

© 2008 by Yimu Yin. All rights reserved.

SETS, MODELS, AND VALUED FIELDS

BY

YIMU YIN

Department of Philosophy
Carnegie Mellon University

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Logic
in the College of Humanities and Social Sciences of
Carnegie Mellon University, 2008

Pittsburgh, Pennsylvania, U.S.A.

Acknowledgments

I would like to thank my advisor, Jeremy Avigad, for his guidance and constant support during the past five years. I am especially grateful for his patience while listening to me, very often having to endure nonsense.

I wish to thank the following who have helped me obtain the results contained in this thesis: Uri Abraham, James Cummings, Rami Grossberg, Thomas Hales, and Chris Miller.

A special thanks goes to Haim Gaifman, who inspired me to study logic.

It has been a real pleasure to work in the Department of Philosophy at Carnegie Mellon University. I thank everyone there for providing such a free intellectual environment.

This thesis is based upon work partially supported by the National Science Foundation under Grant No. 0700174.

Table of Contents

Introduction	1
Chapter 1 The Engelking-Karłowicz theorem: A case study in infinitary combinatorics	2
1.1 Introduction	2
1.2 $\chi(G(\omega, \{\beth_\omega\})) > \beth_\omega$ (and similar results)	5
1.3 Ladder graphs	9
1.4 Graphs of the form $G(\omega, \mu)$	13
Chapter 2 Combinatorics in model theory: Quantifier elimination tests	15
2.1 Introduction	15
2.2 Some QE tests	16
2.3 The base case of the induction	20
2.4 The inductive step	24
Chapter 3 Quantifier elimination for the reals with a predicate for the powers of two .	28
3.1 Introduction	28
3.2 The first step	30
3.3 Reasoning about powers of two	32
3.4 Eliminating a quantifier over powers of two	38
3.5 Complexity analysis	41
3.6 Adding a predicate for the Fibonacci numbers	46
Chapter 4 Henselianity and the Denef-Pas language	50
4.1 Introduction	50
4.2 Preliminaries	51
4.3 Overview of the proof	52
4.4 Henselianity and Denef-Pas style languages	54
4.4.1 Step 1: Open sets	56
4.4.2 Step 3 and 4: Omitting a type	58
Chapter 5 Grothendieck homomorphisms in algebraically closed valued fields	63
5.1 Introduction	63
5.1.1 Outline of the construction	64
5.1.2 Course of the paper	65
5.1.3 Technical differences from the original construction	68
5.2 Logical preliminaries and the theory ACVF	70
5.2.1 The setting of \mathcal{L}_{RV} and ACVF	70
5.2.2 Compactness	76
5.3 Quantifier elimination in ACVF	80
5.4 Basic structural properties	86
5.5 Parametric balls and atomic subsets	92
5.6 Categories of definable subsets	98

5.6.1	Dimensions	98
5.6.2	Categories of definable subsets	99
5.6.3	Grothendieck groups	102
5.7	RV-products and special bijections	104
5.8	2-cells	111
5.9	Lifting functions from RV to VF	115
5.10	Contracting to RV	119
5.11	The kernel of \mathbb{L}	124
5.11.1	Blowups in RV and the congruence relation I_{sp}	124
5.11.2	Blowups and special bijections	130
5.12	Motivic integration	136
References		138

Introduction

Each chapter has its own introduction. There is not a particular theme that runs through these chapters. However, they do show a tendency of moving away from the pure set-theoretic world and gradually making more contact with model-theoretic algebra, or, as some would like to put it, definability theory. Some of the chapters have already been published.

Chapter 1

The Engelking-Karłowicz theorem: A case study in infinitary combinatorics

Abstract: We investigate the chromatic number of infinite graphs whose definition is motivated by the theorem of Engelking and Karłowicz (in [22]). In these graphs, the vertices are subsets of an ordinal, and two subsets X and Y are connected iff for some $a \in X \cap Y$ the order-type of $a \cap X$ is different from that of $a \cap Y$.

In addition to the chromatic number $\chi(G)$ of these graphs we study $\chi_\kappa(G)$, the κ -chromatic number, which is the least cardinal μ with a decomposition of the vertices into μ classes none of which contains a κ -complete subgraph.¹

1.1 Introduction

A celebrated theorem of Engelking and Karłowicz [22] states that if θ and μ are cardinals such that $\mu^{<\theta} = \mu$, then there is a family \mathcal{F} of size 2^μ , consisting of functions from μ into μ , with the following property. For every one-to-one sequence $\langle f_i \in \mathcal{F} \mid i \in \theta^* \rangle$ and sequence $\langle \beta_i \in \mu \mid i \in \theta^* \rangle$, where $\theta^* < \theta$, there exists some $\alpha \in \mu$ such that for all $i \in \theta^*$ $f_i(\alpha) = \beta_i$.

An equivalent formulation takes the following form. Let θ and μ be cardinals such that $\mu^{<\theta} = \mu$. Then there are functions $f_\xi : 2^\mu \rightarrow \mu$, for $\xi < \mu$, such that if $X \subset 2^\mu$, $|X| < \theta$ and $f : X \rightarrow \mu$, then there is $\xi < \mu$ such that $f \subset f_\xi$.

This theorem has diverse applications such as the Hewitt-Marczewski-Pondiczery theorem that the product of 2^μ topological spaces each with a dense subset of cardinality μ has itself a dense subset of cardinality μ . We are interested here in the following corollary used by Shelah in [45] and [46]:

Corollary 1.1.1. *1. If $\mu^{<\theta} = \mu$ and A is any set of cardinality 2^μ , then there is a map $\tau : [A]^{<\theta} \rightarrow \mu$ such that whenever $\tau(M_1) = \tau(M_2)$ then M_1 and M_2 have the same order-type and the order-preserving isomorphism $g : M_1 \rightarrow M_2$ is the identity on $M_1 \cap M_2$.*

2. Thus, if $\mu^\theta = \mu$ and A is any set of cardinality 2^μ , then there is a map $\tau : [A]^\theta \rightarrow \mu$ such that whenever

¹This is joint work with Uri Abraham, published as [3].

$\tau(M_1) = \tau(M_2)$ then M_1 and M_2 have the same order-type and the order-preserving isomorphism $g : M_1 \longrightarrow M_2$ is the identity on $M_1 \cap M_2$.

For example, if $\mu^\theta = \mu$ and $\lambda < (2^\mu)^+$ is any ordinal, then there is a map $\tau : [\lambda]^{<\theta} \longrightarrow \mu$ such that $\tau(M_1) = \tau(M_2)$ implies that M_1 and M_2 are isomorphic with an isomorphism that is the identity on $M_1 \cap M_2$.

Proof. Here $[A]^{<\theta}$ is the collection of all subsets of A of cardinality $< \theta$. Since there are only $\theta \leq \mu$ possible order-types of $M \in [A]^{<\theta}$, it is enough to find a function that works for a specific order-type $\gamma < \theta$ and then to combine these functions into a single τ that works for all $\gamma < \theta$. Note that the requirement that the order-isomorphism $g : M_1 \longrightarrow M_2$ is the identity on $M_1 \cap M_2$ can be expressed by saying that every $a \in M_1 \cap M_2$ has the same place in M_1 as in M_2 (namely the order-types of $a \cap M_1$ and $a \cap M_2$ are the same).

Fix a sequence of functions $f_\xi : 2^\mu \longrightarrow \mu$ for $\xi < \mu$, as in the equivalent formulation of the Engelking and Karłowicz theorem. Since A has cardinality 2^μ , we can have such functions be defined over A and with the same properties. Namely, that if $X \in [A]^{<\theta}$ and $f : A \longrightarrow \mu$ then $f \subset f_\xi$ for some $\xi < \mu$.

For every $M \subset A$ of order-type γ , let $f : M \longrightarrow \gamma$ be its order-preserving collapse. There is some $\xi < \mu$ such that $f \subset f_\xi$ and we define $\tau(M) = \xi$ (say for the least such ξ). Now if $\tau(M_1) = \tau(M_2) = \xi$ then for $a \in M_1 \cap M_2$, $f_\xi(a)$ is the place of a both in M_1 and in M_2 .

The second paragraph of the corollary is obtained by replacing θ with θ^+ . It is this case that will interest us in this paper. □

In this note we want to investigate the extent to which the assumption $\mu^\theta = \mu$ in the second item of the corollary is necessary. We are mainly interested in the case $\mu = \aleph_\omega$ and $\theta = \omega$ and we will prove that if μ is a strong limit singular cardinal of cofinality θ then the conclusion of the corollary does not hold.

In graph theoretic language our problem finds a concise formulation as follows. Let $(A, <)$ be any linearly ordered-set. We say that X and Y subsets of A are *consistent* if and only if there exists an order isomorphism $f : X \longrightarrow Y$ that is the identity on $X \cap Y$ (namely $f(x) = x$ for $x \in X \cap Y$). We say that X and Y are *inconsistent* if they are not consistent. In this paper, we deal only with well-ordered sets $(A, <)$, and in this case X and Y are inconsistent if and only if either the order-type of X is different from that of Y , or else there exists some $\xi \in X \cap Y$ such that $order\text{-type}(X \cap \xi) \neq order\text{-type}(Y \cap \xi)$.

For any ordered-set $(A, <)$, we define a graph with vertices $\mathcal{P}(A)$, the powerset of A , and with edges all pairs (X, Y) where X and Y are inconsistent subsets of A . We will be interested in subgraphs of $\mathcal{P}(\gamma)$ for different ordinals γ 's and ask for their chromatic number. In fact, we will be interested here mainly in the

case in which we take only subsets of γ of some fixed order-type α .

Recall that the chromatic number $\chi(G)$ of a graph $G = (V, E)$ is the least cardinal κ such that there is a function $\tau : V \rightarrow \kappa$ so that $(a, b) \notin E$ whenever $\tau(a) = \tau(b)$. We call such a function “separating”. That is, τ is separating if and only if $\tau(a) \neq \tau(b)$ whenever $(a, b) \in E$. The chromatic number is thus the least cardinality of the range of a separating function.

Let $\text{typ}(\alpha)$ be the class of all sets of ordinals of order-type α . If $X \in \text{typ}(\alpha)$ we say that X is an α -set. (We take sets rather than sequences because we refer to the intersection of two sets when the edges of the graph are defined). For a set B of ordinals, let $\text{typ}(\alpha, B)$ be the collection of all $X \in \text{typ}(\alpha)$ such that $\text{sup}(X) \in B$ (where $\text{sup}(X)$ is the first ordinal greater or equal to all ordinals of X). We will be interested here in two cases: for ordinals $\alpha < \beta$, $\text{typ}(\alpha, \beta)$, is the collection of bounded subsets of β of order-type α , and for a limit ordinal α $\text{typ}(\alpha, \{\beta\})$ is the collection of all unbounded subsets of β of order-type α .

If $a, b \in \text{typ}(\alpha)$ there is a unique order-preserving isomorphism $g : a \rightarrow b$, and in this case a and b are consistent iff g is the identity on $a \cap b$. They are inconsistent otherwise. So a and b are inconsistent if and only if for some $x \in a \cap b$ the order-type of $x \cap a$ differs from that of $x \cap b$.

For two ordinals $\alpha \leq \beta$ let $G(\alpha, \beta)$ be the graph $G = (V, E)$ with set of vertices $V = \text{typ}(\alpha, \beta)$ and edges $(a, b) \in E$ if and only if a and b are inconsistent. Likewise, $G(< \alpha, \beta)$ consists of subsets of β of order-type $< \alpha$, with edges (a, b) defined whenever the order-type of a is different from that of b , or else they have the same order-type but are inconsistent.

For example, $G(\omega, \omega_1)$ has vertices all ω -sets of countable ordinals, and ω -sets X and Y are connected iff some $a \in X \cap Y$ has different position in X and Y . The graph $G(< \omega, \omega_1)$, has vertices all finite subsets of ω_1 , and edges all pairs (a, b) where a and b are inconsistent.

Similarly, $G(\alpha, \{\beta\})$ is the graph with vertices $\text{typ}(\alpha, \{\beta\})$ and edges all pairs (a, b) that are inconsistent α sets (unbounded in β).

The graphs $G(2, \beta)$ were considered by Erdős and Hajnal [25] and called “shift graphs”. So a vertex is a pair $\{a, b\}$ (with $a < b$) and two pairs $a_0 < a_1$ and $b_0 < b_1$ are connected in the graph if and only if $a_1 = b_0$ or $b_1 = a_0$.

We shall be particularly interested in the case $G(\omega, \{\aleph_\omega\})$ which is the graph G with set of vertices all unbounded ω -sets in \aleph_ω , and with edges defined by $(s, t) \in G$ if and only if there is $x \in s \cap t$ such that $|x \cap s| \neq |t \cap x|$.

Our aim is to investigate the chromatic number of these and similarly defined graphs.

In this graph theoretic terminology, Corollary 1.1.1 can be restated as follows:

Corollary 1.1.2. *Suppose that $\mu^{<\theta} = \mu$, and let λ be any ordinal of cardinality $\leq 2^\mu$. Let G be the graph*

with vertices all subsets of λ of cardinality $< \theta$ and edges connecting two vertices if and only if they are inconsistent. Then the chromatic number of G is $\leq \mu$.

In particular, we get the following when we consider θ^+ . Suppose that $\mu^\theta = \mu$, and let λ be any ordinal of cardinality 2^μ . Let G be the graph with vertices all subsets of λ of cardinality θ and edges connecting two vertices if and only if they are inconsistent. Then the chromatic number of G is $\leq \mu$.

Here are a couple of illustrations of the corollary. Since $\mu^{<\aleph_0} = \mu$ for every infinite cardinal μ , we have that $\chi(G(< \omega, 2^\mu)) \leq \mu$. Another example: $\chi(G(< \omega_1, 2^{(2^{\aleph_0})})) \leq 2^{\aleph_0}$. In fact, $\chi(G(< \omega_1, 2^{(2^{\aleph_0})})) = 2^{\aleph_0}$, because already $G(\omega, \{\omega\})$ contains a clique of size 2^{\aleph_0} . To see this, form for every subset X of the even numbers the set $S(X)$ which is the union of X with the set of odd numbers. Then $\{S(X) \mid X \subseteq \text{even}\}$ is a clique.

Corollary 1.1.2 thus says that if $\mu^\theta = \mu$ then for every $\alpha < \theta^+$ there is a separating function from $G(\alpha, 2^\mu)$ into μ and hence $\chi(G(\alpha, 2^\mu)) \leq \mu$. A simple but quotable result of this note is that if \aleph_ω is strong limit, then $\chi(G(\omega, \{\aleph_\omega\})) > \aleph_\omega$. Hence Corollary 1.1.1 does not hold in case $\mu = \aleph_\omega$, $\theta = \omega$, and the cardinal assumption in that corollary is needed. This is the content of the following section.

1.2 $\chi(G(\omega, \{\beth_\omega\})) > \beth_\omega$ (and similar results)

It is convenient to define, for any set of ordinals B , a function $\pi_B : B \rightarrow \text{order-type}(B)$ by

$$\pi_B(a) = \text{order-type}(a \cap B).$$

If $A \subseteq B$, then $\pi_B \upharpoonright A$ is denoted $\pi_{A,B}$. That is, $\pi_{A,B}$ gives the position of a within B for every $a \in A$. So for arbitrary sets of ordinals X and Y , X and Y are consistent if and only if $\pi_{X \cap Y, X} = \pi_{X \cap Y, Y}$.

The following lemma is obvious.

Lemma 1.2.1. *Suppose A and B are α sets of ordinals and $X_0 \subseteq A \cap B$. Then $\pi_{X_0, A} = \pi_{X_0, B}$ if and only if the (unique) order isomorphism $g : A \rightarrow B$ is the identity on X_0 .*

Our first result is

Theorem 1.2.2. $\chi(G(\omega, \{\beth_\omega\})) > \beth_\omega$. *More generally, if λ is a strong limit singular cardinal and $\text{cf}(\lambda) = \kappa$, then $\chi(G(\kappa, \{\lambda\})) > \lambda$.*

Proof. For a simpler exposition we present the proof for the particular case of \beth_ω , but the reader will have no problems in making the obvious changes. Recall that $\beth_0 = \aleph_0$, $\beth_{n+1} = 2^{\beth_n}$, and \beth_ω is the limit of the \beth_n

sequence. Recall also that the graph $G(\omega, \{\beth_\omega\})$ consists of all ω -sets that are unbounded in \beth_ω . Suppose τ is a separating function from $\text{typ}(\omega, \{\beth_\omega\})$ into \beth_ω , and we shall reach a contradiction. So we assume that for every two ω -sets $M_1 \neq M_2$ unbounded in \beth_ω , if $\tau(M_1) = \tau(M_2)$ then M_1 and M_2 are consistent.

Given any $M \in \text{typ}(\omega, \{\beth_\omega\})$ we define the trace of M , denoted t_M , as the following (partial) function on \beth_ω . For every $\alpha < \beth_\omega$, pick (if there is one) some $N \in \text{typ}(\omega, \{\beth_\omega\})$ such that $M \subseteq N$ and $\tau(N) = \alpha$. Then define

$$t_M(\alpha) = \text{range } \pi_{M,N} = \{x \cap N \mid x \in M\}. \quad (1.2.1)$$

In words, $t_M(\alpha)$ is the set of positions occupied by M in N (it is a subset of ω). Notice that $t_M(\alpha)$ does not depend on N : if N' is some other member of $\text{typ}(\omega, \{\beth_\omega\})$ with $M \subseteq N'$ and $\tau(N') = \alpha$, then M occupies the same positions in N as in N' . In fact, the isomorphism between N and N' is the identity on $N \cap N'$ and hence on M .

Now, for every $n < \omega$, the set of all functions from \beth_n to $\mathcal{P}(\omega)$ has cardinality $2^{\beth_{n+1}}$, but the cardinality of the set of ω sequences unbounded in \beth_ω is 2^{\beth_ω} . Hence there are two distinct sets M_n and M'_n in $\text{typ}(\omega, \{\beth_\omega\})$ that begin after \beth_n and are such that

$$t_{M_n} \upharpoonright \beth_n = t_{M'_n} \upharpoonright \beth_n. \quad (1.2.2)$$

Let $K = \bigcup(\{M_n \mid n \in \omega\} \cup \{M'_n \mid n \in \omega\})$. Then K has order-type ω (its intersection with any \beth_n is finite) so that $K \in \text{typ}(\omega, \{\beth_\omega\})$, and $\tau(K) = \alpha$ is defined. Pick $n \in \omega$ with $\alpha < \beth_n$. Then $\alpha \in \text{dom}(t_{M_n}) \cap \text{dom}(t_{M'_n})$, and $t_{M_n}(\alpha) = t_{M'_n}(\alpha)$ by the choice of M_n and M'_n in (1.2.2). Hence M_n occupies in K the same positions as M'_n does, which is impossible since $M_n \neq M'_n$. \square

The following definition motivates much of the research reported in this paper.

Definition 1.2.3. Let κ be a (finite or infinite) cardinal; we say that a graph G has κ -chromatic number μ if and only if μ is the least cardinal such that there is a function τ from the vertices into μ such that $\tau^{-1}\{\alpha\}$ does not contain a clique of cardinality κ . That is, if $\{a_i \mid i \in \kappa\}$ is a set of size κ of vertices with $\tau(a_i) = \tau(a_j)$ for all i and j , then there are $a_i \neq a_j$ in this collection that are not edge connected in the graph. We say that such a function τ is a κ -separating coloring of the graph. We denote with $\chi_\kappa(G)$ the κ chromatic number of G .

For example, $\chi_2(G)$ is the chromatic number of G . $\chi_3(G) = 1$ is the statement that G is triangle free, and $\chi_3(G)$ is the least cardinality of a partition of G into triangle free subsets. So, $\chi_\kappa(G) > \mu$ is equivalent to the statement that any function F from the set of vertices of G to μ has some $\gamma \in \mu$ such that $F^{-1}\{\gamma\}$

contains a clique of cardinality κ .

Clearly, $\chi_2(G) \geq \chi_3(G) \geq \dots \geq \chi_{\aleph_0}(G) \geq \dots$.

A well-known question of Erdős and Hajnal [24] can be expressed in these terms as follows: is there a graph G with $\chi_4(G) = 1$ and $\chi_3(G) > \aleph_0$? (That is, does there exist a graph with no subgraph isomorphic to K_4 which cannot be expressed as a union of \aleph_0 triangle free graphs?)

In an email, A. Hajnal noted that a result in [24] is (in our terminology) that for every regular κ and $2 \leq n < \omega$ there is a graph G such that $\chi_n(G) = \kappa$ but $\chi_{n+1}(G) = 1$. This result was used by our referee to answer a question that we had in a previous draft and to construct, for every n , a graph G such that $\chi_2(G) > \dots > \chi_n(G) > \omega$. The construction of graphs G_i and uncountable cardinals κ_i , for $i = n, n-1, \dots, 1$ is done backwards and so that $\chi_i(G_i) = \kappa_i$ holds. First $\kappa_n = \aleph_1$ (for example) and G_n is chosen so that $\chi_n(G_n) = \kappa_n$ but $\chi_{n+1}(G_n) = 1$. If G_{i+1} and κ_{i+1} are defined, then $\kappa_i > |G_{i+1}|$ is chosen and G_i is defined so that $\chi_i(G_i) = \kappa_i$ and $\chi_{i+1}(G_i) = 1$. Then G is defined as the vertex disjoint union of the G_i 's.

An obvious application of the Erdős-Rado theorem is the following.

Theorem 1.2.4. *For every cardinal λ , for the graph $G = (\omega, \lambda)$, $\chi_{(2^{\aleph_0})^+}(G) = 1$.*

Proof. Suppose that $A \subseteq G$ is a clique of cardinality $(2^{\aleph_0})^+$. Define for $X \neq Y$ in A which are inconsistent $f(X, Y) = \langle n, m \rangle$ if the n th member of X is equal to the m th member of Y and $n \neq m$. As there is no homogenous triple, a contradiction to the Erdős-Rado theorem is obtained. \square

Our aim now is to prove the following.

Theorem 1.2.5. $\chi_{\aleph_0}(G(\omega, \{\beth_\omega\})) > \beth_\omega$.

To prove the theorem, we shall define first a graph $G^*(\omega, \{\beth_\omega\})$ on the set of vertices $\text{typ}(\omega, \{\beth_\omega\})$ but with fewer edges than $G(\omega, \{\beth_\omega\})$. We let (a, b) form an edge in G^* iff there are infinitely many $x \in a \cap b$ such that the order-type of $x \cap a$ is different from that of $x \cap b$. In case no G^* edge connects a and b we say that a and b are “eventually consistent”. So, $a, b \in \text{typ}(\omega, \{\beth_\omega\})$ are eventually consistent if and only if the isomorphism $f : a \rightarrow b$ is the identity on “almost all” members of $a \cap b$.

Define the “almost inclusion” relation $X \subseteq^* Y$ if and only if $Y \setminus X$ is finite, and then define the “almost equal” relation $X =^* Y$ if and only if $X \subseteq^* Y$ and $Y \subseteq^* X$. If $X \subseteq \omega$, then $[X]^*$ denotes the equivalence class of X . That is, the collection of all subsets of ω that are $=^*$ equivalent to X . In case f and g are functions, $f =^* g$ if and only if the domain of f is almost equal to the domain of g and $f(x) = g(x)$ for almost all x 's in intersection of the domains of f and g .

In these notations, $a, b \in \text{typ}(\omega, \{\beth_\omega\})$ are eventually consistent if and only if $\pi_{a \cap b, a} =^* \pi_{a \cap b, b}$.

We first note the following.

Theorem 1.2.6. *The chromatic number of $G^* = G^*(\omega, \{\beth_\omega\})$ is bigger than \beth_ω .*

The proof is similar to that of Theorem 1.2.2. Assume that $\tau : G^* \rightarrow \beth_\omega$ is separating in the sense that $\tau(M_1) = \tau(M_2)$ implies that M_1 and M_2 are eventually consistent. Then $t_M(\alpha)$ is defined, when $M \subseteq N$ for some N with $\tau(N) = \alpha$, as $[\text{range } \pi_{M,N}]_*$, the $=^*$ equivalence class of the set in (1.2.1). Then it follows again that $t_M(\alpha)$ does not depend on the set N chosen: any two such supersets will give equivalent sets of positions. At stage n choose two sets M_n and M'_n that are disjoint and such that $t_{M_n} \upharpoonright \beth_n = t_{M'_n} \upharpoonright \beth_n$. The contradiction is obtained as before.

Theorem 1.2.7. *For $G^* = G^*(\omega, \{\beth_\omega\})$, we have $\chi_{\aleph_0}(G^*) > \beth_\omega$: the \aleph_0 -chromatic number of G^* is bigger than \beth_ω .*

Proof. Suppose on the contrary that $\tau : \text{typ}(\omega, \{\beth_\omega\}) \rightarrow \beth_\omega$ is a \aleph_0 -separating coloring of the graph $G^*(\omega, \{\beth_\omega\})$. Given $M \in \text{typ}(\omega, \{\beth_\omega\})$ we define t_M on \beth_ω as follows. For any $\gamma \in \beth_\omega$ define

$$t_M(\gamma) = \{ [\text{range } \pi_{M \cap B, B}]_* \mid \tau(B) = \gamma \text{ and } M \subseteq^* B \}$$

In words, $t_M(\gamma)$ is the collection of the almost equality equivalence classes of subsets of ω induced by sets of the form $\{ |m \cap B| \mid m \in M \cap B \}$ where $B \in \text{typ}(\omega, \{\beth_\omega\})$ is such that $\tau(B) = \gamma$ and $M \setminus B$ is finite.

We claim that $t_M(\gamma)$ is a finite set (of equivalence classes); this is the content of the following lemma.

Lemma 1.2.8. *For every $M \in \text{typ}(\omega, \{\beth_\omega\})$ and $\gamma \in \beth_\omega$,*

$$\{ [\text{range } \pi_{M \cap B, B}]_* \mid \tau(B) = \gamma \text{ and } M \subseteq^* B \}$$

is finite.

Proof. If not, then there are $B_i \in \text{typ}(\omega, \{\beth_\omega\})$ for $i \in \omega$ such that $\tau(B_i) = \gamma$, $M \subseteq^* B_i$ and

$$\text{range } \pi_{M \cap B_i, B_i} \neq^* \text{range } \pi_{M \cap B_j, B_j} \tag{1.2.3}$$

for all $i \neq j$ (here \neq^* is the negation of $=^*$). We claim that $\{B_i \mid i \in \omega\}$ is a clique, which contradicts the assumption that τ is \aleph_0 -separating. To prove that (B_i, B_j) is an edge in G^* for $i \neq j$, we must find an infinite number of $m \in B_i \cap B_j$ for which $|m \cap B_i| \neq |m \cap B_j|$. But since $M \subseteq^* B_i$ and $M \subseteq^* B_j$, this follows immediately from (1.2.3). \square

-lemma

Thus (continuing the proof of the theorem) t_M takes values essentially in $[\mathcal{P}(\omega)]^{<\omega}$, and hence for every $n \in \omega$ there are not more than 2^{\beth_n} possible functions of the form $t_M \upharpoonright \beth_n$. It follows, for every fixed $n \in \omega$, that we can find M_n^i for $i \in \omega$ such that $M_n^i \neq^* M_n^j$ and $t_{M_n^i} \upharpoonright \beth_n = t_{M_n^j} \upharpoonright \beth_n$ for all i and j . Now find $K \in \text{typ}(\omega, \{\beth_\omega\})$ such that $M_n^i \subseteq^* K$ for all indices, and consider $\gamma = \tau(K)$. Pick some n such that $\gamma < \beth_n$. Since $M_n^i \neq^* M_n^j$ and these sets are almost included in K , $M_n^i \cap K \neq^* M_n^j \cap K$, and so

$$\text{range } \pi_{M_n^i \cap K, K} \neq^* \text{range } \pi_{M_n^j \cap K, K} \text{ for } i \neq j.$$

On one hand we have $[\text{range } \pi_{M_n^i \cap K, K}]_* \in t_{M_n^i}(\gamma)$ by the definition of $t_{M_n^i}(\gamma)$, but on the other hand there is a fixed F such that $t_{M_n^i}(\gamma) = F$ for all $i \in \omega$ (by definition of $\{M_n^i \mid i \in \omega\}$). Hence F is infinite, in contradiction to the lemma. \square

Since any edge of $G^* = G^*(\omega, \{\beth_\omega\})$ is also an edge of $G = G(\omega, \{\beth_\omega\})$, we have that $\chi_{\aleph_0}(G) \geq \chi_{\aleph_0}(G^*)$. That is, Theorem 1.2.5 is proven.

The following remain unresolved.

1. Improve the theorems by finding the exact value of the chromatic number, rather than just saying it is above \beth_ω . For example, is it always 2^{\beth_ω} ?
2. Can we replace \aleph_0 with \aleph_1 in Theorem 1.2.5? I. e, what is the \aleph_1 chromatic number of the graph? Observe that there are no cliques of size $(2^{\aleph_0})^+$ (by Erdős Rado).

1.3 Ladder graphs

The graph $G(\omega, \{\beth_\omega\})$ considered in the previous section has all its sets with the same supremum, namely \beth_ω . Now we consider the other extreme, when all sets have different suprema. These are the ladder graphs.

Let λ be some ordinal and suppose that a “ladder system” $\bar{X} = \langle X_\alpha \mid \alpha \in S_\omega^\lambda \rangle$ is given where $S_\omega^\lambda \subset \lambda$ is the subset of λ of limit ordinals with countable cofinality, and $X_\alpha \subset \alpha$ is unbounded in α and of order-type ω . A ladder graph induced by \bar{X} is a subgraph of $G(\omega, \lambda)$ having the X_α ’s as vertices, and edges all inconsistent pairs. It is easy to have such a graph with no edges at all: just assume that each X_α has the form $\{x_i \mid i \in \omega\}$ an increasing enumeration where each x_i is the i th successor of some limit ordinal. So the question is about constructing such a graph with large chromatic number. We concentrate on Ladder subgraphs of $G(\omega, \omega_1)$ and prove that assuming \diamond_{ω_1} there are such graphs of chromatic number \aleph_1 , but under MA_{\aleph_1} each such graph has countable chromatic number.

Our referee noticed that if $S \subset \omega_1$ is non-stationary, then the ladder graph built on $\langle X_\alpha \mid \alpha \in S \rangle$ has countable chromatic number. To see this, take C club disjoint to S and such that every $\alpha \in C$ is sufficiently closed. Then define the coloring on the interval (α, α') for $\alpha \in C$ by induction on α (where $\alpha' > \alpha$ is the next ordinal in C) so that vertices in (α, α') have different colors. The inductive requirement for $\alpha \in C$ ($\alpha > 0$) is that for every $\beta > \alpha$ there is an infinite number of differently colored $\beta' < \alpha$ with $X_\beta \cap \alpha$ an initial segment of $X_{\beta'}$. Now when a color has to be chosen for X_γ where $\gamma \in (\alpha, \alpha')$ while finitely many colors are to be avoided, an example is taken from some already defined $X_{\beta'}$ that extends $X_\gamma \cap \alpha$.

This situation is reminiscent of the one of the Hajnal–Máté graphs defined on ω_1 (in [29]). These graphs are also defined by means of a ladder system $\langle X_\alpha \mid \alpha \in S_\omega^\lambda \rangle$, by joining $\alpha < \beta$ with an edge if $\alpha \in X_\beta$. It is proven in [29] that the diamond \diamond_{ω_1} implies that there is a Hajnal–Máté graph of chromatic number \aleph_1 , while MA_{\aleph_1} implies that all such graphs have chromatic number $\leq \aleph_0$. Yet, the situation with respect to the continuum hypothesis is clearer with the Hajnal–Máté graphs (see [1] and [2]): we know that it is consistent that CH holds and all of these graphs have countable chromatic number, but we do not know the impact of CH on the ladder graphs defined here.

We first note the following.

Theorem 1.3.1. *If G is any ladder subgraph of $G(\omega, \lambda)$ induced by $\overline{X} = \langle X_\alpha \mid \alpha \in S_\omega^\lambda \rangle$, then $\chi_{\aleph_1}(G) = 1$. That is, there are no uncountable cliques in G .*

Proof. Given $S \subseteq \lambda_\omega^\lambda$ a set of cardinality \aleph_1 , we shall find $\alpha_1 < \alpha_2$ in S such that X_{α_1} and X_{α_2} are consistent (and hence not connected in the graph). Take M a countable elementary substructure of some H_κ rich enough to contain G and S . Let A be the closure of $M \cap \lambda$ in λ . That is, the set of all ordinals that are in $M \cap \lambda$ or are limits of ordinals in $M \cap \lambda$. Then A is countable and we can pick $\alpha_1 \in S \setminus A$. Since $\alpha_1 \notin A$, $F = X_{\alpha_1} \cap M$ is finite. For every $x \in F$, write $n(x) = |x \cap X_{\alpha_1}|$. Since M is an elementary substructure, there is $\alpha_2 \in M \cap S$ such that for every $x \in F$, we have $|x \cap X_{\alpha_2}| = n(x)$. Now $X_{\alpha_1} \cap X_{\alpha_2} \subseteq F$ and $X_{\alpha_1}, X_{\alpha_2}$ are consistent. \square

Theorem 1.3.2. *Assume \diamond_{ω_1} . There is a ladder graph G on $S_\omega^{\omega_1}$ such that $\chi_{\aleph_0}(G) = \aleph_1$.*

Proof. Let $\langle S_\alpha \mid \alpha \in \omega_1 \rangle$ be the assumed diamond sequence. We define a ladder system $\langle X_\alpha \mid \alpha \in S_\omega^{\omega_1} \rangle$, where X_α is defined by induction on limit $\alpha \in \omega_1$, an ω set cofinal in α . At stage α consider S_α and suppose that it encodes a function $f_\alpha : \alpha \rightarrow \omega$. Let $\langle \gamma_i \mid i \in \omega \rangle$ be an ω sequence increasing and cofinal in α . We define $x_i^\alpha \in \alpha$ for $i \in \omega$ by induction with the aim of defining $X_\alpha = \{x_i^\alpha \mid i \in \omega\}$. At stage i of the construction we have defined $k(i) \in \omega$ and the first $k(i)$ members of the sequence, denoted $\langle x_j^\alpha \mid j < k(i) \rangle$. We will choose as follows a finite increasing sequence of the form $x_{k(i)}^\alpha, \dots, x_{k(i+1)-1}^\alpha$ in $\alpha \setminus (\gamma_i \cup \max\{x_0^\alpha, \dots, x_{k(i)-1}^\alpha\}) + 1$

(to ensure that the resulting sequence is increasing and cofinal in α). Let B_i be the collection of all limit $\alpha' \in \alpha \setminus (\gamma_i \cup \max\{x_0^\alpha, \dots, x_{k(i)-1}^\alpha\}) + 1$ such that $x_0^\alpha, \dots, x_{k(i)-1}^\alpha$ is an initial segment of $X_{\alpha'}$ and $f_\alpha(\alpha') = i$. Suppose that $\{X_{\alpha'} \mid \alpha' \in B_i\}$, being a subgraph of the graph constructed so far, contains a finite maximal clique. In this case let $\alpha_0, \dots, \alpha_{k-1} \in B_i$ be the set of indices of a maximal clique enumerated in increasing order. In fact, we take $\{\alpha_0, \dots, \alpha_{k-1}\}$ to be minimal in some well ordering of the finite sets of ordinals. Define $x_{k(i)}^\alpha, \dots, x_{k(i)+k-1}^\alpha$ so that X_α (no matter how it is going to be completed) and X_{α_j} are inconsistent for every $j < k$.

We must prove that the resulting graph has \aleph_0 -chromatic number \aleph_1 . Suppose $f : \omega_1 \rightarrow \omega$ is a coloring (that is, the function taking X_α to $f(\alpha)$ defined on the vertices of the graph is the coloring). We have to prove that for some $i \in \omega$, $f^{-1}\{i\}$ contains an infinite clique. Suppose on the contrary that for every $i \in \omega$ all cliques of $f^{-1}\{i\}$ are finite (actually, the contradiction is derived from the assumption that every $f^{-1}\{i\}$ contains a maximal finite clique).

Let $\langle M_\alpha \mid \alpha \in \omega_1 \rangle$ be an increasing and continuous sequence of countable elementary substructures of H_κ which is large enough to contain the graph and the function f . Find $\alpha \in \omega_1$ so that $f \upharpoonright \alpha$ is encoded by S_α and $\alpha = M_\alpha \cap \omega_1$. Suppose $f(\alpha) = i_0$ and consider stage i_0 in the definition of the sequence x_i^α . The definition of B_{i_0} can be done in M_α and it contains a finite maximal clique. So the maximal clique, subset of B_{i_0} used at stage i_0 is in M_α , and by the construction it turns out that it is not maximal since X_α is inconsistent with each X_{α_i} . This contradiction proves the theorem. \square

Assuming MA_{ω_1} , any $G = \langle X_\alpha \mid \alpha \in \lim \omega_1 \rangle$, a ladder subgraph of $G(\omega, \omega_1)$, has countable chromatic number.

Theorem 1.3.3. *Assume Martin's Axiom and $2^{\aleph_0} > \kappa$. If $G = \langle X_\alpha \mid \alpha \in S_\omega^\kappa \rangle$ is a ladder subgraph of $G(\omega, \kappa)$, then $\chi(G) \leq \omega$.*

Proof. Consider the poset P of all finite approximation to a separating function. That is, $p \in P$ if and only if $\text{dom}(p) \subset S_\omega^\kappa$, $p : \text{dom}(p) \rightarrow \omega$, and for every $\alpha, \beta \in \text{dom}(p)$, if $p(\alpha) = p(\beta)$ then X_α and X_β are consistent. The ordering of P is plain extension. Clearly, any condition can be extended to include any given limit ordinal in its domain, since the range of p is finite. The countable chain condition of P is proved below and so Martin's Axiom applies to yield that the chromatic number of G is countable.

It remains to prove that P satisfies the c.c.c. Let $P_0 \subseteq P$ be uncountable. We may assume that the sets $\{\text{dom}(p) \mid p \in P_0\}$ form a Δ -system with core D_0 . We may even assume that the sets $\text{dom}(p)$ for $p \in P_0$ are pairwise disjoint and $D_0 = \emptyset$ (just replace p with $p \upharpoonright (\text{dom}(p) \setminus D_0)$). Pick a countable $M \prec H_\kappa$ with $G, P, P_0 \in M$. Let \overline{M} be the union of $M \cap \lambda$ with its set of accumulation points in λ . As M is countable \overline{M} is

countable. As P_0 is uncountable, we can find $p \in P_0$ such that $\text{dom}(p) \cap \overline{M} = \emptyset$. Thus for every $x \in \text{dom}(p)$, $C_x \cap M$ is finite. We think of p as a structure with universe $\omega \cup \text{dom}(p) \cup \bigcup_{x \in \text{dom}(p)} C_x$, with predicates the ordering relation and the two binary relations $a \in C_x$ and $\alpha \in \text{dom}(p)$, and with constants all members of $C_x \cap M$ for $x \in \text{dom}(p)$. The function p itself is also part of that structure. Since M is an elementary substructure there is $q \in P_0 \cap M$ such that the structures of p and q are isomorphic with an isomorphism f that does not move the constants.

We claim that p and q are compatible. Suppose not, and $\alpha \in \text{dom}(p)$ and $\beta \in \text{dom}(q)$ are such that $p(\alpha) = q(\beta)$ but X_α and X_β are inconsistent. Recall that $\alpha \notin M$. Say $\alpha' = f(\alpha)$. Then $\alpha' \in \text{dom}(q)$ and $C_\alpha \cap M = C_\alpha \cap C_{\alpha'}$. Moreover, each $t \in C_\alpha \cap C_{\alpha'}$ has the same position in C_α as it has in $C_{\alpha'}$. It also follows that $q(\alpha') = q(\beta)$. Supposedly there is $x \in C_\alpha \cap C_\beta$ that has different positions in C_α and C_β . But then $x \in C_\alpha \cap M$ and so $x \in C_{\alpha'}$ and has the same position there as it has in C_α . Which is impossible since α' and β are both in $\text{dom}(q)$. \square

In view of the last three theorems, we ask: is there (in ZFC) a graph G with $\chi_{\aleph_0}(G) = \aleph_1$ and $\chi_{\aleph_1}(G) = 1$?

Moving one cardinal higher we look at ladder subgraphs of $G(\omega, \omega_2)$. By the previous theorem, under $\text{MA} + 2^{\aleph_0} > \aleph_2$ they all have countable chromatic number. If CH holds then they have \aleph_1 chromatic number (by Corollary 1.1.2, take $\mu = \theta = \aleph_1$). We prove next that for the case that $2^{\aleph_0} = \aleph_2$ it is consistent to have a ladder subgraph of $G(\omega, \omega_2)$ with chromatic number \aleph_2 : just add Cohen reals.

Theorem 1.3.4. *In a model obtained by adding \aleph_2 many Cohen reals there is a ladder subgraph of $G(\omega, \omega_2)$ with chromatic number \aleph_2 .*

Proof. Pick for any $\alpha \in S_{\omega^2}^{\omega}$ an unbounded ω set with an increasing enumeration $C_\alpha = \{C_\alpha(n) \mid n \in \omega\}$. Suppose G is a V -generic filter over the Cohen forcing poset (of finite functions from ω_2 to 2). Let $g = \bigcup G$ be the resulting generic function from ω_2 to 2, and denote for any limit ordinal $\alpha \in \omega_2$ $g_\alpha = g \upharpoonright [\alpha, \alpha + \omega)$. In $V[G]$ define $D_\alpha \subset C_\alpha$ as the subset of C_α obtained by picking only those member of the C_α sequence in positions that are in g_α . That is, $C_\alpha(n) \in D_\alpha$ if and only if $g_\alpha(\alpha + n) = 1$.

We claim that the resulting graph has chromatic number \aleph_2 . Suppose for a contradiction that \tilde{f} is a name forced by every condition to be a function from ω_2 in ω_1 . We shall find a condition (extending a given condition) that forces two vertices to be connected and have the same color under \tilde{f} .

So let r_0 be an arbitrary condition. It is a finite function from a subset of ω_2 to 2. Let M be an elementary substructure of some large enough H_κ and with cardinality \aleph_1 such that $r_0, \tilde{f} \in M$ and $\delta = M \cap \omega_2 > \omega_1$ is of countable cofinality. Let r_1 be an extension of r_0 that forces that $\tilde{f}(\delta) = \xi$ for some $\xi \in \omega_1$. Let n be the cardinality of $\text{dom}(r_1) \cap [\delta, \delta + \omega)$. So r_1 determines which of the first members of C_δ are in D_δ . Let

x_0, \dots, x_{n+1} be the first $n + 2$ members of C_δ . It will be soon evident why we want those two additional members of the sequence, x_n and x_{n+1} .

Let $s = r_1 \upharpoonright M$. Then $s \in M$ and s is also an extension of r_0 . Since M is an elementary substructure, there is in M a condition s_1 that extends s and “reflects” r_1 . That is, there is an isomorphism $i : \text{dom}(r_1) \longrightarrow \text{dom}(s_1)$ such that for $\delta' = i(\delta)$ we have that:

1. x_0, \dots, x_{n+1} are also the first $n + 2$ members of $C_{\delta'}$.
2. s_1 forces that $\check{f}(\delta') = \xi$.

Now extend s_1 to force that $D_{\delta'}$ includes both x_n and x_{n+1} , and extend r_1 to force that D_δ includes x_{n+1} but not x_n . Then these two extensions are compatible in the Cohen poset and they force that $D_{\delta'}$ and D_δ are inconsistent. \square

Suppose the GCH. What is the chromatic number of ladder subgraphs of $G(\omega, \omega_3)$? Certainly $\leq \aleph_2$ (by Corollary 1.1.2). Can we define a ladder graph (in L? with forcing?) so that its chromatic number is \aleph_2 ?

What are the chromatic numbers of ladder subgraphs of $G(\omega, \aleph_{\omega+1})$?

1.4 Graphs of the form $G(\omega, \mu)$

In the previous sections we considered graphs of ω sequences that had all the same supremum or all different suprema. Now we consider graphs of the form $G(\omega, \mu)$ where μ is a cardinal. That is, graphs of all ω sequences in μ with no restriction on their suprema.

One can consider the more general case $G(\alpha, \mu)$ of all subsets of μ of order-type α (edges defined as subgraphs of $\mathcal{P}(\mu)$). The case $G(2, \mu)$ was considered in [25], and here we extend this discussion. They proved that $G(2, (2^\kappa)^+)$ (called there a shift graph) is a triangle free graph with chromatic number $\geq \kappa^+$ such that all its subgraphs of cardinality $\leq 2^\kappa$ have chromatic number $\leq \kappa$. Our example below is different: not only the chromatic number of the graph is greater than κ , but its \aleph_0 -chromatic number is also above κ .

Theorem 1.4.1. *For $G = G(\omega, (2^\kappa)^+)$, $\chi_{\aleph_0}(G) > \kappa$. If $\kappa^{\aleph_0} = \kappa$ then any subgraph of G of smaller cardinality has chromatic number $\leq 2^\kappa$.*

Proof. In the following, G denotes the set of vertices of the graph (all ω -subsets of $(2^\kappa)^+$). Suppose that $\chi_{\aleph_0}(G) \leq \kappa$ and $\tau : G \longrightarrow \kappa$ is an \aleph_0 -separating function. That is, $\tau(X_i) = \tau(X_j)$ for all $i, j \in \omega$ implies that for some $i \neq j$ X_i and X_j are consistent sequences.

For every $\alpha \in (2^\kappa)^+$ define a function $g_\alpha : \kappa \longrightarrow \omega$ as follows. Given $\xi \in \kappa$, define $g_\alpha(\xi) = \max \{ |X \cap \alpha| \mid \alpha \in X \text{ and } \tau(X) = \xi \}$. We claim that $g_\alpha(\xi) \in \omega$. Otherwise there are $X_i \in G$ (for $i \in \omega$) such that

$|X_i \cap \alpha| \neq |X_j \cap \alpha|$ for $i \neq j$ and yet $\tau(X_i) = \xi$ for all i . But this contradicts the property of τ since any two X_i 's are inconsistent.

Since each g_α can be encoded as a subset of κ , there is a set $A \subset (2^\kappa)^+$ of cardinality $(2^\kappa)^+$ and such that $g_\alpha = g_\beta$ for every $\alpha, \beta \in A$. Let $X = \{x_i \mid i \in \omega\}$ be an increasing ω enumeration of ordinals from A . Say $\xi = \tau(X)$. We claim that $m < g_{x_0}(\xi)$ for every $m \in \omega$, and this is a contradiction. Clearly $m \leq g_{x_m}(\xi)$ since $X \cap x_m = \{x_0, \dots, x_{m-1}\}$. But $g_{x_0} = g_{x_m}$ and hence $m \leq g_{x_0}(\xi)$.

The second statement of the theorem is that if G_0 is a subgraph of $G(\omega, (2^\kappa)^+)$ generated by $\leq 2^\kappa$ vertices, then the chromatic number of G_0 is $\leq \kappa$. This follows if we prove for every $\lambda < (2^\kappa)^+$ that the chromatic number of $G(\omega, \lambda)$ is $\leq \kappa$. We use here Corollary 1.1.2 to the Engelking and Karłowicz theorem with $\mu = \kappa$ and $\theta = \aleph_0$. \square

Theorem 1.4.2. *Assume that $\kappa^{\aleph_0} = \kappa$. Then $\chi_{\aleph_1}(G(\omega, (2^\kappa)^+)) \leq \kappa$.*

Proof. Fix for every $\beta < (2^\kappa)^+$ a function $\tau_\beta : \text{typ}(\omega, \{\beta\}) \rightarrow \kappa$ such that if $\tau(M_1) = \tau(M_2)$ then M_1 and M_2 are consistent. This is possible by Corollary 1.1.2 since β has cardinality $\leq 2^\kappa$. Now we define the \aleph_1 -separating function $\tau : \text{typ}(\omega, (2^\kappa)^+) \rightarrow \kappa$ as follows. For any ω -set $X \subset (2^\kappa)^+$, let $\beta = \sup X$, and define $\tau(X) = \tau_\beta(X)$. We prove that τ is \aleph_1 -separating. Suppose that $\{X_i \mid i \in \omega_1\}$ is a collection of \aleph_1 vertices and that for some fixed $\alpha \in \kappa$ we have $\tau(X_i) = \alpha$ for all i . We must prove that this collection is not a clique. Denote $\beta_i = \sup X_i$ for all i . In case, for some $i \neq j$, we have $\beta_i = \beta_j = \beta$, then X_i and X_j are consistent by the property of τ_β . Otherwise, $\{X_i \mid i \in \omega_1\}$ forms a ladder system and is hence not a clique (by Theorem 1.3.1). \square

For example, for $\kappa = 2^{\aleph_0}$ and $G = G(\omega, (2^{2^{\aleph_0}})^+)$ we get by the last two theorems that $\chi_{\aleph_0}(G) > \kappa \geq \chi_{\aleph_1}(G)$. When $\kappa = 2^{\aleph_0}$ is regular, $\chi_{\aleph_1}(G) = \kappa$ because $G(\omega, \omega)$ has a clique of size 2^{\aleph_0} .

Chapter 2

Combinatorics in model theory: Quantifier elimination tests

Abstract: We prove that, for countable languages, two model-theoretic quantifier elimination tests, one proposed by J. R. Shoenfield and the other by L. van den Dries, are equivalent.¹

2.1 Introduction

To facilitate the discussion we first introduce the following terminological and notational conventions.

Definition 2.1.1. Let M be a model and $A \subseteq |M|$. Let N be the model $\langle M, a \rangle_{a \in A}$.

1. The theory $\text{Th}(N)$, denoted by $\text{CD}(A, M)$, is called the *complete diagram of A in M* . If $A = |M|$ we simply write $\text{CD}(M)$.
2. The set of all quantifier-free sentences in $\text{Th}(N)$, denoted by $\text{ED}(A, M)$, is called the *elementary diagram of A in M* . Again if $A = |M|$ we simply write $\text{ED}(M)$.

Obviously if $N \preceq M$ then $\text{CD}(N, M) = \text{CD}(N)$ and if $N \subseteq M$ then $\text{ED}(N, M) = \text{ED}(N)$.

We say that a theory T is *model complete* if and only if, for every pair of models $N, M \models T$, $N \subseteq M$ implies $N \preceq M$. Abraham Robinson showed that under certain conditions a model complete theory admits quantifier elimination (QE for short). This was one of the results that inaugurated the use of model-theoretic methods in the study of QE. Model-completeness has many equivalent formulations:

Fact 2.1.2. *Let T be any theory. The following are equivalent:*

1. T is model complete.
2. For any two models $N, M \models T$ with $N \subseteq M$ there is an $N^* \models T$ such that $N \preceq N^*$ and M can be embedded into N^* over N .
3. For any $M \models T$ the theory $T \cup \text{ED}(M)$ is complete.

¹This chapter was submitted for publication. The referee found an easier proof that avoided the complicated infinitary combinatorics in the inductive step.

4. For any two models $N, M \models T$ with $N \subseteq M$, every existential formula $\varphi(\bar{x})$, and every $\bar{b} \in |N|$, we have $M \models \varphi(\bar{b})$ if and only if $N \models \varphi(\bar{b})$.
5. For every existential formula $\varphi(\bar{x})$ there is a universal formula $\varphi^*(\bar{x})$ such that $T \vdash \varphi(\bar{x}) \leftrightarrow \varphi^*(\bar{x})$.
6. For every formula $\varphi(\bar{x})$ there is a universal formula $\varphi^*(\bar{x})$ such that $T \vdash \varphi(\bar{x}) \leftrightarrow \varphi^*(\bar{x})$.
7. For every formula $\varphi(\bar{x})$ there is a universal formula $\varphi_1(\bar{x})$ and an existential formula $\varphi_2(\bar{x})$ such that $T \vdash \varphi_1(\bar{x}) \leftrightarrow \varphi(\bar{x}) \leftrightarrow \varphi_2(\bar{x})$.

For a proof of this fact see [11] and [43].

However, there are theories which are model complete but do not admit QE. For example, the complete theory of real closed fields in the language of rings is model complete, but the formula $\exists x x \times x = y$ is not equivalent to any quantifier-free formula in this theory. See [11] for details.

Over the years many model-theoretic properties have been proposed to strengthen model-completeness so that QE is implied without any additional assumptions on the theory in question. Some of these properties are logically equivalent to QE; others are strictly stronger than QE. Below we shall prove that two of the stronger ones, one proposed by J. R. Shoenfield and the other by L. van den Dries, are equivalent for countable languages.

2.2 Some QE tests

Let T be any theory. Here are some model-theoretic QE tests that are stronger than model-completeness:

Definition 2.2.1. T is *submodel complete* if and only if for any model $M \models T$ and any $N \subseteq M$ the theory $T \cup \text{ED}(N)$ is complete.

This is a direct strengthening of 2.1.2.3.

Definition 2.2.2. T has the *submodel amalgamation property* (SA-property for short) if and only if for any $M_1, M_2 \models T$ and any $N \subseteq M_1, M_2$ there is an $M^* \models T$ such that $M_1 \preceq M^*$ and M_2 can be embedded into M^* over N via a monomorphism f ; that is, the following diagram

$$\begin{array}{ccc}
 M_1 & \xrightarrow{\preceq} & M^* \\
 \uparrow \subseteq & & \uparrow f \\
 N & \xrightarrow{\subseteq} & M_2
 \end{array}$$

commutes.

This is a direct strengthening of 2.1.2.2.

Definition 2.2.3. T has the *Shoenfield property* (S-property for short) if and only if for any two models $M_1, M_2 \models T$ such that M_2 is $\|M_1\|^+$ -saturated and any isomorphism $f : N_1 \rightarrow N_2$ with $N_1 \subseteq M_1$ and $N_2 \subseteq M_2$, there is a monomorphism $f^* : M_1 \rightarrow M_2$ extending f .

Definition 2.2.4. T has the *strong Shoenfield property* (SS-property for short) if and only if

1. For every two models $M_1, M_2 \models T$ and every two models $N_1 \subseteq M_1$ and $N_2 \subseteq M_2$, if $f : N_1 \rightarrow N_2$ is an isomorphism, then there is an isomorphism $f^* : N_1^* \rightarrow N_2^*$ which is an extension of f , where $N_1^* \subseteq M_1$, $N_2^* \subseteq M_2$, and $N_1^*, N_2^* \models T$;
2. For every two models $N, M \models T$ with $N \subseteq M$, every existential formula $\varphi(\bar{x})$, and every $\bar{b} \in |N|$, we have $M \models \varphi(\bar{b})$ if and only if $N \models \varphi(\bar{b})$. In other words, T is model complete.

When there is no danger of confusion we abuse $L(T)$ to denote both the language of T and the set of all well-formed formulas in the language of T . For two structures N and M in $L(T)$ we say that M is a *T-extension* of N if $|N| \subseteq |M|$ and $M \models T$.

Definition 2.2.5. T has the *van den Dries property* (D-property for short) if and only if

1. For any model N , if there exists a model $M \models T$ such that $N \subseteq M$, then there is a T -closure N^* of N , that is, a model $N^* \models T$ such that $N \subseteq N^*$ and N^* can be embedded over N into any T -extension of N ;
2. If $N, M \models T$ and $N \subsetneq M$, then there is an $a \in |M| \setminus |N|$ such that $N + a$ can be embedded into an elementary extension of N over N , where $N + a$ is the smallest submodel of M that contains $|N| \cup \{a\}$.


The SS-property first appeared in Shoenfield's textbook [47]. He subsequently modified it into the S-property and proved its equivalence to QE in [48]. The D-property was given by van den Dries in [18] and [19], which is a straightforward strengthening of the SS-property. However, the main result Theorem 2.2.7 below shows that, for countable languages, its main advantage over the SS-property is its conceptual concreteness rather than its logical strength.

Theorem 2.2.6. *Let T be a theory in a language with at least one constant symbol. For the following statements,*

1. T is submodel complete,
2. T has the SA-property,

3. T has the S -property,
4. T has the SS -property,
5. T has the D -property,
6. T admits QE ,

these logical implications hold:



Proof. That 1, 2, and 3 are equivalent to QE is well-known. See, for example, [43] and [48]. Here we give proofs to the remaining two implications. We also show directly how the first condition of the SS -property achieves QE on top of model-completeness. This proof is a modification of the standard proof of “2.1.2.4 \Rightarrow 2.1.2.5” in the literature, which establishes a crucial connection between model-theoretic properties and syntactical properties.

4 \Rightarrow 6: Let $\varphi(\bar{x})$ be a formula in $L(T)$. Since T is model complete, by 2.1.2, $\varphi(\bar{x})$ is equivalent to both a universal formula and an existential formula. Hence we may assume that $\varphi(\bar{x})$ is a universal formula. Let $\varphi^*(\bar{x})$ be an existential formula such that $T \vdash \varphi(\bar{x}) \leftrightarrow \varphi^*(\bar{x})$. Let \bar{c} be new constants. Let Γ be a set that contains exactly the following formulas:

- $T \cup \{\varphi(\bar{c})\}$, and
- every quantifier-free formula $\neg\psi(\bar{c})$ such that $T \vdash \forall \bar{x} (\psi(\bar{x}) \rightarrow \varphi(\bar{x}))$.

Suppose for contradiction that Γ is consistent. Take any model $M \models \Gamma$. Let $N \subseteq M$ be the minimal submodel generated by \bar{c} . Note that every element in N can be written as a term that only involves \bar{c} , the constants of $L(T)$, and the functions of $L(T)$. Now, if $T \cup \text{ED}(N)$ does not prove $\varphi(\bar{c})$, then fix a model $M^* \models T \cup \text{ED}(N) \cup \{\neg\varphi(\bar{c})\}$. By the first condition of the SS -property we can find an $N_1 \models T \cup \text{ED}(N)$ in M and an $N_2 \models T \cup \text{ED}(N)$ in M^* such that they are isomorphic over N . Since $\varphi(\bar{x})$ is a universal formula and $M \models \varphi(\bar{c})$, we have $N_1 \models \varphi(\bar{c})$. So $N_2 \models \varphi(\bar{c})$, so $N_2 \models \varphi^*(\bar{c})$, so $M^* \models \varphi^*(\bar{c})$, so $M^* \models \varphi(\bar{c})$, contradiction. So $T \cup \text{ED}(N) \vdash \varphi(\bar{c})$. So there is a quantifier-free formula $\psi(\bar{c}) \in \text{ED}(N)$ such that $T \cup \{\psi(\bar{c})\} \vdash \varphi(\bar{c})$, so $T \vdash \psi(\bar{c}) \rightarrow \varphi(\bar{c})$. But \bar{c} are new constants, so $T \vdash \forall \bar{x} (\psi(\bar{x}) \rightarrow \varphi(\bar{x}))$. So $\neg\psi(\bar{c}) \in \Gamma$, contradiction again.

So Γ is not consistent. This means that there are finitely many quantifier-free formulas $\psi_i(\bar{x})$ such that $T \vdash \forall \bar{x} (\psi_i(\bar{x}) \rightarrow \varphi(\bar{x}))$ for every i and $T \vdash \forall \bar{x} (\varphi(\bar{x}) \rightarrow \bigvee_i \psi_i(\bar{x}))$. So $T \vdash \forall \bar{x} (\varphi(\bar{x}) \leftrightarrow \bigvee_i \psi_i(\bar{x}))$, as desired.

4 \Rightarrow 3: Let $M_1, M_2 \models T$, $N \subseteq M_1, M_2$, and let M_2 be $\|M_1\|^+$ -saturated. By the first condition of the SS -property we can find two T -extensions N_1, N_2 of N in M_1, M_2 respectively that are isomorphic over N .

Let the isomorphism be f . Pick an $a \in |M_1| \setminus |N_1|$ and consider any quantifier-free formula $\varphi(x; \bar{b})$ with $\bar{b} \in |N_1|$ such that $M_1 \models \varphi(a; \bar{b})$. Since $M_1 \models \exists x \varphi(x; \bar{b})$, by the second condition of the SS-property we have $N_1 \models \exists x \varphi(x; \bar{b})$, so $N_2 \models \exists x \varphi(x; f(\bar{b}))$, so $M_2 \models \exists x \varphi(x; f(\bar{b}))$. Hence the quantifier-free type $f(p)$ is realized in M_2 , say, by d , where p is the set of all quantifier-free formulas in $\text{tp}(a/|N_1|, M_1)$. If we set $a \mapsto d$ then we get an induced isomorphism between $N_1 + a$ and $N_2 + d$. Iterating this procedure to exhaust all elements in M_1 we see that M_1 can be embedded into M_2 over N .

