

Definable sets in algebraically closed valued fields: elimination of imaginaries

Deirdre Haskell ^{*} Ehud Hrushovski [†] Dugald Macpherson [‡]

July 19, 2005

2000 Mathematics Subject Classification: 03C60

Abstract.

It is shown that if K is an algebraically closed valued field with valuation ring R , then $\text{Th}(K)$ has elimination of imaginaries if sorts are added whose elements are certain cosets in K^n of certain definable R -submodules of K^n (for all $n \geq 1$). The proof involves the development of a theory of independence for *unary types*, which play the role of 1-types, followed by an analysis of germs of definable functions from unary sets to the sorts.

Contents.

1. Introduction.
2. The geometric sorts and unary types
 - 2.1 Fundamental definitions and elementary properties
 - 2.2 Definable modules
 - 2.3 Unary sets
 - 2.4 Definable functions from Γ
 - 2.5 Independence and orthogonality for unary types
 - 2.6 Sets internal to k
3. Elimination of imaginaries
 - 3.1 Quantifier elimination for the geometric sorts
 - 3.2 Preliminaries on coding
 - 3.3 Coding of functions
 - 3.4 Proof of elimination of imaginaries
 - 3.5 Necessity of the geometric sorts

^{*}Research partially supported by NSF and NSERC grants.

[†]This research was conducted by Ehud Hrushovski as a CMI Prize Fellow, for the Clay Mathematics Institute.

[‡]Research supported partially by an EPSRC grant. All three authors were supported by NSF funding for a semester at the Mathematical Sciences Research Institute, Berkeley.

1 Introduction

The purpose of this paper is to give the foundations of a study of structures which are first-order *interpretable* in an algebraically closed valued field; that is, which live on a quotient of a power of the field by a definable equivalence relation. The paper will have a successor [3]. The latter develops some of the tools of stability theory for algebraically closed valued fields, which, like o-minimal structures, have the strict order property, so are unstable.

A complete multi-sorted theory T is said to have *elimination of imaginaries* if the following holds: for any $M \models T$, any collection M_1, \dots, M_k of sorts in M , any \emptyset -definable $S \subset M_1 \times \dots \times M_k$, and any \emptyset -definable equivalence relation E on S , there is an \emptyset -definable function f from S into a product of sorts, such that for any $a, b \in S$ we have Eab if and only if $f(a) = f(b)$. Thus, $f(a)$ acts as a code in M for the E -class of a (an *imaginary*). The theory of pure algebraically closed fields has elimination of imaginaries, essentially because a Zariski closed set has a unique field of definition. The theory of real closed fields has elimination of imaginaries because the midpoint of an interval is definable from the parameters defining its endpoints, and these endpoints are determined by the interval.

That the theory of algebraically closed valued fields is reasonably tractable was proved by A. Robinson [15], who showed that the theory is model-complete (see Section 2.1), and described the completions. However, it is easy to see that this theory does not eliminate imaginaries in the Robinson language. The original impetus for the work in this and the subsequent paper was to prove elimination of imaginaries to a level suggested in the thesis of J. Holly, that is, relative to the sorts of the open and closed balls. It was shown by Holly [6] that in equi-characteristic 0, definable subsets of the field in one variable are coded in the ball sorts. In fact, it turns out that these sorts are too coarse to code all of the definable sets. Instead, we need some n -dimensional version of balls, and we see what these might be by thinking of balls algebraically. A general ball is a coset of a submodule (over the valuation ring) of the field, and to eliminate imaginaries, we add sorts for some of the *torsors*, that is, cosets of submodules of powers of the field. Our main theorem in this paper is the following.

Theorem 1.0.1 *The theory of algebraically closed valued fields in the sorted language $\mathcal{L}_{\mathcal{G}}$ (defined in Section 3.1) has elimination of imaginaries.*

In addition to the field sort, one requires certain ‘geometric’ sorts: spaces of lattices over the valuation ring, and vector spaces over the residue field, which arise as quotients of these lattices. The proposed language $\mathcal{L}_{\mathcal{G}}$ is not explicit, but is based on a stability-theoretic notion of a generic lattice basis. Elimination of quantifiers holds in this language.

If M is a structure, $\{R_i : i \in I\}$ is a collection of sorts in M^{eq} with $\mathcal{R} := \cup(R_i : i \in I)$, and $A \subset \mathcal{R}$, then an imaginary $i \in M^{\text{eq}}$ is *coded in \mathcal{R} over A* if

there is $e \in \text{dcl}(Ai)$, e a tuple from \mathcal{R} , with $i \in \text{dcl}(Ae)$. Now Theorem 1.0.1 says that if K is an algebraically closed valued field and \mathcal{G} is the collection of sorts for $\mathcal{L}_{\mathcal{G}}$, then every imaginary of K^{eq} is coded in \mathcal{G} over \emptyset . See Hodges [4] for more on the equivalence of this and the previously stated definition of elimination of imaginaries (note that K has two constant symbols, 0 and 1).

We give a more concrete version of this theorem. If M is a model and the definable set $X \subset M^n$ is the solution set of the formula $\varphi(x, a)$ say, where $a \in M^m$, there is an \emptyset -definable equivalence relation E_{φ} on M^n : $E_{\varphi}(y_1, y_2)$ if and only if $M \models \forall x(\varphi(x, y_1) \leftrightarrow \varphi(x, y_2))$. The E_{φ} -class of a is an imaginary which is a *code* for X ; it is unique up to interdefinability.

Theorem 1.0.2 *Let $(K, R, +, \cdot)$ be an algebraically closed valued field, with valuation ring R . Then for every imaginary e of K , there is for some n a definable R -submodule of K^n with a code interdefinable with e .*

In fact, we do not need sorts for all definable modules. It suffices to have a sort for elements of K , a sort S_n (for each n) whose elements are (codes for) R -lattices in K^n , that is, free rank n R -submodules of K^n ; and a sort T_n consisting of elements of $A/\mathcal{M}A$, where $A \in S_n$ and \mathcal{M} is the maximal ideal of R . (We also add sorts for the residue field and value group, for notational convenience.) The coding of e in Theorem 1.0.2 can be done by a tuple abc , where $a \in K^{\ell}$, $b \in T_m$, and $c \in S_n$, for some $\ell, m, n > 0$.

Two key roles in the paper are played by definable R -submodules of K^n and definable R -torsors, that is, cosets in K^n of definable R -submodules. They are used to code imaginaries; and certain specific torsors, namely *1-torsors* (Section 2.3) are used as a generalisation of 1-types.

Our first attempted proof of elimination of imaginaries had a stability-theoretic flavor, and this led us to develop notions of independence more systematically. In fact, the proof we give of elimination of imaginaries uses these ideas of independence only for 1-types (or rather, ‘unary types’), but the point of view turns out to be very helpful. In this first paper, we define both genericity and orthogonality to the value group for unary types, and investigate some of their properties. This is used both for the proof of elimination of imaginaries and to lay the groundwork for the subsequent paper [3], in which we develop the theory for n -types.

The residue field (usually denoted by k) of an algebraically closed valued field is a stably embedded pure algebraically closed field, so strongly minimal. The value group is a stably embedded divisible ordered abelian group, so o-minimal. As might be expected from Ax-Kochen-Ershov style results, the model-theoretic structure of an algebraically closed valued field can be understood in terms of these familiar theories, which can be loosely regarded as rank one geometries. The instability seems to live entirely in the value group. Over a base set of parameters C , an important role is played by an ω -stable structure whose sorts are the C -definable k -internal sets (see Section 2.6). We exploit this to construct

a counterexample to the original conjecture that algebraically closed valued fields eliminate imaginaries to the ball sorts. We exhibit a definable set which is internal to k , of Morley rank 2, but for which the algebraic closure of a generic element contains a unique algebraically closed subset of rank 1. If such a set were coded in the ball sorts, the algebraic closure of a generic element would have to contain at least two distinct rank one algebraically closed subsets.

The outline of the paper is as follows. In Section 2.1 we explain our notation and collect together some known facts about algebraically closed valued fields which we will use. In Section 2.2 we investigate definable modules, and their homomorphisms. Section 2.3 defines the unary types, which replace 1-types for general imaginaries. In Section 2.4 we analyse functions from the value group to the sorts with respect to which we eliminate imaginaries. In Section 2.5 we begin to develop the theory of the notions of independence and orthogonality that come from genericity for unary types, and in Section 2.6 we describe the structure of the sets internal to the residue field.

The proof of elimination of imaginaries is in Section 3. We first give (Section 3.1) a precise description of the first order language in which we work, and prove that the theory of algebraically closed valued fields has quantifier elimination in this language. It turns out that we really use the proof of this theorem, rather than the result itself, in the proof of elimination of imaginaries, but it is reassuring to have the result. In Section 3.2 we give a lemma reducing elimination of imaginaries to the coding of certain functions on unary sets. Then in Section 3.3 we prove that germs of definable functions on unary sets are coded, and deduce in Section 3.4 the elimination of imaginaries (coding of finite sets is first proved). Section 3.5 contains the example described above of a definable set which cannot be coded just in the ball sorts. It also contains a proof of a more general result, that to obtain elimination of imaginaries, one cannot make do with a finite sub-collection of the sorts.

Some results in this paper, principally in Sections 2.3, 2.5, and the beginning of Section 2.4, should hold in a more general setting. This might be the notion of *C-minimality*, developed in [12] and [2], or might be a related notion of minimality associated with a uniformly definable chain of equivalence relations on a set. A valuation on a group or field always provides such a chain. There is also a development of the model theory of ultrametric spaces (without any assumed algebraic structure) in [13].

Although the goal of the present paper is Theorem 1.0.1, it includes some results (such as Lemma 2.5.11 and Lemma 3.4.13) which are not used in this paper. However, they will be used in [3], and seem naturally to belong here.

Acknowledgements. We thank Françoise Delon, Bradd Hart, Franz-Viktor Kuhlmann, Tom Scanlon, Philip Scowcroft, and Tim Mellor for a number of helpful remarks. We particularly thank Marie-Hélène Mourgues and Patrick Simonetta for making us aware of a mistake in an earlier draft of this paper, which

led us to realise the significance of the role played by the torsors. We also thank Assaf Hasson, Moshe Kaminsky, Liat Kessler, Mark Ryten, Alex Usvyatsov, and Yoav Yaffe, for running a seminar in Jerusalem which led to several errors being pointed out.

2 The geometric sorts and unary types

2.1 Fundamental definitions and elementary properties

We work throughout with the following set-up. K is a $(2^{\aleph_0})^+$ -saturated homogeneous algebraically closed valued field, with value group Γ written multiplicatively and valuation map $|\cdot| : K \rightarrow \Gamma$. In fact, it would serve our purposes just to assume that K is sufficiently saturated; reference to automorphisms of K can always be replaced by reference to elementary partial maps. We order Γ so that $|x + y| \leq \text{Max}\{|x|, |y|\}$. For convenience, we adjoin to the value group the symbol 0 , so $0 \leq \gamma$ for all $\gamma \in \Gamma$, and $|x| = 0$ if and only if $x = 0$. The valuation ring is $R := \{x \in K : |x| \leq 1\}$, its maximal ideal is $\mathcal{M} := \{x \in R : |x| < 1\}$, and the residue field is $k = R/\mathcal{M}$. When necessary, we denote the residue map from R to k by res . To begin with, we will work in the language \mathcal{L}_{div} , which has the usual ring language $(+, -, \cdot, 0, 1)$ on K , and the binary predicate div , where $x \text{ div } y$ means $|y| \leq |x|$. We will take K^{eq} to be the sorted structure with respect to this language, which has sorts made up of all the imaginaries, that is, the equivalence classes of $\mathcal{L}_{\text{div}}\text{-}\emptyset$ -definable equivalence relations, and functions from powers of K to the sorts sending a tuple to the class to which it belongs. We emphasise that any language we consider later in the paper will have relations and functions which are definable in \mathcal{L}_{div} , and hence the collection of definable sets and the structure K^{eq} will always be the same, though the different languages have different sorts. As usual, *definable* means *definable with parameters*. Sets of parameters will always be assumed to be subsets of K^{eq} , and will always be taken to be small relative to the size of K , and in particular of cardinality at most 2^{\aleph_0} .

As noted in the introduction, any definable set X can be identified with an imaginary which is unique up to interdefinability, and is denoted $\ulcorner X \urcorner$. We refer to $\ulcorner X \urcorner$ as a *code* for X . We sometimes treat the definable set X as an imaginary by identifying it with $\ulcorner X \urcorner$. Our purpose is to identify a set \mathcal{G} of sorts in K^{eq} , and prove that any definable set X has a code in \mathcal{G} : that is, that there is a sequence e from \mathcal{G} such that $\text{dcl}(\ulcorner X \urcorner) = \text{dcl}(e)$. Sometimes, in particular contexts, we have to work over a parameter set C and prove that X is coded by e over C , that is $\text{dcl}(C \ulcorner X \urcorner) = \text{dcl}(Ce)$. When such C is not specified, it is assumed that such coding is over \emptyset .

We employ some notational conventions fairly consistently. Greek letters α, β, γ range over the value group or residue field. Lower-case letters a, b, c range

over the field, but also more generally over singletons or sequences of imaginaries. We do not usually distinguish notationally between singletons or sequences; more often than not $a = (a_1, \dots, a_n)$ is a sequence, of which the a_i may themselves be sequences. We use juxtaposition ab for concatenation of sequences; by extension we frequently write AB for $A \cup B$, where A and B are sets of (imaginary) parameters. In general, the upper-case letters are for sets, which may have additional structure, in particular, that of a module over the valuation ring. A type will usually be denoted by p or q , and the set of realisations of the type by the corresponding upper-case letter P or Q . Finally, if C is a set of parameters and a, b are tuples, then $a \equiv_C b$ means that $\text{tp}(a/C) = \text{tp}(b/C)$.

If A is a subset of K^{eq} , then $\text{acl}(A)$ is the algebraic closure of A in K^{eq} and $\text{dcl}(A)$ is the definable closure of A in K^{eq} . In general, a subscript denotes the intersection of the set with the specified sort; for example, $A_K := A \cap K$, $\text{acl}_K(A) = \text{acl}(A) \cap K$. We also write $\Gamma(A) := \text{dcl}(A) \cap \Gamma$, and $k(A) := \text{dcl}(A) \cap k$.

The following quantifier elimination results are basic to our theory.

Theorem 2.1.1 *Let K be an algebraically closed valued field.*

- (i) *The theory of K has quantifier elimination in the language L_{div} .*
- (ii) *The theory of K has quantifier elimination in a 2-sorted language with a sort K for the field (equipped with the language of rings), a sort Γ for the value group written multiplicatively (with the language $(\langle, \cdot, 0$ with usual conventions for 0), and a value map $|\cdot| : K \rightarrow \Gamma$ with $|0| = 0$.*
- (iii) *The theory of K has quantifier elimination in a 3-sorted language $\mathcal{L}_{\Gamma k}$ with the sorts and language of (ii) together with a sort k for the residue field, with the language of rings, and a map $\text{Res} : K^2 \rightarrow k$ given by putting $\text{Res}(x, y)$ equal to the residue of xy^{-1} (and taking value $0 \in k$ if $|x| > |y|$).*

We remark that in (i), the theory is model complete if the symbol div is replaced by a predicate for the valuation ring (the definable sets are the same). By [11], any valued field having quantifier-elimination in \mathcal{L}_{div} is algebraically closed. Also, by [15], the complete theory (in for example \mathcal{L}_{div}) of an algebraically closed non-trivially valued field K is determined by the pair $(\text{char}(K), \text{char}(k))$. This can take any of the values $(0, 0)$, $(0, p)$, or (p, p) (where p is a prime).

Proof. This is essentially due to A. Robinson [15], though only the model completeness is stated there. Part (i) is made explicit in [11], but follows quickly from the model completeness and existence of prime models in [15]. Part (ii) is Theorem 3.2 of [16].

Part (iii) follows from results in the Ph.D. thesis of F. Delon, but as far as we know there is no proof in print. We sketch the main steps. The task is the following. Given a large saturated model M of $\text{Th}(K)$ (in $\mathcal{L}_{\Gamma k}$), a substructure S of K , and isomorphism $\varphi : S \rightarrow S'$ with $S' \subset M$, and $c \in K \setminus S$, find $c' \in M$ such that we may extend φ to an isomorphism taking c to c' from the structure generated by $S \cup \{c\}$ to that generated by $S' \cup \{c'\}$. This is done as follows. First,

using the map Res , we may extend φ to the field of fractions of $S \cap K$ to ensure $S \cap K$ is a field. Second, for each $\alpha \in S \cap k$, if there is no element of $S \cap K$ with residue α , add some such element a to S , and a corresponding element $a' \in \alpha'$ to S' , and extend φ so that $\varphi(a) = a'$. If in this situation the algebraic closure of $S \cap K$ contains an element of residue α , choose a to be such an element, and a' to be an element with minimal polynomial (over $S' \cap K$) corresponding to that of a over $S \cap K$. Next, using the fact that extensions of a valuation to a finite normal extension are conjugate in the Galois group (ch. 4.2 of [14]), we may extend φ to the algebraic closure of $S \cap K$. Now, if $\gamma \in S$, we may assume there is $a \in S \cap K$ with $|a| = \gamma$. For otherwise, pick such a , and pick $a' \in M$ with $|a'| = \varphi(\gamma)$, and put $\varphi(a) = a'$. Then φ extends to an isomorphism of structures; for example, if $f, g \in S \cap K[X]$, then the condition $|f(a)| \leq |g(a)|$ is preserved by φ (split f, g as products of linear factors). After all of these reductions have been made, we may assume c is a field element transcendental over $S \cap K$. Extending φ to c is handled just as in [15]. \square

The definable subsets called *balls* play an important role in our theory (though not the one we first envisioned). If $a \in K$, $\alpha \in \Gamma$, then $B_{\leq \alpha}(a)$ is the ‘closed’ ball $\{x \in K : |x - a| \leq \alpha\}$ and $B_{< \alpha}(a)$ is the ‘open’ ball $\{x \in K : |x - a| < \alpha\}$. These balls are said to have *radius* α , and we write $\text{rad}(s)$ for the radius of the ball s . If s is a ball of radius γ , and $\delta > \gamma$, we extend this notation to write $B_{< \delta}(s)$ for the unique open ball (or $B_{\leq \delta}(s)$ for the unique closed ball) of radius δ containing s . We write simply $B_\delta(s)$ if we do not want to specify whether the ball is open or closed. Notice that an element a in K is just a closed ball $B_{\leq 0}(a)$ of radius 0, and the whole field K can be regarded informally as an open ball of infinite radius.

The following theorem of Holly ([5], Theorem 3.26) gives a precise description of definable sets in one variable. With our notation, a *Swiss cheese* is a non-empty set of the form $t \setminus (s_1 \cup \dots \cup s_n)$, where t (the *block*) is a ball of K or the whole of K , and the s_i (the *holes*) are distinct proper sub-balls (remember here that field elements are balls of radius zero). We allow the case when there are no s_i .

Theorem 2.1.2 (Holly) *Each parameter-definable set $X \subset K$ is a union of a unique set $\{S_1, \dots, S_m\}$ of disjoint Swiss cheeses such that no two are trivially nested, that is, for no i, j does the block of S_i equal a hole of S_j .*

The fact that any definable set X can be so expressed is an easy consequence of quantifier elimination in \mathcal{L}_{div} , which gives that any definable subset of K is a Boolean combination of singletons and balls. It is used frequently in the paper. The uniqueness of the expression is used in Section 2.3, and in the proof of Corollary 3.4.9.

If M is a structure and A is a set a -definable in M^{eq} , then A is *stably embedded* in M if, for any r and any definable set D in M^r , $D \cap A^r$ is definable over Aa

(uniformly in the parameters which define D). See (5.6) of [1] for more on this condition.

Proposition 2.1.3 (i) *The value group Γ of K is o -minimal in the sense that every K -definable subset of Γ is a finite union of intervals.*

(ii) *The residue field k is strongly minimal in the sense that any K -definable subset of k is finite or cofinite (uniformly in the parameters).*

(iii) *Γ is stably embedded in K .*

(iv) *If $A \subset K$ then the model-theoretic algebraic closure $\text{acl}(A) \cap K$ of A in the field sort K is equal to the field-theoretic algebraic closure.*

(v) *If $S \subset k$ and $\alpha \in k$ and $\alpha \in \text{acl}(S)$ (in the sense of K^{eq}), then α is in the field-theoretic algebraic closure of S in the sense of k .*

(vi) *k is stably embedded in K .*

Proof. All parts follow from quantifier elimination for algebraically closed valued fields. Parts (i) and (ii) follow immediately from Theorem 2.1.2. Part (iii) comes from Theorem 2.1.1(ii). Quantifier elimination in \mathcal{L}_{div} also yields (iv), and (v) comes from Theorem 2.1.1(iii).

For (vi), we again use Theorem 2.1.1(iii). It suffices to consider an atomic formula $\varphi(x, a)$ in this language, with x a tuple of residue field variables. If this mentions x , it has the form $f(x, \beta) = 0$, where $\beta = (\beta_1, \dots, \beta_n)$ and for each i , $\beta_i = \text{Res}(g_i(a), h_i(a))$, with $g_i(Y), h_i(Y) \in \mathbf{Z}[Y]$. Since the β_i are in k , the result follows. \square

For models of a theory to have property (iii) was originally considered by Shelah as a “stage 0 stability over Γ ”. Over a base model of size λ there can be 2^λ distinct types that do not increase Γ . Hence, in the main sense of the expression, the theory of an algebraically closed valued field is not stable over Γ .

Remark 2.1.4 In $(0, 0)$ or (p, p) characteristic, the algebraic closure of \emptyset (in K) is trivially valued, so no element of $\Gamma \setminus \{0, 1\}$ is definable over it. In any characteristics, if an element of $\Gamma \setminus \{0, 1\}$ is definable over an algebraically closed field C , then by quantifier elimination (in the language $\mathcal{L}_{\Gamma k}$ with sorts K , k , and Γ) some element of C is non-trivially valued. This always happens in mixed characteristic.

Notice that a ball containing 0 is an R submodule of K , and furthermore, every proper definable R submodule of K is such a ball. A ball which does not contain 0 is a coset of an R -module. In general, for $\gamma \in \Gamma$, we will write $\gamma R = B_{\leq \gamma}(0) = \{x \in K : |x| \leq \gamma\}$ and $\gamma \mathcal{M} = B_{< \gamma}(0) = \{x \in K : |x| < \gamma\}$. Then $\gamma R / \gamma \mathcal{M}$ is γ -definably a one-dimensional vector space over k , for every γ . In [5], Holly proved (in equicharacteristic 0) that the definable sets in one variable are coded by balls. It turns out that in order to code the definable sets of tuples, we need what one might call n -dimensional balls, that is, some of the R -modules and their cosets in K^n .

Definition 2.1.5 A *definable torsor* in K^n is a coset in K^n of a definable R -submodule of K^n . If the torsor X is a coset of the submodule U of K^n , then a *subtorsor* of X is a coset (contained in X) of an R -submodule of U . If also Y is a coset of the submodule V of K^m , then we define an *affine homomorphism* to be a pair (g, c) where $g \in \text{Hom}(U, V)$, c is a function from X to Y , and for all $u \in U$ and $x \in X$, $c(x+u) = g(u)+c(x)$. In particular, if (g_1, c) and (g_2, c) are both affine homomorphisms then $g_1 = g_2$, so we often refer to the affine homomorphism as c , with *homogeneous component* g . We denote the set of all affine homomorphisms $X \rightarrow Y$ by $\text{Aff}(X, Y)$. The set $\text{Aff}(X, V)$ of affine homomorphisms to the module V is naturally an R -module: for $g_1, g_2 \in \text{Hom}(U, V)$ and $c_1, c_2 : X \rightarrow V$, $r \in R$, define $(g_1, c_1) + (g_2, c_2) = (g_1 + g_2, c_1 + c_2)$, where, for $x \in X$, $(c_1 + c_2)(x) = c_1(x) + c_2(x)$, and define $r(g_1, c_1) = (rg_1, rc_1)$ where $(rc_1)(x) = r(c_1(x))$. It has an R -submodule C (the constant maps), consisting of pairs $(0, c)$ where c is a constant map $X \rightarrow V$, and the quotient module is isomorphic to $\text{Hom}(U, V)$. In particular, $\text{Aff}(U, V)$ is naturally isomorphic to $\text{Hom}(U, V) \oplus V$.

An R -torsor can be regarded as a pair (U, X) , where U is an R -module and U has a faithful transitive action on X . In this sense, we sometimes talk of interpretable R -torsors living in K^{eq} , but not necessarily as cosets of submodules of K^n . Observe that in the notation above, $\text{Aff}(X, Y)$ is in this sense a torsor of $\text{Aff}(X, V)$.

Definition 2.1.6 We will be using a uniformly definable family of torsors in K^n . For each natural number n , the set S_n consists of the R -sublattices of K^n , that is, the free R -submodules of K^n on n generators. (Formally, S_n consists of codes for lattices, chosen in a uniform way, but we often slur over this distinction.) The elements of S_1 are precisely the modules of the form γR , for $\gamma \in \Gamma$. In general, each element of S_n is definably R -isomorphic to R^n , but not canonically so. We write $\mathcal{S} = \bigcup_{n=1}^{\infty} S_n$. For any $s \in S_n$, we define $\text{red}(s) = s/\mathcal{M}s$ (the reduction of s modulo \mathcal{M}), where $\mathcal{M}s = \{ma : m \in \mathcal{M}, a \in s\}$. Then $\text{red}(s)$ is a set of torsors, and also is an n -dimensional vector space over k . For each n , let $T_n = \bigcup\{s/\mathcal{M}s : s \in S_n\}$ and $\mathcal{T} = \bigcup_{n=1}^{\infty} T_n$. Notice that T_1 contains all of the open balls in K of the form $B_{<|a|}(a)$. Let $\tau_n : T_n \rightarrow S_n$ be defined by $\tau_n(t) = s$ if and only if $t = a + \mathcal{M}s$ for some $a \in s$. We will often write τ for τ_n . Then for each n and for each $s \in S_n$, $\tau^{-1}(s) = s/\mathcal{M}s = \text{red}(s)$ is a definable subset of T_n . In our language (to be defined in Section 3.1) K, k, Γ are sorts, as is each S_n and each T_n . As noted in Section 2.4, each sort S_n can be regarded as a coset space of an \emptyset -definable group by an \emptyset -definable subgroup, and each T_n can be regarded as a finite union of coset spaces. We call these sorts the *geometric* sorts, and write $\mathcal{G} = K \cup \Gamma \cup k \cup \mathcal{S} \cup \mathcal{T}$. Formally, k and Γ are redundant.

The $\text{red}(s)$ notation is occasionally extended: for $a \in s$ we sometimes write $\text{red}(a)$ for $a + \mathcal{M}s \in \text{red}(s)$.

There is inevitably some confusion between definable subsets from K^n and elements of K^{eq} . In order to avoid the notation becoming too thick, we will try to maintain the following convention. Arbitrary torsors will be denoted by capital letters if thought of as sets, and by corresponding (if possible) lower-case letters if thought of as elements of K^{eq} . Each S_n and T_n is a sort in K^{eq} , and their elements will be denoted by small letters s, t . However, we will sometimes write, for example, $A \in S_n$ when we want to consider the module A for which $\ulcorner A \urcorner \in S_n$.

Lemma 2.1.7 *Let C be an algebraically closed valued field. Then any element s of S_n definable over C is C -definably isomorphic to R^n , and in particular, will contain a tuple from C . The torsor $\text{red}(s)$ is C -definably isomorphic to k^n .*

Proof. This holds automatically if C is a model of the theory of algebraically closed non-trivially valued fields, for then s has a free basis consisting of elements of s , and this can be mapped to the standard basis of R^n . Otherwise, by Remark 2.1.4, no element of $\Gamma \setminus \{0, 1\}$ is definable over C , so the only free R -submodule of K^n defined over C is R^n itself. The second part is clear, as $R/\mathcal{M}R$ is isomorphic to k . \square

2.2 Definable modules

In this section we develop the theory of the definable modules and torsors. In particular, we will show that a torsor in K^n is interdefinable with a module in K^{n+1} ; also that any definable R -submodule of K^n is, up to definable isomorphism, a direct sum of copies of K , R , and \mathcal{M} . We begin with a reminder of some standard valuation-theoretic terminology.

An extension field L of a valued field F is *immediate* if F and L have the same value group and residue field. We occasionally refer to *maximal* valued fields, that is, fields with no proper immediate extensions. If λ is a limit ordinal, then a sequence $(a_\alpha : \alpha < \lambda)$ is *pseudo-convergent* if, for all $\mu_1 < \mu_2 < \mu_3 < \lambda$, $|a_{\mu_1} - a_{\mu_2}| > |a_{\mu_2} - a_{\mu_3}|$. An element a of K is a *pseudo-limit* of the pseudo-convergent sequence $(a_\alpha : \alpha < \lambda)$ if $|a - a_\mu| = |a_{\mu+1} - a_\mu|$ for all $\mu < \lambda$. We recall the following theorem of Kaplansky [8].

Theorem 2.2.1 (Kaplansky) *Let (F, v) be a valued field. Then (F, v) is maximal if and only if every pseudo-convergent sequence in F has a pseudo-limit in F .*

The first two lemmas show that a definable homomorphism from either a proper submodule of K or from K itself to a quotient of modules is essentially linear.

Lemma 2.2.2 *Let V be a definable R -submodule of K , and $\beta \in \Gamma$. Then every definable homomorphism $h : \beta\mathcal{M} \rightarrow K/V$ has the form $h(x) = ax + V$ for some $a \in K$, so lifts to a definable homomorphism from βR to K/V .*

Proof. We may suppose that $\beta = 1$, so $\beta R = R$. Clearly, V has the form $\{0\}$, or K , or $\delta\mathcal{M}$, or δR , for some $\delta \in \Gamma$. We shall suppose that $V = \delta R$, the other cases being similar. Now $h(\mathcal{M})$ is a definable R -submodule of $K/\delta R$, so has the form $\varepsilon\mathcal{M}/\delta R$, $\varepsilon R/\delta R$, or $K/\delta R$. Since $\mathcal{M} = \mathcal{M}\mathcal{M}$, we have $h(\mathcal{M}) = \mathcal{M}h(\mathcal{M})$. It follows that if $h(\mathcal{M})$ is finitely generated then by Nakayama's Lemma $h(\mathcal{M}) = 0$, and the lemma is trivial. Thus, we may suppose that $h(\mathcal{M}) = K/\delta R$ or $h(\mathcal{M}) = \varepsilon\mathcal{M}/\delta R$ with $\delta < \varepsilon$ (in which case we may assume $\varepsilon = 1$). Either way, h is a surjection $\mathcal{M} \rightarrow K/\delta R$ or $\mathcal{M} \rightarrow \mathcal{M}/\delta R$. We must extend h to $h^* : R \rightarrow K/\delta R$ (or $h^* : R \rightarrow \mathcal{M}/\delta R$).

Choose a sequence $(x_\lambda : \lambda < \kappa)$ of elements of \mathcal{M} , indexed by a cardinal κ , with $\gamma_\lambda := |x_\lambda| \rightarrow 1$ as $\lambda \rightarrow \kappa$. We may suppose κ is the least cardinality of such a sequence, so is regular. For each $\lambda < \kappa$, choose y_λ such that $h(x_\lambda) = y_\lambda + \delta R$ and put $a_\lambda := y_\lambda x_\lambda^{-1}$. Let $x \in \gamma_\lambda R$. Then $x_\lambda^{-1}x \in R$ and, as $\gamma_\lambda R$ is a cyclic R -module and h is an R -module homomorphism, we have

$$h(x) = h(x_\lambda x_\lambda^{-1}x) = x_\lambda^{-1}xh(x_\lambda) = x_\lambda^{-1}x(a_\lambda x_\lambda + \delta R) = a_\lambda x + \delta R.$$

Hence, if $\mu < \lambda < \kappa$, then $a_\lambda x_\mu + \delta R = a_\mu x_\mu + \delta R$, so $|a_\lambda - a_\mu| \leq \delta\gamma_\mu^{-1}$ (use that x_μ lies in both $\gamma_\lambda R$ and $\gamma_\mu R$).

