a_i be a tuple in M_i. Suppose that $\mathrm{qftp}^{ACF}(a_i/C)$ does not depend on i. Then we can find a tuple a in $\mathfrak{M}$ such that $\mathrm{qftp}^i(a/C) = \mathrm{qftp}^i(a_i/C)$ for every i.*

Proof. Let $C[a_i]$ denote the subring or subfield of M_i generated by C and a_i. By assumption, $C[a_i]$ is isomorphic to $C[a_{i'}]$ as a ring, for every i and i'. Use these isomorphisms to identify all the $C[a_i]$ with each other, getting a single ring $C[a]$ which is isomorphic to $C[a_i]$ for every i. Use these isomorphisms to move the $(T_i)_\forall$ structure from $C[a_i]$ to $C[a]$. Now $C[a]$ is a model of $T_\forall$, and $\mathrm{qftp}^i(a/C) = \mathrm{qftp}^i(a_i/C)$, for every i. As $C = \mathrm{acl}(C)$, C is relatively separably closed in $\mathfrak{M}$, so by Lemma 11.2.9, one can embed $C[a]$ and $\mathfrak{M}$ into a bigger model of T. By model completeness and saturation, $\mathrm{tp}(a/C)$ is already realized in $\mathfrak{M}$. □

Lemma 11.6.9. *Let a, B, C be small subsets of a monster model $\mathfrak{M} \models T$. Suppose $a \not\downarrow^{T_1}_C b$. Then $a \not\downarrow_C b$.*

Proof. By Claim 3.6 in [10] applied to both T_1 and T, we may assume $C = \mathrm{acl}(C)$ and $B = \mathrm{acl}(BC)$. By finite character of forking, we may assume a is finite. For every i, let $\mathfrak{M}_i$ be an even more monstrous model of T_i extending $\mathfrak{M} \restriction \mathcal{L}_i$. As $a \not\downarrow^{T_1}_C B$, within $\mathfrak{M}_1$ we have $a \not\downarrow_C B$. By Lemma 11.1.15, some $\mathcal{L}_1$-formula $\phi(x; B)$ in $\mathrm{tp}(a/BC)$ divides over C. By quantifier-elimination in T_i, we may assume that $\phi(x; y)$ is a quantifier-free $\mathcal{L}_1$-formula. By Lemma 11.1.18, there is a sequence $B = B^1_0, B^1_1, B^1_2, \ldots$ in $\mathfrak{M}_1$ which is indiscernible over C and algebraically independent over C, and such that $\bigwedge_{j=0}^{\infty} \phi(x; B^1_j)$ is k-inconsistent in $\mathfrak{M}_1$, for some k. Thus $\mathrm{qftp}^1(B^1_j/C) = \mathrm{qftp}^1(B/C)$, and in a certain sense

$$\mathrm{qftp}^{ACF}(B^1_0 B^1_1 B^1_2 \cdots / C) = \mathrm{qftp}^{ACF}(B/C) \otimes \mathrm{qftp}^{ACF}(B/C) \otimes \cdots .$$

The right hand side makes sense because C is relatively separably closed in B (Definition 11.1.6), so $\mathrm{qftp}^{ACF}(B/C)$ is stationary.

Meanwhile, for $i > 1$, we can apply Lemma 11.1.16 to $\mathfrak{M}_i$ and $\mathrm{tp}(B/C)$, getting a sequence $B = B^i_0, B^i_1, B^i_2, \ldots$ which is indiscernible over C and algebraically independent

over C. (Note that Lemma 11.1.16 is true even without the restriction that B be finite.) So again, we get $\mathrm{qftp}^i(B^i_j/C) = \mathrm{qftp}^i(B/C)$, and

$$\mathrm{qftp}^{ACF}(B^i_0B^i_1B^i_2\cdots/C) = \mathrm{qftp}^{ACF}(B/C)\otimes\mathrm{qftp}^{ACF}(B/C)\otimes\cdots.$$

In particular, $\mathrm{qftp}^{ACF}(B^i_0B^i_1B^i_2\cdots/C)$ does not depend on i, as i ranges from 1 to n. By Lemma 11.6.8, we can therefore find a sequence $B_0, B_1, \ldots$ in $\mathfrak{M}$ such that

$$\mathrm{qftp}^i(B_0B_1\ldots/C) = \mathrm{qftp}^i(B^i_0B^i_1B^i_2\ldots/C)$$

for every i. In particular, $\mathrm{qftp}^i(B_j/C) = \mathrm{qftp}^i(B^i_j/C) = \mathrm{qftp}^i(B/C)$. Because this holds for all i, $\mathrm{qftp}(B_j/C) = \mathrm{qftp}(B/C)$. Because $B = \mathrm{acl}(B)$, $\mathrm{qftp}(B/C) \vdash \mathrm{tp}(B/C)$ by Corollary 11.2.11. So $\mathrm{tp}(B_j/C) = \mathrm{tp}(B/C)$ for every j. Also,

$$\mathrm{qftp}^1(B_0B_1\ldots/C) = \mathrm{qftp}^1(B^1_0B^1_1\ldots/C)$$

implies that there is an automorphism σ of $\mathfrak{M}_1$ sending $B^1_0B^1_1\ldots$ to $B_0B_1\ldots$. Consequently, $\bigwedge_{j=0}^{\infty}\phi(x;B_j)$ is k-inconsistent in $\mathfrak{M}_1$. Clearly it is also k-inconsistent in $\mathfrak{M}$, because $\mathfrak{M}$ is smaller than $\mathfrak{M}_1$. Since B_0, B_1 is a sequence of realizations of $\mathrm{tp}(B/C)$, we conclude that $\phi(x;B)$ divides over C, in $\mathfrak{M}$. □

Theorem 11.6.10. *Let M be a model of T, and let A, B, C be subsets of M (in the home sort). The following are equivalent:*

- *$A \mathop{\smile\!\!\!\!\!\!|}_C B$, i.e., the type of A over BC does not fork over C.*
- *The type of A over BC does not divide over C.*
- *$A \mathop{\smile\!\!\!\!\!\!|}^{T_i}_C B$ for every $1 \leq i \leq n$.*

Proof. The first two bullet points are equivalent by Theorem 11.6.5. If $A \mathop{\smile\!\!\!\!\!\!|}_C B$, then by Lemma 11.6.9 $A \mathop{\smile\!\!\!\!\!\!|}^{T_1}_C B$. Similarly, $A \mathop{\smile\!\!\!\!\!\!|}^{T_i}_C B$ for every $1 \leq i \leq n$. Conversely, if $A \mathop{\smile\!\!\!\!\!\!|}^{T_i}_C B$ for every $1 \leq i \leq n$, then by Lemma 11.6.7, $a \mathop{\smile\!\!\!\!\!\!|}_C B$ for every finite subset $a \subseteq A$. By finite character of forking, $A \mathop{\smile\!\!\!\!\!\!|}_C B$. □

Chapter 12

Definable Strong Euler Characteristics on Pseudofinite Fields

12.1 Introduction

12.1.1 Euler characteristics

Let M be a structure and R be a ring. Let $\mathrm{Def}(M)$ denote the collection of definable sets in M. Recall the following definitions from [44] and [45]. An *R-valued Euler characteristic* is a function $\chi : \mathrm{Def}(M) \to R$ such that

- $\chi(\emptyset) = 0$
- $\chi(X) = 1$ if X is a singleton
- $\chi(X) = \chi(Y)$ if X and Y are in definable bijection.
- $\chi(X \times Y) = \chi(X) \cdot \chi(Y)$
- $\chi(X \cup Y) = \chi(X) + \chi(Y)$ if X and Y are disjoint.

If the following additional property holds, then χ is called a *strong* Euler characteristic:

> If $f : X \to Y$ is a definable function and $r \in R$ is such that for every $y \in Y$, $\chi(f^{-1}(y)) = r$, then
> $$\chi(X) = r \cdot \chi(Y).$$

An Euler characteristic χ is *definable* if for every definable function $f : X \to Y$ and every $r \in R$, the set $\{y \in Y : \chi(f^{-1}(y)) = r\}$ is definable.

12.1.2 Examples of Euler characteristics

The simplest example of an Euler characteristic is the counting function on a finite structure. If M is a finite structure, there is a $\mathbb{Z}$-valued Euler characteristic given by

$$\chi(X) = |X|$$

where $|X|$ denotes the size of X. This χ is always strong and 0-definable.

Another well-known example is the Euler characteristic on dense o-minimal structures [17]. If $(M, <, \ldots)$ is a dense o-minimal structure, there is a $\mathbb{Z}$-valued Euler characteristic on M, characterized by the fact that $\chi(C) = -1^{\dim C}$ for any open cell C. This Euler characteristic is strong and 0-definable. By work of Kamenkovich and Peterzil [40], it can be extended to M^{eq}. On o-minimal expansions of the reals, $\chi(X)$ agrees with the topological Euler characteristic for compact definable $X \subseteq \mathbb{R}^n$.