5 \Rightarrow 4: Trivially the closure property, that is, the first condition of the D-property, implies the first condition of the SS-property. For the second condition of the SS-property, let $N, M \models T$ with $N \subseteq M$. Consider an existential formula $\exists \bar{x} \varphi(\bar{x}; \bar{b})$ that is satisfied in M , where $\bar{b} \in |N|$ and $\varphi(\bar{x}; \bar{b})$ is quantifier-free. So let \bar{c} be such that $M \models \varphi(\bar{c}; \bar{b})$. We construct the following diagram:

$$\begin{array}{ccccccc}
N_0 & \xrightarrow{\subseteq} & N_0 + a_0 & \xrightarrow{\subseteq} & N_1 & \xrightarrow{\subseteq} & N_1 + a_1 & \xrightarrow{\subseteq} & N_2 & \cdots & \xrightarrow{\subseteq} & M \\
& \searrow \cong & \downarrow & \swarrow f_0 & \searrow \cong & \downarrow & \swarrow f_1 & & \searrow & & & \\
& & N_0^* & & & N_1^* & & & & \cdots & &
\end{array}$$

where $N_0 = N$, each N_{i+1} is the T -closure of $N_i + a_i$ promised by the closure property, each a_i and N_i^* are as described in the second condition of the D-property, all arrows are monomorphisms, and at the limit stage we simply take the union of all previous N_i 's.

Now, let i be the least index such that $\bar{c} \in N_i$. Note that i cannot be a limit ordinal. So $N_i \models \exists \bar{x} \varphi(\bar{x}; \bar{b})$, so $N_{i-1}^* \models \exists \bar{x} \varphi(\bar{x}; \bar{b})$, so $N_{i-1} \models \exists \bar{x} \varphi(\bar{x}; \bar{b})$, etc. If γ is a limit ordinal and $N_\gamma \models \exists \bar{x} \varphi(\bar{x}; \bar{b})$, then there is a $\bar{d} \in |N_\gamma|$ such that $N_\gamma \models \varphi(\bar{d}; \bar{b})$, so by the construction there is a $j < \gamma$ such that $\bar{d} \in |N_j|$, so $N_j \models \varphi(\bar{d}; \bar{b})$, so $N_j \models \exists \bar{x} \varphi(\bar{x}; \bar{b})$. As we trace back in the diagram we see that $N = N_0 \models \exists \bar{x} \varphi(\bar{x}; \bar{b})$. \square

The reason that we have assumed that the language of T has at least one constant symbol is to avoid certain pathology. That is, in the proof of “4 \Rightarrow 6” above, if φ is a sentence and $L(T)$ has no constant symbol, then \bar{c} is the empty sequence and cannot generate any submodel as we do not allow an empty model. The reader should observe that in this case the proof will not go through if we simply use an arbitrary submodel. In the sequel we shall always assume that T has a constant symbol whenever we are in a similar situation.

There are still more model-theoretic tests that are equivalent to QE. They are all more or less variations of the three equivalent tests in the above theorem. See [33] for more details about this. On the other hand, it is tempting to ask if in the above theorem all of the statements are indeed equivalent.

Jeremy Avigad has an example which shows that QE is strictly weaker than the SS-property. Consider the set 2^ω of all binary sequences of length ω . For each $n \in \omega$ let Z_n be a unary predicate such that if

$n = 0$ then $Z_n(\eta)$ for any $\eta \in 2^\omega$, otherwise $Z_n(\eta)$ if and only if $(\eta)_n = 0$. Let $T = \text{Th}(\langle 2^\omega, Z_n \rangle_{n \in \omega})$. Since except equality all predicates in the language are unary, every existential formula $\exists x \varphi(x; \bar{y})$ is equivalent to a formula of the form $\bigvee_i (\theta_i(\bar{y}) \wedge \exists x \phi_i(x; \bar{y}))$, where $\phi_i(x; \bar{y})$ is a conjunction of literals each of which contains x . If the unary predicates in the formula $\exists x \phi_i(x; \bar{y})$ describe a “consistent” finite sequence, then it can be translated into an equivalent quantifier-free formula that only involves \bar{y} . So T proves that every existential formula is equivalent to a quantifier-free formula, which means that T admits QE. Now, it is not hard to see that any dense subset of 2^ω is a model of T . Let $S_0 \subseteq 2^\omega$ be the set of those sequences that have only finitely many 0’s. Let $S_1 \subseteq 2^\omega$ be the set of those sequences that have only finitely many 1’s and the constant sequence $\bar{1}$. So both S_0 and S_1 are models of T . Notice that $\{\bar{1}\}$ is a submodel of both models as there is no function symbol in the language. Clearly there cannot be isomorphic T -extensions of $\{\bar{1}\}$ in S_0 and S_1 .

What about the SS-property and the D-property? First of all it is trivial that if a theory T admits QE then the second condition of the D-property holds, because, by 2.1.2, if $N, M \models T$ and $N \subseteq M$ then M itself is an elementary extension of N . The closure property, however, is much harder to achieve. The rest of this paper is devoted to proving

Theorem 2.2.7. *For countable languages the SS-property and the D-property are equivalent.*

The argument is by a transfinite induction.

2.3 The base case of the induction

We need more concepts and Henkin’s Omitting Type Theorem.

Definition 2.3.1. Let \bar{x} be a sequence of variables and p a T -type in \bar{x} . If there exists a formula $\varphi(\bar{x})$ such that $T \cup \{\varphi(\bar{x})\}$ is consistent and $\varphi(\bar{x}) \vdash p$, then we say that p is *isolated by $\varphi(\bar{x})$ via T* . If in context it is clear that which theory is being discussed then we omit T .

Note that if p is a complete T -type then p is isolated via T if and only if there exists a $\varphi \in p$ such that $\varphi \vdash p$.

Definition 2.3.2. Let $M \models T$ and $A \subseteq |M|$. We say that M is *almost T -primary over A* if there exists an ordinal α and a sequence $\langle (N_i, b_i) : i < \alpha \rangle$ such that

1. N_0 is the minimal submodel of M that contains A ,
2. $b_i \in |M| \setminus |N_i|$ and $N_{i+1} = N_i + b_i$ for each $i < \alpha$ (if $\alpha = \beta + 1$ then b_β is not defined),

3. $N_\beta = \bigcup_{i < \beta} N_i$ if β is a limit ordinal and $\bigcup_{i < \alpha} N_i = M$,
4. the type $\text{tp}(b_j / |N_j|, M)$ is isolated via T_j for every $j < \alpha$, where $T_j = T \cup \text{CD}(N_j, M)$.

The sequence $\langle (N_i, b_i) : i < \alpha \rangle$ is called an *almost isolating sequence for M over A* . The ordinal α is the *length* of the sequence.

For convenience, if $T = \text{Th}(M)$ then we omit T . Also, sometimes we allow an almost isolating sequence to have repeated consecutive b_i 's. Of course in this case we no longer require $b_i \notin |N_i|$ for the repeated occurrences. Note that this definition is a variation of the notion of a primary model, which plays an important role in the proof of Morley's Theorem.

Definition 2.3.3. Let $M \models T$ and $A \subseteq |M|$. We say that M is *T -primary over A* if there exists an ordinal α and an enumeration $\langle b_i : i < \alpha \rangle$ of $|M| \setminus A$ such that the type

$$\text{tp}(b_j / A \cup \{b_i : i < j\}, M)$$

is isolated via T_j for every $j < \alpha$, where $T_j = T \cup \text{CD}(A \cup \{b_i : i < j\}, M)$. The sequence $\langle b_i : i < \alpha \rangle$ is called an *isolating sequence for M over A* . The ordinal α is the *length* of the sequence.

It is not hard to see that if T is submodel complete and $N \subseteq M \models T$ then M is almost T -primary over N if and only if M is T -primary over N . We prefer the concept of an almost primary model below because it is more explicit about what property is being exploited, namely submodel completeness.

Theorem 2.3.4 (Henkin's Omitting Type Theorem). *If $L(T)$ is countable and Γ is a countable collection of T -types such that p is not isolated for every $p \in \Gamma$, then there exists a countable model $M \models T$ that omits all the types in Γ .*

We proceed to develop a couple of technical lemmas. We have the following basic fact about an almost primary model satisfying a submodel complete theory:

Lemma 2.3.5. *Suppose T is submodel complete. Let $N \subseteq M \models T$. Then: if M is almost T -primary over N , then for every model $M^* \models T \cup \text{ED}(N)$ there is an elementary embedding from M into M^* over N .*

Proof. Since T is submodel complete, the theory $T \cup \text{ED}(N)$ is complete. This means that for any formula $\varphi(\bar{x})$ and any $\bar{a} \in |N|$ we have

$$M \models \varphi(\bar{a}) \text{ iff } M^* \models \varphi(\bar{a}).$$

Let $\langle (N_i, b_i) : i < \alpha \rangle$ be an almost isolating sequence for M over N . So by definition $N_0 = N$. In order to prove the lemma it is enough to construct a continuous sequence of monomorphisms $g_i : N_i \longrightarrow M^*$ for $i < \alpha$ such that

1. $g_0 = \text{id}_N$,
2. $N_i \models \varphi(\bar{a})$ iff $M^* \models \varphi(g_i(\bar{a}))$ for each formula $\varphi(\bar{x})$ and each $\bar{a} \in N_i$,
3. if $i < j < \alpha$ then $g_i \subseteq g_j$, and
4. if β is a limit then $g_\beta = \bigcup_{i < \beta} g_i$.

The embedding $g = \bigcup_{i < \alpha} g_i$ is as desired. That g is elementary is because submodel completeness implies model completeness (see 2.1.2 and 2.2.6).

Now we proceed to construct the sequence. Due to the clause 4 all we have to do is to make the successor case work. So suppose we have successfully constructed the sequence up to the ordinal $i < \alpha$. Since the complete type $p_i = \text{tp}(b_i / |N_i|, M)$ is isolated via T_i where $T_i = T \cup \text{CD}(N_i, M)$, there exists a formula $\varphi(x; \bar{a}) \in p_i$ isolating it. By the clause 2 we have

$$\varphi(x; \bar{a}) \vdash p_i \Rightarrow \varphi(x; g_i(\bar{a})) \vdash g_i(p_i). \quad (\star)$$

Since $M \models \varphi(b_i; \bar{a})$, we have $M \models \exists x \varphi(x; \bar{a})$, so $M^* \models \exists x \varphi(x; g_i(\bar{a}))$. Let $c_i \in |M^*|$ such that $M^* \models \varphi(c_i; g_i(\bar{a}))$. So by (\star) c_i realizes the type $g_i(p_i)$. Now define a function g_{i+1} by setting $\tau(b_i) \longmapsto \tau(c_i)$ for each term $\tau(x)$ of $L(T_i)$. It is easy to see that this is a well-defined monomorphism from N_{i+1} into M^* which extends g_i and takes b_i to c_i . That the clause 2 is satisfied is, again, because T is submodel complete. \square

In order to build almost primary models we need the next crucial lemma.

Lemma 2.3.6. *Suppose that $L(T)$ is countable and T has the SS-property. Then for*

1. every model $M \models T$,
2. every countable submodel $N \subseteq M$,
3. every formula $\varphi(x; \bar{y})$ and every $\bar{a} \in |N|$ such that $\exists x \varphi(x; \bar{a}) \in T \cup \text{ED}(N)$ but $M \models \neg \varphi(b; \bar{a})$ for every $b \in |N|$,

there is an element $c \in |M| \setminus |N|$ such that the type $\text{tp}(c / |N|, M)$ is isolated and $M \models \varphi(c; \bar{a})$.

Proof. Fix an M , an N , an \bar{a} , and a $\varphi(x; \bar{y})$ as above. Without loss of generality we may assume M is countable as well. Since T has the SS-property, by 2.2.6, the theory $T \cup \text{ED}(N)$ is complete. So $M \models \exists x \varphi(x; \bar{a})$. So $\varphi(M; \bar{a}) \neq \emptyset$ and, by the third condition, $\varphi(M; \bar{a}) \subseteq |M| \setminus |N|$, where $\varphi(M; \bar{a})$ is the set $\{c \in |M| : M \models \varphi(c; \bar{a})\}$. Also note that T is model complete.

Suppose for contradiction we cannot find an element c in M as required. Define a collection Γ of $T \cup \text{ED}(N)$ -types:

$$\Gamma = \{\text{tp}(c/|N|, M) : c \in |M| \setminus |N| \text{ and } M \models \varphi(c; \bar{a})\}.$$

Since Γ is countable, by Henkin's Omitting Type Theorem there is a model $O \models T \cup \text{ED}(N)$ that omits every type in Γ . But T has the SS-property, so we can find two models $M^* \subseteq M$, $O^* \subseteq O$ of T such that there is an isomorphism $h : M^* \cong O^*$ whose restriction to N is id_N . Since $\exists x \varphi(x; \bar{a}) \in T \cup \text{ED}(N)$, there must be some $c \in |M^*| \setminus |N|$ such that $M^* \models \varphi(c; \bar{a})$. Since T is model complete, we deduce

$$\varphi(x; \bar{a}) \in \text{tp}(c/|N|, M^*) = \text{tp}(c/|N|, M).$$

This means that $h(c)$ realizes the $T \cup \text{ED}(N)$ -type $\text{tp}(c/|N|, M)$ in O , contradicting the choice of O . \square

Note that in the above lemma, if N is not a model of T , then there must exist a formula $\exists x \varphi(x; \bar{a}) \in T \cup \text{ED}(N)$ with $\bar{a} \in |N|$ such that $M \models \neg \varphi(b; \bar{a})$ for every $b \in |N|$, because otherwise N would be a model of T by the Tarski-Vaught Test as $T \cup \text{ED}(N)$ is complete. This property is important for our argument. We shall give it a name:

Definition 2.3.7. Let $M \models T$, $N \subseteq M$, and $\bar{a} \in |N|$. We say that $\varphi(x; \bar{a})$ is *critical* for N if $\exists x \varphi(x; \bar{a}) \in T \cup \text{ED}(N)$ and $\varphi(M; \bar{a}) \subseteq |M| \setminus |N|$.

Now the SS-property enables us to construct almost primary models over countable submodels.

Theorem 2.3.8. *If $L(T)$ is countable and T has the SS-property then, for any model $M \models T$ and any countable submodel $N \subseteq M$, N has a T -closure.*

Proof. Fix $N \subseteq M \models T$ such that N is countable. Again we may assume that M is countable as well. So by Lemma 2.3.5 all we need to do is to build an almost T -primary model N^* over N inside M . For this it is enough to build an almost isolating sequence for some model of T over N . The idea here is of course to find a suitable Skolem hull of N inside M such that the type of each “key” new element we find is isolated over all the previous elements.

To be precise, we want to build an almost isolating sequence $\langle (N_i, b_i) : i < \omega \cdot \omega \rangle$ over N such that for

- each $n < \omega$,
- each $\bar{a} \in N_{\omega \cdot n}$, and
- each formula $\varphi(x; \bar{y})$ such that $M \models \exists x \varphi(x; \bar{a})$,

there is an $m < \omega$ such that $M \models \varphi(\tau(b_{\omega \cdot n+m}); \bar{a})$ for some term $\tau(x)$ in the language $L(T \cup \text{ED}(N_{\omega \cdot n+m}))$. It should be clear that $\bigcup_{i < \omega \cdot \omega} N_i = N^*$ is an elementary submodel of M , and hence is almost T -primary over N .

Now we carry out the construction. Start with $N_0 = N$ of course. Suppose $\langle (N_i, b_i) : i < \omega \cdot n \rangle$ is defined. Let $\langle \varphi_k(x; \bar{a}_k) : k < \omega \rangle$ be an enumeration of all the formulas in $T \cup \text{ED}(N_{\omega \cdot n})$ such that for every $k < \omega$ we have $M \models \exists x \varphi_k(x; \bar{a}_k)$ but $M \models \neg \varphi_k(d; \bar{a}_k)$ for every $d \in N_{\omega \cdot n}$. Now suppose we have extended the sequence all the way up to $(N_{\omega \cdot n+k}, b_{\omega \cdot n+k})$ for some $k < \omega$. Let $N_{\omega \cdot n+k+1} = N_{\omega \cdot n+k} + b_{\omega \cdot n+k}$. If there is a $d \in N_{\omega \cdot n+k+1}$ such that $M \models \varphi_{k+1}(d; \bar{a}_{k+1})$ then let $b_{\omega \cdot n+k+1} = b_{\omega \cdot n+k}$. Otherwise by Lemma 2.3.6 we can pick a $b_{\omega \cdot n+k+1} \in |M| \setminus |N_{\omega \cdot n+k+1}|$ such that $M \models \varphi_{k+1}(b_{\omega \cdot n+k+1}; \bar{a}_{k+1})$ and the type $\text{tp}(b_{\omega \cdot n+k+1} / |N_{\omega \cdot n+k+1}|, M)$ is isolated. \square

2.4 The inductive step

The reader may ask: What is preventing us here from simply extending the above theorem to arbitrary theories and arbitrary submodels? One difficulty is this: We do not know how to extend Henkin's Omitting Type Theorem to uncountable languages and hence are unable to develop an analog of Lemma 2.3.6 for uncountable languages. In fact if we simply drop the countability requirement in Henkin's Omitting Type Theorem then it is false. See [11] for discussions. However, in this last section we will show how to circumvent this difficulty if the language in question is countable. For this we need some basic concepts and facts in infinitary combinatorics, in particular stationary sets and Fodor's Lemma.

Throughout the rest of this section T is a theory in a countable language and has the SS-property. Our strategy is to establish an analog of Lemma 2.3.6 for any submodel. Let $M \models T$ and $N \subseteq M$ such that N is uncountable and is not a model of T . We have two cases to consider, namely $\|N\|$ is regular and $\|N\|$ is singular.

Definition 2.4.1. Let α be an ordinal. A sequence $\langle N_i : i < \alpha \rangle$ is an α -resolution of N if

1. N_i is a submodel of N for all $i < \alpha$,
2. if $i < j < \alpha$ then $N_i \subseteq N_j$,

3. $\bigcup_{i < \alpha} N_i = N$.

If, in addition, $\bigcup_{i < \delta} N_i = N_\delta$ for every limit ordinal $\delta < \alpha$, then the sequence is a *continuous α -resolution of N* .

Lemma 2.4.2. *Suppose $\|N\| = \kappa$ is regular and $\varphi(x; \bar{a})$ is critical for N . Then there is an element $c \in \varphi(M; \bar{a})$ such that the type $\text{tp}(c/|N|, M)$ is isolated.*

Proof. Without loss of generality we may assume $\|M\| = \kappa$. Fix a club $C = \langle \alpha_i : i < \kappa \rangle \subseteq \kappa$ and a continuous κ -resolution $\langle N_i : i < \kappa \rangle$ of N such that

1. for all $\alpha_i, \alpha_j \in C$ and $i < j$ we have $|\alpha_i| \leq |\alpha_j \setminus \alpha_i|$,
2. $\|N_i\| = |\alpha_i|$,
3. $\bar{a} \in N_0$.

By the inductive hypothesis we construct a sequence $\langle b_i \in \varphi(M; \bar{a}) : i < \kappa \rangle$ such that each type $\text{tp}(b_i/|N_i|, M)$ is isolated. Fix an enumeration $\langle \phi_i : i < \kappa \rangle$ of all the formulas in the language of $T \cup \text{ED}(N)$ such that for each $\alpha_i \in C$ we have

$$\{i : \phi_i \text{ is a formula in the language of } T \cup \text{ED}(N_i)\} \subseteq \alpha_i.$$

Now define a function $f : C \rightarrow \kappa$ by letting $f(\alpha_i)$ be the least ordinal such that $\phi_{f(\alpha_i)}$ isolates the type $\text{tp}(b_i/|N_i|, M)$. Since f is a pressing-down function on a stationary subset of κ and κ is regular, by Fodor's Lemma, there is a $\gamma < \kappa$ such that $f^{-1}(\gamma) \subseteq C$ is stationary. Clearly for any $\alpha_i, \alpha_j \in f^{-1}(\gamma)$, if $\alpha_i < \alpha_j$ then $\text{tp}(b_i/|N_j|, M) = \text{tp}(b_j/|N_j|, M)$ as they are both isolated by ϕ_γ . So $\text{tp}(b_i/|N|, M) = \text{tp}(b_j/|N|, M)$ for any $\alpha_i, \alpha_j \in f^{-1}(\gamma)$. And this type is isolated by ϕ_γ as desired. \square

For the case that $\|N\|$ is singular we need to work harder. First we formulate the following concept:

Definition 2.4.3. Let $\langle N_i : i < \alpha \rangle$ be an α -resolution of N . Let $\bar{a} \in N_0$. Let $\varphi(x; \bar{a})$ be critical for N . We say that $\mathbf{F} = \langle \varphi_i(x) : i < \alpha \rangle$ is a *spinal sequence of $\varphi(x; \bar{a})$ for $\langle N_i : i < \alpha \rangle$* if:

1. each $\varphi_i(x)$ is a formula in the language of $T \cup \text{ED}(N_i)$,
2. $\varphi_i(M) \neq \emptyset$ and $\varphi_i(M) \subseteq \varphi(M; \bar{a})$ for each $i < \alpha$,
3. if $b \in \varphi_i(M)$ then the type $\text{tp}(b/|N_i|, M)$ is isolated by $\varphi_i(x)$.

We write $\text{dom}(\mathbf{F})$ for the set

$$\{a \in |N| : a \text{ occurs as a parameter in some } \varphi_i(x) \in \mathbf{F}\}.$$

Lemma 2.4.4. *Suppose $\|N\| = \kappa$ is singular and $\varphi(x; \bar{a})$ is critical for N . Then there is an element $c \in \varphi(M; \bar{a})$ such that the type $\text{tp}(c/|N|, M)$ is isolated.*

Proof. As above we may assume $\|M\| = \kappa$. Let $\lambda = \text{cf}(\kappa) < \kappa$. Let $\langle \mu_i : i < \lambda \rangle \subseteq \kappa$ be a strictly increasing sequence of cardinals such that it is unbounded in κ . Let $\langle N_i : i < \lambda \rangle$ be a λ -resolution of N such that $\bar{a} \in N_0$ and $\|N_i\| = \mu_i$.

Let \mathbf{F}_0 be a spinal sequence of $\varphi(x; \bar{a})$ for $\langle N_i : i < \lambda \rangle$. Note that the existence of such a sequence is guaranteed by the inductive hypothesis. We have $|\text{dom}(\mathbf{F}_0)| \leq \lambda$. Now let $K_0 \subseteq N$ be the submodel generated by $\text{dom}(\mathbf{F}_0) \cup \{\bar{a}\}$. Note that $\varphi(x; \bar{a})$ is critical for K_0 . Since $\|K_0\| \leq \lambda < \kappa$, by the inductive hypothesis there is an element $c_0 \in \varphi(M; \bar{a})$ such that $\text{tp}(c_0/|K_0|, M)$ is isolated by some formula $\sigma_0(x)$ in $L(T \cup \text{ED}(K_0))$. Notice that if $\mathbf{F}_0 \subseteq \text{tp}(c_0/|K_0|, M)$ then we are done: in this case $\sigma_0(x)$ isolates the entire \mathbf{F}_0 and each $\varphi_i(x) \in \mathbf{F}_0$ isolates the type $\text{tp}(c_0/|N_i|, M)$, so the type $\text{tp}(c_0/|N|, M)$ is isolated by $\sigma_0(x)$.

Next, since $\varphi(x; \bar{a}) \wedge \sigma_0(x)$ is critical for N (because it contains $\varphi(x; \bar{a})$ as a conjunct), we can find a spinal sequence \mathbf{F}_1 of $\varphi(x; \bar{a}) \wedge \sigma_0(x)$ for $\langle N_i : i < \lambda \rangle$. Clearly \mathbf{F}_1 is also a spinal sequence of $\varphi(x; \bar{a})$ for $\langle N_i : i < \lambda \rangle$. Let $K_1 \subseteq N$ be the submodel generated by $|K_0| \cup \text{dom}(\mathbf{F}_1)$. Then, similarly, we can find an element $c_1 \in \varphi(M; \bar{a})$ and a formula $\sigma_1(x)$ in $L(T \cup \text{ED}(K_1))$ that isolates the type $\text{tp}(c_1/|K_1|, M)$.

Continuing in this fashion we can construct a sequence $\langle (\mathbf{F}_i, c_i, \sigma_i(x)) : i < \lambda^+ \rangle$ such that

1. $c_i \in \varphi(M; \bar{a})$,
2. \mathbf{F}_{i+1} is a spinal sequence of $\varphi(x; \bar{a}) \wedge \sigma_i(x)$ for $\langle N_i : i < \lambda \rangle$,
3. $\sigma_i(x)$ is a formula in $L(T \cup \text{ED}(K_i))$ which isolates the type $\text{tp}(c_i/|K_i|, M)$, where $K_i \subseteq N$ is the submodel generated by the set $\{\bar{a}\} \cup \bigcup_{j \leq i} \text{dom}(\mathbf{F}_j)$,
4. if i is a limit ordinal then \mathbf{F}_i is not defined.

Let $K = \bigcup_{j < \lambda^+} K_j$. Let

$$S_{\lambda^+}^\lambda = \{\alpha < \lambda^+ : \text{cf}(\alpha) = \lambda\},$$

which is a stationary subset of λ^+ . Fix an enumeration of all the formulas in $L(T \cup \text{ED}(K))$ such that for each $\alpha \in S_{\lambda^+}^\lambda$ we have

$$\{i : \phi_i \text{ is a formula in the language of } T \cup \text{ED}(K_\alpha)\} \subseteq \alpha.$$

So again by Fodor's Lemma there is a $\sigma_j(x)$ and a stationary subset $S \subseteq S_{\lambda^+}^\lambda$ such that for all $\alpha \in S$ the type $\text{tp}(c_\alpha/|K_\alpha|, M)$ is isolated by $\sigma_j(x)$.

For any $\alpha, \beta \in S$ with $\alpha < \beta$, consider $\mathbf{F}_{\alpha+1}$. Since $\sigma_\alpha(x)$ is $\sigma_j(x)$, $\mathbf{F}_{\alpha+1}$ is a spinal sequence of $\varphi(x; \bar{a}) \wedge \sigma_j(x)$ for $\langle N_i : i < \lambda \rangle$. So

$$M \models \exists x (\varphi(x; \bar{a}) \wedge \sigma_j(x) \wedge \varphi_i(x))$$

for all $\varphi_i(x) \in \mathbf{F}_{\alpha+1}$ (this is by the second condition in the definition of a spinal sequence above). Since $\sigma_j(x)$ also isolates the complete type $\text{tp}(c_\beta/|K_\beta|, M)$ and $\text{dom}(\mathbf{F}_{\alpha+1}) \subseteq |K_\beta|$, we must have $\mathbf{F}_{\alpha+1} \subseteq \text{tp}(c_\beta/|K_\beta|, M)$. So $\sigma_j(x)$ isolates $\mathbf{F}_{\alpha+1}$. Since each $\varphi_i(x) \in \mathbf{F}_{\alpha+1}$ determines the type over N_i , we see that $\sigma_j(x)$ isolates the type $\text{tp}(c_\beta/|N|, M)$. \square

With these two lemmas we can now simply proceed to build an almost isolating sequence for some model of T over N much in the same way as in Theorem 2.3.8, only now the length of the almost isolating sequence can go up to $\|N\| \cdot \omega$. This proves Theorem 2.2.7.

We end this paper with a question:

Question 2.4.5. Is there an analog of Theorem 2.2.7 for uncountable languages?

Notice that, if T is a theory in an uncountable language and the SS-property and the D-property are not equivalent for T , then there is an $M \models T$ and an $N \subseteq M$ such that the complete theory $T \cup \text{ED}(N)$ is not totally transcendental. This is because primary models always exist for totally transcendental theories.

Chapter 3

Quantifier elimination for the reals with a predicate for the powers of two

Abstract: We give a procedure for eliminating quantifiers for the theory of real closed ordered fields with a predicate for the powers of two. This result was first obtained by van den Dries [18]. His method is model-theoretic, which provides no apparent bounds on the complexity of a decision procedure. In the last section we give a complete axiomatization of the theory of real closed ordered fields with a predicate for the Fibonacci numbers.¹

3.1 Introduction

It was Tarski who first found a decision procedure for the theory of real closed ordered fields. His method was QE. However, his original proof in [49] ran to several dozens of pages and involved a great deal of complex symbolism. It is a daunting task for anyone to decipher the crucial ideas in the proof. Fortunately many significant simplifications and improvements of Tarski's method have been made since the result was first published. One that is highly recommendable is Kreisel and Krivine's presentation in their textbook [35], though, as far as computational efficiency is concerned, it is really not that far away from Tarski's version.

Here we shall quote two key lemmas from their presentation because many claims in this section are inspired by them. The language of the theory of real closed ordered fields has the symbols $0, 1, +, -, \times, <$. In this theory each quantifier-free formula $\varphi(x)$ can be written in the form

$$\bigwedge_{i < n} p_i(x) = 0 \wedge \bigwedge_{i < m} q_i(x) > 0,$$

where $p_i(x)$ and $q_i(x)$ are terms in the standard form, that is, polynomials. For any polynomial p we write $\deg(x, p)$ for the highest degree of x in p . The *degree in x* of $p_i(x) = 0$ is $\deg(x, p_i(x))$. The *degree in x* of $q_i(x) > 0$ is $\deg(x, q_i(x)) + 1$. The *degree in x* of $\varphi(x)$ is the maximum of the degrees of its atomic components.

¹This is joint work with Jeremy Avigad (except the last section), published as [4].

Lemma 3.1.1. *For any quantifier-free formula $\varphi(x)$ of the form*

$$\bigwedge_{i < n} p_i(x) = 0 \wedge \bigwedge_{i < m} q_i(x) > 0,$$

there is a quantifier-free formula $\psi(x)$ which is equivalent to $\varphi(x)$ such that the degree in x of $\psi(x)$ is less than or equal to the least of the degrees in x of $p_i(x) = 0$ (which we assume is not 0).

Lemma 3.1.2. *Let $\varphi(x)$ be a quantifier-free formula. Let a, b be two variables that does not occur in $\varphi(x)$. Assume $a < b$. Then the formula $\exists x (a < x < b \wedge \varphi(x))$ is equivalent to a quantifier-free formula ψ such that ψ does not contain x , each variable in ψ is a, b , or a variable in φ , and no atomic formula in ψ contains both a and b . (Note that the claim can be rephrased accordingly if a, b are closed terms.)*

Now extend the language of real closed ordered fields with a predicate A which, in the intended interpretation, denotes the powers of two, $2^{\mathbb{Z}}$. Adopting the obvious conventions and abbreviations, add the following axioms:

- $\forall x (A(x) \rightarrow x > 0)$
- $\forall x, y (A(x) \rightarrow (A(y) \leftrightarrow A(xy)))$
- $A(2) \wedge \forall x (1 < x < 2 \rightarrow \neg A(x))$
- $\forall x (x > 0 \rightarrow \exists y (A(y) \wedge y \leq x < 2y))$

The first two imply that the A picks out a multiplicative subgroup of the positive elements. In [18], van den Dries showed that the resulting theory admits quantifier elimination in an expanded language. As a result, it is complete and decidable, and, in particular, axiomatizes the real numbers with a predicate for the powers of two.

The theory we have just described includes not only the theory of real closed ordered fields, but also, via an interpretation of integers as exponents, Presburger arithmetic. Thus, van den Dries's result is particularly interesting in that it subsumes two of the most important decidability results of the twentieth century. In recent years, this result has been extended in various directions (see, for example, [28] and [21]).

To establish QE, van den Dries gave a model-theoretic argument. In particular his argument shows that the theory in question has the D-property and QE follows from Theorem 2.2.6. The proof does not provide an explicit procedure, nor does it provide a bound on the length of the resulting formula. Here, we present a proof that makes use of nested calls to a QE procedure for real closed ordered fields, yielding a procedure that is primitive recursive but not elementary. In particular, it requires time $2^0_{O(n)}$ to eliminate a single block

of existential quantifiers, or even a single existential quantifier, where n is the length of the input formula and 2_k^0 denotes a stack of k exponents. Thus, the best bound we can give on the time complexity of the full QE procedure involves $O(n)$ iterates of the stack-of-twos function. We leave it as an open question as to whether one can avoid such nesting and, say, obtain elementary bounds for the elimination of a single existential quantifier.

In Section 3.2, we describe the extension of the theory above that admits QE. Our method of eliminating an existential quantifier proceeds in two steps: first, we eliminate that quantifier in favor of a multiple existential quantifiers over powers of two (the number of which is bounded by the length of the original formula); then we successively eliminate each of these. The first step is described in Section 3.2. In Section 3.3, we prove a number of lemmas that fill out the relationship between the powers of two and the underlying model of real closed ordered fields in a model of the relevant theory; this contains the bulk of the syntactic and algebraic work. In Section 3.4, we use these results to carry out the second step. Finally, in Section 3.5, we show that our procedure satisfies the complexity bounds indicated above.

3.2 The first step

Expand the language of real closed ordered fields to include a unary function λ and a unary predicate D_n for each $n \geq 1$. Let T be the theory given by the axioms above together with the following:

- $D_n(x) \leftrightarrow \exists y (A(y) \wedge y^n = x)$
- $\forall x (x \leq 0 \rightarrow \lambda(x) = 0)$
- $\forall x (x > 0 \rightarrow A(\lambda(x)) \wedge \lambda(x) \leq x < 2\lambda(x))$

In the standard interpretation, λ maps negative real numbers to 0 and rounds positive reals down to the nearest power of two, and D_n holds of numbers of the form 2^i where i is an integer divisible by n .² Note that A and D_1 are equivalent; we will treat them as the same symbol and use the two notations interchangeably.

Our goal is to prove the following:

Theorem 3.2.1. *T admits QE.*

This is Theorem II of [18]. Henceforth, by “formula,” we mean “formula in the language of T .” We will use \bar{x} to denote a sequence of variables x_0, x_1, \dots, x_{k-1} , and we will use notation like $A(\bar{x})$ to denote $A(x_0) \wedge A(x_1) \dots \wedge A(x_{k-1})$.

²For parsimony, 0 can be defined as $1 - 1$ and $A(x)$ by $x > 0 \wedge \lambda(x) = x$. In the next section, we will see that the division symbol is another inessential addition to the language. But in contrast to QE for real closed ordered fields, one can’t eliminate $-$ in terms of $+$; for example, the quantifier-free formula $A(x - y)$, if replaced by $\exists z (z + y = x \wedge A(z))$, would have no quantifier-free equivalent.

To eliminate quantifiers from any formula it suffices to be able to eliminate a single existential quantifier, that is, transform a formula $\exists x \varphi$, where φ is quantifier-free, to an equivalent quantifier-free formula. Since $\exists x (\varphi \vee \psi)$ is equivalent to $\exists x \varphi \vee \exists x \psi$, we can always factor existential quantifiers through a disjunction. In particular, since any quantifier-free formula can be put in disjunctive normal form, it suffices to eliminate existential quantifiers from conjunctions of atomic formulas and their negations. Also, since $\exists x (\varphi \wedge \psi)$ is equivalent to $\exists x \varphi \wedge \psi$ when x is not free in ψ , we can factor out any formulas that do not involve x . Furthermore, whenever we can prove $\forall x (\theta \vee \eta)$, $\exists x \varphi$ is equivalent to $\exists x (\varphi \wedge \theta) \vee \exists x (\varphi \wedge \eta)$; so we can “split across cases” as necessary. We will use all of these facts freely below.

In [18], van den Dries established quantifier elimination by establishing the D-property. The novelty of this test, as compared to more common ones (see the definitions in Section 2.2), lies in the prover’s right to choose an appropriate b in the second clause (see also the discussion in [19]). This clause implies that any existential formula with parameters from the smaller model N that is true in the T -closure of $N + b$ is true in N ; the test works because this clause can be iterated in a countable model to obtain a sequence of T -extensions $N = N_0 \subseteq N_1 \subseteq N_2 \dots \subseteq M$ that eventually picks up every element of M , so any existential formula with parameters from N true in M is true in N (see Theorem 2.2.6). On the syntactic side, this iteration translates to the simple observation that to eliminate a single existential quantifier from an otherwise quantifier-free formula, it suffices to eliminate additional existential quantifiers from an equivalent existential formula. Thus, our effective proof is based on the following two lemmas:

Lemma 3.2.2. *Every formula of the form $\exists w \psi$, with ψ quantifier-free, is equivalent to a disjunction of formulas of the form $\exists \bar{x} (A(\bar{x}) \wedge \varphi)$, with φ quantifier-free.*

Lemma 3.2.3. *Every formula of the form $\exists x (A(x) \wedge \varphi)$, with φ quantifier-free, is equivalent to a formula that is quantifier-free.*

The remainder of this section is devoted to proving the first of these two lemmas. The next lemma explains why the new existentially quantified variables are helpful.

Lemma 3.2.4. *Every existential formula is equivalent, in T , to an existential formula in which λ does not occur and the predicates D_i are applied only to variables.*

Proof. First, replace $\dots D_i(t) \dots$ by $\exists z (z = t \wedge \dots D_i(z) \dots)$. Then, iteratively simplify terms involving λ , noting that $\psi(\lambda(t))$ is equivalent to

$$(t \leq 0 \wedge \psi(0)) \vee \exists z (A(z) \wedge z \leq t < 2z \wedge \psi(z)),$$

and that the existential quantifier can be brought to the front. \square

Thus to prove Lemma 3.2.2, we are reduced to showing that when ψ is quantifier-free, λ does not occur in ψ , and the predicates D_i occurring in ψ are applied only to variables, the formula $\exists \bar{x} \psi$ is equivalent to one of the form $\exists \bar{x} (A(\bar{x}) \wedge \varphi)$, where φ is quantifier-free. In general, $\exists x \theta(x)$ is equivalent to

$$\exists x > 0 \theta(x) \vee \theta(0) \vee \exists x > 0 \theta(-x).$$

Moreover, assuming $x > 0$, any subformula of the form $D_i(-x)$ is equivalent to falsity. So, across a disjunction, we are reduced to proving the claim for formulas of the form $\exists \bar{x} > 0 \psi(\bar{x})$, where ψ satisfies the criteria above.

In T we can factor out the greatest power of two from any positive x , that is we can prove

$$x > 0 \rightarrow \exists y \exists z (A(y) \wedge 1 \leq z < 2 \wedge x = yz).$$

Since we have $1 \leq z < 2 \leftrightarrow (z = 1 \vee 1 < z < 2)$, we can transform our formula into a disjunction of formulas of the form

$$\exists \bar{y}, \bar{z} (A(\bar{y}) \wedge 1 < \bar{z} < 2 \wedge \psi)$$

where ψ once again meets the criteria above, except that the predicates D_i are applied to expressions of the form yz . When $1 < z < 2$, each $D_i(yz)$ is false, so we can rewrite the formula above as

$$\exists \bar{y} (A(\bar{y}) \wedge \theta \wedge \exists \bar{z} \eta)$$

where θ is a conjunction of predicates of the form $D_n(y)$ and negations of such, and $\exists \bar{z} \eta$ is in the language of real closed ordered fields. We can therefore replace $\exists \bar{z} \eta$ by a quantifier-free formula, using any QE procedure for real closed ordered fields.

3.3 Reasoning about powers of two

Our goal in this section is to establish some general relationships between the powers of two in a model of our theory, T , and the underlying real closed field.

Definition 3.3.1. Let φ be a quantifier-free formula. We say φ is *simple in x* if the following hold:

1. every equality or inequality occurring in φ is either of the form $p(x) = 0$ or $q(x) > 0$, where $p(x), q(x)$

are polynomials in x ; that is, they are of the form $\sum_{i \leq n} s_i x^i$ where each s_i is a term that does not involve x .

2. for every atomic formula $D_n(t)$ occurring in φ , either t does not contain x or t is of the form $2^r x$ for some integer r such that $0 \leq r < n$.

The main goal of this section is to prove the following proposition:

Proposition 3.3.2. *Let φ be any quantifier-free formula. Then there is a quantifier-free formula φ' such that φ' is simple in x and T proves $A(x) \rightarrow (\varphi \leftrightarrow \varphi')$.*

In semantic terms, this says the following: let N be any model of T , let $M \subseteq N$ be a model of T^\forall , that is the universal fragment of T , and let x be a power of two in N . Then the structure of $M + x$ is completely determined by the structure of M , the structure of $M + x$ as an ordered ring, and the divisibility properties of the exponent of x .

First, we need to note some easy facts about λ and the predicates D_i .

Lemma 3.3.3. *For any n , T proves*

$$0 < u < x < 2^n u \wedge A(x) \rightarrow (x = 2\lambda(u) \vee \dots \vee x = 2^n \lambda(u)).$$

Lemma 3.3.4. *For any n , T proves*

$$A(x) \rightarrow D_n(x) \vee D_n(2x) \vee \dots \vee D_n(2^{n-1}x).$$

Although we have not included the division symbol in the language of T , we can define the function r/s by making $x/y = z$ equivalent to $x = yz \vee (y = 0 \wedge z = 0)$. In the proof of Proposition 3.3.2, it will be useful to act as though the division symbol is part of the language. The next few lemmas show that if θ is any quantifier-free formula in the expanded language with division, there is a quantifier-free formula θ' in the language without division such that $T \vdash \theta \leftrightarrow \theta'$.

Lemma 3.3.5. *From the hypotheses $0 < x$ and $0 < y$, T proves*

$$x\lambda(y) < y\lambda(x) \rightarrow \lambda(x/y) = \lambda(x)/2\lambda(y)$$

and

$$x\lambda(y) \geq y\lambda(x) \rightarrow \lambda(x/y) = \lambda(x)/\lambda(y).$$

Proof. An easy calculation shows that if $x/y < \lambda(x)/\lambda(y)$, then $\lambda(x/y) = \lambda(x)/2\lambda(y)$; and otherwise, $\lambda(x/y) = \lambda(x)/\lambda(y)$. \square

Lemma 3.3.6. *If θ is any quantifier-free formula involving the division symbol, there is a quantifier-free formula θ' in which the division symbol does not occur in the scope of λ , such that $T \vdash \theta \leftrightarrow \theta'$.*

Proof. This can be done by iterating the previous lemma. To measure the nesting of λ 's and division symbols, we define the “ λ -depth of the division symbol in t ,” $\Lambda^\dagger(t)$, recursively, as follows:

1. $\Lambda^\dagger(t) = 0$ if the division symbol does not occur in the scope of λ in t ;
2. if t is $t_1 + t_2$, $t_1 - t_2$, $t_1 \times t_2$, or t_1/t_2 , then $\Lambda^\dagger(t) = \max\{\Lambda^\dagger(t_1), \Lambda^\dagger(t_2)\}$;
3. assuming the division symbol occurs in t , $\Lambda^\dagger(\lambda(t)) = \Lambda^\dagger(t) + 1$.

The previous lemma shows that, using a case disjunction over the possibilities for the signs of the numerator and denominator, we can eliminate one term t such that the λ -depth of the division symbol in t is maximal, in favor of terms in which the λ -depth of the division symbol is smaller. Lemma 3.3.6 follows, by a primary induction on this maximal depth, and a secondary induction on the number of terms of this depth. \square

Lemma 3.3.7. $T \vdash A(x) \wedge A(y) \rightarrow (D_n(x/y) \leftrightarrow \bigvee_{i < n} (D_n(2^i x) \wedge D_n(2^i y)))$.

Proof. The right-to-left direction is easy: if $z^n = 2^i x$ and $w^n = 2^i y$ then $(z/w)^n = x/y$. Proving the other direction is not much more difficult, using Lemma 3.3.4. \square

Proposition 3.3.8. *Let θ be any quantifier-free formula involving division. Then there is a quantifier-free formula θ' that does not involve division, such that $T \vdash \theta \leftrightarrow \theta'$.*

Proof. Using Lemma 3.3.6, we can assume that division does not occur in the scope of any λ in θ . So each atomic formula $D_n(t)$ can be put in the form $D_n(r/s)$, where the division symbol does not occur in r and s . Across a case disjunct, we can assume r and s are positive. Then $D_n(r/s)$ is equivalent to

$$\lambda(r/s) = r/s \wedge D_n(\lambda(r/s)).$$

Using Lemma 3.3.5, we can replace $\lambda(r/s)$ by either $\lambda(r)/\lambda(s)$ or $\lambda(r)/2\lambda(s)$. Then using Lemma 3.3.7 we can replace $D_n(\lambda(r)/\lambda(s))$ or $D_n(\lambda(r)/2\lambda(s))$ by a disjunction in which the division symbol does not occur.

Once all divisibility symbols are removed from the λ 's and D_n 's, we can clear division from the remaining equalities and inequalities by multiplying through. \square

It therefore suffices to prove Proposition 3.3.2 where φ' is a quantifier-free formula in the expanded language with the division symbol. The next few lemmas, then, make use of this expanded language.

Lemma 3.3.9. *Let $p(x)$ be the term $\sum_{i \leq n} a_i x^i$. Then there is a sequence of quantifier-free formulas $\theta_0, \dots, \theta_{m-1}$ such that T proves*

$$A(x) \wedge p(x) > 0 \rightarrow \bigvee_{k < m} \theta_k,$$

where each θ_k is of one of the following forms:

- $\lambda(p(x)) = 2^r \lambda(a_i) x^i$ for some $-1 \leq r \leq n$,
- $x^e = \frac{2^r \lambda(a_i)}{\lambda(-a_j)}$ or $x^e = \frac{2^r \lambda(-a_j)}{\lambda(a_i)}$, for some e, i, j , and r such that $1 \leq e \leq n$, $0 \leq i, j \leq n$, and $-(n+1) \leq r \leq (n+1)$.

Proof. Argue in T . Using a disjunction on all possible cases, we can write $p(x)$ as $a_i x^i + a_j x^j + \hat{p}(x)$, where $a_i x^i$ is the largest summand and $a_j x^j$ the least summand. Note that we have $a_i x^i > 0$, $i \neq j$, $p(x) \leq (n+1)a_i x^i$, and

$$p(x) - a_i x^i = a_j x^j + \hat{p}(x) \geq n a_j x^j.$$

We now distinguish between two cases, depending on whether $p(x)$ is roughly the same size as $a_i x^i$ or sufficiently smaller.

In the first case, suppose we have $p(x) \geq (a_i x^i)/2$. This means we have

$$(a_i/2)x^i \leq p(x) \leq (n+1)a_i x^i \leq 2^n a_i x^i$$

and so

$$(\lambda(a_i)/2)x^i \leq \lambda(p(x)) \leq 2^n \lambda(a_i) x^i.$$

This yields a disjunction of clauses of the first type, by Lemma 3.3.3.

In the second case, we have $p(x) < (a_i x^i)/2$. This means that $a_j x^j$ must be negative and roughly comparable to $a_i x^i$ in absolute value; that is $a_j < 0$ and

$$(a_i/2)x^i < a_i x^i - p(x) \leq -n a_j x^j,$$

and so

$$0 < (a_i/(-a_j))x^{i-j} \leq 2n \leq 2^n.$$

Using Lemma 3.3.5 and Lemma 3.3.3 we get a disjunction of clauses of the second type. □

Lemma 3.3.10. *In Lemma 3.3.9, if the assumption is changed to $A(x) \wedge p(x) = 0$, then in the conclusion we can assume that each θ_k is of the second form.*

Proof. This is exactly as in the second case of the previous proof. □

Lemma 3.3.11. *In the conclusion of Lemma 3.3.9, we may demand that each θ_k is of the form $\lambda(p(x)) = sx^i$ for some $0 \leq i \leq n$ and some term s that does not contain x .*

Proof. The proof is by induction on the degree of x in $p(x)$. The lemma is trivial if the degree of x in $p(x)$ is 0.

Now assume that the degree of x in $p(x)$ is n and the lemma holds whenever the degree is less than n . By Lemma 3.3.9, T proves a disjunction $\bigvee \sigma_l$, with σ_l of one of those two forms. Each σ_l of the first form there is already as required. For each σ_l of the second form, consider a new term $\hat{p}(x)$, which is obtained by substituting the right-hand side of σ_l for x^e in $p(x)$. Notice that the degree of x in $\hat{p}(x)$ is less than n , and clearly T proves $p(x) = \hat{p}(x) \wedge \hat{p}(x) > 0$. By the inductive hypothesis we may replace σ_l in $\bigvee \sigma_l$ by a disjunction $\bigvee \theta_k$ which is of the required form. □

As was the case with the division symbol, we will iterately “squeeze” x ’s out from within the λ symbols. Thus we introduce the following definitions:

Definition 3.3.12. Let t be a term. Define the λ -depth of x in t , $\Lambda(x, t)$, recursively, as follows:

1. $\Lambda(x, t) = 0$ if x is not in the scope of any λ ;
2. if t is $t_1 + t_2$, $t_1 - t_2$, $t_1 \times t_2$, or t_1/t_2 , then $\Lambda(x, t) = \max\{\Lambda(x, t_1), \Lambda(x, t_2)\}$;
3. if t is $\lambda(t_1)$ and t_1 contains x , then $\Lambda(x, t) = \Lambda(x, t_1) + 1$.

Definition 3.3.13. Let φ be a formula. Define the λ -depth of x in φ by

$$\Lambda(x, \varphi) = \max\{\Lambda(x, t) : t \text{ is a term that contains } x \text{ and occurs in } \varphi\}.$$

Lemma 3.3.14. *Let φ be any quantifier-free formula. Then there is a quantifier-free formula φ' such that $T \vdash A(x) \rightarrow (\varphi \leftrightarrow \varphi')$, and $\Lambda(x, \varphi') = 0$.*

Proof. The proof is by induction on the λ -depth of x in φ . The lemma is trivial if $\Lambda(x, \varphi) = 0$.

Assume $\Lambda(x, \varphi) = n > 0$ and the lemma holds for every quantifier-free formula ψ if $\Lambda(x, \psi) < n$. Let $\lambda(p_0), \dots, \lambda(p_{m-1})$ be all the different terms in φ with $\Lambda(x, p_i) = 0$ for all $i < m$. Across a case disjunction we can assume $p_i > 0$ for all $i < m$, since otherwise we can replace $\lambda(p_i)$ by 0. By Lemma 3.3.6, we may

assume that each p_i is a polynomial in x . By Lemma 3.3.11, T proves $\varphi \leftrightarrow \bigvee(\tau_l \wedge \sigma_l)$, where each τ_l is of the form $\bigwedge_{i < m} \lambda(p_i(x)) = s_i x^{j_i}$, and each σ_l is obtained by substituting $s_i x^{j_i}$ for $\lambda(p_i)$ in φ . Clearly T proves $\lambda(p_i(x)) = s_i x^{j_i} \leftrightarrow A(s_i) \wedge s_i x^{j_i} \leq p_i(x) < 2s_i x^{j_i}$. Now since $\Lambda(x, \sigma_l) < n$, we may apply the inductive hypothesis to each σ_l and the lemma is proved. \square

Lemma 3.3.15. *Let p be a term such that $\Lambda(x, p) = 0$. Then for any n there is a sequence of terms p_k such that*

- T proves $A(x) \wedge p > 0 \rightarrow (D_n(p) \leftrightarrow \bigvee(p = p_k \wedge D_n(p_k)))$,
- each p_k is of the form $s x^i$, where s is a term that does not contain x .

Proof. Using Lemma 3.3.7, we can assume that p is a polynomial in x . We can replace $D_n(p)$ by $p = \lambda(p) \wedge D_n(\lambda(p))$, and then by Lemma 3.3.11, across a disjunction we may replace $\lambda(p)$ in each disjunct by a term of the form $s x^i$, where s does not contain x . (Note that here no formulas like the τ_l 's in the previous lemma are needed.) \square

Lemma 3.3.16. *Let s be a term that does not contain x . Then for any n, i there is a sequence of formulas θ_k such that T proves*

$$A(x) \rightarrow (D_n(s x^i) \leftrightarrow \bigvee \theta_k),$$

and each θ_k is of the form $D_n(2^w s) \wedge D_n(2^r x)$ for some $0 \leq w, r < n$.

Proof. Since for each n , from the assumption $A(x)$, T proves $\bigvee_{j < n} D_n(2^j x)$, it is straightforward to see that $D_n(s x^i)$ is equivalent to a disjunction each of whose disjuncts is of the specified form. \square

We are finally ready to prove Proposition 3.3.2.

Proof. Given φ , first use Lemma 3.3.14 to eliminate x from the scope of any λ . Then use Lemma 3.3.15 to ensure the atomic formulas involving D_n are in the form $D_n(s x^i)$, where s does not involve x . (This will require splitting across cases depending on whether $p > 0$ or $p \leq 0$; in the latter case, $D_n(p)$ is equivalent to \perp .) Finally, use Lemma 3.3.16 to ensure that all the atomic formulas involving D_n are in the required form. \square

We close with some consideration about the predicates D_n which are analogous to considerations that arise in the context of QE for Presburger arithmetic. Remember that when n is a positive integer and s is a non-negative integer, $D_n(2^s x)$ asserts, in the intended interpretation, that x is equal to 2^t for some integer t , and n divides $s + t$; in other words, the exponent of x is congruent to $-s$ modulo n . Let θ be any boolean

combination of predicates of the form $D_n(2^s x)$, and let M be the least common multiple of these various n . Then in T one can show that there is an x satisfying θ if and only if for any w satisfying $A(w)$ we have

$$\theta(w) \vee \theta(2w) \vee \theta(4w) \vee \dots \vee \theta(2^{M-1}w),$$

and, in particular, if and only if

$$\theta(1) \vee \theta(2) \vee \theta(4) \vee \dots \vee \theta(2^{M-1}).$$

Moreover, T can decide the truth or falsity of this last sentence. So we have:

Lemma 3.3.17. *With θ and M as above, either T proves $\forall x \neg\theta$, or it proves*

$$\forall u (0 < u \rightarrow \exists x (u \leq x < 2^M u \wedge \theta)).$$

3.4 Eliminating a quantifier over powers of two

We are now ready to prove Lemma 3.2.3, which asserts that every formula of the form $\exists x (A(x) \wedge \varphi)$, with φ quantifier-free, is equivalent to a formula that is quantifier-free. By Proposition 3.3.2, we can assume that φ is simple, which is to say, x does not occur in the scope of any λ and all divisibility assertions involving x are of the form $D_n(2^r x)$. Put φ in disjunctive normal form, replace negated equalities $s \neq t$ by $s < t \vee t < s$, and replace negated inequalities $s \not< t$ by $t < s \vee t = s$. Rewrite equalities and inequalities so that they are of the form $p(x) = 0$ and $q(x) > 0$, where $p(x)$ and $q(x)$ are polynomials in x . Factoring existential quantifiers through disjunctions and getting rid of atomic formulas that do not depend on x , we are reduced to eliminating quantifiers of the form $\exists x (A(x) \wedge \varphi)$ where φ is a conjunction of formulas of the following types:

- $p(x) = 0$, where p is a polynomial,
- $q(x) > 0$, where q is a polynomial,
- $D_n(2^r x)$, where $0 \leq r < n$, or
- $\neg D_n(2^r x)$, where $0 \leq r < n$.

Splitting across a disjunction, we can assume that when a conjunct of the form $p(x) = 0$, not all the coefficients are zero. By Lemma 3.3.10, we can assume that one of the conjuncts is of the form $x^e = s$, where x does not occur in s . In that case, each conjunct $D_n(2^r x)$ is equivalent to $D_{ne}(2^{re} x^e)$ and hence

$D_{ne}(2^{re}s)$ (and $A(x)$, in particular, is equivalent to $D_e(s)$). But now x no longer occurs in these formulas, and so they can be brought outside the scope of the existential quantifier. The resulting existential formula is then essentially in the language of real closed ordered fields. By this last phrase we mean that it is of the form $\exists x \alpha(x, t_0, \dots, t_{k-1})$, where $\alpha(x, y_0, \dots, y_{k-1})$ is in the language of real closed ordered fields. Treating the terms t_0, \dots, t_{k-1} in the expanded language as parameters, we can therefore replace it by an equivalent quantifier-free formula using any QE procedure for real closed ordered fields.

We are thus reduced to eliminating an existential quantifier of the form

$$\exists x \left(\bigwedge q_i(x) > 0 \wedge \theta(x) \right) \tag{3.4.1}$$

where θ is a conjunction of formulas of the form $D_n(2^r x)$ and negations of such that includes at least the formula $A(x)$. By Lemma 3.3.17, either T proves that θ is false for every x , or there is a natural number M such that T proves that for any $u > 0$, that θ is satisfied by some x in the interval $[u, 2^M u]$. In the first case, T proves that formula 3.4.1 is false. So we only have to worry about the second case. Fix such an M for the remainder of the discussion.

Arguing in T , suppose formula (3.4.1) holds. There are two possibilities: either there is a “large” interval on which $\bigwedge q_i(x) > 0$, that is, an interval of the form $[u, 2^M u]$; or there is an x satisfying $A(x) \wedge \bigwedge q_i(x) > 0 \wedge \theta$, but it is trapped between a u and a v with $q_i(u) = 0$ for some i , $q_j(v) = 0$ for some j , and $v < 2^M u$. Thus formula (3.4.1) is equivalent to a disjunction of the formula

$$\exists u > 0 \forall x (u \leq x \leq 2^M u \rightarrow \bigwedge q_i(x) > 0)$$

and the formulas

$$\exists u > 0 (q_j(u) = 0 \wedge \exists x (u < x \leq 2^M u \wedge \bigwedge q_i(x) > 0 \wedge \theta(x))$$

for the various j . To see this, note that if formula (3.4.1) holds, then by the previous discussion one of these formulas holds; and conversely, each of these formulas implies (3.4.1).

The first of these formulas is essentially in the language of real closed ordered fields, so these quantifiers can be eliminated. The second formula is equivalent to

$$\exists u_1, u_2 (A(u_1) \wedge 1 \leq u_2 < 2 \wedge q_j(u_1 u_2) = 0 \wedge$$

$$\exists x (u_1 < x \leq 2^M u_1 \wedge \bigwedge q_i(x) > 0 \wedge \theta(x)).$$

In this case, we can replace the inner existential quantifier over x by a disjunction, so that the entire formula is equivalent to a disjunction of formulas of the form

$$\exists u_1, u_2 (A(u_1) \wedge 1 \leq u_2 < 2 \wedge q_j(u_1 u_2) = 0 \wedge \bigwedge \hat{q}_i(u_1) > 0 \wedge \hat{\theta}(u_1)),$$

where each $\hat{q}_i(u_1)$ is $q_i(2^r u_1)$ for some r , and similarly for $\hat{\theta}(u_1)$. In particular, $\hat{\theta}(u_1)$ is a conjunction of formulas of the form $D_i(2^r u_1)$, and their negations.

Think of $q_j(u_1 u_2)$ as a polynomial in u_1 with coefficients of the form $s u_2^n$, where s does not involve u_1 or u_2 . By Lemma 3.3.10, across a disjunction we may add a clause of the form $u_1^e = 2^r \lambda(s u_2^n) / \lambda(t u_2^m)$. Splitting on cases of the form $2^l \leq u_2^h < 2^{l+1}$ we can simplify each of these to an expression of the form $u_1^e = 2^k \lambda(s) / \lambda(t)$ for some integer k . By Lemma 3.3.17, $A(u_1) \wedge \hat{\theta}(u_1)$ is equivalent to a formula $\bar{\theta}$ which now involves neither u_1 nor u_2 , and hence can be brought outside the existential quantifier. We are thus reduced to eliminating quantifiers from a formula of the form

$$\begin{aligned} \exists u_1, u_2 (1 \leq u_2 < 2 \wedge u_1^e = 2^k \lambda(s) / \lambda(t) \wedge 2^l \leq u_2^h < 2^{l+1} \wedge \\ q_j(u_1 u_2) = 0 \wedge \bigwedge \hat{q}_i(u_1) > 0). \end{aligned}$$

We can eliminate these quantifiers using a QE procedure for real closed ordered fields. This completes the proof of Lemma 3.2.3, and hence the proof of our main theorem, Theorem 3.2.1.