We may suppose that if $\lambda < \kappa$ then there is λ' with $\lambda < \lambda' < \kappa$ such that for all λ'' with $\lambda' < \lambda'' < \kappa$ we have $|a_\lambda - a_{\lambda''}| > \delta\gamma_{\lambda''}^{-1}$. For suppose this is false for some λ . Then we may define h^* by putting $h^*(x) = a_\lambda x + \delta R$. For if $x \in \mathcal{M}$, choose $\lambda' > \lambda$ such that $|a_\lambda - a_{\lambda'}| \leq \delta\gamma_{\lambda'}^{-1}$ and $\gamma_{\lambda'} > |x|$. Then $h(x) = a_{\lambda'}x + \delta R = h^*(x)$.

Now construct inductively a subsequence $(b_\lambda : \lambda < \kappa')$ of $(a_\lambda : \lambda < \kappa)$ such that the following hold:

- (i) $b_0 = a_0$ and for some strictly increasing function $f : \kappa' \rightarrow \kappa$, $b_\lambda = a_{f(\lambda)}$ for each $\lambda < \kappa'$;
- (ii) $(b_\lambda : \lambda < \kappa')$ is pseudo-convergent, that is, if $\lambda_1 < \lambda_2 < \lambda_3 < \kappa'$ then $|b_{\lambda_2} - b_{\lambda_1}| > |b_{\lambda_3} - b_{\lambda_2}|$;
- (iii) for all $\lambda_1 < \lambda_2 < \kappa'$ and λ_3 with $f(\lambda_2) < \lambda_3 < \kappa$, we have $|b_{\lambda_1} - b_{\lambda_2}| > |b_{\lambda_2} - a_{\lambda_3}|$.

Suppose that b_μ have been found for all $\mu < \lambda$, and put $\lambda^* = \sup\{f(\mu) : \mu < \lambda\}$. We may suppose $\lambda^* < \kappa$ (otherwise put $\kappa' := \lambda$). As κ is regular, by the last paragraph there is $\nu \in \kappa$ with $\lambda^* < \nu$ such that for all $\mu < \lambda$ and $\nu' \geq \nu$ we have $|b_\mu - a_{\nu'}| > \delta\gamma_{\nu'}^{-1}$. Put $f(\lambda) = \nu$, so $b_\lambda = a_\nu$. Now (ii) holds since (iii) held at the previous stage. Also, (iii) holds, for if $\mu < \lambda$ and $\nu' > \nu$, then

$$|b_\mu - b_\lambda| = |b_\mu - a_\nu| > \delta\gamma_\nu^{-1} \geq |a_{\nu'} - a_\nu| = |a_{\nu'} - b_\lambda|,$$

as required.

It is easy to check that $(b_\alpha : \alpha < \mu)$ is a pseudo-convergent sequence. Hence, by Theorem 2.2.1, there is an algebraically closed immediate extension K' of K such that the sequence $(b_\lambda : \lambda < \kappa')$ has pseudo-limit b' , that is, for all $\lambda < \kappa'$, $|b_{\lambda+1} - b_\lambda| = |b' - b_\lambda|$. If h' is the corresponding function in K' and \mathcal{M}' , R' are the

corresponding maximal ideal and valuation ring, we have $h'(x) = b'x + \delta R'$ for all $x \in \mathcal{M}'$. By Robinson's model-completeness for algebraically closed valued fields, there is $b \in K$ such that $h(x) = bx + \delta R$ for all $x \in \mathcal{M}$. Now define $h^* : R \rightarrow K/\delta R$ (or $h^* : R \rightarrow \mathcal{M}/\delta R$) by putting $h^*(x) = bx + \delta R$ for all $x \in R$. \square

Lemma 2.2.3 *Suppose that V is an R -submodule of K of the form $\alpha\mathcal{M}$ or αR where $\alpha \neq 0$. Then every definable R -homomorphism $f : K \rightarrow K/V$ has the form $f(x) = ax + V$ for some $a \in K$.*

Proof. This is essentially the same as the proof of Lemma 2.2.2, except that the sequence $(x_\alpha : \alpha < \nu)$ satisfies $|x_\alpha| \rightarrow \infty$ as $\alpha \rightarrow \nu$. \square

Lemma 2.2.4 *Let V be a definable R -submodule of K^n . Then V is definably isomorphic to a direct sum of at most n R -modules, each of the form R , \mathcal{M} , or K .*

Proof. We shall show, by induction on n , that there is $g \in GL_n(K)$ such that $g(V) = \bigoplus_{i=1}^n V_i$, where each V_i is of the form $\{0\}$, R , \mathcal{M} , or K . Clearly the induction starts, as any definable R -submodule of K , after multiplying by an element of K , has the required form.

Let $\pi : K^n \rightarrow K$ be the projection onto the first coordinate, and write $V' \subseteq K^{n-1}$ for the R -submodule such that $\{0\} \times V' = \ker(\pi) \cap V$. By induction, there is $g' \in GL_{n-1}(K)$ with $g'(V') := \bigoplus_{i=2}^n V_i$, where each V_i has the form $\{0\}$, \mathcal{M} , R or K . Let $(a_{ij})_{2 \leq i, j \leq n}$ be the matrix for g' (with the standard basis, and the matrix written on the left), and write $A = (a_{ij})_{1 \leq i, j \leq n}$ for the $n \times n$ matrix whose first row and column are all zeroes. We shall define the $n \times n$ matrix $B = (b_{ij})_{1 \leq i, j \leq n}$ for g . In all cases, $b_{1j} = 0$ for $2 \leq j \leq n$ and $b_{ij} = a_{ij}$ for $2 \leq i, j \leq n$. Let $J := \pi(V)$. If $J = \{0\}$, let $V_1 = \{0\}$, $b_{11} = 1$ and $b_{i1} = 0$ for $2 \leq i \leq n$. Then $g(V) = V_1 + g'(V')$.

If $J \neq \{0\}$ then $J \in \{K, \alpha R, \alpha\mathcal{M}\}$ for some $\alpha \in \Gamma$. For each $i = 2, \dots, n$ define an R -homomorphism $\varphi_i : J \rightarrow K/V_i$ as follows: for $x \in J$, $\varphi_i(x) = (Av)_i + V_i$ for any $v \in V$ with $\pi(v) = x$. Since the difference of any two such vectors is in $\{0\} \times V'$, φ_i is well-defined. By Lemmas 2.2.2 and 2.2.3, for each i there is $a_i \in K$ such that $\varphi_i(x) = a_i x + V_i$ for all $x \in J$. Now let $V_1 = \alpha^{-1}J$ (take $\alpha = 1$ if $J = K$), $b_{11} = a^{-1}$ for any $a \in K$ with $|a| = \alpha$, and $b_{i1} = -a_i$ for $i = 2, \dots, n$. Then for any $v \in V$ and $2 \leq i \leq n$, $(Bv)_i = -a_i v_1 + (Av)_i \in V_i$, so $Bv \in \bigoplus_{i=1}^n V_i$ as required. \square

Lemma 2.2.5 *Let V be a definable R -submodule of K^n , and $\beta \in \Gamma$. Then every definable homomorphism $h : \beta\mathcal{M} \rightarrow K^n/V$ lifts to a definable homomorphism βR to K^n/V .*

Proof. By the proof of Lemma 2.2.4 there is $g \in \text{GL}_n(K)$ such that $g(V) \cong \bigoplus_{i=1}^n V_i$, where $V_i = \pi_i(g(V))$ (the projection to the i^{th} coordinate). Thus, g induces an isomorphism $g' : K^n/V \rightarrow \bigoplus_{i=1}^n K/V_i$. This gives a homomorphism $h^* : \beta M \rightarrow \bigoplus_{i=1}^n K/V_i$, which by Lemma 2.2.2 extends to a homomorphism $h' : \beta R \rightarrow \bigoplus K/V_i$. Now apply g'^{-1} . \square

Part (ii) of the next lemma enables us to replace torsors by modules in certain coding arguments.

Lemma 2.2.6 (i) *Let L be a definable R -submodule of K^n . Then there is a definable subtorsor L' of K^{n-1} , a definable R -submodule T of K^{n-1} , and some $\gamma \in \Gamma$, such that $\ulcorner L \urcorner$ is interdefinable over \emptyset with the triple $(\ulcorner L' \urcorner, \ulcorner T \urcorner, \gamma)$.*

(ii) *Let L' be a subtorsor of K^{n-1} . Then there is a R -submodule L of K^n such that $\ulcorner L' \urcorner = \ulcorner L \urcorner$.*

Proof. (i) Let $A := \pi_1(L)$, where $\pi_1 : K^n \rightarrow K$ is projection to the first coordinate, and suppose $\ker(\pi_1) = \{0\} \times T$. Put $B := K^{n-1}/T$. Then L can be regarded as the graph of a homomorphism $h : A \rightarrow B$. We may suppose $A \neq 0$ (otherwise put $L' := T$); so $A = K$, or $A = \gamma R$ or $A = \gamma \mathcal{M}$ for some $\gamma \in \Gamma$.

First suppose that $A = \gamma R$ or $A = \gamma \mathcal{M}$. By Lemma 2.2.5, the restriction map $\text{Hom}(\gamma R, B) \rightarrow \text{Hom}(\gamma \mathcal{M}, B)$ is surjective. Furthermore, since γR is a free R -module, the map $\text{Hom}(\gamma R, K^{n-1}) \rightarrow \text{Hom}(\gamma R, B)$ (obtained by composing each element of $\text{Hom}(\gamma R, K^{n-1})$ with the natural map $K^{n-1} \rightarrow B$) is surjective. Thus, we obtain by composition a surjection $\text{Hom}(\gamma R, K^{n-1}) \rightarrow \text{Hom}(\gamma \mathcal{M}, B)$, and hence, for $A \in \{\gamma R, \gamma \mathcal{M}\}$, we obtain a $(\gamma, \ulcorner T \urcorner)$ -definable surjection $\rho : \text{Hom}(\gamma R, K^{n-1}) \rightarrow \text{Hom}(A, B)$. Let $V := \ker(\rho)$. Since any R -homomorphism $\gamma R \rightarrow K$ is given by multiplication by some uniquely determined $a \in K$, $\text{Hom}(\gamma R, K^{n-1})$ is canonically (over $\gamma, \ulcorner T \urcorner$) R -isomorphic to K^{n-1} , and V to some corresponding submodule V' of K^{n-1} .

We have $\text{Hom}(A, B) \cong \text{Hom}(\gamma R, K^{n-1})/V$. The element h of $\text{Hom}(A, B)$ corresponds to a coset of V in $\text{Hom}(\gamma R, K^{n-1})$, so corresponds to a definable subtorsor L' of $K^{n-1} \cong \text{Hom}(\gamma R, K^{n-1})$, namely, a coset of V' . Now $\ulcorner h \urcorner$ and hence $\ulcorner L \urcorner$ is interdefinable with the triple $(\ulcorner L' \urcorner, \ulcorner T \urcorner, \gamma)$.

The remaining case of the claim is when $A = K$. Again, by Lemmas 2.2.4 and 2.2.3, the natural map $\tau : \text{Hom}(A, K^{n-1}) \rightarrow \text{Hom}(A, B)$ is surjective, so again h is interdefinable with a subtorsor of K^{n-1} . In this case (i) holds with $\gamma = 0$.

(ii) Now $\ulcorner L' \urcorner$ is interdefinable with a code for the subtorsor $L'' := \{1\} \times L'$ of K^n . Let L be the R -submodule of K^n generated by L'' . Since $L'' = L \cap (\{1\} \times K^{n-1})$, we have $\text{dcl}(\ulcorner L \urcorner) = \text{dcl}(\ulcorner L' \urcorner)$. \square

We give an application of Lemma 2.2.4, used in the next section.

Lemma 2.2.7 *Let A be an R -lattice in K^n , let $1 \leq m \leq n-1$, and let $\pi : K^n \rightarrow K^m$ be a projection to the first m coordinates. Then $\pi(A)$ is an R -lattice in K^m .*

Proof. Since A is finitely generated, $\pi(A)$ is finitely generated. By Lemma 2.2.4, $\pi(A)$ is a direct sum of copies of K, R, \mathcal{M} , and so by finite generation, $\pi(A) \cong R^\ell$ for some $\ell \leq m$. Also, the R -module K^n/A is torsion, so $K^m/\pi(A)$ is also torsion. It follows that $\ell = m$, as required. For otherwise, we could complete an R -basis for $\pi(A)$ to a K -basis for K^{n-1} , and the added basis vectors would generate a free R -module, modulo $\pi(A)$. \square

2.3 Unary sets

In our original approach to valued fields, with balls as the basic sorts, we found that we often needed to consider the type of a single imaginary, say $B_\gamma(a)$, as really the type of the pair $(\gamma, B_\gamma(a))$. This led to a dissonance between 1-types and general n -types. To resolve this, we define the unary sets, which will play the role of 1-types. We show in this section that any element of \mathcal{G} can be coded by a sequence in which each element lies in a unary set defined over the previous elements. We will talk of *unary types* as the type of an element of a unary set, the underlying unary set fixed by the context.

A *definable 1-module* is an R -module (living in K^{eq}) which is definably isomorphic to a quotient of one definable R -submodule of K by another. It will be definably isomorphic to one of $\gamma R/\delta R, \gamma R/\delta \mathcal{M}, \gamma \mathcal{M}/\delta R$ or $\gamma \mathcal{M}/\delta \mathcal{M}$, or to $K/\delta R$ or $K/\delta \mathcal{M}$, where $\gamma, \delta \in \Gamma$ with $0 \leq \delta \leq \gamma$ (and in fact we may always assume $\gamma = 1$). In the case when the 1-module, A say, is definably isomorphic to K (that is, to $K/0R$), we actually assume that the 1-module comes equipped with a definable submodule B , and that the definable isomorphism $A \rightarrow K$ maps B to R ; without this it would not be clear below how to define the *radius* of a submodule of A . By allowing $\delta = 0$ we include balls containing 0 as 1-modules. A *definable 1-torsor* is a definable torsor of a definable 1-module. An ∞ -*definable 1-torsor* is an intersection of a chain of definable 1-torsors. A *1-torsor* is a definable or ∞ -definable 1-torsor. If C is a set of parameters, then a *C -1-torsor* is a definable or ∞ -definable 1-torsor for which the parameters come from C ; we do not here require that there be any C -definable isomorphism with, say, $\gamma R/\delta R$.

We will say that a 1-torsor is *closed* if it is definably isomorphic to a torsor of some $\gamma R, \gamma R/\delta R$ or $\gamma R/\delta \mathcal{M}$; it is *open* if it is definably isomorphic to a torsor of a module which is a quotient of \mathcal{M} (and we also regard modules definably isomorphic to quotients of K as open). Notice that if $\gamma < |a|$ then a closed ball $s = B_{\leq \gamma}(a)$ is a closed 1-torsor of the 1-module $\gamma R/0R$. But s is also an element of the 1-module $\gamma' R/\gamma R$, where $|a| = \gamma'$. In Section 2.1, we wrote $\text{red}(s) = s/\mathcal{M}s$ (when $s \in S_1$) for the set of open balls in s of the same radius as that of s . In the same way, if T is a closed 1-torsor of the 1-module A we will write $\text{red}(T)$ for the set of all open 1-torsors of $\mathcal{M}A$ contained in T . Then $\text{red}(T)$ is also a closed 1-torsor (it is a torsor of $A/\mathcal{M}A$). As in the ball case, $\text{red}(T)$ is definably isomorphic to k , hence is strongly minimal. In particular, we have a notion of *generic* for elements of $\text{red}(T)$.

If the elements of a 1-torsor U are subsets of K , that is, U is $\gamma R/\delta R$ or $\gamma R/\delta \mathcal{M}$, etc, then we say that U is a *true* 1-torsor. More generally, if U is a definable C -1-torsor and C is a model, then U will be C -definably isomorphic to a true 1-torsor. Suppose U is a 1-torsor of the 1-module A . We can define the *radius* of definable subtorsors V of U as follows. Suppose first A is closed. By definition, V is a torsor of a definable submodule B of A , and for some unique γ , $B = \gamma RA$ or $\gamma \mathcal{M}A$. Then $\text{rad}(V) := \gamma$. If A is open (but not definably isomorphic to a quotient of K), then a definable submodule has radius γ if it has the form γRA or $\bigcap(\delta RA : \delta > \gamma)$. If U is an intersection of definable 1-torsors $\{U_i : i \in I\}$, we fix any $i_0 \in I$ and define the radius of a subtorsor V of U to be its radius with respect to the fixed U_{i_0} (this ensures that the radius of a definable subtorsor of an ∞ -definable 1-torsor lies in Γ rather than its Dedekind completion). The definition of radius for a subtorsor of a torsor arising from a quotient of K is clear: if A is definably isomorphic to $K/\delta R$ for $\delta > 0$, then a definable submodule D has radius γ if γ is greatest such that $\gamma RD = \{0\}$; if A is definably isomorphic to $K/\delta \mathcal{M}$, then D has radius δ where δ is greatest such that γRD is isomorphic to $\{0\}$ or k ; and if the pair (A, B) is definably isomorphic via an isomorphism f to (K, R) , then the radius of D is exactly the radius of $f(D)$ as a submodule of K (this does not depend on the choice of f).

In all the above cases, if V is a subtorsor of U and $\text{rad}(U) > \gamma' > \text{rad}(V)$, then $B_{\leq \gamma'}(V)$ denotes the closed subtorsor of U of radius γ' containing V , and $B_{< \gamma'}(V)$ the open subtorsor of U containing V ; these are uniquely determined. Also, we can define $|a - b|$ for $a, b \in U$. For $a - b \in A$, hence $a - b$ generates a submodule $(a - b)R$, and $|a - b|$ is the radius of this submodule.

Definition 2.3.1 A *unary set* is a 1-torsor or an interval $[0, \alpha)$ in Γ , where $\alpha \in \Gamma \cup \{\infty\}$. A *C -unary set* is a unary set (possibly ∞ -definable) where the parameters may be chosen from C . A *unary type* over C is the type of an element of a C -unary set.

Remark 2.3.2 Below and in Section 2.5, when considering a C -1-torsor U we frequently assume that the base set of parameters C is algebraically closed in K^{eq} . However, all that is really needed is that any $\text{acl}(C)$ -definable subtorsor of U is C -definable. In Section 3, we will be considering a restricted class \mathcal{U} whose definable subtorsors are coded in the geometric sorts \mathcal{G} . We will then be able to apply all results of Sections 2.3 and 2.5 under the weaker assumption $C = \text{acl}(C) \cap \mathcal{G}$.

If U is a 1-torsor, the notions of *Swiss cheese* and *trivially nested* set of Swiss cheeses from Section 2.1 carry through to definable subsets of U : a Swiss cheese of U is a non-empty set $t \setminus (t_1 \cup \dots \cup t_n)$ where t and t_i are definable subtorsors of U .

Lemma 2.3.3 *Let $C \subset K^{\text{eq}}$ and U be a C -1-torsor.*

(i) *Let X be a definable subset of U . Then X is uniquely expressible as a finite union of Swiss cheeses, no two trivially nested.*

(ii) *Assume $C = \text{acl}(C)$. If $a, b \in U$ and neither of a, b lie in a C -definable proper subtorsor of U , then $a \equiv_C b$. In particular, if there is no C -definable proper subtorsor of U then all elements of U have the same type over C .*

Proof. (i) By expanding C to some C' , we may assume that U is a true 1-torsor. Existence and uniqueness now follow from Theorem 2.1.2.

(ii) If $a \not\equiv_C b$, then by (i), some C -definable Swiss cheese $t \setminus (t_1 \cup \dots \cup t_n)$ contains just one of a, b , and the uniqueness assertion in (i) ensures we may assume t and the t_i are C -definable. At least one of them must be a proper subtorsor of U . \square

Definition 2.3.4 Let $C \subset K^{\text{eq}}$ be a set of parameters. Let U be an $\text{acl}(C)$ -unary set and $a \in U$. Then a is *generic in U over C* if a lies in no $\text{acl}(C)$ -unary proper subset of U .

Remark 2.3.5 (i) By Lemma 2.3.3, if a, b are generic in a unary set over C , then $a \equiv_{\text{acl}(C)} b$. Thus, we may talk of *the generic type of U (over C)* as the type of an element of U which is generic over C . Existence of generic types is by compactness — one has to check that a 1-torsor is not the union of finitely many proper subtorsors.

(ii) If T is a closed 1-torsor then the above notion of genericity for the strongly minimal 1-torsor $\text{red}(T)$ agrees with that from stability theory. Also, suppose T is a C -definable closed 1-torsor, and a is generic in $\text{red}(T)$ over C . Then all elements of a have the same type over C ; for otherwise, some C -definable subset of T intersects infinitely many elements of $\text{red}(T)$ in a proper non-empty subset, contradicting Lemma 2.3.3.

(iii) We adapt slightly the above language, by saying that if $\delta \in \Gamma$, then γ is *generic over C below δ* if for any $\varepsilon \in \Gamma(C)$, if $\varepsilon < \delta$ then $\varepsilon < \gamma$. That is, γ is generic in the unary set $[0, \delta)$.

Lemma 2.3.6 *Suppose $C = \text{acl}(C) \subset K^{\text{eq}}$, and a is an element of a C -unary set U . Then a realises the generic type over C of a unique unary subset of U .*

Proof. We may assume $a \notin C$. Let $\{U_i : i \in I\}$ be the set of infinite C -definable unary subsets of U containing a . This set is clearly totally ordered by inclusion, and a realises the generic type over C of the intersection. \square

Remark 2.3.7 It follows in particular that if $C = \text{acl}(C)$ then any type over C of a field element or ball (of radius in C) is the generic type over C of a unary set.

Lemma 2.3.8 *Let C be any set of parameters in K^{eq} .*

(i) *If p is the generic type over C of a C -definable unary set, then p is definable (over C).*

(ii) *Let $\{U_i : i \in I\}$ be a descending sequence of C -definable subtorsors of some C -1-torsor U , with no least element, and let p be the generic type over K of field elements of $\bigcap(U_i : i \in I)$. Then p is not definable.*

Proof. (i) We assume the unary set is a 1-torsor, as the proof is similar for subsets of Γ . Suppose p is the generic type of the closed 1-torsor U . Then by Theorem 2.1.2, for any formula $\varphi(x, \bar{y})$ there is a natural number N_φ so that for any \bar{c} , $\varphi(x, \bar{c}) \in p$ if and only if $\varphi(x, \bar{c})$ holds on all elements of each torsor in an infinite subset of $\text{red}(U)$, if and only if $\varphi(x, \bar{c})$ holds on all elements of all except at most N_φ torsors in $\text{red}(U)$. This gives the definition of p .

If p is the generic type of an open 1-torsor U , then $\varphi(x, \bar{c}) \in p$ if and only if there is a proper definable sub-torsor U' of U such that $\varphi(x, \bar{c})$ holds throughout $U \setminus U'$. Note here that the collection of definable subtorsors of U is a uniformly definable family.

(ii) This is a special case of the following general fact. Let M be a large sufficiently saturated model of some theory, let \mathcal{F} be a uniformly definable family of definable subsets of M^n (such as the collection of subtorsors of a 1-torsor), and let $(V_i : i \in I)$ be a decreasing sequence (totally ordered by inclusion) of elements of \mathcal{F} with no least element, with $|I|$ small relative to $|M|$. Then if p is a definable type with solution set P in K , it cannot happen that for all $V \in \mathcal{F}$, $V \supset P \Leftrightarrow V$ contains some V_i : indeed, otherwise the partial order consisting of members of \mathcal{F} containing P would be definable and of small infinite cofinality, contrary to saturation of M . \square

Recall from Definition 2.1.6 the notation \mathcal{S}, \mathcal{T} . The following notion of a unary code for an imaginary is the essential idea that allows us to think of elements of $\mathcal{S} \cup \mathcal{T}$ as sequences of elements of unary sets. This enables us frequently to apply results about unary sets to elements of the geometric sorts \mathcal{S} and \mathcal{T} . In [3], we use it to extend the notion of *generic* to a ‘sequential independence’ for n -types.

Definition 2.3.9 Let e be an element of K^{eq} . A sequence (a_1, \dots, a_m) of elements of K^{eq} is a *unary code* for e if $\text{dcl}(e) = \text{dcl}(a_1, \dots, a_m)$, and for each $i = 1, \dots, m$, a_i is an element of a unary set defined over $\text{dcl}(a_j : j < i)$.

Proposition 2.3.10 *Let $s \in \mathcal{G}$. Then s has a unary code whose elements lie in \mathcal{G} .*

Proof. The only cases needing proof are when $s \in S_n \cup T_n$, for some n . We exhibit a canonical filtration of lattices, and a corresponding filtration of elements of T_n . For each $1 \leq i < n$ let π_i be the projection of K^n to the first i coordinates, and π^i the projection to the last $n - i$ coordinates.

Step 1. Suppose first $s \in S_n$, so s codes an R -lattice A from K^n . Let $A_i := \ker(\pi^i|_A)$ (so $A_i = A \cap (K^i \times \{0\}^{n-i})$). Now A/A_{n-1} is isomorphic to $\pi^{n-1}(A)$, which is isomorphic to R by Lemma 2.2.7. Hence, the sequence $0 \rightarrow A_{n-1} \rightarrow A_n \rightarrow R \rightarrow 0$ must split, and as A is isomorphic to R^n , so A_{n-1} is isomorphic to R^{n-1} . Continuing this way, we see that each A_i is isomorphic to R^i , and each quotient A_{i+1}/A_i is isomorphic to R . If we write A'_i for the R -module in K^i with $A_i = A'_i \times \{0\}^{n-i}$, then $\ulcorner A'_i \urcorner$ is in S_i .

Step 2. To obtain a unary code for a lattice A in K^n , we reduce by Step 1 and induction on n to the following. Let $B_{n-1} = \pi_{n-1}(A)$. Then $A'_{n-1} \leq B_{n-1}$. By Step 1, A'_{n-1} is a lattice, and by Lemma 2.2.7, B_{n-1} is also a lattice. Thus, they both have codes in S_{n-1} so by induction have unary codes say c_{n-1}, b_{n-1} from \mathcal{G} . Let $B_1 := \pi^{n-1}(A)$ and $\{0\}^{n-1} \times A''_1 = \ker(\pi_{n-1})$. By Lemma 2.2.7, B_1 is a lattice. Also, as in Step 1, the sequence $0 \rightarrow \ker(\pi_{n-1}) \rightarrow A_n \rightarrow B_{n-1} \rightarrow 0$ splits (as B_{n-1} is free), so A''_1 is also a lattice. Thus $A''_1 \leq B_1$ are both coded in S_1 , so by induction have unary codes c_1, b_1 , say. We claim that $(c_1, b_1, c_{n-1}, b_{n-1}, s)$ is the unary code for s . To show the claim, we need to verify that s is in a $(c_1, b_1, c_{n-1}, b_{n-1})$ -definable unary set. So let $Y(c_1, b_1, c_{n-1}, b_{n-1})$ be the set of codes of lattices C of $K^n = K^{n-1} \times K$ such that $C \cap (K^{n-1} \times \{0\}) = A'_{n-1}$, $\pi_{n-1}(C) = B_{n-1}$, $C \cap (\{0\}^{n-1} \times K) = A''_1$ and $\pi^{n-1}(C) = B_1$. Then $s \in Y(c_1, b_1, c_{n-1}, b_{n-1})$, so we need to show that Y is contained in a unary set.

We claim that there is a $(c_1, b_1, c_{n-1}, b_{n-1})$ -definable R -module D , isomorphic to some module $R/\alpha R$, and a canonical identification of Y with a subset $D' := D \setminus \mathcal{M}D$. The module D is $\text{Hom}_R(B_{n-1}/A'_{n-1}, B_1/A''_1)$, and the subset D' consists of the invertible homomorphisms. The identification takes $f \in D$ to $\{(x, y) \in B_{n-1} \times B_1 : f(x + A'_{n-1}) = y + A''_1\}$. Since $Y(c_1, b_1, c_{n-1}, b_{n-1}) \neq \emptyset$, B_{n-1}/A'_{n-1} and B_1/A''_1 are isomorphic, and each is isomorphic to an R -module $R' := R/\alpha R$ for some $\alpha < 1$ from Γ , i.e. to free R' -modules on one generator. Now $D \cong \text{Hom}_R(R', R') = \text{Hom}_{R'}(R', R') \cong R'$, so is a 1-module, and hence $Y(c_1, b_1, c_{n-1}, b_{n-1})$ is a subset of a 1-module.

Step 3. Finally, we exhibit a unary code for elements $a \in T_n$. Let $V = A/\mathcal{M}A \subset T_n$. With A_i as above, let $V_i := A_i/\mathcal{M}A_i$, to obtain a corresponding filtration $0 \leq V_1 \leq \dots \leq V_n = V$ of k -vector spaces. Now if $a \in V$, we have a sequence $s = \tau(a), a + V_{n-1}, a + V_{n-2}, \dots, a$. Each element $a + V_i$ lies in a torsor of a 1-dimensional k -vector space V_{i+1}/V_i defined over the previous element. Furthermore, $a + V_i \in \text{red}(A/A_i) \cong \text{red}(\pi^i(A))$. Thus, $\ulcorner a + V_i \urcorner$ is an element of T_{n-i} . It follows that if, in the above sequence, s is replaced by a unary code for A , then we have a unary code for V . \square

Note that in the above lemma, the 1-torsors required are all either strongly minimal 1-torsors, or 1-modules.

We close this section with some remarks about invariant types. Let $C = \text{acl}(C)$ and p be a type over C . An *invariant* extension of p is a type p^* over K extending p such that $\text{Aut}(K/C)$, in its action on the set of types over K , fixes

p^* . If p has an invariant extension p^* and $C \subset C' \subset K$ then $p^*|C'$ denotes the restriction of p^* to C' . In general, one would not expect there to be a unique invariant extension of a type p . However, in Section 2.5, we will show that any unary type has a canonical invariant extension (given by the generic type over any parameter set); hence we can just write $p|C'$ for $p^*|C'$, the generic extension of p over C' . In [3] we will extend this to n -types. In Remark 2.11 of [7] it is claimed that invariant extensions of types exist for arbitrary C -minimal structures. However the Remark rests on Lemma 2.2 of that paper, and the proof of Lemma 2.2 is incomplete and the result may well be incorrect (other applications of 2.2 in [7] seem to be unaffected, as problems only arise when $\text{acl} \neq \text{dcl}$).

2.4 Definable functions from Γ

In this section we study definable functions from the value group into \mathcal{G} , and show that they are fairly simple. In order eventually to obtain the technical Lemma 3.4.12, we actually work in the more general setting of a function from a finite cover of Γ , in the following sense.

Definition 2.4.1 A definable surjection $\rho : X \rightarrow Y$ between definable sets X and Y is a *finite cover* of Y if all the fibres $\rho^{-1}(a)$ ($a \in Y$) are finite. We often just refer to ‘the finite cover ρY ’, meaning the triple (ρ, X, Y) .

First, recall that any o-minimal expansion of an ordered abelian group (with at least one \emptyset -definable non-zero element) has definable Skolem functions. For given an interval I other than $\{x : 0 < x\}$ or $\{x : x < 0\}$, one of $(\inf(I) + \sup(I))/2$, $\inf(I) + |\inf(I)|/2$, $\sup(I) - |\sup(I)|/2$, or 0 , lies in I , and is definable from parameters used to define I ; the remaining intervals contain an \emptyset -definable element. In the next few lemmas, we work in a fixed 1-torsor U . So we may talk of the *radius* of subtorsors of U (possibly with respect to some fixed definable $U_i \supset U$), as defined at the beginning of Section 2.3.

