The Mordell-Lang Theorem from the Zilber Dichotomy

by

Christopher Eagle

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Pure Mathematics

Waterloo, Ontario, Canada, 2010

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

We present a largely self-contained exposition of Ehud Hrushovski's proof of the function field Mordell-Lang conjecture beginning from the Zilber Dichotomy for differentially closed fields and separably closed fields. Our account is based on notes from a series of lectures given by Rahim Moosa at a MODNET workshop at Humboldt Universität in Berlin in September 2007. We treat the characteristic 0 and characteristic p cases uniformly as far as is possible, then specialize to characteristic p in the final stages of the proof. We also take this opportunity to work out the extension of Hrushovski's "Socle Theorem" from the finite Morley rank setting to the finite U-rank setting, as is in fact required for Hrushovski's proof of Mordell-Lang to go through in positive characteristic.

Acknowledgements

First, I would like to thank my advisor, Professor Rahim Moosa, for his guidance throughout my time as a Masters student. In addition to providing immeasurable help in the process of learning the mathematical material presented here, the quality of this thesis was greatly improved by his comments throughout the writing process.

I would also like to thank the other members of the Pure Mathematics department at the University of Waterloo, especially Professors David McKinnon and Ruxandra Moraru, for always being available for discussion. Lis d'Alessio, through her patience and knowledge, made my time as a Masters student significantly more pleasant than it otherwise might have been.

I offer thanks also to Professor Ross Willard of the University of Waterloo and Professor Bradd Hart of McMaster University for acting as readers for this thesis and suggesting further areas for improvement.

I appreciate the financial support of the Natural Sciences and Engineering Research Council and the University of Waterloo throughout the program. Finally, I also appreciate the support and encouragement of my family and friends.

Dedication

To Amy, who taught me that lemmas studying tops can have just as much fun as lemmas in a field.

Contents

Chapter 1

Model-Theoretic Preliminaries

We hope to keep the model-theoretic prerequisites of this thesis to a first graduate course in model theory. Toward that end, in this chapter we review and summarise some of the more advanced results from model theory that we will need. A more detailed discussion of all of this material may be found in [25]. The reader desiring additional background could also consult [24] and [9]. Some additional model-theoretic notions will be introduced later, as they are needed, and we provide in Appendix B.1 proofs of some further results which, though elementary, do not seem to be available in the literature. Since much of our work will eventually be carried out in two specific theories, most of the abstract notions discussed here will have rather concrete manifestations when we use them.

1.1 Saturation and the Universal Domain

Let $\mathcal{L}$ be a first-order language, T a complete $\mathcal{L}$-theory. For convenience of exposition we assume that $\mathcal{L}$ is one-sorted, although all of what we describe here can be developed in the general case as well. Let κ be an infinite cardinal, and let $\mathcal{M} \models T$. Recall that $\mathcal{M}$ is *κ-saturated* if, for every $A \subseteq M$ with $|A| < \kappa$ and every $n \geq 1$, every n-type over A is realized in M. $\mathcal{M}$ is called *saturated* if it is $|M|$-saturated. We say that $\mathcal{M}$ is *strongly κ-homogeneous* if whenever $f : A \to M$ is a partial elementary map with $A \subseteq M$ and $|A| < \kappa$ then f extends to an automorphism of M. We say $\mathcal{M}$ is *strongly homogeneous* if it is strongly $|M|$-homogeneous.

Every model $\mathcal{M}$ has elementary extensions which are κ-saturated and strongly κ-homogeneous for any $\kappa \geq |M|$. Saturation implies strong homogeneity (although this is not true of κ-saturation and strong κ-homogeneity), so we often omit mention of strong homogeneity when discussing saturated models.

Now suppose that T is a complete first-order theory. A model $\mathcal{U} \models T$ which is κ-saturated and strongly κ-homogeneous is κ^+*-universal*, in the sense that if $\mathcal{N} \models T$ has $|N| \leq \kappa$ then then there is an elementary embedding of $\mathcal{N}$ into $\mathcal{U}$. In particular, if $\mathcal{U}$ is saturated then every model of smaller or equal cardinality can be identified with an elementary substructure of $\mathcal{U}$. By a *universal domain* for T we mean a model which is κ-saturated and strongly κ-homogeneous for some infinite cardinal κ larger than any parameters or models we are interested in. We usually do not specify in advance how large κ must be; rather, we note that we could keep track of the sizes of every parameter set and model we use, and then later go back to fix an appropriately large κ.

As has become standard in model theory, we work in a fixed universal domain $\mathcal{U}$ of our complete first-order theory T. All parameter sets are assumed to be subsets of $\mathcal{U}$ of cardinality less than κ, and all models of T are assumed to be elementary submodels of $\mathcal{U}$ of of cardinality less than κ, except when we explicitly state otherwise. In particular, types are always types over parameters of size less than κ. By a *global type*, however, we mean a type over $\mathcal{U}$ itself. We usually write $\models \phi(a)$ instead of $\mathcal{U} \models \phi(a)$.

A key benefit of working in a universal domain is that we can detect when an element is algebraic or definable over a set of parameters using automorphism arguments, as follows: Let A be a set of parameters (of cardinality less than κ). Then for any tuple a, we have $a \in \mathrm{acl}(A)$ if and only if the orbit of a under $\mathrm{Aut}_A(\mathcal{U})$ is finite, if and only if $\mathrm{tp}(a/A)$ has only finitely many realizations. Similarly, $a \in \mathrm{dcl}(A)$ if and only if $f(a) = a$ for all $f \in \mathrm{Aut}_A(\mathcal{U})$, if and only if $\mathrm{tp}(a/A)$ has only one realization.

A set is *type-definable over* A if it is the set of a realizations of a type over A. If X is type-definable we will say that $Y \subseteq X$ is *(relatively) definable in* X (or is a *definable subset of* X) if $Y = X \cap D$ where D is definable. Given type-definable sets X and Y, by a *definable map* $f : X \to Y$ we mean a function whose graph is a definable subset of the type-definable set $X \times Y$.

1.2 Elimination of Imaginaries

The automorphism arguments described in the previous section can be extended from tuples to definable sets in theories with elimination of imaginaries. If D is a definable set and d is a tuple then we say that d is a *code* for D if, for all $f \in \mathrm{Aut}(\mathcal{U})$, we have $f(d) = d \iff f(D) = D$. Equivalently, d is a code for D if and only if there is a formula $\phi(x, y)$ such that $\phi(x, d)^{\mathcal{U}} = X$ and for any d' if $\phi(x, d')^{\mathcal{U}} = X$ then $d = d'$. We say that T has *elimination of imaginaries* if every definable set in $\mathcal{U}$ has a code.

We can force elimination of imaginaries by passing to the theory T^{eq}, defined as follows. Let $\mathcal{E}$ be the set of all $\mathcal{L}$-formulae $E(x, y)$ (where x and y are n_E-tuples of variables for some $n_E \in \mathbb{N}$) such that in every model of T, E defines an equivalence relation. We form a many-sorted language L^{eq} by taking the symbols from L as symbols on the sort $S_=$, adding a new sort S_E for each $E \in \mathcal{E}$, and adding function symbols $f_E : S_=^{n_E} \to S_E$ for each $E \in \mathcal{E}$. The theory T^{eq} is then the theory whose axioms are the axioms of T restricted to the sort $S_=$, together with axioms expressing that each f_E is a surjective map from $S_=^{n_E}$ to S_E such that for all a, b we have $f_E(a) = f_E(b) \iff E(a,b)$. We expand $\mathcal{U}$ into a model $\mathcal{U}^{\mathrm{eq}} \models T^{\mathrm{eq}}$ in the natural way, where the sort S_E is interpreted as $\mathcal{U}^n/E$. Then $\mathcal{U}^{\mathrm{eq}}$ is a universal domain for T^{eq}. The theory T^{eq} is complete and has elimination of imaginaries, but gains no structure not already present in T. For example, every automorphism of $\mathcal{U}$ extends uniquely to an automorphism of $\mathcal{U}^{\mathrm{eq}}$. Further, if we identify $\mathcal{U}$ with the sort $S_=$ of $\mathcal{U}^{\mathrm{eq}}$, then any subset of $\mathcal{U}^n$ definable in $\mathcal{U}^{\mathrm{eq}}$ is already definable in $\mathcal{U}$. If T already has elimination of imaginaries (and $\mathrm{dcl}(\emptyset)$ has at least two elements) then for any 0-definable equivalence relation E on $\mathcal{U}^n$ there is a definable function f_E such that uEv if and only if $f_E(u) = f_E(v)$. We can thus identify $f_E(\mathcal{U}^n)$ as the quotient $\mathcal{U}^n/E$, and hence do not need T^{eq}. Nevertheless, we generally work in $\mathcal{U}^{\mathrm{eq}}$ without mentioning it explicitly. In particular, from now on when we write acl or dcl we mean it in the sense of $\mathcal{U}^{\mathrm{eq}}$.

1.3 Stability and Definability of Types

We now describe a hierarchy of tameness properties that a theory may possess, and indicate some relationships between them.

A type-definable set $X \subseteq \mathcal{U}^n$ is called *minimal* if every set definable in X is either finite or cofinite. A complete theory is called *strongly minimal* if the universe of every model is minimal. This condition has very strong consequences for the theory. In particular, in a strongly minimal theory the algebraic closure relation satisfies the Steinitz exchange property: if $a \in \mathrm{acl}(A \cup \{b\}) \setminus \mathrm{acl}(A)$ then $b \in \mathrm{acl}(A \cup \{a\})$. A model of T is then a matroid with acl as the closure relation. This gives rise to the notion of acl-independence: We say that Y is acl-*independent over* A if $a \notin \mathrm{acl}(A \cup (Y \setminus \{a\}))$ for all $a \in Y$. Given any set $Y \subseteq \mathcal{U}$, all maximally acl-independent subsets of Y over A have the same cardinality, and this cardinality is called the acl-*dimension* of Y over A. If $a = (a_1, \ldots, a_n) \in \mathcal{U}^n$ then by acl-dim(a/A) we mean acl-dim$(\{a_1, \ldots, a_n\}/A)$, which is no greater than n. By the acl-dimension of a type we mean the acl-dimension of any of the realizations of that type. In a strongly minimal theory there is, for each n, a unique n-type over A of acl-dim n over A. This type is called the *generic* n-type over A.

Definition 1.1. Let κ be an infinite cardinal. The complete theory T is called κ-*stable* if, for every $\mathcal{M} \models T$ and every $A \subseteq M$ such that $|A| \leq \kappa$, we have $\left|S_n^{\mathcal{M}}(A)\right| \leq \kappa$.

Every strongly minimal theory is κ-stable for all κ. Theories which are $\aleph_0$-stable are referred to as ω-stable. Under the assumption that the language $\mathcal{L}$ is countable such theories also possess a rank function, called the Morley rank, which is ordinal-valued on all definable sets. We omit the definition here, and only mention that it agrees with acl-dimension in the strongly minimal case.

A theory which is κ-stable for all sufficiently large cardinals κ is called *superstable.* As, ω-stability implies κ-stability for all infinite κ, ω-stable theories are superstable. Superstable theories also admit a rank function, defined on complete types, which is ordinal-valued for all complete types. This rank is Lascar's U-rank, which we will define in the next section, and which need not agree with Morley rank in ω-stable theories.

The theory T is called *stable* if it is κ-stable for some infinite κ. Stability has significant consequences for a theory. One useful fact is that if T is stable then T has saturated models of arbitrarily large cardinality. (Hence when we are working in a stable theory we can take a saturated, rather than merely κ-saturated, model for our universal domain.)

A key property of stable theories is that every complete type in a stable theory is definable, in the following sense:

Definition 1.2. Suppose that $p(x) \in S_n(B)$ where $x = (x_1, \ldots, x_n)$, and suppose $A \subseteq B$. We say that p is *A-definable* if, for every $\mathcal{L}$-formula $\phi(x, y)$ with $y = (y_1, \ldots, y_m)$ there is an $\mathcal{L}_A$-formula $d_p\phi(y)$ such that for all b,

$$\phi(x, b) \in p(x) \iff \models d_p\phi(b)$$

We say that p is definable if it is definable over some $A \subseteq B$.

Besides definability of types, the other key, and related, use of stability is that in stable theories there is a very well-behaved notion of independence.

1.4 Independence and U-rank

In this section T is assumed to be stable. Recall that we work in T^{eq}, so in particular acl is to be understood as acl in $\mathcal{U}^{\mathrm{eq}}$. The stability of T gives rise to a useful notion of a tuple a being *independent* of a (small) set B over a (small) set A, denoted by $a \mathop{\smile\!\!\!\!\!|}_A B$ (or $a \mathop{\smile\!\!\!\!\!|} B$ when $A = \emptyset$). The idea is that a is not substantially more related to $B \cup A$ than it is to A. Rather than present the technical definition of independence in terms of Shelah's *non-forking*, we give the following axioms which characterize it. If $b = (b_1, \ldots, b_n)$ is a tuple we write $a \mathop{\smile\!\!\!\!\!|}_A b$ to mean $a \mathop{\smile\!\!\!\!\!|}_A \{b_1, \ldots, b_n\}$.

1. (Finite Character) $a \mathop{\smile\!\!\!\!\!|}_A B \iff a \mathop{\smile\!\!\!\!\!|}_A b$ for all finite tuples b from B.

2. $\mathop{\smile\!\!\!\!\!|}$ is automorphism-invariant.

3. (Symmetry) $a \mathop{\smile}\limits^{\vert}_{A} b \iff b \mathop{\smile}\limits^{\vert}_{A} a$.
4. (Transitivity and Monotonicity) $a \mathop{\smile}\limits^{\vert}_{A} bc \iff a \mathop{\smile}\limits^{\vert}_{A} b$ and $a \mathop{\smile}\limits^{\vert}_{Ab} c$.
5. (Existence) For any a, A, B there is $a' \models \mathrm{tp}(a/A)$ with $a' \mathop{\smile}\limits^{\vert}_{A} B$.
6. (Local Character) For any a, A there is $A_0 \subseteq A$ with $|A_0| \leq |\mathcal{L}|$ such that $a \mathop{\smile}\limits^{\vert}_{A_0} A$.
7. (Stationarity) If $A = \mathrm{acl}(A)$, $B \supseteq A$, $a \mathop{\smile}\limits^{\vert}_{A} B$, $b \mathop{\smile}\limits^{\vert}_{A} B$ and $\mathrm{tp}(a/A) = \mathrm{tp}(b/A)$ then $\mathrm{tp}(a/B) = \mathrm{tp}(b/B)$.

It is a fact that there is a unique such notion in a stable theory, and that it coincides with Shelah's non-forking. If $a \mathop{\smile}\limits^{\vert}_{C} b$ we say that $\mathrm{tp}(a/Cb)$ is a *free* or *non-forking* extension of $\mathrm{tp}(a/C)$. We say that $p \in S_n(A)$ *does not fork* over $B \subseteq A$ if p is a non-forking extension of its restriction to B. Given $p \in S_n(A)$ and $B \supseteq A$ there always exists at least one non-forking extension of p to B, this is axiom (4). The type p is called *stationary* if it has a unique non-forking extension to each $B \supseteq A$. Types over algebraically closed sets are stationary (this is axiom (6)). If T is strongly minimal then $a \mathop{\smile}\limits^{\vert}_{C} b$ if and only if $\text{acl-dim}(a/Cb) = \text{acl-dim}(a/C)$.

There is a natural dimension associated to non-forking in superstable theories, which we can now define:

Definition 1.3. The *U-rank* of a complete type $p \in S_n(A)$ is defined inductively, as follows:

- $U(p) \geq 0$.
- If α is a limit ordinal then $U(p) \geq \alpha$ if and only if $U(p) \geq \beta$ for all $\beta < \alpha$.
- For any ordinal α, $U(p) \geq \alpha + 1$ if and only if there exists $B \supseteq A$ and $q \in S_n(B)$ such that q is a forking extension of p (i.e., an extension which is not non-forking) and $U(q) \geq \alpha$.

We define $U(p) = \alpha$ if $U(p) \geq \alpha$ and $U(p) \not\geq \alpha + 1$. We define $U(p) = \infty$ if $U(p) \geq \alpha$ for every ordinal α. We write $U(a/A)$ for $U(\mathrm{tp}(a/A))$. The U-rank of a partial type Φ is defined to be the supremum of the U-ranks of all of the completions of Φ in $S(B)$.

We commented earlier that when T is superstable every complete type has ordinal-valued U-rank. Moreover, independence can be understood in terms of U-rank: If T is superstable then $a \mathop{\smile}\limits^{\vert}_{C} b$ if and only if $U(a/Cb) = U(a/C)$. A key property of U-rank which we will frequently need are the Lascar inequalities:

Fact 1.4. *For any parameters A and any tuples a, b such that $U(b/Aa), U(a/A), U(ab/A) < \infty$ we have*

$$U(b/Aa) + U(a/A) \leq U(ab/A) \leq U(b/Aa) \oplus U(a/A),$$

where $+$ is the ordinary ordinal sum and $\oplus$ is the natural (or Hessenberg) sum. In particular, when $U(b/Aa)$ and $U(a/A)$ are both finite the ordinal sum and the natural sum agree, and we have

$$U(b/Aa) + U(a/A) = U(ab/A)$$

In Section 1.3 we defined the generic n-type over parameters A in a strongly minimal theory as the (unique) n-type whose realizations are acl-independent over A. We now extend this notion of generic to the superstable setting. Suppose that X is type-definable over A with $U(X) < \infty$. We say that a complete type $p \in S(A)$ extending X is a *generic type of X over A* if $U(p) = U(X)$. Such a generic type always exists. Since non-forking extensions preserve U-rank a non-forking extension of a generic type is again generic. We say that $a \in X$ is *generic in X over A* if $\mathrm{tp}(a/A)$ is generic in X. Observe that if $A \subseteq B$ then a is generic in X over B if and only if a is generic in X over A and $a \mathop{\smile\!\!\!\!\!\!|}_A B$.

1.5 Canonical Bases

In Section 1.2 we described the code of a definable set X as an element (in $\mathcal{U}^{\mathrm{eq}}$) fixed by exactly those automorphisms that fix X as a set. We would like to have an analogous notion for complete types.

Definition 1.5. Let $\mathbf{p} \in S_n(\mathcal{U})$ be a complete global type. We say $A \subseteq \mathcal{U}^{\mathrm{eq}}$ is a *canonical base* for $\mathbf{p}$ if, for all $\sigma \in \mathrm{Aut}(\mathcal{U})$ we have $\sigma(\mathbf{p}) = \mathbf{p} \iff \sigma$ fixes A pointwise.

In a stable theory every global type has a (small) canonical base. Given a global type $\mathbf{p}$ and canonical bases A, B for $\mathbf{p}$ we have $\mathrm{dcl}(A) = \mathrm{dcl}(B)$. We can therefore define $\mathrm{cb}(\mathbf{p}) = \mathrm{dcl}(A)$ where A is any canonical base for $\mathbf{p}$. We can extend the definition of canonical bases from global types to stationary types in a straightforward way: Given a stationary type p, let $\mathbf{p}$ be its unique global non-forking extension. Then we define $\mathrm{cb}(p) = \mathrm{cb}(\mathbf{p})$. We write $\mathrm{cb}(a/A)$ for $\mathrm{cb}(\mathrm{tp}(a/A))$ when $\mathrm{tp}(a/A)$ is stationary. In general we do not have that an automorphism fixes p if and only if it fixes $\mathrm{cb}(p)$, but we do have that an automorphism σ fixes $\mathrm{cb}(p)$ if and only if p and $\sigma(p)$ have a common non-forking extension. The canonical base has many useful properties, which we summarize below:

Fact 1.6. *Let $p \in S_n(A)$ be a stationary type. Then*

1. $\mathrm{cb}(p) \subseteq \mathrm{dcl}(A)$.

2. *For any $B \subseteq A$, $\mathrm{cb}(p) \subseteq \mathrm{acl}(B)$ if and only if p does not fork over B.*

3. *For any $B \subseteq A$, $\mathrm{cb}(p) \subseteq \mathrm{dcl}(B)$ if and only if p does not fork over B and the restriction of p to B is stationary.*

1.6 Stable Groups

In this final section we review some of the theory of stable groups. We continue to work in a complete stable theory T (or, rather, T^{eq}). The reader desiring more information about this material should consult [25, Chapter 1 Section 6] and [29, Chapter 5]. The material on superstable groups is found in [2].

Definition 1.7. A *type-definable group* is a type-definable set G together with definable maps $\cdot : G \times G \to G$ and ${}^{-1} : G \to G$ such that $(G, \cdot, {}^{-1})$ is a group.

We note that in the definition of a type-definable group it might appear more natural to allow the operations $\cdot$ and ${}^{-1}$ to be type-definable, rather than relatively definable. In fact this does not offer a more general setting, as a compactness argument shows that if the group operations are type-definable then they are relatively definable (see Lemma B.3).

Fact 1.8. 1. *If G is a type-definable group over A then there exists an A-definable group H and A-definable subgroups $H_i \leq H$ for each $i \in I$ such that $G = \bigcap_{i \in I} H_i$. Moreover, the H_i's can be chosen "canonically", in the sense that any automorphism fixing G setwise fixes each H_i setwise (see [25, Lemma I.6.18 and Remark I.6.20]).*

2. *If G is a type-definable subgroup of the type-definable group G' then there exist (relatively) definable subgroups H_i of G' such that $G = \bigcap_{i \in I} H_i$. Again the H_i's can be chosen so that every automorphism fixing G setwise fixes each H_i setwise.*

It follows from this fact that in the ω-stable case, in which one has the descending chain condition on definable subgroups, that every type-definable group is, in fact, definable. Another consequence of the above fact is that all cosets of a type-definable subgroup have "canonical parameters":

Proposition 1.9. *Let G be a type-definable group over A, and let $H \leq G$ be a type-definable subgroup. Then for all $a \in G$ there exists $B \supseteq A$ such that $a + H$ is type-definable over B, and for all $\sigma \in \mathrm{Aut}_A(\mathcal{U})$, $\sigma(a + H) = a + H$ if and only if $\sigma|_B = \mathrm{id}$.*

Proof. By Fact 1.8 G is a type-definable subgroup of an A-definable group G', so we may replace G by G', and hence assume that G is definable.

Write $H = \bigcap_{i \in I} H_i$, where each H_i is an A-definable subgroup of G and each H_i is fixed by all automorphisms fixing H. For each $i \in I$ let b_i be a code for $a + H_i$, and let $B = \{b_i : i \in I\}$. Since $a+H = \bigcap_{i \in I} a+H_i$ and each $a+H_i$ is b_i-definable, $a+H$ is type-definable over B. If $\sigma \in \mathrm{Aut}_A(\mathcal{U})$ and $\sigma(a+H) = a+H$ then $\sigma(a+H_i) = \sigma(a)+\sigma(H_i) = \sigma(a)+H_i$, and $\sigma(a)$ and a are in the same coset of H, namely $a+H$. Thus $\sigma(a+H_i) = a+H_i$. Hence $\sigma(b_i) = b_i$, so $\sigma|_B = \mathrm{id}$, as required. □

Definition 1.10. A type-definable group G is called *connected* if it has no proper non-trivial definable subgroup of finite index.

Suppose that G is type-definable over parameters A. The *connected component* of G, G°, is the intersection of all (relatively) A-definable subgroups of finite index in G. It is clearly type-definable over A. In fact, G° is the intersection of all finite index definable subgroups of G over any parameters (see [25, Corollary I.6.14]). In particular, G° is connected.

A (relatively) definable subset $D \subseteq G$ is called *generic* if there exists a finite collection $a_1, \ldots, a_n \in G$ such that $G = a_1 D \cup \ldots \cup a_n D$, where $a_i D$ is the group-theoretic translate of D by a_i. If D is any relatively definable subset of G then either D or $G \setminus D$ is generic in G.

A complete type p extending G is *generic* if every formula in p defines a generic subset of G. If G is connected then for any parameters A over which G is defined G has a unique generic type over A. The following facts describe the interaction between genericity and the group operation:

Fact 1.11. *Let G be a type-definable group over parameters A, and $a \in G$.*

1. $\mathrm{tp}(a/A)$ *is generic if and only if for every $g \in G$ such that $a \mathop{\smile\!\!\!\!\!\!|}_A g$ we have $g \cdot a \mathop{\smile\!\!\!\!\!\!|}_A A \cup \{g\}$.*
2. *If $g \in G$ is generic in G over B and $a \in B$ then $g \cdot a$ and g^{-1} are generic in G over B.*
3. *If G is connected then there exist $g, h \in G$, both generic over $\{a\}$, such that $a = g \cdot h$.*

Next we discuss stabilizers. Suppose that G is a type-definable group over A, and $p(x)$ is a complete type over some $B \supseteq A$ extending the type "$x \in G$" (we say that $p(x)$ *is in* G). If $g \in G \cap \mathrm{dcl}(B)$ then we can define

$$g \cdot p = \left\{ \phi(g^{-1} \cdot x) : \phi \in p(x) \right\}$$

We note that $g \cdot p$ is also a complete type in G over B. We also note that the above definition makes sense even when B is not a small set of parameters. In particular, we will use this in the case when $B = \mathcal{U}$.