Note that there is nothing special about the number 2 in our quantifier elimination procedure: inspection of the proofs shows that the arguments go through unchanged for any real algebraic number $\alpha > 1$. There are various ways to represent the real algebraic numbers; for example, we can represent α by providing a polynomial, $p(x)$, of which it is a root, together by a pair of rational numbers u and v isolating α from the other roots of p . In that case, we simply replace 2 by a new constant, c , in the axioms, and then add the following:

- $p(c) = 0$
- $u < c < v$

As noted in [21], this implies that the resulting theory is decidable. To see this, it suffices to see that any quantifier-free sentence φ is decidable. But we can do this using the decision procedure for real closed ordered fields to iteratively compute the values of $\lambda(t)$ for any t involving the field operations and c , and then to determine the truth of terms of atomic formulas $D_n(t)$. (For explicit algorithms for computing with real algebraic numbers, see [8].)

3.5 Complexity analysis

In this section we establish an upper bound on the complexity of our elimination procedure.

For the theory of real closed ordered fields, the best known upper bound for a QE procedure, in terms of the length of the input formula, is $2^{2^{O(n)}}$. This is originally due to Collins [15], and, independently, Monk and Solovay. There are more precise bounds that depend on various parameters, such as the number of quantifier alternations and the degrees of the polynomials in the formula; see, for example, [7] and [8]. In particular, a block of existential quantifiers can be eliminated in time $2^{O(n)}$. The best lower bound for the full QE procedure is $2^{O(n)}$, by Fischer and Rabin [27], and applies even to just the additive fragment. The best upper bound for Presburger arithmetic is $2_3^{O(n)}$ (see [26] and [52]) and is essentially sharp (see [53]).

Our bounds are far worse. Consider what our procedure does when given a formula with a single block of existential quantifiers:

1. First, replace this by a disjunction of formulas of the form

$$\exists \bar{y} (A(\bar{y}) \wedge \exists \bar{z} (1 < \bar{z} < 2 \wedge \psi))$$

where ψ is in the language of real closed ordered fields.

2. Then, use an elimination procedure for real closed ordered fields to eliminate the quantifiers $\exists \bar{z}$.
3. Successively eliminate the innermost quantifier over a power of two, as follows:

- (a) Call the relevant formula $\exists x (A(x) \wedge \varphi)$. Apply Proposition 3.3.2, to reduce φ to a formula that is simple in x .
- (b) Put the new φ in disjunctive normal form, split across a disjunction, and remove atomic formulas that do not involve x , so that each formula is of the form

$$\exists x (A(x) \wedge \bigwedge p_i(x) = 0 \wedge \bigwedge q_j(x) = 0 \wedge \theta)$$

where θ is a conjunction of formulas of the form $D_n(2^r x)$ and negations of such, and in each disjunction where a disjunct of the form $p(x) = 0$ occurs, we can assume p is not identically 0.

- (c) In each disjunct where a conjunct of the form $p(x) = 0$ occurs, apply Lemma 3.3.10, factor out the divisibility predicates, D_n , and call a QE procedure for real closed ordered fields.
- (d) In the remaining disjuncts, again, split across a disjunct; in one case, we call a QE procedure for real closed fields right away; in another, we expand a bounded existential quantifier into a

disjunction, and then call the elimination procedure for real closed ordered fields.

Note that each iteration of the inner loop, 3, requires at least one call to a QE procedure for real closed ordered fields. Each of these calls can be carried out in time, say, $2^{2^{O(n)}}$, where n is the length of the relevant formula. But then the next iteration of the loop will involve calls to the QE procedure for real closed ordered fields on a formula that is potentially much longer. Thus, part 3 of the procedure requires an exponential stack of Cm twos, for some constant C , where m is the number of existential quantifiers over powers of two that need to be eliminated.

In this section, we will confirm that such an upper bound can be obtained. To that end, it is sufficient to show that each pass of the inner loop is elementary, which is to say, it can be computed in time bounded by some fixed stack of exponents to the base 2. Note that after the first step, the number of quantifiers over powers of two is bounded by the length of the original formula (in fact, it is bounded by the number of A 's and λ 's in the original formula). Thus our procedure for eliminating a block of existential quantifiers runs in time $2_{O(n)}^0$, where n is the length of the original formula.

We have been unable to eliminate this nesting of calls to a procedure for real closed ordered fields. Efficient procedures for this latter theory avoid putting formulas in disjunctive normal form; for example, Collins's cylindrical algebraic decomposition procedure obtains a description of cells, depending on the coefficients, on which a set of polynomials have constant sign. In our setting, suppose we are given a formula $\exists \bar{x} (A(\bar{x}) \wedge \eta \wedge \theta)$, where η contains only equalities and inequalities between polynomials, and θ consists of divisibility conditions D_n on the exponents of the x 's. One might start by applying Collins's procedure to the polynomials occurring in η . Then, given a description of the various cells (depending on the other parameters in the formula), one needs to determine which cells contain points with coordinates that are powers of two, with exponents satisfying the requisite divisibility conditions. For one dimensional cells, our procedure relies on a simple disjunction: if the cell is large enough, one is guaranteed a solution, and otherwise one need only test a finite number of cases. For multidimensional cells, however, the situation is more complex, and we do not see how one can proceed except along the lines we have described above. It is thus an interesting question as to whether it is possible to obtain elementary bounds on a procedure for eliminating a single block of quantifiers. Given our failure to do so, we have not taken great pains to bound the number of exponents in the time bound on the inner loop, which would merely improve the constant bound implicit in the $O(n)$.

For the discussion which follows, we define the *length* of a formula in the language of T to be the number of symbols in a reasonable formulation of the first-order language, with the following exception: we count the length of each symbol D_n as n , rather than, say, one plus the binary logarithm of n . This choice is a

pragmatic one in that it simplifies the analysis, and our results below then imply the corresponding results for the alternative definition of length. A more refined analysis might take both the length of the formula and a bound on the n 's occurring in atomic formulas $D_n(t)$, but that does not seem to help much.

It seems that the most delicate part of our task is showing that one can remove the division symbols, and “squeeze” variables ranging over powers of two out of the λ symbols that are repeatedly introduced after the first step of the procedure, as required in step 3(a). A priori, the procedures described in Section 3.3 look as though they may be non-elementary. The next few lemmas show that this is not the case, by keeping careful track of the terms and formulas that need to be dealt with in the disjunctions.

Lemma 3.5.1. *Let t be a term with length l . Then there is a sequence of terms $\langle t_k : k < 2^l \rangle$ such that*

- $T \vdash \bigvee_{k < 2^l} t = t_k$,
- each t_k is of the form r/s , where r and s are division-free terms, and
- each t_k has length at most 2^l .

Proof. This can be proved by a straightforward induction on terms. Suppose t is of the form $t_1 + t_2$, where the length of t_1 is l_1 and the length of t_2 is l_2 . By the inductive hypothesis, t is equal to one of at most $2^{l_1} 2^{l_2} \leq 2^l$ terms of the form $r_1/s_1 + r_2/s_2$, where $r_1, s_1, r_2,$ and s_2 are division-free, the length of r_1/s_1 is at most 2^{l_1} , and the length of r_2/s_2 is at most 2^{l_2} . But then the length of $(r_1 s_2 + r_2 s_1)/s_1 s_2$ is at most $2(2^{l_1} + 2^{l_2}) < 2^l$, as required.

If t is of the form $\lambda(t_1)$, the claim follows from the inductive hypothesis, using Lemma 3.3.6. The other cases are similar. □

Lemma 3.5.2. *Let φ be a quantifier-free formula with length l . Then there is a quantifier-free division-free formula φ' with length $2^{O(l)}$ such that $T \vdash \varphi \leftrightarrow \varphi'$.*

Proof. Enumerate all the different terms t_0, \dots, t_{m-1} in φ such that, for each $i < m$, s_i is not a proper subterm of any term in φ . Using the above lemma we can have a sequence of quantifier-free formulas φ_j for $j < 2^l$ each of which is obtained by replacing each t_i with an appropriate term and therefore has length less than 2^l . Notice that for each φ_j , as indicated in Lemma 3.3.6, there are some division-free atomic formulas that T used to derive the equalities in question. Clearly for each φ_j there are less than l such atomic formulas, each of which has length less than $2^{O(l)}$. Let σ_j be the conjunction of them all. Let φ' be the formula $\bigvee_{j < 2^l} (\varphi_j \wedge \sigma_j)$. The length of φ' is again bounded by $2^{O(l)}$, and clearly $T \vdash \varphi \leftrightarrow \varphi'$.

Finally, we need to clear denominators from atomic formulas of the form $r/s < t/u$ and $r/s = t/u$, and deal with atomic formulas of the form $D_n(r/s)$. The first two require a disjunction over cases, depending

on whether denominators are positive, negative, or zero. The third set of atomic formulas is handled as described in the proofs of Lemma 3.3.7, 3.3.8. But each atomic formula occurring in a disjunct occurs to an atomic formula in the original formula, φ , and there are at most l of these. It is not hard to verify that the corresponding increase in length can be absorbed into the bound $2^{O(l)}$. \square

Lemma 3.5.3. *Let $\lambda(t)$ be a term, where the length of t is l and x does not occur in the scope of any division symbol in t . Then there is a sequence of terms $\langle t_k : k < 2^{8l^2 \log l} \rangle$ such that*

- $T \vdash A(x) \wedge t > 0 \rightarrow \bigvee_{k < 2^{8l^2 \log l}} (\lambda(t) = t_k)$,
- each t_k is of the form sx^i , where s is a term that does not contain x and $i < l$,
- each t_k has length at most 2^{4l} .

Proof. For any polynomial p in x , clearly the number of possible values of $\lambda(p)$ of the form sx^i , as in Lemma 3.3.11, depends on the degree n of x in p . So let $f(n)$ denote the number of possible values of $\lambda(p)$. Observe that the value of $\lambda(p)$ is determined in the first case of Lemma 3.3.9, and when $e = 1$ in the second case. An calculation shows that there are no more than $(n + 1)(n + 2)$ possibilities in the first case, no more than $2n(2n + 2)$ possibilities in the second case when $e = 1$, and no more than $(n + 1)(n - 1)2(n + 2)$ possibilities for all the remaining values of e . Hence we have the following equation:

$$f(n) \leq (n + 1)(n + 2) + 2n(2n + 2) + (n + 1)(n - 1)2(n + 2)f(n - 1).$$

This can be simplified as $f(n) < 10(n + 2)^3 f(n - 1)$. So we have $f(n) < 2^{8n \log(n+2)}$. Let the length of p be l . Since $n + 2 < l$, we have $f(n) < 2^{8l \log l} < 2^{8l^2 \log l}$.

Now the proof proceeds by induction on the λ -depth of x in t . If $\Lambda(x, t) = 0$, then t is a polynomial in x . So we apply the above analysis to t and obtain no more than $2^{8l \log l}$ possible values of $\lambda(t)$ which are all of the form sx^i for some $i < l$. To compute the length of s , only note that each step of the iteration produces a polynomial whose length is no more than the square of the length of the previous polynomial. So we conclude that the length of s is no more than $l^{2^l} < 2^{2^{4l}}$.

Now suppose the lemma holds for each term s with $\Lambda(x, s) < d$, and suppose $\Lambda(x, t) = d$. Enumerate all the different terms $\lambda(s_0), \dots, \lambda(s_{m-1})$ in t such that $\lambda(s_i)$ is not in the scope of any λ for each $i < m$. Clearly $\Lambda(x, s_i) < n$ for each $i < m$. So by the inductive hypothesis there are less than $2^{8l_i^2 \log l_i}$ possible values for each $\lambda(s_i)$, where l_i is the length of s_i . Since $\sum_{i < m} l_i < l - 1$, there are no more than $2^{8(l-1)^2 \log l}$ possible values for t . Enumerate these possibilities as $\langle t_k : k < 2^{8(l-1)^2 \log l} \rangle$. In each t_k , $\lambda(s_i)$ is replaced by a term of the form sx^j with $j < l_i$. So t_k is a polynomial in x whose degree in x is less than $l - 2$. So there

are $2^{8(l-1)^2 \log l} \cdot 2^{8l \log l} \leq 2^{8l^2 \log l}$ possible values for $\lambda(t)$. The length of each t_k is bounded by $2^{2^{4(l-1)}}$, so the length of each possible value of $\lambda(t)$ is bounded by $2^{2^{4(l-1)}} \cdot l \cdot l^{2^l} < 2^{2^{4l}}$. \square

Lemma 3.5.4. *Let φ be a quantifier-free formula with length l . Assume x does not occur in the scope of any division symbol in φ . Then there is a quantifier-free formula φ' with length at most $2^{2^{O(l)}}$ such that φ' is simple in x and $T \vdash A(x) \rightarrow (\varphi \leftrightarrow \varphi')$.*

Proof. First we claim there is a quantifier-free formula φ^* with length at most $2^{2^{O(l)}}$ such that

- $T \vdash A(x) \rightarrow (\varphi \leftrightarrow \varphi^*)$,
- x does not occur in the scope of any division in φ^* ,
- $\Lambda(x, \varphi^*) = 0$.

The proof is essentially the same as the proof of Lemma 3.5.2, using Lemma 3.5.3 instead of Lemma 3.5.1.

Next we need to deal with atomic formulas of the form $D_n(p)$ in φ^* , as shown in Lemma 3.3.15. So p is a polynomial in x whose degree in x is less than l . So there are at most $2^{2^{O(l)}}$ possible values for $\lambda(p)$, the length of each of which is bounded by $2^{2^{O(l)}}$. So each $D_n(p)$ can be replaced by a disjunction whose length is less than $2^{2^{O(l)}}$. So the bound does not change.

The increase in length in transforming φ^* to a formula that is simple in x , as described in the proof of Lemma 3.3.16, can be absorbed in the bound $2^{2^{O(l)}}$. \square

Lemma 3.5.5. *Let φ be a quantifier-free formula with length l . Then there is a quantifier-free formula φ' with length at most $2_3^{O(l)}$ such that φ' is simple in x and T proves $A(x) \rightarrow (\varphi \leftrightarrow \varphi')$.*

Proof. Immediate by Lemma 3.5.2 and Lemma 3.5.4. \square

Lemma 3.5.6. *Each iteration of step 3 can be performed by an elementary function.*

Proof. It is straightforward to verify that the procedure implicit in Lemmas 3.5.5 runs in time polynomial in its output. As a result, step 3(a) is elementary. Step 3(b) is also clearly elementary. In fact, even though putting a formula in disjunctive normal form can result in exponentially many disjuncts, since each disjunct only involves atomic formulas from the original formula, the length of each disjunct is bounded in the length of the original formula.

After step 3(a), the main increase therefore comes from the handling of the cases in (c) and (d), each of which is easily seen to be elementary. Case (c) involves a call to a QE procedure for real closed ordered fields, with a $\forall\exists$ formula; case (d) involves calls to such a procedure, on existential formulas, across a number of disjuncts that is exponential in the length of the original formula. \square

Theorem 3.5.7. *There is a procedure for eliminating a single block of existential quantifiers in theory T in time $2^0_{O(l)}$, where l is the length of the original formula.*

Proof. Steps 1 and 2 are clearly elementary, after which the procedure performs an elementary operation for each quantifier over a power of two. As noted above, the number of such quantifiers can even be bounded by the number of predicates D_n and λ 's in the original formula. \square

Corollary 3.5.8. *There is a procedure for eliminating quantifiers in theory T that runs in time bounded by $O(l)$ iterations of the stack-of-twos function, where l is the length of the original formula.*

Proof. Put the formula in prenex form, and iteratively apply the previous theorem to eliminate each block of quantifiers. \square

3.6 Adding a predicate for the Fibonacci numbers

In this section we show that the theory of real closed ordered fields with a predicate for the Fibonacci numbers is recursively axiomatizable and hence is decidable. Moreover, the decision procedure described in the last section can be used to decide this theory.

The Fibonacci numbers are a sequence of natural numbers F_n defined by the recurrence relation

$$F_{n+2} = F_{n+1} + F_n$$

for $n > 0$ with $F_1 = F_2 = 1$. It is conventional to define $F_0 = 0$. The first few Fibonacci numbers are 0, 1, 1, 2, 3, 5, 8, 13, 21, ...

Let $\phi = \frac{1+\sqrt{5}}{2}$. Let A be a predicate for the multiplicative subgroup $\phi^{\mathbb{Z}} \subseteq \mathbb{R}^{>0}$. By Binet's Fibonacci number formula the n th Fibonacci number can be computed as follows:

$$F_n = \frac{\phi^n - \frac{(-1)^n}{\phi^n}}{\sqrt{5}}.$$

Therefore we can introduce a predicate F^* for the Fibonacci numbers with the following defining axiom:

$$F^*(x) \leftrightarrow \Gamma^e(x) \vee \Gamma^o(x), \tag{3.6.1}$$

where $\Gamma^e(x)$ and $\Gamma^o(x)$ are the formulas

$$\begin{aligned}\exists y (A(y) \wedge y \geq 1 \wedge \exists z (A(z) \wedge y = z^2) \wedge x = \frac{y - \frac{1}{y}}{\sqrt{5}}), \\ \exists y (A(y) \wedge y \geq 1 \wedge \exists z (A(z) \wedge y = \phi z^2) \wedge x = \frac{y + \frac{1}{y}}{\sqrt{5}}),\end{aligned}$$

respectively.

Now we give a complete axiomatization of the theory of real closed ordered fields with the distinguished Fibonacci numbers. We start with the theory of real closed ordered fields and a new predicate F for the Fibonacci numbers. Let $\Delta(x, y)$ abbreviates the formula

$$F(x) \wedge F(y) \wedge x < y \wedge \forall z (x < z < y \rightarrow \neg F(z)).$$

First we add the following axioms:

$$(A1) \quad F(x) \rightarrow x \geq 0;$$

$$(A2) \quad \Delta(0, 1) \wedge \Delta(1, 2);$$

$$(A3) \quad x > 2 \rightarrow (F(x) \leftrightarrow \exists y, z (\Delta(y, z) \wedge x = y + z \wedge \forall w (z < w < x \rightarrow \neg F(w))));$$

$$(A4) \quad z \geq 0 \rightarrow \exists x, y (\Delta(x, y) \wedge x \leq z < y).$$

Notice the following identities on the Fibonacci numbers:

$$\begin{aligned}F_{2n} &= F_n(2F_{n+1} - F_n), \\ F_{2n+1} &= F_{n+1}^2 + F_n^2.\end{aligned}$$

Generalizing these we let $\Sigma^{(e,o)}(x, y)$ and $\Sigma^{(o,e)}(x, y)$ be the formulas

$$\begin{aligned}\exists w, z (\Delta(w, z) \wedge x = w(2z - w) \wedge y = z^2 + w^2), \\ \exists w, z (\Delta(w, z) \wedge x = z^2 + w^2 \wedge y = z(2w + z)),\end{aligned}$$

respectively and obtain a new axiom:

$$(A5) \quad \Delta(x, y) \leftrightarrow \Sigma^{(e,o)}(x, y) \vee \Sigma^{(o,e)}(x, y).$$

This actually enables us to define the predicate A : let $\Theta^e(x)$ and $\Theta^o(x)$ be the formulas

$$\begin{aligned}\exists y, w, z (\Delta(w, z) \wedge y = w(2z - w) \wedge y = \frac{x - \frac{1}{x}}{\sqrt{5}}), \\ \exists y, w, z (\Delta(w, z) \wedge y = z^2 + w^2 \wedge y = \frac{x + \frac{1}{x}}{\sqrt{5}}),\end{aligned}$$

respectively, then

$$A(x) \leftrightarrow x > 0 \wedge (\Theta^e(x) \vee \Theta^o(x) \vee \Theta^e(\frac{1}{x}) \vee \Theta^o(\frac{1}{x})).$$

Now the idea is this. By the results in the last section there is a complete axiomatization of the theory of $(\mathbb{R}, \phi^{\mathbb{Z}})$ (see the last paragraph of Subsection 3.4), we may use F to define the predicate A and subsequently use A to define the predicate F^* via 3.6.1. Finally we throw in some axioms to guarantee that

- A picks out a suitable multiplicative subgroup and
- F^* and F are the same.

This will axiomatize a complete, hence decidable, theory with a predicate for the Fibonacci numbers.

Let K be an ordered field with a valuation v . Two nonzero elements a, b of K are in the same *Archimedean class* if $\frac{1}{n} < v(\frac{a}{b}) < n$ for some positive integer n . Let us say that a “local relation” is a relation that holds only among elements in the same Archimedean class and a “global relation” is a relation that is not local. Some classic identities on the Fibonacci numbers can, when generalized, control the behaviors of the predicate A , though only locally. For example, it is not hard to deduce the following:

$$\forall x, y (\Delta(x, y) \rightarrow y^2 - yx - x^2 = 1) \leftrightarrow \forall z (A(z) \leftrightarrow A(\phi z)).$$

But one should not think that such local identities are sufficient when elements in different Archimedean classes are involved. In fact an axiom is needed for the predicate A 's multiplicative closure:

$$(A6) \quad A(x) \wedge A(y) \rightarrow A(xy).$$

It is not hard to see that (A4) and (A6) together prove that

$$y > 0 \rightarrow \exists x (A(x) \wedge x \leq y < \phi x).$$

Finally we stipulate that

$$(A7) \quad F(x) \leftrightarrow F^*(x).$$

One may of course recast some of the axioms above into a form that is more explicit about the Fibonacci numbers. The calculations are easy but tedious. We shall not include them here. For example (A7) can be transformed into

$$F(x) \leftrightarrow \exists a, b, c, d (\Delta(a, b) \wedge \Delta(c, d) \\ \wedge ((x = a(2b - a) \wedge P_1(a, b, c, d)) \vee (x = a^2 + b^2 \wedge P_2(a, b, c, d))))),$$

where $P_1(a, b, c, d)$ and $P_2(a, b, c, d)$ are the polynomials

$$(5c^4(2d - c)^4 + 4c^2(2d - c)^2 - a^2(2b - a)^2) \\ (5(c^2 + d^2)^4 - 4(c^2 + d^2)^2 - a^2(2b - a)^2) = 0$$

and

$$(5c^4(2d - c)^4 + 4c^2(2d - c)^2 - 5(a^2 + b^2)c^2(2d - c)^2 + (a^2 + b^2 - 1)^2) \\ (5(c^2 + d^2)^4 - 4(c^2 + d^2)^2 - 5(a^2 + b^2)(c^2 + d^2)^2 + (a^2 + b^2 + 1)^2) = 0,$$

respectively. But these do not seem to be more natural than the ones that are listed above, even though it is rather curious why these polynomials are sufficient to determine the complete theory.

Chapter 4

Henselianity and the Denef-Pas language

Abstract: We prove that if an equicharacteristic valued field has a \mathbb{Z} -group as its value group and admits quantifier elimination in the main sort of the prototypical Denef-Pas style language then it is henselian. In fact the proof of this suggests that a reasonable class of Denef-Pas style languages is natural with respect to henselianity.¹

4.1 Introduction

Tarski's theorem says that the theory RCF of real closed fields, as formulated in the language \mathcal{L}_{OR} of ordered rings, admits quantifier elimination (QE). It is natural to ask whether any other ordered fields admit QE in \mathcal{L}_{OR} . There is a good answer to this:

Theorem 4.1.1 (Macintyre, McKenna, van den Dries). *Let K be an ordered field such that the theory of K in \mathcal{L}_{OR} admits QE. Then K is real closed.*

This is a prototypical example of a “converse QE” result; it shows that for the class of ordered fields, real-closedness is equivalent to QE.

There are analogous results in the class of valued fields. In the forward direction, the first result is due to Macintyre [37], who showed that the theory of p -adic fields, as formulated in the language \mathcal{L}_{Mac} , admits QE. In this case, we only have a partial converse:

Theorem 4.1.2 (Macintyre, McKenna, van den Dries). *Let K be a p -field such that the theory of K in \mathcal{L}_{Mac} admits QE. Then K is p -adically closed.*

The definition of a p -field is rather special: it is a substructure of a p -adically closed field (of p -rank 1) with respect to \mathcal{L}_{Mac} . Let K be a p -field and L a p -adically closed field such that K is an \mathcal{L}_{Mac} -substructure of L . The point is that, as L is henselian, each n th power predicate P_n defines a clopen subset of K in the valuation topology of K , which is essential to the proof of the theorem. This way to interpret each P_n is

¹This chapter has an unpolished longer version that contains more results on the Macintyre language of valued fields; see [54].

not very satisfying since an element in P_n may not be an n th power at all in K . Hence it is asked in [38] to extend the result to the class of valued fields where P_n is simply interpreted as the group of n th powers. In [54] this is established for a subclass of such structures.

In this paper we shall prove a converse QE theorem for a different kind of language for valued fields:

Theorem 4.1.3. *Let $S = \langle K, \overline{K}, \Gamma \cup \{\infty\}, v, \overline{\text{ac}} \rangle$ be a structure of the Denef-Pas style language \mathcal{L}_{RRP} such that*

1. K and \overline{K} are fields such that $\text{char } K = \text{char } \overline{K}$,
2. $v : K \rightarrow \Gamma$ is a valuation map and $\overline{\text{ac}} : K \rightarrow \overline{K}$ is an angular component map,
3. the value group Γ is a \mathbb{Z} -group,
4. the theory $\text{Th}(S)$ admits QE in the K -sort.

Then the valuation v is henselian.

This answers a question mentioned in [12]. This result also holds in slightly more general settings; see Remark 4.4.11 and Remark 4.4.13. The relevant definitions will be given in the next section.

We thank the referee for several helpful suggestions.

4.2 Preliminaries

In this paper all valued fields are equicharacteristic and all valuation rings are proper subrings. We use \mathcal{O} , \mathcal{O}_1 , etc. and \mathcal{M} , \mathcal{M}_1 , etc. to denote valuation rings and their maximal ideals, respectively. Valuation maps are denoted by v , v_1 , etc. If v is a valuation of K then vK , \overline{K} stand for the corresponding value group and residue field, respectively.

Next we describe the Denef-Pas style language for valued fields.

Definition 4.2.1. Let K be a valued field and \overline{K} its residue field. An *angular component map* is a function $\overline{\text{ac}} : K \rightarrow \overline{K}$ such that

1. $\overline{\text{ac}} 0 = 0$,
2. the restriction $\overline{\text{ac}} \upharpoonright K^\times$ is a group homomorphism $K^\times \rightarrow \overline{K}^\times$,
3. the restriction $\overline{\text{ac}} \upharpoonright (\mathcal{O} \setminus \mathcal{M})$ is the projection map, that is, $\overline{\text{ac}} u = u + \mathcal{M}$ for all $u \in \mathcal{O} \setminus \mathcal{M}$.

The template for Denef-Pas style language has three sorts: the field sort, the residue field sort, and the value group sort. These are usually denoted by K , \overline{K} , and Γ , respectively. Sometimes we shall refer to the K -sort as the “main sort”. The K -sort and \overline{K} -sort use the language \mathcal{L}_R of rings. The Γ -sort uses the language \mathcal{L}_{OG} of ordered groups, $\{+, <, 0\}$, and an additional symbol ∞ that designates the top element in the ordering. There are two cross-sort function symbols: $v : K \rightarrow \Gamma$, which stands for the valuation, and $\overline{ac} : K \rightarrow \overline{K}$, which stands for an angular component map.

Any language that expands this template is a Denef-Pas language. A prototypical example is the language \mathcal{L}_{RRP} used in [41], in which the field sort and the residue field sort use the language \mathcal{L}_R and the Γ -sort uses the language $\mathcal{L}_{Pr\infty} = \mathcal{L}_{Pr} \cup \{\infty\}$, where \mathcal{L}_{Pr} is the Presburger language $\{+, <, 0, 1\} \cup \{D_n : n > 1\}$. Let $S = \langle K, \overline{K}, \Gamma \cup \{\infty\}, v, \overline{ac} \rangle$ be a structure of \mathcal{L}_{RRP} . One of the main results of [41] is that if K is henselian and both K and \overline{K} are of characteristic 0 then $\text{Th}(S)$ admits QE in the K -sort; that is, for every formula φ in \mathcal{L}_{RRP} there is a formula φ^* in \mathcal{L}_{RRP} that does not contain K -quantifiers such that $S \models \varphi \leftrightarrow \varphi^*$. Hence Theorem 4.1.3 contains a converse of this result with respect to henselianity under the additional assumption that Γ is a \mathbb{Z} -group.

The following notions are formulated for any Denef-Pas language \mathcal{L} , where we use \mathcal{L}_K , $\mathcal{L}_{\overline{K}}$, and $\mathcal{L}_{\Gamma\infty}$ to denote the languages used by the three sorts.

Definition 4.2.2. A formula φ in \mathcal{L} is *simple* if φ does not contain any K -quantifiers.

Definition 4.2.3. A formula φ in $\mathcal{L}_K \cup \mathcal{L}_{\Gamma\infty}$ is a Γ -*formula* if it does not contain K -quantifiers or atomic formulas in \mathcal{L}_K . Similarly a formula φ in $\mathcal{L}_K \cup \mathcal{L}_{\overline{K}}$ is a \overline{K} -*formula* if it does not contain K -quantifiers or atomic formulas in \mathcal{L}_K .

Note that Γ -formulas and \overline{K} -formulas may contain \mathcal{L}_K -terms. We shall simplify our terminology for these formulas. For example, a literal Γ -formula shall be called a “ Γ -literal”. Similarly for \overline{K} -formulas.

4.3 Overview of the proof

The proof relies on the approximation technique devised in [38]. In general this technique consists of the following steps. Let \mathcal{L} be a language for valued fields in which henselianity is first-order expressible. The main sort of \mathcal{L} is the field sort. Let (K, v) be a valued field such that $\text{Th}(K)$ admits QE in the main sort of \mathcal{L} , where $\text{Th}(K)$ denotes the theory of K as a structure of \mathcal{L} . Let \mathcal{O}, \mathcal{M} be its valuation ring and maximal ideal.

- **Step 1.** Show that, except equations in the field, all formulas without quantifier ranging over the main sort define open sets in (the product of) the valuation topology. Note that, for each formula $\varphi(X)$,

the assertion that it defines an open set can be expressed by a first-order sentence:

$$\forall X (\varphi(X) \rightarrow \exists Y (v(Y) > v(X) \wedge \forall Z (v(Z) > v(Y) \rightarrow \varphi(X + Z))).$$

- **Step 2.** Let $F(X, \bar{a}) \in \mathcal{O}[X]$ be a monic polynomial of degree n , where \bar{a} are the coefficients. Suppose for contradiction that $F(X, \bar{a})$ is a counterexample to a version of Hensel's Lemma: there is an $s \in \mathcal{O}$ such that $F(s, \bar{a}) \in \mathcal{M}$ and $F'(s, \bar{a}) \notin \mathcal{M}$ but $F(X, \bar{a})$ has no root in K . We may assume that $F(X, \bar{a})$ is irreducible over K . Let φ be the formula that defines the nonempty set of the tuples of coefficients of all such counterexamples. By assumption, φ is equivalent to a formula that is quantifier-free in the main sort. Without loss of generality we may assume that φ is in disjunctive normal form. Using the fact that the Vandermonde matrix of $F(X, \bar{a})$ is invertible, we may construct polynomials $F_1(\bar{Y}), \dots, F_n(\bar{Y}) \in \mathcal{O}[\bar{Y}]$, where \bar{Y} is a tuple of variables Y_1, \dots, Y_n , such that

1. they are algebraically independent over K ,
2. $F(X, F_1(\bar{b}), \dots, F_n(\bar{b}))$ has no root in K for every $\bar{b} = (b_1, \dots, b_n) \in K^n$ with $b_i \neq 0$ for some $i > 1$,
3. $F(X, F_1(0, 1, 0, \dots), \dots, F_n(0, 1, 0, \dots)) = F(X, \bar{a})$.

By continuity of polynomial maps, there is an open neighborhood U of $(0, 1, 0, \dots)$ in the product topology on K^n such that $(F_1(\bar{b}), \dots, F_n(\bar{b}))$ satisfies φ for every $\bar{b} \in U$. Since U is not contained in any proper Zariski closed subset of K^n , there must be a disjunct φ_0 of φ that lacks equational conditions and hence, by Step 1, defines a nonempty open subset of K^n . Without loss of generality $\bar{a} \in \varphi_0(K^n)$. For details see [38, Theorem 1, 4].

- **Step 3.** If K is dense in its henselization K^h then the approximation can be carried out as follows: Choose a root $r \in K^h$ of $F(X, \bar{a})$ and write

$$F(X, \bar{a}) = (X - r)F^*(X, \bar{b}),$$

where $\bar{b} \in K^h$ are the coefficients of F^* . Let $U \subseteq \varphi_0(K^n)$ be an open neighborhood of \bar{a} , where φ_0 is as in Step 2. Now we can choose $r', \bar{b}' \in K$ that are arbitrarily close to r, \bar{b} with respect to the valuation. Write

$$F(X, \bar{a}') = (X - r')F^*(X, \bar{b}').$$

So \bar{a}' are arbitrarily close to \bar{a} and hence $\bar{a}' \in U$. But then $F(X, \bar{a}')$ have a root in K , contradicting

the choice of U .

- **Step 4.** However, in general K is not dense in its henselization. The solution to this in [38] is rather specialized. Dickmann [17] uses a more general method to get around this problem. Using the Omitting Types Theorem, another valued field (L, w) may be constructed such that (L, w) is elementarily equivalent to (K, v) with respect to \mathcal{L} and w is of rank 1 (that is, wL is a subgroup of the additive group of \mathbb{R} with the canonical ordering). It is well-known that if the valuation w for L is of rank 1 then L is dense in its henselization; see the discussion in [23, p. 53]. Hence the argument outlined above can be used to show that (L, w) is henselian. Consequently (K, v) is henselian.

Note that Step 2 can always be implemented for any valued field that is not henselian. So the bulk of the work in the sequel will concentrate on Step 1, 3, and 4.

4.4 Henselianity and Denef-Pas style languages

We shall prove Theorem 4.1.3 in this section. The proof of this theorem can be adapted for other Denef-Pas style languages as well, provided that the value group satisfies certain mild conditions; see Remark 4.4.11.

Throughout this section let $S = \langle K, \overline{K}, \Gamma \cup \{\infty\}, v, \overline{ac} \rangle$ be a structure of \mathcal{L}_{RRP} that satisfies the assumptions of Theorem 4.1.3. We shall work in S .

In this section the following notational conventions are adopted. We use X, Y , etc. for K -sort variables, M, N , etc. for Γ -sort variables, and Ξ, Λ , etc. for \overline{K} -sort variables. The lowercase of these letters stands for closed terms or elements in the corresponding sorts. Unless indicated otherwise, all these letters stand for tuples of variables whenever they appear in a formula. We use $\text{lh } X$ to denote the length of X . Let \mathbb{Z} be the subring generated by 1 in K and \mathbb{Z}_Γ the subgroup generated by 1 in Γ .

Remark 4.4.1. The theory of \mathbb{Z} -groups with a top element in $\mathcal{L}_{\text{Pr}\infty}$ admits QE. This follows from a straightforward generalization of [41, Lemma 5.4, 5.5] to \mathbb{Z} -groups.

The following lemma is slightly more general than [41, Lemma 5.3].

Lemma 4.4.2. *Let φ be a simple formula in \mathcal{L}_{RRP} . Then φ is equivalent to a formula of the form*

$$\bigvee_i (\sigma_i \wedge \chi_i \wedge \theta_i)$$

where σ_i is a quantifier-free formula in \mathcal{L}_K , χ_i a \overline{K} -formula, and θ_i a Γ -formula.

Proof. We can write φ in its prenex normal form $Q_1 \dots Q_k \psi$ where each Q_j is either a Γ -quantifier or a \overline{K} -quantifier and ψ is a quantifier-free formula. We proceed by induction on the number k of quantifiers.

If $k = 0$ then φ is quantifier-free. Since there are no symbols in \mathcal{L}_{RRP} relating the \overline{K} -sort and the Γ -sort, φ can be written in its disjunctive normal form

$$\bigvee_i (\sigma_i \wedge \chi_i \wedge \theta_i)$$

where σ_i is a conjunction of literals in \mathcal{L}_K , χ_i a conjunction of \overline{K} -literals, and θ_i a conjunction of Γ -literals.

This proves the base case.

Suppose now $k = l + 1$ and Q_1 is $\exists N$. So by the inductive hypothesis φ can be written in the form

$$Q_1 \bigvee_i (\sigma'_i \wedge \chi'_i \wedge \theta'_i)$$

where σ'_i is a quantifier-free formula in \mathcal{L}_K , χ'_i a \overline{K} -formula, and θ'_i a Γ -formula. Now we can simply push the quantifier in and write φ as

$$\bigvee_i (\sigma'_i \wedge \chi'_i \wedge \exists N \theta'_i).$$

If Q_1 is $\forall N$ then we can rewrite $\bigvee_i (\sigma'_i \wedge \chi'_i \wedge \theta'_i)$ in its conjunctive normal form and then push the quantifier in. The other two cases of Q_1 being $\exists \Xi$ or $\forall \Xi$ are treated in the same way. \square

Simple formulas play an important role in this section. Let φ be a simple formula. By Lemma 4.4.2, φ can be written as a disjunction of conjunctions of formulas of the following forms:

- Type I: $F(X) \square 0$, where \square is either $=$ or \neq and $F(X) \in \mathbb{Z}[X]$.
- Type II: Γ -formulas. Suppose that $F_i(X) \in \mathbb{Z}[X]$ run through all the distinct polynomials that appear in a formula of this type in the form $vF_i(X)$. For every i , since the formulas $vF_i(X) = \infty$ and $vF_i(X) \neq \infty$ are equivalent to the formulas $F_i(X) = 0$ and $F_i(X) \neq 0$ respectively and the latter ones can be assimilated into Type I, we may assume that $vF_i(X) = \infty$ and $vF_i(X) \neq \infty$ do not occur in φ and $F_i(X) \neq 0$ is a conjunct in each disjunct of φ in which $F_i(X)$ appears in a formula of this type.
- Type III: \overline{K} -formulas. Suppose that $F_i(X) \in \mathbb{Z}[X]$ run through all the distinct polynomials that appear in a formula of this type in the form $\overline{ac}F_i(X)$. Similar to Type II, for every i , since the formulas $\overline{ac}F_i(X) = 0$ and $\overline{ac}F_i(X) \neq 0$ are equivalent to the formulas $F_i(X) = 0$ and $F_i(X) \neq 0$, we may assume that $\overline{ac}F_i(X) = 0$ and $\overline{ac}F_i(X) \neq 0$ do not occur in φ and $F_i(X) \neq 0$ is a conjunct in each disjunct of φ in which $F_i(X)$ appears in a formula of this type.

4.4.1 Step 1: Open sets

Since Step 2, 3, and 4 in Section 4.3 do not involve formulas that contain free \overline{K} -variables or free Γ -variables, we may limit our attention to such formulas of Type I, II, and III. We shall show that such formulas, except equalities in the K -sort, define open sets in the corresponding product of the valuation topology. This takes care of Step 1.

Since polynomials are continuous maps with respect to the valuation topology, it is clear that disequalities in the K -sort define open sets.

Lemma 4.4.3. *Let $\varphi(X)$ be a formula of Type II. Then φ defines an open set.*

Proof. First note that, for $m \in \Gamma$, sets of the forms $\{x : v(x) \square m\}$, where \square is one of the symbols $=, \neq, <, \geq$, are all open in the valuation topology. See [23, Remark 2.3.3].

Let $F_i(X) \in \mathbb{Z}[X]$ run through all the distinct polynomials that appear in $\varphi(X)$ in the form $vF_i(X)$. Let $\varphi^*(M)$ be the formula obtained from $\varphi(X)$ by replacing each $vF_i(X)$ with a new variable M_i . Let B be the set

$$\{\langle m_1, \dots, m_d \rangle \in \Gamma^d : S \models \varphi^*(m_1, \dots, m_d)\},$$

where $d = \text{lh } M$. For each $m = \langle m_1, \dots, m_d \rangle \in \Gamma^d$ let

$$A_m = \left\{ x \in K^e : \bigwedge_{i=1}^d vF_i(x) = m_i \right\},$$

where $e = \text{lh } X$. Since polynomial maps are continuous, each A_m is open in the valuation topology. So $\varphi(K^e) = \bigcup_{m \in B} A_m$ is open. \square

Let \mathcal{O}, \mathcal{M} be the valuation ring and its maximal ideal that correspond to v . The following lemma establishes a crucial relation between the valuation and the angular component map.

Lemma 4.4.4. *For nonzero $x, y \in K$ with $v(x) = v(y) = m \in \Gamma$, $\overline{a}x = \overline{a}y$ if and only if $v(x - y) > m$.*

Proof. If $x = y$ then the lemma is trivial. So we assume further that $x \neq y$.

For the ‘‘only if’’ direction, suppose for contradiction that $\overline{a}x = \overline{a}y$ but $v(x - y) = m$. So $(x - y)/x$ is

a unit. So

$$\begin{aligned}
\overline{\text{ac}} \frac{x-y}{x} &= \left(1 - \frac{y}{x}\right) + \mathcal{M} \\
&= (1 + \mathcal{M}) - \left(\frac{y}{x} + \mathcal{M}\right) \\
&= (1 + \mathcal{M}) - \overline{\text{ac}} \frac{y}{x} \\
&= (1 + \mathcal{M}) - \frac{\overline{\text{ac}} y}{\overline{\text{ac}} x} \\
&= 0.
\end{aligned}$$

So $(x - y)/x = 0$, so $x = y$, contradiction.

For the “if” direction, suppose for contradiction that $v(x - y) > m$ but $\overline{\text{ac}} x \neq \overline{\text{ac}} y$. If $m = 0$, that is, x and y are units in the valuation ring, then

$$x + \mathcal{M} = \overline{\text{ac}} x \neq \overline{\text{ac}} y = y + \mathcal{M}.$$

So $x - y$ is a unit in the valuation ring, that is, $v(x - y) = 0$, contradiction. In general we may consider $1 - y/x$: since $v(1 - y/x) > 0$ and y/x is a unit, we get $\overline{\text{ac}} 1 = \overline{\text{ac}}(y/x)$ by the previous two sentences, so $\overline{\text{ac}} x = \overline{\text{ac}} y$. \square

Lemma 4.4.5. *Let $\lambda \in \overline{K}^\times$ and $F(X) \in \mathbb{Z}[X]$. The set*

$$A_\lambda = \{x \in K^e : \overline{\text{ac}} F(x) = \lambda\}$$

is open, where $e = \text{lh } X$.

Proof. For any $x \in K^e$ such that $F(x) \neq 0$ we consider the open set

$$U = \{F(x) + z : z \in K \text{ and } v(z) > vF(x)\}.$$

Since F is continuous, there is an open neighborhood U_x of x such that $F(U_x) \subseteq U$. Since $vF(y) = vF(x)$ and $v(F(y) - F(x)) > vF(x)$ for every $y \in U_x$, by Lemma 4.4.4, $\overline{\text{ac}} F(y) = \overline{\text{ac}} F(x)$. So $A_\lambda = \bigcup_{x \in A_\lambda} U_x$ is open. \square

Lemma 4.4.6. *Let $\varphi(X)$ be a formula of Type III. Then φ defines an open set.*

Proof. Let $F_i(X) \in \mathbb{Z}[X]$ run through all the distinct polynomials that appear in $\varphi(X)$ in the form $\overline{\text{ac}} F_i(X)$. Let $\varphi^*(\Lambda)$ be the formula obtained from $\varphi(X)$ by replacing each $\overline{\text{ac}} F_i(X)$ with a new variable Λ_i . Let B be

the set

$$\left\{ \langle \lambda_1, \dots, \lambda_d \rangle \in (\overline{K}^\times)^d : S \models \varphi^*(\lambda_1, \dots, \lambda_d) \right\},$$

where $d = \text{lh } \Lambda$. For each $\lambda = \langle \lambda_1, \dots, \lambda_d \rangle \in (\overline{K}^\times)^d$ let

$$A_\lambda = \left\{ x \in K^e : \bigwedge_{i=1}^d \overline{\text{ac}} F_i(x) = \lambda_i \right\},$$

where $e = \text{lh } X$. By Lemma 4.4.5 each A_λ is open. So $\varphi(K^e) = \bigcup_{\lambda \in B} A_\lambda$ is open. \square

4.4.2 Step 3 and 4: Omitting a type

If K is dense in its henselization then, combining the argument in Step 3 and the results in the last section, we see that the conclusion of Theorem 4.1.3 holds. If K is not dense in its henselization then we need to carry out Step 4. Thus, we shall show:

Theorem 4.4.7. *There is a structure $S_1 = \langle K_1, \overline{K}_1, \Gamma_1 \cup \{\infty\}, v_1, \overline{\text{ac}}_1 \rangle$ of \mathcal{L}_{RRP} such that $S_1 \equiv S$ and v_1 is of rank 1.*

For the rest of this section let X, Y be two single variables. For $r, t \in \mathcal{O}$ we say that they are *comparable*, written as $r \asymp t$, if there is a natural number n such that either $v(r^n) \leq v(t) \leq v(r^{n+1})$ or $v(t^n) \leq v(r) \leq v(t^{n+1})$. They are *incomparable* if they are not comparable. We write $r \ll t$ if r, t are incomparable and $v(r) < v(t)$.

By the Omitting Types Theorem, Theorem 4.4.7 may be proved by omitting the 2-type

$$\Phi(X, Y) = \{0 < v(X^l) < v(Y) \wedge Y \neq 0 : l \geq 1\}.$$

Thus it suffices to show that this type is not isolated modulo $\text{Th}(S)$. To that end, suppose for contradiction that there is a formula $\pi(X, Y)$ such that the sentence $\exists X, Y \pi(X, Y)$ is in $\text{Th}(S)$ and $\pi(X, Y) \vdash \Phi(X, Y)$ modulo $\text{Th}(S)$. Since $\text{Th}(S)$ admits QE in the main sort, by Lemma 4.4.2, $\pi(X, Y)$ is equivalent to a disjunction of conjunctions of formulas of Type I, II, and III. Without loss of generality we may assume that $\pi(X, Y)$ is just a conjunction of formulas of those three types.

The following lemma shows that in fact $\pi(X, Y)$ does not contain equations in the K -sort.

Lemma 4.4.8. *For any nonzero $r, t \in \mathcal{M}$ with $r \ll t$ and any nonzero polynomial $F(X, Y) \in \mathbb{Z}[X, Y]$, $F(r, t) \neq 0$.*

Proof. Suppose for contradiction $F(r, t) = 0$. Write $F(X, Y)$ as

$$Y^d(F_l(X)Y^l + \dots + F_0(X)), \quad (4.4.1)$$

where $F_0(X), \dots, F_l(X) \in \mathbb{Z}[X]$. If $F(X, Y)$ is a monomial in Y then it can be written as

$$(e_k X^k + \dots + e_0) Y^i \quad (4.4.2)$$

for some $i \geq 0$, where $e_0, \dots, e_k \in \mathbb{Z}$. But no two summands in $e_k r^k + \dots + e_0$ have the same valuation, for otherwise we would have $v(r) = 0$. Hence $v(e_k r^k + \dots + e_0) < \infty$, contradiction.

So we may assume that $F(X, Y)$ has at least two nonzero monomial summands. Now for some $i > j \geq 0$ we have $v(F_i(r)t^i) = v(F_j(r)t^j)$. So

$$v(t^{i-j}) = v(F_j(r)/F_i(r)).$$

But again, in either $F_i(r)$ or $F_j(r)$, no two summands have the same valuation, so either $F_j(r)/F_i(r) \ll r$ or $F_j(r)/F_i(r) \asymp r$. So either $t \ll r$ or $t \asymp r$, contradiction again. \square

Remark 4.4.9. This lemma is well-known. It is a corollary of the fundamental dimension inequality in the theory of valued fields; see [23, Theorem 3.4.3]. We prefer to give an elementary proof here to make clear that its failure in valued fields of mixed characteristics is the main reason that Theorem 4.1.3 has not been extended to such fields in general. On the other hand, the above lemma clearly may be applied to the case $r = \text{char}(\overline{K}) > 0$. So Theorem 4.1.3 does hold for a particular subclass of valued fields of mixed characteristics, namely tight valued fields; see Remark 4.4.13.

Lemma 4.4.10. *Let $\varphi(X, Y)$ be a conjunction of formulas of Type II, where X, Y are the only free variables. Let $x, y \in \mathcal{M}$ be nonzero such that $x \ll y$ and $S \models \varphi(x, y)$. Then for every natural number k there is an $m \in \Gamma$ with $v(x^k) < m < v(x^l)$ for some $l > k$ such that for every $t \in \mathcal{M}$ with $v(t) = m$ we have $S \models \varphi(x, t)$.*

Proof. Let $F_i(X, Y) \in \mathbb{Z}[X, Y]$ run through all the distinct polynomials that appear in $\varphi(X, Y)$ in the form $vF_i(X, Y)$. We may assume that each $F_i(X, Y)$ is written in the form (4.4.1) in Lemma 4.4.8. It is not hard to see that if we choose a $k_0 > 0$ that is larger than the sum of all the exponents of X that appear in all the polynomials $F_i(X, Y)$, then for each $F_i(X, Y)$ there are integers e_i, d_i with $e_i < k_0$ such that for each $t \in \mathcal{M}$, if $v(t) > v(x^{k_0})$ then

$$vF_i(x, t) = v(x^{e_i} t^{d_i}). \quad (4.4.3)$$

Now substituting $v(x^{e_i} t^{d_i})$ for $vF_i(x, t)$ in $\varphi(x, t)$ and then substituting two variables N_1, N_2 for $v(x), v(t)$

respectively in the resulting formula we obtain an $\mathcal{L}_{\Gamma\infty}$ -formula $\varphi^*(N_1, N_2)$ from $\varphi(x, t)$ such that for all $t \in \mathcal{M}$ with $v(t) > v(x^{k_0})$

$$S \models \varphi(x, t) \text{ if and only if } \Gamma \cup \{\infty\} \models \varphi^*(v(x), v(t)). \quad (4.4.4)$$

In particular we have $\Gamma \cup \{\infty\} \models \varphi^*(v(x), v(y))$. Let $v(x) = n$. Let $\Gamma(n)$ be the smallest \mathbb{Z} -group generated by n in Γ . It is easy to see that the set $\{kn : k \in \mathbb{N}\}$ is cofinal in $\Gamma(n)$. By Remark 4.4.1, $\Gamma(n) \cup \{\infty\}$ is an elementary substructure of $\Gamma \cup \{\infty\}$. So for every natural number $k \geq k_0$ we have

$$\Gamma(n) \cup \{\infty\} \models \exists N (kn < N < \infty \wedge \varphi^*(n, N)).$$

So for some $m \in \Gamma(n)$ and some $l > k$ we have

$$\Gamma(n) \cup \{\infty\} \models kn < m < ln \wedge \varphi^*(n, m).$$

So for every $t \in \mathcal{M}$ with $v(t) = m$ we have $\Gamma \cup \{\infty\} \models \varphi^*(n, v(t))$. By (4.4.4) we have $S \models \varphi(x, t)$, as desired. \square

Remark 4.4.11. A close examination of the proof of Lemma 4.4.3 and Lemma 4.4.6 shows that, regardless of what languages the group Γ and the field \overline{K} use and what additional structures they have, formulas of Type II and III always define open sets. Therefore Lemma 4.4.10 is actually the only place where we need to use some special properties that hold in \mathbb{Z} -groups, namely

1. for any element n in the Γ -sort the set $\{kn : k \in \mathbb{N}\}$ is cofinal in the submodel generated by n ;
2. the theory of the Γ -sort in $\mathcal{L}_{\Gamma\infty}$ is model-complete.

So our converse QE result holds for any group Γ , any field \overline{K} , and any languages $\mathcal{L}_{\Gamma\infty}, \mathcal{L}_{\overline{K}}$, provided that these two properties are satisfied.

Lemma 4.4.12. *Let $\varphi(X, Y)$ be a formula of Type III, where X, Y are the only free variables. Let $x, y \in \mathcal{M}$ be nonzero such that $x \ll y$ and $S \models \varphi(x, y)$. For every $t \in \mathcal{M}$, if $v(t)$ is sufficiently large and $\overline{ac}t = \overline{ac}y$, then $S \models \varphi(x, t)$.*

Proof. Let $F_i(X, Y) \in \mathbb{Z}[X, Y]$ run through all the distinct polynomials that appear in $\varphi(X, Y)$ in the form $\overline{ac}F_i(X, Y)$. As in the previous lemma we choose a $k_0 > 0$ that is larger than the sum of all the exponents of X that appear in all the polynomials $F_i(X, Y)$. So for each $t \in \mathcal{M}$, if $v(t) > v(x^{k_0})$ then the equation (4.4.3)

in Lemma 4.4.10 holds for each $F_i(X, Y)$. For such a $t \in \mathcal{M}$, if $F_i(X, Y)$ is written in the form (4.4.1) in Lemma 4.4.8, then we have

$$v(F_l(x)t^l + \dots + F_0(x)) = vF_0(x)$$

and

$$v(F_l(x)t^l + \dots + F_1(x)t) > vF_0(x)$$

if $l > 0$. Let $F_0(X)$ be written as $X^b(e_j X^j + \dots + e_0)$, with $e_0, \dots, e_j \in \mathbb{Z}$ and e_0 nonzero. So by Lemma 4.4.4 we have

$$\overline{\text{ac}} F_i(x, t) = (\overline{\text{ac}} t)^d \cdot \overline{\text{ac}} F_0(x) = (\overline{\text{ac}} t)^d \cdot (\overline{\text{ac}} x)^b \cdot \overline{\text{ac}} e_0.$$

In particular, since $x \ll y$, we have

$$\overline{\text{ac}} F_i(x, y) = (\overline{\text{ac}} y)^d \cdot (\overline{\text{ac}} x)^b \cdot \overline{\text{ac}} e_0.$$

Now if $\overline{\text{ac}} t = \overline{\text{ac}} y$ then we have

$$\overline{\text{ac}} F_i(x, t) = (\overline{\text{ac}} t)^d \cdot (\overline{\text{ac}} x)^b \cdot \overline{\text{ac}} e_0 = (\overline{\text{ac}} y)^d \cdot (\overline{\text{ac}} x)^b \cdot \overline{\text{ac}} e_0 = \overline{\text{ac}} F_i(x, y).$$

So clearly $S \models \varphi(x, t)$, as desired. □

Proof of Theorem 4.4.7. Let $x \ll y$ be such that $S \models \pi(x, y)$. We shall show that there is a $t \in \mathcal{M}$ with $x \asymp t$ such that $S \models \pi(x, t)$. This shows that the type $\Phi(X, Y)$ is not isolated by $\pi(X, Y)$ modulo $\text{Th}(S)$.

By Lemma 4.4.8, $\pi(X, Y)$ cannot contain equalities in the K -sort. Clearly, for sufficiently large k , if $t \in \mathcal{M}$ is nonzero and $v(t) \geq v(x^k)$ then the pair (x, t) satisfies the disequalities in the K -sort that appear in $\pi(X, Y)$. Finally, by Lemma 4.4.10 and 4.4.12 we can choose a sufficiently large k and a $t \in \mathcal{M}$ with $v(x^k) < v(t) < v(x^l)$ for some $l > k$ and $\overline{\text{ac}} t = \overline{\text{ac}} y$ such that $S \models \pi(x, t)$, as desired. □

Remark 4.4.13. If we replace the first assumption of Theorem 4.1.3 with $\text{char } K = 0$ and $\text{char}(\overline{K}) = p > 0$ and then add another assumption that the valued field is *tight*, that is, $v(p)$ is contained in the smallest nontrivial convex subgroup \mathbb{Z}_Γ of Γ , then the argument above can be quite easily adapted to show that the theorem still holds. To see this, first note that for some $n \in \mathbb{Z}_\Gamma$ the sentence $v(p) = n$ is in $\text{Th}(S)$. Next, we leave Step 1, 2, and 3 unchanged. For Step 4, it is enough to show that the 1-type

$$\Phi(Y) = \{0 < v(p^l) < v(Y) \wedge Y \neq 0 : l \geq 1\}.$$

is not isolated modulo $\text{Th}(S)$. To that end, suppose for contradiction that there is a formula $\pi(Y)$ such that the sentence $\exists Y \pi(Y)$ is in $\text{Th}(S)$ and $\pi(Y) \vdash \Phi(Y)$ modulo $\text{Th}(S)$. By QE in the K -sort, $\pi(Y)$ is equivalent to, without loss of generality, a conjunction of formulas of Type I, II, and III with only one free K -sort variable Y . Clearly Lemma 4.4.8 holds with $r = p$ and Lemma 4.4.10, Lemma 4.4.12 hold with $x = p$. Now the contradiction is that we can find an element $t \in \mathcal{M}$ such that $v(t) \in \mathbb{Z}_\Gamma$ and $S \models \pi(t)$. Finally, observe that the tightness condition is necessary for Step 4, since, otherwise, the sentences $v(p) > 1, v(p) > 2, \dots$ are all in $\text{Th}(S)$.

Chapter 5

Grothendieck homomorphisms in algebraically closed valued fields

Abstract: We give a presentation of the construction of motivic integration, that is, a homomorphism between Grothendieck semigroups that are associated with a first-order theory of algebraically closed valued fields, in the fundamental work of Hrushovski and Kazhdan [34]. We limit our attention to a simple major subclass of V -minimal theories of the form ACVF_S^0 , that is, the theory of algebraically closed valued fields of pure characteristic 0 expanded by a (VF, Γ) -generated substructure S in the language \mathcal{L}_{RV} . The main advantage of this subclass is the presence of syntax. It enables us to simplify the arguments with many new technical details while following the major steps of the Hrushovski-Kazhdan theory.

5.1 Introduction

The theory of motivic integration in valued fields has been progressing rapidly since its first introduction by Kontsevich. Early developments by Denef and Loeser et al. have yielded many important results in many directions. The reader is referred to [30] for an excellent introduction to the construction of motivic measure. There have been different approaches to motivic integration. The comprehensive study in Cluckers-Loeser [14] has successfully united the major ones on a general foundation. Their construction may be applied in general to the field of formal Laurent series over a field of characteristic 0 but heavily relies on the Cell Decomposition Theorem of Denef-Pas [16, 41], which is only achieved for valued fields of characteristic 0 that are equipped with an angular component map. However, an angular component map is not guaranteed to exist for just any valued field, for example, algebraically closed valued fields. The Hrushovski-Kazhdan integration theory [34] is a major development that does not require the presence of an angular component map and hence is of great foundational importance. Its basic objects of study are models of V -minimal theories. This class of theories encompasses a wide range of first-order expansions of the theory of algebraically closed valued fields of pure characteristic 0 that have been shown to have nice geometrical behaviors. Moreover, by compactness, when integrating a definable object, the theory may be applied to valued fields with large positive residue characteristics.

In this paper, following the major steps of the construction of Grothendieck homomorphisms, that is, homomorphisms between Grothendieck semigroups, but supplying new technical lemmas, we give a presentation of the materials in the first eight chapters of [34]. In doing so, we limit our attention to a simple major subclass of V -minimal theories, namely the theory of algebraically closed valued fields of pure characteristic 0 in the language \mathcal{L}_{RV} with parameters from the field sort and the (imaginary) value group sort allowed. The main technical differences from the original construction are all results of this restriction. Our principal aim is to reconstruct the Grothendieck homomorphisms in [34, Theorem 8.8]. Due to technical reason, which we shall describe below, our reconstructed isomorphisms are actually quotients of the isomorphisms in [34, Theorem 8.8]. Other similar homomorphisms that involve differential calculus are completely left out. They will be presented in a sequel to this paper that is devoted to the study of Fourier transform in the Hrushovski-Kazhdan theory and its extension to the adelic setting via Weispfenning's fundamental work on the model theory of boolean products [50].

5.1.1 Outline of the construction

The method of the Hrushovski-Kazhdan integration theory is based on a fine analysis of definable subsets up to definable bijections in a Basarab-Kuhlmann style language \mathcal{L}_{RV} for valued fields. This language has two sorts: the VF-sort and the RV-sort. One of the main features of \mathcal{L}_{RV} is that the residue field and the value group are wrapped together in one sort RV; see Section 5.2 for details. Let (K, val) be a valued field and \mathcal{O} , \mathcal{M} , \overline{K} the corresponding valuation ring, its maximal ideal, and the residue field. Let $\text{RV}(K) = K^\times / (1 + \mathcal{M})$ and $\text{rv} : K^\times \rightarrow \text{RV}(K)$ the quotient map. Note that, for each $a \in K$, val is constant on the subset $a + a\mathcal{M}$ and hence there is a naturally induced map v_{rv} from $\text{RV}(K)$ onto the value group Γ . The situation is illustrated in the following commutative diagram

$$\begin{array}{ccccc}
 \mathcal{O} \setminus \mathcal{M} & \hookrightarrow & K^\times & & \\
 \text{quotient} \downarrow & & \downarrow \text{rv} & \searrow \text{val} & \\
 \overline{K}^\times & \hookrightarrow & \text{RV}(K) & \xrightarrow{v_{\text{rv}}} & \Gamma
 \end{array}$$

where the bottom sequence is exact. Note that the existence of an angular component $\overline{ac} : K^\times \rightarrow \overline{K}^\times$ is equivalent to the existence of a group homomorphism from $\text{RV}(K)$ onto \overline{K}^\times in the diagram. For each $\gamma \in \Gamma$, the fiber $v_{\text{rv}}^{-1}(\gamma)$ has a natural one-dimensional \overline{K} -affine structure, which is denoted as \overline{K}_γ . The direct sum $\bigoplus_{\gamma \in \Gamma} \overline{K}_\gamma$ may be viewed as a generalized residue field.