Lemma 2.4.2 *Let U be a 1-torsor and \mathcal{Y} be the set of subtorsors of U . Let $E \subset \mathcal{Y} \times \Gamma$ be definable, and suppose that the second projection $p_2 : E \rightarrow \Gamma$ is finite-to-one. Then $p_1(E)$, the projection of E to \mathcal{Y} , contains only finitely many elements of \mathcal{Y} of any given radius.*

Proof. The hypothesis on E persists to definable subsets, so we may suppose $p_1(E)$ consists of subtorsors of equal radius δ , and must show $p_1(E)$ is finite. So suppose $p_1(E)$ is infinite. We may suppose that $p_1(E)$ is contained in the set of closed subtorsors in \mathcal{Y} , or that $p_1(E)$ is contained in the set of open subtorsors in \mathcal{Y} . By Theorem 2.1.2 applied to the union in K of the subtorsors in $p_1(E)$, there is a closed 1-torsor $t \in \mathcal{Y}$ of radius $\gamma > \delta$ (or $\gamma \geq \delta$ if the elements of $p_1(E)$ are open) whose sub-torsors of radius δ of the appropriate type all lie in

$p_1(E)$. Let $V := \{t' \in p_1(E) : t' \subset t\}$, and let \sim be the equivalence relation on V , two balls equivalent if they lie in the same element of $\text{red}(t)$. Then V/\sim is in definable bijection with $\text{red}(t)$ so is strongly minimal (and infinite). Furthermore, using definable Skolem functions (over an extra constant) we can find a definable finite-to-one map $f : V/\sim \rightarrow \Gamma$: define f so that if $v \in V/\sim$ then for some representative t' of v , $(t', f(v)) \in E$. This is a contradiction, since there cannot be a definable finite-to-one map from a strongly minimal set to a totally ordered set. \square

Lemma 2.4.3 *Let $(t_i : i \in I)$ be a definable chain of subtorsors of a 1-torsor U , the chain totally ordered by inclusion. Then there is $t \in U$ with $t \in t_i$ for every $i \in I$.*

Proof. By adding parameters, we may suppose U is a true 1-torsor. Let K' be a maximal algebraically closed immediate extension of K , and let $(t_j : j \in J)$ be the chain of subtorsors defined by the same formula in K' . For each $j \in J$, choose $a_j \in K'$ with $a_j \in t_j \setminus \bigcup(t_k : k < j)$. Then any well-ordered cofinal subsequence of $(a_j : j \in J)$ is pseudo-convergent, so (by Theorem 2.2.1) has a pseudo-limit $a' \in K'$. Then a' lies in t_j for each $j \in J$. As $K \prec K'$, there is a corresponding $a \in K$. Let t be the element of U containing a . \square

Proposition 2.4.4 *Let $B \subset K^{\text{eq}}$ be a set of parameters, U a B -1-torsor, $\alpha, \gamma \in \Gamma$, and t be a subtorsor of U of radius γ (possibly 0) with $t \in \text{acl}(B\alpha) \setminus \text{acl}(B)$. Then $\gamma \in \text{dcl}(B\alpha)$ and there is an $s \in \text{acl}(B)$ (a subtorsor of U) with $\text{rad}(s) < \gamma$, such that $t \in \{B_{\leq \gamma}(s), B_{< \gamma}(s)\}$.*

Proof. Let \mathcal{U} be the set of definable subtorsors of U . There is a B -definable set $E \subset \mathcal{U} \times \Gamma$ containing (t, α) such that the projection p_2 to the second coordinate is finite-to-one. Put $D := p_1(E)$. It follows from Lemma 2.4.2 that for any $\delta \in \Gamma$, $D^{\text{cl}}(\delta) := \{B_{\leq \delta}(t') : t' \in D\}$ is finite (we apply the lemma to $\{(B_{\leq \delta}(t'), \gamma) : (t', \gamma) \in E, \text{rad}(t') \leq \delta\}$, for any fixed δ). By compactness and saturation, there is a natural number m such that $|D^{\text{cl}}(\delta)| \leq m$ for each $\delta \in \Gamma$. Likewise, for any δ , $D^{\text{op}}(\delta) := \{B_{< \delta}(t') : t' \in D\}$ is finite.

We claim that there is a positive integer ℓ such that any subset of D which is pairwise incomparable under inclusion has size at most ℓ . For if not, then by saturation there is an infinite antichain $(t_i : i \in I)$ under inclusion, and as each $D(\delta)$ is finite, we may suppose that if $i < j$ then $\text{rad}(t_i) > \text{rad}(t_j)$. Put $\delta_i := \text{rad}(t_i)$ for each $i \in I$. For $\{i, j, k\} \subset I$ with $i < j < k$, colour $\{i, j, k\}$ red if $B_{\leq \delta_i}(t_j) = B_{\leq \delta_i}(t_k)$, and green otherwise. By Ramsey's Theorem and the last paragraph, we may suppose that $I = \omega + 1$ and all triples are red. Let $a \in t_\omega$. For each $i < \omega$, as t_i and t_ω are disjoint, $|x - a|$ takes fixed value, γ_i say, as x ranges through t_i . Hence, as each $D(\delta)$ is finite, if $X := \{x \in U : \exists t' \in D : x \in t' \wedge a \notin t'\}$, then X meets $B_{\leq \gamma_i}(a) \setminus B_{< \gamma_i}(a)$ in a proper subset; for $t_i \subset X$, but

X meets just finitely many elements of $\text{red}(B_{\leq \gamma_i}(a))$, and $D^{\text{op}}(\gamma_i)$ is finite. Thus, X is not a finite union of Swiss cheeses, contrary to Lemma 2.3.3(i).

The semilinearly ordered set D is the union of finitely many chains of subtorsors $\mathcal{C}_1, \dots, \mathcal{C}_\ell$, ordered by inclusion; to see this, choose a maximal antichain $\{t_1, \dots, t_\ell\}$ in D of size ℓ , and let C_i be a maximal chain containing t_i . By removing those elements of D which do not lie in an antichain of size ℓ (a B -definable set) we may suppose (by induction on ℓ) that $\mathcal{C}_1, \dots, \mathcal{C}_\ell$ are disjoint. The relation of non-disjointness is therefore an equivalence relation on D , whose classes are $\mathcal{C}_1, \dots, \mathcal{C}_\ell$; thus each C_i is in $\text{acl}(B)$. By Lemma 2.4.3, for each $i = 1, \dots, \ell$ there is $a_i \in K$ such that $C_i = \{u \in D : a_i \in u\}$. Suppose $t \in \mathcal{C}_1$, and put $s := \bigcap \{u : u \in \mathcal{C}_1\}$. Then a_1 lies in each $u \in \mathcal{C}_1$ and s is definable, $s \in \mathcal{U}$ by the existence part of Lemma 2.3.3(i). Also, $s \in \text{acl}(B)$, and $t = B_{\leq \gamma}(s)$ or $t = B_{< \gamma}(s)$. \square

Corollary 2.4.5 *Let $B \subset K^{\text{eq}}$ be a set of parameters and U be a B -1-torsor with no proper $\text{acl}(B)$ -definable subset. Suppose T, T' are subtorsors of U , both closed or both open, of radii δ, δ' respectively and $\delta \equiv_B \delta'$. Then*

- (i) $T \equiv_B T'$
- (ii) all elements of T have the same type over $B^\Gamma T^\neg$.

Proof. (i) We show that all subtorsors of U of radius δ have the same type over C . If $\delta \neq \delta'$ and the type of torsors over δ' is different from the type over δ , then $\delta \not\equiv_B \delta'$.

Consider the set V of all closed subtorsors of U of radius δ (the open case is similar). This is a $B\delta$ -definable 1-torsor. Suppose T is not generic in V . Then there is an $\text{acl}(B\delta)$ -definable subtorsor S of V containing T . By Proposition 2.4.4, there is a proper $\text{acl}(B)$ -definable subtorsor S' of S . Then $\bigcup S'$ is an $\text{acl}(B)$ -definable subset of U , contrary to hypothesis. So T and likewise T' are generic in V . By Remark 2.3.5, T and T' have the same type over $B\delta$, and hence over B .

(ii) Suppose $u \in T$ is not generic in T over $B^\Gamma T^\neg$. Then there is an $\text{acl}(B^\Gamma T^\neg)$ -definable subtorsor V_T of T . Consider the set $\bigcup \{V_S : S \equiv_{B\delta} T \wedge V_S \subset S\}$. This is a definable subset of K which is not a finite union of Swiss cheeses, contrary to Theorem 2.1.2. \square

Corollary 2.4.6 *Let $B = \text{acl}(B)$ be a set of parameters, U a 1-torsor over B , \mathcal{U} the set of subtorsors of U , $\rho\Gamma$ a B -definable finite cover of Γ , and $f : \rho\Gamma \rightarrow \mathcal{U}$ a B -definable function. Then for each complete type p over B with solution set $P \subset \text{dom}(f)$, there are a B -definable function $g : \rho\Gamma \rightarrow \Gamma$ and a B -definable subtorsor V of U such that for all $\delta \in P$, $f(\delta) \in \{B_{< g(\delta)}(V), B_{\leq g(\delta)}(V)\}$.*

Proof. Let δ realise p , and put $\gamma := \rho(\delta)$. Then $f(\delta) \in \text{acl}(B\gamma)$ and is a subtorsor of U with radius in $\text{dcl}(B\gamma)$, say $g(\delta)$. Now apply Proposition 2.4.4. \square

Notation 2.4.7 Let $T_n(K)$ denote the ring of $n \times n$ upper triangular matrices over K , $B_n(K)$ the group of invertible elements of $T_n(K)$, $U_n(K)$ the group of elements of $B_n(K)$ with ones on the diagonal, and $D_n(K)$ the group of diagonal matrices in $B_n(K)$. We have $B_n(K) = U_n(K)D_n(K)$, with $U_n(K) \trianglelefteq B_n(K)$. Let $T_n(R)$, $B_n(R)$, $U_n(R)$, $D_n(R)$ denote the corresponding objects over R (where inverses are assumed to be over R). Observe that the map $D_n(K) \rightarrow \Gamma^n$ which takes the diagonal matrix (d_1, \dots, d_n) to $(|d_1|, \dots, |d_n|)$ has kernel $D_n(R)$, so $D_n(K)/D_n(R)$ is \emptyset -definably isomorphic to $(\Gamma \setminus \{0\})^n$. For groups G, H with $H < G$ we write G/H for the space of left cosets of H in G . If $a \in D_n(K)$ and $A = aD_n(R) \in D_n(K)/D_n(R)$, define $U_n(R)^A := aU_n(R)a^{-1}$; this is well-defined. Finally, for $A, B \subset G$, A^B denotes $BAB^{-1} := \{bab^{-1} : a \in A, b \in B\}$.

Let $\ell = \binom{n}{2}$ and let ν_1, \dots, ν_ℓ enumerate the pairs (i, j) with $1 \leq i < j \leq n$. For each $m \leq \ell$, let $X_m := \{\nu_1, \dots, \nu_m\}$. We assume the ν_k are enumerated so that if $(i, j) \in X_m$ then $(i', j) \in X_m$ for $i' < i$ and $(i, j') \in X_m$ for $j' > j$. Now for each $m \leq \ell$, let $J_m := \{r \in T_n(K) : r(\nu) \neq 0 \rightarrow \nu \in X_m\}$, where $r(\nu)$ denotes the ν^{th} entry of r . Then J_m is a 2-sided ideal of $T_n(K)$, and $N_m := \{I_n + A : A \in J_m\}$ is a normal subgroup of $B_n(K)$. For $i \leq \ell$, put $G_i := U_n(K)/N_i$ (so in particular, G_ℓ is trivial). Observe that if $i < \ell$ then $N_i \trianglelefteq N_{i+1}$, and $M_i := N_{i+1}/N_i$ is naturally isomorphic to $(K, +)$. There is an exact sequence

$$1 \rightarrow M_i \rightarrow G_i \xrightarrow{\pi_i} G_{i+1} \rightarrow 1.$$

The next lemma shows that, to handle definable functions $\rho\Gamma \rightarrow S_n$, we have to describe definable functions $\rho\Gamma \rightarrow B_n(K)/B_n(R)$. Observe that $R^n \in S_n$, and that if $B \in \text{GL}_n(K)$ then $B(R^n)$, which is the image of the subset R^n of K^n under left multiplication of column vectors by B , is also an R -lattice in K^n ; it has the columns of B as an R -basis.

Lemma 2.4.8 *Let A be an R -lattice in K^n . Then there is $B \in B_n(K)$ such that $A = B(R^n)$. If also $B' \in B_n(K)$, then $B(R^n) = B'(R^n)$ if and only if $B'^{-1}B \in B_n(R)$.*

Proof. For existence, we use the filtration of A in the proof of Proposition 2.3.10. For each $i < n$, let A_i be the kernel of the projection of A to the last $n - i$ coordinates. Choose a basis (u_1, \dots, u_n) of A so that for each i , (u_1, \dots, u_i) is a basis for A_i . Then let u_i^T be the i^{th} column of B .

For uniqueness, suppose $B(R^n) = B'(R^n)$. Then $B'^{-1}B(R^n) = B^{-1}B'(R^n) = R^n$, so $B'^{-1}B$ and $B^{-1}B'$ have entries in R . Since they are upper triangular, $B'^{-1}B \in B_n(R)$. For the converse, observe that if $B'^{-1}B \in B_n(R)$ then $B'^{-1}B(R^n) \subseteq R^n$. If this containment is strict, then $B^{-1}B'(R^n)$ strictly contains R^n , which is impossible. \square

Remark 2.4.9 The above proof can also be thought of in the following way, useful below. Let $\text{TB}(K)$ be the set of triangular bases of K^n , that is, bases (v_1, \dots, v_n) where $v_i \in K^i \times \{0\}^{n-i}$ (that is, the last $n - i$ entries of v_i are zero). An element $a = (v_1, \dots, v_n) \in \text{TB}(K)$ can be identified with an element of $B_n(K)$, with v_i as the i^{th} column. Now $B_n(R)$ acts on $B_n(K) = \text{TB}(K)$ on the right. Two elements M, M' of $\text{TB}(K)$ generate the same R -module precisely if $MB_n(R) = M'B_n(R)$: indeed, M, M' generate the same R -module precisely if there is some $N \in \text{GL}_n(R)$ with $MN = M'$, and as $M, M' \in B_n(K)$, we must have $N \in \text{GL}_n(R) \cap B_n(K) = B_n(R)$. This gives an identification of S_n with $\text{TB}(K)$ modulo the right action of $B_n(R)$, that is, with the set of orbits of $B_n(R)$ on $\text{TB}(K)$. Equivalently, S_n can be identified with the set of left cosets of $B_n(R)$ in $B_n(K)$.

We wish also to treat T_n as a finite union of coset spaces. For each $m = 1, \dots, n$, let $B_{n,m}(k)$ be the set of elements of $B_n(k)$ whose m^{th} column has a 1 in the m^{th} entry and other entries zero. Let $B_{n,m}(R)$ be the set of matrices in $B_n(R)$ which reduce (elementwise) modulo \mathcal{M} to an element of $B_{n,m}(k)$. Let $e \in S_n$, and put $V := \text{red}(e)$. We may put $e = aB_n(R)$ for some $a = (a_1, \dots, a_n) \in \text{TB}(K)$ (so e is the orbit of a under $B_n(R)$, or the left coset $aB_n(R)$ where a is regarded as a member of $B_n(K)$). There is a natural filtration

$$\{0\} = V_0 < V_1 < \dots < V_{n-1} < V_n$$

of V , where V_i is the k -subspace of $\text{red}(e)$ spanned by $\{\text{red}(a_1), \dots, \text{red}(a_i)\}$ (here $\text{red}(a_j) = a_j + \mathcal{M}e$). Let $\text{TB}(V)$ be the set of triangular bases of V , that is, bases (v_1, \dots, v_n) where $v_i \in V_i \setminus V_{i-1}$. Now $B_n(k)$ acts sharply transitively on $\text{TB}(V)$ on the right, with

$$(v_1, \dots, v_n)(a_{ij}) = (a_{11}v_1, a_{12}v_1 + a_{22}v_2, \dots, \sum_{i=1}^n a_{in}v_i).$$

For each $i = 0, \dots, n$, put $O_i(V) = V_i \setminus V_{i-1}$ (so $O_0(V) = \{0\}$). It is easily verified that two elements of $\text{TB}(V)$ are in the same orbit under $B_{n,m}(k)$ precisely if they agree in the m^{th} entry. Thus, $O_m(V)$ can be identified with $\text{TB}(V)/B_{n,m}(k)$, and $V \setminus \{0\}$ with $\bigcup_{m=1}^n \text{TB}(V)/B_{n,m}(k)$.

If M is the triangular basis (a_1, \dots, a_n) of the lattice e , then

$$\text{red}(M) := (\text{red}(a_1), \dots, \text{red}(a_n)) = (a_1 + \mathcal{M}e, \dots, a_n + \mathcal{M}e).$$

From the last two paragraphs, it follows that if $M, M' \in \text{TB}(K)$, then they are $B_{n,m}(R)$ -conjugate (i.e. there is $N \in B_{n,m}(R)$ with $MN = M'$) precisely if they generate the same lattice A , and their reductions $\text{red}(M), \text{red}(M')$ are $B_{n,m}(k)$ -conjugate. This holds precisely if they generate the same lattice, and $\text{red}(M), \text{red}(M')$ have the same element of T_n in the m^{th} entry. The identification of $\text{TB}(K)$ with $B_n(K)$ now yields the following lemma.

Lemma 2.4.10 *For each $n > 0$, there is a \emptyset -definable bijection between T_n and $\cup_{m=1}^n B_n(K)/B_{n,m}(R)$.*

Below, when we say a definable function $f : \rho\Gamma \rightarrow X$ is ‘piecewise $*$ ’, we mean that its domain can be partitioned into finitely many definable pieces, and the restriction of f to each part has the form $*$. By compactness, Corollary 2.4.6 can also be formulated in this way.

Proposition 2.4.11 *Let $i \leq \ell$ and let g be a definable map on a definable subset I of a finite cover $\rho\Gamma$ of Γ , with $g(\gamma)$ a subgroup of G_i for each $\gamma \in I$. Suppose f is also a definable map on I , with $f(\gamma) \in G_i/g(\gamma)$. Then there is a partition of I into finitely many definable subsets I' such that for each I' there is $b \in G_i$ with $f(\gamma) = bg(\gamma)$ for all $\gamma \in I'$.*

Proof. We argue by induction on $\ell - i$. Suppose first that for all $\gamma \in I$, $f(\gamma) \in M_i g(\gamma)/g(\gamma)$. The latter is canonically in bijection with $M_i/M_i \cap g(\gamma)$. Since $M_i \cong (K, +)$, $f(\gamma)$ is a finite union of R -torsors each algebraic over $B\gamma$, where B is a parameter set defining the data in the proposition. The result follows in this case from Proposition 2.4.4.

As a slight extension of this, suppose there is fixed $b_0 \in G_i$ such that $f(\gamma) \in b_0 M_i g(\gamma)/g(\gamma)$. Then if $f'(\gamma) = b_0^{-1} f(\gamma)$, then f' satisfies the assumptions of the last paragraph. Hence, after subdividing I finitely we find b such that (piecewise) $f'(\gamma) = bg(\gamma)$, so

$$f(\gamma) = b_0 f'(\gamma) = b_0 b g(\gamma) = (b_0 b) g(\gamma).$$

For the general case, let $G(\gamma) := \pi_i(g(\gamma))$, a subgroup of G_{i+1} . Let $F(\gamma)$ be the image of $f(\gamma)$ in $G_{i+1}/G(\gamma)$. Using induction and arguing piecewise we may assume there is $B_0 \in G_{i+1}$ with $F(\gamma) = B_0 G(\gamma)$ for each $\gamma \in I$. There is $b_0 \in G_i$ with $B_0 = \pi_i(b_0)$. Then $f(\gamma) \in b_0 M_i g(\gamma)/g(\gamma)$. The result now follows from the last paragraph. \square

Corollary 2.4.12 (i) *Let $f : \rho\Gamma \rightarrow B_n(K)/B_n(R)$ be B -definable, where $\rho\Gamma$ is a finite cover of Γ . Then, piecewise, there is a B -definable function $h : \rho\Gamma \rightarrow D_n(K)/D_n(R)$ and some fixed $b \in U_n(K)$ such that $f(\gamma) = bh(\gamma)B_n(R)$.*

(ii) *Let $\rho\Gamma$ be a finite cover of Γ , let $m \in \{1, \dots, n\}$, and let $f : \rho\Gamma \rightarrow B_n(K)/B_{n,m}(R)$ be definable. Then, piecewise, there is a B -definable function $h : \rho\Gamma \rightarrow D_n(K)/D_n(R)$ and some fixed $b \in U_n(K)$ such that $f(\gamma) = bh(\gamma)B_{n,m}(R)$.*

Proof. (i) To obtain h , suppose $f(\gamma) = b(\gamma)B_n(R) = u(\gamma)d(\gamma)B_n(R)$, where $u = u(\gamma) \in U_n(K)$ and $d = d(\gamma) \in D_n(K)$. If also $f(\gamma) = u'd'B_n(R)$, then $d'^{-1}u'^{-1}ud \in B_n(R)$, which forces that $d'^{-1}d \in D_n(R)$. Thus the map $h(\gamma) = d(\gamma)D_n(R)$ is well-defined (and B -definable).

Now write $f(\gamma) = u(\gamma)h(\gamma)B_n(R)$. Here $u(\gamma)$ is not well-defined, but $f^*(\gamma) := u(\gamma)U_n(R)^{h(\gamma)}$ is: for $uh(\gamma)B_n(R) = u'h(\gamma)B_n(R)$ if and only if

$$u'^{-1}u \in h(\gamma)B_n(R)h(\gamma)^{-1} \cap U_n(K) = h(\gamma)U_n(R)h(\gamma)^{-1}.$$

By applying Proposition 2.4.11 (with $i = 0$) to f^* , there is $b \in U_n(K)$ with $f^*(\gamma) = bU_n(R)^{h(\gamma)}$ (piecewise). Then,

$$f(\gamma) = u(\gamma)h(\gamma)B_n(R) = u(\gamma)U_n(R)^{h(\gamma)}h(\gamma)B_n(R) = bU_n(R)^{h(\gamma)}h(\gamma)B_n(R),$$

which equals $bh(\gamma)B_n(R)$.

(ii) This is similar to (i). □

We now consider definable functions from a finite cover of Γ to \mathcal{G} , again using the notation of Definition 2.1.6.

Theorem 2.4.13 *Let $\rho\Gamma$ be a definable finite cover of Γ , let $f : \rho\Gamma \rightarrow \mathcal{G}$ be a definable function, and let B be a set of parameters over which Γ , f and ρ are defined. Then, piecewise, the following hold.*

- (i) *If $\text{ran}(f) \subset k \cup K$, then f is constant;*
- (ii) *If $\text{ran}(f) \subset \Gamma$, then there are $q \in \mathbf{Q}$ and $\delta \in \Gamma(B)$ with $f(\gamma) = \delta x^q$ for all x and $\gamma \in \rho^{-1}(x)$.*
- (iii) *Suppose $\text{ran}(f) \subset S_n$. Then there is $b \in U_n(K)$, and B -definable $h : \rho\Gamma \rightarrow \Gamma^n$ given in each coordinate by a definable function h_i satisfying (ii), such that for $\gamma \in \rho\Gamma$, $f(\gamma)$ is the lattice spanned by the columns of $bD(\gamma)$. Here, $D(\gamma)$ is any $n \times n$ diagonal matrix over K whose (i, i) -entry has norm $h_i(\gamma)$ for each i .*

(iv) *Suppose $\text{ran}(f) \subset T_n$. Then there are $b, h, D(\gamma)$ as in (iii) and some $m \in \{1, \dots, n\}$, such that for each γ , if a_m is the m^{th} column of $bD(\gamma)$, and $g(\gamma)$ is the lattice spanned by the columns of $bD(\gamma)$, then $f(\gamma) = a_m + \mathcal{M}g(\gamma)$. Also, $g(\gamma) = \tau(f(\gamma))$.*

Proof. As usual, we work piecewise. (i) is immediate for k and follows from Proposition 2.4.4 for K . Part (ii) follows from Proposition 2.1.3(iii) and quantifier elimination for divisible ordered abelian groups.

(iii) By Lemma 2.4.8, there is B -definable $f' : \rho\Gamma \rightarrow B_n(K)/B_n(R)$ such that for $\gamma \in \rho\Gamma$, $f(\gamma) = A(R^n)$ for any $A \in f'(\gamma)$. By Corollary 2.4.12, and arguing piecewise, there are fixed $b \in U_n(K)$ and $h : \rho\Gamma \rightarrow D_n(K)/D_n(R)$ such that for $\gamma \in \rho\Gamma$, $f'(\gamma) = bh(\gamma)B_n(R)$. Thus, $f(\gamma)$ is the lattice with an R -basis given by the columns of $bh(\gamma)$. Regarding h as a function $\rho\Gamma \rightarrow \Gamma^n$ (as mentioned under Notation 2.4.7), we obtain (iii).

(iv) In this case, we apply the identification from Lemma 2.4.10. There is a B -definable function $f' : \rho\Gamma \rightarrow B_n(K)/B_{n,m}(R)$ such that if $f'(\gamma) = AB_{n,m}(R)$ then $g(\gamma) = \tau(f(\gamma))$ is the lattice spanned by the columns of A , and $f(\gamma) = v + \mathcal{M}A$ where v is the m^{th} column of A . By Corollary 2.4.12(ii), there is $b \in U_n(K)$ and B -definable $h : \rho\Gamma \rightarrow D_n(K)/D_n(R)$ such that (piecewise), $f'(\gamma) = bh(\gamma)B_{n,m}(R)$.

Remark 2.4.14 1. We mention another way of viewing definable functions $f : \rho\Gamma \rightarrow T_n$. Given such f , there is definable $g : \rho\Gamma \rightarrow S_n$ with $g(\gamma) = \tau(f(\gamma))$.

Then, piecewise, there are $b, D(\gamma)$ as in Theorem 2.4.13(iii), and a set V such that, for all γ , V lies in a single coset of $\mathcal{M}g(\gamma)$ and $f(\gamma)$ is the element of $\text{red}(g(\gamma))$ containing V .

To see this, argue as follows. The existence of $b, D(\gamma)$ for g is by (iii). By adding parameters for a matrix C mapping the lattice identified with b to R^n , we may suppose that $b = R^n$. Let $h_i(\gamma)$ denote the norm of the (i, i) -element of $D(\gamma)$. Then $g(\gamma) = ((h_1(\gamma)R, \dots, h_n(\gamma)R)$, and $\text{red } g(\gamma) = (\text{red}(h_1(\gamma)R), \dots, \text{red}(h_n(\gamma)R))$. Thus, $f(\gamma) = (f_1(\gamma), \dots, f_n(\gamma))$ with $f_i(\gamma) \in \text{red}(h_i(\gamma)R)$. By Proposition 2.4.4, for each i there is a torsor V_i with $f_i(\gamma)$ equal to the element of $\text{red}(h_i(\gamma)R)$ containing V_i . Put $V := \prod_{i=1}^n V_i$. Then piecewise, V lies in a single coset of $\mathcal{M}g(\gamma)$ and $f(\gamma)$ is the element of $\text{red}(g(\gamma))$ containing V .

2. If f is definable over a parameter set B , then the pieces can be chosen to be B -definable. This is essentially the content of Lemma 3.3.6 below.

We will show later (Proposition 3.3.4) that definable functions $\Gamma \rightarrow \mathcal{G}$ are coded in \mathcal{G} . This is not immediate, since in Theorem 2.4.13 (iii), the element b is not in general determined by the function f .

2.5 Independence and orthogonality to Γ for unary types

Definition 2.5.1 Let $a \in K^{\text{eq}}$ be an element of a unary set, and C, B be sets of parameters with $C = \text{acl}(C) \subset \text{dcl}(B)$. We say that a is *generically independent from B over C* , and write $a \downarrow_C^g B$, if either $a \in \text{acl}(C)$, or, whenever a is generic over C in a C -unary set U , it remains generic in U over B .

In the subsequent paper [3], we shall extend this to obtain a notion of sequential independence for n -tuples, and in particular, for elements of \mathcal{G} . This will give an invariant extension of any type, partly because of the next result, which ensures that unary types have invariant extensions. As in Section 2.3, the frequent assumption $C = \text{acl}(C)$ in this section can be weakened: the most that is needed is that in an ambient C -1-torsor U , any subtorsor algebraic over C is definable over C .

Proposition 2.5.2 Let B, C be sets of parameters with $C = \text{acl}(C) \subset \text{dcl}(B)$, and let p be the type of an element of a C -unary set U . Then there is a unique unary type q over B extending p such that if $\text{tp}(a/B) = q$ then $a \downarrow_C^g B$.

Proof. We may suppose that p is non-algebraic, as otherwise the result is immediate. By Lemma 2.3.6, p is the generic type of a unique C -unary subset of U . The type q must be the generic type of this unary set over B . \square

Lemma 2.5.3 Let $C_0 \subseteq C$ with $C_0 = \text{acl}(C_0)$, and let P be the solution set of a non-algebraic unary type over C_0 . Suppose that $z \in K^{\text{eq}}$, $d \in P$ is generic in P over Cz , and $z \in \text{acl}(Cd)$. Then $z \in \text{acl}(C)$.

Proof. Since $z \in \text{acl}(Cd)$, we may suppose that $z \in F(d)$, where $F(d)$ is a Cd -definable set of size m . Suppose for a contradiction that $z \notin \text{acl}(C)$, and let z_1, \dots, z_{m+1} be conjugates of z over C . Then for each z_i and any $d' \in P$ generic in P over Cz_i we have $d \downarrow_{C_0}^g Cz$ and $d' \downarrow_{C_0}^g Cz_i$, so by Proposition 2.5.2 $dz \equiv_C d'z_i$, and so $z_i \in F(d')$. Choose $d' \in P$ generic in P over z_1, \dots, z_{m+1} . Then $z_1, \dots, z_{m+1} \in F(d')$, which is impossible. \square

Next, we introduce a notion of orthogonality to Γ . At this stage, it is introduced just for unary types, and it is extended to arbitrary types in the subsequent paper.

Definition 2.5.4 Let $C = \text{acl}(C)$, and $a \in K^{\text{eq}}$ be an element of a C -unary set. We write $\text{tp}(a/C) \perp \Gamma$, and say $\text{tp}(a/C)$ is *orthogonal* to Γ if, for any algebraically closed valued field M such that $C \subset \text{dcl}(M)$ and $a \downarrow_C^g M$, we have $\Gamma(M) = \Gamma(Ma)$.

Our first lemma shows that orthogonality to Γ is equivalent to genericity in a closed unary set.

Lemma 2.5.5 *Let $C = \text{acl}(C)$ and $a \in K^{\text{eq}} \setminus C$ lie in a C -unary set U . Then the following are equivalent:*

- (i) a is generic over C in a closed subtorsor of U defined over C .
- (ii) $\text{tp}(a/C) \perp \Gamma$.

Furthermore, if $A = \text{acl}(Ca)$ then condition (iii) $\text{trdeg}(k(A)/k(C)) = 1$ implies both (i) and (ii). If in addition $C = \text{acl}(C \cap K)$, then (i), (ii) are equivalent to (iii).

Proof. (i) \Rightarrow (ii) Suppose that a is generic over C in the closed subtorsor T of U . Let M be an algebraically closed valued field with $C \subset \text{dcl}(M)$ and $a \downarrow_C^g M$, and suppose for a contradiction that there is $\gamma \in \Gamma(Ma) \setminus \Gamma(M)$. Since $\text{acl}_\Gamma(Ma) = \text{dcl}_\Gamma(Ma)$, there is an M -definable function $f : T \rightarrow \Gamma$ with $f(a) = \gamma$, defined on an M -definable set D containing generic elements of T . Since $\gamma \notin \text{dcl}(M)$, f is not generically constant on D . It follows that f is not constant on generic elements of $\text{red}(T)$, since otherwise it would induce a definable generically non-constant function from a strongly minimal set to an o-minimal set. For each generic $V \in \text{red}(T)$, $\{f(x) : x \in V\}$ is a finite union of intervals and singletons of Γ , and for simplicity we suppose it is always an interval, denoted $f(V)$. By considering the corresponding function to the endpoints, the map $V \mapsto f(V)$ from $\text{red}(T)$ is generically constant, with $f(V) = I$ for generic $V \in \text{red}(T)$. It follows that if $\delta \in I$, then the definable set $f^{-1}(\delta)$ meets each generic element of $\text{red}(T)$ in a proper non-empty subset, contrary to Lemma 2.3.3.

(ii) \Rightarrow (i) Suppose (i) is false. Then a is generic in a unary set T which is an open 1-torsor or an ∞ -definable 1-torsor or a subset of Γ . We may assume

the last case does not occur as it clearly contradicts (ii). There is a model M containing C , and containing an element $b \in T$. We may choose M with $a \downarrow_C^g M$. Then $|a - b|$ is in $\Gamma(Ma) \setminus \Gamma(M)$, so $\text{tp}(a/C) \not\perp \Gamma$ (recall the notation $|a - b|$ from the beginning of Section 2.3).