Pseudofinite structures have strong Euler characteristics arising from counting mod n. More precisely, if M is an ultraproduct of finite structures, there is a canonical strong Euler characteristic $\chi_n : \mathrm{Def}(M) \to \mathbb{Z}/n\mathbb{Z}$ defined in the following way. Let M be the ultraproduct $\prod_{i \in I} M_i/\mathcal{U}$, and $X = \phi(M; a)$ be a definable set. Choose a tuple $\langle a_i \rangle_{i \in I} \in \prod_{i \in I} M_i$ representing a. Then define $\chi_n(X)$ to be the ultralimit along $\mathcal{U}$ of the sequence

$$|\phi(M_i; a_i)| + n\mathbb{Z} \in \mathbb{Z}/n\mathbb{Z}$$

This ultralimit exists because $\mathbb{Z}/n\mathbb{Z}$ is finite.

More intuitively, if we take $\mathbb{Z}^* = \prod_{i \in I} \mathbb{Z}/\mathcal{U} \succeq \mathbb{Z}$, then there is a non-standard counting function $\chi^* : \mathrm{Def}(M) \to \mathbb{Z}^*$ assigning to each definable set $X \subseteq M^n$ its non-standard "size" in $\mathbb{Z}^*$. Then χ_n is the composition

$$\mathrm{Def}(M) \xrightarrow{\chi^*} \mathbb{Z}^* \to \mathbb{Z}^*/n\mathbb{Z}^* \xrightarrow{\sim} \mathbb{Z}/n\mathbb{Z}$$

The map χ^* happens to be a strong Euler characteristic itself, though we will not discuss it further.

The mod n Euler characteristics on pseudofinite structures need not be definable. A simple example of this is an ultraproduct of the totally ordered sets $\{0, 1, \ldots, m\}$ as $m \to \infty$. If $(M, <)$ is the resulting ultraproduct, then M admits no definable $\mathbb{Z}/n\mathbb{Z}$-valued Euler characteristics (for $n > 1$). Indeed, if χ is an Euler characteristic on M, consider the function

$$f(a) = \chi([0, a]) \in \mathbb{Z}/n\mathbb{Z}$$

Then $f(b) = f(a) + 1$ when b is the successor of a. The set $f^{-1}(0)$ must therefore contain every nth element of M, and hence cannot be definable, because M is (non-dense) o-minimal.

We will see below (Theorem 12.1.1.1) that this does *not* happen with ultraproducts of finite fields: the χ_n are always definable on ultraproducts of finite fields.

On an ultraproduct M of finite structures, these χ_n maps are compatible in the sense that the following diagram commutes when n divides m:

$$\begin{array}{ccc} \mathrm{Def}(M) & \xrightarrow{\chi_m} & \mathbb{Z}/m\mathbb{Z} \\ & \searrow^{\chi_n} & \downarrow \\ & & \mathbb{Z}/n\mathbb{Z} \end{array}$$

Consequently, they assemble into a map

$$\hat{\chi} : \mathrm{Def}(M) \to \hat{\mathbb{Z}}$$

where $\hat{\mathbb{Z}}$ is the ring $\varprojlim_{n\in\mathbb{N}} \mathbb{Z}/n\mathbb{Z}$.

More generally, if M is any structure, we will say that a map $\chi : \mathrm{Def}(M) \to \hat{\mathbb{Z}}$ is

1. an *Euler characteristic* if all the compositions $\mathrm{Def}(M) \to \hat{\mathbb{Z}} \to \mathbb{Z}/n\mathbb{Z}$ are Euler characteristics
2. a *strong* Euler characteristic if all the compositions $\mathrm{Def}(M) \to \hat{\mathbb{Z}} \to \mathbb{Z}/n\mathbb{Z}$ are strong Euler characteristics
3. a *definable* Euler characteristic if all the compositions $\mathrm{Def}(M) \to \hat{\mathbb{Z}} \to \mathbb{Z}/n\mathbb{Z}$ are definable Euler characteristics.

For 2 and 3, this is an abuse of terminology.

We can repeat this discusison with the p-adics $\mathbb{Z}_p = \varprojlim_k \mathbb{Z}/p^k\mathbb{Z}$. Recall that

$$\hat{\mathbb{Z}} \cong \prod_p \mathbb{Z}_p$$

by the Chinese remainder theorem. Giving an Euler characteristic $\hat{\chi} : \mathrm{Def}(M) \to \hat{\mathbb{Z}}$ is therefore equivalent to giving an Euler characteristic $\chi_p : \mathrm{Def}(M) \to \mathbb{Z}_p$ for every p. Moreover, $\hat{\chi}$ is strong or definable if and only if every χ_p is strong or definable, respectively.

12.1.3 Statement of Results

By definition, a structure is pseudofinite if it is infinite, but elementarily equivalent to an ultraproduct of finite structures. Recall from the work of Ax [4] that a field K is pseudofinite if and only if K satsifies the following three conditions:

- K is perfect
- K is pseudo-algebraically closed: every geometrically integral variety over K has a K-point.

- $\mathrm{Gal}(K) \cong \hat{\mathbb{Z}}$, or equivalently, K has a unique field extension of degree n for each n.

Our first main result can be phrased purely in terms of pseudofinite fields.

Theorem 12.1.1.

1. *Let $K = \prod_i K_i/\mathcal{U}$ be an ultraproduct of finite fields. Then the non-standard counting functions χ_n are $\mathrm{acl}^{eq}(\emptyset)$-definable.*

2. *Every pseudofinite field admits an $\mathrm{acl}^{eq}(\emptyset)$-definable strong Euler characteristic.*

Note that the χ_n are not 0-definable, by Theorem 7.3 of [44].

We can state a more precise result, in terms of a certain type of difference field. Recall that a difference field is a pair (K, σ) where K is a field and σ is an automorphism of K.

Definition 12.1.2. *A* periodic *difference field (PDF) is a difference field (K, σ) such that every element of K has finite orbit under σ.*

PDFs are not an elementary class in the language of difference fields. However, they constitute an elementary class if we regard them as multi-sorted structures $(K_1, K_2, \ldots)$ where K_i is the fixed field of σ^i, with the following structure:

- The difference-field structure on each K_i

- The inclusion map $K_n \to K_m$ for each pair n, m with n dividing m

For notational simplicity, we will write (K, σ) when we really mean $(K_1, K_2, \ldots)$.

These objects were considered by Hrushovski in [33], and we will give an overview of their basic properties in §12.2 below.

For any q, let F_q denote $(\mathbb{F}_q^{alg}, \sigma)$, where σ is the qth power Frobenius. We will call the F_q's *Frobenius PDF*s. Frobenius PDFs are morally finite, in the sense that every definable set is finite. Consequently, ultraproducts of Frobenius PDFs admit $\mathbb{Z}/n\mathbb{Z}$-valued strong Euler characteristics χ_n.

Up to elementary equivalence, the existentially closed PDFs are exactly the non-principal ultraproducts of Frobenius PDFs (see Fact 12.2.3 below). We will write ECPDF as an abbreviation for "existentially closed PDF". ECPDFs are closely related to pseudofinite fields, as explained in Fact 12.2.4.2 below.

Here is the analog of Theorem 12.1.1 for ECPDFs

Theorem 12.1.3.

1. *There is a $\hat{\mathbb{Z}}$-valued 0-definable strong Euler characteristic χ on any ECPDF, uniformly definable across all ECPDFs*

2. *χ is the unique $\hat{\mathbb{Z}}$-valued strong Euler characteristic on (K, σ) satisfying the following property: if C is a smooth geometrically integral curve over K_1, if J is the Jacobian of C, and if q is a prime power, then*

$$\chi_q(C) = 1 - t_1 + t_2,$$

where χ_q is the composition $\mathrm{Def}(K) \xrightarrow{\chi} \hat{\mathbb{Z}} \to \mathbb{Z}/q\mathbb{Z}$, where t_1 is the trace of σ on the q-torsion $J[q]$, and where t_2 is the trace of σ on the q-torsion $\mathbb{G}_m[q]$ in the multiplicative group $\mathbb{G}_m = K^\times$.

3. *If (K, σ) is an ultraproduct of Frobenius PDFs, then χ agrees with the non-standard counting Euler characteristic $\hat{\chi} : \mathrm{Def}(K) \to \mathbb{M}$.*

4. *For any formula $\phi(x; y)$ in the language of PDFs, any $n \in \mathbb{N}$, and any $k \in \mathbb{Z}/n\mathbb{Z}$, there is a formula $\psi_{\phi,n,k}(y)$ such that for any Frobenius PDF F_q and any tuple b from F_q,*

$$|\phi(F_q; b)| \equiv k \mod n \iff F_q \models \psi_{\phi,n,k}(b)$$

It seems highly likely that 2 can be generalized to higher dimensional varieties using ℓ-adic cohomology, though we do not pursue this here.[1] The idea of using ℓ-adic cohomology originated with Hrushovski, who also suggested that these Euler characteristics are definable from the non-standard Frobenius automorphisms (see Krajíček's comments at the end of [44]).

12.2 Review of Periodic Difference Fields

In this section, we review the basic facts about periodic difference fields. All the facts here are due to Ax [4] or Hrushovski [33]. PDFs are my preferred way of establishing the basic theory of pseudofinite fields, so we will sketch some of the proofs.