Let ACVF be the theory of algebraically closed valued fields in \mathcal{L}_{RV} . Let $\text{VF}_*[\cdot]$ and $\text{RV}[\cdot, \cdot]$ be two categories of definable sets with respect to the VF-sort and the RV-sort, respectively. In order to integrate

definable functions with RV-sort parameters, the objects in $\mathrm{VF}_*[\cdot]$ are exactly the definable subobjects of the products $\mathrm{VF}^n \times \mathrm{RV}^m$ and the morphisms are just the definable maps. On the other hand, for technical reasons (particularly for keeping track of dimensions), $\mathrm{RV}[*,\cdot]$ is formulated in a quite complicated way. All this is explained in Section 5.6. One of the main goals of the Hrushovski-Kazhdan integration theory is to construct a canonical homomorphism from the Grothendieck semigroup $\mathbf{K}_+ \mathrm{VF}_*[\cdot]$ to the Grothendieck semigroup $\mathbf{K}_+ \mathrm{RV}[*,\cdot]$ modulo a semigroup congruence relation \mathbf{I}_{sp} on the latter. In fact, it may be turned into an isomorphism if we take quotient with respect to a semigroup congruence relation \mathbf{I}_{bu} on $\mathbf{K}_+ \mathrm{VF}_*[\cdot]$. This construction has three main steps.

- *Step 1.* First we define a lifting map \mathbb{L} from the objects in $\mathrm{RV}[*,\cdot]$ into the objects in $\mathrm{VF}_*[\cdot]$; see Definition 5.6.16. Next we single out a subclass of isomorphisms in $\mathrm{VF}_*[\cdot]$, which are called definable special bijections; see Definition 5.7.5. Then we show that for any object X in $\mathrm{VF}_*[\cdot]$ there is a special bijection T on X and an object Y in $\mathrm{RV}[*,\cdot]$ such that $T(X)$ is isomorphic to $\mathbb{L}(Y)$. This implies that \mathbb{L} hits every isomorphism class of $\mathrm{VF}_*[\cdot]$. Of course, for this result alone we do not have to limit our means to special bijections. However, in Step 3 below, special bijections become an essential ingredient in computing the congruence relation \mathbf{I}_{sp} .
- *Step 2.* For any two isomorphic objects Y_1, Y_2 in $\mathrm{RV}[*,\cdot]$, their lifts $\mathbb{L}(Y_1), \mathbb{L}(Y_2)$ in $\mathrm{VF}_*[\cdot]$ are isomorphic as well. This shows that \mathbb{L} induces a semigroup homomorphism from $\mathbf{K}_+ \mathrm{RV}[*,\cdot]$ into $\mathbf{K}_+ \mathrm{VF}_*[\cdot]$, which is also denoted as \mathbb{L} .
- *Step 3.* In order to invert the homomorphism \mathbb{L} , we need a precise description of the semigroup congruence relation induced by it. The basic notion used in the description is that of a blowup of an object in $\mathrm{RV}[*,\cdot]$; see Definition 5.11.1. We then show that, for any objects Y_1, Y_2 in $\mathrm{RV}[*,\cdot]$, there are piecewise isomorphic parameterized iterated blowups Y_1^\sharp, Y_2^\sharp of Y_1, Y_2 if and only if there are parameterized special bijections T_1, T_2 on Y_1, Y_2 such that $T_1(\mathbb{L}(Y_1)), T_2(\mathbb{L}(Y_2))$ are piecewise isomorphic. The “if” direction contains a form of Fubini’s Theorem and is the most technically involved part of the construction. Its difficulty will be explained further below when we describe the course of the paper.

The inverse of \mathbb{L} thus obtained is a motivic integration; see Theorem 5.12.2.

5.1.2 Course of the paper

A remarkable feature of the Hrushovski-Kazhdan integration theory is that model-theoretic study of definable sets plays a fundamental role and yet no advanced results from model theory, say, beyond the first five

chapters of [40], are used. In section 5.2, after introducing the language \mathcal{L}_{RV} and the theory ACVF, we briefly review some concepts and results in model theory. To suggest how they may be used later, some of these, especially the various incarnations of the Compactness Theorem, are stated specifically for \mathcal{L}_{RV} and ACVF. We also give a syntactical description of what it means to have imaginary elements as parameters in defining sets. In section 5.3, we establish quantifier elimination for ACVF by one of the standard model-theoretic tests. This is not proved in [34] and the reader is referred to [31]. The theme of the latter is elimination of imaginaries and the relevant results use a much more complicated language than \mathcal{L}_{RV} , which do not seem to imply quantifier elimination in ACVF in a straightforward fashion. Our proof, except some fundamental facts in the theory of valued fields, is self-contained. In the following two sections we prove some properties that delineate the basic geography of definable sets in ACVF. These properties are used throughout the rest of the paper. As in [34], the key notion here is C -minimality, which was first introduced in [39] and has been further studied in [32]. The main difference between Section 5.4 and Section 5.5 is that in the former we work at the level of formulas with real parameters and in the latter we work at the level of types with imaginary parameters allowed.

With the preparatory work done, we are now ready to move on to the actual construction of motivic integration. First of all, we discuss various dimensions, mainly VF-dimension and RV-dimension, and describe the relevant categories of definable sets and the formulation of their Grothendieck semigroups in Section 5.6. The fundamental lifting map L between VF-categories and RV-categories and the “dummy” functor E between RV-categories are also introduced here. The central topic of Section 5.7 is RV-products and special bijections on them; see Definition 5.7.3 and Definition 5.7.5. The main result is Proposition 5.7.14, which corresponds to Step 1 above. This section contains the most important technical tool that is not available in [34], namely Proposition 5.7.13. With its presence, many hard lemmas in [34] have been simplified a great deal (for example, [34, Lemma 7.8], which corresponds to Lemma 5.10.2 in this paper) or circumvented (among the most notable ones are [34, Lemma 5.5] and the entire [34, Section 3.3]).

The notion of a 2-cell is introduced in Section 5.8, which corresponds to the notion of a bicell in [14]. This notion may look strange and is, perhaps, only of technical interest. It arises when we try to prove some form of Fubini’s Theorem, such as Lemma 5.11.24. The difficulty is that, although, using C -minimality, the construction of the integration of definable sets of VF-dimension 1 is very functorial (see Lemma 5.10.3), we are unable to extend this construction to higher VF-dimensions. This is the concern of [34, Question 7.9]. It has also occurred in [14] and may be traced back to [20]; see [14, Section 1.7]. Anyway, in this situation, the natural strategy of integrating definable sets of higher VF-dimensions is to use the result for VF-dimension 1 and integrate with respect to one VF-sort variable at a time. As in the classical theory of integration, this

strategy requires some form of Fubini's Theorem: for a well-behaved integration, an integral should give the same value when it is evaluated along different orders of VF-sort variables. By induction, this problem is immediately reduced to the case of two VF-sort variables. A 2-cell is a definable subset of VF^2 with certain symmetrical geometrical structure that satisfies this Fubini type of requirement. Now the idea is that, if we can find a definable partition for every definable subset such that each piece is a 2-cell indexed by some RV-sort parameters, then, by compactness, every definable subset satisfies the Fubini type of requirement. This kind of partition is achieved in Lemma 5.8.8.

Section 5.9 is devoted to showing Step 2 above. The notion of a $\bar{\gamma}$ -polynomial is introduced here, which generalizes the relation between a polynomial with coefficients in the valuation ring and its projection into the residue field. This leads to Lemma 5.9.2, a generalized form of the multivariate version of Hensel's Lemma. Note that in order to apply Lemma 5.9.2 to a given definable set we need to find suitable polynomials with a simple common residue root. This is investigated in Lemma 5.9.4, which does not hold when the substructure in question contains an excessive amount of parameters in the RV-sort. This is the reason why motivic integration is constructed only for theories of the form ACVF_S^0 , where the structure S is (VF, Γ) -generated. There is a straightforward remedy for this limitation. For every substructure S there is a canonical expansion S^* of S such that S^* is (VF, Γ) -generated and may be embedded into every (VF, Γ) -generated substructure that contains S ; see [34, Proposition 3.51]. Then S and S^* are identified for the construction of integration. To keep the conceptual framework simple, we do not include this treatment in the paper.

The key result of Section 5.10, Lemma 5.10.3, says that, modulo special bijections, every definable bijection between two definable sets of VF-dimension 1 is equal to the lift of an isomorphism in $\text{RV}[* , \cdot]$. As has been remarked above, it would be ideal to extend this result to definable sets of all VF-dimensions. Being unable to do this, we introduce the notion of a standard contraction, which gives rise to the Fubini type of problem described above; see Definition 5.10.6. Then in Lemma 5.10.8 we show that an essential part of Lemma 5.10.3 holds for 2-cells, which is good enough for the rest of the construction.

The task of identifying the kernel of \mathbb{L} , that is, Step 3 above, is carried out in Section 5.11. We introduce the notion of a blowup in Definition 5.11.1 and then extend it in Definition 5.11.3 to a parameterized version. The equivalence relation $\text{I}_{\text{sp}}[* , \cdot]$ on $\text{RV}[* , \cdot]$ induced by parameterized blowups is indeed a semigroup congruence relation; see 5.11.1 and Lemma 5.11.12. We also need to parameterize special bijections; see Definition 5.11.19. This induces a semigroup congruence relation I_{bu} on $\text{VF}_*[\cdot]$; see Definition 5.11.25. We conclude this section with Lemma 5.11.26, which says that $\text{I}_{\text{sp}}[* , \cdot]$ is the congruence relation induced by the homomorphism \mathbb{L} modulo I_{bu} . In the last section we assemble everything together and deduce the main

theorem.

5.1.3 Technical differences from the original construction

We emphasize again that, in this paper, we do not work at the level of generality as in [34], that is, the whole class of V -minimal theories. Instead, our construction is specialized for the theories of algebraically closed valued fields of pure characteristic 0 expanded by a substructure S in the language \mathcal{L}_{RV} . As has been discussed above, Step 2 of the construction requires S to be (VF, Γ) -generated, but other parts of the construction in general do not require this restriction. For this subclass of V -minimal theories we are able to work with syntax. Very often, in order to grasp the geometrical content of a definable set X , it is a very fruitful exercise to analyze the logical structure of a typical formula that defines X , especially when quantifier elimination is available. Consequently, in the context of this paper, syntactical analysis affords simplifications of many lemmas in [34]. The main technical differences are described here roughly in the order of their first appearances. For this purpose we fix a theory ACVF_S^0 , where S is (VF, Γ) -generated.

Let \mathcal{L}_v be the two-sorted language for valued fields: one sort for the field and the other for the value group. Every model of ACVF_S^0 may be turned naturally into a structure of \mathcal{L}_v and consequently any definable subset of any product VF^n in ACVF_S^0 is S -definable in \mathcal{L}_v . This translation provides the strategy in Section 5.3 to reduce quantifier elimination in \mathcal{L}_{RV} to that in \mathcal{L}_v , which has been established by Weispfenning (see Theorem 5.2.5). Another notable application of it is in Lemma 5.4.12, whose proof is conceptually much simpler than the corresponding [34, Lemma 3.35].

In almost all sections in the first eight chapters of [34] there are results that we need to more or less reproduce, except [34, Section 3.3], which has been completely dispensed with in this paper. Although, according to [34, Remark (3), p. 34], the lemmas in [34, Section 3.3] are not needed for the construction of integration maps, [34, Lemma 3.26] is used in the very important [34, Lemma 5.5], which is needed for [34, Lemma 5.10], which in turn is directly applied in [34, Lemma 7.24] to settle the Fubini problem described above. Because of Proposition 5.7.13 and its many consequences, we are still able to reproduce [34, Lemma 5.10], namely Lemma 5.8.8, without [34, Lemma 5.5]. More details on Proposition 5.7.13 will be given below.

In Section 5.5 we follow the syntactical treatment of imaginary elements described in Section 5.2. In particular, we are able to show that an atomic closed ball or an atomic thin annulus cannot correspond algebraically to an atomic open ball, which implies that one cannot define an atomic closed ball from an atomic open ball; see Lemma 5.5.9 and Lemma 5.5.10. These and Lemma 5.5.11 yield (trivially) a special case of [34, Lemma 3.46].

In order to bypass the notion of measure-preserving isomorphisms in RV-categories (see [34, Definition 5.21]), which requires a discussion of differential calculus, the very simple notion of the weight of an RV-sort tuple (Definition 5.6.10) is introduced in Section 5.6. This is used to formulate one of the conditions in the definition of a morphism in RV-categories; see Definition 5.6.11. The idea is that, a morphism $F : X \rightarrow Y$ in an RV-category should encode the ordering of the volumes of the lifts $\mathbb{L}(X)$, $\mathbb{L}(Y)$ of X , Y so that F itself may be lifted to the corresponding VF-category. To be more concrete, suppose that $X = \{1\}$ and $Y = \{\infty\}$, then $\mathbb{L}(X) = \text{rv}^{-1}(1) \times \{1\}$ and $\mathbb{L}(Y) = \text{rv}^{-1}(\infty) \times \{\infty\} = \{(0, \infty)\}$ and hence if F is an isomorphism then it is impossible to lift it to an isomorphism. The solution to this is to simply disqualify F as a morphism but allow F^{-1} to be a morphism, which amounts to adopting the alternative definition of RV-categories in [34, Section 3.8.1]. A main advantage of allowing the element ∞ in RV-categories is that it makes the discussion of blowups in Section 5.11 more streamlined.

In [34, Chapter 4], Step 1 of the construction is accomplished through a class of bijections called admissible transformations. Later in [34, Chapter 7] another class of bijections called special bijections are introduced for Step 3. In this paper the two classes are adjusted so that they may be unified into one class and still serve their original purposes; see Definition 5.7.5. Now we come to Proposition 5.7.13, which says that, up to isomorphism classes, a polynomial map on an object in a VF-category may be projected down to a morphism between two objects in the corresponding RV-category. To be more precise, let $f(x_1, \dots, x_n)$ be a polynomial with VF-sort coefficients and X a definable subset of VF^n , then there is a definable special bijection T on X such that there is a function $f_{\downarrow} : \text{RV}^m \rightarrow \text{RV}$ that makes the diagram

$$\begin{array}{ccc} T(X) & \xleftarrow{T} X & \xrightarrow{f} f(X) \\ \text{rv} \downarrow & & \downarrow \text{rv} \\ \text{RV}^m & \xrightarrow{\quad} & \text{RV} \\ & & f_{\downarrow} \end{array}$$

commute. Moreover, this may be carried out simultaneously for any finite number of VF-sort polynomials. Except Section 5.9, the remainder of the paper heavily relies on Proposition 5.7.13. Almost all its applications involve the following procedure. Given a morphism $f : X \rightarrow Y$ in $\text{VF}_*[\cdot]$ that is defined by a formula ϕ , we obtain a special bijection T on X such that for any term in ϕ of the form $\text{rv}(g(\bar{x}))$ there is a commutative diagram as above for $g(\bar{x})$ and hence the morphism $f \circ T^{-1}$ in $\text{VF}_*[\cdot]$ may be projected down to a morphism in $\text{RV}[* , \cdot]$.

In Section 5.8 we give a more detailed treatment of 2-cells than in [34]. The lemmas that lead to Lemma 5.8.8 should make clear the crucial role of Proposition 5.7.13.

Let $t \neq \infty$ be an RV-sort element that is algebraic over some other RV-sort elements. In Lemma 5.9.4,

through analyzing a suitable formula that witnesses this algebraic relation, we find a minimal $\bar{\gamma}$ -polynomial for t . This essentially reduces the task of lifting isomorphisms in $\text{RV}[* , \cdot]$ (Lemma 5.9.6) to the multivariate version of Hensel’s Lemma. The proof of [34, Proposition 6.1] is thus simplified.

Section 5.10 and Section 5.11 more or less correspond to [34, Section 7.2, Section 7.3] and [34, Section 7.4, Section 7.5], respectively. Most of the changes here are made with the hope that the difficult situation may become easier to grasp. For example, unlike in [34, Section 7.5], we do not form additional categories for the computation of the kernel of \mathbb{L} . Instead, we work directly with objects in $\text{VF}_*[\cdot]$ and operations on them called standard contractions, which are a natural conceptual extension of special bijections; see Definition 5.10.6.

There is an important change in this last part of the paper that is of a different nature: since, for any object X in $\text{VF}_*[\cdot]$, we can only establish the desired correspondence between standard contractions and blowups on each piece of a suitable definable partition of X , the congruence relation I_{sp} on $\mathbf{K}_+ \text{RV}[* , \cdot]$ is defined through parameterized blowups instead of blowups. Consequently special bijections are parameterized as well so that we may have a canonical description of the image I_{bu} of I_{sp} under the surjective homomorphism $\mathbb{L} : \mathbf{K}_+ \text{RV}[* , \cdot] \longrightarrow \mathbf{K}_+ \text{VF}_*[\cdot]$, which is a congruence relation on $\mathbf{K}_+ \text{VF}_*[\cdot]$. We note that this complication seems to have been avoided in [34] through the use of [34, Lemma 7.20], which we are unable to reproduce in this paper. In comparison with [34, Theorem 8.8], Theorem 5.12.2 offers a surjective homomorphism instead of an isomorphism, which, of course, may be turned into an isomorphism if we take quotient with respect to I_{bu} . However, no modified version of [34, Corollary 8.9] is obtained.

5.2 Logical preliminaries and the theory ACVF

In this section we review some of the basic concepts and results from model theory that will be used in the construction. In order to make connections with our context as quickly as possible, many of them will be stated in forms that directly involve the language \mathcal{L}_{RV} and the theory ACVF. The main advantage of being particular here is that it allows us to exemplify the many ways to use compactness in [34]. Since a thorough list of them all is not feasible, hopefully these examples may function as a guide so that every usage of compactness below will be seen as an easy variation of one of them.

5.2.1 The setting of \mathcal{L}_{RV} and ACVF

Let us first introduce the Basarab-Kuhlmann style language \mathcal{L}_{RV} for algebraically closed valued fields. This style first appeared in [5] and [6] and has been further investigated in [36] and [44]. Its main feature is the

use of a countable collection of residue multiplicative structures, which are reduced to just one for valued fields of pure characteristic 0.

Definition 5.2.1. The language \mathcal{L}_{RV} has the following sorts and symbols:

1. a VF-sort, which uses the language of rings $\mathcal{L}_{\text{R}} = \{0, 1, +, -, \times\}$;
2. an RV-sort, which uses
 - (a) the group language $\{1, \times\}$,
 - (b) two constant symbols 0 and ∞ ,
 - (c) a unary predicate \overline{K}^\times ,
 - (d) a binary function $+: \overline{K}^2 \longrightarrow \overline{K}$ and a unary function $-: \overline{K} \longrightarrow \overline{K}$, where $\overline{K} = \overline{K}^\times \cup \{0\}$,
 - (e) a binary relation \leq ;
3. a function symbol rv from the VF-sort into the RV-sort.

Technically speaking, the constant 0 and the functions $+$, $-$ in the RV-sort should all be relations. This point of view may be more convenient in some of the statements and arguments below that are of a syntactical nature. For notational convenience, we do not use different symbols for 0 and 1, since which ones are being referred to should always be clear in context.

Notation 5.2.2. The two sorts without the zero elements are denoted as VF^\times and RV , $\text{RV} \setminus \{\infty\}$ is denoted as RV^\times , and $\text{RV} \cup \{0\}$ is denoted as RV_0 . For any structure M of \mathcal{L}_{RV} and any formula ϕ with parameters in M , we write $\phi(M)$ for the subset defined by ϕ in M . In particular, we write $\text{VF}(M)$, $\text{RV}(M)$, $\text{RV}^\times(M)$, $\overline{K}(M)$, etc. for the corresponding subsets of M . These are simply written as VF , RV , RV^\times , \overline{K} , etc. when the structure in question is clear or when the discussion takes place in an ambient monster model (that is, a universal domain that embeds all “small” models that will occur in the discussion). For any subset $X \subseteq \text{VF}(M)^n \times \text{RV}(M)^m$, we write $\bar{a} \in X$ to mean that every element in the tuple \bar{a} is in X . In particular, we often write (\bar{a}, \bar{t}) for a tuple of elements in M with the understanding that $\bar{a} \in \text{VF}$ and $\bar{t} \in \text{RV}$. For such a tuple $(\bar{a}, \bar{t}) = (a_1, \dots, a_n, t_1, \dots, t_m)$, let

$$\begin{aligned} \text{rv}(\bar{a}, \bar{t}) &= (\text{rv}(a_1), \dots, \text{rv}(a_n), \bar{t}) \\ \text{rv}^{-1}(\bar{a}, \bar{t}) &= \{\bar{a}\} \times \text{rv}^{-1}(t_1) \times \dots \times \text{rv}^{-1}(t_m); \end{aligned}$$

similarly for other functions.

Let M be a structure of \mathcal{L}_{RV} . For any subset $A \subseteq M$, the smallest substructure of M containing A is denoted as $\langle A \rangle$. An element $b \in M$ is *A-definable* if there is a tuple $\bar{a} \in A$ such that b is \bar{a} -definable, that is, b is defined by a formula $\phi(\bar{a})$. The definable closure of A in M , which is the smallest substructure of M containing all the A -definable elements, is denoted as $\text{dcl}(A)$. Note that, although in general $\langle A \rangle \neq \text{dcl}(A)$, they may be identified as far as definable sets are concerned. Except in Section 5.3, this is what we shall do below. An element $b \in M$ is *algebraic over A*, or *A-algebraic*, if it is algebraic over some $\bar{a} \in A$, that is, there is a formula $\phi(\bar{a})$ that defines a finite subset of M containing b . The algebraic closure of A in M , which is the smallest substructure of M containing all the $\langle A \rangle$ -algebraic elements, is denoted as $\text{acl}(A)$. A basic fact is that, if M models a complete theory in \mathcal{L}_{RV} , then $\text{acl}(A)$ is the same (up to isomorphism, of course) in any other model of the theory that contains A .

Let M be a structure of \mathcal{L}_{RV} , $D \subseteq \text{VF}(M)^n \times \text{RV}(M)^m$ a definable subset, and E a definable equivalence relation on D . Each equivalence class under E is an *imaginary element* of M and the collection D/E of the equivalence classes is an *imaginary sort* of M . An imaginary element may occur in a formula as a parameter. Semantically, this means taking union of all the subsets defined by formulas $\phi(\bar{a}, \bar{t})$, where the parameters (\bar{a}, \bar{t}) run through all the “real” elements contained in the equivalence class. Syntactically, it corresponds to an extra existential quantifier and the invariance of the subset that is being defined when a different representative of the equivalence class is used. Examples will be given below after the imaginary sorts of values and balls have been defined.

Definition 5.2.3. The *theory of algebraically closed valued fields of characteristic 0 in \mathcal{L}_{RV}* (hereafter abbreviated as ACVF) states the following:

1. $(\text{VF}, 0, 1, +, -, \times)$ is an algebraically close field of characteristic 0;
2. $(\text{RV}^\times, 1, \times)$ is a divisible abelian group, where multiplication \times is augmented by $t \times 0 = 0$ for all $t \in \overline{K}$ and $t \times \infty = \infty$ for all $t \in \text{RV}_0$;
3. $(\overline{K}, 0, 1, +, -, \times)$ is an algebraically closed field;
4. the relation \leq is a preordering on RV with ∞ the top element and \overline{K}^\times the equivalence class of 1;
5. the quotient $\text{RV} / \overline{K}^\times$, denoted as $\Gamma \cup \{\infty\}$, is a divisible ordered abelian group with a top element, where the ordering and the group operation are induced by \leq and \times , respectively, and the quotient map $\text{RV} \longrightarrow \Gamma \cup \{\infty\}$ is denoted as rv ;
6. the function $\text{rv} : \text{VF}^\times \longrightarrow \text{RV}^\times$ is a surjective group homomorphism augmented by $\text{rv}(0) = \infty$ such

that the composite function

$$\text{val} = v_{\text{rv}} \circ \text{rv} : \text{VF} \longrightarrow \Gamma \cup \{\infty\}$$

is a valuation with the valuation ring $\mathcal{O} = \text{rv}^{-1}(\text{RV}^{\geq 1})$ and its maximal ideal $\mathcal{M} = \text{rv}^{-1}(\text{RV}^{> 1})$, where

$$\text{RV}^{\geq 1} = \{x \in \text{RV} : 1 \leq x\},$$

$$\text{RV}^{> 1} = \{x \in \text{RV} : 1 < x\};$$

The set $\mathcal{O} \setminus \mathcal{M}$ of units in the valuation ring is sometimes denoted as \mathcal{U} . In any model of ACVF, the function $\text{rv} \upharpoonright \text{VF}^\times$ may be identified with the quotient map $\text{VF}^\times \longrightarrow \text{VF}^\times / (1 + \mathcal{M})$. Hence an RV-sort element t may be understood as a coset of $(1 + \mathcal{M})$. We occasionally treat t as a set and write $a \in t$ to mean that $a \in \text{rv}^{-1}(t)$.

Although we do not include the multiplicative inverse function in the VF-sort and the RV-sort, we always assume that, without loss of generality, $\text{VF}(S)$ is a field and $\text{RV}^\times(S)$ is a group for a substructure S of a model of ACVF.

Remark 5.2.4. Let \mathcal{L}_v be the natural two-sorted language for valued fields: one sort for the field and the other for the value group. With the imaginary Γ -sort and the valuation map val , \mathcal{L}_{RV} may be viewed as an expansion of \mathcal{L}_v . Each valued field may be turned naturally into an \mathcal{L}_{RV} -structure and hence an \mathcal{L}_v -structure. In fact, it is not hard to see that, under the natural interpretations, two valued fields are isomorphic as \mathcal{L}_{RV} -structures if and only if they are isomorphic as \mathcal{L}_v -structures. Henceforth we shall refer to the two sorts of \mathcal{L}_v as the VF-sort and the Γ -sort.

In Section 5.3 we shall establish quantifier elimination for ACVF. The strategy of the proof is to reduce the problem to the following fundamental result of Weispfenning's [51, Theorem 3.2]:

Theorem 5.2.5. *The theory of algebraically closed valued fields of characteristic 0 as formulated in \mathcal{L}_v admits quantifier elimination.*

It is equivalent to quantifier elimination that, for any substructure S of a model of ACVF, the theory ACVF_S — that is, the union of ACVF and the set of all quantifier-free formulas $\phi(\bar{a})$ with $\bar{a} \in S$ that hold in S — is complete. This implies that, for every integer $p \geq 0$, the theory $\text{ACVF}^p = \text{ACVF} \cup \{\text{char } \bar{K} = p\}$ is complete. It is a basic fact in model theory that monster models (that is, universal domains) are guaranteed to exist for complete theories.

Convention 5.2.6. Henceforth, except in Section 5.3, we assume that everything happens in an ambient monster model \mathfrak{C} of ACVF_S^0 , where S is a fixed “small” substructure of \mathfrak{C} . Accordingly, below, in terms such

as “definable” (that is, “ \emptyset -definable”), \bar{a} -definable, “ $\text{acl}(\emptyset)$ ”, “ \mathcal{L}_{RV} ”, etc. we shall always mean “ S -definable”, “ $\langle S, \bar{a} \rangle$ -definable”, “ $\text{acl}(S)$ ”, “ $\mathcal{L}_{\text{RV}} \cup S$ ”, etc. When the additional parameters are not specified, we will just say “parametrically definable”.

The imaginary sort $\Gamma \cup \{\infty\}$ is called the Γ -sort. We write $\bar{t} \in \bar{\gamma}$ to mean that $v_{\text{rv}}(\bar{t}) = \bar{\gamma}$. For any subset A , the assertion that $v_{\text{rv}}^{-1}(\gamma_i) \subseteq A$ for every γ_i in the tuple $\bar{\gamma}$ is abbreviated as $\bar{\gamma} \in A$. A subset X is $\bar{\gamma}$ -definable if there is a formula $\phi(\bar{z})$ such that $X = \bigcup_{\bar{t} \in \bar{\gamma}} X_{\bar{t}}$, where $X_{\bar{t}}$ is the subset defined by $\phi(\bar{t})$. Syntactically, X is defined by any formula of the form

$$\exists \bar{x} (\text{rv}(\bar{x}) \leq \bar{t} \wedge \text{rv}(\bar{x}) \geq \bar{t} \wedge \phi(\text{rv}(\bar{x}))),$$

where $\bar{t} \in \bar{\gamma}$ and no element in $\bar{\gamma}$ occurs in $\phi(\text{rv}(\bar{x}))$. Accordingly, when a subset $A \subseteq \text{VF} \cup \text{RV} \cup \Gamma$ is used as a source of parameters, the elements in $\Gamma(A)$ can only occur in formulas of the above form. Naturally, the definable closure $\text{dcl}(A)$ of A also contains those elements that are definable with parameters in $\Gamma(A)$. Similarly for the algebraic closure $\text{acl}(A)$ of A .

A substructure S is VF-generated if $S = \text{dcl}(A)$ for some $A \subseteq \text{VF}$. Similarly for RV, Γ , and any combination of the three sorts. From now on, unless specified otherwise, a substructure is always (VF, RV, Γ)-generated.

Notation 5.2.7. Coordinate projection maps are ubiquitous in this paper. To facilitate the discussion, certain notational conventions about them are adopted.

Let $X \subseteq \text{VF}^n \times \text{RV}^m$. For any $n \in \mathbb{N}$, let $\mathcal{I}_n = \{1, \dots, n\}$. First of all, the VF-coordinates and the RV-coordinates of X are indexed separately. It is cumbersome to actually distinguish them notationally, so we just assume that the set of the indices of the VF-coordinates (VF-indices) is \mathcal{I}_n and the set of the indices of the RV-coordinates (RV-indices) is \mathcal{I}_m . This should never cause confusion in context. Let $\mathcal{I}_{n,m} = \mathcal{I}_n \uplus \mathcal{I}_m$, $E \subseteq \mathcal{I}_{n,m}$, and $\tilde{E} = \mathcal{I}_{n,m} \setminus E$. If E is a singleton $\{i\}$ then we always write E as i and \tilde{E} as \tilde{i} . We write $\text{pr}_E X$ for the projection of X to the coordinates in E . For any $\bar{x} \in \text{pr}_{\tilde{E}} X$, the fiber $\{\bar{y} : (\bar{y}, \bar{x}) \in X\}$ is denoted as $\text{fib}(X, \bar{x})$. Note that, for notational convenience, we shall often tacitly identify the two subsets $\text{fib}(X, \bar{x})$ and $\text{fib}(X, \bar{x}) \times \{\bar{x}\}$. Also, it is often more convenient to use simple descriptions as subscripts. For example, if $E = \{1, \dots, k\}$ etc. then we may write $\text{pr}_{\leq k}$ etc. If E contains exactly the VF-indices (respectively RV-indices) then pr_E is written as pVF (respectively pRV). Suppose that E' is a subset of the indices of the coordinates of $\text{pr}_E X$. Then the composition $\text{pr}_{E'} \circ \text{pr}_E$ is written as $\text{pr}_{E', E'}$. Naturally $\text{pr}_{E'} \circ \text{pVF}$ and $\text{pr}_{E'} \circ \text{pRV}$ are written as $\text{pVF}_{E'}$ and $\text{pRV}_{E'}$.

Suppose that X , V , and W are all definable subsets and $X \subseteq V \times W$. Sometimes we shall want to

investigate the fibers of X of the form $\text{fib}(X, \bar{v})$ with $\bar{v} \in V$. Note that $\text{fib}(X, \bar{v})$ is in general not definable. Of course it is $\langle \bar{v} \rangle$ -definable. Many properties and notions below depend on the underlying substructure from which the subsets in question are definable. Hence, below, when we study fibers of X , we shall always assume that the underlying substructure has been expanded in an appropriate way.

We shall frequently need to keep track of the correspondence between the VF-indices and the RV-indices in a subset derived from X . It is unduly complicated to describe a precise indexing scheme that is suitable for this task and hence we shall not attempt it here. Instead, we shall give a few typical examples and then rely on the reader's intuition to figure out the actual indexing in each instance. There is a principle that underlies these examples: coordinates of interest get indices as small as possible. Let

$$\mathbf{c}(X) = \{(\bar{a}, \text{rv}(\bar{a}), \bar{t}) : (\bar{a}, \bar{t}) \in X\} \subseteq \text{VF}^n \times \text{RV}^{n+m}.$$

Clearly X is definably bijective to $\mathbf{c}(X)$ in a canonical way. This bijection is called the *canonical bijection* and is denoted as \mathbf{c} . In $\mathbf{c}(X)$, the set of the new RV-indices created by the map rv is \mathcal{I}_n . Next, let

$$X^* = \bigcup \{\text{rv}^{-1}(\bar{t}) \times \{(\bar{a}, \bar{t})\} : (\bar{a}, \bar{t}) \in X\} \subseteq \text{VF}^{m+n} \times \text{RV}^m.$$

In X^* , the set of the new VF-indices created by the “lifting” map rv^{-1} is \mathcal{I}_m . Lastly, let $f : X \rightarrow \text{VF}^n \times \text{RV}^m$ be a definable function such that, for every $(\bar{a}, \bar{t}) \in X$, $(\text{pr}_{>1} \circ f)(\bar{a}, \bar{t}) = \text{pr}_{>1}(\bar{a}, \bar{t})$. Let $Y = (\text{pr}_{>1} \circ \mathbf{c} \circ f)(X) \subseteq \text{VF}^{n-1} \times \text{RV}^{n+m}$ and $g : Y \rightarrow \text{VF}^{n-1} \times \text{RV}^{n+m}$ a definable function such that, for every $(\bar{a}, \bar{t}) \in Y$, $(\text{pRV} \circ g)(\bar{a}, \bar{t}) = \bar{t}$. Let $Z = (\text{pRV} \circ \mathbf{c} \circ g)(Y) \subseteq \text{RV}^{2n+m-1}$. Among the coordinates of Z there are n special ones that correspond to the VF-coordinates of X , which have been truncated in the transformation from X to Z . These special coordinates are indexed by $1, \dots, n$.

We now turn to the other important kind of imaginary elements: balls. The open balls form a basis of the valuation topology. Basic properties of balls will be explored in Section 5.4.

Definition 5.2.8. A subset \mathfrak{b} of VF is an *open ball* if there is a $\gamma \in \Gamma$ and a $b \in \mathfrak{b}$ such that $a \in \mathfrak{b}$ if and only if $\text{val}(a - b) > \gamma$. It is a *closed ball* if $a \in \mathfrak{b}$ if and only if $\text{val}(a - b) \geq \gamma$. It is an *rv-ball* if $\mathfrak{b} = \text{rv}^{-1}(t)$ for some $t \in \text{RV}$. The value γ is the *radius* of \mathfrak{b} , which is denoted as $\text{rad}(\mathfrak{b})$. If val is constant on \mathfrak{b} — that is, \mathfrak{b} is contained in an rv-ball — then $\text{val}(\mathfrak{b})$ is the *valuative center* of \mathfrak{b} ; if val is not constant on \mathfrak{b} , that is, $0 \in \mathfrak{b}$, then the *valuative center* of \mathfrak{b} is ∞ . The valuative center of \mathfrak{b} is denoted by $\text{vcr}(\mathfrak{b})$.

Note that each point in VF is a closed ball of radius ∞ . Also, we shall regard VF as a clopen ball of radius $-\infty$.

A ball \mathfrak{b} may be represented by a triple $(a, b, d) \in \text{VF}^3$, where $a \in \mathfrak{b}$, $\text{val}(b)$ is the radius of \mathfrak{b} , and $d = 1$ if \mathfrak{b} is open and $d = 0$ if \mathfrak{b} is closed. A set \mathfrak{B} of balls is a subset of VF^3 of triples of this form such that if $(a, b, d) \in \mathfrak{B}$ then for all $a' \in \text{VF}$ with $\text{rv}(a - a') \square_d b$, where \square_d is $>$ if $d = 1$ or \geq if $d = 0$, there is a $b' \in \text{VF}$ with $\text{val}(b) = \text{val}(b')$ such that $(a', b', d) \in \mathfrak{B}$. Clearly two triples $(a, b, d), (a', b', d') \in \mathfrak{B}$ represent two different balls, which may or may not be disjoint, if and only if either $(\text{val}(b), d) \neq (\text{val}(b'), d')$ or, in case that they are the same, $\text{rv}(a - a') \square_d b$ does not hold.

We note the following terminological convention. The union of \mathfrak{B} , sometimes written as $\bigcup \mathfrak{B}$, is actually the subset $\text{pr}_1 \mathfrak{B}$. For any subset $A \subseteq \text{VF}$, the assertion that $\bigcup \mathfrak{B} \subseteq A$ may simply be written as $\mathfrak{B} \subseteq A$. We say that \mathfrak{B} is finite if it contains finitely many distinct balls. A subset of \mathfrak{B} is always a set of balls in \mathfrak{B} . A function f of \mathfrak{B} is always a function on the balls in \mathfrak{B} ; that is, f is a relation between \mathfrak{B} and a subset W such that for every $\mathfrak{b} \in \mathfrak{B}$ there is a unique $w \in W$ between which and every $(a, b, d) \in \mathfrak{b}$ the relation holds. Notice that f may or may not be a function on the triples in \mathfrak{B} .

Remark 5.2.9. In a similar way a ball \mathfrak{b} may be represented by a triple in $\text{VF} \times \text{RV}^2$. This representation is sometimes more convenient. Below we shall not distinguish these two representations.

We have seen above how to use elements in the imaginary Γ -sort as parameters in formulas. The idea is the same for balls. Let \mathfrak{b} be a ball. A subset X is $\langle \mathfrak{b} \rangle$ -definable if there is a formula $\phi(x, y, z)$ such that $X = \bigcup_{(a, b, d) \in \mathfrak{b}} X_{(a, b, d)}$, where $(a, b, d) \in \text{VF}^3$ is a representative of \mathfrak{b} and $X_{(a, b, d)}$ is the subset defined by $\phi(a, b, d)$. Syntactically, X is defined by any formula of the form

$$\exists x, y, z (\text{rv}(x - a) \square_d b \wedge \text{rv}(y) \geq \text{rv}(b) \wedge \text{rv}(y) \leq \text{rv}(b) \wedge z = d \wedge \phi(x, y, z)),$$

where $(a, b, d) \in \text{VF}^3$ is any representative of \mathfrak{b} and no representative of \mathfrak{b} occurs in $\phi(x, y, z)$ and \square_d is $>$ if $d = 1$ or \geq if $d = 0$. Accordingly, if a subset A contains balls and is used as a source of parameters, then the balls in A can only occur in formulas of the above form. With this understanding, the definable closure $\text{dcl}(A)$ and the algebraic closure $\text{acl}(A)$ of A may be defined in the obvious way.

5.2.2 Compactness

The use of the Compactness Theorem in [34] is extensive. Here we prove a few lemmas to illustrate it.

Definition 5.2.10. Let X, Y be definable subsets and $p : X \rightarrow Y$ a definable function. A definable function f is a p -function if there is a $Y' \subseteq Y$ and a partial function \hat{f} on X such that $\text{dom}(\hat{f}) = p^{-1}(Y')$ and $f = p \times \hat{f}$. Let $\Phi(p)$ be a set of p -functions. We say that $\Phi(p)$ is p -closed if for all $f_1, \dots, f_n \in \Phi(p)$ there is an $f \in \Phi(p)$ such that $\text{dom}(f) = \bigcup_i \text{dom}(f_i)$ and, for each $\bar{y} \in Y$ with $p^{-1}(\bar{y}) \subseteq \text{dom}(f)$, there is an

f_i such that $f \upharpoonright p^{-1}(\bar{y}) = \{\bar{y}\} \times (\widehat{f}_i \upharpoonright p^{-1}(\bar{y}))$, where \widehat{f}_i is the partial function such that $f_i = p \times \widehat{f}_i$.

Let X be a definable subset and p a definable function such that $X \subsetneq \text{dom}(p)$. In this situation a p -function with respect to X should always be understood as a $(p \upharpoonright X)$ -function.

Lemma 5.2.11. *Let X, Y be definable subsets, $p : X \rightarrow Y$ a definable function, and $\Phi(p)$ a set of p -functions that is p -closed. Suppose that, for every $\bar{y} \in Y$, there is an $f_{\bar{y}} \in \Phi(p)$ such that $f_{\bar{y}}$ is injective on $p^{-1}(\bar{y})$. Then there is an $f \in \Phi(p)$ such that $f : X \rightarrow Y \times Z$ is an injective function for some definable subset Z .*

Proof. Suppose for contradiction that no $f \in \Phi(p)$ is an injective function on X of the required form. Let $\mathcal{L} = \mathcal{L}_{\text{RV}} \cup \{\bar{c}\}$, where \bar{c} are new constants. Consider the \mathcal{L} -theory T that states the following:

1. everything in ACVF_S^0 ,
2. $\bar{c} \in Y$,
3. every $f \in \Phi(p)$ fails to be injective on $p^{-1}(\bar{c})$.

If T is not consistent then there is a finite list of functions $f_i \in \Phi(p)$ such that, for all $\bar{y} \in Y$, one of the functions f_i is injective on $p^{-1}(\bar{y})$. Since $\Phi(p)$ is p -closed, there is a function $f \in \Phi(p)$ on X such that, for each $\bar{y} \in Y$, there is an f_i such that $f(\bar{x}) = (p(\bar{x}), \widehat{f}_i(\bar{x}))$ for every $\bar{x} \in p^{-1}(\bar{y})$. Clearly f is an injective function on X of the required form, contradiction. So T is consistent and there is a model $N \models T$. Since N is also a model of ACVF_S^0 , we have that $\bar{c}^N \in Y$ and, by assumption, there is an $f_{\bar{c}^N} \in \Phi(p)$ that is injective on $p^{-1}(\bar{c}^N)$, contradiction again. \square

In application, the function p in this lemma is often taken to be the map rv ; see, for example, Lemma 5.4.3. The flexibility of Lemma 5.2.11 is twofold: on the one hand, injectivity may be replaced by other first-order properties and, on the other hand, restrictions may be imposed on the set $\Phi(p)$ so that we can achieve better control over the form of the function f . In the following sections, the phrase “by compactness” often means “by a variation of Lemma 5.2.11”.

Lemma 5.2.12. *Let $\bar{t}, s \in \text{RV}$ and $X \subseteq \text{rv}^{-1}(\bar{t})$ a \bar{t} -definable subset such that, for every $\bar{a} \in X$, $s \in \text{acl}(\bar{a})$. Then $s \in \text{acl}(\bar{t})$.*

Proof. Let $\mathcal{L} = \mathcal{L}_{\text{RV}} \cup \{\bar{t}, s, \bar{c}\}$, where \bar{c} are new constants. Consider the \mathcal{L} -theory T that states the following:

1. everything in $\text{ACVF}_{\langle \bar{t}, s \rangle}^0$,
2. $\bar{c} \in X$,

3. for every \mathcal{L} -formula ϕ that does not contain s and every integer $k > 0$, either the subset defined by ϕ is of size at most k but does not contain s or it is of size greater than k .

By the assumption, T must be inconsistent. Therefore, there are integers k_1, \dots, k_m , \mathcal{L} -formulas $\phi_1(\bar{x}, y), \dots, \phi_m(\bar{x}, y)$ that do not contain s , and subsets X_1, \dots, X_m of X defined by $\phi_1^*, \dots, \phi_m^*$, where ϕ_i^* is the formula

$$\exists y_1, \dots, y_{k_i} \forall y \left(\phi_i(\bar{x}, y) \rightarrow \bigvee_{1 \leq j \leq k_i} y = y_j \right),$$

such that $\bigcup_i X_i = X$ and, for every $\bar{a} \in X_i$, the formula $\phi_i(\bar{a}, y)$ defines a finite subset $U_{\bar{a}}$ containing s of size at most k_i . Without loss of generality, we may assume that X_1, \dots, X_m are pairwise disjoint. Then $\bigcap_{\bar{a} \in X} U_{\bar{a}}$ is a \bar{t} -definable finite subset that contains s . \square

For the proof of the next lemma we need to assume quantifier elimination, which is to be established in Section 5.3.

Lemma 5.2.13. *The exchange principle holds in both sorts:*

1. For any $a, b \in \text{VF}$, if $a \in \text{acl}(b) \setminus \text{acl}(\emptyset)$ then $b \in \text{acl}(a)$.
2. For any $t, s \in \text{RV}$, if $t \in \text{acl}(s) \setminus \text{acl}(\emptyset)$ then $s \in \text{acl}(t)$.

Proof. For the first item, let ϕ be a quantifier-free formula in disjunctive normal form that witnesses $a \in \text{acl}(b)$. For any term $\text{rv}(g(x))$ in ϕ , where $g(x) \in \text{VF}(\langle b \rangle)[x]$, and any $d \in \text{VF}$, if $\text{val}(d - a)$ is sufficiently large then $\text{rv}(g(a)) = \text{rv}(g(d))$. On the other hand, clearly VF-sort disequalities cannot define nonempty finite subset. Therefore every irredundant disjunct of ϕ has a conjunct of the form $f(x, b) = 0$, where $f(x, b) \in \text{VF}(\langle b \rangle)[x]$. If $f(a, b) = 0$ then, since $a \notin \text{acl}(\emptyset)$, we must have that $f(x, b) \notin \text{VF}(\langle \emptyset \rangle)[x]$. So the item follows from the exchange principle in field theory.

For the second item, again let ϕ be a quantifier-free formula in disjunctive normal form that witnesses $t \in \text{acl}(s)$. Clearly we may assume that ϕ does not contain any VF-sort literal. So each literal in ϕ may be assumed to be of the form

$$\sum_i (\text{rv}(a_i) \cdot r_i \cdot x^{n_i}) \square \text{rv}(a) \cdot r \cdot x^m \cdot \sum_j (\text{rv}(a_j) \cdot r_j \cdot x^{n_j}),$$

where $a_i, a, a_j \in \text{VF}(\langle s \rangle)$, $r_i, r, r_j \in \text{RV}(\langle s \rangle)$, and \square is one of the symbols $=, \neq, \leq,$ and $>$. It is easily seen that, in ϕ , the inequalities cannot define nonempty finite subset and neither can the disequalities. Therefore every irredundant disjunct of ϕ has an equality conjunct. Since $t \notin \text{acl}(\emptyset)$, again, the item follows from the exchange principle in field theory. \square

Lemma 5.2.14. *Let $f : X \longrightarrow Y$ be a definable surjective function, where $X, Y \subseteq \text{VF}$. Then there are definable disjoint subsets $Y_1, Y_2 \subseteq Y$ with $Y_1 \cup Y_2 = Y$ such that Y_1 is finite, $f^{-1}(b)$ is infinite for each $b \in Y_1$, and the function $f \upharpoonright f^{-1}(Y_2)$ is finite-to-one.*

Proof. For each $b \in Y$, if $f^{-1}(b)$ is infinite then, by compactness, there is an $a \in f^{-1}(b)$ such that $a \notin \text{acl}(b)$. Since $b \in \text{dcl}(a) \subseteq \text{acl}(a)$, by Lemma 5.2.13, we must have that $b \notin \text{acl}(a) \setminus \text{acl}(\emptyset)$ and hence $b \in \text{acl}(\emptyset)$. Let $\mathcal{L} = \mathcal{L}_{\text{RV}} \cup \{c\}$, where c is a new constant. Consider the \mathcal{L} -theory T that states the following:

1. everything in ACVF_S^0 ,
2. $c \in Y$,
3. $|f^{-1}(c)| > k$ for every integer $k > 0$,
4. for every \mathcal{L}_{RV} -formula ϕ and every integer $k > 0$, either the subset defined by ϕ is of size at most k but it does not contain c or it is of size greater than k .

If $N \models T$ then $c^N \in Y$ and $f^{-1}(c^N)$ is infinite and $c^N \notin \text{acl}(\emptyset)$, contradiction. So T is inconsistent. So there is an \mathcal{L}_{RV} -formula ϕ and an integer $k > 0$ such that $\phi(\mathfrak{C})$ is finite and, for every $b \in Y$, if $|f^{-1}(b)| > k$ then $b \in \phi(\mathfrak{C})$. Let $Y_1 = \{b \in Y : f^{-1}(b) \text{ is infinite}\}$ and $Y_2 = Y \setminus Y_1$. Since $Y_1 \subseteq \phi(\mathfrak{C})$ and $\phi(\mathfrak{C})$ is finite, clearly Y_1 is definable and hence Y_2 is definable, as desired. \square

Lemma 5.2.15. *Let $f : X \longrightarrow Y$ be a definable function, where $X, Y \subseteq \text{VF}$. For every $a \in X$ let Z_a be the intersection of all definable subsets that contain a . Suppose that $f \upharpoonright Z_a$ is injective for every $a \in X$. Then there is a finite definable partition X_1, \dots, X_n of X such that $f \upharpoonright X_i$ is injective for every i .*

Proof. Let $\mathcal{L} = \mathcal{L}_{\text{RV}} \cup \{c_1, c_2\}$, where c_1, c_2 are new constants. Consider the \mathcal{L} -theory T that states the following:

1. everything in ACVF_S^0 ,
2. $c_1, c_2 \in X$ and $c_1 \neq c_2$,
3. $f(c_1) = f(c_2)$,
4. for every \mathcal{L}_{RV} -formula ϕ , either the subset defined by ϕ contains both c_1 and c_2 or it does not contain either of them.

If $N \models T$ then c_1^N, c_2^N are distinct elements in X and $c_1^N \in Z_{c_2^N}$ and $f(c_1^N) = f(c_2^N)$, contradiction. So T is inconsistent. So there are \mathcal{L}_{RV} -formulas ϕ_1, \dots, ϕ_n such that, for every two distinct elements $a_1, a_2 \in X$, if

$f(a_1) = f(a_2)$ then $\phi_i(\mathfrak{C})$ separates a_1, a_2 for some i . So the partition on X induced by $\phi_1(\mathfrak{C}), \dots, \phi_n(\mathfrak{C})$ is as desired. \square

Naturally injectivity may be replaced by other first-order properties in this lemma.

5.3 Quantifier elimination in ACVF

We shall show in this section that ACVF admits quantifier elimination. The following model-theoretic test for quantifier elimination will be used; see [48] for a proof.

Fact 5.3.1. *For any first-order theory T in a language that has at least one constant symbol, the following are equivalent:*

1. T admits quantifier elimination.
2. For any two models $M_1, M_2 \models T$ such that M_2 is $\|M_1\|^+$ -saturated and any isomorphism f between two substructures $N_1 \subseteq M_1$ and $N_2 \subseteq M_2$, there is a monomorphism $f^* : M_1 \rightarrow M_2$ extending f .

Recall that our strategy is to establish the second item in this test for ACVF via reduction to Theorem 5.2.5; see Remark 5.2.4.

Lemma 5.3.2. *Let $B \subseteq M \models \text{ACVF}$ and $b_0, \dots, b_n \in \text{VF}(B)$. Let $\bar{F}(X) = \sum_{0 \leq i \leq n} t_i X^i$ be a nonzero polynomial with coefficients in RV_0 such that $t_i = \text{rv}(b_i)$ if $t_i \neq 0$. Let $F(X) = \sum_{0 \leq i \leq n} b_i X^i$. For every $t \in \text{RV}(M)$, if $\bar{F}(t) = 0$ and $v_{\text{rv}}(\text{rv}(b_i)t^i) > 0$ for all $t_i = 0$, then there is a $b \in \text{rv}^{-1}(t)$ such that $F(b) = 0$.*

Proof. Fix a $t \in \text{RV}(M)$ with $\bar{F}(t) = 0$ and $v_{\text{rv}}(\text{rv}(b_i)t^i) > 0$ for all $t_i = 0$. Note that, since such a t exists and $\bar{F}(X)$ is not the zero polynomial, we must have that $\bar{F}(X)$ is not a monomial and $t \neq \infty$. Let $m < n$ be the least number such that $t_m \neq 0$. Let $r_1, \dots, r_n \in \text{VF}(M)$ be the (possibly repeated) roots of $F(X)$. Let $F^*(X) = \sum_{t_i \neq 0} b_i X^i$. For any $b \in t$, if $\text{rv}(b) \neq \text{rv}(r_i)$ for every i then $\text{val}(b - r_i) = \text{val}(b)$ if $\text{val}(b) < \text{val}(r_i)$ and $\text{val}(b - r_i) = \text{val}(r_i)$ if $\text{val}(b) \geq \text{val}(r_i)$. So $\prod_i \text{val}(b - r_i) \leq \text{val}(b_m b^m / b_n)$ and hence $\text{val}(F(b)) \leq \text{val}(b_m b^m) = 0$. Since $\text{val}(b_i b^i) > 0$ for all $t_i = 0$, we have that $\text{val}(F^*(b)) = 0$, contradicting the choice of t . So $t = \text{rv}(b) = \text{rv}(r_i)$ for some i . \square

Notation 5.3.3. For a polynomial $F(X) = \sum_i t_i X^i$ with coefficients $t_i \in \text{RV}_0$ it is often convenient to choose a $b_i \in t_i$ for each nonzero t_i and write $F(X)$ as $\sum_i \text{rv}(b_i) X^i$. Below, whenever $F(X)$ is written in this form, it should be understood that b_i is chosen only if $t_i \neq 0$.

For the rest of this section, we fix two models $M_1, M_2 \models \text{ACVF}$ such that M_2 is $\|M_1\|^+$ -saturated. Let $S_1 \subseteq M_1$ and $f_1 : S_1 \rightarrow M_2$ a monomorphism.

For any $A \subseteq M \models \text{ACVF}$, we write $\text{VF}(A)^{\text{ac}}$, $\overline{K}(A)^{\text{ac}}$, etc. for the corresponding field-theoretic algebraic closures.

Lemma 5.3.4. *There is a $P \subseteq M_1$ and a monomorphism $g : P \rightarrow M_2$ extending f_1 such that*

1. $\text{VF}(P) = \text{VF}(S_1)$,
2. $\overline{K}(P)$ is the algebraic closure of $\overline{K}(S_1)$,
3. $\Gamma(P)$ is the divisible hull of $\Gamma(S_1)$.

Proof. First of all, there is a field homomorphism $g_1 : \overline{K}(S_1)^{\text{ac}} \rightarrow \overline{K}(M_2)$ extending $f_1 \upharpoonright \overline{K}(S_1)$. Let $\langle \overline{K}(S_1)^{\text{ac}}, \text{RV}(S_1) \rangle = S_2$ and $g_2 : S_2 \rightarrow M_2$ be the monomorphism determined by

$$ts \mapsto g_1(t)f_1(s) \text{ for all } t \in \overline{K}(S_1)^{\text{ac}} \text{ and } s \in \text{RV}(S_1).$$

Next, let $n > 1$ be the least integer such that there is a $t_1 \in \text{RV}(M_1)$ with $t_1^n \in \text{RV}(S_2)$ but $v_{\text{rv}}(t_1) \notin \Gamma(S_2)$ for every $0 < i < n$. Let $t_2 \in \text{RV}(M_2)$ such that $g_2(t_1^n) = t_2^n$. Let $g_3 : \langle S_2, t_1 \rangle \rightarrow M_2$ be the monomorphism determined by

$$t_1 s \mapsto t_2 g_2(s) \text{ for all } s \in S_2.$$

Iterating this procedure the lemma follows. □

In the light of this lemma, without loss of generality, we may assume that $\overline{K}(S_1)$ is algebraically closed and $\Gamma(S_1)$ is divisible.

Let $S \subseteq M_1$ be a VF-generated substructure such that

1. $\text{VF}(S_1) \subseteq \text{VF}(S)$,
2. $\text{RV}(S) \subseteq \text{RV}(S_1)$,
3. there is a monomorphism $f : S \rightarrow M_2$ with $f \upharpoonright (S \cap S_1) = f_1 \upharpoonright (S \cap S_1)$.

Fix an $e \in \text{VF}(M_1)$ such that $\text{rv}(e) \in \text{RV}(S_1) \setminus \text{RV}(S)$. In the next few lemmas, under various assumptions, we shall prove the following claim:

Claim (*). $\text{RV}(\langle S, e \rangle) \subseteq \text{RV}(S_1)$ and f may be extended to a monomorphism $f^* : \langle S, e \rangle \rightarrow M_2$ such that $f^* \upharpoonright \text{RV}(\langle S, e \rangle) = f_1 \upharpoonright \text{RV}(\langle S, e \rangle)$.

Lemma 5.3.5. *Let $\overline{F}(x) = x^n + \sum_{0 \leq i < n} \text{rv}(a_i)x^i \in \overline{K}(S)[x]$ be an irreducible polynomial with $\text{rv}(a_0) \neq 0$. Suppose that $e \in \mathcal{U}(M_1)$ is a root of the polynomial $F(x) = x^n + \sum_{0 \leq i < n} a_i x^i \in \mathcal{O}(S)[x]$. If the valued field $(\text{VF}(S), \mathcal{O}(S))$ is henselian, then Claim (\star) holds.*

Proof. Obviously $\text{rv}(e)$ is a root of $\overline{F}(x)$. Also, note that $F(x)$ is irreducible over $\text{VF}(S)$. The polynomial

$$f_1(\overline{F}(x)) = x^n + \sum_{0 \leq i < n} f_1(\text{rv}(a_i))x^i \in f_1(\overline{K}(S))[x]$$

is irreducible over $f_1(\overline{K}(S))$ and $f_1(\text{rv}(e))$ is a root of $f_1(\overline{F}(x))$. By Lemma 5.3.2, there is a root $d \in \text{VF}(M_2)$ of $f_1(\overline{F}(x))$ such that $\text{rv}(d) = f_1(\text{rv}(e))$. By Remark 5.2.4, Theorem 5.2.5, and Fact 5.3.1, there is an \mathcal{L}_v -monomorphism $f^* : \langle S, e \rangle \rightarrow M_2$ extending f . Since $(\text{VF}(S), \mathcal{O}(S))$ is henselian, without loss of generality, we may assume that $f^*(e) = d$. By Remark 5.2.4 again, f^* may be treated as an \mathcal{L}_{RV} -monomorphism extending f with $f^*(\text{rv}(e)) = f_1(\text{rv}(e))$.

Now, since $[\overline{K}(\langle S, e \rangle) : \overline{K}(S)] = [\text{VF}(\langle S, e \rangle) : \text{VF}(S)]$, by the fundamental inequality of valuation theory (see [23, Theorem 3.3.4]), we have that

$$\begin{aligned} \overline{K}(\langle S, e \rangle) &= \overline{K}(S)(\text{rv}(e)) \subseteq \overline{K}(S_1) \\ \Gamma(\langle S, e \rangle) &= \Gamma(S). \end{aligned}$$

Therefore, $\text{RV}(\langle S, e \rangle) = \text{RV}(\langle \text{RV}(S), \text{rv}(e) \rangle) \subseteq \text{RV}(S_1)$, which clearly implies that $f^* \upharpoonright \text{RV}(\langle S, e \rangle) = f_1 \upharpoonright \text{RV}(\langle S, e \rangle)$. \square

Lemma 5.3.6. *Suppose that $e \notin \mathcal{U}(M_1)$, $e^n = a \in \text{VF}(S)$ for some integer $n > 1$, and $\text{val}(e^i) \notin \Gamma(S)$ for all $0 < i < n$. If $(\text{VF}(S), \mathcal{O}(S))$ is henselian, then Claim (\star) holds.*

Proof. By the fundamental inequality of valuation theory and the assumption, we have that

$$n \leq [\Gamma(\langle S, e \rangle) : \Gamma(S)] \leq [\text{VF}(\langle S, e \rangle) : \text{VF}(S)] \leq n.$$

So $n = [\Gamma(\langle S, e \rangle) : \Gamma(S)]$ and $\overline{K}(\langle S, e \rangle) = \overline{K}(S)$. Since $\Gamma(S_1)$ is divisible, we have that $\text{val}(e) \in \Gamma(S_1)$ and $\Gamma(\langle S, e \rangle) \subseteq \Gamma(S_1)$.