(iii) \Rightarrow (i). Suppose (iii) holds, and let $p := \text{tp}(a/C)$, with solution set P . Let $s \in k(Ca) \setminus k(C)$, and let $s = s_1, \dots, s_m$ be the conjugates of s over Ca . Then there is a C -definable function with domain containing P and with range in the set of m -element subsets of k , with $f(a) = \{s_1, \dots, s_m\}$. Since finite sets are coded in the field k by tuples, we may suppose that $m = 1$. If now (i) is false, then p is the generic type of an open 1-torsor or an ∞ -definable 1-torsor or a subset of Γ and f is a definable function from P taking infinitely many distinct values in the strongly minimal set k , which is clearly impossible.

(i) \Rightarrow (iii). Suppose $C = \text{acl}(C \cap K)$ and a is generic in the closed 1-torsor T . Then there is a C -definable bijection between $\text{red}(T)$ and k . The element of $\text{red}(T)$ containing a is thus interdefinable over C with an element of $k(A) \setminus k(C)$. Thus, $\text{trdeg}(k(A)/k(C)) \geq 1$. Since T is in C -definable bijection with the true 1-torsor R , there is $a' \in A \cap K$ with $A = \text{acl}(Ca')$, and so we have equality. \square

Despite the fact that a generic element of an open torsor or an ∞ -definable torsor is not orthogonal over the parameters to Γ , it still need not increase the value group.

Lemma 2.5.6 *Let $C \subset K^{\text{eq}}$, and let T be a C -1-torsor which is not closed. Then the following are equivalent:*

- (i) *no proper subtorsor T' of T is algebraic over C ;*
- (ii) *for all a generic in T , $\Gamma(C) = \Gamma(Ca)$.*

Proof. For the direction (i) \Rightarrow (ii), suppose for contradiction that a is generic in T and $\delta \in \Gamma(Ca) \setminus \Gamma(C)$. Then there is a C -definable function $f : T \rightarrow \Gamma$ with $f(a) = \delta$, and $f^{-1}(\delta)$ is a proper subset of T . By Lemma 2.3.3, there is a proper subtorsor T_δ of T (possibly a field element) algebraic over $C\delta$, and we may choose T_δ to be definable over $C\delta$. By Proposition 2.4.4, T_δ is a neighborhood of some T' algebraic over C , contrary to the hypothesis.

For (ii) \Rightarrow (i), let T' be a proper unary subset of T algebraic over C . If a is generic in T over C , then $|a - T'|$ is Ca -definable (it is the constant value of $|a - c|$ as c ranges over T'); it is not in $\Gamma(C)$. \square

The definition of $\text{tp}(a/C) \perp \Gamma$ says that a does not increase the value group of a model from which a is independent. The next lemma shows that this is true for C itself. Its converse is false. For if C is the algebraic closure (in K^{eq}) of an algebraically closed valued field which is not maximal, then there is a C - ∞ -definable 1-torsor T which is not C -definable such that $\{x \in T\}$ determines a complete type p over C . If a realises p , then $\Gamma(C) = \Gamma(Ca)$ by Lemma 2.5.6, but $\text{tp}(a/C) \not\perp \Gamma$ by Lemma 2.5.5.

Lemma 2.5.7 *Suppose $C = \text{acl}(C)$ and $\text{tp}(a/C) \perp \Gamma$. Then $\Gamma(C) = \Gamma(Ca)$.*

Proof. By Lemma 2.5.5, a is chosen generically over C in a closed 1-torsor. The proof of 2.5.5 (i) \Rightarrow (ii) easily yields $\Gamma(C) = \Gamma(Ca)$ (the fact that M is a model is not used here). \square

Next, we give an easy lemma on closed 1-torsors, which shows that when we choose a sequence of elements generically from a sequence of closed 1-torsors, the order of the sequence does not affect the genericity. It will be used without explicit reference in the remainder of the paper. Generalisations will appear in the subsequent paper.

Lemma 2.5.8 *Let T_1, \dots, T_n be closed 1-torsors defined over a parameter set C , and suppose that for each $i = 1, \dots, n$, a_i is generic in T_i over $Ca_1 \dots a_{i-1}$. Then for each i , a_i is generic in T_i over $C \cup \{a_j : j \neq i\}$.*

Proof. We prove the result by induction on n . For convenience we suppose $i = 1$. So suppose that a_1 is generic in T_1 over $Ca_2 \dots a_{n-1}$ but not over $Ca_2 \dots a_n$. Let $S \in \text{red}(T_1)$ contain a_1 . Then as $\text{red}(T_1)$ is strongly minimal, there is an algebraic formula $\varphi(u, a_n)$ over $Ca_1 \dots a_{n-1}$ such that $\varphi(S, a_n)$ holds. Hence, as $\text{red}(T_n)$ is strongly minimal, for all elements $S' \in \text{red}(T_n)$ except for finitely many, and all $y \in S'$, the formula $\varphi(u, y)$ is algebraic and $\varphi(S, y)$ holds. This contradicts that a_1 is generic in T_1 over $Ca_2 \dots a_{n-1}$. \square

We conclude this section with a lemma which gives symmetry of \downarrow^g , under weaker conditions than in Lemma 2.5.8. It will be used in the subsequent paper, when \downarrow^g -independence is extended to n -types.

Definition 2.5.9 If $C = \text{acl}(C)$, and $a \in K^{\text{eq}}$ is an element of a unary set, we shall say that $\text{tp}(a/C)$ is *order-like* if a is generic over C in a C -unary set which is either (i) contained in Γ , or (ii) an open 1-torsor, or (iii) an ∞ -definable 1-torsor which contains a proper C -unary subset.

Remark 2.5.10 If $\text{tp}(a/C)$ is order-like (and if $C = \text{acl}(C \cap K)$ in case (ii)) then the second part of Lemma 2.5.6 will apply and $\Gamma(C) \neq \Gamma(Ca)$. Conversely, if $\text{tp}(a/C)$ is not order-like but a is an element of a unary set, then either a is generic in a closed 1-torsor or a is generic in an ∞ -definable 1-torsor which does not contain a proper C -unary subset. By Lemma 2.5.5 in the first case, and by the first part of Lemma 2.5.6 in the second case, $\Gamma(C) = \Gamma(Ca)$.

Lemma 2.5.11 *Let $C = \text{acl}(C)$. Let $a, b \in K^{\text{eq}}$ be elements of C -1-torsors U and V respectively, and put $A = \text{acl}(Ca)$ and $B = \text{acl}(Cb)$. Assume that at least one of $\text{tp}(a/C)$, $\text{tp}(b/C)$ is not order-like. Then $a \downarrow_C^g B$ if and only if $b \downarrow_C^g A$.*

Proof. By Proposition 2.5.2, it suffices to find a 2-type $r(x, y) \supset \text{tp}_x(a/C) \cup \text{tp}_y(b/C)$ such that if $r(a', b')$ holds then $a' \downarrow_C^g b'$ and $b' \downarrow_C^g a'$.

Suppose first that one of a, b , say a , is generic in a closed 1-torsor T of C , and write $[a]$ for the element of $\text{red}(T)$ which contains a . Now $a \downarrow_C^g b$ if and only if $[a] \notin \text{acl}(Cb)$. Thus, it suffices in this case to show that $b \downarrow_C^g a$ if and only if $[a] \notin \text{acl}(Cb)$. For this, in one direction, suppose $[a]$ has n conjugates over Cb , and argue as in Lemma 2.5.3. Choose $a_1, \dots, a_{n+1} \models \text{tp}(a/C)$ with the $[a_i]$ all distinct, and choose $b' \equiv_C b$ with $b' \downarrow_C^g a_1 \dots a_{n+1}$. We cannot have $\text{tp}(b'a_i/C) = \text{tp}(ba/C)$, so $b \not\downarrow_C^g a$ by the uniqueness in Proposition 2.5.2. Conversely, if $b \not\downarrow_C^g a$ then there is a proper subtorsor S of V algebraic over Ca and containing b . We can assume V is the intersection of a chain $(V_i)_{i \in I}$ of C -1-torsors, possibly with a least element. Fix $i_0 \in I$ and take radius to be defined with respect to V_{i_0} . Let $\gamma := \inf\{\text{rad}(W) : W \text{ is a } C\text{-subtorsor of } V_{i_0} \text{ containing } b\}$, and $\delta = \text{rad}(S)$. Suppose for contradiction that $[a] \notin \text{acl}(Cb)$. Then this situation holds for all generic elements of $\text{red}(T)$. Hence there is a C -definable partial function $f : T \rightarrow \Gamma$, defined generically on T , with $f(a) = \delta$. Now for each $u \in \text{red}(T)$ generic over Cb , put $\hat{f}(u) := \sup\{f(x) : x \in u\}$. Since $\text{red}(T)$ is strongly minimal, \hat{f} is generically constant with value $\hat{\gamma}$, say. Then there is a C -unary set containing b of radius $\hat{\gamma}$, hence $\hat{\gamma} = \gamma$, and I has a least element. It follows that if δ' is chosen generically below γ , then $f^{-1}(\delta')$ contains some, but not all, elements of infinitely many members of $\text{red}(T)$. Since $f^{-1}(\delta')$ is definable, this contradicts Theorem 2.1.2. The lemma is thus proved in the case when either of a or b is generic in a closed 1-torsor.

Suppose now that $\text{tp}(b/C)$ is order-like. Then by our assumption, $\text{tp}(a/C)$ is not order-like, so by the last paragraph, we may suppose that $\text{tp}(a/C)$ is the intersection E of a chain $\{U_i : i \in I\}$ of 1-torsors with no least element, such that there is no C -definable proper unary subset of E . By Lemma 2.5.6, $\Gamma(C) = \Gamma(A)$. We shall show that in this situation, $b \downarrow_C^g a$. This suffices, by the first paragraph of the proof, and Proposition 2.5.2; for we can choose $a' \equiv_C a$ with $a' \downarrow_C^g b$, and our argument shows that $b \downarrow_C^g a'$, and we have $ab \equiv_C a'b$.

Suppose first that b is generic over C in a chain of 1-torsors $(V_i : i \in I)$ which contains a proper subtorsor W all defined over C . If $b \not\downarrow_C^g a$ then there is a sub-torsor T of the V_i which contains b and lies in $\text{acl}(Ca)$. It follows that $|T - W| = \sup\{|x - y| : x \in T, y \in W\}$ lies in $\Gamma(Ca) \setminus \Gamma(C)$, which is impossible, by Lemma 2.5.6.

Next, suppose b is generic over C in an open 1-torsor S . We may suppose there is no C -definable proper subtorsor of S , since otherwise the above argument works. If $b \not\downarrow_C^g a$, then again there is a subtorsor T of S , algebraic over Ca and containing b . Since $\Gamma(C) = \Gamma(Ca)$, $\delta := \text{rad}(T) \in \Gamma(C)$. Also, since we may replace T by the smallest 1-torsor containing all its conjugates over Ca , we may suppose $T \in \text{dcl}(Ca)$, with $T = f(a)$ for some C -definable function f . Now, the domain of f contains E . Let D be the set of closed subtorsors of S of radius δ . Then by Lemma 2.4.5(i), D is a 1-type over C so we may suppose the range of

f is exactly D and hence that $\text{dom}(f) = U_i$. Since D is a 1-type over C , for any $T' \in D$, $f^{-1}(T')$ contains some but not all elements of $U_j \setminus U_k$ for any $j, k \in I$ with $i < j < k$. This contradicts Lemma 2.3.3. This concludes the case when $\text{tp}(b/C)$ is order-like.

Thus, we may suppose that neither of $\text{tp}(a/C)$, $\text{tp}(b/C)$ is order-like or generic in a closed 1-torsor. Thus, a is generic in an intersection E of a chain of subtorsors of U over C whose radii (with respect to a fixed element of the chain) have infimum $\text{rad}(E)$ (a cut in Γ), and b is generic in an intersection F of a chain of subtorsors of V over C whose radii have infimum $\text{rad}(F)$. Furthermore, by Lemma 2.5.6, $\Gamma(C) = \Gamma(A)$ and $\Gamma(C) = \Gamma(B)$.

Claim. For $a' \in E$ and $b' \in F$, $a' \downarrow_C^g b'$ if and only if for all $\text{acl}(Cb')$ -subtorsors U' of U and for all $\alpha \in \Gamma(C)$ if $\alpha < \text{rad}(E)$ then for some $x \in U'$, $|a' - x| > \alpha$.

Proof of Claim. The direction \Rightarrow is immediate. Conversely, if $a' \not\downarrow_C^g b'$, then there is a Cb' -algebraic closed 1-torsor U' contained in E and containing a' , and we may put $\alpha = \text{rad}(U')$ (so $\alpha \in \Gamma(C)$ as $\Gamma(C) = \Gamma(Cb')$).

Suppose there is no type $r(x, y)$ as at the beginning of the proof. Then by compactness and the claim, there are $\alpha, \beta \in \Gamma(C)$ with $\alpha < \text{rad}(E)$ and $\beta < \text{rad}(F)$, and formulas $\varphi(x, u)$, $\psi(v, y)$ over C such that $\varphi(a, u)$ and $\psi(v, b)$ each have finitely many solutions, and such that the following holds: there is no pair (a', b') with $a' \in E$ and $b' \in F$, so that $\forall x \in d(|a' - x| > \alpha)$ for each d satisfying $\psi(d, b')$, and $\forall y \in c(|b' - y| > \beta)$ for each c satisfying $\varphi(a', c)$. By compactness, E and F can be replaced by closed 1-torsors T, S (containing E, F respectively) such that the same statement holds. This implies there do not exist a'', b'' in U , generic in T, S respectively, with $a'' \downarrow_C^g b''$ and $b'' \not\downarrow_C^g a''$. However, this contradicts Lemma 2.5.8. \square

2.6 Sets internal to k

By Proposition 2.1.3, the residue field k is stably embedded and strongly minimal. This enables us to construct, over any base set of parameters, a part of the structure which inherits stability-theoretic properties from k , and plays a crucial role later. We give a tidy description of this part of the structure, denoted $\text{Int}_{k,C}$ where we work over parameters C , prove that members of any k -internal set are coded in it, and that it has elimination of imaginaries. In [3] it will determine independence (for types orthogonal to Γ), and motivates the development of *stable domination*.

Recall that $\mathcal{G} = K \cup k \cup \Gamma \cup \mathcal{S} \cup \mathcal{T}$ is the union of the geometric sorts.

Definition 2.6.1 A definable set D is *k -internal* if there is a finite $F \subset \mathcal{G}$ with $D \subseteq \text{dcl}(kF)$.

Lemma 2.6.2 *Let $C \subset K^{\text{eq}}$ and let $D \subset \mathcal{G}^\ell$ be C -definable. Then the following are equivalent.*

- (i) D is k -internal.
- (ii) D , expanded by predicates for C -definable relations, has finite Morley rank.
- (iii) D (with the induced C -definable structure as in (ii)) does not have the strict order property.
- (iv) For any k , there is no definable surjective map from D^k to an infinite interval in Γ .
- (v) D is finite or (possibly after a permutation of coordinates) is contained in a finite union of sets of the form $\text{red}(s_1) \times \dots \times \text{red}(s_m) \times F$ where s_1, \dots, s_m are $\text{acl}(C)$ -definable elements of \mathcal{S} and F is a C -definable finite set of tuples from \mathcal{G} .
- (vi) $D \subset \text{dcl}(kE)$ for some finite $E \subset D$.
- (vii) For $i = 0, \dots, n$ there are definable sets $D_i \subset \mathcal{G}^{\ell_i}$ with D_0 finite and $D \subseteq \text{dcl}(D_n)$ and for $i = 1, \dots, n$ there is a definable map $f_i : D_{i+1} \rightarrow D_i$ whose fibres are stably embedded and k -internal (that is, D is k -analysable).

Proof. The implication (i) \Rightarrow (ii) holds since k is strongly minimal and stably embedded, and (ii) \Rightarrow (iii), (iii) \Rightarrow (iv) are trivial.

(iv) \Rightarrow (v) Let a be any coordinate of an element of D . If $a \in \Gamma$ then by (iv), $a \in \text{dcl}(C)$. We show that if $a \in S_n$ then $a \in \text{acl}(C)$ (an easier argument shows the same if $a \in K$). Then, if $a \in T_n$, we have $\tau(a) \in \text{acl}(C)$.

So suppose $a \in S_n$, and let (a_1, \dots, a_r) be a unary code for a with elements from \mathcal{G} , as given in Steps 1 and 2 of Proposition 2.3.10. We show inductively that $a_i \in \text{acl}(C)$ for each i . Suppose it holds for all $j < i$. We may suppose that a_i is chosen in a unary set T in definable bijection with $R/\gamma R$ for some $\gamma \in \Gamma$; otherwise, by inspection of the proof of 2.3.10, $T \subset \Gamma$, and clearly $a_i \in \text{acl}(C)$. If $a_i \notin \text{acl}(C)$, then there is a definable surjection from D to an infinite subtorsor of $R/\gamma R$. By composing this with the map $x \mapsto |x - b|$ (for some parameter b) we obtain a contradiction to (iv).

(v) \Rightarrow (vi) We may suppose by (v) that D is a subset of $\text{red}(s_1) \times \dots \times \text{red}(s_m) \times F$ with the s_i and F as above. Let s'_i be the projection of D to s_i , and let E_i be a maximal linearly independent (over k) subset of s'_i . Then $s'_i \subset \text{dcl}(kE_i)$. Thus, we may choose E to be any finite set which projects onto F and onto each E_i .

The implication (vi) \Rightarrow (i) is trivial. Also, (vi) \Rightarrow (vii) is trivial, and (vii) \Rightarrow (ii) is an easy induction on n . \square

Remark 2.6.3 The above lemma yields that for C -definable $D \subset \mathcal{G}^\ell$, D is k -internal if and only if it is stably embedded and stable (when equipped with its C -definable structure). Indeed, (vi) above yields stable embeddedness, whilst in the other direction, any C -definable stable set satisfies (iii), and hence (i). The same result will follow for *arbitrary* D from elimination of imaginaries (see Proposition 3.4.11 below).

In (v) the lattices and tuples are $\text{acl}(C)$ -definable. In Proposition 3.4.11 it will be shown that they can be taken to be C -definable.

For any parameter set C , we denote by $\text{Int}_{k,C}$ a many-sorted structure whose sorts are the k -vector spaces $\text{red}(s)$ where $s \in \text{dcl}(C) \cap \mathcal{S}$. Each sort $\text{red}(s)$ is equipped with its k -vector space structure, along with any other C -definable relations as \emptyset -definable relations. As in condition (vi) of Lemma 2.6.2, we have that if $s \in S_n$ is C -definable then there is finite $E \subset \text{red}(s)$ with $\text{red}(s) \subset \text{dcl}(kE)$. It follows that $\text{Int}_{k,C}$ is stably embedded in K^{eq} , and is stable. Proposition 3.4.11 below, which rests both on elimination of imaginaries and on some of the lemmas for coding finite sets, will clarify the role of $\text{Int}_{k,C}$, which will be central in [3].

Below we prove elimination of imaginaries for $\text{Int}_{k,C}$. The first step is to prove that subspaces of its sorts are coded in $\mathcal{T} \cup k$.

Lemma 2.6.4 *Let $R_{n,\ell}$ be the sort consisting of all ℓ -dimensional subspaces of the k -spaces $\text{red}(A)$ where $A \in S_n$. Then every member of $R_{n,\ell}$ is coded in $\mathcal{T} \cup k$.*

Proof.

(1) Let $N = \binom{n}{\ell}$. Then K^N can be identified with $\Lambda^\ell(K^n)$, the ℓ^{th} exterior power, via the standard basis $\{e_1, \dots, e_n\}$ for K^n and the standard basis $\{e_{i_1} \wedge \dots \wedge e_{i_\ell} : i_1 < \dots < i_\ell\}$ for $\Lambda^\ell(K^n)$. We have an alternating multilinear map $c_\ell : (K^n)^\ell \rightarrow K^N$. If $A \in S_n$, let $\Lambda^\ell(A) = c_\ell(A^\ell)$. Now A is a free R -module on some a_1, \dots, a_n ; this is equally a basis for K^n , and so clearly the various wedges $a_{i_1} \wedge \dots \wedge a_{i_\ell}$, $i_1 < \dots < i_\ell$, form a basis for the exterior power $\Lambda^\ell(K^n)$, and also a free basis for the R -module $\Lambda^\ell(A)$. Moreover, c_ℓ induces a (canonical) k -vector space isomorphism $\Lambda^\ell(\text{red}(A)) \rightarrow \text{red}(\Lambda^\ell(A))$.

We may canonically identify K^{nm} with the K -vector space $K^n \otimes K^m$, via the standard basis. Hence, given $A \in S_n$ and $B \in S_m$, we find $C \in S_{nm}$ and a canonical isomorphism $A \otimes_R B \rightarrow C$. Identify K^n with its dual space $(K^n)^*$, again via the standard basis. Given $A \in S_n$, define

$$A^* := \{f \in (K^n)^* : \text{for all } a \in A, f(a) \in R\}.$$

It is easy to see that A^* is indeed isomorphic to $\text{Hom}_R(A, R)$, and $A^* \in S_n$ (via the above identification).

(2) It follows that if $A \in S_n$ and $B \in S_m$, then there exists $C \in S_{nm}$ and a canonical isomorphism $\text{Hom}_R(A, B) \rightarrow C$. Namely,

$$\text{Hom}_R(A, B) \cong \text{Hom}_R(A, R) \otimes_R B = A^* \otimes_R B$$

(for the isomorphism see for example Corollary 5.5 on p. 580 of [9]).

(3) If $A \in S_n$ and $B \in S_m$ then there is an isomorphism $\varphi : \text{red}(\text{Hom}_R(A, B)) \rightarrow \text{Hom}_k(\text{red}(A), \text{red}(B))$: for $f \in \text{Hom}_R(A, B)$ and $a \in A$ define

$$\varphi(f + \mathcal{M} \text{Hom}_R(A, B))(a + \mathcal{M}A) = f(a) + \mathcal{M}B.$$

An ℓ -dimensional subspace of $\text{red}(A)$ can be coded by a 1-dimensional subspace of $\Lambda^\ell(\text{red}(A))$ (namely the space generated by c_1, \dots, c_ℓ is coded by the space generated by $c_1 \wedge \dots \wedge c_\ell$.) Thus, by the identification of $\Lambda^\ell(\text{red}(A))$ with $\text{red}(\Lambda^\ell(A))$ from (1), we see that the union of all the sorts $R_{n,1}$ suffices to code all the $R_{n,\ell}$. We shall prove by induction on n that all the sorts $R_{n,1}$ are coded in $\mathcal{T} \cup k$.

Let $0 \rightarrow A \rightarrow B \rightarrow C \rightarrow 0$ be an exact sequence of R -modules, with C free. Then $0 \rightarrow MA \rightarrow MB \rightarrow MC \rightarrow 0$ and $0 \rightarrow \text{red}(A) \rightarrow \text{red}(B) \rightarrow \text{red}(C) \rightarrow 0$ are also exact. (For since C is free, the first sequence splits, and is isomorphic to $0 \rightarrow A \rightarrow (A \oplus C) \rightarrow C \rightarrow 0$; for this sequence the result is obvious.)

Let $A \in S_n$, and let $H \in R_{n,1}$ be a 1-dimensional subspace of $\text{red}(A)$. We aim to code H , and argue by induction on n . Let $h : K^n = (K \times K^{n-1}) \rightarrow K$ be the first coordinate projection, and let A' be so that $\{0\} \times A' = \ker(h)$. Then

$$0 \rightarrow A' \rightarrow A \xrightarrow{h} hA \rightarrow 0$$

is exact. As $hA \subset K$ is a finitely generated R -submodule of K , it is free, so by the last paragraph,

$$0 \rightarrow \text{red}(A') \xrightarrow{f} \text{red}(A) \xrightarrow{\text{red}(h)} \text{red}(hA) \rightarrow 0$$

is exact. If $\text{red}(h)$ vanishes on $H \leq \text{red}(A)$, then $H = f(H')$ for a unique 1-dimensional $H' \leq \text{red}(A')$, and to code H it suffices to code H' . The latter is possible by induction on n . If on the other hand $\text{red}(h) : \text{red}(A) \rightarrow \text{red}(hA)$ is injective on H , then H is inter-definable with an element H^* of $\text{Hom}_k(\text{red}(hA), \text{red}(A))$; namely, H^* is the unique homomorphism on $\text{red}(hA)$ with image H and such that $hH^* = \text{id}_{\text{red}(hA)}$. But $\text{Hom}_k(\text{red}(hA), \text{red}(A))$ is canonically isomorphic to $\text{red}(\text{Hom}_R(hA, A))$ by (3); and by (2) $\text{Hom}_R(hA, A)$ is canonically isomorphic to some element B of S_n . As $\text{red}(B) \subset T_n$, H is coded (in T_n). \square

Proposition 2.6.5 *Let $C \subset K^{\text{eq}}$. Then $\text{Int}_{k,C}$ has elimination of imaginaries.*

Proof. We first reduce to coding subsets of $\text{red}(s)$, where $s \in \mathcal{S}$ is C -definable. For this, note that any definable $U \subset \text{red}(s_1)^{i_1} \times \dots \times \text{red}(s_k)^{i_k}$ (where the s_i are C -definable) is interdefinable over C with some $U' \subset \text{red}(s_1^{i_1} \times \dots \times s_k^{i_k})$.

Observe that the collection of sorts $\text{red}(s)$ in $\text{Int}_{k,C}$ is closed under duals and tensors, by (1), (2) of the proof of Lemma 2.6.4. We suppose $A \in S_n \cap \text{dcl}(C)$, and $V = \text{red}(A)$, and that Y is a definable subset of V . If a basis of V is fixed, then V may be identified with k^n , and so we may talk of Zariski closed subsets of V ; furthermore, this notion is independent of the choice of basis. Since any definable subset of V is a Boolean combination of Zariski closed sets defined over the same parameters, we may suppose Y is Zariski closed. (This reduction to the Zariski closed case appears to require a coding of finite sets, in order to code the

Boolean combination; this is not a problem, as finite sets of tuples *are* Zariski closed.)

If a basis of V is fixed, then V^* has a corresponding (dual) basis, and hence there is an identification of the polynomial ring $k[X_1, \dots, X_n]$ with the ring $S(V) := k \oplus V^* \oplus \Sigma_{i=2}^{\infty} \text{Sym}^i(V^*)$, where $\text{Sym}^i(W)$ denotes the symmetric i^{th} power of W . Furthermore, elements of $S(V)$ induce functions $V \rightarrow k$ independently of the choice of basis of V . Also, the ideal in $S(V)$ which vanishes on Y is independent of basis, and it follows from the above identification that this ideal determines Y . As $S(V)$ is noetherian, this ideal is determined by its intersection with the vector space $S^m(V) := k \oplus V^* \oplus \Sigma_{i=2}^m \text{Sym}^i(V^*)$ for some sufficiently large finite m . This intersection is a subspace U of $S^m(V)$. Let U' be the pullback of U in $T^m(V) := k \oplus V^* \oplus \Sigma_{i=2}^m \otimes^i(V^*)$. Now as $\text{Int}_{k,C}$ is closed under duals and tensors, $T^m(V)$ is a union of sorts in $\text{Int}_{k,C}$. It follows from the first paragraph of the proof and Lemma 2.6.4 that U' is coded in $\text{Int}_{k,C}$, and hence so are U and Y , as required. \square

Part (i) of the next lemma will be crucial in Section 3.

Lemma 2.6.6 (i) *Every definable R -subtorsor of K^n is coded in \mathcal{G} .*

(ii) *If C is any set of parameters, and A is any C -definable R -submodule of K^n , then the elements of $\text{red}(A)$ are coded in $\text{Int}_{k,C}$.*

Proof. (i) Let A be a definable subtorsor of K^n . By Lemma 2.2.6, there is an \emptyset -definable R -submodule of K^{n+1} interdefinable with A , so we may reduce to the case when A is an R -submodule (of K^n).

Next, we reduce to the case when A contains no K -vector spaces of dimension greater than 0. Let $U := \{a \in A : Ka \subset A\}$. Then U is a K -subspace of K^n . Pick a basis I_0 for U , and find a subset I_1 of the standard basis of K^n such that $I_0 \cup I_1$ is a basis of K^n . Let U' be the subspace of K^n generated by I_1 . Then $K^n = U \oplus U'$; let $\pi : K^n \rightarrow U'$ be the corresponding projection. Since I_1 is chosen from the standard basis, $\pi(A)$ is $\ulcorner A \urcorner$ -definable. Also $U \subset A$, so we have $A = \pi^{-1}(\pi(A))$. Thus, it suffices to code $\pi(A) \subset U'$. However, U' is \emptyset -definably isomorphic to K^m for some m . Since $\pi(A)$ contains no positive-dimensional K -spaces, we have made the reduction.

Now let $B := \{a \in K^n : \mathcal{M}a \subset A\}$. Then by the last paragraph, B contains no copies of K . Furthermore, for any $c \in B$, $\{r \in K : rc \in B\}$ is a definable R -module, and is of the form γR for some $\gamma \in \Gamma$. Thus, B has no direct summand isomorphic to \mathcal{M} . Thus, by Lemma 2.2.4, B is definably R -isomorphic to R^ℓ for some $\ell \leq n$. Let KB be the K -subspace of K^n generated by B , so $\dim_K(KB) = \ell$. There is a coordinate projection $\pi : K^n \rightarrow K^\ell$ which is injective on KB . Now KB is coded in \mathcal{G} by elimination of imaginaries in the *pure* algebraically closed field K . Also, $B = KB \cap \pi^{-1}(\pi(B))$ and $A = KB \cap \pi^{-1}(\pi(A))$. Thus, since $\ulcorner B \urcorner \in \text{dcl}(\ulcorner A \urcorner)$, we may if necessary replace A and B by their images $\pi(A)$ and

$\pi(B)$. Hence we may suppose that $n = \ell$, so B is definably R -isomorphic to R^n , that is, $B \in S_n$.

Now $\mathcal{M}B \subset A$, and A is determined by $\mathcal{M}B \in T_n$ and the image iA of A in $B/\mathcal{M}B$. Since $B \in S_n \subset \mathcal{G}$, and $B \in \text{dcl}(\ulcorner A \urcorner)$, it suffices to show that iA is coded in \mathcal{G} . The latter is a subspace of $\text{red}(B)$, so the result follows from Lemma 2.6.4.

(ii) We apply the proof of (i) to A . If U, π are as in the second paragraph, then elements of $\text{red}(A)$ are interdefinable with elements of $\text{red}(\pi(A))$. Thus, we may suppose that A has no K -subspaces of positive dimension. Part (ii) then follows, for $B \in S_n \cap \text{dcl}(C)$, and $\mathcal{M}A \subset A \subset B$, so $\text{red}(A) \subset \text{red}(B) \subset \text{Int}_{k,C}$. \square

Finally, we give a lemma which is required not for the proof of elimination of imaginaries, but for Proposition 3.4.11.

Lemma 2.6.7 *Let $C \subset K^{\text{eq}}$ and let $e \in \text{dcl}(C) \cap S_n$, and $V = \text{red}(e)$. Suppose that $F \subset V$ is C -definable and finite. For each $a \in F$, let $s_a \in \text{dcl}(Ca) \cap S_m$ (so the map $a \mapsto s_a$ is C -definable). Then $\text{red}(s_a) \subset \text{dcl}(C \cup \text{Int}_{k,C})$ for each $a \in F$.*

Proof of Lemma 2.6.7. Let $F = \{a_1, \dots, a_m\}$.

Claim. We may suppose that F is linearly independent over k .

Proof of Claim. First, we may suppose that $0 = 0_V \notin F$, since otherwise s_0 is C -definable, so $\text{red}(s_0) \subset \text{Int}_{k,C}$. We shall show that for some n' there is C -definable $e' \in S_{n'}$ and a C -definable injection $h : V \rightarrow V'$ (where $V' := \text{red}(e')$), such that $h(F)$ is a linearly independent subset of V' . We shall show, by induction on ℓ , that for every $\ell \leq |F|$ there are $n' = n'_\ell$, $V' := V'_\ell$ and $h := h_\ell : V \rightarrow V'$ such that every ℓ -subset of $h_\ell(F)$ is linearly independent. To construct these, suppose inductively that for all $\ell' < \ell$, every ℓ' -subset of F is linearly independent. Put $V' := V \times (V \otimes V)$, and $h_\ell(v) := (v, v \otimes v)$. It is left to the reader to check that every ℓ -subset of $h_\ell(F)$ is linearly independent.