Definition 1.12. Suppose that $p(x)$ is a complete stationary type in G. Let $\mathbf{p}$ denote the unique global non-forking extension of p to the universal domain $\mathcal{U}$. The *(model-theoretic) stabilizer* of p is

$$\operatorname{stab}(p) = \{g \in G : g \cdot \mathbf{p} = \mathbf{p}\} \leq G.$$

It turns out, because of definability of types, that the stabilizer of $p(x)$ is a type-definable subgroup of G, defined over $\operatorname{cb}(p)$. Stabilizers can also be used to characterize the generic types in G: $p(x)$ is generic in G if and only if $\operatorname{stab}(p) = G^{\circ}$.

Proposition 1.13. *Let $p(x)$ be a complete stationary type in G over $B \supseteq A$. Then the following are equivalent:*

1. *$g \in \operatorname{stab}(p)$.*
2. *For some (any) $a \models p$ such that $a \mathop{\smile\!\!\!\!\!\!\mid}_B g$, $g \cdot a \models p$ and $g \cdot a \mathop{\smile\!\!\!\!\!\!\mid}_B g$.*

Proof. Let $\mathbf{p}$ denote the unique global non-forking extension of p, and let p' be the unique non-forking extension of p to Bg.

(1) $\Rightarrow$ (2) : By stationarity and the hypothesis $a \mathop{\smile\!\!\!\!\!\!\mid}_B g$ we see $p' = \operatorname{tp}(a/Bg)$, and so $\mathbf{p}|_{Bg} = p'$. Since $g \cdot \mathbf{p} = \mathbf{p}$ we have that $\mathbf{p}$ extends $g \cdot p' \in S(Bg)$ also. Hence $g \cdot p' = p'$, and we get that $g \cdot a \models p'$. That is, $g \cdot a \models p$ and $g \cdot a \mathop{\smile\!\!\!\!\!\!\mid}_B g$.

(2) $\Rightarrow$ (1) : The statement (2) says that $g \cdot p' = p'$. Note that the global non-forking extension $\mathbf{p}$ of p is also the unique global non-forking extension of p'. But $g \cdot \mathbf{p}$ is also a global non-forking extension of $g \cdot p' = p'$, so $g \cdot \mathbf{p} = \mathbf{p}$. Hence $g \in \operatorname{stab}(p)$. □

We conclude with a few remarks concerning the superstable case when $U(G) < \infty$. First, the notion of generic given here agrees with the one introduced in Section 1.4. That is, p is generic in G if and only if $U(p) = U(G)$. For a (relatively) definable subgroup H of G, $U(H) = U(G)$ if and only if $[G : H]$ is finite. In particular, since $U(G^{\circ})$ is an intersection of definable subgroups of G of finite index, $U(G^{\circ}) = U(G)$. When $H \trianglelefteq G$ is a definable subgroup the Lascar inequalities described earlier take on the particularly convenient form

$$U(H) + U(G/H) \leq U(G) \leq U(H) \oplus U(G/H).$$

Finally, in the superstable setting we also have the following stronger characterisation of stabilizers:

Proposition 1.14. *Suppose that $U(G) < \infty$, and let $p(x)$ be a complete stationary type in G over $B \supseteq A$. Then the following are equivalent:*

1. *$g \in \operatorname{stab}(p)$.*

2. For some (any) $a \models p$ *such that* $a \mathop{\smile\!\!\!\!\!\!|}_{B} g$, $g \cdot a \models p$.

Proof. We actually show that if $a \mathop{\smile\!\!\!\!\!\!|}_{B} g$ and $\mathrm{tp}(g \cdot a/B) = \mathrm{tp}(a/B)$ then $g \cdot a \mathop{\smile\!\!\!\!\!\!|}_{B} g$. The result then follows from Proposition 1.13. So suppose $a \mathop{\smile\!\!\!\!\!\!|}_{B} g$ and $\mathrm{tp}(g \cdot a/B) = \mathrm{tp}(a/B)$. Then

$$\begin{aligned} U(g \cdot a/Bg) &= U(a/Bg) \\ &= U(a/B) && \text{since } a \mathop{\smile\!\!\!\!\!\!|}_{B} g \\ &= U(g \cdot a/B) && \text{since } \mathrm{tp}(g \cdot a/B) = \mathrm{tp}(a/B). \end{aligned}$$

So $g \cdot a \mathop{\smile\!\!\!\!\!\!|}_{B} g$. □

Chapter 2

Introduction

In this thesis we describe a striking application of the tools of model theory to a problem in diophantine geometry. Ehud Hrushovski's 1996 proof of the function field Mordell-Lang conjecture differs from previous applications of model theory both because no purely algebraic proof of the result is known, and because of the depth of the model-theoretic methods and results it employs. This thesis is devoted to presenting Hrushovski's argument, starting from the Zilber Dichotomy. Except where explicitly noted otherwise, our exposition is an elaboration on the notes from a series of lectures given by Rahim Moosa ([21]). The other exception is Appendix A, which contains new material. We assume that the reader has a background in model theory and algebraic geometry equivalent to a first graduate course in each of these subjects.

2.1 Motivation : From Mordell to Mordell-Lang

We begin with a conjecture formulated by Mordell in the 1920's for curves over $\mathbb{Q}$, and proved by Faltings for curves over any number field in the 1980's.

Theorem 2.1 (Mordell Conjecture)**.** *Let k be a number field, and let C be an algebraic curve over k. If* $\mathrm{genus}(C) \geq 2$ *then $C(k)$ is finite.*

To facilitate the generalisations we are interested in we reformulate the theorem. Recall that an *abelian variety* is a connected projective algebraic group; that is, an irreducible projective variety A equipped with regular morphisms $\cdot : A \times A \to A$ and ${}^{-1} : A \to A$ which make A into a group. See [20] or [14] for detailed discussions of abelian varieties.

Theorem 2.2 (Reformulated Mordell Conjecture)**.** *Let A be an abelian variety over a number field k, and let $C \subseteq A$ be an algebraic curve over k. Then $C(k)$ is a finite union of translates of subgroups of $A(k)$.*

Proposition 2.3. *Theorem 2.1 is equivalent to Theorem 2.2.*

Proof. First, a word about the Jacobian varieties. For every curve of genus > 0 there is an abelian variety $J(C)$ such that $\dim(J(C)) = \text{genus}(C)$, and a canonical regular embedding $f : C \to J(C)$ such that the following universal property holds: Let A be any abelian variety, and let $h : C \to A$ be a rational map. Then there is a unique morphism of algebraic groups $\alpha : J(C) \to A$ and a constant a such that $h = \alpha f + a$.

$$\begin{tikzcd}
C \arrow[rr, "h-a"] \arrow[dr, "f"'] & & A \\
& J(C) \arrow[ur, dashed, "\alpha"'] &
\end{tikzcd}$$

Moreover, if C is defined over a field k and has a k-rational point, then $J(C)$ and f can be chosen to be defined over k. If the map $h : C \to A$ is also defined over k then α is defined over k and a is a k-rational point. See [14, Theorems 2.2.8 and 2.2.9].

Suppose that the reformulated Mordell conjecture holds. We may assume that C is irreducible, since if not then we carry out the following proof for each of the (finitely many) irreducible components of C. Suppose that $C(k)$ is infinite. We must show that genus(C) is either 0 or 1, so we suppose that genus$(C) \neq 0$ and show that genus$(C) = 1$. Recall that having genus 1 is equivalent to having an algebraic group structure. Now embed C in its Jacobian J. As noted above, since we are assuming C has (infinitely many) k-rational points, J is an abelian variety defined over k. Then by the reformulated Mordell conjecture there exist $a_1, \ldots, a_n \in J(k)$ and $G_1, \ldots, G_n$ subgroups of $J(k)$ such that $C(k) = \cup_{i=1}^{n} a_i + G_i$. Since $C(k)$ is infinite it is Zariski-dense in C. Thus we have

$$\begin{aligned}
C &= \overline{C(k)} \\
&= \overline{\bigcup_{i=1}^{n} a_i + G_i} \\
&= \bigcup_{i=1}^{n} \overline{a_i + G_i} \\
&= \overline{a_{i_0} + G_{i_0}} && \text{for some } i_0\text{, since } C \text{ is irreducible} \\
&= a_{i_0} + \overline{G_{i_0}}
\end{aligned}$$

Since $G_{i_0} \leq J(k)$ it follows from the general theory of algebraic groups that $\overline{G_{i_0}}$ is an algebraic subgroup of A. So we have seen that C is a translate of an algebraic group, and hence can be given the structure of an algebraic group itself. Hence genus$(C) = 1$ as required.

For the converse, suppose that the Mordell conjecture holds, and $C \subseteq A$ is a curve contained in an abelian variety, both over k. The case $C(k) = \emptyset$ is trivial, so we assume that C has a k-rational point. If $\text{genus}(C) = 0$ then C is birationally equivalent to $\mathbb{P}^1$, which cannot be embedded into any abelian variety (see [20, Proposition III.3.9]). So the hypotheses of the reformulated Mordell conjecture force $\text{genus}(C) > 0$. If $\text{genus}(C) \geq 2$ then, by the Mordell conjecture, $C(k)$ is finite, and hence a finite union of translates of the subgroup $\{0\}$ of A. Otherwise, if $\text{genus}(C) = 1$ then $C = J(C)$, in which case by the universal property of Jacobians there exists a such that $J(C) + a$ is a subgroup of A. □

Definition 2.4. Let A be an abelian variety over $\mathbb{C}$, and let $\Gamma \leq A(\mathbb{C})$ be any subgroup. Let $\text{div}(\Gamma) = \{g \in A(\mathbb{C}) : ng \in \Gamma \text{ for some } n > 0\}$ be the *divisible hull* of Γ. We say that Γ has *finite rank* if there exists a finitely generated group $\Gamma' \leq A(\mathbb{C})$ such that $\Gamma \leq \text{div}(\Gamma')$.

Example 2.5. It is clear that if Γ is itself finitely generated then Γ is of finite rank. The Mordell-Weil Theorem says that $A(k)$ is a finitely generated subgroup of $A(\mathbb{C})$, and so is an example of a finite rank group.

Example 2.6. Let $\text{Tor}(A) = \{g \in A(\mathbb{C}) : \exists n > 0 \text{ such that } ng = 0\}$. Then $\text{Tor}(A) = \text{div}(\{0\})$, so $\text{Tor}(A)$ is of finite rank. $\text{Tor}(A)$ is not finitely generated, however, as we now show. Suppose to the contrary that $\text{Tor}(A)$ is finitely generated, say by $a_1, \ldots, a_m$. So any element of $\text{Tor}(A)$ is of the form $n_1a_1 + \ldots + n_ma_m$ for some $n_1, \ldots, n_m \in \mathbb{Z}$. Let $|a_i|$ denote the order of a_i, and let $l = \text{lcm}(|a_1|, \ldots, |a_m|)$. Observe that $la_i = 0$ for all i, since $l \geq |a_i|$ for all i. Thus $l(n_1a_1 + \ldots + n_ma_m) = n_1(la_1) + \ldots + n_m(la_m) = 0$. It follows that $|n_1a_1 + \ldots + n_ma_m| \leq l$. But an abelian variety has torsion elements of arbitrarily high order, so there is an element of $\text{Tor}(A)$ of order at least $l + 1$, contradicting $\text{Tor}(A)$ being generated by $a_1, \ldots, a_m$.

We can now generalize Theorem 2.2 in several ways:

Generalize the geometric object: Rather than considering only curves C we can consider arbitrary, perhaps higher-dimensional, subvarieties of A.

Generalize the arithmetic object: Observe that $C(k) = C \cap A(k)$, so Theorem 2.2 can be seen as describing the structure of the intersection of C with the finitely generated subgroup $A(k)$. We saw in Example 2.5 that $A(k)$ is a finite rank subgroup of A. Instead of $A(k)$ we can consider intersections with arbitrary finite rank subgroups of A.

Generalize the ambient algebraic group: We can replace the abelian variety A by a semiabelian variety S. A *semiabelian variety* is a connected commutative algebraic group over a field k such that there exists a short exact sequence

$$\begin{tikzcd} 0 \arrow[r] & T \arrow[r] & S \arrow[r] & A \arrow[r] & 0 \end{tikzcd}$$

where A is an abelian variety, and $T \cong \mathbb{G}_m^s$ for some $s \in \mathbb{N}$, where $\mathbb{G}_m$ denotes the multiplicative group of k. Some details about semiabelian varieties can be found in Appendix B.2.

Generalize the ground field: In characteristic 0 we can replace $\mathbb{C}$ by any algebraically closed field. We will see that this does not quite work in characteristic p, but in that context we will weaken the conclusion to get the Mordell-Lang statement that we will prove in these notes.

Performing the above generalizations in characteristic 0 gives the following statement, proved by the work of Faltings, McQuillen, Raynaud, Vojta, and others:

Theorem 2.7 (Absolute Mordell-Lang in Characteristic 0)**.** *Let F be an algebraically closed field of characteristic* 0*. Let S be a semiabelian variety over F, $X \leq S$ a subvariety also over F. Let $\Gamma \leq S(F)$ be a finite rank subgroup. Then $X(F) \cap \Gamma$ is a finite union of translates of subgroups of* Γ*.*

As observed by Abramovich and Voloch in [1], this statement fails in positive characteristic:

Proposition 2.8. *Let F be an algebraically closed field of characteristic $p > 0$ such that $F \neq \mathbb{F}_p^{\text{alg}}$. Let C be a curve over $\mathbb{F}_p$ with an $\mathbb{F}_p$-point and with* $\text{genus}(C) > 1$*. Let A denote the Jacobian variety of C, and let $K = \mathbb{F}_p(t)$ for some $t \in C(F) \setminus C(\mathbb{F}_p^{\text{alg}})$. Then $A(K)$ is a finite rank subgroup of $A(F)$, but $C(K) = C(F) \cap A(K)$ is not a finite union of translates of subgroups of K.*

Proof. A is defined over $\mathbb{F}_p$, and K is a finitely generated extension of $\mathbb{F}_p$, so the Lang-Néron theorem (see [15, Theorem 6.1]) says that $A(K)$ is finitely generated. Consider the Frobenius automorphism $\text{Fr} : F \to F$ given by $\text{Fr}(x) = x^p$. Since A and C are defined over $\mathbb{F}_p$ Fr acts on both. Since $t \notin C(\mathbb{F}_p^{\text{alg}})$ we get an infinite collection of distinct points $t, \text{Fr}(t), \text{Fr}^2(t), \ldots \in C(K)$, so $C(K)$ is infinite. Suppose for contradiction that $C(K)$ is a finite union of translates of subgroups of $A(K)$. Then since $C(K)$ is infinite it is Zariski-dense in C, so as before, C admits an algebraic group structure, and hence $\text{genus}(C) = 1$, contradicting our hypothesis. □

So the absolute Mordell-Lang conjecture fails in positive characteristic. We can, however, weaken the conclusion to yield a correct statement.

Definition 2.9. Let S be a semiabelian variety over an algebraically closed field F. Let $X \subseteq S$ be a subvariety over F, and let $k \subseteq F$ be an algebraically closed subfield. Suppose that there exist a semiabelian variety S_0 over k, a subvariety X_0 also over k, an algebraic subgroup $S' \subseteq S$, a surjective morphism of algebraic groups $h : S' \to S_0$ over F and a point $c \in S(F)$ such that $X = c + h^{-1}(X_0)$. Then we say that X is *k-special.* If X is $\mathbb{F}^{\mathrm{alg}}$-special, where $\mathbb{F}$ is the prime field, then we say that X is *special.*

We observe that special subvarieties generalize translates of algebraic subgroups:

Lemma 2.10. *Let S be a semiabelian variety over an algebraically closed field F, and let $X = d + S'$ where S' is an algebraic subgroup of S over F. Then X is special.*

Proof. Take $c = d$ and $h : S' \to 0$ in the definition of special. □

We consider two extreme cases:

Example 2.11. Let S be a semiabelian variety over an algebraically closed field F with prime field $\mathbb{F}$. Suppose that S has $\mathbb{F}^{\mathrm{alg}}$-trace 0, that is, no algebraic subgroup of S has a non-trivial regular homomorphic image over $\mathbb{F}^{\mathrm{alg}}$. Then for any subvariety $X \subseteq S$, X is special $\iff$ X is a translate of an algebraic subgroup of S.

Proof. If X is a translate of an algebraic subgroup of S then X is special by Lemma 2.10. Suppose that X is special. Then by hypothesis the $h : S' \to S_0$ in the definition of special must be the trivial map onto 0, and so $h^{-1}(X_0) = S'$, an algebraic subgroup of S. Then $X = c + h^{-1}(X_0) = c + S'$ is a translate of an algebraic subgroup of S. □

Example 2.12. Let F be an algebraically closed field with prime field $\mathbb{F}$. Let S be a semiabelian variety defined over $\mathbb{F}^{alg}$. Then for any subvariety $X \subseteq S$, X is special $\iff$ X is the translate of a subvariety of S which is also defined over $\mathbb{F}^{\mathrm{alg}}$.

Proof. If X is the translate of a subvariety of S over $\mathbb{F}^{\mathrm{alg}}$, say $X = c + X'$, then let $S_0 = S' = S$, $X_0 = X'$, and $h = \mathrm{id}$. Conversely, suppose that X is special. Then $X = c + h^{-1}(X_0)$ as in the definition of special, where X_0 is over $\mathbb{F}^{\mathrm{alg}}$. Since h is a morphism of algebraic groups the graph of h, $\Gamma(h)$, is an algebraic subgroup of $S \times S_0$. We show in Appendix B.2 that semiabelian varities are rigid (Lemma B.8), so since $S \times S_0$ is a semiabelian variety over $\mathbb{F}^{\mathrm{alg}}$ $\Gamma(h)$ is defined over $\mathbb{F}^{\mathrm{alg}}$, and so h is over $\mathbb{F}^{\mathrm{alg}}$. Then $h^{-1}(X_0)$ is also a subvariety of S defined over $\mathbb{F}^{\mathrm{alg}}$, so X is a translate of a subvariety of S over $\mathbb{F}^{\mathrm{alg}}$, as required. □

We now state the version of the Mordell-Lang Conjecture that Hrushovski proved in 1996 in [10], and which is the subject of the remainder of this thesis.

Theorem 2.13 (Relative Mordell-Lang Conjecture)**.** *Let S be a semiabelian variety over an algebraically closed field F of characteristic p (either prime or 0). Let $X \subseteq S$ be a subvariety over F. Let $\Gamma' \leq S(F)$ be a finitely generated group. Define*

$$\operatorname{div}_p(\Gamma') = \begin{cases} \operatorname{div}(\Gamma') = \{g \in S(F) : ng \in \Gamma' \textit{ for some } n > 0\} & \textit{if } p = 0 \\ \{g \in S(F) : ng \in \Gamma' \textit{ for some } n \textit{ such that } p \nmid n\} & \textit{if } p > 0 \end{cases}$$

Let $\Gamma \leq \operatorname{div}_p(\Gamma')$. Then for some $l \in \mathbb{N}$ there exist special subvarieties $X_1, \ldots, X_l$ of S such that $X_i \subseteq X$ for all i, and

$$X(F) \cap \Gamma = \bigcup_{i=1}^{l} (X_i(F) \cap \Gamma)$$

Observe that, by Example 2.11, if S has $\mathbb{F}^{\text{alg}}$-trace 0 then the conclusion of Theorem 2.13 says that the X_i's are translates of algebraic subgroups of S, so the conclusion is exactly the same as the conclusion of Theorem 2.7. That is, even in characteristic p, when S has $\mathbb{F}_p^{\text{alg}}$-trace 0 we get the full strength conclusion. In the other extreme, when S is defined over $\mathbb{F}_p^{\text{alg}}$, Theorem 2.13 only reduces the problem of describing $X(F) \cap \Gamma$ to the case when X is (up to translation) also defined over $\mathbb{F}_p^{\text{alg}}$. See [22] for an analysis of that case.

2.2 Model-Theoretic Framework

In this section we discuss the framework for the proof of Mordell-Lang. In characteristic 0 the appropriate first-order theory is differentially closed fields, while in characteristic $p > 0$ we use separably closed fields. Since good references for the material we will need are readily available we present them without proof. For the proofs in the case of differentially closed fields, see [17] and [31]. For separably closed fields, see [7] and [18]. It is possible to unify these two theories by using the theory of Hasse closed fields (see [19]), but we have chosen to avoid developing the Hasse formalism in favour of using the more familiar differentially and separably closed fields. We will still be able to proceed in a characteristic-free way for most of the proof, since the properties we will use are shared by these two theories.

2.2.1 Characteristic 0 - Differentially Closed Fields

All of the results in this section can be found in [31].

Let R be a ring. A *derivation* on R is an additive map $\delta : R \to R$ that satisfies the product rule $\delta(xy) = x\delta(y) + y\delta(x)$. A *differential ring* is a pair (R, δ) where R is a ring and δ is a derivation on R. A *differential field* is a differential ring (R, δ) where R is a field. The *constants* of R are the elements of $k_R = \{x \in R : \delta(x) = 0\}$. The constants form a subring of R. The *differential polynomial ring* over R in the (differential) variable X is the differential ring $(R\{X\}, \delta_0)$ where $R\{X\} = R[X_0, X_1, \ldots]$ and δ_0 extends δ by $\delta_0(X_n) = X_{n+1}$. We identify δ with δ_0, X_0 with X, and X_n with $\delta^n(X)$. For a differential polynomial $f \in R\{X\}$ we define the *order* of f to be -1 if $f \in R$, and otherwise to be the largest n such that $\delta^n(X)$ appears in f.

Let $\mathcal{L}$ be the language of rings together with a new unary function symbol δ.

Definition 2.14. The *theory of differentially closed fields of characteristic* 0 (DCF_0) is axiomatized as follows:

1. The axioms for ACF_0.
2. $\forall x, y\ \delta(x + y) = \delta(x) + \delta(y)$.
3. $\forall x, y\ \delta(xy) = x\delta(y) + y\delta(x)$.
4. For any non-constant differential polynomials $f(X), g(X)$ such that the order of g is less than the order of f there exists x such that $f(x) = 0$ and $g(x) \neq 0$.

DCF_0 admits quantifier elimination and elimination of imaginaries. It is a complete, model-complete, and ω-stable theory. It is the model-completion of the theory of differential fields in characteristic 0, so in particular any differential field in characteristic 0 extends to a differentially closed field.

The constant field k in a model $L \models \mathrm{DCF}_0$ is a definable algebraically closed field. By stability of DCF_0 k is stably embedded in the sense that any subset of k^n definable using parameters from L can be defined using parameters from k. It is also a pure field, meaning that any subset of k^n definable in the language $\mathcal{L}$ of differential rings is also definable in the language of rings, which does not include the derivation. It follows that k is a strongly minimal set in L. Moreover, up to definable isomorphism k is the unique infinite minimal field type-definable in L.

Independence in DCF_0, in the sense Section 1.4, can be understood in purely algebraic terms. If $F < K$ are differential subfields of a saturated model $L \models \mathrm{DCF}_0$ and a is a finite tuple from L, then $a \mathop{\smile\!\!\!\!\!\!\,|}_F K$ if and only if $F\langle a\rangle_\delta = F(a, \delta(a), \delta^2(a), \ldots)$ is algebraically disjoint from K over F.

Although we will not use it directly, it is worth noting that there is a well-developed algebraic geometry over differential fields (see [17]).

2.2.2 Characteristic p - Separably Closed Fields

In this section we work in some fixed characteristic $p > 0$. We review some basic facts about fields in positive characteristic in order to define the theory of separably closed fields. Throughout Fr denotes the Frobenius automorphism $x \mapsto x^p$. All of the results of this section can be found in [7].

Definition 2.15. Let K be a field, and let $K^p = \{x^p : x \in K\}$, a subfield of K. We say K is *perfect* if $K = K^p$.

We note that it will be clear from context when we use K^p whether we mean the set of p-tuples of elements from K or $\{x^p : x \in K\}$. If K is a field then for any $x \in K$ there exists a unique (since Fr is injective) $y \in K^{\text{alg}}$ such that $y^{p^n} = x$. We denote $y = x^{p^{-n}}$, and let $K^{p^{-n}} = \left\{x^{p^{-n}} : x \in K\right\}$. We set $K^{p^{-\infty}} = \bigcup_{i=1}^{\infty} K^{p^{-n}}$, and call $K^{p^{-\infty}}$ the *perfect closure* of K.