Any element $b \in \text{VF}(\langle S, e \rangle)$ may be written as a quotient of two elements of the form $\sum_{0 \leq i \leq m} b_i e^i$, where $b_i \in \text{VF}(S)$. Since $e^n = a \in \text{VF}(S)$, we may assume that $0 \leq m < n$.

Claim. For some $t \in \text{RV}(S)$ and some integer $0 \leq k \leq m$, $\text{rv}(b) = t \cdot \text{rv}(e^k)$.

Proof. We do induction on m . Without loss of generality, we may assume that $b_m, b_0 \neq 0$. We claim that $\text{val}(b_0) \neq \text{val}(e \sum_{j=1}^m b_j e^{j-1})$. Suppose for contradiction that this is not the case. By the inductive hypothesis, $\text{rv}(\sum_{j=1}^m b_j e^{j-1}) = t \cdot \text{rv}(e^k)$ for some $t \in \text{RV}(S)$ and some integer $0 \leq k \leq m-1$. So

$$\text{val} \left(e \sum_{j=1}^m b_j e^{j-1} \right) = v_{\text{rv}}(t \cdot \text{rv}(e^{k+1})) = v_{\text{rv}}(\text{rv}(b_0)).$$

So $\text{val}(e^{k+1}) \in \Gamma(S)$, which is a contradiction because $0 < k+1 < n$. Now, since $\text{val}(b_0) \neq \text{val}(e \sum_{j=1}^m b_j e^{j-1})$, either $\text{rv}(b) = \text{rv}(b_0)$ or $\text{rv}(b) = \text{rv}(e \sum_{j=1}^m b_j e^{j-1})$ and hence $\text{rv}(b)$ is of the desired form by the inductive hypothesis. \square

Therefore, $\Gamma(\langle S, e \rangle) = \Gamma(\langle \Gamma(S), \text{val}(e) \rangle)$ and $\text{RV}(\langle S, e \rangle) = \text{RV}(\langle \text{RV}(S), \text{rv}(e) \rangle) \subseteq \text{RV}(S_1)$.

Note that, since the roots of $F(x) = x^n - a$ are all of the same value, by the assumption on $\text{val}(e)$, $F(x)$ is irreducible over $\text{VF}(S)$. Let a_1, \dots, a_n be the distinct roots of $F(x)$ in M_1 . We consider the symmetric polynomial

$$G(y_1, \dots, y_n) = \prod_{i=1}^n \prod_{j=1}^n \left(y_j - \frac{\text{rv}(a_i)}{\text{rv}(a_j)} \right).$$

In the expansion of $G(y_1, \dots, y_n)$, the coefficient of each monomial is a sum of elements in the residue field and hence may be written as a quotient of two terms:

$$\frac{\text{rv}(I(a_1, \dots, a_n))}{\text{rv}(J(a_1, \dots, a_n))},$$

where $I(a_1, \dots, a_n)$ is a symmetric VF-sort term and hence may be written as $I(a)$. Moreover, if we substitute $y/\text{rv}(a_j)$ for y_j in each monomial then the denominator of its coefficient becomes $\text{rv}(\prod_i a_i)^n = \text{rv}(a^n)$. So the term $G(y/\text{rv}(a_1), \dots, y/\text{rv}(a_n))$ may be written as a summation $G(y, a)$ of terms of the form

$$\frac{\text{rv}(I(a))y^m}{\text{rv}(a^n)},$$

where $m \leq n^2$. Since $\text{RV}(\langle S, e \rangle) \subseteq \text{RV}(S_1)$, it makes sense to write

$$S_1 \models G(\text{rv}(e), a) = 0$$

and hence

$$f_1(S_1) \models G(f_1(\text{rv}(e)), f(a)) = 0.$$

So, by Lemma 5.3.2, there is a root $d \in \text{VF}(M_2)$ of the polynomial $x^n - f(a)$ such that $\text{rv}(d) = f_1(\text{rv}(e))$.

As in the previous lemma, there is an \mathcal{L}_v -monomorphism $f^* : \langle S, e \rangle \longrightarrow M_2$ extending f with $f^*(e) = d$, which may be treated as an \mathcal{L}_{RV} -monomorphism extending f with $f^*(\text{rv}(e)) = f_1(\text{rv}(e))$. Since $\text{RV}(\langle S, e \rangle) = \text{RV}(\langle \text{RV}(S), \text{rv}(e) \rangle)$, we must have that $f^* \upharpoonright \text{RV}(\langle S, e \rangle) = f_1 \upharpoonright \text{RV}(\langle S, e \rangle)$. \square

Lemma 5.3.7. *Suppose that $\text{rv}(e) \in \overline{K}(S_1)$ is transcendental over $\overline{K}(S)$. If $\Gamma(S)$ is divisible, then Claim (\star) holds.*

Proof. Clearly $\text{rv}(e)$ does not contain any element that is algebraic over $\text{VF}(S)$; in particular, e is transcendental over $\text{VF}(S)$. Similarly $f_1(\text{rv}(e))$ does not contain any element that is algebraic over $f(\text{VF}(S))$. Fix a $d \in \text{VF}(M_2)$ with $\text{rv}(d) = f_1(\text{rv}(e))$.

By the dimension inequality of valuation theory (see [23, Theorem 3.4.3]), the rational rank of $\Gamma(\langle S, e \rangle)/\Gamma(S)$ is 0. Since $\Gamma(S)$ is divisible, we actually have that $\Gamma(\langle S, e \rangle) = \Gamma(S)$. So for every $b \in \text{VF}(\langle S, e \rangle)$ there is an $a \in \text{VF}(S)$ such that $\text{val}(b/a) = 0$. Let $b = \sum_{0 \leq i \leq m} b_i e^i \in \text{VF}(\langle S, e \rangle)$, where $b_i \in \text{VF}(S)$, and $b^* = \sum_{0 \leq i \leq m} f(b_i) d^i \in \text{VF}(\langle f(S), d \rangle)$.

Claim. If $\text{val}(b) = 0$ then

1. $\text{rv}(b) \in \overline{K}(S)[\text{rv}(e)]$ and $\text{rv}(b^*) \in \overline{K}(f(S))[\text{rv}(d)]$,
2. $\text{val}(b^*) = 0$.

Proof. We do induction on m . Without loss of generality, we may assume that $b_m, b_0 \neq 0$. First of all, suppose that $\text{val}(b_0) \neq \text{val}(e \sum_{j=1}^m b_j e^{j-1})$. Then either $\text{val}(b) = \text{val}(b_0) = 0$ and $\text{val}(\sum_{j=1}^m b_j e^{j-1}) > 0$ or $\text{val}(b) = \text{val}(\sum_{j=1}^m b_j e^{j-1}) = 0$ and $\text{val}(b_0) > 0$. In the former case, let $a \in \text{VF}(S)$ be such that $\text{val}(a) = \text{val}(\sum_{j=1}^m b_j e^{j-1})$. By the inductive hypothesis, $\text{val}(\sum_{j=1}^m f(b_j/a) d^{j-1}) = 0$ and hence $\text{val}(d \sum_{j=1}^m f(b_j) d^{j-1}) > 0$. So $\text{val}(b^*) = \text{val}(f(b_0)) = 0$ and $\text{rv}(b^*) = \text{rv}(f(b_0)) \in \overline{K}(f(S))[\text{rv}(d)]$. In the latter case, by the inductive hypothesis, we have that $\text{val}(d \sum_{j=1}^m f(b_j) d^{j-1}) = 0$ and $\text{rv}(\sum_{j=1}^m f(b_j) d^{j-1}) \in \overline{K}(f(S))[\text{rv}(d)]$, which immediately imply that $\text{val}(b^*) = 0$ and

$$\text{rv}(b^*) = \text{rv} \left(d \sum_{j=1}^m f(b_j) d^{j-1} \right) \in \overline{K}(f(S))[\text{rv}(d)].$$

Similarly, for $\text{rv}(b)$, since either $\text{rv}(b) = \text{rv}(b_0)$ or $\text{rv}(b) = \text{rv}(e \sum_{j=1}^m b_j e^{j-1})$, clearly $\text{rv}(b)$ is of the desired form. Next, if $\text{val}(b_0) = \text{val}(e \sum_{j=1}^m b_j e^{j-1}) < 0$ then, since $\text{val}(b/b_0) > 0$, we have that $\text{val}(e \sum_{j=1}^m b_j e^{j-1}/b_0 + 1) > 0$ and hence

$$\text{rv}(e) \text{rv} \left(\sum_{j=1}^m \frac{b_j e^{j-1}}{b_0} \right) + 1 = 0.$$

By the inductive hypothesis, $\text{rv}(\sum_{j=1}^m b_j e^{j-1}/b_0) \in \overline{K}(S)[\text{rv}(e)]$. So the equality implies that $\text{rv}(e)$ is algebraic over $\overline{K}(S)$, contradiction. Now the only possibility left is that $\text{val}(b_0) = \text{val}(e \sum_{j=1}^m b_j e^{j-1}) = 0$. In this case,

$$\text{rv}(b) = \text{rv}(e) \text{rv} \left(\sum_{j=1}^m b_j e^{j-1} \right) + \text{rv}(b_0) \in \overline{K}(S)[\text{rv}(e)]$$

by the inductive hypothesis. For the second item, since $\text{val}(\sum_{j=1}^m f(b_j) d^{j-1}) = 0$ and $\text{rv}(\sum_{j=1}^m f(b_j) d^{j-1}) \in \overline{K}(f(S))[\text{rv}(d)]$, if $\text{val}(b^*) > 0$ then

$$\text{rv}(d) \text{rv} \left(\sum_{j=1}^m f(b_j) d^{j-1} \right) + \text{rv}(f(b_0)) = 0$$

and hence $\text{rv}(d)$ is algebraic over $\overline{K}(f(S))$, contradiction. So $\text{val}(b^*) = 0$ and hence

$$\text{rv}(b^*) = \text{rv}(d) \text{rv} \left(\sum_{j=1}^m f(b_j) d^{j-1} \right) + \text{rv}(f(b_0)) \in \overline{K}(f(S))[\text{rv}(d)]. \quad \square$$

Note that, symmetrically, the claim still holds if b is replaced by b^* . It follows that the embedding of the field $\text{VF}(\langle S, e \rangle)$ into the field $\text{VF}(M_2)$ determined by $e \mapsto d$ induces an \mathcal{L}_v -monomorphism $f^* : \langle S, e \rangle \rightarrow M_2$ extending f . As in the previous lemmas, f^* may be identified as an \mathcal{L}_{RV} -monomorphism. Since $f^*(\text{rv}(e)) = f_1(\text{rv}(e))$ and, by the claim, $\text{RV}(\langle S, e \rangle) = \text{RV}(\langle \text{RV}(S), \text{rv}(e) \rangle) \subseteq \text{RV}(S_1)$, we must have that $f^* \upharpoonright \text{RV}(\langle S, e \rangle) = f_1 \upharpoonright \text{RV}(\langle S, e \rangle)$. \square

Lemma 5.3.8. *Suppose that e is transcendental over $\text{VF}(S)$ and $\text{val}(e)$ is of infinite order modulo $\Gamma(S)$. Then for any $b = \sum_{0 \leq i \leq m} b_i e^i \in \text{VF}(\langle S, e \rangle)$, where $b_i \in \text{VF}(S)$, if $b \neq 0$ then $\text{val}(b) = \min \{ \text{val}(b_i e^i) : 0 \leq i \leq m \}$. Also, $\Gamma(\langle S, e \rangle)$ is the direct sum of $\Gamma(S)$ and the cyclic group generated by $\text{val}(e)$: $\Gamma(\langle S, e \rangle) = \Gamma(S) \oplus \langle \mathbb{Z} \cdot \text{val}(e) \rangle$.*

Proof. This is well-known; see, for example, [42, Lemma 4.8]. \square

Lemma 5.3.9. *If $\overline{K}(S) = \overline{K}(S_1)$ and $\Gamma(S)$ is divisible, then Claim (\star) holds.*

Proof. Note that, by the assumption, $e \notin \mathcal{U}(M_1)$, $\overline{K}(S)$ is algebraically closed, and $\text{val}(e) \notin \Gamma(S)$. Since $\Gamma(S)$ is divisible, clearly $\text{val}(e)$ is of infinite order modulo $\Gamma(S)$ and hence e is transcendental over $\text{VF}(S)$. Choose a $d \in \text{VF}(M_2)$ with $\text{rv}(d) = f_1(\text{rv}(e))$. Then d is transcendental over $f(\text{VF}(S))$. It is not hard to see that, by Lemma 5.3.8, the embedding of the field $\text{VF}(\langle S, e \rangle)$ into the field $\text{VF}(M_2)$ determined by $e \mapsto d$ induces an \mathcal{L}_v -monomorphism $f^* : \langle S, e \rangle \rightarrow M_2$ extending f , which, as above, is identified as an \mathcal{L}_{RV} -monomorphism $f^* : \langle S, e \rangle \rightarrow M_2$ extending f . Now, since the rational rank of $\Gamma(\langle S, e \rangle)/\Gamma(S)$ is nonzero and $\overline{K}(S)$ is

algebraically closed, by the dimension inequality of valuation theory, we have that $\overline{K}(\langle S, e \rangle) = \overline{K}(S)$. By Lemma 5.3.8 again, $\Gamma(\langle S, e \rangle) = \Gamma(S) \oplus (\mathbb{Z} \cdot \text{val}(e))$, that is, $\Gamma(\langle S, e \rangle) = \Gamma(\langle \Gamma(S), \text{val}(e) \rangle)$. So $\text{RV}(\langle S, e \rangle) = \text{RV}(\langle \text{RV}(S), \text{rv}(e) \rangle) \subseteq \text{RV}(S_1)$ and hence $f^* \upharpoonright \text{RV}(\langle S, e \rangle) = f_1 \upharpoonright \text{RV}(\langle S, e \rangle)$. \square

Proposition 5.3.10. *There is a monomorphism $f_1^* : M_1 \rightarrow M_2$ extending f_1 .*

Proof. First of all, since the henselization L of $(\text{VF}(S_1), \mathcal{O}(S_1))$ in M_1 is an immediate extension (in the sense of valuation theory), we have that $\text{RV}(\langle L, S_1 \rangle) = \text{RV}(S_1)$. So we may assume that f_1 is a monomorphism from $\langle L, S_1 \rangle$ into M_2 . Now we use Lemma 5.3.5 to extend $f_1 \upharpoonright L$ to $f_2 : S_2 \rightarrow M_2$ by adding all the elements in $\overline{K}(S_1)$ that are algebraic over $\overline{K}(L)$. Manifestly $\overline{K}(S_2)$ is algebraically closed. Then, starting with the least n such that there is a $\gamma \in \Gamma(S_2)$ that is not divisible by n , we use Lemma 5.3.6 to extend f_2 to $f_3 : S_3 \rightarrow M_2$ such that $\Gamma(S_3)$ is divisible. Note that, by the proof of Lemma 5.3.6, $\overline{K}(S_3) = \overline{K}(S_2)$. Next, we use Lemma 5.3.7 to extend f_3 to $f_4 : S_4 \rightarrow M_2$ by adding an element in $\overline{K}(S_1)$ that is transcendental over $\overline{K}(S_3)$. Iterating this procedure we may exhaust all elements in $\overline{K}(S_1)$ and hence obtain a monomorphism $f_5 : S_5 \rightarrow M_2$ such that S_5 satisfies the assumption of Lemma 5.3.9. Then, a combined application of henselization, Lemma 5.3.6, and Lemma 5.3.9 eventually brings a monomorphism $f^* : S^* \rightarrow M_2$ such that $f_1 \subseteq f^*$ and S^* is VF-generated. In this case, the proposition follows from Remark 5.2.4, Theorem 5.2.5, and Fact 5.3.1. \square

This proposition and Fact 5.3.1 immediately yields:

Theorem 5.3.11. *The theory ACVF admits quantifier elimination.*

5.4 Basic structural properties

From this section forth the background assumption is resumed: we work in a monster model \mathfrak{C} of ACVF_S^0 , where S is a fixed “small” substructure of \mathfrak{C} .

Although its proof only involves elementary calculations, the following simple lemma is vital to the inductive arguments below. Its failure when $\text{char } \overline{K} > 0$ is one of the major obstacles for generalizing the Hrushovski-Kazhdan integration theory to valued fields of positive residue characteristics.

Lemma 5.4.1. *Let $c_1, \dots, c_k \in \text{VF}$ be distinct elements of the same value α such that their average is 0. Then for some $c_i \neq c_j$ we have that $\text{val}(c_i - c_j) = \alpha$ and hence rv is not constant on the set $\{c_1, \dots, c_k\}$.*

Proof. Suppose for contradiction that $\text{val}(c_i - c_j) > \alpha$ for all $c_i \neq c_j \in A$. Since $c_1 = -(c_2 + \dots + c_k)$ and

$\text{char } \overline{K} = 0$, we have that

$$\alpha = \text{val}(kc_1) = \text{val}((k-1)c_1 - (c_2 + \dots + c_k)) = \text{val}\left(\sum_{i=2}^k (c_1 - c_i)\right) > \alpha,$$

contradiction. \square

Definition 5.4.2. Let A be a definable subset of VF^m . A *definable auxiliary projection* of A is a definable function of A of the form

$$(x_1, \dots, x_m) \mapsto (\text{rv}(g_1), \dots, \text{rv}(g_k)),$$

where each $g_i : A \rightarrow \text{VF}$ is a definable function.

Lemma 5.4.3. *Let A be a definable finite subset of VF^n . Then there is a definable injective auxiliary projection of A .*

Proof. We do double induction on n and the number k of elements in A . For $n = 1$, let $A = \{c_1, \dots, c_k\} \subseteq \text{VF}$. Let $c = (\sum_{i=1}^k c_i)/k$ be the average of A . Then there is a definable bijective function from A onto $\{c_1 - c, \dots, c_k - c\}$. So we may assume that the average of A is 0. Since the set $\text{val}(A)$ is finite, for each $\gamma \in \text{val}(A)$, the set $A \cap \text{val}^{-1}(\gamma)$ is definable. So by the inductive hypothesis we may also assume that val is constant on A ; say, $\text{val}(c_i) = \alpha$ for all $c_i \in A$. By Lemma 5.4.1, rv is not constant on A , that is, $1 < |\text{rv}(A)| \leq k$. So $1 \leq |\text{rv}^{-1}(t) \cap A| < k$ for each $t \in \text{rv}(A)$. By the inductive hypothesis there is a $\langle t \rangle$ -definable injective auxiliary projection f_t of $\text{rv}^{-1}(t) \cap A$ for each $t \in \text{rv}(A)$. It is easy to see that for each f_t there is a definable rv -function f_t^* on a subset of A such that $f_t^*(c_i) = (t, f_t(c_i))$ for each $c_i \in \text{rv}^{-1}(t) \subseteq \text{dom}(f_t^*)$. Also, the collection of rv -functions f of A with $\text{ran}(f) \subseteq \text{RV}^m$ for some m is rv -closed. Applying Lemma 5.2.11 we obtain a definable injective auxiliary projection of A .

Now suppose that $n > 1$. By the inductive hypothesis, there is a definable injective auxiliary projection g of $\text{pr}_n(A)$ and, for each $c \in \text{pr}_n(A)$, a c -definable injective auxiliary projection f_c of $\text{fib}(A, c)$. As above, for each f_c ,

1. there is a definable $(g \circ \text{pr}_n)$ -function f_c^* on a subset of A such that $f_c^*(\bar{c}_i) = ((g \circ \text{pr}_n)(\bar{c}_i), f_c(\bar{c}_i))$ for each $\bar{c}_i \in \text{fib}(A, c)$,
2. the collection of $(g \circ \text{pr}_n)$ -functions f of A with $\text{ran}(f) \subseteq \text{RV}^m$ for some m is $(g \circ \text{pr}_n)$ -closed.

Applying Lemma 5.2.11 we obtain a definable injective auxiliary projection of A . \square

Note that this proof has nothing to do with algebraic closedness and hence works for the theory of valued fields as naturally formulated in \mathcal{L}_{RV} .

The role of balls in a motivic measure on a valued field is similar to that of intervals in the Lebesgue measure on the real line. We begin the study of balls with a list of easily seen properties.

Remark 5.4.4. Let \mathfrak{a} be an open ball and \mathfrak{b} a ball.

1. For any $c \in \text{VF}$, the subset $\mathfrak{a} - c = \{a - c : a \in \mathfrak{a}\}$ is an open ball. If $c \in \mathfrak{a}$ then $\text{vcr}(\mathfrak{a} - c) = \infty$ and $\text{rad}(\mathfrak{a} - c) = \text{rad}(\mathfrak{a})$ and $\mathfrak{a} - c$ is a union of rv-balls. If $c \notin \mathfrak{a}$ and $\text{val}(c) \leq \text{rad}(\mathfrak{a})$ then $\text{vcr}(\mathfrak{a} - c) \leq \text{rad}(\mathfrak{a} - c) = \text{rad}(\mathfrak{a})$. If $c \notin \mathfrak{a}$ and $\text{val}(c) > \text{rad}(\mathfrak{a})$ then $\mathfrak{a} - c = \mathfrak{a}$.
2. $0 \notin \mathfrak{a}$ if and only if \mathfrak{a} is contained in an rv-ball if and only if $\text{vcr}(\mathfrak{a}) \neq \infty$ if and only if $\text{rad}(\mathfrak{a}) \geq \text{vcr}(\mathfrak{a})$.
3. The average of finitely many elements in \mathfrak{a} is in \mathfrak{a} , which fails if $\text{char}(\overline{K}) > 0$.
4. For any $c_1, c_2 \in \text{VF}$, $(\mathfrak{a} - c_1) \cap (\mathfrak{a} - c_2) \neq \emptyset$ if and only if $\mathfrak{a} - c_1 = \mathfrak{a} - c_2$ if and only if $\text{val}(c_1 - c_2) > \text{rad}(\mathfrak{a})$.
5. If $\mathfrak{a} \cap \mathfrak{b} = \emptyset$ then $\text{val}(a - b) = \text{val}(a' - b')$ for all $a, a' \in \mathfrak{a}$ and $b, b' \in \mathfrak{b}$. The subset $\mathfrak{a} - \mathfrak{b} = \{a - b : a \in \mathfrak{a} \text{ and } b \in \mathfrak{b}\}$ is a ball that does not contain 0. In fact, for any $a \in \mathfrak{a}$ and $b \in \mathfrak{b}$, either $\mathfrak{a} - \mathfrak{b} = \mathfrak{a} - b$ or $\mathfrak{a} - \mathfrak{b} = a - \mathfrak{b}$.
6. Suppose that $\mathfrak{a} \cap \mathfrak{b} = \emptyset$. Let \mathfrak{c} be the smallest closed ball that contains \mathfrak{a} . Clearly $\text{vcr}(\mathfrak{c}) = \text{vcr}(\mathfrak{a})$ and $\text{rad}(\mathfrak{c}) = \text{rad}(\mathfrak{a})$. If \mathfrak{b} is a maximal open subball of \mathfrak{c} , that is, if \mathfrak{b} is an open ball contained in \mathfrak{c} with $\text{rad}(\mathfrak{b}) = \text{rad}(\mathfrak{c})$, then $\mathfrak{a} - \mathfrak{b}$ is an rv-ball $\text{rv}^{-1}(t)$ with $\text{val}(t) = \text{rad}(\mathfrak{a})$. This means that the collection of maximal open subballs of \mathfrak{c} admits a \overline{K} -affine structure.
7. Let $f(x)$ be a polynomial with coefficients in VF and d_1, \dots, d_n the roots of $f(x)$. Suppose that \mathfrak{a} is contained in an rv-ball and does not contain any d_i . Then each $\mathfrak{a} - d_i$ is contained in an rv-ball and hence $f(\mathfrak{a})$ is contained in an rv-ball, that is, $(\text{rv} \circ f)(\mathfrak{a})$ is a singleton.

Similar properties are available if \mathfrak{a} is a closed ball.

Definition 5.4.5. A subset X of VF is a *punctured (open, closed, rv-) ball* if $X = \mathfrak{b} \setminus \bigcup_{i=1}^n \mathfrak{h}_i$, where \mathfrak{b} is an (open, closed, rv-) ball, $\mathfrak{h}_1, \dots, \mathfrak{h}_n$ are disjoint balls, and $\mathfrak{h}_i, \dots, \mathfrak{h}_n \subseteq \mathfrak{b}$. Each \mathfrak{h}_i is a *hole* of X . The *radius* and the *valuative center* of X are those of \mathfrak{b} . A subset Y of VF is a *simplex* if it is a finite union of disjoint balls and punctured balls of the same radius and the same valuative center, which are defined to be the *radius* and the *valuative center* of Y and are denoted by $\text{rad}(Y)$ and $\text{vcr}(Y)$.

A special kind of simplex is called a *thin annulus*: it is a punctured closed ball \mathfrak{b} with a single hole \mathfrak{h} such that \mathfrak{h} is a maximal open ball contained in \mathfrak{b} . For example, an element $\gamma \in \Gamma$ may be regarded as a thin annulus: it is the punctured closed ball with radius γ and valuative center ∞ and the special maximal open ball containing 0 removed.

Remark 5.4.6. The theory ACVF^0 is C -minimal; that is, every parametrically definable subset of VF is a boolean combination of balls. This basically follows from [39, Theorem 4.11] and the easy fact that any subset of VF that is parametrically definable in \mathcal{L}_{RV} is also parametrically definable in the two-sorted language \mathcal{L}_v for valued fields. Hence, for any parametrically definable subset X of VF , there are disjoint balls and punctured balls $\mathbf{a}_1, \dots, \mathbf{a}_l$ obtained from a unique set of balls $\mathbf{b}_1, \dots, \mathbf{b}_n, \mathbf{h}_1, \dots, \mathbf{h}_m$ such that $X = \bigcup_i \mathbf{b}_i \setminus \bigcup_j \mathbf{h}_j$. If we group $\mathbf{a}_1, \dots, \mathbf{a}_l$ by their radii and valuative centers then X may also be regarded as the union of a unique set of disjoint parametrically definable simplexes. Each \mathbf{b}_i is a *positive boolean component* of X and each \mathbf{h}_j is a *negative boolean component* of X . It follows that, as imaginary definable subsets, Γ is o -minimal and the set of maximal open balls contained in a closed ball is strongly minimal.

Definition 5.4.7. Let $\mathbf{b}_1, \dots, \mathbf{b}_n$ be the positive boolean components of a subset $X \subseteq \text{VF}$. The *positive closure* of X is the set of the minimal closed balls $\{\mathbf{c}_1, \dots, \mathbf{c}_m\}$ such that each \mathbf{c}_i contains some \mathbf{b}_j .

Note that, if $X \subseteq \text{VF}$ is definable from a set of parameters then its positive closure is definable from the same set of parameters.

Lemma 5.4.3 is of fundamental importance in the Hrushovski-Kazhdan theory. Other structural properties of functions between or within the two sorts will also be needed below. For example:

Lemma 5.4.8. *Let W be a definable subset of RV^m and $f : W \rightarrow \text{VF}^n$ a definable function. Then $f(W)$ is finite.*

Proof. The proof is by induction on n . For the base case $n = 1$, suppose for contradiction that $f(W)$ is infinite. By C -minimality, $f(W)$ is a union of disjoint balls and punctured balls $\mathbf{b}_1, \dots, \mathbf{b}_l$ such that $\text{rad } \mathbf{b}_i < \infty$ for some i , say \mathbf{b}_1 . Let ϕ be a formula that defines f . By quantifier elimination, ϕ may be assumed to be a disjunction of conjunctions of literals. Since $f(W)$ is infinite, there is at least one disjunct in ϕ , say ϕ^* , that does not have an irredundant VF -sort equality as a conjunct. Fix a $b \in \mathbf{b}_1$ and a $\bar{t} \in W$ such that the pair (\bar{t}, b) satisfies ϕ^* . For any term $\text{rv}(g(x))$ in ϕ^* , where $g(x) \in \text{VF}(\langle \emptyset \rangle)[x]$, and any $d \in \text{VF}$, if $\text{val}(d - b)$ is sufficiently large then $\text{rv}(g(b)) = \text{rv}(g(d))$. So there is a $d \in \mathbf{b}_1$ such that the pair (\bar{t}, d) also satisfies ϕ^* , which is a contradiction as f is a function. In general, if $n > 1$, by the inductive hypothesis both $\text{pr}_1 \circ f(W)$ and $\text{pr}_{>1} \circ f(W)$ are finite, hence $f(W)$ is finite. \square

Lemma 5.4.9. *Let $\mathbf{b} \subseteq \text{VF}$ be a ball such that $\mathbf{b} \cap \text{VF}(\text{acl}(\emptyset)) = \emptyset$. For any definable function $f : X \rightarrow \text{RV}^n$ with $\mathbf{b} \subseteq X$, $f \upharpoonright \mathbf{b}$ is constant.*

Proof. Clearly it is enough to show the case $n = 1$. Let ϕ be a quantifier-free formula in disjunctive normal form that determines $f \upharpoonright \mathbf{b}$. We may assume that no disjunct of ϕ is redundant and hence ϕ does not

contain any VF-sort literal. For any term $\text{rv}(g(x))$ in ϕ , where $g(x) \in \text{VF}(\langle\emptyset\rangle)[x]$, and any root b of $g(x)$, since $b \in \text{VF}(\text{acl}(\emptyset))$, we have that $b \notin \mathfrak{b}$ and hence there is a $t \in \text{RV}$ such that $\mathfrak{b} - b \subseteq \text{rv}^{-1}(t)$. So $\text{rv}(g(a_1)) = \text{rv}(g(a_2))$ for all $a_1, a_2 \in \mathfrak{b}$. It follows that $|f(\mathfrak{b})| = 1$. \square

Lemma 5.4.10. *Let $f : X \rightarrow Y$ be a definable surjective function, where $X, Y \subseteq \text{VF}$. Then there is a definable function $P : X \rightarrow \text{RV}^m$ such that, for each $\bar{t} \in \text{ran } P$, $f \upharpoonright P^{-1}(\bar{t})$ is either constant or injective.*

Proof. Let Y_1, Y_2 be a partition of Y given by Lemma 5.2.14. By Lemma 5.4.3, there is an injective function from Y_1 into RV^l for some l . The same holds for every $f^{-1}(b)$ with $b \in Y_2$. So the lemma follows from compactness. \square

Definition 5.4.11. Let \mathfrak{B} be a finite definable set of (open, closed, rv-) balls $\mathfrak{b}_1, \dots, \mathfrak{b}_n$. We call \mathfrak{B} an *algebraic set of balls*, $\bigcup \mathfrak{B}$ an *algebraic union of balls*, and each \mathfrak{b}_i an *algebraic (open, closed, rv-) ball*. If there is a definable subset C of $\bigcup \mathfrak{B}$ and a definable surjective function $f : \mathfrak{B} \rightarrow C$ such that $f(\mathfrak{b}_i) \in \mathfrak{b}_i$ for every $\mathfrak{b}_i \in \mathfrak{B}$ then we say that \mathfrak{B} has *definable centers* and C is an *definable set of centers of \mathfrak{B}* .

It is not hard to see that, if S is VF-generated and X is a $\bar{\gamma}$ -definable subset of VF^n , then X is $\bar{\gamma}$ -definable in the two-sorted language \mathcal{L}_v .

Lemma 5.4.12. *Suppose that S is VF-generated and $\bar{\gamma} \in \Gamma$. Let X be a $\bar{\gamma}$ -algebraic union of disjoint balls $\mathfrak{b}_1, \dots, \mathfrak{b}_n$. Then there is a disjunction of VF-sort equalities $\bigvee_j F_j(x) = 0$, where $F_j(x) \in \text{VF}(\langle\emptyset\rangle)[x]$, such that $(\bigvee_j F_j(\text{VF}) = 0) \cap \mathfrak{b}_i \neq \emptyset$ for each \mathfrak{b}_i .*

Proof. Without loss of generality we may assume that $\text{rad } \mathfrak{b}_i < \infty$ and $0 \notin \mathfrak{b}_i$ for each \mathfrak{b}_i , that is, each \mathfrak{b}_i is an infinite subset and is contained in an rv-ball. Let ϕ be an \mathcal{L}_v -formula such that $\phi(\bar{a}, \bar{\gamma})$ defines X , where $\bar{a} \in \text{VF}(\langle\emptyset\rangle)$. By Theorem 5.2.5, we may assume that ϕ is quantifier-free and is written in disjunctive normal form. If ϕ does not contain any Γ -sort literal then each disjunct of ϕ must contain a VF-sort equality. In this case the lemma is clear. So let us assume that some disjunct of ϕ contains an irredundant Γ -sort literal and also lacks VF-sort equality. Let $\Gamma_{\bar{\gamma}}$ be the substructure of Γ generated by $\bar{\gamma}$. Each Γ -sort literal in ϕ is of the form

$$\text{val } F(x) \square \text{val } G(x) + \xi,$$

where $F(x), G(x) \in \text{VF}(\langle\emptyset\rangle)[x]$, $\xi \in \Gamma_{\bar{\gamma}}$, and \square is one of the symbols $=, \neq, \leq,$ and $>$. Let $F_j(x)$ enumerate all polynomials in $\text{VF}(\langle\emptyset\rangle)[x]$ that occur in the literals in ϕ .

We claim that $\bigvee_j F_j(x) = 0$ is as required. Suppose for contradiction that this is not the case, say $(\bigvee_j F_j(\text{VF}) = 0) \cap \mathfrak{b}_1 = \emptyset$. Let R_j be the set of the roots of $F_j(x)$. For each $r \in \bigcup R_j$, since $r \notin \mathfrak{b}_1$, we have

that $\text{vcr}(\mathfrak{b}_1 - r) < \text{rad } \mathfrak{b}_1 \leq \infty$ if \mathfrak{b}_1 is a closed ball and $\text{vcr}(\mathfrak{b}_1 - r) \leq \text{rad } \mathfrak{b}_1 < \infty$ if \mathfrak{b}_1 is an open ball. So there is a $d \in \text{VF} \setminus X$ such that

1. $\text{val}(d) = \text{vcr}(\mathfrak{b}_1)$,
2. $\max \{ \text{vcr}(\mathfrak{b}_1 - r) : r \in \{0\} \cup \bigcup R_j \} \leq \text{vcr}(\mathfrak{b}_1 - d) \leq \text{rad } \mathfrak{b}_1$,
3. $\text{vcr}(\mathfrak{b}_1 - r) = \text{val}(d - r)$ for each $r \in \bigcup R_j$,
4. d satisfies all VF-sort disequalities in ϕ .

Since \mathfrak{b}_1 is an infinite subset, there is a $b \in \mathfrak{b}_1$ such that b satisfies a disjunct ϕ' of ϕ and ϕ' lacks VF-sort equality. Then d also satisfies ϕ' , contradiction. \square

Corollary 5.4.13. *Suppose that S is VF-generated. If $\Gamma(\text{acl}(\emptyset))$ is nontrivial then $\text{acl}(\emptyset)$ is a model of ACVF_S^0 .*

Lemma 5.4.14. *Suppose that S is VF-generated. Let $\bar{\gamma} \in \Gamma$ and \mathfrak{B} a $\bar{\gamma}$ -algebraic set of balls $\mathfrak{b}_1, \dots, \mathfrak{b}_n$. Then \mathfrak{B} has $\bar{\gamma}$ -definable centers.*

Proof. The set \mathfrak{B} may be partitioned into subsets $\mathfrak{B}_1, \dots, \mathfrak{B}_m \subseteq \mathfrak{B}$ such that each \mathfrak{B}_i is an $\bar{\gamma}$ -algebraic set of disjoint balls. So without loss of generality we may assume that \mathfrak{B} is a set of disjoint balls. By Lemma 5.4.12, there is an algebraic subset C of VF such that $C \cap \mathfrak{b}_i \neq \emptyset$ for every i . So the set \mathfrak{B} gives rise to a partition of C and the set of the averages of the parts of this partition is $\bar{\gamma}$ -definable. Since $\text{char } \bar{K} = 0$, the corresponding average remains in each \mathfrak{b}_i . \square

Lemma 5.4.15. *If \mathfrak{B} is a parametrically definable infinite set of closed balls then there is a parametrically definable map of \mathfrak{B} onto a proper interval of Γ .*

Proof. Since Γ is o -minimal, any parametrically definable infinite subset of Γ contains an interval. Therefore it suffices to show that there is a parametrically definable map of \mathfrak{B} into Γ whose image is infinite. If either the subset $\{\text{rad } \mathfrak{b} : \mathfrak{b} \in \mathfrak{B}\}$ is infinite or the subset $\{\text{vcr } \mathfrak{b} : \mathfrak{b} \in \mathfrak{B}\}$ is infinite then clearly such a map exists. So, without loss of generality, we may assume that both rad and vcr are constant on \mathfrak{B} . Since \mathfrak{B} is infinite, obviously $\text{vcr } \mathfrak{B} \neq \infty$. Now, by C -minimality, the subset $\text{pr}_1 \mathfrak{B}$ is a finite union of disjoint balls $\mathfrak{b}_1, \dots, \mathfrak{b}_n$, some of which may be punctured. Clearly $\text{vcr } \mathfrak{b}_i = \text{vcr } \mathfrak{B}$ for every \mathfrak{b}_i . Since every $\mathfrak{b} \in \mathfrak{B}$ is closed and \mathfrak{B} is infinite, we must have that $\text{rad } \mathfrak{b} > \text{rad } \mathfrak{b}_i$ for some \mathfrak{b}_i , say \mathfrak{b}_1 . Choose a $c \in \mathfrak{b}_1$ such that the open ball $\{x \in \text{VF} : \text{val}(x - c) > \text{rad } \mathfrak{b}_1\}$ is contained in \mathfrak{b}_1 . Clearly the subset

$$\{\text{vcr}(\mathfrak{b} - c) : \mathfrak{b} \in \mathfrak{B} \text{ and } \mathfrak{b} \subseteq \mathfrak{b}_1\}$$

is infinite. Hence the parametrically definable map of \mathfrak{B} given by $\mathfrak{b} \mapsto \text{vcr}(\mathfrak{b} - c)$ is as desired. \square

Lemma 5.4.16. *Suppose that S is (VF, Γ) -generated. Let $\bar{t} = (t_1, \dots, t_n) \in \text{RV}$ and \mathfrak{B} a \bar{t} -algebraic set of closed balls. Then \mathfrak{B} has \bar{t} -definable centers.*

Proof. The proof is by induction on n . The base case $n = 0$ is covered by Lemma 5.4.14. We proceed to the inductive step. First note that for any $\gamma \in \Gamma$ the subset $A_\gamma = \{t \in \text{RV} : \text{v}_{\text{rv}}(t) = \gamma\}$ is strongly minimal. Let ϕ be a formula that defines \mathfrak{B} . Let $\text{v}_{\text{rv}}(t_1) = \gamma_1$. For any $\bar{s} = (s_1, t_2, \dots, t_{n+1})$ with $\text{v}_{\text{rv}}(s_1) = \gamma_1$, let $W_{\bar{s}} \subseteq \text{VF}^3$ be the subset defined by $\phi(\bar{s})$. Let $\mathfrak{B}_{\bar{s}} = W_{\bar{s}}$ if $W_{\bar{s}}$ is a finite set of closed balls; otherwise $\mathfrak{B}_{\bar{s}} = \emptyset$. Consider the set of closed balls $\mathfrak{D} = \bigcup \mathfrak{B}_{\bar{s}}$, which contains \mathfrak{B} , and the subset

$$D = \bigcup \{ \{\bar{s}\} \times \mathfrak{B}_{\bar{s}} : \bar{s} \in \gamma_1 \times \{(t_2, \dots, t_{n+1})\} \},$$

both of which are $\langle \gamma_1, t_2, \dots, t_{n+1} \rangle$ -definable. We claim that \mathfrak{D} is finite. Suppose for contradiction that \mathfrak{D} is infinite. Since any two disjoint parametrically definable infinite subsets of \mathfrak{D} would give rise to two disjoint parametrically definable infinite subsets of A_{γ_1} , which is a contradiction as A_{γ_1} is strongly minimal, we deduce that \mathfrak{D} is strongly minimal. By Lemma 5.4.15, there is a parametrically definable map of \mathfrak{D} onto an interval of Γ , which must be strongly minimal as well. However, the ordering of Γ is linear and dense, and hence no interval of Γ is strongly minimal, contradiction. So \mathfrak{D} is finite. Applying the inductive hypothesis with respect to the substructure $\langle \gamma_1 \rangle$ and the tuple (t_2, \dots, t_{n+1}) , we conclude that \mathfrak{B} has \bar{t} -definable centers. \square

Lemma 5.4.17. *For any $t \in \text{RV}$, if $\text{rv}^{-1}(t)$ has a definable proper subset then it has definable center.*

Proof. Let X be a definable proper subset of $\text{rv}^{-1}(t)$. Let $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ be the positive boolean components of X and $\mathfrak{h}_1, \dots, \mathfrak{h}_m$ the negative boolean components of X . Since X is a proper subset of $\text{rv}^{-1}(s)$, at least one of these balls is a proper subball of $\text{rv}^{-1}(s)$ and hence its positive closure is also a proper subball of $\text{rv}^{-1}(s)$. Then, by Lemma 5.4.16, there is a definable finite subset of $\text{rv}^{-1}(s)$ and hence, by taking the average, a definable point in $\text{rv}^{-1}(s)$. \square

5.5 Parametric balls and atomic subsets

In this section let Q be a set of parameters that consists of balls of radius $< \infty$. Without loss of generality, we may assume that no ball in Q is definable.

Definition 5.5.1. A subset X generates a complete Q -type if for all Q -definable subset Y either $X \subseteq Y$ or $X \cap Y = \emptyset$. An Q -definable subset X is *atomic over* $\langle Q \rangle$ if it generates a complete Q -type.

Lemma 5.5.2. *Let T be an Q -definable set of balls and ϕ a formula such that, for all $t_1 \neq t_2 \in T$, $\phi(t_1)$ and $\phi(t_2)$ define two disjoint balls \mathfrak{b}_{t_1} and \mathfrak{b}_{t_2} . For each $t \in T$, if \mathfrak{b}_t is not Q -algebraic then it is atomic over $\langle Q, t \rangle$.*

Proof. Suppose for contradiction that there is a non- Q -algebraic \mathfrak{b}_s and a formula ψ such that $\psi(s)$ defines a proper subset of \mathfrak{b}_s . For each $t \in T$, let X_t be the set defined by $\psi(t)$ if it is a proper subset of \mathfrak{b}_t and $X_t = \emptyset$ otherwise. Set $X = \bigcup_{t \in T} X_t$, which is Q -definable. By C -minimality, X is a boolean combination of some balls $\mathfrak{d}_1, \dots, \mathfrak{d}_n$. Since the balls \mathfrak{b}_t are pairwise disjoint, there are only finitely many balls \mathfrak{b}_t that contain some \mathfrak{d}_i . Note that this finite collection of balls is Q -definable, which does not contain \mathfrak{b}_s since \mathfrak{b}_s is not Q -algebraic. On the other hand, since $\mathfrak{b}_s \cap X \neq \emptyset$, we must have that $\mathfrak{b}_s \subseteq X$. This is a contradiction because the balls \mathfrak{b}_t being pairwise disjoint implies that $\mathfrak{b}_s \cap X$ is a proper subset of \mathfrak{b}_s . \square

Lemma 5.5.3. *Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be atomic over $\langle Q \rangle$ and $\bar{\gamma} \in \Gamma$. Then X is atomic over $\langle Q, \bar{\gamma} \rangle$.*

Proof. By induction this is immediately reduced to the case that the length of $\bar{\gamma}$ is 1. Suppose for contradiction that there is a formula $\psi(\gamma)$ that defines a proper subset of X . Then the subset $\Delta = \{\gamma \in \Gamma : \psi(\gamma) \text{ defines a proper subset of } X\}$ is nonempty and is Q -definable. By σ -minimality, some $\alpha \in \Delta$ is Q -definable, contradicting the assumption that X is atomic over $\langle Q \rangle$. \square

Definition 5.5.4. Let \mathfrak{b}_1 and \mathfrak{b}_2 be two (punctured) balls. We say that they are of the same *haecceitistic type* if

1. $\text{rad}(\mathfrak{b}_1) = \text{rad}(\mathfrak{b}_2)$ and $\text{vcr}(\mathfrak{b}_1) = \text{vcr}(\mathfrak{b}_2)$,
2. they are both open balls or both closed balls or both thin annuli.

Lemma 5.5.5. *Let $X \subseteq \text{VF}$ be atomic over $\langle Q \rangle$. Then X is the union of disjoint balls $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ of the same haecceitistic type.*

Proof. By C -minimality, X is a union of disjoint balls $\mathfrak{b}_1, \dots, \mathfrak{b}_n$, some of which may be punctured. First of all, since X is atomic, both vcr and rad must be constant on $\{\mathfrak{b}_1, \dots, \mathfrak{b}_n\}$, because otherwise there would be an Q -definable proper subset of X according to $\min\{\text{vcr}(\mathfrak{b}_1), \dots, \text{vcr}(\mathfrak{b}_n)\}$ or $\min\{\text{rad}(\mathfrak{b}_1), \dots, \text{rad}(\mathfrak{b}_n)\}$. Similarly either $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ are all closed balls or are all open balls. Also, since the subset of X that contains exactly every unpunctured ball \mathfrak{b}_i is definable, we have that either $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ are all punctured or are all unpunctured.

So it is enough to show that if \mathfrak{b}_i is punctured then it must be a thin annulus. By atomicity again, if $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ are punctured then each \mathfrak{b}_i must contain the same number of holes. If \mathfrak{b}_i has a hole \mathfrak{h} with

$\text{rad}(\mathfrak{h}) < \text{rad}(\mathfrak{b}_i)$ then $\mathfrak{b}_i \setminus \mathfrak{h}^*$ is nonempty, where \mathfrak{h}^* is the closed ball that has radius $(\text{rad}(\mathfrak{b}_i) + \text{rad}(\mathfrak{h}))/2$ and contains \mathfrak{h} . The collection of all such holes $\mathfrak{h}_1, \dots, \mathfrak{h}_m$ is Q -definable and hence, if it is not empty, then there would be a proper subset of X that is Q -defined by replacing each \mathfrak{h}_i with \mathfrak{h}_i^* . So each hole in each \mathfrak{b}_i is a maximal open ball in \mathfrak{b}_i . Suppose for contradiction that \mathfrak{b}_1 contains more than one holes $\mathfrak{h}_1, \dots, \mathfrak{h}_m$. Without loss of generality we may assume that $0 \notin \mathfrak{h}_1$. Since the subset $\mathfrak{h}_2 - \mathfrak{h}_1$ is an rv-ball and $1 \cdot (\mathfrak{h}_2 - \mathfrak{h}_1), \dots, (m+1) \cdot (\mathfrak{h}_2 - \mathfrak{h}_1)$ are distinct rv-balls, for some $1 \leq k \leq m+1$ we have that $\mathfrak{h}_1 + k \cdot (\mathfrak{h}_2 - \mathfrak{h}_1)$ is a maximal open ball in \mathfrak{b}_1 and is disjoint from $\bigcup_i \mathfrak{h}_i$. This means that there is a finite Q -definable set of maximal open balls in $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ that strictly contains the set of holes in $\mathfrak{b}_1, \dots, \mathfrak{b}_n$. This readily implies that X has a nonempty proper Q -definable subset, contradiction. \square

Note that, in the above lemma, if $Q = \emptyset$ then X cannot be a disjoint union of closed balls of radius $< \infty$, because in that case, by Lemma 5.4.16, the closed balls would have definable centers. Now, if $X \subseteq \text{VF}$ is atomic over $\langle Q \rangle$ then the *radius* and the *valuative center* of X are well-defined quantities: they are respectively the radius and the valuative center of the balls $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ in the above lemma. These are also denoted by $\text{rad}(X)$ and $\text{vcr}(X)$. The balls $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ are called the *haecceitistic components* of X .

Corollary 5.5.6. *If $X \subseteq \text{VF}$ is atomic over $\langle Q \rangle$ and $\mathfrak{b} \subseteq X$ is an open (closed) ball then every $a \in X$ is contained in an open (closed) ball $\mathfrak{d}_a \subseteq X$ with $\text{rad}(\mathfrak{d}_a) = \text{rad}(\mathfrak{b})$.*

Lemma 5.5.7. *Let $X \subseteq \text{VF}$ be atomic over $\langle Q \rangle$ and $f : X \rightarrow \text{VF}$ an Q -definable injective function. If X has only one haecceitistic component then $f(X)$ also has only one haecceitistic component.*

Proof. Let $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ be the haecceitistic components of $f(X)$ given by Lemma 5.5.5. Suppose that X is an open ball or a closed ball or a thin annulus. Suppose for contradiction that $n > 1$. Then there is exactly one of the components $\mathfrak{b}_1, \dots, \mathfrak{b}_n$, say \mathfrak{b}_1 , such that $f^{-1}(\mathfrak{b}_1)$ contains the punctured ball $X \setminus \bigcup_j \mathfrak{h}_j$ for some holes \mathfrak{h}_j . Consequently, since $\text{rad}(f(X))$ is Q -definable, the ball \mathfrak{b}_1 and $f^{-1}(\mathfrak{b}_1)$ are Q -definable, contradicting the assumption that X is atomic. \square

Lemma 5.5.8. *Let $X \subseteq \text{VF}$ be atomic over $\langle Q \rangle$ with haecceitistic components $\mathfrak{b}_1, \dots, \mathfrak{b}_n$ and \mathfrak{b} an open (closed) ball properly contained in some \mathfrak{b}_i . Set $\gamma = \text{rad}(\mathfrak{b})$. Then \mathfrak{b} is atomic over $\langle Q, \gamma, \mathfrak{b} \rangle$.*

Proof. We assume that \mathfrak{b} is an open ball, since the proof for closed balls is identical. By Lemma 5.5.3, X is atomic over $\langle Q, \gamma \rangle$. Since the infinite set of pairwise disjoint balls

$$\mathfrak{D} = \{\mathfrak{d} \subseteq X : \mathfrak{d} \text{ is an open subball of } X \text{ with } \text{rad}(\mathfrak{d}) = \gamma\}$$

is $\langle Q, \gamma \rangle$ -definable and $\bigcup \mathfrak{D} = X$, clearly no $\mathfrak{d} \in \mathfrak{D}$ is $\langle Q, \gamma \rangle$ -algebraic. So, by Lemma 5.5.2, every $\mathfrak{d} \in \mathfrak{D}$ is atomic over $\langle Q, \gamma, \mathfrak{d} \rangle$. \square

Lemma 5.5.9. *Let \mathfrak{o} be an open ball and \mathfrak{l} a close ball or a thin annulus such that both \mathfrak{o} and \mathfrak{l} are atomic over $\langle Q \rangle$. If $X \subseteq \mathfrak{o} \times \mathfrak{l}$ is Q -definable then the projection $\text{pr}_1 \upharpoonright X$ cannot be finite-to-one.*

Proof. We assume that \mathfrak{l} is a closed ball, since the proof for thin annuli is identical. Suppose for contradiction that there is an Q -definable $X \subseteq \mathfrak{o} \times \mathfrak{l}$ such that the first coordinate projection on X is finite-to-one. Note that, since \mathfrak{o} and \mathfrak{l} are atomic, we must have that $\text{pr}_1 X = \mathfrak{o}$ and $\text{pr}_2 X = \mathfrak{l}$. Let \mathfrak{M} be the set of maximal open subballs of \mathfrak{l} , which is Q -definable. For any $\mathfrak{r} \in \mathfrak{M}$, let $A_{\mathfrak{r}} = \text{pr}_1((\text{pr}_2 \upharpoonright X)^{-1}(\mathfrak{r}))$. By C -minimality each $A_{\mathfrak{r}}$ is a boolean combination of balls. In fact, for any $\mathfrak{r}, \mathfrak{h} \in \mathfrak{M}$, $A_{\mathfrak{r}}$ and $A_{\mathfrak{h}}$ must have the same number of boolean components, because otherwise there would be an Q -definable proper subset of \mathfrak{l} . Let this number be k .

For any $\mathfrak{r} \in \mathfrak{M}$, suppose that $\mathfrak{B} = \{\mathfrak{b}_1, \dots, \mathfrak{b}_k\}$ is the set of the boolean components of $A_{\mathfrak{r}}$, we let $\lambda_{\mathfrak{r}} = \min\{\text{rad}(\mathfrak{b}_1), \dots, \text{rad}(\mathfrak{b}_k)\}$. Moreover, for any $\mathfrak{b}_i, \mathfrak{b}_j \in \mathfrak{B}$, if $\mathfrak{b}_i \cap \mathfrak{b}_j \neq \emptyset$ then let $\rho(\mathfrak{b}_i, \mathfrak{b}_j) = \min\{\text{rad}(\mathfrak{b}_i), \text{rad}(\mathfrak{b}_j)\}$, otherwise let $\rho(\mathfrak{b}_i, \mathfrak{b}_j) = \text{val}(\mathfrak{b}_i - \mathfrak{b}_j)$. Let $\rho_{\mathfrak{r}} = \min\{\rho(\mathfrak{b}_i, \mathfrak{b}_j) : (\mathfrak{b}_i, \mathfrak{b}_j) \in \mathfrak{B}^2\}$. Note that the subsets $\Lambda = \{\lambda_{\mathfrak{r}} : \mathfrak{r} \in \mathfrak{M}\} \subseteq \Gamma$ and $\Delta = \{\rho_{\mathfrak{r}} : \mathfrak{r} \in \mathfrak{M}\} \subseteq \Gamma$ are both Q -definable. Since \mathfrak{l} is atomic, we must have that both Λ and Δ are singletons, say $\Lambda = \{\lambda\}$ and $\Delta = \{\rho\}$. Also, we claim that $\lambda < \text{rad}(\mathfrak{o})$. To see this, suppose for contradiction $\lambda_{\mathfrak{r}} = \text{rad}(\mathfrak{o})$ for every $\mathfrak{r} \in \mathfrak{M}$. This means that $A_{\mathfrak{r}}$ has \mathfrak{o} as a positive boolean component for every $\mathfrak{r} \in \mathfrak{M}$. Since \mathfrak{o} is open, we have that for any n and any $\mathfrak{r}_1, \dots, \mathfrak{r}_n \in \mathfrak{M}$ there is an $a \in \bigcap_{i \leq n} A_{\mathfrak{r}_i}$ and hence there is a $b_i \in \mathfrak{r}_i$ for every $i \leq n$ such that $(a, b_i) \in X$. Therefore, by compactness, there is an $a \in \mathfrak{o}$ such that the fiber $\{b : (a, b) \in X\}$ is infinite, contradicting the assumption that $\text{pr}_1 \upharpoonright X$ is finite-to-one.

Now, fix an $\mathfrak{r} \in \mathfrak{M}$. Again, since \mathfrak{o} is open, there is a proper open subball \mathfrak{z} of \mathfrak{o} that properly contains $A_{\mathfrak{r}}$. Let $B_{\mathfrak{z}} = \text{pr}_2((\text{pr}_1 \upharpoonright X)^{-1}(\mathfrak{z}))$. Since $B_{\mathfrak{z}}$ properly contains the maximal open subball \mathfrak{r} of \mathfrak{l} , by C -minimality, either \mathfrak{r} is a boolean component of $B_{\mathfrak{z}}$ that is disjoint from any other boolean component of $B_{\mathfrak{z}}$ or \mathfrak{l} is a positive boolean component of $B_{\mathfrak{z}}$. However, the former is impossible, because in that case $B_{\mathfrak{z}}$ could only have finitely many maximal open subballs of \mathfrak{l} as its positive boolean components and consequently, since $\Lambda = \{\lambda\}$ is a singleton, \mathfrak{z} could not be an open ball, contradiction. So we must have that \mathfrak{l} is a positive boolean component of $B_{\mathfrak{z}}$. This means that, by C -minimality, $B_{\mathfrak{z}}$ can only have finitely many maximal open subballs of \mathfrak{l} as its negative boolean components, say $\mathfrak{r}_1, \dots, \mathfrak{r}_n$. Again, since $\Lambda = \{\lambda\}$ and $\lambda < \text{rad}(\mathfrak{o})$, $\bigcup_{i \leq n} A_{\mathfrak{r}_i}$ must be a proper subset of $\mathfrak{o} \setminus \mathfrak{z}$ and hence there is a $\mathfrak{h} \in \mathfrak{M}$ such that $\mathfrak{h} \subseteq B_{\mathfrak{z}}$ and $A_{\mathfrak{h}}$ has a boolean component contained in \mathfrak{z} and another boolean component disjoint from \mathfrak{z} . This implies that $\rho_{\mathfrak{h}} \leq \text{rad}(\mathfrak{z})$. On the other hand, since $A_{\mathfrak{r}} \subseteq \mathfrak{z}$, we have that $\rho_{\mathfrak{r}} > \text{rad}(\mathfrak{z})$. This is a contradiction since Δ is a singleton. \square

Lemma 5.5.10. *Let \mathfrak{q} be an open ball such that it is atomic over $\langle \mathfrak{q} \rangle$. Let $X \subseteq \text{VF}$ be atomic over $\langle \mathfrak{q} \rangle$ such that it only has one haecceitistic component. If X is infinite then it is either an open ball or a thin annulus.*

Proof. Suppose for contradiction that X is a closed ball of radius $< \infty$. Let ψ be a quantifier-free formula in disjunctive normal form that defines X . Note that, by Lemma 5.4.16, \mathfrak{q} must occur in ψ . Without loss of generality, \mathfrak{q} is represented in ψ by some $q \in \mathfrak{q}$. We claim that any disjunct in ψ that contains a nontrivial VF-sort equality $f(x) = 0$ as a conjunct is redundant: if q does not occur in $f(x)$ then, since X is atomic, $f(a) \neq 0$ for any $a \in X$; if q does occur in $f(x)$ then we still have that $f(a) \neq 0$ for any $a \in X$, because otherwise there would be an \mathfrak{q} -definable $Y \subseteq \mathfrak{q} \times X$ with $\text{pr}_1 \upharpoonright Y$ finite-to-one, contradicting Lemma 5.5.9. Dually, we may also assume that no disjunct in ψ contains VF-sort disequality. Similarly, for any term $\text{rv}(g(x))$ in ψ with $g(x)$ nonconstant, we have that $g(a) \neq 0$ for any $a \in X$. It is not hard to see that, since all the roots of all the nonconstant polynomials $g(x)$ in all the terms of the form $\text{rv}(g(x))$ in ψ lie outside X and X is a ball, there is a $b \notin X$ such that $\text{rv}(g(b)) = \text{rv}(g(a_1)) = \text{rv}(g(a_2))$ for any $a_1, a_2 \in X$. So b also satisfies ψ , contradiction. \square

Lemma 5.5.11. *Let $X \subseteq \text{VF}$ be an open ball atomic over $\langle Q \rangle$ and $f : X \rightarrow \text{VF}$ an Q -definable injective function. Then $f(X)$ is also an atomic open ball.*

Proof. By Lemma 5.5.7, $f(X)$ is an open ball or a closed ball or a thin annulus. Then, according to Lemma 5.5.9, $f(X)$ must be an open ball. \square

Lemma 5.5.12. *Let $X \subseteq \text{VF}$ generate a complete type. Let $f : \text{VF} \rightarrow \text{VF}$ be a definable function such that $f \upharpoonright X$ is injective. Then for every open ball $\mathfrak{b} \subseteq X$ the image $f(\mathfrak{b})$ is also an open ball.*

Proof. Fix an open ball $\mathfrak{b} \subseteq X$ and set $\gamma = \text{rad}(\mathfrak{b})$. We claim that \mathfrak{b} is not γ -algebraic. To see this, suppose for contradiction that there is a formula $\psi(\gamma)$ in disjunctive normal form that defines a finite set $\mathfrak{B} = \{\mathfrak{b}_1, \dots, \mathfrak{b}_n\}$ of balls such that $\mathfrak{b}_1 = \mathfrak{b}$. Without loss of generality we may assume that every \mathfrak{b}_i is an open ball of radius γ and, since \mathfrak{B} is finite, $\bigcup \mathfrak{B} \subseteq X$. So all VF-sort literals in $\psi(\gamma)$ are redundant. For any term $\text{rv}(g(x))$ in ψ with $g(x) \in \text{VF}(\langle \emptyset \rangle)[x]$, clearly if $g(x)$ is not a constant polynomial then $g(b) \neq 0$ for any $b \in \mathfrak{b}_i$. Since all the roots of all the nonconstant polynomials $g(x)$ in all the terms of the form $\text{rv}(g(x))$ in $\psi(\gamma)$ lie outside $\bigcup \mathfrak{B}$ and \mathfrak{B} is finite, there is an $a \notin \bigcup \mathfrak{B}$ such that $\text{rv}(g(a)) = \text{rv}(g(b_1)) = \text{rv}(g(b_2))$ for any $b_1, b_2 \in \mathfrak{b}$. Therefore a also satisfies $\psi(\gamma)$, contradiction.