Given the claim, let A be the R -submodule of $K^n \otimes K^m$ generated by

$$\{v \otimes w : v \in e \text{ with } \text{red}(v) = a \in F, w \in s_a\}.$$

Clearly A is C -definable. Put $F = \{a_1, \dots, a_t\}$, and for each i choose $c_i \in e$ with $a_i = c_i + \mathcal{M}e$. Then c_1, \dots, c_t are K -linearly independent: for given a dependence relation $d_1c_1 + \dots + d_tc_t = 0$ with $d_1, \dots, d_t \in K$, one can arrange that all the d_i lie in R but some d_i does not lie in \mathcal{M} , and by the reduction map this contradicts k -linear independence of F . Let A' be the R -module generated by $\bigcup_{i=1}^t (c_i \otimes s_{a_i})$ and A'' be generated by $\mathcal{M}e \otimes \sum_{i=1}^t s_{a_i}$. Then A is generated as an R -module by $A' + A''$, and $\mathcal{M}A = \mathcal{M}A'' = A''$. Also, as c_1, \dots, c_t are K -linearly independent, each element of A' is *uniquely* expressible in the form $\sum_{i=1}^t c_i \otimes u_i$ for $u_i \in s_{a_i}$. It follows that $A' \cap A'' \subseteq \mathcal{M}A'$. Hence the map $A'/\mathcal{M}A' \rightarrow (A' + A'')/\mathcal{M}(A' + A'') = A/\mathcal{M}A$ given by $x + \mathcal{M}A' \mapsto x + \mathcal{M}A$ is an isomorphism.

Let $\varphi_i : s_{a_i} \rightarrow A/\mathcal{M}A$ be the homomorphism given by $y \mapsto (c_i \otimes y) + \mathcal{M}A$. Now for $y \in s_{a_i}$, $c_i \otimes y \in \mathcal{M}A \Leftrightarrow c_i \otimes y \in \mathcal{M}A' \Leftrightarrow y \in \text{red}(s_{a_i})$, so φ_i has kernel $\mathcal{M}s_{a_i}$. The map φ_i clearly does not depend on the choice of the c_i , so is $(C \cup \text{Int}_{k,C})$ -definable, as $a_i \in F \subset V \subset \text{Int}_{k,C}$. Thus, $\text{red}(s_{a_i}) \subset \text{dcl}(C \cup \text{Int}_{k,C} \cup \text{red}(A))$ for each i . Since A is C -definable, $\text{red}(A) \subset \text{dcl}(C \cup \text{Int}_{k,C})$ by Lemma 2.6.6(ii), and the claim follows. \square

3 Elimination of imaginaries

3.1 Quantifier elimination for the geometric sorts

We shall need repeatedly a notion of *generic basis* for an element of S_n . In fact, we need a *generic sequence of bases* for a sequence (s_1, \dots, s_m) of lattices, but as $s_1 \times \dots \times s_m$ is a lattice and $\text{red}(s_1) \times \dots \times \text{red}(s_m)$ is naturally isomorphic to $\text{red}(s_1 \times \dots \times s_m)$, we focus on a single lattice $s \in S_n$, and work over a set C of parameters, with $s \in \text{dcl}(C)$.

Let $B(s) := \{a \in (K^n)^n : a = (a_1, \dots, a_n), s = Ra_1 + \dots + Ra_n\}$, the set of all bases of s . We shall describe an invariant extension q_s of the partial type $B(s)$ over C . As $\text{red}(s)^n$ is a definable set of Morley rank n^2 and degree 1 in the structure $\text{Int}_{k,C}$, it has a unique generic type (in the sense of stability theory) $q_{\text{red}(s)^n}$. If $C' \supset C$ then $a = (a_1, \dots, a_n) \models q_s|C'$ if and only if $(\text{red}(a_1), \dots, \text{red}(a_n)) \models q_{\text{red}(s)^n}|C'$. To show that q_s is complete, choose $C' \supset C$ such that there is a C' -definable isomorphism $s \rightarrow R^n$. Then we may suppose $s = R^n$. Now $q_{\text{red}(s)^n}|C'$ is just the type over C' of a generic element $(\beta_1, \dots, \beta_{n^2})$ of k^{n^2} . It follows easily from Remark 2.3.5(ii) that for such a sequence, any two tuples (b_1, \dots, b_{n^2}) , where $\text{res}(b_i) = \beta_i$ for each i , have the same type. This gives completeness of q_s , and, along with the invariance of $q_{\text{red}(s)^n}$, yields invariance of q_s . We call a realisation of q_s a *generic resolution of s* , and also talk of a generic resolution of a sequence s_1, \dots, s_m of lattices, over a given set of parameters. The order of the sequence is irrelevant.

Remark 3.1.1 In the notation above, if $D \supset C$ and $a \models q_s|D$, then $\Gamma(D) = \Gamma(Da)$. To see this, suppose $\gamma \in \Gamma(Da)$, so there is a D -definable function f with $f(a) = \gamma$. Choose $\gamma' \equiv_D \gamma$. We must show $\gamma' = \gamma$, for then γ is definable over D . Choose a model $M \supset D$ with $a \models q_s|M$ (to do this, first choose an arbitrary model $M' \supset D$ and $a' \equiv_D a$ with $a' \models q_s|M'$, and apply an automorphism over D taking a' to a). Then as in the last paragraph, a is interdefinable over M with a generic sequence of length n^2 from the closed 1-torsor R . Each of the n^2 steps does not add to the value group, by Lemma 2.5.5, so $\Gamma(Ma) = \Gamma(M)$. Hence $\gamma \in M$, so in particular, $a \models q_s|D\gamma$. Now choose $a' \equiv_D a$ with $a' \models q_s|D\gamma\gamma'$. Then $a\gamma \equiv_D a'\gamma' \equiv_D a'\gamma$. Since f is D -definable and $f(a) = \gamma$, $f(a') = \gamma' = \gamma$, as required.

The language $\mathcal{L}_{\mathcal{G}}$ in which we prove elimination of imaginaries is richer than \mathcal{L}_{div} in order to give quantifier elimination with the sorts \mathcal{G} . We emphasise that the geometric sorts and the functions and relations of the language are definable in the standard language \mathcal{L}_{div} for valued fields, hence all of the results of Section 2 remain valid for our new language. As usual, there is a sort K consisting of field elements (but if the field is called F , say, we shall refer to *the sort F*), and the usual ring language $(+, -, \cdot, 0, 1)$ on it. We have a sort Γ for the value group with an extra symbol 0 . We write the operation as multiplication, and the language is the language of ordered groups, with an additional constant symbol $(\cdot, ^{-1}, 1, <, 0)$. Here, 0 is not part of the group structure, and for $x \in \Gamma \setminus \{0\}$ we assume $0 < x$, and $0 \cdot x = x \cdot 0 = 0$. There is also a sort for the residue field k , with the usual ring language $(+, -, \cdot, 0, 1)$. In general, the use of the same symbols for the languages on the different sorts should not cause confusion. For functions between these sorts there is the valuation function $|\cdot| : K \rightarrow \Gamma$ (the ultrametric axiom becomes $|x + y| \leq \text{Max}\{|x|, |y|\}$ for $x, y \in K$, and we have $|0| = 0$), and the residue function $\text{Res} : K \times K \rightarrow k$, where $\text{Res}(x, y)$ is the residue in k of x/y , and is 0 if $|x| > |y|$. We saw in Theorem 2.1.1(iii) that the theory of K eliminates quantifiers in this three-sorted language $\mathcal{L}_{\Gamma k}$.

In addition, we have the sorts S_n and T_n for all $n > 0$, with additional structure on them. We first describe relations involving just k and \mathcal{T} . For each $s \in S_n$, $\tau_n^{-1}(s)$ has definably the structure of a k -vector space. For the moment consider the structure N consisting of k , and for each n the sort T_n and the map $\tau_n : T_n \rightarrow S_n$ with each S_n regarded as a pure set. For each n we also have a k -vector space structure (given from the original structure K) on each fibre of τ_n , with addition given by a partial function $+ : T_n^2 \rightarrow T_n$, and a scalar multiplication $k \times T_n \rightarrow T_n$, interpreted naturally and in the signature of N . For any choice $n_1, \dots, n_r \in \mathbf{N}$ and s_1, \dots, s_r with $s_i \in S_{n_i}$, put $V_i := \tau^{-1}(s_i)$ for each $i = 1, \dots, r$. If we choose a basis for each V_i , then $k \times V_1 \times \dots \times V_\ell$ becomes identified with a power of k , so we can talk of the Zariski closed subsets of powers of $k \times V_1 \times \dots \times V_\ell$. Furthermore, the notion *Zariski closed* is independent of the choice of bases. For the structure on $k \cup \mathcal{T}$, we impose relation symbols (involving lattice variables for s_1, \dots, s_ℓ , uniformly) for all the Zariski closed sets in $k \times V_1 \times \dots \times V_\ell$ which are $\{s_1, \dots, s_\ell\}$ -definable in the above structure N . We will not have symbols for the vector space addition and scalar multiplication.

This structure on $k \times V_1 \times \dots \times V_\ell$ (over s_1, \dots, s_ℓ) has quantifier elimination in this language (uniformly in the s_i), as a result of the usual quantifier elimination for algebraically closed fields. For suppose C is a projection of a quantifier-free \emptyset -definable set. Then C is \emptyset -definable in this structure, so its Zariski closure \bar{C} is \emptyset -definable, and $\bar{C} \setminus C$ is \emptyset -definable of lower Zariski dimension. Hence, by induction, \bar{C} is a Boolean combination of \emptyset -definable Zariski closed sets, and hence so is C . It follows in particular that if V is one of the vector spaces, and t_1, \dots, t_n is a basis for V , and $\alpha_1, \dots, \alpha_n \in k$ with $t = \sum_{i=1}^n \alpha_i t_i$, then the atomic type of either (t_1, \dots, t_n, t) or $(t_1, \dots, t_n, \alpha_1, \dots, \alpha_n)$ determines that of

$(t_1, \dots, t_n, t, \alpha_1, \dots, \alpha_n)$.

In the language for $K \cup \Gamma \cup k \cup \mathcal{S} \cup \mathcal{T}$ there is, for each n , a relation $\in_n \subset K^n \times S_n$, defined by $\in_n(a_1, \dots, a_n, s)$ if and only if $(a_1, \dots, a_n) \in s$. We have also the functions $\tau_n : T_n \rightarrow S_n$ defined by $\tau_n(t) = s$ if and only if $t \in \text{red}(s)$ and partial functions $\nu_n : K^n \times S_n \rightarrow T_n$ defined by $\nu_n(a, s) = a + \mathcal{M}s$ if and only if $a \in_n s$.

We are not quite finished with the language. Suppose $\varphi(X_1, \dots, X_r)$ is an atomic formula where each X_i is an n_i^2 -tuple of field variables (possibly some other variables are not listed). We introduce a new relation symbol ${}^*\varphi$ with the same other variables, where ${}^*\varphi(s_1, \dots, s_r)$ holds if and only if $\varphi(a_1, \dots, a_r)$ holds for (a_1, \dots, a_r) any generic resolution of (s_1, \dots, s_r) over the other parameters. This is well-defined because the type $q_{s_1 \times \dots \times s_r}$ is complete. The above symbols together constitute the language $\mathcal{L}_{\mathcal{G}}$.

Theorem 3.1.2 *The theory of algebraically closed valued fields in the sorts K, k, Γ, S_n, T_n (for $n > 0$) has elimination of quantifiers in $\mathcal{L}_{\mathcal{G}}$.*

Proof. We show that the quantifier-free type of any finite set F implies the complete type. We may suppose F is closed under the τ_n , and for ease of notation we suppose it contains a single lattice $s \in S_n$. We add parameters from K^{n^2} for a generic resolution $a = (a_1, \dots, a_n)$ for s (so $a_i \in K^n$ for each i). If $\psi(x)$ is any atomic formula, then $\psi(a)$ holds if and only if ${}^*\psi(s)$ holds. Since ${}^*\psi$ is in the language, the quantifier-free type of F (listed as a tuple) implies that of Fa . In the structure generated by Fa (using the ν_n), $\text{red}(s)$ has a basis t_1, \dots, t_n , where $t_i := a_i + \mathcal{M}s$. If $t \in \text{red}(s) \cap Fa$, then $t = \sum_{i=1}^n \alpha_i t_i$ for some $\alpha_1, \dots, \alpha_n \in k$, and the atomic type of Fa determines that of $F_1 := Fa\alpha_1, \dots, \alpha_n$, as noted in the discussion before the theorem. If F_2 is obtained from F_1 by deleting the elements of $\text{red}(s)$, and F_3 is obtained from F_2 by deleting s , then F_3 is in $K \cup k \cup \Gamma$. By the quantifier elimination for these sorts in Theorem 2.1.1(iii), the quantifier-free type of F_3 determines its complete type. Since s is definable over F_3 , and the elements of $\text{red}(s)$ are definable over F_2 , the atomic type of F_3 determines the complete type of F_1 , and hence of F . \square

In fact, Theorem 3.1.2 is not used in our proof of elimination of imaginaries, though ideas from its proof are.

3.2 Preliminaries on coding

The following lemma is central to our proof of elimination of imaginaries.

Lemma 3.2.1 *Let M be a sufficiently saturated homogeneous structure (in a sorted language, with at least one \emptyset -definable symbol which for convenience we will write ∞), and suppose that M has an $\text{Aut}(M)$ -invariant family \mathcal{V} of definable sets with the following property: for every $a \in M^n$ there is a sequence (a_1, \dots, a_m) from M^{eq} such that $\text{dcl}(a) = \text{dcl}(a_1, \dots, a_m)$ and for each $i \leq m$, there is an*

$a_1 \dots a_{i-1}$ -definable set U which contains a_i and is in $a_1 \dots a_{i-1}$ -definable bijection with some $U' \in \mathcal{V}$. Suppose also that whenever g is a definable function from a set U' in \mathcal{V} to M , the function g is coded in M over $\ulcorner U' \urcorner$. Then all definable subsets of M^n are coded in M , so $\text{Th}(M)$ has elimination of imaginaries.

Proof. Let $X \subset M^n$ be a definable set which needs to be coded in M (over an arbitrary base set B of parameters). By assumption, for each element a of X , there is a tuple (a_1, \dots, a_m) from M^{eq} such that $\text{dcl}(a) = \text{dcl}(a_1, \dots, a_m)$ and for each $i \leq m$, there is a pair (U, U') as in the statement of the lemma. By compactness, we can assume m is independent of a and the ‘coding’ is uniform. Let X' be the definable set of such tuples. Then $\text{dcl}(\ulcorner X \urcorner) = \text{dcl}(\ulcorner X' \urcorner)$, so $\text{dcl}(B \ulcorner X \urcorner) = \text{dcl}(B \ulcorner X' \urcorner)$ and so it suffices to code X' over B . We argue by induction on m . So we assume that over any base set F of parameters, if $l < m$ and X^* is a definable set of l -tuples (b_1, \dots, b_l) from M with each b_i in an $(F \cup \{b_j : j < i\})$ -definable set which is in $(F \cup \{b_j : j < i\})$ -definable bijection with a member of \mathcal{V} , then X^* is coded in M over F .

To start the induction, observe that if $m = 1$ then by compactness there are finitely many B -definable sets U_1, \dots, U_r , each U_i in B -definable bijection with some $U'_i \in \mathcal{V}$, such that $X' \subset U_1 \cup \dots \cup U_r$. Let $g_i : U_i \rightarrow U_i \cup \{\infty\}$ be given by $g_i(x) = x$ if $x \in X' \cap U_i$, and $g_i(x) = \infty$ otherwise. By assumption (and using the bijections), each function g_i is coded over B by some sequence e_i from M , and (e_1, \dots, e_r) codes X' over B .

Now assume the result for $m - 1$. Let Y be the set of first coordinates of tuples from X' . Each such a_1 lies in an \emptyset -definable set in \emptyset -definable bijection with a member of \mathcal{V} , and again by compactness, we can assume they all lie in the same such set U . For each $a \in Y$, let $X'(a) := \{x : (a, x) \in X'\}$. By induction, each $X'(a)$ is coded in M over Ba by a sequence $c_a = (c_a^1, \dots, c_a^l) \in M^l$. By compactness, we may suppose l is fixed. By assumption, each coordinate function $a \mapsto c_a^i$ is coded over B , and a tuple listing these codes is a code for X' over B (in M).

The final assertion of the lemma (elimination of imaginaries for $\text{Th}(M)$) now follows by saturation of M . \square

Remark 3.2.2 Lemma 3.2.1 is a refinement of an earlier version, which has the following easier statement, and a similar proof. Let M be a structure, and $\{R_i : i \in I\}$ be a collection of sorts from M^{eq} , with $R_0 = M$ and $\mathcal{R} := \bigcup_{i \in I} R_i$. Assume that for every definable subset U of M , every $i \in I$, and every definable function $f : U \rightarrow R_i$, the pair (U, f) is coded by some tuple from \mathcal{R} . Then every element of M^{eq} is coded in \mathcal{R} .

Proof of Remark 3.2.2. We show by induction that every n -ary relation on M is coded. The case $n = 1$ holds by assumption. Suppose that $X \subset M^{n+1} = M \times M^n$ is definable, and let Y be the projection of X to the first coordinate.

For each $a \in Y$, let $X(a) := \{x : (a, x) \in X\}$. By the inductive assumption, each $X(a)$ is coded by some tuple $h(a)$ in \mathcal{R} . By compactness, the function h is definable, and Y can be partitioned into finitely many pieces U_1, \dots, U_k such that for each $i = 1, \dots, k$, $h|_{U_i}$ is to some product of the R_i (and i is determined by the product). By assumption, each pair $(U_i, h|_{U_i})$ is coded by some tuple c_i in \mathcal{R} . Now X is coded by (c_1, \dots, c_k) . \square

In our case, the structure, of course, is an algebraically closed valued field with sorts \mathcal{G} , and we take the family \mathcal{V} of Lemma 3.2.1 to consist of those members of the following family \mathcal{U} which are definable, i.e. which satisfy (i) or (ii) below.

Definition 3.2.3 Let \mathcal{U} be the collection of unary sets of the following kinds:

- (i) intervals in Γ ,
- (ii) for definable R -submodules $M < N$ of K^n , 1-torsors which are cosets of N/M in K^n/M ;
- (iii) the ∞ -definable 1-torsors which are intersections of chains of 1-torsors in (ii).

Lemma 3.2.4 *The collection \mathcal{V} of definable unary sets in \mathcal{U} satisfies the hypotheses of Lemma 3.2.1.*

Proof. Apply Proposition 2.3.10 and its proof. The main point is that, in Step 2 (using its notation), there is a bijection, defined over $\ulcorner B_{n-1} \urcorner, \ulcorner A'_{n-1} \urcorner, \ulcorner B_1 \urcorner, \ulcorner A''_1 \urcorner$, between $D' = \text{Hom}_R(B_1/A''_1, B_{n-1}/A'_{n-1})$ and a member of \mathcal{V} . Here, $A''_1 < B_1$ are R -submodules of K , and $A'_{n-1} < B_{n-1}$ are R -submodules of K^{n-1} . To see the existence of this bijection, note that there is some $\alpha < 1$ such that each of B_{n-1}/A'_{n-1} and B_1/A''_1 is isomorphic (not canonically) to $R' = R/\alpha R$. Then D' is *canonically* isomorphic to $\text{Hom}(B_1, B'_{n-1}/A'_{n-1})$, where $B'_{n-1} = \{x \in B_{n-1} : ax \in A'_{n-1} \text{ whenever } |a| = \alpha\}$. Write $B_1 := d^{-1}R$ (so d is not canonical, but $|d|$ is). Then there is also a canonical isomorphism $\varphi : dB'_{n-1}/dA'_{n-1} \rightarrow \text{Hom}(B_1, B'_{n-1}/A'_{n-1})$: for $x \in B'_{n-1}$, $\varphi(dx + dA'_{n-1})(y) = dxy + A'_{n-1}$. Thus, D' and dB'_{n-1}/dA'_{n-1} are canonically isomorphic, over $\ulcorner B_{n-1} \urcorner, \ulcorner A'_{n-1} \urcorner, \ulcorner B_1 \urcorner, \ulcorner A''_1 \urcorner$. Finally, observe that $dB'_{n-1}/dA'_{n-1} \in \mathcal{V}$. \square

Lemma 3.2.5 (i) *Any definable subtorsor of a torsor in \mathcal{U} is coded in \mathcal{G} .*

(ii) *Any element of a torsor in \mathcal{U} is coded in \mathcal{G} .*

Proof. This is immediate from Lemma 2.6.6(i). \square

Remark 3.2.6 1. By Lemma 3.2.5 and Remark 2.3.2, for unary sets in \mathcal{U} we can apply all results from Sections 2.3 and 2.5 over a base C with $C = \text{acl}_{\mathcal{G}}(C)$ (rather than $C = \text{acl}(C)$).

2. It would be possible to use Remark 3.2.2 rather than Lemma 3.2.1, and prove elimination of imaginaries by coding functions from K rather than from any

unary set. This would be marginally simpler, but the methods given here have other uses; in particular, they support Lemma 3.4.13, which will be important in [3].

3.3 Coding of functions

In this section we prove coding results for definable functions on unary sets. The key result, Proposition 3.3.9, gives a kind of weak coding. The remaining ingredient is the coding of finite sets. This rests on Theorem 3.3.2, and is in the following section.

Suppose that T is any complete theory, $M \models T$, p is a type over M with solution set P such that p is definable over some $B \subset M$, and f is an M -definable function whose domain contains P . Suppose that $f = f_a$ is defined by the formula $\varphi(x, y, a)$ (so $f_a(x) = y$). We say that $f_a, f_{a'}$ *have the same germ on P* , or the same p -*germ*, if the formula $f_a(x) = f_{a'}(x)$ lies in p . By the definability of p , the equivalence relation ‘has the same germ’ is definable over B . Hence, the *germ of f on P* (which is defined to be the equivalence class of φ -definable functions with the same germ), lies in M^{eq} . Furthermore, up to interdefinability over B this germ is independent of the choice of φ . We shall not always mention the base set B explicitly. For example, if p is the generic type of a closed ball b , then B could be taken to be the singleton $\ulcorner b \urcorner$.

We say that a code c in M for the germ of f on P is *strong* if there is a c -definable function g such that the formula $f(x) = g(x)$ lies in p .

We also may talk of an M -definable function g having the same germ on P as f , even if g is not φ -definable. By this we mean again that the formula $f(x) = g(x)$ is in p . If the type p is not definable, then the equivalence relation ‘has the same germ on P ’ still makes sense, but we avoid talking of the ‘germ of f on P ’, as this is not an interpretable object.

Remark 3.3.1 1. Suppose that T is an arbitrary complete theory, M is a model of T , and p is a type over M defined over $B \subset M$. If f is an M -definable function with domain containing P , and c is a code for the germ of f on P , we could say the code c is *strong over B* if there is a Bc -definable function with the same germ as f on P . If T is stable, then any code for f on P is strong over B , by the following argument. We may suppose M is sufficiently saturated. Let c be a code for the germ of $f = f_a$ on P . Let $q := \text{tp}(a/Bc)$, with Q its set of realisations (in the monster model). Let $p|Bc$ denote the restriction of p to parameters in Bc , and $d \models p|Bc$. If $a, b \in Q \cap M$ with $d \downarrow_{Bc} ab$ (in the sense of stable non-forking), then $f_a(d) = f_b(d)$, and it follows that if $a, b \in Q$ and $d \downarrow_{Bc} a$ and $d \downarrow_{Bc} b$ then $f_a(d) = f_b(d)$ (for we may choose $e \in Q \cap M$ with $e \downarrow_{Bc} abd$). Hence, if $a \in Q$ and $d \downarrow_{Bc} a$, then $f_a(d) \in \text{dcl}(Bcd)$. It follows by compactness that there is a Bc -definable function g such that for all $x \models p|Bc$, and $a \in Q$ with $a \downarrow_{Bc} x$ we have $g(x) = f_a(x)$.

2. In the situation of the paper, $M = K_{\mathcal{G}}$ for an algebraically closed valued field K . If t is a definable 1-torsor, then by Lemma 2.3.8, the generic type p of t over $K_{\mathcal{G}}$ is definable over $\ulcorner t \urcorner$. If f is a definable function on t , then the *germ* of f on t is the germ of f on P (and so is an element of K^{eq}). Similarly, for a definable function on Γ and $\gamma \in \Gamma$ we can talk of the germ of f on $\{x \in \Gamma : x < \gamma\}$, meaning the germ of f on the generic type of elements of Γ immediately below γ . It turns out that if p is a generic type of Γ , or of an open 1-torsor, then germs of functions on P may not be strong. For example, let c be generic in R over \emptyset , and let $f : \mathcal{M} \rightarrow \mathcal{B}^{\text{cl}}$ be the function $x \mapsto B_{|x|}(c)$. Then the germ of f is coded by the ball $\mathcal{M} + c$, but this germ is not strongly coded. We show below that this problem does not arise for closed 1-torsors.

By Lemma 3.2.1 and Corollary 3.4.8 below, our proof of Theorem 1.0.1 reduces to showing the following: if $U \in \mathcal{V}$ and $f : U \rightarrow \mathcal{G}$ is definable, and $B = \text{acl}_{\mathcal{G}}(\ulcorner f \urcorner)$, then f is definable over B . We use compactness and consider the restriction of f to types over B . The key is the next theorem, which shows that the germ of a function on the generic type of a closed 1-torsor has a strong code. We then consider the germ of f on the generic type of U , an open 1-torsor or the intersection of a chain of 1-torsors, and approximate U from inside by closed 1-torsors, piecing together the corresponding functions defined by strong codes, to obtain Proposition 3.3.9.

Theorem 3.3.2 *Let U be a unary set in \mathcal{V} which is a closed 1-torsor. Let $f : U \rightarrow \mathcal{G}$ a definable function. Then, working over the parameter $\ulcorner U \urcorner$,*

- (i) *the germ of f on U is coded in \mathcal{G} , and*
- (ii) *the code in \mathcal{G} for the germ of f on U is strong.*

Proof. By Lemma 3.2.5 (ii), we may suppose that the elements of \mathcal{V} are sequences from \mathcal{G} , and by compactness we may suppose that this coding is uniform.

(i) We first replace U by a definable set U' of tuples with entries in $K \cup k$, to make an application of quantifier elimination easier. For convenience of notation we suppose that $U \subset S_m \cup T_m$ (the general case $U \subset \mathcal{G}^r$ is not much harder). There is an \emptyset -definable $V_1 \subset (K^m)^m$ and an \emptyset -definable surjection $h_1 : V_1 \rightarrow S_m$, such that for $(a_1, \dots, a_m) \in V_1$ we have $h_1(a_1, \dots, a_m) = a_1R + \dots + a_mR$. Also, there is \emptyset -definable $V_2 \subset V_1 \times k^m$ and an \emptyset -definable surjection $h_2 : V_2 \rightarrow T_m$, with $h_2(a_1, \dots, a_m, \beta_1, \dots, \beta_m) = \sum_{i=1}^m \beta_i(a_i + \mathcal{M}h_1(a_1, \dots, a_m))$. We may replace U by $U' = h_1^{-1}(U)$ if $U \subset S_m$ or $U' = h_2^{-1}(U)$ if $U \subset T_m$ and f by the corresponding composition. Notice that U' is no longer a unary set. If $U \subset S_m$, we say that a is generic in U' if $h_1(a)$ is generic in U and a is a generic resolution of $h_1(a)$ in the sense of Section 3.2 (over a given parameter set). Similarly, if $U \subset T_m$, we may talk of a generic element of U' . Observe that if M is a model over which U is defined, and a is generic in U' over M , then $\Gamma(M) = \Gamma(Ma)$. Also, the generic type of U' is definable. Finally, the germ of

f on U has a code in \mathcal{G} if and only if the germ of $f \circ h_1$ on U' is coded in \mathcal{G} (if $U \subset S_n$), or if and only if the germ of $f \circ h_2$ on U' is coded in \mathcal{G} (if $U \subset T_n$). We shall in fact assume $U \subset T_m$, so $U' = h_2^{-1}(U)$, as this is the more involved case. Write F for the function $f \circ h_2 : U' \rightarrow \mathcal{G}$.

We shall consider the cases when $\text{ran}(F)$ is a subset of S_n or T_n , as the cases when the range lies in K or k are easier. Let $B := \text{dcl}_G(\Gamma \text{germ}(F)^{\top} U^\top)$, let F be c -definable where c is a tuple from K , and put $q := \text{tp}(c/B)$, with solution set Q . For any $c' \in Q$, write F' for the function defined by the same formula as F with parameters c' . Observe that since U is definable over B , so is U' . Let p be the generic type over B of elements of U' .

Claim 1. Let $c' \in Q$. Then for all $a \in U'$ generic over Bcc' , we have $F(a) = F'(a)$.

This claim already yields that F, F' have the same germ on p , so the germ of F on p is definable over B , and hence (i) holds.

Proof of Claim. We fix one notational convention: if $e = (e_1, \dots, e_m) \in K^m$ and $\beta = (\beta_1, \dots, \beta_n) \in k^n$, and $P(X, Y) \in K[X, Y]$, write $|P(e, \beta)| < 1$ if for all $b_1 \in \beta_1, \dots, b_n \in \beta_n$ we have $|P(e, b_1, \dots, b_n)| < 1$. Here we do not assign a value to $|P(e, \beta)|$ – just a truth value to $|P(e, \beta)| < 1$.

Let M be any small elementary submodel of K containing Bcc' . We shall suppose $a = \bar{a}\bar{\alpha}$, where \bar{a} is a tuple in K and $\bar{\alpha}$ is a tuple in k .

Suppose first $F : U' \rightarrow S_n$. Choose a generic resolution over Ma (in the sense of Section 3.1) \bar{d} of $F(a)$ and \bar{d}' of $F'(a)$. We must show that $\bar{a}\bar{\alpha}\bar{d} \equiv_M \bar{a}\bar{\alpha}\bar{d}'$, for then $aF(a) \equiv_M aF'(a)$, which (as F and F' are M -definable) forces $F(a) = F'(a)$. If a is a generic element of U' over M , then, as observed above, $\Gamma(Ma) = \Gamma(M)$, so $\Gamma(Ma\bar{d}) = \Gamma(M)$. Hence, by the quantifier elimination for (K, k, Γ) proved in Theorem 2.1.1(iii), $\text{tp}(\bar{a}\bar{\alpha}\bar{d}/M)$ is determined by expressions of the form

$$|g(\bar{a}, \bar{d})| = \gamma \tag{1}$$

together with those of form

$$h(\text{res}(p_1(\bar{a}, \bar{d}), q_1(\bar{a}, \bar{d})), \dots, \text{res}(p_r(\bar{a}, \bar{d}), q_r(\bar{a}, \bar{d})), \bar{\alpha}) = 0.$$

Here $g(X, Y), p_i(X, Y), q_i(X, Y) \in M[X, Y]$, and $h(W, V) \in M_k[W, V]$, where $W = (W_1, \dots, W_r)$. (Recall that M_k denotes $M \cap k$.) Since $\Gamma(Ma\bar{d}) = \Gamma(M)$, we may multiply p_i, q_i by elements of M to ensure $p_i(\bar{a}, \bar{d})$ and $q_i(\bar{a}, \bar{d})$ have norm 1; hence, after multiplying out, we can replace the above equation involving h by one of the form $h'(\text{res}(p'_1(\bar{a}, \bar{d})), \dots, \text{res}(p'_s(\bar{a}, \bar{d})), \bar{\alpha}) = 0$. Lifting h' to a polynomial over R and composing with the p_i , we replace this by one of form:

$$|P(\bar{a}, \bar{d}, \bar{\alpha})| < 1 \tag{2}$$

where $P(X, Y, Z) \in M[X, Y, Z]$.