Definition 2.16. Let K be a field. A polynomial over K is called *separable* if each of its irreducible factors has distinct roots. Let x be algebraic over K. We say x is *separable over* K if its minimal polynomial over K is separable. An algebraic extension $K \subseteq L$ is *separable* if every element is separable. We say that x is *purely inseparable over* K if its minimal polynomial is of the form $X^{p^n} - a$ for some $a \in K \setminus K^p$.

The *separable closure* of K is $K^{\text{sep}} = \left\{x \in K^{\text{alg}} : x \text{ is separable over } K\right\}$. K is *separably closed* if $K = K^{\text{sep}}$.

Definition 2.17. Let K be a field of characteristic $p > 0$, $A, B \subseteq K$ some subsets, and $x \in K$. x is *p-independent over A in K* if $x \notin K^p(A)$. B is *p-free over A* if every $b \in B$ is p-independent over $A \cup B \setminus \{b\}$ in K. If we only say x is p-independent or p-free we mean p-independent or p-free over $\emptyset$. B is said to *p-generate* K if $K \subseteq K^p(B)$.

Let K be a field of characteristic $p > 0$ and $B \subseteq K$. Then the following are equivalent: B is a minimal p-generating set; B is a maximal p-free set; B is p-free and p-generates K. In this case we say that B is a *p-basis* of K. Any two p-bases have the same cardinality, and this cardinal is called the *degree of imperfection* of K. We can use p-bases to form linear bases for K as a K^p vector space, as follows. Suppose that $B = \{b_i : i \in I\} \subseteq K$ is a p-basis. For any $j : I \to \{0, 1, \ldots, p-1\}$ with finite support let $m_j = \prod_{i \in I} b_i^{j(i)}$. The set of all such m_j is a linear basis for K as a K^p vector space. Hence any $x \in K$ has a unique expansion as a K^p-linear combination of the m_j's. We call the coefficients of such an expansion the *p-components* of x (with respect to B).

Now fix a $\nu \in \mathbb{N}$. Let $\mathcal{L}$ denote the language of rings, $\mathcal{L}_\nu = \mathcal{L} \cup \{b_1, \ldots, b_\nu\}$ for some new constant symbols $b_1, \ldots, b_\nu$, and $\mathcal{L}_{p,\nu} = \mathcal{L}_\nu \cup \{\lambda_i : i < p^\nu\}$ where the λ_i's are new unary function symbols. $\mathrm{SC}_{p,\nu}$ is the theory of separably closed fields of characteristic p and degree of imperfection ν in the language $\mathcal{L}$ of rings. In the language $\mathcal{L}_\nu$ let $\mathrm{S}_{p,\nu}$ be the theory $\mathrm{SC}_{p,\nu}$ together with axioms expressing that the b_i's form a p-basis. In $\mathcal{L}_{p,\nu}$ let $\mathrm{SCF}_{p,\nu}$ be $\mathrm{S}_{p,\nu}$ together with axioms expressing that each λ_i maps x to its ith p-component (under some fixed ordering of the monomials m_j). Then $\mathrm{SCF}_{p,\nu}$ is complete, model-complete, and stable. It admits quantifier elimination and elimination of imaginaries.

Unlike the case of differentially closed fields, the theory $\mathrm{SCF}_{p,\nu}$ is not ω-stable (or even superstable). As a consequence, since we wish to present a characteristic-free proof of Mordell-Lang, we will be making use of general stability machinery even though DCF_0 has the much stronger property of ω-stability.

There is a notion of the constants of a model of $\mathrm{SCF}_{p,\nu}$:

Definition 2.18. Let $K \models \mathrm{SCF}_{p,\nu}$. Let $k = K^{p^\infty} = \bigcap_{i=1}^{\infty} K^{p^n}$. We call k the *constants* of K.

Since each K^{p^n} is a 0-definable set in K it follows that the constants form a type-definable set in K over $\emptyset$. They share many model-theoretic properties with the constants in DCF_0:

Fact 2.19. *1. Let $K \models \mathrm{SCF}_{p,\nu}$ be sufficiently saturated. Then the constants k form a stably embedded pure algebraically closed subfield of K. It is a minimal type-definable field, and is the unique (up to definable isomorphism) minimal field type-definable in K.*

2. Let F, K be elementary submodels of a saturated model $L \models \mathrm{SCF}_{p,\nu}$ with $F \subseteq K$, and let $a \in L$. Let $F\langle a\rangle_\lambda$ denote the field generated by F and all of the p-components of a. Then $a \mathop{\smile\!\!\!\!\!\mid}_F K$ if and only if $F\langle a\rangle_\lambda$ and K are algebraically disjoint over K.

Chapter 3

The Zilber Dichotomy

In this thesis we show how the relative Mordell-Lang statement (Theorem 2.13) follows from a model-theoretic result known as the Zilber Dichotomy. In this chapter we will explain what the Zilber Dichotomy says for differentially and separable closed fields. For an account of Zariski geometries, which are the general setting in which the Zilber Dichotomy holds, see [16]. The proof of the dichotomy theorem in that setting is due to Hrushovski and Zilber, and may be found in [12].

Throughout this chapter the following conventions will be in force. T will denote either DCF_0 or $\mathrm{SCF}_{p,\nu}$, where p is prime and ν is a positive integer. We work in a universal domain L of T, and all parameter sets are thus assumed to be of cardinality strictly less than $|L|$. $k \subseteq L$ denotes the field of constants in L - that is, $k = \{x \in L : \delta(x) = 0\}$ in the case of DCF_0 or $k = L^{p^\infty}$ in the case of $\mathrm{SCF}_{p,\nu}$. $\mathbb{F} \subseteq L$ denotes the prime field.

3.1 The Dichotomy in Differentially and Separably Closed Fields

Before stating the dichotomy we will need to develop the model-theoretic notions of orthogonality and one-basedness.

3.1.1 One-Basedness

Definition 3.1. Let X be a type-definable set over A. We say that X is *one-based* if for all $a \in \mathrm{dcl}(X \cup A)$ and any set $B \supseteq A$, $\mathrm{cb}(a/B) \subseteq \mathrm{acl}(Aa)$.

The previous definition appears to depend on the parameter set A, but this dependence is illusory (see [25, Section 4.4 and Remark 4.1.8]). It is useful to introduce the notion of linearity as a means of explaining one-basedness for minimal sets, including the constant field k:

Definition 3.2. Let X be a minimal type-definable set. We say that X is *linear* if whenever $(a, b) \in X \times X$ and $C = \mathrm{cb}(ab/C)$ are such that $U(ab/C) = 1$ then $U(C) \leq 1$.

A minimal set X is one-based if and only if it is linear (see [32, Theorems 5.12 and 5.14]). Viewing $\mathrm{tp}(ab/C)$ in the definition of linearity as a "plane curve" in X^2, we can thus understand one-basedness as expressing that X fails to have any rich definable families of plane curves - all definable families of plane curves are "one-parameter" families.

Proposition 3.3. *The constant field k is not one-based.*

Proof. Let $c_1, c_2 \in k$ be algebraically independent elements transcendental over $\mathbb{F}^{\mathrm{alg}}$. Consider the line $Y = c_1X + c_2$. Intuitively this should contradict linearity, since $Y = c_1X + c_2$ is a two-parameter definable family of plane curves. Let (a, b) be a generic solution to $Y = c_1X + c_2$. Then $\mathrm{tp}(ab/c_1c_2)$ is of U-rank 1.

We show that $C = c_1c_2$ is the canonical base of $p = \mathrm{tp}(ab/C)$. Let α be any automorphism of L. Clearly if α fixes C pointwise then α fixes p^L setwise. Conversely, suppose that p and $\alpha(p)$ have a common non-forking extension q to some set D. We have "$Y = c_1X+c_2$" $\in p$ and "$Y = \alpha(c_1)X + \alpha(c_2)$" $\in \alpha(p)$, so both of these formulae are in q. Thus any realization of q must satisfy $(c_1 - \alpha(c_1))X + (c_2 - \alpha(c_2)) = 0$. Since q is realized (by saturation of L), we compare coefficients to conclude that $c_1 = \alpha(c_1)$ and $c_2 = \alpha(c_2)$, so α fixes C pointwise, and hence C is a canonical base of p.

But $U(C) = \dim(c_1c_2) = 2$ since c_1, c_2 are algebraically independent. This contradicts the definition of linearity. (One could also show that $c_1c_2 \notin \mathrm{acl}(a, b)$ and thus contradict one-basedness directly.) □

One-basedness for groups has very strong consequences, which we will describe in the next section.

3.1.2 Full Orthogonality

Definition 3.4. Let X and Y be type-definable sets. We say that X is *fully orthogonal* to Y, denoted $X \perp Y$, if for every $a \in X$, every $b \in Y$, and every parameter set A over which both X and Y are defined, we have $a \mathop{\smile\!\!\!\!\!\!|}_A b$.

The above definition is not entirely standard. It appears in Hrushovski's proof of Mordell-Lang ([10]), and is well-suited to our purposes. To avoid confusion with other notions of orthogonality we will use the above notion of full orthogonality exclusively throughout this thesis; that is, whenever we speak of types being orthogonal or non-orthogonal, we mean it in the above sense. Full orthogonality is a way of expressing that the sets X and Y are very much unrelated. In fact, $X \perp Y$ implies that every tuple from X is independent of every tuple from Y (over any set of parameters defining X and Y):

Lemma 3.5. *Let $X_1, \ldots, X_m$ and $Y_1, \ldots, Y_n$ be type-definable sets such that $X_i \perp Y_j$ for all $1 \leq i \leq n, 1 \leq j \leq m$. Then $X_1 \times \cdots \times X_m \perp Y_1 \times \cdots \times Y_n$.*

Proof. We begin by observing that the relation $\perp$ is symmetric since non-forking independence is symmetric. It thus suffices to consider the case when $m = 1$. The proof is by induction on n. The base case $n = 1$ is exactly the hypothesis of the lemma. Suppose that $X \perp Y_1 \times \cdots \times Y_l$ for all $l < n$. Consider any $x \in X$ and any $(y_1, \ldots, y_n) \in Y_1 \times \cdots \times Y_n$, and any parameter set A over which X and the Y_i's are defined. We know $x \mathop{\smile\!\!\!\!\!\!\mid}_A y_n$, and X and Y_i are defined over Ay_n for all $1 \leq i \leq n-1$, so the induction hypothesis gives $x \mathop{\smile\!\!\!\!\!\!\mid}_{Ay_n} (y_1, \ldots, y_{n-1})$. By transitivity of non-forking independence $x \mathop{\smile\!\!\!\!\!\!\mid}_A (y_1, \ldots, y_n)$ as required. $\square$

We now present a couple of results about orthogonality that further illuminate its meaning. First, among minimal sets non-orthogonality is a transitive relation. Next, we give several alternative formulations of full orthogonality:

Lemma 3.6. *Let X and Y be type-definable sets in some sufficiently saturated model. Then the following are equivalent:*

1. $X \perp Y$
2. *For any set $A = \operatorname{acl}(A)$ over which X and Y are defined, and any $a \in X$, $b \in Y$,* $\operatorname{tp}(a/A) \cup \operatorname{tp}(b/A) \vdash \operatorname{tp}(ab/A)$
3. *For any set $A = \operatorname{acl}(A)$ over which X and Y are defined, and any $a \in X$, $b \in Y$, the only extension of* $\operatorname{tp}(a/A)$ *to a complete type over Ab is* $\operatorname{tp}(a/Ab)$.

Proof. (1) ⇒ (3): Suppose that $X \perp Y$, and let $A = \operatorname{acl}(A)$ be parameters over which X and Y are defined. Take any $a \in X$, and any $b \in Y$. Let q be a complete type over Ab extending $\operatorname{tp}(a/A)$. Since we are working in a sufficiently saturated model, $q = \operatorname{tp}(c/Ab)$ for some c. Since q extends $\operatorname{tp}(a/A)$ we have $\operatorname{tp}(c/A) = \operatorname{tp}(a/A)$, so in particular $c \in X$. By definition of $X \perp Y$ we have $c \mathop{\smile\!\!\!\!\!\!\mid}_A b$, so $q = \operatorname{tp}(c/Ab)$ is a non-forking extension of $\operatorname{tp}(a/A)$

to Ab. This shows that every extension of $\mathrm{tp}(a/A)$ to a complete type over Ab is non-forking. Since A is algebraically closed there is a unique non-forking extension of $\mathrm{tp}(a/A)$ to Ab, and hence a unique extension of $\mathrm{tp}(a/A)$ to Ab, which then must be $\mathrm{tp}(a/Ab)$.

(3) ⇒ (1): Let A be parameters over which X and Y are defined, and take any $a \in X$, $b \in Y$. By forking calculus it suffices to consider the case when $A = \mathrm{acl}(A)$. Then by (3) and existence of non-forking extensions we get that $\mathrm{tp}(a/Ab)$ is a non-forking extension of $\mathrm{tp}(a/A)$, that is, $a \mathop{\smile\!\!\!\!\!\mid}_A b$, as required to show $X \perp Y$.

(2) ⇒ (3): Let $A = \mathrm{acl}(A)$ be parameters over which X and Y are defined, $a \in X$, $b \in Y$. Let p be any extension of $\mathrm{tp}(a/A)$ to Ab, and let c be a realisation of p, so $\mathrm{tp}(c/Ab) = p$. As p extends $\mathrm{tp}(a/A)$, c also realises $\mathrm{tp}(a/A)$. Hence by (2), we have $cb \models \mathrm{tp}(ab/A)$, that is, $\mathrm{tp}(cb/A) = \mathrm{tp}(ab/A)$. We have, for any $\mathcal{L}_A$-formula $\phi(x, y)$,

$$\begin{aligned}\phi(x,b) \in \mathrm{tp}(c/Ab) &\iff \phi(x,y) \in \mathrm{tp}(cb/A)\\ &\iff \phi(x,y) \in \mathrm{tp}(ab/A)\\ &\iff \phi(x,b) \in \mathrm{tp}(a/Ab)\end{aligned}$$

So $p = \mathrm{tp}(c/Ab) = \mathrm{tp}(a/Ab)$.

(3) ⇒ (2): Let $A = \mathrm{acl}(A)$ be parameters over which X and Y are defined, $a \in X$, $b \in Y$, c, d such that $c \models \mathrm{tp}(a/A)$ and $d \models \mathrm{tp}(b/A)$. We must show $cd \models \mathrm{tp}(ab/A)$. Since $\mathrm{tp}(b/A) = \mathrm{tp}(d/A)$ there exists $\sigma \in \mathrm{Aut}_A(\mathcal{M})$ such that $\sigma(b) = d$. Since σ fixes A pointwise $\sigma(a) \in X$ and $\mathrm{tp}(\sigma(a)/A) = \mathrm{tp}(a/A) = \mathrm{tp}(c/A)$. Then $\mathrm{tp}(c/Ad)$ and $\mathrm{tp}(\sigma(a)/Ad)$ are both extensions of $\mathrm{tp}(c/A)$ to Ad, so by (3) $\mathrm{tp}(c/Ad) = \mathrm{tp}(\sigma(a)/Ad)$. Thus there exists an automorphism $\tau \in \mathrm{Aut}_{Ad}(\mathcal{M})$ such that $\tau(\sigma(a)) = c$. We observe that $\tau \circ \sigma \in \mathrm{Aut}_A(\mathcal{M})$, and:

$$\begin{aligned}\tau \circ \sigma(a, b) &= (\tau(\sigma(a)), \tau(\sigma(b)))\\ &= (\tau(\sigma(a)), \tau(d)) && \text{by our choice of } \sigma.\\ &= (c, d) && \text{by our choice of } \tau.\end{aligned}$$

So we found an automorphism fixing A and sending ab to cd, so $\mathrm{tp}(ab/A) = \mathrm{tp}(cd/A)$ as required. □

Proposition 3.7. *Let X and Y be type-definable sets in a sufficiently saturated model $\mathcal{U}$, and suppose that $X \perp Y$. Then any definable subset $R \subseteq X \times Y$ is of the form $R = \bigcup_{i=1}^{n} X_i \times Y_i$ where the X_i's and Y_i's are definable in X and Y, respectively.*

Proof. Suppose that X and Y are both defined over some parameter set A, by types p and q, respectively. Now take any definable $R \subseteq X \times Y$, defined over $B = \mathrm{acl}(B) \supseteq A$. Let $\Sigma = p \cup q \cup \{\phi\}$ be such that $R = \Sigma^{\mathcal{U}}$.

Take any $(a, b) \in R$. Then since $X \perp Y$, the equivalence in Lemma 3.6 gives that $\mathrm{tp}(a/B) \cup \mathrm{tp}(b/B) \vdash \mathrm{tp}(ab/B)$. In particular, $\mathrm{tp}(a/B) \cup \mathrm{tp}(b/B) \vdash \Sigma$. By compactness there are $\psi_1 \in \mathrm{tp}(a/B)$ and $\psi_2 \in \mathrm{tp}(b/B)$ such that $p \cup q \cup \{\psi_1, \psi_2\} \vdash \Sigma$. Let $R_1 = \psi_1^{\mathcal{U}}, R_2 = \psi_2^{\mathcal{U}}$. Then we have that R_1 is definable in X, and R_2 is definable in Y, and the above shows that $R_1 \times R_2 \subseteq R$.

Repeating the above process for each $(a, b) \in R$ gives a cover $R = \bigcup_{i \in I} X_i \times Y_i$ where the X_i's and Y_i's are definable in X and Y, respectively. By compactness we get a finite subcover $R = \bigcup_{i=1}^{n} X_i \times Y_i$, as desired. □

Proposition 3.8. *Let X and Y be infinite type-definable sets in a strongly minimal theory T. Then $X \not\perp Y$.*

Proof. Let A be a set of parameters over which both X and Y are defined. Let $\mathcal{M} \models T$ be a sufficiently large saturated model, and suppose that $X \subseteq M^n$, $Y \subseteq M^m$. We recall that, since T is strongly minimal, forking is determined by acl-dimension.

We first show that there is a tuple $(a_1, \ldots, a_n) \in X$ such that some $a_i \in M$ is generic over A. Since T is strongly minimal we have that $c \in M$ is generic over A if and only if $c \notin \mathrm{acl}(A)$. Let $\Phi(x_1, \ldots, x_n)$ be the type defining X. Let $\Psi(x_1, \ldots, x_n)$ be the collection consisting of the negations of all the $\mathcal{L}_A$-formulae with only finitely many realizations. Then since X is infinite $\Phi \cup \Psi$ is finitely satisfiable, and hence satisfiable by compactness. Let $(a_1, \ldots, a_n)$ be a realization. For a contradiction suppose that $a_1, \ldots, a_n$ are all algebraic over A, and let $\phi_1(x_1), \ldots, \phi_n(x_n)$ be L_A-formulae witnessing this. Then $(a_1, \ldots, a_n)$ realizes $\wedge_{i=1}^{n} \phi_i(x_i)$, which has only finitely many realizations, contradicting our choice of $(a_1, \ldots, a_n)$. So for some i we have that $a_i \notin \mathrm{acl}(A)$. Since changing the order of variables is an automorphism we may assume that $i = 1$.

In light of the above, let $a = (a_1, \ldots, a_n) \in X$ be such that (without loss of generality) $a_1 \notin \mathrm{acl}(A)$. Similarly, let $b = (b_1, \ldots, b_m) \in Y$ be such that $b_i \notin \mathrm{acl}(A)$. By the uniqueness of generic types $\mathrm{tp}(a_1/A) = \mathrm{tp}(b_i/A)$. There is thus some A-automorphism α such that $\alpha(a_1) = b_i$. For each $2 \leq j \leq n$ let $a'_j = \alpha(a_j)$. Since X is defined over A we have $\alpha(X) = X$, and so $(b_i, a'_2, \ldots, a'_n) \in X$. Moreover, $(b_i, a'_2, \ldots, a'_n)$ is also non-algebraic over A. Let $(b_i, c_1, \ldots, c_r)$ be an acl-basis for $(b_i, a'_2, \ldots, a'_n)$ over A. Then $\text{acl-dim}((b_i, a'_1, \ldots, a'_n)/A) = r$, while $\text{acl-dim}((b_i, a'_1, \ldots, a'_n)/A(b_1, \ldots, b_m)) \leq r - 1$, and hence $(b_i, a'_1, \ldots, a'_n) \not\downarrow_A (b_1, \ldots, b_m)$, so $X \not\perp Y$. □

3.1.3 The Dichotomy Statement

With the notions of one-basedness and full orthogonality at our disposal, we can now state the Zilber Dichotomy Theorem for our separably closed or differentially closed field

L. Recall that the constants k are not one-based. The dichotomy says that, up to non-orthogonality, k is the only non-one-based minimal set:

Theorem 3.9 (Zilber Dichotomy)**.** *Let X be a minimal type-definable set in L. Then either X is one-based or $X \not\perp k$.*

For a proof of Theorem 3.9 see [10, Lemma 5.4] for the case $L \models \mathrm{SCF}_{p,\nu}$ and [11, Corollary 1.20] for the case $L \models \mathrm{DCF}_0$. We end this chapter by showing how the dichotomy theorem can be extended to the wider class of semiminimal sets.

Definition 3.10. A type-definable set E is *semiminimal* if it is infinite and there is a minimal set D and a finite set F such that $E \subseteq \mathrm{acl}(F \cup D)$.

Corollary 3.11. *Let E be a semiminimal type-definable set in L. Then either E is one-based or $E \not\perp k$.*

Proof. Let $E \subseteq \mathrm{acl}(F \cup D)$ with F finite and D minimal. We proceed by proving two claims.

Claim 3.11.1. *$E \perp k$ if and only if $D \perp k$.*

Proof of Claim 3.11.1 First, suppose that $E \perp k$. Suppose towards a contradiction that $D \not\perp k$, so there exists $d \in D$, $a \in k$, and some set of parameters A such that $d \not\mathrel{\unicode{x2AEB}}_A a$. By taking non-forking extensions of $\mathrm{tp}(d/A)$ and $\mathrm{tp}(a/A)$ to $A \cup F$ we may assume that $F \subseteq A$. Since D is minimal, $d \not\mathrel{\unicode{x2AEB}}_A a$ says $d \in \mathrm{acl}(Aa) \setminus \mathrm{acl}(A)$. In particular, d is generic in D over A. Minimality of D implies that D has a unique generic type over A, so given any generic $d' \in D$ over A there is an automorphism fixing A and sending d to d'. Let a' be the image of a under such an automorphism. Then $d' \in \mathrm{acl}(Aa')$ and $a' \in k$.

Let $e \in E$ be generic over A. Then since $E \subseteq \mathrm{acl}(F \cup D)$ and $F \subseteq A$ there is a tuple $\overline{d}$ from D such that $e \in \mathrm{acl}(A\overline{d})$. Then since e is generic over A, $e \notin \mathrm{acl}(A)$. Now each d_i from $\overline{d}$ is either in $\mathrm{acl}(A)$ or is generic in D over A, and hence is in $\mathrm{acl}(Aa_i)$ for some a_i from k. Let $\overline{a}$ be the tuple of a_i's. Then $e \in \mathrm{acl}(A\overline{a}) \setminus \mathrm{acl}(A)$, which implies $e \not\mathrel{\unicode{x2AEB}}_A \overline{a}$. By Lemma 3.5, $E \not\perp k$.

For the converse, suppose that $D \perp k$. Let A be parameters over which E and k are defined, and take $e \in E, a \in k$. By taking non-forking extensions of $\mathrm{tp}(e/A)$ and $\mathrm{tp}(a/A)$ to $A \cup F$ we may assume $F \subseteq A$. Since $E \subseteq \mathrm{acl}(D \cup F)$ there is a tuple d from D such that $e \in \mathrm{acl}(d \cup F)$. Then since $D \perp k$ we have $d \mathrel{\unicode{x2AEB}}_A a$. Clearly $F \mathrel{\unicode{x2AEB}}_A a$. Hence by transitivity and symmetry for non-forking we get $a \mathrel{\unicode{x2AEB}}_A dF$, and so $a \mathrel{\unicode{x2AEB}}_A e$ as well. So $E \perp k$. $\dashv$

Claim 3.11.2. *If D is one-based then E is one-based.*

Proof of Claim 3.11.2 First, for notational convenience, we name the set F, along with any parameters used to define D and E, to the language, and hence assume that $F = \emptyset$ and D, E are defined over $\emptyset$. Take any $a \in \operatorname{dcl}(E)$, and let C be such that $C = \operatorname{cb}(a/C)$. We want to show that $C \subseteq \operatorname{acl}(a)$.