Now, since \mathfrak{b} is not γ -algebraic, by Lemma 5.5.2, \mathfrak{b} is atomic over $\langle \gamma, \mathfrak{b} \rangle$ and hence, by Lemma 5.5.11, $f(\mathfrak{b})$ is an open ball. \square

Proposition 5.5.13. *Let $X, Y \subseteq \text{VF}$ be definable and $f : X \rightarrow Y$ a definable bijection. Then there are definable disjoint subsets $X_1, \dots, X_n \subseteq X$ with $\bigcup X_i = X$ such that, for any open balls $\mathbf{a} \in X_i$ and $\mathbf{b} \in f(X_i)$, both $f(\mathbf{a})$ and $f^{-1}(\mathbf{b})$ are open balls.*

Proof. For every $a \in X$ let $Z_a \subseteq X$ be the intersection of all definable subsets of X that contains a . So Z_a generates a complete type. By Lemma 5.5.12, for every open ball $\mathbf{a} \subseteq Z_a$, the image $f(\mathbf{a})$ is an open ball. This open-to-open property may be rephrased as follows: for every $b \in Z_a$ and $t \in \text{RV}$ let $\mathfrak{o}(b, t)$ be the open ball that contains b and has radius $v_{\text{rv}}(t)$, if $\mathfrak{o}(b, t) \subseteq Z_a$ then $f(\mathfrak{o}(b, t))$ is an open ball. Therefore, by compactness, there is a definable subset $D_a \subseteq X$ containing a such that $f \upharpoonright D_a$ has this open-to-open property. By compactness again, there are definable subsets $X_1, \dots, X_m \subseteq X$ with $\bigcup X_i = X$ such that each $f \upharpoonright X_i$ has this open-to-open property. Similarly there are definable subsets $Y_1, \dots, Y_l \subseteq Y$ with $\bigcup Y_i = Y$ such that each $f^{-1} \upharpoonright Y_i$ has this open-to-open property. The partition of X determined by $X_1, \dots, X_m, f^{-1}(Y_1), \dots, f^{-1}(Y_l)$ is as desired. \square

Let $X \subseteq \text{VF}^n \times \text{RV}^m$ and $i \in \{1, \dots, n\}$. A subset $Y \subseteq X$ is an *open ball contained in $X[i]$* if Y is of the form $\mathbf{b} \times \{\bar{x}\}$, where \mathbf{b} is an open ball and $\bar{x} \in \text{pr}_{\bar{i}} X$. Of course, if Y is an open ball contained in $X[i]$ and $\text{pr}_{\bar{i}} X$ is a singleton then we simply say that Y is an open ball contained in X . The same goes to closed balls, rv-balls, simplexes, etc.

Definition 5.5.14. Let $X \subseteq \text{VF}^{n_1} \times \text{RV}^{m_1}$, $Y \subseteq \text{VF}^{n_2} \times \text{RV}^{m_2}$, and $f : X \rightarrow Y$ a bijection. Let $i \in \{1, \dots, n_1\}$ and $j \in \{1, \dots, n_2\}$. For any $\bar{a} \in \text{pr}_{\bar{i}} X$ and any $\bar{b} \in \text{pr}_{\bar{j}} Y$, let $f_{\bar{a}, \bar{b}} = f \upharpoonright (\text{fib}(X, \bar{a}) \cap f^{-1}(\text{fib}(Y, \bar{b})))$. We say that f has the (i, j) -*open-to-open* property if, for every $\bar{a} \in \text{pr}_{\bar{i}} X$ and every $\bar{b} \in \text{pr}_{\bar{j}} Y$, $f_{\bar{a}, \bar{b}}$ has the open-to-open property described in Proposition 5.5.13. If f has the (i, j) -open-to-open property for every $(i, j) \in \{1, \dots, n_1\} \times \{1, \dots, n_2\}$ then f has the *open-to-open* property.

With this understanding, Proposition 5.5.13 may be easily generalized as follows:

Proposition 5.5.15. *Let $X \subseteq \text{VF}^{n_1} \times \text{RV}^{m_1}$, $Y \subseteq \text{VF}^{n_2} \times \text{RV}^{m_2}$ be definable subsets and $f : X \rightarrow Y$ a definable bijection. Then there are definable disjoint subsets $X_1, \dots, X_n \subseteq X$ with $\bigcup X_i = X$ such that $f \upharpoonright X_i$ has the open-to-open property for every i .*

Proof. First observe that if f has the (i, j) -open-to-open property then, for every subset $X^* \subseteq X$, $f \upharpoonright X^*$ has the (i, j) -open-to-open property. Next, by Proposition 5.5.13, for any $\bar{a} \in \text{pr}_{>1} X$ and $\bar{b} \in \text{pr}_{>1} Y$ there is a (\bar{a}, \bar{b}) -definable finite partition V_1, \dots, V_n of $\text{dom}(f_{\bar{a}, \bar{b}})$ such that each $f_{\bar{a}, \bar{b}} \upharpoonright V_i$ has the open-to-open property. Since V_1, \dots, V_n may be extended into a definable partition V_1^*, \dots, V_n^* of X such that $V_i^* \cap \text{dom}(f_{\bar{a}, \bar{b}}) = V_i$ and for any finite collection of partitions P_1, \dots, P_m of X there is a partition P of X such that P is finer

than each P_i , by compactness, we obtain a definable partition $V_{1,1}, \dots, V_{1,n}$ of X such that each $f \upharpoonright V_{1,i}$ has the (1,1)-open-to-open property. Iterating this procedure for each $(i, j) \in \{1, \dots, n_1\} \times \{1, \dots, n_2\}$ on each piece of the partition obtained in the previous step, we eventually get a partition of X that is as desired. \square

5.6 Categories of definable subsets

Motivic integrals will be constructed as homomorphisms between the Grothendieck semigroups (or semirings) of various categories associated with the theory ACVF_S^0 .

5.6.1 Dimensions

Before we introduce the categories and their Grothendieck groups, two notions of dimension with respect to the two different sorts are needed.

Definition 5.6.1. Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be a definable subset. The *VF-dimension* of X , denoted as $\dim_{\text{VF}} X$, is the smallest number k such that there is a definable finite-to-one function $f : X \rightarrow \text{VF}^k \times \text{RV}^l$.

Lemma 5.6.2. Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be a definable subset. Then $\dim_{\text{VF}} X \leq k$ if and only if there is a definable injection $f : X \rightarrow \text{VF}^k \times \text{RV}^l$ for some l .

Proof. Suppose that $\dim_{\text{VF}} X \leq k$. Let $g : X \rightarrow \text{VF}^k \times \text{RV}^l$ be a definable finite-to-one function. For every $(\bar{a}, \bar{t}) \in g(X)$, since $g^{-1}(\bar{a}, \bar{t})$ is finite, by Lemma 5.4.3, there is an (\bar{a}, \bar{t}) -definable injection $h_{\bar{a}, \bar{t}} : g^{-1}(\bar{a}, \bar{t}) \rightarrow \text{RV}^j$ for some j . By compactness, there is a definable function $h : X \rightarrow \text{RV}^j$ for some j such that $h \upharpoonright g^{-1}(\bar{a}, \bar{t})$ is injective for every $(\bar{a}, \bar{t}) \in g(X)$. Then the function f on X given by $(\bar{b}, \bar{s}) \mapsto (g(\bar{b}, \bar{s}), h(\bar{b}, \bar{s}))$ is as desired. The other direction is trivial. \square

Lemma 5.6.3. Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be a definable subset and $f : X \rightarrow \text{RV}^l$ a definable function. Then $\dim_{\text{VF}} X = \max \{ \dim_{\text{VF}} f^{-1}(\bar{t}) : \bar{t} \in \text{ran}(f) \}$.

Proof. Let $\max \{ \dim_{\text{VF}} f^{-1}(\bar{t}) : \bar{t} \in \text{RV}^l \} = k$. By Lemma 5.6.2, for every $\bar{t} \in \text{ran}(f)$, there is a \bar{t} -definable injective function $h_{\bar{t}} : f^{-1}(\bar{t}) \rightarrow \text{VF}^k \times \text{RV}^j$ for some j . By compactness, there is a definable function $h : X \rightarrow \text{VF}^k \times \text{RV}^j$ for some j such that $h \upharpoonright f^{-1}(\bar{t})$ is injective for every $\bar{t} \in \text{ran}(f)$. Then the function on X given by $(\bar{b}, \bar{s}) \mapsto (h(\bar{b}, \bar{s}), f(\bar{b}, \bar{s}))$ is injective and hence $\dim_{\text{VF}} X \leq k$. The other direction is trivial. \square

Lemma 5.6.4. Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be a definable subset. Suppose that there is an $(\bar{a}, \bar{t}) \in X$ such that the transcendental degree of $\text{VF}(\langle \bar{a} \rangle)$ over $\text{VF}(\langle \emptyset \rangle)$ is k . Then $\dim_{\text{VF}} X \geq k$.

Proof. Suppose for contradiction that the transcendental degree of $\text{VF}(\langle \bar{a} \rangle)$ over $\text{VF}(\langle \emptyset \rangle)$ is k for some $(\bar{a}, \bar{t}) \in X$ but $\dim_{\text{VF}} X = i < k$. By Lemma 5.6.2, there is a definable bijection $f : X \rightarrow Y \subseteq \text{VF}^i \times \text{RV}^l$ for some l . Let $f(\bar{a}, \bar{t}) = (\bar{b}, \bar{s})$. By quantifier elimination, we have that $\text{VF}(\langle \bar{a} \rangle)^{\text{ac}} \subseteq \text{VF}(\langle \bar{b} \rangle)^{\text{ac}}$. So the transcendental degree of $\text{VF}(\langle \bar{a} \rangle)$ over $\text{VF}(\langle \emptyset \rangle)$ is at most i , contradiction. \square

Corollary 5.6.5. *Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be a definable subset that contains a subset of the form $\{(0, \dots, 0)\} \times \text{rv}^{-1}(\bar{t}) \times \{\bar{s}\}$ for some $\bar{t} \in (\text{RV}^\times)^k$. Then $\dim_{\text{VF}} X \geq k$.*

Definition 5.6.6. Let $X \subseteq \text{RV}^m$ be a definable subset. The *RV-dimension* of X , denoted as $\dim_{\text{RV}} X$, is the smallest number k such that there is a definable finite-to-one function $f : X \rightarrow \text{RV}^k$.

Definition 5.6.7. Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be a definable subset. The *RV-fiber dimension* of X , denoted as $\dim_{\text{RV}}^{\text{fib}} X$, is $\max \{\dim_{\text{RV}}(\text{fib}(X, \bar{a})) : \bar{a} \in \text{pVF } X\}$.

Lemma 5.6.8. *Let $X \subseteq \text{VF}^{n_1} \times \text{RV}^{m_1}$ be a definable subset and $f : X \rightarrow \text{VF}^{n_2} \times \text{RV}^{m_2}$ a definable injection. Then $\dim_{\text{RV}}^{\text{fib}} X = \dim_{\text{RV}}^{\text{fib}} f(X)$.*

Proof. Let $\dim_{\text{RV}}^{\text{fib}} X = k_1$ and $\dim_{\text{RV}}^{\text{fib}} f(X) = k_2$. Since for every $\bar{b} \in \text{pVF } f(X)$ there is a \bar{b} -definable finite-to-one function $h_{\bar{b}} : \text{fib}(f(X), \bar{b}) \rightarrow \text{RV}^{k_2}$, by compactness, there is a definable function $h : f(X) \rightarrow \text{RV}^{k_2}$ such that $h \upharpoonright \text{fib}(f(X), \bar{b})$ is finite-to-one for every $\bar{b} \in \text{pVF } f(X)$. For every $\bar{a} \in \text{pVF } X$, by Lemma 5.4.8, the subset $(\text{pVF} \circ f)(\text{fib}(X, \bar{a}))$ is finite. So the function $g_{\bar{a}}$ on $\text{fib}(X, \bar{a})$ given by $(\bar{a}, \bar{t}) \mapsto (h \circ f)(\bar{a}, \bar{t})$ is \bar{a} -definable and finite-to-one. So $k_1 \leq k_2$. Symmetrically we also have $k_1 \geq k_2$ and hence $k_1 = k_2$. \square

5.6.2 Categories of definable subsets

The class of the objects and the class of the morphisms of any category \mathcal{C} are denoted as $\text{Ob } \mathcal{C}$ and $\text{Mor } \mathcal{C}$, respectively.

Definition 5.6.9 (VF-categories). The objects of the category $\text{VF}[k, \cdot]$ are the definable subsets of VF-dimension $\leq k$. The morphisms in this category are the definable functions between the objects.

The category $\text{VF}[k]$ is the full subcategory of $\text{VF}[k, \cdot]$ of the definable subsets that have RV-fiber dimension 0 (that is, all the RV-fibers are finite). The category $\text{VF}_*[\cdot]$ is the union of the categories $\text{VF}[k, \cdot]$. The category VF_* is the union of the categories $\text{VF}[k]$.

Note that, for any definable subset X , by Lemma 5.4.3 and Lemma 5.6.4, $\text{fib}(X, \bar{t})$ is finite for any $\bar{t} \in \text{pRV } X$ if and only if $X \in \text{Ob } \text{VF}[0, \cdot]$.

Definition 5.6.10. For any tuple $\bar{t} = (t_1, \dots, t_n) \in \text{RV}$, the *weight* of \bar{t} is the number $|\{i \leq n : t_i \neq \infty\}|$, which is denoted as $\text{wgt } \bar{t}$.

Definition 5.6.11 (RV-categories). The objects of the category $\text{RV}[k, \cdot]$ are the definable pairs (U, f) , where $U \subseteq \text{RV}^m$ for some m and $f : U \rightarrow \text{RV}^k$ is a function (RV^0 is taken to be the singleton $\{\infty\}$). We often denote the projections $\text{pr}_i \circ f$ as $f_{|i}$ and write f as $(f_{|1}, \dots, f_{|k})$. The *companion* U_f of (U, f) is the subset $\{(f(\bar{u}), \bar{u}) : \bar{u} \in U\}$.

For any two objects $(U, f), (U', f')$ in this category and any function $F : U \rightarrow U'$, the *companion* $F_{f, f'} : U_f \rightarrow U'_{f'}$ of F is the function given by

$$(f(\bar{u}), \bar{u}) \mapsto ((f' \circ F)(\bar{u}), F(\bar{u})).$$

If, for every $\bar{u} \in U$, $\text{wgt } f(\bar{u}) \leq \text{wgt}(f' \circ F)(\bar{u})$, then we say that F is *volumetric*. If F is definable, volumetric, and, for every $\bar{t} \in \text{ran}(f)$, the subset

$$(\text{pr}_{\leq k} \circ F_{f, f'})(\{\bar{t}\} \times f^{-1}(\bar{t}))$$

is finite, then it is a morphism in $\text{Mor } \text{RV}[k, \cdot]$.

The category $\text{RV}[k]$ is the full subcategory of $\text{RV}[k, \cdot]$ of the pairs (U, f) such that $f : U \rightarrow \text{RV}^k$ is finite-to-one.

Direct sums over these categories are formed naturally:

$$\text{RV}[* , \cdot] = \coprod_{0 \leq k} \text{RV}[k, \cdot], \quad \text{RV}[*] = \coprod_{0 \leq k} \text{RV}[k].$$

Notation 5.6.12. We just write X for the object $(X, \text{id}) \in \text{RV}[k, \cdot]$. For an object $(U, \text{pr}_E) \in \text{RV}[k, \cdot]$ with $E \subseteq \mathbb{N}$ and $|E| = k$, it is often much more convenient to assume that $E = \{1, \dots, k\}$ and hence write (U, pr_E) as $(U, \text{pr}_{\leq k})$. This should not cause any confusion in context.

Remark 5.6.13. One of the main reasons for the peculiar forms of the objects and the morphisms in the RV-categories is that each isomorphism class in these categories may be “lifted” to an isomorphism class in the corresponding VF-category. See Proposition 5.9.6 and Corollary 5.9.7 for details.

A *subobject* of an object X of a VF-category is just a definable subset. A *subobject* of an object (U, f) of an RV-category is a definable pair (X, g) with X a definable subset of U and $g = f \upharpoonright X$. Note that the inclusion map is a morphism in both cases.

Notice that the cartesian product of two objects $X, Y \in \text{VF}[k, \cdot]$ may or may not be in $\text{VF}[k, \cdot]$. On the other hand, the cartesian product of two objects $(U, f), (U', f') \in \text{RV}[k, \cdot]$ is the object $(U \times U', f \times f') \in \text{RV}[2k, \cdot]$, which is definitely not in $\text{RV}[k, \cdot]$ if $k > 0$. Hence, in $\text{RV}[* , \cdot]$ or $\text{RV}[*]$, multiplying with a singleton

in general changes isomorphism class.

Remark 5.6.14. The categories $\mathrm{VF}_*[\cdot]$ and VF_* are formed through union instead of direct sum or other means that induces more complicated structure. The reason for this is that the main goal of the Hrushovski-Kazhdan integration theory is to assign motivic volumes, that is, elements in the Grothendieck groups of the RV-categories, to the definable subsets, or rather, the isomorphism classes of the definable subsets, in the VF-categories, and the simplest categories that contain all the definable subsets that may be “measured” in this motivic way are $\mathrm{VF}_*[\cdot]$ and VF_* . In contrast, the unions of the RV-categories are naturally endowed with the structure of direct sum, which gives rise to graded Grothendieck semirings.

Definition 5.6.15. For any $(U, f) \in \mathrm{Ob} \mathrm{RV}[k, \cdot]$ and any $F \in \mathrm{Mor} \mathrm{RV}[k, \cdot]$, let $\mathbb{E}_k(f)$ be the function on U given by $\bar{u} \mapsto (f(\bar{u}), \infty)$, $\mathbb{E}_k(U, f) = (U, \mathbb{E}_k(f))$, and $\mathbb{E}_k(F) = F$. Obviously $\mathbb{E}_k : \mathrm{RV}[k, \cdot] \longrightarrow \mathrm{RV}[k+1, \cdot]$ is a functor that is faithful, full, and injective on objects. For any $i < j$ let $\mathbb{E}_{i,j} = \mathbb{E}_{j-1} \circ \dots \circ \mathbb{E}_i$ and $\mathbb{E}_{i,i} = \mathrm{id}$.

Motivic integrals shall be induced by the following fundamental maps.

Definition 5.6.16. For any $(U, f) \in \mathrm{Ob} \mathrm{RV}[k, \cdot]$, let

$$\mathbb{L}_k(U, f) = \bigcup \{ \mathrm{rv}^{-1}(f(\bar{u})) \times \{ \bar{u} \} : \bar{u} \in U \}.$$

The map $\mathbb{L}_k : \mathrm{Ob} \mathrm{RV}[k, \cdot] \longrightarrow \mathrm{Ob} \mathrm{VF}[k, \cdot]$ is called the *kth canonical RV-lift*. The map $\mathbb{L}_{\leq k} : \mathrm{Ob} \mathrm{RV}[\leq k, \cdot] \longrightarrow \mathrm{Ob} \mathrm{VF}[k, \cdot]$ is given by

$$((U_1, f_1), \dots, (U_k, f_k)) \mapsto \bigoplus_{i \leq k} (\mathbb{L}_k \circ \mathbb{E}_{i,k})(U_i, f_i).$$

The map $\mathbb{L} : \mathrm{Ob} \mathrm{RV}[* , \cdot] \longrightarrow \mathrm{Ob} \mathrm{VF}_*[\cdot]$ is simply the union of the maps $\mathbb{L}_{\leq k}$.

For notational convenience, when there is no danger of confusion, we shall drop the subscripts and simply write \mathbb{E} and \mathbb{L} for these maps.

Remark 5.6.17. Observe that if $(U, f) \in \mathrm{Ob} \mathrm{RV}[k]$ then $\mathbb{L}(U, f) \in \mathrm{Ob} \mathrm{VF}[k]$ and hence the restriction $\mathbb{L} : \mathrm{Ob} \mathrm{RV}[k] \longrightarrow \mathrm{Ob} \mathrm{VF}[k]$ is well-defined. Similarly we have the maps $\mathbb{L} : \mathrm{Ob} \mathrm{RV}[\leq k] \longrightarrow \mathrm{Ob} \mathrm{VF}[k]$ and $\mathbb{L} : \mathrm{Ob} \mathrm{RV}[*] \longrightarrow \mathrm{Ob} \mathrm{VF}_*$.

Note that $\mathrm{rv}(\mathbb{L}(U, f)) = U_f$ for $(U, f) \in \mathrm{Ob} \mathrm{RV}[k, \cdot]$.

Lemma 5.6.18. *Let $(U, f), (U', f') \in \mathrm{Ob} \mathrm{RV}[k, \cdot]$ and $F : U \longrightarrow U'$ a definable volumetric function. Suppose*

that there is a definable function $F^\dagger : \mathbb{L}(U, f) \longrightarrow \mathbb{L}(U', f')$ such that the diagram

$$\begin{array}{ccccc}
\mathbb{L}(U, f) & \xrightarrow{\text{rv}} & U_f & \xrightarrow{\text{pr}_{>k}} & U \\
F^\dagger \downarrow & & \downarrow F_{f, f'} & & \downarrow F \\
\mathbb{L}(U', f') & \xrightarrow{\text{rv}} & U'_{f'} & \xrightarrow{\text{pr}_{>k}} & U'
\end{array}$$

commutes. Then F is a morphism in $\text{RV}[k, \cdot]$.

Proof. It is enough to show that, for every $\bar{u} \in U$ and every $i \leq k$, $((f')_i \circ F)(\bar{u}) \in \text{acl}(f(\bar{u}))$, which is equivalent to $(\text{pr}_i \circ F_{f, f'})(f(\bar{u}), \bar{u}) \in \text{acl}(f(\bar{u}))$. To that end, fix a $\bar{u} \in U$. Let $\bar{a} \in f(\bar{u})$ and $F^\dagger(\bar{a}, \bar{u}) = (b_1, \dots, b_k, \bar{u}')$. By an argument similar to the one in the proof of Lemma 5.4.8, we deduce that $b_i \in \text{acl}(\bar{a})$ and hence $\text{rv}(b_i) \in \text{acl}(\bar{a})$ for each $i \leq k$. By Lemma 5.2.12, we conclude that $\text{rv}(b_i) \in \text{acl}(f(\bar{u}))$. \square

Remark 5.6.19. In Lemma 5.6.18, if both F and F^\dagger are bijections then we may drop the assumption that F is volumetric, since it is guaranteed by the commutative diagram and Corollary 5.6.5.

5.6.3 Grothendieck groups

We now introduce the Grothendieck groups associated with the categories defined above. The construction is of course the same for any reasonable category of definable sets of a first-order theory. For concreteness, we shall limit our attention to the present context.

Convention 5.6.20. Let f_1, \dots, f_n be definable functions on subsets X_1, \dots, X_n , respectively. Padding with elements in $\text{dcl}(\emptyset)$ if necessary, we may glue f_1, \dots, f_n together to form one definable function $f : \biguplus_i X_i \longrightarrow \biguplus_i f_i(X_i)$ in the obvious way. Below, when functions or other kinds of subsets are glued together in this way, we shall always tacitly assume that sufficient padding work has been performed.

Let \mathcal{C} be a VF-category or an RV-category. For any $X \in \text{Ob } \mathcal{C}$, let $[X]$ denote the isomorphism class of X . The *Grothendieck semigroup* of \mathcal{C} , denoted as $\mathbf{K}_+ \mathcal{C}$, is the semigroup generated by the isomorphism classes $[X]$ of elements $X \in \text{Ob } \mathcal{C}$, subject to the relation

$$[X] + [Y] = [X \cup Y] + [X \cap Y].$$

It is easy to check that $\mathbf{K}_+ \mathcal{C}$ is actually a commutative monoid, the identity element being $[\emptyset]$ or $([\emptyset], \dots)$. Since \mathcal{C} always has disjoint unions, the elements of $\mathbf{K}_+ \mathcal{C}$ are precisely the isomorphism classes of \mathcal{C} . If \mathcal{C} is one of the categories $\text{VF}_*[\cdot]$, VF_* , $\text{RV}[*]$, and $\text{RV}[*]$ then it is closed under cartesian product. In this case,

$\mathbf{K}_+ \mathcal{C}$ has a semiring structure with multiplication given by

$$[X][Y] = [X \times Y].$$

Since the symmetry isomorphisms $X \times Y \longrightarrow Y \times X$ and the association isomorphisms $(X \times Y) \times Z \longrightarrow X \times (Y \times Z)$ are always present in these categories, $\mathbf{K}_+ \mathcal{C}$ is always a commutative semiring.

Remark 5.6.21. If \mathcal{C} is either $\text{VF}_*[\cdot]$ or VF_* then the isomorphism class of definable singletons is the multiplicative identity element of $\mathbf{K}_+ \mathcal{C}$. If \mathcal{C} is $\text{RV}[* , \cdot]$ then we adjust multiplication when $\text{RV}[0, \cdot]$ is involved as follows. For any $(U, f) \in \text{RV}[0, \cdot]$ and $(X, g) \in \text{RV}[k, \cdot]$, let $(U, f) \boxtimes (X, g) = (X, g) \boxtimes (U, f) = (U \times X, g^*)$, where g^* is the function on $U \times X$ given by $(\bar{t}, \bar{s}) \mapsto g(\bar{s})$. Let

$$[(U, f)][(X, g)] = [(U, f) \boxtimes (X, g)].$$

It is easily seen that, with this adjustment, $\mathbf{K}_+ \text{RV}[* , \cdot]$ becomes a filtrated semiring and its multiplicative identity element is the isomorphism class of (∞, id) in $\text{RV}[0, \cdot]$. Multiplication in $\mathbf{K}_+ \text{RV}[*]$ is adjusted in the same way.

Definition 5.6.22. A *semigroup congruence relation* on $\mathbf{K}_+ \mathcal{C}$ is a sub-semigroup R of the semigroup $\mathbf{K}_+ \mathcal{C} \times \mathbf{K}_+ \mathcal{C}$ such that R is an equivalence relation on $\mathbf{K}_+ \mathcal{C}$. Similarly, a *semiring congruence relation* on $\mathbf{K}_+ \mathcal{C}$ is a sub-semiring R of the semiring $\mathbf{K}_+ \mathcal{C} \times \mathbf{K}_+ \mathcal{C}$ such that R is an equivalence relation on $\mathbf{K}_+ \mathcal{C}$.

Let R be a semigroup congruence relation on $\mathbf{K}_+ \mathcal{C}$ and $(x, y), (v, w) \in R$. Then $(x + v, y + v), (y + v, y + w) \in R$ and hence $(x + v, y + w) \in R$. Therefore the equivalence classes of R has a semigroup structure that is induced by that of $\mathbf{K}_+ \mathcal{C}$. This semigroup is denoted as $\mathbf{K}_+ \mathcal{C}/R$ and is also referred to as a Grothendieck semigroup. Similarly, if R is a semiring congruence relation on $\mathbf{K}_+ \mathcal{C}$ then $\mathbf{K}_+ \mathcal{C}/R$ is actually a Grothendieck semiring.

Remark 5.6.23. Let R be an equivalence relation on the semiring $\mathbf{K}_+ \mathcal{C}$. If for every $(x, y) \in R$ and every $z \in \mathbf{K}_+ \mathcal{C}$ we have that $(x + z, y + z) \in R$ and $(xz, yz) \in R$ then R is a semiring congruence relation.

Let $(\mathbb{Z}^{\mathbf{K}_+ \mathcal{C}}, \oplus)$ be the free abelian group generated by the elements of $\mathbf{K}_+ \mathcal{C}$ and C the subgroup of $(\mathbb{Z}^{\mathbf{K}_+ \mathcal{C}}, \oplus)$ generated by all elements of $(\mathbb{Z}^{\mathbf{K}_+ \mathcal{C}}, \oplus)$ of the types

$$(1 \cdot x) \oplus ((-1) \cdot x)$$

$$(1 \cdot x) \oplus (1 \cdot y) \oplus ((-1) \cdot (x + y)),$$

where $x, y \in \mathbf{K}_+ \mathcal{C}$. The *Grothendieck group* of \mathcal{C} , denoted as $\mathbf{K} \mathcal{C}$, is the formal groupification $(\mathbb{Z}^{(\mathbf{K}_+ \mathcal{C})}, \oplus) / \mathcal{C}$ of $\mathbf{K}_+ \mathcal{C}$, which is essentially unique by the universal mapping property. Clearly $\mathbf{K}_+ \mathcal{C}$ is canonically isomorphic to a sub-semigroup of $\mathbf{K} \mathcal{C}$. If $\mathbf{K}_+ \mathcal{C}$ is a semiring then $\mathbf{K} \mathcal{C}$ is a commutative ring.

Remark 5.6.24. It is easily checked that \mathbb{E}_k induces an injective semigroup homomorphisms $\mathbf{K}_+ \text{RV}[k, \cdot] \longrightarrow \mathbf{K}_+ \text{RV}[k+1, \cdot]$, which is also denoted as \mathbb{E}_k .

Notation 5.6.25. For any definable subset $X \subseteq \text{RV}^n$, we write $[X]_n$ for the isomorphism class $[(X, \text{id})] \in \mathbf{K}_+ \text{RV}[n, \cdot]$. For any subset $E \subseteq \mathbb{N}$ with $|E| = k$, we write $[X]_E$ for the isomorphism class $[(X, \text{pr}_E)] \in \mathbf{K}_+ \text{RV}[k, \cdot]$. If $E = \{1, \dots, k\}$ etc. then we may write $[X]_{\leq k}$ etc. If X is a singleton then we just write $[1]_k$ for the isomorphism class $[(X, f)] \in \mathbf{K}_+ \text{RV}[k, \cdot]$.

5.7 RV-products and special bijections

Convention 5.7.1. Since definably bijective subsets are to be identified, we shall tacitly substitute $\mathbf{c}(X)$ for a subset X in the discussion if it is necessary or is just more convenient.

Definition 5.7.2. A subset \mathfrak{p} is an (*open, closed, rv-*) *polyball* if it is of the form $\prod_{i \leq n} \mathfrak{b}_i \times \bar{t}$, where each \mathfrak{b}_i is an (open, closed, rv-) ball and $\bar{t} \in \text{RV}$. In this case, the *radius* of \mathfrak{p} , denoted as $\text{rad}(\mathfrak{p})$, is $\min \{\text{rad}(\mathfrak{b}_i) : i \leq n\}$.

For any definable subset X , both the subset of X that contains all the rv-polyballs contained in X and the superset of X that contains all the rv-polyballs with nonempty intersection with X are definable.

Definition 5.7.3. For any subset $U \subseteq \text{VF}^n \times \text{RV}^m$, the *RV-hull* of U , denoted by $\text{RVH}(U)$, is the subset $\bigcup \{\text{rv}^{-1}(\bar{t}) \times \{\bar{s}\} : (\bar{t}, \bar{s}) \in \text{rv}(U)\}$. If $U = \text{RVH}(U)$, that is, if U is a union of rv-polyballs, then we say that U is an *RV-product*.

Lemma 5.7.4. *Let $X \subseteq (\text{VF}^{n_1} \times \text{RV}^{m_1}) \times (\text{VF}^{n_2} \times \text{RV}^{m_2})$ be a definable subset such that, for each $(\bar{a}, \bar{t}) \in \text{pr}_{n_1+m_1} X$, $\text{fib}(X, (\bar{a}, \bar{t}))$ is finite. Suppose that $Y \subseteq \text{VF}^{n_1+n_2} \times \text{RV}^m$ is an RV-product that is definably bijective to X . Then, for any rv-polyball*

$$\text{rv}^{-1}(t_1, \dots, t_{n_1+n_2}) \times \{(t_1, \dots, t_{n_1+n_2}, \bar{s})\} \subseteq Y,$$

the weight of $(t_1, \dots, t_{n_1+n_2})$ is at most n_1 .

Proof. Clearly we have that $\dim_{\text{VF}} X = \dim_{\text{VF}}(\text{pr}_{n_1+m_1} X) \leq n_1$. Suppose for contradiction that there is an rv-polyball contained in Y such that the weight of the tuple in question is greater than n_1 . By Corollary 5.6.5, $\dim_{\text{VF}} Y > n_1$ and hence $\dim_{\text{VF}} X > n_1$, contradiction. \square

Definition 5.7.5. Let $X \subseteq \text{VF} \times \text{VF}^n \times \text{RV}^m$. Let $C \subseteq \text{RVH}(X)$ be an RV-product and $\lambda : \text{pr}_{>1}(C \cap X) \longrightarrow \text{VF}$ a function such that $(\lambda(\bar{a}_1, \bar{t}), \bar{a}, \bar{t}) \in C$ for every $(\bar{a}_1, \bar{t}) = (\bar{a}_1, t_1, \dots, t_m) \in \text{pr}_{>1}(C \cap X)$. Let

$$C^\sharp = \bigcup_{(\bar{a}_1, \bar{t}) \in \text{pr}_{>1} C} \left(\left(\bigcup \{ \text{rv}^{-1}(t) : v_{\text{rv}}(t) > v_{\text{rv}}(t_1) \} \right) \times (\bar{a}_1, \bar{t}) \right),$$

$$\text{RVH}(X)^\sharp = C^\sharp \uplus (\text{RVH}(X) \setminus C).$$

The *centripetal transformation* $\eta : X \longrightarrow \text{RVH}(X)^\sharp$ with respect to λ is defined by

$$\eta(a_1, \bar{a}_1, \bar{t}) = (a_1 - \lambda(\bar{a}_1, \bar{t}), \bar{a}_1, \bar{t})$$

on $C \cap X$ and $\eta = \text{id}$ on $X \setminus C$. Note that η is injective. The inverse of η is naturally called the *centrifugal transformation with respect to λ* . The function λ is called a *focus map of X* . The RV-product C is called the *locus* of λ . A *special bijection* T is an alternating composition of centripetal transformations and the canonical bijection. The *length* of a special bijection T , denoted by $\text{lh} T$, is the number of centripetal transformations in the composition of T . The image $T(X)$ is sometimes denoted as X^\sharp .

Note that we should have included the index of the targeted VF-coordinate as a part of the data of a focus map. Since it should not cause confusion, below, we shall suppress mentioning it for notational ease.

Clearly if X is an RV-product and T is a special bijection on X then $T(X)$ is an RV-product. Notice that a special bijection T on X is definable if X and all the focus maps involved are definable. Since we are only interested in definable subsets and definable functions on them, we further require a special bijection to be definable.

Example 5.7.6. Let $\mathfrak{b} \subseteq \text{VF}$ be a definable open ball properly contained in $\text{rv}^{-1}(t)$. By Convention 5.7.1, \mathfrak{b} is identified with the subset $\mathfrak{b} \times \{t\}$. By Lemma 5.4.17, $\text{rv}^{-1}(t)$ contains a definable element a , which may or may not be in \mathfrak{b} . Let λ be the focus map $t \longmapsto a$. Then the centripetal transformation on \mathfrak{b} with respect to λ is given by $(b, t) \longmapsto (b - a, t)$.

Let $F \subseteq \text{rv}^{-1}(t) \times \text{rv}^{-1}(s) \subseteq \text{VF}^2$ be a definable finite-to-one function, which may be regarded as a focus map of itself whose locus is $\text{rv}^{-1}(t) \times \text{rv}^{-1}(s)$. Let η_1 be the corresponding centripetal transformation. Then $\eta_1(F) = \text{dom}(F) \times \{0\}$. For each $b \in \text{ran}(F)$ let b^* be the average of $F^{-1}(b)$. Note that, by compactness, the subset $\{b^* : b \in \text{ran}(F)\}$ is definable. Let $\lambda_2 : \text{ran}(F) \longrightarrow \text{dom}(F)$ be the focus map given by $(a, b) \longmapsto (b^*, b)$. Let η_2 be the corresponding centripetal transformation and $F^* = (c \circ \eta_2)(F)$. Notice that, by Lemma 5.4.1, rv is not constant on the subset $F^{-1}(b) - b^*$. Hence F^* is a function from $\text{pr}_1 F^*$ onto $\text{pr}_2 F^*$ such that the maximum size of its fibers on the first VF-coordinate is strictly smaller than that of F . This phenomenon

will be the basis of many inductive arguments below.

Definition 5.7.7. A subset X is a *deformed RV-product* if there is a special bijection T such that $T(X)$ is an RV-product. In that case, if T is definable then we say that X is a *definable deformed RV-product*.

Lemma 5.7.8. *Every definable subset $X \subseteq \text{VF} \times \text{RV}^m$ is a definable deformed RV-product.*

Proof. By compactness, it is enough to show that, for every $(a, \bar{t}) \in X$, there is a special bijection T on X such that $T(a, \bar{t})$ is contained in an rv-polyball $\mathfrak{p} \subseteq T(X)$. Fix an $(a, \bar{t}) \in X$. Let Z be the union of the rv-polyballs contained in X , which is a definable RV-product. If $(a, \bar{t}) \in Z$ then the canonical bijection is as required. So, without loss of generality, we may assume that $Z = \emptyset$. By Convention 5.7.1, the canonical bijection has been applied to X and hence, for any $\bar{s} = (s_1, \dots, s_m) \in \text{pRV } X$, the \bar{s} -definable subset $\text{fib}(X, \bar{s})$ is properly contained in the rv-ball $\text{rv}^{-1}(s_1)$.

By C -minimality, $\text{fib}(X, \bar{t})$ is a disjoint union of \bar{t} -definable simplexes. Let \mathfrak{s} be the simplex that contains (a, \bar{t}) . Let $\mathfrak{b}_1, \dots, \mathfrak{b}_l, \mathfrak{h}_1, \dots, \mathfrak{h}_n$ be the boolean components of \mathfrak{s} , where each \mathfrak{b}_i is positive and each \mathfrak{h}_i is negative. The proof now proceeds by induction on n .

For the base case $n = 0$, \mathfrak{s} is a disjoint union of balls $\mathfrak{b}_1, \dots, \mathfrak{b}_l$ of the same radius and valuative center. Without loss of generality, we may assume that $a \in \mathfrak{b}_1$. Let $\{\mathfrak{c}_1, \dots, \mathfrak{c}_k\}$ be the positive closure of \mathfrak{s} . Note that this closure is also \bar{t} -definable. We now start a secondary induction on k . For the base case $k = 1$, by Lemma 5.4.16, there is a \bar{t} -definable point $c \in \mathfrak{c}_1$. Clearly $\mathfrak{c}_1 - c \subseteq \text{rv}^{-1}(\text{rv}(a)) - c$ is a union of rv-balls. Let C be a definable subset of $\text{RVH}(X)$ and $\lambda : \text{pr}_{>1}(C \cap X) \rightarrow \text{VF}$ a definable focus map such that $(a, \bar{t}) \in C$ and $\lambda(\bar{t}) = c$. Then the centripetal transformation η with respect to λ is as desired. For the inductive step of the secondary induction, by Lemma 5.4.16 again, there is a \bar{t} -definable set of centers $\{c_1, \dots, c_k\}$ with $c_i \in \mathfrak{c}_i$. Let c be the average of c_1, \dots, c_k . Let λ, η be as above such that $\lambda(\bar{t}) = c$. If $c \in \mathfrak{b}_1$ then, as above, the centripetal transformation η with respect to λ is as desired. So suppose that $c \notin \mathfrak{b}_1$. Note that if val is not constant on the set $\{c_1 - c, \dots, c_k - c\}$ then rv is not constant on it and if val is constant on it then, by Lemma 5.4.1, rv is still not constant on it. Consider the special bijection $\sigma = \mathbf{c} \circ \eta$. We have that $\sigma(a, \bar{t}) = (a - c, r, \bar{t}) \in \sigma(X)$, where $r = \text{rv}(a - c)$. Observe that the positive closure of the (r, \bar{t}) -definable subset $\text{fib}(\sigma(X), (r, \bar{t}))$ is a proper subset of the set $\{\mathfrak{c}_1 - c, \dots, \mathfrak{c}_k - c\}$ of closed balls. Hence, by the inductive hypothesis, there is a special bijection T on $\sigma(X)$ such that $T(a - c, r, \bar{t})$ is contained in an rv-polyball $\mathfrak{p} \subseteq T \circ \sigma(X)$. So $T \circ \sigma$ is as required. This completes the base case $n = 0$.

We proceed to the inductive step. Note that, since $\mathfrak{b}_1, \dots, \mathfrak{b}_l$ are pairwise disjoint, the holes $\mathfrak{h}_1, \dots, \mathfrak{h}_n$ are also pairwise disjoint. Without loss of generality we may assume that all the holes $\mathfrak{h}_1, \dots, \mathfrak{h}_n$ are of the same radius. Let $\{\mathfrak{c}_1, \dots, \mathfrak{c}_k\}$ be the positive closure of $\bigcup_i \mathfrak{h}_i$. The secondary induction on k above may be

carried out here almost verbatim. Only note that, in the inductive step, after applying the special bijection σ , the number of holes in the fiber that contains $\sigma(a, \bar{t})$ decreases and hence the inductive hypothesis may be applied. \square

Corollary 5.7.9. *Let $f : X \rightarrow Y$ be a definable surjective function, where $X, Y \subseteq \text{VF}$. Then there is a definable function $P : X \rightarrow \text{RV}^m$ such that, for each $\bar{t} \in \text{ran } P$, $P^{-1}(\bar{t})$ is an open ball or a point and $f \upharpoonright P^{-1}(\bar{t})$ is either constant or injective.*

Proof. Let $P_1 : X \rightarrow \text{RV}^l$ be a function given by Lemma 5.4.10. Applying Lemma 5.7.8 to each fiber $P_1^{-1}(\bar{t})$, we see that desired function exists by compactness. \square

Remark 5.7.10. Corollary 5.7.9 and Lemma 5.4.8 imply that the theory ACVF_S^0 is b -minimal, in the sense of [13].

Lemma 5.7.11. *Let $X \subseteq \text{VF} \times \text{RV}^m$ be a definable subset and T a special bijection on X such that $T(X)$ is an RV -product. Then there is a definable function $\epsilon : (\text{pRV} \circ T)(X) \rightarrow \text{VF}$ such that, for every $(t, \bar{s}) \in (\text{pRV} \circ T)(X)$, we have that*

$$(\text{pVF} \circ T^{-1})(\text{rv}^{-1}(t) \times \{(t, \bar{s})\}) = \text{rv}^{-1}(t) + \epsilon(t, \bar{s}).$$

Proof. We do induction on the length of T . For the base case $\text{lh } T = 1$, let $T = \mathbf{c} \circ \eta$, where η is a centripetal transformation. Let λ and $C \subseteq \text{RVH}(X)$ be the corresponding focus map and its locus. For each $(t, \bar{s}) \in (\text{pRV} \circ T)(X)$, if $\bar{s} \in \text{dom}(\lambda)$ then set $\epsilon(t, \bar{s}) = \lambda(\bar{s})$, otherwise set $\epsilon(t, \bar{s}) = 0$. Clearly ϵ is as required.

We proceed to the inductive step $\text{lh } T = n > 1$. Let $T = \mathbf{c} \circ \eta_n \circ \dots \circ \mathbf{c} \circ \eta_1$ and $T_1 = \mathbf{c} \circ \eta_n \circ \dots \circ \mathbf{c} \circ \eta_2$. By the inductive hypothesis, for the special bijection T_1 , there is a function $\epsilon_1 : (\text{pRV} \circ T_1)((\mathbf{c} \circ \eta_1)(X)) \rightarrow \text{VF}$ as required. Let λ and $C \subseteq \text{RVH}(X)$ be the focus map and its locus for the centripetal transformation η_1 . For each $(t, \bar{s}) \in (\text{pRV} \circ T)(X)$, if $(\text{pRV} \circ T_1^{-1})(\text{rv}^{-1}(t) \times \{(t, \bar{s})\}) = (r, \bar{u})$ and $\bar{u} \in \text{dom}(\lambda)$ then set $\epsilon(t, \bar{s}) = \epsilon_1(t, \bar{s}) + \lambda(\bar{u})$, otherwise set $\epsilon(t, \bar{s}) = \epsilon_1(t, \bar{s})$. Then ϵ is as required. \square

Remark 5.7.12. Note that, in Lemma 5.7.11, since $\text{dom}(\epsilon) \subseteq \text{RV}^l$ for some l , by Lemma 5.4.8, $\text{ran}(\epsilon)$ is actually finite.

The following technical result is very important for the rest of the construction.

Proposition 5.7.13. *Let $f_i(\bar{x}) = f_i(x_1, \dots, x_n) \in \text{VF}(\langle \emptyset \rangle)[\bar{x}]$ be a finite list of polynomials and $\bar{t} = (t_1, \dots, t_n) \in \text{RV}$ a definable tuple. Then there is a special bijection T on $\text{rv}^{-1}(\bar{t})$ such that, for every rv -polyball $\mathbf{p} \subseteq T(\text{rv}^{-1}(\bar{t}))$ and every $f_i(\bar{x})$, the subset $f_i(T^{-1}(\mathbf{p}))$ is contained in an rv -ball.*

Proof. We do induction on n . For the base case $n = 1$, we write t, x for \bar{t}, \bar{x} , respectively. By compactness, it is enough to show that for any $a \in \text{rv}^{-1}(t)$ there is a special bijection T on $\text{rv}^{-1}(t)$ such that the image of $\text{RVH}(T(a))$ under every composite map

$$\text{RVH}(T(a)) \xrightarrow{T^{-1}} \text{rv}^{-1}(t) \xrightarrow{f_l} \text{VF} \xrightarrow{\text{rv}} \text{RV}$$

is a singleton. So fix an $a \in \text{rv}^{-1}(t)$. For any special bijection T on $\text{rv}^{-1}(t)$, let $k(T)$ be the number of elements $(b, \bar{s}) \in \text{RVH}(T(a))$ such that $f_l(T^{-1}(b, \bar{s})) = 0$ for some l . It is sufficient to prove the following:

Claim. For every special bijection T on $\text{rv}^{-1}(t)$ there is a special bijection T^* on $T(\text{rv}^{-1}(t))$ such that the image of $\text{RVH}((T^* \circ T)(a))$ under every composite map

$$\text{RVH}((T^* \circ T)(a)) \xrightarrow{(T^* \circ T)^{-1}} \text{rv}^{-1}(t) \xrightarrow{f_l} \text{VF} \xrightarrow{\text{rv}} \text{RV}$$

is a singleton.

Proof. We do induction on $k(T)$. For the base case $k(T) = 1$, there is a definable subset $Y \subseteq T(\text{rv}^{-1}(t))$ such that Y is the union of those rv -polyballs that contain exactly one $d \in \text{VF}$ with $f_l(T^{-1}(d)) = 0$ for some l . So there is a definable focus map $\lambda : \text{pRV} Y \rightarrow \text{VF}$ such that for every $\text{rv}^{-1}(s) \times \{(s, \bar{r})\} \subseteq Y$ we have that $f_l(T^{-1}(\lambda(s, \bar{r}))) = 0$ for some l . Clearly the special bijection T^* given by

$$(b, s, \bar{r}) \mapsto (b - \lambda(s, \bar{r}), \text{rv}(b - \lambda(s, \bar{r})), s, \bar{r})$$

for $(b, s, \bar{r}) \in Y$ is as required. For the inductive step $k(T) > 1$, there is a definable subset $Y \subseteq T(\text{rv}^{-1}(t))$ such that Y is the union of those rv -polyballs that contain exactly $k(T)$ elements $d \in \text{VF}$ with $f_l(T^{-1}(d)) = 0$ for some l . Let $\text{rv}^{-1}(s) \times \{(s, \bar{r})\} \subseteq Y$ and $d_1, \dots, d_{k(T)} \in \text{rv}^{-1}(s)$ enumerate these $k(T)$ elements. The focus map λ defined above is now modified so that $\lambda(s, \bar{r})$ is the average d of $d_1, \dots, d_{k(T)}$. Let T^* be defined as above with respect to this λ . By Lemma 5.4.1, rv is not constant on the set $\{d_1 - d, \dots, d_{k(T)} - d\}$. Hence $k(T^* \circ T) < k(T)$ and the inductive hypothesis may be applied. \square

This completes the base case of the induction.

We now proceed to the inductive step. Let $\bar{x}_1 = (x_2, \dots, x_n)$ and $\bar{t}_1 = (t_2, \dots, t_n)$. For every $a \in \text{rv}^{-1}(t_1)$, by the inductive hypothesis, there is an a -definable special bijection T_a on $\text{rv}^{-1}(\bar{t}_1)$ such that every function $\text{rv}(f_l(a, \bar{x}_1))$ is constant on every subset $T_a^{-1}(\mathfrak{p})$, where \mathfrak{p} is an rv -polyball contained in $T_a(\text{rv}^{-1}(\bar{t}_1))$. By compactness, there are definable disjoint subsets $Y_1, \dots, Y_m \subseteq \text{rv}^{-1}(t_1)$ with $\bigcup_i Y_i = \text{rv}^{-1}(t_1)$ and formulas

$\phi_1(x_1), \dots, \phi_m(x_1)$ such that, for every $a \in Y_i$, $\phi_i(a)$ defines a special bijection T_a on $\text{rv}^{-1}(\bar{t}_1)$ such that the property just described holds with respect to T_a . Applying Lemma 5.7.8 repeatedly, we obtain a special bijection T_1 on $\text{rv}^{-1}(t_1)$ such that each $T_1(Y_i)$ is an RV-product.

Now, for every $a \in \text{rv}^{-1}(t_1)$, each locus C involved in T_a is determined by an a -definable subset U_C of RV^k for some k . Let $\chi(x_1)$ be the formula that defines U_C . Note that if T_a is defined by $\phi_i(a)$ then $\chi(x_1)$ may be taken as a subformula of $\phi_i(x_1)$. Let $\chi^*(x_1, \bar{z})$ be a quantifier-free formula in disjunctive normal form that is equivalent to the formula $\chi(T_1^{-1}(x_1, \bar{z}))$, where \bar{z} are RV-sort variables. By compactness and the base case above, there is a special bijection ρ_i on $T_1(Y_i)$ such that each term $\text{rv}(g(x_1))$ that occurs in $\chi^*(x_1, \bar{z})$ is constant on every subset $\rho_i^{-1}(\mathbf{p})$, where \mathbf{p} is an rv-polyball contained in $(\rho_i \circ T_1)(Y_i)$. Hence, for each $a \in (\rho_i \circ T_1)^{-1}(\mathbf{p})$, $\chi(a)$ defines the same loci for the corresponding centripetal transformations. Consequently, by compactness again, we obtain a special bijection ρ on $T_1(\text{rv}^{-1}(t_1))$ such that $(\rho \circ T_1)(Y_i)$ is an RV-product for each i and, for each rv-polyball $\mathbf{p} \subseteq (\rho \circ T_1)(Y_i)$ with $\text{pRV } \mathbf{p} = \bar{s}$, the formula $\phi_i((\rho \circ T_1)^{-1}(x_1, \bar{s}))$ defines a special bijection on $\mathbf{p} \times \text{rv}^{-1}(\bar{t}_1)$. So $\phi_i((\rho \circ T_1)^{-1}(x_1, \bar{s}))$ defines a special bijection on $(\rho \circ T_1)(\text{rv}^{-1}(t_1)) \times \text{rv}^{-1}(\bar{t}_1)$, which is denoted as ϕ_i .

It is not hard to see that the special bijections $\phi_1 \circ \rho \circ T_1, \dots, \phi_m \circ \rho \circ T_1$ actually form one special bijection T_2 on $\text{rv}^{-1}(\bar{t})$. Let $\text{rv}^{-1}(\bar{s}) \times \{(\bar{s}, \bar{r})\} \subseteq T_2(\text{rv}^{-1}(\bar{t}))$, where $\bar{s} = (s_1, \dots, s_n)$. Let $\bar{s}_1 = (s_2, \dots, s_n)$. By the construction of T_2 , for each $a_1 \in \text{rv}^{-1}(s_1)$, every function $\text{rv}(f_l(\bar{x}))$ is constant on the subset

$$T_2^{-1}(\{a_1\} \times \text{rv}^{-1}(\bar{s}_1) \times \{(\bar{s}, \bar{r})\}).$$

Let this constant value be $u_{a_1}^l$. So the function $h_l : \text{rv}^{-1}(s_1) \rightarrow \text{RV}$ given by $a_1 \mapsto u_{a_1}^l$ is (\bar{s}, \bar{r}) -definable. For each l , let $\psi_l(x_1, z)$ be a quantifier-free formula in disjunctive normal form that defines the function h_l , where z is an RV-sort variable. We may assume that some conjunct in each disjunct of $\psi_l(x_1, z)$ is an RV-sort equality. Let $g_i(x_1)$ enumerate all the polynomials that occur in a term of the form $\text{rv}(g_i(x_1))$ in some $\psi_l(x_1, z)$. By the base case, there is an (\bar{s}, \bar{r}) -definable special bijection T_{s_1} on $\text{rv}^{-1}(s_1)$ such that, for each rv-polyball $\mathbf{p} \subseteq T_{s_1}(\text{rv}^{-1}(s_1))$, every term $\text{rv}(g_i(x_1))$ is constant on $T_{s_1}^{-1}(\mathbf{p})$ and hence every function h_l is constant on $T_{s_1}^{-1}(\mathbf{p})$. We may identify T_{s_1} with the function it naturally induces on $\text{rv}^{-1}(\bar{s}) \times \{(\bar{s}, \bar{r})\}$. Therefore, every function $\text{rv}(f_l(\bar{x}))$ is constant on the subset

$$(T_{s_1} \circ T_2)^{-1}(\mathbf{p} \times \text{rv}^{-1}(\bar{s}_1) \times \{(\bar{s}, \bar{r})\}).$$

By compactness, we obtain a special bijection T_3 on $T_2(\text{rv}^{-1}(\bar{t}))$ such that the property just described holds for every rv-polyball contained in $(T_3 \circ T_2)(\text{rv}^{-1}(\bar{t}))$. This completes the inductive step. \square

Lemma 5.7.8 is easily generalized as follows.

Proposition 5.7.14. *Every definable subset $X \subseteq \text{VF}^n \times \text{RV}^m$ is a definable deformed RV-product.*

Proof. This is by induction on n . The base case $n = 1$ is just Lemma 5.7.8. For the inductive step, by compactness, without loss of generality, we may assume that $\text{pRV } X$ is a singleton $\bar{t} = (t_1, \dots, t_n) \in \text{RV}$. It is not hard to see that, if we replace the conclusion of Proposition 5.7.13 with the desired property here, then a simpler version of the argument in the inductive step of the proof of Proposition 5.7.13 works almost verbatim. It is simpler because the last part of that argument is not needed here. \square

Corollary 5.7.15. *The map $\mathbb{L} : \text{Ob RV}[k, \cdot] \longrightarrow \text{Ob VF}[k, \cdot]$ is surjective on the isomorphism classes of $\text{VF}[k, \cdot]$.*

Corollary 5.7.16. *Let $f_l(\bar{x}) \in \text{VF}(\langle \emptyset \rangle)[\bar{x}]$ be a finite list of polynomials and X a definable subset of $\text{VF}^n \times \text{RV}^m$. Then there is a special bijection T on X such that $T(X)$ is an RV-product and, for every rv-polyball $\mathfrak{p} \subseteq T(X)$, every subset $f_l(T^{-1}(\mathfrak{p}))$ is contained in an rv-ball.*

Proof. By Proposition 5.7.14, there is a special bijection T_1 on X such that $T_1(X)$ is an RV-product. Let $\mathfrak{p} = \text{rv}^{-1}(\bar{t}) \times \{(\bar{t}, \bar{s})\}$ be an rv-polyball contained in $T_1(X)$. For each l , let ψ_l be a quantifier-free formula in disjunctive normal form that defines the function $\text{rv}(f_l(T_1^{-1}(\bar{x}, \bar{t}, \bar{s})))$ on \mathfrak{p} . Clearly we may assume that some conjunct in each disjunct of any ψ_l is an RV-sort equality. By Proposition 5.7.13, there is a (\bar{t}, \bar{s}) -definable special bijection $T_{\bar{t}, \bar{s}}$ on \mathfrak{p} such that each term $\text{rv}(g(\bar{x}))$ that occurs in some ψ_l is constant on every subset $T_{\bar{t}, \bar{s}}^{-1}(\mathfrak{q})$, where \mathfrak{q} is an rv-polyball contained in $T_{\bar{t}, \bar{s}}(\mathfrak{p})$, and hence $\text{rv}(f_l(\bar{x}))$ is constant on $(T_{\bar{t}, \bar{s}} \circ T_1)^{-1}(\mathfrak{q})$. By compactness, there is a special bijection T_2 on $T_1(X)$ such that the property just described holds for every rv-polyball contained in $(T_2 \circ T_1)(X)$. \square

Proposition 5.7.17. *Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be a definable subset. If $\text{pVF} \upharpoonright X$ is finite-to-one then there is a $Y \subseteq \text{RV}^l$ such that $\text{pr}_{\leq n} \upharpoonright Y$ is finite-to-one and $\mathbb{L}(Y, \text{pr}_{\leq n})$ is definably bijective to X .*

Proof. By Proposition 5.7.14, there is a $Y \subseteq \text{RV}^l$ such that there is a definable bijection $T : X \longrightarrow \mathbb{L}(Y, \text{pr}_{\leq n})$. Suppose for contradiction that there is a $\bar{t} \in \text{pr}_{\leq n} Y$ such that the subset $\text{fib}(Y, \bar{t})$ is infinite. Fix a tuple $\bar{a} \in \bar{t}$. Consider the \bar{a} -definable function $\text{pVF} \circ T^{-1} : \{\bar{a}\} \times \text{fib}(Y, \bar{t}) \longrightarrow \text{pVF } X$. By Lemma 5.4.8, $\text{ran}(\text{pVF} \circ T^{-1})$ is finite. Since $\text{pVF} \upharpoonright X$ is finite-to-one, we must have that the subset $T^{-1}(\{\bar{a}\} \times \text{fib}(Y, \bar{t}))$ is finite and hence $\{\bar{a}\} \times \text{fib}(Y, \bar{t})$ is finite, contradiction. \square

Corollary 5.7.18. *The map $\mathbb{L} : \text{Ob RV}[k] \longrightarrow \text{Ob VF}[k]$ is surjective on the isomorphism classes of $\text{VF}[k]$.*

5.8 2-cells

For functions between subsets that have only one VF-coordinate, composing with special bijections on the right and their inverses on the left preserves the open-to-open property.

Lemma 5.8.1. *Let $X, Y \subseteq \text{VF}$ be definable and $f : X \rightarrow Y$ a definable bijection. Then there is a special bijection T on X such that $T(X)$ is an RV-product and, for each rv-polyball $\mathfrak{p} \subseteq T(X)$, $f \upharpoonright T^{-1}(\mathfrak{p})$ has the open-to-open property.*

Proof. By Proposition 5.5.13, there is a definable finite partition of X such that the restriction of f to each piece has the open-to-open property. Applying Proposition 5.7.14 to each piece or its subsequent image yields the desired special bijection. \square

Lemma 5.8.2. *Let $X, Y \subseteq \text{VF}$ be definable open balls and $f : X \rightarrow Y$ a definable bijection that has the open-to-open property. Let $\alpha \in \Gamma$ be definable. Then there is a special bijection T on X such that $T(X)$ is an RV-product and, for each rv-polyball $\mathfrak{p} \subseteq T(X)$, the set*

$$\{\text{rad}(\mathfrak{b}) : \mathfrak{b} \text{ is an open ball contained in } T^{-1}(\mathfrak{p}) \text{ with } \text{rad}(f(\mathfrak{b})) = \alpha\}$$

is a singleton.