Recall that we are secretly thinking of a periodic difference field (K, σ) as a multi-sorted structure $(K_1, K_2, \ldots)$ where K_n is the fixed field of σ^n, with the following structure:

- The inclusion maps $K_n \hookrightarrow K_m$ when n divides m
- The difference field structure on each K_n

Fact 12.2.1. *A PDF (K, σ) is existentially closed if and only if K is algebraically closed, $\mathrm{Gal}(K/K_1) \cong \hat{\mathbb{Z}}$, and K_1 is pseudo-algebraically closed (PAC): every geometrically irreducible variety over K_1 has a K_1-rational point.*

[1] Specifically, if V is a geometrically integral variety over K_1, and ℓ is prime to the characteristic of K, then the $\mathbb{Z}_\ell$ component of $\chi(V)$ is probably an alternating sum of the traces of the action of σ on the ℓ-adic cohomology groups $H^n(V \otimes_{K_1} K; \mathbb{Z}_\ell)$. We avoid this line of proof, because it seems unnecessarily advanced, doesn't handle the case where $\ell = \mathrm{char}\, K$, and may run into problems in positive characteristic due to lack of resolution of singularities.

Proof. First suppose (K, σ) is existentically closed. Extending σ arbitrarily to K^{alg}, we get a bigger PDF (K^{alg}, σ), which contradicts existential closedness unless $K = K^{alg}$. Let V be geometrically integral over K_1. The field $K(V)$ can be thought of as the definable closure of K and a generic point p on the variety V. Extend σ to $K(V)$ by having σ fix p. Then $K(V)$ is periodic (as a difference field), and its fixed field contains a point on V, so $V(K_1)$ must be non-empty by existential closedness. The statement that $\mathrm{Gal}(K/K_1) \cong \hat{\mathbb{Z}}$ amounts to the statement that for every $n \in \mathbb{N}$, there is an element in K whose orbit under σ has size exactly n. This follows by existential closedness and $K(t_1, \ldots, t_n)$ with σ extended to map $t_1 \mapsto t_2 \mapsto \cdots \mapsto t_n \mapsto t_1$.

Conversely, suppose $K = K^{alg}$, $\mathrm{Gal}(K/K_1) \cong \hat{\mathbb{Z}}$, and K is PAC. Let (L, σ) extend (K, σ). If $\vec{t}$ is a basis of K_n over K, then the orbit $\sigma^{\mathbb{Z}}(\vec{t})$ has size n, hence also has size n inside L. This ensures that $[L(\vec{t}) : L] \geq n$, hence $L(\vec{t}) = L_n$. It follows that $L_n = L_1 \otimes_{K_1} K_n$ for every n, and therefore $L = L_1 \otimes_{K_1} K$. As $K = K_1^{alg}$, the extension L_1/K_1 is regular, so by PAC, K_1 is relatively existentially closed in L_1. This implies that L_1 embeds into an ultrapower K_1^* of K_1. If $K^* = (K_1^*, K_2^*, \ldots)$ denotes the ultrapowers of the structure (K, σ) using the same ultrafilter, then by the argument we gave for L, $K^* = K_1^* \otimes_{K_1} K$. So, the embedding $L_1 \hookrightarrow K_1^*$ induces an embedding

$$L = L_1 \otimes_{K_1} K \hookrightarrow K_1^* \otimes_{K_1} K = K^*$$

of difference fields. Because L embeds into an ultrapower of K, it follows that K is relatively existentially closed in L. □

The conditions in Fact 12.2.1 are first order (in spite of appearances to the contrary), so the class of PDFs has a model companion.

If (K, σ) is a PDF, let $\mathrm{Abs}(K)$ denote the "absolute numbers," the algebraic closure within K of the prime field. Note that $\mathrm{Abs}(K)$ is a sub-PDF of K.

Fact 12.2.2. *Two ECPDFs K_1 and K_2 are elementarily equivalent if and only if $\mathrm{Abs}(K_1) \cong \mathrm{Abs}(K_2)$. More generally, if K_1 and K_2 are two ECPDFs, a partial map f from $S \subseteq K_1$ to K_2 is a partial elementary map if and only if it can be extended to an isomorphism of PDFs from S^{alg} to $f(S)^{alg}$.*

Proof. The proof is the same as in ACFA ([7], Theorem 1.3), and boils down to the fact that we can amalgamate PDFs over algebraically closed PDFs, by taking tensor products. □

Recall the Frobenius PDFs $F_q = (\mathbb{F}_q^{alg}, (\cdot)^q)$.

Fact 12.2.3. *A PDF (K, σ) is existentially closed if and only if it is a non-principal ultraproduct of Frobenius PDFs.*

Proof. Suppose K is a non-principal ultraproduct of Frobenius PDFs. Frobenius PDFs satisfy the axioms that K is algebraically closed and $\mathrm{Gal}(K/K_1) \cong \hat{\mathbb{Z}}$, so we only need to check that K_1 is PAC. This follows in the usual way from the Weil conjectures for curves; see §6 of [4].

Conversely, suppose K is an ECPDF. If K has characteristic 0, we can use the Chebotarev density theorem (as in [4]) to find a sequence of primes $p_1, p_2, \ldots$ such that,in the limit, the non-standard Frobenius belongs to the same conjugacy class of $\mathrm{Gal}(\mathbb{Q}^{alg}/\mathbb{Q})$ as $\sigma|\,\mathrm{Abs}(K)$. By the previous paragraph, the ultraproduct of the F_{p_i}'s will be an ECPDF, and it will be elementarily equivalent to K by Fact 12.2.2. The positive characteristic case is similar, but significantly easier. □

Putting things together, we see that

Fact 12.2.4.

1. *A field K is pseudofinite if and only if it has absolute Galois group $\hat{\mathbb{Z}}$, it is perfect, and it is PAC*
2. *If K is a pseudofinite field and σ is any topological generator of $\mathrm{Gal}(K)$, then (K^{alg}, σ) is an ECPDF, and all ECPDFs arise this way.*
3. *The ECPDFs are exactly the non-principal ultraproducts of Frobenius PDFs, up to elementary equivalence*
4. *Let K be a pseudofinite field. Then (K^{alg}, σ) is K-interpretable in K, where by (K^{alg}, σ) we really mean $(K_1, K_2, \ldots)$ as above. In particular, every (K^{alg}, σ)-definable subset of K is also K-definable in the pure field language on K.*

Note that 4 holds because each K_n is just the unique degree n extension of $K = K_1$, which we can identify with K^n after choosing an irreducible polynomial for a generator.

We will use the following additional fact about pseudofinite fields and ECPDFs:

Fact 12.2.5. *Pseudofinite fields and ECPDFs are supersimple of finite SU rank.*

Finite rank supersimplicity of pseudofinite fields is proven in [7], and follows in ECPDF by Fact 12.2.4.4.

12.3 Some technical lemmas

Because PDFs can be amalgamated over algebraically closed PDFs using tensor products, the field theoretic and model theoretic notions of algebraic closure agree in ECPDFs, as well as in pseudofinite fields. Now pseudofinite fields have SU-rank 1, so $SU(a/S) = \mathrm{tr.deg}(a/S)$ when $|a| = 1$. By Lascar inequalities, we indeed see $SU(a/S) = \mathrm{tr.deg}(a/S)$ when a is a tuple of any length.

Lemma 12.3.1. *Let K be a pseudofinite field, and $X \subseteq K^n$ be quantifier-free definable, of SU-rank 1. Then there exists a collection $\{C_1, \ldots, C_m\}$ of smooth projective geometrically irreducible curves over K, and a definable bijection between a cofinite subset of X and a cofinite subset of $\coprod_{i=1}^m C_i(K)$.*

Proof. Let K^* be a sufficiently saturated elementary extension of K. There is no tuple $a \in X(K^*)$ having transcendence degree 2 or higher, over K. By compactness, X must be contained in a finite union of curves over K. Therefore X contains only finitely many non-algebraic quantifier free types $p_1, \ldots, p_m$. Note that each p_i must be wholly contained in X, as X is quantifier-free definable.

Choose elements $a_1, \ldots, a_m \in X(K^*)$ representing these types. In an ambient model of ACF, $\mathrm{tp}_{ACF}(a_i/K)$ is stationary, as K is relatively algebraically closed in K^*. Therefore there is an absolutely irreducible Zariski-closed curve $D_i \subseteq \mathbb{A}^n$, defined over K, such that p_i is the generic type of D_i. As X is quantifier-free definable, $D_i \setminus X$ must be finite. Moreover, $\bigcup_{i=1}^m D_i$ contains all the quantifier-free types in X of positive rank, so $X \setminus \bigcup_{i=1}^m D_i$ is also finite.