To handle expressions of the form (1), let

$$J_F := \{P(X, Y) \in K[X, Y] : \text{for generic } a \in U', \text{ generic } \bar{d} \in F(a), |P(\bar{a}, \bar{d})| \leq 1\}.$$

For each $\ell > 0$ let J_F^ℓ consist of the polynomials in J_F of total degree at most ℓ . If we identify each member of J_F^ℓ with a tuple of coefficients, we see that J_F^ℓ is an R -module. Define $J_{F'}, J_{F'}^\ell$ correspondingly. Now J_F^ℓ is definable and so is coded in \mathcal{G} , by Lemma 2.6.6; hence, as $\text{dcl}_{\mathcal{G}}(B^\Gamma \text{germ}(F)^\Gamma) = B$, J_F^ℓ is definable over B . It follows that any B -automorphism taking F to F' fixes the J_F^ℓ , so $J_F^\ell = J_{F'}^\ell$ for each ℓ .

Now suppose a is generic in U' over M and $|P(\bar{a}, \bar{d})| = \varepsilon > 0$, where $P(X, Y) \in M[X, Y]$. Pick $e \in M$ with $|e| = \varepsilon$. Then $e^{-1}P(X, Y) \in J_F^\ell$ for some ℓ , whence $e^{-1}P(X, Y) \in J_{F'}^\ell$, so $|P(\bar{a}, \bar{d}')| \leq \varepsilon$. Reversing F, F' we get $|P(\bar{a}, \bar{d}')| = \varepsilon$. If $|P(a, F(a))| = 0$, apply a similar argument; note here that the set of polynomials $P \in K[X, Y]$ of degree at most n such that $|P(a, F(a))| = 0$ for generic $a \in U'$ corresponds to a definable set in K (as a pure algebraically closed field); hence it is coded in \mathcal{G} .

For expressions of the form (2), argue similarly. This time, we define J'_F to consist of

$$\{P(X, Y, Z) \in K[X, Y, Z] : \text{for generic } a \in U', \text{ generic } \bar{d} \in F(a) (|P(\bar{a}, \bar{d}, \bar{\alpha})| < 1)\}.$$

Again, J'_F is a collection of definable modules, each coded in \mathcal{G} (by Lemma 2.6.6) and so definable over B , so $J'_F = J'_{F'}$. It follows that if $P \in M[X, Y, Z]$, then

$$|P(\bar{a}, \bar{d}, \bar{\alpha})| < 1 \leftrightarrow |P(\bar{a}, \bar{d}', \bar{\alpha})| < 1.$$

Thus, in the case when F is a map to S_n , $aF(a) \equiv_M aF'(a)$, as required.

Suppose next that F is a map to T_n . Now $\tau_n \circ F$ is a map $U' \rightarrow S_n$, so by the above, $\tau_n \circ F(a) = \tau_n \circ F'(a)$ for a generic in U' over M . Thus, by the above case we may suppose there is a map $g : U' \rightarrow S_n$ with germ definable over B , so that for all $a \in U'$, $F(a) \in \text{red}(G(a))$, where $G(a)$ is the lattice coded by $g(a)$. We must show that $\text{tp}(\bar{a}\bar{\alpha}g(a)F(a)/M)$ is determined as above by definable modules.

We fix some notation. Suppose now $g \in S_n$ and $h \in \text{red}(g)$. Given a basis \bar{d} for g with induced basis $\text{red}(\bar{d}) := (\text{red}(d_1), \dots, \text{red}(d_n))$ for $g/\mathcal{M}g$, let $\lambda = (\lambda_1, \dots, \lambda_n) = \lambda(\bar{d}, h)$ be the unique element of k^n such that $\sum_{i=1}^n \lambda_i \text{red}(d_i) = h$. Now let J^F be the set of polynomials $P(X, W, \Xi, \Lambda) \in K[X, W, \Xi, \Lambda]$ such that for generic $a \in U'$ (over K) and generic basis \bar{d} for $G(a)$, and for $\lambda = \lambda(\bar{d}, F(a)) \in k^n$, we have $|P(\bar{a}, \bar{d}, \bar{\alpha}, \lambda)| < 1$. Then J^F is a collection of R -modules which are coded in \mathcal{G} (by Lemma 2.6.6) and definable over $\text{germ}(F)$ and hence over B . Thus, $J^F = J^{F'}$.

As in the argument around (2) above, we show that

$$\text{tp}(\bar{a}\bar{\alpha}g(a)F(a)/M) = \text{tp}(\bar{a}\bar{\alpha}g(a)F'(a)/M).$$

Let \bar{d} be a generic basis of $G(a)$, let $\lambda := \lambda(\bar{d}, F(a))$ and $\lambda' := \lambda(\bar{d}, F'(a))$. It suffices (by the quantifier elimination for K, k, Γ) to show that for any tuple of polynomials $H \in M[X, Y]^\ell$ and any $h \in M_k[W_1, \dots, W_\ell, V]$, we have $h(\text{res}(H(a, b), \delta)) = 0$ if and only if $h(\text{res}(H(a, b), \delta')) = 0$. Lifting h to a polynomial over R , and composing with H , we find that this follows from the equality $J^F = J^{F'}$.

It remains (for Claim 1) to check that if a is generic in U' over Bcc' (rather than over M), then $F(a) = F'(a)$. However, if this is false for some a , then we may choose a' generic in U' over M such that $F(x) \neq F'(x)$. We have just shown that the latter is impossible. This finishes the proof of Claim 1, and hence of (i).

(ii) We now need to show that the germ of f on U is strongly coded (with the original f, U of the theorem). For $c, c' \in Q$, let $A(c, c') := \{x \in U : f_c(x) \neq f_{c'}(x)\}$. Suppose that U is a torsor of the definable 1-module A , and recall that $\text{red}(U) := \{x + \mathcal{M}A : x \in U\}$. Then $\text{red}(U)$ is a strongly minimal set. Let

$$Z(c, c') := \{u \in \text{red}(U) : u \cap A(\bar{c}, \bar{c}') \neq \emptyset\}.$$

Since $\text{red}(U)$ is strongly minimal, it follows from the last paragraph that $Z(c, c')$ is finite. Also, for any $c, c', c'' \in Q$,

$$A(c, c') \subset A(c, c'') \cup A(c', c''), \text{ so}$$

$$Z(c, c') \subset Z(c, c'') \cup Z(c', c'').$$

Claim 2. Let $c' \in Q$, and $a \in U$ be generic over Bc and over Bc' . Then $f_c(a) = f_{c'}(a)$.

Proof. Let $z \in \text{red}(U)$ with $a \in z$. We shall show there is $c'' \in Q$ with $z \notin \text{acl}(Bcc'') \cup \text{acl}(Bc'c'')$. For then $z \notin Z(c, c'') \cup Z(c', c'')$, so $a \notin A(c, c'') \cup A(c', c'')$, so $a \notin A(c, c')$, and hence $f_c(a) = f_{c'}(a)$.

To find $c'' = (c_1, \dots, c_n)$, we inductively find c_i in the required type, R_i say, over $\text{acl}_G(Bc_1 \dots c_{i-1})$. To start, we clearly have $z \notin \text{acl}(Bc) \cup \text{acl}(Bc')$, and we may suppose

$$z \notin \text{acl}(Bcc_1 \dots c_{i-1}) \cup \text{acl}(Bc'_1 \dots c_{i-1}).$$

We may suppose that $c_i \notin \text{acl}(Bc_1 \dots c_{i-1})$, as otherwise there is no problem. Choose c_i generically (in R_i) over $Bcc'_1 \dots c_{i-1}z$. We apply Lemma 2.5.3 twice, both times with $C_0 := \text{acl}_G(Bc_1 \dots c_{i-1})$, the first time with $C := C_0c$ and the second time with $C := C_0c'$. This gives $z \notin \text{acl}(Bcc_1 \dots c_i) \cup \text{acl}(Bc'_1 \dots c_i)$, as required.

For any B -conjugate f' of f , let $\ulcorner f' \urcorner$ denote the corresponding canonical parameter. It follows from Claim 2 by compactness that there is a $B^\ulcorner f' \urcorner$ -definable set $W(\ulcorner f' \urcorner) \subset U$ (a finite union of elements of $\text{red}(U)$), such that for any conjugate f' of f over B , if $a \in U \setminus (W(\ulcorner f' \urcorner) \cup W(\ulcorner f' \urcorner))$ then $f(a) = f'(a)$. Now define a function g on a subset of U as follows: if $x \in U$ and there is $f' \equiv_B f$ with

$x \notin W(f')$, define $g(x) = f'(x)$. By the above, this definition is independent of the choice of f . Also, since $W(f)$ is a finite union of elements of $\text{red}(U)$, g is definable, so as g is B -invariant, g is B -definable by compactness. Clearly g and f have the same germ on U . There is a tuple d in B such that both g and the germ of f on U are B -definable, and such d is a strong code for the germ of f . \square

Remark 3.3.3 The proof in Part (i) generalises. Suppose that f is a definable function with domain $X \subset G^n$, and that p is a C -definable type over K whose realisations lie in X . Suppose also that for any model M containing $C, \ulcorner f \urcorner$ and any $a \models p|_M$, we have $\Gamma(M) = \Gamma(Ma)$. Then the germ of f on p is coded in \mathcal{G} over C . Here, if f has range in \mathcal{G}^m , then the germ of f is interdefinable with the tuples of germs of its coordinate functions. This will be used in the proof that finite sets are coded (Theorem 3.4.1, via Lemma 3.4.7).

Next, we use the results from Section 2.4 (in particular Theorem 2.4.13, applied with $\rho\Gamma = \Gamma$) to show that definable functions from Γ are coded in \mathcal{G} .

Proposition 3.3.4 *Let $f : \Gamma \rightarrow \mathcal{G}$ be a definable function.*

- (i) *The function f is coded in \mathcal{G} .*
- (ii) *Let $\gamma_0 \in \Gamma \cup \{\infty\}$. Then the germ of f below γ_0 is coded in \mathcal{G} over γ_0 .*

The main problems in the proof arise with functions $\Gamma \rightarrow S_n$ and $\Gamma \rightarrow T_n$. We need two lemmas. We identify S_n with $B_n(K)/B_n(R)$, using Lemma 2.4.8. A definable function $\Gamma \rightarrow D_n(K)/D_n(R)$ will be called *affine* if, when $D_n(K)/D_n(R)$ is identified canonically with Γ^n , each of the n coordinate functions has the form $x \mapsto \delta_i x^{q_i}$ for $\delta_i \in \Gamma$ and $q_i \in \mathbf{Q}$. We say that a (partial) function $f : \Gamma \rightarrow S_n$ defined on an interval $I \subset \Gamma$ has *canonical form* if it has the form $x \mapsto uh(x)B_n(R)$, where $u \in U_n(K)$ and h is an affine map $I \rightarrow D_n(K)/D_n(R)$. Likewise, if $f : \Gamma \rightarrow T_n$, it has canonical form on I if for some $m \leq n$ it has the form $x \mapsto uh(x)B_{n,m}(R)$, for some affine h on I . Recall also from Section 2.4 our convention that if g, h are elements of a group G , then $g^h := hgh^{-1}$.

Lemma 3.3.5 *Let G be a soluble linear algebraic group, let I be an interval in Γ , and let $(B_\gamma : \gamma \in I)$ be a sequence of cosets of subgroups of G , such that $\delta < \gamma$ implies $B_\delta \subseteq B_\gamma$, and such that the function $\gamma \mapsto B_\gamma$ is definable (in the algebraically closed valued field K). Then $\bigcap_{\gamma \in I} B_\gamma \neq \emptyset$.*

Proof. We may suppose G is connected, and we argue by induction on the Zariski dimension of G . For the inductive step, suppose $\dim(G) \geq 2$. Since G is conjugate in $\text{GL}_n(K)$ to a subgroup of $B_n(K)$, G has a normal subgroup N such that G/N has dimension 1. By induction, the images $B_\gamma N/N$ in G/N of the B_γ have a point of intersection cN . Then $c^{-1}B_\gamma \cap N \neq \emptyset$ for each γ . So $(c^{-1}B_\gamma \cap N : \gamma \in I)$ is a definable sequence of cosets of subgroups of N , so by induction has a point of intersection d . Then $cd \in \bigcap (B_\gamma : \gamma \in I)$.

To start the induction, note that if $\dim(G) = 1$ then G is definably isomorphic to the additive group G_a or the multiplicative group G_m of K . If $G \cong G_a$ then every coset of every definable subgroup is a finite union of balls; the result follows in this case from Lemma 2.4.3. Finally, the group G_m has normal subgroups $H_1 < H_2$ with $H_1 = 1 + \mathcal{M}$, $H_2/H_1 \cong (k \setminus \{0\}, \cdot)$, and $G/H_2 \cong \Gamma$. As in the last paragraph, it suffices to prove the result for each of these quotients. The group H_2/H_1 and G/H_2 are strongly minimal and o-minimal respectively, so have no proper infinite definable subgroups. The group $1 + \mathcal{M}$ is handled like the group G_a . \square

Lemma 3.3.6 *Let $f : \Gamma \rightarrow B_n(K)/B_n(R)$ be a definable function. Then there is a unique finite sequence $\gamma_1 > \dots > \gamma_m$ in Γ such that (with $\gamma_0 := \infty$ and $\gamma_{m+1} = 0$):*

- (i) on each interval $I_j := (\gamma_{j+1}, \gamma_j)$, $f|_{I_j}$ has canonical form;
- (ii) for each $i = 0, \dots, m-1$ and for any $\delta < \gamma_{i+1}$, $f|_{(\delta, \gamma_i)}$ does not have canonical form.

Proof. The existence of some $\gamma'_1, \dots, \gamma'_{m'}$ satisfying (i) comes from Theorem 2.4.13 (iii). Now define $\gamma_1, \dots, \gamma_m$ inductively: for each $i > 0$, $\gamma_i := \inf\{\delta : f|_{(\delta, \gamma_{i-1})}$ is canonical $\}$. Then inductively, $\gamma_i < \gamma'_i$ for each i , so $\gamma_{m+1} = 0$ for some $m \leq m'$. It remains to verify the following claim, which ensures also that the γ_i are infima of definable sets, so lie in Γ not just its Dedekind completion.

Claim. If $\delta_1 > \delta_2$ lie in Γ , and $f|_{(\delta, \delta_1)}$ is canonical for all δ with $\delta_1 > \delta > \delta_2$, then $f|_{(\delta_2, \delta_1)}$ is canonical.

Proof of Claim. We may write $f(x) = u(x)h(x)B_n(R)$ on (δ_2, δ_1) , with $h(x) \in D_n(K)/D_n(R)$, and $u(x) \in U_n(K)$. Now the function $x \mapsto h(x)$ is affine on (δ, δ_1) for all $\delta > \delta_2$, so must be affine on (δ_2, δ_1) . For each $\delta > \delta_2$, there is $u_\delta \in U_n(K)$ such that for all $x \in (\delta, \delta_1)$, $u(x)h(x)B_n(R) = u_\delta h(x)B_n(R)$. Put $C(\delta) := \bigcap_{\delta < x < \delta_1} B_n(R)^{h(x)}$. The cosets $u_\delta C(\delta)$ form a decreasing chain. By Lemma 3.3.5, there is $u \in \bigcap_{\delta_2 < \delta < \delta_1} u_\delta C(\delta)$. Then $f(x) = uh(x)B_n(R)$ for $x \in (\delta_2, \delta_1)$. \square

The analogue of the last lemma also holds for functions $f : \Gamma \rightarrow B_n(K)/B_{n,m}(R)$, where $1 \leq m \leq n$.

The lemma below, and its proof, are phrased in order to handle simultaneously functions $\Gamma \rightarrow S_n$ and $\Gamma \rightarrow T_n$.

Lemma 3.3.7 *Let $n > 0$, let C be a parameter set, T be a C -definable subgroup of $D_n(K)$, and let H be one of the groups $B_n(R)$ or $B_{n,m}(R)$. Put $G(T) := \bigcap_{t \in T} H^t$. Then each element of $B_n(K)/G(T)$ is coded in \mathcal{G} over C .*

Proof. Let $\text{diag}(t)$ be the diagonal matrix with $t = (t_1, \dots, t_n) \in (K \setminus \{0\})^n$ as its sequence of diagonal entries, and put $G(t) := \text{diag}(t)B_n(R)\text{diag}(t)^{-1}$. Observe first that $G(t)$ depends only on the sequence $\gamma_1 = |t_1|, \dots, \gamma_n = |t_n|$.

Thus, for $\gamma = (\gamma_1, \dots, \gamma_n)$, we may define $G(\gamma) := G(t)$, and may write $G(\gamma) = \text{diag}(\gamma)B_n(R)\text{diag}(\gamma)^{-1}$.

We wish to treat expressions $x < \gamma$, $x \leq \gamma$, $x = 0$, $x < \infty$ (for $x \in \Gamma$) on the same footing, in order to handle functions to S_n and to T_n simultaneously. So for convenience we write $x \leq \gamma^-$ for $x < \gamma$. We write $\bar{\Gamma}$ for the set of expressions $\gamma, \gamma^-, 0, \infty$ (for $\gamma \in \Gamma$). There are obvious rules for the min and inf of a subset of $\bar{\Gamma}$, and for the product of two elements (with $0 \times \infty = 0$).

Let Σ be the set of all functions

$$\sigma : \{(i, j) : 1 \leq i < j \leq n\} \rightarrow \bar{\Gamma}$$

with the property that for $1 \leq i < j < k \leq n$ we have $\sigma(i, k) \geq \sigma(i, j) \times \sigma(j, k)$. For $\sigma \in \Sigma$, put

$$G[\sigma] := \{(a_{ij}) \in B_n(K) : |a_{ii}| = 1, |a_{ij}| \leq \sigma(i, j) \text{ for each } i < j\}.$$

Observe the following.

- (a) If $\sigma \in \Sigma$ then $G[\sigma]$ is a subgroup of $B_n(K)$.
- (b) If Φ is a C -definable subset of Σ , define $\text{inf}(\Phi)$ by $\text{inf}(\Phi)(i, j) := \text{inf}\{\sigma(i, j) : \sigma \in \Phi\}$ (these infima are defined, by o-minimality of Γ). Then $\text{inf}(\Phi) \in \Sigma$, and if Φ is C -definable, so is $\text{inf}(\Phi)$. Also, $G[\text{inf}(\Phi)] := \bigcap_{\sigma \in \Phi} G[\sigma]$.
- (c) If $\gamma = (\gamma_1, \dots, \gamma_n) \in \bar{\Gamma}^n$ and $\sigma \in \Sigma$ with $\sigma(i, j) = \gamma_i \gamma_j^{-1}$ for each $i < j$, then $G[\sigma] = G(\gamma)$.
- Now, let Σ^c be the collection of C -definable $\sigma \in \Sigma$ such that all elements of $B_n(K)/G[\sigma]$ are coded.
- (d) If $\sigma \in \Sigma$ only takes values in $\{0, \infty\}$, then $G[\sigma]$ is an algebraic group, so $\sigma \in \Sigma^c$ by elimination of imaginaries for algebraically closed fields.
- (e) If $\sigma_1, \dots, \sigma_\ell \in \Sigma^c$, then all elements of $B_n(K)/\bigcap_{i=1}^\ell G[\sigma_i]$ are coded: indeed, $B_n(K)/\bigcap_{i=1}^\ell G[\sigma_i]$ embeds naturally into $B_n(K)/G[\sigma_1] \times \dots \times B_n(K)/G[\sigma_\ell]$.
- (f) Suppose that $\sigma \in \Sigma$ is C -definable and $\sigma(i, j) = 1$ or $\sigma(i, j) = 1^-$ for each $i < j$. Then $\sigma \in \Sigma^c$. Indeed, let $B^* := B_{n,1}(R) \cap \dots \cap B_{n,n}(R)$. Then $B_n(K)/B^*$ embeds into $B_n(K)/B_{n,1}(R) \times \dots \times B_n(K)/B_{n,n}(R)$, so elements of $B_n(K)/B^*$ are coded in \mathcal{G} , via the sort T_n . Now $B^* \leq G[\sigma] \leq B_n(R) \leq B_n(K)$. Thus, there is a map $\varphi : B_n(K)/G[\sigma] \rightarrow B_n(K)/B_n(R) = S_n$. If $s \in S_n$, then each element of $\varphi^{-1}(s)$ is an imaginary of $\text{red}(s)$, so is coded in \mathcal{G} over $C^{\ulcorner s \urcorner}$ by Proposition 2.6.5, and hence is coded in \mathcal{G} over C . (We here use that, by compactness, $\text{Int}_{k, C^{\ulcorner s \urcorner}}$ has elimination of imaginaries uniformly in s .)

Now let H be one of the groups $G[\sigma]$ in (f). For $t = (t_1, \dots, t_n)$ and $\gamma_i = |t_i|$, the coset $H \text{diag}(t)$ depends only on $(\gamma_1, \dots, \gamma_n)$.

(g) Let γ be a C -definable element of Γ^n , and define $\sigma \in \Sigma$ by $\sigma(i, j) = \gamma_i \gamma_j^{-1}$ for each $i < j$. Then $\sigma \in \Sigma^c$, and indeed, if $\sigma'(i, j)$ is obtained from $\sigma(i, j)$ by putting a $-$ whenever H has a $-$, then $\sigma' \in \Sigma^c$. To see the latter, observe that

as in (c), $G[\sigma'] = G(\gamma) = \text{diag}(\gamma)H \text{diag}(\gamma^{-1})$ (with a slight abuse of notation). Thus, there is a C -definable 1-1 correspondence between left cosets of $G[\sigma']$ in $B_n(K)$ and left cosets of H , given by right multiplication by $\text{diag}(\gamma)H$: indeed,

$$eG[\sigma'] \text{diag}(\gamma)H = e \text{diag}(\gamma)H \text{diag}(\gamma^{-1}) \text{diag}(\gamma)H = e \text{diag}(\gamma)H.$$

Thus, as elements of $B_n(K)/H$ are coded in \mathcal{G} over C (by (f)), so are elements of $B_n(K)/G[\sigma']$.

Now let

$$\Delta := \{(\gamma_1, \dots, \gamma_n) : \exists t = \text{diag}(t_1, \dots, t_n) \in T \text{ with } |t_i| = \gamma_i \text{ for each } i\}.$$

Let Φ be the set of all σ (or possibly the set of all σ' as in the last paragraph (g)) which arise from some $\gamma \in \Delta$. That is, $\sigma(i, j) = \gamma_i \gamma_j^{-1}$ (or possibly $(\gamma_i \gamma_j^{-1})^-$) for each $i < j$. Put $\tau := \inf(\Phi)$. We shall first assume that there is a C -definable non-zero element of Γ , and then modify the argument in the remaining case. Under this assumption, the C -definable structure on Γ has definable Skolem functions, by o-minimality.

By (b), it suffices to show that $\tau \in \Sigma^c$, so the lemma follows from the following claim.

Claim. (i) For $\ell = \binom{n-1}{2}$, there are C -definable $\sigma_1, \dots, \sigma_\ell \in \Sigma$ such that $\tau = \min\{\sigma_1, \dots, \sigma_\ell\}$. Furthermore, each σ_i either corresponds to some C -definable $(\gamma_1, \dots, \gamma_n) \in \Gamma^n$ as in (g), or only takes values in $\{0, \infty\}$, as in (d) above.

(ii) $\tau \in \Sigma^c$.

Proof of Claim. (i) For each (i, j) entry, we put an element into the list $\sigma_1, \dots, \sigma_\ell$. Fix $i < j$, and let $\tau_{ij} := \tau(i, j)$. For each $\delta > \tau_{ij}$, there is $\mu(\delta) \in \Phi$ with $(\mu(\delta))(i, j) < \delta$. By Skolemisation of Γ , we may suppose that μ is a C -definable function from the interval (τ_{ij}, ∞) to Φ . For each such δ , there is $\gamma(\delta) \in \Gamma^n$ such that for all $i' < j'$, $(\mu(\delta))(i', j') = \gamma_{i'}(\delta)(\gamma_{j'}(\delta))^{-1}$ (or, as usual, $\mu(\delta)(i', j')$ may take value $(\gamma_{i'}(\delta)(\gamma_{j'}(\delta))^{-1})^-$). Here, each $\gamma_{i'}$ is a C -definable function from (τ_{ij}, ∞) to Γ .

We have $\tau_{ij} \in \bar{\Gamma}$. Suppose first that $\tau_{ij} > 0$. As definable functions on Γ are affine, $\lim_{\delta \rightarrow \tau_{ij}^+} (\gamma_{i'}(\delta))$ is finite for each $i' \in \{1, \dots, n\}$. It follows that $\lim_{\delta \rightarrow \tau_{ij}^+} (\mu(\delta))$ is a C -definable element of Φ . Furthermore, $\inf(\Phi) \leq \lim_{\delta \rightarrow \tau_{ij}^+} (\mu(\delta))$, and $\lim_{\delta \rightarrow \tau_{ij}^+} (\mu(\delta))(i, j) = \inf(\Phi)(i, j)$. Thus, in this case we put $\lim_{\delta \rightarrow \tau_{ij}^+} (\mu(\delta))$ into the list $\sigma_1, \dots, \sigma_\ell$, and have taken care of the (i, j) entry.

Suppose now that $\tau_{ij} = 0$. Consider the behaviour of $\mu(\delta)(i', j') = \gamma_{i'}(\delta)\gamma_{j'}(\delta)^{-1}$ (possibly with $-$) as $\delta \rightarrow 0$. As the functions $\gamma_{i'}$ are affine, for some such entries (i', j') , the function $\mu(\delta)(i', j')$ will be constant as $\delta \rightarrow 0$, and for other (i', j') , $\lim_{\delta \rightarrow 0} \mu(\delta)(i', j')$ will equal 0 or ∞ . Let $\rho : \{(i', j') : i' < j'\} \rightarrow \bar{\Gamma}$ take value ∞ on (i', j') if $\lim_{\delta \rightarrow 0} \mu(\delta)(i', j')$ is non-zero, and take value 0 otherwise. Then $\rho \in \Sigma$ (this is easily checked), and is C -definable of type (d) above, so $\rho \in \Sigma^c$.

Also, $\rho(i', j') \geq \inf(\Phi)(i', j')$ for each (i', j') , and $\rho(i, j) = \inf(\Phi)(i, j) = 0$. Thus, in this case we put ρ into the list $\sigma_1, \dots, \sigma_\ell$ and so take care of the (i, j) entry.

(ii) By (b) and part (i) of the claim, $G[\tau] = \bigcap_{i=1}^\ell G[\sigma_i]$. By (d) and (g), each $\sigma_i \in \Sigma^c$. It follows by (e) that $\tau \in \Sigma^c$.

In the case when there are no non-zero C -definable elements of Γ , we again consider $\tau := \inf(\Phi)$ as in the claim. For each $i < j$, $\tau(i, j)$ is one of $1, 1^-$, or 0 . Let σ_1 be obtained from τ by putting $\sigma_1(i, j) = \infty$ whenever $i < j$ and $\tau(i, j) \in \{1, 1^-\}$. Also let σ_2 be the modification of τ arranged by putting $\sigma_2(i, j) = 1^-$ whenever $i < j$ and $\tau(i, j) = 0$. Then $\sigma_1, \sigma_2 \in \Sigma^c$ (using (d) and (f)). Also, $\tau = \min\{\sigma_1, \sigma_2\}$, so $\tau \in \Sigma^c$ by (e).

Proof of Proposition 3.3.4. (i) The cases when the range of f is in $K \cup k \cup \Gamma$ are handled by Theorem 2.4.13 (i) and (ii). For example, if $f : \Gamma \rightarrow \Gamma$, then there are $\ulcorner f \urcorner$ -definable $\gamma_0 = \infty > \gamma_1 > \dots > \gamma_{m+1} = 0$ such that on each (γ_{i+1}, γ_i) , f has the form $x \mapsto \delta_i x^{q_i}$; then $(\gamma_1, \dots, \gamma_m, \delta_0, \dots, \delta_m)$ is a code for f .

Suppose that $f : \Gamma \rightarrow S_n$. There are $\gamma_1, \dots, \gamma_m$ and I_0, \dots, I_m so that (i) and (ii) of Lemma 3.3.6 hold. Clearly these γ_i lie in $\text{dcl}(\ulcorner f \urcorner)$. We shall fix some $j \in \{0, \dots, m\}$, and show that if $I = I_j$, then $f|_I$ is coded. On I , f has canonical form $x \mapsto uh(x)B_n(R)$.

Clearly the function $h|_I$ is coded. The element u is not uniquely determined by $f|_I$, but the set $\{ug : g \in \bigcap_{x \in I} B_n(R)^{h(x)}\}$ is uniquely determined, and together with $h|_I$, determines $f|_I$. Furthermore, by Lemma 3.3.7, $\{ug : g \in \bigcap_{x \in I} B_n(R)^{h(x)}\}$ is coded in \mathcal{G} .

If $f : \Gamma \rightarrow T_n$, then for some $m \in \{1, \dots, n\}$ the function f is of the form $\Gamma \rightarrow B_n(K)/B_{n,m}(R)$. This case is handled as above (since Lemma 3.3.7 covers this case too).

(ii) We let $C := \ulcorner \gamma_0 \urcorner$, and show that the germ is coded over C . Let p be the generic type below γ_0 . We shall suppose that the germ of f below γ_0 is the set $\{f^{(\lambda)} : \lambda \in \Lambda\}$ of functions. If f is constant on p with value a , its germ below γ_0 has code (γ_0, a) . If $f : \Gamma \rightarrow \Gamma$ has form $f(x) = \delta x^q$, the germ has code (γ_0, δ) .

Next, suppose that f has range in S_n , identified as usual with $B_n(K)/B_n(R)$. For each $\lambda \in \Lambda$ there is $\gamma^{(\lambda)} < \gamma_0$ with $\gamma^{(\lambda)} \in \text{dcl}(\ulcorner f^{(\lambda)} \urcorner)$, such that $f^{(\lambda)}$ has canonical form $f^{(\lambda)}(x) = u^{(\lambda)}h^{(\lambda)}(x)B_n(R)$ on $(\gamma^{(\lambda)}, \gamma_0)$. It is easily checked that the $h^{(\lambda)}$ are all equal, say $h^{(\lambda)} = h$; these will be part of the code of the germ, as is γ_0 . We may suppose also that $u^{(\lambda)} = u$ for all i , by Lemma 3.3.5 (but u is not uniquely determined). Hence, for each $\lambda \in \Lambda$ there is a smallest $\gamma'^{(\lambda)} < \gamma^{(\lambda)}$ such that $f^{(\lambda)}(x) = uh(x)B_n(R)$ on $(\gamma'^{(\lambda)}, \gamma_0)$.

In the argument in (i), to code $f_{(\gamma'^{(\lambda)}, \gamma_0)}$, we had to code the coset $uH(\gamma'^{(\lambda)})$, where $H(\gamma'^{(\lambda)}) := \bigcap_{\gamma'^{(\lambda)} < x < \gamma_0} B_n(R)^{h(x)}$. Clearly, if $\gamma'^{(\lambda)} < \gamma'^{(\mu)}$ then $H(\gamma'^{(\lambda)}) \leq H(\gamma'^{(\mu)})$, and hence $uH(\gamma'^{(\lambda)}) \subseteq uH(\gamma'^{(\mu)})$. Thus, to code the germ of f below γ_0 , we must code $u \bigcup_{\lambda \in \Lambda} H(\gamma'^{(\lambda)})$.

If S is a set of pairs (i, j) with $1 \leq i < j \leq n$, let G_S be the set of matrices $A = (a_{ij}) \in B_n(R)$ such that $|a_{ij}| < 1$ whenever $(i, j) \in S$. Then, as in Lemma 3.3.7,

G_S is a group precisely if S has the property that whenever $i < k < j$ and $(i, j) \in S$, then $(i, k) \in S$ or $(k, j) \in S$. Let \mathcal{S} be the collection of all sets S with this property.