Proof. Let \mathfrak{B} be the collection of all open balls $\mathfrak{b} \subseteq X$ with $\text{rad}(f(\mathfrak{b})) = \alpha$. Let $\psi(x)$ be a quantifier-free formula in disjunctive normal form that defines the radius function rad on \mathfrak{B} , where x is the VF-sort variable. By Corollary 5.7.16, there is a special bijection T on X such that $T(X)$ is an RV-product and each term $\text{rv}(g(x))$ that occurs in $\psi(x)$ is constant on every subset $T^{-1}(\mathfrak{p})$, where \mathfrak{p} is an rv-polyball contained in $T(X)$. So T is as required. \square

Lemma 5.8.3. *Let $X \subseteq \text{VF}^2$ be a definable subset such that $\text{pr}_1 X$ is an open ball. Let $f : \text{pr}_1 X \rightarrow \text{pr}_2 X$ be a definable bijection that has the open-to-open property. Suppose that for each $a \in \text{pr}_1 X$ there is a $t_a \in \text{RV}$ such that*

$$\text{fib}(X, a) = \text{rv}^{-1}(t_a) + f(a).$$

Then there is a special bijection T on $\text{pr}_1 X$ such that $T(\text{pr}_1 X)$ is an RV-product and, for each rv-polyball $\mathfrak{p} \subseteq T(\text{pr}_1 X)$, the set

$$\{\text{rv}(a - f^{-1}(b)) : a \in T^{-1}(\mathfrak{p}) \text{ and } b \in \text{fib}(X, a)\}$$

is a singleton.

Proof. For each $a \in \text{pr}_1 X$, let \mathfrak{b}_a be the minimal closed ball that contains $\text{fib}(X, a)$. Since $\text{fib}(X, a) - f(a) = \text{rv}^{-1}(t_a)$, we have that $f(a) \in \mathfrak{b}_a$ but $f(a) \notin \text{fib}(X, a)$ if $t_a \neq \infty$. Hence $a \notin f^{-1}(\text{fib}(X, a))$ if $t_a \neq \infty$ and $\{a\} = f^{-1}(\text{fib}(X, a))$ if $t_a = \infty$. Since $f^{-1}(\text{fib}(X, a))$ is a ball, in either case, the function $\text{rv}(a - x)$ is constant on the subset $f^{-1}(\text{fib}(X, a))$. The function $h : \text{pr}_1 X \rightarrow \text{RV}$ given by $a \mapsto \text{rv}(a - f^{-1}(\text{fib}(X, a)))$ is definable. Now we may proceed as in Lemma 5.8.2. \square

Definition 5.8.4. Let $X \subseteq \text{VF}^2$ be such that $\text{pr}_1 X$ is an open ball. Let $f : \text{pr}_1 X \rightarrow \text{pr}_2 X$ be a bijection that has the open-to-open property. We say that f is *trapezoidal in X* if there are $t_1, t_2 \in \text{RV}$ such that, for each $a \in \text{pr}_1 X$,

1. $\text{fib}(X, a) = \text{rv}^{-1}(t_2) + f(a)$,
2. $f^{-1}(\text{fib}(X, a)) = a - \text{rv}^{-1}(t_1)$.

The elements t_1, t_2 are called the *paradigms* of X .

Remark 5.8.5. Let f be trapezoidal in X with respect to $t_1, t_2 \in \text{RV}$. Let $a \in \text{pr}_1 X$, \mathfrak{b} the minimal closed ball that contains $\text{fib}(X, a)$, and \mathfrak{a} the minimal closed ball that contains $f^{-1}(\text{fib}(X, a))$. The following properties are easily deduced:

1. $f(a) \notin \text{fib}(X, a)$ and hence $a \notin f^{-1}(\text{fib}(X, a))$.
2. $\text{v}_{\text{rv}}(t_1) = \text{rad}(\mathfrak{a}) > \text{rad}(\text{pr}_1 X)$ and $\text{v}_{\text{rv}}(t_2) = \text{rad}(\mathfrak{b}) > \text{rad}(\text{pr}_2 X)$.
3. $f(a) \in \mathfrak{b} \subseteq \text{pr}_2 X$ and $a \in \mathfrak{a} \subseteq \text{pr}_1 X$.
4. Let $\mathfrak{o}_a, \mathfrak{o}_{f(a)}$ be the maximal open subballs of $\mathfrak{a}, \mathfrak{b}$ that contains $a, f(a)$, respectively. We have that, for every $a^* \in f^{-1}(\mathfrak{o}_{f(a)})$,

$$\text{fib}(X, a^*) = \text{rv}^{-1}(t_2) + f(a^*) = \text{rv}^{-1}(t_2) + f(a) = \text{fib}(X, a)$$

and hence $a^* - \text{rv}^{-1}(t_1) = a - \text{rv}^{-1}(t_1)$; so $a^* \in \mathfrak{o}_a$. Symmetrically, for every $b^* \in f(\mathfrak{o}_a)$,

$$\begin{aligned} f^{-1}(\text{fib}(X, f^{-1}(b^*))) &= f^{-1}(b^*) - \text{rv}^{-1}(t_1) \\ &= a - \text{rv}^{-1}(t_1) \\ &= f^{-1}(\text{fib}(X, a)) \end{aligned}$$

and hence $\text{rv}^{-1}(t_2) + b^* = \text{rv}^{-1}(t_2) + f(a)$; so $b^* \in \mathfrak{o}_{f(a)}$. So actually $f(\mathfrak{o}_a) = \mathfrak{o}_{f(a)}$.

5. Let \mathfrak{A} , \mathfrak{B} be the sets of maximal open subballs of \mathfrak{a} , \mathfrak{b} , respectively. Then f induces a bijection $f_{\downarrow} : \mathfrak{A} \longrightarrow \mathfrak{B}$.

6. For each $\mathfrak{o} \in \mathfrak{A}$, each $c \in \mathfrak{o}$, and each $d \in \text{fib}(X, c)$, we have that

$$\begin{aligned}\text{fib}(X, c) &= \text{rv}^{-1}(t_2) + f_{\downarrow}(\mathfrak{o}), \\ \text{fib}(X, d) &= f^{-1}(d) + \text{rv}^{-1}(t_1) = \mathfrak{o}.\end{aligned}$$

So $\mathfrak{o} - f^{-1}(\text{fib}(X, c)) = \text{rv}^{-1}(t_1)$ and $f_{\downarrow}(\mathfrak{o}) - \text{fib}(X, c) = -\text{rv}^{-1}(t_2)$. (Hence f is “trapezoidal”.)

7. The subset X is symmetrical in the following way:

$$\begin{aligned}& \bigcup \{ \mathfrak{o} \times (\text{rv}^{-1}(t_2) + f_{\downarrow}(\mathfrak{o})) : \mathfrak{o} \in \mathfrak{A} \} \\ &= \bigcup \{ (f_{\downarrow}^{-1}(\mathfrak{o}) + \text{rv}^{-1}(t_1)) \times \mathfrak{o} : \mathfrak{o} \in \mathfrak{B} \} \\ &= X \cap (\mathfrak{a} \times \text{VF}) \\ &= X \cap (\text{VF} \times \mathfrak{b}) \\ &= X \cap (\mathfrak{a} \times \mathfrak{b}).\end{aligned}$$

Definition 5.8.6. We say that a subset X is a *1-cell* if it is either an open ball contained in one rv-ball or a point in VF. We say that X is a *2-cell* if

1. $X \subseteq \text{VF}^2$ is contained in one rv-polyball and $\text{pr}_1 X$ is a 1-cell,
2. there is a function $\epsilon : \text{pr}_1 X \longrightarrow \text{VF}$ and a $t \in \text{RV}$ such that, for each $a \in \text{pr}_1 X$, $\text{fib}(X, a) = \text{rv}^{-1}(t) + \epsilon(a)$,
3. one of the following three possibilities occurs:
 - (a) ϵ is constant,
 - (b) ϵ is injective, has the open-to-open property, and $\text{rad}(\epsilon(\text{pr}_1 X)) \geq v_{\text{rv}}(t)$,
 - (c) ϵ is trapezoidal in X .

The function ϵ is called the *positioning function* of X and the element t is called the *paradigm* of X .

Remark 5.8.7. A subset $X \subseteq \text{VF} \times \text{RV}^m$ is a *1-cell* if for each $\bar{t} \in \text{pRV } X$ the fiber $\text{fib}(X, \bar{t})$ is a 1-cell in the sense of Definition 5.8.6. The concept of a 2-cell is generalized in the same way. A unit is definable if all the relevant ingredients are definable.

Suppose that X is a 2-cell. Clearly if its paradigm t is ∞ then X and its positioning function ϵ are identical. It is also easy to see that, if $t \neq \infty$ and ϵ is not trapezoidal, then X is actually an open polyball.

Notice that Lemma 5.7.8 implies that for every definable subset $X \subseteq \text{VF} \times \text{RV}^m$ there is a definable function $P : X \rightarrow \text{RV}^l$ such that, for each $\bar{s} \in \text{ran } P$, the fiber $P^{-1}(\bar{s})$ is a 1-cell. The same holds for 2-cell:

Lemma 5.8.8. *For every definable subset $X \subseteq \text{VF}^2$ there is a definable function $P : X \rightarrow \text{RV}^m$ such that, for each $\bar{s} \in \text{ran } P$, the fiber $P^{-1}(\bar{s})$ is a 2-cell.*

Proof. Without loss of generality, we may assume that X is contained in one rv-polyball. For any $a \in \text{pr}_1 X$, by Lemma 5.7.8, there is an a -definable special bijection T_a on $\text{fib}(X, a)$ such that $T_a(\text{fib}(X, a))$ is an RV-product. By Lemma 5.7.11, there is an a -definable function $\epsilon_a : (\text{pRV} \circ T_a)(\text{fib}(X, a)) \rightarrow \text{VF}$ such that, for every $(t, \bar{s}) \in (\text{pRV} \circ T_a)(\text{fib}(X, a))$, we have that

$$(\text{pVF} \circ T_a^{-1})(\text{rv}^{-1}(t) \times \{(t, \bar{s})\}) = \text{rv}^{-1}(t) + \epsilon_a(t, \bar{s}).$$

By compactness, we may assume that there is a definable subset $X' \subseteq \text{pr}_1 X \times \text{RV}^l$ and a definable function $\epsilon : X' \rightarrow \text{VF}$ such that, for each $a \in \text{pr}_1 X$, $\text{fib}(X', a) = (\text{pRV} \circ T_a)(\text{fib}(X, a))$ and $\epsilon \upharpoonright \text{fib}(X', a) = \epsilon_a$. Since, for each $(t, \bar{s}) \in \text{pRV } X'$, $\epsilon \upharpoonright \text{fib}(X', (t, \bar{s}))$ may be regarded as a (t, \bar{s}) -definable function from VF into VF , by Lemma 5.4.10, we are reduced to the case that each $\epsilon \upharpoonright \text{fib}(X', (t, \bar{s}))$ is either constant or injective. If no $\epsilon \upharpoonright \text{fib}(X', (t, \bar{s}))$ is injective then we can finish by applying Lemma 5.7.8 to each $\text{fib}(X', (t, \bar{s}))$ and then compactness.

Suppose that $\epsilon_{(t, \bar{s})} = \epsilon \upharpoonright \text{fib}(X', (t, \bar{s}))$ is injective. By Lemma 5.8.1, we are reduced to the case that $\text{fib}(X', (t, \bar{s}))$ is an open ball and $\epsilon_{(t, \bar{s})}$ has the open-to-open property. Note that, if $\text{rad}(\text{ran } \epsilon_{(t, \bar{s})}) < v_{\text{rv}}(t)$, then

$$\text{ran } \epsilon_{(t, \bar{s})} = \bigcup_{a \in \text{fib}(X', (t, \bar{s}))} (\text{rv}^{-1}(t) + \epsilon_{(t, \bar{s})}(a)).$$

By Lemma 5.8.2, we are further reduced to the case that, if $\text{rad}(\text{ran } \epsilon_{(t, \bar{s})}) < v_{\text{rv}}(t)$, then there is a $\gamma \in \Gamma$ such that $\text{rad}(\epsilon_{(t, \bar{s})}^{-1}(\mathbf{b})) = \gamma$ for every open ball $\mathbf{b} \subseteq \text{ran } \epsilon_{(t, \bar{s})}$ with $\text{rad}(\mathbf{b}) = v_{\text{rv}}(t)$. By Lemma 5.8.3, we are finally reduced to the case that, if $\text{rad}(\text{ran } \epsilon_{(t, \bar{s})}) < v_{\text{rv}}(t)$, then there is an $r \in \text{RV}$ such that, for every $a \in \text{fib}(X', (t, \bar{s}))$,

$$\text{rv}(a - f^{-1}(\text{rv}^{-1}(t) + \epsilon_{(t, \bar{s})}(a))) = r$$

and hence

$$f^{-1}(\text{rv}^{-1}(t) + \epsilon_{(t, \bar{s})}(a)) = a - \text{rv}^{-1}(r).$$

So, in this case, $\epsilon_{(t, \bar{s})}$ is trapezoidal. Now we are done by compactness. \square

5.9 Lifting functions from RV to VF

We shall show that the map \mathbb{L} actually induces homomorphisms between various Grothendieck semigroups when S is a (VF, Γ) -generated substructure.

Any polynomial in $\mathcal{O}[\bar{x}]$ corresponds to a polynomial in $\overline{K}[\bar{x}]$ via the canonical quotient map. The following definition generalizes this phenomenon.

Definition 5.9.1. Let $\bar{\gamma} = (\gamma_1, \dots, \gamma_n) \in \Gamma$. A polynomial $f(\bar{x}) = \sum_{\bar{i}j} a_{\bar{i}j} \bar{x}^{\bar{i}}$ with coefficients $a_{\bar{i}j} \in \text{VF}$ is a $\bar{\gamma}$ -polynomial if there is an $\alpha \in \Gamma$ such that $\alpha = \text{val}(a_{\bar{i}j}) + i_1 \gamma_1 + \dots + i_n \gamma_n$ for each $\bar{i}j = (\bar{i}, j) = (i_1, \dots, i_n, j)$. In this case we say that α is a *residue value* of $f(\bar{x})$ (with respect to $\bar{\gamma}$). For a $\bar{\gamma}$ -polynomial $f(\bar{x})$ with residue value α and a $\bar{t} \in \text{RV}$ with $v_{\text{rv}}(\bar{t}) = \bar{\gamma}$, if $\text{val} f(\bar{a}) > \alpha$ for all $\bar{a} \in \text{rv}^{-1}(\bar{t})$ then \bar{t} is a *residue root* of $f(\bar{x})$. If $\bar{t} \in \text{RV}$ is a common residue root of the $\bar{\gamma}$ -polynomials $f_1(\bar{x}), \dots, f_n(\bar{x})$ but is not a residue root of the $\bar{\gamma}$ -polynomial $\det \partial(f_1, \dots, f_n) / \partial(x_1, \dots, x_n)$, then we say that $f_1(\bar{x}), \dots, f_n(\bar{x})$ are *minimal* for \bar{t} and \bar{t} is a *simple common residue root* of $f_1(\bar{x}), \dots, f_n(\bar{x})$.

Therefore, according to this definition, every polynomial in $\overline{K}[\bar{x}]$ is the projection of some $(0, \dots, 0)$ -polynomial $f(\bar{x})$ with residue value 0.

Hensel's Lemma is generalized as follows.

Lemma 5.9.2 (Generalized Hensel's Lemma). *Let $f_1(\bar{x}), \dots, f_n(\bar{x})$ be $\bar{\gamma}$ -polynomials with residue values $\alpha_1, \dots, \alpha_n$, where $\bar{\gamma} = (\gamma_1, \dots, \gamma_n) \in \Gamma$. For every simple common residue root $\bar{t} = (t_1, \dots, t_n) \in \text{RV}$ of $f_1(\bar{x}), \dots, f_n(\bar{x})$ there is a unique $\bar{a} \in \text{rv}^{-1}(\bar{t})$ such that $f_i(\bar{a}) = 0$ for every i .*

Proof. Fix a simple common residue root $\bar{t} = (t_1, \dots, t_n) \in \text{RV}$ of $f_1(\bar{x}), \dots, f_n(\bar{x})$. Choose a $c_i \in \text{rv}^{-1}(t_i)$. Changing the coefficients accordingly we may rewrite each $f_i(\bar{x})$ as $f_i(x_1/c_1, \dots, x_n/c_n)$. Write y_i for x_i/c_i . Note that, for each i , the coefficients of the $(0, \dots, 0)$ -polynomial $f_i(\bar{y})$ are all of the same value α_i . For each i choose an $e_i \in \text{VF}$ with $\text{val}(e_i) = -\alpha_i$. We have that each $(0, \dots, 0)$ -polynomial $f_i^*(\bar{y}) = e_i f_i(\bar{y})$ has residue value 0 (that is, the coefficients of $f_i^*(\bar{y})$ is of value 0). Clearly $(t_1/\text{rv}(c_1), \dots, t_n/\text{rv}(c_n)) = (1, \dots, 1)$ is a common residue root of $f_1^*(\bar{y}), \dots, f_n^*(\bar{y})$; that is, for every $\bar{a} \in \text{rv}^{-1}(1, \dots, 1)$ and every i we have that $\text{val} f_i^*(\bar{a}) > 0$. It is actually a simple root because for every $\bar{a} \in \text{rv}^{-1}(1, \dots, 1)$ we have that

$$\det \partial(f_1^*, \dots, f_n^*) / \partial(y_1, \dots, y_n)(\bar{a}) = \left(\prod_i e_i c_i \right) \cdot \det \partial(f_1, \dots, f_n) / \partial(x_1, \dots, x_n)(\bar{a}c),$$

where $\bar{a}c = (a_1c_1, \dots, a_nc_n)$, and hence

$$\text{val}(\det \partial(f_1^*, \dots, f_n^*) / \partial(y_1, \dots, y_n)(\bar{a})) = \sum_i (-\alpha_i + \gamma_i) + \sum_i \alpha_i - \sum_i \gamma_i = 0.$$

Now the lemma follows from the multivariate version of Hensel's Lemma (e.g. see [9, Corollary 2, p. 224]). \square

Definition 5.9.3. Let X, Y be two RV-products, F a subset of $X \times Y$, and A a subset of $\text{rv}(X \times Y)$. We say that F is a (X, Y) -lift of A from RV to VF, or just a lift of A for short, if $F \cap (\mathfrak{p} \times \mathfrak{q})$ is a bijective function from \mathfrak{p} onto \mathfrak{q} for every rv-polyball $\mathfrak{p} \subseteq X$ and every rv-polyball $\mathfrak{q} \subseteq Y$ with $\text{rv}(\mathfrak{p} \times \mathfrak{q}) \in A$. A partial lift of A is a lift of any subset of A .

It would be ideal to lift all definable subsets of $\text{RV}^n \times \text{RV}^n$ with finite-to-finite correspondence for any substructure S . However, the following crucial lemma fails when S is not (VF, Γ) -generated.

Lemma 5.9.4. Suppose that S is a (VF, Γ) -generated substructure. Let $\bar{t} = (t_1, \dots, t_n) \in \text{RV}$ be such that $t_n \neq \infty$ and $t_n \in \text{acl}(t_1, \dots, t_{n-1})$. Let $v_{\text{rv}}(\bar{t}) = (\gamma_1, \dots, \gamma_n) = \bar{\gamma}$. Then there is a $\bar{\gamma}$ -polynomial $f(x_1, \dots, x_n) = f(\bar{x})$ with coefficients in $\text{VF}(\langle \emptyset \rangle)$ such that the subset $\{r \in \text{RV} : (t_1, \dots, t_{n-1}, r) \text{ is a residue root of } f(\bar{x})\}$ is finite and \bar{t} is a residue root of $f(\bar{x})$ but is not a residue root of $\partial f(\bar{x}) / \partial x_n$.

Proof. Write (t_1, \dots, t_{n-1}) as \bar{t}_n . Let $\phi(\bar{x})$ be a formula such that $\phi(\bar{t}_n, x_n)$ defines a finite subset that contains t_n . By quantifier elimination, there is a conjunction $\psi(\bar{x})$ of RV-sort literals such that $\psi(\bar{x})$ implies $\phi(\bar{x})$ and $\psi(\bar{t})$ holds. By C -minimality, we may assume that some conjunct $\theta(\bar{x})$ in $\psi(\bar{x})$ is an RV-sort equality such that $\theta(\bar{t}_n, x_n)$ defines a finite subset. Since S is (VF, Γ) -generated, we may assume that $\theta(\bar{x})$ does not contain parameters from $\text{RV}(\langle \emptyset \rangle) \setminus \text{rv}(\text{VF}(\langle \emptyset \rangle))$. Hence it is of the form

$$\bar{x}^{\bar{k}} \cdot \sum_{\bar{i}} (\text{rv}(a_{\bar{i}}) \cdot \bar{x}^{\bar{i}}) = \text{rv}(a) \cdot \bar{x}^{\bar{l}} \cdot \sum_{\bar{j}} (\text{rv}(a_{\bar{j}}) \cdot \bar{x}^{\bar{j}}),$$

where $a_{\bar{i}}, a, a_{\bar{j}} \in \text{VF}(\langle \emptyset \rangle)$. Fix an $s \in \text{RV}$ such that $v_{\text{rv}}(s \cdot \bar{t}^{\bar{k}}) = v_{\text{rv}}(s \cdot \text{rv}(a) \cdot \bar{t}^{\bar{l}}) = 0$. Let $v_{\text{rv}}(s) = \delta$. Note that δ is \bar{t}_n -definable. Let $h_1(\bar{x}, s) = \sum_{\bar{i}} (s \cdot \text{rv}(a_{\bar{i}}) \cdot \bar{x}^{\bar{i}+\bar{k}})$ and $h_2(\bar{x}, s) = \sum_{\bar{j}} (-s \cdot \text{rv}(aa_{\bar{j}}) \cdot \bar{x}^{\bar{j}+\bar{l}})$. Consider the RV-sort polynomial $H(\bar{x}, s) = h_1(\bar{x}, s) + h_2(\bar{x}, s)$. For any $r \in \text{RV}$, $H(\bar{t}_n, s, r) = 0$ if and only if either

$$\sum_{\bar{i}} (\text{rv}(a_{\bar{i}}) \cdot (\bar{t}_n, r)^{\bar{i}}) = \sum_{\bar{j}} (\text{rv}(a_{\bar{j}}) \cdot (\bar{t}_n, r)^{\bar{j}}) = 0$$

or

$$\text{rv}(h_1(\bar{t}_n, s, r)/s) = \text{rv}(-h_2(\bar{t}_n, s, r)/s).$$

Since $r \neq t_n$ in the former case, by C -minimality again, the equation $H(\bar{t}_n, s, x_n) = 0$ defines a finite subset that contains the subset defined by $\theta(\bar{t}_n, x_n)$ and is actually \bar{t}_n -definable. Let m be the maximal exponent of x_n in $H(\bar{x}, s)$. For each $i \leq m$ let $H_i(\bar{x}, s)$ be the sum of all the monomials $h(\bar{x}, s)$ in $H(\bar{x}, s)$ such that the exponent of x_n in $h(\bar{x}, s)$ is i . Replacing s with a variable y and each $\text{rv}(a)$ with a in $H_i(\bar{x}, s)$, we obtain a VF-sort polynomial $H_i^*(\bar{x}, y)$ for each $i \leq m$. Let

$$E = \{i \leq m : v_{\text{rv}}(H_i^*(\bar{b}, c)) = 0 \text{ for all } (\bar{b}, c) \in \text{rv}^{-1}(\bar{t}, s)\}.$$

Since $H(\bar{t}, s) = 0$, clearly $|E| \neq 1$. We claim that $|E| > 1$. To see this, suppose for contradiction that $E = \emptyset$. Write $H_i^*(\bar{x}, y)$ as $yx_n^i G_i(\bar{x}_n)$, where $\bar{x}_n = (x_1, \dots, x_{n-1})$. Let $\bar{\gamma}_n = (\gamma_1, \dots, \gamma_{n-1})$. Clearly each $G_i(\bar{x}_n)$ is a $\bar{\gamma}_n$ -polynomial with residue value $-\delta - i\gamma_n$. Since $v_{\text{rv}}(cb_n^i G_i(\bar{b}_n)) > 0$ for all $c \in \text{rv}^{-1}(s)$, $b_n \in \text{rv}^{-1}(t_n)$, and $\bar{b}_n \in \text{rv}^{-1}(\bar{t}_n)$, we have that (\bar{t}_n) is a residue root of $G_i(\bar{x}_n)$. So for all $r \in \text{RV}$ with $v_{\text{rv}}(r) = v_{\text{rv}}(t_n) = \gamma_n$ we have that $v_{\text{rv}}(cd^i G_i(\bar{b}_n)) > 0$ for all $c \in \text{rv}^{-1}(s)$, $d \in \text{rv}^{-1}(r)$, and $\bar{b}_n \in \text{rv}^{-1}(\bar{t}_n)$ and hence $H_i(\bar{t}_n, s, r) = 0$. So $H(\bar{t}_n, s, r) = 0$ for all $r \in \text{RV}$ with $v_{\text{rv}}(r) = v_{\text{rv}}(t_n) = \gamma_n$, which is a contradiction because the equation $H(\bar{t}_n, s, x_n) = 0$ defines a finite subset.

Let $H^*(\bar{x}, y) = \sum_{i \in E} H_i^*(\bar{x}, y) = \sum_{i \in E} (yx_n^i G_i(\bar{x}_n)) = yG(\bar{x})$. Since (\bar{t}, s) is a residue root of $H^*(\bar{x}, y)$, clearly $G(\bar{x})$ is a $\bar{\gamma}$ -polynomial with residue value $-\delta$ such that \bar{t} is a residue root of it. Also, \bar{t}_n is not a residue root of any $G_i(\bar{x}_n)$. It follows that, for some $k < \max E$, \bar{t} is a residue root of the $\bar{\gamma}$ -polynomial $\partial G(\bar{x})/\partial^k x_n$ but is not a residue root of the $\bar{\gamma}$ -polynomial $\partial G(\bar{x})/\partial^{k+1} x_n$. \square

For definable subsets of the residue field, the situation may be further simplified. The following lemma shows that the geometry of definable subsets over the residue field coincides with its algebraic geometry; in other words, each definable subset over the residue field is a constructible subset (in the sense of algebraic geometry) of the Zariski topological space $\text{Spec } \bar{K}(S)[x_1, \dots, x_n]$.

Lemma 5.9.5. *If $X \subseteq \bar{K}^n$ is definable then it is a boolean combination of subsets defined by equalities with coefficients in $\bar{K}(S)$.*

Proof. Let ϕ be a quantifier-free formula in disjunctive normal form that defines X and $\bar{\gamma} = (\gamma_1, \dots, \gamma_m)$ the Γ -sort parameters in ψ . Without loss of generality $\gamma_i \notin \text{acl}(\text{VF}(\langle \emptyset \rangle), \text{RV}(\langle \emptyset \rangle))$ for all i . Since $X \subseteq \bar{K}^n$, each conjunct in each disjunct of ϕ may be assumed to be of the form

$$\sum_{\bar{i}} (\text{rv}(a_{\bar{i}}) \cdot r_{\bar{i}} \cdot \bar{x}^{\bar{i}}) \square \text{rv}(a) \cdot r \cdot \sum_{\bar{j}} (\text{rv}(a_{\bar{j}}) \cdot r_{\bar{j}} \cdot \bar{x}^{\bar{j}}),$$

where $a_{\bar{i}}, a, a_{\bar{j}} \in \text{VF}(\langle \emptyset \rangle)$, $r_{\bar{i}}, r, r_{\bar{j}} \in \text{RV}(\langle \emptyset \rangle)$, and \square is one of the symbols $=, \neq, \leq,$ and $>$. It is easily seen

that the literals involving \leq or $>$ are redundant. So each conjunct in ϕ is either an RV-sort equality or an RV-sort disequality. Now the proof proceeds by induction on m . The base case $m = 0$ is clear. For the inductive step, if one of the conjuncts in ϕ is an equality and contains some γ_i as an irredundant parameter then, since $X \subseteq \overline{K}^n$, actually γ_i may be defined from other parameters in ϕ and hence, by the inductive hypothesis, the lemma holds. By the same reason, we see that no nontrivial equality with parameters in $\langle \text{VF}(\langle \emptyset \rangle), \text{RV}(\langle \emptyset \rangle) \rangle$ may hold between $\overline{\gamma}$ and any $\overline{s} \in X$. So each disequality in ϕ that contains some γ_i as an irredundant parameter must define either the empty set or a superset of X and hence is redundant. \square

Proposition 5.9.6. *Suppose that the substructure S is (VF, Γ) -generated. Let $C \subseteq (\text{RV}^\times)^n \times (\text{RV}^\times)^n$ be a definable subset such that both $\text{pr}_{\leq n} \upharpoonright C$ and $\text{pr}_{> n} \upharpoonright C$ are finite-to-one. Then there is a definable subset $C^\uparrow \subseteq \text{VF}^n \times \text{VF}^n$ that lifts C .*

Proof. By compactness, the lemma is reduced to showing that for every $(\overline{t}, \overline{s}) \in C$ there is a definable lift of some subset of C that contains $(\overline{t}, \overline{s})$. Fix a $(\overline{t}, \overline{s}) \in C$ and set $(\overline{\gamma}, \overline{\delta}) = \text{v}_{\text{rv}}(\overline{t}, \overline{s})$. Let $\phi(\overline{x}, \overline{y})$ be a formula that defines C . Consider the formulas $\exists \overline{y}_i \phi(\overline{x}, \overline{y})$ and $\exists \overline{x}_i \phi(\overline{x}, \overline{y})$, where $\overline{y}_i = \overline{y} \setminus y_i$ and $\overline{x}_i = \overline{x} \setminus x_i$. By Lemma 5.9.4, for each y_i there is a $(\overline{\gamma}, \delta_i)$ -polynomial $\mu_i(\overline{x}, y_i)$ with coefficients in $\text{VF}(\langle \emptyset \rangle)$ such that (\overline{t}, s_i) is a residue root of $\mu_i(\overline{x}, y_i)$ but is not a residue root of $\partial \mu_i(\overline{x}, y_i) / \partial y_i$. Similarly we obtain such a $(\gamma_i, \overline{\delta})$ -polynomial $\nu_i(x_i, \overline{y})$ for each x_i . For each i , let $a_i(\overline{xy})^{\overline{k}_i}$ and $b_i(\overline{xy})^{\overline{l}_i}$ be two monomials with $a_i, b_i \in \text{VF}(\langle \emptyset \rangle)$ such that

$$\mu_i^*(\overline{x}, \overline{y}) + \nu_i^*(\overline{x}, \overline{y}) = a_i(\overline{xy})^{\overline{k}_i} \mu_i(\overline{x}, y_i) + b_i(\overline{xy})^{\overline{l}_i} \nu_i(x_i, \overline{y})$$

is a $(\overline{\gamma}, \overline{\delta})$ -polynomial. Let α_i be the residue value of $\mu_i^*(\overline{x}, \overline{y}) + \nu_i^*(\overline{x}, \overline{y})$. Note that for any $(\overline{a}, \overline{b}) \in \text{rv}^{-1}(\overline{t}, \overline{s})$ we have

$$\text{val}(\partial \mu_i^* / \partial y_i)(\overline{a}, \overline{b}) = \text{val}(a_i(\overline{ab})^{\overline{k}_i}) + \text{val}(\partial \mu_i / \partial y_i)(\overline{a}, \overline{b}) = \alpha_i - \delta_i$$

and for $j \neq i$ we have

$$\begin{aligned} \text{val}(\partial \mu_i^* / \partial y_j)(\overline{a}, \overline{b}) &= \text{val}(a_i) + \text{val}(\partial (\overline{xy})^{\overline{k}_i} / \partial y_j)(\overline{a}, \overline{b}) + \text{val} \mu_i(\overline{a}, b_i) \\ &> \alpha_i - \delta_j. \end{aligned}$$

Therefore,

$$\begin{aligned} \text{val} \det(\partial(\mu_1^*, \dots, \mu_n^*) / \partial(y_1, \dots, y_n))(\overline{a}, \overline{b}) &= \text{val} \prod_i (\partial \mu_i^* / \partial y_i)(\overline{a}, \overline{b}) \\ &= \sum_i \alpha_i - \sum_i \delta_i. \end{aligned}$$

This shows that \bar{s} is a simple common residue root of $\mu_1^*(\bar{a}, \bar{y}), \dots, \mu_n^*(\bar{a}, \bar{y})$ for any $\bar{a} \in \text{rv}^{-1}(\bar{t})$. Similarly \bar{t} is a simple common residue root of $\nu_1^*(\bar{x}, \bar{b}), \dots, \nu_n^*(\bar{x}, \bar{b})$ for any $\bar{b} \in \text{rv}^{-1}(\bar{s})$. Now, it is not hard to see that we may choose integers d_i, e_i and form a $(\bar{\gamma}, \bar{\delta})$ -polynomial

$$\tau_i(\bar{x}, \bar{y}) = d_i \mu_i^*(\bar{x}, \bar{y}) + e_i \nu_i^*(\bar{x}, \bar{y})$$

such that \bar{s} is a simple common residue root of $\tau_1(\bar{a}, \bar{y}), \dots, \tau_n(\bar{a}, \bar{y})$ for any $\bar{a} \in \text{rv}^{-1}(\bar{t})$ and \bar{t} is a simple common residue root of $\tau_1(\bar{x}, \bar{b}), \dots, \tau_n(\bar{x}, \bar{b})$ for any $\bar{b} \in \text{rv}^{-1}(\bar{s})$. By the generalized Hensel's Lemma 5.9.2, for each $\bar{a} \in \text{rv}^{-1}(\bar{t})$ there is a unique $\bar{b} \in \text{rv}^{-1}(\bar{s})$ such that $\bigwedge_i \tau_i(\bar{a}, \bar{b}) = 0$, and vice versa. \square

Corollary 5.9.7. *Suppose that the substructure S is (VF, Γ) -generated. The map \mathbb{L} induces homomorphisms between various Grothendieck semigroups: $\mathbf{K}_+ \text{RV}[k, \cdot] \longrightarrow \mathbf{K}_+ \text{VF}[k, \cdot]$, $\mathbf{K}_+ \text{RV}[k] \longrightarrow \mathbf{K}_+ \text{VF}[k]$, etc.*

Proof. For any $\text{RV}[k, \cdot]$ -isomorphism $F : (U, f) \longrightarrow (V, g)$ and any $\bar{u} \in U$, by definition, $\text{wgt } f(\bar{u}) = \text{wgt}(g \circ F)(\bar{u})$. Therefore, $\mathbb{L}(U, f)$ and $\mathbb{L}(V, g)$ are $\text{VF}[k, \cdot]$ -isomorphic by Proposition 5.9.6. \square

5.10 Contracting to RV

Definition 5.10.1. Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be an RV-product and $f : X \longrightarrow Y$ a function, where Y is also an RV-product. We say that f is *contractible* if for every rv-polyball $\mathfrak{p} \subseteq X$ the subset $f(\mathfrak{p})$ is contained in an rv-polyball.

Clearly, for two (definable) RV-products X and Y , if $f : X \longrightarrow Y$ is an (definable) contractible function, then there is a unique (definable) function $f_{\downarrow} : \text{rv}(X) \longrightarrow \text{rv}(Y)$ such that the diagram

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ \text{rv} \downarrow & & \downarrow \text{rv} \\ \text{rv}(X) & \xrightarrow{f_{\downarrow}} & \text{rv}(Y) \end{array}$$

commutes. Note that, in this case, if both f and f_{\downarrow} are bijective then f is a lift of f_{\downarrow} . Equivalently, if f is bijective and both f and f^{-1} are contractible then f is a lift of f_{\downarrow} .

Lemma 5.10.2. *Let $X \subseteq \text{VF}^{n_1} \times \text{RV}^{m_1}$ and $Y \subseteq \text{VF}^{n_2} \times \text{RV}^{m_2}$ be definable subsets and $f : X \longrightarrow Y$ a definable function. Then there exist special bijections T_X, T_Y on X, Y such that the function $T_Y \circ f \circ T_X^{-1}$ is contractible.*

Proof. Recall that by Convention 5.7.1 the canonical bijection is automatically applied to all subsets. By Proposition 5.7.14, there is a special bijection T_Y on Y such that $T_Y(Y)$ is an RV-product. So we may assume that Y is an RV-product. Fix a sequence of quantifier-free formulas $\psi_1, \dots, \psi_{m_2}$ that define the functions $f_i = \text{pr}_i \circ \text{pRV} \circ T_Y \circ f$ for $1 \leq i \leq m_2$. Let $g_i(\bar{x})$ enumerate all the VF-sort polynomials that occur in $\psi_1, \dots, \psi_{m_2}$ in the form $\text{rv}(g_i(\bar{x}))$. By Proposition 5.7.14 and Corollary 5.7.16, there is a special bijection T_X on X such that $T_X(X)$ is an RV-product and the function $\text{rv}(g_i(T_X^{-1}(\bar{x})))$ is constant on every rv-polyball $\mathfrak{p} \subseteq T_X(X)$ for every i . So on such an rv-polyball every $f_i \circ T_X^{-1}$ is constant. \square

Lemma 5.10.3. *Let $X \subseteq \text{VF} \times \text{RV}^{m_1}$ and $Y \subseteq \text{VF} \times \text{RV}^{m_2}$ be definable subsets and $f : X \rightarrow Y$ a definable bijection. Then there exist special bijections $T_X : X \rightarrow X^\sharp$ and $T_Y : Y \rightarrow Y^\sharp$ such that, in the commutative diagram*

$$\begin{array}{ccccc} X & \xrightarrow{T_X} & X^\sharp & \xrightarrow{\text{rv}} & \text{rv}(X^\sharp) \\ f \downarrow & & f^\sharp \downarrow & & \downarrow f^\sharp_\downarrow \\ Y & \xrightarrow{T_Y} & Y^\sharp & \xrightarrow{\text{rv}} & \text{rv}(Y^\sharp) \end{array}$$

f^\sharp_\downarrow is bijective and hence f^\sharp is a lift of it.

Proof. By Proposition 5.5.15, there is a definable partition X_1, \dots, X_n of X such that each $f \upharpoonright X_i$ has the open-to-open property. Therefore, applying Lemma 5.10.2 to each $f \upharpoonright X_i$ or its subsequent image, we may assume that X, Y are RV-products and f is contractible and has the open-to-open property. In particular, for each rv-polyball $\mathfrak{p} \subseteq X$, $f(\mathfrak{p})$ is an open ball contained in an rv-polyball $\mathfrak{p}^* \subseteq Y$. By Lemma 5.10.2 again, there is a special bijection $T_Y : Y \rightarrow Y^\sharp$ such that $(T_Y \circ f)^{-1}$ is contractible. Let $T_Y = \mathbf{c} \circ \eta_n \circ \dots \circ \mathbf{c} \circ \eta_1$, where each η_i is a centripetal transformation and \mathbf{c} is the canonical bijection.

Now, by induction on n , we construct a special bijection $T_X = \mathbf{c} \circ \eta_n^* \circ \dots \circ \mathbf{c} \circ \eta_1^*$ on X such that, for each i , both $L_i \circ f \circ (L_i^*)^{-1}$ and $(T_Y \circ f \circ L_i^*)^{-1}$ are contractible, where $L_i = \mathbf{c} \circ \eta_i \circ \dots \circ \mathbf{c} \circ \eta_1$ and $L_i^* = \mathbf{c} \circ \eta_i^* \circ \dots \circ \mathbf{c} \circ \eta_1^*$. Then T_X, T_Y will be as desired. To that end, suppose that η_i^* has been constructed for each $i \leq k < n$. Let $Z_k = L_k^*(X)$ and $Z_k^\sharp = L_k(Y)$. Let $C \subseteq Z_k^\sharp$ be the locus of η_{k+1} and λ the corresponding focus map. Since $L_k \circ f \circ (L_k^*)^{-1}$ is contractible and has the open-to-open property, each rv-polyball $\mathfrak{p} \subseteq Z_k^\sharp$ is the union of disjoint subsets of the form $(L_k \circ f \circ (L_k^*)^{-1})(\mathfrak{q})$, where $\mathfrak{q} \subseteq Z_k$ is an rv-polyball. For each $\bar{t} = (t_1, \bar{t}_1) \in \text{dom}(\lambda)$, let

$$O_{\bar{t}} = \{ \mathfrak{q} \subseteq Z_k : \mathfrak{q} \text{ is an rv-polyball and } (L_k \circ f \circ (L_k^*)^{-1})(\mathfrak{q}) \subseteq \text{rv}^{-1}(t_1) \times \{\bar{t}_1\} \}.$$

Then, for each $\bar{t} = (t_1, \bar{t}_1) \in \text{dom}(\lambda)$, there is an open subball $\mathfrak{o}_{\bar{t}} \subseteq \text{rv}^{-1}(t_1) \times \{\bar{t}_1\} \subseteq C$ and a $\mathfrak{q}_{\bar{t}} \in O_{\bar{t}}$ such that $(\lambda(\bar{t}), \bar{t}) \in \mathfrak{o}_{\bar{t}}$ and $(L_k \circ f \circ (L_k^*)^{-1})(\mathfrak{q}_{\bar{t}}) = \mathfrak{o}_{\bar{t}}$. Let $C^* = \bigcup \{ \mathfrak{q}_{\bar{t}} : \bar{t} \in \text{dom}(\lambda) \} \subseteq Z_k$ and, for each

$\bar{t} \in \text{dom}(\lambda)$,

$$a_{\bar{t}} = (L_k \circ f \circ (L_k^*)^{-1})^{-1}(\lambda(\bar{t}), \bar{t}) \in \mathfrak{q}_{\bar{t}}.$$

Let $\lambda^* : \text{pr}_{>1} C^* \longrightarrow \text{VF}$ be the corresponding focus map given by $\lambda^*(\text{pr}_{>1} \mathfrak{q}_{\bar{t}}) = a_{\bar{t}}$. Note that both C^* and λ^* are definable. Let η_{k+1}^* be the centripetal transformation determined by C^* and λ^* . For each $\bar{t} \in \text{dom}(\lambda)$, the restriction of $L_{k+1} \circ f \circ (L_{k+1}^*)^{-1}$ to $\mathfrak{c}(\mathfrak{q}_{\bar{t}} - a_{\bar{t}})$ is a bijection between the RV-products $\mathfrak{c}(\mathfrak{q}_{\bar{t}} - a_{\bar{t}})$ and $\mathfrak{c}(\mathfrak{o}_{\bar{t}} - \lambda(\bar{t}))$ that is contractible in both ways. So, by the construction of L_k^* , $(T_Y \circ f \circ L_{k+1}^*)^{-1}$ is contractible. Also, for each $\bar{t} \in \text{dom}(\lambda)$ and any $\mathfrak{q} \in O_{\bar{t}}$ with $\mathfrak{q} \neq \mathfrak{q}_{\bar{t}}$, $(L_{k+1} \circ f \circ (L_{k+1}^*)^{-1})(\mathfrak{c}(\mathfrak{q}))$ is an open polyball contained in an rv-polyball. So $L_{k+1} \circ f \circ (L_{k+1}^*)^{-1}$ is contractible. \square

Corollary 5.10.4. *Let $(X_1, g_1), (X_2, g_2) \in \text{Ob RV}[1, \cdot]$ be such that $\mathbb{L}(X_1, g_1)$ is definably bijective to $\mathbb{L}(X_2, g_2)$. Then there are special bijections T_1, T_2 on $\mathbb{L}(X_1, g_1), \mathbb{L}(X_2, g_2)$ such that (X_1^*, pr_1) and (X_2^*, pr_1) are isomorphic, where*

$$(X_1^*, \text{pr}_1) = ((\text{pRV} \circ T_1)(\mathbb{L}(X_1, g_1)), \text{pr}_1),$$

$$(X_2^*, \text{pr}_1) = ((\text{pRV} \circ T_2)(\mathbb{L}(X_2, g_2)), \text{pr}_1).$$

Proof. By Lemma 5.10.3, there are special bijections T_1, T_2 on $\mathbb{L}(X_1, g_1), \mathbb{L}(X_2, g_2)$ such that there are definable bijections

$$F : (\text{rv} \circ T_1)(\mathbb{L}(X_1, g_1)) \longrightarrow (\text{rv} \circ T_2)(\mathbb{L}(X_2, g_2))$$

$$F^\uparrow : T_1(\mathbb{L}(X_1, g_1)) \longrightarrow T_2(\mathbb{L}(X_2, g_2))$$

and F^\uparrow is a lift of F . Since F is a bijection between the companions of (X_1^*, pr_1) and (X_2^*, pr_1) , by Remark 5.6.19, the natural projection of F is an isomorphism between the two. \square

Definition 5.10.5. Let $X \subseteq \text{VF}^n \times \text{RV}^{m_1}$ and $Y \subseteq \text{VF}^n \times \text{RV}^{m_2}$ and $f : X \longrightarrow Y$ a bijection. Let $E \subseteq \mathbb{N}$ be the set of the indices of the VF-coordinates. We say that f is *relatively unary* if there is an $i \in E$ such that $(\text{pr}_{E_i} \circ f)(\bar{x}) = \text{pr}_{E_i}(\bar{x})$, where $E_i = E \setminus \{i\}$. In this case we say that f is *unary relative to the coordinate i* . If, in addition, $f \upharpoonright \text{fib}(X, \bar{a})$ is a special bijection on $\text{fib}(X, \bar{a})$ for every $\bar{a} \in \text{pr}_{E_i} X$ then we say that f is *special relative to the coordinate i* .

Obviously the inverse of a relatively unary bijection is a relatively unary bijection.

Let $X \subseteq \text{VF}^n \times \text{RV}^m$, $C \subseteq \text{RVH}(X)$ an RV-product, λ a focus map on $\text{pr}_{>1} C$, and η the centripetal transformation with respect to λ . Let $X_C = X \cap C$. Clearly $\eta \upharpoonright X_C$ is a special bijection relative to the

coordinate 1. It follows that, for every special bijection T on X , $T \upharpoonright X$ is a composition of relatively special bijections. Suppose that X is definable. Let $i \leq n$ and $E_i = \{1, \dots, n\} \setminus \{i\}$. By Proposition 5.7.14, for every $\bar{a} \in \text{pr}_{E_i} X$, there is an \bar{a} -definable special bijection $I_{\bar{a}}$ such that $I_{\bar{a}}(\text{fib}(X, \bar{a}))$ is an RV-product and hence, by compactness, there is a special bijection I_i relative to the coordinate i such that $I_i(\text{fib}(X, \bar{a}))$ is an RV-product for every $\bar{a} \in \text{pr}_{E_i} X$. Let

$$X_i = \{(\bar{a}_i, (\text{pRV} \circ I_i)(a_i, \bar{a}_i, \bar{t}), \bar{t}) : (a_i, \bar{a}_i, \bar{t}) \in X\} \subseteq \text{VF}^{n-1} \times \text{RV}^{m+1}.$$

We write $\widehat{I}_i : X \rightarrow X_i$ for the function induced by I_i . Let $j \leq n$ with $j \neq i$. Repeating the above procedure for X_i with respect to j , we obtain a subset $X_j \subseteq \text{VF}^{n-2} \times \text{RV}^{m+2}$ and a function $\widehat{I}_j : X_i \rightarrow X_j$, which depend on the relatively special bijection I_j . Continuing this procedure, we see that, for any permutation σ of $\{1, \dots, n\}$, there is a sequence of relatively special bijections $I_{\sigma(1)}, \dots, I_{\sigma(n)}$ and a corresponding function $\widehat{I}_\sigma : X \rightarrow \text{RV}^{m+n}$ such that there are an $E \subseteq \mathbb{N}$ with $|E| = n$ and a special bijection $I_\sigma = I_{\sigma(n)} \circ \dots \circ I_{\sigma(1)} : X \rightarrow \mathbb{L}(\widehat{I}_\sigma(X), \text{pr}_E)$. As before, after a permutation of indices, we may always assume that $E = \{1, \dots, n\}$.

Definition 5.10.6. The function \widehat{I}_σ is called a *standard contraction* of X .

Let $X \subseteq \text{VF}^n \times \text{RV}^m$ and \widehat{I}_{id} a standard contraction of X such that $I_{\text{id}}(X)$ is of the form $\text{rv}^{-1}(t_i) \times \{(\bar{0}, t_i, \overline{\infty}, \bar{s})\}$, where $\bar{0}$ is a tuple of 0 of length $n-1$ and $\overline{\infty}$ is a tuple of ∞ of length $n-1$. Let $I_{\text{id}} = I_n \circ \dots \circ I_1$ and $I_{\leq i} = I_i \circ \dots \circ I_1$. Clearly $I_{\leq i}(X)$ is of the form $\text{rv}^{-1}(t_i) \times \{(\bar{0}, \bar{a}, t_i, \overline{\infty}, \bar{s})\}$, where $\bar{0}$ is a tuple of 0 of length $i-1$, $\bar{a} \in \text{VF}$ is a tuple of length $n-i$, and $\overline{\infty}$ is a tuple of ∞ of length $i-1$. So for any distinct $(a, \bar{0}, t_i, \overline{\infty}, \bar{s}), (b, \bar{0}, t_i, \overline{\infty}, \bar{s}) \in I_{\text{id}}(X)$ we have that

$$(\text{pVF}_i \circ I_{\text{id}}^{-1})(a, \bar{0}, t_i, \overline{\infty}, \bar{s}) \neq (\text{pVF}_i \circ I_{\text{id}}^{-1})(b, \bar{0}, t_i, \overline{\infty}, \bar{s}).$$

This simple observation is used to prove the following:

Lemma 5.10.7. *Let $X \subseteq \text{VF}^n \times \text{RV}^{m_1}$, $Y \subseteq \text{VF}^n \times \text{RV}^{m_2}$ be definable subsets and $f : X \rightarrow Y$ a definable bijection. Then there is a definable partition X_1, \dots, X_k of X such that each $f \upharpoonright X_i$ is a composition of definable relatively unary bijections.*

Proof. We do induction on n . Since the base case $n = 1$ holds vacuously, we proceed to the inductive step directly. By Lemma 5.7.14, for each $\bar{a} = (a_1, \dots, a_{n-1}) \in \text{pVF}_{<n} X$, there is an \bar{a} -definable standard contraction $\widehat{I}_{\text{id}, \bar{a}}$ on $f(\text{fib}(X, \bar{a}))$ such that $(I_{\text{id}, \bar{a}} \circ f)(\text{fib}(X, \bar{a})) = Z_{\bar{a}}$ is an RV-product. By Lemma 5.7.4, in each tuple $(\bar{t}, \bar{s}) = (t_1, \dots, t_n, \bar{s}) \in \text{pRV } Z_{\bar{a}}$, there is at most one $i \leq n$ such that $t_i \neq \infty$, that is, each rv-polyball contained in $Z_{\bar{a}}$ is of the form $\text{rv}^{-1}(t_i) \times \{(\bar{0}, t_i, \overline{\infty}, \bar{s})\}$ for some $i \leq n$. So there is an

\bar{a} -definable partition A_1, \dots, A_n of $\text{fib}(X, \bar{a})$ such that if $(\bar{a}, a_n, \bar{r}) \in A_i$ then $(I_{\text{id}, \bar{a}} \circ f)(\bar{a}, a_n, \bar{r})$ is of the form $(b_i, \bar{0}, t_i, \overline{\infty}, \bar{s})$. By the observation above, if $(b_i, \bar{0}, t_i, \overline{\infty}, \bar{s}), (b'_i, \bar{0}, t_i, \overline{\infty}, \bar{s})$ are distinct elements in $Z_{\bar{a}}$, then

$$(\text{pVF}_i \circ I_{\text{id}, \bar{a}}^{-1})(b_i, \bar{0}, t_i, \overline{\infty}, \bar{s}) \neq (\text{pVF}_i \circ I_{\text{id}, \bar{a}}^{-1})(b'_i, \bar{0}, t_i, \overline{\infty}, \bar{s}).$$

Let $g_{\bar{a}, i}$ be the function on A_i given by

$$(\bar{a}, a_n, \bar{r}) \mapsto (\bar{a}, d_i, \bar{r}, t_i, \overline{\infty}, \bar{s}),$$

where $(\text{pRV} \circ I_{\text{id}, \bar{a}} \circ f)(\bar{a}, a_n, \bar{r}) = (t_i, \overline{\infty}, \bar{s})$ and $(\text{pVF}_i \circ f)(\bar{a}, a_n, \bar{r}) = d_i$. Therefore, after reindexing the VF-coordinates in each A_i separately, each $g_{\bar{a}, i}$ is an \bar{a} -definable unary bijection on A_i relative to the coordinate i such that $\text{pVF}_i \circ f = \text{pVF}_i \circ g_{\bar{a}, i}$. By compactness, there are a definable partition B_1, \dots, B_n of X and definable unary bijections g_i on B_i relative to the coordinate i such that $\text{pVF}_i \circ f = \text{pVF}_i \circ g_i$.

For each $i \leq n$ let h_i be the function on $g_i(B_i)$ such that $f \upharpoonright B_i = h_i \circ g_i$. For each $a \in (\text{pVF}_i \circ g_i)(B_i)$, since $h_i(\text{fib}(g_i(B_i), a)) = \text{fib}(f(B_i), a)$, by the inductive hypothesis, there is an a -definable partition D_1, \dots, D_l of $\text{fib}(g_i(B_i), a)$ such that each $h_i \upharpoonright D_j$ is a composition of a -definable relatively unary bijections. So the inductive step holds by compactness. \square

Lemma 5.10.8. *Let $X \subseteq \text{VF}^2$ be a definable 2-cell. Let 12, 21 denote the permutations of $\{1, 2\}$. Then there are standard contractions \widehat{I}_{12} and \widehat{J}_{21} of X such that $(\widehat{I}_{12}(X), \text{pr}_{\leq 2})$ and $(\widehat{J}_{21}(X), \text{pr}_{\leq 2})$ are isomorphic.*

Proof. Let ϵ be the positioning function of X and $t \in \text{RV}$ the paradigm of X . If $t = \infty$ then X is the function $\epsilon : \text{pr}_1 X \rightarrow \text{pr}_2 X$, which is either a constant function or a bijection. In the former case, since X is essentially just an open ball, the lemma simply follows from Lemma 5.7.8. In the latter case, there are special bijections I_2, J_1 on X relative to the coordinates 2, 1 such that $I_2(X) = (\text{pr}_1 X) \times \{(0, \infty)\}$ and $J_1(X) = \{0\} \times (\text{pr}_2 X) \times \{\infty\}$. So the lemma simply follows from Lemma 5.10.3. For the rest of the proof we assume that $t \neq \infty$.

If ϵ is not trapezoidal in X then X is an open polyball, that is, $X = (\text{pr}_1 X) \times (\text{pr}_2 X)$, where $\text{pr}_1 X, \text{pr}_2 X$ are definable open balls. By Lemma 5.7.8, there are special bijections T_1, T_2 on $\text{pr}_1 X, \text{pr}_2 X$ such that $T_1(\text{pr}_1 X), T_2(\text{pr}_2 X)$ are RV-products. Trivially, the standard contractions determined by (T_1, T_2) and (T_2, T_1) are the same.

Suppose that ϵ is trapezoidal in X . Let r be the other paradigm of X . Recall that $\epsilon : \text{pr}_1 X \rightarrow \text{pr}_2 X$ is again a bijection. Let I_2 be the special bijection on X relative to the coordinate 2 given by $(a, b) \mapsto (a, b - \epsilon(a))$ and J_1 the special bijection on X relative to the coordinate 1 given by $(a, b) \mapsto (a - \epsilon^{-1}(b), b)$,

where $(a, b) \in X$. Clearly $I_2(X) = (\text{pr}_1 X) \times \text{rv}^{-1}(t) \times \{t\}$ and $J_1(X) = \text{rv}^{-1}(r) \times (\text{pr}_2 X) \times \{r\}$. So, again, the lemma follows from Lemma 5.10.3. \square

Corollary 5.10.9. *Let $X \subseteq \text{VF}^2 \times \text{RV}^m$ be a definable subset. Then there is a definable bijection $f : X \rightarrow \text{VF}^2 \times \text{RV}^l$ such that f is unary relative to both coordinates and there are standard contractions \widehat{I}_{12} and \widehat{J}_{21} of $f(X)$ such that, for every $\bar{t} \in \text{pRV } f(X)$, $(\widehat{I}_{12}(\text{fib}(f(X), \bar{t})), \text{pr}_{\leq 2})$ and $(\widehat{J}_{21}(\text{fib}(f(X), \bar{t})), \text{pr}_{\leq 2})$ are isomorphic.*

Proof. By Lemma 5.8.8, there is a definable function $f : X \rightarrow \text{VF}^2 \times \text{RV}^l$ such that, for each $(\bar{a}, \bar{t}) \in X$, $f(\bar{a}, \bar{t}) = (\bar{a}, \bar{t}, \bar{s})$ for some $\bar{s} \in \text{RV}^{l-m}$, and $\text{fib}(f(X), (\bar{t}, \bar{s}))$ is a 2-cell for every $(\bar{t}, \bar{s}) \in \text{pRV } f(X)$. Now the corollary follows from Lemma 5.10.8 and compactness. \square

5.11 The kernel of \mathbb{L}

We identify all the semigroup homomorphisms induced by \mathbb{L} with \mathbb{L} . We shall show that the kernel of \mathbb{L} , that is, the semigroup (semiring) congruence relation induced by \mathbb{L} on the domain of \mathbb{L} , is in effect definable and hence the inverse of \mathbb{L} modulo the congruence relation is definable.

5.11.1 Blowups in RV and the congruence relation I_{sp}

Definition 5.11.1. Let $(Y, f) \in \text{Ob RV}[k, \cdot]$ be such that, for all $\bar{t} \in Y$, $f_{|k}(\bar{t}) \in \text{acl}(f_{|1}(\bar{t}), \dots, f_{|k-1}(\bar{t}))$ and $f_{|k}(\bar{t}) \neq \infty$. Let $(Y, f)^\sharp = (Y^\sharp, f^\sharp) \in \text{Ob RV}[k, \cdot]$ be such that $Y^\sharp = Y \times \text{RV}^{>1}$ and, for any $(\bar{t}, s) \in Y^\sharp$, $f_{|i}^\sharp(\bar{t}, s) = f_{|i}(\bar{t})$ if $1 \leq i < k$ and $f_{|k}^\sharp(\bar{t}, s) = s f_{|k}(\bar{t})$. The object $(Y, f)^\sharp$ is an *elementary blowup* of (Y, f) . An elementary blowup of any subobject of (Y, f) is an *elementary sub-blowup* of (Y, f) .

Let $(X, g) \in \text{Ob RV}[k, \cdot]$ and $(C, g \upharpoonright C) \in \text{Ob RV}[k, \cdot]$ a subobject of (X, g) . Let $F : (Y, f) \rightarrow (C, g \upharpoonright C)$ be an isomorphism. Then

$$(Y, f)^\sharp \uplus (X \setminus C, g \upharpoonright (X \setminus C)) = (Y^\sharp \uplus (X \setminus C), f^\sharp \uplus (g \upharpoonright (X \setminus C)))$$

is the *blowup of (X, g) via F* , written as $(X, g)^\sharp_F$, where the subscript F may be dropped when it is not needed. The subset C is called the *blowup locus* of $(X, g)^\sharp_F$. Let $(Z, h) \in \text{Ob RV}[k, \cdot]$ be isomorphic to a subobject of (X, g) . Then the blowup of (Z, h) induced by F , that is, the disjoint union of an elementary sub-blowup of (Y, f) and a subobject of (Z, h) , is a *sub-blowup of (X, g) via F* .

An *iterated blowup* is a composition of finitely many blowups. The *length* of an iterated blowup is the length of the composition, that is, the number of the blowups involved.

Note that, for any $(Y, f) \in \text{Ob RV}[k, \cdot]$ and a coordinate of $f(Y)$, if there is an elementary blowup of (Y, f) with respect to that coordinate then it is unique. We should have included the index of the “blown up” coordinate as a part of the data for an elementary blowup. Since, in context, either this is clear or it does not need to be spelled out, we shall suppress mentioning it below for notational ease.