Consequently, X and $\bigcup_{i=1}^m D_i$ differ by a finite set. Let C_i be a smooth projective model of D_i for each i. We end up with the following chain of sets, each of which is in definable bijection with the next set (up to finitely many exceptional points):

$$X(K) \approx \bigcup_{i=1}^m D_i(K) \approx \coprod_{i=1}^m D_i(K) \approx \coprod_{i=1}^m C_i(K)$$

Consequently, we get a definable bijection between $X(K)$ and $\coprod_{i=1}^m C_i(K)$, except at finitely many points. □

Lemma 12.3.2. *Let (K, σ) be an ECPDF and let Y be a quantifier-free definable set. Then there is some definable set $X \subseteq (K_1)^m$ in definable bijection with Y, such that X is quantifier-free definable in K_1 in the pure field language (with parameters from K_1).*

Proof. Choose d large enough that Y lives in some power of K_d, and K_d contains all the parameters used to define Y. Then all the variables and terms in the definition of Y live within K_d. Choosing a basis of K_d over K_1, we can identify K_d with $(K_1)^d$. Under this identification, the addition, multiplication, and σ maps on K_d are all given by polynomial maps on K_1, with coefficients from K_1. Consequently, Y becomes identified with a definable subset of a power of K_1, whose definition is quantifier-free using only constants from K_1 and the langauge of rings. □

Lemma 12.3.3. *Let (K, σ) be an ECPDF and let X be a definable set. Then there exists a quantifier-free definable set Y and a definable surjection $Y \to X$ with finite fibers.*

Proof. Let $\mathcal{P}$ be the collection of all sets of the form $\pi(Y(K))$ where Y is a quantifier-free definable set and π is a coordinate projection, such that the map $Y(K) \to \pi(Y(K))$ has finite fibers.

The collection $\mathcal{P}$ is closed under finite unions and finite intersections.

Claim 12.3.4. *Let K^* be a sufficiently saturated elementary extension of K. Let a and b be two tuples from the same product of sorts. Suppose that $a \in P \implies b \in P$ for every $P \in \mathcal{P}$. Then $\mathrm{tp}(a/K) = \mathrm{tp}(b/K)$.*

Proof. We claim that we can find an embedding f of PDFs from $K(a)^{alg}$ to K^* mapping a to b and fixing K pointwise. This follows by compactness, unless the diagram of $K(a)^{alg}$ is inconsistent with b. In this case there would be a finite tuple c from $K(a)^{alg}$, and some quantifier-free formula $\phi(z;x)$ over K such that $\phi(c;a)$ holds, but $\phi(c';b)$ does not hold for any c', i.e.,

$$K^* \models \forall z \neg\phi(z;b) \tag{12.1}$$

We may as well add to ϕ the statement witnessing that c is algebraic over a. Let Y be the set

$$\{(z;y)|K^* \models \phi(z;y)\}$$

and let π be the projection sending $(z;y)$ to y. Then Y is quantifier-free definable over K, and the fibers of $Y \to \pi(Y)$ are finite, so $\pi(Y) \in \mathcal{P}$. The existence of c ensures that $a \in \pi(Y)$. The assumption connecting a and b then implies that $b \in \pi(Y)$, contradicting (12.1).

So we can find an embedding of $K(a)^{alg}$ into K^* mapping a to b and fixing K pointwise. The image must of course be $K(b)^{alg}$. Now by Fact 12.2.2, $\mathrm{tp}(a/K) = \mathrm{tp}(b/K)$. □

Now if p is any type over K consistent with X, let Σ_p be the collection of all sets in $\mathcal{P}$ consistent with p. By the Claim, $\Sigma_p \vdash p \vdash X$. By compactness, and the fact that $\mathcal{P}$ is closed under finite intersections, some $P \in \Sigma_p$ implies X.

Consequently, we can write X as a union of sets in $\mathcal{P}$. By compactness of the type space, X is a finite union of sets in $\mathcal{P}$. As $\mathcal{P}$ is closed under finite unions, X is in $\mathcal{P}$. □

12.4 Uniqueness and definability

We will use the following form of Beth implicit definability:

Theorem 12.4.1. *Let $L^+ \supseteq L^-$ be languages. Let T^- be an L^- theory and T^+ be an L^+ theory extending T^-. Let $\phi(x)$ be an L^+ formula. Suppose that whenever $N \models T^-$, and M_1^+ and M_2^+ are two expansions of N to a model of T^+, that $\phi(M_1^+) = \phi(M_2^+)$. Then there is an L^--formula $\psi(x)$ such that $T^+ \vdash \phi \leftrightarrow \psi$.*

This is Theorem 6.6.4 in [29]. We would like to apply this to the conditions in part 2 of Theorem 12.1.3. The language L^- will be the language of PDFs discussed in the previous section. The language L^- will be an expansion of L^- by some new predicates $P_{\phi,n,k}(y)$ for every formula $\phi(x;y)$, every $n \in \mathbb{N}$, and every $k \in \mathbb{Z}/n\mathbb{Z}$. The theory T^- will be ECPDF, and T^+ will be T^- with the following additions:

1. The statement that for every ϕ, n, and b, there is a unique k such that $P_{\phi,n,k}(b)$ holds. Let $f_{\phi,n}(b)$ denote this unique k.

2. The statement that if $\phi(K;b) = \phi'(K;b')$, then $f_{\phi,n}(b) = f_{\phi',n}(b')$, for each n. Let $\chi_n(X)$ denote $f_{\phi,n}(b)$ for any representation of X as $\phi(K;b)$.

3. The statement that χ_n is a strong Euler characteristic for each n.

4. The statement that the diagrams

$$\begin{array}{ccc} \mathrm{Def}(M) & \xrightarrow{\chi_m} & \mathbb{Z}/m\mathbb{Z} \\ & \searrow^{\chi_n} & \downarrow \\ & & \mathbb{Z}/n\mathbb{Z} \end{array}$$

   commute when n divides m.

5. The statement that if q is a prime-power, if C is a smooth projective genus g curve over K_1, if J is a projective group variety over K_1, and if J is birationally equivalent (over K_1) to $\mathrm{Sym}^g C$, then $\chi_q(C)$ is given by the formula $1 - t_1 + t_2$ where t_1 is the trace of σ on the q-torsion on $J(K)$ and t_2 is the trace of the σ on the q-torsion of $\mathbb{G}_m(K)$.

Points 1-4 encode the statement that χ is a $\hat{\mathbb{Z}}$-valued strong Euler characteristic. Note that the conditions in 5 correctly ensure that J is the Jacobian of C, because by §V.1 of [65] there is a birational equivalence between the Jacobian of C and $\mathrm{Sym}^g C$, and by [52] I.3.7 any birational map between two projective group varieties extends to an isomorphism.[2]

Our first goal is to show that on any ECPDF (K, σ), there is at most one $\hat{\mathbb{Z}}$-valued Euler characteristic satisfying the axioms of T^+.

Lemma 12.4.2. *Fix $n \in \mathbb{N}$. Let (K, σ) be an ECPDF, admitting two expansions to a model of T^+. Let χ and χ' be the corresponding $\hat{\mathbb{Z}}$-valued strong Euler characteristics. Then $\chi(X) = \chi'(X)$ for all definable $X \subseteq (K_1)^n$.*

Proof. We proceed by induction on n. First suppose $n = 1$. Say that a set is *good* if $\chi(X) = \chi'(X)$. Finite sets are good. If X is in definable bijection with Y and X is good, then so is Y. A disjoint union of two good sets is good. If S is a cofinite subset of X, then S is good if and only if X is good. Consequently, if a cofinite subset of X is in definable bijection with a cofinite subset of Y, then X is good if and only if Y is good.

If C is a projective smooth curve over K_1, then $C(K_1)$ is good, because of the explicit definition involving Jacobians (Axiom 5 of T^+). Any disjoint union of sets of this form is also good. By Lemmas 12.3.1 and 12.3.2, this establishes goodness for any quantifier-free definable set of rank 1.

Now let X be a definable subset of $(K_1)^1$. By Lemma 12.3.3, there is a quantifier-free definable set Y and a definable surjection $Y \twoheadrightarrow X$ given by a coordinate projection. Let

[2] We are sweeping a lot of computational algebraic geometry and "definability" results under the rug, though Chapter 10 helps considerably. Smoothness of C can be checked by calculating the rank of the matrix of partial derivatives of the functions defining the curve. We can witness that C has some genus g by exhibiting a section of the cotangent bundle, and counting its poles and zeros to verify that the cotangent bundle has degree $2g - 2$–or by exhibiting a birational map to $\mathbb{P}^1$ for genus 0.

Y_n denote the n-fold fiber product $Y \times_X Y \times_X \cdots \times_X Y$. So an element of Y_n is a tuple $(y_1, \ldots, y_n, x)$ such that $(y_i, x) \in Y$ for each Y. Because Y is quantifier-free definable, so is each Y_n.

Each of the maps $Y_n \to X$ has finite fibers. By Lascar inequalities, it follows that $SU(Y_N) \leq SU(X) \leq SU(K_1) = 1$. Therefore, each Y_n is good.

Let m be a bound on the size of the fibers of $Y \to X$. For $1 \leq k \leq m$, let X_k denote the set of $a \in X$ such that $f^{-1}(a)$ has size m. Let α_k and β_k denote $\chi(X_k)$ and $\chi'(X_k)$.