Since $a \in \operatorname{acl}(D)$ we have $a \in \operatorname{acl}(c_1, \ldots, c_l)$ for some $c_1, \ldots, c_l$ from D. Since D is minimal some subset of $\{c_1, \ldots, c_l\}$ is an acl-basis for $c_1, \ldots, c_l$ over a. Reindexing if necessary, suppose that $c_1, \ldots, c_r$ is such a basis. Let $B = \{c_1, \ldots, c_r\}$, and let $b = (c_{r+1}, \ldots, c_l)$. Then $a \mathop{\smile\!\!\!\!\!\!\mid}\; B$ by definition of acl-basis, and we have, by the choice of $c_1, \ldots, c_l$, that $a \in \operatorname{acl}(Bb)$. Conversely, by definition of acl-basis we have $b \in \operatorname{acl}(Ba)$. So $\operatorname{acl}(Ba) = \operatorname{acl}(Bb)$.

Let A realize a non-forking extension of $\operatorname{tp}(B/a)$ to Ca. Then we still have $a \mathop{\smile\!\!\!\!\!\!\mid}\; A$, and, constructing b' in the same manner as b above, we again get $\operatorname{acl}(Aa) = \operatorname{acl}(Ab')$, but now also $A \mathop{\smile\!\!\!\!\!\!\mid}_a C$, so by forking calculus $A \mathop{\smile\!\!\!\!\!\!\mid}_C a$. Since $C = \operatorname{cb}(a/C)$ we thus have $C = \operatorname{cb}(a/AC)$ as well. Moreover, since a and b' are interalgebraic over A we also get that $\operatorname{tp}(b'/AC)$ does not fork over C. It follows from this, since canonical bases are the minimal non-forking base, that $\operatorname{acl}(C) = \operatorname{acl}(\operatorname{cb}(b'/AC))$. Hence $\operatorname{acl}(C) = \operatorname{acl}(\operatorname{cb}(b'/C))$ Now D is one-based, so $\operatorname{acl}(C) \subseteq \operatorname{acl}(b) \subseteq \operatorname{acl}(Ab') = \operatorname{acl}(Aa)$. In particular, $C \subseteq \operatorname{acl}(Aa)$. Since $F \mathop{\smile\!\!\!\!\!\!\mid}_a C$, we have $C \subseteq \operatorname{acl}(a)$. Hence E is one-based. $\dashv$

The dichotomy for semiminimal sets now follows immediately from the dichotomy for minimal sets. This proves Corollary 3.11. □

3.2 Consequences of One-Basedness

In this section we describe some consequences of one-basedness for type-definable groups. In the next section we will consider the consequences of the other half of the dichotomy, being non-orthogonal to the constants. Unlike the consequences of non-orthogonality to the constants, the results in this section are a standard part of the study of stable groups, so we omit many of the proofs. The reader desiring further background on one-based groups should consult [25] or [29]. The results of this section do not make use of separably or differentially closed fields, so we work temporarily in an arbitrary stable theory T. Let G be a group type-definable over parameters A. For convenience we assume that G is commutative, though the results stated here can be developed for non-commutative groups. We state first a useful characterization of one-basedness for type-definable groups.

Fact 3.12. *The following are equivalent:*

1. *G is one-based.*

2. *For every $n < \omega$, every definable subset of G^n is a boolean combination of cosets of* $\mathrm{acl}(A)$*-definable subgroups of G^n.*

3. *Given two complete types p and q in G, if p and q contain the same cosets of* $\mathrm{acl}(A)$*-definable subgroups of G^n then $p = q$.*

Corollary 3.13. *If G is one-based then every definable subgroup of G is also one-based, and defined over* $\mathrm{acl}(A)$*.*

What Fact 3.12 tells us is that the induced structure on a one-based group G comes entirely from the definable subgroups of G and the group structure. One particular, and somewhat technical, formulation of this is the following theorem, which states that a type in a one-based group is determined by its stabilizer (see Section 1.6 for a discussion of stabilizers).

Theorem 3.14. *Suppose that G is one-based, $B \supseteq A$, and $p(x) \in S(B)$ is a complete stationary type in G. Then* $\mathrm{stab}(p)$ *is type-definable over* $\mathrm{acl}(A)$*, and p is the generic type of a B-definable translate of* $\mathrm{stab}(p)$*. In particular,* $\mathrm{stab}(p)$ *is connected.*

A proof of the above theorem can be found in [25]. We end with the following useful fact:

Proposition 3.15. *Suppose that A and B are fully orthogonal one-based type-definable subgroups of G. Then $A + B$ is one-based.*

Proof. We begin by showing that $A \times B \leq G \times G$ is one-based, and later will use the map $+ : A \times B \to A + B$ to deduce that $A + B$ is as well. For convenience we name the parameters used in defining A, B, and G to the language.

By Proposition 3.7, since $A \perp B$, every definable subset $X \subseteq A \times B$ is a finite union of sets of the form $A' \times B'$ where $A' \subseteq A$ and $B' \subseteq B$ are definable. Since A and B are one-based groups Fact 3.12 gives that each A' and B' are boolean combinations of cosets of $\mathrm{acl}(\emptyset)$-definable subgroups of A and B, respectively. A straightforward induction on the complexity of the boolean combination shows that each $A' \times B'$ is then a boolean combination of cosets of $\mathrm{acl}(\emptyset)$-definable subgroups of $A \times B$, so $A \times B$ is one-based.

Now take any distinct complete types p and q in $A+B$. Let $a_1, a_2 \in A$ and $b_1, b_2 \in B$ be such that $+(a_1, b_1) \models p$ and $+(a_2, b_2) \models q$. Then $\mathrm{tp}(a_1b_1) \neq \mathrm{tp}(a_2b_2)$ since $+$ is definable and $p \neq q$. As $A \times B$ is one-based by the characterization in Fact 3.12 there is some $\mathrm{acl}(\emptyset)$-definable subgroup $H \leq A \times B$ and some r such that $r + H$ distinguishes $\mathrm{tp}(a_1b_1)$ from $\mathrm{tp}(a_2b_2)$. Then $+(r + H)$ is a coset of an $\mathrm{acl}(\emptyset)$-definable subgroup of $A + B$, namely $+(H)$, and it distinguishes p and q, so $A + B$ is one-based. □

3.3 Consequences of Non-orthogonality to the Constants

In this section we give concrete consequences of non-orthogonality to the constants, first for semiminimal sets, then for semiminimal groups, and finally for semiminimal subgroups of semiabelian varieties. We return to working in a universal domain L for DCF_0 or $\mathrm{SCF}_{p,\nu}$, with k the constant field.

3.3.1 Semiminimal Sets

Proposition 3.16. *Let Y be a semiminimal set type-definable over A, with $Y \not\perp k$, and p a nonalgebraic type in Y over A. Then there exists a set $B \supseteq A$, a non-forking extension $q \in S(B)$ of p, and a B-definable function $f : q^L \to k^n$ with finite fibres.*

Proof. Since Y is semiminimal there exists a finite set F and a minimal set X such that $Y \subseteq \mathrm{acl}(F \cup X)$.

Claim 3.16.1. *There is a parameter set B' such that $Y \subseteq \mathrm{acl}(B' \cup k)$.*

Proof of Claim 3.16.1 By Claim 3.11.1 we have $X \not\perp k$. Exactly as in the proof of that Claim we get a parameter set B', which we may assume contains F, such that if $s \in X$ is generic in X over B' then there is a tuple $\overline{a}$ from k such that $s \in \mathrm{acl}(B'\overline{a}) \subseteq \mathrm{acl}(B' \cup k)$. On the other hand, if $s \in X$ is not generic over B' then, since X is minimal, $s \in \mathrm{acl}(B')$. Hence $X \subseteq \mathrm{acl}(B' \cup k)$. Then since $B' \supseteq F$ and $Y \subseteq \mathrm{acl}(X \cup F)$ we also have $Y \subseteq \mathrm{acl}(B' \cup k)$. $\dashv$

Let a realize a non-forking extension of p to B'. Then $a \in \mathrm{acl}(B'c_1 \dots c_l)$ for some $c_1, \dots, c_l \in k$ by Claim 3.16.1. Since a realizes a non-forking extension of p to B', $a \notin \mathrm{acl}(B')$, and so $l \geq 1$. Since k is minimal, some subset of $c_1, \dots, c_l$ is an acl-basis for $\{c_1, \dots, c_l\}$ over $B'a$. After reindexing, suppose that $c_1, \dots, c_r$ is such a basis.

Let $c = (c_{r+1}, \dots, c_l)$, and let $B = B' \cup \{c_1, \dots, c_r\}$. Then $a \perp_{B'} B$, so $q = \mathrm{tp}(a/B)$ is a non-forking extension of p, and $\mathrm{acl}(Ba) = \mathrm{acl}(Bc)$.

Let $d_1, \dots, d_m$ be the Ba-conjugates of c (under automorphisms of L). Since any automorphism of L fixes k setwise, $d_1, \dots, d_m$ are tuples from k. Since $k \models \mathrm{ACF}_p$, which has elimination of imaginaries, there is some tuple $\widehat{c} \in k^n$ which is a code for $\{d_1, \dots, d_m\}$. Fix $\alpha \in \mathrm{Aut}_{Ba}(L)$. Then $\alpha(\{d_1, \dots, d_m\}) = \{d_1, \dots, d_m\}$ since $\{d_1, \dots, d_m\}$ is the orbit of c under $\mathrm{Aut}_{Ba}(L)$. Hence also $\alpha(\widehat{c}) = \widehat{c}$. So $\widehat{c} \in \mathrm{dcl}(Ba)$. Also, $c \in \mathrm{acl}(\widehat{c})$ since $c \in \{d_1, \dots, d_m\}$ and $\widehat{c}$ codes $\{d_1, \dots, d_m\}$. Hence $a \in \mathrm{acl}(B\widehat{c})$.

It follows that there is a B-definable map f with finite fibres such that $f(a) = \widehat{c}$. Since $a \in \operatorname{dom} f$, $\operatorname{tp}(a/B)^L \subseteq \operatorname{dom} f$. Let a' realise $q := \operatorname{tp}(a/B)$, so there exists an automorphism $\alpha \in \operatorname{Aut}_B(L)$ such that $\alpha(a) = a'$. Then

$$\begin{aligned} f(a') &= f(\alpha(a)) & \\ &= \alpha(f(a)) & \text{since } \alpha \text{ fixes } B \\ &= \alpha(\widehat{c}) & \\ &\in k^n & \text{since } \widehat{c} \in k^n \text{ and } \alpha \text{ restricts to an automorphism of } k \end{aligned}$$

Thus $f\mid_{q^L}: q^L \to k^n$ is the desired map. □

3.3.2 Semiminimal Groups

We now specialize further, to the case when our semiminimal set is also a type-definable group.

Proposition 3.17. *Let H be a semiminimal connected commutative type-definable group such that $H \not\perp k$. Then there exists a group G, definable in $(k, +, -, \cdot, 0, 1)$ and a definable surjective group homomorphism $h : H \to G$ with finite kernel.*

Proof. From Proposition 3.16 we have a set B over which H is defined, and a B-definable map with finite fibres $f : p^L \to L^n$ where p is the (unique, since H is connected) generic type of H over B. Note that p is stationary, since non-forking extensions of generic types are generic.

We first extend f to all of H: Let $D := \operatorname{dom} f \cap H$. Then we have $f : D \to L^n$ and $f(p^L) \subseteq k^n$. D is a generic definable subset of H as it contains p^L. So there exist $h_1, \dots, h_t \in H$ such that $H = (h_1 + D) \cup \ldots \cup (h_t + D)$. On each $h_i + D$ we have the map given by $f_i(x) = f(x - h_i)$, so we have definable maps with finite fibres $f_i : h_i + D \to L^n$. We can thus define a new $f : H \to L^n$ by $f(x) = f_i(x)$ where i is least such that $x \in h_i + D$. It is clear from the definition of the f_i's that f is well-defined, has finite fibres, and is definable over $Bh_1 \dots h_r$. We replace the f given by Lemma 3.16 with this f, and replace B by $Bh_1 \dots h_r$ to get the desired map defined on all of H. We note that, in general, f takes values in L^n, though it generically takes values in k^n.

We consider the following set:

$$N := \{h \in H : \text{for some } a \text{ generic over } Bh,\ f(a + h) = f(a)\}$$

Suppose that $h \in N$, so for some generic a over h we have $f(a + h) = f(a)$. Then since f is B-definable, we have that "$f(x + h) = f(x)$" $\in \operatorname{tp}(a/Bh)$. In particular, for any generic a' over Bh we have $\operatorname{tp}(a'/Bh) = \operatorname{tp}(a/Bh)$, and so $f(a' + h) = f(a')$. Thus we have

$$N = \{h \in H : \text{for all } a \text{ generic over } Bh,\ f(a + h) = f(a)\}$$

Claim 3.17.1. *N is a finite subgroup of H defined over B.*

Proof of Claim 3.17.1 For any a $f(a+0) = f(a)$, so $0 \in N$. Suppose that $h_1, h_2 \in N$, and let a be generic over Bh_1h_2. Then $a + h_1$ is generic over Bh_2. So we have

$$\begin{aligned} f(a+(h_1+h_2)) &= f((a+h_1)+h_2) \\ &= f(a+h_1) && \text{since } a+h_1 \text{ is generic over } Bh_2 \text{ and } h_2 \in N \\ &= f(a) && \text{since } a \text{ is generic over } Bh_1 \text{ and } h_1 \in N. \end{aligned}$$

Hence $h_1 + h_2 \in N$. Next, if $h \in N$, let a be generic over $B \cup \{-h\}$. Then a is generic over Bh and $a - h$ is generic over Bh as well. We have

$$\begin{aligned} f(a) &= f((a-h)+h) \\ &= f(a-h) && \text{since } a-h \text{ is generic over } Bh \text{ and } h \in N. \end{aligned}$$

So $-h \in N$.

To see that N is finite, let $m \in \mathbb{N}$ be the size of a largest fibre of f. Such an m exists by a standard compactness argument since every fibre is finite and we are in a saturated model. Consider any $h_0, \ldots, h_m \in N$, and let a be generic over $Bh_0, \ldots, h_m$. Then by definition of N, $f(a+h_1) = \ldots = f(a+h_m) = f(a)$, so $a+h_1, \ldots, a+h_m, a$ all lie in the same fibre of f. Since the largest fibre of f has size m, it follows that not all of the h_i's are distinct, and hence $|N| \leq m < \omega$.

Finally, we show that N is defined over B. Since N is finite it is definable, so we use an automorphism argument. Take any $\alpha \in \mathrm{Aut}_B(L)$, and any $h \in H$. Let a be generic over $B \cup \{\alpha(h)\}$. Then $b = \alpha^{-1}(a)$ is generic over Bh. We have

$$\begin{aligned} f(\alpha(h)+a) &= f(\alpha(h)+\alpha(b)) \\ &= f(\alpha(h+b)) \\ &= \alpha(f(h+b)) && \text{since } f \text{ is } B\text{-definable} \\ &= \alpha(f(b)) && \text{since } h \in N \\ &= f(\alpha(b)) \\ &= f(a) \end{aligned}$$

So $\alpha(h) \in N$. As $\alpha \in \mathrm{Aut}_B(L)$ was arbitrary, N is defined over B. $\dashv$

Since H is semiminimal it has finite U-rank, say $U(H) = r$. Let $h_0, \ldots, h_{2r}$ be independent generic elements of H over B. Define $\overline{f} : H \to (L^n)^{(2r+1)}$ by

$$\overline{f}(h) = (f(h+h_0), \ldots, f(h+h_{2r}))$$

For each $h \in H$ define $N_h := \left\{\overline{f}(h+d) : d \in N\right\} \subseteq (L^n)^{(2r+1)}$. Clearly each N_h is definable over $Bh_0, \ldots, h_{2r}$.

Claim 3.17.2. *For all $h, h' \in H$, if $\overline{f}(h) = \overline{f}(h')$ then $h - h' \in N$*

Proof of Claim 3.17.2 Suppose $\overline{f}(h) = \overline{f}(h')$. For any $b_0, b_1 \in H$, since $U(H) = r$, we have $U(b_0 b_1) \leq 2r$. Since we have $2r + 1$ independent generics h_i over B, it follows that at least one of them remains generic over Bb_0b_1. To see this, we first note that $U(h_i/B) = r$ for all i since each h_i is generic. Since the h_i's are independent over B, we have $U(h_0h_1 \dots h_i/B) = ir$ for each $0 \leq i \leq 2r$ by the Lascar equality. If none of the h_i's are generic over Bb_0b_1, then, by repeated application of the Lascar equality, we get $U(b_0b_1/Bt_0 \dots t_i) \leq 2r - i - 1$ for each $0 \leq i \leq 2r$. In particular, $U(b_0b_1/Bt_1 \dots t_{2r}) \leq 2r - 2r - 1 = -1$, a contradiction.

In particular, some h_i is generic over Bhh'. Then $h_i + h$ is generic over Bh', and hence over $B \cup \{h - h'\}$ since $h - h' \in \mathrm{dcl}(h, h')$, and so

$$\begin{aligned} f(h_i + h + h' - h) &= f(h_i + h') \\ &= f(h_i + h) && \text{since } \overline{f}(h) = \overline{f}(h') \end{aligned}$$

So $h - h' \in N$ by definition of N. ⊣

Claim 3.17.3. $N_h = N_{h'} \iff h - h' \in N$.

Proof of Claim 3.17.3 First,

$$\begin{aligned} N_h = N_{h'} &\Rightarrow \{\overline{f}(h + d) : d \in N\} = \{\overline{f}(h' + d) : d \in N\} \\ &\Rightarrow \text{There exists } d \in N \text{ such that } \overline{f}(h) = \overline{f}(h' + d) \\ &\Rightarrow \text{There exists } d \in N \text{ such that } h - h' - d \in N && \text{by Claim 3.17.2} \\ &\Rightarrow h - h' \in N && \text{since } d \in N \end{aligned}$$

Conversely, suppose $h - h' \in N$. Then for any $d \in N$, $\overline{f}(h + d) = \overline{f}(h' + h - h' + d)$, where $h - h' + d \in N$. Thus $\{\overline{f}(h + d) : d \in N\} \subseteq \{\overline{f}(h' + d) : d \in N\}$. Equality follows by symmetry. So $N_h = N_{h'}$. ⊣

Since both separably closed fields and differentially closed fields have elimination of imaginaries and the family $\{N_h : h \in H\}$ is a definable family of definable sets, there exists a definable function $g : H \to L^m$ for some m such that $g(h)$ is a code for N_h for every $h \in H$, and $g(h) = g(h')$ if and only if $N_h = N_{h'}$ (see Lemma B.1). We then have $g(h) = g(h')$ if and only if $h - h' \in N$ by Claim 3.17.3.

Thus $g(H) \subseteq L^m$ is in definable bijection, by the map $g(h) \mapsto h \bmod N$, with H/N, which is type-definable by elimination of imaginaries (see Lemma B.2). We use this bijection to put a group structure on $g(H)$. In general we do not have $g(H) \subseteq k^m$, but we can show that it is so generically:

Claim 3.17.4. *If $a \in H$ is generic over $Bh_0 \dots h_{2r}$ then $g(a) \in k^m$.*

Proof of Claim 3.17.4 Let $a \in H$ be generic over $Bh_0 \dots h_{2r}$. We recall that $g(a)$ is a code for $\{(f(a+d+h_0), \dots, f(a+d+h_{2r})) : d \in N\}$. Now consider any $d \in N$ and any h_i. Then since a is generic over Bh_i, $a+h_i$ is generic over B. Since N is finite and defined over B (Claim 3.17.1), $d \in \operatorname{acl}(B)$. So $a+h_i \mathop{\smile\hspace{-0.9em}\mid}_B d$, and hence $a+h_i+d$ is also generic in H over B. Thus $f(a+d+h_i) \in k^n$, and so $\left\{\overline{f}(a+d) : d \in N\right\} \subseteq (k^n)^{(2r+1)}$. By elimination of imaginaries for ACF, and the stable embeddedness of k, the code for this finite subset is itself in $\operatorname{dcl}(k)$. It follows that $g(a) \in k^m$. $\dashv$

Now replace B by $Bh_0 \dots h_{2r}$. We then have produced a B-definable surjective group homomorphism $g : H \to G'$ where $G' \subseteq L^m$ is a type-definable connected group, namely $G' = g(H)$. Since g is a homomorphism the fact that $g(h) = g(h') \iff h - h' \in N$ says that $\ker(g) = N$, which is finite by Claim 3.17.1. What is missing is that G' lies in L rather than in k, but we have also seen that the generics of G' are contained in k^m. We now show that we can recover all of G' from its generic type. Let q be the generic type of G' over B. Recall that, for any $x \in G'$, there exist $a, b \in q^L$ such that $x = a + b$. Let $\mathcal{R}$ be the definable equivalence relation on $q^L \times q^L$ given by

$$(x,y)\mathcal{R}(x',y') \iff x + y = x' + y'$$

And define

$$G := (q^L \times q^L)/\mathcal{R}$$

Then since $q^L \subseteq k^m$ and k is stably embedded, it follows from elimination of imaginaries for ACF_p that G is type-definable in $(k, +, -, \cdot, 0, 1)$ (see Lemma B.2). Since ACF_p is ω-stable G is in fact definable in $(k, +, -, \cdot, 0, 1)$.

Claim 3.17.5. *Define $g' : G' \to G$ by $g'(x) = (a, x-a) \bmod \mathcal{R}$, where a is generic in G' over Bx. Then g' is a B-definable bijective group homomorphism.*

Proof of Claim 3.17.5 Note that g' is well-defined, since $a + (x-a) = x = a' + (x - a')$, so $(a, x-a)\mathcal{R}(a', x-a')$.

Next we wish to see that g' is B-definable. It suffices to show that

$$\left\{(x, a, x-a) \in G' \times q^L \times q^L : x \mathop{\smile\hspace{-0.9em}\mid}_B a\right\}$$

is type-definable over B. Indeed, this implies that the graph of g is type-definable over B, and hence B-definable (see Lemma B.3). Since $-$ is definable it is enough to show that $\left\{(x,a) \in G' \times q^L : x \mathop{\smile\hspace{-0.9em}\mid}_B a\right\}$ is type-definable over B, which follows from definability of types (see Lemma B.4).

To see that g' is a group homomorphism, consider any $x, y \in G'$, and let a, b, c be generics in G such that $a \mathop{\smile\hskip-0.9em ^| \ } x$, $b \mathop{\smile\hskip-0.9em ^| \ } y$, $c \mathop{\smile\hskip-0.9em ^| \ } (x+y)$. Then $g'(x+y) = (c, x+y-c) \bmod \mathcal{R}$, and

$$\begin{aligned} g'(x) + g'(y) &= (a, x-a) + (b, y-b) \bmod \mathcal{R} \\ &= (a+b, x+y-a-b) \bmod \mathcal{R} \\ &= (c, x+y-c) \bmod \mathcal{R} && \text{by definition of } \mathcal{R} \end{aligned}$$

So g' is a group homomorphism. Finally, the map $(x, y) \bmod R \mapsto x + y$ is inverse to g, so g is bijective. $\dashv$ We thus have a definable group G in $(k, +, -, \cdot, 0, 1)$ and a definable group isomorphism $g' : G' \to G$.

Define $h : H \to G$ by $h = g' \circ g$. Since g' is a group isomorphism and g is a surjective homomorphism with finite kernel, h is also a surjective homomorphism with finite kernel, and hence is the required map. □

3.3.3 Semiminimal Subgroups of Semiabelian Varieties

Recall that if F is a field of characteristic $p > 0$ we have the *Frobenius morphism* given by $x \mapsto x^p$. If X is a variety over F then this lifts to a morphism $\mathrm{Fr} : X \to X$, also called the Frobenius morphism on X.

Definition 3.18. Let X and Y be varieties over a field of characteristic $p > 0$. We say that a map $g : X \to Y$ is *p-regular* if $g = \mathrm{Fr}^{-n} \circ f$ for some $n \in \mathbb{N}$ and some regular map f. In characteristic 0 a p-regular map is just a regular map.

If $f : G \to H$ is a definable homomorphism between algebraic groups G and H then f is p-regular (see [26, Lemma 4.9]). In particular, if f is regular (or p-regular) and bijective then f^{-1} is definable, and hence p-regular as well.