Remark 5.11.2. Let (Y^\sharp, f^\sharp) be an elementary blowup of $(Y, f) \in \text{Ob RV}[k, \cdot]$. Since $r_k \in \text{acl}(\bar{r}_k)$ for each $(\bar{r}_k, r_k) \in f(Y)$, by compactness, the fiber $\text{fib}(f(Y), \bar{r}_k)$ is finite for each $\bar{r}_k \in (\text{pr}_{<k} \circ f)(Y)$. By compactness again, we see that actually the set $\{|\text{fib}(f(Y), \bar{r}_k)| : \bar{r}_k \in (\text{pr}_{<k} \circ f)(Y)\}$ is bounded. By definition, for any $(\bar{r}_k, u) \in f^\sharp(Y^\sharp)$ and $(\bar{t}, s) \in (f^\sharp)^{-1}(\bar{r}_k, u)$, $\bar{r}_k = (f_{|1}(\bar{t}), \dots, f_{|k-1}(\bar{t}))$ and $u = sf_{|k}(\bar{t})$, where $f_{|k}(\bar{t}) \in \text{fib}(f(Y), \bar{r}_k)$. So the projection map $\text{pr}_{\leq k} : Y^\sharp \rightarrow Y$ is an $\text{RV}[k, \cdot]$ -morphism. Also, since

$$(f^\sharp)^{-1}(\bar{r}_k, u) = \bigcup_{r_k \in \text{fib}(f(Y), \bar{r}_k)} \{f^{-1}(\bar{r}_k, r_k) \times \{s\} : u = sr_k\},$$

clearly if $(Y, f) \in \text{RV}[k]$ then $(Y^\sharp, f^\sharp) \in \text{RV}[k]$. So any iterated blowup of an object in $\text{Ob RV}[k]$ is an object in $\text{Ob RV}[k]$.

Definition 5.11.1 is stated relative to the underlying substructure S . If an object (X, f) is \bar{a} -definable for some extra parameters \bar{a} , then the iterated blowups of (X, f) should be \bar{a} -definable.

Let $(X, g) \in \text{Ob RV}[k, \cdot]$ and $p : X \rightarrow \text{RV}^m$ a definable function. Let $Z = \{(\bar{t}, p(\bar{t})) : \bar{t} \in X\}$ and h the function on Z given by $(\bar{t}, p(\bar{t})) \mapsto g(\bar{t})$. Clearly (Z, h) is isomorphic to (X, g) . For each $\bar{s} \in \text{RV}^m$ let $(X_{\bar{s}}^\sharp, g_{\bar{s}}^\sharp)$ be an \bar{s} -definable blowup of $(p^{-1}(\bar{s}), g \upharpoonright p^{-1}(\bar{s}))$. Let

$$(Z^\sharp, h^\sharp) = \bigcup_{\bar{s} \in \text{RV}^m} (X_{\bar{s}}^\sharp \times \{\bar{s}\}, g_{\bar{s}}^\sharp \times \{\bar{s}\}).$$

By compactness, $(Z^\sharp, h^\sharp) \in \text{Ob RV}[k, \cdot]$. The projection of Z^\sharp to the last m coordinates is identified with the function p .

Definition 5.11.3. The object (Z^\sharp, h^\sharp) is a *parameterized blowup of (X, g) with respect to p* , or a *p -blowup* for short. An *iterated parameterized blowup* (X^\sharp, g^\sharp) of (X, g) is a composition of finitely many parameterized blowups with respect to a sequence of functions of the form $p_1, p_1 \times p_2, \dots, p_1 \times \dots \times p_n$. To specify the functions involved, we also say that it is a (p_1, \dots, p_n) -*blowup*. An iterated parameterized blowup (X^\sharp, g^\sharp) of (X^\sharp, g^\sharp) is a *continuation* of (X^\sharp, g^\sharp) if it starts with a function of the form $p_1 \times \dots \times p_n \times p_{n+1}$. The *length* of an iterated parameterized blowup is the length of the composition.

We could have allowed each $(X_{\bar{s}}^\sharp, g_{\bar{s}}^\sharp)$ to be an \bar{s} -definable iterated blowup, but it is easily seen to be equivalent to the above formulation. Note that if (X^\sharp, g^\sharp) is a (p_1, \dots, p_n) -blowup of $(X, g) \in \text{Ob RV}[k, \cdot]$

and $\text{ran}(\coprod \bar{p}) \subseteq \text{RV}^m$ then $(X^\sharp, g^\sharp) \in \text{Ob RV}[k+m, \cdot]$.

The results below will be stated only for the more general categories $\text{RV}[k, \cdot]$, $\text{RV}[* , \cdot]$, etc. But, by Remark 5.11.2, they are easily seen to hold when restricted to $\text{RV}[k]$, $\text{RV}[*]$, etc. as well.

Lemma 5.11.4. *Let $(Y_1, f_1), (Y_2, f_2) \in \text{Ob RV}[k, \cdot]$ and $(Y_1, f_1)^\sharp, (Y_2, f_2)^\sharp$ two elementary blowups. If $(Y_1, f_1), (Y_2, f_2)$ are isomorphic then $(Y_1, f_1)^\sharp, (Y_2, f_2)^\sharp$ are isomorphic.*

Proof. Let $F : (Y_1, f_1) \longrightarrow (Y_2, f_2)$ be an isomorphism. Let $F^\sharp : Y_1^\sharp \longrightarrow Y_2^\sharp$ be the bijection given by $(\bar{t}, s) \longmapsto (F(\bar{t}), s)$. We claim that F^\sharp is an isomorphism.

We first check the condition of finite-to-finite correspondence. By compactness, for any $\bar{r}_k = (r_1, \dots, r_{k-1}) \in (\text{pr}_{<k} \circ f_1)(Y_1)$, the fiber $\text{fib}(f_1(Y_1), \bar{r}_k)$ is finite and does not contain ∞ . For any $(\bar{r}_k, u) \in f_1^\sharp(Y_1^\sharp)$ and any $(\bar{t}, s) \in Y_1^\sharp = Y_1 \times \text{RV}^{>1}$ with $f_1^\sharp(\bar{t}, s) = (\bar{r}_k, u)$, there is an $r_k \in \text{fib}(f_1(Y_1), \bar{r}_k)$ such that $f_{1|k}(\bar{t}) = r_k$ and $u = sr_k$. Let

$$A = \{r_k \in \text{fib}(f_1(Y_1), \bar{r}_k) : \text{there is an } s \in \text{RV}^{>1} \text{ such that } sr_k = u\}.$$

We have that

$$\begin{aligned} & (f_2^\sharp \circ F^\sharp \circ (f_1^\sharp)^{-1})(\bar{r}_k, u) \\ &= (f_2^\sharp \circ F^\sharp) \left(\bigcup_{r_k \in A} \left(f_1^{-1}(\bar{r}_k, r_k) \times \left\{ \frac{u}{r_k} \right\} \right) \right) \\ &= f_2^\sharp \left(\bigcup_{r_k \in A} \left((F \circ f_1^{-1})(\bar{r}_k, r_k) \times \left\{ \frac{u}{r_k} \right\} \right) \right) \\ &= \bigcup_{r_k \in A} \left\{ f_2^\sharp \left(\bar{t}, \frac{u}{r_k} \right) : \bar{t} \in (F \circ f_1^{-1})(\bar{r}_k, r_k) \right\} \\ &= \bigcup_{r_k \in A} \left\{ \left((\text{pr}_{<k} \circ f_2)(\bar{t}), \frac{uf_{2|k}(\bar{t})}{r_k} \right) : \bar{t} \in (F \circ f_1^{-1})(\bar{r}_k, r_k) \right\}. \end{aligned}$$

Since the subset $(f_2 \circ F \circ f_1^{-1})(\bar{r}_k, r_k)$ is finite for each $r_k \in A$, it follows that the subset $(f_2^\sharp \circ F^\sharp \circ (f_1^\sharp)^{-1})(\bar{r}_k, u)$ is also finite. Similarly for the other direction $f_1^\sharp \circ (F^\sharp)^{-1} \circ (f_2^\sharp)^{-1}$.

Next we check that F^\sharp is volumetric, that is, the condition on weight. For any $(\bar{t}, s) \in Y_1^\sharp$, if $s \neq \infty$ then $\text{wgt } f_1^\sharp(\bar{t}, s) = \text{wgt } f_1(\bar{t})$ and $\text{wgt}(f_2^\sharp \circ (F^\sharp)^{-1})(\bar{t}, s) = \text{wgt}(f_2 \circ F)(\bar{t})$. Since $\text{wgt } f_1(\bar{t}) = (f_2 \circ F)(\bar{t})$, we deduce that $\text{wgt } f_1^\sharp(\bar{t}, s) = \text{wgt}(f_2^\sharp \circ (F^\sharp)^{-1})(\bar{t}, s)$. If $s = \infty$ then $\text{wgt } f_1^\sharp(\bar{t}, s) = \text{wgt } f_1(\bar{t}) - 1$ and $\text{wgt}(f_2^\sharp \circ (F^\sharp)^{-1})(\bar{t}, s) = \text{wgt}(f_2 \circ F)(\bar{t}) - 1$ and hence $\text{wgt } f_1^\sharp(\bar{t}, s) = \text{wgt}(f_2^\sharp \circ (F^\sharp)^{-1})(\bar{t}, s)$. \square

Corollary 5.11.5. *Let $(X_1, g_1), (X_2, g_2) \in \text{Ob RV}[k, \cdot]$ be isomorphic. Let $(X_1, g_1)^\sharp, (X_2, g_2)^\sharp$ be two blowups of $(X_1, g_1), (X_2, g_2)$ with isomorphic blowup loci. Then $(X_1, g_1)^\sharp, (X_2, g_2)^\sharp$ are isomorphic.*

Lemma 5.11.6. *Let $(X_1, g_1), (X_2, g_2) \in \text{Ob RV}[k, \cdot]$ be isomorphic. Let $(Z_1, h_1), (Z_2, h_2)$ be two iterated blowups of $(X_1, g_1), (X_2, g_2)$ of length l_1, l_2 , respectively. Then there are isomorphic iterated blowups $(Z_1^*, h_1^*), (Z_2^*, h_2^*)$ of $(Z_1, h_1), (Z_2, h_2)$ of length l_2, l_1 .*

Proof. Fix an isomorphism $I : (X_1, g_1) \longrightarrow (X_2, g_2)$. We do induction on the sum $l = l_1 + l_2$. For the base case $l = 1$, without loss of generality, we assume that $l_2 = 0$. Let C be the blowup locus of (Z_1, h_1) . Clearly (X_2, g_2) may be blown up by using the same elementary blowup as (Z_1, h_1) , where the blowup locus is changed to $I(C)$. So the base case holds.

We proceed to the inductive step. Let $(X_1, g_1)^\sharp, (X_2, g_2)^\sharp$ be the first blowups in $(Z_1, h_1), (Z_2, h_2)$ and C_1, C_2 their blowup loci, respectively. Let $(Y_1, f_1)^\sharp, (Y_2, f_2)^\sharp$ be the corresponding elementary blowups. If, say, $l_2 = 0$, then by the argument in the base case (X_2, g_2) may be blown up to an object that is isomorphic to $(X_1, g_1)^\sharp$ and hence the inductive hypothesis may be applied. So let us assume that $l_1, l_2 > 0$. Let $A_1 = C_1 \cap I^{-1}(C_2)$ and $A_2 = I(C_1) \cap C_2$. Since $(A_1, g_1 \upharpoonright A_1)$ and $(A_2, g_2 \upharpoonright A_2)$ are isomorphic, by Lemma 5.11.4, the elementary sub-blowups of $(Y_1, f_1)^\sharp, (Y_2, f_2)^\sharp$ that correspond to $(A_1, g_1 \upharpoonright A_1)$ and $(A_2, g_2 \upharpoonright A_2)$ are isomorphic. Then, it is not hard to see that the blowup $(X_1, g_1)^{\sharp\sharp}$ of $(X_1, g_1)^\sharp$ using the locus $I^{-1}(C_2) \setminus C_1$ and its corresponding elementary sub-blowup of $(Y_2, f_2)^\sharp$ and the blowup $(X_2, g_2)^{\sharp\sharp}$ of $(X_2, g_2)^\sharp$ using the locus $I(C_1) \setminus C_2$ and its corresponding elementary sub-blowup of $(Y_1, f_1)^\sharp$ are isomorphic.

Applying the inductive hypothesis to the iterated blowups $(X_1, g_1)^{\sharp\sharp}, (Z_1, h_1)$ of $(X_1, g_1)^\sharp$, we obtain an iterated blowup (X_1^*, g_1^*) of $(X_1, g_1)^{\sharp\sharp}$ of length $l_1 - 1$ and a blowup $(Z_1, h_1)^\sharp$ of (Z_1, h_1) such that (X_1^*, g_1^*) and $(Z_1, h_1)^\sharp$ are isomorphic. Similarly, we obtain an iterated blowup (X_2^*, g_2^*) of $(X_2, g_2)^{\sharp\sharp}$ of length $l_2 - 1$ and a blowup $(Z_2, h_2)^\sharp$ of (Z_2, h_2) such that (X_2^*, g_2^*) and $(Z_2, h_2)^\sharp$ are isomorphic. Now, applying the inductive hypothesis to the iterated blowups $(X_1^*, g_1^*), (X_2^*, g_2^*)$ of $(X_1, g_1)^{\sharp\sharp}, (X_2, g_2)^{\sharp\sharp}$, we obtain an iterated blowup (X_1^{**}, g_1^{**}) of (X_1^*, g_1^*) of length $l_2 - 1$ and an iterated blowup (X_2^{**}, g_2^{**}) of (X_2^*, g_2^*) of length $l_1 - 1$ such that (X_1^{**}, g_1^{**}) and (X_2^{**}, g_2^{**}) are isomorphic. Finally, applying the inductive hypothesis to the iterated blowups $(X_1^{**}, g_1^{**}), (Z_1, h_1)^\sharp$ of $(X_1^*, g_1^*), (Z_1, h_1)^\sharp$ and the iterated blowups $(X_2^{**}, g_2^{**}), (Z_2, h_2)^\sharp$ of $(X_2^*, g_2^*), (Z_2, h_2)^\sharp$, we obtain an iterated blowup (Z_1^*, h_1^*) of $(Z_1, h_1)^\sharp$ of length $l_2 - 1$ and an iterated blowup (Z_2^*, h_2^*) of $(Z_2, h_2)^\sharp$ of length $l_1 - 1$ such that $(X_1^{**}, g_1^{**}), (Z_1^*, h_1^*)$ are isomorphic and $(X_2^{**}, g_2^{**}), (Z_2^*, h_2^*)$ are isomorphic.

This process is illustrated as follows:

$$\begin{array}{ccccccc}
(X_1, g_1) & \xrightarrow{1} & (X_1, g_1)^\# & \xrightarrow{l_1-1} & (Z_1, h_1) & \xrightarrow{1} & (Z_1, h_1)^\# & \xrightarrow{l_2-1} & (Z_1^*, h_1^*) \\
\parallel & & \searrow 1 & & \parallel & & \parallel & & \parallel \\
& & (X_1, g_1)^\#\# & \xrightarrow{l_1-1} & (X_1^*, g_1^*) & \xrightarrow{l_2-1} & (X_1^{**}, g_1^{**}) & & \\
& & \parallel & & \parallel & & \parallel & & \\
& & (X_2, g_2)^\#\# & \xrightarrow{l_2-1} & (X_2^*, g_2^*) & \xrightarrow{l_1-1} & (X_2^{**}, g_2^{**}) & & \\
& & \parallel & & \parallel & & \parallel & & \\
(X_2, g_2) & \xrightarrow{1} & (X_2, g_2)^\# & \xrightarrow{l_2-1} & (Z_2, h_2) & \xrightarrow{1} & (Z_2, h_2)^\# & \xrightarrow{l_1-1} & (Z_2^*, h_2^*)
\end{array}$$

So (Z_1^*, h_1^*) and (Z_2^*, h_2^*) are as desired. \square

Definition 5.11.7. Two objects $(X_1, g_1), (X_2, g_2) \in \text{Ob RV}[k, \cdot]$ are *parametrically isomorphic* if there are definable functions $p_1 : X_1 \rightarrow \text{RV}^m, p_2 : X_2 \rightarrow \text{RV}^m$ such that for every $\bar{s} \in \text{RV}^m$ there is an \bar{s} -definable isomorphism between the subobjects $(p_1^{-1}(\bar{s}), g_1 \upharpoonright p_1^{-1}(\bar{s})), (p_2^{-1}(\bar{s}), g_2 \upharpoonright p_2^{-1}(\bar{s}))$.

As usual, by compactness, if two objects are parametrically isomorphic then there is a uniform formula that witnesses the fiberwise isomorphisms.

Lemma 5.11.8. *Let $(X_1, g_1), (X_2, g_2), (X_3, g_3) \in \text{Ob RV}[k, \cdot]$ be such that the first two are parametrically isomorphic with respect to the functions p_1, p_2 and the last two are parametrically isomorphic with respect to the functions q_2, q_3 . Then $(X_1, g_1), (X_3, g_3)$ are parametrically isomorphic.*

Proof. Without loss of generality, we may suppose that $\text{ran } p_1, \text{ran } p_2, \text{ran } q_2, \text{ran } q_3$ are all subsets of RV^m . For every $\bar{r}, \bar{s} \in \text{RV}^m$ let $Y_{\bar{r}, \bar{s}}$ be the intersection of $p_2^{-1}(\bar{r})$ and $q_2^{-1}(\bar{s})$. Since the (\bar{r}, \bar{s}) -definable isomorphic images of $(Y_{\bar{r}, \bar{s}}, g_2 \upharpoonright Y_{\bar{r}, \bar{s}})$ in $(X_1, g_1), (X_3, g_3)$ are manifestly isomorphic, the lemma follows from compactness. \square

So “parametrically isomorphic” is an equivalence relation on $\text{Ob RV}[k, \cdot]$. For any $(X, g) \in \text{RV}[k, \cdot]$ the parametrical isomorphism class of (X, g) is denoted as $[[X, g]]$. Notation 5.6.25 is adapted for parametrical isomorphism classes in the obvious way.

Lemma 5.11.9. *Let $(X_1, g_1), (X_2, g_2) \in \text{Ob RV}[k, \cdot]$ be parametrically isomorphic with respect to o_1, o_2 . Let (Z_1, h_1) be a (p_1, \dots, p_{l_1}) -blowup of (X_1, g_1) and (Z_2, h_2) a (q_1, \dots, q_{l_2}) -blowup of (X_2, g_2) . Then there are parametrically isomorphic continuations $(Z_1^*, h_1^*), (Z_2^*, h_2^*)$ of $(Z_1, h_1), (Z_2, h_2)$ of length l_2, l_1 .*

Proof. Applying Lemma 5.11.6, it is not hard to see that, in the case $l_1 = l_2 = 1$, such continuations exist and they are parametrically isomorphic with respect to $o_1 \times p_1 \times q_1^*$ and $o_2 \times p_1^* \times q_1$, where q_1^*, p_1^* are functions on X_1, X_2 induced by o_1, o_2 . Then the proof is completely analogous to that of Lemma 5.11.6.

The additional bookkeeping on the parameterizing functions is carried along inductively in a straightforward way. \square

Definition 5.11.10. Let $\mathbf{I}_{\text{sp}}[k, \cdot]$ be the subclass of $\text{Ob RV}[k, \cdot] \times \text{Ob RV}[k, \cdot]$ of those pairs $((X_1, g_1), (X_2, g_2))$ such that there are a \bar{p} -blowup (X_1^\sharp, g_1^\sharp) of (X_1, g_1) and a \bar{q} -blowup (X_2^\sharp, g_2^\sharp) of (X_2, g_2) such that they are parametrically isomorphic with respect to $\prod \bar{p}$ and $\prod \bar{q}$. Let

$$\mathbf{I}_{\text{sp}}[* , \cdot] = \prod_{0 \leq k} \mathbf{I}_{\text{sp}}[k, \cdot],$$

$$\mathbf{I}_{\text{sp}}[k] = \mathbf{I}_{\text{sp}}[k, \cdot] \cap (\text{Ob RV}[k] \times \text{Ob RV}[k]),$$

$$\mathbf{I}_{\text{sp}}[*] = \mathbf{I}_{\text{sp}}[* , \cdot] \cap \prod_{0 \leq k} (\text{Ob RV}[k] \times \text{Ob RV}[k]).$$

We will just write \mathbf{I}_{sp} for all these classes if there is no danger of confusion. When the underlying substructure S is expanded with some extra parameters \bar{a} we shall write $\mathbf{I}_{\text{sp}}\langle \bar{a} \rangle$ for the accordingly expanded classes.

Corollary 5.11.11. *The class $\mathbf{I}_{\text{sp}}[k, \cdot]$ is transitive.*

Proof. This is immediate by Lemma 5.11.8 and Lemma 5.11.9. \square

Let $(X_1, g_1), (Y_1, f_1), (X_2, g_2), (Y_2, f_2) \in \text{Ob RV}[k, \cdot]$ such that the first two and the last two are parametrically isomorphic. Then certainly we have that $((X_1, g_1), (Y_1, f_1)), ((X_2, g_2), (Y_2, f_2)) \in \mathbf{I}_{\text{sp}}$. If $((X_1, g_1), (X_2, g_2)) \in \mathbf{I}_{\text{sp}}$ then, by Corollary 5.11.11, we have that $((Y_1, f_1), (Y_2, f_2)) \in \mathbf{I}_{\text{sp}}$. So we may treat \mathbf{I}_{sp} as a binary relation on parametrical isomorphism classes and hence on isomorphism classes. In fact,

Lemma 5.11.12. *$\mathbf{I}_{\text{sp}}[k, \cdot]$ is a semigroup congruence relation and $\mathbf{I}_{\text{sp}}[* , \cdot]$ is a semiring congruence relation.*

Proof. Clearly $\mathbf{I}_{\text{sp}}[k, \cdot]$ is reflexive and symmetric. It is also transitive by Corollary 5.11.11 and hence is an equivalence relation. Suppose that $([(X_1, g_1)], [(X_2, g_2)]) \in \mathbf{I}_{\text{sp}}[k, \cdot]$. For any $[(Z, h)] \in \mathbf{K}_+ \text{RV}[k, \cdot]$, it is easily checked that

$$([(X_1, g_1) \uplus (Z, h)], [(X_2, g_2) \uplus (Z, h)]) \in \mathbf{I}_{\text{sp}}[k, \cdot],$$

$$([(X_1, g_1) \times (Z, h)], [(X_2, g_2) \times (Z, h)]) \in \mathbf{I}_{\text{sp}}[* , \cdot].$$

So the lemma follows from Remark 5.6.23. \square

5.11.2 Blowups and special bijections

Lemma 5.11.13. *Let $(Y, f) \in \text{Ob RV}[k, \cdot]$ and η a centripetal transformation on $\mathbb{L}(Y, f)$ with respect to a focus map λ whose locus is $\mathbb{L}(Y, f)$. Let $Z = (\text{pRV} \circ \mathbf{c} \circ \eta)(\mathbb{L}(Y, f))$. Then $(Z, \text{pr}_{\leq k}) \in \text{Ob RV}[k, \cdot]$ is isomorphic to an elementary blowup of (Y, f) .*

Proof. Suppose that $\text{dom}(\lambda) = \text{pr}_{>1} \mathbb{L}(Y, f)$. Without loss of generality, we may assume that $0 \notin \text{ran}(\lambda)$, that is, $\infty \notin \text{pr}_1 f(Y)$. Since λ is a function, for every $(r_1, \bar{r}_1) \in f(Y)$ and every $\bar{a}_1 \in \text{rv}^{-1}(\bar{r}_1)$ we have that $r_1 \in \text{acl}(\bar{a}_1)$ and hence, by Lemma 5.2.12, $r_1 \in \text{acl}(\bar{r}_1)$. So the elementary blowup $(Y^\#, f^\#)$ of (Y, f) with respect to the first coordinate of $f(Y)$ does exist. Note that, by Convention 5.7.1, $\text{pr}_{>k} Z$ is the companion Y_f of (Y, f) . Clearly the function $F : Z \rightarrow Y^\#$ given by

$$(r_1, \bar{r}_1, f(t_1, \bar{t}_1), t_1, \bar{t}_1) \mapsto (t_1, \bar{t}_1, r_1/f_{|1}(t_1, \bar{t}_1))$$

is an isomorphism between $(Z, \text{pr}_{\leq k})$ and $(Y^\#, f^\#)$, where $(t_1, \bar{t}_1) \in Y$ and $f(t_1, \bar{t}_1) = (f_{|1}(t_1, \bar{t}_1), \bar{r}_1)$. \square

Corollary 5.11.14. *Let $(X, g) \in \text{Ob RV}[k, \cdot]$ and T a special bijection on $\mathbb{L}(X, g)$. Let $(\text{pRV} \circ T)(\mathbb{L}(X, g)) = Z$. Then $(Z, \text{pr}_{\leq k}) \in \text{Ob RV}[k, \cdot]$ is isomorphic to an iterated blowup of (X, g) .*

Proof. By induction on the length $\text{lh} T$ of T and Lemma 5.11.6, this is immediately reduced to the case $\text{lh} T = 1$, which follows from Lemma 5.11.13. \square

Corollary 5.11.15. *Let $(X_1, g_1), (X_2, g_2) \in \text{Ob RV}[1, \cdot]$ be such that $\mathbb{L}(X_1, g_1)$ is definably bijective to $\mathbb{L}(X_2, g_2)$. Then $([(X_1, g_1)], [(X_2, g_2)]) \in \mathbf{I}_{\text{sp}}$.*

Proof. Let $(\text{pRV} \circ T_1)(\mathbb{L}(X_1, g_1)) = Z_1$ and $(\text{pRV} \circ T_2)(\mathbb{L}(X_2, g_2)) = Z_2$. By Corollary 5.10.4 and Remark 5.6.19, there are special bijections T_1, T_2 on $\mathbb{L}(X_1, g_1), \mathbb{L}(X_2, g_2)$ such that (Z_1, pr_1) and (Z_2, pr_1) are isomorphic. So the corollary follows from Corollary 5.11.14. \square

Lemma 5.11.16. *Suppose that the substructure S is (VF, Γ) -generated. Let $(Y^\#, f^\#)$ be an elementary blowup of $(Y, f) \in \text{Ob RV}[k, \cdot]$. Then there is a special bijection T of length 1 on $\mathbb{L}(Y, f)$ such that there is a commutative diagram*

$$\begin{array}{ccccc} \mathbb{L}(Y, f) & \xrightarrow{T} & T(\mathbb{L}(Y, f)) & \xrightarrow{\text{pRV}} & Z \\ & & \downarrow F & & \downarrow F_{\downarrow} \\ & & \mathbb{L}(Y^\#, f^\#) & \xrightarrow{\text{pRV}} & Y^\# \end{array}$$

where $Z = (\text{pRV} \circ T)(\mathbb{L}(Y, f))$ and both F and F_{\downarrow} are definably bijective.

Proof. For any $\bar{t} = (\bar{t}_k, t_k) = (t_1, \dots, t_{k-1}, t_k) \in f(Y)$ and any centripetal transformation η on $\text{rv}^{-1}(\bar{t})$ with respect to a focus map λ on $\text{rv}^{-1}(\bar{t}_k)$, the function

$$F_{\bar{t}} : (\mathbf{c} \circ \eta)(\text{rv}^{-1}(\bar{t}) \times f^{-1}(\bar{t})) \longrightarrow \mathbb{L}(f^{-1}(\bar{t}) \times \text{RV}^{>1}, f^\sharp \upharpoonright (f^{-1}(\bar{t}) \times \text{RV}^{>1}))$$

given by

$$(\mathbf{c} \circ \eta)(\bar{a}_k, a_k, \bar{s}) \longmapsto (\bar{a}_k, a_k - \lambda(\bar{a}_k), \bar{s}, \text{rv}(a_k - \lambda(\bar{a}_k))/t_k)$$

is a bijection as required. So, by compactness, it is enough to show that there is a \bar{t} -definable focus map λ such that $\text{rv}^{-1}(\bar{t}_k) \times f^{-1}(\bar{t}) \subseteq \text{dom}(\lambda)$. Let $v_{\text{rv}}(\bar{t}) = (\gamma_1, \dots, \gamma_k) = \bar{\gamma}$. Since $t_k \in \text{acl}(t_1, \dots, t_{k-1})$ and $t_k \neq \infty$, by Lemma 5.9.4, there is a $\bar{\gamma}$ -polynomial $p(x_1, \dots, x_k) = p(\bar{x})$ with coefficients in $\text{VF}(\langle \emptyset \rangle)$ such that \bar{t} is a residue root of $p(\bar{x})$ but is not a residue root of $\partial p(\bar{x})/\partial x_k$. This means that, for every $\bar{a}_k \in \text{rv}^{-1}(\bar{t}_k)$, t_k is a simple residue root of the γ_n -polynomial $p(\bar{a}_k, x_k)$ and hence, by the generalized Hensel's Lemma 5.9.2, there is a unique $a_k \in \text{rv}^{-1}(t_k)$ such that $p(\bar{a}_k, a_k) = 0$. So there exists a focus map as desired. \square

Remark 5.11.17. By the conclusion of Lemma 5.11.16, F is a lift of F_\downarrow and hence, by Remark 5.6.19, $F_\downarrow \in \text{Mor RV}[k, \cdot]$. This gives an alternative proof of Lemma 5.11.4 for the case that the substructure S is (VF, Γ) -generated.

Corollary 5.11.18. *Suppose that the substructure S is (VF, Γ) -generated. Let $(X, g), (Y, f) \in \text{RV}[k, \cdot]$ be isomorphic and (Y^\sharp, f^\sharp) an iterated blowup of (Y, f) of length l . Then there is a special bijection T of length l on $\mathbb{L}(X, g)$ such that $((\text{pRV} \circ T)(\mathbb{L}(X, g)), \text{pr}_{\leq k}), (Y^\sharp, f^\sharp)$ are isomorphic.*

Proof. By induction this is immediately reduced to the case $l = 1$, which follows from Corollary 5.11.5 and Lemma 5.11.16. \square

Let $X \in \text{Ob VF}[k, \cdot]$ be an RV-product and $p : X \longrightarrow \text{RV}^m$ a definable contractible function. For each $\bar{s} \in \text{RV}^m$ let $\eta_{\bar{s}}$ be an \bar{s} -definable centripetal transformation on $p^{-1}(\bar{s}) \times \text{rv}^{-1}(\bar{s})$ with respect to a VF-coordinate of X . Let

$$Z = \bigcup \{p^{-1}(\bar{s}) \times \text{rv}^{-1}(\bar{s}) : \bar{s} \in \text{RV}^m\}$$

and η the centripetal transformation on Z given by $(\bar{a}, \bar{b}) \longmapsto \eta_{\text{rv}(\bar{b})}(\bar{a}, \bar{b})$. By compactness, η is definable. We identify p with the function on $(\mathbf{c} \circ \eta)(Z)$ induced by p and projection.

Definition 5.11.19. The subset $(\mathbf{c} \circ \eta)(Z)$ is a *propagation of X with respect to p* , or a *p -propagation* for short. An *iterated propagation X^\sharp of X* is a composition of finitely many propagations with respect to a sequence of functions of the form $p_1, p_1 \times p_2, \dots, p_1 \times \dots \times p_n$. To specify the functions involved, we also say

that it is a (p_1, \dots, p_n) -propagation. An iterated propagation X^\sharp of X^\sharp is a *continuation* of X^\sharp if it starts with a function of the form $p_1 \times \dots \times p_n \times p_{n+1}$. The *length* of an iterated propagation is the length of the composition.

Note that if X^\sharp is a (p_1, \dots, p_n) -propagation of $X \in \text{Ob VF}[k, \cdot]$ and $\text{ran}(\prod \bar{p}) \subseteq \text{RV}^m$ then $X^\sharp \in \text{Ob VF}[k+m, \cdot]$.

Lemma 5.11.20. *Let $X \in \text{Ob VF}[k, \cdot]$ be an RV-product. Let X^\sharp be a (p_1, \dots, p_l) -propagation of X with $\text{ran}(\prod \bar{p}) \subseteq \text{RV}^m$. Then there is a (q_1, q_2, \dots, q_l) -blowup (Z^\sharp, h^\sharp) of $(\text{pRV } X, \text{pr}_{\leq k})$ such that (Z^\sharp, h^\sharp) , $(\text{pRV } Y^\sharp, \text{pr}_{\leq k+m})$ are parametrically isomorphic with respect to $\prod \bar{q}$, $\prod \bar{p}$.*

Proof. By induction on l and Lemma 5.11.9, this is immediately reduced to the case $l = 1$, which follows from compactness and a fiberwise application of Lemma 5.11.13, as in Corollary 5.11.14. \square

Lemma 5.11.21. *Suppose that the substructure S is (VF, Γ) -generated. Let $(X, g), (Y, f) \in \text{Ob RV}[k, \cdot]$ be parametrically isomorphic with respect to o_1, o_2 , where $\text{ran } o_1, \text{ran } o_2 \subseteq \text{RV}^n$. Let (Y^\sharp, f^\sharp) be a (p_1, \dots, p_l) -blowup of (Y, f) with $\text{ran}(\prod \bar{p}) \subseteq \text{RV}^m$. Then there is an $(o_1 \times q_1, q_2, \dots, q_l)$ -propagation $\mathbb{L}(X, g)^\sharp$ of $\mathbb{L}(X, g)$ such that (Y^\sharp, f^\sharp) , $(\text{pRV}(\mathbb{L}(X, g)^\sharp), \text{pr}_{\leq k+n+m})$ are parametrically isomorphic with respect to $o_2 \times \prod \bar{p}$, $o_1 \times \prod \bar{q}$.*

Proof. By induction this is immediately reduced to the case $l = 1$. In that case, let $q : X \rightarrow \text{RV}^m$ be the function induced by p_1 and the fiberwise isomorphisms. Then the lemma follows from compactness and a fiberwise application of Corollary 5.11.5 and Lemma 5.11.16, as in Corollary 5.11.18. \square

Lemma 5.11.22. *Let $X_1 \subseteq \text{VF}^n \times \text{RV}^{m_1}, X_2 \subseteq \text{VF}^n \times \text{RV}^{m_2}$ be two definable subsets such that $\text{pVF } X_1 = \text{pVF } X_2 = A$. Suppose that there is an $E \subseteq \mathbb{N}$ with $|E| = k$ such that $([\text{fib}(X_1, \bar{a})]_E, [\text{fib}(X_2, \bar{a})]_E) \in \text{I}_{\text{sp}}(\bar{a})$ for every $\bar{a} \in A$. Let $\widehat{I}_\sigma, \widehat{J}_\sigma$ be two standard contractions of X_1, X_2 and $E' = E \cup \{1, \dots, n\}$. Then*

$$([\widehat{I}_\sigma(X_1), \text{pr}_{E'}], [\widehat{J}_\sigma(X_2), \text{pr}_{E'}]) \in \text{I}_{\text{sp}}.$$

Proof. By induction on n this is immediately reduced to the case $n = 1$. By an argument similar to the one in the proof of Lemma 5.10.2, there is a special bijection T_A on A such that the following hold.

1. $T_A(A) = A^\sharp$ is an RV-product.
2. For every rv-polyball $\mathfrak{p} \subseteq A^\sharp$ and every $a_1, a_2 \in A$ with $T_A(a_1), T_A(a_2) \in \mathfrak{p}$, $\text{fib}(X_1, a_1) = \text{fib}(X_1, a_2)$ and $\text{fib}(X_2, a_1) = \text{fib}(X_2, a_2)$.

3. Let $h_1 = T_A \circ (\text{pVF} \upharpoonright X_1)$ and $h_2 = T_A \circ (\text{pVF} \upharpoonright X_2)$. For any rv-polyball $\mathfrak{p} \subseteq A^\sharp$ let $h_1^{-1}(\mathfrak{p}) = A_{\mathfrak{p}} \times U_{\mathfrak{p},1}$ and $h_2^{-1}(\mathfrak{p}) = A_{\mathfrak{p}} \times U_{\mathfrak{p},2}$, where $A_{\mathfrak{p}} \subseteq A$ and $U_{\mathfrak{p},1} = \text{fib}(X_1, a)$, $U_{\mathfrak{p},2} = \text{fib}(X_2, a)$ for any $a \in A_{\mathfrak{p}}$. There is a formula ϕ such that, for any $a \in A_{\mathfrak{p}}$, $\phi(a)$ defines the same iterated parameterized blowups that witness $([U_{\mathfrak{p},1}]_E, [U_{\mathfrak{p},2}]_E) \in \text{I}_{\text{sp}}\langle a \rangle$, and hence $([U_{\mathfrak{p},1}]_E, [U_{\mathfrak{p},2}]_E) \in \text{I}_{\text{sp}}\langle \text{pRV } \mathfrak{p} \rangle$.

Let $Y_1 = \mathbb{L}(\widehat{I}_\sigma(X_1), \text{pr}_1)$ and $Y_2 = \mathbb{L}(\widehat{J}_\sigma(X_2), \text{pr}_1)$. Let

$$f_1 = T_A \circ \text{pVF} \circ I_\sigma^{-1} : Y_1 \longrightarrow A^\sharp$$

and

$$f_2 = T_A \circ \text{pVF} \circ J_\sigma^{-1} : Y_2 \longrightarrow A^\sharp.$$

By Lemma 5.10.2, there are special bijections T_1, T_2 on Y_1, Y_2 such that $f_1 \circ T_1^{-1}, f_2 \circ T_2^{-1}$ are contractible.

Let $\mathfrak{p} \subseteq A^\sharp$ be an rv-polyball and $T_{\mathfrak{p}}$ a special bijection on $A_{\mathfrak{p}}$ such that $T_{\mathfrak{p}}(A_{\mathfrak{p}}) = A_{\mathfrak{p}}^\sharp$ is an RV-product. Since $\text{I}_{\text{sp}}\langle \text{pRV } \mathfrak{p} \rangle$ is a semiring congruence relation, by the third item above, clearly we have that

$$((\text{pRV } A_{\mathfrak{p}}^\sharp) \times U_{\mathfrak{p},1}]_{E'}, [(\text{pRV } A_{\mathfrak{p}}^\sharp) \times U_{\mathfrak{p},2}]_{E'} \in \text{I}_{\text{sp}}\langle \text{pRV } \mathfrak{p} \rangle.$$

For any $\bar{t} \in \text{pr}_E U_{\mathfrak{p},1}$, $\text{fib}((T_1 \circ f_1^{-1})(\mathfrak{p}), \bar{t})$ is an RV-product that is $(\text{pRV } \mathfrak{p}, \bar{t})$ -definably bijective to the RV-product $\text{fib}(A_{\mathfrak{p}}^\sharp \times U_{\mathfrak{p},1}, \bar{t})$. By Corollary 5.11.15, we have that

$$([\text{pRV } \text{fib}((T_1 \circ f_1^{-1})(\mathfrak{p}), \bar{t})]_1, [\text{pRV } \text{fib}(A_{\mathfrak{p}}^\sharp \times U_{\mathfrak{p},1}, \bar{t})]_1) \in \text{I}_{\text{sp}}\langle \text{pRV } \mathfrak{p}, \bar{t} \rangle$$

and hence, by compactness,

$$([\text{pRV } \circ T_1 \circ f_1^{-1})(\mathfrak{p})]_{E'}, [(\text{pRV } A_{\mathfrak{p}}^\sharp) \times U_{\mathfrak{p},1}]_{E'} \in \text{I}_{\text{sp}}\langle \text{pRV } \mathfrak{p} \rangle.$$

Symmetrically we have that

$$([\text{pRV } \circ T_2 \circ f_2^{-1})(\mathfrak{p})]_{E'}, [(\text{pRV } A_{\mathfrak{p}}^\sharp) \times U_{\mathfrak{p},2}]_{E'} \in \text{I}_{\text{sp}}\langle \text{pRV } \mathfrak{p} \rangle$$

and hence

$$([\text{pRV } \circ T_1 \circ f_1^{-1})(\mathfrak{p})]_{E'}, [(\text{pRV } \circ T_2 \circ f_2^{-1})(\mathfrak{p})]_{E'} \in \text{I}_{\text{sp}}\langle \text{pRV } \mathfrak{p} \rangle.$$

Since $f_1 \circ T_1^{-1}$, $f_2 \circ T_2^{-1}$ are contractible, we deduce that

$$([(pRV \circ T_1)(Y_1)]_{E'}, [(pRV \circ T_2)(Y_2)]_{E'}) \in \mathbf{I}_{\text{sp}}.$$

For any $\bar{t} \in \text{pr}_E \widehat{I}_\sigma(X_1)$, since T_1, T_2 are special bijections, $\text{fib}(Y_1, \bar{t})$ is an RV-product that is \bar{t} -definably bijective to the RV-product $\text{fib}(T_1(Y_1), \bar{t})$. As above, by Corollary 5.11.15 and compactness, we conclude that

$$([(pRV \circ T_1)(Y_1)]_{E'}, [\widehat{I}_\sigma(X_1)]_{E'}) \in \mathbf{I}_{\text{sp}}.$$

and, symmetrically,

$$([(pRV \circ T_2)(Y_2)]_{E'}, [\widehat{J}_\sigma(X_2)]_{E'}) \in \mathbf{I}_{\text{sp}}.$$

The claim follows. □

Corollary 5.11.23. *Let $X_1 \subseteq \text{VF}^n \times \text{RV}^{m_1}$, $X_2 \subseteq \text{VF}^n \times \text{RV}^{m_2}$ be two definable subsets and $f : X_1 \rightarrow X_2$ a unary bijection relative to the coordinate i . Then for any permutation σ of $\{1, \dots, n\}$ with $\sigma(1) = i$ and any standard contractions $\widehat{I}_\sigma, \widehat{J}_\sigma$ of X_1, X_2 ,*

$$([\widehat{I}_\sigma(X_1)]_{\leq n}, [\widehat{J}_\sigma(X_2)]_{\leq n}) \in \mathbf{I}_{\text{sp}}.$$

Proof. Let $E = \{1, \dots, n\} \setminus \{i\}$. For any $\bar{a} \in \text{pr}_E X_1 = \text{pr}_E X_2$ and any \bar{a} -definable standard contractions \widehat{I}, \widehat{J} of $\text{fib}(X_1, \bar{a}), \text{fib}(X_2, \bar{a})$, by Corollary 5.11.15, we have that

$$([\widehat{I}(\text{fib}(X_1, \bar{a}))]_1, [\widehat{J}(\text{fib}(X_2, \bar{a}))]_1) \in \mathbf{I}_{\text{sp}}(\bar{a}).$$

Then the corollary follows from Lemma 5.11.22. □

Lemma 5.11.24. *Let $X \subseteq \text{VF}^n \times \text{RV}^m$ be a definable subset. Let $i, j \in \{1, \dots, n\}$ be distinct and σ_1, σ_2 two permutations of $\{1, \dots, n\}$ such that $\sigma_1(1) = \sigma_2(2) = i$, $\sigma_1(2) = \sigma_2(1) = j$, and $\sigma_1 \upharpoonright \{3, \dots, n\} = \sigma_2 \upharpoonright \{3, \dots, n\}$. Then, for any standard contractions $\widehat{I}_{\sigma_1}, \widehat{I}_{\sigma_2}$ of X ,*

$$([\widehat{I}_{\sigma_1}(X)]_{\leq n}, [\widehat{I}_{\sigma_2}(X)]_{\leq n}) \in \mathbf{I}_{\text{sp}}.$$

Proof. Let ij, ji denote the permutations of $\{i, j\}$ and $E = \{1, \dots, n\} \setminus \{i, j\}$. By compactness and Lemma 5.11.22, it is enough to show that, for any $\bar{a} \in \text{pr}_E X$ and any standard contractions $\widehat{I}_{ij}, \widehat{I}_{ji}$ of

$\text{fib}(X, \bar{a})$,

$$([\widehat{I}_{ij}(\text{fib}(X, \bar{a}))]_{\leq 2}, [\widehat{I}_{ji}(\text{fib}(X, \bar{a}))]_{\leq 2}) \in \text{I}_{\text{sp}}\langle \bar{a} \rangle.$$

To that end, fix an $\bar{a} \in \text{pr}_E X$ and let $Y = \text{fib}(X, \bar{a})$. By Corollary 5.10.9, there are a definable bijection $f : Y \rightarrow \text{VF}^2 \times \text{RV}^k$ that is unary relative to both coordinates and two standard contractions $\widehat{J}_{ij}, \widehat{J}_{ji}$ of $f(Y)$ such that, for every $\bar{t} \in \text{pRV } f(Y)$,

$$([\widehat{J}_{ij}(\text{fib}(f(Y), \bar{t}))]_{\leq 2}, [\widehat{J}_{ji}(\text{fib}(f(Y), \bar{t}))]_{\leq 2}) \in \text{I}_{\text{sp}}\langle \bar{a}, \bar{t} \rangle.$$

and hence

$$([\widehat{J}_{ij}(f(Y))]_{\leq 2}, [\widehat{J}_{ji}(f(Y))]_{\leq 2}) \in \text{I}_{\text{sp}}\langle \bar{a} \rangle.$$

Now the desired property follows from Corollary 5.11.23. \square

Definition 5.11.25. Let $\text{I}_{\text{bu}}[k, \cdot]$ be the subclass of $\text{Ob VF}[k, \cdot] \times \text{Ob VF}[k, \cdot]$ of those pairs (X, Y) such that

1. X, Y are RV-products,
2. there are a \bar{p} -propagation X^\sharp of X and a \bar{q} -propagation Y^\sharp of Y , where $\text{ran } p, \text{ran } q \subseteq \text{RV}^m$, such that X^\sharp, Y^\sharp are isomorphic over their projections to RV^m , that is, for every $\bar{s} \in \text{RV}^m$, $\text{fib}(X^\sharp, \bar{s})$ and $\text{fib}(Y^\sharp, \bar{s})$ are \bar{s} -definably bijective.

Let

$$\text{I}_{\text{bu}}[* , \cdot] = \bigcup_{k \geq 0} \text{I}_{\text{bu}}[k, \cdot],$$

$$\text{I}_{\text{bu}}[k] = \text{I}_{\text{bu}}[k, \cdot] \cap (\text{Ob VF}[k] \times \text{Ob VF}[k]),$$

$$\text{I}_{\text{bu}}[*] = \bigcup_{k \geq 0} \text{I}_{\text{bu}}[k].$$

We will just write I_{bu} for all these classes if there is no danger of confusion.

If the substructure S is (VF, Γ) -generated then the congruence relation I_{sp} is the congruence relation induced by \mathbb{L} modulo I_{bu} :

Proposition 5.11.26. *Suppose that the substructure S is (VF, Γ) -generated. Let $(X, g), (Y, f) \in \text{Ob RV}[k, \cdot]$.*

Then

$$(\mathbb{L}(X, g), \mathbb{L}(Y, f)) \in \text{I}_{\text{bu}} \text{ if and only if } ((X, g), (Y, f)) \in \text{I}_{\text{sp}}.$$

Proof. The “if” direction follows immediately from Lemma 5.11.21 and Corollary 5.9.7.

For the “only if” direction, let $\mathbb{L}(X, g)^\sharp$ be a \bar{p} -propagation of $\mathbb{L}(X, g)$ and $\mathbb{L}(Y, f)^\sharp$ a \bar{q} -propagation of $\mathbb{L}(Y, f)$, where $\text{ran } p, \text{ran } q \subseteq \text{RV}^m$, such that they are isomorphic over their projections to RV^m . By Lemma 5.11.20, Lemma 5.11.9, and compactness, it is enough to show that, for every $\bar{s} \in \text{RV}^m$,

$$((\text{fib}(X^\sharp, \bar{s}), \text{pr}_{\leq k+m}), (\text{fib}(Y^\sharp, \bar{s}), \text{pr}_{\leq k+m})) \in \mathbf{I}_{\text{sp}}\langle \bar{s} \rangle.$$

To that end, suppose that $F : \text{fib}(X^\sharp, \bar{s}) \longrightarrow \text{fib}(Y^\sharp, \bar{s})$ is an \bar{s} -definable bijection. By Lemma 5.10.7, there is a \bar{s} -definable partition X_1, \dots, X_n of $\text{fib}(X^\sharp, \bar{s})$ such that each $F_i = F \upharpoonright X_i$ is a composition of relatively unary bijections. By Lemma 5.7.14, there are \bar{s} -definable special bijections T_1, T_2 on $\text{fib}(X^\sharp, \bar{s}), \text{fib}(Y^\sharp, \bar{s})$ such that $T_1(X_i), (T_2 \circ F)(X_i)$ are RV-products for each i . Let

$$G_i = (T_2 \upharpoonright F(X_i)) \circ F_i \circ (T_1^{-1} \upharpoonright T_1(X_i)).$$

Note that each G_i is a composition of relatively unary bijections. By Corollary 5.11.14, it is enough to show that, for each i ,

$$([(p\text{RV} \circ T_1)(X_i)]_{\leq k}, [(p\text{RV} \circ T_2 \circ F)(X_i)]_{\leq k}) \in \mathbf{I}_{\text{sp}}.$$

This follows from Corollary 5.11.23 and Lemma 5.11.24. \square

Remark 5.11.27. Since \mathbf{I}_{sp} is a semigroup (semiring) congruence relation, by Proposition 5.7.14 and Proposition 5.11.26, we may also treat \mathbf{I}_{bu} as a semigroup (semiring) congruence relation.

5.12 Motivic integration

In this section we assume that the substructure S is (VF, Γ) -generated.

As before, the results will be stated for the more general categories $\text{RV}[k, \cdot], \text{RV}[\ast, \cdot]$, etc. By Remark 5.11.2, analogous results for the restricted categories $\text{RV}[k], \text{RV}[\ast]$, etc are easily seen to hold by accordingly restricted arguments.

Proposition 5.12.1. *For each $k \geq 0$ there is a canonical surjective homomorphism of Grothendieck semi-groups*

$$\int_+ : \mathbf{K}_+ \text{VF}[k, \cdot] \longrightarrow \mathbf{K}_+ \text{RV}[k, \cdot] / \mathbf{I}_{\text{sp}}$$

such that

$$\int_+ [X] = [(U, f)] / \mathbf{I}_{\text{sp}} \text{ if and only if } [X] / \mathbf{I}_{\text{bu}} = [\mathbb{L}(U, f)] / \mathbf{I}_{\text{bu}}.$$

Proof. Let $\mathbb{Q} : \mathbf{K}_+ \text{VF}[k, \cdot] \longrightarrow \mathbf{K}_+ \text{VF}[k, \cdot] / \mathbf{I}_{\text{bu}}$ be the quotient map. By Corollary 5.9.7, \mathbb{L} induces a canonical semigroup homomorphism

$$\mathbb{Q} \circ \mathbb{L} : \mathbf{K}_+ \text{RV}[k, \cdot] \longrightarrow \mathbf{K}_+ \text{VF}[k, \cdot] / \mathbf{I}_{\text{bu}}.$$

By Corollary 5.7.15, $\mathbb{Q} \circ \mathbb{L}$ is surjective. By Proposition 5.11.26, the semigroup congruence relation on $\mathbf{K}_+ \text{RV}[k, \cdot]$ induced by $\mathbb{Q} \circ \mathbb{L}$ is precisely \mathbf{I}_{sp} and hence $\mathbf{K}_+ \text{RV}[k, \cdot] / \mathbf{I}_{\text{sp}}$ is canonically isomorphic to $\mathbf{K}_+ \text{VF}[k, \cdot] / \mathbf{I}_{\text{bu}}$. Inverting this isomorphism and composing it with \mathbb{Q} , we obtain the desired homomorphism. \square

Putting together Proposition 5.12.1 for all k , we obtain:

Theorem 5.12.2. *There is a canonical surjective homomorphism of Grothendieck semirings*

$$\int_+ : \mathbf{K}_+ \text{VF}_*[\cdot] \longrightarrow \mathbf{K}_+ \text{RV}[* , \cdot] / \mathbf{I}_{\text{sp}}$$

such that

$$\int_+ [X] = [(U, f)] / \mathbf{I}_{\text{sp}} \text{ if and only if } [X] / \mathbf{I}_{\text{bu}} = [\mathbb{L}(U, f)] / \mathbf{I}_{\text{bu}}.$$

References

- [1] U. Abraham, *Proper forcing*, in “Handbook of Set Theory” (M. Foreman, A. Kanamori, M. Magidor, eds.).
- [2] Uri Abraham, Keith Devlin, and Saharon Shelah, *The consistency with CH of some consequences of Martin’s Axiom plus $2^{\aleph_0} > \aleph_1$* , Israel Journal of Mathematics **31** (1978), no. 1, 19–33.
- [3] Uri Abraham and Yimu Yin, *A note on the Engelking-Karłowicz theorem*, Acta Mathematica Hungarica (2008), <http://dx.doi.org/10.1007/s10474-008-7160-4>.
- [4] Jeremy Avigad and Yimu Yin, *Quantifier elimination for the reals with a predicate for the powers of two*, Theoretical Computer Science **370** (2007), no. 1-3, 48–59, cs.LO/0610117.
- [5] S. A. Basarab, *Relative elimination of quantifiers for Henselian valued fields*, Annals of Pure and Applied Logic **53** (1991), 51–74.
- [6] S. A. Basarab and F.-V. Kuhlmann, *An isomorphism theorem for Henselian algebraic extensions of valued fields*, Manuscripta Mathematica **77** (1992), 113–126.
- [7] S. Basu, *New results on quantifier elimination over real closed fields and applications to constraint databases*, J. ACM **46** (1999), no. 4, 537–555.
- [8] S. Basu, R. Pollack, and M. F. Roy, *Algorithms in real algebraic geometry*, Algorithms and Computation in Mathematics, vol. 10, Springer-Verlag, Berlin, 2003.
- [9] Nicolas Bourbaki, *Elements of mathematics, commutative algebra, chapters 1-7*, Springer-Verlag, Berlin, 2003.
- [10] B. F. Caviness and J. R. Johnson (eds.), *Quantifier elimination and cylindrical algebraic decomposition*, Texts and Monographs in Symbolic Computation, Springer-Verlag, Vienna, 1998.
- [11] C. C. Chang and H. J. Keisler, *Model theory*, third ed., Studies in logic and the foundations of mathematics, vol. 73, North-Holland, Amsterdam, 1990.
- [12] Raf Cluckers, *Model theory of valued fields*, Quaderni di matematica, vol. 11, pp. 73–85, Dipartimento di Matematica della Seconda Università di Napoli, 2002.
- [13] Raf Cluckers and François Loeser, *b-minimality*, math.LO/0610183.
- [14] ———, *Constructible motivic functions and motivic integration*, Inventiones Mathematicae **173** (2008), no. 1, 23–121, math.AG/0410203.
- [15] G. E. Collins, *Quantifier elimination for real closed fields by cylindrical algebraic decomposition*, Lecture Notes in Computer Science, vol. 33, pp. 134–183, Springer, Berlin, 1975, reprinted in [10].
- [16] Jan Denef, *p -adic semi-algebraic sets and cell decomposition*, Journal für die reine und angewandte Mathematik **369** (1986), 154–166.

- [17] M. A. Dickmann, *Elimination of quantifiers for ordered valuation rings*, The Journal of Symbolic Logic **52** (1987), 116–128.
- [18] Lou van den Dries, *The field of reals with a predicate for the powers of two*, Manuscripta Mathematica **54** (1985), 187–195.
- [19] ———, *Alfred Tarski's elimination theory of real closed fields*, The Journal of Symbolic Logic **53** (1988), no. 1, 7–19.
- [20] ———, *Tame topology and o-minimal structures*, LMS Lecture Note Series, vol. 248, Cambridge University Press, Cambridge, UK, 1998.
- [21] Lou van den Dries and Ayhan Günaydin, *The fields of real and complex numbers with a small multiplicative group*, Proc. London Math. Soc. **93** (2006), no. 1, 43–81.
- [22] R. Engelking and M. Karłowicz, *Some theorems of set theory and their topological consequences*, Fundamenta Mathematicae **57** (1965), 275–285.
- [23] Antonio J. Engler and Alexander Prestel, *Valued fields*, Springer-Verlag, Berlin, 2005.
- [24] Paul Erdős and András Hajnal, *On decomposition of graphs*, Acad. Sci. Hungar. **18** (1967), 359–377.
- [25] ———, *On chromatic number of finite graphs*, pp. 83–98, Academic Press, New York, 1968.
- [26] J. Ferrante and C. W. Rackoff, *The computational complexity of logical theories*, Lecture Notes in Mathematics, vol. 718, Springer, Berlin, 1979.
- [27] M. J. Fischer and M. O. Rabin, *Super-exponential complexity of Presburger arithmetic*, SIAM-AMS Proc., vol. VII, pp. 27–41, Amer. Math. Soc., Providence, R.I., 1974, reprinted in [10].
- [28] Harvey Friedman and Chris Miller, *Expansions of o-minimal structures by fast sequences*, Journal of Symbolic Logic **70** (2005), no. 2, 410–418.
- [29] A. Hajnal and A. Máté, *Set mappings, partitions, and chromatic numbers*, Logic colloquium '73 : proceedings of the Logic Colloquium, Bristol, July 1973 / ed. by H. E. Rose and J. C. Shepherdson, North-Holland Pub. Co., 1975, pp. 347–379.
- [30] Thomas C. Hales, *What is motivic measure?*, Bulletin of the American Mathematical Society **42** (2005), no. 2, 119–135, math.AG/0312229.
- [31] Deirdre Haskell, Ehud Hrushovski, and Dugald Macpherson, *Definable sets in algebraically closed valued fields: elimination of imaginaries*, Journal für die reine und angewandte Mathematik **597** (2006), 175–236.
- [32] Deirdre Haskell and Dugald Macpherson, *Cell decomposition of C-minimal structures*, Annals of Pure and Applied Logic **66** (1994), 113–162.
- [33] W. Hodges, *Model theory*, Cambridge University Press, Cambridge, UK, 1993.
- [34] Ehud Hrushovski and David Kazhdan, *Integration in valued fields*, math.AG/0510133.
- [35] G. Kreisel and J. L. Krivine, *Elements of mathematical logic*, North-Holland Publishing Company, Amsterdam, 1971.
- [36] Franz-Viktor Kuhlmann, *Quantifier elimination for henselian fields relative to additive and multiplicative congruences*, Israel J. Math. **85** (1994), no. 1-3, 277–306.
- [37] Angus Macintyre, *On definable subsets of p-adic fields*, Journal of Symbolic Logic **41** (1976), 605–610.
- [38] Angus Macintyre, Kenneth McKenna, and Lou van den Dries, *Elimination of quantifiers in algebraic structures*, Advances in Mathematics **47** (1983), 74–87.

- [39] Dugald Macpherson and Charles Steinhorn, *On variants of o-minimality*, Annals of Pure and Applied Logic **79** (1996), no. 2, 165–209.
- [40] David Marker, *Model theory: an introduction*, Graduate Texts in Mathematics, vol. 217, Springer, New York, 2002.
- [41] Johan Pas, *Uniform p -adic cell decomposition and local zeta functions*, Journal für die reine und angewandte Mathematik **399** (1989), 137–172.
- [42] Alexander Prestel and Peter Roquette, *Formally p -adic fields*, Lecture Notes in Mathematics, vol. 1050, Springer-Verlag, Berlin, 1984.
- [43] Gerald E. Sacks, *Saturated model theory*, W. A. Benjamin, Inc., Reading, Massachusetts, 1972.
- [44] Thomas Scanlon, *Quantifier elimination for the relative Frobenius*, Valuation Theory and Its Applications, Fields Institute Communications Series, vol. II, AMS, Providence, 2003, conference proceedings of the International Conference on Valuation Theory (Saskatoon, 1999), Franz-Viktor Kuhlmann, Salma Kuhlmann, and Murray Marshall, eds., pp. 323–352.
- [45] Saharon Shelah, *Remarks on Boolean algebras*, Algebra Universalis **11** (1980), 77–89.
- [46] ———, *Cellularity of free products of Boolean algebras (or topologies)*, Fundamenta Mathematicae **166** (2000), 153–208.
- [47] Joseph R. Shoenfield, *Mathematical logic*, Addison-Wesley, Reading, 1967.
- [48] ———, *A theorem on quantifier elimination*, Symposia Mathematica **5** (1971), 173–176, INDAM, Rome, 1969/1970, Academic Press, London.
- [49] Alfred Tarski, *A decision method for elementary algebra and geometry*, second ed., University of California Press, Berkeley, CA, 1951, prepared for publication with the assistance of J. C. C. McKinsey.
- [50] Volker Weispfenning, *Model theory of lattice products*, Habilitationsschrift, Universität Heidelberg, 1978.
- [51] ———, *Quantifier elimination and decision procedures for valued fields*, Models and Sets, LNM, vol. 1103, Springer, 1983, Proc. Logic Coll. '83, Aachen, part 1, pp. 419–472.
- [52] ———, *The complexity of almost linear Diophantine problems*, J. Symbolic Computation **10** (1990), no. 5, 395–403.
- [53] ———, *Complexity and uniformity of elimination in Presburger arithmetic*, pp. 48–53, ACM, New York, 1997, electronic.
- [54] Yimu Yin, *On logical characterization of henselianity*, arXiv:0705.0395.