Because χ and χ' are strong Euler characteristics,

$$\chi(Y_n) = \sum_{k=1}^{m} \alpha_k k^n$$
$$\chi'(Y_n) = \sum_{k=1}^{m} \beta_k k^n$$

for all n. As the Y_n's are good,

$$\sum_{k=1}^{m} \alpha_k k^n = \sum_{k=1}^{m} \beta_k k^n$$

for $n = 1, \ldots, m$. By invertibility of the Vandermonde matrix $\langle k^n \rangle_{1 \leq k \leq m,\ 1 \leq n \leq m}$, and the fact that $\hat{\mathbb{Z}}$ has no $\mathbb{Z}$-torsion, it follows that $\alpha_k = \beta_k$ for all k. Consequently,

$$\chi(X) = \sum_{k=1}^{m} \alpha_k = \sum_{k=1}^{m} \beta_k = \chi'(X). \tag{12.2}$$

Therefore X is good.

This completes the base case. Now suppose $n > 1$. Let $\phi(x; y; z)$ be a formula with $(x; y)$ an n-tuple in K_1, and $|y| = 1$ (so $|x| = n - 1$). For any m and any $k \in \mathbb{Z}/m\mathbb{Z}$, let $P_{m,k}(y; z)$ denote the predicate in L^+ indicating

$$P_{m,k}(b; c) \iff \chi_m(\phi(K; b; c)) = k + m\mathbb{Z}$$

The sets $\phi(K; b; c)$ are in $(K_1)^{n-1}$, so by induction $P_{m,k}$ agrees in two models of T^+ sharing the same underlying model of T^-. By Beth implicit definability (Theorem 12.4.1), it follows that $P_{m,k}$ is definable across all models of T^+ by an L^--formula $\psi_{m,k}(y; z)$.

Now let (K, σ) be a specific ECPDF, let χ and χ' be two strong Euler characteristics satisfying T^+, and let X be a definable subset of $(K_1)^n$ of the form $\phi(K; K; c)$. We will show that $\chi(X) = \chi'(X)$. If not, then $\chi_m(X) \neq \chi'_m(X)$ for some m. For $k \in \mathbb{Z}/m\mathbb{Z}$, let $Y_k = \psi_{m,k}(K; c)$. Thus, for any $b \in K_1$,

$$\chi_m(\phi(K; b; c)) \equiv k \iff b \in Y_k \iff \chi'_m(\phi(K; b; c)) \equiv k \tag{12.3}$$

Let X_k be the $(x; y) \in X$ such that $y \in Y_k$. Equation (12.3) says that every fiber of $X_k \to Y_k$ has Euler characteristic k with respect to both χ_m and χ'_m. By definition of strong Euler

characteristic, it follows that

$$\chi_m(X) = \sum_{k\in\mathbb{Z}/m\mathbb{Z}} \chi_m(X_k) = \sum_{k\in\mathbb{Z}/m\mathbb{Z}} k\chi_m(Y_k)$$
$$\chi'_m(X) = \sum_{k\in\mathbb{Z}/m\mathbb{Z}} \chi'_m(X_k) = \sum_{k\in\mathbb{Z}/m\mathbb{Z}} k\chi'_m(Y_k)$$

As each Y_k is a definable subset of K_1, it follows by induction that $\chi_m(Y_k) = \chi'_m(Y_k)$, so the right hand sides are equal. Therefore, $\chi_m(X) = \chi'_m(X)$, completing the inductive step and the proof. □

Proposition 12.4.3. *Let (K, σ) be an ECPDF. There is at most one expansion of (K, σ) to a model of T^+.*

Proof. Let χ and χ' be two strong Euler characteristics both satisfying the axioms of T^+. Let X be an arbitrary definable set; we will show $\chi(X) = \chi'(X)$. Each sort K_n in (K, σ) is in definable bijection with $(K_1)^n$, so we may find a definable bijection between X and a definable subset $Y \subseteq (K_1)^n$ for some n. Then

$$\chi(X) = \chi(Y) = \chi'(Y) = \chi'(X) \tag{12.4}$$

where the middle equivalence is Lemma 12.4.2. □

By 12.4.1,

Corollary 12.4.4. *Let (K, σ) be an ECPDF. If χ is a strong Euler characteristic on (K, σ) satisfying the axioms of T^+, then χ is a 0-definable Euler characteristic.*

12.5 Some algebraic geometry

In this section, we review some algebraic geometry we will need, culminating in Theorems 12.5.8 and 12.5.10. These facts are well-known, but I had trouble finding direct references, so I explain some of the proofs.

If $f : A \to A$ is an isogeny on an abelian variety A, then f has a well-defined degree $\deg(f)$, which can be written as a product

$$\deg(f) = \deg_s(f) \cdot \deg_i(f)$$

where $\deg_s(f)$ and $\deg_i(f)$ denote the separable and inseparable degrees of f. Moreover, $\deg_s(f)$ is the size of the set theoretic kernel of f. These facts are proven in ([55] §6 Application 3). Since isogenies are finite and flat, $\deg(f)$ is also the length of the scheme-theoretic kernel of f. Thus $\deg_i(f)$ is the length of the connected component $(\ker f)_0$ of the scheme-theoretic kernel of f.

Lemma 12.5.1. *Let G be a finite connected group scheme of length n over $\mathbb{F}_q$. If $n < q$ then the qth-power Frobenius annihilates G.*

Proof. We can write G as $\operatorname{Spec} A$ for some local Artinian n-dimensional $\mathbb{F}_q$-algebra A. Let $\mathfrak{m}$ be the maximal ideal of A; then $\mathfrak{m}$ is also the nilradical, so every $x \in \mathfrak{m}$ is nilpotent. In fact, $x^q = 0$ for all $x \in \mathfrak{m}$. Otherwise, the descending chain of ideals

$$A \supsetneq (x) \supsetneq (x^2) \supsetneq \cdots \supsetneq (x^q) \supsetneq (0)$$

would contradict length $\leq q$.

So the qth power homomorphism on A annihilates $\mathfrak{m}$, and must therefore be

$$A \twoheadrightarrow A/\mathfrak{m} \xrightarrow{\sim} \mathbb{F}_q \hookrightarrow A$$

Thus the qth power Frobenius on G must be $G \to \operatorname{Spec} \mathbb{F}_q \to G$, the zero map. □

Recall the Tate modules $T_\ell(A)$ and $T_p(A)$ of an Abelian variety (defined in §18 of [55]). If g denotes the dimension of A, then $T_\ell(A)$ is a free $\mathbb{Z}_\ell$-module of rank $2g$, and $T_p(A)$ is a free $\mathbb{Z}_p$-module of rank r, for some $r \leq g$ known as the p-rank of A. If ϕ is an endomorphism of an Abelian variety, then we can talk about the determinants and traces of the induced maps $T_\ell(\phi)$ and $T_p(\phi)$.

The following fact is noted in the proof of Theorem 19.4 in [55]:

Fact 12.5.2. *For any ℓ (possibly $\ell = p$),*

$$v_\ell(\det T_\ell(\phi)) = v_\ell(|\ker \phi|) = v_\ell(\deg_s(\phi))$$

Moreover, Theorem 19.4 in [55] yields the following:

Fact 12.5.3. *For any endomorphism ϕ, and $\ell \neq p$,*

$$\deg(\phi) = \det T_\ell(\phi)$$

and therefore, if $\alpha_1, \ldots, \alpha_{2g}$ denote the eigenvalues of $T_\ell(\phi)$, then

1. *The α_i's are algebraic integers that do not depend on ℓ*
2. *For any polynomial $P(X)$, we have*

$$\deg(P(\phi)) = \prod_{i=1}^{2g} P(\alpha_i)$$

The α_i's are called the *characteristic roots* of the endomorphism ϕ.

Now we need two technical lemmas

Lemma 12.5.4. *Suppose k is sufficiently large relative to $i, p, and g$. Let A be a g-dimensional variety over $\mathbb{F}_q$, for $q = p^k$. Let r be the p-rank of A. Let ϕ_q denote the qth power Frobenius endomorphism of A. Then $\deg(\phi_q - p^i)$ is divisible by $p^{i(2g-r)}$, where r is the p-rank of A.*

Proof. By Fact 12.5.3, $\deg(p^i) = p^{2gi}$ because $T_\ell(A)$ has dimension $2g$. Let K denote the scheme-theoretic p^i-torsion in A, i.e., the kernel of multiplication by p^i. Then K is a finite group scheme of length $\deg(p^i) = p^{2gi}$. By definition of p-rank, $K(\mathbb{F}_q^{alg}) \cong (\mathbb{Z}/p^i\mathbb{Z})^r$, so

$$|K(\mathbb{F}_q^{alg})| = p^{ir}$$

and therefore the connected component K_0 has length $p^{2gi}/p^{ir} = p^{i(2g-r)}$. As q is sufficiently big relative to i and $2g$, $q > p^{i(2g-r)}$, and therefore ϕ_q annihilates K_0 by Lemma 12.5.1.