For each $i = 1, \dots, n$, let $\delta_i(x)$ be the norm of the (i, i) -entry of the diagonal matrix $h(x)$. Then, $\delta_i(x) = \varepsilon_i x^{q_i}$ for some $\varepsilon_i \in \Gamma$ and $q_i \in \mathbf{Q}$. Now let $S := \{(i, j) : i < j \text{ and } q_i > q_j\}$. Then $S \in \mathcal{S}$. It is now easily verified that $\bigcup_{\lambda \in \Lambda} H(\gamma'^{(\lambda)}) = X \cap G_S^{h(\gamma_0)}$ for some algebraic group X . (The point here is that if $q_i > q_j$, then for $x < \gamma_0$, the (i, j) entry of any element of $B_n(R)^{h(x)}$ has norm at most $\delta_i \delta_j^{-1} x^{q_i - q_j}$ which is increasing below γ_0 so the norm of the (i, j) entry of elements of $\bigcup_{\lambda \in \Lambda} H(\gamma'^{(\lambda)})$ is less than $\delta_i \delta_j^{-1} \gamma_0^{q_i - q_j}$; the group X arises because possibly $\gamma'^{(\lambda)} = 0$ for all i , and hence an (i, j) -entry of $\bigcup_{\lambda \in \Lambda} H(\gamma'^{(\lambda)})$ may be zero even though $\delta_i \delta_j^{-1} \gamma_0^{q_i - q_j} \neq 0$.) Thus, we must code $uG_S^{h(\gamma_0)}$. This is done exactly as in part (g) in the proof of Lemma 3.3.7.

Finally, we code germs of functions $\Gamma \rightarrow T_n$, and as in (i) we treat these as functions $\Gamma \rightarrow B_n(K)/B_{n,m}(R)$. We argue almost exactly as with germs of functions $\Gamma \rightarrow S_n$, except that the set S may be slightly larger. Indeed the set S may contain certain additional pairs (i, m) , arising from certain $<$ -inequalities obtained before the union process. \square

Lemma 3.3.8 *Let $U \in \mathcal{U}$ be an open 1-torsor. Let $f : U \rightarrow \mathcal{G}$ be a definable function. Then the germ of f on U is coded in \mathcal{G} .*

Proof. We handle the case when U is not definably isomorphic to a quotient of K ; the other case is similar. As in Theorem 3.3.2, we assume for convenience that $C = \emptyset$.

We may suppose U is a torsor of the open 1-module A . Let $u \in U$. Then the u -definable map $g_u : x \rightarrow x - u$ maps U bijectively to A . Furthermore, for sufficiently large $\gamma < 1$, A has a γ -definable closed submodule $A_\gamma := \bigcap (\delta R A : \gamma < \delta < 1)$, and the A_γ form a chain under inclusion with union A . For each sufficiently large $\gamma < 1$, let $U_{\gamma, u} := g_u^{-1}(A_\gamma)$. Then $U_{\gamma, u}$ is a closed unary set, so by Lemma 3.3.2 there is a code $c(u, \gamma) \in \mathcal{G}^n$ (for some n) for the germ of f on $U_{\gamma, u}$. By compactness we may suppose that $c(u, \gamma)$ is uniform in u, γ . For each $u \in U$ the function $c_u : \Gamma \rightarrow \mathcal{G}^n$ given by $c_u(\gamma) := c(u, \gamma)$ is definable, so by Proposition 3.3.4 its germ below 1 is coded (uniformly in u) by some $c'(u)$ in \mathcal{G} . Now for $u, u' \in U$, $c'(u) = c'(u')$, since $U_{\gamma, u} = U_{\gamma, u'}$ for sufficiently large γ . Thus, if $c' := c'(u)$, then c' is a code for the germ for f on U . \square

We must clarify the notion of *germ* of a function on a unary set with non-definable generic type. Let B be a base set of parameters, and $(t_i : i \in I)$ be a chain of 1-torsors, strictly ordered under reverse inclusion. Put $E := \bigcap (t_i : i \in I)$, and let p be the generic type of E over B . Suppose first that the definable functions f, g on E have the same germ on P . Then by compactness, there is

$n \in I$ such that $\{x \in t_n : f(x) \neq g(x)\}$ lies in a proper subtorsor of t_i for each $i \in I$. If this holds, then they have the same germ on t_i for each $i \geq n$. Conversely, suppose f, g are definable functions on E and for some n they have the same germ on t_i for each $i \geq n$. Then by Lemma 2.3.3, $\{x \in t_n : f(x) \neq g(x)\}$ is a Boolean combination of subtorsors of t_n , and meets each t_i for $i > n$ in a proper subtorsor. By adding parameters to identify t_n with a true 1-torsor, we see that f, g have the same germ on P . Thus, f, g have the same germ on P if and only if for sufficiently large i , f and g have the same germ on t_i .

Proposition 3.3.9 *Let $U \in \mathcal{U}$, f be a definable function to \mathcal{G} with domain containing U , and $B \subset \mathcal{G}$ with $B = \text{acl}_{\mathcal{G}}(B^{\Gamma} f^{\neg})$. Suppose that U is ∞ -definable over B . Then there is a B -definable function g with the same germ on U as f .*

We emphasise that the assumption $B = \text{acl}_{\mathcal{G}}(B^{\Gamma} f^{\neg})$ ensures that the \mathcal{G} -part of the code for f lies in B . In the case when U is an open unary set, this proposition does not say that the germ of f on U has a *strong* code (this is false in general by Remark 3.3.1(2)). This is because g may not be definable from a code for the germ of f .

Proof. First, suppose that U is a closed 1-torsor. Then by Theorem 3.3.2, B contains a strong code for the germ of f on U , and the lemma follows. Next, if U is a unary subset of Γ , the result follows from Proposition 3.3.4(i).

We suppose now that U is an open 1-torsor or an intersection of a chain $(t_i : i \in I)$ of B -definable open 1-torsors. For uniformity of notation, we suppose in the first case that $U = t_{i_0}$ and in the second case that i_0 is some fixed element of I . Suppose each t_i is a torsor of the module e_i , and that U is a torsor of $e := \bigcap (e_i : i \in I)$. Let $\Delta := \{\gamma \in \Gamma : \gamma Re_{i_0} \subseteq e\}$ (so $\Delta = (0, 1)$ when e is definable). Let $\delta := \sup \Delta$, a cut in Γ . For each $\gamma < \delta$, let $s_{\gamma} := \bigcap (\varepsilon Re_{i_0} : \delta > \varepsilon > \gamma)$, a closed submodule of e . We refer to cosets of s_{γ} as *closed subtorsors of radius γ* , and cosets of γRe_{i_0} as *open subtorsors of radius γ* . For any $u \subset U$, write $B_{\leq \gamma}(u)$ for the closed subtorsor of radius γ containing u , and $B_{< \gamma}(u)$ for the open one. The argument splits into two cases.

Case 1. U contains a B -definable element or subset s .

Case 2. Not Case 1 (in which case, by Lemma 2.3.3, U is a complete type over B).

Fix $\gamma \in \Gamma$ with $\gamma < \delta$. The equivalence relation $x - y \in s_{\gamma}$ partitions U into a set $S(\gamma)$ of closed 1-torsors t of radius γ . For each such t , by Theorem 3.3.2 there is a strong code $c(t)$ for the germ of f on t , and a $c(t)$ -definable function g_t with the same germ as f on t . Let $X(t) := \{x \in t : f(x) \neq g_t(x)\}$. If $X(\gamma) := \bigcup (X(t) : t \in S(\gamma))$, then for $t \in S(\gamma)$, $X(\gamma) \cap t$ is a proper subset of t , so $X(\gamma)$ is contained in a proper subtorsor $X'(\gamma)$ of U . By choosing $X'(\gamma)$ as small as possible, we may ensure that $X'(\gamma)$ is $B^{\Gamma} f^{\neg} \gamma$ -definable. Thus, the function $\gamma \mapsto X'(\gamma)$ is $B^{\Gamma} f^{\neg}$ -definable on some interval (δ_1, δ_2) , with $\delta_1 < \delta < \delta_2$. By

Proposition 3.3.4 it is coded in \mathcal{G} so is B -definable. It follows by Corollary 2.4.6 that either $X'(\gamma) = \emptyset$, or there is a proper B -definable subset s of U with the following property: for some B -definable function $h : \Gamma \rightarrow \Gamma$, and all generic γ below δ , $X'(\gamma)$ is the closed or open subtorsor of radius $h(\gamma)$ containing s . In the latter case, we must be in Case 1.

Proof in Case 1. There is B -definable $\delta' > \delta$ so that $c(B_{\leq \gamma}(s))$ (and hence $g_{B_{\leq \gamma}(s)}$) is definable for each γ with $\text{rad}(s) < \gamma < \delta'$. For such γ , put $c'(\gamma) := c(B_{\leq \gamma}(s))$. Then the function c' from Γ is $B^\Gamma f^\neg$ -definable, and coded in \mathcal{G} by Proposition 3.3.4, so is B -definable.

For $x \notin s$, let $|x - s|$ denote the radius of the smallest submodule of U containing $\{x - y : y \in s\}$. Let $B(x) := B_{\leq |x-s|}(x)$. Since c' and the function $x \mapsto |x - s|$ are B -definable, so is the function $x \mapsto c(B(x))$. For sufficiently large $\gamma < \delta$, f and $g_{B_{\leq \gamma}(s)}$ agree generically on $B_{\leq \gamma}(s)$, so agree on $B_{\leq \gamma}(s) \setminus B_{< \gamma}(s)$ (since the generic type of U over B is the type of all elements of $B_{\leq \gamma}(s) \setminus B_{< \gamma}(s)$ for sufficiently large $\gamma < \delta$). Since $g_{B(x)}$ is $c(B(x))$ -definable, there is a B -definable function H (with domain $B_{< \delta'}(s) \setminus s$) given by $H(x) := g_{B(x)}(x)$ for all $x \in B(x) \setminus B_{< |x-s|}(s)$. The function H has the same germ on U as f .

Proof in Case 2. Now, by the argument before Case 1, for generic $\gamma < \delta$ and a closed subtorsor t of U of radius γ , the functions g_t and f agree on t . We must show that for any conjugate f' of f over B , the functions f and f' agree on U . So suppose not, for some f' . By Lemma 3.3.8 (if U is an open torsor) or Theorem 3.3.2 applied to the $B_{\leq \text{rad}(t_i)}(t_i)$ (if I has no least element), f and f' have the same germ on U ; hence, $\{x \in U : f(x) \neq f'(x)\}$ lies in a proper subtorsor s of U . Pick $a \in s$, and choose $a' \in U$ so that $af \equiv_B a'f'$. Choose $\gamma < \delta$ generically over $Bafaf'$, and let $s' := B_{\leq \gamma}(s)$. Then f, f' have the same germ on s' . Furthermore, $af\gamma \equiv_B a'f'\gamma$, so there is a B -automorphism σ with $\sigma(af\gamma) = a'f'\gamma$. Then σ fixes $s' := B_{\leq \gamma}(a) = B_{\leq \gamma}(a')$. Hence, as f and $\sigma(f)$ have the same germ on s' , σ fixes this germ, so σ fixes $c(s')$ (here the full force of the notion of *strong code* is used). Thus, σ fixes $g_{s'}$. Now $f|_{s'} = g|_{s'} = \sigma(g|_{s'}) = \sigma(f|_{s'}) = f'|_{s'}$, the first equality coming by the discussion before the proof in Case 1. Hence f, f' agree on $s \subset s'$, and so on all of U . \square

3.4 Proof of elimination of imaginaries

We begin with a proof that finite sets (of sequences from \mathcal{G}) are coded in \mathcal{G} . This, via Theorem 2.1.2, easily implies that definable subsets of K are coded in \mathcal{G} . The latter was proved by Holly in [5] and [6] in equi-characteristic 0. We give a different proof, which applies in all characteristics. Elimination of imaginaries then follows rapidly from this and Proposition 3.3.9.

Proposition 3.4.1 *For each $r \in \mathbf{N}$, every finite subset of \mathcal{G}^r is coded in \mathcal{G} .*

Before proving this, we give a sequence of lemmas. We shall say that a finite set F is *primitive* over A if there is no proper non-trivial $(\ulcorner F \urcorner \cup A)$ -definable equivalence relation on F (so the permutation group induced on F by its setwise stabiliser in $\text{Aut}(K/A)$ is primitive). Let S_n^* denote the set of all cosets in K^n of elements of S_n , and T_n^* denote the set of cosets in K^n of elements of T_n . Also, let $\mathcal{M}S_n$ denote $\{\mathcal{M}A : A \in S_n\}$.

Lemma 3.4.2 *Let $C \subset K^{\text{eq}}$, and let F be a C -definable finite primitive subset of S_n . Then $F \subset \text{dcl}(C \cup \text{Int}_{k,C})$.*

Proof. Let $F := \{s_1, \dots, s_m\}$. We may assume $m > 1$. By primitivity, there is fixed n such that $s_i \in S_n$ for each i . Let s be the subgroup of K^n generated by $s_1 \cup \dots \cup s_m$. Then s is a finitely generated R -submodule of K^n . It follows, by Lemma 2.2.4, that s is a lattice, i.e., $s \in S_n$. Clearly s is C -definable. Also, $s_i \neq s$ for each i .

Let $\varphi : s \rightarrow \text{red}(s) = s/\mathcal{M}s$ be the reduction map. By Nakayama's Lemma, as $s_i \neq s$, also, $\varphi(s_i) \neq \varphi(s)$. On the other hand, $\text{red}(s) = \varphi(s) = \langle \varphi(s_1), \dots, \varphi(s_m) \rangle$. Hence, the $\varphi(s_i)$ are not all equal, so do not all have the same type over $\text{red}(s) \subset \text{Int}_{k,C}$. Hence neither do the s_i , so by primitivity all the s_i have distinct types over $C \cup \text{Int}_{k,C}$. \square

Lemma 3.4.3 *Every finite primitive set of balls is coded in \mathcal{G} .*

Proof. We shall assume the balls are all open. If they are all closed, then the same argument applies, with $<$ and \leq interchanged.

Let $F = \{t_1, \dots, t_m\} \subset T_1^*$ be a finite set of open balls. By primitivity (or just transitivity) of F , all the t_i have the same radius, γ say, and by primitivity, there is some $\delta > \gamma$ such that if $i \neq j$ and $x \in t_i, y \in t_j$ then $|x - y| = \delta$. Let $T := t_1 \cup \dots \cup t_m$ (regarding the t_i as subsets of K).

Let J^F be the set consisting of one variable polynomials

$$\{Q \in K[X] : \deg(f) \leq m \wedge \forall x \in T (|f(x)| < \delta^{m-1}\gamma)\}.$$

Then J^F is a definable R -submodule of K^{m+1} , so is coded in \mathcal{G} by Lemma 2.6.6. Also, J^F , together with γ and δ , are definable from $\ulcorner F \urcorner$. We must check that F is recoverable from J^F, γ, δ . For this, it suffices to check that if $f \in K[X]$ is monic of degree m , then $f \in J^F$ if and only if f has a root in each t_i .

In one direction, suppose that f has a root α_i in each t_i . Then $f(X) = \prod_{i=1}^m (X - \alpha_i)$. Suppose $x \in T$, with say $x \in t_1$. Then $|x - \alpha_1| < \gamma$, and $|x - \alpha_i| = \delta$ for $i = 2, \dots, m$. Hence $|f(x)| < \delta^{m-1}\gamma$.

In the other direction, suppose that $f \in J^F$ is monic of degree m and has roots β_1, \dots, β_m (listing repeated roots according to multiplicity). Then for all $j = 1, \dots, m$ there is i such that β_i lies at distance less than δ from (all elements

of) t_j : indeed, otherwise there is some t_j so that all β_i are at distance at least δ from t_j ; then if $x \in t_j$, we have $|f(x)| \geq \delta^m$, a contradiction. Hence, after relabelling, we may assume that for each i and all $x \in t_i$, $|\beta_i - x| < \delta$. Thus, if $i \neq j$ and $x \in t_j$, we have $|\beta_i - x| = \delta$. Now choose $x \in T$, with $x \in t_i$ say. Then $|f(x)| = \prod_{i=1}^m |x - \beta_i| = \delta^{m-1} |\beta_i - x|$. As $f \in J^F$, this forces $|\beta_i - x| < \gamma$, and hence $\beta_i \in t_i$, as required. \square

For any $a + s \in S_n^*$ (with $s \in S_n$), write $\text{red}(a + s)$ for $\{a + t : t \in \text{red } s\}$. This is a torsor of the k -vector space $\text{red}(s)$. Its elements are torsors of $\mathcal{M}s$.

Lemma 3.4.4 (i) Let $s \in S_n$ and $a \in K^n$, so $a + s \in S_n^*$. Let s' be the lattice in S_{n+1} which codes $a + s$, (as in Lemma 2.2.6); that is, s' is the R -submodule of K^{n+1} generated by $\{1\} \times (a + s)$. Then there is an $\ulcorner s' \urcorner$ -definable injection $g : \text{red}(a + s) \rightarrow \text{red}(s')$.

(ii) Let s be an R -lattice in K^n , let $a \in K^n$, and let $1 \leq r < n$, and let $\pi : K^n \rightarrow K^r$ be the projection to the first r coordinates. Then $\pi(\mathcal{M}s + a) \in T_r^*$. In particular, if $r = 1$ then $\pi(\mathcal{M}s + a)$ is an open ball.

(iii) In the notation of (ii), let $b \in \pi(\mathcal{M}s + a)$ and put $E_b := \{x - y : (b, x), (b, y) \in \mathcal{M}s + a\}$. Then $E_b = \mathcal{M}S_{n-r}$, and is independent of the choice of b .

Proof. (i) The map g takes $b + \mathcal{M}s$ to $(1, b) + \mathcal{M}s'$.

(ii) First, $\pi(\mathcal{M}s + a)$ is a coset of $\pi(\mathcal{M}s)$. As in Proposition 2.3.10, $\pi(s)$ is a lattice $s' \in S_r$. Now $\pi(\mathcal{M}s) = \mathcal{M}s'$, and the result follows.

(iii) We omit this. The main point is the observation (see Step 1 of the proof of Proposition 2.3.10) that kernels of projections of lattices are lattices. \square

Lemma 3.4.5 Let F be a finite set of balls, and $f : F \rightarrow \mathcal{G}$ be a function. Suppose that F is primitive over $\ulcorner f \urcorner$.

- (i) For any parameter set C , if $f(F) \subset \text{Int}_{k,C}$, then f is coded in \mathcal{G} over C .
- (ii) If $F' := f(F) \subset S_n^*$, then f is coded in \mathcal{G} over $\ulcorner F \urcorner$.
- (iii) If $F' := f(F) \subset T_n^*$, then f is coded in \mathcal{G} over $\ulcorner F' \urcorner$.
- (iv) If $f(F) \subset K$, then f is coded in \mathcal{G} .

Proof. (i) By Lemma 3.4.3, the set F is coded in \mathcal{G} . Let U be the smallest closed ball containing the elements of F . Then the elements of F lie in distinct elements of $\text{red}(U)$, by primitivity. Let $h : F \rightarrow \text{red}(U)$ be the map taking each member of F to the member of $\text{red}(U)$ which contains it. Let s be the element of S_2 which codes U , and $g : \text{red}(U) \rightarrow \text{red}(s)$ the injection given by Lemma 3.4.4 (i). Put $F^* := g(F) \subset \text{red}(s)$. By Proposition 2.6.5, $\text{Int}_{k,C \ulcorner s \urcorner}$ has elimination of imaginaries, so $f \circ h^{-1} \circ g^{-1}|_{F^*}$ is coded over $C \ulcorner s \urcorner$; a code for this, together with s and a code for F , codes f over C .

(ii) We adopt the notation and proof of Lemma 3.4.2. We may suppose that f is injective, and (by the usual identification of elements of S_n^* with elements

of S_{n+1}) that $F' \subset S_{n+1}$. As in 3.4.2, let $F' = \{s_1, \dots, s_m\}$, let s be the lattice generated by the s_i , put $C := \ulcorner s \urcorner$, and let $\varphi : s \rightarrow \text{red}(s)$. For each $i = 1, \dots, m$, there is $u_i \in \text{Int}_{k,C}$ interdefinable over $C \ulcorner F' \urcorner$ with s_i : write $u_i = g(s_i)$. Then by (i), $g \circ f$ is coded over C , and this code, together with s , codes f over $\ulcorner F' \urcorner$.

(iii) Let $F' := \{t_1, \dots, t_m\}$. For each i , there is $s_i \in S_n$ and $a_i \in K$ with $t_i = a_i + \mathcal{M}s_i$. The elements $a_1 + s_1, \dots, a_m + s_m$, lie in S_n^* , and are all equal or all distinct, by primitivity. If they are all distinct, then f is coded over $\ulcorner F' \urcorner$ by a code for the corresponding map to S_n^* , which exists by (ii).

So suppose that the $a_i + s_i$ all equal the same element s^* of S_n^* . Let s be the element of S_{n+1} which codes the torsor s^* . Then by Lemma 3.4.4, there is an s -definable injection $g : \text{red}(s^*) \rightarrow \text{red}(s)$. Now $g \circ f$ is coded over s by (i), and hence, so is f . Since $s \in \text{dcl}(F')$, f is coded over $\ulcorner F' \urcorner$.

(iv) Let U be the smallest closed ball containing all elements of $f(F)$. By primitivity, the elements of $f(F)$ all lie in distinct elements of $\text{red}(U)$. Let c be a code in K for $f(F)$ (this exists by elimination of imaginaries in algebraically closed fields). Let h be the map which takes each element of U to the element of $\text{red}(U)$ containing it. Then $h \circ f$ is coded over c by (iii), and f is coded over \emptyset by c and a code for $h \circ f$. \square

Lemma 3.4.6 *Let F be a finite subset of S_r , and $f : F \rightarrow S_n^* \cup T_n^*$ be a function. Suppose that F is primitive over $\ulcorner f \urcorner$, and put $F' := f(F)$. Then f is coded over $\ulcorner F \urcorner \ulcorner F' \urcorner$.*

Proof. This is just a slight variant of Lemma 3.4.5(iii). We use that if $F = \{s_1, \dots, s_m\}$ and s is the member of S_r generated by the s_i , then the s_i are coded over $\ulcorner s \urcorner \ulcorner F \urcorner$ by sequences from $\text{Int}_{k, \ulcorner s \urcorner}$. \square

Next, we give an application of Remark 3.3.3, and a slight extension of Lemma 3.3.8.

Lemma 3.4.7 *Let $\{e_1, \dots, e_n\}$ be a transitive C -definable set of balls. Put $W := \{\{c_1, \dots, c_n\} : c_i \in e_i\}$ (so the elements of W are n -element subsets of K , so coded by elements of K^n). Then there is a definable type q on W such that any definable function $f : W \rightarrow \mathcal{G}$ has germ on q coded in \mathcal{G} over C .*

Proof. By transitivity, the balls e_i all have the same radius $\gamma \in \text{dcl}(C)$.

Case 1. *The balls e_i are all closed.* In this case, for any $(x_1, \dots, x_n) \in e_1 \times \dots \times e_n$, let $x^* \in K^n$ be a code for the set $\{x_1, \dots, x_n\}$; here x^* lists the coefficients of the polynomial $\prod_{i=1}^n (X - x_i)$. There is a C -definable type p (over C) realised by (x_1, \dots, x_n) where x_i is chosen generically in e_i over $C \cup \{x_j : j < i\}$ (see Lemma 2.5.8). Let q be the corresponding type of the set $\{x_1, \dots, x_n\}$, and q^* the type (over C) of x^* . By Remark 3.3.3, any definable function has germ on q^* coded over C , and the result follows.

Case 2. The balls e_i are all open. Let r_γ be the generic type in Γ below γ . For $a := \{a_1, \dots, a_n\} \in W$, and $\delta \models r_\gamma$, let $c(a, \delta) = \{c_1(a), \dots, c_n(a)\}$, the n -set of closed balls around the c_i of radius δ . Let q_δ^a be the symmetrised generic type of $c_1(a) \times \dots \times c_n(a)$, i.e. the type of an imaginary coding the set $\{x_1, \dots, x_n\}$, where (x_1, \dots, x_n) is generic in $c_1(a) \times \dots \times c_n(a)$. Let q^a be the definable type such that, for any model M containing C , a , γ , we have $d \models q^a$ if and only if there is $\delta \models r_\gamma|_M$ such that $d \models q_\delta^a|M\delta$. The type q^a does not depend on the choice of a . Indeed, if $a'_i \in e_i$ for each i (a'_i chosen in M) and $a' = \{a'_1, \dots, a'_n\}$, and $\delta \models r_\gamma|_M$, then $c(a, \delta) = c(a', \delta)$, so $q_\delta^a = q_\delta^{a'}$. Thus, we may write $q = q^a$.

Now let $f_b(w) = f(w, b)$ be a Cb -definable function on W . We must show that the q -germ of f_b is coded in \mathcal{G} .

First, by Case (i), if $a \in W$ and $\delta \models r_\gamma|_Ca$, then the q_δ^a -germ of f_b is coded in \mathcal{G} over Ca , by say $g(a, b, \delta)$. Now g is a function of $b, \delta, c_a(\delta)$, so write $g = g'(b, \delta, c_a(\delta))$. If we view a, b as parameters, then g is a function on r_γ , and its germ on r_γ is coded over C , by Lemma 3.3.4(ii); that is, there is a code $h = h(a, b)$ such that $g(a, b, \delta)$ has the same r_γ -germ as $g(a', b', \delta)$ if and only if $h(a, b) = h(a', b')$. However, $h(a, b)$ does not depend on a : given $a, a' \in W$ and b , for $\delta \models r_\gamma$ we have $c(a, \delta) = c(a', \delta)$, so $g(a, b, \delta) = g(a', b, \delta)$, and hence $h(a, b) = h(a', b)$. Thus we can write $h(b) = h(a, b)$ for any $a \in W$, and $h(b)$ is a code for f_b over C . \square

Proof of Proposition 3.4.1. We shall prove by simultaneous induction on the natural number m the following two statements.

(I) $_m$ For each $r > 0$, every subset of \mathcal{G}^r of size m is coded in \mathcal{G} .

(II) $_m$ For any $F \subset G$ of size m , every function $f : F \rightarrow \mathcal{G}$ is coded in \mathcal{G} .

Both statements are trivial for $m = 1$.

We first prove (I) $_m$, for $r = 1$, assuming that (I) $_\ell$ and (II) $_\ell$ hold for all $\ell < m$, and that $m > 1$. We then prove (II) $_m$, and finally deduce (I) $_m$ for all r . In the proof of (I) $_m$, we may assume that F is primitive. For suppose that E is a proper non-trivial $\ulcorner F \urcorner$ -definable equivalence relation on F . Let C_1, \dots, C_ℓ be the E -classes. Then for each i there is by induction on m a code c_i in \mathcal{G} for C_i , and also a code e in \mathcal{G} for $\{c_1, \dots, c_\ell\}$. Then e is a code for F .

By elimination of imaginaries for algebraically closed fields, every finite subset of K is coded. Thus, it suffices to show that finite primitive subsets (of size m) of T_n and of S_n are coded. In fact, we prove by induction on n that primitive subsets of T_n^* of size m are coded. The case $n = 1$ is handled by Lemma 3.4.3. We omit the proof of the corresponding result for primitive subsets of S_n^* , as the proof is very similar (but slightly easier, with, for example, open balls replaced by closed balls).

So let $F = \{Z_1, \dots, Z_m\} \subset T_n^*$. Then each Z_i is a coset of $\mathcal{M}L_i$ for some lattice $L_i \in S_n$. By our primitivity assumption and the corresponding result for primitive subsets of S_n^* , $\{L_1, \dots, L_m\}$ is coded by some sequence λ from \mathcal{G} . Let A_i be the projection of Z_i to the first coordinate. Then, by Lemma 3.4.4(ii), A_i is an open ball, a torsor of some $B_i \in \mathcal{M}S_1$. For $a \in A_i$, let $E_{i,a} := \{x - y :$

$(a, x), (a, y) \in L_i\}$. Then $E_i = E_{i,a}$ is independent of a , and is an R -submodule of K^{n-1} lying in \mathcal{MS}_{n-1} , by Lemma 3.4.4(iii). For each $i = 1, \dots, m$, and each $x \in A_i$, let $h_i(x) := \{y : (x, y) \in Z_i\}$. Then $h_i(x)$ is a coset of E_i , so is a member of T_{n-1}^* .

Observe that Z_i is interdefinable (over \emptyset) with the triple consisting of $\ulcorner L_i \urcorner, \ulcorner A_i \urcorner$, and the germ of h_i on A_i . Indeed, if $c + \mathcal{M}L_i, c' + \mathcal{M}L_i$ are distinct, then they are disjoint, so either they have distinct projections to the first coordinate, or the corresponding functions do not agree anywhere on this projection.

By primitivity, the A_i are all equal or all distinct, as are the L_i .

Case 1. $A_1 = \dots = A_m = A$ and $L_1 = \dots = L_m = L$. Now, for $x \in A$, let $h(x)$ be a code in \mathcal{G} for the set $\{h_1(x), \dots, h_m(x)\}$; this exists by induction, as the $h_i(x)$ lie in T_{n-1}^* . By Lemma 3.3.8, the germ of h on the open ball A is coded in \mathcal{G} , by $\text{germ}(h)$ say, and $(\ulcorner A \urcorner, \ulcorner L \urcorner, \text{germ}(h))$ is a code for F .

Case 2. $A_1 = \dots = A_m = A$, and the L_i are all distinct. In this case, for each $x \in A$ there is a code in \mathcal{G} , denoted $h(x)$, for the map $\{L_1, \dots, L_m\} \rightarrow \{h_1(x), \dots, h_m(x)\}$ taking each L_i to $h_i(x)$; indeed, this map is coded by Lemma 3.4.6 — we may take the code to be over \emptyset since λ codes $\{L_1, \dots, L_m\}$ and $\{h_1(x), \dots, h_m(x)\}$ is coded in \mathcal{G} by induction. Again, the germ of h on A is coded in \mathcal{G} by $\text{germ}(h)$, say, and $(\ulcorner A \urcorner, \text{germ}(h))$ is a code in \mathcal{G} for F .

Case 3. A_1, \dots, A_m are all distinct. We shall assume that the L_i are also distinct, but the other case is easy.

For $(x_1, \dots, x_m) \in A_1 \times \dots \times A_m$, let x^* be a code in K for $\{x_1, \dots, x_m\}$. For each such x^* , consider the map $f(x^*) : \{A_1, \dots, A_m\} \rightarrow \{x_1, \dots, x_m\}$ which takes each A_i to x_i , the map $h(x^*) : \{A_1, \dots, A_m\} \rightarrow \{h_1(x_1), \dots, h_m(x_m)\}$ taking each A_i to $h_i(x_i)$, and the map $\ell(x^*) : \{A_1, \dots, A_m\} \rightarrow \{L_1, \dots, L_m\}$ which takes each A_i to L_i . By Lemma 3.4.5 (and induction, in the case of h , where we need that $\{h_1(x_1), \dots, h_m(x_m)\}$ is coded in \mathcal{G}), each of $f(x^*)$, $h(x^*)$, and $\ell(x^*)$ is coded in \mathcal{G} . Let $g(x^*)$ be a code in \mathcal{G} for the triple $(f(x^*), h(x^*), \ell(x^*))$.

We consider the germ of the function g on the definable type q , provided by Lemma 3.4.7 (Case 2, in the proof). Let c be a code in \mathcal{G} for $\{A_1, \dots, A_m\}$; this exists by Lemma 3.4.3. The above germ of g is coded in \mathcal{G} over c , and its code, together with c , is a code for F . This completes the proof of $(I)_m$ for $r = 1$.

We now prove $(II)_m$, that if $F \subset \mathcal{G}$ with $|F| = m$, then any function $f : F \rightarrow \mathcal{G}$ is coded. For each n , let τ_n denote the natural map $T_n \rightarrow S_n$. By an easy induction, we may suppose that F and F' are primitive over $\ulcorner f \urcorner$. By $(I)_m$, F and $F' := f(F)$ are coded in \mathcal{G} .

Suppose first that $F \subset T_n$, and $F' \subset T_{n'}$. By primitivity, $\tau_n|_F$ is injective or constant; likewise for $\tau_{n'}|_{F'}$. In all the four cases, Proposition 2.6.5 and Lemma 3.4.2 ensures that f is coded.

If F (or similarly F') is a subset of K , then there is a canonical map $F \rightarrow T_1$ (take the map to the reduction of the smallest closed ball containing F). Likewise, if $F \subset S_n$, then by Lemma 3.4.2 there is an $\ulcorner F \urcorner$ -definable injection $F \rightarrow \text{Int}_{k, \ulcorner F \urcorner}$. Thus, we can reduce all cases to that handled in the last paragraph.