Proposition 3.19. *Let S be a semiabelian variety over L, and let $H \leq S(L)$ be a semiminimal connected subgroup such that $H \not\perp k$. Then there exists a semiabelian variety S_0 over k and a bijective regular homomorphism $g : \overline{H} \to S_0$ such that $g(H) = S_0(k)$.*

Proof. By Proposition 3.17 we get a definable surjective homomorphism with finite kernel $h : H \to G$ where G is definable in $(k, +, -, \cdot, 0, 1)$. The first step is to reverse the homomorphism h. Let $n = |\ker(h)|$.

Define $f : G \to H$ by $f(x) = ny$ where y is such that $h(y) = x$, which exists since h is surjective. Since h is definable so is f. We observe that

$$\begin{aligned} h(y') = h(y) &\Rightarrow y - y' \in \ker(h) \\ &\Rightarrow n(y - y') = 0 && \text{since } \ker(h) \text{ is a group of order } n \\ &\Rightarrow ny = ny' \end{aligned}$$

Hence f is well-defined. Clearly f is a group homomorphism.

Next we show that f is surjective. Since H is semiminimal it has finite U-rank. Since H is a subgroup of a semiabelian variety, and semiabelian varieties have finite m-torsion for all m (see Lemma B.6), H in particular has finite n-torsion. So the map $[n] : H \to H$ has finite kernel. In particular, $U(\ker[n]) = 0$, so from the Lascar equality for groups we have $U(H) = U(nH) + U(\ker[n]) = U(nH)$. Hence $[H : nH]$ is finite. Since H is connected $H = nH$. Now for any $y \in H$, let $y' \in H$ be such that $ny' = y$. Then $f(h(y')) = ny' = y$, so f is surjective.

Now let $G_1 = G/\ker(f)$. Then since G is definable in $(k, +, -, \cdot, 0, 1)$ so is G_1, by elimination of imaginaries. Let $f_1 : G_1 \to H$ be the bijective homomorphism induced by f. Then f_1 is definable in the language of DCF_0 or $\mathrm{SCF}_{p,\nu}$ since f is. By the Weil-van den Dries-Hrushovski Theorem (see [3]) we get that, up to definable isomorphism, $G_1 = T(k)$, where T is an algebraic group over k. Since $T(k)$ is definable in the field language and we only look at constant points we get that f_1 agrees on $T(k)$ with a function g definable in the field language.

We next want to extend f_1 to Zariski closures in L^{alg}.

$$\begin{tikzcd}
T(L^{\mathrm{alg}}) \arrow[r, dashed, "f_2"] & \overline{H}(L^{\mathrm{alg}}) \\
T(k) \arrow[u] \arrow[r, "f_1"] & H \arrow[u]
\end{tikzcd}$$

We first note that $\overline{T(k)} = T(L^{\mathrm{alg}})$. Next, since H is an algebraic subgroup of $S(L)$ its Zariski closure (in L^{alg}) is the set of L^{alg} points of an algebraic subgroup of S, which we will also call $\overline{H}$. To extend f_1 we in fact extend g. Since g is definable in the field language it is p-rational. Since $k \preceq L^{\mathrm{alg}}$ as structures in the field language we have that g is also defined on $T(L^{\mathrm{alg}})$. Let f_2 be the restriction of g to $T(L^{\mathrm{alg}})$, so f_2 is a field-language definable, hence p-rational, map extending f_1. Moreover, since the property of being a group homomorphism is expressible (in the field language) f_2 is a group homomorphism. Since f_2 is a Zariski-continuous extension of f_1 from $T(k)$ to its Zariski closure it follows from general topology that the image of f_2 is contained in $\overline{H}$, and since f_1 was surjective it follows also that f_2 is surjective. We may not have that f_2 is bijective, but $\ker(f_2)(k) = \ker(f_1)$, so since f_1 is bijective $\ker(f_2)(k) = 0$.

Claim 3.19.1. $\ker f_2$ *is defined over* k.

Proof of Claim 3.19.1 From the theory of algebraic groups we know that there exists an algebraic subgroup $R \leq T$ such that T/R is an abelian, and hence semiabelian, variety. In fact, it can be shown that there is a unique minimal R such that T/R is semiabelian, and that this R is defined over k (see [30, Theorem 16]). Now $T/\ker(f_2) \cong \overline{H}$, and $\overline{H}$ is a semiabelian variety, so $R \leq \ker(f_2)$. Since R is over k, $R(k)$ is Zariski-dense in R. But $R(k) \leq \ker(f_2)(k) = 0$, so $R = 0$. Thus T is already a semiabelian variety. Since T is defined over k and $\ker(f_2) \leq T$ we get by the rigidity of semiabelian varieties (see Proposition B.8) that $\ker(f_2)$ is defined over k. $\dashv$

Since $\ker(f_2)$ is defined over k, $\ker(f_2)(k)$ is Zariski-dense in $\ker(f_2)$. But $\ker(f_2)(k) = 0$, so $\ker(f_2) = 0$, and combining this with our earlier observations about f_2 gives that $f_2 : T \to \overline{H}$ is a definable bijective p-regular homomorphism. Since f_2 is bijective and definable in $(L^{\text{alg}}, +, -, \cdot, 0, 1)$ so is f_2^{-1}, and so $f_2^{-1} : \overline{H} \to T$ is a bijective p-regular homomorphism. Let $g : \overline{H} \to T^{(p^n)}$ be given by $g = \mathrm{Fr}^n \circ f_2^{-1}$, where n is the appropriate power of the Frobenius map Fr that makes g regular. If $p = 0$ then f_2^{-1} is already regular, so this step is unnecessary.

$$\begin{tikzcd}
\overline{H} \arrow[rr, "g"] \arrow[dr, "f_2^{-1}"'] & & T^{(p^n)} \\
& T \arrow[ur, "\mathrm{Fr}^n"'] &
\end{tikzcd}$$

Then g is a bijective regular homomorphism. Let $S_0 = T^{(p^n)}$. Since T is a semiabelian variety over k, so is S_0 (see Lemma B.10). All that remains to show is that $g(H) = S_0(k)$. We compute

$$\begin{aligned}
g(H) &= \mathrm{Fr}^n(f_2^{-1}(H)) \\
&= \mathrm{Fr}^n(T(k)) && \text{since } f_2 \text{ extends the bijection } f_1 : T(k) \to H \\
&= S_0(k) && \text{by definition of } S_0.
\end{aligned}$$

So S_0 and g are the required semiabelian variety and bijective rational homomorphism. □

The above Proposition gives rise to our first statement that resembles the Mordell-Lang statement.

Corollary 3.20. *Let S be a semiabelian variety over L, and let $H \leq S(L)$ be a connected semiminimal group such that $H \not\leq k$. Let $X \subseteq S$ be a subvariety over L. Then there exist $X_1, \ldots, X_l$ such that each X_i is a k-special subvariety of X, and*

$$X(L) \cap H = \bigcup_{i=1}^{l} X_i(L) \cap H$$

.

Proof. Let $X_1, \ldots, X_l$ be the irreducible components of $\overline{X(L) \cap H}$. Then

$$\begin{aligned} \bigcup_{i=1}^{l} X_i(L) \cap H &= \left(\bigcup_{i=1}^{l} X_i(L) \right) \cap H \\ &= \overline{X(L) \cap H} \cap H \\ &= X(L) \cap H \end{aligned}$$

So it remains to show that each X_i is k-special. Fix some $1 \leq i \leq l$.

By Proposition 3.19 there is a semiabelian variety S_0 over k and a bijective regular homomorphism $g : \overline{H} \to S_0$ such that $g(H) = S_0(k)$.

Let X_0 be the Zariski closure of $g(X_i(L) \cap H)$ in S_0. By automorphisms, since $g(X_i(L) \cap H) \subseteq S_0(k)$, X_0 is (model-theoretically) defined over k. Since $k = k^{\text{alg}}$ the model-theoretic and algebraic-geometric notions of being defined over k agree.

Note that by definition of X_i, $X_i(L) \cap H$ is Zariski-dense in X_i. By definition of X_0, we have $X_i(L) \cap H \subseteq g^{-1}(X_0)$, and so $\overline{X_i(L) \cap H} = X_i \subseteq g^{-1}(X_0)$. Conversely, g is a regular bijection, so its inverse is p-regular, and hence is also continuous for the Zariski topology. Hence we have

$$\begin{aligned} g^{-1}(X_0) &= g^{-1}(\overline{g(X_i(L) \cap H}) \\ &= \overline{g^{-1}(g(X_i(L) \cap H))} \\ &= \overline{X_i(L) \cap H} \\ &= X_i \end{aligned}$$

So $X_i = g^{-1}(X_0)$, and hence X_i is k-special by definition. □

Chapter 4

The Proof of Mordell-Lang

In this chapter we use the consequences of the Zilber Dichotomy discussed in the previous chapter to prove the Mordell-Lang statement. We begin with a proof in the case of semipluriminimal type-definable subgroups of semiabelian varieties, then generalize to finite U-rank subgroups of semiabelian varieties, and finally leave the definable category altogether and prove the positive characteristic Mordell-Lang statement (Theorem 2.13).

The conventions of Chapter 3 are still in force : We are working in a universal domain L for DCF_0 or $\mathrm{SCF}_{p,\nu}$, $k \subseteq L$ is the field of constants, and $\mathbb{F} \subseteq L$ is the prime field.

4.1 Semipluriminimal Subgroups of Semiabelian Varieties

Definition 4.1. Let Y be a type-definable set. If there exist a finite set F and minimal sets $X_1, \ldots, X_l$ such that $Y \subseteq \mathrm{acl}(F \cup X_1 \cup \ldots \cup X_l)$ then we say that Y is *semipluriminimal*.

The key fact about semipluriminimal groups that we will need is the following orthogonal decomposition theorem:

Proposition 4.2. *Let H be a connected semipluriminimal type-definable group. Then there exist $H_1, \ldots, H_l$, connected type-definable semiminimal subgroups of H, such that $H = H_1 + \ldots + H_l$, and for each $1 \leq i \leq l$, $H_i \perp \sum_{j \neq i} H_j$.*

The above proposition follows directly from the basic theory of the socle of a finite U-rank group, see Appendix A. We need the notion of the geometric stabilizer of a subvariety of a semiabelian variety:

Definition 4.3. Let S be a semiabelian variety, $X \subseteq S$ a subvariety. The *(geometric) stabilizer* of X is the subgroup $\operatorname{stab}_g(X) = \{a \in S : a + X = X\} \leq S$.

In general this stabilizer is distinct from the model-theoretic stabilizer introduced in Section 3.2. The two stabilizers are related, however, as the next lemma shows.

Lemma 4.4. *Let S be a semiabelian variety over L. Let $H \leq S(L)$ be a connected type-definable group of finite U-rank. Let $X \subseteq S$ be an irreducible subvariety over L such that $X(L) \cap H$ is Zariski-dense in X. Then there exists a complete type p in $X(L) \cap H$ over some parameter set $A = \operatorname{acl}(A)$ such that $p(x)^L$ is Zariski-dense in X. Moreover, if p in $X(L) \cap H$ is taken to be of minimal U-rank with p^L Zariski-dense in X, then $\operatorname{stab}(p) \subseteq \operatorname{stab}_g(X)$.*

Proof. We will need the following claim to build our complete type p:

Claim 4.4.1. *Let $D \subseteq S$ be any set, and let Z be Zariski-dense in X. Then either $Z \cap D$ or $Z \setminus D$ is Zariski-dense in X.*

Proof of Claim 4.4.1 We have

$$\begin{aligned} X &= \overline{Z} \\ &= \overline{(Z \cap D) \cup (Z \setminus D)} \\ &= \overline{Z \cap D} \cup \overline{Z \setminus D} \end{aligned}$$

Since X is irreducible, it follows that $X = \overline{Z \cap D}$ or $X = \overline{Z \setminus D}$, as claimed. ⊣

Now we find the desired complete type p as follows. First, we want $p^L \subseteq X(L) \cap H$, so let p_0 be the type defining $H \cap X(L)$. So $\overline{p_0^L} = \overline{X(L) \cap H} = X$, but p_0 may not be complete. Choose $A = \operatorname{acl}(A)$ a countable set of parameters over which H and X are defined. Enumerate the set of all L_A-formulae as $\theta_1, \theta_2, \ldots$. We now construct a sequence of types $p_0 \subseteq p_1 \subseteq \ldots$ inductively, starting from p_0 as already defined. At the ith step suppose that we have a type p_i in $X(L) \cap H$ whose realizations are Zariski-dense in X. Then either $(p_i \cup \{\theta_{i+1}\})^L$ or $(p_i \cup \{\neg\theta_{i+1}\})^L$ is Zariski-dense in X by Claim 4.4.1, so set $p_{i+1} = p_i \cup \{\theta_{i+1}\}$ or $p_{i+1} = p_i \cup \{\neg\theta_{i+1}\}$ accordingly. Then $p = \bigcup_{i=1}^{\infty} p_i$ is the required complete type in $X(L) \cap H$ whose solution set is Zariski-dense in X.

To finish the proof, suppose that p is as above, and has minimal U-rank amongst all such complete types. Let $h \in \operatorname{stab}(p)$, and let $\mathbf{p}$ be the unique non-forking extension of p to L. Then $\mathbf{p}(x)$ extends both $p(x)$ and $h + p(x)$ by definition of the stabilizer. Let $Y = p^L$. So $U(Y \cap (h + Y)) = U(Y) = U(\mathbf{p})$. Now $Y \cap (h + Y) \subseteq Y \cap (h + X(L) \cap H) \subseteq Y$, so

$$\begin{aligned} U(Y) &= U(Y \cap (h + Y)) \\ &\leq U(Y \cap (h + X(L) \cap H)) \\ &\leq U(Y) \end{aligned}$$

So $U(Y \cap (h + X(L) \cap H)) = U(Y)$, and the type defining $Y \cap (h + X(L) \cap H)$ is a partial non-forking extension of p to Ah. Since p is stationary it has a unique complete non-forking extension to Ah, which must be a completion of $Y \cap (h + X(L) \cap H)$. In particular, the type defining $Y \setminus (h + X(L) \cap H)$ must be a forking extension of p to Ah, so $U(Y \setminus (h + X(L) \cap H)) < U(Y)$. Then by the minimality of $U(Y)$, $Y \setminus (h + X(L) \cap H)$ is not Zariski-dense in X, and hence by Claim 4.4.1 $Y \cap (h + X(L) \cap H)$ is Zariski-dense in X. We thus have

$$\begin{aligned} X &= \overline{Y \cap (h + X(L) \cap H)} \\ &\subseteq \overline{Y} \cap \overline{h + X(L) \cap H} \\ &= X \cap (h + \overline{X(L) \cap H}) \\ &= X \cap (h + X) \end{aligned}$$

Hence $X \subseteq h + X$, from which it follows that $X = h + X$, so $h \in \operatorname{stab}_g(X)$. □

Lemma 4.5. *Let S be a semiabelian variety over L, $X \subseteq S$ a subvariety also over L. If the image of X under the quotient map $\pi : S \to S/\operatorname{stab}_g(X)$ is k-special then X is k-special.*

Proof. Let $S' = S/\operatorname{stab}_g(X)$. Then since quotients of semiabelian varieties naturally carry the structure of semiabelian varieties themselves (see Proposition B.9), S' is a semiabelian variety over L, and the canonical projection $\pi : S \to S'$ is a rational map. Let $X' = \pi(X)$. Since X' is k-special there exists $c' \in S'$, $S^* \leq S'$, S_0 a semiabelian variety over k, $X_0 \subseteq S_0$ a subvariety also over k, and a surjective rational homomorphism $h : S^* \to S_0$ such that $X' = c' + h^{-1}(X_0)$.

Let $\widehat{S} = \pi^{-1}(S^*)$. Let $g : \widehat{S} \to S_0$ be $g = h \circ \pi$. Since h and π are both rational and surjective, g is rational and surjective. Let $c \in S$ be such that $\pi(c) = c'$. We will show that $X = c + g^{-1}(X_0)$. Indeed, first observe that $c + g^{-1}(X_0) = c + \pi^{-1}(h^{-1}(X_0)) = \pi^{-1}(c' + h^{-1}(X_0))$. Next we see that $\pi^{-1}(\pi(X)) = X$. The direction $X \subseteq \pi^{-1}(\pi(X))$ is immediate. For the converse, suppose $x \in \pi^{-1}(\pi(X))$. Then $\pi(x) \in \pi(X)$, so there exists $y \in X$ such that $\pi(x) = \pi(y)$. Then $x - y \in \ker \pi = \operatorname{stab}_g(X)$, and so $x = y + (x - y) \in X + \operatorname{stab}_g(X) = X$. Finally, we have $\pi(X) = c' + h^{-1}(X_0)$, so taking preimages on both sides and using the two preceding observations gives $X = c + g^{-1}(X_0)$. □

In the proof of Theorem 4.6 below the role of the Zilber Dichotomy in Mordell-Lang begins to become clear, as we combine the results on one-basedness from Section 3.2 with the consequences of non-orthogonality developed in the previous chapter.

Theorem 4.6. *Let S be a semiabelian variety over L. Let $H \leq S(L)$ be a semipluriminimal connected type-definable group. Let $X \subseteq S$ be a subvariety over L. Then there exist*

$X_1, \ldots, X_l$, k-special subvarieties over X, such that

$$X(L) \cap H = \bigcup_{i=1}^{l} X_i(L) \cap H$$

Proof. Let $X_1, \ldots, X_l$ be the irreducible components of $\overline{X(L) \cap H}$. Then, as in the proof of Corollary 3.20, $X(L) \cap H = \bigcup_{i=1}^{l} X_i(L) \cap H$. So it remains to show that each X_i is k-special. Fix $1 \leq r \leq l$. Then X_r is, by definition, irreducible, and $\overline{X_r(L) \cap H} = X_r$. We saw in Lemma 4.5 that if the image of X_r in $S/\operatorname{stab}_g(X_r)$ is k-special then X_r is itself k-special, so by taking the quotient by $\operatorname{stab}_g(X_r)$ we may assume that $\operatorname{stab}_g(X_r) = \{0\}$. If $\dim(X_r) = 0$ then X_r is finite and hence k-special. So we may also assume that $\dim(X_r) > 0$.

Note that H has finite U-rank since it is semipluriminimal, so using Lemma 4.4 let $p(x)$ be a complete type in $X_r(L) \cap H$ over some parameters $A = \operatorname{acl}(A)$ such that $Y = p(x)^L \subseteq X_r(L) \cap H$ is Zariski-dense in X_r. Taking p to be of minimal U-rank, the Lemma also gives that $\operatorname{stab}(p) = \{0\}$.

Claim 4.6.1. *H is not one-based.*

Proof of Claim 4.6.1 Suppose to the contrary that H is one-based. Then by Theorem 3.14 p is the the generic type of a translate of $\operatorname{stab}(p) = \{0\}$. It follows that p^L is a singleton, contradicting that p^L is Zariski-dense in X_r. $\dashv$

Let $H_1, \ldots, H_m$ be an orthogonal decomposition of H, so each H_i is fully orthogonal to the sum of the previous ones, and each is semiminimal and connected. By Proposition 3.15 if each H_i is one-based then so is H. As H is not one-based, some H_i is not one-based. In fact, more is true:

Claim 4.6.2. *Exactly one H_i is not one-based.*

Proof of Claim 4.6.2 We have already seen that at least one H_i is not one-based. We now show that at most one H_i is not one-based. Suppose that H_i and H_j are both not one-based. By the Zilber Dichotomy both are not fully orthogonal to k. Hence by Proposition 3.19 each is definably isomorphic to some group definable in $(k, +, -, \cdot, 0, 1)$. But k is strongly minimal, so by Proposition 3.8 those two groups are not fully orthogonal. Hence $H_i \not\perp H_j$, and so we must have $H_i = H_j$. $\dashv$

Reindexing the H_i's if necessary, we may assume that H_m is not one-based. Let $B = H_1 + \ldots H_{m-1}$. Then we have that B is one-based (by Lemma 3.15). We noted earlier that $H_m \perp H_1 + \ldots + H_{m-1}$. That is, $H_m \perp B$.

Claim 4.6.3. *There exist complete types q_1 in B, q_2 in H_m, both over A, such that $Y = q_1^L + q_2^L$.*

Proof of Claim 4.6.3 Recall that Y is defined by the complete type p. Consider any $y \in Y$. Since p is complete, $p = \mathrm{tp}(y/A)$. Since $Y \subseteq B + H_m$ we can write $y = u + v$ for some $u \in B$ and some $v \in H_m$. Let $q_1 = \mathrm{tp}(u/A)$ and $q_2 = \mathrm{tp}(v/A)$. Since $u \in B$ and B is type-definable, $q_1^L \subseteq B$. Similarly, $q_2^L \subseteq H_m$. We show that $Y = q_1^L + q_2^L$.

First, suppose that $z \in q_1^L + q_2^L$. So $z = a + b$ where $a \models q_1$ and $b \models q_2$. Then since $B \perp H_m$ we have (by the second equivalent version of full orthogonality in Lemma 3.6) $\mathrm{tp}(u/A) \cup \mathrm{tp}(v/A) \vdash \mathrm{tp}(uv/A)$. So by our choice of q_1 and q_2, $ab \models \mathrm{tp}(uv/A)$. In particular $a + b \models \mathrm{tp}(u + v/A) = p$. That is, $z = a + b \in Y$. So we have $q_1^L + q_2^L \subseteq Y$.

Conversely, suppose that $a \in Y$. Then $\mathrm{tp}(a/A) = p = \mathrm{tp}(u + v/A)$, so there exists $\sigma \in \mathrm{Aut}_A(L)$ such that $\sigma(u+v) = a$. So $\sigma(u)+\sigma(v) = a$. Since σ fixes A pointwise we have $q_1 = \mathrm{tp}(u/A) = \mathrm{tp}(\sigma(u)/A)$ and $q_2 = \mathrm{tp}(v/A) = \mathrm{tp}(\sigma(v)/A)$. So $a = \sigma(u) + \sigma(v) \in q_1^L + q_2^L$ as required. $\dashv$

Now since B is one-based q_1 is the generic type of a translate of $\mathrm{stab}(q_1)$ (Theorem 3.14). We then have, since $q_1^L \subseteq Y$, $\mathrm{stab}(q_1) \subseteq \mathrm{stab}(p) = \{0\}$, so $\mathrm{stab}(q_1) = \{0\}$. Thus $q_1^L = \{u\}$ for some element $u \in H$. Then $Y = u + q_2^L$ is a translate of q_2^L. Since k-specialness is preserved by translation, translating everything by $-u$ we may assume that $Y = q_2^L \subseteq H_m$. As Y is Zariski-dense in X_r and $Y \subseteq X_r(L) \cap H_m$, we have that $X_r(L) \cap H_m$ is also Zariski-dense in X_r. By the Zilber Dichotomy as H_m is not one-based we must have $H_m \not\perp k$. By Corollary 3.20 there exist $Z_1, \dots, Z_d$, k-special subvarieties of X_r, such that $X_r(L) \cap H_m = \bigcup_{i=1}^d Z_i(L) \cap H_m$. Then

$$\begin{aligned} X_r &= \overline{X_r(L) \cap H_m} \\ &= \overline{\bigcup_{i=1}^{d} Z_i(L) \cap H_m} \\ &= \bigcup_{i=1}^{d} \overline{Z_i(L) \cap H_m} \\ &= \bigcup_{i=1}^{d} Z_i \end{aligned}$$

Then since X_r is irreducible, $X_r = Z_j$ for some $1 \leq j \leq d$, and so X_r is k-special, as required. □

4.2 Finite U-Rank Subgroups of Semiabelian Varieties

In this section we extend Theorem 4.6 to all type-definable subgroups of finite U-rank. The following theorem is the key model-theoretic fact we need to reduce to the semipluriminimal case of the previous section. It says that a type-definable group H has a semipluriminimal subgroup $\mathrm{socle}(H)$ which, under the assumption of rigidity, controls all of H.

Definition 4.7. Let H be a type-definable group of finite U-rank. The *socle* of H, denoted $\mathrm{socle}(H)$, is the subgroup of H generated by all of the connected type-definable semiminimal subgroups of H.

If H is a type-definable group over parameters A we say that H is *rigid* if every relatively definable subgroup of H is defined over $\mathrm{acl}(A)$.

Theorem 4.8. *Let H be a type-definable group of finite U-rank. Then*

1. $\mathrm{socle}(H)$ *is the unique maximal connected type-definable semipluriminimal subgroup of H.*

2. *Suppose that* $\mathrm{socle}(H)$ *is rigid. Then every complete stationary type in H with finite stabilizer is contained in a single coset of* $\mathrm{socle}(H)$.