Let K' denote the kernel of $\phi_q - p^i$. Then $K_0 \subseteq K'$, so the length of K_0 (which is $p^{i(2g-r)}$) divides the length of K', which is $\deg(\phi_q - p^i)$. □

Lemma 12.5.5. *There is some function $h(d, d', p, s)$ such that if $Q(X)$ is a monic polynomial of degree d such that $v_p(Q(p^i)) \geq v_p(p^{id'})$ for all $i \leq h(d, d', p, s)$, then at least d' roots of $Q(X)$ (counting multiplicities) have p-adic valuation at least s.*

Proof. Identify the space of $Q(X)$'s with the Stone space $\mathbb{Z}_p^d$. The set of $Q(X)$ such that $v_p(Q(p^i)) \geq v_p(p^{id'})$ is closed (even clopen). The set of $Q(X)$ such that at least d' roots of $Q(X)$ have p-adic valuation at least s is also clopen, using Newton polygons. By compactness, it suffices to show that if $v_p(Q(p^i)) \geq v_p(p^{id'})$ for all i, then at least d' roots of $Q(X)$ are 0, hence have p-adic valuation at least s. Indeed, if we let the roots be $\rho_1, \ldots, \rho_d$, then

$$\lim_{\epsilon \to 0} \frac{v_p(Q(\epsilon))}{v_p(\epsilon)}$$

counts exactly how many of the ρ_i are 0. The assumption that $v_p(Q(p^i)) \geq d' \cdot v_p(p^i)$ implies this limit is at least d'. □

Corollary 12.5.6. *Let A be a g-dimensional abelian variety over $\mathbb{F}_q$, for $q = p^k$. Let ϕ_q denote the qth power Frobenius on A.*

If q is sufficiently large relative to p, g, j, then, counting multiplicities, $2g - r$ of the characteristic roots of ϕ_q have p-adic valuation greater than j.

We also make the following observation:

Lemma 12.5.7. *Let A be an abelian variety over $\mathbb{F}_q$ for $q = p^k$. Let $\beta_1, \ldots, \beta_r$ be the eigenvalues of $T_p(\phi_q)$, for ϕ_q the qth power Frobenius on A. Then $\{\beta_i\}$ is a submultiset of the characteristic roots $\{\alpha_i\}$ of ϕ_q.*

Proof. By Lemma 12.5.2,

$$v_p\left(\prod_{i=1}^{r} P(\beta_i)\right) = v_p(\det T_p(P(\phi_q))) = v_p(\deg_s(P(\phi_q))) \leq v_p(\deg(P(\phi_q))) = v_p\left(\prod_{i=1}^{2g} P(\alpha_i)\right)$$

for any polynomial $P(X) \in \mathbb{Z}[X]$. By continuity, the outer inequality

$$v_p\left(\prod_{i=1}^{r} P(\beta_i)\right) \leq v_p\left(\prod_{i=1}^{r} P(\beta_i)\right) \tag{12.5}$$

holds for all $P(X) \in \mathbb{Z}_p[X]$. For each β_i, if we let $Q_i(X)$ denote the minimal polynomial over $\mathbb{Z}_p$, then applying (12.5) with $P(X) = Q_i(X + \epsilon)$ and sending $\epsilon \to 0$ implies that the multiplicity of β_i in the α's is greater than or equal to the multiplicity of $\beta_i$in the β's. □

Here is a summary statement about abelian varieties:

Theorem 12.5.8. *Let A be a g-dimensional abelian variety over $\mathbb{F}_q$ for $q = p^k$, and let ϕ_q denote the qth power Frobenius map. Then*

$$|A(\mathbb{F}_{q^n})| = \prod_{i=1}^{2g}(1 - \alpha_i^n)$$

where $\alpha_1, \ldots, \alpha_{2g}$ are algebraic integers satisfying the following properties:

1. *For $\ell \neq p$, the α_i are the eigenvalues of $T_\ell(\phi_q)$*
2. *If r is the p-rank of A, then r of the α_i are the eigenvalues of $T_p(\phi_q)$ (hence have p-adic valuation 0, as $T_p(\phi_q)$ is an automorphism). If q is sufficiently large relative to g, p, ϵ, then the other $2g - r$ of the α_i have p-adic valuation less than ϵ.*
3. *Choosing some embedding into the complex numbers, each α_i has magnitude $|\alpha_i| = \sqrt{q}$.*

Proof. This is Theorem III.4 in [55], except for 2, which is Corollary 3 and Lemma 12.5.7. □

Next, we would like to tie things to curves. Recall the Riemann hypothesis for curves:

Fact 12.5.9. *Let C be a genus g curve over a finite field $\mathbb{F}_q$. Then there exist $\alpha_1, \ldots, \alpha_{2g} \in \mathbb{C}$, each of magnitude $\sqrt{q}$, such that for every n,*

$$|C(\mathbb{F}_{q^n})| = 1 - \left(\sum_{i=1}^{2g} \alpha^n\right) + q^n$$

Moreover, if J is the Jacobian of C, then $J(\mathbb{F}_q)$ has size

$$|J(\mathbb{F}_q)| = \prod_{i=1}^{2g}(1 - \alpha_i) \tag{12.6}$$

Proof. This is well-known; see §3.5 of [22] and Exercise C.5.7 in [25]. For the "moreover" clause, recall the following formula from §3.5 of [22]:

$$Z(t) = \Phi(t) + \frac{h}{q^g}q - 1 \cdot \frac{t^{2g-1}}{1 - qt} - \frac{h}{q - 1} \cdot \frac{t^{2g-1}}{1 - t}$$

where $\Phi(t)$ is some polynomial, h is the class number of C, (i.e., $|J(\mathbb{F}_q)|$), and where

$$Z(t) = \frac{\prod_{i=1}^{2g}(1-\alpha_i t)}{(1-t)(1-qt)}$$

. Clearing denominators, we see that

$$\prod_{i=1}^{2g}(1-\alpha_i t) = \Phi(t)(1-qt)(1-t) + \frac{hq^g}{q-1}t^{2g-1}(1-t) - \frac{h}{q-1}\cdot t^{2g-1}(1-qt)$$

Substituting in $t = 1$ gives (12.6). □

Using this, we verify the following well-known result:

Theorem 12.5.10. *Let C be a curve over a finite field $\mathbb{F}_q$, and let J be its Jacobian. Then*

$$|C(\mathbb{F}_q)| = 1 - \left(\sum_{i=1}^{2g}\alpha_i\right) + q$$

where the α_i are the characteristic roots of the qth power Frobenius on J (the numbers from Theorem 12.5.8).

Proof. For each m, we can consider C as a curve over $\mathbb{F}_{q^m}$. By Fact 12.5.9, there are $\alpha_{1,m}, \ldots, \alpha_{2g,m}$ such that

$$|C(\mathbb{F}_{q^{nm}})| = 1 - \left(\sum_{i=1}^{2g}\alpha_{i,m}^n\right) + q^{nm} \tag{12.7}$$

for all n, and

$$|J(\mathbb{F}_{q^m})| = \prod_{i=1}^{2g}(1-\alpha_{i,m}) \tag{12.8}$$

From (12.7) one can easily verify that, after applying permutations for each m, $\alpha_{i,m} = \alpha_i^m$ for each m, for some fixed α_i's. Plugging this into (12.8) then yields

$$|J(\mathbb{F}_{q^n})| = \prod_{i=1}^{2g}(1-\alpha_i^n)$$

for all i.

Now let $\beta_1, \ldots, \beta_{2g}$ be the characteristic roots of the qth power Frobenius on J. By Theorem 12.5.8,

$$\prod_{i=1}^{2g}(1-\alpha_i^n) = |J(\mathbb{F}_{q^n})| = \prod_{i=1}^{2g}(1-\beta_i^n)$$

This implies that the two multisets

$$\{\prod_{i\in S} \alpha_i : |S| \text{ odd}\}$$
$$\{\prod_{i\in S} \beta_i : |S| \text{ odd}\}$$

are equal. Filtering out the elements with complex absolute value $\sqrt{q}$, we see that the α's and β's must be the same things. □

12.6 Verifying the implicit definition

Next, we show that on any non-principal ultraproduct K of Frobenius PDFs, the nonstandard counting Euler characteristic $\chi : \mathrm{Def}(K) \to \hat{\mathbb{Z}}$ does in fact satisfy the axioms of T^+.