Finally, we deduce $(I)_m$ for all r . Let $r > 1$, and let $F \subset \mathcal{G}^r$ have size m , and for $i = 1, \dots, r$ let π_i be the projection $\mathcal{G}^r \rightarrow \mathcal{G}$ to the i^{th} coordinate. As in the $r = 1$ case, by induction on m we may suppose that F is primitive. Then by primitivity, each π_i is constant or injective. As $|F| > 1$, there is some projection, say π_1 , which is injective. But now F is coded by a code for $\pi_1(F) \subset \mathcal{G}$, together with the sequence of codes for functions $\pi_i \circ \pi_1^{-1} : \mathcal{G} \rightarrow \mathcal{G}$. \square

The following corollary enables us to work with a weaker-looking definition of *coding*.

Corollary 3.4.8 *If $i \in K^{\text{eq}}$ and there is a tuple e in $\text{acl}(Ai) \cap \mathcal{G}$ such that $i \in \text{dcl}(Ae)$ then i is coded in \mathcal{G} over A .*

Proof. Let S be the set of conjugates of e over Ai . Then $\ulcorner S \urcorner \in \text{dcl}(Ai)$, and $i \in \text{dcl}(A\ulcorner S \urcorner)$. Furthermore, since finite subsets of \mathcal{G}^n can be coded by Theorem 3.4.1, S has a code e' in \mathcal{G} over A , and e' is a code of i over A . \square

Corollary 3.4.9 *Let $U \in \mathcal{V}$ be definable, let $C \subset \mathcal{G}$ be such that there is a C -definable injection $U \rightarrow \mathcal{G}$, let $f : U \rightarrow \mathcal{G}$ be a definable function, and let $B = \text{acl}_{\mathcal{G}}(C\ulcorner f \urcorner)$. Then $f \in \text{dcl}(B)$.*

Proof. Consider $\Sigma := \{D \subset U : D, f|_D \text{ both definable over } B\}$. If $\bigcup \Sigma = U$, then by compactness, f is B -definable, so we may suppose $\bigcup \Sigma \neq U$. Then there is a complete type p over B whose set P of realisations lie in $U \setminus \bigcup \Sigma$. By Lemma 2.3.6, p is the generic type of a unary set V over B . As V is a subtorsor of U , $V \in \mathcal{U}$. By Proposition 3.3.9, there is a B -definable function g with the same germ on V as f . If $X := \{x \in U : f(x) = g(x)\}$, then $X \cap P \neq \emptyset$. By Lemma 2.3.3, X is uniquely a finite set of Swiss cheeses no two trivially nested. By Theorem 3.4.1, as subtorsors of U are coded, each of the Swiss cheeses is coded in \mathcal{G} , and hence X is coded in \mathcal{G} by Theorem 3.4.1 again. But X is $B\ulcorner f \urcorner$ -definable, so B -definable, as $B = \text{acl}_{\mathcal{G}}(\ulcorner f \urcorner)$. Hence, as p is a complete type over B , $X \supseteq P$. But as g is B -definable, $f|_X$ is B -definable, so $X \in \Sigma$, a contradiction. \square

Theorem 3.4.10 *The theory $T_{\mathcal{G}}$ in the language $\mathcal{L}_{\mathcal{G}}$ has elimination of imaginaries.*

Proof. By Lemma 3.2.1 and Remark 3.2.6 it suffices to code definable functions f from sets in \mathcal{V} to \mathcal{G} . By Corollary 3.4.8, it suffices to show that $\ulcorner f \urcorner \in \text{dcl}(\text{acl}_{\mathcal{G}}(\ulcorner f \urcorner))$. This is precisely what Corollary 3.4.9 says. \square

We also justify the more concrete version of elimination of imaginaries stated in the Introduction.

Proof of Theorem 1.0.2. Let e be an imaginary in the algebraically closed valued field K . By Theorem 3.4.10, there is a sequence $\bar{a}\bar{b}\bar{c}$ interdefinable with e , with \bar{a} a tuple of field elements, \bar{b} a tuple from \mathcal{T} , and \bar{c} a tuple from \mathcal{S} . (We identify Γ with S_1 and k with $\text{red}(R) \in T_1$.) If $\bar{a} \in K^n$ then it can be regarded as a torsor of the trivial submodule of K^n , hence as a submodule of K^{n+1} . We may identify \bar{c} with a single lattice c (the product of the entries of \bar{c}). Likewise, \bar{b} is identifiable with a singleton element of \mathcal{T}_m , for some m , and hence, by Lemma 2.2.6, with an R -submodule of K^{m+1} . The product of the three modules obtained is an R -module which codes e . \square

We give now a corollary of Theorem 3.4.10 for k -internal sets. This extends the results of Section 2.6, and will be essential in [3].

Proposition 3.4.11 *Let D be a C -definable subset of K^{eq} . Then*

- (i) *D is k -internal if and only if D is stable and stably embedded;*
- (ii) *D is k -internal if and only if $D \subset \text{dcl}(C \cup \text{Int}_{k,C})$.*

In the proof of (ii) we shall freely use the fact that any tuple of $\text{Int}_{k,C}$ is interdefinable over C with an element of $\text{Int}_{k,C}$ (using the product of the corresponding lattices). The ingredients of the proof are elimination of imaginaries, Lemma 3.4.2 (to handle finite primitive sets of lattices) and Lemma 2.6.7 (to reduce finite sets of lattices to the primitive case).

Proof. (i) By elimination of imaginaries for ACVF, we may suppose that the entries of D lie in \mathcal{G} . Then (i) follows from Lemma 2.6.2.

(ii) The right-to-left direction is trivial. For the other direction, let D be k -internal. As in (i) above, we may suppose that the entries of D lie in \mathcal{G} . By Lemma 2.6.2 (i) \Leftrightarrow (v), it suffices to show

- (a) if F is a C -definable finite subset of \mathcal{G} , then $F \subset \text{dcl}(C \cup \text{Int}_{k,C})$, and
- (b) if s is a lattice with $s \in \text{acl}(C)$, then $\text{red}(s) \subset \text{dcl}(C \cup \text{Int}_{k,C})$.

We first prove (a). The proof is by induction on $|F|$. We may suppose that F is a single orbit over C . Suppose first that F is imprimitive over C , with an equivalence relation having non-trivial classes F_1, \dots, F_r , say. By elimination of imaginaries each F_i has a canonical parameter b_i (a tuple from \mathcal{G}), and by induction (as $\{b_1, \dots, b_r\}$ is C -definable and $r < |F|$), $b_1, \dots, b_r \in \text{dcl}(C \cup \text{Int}_{k,C})$. Let t_i be chosen canonically in $\text{Int}_{k,C}$ such that b_i is Ct_i -definable. We may suppose that there is a C -definable lattice s such that $\{t_1, \dots, t_r\}$ is a C -definable subset of $\text{red}(s)$. Put $C_i := Ct_i$. By induction, $F_i \subset \text{dcl}(C_i \cup \text{Int}_{k,C_i})$. Thus, there is a C_i -definable lattice s_i (for $i = 1, \dots, r$) and $u_i \in \text{red}(s_i)$ such that $F_i \subset \text{dcl}(Ct_i u_i)$. Now, by Lemma 2.6.7, $\text{red}(s_i) \subset \text{dcl}(C \cup \text{Int}_{k,C})$, and it follows that $F \subset \text{dcl}(C \cup \text{Int}_{k,C})$.

Thus, we may suppose that F is primitive over C . If $F \subset S_n$, then $F \subset \text{dcl}(C \cup \text{Int}_{k,C})$ by Lemma 3.4.2. If $F \subset K$, then by primitivity there is a C -definable closed ball U such that all elements of F lie in distinct members of

$\text{red}(U)$. By Lemma 3.4.4 (i), $\text{red}(U) \subset \text{dcl}(C \cup \text{Int}_{k,C})$, and the result follows, as $F \subset \text{dcl}(\text{red}(U) \cup \{\ulcorner F \urcorner\})$. Finally, suppose that $F \subset T_n$, and let $\tau : T_n \rightarrow S_n$ be the natural map. If τ takes constant value s on F , then $F \subset \text{dcl}(\text{red}(s)) \subset \text{dcl}(C \cup \text{Int}_{k,C})$. Otherwise, by primitivity, τ is injective on F . Let $s_i := \tau(t_i)$. Then (by the S_n -case just above), $s_1, \dots, s_m \in \text{dcl}(C \cup \text{Int}_{k,C})$. Hence, there is a C -definable lattice s , and $u_1, \dots, u_m \in \text{red}(s)$ such that s_i is Cu_i -definable. By elimination of imaginaries in $\text{Int}_{k,C}$ we may suppose that the u_i are chosen canonically, so $\{u_1, \dots, u_m\}$ is C -definable. Then, by Lemma 2.6.7, $\text{red}(s_i) \subset \text{dcl}(C \cup \text{Int}_{k,C})$ for each i , so $F \subset \text{dcl}(C \cup \text{Int}_{k,C})$.

We now prove (b). So let $F = \{s_1, \dots, s_m\}$ be a C -definable subset of S_n . By (a), $F \subset \text{dcl}(C \cup \text{Int}_{k,C})$, so we may suppose there is a C -definable lattice s , and a C -definable subset $\{t_1, \dots, t_m\}$ of $\text{red}(s)$, such that s_i is Ct_i -definable. but now, by Lemma 2.6.7, $\text{red}(s_i) \subset \text{dcl}(C \cup \text{Int}_{k,C})$, as required. \square

Finally, we give two results which use the ideas of the last section, and are crucial to the independence theory developed in [3]; they are used particularly for the existence of invariant extensions of arbitrary types, and for the maximum modulus principle.

Lemma 3.4.12 *If $B \supseteq \text{acl}(B \cap \mathcal{G})$, and $\alpha \in \Gamma$, then $\text{acl}(B\alpha) \cap \mathcal{G} = \text{dcl}(B\alpha) \cap \mathcal{G}$.*

Proof. If $a \in \text{acl}(B\alpha) \cap \mathcal{G}$, then a lies in a fibre above α of a B -definable finite cover $\rho\Gamma$ of Γ . We apply the results of Section 2.4, together with Lemmas 3.3.5 and 3.3.6, to the identity function id on $\rho\Gamma$. We suppose that ρ has fibres of size r .

We shall consider the case when $a \in S_n$, as the other cases are similar. First, if I is an interval of Γ , we say that id has *canonical form* on I if there are affine functions $h_i : I \rightarrow D_n(K)/D_n(R)$ and $u_i \in U_n(K)$ (for $i = 1, \dots, r$) such that if $\rho(x) = y \in I$ then $\text{id}(x) = x \in \{u_i h_i(y) B_n(R) : 1 \leq i \leq r\}$. Observe that as the identity function is injective, for each such $x = \rho^{-1}(y)$, if $i \neq j$ then $u_i h_i(y) B_n(R) \neq u_j h_j(y) B_n(R)$. As in Lemmas 3.3.5 and 3.3.6, one can partition Γ into finitely many B -definable intervals, on each of which id has canonical form. Suppose I is such an interval. Then for all $y \in I$ and $x \in \rho^{-1}(y)$, $x \in \{u_i h_i(y) B_n(R) : 1 \leq i \leq r\}$. Now each of the functions $x \mapsto u_i h_i(y) B_n(R)$ is coded in \mathcal{G} , and they are algebraic over B , so each is definable over B (as $B \supseteq \text{acl}(B \cap \mathcal{G})$). Hence, for each $i = 1 \dots, r$, $\{x \in \rho^{-1}(I) : x = u_i h_i(\rho(x)) B_n(R)\}$ is definable over B . This set intersects each fibre of ρ in a singleton, and it follows that for each $x \in \rho^{-1}(I)$, $x \in \text{dcl}(B\rho(x))$, as required. \square

Proposition 3.4.13 *Let $B \subset K^{\text{eq}}$ with $\text{acl}(B) = B$, and let U be a unary set over B . Let f be a definable function (not necessarily B -definable) with range in \mathcal{G} such that for all $x \in U$ we have $f(x) \in \text{acl}(Bx)$. Then there is a B -definable function h with the same germ on U as f .*

Remark. By elimination of imaginaries to \mathcal{G} (Theorem 3.4.10), the assumption $B = \text{acl}(B)$ could be replaced by $B \supseteq \text{acl}_{\mathcal{G}}(B)$.

Proof. In the case when $\text{dom}(f) \subset \Gamma$, the hypothesis implies by Lemma 3.4.12 that f (or its restriction to some set containing U) is itself B -definable. So we shall suppose U is a 1-torsor.

Let n be the number of conjugates of $f(x)$ over Bx , for $x \in U$. Suppose first that p is the generic type of an open or closed 1-torsor defined over B . Then by Lemma 2.3.8, the germ of f on P is definable, and we claim that it is B -definable. For suppose not. Then as $B = \text{acl}(B)$, the germ of f on P is not in $\text{acl}(B)$, so there are conjugates $f = f_0, \dots, f_n$ of f with distinct germs. Now let $a \in P$ be generic over Bf_0, \dots, f_n . Then the $f_i(a)$ are pairwise distinct, which is a contradiction.

The lemma follows if p is the generic type of a closed 1-torsor, for B contains a code c for the germ of f on P , and by Theorem 3.3.2, this code is strong.

Suppose now that U is either a B -definable open 1-torsor or the intersection of a chain $(t_i : i \in I)$ of B -definable closed 1-torsors. We first suppose that U has a proper B -definable subtorsor s . We adopt the notation $(i_0, \delta, e_i, e, B_{\leq \gamma}(s), \text{etc.})$ of the proof of Proposition 3.3.9. For each $\gamma \in \Gamma$ with $\text{rad}(s) < \gamma < \delta$, consider $s_\gamma := B_{\leq \gamma}(s)$. By the closed subtorsor case above, for each such $\gamma \in \Gamma$ there is a function g_γ on s_γ , defined over $\text{acl}_{\mathcal{G}}(B\gamma)$ and agreeing with f generically on s_γ .

Furthermore, by compactness g_γ is definable uniformly in γ . Now by Lemma 3.4.12, $\text{acl}_{\mathcal{G}}(B\gamma) \subset \text{dcl}(B\gamma)$, so g_γ is $B\gamma$ -definable. We now argue as in the proof of Case 1 of Proposition 3.3.9. For sufficiently large $\gamma < \delta$, g_γ and f agree on $s_\gamma \setminus B_{< \gamma}(s)$. Define h to agree with $g_\gamma(x)$ on $s_\gamma \setminus B_{< \gamma}(s)$ for all $\gamma > \text{rad}(s)$. Such a function h can be chosen to be B -definable, and if U is an intersection of a chain $(t_i : i \in I)$ of 1-torsors, then the domain of h will contain the generic type of one of the t_i . Now h and f have the same germ on P , as required.

Finally, suppose that U has no B -definable subtorsor. Then by Lemma 2.3.3, U is the solution set of a complete type p over B . By Corollary 2.4.5, for generic $\gamma < \delta$, all closed subtorsors t of U of radius γ have the same type, and indeed, all elements of t have the same type over Bt . For each $x \in U$, let D_x denote the set of conjugates of $f(x)$ over Bx , a Bx -definable set of size n .

Let $\gamma < \delta$, and t be a closed subtorsor of U of radius γ . By the closed torsor case, there is an $\text{acl}_{\mathcal{G}}(Bt)$ -definable function g on t agreeing generically with f on t . Let g_1, \dots, g_m be the conjugates of g over Bt . Since the elements of t all have the same type over Bt , there are no Bt -definable proper subtorsors of t , and hence no $\text{acl}(Bt)$ -definable proper subtorsors of t (otherwise, take unions of conjugates). It follows that for any $i, j \leq m$, $\{x : g_i(x) = g_j(x)\}$ is empty or equals t , and the former must hold if $i \neq j$. From this a short argument shows that $m = n$ and for $x \in t$, $D_x = \{g_i(x) : 1 \leq i \leq n\}$. Furthermore, if $\gamma' < \delta$ is chosen generically over Bt (so $\gamma' > \gamma$), and t' is the closed subtorsor of radius γ' containing t , and g'_1, \dots, g'_m are analogues of g_1, \dots, g_m for t' , then

$$\{g'_i|_t : 1 \leq i \leq m\} = \{g_i : 1 \leq i \leq m\}.$$

Now define as follows an equivalence relation \sim on the set of conjugates of (γ, t, g) : (γ', t', g') is equivalent to (γ'', t'', g'') if for generic (over the above data) $\gamma''' < \delta$, the ball t''' of radius γ''' containing t' and t'' , has a function g''' such that $\text{tp}(\gamma', t', g') = \text{tp}(\gamma'', t'', g'') = \text{tp}(\gamma''', t''', g''')$ and $g'''|_{t'} = g'$, $g'''|_{t''} = g''$. By the last paragraph, the relation \sim has m classes. Furthermore, \sim is the restriction to $\text{tp}(\gamma, t', g')$ of a B -definable equivalence relation with m classes. As $B = \text{acl}(B)$, each class is definable over B . Now there is (γ, t, g) such that g and $f|_t$ have the same germ on t , and the union of all g' with $(\gamma', t', g') \sim (\gamma, t, g)$ is the required function h . \square

3.5 Necessity of the geometric sorts

In this section we show that the main theorem is in a sense optimal, that is, elimination of imaginaries could not be proved with very much simpler sorts. The first result shows that we could not make do with the S_n and just finitely many T_n , in order to obtain elimination of imaginaries. Given any base C of parameters, let $\text{Int}_{k,C}^n$ be the many-sorted substructure of $\text{Int}_{k,C}$ consisting of sorts $\text{red}(s)$ for $s \in \text{dcl}(C) \cap S_m$ for all $m \leq n$ (with the induced C -definable structure). The result shows that in general $\text{Int}_{k,C}^n$ does not even interpret the whole of $\text{Int}_{k,C}$.

Proposition 3.5.1 *Let $n \in \mathbf{N}$, with $n > 1$.*

(i) *There is a base C and $s \in \text{dcl}(C) \cap S_{n+2}$ such that $\text{red}(s)$ is not a subset of $\text{dcl}(C \cup \text{Int}_{k,C}^n)$.*

(ii) *The theory of an algebraically closed valued field K does not admit elimination of imaginaries to sorts K, k, Γ, S_m ($m \in \mathbf{N}$) and T_m ($m \leq n$).*

Proof. (i) First observe that if $s \in S_n \cap C$, then the group of automorphisms of $V = \text{red}(s)$ induced by the subgroup of $\text{Aut}(K)$ which fixes $k \cup C$ pointwise preserves the k -vector space structure on V . It also preserves the filtration on V (that is, the filtration used for example in Step 3 of the proof of Proposition 2.3.10); hence it embeds in $B_n(k)$, the group of upper triangular matrices over k , so is soluble of derived length at most n . Thus, the group induced on $\text{Int}_{k,C}^n$ by $\text{Aut}(K/k \cup C)$ (the pointwise stabiliser of $k \cup C$) is soluble of derived length at most n . In particular, if $n' > n$ and $s \in S_{n'} \cap \text{dcl}(C)$ and $\text{red}(s) \subset \text{dcl}(C \cup \text{Int}_{k,C}^n)$, then $\text{Aut}(K/k \cup C)$ induces a soluble group of derived length at most n on $\text{red}(s)$.

On the other hand, for any $m > 0$ let $s \in S_m$, and let $C = \text{dcl}_G(s)$. We show that if s is chosen carefully then the group of automorphisms induced on $\text{red}(s)$ by $\text{Aut}(K/k \cup C)$ has derived length $m - 1$; this can be arbitrarily large, contrary to the last paragraph.

To see this, first observe that if $\gamma_1 < \dots < \gamma_t$ is a sequence of elements of Γ , with each γ_i chosen generically large over the previous γ_j , and $V_i := \{x \in K :$

$|x| = \gamma_i\}$, then $V_1 \times \dots \times V_t$ is a complete type over $k \cup \{\gamma_1, \dots, \gamma_t\}$. Now choose a lower unitriangular matrix $B = (b_{ij})$ over K , with

$$1 < |b_{21}| < \dots < |b_{m1}| < |b_{32}| < \dots < |b_{m2}| < \dots < |b_{m,m-1}|.$$

We also assume each $|b_{ij}|$ is chosen generically large over the previous $|b_{i'j'}|$ in the above sequence. Let A be any lower unitriangular matrix over R . The genericity (and the fact that corresponding elements of B and AB have the same norms) ensures there is $\sigma \in \text{Aut}(K/k)$ with $\sigma(B) = AB$.

Let L be the lattice generated by the rows of B . Then σ takes these rows to the rows of AB , so fixes L , and induces an automorphism of $V = L/\mathcal{M}L$. Also, σ fixes $C := \text{dcl}_{\mathcal{G}}(\ulcorner L \urcorner)$. As σ fixes k , this is a k vector space automorphism of V . Furthermore, with respect to the basis of V consisting of the reductions of the rows of B , σ is represented by the matrix $\text{red}(A)^T$ (acting by left multiplication). Thus, left multiplication by any element of $U_m(k)$ gives an automorphism of V induced by $\text{Aut}(K/k \cup C)$. The derived length of $U_m(k)$ is $m - 1$, so (i) follows provided $m \geq n + 2$.

(ii) Let $m = n + 2$, and $V = L/\mathcal{M}L$ as in (i). In (i), as $U_m(k)$ has no definable proper subgroups of finite index, its action on V is induced by $\text{Aut}(K/k \cup \text{acl}(C))$. Choose $g \in U_m(k)^{(m-1)} \setminus \{1\}$, that is, non-trivially in the penultimate term of the derived series. Let $\sigma \in \text{Aut}(K/k \cup \text{acl}(C))$ induce g ; we may suppose that σ can be expressed as a product of a sequence τ of elements of $\text{Aut}(K/k \cup \text{acl}(C))$ so as to witness that $g \in U_m(k)^{(m-1)}$. There is $v \in V$ with $\sigma(v) \neq v$. Let $c = (c_1, \dots, c_r)$ be a code for v in $\text{Int}_{k,C}^n$. Then each c_i lies in a k -internal C -definable set. If c_i is a lattice or field element, then by Lemma 2.6.2 $c_i \in \text{acl}(C)$, so σ fixes c_i . Otherwise, $c_i \in \text{red}(s)$ where $s \in \text{acl}(C) \cap S_\ell$ for some $\ell \leq n$. Then by the first paragraph of the proof of (i), the elements of τ fix s , and σ fixes c_i . It follows that $v \notin \text{dcl}(c)$, a contradiction. \square

The next result gives an alternative proof that the original conjecture (that elimination of imaginaries holds to sorts consisting of open and closed balls) is false. This fact is implied by Proposition 3.5.1; for all balls are coded in $S_1 \cup S_2 \cup T_1 \cup T_2$ (as in the last paragraph of the proof of Lemma 2.2.6). However, the next result also contains slightly more delicate information about k -internal sorts. It shows that in general, over a base C , we do not have elimination of imaginaries for the multi-sorted structure with a sort $\text{red}(u)$ for each C -definable closed ball.

Proposition 3.5.2 (i) There is a parameter set C such that the multisorted k -internal structure $\text{Int}_{k,C}^{\text{op}}$ does not have elimination of imaginaries. Here, $\text{Int}_{k,C}^{\text{op}}$ is the structure which has a sort for each set $\text{red}(t)$ (t a C -definable closed ball), and its \emptyset -definable relations are those induced by the C -definable relations of K^{eq} .

(ii) The theory of an algebraically closed valued field does not have elimination of imaginaries to the level of sorts consisting of field elements and open and closed balls.

Proof. (i) We work over an arbitrary parameter set C_0 . Pick generic $\varepsilon < 1$ (in Γ). Then choose b_1 generic in R over $C_0\varepsilon$ and b_2 generic in R over $C_0\varepsilon b_1$, and put $U_i := B_{\leq \varepsilon}(b_i)$ for each i . Let $V := \text{red}(\varepsilon R)$, a 1-dimensional k -space, and for $i = 1, 2$ let $A_i := \text{red}(U_i) = U_i/\varepsilon\mathcal{M}$, a torsor of V . Let $C := \text{acl}(C_0 \ulcorner U_1 \urcorner U_2 \urcorner)$. Let $\text{Aff}(A_1, A_2)$ be the set of affine homomorphisms $A_1 \rightarrow A_2$. This is clearly a C -definable k -internal set of Morley rank 2: a generic affine homomorphism h is determined by the induced element of $\text{Hom}(V, V)$ (a Morley rank one set), and, for any fixed $a \in A_1$, the image $h(a)$. If elimination of imaginaries to balls held in $\text{Int}_{k,C}^{\text{op}}$, then each generic element of $\text{Aff}(A_1, A_2)$ would be coded over C in $\text{Int}_{k,C}^{\text{op}}$ by an independent (over C) pair of elements of C -definable strongly minimal sets of the form $\text{red}(e)$, where e is a C -definable closed ball. In particular, if h were generic in $\text{Aff}(A_1, A_2)$ then $\text{acl}(h)$ would contain two distinct rank 1 algebraically closed subsets (in $\text{Int}_{k,C}^{\text{op}}$). We now work over C , so omit reference to parameters from C .

Claim 1. The action of $V \times V$ on $A_1 \times A_2$ by translation is elementary over k, V .

Proof. It suffices to check that for $(a_1, a_2) \in A_1 \times A_2$ and $(v_1, v_2) \in V \times V$, $\text{tp}(a_1, a_2/kV) = \text{tp}(a_1 + v_1, a_2 + v_2/kV)$. This follows from the generic choice of U_1 and U_2 . As a first step, observe that if $\text{tp}(a_1/kV) \neq \text{tp}(a_1 + v_1/kV)$, then for each generic $U \in R/\varepsilon R$ there is a finite non-empty subset U^* of $\text{red}(U)$ definable over $C_0 \ulcorner U \urcorner U_2 \urcorner \bar{v}$ (\bar{v} from kV), with $U^* := f(\ulcorner U \urcorner, \bar{v})$ say. Let

$$g(\ulcorner U \urcorner) := \{\bar{v}' \in kV : f(\ulcorner U \urcorner, \bar{v}') \text{ is a finite non-empty subset of } \text{red}(U)\}.$$

Then g is an $C_0 \ulcorner U_2 \urcorner$ -definable function from $R/\varepsilon R$ into the stable structure $\text{Int}_{k,C}$. It follows easily that g is constant on an infinite subtorsor W of $R/\varepsilon R$ containing U . Hence, $\{f(\ulcorner U' \urcorner, \bar{v}) : U' \in W\}$ is a definable subset of W which is not a finite union of Swiss cheeses, contrary to Lemma 2.3.3. Thus $\text{tp}(a_1/kV) = \text{tp}(a_1 + v_1/kV)$. A similar argument shows that all elements of U_2 have the same type over kVA_1 , and completes the proof of the claim.

Part (i) of the proposition now follows immediately from the following claim.

Claim 2. Let h be a generic element of $\text{Aff}(A_1, A_2)$. Then $\text{acl}(h)$ contains a unique algebraically closed subset of Morley rank 1.

Proof. There is a natural map $\pi : \text{Aff}(A_1, A_2) \rightarrow \text{Hom}(V, V)$, where for $h \in \text{Aff}(A_1, A_2)$ and $a_1 \in A_1$, $v \in V$, we have $h(a_1 + v) = h(a_1) + \pi(h)(v)$. We shall show that if h is generic in $\text{Aff}(A_1, A_2)$ and $b \in \text{acl}(h)$ with $\text{rk}(b) = 1$, then $b \in \text{acl}(\pi(h))$ (so $\text{acl}(\pi(h))$ is the claimed rank 1 algebraically closed subset). Suppose this is false, and choose b as above but with $b \notin \text{acl}(\pi(h))$, so b is independent from $\pi(h)$ over \emptyset . Let h' be an independent conjugate of h over $\text{acl}(b)$, and $g := \pi(h)$, $g' := \pi(h')$. Then as $\text{Hom}(V, V)$ has rank one, g, g' are independent elements of $\text{Hom}(V, V)$, so in particular are distinct. Now since $\text{Aff}(A_1, A_2)$ is invariant under the action of $V \times V$, the induced action on

$\text{Aff}(A_1, A_2)$ is elementary over CkV , by Claim 1. For any $f \in \text{Hom}(V, V)$, let $\Delta(f)$ be the graph of f and let $A(f) := \pi^{-1}(f)$. Then $V \times V$ fixes $A(f)$ setwise, and so acts on $A(f)$ with kernel $\Delta(f)$ (this is easily checked). Thus, $V \times V$ acts on $A(g') \times A(g)$ with kernel $\Delta(g) \cap \Delta(g') = \{0\}$, so the action is faithful. It is also easily checked that $V \times V$ is transitive on $A(g)$, so $\Delta(g')$ is transitive on $A(g)$, and likewise $\Delta(g)$ is transitive on $A(g')$. Thus, $V \times V = \Delta(g') \oplus \Delta(g)$ is transitive on $A(g) \times A(g')$. In particular, some generic $(h_1, h_2) \in A(g) \times A(g')$ has the same type (over kV) as (h, h') . Now as (g, g') is generic in $\text{Hom}(V, V)$, in fact (h_1, h_2) is generic in $\text{Aff}(A_1, A_2)$. In particular, $\text{acl}(h_1) \cap \text{acl}(h_2) = \text{acl}(\emptyset)$. Since $\text{tp}(h_1 h_2) = \text{tp}(hh')$, this contradicts the fact that $b \in \text{acl}(h) \cap \text{acl}(h')$. \square

(ii) Suppose \mathcal{G}^* is a collection of sorts consisting of a sort for open balls and a sort for closed balls. (We can add sorts for K , k , and Γ , but these are redundant - for example elements of K are closed balls of radius zero.) Much as in the proof of Lemma 2.6.2, it can be shown that if $C = \text{acl}(C)$, then any C -definable k -internal subset of $(\mathcal{G}^*)^n$ is a subset of a finite union of sets $\text{red}(u_1) \times \dots \times \text{red}(u_m) \times \{c\}$, where c is a tuple in C and the u_i are C -definable closed balls. The result now follows from (i). \square

References

- [1] *Z. Chatzidakis, E. Hrushovski*, Model theory of difference fields, *Trans. Amer. Math. Soc.* 351 (1999), 2997–3051.
- [2] *D. Haskell, H. D. Macpherson*, Cell decompositions of C -minimal structures, *Ann. Pure Appl. Logic* 66 (1994), 113–162.
- [3] *D. Haskell, E. Hrushovski, H.D. Macpherson*, Definable sets in algebraically closed valued fields: stable domination and independence, preprint.
- [4] *W. Hodges*, *Model Theory*, Cambridge University Press, Cambridge, 1993.
- [5] *J.E. Holly*, Canonical forms of definable subsets of algebraically closed and real closed valued fields, *J. Symb. Logic* 60 (1995), 843–860.
- [6] *J.E. Holly*, Prototypes for definable subsets of algebraically closed valued fields, *J. Symb. Logic* 62 (1997), 1093–1141.
- [7] *A.A. Ivanov, H.D. Macpherson*, Strongly determined types, *Ann. Pure Appl. Logic* 99 (1999), 197–230.
- [8] *I. Kaplansky*, Maximal fields with valuations I, *Duke Math. J.* 9 (1942), 303–321.
- [9] *S. Lang*, *Algebra*, Addison-Wesley, Menlo Park, 2nd edition, 1984.

- [10] *D. Lascar*, *Stability in Model Theory*, Longman, 1987.
- [11] *A.J. Macintyre*, *K. McKenna*, *L. van den Dries*, Elimination of quantifiers in algebraic structures, *Adv. in Math.* 47 (1983), 74–87.
- [12] *H.D. Macpherson*, *C. Steinhorn*, On variants of o-minimality, *Ann. Pure Appl. Logic* 79 (1996), 165–209.
- [13] *B. Poizat*, *A course in model theory. An introduction to contemporary mathematical logic*. Springer, New York, 2000.
- [14] *P. Ribenboim*, *The theory of classical valuations*, Springer, New York, 1998.
- [15] *A. Robinson*, *Complete theories*, North-Holland, Amsterdam, 1956.
- [16] *V. Weispfenning*, Quantifier elimination and decision procedures for fields, in: G.H.Müller and M.M. Richter, eds., *Models and Sets*, *Lecture Notes in Mathematics* 1103 (Springer, Berlin, 1984), 419–472.

Deirdre Haskell,
 Department of Mathematics and Statistics,
 McMaster University,
 Hamilton, Ontario L8S 4K1, Canada,
 haskell@math.mcmaster.ca

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

Dugald Macpherson,
 Department of Pure Mathematics,
 University of Leeds,
 Leeds LS2 9JT, England,
 h.d.macpherson@leeds.ac.uk