Proof. For the case when H is of finite Morley rank this is an immediate consequence of Hrushovski's "socle theorem" which appears in [10, Proposition 4.3]. In the positive characteristic $\mathrm{SCF}_{p,\nu}$ setting we will be applying this theorem to a group of finite U-rank which does not have finite Morley rank, so we give a proof of the socle theorem in the more general finite U-rank case in Appendix A. Theorem 4.8 then follows from Proposition A.6 and Theorem A.8, since every type-definable subgroup of H can be written as an intersection of definable subgroups of H. □

Lemma 4.9. *Let H be a type-definable group over B, and let X, Y be fully orthogonal type-definable subgroups of H over $A \supseteq B$. Then $X + Y$ is rigid.*

Proof. We first note that any definable subgroup of $X + Y$ is the image of a definable subgroup of $X \times Y$ under the A-definable map $+ : X \times Y \to X + Y$, so it suffices to show that $X \times Y$ is rigid. By Proposition 3.7, since $X \perp Y$, every definable subset of $X \times Y$ is a finite union of sets $X' \times Y'$ with X' definable in X and Y' definable in Y. So a connected definable subgroup of $X \times Y$ is a product of a definable subgroup of X and a definable subgroup of Y. It follows that if Z is a definable subgroup of $X \times Y$ then there exist X_1, X_2 definable subgroups of X and Y_1, Y_2 definable subgroups of Y such that

$X_1 \times Y_1 \leq Z \leq X_2 \times Y_2$ and $[X_2 \times Y_2 : X_1 \times Y_1] < \omega$. Since X and Y are rigid, $X_1 \times Y_1$ and $X_2 \times Y_2$ are over $\operatorname{acl}(A)$, so Z is a finite union of $\operatorname{acl}(A)$-definable translates of $X_1 \times Y_1$, and hence is itself over $\operatorname{acl}(A)$. So $X \times Y$ is rigid, and hence so is $X + Y$. □

Theorem 4.10. *Let S be a semiabelian variety over L. Let $H \leq S(L)$ be a type-definable subgroup of finite U-rank. Let $X \subseteq S$ be a subvariety over L. Then there exist $X_1, \ldots, X_l$, k-special subvarieties of X, such that*

$$X(L) \cap H = \bigcup_{i=1}^{l} X_i(L) \cap H$$

Proof. As in previous proofs, we let $X_1, \ldots, X_l$ be the irreducible components of $\overline{X(L) \cap H}$. We fix some $1 \leq r \leq l$ and show that X_r is k-special. As in the previous section we may assume that $\operatorname{stab}_g(X_r) = \{0\}$ by working in $S/\operatorname{stab}_g(X_r)$. Let p be a complete type in $X_r(L) \cap H$ over $A = \operatorname{acl}(A)$ such that everything is over A, $\operatorname{stab}(p) = \{0\}$ and $Y = p(x)^L$ is Zariski-dense in X_r. Such a p exists by Lemma 4.4.

Claim 4.10.1. $\operatorname{socle}(H)$ *is rigid.*

Proof of Claim 4.10.1 As $\operatorname{socle}(H)$ is semipluriminimal it has an orthogonal decomposition $\operatorname{socle}(H) = H_1 + \ldots + H_m$. The sum of two fully orthogonal rigid subgroups is rigid by the previous lemma, so by induction it suffices to show that each H_i is rigid. So fix some $1 \leq i \leq m$. If H_i is not one-based, then by Proposition 3.19 H_i is definably isomorphic to the k-points of a semiabelian variety over k. Semiabelian varieties are rigid over k (see Lemma B.7), and so H_i is rigid. Otherwise H_i is one-based, but one-based groups are always rigid. ⊣

By the socle theorem $p(x)^L$ is contained in a coset of $\operatorname{socle}(H)$. By translating everything, we may assume that $Y = p(x)^L \subseteq \operatorname{socle}(H)$. Then

$$\begin{aligned} X_r &= \overline{Y} \\ &\subseteq \overline{X_r(L) \cap \operatorname{socle}(H)} \\ &\subseteq \overline{X_r(L)} \\ &= X_r \end{aligned}$$

So $X_r(L) \cap \operatorname{socle}(H)$ is Zariski-dense in X_r. Since $\operatorname{socle}(H)$ is semipluriminimal we can apply Theorem 4.6 to get that there exist $Z_1, \ldots, Z_t$, k-special subvarieties of X_r, such that

$$X_r(L) \cap \operatorname{socle}(H) = \bigcup_{i=1}^{t} Z_i(L) \cap \operatorname{socle}(H)$$

Then

$$\begin{aligned}
X_r &= \overline{X_r(L) \cap \operatorname{socle}(H)} \\
&= \overline{\bigcup_{i=1}^{t} Z_i(L) \cap \operatorname{socle}(H)} \\
&= \bigcup_{i=1}^{t} \overline{Z_i(L) \cap \operatorname{socle}(H)} \\
&= \bigcup_{i=1}^{t} Z_i
\end{aligned}$$

So as X_r is irreducible, $X_r = Z_j$ for some j, and hence X_r is k-special as required. □

4.3 Mordell-Lang in Positive Characteristic

Theorem 4.10 was the Mordell-Lang statement for finite U-rank type-definable subgroups of semiabelian varieties. Now we leave the definable category to prove the arithmetic Mordell-Lang statement in characteristic $p > 0$. The characteristic 0 case follows a similar strategy, but we focus on the positive characteristic case for two reasons: First, the positive characteristic case is the part of Hrushovski's proof that was not already known, and second, the characteristic 0 case is explained in detail in [4]. We drop our standing assumptions about the fields L and k.

Theorem 4.11. *Let F be an algebraically closed field of characteristic $p > 0$. Let S be a semiabelian variety over F, and let $X \subseteq S$ be a subvariety also over F. Let $\Gamma' \leq S$ be a finitely generated group. Let $\operatorname{div}_p(\Gamma') = \{y \in S(F) : ny \in \Gamma'$ for some n such that $p \nmid n\}$. Then for any $\Gamma \leq \operatorname{div}_p(\Gamma')$ there exist $X_1, \ldots, X_l$, special subvarieties of S, such that $X_i \subseteq X$ for all i, and*

$$X(F) \cap \Gamma = \bigcup_{i=1}^{l} X_i(F) \cap \Gamma$$

Proof. As usual, we let $X_1, \ldots, X_l$ be the irreducible components of $\overline{X(F) \cap \Gamma}$, fix $1 \leq r \leq l$ and show that X_r is special.

Let $k = \mathbb{F}_p^{\text{alg}}$, and let K be a finitely generated extension of k over which S and X are defined, and such that the coordinates of the generators for Γ' are all from K. Such an extension exists because Γ' is finitely generated, and each of S, X is defined by finitely many polynomial equations, each having only finitely many coefficients. Moreover, since the group operation on S is defined by polynomial equations (with, again, finitely many coefficients), we may also assume that $+$ and $-$ are defined over K. Let $L = K^{\text{sep}}$.

Claim 4.11.1. *$k = L^{p^\infty}$.*

Proof of Claim 4.11.1 Since k is algebraically closed and K/k is finitely generated there exists a transcendence basis $\alpha_1, \ldots, \alpha_m$ for K over k such that $K/k(\alpha_1, \ldots, \alpha_m)$ is a separably algebraic extension (that such a transcendence basis exists is a standard fact from algebra - see [13, Corollary 2.12]).

Let $\nu(E)$ denote the degree of imperfection of a field E. We recall that, for a purely transcendental extension E'/E we have $\nu(E') = \nu(E) + \text{trdeg}(E'/E)$, while for a separably algebraic extension E'/E we have $\nu(E') = \nu(E)$. We thus have

$$\begin{aligned}\nu(L) &= \nu(K)\\ &= \nu(k(\alpha_1, \ldots, \alpha_m))\\ &= \nu(k) + \text{trdeg}(k(\alpha_1, \ldots, \alpha_m))\\ &= 0 + m && \text{since } k \text{ is perfect and } k(\alpha_1, \ldots, \alpha_m)/k \text{ is purely transcendental}\\ &= m\end{aligned}$$

Hence $L \models \text{SCF}_{p,m}$, and so its constant field L^{p^∞} is algebraically closed (see Section 2.2.2). We then compute

$$\begin{aligned}\nu(L) &= \nu(L^{p^\infty}) + \text{trdeg}(L/L^{p^\infty})\\ &= 0 + \text{trdeg}(L/L^{p^\infty}) && \text{since } L^{p^\infty} \text{ is perfect.}\end{aligned}$$

Hence $m = \text{trdeg}(L/L^{p^\infty})$. As $k = k^{p^\infty} \subseteq L^{p^\infty}$ are algebraically closed, and $\text{trdeg}(L/k) = m$, we have $k = L^{p^\infty}$. $\dashv$

We will show that X_r is k-special.

Claim 4.11.2. *$\Gamma \leq S(L)$.*

Proof of Claim 4.11.2 By our choice of L, S and its group operations are defined over L and the generators of Γ' are elements of $S(L)$, so $\Gamma' \leq S(L)$. Note that for $n \in \mathbb{N}$ such that $\gcd(n, p) = 1$ the multiplication map $[n] : S \to S$ is a separable morphism. Suppose $g \in \Gamma$. As $\Gamma \leq \text{div}_p(\Gamma')$ there is some $n \in \mathbb{N}$ such that $\gcd(n, p) = 1$ and $ng \in \Gamma'$. Since $[n]$ is separable and $ng \in S(L)$ we get that $g \in S(L^{\text{sep}}) = S(L)$. Hence $\Gamma \leq S(L)$. $\dashv$

So $X_r(F) \cap \Gamma = X_r(L) \cap \Gamma$, and so $X_r(L) \cap \Gamma$ is Zariski-dense in X_r. Note that we are approaching the setting of the previous chapters. We have $L \models \text{SCF}_{p,m}$, k is the field of constants, and $X_r(L) \cap \Gamma$ is Zariski-dense in X_r. However, Γ is not type-definable, and L may not be saturated. We first deal with replacing L by a saturated model.

Let L^* be a saturated elementary extension of L (in the sense of $\mathrm{SCF}_{p,m}$). Let $k^* = (L^*)^{p^\infty}$. Let $\mathcal{L}' = \{+, -, \cdot, 0, 1, k_0\}$ where k_0 is a new unary relation symbol. We will be using k_0 to represent the constant field in a larger field, so we abbreviate $(L^{\mathrm{alg}}, +, -, \cdot, 0, 1, k)$ by (L^{alg}, k).

Claim 4.11.3. *As structures in the language $\mathcal{L}'$, $(L^{\mathrm{alg}}, k) \preceq ((L^*)^{\mathrm{alg}}, k^*)$.*

Proof of Claim 4.11.3 Since $L^* \models \mathrm{SCF}_{p,m}$ $k^* = (L^*)^{p^\infty} \models \mathrm{ACF}_p$. We are therefore considering pairs of models of ACF_p. It is clear from the fact that $L \subseteq L^*$ that $(L^{\mathrm{alg}}, k) \subseteq ((L^*)^{\mathrm{alg}}, k^*)$. We wish to use the fact that $(L^{\mathrm{alg}}, k) \preceq ((L^*)^{\mathrm{alg}}, k^*)$ if and only if $(L^{\mathrm{alg}}, k) \subseteq ((L^*)^{\mathrm{alg}}, k^*)$ and k^* is algebraically disjoint from L^{alg} over k (see [28, Theorem 8]). Since we are working with algebraically closed fields algebraic disjointedness is equivalent to linear disjointedness. So it remains to show that k^* is linearly disjoint from L^{alg} over k.

We first show that $(L^*)^{p^n}$ is linearly disjoint from L^{alg} over L^{p^n} for each n. Let $v_1, \ldots, v_m \in F$ be linearly independent over L^{p^n}. So for all $\alpha_1, \ldots, \alpha_m \in L^{p^n}$ we have $\sum_{i=1}^m \alpha_i v_i = 0 \Rightarrow \alpha_1 = \ldots = \alpha_m = 0$. Since L^{p^n} is definable in $(L, 0, 1, +, -, \cdot)$ the previous sentence is expressible as a first-order sentence, and hence remains true in the elementary extension L^*. That is, $v_1, \ldots, v_m$ remain linearly independent over $(L^*)^{p^n}$.

Now we show that k^* is linearly disjoint from L^{alg} over k. Let $v_1, \ldots, v_n \in k^*$ be linearly dependent over L^{alg}. We must show that they are also linearly dependent over k. We may assume that every subset $v_{i_1}, \ldots, v_{i_m}$ is linearly independent over L^{alg}, since if not we can replace the v_i's with a subset having this property. Let $\alpha_1, \ldots, \alpha_n \in L^{\mathrm{alg}}$ be not all 0 and such that $\sum_{i=1}^n \alpha_i v_i = 0$. Then since every proper subset of $\{v_1, \ldots, v_n\}$ is linearly independent we get that for all i $\alpha_i \neq 0$. Further, since $\{v_2, \ldots, v_n\}$ are linearly independent and $v_1 \in \operatorname{span}\{v_2, \ldots, v_n\}$ there is a unique way of expressing v_1 as a linear combination of $v_2, \ldots, v_n$. Hence, by dividing by α_1, we may assume that $\alpha_1 = 1$, and this uniquely determines the remaining α_i's. Consider any $r \geq 1$. We have that $v_1, \ldots, v_n \in (L^*)^{p^r}$ are linearly dependent over L^{alg}, so since we have already shown that $(L^*)^{p^r}$ is linearly disjoint from L^{alg} over L^{p^r} it follows that $v_1, \ldots$ v_n are linearly dependent over $L^{p^r} \subseteq L^{\mathrm{alg}}$. Let $\beta_1, \ldots, \beta_n \in L^{p^r}$ be not all 0 such that $\sum_{i=1}^n \beta_i v_i = 0$. Then we again have that $\beta_i \neq 0$ for all i, and we can assume that $\beta_1 = 1$. Then $\sum_{i=2}^n \alpha_i v_i = \sum_{i=2} \beta_i v_i$, so since $\{v_2, \ldots, v_n\}$ is linearly independent we have $\alpha_i = \beta_i \in L^{p^r}$ for all i. Hence for all i we have $\alpha_i \in L^{p^\infty} = k$, so $v_1, \ldots, v_n$ are linearly dependent over k. ⊣

Claim 4.11.4. *If X_r is k^*-special then X_r is k-special.*

Proof of Claim 4.11.4 Suppose that X_r is k^*-special. This means that there is an algebraic subgroup $S' \leq S$, a semiabelian variety S_0 over k^*, a surjective rational homomorphism $h : S' \to S_0$, a subvariety $X_0 \subseteq S_0$ also over k^*, and a c such that $X_r = c + h^{-1}(X_0)$. We first note that, being a variety over k^*, S_0 is definable in $((L^*)^{\mathrm{alg}}, k^*)$, say by a formula ϕ_{S_0}.

As a semiabelian variety there is a short exact sequence $0 \to T_0 \to^{f_1} S_0 \to^{f_2} A_0 \to 0$ for some algebraic torus T_0 and some abelian variety A_0. Again $\mathbb{G}_m^t$ and A_0 are definable say by formulae ϕ_{T_0} and ϕ_{A_0}. Since the maps f_1 and f_2 are morphisms they are definable by formulae ϕ_{f_1} and ϕ_{f_2}. Then the statement "ϕ_{f_1} defines an injective group homomorphism, ϕ_{f_2} defines a surjective group homomorphism, and $\ker f_2 = \operatorname{im} f_1$" can be expressed as a first-order sentence, as can the statements "ϕ_{T_0} defines a torus" and "ϕ_{A_0} is a projective algebraic group". Combining these sentences yields a sentence which says "ϕ_{S_0} defines a semiabelian variety". So ϕ_{S_0} interpreted in (L, k) defines a semiabelian variety S_1 over k. Similarly there is a formula $\phi_{S'}$ defining S', and "$\phi_{S'}$ defines an algebraic subgroup of S" is expressible as a first-order sentence, so $\phi_{S'}$ interpreted in (L, k) gives an algebraic subgroup $S_1' \leq S$. The function h, being rational, is definable by a formula ϕ_h. Then "ϕ_h defines a surjective homomorphism" is a first-order sentence, so ϕ_h defines a surjective rational homomorphism $h_1 : S_1' \to S_1$ when interpreted in (L, k). Finally, letting ϕ_{X_0} be the formula defining X_0, ϕ_{X_0} defines a subvariety X_1 of S_1, and there is a first-order sentence expressing "There exists c such that $X_r = c + h^{-1}(X_0)$. When interpreted in (L, k) this means that there is some c_1 such that $X_r = c_1 + h_1^{-1}(X_1)$, so X_r is k-special. ⊣

In light of Claim 4.11.4 we relabel to assume that L is itself a saturated model of SCF_p. Of course, now $k \neq \mathbb{F}_p^{\mathrm{alg}}$.

Finally, and this is the main step, we need to relate Γ to a type-definable group of finite U-rank. The goal will be to replace Γ by the type-definable group $\bigcap_n p^n S(L)$. The first thing to do is to show that $X_r \cap (\bigcap_n p^n S(L))$ is Zariski-dense in X_r.

Claim 4.11.5. *For all $n \geq 0$ the quotient $\operatorname{div}_p(\Gamma')/p^n \operatorname{div}_p(\Gamma')$ is finite.*

Proof of Claim 4.11.5 Fix $n \geq 0$. As Γ' is finitely generated the quotient $\Gamma'/p^n\Gamma'$ is a finitely generated $\mathbb{Z}/p^n\mathbb{Z}$-module, and hence is in fact finite. The inclusion $\Gamma' \leq \operatorname{div}_p(\Gamma')$ induces a map $\theta : \Gamma'/p^n\Gamma' \to \operatorname{div}_p(\Gamma')/p^n \operatorname{div}_p(\Gamma')$ given by $\theta(a + p^n\Gamma') = a + p^n \operatorname{div}_p(\Gamma')$. We show that θ is surjective.

Consider any coset $g + p^n \operatorname{div}_p(\Gamma') \in \operatorname{div}_p(\Gamma')/p^n \operatorname{div}_p(\Gamma')$. Since $g \in \operatorname{div}_p(\Gamma')$ there exists $l \in \mathbb{Z}$ such that $\gcd(l, p) = 1$ and $lg \in \Gamma'$. Then also $\gcd(l, p^n) = 1$, so there exist $x, y \in \mathbb{Z}$ such that $lx + p^n y = 1$. So

$$\begin{aligned} g - p^n yg &= g - (1 - lx)g \\ &= g - g + lxg \\ &= lxg \end{aligned}$$

Hence $lxg \equiv g (\operatorname{mod} p^n \operatorname{div}_p(\Gamma'))$. Note that $lxg = x(lg) \in \Gamma'$. So $\theta(lxg + p^n\Gamma') = g + p^n \operatorname{div}_p(\Gamma')$. Hence θ is surjective. ⊣

Claim 4.11.6. *For all $n < \omega$ there exists $a_n \in \operatorname{div}_p(\Gamma')$ such that $(a_n + p^n \operatorname{div}_p(\Gamma')) \cap X_r(L)$ is Zariski-dense in X_r.*

Proof of Claim 4.11.6 Fix $n < \omega$. By Claim 4.11.5 $\operatorname{div}_p(\Gamma')/p^n \operatorname{div}_p(\Gamma')$ is finite, say having size $t < \omega$. There thus exist $b_1, \ldots, b_t \in \operatorname{div}_p(\Gamma')$ such that $\operatorname{div}_p(\Gamma') = (b_1 + p^n \operatorname{div}_p(\Gamma')) \cup \ldots \cup (b_t + p^n \operatorname{div}_p(\Gamma'))$. Recalling that $X_r(L) \cap \Gamma$ is Zariski-dense in X_r, and that $\Gamma \leq \operatorname{div}_p(\Gamma')$, we thus have

$$\begin{aligned} X_r &= \overline{\operatorname{div}_p(\Gamma') \cap X_r(L)} \\ &= \overline{\left(\bigcup_{i=1}^{t} b_i + p^n \operatorname{div}_p(\Gamma') \right) \cap X_r(L)} \\ &= \bigcup_{i=1}^{t} \overline{(b_i + p^n \operatorname{div}_p(\Gamma')) \cap X_r(L)} \end{aligned}$$

Since X_r is irreducible, there is some $1 \leq j \leq t$ such that $X_r = \overline{b_j + p^n \operatorname{div}_p(\Gamma')}$. So $a_n = b_j$ is the required element of $\operatorname{div}_p(\Gamma')$. ⊣

Let $p^\infty \operatorname{div}_p(\Gamma') = \bigcap_{n=0}^{\infty} p^n \operatorname{div}_p(\Gamma')$.

Claim 4.11.7. *There exists $g \in S(L)$ such that $X_r(L) \cap (g + p^\infty \operatorname{div}_p(\Gamma'))$ is Zariski-dense in X_r.*

Proof of Claim 4.11.7 For each $n < \omega$ let $\theta_n(x)$ be the type expressing "$(x + p^n \operatorname{div}_p(\Gamma')) \cap X_r(L)$ is Zariski-dense in X_r". Let $\Theta(x) = \bigcup_{n=0}^{\infty} \theta_n(x)$, and let $\Theta'(x)$ be any finite subset of $\Theta(x)$. Let $N < \omega$ be the largest value of n such that $\theta_n(x) \subseteq \Theta'(x)$. By Claim 4.11.6 there exists g_N such that $\models \theta_N(g_N)$. Then for any $m < N$, we have $p^m \operatorname{div}_p(\Gamma') \supseteq p^N \operatorname{div}_p(\Gamma')$, and so

$$\begin{aligned} X_r &\supseteq \overline{(g_N + p^m \operatorname{div}_p(\Gamma')) \cap X_r} \\ &\supseteq \overline{(g_N + p^N \operatorname{div}_p(\Gamma')) \cap X_r} \\ &= X_r \end{aligned}$$

So we have $\models \theta_m(g_N)$. Thus $\models \Theta'(g_N)$. By compactness $\Theta(x)$ is satisfiable, and by saturation we thus have $g \in S(L)$ such that $\models \Theta(g)$, so $g + p^n \operatorname{div}_p(\Gamma') \cap X_r(L)$ is Zariski-dense in X_r for all $n < \omega$. Then, as it follows from saturation that a definable intersection of nested Zariski-dense subsets of X_r is again Zariski-dense in X_r, $\overline{(g + p^\infty \operatorname{div}_p(\Gamma')) \cap X_r(L)} = X_r$ as required. ⊣

After applying Claim 4.11.7 we can translate to assume that $g = 0$; that is, we may assume that $X_r(L) \cap p^\infty \operatorname{div}_p(\Gamma')$ is Zariski-dense in X_r. Letting $p^\infty S(L) = \bigcap_n p^n S(L)$ we then have that $X_r(L) \cap p^\infty S(L)$ is Zariski-dense in X_r.

Fact 4.12. *$p^\infty S(L)$ is a type-definable subgroup of $S(L)$ of finite U-rank.*

Proof. $p^\infty S(L)$ is visibly a type-definable subgroup of $S(L)$. For the fact that it is of finite U-rank, see [8, Proposition 5.8] □

Now by Theorem 4.10 applied to $p^\infty S(L)$, we have that there exist $Z_1, \dots, Z_m$, k-special subvarieties of X_r, such that

$$X_r(L) \cap p^\infty S(L) = \bigcup_{i=1}^{m} Z_i(L) \cap p^\infty S(L)$$

Then

$$\begin{aligned} X_r &= \overline{X_r(L) \cap p^\infty S(L)} \\ &= \overline{\bigcup_{i=1}^{m} Z_i(L) \cap p^\infty S(L)} \\ &= \bigcup_{i=1}^{m} \overline{Z_i(L) \cap p^\infty S(L)} \\ &= \bigcup_{i=1}^{m} Z_i \end{aligned}$$

So since X_r is irreducible there is some i such that $X_r = Z_i$, and hence X_r is k-special as required. □

APPENDICES

Appendix A

Groups of Finite U-Rank : The Socle Theorem

In this appendix we present a proof of the finite U-rank version of Hrushovski's socle theorem. There are several well-known expositions of the socle theorem in the case of finite Morley rank, including [4, Proposition 2.10], [27, Proposition 5.29], and Hrushovski's original presentation in [10, Proposition 4.3]. However, in positive characteristic, Hrushovski's proof of Mordell-Lang requires a version of this theorem for groups of finite U-rank. It appears to be generally known amongst experts that Hrushovski's argument can be modified to work in this setting, though to our knowledge these details have not been written down. We therefore take this opportunity to do so.

Throughout $\mathcal{U}$ is a universal domain for a complete stable theory T. G is a commutative type-definable group in $\mathcal{U}$ with $U(G) < \omega$. We write G additively. As usual we work throughout in $\mathcal{U}^{\mathrm{eq}}$ without mention.