Proposition 12.6.1. *Let (K, σ) be a non-principal ultraproduct of Frobenius PDFs. Then the nonstandard counting function $\chi : \mathrm{Def}(K) \to \hat{\mathbb{Z}}$ satisfies the axioms of T^+.*

Proof. We need to show that each instance of the axioms holds in all but finitely many Frobenius PDFs. Axioms 1 through 4 hold in any finite structure, hence in any Frobenius PDF. This leaves Axiom 5, which says that $C(\mathbb{F}_q)$ is congruent mod ℓ^k to $1 - t_1 + t_2$, where t_1 is the trace of the action of ϕ_q on the ℓ^k-torsion in the Jacobian J, and t_2 is the trace of the action of ϕ_q on the ℓ^k torsion in $\mathbb{G}_m$. Note that t_1 and t_2 are just the mod ℓ^k reductions of the traces of the actions on the full ℓ-adic Tate modules of J and $\mathbb{G}_m$. When ℓ is not the characteristic, we know that $C(\mathbb{F}_q)$ is exactly $1 - t_1 + t_2$ with no caveats:

- The trace of the action on $T_\ell J$ is the sum of the α_i's from Theorem 12.5.8.
- The qth power Frobenius acts as multiplication by q on $\mathbb{G}_m$, hence on $T_\ell \mathbb{G}_m$, so the sole eigenvalue is q. Thus t_2 is q mod ℓ^k.
- By Theorem 12.5.10,

$$|C(\mathbb{F}_q)| = 1 - \alpha_1 - \cdots - \alpha_{2g} + q = 1 - t_1 + t_2$$

So any counterexamples will have to happen in the bad characteristic $\ell = p$. In this case, Theorem 12.5.8 tells us that we can order the α_i's in such a way that $\alpha_1, \ldots, \alpha_r$ are the eigenvalues of $T_p(\phi_q)$, and $\alpha_{r+1}, \ldots, \alpha_{2g}$ are p-adically small (if q is sufficiently large). So, for fixed p^k and sufficiently large q,

$$|C(\mathbb{F}_q)| = 1 - \alpha_1 - \cdots - \alpha_{2g} + q \equiv 1 - \alpha_1 - \cdots - \alpha_r = 1 - t_1 = 1 - t_1 + t_2$$

mod p^k. Note that $t_2 = 0$ because $T_p(\mathbb{G}_m)$ vanishes in characteristic p. □

12.7 Completing the proofs

We now complete the proofs of Theorems 12.1.1 and 12.1.3.

Proof (of Theorems 12.1.1 and 12.1.3). By Proposition 12.6.1, any ultraproduct of Frobenius PDFs admits a strong Euler characteristic χ satisfying T^+. By Corollary 12.4.4, this strong Euler characteristic is 0-definable. An arbitary ECPDF (K, σ) admits an elementary embedding into such an ultraproduct (K', σ). It is a general fact that if $N \preceq M$, any 0-definable Euler characteristic on M induces a 0-definable Euler characteristic on N, having the same first order properties. Applying this to $(K, \sigma) \preceq (K', \sigma)$, we see that (K, σ) has a 0-definable Euler characteristic that is also strong, and in fact satisfies the conditions of T^+.

Consequently, every ECPDF can be expanded to a model of T^+. By Beth implicit definability, T^+ is a definitional expansion of ECPDF. This establishes part 1 of Theorem 12.1.3, and the existence part of 12.1.3.2. The uniqueness part of 12.1.3.2 was established in Proposition 12.4.3. Part 3 of Theorem 12.1.3 follows by Proposition 12.6.1. To prove part 4 of 12.1.3, take the uniform definitions of χ for ECPDFs from part 1 and apply them to Frobenius PDFs. By part 3, each definition must agree with the standard counting Euler characteristic, except on finitely many Frobenius PDFs. The definitions can be modified to correctly handle the finite set of exceptions.

Finally, we turn to Theorem 12.1.1. For part 1, given an ultraproduct $K = \prod_i \mathbb{F}_{q_i}/\mathcal{U}$, let $(L, \sigma) = \prod_i F_{q_i}/\mathcal{U}$ be the corresponding ultraproduct of Frobenius PDFs. Then $K \cong L_1$. The non-standard counting function on L is definable, so its restriction to K is also definable. Because L is interpretable within K (by Fact 12.2.4.4), the non-standard counting function on K itself is definable.

Finally, for part 2 of Theorem 12.1.1, given a pseudofinite field K, choose some topological generator σ of $\mathrm{Gal}(K)$. Then (K^{alg}, σ) is bi-interpretable (after naming parameters) with K, and the larger structure (K^{alg}, σ) admits a strong definable Euler characteristic, which induces a strong definable Euler characteristic on K. This Euler characteristic is $\mathrm{acl}^{eq}(0)$-definable, because it has a bounded number of conjugates under $\mathrm{Aut}(K)$, owing to the bounded set of choices for σ. □

12.8 Conclusion

We have relied heavily on algebraic geometry and number theory to prove a relatively simple model-theoretic fact. One could dream of turning around the process and applying the model-theoretic result to yield new insights in number theory. There are sources of pseudofinite fields other than ultraproducts of finite fields. For example, $(\mathbb{Q}, \sigma)$ is an ECPDF, for random $\sigma \in \mathrm{Gal}(\mathbb{Q})$ (with probability 1) by results in [22]. Perhaps one could prove non-trivial facts by reasoning about non-standard sizes of definable sets in these structures.

Unfortunately, from a number-theoretic point of view we have probably done nothing interesting. The nonstandard sizes we have defined are probably a simple artifact of etale

cohomology, which is already well understood. The strong Euler characteristic condition—the ability to fiber a space E over a base B and calculate the "size" of E by "integrating" the "sizes" of the fibers over B—is probably a disguised way of saying that the cohomology of E can be calculated in terms of the cohomology of B with coefficients in the cohomology of the fibers (the Leray spectral sequence).

One tool which might be new on the model-theoretic side is elimination of imaginaries, which holds in ECPDFs by work of Hrushovski [33]. One could imagine proving non-trivial facts about etale cohomology by defining interesting equivalence relations on pseudo-finite fields, and doing combinatorics on the quotients. (For example, this approach shows that if A is a definable abelian group in a pseudofinite field, and A has n-torsion, then $\chi(A)$ is divisible by n.)

One interesting thing to note is that the strong Euler characteristic condition is related to p-adic integration, in a manner we now explain. Suppose we have a definable strong $\mathbb{Z}_p$-valued Euler characteristic χ in some structure. Then χ induces a p-adic measure on the type space of any definable set Y. Moreover, if $f : X \twoheadrightarrow Y$ is a surjection, then the map $y \mapsto \chi(f^{-1}(y))$ is continuous on this type space, and its p-adic integral with respect to the measure is exactly $\chi(X)$. In this way, model theory could conceivably offer a new pseudofinite perspective on p-adic integration and p-adic L-functions. I lack the expertise to pursue this connection further, however.

Bibliography

[1] Hans Adler. *Strong theories, burden, and weight*. Draft. 2007.

[2] Hans Adler. "Theories controlled by formulas of Vapnik-Chervonenkis codimension 1". Draft, `http://www1.maths.leeds.ac.uk/~adler/docs/vcm.pdf`. July 2008.

[3] Matthias Aschenbrenner et al. "Vapnik-Chervonenkis density in some theories without the independence property, I". In: *Trans. Amer. Math. Soc.* 368 (2016), pp. 5889–5949.

[4] James Ax. "The elementary theory of finite fields". In: *The Annals of Mathematics* 88 (2 Sept. 1968), pp. 239–271.

[5] E. Becker. "Euklidische Körper und euklidische Hüllen von Körpern". In: *J. reine angew. Math.* 268/269 (1974), pp. 41–52.

[6] N. Bourbaki. *Algèbre Commutative*. Éléments de mathématique. Masson, 1985.

[7] Zoé Chatzidakis and Ehud Hrushovski. "Model Theory of Difference Fields". In: *Trans. Amer. Math. Soc.* 351.8 (Apr. 1999), pp. 2997–3071.

[8] Gregory Cherlin and Saharon Shelah. "Superstable fields and groups". In: *Annals of Mathematical Logic* 18 (3 Aug. 1980), pp. 227–270.

[9] Artem Chernikov. "Theories without the Tree Property of the Second Kind". In: *Annals of Pure and Applied Logic* 165.2 (Feb. 2014), pp. 695–723.

[10] Artem Chernikov and Itay Kaplan. "Forking and Dividing in NTP_2 Theories". In: *Journal of Symbolic Logic* 77.1 (2012), pp. 1–20.

[11] Artem Chernikov, Itay Kaplan, and Pierre Simon. "Groups and Fields with NTP_2". In: *Proc. Amer. Math. Soc.* 143 (2015), pp. 395–406.

[12] Artem Chernikov and Pierre Simon. "Henselian valued fields and inp-minimality". arXiv:1601.07313 [math.LO]. 2016.

[13] Françoise Delon. "C-minimal structures without density assumption". In: *Motivic Integration and its Interactions with Model Theory and Non-Archimedean Geometry*. Ed. by Raf Cluckers, Johannes Nicaise, and Julien Sebag. Vol. 2. Cambridge University Press, Dec. 2011, pp. 51–86.

[14] Jan Denef and Lou van den Dries. "p-adic and Real Subanalytic Sets". In: *The Annals of Mathematics*. 2nd ser. 128.1 (July 1988), pp. 79–138.

[15] Alfred Dolich, John Goodrick, and David Lippel. "Dp-minimality: basic facts and examples". In: *Notre Dame Journal of Formal Logic* 52.3 (2011), pp. 267–288.

[16] Laurentius Petrus Dignus van den Dries. "Model Theory of Fields: Decidability and Bounds for Polynomial Ideals". PhD thesis. Rijksuniversiteit Utrecht, 1978.

[17] Lou van den Dries. *Tame Topology and O-minimal Structures*. London Mathematical Society Lecture Note Series 248. Cambridge University Press, 1998.

[18] Antoine Ducros. "Espaces de Berkovich, polytopes, squelettes et théorie des modèles". In: *Confluentes Mathematici* 4.4 (2012).

[19] Antonio J. Engler and Alexander Prestel. *Valued Fields*. Springer, 2005.

[20] Yuri L. Ershov. *Multi-valued Fields*. Springer, 2001.

[21] Antongiulio Fornasiero. "Definable compactness for topological structures". Unpublished. 2015.

[22] Michael D. Fried and Moshe Jarden. *Field Arithmetic*. Ergebnisse der Mathematik und ihrer Grenzgebiete. Springer, 2008.

[23] John Goodrick. "A monotonicity theorem for dp-minimal densely ordered groups". In: *Journal of Symbolic Logic* 75.1 (2010), pp. 221–238.

[24] Vince Guingona. "On VC-minimal fields and dp-smallness". In: *Archive for Mathematical Logic* 53.5-6 (Aug. 2014), pp. 503–517.

[25] Robin Hartshorne. *Algebraic Geometry*. Graduate Texts in Mathematics 52. Springer, 1977.

[26] Deirdre Haskell, Ehud Hrushovski, and Dugald Macpherson. "Definable sets in algebraically closed valued fields: elimination of imaginaries". In: *J. reine angew. Math.* 597 (2006), pp. 175–236.

[27] Deirdre Haskell, Ehud Hrushovski, and Dugald Macpherson. *Stable domination and independence in algebraically closed valued fields*. Lectures Notes in Logic 30. Cambridge University Press, 2008.

[28] Deirdre Haskell and Dugald Macpherson. "Cell decompositions of C-minimal structures". In: *Annals of Pure and Applied Logic* 66 (1994), pp. 113–162.

[29] Wilfrid Hodges. *Model Theory*. Encyclopedia of mathematics and its applications. Cambridge University Press, 1993.

[30] Jan E. Holly. "Canonical forms for definable subsets of algebraically closed and real closed valued fields". In: *J. Symbolic Logic* 60.3 (Sept. 1995), pp. 843–860.

[31] Ehud Hrushovski. "A new strongly minimal set". In: *Annals of Pure and Applied Logic* 62.2 (1993), pp. 147–166.

[32] Ehud Hrushovski. "Imaginaries and definable types in algebraically closed valued fields". arXiv:1403.7326. 2014.

[33] Ehud Hrushovski. "Pseudo-finite fields and related structures". In: *Model theory and Applications*. Ed. by Luc Bélair. Vol. 11. Quaderni di matematica. Aracne, 2002.

[34] Ehud Hrushovski and François Loeser. *Non-Archimedean Tame Topology and Stably Dominated Types*. Annals of Mathematics Studies. Princeton University Press, Mar. 2016.

[35] Ehud Hrushovski, Ya'acov Peterzil, and Anand Pillay. "Groups, measures, and the NIP". In: *J. Amer. Math. Soc.* 21.2 (Apr. 2008), pp. 563–596.

[36] Ehud Hrushovski and Anand Pillay. "On NIP and invariant measures". In: *J. Eur. Math. Soc.* 13.4 (2011), pp. 1005–1061.

[37] Ehud Hrushovski and Boris Zilber. "Zariski geometries". In: *J. Amer. Math. Soc.* 9.1 (1996), pp. 1–56.

[38] Franziska Jahnke and Jochen Koenigsmann. "Uniformly defining p-henselian valuations". In: *Annals of Pure and Applied Logic* 166.7-8 (July 2015), pp. 741–754.

[39] Franziska Jahnke, Pierre Simon, and Erik Walsberg. "Dp-minimal valued fields". arXiv: 1507.03911v1 [math.LO], to appear in JSL. 2015.

[40] Sofya Kamenkovich and Ya'acov Peterzil. "Euler Characteristic of imaginaries in o-minimal structures". Preprint, available at `http://math.haifa.ac.il/kobi/paper.3.pdf`. Sept. 2014.

[41] Itay Kaplan, Alf Onshuus, and Alexander Usvyatsov. "Additivity of the dp-Rank". In: *Trans. Amer. Math. Soc.* 365.11 (Nov. 2013), pp. 5783–5804.

[42] Itay Kaplan, Thomas Scanlon, and Frank O. Wagner. "Artin-Schreier extensions in dependent and simple fields". In: *Israel Journal of Mathematics* 185.1 (Oct. 2011), pp. 141–153.

[43] H. J. Kowalsky and H. Dürbaum. "Arithmetische Kennzeichnung von Körpertopologien". In: *J. reine angew. Math.* 191 (1953), pp. 135–152.

[44] Jan Krajíček. "Uniform families of polynomial equations over a finite field and structures admitting an euler characteristic of definable sets". In: *Proc. London Math. Soc.* 3.81 (2000), pp. 257–284.

[45] Jan Krajíček and Thomas Scanlon. "Combinatorics with definable sets: Euler characteristics and Grothendieck Rings". In: *The Bulleting of Symbolic Logic* 6.3 (Sept. 2000), pp. 311–330.

[46] Krzysztof Krupiński and Anand Pillay. "On stable fields and weight". In: *Journal of the Institute of Mathematics of Jussieu* 10 (2011), pp. 349–358.

[47] Franz-Viktor Kuhlmann. "Henselian Function Fields". PhD thesis. Heidelberg, 1989.

[48] Franz-Viktor Kuhlmann. "Quantifier elimination for henselian fields relative to additive and multiplicative congruences". In: *Israel Journal of Mathematics* 85 (1994), pp. 277–306.

[49] Dugald Macpherson. *Model Theory of Valued Fields*. `http://www1.maths.leeds.ac.uk/Pure/staff/macpherson/modnetval4.pdf`. May 2008.

[50] Dugald Macpherson, David Marker, and Charles Steinhorn. "Weakly o-minimal structures and real closed fields". In: *Trans. Amer. Math. Soc.* 352.12 (2000), pp. 5435–5483.

[51] Dugald Macpherson and Charles Steinhorn. "On variants of *o*-minimality". In: *Annals of Pure and Applied Logic* 79 (1996), pp. 165–209.

[52] James S. Milne. *Abelian Varieties (v2.00)*. Available at www.jmilne.org/math/. 2008.

[53] James S. Milne. *Étale cohomology*. Vol. 33. Princeton Mathematical Series. Princeton University Press, 1980.

[54] Samaria Montenegro-Guzmán. "Théorie des modèles des corps pseudo-réels clos et pseudo-p-adiquement clos". PhD thesis. Université Paris Diderot, 2015.

[55] David Mumford. *Abelian Varieties*. Studies in Mathematics 5. Tata Institute of Fundamental Research, 1970.

[56] Alf Onshuus and Alexander Usvyatsov. "On dp-minimality, strong dependence, and weight". In: *Journal of Symbolic Logic* 76.3 (2011), pp. 737–758.

[57] Anand Pillay. *Geometric Stability Theory*. Oxford Logic Guides 32. Oxford Science Publications, 1996.

[58] Anand Pillay. "Model theory of algebraically closed fields". In: *Model Theory and Algebraic Geometry*. Ed. by Elisabeth Bouscaren. Springer, 1998, pp. 61–84.

[59] Anand Pillay. "Some remarks on definable equivalence relations in o-minimal structures". In: *Journal of Symbolic Logic* 51.3 (Sept. 1986).

[60] Bruno Poizat. "Une théorie de Galois imaginaire". In: *The Journal of Symbolic Logic* 48 (4 Dec. 1983), pp. 1151–1170.

[61] Alexander Prestel and Martin Ziegler. "Model theoretic methods in the theory of topological fields". In: *Journal für die reine und angewandte Mathematik* 299-300 (1978), pp. 318–341.

[62] Janak Ramakrishnan, Ya'acov Peterzil, and Pantelis Eleftheriou. "Interpretable Groups are Definable". arXiv:1110.6581v1 [math.LO]. 2011.

[63] Abraham Robinson. *Complete theories*. Studies in Logic and the Foundations of Mathematics. North-Holland, 1956.

[64] Robert S. Rumely. "Arithmetic over the ring of all algebraic integers". In: *Journal für die reine und angewandte Mathematik* 368 (1986), pp. 127–133.

[65] Jean-Pierre Serre. *Algebraic Groups and Class Fields*. Graduate Texts in Mathematics. Springer, 1988.

[66] Saharon Shelah. "Strongly dependent theories". In: *Israel Journal of Mathematics* 204.1 (Oct. 2014), pp. 1–83.

[67] Pierre Simon. *A guide to NIP theories*. Lecture Notes in Logic. Cambridge University Press, July 2015.

[68] Pierre Simon. "On dp-minimal ordered structures". In: *Journal of Symbolic Logic* 76.2 (2011), pp. 448–460.

[69] A. L. Stone. "Nonstandard analysis in topological algebra". In: *Applications of model-theory to algebra, analysis, and probability theory*. Ed. by Luxemburg. Holt, Rinehart, and Winston, 1969.

[70] Alexander Usvyatsov. "On generically stable types in dependent theories". In: *Journal of Symbolic Logic* 74.1 (Mar. 2009), pp. 216–250.

[71] Itaï Ben Yaacov and Artem Chernikov. "An Independence Theorem for NTP_2 Theories". In: *Journal of Symbolic Logic* 79 (1 Mar. 2014), pp. 135–153.

[72] M. Ziegler. "A language for topological structures which satisfies a Lindström-theorem". In: *Bulletin of the AMS* 82.4 (July 1976).