In what follows we will describe the socle of G, which will turn out to be a type-definable subgroup of G. Some of the difficulty in extending the Socle Theorem to the finite U-rank setting will be that the socle need not be (relatively) definable in G. For example, we would like to work in $G/\operatorname{socle}(G)$, but in general quotients of type-definable sets by type-definable equivalence relations are not even type-definable. Nevertheless, the next result shows that we can still view cosets of connected subgroups as elements, though we cannot necessarily collect them into a type-definable set.

Proposition A.1. *Let $H \leq G$ be a connected type-definable subgroup, and let A be parameters over which G and H are defined. Then for any $a \in G$ the coset $a + H$ has a "code over A". That is, there is a finite tuple b such that $a + H$ is type-definable over $A \cup \{b\}$ and for any $\sigma \in \operatorname{Aut}_A(\mathcal{U})$, $\sigma(a + H) = a + H$ if and only if $\sigma(b) = b$.*

Proof. We saw in Proposition 1.9 that it follows just from stability that there exists $B \supseteq A$ such that $a + H$ is type-definable over B and for all $\sigma \in \mathrm{Aut}_A(\mathcal{U})$, $\sigma(a + H) = a + H$ if and only if $\sigma|_B = \mathrm{id}$. We now use superstability to find a finite tuple b from B such that $B \subseteq \mathrm{dcl}(A \cup \{b\})$. This b then satisfies the conclusions of the proposition to be proved.

Let $a' \in a + H$ be generic over B, and let $p = \mathrm{tp}(a'/B)$.

Claim A.1.1. *p is a stationary type.*

Proof of Claim A.1.1 Take any $C \supseteq B$, and suppose that q_1, q_2 are two non-forking extensions of p to $C \cup \{a\}$. Then both $-a + q_1$ and $-a + q_2$ extend H, and moreover $U(-a + q_1) = U(-a + q_2) = U(p)$, so both are generic in H. Since H is connected it has a unique generic type over $C \cup \{a\}$, so $-a + q_1 = -a + q_2$, and hence $q_1 = q_2$ as required. $\dashv$

By superstability (in this case, finite U-rank) there exists a finite tuple b from B such that $\mathrm{cb}(p) = \mathrm{dcl}(b)$. Take any $\sigma \in \mathrm{Aut}_A(\mathcal{U})$ and suppose that $\sigma(b) = b$. Since σ fixes A pointwise and H was defined over A, σ must permute the cosets of H in G. Since p is stationary and $\mathrm{cb}(p) = \mathrm{dcl}(b)$, p and $\sigma(p)$ have a common non-forking extension. It follows that $a + H$ and $\sigma(a + H) = \sigma(a) + H$ are not disjoint, and hence $\sigma(a + H) = a + H$. By the choice of B this implies $\sigma|_B = \mathrm{id}$. That is, $B \subseteq \mathrm{dcl}(A \cup \{b\})$, as desired. □

A.1 Zilber Indecomposibility

In this section we present some results about groups of finite U-rank. In particular, we describe a generalization of Zilber's Indecomposibility Theorem on groups of finite Morley rank to the finite U-rank setting, which is due to Chantal Berline and Daniel Lascar in [2]. Berline and Lascar, in fact, develop a theory of indecomposibility in the setting of superstable groups; we present here only the special case when $U(G) < \omega$.

Definition A.2. Let X be a type-definable set in G. Then X is *indecomposible* if, for each definable subgroup $H \leq G$, either X is contained in a single coset of H or X meets infinitely many distinct cosets of H.

Intuitively the notion of indecomposibility is a kind of "connectedness" for subsets of G which might not be subgroups. Recall that a type-definable group is *connected* if it has no finite index proper definable subgroups. Connectedness and indecomposibility agree for subgroups of G.

For another example of indecomposibility, the set of realizations of a complete stationary type is indecomposible. Translates of indecomposible sets are also indecomposible.

The following theorem is a generalization of Zilber's Indecomposibility Theorem on groups of finite Morley rank, and is the case $\alpha = \beta = 0$ of [2, Theorem V.3.1].

Theorem A.3 (Zilber Indecomposibility)**.** *Suppose that $\{X_i : i \in I\}$ is a set of indecomposible type-definable subsets of G such that $0_G \in X_i$ for each i. Let $H = \langle X_i : i \in I\rangle$ be the (abstract) subgroup generated by the X_i's. Then H is type-definable and connected, and there exist $i_1, \ldots, i_n \in I$ such that*

$$H = X_{i_1} + X_{i_2} + \cdots + X_{i_n}$$

In particular, if the X_i's are (relatively) definable in G then H is also.

A.2 The Socle

Definition A.4. The *socle* of G, $\mathrm{socle}(G)$, is the subgroup of G generated by all connected type-definable semiminimal subgroups of G.

Lemma A.5. $\mathrm{socle}(G)$ *is a connected, type-definable, semipluriminimal subgroup of G. Moreover, there exist $G_1, \ldots, G_l$, connected, type-definable, semiminimal subgroups of G, such that* $\mathrm{socle}(G) = G_1 + \ldots + G_l$ *and for each $1 \leq i \leq n$, $G_i \perp \sum_{j \neq i} G_j$.*

Proof. Let $\mathcal{F} = \{H \leq G : H \text{ is connected, type-definable, and semiminimal}\}$. So $\mathrm{socle}(G) = \langle \mathcal{F} \rangle$. Then any $H \in \mathcal{F}$ is indecomposable and contains 0_G, so we can apply the finite U-rank version of the Zilber Indecomposibility Theorem (Theorem A.3) to $\mathcal{F}$. We get that $\mathrm{socle}(G)$ is connected and type-definable in G, and $\mathrm{socle}(G) = H_1 + \cdots + H_l$ for some $H_1, \ldots, H_l \in \mathcal{F}$. We choose l to be minimal with the property that such a decomposition of $\mathrm{socle}(G)$ exists. Suppose that $H_i \subseteq \mathrm{acl}(X_i \cup F_i)$. Let $F = \cup_{i=1}^{l} F_i$. Then $\mathrm{socle}(G) \subseteq \mathrm{acl}(X_1 \cup \ldots \cup X_l \cup F)$, so $\mathrm{socle}(G)$ is semipluriminimal. It remains to show that we can choose the H_i's to have the desired orthogonality property.

Fix a minimal set X, and let $\mathcal{H}_X$ denote the collection of all connected type-definable subgroups H of G such that $H \subseteq \mathrm{acl}(X \cup F)$ for some finite set F. By the Zilber Indecomposibility Theorem $\langle \mathcal{H}_X \rangle \in \mathcal{H}_X$. For each $1 \leq i \leq l$ let $G_i = \langle \mathcal{H}_{X_i} \rangle$, so $G_i \supseteq H_i$ for all i. By definition of $\mathrm{socle}(G)$ we have $G_i \subseteq \mathrm{socle}(G)$ for all i, so the decomposition $\mathrm{socle}(G) = H_1 + \ldots + H_l$ implies $\mathrm{socle}(G) = G_1 + \ldots + G_l$ as well. If $X \not\perp Y$ are minimal then they are, up to finite sets, interalgebraic, and so $\langle \mathcal{H}_X \rangle = \langle \mathcal{H}_Y \rangle$. So it follows from minimality of l that $G_i \perp G_j$ for all $i \neq j$.

Now fix any $1 \leq i \leq n$. Note that by the Zilber Indecomposibility Theorem, since the G_i's are connected, $\sum_{j \neq i} G_j$ is indeed type-definable, so the claim $G_i \perp \sum_{j \neq i} G_j$ makes sense. By Lemma 3.5, since $G_i \perp G_j$ for all $i \neq j$, we have $G_i \perp \prod_{j \neq i} G_j$. Take any $a \in G_i$, $b \in \sum_{j \neq i} G_j$, and A parameters over which all the G_i's are defined. Write $b = \sum_{j \neq i} b_j$ for some $b_j \in G_j$. Then $a \mathop{\smile\!\!\!\!\!\!|}_A (b_j)_{j \neq i}$, and hence $a \mathop{\smile\!\!\!\!\!\!|}_A b$ as required. □

Theorem A.6. socle(G) *is the unique maximal connected type-definable semiplurminimal subgroup of* G.

Proof. The previous lemma showed the socle(G) is a connected type-definable semipluriminimal subgroup of G. Note also that the sum $H_1 + H_2$ of type-definable, connected, semipluriminimal subgroups is again type-definable, connected, and semipluriminimal. Indeed, $\langle H_1, H_2\rangle = H_1 + H_2$, and by Theorem A.3 $\langle H_1, H_2\rangle$ is type-definable and connected. It is clear that the sum of semipluriminimal groups is again semipluriminimal. So maximality will imply unique maximality. All that is left is to see that socle(G) is maximal amongst connected type-definable semiplurminimal subgroups of G.

Towards a contradiction, suppose that H is a connected type-definable semipluriminimal subgroup of G such that socle$(G) \subsetneq H$. As a type-definable subgroup of H, socle(G) is an intersection of definable subgroups of H (see Fact 1.8). As socle$(G) \neq H$ some of these definable subgroups are proper, and hence of strictly smaller U-rank than H since H is connected. So $U(\text{socle}(G)) < U(H)$. We will find a type-definable connected semiminimal subgroup of H that meets infinitely many cosets of socle(G), contradicting the definition of socle(G).

We have $H \subseteq \text{acl}(Y_1 \cup \ldots \cup Y_d \cup A)$ where $Y_1, \ldots, Y_d$ are minimal and A is a finite set. Expand A if necessary so that $G, \text{socle}(G), H, Y_1, \ldots, Y_d$ are all defined over A. Let $a \in H$ be generic over A. Let $\overline{a}$ be a finite tuple which is a code for $a + \text{socle}(G)$ over A, as in Proposition A.1. If $\overline{a} \in \text{acl}(A)$ then $a + \text{socle}(G)$ is type-definable over acl(A), and hence we can compute:

$$\begin{aligned} U(H) &= U(a/A) \\ &= U(a/\operatorname{acl}(A)) \\ &\leq U(a + \text{socle}(G)) \\ &= U(\text{socle}(G)) \end{aligned}$$

This contradicts the above observation that $U(H) > U(\text{socle}(G))$, so $\overline{a} \notin \text{acl}(A)$.

Since $a \in \text{acl}(Y_1 \cup \ldots \cup Y_d \cup A)$ there is some finite $Y \subseteq Y_1 \cup \ldots \cup Y_d$ such that $a \in \text{acl}(Y \cup A)$. We may choose Y to be minimal among all such subsets. Note that this means that Y is an acl-independent set over A. Let $Y' \subseteq Y$ be such that $\overline{a} \notin \text{acl}(Y' \cup A)$, and take Y' to be maximal among all such subsets of Y. We know that such Y' exist since $\overline{a} \notin \text{acl}(A)$. Also, as $\overline{a} \in \text{acl}(Y \cup A)$, $Y' \neq Y$.

Claim A.6.1. *Fix* $y_0 \in Y \setminus Y'$. *Both* a *and* $\overline{a}$ *are interalgebraic with* y_0 *over* $A \cup (Y \setminus \{y_0\})$.

Proof of Claim A.6.1 We already have $\overline{a} \in \text{acl}(A \cup \{a\}) \subseteq \text{acl}(A \cup Y)$. By the maximal choice of Y' we have $\overline{a} \in \text{acl}(A \cup Y' \cup \{y_0\}) \setminus \text{acl}(A \cup Y')$. So $\overline{a} \not\perp_{A \cup Y'} y_0$. Then since

y_0 comes from some Y_i, which is minimal, symmetry gives $y_0 \in \operatorname{acl}(A \cup Y' \cup \{\overline{a}\}) \subseteq \operatorname{acl}(A \cup (Y \setminus \{y_0\}) \cup \{\overline{a}\})$. So $\overline{a}$ and y_0 are interalgebraic over $A \cup (Y \setminus \{y_0\})$.

Next we consider a. We again already have $a \in \operatorname{acl}(A \cup Y)$. On the other hand, since $\overline{a} \not\perp\!\!\!\perp_{A\cup Y'} y_0$ and $\overline{a} \in \operatorname{dcl}(A \cup a)$, we also have $a \not\perp\!\!\!\perp_{A\cup Y'} y_0$. So as above symmetry and minimality of the Y_i's gives $y_0 \in \operatorname{acl}(A \cup Y' \cup \{a\})$. So a is interalgebraic with y_0 over $A \cup (Y \setminus Y')$. $\dashv$

Let $B = A \cup (Y \setminus \{y_0\})$, $p = \operatorname{tp}(\overline{a}/\operatorname{acl}(B))$, and $q = \operatorname{tp}(a/\operatorname{acl}(B))$. Then since $U(y_0/B) = 1$ the above shows that $U(p) = U(q) = 1$. It follows that $X = q^{\mathcal{U}}$ is minimal and indecomposible, and hence $\langle X - a\rangle$ is a connected type-definable semiminimal subgroup of H. Hence $\langle X - a\rangle \subseteq \operatorname{socle}(G)$ by definition of $\operatorname{socle}(G)$. On the other hand, since $U(p) > 0$, $\overline{a}$ has infinitely many conjugates under automorphisms fixing B, and so $a + \operatorname{socle}(G)$ also has infinitely many such conjugates, each of which is a coset of $\operatorname{socle}(G)$. Moreover, as X is type-definable over B and $a + \operatorname{socle}(G)$ intersects X, each of these conjugates of $a + \operatorname{socle}(G)$ also intersects X. Thus X intersects infinitely many cosets of $\operatorname{socle}(G)$, and hence so does $\langle X - a\rangle$. This is the desired contradiction. □

Corollary A.7. *Suppose that G is connected and semipluriminimal. Then there exist connected, type-definable, semiminimal subgroups $G_1, \ldots, G_l$ such that $G = G_1 + \ldots + G_l$ and for each $1 \le i \le l$, $G_i \perp \sum_{j\neq i} G_j$.*

Proof. The above theorem shows that if G is semipluriminimal then $G = \operatorname{socle}(G)$, so the result follows immediately from Lemma A.5. □

A.3 The Socle Theorem

Theorem A.8. *Let $A = \operatorname{acl}(A)$ be parameters over which G and* $\operatorname{socle}(G)$ *are defined. Take $a \in G$, and let $p(x) = \operatorname{tp}(a/A)$. Suppose that*

1. *Every connected type-definable subgroup of* $\operatorname{socle}(G)$ *is defined over A.*
2. $\operatorname{stab}(p)$ *is finite. (See Section 3.2 for the definition of* $\operatorname{stab}(p)$*.)*

Then all the realizations of p are contained in a single coset of $\operatorname{socle}(G)$*.*

Proof. The proof is by contradiction. We aim to produce a connected type-definable subgroup of G which is semipluriminimal and not contained in $\operatorname{socle}(G)$, contradicting Proposition 3.17.

Let $\overline{a}$ be a finite tuple such that $\overline{a}$ codes $a + \mathrm{socle}(G)$ over A (see Proposition A.1). Let $X = p^{\mathcal{U}}$. We write $G_{\overline{a}}$ for $a + \mathrm{socle}(G)$. As $A = \mathrm{acl}(A)$, if $\mathrm{tp}(\overline{a}/A)$ is algebraic then $\overline{a} \in A$, and so $X \subseteq G_{\overline{a}} = a + \mathrm{socle}(G)$, as desired. So we may assume that $\mathrm{tp}(\overline{a}/A)$ is non-algebraic. Let $X_{\overline{a}} = G_{\overline{a}} \cap X$.

The next step of the proof is to show that there is a complete algebraic type in $G_{\overline{a}}$ over parameters from $\mathrm{socle}(G) \cup A \cup \{\overline{a}\}$. Consider $\mathcal{F}$, the collection of all complete global types $\mathbf{q}$ extending $G_{\overline{a}}$ and such that $\mathrm{cb}(\mathbf{q}) \subseteq \mathrm{acl}(\mathrm{socle}(G) \cup A \cup \{\overline{a}\})$. Note that $G_{\overline{a}}$ is a partial type over $A \cup \{a\}$, so any global non-forking extension of any completion of $G_{\overline{a}}$ is in $\mathcal{F}$; in particular, $\mathcal{F} \neq \emptyset$. Let $\mathbf{q} \in \mathcal{F}$ be of minimal U-rank amongst the elements of $\mathcal{F}$, and let q be the restriction of $\mathbf{q}$ to $C = \mathrm{acl}(\mathrm{cb}(\mathbf{q}) \cup A \cup \{\overline{a}\})$. Let $Q = q^{\mathcal{U}}$, $H = \mathrm{stab}(q) \cap \mathrm{socle}(G)$, and H° be the connected component of H.

Claim A.8.1. *For any $c, c' \models q$ and any parameter set D such that $C \subseteq D \subseteq \mathrm{acl}(\mathrm{socle}(G) \cup A \cup \{\overline{a}\})$ we have $c \mathop{\smile\!\!\!\!\!\mid}_C D$ and $\mathrm{tp}(c/D) = \mathrm{tp}(c'/D)$.*

Proof of Claim A.8.1 Let q' be any extension of q to D, and let $\mathbf{q}'$ be a global non-forking extension of q'. Then $\mathbf{q}' \in \mathcal{F}$, and $U(\mathbf{q}) = U(q) \leq U(q') = U(\mathbf{q'})$, so by minimality of $U(\mathbf{q})$ we get $U(q') = U(q)$, that is, q' is a non-forking extension of q. So every extension of q to D is non-forking, and so $c \mathop{\smile\!\!\!\!\!\mid}_C D$.

The above shows that $\mathrm{tp}(c/D)$ and $\mathrm{tp}(c'/D)$ are both non-forking extensions of q to D. Since q is over an algebraically closed set of parameters it is stationary. In particular it has a unique non-forking extension to D, so $\mathrm{tp}(c/D) = \mathrm{tp}(c'/D)$. ⊣

Claim A.8.2. *Q is invariant under translation by elements of H.*

Proof of Claim A.8.2 Take any $c \in Q, h \in H$. We must show $h + c \in Q$. By Claim A.8.1 we have $c \mathop{\smile\!\!\!\!\!\mid}_C h$. Since $h \models \mathrm{stab}(q)$ and $c \models q$ we have $h + c \models q$ (see Fact 1.14), as desired. ⊣

Claim A.8.3. *$X_{\overline{a}}$ is invariant under translation by elements of H.*

Proof of Claim A.8.3 Fix any $a' \in X_{\overline{a}}, h \in H$. We show that $a' + h \in X_{\overline{a}}$. Let $c \in Q$ be arbitrary. Since $\mathbf{q}$ extends $G_{\overline{a}} = a + \mathrm{socle}(G)$ we have $c = a + t_1$ for some $t_1 \in \mathrm{socle}(G)$. Also $X_{\overline{a}} \subseteq a + \mathrm{socle}(G)$, so $a' = a + t_2$ for some $t_2 \in \mathrm{socle}(G)$. Let $s = t_2 - t_1 \in \mathrm{socle}(G)$. Then $a' = s + c$.

As $c \in Q$ and $h \in H$ Claim A.8.2 gives $h + c \in Q$. By Claim A.8.1 we get $\mathrm{tp}(c/C \cup \{s\}) = \mathrm{tp}(h + c/C \cup \{s\})$. Let $\alpha \in \mathrm{Aut}_{C \cup \{s\}}(\mathcal{U})$ be such that $\alpha(c) = h + c$. Then $\alpha(a') = \alpha(s + c) = s + \alpha(c) = s + h + c = a' + h$. As α fixes $A \cup \{\overline{a}\}$ pointwise, $a' + h \in X_{\overline{a}}$. ⊣

Claim A.8.4. $H^{\circ} \subseteq \mathrm{stab}(p)$.

Proof of Claim A.8.4 By assumption (1) the type-definable subgroup H° is defined over A. Fix $d \in H^\circ$. Let d' realize a non-forking extension of $\mathrm{tp}(d/A)$ to $A \cup \{\overline{a}\}$, so $d' \mathop{\smile\!\!\!\!\!\!\mid}_A \overline{a}$. As $\mathrm{stab}(p)$ is defined over A it suffices to show that $d' \in \mathrm{stab}(p)$. Pick $a' \in X_{\overline{a}}$ such that $a' \mathop{\smile\!\!\!\!\!\!\mid}_{A\cup\{\overline{a}\}} d'$. So $a' \mathop{\smile\!\!\!\!\!\!\mid}_A d'$ by transitivity. Also $a' \models p$ since $a' \in X_{\overline{a}} \subseteq X$. Since $X_{\overline{a}}$ is H° translation-invariant $d' + a' \in X_{\overline{a}}$, and hence $(d' + a') \models p$. By Fact 1.14 this shows that $d' \in \mathrm{stab}(p)$. $\dashv$

By Claim A.8.4 and the hypothesis that $\mathrm{stab}(p)$ is finite we have that H° is finite (hence actually $H^\circ = \{0\}$), and so H is finite since $U(H^\circ) = U(H)$. Now fix any $c \in Q$. If $z \in Q$ then, since $Q \subseteq a + \mathrm{socle}(G)$, we have $z - c \in \mathrm{socle}(G)$. By Claim A.8.1 $c \mathop{\smile\!\!\!\!\!\!\mid}_A z - c$. As $(z - c) + c = z \models q$, this gives $z - c \in \mathrm{stab}(q)$. So $z - c \in H$. We have shown $Q \subseteq H + c$. Hence Q is finite. In fact, as q is stationary Q is a singleton, $Q = \{c\}$.

So we have found $c \in G_{\overline{a}}$ such that $c \in \mathrm{acl}(\mathrm{socle}(G) \cup A \cup \{\overline{a}\})$. Let $A_0 \subseteq A$ be a finite set, $\phi(x, y, z)$ an L_{A_0}-formula, and e a tuple from $\mathrm{socle}(G)$ be such that $\phi(x, \overline{a}, e)$ defines a finite subset of $G_{\overline{a}}$ containing c. Since $\overline{a}$ is not algebraic over A there is a minimal type r extending $\mathrm{tp}(\overline{a}/A)$. Without loss of generality, we may assume that $r = \mathrm{tp}(\overline{a}/B)$ for some algebraically closed $B \supseteq A$.

Since $\overline{a}$ codes $a+\mathrm{socle}(G) = c+\mathrm{socle}(G)$ over A, and $c+\mathrm{socle}(G)$ is defined over $A\cup\{c\}$, we have that $\overline{a} \in \mathrm{dcl}(A\cup\{c\})$. Let f be a (partial) A-definable function such that $f(c) = \overline{a}$. Let $W = \mathrm{tp}(c/B)^{\mathcal{U}}$. We have that $\phi(x, f(c), e)$ defines a finite subset of $c + \mathrm{socle}(G)$, and $\models \phi(c, f(c), e)$. Now suppose that $c' \in W$. Then there is some $\alpha \in \mathrm{Aut}_B(\mathcal{U})$ such that $\alpha(c) = c'$. Note that $f(c') \models r$. Indeed, $f(c) = \overline{a} \models r$, and we have:

$$\begin{aligned} f(c') &= f(\alpha(c)) \\ &= \alpha(f(c)) \\ &= \alpha(\overline{a}) \end{aligned}$$

So $f(c') \models r$ as well. Moreover, $\phi(x, f(c'), \alpha(e))$ defines a finite subset of $c' + \mathrm{socle}(G)$, and $\models \phi(c', f(c'), \alpha(e))$. Since $\alpha(e)$ is also a tuple from $\mathrm{socle}(G)$, this shows that $W \subseteq \mathrm{acl}(A_0 \cup \mathrm{socle}(G) \cup r^{\mathcal{U}})$. Since r is minimal, $\mathrm{socle}(G)$ is semipluriminimal, and A_0 is finite, this shows that W is semipluriminimal.

We next see that W meets infinitely many cosets of $\mathrm{socle}(G)$ in G. We have that $f(c)$ is a code for $c + \mathrm{socle}(G)$ over A, and $f(c) \models r$. Since $U(r) = 1$ this means, in particular, that $f(c) \notin \mathrm{acl}(B)$. So there are infinitely many conjugates of $f(c)$ under automorphisms fixing B pointwise. Hence there are infinitely many conjugates of $c + \mathrm{socle}(G)$ under automorphisms fixing B pointwise. Each such automorphism permutes cosets of $\mathrm{socle}(G)$ and preserves W, so W meets infinitely many cosets of $\mathrm{socle}(G)$.

Since B is algebraically closed W is the set of realizations of a complete stationary type, and hence is indecomposible. Let W' be a translate of W such that $0_G \in W'$. Then $\langle W' \rangle$

is a connected type-definable semipluriminimal subgroup of G, and so $\langle W' \rangle \subseteq \mathrm{socle}(G)$ by Theorem A.6. This contradicts the fact that $\langle W' \rangle$ meets infinitely many cosets of socle(G). □

Appendix B

Miscellany

B.1 Model Theory

In this section we prove some model-theoretic facts which were needed in the proof of Mordell-Lang. Throughout we assume that we are working inside a universal domain $\mathcal{U}$ of a complete theory T which admits elimination of imaginaries. Every group we encounter in the proof of Mordell-Lang is commutative, so we assume that all groups in this chapter are commutative, and we write them additively.

Lemma B.1. *Let $\{N_h : h \in H\}$ be a definable family of definable sets. Then there is a definable function $g : H \to \mathcal{U}^m$ such that for each $h \in H$ $g(h)$ is a code for N_h, and $g(h) = g(h') \iff N_h = N_{h'}$.*

Proof. Let A be a parameter set over which H is defined and such that for each $h \in H$, N_h is defined over Ah. Consider any $h \in H$, and let c be a code for N_h with the formula θ, so $\theta(y, c)^{\mathcal{U}} = N_h$, and for any $z \neq c$ we have $\theta(y, z)^{\mathcal{U}} \neq N_h$. Then $c \in \mathrm{dcl}(Ah)$, so there exists an A-definable function g_h on an A-definable neighbourhood of h such that $c = g_h(h)$. Define

$$U_h = \{h' \in H : \mathrm{tp}(h'/A) = \mathrm{tp}(h/A), g_h \text{ is defined on } h', \text{ and } g_h(h') \text{ with } \theta(y, z) \text{ codes } N_{h'}\}\,.$$

Then since g_h is A-definable we have that U_h is also A-definable, and by definition we have that $h \in U_h$ and g_h is defined on U_h. By elimination of imaginaries for T we have that g_h maps into $\mathcal{U}^{r_h}$ for some r_h.

Consider now the collection $\mathcal{F} = \{U_h : h \in H\}$. We have just seen that $\mathcal{F}$ is an A-definable cover of H. That is, $H = \bigcup_{h \in H} U_h$. By saturation we actually only need finitely many of the U_h's, so for some $h_1, \ldots, h_n \in H$ we have $H = U_{h_1} \cup \ldots \cup U_{h_n}$. Let $m =$

$r_{h_1}+\cdots+r_{h_n}$. We view g_{h_i} as a map into $\mathcal{U}^m$ by setting the first $r_{h_1}+\cdots+r_{h_{i-1}}$ coordinates to 0, then the next r_{h_i} coordinates to be g_{h_i}, and then the remaining coordinates to be 0. We still have that $g_{h_i}(h)$ is a code for N_h for any $h \in U_{h_i}$. Define $g : H \to \mathcal{U}^m$ by, for each $h \in H$, $g(h) = g_{h_i}(h)$, where i is minimal such that $h \in U_{h_i}$. Then g is $Ah_1 \dots h_n$-definable, and for each $h \in H$ $g(h)$ is a code for N_h by the definition of the g_{h_i}'s. Note that if $g(h) = g(h')$ then by the construction of g we get that $g_{h_i}(h) = g_{h_i}(h')$ for some i. In particular, $h, h' \in U_{h_i}$, so $\mathrm{tp}(h/A) = \mathrm{tp}(h'/A)$. Let $\sigma \in \mathrm{Aut}_A(\mathcal{U})$ be such that $\sigma(h) = h'$. Then $\sigma(N_h) = N_{h'}$. Also, $\sigma(g(h)) = g(h') = g(h)$, so $\sigma(N_h) = N_h$. Hence $N_h = N_{h'}$. □

Lemma B.2. *Let X be a type-definable set, and let E be a relatively definable equivalence relation on X. Then there exists a type-definable set Y and a relatively definable surjection $f : X \to Y$ such that $f(x) = f(y) \iff xEy$. (We think of Y as X/E.)*

Proof. Let $\Gamma(x) = \{\phi_i(x) : i \in I\}$ be a type such that $\Gamma(x)^{\mathcal{U}} = X$. We may assume that Γ is closed under finite conjunctions.

We first show that there is a definable set $A \supseteq X$ such that E is an equivalence relation on A. Suppose to the contrary that there is no such A. For each i let $\psi_i(x, y, z) = \phi_i(x) \wedge \phi_i(y) \wedge \phi_i(z)$. Let $\Psi(x, y, z) = \{\neg(xEx) \vee \neg(xEy \leftrightarrow yEx) \vee \neg(xEy \wedge yEz \to xEz)\} \cup \{\psi_i(x, y, z) : i \in I\}$. Since Γ is closed under finite conjunctions any finite subset of Ψ is equivalent to $(\neg(xEx) \vee \neg(xEy \leftrightarrow yEx) \vee \neg(xEy \wedge yEz \to xEz)) \wedge \psi_i$ for some i. Since ϕ_i does not define a set on which E is an equivalence relation this formula is realized. Hence by saturation Ψ is realized, contradicting that E is an equivalence relation on X. We let θ be the formula defining the set A whose existence we have just proved.

For each $i \in I$ let $X_i = \phi_i(x)^{\mathcal{U}}$. Since $X = \bigcap_{i \in I} X_i \subseteq A$ we may replace each ϕ_i by $\phi_i \wedge \theta$, and hence we may assume $X_i \subseteq A$ for all i. By elimination of imaginaries there is a definable set A' and a definable surjection $\pi : A \to A'$ such that $\pi(x) = \pi(y) \iff xEy$. Let $Y = \pi(X)$ and $f = \pi|_X$. All that remains is to see that Y is type-definable.

Define $Y_i = \pi(X_i)$ for each $i \in I$. Then each Y_i is a definable set in the same sort as A'. We will show $Y = \bigcap_{i \in I} Y_i$. First suppose $a \in Y$. Then there is $b \in X$ such that $a = \pi(b)$. Since $b \in X$ we have $b \in X_i$ for each i, so $a \in Y_i$ for each i as well. Conversely, suppose that $a \in \bigcap_{i \in I} Y_i$. Then for each $i \in I$ there is $b_i \in X_i$ such that $\pi(b_i) = a$. For each $i \in I$ let $\psi_i(x) =$ "$x \in X_i$" $\wedge$ "$\pi(x) = a$", and let $\Psi(x) = \{\psi_i(x) : i \in I\}$. Since Γ is closed under finite conjunctions any finite subset of Ψ is equivalent to a formula $\psi_i(x)$ for some $i \in I$, and hence is satisfied by b_i. So by saturation Ψ is realized, say by b. Then by definition of Ψ we have $b \in X_i$ for all $i \in I$, so $b \in X$, and $\pi(b) = a$. So $a \in Y$. □

Lemma B.3. *Let X and Y be type-definable sets, and let $f : X \to Y$ be a type-definable function. Then f is relatively definable.*

Proof. Let A be parameters over which X, Y, f are defined. Let $\Phi(x, y) = \{\phi_i(x, y) : i \in I\}$ be a type which defines f. We may assume that Φ is closed under finite conjunctions.

Consider any $a \in X$. It follows from saturation that there is some $\phi \in \Phi$ such that there is a unique solution to $\phi(a, y)$. Indeed, suppose that no such ϕ exists. Let $\Psi(y, z) = \{y \neq z\} \cup \{\phi_i(a, y) \wedge \phi_i(a, z) : i \in I\}$. Let Ψ' be a finite subset of Ψ. Since Φ is closed under finite conjunctions Ψ' is equivalent to the formula $\phi_i(a, y) \wedge \phi_i(a, z) \wedge y \neq z$ for some $i \in I$. Since there is no ϕ such that $\phi(a, y)$ has a unique solution this formula is satisfiable, and hence Ψ is realized. But any $(b, c) \models \Psi$ has, in particular, $(a, b) \models \Phi$ and $(a, c) \models \Phi$, and $b \neq c$. This contradicts that Φ defines the graph of a function.

For each $a \in X$ let $\phi_a \in \Phi$ be such that $\phi_a(a, y)$ has a unique solution. Let $U_a = \{b : \phi_a(b, y) \text{ has a unique solution}\}$, and let $\psi_a(x, y) = \phi_a(x, y) \wedge$ "$x \in U_a$". Then for any $b \in X \wedge U_a$, $\psi_a(b, y)$ has a unique solution, which must be $f(b)$. So we have a definable function g_a on U_a such that $g_a|_{U_a \cap X} = f|_{U_a \cap X}$. Clearly the U_a's cover X, and are definable over A. By saturation only finitely many of the U_a's are required to cover X, say $U_{a_i}, \ldots, U_{a_n}$. As we have seen before (for example, in the proof of Lemma B.1) we can patch the functions $g_{a_1}, \ldots, g_{a_n}$ to get a definable function g on $\bigcup_{i=1}^{n} U_{a_i}$ such that $g|_X = f$. □

Lemma B.4. *Suppose that T is stable. Let $p \in S_n(A)$, $q \in S_m(B)$, with p stationary. Then $X = \left\{(a, b) \in \mathcal{U}^n \times \mathcal{U}^m : a \models p, b \models q, a \mathop{\smile\!\!\!\!\!\!\mid}_A b\right\}$ is type-definable over $A \cup B$.*

Proof. We note that this is clear from the definition of non-forking independence, as $(a, b) \in X$ if and only if $a \models p, b \models q$, and b does not realize any forking formula over Aa. Nevertheless, we have thus far chosen to avoid discussing the definition of forking, so we present an alternative proof using definability of types.

Since p is stationary it has a unique non-forking extension to Ab, which we denote by $p \upharpoonright Ab$. Since T is stable the type $p \upharpoonright Ab$ is definable. That is, for any $\mathcal{L}_A$-formula $\phi(x, y)$, where $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_m)$, there is an $\mathcal{L}_A$-formula $d_p\phi(y)$ such that $\models d_p\phi(b) \iff \phi(x, b) \in p \upharpoonright Ab$.

First, by stationarity of p, observe that $(a, b) \in X$ if and only if $b \models q, a \models p$ and $\mathrm{tp}(a/Ab) = p \upharpoonright Ab$.

Claim B.4.1. *For any a, b the following are equivalent:*

1. $\mathrm{tp}(a/Ab) = p \upharpoonright Ab$.
2. *For L_A-formulae $\phi(x, y)$, $\models \phi(a, b) \iff \models d_p\phi(b)$.*

Proof of Claim B.4.1 First suppose $\mathrm{tp}(a/Ab) = p \upharpoonright Ab$, and let $\phi(x, y)$ be an L_A-formula.

Then

$$\begin{aligned} \models \phi(a,b) &\iff \phi(x,b) \in \mathrm{tp}(a/Ab) \\ &\iff \phi(x,b) \in p \upharpoonright Ab \\ &\iff \models d_p\phi(b) && \text{by definition of } d_p\phi(y). \end{aligned}$$

Conversely, suppose that for any L_A-formula $\phi(x,y)$ we have $\models \phi(a,b) \iff \models d_p\phi(b)$, and let $\phi(x,y)$ be any L_A-formula. Then

$$\begin{aligned} \phi(x,b) \in p \upharpoonright Ab &\iff \models d_p\phi(b) && \text{by definition of } d_p\phi(y). \\ &\iff \models \phi(a,b) \\ &\iff \phi(x,b) \in \mathrm{tp}(a/Ab) \end{aligned}$$

$\dashv$

By the claim $(a,b) \in X$ if and only if $a \models p, b \models q$, and $\models \phi(a,b) \iff \models d_p\phi(b)$ for all L_A-formulae ϕ. So X is defined by $\Gamma(x,y) = p(x) \cup q(y) \cup \{\phi(x,y) \leftrightarrow d_p\phi(y) : \phi \text{ an } L_A\text{-formula.}\}$.

□

B.2 Semiabelian Varieties

Here we collect various results about semiabelian varieties which are useful, but which are not readily available in the literature. Many of the results that we prove in this section generalize similar statements for abelian varieties; the reader desiring background on abelian varieties should consult [14] or [20].

Definition B.5. Let S be a connected commutative algebraic group over a field k. We say that S is a *semiabelian variety* if there exists an exact sequence

$$\begin{tikzcd} 1 \arrow[r] & T \arrow[r] & S \arrow[r] & A \arrow[r] & 0 \end{tikzcd}$$

where A is an abelian variety, and $T \cong \mathbb{G}_m^s$ over some algebraically closed $L \geq k$ for some $s \in \mathbb{N}$.

Throughout this appendix we will make the usual identification of T with its image in S, and hence also identify A with S/T.

Lemma B.6. *Let S be a semiabelian variety over k. Then for all $n \in \mathbb{N}$, S has finite n-torsion.*

Proof. Recall that $G[n]$ denotes the n-torsion elements of an algebraic group G. Since S is a semiabelian variety we have a short exact sequence as in the definition:

$$\begin{tikzcd} 1 \arrow[r] & T \arrow[r] & S \arrow[r, "\pi"] & A \arrow[r] & 0 \end{tikzcd}$$

We have that $\pi(S[n]) \subseteq A[n]$ and $\ker(\pi|_{S[n]}) \subseteq T[n]$. So we have another exact sequence

$$\begin{tikzcd} 1 \arrow[r] & \ker(\pi|_{S[n]}) \arrow[r] & S[n] \arrow[r] & \pi(S[n]) \arrow[r] & 0 \end{tikzcd}$$

The abelian variety A has finite n-torsion (see [14, Corollary IV.2.1]), and it is clear that T does as well. So $S[n]$ is finite. □

Lemma B.7. *Let S be a semiabelian variety over an algebraically closed field L, and let $R \leq S$ be a connected algebraic subgroup of S. Then R is itself a semiabelian variety over L.*

Proof. We have the short exact sequence from the definition of S being semiabelian:

$$\begin{tikzcd} 0 \arrow[r] & T \arrow[r] & S \arrow[r, "\pi"] & A \arrow[r] & 0 \end{tikzcd}$$

Now let $R \leq S$ be a connected algebraic subgroup. Then since π is a morphism of algebraic groups $\pi(R)$ is an algebraic subgroup of A, and hence is also projective, so an abelian variety. Also, $R \cap T$ is a connected algebraic subgroup of T. Since L is algebraically closed T is isomorphic over L to $\mathbb{G}_m^s$ for some s. Since any connected algebraic subgroup of $\mathbb{G}_m^s$ is equal to $\mathbb{G}_m^t$ for some $t \leq s$ (see [23, Theorem III.5]), the same isomorphism shows that $R \cap T$ is isomorphic over L to $\mathbb{G}_m^t$ for some $t \leq s$. Since the original sequence is exact so is the sequence

$$\begin{tikzcd} 0 \arrow[r] & R \cap T \arrow[r] & R \arrow[r, "\pi"] & \pi(R) \arrow[r] & 0 \end{tikzcd}$$

So R is semiabelian. □

Proposition B.8. *Let S be a semiabelian variety over a field k. Then any algebraic subgroup of S defined over some algebraically closed $L \geq k$ is in fact defined over k^{alg}.*

Proof. Note that the torsion points of S are all in $S(k^{\mathrm{alg}})$ since S is over k. Moreover, as noted in [5, Remark 3.4], the torsion subgroup of S, S_{tor}, is Zariski-dense in S. Let $H \leq S$ be a connected algebraic subgroup defined over L, so H is again a semiabelian variety. Then $H_{\mathrm{tor}} \subseteq S_{\mathrm{tor}} \subseteq S(k^{\mathrm{alg}})$, and so in fact $H_{\mathrm{tor}} \subseteq H(k^{\mathrm{alg}})$. Since H_{tor} is Zariski-dense in H we have that any $\sigma \in \mathrm{Aut}_{k^{\mathrm{alg}}}(L)$ must fix H setwise, and so H is defined over k^{alg}. □

Proposition B.9. *Let S be a semiabelian variety over an algebraically closed field L, and let $R \leq S$ be an algebraic subgroup of S. Then S/R is a semiabelian variety over L, and the canonical projection $p : S \to S/R$ is a rational map.*

Proof. Since S is semiabelian, we have a short exact sequence:

$$\begin{tikzcd} 1 \arrow[r] & T \arrow[r] & S \arrow[r, "\pi"] & A \arrow[r] & 0 \end{tikzcd}$$

Let $R_T = R \cap T$ and $R_A = \pi(R)$.

Claim B.9.1. *The sequence*

$$\begin{tikzcd} 1 \arrow[r] & T/R_T \arrow[r, "\iota^*"] & S/R \arrow[r, "\pi^*"] & A/R_A \arrow[r] & 0 \end{tikzcd}$$

is exact, where $\iota^*(t + R_T) = t + R$ and $\pi^*(s + R) = \pi(s) + R_A$.

Proof of Claim B.9.1 It is easy to check that ι^* is injective and π^* is surjective. All that remains is to see that $\mathrm{im}(\iota^*) = \ker(\pi^*)$.

First, suppose that $x + R \in \mathrm{im}(\iota^*)$. Then for some $z \in T$, $x + R = z + R$ by definition of ι^*. Hence $\pi^*(x + R) = \pi^*(z + R) = \pi(z) + R_A = 0 + R_A$ so $x + R \in \ker(\pi^*)$.

Conversely, suppose that $a + R \in \ker(\pi^*)$. Then $\pi^*(a + R) = \pi(a) + R_A = 0 + R_A$, so $\pi(a) \in R_A = \pi(R)$. So there exists $r \in R$ such that $\pi(a) = \pi(r)$. Hence $\pi(a - r) = 0$, so as $\ker(\pi) = T$ there exists $t \in T$ such that $a - r = t$. Then $a - t = r \in R$, so $a + R = t + R = \iota^*(t + R_T)$, and so $a + R \in \mathrm{im}(\iota^*)$. Thus $\mathrm{im}(\iota^*) = \ker(\pi^*)$ and the sequence is exact. $\dashv$

That A/R_A has the structure of an abelian variety is shown in [6]. As we have noted earlier, identifying T with $\mathbb{G}_m^s$, we get that $R_T = \mathbb{G}_m^t$ for some $t \leq s$. Hence $T/R_T = \mathbb{G}_m^{s-t}$ is again an algebraic torus, and S/R is semiabelian. □

Lemma B.10. *Let L be an algebraically closed field of characteristic $p > 0$, and let S be a semiabelian variety over L. Then, for any $n \in \mathbb{N}$, $S^{(p^n)}$ is also a semiabelian variety over L.*

Proof. Since S is defined over L so is $S^{(p^n)}$. As S is an algebraic group, so is $S^{(p^n)}$. It remains only to see that $S^{(p^n)}$ is semiabelian. Since S is semiabelian there is a short exact sequence

$$\begin{tikzcd} 1 \arrow[r] & T \arrow[r] & S \arrow[r, "\pi"] & A \arrow[r] & 0 \end{tikzcd},$$

where $T \cong \mathbb{G}_m^s$ over L for some s and A is an abelian variety over L. Hence there is a short exact sequence

$$\begin{tikzcd} 1 \arrow[r] & T^{(p^n)} \arrow[r] & S^{(p^n)} \arrow[r, "\pi"] & A^{(p^n)} \arrow[r] & 0 \end{tikzcd}$$

The Frobenius automorphism $\mathrm{Fr} : k \to k$ lifts to an injective algebraic group homomorphism from A to $A^{(p^n)}$. As the homomorphic image of an abelian variety is again an abelian variety, $A^{(p^n)}$ is an abelian variety. On the other hand, $\mathbb{G}_m$ is fixed by the Frobenius, and so $(\mathbb{G}_m^s)^{(p^n)} = \mathbb{G}_m^s$. In particular, $T^{(p^n)} \cong T$, and so the above exact sequence shows that $S^{(p^n)}$ is a semiabelian variety. □

References

[1] Dan Abramovich and José Felipe Voloch. Toward a proof of the Mordell-Lang conjecture in characteristic p. *International Mathematics Research Notices*, (5):103 – 115, 1992.

[2] Chantal Berline and Daniel Lascar. Superstable groups. *Annals of Pure and Applied Logic*, 30:1 – 43, 1986.

[3] Elisabeth Bouscaren. Model theoretic versions of Weil's theorem on pregroups. In *The Model Theory of Groups*, number 11 in Notre Dame Mathematical Lectures. University of Notre Dame Press, 1989.

[4] Elisabeth Bouscaren. Proof of the Mordell-Lang conjecture for function fields. In Elisabeth Bouscaren, editor, *Model Theory and Algebraic Geometry: An Introduction to E. Hrushovski's Proof of the Geometric Mordell-Lang Conjecture*, number 1696 in Lecture Notes in Mathematics. Springer, 1998.

[5] Elisabeth Bouscaren and Francoise Delon. Minimal groups in separably closed fields. *The Journal of Symbolic Logic*, 67(1):239 – 259, March 2002.

[6] Wei-Liang Chow. On the quotient variety of an abelian variety. *Proceedings of the National Academy of Sciences of the United States of America*, 38(12):1039 – 1044, December 1952.

[7] Françoise Delon. Separably closed fields. In Elisabeth Bouscaren, editor, *Model Theory and Algebraic Geometry: An Introduction to E. Hrushovski's Proof of the Geometric Mordell-Lang Conjecture*, number 1696 in Lecture Notes in Mathematics. Springer, 1998.

[8] Françoise Delon. Model theory of Hasse closed fields. 2007.

[9] Wilfrid Hodges. *Model Theory.* Number 42 in Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1993.

[10] Ehud Hrushovski. The Mordell-Lang conjecture for function fields. *Journal of the American Mathematical Society*, 9:667 – 690, 1996.

[11] Ehud Hrushovski and Zeljko Sokolovic. Minimal subsets of differentially closed fields (unpublished manuscript).

[12] Ehud Hrushovski and Boris Zilber. Zariski geometries. *Journal of the American Mathematical Society*, 9(1):1 – 56, Jan. 1996.

[13] Thomas Hungerford. *Algebra*. Number 73 in Graduate Texts in Mathematics. Springer, 1974.

[14] Serge Lang. *Abelian Varieties*. Springer-Verlag, 1983.

[15] Serge Lang. *Fundamentals of Diophantine Geometry*. Springer-Verlag, 1983.

[16] David Marker. Zariski geometries. In Elisabeth Bouscaren, editor, *Model Theory and Algebraic Geometry: An Introduction to E. Hrushovski's Proof of the Geometric Mordell-Lang Conjecture*, number 1696 in Lecture Notes in Mathematics. Springer, 1998.

[17] David Marker. Model theory of differential fields. In *Model Theory of Fields*. Association for Symbolic Logic, 2nd edition, 2006.

[18] Margit Messmer. Some model theory of separably closed fields. In *Model Theory of Fields*. Association for Symbolic Logic, 2nd edition, 2006.

[19] Margit Messmer and Carol Wood. Separably closed fields with higher derivations. *The Journal of Symbolic Logic*, 60(3), Sept. 1995.

[20] James Milne. Abelian varieties. http://www.jmilne.org/math/CourseNotes/av.html.

[21] Rahim Moosa. Six lectures on Hrushovski's proof of the function field Mordell-Lang conjecture. Unpublished notes, private correspondance, 2007.

[22] Rahim Moosa and Thomas Scanlon. F -structures and integral points on semiabelian varieties over finite fields. *American Journal of Mathematics*, 126(3):473 – 522, June 2004.

[23] A. L. Onishchik and E. B. Vinberg. *Lie Groups and Algebraic Groups*. Springer-Verlag, 1990.

[24] Anand Pillay. *An Introduction to Stability Theory*. Number 8 in Oxford Logic Guides. Oxford University Press, 1983.

[25] Anand Pillay. *Geometric Stability Theory*. Number 32 in Oxford Logic Guides. Oxford University Press, 1996.

[26] Anand Pillay. Model theory of algebraically closed fields. In Elisabeth Bouscaren, editor, *Model Theory and Algebraic Geometry: An Introduction to E. Hrushovski's Proof of the Geometric Mordell-Lang Conjecture*, number 1696 in Lecture Notes in Mathematics. Springer, 1998.

[27] Anand Pillay. Lecture notes - stability theory. September 2003.

[28] Bruno Poizat. Paires de structures stables. *The Journal of Symbolic Logic*, 48(2), June 1983.

[29] Bruno Poizat. *Stable Groups*. Number 87 in Mathematical Surveys and Monographs. American Mathematical Society, 1987.

[30] Maxwell Rosenlicht. Some basic theorems on algebraic groups. *American Journal of Mathematics*, 78(2):401 – 443, April 1956.

[31] Carol Wood. Differentially closed fields. In *Model Theory and Algebraic Geometry: An Introduction to E. Hrushovski's Proof of the Geometric Mordell-Lang Conjecture*. Springer.

[32] Martin Ziegler. Stability and Morley rank. In Elisabeth Bouscaren, editor, *Model Theory and Algebraic Geometry: An Introduction to E. Hrushovski's Proof of the Geometric Mordell-Lang Conjecture*, number 1696 in Lecture Notes in Mathematics. Springer, 1998.